



runZero

User Guide

v4.0.240425.0

Updated on 2024-04-25

Contents

- User Guide
- Quick start
- Deployment plan
- Creating an account
 - Activating your account
 - Changing your password
 - Updating your profile picture
- Installing an Explorer
 - Installation
 - System requirements
 - Web screenshots
 - Configuration
 - Network communication
 - Restarting an Explorer
 - Removing an Explorer
 - Log management
 - Certificate Authorities (CAs)
 - Manual mode
 - Storage locations
 - Container installations
 - Automated installations
 - AWS EC2 installations
 - runZero binary verification
 - Automated MSI deployments
 - Installing on a Raspberry Pi
- Full-scale deployment
 - Planning
 - Initial configuration
 - Analysis
 - Advanced configuration / Optimization
 - Automation
 - Rollout
 - Additional Resources
 - Getting help
 - Types of networks
- Organizations
 - Use cases for organizations
 - Creating organizations
 - Organization configurations
 - Projects
 - Organization and project details
- Sites
 - Use cases for sites
 - Creating sites
 - Site configurations
 - Importing and exporting sites

- Sites and Explorers
- Self-hosting runZero
 - Background
 - Requirements
 - Offline mode
 - Installation steps
 - Installation with your own PostgreSQL database
 - runZero updates
 - Managing users
 - CLI service management
 - CLI update management
 - CLI user management
 - CLI organization management
 - Advanced configuration
 - Permissions
 - Backup and restoration
 - Support and debugging
 - Manual migrations
 - Offline mode configuration
 - High-availability configuration
 - Self-hosted troubleshooting
- Data retention
 - Stale asset expiration
 - Offline asset expiration
 - Scan data expiration
 - Event records
- User details page
- Default roles
 - Superusers
 - Administrators
 - Users
 - Billing
 - Annotators
 - Viewers
 - No Access
- Account settings
 - Single-sign on (SSO)
 - Multi-factor authentication (MFA)
 - Disabling support access
 - Idle times and login duration
 - Account API keys
 - License information
 - Entity information
- Managing user groups
 - Creating user groups
 - Adding users to user groups
 - Setting an expiration date for a user group
 - Viewing users in a user group
 - Viewing user groups assigned to a user
 - Removing users from a user group

- Deleting user groups
- Searching for users and user groups
- Bulk importing users
 - Creating the CSV file for importing users
 - Importing users into runZero
 - Verifying users have registered
- Managing external users
 - Inviting an external user into your account
 - Switching clients and organizations as an external user
 - Removing an external user from your account
- Implementing SSO
 - Specific SSO providers
 - Identity provider settings
 - SSO walkthrough
 - Service provider information
 - Common problems
 - Managing SSO group mappings
 - Setting up Azure AD SSO
 - Setting up Okta SSO
- Managing licenses
 - How do I view my license?
 - What count as recent assets?
 - When does my subscription expire?
 - How do I renew my subscription?
 - How do I convert to the Community Edition?
 - How do I find my invoices?
 - How do I change or cancel my subscription?
- Active scanning
- Traffic sampling
- Inbound integrations
- Discovery scanning
 - Site
 - Explorer
 - Hosted zone
 - Discovery scope
 - Scan name
 - Scan speed
 - Schedule
 - Scheduling grace period
 - Advanced scan options
 - Running initial scans
 - Scanning SNMP
 - Leveraging scan probes
- Managing Explorers
 - Viewing all Explorers
 - Screenshot capabilities
 - Searching for Explorers
 - Explorer actions
 - Bulk management operations
 - Viewing Explorer details

- Upgrading npcap
- Managing scan templates
 - Creating a scan template
 - Applying a scan template to a scan
 - Editing a scan template
 - Deleting a scan template
- Identifying gaps in scanning
 - Background
- Credentials
 - Credential settings
 - Credential verification
 - Credential management
- Using custom fingerprints
 - Create new fingerprints
 - Add new fingerprints to your self-hosted runZero instance
 - Restart the runZero service
 - Verify your fingerprints
- Inbound integrations
 - Enriching runZero results with data from other tools
 - Supported integrations
 - Scan probes or connector tasks
 - Importing integration data
 - Automatic asset merge
 - Removing an integration data source
 - Source names and IDs
 - Amazon Web Services
 - Azure Active Directory
 - Censys Search & Data
 - CrowdStrike Falcon
 - Google Cloud Platform
 - Google Workspace
 - Microsoft 365 Defender
 - Microsoft Active Directory
 - Microsoft Azure
 - Microsoft Intune
 - Miradore MDM
 - Qualys VMDR
 - Rapid7
 - SentinelOne
 - Shodan
 - Tenable
 - VMware
 - Wiz
- Outbound integrations
 - Using runZero data to enrich other tools
 - Atlassian Insight & Jira Service Management
 - Panther
 - SecurityGate.io
 - ServiceNow Service Graph
 - Splunk Search

- Sumo Logic
- Tines
- Thinkst Canary
- Reviewing results
 - Task details
 - Dashboard & inventory views
 - Insights from queries
 - Reports
 - Alerts
- Data formats
 - Formats
 - Scan data
 - Asset data
 - Change reports
- Using the dashboard
 - Customizing the dashboard
 - Widget library
 - Widget types
 - Drill down
 - Printing and exporting
- Using the inventory
 - Understanding assets
 - Loading assets
 - Connecting to other systems
 - Viewing services
 - Viewing screenshots
 - Viewing software
 - Viewing vulnerabilities
 - Viewing wireless networks
- Understanding assets
 - Asset fields
- Managing asset risk and criticality
 - Defining risk and criticality
 - Assigning asset risk and criticality
 - Asset risk report
- Managing ownership
 - Ownership types
 - Assigning owners to assets and vulnerabilities
- Managing tasks
 - Task status values
 - Tabs
 - Task details
 - Scheduled tasks
 - System tasks
 - Dismissing failed tasks
- Tracking goal progress
 - Goal creation
 - Creating system query or custom query goals
 - Creating asset ownership goals
 - Creating asset risk goals

- Goal progress calculation
- Goal events and notifications
- Understanding network segmentation
 - runZero multi-homed asset detection
 - Using the bridge report
 - Using the asset route pathing report
- Managing alerts
 - Using the rules engine
 - Creating alert templates
- Querying your data
 - Filtering and searching data
 - System and custom queries
 - Creating and editing queries
 - Search query syntax
 - Query examples
 - Inventory keywords
 - Component keywords
 - Automating queries
- Asset route pathing report
 - Terminology
 - Generating the asset route pathing report
 - Analyzing the report
 - Sharing the report
 - Exporting a dotfile
 - FAQs
- Coverage reports
 - RFC1918 coverage report
- Site and organization comparison
 - Generate a site comparison
 - View how assets change over time
 - View how exposure differs between networks
 - Analyze the results in the site comparison report
 - Search the site comparison report
- Organization Overview Report
 - Generating the Organization Overview Report
 - Email notifications
- External Asset Report
 - Generating the External Asset Report
 - Email notifications
- Exporting asset data
- Exporting HP iLO data
 - How to export HP iLO CSV data in runZero Enterprise
 - HP iLO CSV export data
- Leveraging the API
 - API keys and tokens
 - API client credentials
 - API rate limiting
 - Additional documentation
- Using the standalone command-line scanner
 - Performance

- Automatic web screenshots
- Scan outputs
- Raw Scan Data
- Glossary
 - Terms
- Frequently Asked Questions
 - Issues and FAQs
 - Identical assets in inventory
 - Scanning routers
 - Scanning VMWare virtual machines
 - Explorer not capturing screenshots
 - Protocols scanned by runZero
 - Ports scanned by runZero
 - Scanning IoT and OT
 - Browsers supported by the runZero Console
- Latest release notes
 - 4.0.240425.0
 - 4.0.240424.0
 - 4.0.240423.0
 - 4.0.240419.0
 - 4.0.240417.0
 - 4.0.240411.0
 - 4.0.240410.0
 - 4.0.240408.0
 - 4.0.240405.0
 - 4.0.240404.0
 - 4.0.240403.0
 - 4.0.240402.0
 - 4.0.240401.0
 - 4.0.240331.0
 - 4.0.240329.0
 - 4.0.240327.0
 - 4.0.240326.0
 - 4.0.240325.0
 - 4.0.240320.0
 - 4.0.240318.0
 - 4.0.240314.0
 - 4.0.240311.0
 - 4.0.240308.0
 - 4.0.240306.0
 - 4.0.240305.1
 - 4.0.240305.0
 - 4.0.240304.0
 - 4.0.240301.0
 - 4.0.240228.0
 - 4.0.240226.0
 - 4.0.240223.0
 - 4.0.240221.0
 - 4.0.240218.0
 - 4.0.240216.0

- 4.0.240214.0
- 4.0.240213.0
- 4.0.240208.0
- 4.0.240207.0
- 4.0.240206.0
- 4.0.240205.0
- 4.0.240202.0
- 4.0.240131.0
- 4.0.240129.0
- 4.0.240126.0
- 4.0.240124.0
- 4.0.240122.0
- 4.0.240119.0
- 4.0.240117.0
- 4.0.240112.0
- 4.0.240110.0
- 4.0.240109.0
- 4.0.240105.0
- 4.0.240103.0
- 4.0.231222.1
- 4.0.231222.0
- 4.0.231220.0
- 4.0.231218.0
- 4.0.231215.0
- 4.0.231213.0
- 4.0.231211.0
- 4.0.231208.0
- 4.0.231207.0
- 4.0.231205.0
- 4.0.231201.1
- 4.0.231201.0
- 4.0.231130.0
- 4.0.231128.0
- 4.0.231122.0
- 4.0.231121.0
- 4.0.231120.0
- 4.0.231116.0
- 4.0.231115.0
- 4.0.231114.0
- 4.0.231108.1
- 4.0.231108.0
- 4.0.231102.0
- 4.0.231101.0
- 4.0.231031.2
- 4.0.231031.1
- 4.0.231031.0
- 4.0.231030.0
- 4.0.231027.1
- 4.0.231027.0
- 4.0.231023.0

- 4.0.231019.0
- 4.0.231018.0
- 4.0.231017.1
- 4.0.231017.0
- 4.0.231013.0
- 4.0.231012.0
- 4.0.231011.0
- 4.0.231006.0
- 4.0.231005.0
- 4.0.231002.0
- 4.0.230928.0
- 4.0.230927.0
- 4.0.230925.0
- 3.10.230921.0
- 3.10.230920.0
- 3.10.230918.0
- 3.10.230917.1
- 3.10.230917.0
- 3.10.230913.0
- 3.10.230911.0
- 3.10.230908.01
- 3.10.25
- 3.10.24
- 3.10.23
- 3.10.22
- 3.10.21
- 3.10.20
- 3.10.19
- 3.10.18
- 3.10.17
- 3.10.16
- 3.10.15
- 3.10.14
- 3.10.13
- 3.10.12
- 3.10.11
- 3.10.10
- 3.10.9
- 3.10.8
- 3.10.7
- 3.10.6
- 3.10.5
- 3.10.4
- 3.10.3
- 3.10.2
- 3.10.1
- 3.10.0
- 3.9.10
- 3.9.9
- 3.9.8

- 3.9.7
- 3.9.6
- 3.9.5
- 3.9.4
- 3.9.3
- 3.9.2
- 3.9.1
- 3.9.0
- 3.8.14
- 3.8.13
- 3.8.12
- 3.8.11
- 3.8.10
- 3.8.9
- 3.8.8
- 3.8.7
- 3.8.6
- 3.8.5
- 3.8.4
- 3.8.3
- 3.8.2
- 3.8.1
- 3.8.0
- 3.7.11
- 3.7.10
- 3.7.9
- 3.7.8
- 3.7.7
- 3.7.6
- 3.7.5
- 3.7.4
- 3.7.3
- 3.7.2
- 3.7.1
- 3.7.0
- 3.6.19
- 3.6.18
- 3.6.17
- 3.6.16
- 3.6.15
- 3.6.14
- 3.6.13
- 3.6.12
- 3.6.11
- 3.6.10
- 3.6.9
- 3.6.8
- 3.6.7
- 3.6.6
- 3.6.5

- 3.6.4
- 3.6.3
- 3.6.2
- 3.6.1
- 3.6.0
- 3.5.10
- 3.5.9
- 3.5.8
- 3.5.7
- 3.5.6
- 3.5.5
- 3.5.4
- 3.5.3
- 3.5.2
- 3.5.1
- 3.5.0
- 3.4.23
- 3.4.22
- 3.4.21
- 3.4.20
- 3.4.19
- 3.4.18
- 3.4.17
- 3.4.16
- 3.4.15
- 3.4.14
- 3.4.13
- 3.4.12
- 3.4.11
- 3.4.10
- 3.4.9
- 3.4.8
- 3.4.7
- 3.4.6
- 3.4.5
- 3.4.4
- 3.4.3
- 3.4.2
- 3.4.1
- 3.4.0
- 3.3.8
- 3.3.7
- 3.3.6
- 3.3.5
- 3.3.4
- 3.3.3
- 3.3.2
- 3.3.1
- 3.3.0
- 3.2.11

- 3.2.10
- 3.2.9
- 3.2.8
- 3.2.7
- 3.2.6
- 3.2.5
- 3.2.4
- 3.2.3
- 3.2.2
- 3.2.1
- 3.2.0
- 3.1.13
- 3.1.12
- 3.1.11
- 3.1.10
- 3.1.9
- 3.1.8
- 3.1.7
- 3.1.6
- 3.1.5
- 3.1.4
- 3.1.3
- 3.1.2
- 3.1.1
- 3.1.0
- 3.0.24
- 3.0.23
- 3.0.21
- 3.0.20
- 3.0.19
- 3.0.18
- 3.0.17
- 3.0.16
- 3.0.15
- 3.0.14
- 3.0.12
- 3.0.11
- 3.0.10
- 3.0.9
- 3.0.8
- 3.0.7
- 3.0.6
- 3.0.5
- 3.0.4
- 3.0.3
- 3.0.2
- 3.0.1
- 3.0.0
- 2.15.11
- 2.15.10

- 2.15.9
- 2.15.8
- 2.15.7
- 2.15.6
- 2.15.5
- 2.15.4
- 2.15.3
- 2.15.2
- 2.15.1
- 2.15.0
- 2.14.11
- 2.14.10
- 2.14.9
- 2.14.8
- 2.14.7
- 2.14.6
- 2.14.5
- 2.14.4
- 2.14.3
- 2.14.2
- 2.14.1
- 2.14.0
- 2.13.7
- 2.13.6
- 2.13.5
- 2.13.4
- 2.13.3
- 2.13.2
- 2.13.1
- 2.13.0
- 2.12.12
- 2.12.11
- 2.12.10
- 2.12.9
- 2.12.8
- 2.12.7
- 2.12.6
- 2.12.5
- 2.12.4
- 2.12.3
- 2.12.2
- 2.12.1
- 2.12.0
- 2.11.16
- 2.11.15
- 2.11.14
- 2.11.13
- 2.11.12
- 2.11.11
- 2.11.10

- 2.11.9
- 2.11.8
- 2.11.7
- 2.11.6
- 2.11.5
- 2.11.4
- 2.11.3
- 2.11.2
- 2.11.1
- 2.11.0
- 2.10.6
- 2.10.5
- 2.10.4
- 2.10.3
- 2.10.2
- 2.10.1
- 2.10.0
- 2.9.14
- 2.9.13
- 2.9.12
- 2.9.11
- 2.9.10
- 2.9.9
- 2.9.8
- 2.9.7
- 2.9.6
- 2.9.5
- 2.9.4
- 2.9.3
- 2.9.2
- 2.9.1
- 2.9.0
- 2.8.14
- 2.8.13
- 2.8.12
- 2.8.11
- 2.8.10
- 2.8.9
- 2.8.8
- 2.8.7
- 2.8.6
- 2.8.5
- 2.8.4
- 2.8.3
- 2.8.2
- 2.8.1
- 2.8.0
- 2.7.11
- 2.7.10
- 2.7.9

- 2.7.8
- 2.7.7
- 2.7.6
- 2.7.5
- 2.7.4
- 2.7.3
- 2.7.2
- 2.7.1
- 2.7.0
- 2.6.4
- 2.6.3
- 2.6.2
- 2.6.1
- 2.6.0
- 2.5.8
- 2.5.7
- 2.5.6
- 2.5.5
- 2.5.4
- 2.5.3
- 2.5.2
- 2.5.1
- 2.5.0
- 2.4.4
- 2.4.3
- 2.4.2
- 2.4.1
- 2.4.0
- 2.3.5
- 2.3.4
- 2.3.3
- 2.3.2
- 2.3.1
- 2.3.0
- 2.2.5
- 2.2.4
- 2.2.3
- 2.2.2
- 2.2.1
- 2.2.0
- 2.1.7
- 2.1.6
- 2.1.5
- 2.1.4
- 2.1.3
- 2.1.2
- 2.1.1
- 2.1.0
- 2.0.14
- 2.0.13

- 2.0.12
- 2.0.11
- 2.0.10
- 2.0.9
- 2.0.8
- 2.0.7
- 2.0.6
- 2.0.5
- 2.0.4
- 2.0.3
- 2.0.2
- 2.0.1
- 2.0.0
- 1.16.6
- 1.16.5
- 1.16.4
- 1.16.3
- 1.16.2
- 1.16.1
- 1.16.0
- 1.15.6
- 1.15.5
- 1.15.4
- 1.15.3
- 1.15.2
- 1.15.1
- 1.15.0
- 1.14.9
- 1.14.8
- 1.14.7
- 1.14.6
- 1.14.5
- 1.14.4
- 1.14.3
- 1.14.2
- 1.14.1
- 1.14.0
- 1.13.11
- 1.13.10
- 1.13.9
- 1.13.8
- 1.13.7
- 1.13.6
- 1.13.5
- 1.13.4
- 1.13.3
- 1.13.2
- 1.13.1
- 1.13.0
- 1.12.9

- 1.12.8
- 1.12.7
- 1.12.6
- 1.12.5
- 1.12.4
- 1.12.3
- 1.12.2
- 1.12.1
- 1.12.0
- 1.11.9
- 1.11.8
- 1.11.7
- 1.11.6
- 1.11.5
- 1.11.4
- 1.11.3
- 1.11.2
- 1.11.1
- 1.11.0
- 1.10.8
- 1.10.7
- 1.10.6
- 1.10.5
- 1.10.4
- 1.10.3
- 1.10.2
- 1.10.1
- 1.10.0
- 1.9.10
- 1.9.9
- 1.9.8
- 1.9.7
- 1.9.6
- 1.9.5
- 1.9.4
- 1.9.3
- 1.9.2
- 1.9.1
- 1.9.0
- 1.8.14
- 1.8.13
- 1.8.12
- 1.8.11
- 1.8.10
- 1.8.9
- 1.8.8
- 1.8.7
- 1.8.6
- 1.8.5
- 1.8.4

- [1.8.3](#)
- [1.8.2](#)
- [1.8.1](#)
- [1.8.0](#)
- [1.7.13](#)
- [1.7.12](#)
- [1.7.11](#)
- [1.7.10](#)
- [1.7.9](#)
- [Archived release notes](#)
- [Older platform release notes](#)
 - [Platform release notes](#)
- [Older Explorer release notes](#)
 - [Explorer release notes \(agents\)](#)
- [Older scanner release notes](#)
 - [Scanner release notes](#)

Getting started

To get started, you'll need to sign up for a runZero account. The default account is a trial of the full runZero Platform. After the trial expires, you will have the option to convert to the free Community Edition.

[Sign up for a runZero account](#)

Read up on [creating an account](#) for help activating your account, changing your password, and adding a profile picture.

Once your account is set up, there are a couple of paths you can take to deploy runZero.

Quick start

The quick start path is recommended for testing out runZero. You will jump straight into deploying an Explorer for discovery, running your first scan, and onboarding users. If you are looking for more to test out after finishing these tasks, you can jump to the [deployment plan](#) to dive deeper.

- **Deploy Explorers:** runZero Explorers are the scanners that will be used to discover your network. [Deploy an Explorer](#) under the Deploy menu. Make sure to review the system requirements and dependencies for web screenshots as you plan where to install the Explorer.
- **Run initial scans:** After deploying your Explorers, [you can run your first scan](#). A discovery scan finds, identifies, and builds an inventory of all the connected assets on your internal network. The Explorer and discovery scope are the most important aspects of configuration, followed by SNMP community information to support the generation of switch topology.
- **Review results:** After running the network discovery scan, you can see what was found by [reviewing your results](#).
- **Create users:** The person who signs up for runZero service will be given superuser permissions. See the [“Managing Your Team”](#) page to learn how to create additional users, set up single sign-on and multi-factor authentication options, and discover the user roles available.

Deployment plan

The deployment plan path is meant for a full production deployment of runZero. You will do more planning and configurations prior to getting started on scanning.

- **Review the glossary:** The [glossary](#) covers all of the runZero terms you will want to know.
- **Start the deployment plan:** The [deployment plan](#) covers the steps you will want to do for a production deployment of runZero.

Creating an account

To get started, you'll need to sign up for a runZero account. The default account is a trial of the full runZero Platform. After the trial expires, you will have the option to convert to the free Community Edition.

[Sign up for a runZero account](#)

Activating your account

After you sign up for an account, we'll email you a link to activate your account. If you don't see an email from us, check your spam folder.

Open the link in the email to go to the **Activation** page. Follow the instructions on the page to activate your account. You'll need to provide your name, set up a password, specify your location, and accept our [privacy policy](#) and [terms of service](#).

After activating your account, you'll be taken directly to the runZero Console. Your new account has administrative access, so you will be able to manage [sites](#), [organizations](#), [users](#), and [Explorers](#).

If you have any trouble creating your account, please contact [support](#).

Changing your password

To change your password, go to your [account settings](#). You'll need to provide your current password before you can enter a new one.

All passwords must contain:

- At least 8 characters
- At least 1 uppercase character
- At least 1 lowercase character
- At least 1 number

Updating your profile picture


User profile images are managed through [Gravatar](#) and associated with your email address. If you don't have an account, [sign up for one](#).

Installing an Explorer

runZero requires the use of at least one Explorer within your environment to enable network discovery. The Explorer should be installed on a system with reliable connectivity to the network you want to discover. For internal networks, runZero works best when installed on a system with a wired (vs wireless) connection.

For external network discovery, nearly any cloud provider with a reliable connection should do. If the runZero Explorer is installed in a container or virtualized system, ensure that it has direct access to the network (host networking in Docker, bridged networking in VMware, etc).

[View an interactive version of this diagram](#)



Installation

To install the runZero Explorer, log in to the runZero Console and switch to the Organization that should be associated with the Explorer. Explorer downloads are then available by selecting Deploy in the left navigator and choosing the Deploy Explorers sub-menu.

Note: The Explorer download link is specific to your active organization and **using the wrong link can result a new Explorer being associated with the wrong organization.**

Download the correct binary for your system from the [Explorer download page](#). For most systems, select the 64-bit (x86_64) architecture. For macOS, you will need to select 64-bit Intel (x86_64) or ARM (Apple M*), depending on your hardware. For embedded devices, such as the Raspberry Pi 3+, choose the ARM7 architecture. Windows binaries are signed with a valid Authenticode signature, which should be validated before the executable is launched.

The Explorer installation process requires administrative privileges. On Windows, a UAC prompt may be displayed. On Linux and macOS the downloaded binary should be made executable (`chmod u+x runzero-agent.bin`) and then executed with root privileges (`sudo` or from root shell). In either case, the Explorer should install itself as a system service and start immediately, displaying a new entry in the [Explorers page](#).

System requirements

Windows

- Windows Server 2012 R2+ or Windows 10 Build 1604+
- Processor running at 2.0 GHz or faster
- At least 16GiB of memory (8GiB for small environments)
- At least 1GB of free storage space

Limitations

- The Trellix agent for Windows appears to interfere with network scans and traffic capture. Please use a system without the Trellix agent for active scans or passive network sampling.
- Windows Server 2008, Windows Server 2012, Windows 7, and Windows 8 may be able to run the Explorer in a pinch, but are not officially supported.
- Windows Explorers are limited to a single concurrent scan task due to performance limitations of the raw packet driver.

Linux

- Kernel version 2.6.23 or later
- Processor running at 2.0 GHz or faster
- At least 16GiB of memory (8GiB for small environments)
- At least 1GB of free storage space

Linux ARM devices with limited processing power and memory, such as the Raspberry Pi, can run the runZero Explorer, but may have trouble scanning larger networks, or running integrations with large amounts of data.

MacOS

- macOS 10.11 (El Capitan) or newer
- Processor running at 2.0 GHz or faster
- At least 16GiB of memory (8GiB for small environments)
- At least 1GB of free storage space

macOS systems running Catalina (10.15) or newer need to use the **curl** download method to avoid issues with the new Notary requirements.

BSD variants

- Processor running at 2.0 GHz or faster
- At least 16GiB of memory (8GiB for small environments)
- At least 1GB of free storage space

Requires root access to a system running a recent version of the operating system. FreeBSD 11.2 or newer, recent versions of NetBSD/DragonFly/OpenBSD.

Web screenshots

[Google Chrome](#) should be installed on the Explorer system to enable [web screenshots](#). Please note that “snap”-based Chromium installs (Ubuntu 20.04 and newer) don’t appear to work properly in headless mode and the official Chrome packages should be used instead with the following commands:

```
$ curl -o chrome.deb https://dl.google.com/linux/direct/google-chrome-stable_current_amd64.deb
sudo apt install ./chrome.deb
```

Configuration

Note: Existing installations may refer to `Rumble` in the directory name, service name, and binary names. `Rumble, Inc` was the previous name of our business and new installations will only refer to `runZero`, our new name. going forward.

The Explorer can be configured by setting variables in a `config.txt` file located in the same directory as the executable. On Windows this file should be created in `C:\Program Files\runZero\config.txt`, while other platforms should use `/opt/runzero/bin/config.txt`. In addition to `config.txt`, the legacy `.env` name is also supported. If the file `/etc/runzero/config` is present, this will be preferred as the configuration settings for the Explorer.

The format of this file is `VAR=VAL` with one variable per line.

Configuration file locations

Windows: `C:\Program Files\runZero\config.txt`

Other Platforms: `/opt/runzero/bin/config.txt`

Network communication

The Explorer connects to the `console.runzero.com` host on TCP port 443 using TLS and two static IPv4 addresses (13.248.161.247, 76.223.34.198) and two static IPv6 addresses (2600:9000:a415:cd87:fbe5:476a:3533:69f2, 2600:9000:a716:ee91:85f9:3c9:48c9:59b9). This connection is used for Explorer registration, job scheduling, status messages, and submission of completed scan jobs. For completely offline environments, the [runZero Scanner](#) can be used to create scan data files that can be uploaded later via the [Inventory Import](#) action. The host `console.runzero.com` is used for automatic updates of the Explorer executable.

Proxy support

Please note that certain web proxies that perform TLS inspection do not handle WebSocket communication properly and TLS inspection will need to be disabled for the runZero Explorer to successfully connect. The most popular product with this problem is the Sophos (previously Cyberoam) security appliance. Websense users may need to add a [bypass rule](#) for `console.runzero.com`.

Proxy support is handled automatically in most cases. On the Windows platform, proxy information is read from the registry keys (used by Chrome, Edge, and IE).

The proxy can also be configured by setting the `HTTPS_PROXY` environment variable. The value of the `HTTPS_PROXY` environment variable should be a hostname and port (`proxy.example.com:8080`) or just a hostname (`proxy`). Environment variables are read from your configuration file. Please view the [Configuration](#) section to see how to set environment variables. The common option is to create a file named `config.txt` in the same directory as the Explorer binary and set the environment variables in the format described below.

It's also possible to use a SOCKS proxy by setting `HTTPS_PROXY` to a socks5 URL, for example `HTTPS_PROXY=socks5://socks.example.com:1080`.

The Explorer will attempt to use the configured proxy for each probe. If it doesn't succeed, it will try making a direct connection. This means both proxy and non-proxy connection attempts may appear in logs. Using only proxies to try to hide or anonymize Explorer connections is not supported.

Restarting an Explorer

The easiest way to restart an Explorer is to force a software update from the cloud console. Otherwise, you can find the service on the host machine and restart it by hand.

On Linux or Mac, you can run `/opt/runzero/bin/runzero-agent-[uuid] restart` where `[uuid]` is the ID of the organization the Explorer belongs to.

On Linux systems you can also use `systemd` to restart Explorers. First obtain the name of the Explorer (`runzero-agent-[uuid]`) service:

```
$ systemctl | grep runzero-agent
```

Then restart the service using this name:

```
$ systemctl restart runzero-agent-[uuid]
```

A `kill -9` of the Explorer pid should cause a restart as well.

On macOS, you can use `launchctl` to restart the Explorer:

```
launchctl kickstart -k runzero-agent-[uuid]
```

As with Linux, the `[uuid]` is the organization UUID, which you can find by looking at the `runzero-agent-*` filename in `/opt/runzero/bin`

On Windows, you can use the Services console to restart the Explorer like any other background service.

Removing an Explorer

The easiest way to remove an Explorer is to use the [Explorers](#) page Manage menu and choose the Remove Explorer option. This will remove the service and terminate the current Explorer process. If you would like to remove the Explorer without using the runZero Console, there are a couple of options.

On the Windows platform, each Explorer will be listed in Programs and Features (as the runZero Agent), and can be uninstalled like any other application.

On all platforms, including Windows, the Explorer can uninstall itself if run with the `uninstall` argument from a root or Administrator shell:

Windows

```
c:\Program Files\runZero\runzero-agent-[uuid].exe uninstall
```

Other Platforms

```
/opt/runzero/bin/runzero-agent-[uuid] uninstall
```

Log management

The Explorer logs to a file and to standard output by default. On Windows the default log file location is the installation directory (`C:\Program Files\runZero`) while other platforms log to the files `/var/log/runzero.log` and `/var/log/runzero.err`. The default configuration limits

log files to 100MiB, creates three backups, and expires logs after 90 days. These defaults can be changed by setting the following values in the \$BIN/config.txt file:

- The RUNZERO_AGENT_LOG_MAX_SIZE setting controls the maximum log size in mibibytes. The default is **100**.
- The RUNZERO_AGENT_LOG_MAX_BACKUPS setting controls the number of backup files created by log rotation. The default is **3**.
- The RUNZERO_AGENT_LOG_MAX_AGE setting controls the maximum age in days, this applies to all files, including backups. The default is **90**.
- The RUNZERO_AGENT_LOG_COMPRESS setting determines whether to gzip compress the backups. The default is **false**.
- The RUNZERO_AGENT_LOG_STDOUT setting determines whether to write logs to standard output. The default is **true**. On Linux this results in logs being written to the system log when the Explorer is started by systemd or upstart. On macOS this results in separate logs viewable in the Console application under “Log Reports” when the Explorer is started by launchd.
- The RUNZERO_AGENT_LOG_FILE setting determines whether to write logs to a log file as described above. The default is **true**, set to false to disable log file writing.

The Explorer must be restarted for these settings to take effect.

Certificate Authorities (CAs)

The runZero Explorer uses the system-installed certificate authorities to validate TLS connections in addition to an internal CA certificate bundle (derived from Debian 10). By default, both the system certificate roots, and the bundled roots are considered for all secure TLS connections. This behavior can be controlled via environment variables (set in the \$BIN/config.txt file or at the system level):

- The RUNZERO_TLS_IGNORE_SYSTEM_ROOTCA setting can be set to **true** to ignore the system CA roots.
- The RUNZERO_TLS_IGNORE_EMBEDDED_ROOTCA setting can be set to **true** to ignore the bundled CA roots.
- The RUNZERO_TLS_ADDITIONAL_ROOTCA setting can be set to a file path containing additional CA roots in PEM format.

Manual mode

If a supported system service manager, such as systemd or upstart, is not detected, the runZero Explorer will switch to manual mode, running in the foreground, and replacing and re-executing its own binary as new updates become available. For temporary Explorer installations or to run the Explorer in a container environment, the argument “manual” can be specified:

```
$ sudo ./runzero-agent.bin manual
```

Storage locations

The runZero Explorer installs into %PROGRAMFILES%\runZero on Windows and /opt/runzero/bin on all other platforms. Temporary files are stored in the default operating system locations. These locations can be overridden using the config.txt file (see the above [Configuration](#) section). Note that the Explorer service needs to be restarted (or force updated) for these changes to take effect. Older installations may still refer to the Rumble versions of the previously mentioned directories and .env instead of config.txt for environment overrides.

On Windows, the temporary file location is chosen from the first non-empty environment value of TMP, TEMP, or USERPROFILE, falling back to the Windows directory. To override this location, set an entry in config.txt like the following:

```
TMP=D:\Storage\runZero
```

On all other platforms, the temporary file location is chosen based on the value of TMPDIR, falling back to /tmp otherwise. To override this location, set an entry in config.txt like the following:

```
TMPDIR=/home/storage/runzero
```

Any scans that fail to upload are stored in the runZero Explorer installation directory and can be imported into the platform manually or using the runZero Scanner's --import and --upload options.

Container installations

The runZero Explorer can run in standard container environments, but may require additional configuration. To run as a standalone executable, the Explorer can be run with the argument manual. For non-persistent containers an Explorer identifier needs to be persisted through an environment variable. This can be done by setting the variable RUNZERO_AGENT_HOST_ID to a 32-character hexadecimal string. This identifier is used to uniquely identify the Explorer within an organization.

To generate a suitable identifier, the openssl tool may be used:

```
$ openssl rand -hex 16
01b0283809b24511929d0b062bd36109
```

Here is a sample Containerfile you can edit and use:

```
#
# Sample Containerfile for running the runZero Explorer in a container, with
# screenshot support.
#
```

```
FROM debian:stable-slim

WORKDIR /opt/runzero

# Ensure curl is available and install tools for wireless scanning.
#
RUN apt update && apt install -y curl wireless-tools

# Install Chrome for screenshots.
#
RUN curl -o chrome.deb https://dl.google.com/linux/direct/google-chrome-stable_current_amd64.deb
    apt install -y ./chrome.deb

# Set AGENT_URL to be the download URL for your Linux runZero Explorer. To
# find your URL, go to https://console.runzero.com/deploy/download/explorers
# and click on the first URL box to copy it to the clipboard.
#
ENV AGENT_URL=https://console.runzero.com/download/explorer/DT[uniqueToken]/[versionID]/runzero

# This ID is used to track the Explorer even if the container is rebuilt.
# Set it to a unique 32 character hex ID. You can generate one via:
#
# $ openssl rand -hex 16
#
ENV RUNZERO_AGENT_HOST_ID=[UNIQUE-ID]

# If you need to set environment variables to change the Explorer behavior,
# you can do so via the ENV directive. Example:
#
# ENV RUNZERO_AGENT_LOG_DEBUG=true

ADD ${AGENT_URL} runzero-agent.bin

RUN chmod +x runzero-agent.bin

# For full functionality the runZero scanner needs to send and receive raw
# packets, which requires elevated privileges.
USER root

# The argument `manual` tells runZero not to look for SystemD or upstart.
ENTRYPOINT [ "/opt/runzero/runzero-agent.bin", "manual" ]
```

This containerfile works with podman as well as Docker. Note that because of the requirement for root privileges, you should start the container as root. For the best results, run the container with the `--privileged` option to allow the Explorer to listen to network traffic.

Automated installations

The Explorer will automatically install when executed if root or administrative privileges are available.

On Linux and BSD systems, automatic installation depends on the presence of a supported init service like systemd or upstart. If no supported init service is found, the Explorer will instead run in manual mode, automatically overwriting and re-executing itself with each update. To automatically deploy an Explorer on systems without a supported init service, the Explorer should be executed in the background and with the nohup wrapper.

On Windows systems, the Explorer will automatically install when run interactively or when the `updater` parameter is passed to the binary. For environments where MSIs are required, the [Explorer MSI](#) wrapper can be used to deploy an Explorer from the runZero Console or a local mirror.

AWS EC2 installations

The runZero Explorer can be run in an AWS EC2 instance. However, there are a number of configuration changes required to avoid packet loss when scanning.

- DNS resolution from EC2 to the AWS DNS server has a [fixed cap of 1024 packets per second](#) which cannot be increased. To avoid this, set a custom list of non-AWS nameservers in the scan configuration advanced section.
- Any Security Group without a `0/0` rule results in [connection tracking](#), which has an undocumented limit on connections per instance type. Avoid this by adding `0/0 allow` for ICMP/UDP inbound and outbound for the Explorer instance. Also add a `0/0 allow` for outbound TCP connections. (Inbound TCP is not currently required for runZero scans, but may be needed in the future for callback protocols.)
- Overall packet rates have [undocumented limits which depend on instance type](#). You will need to experiment with sizing your Explorer instances until scans are consistent for a given scan rate. We hope to gather and share data on appropriate instance sizes soon.

The Explorer should work well deployed to a memory optimized, compute optimized, or general compute instance. Since the Explorer can make full use of multi-core systems, you may want to target the number of cores to the number of simultaneous scans. You probably won't want to run larger scans on an instance with less than 32GiB of RAM.

runZero binary verification

runZero uses dynamically generated binaries for the runZero Scanner and runZero Explorer downloads. Although Windows binaries have a valid Authenticode signature, all binaries also contain a secondary, internal signature. Dynamic binaries make it easy to deploy Explorers that connect back to the right organization, but present a challenge for independent integrity validation. To enable verification of the internal signature, we offer the **runZero Verifier**. This

verification tool can confirm whether a given binary contains a valid internal signature, in addition to any existing Authenticode signatures.

To get started, download the latest version of the verifier from the bottom of this page along with the PGP signature file for the selected architecture.

The runZero Verifier is always signed by PGP Key ID [60EBAAE9AEF08C6D](#).

To validate the signature of the runZero Verifier for Windows 64-bit, you will need a [GPG](#) client and to run the following commands:

```
C:\> curl -s https://www.runzero.com/.well-known/security.pub.asc | gpg --import
C:\> gpg --verify runzero-verifier-3.0.0-windows-amd64.exe.asc
```

Successful validation will show a valid signature by key ID 9B5DAFF7D43349298A3039BD60EBAAE9AEF08C6D.

```
gpg: Signature made Sun 07 Aug 2022 11:33:15 AM CDT
gpg:          using RSA key 9B5DAFF7D43349298A3039BD60EBAAE9AEF08C6D
gpg:          issuer "security@runzero.com"
gpg: Good signature from "runZero Security <security@runzero.com>" [unknown]
```

The warning below is expected and does not indicate a problem with the signature:

```
gpg: WARNING: This key is not certified with a trusted signature!
gpg: There is no indication that the signature belongs to the owner.
```

Once the runZero Verifier itself has been validated, it can be used to check the signature of any runZero binary:

```
C:\> runzero-verifier-3.0.0-windows-amd64.exe runzero-explorer-3.0.0-windows-amd64.exe
runzero-explorer-3.0.0-windows-amd64.exe: VALID SIGNATURE
```

A failed validation will show the error Invalid or missing signature and the verifier will set exit status to 1.

Binary downloads

Windows

Build	PGP sig	SHA hash
runZero Verifier x86 64-bit	pgp signature	sha-256
runZero Verifier x86 32-bit	pgp signature	sha-256

Linux

Build	PGP sig	SHA hash
runZero Verifier x86 64-bit	pgp signature	sha-256
runZero Verifier x86 32-bit	pgp signature	sha-256

Additional Linux builds

Build	PGP sig	SHA hash
runZero Verifier ARM v5 32-bit	pgp signature	sha-256
runZero Verifier ARM v6 32-bit	pgp signature	sha-256
runZero Verifier ARM v7 32-bit	pgp signature	sha-256
runZero Verifier ARM 64-bit (aarch64)	pgp signature	sha-256
runZero Verifier PPC 64-bit Little Endian	pgp signature	sha-256
runZero Verifier MIPS 32-bit Big Endian	pgp signature	sha-256
runZero Verifier MIPS 32-bit Little Endian	pgp signature	sha-256
runZero Verifier MIPS 64-bit Big Endian	pgp signature	sha-256
runZero Verifier MIPS 64-bit Little Endian	pgp signature	sha-256
runZero Verifier S390X	pgp signature	sha-256

MacOS

Build	PGP sig	SHA hash
runZero Verifier x86 64-bit	pgp signature	sha-256
runZero Verifier ARM 64-bit	pgp signature	sha-256

BSD Variants

FreeBSD

Build	PGP sig	SHA hash
runZero Verifier x86 64-bit	pgp signature	sha-256
runZero Verifier x86 32-bit	pgp signature	sha-256
runZero Verifier ARM v6 32-bit	pgp signature	sha-256
runZero Verifier ARM v7 32-bit	pgp signature	sha-256

NetBSD

Build	PGP sig	SHA hash
runZero Verifier x86 64-bit	pgp signature	sha-256
runZero Verifier x86 32-bit	pgp signature	sha-256
runZero Verifier ARM v5 32-bit	pgp signature	sha-256
runZero Verifier ARM v6 32-bit	pgp signature	sha-256
runZero Verifier ARM v7 32-bit	pgp signature	sha-256

Dragonfly

Build	PGP sig	SHA hash
runZero Verifier 64-bit	pgp signature	sha-256

OpenBSD

Build	PGP sig	SHA hash
runZero Verifier 64-bit	pgp signature	sha-256

Automated MSI deployments

runZero uses dynamically generated binaries for the runZero Explorer downloads and this doesn't always play well with MSI-based installation methods.

To work around this issue, we have provided a shim MSI package that can be used with automated installers. This package has a valid Authenticode signature and can also be verified using the [runZero Verifier](#).

Note: Some components of the application still reference the name "rumble" for backwards compatibility. All new installations will use `runzero` instead of `rumble` for directory, file, and user names.

To use this package, deploy it with the URL parameter specified as the organization-specific download URL from the runZero Console Explorers section.

```
msiexec /i runzero-explorer-installer-amd64.msi URL=https://console.runzero.com/download/explorer/DT[uniqueToken]/[versionID]/runzero-explorer-windows-amd64.exe
```

Note: the above command should be entered as one long line with a single URL parameter, but is shown wrapped here.

The MSI shim will verify that the URL contains a valid runZero Explorer binary and install it normally.

Warning: Note that installing with the `/a` parameter will not work and `/i` must be used instead.

Binary downloads

Build	SHA256
runZero Explorer Installer MSI x86 64-bit	sha-256
runZero Explorer Installer MSI x86 32-bit	sha-256

Installing on a Raspberry Pi

The runZero Explorer enables discovery scanning. In most cases, you can deploy an Explorer on an existing system that has connectivity to the network you want to discover. However, there may be times when the traditional deployment model may not work for you. Some locations, like retail stores or customer sites, may not have staff or hardware available to install the Explorer, making remote deployment a bit tricky.

In these types of scenarios, you can install a runZero Explorer on a Raspberry Pi and send the device to the location for them to plug into their network.

What you'll need

- Raspberry Pi 4 Model B (4GB or 8GB), Raspberry Pi Compute Module 4, or Raspberry Pi 400
- At least 1GB free of storage space on your Raspberry Pi's MicroSD card after installing the operating system

Connecting to your Raspberry Pi via SSH

In this section, we're going to show you how to SSH to your Raspberry Pi and install the Explorer from your terminal.

Step 1. Enable SSH on your Raspberry Pi

Before you can connect to your Raspberry Pi, make sure to [enable SSH](#) on the device.

You can add a file called `ssh.txt` at the root of the SD card and reboot the Raspberry Pi. The contents of the file can be empty. On reboot, the Pi looks for the SSH file and enables SSH if it finds the file.

Step 2. Update the password for your Raspberry Pi

If you enable SSH on your Raspberry Pi, you must [update your credentials](#).

Log in as the pi user and use the passwd command to change the default password. Entering the passwd command will prompt you for your current password to authenticate before you can change it.

Step 3. Copy the Explorer instructions

- Go to your [console](#). Verify you are in the right organization. runZero keys your Explorer's download link to the organization you are currently viewing to associate them together.
- Go to the [Explorer deployment page](#) and select **Linux Distributions**, and then choose Linux ARM 32-bit V7. Note that Raspbian uses a 32-bit kernel by default, even on 64-bit Raspberry Pi hardware.
- Click the instructions at the bottom of the Linux installation page to place the commands into your clipboard.

Step 4. Install the Explorer

Note: Some components of the application still reference the name "Rumble" for backwards compatibility. The documentation will be updated as these are changed.

Paste the installation command into the terminal after connecting to the Raspberry Pi via SSH. The URL in the command links the installation to your active organization and will be different from the example below.

```
me@mac ~ % ssh pi@hostIP
pi@host's password:
pi@raspberrypi:~ $ sudo bash
root@raspberrypi:/home/pi# cd Downloads
root@raspberrypi:/home/pi# curl -o runzero-explorer.bin https://console.runzero.com/download/ex
```

The Explorer automatically installs itself in /opt/runzero and sets up a systemd service with the name runzero-agent-. The service automatically starts on boot.

Next steps

Your Raspberry Pi is now set up to use as a runZero Explorer box and ready to be sent wherever you need.

You can always view and manage your [registered Explorers](#) from your console.

Full-scale deployment

As you get started with runZero, we recommend kicking off with our standard deployment plan and adding tasks as needed. The standard deployment plan is broken out into six stages which will help you plan out your requirements, execute the deployment, and optimize your environment based on runZero's best practices.

Planning

This first set of tasks will help your team identify target results, get ahead of potential blockers, and help you avoid misconfigurations within runZero.

Tasks

- Determine your success outcomes using runZero. Some common success outcomes include:
 - [Reducing gaps in asset visibility](#)
 - Key results
 - Scan all assets in days rather than weeks
 - Integrated with all cloud providers and other relevant tools
 - Resources
 - [Achieving RFC1918 coverage](#)
 - [Reducing investigation times](#)
 - Key results
 - Find any asset in your environment in seconds
 - Able to review all services an asset runs in minutes
 - Understanding of potential exposure on new vulnerabilities
 - Resources
 - [Building a complete asset inventory](#)
 - [Search syntax](#)
 - [Training](#)
 - [Reducing asset risk](#)
 - Key results
 - Eliminate misconfigurations
 - Reduce gaps in endpoint protection
 - Reduce gaps in vulnerability scanning
 - Eliminate unmanaged assets
 - Number of unauthorized assets found and removed
 - Resources
 - [Alerting on asset and service changes](#)
 - [Tracking hardware and operating system retirement](#)
 - [Finding gaps in endpoint protection](#)
 - [Finding gaps in vulnerability scanning](#)
- Identify key organizational stakeholders
 - Administrator(s) - who will be setting up runZero?
 - Integration owner(s) - who will provide credentials for each integration?
- All users take the [runZero 101 training](#)

- Administrators take the [runZero 201 training](#)
- Determine whether self-hosting is required ([docs](#) | [video](#))
- Identify known networks and subnets for discovery and other inventory sources ([docs](#) | [video](#))
- Define organizations based on RBAC requirements ([docs](#) | [video](#))
- Determine [Explorer](#) deployment location(s)

Initial configuration

Once you have your plan in place, it's time to execute and run your initial scans. **Please note** that these configuration tasks are in a prioritized order to help you avoid having to reconfigure things down the road.

Tasks

- Deploy self-hosted console (if required) ([docs](#) | [video](#))
- Setup [organizations](#)
- Set up [sites](#), and [define subnets](#) for discovery ([video](#))
 - **Sites do not necessarily correspond to physical locations within runZero.** Sites are used to represent distinct networks that may have overlapping IP space
- Install [Explorer\(s\)](#) ([video](#))
- Run initial scan ([docs](#) | [video](#))

Analysis

Now that you have done some initial discovery, it's time to review the results. Reviewing the results and leveraging our reports will help you expand scan scope, better understand your network, as well as help you identify key issues such as misconfigurations.

Tasks

- [Review results of initial scan](#)
 - Review dashboard ([docs](#) | [video](#))
 - Review the asset inventory ([docs](#) | [video](#))
 - Review the asset detail view ([docs](#) | [video](#))
- Identify risky assets using the Queries library ([docs](#) | [clickthrough](#))
 - [Learn query syntax](#)
 - Apply vulnerability records to queries ([docs](#))
 - Track query results on the dashboard ([clickthrough](#))
- [Track long-term initiatives with Goals](#)
- [Review reporting](#)
 - Identify gaps in scanning ([docs](#) | [video](#))
 - Understand network segmentation ([docs](#) | [video](#))

Advanced configuration / Optimization

After you've done your initial analysis, you will want to optimize your scans and configure integrations to further build your complete asset inventory.

Tasks

- [Configure inbound integration connections](#)
 - [Cloud inventory sources](#)
 - [Endpoint protection](#)
 - [Vulnerability management](#)
 - [Directory services](#)
- [Configure SNMP credentials \(video\)](#)
- Optimize scans by adjusting scan rates and other configurations ([docs](#) | [video](#))
 - See our clickthrough of some key additional configuration options as well

Automation

Now that you have optimized your scans and have analyzed your runZero data, you can automate these tasks to avoid manual effort. You can leverage this automation to run scans on a recurring basis, automate queries, and generate alerting for the team.

Tasks

- Schedule recurring scan [tasks](#) and any inbound [integration tasks](#)
- [Automate queries](#) and [configure alerts](#) to align with use cases ([video](#))
- Configure outbound [integration](#) connections to enrich other IT and security tools
 - [CMDB](#)
 - [SIEM](#)
 - [SOAR](#)
 - **Note:** If you're utilizing a solution that runZero does not offer a standard outbound integration for at this time, be sure to review our [API documentation](#) to learn about how to export runZero data.

Rollout

As your runZero deployment comes to a close, you will want to ensure all users have gone through training and ensure anyone that would get value from runZero has access to the platform.

Tasks

- [Add users](#)
-

- Ensure all users are trained on runZero
 - Training and key documentation
 - [101 user training](#)
 - [201 administrator training](#)
 - [Search](#)
 - [Reporting](#)
 - [Exporting assets](#)
 - [runZero playbooks](#)
- Identify other teams interested in the asset inventory data, such as:
 - [Enterprise security team](#)
 - runZero is typically used by security teams to achieve a complete asset inventory, find gaps in their vulnerability scanning and endpoint protection, as well as discover potential vulnerabilities.
 - [IT Operations team](#)
 - runZero is typically used by IT Operations teams to achieve a complete inventory of all assets across on-premise and cloud-based infrastructure. This allows the team to identify misconfigurations as well as report on assets in the environment by leveraging our searching and reporting capabilities.
 - [Penetration testing team](#)
 - runZero is typically used by penetration testing teams for conducting reconnaissance both internally and externally, identifying vulnerable targets, and finding ways to get to these vulnerable targets by using our reporting and searching capabilities.

Additional Resources

Now that runZero has been deployed and users have been trained on the platform, please review some of our additional resources to help answer questions you might have as well as maximize the value of runZero:

- [runZero playbooks](#)
- [Leveraging the API](#)
- [Glossary](#)
- [runZero FAQs](#)


Getting help

If you need assistance at any point in this process, you can [book a session with a runZero Customer Success Engineer](#) to discuss further.

Types of networks

It is often helpful to use network examples as a starting point for planning your runZero implementation. This document breaks down a few standard network types and provides potential configurations for each. With that being said, every network has nuance, so it's likely there will be some differences for your implementation.

This is a basic overview of how discovery will be done using Explorers and scanners. By default, one Explorer will be deployed with the goal of running discovery on as much of the network as possible. If needed, more Explorers can be added for areas the primary Explorer cannot get to. You can also use a scanner for offline environments where there is no internet connectivity.



SaaS company

Network characteristics

- Flat on-premises network for offices
- Multi-cloud environment
- Hybrid workforce remote and in office

Sample runZero implementation

- Explorers
 - Corporate network Explorer that is able to get all on-premises network and cloud services private connectivity
 - One Explorer per VPC that isn't routable from on-premises networks
- Organizations
 - Single organization
- Sites
 - Primary corporate site
 - One site per VPC

Large corporation

Network characteristics

- Corporate offices with many connected branches
- Multiple M&A transactions happening at any given time, onboarding new offices every year
- Multiple data centers for internal and externally-facing applications

Sample runZero implementation

- Explorers
 - Corporate network Explorer that is able to get all on-premises networks
 - One Explorer per site with low bandwidth or legacy firewalls, proxies, etc.
 - One Explorer for each M&A transaction to understand new risks and keep data segmented
- Organizations
 - Corporate assets
 - One project per M&A transaction until the deal is complete and assets are merged
- Sites
 - Primary corporate site
 - Potentially separate sites if there is overlapping IP space in branch offices

Retail company

Network characteristics

- Flat on-premise network for corporate offices
- Data centers
- 100s of retail locations

Sample runZero implementation

- Explorers
 - Corporate network Explorer that is able to get all on-premise networks
 - One Explorer per retail location with low bandwidth or legacy firewalls, proxies, etc.
 - One Explorer per data center
- Organizations
 - Single organization
- Sites
 - Primary corporate site
 - One per retail site if they have overlapping IP space

Manufacturing company

Network characteristics

- Flat on-premise network for corporate offices
- OT environment is completely disconnected from the internet
- Field service agents VPN in from varying locations

Sample runZero implementation

- Explorers
 - Corporate network Explorer that is able to get all on-premise networks
- Scanners
 - Manufacturing plant that is not connected to the corporate networks
- Organizations
 - Single organization
- Sites
 - Primary corporate site
 - One per OT site if there is overlapping IP space

Telecommunications company

Network characteristics

- Flat on-premise network for corporate offices
- Segmented data centers
- IoT devices scattered across the country
- Field service agents VPN in from varying locations

Sample runZero implementation

- Explorers
 - Corporate network Explorer that is able to get all on-premise networks
 - One for each segmented lab
 - One for each data center
- Organizations
 - Corporate network
 - Production network
 - One per lab
- Sites
 - One primary site per organization

Managed security service provider

Network characteristics

- Many customers spinning up/down in engagements
- Mostly small, flat networks
- Some customers have OT environments that are sensitive

Sample runZero implementation

- Explorers
 - One per customer
- Scanners
 - Only needed for customers with offline networks
- Organizations
 - Project per customer initially
 - Promote project to organization for long-term use
- Sites
 - One per customer but potentially multiple if a customer has overlapping IP space

Academic organization

Network characteristics

- Many buildings/networks spread around the campus with varying connectivity
- Multiple data centers managed by different departments
- Multiple labs for research and development that are disconnected from the rest of the network and the internet

Sample runZero implementation

- Explorers
 - Main network Explorer that is able to get all on-premise networks
 - One Explorer per building with low bandwidth or legacy firewalls, proxies, etc
 - One Explorer per data center
- Scanners
 - One for each disconnected lab
- Organizations
 - Main assets
 - One for each lab
- Sites
 - One primary per organization but potentially multiple if buildings have overlapping IP space

Organizations

Professional Community Platform

An organization represents a distinct entity; this can be your business, a specific department within your business, or one of your customers. All actions, tasks, Explorers, scans, and other objects managed by runZero are tied to specific organizations and isolated from each other.

Your active organization can be switched by using the dropdown selector at the top right of the runZero Console. If your default role is Viewer or higher, you can select **All Organizations** from the dropdown to view the data for all of your organizations in the Dashboard and Inventory. The Queries page also supports the All Organizations view. Pages that are not compatible with the All Organizations view will be hidden until a single organization is selected.

runZero Community Edition is limited to a single organization.

Use cases for organizations

It is common to segment your assets to multiple organizations based on a few common requirements:

- **RBAC:** Users will only be able to see the assets in the organizations or the children of organizations they're added to.
- **Environments:** It is possible you do not want your development, production, or other assets in one pool and would prefer them to be segmented.
- **Customers:** Service providers generally have one organization per customer.

Creating organizations

Organizations can be created, modified, and deactivated by going to the organizations section within the console. Click the Organizations link under Global Settings.

Organization configurations

- **Name:** Name of the organization.
- **Parent organization:** Allows for sub-organizations, limited to three levels of nesting. Permissions will be inherited based on most privilege.
- **Description:** description of the organization.
- **Stale asset expiration:** Sets the number of days before stale assets are removed from the inventory. This applies to online and offline assets not seen within the configured number of days.
- **Offline asset expiration:** Sets the number of days before offline assets are removed from the inventory. This applies to offline assets not seen within the configured number of days.

- **Scan data expiration:** Sets the number of days before completed scans are removed from the platform. This applies to the scan task metadata as well as the raw data files and change reports.

Projects

A project is a special type of organization designed for temporary use. They behave like organizations, and can have sites defined within them. The important difference is that projects cannot be parents of organizations or other projects, automatically become read only after 30 days, and are automatically deleted after 90 days. While organizations allow you to customize expiration times, the expiration times for projects are fixed.

Your runZero license includes the addition of up to five times as many project assets as the total number of licensed live assets. For example, if your license includes 1,000 live assets, you can have up to 5,000 project assets total. The total number of assets in organizations and projects can be seen on your [license information](#) page in the runZero console.

If you decide that you want to keep a project indefinitely, it can be converted into an organization.

Note that organizations support custom [data retention](#) settings that you can edit to your requirements.

Organization and project details

Click on the name of an organization or project to get to the details page. This page includes basic details about the organization or project, as well as its place in the organization hierarchy.

Sites

By default, your account includes a single organization, which itself contains a single site, named Primary. If the only site in an organization is deleted, a replacement will be created automatically. Similarly, if the last organization is removed, a replacement will be created. You can rename organizations and sites at any time.

Every organization has at least one site, but may have multiple sites. A site represents a distinct network segment, usually defined by addressing or accessibility. Sites in runZero do not necessarily correspond to physical sites or locations. Instead, they are used to represent distinct networks that may have overlapping address space. This allows for multiple sites to use the same RFC1918 space, something common in retail, while still being possible to differentiate their assets within the inventory.

Because sites represent separate networks, if you set up two sites and scan the same devices from both, you will end up with two copies of the resulting assets, one set for each site's network.

All analysis actions within runZero occur at the site level. For example, reports such as the switch topology report analyze a single site's devices, so you will likely want to avoid splitting routers and non-router assets into separate sites unless they are truly on separate disconnected networks.

Use cases for sites

For flat networks, where every IP address can reach any other address on the network, a single site is usually enough, and avoids the possibility of accidentally creating duplicate assets by scanning the same devices from multiple sites. Sites are recommended for complicated, sprawling, and highly-segmented environments.

Two circumstances that could lead to multiple sites:

- **Overlapping IP space:** sites will allow you to differentiate identical IPs that are actually different machines if you have overlapping IP space.
- **Highly complex network:** sites would not be required in this case, but they can be used as an organizational tool.

Creating sites

To create a new site, click the **New Site** button on the top of the [sites](#) page.

Site configurations

- **Name:** The name of your site.
- **Description:** The description can help identify the purpose of the site.

- **Default scan scope:** The default scan scope will be pre-populated when creating scans for this site.
- **Default scan exclusions:** The default scan exclusions will be pre-populated when creating scans for this site.
- **Registered subnets:** Registered subnets can be used to automatically tag assets, services, screenshots, and software that fall within each subnet.

Importing and exporting sites

Site configurations can be created or updated based on a CSV file import. [Import your CSV](#) from the sites page. The CSV format should include the following:

```
name,description,scope,exclusion,subnet_ranges,subnet_tags,subnet_descriptions
```

Your site configurations can also be exported as a CSV from the sites page.

Sites and Explorers

Sites can be tied to specific Explorers, which can help limit traffic between low-bandwidth segments. The site configuration allows a default scan scope to be defined, along with an optional list of excluded scan scopes. These fields can be used to set the scan scope for scans of the site.

If you would like to tie an Explorer to a site, navigate to the [Explorers](#) page, click the Explorer you would like to tie to the site, and then click configure. After that, you will see a dropdown with your site options.

Self-hosting runZero

Platform

Background

The self-hosted version of runZero allows you to run the entire platform on-premises or within your own cloud environment. This platform is functionally identical to the hosted service, provides a fully-offline mode, and does not send any inventory data back to runZero.

While self-hosting is less common, here are a few reasons your company might choose to:

- ISO compliance requirement
- Other compliance requirement
- Prefer data on-premise

Self-hosting requires an enterprise license and must be explicitly enabled for your account. Please contact your runZero sales representative for further information.

Requirements

Before you get started, make sure your system meets the following requirements.

Hardware requirements

Recommended production system requirements:

- 12 CPU Cores at 2 GHz or faster
- 1TB of local disk storage
- 128 GB of RAM

Minimum production system requirements:

- 4 CPU Cores at 2 GHz or faster
- 100 GB of local disk storage
- 32 GB of RAM (more for large sites)

Minimum testing system requirements:

- 2 CPU Cores at 2 GHz or faster
- 20 GB of local disk storage
- 16 GB of RAM (more for large sites)

Software requirements

- PostgreSQL 13 or newer (our installer can install this for you)

Supported operating systems

- Ubuntu 18.04 and newer running on x86_64
- Red Hat Enterprise Linux 7.x and newer running on x86_64
- CentOS Linux 7.x and newer running on x86_64
- Oracle Linux 7.x and newer running on x86_64
 - 8.x must be 8.4+ with UEK 5.4+ or kernel 4.18+
 - 7.x must be 7.9+ with UEK 5.4+ or kernel 3.10+
- Debian Linux 9.x and newer running on x86_64

Windows Subsystem for Linux is not supported.

Note about Debian

By default, Debian Linux uses the `su` command instead of `sudo`. It also requires that you use the command `su -` to become root, not just `su`, in order to update the `PATH` to include system administration commands such as `useradd`. If you receive an error during installation saying `executable file not found in $PATH`, this is the most likely reason.

Connectivity

The self-hosted runZero platform requires the following connectivity for online updates:

- Outbound access to **console.runzero.com (13.248.161.247, 76.223.34.198, 2600:9000:a415:cd87:fbe5:476a:3533:69f2, 2600:9000:a716:ee91:85f9:3c9:48c9:59b9)** on TCP port 443 for HTTP over TLS.

The system running the runZero platform should have a static IP address.

The Explorers need to be able to validate their HTTPS connection to the platform via a TLS certificate. By default, the runZero platform installer will set up a self-signed certificate for the console's IP address, and downloaded Explorers will be preconfigured with the appropriate URL and certificate so they can make a verified connection.

Note: Some components of the application still reference the name "Rumble" for backwards compatibility. The documentation will be updated as these are changed.

If instead you plan to use a real certificate from an internal or public CA for your runZero console deployment, you will want to assign the console server a DNS name, and edit `RUNZERO_CONSOLE` in the configuration to contain the fully qualified domain name. The `TLS_CERT` and `TLS_KEY` should be set to point at the appropriate certificate and private key files. The console should then be restarted via `runzeroctl restart`. Newly downloaded Explorers will then be preconfigured with the right URL and certificate.

Because Explorers are preconfigured with the console address during download, if the IP address and/or DNS name are changed at a later date, the Explorers may need to be

redeployed to receive the updated address.

Offline mode

The self-hosted version of runZero has the ability to run in [offline mode](#). In this mode, the runZero update service is not used and [offline updates](#) must be applied manually. Enable this mode if you're in an isolated network or you don't want your self-hosted runZero console to make any connections to the internet. In addition to disabling online updates, offline mode also disables certain DNS probes that could reflect responses to the internet during a scan.

Installation steps

For offline installs please see [offline installation](#).

For installs that use your own database credentials see [Installation with your own PostgreSQL database](#).

Here's what the installation process will do:

- Set up PostgreSQL and create a passworded user.
- Generate TLS certificates for your IP address located in `/etc/runzero/certs`.
- Generate a configuration file at `/etc/runzero/config` and set some defaults.
- Create a systemd service for the runZero platform.
- Create all the necessary cron jobs required for the runzero platform.

Step 1: Download and run the installer

- Go to <https://console.runzero.com/deploy/download/platform>.
- Copy the command directly from the download page and run it in your terminal. This will download the installer and mark it as executable. The download path for the installer is uniquely keyed.

Step 2: Initialize the admin user

After you've installed the runZero platform, you will have access to the runZero CLI `runzeroctl`.

To initialize an admin user, run:

```
runzeroctl initial [email address]
```

Step 3: Log in to your self-hosted console

If everything is set up correctly, you can log in to your console at `https://YourInternalIPAddress`.

Note that you may need to enable HTTPS to pass through the Linux system firewall. Example commands:

Ubuntu Linux: `sudo ufw allow https/tcp`

RHEL/CentOS/Oracle: `sudo firewall-cmd --add-service=https`

To make a `firewall-cmd` change permanent across reboots, run the command a second time with the `--permanent` flag added.

Installation with your own PostgreSQL database

runZero defaults to installing and configuring a PostgreSQL user and database for you. If you would like to provide your own details this option will allow you to override that behavior.

Requirements

- PostgreSQL 13 or newer.
- [Password authentication](#) must be enabled.
- Two extensions are required: [pg_trgm](#) and [uuid-osspl](#). These may be supplied as part of a `contrib` package rather than as part of the main PostgreSQL server install, depending on where you obtain your PostgreSQL packages.

PostgreSQL example to enable extensions and add a database/user

```
CREATE DATABASE rumble;
CREATE USER runzero WITH PASSWORD 'YOURPASSWORD';
GRANT ALL PRIVILEGES ON DATABASE runzero TO rumble;
\connect rumble;
CREATE EXTENSION IF NOT EXISTS pg_trgm;
CREATE EXTENSION IF NOT EXISTS "uuid-osspl";
```

Steps to install the self-hosted runZero platform with your own database credentials

1. Run this install command as root:

```
./runzero-platform-[VERSION]-linux-amd64.bin install --manual-database
```

2. Edit your runZero configuration and add your database details in `/etc/runzero/config`.
The line you need to edit is:

```
DATABASE_URL=postgres://rumble:{DB_PASSWORD}@127.0.0.1:5432/rumble?sslmode=disable
```

Change this to match your credentials. You need to set the user, password, host, port, and database name. Here is the format:

```
DATABASE_URL=postgres://{DB_USER}:{DB_PASSWORD}@{DB_HOST}:{DB_PORT}/{DB_NAME}?{DB_OPTIONS}
```

4. Verify the self-hosted runZero platform can connect to your database with this command. `sudo runzeroctl database verify`
5. Once your database is configured and verified you can restart the self-hosted runZero platform service `sudo systemctl restart runzero-console`

runZero updates

For offline updates please see [CLI update with offline mode](#).

The self-hosted runZero platform must be updated prior to first use. The `runzeroctl` command can be used to download the update and then restart the service after the update is complete.

You can update the platform and scanners at the same time or separately with the [CLI update management commands](#).

Managing users

You can manage users inside your self-hosted runZero platform console at <https://YourInternalIPAddress/team> or via the [runZero CLI](#)

Some things you can manage:

- Adding, deleting, and listing users
- Resetting passwords
- Changing default roles
- Viewing details
- Changing organization roles

CLI service management

Start the runZero service

Starts the runZero platform service.

```
runzeroctl start
```

Stop the runZero service

Stops the runZero platform service.

```
runzeroctl stop
```

Restart the runZero service

Restarts the runZero platform service.

```
runzeroctl restart
```

Install the runZero platform

Install the runZero platform service and all required dependencies such as PostgreSQL. Creates a systemd service, generates cron jobs, and generates a configuration file in `/etc/runzero/config`.

```
runzeroctl install
```

Uninstall the runZero platform

Stop and remove the runZero platform service from systemd and removes the generated cron jobs. This does not remove your PostgreSQL database, and it retains your data.

```
runzeroctl uninstall
```

Purge the runZero platform

Stop and remove the runZero platform service from systemd and removes the generated cron jobs. This will delete your runZero database and remove all the runZero directories `/etc/runzero` and `/opt/runzero`.

```
runzeroctl purge
```

You can uninstall and purge everything except the database and your PostgreSQL settings with this flag:

```
runzeroctl uninstall --purge --ignore-database
```

Run the runZero platform manually

Starts the runZero platform manually. Logs will be written to standard output.

```
runzeroctl server
```

Verify your database is reachable

Attempts to connect to your database using your self-hosted runZero platform configuration. It will either succeed or display an error.

```
runzeroctl database verify
```

CLI update management

Update the runZero platform and scanners

Updates the runZero platform service and runZero scanners. You can use the optional parameter `force` to force the update even if the current installation is the latest version.

```
runzeroctl update [--force]
```

Update the runZero platform

Updates just the runZero platform service. You can use the optional parameter `force` to force the update even if the current install is the latest version.

```
runzeroctl update-platform [--force]
```

Update the runZero scanners

Updates just the runZero scanners. You can use the optional parameter `force` to force the update even if the current installation is the latest version.

```
runzeroctl update-scanner [--force]
```

CLI user management

Create the initial administrator account

Creates the initial admin user for a new installation. You must provide an email address.

```
runzeroctl initial <email>
```

List user accounts

Lists all the users along with their email address, full name, and current roles.

```
runzeroctl user list
```

Add a user account

Creates a new user account under the initial administrator user. You must provide an email address.

```
runzeroctl user add <email>
```

Delete a user account

Deletes a user account. You must provide an email address. This cannot be undone.

```
runzeroctl user delete <email>
```

Get user details

Gets the details for a user account. You must provide an email address. Provides information such as full name, date created, last login IP, last login time, last activity, default organization role, and their current roles. You must provide an email address.

```
runzeroctl user details <email>
```

Set a user role

Sets a user's role to the role provided. Email and role must be provided. The organization is optional. If the organization isn't provided this sets their default role.

```
runzeroctl user set-role <email> [organization name or organization ID]:<role>
```

Reset a user password and MFA

This will generate and apply a new password for the specified user. The password will be printed to the terminal. You must provide an email address. This reset process also clears any associated MFA tokens.

```
runzeroctl user reset <email>
```

CLI organization management

List all organizations

Lists all the organizations by their name and ID.

```
runzeroctl organization list
```

Advanced configuration

The file at `/etc/runzero/config` can be modified to support a wide variety of configurations.

After making changes, apply them by running `runzeroctl restart`.

Email server

runZero uses an SMTP server for user account invitations and notifications. The default configuration assumes that a SMTP server is available on localhost that does not require authentication:

```
SMTP_SERVER=127.0.0.1:25
SMTP_AUTH_METHOD=none
```

runZero will automatically use STARTTLS with plaintext SMTP servers and validate the certificate. In internal environments where the SMTP server is not using a valid TLS certificate, verification can be disabled by setting:

```
SMTP_TLS_NOVERIFY=true
```

Transport-layer TLS (instead of STARTTLS) can be configured with:

```
SMTP_TLS=true
```

If authentication is required, the following three settings should be configured:

```
SMTP_AUTH_METHOD=plain
SMTP_AUTH_USER=YourUsername
SMTP_AUTH_PASS=YourPassword
```

Emails are sent from `noreply@rumble.run` by default, but this can be changed by setting the `FROM_EMAIL` option:

```
FROM_EMAIL=runzero@yourcompany.int
```

Hostname and port

The `RUNZERO_CONSOLE` variable is used for creating inbound links, configuring deployed Explorers, and generating the default self-signed TLS certificate. This setting is how both users and deployed Explorers connect to the platform.

```
RUNZERO_CONSOLE=https://{IP ADDRESS OR HOSTNAME}:443
```

Changing this setting may require regeneration of the TLS certificate and redeployment of Explorers.

```
runzeroctl generate-certificate
```

runZero can be configured to run on a different port with the `CONSOLE_PORT` setting. This port defines where the console listens, but users and Explorers still connect to the `RUNZERO_CONSOLE` value. In most cases this should match the port specified by the `RUNZERO_CONSOLE`.

```
CONSOLE_PORT=443
```

TLS configuration

runZero will generate a self-signed TLS certificate and serve all web requests using HTTP over TLS. The standard configuration uses a self-signed certificate stored in the filesystem:

```
TLS=true
TLS_CERT=/etc/runzero/certs/cert.pem
TLS_KEY=/etc/runzero/certs/key.pem
```

The certificate and key file are PEM encoded and can be replaced with any valid certificate. Please ensure that any new certificate lists the value of `RUNZERO_CONSOLE` in the list of Subject Alternative Names.

Note that the key file must be unencrypted (not password protected). If you have a passphrase set, you can remove it using OpenSSL. For example:

```
# RSA key
openssl rsa -in key.pem -out newkey.pem
# DSA key
openssl dsa -in key.pem -out newkey.pem
```

If a TLS-terminating reverse proxy is used (AWS ELB, nginx, etc), TLS can be disabled at the application level:

```
TLS=false
```

Please note that while the web interface can be accessed over plain HTTP in this scenario, Explorers will refuse to connect to a plain HTTP port, and features like WebAuthn MFA will not work unless the site is accessed through TLS.

Specific TLS versions and ciphers can be configured.

TLS versions are chosen by minimum and maximum:

```
TLS_VERSION_MIN=1.2
TLS_VERSION_MAX=1.3
```

TLS ciphers may be chosen by profile name:

- **default:** A set of strong ciphers, great for most configurations
- **nist80052:** A set of strong ciphers, approved in [NIST 800-52r2](#).

- **nist80052-aes256**: A set of strong ciphers, approved in [NIST 800-52r2](#), restricted to AES-256 variants

Please note that TLS 1.3 ciphers work differently and if a specific set of ciphers is required, both `TLS_VERSION_MIN` and `TLS_VERSION_MAX` should be set to 1.2. For example, to restrict runzero to **only** NIST 800-52r2 approved ciphers using AES-256, the following configuration should be used:

```
TLS_VERSION_MIN=1.2
TLS_VERSION_MAX=1.2
TLS_CIPHERS=nist80052-aes256
```

TLS ciphers may also be chosen using comma-separated list of [cipher constants](#).

Database configuration

runZero uses a PostgreSQL database for all platform data, except for raw scan files, change reports, and images processed from scans. By default, runZero will configure a local PostgreSQL server on the same system, with a random password, and without TLS encryption:

```
DATABASE_URL=postgres://rumble:{DB_PASSWORD}@127.0.0.1:5432/rumble?sslmode=disable
```

If separate database is preferred, any PostgreSQL server running 12.x or newer should work. TLS (`sslmode=require`) should be enabled when a non-local database server is configured.

The default database pool (connection count) can be modified for high throughput environments:

```
DATABASE_POOL_COUNT=50
```

Proxy configuration

runZero makes outbound connections to receive platform updates (in online mode), to connect to third-party APIs, and to delivery webhooks for notifications. If a proxy server is required, it can be configured with this setting:

```
HTTPS_PROXY=host:port
```

Storage configuration

runZero uses local file storage to store raw scan data, change reports, and images retrieved from assets. This storage directory must be owned by the `rumble` user and be mounted below the `/opt/runzero` path.

```
RUNZERO_STORAGE_MODE=local
RUNZERO_STORAGE_PATH=/opt/runzero/storage
```

Files within the storage directory are split up into two groups, assets and scans. The names of these can be changed by setting:

```
ASSET_BUCKET=assets
SCAN_BUCKET=scans
```

To use AWS S3 for file storage instead, the following configuration can be set:

```
RUNZERO_STORAGE_MODE=s3
ASSET_BUCKET=company-runzero-assets
SCAN_BUCKET=company-runzero-scans
```

If a non-AWS backend is used that is compatible with the S3 API, use the same AWS and bucket variables above but override `AWS_REGION` and set the `AWS_ENDPOINT_URL_S3` or `RUNZERO_STORAGE_ENDPOINT` parameter to the endpoint as appropriate. Reach out to runZero support if you run into any issues with the endpoint configuration.

If the S3 buckets are in a different region than the environment, set the `RUNZERO_STORAGE_REGION` to the correct region for the buckets.

If S3 is used, AWS must also be [configured](#), with at least the `AWS_REGION` variable set, even if a non-AWS backend is enabled.

Secret configuration

runZero uses three randomly generated secret tokens to secure the platform. These keys are hexadecimal strings generated by 16 bytes of random. Compatible values can be generated by OpenSSL:

```
$ openssl rand -hex 16
```

The authentication key used for local storage HMAC operations. This key can be rotated as long as the service is restarted afterwards:

```
RUNZERO_STORAGE_KEY_SECRET={SECRET_32_HEX}
```

The session secret key is used to sign and validate browser session cookies. This key can be rotated, but doing so will invalidate all existing web logins:

```
SESSION_SECRET={SECRET_32_HEX}
```

The DB key is used for encryption of sensitive fields (user password bcrypt hashes). This key cannot be rotated, as password authentication will no longer work. If this key is changed, users must reset their password from the command-line or web interface using email before they can sign back in:

```
DB_KEY={SECRET_32_HEX}
```

AWS configuration

The AWS region is required:

```
AWS_REGION=us-east-1
```

The Access Key ID and Secret must be valid and correlated to a user with read-write access to the S3 buckets and read-only access to Secrets Manager.

```
AWS_ACCESS_KEY_ID=AKIA....
```

```
AWS_SECRET_ACCESS_KEY=SECRET....
```

AWS Secrets Manager can be used to retrieve almost any configuration setting at startup. The Secrets Manager entries should match the key names of the configuration file. The secret name can be defined with:

```
AWS_SECRETS_MANAGER_KEY=rumble/production
```

In addition the variables above, most [AWS CLI environment variables](#) are also available for specific tuning.

The location of the Explorer and scanner binaries can be changed with these settings. Note that these should still live under `/opt/runzero` or the service will not be able to load them:

```
RUNZERO_RELEASE_DIR=/opt/runzero/agent/
```

```
RUNZERO_SCANNER_RELEASE_DIR=/opt/runzero/scanner/
```

Content security policy

In the case of a non-standard S3 configuration (or S3-like deployment, such as minio), the Content Security Policy headers need to be configured to allow external image loads.

The `CSP_IMAGES` setting can be used to specify one or more (comma-delimited) external image sources:

```
CSP_IMAGES=https://*.custom-storage-backend.lan
```

In addition to `CSP_IMAGES`, the following additional CSP settings are available:

```
CSP_SCRIPTS=https://*.myscripts.lan
```

```
CSP_FONTS=https://*.myfonts.lan
```

```
CSP_STYLES=https://*.mystyles.lan
```

Logging

The self-hosted runZero console sends its logging output to standard output. On Linux this is picked up by Linux systemd, and stored in the journal where it can be queried with the

journalctl tool. For example, to view the most recent hour of logs, with most recent messages first, you can run the following command:

```
journalctl --unit=runzero-console --since=-1hour --reverse
```

The journalctl output can be piped to a text file to send to runZero support.

The systemd logging subsystem can be configured to send log messages to a local syslog daemon as well as the journal. Once logs are in syslog, they can be forwarded across the network to remote logging servers using the standard syslog protocol. Some Linux distributions, such as RHEL, are set up to forward logs from systemd to syslog by default. Other distributions, such as Ubuntu and Debian, don't include a syslog daemon in their default minimal server installs.

To configure systemd to send logs to syslog, you can use the [ForwardToSyslog option](#) in `/etc/systemd.conf`. Alternatively, some syslog daemons have an option to read the systemd journal; for example, rsyslog has [imjournal](#).

runZero's logs are output in [CEE-enhanced](#) JSON format. This is compatible with rsyslog and [syslog-ng](#), DataDog, and other common logging tools. For rsyslog, the [mmjsonparse](#) module can be used to filter the logs based on individual JSON fields, and [forward them to Elasticsearch](#) or other JSON databases.

The environment variable `LOG_MAX_LENGTH` can be set in runZero's config file to apply a limit to the length of each log line, in bytes of UTF-8 text. A value of 0 means no limit, other values below 480 are treated as 480. runZero will attempt to preserve the most valuable logging fields when truncating log output, and ensure that the result is still valid JSON. Note that the length limit is applied before any additional information systemd or syslog adds to the start of the line.

The environment variable `LOG_FORMAT` can be set to text to disable the CEE and JSON format, and log in plain text. For example:

```
LOG_FORMAT=text  
LOG_MAX_LENGTH=512
```

HTTP timeouts

The self-hosted runZero Console allows you to change the built-in web server's HTTP timeouts. These can be changed through three configuration variables.

You are able to set the HTTP idle, read, and write timeout. The default settings are below and don't usually need to be changed.

```
HTTP_IDLE_TIMEOUT_MINUTES=3  
HTTP_READ_TIMEOUT_MINUTES=60  
HTTP_WRITE_TIMEOUT_MINUTES=720
```

Concurrent processing

The runZero Console can process more than one completed task at a time if the `RUNZERO_CRUNCHER_INSTANCES` option is set to a value greater than 1. Tasks are only processed concurrently when they exist within different organizations, or the tasks are within the same organization, but different sites, and the task itself does not require cross-site asset merging. Most third-party connectors and integrations require cross-site merging and are not able to take advantage of concurrent site processing within the same organization. Please note that the resource requirements for concurrent task processing scale linearly with the specified instance count.

The example below configures the console to process up to four concurrent tasks across all organizations:

```
RUNZERO_CRUNCHER_INSTANCES=4
```

Custom javaScript

The self-hosted runZero Console allows you to include custom arbitrary JavaScript to be executed on the various runZero Console web pages. To use this feature, add your JavaScript to `/opt/runzero/etc/custom.js`. If this is your first time configuring this feature, the `custom.js` file will not exist. You will need to create the file `custom.js` inside the `/opt/runzero/etc` folder.

To enable the feature set the environment variable below to your configuration.

```
RUNZERO_CUSTOM_JS=true
```

HTTP headers

The security headers sent by the runZero Console can be disabled as needed using the following options:

```
# Disable Strict-Transport-Security
```

```
RUNZERO_DISABLE_HSTS=true
```

```
# Disable X-Frame-Options
```

```
RUNZERO_DISABLE_HXFO=true
```

```
# Disable X-Content-Type-Options
```

```
RUNZERO_DISABLE_HXCTO=true
```

```
# Disable X-XSS-Protection
```

```
RUNZERO_DISABLE_HXXP=true
```

```
# Disable Referrer-Policy
```

```
# Disable Content-Security-Policy
RUNZERO_DISABLE_HCSP=true
```

Unofficial CPEs

When runZero can successfully fingerprint an asset's operating system, a CPE will be generated. In cases where the NIST database does not contain an official match, runZero will generate an unofficial CPE by default. This behavior can be disabled by setting `RUNZERO_GENERATE_UNOFFICIAL_CPE` to `false`:

```
RUNZERO_GENERATE_UNOFFICIAL_CPE=false
```

When unofficial CPEs are generated by runZero, they include `r0_unofficial` in the other field of the CPE by default. This value can be changed to any alphanumeric-constrained tag (limited to 32 characters):

```
RUNZERO_UNOFFICIAL_CPE_TAG=custom_unofficial_tag
```

Permissions

The self-hosted platform requires **root** access to install and manage from the command-line.

The platform service (`runzero-console`) runs as **root** and spawns a worker subprocess that runs as the **runzero** user account inside of a chroot environment (`/opt/runzero`). All substantive work happens within this isolated subprocess. Please note that older installations will use `rumble` instead of `runzero` in directory, file, and user names.

The following filesystem locations are used by the self-hosted platform:

/etc/runzero

Path	Owner	Permission	Notes
<code>/etc/runzero</code>	root	0700	Configuration files and certificates
<code>/etc/runzero/config</code>	root	0600	A plain-text configuration file
<code>/etc/runzero/certs</code>	root	0700	A directory containing the TLS certificate and key
<code>/etc/runzero/certs/cert.pm</code>	root	0600	The TLS certificate in PEM format
<code>/etc/runzero/certs/key.pm</code>	root	0600	The TLS certificate private key in PEM format

/opt/runzero

Path	Owner	Permission	Notes
/opt/runzero/tmp	runzero	0755	A temporary directory
/opt/runzero/storage	runzero	0700	Contains asset and scan artifacts
/opt/runzero/console	root	0755	Contains the platform executable
/opt/runzero/console/runzero-console.bin	root	0755	The platform executable
/opt/runzero/agent	root	0755	Contains the Explorer binaries
/opt/runzero/agent/runzero-agent-*	root	0755	The Explorer binaries
/opt/runzero/scanner	root	0755	Contains the Scanner binaries
/opt/runzero/agent/runzero-scanner-*	root	0755	The CLI Scanner binaries
/opt/runzero/proc	root	0755	Contains copies of system /proc files
/opt/runzero/proc/cpuinfo	root	0644	A copy of /proc/cpuinfo
/opt/runzero/proc/meminfo	root	0644	A copy of /proc/meminfo
/opt/runzero/proc/version	root	0644	A copy of /proc/version
/opt/runzero/etc	root	0755	Contains copies of system files
/opt/runzero/etc/resolv.conf	root	0644	A copy of /etc/resolv.conf
/opt/runzero/etc/hosts	root	0644	A copy of /etc/hosts
/opt/runzero/etc/ca-certificates.crt	root	0644	A copy of the system root CA store
/opt/runzero/etc/runzero	runzero	0700	Contains instance identifiers
/opt/runzero/etc/runzero/cruncher.id	runzero	0700	A unique ID to identify the cruncher instance
/opt/runzero/etc/runzero/hub.id	runzero	0700	A unique ID to identify the hub instance

Backup and restoration

Your runZero installation and data can be backed up and restored to preserve your configuration.

runZero data backup

A backup of a self-hosted installation can be obtained by archiving the file system and database.

The file system archive includes the following paths:

- /etc/runzero
- /opt/runzero
- /lib/systemd/system/runzero-console.service
- /etc/systemd/system/multi-user.target.wants/runzero-console.service
- /usr/bin/runzeroctl

A sample file system backup command is:

```
# tar zcvf runzero-backup-fs.tar.gz /etc/runzero/ /opt/runzero/ \
/lib/systemd/system/runzero-console.service \
/etc/systemd/system/multi-user.target.wants/runzero-console.service \
/usr/bin/runzeroctl
```

The PostgreSQL database must be backed up separately. A sample command is shown below:

```
# sudo su - postgres
$ pg_dumpall -f runzero.sql && gzip runzero.sql
```

runZero data restoration

To restore the runZero install, follow these steps.

1. Stop any running runZero service:

```
# runzeroctl stop
```

2. Unpack the filesystem archive:

```
# tar -C / -zxvf /path/to/runzero-backup-fs.tar.gz
```

3. Restore the PostgreSQL database:

```
# sudo su - postgres
$ dropdb rumble; gzip -dc rumble.sql.gz | psql
```

- Restart the runZero service:

```
sudo systemctl restart runzero-console
```

Support and debugging

The `runzeroctl` command includes a debugging tool which can collect diagnostics from your server and assemble them into a zip file which you can send to support. If you are asked to do this, the command is:

```
runzeroctl diagnostics run-script
```

Data is written to `/opt/runzero/collector`.

Alternatively you can save a copy of the script to the current directory so that you can examine it before running it:

```
runzeroctl diagnostics write-script
```

Manual migrations

Starting with version `4.0.240221.0`, the runZero self-hosted upgrade process will run migrations before restarting the service. If you are running an older version of the software and would like to prevent downtime during the upgrade of a single-node self-hosted installation, the following steps can be used:

- Obtain the self-hosted download link from the [runZero SaaS](#).
- Download this file manually to your self-hosted systems:

```
$ curl -o platform.bin https://console.runzero.com/...../runzero-platform.bin
```

- Mark this file as executable and run it with the `task db:migrate` parameter:

```
$ chmod u+x platform.bin; ./platform.bin task db:migrate
```

- Once the migrations are done, install the update as usual:

```
$ runzeroctl update
```

Offline mode configuration

Offline installation

The self-hosted runZero platform comes with a few options for your installation. You can utilize these options by adding flags to the `install` command. The current flags available are --

offline, `--distro-packages-only`, and `--postgres-rpm-directory`.

`--offline`

Configures the installation for operation without internet access. Network access may still be necessary during the installation to acquire dependencies. The installer uses the upstream postgresql.org packages by default and this can be disabled by specifying the options below.

`--distro-packages-only`

Install the platform without using third-party repositories (not available on RHEL/CentOS 7). The install will try to acquire the PostgreSQL package from the configured operating system repositories (local or remote). If no repository is reachable, these packages can also be specified as a directory on the file system using the option below.

`--postgres-rpm-directory [directory]`

Install using supplied PostgreSQL RPM files (requires RHEL or CentOS). The installer will use the RPMs in the specified directory to satisfy the PostgreSQL dependencies.

PostgreSQL RPMs required for `--postgres-rpm-directory`

RHEL/CentOS 8 RPMs can be found [here](#).

RHEL/CentOS 7 RPMs can be found [here](#).

There are four RPMs we require for installing PostgreSQL 13:

- postgresql13
- postgresql13-server
- postgresql13-contrib
- postgresql13-libs

Example install commands for offline mode

The first step is to download the [runZero platform](#)

RHEL/CentOS 7 or 8

1. Download the required RPMs above and store them in a directory.
2. Run this install command as root:

```
./runzero-platform-[VERSION]-linux-amd64.bin install --offline
--postgres-rpm-directory [RPM_DIRECTORY]
```

Ubuntu 18.04+, Debian 10+, or RHEL/CentOS 8

1. Run this install command as root.

```
./runzero-platform-[VERSION]-linux-amd64.bin install --offline --distro-packages-only
```

Remember to use `su -` on Debian, not just `su`.

Note: Some components of the application still reference the name "Rumble" for backwards compatibility. The documentation will be updated as these are changed.

Enabling offline mode for existing installs

- Open `/etc/runzero/config` with an editor of your choice.
- Look for `OFFLINE=` and change it to `OFFLINE=true`.
- Restart the runZero service `runzeroctl restart`.

CLI update with offline mode

The self-hosted runZero platform must be updated prior to first use.

Update the runZero platform and scanners with an offline update

- Go to <https://console.runzero.com/deploy/download/platform>.
- Copy the command directly from the download page and run it in your terminal, or you can use one of the following commands to update using the zip archive you downloaded.


```
runzeroctl update runzero-platform-update-[VERSION].zip
```

```
runzeroctl update --offline --zip-file-path runzero-platform-update-[VERSION].zip
```

You will need to change the version to match the zip archive you downloaded.

High-availability configuration

Self-hosted installations of runZero can be configured for high-availability. For this configuration, a load balancer is used to direct traffic to multiple console servers, which use a shared PostgreSQL cluster and storage backend. The following diagram illustrates an example architecture of a high-availability installation using AWS with two availability zones.



In this diagram, an application load balancer (ALB) is terminating TLS and pointing to a target group that consists of two runZero servers, each in separate availability zones. The runZero servers use a multi-availability-zone PostgreSQL RDS instance, also configured for the same two availability zones, and both servers point to the same set of S3 storage buckets.

Installation steps

These are the general steps for installing and configuring a high-availability instance of runZero:

1. [Deploy an application load balancer.](#)
2. [Install and configure the first server node.](#)
3. [Install and configure subsequent server nodes.](#)

Deploy a load balancer with TLS termination

In order to provide a high-availability interface to the runZero cluster, an application load balancer should be used that receives TLS connections and forwards the HTTP requests to the available servers. This load balancer should be highly-available on its own. Any load balancer that can proxy large HTTP requests and websockets should be supported. The load balancer should add an X-Forwarded-For header to the HTTP requests sent onward to the runZero servers. The load balancer should also support detection of unhealthy nodes in order to handle automatic failover. The /health endpoint on the runZero servers can be used to implement health checks on the load balancer.

Prepare the first runZero server node

The general steps for installing and configuring the first runZero node are:

1. Install the runZero platform.
2. Configure the console URL and TLS/XFF settings.

3. Configure the database settings.
4. Configure the shared storage settings.
5. Verify your configuration.
6. Create the initial user account.
7. Add the runZero server to the load balancer and verify the connection.

Install the runZero platform

Note: Some components of the application still reference the name "Rumble" for backwards compatibility. The documentation will be updated as these are changed.

1. Go to [platform download page](#) in the runZero Console.
2. Copy the URL to the platform installer from the download page. The download path for the installer is uniquely keyed to your license.
3. Download the install using Wget or cURL:

- `wget -O runzero-platform.bin https://console.runzero.com/download/platform-combined/<UNIQUE KEY>/62e41615/runzero-platform-vX.X.X-linux-amd64.bin`
- `curl -o runzero-platform.bin https://console.runzero.com/download/platform-combined/<UNIQUE KEY>/62e41615/runzero-platform-vX.X.X-linux-amd64.bin`

4. Run the following command to make the installer file executable:

```
# chmod u+x runzero-platform.bin
```

5. Run the install command with the manual database option:

```
# ./runzero-platform.bin install --manual-database
```

Configure the console URL and TLS/XFF

To configure your self-hosted installation for high-availability, the settings below should be configured in `/etc/runzero/config`, starting with the console URL and the TLS/XFF settings.

- Set `RUNZERO_CONSOLE` variable to the hostname or IP address of the load balancer:

```
RUNZERO_CONSOLE=https://{IP ADDRESS OR HOSTNAME}:443
```

- Disable TLS and configure the server to trust X-Forwarded-For headers with the following two lines:

```
TLS=false
RUNZERO_TRUST_XFF=true
```

Configure the database

The next step is to configure the database to be used by runZero.

- To do this, edit the `DATABASE_URL` value to match this format:

```
DATABASE_URL=postgres://{DB_USER}:{DB_PASSWORD}@{DB_HOST}:{DB_PORT}/{DB_NAME}?{DB_OPTIONS}
```

An example DATABASE_URL for AWS RDS will look like:

```
DATABASE_URL=postgres://rumbler:[password]@rumbledb.rds-id.us-east-1.rds.amazonaws.com:5432/rumbler
```

To use a replica for read-only queries, specify DATABASE_REPLICA_URL with the same syntax. For RDS Aurora, specify the read-only endpoint as the replica URL. If no replicas are in use, only the DATABASE_URL is required, and this should point to the writer endpoint for RDS Aurora instances.

If a different type of PostgreSQL clustering is in use, make sure that the DATABASE_URL points to the highly-available endpoint for the active writer.

Configure shared storage

Next, configure the shared storage backend. This is used for storing asset data, raw scan data, and reports. runZero supports both mounted NFS and object storage.

- To use NFS, mount a read-write NFS share to /opt/runzero/storage and specify local storage mode with the following configuration options. The bucket names will be created under the NFS mount base directory.

```
RUNZERO_STORAGE_MODE=local
RUNZERO_STORAGE_PATH=/opt/runzero/storage
ASSET_BUCKET=runzero-assets
SCAN_BUCKET=runzero-scans
```

- To use AWS S3 as the shared storage, either configure an IAM instance role with read-write access to two S3 buckets, or specify the AWS credentials in the configuration file.
 - To specify AWS credentials specifically for runZero, use the following syntax:

```
AWS_ACCESS_KEY_ID=AKIA...
AWS_SECRET_ACCESS_KEY=[SECRET]
```

- Specify that the S3 storage mechanism should be used and set the name of the two buckets:

```
AWS_REGION=<your-region>
RUNZERO_STORAGE_MODE=s3
ASSET_BUCKET=<company>-runzero-assets
SCAN_BUCKET=<company>-runzero-scans
```

If a non-AWS backend is used that is compatible with the S3 API, use the same AWS and bucket variables above but override AWS_REGION and set the AWS_ENDPOINT_URL_S3 parameter to the endpoint as appropriate. Reach out to runZero support if you run into any issues with the endpoint configuration.

In addition the variables above, most [AWS CLI environment variables](#) are also available for specific tuning.

Verify the configuration

1. After saving your configuration file, use this command to verify the self-hosted runZero platform can connect to your database:

```
sudo runzeroctl database verify
```

2. Once your settings are configured and verified you can restart the self-hosted runZero platform service:

```
sudo systemctl restart runzero-console
```

Create the initial user account

- Use the `runzeroctl initial <your-email>` command to create the initial user account.

Add the runZero server to the load balancer and verify the connection

- Point the load balancer to the new runZero server's IP on port 80 (unless modified).
- Browse to the load balancer URL over HTTPS using the configured DNS name.
- Verify that the user account configured previously can authenticate correctly.

Installing and configuring subsequent nodes

- Download and install the runZero platform binary on each node as above with the `--manual-database` option specified.
- Copy the configuration file from the first node to `/etc/runzero/config` on each subsequent node.
- After saving your configuration file, use this command to verify that the self-hosted runZero platform can connect to your database:

```
sudo runzeroctl database verify
```

- Restart the platform service by running:

```
sudo systemctl restart runzero-console
```

runZero nodes should be spread across multiple availability zones and added to the load balancer's target group. Verify that requests are reaching each individual node by reviewing the syslog entries on that node.

AWS advanced configuration

AWS Secrets Manager can be used to retrieve configuration settings at startup. The Secrets Manager entries should match the key names of the configuration file. The secret name can be defined with:

`AWS_SECRETS_MANAGER_KEY=rumble/production`

To ensure that runZero can access the Secrets Manager, ensure that an `AWS_REGION` is set in the configuration file, and that the instance has an IAM instance role with permission to read the Secrets Manager key or the AWS credentials have been specified directly in the configuration file. With a Secrets Manager configuration, only the AWS options need to be present in the configuration file and the rest of the settings will be loaded from the Secrets Manager key. This will require copying and defining various settings from the generated configuration file to the Secrets Manager key-value pairs.

Note that many settings need to be identical across the runZero cluster for web requests and authentication to work correctly (including session secrets, DB encryption key, and the console URL).

Self-hosted troubleshooting

The runZero console includes a diagnostics collection script inspired by the need to troubleshoot a self-hosted environment. Collecting the necessary performance statistics, log files, system configuration, and profile debug capture was difficult for customers since there are many different commands and files involved. After checking permissions and dependencies, the diagnostic script gathers and compresses troubleshooting data into a convenient archive file that can be provided to runZero support. The script provides a consistent method for data collection whether it be from a system running your self-hosted console or an Explorer.

Requirements

The diagnostic script must be run from a Linux system. Additionally, the following dependencies must be installed for the script to run successfully:

Dependency	To install (Ubuntu)	To install (RHEL, CentOS, Fedora, Oracle Linux)
zip	<code>apt install zip</code>	<code>dnf install zip</code>
unzip	<code>apt install unzip</code>	<code>dnf install unzip</code>
lsscsi	<code>apt install lsscsi</code>	<code>dnf install lsscsi</code>
iostat	<code>apt install sysstat</code>	<code>dnf install sysstat</code>
curl	<code>apt install curl</code>	<code>dnf install curl</code>

The script must be run as `root` or with `sudo` in order to gather information from the system files and the `/opt/runzero` and `/etc/runzero` directories.

Running the script

Using runzeroctl

To run the script via the `runzeroctl` command, use:

```
runzeroctl diagnostics run-script
```

Output will be placed in `/opt/runzero/collector`.

Manually

To run the script manually, use the command below to write a copy of the script to the current directory:

```
runzeroctl diagnostics write-script
```

The script will be written as `runZero-collector.sh`. To run it:

```
sudo ./runZero-collector.sh
```

The script also has a `--help` option.

This script will perform the following tasks:

- Check for the required permissions.
- Check for the dependency commands (`lsscsi / zip / unzip / iostat / curl`).
- Check that the operating system is Linux.
- Check that either a runZero Explorer or runZero console is installed.
- Collect system configuration information.
- Collect system performance statistics.
- Collect database configuration information.
- Collect some basic statistics from the database.
- Collect runZero debug information from the console or log files from the Explorer.
- Create a zip file of the data collected and the log from the script.

Logging

The script will log all actions in `/tmp` and will provide the file name. Upon completion the log file is moved to the `/opt/runzero/collector` directory. This directory does not exist with the installation of a runZero Explorer or runZero console. The script will create the directory if it doesn't exist and place all files into that directory.

Data retention

runZero allows the data retention periods to be configured at the organization level. The organization settings page provides three ways to control how runZero manages your asset and scan data. Data expiration is processed as a nightly batch job based on the current settings for each organization in your account.

By default, data is retained for up to one year in the runZero Platform. This can be increased to up to three years. The runZero Community edition is limited to 30 days of retention.

Stale asset expiration

Stale asset expiration allows you to make sure that out-of-date information gets purged from runZero. Assets which have not been seen in the specified number of days are automatically deleted. The automatic deletion task runs once per day.

The stale asset expiration limit applies to both online and offline assets. For example, if an asset was online when last scanned, but hasn't been scanned (or imported via a connector) in the specified number of days, then the information is stale and the asset will be deleted. It's like running an asset search query of `last_seen:>Xdays` (where X is replaced by a number), and deleting the resulting assets.

You can disable stale asset expiration by setting the number of days to zero.

Offline asset expiration

Offline asset expiration allows you to clear out information about assets that no longer exist on the network. Assets which are offline and have not been seen in the specified number of days are automatically deleted daily. Online assets are not touched. The automatic deletion task runs once per day.

This feature is particularly useful for guest wireless networks, or networks with short DHCP leases. It's like running an asset search query of `last_seen:>Xdays AND alive:false` (where X is replaced by a number), and deleting the resulting assets.

You can disable offline asset expiration by setting the number of days to zero.

Scan data expiration

This setting controls how long runZero will retain the raw data and metadata for Scan and Import tasks. If your organization has a maximum data retention policy for cloud services, this setting can be used to force the removal of any data older than a given number of days. The default setting of 365 preserves one year of historical scan data.

Event records

Account-level event records (found under Alerts → Events) are always retained for one year. If you need to keep copies of these for longer than one year, these records can be exported and archived through the UX and API.

Managing access

runZero supports multiple concurrent users with a variety of roles. To add a team member, access the [Your Team](#) page, and use the *Invite User* button to send an invitation.

The **Your team** menu entry has four submenus. The first, *Users*, shows all users in the current client account. The second tab, *Groups*, lists the [user groups](#) available; the groups define the access and permissions users have within each organization. The third, *Restricted*, lists users who only have access to a single organization. The last tab lists users who have access to the current organization, as selected from the organization selector at the top of the screen.

User details page

From the Users page, click on a user to view their details page. The user details include a list of their effective access to each organization. While editing user permissions, the explicit assigned roles are displayed. View the user details page to understand effective access, even where explicit roles are not assigned.

Default roles

runZero allows roles to be defined per-user on both a default and per-organization basis. The standard privilege levels are administrator, user, billing, annotator, viewer, and no access. There is also a superuser role available to manage global settings. Where a per-organization role is defined, access will be provisioned based on most privilege.

Superusers

The first user created within the runZero console is considered a superuser. This role allows management of global settings like subscriptions and SSO parameters, and is shown as an access level of “everything”.

If you are a superuser, you can promote someone else to be a superuser. To do this, check the row listing them, and click the *Promote to superuser* button.

If you are using [SSO authentication](#), you should configure at least one superuser with a strong password and MFA that can be used as a backup if SSO settings need to be changed in the future.

We strongly recommend having more than one superuser, particularly if you are using MFA. That way if an MFA token is lost or a superuser leaves your organization, another superuser can fix the problem.

Administrators

Administrators can modify any aspect of an organization and have the unique ability to permanently delete bulk data, create additional organizations, and reset settings for other users.

Users

Users have full access to an organization and can update sites, modify assets, schedule scans, and generally use most functionality. Users are not permitted to reset other users' security credentials, bulk delete data, or delete an organization.

Billing

Billing users are unable to see any asset data, but can manage the licensing, billing, and entity settings for the account.

Annotators

Annotators have the same permissions as a viewer, except they have the ability to add tags to assets. Annotators do not have any other write-access within an organization, so they are unable to modify or remove existing tags. Modifications to existing tags must be made by a runZero user or administrator.

Viewers

Viewers have read-only access to an organization. This includes all inventory data and reports. Viewers are not allowed to interact with tasks, modify settings, or update assets. Viewers may not download the command-line runZero scanner or install runZero Explorers, and they do not have access to view API tokens or export tokens.

No Access

Accounts with no access, which is set in the global role, are limited to those organizations where they have been granted access. If no organizations are allowed, the user is limited to managing their own account settings.

These accounts can only see other users that share their authorized organizations. The no access global role can be used to create a single-organization user, such as a customer or third-party that needs access to the inventory for a specific organization. For consulting use cases, a single-organization user is a way to provide clients with visibility into their environment at no additional cost.

Account settings

The *Account* page is available to superusers. It contains settings which apply to all users and organizations within the account.

Single-sign on (SSO)

runZero supports the [implementation of SSO](#) through SAML2. If you use a SAML2-compatible single sign-on (SSO) implementation, the [SSO Settings](#) page can be used to configure an Identity Provider (IdP) and allow permitted users to login to the runZero console.

Multi-factor authentication (MFA)

runZero supports multi-factor authentication, also known as two-factor authentication or 2FA. Physical hardware keys such as [Google TitanKey](#) and [Yubico YubiKey](#) are supported via the [WebAuthn](#) standard.

You can configure MFA policies for your account via the [Account settings page](#). If multi-factor authentication is required, users who do not have an MFA token set up will be required to set one up when they next log in. You can choose between requiring this for all users, or only requiring it for non-SSO users. The latter option is useful if your SSO server enforces MFA use.

Once a user registers one or more MFA tokens, they will be required to use one of the tokens every time they log in.

Note that **changing the account settings to not require MFA will not alter the MFA status of existing accounts**. Existing accounts will keep any existing MFA tokens they have registered, and will still be required to use one to log in. To disable MFA for a user, the user must clear the MFA token registration. To do this, they can go to their [user settings page](#) and click the red “Unlink” text next to the token ID in the bottom right.

Disabling support access

If you check the box labeled *Disable support access to your account*, runZero support staff will not be allowed to switch to your account.

If you choose to disable support access, this may make it harder for runZero support to answer any questions you have. In some cases we may need you to turn support access back on so that we can help you.

Idle times and login duration

You can set a maximum idle session time in minutes. If set, users whose web browsers don't access runZero for the specified time period will be considered idle, and logged out.

You can also specify a maximum login duration. If set, users will be forced to log in again regularly, at least once every specified period.

Account API keys

The Account API is a REST API which allows account-level operations such as adding and removing organizations and sites, adding users, and accessing the system event log. The *Generate API Key* button on the Account page can be used to generate a token which will allow access to the Account API.

License information

The *License* tab shows information about your runZero software license, including how many assets you are licensed for, how many assets you have across all organizations, and when your license renews.

Entity information

The *Entity* tab allows you to update information about the legal entity runZero is licensed to. You should ensure that this information is kept up-to-date if your company changes name or location, as we use the information to calculate taxes and ensure compliance with appropriate regulations.

Managing user groups

User groups help streamline the management of users who need the same set of permissions. A user group explicitly sets the organizational [role](#) for users, which determines the tasks they can perform within each organization. You can assign roles at a per-organization level or assign a single role across all organizations. Single sign-on settings can also be applied to groups through [SSO group mappings](#).

What happens if there are conflicting permissions?

runZero will always grant the role with the highest permissions level. For example, let's say an account has a viewer role for all organizations, but they've been added to a user group that has a user role for all organizations. This user will now have user-level permissions for all organizations. If the user group expires, the user's role reverts back to their account-level role.

User groups can also have an optional expiration date, which sets time-bound access to organizations within runZero for specific users. When the expiration date elapses, the user reverts back to their account-level permissions. If no expiration date is set, the user group settings will be persistent.

Creating user groups

runZero administrators and super users can create user groups.

1. Go to [Your team > Groups](#) and click Add Group.
2. Enter a name for the user group.
3. Choose the default role you'd like to assign to the group. This setting will establish the access level for every organization you have.
4. Set the per-organization roles, if you need to provide different access levels to specific organizations.
5. Set an expiration date for the user group, if you need to time-bound the permissions. When the expiration date elapses, the role for users part of the group will revert back to their user-level permissions. Otherwise, if you don't specify an expiration date, the user group will be persistent.
6. Go to the **Add users** tab and search for the users you want to add to the group. You can search by **username** or **email**.
7. When you are done adding users, save the user group. The user group will be listed on the [Groups page](#).

Adding users to user groups

To add users to multiple groups at the same time, you can use the [Edit group membership](#) button on the [Users page](#). The [Edit group membership](#) window will list every user group each

user is currently in. Making changes from the Edit group membership window will apply to all users you have selected. Only runZero administrators and superusers can add users to user groups.

1. Go to the [Users page](#).
2. Select the users you want to add to a group.
3. Click the **Edit group membership** button.
4. Choose the user groups you want to add the users to.
5. Save your changes.

Setting an expiration date for a user group

runZero administrators and super users can set an expiration date for a user group.

1. Go to [Your team > Groups](#).
2. Find the user group you want to assign an expiration date. Click the name to open the config page.
3. Set an expiration date for the user group. After the expiration date, the user's role will revert back to their account-level permissions.
4. Save the user group. You'll be able to see the user group's expiration date from the [Groups page](#).

Viewing users in a user group

1. Go to [Your team > Groups](#).
2. Find the Users column in the User groups table, which shows the user count for the group.
3. Click the user count to query and display the users assigned to the group.

Viewing user groups assigned to a user

1. Go to [Your team > Users](#).
2. In the Groups column for each user is the number of groups that user is a member of.
3. Click on a number to show the corresponding list of groups.

Removing users from a user group

runZero administrators and super users can remove users from a user group.

1. Go to [Your team > Groups](#).
2. Find the user group you want to modify. Click the name to open the config page.
3. Go to the Users tab.
4. Remove the users you no longer want assigned to the group. Their permissions will revert back to their account-level ones.
5. Save your changes.

Deleting user groups

runZero administrators and super users can delete user groups.

1. Go to [Your team > User groups](#).
2. Select the user group you want to delete.
3. Click the **Delete** button.
4. Confirm you want to delete the user group.

Searching for users and user groups

When you are on [Users page](#) or [Groups page](#), you can use the following keywords to search in the table:

Keyword	Description	Example
id	User's ID.	id:123456789
name	User's name.	name:john
expires_at	Time or date the user group expires.	expires_at:>2weeks
created_at	Time or date the user group was created.	created_at:>2weeks
updated_at	Time or date the user group was last updated.	updated_at:>1year
has_expiration	Whether the group has an expiration date.	has_expiration:true
created_by_id	ID of user who created the user group.	created_by_id:123456789
created_by_email	Email of the user who created the user group.	created_by_email:user@example.com
group_id	The user group ID.	group_id:123456789
group_name	The user group's name.	group_name:group1

Group IDs can be found in the URL for the group config page
<https://console.runzero.com/groups/<groupid>/edit>.

The `group_id` keyword is only available for the users table; for the groups table, use `id`.

Bulk importing users

Professional

Community

Platform

Instead of manually adding users one at a time, runZero administrators can add multiple users via bulk import. To bulk import users, you will need to create a CSV (comma separated values) file that contains the user information, such as their first name, last name, email, role, and organizational access.

Bulk imports will only add new users; it will not update existing users. If the file contains a user that already exists in the system, the import will not complete. You'll need to remove all duplicate users from your CSV file and import the file again.

Here are the high-level steps for bulk-importing users:

1. [Create a CSV file](#) with your user information.
2. [Import the CSV file](#) into runZero.
3. [Verify users have registered](#) their accounts.

Creating the CSV file for importing users

To create the CSV file, you can use the provided [template](#) or create a file from scratch. You'll need to follow the guidelines and include the fields outlined below.

- The first row of the CSV file is the header row and it must be included in the file. It contains the required fields needed for importing users, and those fields must follow a specific order (`first_name`, `last_name`, `email`, `all_orgs_role`, `org_roles`).
- runZero requires the `first_name`, `last_name`, and `email` fields. The organization-level fields, `all_orgs_role` and `org_roles`, are optional.
- To define the `org_roles` field, use the following format: `orgName=role`. You can use pipes to separate multiple definitions.
- Excluding the `all_orgs_role` field from the CSV will result in users being assigned a no-access role in all organizations.
- Excluding the `orgs_role` field from the CSV will result in users being assigned the role defined in the `all_orgs_role` field. If both fields are excluded from the file, all users will be assigned the no-access role. They will need to contact the runZero administrator to modify their account access.

User fields for the CSV file

The following table lists the fields you can include in your CSV file:

Field	Description	Example
<code>first_name</code>	The user's first name.	sarah
<code>last_name</code>	The user's last name.	smith

email	The user's email address.	ssmith@company.org
all_orgs_role	The role assigned to the user for all organizations. This is optional.	user
org_roles	The role assigned to the user for a specific site. Use pipes to separate the organizations. This field is optional.	org1=user

Your resulting CSV file should look like the following example:

```
first_name,last_name, email,all_orgs_role,orgs_role
Sarah,Smith,ssmith@company.org,user,org1=annotator|org2=annotator
```

Importing users into runZero

1. Go to your [Team Import](#) page.
2. Choose the CSV file that contains your user information and upload it to runZero.
3. Import the file.

runZero will alert you if there are issues with your file. For example, if the file contains an existing user, the import will not complete. You'll need to remove all duplicate users from your file and import it again.

Verifying users have registered

After you add new users via import, each user receives an email invitation to join your team. The email contains a link to register their email address, which prompts them to enter their name and password for their account. After they create an account, they'll be able to log in to the runZero console.

You can visit the [Teams](#) page to find the status and last login for each user. Users who have completed the registration process will show an Activated status.

Managing external users

Professional Community Platform

You can invite external users to join your runZero instance and view the organizational data available to them. The ability to add external users is useful for consultants, value-added resellers, and managed service providers who want to be able to share data from runZero with external partners and clients.

If you are a superuser, you can invite any user who has a runZero account to join your account. When you invite the user, they will receive an invitation to join your account via email. After they accept the invitation and log in to runZero, they will see a new menu next to their organization switcher that lists all the clients they can access. Your client name will display in the list.

Note: If the person you want to invite does not have a runZero client account, they will need to sign up for one before you can invite them as an external user. A free runZero Community Edition account can be invited as an external user.

External users are only viewable from the [External users page](#), separate from your other user lists. These users can use their same credentials to log in and switch between clients and the organizations they've been granted access to.

Inviting an external user into your account

1. Go to **Your team > External** and click [Invite external users](#).
2. Assign the default role you want to assign the user for all organizations.
3. Enter the email address for the user. The user must have an existing runZero account.
4. Customize access for the user at the per-organization level. runZero grants the highest permission levels, when there is a conflict between the global and per-organization permissions.
5. Add a custom email subject or body, if needed. runZero uses a default email template that informs the user you have invited them to your account and provides a link to activate the invitation. The email sender will be your name and the subject line will be "Invitation to join [client] by [your name]."
6. Send the email.

After they activate their invitation, they will be able to log in to runZero, switch clients, and view your organizational data.

Switching clients and organizations as an external user

External users who have access to another account will see a client switcher, or client dropdown menu, located next to their organization switcher at the top right of the runZero web interface. Changing the client will change the organizations you can access and view.

The account that granted you access has configured the organizations and permissions you have within the client account.

Removing an external user from your account

Removing an external user from your account will remove their access to your data. The user will no longer be able to choose your client. To remove an external user, go to the [External users page](#), select the user account you want to remove, and click **Remove user account**. Confirm you want to remove the user. They will be reverted to their primary client account and will no longer have the option to access your runZero data.

Implementing SSO

If you use a SAML2-compatible single sign-on (SSO) implementation, the [SSO Settings](#) page can be used to configure an SSO Identity Provider (IdP) and allow permitted users to login to the runZero console.

runZero's SSO implementation is designed to work with common SAML providers with minimal configuration, but there are a few requirements:

- Your users need to authenticate to a single domain such as `example.com`, not to multiple domains or a domain with many subdomains.
- The domain name needs to be configured in the SSO identity provider settings in runZero. This is true even for self-hosted runZero deployments.
- Your SAML IdP should provide something that looks like an email address in the `NameID` parameter. It doesn't need to be a valid email address, but it should be a unique value that has the same syntax as an email address (`user@example.com`).
- If the `NameID` does not look like an email address, runZero will check the fields `email`, `user.email`, `emailaddress` and `email address` for a suitable ID.
- runZero will check for the user's full name in the fields `name`, `gecos`, `user.name` and `displayname`. If no full name field is found, runZero will proceed to check for a first name in `first_name`, `firstname`, `given_name`, `user.firstname`, `givenname` or `first name`; and for a last name in `last_name`, `lastname`, `family_name`, `user.lastname`, `surname`, `sn`, or `last name`. These attributes are case insensitive.

Note that you must be a [superuser](#) to manage runZero SSO settings.

Note: If a user first registers using SSO, they are marked as an SSO-only user, and cannot bypass SSO by setting up a regular password. For this reason, we strongly recommend that you set up and keep a non-SSO superuser account, so that you can update the settings if SSO stops working for any reason.

Specific SSO providers

You can refer to the [Azure AD](#) or [Okta](#) pages for details on configuring these providers. The basic steps for configuration are:

1. Add runZero as an application.
2. Set up SSO in runZero.
3. Provision users to the runZero app.

For other SSO providers, the following information will help you to configure things.

Identity provider settings

To get to the identity provider settings, choose *Your team* from the left navigator, then click the SSO Settings button at the top right of the page.

The identity provider settings form is where you enter information about your SSO identity provider.

The easiest way to configure SSO is to use XML metadata, if your SSO service provides it. This will usually require that you set up runZero as an application on your SSO provider, then download an XML file. The XML can be then opened in a text editor, and copied and pasted into the box at the bottom of the runZero identity provider settings form. If the XML decodes successfully, the key fields on the form will then be set automatically.

Single Sign On mode

You can choose whether to allow regular runZero accounts, SSO-provisioned accounts, or both.

If both are allowed, you can set up backup administrator accounts in case your SSO provider is unavailable.

Domain name

The domain name is used by runZero to locate the correct SSO IdP settings when users choose to log in via SSO. This is so that it can redirect them to the correct SSO provider to log in. For example, if the domain is `example.com`, the runZero login page will expect users to begin logging in by entering an address of the form `username@example.com`.

The domain is also passed back by the SSO provider after login, and used to fetch the SSO IdP settings so that runZero can verify and process the data.

Issuer URL

The issuer URL, also often referred to as the entity ID, is a URL your SSO provider uses to uniquely identify itself in data passed back to runZero. You need to obtain this value from your SSO provider.

Login URL

The login URL is the URL users should be redirected to in order to begin the login process for your SSO provider.

Certificate

The certificate is the PEM encoded CA certificate runZero will use to verify the signature on the data from your SSO provider.

If you paste multiple CA certificates into the box, your SSO provider will need to include a KeyInfo element in data passed back to runZero, to specify which certificate to check the signature against.

SSO walkthrough

To configure and debug SSO settings, it helps to understand how the SAML login process works.

SAML SSO login is carried out using messages passed by the user's web browser. No direct connection occurs between runZero and the SSO identity provider (IdP) in either direction.

The normal login process starts with the user going to the runZero console login page, choosing **Login via SSO**, and entering their email address. The domain name of the email address is used to find the correct SSO identity provider settings in the runZero database.

runZero then redirects the user's web browser to the Login URL specified in those SSO IdP settings. Some additional parameters are added to the URL, according to the SAML specifications. These are verified by the SSO IdP.

The user logs in on the SSO IdP website. The IdP then redirects the user's web browser back to runZero, to the Assertion Consumer URL (ACS URL). The redirection includes a digitally signed blob of encoded data about the authenticated user. The ACS URL is shown on the service provider information tab of the SSO settings. Note that it includes the domain name, so that runZero can find the correct set of SSO settings to use to process the data.

runZero then decodes the encoded data, and checks its digital signature using the certificate that was provided as part of the identity provider settings. Assuming the certificate is valid, the decoded information about the user is then trusted by runZero. If necessary, a runZero user account is provisioned. The user is then logged in to begin a new runZero session.

Service provider information

After filling out the identity provider settings, you can switch to the service provider information tab to obtain information you need to provide to your SSO identity provider.

The Assertion Consumer URL, sometimes called the SSO URL, is the URL a person's web browser will be redirected to after they log in successfully on the SSO identity provider. The URL contains the domain that was set up in the identity provider settings. runZero uses the domain to find the right set of SSO settings, so it can verify the data from the SSO provider. This means that if you change the domain in the identity provider settings for any reason, the ACS URL will change, and your SSO provider will need to be given the new URL.

The user login URL is a URL that you can visit to begin the login process. It displays the computed identity provider URL to aid in debugging.

Common problems

A user logs in on the SSO provider, but runZero refuses the login

The most common cause of this problem is that the IdP has been configured with an Assertion Consumer URL that does not include the domain name that was set in the SSO identity provider settings in runZero, or includes the wrong domain name.

A correct ACS URL will look like

`https://console.runzero.com/auth/example.com/saml20/process` where `example.com` is the domain name configured in the runZero SSO identity provider settings.

Missing x509 Element

The full error message is “invalid SAML response: error validating response: Missing x509 Element”

This error can occur if you have entered multiple CA certificates into the certificate field of the SSO settings in runZero, but your SSO IdP didn't include a KeyInfo element to tell runZero which certificate to check the signature against. You will need to either configure your IdP to include information about which of the CA certificates to use as KeyInfo, or use a single CA certificate.

Managing SSO group mappings

Platform

Only runZero administrators can automatically map users to user groups using SSO attributes and custom rules.

SSO group mapping allows you to map your SAML attributes to user groups in runZero. In runZero, user groups explicitly set the organizational [role](#) and determines the tasks users can perform within each organization. When you set up SSO group mappings, you explicitly define the SSO attribute and value you want to use for mapping. If there is a match, runZero will apply the group settings for the user. As a result, you can ensure that SSO users are mapped to their respective groups in runZero.

For example, your IT team may need to be part of a group with administrator privileges. In this case, you can create a user group with an administrator role and then create an SSO group mapping that maps the SAML attribute that identifies your IT team to the user group. When

someone from your IT team logs in to runZero, they will automatically be added with the appropriate access and permissions, all without pre-provisioning their account. After evaluating all SSO group mapping rules, runZero grants the user the highest privilege assigned for each organization.

Creating SSO group mappings

Before you create your SSO group mapping, make sure that you have [set up SSO](#) for your organization and [created user groups](#). Both must be set up in order to successfully create SSO group mappings.

Only runZero super users can create SSO group mappings.

1. Go to [Your team > SSO settings > Group mappings > Add group mapping](#).
2. In the **SSO attribute** field, enter the attribute you want to check for matching values. These values are defined in your SSO configuration.
3. In the **SSO value** field, provide a comma separated list of values that the attribute could match. When there is a match, runZero will apply the group permissions.
4. Click the **Group** dropdown and choose the user group that will be assigned if there is a match. The dropdown will list all user groups that have been created.
5. Save the SSO group mapping. These settings will apply the next time the user logs in to runZero.

Changes will not apply to users currently logged in. They will need to log out and log back in for the changes to take effect. You can force log out users to apply the SSO group mappings immediately.

Forcing a user to log out

Changes to user permissions will not apply until the user logs out and logs back in to runZero. If you need to apply permissions immediately after setting up the SSO group mappings, you can force log out users. This will log users out of their current session and require them to log back in again. After they log in, their updated permissions will be applied. Only super users and admins who have access to all organizations can force logouts.

To force log out users, go to the [Teams page](#) and select the users you want to log out. Click the sign-out button to log these users out.

Viewing SSO group mappings

To view all SSO group mappings that have been created, you can go to the [Group mappings page](#). From this page, you can create, edit, or delete group mappings as needed.

Viewing SSO group mapping assignments

To see the SSO groups that a user has been assigned to, go to the [Users page](#). From the Groups column, you can see the number of user groups and SSO groups the user is a part of.

The number of SSO groups will be in parentheses.

Clicking the gear icon under actions will open the user settings for the user. The access summary tab will then display all of the organizations and roles they have assigned.

Deleting a group mapping

1. Go to the [Group mappings page](#).
2. Select the group you want to delete and click the **Delete** button. All users provisioned through the group mapping will revert back to their account-level permissions.

Searching for SSO group mappings

When you are on the [Group mappings page](#), you can use the following keywords to search in the table:

Keyword	Description	Example
id	User's ID.	id:123456789
sso_attribute	User's SSO attribute.	sso_attribute:department
sso_value	User SSO attribute value.	sso_value:securityteam
created_at	Time or date user group was created.	created_at:>2weeks
updated_at	Time or date user group was last updated.	updated_at:>1year
created_by_email	Email of user who created the group.	created_by_email:user@example.com
group_id	User group ID.	group_id:123456789
group_name	User group's name.	group_name:group1

Group IDs can be found in the URL for the group config page (<https://console.runzero.com/groups//edit>).

The group_id keyword is only available for the Users table; for the groups table, use id.

The group_name keyword is only available for the Users table.

Setting up Azure AD SSO

Superusers can configure single sign-on to the runZero Console using an external identity provider (IdP), which enables authentication and user access control to the runZero Console from your single sign-on (SSO) solution. By default, runZero has SSO functionality available,

but it's not a requirement to log in to the console. You can make it a requirement or disable it completely.

Here are the high-level steps to set up SSO using Azure AD to authenticate and manage user access to runZero:

1. [Add and configure runZero as an Azure AD app.](#)
2. [Download the SSO configuration metadata in XML format.](#)
3. [Set up SSO in runZero.](#)
4. [Add users to your runZero app in Azure AD.](#)

Requirements

Before you can set up SSO for Azure AD:

- Verify that you have administrator privileges for Azure AD.
- Verify that you are a [superuser in runZero](#). Look for the yellow star in your account status.

Step 1: Add and configure runZero as an Azure app

The first thing you need to do is [add runZero as a non-gallery application](#) to your Azure AD setup and to [configure the settings for runZero](#) as an Azure AD application.

1. In Azure, go to **Enterprise Applications > New Application > Create your own application**.
2. Under the **What are you looking to do with your application?** section, choose the **Non-gallery application** option.
3. Name your application something like runZero, and then add it.
4. Go to **Azure Active Directory > Enterprise applications** and open the newly created runZero application.
5. Select the **Single sign-on** tab, and then choose **SAML** as the sign-on method.
6. For the fields on the **Configure App Settings** page, go to <https://console.runzero.com/team/sso/sp> and copy the necessary service provider details:
 - Entity ID
 - Single sign-on URL
 - SSO callback (ACS) URL
7. Enter the values into the relevant fields in the Azure AD portal.
8. **Do not** set a value for "Sign on URL (Optional)" or "Relay State (Optional)".

Step 2: Download the SSO configuration metadata

While editing your application settings, you can get the download link to obtain the SSO configuration metadata in XML. You'll need this information to set up SSO in runZero.

1. On the **Configure App Settings** page, find the **SAML Signing Certificate** section.
2. Locate the XML download link under the [Federation Metadata URL](#).

3. Download the file. You'll need the contents of this file for the next step.

Step 3: Set up Azure AD SSO in runZero

Now that you have the SSO configuration metadata in XML, you can configure Azure AD SSO settings in runZero.

1. Go to <https://console.runzero.com/team/sso/idp> to access the SSO IdP provider settings page in runZero.
2. Choose one of the following modes to enable SSO:
 - **Allowed** - Enables SSO, but users still have the option to login without SSO.
 - **Required** - Requires users to log in with SSO. Only superusers can log in without SSO.
3. Enter the domain name that is associated with SSO authentication. This is likely your company domain (companyabc.com).
4. Choose a default role for SSO users. This is the [role](#) all new users will be assigned when their account is created.
5. Copy the XML you downloaded from Azure and paste it into the Metadata XML field on the runZero SSO IdP page.
6. Apply your SSO settings. The remaining IdP fields will auto-configure for you.
 - The issuer URL will look something like `https://sts.windows.net/00000000-0000-0000-0000-000000000000/` where the UUID at the end is your unique Microsoft Active Directory (tenant) ID, listed under **App registrations > Overview > Endpoints**.
 - The login URL will be something like `https://login.microsoftonline.com/00000000-0000-0000-0000-000000000000/saml2` with the zero UUID replaced with your unique tenant ID.
 - The certificate will be Microsoft's PEM encoded certificate, which will be extracted automatically from the XML.
 - On the Microsoft side, the redirection URL for runZero should be `https://console.runzero.com/auth/<domain>/saml20/process`, where `<domain>` is replaced with the domain specified in the runZero SSO settings.

Step 4: Add users to the runZero app in Azure

Now that you've completed the set up, you can go to the runZero app in Azure portal to [add users and assign their access](#). Any users you add to the runZero app will be viewable from the [Team members page](#) in runZero, once they have logged into runZero.

Step 5: Update SSO group mappings to match any configured Azure groups (if applicable)

If you have created user groups within Azure, you will need to update your [SSO group mappings](#) in runZero to associate the groups created in Azure with user groups in runZero. This will ensure that the appropriate access and permissions are added to your users when they log in to runZero.

Setting up Okta SSO

Superusers can configure single sign-on to the runZero Console using an external SAML identity provider (IdP), such as Okta, which enables authentication and user access control to the runZero Console without typing in credentials.

Here are the high-level steps to set up single sign-on (SSO) using Okta to authenticate and manage user access to runZero:

- [Add runZero as an application in Okta.](#)
- [Set up SSO in runZero.](#)
- [Add users to the runZero app in Okta.](#)

Requirements

Before you can set up Okta SAML:

- Verify that you have administrator privileges for Okta.
- Verify that you are a [superuser in runZero](#). Look for the yellow star in your account status.

Step 1: Add and configure runZero as an Okta app

1. Go to **Okta > Applications > Create App Integration**. When the **Create a new app integration** window appears, select SAML 2.0 for your sign-in method.

- OIDC - OpenID Connect**
Token-based OAuth 2.0 authentication for Single Sign-On (SSO) through API endpoints. Recommended if you intend to build a custom app integration with the Okta Sign-In Widget.
- SAML 2.0**
XML-based open standard for SSO. Use if the Identity Provider for your application only supports SAML.
- SWA - Secure Web Authentication**
Okta-specific SSO method. Use if your application doesn't support OIDC or SAML.
- API Services**
Interact with Okta APIs using the scoped OAuth 2.0 access tokens for machine-to-machine authentication.

2. For the general settings, you'll need to provide a name for the app. Call the app runZero. You can also add a logo to make it easier for users to identify the runZero app.
3. For the SAML settings, you'll need to go to the [service provider information page](#)

in runZero to find the relevant URLs.

- **Single sign on URL** - In runZero, this is the assertion consumer service (ACS) URL.
 - **Audience URI or SIP Entity ID** - In runZero, the entity ID, or SAML audience, will be `https://console.runzero.com`.
 - **NOTE:** make sure to include the leading 'https://' when entering this field in Okta.
4. For the remaining settings, like the attribute statements, visit the Okta documentation to [learn how to configure them](#).
 5. When you finish configuring the SAML settings, Okta will prompt you for some feedback on how you will be using the app. You can opt to provide feedback or skip to complete the set up.
 6. After Okta creates the app, you will need to view the SAML 2.0 instructions to complete the set up. Go to the **Sign On** tab for the runZero app and view the SAML 2.0 instructions. You'll need these details for the next step.

☰

SAML 2.0 is not configured until you complete the setup instructions.

[View Setup Instructions](#)

Identity Provider metadata is available if this application supports dynamic configuration.

Step 2: Set up SSO in runZero

1. Go to the [SSO setup page](#) in runZero. Choose one of the following modes to enable SSO:
 - **Allowed** - Enables SSO, but users still have the option to login without SSO.
 - **Required** - Requires users to log in with SSO. Only superusers can log in without SSO.
2. Enter the domain name that is associated with SSO authentication. This is likely your company domain (companyabc.com).
3. Choose a default role for SSO users. This is the [role](#) all new users will be assigned when their account is created.
4. Copy the fields from Okta into runZero.
 - **Issuer URL** - In this field, enter the **Identity Provider Issuer URL** from Okta. This will look something like `http://www.okta.com/<ID>`
 - **Login URL** - In this field, enter the **Identity Provider Single Sign-On URL** from Okta. This will look something like `http://<okta-instance>/app/<app-name>/<ID>/sso/saml`.
 - **Certificate** - Copy and paste the entire contents of the X.509 certificate from Okta.
5. Apply your SSO settings.

Step 3: Add users to the runZero App in Okta

Now that you've completed the set up, you can go to the runZero app in Okta to [add and manage user access](#). After you've completed this step, your users will be able to go to your SSO sign-in URL to access runZero.

Step 4: Update SSO group mappings to match any configured Okta groups (if applicable)

If you have created user groups within Okta, you will need to update your [SSO group mappings](#) in runZero to associate the groups created in Okta with user groups in runZero. This will ensure that the appropriate access and permissions are added to your users when they log in to runZero.

Managing licenses

As a runZero superuser or billing user, you can access and manage your organization's licensing, plan, and billing information.

Not a superuser or billing user? Please contact your organization's [superuser or billing user](#) to get help with licensing and billing information. These users are tagged with a yellow star.

How do I view my license?

If you're a superuser or billing user, go to [Account > License](#) to view your runZero licensing information.

Your licensing information will show:

- The number of organizations, sites, Explorers, and user accounts you have.
- The number of recent assets you have across all organizations.
- The number of project assets you have.
- The total number of recent assets and project assets your license supports.

What count as recent assets?

runZero calculates the number of recent assets based on how many have been seen during the last 30 days.

For assets that have been scanned by runZero, if they have been seen by a scan in the last 30 days, they are counted as recent assets.

If third party data indicates a device has been seen in the last 30 days, that will also cause the asset to count as recently seen.

License limits are uniform across all asset types. It doesn't make any difference whether the asset is online or offline, scanned or unscanned, or where the asset data was sourced from. For assets with multiple data sources, with different dates that the device was last seen at, the most recent date applies.

The number of recent assets is calculated in the background. If you delete assets it may take a short while for the number to recalculate.

When does my subscription expire?

To see when your subscription or license expires, go to [Account > License](#).

Find the line: This is a runZero [edition] subscription that expires at [date and time].

If your subscription has expired, you will see: This is a runZero [edition] subscription that expired on [date and time]. You will no longer be able to run discovery scans. To continue using runZero, you will need to [renew your subscription](#).

Does the Community license expire? Nope, it's part of our free tier and has no expiration date. You can continue to use it as long as you don't have more than 100 recent assets.

How do I renew my subscription?

For help renewing your subscription, please contact us by [email](#) or by using this [online form](#).

How do I convert to the Community Edition?

No longer need the full platform? You can [convert it to the Community Edition](#). This plan is on our free tier and supports up to 100 recent assets.

How do I find my invoices?

Super users and billing users can access invoices from the runZero Console. Go to the [License page](#) to see all invoices.

How do I change or cancel my subscription?

Please [contact us](#) if you need to modify or cancel your subscription.

Gathering data

runZero can gather asset data through unauthenticated active scanning, passive traffic sampling, and inbound integrations.

Active scanning

The runZero Explorer and scanner perform [unauthenticated active scanning](#) of your specified networks based on the configurations you set. They leverage various network protocols to discover and fingerprint assets connected to the network.

Active scans can be configured to run once or on a schedule. Scan templates can also be used to ensure consistency across multiple scan tasks.

Traffic sampling

Community Platform

Explorers can be configured to perform passive traffic sampling in your environment in order to identify assets communicating across your networks. [Traffic sampling](#) is configured on a per-Explorer basis on the Explorer details page. Explorers will automatically pause traffic sampling activities in order to handle scan tasks, continuing to parse sniffed traffic once the scan completes.

The traffic sampling feature can also ingest a PCAP file to glean useful asset data.

Inbound integrations

[Inbound integrations](#) for many different platforms can be configured to enrich your runZero asset context and provide important insights into coverage and capabilities. Integrations can be configured either as connectors or scan tasks.

Discovery scanning

A discovery scan finds, identifies, and builds an inventory of all the connected devices and assets on your internal network. Running a discovery scan routinely will help you keep track of and know exactly what is on your network.

Discovery scans are configured by site, Explorer, and scope. In order to run a scan against a specific site, an Explorer must be activated and either assigned to that site or configured for all sites.

When creating a new scan, you have multiple parameters you can set, ranging from scheduling a date to more advanced options. To launch a discovery scan, browse to the [Inventory](#) page, click the **Scan** menu in the upper right, and select **Standard Scan**.

Site

runZero organizes information into organizations and sites. Organizations are distinct entities that are useful for keeping data separate and contain a collection of sites. Sites are used to model segmented networks, particularly independent networks which use the same private IP address ranges.

For example, you might have multiple physical locations with their own local networks, all using the 10.0.0.0/8 private IP range. By defining them as sites, you can set up an Explorer for each, and the networks and assets will be treated as completely independent even if similar systems are seen at the same IP addresses in each.

Since scan analysis occurs at the site level, the boundaries you define for a site set the default scope for scans for that site.

Explorer

Select the Explorer to run the scan from, chosen from the set of registered Explorers for the site. The Explorer you choose must be able to directly communicate with the networks and addresses you define for the discovery scope.

The chosen Explorer should ideally be able to reach all addresses in the scope directly, without a firewall in the way. Stateful firewalls and VPN gateways may interfere with the discovery process.

Hosted zone

[Community](#) [Platform](#)

runZero Platform users can perform scans of public IP space using runZero-hosted scanners. When creating a scan, set the Explorer to **None** and choose a hosted zone from which to scan. When using this option, the discovery scope must use public IP addresses or ranges, or resolve to public IP space.

Discovery scope

The discovery scope defines the IP addresses that will be scanned. The scope uses the site settings when specified as they keyword “defaults”, but may be changed on a per scan basis as well. The scope should include at least one IP address or hostname. IPv4 address ranges can be specified in most standard formats:

- 10.0.0.1
- 10.0.0.0/24
- 10.0.0.0/255.255.255.0
- 10.0.0.1-10.0.0.255

IPv6 addresses can be specified individually, but IPv6 ranges are not supported.

Hostnames specified in the scope will be resolved at runtime by the assigned Explorer. If the hostname returns multiple IP addresses, all addresses in the response will be scanned. Hostnames can also have masks applied, indicating that the mask should expand to each resolved address of the hostname. For example, if `example.com` resolves to both `1.2.3.4` and `5.6.7.8`, the input of `example.com/24` would become `1.2.3.0/24` and `5.6.7.0/24`. IPv6 addresses returned from hostname resolution will be scanned if the Explorer has a valid IPv6 address and route to the target.

Note that the Explorer scans addresses in random order. Subnets are scanned in a random order, and within each subnet the IP addresses are also scanned in a random order. This is done to avoid concentrating traffic in particular parts of the network.

Discovery keywords

The following keywords are supported for both scan scopes and exclusions.

- **asn4**: The `asn4:<AS number>` keyword can be used to specify IPv4 ranges associated with a given AS number.
- **country4**: The `country4:<ISO code>` keyword can be used to specify IPv4 ranges associated with a given two-character country code.
- **public** and **private**: The `public:<mode>` and `private:<mode>` keywords can be used to specify IPv4 and IPv6 addresses associated with assets in the current organization. The mode parameter can be set to `all`, `primary`, or `secondary` to indicate which IP addresses are used. The `public` keyword selects all non-reserved IP addresses associated with organization assets. The `private` keyword selects all RFC-1918 and private use IP addresses associated with organization assets.
- **domain**: The `domain:<domain>` keyword is available to cloud-hosted users and uses the syntax `domain:<domain name>` to automatically select publicly-known hostnames for a given domain name.

Scan name

You can assign a name to your Scan task to make it easier to keep track of.

Scan speed

Specify the maximum packet rate for the overall discovery process, in network packets per second. 500 is conservative, 3000 works for most LANs including WiFi, 10000 or more may be helpful for large sites with fast connectivity.

The scan speed directly affects how long the scan will take to complete. An approximate formula is:

$$\text{time in seconds} = \text{hosts} \times \text{ports} \times \text{attempts} \div \text{scan speed}$$

The number of hosts scanned is primarily determined by the discovery scope. The number of ports is around 500 by default, and three attempts are made to connect.

The number of hosts and ports scanned can be affected by the advanced scan options, and speed can also be impacted by maximum host rate and group size; see the descriptions of the advanced scan options below.

Note also that this formula doesn't take into account time taken to take screenshots, follow web server redirects, or process the scan data.

Schedule

You can set a date and frequency for your scan task. Dates and times take into account your browser's advertised timezone.

Scans scheduled to start in the past will be launched immediately and then repeated at the specified time based at the frequency selected.

Scheduling grace period

Specify the number of hours to wait for an available Explorer before giving up on this scan. A zero or negative value will result in the scan retrying indefinitely until an Explorer becomes available.

Advanced scan options

The Advanced tab can be used to display and modify additional scan settings, such as network exclusions, scan speed, the ports covered by the TCP scan, and which probes are enabled. The default settings should work for most organizations but may need to be tweaked for slow networks or unreliable links.

Maximum host rate

As well as setting an overall scan rate in packets per second, you can also control the maximum rate at which packets are sent to any single host IP address. This is useful when you have devices which are easily overloaded by network traffic. The default should be safe for most systems.

Max group size

When runZero scans your network, it spreads the scan load across many IP addresses at once. The max group size determines how many IP addresses can be actively scanned at once – allowing for the fact that hosts may take some time to respond to probes. The max group size needs to be at least as large as the overall scan speed, or else it would limit the speed of the scan to below the set value. If you provide a value that's lower than the overall scan speed, it will be increased automatically at scan time.

The max group size is mostly useful when dealing with stateful network devices that can only track a limited number of connections at once, as a way to restrict how many active TCP sessions will result from a runZero scan.

Max TTL

The IP standards define a maximum hop count for packets. In IPv4, this is called the Time To Live or TTL, while on IPv6 this is called the Hop Limit. Every device processing a packet must decrease the TTL or Hop Limit one. If this value reaches zero, the route receiving the packet must discard the packet. This setting can be used to set the maximum hop limit for scan traffic.

ToS

The IP standards define a Type of Service or ToS for packets. In IPv4, this is called the Type of Service or ToS, while on IPv6 this is called the Traffic Class or TC. The ToS or Traffic Class is used by switches and routers to prioritize network traffic. The lower bits of the IPv4 ToS are also used for congestion controller. This setting can be used to set the ToS or Traffic Class for scan traffic. Please note that the ToS/Traffic Class settings do not apply to all traffic sent by runZero, but instead are limited to the basic discovery probes. Some protocols, such as SNMP, and integrations, such as VMware, do not set the ToS/Traffic Class fields on their corresponding packets. If all scan traffic must be consistently tagged with the correct ToS or Traffic Class, this can be accomplished through settings on the managed switch port instead.

TCP ports

The **Included TCP ports** and **Excluded TCP ports** fields can be used to override the default scan ports. The string "defaults" will lookup the current default port list at scan time. The current port list is:

runZero User Guide

1	7	9	13	17	19	21	22	23	25	37
42	43	49	53	69	70	79	80	81	82	83
84	85	88	102	105	109	110	111	113	119	123
135	137	139	143	161	179	222	264	280	384	389
402	407	442	443	444	445	465	500	502	512	513
515	523	524	540	548	554	587	617	623	631	636
664	689	705	717	743	771	783	830	873	888	902
903	910	912	921	990	993	995	998	1000	1024	1030
1035	1080	1083	1089	1090	1091	1098	1099	1100	1101	1102
1103	1128	1129	1158	1199	1211	1220	1234	1241	1260	1270
1300	1311	1352	1433	1440	1443	1468	1494	1514	1521	1530
1533	1581	1582	1583	1604	1610	1611	1723	1755	1801	1811
1830	1883	1900	2000	2002	2021	2023	2049	2068	2074	2082
2083	2100	2103	2105	2121	2181	2199	2207	2222	2224	2323
2362	2375	2376	2379	2380	2381	2443	2525	2533	2598	2601
2604	2638	2809	2947	2967	3000	3001	3003	3033	3037	3050
3057	3071	3083	3128	3200	3217	3220	3260	3268	3269	3273
3299	3300	3306	3311	3312	3351	3389	3460	3500	3502	3628
3632	3690	3780	3790	3817	3871	3872	3900	4000	4092	4322
4343	4353	4365	4366	4368	4369	4433	4443	4444	4445	4567
4659	4679	4730	4786	4840	4848	4949	4950	4987	5000	5001
5007	5022	5037	5038	5040	5051	5060	5061	5093	5168	5222
5247	5250	5275	5347	5351	5353	5355	5392	5400	5405	5432
5433	5498	5520	5521	5554	5555	5560	5580	5601	5631	5632
5666	5671	5672	5683	5800	5814	5900	5901	5902	5903	5904
5905	5906	5907	5908	5909	5910	5911	5920	5938	5984	5985
5986	5988	5989	6000	6001	6002	6050	6060	6070	6080	6082
6101	6106	6112	6161	6262	6379	6405	6443	6502	6503	6504
6514	6542	6556	6660	6661	6667	6905	6988	7000	7001	7002
7021	7070	7071	7077	7080	7100	7144	7181	7210	7373	7443
7474	7510	7547	7579	7580	7676	7700	7770	7777	7778	7787
7800	7801	7879	7902	8000	8001	8003	8006	8008	8009	8010
8012	8014	8020	8023	8028	8030	8080	8081	8082	8083	8086
8087	8088	8089	8090	8095	8098	8099	8100	8123	8127	8161
8172	8180	8181	8182	8205	8222	8300	8303	8333	8400	8443
8444	8445	8471	8488	8500	8503	8530	8531	8545	8649	8686
8787	8800	8812	8834	8850	8871	8880	8883	8888	8889	8890
8899	8901	8902	8903	8983	9000	9001	9002	9042	9060	9080
9081	9084	9090	9091	9092	9099	9100	9111	9152	9160	9200
9300	9380	9390	9391	9401	9418	9440	9443	9471	9495	9524
9527	9530	9593	9594	9595	9600	9809	9855	9999	10000	10001
10008	10050	10051	10080	10098	10162	10202	10203	10250	10255	10257
10259	10443	10616	10628	11000	11099	11211	11234	11333	12174	12203
12221	12345	12379	12397	12401	13364	13500	13778	13838	14330	15200
15671	15672	16102	16443	16992	16993	17185	17200	17472	17775	17776
17777	17778	17781	17782	17783	17784	17790	17791	17798	18264	18881
19300	19810	19888	20000	20010	20031	20034	20101	20111	20171	20222
20293	22222	23472	23791	23943	25000	25025	25565	25672	26000	26122
27000	27017	27018	27019	27080	27888	28017	28222	28784	30000	31001
31099	32764	32844	32913	33060	34205	34443	34962	34963	34964	37718

```

37777 37890 37891 37892 38008 38010 38080 38102 38292 40007 40317 41025
41080 41523 41524 44334 44343 44818 45230 46823 46824 47001 47002 47290
48899 49152 50000 50013 50021 50051 50070 50090 50121 51443 52302 52311
54321 54921 54922 54923 55553 55580 57772 61614 61616 62078 62514 65002
65535

```

Prescan modes for large IP spaces

Sometimes, the scope of your IP space is unknown, subnet usage is unknown, and the total number of assets is unknown. These unknowns can make it challenging to optimize your discovery scans for efficiency and speed. And when your IP space is large, like a /16 space with a few thousand IPs in use, a full discovery scan can take more time to complete, since it looks at more than 500 TCP ports and 15 UDP ports on every address. In these types of cases, you may want to tune your scan settings to prefilter ranges and IP addresses before a full scan.

runZero has two prescan modes that you can use to run a faster scan: subnet sampling and host ping.

Subnet sampling

Professional Community Platform

To speed up scans of large subnets you can use the **“Only scan subnets with active hosts”** advanced scan option. If this option is on, a prescan runs against the target space to identify the subnets with an active host. This mode leverages heuristics runZero has collected to identify addresses that are more likely to be responsive across subnets. This process allows runZero to quickly scan larger spaces by identifying the subnets that are in use, before starting full probes. All subnets that are identified as having active hosts are then fully scanned – unless you enable host pings.

There are two tweakable parameters for subnet sampling. The sample rate determines what percentage of addresses in each subnet are prescanned to determine if the subnet should be scanned. The subnet size determines how many IP addresses are in each subnet. By default, the subnet size is 256 addresses, corresponding to a /24 subnet, and 3% of the addresses in each subnet are prescanned.

Host ping

After you have some insights on the subnets that are in use, you may want to limit the full scan to only addresses that respond to the most common ping methods, such as ICMP and some TCP and UDP ports. If you choose the **“Limit scans to pingable hosts”** advanced scan option, only hosts that respond to a ping request will be fully scanned.

The runZero Explorer uses multiple protocols for ping scans:

- Conventional ICMP ping, performed by sending an ICMP echo request and looking for an ICMP echo reply.
- TCP ping, performed by sending a TCP SYN packet to a series of common ports and seeing whether the host responds with RST or TCP SYN/ACK.

- UDP ping, performed by sending a packet to port 65535 and checking for an ICMP response of port unreachable.

The set of ports used for TCP and UDP ping can be adjusted in the LAYER2 section of the Probes and SNMP tab when setting up a scan task.

Note that it is relatively common for enterprise firewalls to be set up to block ping, or for hosts to be set up not to respond to ping requests. Limiting scans to pingable hosts can therefore result in assets being missed entirely, even if their IP addresses are probed. If your goal is to speed up scan times, subnet sampling is usually the better option.

It's possible to use both subnet sampling and limiting scans to pingable hosts at the same time, but this is not recommended except as a last resort for reducing scan times.

Running initial scans

Background

Once you have an Explorer installed, you can start using it for network discovery. While our goal is to configure scheduled scans that we set and forget, we need to go about our first scans in a more structured manner.

The goals of our first scans are to:

- Verify the Explorer is setup properly and has everything installed
- Validate Explorer connectivity to varying parts of the network
- Determine how long scans will take at varying sizes to help with future scheduling

Your first few scans

To get started, you will want to scan a few smaller ranges to make sure everything is working as expected. Start with a few /24 network blocks from each of the RFC 1918 ranges to make sure everything looks good.

For setting up the first scan:

1. Navigate to Sites > New Site > Create a new temporary site within the Organization
2. Navigate to Tasks > Scan > Standard Scan to create a scan task
3. Chose the new site you created in step 1
4. Include a range of the RFC1918 IP addresses in the Discovery Scope, plus a small network or two that you know is in use. A suggested value for the RFC1918 range includes:
10.0.0.0/24, 10.0.255.0/24, 10.64.0.0/24, 10.64.255.0/24, 10.128.0.0/24, 10.128.255.0/24, 10.192.0.0/24, 10.192.255.0/24
<your networks here>
5. On the Advanced tab, enable the Subnet sampling option
6. Click on Initialize Scan

After these scans are complete, you will want to check for these things:

- Check the `ipv4.traceroute` value for assets in each RFC1918 range to verify you aren't sending traffic to an edge router or firewall.
 - Unused private IPs should have routes stubbed out to prevent traffic from being sent to the default gateway which can create a loop. You can also verify this with traceroutes from the Explorer.
- If your scan results have a large series of somewhat sequential IPs that have only ICMP or a very small number of similar ports open on them, that's probably a proxy or firewall. Check out those IPs to see if any are real. To find assets with only ICMP enabled use the inventory query `alive:t AND service_count:=1 AND service_count_icmp:=1`
 - You can add an allow rule for the Explorer IP to properly scan devices on the other side
 - Another option is to add a second Explorer on the other side of the proxy or firewall
- If you receive reports or alerts about service outages, check for session aware devices such as routers, firewalls, and proxies that are having issues handling the session load. If you run into this, there are multiple ways to approach solving the issue.
 - The simplest solution is to set up another Explorer on the other side of the device and run scans separately
 - Another option is to segment your scans on the existing Explorer, and run smaller, separate scan tasks for the network ranges on the other side of the device with lower packet per second and max group sizes to minimize the number of IPs that will be scanned at once
- Check how long each scan took to get an idea for how long larger scans would take
- Verify you see screenshots on ports that accept HTTP/HTTPS requests
 - If you don't see any, you likely need to install Chrome on the machine
- Check for MAC addresses
 - If you aren't seeing them you should configure SNMP
 - If SNMP is configured, you should verify community strings and check for unmanaged switches

Note: Once you have verified that your first scan ran successfully, you can delete the temporary site and set up a real scan.

Full RFC 1918 scans

Once you have completed initial test scans, it's time to expand scanning to cover all subnets with live assets. One method of discovering all subnets with live assets is to run full RFC 1918 scans.

runZero offers a **Full RFC 1918 discovery** scan option that will discover assets across the following private address ranges as a single task.

- 10.0.0.0/8 or 10.0.0.0-10.255.255.255
- 172.16.0.0/12 or 172.16.0.0-172.31.255.255
- 192.168.0.0/16 or 192.168.0.0-192.168.255.255

The **Full RFC 1918 discovery** scan option is only recommended for small networks with limited complexity and should only be leveraged in a single site configuration.

Discovering the entire RFC 1918 private address space in a single scan can take days, if not weeks, to complete in a large complex network. runZero recommends reviewing the [Achieving RFC 1918 coverage](#) playbook for more information on scanning the entire RFC 1918 private address space.

Once you've completed scans of your private address space, review [Identifying gaps in scanning](#) to learn more about some of our built-in reports that we offer to help you get a better understanding of your network.

Scanning SNMP

SNMP is an open standard network protocol for collecting information about devices on a network. There are three main versions of the protocol.

SNMP versions 1 and 2

SNMP version 1 was designed in the 1980s as an interim protocol, intended to be replaced by ISO CMIP. It was built to be used across any network common at the time, not just TCP/IP networks, so security was left up to the host network. The protocol defined a community string for arbitrary organization of groups of assets, but didn't specify how access should be granted.

SNMP version 2 attempted to introduce a security model based on "parties," but it wasn't widely adopted so a revised standard was issued as SNMPv2c. SNMPv2c removed party-based security and went back to just using community strings. The previous SNMPv2 specification is considered obsolete.

Unfortunately, both SNMPv1 and SNMPv2 send community strings in plain text across the network. There's no encryption of SNMPv1 or v2 data packets, and because the protocol is based on UDP there's no way to simply add TLS to make a secure connection.

SNMP version 3

SNMPv3 fixes the security problems of SNMPv2 by supporting both password-based authentication and encryption (referred to as privacy protection) as part of the standard. Unfortunately there are multiple algorithms supported for both the encryption and the password authentication, and you have to know which ones your network devices use. Unlike SSH for example, the protocol doesn't include any kind of negotiation of encryption methods.

SNMPv3 devices can operate in three different modes:

- noAuthNoPriv, meaning no authentication is required and there's no privacy protection (no encryption),

- `authNoPriv`, meaning authentication *is* required but there's still no privacy protection, and
- `authPriv`, meaning both authentication and privacy protection are required.

runZero's SNMP support

runZero supports SNMPv1, SNMPv2 (the SNMPv2c variant), and SNMPv3. Scans can be performed using only v1/v2, only v3, or both.

SNMP scanning is on by default. You can turn it off or customize it using the SNMP tab when setting up a scan or a scan template.

SNMPv1/v2 scanning

There are two ways to set up community strings for SNMPv1/v2 scanning. The first is to enter them as a comma-separated list on the SNMP tab. By default, runZero supplies the community strings `public` and `private`, as these are common defaults on network-enabled hardware such as printers and NAS servers. If you remove those defaults, runZero will not probe with them.

The second way to set up community strings is to enter them using the [credentials feature](#) as the credential type *SNMP v2 Communities*. While communities technically shouldn't be used as credentials because they're sent in plain text across the network, in practice many networks use them that way. By entering your community strings as credentials, you can use the CIDR allow list feature to control which parts of the network the community strings will be sent to, reducing the risk of their being captured by rogue devices.

If you remove all community strings from the SNMP tab and do not set up any SNMP v2 Communities as credentials, no SNMPv1/v2 scanning will be performed.

SNMPv3 scanning

Most devices which provide potentially sensitive information such as serial numbers and software versions will only do so in response to an *authenticated* SNMPv3 query. To perform authenticated SNMPv3 scanning, you will need to set up an *SNMP v3 Credential* record in runZero's credentials feature.

The authentication protocol determines the hashing algorithm used to process the authentication passphrase – that is, how the runZero Explorer logs in to the remote device.

The most common default algorithm that devices use for authentication is sha (HMAC-SHA-96), which is required to be supported by the SNMPv3 standards. Newer devices may support more secure variants such as SHA-256, and runZero supports up to SHA-512. HMAC-MD5-96 is also supported, as per the standards, but is best not used on your network because the MD5 algorithm is known to be insecure.

The privacy protocol determines how the data sent to and from the remote device is encrypted to prevent eavesdropping. The privacy passphrase is used as seed data to initialize the encryption.

The most common privacy algorithm, required by the original SNMPv3 standards, is CBC-DES-128. It's selected as `des` in the runZero user interface. Later RFCs added AES-128, represented as `aes`. Again, runZero supports more secure variants that may be used by newer devices, such as AES256.

As [mentioned above](#), devices may require only authentication, both authentication and privacy, or neither. For some devices there is only space for a single passphrase in the device configuration. This usually means that the same passphrase is used for both the privacy passphrase and the authentication passphrase, if the device is running in `authPriv` mode.

As with SNMPv2 community strings, you can set a CIDR allow list to determine which IP addresses SNMPv3 credentials will be sent to. This is particularly recommended if your network is not set up to require privacy passphrases.

In addition to credentialed scanning, runZero will send uncredentialed SNMPv3 probes. Some devices respond to SNMPv3 with useful information without needing any credentials.

To perform SNMPv3 scanning across Cisco switches, modify the ACLs to include the following rule:

- Newer IOS versions: `snmp-server group YourGroupName v3 auth context vlan-match prefix`
- Older IOS versions: `snmp-server group YourGroupName v3 auth context vlan-1` (repeated for every VLAN)

Choosing which versions of SNMP are scanned

To disable SNMPv1/v2 and only scan with SNMPv3, remove the community strings from the SNMP tab when setting up a scan or template and ensure that any SNMPv2 community strings stored as credentials are disabled on the Credentials tab.

To disable authenticated SNMPv3 scanning and only scan with SNMPv1/v2, ensure that any SNMPv3 credentials are disabled on the Credentials tab and make sure you have community strings specified, either on the SNMP tab or a credential.

To disable all SNMP scanning, including uncredentialed SNMPv3, switch off the toggle switch on the SNMP tab labeled *Use the SNMP protocol for discovery*.

How scans are performed

runZero will probe devices using all configured SNMP versions and all community strings. This is because devices often respond with different levels of data for SNMPv3 versus SNMPv2, and may respond with different information depending on community string.

SNMP is based on UDP, and UDP doesn't guarantee the order of data packets, so it's impossible to guarantee in what order community strings will be received by remote devices.

SNMPv3 authentication errors are not treated as scan failures. This is because it's common to have multiple sets of SNMPv3 credentials used on different parts of a network, as well as misconfigured devices with incorrect passphrases.

Checking SNMPv3 credentials and SNMPv1/v2 authentication

If you are not getting SNMP data in your assets and think that you should be, a common cause is network firewalls, which often block SNMP traffic entirely. If possible, deploy an Explorer on the other side of the firewall, on the network segment you want to scan with SNMP.

The second most common reason for not getting the expected SNMP data in assets is incorrect SNMPv3 credentials.

- If SNMPv3 authentication was attempted, you will see an *SNMP engine ID* recorded in the asset as `snmp.engineID.raw`. The engine ID is usually a long string of hex digits and is used to calculate access keys.
- If SNMPv3 authentication failed, you will see an error reported in the attribute `snmp.failedAuth`. If the error is a request timeout, that likely indicates that authentication succeeded but the encrypted data channel could not be set up, which means you may have the wrong privacy password or privacy algorithm.
- If SNMP authentication was attempted successfully (v1/v2 or v3), you will see an `snmp.credentials` attribute with the UUID of any credentials record that succeeded.
- If SNMPv1/v2 authentication via community string succeeded, you will see `snmp.secretCommunities` or `snmp.defaultCommunities` attributes, depending on the type of community string.
- If SNMPv1/v2 authentication succeeded using community strings entered directly into the task request rather than credentials records, the `snmp.secretCommunities` and/or `snmp.defaultCommunities` attributes will still be present, but `snmp.credentials` will be a zero UUID.

Debugging credential issues

It can be frustrating to debug SNMPv3 credential problems by trial and error using multiple scan attempts. Instead, you can use the `snmpwalk` utility. It's part of the [net-snmp](#) open source project, packaged for most Linux distributions and available for macOS via Homebrew. There are unofficial builds for Windows, or you can install it in the Windows Subsystem for Linux (WSL).

The `snmpwalk` utility will connect to an IP address and dump all of the information it can retrieve. Here's an example of how to use it to connect in `authPriv` mode with both authentication and privacy passphrases:

```
snmpwalk -v3 -l authPriv -a SHA -A "authentication passphrase"
-x AES -X "privacy passphrase" -u username 10.0.1.25
```

(Text above has been line-wrapped, but it's one long command.) The parameters are:

- `-v3`: switch to SNMPv3 mode.
- `-l authPriv`: run in `authPriv` mode, with both authentication and privacy required.
- `-a SHA`: authentication algorithm is SHA (SHA-128).
- `-A`: specify the authentication passphrase as the next argument.

- -x AES: privacy (encryption) algorithm is AES (AES-128).
- -X: specify the privacy passphrase as the next argument.
- -u username: specify the login username.
- 10.0.1.25: the IP address to probe.

If you have the right passphrases and algorithms, you'll get data back, probably a lot of it. If not, you'll get an error message which may help identify the cause.

To check SNMPv2 community strings, you can use a command like this:

```
snmpwalk -v2c -c commstring 10.0.1.25
```

The `-v2c` argument chooses SNMPv2 (SNMPv2c) mode, and the `-c` argument specifies the community string. If you have the wrong community string, `snmpwalk` will run for a while and then print a timeout error.

Reviewing discovered SNMP services

When SNMP services are found during a scan, their protocol versions are tracked at the asset level where you will see `snmp` in addition to `snmp1`, `snmp2`, and `snmp3` depending on which version responded. This makes for performant queries via the [asset inventory](#), such as when querying for `protocol:snmp2`.

Additionally, SNMP services track how they authenticated and which protocols they used. If a commonly used value for a SNMPv2 community was used (such as `public`, `private`, some vendor defaults, and common values like `password`, `cisco`, and `community`), these will be listed in cleartext in the service details under `snmp.defaultCommunities` and asset details under `snmp.v2DefaultCommunities` with the list of communities that responded. If the SNMPv2 community is not a common value, this will be reported as `snmp.secretCommunities` with a value of `true` in the service details and in an attribute called `snmp.v2SecretCommunities` with a value of `true` in the asset details. If SNMPv3 is used, then `snmp.v3Usernames` will be populated at the asset level. Lastly, an asset attribute called `snmp.auth` is populated indicating whether `v2DefaultCommunity`, `v2SecretCommunity`, or `v3Username` was successful to authenticate to the asset.

Object Identifiers (OIDs)

While some of these OIDs are used only for specific vendors, this is a complete list of all OIDs potentially used on a given asset.

System attributes

- 1.3.6.1.2.1.1.*.0

IP addresses

- .1.3.6.1.2.1.4.20.1.3.*

MAC addresses

- .1.3.6.1.2.1.2.2.1.6.*

ARP caches

- .1.3.6.1.2.1.4.35.1.4.*
- .1.3.6.1.2.1.3.1.1.2.*
- .1.3.6.1.2.1.4.22.1.2.*

Routes

- .1.3.6.1.2.1.4.24.7.1.7.*

Port (CAM) tables

- .1.3.6.1.2.1.17.7.1.2.2.1.2.*
- .1.3.6.1.2.1.17.4.3.1.2.*

VLANs

- .1.3.6.1.2.1.17.7.1.4.3.1.1.*
- .1.3.6.1.4.1.9.9.46.1.3.1.1.4.1.*

Interface names

- .1.3.6.1.2.1.17.1.4.1.2.*
- .1.3.6.1.2.1.31.1.1.1.1.*

Serial numbers

- .1.3.6.1.2.1.47.1.1.1.1.11.*

Device models

- .1.3.6.1.2.1.47.1.1.1.1.13.*

Hostnames

- .1.3.6.1.4.1.9.2.1.3.*

Model names from Juniper switches

- .1.3.6.1.4.1.2636.3.1.2.0.*

Serial numbers from A10 devices

- .1.3.6.1.4.1.22610.2.4.1.6.2

Please note that Cisco Catalyst & Nexus switches require per-VLAN enumeration using either community indexing (v2) or contexts (v3). runZero automatically enumerates all ARP caches and port tables for each discovered VLAN.

runZero gathers information from SNMP v3 systems even when credentials are not available. This information includes the engineID as well as values from the following OIDs

- .1.3.6.1.6.3.15.1.1.1.0: usmStatsUnsupportedSecLevels
- .1.3.6.1.6.3.15.1.1.2.0: usmStatsNotInTimeWindows
- .1.3.6.1.6.3.15.1.1.3.0: usmStatsUnknownUserNames
- .1.3.6.1.6.3.15.1.1.4.0: usmStatsUnknownEngineIDs
- .1.3.6.1.6.3.15.1.1.5.0: usmStatsWrongDigests
- .1.3.6.1.6.3.15.1.1.6.0: usmStatsDecryptionErrors

Leveraging scan probes

When performing a scan, runZero Explorers and scanners use probes to extract information from open scanned ports. Many probes can be configured using the Probes and SNMP tab of a scan task configuration. All the ports included in the scan scope with an enabled probe will be sent a request and the response will be collected. Some probes will perform additional checks when certain protocols, services, or software are detected in the response.

Most probes can be configured differently for TCP and UDP ports:

- **TCP ports:** These scan probes will run automatically when the specific protocol is detected in the response from a TCP port.
- **UDP ports:** Most scan probes allow you to disable performing requests against UDP ports.

Probes

The table below breaks out the following for each scan probe:

- **Probe:** Name of the protocol, service, or software, linked to a source for additional details.
- **Port(s):** Ports the scanner automatically run these probes on. This is configurable for many UDP probes.
- **UDP supported:** Allows you to disable the probe on UDP, but TCP will still occur if the port is in the scan configuration.
- **Additional details:** Any additional context about the probe.

Probe	Port	UDP supported	Additional details
activemq			Banner match.
adb			Banner match, and any key/value pairs returned by the service.
airplay			
ajp			Detection of the AJP protocol via the Shodan integration .

arp			
aws- instances		Y	Integration with Amazon Web Services (AWS) APIs, can be run as scan probe or a connector task.
azure		Y	Integration with Microsoft Azure APIs, can be run as scan probe or a connector task.
azuread		Y	Integration with Microsoft Azure Active Directory (Azure AD) via Azure AD API, can be run as scan probe or a connector task.
backupexec			
bacnet	468,084,780,848,808	Y	
bedrock		Y	
bitdefender- app			
cassandra			
censys		Y	Integration with Censys Search and Censys Universal Internet Dataset, can be run as scan probe or a connector task.
checkmk	6556		Discovers Checkmk services and reports agent version, asset OS, running services, and process names.
chromecast			
citrix			
click			
coap		Y	
couchdb			
crestron	41794	Y	
crowdstrike		Y	Integration with CrowdStrike Falcon, can be run as scan probe or a connector task.
dahua-dhip		Y	Probes Dahua assets with the proprietary binary DHIP protocol and reports any details received (version, serial number, addresses)
dcerpc			
defender365		Y	Integration with Microsoft 365 Defender, can be run as scan probe or a connector task
dns	53	Y	

			details
dotnet-remoting			
drbd			
drobo-nasd			Banner match
dtls	44,333,914,433,524,600	Y	
echo			Banner match
echo (icmp)		Y	
elasticsearch			
epm			
ftp			
gcp		Y	Integration with Google Cloud Platform (GCP), can be run as scan probe or a connector task
genudp		Y	Generic UDP module for internal research purposes
giop			Banner match
googleworkspace		Y	Integration with Google Workspace, can be run as scan probe or a connector task
gpsd			Banner match
http			
http2			
icmp	623		
ike	500	Y	
imap			
infinispan			
influxdb			
insightvm		Y	Integration with Rapid7 InsightVM, can be run as scan probe or a connector task; Nexpose data can be pulled into runZero by importing files that were exported from a Nexpose instance
intune		Y	Integration with Microsoft Intune, can be run as scan probe or a connector task
ipmi		Y	Discovers and attempts to pull details from assets responding to IPMI

irc			Banner match
java-rmi			Banner match
jdbc-hsqldb			Banner match
jdwp			Banner match
jetdirect			
jms			Banner match
kerberos	88	Y	
knxnet		Y	
l2t	2228	Y	
l2tp		Y	
landesk			Detect and query limited version and config info from a LANDESK agent
lantronix	30718	Y	
layer2	22,80,135,179,443,3389,5040,7547,62078	Y	
ldap		Y	
lockdownd			
lpd			
mdns	5353	Y	
memcache	11211	Y	
minecraft			
miradore		Y	Integration with Miradore MDM via Miradore API, can be run as scan probe or a connector task
mongodb			Banner match
mountd			
mssql	1434	Y	
munin			
mysql			
mysqlx			Gathers some of the supported x-plugin capabilities of the remote service,

ndmp	10000		Gathers some basic info, like version and status, from the service.
neo4j	7474		Gathers any JSON key/value data available from the service.
nessus		Y	
netbios	137	Y	
nfs			
ntp	123	Y	
openvpn	1194	Y	
oracledb			
pca	5632	Y	
pop3			
postgresql			
pptp			
psdisco	987, 9302	Y	Sony PlayStation Discovery Protocol, our probe tried to learn things like the type of PS unit, OS version, ID, hostname, MAC address
qualys		Y	Integration with Qualys Vulnerability Management, Detection, and Respons (VMDR) via Qualys KnowledgeBase API, can be run as scan probe or a connector task
rdp			
rdns		Y	
redis			
rexec			
riak			
riak-http			
rpcbind	111, 2049	Y	
rsyncd			
rtsp			
sentinelone		Y	Integration with SentinelOne API, can be run as scan probe or a connector task
shodan		Y	Integration with Shodan API, can be run as scan probe or a connector task
sip	5060	Y	

smb2			
smb3			
smtp			
snmp	161	Y	
sonicwall-sgms			
spice			
spotify-connect			
ssdp	1900	Y	
ssh	22	Y	
subversion			
sunrpc			
SYN TCP port scan		Y	
teamviewer			
telnet			
tenable		Y	Integration with Tenable.io and Tenable Nessus Professional via Tenable cloud API , can be run as scan probe or a connector task; all versions of Tenable Nessus are supported through .nessus file imports
tftp	69	Y	
tls	443, 3389, 8443, others		
ubnt	10001	Y	Discovery protocol used by EdgeRouter devices (and possible other Ubiquiti devices)
upnp	1900, others		
vmauthd			Banner match
vmware		Y	Integration with VMware vCenter and ESXi instances, can be run as scan probe or a connector task
vnc			
wiz		Y	Integration with Wiz API, can be run as scan probe or a connector task

wlan-list		Y	Discovers WiFi networks within range of Explorers or scanners with wireless NICs
wsd	3702	Y	Discovers wsd services, which can include ONVIF data and additional ports to probe
wsman			Reports any data discovered from a wsman endpoint
zabbix-agent			Captures the version of discovered Zabbix agents
zookeeper			Banner match

Managing Explorers

The runZero Explorer is a lightweight scan engine that enables network and asset discovery. You should have at least one Explorer deployed. After deployment, you can manage your Explorers from the [Deploy page](#) in your runZero web console.

Viewing all Explorers

For each Explorer, you can see:

- The Explorer status (whether it is communicating with runZero)
- The OS it is running on
- Its name
- Any site it is associated with
- Its IP addresses
- The software version it is running
- Whether the version of npcap installed is up-to-date, if the OS is Windows (see [upgrading npcap](#) below)
- The CPU architecture of the host machine
- Any tags associated with the Explorer
- The status of its last scan
- Its capabilities, like Chrome support

Screenshot capabilities

To capture screenshots, Chrome must be installed. You can check if an Explorer has screenshot capabilities by looking for the Chrome icon in the Capabilities column.

Here's what each icon means:

- Green icon - The Explorer has access to a Google Chrome binary and can take screenshots.
- Red icon - No suitable Chrome binary was found.

Searching for Explorers

You can use the search bar to find Explorers. The query syntax is similar to other search bars in runZero, with keywords to filter by specific fields:

Keyword	Search by	Example
arch:	CPU architecture	arch:amd64
name:	assigned name	name:scanner.local
address:	IP address	address:10.0.1.200

capability:	capabilities	capability:screenshot or capability:aws
tag:	assigned tag	tag:dev
npcap_version:	npcap version	npcap_version:1.31

Explorer actions

Each Explorer has a set of action buttons that allow you to:

- **Reinstall an Explorer** - Performs a reinstall or upgrade of the Explorer. The current Explorer will download the latest Explorer code from runZero, and then run the install process.
- **Configure an Explorer** - You can associate the Explorer with a specific site, and add tags to it. You can also set the maximum number of concurrent scans allowed. A single Explorer can be configured to run multiple tasks at once.
- **Reassign an Explorer** - You can reassign an Explorer to a different organization within your account or even to a different runZero client account entirely.
- **Remove an Explorer** - If the Explorer is running, the Explorer will be asked to uninstall itself from the host machine. If the Explorer is not running, you can still tell runZero to forget about it. This is useful if you have decommissioned the machine the Explorer was running on or uninstalled the Explorer manually. If the Explorer runs again after runZero has been told to forget it, it will be readded to the registered Explorers list.

Bulk management operations

Bulk operations allow you to perform a set of actions to multiple Explorers at one time. Bulk actions are available from the Manage all Explorers menu.

You can bulk:

- **Update all online Explorers** - Tells all Explorers—that are up and communicating with runZero—to upgrade their software.
- **Forget all offline Explorers** - Clears all Explorers currently offline, and makes runZero forget them. No data will be lost. If any of the Explorers are reactivated, they will be added back to the active list.
- **Uninstall all online Explorers** - Tells all online Explorers to uninstall themselves from their host systems.
- **Automatically assign sites** - Runs through all of the Explorers that are not currently assigned to a specific site. It checks their IP address against the CIDR IP ranges of the registered subnets of all sites in the current organization. If the Explorer's IP address only matches a single site, the Explorer is assigned to that site.

Viewing Explorer details

Clicking on an Explorer's name takes you to a page showing the diagnostic information for that Explorer, including its software version, available memory, and network interfaces.

At the bottom of the page is a diagnostics text area. Clicking the Update Diagnostics button will fetch an updated list of all sub-processes active within the Explorer. This is useful to send to runZero support if you are having problems with a particular Explorer.

Traffic sampling

Community Platform

The Explorer details page is also where users can configure traffic sampling.

1. From the [Registered Explorers](#) page, select the Explorer you wish to configure to perform traffic sampling.
2. In the traffic sampling card, configure the following options:
 - **Site:** Specify the site the assets discovered as a result of Traffic Sampling will be added to.
 - **Discovery scope:** List the IP addresses or CIDR networks that traffic sampling will observe on this Explorer.
 - **Asset tags** (optional): List the tags you want applied to assets discovered through traffic sampling.
 - **Excluded hosts** (optional): List the IP addresses or CIDR networks that traffic sampling will exclude from the results.
 - **Interfaces:** Toggle the switches for the interfaces you want this Explorer to listen on.
3. Click **Save** to save your configuration and initiate the traffic sampling task.

Once configured, traffic sampling can be disabled by returning to this page and toggling off the selected interfaces. Upon saving, the traffic sampling tasks will automatically stop.

Upgrading npcap

On Windows, runZero uses a licensed third-party library called [npcap](#) for access to raw network traffic. Other software installed on the Explorer's host machine may also use npcap, and sometimes will have installed obsolete versions of the software. This can cause reliability problems.

runZero will alert you to obsolete versions of npcap by displaying a warning icon in the list of Explorers.

However, runZero cannot yet reliably upgrade npcap for you. runZero can't automatically upgrade npcap/winpcap, as it tends to be shared between applications, and forcing an upgrade from the runZero side can break other services (EDRs, Wireshark, etc).

To upgrade npcap manually:

1. Stop any running runZero services. This can be done using the Windows Services app. You'll need to look for "runZero Network Discovery Explorer".
2. Stop any other running software which uses npcap.
3. Uninstall Winpcap and any npcap installations via the Windows Control Panel.
4. Reboot the computer.

runZero will restart automatically, and install the latest npcap.

Managing scan templates

A scan template is a predefined set of scan options and settings. If you have a scan configuration you use often, you can create a scan template to save those settings. The next time you create a scan, you can choose a template instead of manually configuring your settings. Each update you make to the scan template is automatically applied to new and recurring scans based on the template, as well as any queued scans which were set up using the template but have not started yet. With scan templates, you can save time and reduce the likelihood of misconfiguring a scan.

When you create a scan based on a scan template, you will not be able to edit the fields set by the template. However, you can configure the site, scan name, discovery scope, Explorer, and scan schedule, since these are not defined in the template.

Each time a scan runs using values from a template, the scan task is saved with a copy of the parameters. This means the task will list the values used for the scan, even if the template is modified after the scan completes.

Creating a scan template

Scan templates can be created in a few ways in runZero:

- By going to **Tasks > Task library**
- From an existing scan task
- Based on an existing scan template

Creating a scan template from scratch

1. Go to **Tasks > Task library** to view the [task library](#).
2. From the **Task library**, click **Add template**.
3. When the scan configuration form appears, enter a name for the template. Use this name to search and assign the template to a new scan.
4. [Configure the scan](#) as you normally would. The following tabs are available for you to access and configure settings:
 - **Standard** - Provide a name, description, and scan rate for the template.
 - **Advanced** - Configure excluded hosts, included ports, tags, host rates, group size, UDP probe max attempts, max TTL for all scan probe packets, ToS, screenshots, and subnet sampling.
 - **Probes and SNMP** - Enable SNMP and additional probes, and provide credentials. When you enable the SNMP probe, the scan will identify devices that are using the SNMP protocol, using the port and credentials provided to find anything with SNMP enabled.
 - **Credentials** - Enable any preconfigured [credentials](#) you want to use for the scan. The scan uses the credentials if the defined CIDR scope matches the specific scan target.
5. Save the template.

To use this template, go to the Templates tab when you [configure a scan](#) or [create a Template scan](#) from the scan menu on the [inventory](#) or [tasks](#) pages.

Creating a scan template from a scan task

1. Go to **Tasks** to open the [Tasks overview page](#).
2. Find the scan task you want to use as the basis for your scan template and click the name to view the task details.
3. From the task details, click the **Copy** dropdown menu and choose **Copy for new template**.
4. When the scan configuration form appears, enter a name for the template. Use this name to search and assign the template to a new scan.
5. [Configure the scan](#) as you normally would. The following tabs are available for you to access and configure settings:
 - o **Standard** - Provide a name, description, and scan rate for the template.
 - o **Advanced** - Configure excluded hosts, included ports, tags, host rates, group size, UDP probe max attempts, max TTL for all scan probe packets, ToS, screenshots, and subnet sampling.
 - o **Probes and SNMP** - Enable SNMP and additional probes, and provide credentials. When you enable the SNMP probe, the scan will identify devices that are using the SNMP protocol, using the port and credentials provided to find anything with SNMP enabled.
 - o **Credentials** - Enable any preconfigured [credentials](#) you want to use for the scan. The scan uses the credentials if the defined CIDR scope matches the specific scan target.
6. Save the template.

To use this template, go to the Templates tab when you [configure a scan](#) or [create a Template scan](#) from the scan menu on the [inventory](#) or [tasks](#) pages.

Creating a scan template based on a template

1. Go to the **Tasks > Task library** to view the [task library](#).
2. From the **Task library**, find the template you want to use as the basis for a new template and click the name to open it.
3. From the scan template configuration page, click **Copy**. runZero will duplicate and create a new scan template.
4. When the scan configuration form appears, enter a name for the template. Use this name to search and assign the template to a new scan.
5. [Configure the scan](#) as you normally would. The following tabs are available for you to access and configure settings:
 - o **Standard** - Provide a name, description, and scan rate for the template.
 - o **Advanced** - Configure excluded hosts, included ports, tags, host rates, group size, UDP probe max attempts, max TTL for all scan probe packets, ToS, screenshots, and subnet sampling.
 - o **Probes and SNMP** - Enable SNMP and additional probes, and provide credentials. When you enable the SNMP probe, the scan will identify devices that are using

- the SNMP protocol, using the port and credentials provided to find anything with SNMP enabled.
 - **Credentials** - Enable any preconfigured [credentials](#) you want to use for the scan. The scan uses the credentials if the defined CIDR scope matches the specific scan target.
6. Save the template.

To use this template, go to the Templates tab when you [configure a scan](#) or [create a Template scan](#) from the scan menu on the [inventory](#) or [tasks](#) pages.

Applying a scan template to a scan

Instead of manually configuring a scan, you can choose to use a template instead. You can find all templates available for a scan on the **Templates** tab in your scan configuration.

You can also go to the **Task library** page and choose to **Create a scan based on this template** from the template's actions.

Editing a scan template

All updates to a scan template affect every scan that uses it. After you make a change to a template, the next scan will automatically update to use the new settings. Any scan queued to run will use the new settings.

Editing a scan that uses a scan template

Any changes to a template will affect all scans that use them. If a scan uses a template, you will need to edit the template in order to modify the scan settings. Otherwise, you will need to delete the scan and create a new scan configuration with your modifications.

When you create a scan based on a scan template, you will not be able to edit the fields set by the template. However, you can configure the standard fields, such as the site, scan name, discovery scope, Explorer, and scan schedule, since these are not defined in the template.

Deleting a scan template

Deleting a template will remove it permanently from runZero. All configurations will no longer be accessible. To delete a scan template, go to [the Task library](#), find the template you want to delete, and click **Delete**. A modal will prompt you to confirm that you want to delete the template, and the action will be irreversible.

Identifying gaps in scanning

Background

After you have run a full network discovery scan, you can start to better understand your coverage and begin to optimize. By the end of this guide, you will understand how to use the out of the box reports in runZero to understand your gaps in network coverage.

RFC 1918 coverage

The first report to look at is the **RFC 1918 coverage** report. This report shows you which internal IPv4 subnets have been scanned, which likely contain assets, and which are still unknowns.

The scan coverage maps show all the addresses scanned within the 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 ranges. See the legend to understand what percentage of each address space has been scanned. Clicking into any of the scanned subnets gives you access to the subnet grid for deeper asset analysis.

Identify scanned and un-scanned areas with the coverage map: On the flip side, red outlines indicate that there are un-scanned addresses runZero has indirect knowledge of that haven't been scanned directly. For example, this can happen when runZero finds a secondary IP address on a multi-homed device within a scanned subnet. The red boxes highlight the subnets most likely to be in use, but un-scanned.

Scan missing subnets: From the coverage report, you can launch a scan for any missing subnets in a given RFC1918 block – look for the binocular icon.

Scan missed subnets: The missing subnets will be shown in the scan scope and the subnet ping will be enabled by default. You can tune the scan configuration as needed for your environment.

Subnet utilization

The **Subnet utilization** report can provide similar visibility into your network to the RFC 1918 coverage report, but with emphasis on the subnets that contain live assets. This report will enumerate each of the subnets defined in your site definitions, and provide a count of live assets for that site and subnet, along with a utilization percentage. If there are live assets that are outside any site subnet, they will be aggregated into an inferred subnet based on the network mask size you select.

From this report, you can pivot to the asset inventory for a given subnet or initiate a new scan of a subnet. Another benefit of this report is that you can export the results as a CSV. This can be helpful for more complex data analysis and for scheduling recurring scans.

Switch topology

View layer-2 link information extracted from SNMP-enabled switches. This report can be used to find unmapped assets and investigate why they aren't showing up in your scans.

Configuration for this report: The Switch Topology report uses data enumerated via SNMP to map switch ports to assets. In environments where SNMP v1 or v2 with default public or private communities are in use, this enumeration happens automatically. Non-default communities for SNMP v1/v2/v3 can also be provided in the scan configuration. Clicking on a node in this report will expand it to show its connections.

Finding unmapped MACs: This topology view is helpful when trying to understand how a given asset or switch is connected, but also provides a critical data point related to risk; the number of unmapped assets. An unmapped asset is a MAC address connected to a switch, but not found in an ARP cache or through any of the other techniques runZero uses for remote MAC address discovery.

Re-scan to properly map MAC addresses: For environments where a runZero agent is connected to each network segment, unmapped MACs may highlight VLANs or network segments that are missing from the scan scope. In environments where runZero is scanning assets multiple hops away, the unmapped asset count can provide an estimate of how well the remote segment is being identified.

Unmapped MACs

If you prefer a condensed view to see all unmapped macs, you can also use the Unmapped MACs report. This report shows all the unmapped MACs in your organization, organized by switch, and again by port. This data can be used similarly to the unmapped MACs data from the Switch Topology report.

Credentials

The Credentials page provides a single place to store any secure credentials needed by runZero, including:

- SNMPv3 credentials
- Access secrets for cloud services like AWS and Azure
- API keys for services such as Censys and Miradore

Credentials are stored in encrypted form in the runZero database. Credentials, such as SNMP passwords, are used by runZero Explorers and are transmitted to them in encrypted form. For security reasons, the secret part of any credential cannot be viewed once entered.

When adding a credential, you can choose to make it a global credential that can be used for all organizations or to allow access only by specific organizations. The **Allow all** or **Disallow all** buttons let you quickly apply the same setting across all organizations. Individual organizations can also be toggled to allow or disallow access.

Most credential fields can be edited after the credential is saved. Some fields, like URLs, cannot be edited after saving for security reasons. Sensitive fields, such as passwords or access keys, will be hidden but can be overwritten.

Credential settings

The specific fields and options for a credential depend on the type of credential.

VMware and SNMP credentials, which are used by the runZero Explorer, allow a CIDR allow list to be specified. This can be used to limit which scanned IP addresses the credential will be used with. This feature allows you to avoid sending SNMP or VMware credentials to all scanned hosts on the network, and instead limit them to specific IP addresses or ranges.

Credential verification

Credentials can be verified when created or edited to ensure they can successfully authenticate. Choose **Verify & save** when creating or editing a credential to run the verification before saving. If verification fails, it will display an error message and then give you the option to **Save anyway**. **Save anyway** will save the most recent verification status.

Credential management

Users must have administrator-level permissions to manage credentials. Users with **Administrator** as their [default role](#) can fully manage all credentials. Users with [per-organization permissions](#) do not have access to global credentials, and are only able to manage credentials in organizations where they have administrator permissions. A shared credential cannot be deleted by an organization administrator.

Using custom fingerprints

Professional

Community

Platform

Customers running a self-hosted instance or using the standalone scanner have the ability to use custom-written fingerprints. This can be useful in adding new fingerprint coverage for very unique or custom assets and services, such as device prototypes or proprietary applications/services. Custom fingerprints can also be used to override existing, similar runZero fingerprints by using a same-or-higher certainty value.

Note: when using the runZero standalone scanner with custom fingerprints, you'll need to use the `RUNZERO_EXTERNAL_FINGERPRINTS` value as an environment variable when launching the scanner.

Create new fingerprints

Custom fingerprints follow the structure and format of the [open-source Recog fingerprint database](#). You can author your own fingerprint XML entries in files of similar name and format to [those found in Recog](#). For cases where an asset or service matches both a built-in runZero fingerprint and a custom fingerprint of the same kind, preference will be given to the fingerprint with higher “certainty” value(s) (e.g. `hw.certainty`, `os.certainty`, `service.certainty`). In the event of a certainty “tie” (i.e. same certainty value(s)), the custom fingerprint will be given preference.

Add new fingerprints to your self-hosted runZero instance

To ensure the self-hosted instance of runZero can properly access your custom fingerprints, they will need to exist within the runZero installation directory (`/opt/runzero` by default). The following steps will get your custom fingerprints setup for use by runZero:

- Create a new directory within the runZero install directory (e.g. `mkdir /opt/runzero/myfingerprints`)
- Update your `/etc/runzero/config` file with the new directory location (see below)
- Copy your custom fingerprint XML files into the new directory (e.g. `cp *.xml /opt/runzero/myfingerprints`)

The `/etc/runzero/config` file needs to be updated with the location of your custom fingerprints directory, which you can do by using your favorite editor to add the following line to the end of your config file (swapping `myfingerprints` with the name of the directory you added):

```
RUNZERO_EXTERNAL_FINGERPRINTS=myfingerprints
```

Note that you only need to do the directory creation and update of `/etc/runzero/config` file once. After that, you can add, remove, or modify your custom fingerprints in the directory as needed and then restart the runZero application to reload the current custom fingerprints.

Restart the runZero service

The runZero self-hosted instance will need to be [restarted](#) when custom fingerprints are added, removed, or updated. On restart, the runZero log file can be monitored to verify if the custom fingerprints were successfully applied or if an error occurred. On successful load of custom fingerprints, a log message like the following should be present:

```
@cee:{"level":"info","msg":"loaded (5) external fingerprints","time":"2022-09-12T19:51:49Z"}
```

If issues are encountered when loading or processing custom fingerprints, or if a 10 second timeout is reached, a warning message will be logged and the runZero application will continue running without any custom fingerprints.

Verify your fingerprints

Once your custom fingerprints have been added to your self-hosted runZero instance and the instance restarted, you can verify that the custom fingerprints are performing as-expected in one of the following ways:

- Running a scan task to go scan a relevant asset/service, or
- Importing an existing scan data file of the relevant asset/service

Following completion of the scan task OR the import of scan data, depending on which method you chose, you can then navigate to a [specific asset](#) or [specific service](#) and verify that your custom fingerprints are acting as-expected.

Integrating with runZero

Inbound integrations

Enriching runZero results with data from other tools

The runZero platform offers integrations with several sources of asset data, allowing users to enrich their asset inventory and identify assets and subnets that are not effectively managed or protected. By leveraging product APIs and export/import functionality, runZero can pull data from many IT and security tools to extend visibility across your organization's network.

Note: For best results, scan devices with runZero before setting up an inbound integration to pull in data for those assets. A runZero scan will provide more data than many integrations, allowing for better matching of incoming data with existing assets.

Supported integrations

Cloud and virtualization

- [Amazon Web Services](#)
- [Google Cloud Platform](#)
- [Microsoft Azure](#)
- [VMware](#)
- [Wiz](#)

Endpoint protection

- [CrowdStrike Falcon](#)
- [Microsoft 365 Defender](#)
- [Microsoft Intune](#)
- [Miradore MDM](#)
- [SentinelOne](#)

Asset and identity management

- [Azure Active Directory](#)
- [Google Workspace](#)
- [Microsoft Active Directory](#)

Vulnerabilities and risk

- [Censys Search & Data](#)
- [Qualys VMDR](#)
- [Rapid7](#)
 - [InsightVM](#)
 - [Nexpose](#)
- [Shodan](#)
- [Tenable](#)
 - [Tenable Vulnerability Management](#)
 - [Tenable Nessus Professional](#)
 - [Tenable Security Center](#)
 - [Tenable Nessus \(file import\)](#)

Custom integrations

Platform

If the solution you want to draw data from isn't available as a current runZero integration, Platform users can leverage the [custom integrations](#) feature to add asset data from custom sources. Adding custom asset sources can be accomplished through the API or by leveraging the [runZero Python SDK](#).

Scan probes or connector tasks

Most integrations can be run either as a scan probe or a connector task.

Scan probes run as part of a scan task. The scan task can be used to scan your environment and sync integrations at the same time. To run an integration as a scan probe:

1. Configure a scan task from the Scan menu in your inventory or tasks page.
2. Activate the integration under the Probes tab.
3. Activate the correct credential under the Credentials tab.
4. Configure, activate, or deactivate other scan task configuration options as preferred.

Connector tasks run independent of scan tasks in order to allow more finely tuned scheduling of integration syncs and asset scans. Connector tasks are run from the runZero cloud by default, but can be configured to run from an Explorer in your organization if preferred. To run an integration as a connector task:

1. Configure a connector task from the [Integrations page](#) or the Integrate menu in your inventory or tasks page.
2. Select an Explorer from the Explorer menu (optional).
3. Configure, activate, or deactivate other connector task configuration options as preferred.

Importing integration data

Some integrations can be used by importing data from that platform into runZero. For example, .nessus files from [Tenable Nessus](#) and .xml files from [Rapid7 Nexpose](#) can both be ingested without requiring a connection to their APIs.

Automatic asset merge

How runZero maps integration assets to assets:

- For hosts that can be matched to an existing runZero asset, asset-level attributes will be updated, and integration-specific attributes will be added.
- For hosts that cannot be matched with an existing runZero asset, a new asset will be created in the site specified when the integration task is set up.

runZero is able to merge integration data into existing assets by the following, in priority order:

1. MAC address
2. IP address (3-day window)
3. Hostname

Assets from integrations can also be manually merged into runZero assets using the Merge button on the [Asset Inventory](#) page.

Removing an integration data source

When an integration is removed as a data source, the associated attributes are removed from your runZero assets. Since some asset attribute fields are merged, it is possible that attributes populated by both runZero scans and the integration could be deleted. Rescanning the affected assets will resolve this issue.

Source names and IDs

The table below maps the source name to the source ID for querying assets and vulnerabilities.

ID	Name	Description
1	runzero	runZero
2	miradore	Miradore
3	aws	Amazon Web Services
4	crowdstrike	CrowdStrike
5	azure	Microsoft Azure

6	censys	Censys
7	vmware	VMWare
8	gcp	Google Cloud Platform
9	sentinelone	SentinelOne
10	tenable	Tenable.io & Nessus
12	rapid7	Rapid7 Nexpose & InsightVM
14	qualys	Qualys VMDR
15	shodan	Shodan
16	azuread	Azure AD
17	ldap	Active Directory (LDAP)
18	ms365defender	Microsoft 365 Defender
19	intune	Microsoft Intune
20	googleworkspace	Google Workspace
21	packet	runZero traffic sampling
22	tenablesecuritycenter	Tenable Security Center

Amazon Web Services

Professional Community Platform

runZero integrates with Amazon Web Services (AWS) to provide better visibility across your cloud environment. This integration imports data from each applicable API to add detailed information to your asset inventory:

- [AWS EC2 API](#)
- [AWS RDS API](#)
- [AWS ELBv1 API](#)
- [AWS ELBv2 API](#)
- [AWS Lambda API](#)

Syncing with AWS allows you to quickly identify the number of EC2 instances, elastic load balancers, and relational database services you have running, as well as their region, account, and more.

This integration supports the import of all running EC2 instances, RDS instances, and active application, network, gateway, or classic load balancers. It can be configured to connect to a single AWS account or to all accounts in your organization and imports data across multiple regions.

Getting started

The following AWS resource types are supported:

- EC2 instances
- Elastic load balancers
- RDS instances
- Lambda instances

To set up the AWS integration, you'll need to:

1. Configure AWS to allow API access through runZero.
2. Add the AWS credential to runZero, which includes the access key and secret key.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the AWS integration to sync your data with runZero.

Requirements

Before you can set up the AWS integration:

- Make sure you have access to the AWS console.
- Make sure you are using [AWS Organizations](#) if you want to connect to multiple accounts.

Step 1: Configure AWS to allow API access through runZero

To connect to a single AWS account:

1. Log into the AWS console.
2. Go to **Identity and Access Management (IAM) > Users** and select the user that will provide API access to runZero.
3. Click **Add permissions > Attach existing policies directly**. Search for and attach the following policies based on the services you would like to sync:
 - AmazonEC2ReadOnlyAccess to sync EC2 and ELB resources
 - AmazonRDSReadOnlyAccess to sync RDS instances
 - AWSLambda_ReadOnlyAccess to sync Lambda functions
4. From the user summary screen, open the **Security credentials** tab and click on **Create access key**.
5. Save the **Access key ID** and **Secret access key**.

To connect to all accounts in your organization:

1. Log into the AWS console.
2. For each account in your organization, create a role and assign the AmazonEC2ReadOnlyAccess, AmazonRDSReadOnlyAccess, and AWSLambda_ReadOnlyAccess

policies. You can do this one at a time for each account or use StackSets to deploy the role if you have a large number of accounts:

Option 1: Create a role using IAM

1. Go to **Identity and Access Management (IAM) > Roles** and click **Create role**.
2. Choose **Another AWS Account** for the type of trusted entity.
3. For Account ID, enter the ID for your organization's management account.
4. Click **Next: Permissions**.
5. Attach the AmazonEC2ReadOnlyAccess policy if you want to sync EC2 and ELB resources.
6. Attach the AmazonRDSReadOnlyAccess policy if you want to sync RDS instances.
7. Attach the AWSLambda_ReadOnlyAccess policy if you want to sync Lambda functions.
8. Click **Next: Tags** and add tags optionally.
9. Click **Next: Review** and provide a name for the role. (The role must be named the same for all accounts)
10. Click **Create role**.

Option 2: Create and deploy a role to multiple accounts using StackSets

1. While logged into your organization's management account, go to **CloudFormation > StackSets** and click **Create StackSet**.
 2. Select **Template is ready** and upload a file with the JSON template for the stackset provided below. Replace <accountID> with the account ID where the role was created, <userName> with the name of the user, and <roleName> with the name of the role. Click next.
 3. Enter a name for the StackSet. Click next.
 4. Optionally set tags. Click next.
 5. Set the deployment options. Click next.
 6. Review and create the StackSet.
3. For your organization's management account, create an inline policy to allow the STS AssumeRole action.
 1. Go to **Identity and Access Management (IAM) > Users** and select the user that will provide API access to runZero.
 2. Click **Add inline policy**.
 3. In the JSON tab, enter the inline policy text from below, replacing <rolename> with the role name.
 4. From the user summary screen, open the **Security credentials** tab and click on **Create access key**.
 5. Save the **Access key ID** and **Secret access key**.
 6. From the user summary screen, open the **Permissions** tab and click **Add permissions**. Attach the AWSOrganizationsReadOnlyAccess policy.
 7. (Optional) Attach the AmazonEC2ReadOnlyAccess policy to your organization's management account if it has EC2 or ELB instances you would like to sync.

8. (Optional) Attach the `AmazonRDSReadOnlyAccess` policy to your organization's management account if it has RDS instances you would like to sync.
9. (Optional) Attach the `AWSLambda_ReadOnlyAccess` policy to your organization's management account if it has Lambda functions you would like to sync.

StackSet template

```

{
  "Resources": {
    "IAMrunZeroRole": {
      "Type": "AWS::IAM::Role",
      "Properties": {
        "AssumeRolePolicyDocument": {
          "Statement": [ {
            "Action": "sts:AssumeRole",
            "Effect": "Allow",
            "Principal": {
              "AWS": "arn:aws:iam::<accountID>:user/<userName>"
            }
          } ],
          "Version": "2012-10-17"
        },
        "ManagedPolicyArns": [
          { "Fn::Join": [ "",
            [
              "arn:",
              { "Ref": "AWS::Partition" },
              ":iam::aws:policy/AmazonEC2ReadOnlyAccess"
            ]
          ] },
          { "Fn::Join": [
            "",
            [
              "arn:",
              { "Ref": "AWS::Partition" },
              ":iam::aws:policy/AmazonRDSReadOnlyAccess"
            ]
          ] }
        ],
        "RoleName": "<roleName>"
      }
    }
  }
}

```

Inline policy template

```

{
  "Version": "2012-10-17",

```

```
{
  "Effect": "Allow",
  "Action": "sts:AssumeRole",
  "Resource": "arn:aws:iam::*:role/<rolename>"
}
]
```

Step 2: Add the AWS credential to runZero

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like AWS_EC2.
3. Choose **AWS Access & Secret** from the list of credential types.
4. Provide the following information:
 - o **AWS access key** - Access key ID obtained from **User > Security credentials > Create access key**.
 - o **AWS secret access key** - Secret access key obtained from **User > Security credentials > Create access key**.
 - o **AWS role** - Assumed role used to connect to other accounts in your organization. It should be named the same across accounts.
 - o Select the region(s) that you want to sync.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
6. Save the credential. You're now ready to set up and activate the connection to bring in data from AWS.

Step 3: Choose how to configure the AWS integration

The AWS integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the AWS integration to sync data

After you add your AWS credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the AWS integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new AWS-only assets are created.

1. Activate a [connection to AWS](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).

2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like `AWS sync`.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**:
 - Choose the site you want to add your assets to. All newly discovered assets will be stored in this site. You can also choose to **Automatically create a new site per VPC** or **Automatically create a new site per account**, and runZero will take care of creating the sites for newly discovered assets or accounts.
 - Choose whether to **automatically delete stale AWS assets**. If you check this option, runZero will automatically delete AWS assets previously seen in AWS that were not found in the most recent sync.
 - Choose whether to include AWS assets that are not currently running. If you check this option, runZero will import AWS asset data for assets that are not in a running state.
6. Under **Service options**, select the services you would like to sync data from. You must choose at least one.
7. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
8. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the AWS integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the AWS toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the AWS toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View AWS assets

After a successful sync, you can [go to your inventory](#) to view your AWS assets. These assets will have an AWS icon listed in the **Source** column.

To filter by AWS assets, consider running the following queries:

- [View all AWS assets:](#)

```
source:aws
```

- [View all AWS EC2 instances:](#)

```
source:aws AND has:"@aws.ec2.instanceID"
```

- [View all AWS Elastic Load Balancers:](#)

```
source:aws AND (has:"@aws.elb.loadBalancerArn" OR has:"@aws.elb.loadBalancerName")
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by the AWS APIs.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Amazon Web Services integration unable to connect?

1. Are you getting any data from the AWS integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the AWS integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Enriching scans with EC2

Professional Community Platform

As part of a discovery scan, runZero will automatically enrich scanned assets with data from the AWS EC2 API when available. runZero assets will be updated with internal IP addresses, external IP addresses, hostnames, MAC addresses, and tags, along with other EC2-specific attributes, such as the account ID and instance type.

No additional configuration is needed in runZero to get this data enrichment. However, you may need to [modify the permissions](#) associated with the instance's IAM role.

Find Explorers with EC2 enrichment capabilities

To use the EC2 enrichment capabilities, the runZero Explorer must be running on an instance with permissions to describe your EC2 instance list. This can be configured through an IAM role associated with the instance as well as by configuring the AWS credentials for the root user account.

To identify the Explorers with this capability, view your [registered Explorers](#). Any Explorer with the cloud icon indicates that it can enumerate EC2 instances.

Scans run from these Explorers will merge any EC2 instance fields into the asset automatically for any in-scope targets matched against the instance list.

Add permissions to describe instances

To allow for EC2 instance data enrichment, you will need to add the EC2 [ec2:DescribeInstances](#) permission for the instance role configured for your instance. Visit the [Amazon docs](#) to learn how to create and update policies.

Basically, your policy will look like:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "VisualEditor0",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances",
    ],
    "Resource": "*"
  }]
}
```

From the IAM UI, go to **Roles > Permissions > Attach policies** and search for the EC2 service. From the actions, select `DescribeInstances`, which is located under `List`.

You can also configure credentials on the instance by running `aws configure` as root, instead of using the IAM instance role.

After you save your policy, you'll need to restart your Explorer. The easiest way to do this is to force an update from the Explorer menu.

Keep in mind if your configuration uses one region, but the instance is located in another, it will use the instance's region instead for all API requests.

Attributes runZero gets from the EC2 API

When runZero determines that an IP address is also an EC2 instance, it will enrich the existing runZero data with EC2 metadata. runZero will be able to pull in more hostnames based on AWS asset tags, MAC addresses, internal and external IPs.

Some attributes that runZero is able to get from the EC2 API metadata:

- `aws.accountID`
- `aws.architecture`
- `aws.availabilityZone`
- `aws.hypervisor`
- `aws.imageID`

- aws.instanceID
- aws.instanceType
- aws.ipv4
- aws.ipv6
- aws.keyName
- aws.launchTimeTS
- aws.macs
- aws.privateDNS
- aws.privateIP
- aws.publicDNS
- aws.publicIP
- aws.region
- aws.rootDeviceName
- aws.rootDeviceType
- aws.scanner.instanceID
- aws.scanner.instanceType
- aws.state
- aws.subnetID
- aws.tags
- aws.tenancy
- aws.virtualizationType
- aws.vpcID

Additionally, runZero will also report on other things that respond to the scan, but don't match an EC2 entry. You may see things like Amazon RDS, temporary ELBs, and Lambdas in your inventory, but you will only see EC2 metadata on EC2 assets.

Azure Active Directory

Professional Community Platform

runZero integrates with Azure AD to allow you to sync and enrich your asset inventory, as well as gain visibility into Azure AD users and groups. Adding your Azure AD data to runZero makes it easier to find assets that are not part of your domain.

Getting started

To set up the Azure AD integration, you'll need to:

1. Configure Azure AD to allow API access through runZero.
2. Add the Azure AD credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the Azure AD integration to sync your data with runZero.

Requirements

Before you can set up the Azure AD integration, make sure you have access to the Microsoft Azure portal.

Step 1: Register an Azure application for Azure AD API access

runZero can authenticate to the Azure AD API using either a username and password or a client secret. Register an application to configure Azure AD API access.

1. Log into the Microsoft Azure portal.
2. Go to **Azure Active Directory > App registrations** and click on **New registration**.
 - Provide a name.
 - Select the supported account types.
 - Optionally add a redirect URI.
3. Click **Register** to register the application.
4. Once the application is created, you should see the **Overview** dashboard. Note the following information:
 - Application (client) ID
 - Directory (tenant) ID
5. From the application's details page, go to Authentication. Set Allow public client flows to Yes and then save the configuration.
6. Go to API permissions > Add a permission.
7. Select Microsoft Graph from the list of Microsoft APIs.
8. Select the correct permissions type for your needs:
 - **Username & password**: select Delegated permissions
 - **Client secret**: select Application permissions
9. Search for and select the following required permissions:
 - Device.Read.All
 - Group.Read.All
 - User.Read.All
10. Click **Add permissions** to save the permissions to the application.
11. If using a client secret, also perform the following steps:
 - Navigate to **Azure Active Directory > App registrations** and select the application you created.
 - Go to **Certificates & secrets** and click on **New client secret**.
 - Enter a description.
 - Select the expiration.
 - Click **Add** to create the client secret and save the client secret value.

Add the Azure AD credential to runZero

Step 2a: Add an Azure Username & Password credential to runZero

1. Go to the [Credentials](#) page in runZero and click **Add Credential**.
2. Provide a name for the credential, like Azure User/Pass.
3. Choose **Azure Username & Password** from the list of credential types.
4. Provide the following information:
 - **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.

- **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure username** - The username for your Azure cloud account. This cannot be a federated user account.
 - **Azure password** - The password for your Azure cloud account.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
 6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 2b: Add an Azure Client Secret credential to runZero

This type of credential can be used to sync all resources in a single directory (across multiple subscriptions).

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like **Azure Client Secret**.
3. Choose **Azure Client Secret** from the list of credential types.
4. Provide the following information:
 - **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure client secret** - To generate a client secret, go to **Azure Active Directory > App registrations**, select your application, go to **Certificates & secrets** and click on **New client secret**.
 - **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - Select the **Access all subscriptions in this directory (tenant)** option to sync all resources in your directory. Otherwise, specify the **Azure subscription ID** - The unique ID for the subscription that you want to sync. This can be found in the Azure portal if you go to **Subscriptions** and select the subscription.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 3: Choose how to configure the AzureAD integration

The Azure AD integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the Azure AD integration to sync data

After you add your Azure AD credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the Azure AD integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Azure AD-only assets are created.

1. Activate a [connection to Azure AD](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure it has access to the organization you are currently in.
3. Enter a name for the task, like Azure AD sync.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. If you want to include assets in your integration import that the Azure AD account has marked as inactive, switch the **Include inactive assets** toggle to Yes. By default, the integration will not include assets that are marked as inactive.
8. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Azure AD integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Azure AD toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the Azure AD toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View Azure AD assets

After a successful sync, you can [go to your inventory](#) to view your Azure AD assets. These assets will have an Active Directory icon listed in the **Source** column.

To filter by Azure AD assets, consider running the following queries:

- [View all Azure AD assets:](#)

```
source:azuread
```

- [View runZero assets not connected to Azure AD:](#)

```
source:runzero AND NOT source:azuread
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by Azure AD.

[Community](#) [Platform](#)

The Azure AD integration provides details about users and groups in addition to enriching asset inventory data. Go to [Inventory > Users](#) or [Inventory > Groups](#) to view the data provided by Azure AD.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Azure Active Directory integration unable to connect?

1. Are you getting any data from the Azure AD integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Azure AD integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

How do I solve the following Azure AD errors?

- (invalid_client) AADSTS7000218: The request body must contain the following parameter: 'client_assertion' or 'client_secret'

1. This error can be corrected by enabling Allow Public Client Flows in Azure AD. This can be accomplished by entering the Application details page under **Authentication > Advanced Settings**. From here you can toggle the Allow Public Client Flows setting to Yes.
 2. It can also be helpful to ensure that application permissions were granted correctly when registering the Azure application for Azure AD API access. To do this, navigate to the API permissions settings page and ensure that each of the API/Permissions have the type set to application. Also make sure that the permission granted is Grant Admin Consent for Default Directory.
- failed to get Azure AD groups: invalid response: 403 (403 Forbidden)
 1. This error is likely due to an issue with credentialing. Please review the documentation and check your credentials to ensure everything was entered correctly and no steps were accidentally skipped.

Censys Search & Data

Community Platform

runZero supports importing assets from the [Censys Search API](#) and the [Censys Internet Dataset](#).

- [Importing assets from the Censys Search API](#)
- [Importing assets from the Censys Universal Internet Dataset](#)

Censys Search API

To get started with the Censys Search API, you will need to register for a [Censys Search account](#). Once you have done so, you can find your API credentials in the [My Account](#) section.

In runZero, go to the Credentials page, and click [Add Credential](#). Select *Censys Search API Key* as the credential type, and enter your API ID and API secret.

You can now go to your [asset inventory](#), click the **Connect** button, and choose **Censys Search API**. Select the credential you just created from the Censys Search credential dropdown.

Configuration

There are two modes for connecting runZero to the Censys Search API.

- **Custom Query** mode - runZero runs a Censys search query you specify, and then imports all of the results into runZero. The search query should be in [Censys Search Language](#). It is a good idea to test your query using the main [Censys Search 2.0](#) interface before running an import task.
- **All Assets** mode - runZero assembles a list of public IP addresses from all of the assets in the selected site, and then uses the API to find Censys Search information for those

addresses. The information found is imported into runZero and merged into the appropriate assets.

As with a runZero scan, you'll need to select a site to contain the scan data. The usual task scheduling options are available.

When you have finished editing the Censys Search configuration, click **Activate Connection**.

Censys Universal Internet Dataset

To get started with the Censys Universal Internet Dataset API, you will need a paid Censys Data account and the associated API credential. You can find your API credentials in the [My Account](#) section.

The dataset can be downloaded by following the instructions in the [Censys documentation](#). The Search API is used to get a list of files for a given date and those individual files should be downloaded into a local directory backed by SSD or NVMe storage.

Creating the database

The raw files are in Apache Avro format and need to be converted into a database for efficient queries.

To process Censys data files, you use the runZero CLI scanner's `censys-db-convert` command. This command takes two parameters:

- The path to a directory containing the `.avro` files from Censys
- The path to write the computed database

```
$ nice runzero-scanner censys-db-convert /home/censys/avro /home/censys/db
```

The default configuration requires substantial computing resources:

- At least 8 CPU cores, but 16 or more is better
- At least 64GiB of RAM, but more is better
- At least 3Tb of storage backed by SSD or NVMe (1Tb+2Tb or single volume)

An AWS `m5.4xlarge` with a 3Tb GP2 SSD volume meets these requirements and can process a full dataset (single day) in about 13 hours. The resulting database is about twice the size of the source data (1.3TiB database from 640GiB of Avro). Using the database requires additional disk overhead and over provisioning the storage also improves throughput.

Querying the database

After the Avro files have been converted to a local database, the `censys-db` command can be used to import data into runZero.

The scanner queries the local database, and writes a file in runZero scan format containing the appropriate host records. By default, the file has a name matching `censys-*.rumble.gz`

and is written to the current directory. Alternatively you can specify an output filename with the `--output-raw` option, as if performing a runZero scan.

The runZero scan file can be uploaded to the runZero console like any other scan file.

If you have more IP addresses or CIDRs than will fit on a command line, you can use the `--input-targets` option to specify that the scanner should read them from a file. The file is expected to be ASCII text, and contain CIDRs or IP addresses separated by whitespace (which can include newlines).

You can also use the scanner to process data, upload it, and then delete the scan data file if everything succeeded. For example:

```
% runzero-scanner censys-db /home/censys/db \
  12.216.190.0/24 --upload --api-key=YOUR_ORGANIZATION_API_KEY \
  --upload-site="Primary site"
```

If you are using self-hosted runZero, you can use the `--api-url` option to specify your console's API endpoint.

The censys command also supports the `--verbose` option, which will make it list host addresses as they are written to the output file.

Creating a local Censys Search API server

The computed database can also be used to serve a limited, local version of the Censys Search API using the runZero scanner's `censys-db-server` command.

Due to the size of the database, the system `vm.max_map_count` may need to be increased to avoid a memory map error. The memory map count can be increased by adding the following line to `/etc/sysctl.conf`:

```
vm.max_map_count=262144
```

Once this line is added, reload the `sysctl.conf` with the following command:

```
$ sudo sysctl -p /etc/sysctl.conf
```

After the `vm.max_map_count` has been updated, start the Censys DB Server with the following command:

```
$ runzero-scanner censys-db-server /home/censys/db
```

This will start a local web service on port 55555 by default (changeable via `--port <val>`) that responds to the `/api/v2/hosts/search` and `/api/v2/hosts/<ip>` endpoints. Once this server is running, it can be queried using the runZero Censys Search API connector, and through other HTTP clients, such as `curl`:

```
$ curl http://127.0.0.1:55555/api/v2/hosts/search?q=ip%3A8.8.8.0/24
$ curl http://127.0.0.1:55555/api/v2/hosts/8.8.8.8
```

Querying the raw Avro files without database processing

runZero also supports direct queries of the unprocessed Avro files. These queries are slow and may take hours or days to complete depending on the query and local storage speed. To query the raw Avro files, you use the runZero scanner's `censys` command. It takes any number of arguments, which can be:

- Names of Avro files, which must end in `.avro`
- CIDRs or IP addresses to search for in the files

The scanner reads the Avro files specified, and writes a file in runZero scan format containing the appropriate host records. By default, the file has a name matching `censys-*.rubble.gz` and is written to the current directory. Alternatively you can specify an output filename with the `--output-raw` option, as if performing a runZero scan.

The runZero scan file can be uploaded to the runZero console like any other scan file.

If you have more IP addresses or CIDRs than will fit on a command line, you can use the `--input-targets` option to specify that the scanner should read them from a file. The file is expected to be ASCII text, and contain CIDRs or IP addresses separated by whitespace (which can include newlines).

You can also use the scanner to process data, upload it, and then delete the scan data file if everything succeeded. For example:

```
% runzero-scanner censys universal-internet-dataset-20210923-000000000000.avro \
  12.216.190.0/24 --upload --api-key=YOUR_ORGANIZATION_API_KEY \
  --upload-site="Primary site"
```

If you are using self-hosted runZero, you can use the `--api-url` option to specify your console's API endpoint.

The `censys` command also supports the `--verbose` option, which will make it list host addresses as they are written to the output file.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Censys Search integration unable to connect?

1. Are you getting any data from the Censys Search integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.

2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this integration documentation and follow the steps exactly.
3. If the Censys Search integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - o 500 - server error, unable to connect to the endpoint
 - o 404 - hitting an unknown endpoint on the server
 - o 403 - not authorized, likely a credential issue

CrowdStrike Falcon

Community Platform

runZero integrates with CrowdStrike by importing data through the [CrowdStrike Falcon API](#). This integration allows you to sync and enrich your asset inventory, as well as ingesting vulnerability data from Falcon Spotlight and software data from Falcon Discover. Adding your CrowdStrike data to runZero makes it easier to find things like endpoints that are missing an EDR agent.

Getting started

To set up the CrowdStrike integration, you'll need to:

1. Configure CrowdStrike to allow API access through runZero.
2. Add the CrowdStrike credentials, which will include the client ID and client secret, and CrowdStrike base API URL in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the CrowdStrike integration to sync your data with runZero.

Requirements

Before you can set up the CrowdStrike integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the CrowdStrike admin portal.

Step 1: Configure CrowdStrike to allow API access to runZero

1. Log in to CrowdStrike.
2. Go to **Support > API Clients and Keys**. When the **API Key** page appears, choose to add a new API client.
3. Provide the following details for the API client:
 - o **Client name** - API client name, such as runZero.
 - o **API scope** -
 - To ingest host details, include read permissions for **Hosts** and **Host Groups**
 - To ingest vulnerability data, include read permissions for **Vulnerabilities**.
 - To ingest software data, include read permissions for **Assets**.

4. When you are done, add the client. An API client created window appears and shows you the client ID and client secret. You'll need them to configure the integration in runZero.
5. Copy the client ID and client secret now. You may not be able to get them later.

Step 2: Add the CrowdStrike credentials to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like CrowdStrike Falcon.
2. Choose **CrowdStrike Falcon API key** from the list of credential types.
3. Provide the following information:
 - **CrowdStrike client ID and CrowdStrike client secret** - To generate your client ID and client secret, go to **Support > API Clients and Keys > OAuth2 API clients > Add new API Client** in your CrowdStrike portal.
 - **CrowdStrike API URL** - Your organization-specific base URL, which will depend on your account type. You can find this in the CrowdStrike API Swagger documentation.
 - For a **US-1** account use `api.crowdstrike.com`
 - For a **US-2** account use `api.us-2.crowdstrike.com`
 - For a **US-GOV-1** account use `api.laggar.gcw.crowdstrike.com`
 - For a **EU-1** account use `api.eu-1.crowdstrike.com`
4. If you want other organizations to be able to use these credentials, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
5. Save the credentials. You're now ready to set up and activate the connection to bring in data from CrowdStrike.

Step 3: Choose how to configure the CrowdStrike integration

The CrowdStrike integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the CrowdStrike integration to sync data

After you add your CrowdStrike credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the CrowdStrike integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new CrowdStrike-only assets are created.

1. Activate a [connection to CrowdStrike](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure that it has access to the organization you are currently in.
3. Set the severity and risk levels you want to import (optional).
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Add an [FQL filter](#) for imported assets (optional).
 - If access to the CrowdStrike Spotlight API is configured, vulnerability data will only be imported for the assets included in the FQL filter results.
6. Enter a name for the task, like CrowdStrike sync (optional).
7. Choose the Explorer to perform this connector task from (optional).
8. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
9. Enter a description for the task (optional).
10. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
11. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
12. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the CrowdStrike integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the CrowdStrike toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the CrowdStrike toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View CrowdStrike assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your CrowdStrike assets and vulnerabilities. These will have a CrowdStrike icon listed in the **Source** column.

To filter by CrowdStrike attributes, consider running the following queries:

- [View all CrowdStrike assets:](#)

```
source:crowdstrike
```

- [Find assets that have a CrowdStrike EDR agent installed:](#)

```
edr.name:crowdstrike
```

- Find Windows assets, excluding servers, that are missing a CrowdStrike EDR agent:

`os:windows and not type:server and not edr.name:CrowdStrike`

- View all CrowdStrike vulnerabilities:

`source:crowdstrike`

- View all CrowdStrike software results:

`source:crowdstrike`

Click into each asset or vulnerability to see its individual attributes. runZero will show you the attributes returned by the CrowdStrike API, with the exception of policies.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the CrowdStrike integration unable to connect?

1. Are you getting any data from the CrowdStrike integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the integration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the CrowdStrike integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. If the integration endpoint is on-premises, verify they are running the integration task from an Explorer with access to the CrowdStrike host.

How can I solve the following CrowdStrike errors:

- unable to collect software data for CrowdStrike devices: invalid response 403 Forbidden?
1. This error is due to missing CrowdStrike Discover permissions, which means that runZero cannot collect software information. Asset and vulnerability data should be collected just fine, though. This can be remedied by returning to [step 1](#) of the CrowdStrike documentation above and enabling Falcon Discover to collect software data.

CrowdStrike Falcon Filtering

Community Platform

runZero integrates with CrowdStrike by importing data through the [CrowdStrike Falcon API](#). The integration will merge existing assets with Falcon data when the MAC address or hostname matches and create new assets where there is not a match. By default, the integration will import all Falcon hosts. runZero provides a filter option that allows you to scope the import to a subset of hosts based on foreign attributes provided by CrowdStrike. The filter follows the Falcon Query Language (FQL) syntax.

Falcon Query Language

Falcon Query Language (FQL) follows the syntax `<property>: [operator] <value>`. Multiple expressions can be combined for more complex filtering by adding a + between expressions. An OR expression can also be leveraged with comma separated expressions. See [Sample Filters](#) for examples.

Properties

The following are some useful CrowdStrike properties that can be used in an FQL expression to filter assets. Details on additional attributes that are available for filtering can be found by reviewing [CrowdStrike's API documentation](#).

CrowdStrike Property	runZero Attribute	Description	Example
external_ip	externalIP	The external IP address of the device	18.191.169.203
first_seen	firstSeen	A timestamp of when the device was first seen by CrowdStrike	2022-01-08T19:42:34Z
hostname	hostname	The hostname of the device	EXPLORER-01
last_seen	lastSeen	The timestamp of when the device was last seen by CrowdStrike	2022-09-13T19:14:30Z
local_ip	localIP	The local IP address of the device	192.168.1.100
mac_address	macAddress	The mac address of the interface communication with CrowdStrike	0a-6e-20-4a-e6-56
os_version	osVersion	The operation system version of the device	Ubuntu 20.04
platform_name	platformName	The platform running on the device	Linux

product_type_desc	productTypeDesc	The type of device	Server
-------------------	-----------------	--------------------	--------

Operators

The following operators can be used in an FQL expression to filter assets.

Operator	Description
!	Not equal to
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
~	Text match. Tokenizes the string, ignoring spaces, case and punctuation
!~	Does not text match. Tokenized the string, ignoring spaces, cases and punctuation
*	Wildcard matching. Matches one or more characters

Sample Filters

The following are examples of filters that can be applied to the CrowdStrike sync.

Search Filter	Description
hostname: 'WIN10*'	Import all devices where the hostname starts with WIN10
platform_name: 'Linux'	Import all Linux devices
platform_types_desc: 'Server'	Import all devices that CrowdStrike identifies as a Server
hostname: 'PROD*' + platform_name: 'Linux'	Import all Linux devices with a hostname that starts with PROD
local_ip: '192.168.1.100'	Only import the device with a local IP address of 192.168.1.100
local_ip: '!192.168.1.100'	Import all devices, excluding 192.168.1.100
local_ip.raw: '*192.168.1.*'	Import all devices with a local IP address in the 192.168.1.0/24 range
(local_ip.raw: '*192.168.1.*'), (local_ip.raw: '*192.168.2.*')	Import all devices with a local IP address in the 192.168.1.0/24 or 192.168.2.0/24 range

<code>local_ip.raw:!*'192.168.1.*'</code>	Import all devices, excluding devices with a local IP address in the 192.168.1.0/24 range
<code>local_ip.raw:!*'192.168.1.*'+local_ip.raw:!*'192.168.2.*'</code>	Import all devices, excluding devices with a local IP in the 192.168.1.0/24 and 192.168.2.0/24 ranges
<code>last_seen:>'now-24h'</code>	Only import devices that have been seen in the last 24 hours
<code>last_seen:!'<'now-7d'</code>	Exclude devices that have not been seen in the last 7 days

Google Cloud Platform

Professional

Community

Platform

The Google Cloud Platform (GCP) integration provides visibility into your cloud assets by synchronizing your GCP cloud inventories with runZero. runZero also integrates with other cloud providers, such as [Microsoft Azure](#) and [Amazon AWS](#). Similarly to other integrations, you will need to add the [credentials](#) needed to authenticate to GCP and set up a connector in runZero. runZero will pull in GCP compute instance VMs, pulling in GCP attributes that will be viewable from each asset.

The following GCP asset types are supported:

- Compute Engine instances
- Load balancers
- Cloud SQL

Requirements

- Verify you have a [Google Cloud service account](#) with the Compute Network Viewer and Cloud SQL Viewer [roles](#).
 - This service account will need to be granted access to each project that you want the integration to gather data from.
- [Download a key](#) for the GCP service account.
- Verify you have these [GCP APIs enabled](#) on each project:
 - Compute Engine
 - Cloud SQL Admin

How to set up the Google Cloud Platform integration

Here are the high-level steps to set up the Google Cloud Platform integration:

- [Create a Google Cloud Platform credential](#) in runZero.
- Choose whether to configure the integration as [a scan probe or connector task](#).
- [Activate the integration](#) for Google Cloud Platform.
- [View your GCP assets](#).

Step 1: Create Google Cloud Platform credentials

1. Go to the [Credentials page](#) and click **Add Credential**.
2. From the **Credentials type** dropdown, choose **GCP Service Account Key**.
3. Provide a name for the credential, like GCP.
4. Set the **Include all projects** toggle to Yes if you want runZero to gather asset data from all GCP projects that the service account has access to. If set to No, the integration will only gather asset data from the project specified in the key file.
5. Click **Choose file** to upload the service account key file you downloaded from GCP.
6. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
7. Save the credential. You're now ready to set up and activate the connection to bring in data from Google Cloud Platform.

Step 2: Choose how to configure the GCP integration

The GCP integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 3: Activate the Google Cloud Platform integration

After you add your GCP credential, you'll need to set up a connector task or scan probe to sync your data.

Step 3a: Configure the GCP integration as a connector task

1. [Activate a connection to GCP](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like `Google Cloud Platform sync`.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.

5. To organize your assets logically, choose the site you'd like to use to add your assets to. You can choose an existing site or add them to a new site when the sync occurs. Assigning your assets to a site helps organize and group your assets. You can automatically generate a new site per GCP project by selecting this option from the task configuration.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can check the [Scheduled tasks](#) to see when the next sync will occur.

Step 3b: Configure the GCP integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the GCP toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the GCP toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 4: View your Google Cloud Platform assets

After a successful sync, you can [go to your inventory](#) to view your GCP assets. These assets will have a Google icon listed in the **Source** column.

To filter by GCP assets, consider running the following queries:

- [View all GCP assets](#):

```
source:gcp
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by GCP.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Google Cloud Platform integration unable to connect?

1. Are you getting any data from the GCP integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.

2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the GCP integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Google Workspace

Professional

Community

Platform

runZero integrates with Google Workspace to allow you to sync and enrich your asset inventory, as well as gain visibility into users and groups. Adding your Google Workspace data to runZero makes it easier to find unmanaged assets on your network. The Google Workspace integration supports ChromeOS, Mobile, and Endpoint [registered asset types](#).

Requirements

- Verify or create a new [Google service account](#) in whichever project is most suitable.
- [Create and download a key](#) for the Google service account. Save this JSON file.
- Verify that you have the Admin SDK and Cloud Identity APIs enabled for the project. Use the search box in the [API Library](#) to find each API and then enable it.
- Enable domain-wide delegation in the [Google Workspace console](#)
 - Add a new API client using the unique numeric ID of service account as the Client ID
 - Enable the following OAuth scopes for this API client:
`https://www.googleapis.com/auth/admin.directory.user.readonly,https://www.googleapis.com/auth/admin.directory.devices.readonly`
 - Optionally, enter each OAuth scope individually:
 - `https://www.googleapis.com/auth/admin.directory.user.readonly`
 - `https://www.googleapis.com/auth/admin.directory.group.readonly`
 - `https://www.googleapis.com/auth/admin.directory.device.mobile.readonly`
 - `https://www.googleapis.com/auth/admin.directory.device.chromeos.readonly`
 - `https://www.googleapis.com/auth/cloud-identity.devices.readonly`

How to set up the Google Workspace integration

These are the high-level steps to set up the Google Cloud Platform integration:

- [Create a Google Workspace credential](#) in runZero.
- Choose whether to configure the integration as [a scan probe or connector task](#).
- [Activate the connection](#) for Google Workspace.
- [View your results](#).

Step 1: Create a Google Workspace credential

1. Go to the [Credentials page](#) and click **Add Credential**.
2. From the **Credentials type** dropdown, choose **Google Workspace Client Secret**.
3. Provide a name for the credential, like Google Workspace.
4. In the Admin account email field, provide the email address of an administrator account with access to the assets, users, or groups you wish to import.
5. If you want to import from an organization other than the one your administrator account belongs to, provide a Customer ID. By default, runZero will use the Customer ID associated with the service account. (Optional)
6. Click **Choose file** to upload the service account key file you downloaded from Google Workspace.
7. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
8. Save the credential. You're now ready to set up and activate the connection to bring in data from Google Workspace.

Step 2: Choose how to configure the Google Workspace integration

The Google Workspace integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync. Configuring the integration as a scan probe is useful if you are running self-hosted runZero Platform and your console cannot access Google Workspace. For most situations it will be easier to set up a scheduled connection to sync your data from Google Workspace.

Step 3: Activate the Google Workspace integration

After you add your GCP credential, you'll need to set up a connector task or scan probe to sync your data.

Step 3a: Configure the Google Workspace integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where the data is organized.

1. [Activate a connection to Google Workspace](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like Google Workspace sync.

4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. To organize your assets logically, choose the site you'd like to use to add your assets to. You can choose an existing site or add them to a new site when the sync occurs. Assigning your assets to a site helps organize and group your assets.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can check the [Scheduled tasks](#) to see when the next sync will occur.

Step 3b: Configure the Google Workspace integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the GoogleWorkspace toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the GoogleWorkspace toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 4: View Google Workspace assets

After a successful sync, you can [go to your inventory](#) to view your Google Workspace assets. These assets will have a Google Workspace icon listed in the **Source** column.

To filter by Google Workspace assets, consider running the following queries:

- [View all Google Workspace assets:](#)
`source:googleworkspace`
- [View runZero assets not connected to Google Workspace:](#)
`source:runzero AND NOT source:googleworkspace`

Click into each asset to see its individual attributes. runZero will show you the attributes returned by Google Workspace.

Community Platform

The Google Workspace integration provides details about users and groups in addition to enriching asset inventory data. Go to [Inventory > Users](#) or [Inventory > Groups](#) to view the data provided by Google Workspace. Use the query `source:googleworkspace` to filter your results.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Google Workspace integration unable to connect?

1. Are you getting any data from the Google Workspace integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Google Workspace integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Microsoft 365 Defender

Community Platform

runZero integrates with Microsoft 365 Defender to allow you to sync and enrich your asset inventory. Adding your Microsoft 365 Defender data to runZero makes it easier to find assets missing EDR protection.

Getting started

To set up the Microsoft 365 Defender integration, you'll need to:

1. Configure Microsoft 365 Defender to allow API access through runZero.
2. Add the Microsoft 365 Defender credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the Microsoft 365 Defender integration to sync your data with runZero.

Requirements

Before you can set up the Microsoft 365 Defender integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Microsoft Azure portal.

Step 1: Register an Azure application for Microsoft 365 Defender API access

runZero can authenticate to the Microsoft 365 Defender API using a client secret. Register an application to configure Microsoft 365 Defender API access.

1. Log into the Microsoft Azure portal.
2. Go to **Azure Active Directory > App registrations** and click on **New registration**.
 - o Provide a name.
 - o Select the supported account types.
 - o Optionally add a redirect URI.
3. Click **Register** to register the application.
4. Once the application is created, you should see the **Overview** dashboard. Note the following information:
 - o Application (client) ID
 - o Directory (tenant) ID
5. From the application's details page, go to API permissions > Add a permission.
6. Select the second tab called **APIs my organization uses** to view available APIs.
7. Select WindowsDefenderATP from the list of Microsoft APIs.
8. Select the permissions type Application permissions to configure a **client secret**.
9. Search for and select the following required permissions:
 - o Windows Defender ATP API permissions:
 - Machine.Read.All
10. Click **Add permissions** to save the permissions to the application.
11. Navigate to **Azure Active Directory > App registrations** and select the application you created.
12. Go to **Certificates & secrets** and click on **New client secret**.
 - o Enter a description.
 - o Select the expiration.
13. Click **Add** to create the client secret and save the client secret value.

Step 2: Add an Azure Client Secret credential to runZero

This type of credential can be used to sync all resources in a single directory (across multiple subscriptions).

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like Azure Client Secret.
3. Choose **Azure Client Secret** from the list of credential types.
4. Provide the following information:
 - o **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - o **Azure client secret** - To generate a client secret, go to **Azure Active Directory > App registrations**, select your application, go to **Certificates & secrets** and click on **New client secret**.

- **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - Select the **Access all subscriptions in this directory (tenant)** option to sync all resources in your directory. Otherwise, specify the **Azure subscription ID** - The unique ID for the subscription that you want to sync. This can be found in the Azure portal if you go to **Subscriptions** and select the subscription.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
 6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 3: Choose how to configure the Microsoft 365 Defender integration

The Microsoft 365 Defender integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the Microsoft 365 Defender integration to sync data

After you add your Microsoft 365 Defender credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the Microsoft 365 Defender integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Microsoft 365 Defender-only assets are created.

1. Activate a [connection to Microsoft 365 Defender](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure it has access to the organization you are currently in.
3. Enter a name for the task, like **Microsoft 365 Defender sync**.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. If you want to include assets that have stopped reporting to the Microsoft 365 Defender service in your integration import, switch the **Include inactive assets** toggle to Yes. By

default, the integration will not include assets that are marked as inactive.

8. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Microsoft 365 Defender integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Defender365 toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the MS365Defender toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View Microsoft 365 Defender assets

After a successful sync, you can [go to your inventory](#) to view your Microsoft 365 Defender assets. These assets will have an Active Directory icon listed in the **Source** column.

To filter by Microsoft 365 Defender assets, consider running the following queries:

- [View all Microsoft 365 Defender assets:](#)

```
source:ms365defender
```

- [View runZero assets not connected to Microsoft 365 Defender:](#)

```
source:runzero AND NOT source:ms365defender
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by Microsoft 365 Defender.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Microsoft 365 Defender integration unable to connect?

1. Are you getting any data from the Microsoft 365 Defender integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this

documentation and follow the steps exactly.

3. If the Microsoft 365 Defender integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Microsoft Active Directory

Professional

Community

Platform

runZero integrates with Microsoft Active Directory (AD) via LDAP to allow you to sync and enrich your asset inventory, as well as gain visibility into domain users and groups. Adding your AD data to runZero makes it easier to find assets that are not part of your domain.

Getting started

To set up the Active Directory integration, you'll need to:

1. Add the AD credential in runZero.
2. Choose whether to configure the integration as [a scan probe or connector task](#).
3. Activate the Active Directory integration to sync your data with runZero.

Requirements

Before you can set up the LDAP integration, make sure you have credentials for an LDAP account.

Step 1: Add the LDAP credential to runZero

1. Go to the [Add credential page](#) in runZero. Provide a name for the credentials, like LDAP.
2. Choose **LDAP Username & Password** from the list of credential types.
3. Provide the following information:
 - **LDAP username** - The username you want to use with the LDAP integration. The account used for this integration does not require any special permissions. The following username formats are accepted:
 - Distinguished Name (DN): CN=[username],CN=Users,DC=[domain],DC=[tld]
 - User Principle Name (UPN): [username]@[domain].[tld]
 - Domain\Username: [domain]\[username]
 - **LDAP password** - The password for the username to be used with the LDAP integration.
 - **LDAP base DN** - The base distinguished name for LDAP searches. This field requires distinguished name format: DC=[domain],DC=[tld]. Note that only entities underneath the specified base DN will be imported into runZero.
 - **LDAP URL** - The URL for your LDAP server. This field supports IP[:port] notation as well as hostname.domain.tld[:port]. This field requires that the URL entered

- begins with `ldap://` (for insecure LDAP connections) or `ldaps://` (for secure LDAP connections). For example: `ldaps://ad.example.com:636`
- **LDAP insecure** - Set this to Yes if you want to attempt authentication without a verified thumbprint. By default, runZero will attempt to connect with LDAPS but will fall back to LDAP+StartTLS then LDAP. LDAP without StartTLS will only work if this toggle is set to Yes.
 - **LDAP thumbprints** (optional) - A set of `IP[:port]=SHA256:B64HASH` or `hostname.domain.tld=SHA256:B64HASH` pairs to trust for authentication.
 - You will need to scan your LDAP server with runZero in order to obtain the TLS thumbprint. The [TLS fingerprints service attribute report](#) lists all previously seen fingerprints. The TLS thumbprints used for self-signed certificates will only work with LDAPS. If you want to use LDAP+StartTLS with a self-signed certificate, you will need to set the **Insecure** option to Yes.
 - If LDAP insecure is set to No and no thumbprints are provided:
 - With a self-signed certificate, the connection will fail because the certificate chain cannot be verified.
 - With a valid certificate from a public CA, the connection will work without thumbprints.
4. If you want all other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
 5. Save the credential. You're now ready to set up and activate the connection to bring in data from LDAP.

Step 2: Choose how to configure the Active Directory integration

The Active Directory integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 3: Set up and activate the Active Directory integration to sync data

After you add your Active Directory credential, you'll need to set up a connector task or scan probe to sync your data.

Step 3a: Configure the Active Directory integration as a connector task

After you add your LDAP credential, you'll need to set up a connection to sync your data from LDAP. A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new LDAP-only assets are created.

1. Activate a [connection to Active Directory](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Enter a name for the task, like LDAP sync.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 3b: Configure the Active Directory integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the LDAP toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the LDAP toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 4: View Active Directory assets

After a successful sync, you can [go to your inventory](#) to view your LDAP assets. These assets will have an Active Directory icon listed in the **Source** column.

To filter by LDAP assets, consider running the following queries:

- [View all LDAP assets:](#)
source:ldap
- [View runZero assets not connected to LDAP:](#)
source:runzero AND NOT source:ldap

Click into each asset to see its individual attributes. runZero will show you the attributes returned by LDAP.

Community Platform

The LDAP integration provides details about users and groups in addition to enriching asset inventory data. Go to [Inventory > Users](#) or [Inventory > Groups](#) to view the data

provided by LDAP.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Microsoft Active Directory integration unable to connect?

1. Are you getting any data from the Microsoft Active Directory integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Microsoft Active Directory integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Microsoft Active Directory host.

Microsoft Azure

Professional

Community

Platform

runZero integrates with Microsoft Azure to deliver greater visibility into your cloud assets. This integration imports data through each applicable API to enrich your asset inventory:

- [Virtual Machines API](#)
- [Virtual Machine Scale Sets API](#)
- [Load Balancers API](#)
- [AzureSQL API](#)
- [Web Apps API](#)

Syncing with Azure allows you to view information about your asset's OS profile, storage profile, and more. This integration imports assets that are in a running state.

Getting started

The following Azure resource types are supported:

- Virtual Machines
- Virtual Machine Scale Sets

- Azure SQL
- Azure Cosmos DB
- Load Balancers
- Function Apps

To set up the Azure integration, you'll need to:

1. Configure Azure to allow API access through runZero.
2. Add an Azure credential to runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the Azure integration to sync your data with runZero.

Requirements

Before you can set up the Azure integration, make sure you have access to the Microsoft Azure portal.

Step 1: Configure Azure to allow API access through runZero

1. Log into the Microsoft Azure portal.
2. Go to **Azure Active Directory > App registrations** and click on **New registration**.
 - Provide a name.
 - Select the supported account types.
 - Optionally add a redirect URI.
3. Click **Register** to register the application.
4. Once the application is created, you should see the **Overview** dashboard. Save the following information:
 - Application (client) ID
 - Directory (tenant) ID
5. Give the client access to the subscriptions you want to sync. From the subscription details page, go to **Access Control (IAM)** and select **Add > Add role assignment**. Enter the following:
 - Role: Reader
 - Assign access to: User, group, or service principal
 - Under **Select**, search for the name of the application you created. Click on your application to add it to the **Selected members** list below.
6. Click **Save** to save the role assignment.
7. Navigate to **Azure Active Directory > App registrations** and select the application you created.
8. Go to **Certificates & secrets** and click on **New client secret**.
 - Enter a description.
 - Select the expiration.
9. Click **Add** to create the client secret. Save the following information:
 - Client secret value

Step 2: Add the Azure credential to runZero

The credential used for the Azure integration can be either a client secret or a username & password.

Step 2a: Add an Azure Client Secret credential to runZero

This type of credential can be used to sync all resources in a single directory (across multiple subscriptions).

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like Azure Client Secret.
3. Choose **Azure Client Secret** from the list of credential types.
4. Provide the following information:
 - **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure client secret** - To generate a client secret, go to **Azure Active Directory > App registrations**, select your application, go to **Certificates & secrets** and click on **New client secret**.
 - **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - Select the **Access all subscriptions in this directory (tenant)** option to sync all resources in your directory. Otherwise, specify the **Azure subscription ID** - The unique ID for the subscription that you want to sync. This can be found in the Azure portal if you go to **Subscriptions** and select the subscription.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 2b: Add an Azure Username & Password credential to runZero

This type of credential can be used to sync all resources across directories. Alternatively, you can add one **Azure Client Secret** credential for each Azure directory you want to sync.

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like Azure User/Pass.
3. Choose **Azure Username & Password** from the list of credential types.
4. Provide the following information:
 - **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.

- **Azure username** - The username for your Azure cloud account. This cannot be a federated user account.
 - **Azure password** - The password for your Azure cloud account.
5. If you want other organizations to be able to use this credential, select the [Make this a global credential](#) option. Otherwise, you can configure access on a per organization basis.
 6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 3: Choose how to configure the Azure integration

The Azure integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the Azure integration to sync data

After you add your Azure credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the Azure integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Azure-only assets are created.

1. Activate a [connection to Azure](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like `Azure sync`.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
6. Under **Service options**, select the services you would like to sync data from. You must choose at least one.
7. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
8. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Azure integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Azure toggle to Yes, and change any of the default options if needed.

4. On the Credentials tab, set the Azure toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View Azure assets

After a successful sync, you can [go to your inventory](#) to view your Azure assets. These assets will have an Azure icon listed in the **Source** column.

To filter Azure assets, consider running the following queries:

- [View all Azure assets:](#)
`source:azure`
- [View all Azure VMs](#)
`has:"@azure.vm.vmID"`
- [View all Azure virtual machine scale set VMs](#)
`has:"@azure.vmss.vmID"`
- [View all Azure load balancers](#)
`has:"@azure.lb.id"`
- [View all AzureSQL instances](#)
`has:"@azure.azsql.id"`
- [View all Azure Cosmos DB instances](#)
`has:"@azure.cosmos.id"`
- [View all Azure Function Apps](#)
`@azure.functionapp.kind:"functionapp"`

Click into each asset to see its individual attributes. runZero will show you the attributes returned by the Azure APIs.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Microsoft Azure integration unable to connect?

- 1.

1. Are you getting any data from the Microsoft Azure integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Microsoft Azure integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Microsoft Intune

Community Platform

runZero integrates with Microsoft Intune to allow you to sync and enrich your asset inventory. Adding your Microsoft Intune data to runZero makes it easier to find unmanaged assets on your network. Data added includes the [discovered apps](#) from Intune. Managed apps (those pushed to devices by Intune) are not currently reported.

Getting started

To set up the Microsoft Intune integration, you'll need to:

1. Configure Microsoft Intune to allow API access from runZero.
2. Add the Microsoft Intune credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the Microsoft Intune integration to sync your data with runZero.

Requirements

Before you can set up the Microsoft Intune integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Microsoft Azure portal.

Step 1: Register an Azure application for Microsoft Intune API access

runZero can authenticate to the Microsoft Intune API using either a username and password or a client secret. Register an application to configure Microsoft Intune API access.

1. Log into the Microsoft Azure portal.
2. Go to **Azure Active Directory > App registrations** and click on **New registration**

- Provide a name.
 - Select the supported account types.
 - Optionally add a redirect URI.
3. Click **Register** to register the application.
 4. Once the application is created, you should see the **Overview** dashboard. Note the following information:
 - Application (client) ID
 - Directory (tenant) ID
 5. From the application's details page, go to API permissions > Add a permission.
 6. Select Microsoft Graph from the list of Microsoft APIs.
 7. Select the correct permissions type for your needs:
 - **Username & password**: select Delegated permissions
 - **Client secret**: select Application permissions
 8. Search for and select the following required permission:
 - DeviceManagementManagedDevices.Read.All
 - User.Read.All
 9. Click **Add permissions** to save the permissions to the application.
 10. Click **Grant admin consent** to grant consent for the permissions to the application.
 11. If using a client secret, also perform the following steps:
 - Navigate to **Azure Active Directory > App registrations** and select the application you created.
 - Go to **Certificates & secrets** and click on **New client secret**.
 - Enter a description.
 - Select the expiration.
 - Click **Add** to create the client secret and save the client secret value.

Step 2: Add the Microsoft Intune credential to runZero

Adding the Microsoft Intune credential requires adding an Azure username and password to and an Azure Client Secret to runZero. The following sub-steps breaks down each task.

Step 2a: Add an Azure Username & Password credential to runZero

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like Azure User/Pass.
3. Choose **Azure Username & Password** from the list of credential types.
4. Provide the following information:
 - **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - **Azure username** - The username for your Azure cloud account. This cannot be a federated user account.
 - **Azure password** - The password for your Azure cloud account.

5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 2b: Add an Azure Client Secret credential to runZero

This type of credential can be used to sync all resources in a single directory (across multiple subscriptions).

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like **Azure Client Secret**.
3. Choose **Azure Client Secret** from the list of credential types.
4. Provide the following information:
 - o **Azure application (client) ID** - The unique ID for the registered application. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
 - o **Azure client secret** - To generate a client secret, go to **Azure Active Directory > App registrations**, select your application, go to **Certificates & secrets** and click on **New client secret**.
 - o **Azure directory (tenant) ID** - The unique ID for the tenant. This can be found in the Azure portal if you go to **Azure Active Directory > App registrations** and select the application.
5. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per organization basis.
6. Save the credential. You're now ready to set up and activate the connection to bring in data from Azure.

Step 3: Choose how to configure the Microsoft Intune integration

The Microsoft Intune integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the Microsoft Intune integration to sync data

After you add your Microsoft Intune credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the Microsoft Intune integration as a connector task

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Microsoft Intune-only assets are created.

1. Activate a [connection to Microsoft Intune](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure it has access to the organization you are currently in.
3. Enter a name for the task, like Microsoft Intune sync.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Microsoft Intune integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Intune toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the Intune toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View Microsoft Intune assets

After a successful sync, you can [go to your inventory](#) to view your Microsoft Intune assets. These assets will have an Active Directory icon listed in the **Source** column.

To filter by Microsoft Intune assets, consider running the following queries:

- [View all Microsoft Intune assets:](#)
source:intune
- [View runZero assets not connected to Microsoft Intune:](#)
source:runzero AND NOT source:intune

Click into each asset to see its individual attributes. runZero will show you the attributes returned by Microsoft Intune.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Microsoft Intune integration unable to connect?

1. Are you getting any data from the Microsoft Intune integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Microsoft Intune integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

How do I solve the following Microsoft Intune error:

- (invalid_client) AADSTS7000218: The request body must contain the following parameter: 'client_assertion' or 'client_secret'

This error means that you need to enable Public Client Flows in Azure. To do so, follow these steps:

1. Navigate to the App Registration page in the Azure portal
2. Choose **Authentication** from the left navigation
3. Select **Advanced Settings**
4. Toggle the Allow Public Client Flows switch at the bottom of the page to Yes

Miradore MDM

Community Platform

runZero integrates with Miradore mobile device management (MDM) to deliver greater visibility into your mobile assets. This integration imports data from the [Miradore API](#) to enrich your asset inventory. Syncing with Miradore allows you to view information about device hardware, OS version, associated user, and more. This integration imports all enrolled devices.

Getting started

To set up the Miradore integration, you'll need to:

1. Login to your Miradore web portal and create a new API key.
2. Add the Miradore credential to runZero, which includes the endpoint hostname and API key.
3. Choose whether to configure the integration as a [scan probe](#) or [connector task](#).
4. Activate the Miradore integration to sync your data with runZero.

Requirements

Before you can set up the Miradore integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Miradore MDM portal as an administrator.

Step 1: Create a Miradore API key

1. Log into your Miradore portal as an administrator.
2. Click on the **System > Infrastructure diagram** link from the left-side navigation.
3. Scroll down to find the **API** node in the diagram. Click this and choose **Create API key**.
4. Give this new key a name and copy the secret value.

Step 2: Add the Miradore credential to runZero

1. Go to the [Credentials page](#) in runZero and click **Add Credential**.
2. Provide a name for the credential, like Miradore MDM.
3. Choose **Miradore MDM API Key** from the list of credential types.
4. Provide the following information:
 - **Name** - Give this credential a unique name (ex: Miradore)
 - **Miradore endpoint hostname** - The URL for your Miradore portal.
 - **Miradore API key** - The API key created in step 1.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from Miradore.

Step 3: Choose how to configure the Miradore integration

The Miradore integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the Miradore MDM integration to sync data

After you add your Miradore credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the Miradore integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Miradore-only assets are created.

1. Activate a [connection to Miradore](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like `Miradore sync`.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Miradore integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Miradore toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the Miradore toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View Miradore assets

After a successful sync, you can [go to your inventory](#) to view your Miradore assets. These assets will have a Miradore icon listed in the **Source** column.

To filter by Miradore assets, consider running the following queries:

- [View all Miradore assets](#):

```
source:miradore
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by the Miradore API.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Miradore integration unable to connect?

1. Are you getting any data from the Miradore integration?
 - o Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - o In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Miradore integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - o 500 - server error, unable to connect to the endpoint
 - o 404 - hitting an unknown endpoint on the server
 - o 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Miradore host if it is on-premises.

Qualys VMDR

Community Platform

runZero integrates with Qualys VMDR by importing data from the Qualys KnowledgeBase [API](#).

Asset inventory

There is a column on the asset inventory page showing the count of vulnerabilities detected by Qualys for each asset. When a single asset is selected, the vulnerabilities table lists all the results related to that asset. The vulnerability count can be impacted by the type of vulnerability scan as well as the import settings selected.

Vulnerabilities table

The **Vulnerabilities** tab of the inventory lists all vulnerability results that have been imported from Qualys. The table lists every result, and selecting a result will take you to the page for the impacted asset.

Severity and risk scores

Qualys assigns all vulnerabilities a severity rating (Minimal, Medium, Serious, Critical, Urgent). runZero normalizes the severities shown in the vulnerability inventory to be consistent across the runZero Console.

runZero Severity	Qualys Severity
Info	1 / Minimal
Low	2 / Medium
Medium	3 / Serious
High	4 / Critical
Critical	5 / Urgent

runZero will also normalize risk scores assigned by Qualys. A risk score of 0.0 will be shown as none in the runZero Console, and all other risk scores will match the assigned severity level.

Getting started with Qualys

To set up the Qualys VMDR integration, you'll need to:

1. Create or obtain user credentials with access to the Qualys API.
2. Configure [CVSS scoring in Qualys](#).
3. Add the Qualys API username, password, and account API URL in runZero.
4. Choose whether to configure the integration as [a scan probe or connector task](#).
5. Activate the Qualys integration to pull your data into runZero.

Requirements

Before you can set up the Qualys VMDR integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Qualys Cloud Platform portal.

Step 1: Add the Qualys credentials to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like Qualys.
2. Choose **Qualys Username & Password** from the list of credential types.
3. Provide the following information:
 - **Qualys username** - the username you want to use to connect to the Qualys API.
 - **Qualys password** - the password for your Qualys API username.
 - **Qualys account API URL** - the URL of the Qualys API for the relevant account. The expected format is `https://ip:port` or `https://domain.tld:port`. This URL is unique for [each Qualys user](#).

4. If you want other organizations to be able to use this credential, select the `Make this a global credential` option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential. You're now ready to set up and activate the connection to bring in data from Qualys VMDR.

Note: The role of the Qualys API user will determine which VM scanned hosts are visible to the integration. Managers may view all VM scanned hosts in the subscription. Auditors have no permission to view VM scanned hosts. Unit Managers may view VM scanned hosts in the user's assigned business unit. Scanners and Readers may view VM scanned hosts in the user's account, but must be assigned to the assets through asset groups in VM/VMDR.

Step 2: Choose how to configure the Qualys integration

The Qualys integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 3: Set up and activate the Qualys integration to sync data

After you add your Qualys credential, you'll need to sync your data.

Note: The Qualys Cloud Platform enforces limits on the API calls subscription users can make. Both API controls are limited per subscription based on your service level. The tasks generated by this integration may experience slow performance or failures as a result of the enforced API limits. The Qualys API documentation on this topic can be found [here](#).

Step 3a: Configure the Qualys scan probe

You can run the Qualys VMDR integration as a scan probe so that the runZero Explorer will pull your vulnerability data into the runZero Console.

In a new or existing scan configuration:

- Ensure that the `QUALYS` option is set to `Yes` in the `Probes` and `SNMP` tab and change any of the default options if needed.
- Set the correct Qualys credential to `Yes` in the `Credentials` tab.
- Optionally, set the [severity and risk levels](#) for ingested vulnerability scan results.

Step 3b: Configure the Qualys connector

A connection requires you to set a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Qualys-only assets are

created.

1. Activate a [connection to Qualys](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Set the [severity and risk levels](#) you want to import (optional).
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Enter a name for the task, like Qualys sync (optional).
6. Choose the Explorer to perform this connector task from (optional).
7. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
8. Enter a description for the task (optional).
9. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
10. If you want to include assets that have not been assessed for vulnerabilities, switch the **Include unscanned assets** toggle to Yes.
11. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
12. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4: View Qualys assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Qualys assets. These assets will have a Qualys icon listed in the **Source** column.

The Qualys integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to Inventory > [Vulnerabilities](#) to view the vulnerability data provided by Qualys VMDR.

To filter by Qualys assets, consider running the following query:

- [View all Qualys assets](#):

```
source:Qualys
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Qualys VMDR scan data.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Qualys integration unable to connect?

1. Are you getting any data from the Qualys integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the integration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Qualys integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Qualys host if it is on-premises.

Rapid7

Community Platform

runZero integrates with Rapid7's [InsightVM](#) and [Nexpose](#) to enrich your asset inventory and gain visibility into vulnerabilities detected in your environment.

Asset inventory

There is a column on the asset inventory page showing the count of vulnerabilities detected by Rapid7 for each asset. When a single asset is selected, the vulnerabilities table lists all the results related to that asset. The vulnerability count can be impacted by the type of vulnerability scan as well as the import settings selected.

Vulnerabilities table

The **Vulnerabilities** tab of the inventory lists all vulnerability results that have been imported from Rapid7. The table lists every result, and selecting a result will take you to the page for the impacted asset.

Severity and risk scores

Rapid7 assigns all vulnerabilities a severity rating (Moderate, Severe, or Critical) based on the vulnerability's CVSSv2 score. runZero normalizes the severities shown the vulnerability inventory to be consistent across the runZero Console.

runZero Severity	Rapid7 Severity	CVSS Range
Info	Moderate	0.0
Medium	Moderate	0.1 - 3.4
High	Severe	3.5 - 7.4
Critical	Critical	7.5 - 10.0

runZero will also normalize risk scores assigned by Rapid7. A risk score of 0.0 will be shown as none in the runZero Console, and all other risk scores will match the assigned severity level.

InsightVM

Community Platform

runZero integrates with Rapid7 InsightVM by importing data from the [InsightVM API](#).

Both Rapid7 InsightVM Cloud and on-premises InsightVM are supported. For on-premises use you will need to use the InsightVM connector [as a scan probe](#) from a runZero Explorer which has network access to the InsightVM deployment.

The Insight Platform API is distinct from the InsightVM API, and is not supported.

Getting started with InsightVM

To set up the InsightVM integration, you'll need to:

1. Create or obtain user credentials to use with the InsightVM API.
2. Add the InsightVM API username, password, and API URL in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the InsightVM integration to pull your data into runZero.

Requirements

Before you can set up the InsightVM integration:

- Verify that you have [runZero Enterprise](#).
- Obtain credentials for an InsightVM user with administrator access to the InsightVM portal.
- Scan your InsightVM with a runZero Explorer if you want to use trusted authentication (optional).

Step 1: Add the InsightVM credentials to runZero

1. [Create a new credential](#) via the runZero [Credentials page](#).
2. Provide a name for the credentials, like InsightVM.
3. Choose **InsightVM Username & Password** from the list of credential types.
- 4.

4. Provide the following information:
 - **InsightVM username** - The username you want to use to connect to the InsightVM API. This account requires the [User role](#) or greater permissions in InsightVM.
 - **InsightVM password** - The password for your InsightVM username.
 - **InsightVM API URL** - The URL of your InsightVM API instance. By default, the InsightVM API uses port 3780. The expected format is `https://ip:3780` or `https://domain.tld:3780`. (Note: runZero automatically appends `api/3/sites?size=500` to this URL.)
 - **InsightVM insecure** - Set this to Yes if you want to attempt authentication without a verified thumbprint.
 - **InsightVM thumbprints** (optional) - A set of IP=SHA256:B64HASH or domain.tld=SHA256:B64HASH pairs to trust for authentication.
 - You will need to scan your InsightVM instance with runZero in order to obtain the TLS thumbprint. The [TLS fingerprints service attribute report](#) lists all previously seen fingerprints.
 - If InsightVM insecure is set to No and no thumbprints are provided:
 - With a self-signed certificate, the connection will fail because the certificate chain cannot be verified.
 - With a valid certificate from a public CA, the connection can work without thumbprints.
5. If you want all other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
6. Save the credential.

You're now ready to set up and activate the connection to bring in data from InsightVM.

Step 2: Choose how to configure the Rapid7 integration

The Rapid7 InsightVM integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync. Scan probes will be the right option for most users. Setting up a connector will only work for if you're self-hosting runZero or your InsightVM instance is publicly accessible.

Step 3: Sync your InsightVM data

After you add your InsightVM credential, you'll need to activate the integration to sync your data.

Step 3a: Configure the InsightVM scan probe

You can run the InsightVM integration as a scan probe so that the runZero Explorer will pull your vulnerability data into the runZero Console.

In a new or existing scan configuration:

- Ensure that the INSIGHTVM option is set to Yes in the Probes and SNMP tab and change any of the default options if needed.
- Set the correct InsightVM credential to Yes in the Credentials tab.
- Optionally, set the [severity and risk levels](#) for ingested vulnerability scan results.

Step 3b: Configure the InsightVM connector

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new InsightVM-only assets are created.

1. Activate a [connection to InsightVM](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Set the [severity and risk levels](#) you want to import (optional).
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Enter a name for the task, like InsightVM sync (optional).
6. Choose the Explorer to perform this connector task from (optional).
7. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
8. Enter a description for the task (optional).
9. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
10. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
11. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4: View InsightVM assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your InsightVM assets. These assets will have a Rapid7 icon listed in the **Source** column.

The InsightVM integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to Inventory > [Vulnerabilities](#) to view the vulnerability data provided by InsightVM.

To filter by Rapid7 assets, consider running the following queries:

- [View all Rapid7 assets](#):

```
source:Rapid7
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Rapid7 scan data.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Rapid7 InsightVM integration unable to connect?

1. Are you getting any data from the Rapid7 InsightVM integration?
 - o Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - o In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Rapid7 InsightVM integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - o 500 - server error, unable to connect to the endpoint
 - o 404 - hitting an unknown endpoint on the server
 - o 403 - not authorized, likely a credential issue

How can I get a TLS thumbprint for the InsightVM credential?

Here is a set of example commands that would calculate your TLS thumbprint, where is the IP address of your InsightVM instance:

```
$ (echo|openssl s_client -connect <IP address>:3780 -showcerts 2>/dev/null) | openssl x509 -inf
```

An example value returned by this set of commands would look something like this:

X1NWttnkIQprK6zSre/VweKpbRlj7Dt4M6hNfUacytE= Use the following entry for the corresponding TLS thumbprint:

192.168.0.3:3780=SHA256:X1NWttnkIQprK6zSre/VweKpbRlj7Dt4M6hNfUacytE=

Nexpose

Community Platform

runZero integrates with Rapid7 Nexpose by importing files that were exported from your Nexpose instance.

Getting started with Rapid7 Nexpose

To use the Rapid7 Nexpose integration, you'll need to:

1. Download an XML Export or XML Export 2.0 report from Nexpose.
2. Import the Nexpose files through the inventory pages.

Requirements

Before you can set up the Nexpose integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Nexpose portal.

Step 1: Export Nexpose vulnerability scan report

1. Log in to Nexpose with the account being used for the runZero integration.
2. Go to the Reports page and select Create a report.
3. From the Export tab, select either XML Report or XML Report 2.0.
4. Set the scan, asset, asset group, or site scope.
5. Click Save & Run the Report.
6. When the report completes, save the report to a local file.

Step 2: Import the Nexpose files into runZero

1. Go to the [Inventory page](#) in runZero.
2. Choose **Import > Nexpose XML Export (.xml)** from the list of import types.
3. On the import data page:
 - Choose the site you want to add your assets to.
 - Set tags to apply to the imported assets (optional).
 - Set the [severity and risk levels](#) to ingest (optional).
 - Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).

Step 3: View Nexpose assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Nexpose assets. These assets will have a Rapid7 icon listed in the **Source** column.

The Nexpose integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to Inventory > [Vulnerabilities](#) to view the vulnerability data provided by Nexpose.

To filter by Rapid7 assets, consider running the following queries:

- [View all Rapid7 assets](#):

```
source:Rapid7
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Nexpose scan file.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Rapid7 Nexpose integration unable to connect?

1. Are you getting any data from the Rapid7 Nexpose integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Rapid7 Nexpose integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Nexpose host.

SentinelOne

Community Platform

runZero integrates with SentinelOne by importing data from the [SentinelOne API](#). This integration allows you to sync and enrich your asset inventory, as well as gain visibility into the software installed on SentinelOne assets. Adding your SentinelOne data to runZero makes it easier to find things like endpoints that are missing required software.

Any IP address reported by SentinelOne will be treated as a secondary address, not a primary address, since these IPs can be stale and may not be associated with a specific network or site.

Getting started

To set up the SentinelOne integration, you'll need to:

1. Configure SentinelOne to allow API access through runZero.
2. Add the SentinelOne API key and SentinelOne base API URL in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the SentinelOne integration to sync your data with runZero.

Requirements

Before you can set up the SentinelOne integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the SentinelOne admin portal.

Step 1: Configure SentinelOne to allow API access to runZero

1. Log in to SentinelOne with the account being used for the runZero integration.
2. Go to **User > My User**.
3. **Generate** the API token, then download or copy it. This API key expires and will need to be regenerated every six months.

Step 2: Add the SentinelOne credential to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like SentinelOne.
2. Choose **SentinelOne API key** from the list of credential types.
3. Provide the following information:
 - **SentinelOne API URL** - Your organization-specific base URL, which will depend on your account type. It will be something like `organization.sentinelone.net`.
 - **SentinelOne API key** - To generate your API key, go to **User > My User** in your SentinelOne portal. From there, a key can be generated, regenerated, or revoked.
4. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from SentinelOne.

Step 3: Choose how to configure the SentinelOne integration

The SentinelOne integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the SentinelOne integration to sync data

After you add your SentinelOne credential, you'll need to set up a connector task or scan probe to sync your data.

Step 4a: Configure the SentinelOne integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new SentinelOne-only assets are created.

1. Activate a [connection to SentinelOne](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).

2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Enter a name for the task, like SentinelOne sync.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the SentinelOne integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).
2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the SentinelOne toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the SentinelOne toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 5: View SentinelOne assets and software

After a successful sync, you can [go to your inventory](#) to view your SentinelOne assets. These assets will have a SentinelOne icon listed in the **Source** column.

The SentinelOne integration gathers details about installed software in addition to enriching asset inventory data. Go to Inventory > [Software](#) to view the software data provided by SentinelOne.

To filter by SentinelOne assets, consider running the following queries:

- [View all SentinelOne assets:](#)

```
source:SentinelOne
```

- [Find assets that have a SentinelOne agent installed:](#)

```
edr.name:SentinelOne
```

- [Find Windows assets, excluding servers, that are missing a SentinelOne agent:](#)

```
os:windows and not type:server and not edr.name:SentinelOne
```

Click into each asset to see its individual attributes. runZero will show you the attributes returned by the SentinelOne API, with the exception of policies.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the SentinelOne integration unable to connect?

1. Are you getting any data from the SentinelOne integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the SentinelOne integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the SentinelOne host if it is on-premises.

Shodan

Community Platform

runZero integrates with Shodan by importing data from the [Shodan API](#). This integration allows you to sync data about your externally-facing assets and services from Shodan to provide better visibility of your internet footprint and cyber hygiene.

Getting started

To set up the Shodan integration, you'll need to:

1. Add the Shodan API key in runZero.
2. Choose whether to configure the integration as [a scan probe or connector task](#).
3. Activate the Shodan integration to sync your data with runZero.

Requirements

Before you can set up the Shodan integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have a Shodan account with the correct license to meet your needs.

Step 1: Add the Shodan credential to runZero

1. Go to the [new credential page](#) in runZero. Provide a name for the credential, like Shodan.
2. Choose **Shodan Search API key** from the list of credential types.
3. Provide your **Shodan Search API key** - To view your API key, go to your Account page in the Shodan portal. Your API key is available on that page and can be reset if needed.
4. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential. You're now ready to set up and activate the connection to bring in data from Shodan.

Step 2: Choose how to configure the Shodan integration

The Shodan integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 3: Set up and activate the Shodan integration to sync data

After you add your Shodan credential, you'll need to set up a connection or a scan probe to sync your data from Shodan.

Step 3a: Configure the Shodan integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Shodan-only assets are created.

1. Activate a [connection to Shodan](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credential you added earlier. If you don't see the credential listed, make sure the credential has access to the organization you are currently in.
3. Enter a name for the task, like Shodan sync.
4. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start at the date and time you have set.
5. Under **Task configuration**, choose the site you want to add your assets to.
6. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
7. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [tasks](#) page to see when the next sync will occur.

Step 3b: Configure the Shodan integration as a scan probe

1. Create a new scan task or select a future or recurring scan task from your [Tasks page](#).

2. Add or update the scan parameters based on any additional requirements.
3. On the Probes and SNMP tab, choose which additional probes to include, set the Shodan toggle to Yes, and change any of the default options if needed.
4. On the Credentials tab, set the Shodan toggle for the credential you wish to use to Yes.
5. Click **Initialize scan** to save the scan task and have it run immediately or at the scheduled time.

Step 4: View Shodan assets and services

After a successful sync, you can [go to your inventory](#) to view your Shodan assets. These assets will have a Shodan icon listed in the **Source** column.

The Shodan integration gathers details about services in addition to enriching asset inventory data. Go to Inventory > [Services](#) to view the service data provided by Shodan.

To filter by Shodan assets or services, consider running the following queries:

- [View all Shodan assets:](#)

```
source:shodan
```

- [View all Shodan services:](#)

```
source:shodan
```

Click into each asset or service to see its individual attributes. runZero will show you the attributes returned by the Shodan Search API.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Shodan integration unable to connect?

1. Are you getting any data from the Shodan integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Shodan integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Tenable

Community Platform

runZero integrates with Tenable Vulnerability Management (previously Tenable.io), Tenable Nessus, and Tenable Security Center to enrich your asset inventory and gain visibility into vulnerabilities detected in your environment. The [Tenable Vulnerability Management](#), [Nessus Professional](#), and [Tenable Security Center](#) integrations pull data from the Tenable API, while all versions of [Tenable Nessus](#) and Tenable Security Center (previously Tenable.sc) are also supported through Nessus v2 file imports (.nessus).

Note that at this time, only the main [Tenable Vulnerability Management](#) cloud API endpoint at <https://cloud.tenable.com> is supported as an API integration.

Asset inventory

There is a column on the asset inventory page showing the count of vulnerabilities detected by Tenable for each asset. When a single asset is selected, the vulnerabilities table lists all the results related to that asset. The vulnerability count can be impacted by the type of vulnerability scan as well as the import settings selected.

Vulnerabilities table

The **Vulnerabilities** tab of the inventory lists all vulnerability results that have been imported from Tenable. The table lists every result, and selecting a result will take you to the page for the impacted asset.

Severity and risk scores

Tenable uses the third-party Common Vulnerability Scoring System (CVSS) values from the National Vulnerability Database (NVD) to describe severity associated with vulnerabilities. Tenable assigns all vulnerabilities a severity rating (Info, Low, Medium, High, or Critical) based on the vulnerability's static CVSSv2 or CVSSv3 score. By default, Tenable uses CVSSv2, however users can opt to use CVSSv3 instead. runZero displays the CVSSv2-based severity in the Inventory and Asset views.

To supplement the severity ratings, Tenable calculates a dynamic Vulnerability Priority Rating (VPR) for most vulnerabilities. The VPR is a dynamic companion to the data provided by the vulnerability's CVSS score, since Tenable updates the VPR to reflect the current threat landscape. VPR values range from 0.1-10.0, with a higher value representing a higher likelihood of exploit. runZero displays the VPR-based risk in the Inventory and Asset views. For more details about how Tenable handles severity and risk, please refer to their [documentation](#).

Tenable Vulnerability Management

Community Platform

runZero integrates with Tenable Vulnerability Management (previously Tenable.io) by importing data from the Tenable [API](#).

Getting started with Tenable Vulnerability Management

To set up an integration with Tenable Vulnerability Management, you'll need to:

1. Create an Administrator API key in an access group with Can View [permission](#) to Manage Assets.
2. Configure the Tenable Vulnerability Management credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the integration to pull your data into runZero.

Requirements

Before you can set up the Tenable Vulnerability Management integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have administrator access to the Tenable portal.

Step 1: Create an Administrator API key

1. Log in to Tenable Vulnerability Management with the Administrator account being used for the runZero integration.
2. Go to **My Profile > My Account > API Keys**.
3. **Generate** the API token, and then download or copy it.

Step 2: Add the Tenable credential to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like Tenable Vulnerability Management.
2. Choose **Tenable.io Access & Secret** from the list of credential types.
3. Generate your Tenable access and secret keys via your account page in the Tenable portal, and then provide the following information:
 - **Access key** - Your 64-character Tenable access key.
 - **Secret key** - Your 64-character Tenable secret key.
4. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from Tenable Vulnerability Management.

Step 3: Choose how to configure the Tenable integration

The Tenable Vulnerability Management integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync. Setting up a connector will work if you're self-hosting runZero or integrating with Tenable Vulnerability Management.

Step 4: Set up and activate the integration to sync data

After you add your Tenable credential, you'll need to sync your data from Tenable Vulnerability Management.

Step 4a: Configure the Tenable integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Tenable-only assets are created.

1. Activate a connection to [Tenable Vulnerability Management](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Set the [severity and risk levels](#) you want to import (optional). **Note:** Much of the host information provided by Tenable is from Info-level plugins, so if you only import higher levels of severity you may not see much information about assets not scanned by runZero.
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Enter a name for the task, like Tenable Vulnerability Management sync (optional).
6. Choose the Explorer to perform this connector task from (optional).
7. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
8. Enter a description for the task (optional).
9. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
10. If you want to include assets that have not been assessed for vulnerabilities, switch the **Include unscanned assets** toggle to Yes.
11. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
12. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Tenable integration as a scan probe

You can run the Tenable Vulnerability Management integration as a scan probe so that the runZero Explorer will pull your vulnerability data into the runZero Console.

In a new or existing scan configuration:

- Ensure that the TENABLE option is set to Yes in the Probes and SNMP tab and change any of the default options if needed.
- Optionally, set the [severity and risk levels](#) for ingested vulnerability scan results.
- Set the correct Tenable credential to Yes in the Credentials tab.

Step 5: View Tenable assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Tenable assets. These assets will have a Tenable icon listed in the **Source** column.

The Tenable integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to **Inventory > Vulnerabilities** to view the vulnerability data provided by Tenable Vulnerability Management.

To filter by Tenable assets, consider running the following queries:

- [View all Tenable assets:](#)

```
source:Tenable
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Tenable scan data.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Tenable Vulnerability Management integration unable to connect?

1. Are you getting any data from the Tenable Vulnerability Management integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Tenable Vulnerability Management integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Tenable host if it is on-premises.

How do I solve the following error in Tenable Vulnerability Management:

```
"error-message":"no tenable assets match import criteria","level":"error","msg":"could not load"
```

This is an error we have seen intermittently from Tenable. A solution that usually works is to enable the `Include Unscanned Assets` toggle in the Tenable task configuration. This will disable the filters we apply for live assets that were scanned in the last 30 days.

Nessus Professional

Community Platform

runZero integrates with Nessus Professional by importing data from the Tenable [API](#).

Getting started with Nessus Professional

To set up an integration with Nessus Professional, you'll need to:

1. Create an Administrator API key in an access group with `Can View permission` to Manage Assets.
2. Configure the Nessus Professional credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the integration to pull your data into runZero.

Requirements

Before you can set up the Nessus Professional integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have administrator access to the Nessus Professional portal.

Step 1: Create an Administrator API key

1. Log in to Nessus Professional with the Administrator account being used for the runZero integration.
 2. Go to **My Profile > My Account > API Keys**.
 3. **Generate** the API token, and then download or copy it.
- You will either need to configure the Tenable credential to skip TLS verification, or provide the TLS thumbprint when creating the runZero credential.

Step 2: Add the Nessus Professional credential to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like Nessus Professional.
2. Choose **Nessus Professional Access & Secret** from the list of credential types.
- 3.

3. Generate your Tenable access and secret keys via your account page in the Tenable portal, and then provide the following information:
 - **Access key** - Your 64-character Tenable access key.
 - **Secret key** - Your 64-character Tenable secret key.
 - **Nessus API URL** - The API URL for your Nessus Professional instance. The expected format is `https://ip:port` or `https://domain.tld:port`. The default port used by Nessus Professional is 8834.
 - **Nessus insecure** - Set this to Yes if you want to attempt authentication without a verified thumbprint.
 - **Nessus thumbprints** (optional) - A set of IP=SHA256:B64HASH pairs to trust for authentication.
 - You will need to scan your Nessus instance with runZero in order to obtain the TLS thumbprint. The [TLS fingerprints service attribute report](#) lists all previously seen fingerprints.
 - If `Nessus insecure` is set to No and no thumbprints are provided:
 - With a self-signed certificate, the connection will fail because the certificate chain cannot be verified.
 - With a valid certificate from a public CA, the connection can work without thumbprints.
4. If you want other organizations to be able to use this credential, select the `Make this a global credential` option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from Nessus Professional.

Step 3: Choose how to configure the Nessus Professional integration

The Nessus Professional integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync. Setting up a connector will work if you're self-hosting runZero or your Nessus Professional instance is publicly accessible.

Step 4: Set up and activate the integration to sync data

After you add your Tenable credential, you'll need to sync your data from Nessus Professional.

Step 4a: Configure the Nessus Professional integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Tenable-only assets are created.

1. Activate a connection to [Nessus Professional](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.

3. Set the [severity and risk levels](#) you want to import (optional). **Note:** Much of the host information provided by Tenable is from Info-level plugins, so if you only import higher levels of severity you may not see much information about assets not scanned by runZero.
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Enter a name for the task, like Nessus Professional sync (optional).
6. Choose the Explorer to perform this connector task from (optional).
7. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
8. Enter a description for the task (optional).
9. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
10. If you want to include assets that have not been assessed for vulnerabilities, switch the **Include unscanned assets** toggle to Yes.
11. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
12. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Nessus Professional integration as a scan probe

You can run the Nessus Professional integration as a scan probe so that the runZero Explorer will pull your vulnerability data into the runZero Console.

In a new or existing scan configuration:

- Ensure that the NESSUS option is set to Yes in the Probes and SNMP tab and change any of the default options if needed.
- Set the correct Nessus credential to Yes in the Credentials tab.
- Optionally, set the [severity and risk levels](#) for ingested vulnerability scan results.

Step 5: View Tenable assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Tenable assets. These assets will have a Tenable icon listed in the **Source** column.

The Tenable integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to Inventory > [Vulnerabilities](#) to view the vulnerability data provided by Nessus Professional.

To filter by Tenable assets, consider running the following queries:

- [View all Tenable assets:](#)

```
source:Tenable
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Tenable scan data.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Nessus Professional integration unable to connect?

1. Are you getting any data from the Nessus Professional integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the Nessus Professional integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the Tenable host if it is on-premises.

How do I solve the following error in Nessus Professional:

- `"error-message":"no tenable assets match import criteria","level":"error","msg":"could not load scan result data to writer" {#nessuspro-import-error}`

This is an error we have seen intermittently from Tenable. A solution that usually works is to enable the Include Unscanned Assets toggle in the Tenable task configuration. This will disable the filters we apply for live assets that were scanned in the last 30 days.

Tenable Nessus

Community Platform

runZero integrates with Tenable Nessus using two methods. For all versions of Nessus, runZero can import Nessus files (.nessus) that were exported from your Nessus instance. Exports from Tenable Security Center are also supported. For Nessus Professional users, the runZero integration can pull scan data from the [Nessus Professional API](#).

Getting started with Tenable Nessus

To use the Tenable Nessus integration, you'll need to:

1. Export vulnerability scan results as Nessus files.
2. Import the Nessus files through the inventory pages.

Requirements

Before you can set up the Nessus integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have access to the Nessus portal.

Step 1: Export vulnerability scan results

1. Log in to Nessus with the account being used for the runZero integration.
2. Open the scan results you want to be able to import into runZero.
3. Choose **Export** > **Nessus** to download the scan results.

Step 2: Import the Nessus files into runZero

1. Go to the [Inventory page](#) in runZero.
2. Choose **Import** > **Nessus scan (.nessus)** from the list of import types.
3. On the import data page:
 - Choose the site you want to add your assets to, and
 - Set the [severity levels and minimum risk level](#) to ingest. (**Note:** much of the host information provided by Tenable is from Info-level plugins, so if you only import higher levels of severity you may not see much information about assets not scanned by runZero.)

Step 3: View Nessus assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Nessus assets. These assets will have a Tenable icon listed in the **Source** column.

The Nessus integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to Inventory > [Vulnerabilities](#) to view the vulnerability data provided by Nessus.

To filter by Nessus assets, consider running the following queries:

- [View all Nessus assets:](#)

```
source:Tenable
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from the Nessus scan file.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Tenable Nessus integration unable to connect?

- 1.

1. Are you getting any data from the Tenable Nessus integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.

Tenable Security Center

Community Platform

runZero integrates with Tenable Security Center (previously Tenable.sc) by importing data from the Tenable Security Center [API](#).

Getting started with Tenable Security Center

To set up an integration with Tenable Security Center, you'll need to:

1. Create an API key for a user that has access to view and query vulnerabilities in Tenable Security Center.
2. Configure the Tenable Security Center credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the integration to pull your data into runZero.

Requirements

Before you can set up the Tenable Security Center integration:

- Verify that you have [runZero Enterprise](#).
- Make sure you have administrator access to the Tenable Security Center portal.

Step 1: Create an API key

1. Log in to Tenable Security Center with an Administrator account.
2. Make sure [API key authentication is enabled](#)
3. Go to **Users > Users**.
4. Check the box for the user you want to create an API key for. **Note:** The API key will have the same access as the user you select. Make sure the user has access to view and query vulnerabilities in the desired organization.
5. At the top of the table, click the **API Keys > Generate API Key** option.
6. Click **Generate** to create the API token, and then download or copy it.

Step 2: Add the Tenable Security Center credential to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like Tenable Security Center.
2. Choose **Tenable Security Center Access & Secret** from the list of credential types.

3. Generate your Tenable Security Center API key as directed in [Step 1](#), and then provide the following information:
 - **Access key** - Your 64-character Tenable Security Center access key.
 - **Secret key** - Your 64-character Tenable Security Center secret key.
4. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from Tenable Security Center.

Step 3: Choose how to configure the Tenable Security Center integration

The Tenable Security Center integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync. If you are integrating with an internal Tenable Security Center instance, we recommend setting up a connector to run from one of your explorers. Otherwise, if you are integrating with an external-facing Tenable Security Center instance, you can set up a connector to run from the cloud. If you are self-hosting runZero, you can run the connector from an explorer or from your runZero host, whichever can reach your Tenable Security Center install.

Step 4: Set up and activate the integration to sync data

After you add your credential, you'll need to sync your data from Tenable Security Center.

Step 4a: Configure the Tenable Security Center integration as a connector task

A connection requires you to specify a schedule which determines when the sync occurs.

1. Activate a connection to [Tenable Security Center](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Configure the Tenable Security Center query mode setting (optional).
 - Select **Define filters** to define a filter based on [vulnerability severity and risk level](#). **Note:** Much of the host information provided by Tenable is from Info-level plugins, so if you only import higher levels of severity you may not see much information about assets.
 - Select **Use existing query ID** to provide the Tenable Security Center query to use. **Note:** The query must be the Vulnerability type and use the Vulnerability Detail List tool.
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).
5. Enter a name for the task, like **Tenable Security Center sync** (optional).
6. Choose the Explorer to perform this connector task from (optional).

7. Choose the site you want to configure the connector for.
8. Enter a description for the task (optional).
9. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
10. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Tenable Security Center integration as a scan probe

You can run the Tenable Security Center integration as a scan probe so that the runZero Explorer will pull your vulnerability data into the runZero Console.

In a new or existing scan configuration:

- Ensure that the TENABLESECURITYCENTER option is set to Yes in the Probes and SNMP tab and change any of the default options if needed.
- Optionally, set the [severity and risk levels](#) for ingested vulnerability scan results or provide a query ID.
- Set the correct TenableSecurityCenter credential to Yes in the Credentials tab.

Step 5: View Tenable Security Center assets and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Tenable Security Center assets. These assets will have a Tenable icon listed in the **Source** column.

The Tenable Security Center integration gathers details about vulnerabilities detected in addition to enriching asset inventory data. Go to **Inventory > Vulnerabilities** to view the vulnerability data provided by Tenable Security Center.

To filter by Tenable Security Center assets, consider running the following queries:

- [View all Tenable Security Center assets:](#)
source:tenablesecuritycenter
- [View all Tenable Security Center vulnerabilities:](#)
source:tenablesecuritycenter

Click into each asset or vulnerability to see its individual attributes. runZero will show you the attributes gathered from the Tenable Security Center API.

VMware

Community Platform

runZero Platform supports synchronization of VMware vCenter and ESXi virtual machine inventories.

Setting up VMware credentials

Unlike other APIs, the VMware synchronization process is configured as part of your regular runZero Explorer scans. The first step is to set up a set of VMware credentials.

On the [Credentials](#) page, click *Add Credential* and choose a credential type of *VMware vCenter/ESXi Username and Password*, and enter the appropriate username and password. The correct username syntax in most cases is `user@domain.com`. The VMware account used requires at least **read-only** access.

The CIDR allow list field can be used to limit which addresses the credentials should be sent to. This helps ensure that they are not passed to unexpected VMware systems that the runZero Explorer encounters on the network.

If runZero has previously found VMware API endpoints, the optional *VMware thumbprints* field will list their IP addresses and TLS fingerprints. You can edit this list to remove any systems you do not want to trust with your VMware credentials. Alternatively, if you do not want to limit authentication to the set list of IP addresses and TLS certificates, you can set *VMware insecure* to Yes to approve authenticating with untrusted endpoints.

The organization access for the credentials can be set as for any other stored credentials; see [Credentials](#).

Performing VMware synchronization

Once you have defined a set of VMware credentials, the second step is to enable VMware synchronization as part of a scan task. Any task which includes scanning the VMware host systems can be used to synchronize VMware VM data.

The Probes tab of the scan setup has a section for enabling and disabling the VMware probe. The probe must be enabled for VMware synchronization to work; it is enabled by default.

On the Credentials tab of the scan setup, use the toggle switch to enable the appropriate set of VMware vCenter/ESXi credentials.

When the scan runs, the Explorer will use the credentials to authenticate with any VMware ESXi or vCenter hosts it finds that the credentials are configured to trust. Data about assets which are VMware VMs will be imported into runZero automatically, and merged with the other information runZero finds by scanning.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the VMware integration unable to connect?

1. Are you getting any data from the VMware integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
3. If the VMware integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue
4. Verify you are running the integration task from an Explorer with access to the VMware host if it is on-premises.

Wiz

Community Platform

runZero integrates with Wiz by importing data from the [Wiz API](#). This integration allows you to sync data about your cloud assets, software, and vulnerabilities from Wiz to provide better visibility of your cloud assets and security posture.

Getting started with Wiz

To set up an integration with Wiz, you'll need to:

1. Create a Service Account in Wiz with permissions to read graph resources, read reports, and create reports.
2. Configure the Wiz credential in runZero.
3. Choose whether to configure the integration as [a scan probe or connector task](#).
4. Activate the integration to pull your data into runZero.

Requirements

Before you can set up the Wiz integration:

- Make sure you have administrator access to the Wiz portal.

Step 1: Create a Service Account in Wiz

1. Log in to Wiz with an Administrator account.
2. Go to **Settings > Access Management > Service Accounts > Add Service Account**.

3. Enter a descriptive name in the **Name** field.
4. Select Custom Integration (GraphQL API) for the **Type**.
5. Leave the **Projects** section blank so the Service Account has access to all projects.
6. Enable `read:resources`, `read:reports`, and `create:reports` for the **API Scopes**.
7. Click **Add Service Account** and copy the **Client ID** and **Client Secret**.
8. Go to **User Settings > Tenant** and note the **API Endpoint URL** in the format:
`https://api.{{region}}.app.wiz.io/`.

Step 2: Add the Wiz credential to runZero

1. Go to the [Credentials page](#) in runZero. Provide a name for the credentials, like Wiz.
2. Choose **Wiz Client Secret** from the list of credential types.
3. Create your Wiz service account via the settings page in the Wiz portal, and then provide the following information:
 - o **Wiz Client ID** - The client ID of your Wiz service account.
 - o **Wiz Client Secret** - The client secret of your Wiz service account.
 - o **Wiz Auth URL**- The URL used to authenticate the Wiz service account.
 - o **Wiz API URL** - The API Endpoint URL used to access the Wiz API.
4. If you want other organizations to be able to use this credential, select the **Make this a global credential** option. Otherwise, you can configure access on a per-organization basis.
5. Save the credential.

You're now ready to set up and activate the connection to bring in data from Wiz.

Step 3: Choose how to configure the Wiz integration

The Wiz integration can be configured as either a [scan probe](#) or a [connector task](#). Scan probes gather data from integrations during scan tasks. Connector tasks run independently from either the cloud or one of your Explorers, only performing the integration sync.

Step 4: Set up and activate the integration to sync data

After you add your Wiz credential, you'll need to sync your data from Wiz.

Step 4a: Configure the Wiz integration as a connector task

A connection requires you to specify a schedule and choose a site. The schedule determines when the sync occurs, and the site determines where any new Wiz-only assets are created.

1. Activate a connection to [Wiz](#). You can access all available third-party connections from the [integrations page](#), your [inventory](#), or the [tasks page](#).
2. Choose the credentials you added earlier. If you don't see the credentials listed, make sure the credentials have access to the organization you are currently in.
3. Set the severity and risk levels you want to import (optional).
4. Set the **Fingerprint only** toggle to Yes if you want vulnerability records to be ingested for fingerprint analysis but not stored in your runZero vulnerability inventory (optional).

5. Enter a name for the task, like Wiz Sync (optional).
6. Choose the Explorer to perform this connector task from (optional).
7. Choose the site you want to add your assets to. All newly discovered assets will be stored in this site.
8. Enter a description for the task (optional).
9. If you want to exclude assets that have not been scanned by runZero from your integration import, switch the **Exclude unknown assets** toggle to Yes. By default, the integration will include assets that have not been scanned by runZero.
10. If you want to exclude assets that have not been assessed for vulnerabilities, switch the **Include unscanned assets** toggle to No.
11. Schedule the sync. A sync can be set to run on a recurring schedule or run once. The schedule will start on the date and time you have set.
12. Activate the connection when you are done. The sync will run on the defined schedule. You can always check the [Scheduled tasks](#) to see when the next sync will occur.

Step 4b: Configure the Wiz integration as a scan probe

You can run the Wiz integration as a scan probe so that the runZero Explorer will pull your Wiz assets into the runZero Console.

In a new or existing scan configuration:

- Ensure that the WIZ option is set to Yes in the Probes and SNMP tab and change any of the default options if needed.
- Optionally, set the severity and risk levels for ingested vulnerability results.
- Set the correct Wiz credential to Yes in the Credentials tab.

Step 5: View Wiz assets, software, and vulnerabilities

After a successful sync, you can [go to your inventory](#) to view your Wiz assets. These assets will have a Wiz icon listed in the **Source** column.

The Wiz integration gathers details about software and vulnerabilities detected in addition to enriching asset inventory data. Go to **Inventory > Software** or **Inventory > Vulnerabilities** to view the software and vulnerability data provided by Wiz.

To filter by Wiz assets, consider running the following queries:

- [View all Wiz assets](#):

```
source:Wiz
```

Click into each asset to see its individual attributes. runZero will show you the attributes gathered from Wiz.

Troubleshooting

If you are having trouble using this integration, the questions and answers below may assist in your troubleshooting.

Why is the Wiz integration unable to connect?

1. Are you getting any data from the Wiz integration?
 - Make sure to query the inventory rather than look at the task details to review all the data available from this integration.
 - In some cases, integrations have a configuration set that limits the amount of data that comes into the runZero console.
2. Some integrations require very specific actions that are easy to overlook. If a step is missed when setting up the intergration, it may not work correctly. Please review this documentation and follow the steps exactly.
 - Double-check the API Scopes assigned to the Wiz service account. A valid service account that is missing the required permissions will result in a failed import.
3. If the Wiz integration is unable to connect be sure to check the task log for errors. Some common errors include:
 - 500 - server error, unable to connect to the endpoint
 - 404 - hitting an unknown endpoint on the server
 - 403 - not authorized, likely a credential issue

Outbound integrations

Using runZero data to enrich other tools

In addition to being able to enrich your runZero inventory with data from your [other IT and security tools](#), the runZero platform offers [egress integrations with several platforms](#). By leveraging product APIs and export/import functionality, runZero can provide additional asset context in other IT and security tools.

The following integrations are available to send your runZero data into other platforms:

IT service management

- [Atlassian Insight & Jira Service Management](#)
- [ServiceNow CMDB](#)

Detection and investigation

- [Panther](#)
- [Splunk Search](#)
- [Sumo Logic](#)
- [Tines](#)
- [Thinkst Canary](#)

Vulnerabilities and risk

- [SecurityGate.io](#)

Atlassian Insight & Jira Service Management

All runZero editions integrate with Jira Service Management via an import in Atlassian Insight. runZero asset data is then imported into the CMDB.

Follow these steps to perform a basic import.

Step 1: Export runZero asset data

You can export data using the **Export** button from the runZero inventory or the Export API.

The following are sample commands for the export API that include common export fields but omit the tags field. You must replace the token ETxxx... with your account's export token from the [Inventory export API page](#).

Create import configuration

9. Scroll down and click **Next**.

10. Choose the file you have previously exported from runZero.

Create import configuration

11. Click **Save Import Configuration**.

12. Click **Create predefined Insight structure** and wait for the task to finish.

13. Click **Create predefined configuration** and wait for the task to finish.

Create Insight structure or Import configuration?

Insight can help you create the insight structure or import configuration to quickly get started importing.

Create predefined structure



You can create a template structure in the current object schema to get started in no time. Insight will try to create the structure depending on the import module.

Object Type*
None

The object type where the structure will be created

Create predefined Insight structure

Create predefined configuration



You can create a template configuration to get started in no time. Insight will look into the object schema to see if any object types are suitable for this import module.

Object Type*
None

The object type where the configuration will be created

Create predefined configuration

Make sure you have a valid Insight structure (predefined or created manually) before creating the predefined configuration.

Close

14. Close the window. You will see a new entry with your configured import on the page. Click the settings cog and select **Execute Import**.



15. Click **Import**.

Confirm importing

i **Confirm that you want to execute import**

Confirm that you want to execute import with **Import CSV** configuration which will affect **1** object types. Any already imported objects will be ignored.

Import Cancel

16. Your import was successful. Close the window.

Finished process in action

✓ Successfully imported objects.

Import took **20** second(s).

Executed as user **CK** Chris Kirsch

Execution type **MANUAL**

Assets (#50)	
Entries in external source	532

[Close](#)

If you are seeing the following error, one of your fields is longer than 255 characters. Manually trim the field to 255 characters or delete the field to successfully import your assets.

Finished process in action



Error occurred during import.

Import took 1 second(s).

Error messages:


Error(s) for object type: Assets

- Error: Text attribute (id: 802, name: macs) value must be less than 255 characters long. [Object (Label: 461193c4-2690-466a-8c20-3e30c4fb5777, Id: Unknown)].
- Error: Text attribute (id: 802, name: macs) value must be less than 255 characters long. [Object (Label: d7a87e93-0bf3-43b1-a97d-096b7e5aa235, Id: Unknown)].
- Error: Text attribute (id: 802, name: macs) value must be less than 255 characters long; Text attribute (id: 798, name: addresses_extra) value must be less than 255 characters long. [Object (Label: 32432fa8-46d8-42bc-b858-ecef2df2ed54, Id: Unknown)].
- Error: Text attribute (id: 802, name: macs) value must be less than 255 characters long. [Object (Label: b5c64020-4e68-4c8d-a25d-c8db28b81472, Id: Unknown)].
- Error: Text attribute (id: 804, name: names) value must be less than 255 characters long. [Object (Label: 369d2852-b7b3-4125-8ffa-89851cb8b1e4, Id: Unknown)].
- Error: Text attribute (id: 802, name: macs) value must be less than 255 characters long; Text attribute (id: 798, name: addresses_extra) value

[Close](#)

Step 3: View your imported assets

1. In Jira Service Management, click on the Insights.
2. Click on your object schema.
3. Click on assets to view details.



Panther

runZero data can be imported into your [Panther](#) instance for enhanced logging and alerting.

Requirements

- A Panther account with the required permissions,
- An AWS S3 bucket, and
- Exported .jsonl files from runZero that have been uploaded into your AWS S3 bucket.

Step 1: Adding a custom schema

1. Go to **Configure > Schemas** and select **Create New**.
2. Add a name.
3. Upload a sample log to automatically parse the runZero output schema.

Step 2: Adding a custom log source

1. Go to **Configure > Log Sources** and select **Create New**.
2. Complete the **Basic Information** section.
3. Opt to configure S3 prefixes and schemas now and select the custom schema you created.
4. Configure the IAM role:
 - Opt to configure **Using the AWS Console UI**.
 - Click **Launch Console UI**.
 - Review the stack in AWS, then check the box to approve, and click to deploy the stack.
 - When the deployment completes, navigate to the **Resources** tab and select the **LogProcessingRole** that was created.
 - Copy the ARN from that role into the field on the Panther console.
5. Configure an alarm if logs are not processed (optional).

Once completed, any .jsonl files added to the specified AWS S3 bucket will be automatically ingested and processed by Panther.

SecurityGate.io

All runZero editions integrate with SecurityGate.io to enrich asset visibility in support of your risk assessment program. Setting up the integration requires a few steps in your SecurityGate.io console.

Requirements

- Configuring the SecurityGate.io integration requires a [runZero API key](#).

The SecurityGate.io integration will pull runZero asset data from across all organizations.

Integrate runZero with SecurityGate.io

1. Log in to your SecurityGate.io console.
2. Go to **My Account** under the **Hello** dropdown menu.
3. Click on **Integration Manager**.
4. Select **Add New Integration**.
5. Choose *Rumble* from the **Integration Partner** dropdown menu.
6. Provide the **Account API Key**, **Server URL**, and **API Version**.
7. Click **Test Connection**. If the test is successful, click **OK** to save the configuration.

Viewing runZero data in SecurityGate.io

After the integration is enabled within SecurityGate.io, runZero data will be available through the Asset Inventory page.

ServiceNow Service Graph

Platform

The Service Graph connector for runZero allows you to bring runZero assets into your ServiceNow CMDB as CIs, and optionally periodically update the CIs with fresh information from runZero scans.

The Service Graph Connector fetches and transforms data using ServiceNow IntegrationHub ETL, and passes it through the Identification and Reconciliation Engine (IRE). This allows specific fields and CI class mappings to be fine-tuned from the ServiceNow console. You can also specify a runZero search query to determine which assets get brought in by the connector.

Important notes

- There is no charge from runZero for use of the connector. However, CI resources created in CMDB by the connector will increase ServiceNow Subscription Unit consumption. CIs created by Service Graph Connectors should be charged at a lower rate than CIs created via other means; you will need to consult ServiceNow to confirm this and obtain current pricing.
- If you are a self-hosted runZero customer, you will need to set up a MID Server to enable ServiceNow to connect to your runZero console. If your console uses a self-signed TLS certificate, or a TLS certificate signed by your own internal CA, you will need to configure the MID server to accept the CA or certificate as trusted. For more information, see [ServiceNow Knowledge Base article KB0863673](#).

- While the Service Graph Connector is packaged and available from the ServiceNow Application Store, setting up connectors to integrate with your existing ServiceNow configuration is unlikely to be a one-click operation. You are strongly advised to have a ServiceNow Consulting and Implementation partner available to assist with the process.

Prerequisites

To use the Service Graph connector for runZero, you need the following:

- A Platform license for runZero.
- A ServiceNow ITOM license including ITOM Discovery and ITOM Visibility.
- Integration Commons for CMDB installed.

The ServiceNow dependencies are checked as part of the connector installation process.

Installation

The first step of the installation is to locate the Service Graph connector for runZero in the ServiceNow Application Store, and follow the usual process to install it into your ServiceNow instance.

The Service Graph connector adds a new menu entry **Service Graph connector for runZero** to the ServiceNow main menu. The module contains entries for **Setup, Data Sources, Scheduled Imports, Support, and System Import Sets**. You must have the **admin** role in ServiceNow to configure the connector.

Setup

The **Setup** menu offers a **Guided Setup** process to help you get up and running with the connector. There are three main stages to the setup process:

1. Configure the connection
2. Configure the mappings (optional)
3. Configure the scheduled import (optional)

Step 1: Configure the connection

Configuring the connection consists of two steps: entering a runZero API key as credentials, and configuring connection options.

Entering the API key

The connector can use a runZero **Organization API** or **Export API** key. In either case, you can obtain the key by going to the **Organizations** page in runZero and clicking on the organization containing the assets you want to bring into ServiceNow CMDB. Organization tokens begin with OT, export tokens begin with ET.

In the **API Key Credentials** form, paste the runZero token into the field labeled **API Key**, and click **Update**.

Configuring connection options

For runZero cloud users, the default **HTTP connection** options should be correct.

You can optionally add a **Search** string to the **Attributes** for the connection at the bottom of the form. If you do, this search will be passed to runZero to determine the set of assets to bring in to CMDB via the connector. Any search string should be in runZero [search query format](#). You are strongly encouraged to test your search string in your [runZero asset inventory](#).

The **Connection** alias and **Base path** fields should be left as-is. (The connection alias groups together the connection settings and the API key credentials for the connector to use.)

If you are self-hosting runZero, the **connection options** form is where you configure the MID server to use and specify the URL of your runZero console. Your console's fully-qualified hostname should be placed in the **Host** field. The hostname should match the hostname in the console's TLS certificate in order for the MID server to trust it, even if the TLS certificate is trusted by the MID server. Check the **Use MID server** box and select the MID server to use.

Step 2: Configure the mappings (optional)

The second step of setting up the connector allows you to customize the mapping of runZero data to CI classes, using the IntegrationHub ETL Transform Map Assistant.

This stage can be skipped if you do not have special requirements for data mapping.

Step 3: Configure the Scheduled Import (optional)

The third step of setting up the connector is to set up a schedule for data import. Initially, you will want to test the connector setup with a one-off import. Once everything is configured correctly, you can create recurring imports.

Scheduled imports are configured as regular ServiceNow tasks. This step can be skipped if you don't want to set up a scheduled import at this time.

Notes on data mapping

The object data models for ServiceNow CMDB and runZero are not an exact match. Some runZero data does not fit in any standard attributes available in CMDB, and in other cases runZero does not have data that CMDB expects. Review this list to understand how the connector will map runZero attributes to ServiceNow CMDB attributes:

- IP addresses: ServiceNow CIs have an attribute which takes a single IP address. The connector places the first IP address in the CI attribute IP Address. The remaining addresses are created as IP Address CIs owned by the device CI. Each IP Address CI is given a Description indicating whether it has been scanned by runZero or not. **Note:**

By default, ServiceNow does not show IP addresses owned by a device CI. It only shows IP addresses owned by a network adapter CI which is owned by the device CI. This is a known limitation of ServiceNow's default CMDB forms. runZero doesn't associate IP addresses with network adapters because the necessary information isn't generally known. In that situation, the CMDB connector development guide states that the IP addresses should be associated with the device CI.

- **Names:** ServiceNow CIs have a single Name attribute, while runZero assets can have any number of names with runZero attempting to guess which are the best names and put them at the front of the list. The connector places the first name in the CI Name attribute. The first name which looks like a fully-qualified domain name is placed in the Fully qualified domain name attribute, and is used to compute a DNS Domain for CIs descended from Network Gear. The full list of names is placed in the Description attribute for reference.
- **Serial numbers:** ServiceNow CIs have an attribute which takes a single serial number. The connector picks the first serial number for the Serial Number attribute, and then also performs a Serial Number Lookup operation to add the complete list of serial numbers with their associated type.
- **Tags:** runZero assets can have any number of tags, which are either single tags such as tag_name, or tags with values in the format tag=value. These are placed in the Key Value class, and this class is then associated with the Hardware class.
- **MAC addresses/Network Adapters:** As with IP addresses, ServiceNow has both a single-value MAC Address attribute, and an option to look up named MAC addresses to create Network Adapter objects. A runZero asset can have any number of MAC addresses, but since runZero does not run agent software on the scanned systems, it has no way to know the interface name associated with a given MAC address. The first MAC address is imported as the MAC Address attribute, and the entire set is imported via lookup as Network Adapter objects, with each being given a name made up by the MAC address and IP address (e.g. 00:00:00:00:00:00-127.0.0.1) or the MAC address and hostname if the IP address is not available (e.g. 00:00:00:00:00:00-hostname). If neither the IP address or hostname are available, the MAC address alone is used as the network adapter name. In some cases, runZero may not know the MAC address(es) of an asset. In this case, a zero MAC address (00:00:00:00:00:00) is added to satisfy ServiceNow requirements for some CI classes.
- **Sites:** runZero sites represent distinct networks, which may or may not correspond to physical sites. They are imported as CMDB Network Site objects.
- **Organizations:** The runZero organization has no obvious corresponding attribute in network hardware CIs. The organization name is incorporated into the Description of CIs.
- **SNMP data:** The snmp.sysDesc attribute from runZero, if it exists, is also placed in the Description.

Splunk Search

Community Platform

runZero integrates with Splunk using a dedicated [Splunk Addon](#), compatible with Splunk 7, Splunk 8, and Splunk Cloud. With this add-on, you'll be able to pull new or updated hosts into a Splunk index, where you'll be able to analyze, visualize, and monitor them there.

This add-on uses the [Splunk API](#) from the runZero Network Discovery platform. It supports syncing assets into Splunk, with multiple inputs supported, global API key management, and optional search filters for each input. For example, you can track new assets as one input, and SMBv1 enabled assets as another input.

To set up this add-on, you'll need an Export API or Organization API key, which you can generate from your Organization page in the runZero Console.

Get the runZero add-on for Splunk

1. Log in to Splunk.
2. Go to **Find More Apps**.
3. Search for runZero Network Discovery.
4. Install the add-on for runZero.
5. Splunk will prompt you to log in again. After you log back in again, the add-on will be installed. You'll be able to open the runZero Asset Sync app. Splunk might also prompt you to restart your server.

Asset sync modes

Two asset sync modes are available: New Assets Only and All Updated Assets. You can export asset inventory that contains newly discovered assets or updated assets, since the last poll, in a sync-friendly format for Splunk. You can leverage the same capabilities from the Asset Sync API to pull data in Splunk, such as search filters, fields, and time-based checkpoints.

Once data is pulled into Splunk, you can create Splunk inputs with filters. This allows you to sync specific assets with a certain protocol, discovery date, or open service.

Sumo Logic

All runZero editions integrate with Sumo Logic to enrich asset visibility and help you visualize your asset data. Setting up the integration requires a few steps in your Sumo Logic console. The integration can be set up to support two distinct purposes:

- [Complete asset visibility](#)
- [Targeted alerting and visualization](#)

Requirements

- A Sumo Logic account
- A runZero account and [API key](#).

Sumo Logic asset export

runZero integrates with Sumo Logic to make your asset inventory available directly in Sumo Logic. This article will show you how to export your runZero inventory into Sumo Logic for use within the SIEM.

Integrating runZero with Sumo Logic

Setting up the connection between Sumo Logic and runZero has three options with different configuration steps.

Option A: Local script

1. Create a Sumo Logic HTTP Source.
2. Configure your host to run the provided script.

Option B: AWS Lambda function

1. Create a Sumo Logic HTTP Source.
2. Configure the AWS Lambda function to run the provided script.

Option C: Sumo Logic script source

1. Install a Sumo Logic collector.
2. Create a Sumo Logic script source.

Once your data is flowing into Sumo Logic, you can [start using](#) the data in Sumo Logic.

Option A: Local script

Step 1: Create a Sumo Logic HTTP Source

1. After logging in to Sumo Logic, navigate to **Manage Data > Collection**.
2. Click **Add Collector** then select **Hosted Collector**.
 - Provide a name, such as runZero Collector and click **Save**.
3. If prompted to add a data source, click **OK**. Otherwise, find your Collector in the list and click **Add Source**.
4. Select the **HTTP Logs and Metrics** source.
 - Provide a name, such as runZero Assets, then click **Save**.
5. Copy the URL provided to use in [step 2](#).

Step 2: Configure your host to run the provided script

1. Identify the host you would like to run the script from.
2. Ensure the host has Python3 and Pipenv installed.
3. Save the script below to the host it will be run from.

```
#!/usr/bin/env python3
import json
import requests
import os

# RUNZERO CONF
RUNZERO_EXPORT_TOKEN = os.environ['RUNZERO_EXPORT_TOKEN']
HEADERS = {'Authorization': f'Bearer {RUNZERO_EXPORT_TOKEN}'}
BASE_URL = 'https://console.runZero.com/api/v1.0'

# SUMO LOGIC CONF
HTTP_ENDPOINT = os.environ['SUMO_HTTP_ENDPOINT']

def main():
    url = BASE_URL + '/export/org/assets.json'
    assets = requests.get(url, headers=HEADERS)
    batchsize = 500
    for i in range(0, len(assets.json()), batchsize):
        batch = assets.json()[i:i+batchsize]
        f = open('upload.txt', 'w')
        f.truncate(0)
        for a in batch:
            json.dump(a, f)
            f.write('\n')
        f.close()
        r = open('upload.txt')
        requests.post(HTTP_ENDPOINT, data=r.read())
        r.close()

if __name__ == '__main__':
    main()
```

4. Create your environment variables by running the following commands:

- `export RUNZERO_EXPORT_TOKEN=XXX`: Use your runZero export API token, which can be obtained in your runZero console on an [organization detail page](#). Select the organization you wish to export data from, then click **Edit organization** to view the export API token.
- `export SUMO_HTTP_ENDPOINT=XXX`: Use the Sumo Logic token obtained in [step 1](#).

5. Create your virtual environment to run the script by running `pipenv --python /path/to/python3`.

6.

Install the requests library in your virtual environment for making API calls:

- pipenv shell
- pip install requests

7. Test the script by running your script from the virtual environment.

- Use the location from the pipenv output to start.
- Append /bin/python3 to use Python in the virtual environment.
- Use the full path to the script.

```
my-server:~/ $ /home/user/.local/share/virtualenvs/runZero-scripts-mVQtFLD0/bin/python3 /h
```

8. Configure a crontab task to run at the desired cadence.

- On the hour: `0 * * * * RUNZERO_EXPORT_TOKEN=XXX SUMO_HTTP_ENDPOINT=XXX /path/to/virtual/env/python3 /path/to/script.py`
- Daily at midnight: `0 0 * * * RUNZERO_EXPORT_TOKEN=XXX SUMO_HTTP_ENDPOINT=XXX /path/to/virtual/env/python3 /path/to/script.py`
- Weekly at midnight on Monday: `0 0 * * 1 RUNZERO_EXPORT_TOKEN=XXX SUMO_HTTP_ENDPOINT=XXX /path/to/virtual/env/python3 /path/to/script.py`

Option B: AWS Lambda function

Step 1: Create a Sumo Logic HTTP Source

1. After logging in to Sumo Logic, go to **Manage Data > Collection**.
2. Click **Add Collector** then select **Hosted Collector**.
 - Provide a name, such as runZero Collector and click **Save**.
3. If prompted to add a data source, click **OK**. Otherwise, find your Collector in the list and click **Add Source**.
4. Select the **HTTP Logs and Metrics** source.
 - Provide a name, such as runZero Assets, then click **Save**.
5. Copy the URL provided to use in [step 2](#).

Step 2: Configuring the AWS Lambda function to run the provided script

1. Go to your AWS Console and navigate to the [Lambda page](#).
2. Click **Create a function**.
3. Give your function a **name**.
4. Select Python 3.9 as the **runtime**.
5. Everything else can be left with the default setting. Click **Create function** to move to the next page.
6. Click **Add Trigger** to set up a cron job.

7. Select EventBridge to set up a schedule.
8. Use an existing rule or select **Create new rule**.
 - Give it a name and set **Rule type** to Schedule expression.
 - Use one of these options or create your own based on desired cadence:
 - Daily: rate(1 day)
 - Every 12 hours: rate(12 hours)
 - Every 3 hours: rate(3 hours)
 - Click **Add** to return to the main Lambda configuration page.
9. Under **Configuration** select Environment variables.
10. Enter these two environment variables:
 - RUNZERO_EXPORT_TOKEN which can be obtained in your runZero console on an [organization detail page](#). Select the organization you wish to export data from, then click **Edit organization** to view the export API token.
 - SUMO_HTTP_ENDPOINT which was obtained in [step 1](#).
11. Click Save to return to the main Lambda configuration page.
12. Click the Code tab and replace the default code with this script.

```
import json
import urllib3
import os

# RUNZERO CONF
RUNZERO_EXPORT_TOKEN = os.environ['RUNZERO_EXPORT_TOKEN']
HEADERS = {'Authorization': f'Bearer {RUNZERO_EXPORT_TOKEN}'}
BASE_URL = 'https://console.runZero.com/api/v1.0'

# SUMO LOGIC CONF
HTTP_ENDPOINT = os.environ['SUMO_HTTP_ENDPOINT']

def lambda_handler(event, context):
    http = urllib3.PoolManager()
    url = BASE_URL + '/export/org/assets.json'
    response = http.request('GET', url, headers=HEADERS)
    data = response.data
    assets = json.loads(data)
    batchsize = 500
    for i in range(0, len(assets), batchsize):
        batch = assets[i:i+batchsize]
```

```
f.truncate(0)
    for a in batch:
        json.dump(a, f)
        f.write('\n')
    f.close()
    r = open('upload.txt')
    http.request('POST', HTTP_ENDPOINT, body=r.read())
    r.close()
```

13. Click **Deploy** to update the code.

14. Click **Test** to verify the code works.

Your asset data export will now be posted to Sumo Logic at the cadence you configured.

Option C: Sumo Logic script source

Step 1: Installing a Sumo Logic collector

Follow the [Sumo Logic documentation](#) in order to install a collector.

Step 2: Creating a Sumo Logic script source

Sumo Logic has documentation on [script sources](#) as well. Here are the steps to follow to set up the script source once your collector is installed.

1. Navigate to the [Collection](#) page in Sumo Logic.
2. Find your collector and click **Add > Add Source**.
3. Select Script as the **source type**.
4. Input a Name and Source Category.
5. Select a Frequency.
6. Select **Command** type /usr/bin/python.
7. Add the following script in the **Script** field.

```
#!/usr/bin/python
import json
import requests
import os

# RUNZERO CONF
RUNZERO_EXPORT_TOKEN = os.environ['RUNZERO_EXPORT_TOKEN']
HEADERS = {'Authorization': 'Bearer ' + RUNZERO_EXPORT_TOKEN}
BASE_URL = 'https://console.runZero.com/api/v1.0'
```

```
def main():
    url = BASE_URL + '/export/org/assets.json'
    assets = requests.get(url, headers=HEADERS)
    for a in assets.json():
        print(json.dumps(a))

if __name__ == '__main__':
    main()
```

8. Click **Save** to allow the source to start working.

Working with the asset data in Sumo Logic

Once your asset data is in Sumo Logic, you can use it in any way you would use any other log source. Here are some sample searches that you could use to create scheduled searches and dashboards.

Search distinct assets

```
_sourceCategory="runzero"
| json field=_raw "id"
| count_distinct(id) as distinct_assets
```

Search assets with more than 3 services running

```
_sourceCategory="runzero"
| json field=_raw "addresses_extra"
| json field=_raw "addresses"
| json field=_raw "id"
| concat("https://console.runzero.com/inventory/", id) as runzero_link
| json field=_raw "service_count"
| where service_count > 3
| count addresses, addresses_extra, service_count, runzero_link
```

Determine counts of different operating systems

```
_sourceCategory="runzero"
| json field=_raw "os"
| where !isEmpty(os)
| json field=_raw "id"
| count os, id
| count os
```

Sumo Logic alerting

runZero integrates with Sumo Logic to help you visualize your asset data. This helps you track your progress on reducing risk in your asset inventory over time.

Setting up the connection between Sumo Logic and runZero requires:

1. [Creating a Sumo Logic HTTP Source](#)
2. [Creating a runZero alert template](#)
3. [Creating a rule in runZero](#)
4. [Handling runZero data in Sumo Logic](#)
5. [Creating a Sumo Logic dashboard \(optional\)](#)

Step 1: Create a Sumo Logic HTTP Source

1. After logging in to Sumo Logic, navigate to **Manage Data > Collection**.
2. Click **Add Collector** select **Hosted Collector**, provide a name, such as runZero Collector and click **save**.
3. If prompted to add a data source, click **OK**. Otherwise, find your Collector in the list and click **Add Source**.
4. Select the **HTTP Logs and Metrics** source, provide a name, such as runZero Alerts, and then click **save**.
5. Copy the URL provided to use in [step 2](#).

Step 2: Create a runZero alert template

1. Create an [alert template](#) in runZero and provide the following details:
 - o **Name:** Name for template
 - o **Template type:** JSON
 - o **Subject line for message:** Leave empty
 - o **Body of message:** The following JSON example will include the rule name and the search URL in the alert message body

```
{"rule_name":"{{rule.name}}","search_url":"{{search.url}}","found": "{{search.found}}
```

2. Create an [alert channel](#) in runZero and provide the following details:
 - o **Name:** Name for alert channel
 - o **Channel type:** Webhook
 - o **Webhook URL:** The webhook URL you copied from Sumo Logic

Step 3: Create a rule in runZero

Now that you have your alert template and channel created, you will want to identify the triggers to alert on. Some common examples are:

- Asset query results - When there is a match on a query in runZero after a scan completes
- New assets found - When a scan completes with new assets

- Agent offline - When your runZero Explorer stops checking in to the console
- Task failed - When a task fails for any reason.

We will use the **asset query results** selection as an example for the rest of the steps. Review [example queries](#) for ideas on what queries you could create.

1. Create a [new alert rule](#).
2. Select asset-query-results and click **Configure rule**.
3. Input values for the rule:
 - **Name:** name of the rule.
 - **Conditions:** optional parameters that will trigger the alert when all conditions match.
 - **Query:** The query the assets must match.
 - **Number of matches:** The numeric comparison logic for the **value**.
 - **Value:** The threshold of matches to trigger the rule.
 - **Limit to organization:** Allows you to limit the alert to a specific organization.
 - **Limit to site:** Allows you to limit the alert to a specific site.
 - **Action:** Notify
 - **Notification channel:** Name of the alert channel you created in step 2.
 - **Notification template:** Name of the alert template you created in step 2.

Step 4: Handle runZero data in Sumo Logic

This search will display the raw runZero data

```
_source="runZero Alerts" and _collector="runZero"
```

This search will show alerts matching the runZero rule name

```
_source="runZero Alerts" and _collector="runZero"
| json field=_raw "found" nodrop
| json field=_raw "rule_name" nodrop
| where rule_name = "<RULE NAME>"
```

This search will create a graph of the data matching the runZero rule name

```
_source="runZero Alerts" and _collector="runZero"
| json field=_raw "assets_new" nodrop
| toLong(assets_new)
| json field=_raw "found" nodrop
| json field=_raw "search_url" nodrop
| json field=_raw "rule_name" nodrop
| where rule_name = "<RULE NAME>"
| timeslice 1m
| sum(found) by _timeslice
| order by _timeslice
```

Step 5: Create a dashboard in Sumo Logic (optional)

Now that you know how to look at the data in Sumo Logic and make a graph, you can follow these steps to create a dashboard. You will first create four rules in runZero, then you will import the [Sample Sumo Logic Dashboard](#) below.

Create the runZero rules

You will follow the same actions from step 3 to create each of these rules using the form inputs provided.

Assets running a TLS service

- **Rule type:** asset-query-results
- **Name:** Assets running a TLS service - Sumo
- **Query:** alive:t protocol:tls
- **Number of matches:** is greater than 0
- **Notification channel:** alert channel created in [Step 2](#)
- **Notification template:** alert template created in [Step 2](#)

Multihomed assets - Sumo

- **Rule type:** asset-query-results
- **Name:** Multihomed assets - Sumo
- **Query:** alive:t AND multi_home:t
- **Number of matches:** is greater than 0
- **Notification channel:** alert channel created in [Step 2](#)
- **Notification template:** alert template created in [Step 2](#)

Assets with OpenSSL - Sumo

- **Rule type:** asset-query-results
- **Name:** Assets with OpenSSL - Sumo
- **Query:** alive:t product:openssl
- **Number of matches:** is greater than 0
- **Notification channel:** alert channel created in [Step 2](#)
- **Notification template:** alert template created in [Step 2](#)

New assets to Sumo Logic

- **Rule type:** new-assets-found
- **Name:** New assets to Sumo Logic
- **Number of matches:** is greater than 0
- **Notification channel:** alert channel created in [Step 2](#)
- **Notification template:** alert template created in [Step 2](#)

Sample Sumo Logic Dashboard

1. Navigate to the library in Sumo Logic.
2. Click the **options** button on the folder you'd like to import to and click Import.
 - **Name:** runZero Alert Metrics
 - **JSON:** Copy from the sample below
3. Click Import to see this dashboard under the folder it was imported to.

```
{
  "type": "DashboardV2SyncDefinition",
  "name": "runZero Asset Metrics",
  "description": "",
  "title": "runZero Asset Metrics",
  "theme": "Dark",
  "topologyLabelMap": {
    "data": {}
  },
  "refreshInterval": 0,
  "timeRange": {
    "type": "BeginBoundedTimeRange",
    "from": {
      "type": "RelativeTimeRangeBoundary",
      "relativeTime": "-3d"
    },
    "to": null
  },
  "layout": {
    "layoutType": "Grid",
    "layoutStructures": [
      {
        "key": "panelPANE-AC8FB3DCBD32DA48",
        "structure": "{\"height\":6,\"width\":12,\"x\":0,\"y\":0}"
      },
      {
        "key": "panel3D084A3284252A4E",
        "structure": "{\"height\":6,\"width\":12,\"x\":12,\"y\":0}"
      },
      {
        "key": "panelPANE-4389DBF794B13B44",
        "structure": "{\"height\":6,\"width\":12,\"x\":0,\"y\":6}"
      },
      {
        "key": "panelPANE-FBE08549B2123A4A",
        "structure": "{\"height\":6,\"width\":12,\"x\":12,\"y\":6}"
      }
    ]
  }
}
```

```

]
  },
  "panels": [
    {
      "id": null,
      "key": "panelPANE-AC8FB3DCBD32DA48",
      "title": "New assets found",
      "visualSettings": "{ \"general\": { \"mode\": \"timeSeries\", \"type\": \"line\", \"displayType\": \"default\", \"markerSize\": 5, \"lineDashType\": \"solid\", \"markerType\": \"none\", \"lineThickness\": 1 }, \"title\": { \"fontSize\": 14 }, \"axes\": { \"axisX\": { \"titleFontSize\": 12, \"labelFontSize\": 12 }, \"axisY\": { \"titleFontSize\": 12, \"labelFontSize\": 12, \"logarithmic\": false } }, \"legend\": { \"enabled\": true, \"verticalAlign\": \"bottom\", \"fontSize\": 12, \"maxHeight\": 50, \"showAsTable\": false, \"wrap\": true }, \"color\": { \"family\": \"Categorical Default\" }, \"series\": { }, \"override\": [] }",
      "keepVisualSettingsConsistentWithParent": true,
      "panelType": "SumoSearchPanel",
      "queries": [
        {
          "transient": false,
          "queryString": "_source=\"runZero Alerts\" and _collector=\"runZero\"\n| json field=_raw \"assets_new\" nodrop\n| json field=_raw \"found\" nodrop\n| json field=_raw \"search_url\" nodrop\n| json field=_raw \"rule_name\" nodrop\n| where rule_name = \"New Assets to Sumo Logic\"\n| timeslice 1m\n| sum(assets_new) by _timeslice\n| order by _timeslice",
          "queryType": "Logs",
          "queryKey": "A",
          "metricsQueryMode": null,
          "metricsQueryData": null,
          "tracesQueryData": null,
          "spansQueryData": null,
          "parseMode": "Auto",
          "timeSource": "Message",
          "outputCardinalityLimit": 1000
        }
      ],
      "description": "",
      "timeRange": null,
      "coloringRules": null,
      "linkedDashboards": []
    },
    {
      "id": null,
      "key": "panel3D084A3284252A4E",

```

```

"visualSettings": "{ \"general\": { \"mode\": \"timeSeries\", \"type\": \"line\", \"displayType\": \"default\", \"markerSize\": 5, \"lineDashType\": \"solid\", \"markerType\": \"none\", \"lineThickness\": 1}, \"title\": { \"fontSize\": 14}, \"axes\": { \"axisX\": { \"titleFontSize\": 12, \"labelFontSize\": 12}, \"axisY\": { \"titleFontSize\": 12, \"labelFontSize\": 12, \"logarithmic\": false}}, \"legend\": { \"enabled\": true, \"verticalAlign\": \"bottom\", \"fontSize\": 12, \"maxHeight\": 50, \"showAsTable\": false, \"wrap\": true}, \"color\": { \"family\": \"Categorical Default\"}, \"series\": {}, \"overrides\": []}",
    "keepVisualSettingsConsistentWithParent": true,
    "panelType": "SumoSearchPanel",
    "queries": [
      {
        "transient": false,
        "queryString": "_source=\\runZero Alerts\\ and _collector=\\runZero\\ \\n| json field=_raw \\assets_new\\ nodrop\\n| toLong(assets_new)\\n| json field=_raw \\found\\ nodrop\\n| json field=_raw \\search_url\\ nodrop\\n| json field=_raw \\rule_name\\ nodrop\\n| where rule_name = \\Multihomed Assets\\ or rule_name = \\Multihomed Assets - Sumo\\ \\n| timeslice 1m\\n| sum(found) by _timeslice\\n| order by _timeslice",
        "queryType": "Logs",
        "queryKey": "A",
        "metricsQueryMode": null,
        "metricsQueryData": null,
        "tracesQueryData": null,
        "spansQueryData": null,
        "parseMode": "Auto",
        "timeSource": "Message",
        "outputCardinalityLimit": 1000
      }
    ],
    "description": "",
    "timeRange": null,
    "coloringRules": null,
    "linkedDashboards": []
  },
  {
    "id": null,
    "key": "panelPANE-4389DBF794B13B44",
    "title": "Assets with a TLS service",
    "visualSettings": "{ \"general\": { \"mode\": \"timeSeries\", \"type\": \"line\", \"displayType\": \"default\", \"markerSize\": 5, \"lineDashType\": \"solid\", \"markerType\": \"none\", \"lineThickness\": 1}, \"title\": { \"fontSize\": 14}, \"axes\": { \"axisX\": { \"titleFontSize\": 12, \"labelFontSize\": 12}, \"axisY\": { \"titleFontSize\": 12, \"labelFontSize\": 12, \"logarithmic\": false}}, \"legend\": { \"enabled\": true, \"vert

```

```

ap\":true},\"color\":{\"family\": \"Categorical Default\"},\"series\": {},\"override
s\": []}\",
    \"keepVisualSettingsConsistentWithParent\": true,
    \"panelType\": \"SumoSearchPanel\",
    \"queries\": [
      {
        \"transient\": false,
        \"queryString\": \"_source=\\\"runZero Alerts\\\" and _collector=\\\"ru
nZero\\\"\\n| json field=_raw \\\"assets_new\\\" nodrop\\n| toLong(assets_new)\\n| json fie
ld=_raw \\\"found\\\" nodrop\\n| json field=_raw \\\"search_url\\\" nodrop\\n| json field=_r
aw \\\"rule_name\\\" nodrop\\n| where rule_name = \\\"Assets running a TLS service\\\" or r
ule_name = \\\"Assets running a TLS service - Sumo\\\"\\n| timeslice 1m\\n| sum(found) b
y _timeslice\\n| order by _timeslice\",
        \"queryType\": \"Logs\",
        \"queryKey\": \"A\",
        \"metricsQueryMode\": null,
        \"metricsQueryData\": null,
        \"tracesQueryData\": null,
        \"spansQueryData\": null,
        \"parseMode\": \"Auto\",
        \"timeSource\": \"Message\",
        \"outputCardinalityLimit\": 1000
      }
    ],
    \"description\": \"\",
    \"timeRange\": null,
    \"coloringRules\": null,
    \"linkedDashboards\": []
  },
  {
    \"id\": null,
    \"key\": \"panelPANE-FBE08549B2123A4A\",
    \"title\": \"Assets running OpenSSL\",
    \"visualSettings\": \"{\\\"general\\\":{\\\"mode\\\":\\\"timeSeries\\\",\\\"type\\\":\\\"li
ne\\\",\\\"displayType\\\":\\\"default\\\",\\\"markerSize\\\":5,\\\"lineDashType\\\":\\\"solid\\\",\\\"mar
kerType\\\":\\\"none\\\",\\\"lineThickness\\\":1},\\\"title\\\":{\\\"fontSize\\\":14},\\\"axes\\\":{\\\"ax
isX\\\":{\\\"titleFontSize\\\":12,\\\"labelFontSize\\\":12},\\\"axisY\\\":{\\\"titleFontSize\\\":1
2,\\\"labelFontSize\\\":12,\\\"logarithmic\\\":false}},\\\"legend\\\":{\\\"enabled\\\":true,\\\"vert
icalAlign\\\":\\\"bottom\\\",\\\"fontSize\\\":12,\\\"maxHeight\\\":50,\\\"showAsTable\\\":false,\\\"wr
ap\\\":true},\\\"color\\\":{\\\"family\\\": \"Categorical Default\"},\\\"series\\\": {},\"override
s\": []}\",
    \"keepVisualSettingsConsistentWithParent\": true,
    \"panelType\": \"SumoSearchPanel\",
    \"queries\": [

```

```

"transient": false,
  "queryString": "_source=\"runZero Alerts\" and _collector=\"runZero\"\\n| json field=_raw \"assets_new\" nodrop\\n| toLong(assets_new)\\n| json field=_raw \"found\" nodrop\\n| json field=_raw \"search_url\" nodrop\\n| json field=_raw \"rule_name\" nodrop\\n| where rule_name = \"Assets with OpenSSL\" or rule_name = \"Assets with OpenSSL - Sumo\"\\n| timeslice 1m\\n| sum(found) by _timeslice\\n| order by _timeslice",
  "queryType": "Logs",
  "queryKey": "A",
  "metricsQueryMode": null,
  "metricsQueryData": null,
  "tracesQueryData": null,
  "spansQueryData": null,
  "parseMode": "Auto",
  "timeSource": "Message",
  "outputCardinalityLimit": 1000
},
  "description": "",
  "timeRange": null,
  "coloringRules": null,
  "linkedDashboards": []
},
"variables": [],
"coloringRules": []
}

```

Tines

runZero integrates with Tines to help you automate workflows related to your asset data. This helps teams leverage runZero to the fullest while optimizing the team's workflows with automation. A [video demo](#) is available to show the final outcome of these instructions.

Requirements

- A Tines account
- runZero Export API and Organization API tokens

There are two ways to integrate runZero and Tines:

- Follow the steps to create a [custom story in Tines](#), or
- Use the [runZero sample story](#) to begin with a story outline.

Tines custom story

A Tines story is a collection of actions that work together towards a specific goal, like a playbook. Tines has a Story Library that contains ready-made automated playbooks, or you can create a your own custom story if they don't have one that matches your needs. That's what we'll need to do here.

Step 1: Creating a Tines story and adding runZero API credentials

1. After logging in to Tines, create a new story.
2. Use the + in the **credentials** section to add a **text credential**.
 - **Name:** specify a name for the credential, for example `runzero_export_token`.
 - **Value:** your [runZero export API token](#), which can be obtained from the desired [runZero organization](#) page. Export API tokens start with ET.
3. Use the + in the **credentials** section to add another **text credential**.
 - **Name:** specify a name for the credential, for example `runzero_org_token`.
 - **Value:** your [runZero organization API token](#), which can be obtained from the desired [runZero organization](#) page. Organization API tokens start with OT.

Step 2: Creating a Tines webhook action

1. Add a **Webhook** action to your story.
2. Click the **Webhook** action and copy the webhook URL, which will look like `https://<tenant-name>.tines.com/webhook/<guid>`.

Step 3: Creating a runZero alert template

1. Create an **alert template** in runZero:
 - **Name:** name for template
 - **Template type:** JSON
 - **Subject line for message:** leave empty
 - **Body of message:** the following JSON example will include the rule name and the search URL in the alert message body: `{"rule_name": "{{rule.name}}", "search_url": "{{search.url}}"}"`
2. Create an **alert channel** in runZero:
 - **Name:** name for alert channel
 - **Channel type:** webhook
 - **Webhook URL:** the webhook URL you copied from Tines

Step 4: Creating a rule in runZero

Now that you have your template ready to go, you will want to identify which triggers to alert on. Some common examples are:

- **Asset query results** - When there is a match on a query in runZero after a scan completes
- **New assets found** - When a scan completes with new assets

- **Agent offline** - When your runZero Explorer stops checking in to the console
- **Task failed** - When a task fails for any reason

We will use the **asset query results** selection for the rest of the steps. Our example will be any asset that has an open Telnet port port:23. You can see more example queries [here](#).

1. Create an [alert rule](#).
2. Select asset-query-results and click Configure rule.
3. Input values for the rule:
 - **Name:** name of the rule.
 - **Conditions:** optional parameters that will trigger the alert when all conditions match.
 - **Query:** query the assets must match, such as port:23 for our example.
 - **Number of matches:** the numeric comparison logic for the **value**.
 - **Value:** the threshold of matches to trigger the rule.
 - **Limit to organization:** allows you to limit the alert to a specific organization in runZero
 - **Limit to site:** allows you to limit the alert to a specific site in runZero
 - **Action:** notify
 - **Notification channel:** name of the alert channel you created in step 2.
 - **Notification template:** name of the alert template you created in step 2.

Step 5: Handling the data in Tines

1. The [Webhook action](#) created in [Step 2](#) will be the entry point for the runZero alert in Tines.
2. Add an [Event Transform](#) to parse the search from the URL provided in the runZero alert. If you connect the **Webhook** action to this one, some parameters will automatically populate.
 - Mode: extract
 - Matchers:
 - Path: path_to_alert_search, for example receive_alerts_from_runzero.body.search_url if your **Webhook** action is named "RECEIVE alerts from runZero"
 - Regex: \?.*
 - Extract to: search
3. Add an [Event Transform](#) that takes the output of the previous step. If you connect the previous **Event Transform** action with this one, some parameters will automatically populate.
 - Mode: explode
 - Matchers:
 - Path: path_to_extraction_transform, for example get_search_from_url.search if your previous **Event Transform** action is named "GET search from URL"
 - To: individual_item
4. Add an [HTTP Request](#) to make an API call to the runZero Export API.
 - URL: the path_to_explosion_transform added to the end of the runZero API URL endpoint, for example

- `https://console.runZero.com/api/v1.0/export/org/assets.json<<get_search_from_regex.i`
 if your previous **Event Transform** action is named “GET search from regex”
- Content Type: JSON
 - Method: get
 - Use the **+ Option** button to add **Headers**:
 - Change header to Authorization
 - Change value to Bearer CREDENTIAL.runzero_export_token
5. Add an **Event Transform** that takes the runZero Export API output and loops through each value.
- Mode: explode
 - Path: path_to_HTTP_request, for example get_assets_from_runzero.body if your **HTTP Request** action is named “GET assets from runZero”
 - To: individual_item

Step 6: Personalized automation in Tines

Now that you have received the alert, parsed the search, and obtained the assets from the runZero Export API, it's time to add your own flare. While you will likely want to do something more crafty, the sample story provided includes these two actions as examples.

1. **Send an email** to the destination of your choice for each asset. The sample JSON shows how you might customize it to put the context you'd like in the Subject and Body of the email.
2. The **HTTP Request** action shows how you might reach back into runZero to add tags to the assets after other data is gathered.

runZero sample story

The [sample story below](#) can be imported to Tines to do all of the actions outlined in the steps above. Simply save a JSON file with the contents below, and use the **Import** button in Tines to upload it. After importing, you will need to complete the following steps.

Step 1: Update credentials in Tines

1. Replace the runzero_export_token credential:
 - **Name**: specify a name for the credential, such as runzero_export_token
 - **Type**: text
 - **Value**: your [runZero export API token](#), which can be obtained from the desired [runZero organization](#) page. Export API tokens start with ET.
2. Replace the runzero_org_token credential:
 - **Name**: specify a name for the credential, such as runzero_org_token
 - **Type**: text
 - **Value**: your [runZero organization API token](#), which can be obtained from the desired [runZero organization](#) page. Organization API tokens start with OT.

Step 2: Create runZero alert template

1. Create an [alert template](#)

in runZero:

- **Name:** name for template
 - **Template type:** JSON
 - **Subject line for message:** leave empty
 - **Body of message:** the following JSON example will include the rule name and the search URL in the alert message body: `{"rule_name": "{{rule.name}}", "search_url": "{{search.url}}"}"`
2. Create an [alert channel](#) in runZero:
 - **Name:** name for alert channel
 - **Channel type:** webhook
 - **Webhook URL:** the webhook URL you copied from Tines

Step 3: Create the rule in runZero

Now that you have your template ready to go, you will want to identify which triggers to alert on. Some common examples are:

- Asset query results: when there is a match on a query in runZero after a scan completes.
- New assets found: when a scan completes with new assets.
- Agent offline: when your runZero Explorer stops checking in to the console.
- Task failed: when a task fails for any reason.

We will use the **asset query results** selection for the rest of the steps. Our example will be any asset that has an open Telnet port `port:23`. You can see more example queries [here](#).

1. Create a [new alert rule](#).
2. Select `asset-query-results` and click `Configure rule`.
3. Input values for the rule:
 - **Name:** name of the rule.
 - **Conditions:** optional parameters that will trigger the alert when all conditions match.
 - **Query:** query the assets must match, such as `port:23` for our example.
 - **Number of matches:** the numeric comparison logic for the **value**.
 - **Value:** the threshold of matches to trigger the rule.
 - **Limit to organization:** allows you to limit the alert to a specific organization in runZero
 - **Limit to site:** allows you to limit the alert to a specific site in runZero
 - **Action:** notify
 - **Notification channel:** name of the alert channel you created in step 2.
 - **Notification template:** name of the alert template you created in step 2.

runZero sample story JSON

```
{
  "schema_version": 4,
  "standard_lib_version": 6,
```

```

"description": null,
"guid": "7706b502e51b4c68f0dbe9721c88d665",
"slug": "runzero_sample_story",
"exported_at": "2022-10-27T22:22:37Z",
"agents": [
  {
    "type": "Agents::WebhookAgent",
    "name": "RECEIVE alerts from runZero",
    "disabled": false,
    "guid": "ab1f5bf7be49c943d893ab993ade9421",
    "options": {
      "path": "da6e90c0b276bdda97e6bfa31ad50787",
      "secret": "5e290305e0a483bf843f1213f0f21dda",
      "verbs": "get,post"
    },
    "reporting": {
      "time_saved_value": 0,
      "time_saved_unit": "minutes"
    },
    "monitoring": {
      "monitor_all_events": false,
      "monitor_failures": false,
      "monitor_no_events_emitted": null
    },
    "width": null
  },
  {
    "type": "Agents::HTTPRequestAgent",
    "name": "GET assets from runZero",
    "disabled": false,
    "guid": "d54d8f3f5b0ea857308f45f61a224547",
    "options": {
      "url": "https://console.runZero.com/api/v1.0/export/org/assets.json<<get_s
earch_from_regex.individual_item>>",
      "content_type": "application_json",
      "method": "get",
      "headers": {
        "Authorization": "Bearer <<CREDENTIAL.runzero_export_token>>"
      }
    },
    "reporting": {
      "time_saved_value": 0,
      "time_saved_unit": "minutes"
    },
  },

```

```

"monitor_all_events": false,
  "monitor_failures": false,
  "monitor_no_events_emitted": null
},
"width": null,
"schedule": []
},
{
  "type": "Agents::EventTransformationAgent",
  "name": "GET search from URL",
  "disabled": false,
  "guid": "d62ddd189bb9080d778d900f2292504d",
  "options": {
    "mode": "extract",
    "matchers": [
      {
        "path": "=receive_alerts_from_runzero.body.search_url",
        "regexp": "\\?.*",
        "to": "search"
      }
    ]
  },
  "reporting": {
    "time_saved_value": 0,
    "time_saved_unit": "minutes"
  },
  "monitoring": {
    "monitor_all_events": false,
    "monitor_failures": false,
    "monitor_no_events_emitted": null
  },
  "width": null,
  "schedule": null
},
{
  "type": "Agents::EventTransformationAgent",
  "name": "GET search from regex",
  "disabled": false,
  "guid": "41afc5c4d4c0fdb694d14e1eb380688a",
  "options": {
    "mode": "explode",
    "path": "=get_search_from_url.search",
    "to": "individual_item"
  },
}

```

```

"time_saved_value": 0,
  "time_saved_unit": "minutes"
},
"monitoring": {
  "monitor_all_events": false,
  "monitor_failures": false,
  "monitor_no_events_emitted": null
},
"width": null,
"schedule": null
},
{
  "type": "Agents::EventTransformationAgent",
  "name": "LOOP through assets for follow up actions",
  "disabled": false,
  "guid": "b9ada61162bdf006f09d264845ebd304",
  "options": {
    "mode": "explode",
    "path": "=get_assets_from_runzero.body",
    "to": "individual_item"
  },
  "reporting": {
    "time_saved_value": 0,
    "time_saved_unit": "minutes"
  },
  "monitoring": {
    "monitor_all_events": false,
    "monitor_failures": false,
    "monitor_no_events_emitted": null
  },
  "width": null,
  "schedule": null
},
{
  "type": "Agents::EmailAgent",
  "name": "SEND Email Action",
  "disabled": false,
  "guid": "0e9647b114f315d480f1f85c58e481d3",
  "options": {
    "recipients": "youremail@email.com",
    "reply_to": "youremail@email.com",
    "sender_name": "Your Name",
    "subject": "<<receive_alerts_from_runzero.body.rule_name>>",
    "body": "Alert: <<receive_alerts_from_runzero.body.rule_name>>\n\\<br />\n"
  }
}

```

```

ts_for_follow_up_actions.individual_item.id>>\n\\<br />\n\\<br />\nAddresses: <<lo
op_through_assets_for_follow_up_actions.individual_item.addresses>>\n\\<br />\n\\<
br />\nServices:\n<<NEWLINE_TO_BR(NEAT_JSON(loop_through_assets_for_follow_up_acti
ons.individual_item.services))>>"
  },
  "reporting": {
    "time_saved_value": 0,
    "time_saved_unit": "minutes"
  },
  "monitoring": {
    "monitor_all_events": false,
    "monitor_failures": false,
    "monitor_no_events_emitted": null
  },
  "width": null,
  "schedule": null
},
{
  "type": "Agents::HTTPRequestAgent",
  "name": "TAG assets in runZero",
  "disabled": false,
  "guid": "c3171fed91ec40d39572260295a63ae0",
  "options": {
    "url": "https://console.runZero.com/api/v1.0/org/assets/<<loop_through_ass
ets_for_follow_up_actions.individual_item.id>>/tags",
    "content_type": "application_json",
    "method": "patch",
    "headers": {
      "Authorization": "Bearer <<CREDENTIAL.runzero_org_token>>"
    },
    "payload": {
      "tags": "hello=from_tines"
    }
  },
  "reporting": {
    "time_saved_value": 0,
    "time_saved_unit": "minutes"
  },
  "monitoring": {
    "monitor_all_events": false,
    "monitor_failures": false,
    "monitor_no_events_emitted": null
  },
  "width": null,

```

```

}
],
"diagram_notes": [
  {
    "content": "Hello! Thanks for importing the runZero Sample Story. \n\nThis s
tory is a basic example of what you can do with runZero and Tines. \n\nPrerequisit
es:\n1. Create a text [credential](https://www.tines.com/docs/credentials/text) ca
lled `runzero_export_token`. Set the value as your [runZero export token](docs/lev
eraging-the-api.md), which can be obtained from the desired [runZero organization]
(https://console.runzero.com/organizations) page. Export tokens start with `ET`. \n
2. Create a text [credential](https://www.tines.com/docs/credentials/text) called
`runzero_org_token`. Set the value as your as your [organization API token](docs/l
everaging-the-api.md), which can be obtained from the desired [runZero organizatio
n](https://console.runzero.com/organizations). Organization API tokens start with
`OT`. \n\nThe steps are as follows:\n\n1. **RECEIVE alerts** provides the webhook
destination for your runZero alerts \n\n2. **GET search from URL** parses out the
search string from the URL provided in the runZero alert \n\n3. **GET search from
regex** takes the list provided in step 2 and allows you to use the individual val
ue \n\n4. **GET assets from runZero** uses the runZero Export API to get the list
of assets related to the alert \n\n5. **LOOP through assets for follow up actions*
* takes the list of assets and sends each to the next steps individually \n\n6. **
NOTE**: this is the step that you could implement more custom logic in most cases.
All of the initial runZero and Tines data transfer is done, but it's your chance t
o customize this story to fit your use case. \n\n7. **SEND email action** simply s
ends an email to the destination of your choice \n\n8. **TAG assets in runZero** a
dds a tag in runZero to each asset showing an example of how you might reach back
into runZero after doing other automated activities ",
    "position": [
      165.0,
      180.0
    ],
    "guid": "0e4b9b331dc65316849f7b089bf7bde2",
    "width": 375
  }
],
"links": [
  {
    "source": 0,
    "receiver": 2
  },
  {
    "source": 1,
    "receiver": 4
  },
],

```

```

"source": 2,
  "receiver": 3
},
{
  "source": 3,
  "receiver": 1
},
{
  "source": 4,
  "receiver": 6
},
{
  "source": 4,
  "receiver": 5
}
],
"diagram_layout": "{\"ab1f5bf7be49c943d893ab993ade9421\": [555, 180], \"d54d8f3f5b0ea857308f45f61a224547\": [555, 450], \"d62ddd189bb9080d778d900f2292504d\": [555, 270], \"41afc5c4d4c0fdb694d14e1eb380688a\": [555, 360], \"b9ada61162bdf006f09d264845ebd304\": [555, 555], \"0e9647b114f315d480f1f85c58e481d3\": [555, 690], \"c3171fed91ec40d39572260295a63ae0\": [780, 690]}\",
"send_to_story_enabled": false,
"entry_agent_guid": null,
"exit_agent_guids": [],
"exit_agent_guid": null,
"keep_events_for": 604800,
"reporting_status": true,
"send_to_story_access": null,
"send_to_stories": [],
"form": null,
"forms": []
}

```

Thinkst Canary

All runZero editions integrate with Thinkst Canary by providing quick access from the Canary console to your asset data in the runZero Console. Setting up the integration is as simple as one change to your Canary settings.

Integrate runZero with Thinkst Canary

1. Log in to your runZero console.
2. Log in to your Canary console.
3. Go to **Global Settings** under the gear icon.
4. Click on **Integrations**.

5. Toggle the runZero switch.

Accessing runZero data from Canary

After the integration is enabled within the Canary settings, runZero data will be available through any incident. When viewing an incident, click the magnifying glass icon under source IP or reverse IP lookup to open a search in your runZero console.

Analyzing results

Reviewing results

Task details

After each scan task completes, the task details page will list a summary of how many assets were affected. To understand the numbers, it's important to remember that runZero doesn't just rely on IP addresses. Instead, it fingerprints the assets based on how they respond to probes, and tries to catch situations where known assets change IP addresses.

The change summary box on the task details page includes the following statistics:

- Asset changes:
 - **Newly discovered assets** are devices that were found during the task for which no device with matching fingerprints was previously seen.
 - **Assets marked offline** are assets that runZero has previously seen on the scanned network, but that didn't respond on any of the IP addresses during this scan. When this happens, the asset is marked offline. The offline status is a flag on the asset, and doesn't count as a change to the asset. Assets may be marked offline because the device was powered down or disconnected, or because of network problems.
 - **Assets back online** are assets that were marked offline at some point in the past, but the runZero Explorer got a response from them during this scan. The online status is a flag on the asset, and doesn't count as a change to the asset.
 - **Assets changed** is the number of assets where some property of the asset was modified, other than its online status. Examples include changes to the device's IP addresses or hostname, or responses from new ports or protocols.
 - **Assets unchanged** is the number of assets that were seen exactly where runZero found them in the last scan, with no changes to their responses.
 - **Assets ignored** is the number of occasions where the Explorer got a response from probing an IP address, but it turned out to be bogus in some way. This typically happens when a web proxy, stateful firewall, or SIP gateway responds as if it is the asset at every address on a subnet.
 - **Assets updated by task** is the total of **Assets changed** plus **Assets unchanged**. It indicates the number of asset records that are now up-to-date.
- User changes:
 - **Newly discovered users** are users that were seen for the first time during the integration sync.
 - **Users changed** are users that had attributes change during the integration sync.
 - **Users unchanged** are users that did not change during the integration sync.
 - **Users updated by task** is the total number of **Users changed** and **Users unchanged**, indicating how many user records are now up-to-date.
- Group changes:
 - **Newly discovered groups** are groups that were seen for the first time during the integration sync.

- **Groups changed** are groups that had attributes change during the integration sync.
- **Groups unchanged** are groups that did not change during the integration sync.
- **Groups updated by task** is the total number of **Groups changed** and **Groups unchanged**, indicating how many group records are now up-to-date.

Dashboard & inventory views

The dashboard will be populated with results after the first scan completes. The dashboard provides trend data and insights that will help you assess how your inventory is changing over time. You can select a time period and site for the trend data using the selectors at the top right of the dashboard page.

The main asset trends graph shows the number of assets in each of the four main states – live, offline, scanned and unscanned. Beneath the graph are additional asset breakdowns, each of which shows a top 10 of an asset category – asset type, operating system, hardware and tags.

The service trends graph shows how many total services were found in your asset inventory, along with breakdowns for ARP, ICMP, TCP and UDP. Below the service trends graph are breakdowns of the top 10 TCP ports, UDP ports, protocols and products detected.

Clicking the menu button at the top right of each table and selecting “View more” shows a more detailed inventory by category.

Insights from queries

Queries and reports can help you gain valuable insights, but you may wonder where to get started. We recommend trying the pre-built queries in the [Query Library](#) first. Some of these queries are a result of runZero’s [Rapid Response](#) to emerging threats and are described on our [blog](#).

runZero’s query language allows you to search and filter your asset inventory based on asset fields and value pairs. See the [documentation](#) about querying your data. Once you are familiar with the query language you can write your own queries.

You can set queries to run automatically by opening the query and setting the “Automatically track query results on the dashboard”. The query will run when scans complete, and you will be notified of any resulting insights on the dashboard page.

Sample Queries

Asset inventory

- Equipment that is likely 8+ years old: `alive:t mac_age:>8years`
- Assets with end-of-life OS: `os_eol:<now`
- Virtual machines: `has:virtual`
- Devices acting as a router: `router:true`

- Devices that may be bridging: `has_public:t` and `has_private:t`

Service inventory

- Protocol on a non-standard port example: `protocol:ssh not port:22`
- Publicly addressed assets running RDP or VNC: `has_public:t` and `(protocol:rdp or protocol:vnc)`
- Authenticated web services that are not encrypted: `(_asset.protocol:http AND not _asset.protocol:tls) AND (html.inputs:"password:" OR last.html.inputs:"password:" OR has:http.head.wwwAuthenticate OR has:last.http.head.wwwAuthenticate)`
- Older TLS versions in use: `alive:t AND protocol:"=tls" AND (tls.versionName:"=TLS 1.0" OR tls.versionName:"=TLS 1.1")`

Some other sample queries are described in our blog entries:

- [Finding duplicate SSH host keys](#)
- [Identifying rogue remote access solutions](#)
- [Finding device serial numbers](#)

Reports

After viewing the dashboard and inventory, your next stop should be the runZero [Reports](#) page.

Switch topology

This report uses SNMP information to map how the switches on your network are connected. Each switch displays its IP address, name, and the number of assets connected to it. If runZero detected MAC addresses that were not found as part of the scan scope, you will see a number of unmapped assets indicated below the switch.

You can click on a switch to see a pop-up with the number of identified and unmapped assets. From there, you can click to view the unmapped assets, and be taken to a table of unmapped MACs by switch port.

Double-clicking on a switch will expand that part of the diagram and show the individual assets connected to the switch.

The switch topology report won't always be entirely accurate as it's based on which switch claims to have seen each MAC address, and this may not always be the nearest access switch. Our algorithm looks for the switch port with the least number of shared MACs to find the best match, but this may not give the answer you expect, depending on switch cache timeouts and how the switches were scanned.

Subnet utilization

The subnet utilization report lists the subnets scanned on your network, and what percentage of each is in use. For example, if you have scanned 10.0.1.0/24 and found 25 assets, the report will show that 10% of the available IP addresses in the subnet are in use.

Network bridges

The network bridges report is a way to find devices that bridge multiple network segments. It can be useful to locate unintentional bridging between your internal networks and the Internet.

The report shows your internal networks in green, and external networks in red. It then shows you the multihomed assets which bridge an internal network to an external one.

RFC 1918 coverage

The RFC 1918 coverage report is a way to view how much of the private internal network address space has been scanned for assets. It can help you discover rogue assets, unscanned subnets, and secondary interfaces on scanned devices. More information is in the section on [coverage reports](#).

Unmapped MACs

This report uses SNMP information to list MAC addresses runZero found evidence for, but which weren't encountered as addresses of assets during the network scan. The MAC addresses are grouped by the switch that reported them, along with information about the vendor, manufacture date and switch port of the possible asset, to help identify them.

Outliers

The outliers reports allow you to obtain a summary of how often different values occur in specific attributes of assets and services. The values are sorted from most frequent to least frequent.

For example, the HTTP servers outliers report will list all of the HTTP servers encountered by runZero, starting with the most common.

As well as the one-click outliers reports, you can produce an outliers report for any asset or service attribute.

- Switch topology to identify how your assets are connected and find "unmapped" MAC addresses (in red) that were not included in your scan scope (a summary of which is in the Unmapped MACs report)
- Bridging to visualize what hosts may have both public and private connections

- RFC 1918 coverage that can identify potential blindspots on your network like missing (unscanned) subnets, rogue devices, and “hinted” IPs that are secondary interfaces on unscanned network ranges.
- See the “View all” button at the top right for a list of other reports to investigate outliers

Domain membership report

The domain membership report lists the Active Directory domains encountered by runZero, and lists how many assets are in each.

Analysis reports

Platform

Analysis reports are more advanced reports. They may run as tasks, rather than being generated on-the-fly.

The first analysis report is [Compare Sites](#), which generates searchable reports of the differences between two sites.

The Outlier Overview Report analyzes assets across the organization and summarizes the most unusual values for an assortment of key attributes such as hardware type and SNMP enterprise ID.

The Specific Outlier Report allows you to select an attribute which has outlying values and get a detailed breakdown of those values.

The Organization Overview Report builds a high level summary report of the entire organization. It can optionally include lists of assets found.

Alerts

Professional

Platform

As well as manually generated reports and queries, runZero also supports [automatic alerts](#) to designated channels for post-scan inventory queries, asset changes, Explorer and scan issues, security operations, or API events.

Available channels are internal notifications in the runZero web console, email, or webhooks that can enable integration with services such as Slack or Mattermost. Alerts use the same query language as the sample queries above, so this is a good way to automate proactive notification for critical events.

Data formats

runZero consumes and produces a handful of data formats. This page provides examples of these formats and describes the fields and use cases for each.

Formats

- [Scan data \(sample\)](#)
- [Asset data \(sample\)](#)
- [Change reports \(sample\)](#)

Scan data

The raw output produced by the runZero Explorer and the runZero Scanner is the **scan data**. This is newline-delimited JSON – JSONL – that represents the unprocessed output of the scan engine. This format is returned when downloading the task data for an Explorer-run scan and correlates to the scan.runzero.gz file created by the command-line scanner. The runZero Inventory view is built by processing scan data in chronological order to create the current state at a given point in time.

Scan data can be imported into an existing site through the Inventory Import menu of the web console and through the `--import` parameter of the command-line scanner. Each line of the file is a JSON object that specifies a type and a 64-bit Unix timestamp.

The example below is the raw scan data for a single Apple Mac Mini:

```
{
  "type": "config",
  "ts": 1597259738842951567,
  "probes": [
    "arp",
    "bacnet",
    "dns",
    "dtls",
    "echo",
    "ike",
    "ipmi",
    "mdns",
    "memcache",
    "mssql",
    "natpmp",
    "netbios",
    "ntp",
    "openvpn",
    "pca",
    "rdns",
    "rpcbind",
    "sip",
    "snmp",
    "ssdp",
    "syn",
    "tftp",
    "ubnt",
    "wlan-list",
    "wsd"
  ],
  "addresses": [
    "192.168.0.1",
    "192.168.30.1",
    "192.168.40.1"
  ],
  "networks": [
    "192.168.0.1/24",
    "192.168.30.1/24",
    "192.168.40.1/24"
  ],
  "params": {
    "arp-fast": "false",
    "bacnet-port": "47808",
    "clock-offset": "0",
    "dns-port": "53",
    "dns-resolve-name": "www.google.com",
    "dns-trace-domain": "helper.rumble.network",
    "dtls-ports": "443,3391,4433,5246,5349,5684",
    "excludes": "",
    "ike-port": "500",
    "ipmi-port": "623",
    "max-group-size": "4096",
    "max-host-rate": "40",
    "max-sockets": "512",
    "mdns-port": "5353",
    "memcache-port": "11211",
    "mssql-port": "1434",
    "nameservers": "",
    "natpmp-port": "5351",
    "netbios-port": "137",
    "nopcap": "false",
    "ntp-port": "123",
    "openvpn-ports": "1194",
    "passes": "1",
    "pca-port": "5632",
    "probes": "arp,bacnet,dns,dtls,echo,ike,ipmi,mdns,memcache,mssql,natpmp,netbios,ntp,openvpn,pca,rdns,rpcbind,sip,snmp,ssdp,syn,connect,tftp,ubnt,wlan-list,wsd",
    "rate": "1000",
    "rdns-max-concurrent": "64",
    "rpcbind-port": "111",
    "rpcbind-port-nfs": "2049",
    "screenshots": "true",
    "sip-port": "5060",
    "skip-broadcast": "true",
    "snmp-comms": "public,private",
    "snmp-poll-interval": "300",
    "snmp-port": "161",
    "snmp-timeout": "5",
    "snmp-v3-auth-passphrase": "",
    "snmp-v3-auth-protocol": "none",
    "snmp-v3-context": ""
  }
}
```

```

sername":"","ssdp-port":"1900","syn-max-retries":"2","syn-udp-trace-port":"6553
5","tcp-ports":"1300,5554,8020,20034,47001,41080,2601,2604,2638,5060,7181,10202,46
79,2181,34205,13,2323,5601,18881,50070,139,1129,2199,2375,4444,902,1440,2103,3291
3,1311,9524,8028,8883,13364,37718,512,3200,5683,10203,81,1091,5222,8081,13838,3777
7,1,5672,8095,65535,21,540,548,1102,27080,28017,34443,40007,6060,6542,8300,27888,4
786,9443,2049,3050,5984,46823,12221,1352,6405,26122,7210,41025,1103,1530,1883,883
4,443,9100,45230,1234,3128,5432,12397,111,993,3780,5250,6112,524,5247,20031,1211,1
755,5985,6070,8880,1241,3690,6002,1035,4000,8080,9081,2362,23,587,921,8903,31001,1
43,2598,3273,6101,8812,10628,25,113,513,1720,2533,6905,32764,38080,5040,20010,600
1,6660,8471,82,2222,5093,6262,6379,8545,384,5168,20222,7579,998,3057,3217,6106,939
1,9,2380,5520,9060,19300,30718,49,84,161,5900,10001,8009,19,617,2100,5580,38292,8
5,6667,10443,42,2121,5986,23791,515,1199,10008,16993,631,2083,8443,9527,13500,2701
7,30000,41523,554,5061,4659,8333,9855,5355,7001,1000,1220,5521,11234,20000,6988,33
51,7547,7,1900,7778,9160,31099,1030,10616,7902,8090,12174,1533,135,5631,623,5038,9
300,19888,14330,109,1433,15672,1581,3790,5632,9999,80,2381,4840,7800,61616,1101,11
28,1494,3311,9092,11000,110,995,1098,5800,523,8087,10098,28784,407,7777,9090,1981
0,34963,50000,502,1100,8161,8180,9152,11099,2379,8023,88,1582,20101,16102,16992,15
83,5814,5938,20111,11211,636,27000,1158,5400,5920,7443,9530,20171,8800,9099,7474,8
222,10000,2082,8902,50013,689,771,7080,8098,8686,22,1024,12345,105,9080,9111,4700
2,4433,44818,4848,6080,7071,8303,62078,705,873,6000,7077,8503,9495,34964,5666,1720
0,5433,7801,11333,12401,12203,25025,264,2525,3628,9809,26000,50090,9000,2967,137,4
730,5051,8899,10050,52302,8400,53,389,402,4443,7700,62514,1090,5353,6082,6661,4031
7,8089,17185,912,5405,28222,465,4445,6503,8014,57772,23472,1080,7021,8088,22222,50
00,9084,18264,8888,6050,7144,41524,69,1811,44334,102,6502,1521,2809,9471,888,5351,
5498,123,34962,8205,9042,2947,3389,5555,10080,10162,9200,500,1089,7510,15200,2207,
3500,8008,9418,3460,6504,7770,25000,55553,179,783,8012,9595,46824,7580,5560,49152,
83,903,1604,3632,4322,4567,445,8030,9390,3817,8901,10051,9002,27019,910,1099,3000,
3299,222,7787,37,2000,3306,48899,7879,79,1723,3037,3312,8000,23943","tftp-port
s":"69","ubnt-port":"10001","verbose":"true","wlan-list-poll-interval":"300","wsd-
port":"3702"},"scan_targets":{"networks":["192.168.0.5/32"],"enable_dns":true,"ena
ble_ip6":false,"inputs":["192.168.0.5"],"dns_timeout":2000000000,"concurrency":2
4},"version":"1.10.0 (build 20200804052508) [eae4e551f9f0ce5ab3bf0a1410b2ed5098db0
97e]"}
{"type":"status","ts":1597259758851278069,"level":"info","source":"connect","ms
g":"waiting on TCP probes to complete"}
{"type":"stats","ts":1597259769858733899,"stats":{"elapsed":31,"progress":94,"rate
LimitTime":24246445895,"recv":546,"recvBytes":37404,"recvError":10,"recvRate":1
7,"resultCount":22,"secondsLeft":1,"sent":503,"sentBytes":36833,"sentError":0,"sen
tRate":16,"startTime":1597259738843056937}}
{"type":"result","ts":1597259738844372618,"host":"192.168.0.5","port":"0","prot
o":"icmp","probe":"echo","info":{"icmp.addr":"192.168.0.5","icmp.rtt":"54121
4","ip.flags":"DF","ip.id":"0","ip.tos":"0","ip.ttl":"64"}}
{"type":"result","ts":1597259739077542717,"host":"192.168.0.5","port":"137","nam
e":"MACMINI-EE7C7B","proto":"udp","probe":"netbios","info":{"netbios.domain":"WORK

```

```
ios.macVendor":"Apple, Inc."}}
{"type":"result","ts":1597259739219939211,"host":"192.168.0.5","port":"137","proto":"udp","probe":"netbios","info":{"netbios.addr":"192.168.0.5"}}
{"type":"result","ts":1597259739471688048,"host":"192.168.0.5","port":"0","proto":"arp","probe":"arp","info":{"arp.mac":"f0:18:98:ee:7c:7b","arp.macDateAdded":"2017-12-23","arp.macVendor":"Apple, Inc.,"source":"arp"}}
{"type":"result","ts":1597259739826114224,"host":"192.168.0.5","port":"5353","name":"Developers-Mac-mini","proto":"udp","probe":"mdns","info":{"mdns.replies":"5.0.168.192.in-addr.arpa.=PTR,Developers-Mac-mini.local."}}
{"type":"result","ts":1597259740197709484,"host":"192.168.0.5","port":"445","proto":"tcp","probe":"connect","info":{"ntlmssp.dnsComputer":"Developers-Mac-mini.local","ntlmssp.dnsDomain":"local","ntlmssp.negotiationFlags":"0x62898235","ntlmssp.netbiosComputer":"DEVELOPERS-MAC-MINI","ntlmssp.netbiosDomain":"MACMINI-EE7C7B","ntlmssp.ntlmRevision":"15","ntlmssp.targetName":"MACMINI-EE7C7B","ntlmssp.timestamp":"0x01d670dd06500880","ntlmssp.version":"6.1.7600","protocol":"smb1\tsmb2\tsmb3","smb.capabilities":"0x00000066","smb.dialect":"0x0302","smb.guid":"ff12583f-5ba1-53b8-8ff3-a48a394056f7","smb.nativeLM":"@(#)PROGRAM:smbd PROJECT:smbx-499.60.1","smb.nativeOS":"Darwin","smb.sessionID":"0x9551b7cb00000001","smb.signing":"required","source":"mdns"}}
```

The data contains four types of object:

- **Scan config:** The {"type":"config"} object contains the full set of parameters for the scan as well as the version of the scan engine, and on Windows, the version of npcap installed. This record is used to determine the scan targets, which is used by the analysis engine to determine whether a given IP address was in scope.
- **Scan status:** The {"type":"status"} object contains diagnostic output from the scan engine. This can highlight issues that occurred while the scan was running.
- **Scan stats:** The {"type":"stats"} object represents point in time statistics for the scan. This will include the number of packets sent, received, and the progress estimate.
- **Scan result:** The {"type":"result"} object is target response for a specific probe. This can include TCP SYN+ACK replies, ICMP replies, or the result of application-layer probes, such as SNMP query responses, or HTTP screenshots. Scan Results are analyzed and correlated to create to the Asset Data format.

The scan stats sub-fields are defined below:

Field	Description
cpu	CPU Core Percent * 100. 100% of one core would be 100000.
elapsed	The number of seconds since the scan started.
fdcount	The number of open file descriptors.
memory	The current memory usage in bytes.

rateLimitTime	The number of Unix nanoseconds spent idling in the rate limiter.
recv	The number of packets received from the network.
recvBytes	The number of bytes received from the network.
recvError	The number of errors receiving from the network.
recvRate	The average packet receive rate for the scan.
resultCount	The total number of findings from the scan.
routines	The number of internal goroutines in the scan engine.
secondsLeft	The estimated seconds left to complete the scan.
sent	The number of packets sent the network.
sentBytes	The number of bytes sent the network.
sentError	The number of errors sending to the network.
sentRate	The average packet send rate for the scan.
startTime	The Unix timestamp in nanoseconds of when the scan started.

The scan result object type contains the following fields in addition to type and ts:

Field	Description
host	The IP address associated with the response.
name	An optional hostname returned as part of this response.
port	The TCP or UDP port or zero for other protocols.
proto	The transport protocol, one of arp, icmp, tcp, or udp.
probe	The specific internal probe name that returned this response.
info	The result details object where all keys and values are strings.

The info object contains probe-specific response data. The key names are typically in the format of probe.subfield, with a few exceptions, and the values are always strings, even for numeric and array content. Multiple values for a key are represented as a tab-delimited array. Empty values are never reported for info keys. A given scan may return multiple result objects for a single probe, sometimes with duplicate values. These responses are correlated, deduplicated, and merged during the next phase of processing.

Asset data

The correlated and fingerprinted assets shown in the web console Inventory view and in the assets.jsonl file produced by the runZero Scanner are the **asset data**. This data represents the state of each unique asset at a point in time and is built up by processing one or more sets of scan data.

runZero supports a few variants of the asset data, including line-delimited JSON (JSONL), standard JSON documents, and a simplified CSV export. The JSONL format is the easiest to work with as it supports incremental processing without having to load the entire response into memory.

The example below is the correlated asset data for a scan of a single Apple Mac Mini:

```
{
  "id": "b73f8e09-78a6-4d2b-979d-e63908f28251",
  "created_at": 1597259778,
  "updated_at": 1597259778,
  "organization_id": "b7fb13a7-701d-4ca5-b0e6-6f28f06cc866",
  "site_id": "52d60c51-8dee-4f09-94e5-2dee30050a25",
  "alive": true,
  "last_seen": 1597259750,
  "first_seen": 1597259738,
  "detected_by": "arp",
  "type": "Desktop",
  "os": "Apple macOS",
  "os_version": "10.15",
  "hw": "Apple Mac Mini (Late 2018)",
  "addresses": ["192.168.0.5"],
  "addresses_extra": [{"fe80::1c9d:c567:8db1:d79b"}, {"f0:18:98:ee:7c:7b"}],
  "macs": [{"f0:18:98:ee:7c:7b"}],
  "mac_vendors": ["Apple, Inc."],
  "names": ["MACMINI-EE7C7B", "DEVELOPERS-MAC-MINI", "DEVELOPERS-MAC-MINI.LOCAL"],
  "tags": {},
  "domains": [],
  "services": {
    "192.168.0.5/0/arp/": {
      "arp.mac": "f0:18:98:ee:7c:7b",
      "arp.macDateAdded": "2017-12-23",
      "arp.macVendor": "Apple, Inc.",
      "source": "arp",
      "ts": "1597259739"
    },
    "192.168.0.5/0/icmp/": {
      "icmp.addr": "192.168.0.5",
      "icmp.rtts": "541214",
      "ip.flags": "DF",
      "ip.id": "0",
      "ip.tos": "0",
      "ip.ttl": "64",
      "ts": "1597259738"
    },
    "192.168.0.5/137/udp/": {
      "netbios.addr": "192.168.0.5",
      "netbios.domain": "WORKGROUP",
      "netbios.mac": "f0:18:98:ee:7c:7b",
      "netbios.macDateAdded": "2017-12-23",
      "netbios.macVendor": "Apple, Inc.",
      "protocol": "netbios",
      "ts": "1597259739"
    },
    "192.168.0.5/22/tcp/": {
      "banner": "SSH-2.0-OpenSSH_7.9",
      "ip.flags": "DF",
      "ip.id": "0",
      "ip.tos": "0",
      "ip.ttl": "64",
      "protocol": "ssh",
      "service.cpe23": "cpe:/a:openbsd:openssh:7.9",
      "service.family": "OpenSSH",
      "service.vendor": "OpenBSD",
      "service.version": "7.9",
      "source": "mdns",
      "ssh.hostKey.data": "AAAAB3NzaC1yc2EAAAADAQABAAQcJGYTFcSp2Fs/R8dboLYiQ6PPruLZYanYH3SCYYr5QgC1SIF3AURGYTMnUDAS+tTI/Pquwowkgiq3rtfsQMasCrahbPahwiOLTupsuLNp3evXYYSf8ZQFyBN8iz5cys06u+yczqWG7Fu8mgpS8zwCwN7yRrbFwd8+Hp6GgFUU4Z6jUQoZu7iajpbSXlTA90YKXQIZ0m8qc4mPLT/uHw9nxNmExWA1V/2ZeoS59NGSV8zFMKb52SOXKhkvHAIUVh5NJD AudxK4uP4eG6dxr8btYtVKIOYKlsLdSBSfHvSCvVvlb7DKJBiMXG+qspt33Zd73o4S9ICh20aSbVt7h/NZ3",
      "ssh.hostKey.md5": "75:9b:a2:e6:10:da:72:8a:11:91:3f:a1:43:14:7f:2e",
      "ssh.hostKey.sha256": "SHA256:xVJfddKBJ9E5jstVCj0zY8763Rnxy2ppzaLZX0+cHc",
      "ssh.hostKey.type": "ssh-rsa",
      "syn.rtt": "542019",
      "tcp.options": "MSS:05b4",
      "tcp.ts": "2009546838",
      "tcp.urg": "0",
      "tcp.win": "65535",
      "ts": "1597259740"
    },
    "192.168.0.5/3031/tcp/": {
      "source": "mdns",
      "ts": "1597259740"
    },
    "192.168.0.5/3283/tcp/": {
      "source": "mdns",
      "ts": "1597259740"
    },
    "192.168.0.5/445/tcp/": {
      "ip.flags": "DF",
      "ip.id": "0",
      "ip.tos": "0",
      "ip.ttl": "64",
      "ntlmssp.dnsComputer": "Developers-Mac-mini.local",
      "ntlmssp.dnsDomain": "local",
      "ntlmssp.negotiationFlags": "0x62898235",
      "ntlmssp.netbiosComputer": "DEVELOPERS-MAC-MINI",
      "ntlmssp.netbiosDomain": "MACMINI-EE7C7B",
      "ntlmssp.ntlmRevision": "15",
      "ntlmssp.targetName": "MACMINI-EE7C7B",
      "ntlmssp.timestamp": "0x01d670dd06500880",
      "ntlmssp.version": "6.1.7600",
      "protocol": "smb1\\tsmb2\\tsmb3",
      "smb.capabilities": "0x00000066",
      "smb.dialect": "0x0302",
      "smb.guid": "ff12583f-5ba1-53b8-8ff3-a48a394056f7",
      "smb.nativeLM": "@(#)PROGRAM:smbd PROJECT:smbx-499.60.1",
      "smb.nativeOS": "Darwin",
      "smb.sessionID": "0x9551b7cb00000001",
      "smb.signing": "required",
      "source": "mdns",
      "syn.rtt": "624414",
      "tcp.options": "MSS:05b4",
      "tcp.ts": "2009544186",
      "tcp.urg": "0",
      "tcp.win": "65535",
      "ts": "1597259740"
    },
    "192.168.0.5/5353/udp/": {
      "hw.device": "Desktop",
      "hw.family": "Apple Mac Mini (Late 2018)"
    }
  }
}
```

```

y": "Mac mini", "hw.product": "Mac mini (Late 2018)", "hw.vendor": "Apple", "mdns.address": "fe80::1c9d:c567:8db1:d79b\t192.168.0.5", "mdns.device.model": "Macmini8,1", "mdns.device.osxvers": "19", "mdns.ports": "eppc/tcp=3031\tnet-assistant/udp=3283\ttrfb/tcp=5900\tstftp-ssh/tcp=22\t smb/tcp=445\tssh/tcp=22", "mdns.replies": "5.0.168.192.in-addr.arpa.=PTR,Developers-Mac-mini.local.\tDeveloper\\226\\128\\153s\\ Mac\\ mini._device-info._tcp.local.=TXT,model=Macmini8,1 osxvers=19\tDevelopers-Mac-mini.local.=A,192.168.0.5\tDevelopers-Mac-mini.local.=AAAA,fe80::1c9d:c567:8db1:d79b\t_eppc._tcp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._eppc._tcp.local.\t_net-assistant._udp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._net-assistant._udp.local.\t_rfb._tcp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._rfb._tcp.local.\t_sftp-ssh._tcp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._sftp-ssh._tcp.local.\t_smb._tcp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._smb._tcp.local.\t_ssh._tcp.local.=PTR,Developer\\226\\128\\153s\\ Mac\\ mini._ssh._tcp.local.", "mdns.services": "ssh/tcp\tstftp-ssh/tcp\teppc/tcp\ttrfb/tcp\t smb/tcp\t net-assistant/udp", "os.cpe23": "cpe:/o:apple:mac_os_x:10.15", "os.family": "Mac OS X", "os.product": "Mac OS X", "os.vendor": "Apple", "os.version": "10.15", "protocol": "mdns", "ts": "1597259740"}, {"192.168.0.5/5900/tcp/": {"ip.flags": "DF", "ip.id": "0", "ip.tos": "0", "ip.ttl": "64", "protocol": "vnc", "source": "mdns", "syn.rtt": "625419", "tcp.options": "MSS:05b4", "tcp.ts": "2009549134", "tcp.urg": "0", "tcp.win": "65535", "ts": "1597259749", "vnc.version": "RFB 003.889"}}, {"credentials": {}, "rtts": {"icmp/echo": [541214]}, "attributes": {"_macs.ipmap": "f0:18:98:ee:7c:7b=192.168.0.5", "ip.ttl.hops": "0", "ip.ttl.host": "192.168.0.5", "ip.ttl.port": "22", "ip.ttl.source": "64", "ip.ttl.source.icmp": "64", "ip.ttl.win": "65535", "match.db": "mdns-device-info-txt", "match.score": "90", "ntlmssp.dnsComputer": "Developers-Mac-mini.local", "ntlmssp.dnsDomain": "local", "ntlmssp.version": "6.1.7600", "os.cpe23": "cpe:/o:apple:mac_os_x:10.15", "os.family": "Mac OS X", "os.product": "Mac OS X", "os.vendor": "Apple", "os.version": "10.15", "smb.guid": "ff12583f-5ba1-53b8-8ff3-a48a394056f7", "smb.nativeLM": "@(#)PROGRAM:smbd PROJECT:smbx-499.60.1", "smb.nativeOS": "Darwin"}, "service_count": 9, "service_count_tcp": 5, "service_count_udp": 2, "service_count_arp": 1, "service_count_icmp": 1, "lowest_ttl": 0, "lowest_rtt": 541214, "last_agent_id": "ca811190-329c-4da3-8cbe-3fd2ddff2663", "last_task_id": "de5a4176-3614-4b71-8939-95b9108124aa", "newest_mac": "f0:18:98:ee:7c:7b", "newest_mac_vendor": "Apple, Inc.", "newest_mac_age": 1513987200000000000, "comments": null, "service_ports_tcp": ["22", "445", "3031", "3283", "5900"], "service_ports_udp": ["137", "5353"], "service_protocols": ["mdns", "netbios", "smb1", "smb2", "smb3", "ssh", "vnc"], "service_products": ["openbsd openssh"], "org_name": "Test Lab", "site_name": "MAC", "agent_name": "TENTACULAR"}

```

Asset Data uses a number of data types for top-level fields, including string arrays, objects, strings, and integers. runZero tracks multiple IP addresses and MACs per asset and these are represented as arrays. For asset-level attributes and services, these are stored as objects with additional structure. Assets are uniquely identified by the `id` field (a V4 UUID) and nearly every other field can be changed between scans, as assets move around the network, change IPs, and open and close services.

Every asset belongs to an organization and a site within that organization.

The core asset data fields are defined below.

Field	Description
id	The unique ID of this asset defined as a v4 UUID.
created_at	The asset created time represented as a 64-bit Unix timestamp in seconds.
updated_at	The asset last update time represented as a 64-bit Unix timestamp in seconds.
organization_id	The organization identifier defined as a v4 UUID.
site_id	The site identifier defined as a v4 UUID.
alive	A boolean indicating whether this asset was found during the last scan of the site.
last_seen	The time the asset last responded represented as a 64-bit Unix timestamp in seconds.
first_seen	The time the asset first responded represented as a 64-bit Unix timestamp in seconds.
detected_by	The protocol used to first detect that this asset was alive during the last scan.
type	A classification that represents a guess of the asset's purpose.
os_vendor	The operating system vendor name as determined by the fingerprinting engine.
os_product	The operating system product name as determined by the fingerprinting engine.
os_version	The operating system version as determined by the fingerprinting engine.
os	The operating system name as determined by the fingerprinting engine.
hw_vendor	The hardware vendor name as determined by the fingerprinting engine.
hw_product	The hardware product name as determined by the fingerprinting engine.
hw_version	The hardware version as determined by the fingerprinting engine.
hw	The hardware definition as determined by the fingerprinting engine.
addresses	An array of IP (v4/v6) addresses for the asset that were within the scan scope.
addresses_extra	An array of IP (v4/v6) addresses for the asset that were outside the scan scope.
macs	An array of MAC addresses associated with this asset.
mac_vendors	An array of MAC address vendors associated with this asset.
names	An array of unique hostnames associated with this asset (uppercase).
domains	An array of unique domain names associated with this asset (uppercase).

tags	A text representation of the user-specified tags associated with this asset.
attributes	An object containing a map of key-value string attributes for this asset.
services	An object containing each associated service with the key representing the service description.
credentials	An object containing a map of any associated credentials (SNMP v2/v3).
rtts	An object containing a map of round-trip measurement times in milliseconds.
service_count	A count of TCP, UDP, ARP, and ICMP services.
service_count_tcp	A count of TCP services.
service_count_udp	A count of UDP services.
service_count_arp	A count of ARP services (0 or 1).
service_count_icmp	A count of ICMP services.
software_count	A count of software results.
vulnerability_count	A count of vulnerability results.
lowest_ttl	The lowest observed source TTL for this asset.
lowest_rtt	The lowest observed source RTT for this asset.
last_agent_id	The v4 UUID of the Explorer responsible for the last scan of this asset.
last_task_id	The v4 UUID of the task responsible associated with the last scan of this asset.
last_task_id	The v4 UUID of the task responsible associated with the last scan of this asset.
newest_mac	The “newest” MAC address by registration date.
newest_mac_vendor	The “newest” MAC address vendor by registration date.
newest_mac_age	The “newest” MAC address registration date as a Unix timestamp in nanoseconds.
comments	User-specified comments associated with this asset.
service_ports_tcp	An array of strings representing the unique TCP ports found on this asset.
service_ports_udp	An array of strings representing the unique UDP ports found on this asset.
service_protocols	An array of strings representing the unique protocols found on this asset.

	asset.
scanned	A TRUE or FALSE value indicating whether the asset has been scanned by runZero.
source_ids	The ID of the data source, mapped to this table .
eol_os	The operating system End-of-Life time represented as a 64-bit Unix timestamp in seconds.
eol_os_ext	The operating system extended End-of-Life time represented as a 64-bit Unix timestamp in seconds.
outlier_score	The 0-5 score range indicating how unusual an asset is compared to the rest of the inventory.
outlier_raw	The heuristic score indicating how unusual an asset is compared to the rest of the inventory.
sources	The name of the data source, mapped to this table .
org_name	The name of the organization associated with this asset.
site_name	The name of the site associated with this asset.
agent_name	The name of the Explorer associated with this asset.
agent_external_ip	The external IP address of the Explorer associated with this asset.
hosted_zone_name	The name of the hosted zone associated with this asset.
subnets	The registered subnets associated with the site this asset is in.

The `services` field contains string keys that contain the unique service identifier with values stored as strings. Multiple values may be stored as tab-delimited strings in the service values. A typical service key looks like `192.168.0.5/22/tcp/`. The components of the service key name consist of address, port, transport, and virtual host (which can be blank).

Change reports

The runZero platform calculates a **change report** after processing each scan. This is a JSON document available for download from the Task Details page with the following structure.

```
{
  "assets":{
    "new":{ "<asset-UUID>": { "Asset Data Fields":"" } },
    "online":{ "<asset-UUID>": { "Asset Data Fields":"" } },
    "offline":{ "<asset-UUID>": { "Asset Data Fields":"" } },
    "changed":{ "<asset-UUID>": { "Asset Data Fields":"" } },
    "summary":{
      "changed":#,
      "new":#,

```

```

"unchanged":#
  } },
"directory_users":{
  "new":{ "<user-UUID>": { "User Data Fields":"" } },
  "changed":{ "<user-UUID>": { "User Data Fields":"" } },
  "summary":{
    "changed":#,
    "new":#,
    "total":#,
    "unchanged":#
  } },
"directory_groups":{
  "new":{ "<group-UUID>": { "Group Data Fields":"" } },
  "changed":{ "<group-UUID>": { "Group Data Fields":"" } },
  "summary":{
    "changed":#,
    "new":#,
    "total":#,
    "unchanged":#
  } },
  "truncated": <true/false>
}

```

The new, online, offline, and changed objects each contain keys consisting of the modified asset IDs with the values represented in the **asset data** format. The summary field indicates overall change statistics for this task. The truncated field is set to true if the change report is incomplete due to reaching the maximum change threshold (1000 asset changes today).

Using the dashboard

The dashboard is a customizable, visual view into your asset inventory. The dashboard has a variety of visualization [widgets](#) that show operational information, trends, insights, goals, sources, and most and least seen graphs.

You can filter this information based on sites and time buckets based on your needs by using the dropdown buttons at the top of the dashboard and selecting the timeframe and/or site you wish to filter by.

You can also [customize](#) the arrangement and size of each widget on the dashboard to suit your particular needs, and access the [widget library](#) to add or remove widgets as desired.

Customizing the dashboard

Each of the widgets on the dashboard can be resized according to your needs by clicking and dragging the bottom right corner of the widget.

Widgets can also be rearranged within the dashboard by clicking and dragging on the widget heading area as indicated by the “cross-hair” mouse cursor icon.

Your customization settings will be saved and applied any time you visit the dashboard so that you can quickly get to the information you’d like to see.

If you wish to revert your dashboard to its default state, click the menu button at the top of the dashboard and select the “Reset dashboard” option. This will reset the visibility, arrangement, and size of all widgets to reflect the original view of the dashboard prior to customizing it.

Widget library

The dashboard widget library is used to add and remove widgets on the dashboard, and can be accessed by clicking the “Widgets” button at the top right of the dashboard.

When you open the widget library, a list of all the available widgets is presented for you to select from.

Widgets included on the dashboard will be marked with a checkmark icon. When clicked, the widget will be deselected and removed from the dashboard when you click the “Save” button.

Conversely, widgets that are not currently visible on the dashboard can be added by clicking on the widget to select it, and then saving the changes.

Widget types

Operational information

- Live assets: number of assets currently alive based on the latest scans
- Active scans: number of scans currently in progress
- Explorers online: number of Explorers reporting as healthy
- Accounts: number of users with access to your console

Trends

- Asset trends: timeline view of asset counts broken out by live, recent, and a couple other states
- Service trends: timeline view of network services seen in your environment

Insights

- Latest insights: automated queries that were ran after your latest scan with result counts

Most and least seen graphs

A variety of asset views by different asset properties. All of these graphs can be toggled between most and least seen based on the counts.

- Asset types
- Operating systems
- Hardware
- RRT latency
- MAC vendors
- Newest MAC age
- TCP ports
- UDP ports
- Products
- Protocols
- Address count
- Extra address count
- Risk
- Criticality
- Tags

Goals

- Goals overview: summary view of all goals that are pinned to display on the dashboard

Asset count and trend tables

A view into the latest count and historical timeline for your assets related to ownership coverage, source types, and custom integrations.

- Asset ownership: view of all ownership types including assets missing owners
- Asset sources: view of all sources of asset data from supported integrations
- Custom integrations: view of all custom sources of asset data from API integrations

Drill down

Except for timeline-based charts, each widget's distinct data points can be viewed in detail by clicking on the relevant table row links or buttons, or by clicking on individual segments of donut, bar or column charts.

Printing and exporting

All chart-based widgets on the dashboard support exporting to SVG, PNG, and JPEG file formats, and table-based widgets can be exported to CSV format, by using the menu button at the top of the widget and selecting the file format you wish to export.

Using the inventory

The [inventory](#) page is the heart of runZero Network Discovery and the key to understanding what is on your network. The inventory displays all assets within the Organization and can be sorted, filtered, and exported to obtain specific views of the environment.

Understanding assets

An [asset](#) within runZero is defined as a unique network entity. Assets may have multiple IP addresses and MAC addresses and these addresses may change as the environment is updated. runZero tracks assets based on several heuristics, including MAC address, IP address, hostnames, and fingerprint results for the operating system and running services.

In most cases, runZero can accurately follow assets over time in environments using DHCP, even across remote subnets. For external networks, scans that are initiated with fully qualified hostnames will consolidate assets based on the hostname, which allows for consistent asset tracking for cloud-based external systems with dynamic IP addresses.

Within an organization, assets are isolated by site, and each site can have address space that overlaps with other sites. Sorting the Inventory view based on the site column can help in these scenarios, as can filtering the Inventory based on a specific site name.

The search field allows the inventory to be filtered based on the specified criteria. Please see the [search query syntax](#) documentation for specific details.

In addition to viewing assets, the Inventory page provides [data export](#) functionality, along with the ability to select assets, and specify the comments field. The *Rescan* action can be used to selectively rescan specific systems from the inventory, while the *Remove Assets* and *Purge Assets* can be used to permanently remove data from the inventory view.

The *Reports* button provides quick access to key reports from the runZero [reports page](#).

Loading assets

Data is loaded into the inventory using the **Scan** and **Import** buttons. The results are analyzed and merged, updating asset information as necessary.

The **Scan** button has two options: Standard Scan and Full RFC 1918 Discovery. The latter is an easy way to set up a fast scan of all private range IP addresses. You can then use the [coverage reports](#) to check for assets in unexpected private address ranges.

The **Import** button has two options. Importing runZero scan data allows you to import data. This means you can scan networks that have no connectivity to the internet, and still view the results in the runZero console. It's also useful for reprocessing old scan data so that you can use the [site compare](#) feature to see how assets have changed over time.

Bulk asset update

The bulk asset update feature allows you to modify assets by [exporting a CSV](#) using the **Export** button, making changes to the data in a spreadsheet program or text editor, and then importing the result back into runZero with the **Import** button. This feature will update existing assets that have a matching `id` value in the organization.

The fields listed below can be updated through the bulk asset update:

- Type
- Operating system
- OS version
- Hardware
- Comments
- Tags
- Owner
- Names
- Domains

The `type`, `os`, `os_version`, and `hardware` fields only accept a single value. The `comments`, `tags`, `owner`, `names`, and `domains` fields each accept multiple values, and a space-delimited list of `field=value` pairs is the standard syntax. The `tags` field can also be specified without `tag=` as just a space-delimited list of values.

Only modifications to the `tags`, `comments`, and `owners` fields will be retained through subsequent scans, any changes to the other supported fields will be overwritten by the latest scan data.

Connecting to other systems

Professional Community Platform

The **Connect** button lets you connect runZero to other systems. The [integrations](#) you're able to connect depends on your license level, but may include tools like cloud and virtualization platforms, endpoint protection solutions, identity and access management tools, and vulnerability and risk platforms. These inbound integrations can also be configured as [scan probes](#) if required.

Viewing services

The Inventory page has a submenu labeled *Services*. This changes the table of data from an asset-focused view to a service-focused view. For each asset, you will see one row for each service runZero detected.

Like the main asset view, the services view has a full search interface. You can filter services by protocol, port, and [many other criteria](#), using the [runZero search language](#).

Viewing screenshots

If the runZero Explorer has access to Google Chrome, it will attempt to take screenshots of web pages it finds while scanning your network. (This feature can be disabled in the scan options when setting up the scan.)

You can view the screenshots for all of your assets via the *Screenshots* submenu, and click through to the asset records for full details.

Viewing software

The inventory page has a submenu labeled *Software*. This flips the table of data from an asset-focused view to a software-focused view. For each asset, you will see one row for each software detected by runZero or a supported integration.

Like the main asset view, the software view has a full search interface. You can filter software by vendor, product, and [many other criteria](#), using the [runZero search language](#).

Viewing vulnerabilities

The inventory page has a submenu labeled *Vulnerabilities*. This flips the table of data from an asset-focused view to a vulnerability-focused view. For each asset, you will see one row for each vulnerability detected by a supported integration.

Like the main asset view, the vulnerability view has a full search interface. You can filter vulnerabilities by CVSS score, name, CVE, and [many other criteria](#), using the [runZero search language](#).

Viewing wireless networks

If the machine running the runZero Explorer has a working WiFi adapter and appropriate system tools installed, the Explorer will attempt to scan for nearby wireless networks. The *Wireless* submenu will show the results of the scan.

The tools required are:

- Windows: netsh.exe (part of modern Windows releases)
- macOS: Airport Utility
- Linux: iwlist, often available via the wireless-tools package.

Understanding assets

runZero treats assets as unique network entities from the perspective of the system running the Explorer. An asset may have multiple IP addresses, MAC addresses, and hostnames and it may move around the network as these attributes are updated. runZero tries hard to follow assets by correlating new scan data with the existing inventory, using multiple attributes.

An asset is always associated with a single site. If the same system happens to be covered by multiple sites, these will be treated as different assets, and will only be correlated against assets within their respective site. This separation by site allows the same network to be scanned from multiple perspectives and compared in a single view within the organization.

After each scan, all assets within the corresponding site are updated. If a system is identified that doesn't match an existing asset, a new asset will be created. If an asset is part of the site and it is not found during a scan, it will be marked as offline. If an asset is not correlated, due to substantial changes to the fingerprint (for example, a new network adapter was installed and the firewall was enabled), the previous asset will be marked as offline, and a new asset will be created to track the new configuration. This can lead to some level of duplication within a site, but these duplicates are usually marked as offline, and can be safely ignored or removed from the inventory by hand.

Asset fields

The following asset fields are available.

Primary addresses

runZero will report at least one and often multiple primary IP addresses for a given asset. These addresses can encompass multiple network interfaces but will only be displayed as a primary address if runZero has scanned it. This requires that the address is within the scan scope of one or more runZero scans.

Secondary addresses

runZero may report one or more secondary addresses, based on network response probes. These are IP addresses that were detected on the asset but were not within the scan scope. Secondary address detection is critical when trying to identify systems that bridge networks that should be isolated.

Hostnames

runZero may report one or more hostnames. These names can be obtained from the initial DNS lookup (when hostnames are provided in the scan scope), from DNS PTR lookups during the scan, and by extracting names advertised within network probe responses.

Operating System (OS)

runZero attempts to fingerprint, and failing that, guess at the operating system running on each asset. If limited information is available, this field may be empty.

Type

runZero attempts to determine the general device type through analysis of fingerprints and running services.

Hardware

runZero attempts to determine the physical (or virtual) hardware if enough information is present.

MAC addresses

runZero may be able to enumerate one or more MAC addresses from the asset. MAC addresses are pulled from ARP if available, but also several network services that can return MAC address information across routed segments.

Services

runZero tries to detect approximately 100 TCP services by default, along with several useful UDP services. These services are in addition to ARP and ICMP. The services field contains a list of the most recently recorded services for the asset.

Round Trip Time (RTT)

runZero records the amount of time certain probes take in order to get a rough sense of the latency between the Explorer and the asset.

Detected by

runZero records which probe was used to identify an asset. For assets that are on remote subnets and have firewalls in place, this field indicates what service was used to obtain a response.

Alive status

runZero tracks whether a given asset was found during the most recent scan where its site was in scope. If the asset was not found, it will be marked as offline until a following scan

detects it again.

First seen

runZero tracks the initial timestamp when an asset was first identified.

Last seen

runZero tracks the last timestamp when an asset responded to a probe during a scan.

Outlier score

runZero computes an outlier score for all assets in your inventory. The outlier score has a value from 0 to 5 (inclusive). It is a heuristic that aims to indicate how unusual the asset is, compared to all of the others in the inventory.

Outlier scores are computed by examining key properties of the asset and its services, working out which values are unusual (infrequent) across the organization, and then computing how many unusual properties each asset has. The more unusual properties, the higher the asset's outlier score will be.

Managing asset risk and criticality

[Professional](#)[Community](#)[Platform](#)

runZero is able to help users assign and evaluate risk and criticality levels to the assets in their inventory. This can help prioritize risk mitigation or vulnerability remediation efforts by allowing users to quickly identify the assets in their organization with the highest levels of risk or criticality.

Defining risk and criticality

The **risk** level assigned automatically to assets in your inventory is inferred from the risk associated with vulnerabilities or risky configurations on that asset and defaults to the value none. Vulnerability risk level may be defined by the vulnerability management solution the vulnerability records are ingested from, or by the risk level assigned to a query vulnerability. The risk level can be overridden, in which case the override is retained until the asset or vulnerability is deleted. For vulnerabilities ingested from integrations, this may occur when the source no longer reports the vulnerability on that asset.

The **criticality** level is assigned manually and defaults to the value unset. This value is intended to be used to denote the criticality or importance of an asset to your organization. As an example, you may choose to assign business-critical systems such as database and web servers a critical level, but normal enduser systems a medium level.

Assigning asset risk and criticality

Both asset risk and criticality can be assigned via the asset inventory. Asset criticality can also be assigned with alert rules.

Superusers, administrators, and users can add or modify asset risk and criticality levels, and can reset risk assignment or remove criticality assignment from assets.

Risk and criticality in the asset inventory

Follow these steps to set risk and criticality through the asset inventory:

1. Select all the assets you wish to update, applying a query filter if needed.
2. Click the `Modify asset risk` or `Modify asset criticality` button to open the relevant popup.
3. Select the level of risk or criticality you wish to apply to the asset(s).
4. Click `Override risk` or `Set criticality` to apply your changes.

Applying criticality with rules

To automatically apply asset criticality values to assets after a scan, [create an alert rule](#) by going to **Alerts > Rules** and clicking the **Create rule** button:

1. Select an inventory query you wish to use, such as the `asset-query-results` rule type, then click **Configure rule**.
2. Configure any desired settings.
3. Set the **Action** to `Modify asset`.
4. Select an option from the `Asset criticality` menu.
5. Save the rule.

This rule will now add the specified asset criticality level to all assets that match the rule when a scan completes.

Asset risk report

The **Asset risk report** provides visibility into the risk and criticality levels across your asset inventory. To run the Asset risk report, go to **Reports > Asset risk report** and click the **Asset risk report** button. Configure the following fields:

1. **Sites:** Select a site of assets to include in the report, or leave the default `All Sites`.
2. **Minimum risk:** Choose the minimum asset risk level to include.
3. **Minimum criticality:** Choose the minimum asset criticality to include.
4. **Top vulnerabilities per asset:** Set this field to an integer between 0 and 20. If the value is set to an integer between 1 and 20, the report will list up to that number of the top vulnerabilities detected on each asset. The top vulnerabilities are identified by sorting the vulnerability results for each asset by risk rank, then risk score, then severity rank, then severity score.
5. Click **Create report** to generate the results.

The resulting report is grouped by asset criticality level and then sorted by risk level. The results can be exported as JSON Lines (.jsonl), a JSON document (.json), or CSV (.csv).

Managing ownership

Platform

runZero is able to help users track ownership with the ability to configure different types of owners and assign owners to runZero assets and vulnerability records. Ownership coverage can also be tracked as a [goal](#).

Ownership types

Superusers can manage the available types of ownership on the [Account > Ownership types](#) page. Custom ownership types can be configured to meet your needs. Some common ownership types may include **Security owner**, **IT owner**, or **Compliance owner**.

The ownership type requires configuring three fields:

- **Name:** the name of the ownership type.
- **Reference:** whether the ownership type should be correlated with the user inventory, group inventory, or neither.
- **Visibility:** whether the ownership type is visible through the asset inventory and asset details pages.

The default Asset Owner ownership type, when visible, will be automatically populated with ownership-related data that runZero can glean from your configured integrations. The name of this ownership type can be changed by a superuser.

The list of ownership types can be prioritized by dragging the types into the preferred order. This will dictate the order in which the types are displayed in the inventory and asset details pages. Only types marked visible will be displayed.

Assigning owners to assets and vulnerabilities

Once created, custom owners can be assigned via the inventory or through an alert rule.

Superusers, administrators, and users can add or modify owner values, and can remove owners from assets or vulnerability records. Annotators can only add owner values, but cannot modify or remove owners.

Ownership in the inventory

Follow these steps to assign owners through the asset or vulnerability inventory:

1. Select all the assets or vulnerability records you wish to update, applying a query filter if needed.
2. Click the Manage asset ownership or Manage vulnerability ownership button to open the ownership popup. **Note:** Ownership values applied to an asset will be inherited by

unowned vulnerability records on that asset. Vulnerability records with owners defined will not inherit the ownership value assigned to the asset.

3. Click **Add ownership type** and choose which type(s) of owner you wish to apply to the selected assets or vulnerability records.
4. Add the owner value to the field.
5. Click **Save** to apply your changes.

Applying owners with rules

To automatically apply ownership values to assets after a scan, [create an alert rule](#) by going to **Alerts > Rules** and clicking the **Create rule** button:

1. Select an inventory query you wish to use, such as the `asset-query-results` rule type, then click **Configure rule**.
2. Configure any desired settings.
3. Set the **Action** to **Modify asset**.
4. Specify a value for the **Set [ownership type]** field for the ownership type(s) you wish to apply. **Note:** Ownership values applied to an asset will be inherited by unowned vulnerability records on that asset. Vulnerability records with owners defined will not inherit the ownership value assigned to the asset.
5. Save the rule.

This rule will now add the specified owner type and value to all assets that match the rule when a scan completes.

Managing tasks

You can view and manage discovery scans and other background actions from the [Tasks overview page](#). The *Active* and *Completed* task sections will show standard tasks, such as scans and imports, along with their current progress and summarized results. You can [search](#) or [filter](#) the tasks using different attributes.

Task status values

Tasks can have the following status values:

- **New:** The task has been created and is waiting to be picked up by a scanner or connector.
- **Active:** A scan task is in progress and the Explorer is scanning the network.
- **Scanned:** The network scan part of a scan task has completed.
- **Connecting:** A connector task is connecting to the remote system and downloading data.
- **Connected:** A connector task has finished downloading and the data is waiting to be processed.
- **Processing, queued:** Task data has been collected from a scan or a connector task, and is queued for processing.
- **Processing:** Task data is being processed.
- **Processed:** Task data has been processed and runZero data updated.
- **Stopping:** The task was requested to stop and is in the process of doing so.
- **Stopped:** The task successfully stopped.
- **Canceled:** An error occurred which meant that the task could not continue.
- **Paused:** A repeating task has been paused.

Tabs

Tabs on the Tasks overview page allow you to view a filtered subset of active, processing, scheduled, failed, completed, or recurring tasks. Each tab includes a search bar so that you can [search your complete task history](#).

- **Active** shows all tasks currently scanning or connecting
- **Processing** shows all tasks currently processing
- **Scheduled** shows all non-recurring tasks that are scheduled to run
- **Failed** includes all tasks that were unable to finish due to an error or cancellation
- **Completed** lists all completed tasks in chronological order based on when the task finished processing
- **Recurring** shows all recurring tasks

Task details

Clicking on a task will open a card at the top of the page with details. The concentric circles show the progress of the task, with the outer circle tracking the scan or connection progress

and the inner circle tracking the processing progress. The card also includes relevant timestamps and user creation data, as well as site, Explorer, and hosted zone settings. Completed tasks will also include the change summary, and recurring tasks will include a list of past runs.

Scheduled tasks

Scheduled tasks are one-off actions that will be started in the future, while recurring tasks are actions that happen on a regular basis. Recurring scan tasks generate a new standard task on each iteration of their schedule. Both scheduled and recurring scans will only launch if their associated Explorer is online and no other scan tasks are running.

The Tasks page allows *Scheduled* and *Recurring* tasks to be removed and *Active* tasks to be stopped. Please note that stopping an active *Scan* task may take a few moments, as the status of the task is not updated until the Explorer confirms that the scan was terminated.

System tasks

runZero has system tasks that you may notice while searching through your tasks. Here are some examples you might see:

- **Report-based:** a few reports will create tasks when initialized, including:
 - Organization overview report
 - Site and organization comparison
 - External assets report
- **Outlier calculation:** this generates the outlier score seen in the asset details page and will show up with the name `Outlier calculation`
- **Query:** all queries that are pinned to the Dashboard as Insights will show up with the name `Query`

Dismissing failed tasks

Tasks that were unable to finish due to an error or cancellation can be removed from the runZero Console by viewing the task details, clicking the “Dismiss” button and confirming the desired action. Dismissing failed tasks will archive the tasks so that they are no longer accessible.

Tracking goal progress

Professional Community Platform

With runZero goals, users are able to create and monitor progress toward achieving security initiatives. All goal types are supported by the robust query language on the backend. [All types of inventory queries](#) are supported by the goal tracking feature.

There are four types of goals:

- System query
- Custom query
- Asset ownership
- Asset risk

Goal creation

New goals can be created by users whose [default role](#) is administrator or greater from the [Goals](#) page in the console. Users with viewer-level or greater access will be able to view the goals page and see the goals that apply to organizations they have access to.

To create a new goal, click the [New goal](#) button on the [Goals](#) page in the console.

Creating system query or custom query goals

1. On the Name and permissions tab, complete the necessary fields:
 - **Type:**
 - **System query:** Select this type if you wish to use a runZero system query for your goal.
 - **Custom query:** Select this type if you wish to use a custom query for your goal.
 - **Name:** Provide a name for your goal.
 - **Permissions:** By default, goals are visible and applicable to all organizations in your account. To configure the goal on a per-organization basis, disable this toggle and enable individual organization toggles to determine which organizations the goal applies to.
2. On the Optional information tab, complete any desired fields:
 - **Description:** Provide a description for the goal.
 - **Notes:** Add relevant notes about this goal.
 - **Target date:** Provide a target date for achieving this goal.
 - **Pin to dashboard:** Enable this toggle if you wish to include this goal in the “Goals overview” widget on the dashboard. If enabled, the goal will be visible in the dashboard widget for all organizations that the goal applies to.
3. On the Target settings tab, configure the following:
 - **Query:** Select the query you want to use to track progress for this goal by checking. If needed, use the search field to find the query you wish to use.

- **Comparator:** Select whether goal progress will be calculated as Greater than or equal or Less than or equal to the configured threshold value.
 - **Threshold:** Select whether goal progress will be calculated as a percent or fixed number, and set the threshold value.
Note: When a goal applies to multiple organizations, progress is calculated per-organization and extrapolated across all applicable organizations.
4. On the Review and save tab, review the configurations of the goal. Click Save when you're ready, or make changes on previous tabs as needed.

Creating asset ownership goals

1. On the Name and permissions tab, complete the necessary fields:
 - **Type:** Select Asset ownership to measure progress toward assigning owners to assets in your inventory.
 - **Name:** Provide a name for your goal.
 - **Permissions:** By default, goals are visible and applicable to all organizations in your account. To configure the goal on a per-organization basis, disable this toggle and enable individual organization toggles to determine which organizations the goal applies to.
2. On the Optional information tab, complete any desired fields:
 - **Description:** Provide a description for the goal.
 - **Notes:** Add relevant notes about this goal.
 - **Target date:** Provide a target date for achieving this goal.
 - **Pin to dashboard:** Enable this toggle if you wish to include this goal in the "Goals overview" widget on the dashboard. If enabled, the goal will be visible in the dashboard widget for all organizations that the goal applies to.
3. On the Target settings tab, configure the following:
 - **Ownership type:** Select the ownership type you wish to track for this goal. New ownership types can be created by users with a default role of administrator or superusers from the Ownership page.
 - **Target coverage:** Set the target coverage for the selected ownership type.
Note: When a goal applies to multiple organizations, progress is calculated per-organization and extrapolated across all applicable organizations.
4. On the Review and save tab, review the configurations of the goal. Click Save when you're ready, or make changes on previous tabs as needed.

Creating asset risk goals

1. On the Name and permissions tab, complete the necessary fields:
 - **Type:** Select Asset risk to measure progress toward reducing the number of assets with a minimum risk level in your inventory.
 - **Name:** Provide a name for your goal.
 - **Permissions:** By default, goals are visible and applicable to all organizations in your account. To configure the goal on a per-organization basis, disable this toggle and enable individual organization toggles to determine which organizations the goal applies to.
2. On the Optional information

- tab, complete any desired fields:
- **Description:** Provide a description for the goal.
 - **Notes:** Add relevant notes about this goal.
 - **Target date:** Provide a target date for achieving this goal.
 - **Pin to dashboard:** Enable this toggle if you wish to include this goal in the “Goals overview” widget on the dashboard. If enabled, the goal will be visible in the dashboard widget for all organizations that the goal applies to.
3. On the **Target** settings tab, configure the following:
- **Minimum risk:** Select the minimum level of risk you wish to track for this goal.
 - **Minimum criticality:** Set the minimum criticality of the assets tracked by this goal (optional).
 - **Target percent:** Set the percent threshold you aim to be under. As an example, you might define a goal to have less than 25% of your critical assets with a risk level of medium or higher by setting **target percent** to 25%, **minimum criticality** to critical, and **minimum risk** to medium.
Note: When a goal applies to multiple organizations, progress is calculated per-organization and extrapolated across all applicable organizations.
4. On the **Review** and **save** tab, review the configurations of the goal. Click **Save** when you’re ready, or make changes on previous tabs as needed.

Goal progress calculation

Once a goal has been created, the progress will be calculated once a day and when tasks complete. As a result, you may see a different number of metric calculations in the historical line chart depending on the number of tasks that are completed on a given day.

The overall status noted in the Goal information card on a goal details page reflects the progress toward goal completion for all of the organizations that you have access to. This may differ for a user with more or less per-organization access.

Goal events and notifications

The events system includes the following events related to goals:

- **goal-created:** This event fires when a goal is created and saved.
- **goal-updated:** This event fires when a goal is modified and saved.
- **goal-completed:** This event fires when the progress of a goal crosses the completion threshold.
- **goal-lapsed:** This event fires when the progress for a previously completed goal falls below the completion threshold.
- **goal-removed:** This event fires when a goal is deleted.

These events can be used with [channels](#) and [rules](#) to receive notifications about a change in goal status. After a [notification channel](#) has been created, [create a rule](#) for the desired event and select the channel you wish to receive the notification through.

Understanding network segmentation

runZero multi-homed asset detection

Network segmentation is a critical security control for many businesses, but verifying that segmentation is working correctly can be challenging, especially across large and complex environments. Common techniques to validate segmentation, such as reviewing firewall rules and spot testing from individual systems can only go so far, and comprehensive testing, such as running full network scans from every segment to every segment, can be time intensive and are hard to justify on a regular basis.

For businesses subject to the PCI DSS requirements, validating cardholder data environment (CDE) segmentation is an important part of the security audit process. The [PCI guidance on scoping and segmentation](#) describes a common CDE administration model.

The network bridge detection in runZero is opportunistic and far from perfect, but it may highlight areas where segmentation is broken, and can cut down on the number of surprises encountered in a future security audit.

Using the bridge report

The bridge report shows external networks in red and internal networks in green. This view is not a typical network map, but instead shows possible paths that can be taken through the network by traversing multi-homed assets. Assets where runZero only detected a single IP address are not shown in order to keep the graph readable.

Zooming in will show asset and subnet details. Clicking a bridged node once will highlight the networks it is connected to, and clicking it a second time will either take you to the asset page. Clicking a network once will highlight the connections to bridged nodes, and clicking a second time will perform a CIDR-based inventory search.

Bridge detection is useful when validating network segmentation and ensuring that an attacker can't reach a sensitive network from an untrusted network or asset. Examples of this include laptops plugged into the internal corporate network that are also connected to a guest wireless segment and systems connected to an untrusted network, such as a coffee shop's wireless network that also have an active VPN connection to the corporate network.

runZero detects network bridges by looking for extra IP addresses in responses to common network probes (NetBIOS, SNMP, MDNS, UPnP, and others) and only reports bridges when there is at least one asset identified with multiple IP addresses. Typical hardening steps, such as desktop firewalls and disabled network services will usually prevent multi-homed assets from being detected by runZero. The screenshot below shows how to search for multi-homed assets in the runZero inventory.

Using the asset route pathing report

Platform

Network segmentation is a foundational security control that can be easily undermined by network misconfigurations and multi-homed machines. runZero Platform users can now visualize potential network paths between any two assets in an organization using the asset route pathing report.

This report generates a graph of multiple potential paths by analyzing IPv4 and IPv6 traceroute data in combination with subnet analysis of detected multi-homed assets—without requiring access to the hosts or network equipment. This unique methodology identifies surprising and unexpected paths between assets that may not be accounted for by existing security controls or reviews.

With a view of potential paths, security professionals can verify whether a low-trust asset, such as a machine on a wireless guest network, can reach a high-value target, such as a database server within a cardholder data environment (CDE). The new feature highlights potential network segmentation violations and opportunities for an attacker to move laterally from one segment to another.

Managing alerts

Professional Community Platform

runZero can trigger automatic alerts when certain events occur through a combination of Channels and Rules.

runZero currently supports Internal, Email, and Webhook channel types.

Internal channels store events within the [Alerts](#) list within the runZero Console. Internal alerts support explicit acknowledgement. Internal channels can be bulk acknowledged and cleared from within the runZero Console.

Email channels can be configured to deliver mail to one or more recipients. These email messages contain a summary of the alert and a link to the specifics within the runZero Console. Email is sent from the runZero infrastructure using the Sendgrid service.

Webhook channels allow runZero to post alerts to internet-reachable web services. The post request contains a standard text message for use with platforms like Slack and Mattermost, but also additional fields containing the full alert details. Webhooks are a great way to tie runZero alerts into third-party platforms.

To trigger an alert on a channel, a [Rule](#) must be created. Rules define which events lead to alert on which channels. The name of the rule will be included in the alert content and should describe the type of event that it monitors.

The following event types can be used to create rules:

- Scan completed
- New assets found
- Assets back online
- Assets now offline
- Assets changed

Scan completion and *assets changed* rules can be noisy but may be useful to keep a running log of network changes over time. For a typical monitoring use case, a rule would be created to trigger on *Assets now offline*, *Assets back online*, and *New assets found*, automatically alerting an email alias or a Slack channel.

Alert rules, when combined with recurring scans, can be a simple way to track network changes over time.

Using the rules engine

Professional Community Platform

The Rules Engine is an automation framework for monitoring, alerts, and workflow management. You can use the Rules Engine to customize alerts for the events that matter most to your organization and automate repetitive tasks. At the heart of the Rule Engine are

rules. A rule defines the action that is taken based on a set of conditions. You can create rules to proactively alert your team when there are changes to things like Explorers, assets, scans, organizations, and sites. You can also automate tagging and modification of asset fields based on the results of a query.

Some ways you can use the Rules Engine to help automate your workflow:

- Alert your team when new policy violations are identified.
- Modify asset fields when the assets match specific criteria.
- Bulk tag assets that match a specific query.
- Get a Slack notification when a query returns new results.
- Monitor when an Explorer goes offline in the runZero console.
- Know when there are changes to organizations, sites, and users.

Key concepts

Rules can help you stay on top of events as they happen and get better visibility across your network, assets, and your runZero deployment. To build a rule, you need to define three things: events, conditions, and actions. A rule determines that when a specific event happens, and certain conditions are met, the system will automatically perform the configured action.

Events

Each rule begins with an event. The event sets off the trigger and puts your rule into motion. An event can be based on a query or a system-defined event. runZero offers a library of system-defined events you can use to create your rules. Choosing any of these events will show the conditions and actions available.

Conditions

A condition narrows the scope of your rule. Unless the condition is met, the rule will not execute the action. You will only see conditions that apply for the event you have chosen. Generally, conditions specify sites, organizations, and asset attributes for the event.

Actions

An action executes your rule, if the event occurs, and the conditions meet all the criteria. An action can be a notification to a channel or it can be a modification to an asset. What you will need to configure depends on the action type. For notifications, you'll need to specify the notification channel and template. For asset modification, you can edit field like the OS vendor, OS product, OS version, hardware vendor, hardware product, hardware version, asset tags, and asset type.

Channels

A channel provides a way for you to communicate when a specific event has occurred. You can create multiple channels to support different types of communication needs. For example, you may want to create a Slack channel for one team, and an email list for another. It depends on what communication channels you prefer, and who you are trying to reach.

The body of the message uses default text from runZero. Customizations for messaging is currently unavailable.

Create a rule

Rules set the criteria for actions to take place. To create a rule, you need to choose an event, define the conditions, and choose a resulting action.

Step 1: Open the Rules Engine

- From the Alerts menu, select the Rules submenu.
- Click **Create Rule** to open the editor.

Step 2: Choose an event

- Choose an event you want to use as your trigger.
- You can browse the list of available predefined events. Use the search to quickly filter by keyword.
- Choosing 'asset-query-results' or 'service-query-results' will allow you to modify the fields for the resulting assets.
- After you've chosen an event, click **Configure rule**.

Step 3: Define the conditions

- Provide a descriptive name for the rule. Something that quickly tells you what the rule does.
- The conditions you can configure depend on the event you have selected.
- If you have an asset or service based query selected, you'll need to [provide a query](#) for the rule. This query will run against the site after the scan completes. Note that assets with data from non-runZero sources must be recent (seen in the last 30 days) to be included in the scope of the search, and runZero-scanned assets must be live.
- You will also need to set the scope to an organization and site, and sometimes, depending on the event, minimum asset counts.

Step 4: Choose an action

- Actions can execute a notification to the channel of your choice or modify assets. For example, you can choose to send notifications via email when orphaned devices are found.

Step 5: Turn on and save the rule

- Turn the rule on if you want to activate it immediately. Otherwise, you can save the rule and turn it on later.
- Save the rule when you're done.

Keep in mind

Using scan and asset event types can be noisy, but they are useful for tracking network changes over time. To help you focus on the events that matter most, track assets that go offline, assets that come back online, and newly discovered assets.

Monitoring the status of rules

The rules submenu of the Alerts page displays a list of all rules that have been created. For each rule, you can see:

- Whether the rule is enabled.
- The event that triggers processing of the rule.
- The organization the rule applies to, if the rule has been limited to a specific organization.
- When the rule was last triggered.
- Whether the rule resulted in an action being processed or not.

A status of “skipped” means that last time the rule was processed, its preconditions weren’t met, so no action was taken. A status of “processed” means that the rule’s preconditions were met, and its action was processed.

If there is an error processing a rule or sending a notification, the action status of the rule will be set to “error”. The error message can be seen as a tooltip on the error status.

Creating alert templates

Professional Community Platform

With the [Rules Engine](#), you can define rules that alert you on specific events, such as changes to scans, assets, and Explorers. To customize the alert messages, you can create custom templates to standardize and format alerts triggered from rules. With custom templates, you can include more context and data for your alerts.

Templates can output in HTML, JSON, and text for use in emails, internal notifications, or webhooks. You can customize the contents of these templates as needed.

Template building basics

You can define the contents of an alert message using the [Mustache](#) templating language. As long as you know a little bit about how the [Mustache syntax](#) works, you can build custom HTML and JSON templates to reference and pull in runZero data.

- For Slack notifications, you can use Slack’s [Block Kit](#) and their interactive [Block Kit Builder](#) to construct a rich message in JSON format, and then use Mustache to insert the relevant data.

- For Microsoft Teams, you can use the [Adaptive Cards](#) format to build rich messages, and again insert data via Mustache.

Our templates have two fields for template data: subject and body. Both subject and body can be customized using the Mustache syntax.

Inserting a data in a template

A standard set of objects is passed to the `{{template engine}}`, you just have to indicate the fields from the objects you want to insert into your template. Use the Mustache template syntax, `{{variable}}`, to include alert values when a rule matches certain conditions.

For HTML and JSON templates, values inserted using `{{ }}` are automatically escaped according to the appropriate rules. To avoid escaping a particular value, use triple curly braces, like this: `{{{task.name}}}`.

Inserting a value

To insert a value, put the variable name in double curly brackets, like this `{{variablename}}`.

The following example shows how to insert the console address:

The runZero Console is at `{{console}}`.

Inserting a value from an object

To insert a value from a specific object, separate the object name and field name with a dot, like this `{{object.fieldname}}`.

The following example shows how to insert the organization name:

The organization name is `{{organization.name}}`.

Inserting multiple values from an object

To insert multiple values from an object easily, use a section. You will need to start the section with `{{#objectname}}` and close it with a matching `{{/objectname}}` The following example shows how to insert the results from a scan that include total assets and number of assets changed:

Here are the results:

```
{{#scan}}
Scan found {{assets_total}} assets and changed {{assets_changed}} of them.
{{/scan}}
```

What happens if a field contains multiple values?

If a section refers to a field which contains multiple values, the template engine will loop through the values in the field, processing the section inside for each individual value.

The following example loops through all of the assets in `report.new`, and for each one, outputs its names and addresses fields. If there is nothing stored in `report.new`, the section between the tags will not be rendered.

```
{{#report.new}}
{{names}} {{addresses}}
{{/report.new}}
```

Note that within a loop, you can still refer to values from the outer object. If a named value isn't found in the current loop object, the template engine will check the outer object. For example, this can be useful for referring to the `{{console}}` variable, which provides the root runZero console URL.

Using boolean values

You can use boolean values with the `{{#field}}{{/field}}` tags. If the value of the field is false, the section between the tags is not rendered.

For example:

```
{{#query.truncated}}
(Additional results were found but not included in this report)
{{/query.truncated}}
```

Objects and fields reference

To include runZero data and details in your alerts, you can build your template using the following objects and fields.

globals

Field	Contents	Example
<code>console</code>	The base URL of the runZero web console.	https://console.runzero.com

event

The following fields are available in the event object:

Field	Contents	Example
<code>action</code>	The action which triggered the event	task-completed
<code>created_at</code>	When the event was created	2021-04-02 12:50:26 -0500
<code>id</code>	The UUID of the event	b4b871db-bdf1-4a42-b82d-18ae99972228

source_name	Name of the thing which caused the event	Weekly security scan
source_type	Type of thing which caused the event	task
success	Whether the event succeeded	true
target_name	Name of the object targeted by the event	Head Office
target_type	The type of the object targeted by the event	site

task

For events triggered by a task, the following fields are available in the task object:

Field	Contents	Example
created_at	When the task was created	2021-04-02 12:50:26 -0500
created_by	The user who created the task	user@example.com
description	The description of the task	Weekly scan of main network
error	The text of any error message for the task	explorer unavailable after 4h
id	The UUID of the task	b4b871db-bdf1-4a42-b82d-18ae99972228
name	The name of the task	Weekly Scan
start_time	When the task started	2021-04-02 12:50:26 -0500
status	The status of the task	processed
type	The type of task	scan
updated_at	When the task was last updated	2021-04-02 12:50:26 -0500
url	A URL linking to the task details	

organization

The following fields are available in the organization object:

Field	Contents	Example
id	The UUID of the organization	86f12ee1-f0f1-419a-8799-63fff555777a
name	The name of the organization	IT Dept.

site

The following fields are available in the site object:

Field	Contents	Example
id	The UUID of the site	49f9323a-fea1-4afc-b490-2414c3aaaaeee
name	The name of the site	Head Office

rule

The following fields are available in the rule object:

Field	Contents	Example
action	The action the rule said to take	notify
created_at	When the rule was created	2021-03-08 12:43:59 -0600
created_by	The user who created the rule	user@example.com
event	The event triggering the rule	scan-completed
id	The UUID of the rule	b2269a2c-69a1-4652-bcdf-899938886c17
name	The name of the Rule	Alert on scan
updated_at	When the rule was last updated	2021-04-01 17:09:40 -0500

scan

For events triggered by a scan task, the following additional fields are available in a scan object:

Field	Contents	Example
explorer_id	The UUID of the runZero Explorer which carried out the scan	0fd44a62-d827-41c0-b26c-4837222d8888
assets_changed	The number of assets changed as a result of the scan	9
assets_ignored	The number of assets ignored by the scan	2
assets_new	The number of new assets detected by the scan	2
assets_offline	The number of assets which were previously online but now offline	1

assets_online	The number of assets previously offline but now online again	2
assets_total	The total number of assets for the site scanned (including offline)	11
assets_unchanged	The number of assets unchanged by the scan	1
assets_updated	The total number of assets up-to-date as a result of the scan (changed + unchanged)	15
duration	Duration of the scan in seconds	26
end_time	When the scan ended	2021-04-02 12:50:26-0500
excludes	Any IP addresses excluded from the scan	10.0.1.123
id	The UUID of the scan task	894a112c-3fb9-4301-8da7-8ce7fffb4443
name	The name of the scan task	Weekly security scan
rate	The scan rate	1000
recv_bytes	How many bytes were received during the scan	45176
recv_error	How many receive errors were detected	1
recv_packets	How many data packets were received during the scan	555
scheduled_time	When the scan was scheduled to run	2021-04-02 12:48:00-0500
sent_bytes	How many bytes were sent during the scan	44740
sent_error	How many send errors were detected	0
sent_packets	How many data packets were sent during the scan	577
start_time	When the scan actually started	2021-04-02 12:49:55-0500
tags	An array of tags associated with the scan task	
targets	The CIDR address ranges scanned	10.0.1.0/24
type	The type of operation	scan

explorer

For events triggered by a scan task, the following additional fields are available in an explorer object (runZero 2.1.8+):

Field	Contents	Example
id	The UUID of the runZero Explorer which carried out the scan	0fd44a62-d827-41c0-b26c-4837222d8888
name	The name of the Explorer which carried out the scan	MM34B-2
internal_ip	The internal IP address of the Explorer which carried out the scan	10.0.1.200

search

For events triggered by a search query rule, the following additional fields are available in a search object:

Field	Contents	Example
url	A link to perform the same search	https://console.runzero.com/inventory/?search=...
found	The number of matches found	3
comparator	The operation used to compare the number of matches	>=
value	The value the number of matches was compared against	1

report

For “scan completed” events, a report object contains the following results from the scan:

Field	Contents
truncated	Whether the set of objects was truncated due to large numbers of assets
changed	An array of changed assets (see below)
new	An array of new assets (see below)
offline	An array of assets now offline (see below)
online	An array of assets now online (see below)

The limit on the number of objects passed to the template is 25 for email notifications, 10 for Webhook notifications.

asset (from a scan report)

Each asset returned as part of a scan report has the following fields:

Field	Contents	
addresses	The IP address(es) of the asset	10.0.1.123
alive	Whether the asset responded to probes	true
created_at	The timestamp when the asset record was created	2021-04-02 12:50:26 -0500
detected_by	The method by which the asset was detected	arp
domains	Any domains the asset was found in	WORKGROUP
first_seen	The timestamp when the asset was first seen	2021-04-02 12:50:26 -0500
hw	A summary of the asset hardware	HP LaserJet Pro
id	The UUID of the asset	b38295fb-bef1-fa42-b82d-18ae99972228
last_seen	The timestamp when the asset was most recently seen	2021-04-02 12:50:26 -0500
names	Any names detected for the asset	LP538N
os	A summary of the asset's operating system	Linux
service_count	How many services the asset is running	4
type	The type of asset	Router
updated_at	The timestamp when the asset record was last updated	2021-04-02 12:50:26 -0500

query

For query events, the following data is provided in the query variable:

Field	Contents
count	The number of rows matching the search query
assets	If it was an asset search, the array of assets matching the query

services	If it was a service search, the array of services matching the query
wlans	If it was a wireless network search, the array of wireless networks matching the query
truncated	Whether the set of assets was truncated due to the query returning a large number

The limit on the number of assets passed to the template from a query is 25 for email notifications, 10 for Webhook notifications.

asset (from a query)

For assets returned from a query rule, the following fields are available:

Field	Contents	
addresses	The IP address(es) of the asset	10.0.1.123
alive	Whether the asset responded to probes	true
comments	Any comments set in the asset record	
created_at	The timestamp when the asset record was created	2021-04-02 12:50:26 -0500
detected_by	The method by which the asset was detected	arp
domains	Any domains the asset was found in	WORKGROUP
first_seen	The timestamp when the asset was first seen	2021-04-02 12:50:26 -0500
hw	A summary of the asset hardware	ThinkPad X1
id	The UUID of the asset	b38295fb-bef1-fa42-b82d-18ae99972228
last_seen	The timestamp when the asset was most recently seen	2021-04-02 12:50:26 -0500
macs	The list of MAC addresses associated with the asset	[F4:F5:E8:89:92:31, 00:D0:2D:9F:47:77]
names	Any names detected for the asset	laptop.local
organization	The organization the asset belongs to	IT
os	A summary of the asset's operating system	Windows 10
service_count	How many services the asset is running	4

tags	An array of tags set on the asset	
type	The type of asset	Thermostat
updated_at	The timestamp when the asset record was last updated	2021-04-02 12:50:26 -0500

service (from a query)

For services returned from a query rule, the following fields are available, some of which are taken from the associated asset:

Field	Contents	Example
id	The UUID of the service (not the asset)	b38efadb-61f1-f332-b92d-18ae99972228
created_at	When the service record was created	2021-04-02 12:50:26 -0500
summary	A summary of the service	ciscoSystems
port	The TCP/UDP port the service is on	53
vhost	The vhost of the service	ftp.example.com
address	The TCP/IP address of the service	192.168.33.44
transport	The transport	udp
protocol	The name of the protocol, if known	ssh
organization	The name of the organization the service's asset belongs to	HR
site	The name of the site the service's asset belongs to	Lab
alive	Whether the asset offering the service was alive	true
last_seen	When the asset was last seen	2021-04-02 12:50:26 -0500
first_seen	When the asset was first seen	2021-02-11 09:38:17 -0500
type	The asset type offering the service	Laptop
os	The OS offering the service	Linux
hw	The hardware offering the service	APC UPS
addresses	A list of other IP addresses associated with the asset offering the service	[192.168.0.2, 192.168.0.3]
macs	The list of MAC addresses associated with the asset	[F4:F5:E8:89:92:31, 00:D0:2D:9F:47:77]

names	Any names associated with the asset	[fw-3, fw-3a]
tags	Tags set on the asset	
domains	Any domain associated with the asset	LOCAL
service_count	A count of how many services the asset offers	4
comments	Any comments set on the asset	

wlan (wireless LAN, from a query)

For wireless lans (wlans) returned from a query rule, the following fields are available:

Field	Contents	Example
id	The UUID of the wireless LAN in runZero's database	f938934b-ae23-f112-b23d-18ae99972228
last_seen	When the wlan was last seen	2021-04-02 12:50:26 -0500
ssid	The ESSID of the network	Free WiFi
bssid	The ESSID of the network	c4:41:1e:99:88:77
type	The type of wireless network	infrastructure
authentication	The authentication used to access the network	WPA2-PSK
encryption	The encryption used to protect data	AES
signal	The signal strength as a percentage	86
channels	The channel of the network	11
organization	The name of the organization the service's asset belongs to	HR
site	The name of the site the service's asset belongs to	Lab

Example: Alert when scan completes

Let's take an example of something you might want to get alerted on: completed scans. We can create a rule that emails us when a scan completes, and provides us with some details, such as the number of new, online, offline, and modified assets. We'll build these details into our template.

Step 1. Create a template.

You can create a template from the [Alerts](#) page.

- Name the template something like Email the team when a scan completes.
- For the template type, choose HTML, since we want to use this template for emails.
- For the subject line, enter something that's descriptive, like runZero scan `{{scan.name}}` completed at `{{scan.end_time}}`. You can use the Mustache syntax for the subject.

Step 2. Create the body message

Now, let's create the email body. We want the email to tell us how many new, online, offline, and modified assets there are, as well as give us details on the new assets discovered.

The body looks like this:

```
<h1>{{site.name}}</h1>

<h2>Scan Results</h2>
{{#scan}}
<ul>
<li>{{assets_new}} new assets</li>
<li>{{assets_online}} online assets</li>
<li>{{assets_offline}} offline assets</li>
<li>{{assets_changed}} modified assets</li>
</ul>
{{/scan}}

<h2>New assets</h2>
<ul>
{{#report.new}}
<li>{{names}} {{addresses}} {{os}}</li>
{{/report.new}}
{{^report.new}}
<li>No new assets were discovered.</li>
{{/report.new}}
</ul>

<p><a href="{{task.url}}">View the scan results</a></p>
```

See that caret (^)? It represents inverted sections. These sections only render when the list is empty or the value is empty or false.

Step 3. Save the template and create a rule

The template is now available for you to choose when you [create a rule](#).

Sample JSON templates

All available attributes

```
{
  "event": {
    "action": "{{event.action}}",
    "created_at": "{{event.created_at}}",
    "id": "{{event.id}}",
    "source_name": "{{event.source_name}}",
    "source_type": "{{event.source_type}}",
    "success": "{{event.success}}",
    "target_name": "{{event.target_name}}",
    "target_type": "{{event.target_type}}"
  },
  "organization": {
    "name": "{{organization.name}}",
    "id": "{{organization.id}}"
  },
  "site": {
    "name": "{{site.name}}",
    "id": "{{site.id}}"
  },
  "report": {
    "truncated": "{{report.truncated}}",
    "changed": "{{report.changed}}",
    "new": "{{report.new}}",
    "offline": "{{report.offline}}",
    "online": "{{report.online}}"
  },
  "rule": {
    "action": "{{rule.action}}",
    "created_at": "{{rule.created_at}}",
    "created_by": "{{rule.created_by}}",
    "event": "{{rule.event}}",
    "id": "{{rule.id}}",
    "name": "{{rule.name}}",
    "updated_at": "{{rule.updated_at}}"
  },
  "scan": {
    "explorer_id": "{{scan.explorer_id}}",
    "assets_changed": "{{scan.assets_changed}}",

```

```

"assets_new": "{{scan.assets_new}}",
  "assets_offline": "{{assets_offline}}",
  "assets_online": "{{scan.assets_online}}",
  "assets_total": "{{scan.assets_total}}",
  "assets_unchanged": "{{scan.assets_unchanged}}",
  "assets_updated": "{{scan.assets_updated}}",
  "duration": "{{scan.duration}}",
  "end_time": "{{scan.end_time}}",
  "excludes": "{{scan.excludes}}",
  "id": "{{scan.id}}",
  "name": "{{scan.name}}",
  "rate": "{{scan.rate}}",
  "recv_bytes": "{{scan.recv_bytes}}",
  "recv_error": "{{scan.recv_error}}",
  "recv_packets": "{{scan.recv_packets}}",
  "scheduled_time": "{{scan.scheduled_time}}",
  "sent_bytes": "{{scan.sent_bytes}}",
  "sent_error": "{{scan.sent_error}}",
  "sent_packets": "{{scan.sent_packets}}",
  "start_time": "{{scan.start_time}}",
  "tags": "{{scan.tags}}",
  "targets": "{{scan.targets}}",
  "type": "{{scan.type}}",
},
"explorer": {
  "id": "{{explorer.id}}",
  "name": "{{explorer.name}}",
  "internal_ip": "{{explorer.internal_ip}}",
},
"search": {
  "url": "{{search.url}}",
  "found": "{{search.found}}",
  "comparator": "{{search.comparator}}",
  "value": "{{search.value}}",
},
"query": {
  "count": "{{query.count}}",
  "assets": "{{query.assets}}",
  "services": "{{query.services}}",
  "wlans": "{{query.wlans}}",
  "truncated": "{{query.truncated}}",
}
}

```

Event attributes only

```
{
  "action": "{{event.action}}",
  "created_at": "{{event.created_at}}",
  "id": "{{event.id}}",
  "source_name": "{{event.source_name}}",
  "source_type": "{{event.source_type}}",
  "success": "{{event.success}}",
  "target_name": "{{event.target_name}}",
  "target_type": "{{event.target_type}}",
}
```

Organization attributes only

```
{
  "name": "{{organization.action}}",
  "id": "{{organization.id}}",
}
```

Site attributes only

```
{
  "name": "{{site.action}}",
  "id": "{{site.id}}",
}
```

Report attributes only

```
{
  "truncated": "{{report.truncated}}",
  "changed": "{{report.changed}}",
  "new": "{{report.new}}",
  "offline": "{{report.offline}}",
  "online": "{{report.online}}",
}
```

Rule attributes only

```
{
  "action": "{{rule.action}}",
  "created_at": "{{rule.created_at}}",
  "created_by": "{{rule.created_by}}",
  "event": "{{rule.event}}",
}
```

```
"name": "{{rule.name}}",
  "updated_at": "{{rule.updated_at}}"
}
```

Scan attributes only

```
{
  "explorer_id": "{{scan.explorer_id}}",
  "assets_changed": "{{scan.assets_changed}}",
  "assets_ignored": "{{scan.assets_ignored}}",
  "assets_new": "{{scan.assets_new}}",
  "assets_offline": "{{assets_offline}}",
  "assets_online": "{{scan.assets_online}}",
  "assets_total": "{{scan.assets_total}}",
  "assets_unchanged": "{{scan.assets_unchanged}}",
  "assets_updated": "{{scan.assets_updated}}",
  "duration": "{{scan.duration}}",
  "end_time": "{{scan.end_time}}",
  "excludes": "{{scan.excludes}}",
  "id": "{{scan.id}}",
  "name": "{{scan.name}}",
  "rate": "{{scan.rate}}",
  "recv_bytes": "{{scan.recv_bytes}}",
  "recv_error": "{{scan.recv_error}}",
  "recv_packets": "{{scan.recv_packets}}",
  "scheduled_time": "{{scan.scheduled_time}}",
  "sent_bytes": "{{scan.sent_bytes}}",
  "sent_error": "{{scan.sent_error}}",
  "sent_packets": "{{scan.sent_packets}}",
  "start_time": "{{scan.start_time}}",
  "tags": "{{scan.tags}}",
  "targets": "{{scan.targets}}",
  "type": "{{scan.type}}"
}
```

Explorer attributes only

```
{
  "id": "{{explorer.id}}",
  "name": "{{explorer.name}}",
  "internal_ip": "{{explorer.internal_ip}}"
}
```

Search attributes only

```
{
  "url": "{{search.url}}",
  "found": "{{search.found}}",
  "comparator": "{{search.comparator}}",
  "value": "{{search.value}}"
}
```

Query attributes only

```
{
  "count": "{{query.count}}",
  "assets": "{{query.assets}}",
  "services": "{{query.services}}",
  "wlans": "{{query.wlans}}",
  "truncated": "{{query.truncated}}"
}
```

Asset, service, or wireless query alert to SIEM or SOAR

```
{
  "organization": {
    "name": "{{organization.name}}",
    "id": "{{organization.id}}"
  },
  "site": {
    "name": "{{site.name}}",
    "id": "{{site.id}}"
  },
  "rule": {
    "action": "{{rule.action}}",
    "created_at": "{{rule.created_at}}",
    "created_by": "{{rule.created_by}}",
    "event": "{{rule.event}}",
    "id": "{{rule.id}}",
    "name": "{{rule.name}}",
    "updated_at": "{{rule.updated_at}}"
  },
  "search": {
    "url": "{{search.url}}",
    "found": "{{search.found}}",
    "comparator": "{{search.comparator}}",

```

```
}
}
```

Asset, service, or wireless query alert to Slack

```
{
  "blocks": [
    {
      "type": "section",
      "text": {
        "type": "mrkdn",
        "text": ":\red_circle: *runZero Alert* - {{rule.name}}"
      }
    },
    {
      "type": "divider"
    },
    {
      "type": "section",
      "text": {
        "type": "mrkdn",
        "text": "*Rule information*\n\n_Name_: {{rule.name}}\n_Type_: {{rule.event}}\n_Link_: https://console.runzero.com/alerts/rule/{{rule.id}}"
      }
    },
    {
      "type": "section",
      "text": {
        "type": "mrkdn",
        "text": "*Match information*\n\n_Organization_: {{organization.name}}\n_Site_: {{site.name}}\n_Match count_: {{search.found}}\n_Search_: {{search.value}}\n_Link_: https://console.runzero.com/alerts/rule/{{search.url}}"
      }
    }
  ]
}
```

Asset, service, or wireless query alert to Microsoft Teams

```
{
  "$schema": "http://adaptivecards.io/schemas/adaptive-card.json",
  "type": "AdaptiveCard",
  "version": "1.5",
  "body": [
```

```

"type": "TextBlock",
  "text": "runZero Alert - {{rule.name}}"
},
{
  "type": "TextBlock",
  "text": "***Rule information**\n\n_Name_: {{rule.name}}\n\n_Type_: {{rule.event}}\n\n_Link_: [Rule](https://console.runzero.com/alerts/rule/{{rule.id}})",
  "wrap": true
},
{
  "type": "TextBlock",
  "text": "***Match information**\n\n_Organization_: {{organization.name}}\n\n_Site_: {{site.name}}\n\n_Match count_: {{search.found}}\n\n_Search_: {{search.value}}\n\n_Link_: [Search](https://console.runzero.com/alerts/rule/{{search.url}})",
  "wrap": true,
  "spacing": "Medium"
}
]
}

```

Data type accepted by each channel

When you create a rule, the channel you select must accept the data type used by the template. For example, if you want to send a Slack notification, the template must be in plain text or JSON.

Here are the data types accepted by each channel:

Channel	Data type
Email	Plain text, HTML, JSON
Webhook	Plain text, JSON
Internal notification	Plain text

Managing templates

Go to [Alerts > Templates](#). From this page, you can create, edit, and delete templates. Before deleting templates, make sure they are not in use by any rules. If you delete a template that is in use by a rule, the rule will revert to the default template.

Querying your data

runZero provides many ways to query your data. Generally, queries can be broken into two concepts:

- Filters or parameters used in the search bars on pages across the console, or
- System and custom queries for which match metrics are calculated as tasks complete.

Both allow you to leverage the extensive query language to quickly find the information you're looking for.

Filtering and searching data

The various inventory pages are likely the main place you'd look to use these queries, but many other pages include the same type of search bar that can be used to filter results. The following documentation pages will help you craft a query that meets your needs:

- [Syntax for crafting compound queries](#)
- [Examples to guide query development](#)
- [Keywords that can be used across the inventory pages](#)
- [Keywords that can be used across other components of the console](#)

System and custom queries

The [Queries page](#) serves as an inventory of all your saved queries. Saved queries can either be system queries published by runZero or custom queries created by you or your team.

Standard query attributes

Queries have the following standard attributes:

- **Name:** A name for the saved query.
- **Description:** A description of the saved query.
- **Search query:** The query or parameters to search for matches to.
- **Search live assets only:** When toggled to Yes (default), `alive:t` is added to the search query to only include assets marked alive.
- **Type:** The inventory the query will search.
- **Category:** The category the query falls under.
- **Severity:** A severity level for the query.
- **Automatically track query results on the dashboard:** When toggled to Yes, the query and count of matches will be included in a component on the dashboard.

Vulnerability record attributes

Professional Community Platform

System and custom queries can also be used to create and associate vulnerability records with matching assets. This will be enabled by default on some system queries. To have a query create vulnerability records, switch the **Apply a vulnerability record to matching assets** toggle to Yes, then complete the following fields:

- **Vulnerability ID:** Choose a unique ID to track this vulnerability within runZero.
- **CVEs:** Include a list of CVEs relevant to this vulnerability record (optional).
- **Solution:** Provide context for how this vulnerability could be remediated on assets (optional).
- **Risk:** Select a risk level to associate with the vulnerability. This impacts the [asset risk](#).
- **Exploitable:** Specify whether an exploit is available for the vulnerability.
- **CVSS v3 base score:** The CVSS v3 base score (0.00 to 10.00).
- **CVSS v3 temporal score:** The CVSS v3 temporal score (0.00 to 10.00).
- **CVSS v2 base score:** The CVSS v2 base score (0.00 to 10.00).
- **CVSS v2 temporal score:** The CVSS v2 temporal score (0.00 to 10.00).
- **CPE 2.3 identifier:** Specify a relevant Common Platform Enumeration identifier in URI format (v2.3) to associate with the reported vulnerability (optional).

Creating and editing queries

Custom queries can be created by users with the default role of administrator or higher from the [Queries page](#) by clicking the **New query** button. Once created, custom queries can be edited and copied. System queries cannot be edited directly, but can be copied if you wish to make changes.

Search query syntax

runZero supports a deep searching across the Asset, Service, and Wireless Inventory, across organizations and sites, and through the Query Library. The runZero Export API uses the same inventory search syntax to filter results.

Query syntax

Boolean operators

Search queries can be combined through AND and OR operators and be grouped using parenthesis.

AND

For example, a Asset Inventory query of `os:"Windows 10" AND protocols:http AND protocols:smb2` will show only those assets where Windows 10 was identified and both SMB and a web server were discovered. Search values that contain spaces must be placed in double quotes.

OR

By contrast, the example query of `os:"Windows 10" AND protocols:http OR protocols:smb2` will search for Windows 10 running a web server or any assets with the SMB service exposed. In addition to AND and OR, the NOT operator can be used to filter a query. For example, the query `os:"Windows 10" AND NOT protocols:http` will show Windows 10 systems without a web server. If the negation should happen as the first term the AND should be dropped. The query `NOT protocol:http AND os:"Windows 10"` is equivalent to the previous search, with the terms reversed.

Wildcard and fuzzy searches

Most keywords are a fuzzy match by default. To force an exact match, prefix match, or suffix match, the `=` prefix can be applied to the search term, with the `%` character used as a wildcard. To search an operating system name of just Windows, the Asset Inventory query would be `os:="Windows"`, while to specify a prefix match of Ubuntu Linux, the query `os:="Ubuntu Linux%"` can be used.

Time and date values

Time and date (timestamp) fields can be searched using `<` (less than) and `>` (greater than) operators to compare against the current time. You can also use `-` to compare to a relative time in the past, for example `<-3years` would look for timestamps that occurred before three years ago. Supported units:

- hours
- minutes
- seconds
- months
- years

A special value of `now` can also be used.

For example, an asset search of `first_seen:<1year` would search for assets first detected this year. Other examples:

```
first_seen:<3days
```

```
first_seen:>2019-08-01
```

```
first_seen:>8/1/2019
```

```
last_seen:<1week
```

```
last_seen:<2months
```

```
last_seen:<1year
```

`created_at:>2weeks`

`created_at:<30minutes`

`updated_at:>1year`

`updated_at:<12hours`

`os_eol:<now`

`os_eol:>4weeks`

`os_eol:<-2years`

`os_eol_extended:>now`

`os_eol_extended:>90days`

Empty values

To search for an empty value, the `=` prefix can be used with no value after. For example, the query `os:=` will find assets with no identified operating system.

Note that this only works for single-valued attributes such as `os` and `type`; it won't work for multi-value attributes such as `names` or `addresses`.

Asset and service inventory searches

Asset and Service attributes support two special search types in addition to the documented keywords:

- Asset Inventory searches treat unknown keywords as filters against individual Asset **attributes**.
- Service Inventory searches treat unknown keywords as filters against individual Service **data** values.

In situations where an Asset keyword conflicts with a Service data key, or an Asset attribute conflicts with a Service keyword, the prefixes `_asset.` and `_service.` can be used to disambiguate.

Searches are handled slightly differently. Service queries can filter against Asset attributes (`os:linux`) and Service attributes (`banner:Password`), but the Asset queries are limited to summary information about services (`protocol:ssh`).

Query examples

There are endless ways to combine terms and operators into effective queries, and the examples below can be used as-is or adjusted to meet your needs.

Network configurations and access

- Multihomed assets with public and private IP addresses:

`alive:t AND has_public:t AND has_private:t`

- Multihomed assets connected only to private networks

`multi_home:t AND has_public:f`

- Default SSH configuration using passwords for authentication:

`alive:t AND protocol:"ssh" AND ssh.authMethods:"=password"`

- Microsoft FTP servers:

`alive:t AND protocol:"ftp" AND banner:"=%Microsoft FTP%"`

- Remote access services/protocols:

`protocol:rdp OR protocol:vnc OR protocol:teamviewer`

- Assets with public IPs running remote access services:

`has_public:t OR has_public:t AND alive:t AND (protocol:rdp OR protocol:vnc OR protocol:teamviewer)`

- Open ports associated with cleartext protocols:

`port:21 OR port:23 OR port:80 OR port:443 OR port:139 OR port:445 OR port:3306 OR port:1433 OR`

- Telnet on nondefault ports:

`protocol:telnet AND NOT port:23`

- Windows assets offering SMB services:

`os:windows AND protocol:smb1 OR protocol:smb2`

- Switch assets accepting Username and Password authentication:

`type:switch AND (_asset.protocol:http AND NOT _asset.protocol:tls) AND (html.inputs:"password:"`

- Assets more than 8 hops away:

`attribute:"ip.ttl.hops" AND ip.ttl.hops:>"8`

Asset lifecycle and hardware

- Assets created as a result of arbitrary responses:

has_mac:f AND has_name:f AND os:= AND hardware:= AND detected_by:icmp AND service_count:<2

- End of Life assets:

os_eol:<now

- Assets where both OS support and extended support are expired:

os_eol:<now AND os_eol_extended:<now

- Assets where OS support is EOL but still covered by extended support:

os_eol:<now AND os_eol_extended:>now

- EOL Linux operating systems:

os:linux AND os_eol:<now

- EOL Windows operating systems:

os:windows AND os_eol:<now

- Assets discovered within the past two weeks:

first_seen:"<2weeks"

- All available serial number sources

protocol:snmp has:snmp.serialNumbers OR hw.serialNumber:t OR ilo.serialNumber:t

- Asset serial numbers from SNMP:

protocol:snmp has:snmp.serialNumbers

- Older Windows OSes:

os:"Windows Server 2012" OR os:"Windows 7"

- Older Linux OSes:

OS:linux AND os_eol:<now

- BACnet devices:

type:bacnet

- Hikvision DVRs:

type:dvr AND os:hikvision

- IoT Devices:

```
type:"IP Camera" OR type:"thermostat" OR type:"Amazon Device" OR hw:"Google Chromecast" OR type:
```

- Video-related assets:

```
type:"IP Camera" OR type:"DVR" OR type:"Video Encoder"
```

Misconfigurations

- SMBv1:

```
protocol:"smb1"
```

- Remote access with common services:

```
protocol:rdp OR protocol:vnc OR protocol:teamviewer OR protocol:spice OR protocol:pca
```

- Switches with default configurations for web access:

```
type:switch AND (_asset.protocol:http AND NOT _asset.protocol:tls) AND ( html.inputs:"password:
```

- Default SSH configurations using passwords for authentication:

```
alive:t AND protocol:"ssh" AND ssh.authMethods:"=password"
```

- Switches using Telnet or HTTP for remote access:

```
type:switch AND protocol:telnet OR protocol:http
```

- Microsoft FTP servers:

```
alive:t AND protocol:"ftp" AND banner:"=%Microsoft FTP%"
```

- Virtual machines that are not syncing time with the host:

```
@vmware.vm.config.tools.syncTimeWithHost:"False"
```

Weak configurations

- Telnet (vs. SSH):

```
protocol:telnet
```

- FTP on ports 10-21 (vs. FTPS on port 990):

```
protocol:ftp
```

- FTP on ports 20-21 (vs. SCP on port 22):

protocol:ftp

- HTTP on port 80 (vs. HTTPS on port 443):

protocol:http

- SSH versions < 2.0:

protocol:ssh AND NOT banner:"SSH-2.0"

- TLS:

tls.versionName:"=TLSv1.3" OR tls.versionName:"=TLSv1.2" OR tls.versionName:"=TLSv1.1" OR tls.v

- LDAP on port 389 (vs. LDAPS on port 636):

protocol:ldap OR port:389

- Wireless access points without WPA authentication:

not authentication:WPA

- Online assets with SSH accepting password authentication:

alive:t AND has:"ssh.authMethods" AND protocol:"ssh" AND (ssh.authMethods:"=password" OR ssh.au

- Detect OpenSSL version 3.0 - 3.0.6:

product:openssl AND version:3.0

EDR / MDM

- CrowdStrike coverage gaps:

not edr.name:crowdstrike AND (type:server OR type:desktop OR type:laptop)

- Assets with CrowdStrike Agent status "Not Provisioned":

@crowdstrike.dev.provisionStatus:"NotProvisioned"

- Assets with CrowdStrike Agent mode "Reduced Functionality":

@crowdstrike.dev.reducedFunctionalityMode:"yes"

- Assets with CrowdStrike Agent status "Normal":

@crowdstrike.dev.status:"normal"

- SentinelOne coverage gaps:

not edr.name:Sentinelone AND (type:server OR type:desktop OR type:laptop)

- Assets with SentinelOne Agent requiring patch:

(alive:t OR scanned:f) AND has:"@sentinelone.dev.appsVulnerabilityStatus" AND @sentinelone.dev.

- Assets missing either CrowdStrike or SentinelOne EDR agents:

NOT edr.name:crowdstrike AND (type:server OR type:desktop OR type:laptop) OR NOT edr.name:senti

- Miradore coverage gaps:

not source:Miradore AND (os:google android OR os:apple ios) AND type:mobile

- Microsoft Defender coverage gaps:

not edr.name:"Defender" AND os:Windows

- Assets not managed by a Microsoft product:

source:runzero AND NOT (source:ms365defender OR source:intune OR source:azuread)

- Find mobile devices on the network:

(os:google ANDroid OR os:apple ios) AND type:mobile

- Known FCC security threats, like Kaspersky:

alive:t AND edr.name:Kaspersky

Virtual machine configurations

- Virtual machines with less than 8 GB of memory:

@vmware.vm.config.hardware.memoryMB:<"8192"

- VMs with less than 16GB of memory:

@vmware.vm.runtime.maxMemoryUsage:"16384"

- Virtual machines that are not syncing time with the host:

@vmware.vm.config.tools.syncTimeWithHost:"False"

- Virtual machines that are configured with floppy drives:

@vmware.vm.config.extra.floppy0.autodetect:"true"

- Virtual machines running VMware tools:

```
@vmware.vm.config.extra.guestinfo.vmtools.versionString:"_"
```

- Virtual machines running Windows:

```
source:VMware AND os:Windows
```

- Virtual machines running Linux:

```
source:VMware AND os:Linux
```

Vulnerability concerns

- Rapid7 - fails PCI compliance:

```
test.pciComplianceStatus:"fail"
```

- Tenable - High and Critical severity vulnerabilities that are on CISA's Known Exploited list:

```
plugin.xrefs.type:"CISA-KNOWN-EXPLOITED" AND (severity:high OR severity:critical)
```

- Tenable - Critical severity vulnerabilities where exploits are available:

```
plugin.exploitabilityEase:"Exploits are available" AND severity:critical
```

- Tenable - High and Critical severity vulnerabilities where exploits are not required

```
plugin.exploitabilityEase:"No exploit is required" AND (severity:critical OR severity:high)
```

Wireless results

- Search ESSID for authentication exceptions:

```
ssid:"<ESSID>" AND NOT authentication:"wpa2-enterprise"
```

- Find unknown BSSIDs broadcasting known ESSID (exclude known BSSIDs in query for gap analysis)

```
ssid:="<ESSID>" AND NOT bssid:"<MAC address>"
```

Inventory keywords

The data across your runZero inventories can be queried and filtered using the [search syntax](#) in conjunction with the available inventory keywords. Keywords and example values are documented for the following inventories:

- [Assets](#)
- [Services](#)

- [Software](#)
- [Vulnerabilities](#)
- [Wireless](#)
- [Users](#)
- [Groups](#)

Asset inventory

When viewing assets, you can use the following keywords to search and filter.

User-specified fields

Comments

Use the syntax `comment:<text>` to search comments on an asset.

```
comment:"contractor laptop"
```

```
comment:"imaging server"
```

Tags

Use the syntax `tag:<term>` to search tags added to an asset. The term can be the tag name, or the tag name followed by an equal sign and the tag value. Tag value matches must be exact.

```
tag:"group"
```

```
tag:"group=production"
```

Organization name or ID

Use the syntax `organization:<term>` to filter by organization name or ID.

```
organization:runZero
```

```
organization:"Temporary Project"
```

```
organization:f1c3ef6d-cb41-4d55-8887-6ed3cfb3d42d
```

Site name or ID

Use the syntax `site:<term>` to filter by site name or ID.

```
site:Primary
```

```
site:"Branch Office"
```

```
site:ad67d649-041b-439d-af59-f200053a8899
```

Explorer name or ID

Use the syntax `explorer:<term>` to filter by Explorer name or ID.

```
explorer:DESKTOP-AB451F
```

```
explorer:8b927a8e-d405-40e9-aa47-d6afc9bff237
```

Hosted zone

Use the syntax `hosted_zone:<zone name>` to filter by the hosted runZero Explorer that found the asset. Using this filter after a hosted scan can be a good way to locate externally facing assets.

Owner

Use the syntax `owner:<term>` to filter by owner name.

```
owner:user@runzero.com
```

```
owner:"Security Team"
```

Ownership status

Use the syntax `owner_count:<number>` to filter by owner count. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
owner_count:>0
```

```
owner_count:0
```

Use the syntax `has_owner:<boolean>` to find assets with owners or assets that are missing owners.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
has_owner:t
```

```
has_owner:f
```

Use the syntax `ownership_type:<term>` to filter assets by ownership type name. This will return assets that have an owner assigned for the specified ownership type.

```
ownership_type:"Asset Owner"
```

```
ownership_type:"Security Owner"
```

Asset fields

Asset ID

The ID field is the unique identifier for a given asset, written as a UUID. Use the syntax `id:<uuid>` to filter by ID field.

```
id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

Operating system

The operating system field is a string describing the detected operating system software. This field is searched using the syntax `os:<text>`. The OS version, if available, can be searched using `os_version:<number>`.

```
os:"Windows"
```

```
os:"Ubuntu Linux"
```

```
os_version:8
```

OS CPE

The operating system Common Platform Enumeration (CPE) field is a string describing the detected operating system software aligned to the CPE naming scheme. This field is searched using the syntax `os.cpe23:<text>`. In cases where runZero was able to fingerprint the operating system but the NIST database does not contain an official matching entry, an unofficial CPE will be generated and include `r0_unofficial` in the other field of the CPE.

```
os.cpe23:"ubuntu"
```

```
os.cpe23:="cpe:/o:canonical:ubuntu_linux:22.04.1"
```

```
os.cpe23:="cpe:/o:alma:linux:-::~~~~~r0_unofficial"
```

Type

The type field is a string describing the detected system type, such as Desktop, Laptop, Server, BMC, or Mobile. Use the syntax `type:<text>` to search this field.

```
type:Desktop
```

```
type:BMC
```

```
type:"Game Console"
```

Hardware

The hardware field is a string describing the detected physical hardware, such as macMini or Nintendo Switch. Use the syntax `hardware:<text>` to search this field.

hardware:Switch

hardware:macMini

Hostnames

The hostnames associated with an asset are obtained from DNS and exposed services. Use the syntax `name:<text>` to search these names.

`name:"www"`

`name:"TV"`

To search an asset where any asset has a specific prefix or suffix, use the `:=` exact match operator, and use `%` as a wildcard:

`name:="FTP.%"`

`name:="%-09"`

Use the syntax `name_count:<number>` to search the hostname count. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

`name_count:>1`

Use the syntax `name_overlap:<boolean>` to find assets sharing the same name. The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

`name_overlap:t`

Domains

The domains associated with an asset are obtained from DNS and exposed services. Use the syntax `domain:<domainname>` to search the domain names.

`domain:"amazon.com"`

`domain:"corp.lan"`

`domain:"WORKGROUP"`

The domain count can be searched using the syntax `domain_count:<number>`. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

`domain_count:>1`

Addresses

Use the syntax `address:<ip>` to search the addresses (both primary and secondary) associated with an asset, `primary_address:<ip>` to search only the primary addresses associated with an asset, or `secondary_address:<ip>` to search only the secondary addresses associated with an asset. These keywords also allow for CIDR mask matching, as well as wildcard matches using `%`. A comma-separated list of addresses will be used as an efficient multiple-match.

```
address:192.168.0.1
```

```
address:10.0.0
```

```
address:10.1.2.0/24
```

```
address:%.0.1
```

```
address:10.%.254
```

```
address:10.0.0.1,10.0.0.2,10.0.0.3
```

Use the syntax `address_count:<term>` and `address_extra_count:<number>` to search address primary and secondary counts. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
address_extra_count:0
```

Use the syntax `address_overlap:<boolean>` to find assets sharing primary IP addresses. This can be further filtered to single sites using the `site` keyword. The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
address_overlap:t
```

Use the syntax `address_extra_overlap:<boolean>` to find assets sharing secondary IP addresses. This can be further filtered to single sites using the `site` keyword. The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
address_extra_overlap:t
```

Networks

Use the syntax `net:<cidr>` to search the addresses (both primary and secondary) associated with an asset by CIDR mask.

```
net:192.168.0.0/24
```

Default SNMP communities

Use the syntax `has:snmp.v2DefaultCommunities` to search for assets with a default SNMP community (public, private, and other defaults).

```
has:snmp.v2DefaultCommunities  
snmp.v2DefaultCommunities:public
```

Public address

Use the keyword `has_public` and syntax `has_public:<boolean>` to locate any asset with a non-reserved IP address. This often corresponds to public-facing systems, though public IPs can also be used internally behind a firewall. Note that public IPv6 addresses are included by this filter; to search for only public IPv4 addresses, you can use `has_public_v4`.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
has_public:true
```

Private address

Use the keyword `has_private` and syntax `has_private:<boolean>` to locate any asset with a private IP address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
has_private:false
```

IPv6 address

Use the keyword `has_ipv6` and the syntax `has_ipv6:<boolean>` to locate any asset with an identified IPv6 address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
has_ipv6:false
```

Link-local IPv6 address

Use the keyword `has_link_local` and syntax `has_link_local:<boolean>` to locate any asset with an identified IPv6 link local (`fe80::`) address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
has_link_local:true
```

MAC address

Use the syntax `mac:<term>` to search MAC addresses associated with an asset.

```
mac:00:5c:04
```

```
mac:00:00:1c
```

Use the syntax `mac_count:<number>` to search the MAC address count. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
mac_count:>2
```

If you use exact search (`:=`) you can also search for full MAC addresses in Cisco format or dash-separated format:

```
mac:=00-10-fa-c2-bf-d5
```

```
mac:=0010.fac2.bfd5
```

Use the syntax `mac_overlap:<boolean>` to find assets sharing the same MAC address. The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
mac_overlap:t
```

MAC address vendors

The vendor associated with the MAC addresses of an asset can be searched using the syntax `mac_vendor:<text>`.

```
mac_vendor:Apple
```

```
mac_vendor:"Intel Corporate"
```

To search only the vendor associated with the newest MAC address, use the syntax `newest_mac_vendor:<text>`

```
newest_mac_vendor:Apple
```

The MAC address vendor count can be searched using the syntax `mac_vendor_count:<number>`. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
mac_vendor_count:0
```

MAC address age

Use the syntax `mac_age:<term>` to search the allocation date of the newest MAC address associated with an asset. The term supports the standard runZero [time comparison syntax] [time].

```
mac_age:>1year
```

```
mac_age:<6months
```

```
mac_age:2019-12-31
```

Outlier score

Use the syntax `outlier_score:<value>` to search the calculated outlier score of assets. The outlier score is in the range 0 to 5 inclusive. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
outlier_score:>2
```

```
outlier_score:0
```

Upstream switch IP address

Use the syntax `switch.ip:<address>` to search the IP address of the upstream switch assets are connected to.

```
switch.ip:192.168.1.1
```

```
switch.ip:fe80::81f2:1c9d:5ac9:5420
```

Upstream switch name

Use the syntax `switch.name:<hostname>` to search the hostname of the upstream switch assets are connected to.

```
switch.name:"SWITCH-1"
```

```
switch.name:office
```

Upstream switch port

Use the syntax `switch.port:<address>-<port number string>` to search the port on the upstream switch assets that are connected to.

```
switch.port:192.168.1.1-25
```

```
switch.port:10.1.2.3-0/1/2
```

Upstream switch shared port

Use the syntax `attribute:switch.portShared` to find assets which connect to a switch port that reports multiple MAC addresses.

```
attribute:switch.portShared
```

Attributes

Use the syntax `attribute:<term>` to search the asset attribute fields, such as the port used to detect the TTL.

```
attribute:"ip.ttl.port"
```

```
attribute:"cpe:/a:isc:bind:9.11.3"
```

```
attribute:"9.11.3"
```

To determine if an asset has any attribute defined, use the `has:<attribute-name>` keyword. The `has` keyword can be inverted to find missing fields with `not has:<term>`.

```
has:"ip.ttl.port"
```

```
not has:"rdns.names"
```

In addition to the standard fields, the following special attributes are available:

- `has:screenshot` returns assets where at least one screenshot was obtained.
- `has:icons` returns assets where at least one icon was obtained (HTTP, UPnP, or similar).
- `has:uplink` returns assets seen in the CAM table of a network switch.
- `has:downlink` returns assets where the CAM table was queried at least one other asset was connected.
- `has:unmapped` returns assets where the CAM table was queried at least one other asset was connected but not identified by IP.

The attribute can be specified as a term directly. If the attribute name conflicts with an existing term, the prefix `_asset.` can be specified to disambiguate the query.

```
ip.ttl.port:80
```

```
rdns.names:"router"
```

```
_asset.ip.ttl.hops:"1"
```

Foreign attributes from third-party inbound integrations can be queried using the syntax `@<integration>.<source>.<attribute>:<term>`. The table below includes the correct prefix for each integration.

Integration	Prefix
Miradore	@miradore.dev.
AWS EC2	@aws.ec2.
AWS ELB & ELBv2	@aws.elb.
AWS RDS	@aws.rds.
CrowdStrike	@crowdstrike.dev.
Azure Load Balancer	@azure.vm.
Azure VM	@azure.vm.
Azure Scale Set VM	@azure.vms.
Censys	@censys.host.
VMWare	@vmware.vm.
GCP Load Balancer	@gcp.lb.
GCP E2-Micro VM	@gcp.vm.
GCP CloudSQL	@gcp.cloudsql.
SentinelOne	@sentinelone.dev.
Tenable.io & Nessus	@tenable.dev.
Rapid7 Nexpose & InsightVM	@rapid7.dev.
Qualys VMDR	@qualys.dev.
Shodan	@shodan.dev.
Azure AD	azuread
Active Directory (LDAP)	@ldap.computer.
Microsoft 365 Defender	@ms365defender.dev.
Microsoft Intune	@intune.dev.
Google Workspace ChromeOS	@googleworkspace.chromeos.
Google Workspace Endpoint	@googleworkspace.endpoint.

Google Workspace Mobile	@googleworkspace.mobile.
-------------------------	--------------------------

```
@aws.ec2.region:="us-east-2"
```

```
@crowdstrike.dev.agentVersion:="6.49.16201.0"
```

```
@googleworkspace.chromeos.model:="HP Chromebook"
```

Asset services

Service ports

The TCP and UDP services associated with an asset can be searched by port number using the syntax `port:<number>`.

```
port:80
```

```
port:161
```

Service TCP ports

Use the syntax `tcp:<number>` to search the TCP services associated with an asset by port number.

```
tcp:443
```

To search for assets with a specific list of TCP ports open, you can use the syntax `service_ports_tcp:=<list>`. Values should be in ascending numerical order, and separated by commas.

```
service_ports_tcp:=80,443
```

Service UDP ports

Use the syntax `udp:<number>` to search UDP services associated with an asset by port number.

```
udp:53
```

To search for assets with a specific list of UDP ports open, you can use the syntax `service_ports_udp:=<list>`. Values should be in ascending numerical order, and separated by commas.

```
service_ports_udp:=53,123
```

Service protocols

Use the syntax `service_protocols:<term>` (or `protocol:<term>` for short) to search the identified service protocols associated with an asset.

```
protocol:http
```

```
service_protocol:telnet
```

The protocol count can be searched using the syntax `protocol_count:<number>`. This search supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
protocol_count:>1
```

Service products

Use the syntax `service_products:<term>` (or `product:<term>` for short) to search for the identified service products associated with an asset.

```
product:openssh
```

```
service_products:nginx
```

The product count can be searched using the syntax `product_count:<number>`. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
product_count:>3
```

Service counts

Use the following keywords to search the number of services associated with an asset can be searched by port number:

- `service_count_tcp:<number>`
- `service_count_udp:<number>`
- `service_count_icmp:<number>`
- `service_count_arp:<number>`

These keywords support numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

Examples include:

```
service_count_tcp:>=5
```

```
service_count_arp:0
```

```
service_count_udp:<=1
```

Asset tracking fields

Timestamps

Use the following syntaxes to search the asset timestamp fields (`first_seen`, `last_seen`, `created_at`, `updated_at`, `os_eol`, `os_eol_extended`):

- `first_seen:<term>`
- `last_seen:<term>`
- `created_at:<term>`
- `updated_at:<term>`
- `os_eol:<term>`
- `os_eol_extended:<term>`

The term supports the standard runZero [time comparison syntax][time].

`first_seen:<3days`

`first_seen:>2019-08-01`

`first_seen:>8/1/2019`

`last_seen:<1week`

`last_seen:<2months`

`last_seen:<1year`

`created_at:>2weeks`

`created_at:<30minutes`

`updated_at:>1year`

`updated_at:<12hours`

`os_eol:<now`

`os_eol:>4weeks`

`os_eol_extended:>now`

`os_eol_extended:>90days`

Online status

Use the syntax `online:<boolean>` or the inverse syntax `offline:<boolean>` to search the online status of an asset.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

online:t

online:1

offline:0

Operating system support status

The syntax `os_eol_expired:<boolean>` can be used to find identify assets based on whether their operating systems are End of Life (EOL). This field evaluates both the `os_eol` and `os_eol_extended` values to only return assets with expired coverage.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

os_eol_expired:t

os_eol_expired:1

os_eol_expired:no

Detection method

The detected by attribute of an asset can be searched using the syntax `det:<term>` or `detected_by:<term>`. The term is one of `arp`, `icmp`, `<portnumber>-tcp`, or `<portnumber>-udp`. In the case of multiple detections, the priority goes `arp`, `icmp`, and then the first detected service.

det:arp

detected_by:80-tcp

det:53-udp

Time to Live (TTL) comparisons

Use the syntax `ttl:<term>` and `lowest_ttl:<term>` to search the lowest TTL of an asset. TTL is the estimated number of hops between the scan source and the asset.

This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

lowest_ttl:>3

Round Trip Time (RTT) comparisons

Use the syntax `rtt:<term>` and `lowest_rtt:<term>` to search the lowest RTT for an asset. RTT is the round-trip response time of a given probe measured in nanoseconds (1,000,000 ==

1ms).

This search term supports numerical comparison operators (>, >=, <, <=, =).

```
lowest_rtt:>50000000
```

Multiple MAC address status

Use the syntax `multi_mac:<boolean>` to determine if an asset has multiple MAC addresses.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

```
multi_mac:t
```

Any MAC address status

Use the syntax `has_mac:<boolean>` to find assets with any MAC addresses.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

```
has_mac:yes
```

```
has_mac:f
```

Multiple IP address status

Use the syntax `multi_home:<boolean>` to determine if an asset has multiple IP addresses.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

```
multi_home:t
```

Multiple hostname status

Use the syntax `multi_name:<boolean>` to find assets with multiple hostnames.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

```
multi_name:yes
```

```
multi_name:false
```

Software installations

Use the syntax `software:<term>>` to find assets with associated software.

The term has three forms:

- `software:<product>` will look for any assets with a software product that matches the term.
- `software:<product>/<version>` will look for any assets with a software product and version that matches the term exactly.
- `software:<vendor>/<product>/<version>` will look for any assets with a software vendor, product, and version that matches the term exactly.

All three forms allow the use of % as a wildcard (beginning, middle, or end of the term).

```
software:IIS
```

```
software:Microsoft/IIS/10.0
```

Service inventory

When viewing services, you can use the keywords in this section to search and filter.

Ports

The TCP and UDP services associated with a service can be searched by port number using the syntax `port:<number>`. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
port:<=25
```

TCP ports

Use the syntax `tcp:<number>` to search TCP service associated with a service by port number.

```
tcp:53
```

To search for all services on assets with a specific list of TCP ports open, you can use the syntax `service_ports_tcp:=<list>`. Values should be in ascending numerical order, and separated by commas.

```
service_ports_tcp:=80,443
```

UDP ports

Use the `udp:<number>` syntax to search UDP services associated with a service by port number.

```
udp:443
```

To search for all services on assets with a specific list of UDP ports open, you can use the syntax `service_ports_udp:=<list>`. Values should be in ascending numerical order, and separated by commas.

```
service_ports_udp:=53,123
```

Transport

Use the syntax `transport:<term>` to search the transport associated with a service by name.

```
transport:tcp
```

```
transport:udp
```

```
transport:icmp
```

Protocol

Use the syntax `service_protocols:<term>` (or `protocol:<term>` for short) to search the protocols associated with services.

```
protocol:http
```

```
protocol:telnet
```

Assets with product

Use the syntax `service_products:<term>` (or `product:<term>` for short) to search for the identified service products associated with an asset, and return all services for the matching assets.

```
product:openssh
```

```
service_products:nginx
```

Virtual Host (vHost)

Use the syntax `vhost:<text>` to search for virtual hosts associated with a service by name .

```
vhost:"www"
```

Address

Use the keyword `service_address` to match against the service IP address.

```
service_address:192.168.0.1
```

Public address

Use the keyword `service_has_public` and syntax `service_has_public:<boolean>` to locate any service with a non-reserved I address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
service_has_public:true
```

Private address

Use the keyword `service_has_private` and syntax `service_has_private:<boolean>` to locate any service with a private IP address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
service_has_private:false
```

IPv6 address

Use the keyword `service_has_ipv6` and the syntax `service_has_ipv6:<boolean>` to locate any service with an identified IPv6 address.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

```
service_has_ipv6:false
```

Link-local IPv6 address

Use the keyword `service_has_link_local` and syntax `service_has_link_local:<boolean>` to locate any service with an identified IPv6 link local (`fe80::`) address.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

```
service_has_link_local:true
```

Assets with outlier score

You can use the syntax `outlier_score:<value>` to search the calculated outlier score of assets, and return all services on those assets. The outlier score is in the range 0 to 5 inclusive. This search term supports numerical comparison operators (>, >=, <, <=, =).

```
outlier_score:>2
```

```
outlier_score:0
```

Assets with MAC address vendors

To search the vendors associated with the MAC addresses of an asset, and return all services on those assets, use the syntax `mac_vendor:<text>`.

```
mac_vendor:Apple
```

```
mac_vendor:"Intel Corporate"
```

To search only the vendor associated with the newest MAC address, use the syntax `newest_mac_vendor:<text>`

```
newest_mac_vendor:Apple
```

Assets with MAC address age

To search the ages of the newest MAC addresses associated with each asset, and return all services associated with those assets, use the syntax `mac_age:<term>`. The term supports the standard runZero [time comparison syntax][time].

```
mac_age:>1year
```

```
mac_age:<6months
```

```
mac_age:2019-12-31
```

Attributes

You can search all service attributes with the syntax `<attribute>:<term>`. This search term supports numerical comparison operators (>, >=, <, <=, =).

If the attribute name conflicts with an existing term, the prefix `_service.` can be added to disambiguate the query.

Note that service attributes can be slow and it is often better to prefix `_asset.protocol:` `<term>` filter in front of the service attribute query. For example, to search for SSH banners, use the syntax `_asset.protocol:ssh AND banner:<term>`.

```
banner:password
```

```
service.product:"OpenSSH"
```

```
html.title:"Apache2 Ubuntu Default Page"
```

```
http.code:>=500
```

```
screenshot.image.size:>=100000
```

```
_service.arp.macVendor:Xerox
```

To determine if a service has an attribute at all, use the `has` keyword. The `has` keyword can be inverted to find missing fields, with `not has:<term>`.

```
has:"http.head.server"
```

```
not has:"html.title"
```

Software inventory

When viewing software groups, you can use the keywords in this section to search and filter.

Vendor

The vendor associated with a software can be searched by name using the syntax `vendor:<name>`.

```
vendor:oracle
```

Product

The product associated with a software can be searched by name using the syntax `product:<name>`.

```
product:java
```

Version

The version associated with a software can be searched by name using the syntax `version:<name>`.

`version:1.2.3`

Software instance inventory

When viewing software instances on assets, you can use the keywords in this section to search and filter.

Source

The source reporting the software installed can be searched or filtered by name using the syntax `source:<name>`.

`source:runzero`

Vendor

The vendor associated with a software can be searched by name using the syntax `vendor:<name>`.

`vendor:oracle`

Product

The product associated with a software can be searched by name using the syntax `product:<name>`.

`product:java`

Vulnerability inventory

When viewing the Vulnerabilities inventory, you can use the following keywords to search and filter information.

Vulnerability ID

The ID field is the unique identifier for a given vulnerability, written as a UUID. Use the syntax `id:<uuid>` to filter by the ID field.

`id:a124a141-e518-4735-9878-8e89c575b1d2`

Source

The source reporting the vulnerability detected can be searched or filtered by name using the syntax `source:<name>`.

```
source:tenable
```

Severity

The severity field can be searched using the syntax `severity:<term>`.

```
severity:info
```

```
severity:medium
```

Severity score

The severity score can be searched using the syntax `severity_score:<term>`. The term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
severity_score:<5.0
```

```
severity_score:>=9.0
```

Risk

The risk field can be searched using the syntax `risk:<term>`.

```
risk:none
```

```
risk:critical
```

Risk score

The risk score can be searched using the syntax `risk_score:<term>`. The term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
risk_score:>7.0
```

```
risk_score:=10.0
```

Category

The category field can be searched using the syntax `category:<term>`.

```
category:Local
```

```
category:Remote
```

Name

The name field can be searched using the syntax name:<term>.

```
name:"Cisco IOS Software DHCP Remote Code Execution Vulnerability"
```

```
name:"PHP < 5.3.12 / 5.4.2 CGI Query String Code Execution"
```

Description

The description field can be searched using the syntax description:<term>.

```
description:"The remote device is missing a vendor-supplied security patch."
```

```
description:"remote code execution"
```

Solution

The solution field can be searched using the syntax solution:<term>.

```
solution:patch
```

```
solution:upgrade
```

CVE

The CVE field can be searched using the syntax cve:<term>.

```
cve:CVE-2021-44228
```

```
cve:CVE-2016-2183
```

CVSSv2 score

The CVSSv2 fields, cvss2_base_score and cvss2_temporal_score, can be searched using the syntax cvss2_base_score:<term> and cvss2_temporal_score:<term>. The term supports numerical comparison operators (>, >=, <, <=, =).

```
cvss2_base_score:>6.5
```

```
cvss2_base_score:<=3.0
```

```
cvss2_temporal_score:=10.0
```

```
cvss2_temporal_score:<5.0
```

CVSSv3 score

The CVSSv3 fields, `cvss3_base_score` and `cvss3_temporal_score`, can be searched using the syntax `cvss3_base_score:<term>` and `cvss3_temporal_score:<term>`. The term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
cvss3_base_score:>6.5
```

```
cvss3_base_score:<=3.0
```

```
cvss3_temporal_score:=10.0
```

```
cvss3_temporal_score:<5.0
```

Address

The address field can be searched using the syntax `address:<term>`.

```
address:192.168.0.1
```

Transport

The transport field can be searched using the syntax `transport:<term>`.

```
transport:tcp
```

```
transport:udp
```

Port

The port can be searched using the syntax `port:<term>`. The term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
port:22
```

```
port:443
```

Operating system support status

The syntax `os_eol_expired:<boolean>` can be used to find identify vulnerabilities on assets based on whether their operating systems are End of Life (EOL). This field evaluates both the `os_eol` and `os_eol_extended` values to only return vulnerabilities on assets with expired coverage.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

os_eol_expired:t

os_eol_expired:1

os_eol_expired:no

Detection timestamps (first detected at, last detected at)

The timestamp fields, `first_detected_at` and `last_detected_at`, can be searched using the syntax `first_detected_at:<term>` and `last_detected_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

`first_detected_at:>2weeks`

`first_detected_at:<30minutes`

`last_detected_at:>1month`

`last_detected_at:2hours`

Publication timestamps (published at)

The timestamp field, `published_at`, can be searched using the syntax `published_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

`published_at:>2weeks`

`published_at:<30minutes`

Timestamps (created at, updated at)

The timestamp fields, `created_at` and `updated_at`, can be searched using the syntax `created_at:<term>` and `updated_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

`created_at:>2weeks`

`created_at:<30minutes`

`updated_at:>1month`

`updated_at:2hours`

Attributes

You can search all vulnerability attributes with the syntax `<attribute>:<term>`. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

If the attribute name conflicts with an existing term, the prefix `_vulnerability.` can be added to disambiguate the query.

```
plugin.hasPatch:true
```

```
severityID:3
```

```
_vulnerability.state:REOPENED
```

To determine if a vulnerability has any attribute defined, use the `has:<term>` keyword. The `has` keyword can be inverted to find missing fields with `not has:<term>`.

```
has:plugin.vpr.score
```

```
not has:output
```

Wireless inventory

When viewing WiFi networks, you can use the keywords in this section to search and filter.

SSID and ESSID

The SSID/ESSID field can be searched using the syntax `ssid:<text>`.

```
ssid:"Guest Network"
```

```
ssid:"Corporate"
```

BSSID (MAC)

The BSSID field can be searched using the syntax `bssid:<text>` or `mac:<text>`.

```
bssid:"00:01:02:03:04:05"
```

```
mac:"00:01:%"
```

Vendor

The vendor field can be searched using the syntax `mac_vendor:<text>`.

```
mac_vendor:"Google"
```

```
mac_vendor:"Netgear"
```

```
mac_vendor:"Cisco"
```

Family

The family field can be searched using the syntax `family:<term>`.

```
family:"010304"
```

Channels

The channels field can be searched using the syntax `channel:<term>`.

```
channel:"11"
```

Type

The network type field can be searched using the syntax `type:<text>`.

```
type:"infrastructure"
```

Interface

The network interface field can be searched using the syntax `interface:<text>`.

```
interface:"wlan0"
```

Encryption

The encryption field can be searched using the syntax `encryption:<term>`.

```
encryption:"aes"
```

```
encryption:"none"
```

Authentication

The authentication field can be searched using the syntax `authentication:<term>`.

```
authentication:"wpa2-psk"
```

```
authentication:"open"
```

Timestamps

The timestamp fields (`first_seen`, `last_seen`, `created_at`) timestamps can be searched using the syntax `first_seen:<term>`, `last_seen:<term>` and `created_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

```
first_seen:<30seconds
```

first_seen:>2019-08-01

last_seen:<1week

last_seen:<2months

created_at:>2weeks

created_at:<30minutes

Signal

The signal field can be searched using the syntax `signal:<number>` or `sig:<number>`. The term can include the operators `>`, `=`, `<=`, and `=`. The default operator is `=`.

signal:">75"

signal:"<=25"

signal:99

Organization name or ID

Use the syntax `organization:<term>` to filter by organization name or ID.

organization:runZero

organization:"Temporary Project"

organization:f1c3ef6d-cb41-4d55-8887-6ed3cfb3d42d

Site name or ID

The site name or ID can be used as a filter with the syntax `site:<term>`

site:Primary

site:"Branch Office"

site:ad67d649-041b-439d-af59-f200053a8899

Explorer name or ID

The Explorer name or ID can be used as a filter with the syntax `explorer:<term>`

explorer:DESKTOP-AB451F

explorer:8b927a8e-d405-40e9-aa47-d6afc9bff237

Wireless ID

The ID field is the unique identifier for a given wireless network, written as a UUID. This field is searched using the syntax `id:<uuid>`.

```
id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

Last task ID

The Last Task ID field defines which task most recently reported the wireless network and is written as a UUID. This field is searched using the syntax `task:<uuid>`.

```
task:39ab0e71-3cf1-4176-b6b0-4ed495288229
```

Wireless attributes

All wireless attributes can be searched using the syntax `<attribute>:<term>`.

```
radio_type:"802.11n"
```

Users inventory

When viewing the Users inventory, you can use the following keywords to search and filter users.

Source

The source reporting the users can be searched or filtered by name using the syntax `source:<name>`.

```
source:ldap
```

Name fields

There are multiple name fields found in the user attributes that can be searched or filtered using the same syntax. Use the syntax `<name_field>:<text>` to search a field for matches.

The following name fields can be searched this way:

- name
- display_name
- first_name
- last_name

```
name:j
```

```
display_name:"john doe"
```

first_name:john

last_name:doe

Description

Use the syntax description:<text> to search the results in this field.

description:shared

Groups

The group_id field is the unique identifier for a given group, written as a UUID. To search for users that are part of a group based on the group's ID, use the syntax group_id:<uuid>.

group_id:f8a26321-cbce-4fb5-a9ca-ffa809489dd5

Email

Use the syntax email:<address> to search for users by email address.

email:john@example.com

Phone

Use the syntax phone:<phone number> to search for users by phone number.

phone:888-555-1234

Title

Use the syntax title:<text> to search for users by title.

title:CISO

Location

Use the syntax location:<text> to search for users by location.

location:TX

Last logon

Use the syntax last_logon_at:<date> to search for users by last logon time. This field is compatible with the [date comparison operators](#).

last_logon_at:<1year

First seen

Use the syntax `first_seen_at:<date>` to search for users by when they were first seen in the organization. This field is compatible with the [date comparison operators](#).

```
first_seen_at:<1day
```

Last seen

Use the syntax `last_seen_at:<date>` to search for users by when they were last seen in the organization. This field is compatible with the [date comparison operators](#).

```
last_seen_at:>2years
```

Site name or ID

Use the syntax `site:<term>` to filter by site name or ID.

```
site:Primary
```

```
site:"Branch Office"
```

```
site:ad67d649-041b-439d-af59-f200053a8899
```

Organization name or ID

Use the syntax `organization:<term>` to filter by organization name or ID.

```
organization:runZero
```

```
organization:"Temporary Project"
```

```
organization:f1c3ef6d-cb41-4d55-8887-6ed3cfb3d42d
```

Groups inventory

When viewing the Groups inventory, you can use the following keywords to search and filter groups.

Source

The source reporting the groups can be searched or filtered by name using the syntax `source:<name>`.

```
source:ldap
```

Name fields

There are two name fields found in the group attributes that can be searched or filtered using the same syntax. Use the syntax `<name_field>:<text>` to search a field for matches.

The following name fields can be searched this way:

- name
- display_name

```
name:admin
```

```
display_name:"Domain Administrators"
```

Description

Use the syntax `description:<text>` to search the results in this field.

```
description:admin
```

Users

The `user_id` field is the unique identifier for a given user, written as a UUID. To search for groups that have a specific member based on the user's ID, use the syntax `user_id:<uuid>`.

```
user_id:f8a26321-cbce-4fb5-a9ca-ffa809489dd5
```

Email

Use the syntax `email:<address>` to search for users by email address.

```
email:john@example.com
```

First seen

Use the syntax `first_seen_at:<date>` to search for users by when they were first seen in the organization. This field is compatible with the [date comparison operators](#).

```
first_seen_at:<1day
```

Last seen

Use the syntax `last_seen_at:<date>` to search for users by when they were last seen in the organization. This field is compatible with the [date comparison operators](#).

```
last_seen_at:>2years
```

Site name or ID

Use the syntax `site:<term>` to filter by site name or ID.

```
site:Primary
```

```
site:"Branch Office"
```

```
site:ad67d649-041b-439d-af59-f200053a8899
```

Organization name or ID

Use the syntax `organization:<term>` to filter by organization name or ID.

```
organization:runZero
```

```
organization:"Temporary Project"
```

```
organization:f1c3ef6d-cb41-4d55-8887-6ed3cfb3d42d
```

Component keywords

The data across your runZero account can be queried and filtered using the [search syntax](#) in conjunction with the available component keywords. Keywords and example values are documented for the following types of components in your console:

- [Scan templates](#)
- [Tasks](#)
- [Analysis reports](#)
- [Explorers](#)
- [runZero users and groups](#)
- [Sites and organizations](#)
- [Credentials](#)
- [Queries](#)
- [Events](#)

Scan templates

When viewing scan templates, you can use the keywords in this section to search and filter.

ID

The ID field is the unique identifier for a given template, written as a UUID. Use the syntax `id:<uuid>` to filter by ID field.

```
id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

Name

Use the syntax `name:<text>` to search by scan template name.

```
name:WiFi
```

```
name:"Data Center"
```

Timestamps

Use the following syntaxes to search the scan template timestamp fields (`created_at`, `updated_at`):

- `created_at:<term>`
- `updated_at:<term>`

The term supports the standard runZero `[time comparison syntax][time]`.

```
created_at:>2weeks
```

```
created_at:<30minutes
```

```
updated_at:>1year
```

```
updated_at:<12hours
```

Created by

The email address for the user that created the template can be searched using the syntax `created_by_email:<term>`.

```
created_by_email:user@example.com
```

Tasks

When viewing all tasks, you can use the keywords in this section to search and filter them.

Name

The Name field can be searched using the syntax `name:<text>`.

```
name:"test scan"
```

Description

The Description field can be searched using the syntax `description:<text>`

```
description:"full scan"
```

Created by

The Created By field can be searched using the syntax `created_by:<term>`.

```
created_by:"admin"
```

Type

The task type can be searched using `type:<text>`.

```
type:scan
```

Status

The task status can be searched using `status:<text>`.

```
status:error
```

Error

The task error message can be searched using `error:<text>`.

```
error:"no disk space"
```

Recurrence frequency

The frequency tasks recur at (the "Freq" column) can be searched using `recur_frequency:<text>` or `freq:<text>`. The term `recurring:<boolean>` or `recur:<boolean>` can be used to search based on whether tasks recur at all.

```
recur_frequency:hourly
```

```
freq:daily
```

```
freq:continuous
```

```
recur:true
```

To search for tasks with a frequency of Nth Weekday of Month, you can use (for example) `freq:nth_weekday,2 freq:monday` to find tasks which repeat on the second monday of each month.

Timestamps (created at, updated at)

The timestamp fields, `created_at` and `updated_at`, can be searched using the syntax `created_at:<term>` and `updated_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

`created_at:>2weeks`

`created_at:<30minutes`

`updated_at:>1month`

`updated_at:2hours`

Next/last run time

You can search by next recurrence and last recurrence using the terms `recur_last:<term>` and `recur_next:<term>`. The term supports the standard runZero [time comparison syntax][time].

`recur_last:<2hours`

`recur_next:>1day`

Start time

You can search by start time using the syntax `start_time:<term>`. The term supports the standard runZero [time comparison syntax][time].

`start_time:<2hour`

Grace period

The [grace period](#) can be searched using the syntax `grace_period:<term>` or just `grace:<term>`. The term supports the standard runZero [time comparison syntax][time].

`grace:<2hour`

Site name or ID

Use the syntax `site:<term>` to filter by site name or ID.

`site:Primary`

`site:"Branch Office"`

`site:ad67d649-041b-439d-af59-f200053a8899`

Template ID

Use the syntax `template_id:<term>` to filter by scan template ID.

`template_id:de657459-041b-439d-af59-ff1f153a7722`

Source

The data source for tasks can be searched using the term `source:<text>` or `source_id:<number>`.

`source:censys`

Sources are:

ID	Name	Description
1	runzero	runZero
2	miradore	Miradore
3	aws	Amazon Web Services
4	crowdstrike	CrowdStrike
5	azure	Microsoft Azure
6	censys	Censys
7	vmware	VMWare
8	gcp	Google Cloud Platform
9	sentinelone	SentinelOne
10	tenable	Tenable.io & Nessus
12	rapid7	Rapid7 Nexpose & InsightVM
14	qualys	Qualys VMDR
15	shodan	Shodan
16	azuread	Azure AD
17	ldap	Active Directory (LDAP)
18	ms365defender	Microsoft 365 Defender
19	intune	Microsoft Intune
20	googleworkspace	Google Workspace

22	tenablesecuritycenter	Tenable Security Center
23	packet	runZero packet capture import
24	wiz	Wiz

Credential ID

You can search for tasks that use a specific set of credentials using `credential_id:<id>`.

`credential_id:d7931a68-6e56-11ec-ad72-f875a414a63a`

Parameters

Tasks can be searched for task parameters using `params:<text>`. This can be useful for searching for scan tasks that had specific probes enabled.

`params:bacnet`

Asset counts

Completed tasks can be searched by the asset counts found in their results. The available search terms are:

- New assets: `assets_new:<number>`
- Assets back online: `assets_back_online:<number>`
- Assets marked offline: `assets_marked_offline:<number>`
- Assets changed: `assets_changed:<number>`
- Assets unchanged: `assets_unchanged:<number>`
- Assets ignored: `assets_ignored:<number>`
- Assets updated: `assets_updated:<number>`

These terms support numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

`assets_new:>0`

`assets_unchanged:>=1`

Analysis reports

When viewing generated analysis reports, you can use the keywords in this section to search and filter.

Name

The Name field can be searched using the syntax `name:<text>`.

`name:"main"`

Description

The Description field can be searched using the syntax `description:<text>`

`description:"compare secondary"`

Type

The report type can be searched using the syntax `type:<text>`

`type:outliers`

Report ID

The ID field is the unique identifier for a given analysis report, written as a UUID. This field is searched using the syntax `id:<uuid>`.

`id:cdb084f9-4811-445c-8ea1-3ea9cf88d536`

Timestamps

The timestamp when a report was generated can be searched using the syntax `created_at:.`

The term supports the standard runZero `[time comparison syntax][time]`.

`created_at:>2019-08-01`

`created_at:<1week`

Created by

The Created By field can be searched using the syntax `created_by:<term>`.

`created_by:jsmith`

Explorers

When viewing deployed Explorers, you can use the keywords in this section to search and filter.

Name

The Name field can be searched using the syntax `name:<text>`.

`name:"main"`

Site

The site can be searched using the syntax `site:<text>`.

```
site:Primary
```

Up

Whether the Explorer is up can be searched using the syntax `up:<boolean>`.

```
up:true
```

Address

The IP address(es) the Explorer is deployed on can be searched using the syntax `address:<IP address>`.

```
address:10.0.1.200
```

Version

The software version of Explorers can be searched using `version:<text>`.

```
version:2.9.7
```

Npcap version

The version of the npcap library for Windows Explorers can be searched using `npcap_version:<text>`.

```
npcap_version:1.60
```

Architecture

The machine architecture Explorers are deployed on can be searched using `architecture:<text>`.

```
architecture:amd64
```

OS

The operating system Explorers are deployed on can be searched using `os:<text>`. Note that macOS is recorded as `darwin`, the underlying Unix core of macOS.

```
os:windows
```

```
os:darwin
```

Capability

The capabilities of the Explorers can be searched using the syntax `capability:<keyword>`. Two keywords are supported:

- `screenshot` for Explorers which can screenshot web pages
- `ec2` for Explorers which can describe AWS EC2 instances

Example:

```
capability:screenshot
```

Tags

Use the syntax `tag:<term>` to search tags added to an Explorer. The term can be the tag name, or the tag name followed by an equal sign and the tag value. Tag value matches must be exact.

```
tag:"admin"
```

```
tag:"group=cloud"
```

runZero users and groups

User search keywords

When viewing users, you can use the keywords in this section to search and filter.

Email

Use the syntax `email:<address>` to search for someone by email address.

```
email:john@example.com
```

Name

Use the syntax `name:<text>` to search for someone by name.

```
name:john
```

```
name:"John Smith"
```

Superuser

To search for people based on whether they have superuser access, use the term `superuser:<boolean>`.

```
superuser:true
```

superuser:f

Access

Use the syntax `access:<term>` to search for users with a specific access level. Possible access levels are `admin`, `user`, `annotator`, `viewer`, `billing` and `none`.

`access:admin`

Status

To search for users by invitation status, use the term `status:<text>`. Possible status values are `activated`, `pending` and `expired`.

`status:pending`

SSO

To search for people based on whether they can only log in via SSO, use the term `sso:<boolean>`.

`sso:true`

MFA

To search for people based on whether they have enrolled an MFA token, use the term `mfa:<boolean>`.

`mfa:f`

Group ID

The `group_id` field is the unique identifier for a given group, written as a UUID. To search for users that are part of a group based on the group's ID, use the syntax `group_id:<uuid>`.

`group_id:cdb084f9-4811-445c-8ea1-3ea9cf88d536`

Group name

To search for users that are part of a group based on the group's name, use the syntax `group_name:<text>`.

`group_name:administrators`

`group_name:"Temp annotators"`

User groups

To search for users based on whether they are part of a group or not, use the term `has_group:<boolean>`. Use boolean value `t` or `true` to show all users who are members of a group.

```
has_group:t
```

Conversely, use boolean value `f` or `false` to show users who are not members of a group.

```
has_group:f
```

Group count

Use the syntax `group_count:<number>` to search the group membership count. This search term supports numerical comparison operators (`>`, `>=`, `<`, `<=`, `=`).

```
group_count:>1
```

```
group_count:=0
```

Group search keywords

When viewing your groups, you can use the keywords in this section to search and filter.

ID

The ID field is the unique identifier for a given group, written as a UUID. Use the syntax `id:<uuid>` to filter by ID field.

```
id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

Name

Use the syntax `name:<text>` to search by group name.

```
name:administrators
```

```
name:"Temp annotators"
```

Access

Use the syntax `access:<term>` to search for groups with a specific access level. Possible access levels are `admin`, `user`, `annotator`, `viewer`, `billing` and `none`.

```
access:admin
```

Timestamps (created at, updated at)

Filter groups by their timestamp fields, `created_at` and `updated_at`, using the syntax `created_at:<term>` and `updated_at:<term>`. The terms support the standard runZero [time comparison syntax][time].

`created_at:<30days`

`updated_at:<1week`

Expiration

Filter groups by their expiration timestamp, `expires_at`, using the syntax `expires_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

`expires_at:<30days`

`expires_at:>8/1/2019`

The `expired` property describes whether or not a group has expired. Search this property using the `expired:<boolean>` syntax.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

`expired:true`

`expired:0`

Use the syntax `has_expiration:<term>` to find any assets with an expiration date.

The term is a boolean value:

- `true`, `t`, `1`, and `yes` represent *true*
- `false`, `f`, `0`, and `no` represent *false*

`has_expiration:true`

`has_expiration:0`

Email

The `created_by_email` property holds the email address for the user that created the group. It can be searched using the syntax `created_by_email:<term>`.

`created_by_email:user@runzero.com`

Group mapping search keywords

When viewing your SSO group mappings, you can use the keywords in this section to search and filter.

ID

The ID field is the unique identifier for a given group mapping, written as a UUID. Use the syntax `id:<uuid>` to filter by ID field.

```
id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

SSO attribute

The `sso_attribute` is the name of the attribute field to check for matching values. Use the syntax `sso_attribute:<text>` to search by `sso_attribute`.

```
sso_attribute:department
```

SSO value

The `sso_value` is the value or comma-separated list of values to match. Use the syntax `sso_value:<text>` to search by `sso_value`.

```
sso_value:security
```

```
sso_value:"admins, administrators"
```

Group ID

The `group_id` field is the unique identifier for a given group, written as a UUID. To search for group mappings related to a group based on the group's ID, use the syntax `group_id:<uuid>`.

```
group_id:cdb084f9-4811-445c-8ea1-3ea9cf88d536
```

Group name

To search for group mappings related to a group based on the group's name, use the syntax `group_name:<text>`.

```
group_name:administrators
```

```
group_name:"Temp annotators"
```

Timestamps (created at, updated at)

Filter group mappings by their timestamp fields, `created_at` and `updated_at`, using the syntax `created_at:<term>` and `updated_at:<term>`. The terms support the standard runZero [time comparison syntax][time].

created_at:<30days

updated_at:<1week

Email

The created_by_email property holds the email address for the user that created the group. It can be searched using the syntax created_by_email:<term>.

created_by_email:user@runzero.com

Sites and organizations

Site search keywords

When viewing sites, you can use the keywords in this section to search and filter.

Name

The Name field can be searched using the syntax name:<text>.

name:"Primary"

Description

The Description field can be searched using the syntax description:<text>.

description:"wireless"

description:"vlan 50"

Scope

The Scope field can be searched using the syntax scope:<term> .

scope:"10.10.10."

Excludes

The Excludes field can be searched using the syntax excludes:<term> .

excludes:"192.168.0."

Timestamps (created at, updated at)

The timestamp fields (created_at, updated_at) timestamps can be searched using the syntax created_at:<term> and updated_at:<term>. The term supports the standard runZero [time comparison syntax][time].

created_at:>2weeks

created_at:<30minutes

updated_at:>1month

updated_at:2hours

Organization search keywords

Name

The Name field can be searched using the syntax name:<text>.

name:"main"

Description

The Description field can be searched using the syntax description:<text>

description:"branch office"

description:"pci"

Timestamps (created at, updated at)

The timestamp fields (created_at, updated_at) timestamps can be searched using the syntax created_at:<term> and updated_at:<term>. The term supports the standard runZero [time comparison syntax][time].

created_at:>2weeks

created_at:<30minutes

updated_at:>1month

updated_at:2hours

Credentials

When viewing saved credentials, you can use the keywords in this section to search and filter.

Credential fields

Credential ID

The ID field is the unique identifier for a given credential, written as a UUID. This field is searched using the syntax id:<uuid>.

id:cdb084f9-4811-445c-8ea1-3ea9cf88d536

Credential name

The credential name can be searched using the syntax name:<text>.

name:"AWS read-only account"

name:"Miradore API key"

Credential type

The credential type can be searched using the syntax name:<text>.

type:aws_access_secret

type:miradore_api_key_v1

Credential global property

The global property describes the level of access for all organizations. If a credential is global, all organizations have access to it. The global property can be searched using the syntax global:<boolean>.

The term is a boolean value:

- true, t, 1, and yes represent *true*
- false, f, 0, and no represent *false*

global:true

global:0

Credential timestamps

Credential timestamp fields (created_at and last_used_at) can be searched using the syntax:

- created_at:<term>
- last_used_at:<term>

The term supports the standard runZero [time comparison syntax][time].

created_at:<3days

created_at:>2019-08-01

created_at:>8/1/2019

created_at:<1week

created_at:<2months

last_used_at:<1year

last_used_at:>2weeks

last_used_at:<30minutes

last_used_at:>1year

last_used_at:<12hours

last_used_at:0

Credential created by

The created_by_email holds the email address for the user that created the credential. It can be searched using the syntax created_by_email:<term>.

created_by_email:user@example.com

Queries

When viewing saved queries, you can use the keywords in this section to search and filter.

Name

The Name field can be searched using the syntax name:<text>.

name:"smb2"

Description

The Description field can be searched using the syntax description:<text>.

description:"smb version 1"

description:"wep"

Type

The Type field can be searched using the syntax type:<term> .

type:"services"

Category

The Category field can be searched using the syntax category:<term>.

```
category:"security"
```

```
category:"audit"
```

Severity

The Severity field can be searched using the syntax `severity:<term>`.

```
severity:"info"
```

```
severity:"critical"
```

Created by

The Created By field can be searched using the syntax `created_by:<term>`.

```
created_by:"runzero"
```

Timestamps (created at, updated at)

The timestamp fields, `created_at` and `updated_at`, can be searched using the syntax `created_at:<term>` and `updated_at:<term>`. The term supports the standard runZero [time comparison syntax][time].

```
created_at:>2weeks
```

```
created_at:<30minutes
```

```
updated_at:>1month
```

```
updated_at:2hours
```

Events

When viewing system events under alerts, you can use the keywords in this section to search and filter.

Note that event records are retained for one year.

Action

Use the syntax `action:<text>` to search by the action which caused the event.

```
action:agent-reconnected
```

Created timestamp

The timestamp fields `created_at` can be searched using the syntax `created_at:<term>`. The term supports the standard runZero [time comparison syntax](#).

`created_at:>2weeks`

`created_at:<30minutes`

`updated_at:>1month`

`updated_at:2hours`

Details

The details in the event record can be searched using the syntax `details:<text>`. This can be useful for searching for IP addresses.

`details:192.168.0.1`

Source and target name

The source (`src`) column can be searched using the syntax `src:<text>` or `source:<text>`. The target (`tgt`) column can be searched using `tgt:<text>` or `target:<text>`.

`src:crowdstrike`

`target:primary`

Source and target type

The source type (shown at the start of the `src` column) can be searched using the syntax `src_type:<text>` or `source_type:<text>`.

Similarly, the target type can be searched using `tgt_type:<text>` or `target_type:<text>`.

`src_type:task`

`target_type:site`

Organization, site, source and target IDs

The IDs of organizations, sites, sources and targets mentioned in event details can be searched using the following search terms:

- `organization_id:<uuid>`
- `site_id:<uuid>`
- `source_id:<uuid>` or `src_id:<uuid>`

- target_id:<uuid> or tgt_id:<uuid>

The IDs are unique and are written as UUIDs.

organization_id:0eacf412-6e69-11ec-88b9-f875a414a63a

Automating queries

Professional

Community

Platform

runZero's query language allows you to search and filter your asset inventory, based on asset fields and values. runZero includes a query library of prebuilt searches which can be browsed from the Queries page. You can apply these queries after a scan to investigate discovery findings.

In addition to a flexible query language, the same search syntax can be used to track and monitor events across your network, based on any combination of fields. You can save your custom queries to reuse over and over again. Review the [query syntax](#) documentation for a refresher on how to build a query.

If there are certain queries you always want to run after a scan, you can turn on the **Auto** option. After the query runs, you will be able to view its results in the [Queries table](#).

From the Queries table, any filter that has checkmark in the Auto column means that it is active. The Matches column is next to it. You can click on the numbered result in the Matches column to view the results of the query. If there were no matches, the filter will show 0.

Turn on automatic search queries

1. Select Queries from the left navigator.
2. Open or create the query you want to turn on.
3. Turn on the **Automatically track query results on the dashboard** option.
4. Save the query.

Generating reports

Asset route pathing report

Platform

The asset route pathing report generates a visualization of the potential network paths between a source asset and destination asset in an organization. Following the paths, you can see assets connected between the target and source destinations. These assets represent opportunities an attacker could potentially leverage to break into your target asset.

The runZero Explorer performs a traceroute between itself and the source, and then another with the target. runZero then compares the data to infer shared points between the assets. runZero does not get any paths from a direct traceroute. This is runZero's best effort-based on the scan data it has—to identify the assets that it sees as viable points between two assets.

You can share this report with your IT and security teams to highlight assets that could be leveraged as pivot points to your critical assets. Armed with this information, they can identify systems that may need to be hardened. They can assess whether or not the appropriate critical controls are in place to prevent unauthorized access to those assets.

Customers with highly segmented environments can use this report to quickly identify paths from low security assets to critical assets. For example, this report can indicate whether a device in a wireless guest network can reach a system within the PCI cardholder data environment.

Terminology

Before diving into the asset route pathing report, here are some terms you need to know:

- **Hop** - Any node between the source and destination.
- **Node** - Any asset or hop. Nodes can be an IP, an asset, or unknown.
- **Asset** - Any device that is part of your runZero inventory.
- **Network path** - runZero does not get any paths from a direct traceroute. Instead, runZero uses the data it has to identify and display potential network paths between the target and source.

Generating the asset route pathing report

Note: This report is limited to runZero Platform customers.

1. Launch the [asset route pathing report](#).
2. When the Asset route pathing page appears, you will need to select a source asset and a destination asset. Use the search to filter assets by keyword or the table pagination to browse all of your assets.

3. After you have a source and destination asset selected, start the trace.
4. In the generated results, you will see the potential paths between the source and destination asset.

Analyzing the report

Locate your source asset and target asset. If there are hops between the two assets, you should review them and secure the paths between them. Take a look at the services running on those systems that may provide potential entry points for attacks and harden them.

Nodes in the [asset route pathing report](#) are color-coded to help you identify the source asset and destination asset.

The report uses the following colors:

- Green - The source asset
- Red - The destination asset
- Orange - A multi-homed asset that may act as a pivot point
- Blue - A standard layer-3 routing hop
- Gray - Asset is unknown

A hop labeled Unknown indicates an intermediate hop in the layer-3 path that did not respond with ICMP errors for TTL exceeded packets.

Sharing the report

There are a couple of ways to share the results from the [asset route pathing report](#). You can either export a PNG or dotfile of the report or share a direct link to the report.

- **Export a PNG** - A PNG export will take a snapshot of your report. To share the [asset route pathing report](#), click **Export view**. A PNG will download to your computer.
- **Export a dotfile** - You can export a dotfile and feed it into a Graphviz engine or open source visualization tool. The file allows you to render the image in formats like SVG, PSD, and PNG.
- **Share a link** - When viewing the generated asset route pathing report, copy the URL. You can share the URL directly with other team members who have a runZero account and access to the organization.

Exporting a dotfile

A dotfile is a text file that can be fed into Graphviz engines or open source visualization tools. With the dot file, you can render the image in other [formats](#), like SVG, PSD, and PNG.

1. Launch the [asset route pathing report](#).
2. When the Trace path page appears, you will need to select a source asset and a destination asset. Use the search to filter assets by keyword or the table pagination to

browse all of your assets.

3. After you have a source and destination asset selected, generate the report.
4. After the visualization appears, click the **Export report** button. A window appears prompting you to enter a name and download location for the file.
5. Enter a name for the file and choose where you want to save it on your computer.
6. Save the file.

FAQs

Why are there hops in the report that aren't in my inventory?

Not every hop is a runZero asset. runZero fills in IP addresses for some asset hops, based on information that is pulled out of the traceroute data. Sometimes, runZero is able to extract asset information from the traceroute data that isn't part of the inventory, which is why you may be seeing them in the report.

Can runZero determine how two assets are talking to each other?

No, runZero can only identify potential paths—not if they are routable. runZero does not test or validate the paths.

Why isn't the asset route pathing report available?

The asset route pathing report is only available for runZero Platform customers.

Coverage reports

Coverage reports help you understand potential blind spots on your network by identifying which IP spaces have been scanned, which ones contain assets, and which ones still are unknown. With this information, you can find things like missing subnets, rogue devices, and misconfigurations.

To access the [coverage reports](#), go to **Reports** on the main menu and scroll down to **RFC 1918 coverage**.

RFC1918 coverage report

The RFC1918 coverage report helps you better track and identify the subnets that are in use on your internal network, the ones that have been scanned, and the ones that haven't been scanned.

TCP/IP version 4 reserves three ranges of IP addresses for private use. Specified in [RFC1918](#), they are:

CIDR	Address range	Number of addresses
10.0.0.0/8	10.0.0.0 - 10.255.255.255	16,777,216
172.16.0.0/12	172.16.0.0 - 172.31.255.255	1,048,576
192.168.0.0/16	192.168.0.0 - 192.168.255.255	65,536

Most companies use these address ranges for their internal IPv4 networks, connecting them to the Internet via Network Address Translation (NAT).

To help you visualize and assess your RFC1918 coverage, the RFC1918 report includes:

- [Coverage maps](#)
- [Subgrids](#)
- [Statistics](#)

RFC1918 coverage maps

A common network security and administration goal is to scan all of the available private IP addresses, to detect which subnetworks are in use on your internal network. Because of the large number of addresses, this can take a long time, leaving the problem of tracking which addresses have been scanned and which have not.

To solve this problem, runZero's Coverage report shows a graphical map of the RFC1918 private address spaces, showing which pieces have been scanned and to what percentage of completion.

runZero will sometimes detect that a device has additional IP addresses which are not part of the range being scanned. This can indicate that the device is present on an unscanned part of the private IP address space. The coverage reports show this by drawing a red border around the appropriate grid cell.

You can hover the mouse cursor over a cell to see a tooltip showing the CIDR address range the cell represents, how many unscanned hosts are believed to be in that range, and what percentage of the entire range has been scanned.

RFC1918 sub-grids

For the 192.168.0.0/16 map, each cell on the grid represents a /24 (256 addresses). Clicking a cell will take you to the subnet analysis report for that range and list the assets found.

For a map that shows a large address range, such as 10.0.0.0/8, each cell represents an entire /16 range of 65,536 addresses. To help narrow down the search for assets and unscanned hosts, you can click on any cell that represents a /16 range to go a grid map that showing just that range. From a /16 grid, you can use the link at top right to go back to the full range map.

On the /16 sub-grids, each cell is a /24, so clicking one takes you to the subnet analysis for that specific cell's address range, like on the 192.168.0.0/16 map.

Statistics

At the top of the coverage report page, you can see statistics showing how much of the RFC1918 address space you have scanned. Another box breaks the coverage down by the three blocks of reserved addresses.

Clicking the magnifying glass icon in the summary box will create a sample scan task, covering the unscanned address ranges. The refresh buttons create scan tasks to rescan all of the appropriate range.

Site and organization comparison

Platform

The [Site and organization comparison feature](#) lets you generate a side-by-side analysis of two sites, so you can understand:

- [How assets change over time](#) such as their TCP/UDP services, TCP/UDP ports, and service protocols. You can leverage this data to evaluate historical changes for assets for a specific point in time.
- [How exposure changes](#) based on scanning your network from different locations. For example, if you use public IP addresses internally and externally, you may want to scan those addresses from inside and outside your network to understand your potential exposure.

The report provides a summary view of differences. It only captures certain attributes that were added or removed from an asset, such as IP addresses, TCP ports, TCP service counts, UDP ports, UDP service counts, service protocols, and service counts. It does not track every modification to an asset, such as fingerprint or service banner changes.

After you run the report, the data presented in the report will be static. Any changes to your current inventory may result in assets no longer being accessible from the report.

Generate a site comparison

Generating the site comparison requires selecting a current site and a comparison site. The sites can be in different organizations. You can also select “All sites” to compare all sites in an organization with a different site and organization.

When the report runs it will assemble the two sets of assets specified, compare them using runZero’s asset matching algorithms, and generate a set of differences. You can then browse the summarized results in the report.

To generate a site comparison:

1. Verify that your current organization is one containing an inventory of assets you want to compare.
2. Go to the [site comparison page](#).
3. When the site comparison configuration page appears, the current organization will be set to the one you currently have selected. You can change the current site, if needed. For the comparison, choose the organization and site you want to run the analysis against.
4. Run the report. A task will be created to perform the comparisons, and you will be taken to the task page.
5. When the task is complete, the report will appear in the list of recent analysis reports at the top of the Reports page. You can then [view and search the results](#).

View how assets change over time

To analyze how assets have changed over time, you can compare the data from an old scan task with your most recent inventory. Setting a point-in-time comparison requires creating a new project that you can import your old scan data into.

Here's how you can set up a point-in-time analysis:

1. Go to the organization or project that contains the task scan data you'd like to use.
2. Go to your [completed tasks](#) and locate the task that contains the data for the point in time you'd like to compare.
3. From the task page, download the task data. It will be in a file with a name starting scan_ and ending .json.gz. This is the file you'll import into your new project. You don't need to uncompress the file, unless you're curious to look at the JSON data.
4. Next, [create a new project](#) for your import. You can create an organization if you intend to perform the analysis regularly.
5. After you create the project, go to the Inventory page and [import your scan task data](#) into it.

Now, you're ready to compare your current inventory with a previous version of it. Go to the [site comparison page](#). You'll need to select the organization and site for your current inventory as the site to compare against.

After the report runs, it shows a table that [highlights the differences](#) between the two sites, which in this case will represent two points in time, going from past to present.

You can also run the comparison by selecting your current organization first, and then choosing the project with the past data as the comparison site. The results will be the same, but with their sense reversed – that is, services which show as added when going from past to present, will show as removed when going from present to past.

View how exposure differs between networks

You can run the site comparison report to compare how exposure varies based on where you scan your network from. By comparing two inventories from two different perspectives, you can obtain better insights on your attack surface, which can help with active defense or risk reduction. For example, if you use external IPs internally and externally, you may want to scan those addresses from inside your network and outside your network. Then, you can run the site comparison report to compare the results from those two sites.

Here's how you can set up a diff for exposures between networks:

1. Set up a site with an Explorer on one network.
2. Set up another site with an Explorer on a different network. This site can be in the same organization as the first site.
3. Run a scan of the same address range with each Explorer. Verify you have the correct site selected for each scan.
4. After the scans complete, run the [site comparison](#) to generate the diff. The report will show a table that [highlights the differences](#) between the two sites.

Analyze the results in the site comparison report

The site comparison generates a table to show how assets' addresses, names, services, ports, and protocols differ between sites. Red text, denoted by the minus (-) sign, indicates that attributes were removed going from left to right. Green text, denoted by the plus (+) sign, indicates attributes were added.

Here are some things you should know about the results:

- **Address and Other Address** - The first address column contains the asset's addresses for the organization currently selected when the report was requested. The second address column, *Other Address*, contains the asset's addresses for the organization and site that was compared against the current organization.
- **Name and Other Name** - The first name column contains the asset's names in the organization currently selected when the report was generated. The second address column, *Other Name*, contains the asset's names for the organization and site compared against.
- **TCP and UDP services** - These columns show how the total number of TCP and UDP services differ between sites.
- **TCP and UDP ports** - These columns show the ports that have been added or removed between sites.

You can click on the green info (i) icon to view a more detailed comparison of the asset.

Clicking on one of the asset addresses in the report will bring up the full current asset record, if the asset still exists in the appropriate organization and site.

Search the site comparison report

You can search the report using the runZero [search query language](#). In the following descriptions, the main set refers to the assets in the organization that was current when you generated the report (i.e., the address and name columns). The comparison set refers to the assets in the organization and site that you chose to compare against (i.e., the other address and other name columns).

Keyword	Meaning
address:	Search for assets in the main set with a specified IP address.
net: or cidr:	Filter assets by their network CIDR range in the main set.
other_address:	Search for assets in the comparison set with a specified IP address.
other_net: or other_cidr:	Filter assets by their network CIDR range in the comparison set.
id:	Search by asset ID in the main set.

tcp:	Search for a TCP port change by number.
udp:	Search for a UDP port change by number.
protocol:	Search for a TCP or UDP port change by service name.

Organization Overview Report

Professional

Community

Platform

The Organization Overview Report captures a point-in-time snapshot of the asset data within your organization and sites. The report organizes data from your asset inventory into relevant sections and summarizes the major findings. The Organization Overview Report is useful for sharing with teams and leaders who may not have access to runZero and need an at-a-glance look into their network.

The report helps you quickly assess high-level metrics across multiple categories for your organization and sites, such as your asset types, operating systems, hardware, protocols, and products. The report also includes a summary of your RFC 1918 coverage, subnet utilization, and switches. For organizations with less than 50,000 assets, you can include additional information from your inventory via asset details and screenshots.

You can schedule the report to run at a specified start time and at a specified frequency. This report can run once, daily, weekly, monthly, or on the Nth weekday of each month. When a report is configured to run recurrently, one or more email addresses can be notified when the report has run.

runZero generates this report in HTML, but you can use your browser's Print to PDF feature to save the file as a PDF for distribution.

Generating the Organization Overview Report

1. Go to the [Reports page](#).
2. Open the [Organization Overview Report configuration form](#).
3. Enable **Include asset details** if you want additional information about each asset. The asset details will include the asset type, operating system, hardware, outlier score, first seen date, last seen date, site, addresses, names, MAC addresses, protocols, products, and services. This option only applies to organizations with less than 50,000 assets.
4. Enable **Include screenshots** if you want to see captured screenshots in the report. This option only applies to organizations with less than 50,000 assets.
5. Configure **Start time** and **Report frequency** if you want the report to run at a specified time or on a recurring schedule.
6. Provide one or more email addresses in the **Email addresses** field if you want email notifications when the report is run.
7. Generate the report. When the report generation completes, it automatically opens in your browser. Save the report as a PDF to share and distribute as needed.

Email notifications

You can include a comma-separated list of email addresses when configuring the report to run on a schedule. The specified email addresses will receive notifications when the report has finished each run of the configured schedule. The notification email will contain a link to

the report in the runZero Console. Please note that only registered users with access to the relevant organization will be able to access the completed report.

External Asset Report

Platform

The External Asset Report captures a point-in-time snapshot of the external asset data within your organization and sites. The report organizes data from your asset inventory into relevant sections and summarizes the major findings. The External Asset Report is useful for sharing with teams and leaders who may not have access to runZero and need an at-a-glance look into their public-facing assets.

The report helps you quickly assess high-level metrics across multiple categories for external assets in your organization and sites, such as asset types, operating systems, hardware, protocols, and products. The report also includes a summary of top Certificate Authorities and GeolP countries. For organizations with less than 50,000 assets, you can include additional information from your inventory via asset details and screenshots.

You can schedule the report to run at a specified start time and at a specified frequency. This report can run once, daily, weekly, monthly, or on the Nth weekday of each month. When a report is configured to run recurrently, one or more email addresses can be notified when the report has run.

runZero generates this report in HTML, but you can use your browser's Print to PDF feature to save the file as a PDF for distribution.

Generating the External Asset Report

1. Go to the [Reports page](#).
2. Open the [External Asset Report configuration form](#).
3. Enable **Include asset details** if you want additional information about each asset. The asset details will include the asset type, operating system, hardware, outlier score, first seen date, last seen date, site, addresses, names, MAC addresses, protocols, products, and services. This option only applies to organizations with less than 50,000 assets.
4. Enable **Include screenshots** if you want to see captured screenshots in the report. This option only applies to organizations with less than 50,000 assets.
5. Enable **Include unscanned subnet details** if you want to see results from external subnets that have not been scanned by runZero.
6. Enable **Include TLS certificate details** if you want to see details gathered about TLS certificates.
7. Use the **External IP exclusions** if you are using public IP address space internally and do not want those assets included in the report.
8. Configure **Start time** and **Report frequency** if you want the report to run at a specified time or on a recurring schedule.
9. Provide one or more email addresses in the **Email addresses** field if you want email notifications when the report is run.
10. Generate the report. When the report generation completes, it automatically opens in your browser. Save the report as a PDF to share and distribute as needed.

Email notifications

You can include a comma-separated list of email addresses when configuring the report to run on a schedule. The specified email addresses will receive notifications when the report has finished each run of the configured schedule. The notification email will contain a link to the report in the runZero Console. Please note that only registered users with access to the relevant organization will be able to access the completed report.

Exporting data

Exporting asset data

The [inventory](#) view provides a few ways to export asset data. The Export menu offers Export All options in both CSV and JSON format, and when a [search query](#) has been provided, options to export just the search results as both CSV and JSON.

The CSV format can be opened with tools like Microsoft Excel and easily imported into other applications but does not contain the full details of certain fields, such as Services. The JSON format contains a complete export but may take additional processing to use with other tools.

If you have a specific export/import scenario in mind, please contact [support](#) and let us know. The [Export API](#) leverages both the CSV and JSON formats and supports arbitrary search queries in the [same syntax](#) as the inventory.

Exporting HP iLO data

Community Platform

HP Integrated Lights-Out (iLO) provides remote management, configuration, and monitoring capabilities for HP servers. These capabilities centralize operations for your server environments and streamline tasks like rebooting servers, booting into single user mode, and bypassing authentication.

Being able to identify and find the serial number for HP iLO devices is useful for tracking warranties for support and contract management. If you have runZero Enterprise, you can export the runZero data as a CSV to feed into warranty tracking tools.

How to export HP iLO CSV data in runZero Enterprise

In your runZero Console, go to your [inventory](#). From the **Export** menu, choose the **HP iLO CSV format**. This method downloads all HP iLO data from the runZero inventory to a CSV file.

If you want to refine the results in your exported data, you can filter the inventory first. For example, if you only want to export iLOs that have the ProLiant DL360p Gen8 hardware, you can query:

```
alive:t AND ilo.hardware:"=ProLiant DL360p Gen8"
```

And then, from the Export menu, choose the **HP iLO CSV format** from the Export Search Results submenu.

HP iLO CSV export data

When you export HP iLO data, the CSV file will contain the following fields:

- IP address
- MAC address
- Name
- Serial number
- Product ID
- Model
- Version
- Health status
- IRS
- iLO product name
- Serial number
- Boot block
- HW revision
- FW revision
- PWRM
- Auth local

- Auth Kerberos
- Auth LDAP
- Auth license
- Security message
- Alive
- ID
- Site ID
- Last seen

There is a lot of data in the export that you may not need for your warranty tracker. Fields, such as serial numbers, physical hardware information, health status, and firmware version may be the most useful to import into your tracker.

Additional resources

Leveraging the API

runZero provides three primary APIs as well as integration-specific endpoints:

- The Export API provides read-only access to a specific organizations.
- The Organization API provides read-write access to a specific organizations (Professional and Platform licenses).
- The Account API provides read-write access to all account settings and organizations (Platform license).

To get started, you will need an API key / token or API client credentials.

API keys and tokens

Export and Organization APIs

Export API tokens and Organization API keys can be generated from the Organizations section in the runZero web console. You must have administrator access to generate API keys.

From the Organizations page, click on the desired organization to view its details page. Then, click **Edit organization** and scroll down to the export tokens or API tokens section to locate the **Generate** button for the type of API key you want to generate. Clicking the button will generate a secure API key in the form of a long random token.

Account API

Platform

Account API keys require an Platform license and are generated from the [Account settings page](#). To use an Account API key or token with the Organization or Export API, specify the additional parameter `_oid=[organization-id]` in the query parameters.

API client credentials

Account API client credentials are managed from the [API clients page](#). Your REST client should use the OAuth 2.0 authorization type and Client Credentials grant type. See the [OpenAPI specification](#) for the access token details. [Register an API client](#) and use the client ID and client secret to generate an access token as shown in the following example: `curl -X POST -H "Content-Type: application/x-www-form-urlencoded" -d "grant_type=client_credentials&client_id=<CLIENT_ID>&client_secret=<CLIENT_SECRET>" https://console.runzero.com/api/v1.0/account/api/token`

Once you have generated an API key or access token, your REST client should use it with [the Authorization: Bearer standard header](#) to authenticate.

API rate limiting

API calls are rate limited. You can make as many API calls per day as you have licensed assets in your account. For example, if you have 1,000 licensed assets, you can make 1,000 API calls per day. Each API call returns rate limit information in the HTTP headers of the response:

- X-API-Usage-Total - Total number of calls made to the API
- X-API-Usage-Today - Number of calls made to the API today
- X-API-Usage-Limit - Your daily API call limit, shared across all API keys
- X-API-Usage-Remaining - The number of API calls remaining from your daily limit

In addition, there's a limit of 2,000 requests per 5 minutes per source IP address. If this is exceeded, an HTTP 429 response will be sent with an error of "request limit exceeded", and your client code should delay the next request.

Additional documentation

Please see the [Swagger documentation](#) and [runZero OpenAPI specification](#) for details on the individual API calls.

Using the standalone command-line scanner

Platform

Some editions of runZero include the standalone command-line [scanner](#) that can be used to perform network discovery without access to the internet. The scanner has the same options and similar performance characteristics to the Explorer. The scanner output file named `scan.rumble.gz` can be uploaded to the runZero Console through the Inventory Import menu.

The scanner works best with root privileges on Linux/macOS and Administrator privileges on Windows. Although the scanner will function without privileged access, many probe types will be unavailable. The `sudo` command can be used to run the scanner as root on Linux and macOS, while the tool is best run from an elevated command shell on Windows. On the Windows platform, the runZero Scanner will look for an existing nmap installation and try to install it if the software is not found. This behavior can be disabled with the `--nopcap` flag.

Note: Some components of the application still reference the name "Rumble" for backwards compatibility. The documentation will be updated as these are changed.

The runZero Scanner defaults to a semi-interactive terminal interface that writes multiple output files to a directory. The default directory name is `rumble-[current-date]`. To switch to plain text output, use the `--text` option. To skip artifact generation and only produce the raw JSON output file, use the flags `--text -o disable --output-raw scan.rumble`.

Input can be provided as arguments on the command-line or by specifying an input file using the `--input` (or `-i`) parameter. Input can consist of specific IPv4 addresses or IPv4 CIDRs. Supported formats include:

- `10.0.0.1`
- `10.0.0.0/24`
- `10.0.0.0/255.255.255.0`
- `10.0.0.1-10.0.0.255`
- `example.com`
- `example.com/24`

For hostnames, each IPv4 address in the response will be expanded with the optional mask.

The example below downloads and runs the scanner on a Linux x86_64 host. This URL will be different for your installation. The current download links for your organization are available from the [scanner](#) page of the runZero Console.

```
$ wget https://console.runzero.com/download/scanner/[unique-link]/runzero-scanner-linux-amd64.b
$ chmod +x runzero-scanner-linux-amd64.bin
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 -o output-dir
```

Please note that the hexadecimal values in the download URL are specific for your account and organization.

Performance

The default speed of runZero scans is limited to 1,000 packets per second with a single pass. This setting works great for reliable wired networks without stateful firewalls between the scanning system and the destination networks. This rate can be changed via the `--rate` (or `-r`) option, with a reasonable maximum being 10000 for most networks. On slow unreliable networks, a rate of 300 with `--passes` set to 3 may provide better results.

A second parameter, `--max-host-rate` limits how many packets are sent per second to each individual host. This defaults to 40, which is low, but may be necessary when scanning low-power embedded devices. In cases where a small number of hosts (or a single host) should be scanned quickly, the `--max-host-rate` parameter can be increased to match the `--rate`.

Examples

The following example demonstrates a scan of 65,535 TCP ports on all hosts in the 192.168.0.0/24 subnet running at 10,000 packets per second:

```
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 -r 10000 --tcp-ports 1-65535 -o output-dir
```

The following example demonstrates a scan on all hosts in the 192.168.0.0/24 and 10.0.0.0/24 subnets running at 5,000 packets per second:

```
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 10.0.0.0/24 -r 5000 -o output-dir
```

The following example demonstrates a scan on all hosts in the 192.168.0.0/24 and 10.0.0.0/8 subnets running at a max host rate of 20 packets per host:

```
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 10.0.0.0/8 --max-host-rate 20 -o output-dir
```

The following example demonstrates a scan on all hosts in the 192.168.0.0/24 subnet and the domain "example.com" running at 7,500 packets per second:

```
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 example.com -r 7,500 -o output-dir
```

The following example demonstrates a scan on all hosts in the 10.0.0.0/8 subnet and a particular ASN4 value at a default speed of 1,000 packets per second.

```
$ sudo ./runzero-scanner-linux-amd64.bin 10.0.0.0/8 asn4:[ID] -o output-dir
```

The following example demonstrates a scan on all hosts in the 192.168.0.0/24 subnet with the max TTL set at 128 and a scan rate of 2,500 packets per second:

```
$ sudo ./runzero-scanner-linux-amd64.bin 192.168.0.0/24 -r 2,500 --max-ttl 128 -o output-dir
```

The following example demonstrates a scan based on an input file:

```
$ sudo ./runzero-scanner-linux-amd64.bin -i /path/to/input-file.txt -o output dir
```

Here is an example input file:

```
www.example.com  
192.168.0.0/24
```

Automatic web screenshots

The `--screenshots` option defaults to `true` and tells runZero to obtain a screenshot of all web services identified during the scan. This feature depends on the system running the Explorer having a local installation of the Google Chrome or Chromium browsers. The acquired screenshots will be reported as a base64 string, stored in the `"screenshot.image"` field of the containing service scan result.

To disable automatic web screenshots, set the `--screenshots` option to `false` (`--screenshots=false`).

runZero scanner commands and options

The runZero Scanner supports a wide range of options, including the ability to limit scans to specific ports, probes, and SNMP communities. The `--help` output provides basic documentation on the available options.

Commands

`censys`

Import Censys data files

`completion`

Generate the autocompletion script for the specified shell

`help`

Help about any command

`license`

Display license information

`upgrade`

Upgrade to the latest version of the runZero Scanner

`verify`

Perform an internal signature verification

`version`

Print the version number of runZero

Flags

- `--api-key string`
Specify the runZero API key
- `--api-no-verify`
Disable TLS verification for API communication
- `--api-url string`
Specify the runZero API server hostname (default ["https://console.runzero.com/api/v1.0"](https://console.runzero.com/api/v1.0))
- `--arp-fast`
Enables fast mode by ARP scanning at the scan rate vs host rate
- `--aws-instances-access-key string`
The access key for the AWS account
- `--aws-instances-assume-role-name string`
The role to assume for all accounts in the organization for cross-account access
- `--aws-instances-delete-stale`
Automatically delete stale AWS assets
- `--aws-instances-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--aws-instances-include-stopped`
Include assets that are not currently running
- `--aws-instances-regions string`
The comma-separated list of regions for the AWS account
- `--aws-instances-secret-access-key string`
The secret access key for the AWS account
- `--aws-instances-service-options string`
The comma-separated list of services to sync data from (defaults,ec2,elb,elbv2,rds,lambda) (default "defaults")
- `--aws-instances-site-per-account`
Automatically create a new site per account
- `--aws-instances-site-per-vpc`
Automatically create a new site per VPC
- `--aws-instances-token string`
The session token for the AWS account
- `--azure-client-id string`
The application ID (client ID) for the Azure account
- `--azure-client-secret string`
The client secret for the Azure account

- azure-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- azure-multi-subscription
Access all subscriptions in the directory (tenant) for the Azure account
- azure-password string
The password for the Azure account
- azure-service-options string
The comma-separated list of services to sync data from (defaults,vm,vmss,azsql,cosmos,lb,functionapp) (default "defaults")
- azure-site-per-subscription
Automatically create a new site per subscription
- azure-subscription-id string
The subscription ID for the Azure account
- azure-tenant-id string
The directory ID (tenant ID) for the Azure account
- azure-username string
The username for the Azure account
- azuread-client-id string
The application ID (client ID) for the Azure account
- azuread-client-secret string
The client secret for the Azure account
- azuread-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- azuread-include-inactive
Include assets that are marked as inactive in the AzureAD account
- azuread-password string
The password for the AzureAD account
- azuread-service-options string
The comma-separated list of services to sync data from (defaults,dev,user,group) (default "defaults")
- azuread-tenant-id string
The directory ID (tenant ID) for the Azure account
- azuread-username string
The username for the AzureAD account
- bacnet-ports string
The destination ports for BACnet probes (default "46808,47808,48808")
- b, --baseline string
Use the specified file as an asset baseline for tracking
- bedrock-ports string
The destination ports for Bedrock probes (default "19132")

- `--censys-api-url` string
The API endpoint to use for Censys Search (default "<https://search.censys.io>")
- `--censys-client-id` string
The Client ID to use for Censys Search authentication
- `--censys-client-secret` string
The Client Secret to use for Censys Search authentication
- `--censys-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--censys-mode` string
The search mode (assets or query). The assets option queries the scan targets (default "assets")
- `--censys-query` string
The search string to use in query mode
- `--coap-port` uint
The destination port for CoAP probes (default 5683)
- `--cpu` string
Write a cpu profile after the scan completes
- `--crestron-port` uint
The destination port for Crestron probes (default 41794)
- `--crowdstrike-api-url` string
The URL used for the CrowdStrike account's API access
- `--crowdstrike-client-id` string
The client ID for the CrowdStrike account
- `--crowdstrike-client-secret` string
The client secret for the CrowdStrike account
- `--crowdstrike-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--crowdstrike-filter` string
An optional Falcon Query Language (FQL) filter for imported assets
- `--crowdstrike-fingerprint-only`
Import vulnerabilities for fingerprinting purposes only
- `--crowdstrike-risks` string
Minimum risk of imported vulnerabilities (None, Low, Medium, High, Critical) (default "None,Low,Medium,High,Critical")
- `--crowdstrike-severities` string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- `--dahua-dhip-ports` string
The destination ports for Dahua DHIP discovery probes (default "37810")

- defender365-client-id string
The application ID (client ID) for the Azure account
- defender365-client-secret string
The client secret for the Azure account
- defender365-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- defender365-include-inactive
Include assets that have stopped reporting to the Microsoft 365 Defender service
- defender365-tenant-id string
The directory ID (tenant ID) for the Azure account
- disabled-probes string
Specifically exclude these probes, comma-delimited
- dnp3-address-probe-timeout int
Time limit (in seconds) for DNP address discovery. (default 30)
- dnp3-banner-address-discovery string
One of 'require', 'prefer', or 'ignore'. (default "ignore")
- dnp3-destination-address-discovery-range string
A numeric range of addresses to attempt to discover. (default "0-32")
- dnp3-explorer-address int
Source DNP3 address for the explorer. (default -1)
- dns-disable-google-myaddr
Disables resolution of upstream DNS via Google myaddr service
- dns-disable-meraki-detection
Disables detection of Meraki DNS interception
- dns-port uint
The destination port for DNS probes (default 53)
- dns-resolve-name string
The target hostname for DNS queries ('off' to disable) (default "www.google.com")
- dns-trace-domain string
The subdomain to use for trace requests ('off' to disable) (default "helper.rumble.network")
- dtls-ports string
The destination ports for DTLS probes (default "443,3391,4433,5246,5349,5684")
- echo-report-errors
Report errors from intermediate in-scope hosts
- ethernetip-udp-ports string
The destination ports for EtherNet/IP UDP probes (default "44818")
- ethernetip-use-tagged-context string
Set this to true to use a runZero-specific sender context for debugging. (default "false")

- `--exclude string`
Specify scan exclusions
- `--excludefile string`
Read exclusions from an input file
- `--filter-base64`
Filter base64-encoded fields
- `-f, --fingerprints string`
Use the specified directory as an alternate fingerprint database
- `--fingerprints-debug`
Enable debug output for the fingerprint processor
- `--fins-port uint`
The destination port for FINS probes (default 9600)
- `--gcp-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--gcp-key-path string`
Path to GCP service account key file
- `--gcp-service-options string`
The comma-separated list of services to sync data from (defaults,vm,lb,cloudsql)
(default "defaults")
- `--gcp-site-per-project`
Automatically create a new site per project
- `--genudp-payload-base64 string`
The generic udp payload as base64
- `--genudp-payload-hex string`
The generic udp payload as hexadecimal
- `--genudp-payload-text string`
The generic udp payload as plain text
- `--genudp-ports string`
The destination ports for the generic udp probe
- `--googleworkspace-client-email string`
The email address of the service account
- `--googleworkspace-client-id string`
The ID of the service account
- `--googleworkspace-customer-id string`
An optional customer ID for multi-tenant environments (default "my_customer")
- `--googleworkspace-delegate string`
The email address of an admin account with directory access
- `--googleworkspace-exclude-unknown`
Exclude assets that cannot be merged into an existing asset

- googleworkspace-private-key string
The PEM encoded private key
- googleworkspace-private-key-id string
The ID of the private key
- googleworkspace-project-id string
The project ID of the service account
- googleworkspace-service-options string
The comma-separated list of services to sync data from
(defaults,chromeos,mobile,endpoint,user,group) (default "defaults")
- goroutines string
Write a goroutine dump after the scan completes
- heap string
Write a heap profile after the scan completes
- h, --help
help for runZero
- hiddiscoveryd-port uint
The destination port for HID discoveryd probes (default 4070)
- host-ping
Only scan hosts that respond to a ping scan using the host-ping settings
- host-ping-arp-fast
Enables fast mode by ARP scanning at the scan rate vs host rate (host ping)
- host-ping-aws-instances-access-key string
The access key for the AWS account (host ping)
- host-ping-aws-instances-assume-role-name string
The role to assume for all accounts in the organization for cross-account access (host ping)
- host-ping-aws-instances-delete-stale
Automatically delete stale AWS assets (host ping)
- host-ping-aws-instances-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-aws-instances-include-stopped
Include assets that are not currently running (host ping)
- host-ping-aws-instances-regions string
The comma-separated list of regions for the AWS account (host ping)
- host-ping-aws-instances-secret-access-key string
The secret access key for the AWS account (host ping)
- host-ping-aws-instances-service-options string
The comma-separated list of services to sync data from
(defaults,ec2,elb,elbv2,rds,lambda) (host ping) (default "defaults")

- host-ping-aws-instances-site-per-account
Automatically create a new site per account (host ping)
- host-ping-aws-instances-site-per-vpc
Automatically create a new site per VPC (host ping)
- host-ping-aws-instances-token string
The session token for the AWS account (host ping)
- host-ping-azure-client-id string
The application ID (client ID) for the Azure account (host ping)
- host-ping-azure-client-secret string
The client secret for the Azure account (host ping)
- host-ping-azure-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-azure-multi-subscription
Access all subscriptions in the directory (tenant) for the Azure account (host ping)
- host-ping-azure-password string
The password for the Azure account (host ping)
- host-ping-azure-service-options string
The comma-separated list of services to sync data from
(defaults,vm,vmss,azsql,cosmos,lb,functionapp) (host ping) (default "defaults")
- host-ping-azure-site-per-subscription
Automatically create a new site per subscription (host ping)
- host-ping-azure-subscription-id string
The subscription ID for the Azure account (host ping)
- host-ping-azure-tenant-id string
The directory ID (tenant ID) for the Azure account (host ping)
- host-ping-azure-username string
The username for the Azure account (host ping)
- host-ping-azuread-client-id string
The application ID (client ID) for the Azure account (host ping)
- host-ping-azuread-client-secret string
The client secret for the Azure account (host ping)
- host-ping-azuread-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-azuread-include-inactive
Include assets that are marked as inactive in the AzureAD account (host ping)
- host-ping-azuread-password string
The password for the AzureAD account (host ping)
- host-ping-azuread-service-options string
The comma-separated list of services to sync data from (defaults,dev,user,group) (host ping) (default "defaults")

- host-ping-azuread-tenant-id string
The directory ID (tenant ID) for the Azure account (host ping)
- host-ping-azuread-username string
The username for the AzureAD account (host ping)
- host-ping-bacnet-ports string
The destination ports for BACnet probes (host ping) (default "46808,47808,48808")
- host-ping-bedrock-ports string
The destination ports for Bedrock probes (host ping) (default "19132")
- host-ping-censys-api-url string
The API endpoint to use for Censys Search (host ping) (default "<https://search.censys.io>")
- host-ping-censys-client-id string
The Client ID to use for Censys Search authentication (host ping)
- host-ping-censys-client-secret string
The Client Secret to use for Censys Search authentication (host ping)
- host-ping-censys-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-censys-mode string
The search mode (assets or query). The assets option queries the scan targets (host ping) (default "assets")
- host-ping-censys-query string
The search string to use in query mode (host ping)
- host-ping-coap-port uint
The destination port for CoAP probes (host ping) (default 5683)
- host-ping-crestron-port uint
The destination port for Crestron probes (host ping) (default 41794)
- host-ping-crowdstrike-api-url string
The URL used for the CrowdStrike account's API access (host ping)
- host-ping-crowdstrike-client-id string
The client ID for the CrowdStrike account (host ping)
- host-ping-crowdstrike-client-secret string
The client secret for the CrowdStrike account (host ping)
- host-ping-crowdstrike-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-crowdstrike-filter string
An optional Falcon Query Language (FQL) filter for imported assets (host ping)
- host-ping-crowdstrike-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)

- host-ping-crowdstrike-risks string
Minimum risk of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping) (default "None,Low,Medium,High,Critical")
- host-ping-crowdstrike-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping) (default "Info,Low,Medium,High,Critical")
- host-ping-dahua-dhip-ports string
The destination ports for Dahua DHIP discovery probes (host ping) (default "37810")
- host-ping-defender365-client-id string
The application ID (client ID) for the Azure account (host ping)
- host-ping-defender365-client-secret string
The client secret for the Azure account (host ping)
- host-ping-defender365-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-defender365-include-inactive
Include assets that have stopped reporting to the Microsoft 365 Defender service (host ping)
- host-ping-defender365-tenant-id string
The directory ID (tenant ID) for the Azure account (host ping)
- host-ping-dnp3-address-probe-timeout int
Time limit (in seconds) for DNP address discovery. (host ping) (default 30)
- host-ping-dnp3-banner-address-discovery string
One of 'require', 'prefer', or 'ignore'. (host ping) (default "ignore")
- host-ping-dnp3-destination-address-discovery-range string
A numeric range of addresses to attempt to discover. (host ping) (default "0-32")
- host-ping-dnp3-explorer-address int
Source DNP3 address for the explorer. (host ping) (default -1)
- host-ping-dns-disable-google-myaddr
Disables resolution of upstream DNS via Google myaddr service (host ping)
- host-ping-dns-disable-meraki-detection
Disables detection of Meraki DNS interception (host ping)
- host-ping-dns-port uint
The destination port for DNS probes (host ping) (default 53)
- host-ping-dns-resolve-name string
The target hostname for DNS queries ('off' to disable) (host ping) (default "www.google.com")
- host-ping-dns-trace-domain string
The subdomain to use for trace requests ('off' to disable) (host ping) (default "helper.rumble.network")

- host-ping-dtls-ports string
The destination ports for DTLS probes (host ping) (default "443,3391,4433,5246,5349,5684")
- host-ping-echo-report-errors
Report errors from intermediate in-scope hosts (host ping)
- host-ping-ethernetip-udp-ports string
The destination ports for EtherNet/IP UDP probes (host ping) (default "44818")
- host-ping-ethernetip-use-tagged-context string
Set this to true to use a runZero-specific sender context for debugging. (host ping) (default "false")
- host-ping-fins-port uint
The destination port for FINS probes (host ping) (default 9600)
- host-ping-gcp-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-gcp-service-options string
The comma-separated list of services to sync data from (defaults,vm,lb,cloudsql) (host ping) (default "defaults")
- host-ping-gcp-site-per-project
Automatically create a new site per project (host ping)
- host-ping-genudp-payload-base64 string
The generic udp payload as base64 (host ping)
- host-ping-genudp-payload-hex string
The generic udp payload as hexadecimal (host ping)
- host-ping-genudp-payload-text string
The generic udp payload as plain text (host ping)
- host-ping-genudp-ports string
The destination ports for the generic udp probe (host ping)
- host-ping-googleworkspace-client-email string
The email address of the service account (host ping)
- host-ping-googleworkspace-client-id string
The ID of the service account (host ping)
- host-ping-googleworkspace-customer-id string
An optional customer ID for multi-tenant environments (host ping) (default "my_customer")
- host-ping-googleworkspace-delegate string
The email address of an admin account with directory access (host ping)
- host-ping-googleworkspace-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-googleworkspace-private-key string
The PEM encoded private key (host ping)

- host-ping-googleworkspace-private-key-id string
The ID of the private key (host ping)
- host-ping-googleworkspace-project-id string
The project ID of the service account (host ping)
- host-ping-googleworkspace-service-options string
The comma-separated list of services to sync data from
(defaults,chromeos,mobile,endpoint,user,group) (host ping) (default "defaults")
- host-ping-hiddiscoveryd-port uint
The destination port for HID discoveryd probes (host ping) (default 4070)
- host-ping-igel-discovery-ports string
The destination ports for IGEL discovery probes (host ping) (default "30005")
- host-ping-ike-port uint
The destination port for IKE probes (host ping) (default 500)
- host-ping-insightvm-api-url string
The URL used for the InsightVM account's API access (host ping)
- host-ping-insightvm-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-insightvm-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)
- host-ping-insightvm-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (host ping) (default true)
- host-ping-insightvm-password string
The password for the InsightVM account (host ping)
- host-ping-insightvm-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping)
(default "None,Low,Medium,High,Critical")
- host-ping-insightvm-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping)
(default "Info,Low,Medium,High,Critical")
- host-ping-insightvm-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (host ping)
- host-ping-insightvm-username string
The username for the InsightVM account (host ping)
- host-ping-intune-client-id string
The application ID (client ID) for the Azure account (host ping)
- host-ping-intune-client-secret string
The client secret for the Azure account (host ping)
- host-ping-intune-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)

- host-ping-intune-password string
The password for the Intune account (host ping)
- host-ping-intune-tenant-id string
The directory ID (tenant ID) for the Azure account (host ping)
- host-ping-intune-username string
The username for the Intune account (host ping)
- host-ping-ipmi-port uint
The destination port for IPMI probes (host ping) (default 623)
- host-ping-kerberos-port uint
The destination port for kerberos probes (host ping) (default 88)
- host-ping-knxnet-ports string
The destination ports for knxnet probes (host ping) (default "3671")
- host-ping-l2t-port uint
The destination port for L2T probes (host ping) (default 2228)
- host-ping-l2tp-ports string
The destination ports for L2TP probes (host ping) (default "1701")
- host-ping-lantronix-port uint
The destination port for Lantronix probes (host ping) (default 30718)
- host-ping-layer2-add-targets
Set this false to skip scanning discovered targets (host ping) (default true)
- host-ping-layer2-force
Set this to true to force discovery even without local targets (host ping)
- host-ping-layer2-max-retries uint
The desired number of retries (host ping) (default 3)
- host-ping-layer2-tcp-ports string
The TCP ports to ping for local device discovery (host ping) (default "22,80,135,179,443,3389,5040,7547,62078")
- host-ping-layer2-udp-trace-port uint
The UDP port number to use for UDP trace requests (host ping) (default 9)
- host-ping-ldap-base-dn string
The base DN used for LDAP searches (host ping)
- host-ping-ldap-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-ldap-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (host ping)
- host-ping-ldap-legacy-tls
Set this to true to authenticate over legacy TLS versions (< 1.2) (host ping)
- host-ping-ldap-password string
The password for the LDAP account (host ping)

- host-ping-ldap-service-options string
The comma-separated list of services to sync data from (defaults,computer,user,group) (host ping) (default "defaults")
- host-ping-ldap-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (host ping)
- host-ping-ldap-url string
The URL used for the LDAP server (host ping)
- host-ping-ldap-username string
The username for the LDAP account (host ping)
- host-ping-max-attempts int
Set the maximum number of attempts for each probe (default 2)
- host-ping-max-ttl int
Set the default TTL on host-ping probe packets (default 255)
- host-ping-mdns-port uint
The destination port for MDNS probes (host ping) (default 5353)
- host-ping-memcache-port uint
The destination port for memcached probes (host ping) (default 11211)
- host-ping-miradore-api-key string
The API key for the Miradore account (host ping)
- host-ping-miradore-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-miradore-hostname string
The Miradore web console hostname (url) (host ping)
- host-ping-modbus-identification-level string
Identification level, one of 'basic', 'regular', or 'extended'. (host ping) (default "regular")
- host-ping-mssql-port uint
The destination port for MSSQL probes (host ping) (default 1434)
- host-ping-natpmp-port uint
The destination port for NATPMP probes (host ping) (default 5351)
- host-ping-nessus-access-key string
The access key for the Nessus Professional account (host ping)
- host-ping-nessus-api-url string
The URL used for the Nessus Professional account's API access (host ping)
- host-ping-nessus-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-nessus-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)
- host-ping-nessus-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (host ping) (default true)

- host-ping-nessus-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping)
(default "None,Low,Medium,High,Critical")
- host-ping-nessus-secret-key string
The secret key for the Nessus Professional account (host ping)
- host-ping-nessus-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping)
(default "Info,Low,Medium,High,Critical")
- host-ping-nessus-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (host ping)
- host-ping-netbios-port uint
The destination port for NetBIOS Name Service probes (host ping) (default 137)
- host-ping-ntp-port uint
The destination port for NTP probes (host ping) (default 123)
- host-ping-openvpn-ports string
The destination ports for OpenVPN probes (host ping) (default "1194")
- host-ping-passes int
Set the number of passes for the host-ping phase (default 1)
- host-ping-pca-port uint
The destination port for PCAnywhere probes (host ping) (default 5632)
- host-ping-probes string
Launch a subset of the probes for the host-ping, comma-delimited (default
"arp,echo,syn,connect,netbios,snmp,ntp,sunrpc,ike,openvpn,mdns")
- host-ping-psdisco-ports string
The destination ports for playstation discovery probes (host ping) (default "987,9302")
- host-ping-qualys-api-url string
The URL used for the Qualys account's API access (host ping)
- host-ping-qualys-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-qualys-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)
- host-ping-qualys-include-unscanned
Include assets that have not been assessed for vulnerabilities (host ping)
- host-ping-qualys-password string
The password for the Qualys account (host ping)
- host-ping-qualys-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping)
(default "None,Low,Medium,High,Critical")

- host-ping-qualys-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping) (default "Info,Low,Medium,High,Critical")
- host-ping-qualys-username string
The username for the Qualys account (host ping)
- host-ping-rdns-max-concurrent int
The maximum number of concurrent DNS lookups (host ping) (default 64)
- host-ping-rdns-timeout uint
The DNS PTR lookup timeout in seconds (host ping) (default 3)
- host-ping-rpcbind-port uint
The destination port for RPCBind probes (host ping) (default 111)
- host-ping-rpcbind-port-nfs uint
The destination port for NFS probes (host ping) (default 2049)
- host-ping-s7comm-request-extended-information
If true, request extended device information. (host ping)
- host-ping-sample-duration string
Specify the duration in seconds to sample network traffic (or '0' for non-stop) (host ping) (default "300")
- host-ping-sample-excludes string
Specify host exclusions (host ping)
- host-ping-sample-interfaces string
Specify a comma-separated list of network interfaces (or 'all' for everything) (host ping)
- host-ping-sample-targets string
Specify the discovery scope (host ping) (default "10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 169.254.0.0/16")
- host-ping-sentinelone-api-url string
The URL used for the SentinelOne account's API access (host ping)
- host-ping-sentinelone-client-id string
The client ID for the SentinelOne account (host ping)
- host-ping-sentinelone-client-secret string
The client secret for the SentinelOne account (host ping)
- host-ping-sentinelone-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-shodan-api-key string
The key used for the Shodan account's API access (host ping)
- host-ping-shodan-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-shodan-mode string
The search mode (assets or query). The assets option queries the scan targets (host ping) (default "assets")

- host-ping-shodan-query string
The search string to use in query mode (host ping)
- host-ping-sip-port uint
The destination port for SIP probes (host ping) (default 5060)
- host-ping-snmp-comms string
The comma-separated list of SNMP v1/v2c communities (host ping) (default "public,private")
- host-ping-snmp-disable-bulk
If true, do not use bulk walking operations (host ping)
- host-ping-snmp-max-repetitions uint
The maximum number of repetitions in a bulk walk operation (host ping) (default 16)
- host-ping-snmp-max-retries int
The maximum number of retries for an SNMP operation (host ping) (default 1)
- host-ping-snmp-poll-interval uint
The minimum number of seconds between polling each host after initial discovery (host ping) (default 300)
- host-ping-snmp-port uint
The destination port for SNMP probes (host ping) (default 161)
- host-ping-snmp-timeout uint
The maximum number of seconds for each individual SNMP operation (host ping) (default 5)
- host-ping-snmp-v3-auth-passphrase string
The authentication passphrase (host ping)
- host-ping-snmp-v3-auth-protocol string
The authentication protocol (none, md5, sha, sha224, sha256, sha384, sha512) (host ping) (default "none")
- host-ping-snmp-v3-context string
The optional SNMP v3 context to supply (host ping)
- host-ping-snmp-v3-privacy-passphrase string
The privacy passphrase (host ping)
- host-ping-snmp-v3-privacy-protocol string
The privacy protocol (none, des, aes, aes192, aes256, aes192c, aes256c) (host ping) (default "none")
- host-ping-snmp-v3-username string
The username to use for SNMP v3 authentication (host ping)
- host-ping-snmp-walk-timeout uint
The maximum number of seconds for each SNMP walk operation (host ping) (default 60)
- host-ping-ssdp-port uint
The destination port for UPnP/SSDP probes (host ping) (default 1900)

- host-ping-ssh-fingerprint
Enable fingerprinting using partial authentication (host ping) (default true)
- host-ping-ssh-fingerprint-username string
The username to use for partial authentication SSH fingerprinting (host ping) (default "STATUS")
- host-ping-steam-ports string
The destination ports for Steam discovery probes (host ping) (default "27036")
- host-ping-syn-disable-bogus-filter
Disable bogus service detection and filtering (host ping)
- host-ping-syn-forwarding-check
Perform an IP forwarding check as part of the scan (host ping) (default true)
- host-ping-syn-forwarding-check-target string
An external IPv4 address for the forwarding check (default:runzero) (host ping) (default "13.248.161.247")
- host-ping-syn-max-retries uint
The maximum number of retries trace and SYN requests (host ping) (default 2)
- host-ping-syn-report-resets
Set this to true to report RST responses (host ping) (default true)
- host-ping-syn-reset-sessions
Reset middle-box/firewall sessions automatically (host ping) (default true)
- host-ping-syn-reset-sessions-delay uint
Minimum delay in milliseconds between a SYN and a session reset (host ping)
- host-ping-syn-reset-sessions-limit uint
Maximum number of in-flight sessions before forcing session resets (host ping) (default 50)
- host-ping-syn-traceroute
Perform a multi-protocol traceroute as part of the scan (host ping) (default true)
- host-ping-syn-udp-trace-port uint
The UDP port number to use for UDP trace requests (host ping) (default 9)
- host-ping-tcp-ports string
The list of TCP ports to host-ping using the syn and connect probes (default "22,80,135,179,443,3389,5040,7547,62078")
- host-ping-tenable-access-key string
The access key for the Tenable.io account (host ping)
- host-ping-tenable-api-url string
The URL used for the Tenable.io account's API access (host ping)
- host-ping-tenable-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-tenable-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)

- host-ping-tenable-include-unscanned
Include assets that have not been assessed for vulnerabilities (host ping)
- host-ping-tenable-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping)
(default "None,Low,Medium,High,Critical")
- host-ping-tenable-secret-key string
The secret key for the Tenable.io account (host ping)
- host-ping-tenable-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping)
(default "Info,Low,Medium,High,Critical")
- host-ping-tenablesecuritycenter-access-key string
The access key for the Tenable Security Center account (host ping)
- host-ping-tenablesecuritycenter-api-url string
The URL used for the Tenable Security Center account's API access (host ping)
- host-ping-tenablesecuritycenter-batch-size string
The number of records to request at a time. (between 2000 and 10000) (host ping)
(default "2000")
- host-ping-tenablesecuritycenter-exclude-unknown
Exclude assets that cannot be merged into an existing asset (host ping)
- host-ping-tenablesecuritycenter-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (host ping)
- host-ping-tenablesecuritycenter-insecure string
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (host ping)
- host-ping-tenablesecuritycenter-query-id string
The ID of an existing vulnerability query in the Tenable Security Center account (host ping)
- host-ping-tenablesecuritycenter-query-mode string
Set to 'filters' to provide 'severities' and 'risks' values to import. Set to 'query-id' to provide a value for 'query-id'. (host ping) (default "filters")
- host-ping-tenablesecuritycenter-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (host ping)
(default "None,Low,Medium,High,Critical")
- host-ping-tenablesecuritycenter-secret-key string
The secret key for the Tenable Security Center account (host ping)
- host-ping-tenablesecuritycenter-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (host ping)
(default "Info,Low,Medium,High,Critical")
- host-ping-tenablesecuritycenter-sync-since string
Specify an initial date to sync data from. (host ping)

- host-ping-tenablesecuritycenter-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (host ping)
- host-ping-tftp-ports string
The destination ports for TFTP probes (host ping) (default "69")
- host-ping-tos int
Set the default ToS on host-ping probe packets
- host-ping-ubnt-port uint
The destination port for Ubiquiti probes (host ping) (default 10001)
- host-ping-verbose
Display verbose output for the host-ping
- host-ping-very-verbose
Display very verbose output for the host-ping
- host-ping-vmware-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (host ping) (default true)
- host-ping-vmware-password string
The password to use for VMware SDK authentication (read-only) (host ping)
- host-ping-vmware-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (host ping)
- host-ping-vmware-username string
The username to use for VMware SDK authentication (read-only) (host ping)
- host-ping-webmin-ports string
The destination ports for webmin probes (host ping) (default "10000")
- host-ping-wlan-list-poll-interval uint
The minimum number of seconds between polls of the access point list (host ping) (default 300)
- host-ping-wsd-port uint
The destination port for WSD probes (host ping) (default 3702)
- igel-discovery-ports string
The destination ports for IGEL discovery probes (default "30005")
- ike-port uint
The destination port for IKE probes (default 500)
- I, --import stringArray
Import existing scan data from the specified input files ('scan.rumble' format)
- import-pcap stringArray
Import pcap packet capture from the specified input files ('.pcap' or '.pcapng' format)
- i, --input-targets string
Read scan targets from the specified input file
- insightvm-api-url string
The URL used for the InsightVM account's API access

- insightvm-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- insightvm-fingerprint-only
Import vulnerabilities for fingerprinting purposes only
- insightvm-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (default true)
- insightvm-password string
The password for the InsightVM account
- insightvm-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (default "None,Low,Medium,High,Critical")
- insightvm-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- insightvm-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication
- insightvm-username string
The username for the InsightVM account
- intune-client-id string
The application ID (client ID) for the Azure account
- intune-client-secret string
The client secret for the Azure account
- intune-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- intune-password string
The password for the Intune account
- intune-tenant-id string
The directory ID (tenant ID) for the Azure account
- intune-username string
The username for the Intune account
- ipmi-port uint
The destination port for IPMI probes (default 623)
- kerberos-port uint
The destination port for kerberos probes (default 88)
- knxnet-ports string
The destination ports for knxnet probes (default "3671")
- l2t-port uint
The destination port for L2T probes (default 2228)

- l2tp-ports string
The destination ports for L2TP probes (default "1701")
- lantronix-port uint
The destination port for Lantronix probes (default 30718)
- layer2-add-targets
Set this false to skip scanning discovered targets (default true)
- layer2-force
Set this to true to force discovery even without local targets
- layer2-max-retries uint
The desired number of retries (default 3)
- layer2-tcp-ports string
The TCP ports to ping for local device discovery (default "22,80,135,179,443,3389,5040,7547,62078")
- layer2-udp-trace-port uint
The UDP port number to use for UDP trace requests (default 9)
- ldap-base-dn string
The base DN used for LDAP searches
- ldap-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- ldap-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
- ldap-legacy-tls
Set this to true to authenticate over legacy TLS versions (< 1.2)
- ldap-password string
The password for the LDAP account
- ldap-service-options string
The comma-separated list of services to sync data from (defaults,computer,user,group) (default "defaults")
- ldap-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication
- ldap-url string
The URL used for the LDAP server
- ldap-username string
The username for the LDAP account
- max-attempts int
Set the maximum number of attempts for each probe (default 3)
- G, --max-group-size int
Set the maximum number of targets to process in each group (default 4096)
- R, --max-host-rate int
Set the maximum packet rate per target (including ARP broadcast) (default 40)

- `--max-sockets int`
Set the maximum number of concurrent sockets (default 2048)
- `--max-ttl int`
Set the default TTL on probe packets (default 255)
- `--mdns-port uint`
The destination port for MDNS probes (default 5353)
- `--memcache-port uint`
The destination port for memcached probes (default 11211)
- `--miradore-api-key string`
The API key for the Miradore account
- `--miradore-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--miradore-hostname string`
The Miradore web console hostname (url)
- `--modbus-identification-level string`
Identification level, one of 'basic', 'regular', or 'extended'. (default "regular")
- `--mssql-port uint`
The destination port for MSSQL probes (default 1434)
- `--nameservers string`
One or more nameservers to use for DNS resolution
- `--natpmp-port uint`
The destination port for NATPMP probes (default 5351)
- `--nessus-access-key string`
The access key for the Nessus Professional account
- `--nessus-api-url string`
The URL used for the Nessus Professional account's API access
- `--nessus-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--nessus-fingerprint-only`
Import vulnerabilities for fingerprinting purposes only
- `--nessus-insecure`
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
(default true)
- `--nessus-risks string`
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (default
"None,Low,Medium,High,Critical")
- `--nessus-secret-key string`
The secret key for the Nessus Professional account

- `--nessus-severities` string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- `--nessus-thumbprints` string
A set of IP=SHA256:B64HASH pairs to trust for authentication
- `--netbios-port` uint
The destination port for NetBIOS Name Service probes (default 137)
- `--nowait`
Exit the user interface immediately upon completion
- `--ntp-port` uint
The destination port for NTP probes (default 123)
- `--openvpn-ports` string
The destination ports for OpenVPN probes (default "1194")
- `-o, --output` string
Output directory for scan results and analysis ('disable' to skip)
- `--output-raw` string
Set the raw output file for scan data
- `--overwrite`
Overwrite and replace the output directory if it already exists
- `--passes` int
Set the number of passes for each probe (default 1)
- `--pca-port` uint
The destination port for PCAnywhere probes (default 5632)
- `--probes` string
Launch a subset of the probes, comma-delimited (default "defaults")
- `--psdisco-ports` string
The destination ports for playstation discovery probes (default "987,9302")
- `--qualys-api-url` string
The URL used for the Qualys account's API access
- `--qualys-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--qualys-fingerprint-only`
Import vulnerabilities for fingerprinting purposes only
- `--qualys-include-unscanned`
Include assets that have not been assessed for vulnerabilities
- `--qualys-password` string
The password for the Qualys account
- `--qualys-risks` string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (default "None,Low,Medium,High,Critical")

- `--qualys-severities` string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- `--qualys-username` string
The username for the Qualys account
- `-r, --rate` int
Set the maximum packet rate for the overall scan (default 1000)
- `--rdns-max-concurrent` int
The maximum number of concurrent DNS lookups (default 64)
- `--rdns-timeout` uint
The DNS PTR lookup timeout in seconds (default 3)
- `--rpcbind-port` uint
The destination port for RPCBind probes (default 111)
- `--rpcbind-port-nfs` uint
The destination port for NFS probes (default 2049)
- `--s7comm-request-extended-information`
If true, request extended device information.
- `--sample-duration` string
Specify the duration in seconds to sample network traffic (or '0' for non-stop) (default "300")
- `--sample-excludes` string
Specify host exclusions
- `--sample-interfaces` string
Specify a comma-separated list of network interfaces (or 'all' for everything)
- `--sample-targets` string
Specify the discovery scope (default "10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 169.254.0.0/16")
- `-S, --screenshots`
Capture screenshots from scan target web services (default true)
- `--sentinelone-api-url` string
The URL used for the SentinelOne account's API access
- `--sentinelone-client-id` string
The client ID for the SentinelOne account
- `--sentinelone-client-secret` string
The client secret for the SentinelOne account
- `--sentinelone-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--shodan-api-key` string
The key used for the Shodan account's API access

- `--shodan-exclude-unknown`
Exclude assets that cannot be merged into an existing asset
- `--shodan-mode string`
The search mode (assets or query). The assets option queries the scan targets (default "assets")
- `--shodan-query string`
The search string to use in query mode
- `--sip-port uint`
The destination port for SIP probes (default 5060)
- `--snmp-comms string`
The comma-separated list of SNMP v1/v2c communities (default "public,private")
- `--snmp-disable-bulk`
If true, do not use bulk walking operations
- `--snmp-max-repetitions uint`
The maximum number of repetitions in a bulk walk operation (default 16)
- `--snmp-max-retries int`
The maximum number of retries for an SNMP operation (default 1)
- `--snmp-poll-interval uint`
The minimum number of seconds between polling each host after initial discovery (default 300)
- `--snmp-port uint`
The destination port for SNMP probes (default 161)
- `--snmp-timeout uint`
The maximum number of seconds for each individual SNMP operation (default 5)
- `--snmp-v3-auth-passphrase string`
The authentication passphrase
- `--snmp-v3-auth-protocol string`
The authentication protocol (none, md5, sha, sha224, sha256, sha384, sha512) (default "none")
- `--snmp-v3-context string`
The optional SNMP v3 context to supply
- `--snmp-v3-privacy-passphrase string`
The privacy passphrase
- `--snmp-v3-privacy-protocol string`
The privacy protocol (none, des, aes, aes192, aes256, aes192c, aes256c) (default "none")
- `--snmp-v3-username string`
The username to use for SNMP v3 authentication
- `--snmp-walk-timeout uint`
The maximum number of seconds for each SNMP walk operation (default 60)

- ssdp-port uint
The destination port for UPnP/SSDP probes (default 1900)
- ssh-fingerprint
Enable fingerprinting using partial authentication (default true)
- ssh-fingerprint-username string
The username to use for partial authentication SSH fingerprinting (default "STATUS")
- steam-ports string
The destination ports for Steam discovery probes (default "27036")
- subnet-ping
Only scan subnets that have at least one active response using the subnet-ping settings
- subnet-ping-arp-fast
Enables fast mode by ARP scanning at the scan rate vs host rate (subnet ping)
- subnet-ping-aws-instances-access-key string
The access key for the AWS account (subnet ping)
- subnet-ping-aws-instances-assume-role-name string
The role to assume for all accounts in the organization for cross-account access (subnet ping)
- subnet-ping-aws-instances-delete-stale
Automatically delete stale AWS assets (subnet ping)
- subnet-ping-aws-instances-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-aws-instances-include-stopped
Include assets that are not currently running (subnet ping)
- subnet-ping-aws-instances-regions string
The comma-separated list of regions for the AWS account (subnet ping)
- subnet-ping-aws-instances-secret-access-key string
The secret access key for the AWS account (subnet ping)
- subnet-ping-aws-instances-service-options string
The comma-separated list of services to sync data from (defaults,ec2,elb,elbv2,rds,lambda) (subnet ping) (default "defaults")
- subnet-ping-aws-instances-site-per-account
Automatically create a new site per account (subnet ping)
- subnet-ping-aws-instances-site-per-vpc
Automatically create a new site per VPC (subnet ping)
- subnet-ping-aws-instances-token string
The session token for the AWS account (subnet ping)
- subnet-ping-azure-client-id string
The application ID (client ID) for the Azure account (subnet ping)
- subnet-ping-azure-client-secret string
The client secret for the Azure account (subnet ping)

- subnet-ping-azure-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-azure-multi-subscription
Access all subscriptions in the directory (tenant) for the Azure account (subnet ping)
- subnet-ping-azure-password string
The password for the Azure account (subnet ping)
- subnet-ping-azure-service-options string
The comma-separated list of services to sync data from
(defaults,vm,vmss,azsql,cosmos,lb,functionapp) (subnet ping) (default "defaults")
- subnet-ping-azure-site-per-subscription
Automatically create a new site per subscription (subnet ping)
- subnet-ping-azure-subscription-id string
The subscription ID for the Azure account (subnet ping)
- subnet-ping-azure-tenant-id string
The directory ID (tenant ID) for the Azure account (subnet ping)
- subnet-ping-azure-username string
The username for the Azure account (subnet ping)
- subnet-ping-azuread-client-id string
The application ID (client ID) for the Azure account (subnet ping)
- subnet-ping-azuread-client-secret string
The client secret for the Azure account (subnet ping)
- subnet-ping-azuread-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-azuread-include-inactive
Include assets that are marked as inactive in the AzureAD account (subnet ping)
- subnet-ping-azuread-password string
The password for the AzureAD account (subnet ping)
- subnet-ping-azuread-service-options string
The comma-separated list of services to sync data from (defaults,dev,user,group)
(subnet ping) (default "defaults")
- subnet-ping-azuread-tenant-id string
The directory ID (tenant ID) for the Azure account (subnet ping)
- subnet-ping-azuread-username string
The username for the AzureAD account (subnet ping)
- subnet-ping-bacnet-ports string
The destination ports for BACnet probes (subnet ping) (default "46808,47808,48808")
- subnet-ping-bedrock-ports string
The destination ports for Bedrock probes (subnet ping) (default "19132")

- subnet-ping-censys-api-url string
The API endpoint to use for Censys Search (subnet ping) (default ["https://search.censys.io"](https://search.censys.io))
- subnet-ping-censys-client-id string
The Client ID to use for Censys Search authentication (subnet ping)
- subnet-ping-censys-client-secret string
The Client Secret to use for Censys Search authentication (subnet ping)
- subnet-ping-censys-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-censys-mode string
The search mode (assets or query). The assets option queries the scan targets (subnet ping) (default "assets")
- subnet-ping-censys-query string
The search string to use in query mode (subnet ping)
- subnet-ping-coap-port uint
The destination port for CoAP probes (subnet ping) (default 5683)
- subnet-ping-crestron-port uint
The destination port for Crestron probes (subnet ping) (default 41794)
- subnet-ping-crowdstrike-api-url string
The URL used for the CrowdStrike account's API access (subnet ping)
- subnet-ping-crowdstrike-client-id string
The client ID for the CrowdStrike account (subnet ping)
- subnet-ping-crowdstrike-client-secret string
The client secret for the CrowdStrike account (subnet ping)
- subnet-ping-crowdstrike-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-crowdstrike-filter string
An optional Falcon Query Language (FQL) filter for imported assets (subnet ping)
- subnet-ping-crowdstrike-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)
- subnet-ping-crowdstrike-risks string
Minimum risk of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping) (default "None,Low,Medium,High,Critical")
- subnet-ping-crowdstrike-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet ping) (default "Info,Low,Medium,High,Critical")
- subnet-ping-dahua-dhip-ports string
The destination ports for Dahua DHIP discovery probes (subnet ping) (default "37810")
- subnet-ping-defender365-client-id string
The application ID (client ID) for the Azure account (subnet ping)

- subnet-ping-defender365-client-secret string
The client secret for the Azure account (subnet ping)
- subnet-ping-defender365-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-defender365-include-inactive
Include assets that have stopped reporting to the Microsoft 365 Defender service (subnet ping)
- subnet-ping-defender365-tenant-id string
The directory ID (tenant ID) for the Azure account (subnet ping)
- subnet-ping-dnp3-address-probe-timeout int
Time limit (in seconds) for DNP address discovery. (subnet ping) (default 30)
- subnet-ping-dnp3-banner-address-discovery string
One of 'require', 'prefer', or 'ignore'. (subnet ping) (default "ignore")
- subnet-ping-dnp3-destination-address-discovery-range string
A numeric range of addresses to attempt to discover. (subnet ping) (default "0-32")
- subnet-ping-dnp3-explorer-address int
Source DNP3 address for the explorer. (subnet ping) (default -1)
- subnet-ping-dns-disable-google-myaddr
Disables resolution of upstream DNS via Google myaddr service (subnet ping)
- subnet-ping-dns-disable-meraki-detection
Disables detection of Meraki DNS interception (subnet ping)
- subnet-ping-dns-port uint
The destination port for DNS probes (subnet ping) (default 53)
- subnet-ping-dns-resolve-name string
The target hostname for DNS queries ('off' to disable) (subnet ping) (default "www.google.com")
- subnet-ping-dns-trace-domain string
The subdomain to use for trace requests ('off' to disable) (subnet ping) (default "helper.rumble.network")
- subnet-ping-dtls-ports string
The destination ports for DTLS probes (subnet ping) (default "443,3391,4433,5246,5349,5684")
- subnet-ping-echo-report-errors
Report errors from intermediate in-scope hosts (subnet ping)
- subnet-ping-ethernetip-udp-ports string
The destination ports for EtherNet/IP UDP probes (subnet ping) (default "44818")
- subnet-ping-ethernetip-use-tagged-context string
Set this to true to use a runZero-specific sender context for debugging. (subnet ping) (default "false")

- subnet-ping-fins-port uint
The destination port for FINS probes (subnet ping) (default 9600)
- subnet-ping-gcp-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-gcp-service-options string
The comma-separated list of services to sync data from (defaults,vm,lb,cloudsql) (subnet ping) (default "defaults")
- subnet-ping-gcp-site-per-project
Automatically create a new site per project (subnet ping)
- subnet-ping-gendp-payload-base64 string
The generic udp payload as base64 (subnet ping)
- subnet-ping-gendp-payload-hex string
The generic udp payload as hexadecimal (subnet ping)
- subnet-ping-gendp-payload-text string
The generic udp payload as plain text (subnet ping)
- subnet-ping-gendp-ports string
The destination ports for the generic udp probe (subnet ping)
- subnet-ping-googleworkspace-client-email string
The email address of the service account (subnet ping)
- subnet-ping-googleworkspace-client-id string
The ID of the service account (subnet ping)
- subnet-ping-googleworkspace-customer-id string
An optional customer ID for multi-tenant environments (subnet ping) (default "my_customer")
- subnet-ping-googleworkspace-delegate string
The email address of an admin account with directory access (subnet ping)
- subnet-ping-googleworkspace-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-googleworkspace-private-key string
The PEM encoded private key (subnet ping)
- subnet-ping-googleworkspace-private-key-id string
The ID of the private key (subnet ping)
- subnet-ping-googleworkspace-project-id string
The project ID of the service account (subnet ping)
- subnet-ping-googleworkspace-service-options string
The comma-separated list of services to sync data from (defaults,chromeos,mobile,endpoint,user,group) (subnet ping) (default "defaults")
- subnet-ping-hiddiscoveryd-port uint
The destination port for HID discoveryd probes (subnet ping) (default 4070)

- subnet-ping-igel-discovery-ports string
The destination ports for IGEL discovery probes (subnet ping) (default "30005")
- subnet-ping-ike-port uint
The destination port for IKE probes (subnet ping) (default 500)
- subnet-ping-insightvm-api-url string
The URL used for the InsightVM account's API access (subnet ping)
- subnet-ping-insightvm-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-insightvm-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)
- subnet-ping-insightvm-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
(subnet ping) (default true)
- subnet-ping-insightvm-password string
The password for the InsightVM account (subnet ping)
- subnet-ping-insightvm-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping)
(default "None,Low,Medium,High,Critical")
- subnet-ping-insightvm-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet
ping) (default "Info,Low,Medium,High,Critical")
- subnet-ping-insightvm-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (subnet ping)
- subnet-ping-insightvm-username string
The username for the InsightVM account (subnet ping)
- subnet-ping-intune-client-id string
The application ID (client ID) for the Azure account (subnet ping)
- subnet-ping-intune-client-secret string
The client secret for the Azure account (subnet ping)
- subnet-ping-intune-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-intune-password string
The password for the Intune account (subnet ping)
- subnet-ping-intune-tenant-id string
The directory ID (tenant ID) for the Azure account (subnet ping)
- subnet-ping-intune-username string
The username for the Intune account (subnet ping)
- subnet-ping-ipmi-port uint
The destination port for IPMI probes (subnet ping) (default 623)

- subnet-ping-kerberos-port uint
The destination port for kerberos probes (subnet ping) (default 88)
- subnet-ping-knxnet-ports string
The destination ports for knxnet probes (subnet ping) (default "3671")
- subnet-ping-l2t-port uint
The destination port for L2T probes (subnet ping) (default 2228)
- subnet-ping-l2tp-ports string
The destination ports for L2TP probes (subnet ping) (default "1701")
- subnet-ping-lantronix-port uint
The destination port for Lantronix probes (subnet ping) (default 30718)
- subnet-ping-layer2-add-targets
Set this false to skip scanning discovered targets (subnet ping) (default true)
- subnet-ping-layer2-force
Set this to true to force discovery even without local targets (subnet ping)
- subnet-ping-layer2-max-retries uint
The desired number of retries (subnet ping) (default 3)
- subnet-ping-layer2-tcp-ports string
The TCP ports to ping for local device discovery (subnet ping) (default "22,80,135,179,443,3389,5040,7547,62078")
- subnet-ping-layer2-udp-trace-port uint
The UDP port number to use for UDP trace requests (subnet ping) (default 9)
- subnet-ping-ldap-base-dn string
The base DN used for LDAP searches (subnet ping)
- subnet-ping-ldap-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-ldap-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (subnet ping)
- subnet-ping-ldap-legacy-tls
Set this to true to authenticate over legacy TLS versions (< 1.2) (subnet ping)
- subnet-ping-ldap-password string
The password for the LDAP account (subnet ping)
- subnet-ping-ldap-service-options string
The comma-separated list of services to sync data from (defaults,computer,user,group) (subnet ping) (default "defaults")
- subnet-ping-ldap-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (subnet ping)
- subnet-ping-ldap-url string
The URL used for the LDAP server (subnet ping)

- subnet-ping-ldap-username string
The username for the LDAP account (subnet ping)
- subnet-ping-max-attempts int
Set the maximum number of attempts for each probe (default 1)
- subnet-ping-max-ttl int
Set the default TTL on subnet-ping probe packets (default 255)
- subnet-ping-mdns-port uint
The destination port for MDNS probes (subnet ping) (default 5353)
- subnet-ping-memcache-port uint
The destination port for memcached probes (subnet ping) (default 11211)
- subnet-ping-miradore-api-key string
The API key for the Miradore account (subnet ping)
- subnet-ping-miradore-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-miradore-hostname string
The Miradore web console hostname (url) (subnet ping)
- subnet-ping-modbus-identification-level string
Identification level, one of 'basic', 'regular', or 'extended'. (subnet ping) (default "regular")
- subnet-ping-mode string
Set the subnet-ping discovery profile: auto (default "auto")
- subnet-ping-mssql-port uint
The destination port for MSSQL probes (subnet ping) (default 1434)
- subnet-ping-natpmp-port uint
The destination port for NATPMP probes (subnet ping) (default 5351)
- subnet-ping-nessus-access-key string
The access key for the Nessus Professional account (subnet ping)
- subnet-ping-nessus-api-url string
The URL used for the Nessus Professional account's API access (subnet ping)
- subnet-ping-nessus-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-nessus-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)
- subnet-ping-nessus-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
(subnet ping) (default true)
- subnet-ping-nessus-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping)
(default "None,Low,Medium,High,Critical")
- subnet-ping-nessus-secret-key string
The secret key for the Nessus Professional account (subnet ping)

- subnet-ping-nessus-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet ping) (default "Info,Low,Medium,High,Critical")
- subnet-ping-nessus-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (subnet ping)
- subnet-ping-net-size int
Set the subnet size to use for the subnet ping (default 256)
- subnet-ping-netbios-port uint
The destination port for NetBIOS Name Service probes (subnet ping) (default 137)
- subnet-ping-ntp-port uint
The destination port for NTP probes (subnet ping) (default 123)
- subnet-ping-openvpn-ports string
The destination ports for OpenVPN probes (subnet ping) (default "1194")
- subnet-ping-passes int
Set the number of passes for the subnet-ping phase (default 1)
- subnet-ping-pca-port uint
The destination port for PCAnywhere probes (subnet ping) (default 5632)
- subnet-ping-probes string
Launch a subset of the probes for the subnet-ping, comma-delimited (default "arp,echo,syn,connect,netbios,snmp,ntp,sunrpc,ike,openvpn,mdns")
- subnet-ping-psdisco-ports string
The destination ports for playstation discovery probes (subnet ping) (default "987,9302")
- subnet-ping-qualys-api-url string
The URL used for the Qualys account's API access (subnet ping)
- subnet-ping-qualys-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-qualys-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)
- subnet-ping-qualys-include-unscanned
Include assets that have not been assessed for vulnerabilities (subnet ping)
- subnet-ping-qualys-password string
The password for the Qualys account (subnet ping)
- subnet-ping-qualys-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping) (default "None,Low,Medium,High,Critical")
- subnet-ping-qualys-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet ping) (default "Info,Low,Medium,High,Critical")

- subnet-ping-qualys-username string
The username for the Qualys account (subnet ping)
- subnet-ping-rdns-max-concurrent int
The maximum number of concurrent DNS lookups (subnet ping) (default 64)
- subnet-ping-rdns-timeout uint
The DNS PTR lookup timeout in seconds (subnet ping) (default 3)
- subnet-ping-rpcbind-port uint
The destination port for RPCBind probes (subnet ping) (default 111)
- subnet-ping-rpcbind-port-nfs uint
The destination port for NFS probes (subnet ping) (default 2049)
- subnet-ping-s7comm-request-extended-information
If true, request extended device information. (subnet ping)
- subnet-ping-sample-duration string
Specify the duration in seconds to sample network traffic (or '0' for non-stop) (subnet ping) (default "300")
- subnet-ping-sample-excludes string
Specify host exclusions (subnet ping)
- subnet-ping-sample-interfaces string
Specify a comma-separated list of network interfaces (or 'all' for everything) (subnet ping)
- subnet-ping-sample-rate int
Set the sample rate of addresses within each subnet as a percentage (default 4)
- subnet-ping-sample-targets string
Specify the discovery scope (subnet ping) (default "10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 169.254.0.0/16")
- subnet-ping-sentinelone-api-url string
The URL used for the SentinelOne account's API access (subnet ping)
- subnet-ping-sentinelone-client-id string
The client ID for the SentinelOne account (subnet ping)
- subnet-ping-sentinelone-client-secret string
The client secret for the SentinelOne account (subnet ping)
- subnet-ping-sentinelone-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-shodan-api-key string
The key used for the Shodan account's API access (subnet ping)
- subnet-ping-shodan-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-shodan-mode string
The search mode (assets or query). The assets option queries the scan targets (subnet ping) (default "assets")

- subnet-ping-shodan-query string
The search string to use in query mode (subnet ping)
- subnet-ping-sip-port uint
The destination port for SIP probes (subnet ping) (default 5060)
- subnet-ping-snmp-comms string
The comma-separated list of SNMP v1/v2c communities (subnet ping) (default "public,private")
- subnet-ping-snmp-disable-bulk
If true, do not use bulk walking operations (subnet ping)
- subnet-ping-snmp-max-repetitions uint
The maximum number of repetitions in a bulk walk operation (subnet ping) (default 16)
- subnet-ping-snmp-max-retries int
The maximum number of retries for an SNMP operation (subnet ping) (default 1)
- subnet-ping-snmp-poll-interval uint
The minimum number of seconds between polling each host after initial discovery (subnet ping) (default 300)
- subnet-ping-snmp-port uint
The destination port for SNMP probes (subnet ping) (default 161)
- subnet-ping-snmp-timeout uint
The maximum number of seconds for each individual SNMP operation (subnet ping) (default 5)
- subnet-ping-snmp-v3-auth-passphrase string
The authentication passphrase (subnet ping)
- subnet-ping-snmp-v3-auth-protocol string
The authentication protocol (none, md5, sha, sha224, sha256, sha384, sha512) (subnet ping) (default "none")
- subnet-ping-snmp-v3-context string
The optional SNMP v3 context to supply (subnet ping)
- subnet-ping-snmp-v3-privacy-passphrase string
The privacy passphrase (subnet ping)
- subnet-ping-snmp-v3-privacy-protocol string
The privacy protocol (none, des, aes, aes192, aes256, aes192c, aes256c) (subnet ping) (default "none")
- subnet-ping-snmp-v3-username string
The username to use for SNMP v3 authentication (subnet ping)
- subnet-ping-snmp-walk-timeout uint
The maximum number of seconds for each SNMP walk operation (subnet ping) (default 60)
- subnet-ping-ssdp-port uint
The destination port for UPnP/SSDP probes (subnet ping) (default 1900)

- subnet-ping-ssh-fingerprint
Enable fingerprinting using partial authentication (subnet ping) (default true)
- subnet-ping-ssh-fingerprint-username string
The username to use for partial authentication SSH fingerprinting (subnet ping) (default "STATUS")
- subnet-ping-steam-ports string
The destination ports for Steam discovery probes (subnet ping) (default "27036")
- subnet-ping-syn-disable-bogus-filter
Disable bogus service detection and filtering (subnet ping)
- subnet-ping-syn-forwarding-check
Perform an IP forwarding check as part of the scan (subnet ping) (default true)
- subnet-ping-syn-forwarding-check-target string
An external IPv4 address for the forwarding check (default:runzero) (subnet ping) (default "13.248.161.247")
- subnet-ping-syn-max-retries uint
The maximum number of retries trace and SYN requests (subnet ping) (default 2)
- subnet-ping-syn-report-resets
Set this to true to report RST responses (subnet ping) (default true)
- subnet-ping-syn-reset-sessions
Reset middle-box/firewall sessions automatically (subnet ping) (default true)
- subnet-ping-syn-reset-sessions-delay uint
Minimum delay in milliseconds between a SYN and a session reset (subnet ping)
- subnet-ping-syn-reset-sessions-limit uint
Maximum number of in-flight sessions before forcing session resets (subnet ping) (default 50)
- subnet-ping-syn-traceroute
Perform a multi-protocol traceroute as part of the scan (subnet ping) (default true)
- subnet-ping-syn-udp-trace-port uint
The UDP port number to use for UDP trace requests (subnet ping) (default 9)
- subnet-ping-tcp-ports string
The list of TCP ports to subnet-ping using the syn and connect probes (default "22,80,135,179,443,3389,5040,7547,62078")
- subnet-ping-tenable-access-key string
The access key for the Tenable.io account (subnet ping)
- subnet-ping-tenable-api-url string
The URL used for the Tenable.io account's API access (subnet ping)
- subnet-ping-tenable-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-tenable-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)

- subnet-ping-tenable-include-unscanned
Include assets that have not been assessed for vulnerabilities (subnet ping)
- subnet-ping-tenable-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping)
(default "None,Low,Medium,High,Critical")
- subnet-ping-tenable-secret-key string
The secret key for the Tenable.io account (subnet ping)
- subnet-ping-tenable-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet ping)
(default "Info,Low,Medium,High,Critical")
- subnet-ping-tenablesecuritycenter-access-key string
The access key for the Tenable Security Center account (subnet ping)
- subnet-ping-tenablesecuritycenter-api-url string
The URL used for the Tenable Security Center account's API access (subnet ping)
- subnet-ping-tenablesecuritycenter-batch-size string
The number of records to request at a time. (between 2000 and 10000) (subnet ping)
(default "2000")
- subnet-ping-tenablesecuritycenter-exclude-unknown
Exclude assets that cannot be merged into an existing asset (subnet ping)
- subnet-ping-tenablesecuritycenter-fingerprint-only
Import vulnerabilities for fingerprinting purposes only (subnet ping)
- subnet-ping-tenablesecuritycenter-insecure string
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
(subnet ping)
- subnet-ping-tenablesecuritycenter-query-id string
The ID of an existing vulnerability query in the Tenable Security Center account (subnet ping)
- subnet-ping-tenablesecuritycenter-query-mode string
Set to 'filters' to provide 'severities' and 'risks' values to import. Set to 'query-id' to provide a value for 'query-id'. (subnet ping) (default "filters")
- subnet-ping-tenablesecuritycenter-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (subnet ping)
(default "None,Low,Medium,High,Critical")
- subnet-ping-tenablesecuritycenter-secret-key string
The secret key for the Tenable Security Center account (subnet ping)
- subnet-ping-tenablesecuritycenter-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (subnet ping)
(default "Info,Low,Medium,High,Critical")
- subnet-ping-tenablesecuritycenter-sync-since string
Specify an initial date to sync data from. (subnet ping)

- subnet-ping-tenablesecuritycenter-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (subnet ping)
- subnet-ping-tftp-ports string
The destination ports for TFTP probes (subnet ping) (default "69")
- subnet-ping-tos int
Set the default ToS on subnet-ping probe packets
- subnet-ping-ubnt-port uint
The destination port for Ubiquiti probes (subnet ping) (default 10001)
- subnet-ping-verbose
Display verbose output for the subnet-ping
- subnet-ping-very-verbose
Display very verbose output for the subnet-ping
- subnet-ping-vmware-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN) (subnet ping) (default true)
- subnet-ping-vmware-password string
The password to use for VMware SDK authentication (read-only) (subnet ping)
- subnet-ping-vmware-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication (subnet ping)
- subnet-ping-vmware-username string
The username to use for VMware SDK authentication (read-only) (subnet ping)
- subnet-ping-webmin-ports string
The destination ports for webmin probes (subnet ping) (default "10000")
- subnet-ping-wlan-list-poll-interval uint
The minimum number of seconds between polls of the access point list (subnet ping) (default 300)
- subnet-ping-wsd-port uint
The destination port for WSD probes (subnet ping) (default 3702)
- syn-disable-bogus-filter
Disable bogus service detection and filtering
- syn-forwarding-check
Perform an IP forwarding check as part of the scan (default true)
- syn-forwarding-check-target string
An external IPv4 address for the forwarding check (default:runzero) (default "13.248.161.247")
- syn-max-retries uint
The maximum number of retries trace and SYN requests (default 2)
- syn-report-resets
Set this to true to report RST responses (default true)

- syn-reset-sessions
Reset middle-box/firewall sessions automatically (default true)
- syn-reset-sessions-delay uint
Minimum delay in milliseconds between a SYN and a session reset
- syn-reset-sessions-limit uint
Maximum number of in-flight sessions before forcing session resets (default 50)
- syn-traceroute
Perform a multi-protocol traceroute as part of the scan (default true)
- syn-udp-trace-port uint
The UDP port number to use for UDP trace requests (default 9)
- tcp-excludes string
The list of TCP ports to always exclude
- p, --tcp-ports string
The list of TCP ports scan using the syn and connect probes (see below for default)
- tcp-skip-protocol
Set this to skip protocol detection on TCP ports
- tenable-access-key string
The access key for the Tenable.io account
- tenable-api-url string
The URL used for the Tenable.io account's API access
- tenable-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- tenable-fingerprint-only
Import vulnerabilities for fingerprinting purposes only
- tenable-include-unscanned
Include assets that have not been assessed for vulnerabilities
- tenable-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (default "None,Low,Medium,High,Critical")
- tenable-secret-key string
The secret key for the Tenable.io account
- tenable-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- tenablesecuritycenter-access-key string
The access key for the Tenable Security Center account
- tenablesecuritycenter-api-url string
The URL used for the Tenable Security Center account's API access
- tenablesecuritycenter-batch-size string
The number of records to request at a time. (between 2000 and 10000) (default "2000")

- tenablesecuritycenter-exclude-unknown
Exclude assets that cannot be merged into an existing asset
- tenablesecuritycenter-fingerprint-only
Import vulnerabilities for fingerprinting purposes only
- tenablesecuritycenter-insecure string
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
- tenablesecuritycenter-query-id string
The ID of an existing vulnerability query in the Tenable Security Center account
- tenablesecuritycenter-query-mode string
Set to 'filters' to provide 'severities' and 'risks' values to import. Set to 'query-id' to provide a value for 'query-id'. (default "filters")
- tenablesecuritycenter-risks string
Risk levels of imported vulnerabilities (None, Low, Medium, High, Critical) (default "None,Low,Medium,High,Critical")
- tenablesecuritycenter-secret-key string
The secret key for the Tenable Security Center account
- tenablesecuritycenter-severities string
Severity levels of imported vulnerabilities (Info, Low, Medium, High, Critical) (default "Info,Low,Medium,High,Critical")
- tenablesecuritycenter-sync-since string
Specify an initial date to sync data from.
- tenablesecuritycenter-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication
- text
Force text-only mode (no console ui)
- tftp-ports string
The destination ports for TFTP probes (default "69")
- tos int
Set the default ToS on probe packets
- ubnt-port uint
The destination port for Ubiquiti probes (default 10001)
- upload
Automatically upload results to the runZero Console
- u, --upload-site string
Specify the Site ID or Name to upload the raw scan results to if --upload is specified (default "Primary")
- v, --verbose
Display verbose output
- very-verbose
Display very verbose output

- vmware-insecure
Set this to true to authenticate to untrusted endpoints (self-signed or no IP SAN)
(default true)
- vmware-password string
The password to use for VMware SDK authentication (read-only)
- vmware-thumbprints string
A set of IP=SHA256:B64HASH pairs to trust for authentication
- vmware-username string
The username to use for VMware SDK authentication (read-only)
- webmin-ports string
The destination ports for webmin probes (default "10000")
- wlan-list-poll-interval uint
The minimum number of seconds between polls of the access point list (default 300)
- wsd-port uint
The destination port for WSD probes (default 3702)

Standard ports scanned

1	7	9	13	17	19	21	22	23	25	37
42	43	49	53	69	70	79	80	81	82	83
84	85	88	102	105	109	110	111	113	119	123
135	137	139	143	161	179	222	264	280	384	389
402	407	442	443	444	445	465	500	502	512	513
515	523	524	540	548	554	587	617	623	631	636
664	689	705	717	743	771	783	830	873	888	902
903	910	912	921	990	993	995	998	1000	1024	1030
1035	1080	1083	1089	1090	1091	1098	1099	1100	1101	1102
1103	1128	1129	1158	1199	1211	1220	1234	1241	1260	1270
1300	1311	1352	1433	1440	1443	1468	1494	1514	1521	1530
1533	1581	1582	1583	1604	1610	1611	1723	1755	1801	1811
1830	1883	1900	2000	2002	2021	2023	2049	2068	2074	2082
2083	2100	2103	2105	2121	2181	2199	2207	2222	2224	2323
2362	2375	2376	2379	2380	2381	2443	2525	2533	2598	2601
2604	2638	2809	2947	2967	3000	3001	3003	3033	3037	3050
3057	3071	3083	3128	3200	3217	3220	3260	3268	3269	3273
3299	3300	3306	3311	3312	3351	3389	3460	3500	3502	3628
3632	3690	3780	3790	3817	3871	3872	3900	4000	4092	4322
4343	4353	4365	4366	4368	4369	4433	4443	4444	4445	4567
4659	4679	4730	4786	4840	4848	4949	4950	4987	5000	5001
5007	5022	5037	5038	5040	5051	5060	5061	5093	5168	5222
5247	5250	5275	5347	5351	5353	5355	5392	5400	5405	5432
5433	5498	5520	5521	5554	5555	5560	5580	5601	5631	5632
5666	5671	5672	5683	5800	5814	5900	5901	5902	5903	5904
5905	5906	5907	5908	5909	5910	5911	5920	5938	5984	5985
5986	5988	5989	6000	6001	6002	6050	6060	6070	6080	6082
6101	6106	6112	6161	6262	6379	6405	6443	6502	6503	6504

6514	6542	6556	6660	6661	6667	6905	6988	7000	7001	7002	7021
7070	7071	7077	7080	7100	7144	7181	7210	7373	7443	7474	7510
7547	7579	7580	7676	7700	7770	7777	7778	7787	7800	7801	7879
7902	8000	8001	8003	8006	8008	8009	8010	8012	8014	8020	8023
8028	8030	8080	8081	8082	8083	8086	8087	8088	8089	8090	8095
8098	8099	8100	8123	8127	8161	8172	8180	8181	8182	8205	8222
8300	8303	8333	8400	8443	8444	8445	8471	8488	8500	8503	8530
8531	8545	8649	8686	8787	8800	8812	8834	8850	8871	8880	8883
8888	8889	8890	8899	8901	8902	8903	8983	9000	9001	9002	9042
9060	9080	9081	9084	9090	9091	9092	9099	9100	9111	9152	9160
9200	9300	9380	9390	9391	9401	9418	9440	9443	9471	9495	9524
9527	9530	9593	9594	9595	9600	9809	9855	9999	10000	10001	10008
10050	10051	10080	10098	10162	10202	10203	10250	10255	10257	10259	10443
10616	10628	11000	11099	11211	11234	11333	12174	12203	12221	12345	12379
12397	12401	13364	13500	13778	13838	14330	15200	15671	15672	16102	16443
16992	16993	17185	17200	17472	17775	17776	17777	17778	17781	17782	17783
17784	17790	17791	17798	18264	18881	19300	19810	19888	20000	20010	20031
20034	20101	20111	20171	20222	20293	22222	23472	23791	23943	25000	25025
25565	25672	26000	26122	27000	27017	27018	27019	27080	27888	28017	28222
28784	30000	31001	31099	32764	32844	32913	33060	34205	34443	34962	34963
34964	37718	37777	37890	37891	37892	38008	38010	38080	38102	38292	40007
40317	41025	41080	41523	41524	44334	44343	44818	45230	46823	46824	47001
47002	47290	48899	49152	50000	50013	50021	50051	50070	50090	50121	51443
52302	52311	54321	54921	54922	54923	55553	55580	57772	61614	61616	62078
62514	65002	65535									

Scan outputs

The runZero Scanner generates a directory of output files by default. This directory includes the following items.

- `scan.rumble.gz`: The raw scan data compressed via gzip, this can be imported or reprocessed via `--import`
- `assets.jsonl`: The new optimized format for correlated, fingerprinted assets.
- `nmap.xml`: A Nmap XML compatible data file that can be imported into various security tools.
- `urls.txt`: A list of discovered web services in URL format.
- `protocols.csv`: A list of protocols with their ports and URLs.
- `assets.html`: A rudimentary HTML report with screenshots.
- `screenshots`: A directory of raw screenshot images, headers in JSON format, and HTML bodies.
- Various lists including `addresses.txt`, `addresses_all.txt`, `hostnames.txt`, and `domains.txt`

Raw Scan Data

The runZero Scanner raw data is stored in a file named `scan.rumble.gz` within the output directory. This file contains JSONL-formatted records. An example ARP response record is shown below.

```
{
  "type": "result",
  "host": "192.168.0.1",
  "port": "0",
  "proto": "arp",
  "probe": "arp",
  "name": "192.168.0.1",
  "info": {
    "mac": "f0:9f:c2:11:1a:13",
    "macDateAdded": "2014-12-17",
    "macVendor": "Ubiquiti Networks Inc."
  },
  "ts": 1551584126253853200
}
```

The `info` field is a JSON map of strings to strings. Multiple values are encoded using the tab character (0x09), which are otherwise escaped as `\t` (along with `\r` and `\n` for carriage return and line feed bytes and `\x00` for null bytes). runZero scans may return more than one record of the same type for the same host if multiple responses were received.

In addition to the `result` type, there are also records for status messages, `stats`, and an initial `config` type that contains the scan parameters.

Glossary

As you read through the documentation, you will see commonly used terms. These are the definitions of each to familiarize yourself with ahead of time.

Terms

Alerts

Alerts are triggered when a certain event occurs based on **rules** defined in the Rules Engine.

Analysis reports

Analysis reports are reports which run as tasks, rather than being generated on-the-fly. These reports are static, so any changes to your **inventory** may result in assets no longer being accessible from the report.

Asset

An **asset** is a unique network entity from the perspective of the system running the **Explorer**.

Automatic queries

Automatic queries are certain queries you always want to run after a scan. After the query runs, you will be able to view its results in the **queries table**.

Dashboard

The **dashboard** provides trend data and insights that will help you assess how your inventory is changing over time.

Discovery scan

A **discovery scan** finds, identifies, and builds an **inventory** of all the connected devices and assets on your internal network.

Explorer

The **Explorer** is a lightweight scan engine that enables network and asset discovery.

Insights

Insights are queries that run automatically after each scan. They will populate on your **dashboard**.

Inventory

The **inventory** displays all assets within the Organization and can be sorted, filtered,

Organization

An **organization** represents a distinct entity; this can be your business, a specific department within your business, or one of your customers.

Outliers

Outliers show how often different values occur in specific attributes of assets and services.

Queries

Queries are filters that can be applied to your Inventory to find assets of interest.

Rules

A [rule](#) defines the action that is taken based on a set of conditions. You can create rules to proactively alert your team when there are changes to things like Explorers, assets, scans, organizations, and sites.

Scheduled scans

[Scheduled scans](#) allow you to set a date and frequency for your scan task.

Self-hosted

The [self-hosted](#) version runZero allows you to run the entire platform on-premises or within your own cloud environment.

Site

A [site](#) represents a distinct network segment, usually defined by addressing or accessibility.

Frequently Asked Questions

Here you can browse the solutions to some common runZero issues and the answers to some frequently asked questions (FAQs).

For more solutions and FAQs, check out the knowledgebase on [the runZero support portal](#).

Issues and FAQs

- [Why are there so many identical assets in my inventory?](#)
- [How do I run runZero without crashing my router?](#)
- [How do I scan VMware virtual machines without crashing the host?](#)
- [Why didn't the runZero Explorer capture screenshots?](#)
- [What protocols does runZero scan for?](#)
- [What ports does runZero scan?](#)
- [Can I safely scan my IoT or OT environments?](#)
- [Which browsers are supported when accessing the runZero Console?](#)

Still can't find your answer? [Let us know](#).

Identical assets in inventory

Why are there so many identical assets in my inventory?

Some enterprise routers and firewalls, like Cisco ASA devices, are designed to reply to all unexpected attempts on a particular port with a TCP reset (RST). On top of that, some routers listen to SIP traffic on all addresses and automatically respond to it.

runZero will generally detect when a router or firewall is replying to every connection attempt and avoid creating assets based on those responses. However, if you have a network appliance that runZero doesn't detect is spoofing response, there may be a substantial number of identical assets that will appear in your inventory.

Here are a few workarounds if you can't prevent your device from replying to all connections:

- Exclude the ports the device responds to from the scan configuration.
- Exclude all or part of the router's IP address range from the scan.
- Create a post-scan rule to delete any assets within the subnet that have the affected ports open.

These systems will respond to all requests on 1720/tcp and often 5060/tcp as well. runZero tries to automatically detect and avoid most of the SIP helper implementations, but can't always do so without possibly losing real results. If you need help deleting unwanted records, please [contact our support team](#).

Scanning routers

How do I run runZero without crashing my router?

If your router is crashing while being scanned, the likely issue is that your router is stateful and it is keeping track of every connection going through it. Since our scanning process involves thousands of attempted connections, your router likely ran out of available stateful sessions. This usually occurs when a router is using Network Address Translation (NAT) or is acting as a stateful security firewall.

If this happens, here's what you can do:

- Avoid scanning across routed networks (wired and WiFi, multiple VLANs, etc) by deploying additional Explorers.
- Reduce the Max group size in your scan configuration. This limits the number of targets runZero can scan at once, which correlates to the number of connections the router sees. Default is 4096.
- Reduce the scan speed. This will give failed connections more time to expire before new ones are attempted.
- If a router can run in bridge mode, and you don't need its NAT features, bridge mode will likely be more reliable. For example, if you have an ISP-provided router connected to a WiFi mesh system, you will likely want to run the mesh system in bridge mode and let the ISP router handle all routing including NAT; you should then be able to scan across your WiFi network without crashing the base stations.

Scanning VMWare virtual machines

How do I scan VMware virtual machines without crashing the host?

runZero can be used to scan VMware virtual machines. However, there are some precautions you should take.

VMnet interfaces normally use Network Address Translation (NAT) to route traffic between the host system and the virtual machines. The VMware software effectively operates as a stateful router. As explained above, this can cause problems when runZero tries to open thousands of connections.

For scanning VMware systems, the best option is to deploy a runZero Explorer inside VMware, on a virtual machine connected to the VMnet you want to scan. That Explorer should be able to scan all VMs on the same VMnet without VMware needing to track all of the connections.

Explorer not capturing screenshots

Why didn't the runZero Explorer capture screenshots?

The runZero Explorer needs a working install of Google Chrome to obtain screenshots. To check for Google Chrome, the Explorer looks in the following locations on each OS.

Windows

The runZero Explorer looks for Chrome on Windows in:

```
c:\Program Files (x86)\Google\Chrome\Application\chrome.exe
```

The Explorer also checks the following environment variables:

- ProgramFiles(x86)
- ProgramFiles
- ProgramW6432

Each may list another directory, in which case the Explorer looks in \Google\Chrome\Application\chrome.exe under each of those directories as well.

To find what the environment variables are set to, open a Windows command prompt and entering the command set.

For a default Windows 10 install, the default value of ProgramFiles and ProgramW6432 is c:\Program Files, which means the Explorer also checks the following location:

```
c:\Program Files\Google\Chrome\Application\chrome.exe
```

This is the default location for Chrome on a 64 bit Windows 10 system.

MacOS

On macOS, the Explorer checks for Google Chrome in the following locations:

- /Applications/Google Chrome.app/Contents/MacOS/Google Chrome
- /Applications/Google Chrome Canary.app/Contents/MacOS/Google Chrome Canary
- /Applications/Chromium.app/Contents/MacOS/Chromium

Linux

On Linux systems, the Explorer checks for Google Chrome at the following locations:

- /usr/bin/google-chrome
- /usr/bin/google-chrome-beta
- /usr/bin/google-chrome-unstable

- /usr/local/bin/chrome
- /usr/bin/chrome
- /opt/google/chrome/google-chrome
- /usr/bin/chromium
- /usr/bin/chromium-browser

Chrome is installed, but screenshots still don't work

If Google Chrome is installed in one of the standard locations, but isn't being found, it could be a permissions issue. It is also possible for Chrome to fail to run for other reasons, such as a corrupt Chrome profile. The next thing you can do is download the [runZero Scanner](#) and run a scan to disk, which will write a [log file](#) that will have more detail about the scan operation.

Protocols scanned by runZero

What protocols does runZero scan for?

runZero supports the following list of protocols:

acpp	activemq	adb	airplay
ajp	amqp	arp	backupexec
bacnet	bedrock	bgp	bitdefender-app
brother-scanner	cassandra	cdp	ceph
chargen	checkmk	chromecast	ciscosmi
citrix	click	coap	consul
couchdb	crestron	dahua-dhip	daytime
dcerpc	dhcp	dnp3	dns
docker	dotnet-remoting	drbd	drobo-nasd
dtls	echo	elasticsearch	epm
epmd	erldp	etcd2	ethernetip
finger	fins	ftp	ganglia
giop	gpsd	hiddiscoveryd	http
http2	ident	igel	igel-discovery
ike	imap	infinispan	influxdb
intermapper	ipmi	irc	jabber
java-object	java-rmi	jdbc-hsqldb	jdwp
jetdirect	jms	kasa	kerberos
knxnet	l2t	l2tp	landesk
lantronix	ldap	lexmark	lockdownd
lpd	mdns	memcache	mikrotik-bandwidth
minecraft	modbus	mongodb	mountd
mountd	mqtt	mssql	munin
mysql	mysqlx	natpmp	ndmp
neo4j	netbios	netbios-dgm	netbios-ns
nfs	ntp	opcua	openvpn
oracledb	panasonicctv	pca	pop3
postgresql	pptp	printerid	psdisco
qotd	raritan-csc	rdp	redis

rexec	riak	riak-http	rlogin
roomalert	rpcbind	rsyncd	rtsp
s7comm	securemote	sip	smb1
smb2	smb3	smtp	snmp
socks	sonicwall-sgms	spice	spotify-connect
ssdp	ssh	steam	subversion
sunrpc	tcpmux	teamviewer	telnet
tftp	thinprint	time	tls
ubnt	upnp	uscan	uscans
vault	vmauthd	vnc	vsdp
wbsm	webmin	wsd	wsman
zabbix-agent	zookeeper		

Ports scanned by runZero

What TCP ports does runZero scan?

runZero scans the following TCP ports by default:

1	7	9	13	17	19	21	22	23	25	37
42	43	49	53	69	70	79	80	81	82	83
84	85	88	102	105	109	110	111	113	119	123
135	137	139	143	161	179	222	264	280	384	389
402	407	442	443	444	445	465	500	502	512	513
515	523	524	540	548	554	587	617	623	631	636
664	689	705	717	743	771	783	830	873	888	902
903	910	912	921	990	993	995	998	1000	1024	1030
1035	1080	1083	1089	1090	1091	1098	1099	1100	1101	1102
1103	1128	1129	1158	1199	1211	1220	1234	1241	1260	1270
1300	1311	1352	1433	1440	1443	1468	1494	1514	1521	1530
1533	1581	1582	1583	1604	1610	1611	1723	1755	1801	1811
1830	1883	1900	2000	2002	2021	2023	2049	2068	2074	2082
2083	2100	2103	2105	2121	2181	2199	2207	2222	2224	2323
2362	2375	2376	2379	2380	2381	2443	2525	2533	2598	2601
2604	2638	2809	2947	2967	3000	3001	3003	3033	3037	3050
3057	3071	3083	3128	3200	3217	3220	3260	3268	3269	3273
3299	3300	3306	3311	3312	3351	3389	3460	3500	3502	3628
3632	3690	3780	3790	3817	3871	3872	3900	4000	4092	4322
4343	4353	4365	4366	4368	4369	4433	4443	4444	4445	4567
4659	4679	4730	4786	4840	4848	4949	4950	4987	5000	5001
5007	5022	5037	5038	5040	5051	5060	5061	5093	5168	5222
5247	5250	5275	5347	5351	5353	5355	5392	5400	5405	5432
5433	5498	5520	5521	5554	5555	5560	5580	5601	5631	5632
5666	5671	5672	5683	5800	5814	5900	5901	5902	5903	5904
5905	5906	5907	5908	5909	5910	5911	5920	5938	5984	5985
5986	5988	5989	6000	6001	6002	6050	6060	6070	6080	6082
6101	6106	6112	6161	6262	6379	6405	6443	6502	6503	6504
6514	6542	6556	6660	6661	6667	6905	6988	7000	7001	7002
7021	7070	7071	7077	7080	7100	7144	7181	7210	7373	7443

```

7474 7510 7547 7579 7580 7676 7700 7770 7777 7778 7787 7800
7801 7879 7902 8000 8001 8003 8006 8008 8009 8010 8012 8014
8020 8023 8028 8030 8080 8081 8082 8083 8086 8087 8088 8089
8090 8095 8098 8099 8100 8123 8127 8161 8172 8180 8181 8182
8205 8222 8300 8303 8333 8400 8443 8444 8445 8471 8488 8500
8503 8530 8531 8545 8649 8686 8787 8800 8812 8834 8850 8871
8880 8883 8888 8889 8890 8899 8901 8902 8903 8983 9000 9001
9002 9042 9060 9080 9081 9084 9090 9091 9092 9099 9100 9111
9152 9160 9200 9300 9380 9390 9391 9401 9418 9440 9443 9471
9495 9524 9527 9530 9593 9594 9595 9600 9809 9855 9999 10000
10001 10008 10050 10051 10080 10098 10162 10202 10203 10250 10255 10257
10259 10443 10616 10628 11000 11099 11211 11234 11333 12174 12203 12221
12345 12379 12397 12401 13364 13500 13778 13838 14330 15200 15671 15672
16102 16443 16992 16993 17185 17200 17472 17775 17776 17777 17778 17781
17782 17783 17784 17790 17791 17798 18264 18881 19300 19810 19888 20000
20010 20031 20034 20101 20111 20171 20222 20293 22222 23472 23791 23943
25000 25025 25565 25672 26000 26122 27000 27017 27018 27019 27080 27888
28017 28222 28784 30000 31001 31099 32764 32844 32913 33060 34205 34443
34962 34963 34964 37718 37777 37890 37891 37892 38008 38010 38080 38102
38292 40007 40317 41025 41080 41523 41524 44334 44343 44818 45230 46823
46824 47001 47002 47290 48899 49152 50000 50013 50021 50051 50070 50090
50121 51443 52302 52311 54321 54921 54922 54923 55553 55580 57772 61614
61616 62078 62514 65002 65535

```

What UDP ports does runZero scan?

runZero scans the following UDP ports by default:

```

53 69 88 111 123 137 161 443 500 623 987
1194 1434 1701 1900 2049 2228 3391 3671 3702 4433 5060
5246 5349 5351 5353 5632 5683 5684 9302 10000 10001 11211
19132 30718 37810 41794 46808 47808 48808 65535

```

Scanning IoT and OT

Can I safely scan my IoT or OT environments?

Some organizations have IoT or OT equipment sensitive to high traffic rates or malformed packets that may have experienced issues with other scanning tools in the past, resulting in a “don’t scan” rule to be in effect.

runZero is different, and should be able to scan in these environments. runZero provides a lightweight active scan engine called an Explorer that can be deployed almost anywhere. Since the scan is active, there are no tap or span ports that need to be configured, nor device level agents that need to be installed, so you don’t have to modify your environment.

The runZero Explorer was built with sensitive OT environments in mind. It is not based on any other commercial or open source tools such as nmap or masscan. The Explorer only sends

normal traffic, nothing malformed that might potentially crash a fragile system. Some of the controls in place also include:

- packets-per-second scan rates with sensible default values:
 - 1000 packets per second for overall maximum scan rate (adjustable; scan traffic is balanced across all hosts in the scan range)
 - 40 packets per second for per-host maximum scan rate (adjustable)
- IP and TCP port exclusions
- UDP service probes can be enabled or disabled individually
- The scan balances SYN and ACKs and watches for port consumption issues on both the client & target
- Configurable max group size that limits the number of targets runZero can scan at once, which correlates to the number of connections stateful devices such as firewalls or routers receive
- Only those TCP and UDP ports that provide actionable intelligence for fingerprinting a device are checked, not all 65535. This list is adjustable in case specialized equipment runs on a non-standard port (see the [Port List](#)).
- Per port / protocol considerations engineered to avoid issues. Ex: Sending characters to port 9100 on a printer could print "garbage". runZero will collect a banner from some ports such as these but never actively probe them.

Some OT/ICS vendors which runZero can fingerprint upon discovery include:

- Allen-Bradley
- BARIX
- Cisco
- Control Solutions
- Control Techniques
- GE
- GENEREX
- GLC Controls
- Lantronix
- Linor Koda
- Mitsubishi
- Moxa
- PLC
- Pressac
- Rittal
- Rockwell
- Schneider Electric
- Siemens

Many organizations opt to deploy the Explorer to the same system that runs their vulnerability scans since there may already be allow-lists, full network connectivity, and considerations made for session table capacity on any session-aware middle boxes such as firewalls, proxies, or small routers. It may also be advisable to deploy additional Explorers at remote sites to gather additional detail and avoid altogether any need to consider middle boxes.

Browsers supported by the runZero Console

Which browsers are supported when accessing the runZero Console?

We maintain compatibility with the following browser versions or newer:

- *Chrome 110* (released February 6, 2023)
- *Edge 110* (released February 8, 2023)
- *Safari 16.0* (released September 11, 2022)
- *Firefox 115* (released July 3, 2023)
- *Opera 96* (released February 21, 2023)

Internet Explorer is *not* supported.

Mobile browser support is experimental and is not guaranteed to work.

runZero release notes

Latest release notes

4.0.240425.0

2024-04-25

- Fingerprint improvements.

4.0.240424.0

2024-04-24

- Fingerprint improvements.

4.0.240423.0

2024-04-23

- A bug that prevented SSO users from setting a password when SSO was disabled at the runZero account level has been resolved.
- Operating System End of Life (EoL) information is now available for SUSE Enterprise Linux and Apple tvOS.
- Fingerprint improvements.

4.0.240419.0

2024-04-19

- Fingerprinting of assets based on Microsoft 365 Defender data has been improved.
- Fingerprint improvements.

4.0.240417.0

2024-04-17

- Accessibility improvements.
- A bug that could result in errors when deleting a site has been resolved.
- A bug that could cause Wiz tasks to error has been resolved.
- Fingerprint improvements.

4.0.240411.0

2024-04-11

- runZero customers can now sync asset, software, and vulnerability data from Wiz.
- Fingerprint improvements.

4.0.240410.0

2024-04-10

- The runZero dashboard has been improved to better respond to browser window resizing.
- Fingerprint improvements.

4.0.240408.0

2024-04-08

- Data collection from slow SSH services has been improved.
- Fortinet devices are now less likely to cause duplicate assets when traffic is collected using traffic sampling.
- The runZero Explorer now silently skips non-ethernet-like utun (tunnel) interfaces on macOS.
- A bug preventing the "User details" page for external users from loading has been resolved.
- A bug that could lead to errors when changing email address was fixed.
- A bug that could lead to errors when deleting a user was fixed.
- Fingerprint improvements.

4.0.240405.0

2024-04-05

- The profile settings page has been redesigned.
- Names can now be given to multi-factor authentication tokens when enrolling new tokens.

4.0.240404.0

2024-04-04

- Fingerprint improvements.

4.0.240403.0

2024-04-03

- A bug that prevented proper click through from the Query Insights dashboard widget to the appropriate inventory view was fixed.
- Matching of MAC addresses of Fortinet firewall devices was improved.
- Fingerprint improvements.

4.0.240402.0

2024-04-02

- Fingerprint improvements.

4.0.240401.0

2024-04-01

- The layout of the runZero dashboard is now fully customizable.
- The runZero dashboard now supports exporting views as CSV and PNG.
- Fingerprint improvements.

4.0.240331.0

2024-03-31

- Integration task processing is now much faster for assets with large numbers of MAC addresses.
- A bug that could result in assets accumulating link-local IPv6 addresses has been resolved.
- Fingerprint improvements.

4.0.240329.0

2024-03-29

- The “Contact runZero support” menu has been redesigned.
- A bug that could cause the services attribute report to fail has been resolved.
- A bug that could cause hostnames with spaces to be turned into multiple hostnames when imported from the AzureAD connector has been resolved.
- Improved logging for the Intune integration.
- UI improvements.
- Fingerprint improvements.

4.0.240327.0

2024-03-27

- Tenable connector data processing is now significantly faster for devices with large numbers of MAC addresses.
- A bug that could result in the self-hosted updater showing a SQL error during startup has been resolved.
- A bug that could cause scans running on Windows Explorers to accidentally terminate unrelated processes has been resolved.
- Fingerprint improvements.

4.0.240326.0

2024-03-26

- The CrowdStrike connector now only imports actively installed software.
- The CrowdStrike connector now handles large software and vulnerability results reliably.
- The CrowdStrike connector now better filters system accounts from the lastInteractiveUser attribute.
- Fingerprint improvements.

4.0.240325.0

2024-03-25

- Fingerprint improvements.

4.0.240320.0

2024-03-20

- Fingerprint improvements.

4.0.240318.0

2024-03-18

- Task ID is now visible when inspecting a task on the task overview page and on the task details page.
- An issue with calculating mid-scan progress for connector tasks running on Explorers has been resolved.
- A bug that could cause service start issues after upgrading self-hosted runZero instances has been resolved.
- Fingerprint improvements.

4.0.240314.0

2024-03-14

- Colors throughout the product have been tweaked to improve accessibility, legibility, and consistency.
- Tables in the product can now be configured to prefer a mono-spaced variant of the table font.
- Tables throughout the product now allow users to choose text casing preference, available via the “Prefs” dropdown.
- An issue that could prevent updates to Directory Users / Groups has been resolved.
- A bug that could cause the “concurrency” setting on Explorers to be incorrectly changed when editing an Explorer’s settings has been resolved.
- Accessibility improvements.
- Fingerprint improvements.

4.0.240311.0

2024-03-11

- An issue with processing malformed header data from RTSP responses has been resolved.
- The runZero scanner now completes faster for local networks.
- Self-hosted customers can now unbind SSO from a user account using the `runzeroctl user reset` command.
- Self-hosted customers can now change the SSO mode using the `runzeroctl sso-mode mode` command.
- Accessibility improvements.
- Fingerprint improvements.

4.0.240308.0

2024-03-08

- A bug that could cause short keywords to not show any autocomplete suggestions in the query builder has been resolved.
- Long fields in Nmap XML exports of asset data are no longer truncated.
- Probing devices using EtherNet/IP is now supported over UDP.
- Fingerprint improvements.

4.0.240306.0

2024-03-06

- An issue that could prevent new self hosted installations or updating existing installations has been resolved.

- Fingerprint improvements.

4.0.240305.1

2024-03-05

- An issue that could result in incorrect asset merging in certain situations has been resolved.
- An issue that could result in delayed analysis for busy Organizations has been resolved.
- Fingerprint improvements.

4.0.240305.0

2024-03-05

- Fingerprint improvements.

4.0.240304.0

2024-03-04

- Fingerprint improvements.

4.0.240301.0

2024-03-01

- A new “serialNumbers” column has been added to the asset CSV export. This field contains serial numbers observed during scanning, along with the protocol used to discover the serial number.
- An issue that could cause incorrect attack surface assignment to assets discovered by traffic sampling has been fixed.
- A bug which caused some task errors and warnings to fail to display has been fixed.
- Fingerprint improvements.

4.0.240228.0

2024-02-28

- A bug that could prevent sites from being created per project for the Google Cloud Platform integration has been resolved.
- Fingerprint improvements.

4.0.240226.0

2024-02-26

- A bug impacting Operating System End of Life (EoL) assertions for certain versions of Microsoft Windows and Linux distributions has been resolved.

4.0.240223.0

2024-02-23

- A bug that could cause organization statistics to become out of date in organizations with frequent and concurrent tasks has been resolved.
- Operating System End of Life (EoL) information is now available for Apple iOS and iPadOS as well as CentOS Stream.
- Operating System Extended End of Life (EoL) generation has been improved.
- Fingerprint improvements.

4.0.240221.0

2024-02-21

- The vulnerability inventory is now much faster for large organizations.
- Fingerprinting of devices via BGP is now supported.
- Tenable integration performance has been improved.
- A bug that could cause the asset and service attribute reports to fail has been resolved.
- A bug causing some credential form fields to disappear when modifying an existing credential has been resolved.
- An issue with the query format of site-filtered insights has been resolved.
- Fingerprint improvements.

4.0.240218.0

2024-02-18

- Software inventory is now calculated as part of metrics, reducing task processing time.
- A bug that prevented the Organization picker from working on some pages has been resolved.
- Saved queries in the search suggestions menu are now ordered by when they were last updated.
- Improved asset correlation logic for devices with wired and wireless interfaces.
- Improved OS detection logic when considering multiple data sources.
- Fingerprint improvements.

4.0.240216.0

2024-02-16

- Improved correlation behavior for assets with information from NTLMSSP or Qualys.
- Search query and query builder autocomplete results have been improved for shorter sets of input.
- A bug preventing the parent-organization-picker from appearing on the organization create and edit pages has been resolved.
- Fingerprint improvements.

4.0.240214.0

2024-02-14

- Improved protocol detection during traffic sampling.
- The alert event type emitted after a client switch has changed from “login” to “client-switched”.
- The “Site” column has been removed from the software groups table.
- An issue where the software inventory sometimes failed to update after a task has been resolved.

4.0.240213.0

2024-02-13

- The Software Inventory is now much faster for large organizations.
- An issue that could result in stale service attributes persisting through rescans has been resolved.
- The LOG_FORMAT and LOG_MAX_LENGTH configuration values were renamed to RUNZERO_LOG_FORMAT and RUNZERO_LOG_MAX_LENGTH respectively. The old values will continue to work but are deprecated.
- The request timeout for the Qualys integration has been decreased.
- TCP stack based OS fingerprinting has been improved.
- Fingerprint improvements.

4.0.240208.0

2024-02-08

- An issue with adding addresses for Custom Integration assets without MACs has been resolved.
- The request timeout for the Qualys integration has been increased.

4.0.240207.0

2024-02-07

- Additional data points for result count and sent/received data have been added to the Tasks CSV export.
- An issue with the display format of site subnet tags on assets has been resolved.

4.0.240206.0

2024-02-06

- Improved performance on the Software inventory table.
- Additional fields added to Query Builder autocomplete.
- An issue that prevented Site Subnet information from exporting with Assets has been resolved.
- An issue with data missing from the default email template for alerts has been resolved.

4.0.240205.0

2024-02-05

- Filtering of hostnames collected from TLS X.509 certificates has been improved.
- An issue that could cause overlapping subnets to apply another Site's subnet tags has been resolved.
- An issue that could result in incorrect asset correlation between HP iLOs and their servers has been resolved.
- Fingerprint improvements.

4.0.240202.0

2024-02-02

- Performance of Tenable.io connector tasks when only a subset of Severity/Risk values are selected has been improved.
- An issue that allowed users with the Administrator role to downgrade their own permissions has been resolved.
- An issue that could prevent Nessus attributes from being fully hydrated by runZero has been resolved.
- Fingerprint improvements.

4.0.240131.0

2024-01-31

- Fingerprint improvements.

4.0.240129.0

2024-01-29

- A query builder is now available, accessible from most datagrids by clicking the “Query builder” button to the right of the search bar.
- An issue which caused some out-of-date service information to remain on assets has been resolved.
- An issue which caused service information to be incorrectly removed from assets that were offline during a scan has been resolved.

4.0.240126.0

2024-01-26

- Discovery of devices using the DNP3 protocol is now supported.
- Operating System End of Life (EoL) information is now available for Oracle Linux.
- Page break locations in the overview report have been improved.
- Operating System End of Life (EoL) generation for Red Hat Enterprise Linux and CentOS Linux has been improved.
- Assets with no known address are now labeled with “Unknown” for their address rather than “Unscanned”.
- The bundled npcap driver has been updated to version 1.79.
- An issue that could prevent last task details from correctly displaying on the Sites datatable has been resolved.
- An issue that prevented the expansion of dropdown menu sub-menus using keyboard navigation has been resolved.
- An issue that could result in certain OS fingerprinting data not being updated has been resolved.
- An issue that could prevent creating new Azure Credentials via the Azure connector configuration page has been resolved.
- An issue causing Tenable.io integration tasks to import vulnerability data even when no severity or risk levels were selected has been resolved.
- Fingerprint improvements.

4.0.240124.0

2024-01-24

- An issue that could result in hidden fields on the SNMP v3 Credentials form has been resolved.
- Fingerprinting of Red Hat Enterprise Linux derivatives when limited data is available has been improved.

- Additional fingerprint improvements.

4.0.240122.0

2024-01-22

- The datagrid search bar has been improved to show recent queries and available queries from the query library.
- Fingerprinting of Red Hat Enterprise Linux and derivatives from Tenable product data has been improved.

4.0.240119.0

2024-01-19

- Fixed an issue that prevented the “Edit user permissions” modal from working correctly.
- Fingerprint improvements.

4.0.240117.0

2024-01-17

- Fixed an issue where custom integration task data could not be re-imported.
- Fixed an issue where Nessus imports could fail due to Nessus response size.
- Fingerprinting of Red Hat Enterprise Linux derivatives such as CentOS, Rocky Linux, and Oracle Linux has been improved.
- Fingerprint improvements.
- Accessibility improvements.

4.0.240112.0

2024-01-12

- Site column has been added to all tasks lists in the task overview.
- Fingerprint improvements.

4.0.240110.0

2024-01-10

- The Nmap XML export now uses the minimum and maximum asset last_seen timestamps as the start and stop times.
- An issue that could prevent stale services from being cleared from updated Assets has been resolved.
- A resource leak that affects self-hosted customers with transparent huge pages (THP) enabled has been resolved.

- Fingerprint improvements.

4.0.240109.0

2024-01-09

- Tenable Security Center tasks now only retrieve records updated since the previous sync.
- Fingerprint improvements.

4.0.240105.0

2024-01-05

- A bug that prevented the API for creating passive sampling tasks from working as documented was fixed.
- A bug that could cause inventory grids to disappear when using Firefox and resizing the window below a certain point has been resolved.
- Improved error handling for Tenable, Tenable Security Center, and CrowdStrike integrations.
- Fingerprint improvements.

4.0.240103.0

2024-01-03

- Improved correlation for assets sourced from the Censys and Shodan integrations.
- A bug that incorrectly logged certain task failures as 'explorer failed to queue task' has been resolved.
- Fingerprint improvements.

4.0.231222.1

2023-12-22

- A bug that incorrectly set empty private IP ranges instead of default values has been resolved.
- A bug that prevented stale Azure scale set attributes from being cleared has been resolved.

4.0.231222.0

2023-12-22

- A bug causing duplicate line items in CSV exports of site configurations has been resolved.

- Most modal interfaces in the UI can now be dismissed using the escape key.
- MAC address assignment for certain HP servers with iLO devices has been improved.
- Fingerprint improvements.

4.0.231220.0

2023-12-20

- Custom Integration assets now support the inclusion of Service attribute data.
- Improved correlation for Custom Integration assets.
- Fixed a bug in the service inventory table which showed incorrect source icons.
- Fingerprint improvements.

4.0.231218.0

2023-12-18

- Fixed an issue with third-party services matching on existing runZero services.
- Fingerprint improvements.

4.0.231215.0

2023-12-15

- Improved correlation for Tenable sourced assets.
- A bug that could cause a redirect loop when logging in if the user hasn't agreed to the latest terms and MFA is required but not set has been resolved.
- A bug preventing self-hosted customJS from running on all authorized pages has been resolved.
- Fingerprint improvements.

4.0.231213.0

2023-12-13

- The memory requirements for Explorers have been updated to bring them in line with documentation.
- The CVEs column of the Queries datagrid has been made un-sortable.
- A bug that could cause datagrid tables to become un-sortable if an overlay is shown has been resolved.
- Fingerprint improvements.

4.0.231211.0

2023-12-11

- Fingerprint improvements.

4.0.231208.0

2023-12-08

- A bug that could cause the loading state on datagrids to be incorrect when interrupting a running search has been resolved.
- Fingerprint improvements.

4.0.231207.0

2023-12-07

- A bug which could cause tasks to get stuck in “stopping” status was fixed.
- A bug that could prevent accurate goal progress for goals associated with vulnerability queries has been resolved.
- Self-hosted customers can now load configuration items from multiple Secrets Manager keys by separating these with commas in the `AWS_SECRET_ACCESS_KEY` variable.
- Accessibility improvements.
- Fingerprint improvements.

4.0.231205.0

2023-12-05

- A bug that could cause successful tasks to be marked as stale has been resolved.
- A bug where SNMPv3 authentication was not correctly indicated on an asset has been resolved.
- Fingerprint improvements.

4.0.231201.1

2023-12-01

- A bug that could prevent scrollbars from being shown on Datagrids has been resolved.

4.0.231201.0

2023-12-01

- A bug preventing the arm64 versions of FreeBSD, NetBSD, and OpenBSD Explorers and Scanners from being downloaded has been resolved.

- The Datagrid page selector has been improved.
- Datagrid columns can now be manually resized beyond the visible width of the grid.

4.0.231130.0

2023-11-30

- Added support for discovering assets using the S7 protocol.
- Passive detection of RDP authentication methods has been improved.
- Action buttons in the inventory have been redesigned for a better user experience.
- A bug that could cause inventory tables to display too many pages has been resolved.
- A bug that could prevent import of Azure AD users and groups has been resolved.
- Fingerprint improvements.

4.0.231128.0

2023-11-28

- Improved explorer memory usage in some situations.
- A bug preventing the appearance of the task failure reason column has been resolved.
- A bug causing an inaccurate scan progress label has been resolved.
- A bug where strings with multi-byte unicode characters were measured incorrectly by the Custom Integrations SDK has been resolved.
- A bug that could cause an inaccurate asset changed count in the task summary view has been resolved.
- Fingerprint improvements.

4.0.231122.0

2023-11-22

- A bug that incorrectly omitted the OS EOL for some Linux variants has been fixed.
- Improved detection of Fortinet devices.
- Fingerprint improvements.

4.0.231121.0

2023-11-21

- A bug that prevented editing a credential has been fixed.
- The asset JSON export now includes extended ownership data under the 'ownership' field.
- Improved fingerprinting of Cisco devices.
- Fingerprint improvements.

4.0.231120.0

2023-11-20

- A bug causing the Task Library datagrid to be unusable has been resolved.
- A bug preventing some buttons from rendering properly has been resolved.
- A bug preventing the use of the select-all button on some datagrids has been resolved.
- The self-hosted platform now allows users to override the S3 region and endpoint.
- Improved Cisco small business device fingerprints.
- Fingerprint improvements.

4.0.231116.0

2023-11-16

- The performance of the asset attribute report has been improved for third-party sources.
- Ddatagrid column width will now persist when manually resized.
- Ddatagrid column width and visibility can now be managed more granularly through two new dropdown menus on all datagrids.
- Deleting an organization via the API now returns the proper “404” HTTP status code if the organization to be deleted does not exist.
- Fingerprint improvements.

4.0.231115.0

2023-11-15

- Detection of bulk responses from Fortinet network filtering products was improved.
- A bug that prevented asset matching in certain rare cases involving virtual IP addresses or heavy IP reuse has been resolved.
- A bug preventing the datagrid on the Monitor landing from displaying has been resolved.
- A bug causing action buttons on the task datagrids to be visually cut off at the top has been resolved.
- A bug preventing the SSO Group Mapping datagrid from loading has been resolved.
- Fingerprint improvements.

4.0.231114.0

2023-11-14

- A bug that could cause explorer binaries to be inadvertently deleted has been fixed.

4.0.231108.1

2023-11-08

- The AzureAD integration now supports filtering inactive devices.
- SSO group mapping rules that use a DN containing commas are now supported.
- The runZero Explorer can clean up stale files in both the legacy Rumble and runZero install locations.

4.0.231108.0

2023-11-08

- The platform now supports configurable private IP address ranges at the Account and Organization level.
- Added detection of the ThinPrint protocol.
- Added support for the legacy ident protocol.
- Improved scanning of OpenVMS systems.
- Improved OS fingerprinting of data from custom integrations.
- Improved detection of embedded Linux devices.
- A bug that could prevent best-effort parsing of inventory queries with warnings has been resolved.
- A bug that prevented the rendering of the Asset and Service trend charts on the Dashboard has been resolved.
- Fingerprint improvements.

4.0.231102.0

2023-11-02

- A bug that could prevent tasks from processing for large customers has been resolved.

4.0.231101.0

2023-11-01

- A bug that could result in excessive memory consumption when processing asset modification rules has been resolved.
- The scanner now supports the IGEL thin client protocol.
- Fingerprint improvements.

4.0.231031.2

2023-10-31

- A bug that could prevent retrying http requests has been resolved.

4.0.231031.1

2023-10-31

- Improved retry handling in third-party connectors.

4.0.231031.0

2023-10-31

- A bug that could lead to a stall in CrowdStrike connections has been resolved.

4.0.231030.0

2023-10-30

- The scanner and passive engine now supports the Kasa IoT protocol.
- The scanner now reports the Shodan-compatible Murmur3 32-bit hash for favicon files.
- A bug that could prevent users from setting per-organization permissions for the primary organization has been resolved.
- A bug that could display invalid action buttons for External Users has been resolved.
- A bug that could cause inactive users to be omitted from the External Users table has been resolved.
- A bug that could cause the External Users table to omit group permissions from the Org Access summary has been resolved.
- Performance improvements.
- Fingerprint improvements.

4.0.231027.1

2023-10-27

- Fixed an issue where Asset correlations were not being recorded for some existing assets.

4.0.231027.0

2023-10-27

- Added API endpoints to allow setting asset criticality individually or in bulk.
- Added support for updating asset criticality via CSV import.

- A bug that could cause a user's browser to appear to be stuck in a refresh loop on login has been resolved.
- Improved correlation between Azure AD, Microsoft Intune, and Microsoft 365 Defender assets.
- Improved time-based filtering when requesting data from the Tenable Security Center API.
- Improved reliability of the CrowdStrike connector.
- Improved reliability of the Microsoft Intune connector.
- The asset details page has been updated with additional correlation information for attributes, including latest task details.
- Added support for fingerprinting devices based on the SecuRemote protocol.
- Improved probing of certain types of printers and print servers.
- Fingerprint improvements.

4.0.231023.0

2023-10-23

- A bug that could cause query-reported vulnerabilities to be improperly cleared has been resolved.
- A bug that could result in invalid Tenable Security Center assets has been resolved.
- Fingerprint improvements.

4.0.231019.0

2023-10-19

- Improved Crowdstrike vulnerability request handling.
- Fingerprint improvements.

4.0.231018.0

2023-10-18

- Performance improvements.
- Fingerprint improvements.

4.0.231017.1

2023-10-17

- A bug that prevent the type keyword from being processed correctly in task search has been resolved.

4.0.231017.0

2023-10-17

- A button for deleting a single credential from the credentials list was added to each row's actions column.
- Improved merge accuracy for externally scanned AWS assets reported by third-party sources.
- Fingerprint improvements.

4.0.231013.0

2023-10-13

- The asset inventory now shows the Site as the first column on the left.
- Performance improvements.

4.0.231012.0

2023-10-12

- Custom integrations now support the inclusion of vulnerability and software data.
- Improved merge accuracy for attribute based matching and invalid TLS serial IDs.
- Improved MAC vendor accuracy for AWS assets reported via CrowdStrike.
- The `attribute_count` search term has been added to the asset inventory.
- Events created by editing alert channels, rules or templates will now correctly trigger the corresponding alert rules.
- A bug that could display an incorrect resource warning for explorers has been resolved.
- Updated TLS fingerprinting for new Go versions.
- Performance improvements.

4.0.231011.0

2023-10-11

- A bug that could result in a random authorized organization being selected on login has been resolved.
- Performance improvements.

4.0.231006.0

2023-10-06

- A low severity security issue with unsafe request binding has been resolved.

- A bug with mDNS processing that resulted in incorrectly merged assets has been resolved.
- A bug requiring users to re-input the “Authentication passphrase” and “Privacy passphrase” fields when editing an SNMPv3 credential has been resolved.
- Improved performance of third-party user and group imports.
- Long values for asset-level and service-level attributes on the Asset Details page are now truncated for display.
- The self-hosted platform now supports Amazon Linux 2.
- The npcap driver has been updated to version 1.77.
- Improved identification of devices using the Cisco Discovery Protocol.
- Detection of MAC addresses for CradlePoint devices has been improved.
- Fingerprint improvements.

4.0.231005.0

2023-10-05

- A low severity security issue with mass assignment has been resolved.
- A bug causing a “failed to parse query” error when sorting recurring tasks has been resolved.

4.0.231002.0

2023-10-02

- Improved merge accuracy for externally scanned AWS assets reported by third-party sources.
- Fingerprint improvements.

4.0.230928.0

2023-09-28

- Improved detection of spurious and phantom services generated by some network devices during probing.
- Fingerprint improvements.

Important security fix:

- A SQL injection vulnerability was identified and fixed as part of our annual third-party security assessment.

4.0.230927.0

2023-09-27

- Loading time performance on the Scan and Monitor landing pages has been improved.

- Added support for asset detection using the OMRON FINS protocol.
- A bug preventing the self-hosted console from correctly logging certain scan data download errors has been resolved.
- Fingerprint improvements.

4.0.230925.0

2023-09-25

- Welcome to the [runZero Platform version 4.0!](#)

3.10.230921.0

2023-09-21

- A bug which caused a continuous page refresh when logging in via SSO has been resolved.

3.10.230920.0

2023-09-20

- A performance regression when processing third-party assets has been resolved.
- Improved merging of assets with NetBIOS or SMB services.
- The tasks CSV export now includes the template_name column.
- The tasks JSON export and API responses now include the site_name, agent_name, and template_name columns.
- Fingerprint improvements.

3.10.230918.0

2023-09-18

- A bug that could prevent new SSO users from authenticating has been resolved.
- Fingerprint improvements.

3.10.230917.1

2023-09-17

- A bug that could prevent some Windows-based Explorers from connecting with the same ID has been resolved.
- Fingerprint improvements.

3.10.230917.0

2023-09-17

- A bug that could prevent the Explorer from reading the .env configuration file has been resolved.
- A bug that could prevent Tenable Security Center syncs from completing has been resolved.
- A bug that could lead to bogus assets appearing in scans through Fortigate proxies has been resolved.
- A number of small parsing bugs in the protocol parsing engine have been resolved.
- Passive traffic sampling tasks now set source:sample instead of source:passive for assets.
- The self-hosted console now uses the “runZero” brand (and runzeroctl command) by default.
- The self-hosted console now defaults to PostgreSQL 15 and provides an install option to select a version.
- The Explorer now uses the “runZero” brand by default (and matching filesystem/registry locations).
- Fingerprint improvements.

3.10.230913.0

2023-09-13

- Added support for identifying scanners using the eSCL uscan protocol.
- Custom Integrations now support the “exclude unknown” option.
- Improved handling of API request retries for the Microsoft Intune integration.
- Improved detection of spurious services when scanning certain firewalls.
- The Tenable integration has been updated to reduce the possibility of asset and vulnerability export timeouts.
- An issue that could result in login errors for invited users using Single Sign-On has been resolved.
- Fingerprint improvements.

3.10.230911.0

2023-09-11

- The Tenable integration has been updated to reduce the possibility of vulnerability export timeouts.
- Fingerprint improvements.

3.10.230908.01

2023-09-08

- Moved to a new versioning scheme for the Console and Explorers, <major>.<minor>.<yymmdd>.<revision>.
- Added support for fingerprinting devices with the Voice Services Discovery Protocol (VSDP).
- Improved fingerprinting of devices using Spotify Connect.
- Improved detection of BACnet devices.
- Additional Fingerprint improvements.

3.10.25

2023-09-05

- A warning is now recorded for scan tasks if a host is ignored for responding on too many ports.
- The Integrate page now shows active and suggested integrations for the current organization.
- Fingerprint improvements.

3.10.24

2023-08-30

- An issue that could result in an empty dashboard until a metrics recalculation was triggered has been resolved.
- Fingerprint improvements.

3.10.23

2023-08-29

- An issue that could result in stalled scans has been resolved.

3.10.22

2023-08-29

- An issue that could result in an empty dashboard when selecting a single site has been resolved.

3.10.21

2023-08-29

- Discovery of devices using the MODBUS/TCP protocol is now supported.
- The maximum number of ownership types has been increased from 10 to 25.
- A bug that could result in duplicate software entries for some sources has been resolved.
- Improved fingerprinting of devices that provide UPnP information.
- Additional Fingerprint improvements.

3.10.20

2023-08-25

- A bug that prevented SNMPv3 credentials from being saved has been resolved.

3.10.19

2023-08-23

- A bug that could result in an incorrect ts attribute for Azure AD, Google Workspace, and Microsoft Intune has been resolved.
- Fingerprinting of assets based on Microsoft Intune and Microsoft 365 Defender data has been improved.
- Discovery of assets via EtherNet/IP probing is now supported.
- Fingerprint improvements.

3.10.18

2023-08-21

- Event rules now support conditions for Explorer and task type, where relevant.
- Recurring tasks now stop with an error if they use a task template that has been deleted.
- A rotation date for stored credentials is now available through both console and API via a new `secret_updated_at` field.
- CrowdStrike and Azure AD assets will no longer be merged if they have a different globally unique ID. This may lead to more offline assets being generated if devices are frequently reimaged and given new GUIDs.
- A bug preventing some users from being able to manage their user's group membership has been resolved.
- A bug in the user permissions display interface has been resolved.
- A bug that could cause foreign service attributes to be attributed to the wrong source has been resolved.
- Fingerprint improvements.

3.10.17

2023-08-16

- The scanner now supports the MQTT protocol.
- The TCP SYN scanner is now friendlier to stateful firewalls in the network path.
- Tasks in the stopping state are now included in the Processing section of the Tasks overview.
- A bug where users logging in for the first time with SSO would not have access to any organizations from the SSO group mappings has been resolved.
- A bug resulting in incorrect fingerprinting of assets based on AzureAD data has been resolved.
- A bug that could prevent bogus services from certain firewalls from being completely filtered has been resolved.
- Fingerprint improvements.

3.10.16

2023-08-14

- A bug causing tasks in the process of stopping to be seen as dismissible has been resolved, so that only failed and completely stopped tasks can be dismissed.
- A bug causing pending new tasks to be seen as editable has been resolved, so that only new tasks scheduled to start in the future can be modified.
- A bug that prevented download commands from being displayed on the redesigned scanner page has been resolved.
- A bug where existing assets were incorrectly fingerprinted after importing data from Microsoft 365 Defender has been resolved.
- A bug causing incorrect assertion of Microsoft Defender for Endpoint in `edr.name` has been resolved.
- Error logging for the Shodan integration has been improved.
- Fingerprint improvements.

3.10.15

2023-08-10

- The Explorer and scanner download pages have been redesigned for improved UX and performance.
- The Tenable integration has been updated to reduce the possibility of vulnerability export timeouts.
- A bug that could cause scan templates to be hidden when configured with invalid permissions has been resolved.

3.10.14

2023-08-09

- A bug that prevented the scan.explorer_id value from being populated in alert templates has been resolved.
- A bug that prevented the “Find assets in this site” icon from working properly in some cases has been resolved.
- A bug that could prevent queries containing mixed-case search terms from returning results has been resolved.
- A bug that prevented some queries from correctly matching Intune assets has been resolved.
- The scanner now supports a new syn-reset-sessions option that can be used to reduce session usage in middle boxes.
- Fingerprint improvements.

3.10.13

2023-08-07

- A security improvement has been added to clear password reset tokens after a password change or when link-based authentication is requested.
- A bug that could prevent Intune assets from merging with other sources has been resolved.
- A bug causing inconsistent navigation for Explorer configuration editing has been resolved.
- Fingerprint improvements.

3.10.12

2023-08-03

- A bug that could result in the wrong hostname being used in password reset links has been resolved.
- A bug that could cause a memory leak in the Explorer between stopped tasks has been resolved.
- A bug that resulted in the Nmap XML Export having a zero start time has been resolved.
- Fingerprint improvements.

3.10.11

2023-07-31

- A bug that could cause some long-running connection tasks to restart repeatedly has been resolved.
- A bug that could result in invalid Shodan credentials still validating has been resolved.

- A bug that could lead to improper stale service removal on rescan has been resolved.
- A bug that could lead to orphaned tasks when an Explorer is removed has been resolved.
- Fingerprint improvements.

3.10.10

2023-07-27

- A bug causing the user details page to display permissions incorrectly has been resolved.
- A bug causing project expiration to be miscalculated has been resolved.

3.10.9

2023-07-26

- Task name and description can now be modified for tasks created via file imports.
- Improved automatic asset filtering for certain web proxy assets.
- Improved performance when deleting large organizations, projects, or sites.
- The request timeout has been increased for the Microsoft Intune and Azure AD integrations.
- Fingerprinting of SMB 1 endpoints has been improved.
- A bug that could result in an unnecessary screenshot warning for connector tasks has been resolved.
- Additional Fingerprint improvements.

3.10.8

2023-07-24

- Exports of task data now include timestamps which differentiate time spent acquiring data from time spent processing data.
- Task processing times are improved.
- A bug causing app banners to not be visible has been resolved.
- The `api/v1.0/org/sites/{site_id}/import` route now returns the proper 400 http status code error when the request body is empty instead of a status code 500.

3.10.7

2023-07-21

- A bug that could prevent in-scope, unscanned addresses from being cleared on runZero assets has been resolved.
- A bug causing the datepicker to close when navigating by year has been resolved.

3.10.6

2023-07-19

- A bug that could prevent import of wireless networks has been resolved.

3.10.5

2023-07-19

- The asset details screen now has pagination when viewing an asset with more than 30 services.
- A bug that enabled SNMP credentials when modifying or copying existing scan tasks has been resolved.
- Fingerprint improvements.

3.10.4

2023-07-18

- A bug preventing columns from retaining their custom ordering has been resolved.
- A bug that could limit the available information gathered from NFS servers has been fixed.
- Fingerprint improvements.

3.10.3

2023-07-17

- A bug that could cause tasks to be copied with an incorrect discovery scope has been resolved.

3.10.2

2023-07-17

- On-screen text explaining the interaction between a user's default organization role and the granted per-org role is clearer.
- License-based size limits are now applied to file imports.
- An issue that caused the asset details page to load very slowly has been resolved.
- A bug that could result in duplicate service warnings has been resolved.
- A bug that could result in orphaned tasks when removing an explorer has been resolved.
- A bug which could leave SYN and LAYER2 probes in a perpetual error condition loop has been resolved.

- A bug that prevented display of the user permissions table in the User Details screen has been resolved.
- Fingerprint improvements.

3.10.1

2023-07-12

- An update to the runZero Explorer now logs when the host operating system receives an interrupt or terminate signal, such as when the OS reboots.
- An update for improved asset matching for tasks importing both scan and third-party data sources has been added.
- Fingerprint improvements.

3.10.0

2023-07-10

- An integrations page has been added to improve visibility and simplify configuration.
- A bug that showed a discrepancy in organization project status date on tooltip hover has been resolved.
- Fingerprint improvements.

3.9.10

2023-07-06

- A bug that could cause the MDNS probe to panic in limited scenarios has been resolved.

3.9.9

2023-07-06

- An issue that could result in the old Explorer details pages being shown has been resolved.
- A bug preventing Microsoft 365 Defender OAuth Client Credential tokens from accessing Azure government environments has been resolved.
- A bug that could result in invalid Last Seen values for Rapid7 assets has been resolved.
- A bug that could lead to stale service entries has been resolved.
- A bug causing some goals to return an error has been resolved.
- Fingerprint improvements.

3.9.8

2023-07-05

- Fingerprint improvements.

3.9.7

2023-07-03

- Assets with hostnames starting with a numeric prefix are now allowed to merge.
- Inventory searches using keyword organization properly warn that it cannot be used unless either that specific organization, or the All Organizations option, are chosen from the drop-down in the upper right of the console.
- An issue that could prevent alert rule actions from modifying asset ownership based on software, service, or vulnerability query results is resolved.
- An issue where dynamic content did not have the header Cache-Control: no-store has been resolved.
- Fingerprint improvements.

3.9.6

2023-06-28

- Improved detection of various printer models.
- The Explorer Details page has been redesigned.
- Additional bugfixes and performance improvements.
- Fingerprint improvements.

3.9.5

2023-06-27

- Improved database performance for asset, site, and organization delete operations.
- Improved database performance for outlier and vulnerability processing.
- Improved database performance for concurrent integration processing.
- Additional MAC address detection through SSDP and UPnP services.
- Fingerprint improvements.

3.9.4

2023-06-26

- Improved operating system and hardware fingerprinting of Palo Alto Networks devices.
- Trial accounts can now create Custom Integrations.
- Discovery of SSDP services has been improved.
- A bug has been fixed that could cause scans to be dropped with “explorer failed to queue task” when the Explorer was already handling the configured maximum number

- of simultaneous scans.
- Fingerprint improvements.

3.9.3

2023-06-22

- A bug causing the task start time to be shown for the scan start time has been resolved.
- A bug that could prevent the creation of new goals has been resolved.
- A bug that could prevent those with the 'annotator' role from viewing or modifying Asset Ownership has been resolved.
- Fingerprint improvements.

3.9.2

2023-06-20

- Improved handling of email send errors.
- Improved import of assets from Azure Active Directory.
- Asset correlation has been improved for switches with overlapping MAC addresses.
- Improved detection of AIX systems.
- Reduced OS fingerprinting false positives against assets with non-Microsoft SMB stacks.
- An issue that could prevent navigation to the Account settings page has been resolved.
- A bug causing JavaScript errors to be thrown when adding or editing Google Workspace connector tasks has been resolved.
- A bug with thumbprint validation for the LDAP integration has been resolved and the related error messages have been improved.
- A bug where the link to help for query syntax led to a missing page has been resolved.
- A bug preventing the Explorer interface and addresses from being populated has been addressed.
- Fingerprint improvements.

3.9.1

2023-06-14

- Improved handling of login tokens.

3.9.0

2023-06-13

- runZero goals is now generally available. With runZero goals, users are able to create and monitor progress toward achieving security initiatives.
- Improved the goal progress chart display to work in various browser sizes.

- Goals now shows a pending calculation banner when goal metrics have not been calculated yet.
- Added `source_count` and `custom_integration_count` as searchable fields.
- Saved queries can now be created for tasks.
- The search keyword `recur_last_task_status` is now supported on the task pages.
- Improved the display of dashboard charts so that no partial rows, other than the last row, are visible to the user regardless of the number of charts displayed.
- Improved fingerprinting of Fortinet device firmware.
- Optimized database utilization and improved performance.
- A bug causing Cisco 8xx Industrial Routers as well as Catalyst 94xx/95xx switches to be incorrectly merged has been fixed.
- A bug where the autocomplete drop down would not always appear on top of other elements has been resolved.
- A bug where integration sources in dashboard views are displayed as IDs instead of names has been resolved.
- Fingerprint improvements.

3.8.14

2023-06-06

- Improved protocol feature extraction for a variety of protocols.
- A bug where data grid search text would propagate to other data grids has been resolved.
- A bug causing some text inputs to display an autocomplete user experience when it was not intended was resolved.
- Fingerprint improvements.

3.8.13

2023-05-31

- Non-runZero asset sources can now be removed from assets via the asset details or asset inventory pages.
- Equivalent emails are now accepted for email updates.
- Dashboard cards for Asset Source and Custom Integrations should now correctly show only the top 10 counts for each, with a “View more” link added.
- A bug that could allow merging AWS, Azure, and GCP assets has been resolved.
- A bug which omitted some SNMPv3 scan attributes has been resolved.
- A bug which caused some project creations to return a 404 error page has been resolved.
- A bug causing incorrect HTTP response codes for the `/org/metrics/{site_id}` API endpoint has been resolved.
- A bug which cleared the organizations table screen when sorted has been resolved.
- A bug preventing vulnerabilities from sorting correctly on CVSS columns has been resolved.
- A bug where scan tasks on hosted zones couldn't be stopped has been resolved.
- Fingerprint improvements.

3.8.12

2023-05-24

- A warning is now displayed if a Query is not attached to a Goal.
- Users with “Viewer” permission can now see and use the “Sites” page.
- Improved the normalization of certain service attributes.
- A bug preventing vulnerabilities from sorting correctly on CVSS columns has been resolved.
- Fingerprint improvements.

3.8.11

2023-05-22

- A bug that could result in excessive memory usage has been resolved.

3.8.10

2023-05-22

- Improved SentinelOne matching to improve asset merging.
- AWS credential validation now always shows the results for each service.
- A bug that resulted in certain models of Cisco routers being incorrectly merged has been resolved.
- A bug in which AWS probes fail when run outside of an AWS EC2 environment has been resolved.
- A bug which prevented IPv6 UDP SYN scans from working on FreeBSD and OpenBSD systems has been addressed.
- A bug where autocomplete suggestions would not update consistently has been resolved.
- A bug causing the “download task button” to show for tasks without a log has been resolved.
- Fingerprint improvements.

3.8.9

2023-05-19

- A bug that could cause the SNMP probe to panic in rare scenarios has been resolved.

3.8.8

2023-05-19

- Improved reliability of scans so they should stall less frequently.
- A bug that could cause the SNMP probe to stall scans in rare scenarios has been resolved.

3.8.7

2023-05-19

- The activation email should display properly in a broader range of email clients.
- A bug that caused scans to take longer than expected or stall in rare scenarios has been resolved.

3.8.6

2023-05-18

- A bug that could prevent the organization drop down from being clickable has been resolved.

3.8.5

2023-05-17

- A bug that could prevent the rpcbind probe from completing successfully was resolved.

3.8.4

2023-05-17

- A bug with copying some connector tasks has been resolved.
- A bug causing some connectors to be labeled as scans has been resolved.
- Improved operating system fingerprinting via SNMP Installed Software listing.
- The status indicator in the explorer datagrid now has text describing the status.
- Fingerprint improvements.

3.8.3

2023-05-15

- External Asset Report Include screenshots toggle now requires that Include asset details is checked.
- External Asset Report now hides the Top certificate authorities section if Include TLS certificate details is not checked.

- A bug causing the API /org/hosted-zones endpoint to return an empty list of hosted zones has been resolved.
- A bug that could result in an invalid asset ownership assignment has been resolved.
- Improved fingerprinting of Brother scanners.

3.8.2

2023-05-12

- Outlier calculations have been adjusted for performance and now include the TLS stack.
- Event rules that result in asset modifications now complete faster.
- The npcap driver has been updated to version 1.75.
- Added support for legacy Internet time-distribution and traffic testing protocols.
- A bug that could prevent a RUMBLE_CONSOLE override from working in the Explorer configuration has been resolved.
- A bug that prevented sites with more than 1000 subnets from being saved has been resolved.
- A bug that could result in odd dashboard chart behavior has been resolved.
- A bug that required self-hosted users to configure SMTP before setting up their initial account has been resolved.
- A bug that caused some scan task errors to be displayed twice has been resolved.
- Fingerprint improvements.

3.8.1

2023-05-11

- Improved the ability to fingerprint devices using DLNA.
- Improved device type identification of Windows Server assets.
- A bug that could prevent bogus services from certain firewalls from being completely filtered has been resolved.
- A bug where Asset queries for exact strings was performing a fuzzy search has been fixed.
- A bug that could cause malformed auto-populated LDAP thumbprints for LDAP credentials has been resolved.
- A bug that prevented credential validation errors from displaying after verification in the console has been resolved.
- A bug where searching via clicking on a tag would not return the correct results has been resolved.
- A bug where multiple subtasks were incorrectly created for the same parent task has been resolved.
- A bug where filters were not retained when importing a Nessus scan configuration has been resolved.
- A bug that prevented copying of some connector tasks has been resolved.
- A bug with linking to the update page on some connector tasks has been resolved.
- Fingerprint improvements.

3.8.0

2023-05-09

- Risk and criticality levels can now be assigned to assets through third-party integrations, the asset inventory, and custom rules.
- [runZero Preview Program](#): Goal tracking helps users with Professional and Enterprise licenses track progress toward completing their security initiatives. Use built-in goals for asset ownership coverage or system queries, or create goals with custom queries to fit your needs.
- runZero system and custom queries can now be used to create vulnerability records.
- Passwordless authentication is now available, allowing users to request one-time authentication links via email rather than storing a password. This provides a secure authentication alternative when SSO cannot be configured.
- Added support for Azure and Intune GCC, GCC High, and DoD environments.
- Improved compatibility with WireGuard and Tailscale on macOS and *BSD.
- Added support for searching software attributes.
- Alert channels now support more than one email address.
- Asset limit warnings have been updated to be more clear about whether or not scans will be affected.
- Assets will now be merged based on hostname if no other match method succeeds, as long as the hostname is from a trusted source — currently mDNS, NTLM or NetBIOS.
- A bug preventing explorer reassigned to a previous organization and picking up assigned tasks has been resolved.
- A bug causing software search links to navigate to a 404 page has been resolved.
- A bug causing task-failed events to ignore the site restriction has been resolved.
- A bug causing the hostname override tag to not update the hostname displayed has been resolved.
- A bug that prevented clearing Insights from the dashboard has been resolved.
- A bug where the copy scan button was cut off in the recurring tasks tab has been resolved.

3.7.11

2023-05-03

- A bug that could result in a panic while performing a scan has been resolved.
- A bug that could prevent the API from creating valid scan tasks has been resolved.
- A bug that negatively impacted fingerprinting via TLS certificates has been resolved.
- A bug preventing TLS negotiation in some cases has been resolved.

3.7.10

2023-05-02

- A bug that was triggered when submitting Azure credentials for verification with a subscription ID has been resolved.

3.7.9

2023-05-02

- Credential verification is now allowed only after all required fields have been completed.
- A bug that could cause deadlocks in the TCP LDAP probe and Active Directory integration has been resolved.
- A bug that caused an infinite redirect when clicking on site breadcrumbs has been resolved.
- A bug causing recurring tasks to be incorrectly sorted by start time on the tasks page has been resolved.
- A bug allowing “Verify & save” on the credentials update page to error has been resolved.
- A bug where Dell laptops were identified as desktops or servers has been resolved.
- A bug preventing TLS negotiation in some cases has been resolved.
- A bug that caused imported queries to be parsed improperly has been resolved.

3.7.8

2023-04-26

- A bug with the default webhook Slack alert template has been resolved.
- Improved error message when attempting to delete a scan template twice.
- Improved fingerprinting of Brother scanners.
- Fingerprint improvements.

3.7.7

2023-04-24

- A recent update in Explorer and Scanner behavior which could inadvertently trigger CrowdStrike EDR detection has been disabled.

3.7.6

2023-04-24

- Grace period for tasks can now be configured from the task template page.
- Improved asset correlation for multi-source assets.
- A bug regarding Intune rate limiting and intermittent failures has been resolved.
- A bug where certain tasks could not be edited has been resolved.
- A bug regarding erroneously returned results from unscanned runZero assets when searching the asset inventory has been resolved.
- A bug marking assets “unscanned” has been resolved.
- Fingerprint improvements.

3.7.5

2023-04-20

- A bug that resulted in a 500 error when running the asset attribute report has been resolved.
- Miscellaneous bug fixes.

3.7.4

2023-04-19

- Public API endpoints to view hosted zones have been added.
- The API endpoints for managing scan tasks now accept an argument to select a hosted zone.
- Third-party vulnerability integrations now support a more granular risk filter.
- Validation for stored queries has been improved to prevent saving queries with warnings or errors.
- A bug that could prevent custom integration results from merging into existing assets has been resolved.
- Fingerprint improvements.

3.7.3

2023-04-17

- Excerpts of task log messages are now available on the task details page for tasks that are in error status.
- The display of datagrid warning and error messages has been improved.
- A bug that could cause the save button on the credential edit form to be disabled has been resolved.
- Fingerprint improvements.

3.7.2

2023-04-13

- Improved asset processing when FortiGuard endpoints with “Policy Override Authentication” enabled are present.
- Self-hosted installs now support an option to disable TLS validation between Explorers and the console application.
- A bug where clicking links on the Query page of a self-hosted instance may return a 500 has been resolved.
- A bug where clicking links in the Tasks column of the Credentials page would result in an error has been resolved.

- A bug where paginated results could display Viewing 0 - N for the first page has been resolved to now display Viewing 1 - N.

3.7.1

2023-04-12

- Third-party integrations now support more granular vulnerability filters.
- HTTP security headers can now be disabled in self-hosted mode.
- Crowdstrike will now use Connection IP and Connection MAC for asset matching.
- The max-repetitions and disable-bulk parameters have been added to SNMP probes.
- Task failures are now reported in the Task details pane.
- All queries, including runZero-provided system queries, can now be copied.
- The configuration for runZero-provided system queries can be modified.
- A bug that could result in duplicate offline assets has been resolved.
- A bug that prevented CSV exports of assets when using free text search has been resolved.
- A bug where the number of hops could be incorrectly set to zero when ARP is present as a service has been resolved.
- A bug that prevented searching assets using the task search key has been resolved.

3.7.0

2023-04-10

- Customers with an enterprise license can now create custom integrations and import assets from any external asset data source using the runZero Python SDK.
- Improved performance and reliability of metrics calculations.
- Improved performance of the vulnerabilities inventory.
- AWS permission errors are now more detailed to make troubleshooting easier.
- A bug where the asset ownership tag was not able to be changed successfully has been resolved.
- A bug where email addresses were case sensitive on sign in has been resolved.
- A bug where the "Create Organization" button appeared disabled but was still clickable has been resolved.
- A bug preventing the Asset Ownership goals toggle from being clickable has been resolved.
- Upgraded npcap to version 1.73
- Fingerprint improvements.

3.6.19

2023-04-05

- Improved performance of the organization details page.
- A bug preventing asset owners from being updated has been resolved.
- A bug that could result in inaccurate vulnerability counts for assets has been resolved.

- A bug that could prevent a subset of vulnerabilities from being saved for multi-source assets has been resolved.
- A bug that caused errors for CrowdStrike integrations with large amounts of applications has been resolved.
- Fingerprint improvements.

3.6.18

2023-04-03

- Addresses bug where recurring tasks that are “Removed” were still showed in the tasks page after the associated site is deleted.
- Accessibility improvements.
- Client-side timezone updates.
- Fingerprint improvements.

3.6.17

2023-04-02

- API requests to apply tags to one or more assets now complete much faster.
- The scanner now supports the Steam In-Home Streaming discovery protocol.
- Attribute reports now group unique values within a single key.
- The View More link is now accessible for in-progress tasks.
- Asset owner names now suggest auto-complete options.
- A bug that could lead to the self-hosted installer not removing temporary files has been resolved.
- A bug that led to slow SNMP scans of specific Cisco switches has been resolved.
- Fingerprint improvements.

3.6.16

2023-03-29

- A bug that prevented access to runZero canned Queries has been resolved.

3.6.15

2023-03-29

Important security fix:

- A bug that could show cross-tenant Queries and their associated author email addresses was resolved. This issue only applied to the cloud-hosted version of the runZero platform. The affected build was live for slightly more than two hours. Any customers affected by this issue will receive a detailed notice to the email addresses associated with their superuser accounts.

3.6.14

2023-03-29

- A bug that could prevent updating assets with a large number of vulnerabilities has been resolved.
- Fingerprint improvements.

3.6.13

2023-03-29

- Improved reliability of the Tenable integration.
- A bug that could prevent analysis queries from running for directory users and groups has been resolved* A bug that prevented match counts from being displayed on the queries page has been resolved.
- Saved queries can now be created for software, vulnerabilities, and screenshots.
- Fingerprint improvements.

3.6.12

2023-03-28

- The maximum time to complete an SNMP walk is now configurable.
- The default maximum time to complete an SNMP walk has been increased to 5 minutes from 1 minute.
- The maximum results for an SNMP walk have been increased to 8k from 4k.
- Assets owned by a runZero user will now be displayed on the user details page.
- A bug which could result in runZero attributes being removed from Offline assets has been resolved.
- A bug that could prevent subnet stats from being exported has been resolved.
- The API response for a PUT request to /org/sites now returns the details of the new site.
- The Reason column in the failed tasks table will now properly persist the hidden state between page loads.
- Fingerprint improvements.

Note: The upgrade process may take up to an hour for large self-hosted deployments.

3.6.11

2023-03-27

- A bug that could lead to inaccurate asset correlation has been resolved.

3.6.10

2023-03-27

- A bug that could prevent validation of hostname scan targets has been resolved.
- Cisco virtual MAC addresses are now handled more consistently.
- Increased timeouts for the Tenable integration.
- Improved reliability of CrowdStrike credentials verification.
- Datagrids across the UI no longer use the incorrect theme.
- Task WLAN listing functionality has been improved to enforce a timeout if the underlying utility is slow or unresponsive.
- Fingerprint improvements.

3.6.9

2023-03-26

- Fingerprint improvements.

3.6.8

2023-03-23

- Vulnerability inventory now includes an Exploit status, indicating whether the vulnerability is known exploitable. The Exploit status will only be populated for vulnerabilities imported after this release.
- A bug that could prevent the dashboard from loading successfully has been resolved.
- A bug that caused misaligned values when exporting assets to CSV has been resolved.
- A bug that could cause assets to incorrectly merge has been resolved.
- Organizations with Parents set prior to 3.6 have had their Parents reset.
- TLS stack fingerprinting has been improved.
- Fingerprint improvements.

3.6.7

2023-03-22

- A bug that could prevent asset modifications triggered by alert rules has been resolved.

3.6.6

2023-03-22

- Alert rules now support software and vulnerability queries.
- Asset ownership now supports references to runZero users and groups.

3.6.5

2023-03-22

Important security fix:

- A bug that could allow an organization admin to see the names of other organizations in the tenant, even without explicit access, has been resolved.

In addition to the security improvement above, this release includes:

- Asset queries can now surface overlaps in asset names, IP addresses, and MAC addresses across inventory.
- Behavior around parent/child organizations has been improved.
- A change to Chrome which caused web screenshots to fail has been addressed.
- Fingerprint improvements.

3.6.4

2023-03-21

Important security fix:

- A bug that could expose limited information about an organization to cross-tenant users has been resolved. This issue could have allowed an attacker that guessed the v4 UUID of an organization to view the name, description, and top-level statistics (asset count, service count, task count, etc.) without appropriate authorization. This issue was first present in version 3.6.0.

3.6.3

2023-03-20

- Attribute searches and reports are now faster in large organizations.
- It is now possible to download the task log for a failed scan.
- Hosted scans no longer ignore responses from common firewalls.
- Daily asset expiration now records an assets-expired event with the count.
- The task-failed event now includes information about the associated Explorer.
- Scans can now configure specific probes for Subnet and Host pings.
- Validation warnings for internal IPs when using LDAP and InsightVM integrations has been improved.
- Filtering of non-unique MAC addresses has been improved to better support Cisco virtual MAC addresses.
- Fingerprint improvements.

3.6.2

2023-03-16

- A bug that caused misaligned values when exporting assets to CSV has been resolved.
- A bug that could cause the SSO page to render off screen has been resolved.
- Enterprise customers can now scan all ports and up to a /8 at a time using the hosted scan engines.
- The AWS integration now supports the GovCloud partition for assumed roles.
- Fingerprint improvements.

3.6.1

2023-03-15

- Improved quality of errors reported by the CLI Scanner.
- Improved user experience of user management.
- Improved user experience of organization management.
- A race condition that could occur during self-hosted installation has been resolved.
- A bug that could prevent some CrowdStrike software from importing successfully has been resolved.
- Packets sent/received are now visible from the tasks preview.
- A bug that could cause the Tenable connector to fail intermittently for some customers has been resolved.
- A bug that could cause task details not to render on the task overview screen has been resolved.
- A bug that could prevent organization administrators from creating new projects has been resolved.
- Fingerprint improvements.

3.6.0

2023-03-13

This release is a roll-up of the 3.5.x updates in addition to the following changes.

- Organizational hierarchies are available allowing for permissions to be inherited by child organizations based on an established parent.
- Generation of operating system CPEs has been significantly improved.
- Fingerprinting of operating systems imported from the Active Directory and VMware integrations have been significantly improved.
- Asset rescanning now provides the option of including/excluding extra addresses in the scan.
- License warnings can now be hidden by user role.
- Credentials no longer require thumbprints to be provided.
- Asset ownership detection for the Active Directory integration has been improved.
- A bug that could cause paused tasks to fail when unpaused has been resolved.

- A bug that could cause large Qualys imports to fail has been resolved.
- A bug that could cause incorrect credential values to show on the credential edit screen has been resolved.
- Fingerprint improvements.

3.5.10

2023-03-10

- A bug that could cause a browser crash in the latest release of Chromium-based browsers on MacOS has been circumvented.

3.5.9

2023-03-08

- A bug that could cause CrowdStrike tasks to fail when missing software permissions has been resolved.
- A bug that could prevent bogus services from certain firewalls from being completely filtered has been resolved.
- Accessibility improvements.
- Fingerprint improvements.

3.5.8

2023-03-07

- Services, Screenshots, and Software inventory pages now include associated site subnet tags.
- IP addresses reported by CrowdStrike are now considered primary addresses, and will be used for asset correlation.
- CrowdStrike credential verification is now separated by service.
- A bug that could prevent the creation of CrowdStrike assets with software entries containing Unicode escape sequences has been resolved.
- A bug which prevented Chrome on Android users from navigating and submitting the activation form has been resolved.

3.5.7

2023-03-06

- The CrowdStrike connector can now import software data through Falcon Discover.
- Improved error validation for email addresses when setting up an email alert channel.
- A bug where firewalls and similar devices responding to many non-asset IP addresses during scanning would lead to unexpected assets in the inventory has been resolved.
- A bug preventing the active scans dashboard widget from navigating to the associated task has been resolved.

- A bug preventing site subnet tags from appearing in the dashboard Asset tags widget has been resolved.
- Fingerprint improvements.

3.5.6

2023-03-01

- A bug that could prevent the VMWare connector task page from loading has been resolved.
- A bug that could cause duplicate MSDefender attributes on an asset has been resolved.
- Fingerprint improvements.

3.5.5

2023-02-28

- Stability and performance of VMware asset correlation has been improved.
- VMware assets are now merged across sites.
- SNMP protocol versions are now tracked at the asset level.
- SNMP services will now keep track of how they authenticated and which protocols were used.
- The Microsoft Intune integration has been improved to better handle Intune API rate limiting.
- Hostname extraction from malformed subjectAlternativeNames on TLS certificates has been improved.
- Site scopes with subnets ending in /32 (for IPv4) and /128 (for IPv6) are no longer parsed to single IPs and will appear as CIDR entries in the subnets list.
- A bug that could prevent the creation of VMware assets has been resolved.
- Accessibility improvements.
- Fingerprint improvements.

3.5.4

2023-02-22

- A bug that could cause inaccurate asset counts in the Organization Overview report has been resolved.
- A bug that could cause site imports to fail when missing optional fields has been resolved.
- Fingerprint improvements.

3.5.3

2023-02-21

- A new canned query for OpenSSH 9.1 servers which [contain a memory double-free vulnerability](#) has been added.
- Performance of the Microsoft Active Directory (LDAP), Azure AD, and Google Workspace integrations has been improved.
- IP addresses reported by Tenable are now considered primary IP addresses, consistent with the Rapid7 and Qualys integrations.
- A bug causing the dashboard asset trends graph tooltips to appear away from the graph has been resolved.
- A bug causing task page inspection cards to automatically collapse has been resolved.
- A bug that could result in a buildup of frequently recurring tasks has been resolved.
- A bug that could cause extremely large tasks to remain queued for processing indefinitely has been resolved.
- A bug that could prevent export of service attribute reports has been resolved.
- A bug preventing license requirement indicators from being visible on some pages has been resolved.
- A bug preventing saving credentials due to bad org-access settings has been resolved.
- A bug preventing recalculation of the next scheduled run time for a scan has been resolved.
- Fingerprint improvements.

3.5.2

2023-02-14

- The scanner now supports identifying RDP authentication methods, including legacy and NLA, supported by target hosts.
- The scanner now supports the ability to decode ISAKMP/IKEv2 replies.
- A bug that prevented OS fingerprinting and information extraction over RDP has been resolved.
- A bug preventing users from copying or editing connector and analysis tasks has been resolved.
- A bug causing new recurring tasks to display an incorrect first run date has been resolved.
- Several minor bug fixes and UX improvements have been made to the redesigned task page.
- Several minor bug fixes have been made to the redesigned task library page.

3.5.1

2023-02-13

- A bug that could prevent automatic metric calculations from completing has been resolved.
- A bug that could prevent stale assets from being automatically removed on subsequent task runs has been resolved.

- Several minor bug fixes and UX improvements have been made to the redesigned task page.

3.5.0

2023-02-13

This release is a roll-up of the 3.4.x updates in addition to the following changes.

- The Asset Ownership feature allows you to manage asset owners across your asset inventory.
- The task page has been redesigned for improved user experience.
- A new canned query for VMware ESXi servers that could be [targets of the ongoing ESXiArgs ransomware campaign](#) has been added.
- The scanner now parses running processes and services from checkmk.
- Stability and performance of third-party asset correlation has been improved.
- Fingerprint improvements.

3.4.23

2023-02-07

- A new canned query for Lexmark printers which may be [vulnerable to CVE-2023-23560](#) has been added.
- The Qualys integration has been improved to better handle Qualys API rate limiting.
- Dashboard performance has been improved.
- The Service attributes report performance has been improved.
- A bug that could cause early removal of Tenable assets has been resolved.
- A bug that could cause attribute reports to export incorrect data has been resolved.
- A bug that may prevent the mDNS probe from completing has been resolved.
- The scanner now supports the CoAP protocol over UDP.
- The scanner interface selection logic has been improved.
- Fingerprint improvements.

3.4.22

2023-02-01

- The scanner now supports the Minecraft Bedrock protocol.
- Public API endpoints to export directory users and groups have been added.
- The Last checkin column on the Registered Explorers table has been renamed to Online status.
- Task details pages now include a Created by column.
- A bug that would allow you to save an invalid URL in credential forms has been resolved.
- A bug that could prevent creation of new self-hosted clients has been resolved.
- A bug that could result in a 500 error on the dashboard when selecting a site with no existing metrics has been resolved.

- Fingerprint improvements.

3.4.21

2023-01-27

- The consistency of matching based on asset attributes has been improved.
- The ability to capture Telnet banners has been improved.
- A bug that could prevent updating payment information has been resolved.
- A bug that could prevent searching for tasks associated with deleted sites has been resolved.
- A bug that reported an incorrect match .probe for Qualys assets has been resolved.
- Fingerprint improvements.

3.4.20

2023-01-25

- Dashboard performance has been improved.
- The Qualys integration has been improved to better handle Qualys API rate limiting.
- A new canned query for surfacing cloud compute assets with GPU hardware has been added.
- The scanner now supports the NDMP protocol.
- A bug that could display an incorrect value for Explorer architecture in the console has been resolved.
- Fingerprint improvements.

3.4.19

2023-01-18

- A bug that could cause the asset CSV export to fail for some users has been resolved.

3.4.18

2023-01-18

- [runZero Preview Program](#): the Asset Ownership feature lets you manage asset owners across your asset inventory.
- The scanner now supports the Munin protocol.
- Fingerprint improvements.

3.4.17

2023-01-17

- A bug that could intermittently prevent scan completion when scanning IPv6 endpoints with the mDNS probe enabled has been resolved.

3.4.16

2023-01-11

- Queries now include an option to limit results to only live assets.
- A bug that prevented highlighting outlier attributes on the asset details page has been resolved.
- Fingerprint improvements.

3.4.15

2023-01-10

- The InsightVM connector now supports longer timeouts for large sites and slower consoles.
- The CrowdStrike connector now avoids asset duplication when processing large datasets.
- The CrowdStrike connector now handles larger datasets with lower resource usage.
- Asset matching based on attributes is more consistent across lossy networks.
- Excluded scan targets are now no longer matched or marked as offline.
- Site imports now automatically trim trailing whitespace from CIDRs.
- The scanner now supports the MySQL X protocol.
- Fingerprint improvements.

3.4.14

2023-01-09

- A bug that could lead to duplicate assets from CrowdStrike has been resolved.

3.4.13

2023-01-07

- A bug that could prevent very large CrowdStrike syncs from completing has been resolved.

3.4.12

2023-01-05

- A bug that could prevent the creation of new offline assets has been resolved.
- A bug that could cause large InsightVM imports to fail has been resolved.
- A bug that could cause paused tasks to fail when unpaused has been resolved.

3.4.11

2023-01-05

- Credentials can now be edited after being saved.
- Credentials can now be verified against services to provide visibility into upstream configuration issues.
- The Subnet utilization report is now more performant, includes subnets by site, and includes assets outside the scope of a defined site subnet.
- Asset inventory searches with the mac keyword has been improved to support the Cisco MAC address format and additional delimiter characters.
- Filtering of non-unique MAC addresses has been improved, including improvements for Fortinet virtual adapters, Juniper switches, and Project Calico virtual interfaces.
- A bug that could cause CrowdStrike vulnerability imports to fail due to session expiration has been resolved.
- A bug that could cause assets to be duplicated has been resolved.
- A bug that could cause RDNS results to mark assets as online has been resolved.
- Added support for L2TP, Dahua DHIP, KXNnet, Webmin, and Playstation UDP protocols.
- Fingerprint improvements.

3.4.10

2022-12-23

- Fingerprint improvements.

3.4.9

2022-12-22

- A bug that could cause the CrowdStrike connector to fail when missing Spotlight permissions has been resolved.
- Fingerprint improvements.

3.4.8

2022-12-21

- A bug that could prevent manually merging some assets has been resolved.
- A bug that could cause incorrect project expiration information has been resolved.

- Improved fingerprinting of Huawei, Hikvision, Fortinet, and WatchGuard devices over SNMP.
- Fingerprint improvements.

3.4.7

2022-12-19

- Service exports now include a `service_id` field.
- A bug that could cause an excess of temporary files in self-hosted installations has been resolved.
- A bug that could prevent asset attribute imports from reporting all results has been resolved.
- Fingerprint improvements.

3.4.6

2022-12-19

- The dashboard now loads faster for customers with many sites and subnets.

3.4.5

2022-12-18

- New scans and recurring scans can now be saved even when the current license has been exceeded.
- Confirmation dialogs for removal actions are now more consistent across the product interface.
- Modifications to recurring scans with a past start date will no longer immediately launch a task.
- A bug that prevented the latest Windows Explorers from embedding npcap has been resolved.
- Fingerprint improvements.

3.4.4

2022-12-16

- Self-hosted installations of runZero now support inclusion of custom JavaScript in UI web pages.
- Amazon Web Service account IDs are now visible in a new per-asset attribute.
- Support for importing Nessus files without vulnerability details has been improved.
- The organization and client switching dropdown menus can now be filtered if there are more than 5 organizations or clients.
- A bug causing the scan setup page to scroll to the right during the product tour has been resolved.

- A bug where integrations using the Microsoft Graph API may have their token expire between paged responses has been resolved.
- Improved mDNS service discovery.
- Fingerprint improvements.

3.4.3

2022-12-14

- A bug that could lead to asset duplication when importing from Tenable, InsightVM, and Qualys has been resolved.
- Fingerprint improvements.

3.4.2

2022-12-13

- Registered API clients now show the user that created them.
- Microsoft 365 Defender tasks now report details on failed upstream API calls.
- The help text on Google Workspace credentials has been improved.
- The formatting of dashboard category reports has been improved.
- A bug that could prevent successful import from Shodan has been resolved.
- A bug that prevented InsightVM connector options from persisting correctly has been resolved.
- Fingerprint improvements.

3.4.1

2022-12-12

- A bug that could result in incorrect processing of wireless entries has been resolved.
- Fingerprint improvements.

3.4.0

2022-12-12

This release is a roll-up of the 3.3.x updates in addition to the following changes.

- A new canned query for Cisco 7800/8800 series IP phones which may be [vulnerable to CVE-2022-20968](#) has been added.
- The AWS integration now includes an option to automatically remove assets no longer reported by AWS.
- OAuth 2.0 client credentials can now be used to authenticate with runZero APIs.
- The `edr.name` asset attribute is now updated to show when a runZero scan no longer detects the EDR agent.
- Tasks can now be stopped during data gathering and processing phases.

- The site import and export CSV format has been simplified.
- The performance of connector task processing has been improved.
- The performance of tables in the Site comparison report, analysis report results, and SSO group mappings has been improved.
- The fingerprinting coverage of Google Workspace assets has been improved.
- Additional Fingerprint improvements.

3.3.8

2022-12-07

Important security fix:

- A bug that could show cross-tenant “no access” role users in the Your team > Current organization view was resolved. This issue only applied to the cloud-hosted version of the runZero platform. The affected build was live for slightly more than two hours. Any customers affected by this issue will receive a detailed notice to the email addresses associated with their superuser accounts.

3.3.7

2022-12-07

- A bug that could prevent an Explorer from running scans with specific network configurations has been resolved.

3.3.6

2022-12-07

- The CrowdStrike integration now imports vulnerabilities when CrowdStrike Spotlight is enabled for the API key.
- An option to disable the creation of new assets from third-party integrations has been added.
- The performance of the task overview page load time has been improved.
- The consistency in asset terminology has been improved.
- The site import CSV format has been improved.
- Third-party integrations merge assets more consistently.
- The CLI Scanner --api-url parameter handling has been improved.
- The DELETE API method for bulk asset deletion has been deprecated.
- A public API endpoint to check the platform health has been added.
- OS EOL dates are now reported for Windows 11.
- Fingerprinting of HomeKit devices has been improved.
- A new canned query for MegaRAC BMC firmware has been added.
- A bug that could cause recurring tasks to backup has been resolved.
- A bug in the Organization asset export API has been resolved.
- Fingerprint improvements.

3.3.5

2022-11-30

- The import time for third-party data sources was improved.
- A bug that caused the License information page to display an incorrect project asset count was resolved.

3.3.4

2022-11-28

- A bug that could delay concurrent task processing has been resolved.

3.3.3

2022-11-28

- An issue that could cause the command-line scanner to skip LDAP enumeration has been resolved with the `--ldap-thumbprints` flag.
- The scheduler will now delay recurring tasks if the previously completed task has not yet started processing.
- The backend now processes concurrent tasks for separate sites within the same organization when possible.
- Self-hosted customers can configure concurrent task processing with the `RUNZERO_CRUNCHER_INSTANCES` option.
- Third-party integrations now merge more accurately when using IP addresses as the match key.
- Microsoft Intune and Azure Active Directory assets are now fingerprinted more accurately.
- VMware ESXi instances now display OS end-of-life dates based on version.
- Fingerprint improvements.

3.3.2

2022-11-21

- A bug that could prevent tag searches from completing when thousands of tags are in use has been resolved.
- The scanner now supports a configurable ToS/Traffic Class field in the advanced configuration.
- Additional operating system and hardware icons are available in the inventory view.
- Explorer and CLI Scanner binaries are now approximately 5MB smaller.
- New LDAP credentials now auto-populate the discovered port.
- Printer detection has been improved.
- Fingerprint improvements.

3.3.1

2022-11-20

- The Microsoft Defender integration now merges assets more comprehensively.
- The AWS EC2 integration now provides an option to include Stopped instances.
- A bug that could result in partial import of GCP CloudSQL assets was resolved.
- The “All Organizations” view now more accurately handles limited user permissions.
- Searching and sorting is faster when using the asset first seen and last seen columns.
- A bug that could lead to duplicate vulnerabilities when an import was restarted has been resolved.
- Fingerprint improvements.

3.3.0

2022-11-14

This release is a roll-up of the 3.2.x updates in addition to the following changes.

- runZero Professional and Enterprise customers can now sync assets from Google Workspace.
- runZero Platform customers can now sync users and groups from Google Workspace.
- User interface tables were revamped for Organizations, Sites, Explorers, and Teams.
- The “All Organizations” view is now available to restricted users with a filtered scope.
- Live validation is no longer required for Qualys VMDR and InsightVM credentials.
- Fingerprint improvements.

3.2.11

2022-11-08

- The subnet utilization report now supports filtering by site.
- CSV export of assets now includes the same hostname information as the inventory view.
- Up-to-date ARM64 builds of the standalone scanner are now available.
- The account API endpoint for creating organizations now accepts the argument types documented.
- Merging two assets now correctly updates the date of the newest MAC address for the resulting asset.
- Disabling all scan probes now disables the SNMP probe.
- A bug that could prevent the use of third-party credentials when using TLS thumbprints or the insecure connection option with a public URL has been resolved.
- Fingerprint improvements.

3.2.10

2022-11-04

- Service Provider information is now displayed with a default domain before SSO settings are configured.
- Improved performance when scanning from macOS hosts that have certain EDR solutions installed.
- A bug which sometimes prevented GCP imports from completing has been fixed.
- Improved TLS fingerprinting.
- Fingerprint improvements.

3.2.9

2022-11-03

- The AWS integration now includes an option to delete AWS-only assets that were not seen in the most recent import.
- The Qualys integration now includes an option to import unscanned assets and is disabled by default.
- Processing speed for large Qualys imports has been improved.
- Explorers are now ordered alphabetically on the scan configuration and connector configuration pages.
- A bug in how Service Inventory searches were launched from the Asset details page had been resolved.
- Tanium agent detection now sets the `edr.name` attribute.
- Improved fingerprinting of OpenSSL, GnuTLS, and Windows TLS stacks.

3.2.8

2022-11-01

- A bug that could prevent TLS probes from completing has been resolved.

3.2.7

2022-11-01

- A new `tls.stack` attribute that tracks the TLS software provider and version has been added for assets and services.
- A new canned query for OpenSSL 3.0.x with client certificate authentication has been added.
- Improved performance of Intune integration when importing a large number of users and devices.
- runZero users logging in via SSO are now presented with the terms and conditions acceptance dialogue.

- A bug that could prevent updating site metrics has been resolved.
- Improved fingerprinting of OpenSSL versions.
- Apple ecosystem OS Fingerprint improvements.

3.2.6

2022-10-30

- The scanner now reports OpenSSL versions via TLS fingerprinting.
- The scanner now reports Tanium agent instances on the network.
- The scanner now reports additional detail for SSLv3 services.
- A bug that could prevent the Intune integration from completing long-running tasks has been resolved.
- A bug that could prevent the GCP integration from returning all assets has been resolved.
- A bug that could result in a recurring integration running again before the previous task finished has been resolved.
- Fingerprint improvements.

3.2.5

2022-10-26

- GCP credentials can now be configured to import assets from multiple projects.
- A bug that could prevent importing assets from Microsoft Intune has been resolved.
- A bug that could prevent importing assets from Microsoft 365 Defender has been resolved.

3.2.4

2022-10-24

- Scan task processing speed has been improved for SaaS and self-hosted customers.
- The baseline memory usage of Explorers has been reduced.
- A bug that could prevent importing assets from Microsoft 365 Defender has been resolved.
- A bug that could cause broken asset links has been resolved.
- Fingerprint improvements.

3.2.3

2022-10-20

- Error handling of misconfigured fingerprints has been improved to reduce Explorer and scanner crashes.

3.2.2

2022-10-20

- The search keywords `has_os_eol` and `has_os_eol_extended` are now supported on the Assets and Vulnerabilities inventory pages.
- A bug that could cause missing service data for services with conflicting virtual hosts has been resolved.
- A bug that could cause inaccurate user counts for imported directory groups has been resolved.
- A bug that affected tooltip display has been resolved.
- A bug that prevented “open in new tab” navigation using middle/right click has been resolved.
- Fingerprint improvements.

3.2.1

2022-10-18

- The error message indicating that an AWS integration credential has insufficient permissions has been improved.
- The “last seen” link to the most recent scan details has been restored on the asset details page.
- A bug that could prevent Azure AD imports has been resolved.

3.2.0

2022-10-17

This release is a roll-up of the 3.1.x updates in addition to the following changes.

Important security fixes:

- Three stored cross-site scripting vulnerabilities were identified and fixed as part of our annual third-party security assessment.

In addition to the security improvement above, this release includes:

- runZero Platform customers can now sync assets from Microsoft 365 Defender.
- runZero Platform customers can now sync assets from Microsoft Intune.
- The Azure AD integration now imports additional assets and no longer requires a Microsoft Intune license.
- The Azure AD integration can now be configured to optionally import assets, users, and groups.
- The Active Directory integration service options have been adjusted for consistency.
- Directory users and groups can now be included in custom queries.

- The Organization Overview report now contains summary information for directory users and groups when present.
- SNMPv2 options have been moved to the Probes tab (now labeled Probes and SNMP).
- The toggle switch to use or not use SNMP now correctly reflects whether it is overridden by the “Use defaults” option on the Probes tab.
- The asset details pages have been redesigned for improved performance.
- The asset details pages now include a “last loaded” time indicator and the ability to refresh the page data.
- Recent users from Microsoft Intune, SentinelOne, and CrowdStrike are now included on the asset details page.
- Alert notifications, user invitations, and password reset emails are now sent from the runzero.com domain name instead of rumble.run.
- The rumblectl utility now has a diagnostics command to run or save a diagnostic script for self-hosted customers to collect information for runZero support.
- A bug that could prevent repeated import of task data that includes directory users and groups has been resolved.
- A bug that caused subnet sampling and screenshots to be enabled for all scan tasks has been resolved.
- Fingerprint improvements.

3.1.13

2022-10-11

- Inventory pages now offer “all” and “none” column visibility selection options.
- Fingerprint improvements for Fortinet products that may be affected by CVE-2022-40684.

3.1.12

2022-10-07

- The Tenable.io integration now supports a configurable API URL.
- The Active Directory integration now supports optional import of assets, users, and groups.
- The minimum TLS version supported by new Active Directory credentials has been increased from TLS 1.0 to TLS 1.2, with a configurable option to support older TLS versions.
- A bug that could prevent modifying the maximum concurrent scans setting was resolved.
- A bug that could result in an inaccurate task count on the credentials page was resolved.
- A bug that could result in inaccurate searches by credential on the tasks page was resolved.
- A bug that could result in inaccurate reporting of credential reuse was resolved.
- A bug that could cause certain browser extensions to prevent configuring scans was resolved.

3.1.11

2022-10-06

- A bug that could prevent reuse of SNMP credentials for recurring scans was resolved.

3.1.10

2022-10-03

This release contains important security fixes:

- A bug that could lead to stored cross-site scripting in the scan templates view was fixed. This issue could be exploited by an authenticated, but unprivileged user to take over the session of another authenticated user.
- A bug that could lead to stored cross-site scripting in the SSO group mappings view was fixed. This issue could be exploited by an authenticated superuser to take over the session of another authenticated user.

3.1.9

2022-09-30

- A bug that could prevent initializing a scan in some cases was resolved.

3.1.8

2022-09-30

This release contains an important security fix:

- A bug that could lead to stored cross-site scripting in the team view was fixed. This issue could be exploited by an authenticated, but unprivileged user to take over the session of another authenticated user.

In addition to the security improvement above, this release included a separate bug fix:

- A bug that prevented recurring scans from being saved in some cases was resolved.

3.1.7

2022-09-29

- The search keyword `os_eol_expired` is now supported on the Assets inventory.
- The handling of Qualys concurrency and rate limiting has been improved.

- The `rumblectl` command can now be used with self-hosted deployments to configure additional superusers.
- A bug that prevented the `first_seen` timestamp from being set has been fixed.
- A bug that could cause large Qualys imports to fail has been resolved.
- A bug that prevented import of Azure AD users and groups when missing an active Intune license has been resolved.
- A bug that could result in partial import of Azure AD users and groups has been resolved.
- A bug which prevented the `report.changed` value from working in notification rule templates has been fixed.
- A bug that prevented the use of client tokens to authenticate to the API has been fixed.
- Fingerprint improvements.

3.1.6

2022-09-27

- Email notifications are now enabled for non-recurring Organization Overview reports.
- Relative time searches now accept negative numbers.
- Scan tasks and templates now allow empty SNMPv1 and SNMPv2 community strings.
- Credential validation has been improved to prevent common misconfigurations.
- Support for Explorer hosts running virtual machines has been improved.
- MAC vendor display behavior on inventory datagrids has been improved.
- Tooltips on datatable icons have been improved.
- A bug that could cause insight queries for hosted zones to fail has been resolved.
- A bug in the Shodan integration asset-mode query has been resolved.
- A bug that could cause MAC vendor names to be cut off in datagrids has been resolved.

3.1.5

2022-09-22

- The task change report schema has been updated to support changes to directory users and groups.
- Error messages related to API tokens have been improved.
- A bug that could result in missing Shodan services has been resolved.
- A bug that incorrectly imported Active Directory Managed Service accounts as assets has been resolved.
- Fingerprint improvements.

3.1.4

2022-09-19

- Improved performance of asset exports with many subnets.
- Asset exports now filter subnet results to those containing the assets' addresses.
- Improved LDAP connector and probe logging.

- Added group_count keyword to Users search.
- Improved grouping of inputs in connector forms.
- A bug that could cause the Switch Topology report to not show all switches in certain situations has been resolved.
- A bug that could result in a 500 error when exporting assets from sites with many assets and/or subnets has been resolved.
- A bug that could result in UI elements becoming unresponsive has been resolved.
- Fingerprint improvements.

3.1.3

2022-09-14

- A bug that could prevent some service values from being saved has been resolved.

3.1.2

2022-09-14

- Improved loading times of the directory groups inventory page.
- A bug that could result in all subnet tags being applied to exported assets has been resolved.
- A bug that could result in missing Shodan services has been resolved.
- A bug that could cause Azure AD imports to fail for certain configurations has been resolved.
- A bug that could cause excessive export sizes has been resolved.
- A bug that could obscure task errors from the task log has been resolved.
- Fingerprint improvements.

3.1.1

2022-09-13

- Improved loading times of the inventory screens, including multi-page selection.
- Search keyword has_group is now supported on the Users page.
- Fingerprint improvements.

3.1.0

2022-09-12

This release is a roll-up of the 3.0.x updates in addition to the following changes.

- runZero Platform customers can now sync assets from Shodan.
- runZero Platform customers can now sync assets from Azure Active Directory.
- runZero Platform customers can now sync assets from Microsoft Active Directory via LDAP.

- Connector tasks now can optionally be run from an Explorer on a customer's network.
- The Events datatable has been redesigned and is now more performant.
- The Qualys integration now provides a more descriptive error message when rate-limited by the Qualys API.
- Network File System (NFS) protocol detection on TCP ports has been improved.
- A bug that prevented editing certain probe options when configuring a scan has been resolved.
- Fingerprint improvements.

3.0.24

2022-09-10

- A bug that could cause the browser to freeze when viewing assets with many attributes has been resolved.
- Fingerprint improvements.

3.0.23

2022-09-08

- Web screenshots are now limited to a maximum of 16 concurrent processes.
- Web screenshots will now run concurrently on arm64 macOS systems.
- Improved error handling for the GCP integration.
- Improved parsing of input hostnames.
- A bug that could prevent rendering dashboard insights has been resolved.

3.0.21

2022-09-07

- Dashboard insights have been limited to a maximum of three rows.
- Fingerprint improvements.

3.0.20

2022-09-07

- A bug that could result in minimal assets being skipped has been resolved.

3.0.19

2022-09-07

- A bug that could result in the wrong insight counts on the dashboard has been resolved.
- Fingerprint improvements.

3.0.18

2022-09-06

- A bug that could cause attributes and screenshots to be removed from offline assets has been resolved.
- Fingerprint improvements.

3.0.17

2022-09-02

- Fingerprint improvements.

3.0.16

2022-09-02

- Fingerprint improvements.

3.0.15

2022-09-01

- A bug that prevented using certain organization and export tokens has been resolved.
- Fingerprint improvements.

3.0.14

2022-08-31

- A new optional filter has been added to the CrowdStrike connector.
- The performance of the Qualys connector has been improved.
- Event details have been added to alert templates by default.
- A bug that caused the token to be missing from password reset emails has been resolved.
- Fingerprint improvements.

3.0.12

2022-08-29

- Task statistics for asset counts are now included in CSV exports and can be used in task searches.

- The `license-limit-exceeded` event has been added to alert when the live asset count exceeds an accounts license.
- The `ldap.notes` service attribute has been added for providing user-friendly representations of well-known discovered LDAP OIDs.
- A bug that could cause query timeouts has been resolved.
- A bug that could cause large Qualys imports to timeout has been resolved.
- Fingerprint improvements.

3.0.11

2022-08-24

- A bug that prevented Qualys from being fully imported from large sites has been resolved.
- A bug that led to slow exports and job processing has been resolved.
- Fingerprint improvements.

3.0.10

2022-08-22

- The Tenable integration now excludes terminated and deleted assets.
- A bug that affected formatting of `_asset.match` values has been resolved.
- A bug that caused internal tasks for metrics calculation to generate scan-completed events has been resolved.
- A bug that prevented reports for specific asset attributes has been resolved.

3.0.9

2022-08-19

- Dashboard metrics now account for unscanned assets imported from third-party integrations.
- Internal recurring tasks for metrics calculation no longer show in the recurring task count.
- Fingerprint improvements.

3.0.8

2022-08-19

- A bug that could prevent exporting asset attributes has been resolved.
- Fingerprint improvements.

3.0.7

2022-08-17

- A bug that could prevent CrowdStrike tasks from processing has been resolved.

3.0.6

2022-08-17

- Processing performance for foreign asset data has been improved.
- A bug that could prevent the generation of some asset attribute reports has been resolved.
- Fingerprint improvements, including AIX OS and vCenter, Avaya, and Proofpoint appliances.

3.0.5

2022-08-12

- A bug that could cause offline self-hosted platform updates to fail has been resolved.
- The timeout for Qualys connection tasks has been increased from 60 seconds to 5 minutes.
- Fingerprint improvements.

3.0.4

2022-08-11

- A notice was added to the MFA page to inform users that they can continue to use the old `rumble.run` domain until they re-enroll their authenticators for the new `runzero.com` domain.
- Font rendering in Safari browsers now matches Firefox and Chrome.
- UI improvements were made to the queries table.
- A bug that could prevent exporting selected assets and asset search results has been resolved.
- A bug that could prevent starter accounts from setting up recurring tasks has been resolved.
- A bug affecting organization selection when a default organization is set has been resolved.
- A bug that could cause SSH probes to occasionally deadlock has been resolved.

3.0.3

2022-08-09

- Fingerprint improvements.

3.0.2

2022-08-09

- A bug that prevented WebAuthn from registering correctly on console.runzero.com has been resolved.
- A bug that could cause the topology in the asset details page to be mangled has been resolved.

3.0.1

2022-08-08

- Inventory searches now support “runZero” as an asset source type.
- A bug that could affect the default probes selector functionality has been resolved.

3.0.0

2022-08-08

This release is a roll-up of the 2.15.x updates in addition to the following changes.

- Rumble is now runZero and the product UX has been updated to match.
- runZero Platform customers can now sync asset and vulnerability data from Qualys VMDR.
- The Queries datatable has been redesigned and is now more performant.
- The Software and Vulnerabilities datatables now have a “view more details” button.
- Users can now specify a Default Organization in the profile settings page.
- Outlier calculations and insight queries now automatically run as daily analysis tasks.
- Outlier calculations and insight queries can be regenerated on demand using the Metrics menu on the Tasks overview page.
- Merging assets with foreign attributes from the same source now retains all sets of foreign attributes.
- Software entries imported from SentinelOne and Tenable now report their service addresses.
- A custom query to find DrayTek Vigor routers has been added.
- The Asset and Service attributes reports can now be filtered by Site.
- The Organization API now supports asset merging.
- The services view is now up to 40% faster for organizations with large numbers of assets.
- A bug in the AWS Configuration UI causing the “Lambda instances” option to not persist has been resolved.

- A bug that could prevent external users from being directed to their main SSO login page has been resolved.
- Operating system fingerprinting has been improved for cloud assets (AWS, Azure, and GCP).
- Fingerprint improvements.

2.15.11

2022-08-01

- A bug which could cause stale software entries to be retained has been fixed.
- A bug in the Insights table which could render very large buttons has been fixed.
- Improved hostname-based merging for Rapid7 imports.
- Improved fingerprints for some FortiNet, FrontRow, and Synology assets.

2.15.10

2022-07-29

- API keys are now shown hidden by default and can be copied to the clipboard through a click.
- The Route Pathing report is now more performant and aborts early in out-of-memory scenarios.
- A bug that could lead to a 500 error when accessing the users endpoint of the organization API has been resolved.
- A bug that could cause tooltips to persist on the screen has been resolved.
- Printer detection has been improved.
- Fingerprint improvements.

2.15.9

2022-07-28

- A bug that could cause the vulnerabilities table to appear empty when sorted by the details column has been fixed.
- The scanner now fingerprints and reports a much wider range of ePO/McAfee Agent services.
- Fingerprint improvements.

2.15.8

2022-07-26

- A bug that could cause the HTTP probe to abort early has been resolved.

2.15.7

2022-07-26

- Improved support for processing very large Rapid7 imports.
- Software will now be populated from Rapid7 imports.
- OS fingerprinting will now use Rapid7 fingerprints, when Rapid7 is the only data source.
- Rapid7 foreign attributes have been adjusted for clarity.
- Services will now be populated from Censys.
- Vulnerability details are now available on the Vulnerability Inventory screen.
- Improved ability to extract Microsoft Windows information from web services.
- Improved ability to extract information from NetBIOS, including new detection of Domain Controller roles.
- Hosted Zone scan limits have been increased.
- The runZero Explorer now logs configuration file loading and reports any syntax errors.
- The asset tag update and bulk asset tag update APIs now work as documented.
- Fixed a bug which prevented all org admins from deleting other users.
- The User Last Activity date now shows the correct date.
- Fingerprint improvements.

2.15.6

2022-07-21

- Processing speed for large Nexpose and Tenable imports has been improved.
- Hostname identification from LDAP responses has been improved.
- Filtering of non-unique MAC addresses has been improved.
- Inconsistent SNMP data handling has been improved for certain classes of devices.
- A bug that could prevent connector tasks from running in parallel while connecting to third-party APIs has been resolved.
- A bug that prevented organization administrators from deleting other users has been resolved.
- A bug affecting inventory multi-select operations has been resolved.
- A bug preventing inventory column selection has been resolved.
- A bug that could indefinitely stall a task has been resolved.
- Fingerprint improvements.

2.15.5

2022-07-18

- Improved support for processing very large scans.
- Improved performance of the software and vulnerabilities tables.
- Fingerprint improvements.

Note: The upgrade process may take up to an hour for large self-hosted deployments.

2.15.4

2022-07-14

- The API now returns all attributes, sources, and subnets for a single asset.
- The runZero Explorer now runs as a delayed auto start process on Windows to increase reliability after reboots.
- The Organization Overview report now includes navigation links to return to the top of the report.
- A bug affecting license warning banners has been fixed.
- A bug affecting macOS Explorer upgrades on M1 systems has been fixed.
- A bug that prevented importing VMware assets has been fixed.
- Fingerprint improvements.

2.15.3

2022-07-13

- The InsightVM integration now supports larger imports.
- When a templated task fails due to an Explorer being unavailable, copying the failed task now retains the connection to the template.
- A bug in the Overview report which showed blank addresses for Unscanned assets has been resolved.
- A bug that caused scan copies to get assigned to a different site has been resolved.
- A bug that prevented OS icons from showing on inventory tables has been resolved.

2.15.2

2022-07-12

- Click-to-copy functionality has been restored for MAC addresses displayed on inventory pages.
- Asset export query errors now return HTTP 400 status code with descriptive bodies.
- A bug that prevented copying or updating Nessus connector tasks has been resolved.
- Fingerprint improvements.

2.15.1

2022-07-12

- A bug that could lead to an error in the External Asset Report when no assets were present has been resolved.
- A bug that could cause the Export API to return a 500 instead of 400 for invalid queries has been resolved.
- A bug that caused some Explorer updates to fail on Windows has been resolved.

2.15.0

2022-07-11

This release is a roll-up of the 2.14.x updates in addition to the following changes.

- Rumble Enterprise customers can now sync asset and vulnerability data from the InsightVM API and upload data from Nexpose XML Export files.
- Rumble Enterprise customers can now sync asset, software, and vulnerability data from the Nessus Professional API.
- Rumble Enterprise customers can now generate an External Asset report.
- Scan scopes can now be populated using external domains and IP addresses.
- All inventory tables have been redesigned and are now more performant when displaying a large number of assets.
- Integration tasks have been renamed from “Import” to “Connector” and can now run in parallel while connecting to third-party APIs.
- New reports for software and vulnerabilities have been added.
- The self-hosted platform’s web server HTTP timeouts can now be configured using environment variables.
- A bug that affected JSON exports of task information has been resolved.
- Fingerprint improvements.

2.14.11

2022-07-06

- A bug that delayed task processing when many vulnerability records were present has been resolved.

2.14.10

2022-07-01

- An API endpoint was added for Nessus imports.
- The scan engine now sets additional attributes for TLS certificates.
- The site edit page is now much faster for sites with large numbers of subnets.
- A bug has been fixed that prevented assets missing certain date/time fields from being imported from a Nessus file.
- If the start time of a recurring task is changed, the change now saves successfully and queues an immediate run if appropriate.
- Date range queries now work on Tenable / Nessus attributes. (Note that this requires a re-import of the appropriate data.)
- A bug has been fixed that prevented using boolean search terms with the Vulnerabilities table.
- A bug has been fixed that prevented some customers from importing assets from SentinelOne.
- Censys Search data processed using file import is now handled correctly.

- Fingerprints were updated for KVMs, routers, IP cameras, and other general network management equipment.
- Links to query language documentation have been updated throughout the product.

2.14.9

2022-06-27

- Improved merging of AWS, Azure, and GCP assets imported from Tenable.io.
- A bug was fixed with SentinelOne credential creation.
- A bug was fixed which could result in an incorrect agent ID or template ID being written to tasks when they were updated via API.
- A bug was fixed which prevented the asset route pathing report from rendering the source name properly in the report heading

2.14.8

2022-06-24

- A bug that prevented customer address information from being validated has been resolved.
- A bug that prevented SSO settings from being updated correctly has been resolved.

2.14.7

2022-06-24

- The self-hosted platform install now supports Rocky Linux.
- A bug in the scan engine that could lead to unexpected printer output has been resolved.
- A bug in the e-commerce checkout with non-US addresses has been resolved.
- Fingerprints were updated for Nokia SR OS and Cisco RV routers.

2.14.6

2022-06-23

- The Account API now supports scan template management.
- The scan engine discovers additional services, including Elasticsearch, Logstash, and Prometheus.
- Fingerprints were updated for various operating systems and for products by Aruba Networks, Axis, MikroTik, and Nokia.
- Query "address" keywords now support CIDR notation.

2.14.5

2022-06-16

- The scan engine now sets additional attributes for TLS certificates.
- SSO configuration now skips IdP-provided encryption certificates during setup.
- Explorer initialization speed has been optimized.
- A bug that could cause the Explorer service to timeout on startup has been resolved.
- Explorers deployed on macOS can now be transferred between organizations successfully.
- A bug has been fixed that could occasionally cause site deletion to trigger creation of an unnecessary site.
- Last activity dates of users are now shown in the same time zone as the account creation date.
- Fingerprint improvements.

2.14.4

2022-06-15

- The Tenable integration now includes an option to import unscanned assets and is disabled by default.
- The logic for merging Tenable/Nessus assets into the Rumble inventory has been improved.
- The accuracy of `first_seen` and `last_seen` dates has been improved for Nessus imports.
- The Tenable integration now provides more informative error messages for task failures due to invalid API keys and missing permissions.
- New software and vulnerabilities reports have been added.
- A preview of software and vulnerability attributes is now displayed when hovering over the Click To Copy button on the Asset details page.
- The performance of the vulnerabilities table has been improved.
- The scan engine now flags TLS services with self-signed certificates and untrusted CAs.
- CSV exports of asset information now include OS EOL and extended OS EOL dates.
- Credentials limited to specific organizations can no longer be created by an administrator without access to the organizations.
- Improved fingerprinting of Tenable/Nessus assets imported with a severity setting of low or higher.
- Fingerprint improvements.

2.14.3

2022-06-10

- Rumble Enterprise customers can now search inventory by hosted zone.

- A fingerprint for the Cockpit application has been added and includes additional Linux OS fingerprinting capability.
- Fingerprints were updated for various operating systems and for products by 2N, Grandstream, Huawei, and Wago.
- A bug that caused software and vulnerabilities to be ignored when manually merging assets has been fixed.
- A bug that could prevent task configuration from showing the list of available Explorers has been fixed.
- A bug that could prevent using the console to contact support has been fixed.
- A bug that prevented filtering some Fortinet Web Filter replies has been fixed.
- Additional Fingerprint improvements.

2.14.2

2022-06-07

- Additional validation for scan targets has been added to hosted scans.
- Additional validation for Tenable.io credentials has been added.
- Asset inventory column selection and ordering now persist between queries.
- Connector tasks now obey the start time set when they are created.
- A bug that could cause hosted scans to timeout after an hour has been resolved.
- A bug that prevented viewing third-party attributes for assets without Rumble attributes has been resolved.
- A bug that affected the first seen and last seen dates for assets imported from Tenable.io and Nessus has been resolved.
- A bug that persisted a task's error message when copying the task has been resolved.
- Fingerprint improvements.

2.14.1

2022-06-06

- A bug that could prevent the Tenable connector from importing large sites has been resolved.

2.14.0

2022-06-06

This release is a roll-up of the 2.13.x updates in addition to the following changes.

- Rumble Enterprise customers can now run scans using Rumble-hosted Explorers.
- Rumble Enterprise customers can now sync asset, software, and vulnerability data from the Tenable.io API and upload data from Nessus scan files.
- The Azure integration can now import Azure Function Apps.
- The ServiceNow integration can now export asset subnet tags.
- Fingerprint improvements.

2.13.7

2022-05-27

- The size limit for cloud-hosted scans has been increased from 6GiB to 10GiB.
- A bug that prevented searching software by version has been fixed.
- A bug that presented input fields for unselected credential types has been fixed.
- Fingerprint improvements.

2.13.6

2022-05-24

- Scans can now specify `asn4:<id>` and `country4:<2-letter ISO code>` in the scan and site scope and exclusion fields.
- Bogus results caused by firewall interference are now automatically ignored in more cases.
- Single organization administrators can now manage non-global credentials for their organization.
- The Organization Overview report now includes asset tags when available.
- Fingerprint improvements.

2.13.5

2022-05-19

- A bug that prevented the self-hosted installer from completing has been resolved.

2.13.4

2022-05-19

- The Censys integration now reports the `observed_at`, `extended_service_name`, `perspective_id`, and `source_ip` for all services.
- Fingerprint improvements.

2.13.3

2022-05-13

- The size limit for cloud-hosted scans has been increased from 4GiB to 6GiB.
- Fingerprint improvements.

2.13.2

2022-05-11

- A bug affecting some users of scan templates has been fixed to ensure the scan configuration is followed.
- A bug preventing the “TLS serial numbers” report from properly running has been fixed.
- Fingerprint improvements.

2.13.1

2022-05-09

- A bug that could lead to event processing terminating early has been resolved.

2.13.0

2022-05-09

This release is a roll-up of the 2.12.x updates in addition to the following changes.

- Asset information can now be viewed in aggregate across all organizations in the dashboard and asset inventory.
- The AWS integration now imports Lambda instances.
- The Azure integration now imports load balancers and AzureSQL instances.
- The GCP integration now imports load balancers and CloudSQL instances.
- Self-hosted installations now use an in-process task scheduler instead of system cronjobs for maintenance tasks.
- Fingerprint improvements.

2.12.12

2022-05-03

- A bug that could cause the service to reload during task processing has been resolved.
- Fingerprint improvements.

2.12.11

2022-05-02

- Software associated with assets is now tracked in the Rumble inventory.
- Rumble Enterprise customers can now sync asset and software data from the SentinelOne API.
- The AWS integration now imports RDS instances.

- The Azure integration now imports scale set virtual machines.
- AWS credentials enabling STS assume role workflows are simplified to only require a role name.
- Credential access can now be toggled to allow or disallow all organizations during credential creation.
- An intermittent issue that caused some external invitation emails to be missing activation codes has been fixed.
- A performance issue that could cause long load times for the credentials page has been fixed.
- A bug that could prevent group members from being displayed on an organization's users page has been fixed.
- Fingerprint improvements.

2.12.10

2022-04-19

- Tags can now be applied, updated, and deleted in bulk using the API.
- Speed of the RFC 1918 Coverage Report has been greatly improved.
- Fingerprint improvements.

2.12.9

2022-04-18

- Assets with external IP addresses will now be tagged with their geographic location and ASN when available.
- The CrowdStrike and Miradore integrations can now be run as scan probes from the console and scanner CLI.
- A bug that could prevent deleting services from the services inventory has been resolved.
- A bug where certain analysis tasks could error when an asset-query-results rule is enabled has been resolved.
- Fingerprint improvements.

2.12.8

2022-04-14

- A bug that prevented new CrowdStrike credentials from being stored has been fixed.

2.12.7

2022-04-14

- The Organization Overview report can now be generated and emailed to desired recipients on a recurring schedule.

- The organization users table now displays effective access for each user.
- A bug that could prevent adding users to groups has been fixed.
- Fingerprint improvements.

2.12.6

2022-04-12

- A bug that caused VMware instances with non-unique UUIDs to be handled incorrectly has been fixed.
- A bug that allowed IPs not in the scan scope to be used as primary addresses has been fixed.
- A bug that reported Windows OSes incorrectly for VMware has been fixed.
- The CrowdStrike integration now generates downloadable task data that can be used for importing CrowdStrike assets.
- Added ability to truncate syslog to a specified line length.
- Fingerprint improvements.

2.12.5

2022-04-08

- The scan configuration site scope warning now accurately reflects the site default scope.
- A bug that prevented searching for bssid wireless values has been fixed.
- Fingerprint improvements.

2.12.4

2022-04-07

- The Scan menu now provides an option to run a new scan using an existing template.
- The Alert Rules form now handles very long queries in the Test Query action.
- A bug that led to incorrect dashboard stats for multi-site organizations has been fixed.
- A bug that caused reports for certain AWS attributes to show empty results has been fixed.
- A bug that caused stale SNMP credentials to stay associated with an asset has been fixed.
- The self-hosted `rumblectl` update command now also applies content updates.
- Fingerprint improvements.

2.12.3

2022-04-06

- The AWS connector now tags each instance with the associated AWS account email.
- The CrowdStrike connector now handles API service outages more gracefully.

- A bug that prevented the Organization Overview report from being visible in Rumble Professional has been fixed.
- A bug that led to the scan engine logging a debug message related to LDAP has been fixed.
- A bug that led to visual errors on the asset details screen has been resolved.
- Fingerprint improvements.

2.12.2

2022-04-05

- The Scan menu now links to Scan Template selection with a search interface.
- The individual probe options in the Scan Config screen are now consistently sorted.
- A bug that could lead to partial stats being shown in the dashboard for multi-site organizations has been resolved.
- A bug that could prevent content self-hosted content updates from working when /opt was on a different file system from /tmp has been fixed.
- A bug that could lead to duplicate pre-built queries in self-hosted installations has been resolved.
- Fingerprint improvements.

2.12.1

2022-04-05

- A bug that could prevent scan templates from being saved has been resolved.
- Fingerprint improvements.

2.12.0

2022-04-04

This release is a roll-up of the 2.11.x updates in addition to the following changes.

- A new print-friendly Organization Overview report is now available.
- A new integration with the Google Cloud Platform is now available.
- External users may be invited to the cloud console.
- Scans templates are now available.
- Self-hosted instances now sync pre-built queries from the cloud.
- The scan engine now supports the Kerberos and LDAP protocols.
- Fingerprint improvements.

2.11.16

2022-03-29

- A bug that prevented exact = attribute matches from working has been resolved.

2.11.15

2022-03-28

- The self-hosted platform now supports scan imports larger than 4GiB.

2.11.14

2022-03-28

- The inventory search now supports a wider range of UTF-8 input for search patterns.
- The API now handles temporary maintenance-related errors more consistently.
- Fingerprint improvements.

2.11.13

2022-03-25

- Long-running export requests now timeout after 90 minutes, from 30 minutes, for large organizations.
- Fingerprint improvements.

2.11.12

2022-03-22

- Fingerprint improvements.

2.11.11

2022-03-20

- The update process for Explorers on the Windows platform is now more resilient to EDR/AV interference.

2.11.10

2022-03-19

- Fingerprint improvements.

2.11.9

2022-03-17

- A bug that could lead to some metric queries timing out has been resolved.

2.11.8

2022-03-17

- A bug that could lead to concurrent task processing in the same organization has been resolved.
- Fingerprint improvements.

2.11.7

2022-03-16

- A bug that could lead to timeouts with large CrowdStrike imports has been resolved.

2.11.6

2022-03-16

- Fingerprint improvements.

2.11.5

2022-03-15

- The self-hosted CLI now supports setting the superuser role and resetting MFA.
- The scan engine now detects the Veeam Distribution Service API.
- A bug that prevented {{organization.name}} from being used in certain organization alert templates has been resolved.
- A bug that prevented login events from being displayed in the Events view has been resolved.
- A bug that prevented sso-login events from being recorded when the connection was terminated mid-event has been resolved.
- A bug that allowed update requests to interrupt scans has been resolved.
- Incorrect outlier search result links have been fixed.

2.11.4

2022-03-11

- Alert rule queries now take into account assets found through integrations.
- The console now correctly filters bogus ARP replies over a reasonable threshold.

2.11.3

2022-03-11

- The self-hosted `rumblectl set-role` command now also supports setting the superuser role.
- The self-hosted `rumblectl reset` command now also resets the MFA token.
- The Asset CSV import now supports cell widths of up to 16,384 characters (from 1,024).
- A bug that prevented the self-hosted server from restarting in out-of-memory conditions has been resolved.
- A bug that caused the Network Switch report to sometimes show a 500 error has been resolved.
- Fingerprint improvements.

2.11.2

2022-03-10

- The response time when exporting assets via the API has been improved.
- The Asset inventory now displays the subnet tag descriptions on hover of subnet tags.
- The Network Switch report can now be launched from the reports page.
- Autocompletion of search keywords has been added for Sites and Queries.
- A bug that caused the SNMP probe to default to community strings “public,private” when no communities were provided has been resolved.
- A bug that prevented some third-party connectors from recording a task-completed event has been resolved.
- A bug that caused self-hosted Explorer and Scanner offline updates to fail has been resolved.
- Fingerprint improvements.

2.11.1

2022-03-08

- The SSO group mapping form now displays a relevant error when the groups list is empty or no group has been selected.
- The group column in the user table was not meant to be sortable and this has been fixed.
- The group mappings tab no longer shows when a user has SSO disabled.
- A bug that prevented `service_ports_tcp` and `service_ports_udp` search keywords from working as intended has been resolved.

- The dashboard has been updated to display a helpful tooltip icon on the RTT latency chart that defines certain terms and metrics.
- A regression in the VMware connector has been resolved.
- Fingerprint improvements.

2.11.0

2022-03-07

This release is a roll-up of the 2.10.x updates in addition to the following changes.

- The dashboard has been updated and now shows both most and least seen values for most stats.
- The dashboard now has CSV exports for all stats and links to deeper views of each given stat.
- The AWS and Azure integrations are now available to Professional Edition customers.
- The Azure integration can now run from the console, Explorer, or scanner as a probe.
- The Azure integration now identifies VM operating system information using disk image fingerprints.
- The Azure integration now tracks the clientID, tenantID, and subscriptionID as attributes for each asset.
- The Azure integration now creates a site per subscription ID.
- The AWS integration now creates a site per account in addition to the existing site per VPC capability.
- The AWS integration now supports using a provided session token.
- The AWS integration now tracks the account name as an attribute for each asset.
- The Account API now supports group management.
- Asset outliers are now tracked in the inventory and within the asset details page.
- Enterprise Edition customers can now access the Outlier Summary and Specific Outlier reports.
- Enterprise Edition customers can now map users to groups based on SAML attribute rules.
- The Explorer console URL can now be set through the RUMBLE_CONSOLE environment variable.
- The web console now flags under-resourced Explorers in the Deploy view.
- The web console now allows admins to force user logouts from the Team page.
- The web console now allows limited administrators to view users and create new projects.
- The scan engine now spends less time on per-VLAN SNMP enumeration when the device does not support it.
- The scan engine now supports full SNMP v1 enumeration using non-bulk lookups, if necessary.
- The scan engine is now much more conservative on a wider range of ICS ports.
- The scan engine is now much more friendly to fragile Lantronix devices.
- The scan engine now supports the Lantronix device discovery protocol.
- The scan engine now detects the Java Debug Wire Protocol (JDWP).
- The scan engine now detects and uses Qualys Cloud Agent correlation IDs.
- The scan engine now reports more information from NTP services.
- The self-hosted platform now supports a generate-certificate command.

- A bug that could lead to stale asset attributes remaining across scans has been resolved.
- A bug that could prevent the host-ping feature from finding all hosts has been resolved.
- A bug that led to broken search links in the task details page has been resolved.
- Fingerprint improvements.

2.10.6

2022-02-24

- A bug that prevented CIDR addresses in the default scan scope of a Site from being used has been fixed.
- Fingerprint improvements.

2.10.5

2022-02-22

- Asset correlation has been improved for a variety of corner cases, including Cisco Nexus switches.
- A bug that caused stale IPv6 addresses and UDP services to remain between scans has been resolved.
- A bug that prevented single-org admins from seeing users on the team page has been resolved.
- Fingerprint improvements.

2.10.4

2022-02-17

- OS EOL dates are now reported for Red Hat Enterprise Linux, Fedora, and CentOS.
- Fingerprint improvements.

2.10.3

2022-02-15

- The asset route pathing report is out of beta.
- Fingerprint improvements.

2.10.2

2022-02-11

- The Account API now supports group management through new endpoints.

- Asset and service search now supports new keywords for matching primary and secondary addresses.
- The dashboard now tracks how many assets have been seen in the last 30 days across all sources.
- The AWS integration now supports using a provided session token.
- The asset route pathing report is now more accurate.
- A bug that prevented some AWS asset attributes from being populated has been resolved.
- A bug that mangled UTF-8 characters in the subject and message body of e-mail notifications has been fixed.
- Fingerprint improvements.

2.10.1

2022-02-07

- A bug in the HTTP scanner that could prevent images from being captured correctly has been resolved.

2.10.0

2022-02-07

This release contains an important security fix:

- A security issue has been resolved in the SSO SAML handler. This issue was found during internal review and could be abused to trigger a denial of service or limited leak of application internal data by an unauthenticated attacker.

This release is a roll-up of the 2.9.x updates in addition to the following changes.

- The team page now supports user groups, providing more options when managing permissions/roles across your users.
- IPv6 support now includes link-local asset discovery and PTR lookups for the DNS/mDNS probes.
- AWS assets can now be synced from the standalone scanner, as a scan probe in the console, or imported from previous AWS connector tasks.
- A bug where invited users skipped the initial SSO login when joining an organization with required SSO settings has been resolved.
- A bug that prevented public IP addresses from populating an AWS asset's IPv4 attribute has been fixed.
- A bug where stale reverse DNS attributes could persist on rescanned assets has been resolved.
- A bug where the services in an asset view were not properly sorted has been resolved.
- The queries page now displays an Updated column containing the last-modified date and time for each query
- The queries page now supports query execution across all assets, regardless of alive status.
- Asset subnet tags are now included in JSON and XML asset exports.

- Fingerprint improvements.

2.9.14

2022-01-28

- HP iLOs will no longer be merged into their host assets when they share a MAC address.
- A bug that prevented services from displaying after a third-party import has been resolved.
- A bug that prevented Asset Modify rules from updating the HW field has been resolved.
- A bug that could cause the CLI scanner to stack trace has been resolved.
- The CLI Scanner `censys-db` sub-command now requires less memory.
- Fingerprint improvements.

2.9.13

2022-01-27

- Censys Avro files can now be converted to a database for faster lookups.
- Fingerprint improvements.

2.9.12

2022-01-25

- A regression that could lead to login errors after bulk permission updates has been fixed.
- A regression that removed the service names from the asset details page has been fixed.
- AWS internal hostnames are now reported in the asset name list.
- Fingerprint improvements.

2.9.11

2022-01-24

- Nmap XML exports are now much faster.
- Fingerprint improvements.

2.9.10

2022-01-20

- The scan engine now limits the SNMP enumeration speed to the Max Host Rate, reducing CPU usage on older switches.
- The scan engine now ignores additional cases of FortiGate HTTP interception.

- Fingerprint improvements.

2.9.9

2022-01-18

- The scan engine now accepts IPv6 addresses and resolves AAAA records for hostnames.
- The scan engine now skips protocol probes on TCP port 9106.
- A bug that prevented uploading very large scans has been fixed.
- Fingerprint improvements.

2.9.8

2022-01-14

- Added an option to export only selected assets, services, or wireless.
- Added a “Every N Hours” recurring task frequency option.
- Autocompletion of search keywords has been added for Organizations, Tasks, and Events.
- AWS and Azure connectors no longer set asset alive status and no longer are counted as offline or back online in the change report.
- Rules now show when they were last processed, whether they triggered their action, and any error that occurred as a result.
- The coverage report can be filtered by site.
- A bug that could allow duplicate CrowdStrike assets after an import has been fixed. Any resulting duplicates are eliminated on the next CrowdStrike task run.
- A bug which could lead to stalled rule processing has been fixed.
- A bug that prevented importing operating system information from CrowdStrike for some Linux devices has been fixed.
- A bug where scanning of some Lexmark printers interfered with the printer’s job queue has been fixed.
- Fingerprint improvements.

2.9.7

2022-01-10

- The scanner now supports configuration of reverse DNS timeouts and the SSH username.
- Scan tags can now be provided for scan import tasks.
- The `closedPortsMap` field has been removed from JSON exports.
- CrowdStrike connector tasks now move preexisting CrowdStrike-sourced assets into matching scanned assets across sites.
- A bug where task progress (on hover) could exceed 100% has been fixed.
- A bug that caused the Azure integration to occasionally skip public IPs has been fixed.
- A bug that caused a CrowdStrike connector task to send an API request exceeding length limits in specific instances has been fixed.

- Fingerprint improvements.

2.9.6

2021-12-23

- A bug that could lead to some events being processed incorrectly has been resolved.
- Event templates now truncate results correctly.

2.9.5

2021-12-22

- A scan engine bug that could lead to an “invalid exclusions” error has been resolved.

2.9.4

2021-12-21

- The Query search now supports result count selection and remembers the setting between views.
- The scan engine now correctly excludes broadcast addresses from the scan scope.
- The Azure connector now ignores canceled subscriptions automatically.
- The hostname selection logic has been improved for Canon printers.
- The Explorer service now starts up slightly faster on Windows.
- The Censys AVRO importer is now 4 to 8 times faster.
- Fingerprint improvements.

2.9.3

2021-12-15

- The CrowdStrike integration has been updated to improve correlation with existing assets.

2.9.2

2021-12-13

- A bug that prevented some AWS organizations from working with STS AssumeRole has been fixed.
- A bug that persisted service products after asset changes has been fixed.
- A bug that hid the Task Change Report has been fixed.
- Improved product detection for Logstash and Neo4J.
- Fingerprint improvements.

2.9.1

2021-12-08

- The CrowdStrike integration now uses the Scroll API to better support large organizations.

2.9.0

2021-12-06

This release is a roll-up of the 2.8.x updates in addition to the following changes.

- The new Site Comparison report shows differences in assets between two sites, which can be in different organizations.
- The team page now supports bulk user import and bulk permission management.
- The layer 2 topology report has been updated with a search filter, site filter, and visual improvements.
- The network bridges report has been updated with a site filter and visual improvements.
- A bug that caused the `has_public` search filter to flag certain IPv6 addresses has been fixed.
- A bug that caused project deletion to create new, blank organizations has been fixed.
- A bug that caused Censys imports to mark other assets as offline has been fixed.
- The bundled npcap version has been upgraded to 1.60.
- Fingerprint improvements.

2.8.14

2021-12-03

- A bug that prevented scans from running when non-loopback 127.x networks were present has been fixed.

2.8.13

2021-11-24

- A regression in the TLS version enumeration has been fixed.
- Teredo addresses are no longer considered public IPs.
- Improved detection of Chromebooks and ChromeOS.
- Fingerprint improvements.

2.8.12

2021-11-21

- The self-hosted platform now supports internal proxies for external API connections.
- The self-hosted platform now supports internal webhook destinations for alerts.

2.8.11

2021-11-19

- The annotator role is now available.
- Fingerprint improvements.

2.8.10

2021-11-18

- A bug that prevented the Azure connector from working in self-hosted mode has been resolved.
- The last hop calculation for the TCP traceroute is now more accurate.
- Fingerprint improvements.

2.8.9

2021-11-16

- Credentials can now be configured for single IP addresses and IP ranges in addition to CIDRs.
- The scan engine now performs a light traceroute when an open TCP port is found.
- The scan engine now tests for IP forwarding during scans of link-local targets.
- The scan engine now includes Rumble/2 in HTTP user-agent strings.
- The scan engine now limits ARP traffic to the Max Host Rate.
- The OS EOL date for Windows Server 2019 (1809) has been updated.
- Fingerprint improvements.

2.8.8

2021-11-12

- The self-hosted platform now supports custom CSP headers to support external resources.
- Tags with the case-insensitive key of "name" are now treated as additional hostnames.
- Tags containing spaces are now consistently handled by conversion to underscore.
- Meraki DNS interception is now ignored in a wider variety of configurations.
- The CrowdStrike integration now tracks the last 10 recent logins per asset.
- A regression in the Apple macOS end of life calculation has been fixed.
- A few missing icons have been restored to the inventory view.

- A panic in the FreeBSD scan processing chain has been fixed.
- Direct print services on 9002 are no longer fingerprinted.
- TLS versions are now consistently formatted.
- Fingerprint improvements.

2.8.7

2021-11-07

- A bug that caused daily scans to be scheduled incorrectly in UTC forward time zones has been resolved.
- The self-hosted platform now respects proxy settings for external data sources (CrowdStrike, etc.).
- Improved tvOS and Crestron fingerprints.
- Additional UPnP fingerprints.

2.8.6

2021-11-06

- Improved tvOS, homepodOS, and bridgeOS fingerprints.

2.8.5

2021-11-06

- The self-hosted installer now supports manual database configuration.
- The self-hosted platform now includes a database `verify` subcommand.
- Improved iOS device identification.
- Fingerprint improvements.

2.8.4

2021-11-05

- The VMware probe now handles vCenter instances configured with multiple datacenters.
- A race condition that could lead to Explorer updates mid-scan has been resolved.
- The RDP TLS fingerprint will no longer break matching during asset correlation.
- Stopped scans now indicate which user stopped them in the error message.
- Active probes are now disable for some Lantronix and Rockwell PLC ports.
- Fingerprint improvements.

2.8.3

2021-11-04

- SAML SSO now specifies that the required NameID Format is unspecified, for Azure AD compatibility.
- S3 storage operations which fail are retried.
- A bug in Azure credential validation is fixed.
- Additional detail has been added to the field help on the scan form.

2.8.2

2021-11-02

- Recurring tasks no longer schedule jobs when the previous job is still queued.
- A bug that caused VMware-based OS detection to fail has been resolved.
- Fingerprint improvements.

2.8.1

2021-11-02

- A bug that caused out-of-scope assets to be marked as offline during scans has been resolved.

2.8.0

2021-11-01

This release is a roll-up of the 2.7.x updates in addition to the following changes.

- Support for Censys Search API and Censys Data imports (Enterprise).
- Fingerprint improvements.

2.7.11

2021-10-27

- An issue with Windows Explorer updates has been resolved.

2.7.10

2021-10-24

- An issue with restrictive umasks on the self-hosted platform has been resolved.
- Support for VMware vCenter/ESXi virtual machine discovery (Enterprise).
- Asset merging from third-party data sources has been improved.

- TLS fingerprints are now reported as SHA256 hashes (base64).
- Credentials for SNMP v2/v3 and VMware can be managed globally.
- Serial numbers from A10 devices are now collected via SNMP.
- The SNMP v3 probe now supports multiple credentials.
- Scan configuration now has a Credentials tab.
- Detailed task logs can be downloaded.
- Light UX improvements and bug fixes.
- Fingerprint improvements.

2.7.9

2021-10-21

- A permissions issue with the self-hosted platform has been resolved.

2.7.8

2021-10-15

- The runZero Explorer on Windows now explicitly sets the service to automatic start.
- RFC 1918 scans can be launched from the main Scan menu.
- Fingerprint improvements.

2.7.7

2021-10-15

- Credential validation issues with Azure now log detailed errors.
- The self-hosted platform now supports Debian 9.
- Fingerprint improvements.

2.7.6

2021-10-13

- A bug that led to an application error during Azure connector configuration has been resolved.
- Fingerprint improvements.

2.7.5

2021-10-13

- A bug that prevented certain AWS and Azure assets from being imported has been resolved.
- Fingerprint improvements.

2.7.4

2021-10-08

- The self-hosted platform now supports Oracle Linux 7 and 8.
- Fingerprint improvements.

2.7.3

2021-10-07

- The AWS integration is now much faster for large numbers of accounts.

2.7.2

2021-10-06

- A bug that made it difficult to update existing Azure tasks has been resolved.

2.7.1

2021-10-05

- A bug that prevented the date picker from showing arrow icons has been resolved.
- A bug that made it difficult to update existing AWS tasks has been resolved.
- The sidebar is now collapsible using the chevron icon at the top.
- The “Processing” link states are now handled more consistently.

2.7.0

2021-10-05

This release is a roll-up of the 2.6.x updates in addition to the following changes.

Integrations

- The Azure VM connector now supports multi-subscription and multi-directory access.
- The AWS EC2 connector now supports ELB load balancers as importable assets.
- Connector credentials are now automatically validated on save.
- The Splunk add-on now supports self-hosted console endpoints.
- The Splunk add-on now optionally imports asset services.
- The Splunk add-on has been updated to use jQuery 3.5.0.
- CrowdStrike asset merging has been improved.

Self-hosting

- The self-hosted installer now includes the Explorer and scanner binaries.
- The self-hosted console now supports detailed TLS configuration.
- The self-hosted console now runs as an isolated subprocess.
- The self-hosted console no longer enforces API rate limits.

User experience

- Asset and service trends are now shown on the dashboard.
- The Scan configuration view has been overhauled and simplified.
- Imported scans are now tracked for the RFC 1918 coverage report.
- The Services inventory now supports new address-related search keywords.
- The Assets and Services inventory now support wildcard searches of hostnames with anchored patterns.
- The Screenshot inventory is now faster and shows the correct total count.
- A bug in the RFC 1918 coverage report that could lead to skipped IPs was fixed.

Authentication

- Users with standard accounts that authenticate using SSO are now converted into SSO-only accounts.
- SSO is supported for multiple domains using IdP or SP initiated authentication.
- The MFA challenge now provides a Retry button for browsers that require user interaction (Safari).
- The MFA enrollment now supports token or platform authentication mode as separate options.
- Admin users can now set the first and last names of other users.
- The Explorers, scanners, MSI wrapper, and verifier are now signed using a new EV certificate.

Scanner and fingerprinting

- Subnet ping and host ping are now included in the Professional tier.
- OS EOL tracking is now enabled for Windows 10 and APC firmware.
- Windows 10 and Server 2019 OS versions are now tracked by range.
- The Explorer and CLI scanner now detect and report an error when run within the WSL/WSL2 environments.
- The CLI scanner upgrade now supports the `-force` option.
- The scan engine now detects Bitdefender remotely.
- A regression in the ARP probe on newer Windows builds has been resolved.
- Fingerprint improvements

2.6.4

2021-09-16

- The scan engine now detects Azure's OMI WSMAN implementation.
- Fingerprint improvements.

2.6.3

2021-09-16

- The scan engine now detects WSMAN, ADB, and InfluxDB services.
- Fingerprint improvements.

2.6.2

2021-09-14

- A regression in SMB v1 detection has been resolved.
- Fingerprint improvements.

2.6.1

2021-09-08

- A bug that prevented some Azure VMs from being imported has been fixed.

2.6.0

2021-09-07

- Rumble Enterprise customers can now sync virtual machine inventory from the Microsoft Azure cloud.
- The CrowdStrike connector has been overhauled to improve asset merging and avoid duplicates.
- OS end-of-life dates for Windows, macOS, Ubuntu, Debian, and iLO assets are now tracked.
- The self-hosted version of Rumble now supports offline mode & offline updates.
- The self-hosted version of Rumble now supports RHEL 7 in offline mode.
- The scan engine now surfaces NFS exports via discovered mountd services.
- The scan engine now returns details for discovered PPTP services.
- The dashboard loads faster for large organizations.
- The UI now includes new, custom icons.
- Fingerprint improvements.

2.5.8

2021-08-30

- Stale asset expiration now applies to third-party sourced assets.
- The scan engine now reports PPTP services.
- Fingerprint improvements.

2.5.7

2021-08-27

- The CrowdStrike connector has been overhauled to improve merging and avoid duplicates.
- The scan engine now reports NFS exports.
- Fingerprint improvements.

2.5.6

2021-08-23

- The scan engine now implements the Cisco [layer 2 traceroute](#) protocol thanks to [Chris Marget's cisco-l2t](#) project.
- TCP port 1720 is no longer included in the defaults. This may be re-enabled once H.323 is fully implemented.
- The scan engine now handles mangled SNMP responses better.
- The HTTP/2 protocol is now reported at the asset level.
- Fingerprint improvements.

2.5.5

2021-08-19

- A regression in the service attribute report has been fixed.
- The scan engine now reports additional SSH attributes.

2.5.4

2021-08-19

- The Explorer and scanner now support the Windows arm64 platform.
- A scan engine hang in the DCERPC probe has been resolved.
- Fingerprint improvements

2.5.3

2021-08-18

- Third-party data source attributes are now included in all exports.
- Third-party attributes now use the @source.type syntax for search.
- The Merge feature in the asset inventory is now more consistent.
- Large target exclusion lists are now supported for sites and scans.
- Unresolvable hostname excludes are now ignored automatically.
- The scan engine now records more information from McAfee ePO agents.
- Fingerprint improvements

2.5.2

2021-08-12

- Automatic queries are now available to Professional users as well as Enterprise.
- A bug that led to some Windows desktops having the wrong type has been fixed.
- CrowdStrike assets are now matched more accurately against Rumble assets.
- The scanner now skips active protocol detection on port 9999.
- Fingerprint improvements

2.5.1

2021-08-05

- The default TCP port list now includes more SolarWinds products as well as port 7676 for JMS/IMBroker.
- The estimated runtime for scans now takes into account the TCP port list (and excludes).
- Juniper switch fingerprinting now uses a Juniper-specific OID instead of sysDesc.
- Additional bogus SIP ALG services are now ignored by the runZero scanner.
- A bug that prevented offline-agent events from being generated in certain situations has been resolved.
- A bug that could result in Explorers not reconnecting properly after an update has been resolved.
- A bug that showed an incorrect bandwidth calculation in the task view has been resolved.
- A bug that led to an error on CSV export with mixed-source assets has been resolved.
- Fingerprint improvements.

2.5.0

2021-08-03

- Rumble Enterprise customers can now sync AWS EC2 assets across accounts using STS roles.
- Rumble Enterprise customers can now sync asset data from the CrowdStrike Falcon API.
- The scan engine now better differentiates between Windows workstation and server variants.
- The scan engine now detects various asset attributes and services using DCERPC.
- The scan engine now detects multi-homed assets using DCERPC.
- The dashboard can now show stats across all sites or just a specific site.
- The new Unmapped MAC report highlights unscanned assets by switch port.
- The Reports page has been improved with a new layout and inline search.
- The Queries tab has been moved to a new navigation item.
- Fingerprint improvements.

2.4.4

2021-07-26

- The `superuser` role is now available as a default permission for SSO users.
- The scan engine now gathers data from the Windows DCERPC endpoint mapper.
- Fingerprint improvements.

2.4.3

2021-07-18

- Fingerprint improvements.

2.4.2

2021-07-14

- Limited layer-2 topology graphs derived from ARP data are now available for environments using Fortinet switches.
- The Topology and Network Bridges graphs can now be exported as PNG images.

2.4.1

2021-07-13

- This build fixes a [bug in the Go runtime](#) that could allow a remote attacker to cause a recoverable panic in the Rumble services and scan engine (CVE-2021-34558).

2.4.0

2021-07-13

- Rumble Enterprise customers can now sync asset data from Amazon Web Services EC2 and Miradore MDM data sources.
- Rumble Enterprise customers can now self-host the platform on RHEL and CentOS distributions.
- Credentials are now managed at the account level with per-organization access.
- The Rumble self-hosted CLI offers new features and a better user experience.
- Fingerprint improvements.

2.3.5

2021-07-04

- A bug that prevented Explorer upgrades and scan stop requests from processing while a scan was active has been fixed.
- A bug that led to the subnet ping mode missing subnets during large scans cases has been fixed.
- A bug that led to an Explorer showing as offline unexpectedly has been fixed.
- Fingerprint improvements.

2.3.4

2021-06-26

- A regression in the Screenshots inventory tab has been fixed.
- Fingerprint improvements.

2.3.3

2021-06-24

- Tabs, fonts, and styles have seen a light update.
- A number of small UX bugs were addressed.
- Fingerprint improvements.

2.3.2

2021-06-16

- The RFC 1918 coverage report now supports a starting date that can be used to exclude older scans.
- The Your team page is now searchable, sortable, and supports bulk user actions.

- Fingerprint improvements.

2.3.1

2021-06-09

- A bug that prevented single-organization users from viewing sites and tasks has been fixed.
- A bug that led to offline assets not actually being marked as offline has been fixed.
- A bug that could prevent full enumeration of Cisco Catalyst switches has been fixed.
- Fingerprint improvements.

2.3.0

2021-06-08

- The new RFC 1918 coverage report highlights unscanned address space and hinted ranges.
- SNMP v3 enumeration of Cisco Catalyst switches now handles per-vlan port mappings.
- Fingerprint improvements.

2.2.5

2021-05-29

- Tags are now always displayed with = instead of : to match the search engine syntax.
- The Subnet Ping and Host Ping modes are now more reliable on large scans.
- Fingerprint improvements.

2.2.4

2021-05-26

- A bug that led to a stack trace in the rumblectl command for self-hosted mode has been fixed.
- Fingerprint improvements.

2.2.3

2021-05-17

- The self-hosted platform now removes older scanner/Explorer binaries during updates.
- The scan engine now pulls layer-2 information from Force-10 switches.
- The scan engine now ignores CheckPoint SMTP and SIP interception.
- The scan engine now extracts hostnames from Zyxel switches.
- An invalid fingerprint for Cisco IP phones has been fixed.

- Multiple notifications can now trigger from a single event.
- Agent fields are now included in the scan-started event.
- Fingerprint improvements.

2.2.2

2021-05-16

- The scan engine now extracts additional information from Zyxel switches.
- The Explorers page now supports sorting, searching, and tagging.
- Fingerprint improvements.

2.2.1

2021-05-14

- The scan confirmation dialog now warns when a mix of public and private IPs are in the scope.
- The SNMP v3 probe now supports sha224, sha256, sha384, and sha512 authentication.
- The SNMP v3 probe now supports aes192, aes256, aes192c, and aes256c encryption.
- The self-hosted platform now includes a CLI to manage user accounts.
- Fingerprint improvements.

2.2.0

2021-05-11

Rumble 2.2.0 is a roll-up of previous 2.1.x releases along with some additional changes and features.

Web console

- Enterprise customers can now export an HP iLO report which includes serial numbers, physical hardware information, and other fields useful for warranty tracking and server inventory.
- Virtual machines now indicate the virtualization vendor in the asset hardware field and can be searched and filtered by vendor.
- Virtual machines are now labeled with an icon in the asset view, and router icons are displayed with the other icons.
- Virtual machine and legacy MAC prefixes are no longer used for age calculation, as they gave inaccurate results.
- A new report for virtual machine vendors has been added.
- Vendors in the NDAA Section 889 report have been expanded to include Aztech and subsidiaries.
- The Name tag can be used to set a preferred hostname for any asset.
- Tags in inventory views can now be clicked to search the inventory.

- The alert notification templates can now include information about the name and internal IP address of the Explorer which carried out a scan.
- The alert rules list now shows which rules are currently enabled.
- Confirmation dialogs now require a typed response for destructive actions.
- Page layout has been improved for browser window widths between 920 and 1200 pixels.
- If a scan results in too many changes to list in the task report, the report now explains this.
- Progress bars now use standard meter elements for smoother updating and better accessibility.
- Icons and screenshots are now lazy-loaded to speed up initial page rendering.
- Task duration is now rounded up to the nearest minute.

Explorer and scan engine

- A better hostname is chosen for each asset by default.
- VLANs are now tracked on each asset where possible.
- Minecraft servers are now identified on the network.
- HP iLO scans now return additional information.
- Virtual machine hardware is reported if a better fingerprint is not available.
- Pulse Secure VPN devices running newer firmware are now identified correctly.
- Additional CA roots can be set via the RUMBLE_TLS_ADDITIONAL_ROOTCA variable.
- Fingerprint improvements.

Self-hosted platform

- Self-hosted installations now sync license changes during updates.
- Email validation is relaxed for the self-hosted platform.
- Install instructions now use `curl` instead of `wget` for improved robustness.

Bug fixes

- The last seen date for assets is now only updated when they have at least one open port and are therefore “alive”.
- The estimated scan times in recurring task exports are now accurate when default ports are used.
- A problem with password reset requests for accounts with no last name has been resolved.
- A possible scan stall issue caused by endpoint software on macOS has been resolved.
- A memory and CPU leak in the runZero Explorer has been resolved.
- User name validation now works correctly when editing user preferences.
- An issue with user invitations for SSO accounts has been resolved

2.1.7

2021-05-03

- The web console now includes a new HP iLO CSV export for warranty tracking (Enterprise).
- Virtual machine and legacy MAC address prefixes are no longer used for age calculation.
- Self-hosted installations now sync license changes during updates.
- Virtual machines are now represented in the asset HW field.
- VLANs are now tracked on each asset, where possible.
- Tags are now clickable and lead to inventory searches.
- Fingerprint improvements.

2.1.6

2021-04-28

- Rumble Scans on macOS no longer stall when ICMP scans are blocked by endpoint software.
- The web console now better supports browser widths between 920 and 1200 pixels.
- An issue with user invitations from an SSO account has been resolved.

2.1.5

2021-04-27

- The runZero Explorer and runZero Scanner now collect additional information from HP iLO nodes.
- Fingerprint improvements.

2.1.4

2021-04-26

- A memory and goroutine leak in the the runZero Explorer and runZero Scanner has been fixed.
- Fingerprint improvements.

2.1.3

2021-04-23

- The runZero Explorer and runZero Scanner now avoid probes on vendor-specific SunRPC services.
- The latest firmware of Pulse Secure VPNs is now fingerprinted correctly again.
- An issue with the confirmation dialogs on Chrome has been resolved.
- The dashboard now defines next steps for new organizations.

- The bundled npcap version has been upgraded to 1.31.
- Fingerprint improvements.

2.1.2

2021-04-19

- The runZero Explorer and runZero Scanner now detect the Minecraft service.
- Confirmation dialogs now require a typed response for permanent actions.
- Small bugs have been fixed in the default notification templates.
- Fingerprint improvements.

2.1.1

2021-04-14

- This update disables automatic npcap upgrades while we investigate a stall issue.

2.1.0

2021-04-13

Rumble 2.1.0 is a roll-up of previous 2.0.x releases along with some additional changes and features.

Web console

- Custom notification templates are now available and can be configured as Text, HTML, and JSON formats.
- Webhook notification channels can now include arbitrary HTTP headers for authentication.
- Tasks can now be exported as CSV and JSON from the Recurring and Search tabs of the task view.
- Tags be removed in bulk by specifying -tag in the inventory Tag dialog.
- Asset CSV bulk imports now ignore cases where there are extra fields.
- Specific TCP ports can now be excluded from the scan configuration.
- User invitation links can now be copied to the clipboard.
- UX tooltips are now easier to read across all platforms.
- Alert management is no longer organization-specific.
- Exports with complex queries are now much faster.

Integrations

- Device serial numbers can now be exported in Cisco Smart Net Total Care format for warranty checks.
- The Splunk add-on now supports proxy server configuration in version 1.0.11.

Explorer

- The Explorer now rejects scans tasks when there is not enough free disk space for the scan results.
- The Explorer now falls back to the install directory for temporary files if needed.
- The Explorer now tries to upgrade the npcap driver automatically on Windows.

Scanner

- AWS EC2 metadata is used to enrich scan results if the ec2:DescribeInstances permission is available.
- SNMP v2 enumeration of ports and vlan membership now uses community indexing automatically.
- Chromium installations using Snap packages are no longer used for web screenshots.
- An overly aggressive mDNS fingerprint for LG webOS has been fixed.
- EC2 instances now report the instance type as the hardware field.
- Additional bogus SIP helper responses are now ignored.
- LPD fingerprinting is now limited to a status request.
- Fingerprint improvements.

Events

- A single event will only trigger a single notification per unique notification channel.
- The agent-removed event is now generated for all web console agent removal actions.
- The offline-agent event no longer includes recently removed or forgotten agents.
- The task-failed event now includes the full organization and site name in the details.
- The task-failed event is now generated for agent restarts and timeouts.

Self-hosted platform

- SMTP configuration now supports additional TLS settings.
- SMTP errors are now logged to syslog correctly.
- Initial auto configuration is more complete.

2.0.14

2021-04-06

- The runZero Explorer and runZero Scanner will now gather AWS EC2 instance metadata where possible.
- Fingerprint improvements.

2.0.13

2021-04-05

- The runZero Explorer will now fallback to the install directory for temporary files.
- Fingerprint improvements.

2.0.12

2021-04-03

- Exports are now faster across the board, with major speedups for those using deep search queries.
- The Rumble Verifier has been updated and now supports 2.x macOS binaries.
- Various bugs related to the new tooltip implementation have been resolved.
- Chromium will no longer be used for screenshots when installed via Snap.
- Additional bogus SIP helper responses are now automatically ignored.
- Self-hosted installs now log SMTP-related errors.
- Fingerprint improvements.

2.0.11

2021-04-01

- Fingerprint improvements.

2.0.10

2021-03-30

- The runZero Explorer and runZero Scanner now collect additional Cisco-specific SNMP OIDs.
- The Cisco Serial Number export is now available for all licensed clients.
- Fingerprint improvements.

2.0.9

2021-03-30

- The runZero Explorer and runZero Scanner no longer stall in a rare case when enumerating SNMP endpoints.
- Fingerprint improvements.

2.0.8

2021-03-29

- The runZero Explorer now automatically cleans temporary files left by interrupted scans.
- The runZero Explorer now prepends “rumble-” to all temporary files.

2.0.7

2021-03-28

- The runZero Explorer out of disk space error has been improved for legibility.

2.0.6

2021-03-28

- The runZero Explorer now performs a disk space check before running a new scan.
- The runZero Explorer now reports disk related errors more reliably.
- Asset CSV imports are now accepted where the records have extra fields.
- Fingerprint improvements.

2.0.5

2021-03-26

- The runZero Explorer and runZero Scanner now better support enumeration of Catalyst switches.
- Administrators can copy the invite link for direct sharing with team members.
- Issues with image links on the self-hosted version have been resolved.

2.0.4

2021-03-26

- The runZero Explorer and runZero Scanner now handle the LPD protocol more carefully.
- Fingerprint improvements.

2.0.3

2021-03-23

- The runZero Explorer no longer reports an intermittent “no child processes” message on installation.

- The runZero Explorer and runZero Scanner now always return wireless results when iwlist is present.
- The web console now generates agent-removed events for each agent, including for bulk actions.
- The web console now generates task-failed events for agent restarts and timeouts.

2.0.2

2021-03-21

- The runZero Explorer now logs connectivity issues with the websocket protocol in more detail.
- The self-hosted Rumble platform now supports better automatic configuration.
- Fingerprint improvements.

2.0.1

2021-03-20

- A bug that resulted in deleted event rules being processed has been resolved.
- A bug that could lead to scan stalls in Subnet Ping mode has been resolved.
- Fingerprint improvements.

2.0.0

2021-03-16

Rumble 2.0 is a roll-up of the 1.16.x releases, along with the following changes:

Web console

- The new Rules Engine supports advanced alerts and automated asset updates.
- Organization-level Alerts are now managed at the global level.
- The Explore menu item has been renamed to Reports.
- The interface received light cosmetic updates.

Deployment

- runZero Explorers are now runZero Explorers to better indicate their function.
- runZero Explorers and runZero Scanners moved to the Deploy menu.
- Self-hosting of the full platform is now available.

REST API

- The Account API now provides organization, user, and event management.

- Support for ServiceNow ® ITOM integration via three new API endpoints.

Scan engine

- Microsoft Exchange and Outlook Web Access detection has improved.
- A stall bug in the subnet and host ping modes has been fixed.
- The number of bogus results ignored is now reported.
- The npcap driver has been upgraded to v1.20.
- Fingerprint improvements.

1.16.6

2021-03-06

- Layer-2 topology reports now display the best matching port and not all ports where an asset was seen.
- The runZero Explorer and runZero Scanner now handle subnet and host ping modes faster and more accurately.
- The runZero Scanner now processes gzip-compressed imports faster.
- Fingerprint improvements.

1.16.5

2021-02-27

- Assets that were previously identified with through a TCP RST, but otherwise had no services, have been removed from the platform.
- The runZero Explorer and runZero Scanner now collect more data about exposed SSH authentication methods.
- A regression related to asset tracking based on the TCP/IP fingerprint has been resolved.
- Fingerprint improvements.

1.16.4

2021-02-26

- The Rumble scan engine now ignores assets where all TCP ports are closed and no other services are available. This reduces the reporting of bogus hosts when scanning through certain firewalls.
- The task summary and task details now report how many assets were ignored due to having no valid services. This highlights how much network interference may be present.
- The macOS binaries have moved to a new code signing process. Agents for macOS that have been offline for some time may require a manual update.
- The live asset count and project asset count was calculated incorrectly for users with deleted organizations. This has been corrected.

- Search queries that contained string matches against <% and %> now work as expected.
- Fingerprint improvements.

1.16.3

2021-02-25

- Fingerprint improvements.

1.16.2

2021-02-24

- runZero Explorers can now have tags applied in the per-agent settings page.
- Events are now regularly generated for offline agents.
- Fingerprint improvements.

1.16.1

2021-02-21

- The Rumble scan engine now supports a maximum TTL for all scan traffic.
- The Rumble scan engine now supports subnet ping and host ping modes.
- The Rumble scan engine now distributes scan traffic more evenly across subnets.
- The Rumble scan engine now reports and tracks closed TCP ports.
- The Rumble scan engine now reports additional ICMP fields.
- The Rumble scan engine now auto-scales the group size.
- Apple macOS is now partially supported on ARM systems.
- Fingerprint improvements.

1.16.0

2021-02-09

- This includes a roll-up of all 1.15.x point releases.
- VMware ESXi versions are now reported correctly.
- Fingerprint improvements.

1.15.6

2021-01-31

- The Inventory Search, Exports, and Reports are now significantly faster for large organizations.
- Fingerprint improvements.

1.15.5

2021-01-28

- The Agents page will now flag any Windows Agents with an obsolete version of Npcap installed.
- Fingerprint improvements.

1.15.4

2021-01-26

- Fingerprint improvements.

1.15.3

2021-01-23

- The runZero Scanner and runZero Explorer now gather serial numbers from SNMP devices.
- The 169.254.0.0/16 subnet is no longer ignored when processing scan results.
- The runZero Scanner and runZero Explorer now detect the TeamViewer protocol.
- Partial site scans now consider ARP cache data from the entire site.
- The runZero Scanner now supports importing gzip-compressed scan data.
- The runZero Scanner and runZero Explorer now detect the CheckMK service.
- Fingerprint improvements.

1.15.2

2021-01-19

- The dashboard now links to the top 5,000 results for asset types and service details.
- Rumble-provided queries can now be saved as per-account copies and modified.
- Partial site scans will now use ARP cache entries from the rest of the site.
- Fingerprint improvements.

1.15.1

2021-01-16

- This resolves an issue with the Crestron probe that could cause concurrent scans on the same agent to hang.
- Fingerprint improvements.

1.15.0

2021-01-12

- This is a roll-up of all 1.14.x point releases.
- Fingerprint improvements.

1.14.9

2021-01-10

- The runZero Explorer and runZero Scanner now support the Crestron discovery protocol.
- The runZero Explorer and runZero Scanner now capture TLS fields from PostgreSQL services.
- Fingerprint improvements.

1.14.8

2021-01-06

- runZero Explorer proxy usage has been improved to handle additional corner cases.
- Fingerprint improvements.

1.14.7

2021-01-05

- runZero Explorer and runZero Scanner updates now use any proxies configured in the environment.
- Web screenshots now ignore any proxies configured in the environment.
- Fingerprint improvements.

1.14.6

2021-01-04

- A minor memory leak in the runZero Explorer and runZero Scanner's HTTP2 probe has been resolved.
- Web screenshots now ignore any proxy configured for the runZero Explorer communication.
- Web screenshots no longer leave zombies in additional environments.
- Fingerprint improvements.

1.14.5

2020-12-28

- A memory leak in the runZero Explorer and runZero Scanner has been resolved.
- Web screenshots no longer leave zombies in environments without init.
- Fingerprint improvements.

1.14.4

2020-12-24

- Services with empty virtual hosts will be consolidated into the first non-empty virtual host service where applicable.
- Subtasks created by a recurring scan will now carry the “defaults” parameters forward.
- Fingerprint improvements.

1.14.3

2020-12-19

- A memory leak in the runZero Explorer and runZero Scanner has been resolved.

1.14.2

2020-12-17

- The runZero Explorer and runZero Scanner runtime has been upgraded.
- The runZero Explorer and runZero Scanner now use npcap 1.10.
- The site scan API now handles custom probe configurations.
- The task stop API documentation has been updated.
- Fingerprint improvements.

1.14.1

2020-12-14

- Agents can now be automatically assigned to their connected sites from the Manage menu.
- Scan tasks configured through the API now handle the probes parameter correctly.
- PAN-OS virtual MACs are now ignored for asset correlation.
- Scan task parameters are now consistently normalized.
- Fingerprint improvements.

1.14.0

2020-12-08

- SNMP System Description fingerprints now take precedence over SSH-based OS matches.
- Fingerprint improvements.

1.13.11

2020-12-02

- A bug that caused certain API calls to return a 500 instead of a 400 error when called with a non-JSON content type has been resolved.
- Bogus services caused by captive portals, honeypots, and certain firewalls are now automatically ignored.
- Fingerprint improvements.

1.13.10

2020-12-01

- Agent to cloud communication is now limited to the `console.rumble.run` hostname.
- The breadcrumbs are now navigable across the product user interface.
- Fingerprint improvements.

1.13.9

2020-11-25

- Inventory search boxes now have autocompletion for search keywords (`name:`, `hardware:` and so on). Completion can be triggered via keyboard (`tab`, `enter`) or mouse.
- The documentation for search queries has been updated and cleaned up.
- Tag editing dialogs now autocomplete from your top 50 most used tags.
- Some search keywords have been updated to be more consistent.

1.13.8

2020-11-23

- Fingerprint improvements.

1.13.7

2020-11-20

- The fingerprinting engine was updated to support more precise device detection.
- Support for new mDNS fingerprints, including M1-based Apple devices.
- Various security fixes to address findings from an ongoing audit.
- Fingerprint improvements.

1.13.6

2020-11-19

- The task details page now shows the scan speed in the upper left section.
- Fingerprint improvements.

1.13.5

2020-11-13

- The runZero Explorer and Scanner have been updated to work with macOS Big Sur.

1.13.4

2020-11-12

- The runZero Explorer, Scanner, and Console now use the latest Go runtime.
- Fingerprint improvements.

1.13.3

2020-11-10

- Assets with both SMB v1 and v2 enabled are now fingerprinted more consistently.
- The BACnet probe now supports multiple ports per scan.
- Fingerprint improvements.

1.13.2

2020-11-08

- Session and login timeouts can now be configured from the Account Settings page.
- The Subnet Grid report is now faster and supports RTT, TTL, and Age color modes.
- Fingerprint improvements.

1.13.1

2020-11-05

- The new Account Settings page allows MFA to be enforced and provides the ability to block Support access.
- Fingerprint improvements.

1.13.0

2020-11-03

- This is a roll-up of all 1.12.x point releases.
- Fingerprint improvements.

1.12.9

2020-10-25

- The runZero Scanner and runZero Explorer now decompress non-negotiated gzip responses from HTTP services.
- Fingerprint improvements.

1.12.8

2020-10-25

- The Subnet Utilization report now sorts networks by density more accurately.
- The Subnet Grid report is now slightly faster with large networks.
- Asset correlation now ignores bogus replies for SMB and RDP.
- Fingerprint improvements.

1.12.7

2020-10-24

- The Team page now contains a tab showing which users have access to the active organization.
- Users with per-organization admin roles are now able to manage user accounts.
- Fingerprint improvements.

1.12.6

2020-10-22

- Users with the “No Access” permission are now able to manage their own account settings.
- Users can now have a Billing role that is limited to license and entity management.

1.12.5

2020-10-21

- Inventory searches with the `haspublic` keyword now handle multi-homed systems correctly.
- Inventory searches for full and partial IPv4 addresses are now treated as host queries.
- The Subnet Utilization percentage is now calculated correctly for sites with non-default masks.
- Asset tags set from the Inventory page are now additive and will be merged into existing tags.
- Asset tags are now shown on the overview page.
- Small improvements to the Task Search page.

1.12.4

2020-10-21

- The runZero Scanner and runZero Explorer now handle an even wider range of SNMP devices when polling the ARP cache.
- Fingerprint improvements.

1.12.3

2020-10-20

- The runZero Scanner and runZero Explorer now handle a wider range of SNMP devices when polling the ARP cache.
- Fingerprint improvements.

1.12.2

2020-10-17

- The runZero Scanner terminal UI has been cleaned up and generally improved.
- Fingerprint improvements.

1.12.1

2020-10-15

- The runZero Scanner and runZero Explorer now deduplicate overlapping target network ranges.
- Fingerprint improvements.

1.12.0

2020-10-13

- This is a roll-up of all 1.11.x point releases.

1.11.9

2020-10-11

- The runZero Scanner can now automatically update when run with the upgrade argument.
- The runZero Explorer can be installed via a [static MSI wrapper](#).
- Fingerprint improvements.

1.11.8

2020-10-04

- The scan engine is now more consistent through the addition of UDP retries and by pre-warming the ARP cache for each target group.
- Fingerprint improvements.

1.11.7

2020-09-28

- The runZero Explorer and runZero Scanner now include npcap version 1.0.
- Fingerprint improvements.

1.11.6

2020-09-23

- Users with the the Viewer role are now restricted to the Dashboard, Inventory, Explore, and Agents screens.

- Fingerprint improvements.

1.11.5

2020-09-21

- A regression in the license expiration tracking of the runZero Scanner has been resolved.
- Fingerprint improvements.

1.11.4

2020-09-18

- Fingerprint improvements.

1.11.3

2020-09-13

- Fingerprint improvements.

1.11.2

2020-09-11

- Organizations can now be converted to Projects from the settings page.
- The top-level organization navigation has been updated.

1.11.1

2020-09-09

- Accounts with a default organization role of “user”, but a per-organization role of “none” were inadvertently prevented from accessing certain features.
- An intermittent crash in the runZero Explorer was identified and resolved.

1.11.0

2020-09-09

- This is a roll-up of all 1.10.x point releases.

1.10.8

2020-09-08

- Projects are now available as temporary, self-deleting organizations. These are useful for one-off scans and exploring historical data.
- Web screenshots now try additional Chrome locations on the Windows platform.
- Over 10,000 new SNMP fingerprints have been added.
- Small bugfixes and cosmetic improvements.
- Additional Fingerprint improvements.

1.10.7

2020-09-01

- Web screenshots now retry on timeouts and choose the best quality image automatically.
- Web screenshots now use more concurrent Chrome processes on x86 systems, based on available memory.
- The runZero Explorer and runZero Scanner now track CPU and memory usage across the life of a scan.
- Fingerprint improvements.

1.10.6

2020-08-22

- The runZero Explorer, runZero Scanner, and runZero Console now compress raw scan by default. The `scan.rumble` output from the scanner has been renamed to `scan.rumble.gz`. The web console and API can import both compressed and uncompressed versions of this data. Existing scan data will be migrated to the compressed form automatically. This change helps with bandwidth usage by agents and speeds up large imports over the network.
- The Tasks view now links to the inventory search for each associated site.
- The status of agent-run scans is now updated more frequently.

1.10.5

2020-08-21

- The completed task list now shows the task runtime in the information column.
- The task views now also link to the inventory view of each site.
- Fingerprint improvements.

1.10.4

2020-08-17

- A regression in the runZero Scanner that prevented API uploads from succeeding has been resolved.
- Fingerprint improvements.

1.10.3

2020-08-14

- Small bug fixes and dependency updates across the platform.
- Fingerprint improvements.

1.10.2

2020-08-12

- Site exports and imports now include the registered subnets.
- Bulk asset updates are now possible by importing a modified CSV export from the Inventory screen.

1.10.1

2020-08-11

- A bug that made it difficult to query subnet tags with multiple subnets per tag has been resolved.
- Fingerprint improvements.

1.10.0

2020-08-04

- The console user interface received a light update around colors and styles.
- Event logs are now available in the console.
- Fingerprint improvements.

1.9.10

2020-08-03

- A bug that prevented Scan Tags from being shown in the Scan Configuration form has been resolved.
- The DNS and mDNS probes now always report the protocol, even for error responses.
- Fingerprint improvements.

1.9.9

2020-08-02

- Site scopes now automatically convert CIDR input into registered subnets.
- Scan tasks can now have their Scope and Excludes pinned to their associated site using the string “defaults”.
- Scan tasks can now be pinned to the default TCP service list using the string “defaults”.
- A bug that caused non-Windows SMB-enabled services to be identified as Windows has been resolved.
- A bug that caused SMB v1 to be reported incorrectly as been resolved.
- Fingerprint improvements.

1.9.8

2020-07-29

- A bug that prevented the Delete and Merge buttons in the Service Inventory toolbar from working has been resolved.
- A bug that led to the wrong title being shown in the FTP Service Attribute report has been resolved.
- Fingerprint improvements.

1.9.7

2020-07-22

- A bug that prevented the Inventory Import action from recognizing valid scan data has been resolved.
- The runZero Explorer and runZero Scanner are now much more reliable for lossy network environments.

1.9.6

2020-07-21

- The TCP probes have been updated to be less bursty. This resolves an issue where scans consisting of mostly HTTP services can timeout and lose valid responses.

- The TCP fingerprinter now handles unexpected termination more gracefully. This improves the reliability of AWS ELB scans and should help with a reliability across a range of services.

1.9.5

2020-07-20

- All paid plans now support Continuous recurring scans. These scans will run back-to-back and can simplify continuous monitoring. An agent running continuous scans will not run additional scans unless the Concurrency setting is increased beyond 1.
- Out-of-date agents will be upgraded prior to new scans being run. For the few agents where upgrades are impossible (read-only partitions, network filters, etc), this can delay each scheduled scan by up to five minutes.

1.9.4

2020-07-18

- The runZero Explorer and runZero Scanner now include a TLS CA root bundle to work around connectivity issues on older platforms. Bundle selection can be [controlled via the environment](#).
- Fingerprint improvements.

1.9.3

2020-07-16

- The web screenshot features now tries even harder to prevent orphaned Chrome.exe processes.
- The runZero Explorer now removes all agent-related files on uninstall.
- runZero Explorers can now be reassigned to other organizations.

1.9.2

2020-07-12

- The Export API now supports an optional `fields` parameter that determines which fields are exported in JSON/JSONL exports. The `fields` parameter is supported for Assets, Services, Wireless, and Sites.

1.9.1

2020-07-09

- A bug in the scan engine that could cause scans to hang when probing unresponsive SSH daemons has been resolved.
- A bug in the scan engine that could result in SMBv1 being reported erroneously on some NAS devices has been resolved.

1.9.0

2020-07-06

- A bug in the TFTP probe that could lead to missing results in some cases has been fixed.
- The SNMP probe now gathers the route table from many types of switches and routers.
- TCP SYN scans of non-local targets now try harder when there is congestion.
- Fingerprint matches that include a hardware version are now given priority.
- Fingerprint matches for SSH daemons now support more platforms.
- The permanent organization and permanent site can now be deleted/recreated.
- The Scan Configuration page now shows a notice when input validation failed.
- The Scan Configuration now shows SNMP parameters at the top of the form.
- The Network Bridges report now links all external IPs to an internet cloud.
- The Network Bridges report now uses subnet masks from Sites.
- The Subnet Utilization report now provides a Scan link for each network.
- The Subnet Utilization report now uses subnet masks from Sites.
- The Subnet Grid report now handles errors more gracefully.
- A bug that prevented some users from logging in has been resolved.
- Search queries are now slightly faster across assets and services.

1.8.14

2020-07-02

- Tasks are now searchable and sortable via the Search tab.
- A regression in numerical search queries has been resolved.

1.8.13

2020-07-01

- The Scan Configuration page now provides an estimated runtime through a confirmation dialog.
- Trial accounts are now longer limited to scanning a /16 and may now scan a full /8.
- The runZero Explorer now supports log configuration using the environment. See the [documentation](#) for details.
- The runZero Explorer and runZero Scanner now collect SSH pre-auth banners and host keys.
- Bogus service responses from Fortigate helpers on ports 80 and 8008 are now ignored.

- Fingerprint improvements.

1.8.12

2020-06-24

- The runZero Explorer and runZero Scanner now handle a wider range of ppp-based link types on Linux and macOS.
- Bogus service responses from Fortigate helpers on ports 21, 25, 80, 110, 143, 8008, 8010, and 8020 are now ignored.
- Fingerprint improvements.

1.8.11

2020-06-22

- Bogus service responses from Cisco H.323 helpers on port 1720 are now ignored.
- The runZero Explorer now stores additional diagnostics in the raw task data.
- Fingerprint improvements.

1.8.10

2020-06-21

- Bogus service responses from Fortigate SIP ALG helpers on ports 2000 and 5060 are now ignored.
- A regression in HTTP handling with redirects and TLS+HTTP headers has been resolved.
- Fingerprint improvements.

1.8.9

2020-06-20

- The runZero Explorer and runZero Scanner now handle malformed HTTP responses and redirects better.
- Fingerprint improvements.

1.8.8

2020-06-18

- ICMP Echo probes now record the IP header information from the response (useful for Ripple20/Treck detection).
- The Rumble favicon.ico MD5 fingerprint database has been contributed to the Recog project.

1.8.7

2020-06-16

- The runZero Explorer and runZero Scanner now support “cooked” interface types (ppp-based VPNs).
- The scan engine now extracts additional information from Netgear routers.
- Fingerprint coverage for Netgear routers has been improved.

1.8.6

2020-06-15

- Fingerprint improvements.

1.8.5

2020-06-15

- Fingerprint improvements.

1.8.4

2020-06-15

- Fingerprint improvements.

1.8.3

2020-06-14

- Asset Inventory and Search Inventory performance has been improved.
- The bundled npcap driver in the runZero Explorer and runZero Scanner for Windows has been upgraded to version 0.9994.
- Fingerprint improvements.

1.8.2

2020-06-09

- The runZero Scanner CSV output now includes populated UUID values.
- The runZero Scanner now creates a standalone bridges.json file for third-party processing.
- Fingerprint improvements.

1.8.1

2020-06-09

- A bug that could cause agent uninstalls to crash on BSD platforms has been resolved.

1.8.0

2020-06-09

- This release is a roll-up of the 1.7.x changes listed below.

1.7.13

2020-06-08

- Sites now support registered subnets. Assets can be queried via the associated Site subnet tags.
- Tags can be set with empty values and queried more precisely through the Inventory search.
- Asset fingerprinting via favicon.ico hashes has been implemented.
- The runZero Scanner now creates a standalone topology.json file for third-party processing.
- Assets now store the MAC-to-IP relationship in the hidden _macs.ipmap attribute.
- The runZero Explorer and runZero Scanner now support OpenBSD on x86 (64-bit).
- Fingerprint improvements.

1.7.12

2020-06-05

- This release fixes a bug that prevents the runZero Explorer from restarting automatically after an update on certain Debian-based distributions.
- Fingerprint improvements.

1.7.11

2020-06-04

- A reliability bug in the runZero Explorer and runZero Scanner for BSD-based platforms (macOS, FreeBSD, NetBSD, DragonFly BSD) has been resolved. This bug would manifest as missing scan results in the TCP SYN and ARP probe responses.
- The bundled npcap driver in the runZero Explorer and runZero Scanner for Windows has been upgraded to version 0.9992.

1.7.10

2020-06-02

- The runZero Explorer and runZero Scanner now support FreeBSD, NetBSD, and DragonFly BSD. FreeBSD and NetBSD support cover the following architectures: x86 (64-bit, 32-bit), ARM v5, ARM v6, and ARM v7. DragonFly BSD is supported on x86 (64-bit).
- The runZero Explorer and runZero Scanner now support additional Linux architectures. These include x86 (64-bit, 32-bit), ARM v5, ARM v6, ARM v7, ARM 64-bit (aarch64), MIPS (BE/LE), MIPS64 (BE/LE), PowerPC64 (LE), and s390x (IBM Z).
- The runZero Explorer now runs in standalone mode when no supported services backend is detected.
- The runZero Explorer now supports automatic updates in standalone mode on non-Windows platforms.
- The runZero Explorer binary now supports command-line flags (-h, -v, -l) and displays usage.

1.7.9

2020-05-27

- MAC address fingerprints are now live. The initial set includes fingerprints for devices manufactured by Amazon, Google, Honeywell, August, SimpliSafe, TRENDnet, FLIR, Microsoft, Belkin, Meross, LG, Logitech, Hunter, Lutron, Orbit, Arlo, Panasonic, Sony, Vizio, Chameleon, iRobot, SharkNinja, Netatmo, Nintendo, HP, Intel, Lenovo, Dell, and PC Engines. MAC fingerprints are used as a fallback when more precise fingerprinting is not available.
- Microsoft SQL Server versions obtained from the network are now mapped to specific releases and patch levels, enabling queries that look for end-of-life versions and missing patches.
- Chromecast devices now return additional service attributes, including information about the wireless network that they are connected to. Fingerprinting of older Chromecast models (Gen 1) has been improved. MAC addresses and additional IP addresses from the Chromecast web endpoint is now applied to the asset.
- MySQL and MariaDB version detection now also applies the appropriate OS fingerprint, if known.
- HTTP services that return JSON responses now camelCase the attribute names and support a wider range of data types. This impacts JSON-based HTTP interfaces such as Elasticsearch and Riak HTTP.
- OS and Hardware matching is more precise after adjustments were made to the weighting and priorities. The most precise and most confident fingerprint should always be chosen going forward.
- The confidence of the OS match is now reported as the asset-level `match.score` attribute. This may be renamed to `match.os.score` in the future as we accommodate more granular hardware weights.

- NTLMSSP-based OS matching now disqualifies systems that are obviously not Windows (BSD-based stacks, etc).
- Brother printers now use distinct hardware and firmware (OS) fingerprints. This should address cases where the firmware version overrode the hardware model by mistake.
- Release notes are now consolidated across the Platform, Agent, and Scanner.
- Versioning is now shared across the Platform, Agent, and Scanner.

Archived release notes

Prior to version 1.7.9, release notes and versions were split between the Platform, Explorer, and Scanner. You can find these archived release notes at the links below.

- [Archived Release Notes: runZero Platform](#)
- [Archived Release Notes: runZero Explorer](#)
- [Archived Release Notes: runZero Scanner](#)

Older platform release notes

Platform release notes

Starting with version 1.7.9 all release notes have been [consolidated](#) into one page.

v1.7.8

2020-05-23

- The pre-login style has been updated.
- The account registration flow has been updated to provide a smoother activation experience.

v1.7.7

2020-05-17

- Restricted user accounts may now be assigned roles within multiple organizations.

v1.7.6

2020-05-14

- Asset and Service attributes have been normalized. All keys are now camelCase and most service attributes are now prefixed by the protocol name.

v1.7.5

2020-05-13

- SMB v2 OS detection now applies more weight to Windows 10 versus Server 2019 in most situations.
- SMB v1 OS detection has been improved for non-Windows platforms.

v1.7.4

2020-05-08

- SSO now supports a default role of No Access. Accounts created through the SSO login process with this default role must be manually promoted to a usable account role.
- Users now support a No Access role for both All Orgs and specific organizations.

- A bug preventing access to the per-account Demo Organization has been resolved. Agents and Scanner downloads for the Demo Organization have been explicitly disabled. Any agents with the Demo Organization will automatically uninstall.
- All runZero Console confirmation dialogs now list the specific item they are confirming. This can prevent misclicks from removing the wrong item.
- A fingerprinting regression related to non-Windows machines running SMB services has been resolved. This would unintentionally classify some NAS devices as Microsoft Windows.

v1.7.3

2020-05-08

- [ToneLoc export support](#) in the Subnet Grid Report.

v1.7.2

2020-05-07

- Cisco HSRP MAC addresses are now ignored for the purposes of asset correlation.

v1.7.1

2020-05-06

- Agent diagnostic information is now visible from the agent details page. This requires agent version 1.7.1 or later.

v1.7.0

2020-05-04

- The Subnet Grid View report is now available (via the Subnet Report).
- The `ts`, `tls.notBeforeTS`, and `tls.notAfterTS` service fields can now be queried using time comparison operators.

v1.6.10

2020-05-03

- Scan schedules now show the day of the week for first, last, and next timestamps.
- A bug that obscured the validation error for invalid TCP ports has been resolved.
- Alert Rules can now be limited to specific Sites or All Sites.
- Dependency upgrade across the backend and frontend.

v1.6.7

2020-04-26

- A bug in the recurring scan limits of the Starter Edition has been resolved.

v1.6.6

2020-04-23

- A bug in the asset export that resulted in missing service details has been resolved.

v1.6.5

2020-04-22

- runZero now supports a free tier via the Starter Edition.

v1.6.4

2020-04-19

- The dashboard now shows statistics for live assets only and not all assets.
- The license view now ignores demo organizations and sites in the statistics.
- The task view now represents scans pending processing more accurately.

v1.6.3

2020-04-18

- The Asset and Service exports now include the `service.address`, `service.transport`, `service.port`, and `service.vhost` fields (if present) to make them more consistent with the runZero Scanner `assets.jsonl` exports.

v1.6.2

2020-04-12

- The [Analysis Reports](#) section has been added, including the new Domain Membership and Service Attribute reports.
- The [Query Library](#) has been added, including a set of predefined common queries.
- Inventory search queries can now be saved to the Query Library.
- Inventory search action buttons have been slimmed down.
- The [Explore](#) tab provides easy access the Analysis Reports and Query Library.

v1.6.1

2020-04-07

- Accounts that have been temporarily locked due to failed logins may be manually unlocked by superusers from the Team page.

v1.6.0

2020-04-06

- [Data retention](#) controls are now available on the Organization Settings page. These can be used to specify data retention periods and remove stale offline assets after a fixed period of time.
- Inventory, Site, and Organization search now support exact matches with the = prefix. An exact hostname match can be specified as name:=SERVER01.
- Inventory, Site, and Organization search now support matching empty fields using the = prefix. An empty hostname match can be specified as name:=
- Inventory searches now support the mac-vendor_count, macvendor_count, and vendor_count terms.
- Inventory Asset and Service search now support the haspublic, hasprivate, hasipv6, and haslinklocal terms.
- Inventory Asset and Service search now support % as a wildcard for address matching. The query address:%.0.1 will find all assets with an address ending in .0.1.
- A bug that prevented searches for less than one hostname has been fixed (name_count:0 and name_count:<1).
- Active and Failed scan tasks now support the Copy action. Copied scans scheduled in the past will default to the current time to start.
- Single Sign On users can now specify the default organization role of SSO-enrolled users.
- Account lockouts are now enforced after 10 failed password-based login attempts.

v1.5.6

2020-04-01

- The Screenshot Inventory now sorts by address instead of port number by default.

v1.5.5

2020-03-27

- The Asset Inventory now allows attributes to be queried even if they conflict with a keyword by using the prefix `_asset`.
- The Service Inventory now allows attributes to be queried even if they conflict with a keyword by using the prefix `_service`.

v1.5.4

2020-03-24

- Fingerprinting improvements for HTTP, X509, Telnet, SSH, FTP, and SIP.

v1.5.3

2020-03-22

- Scan tasks now have a configurable grace period. Setting this to zero or a negative value will result in the scheduling retrying indefinitely until an agent becomes available.

v1.5.2

2020-03-15

- Web services with multiple asset icons now show each icon correctly.
- A bug that prevented Sites from being imported has been resolved.
- A bug that prevented superusers from logging in directly when SSO was in required mode has been resolved.

v1.5.1

2020-03-09

- Asset icons are now lazy loaded as they scroll into the viewport.
- Asset icons URLs are now more cache friendly.
- Page meta tags have received small tweaks.

v1.5.0

2020-03-04

- The asset inventory now displays any HTTP and UPnP icons acquired during the scan.

- The asset inventory now displays an icon for assets with available screenshots.
- The inventory search no longer triggers an error on some grouped query corner cases.
- The inventory search now supports /32 masks for the net and cidr keywords.
- The inventory search now treats the type keyword as an exact match.
- The inventory search now supports filtering based on service counts.
- The inventory search now supports filtering based on the RTT and TTL fields.
- The inventory search for services now supports filtering based on port ranges.
- The asset detail page now describes upstream and downstream layer 2 links.
- The asset detail page now lists all acquired icons in the top summary.
- The asset detail page now has a Scan action available from the menu.
- The new scan page now supports applying a list of tags to all identified assets.
- The new scan page agent selection now defaults to the first site-specific agent.
- The tasks page now allows recurring scans to be paused and unpaused.
- The console now always shows the live agent count on the navigation menu.
- The console now applies styling to the print view.
- Hostnames consisting of leading digits are now ignored when correlating assets.
- A bug that resulted in the wrong timezone being used after SSO login has been resolved.
- Scans triggered via API now default to the site scope if no targets were specified.
- Scans triggered via API can now specify a list of tags to apply to all identified assets.
- Import tasks with the wrong data supplied format are rejected before task creation.

v1.4.5

2020-02-19

- Completed tasks are now always shown in the order they finished (versus were created).
- A Cisco IOS fallback fingerprint was leading to false positives and has been disabled.

v1.4.4

2020-02-15

- The License page now lists invoices and allows payment methods to be updated.
- The SMB Server GUID attribute is now used to correlate results to assets.
- The SNMP sysName and sysObjectID attributes are now used to unmatch assets that have changed IPs or were mistakenly matched through another attribute (shared bogus MAC addresses or similar).

v1.4.3

2020-02-13

- A bug that could treat active agents as offline has been resolved.
- Windows agents are now limited to a single concurrent scan to avoid resource contention issues.

v1.4.2

2020-02-09

- Changes in source TTLs on external segments no longer disqualify an asset match during correlation.
- Cosmetic updates to email templates and task displays.

v1.4.1

2020-02-05

- Updates to various javascript and backend dependencies.

v1.4.0

2020-02-04

- Version 1.4.0 is a rollup of post-1.3.0 point release work.

v1.3.5

2020-02-02

- Scans now support specifying the SNMP v3 Context for devices that require this (CatOS).
- Scans now support the “Fast” option for the ARP probe, which is on by default in AWS VPC environments.
- Agents now support concurrent scans, with the limit configurable per-agent.
- Asset correlation avoids accidentally grouping when a forged MAC address is present in routed protocol responses (NetBIOS, SNMP, etc).
- Change reports now ignore differences in reverse DNS, hostnames, domain names, and the top-level service count.
- Recurring, scheduled, and queued tasks without an available agent will mark the task as failed after 4 hours. Recurring tasks will attempt to run scans normally during the next scheduled period.
- Scanner binaries for OEM customers now behave according to the license agreement.
- Long task names are now truncated automatically in the task list.

v1.3.4

2020-01-29

- The failed tasks tab now shows a longer history by default.

v1.3.3

2020-01-29

- Minor dependency upgrades (javascript and backend).

v1.3.2

2020-01-28

- Search queries can now use parenthesis to group terms, for example: `os:linux AND (port:22 OR port:80) AND NOT port:3306`.

v1.3.1

2020-01-15

- The Export API now supports an [asset sync backend](#) with support for checkpointing based on `created_at` or `updated_at` fields.

v1.3.0

2020-01-07

- Version 1.3.0 is a rollup of post-1.2.0 point release work.

v1.2.9

2020-01-04

- The [Organization API](#) is now available.

v1.2.8

2019-12-29

- The Teams view is now split between members (access to all organizations) and restricted users (access to single organizations).
- Team member invitations now have a Reply-To header set to the inviting user.
- Team member invitations can now specify custom subject lines and messages.
- Active tasks are now always sorted by their start time, not their last update time.
- Tasks now always show the created by field (for tasks created or updated after v1.2.2).
- The search warnings field is now cleared after each new query.
- The search timestamp fields now default to less-than (<) if no operator is specified.

v1.2.7

2019-12-28

- The Active Tasks tab now offers a Clear Queue button to remove all queued tasks.

v1.2.6

2019-12-27

- Sites and Organizations can now be sorted and searched.

v1.2.5

2019-12-26

- The Beta discount is officially retired.
- End of year discount is valid through 2019

v1.2.4

2019-12-19

- Support for BACnet is now available in new scans.
- Leading and trailing spaces in Site and Organization names are now removed.

v1.2.3

2019-12-13

- Many new fingerprints were added for HTTP and SIP endpoints. These may trigger asset change notifications due to improved device and service recognition.
- Network scans of segments where a device responds to all ARP requests with the same MAC address will now be handled appropriately.
- The SIP protocol is now reported properly on UDP SIP responses.

v1.2.2

2019-12-10

- Tasks now track which user created or last updated them.

v1.2.1

2019-12-05

- SSO configurations with multiple IdP x509 certificates will now validate against any certificate in the list, not just the first one.

v1.2.0

2019-12-01

- Version 1.2.0 is a rollup of post-1.1.0 point release work.

v1.1.9

2019-12-01

- Initial support for the Wireless Inventory using the new wlan-list probe.

v1.1.8

2019-11-27

- An encoding issue in the Inventory display of certain columns has been resolved.

v1.1.7

2019-11-24

- The Scan Max Group Size is now set to a reasonable default when older scans are copied.
- The Copy action will no longer create names with more than one "Copy of" prefix by default.

v1.1.6

2019-11-19

- SNMP v3 credentials may now be specified in the scan configuration.

v1.1.5

2019-11-19

- The [Switch Topology](#) report is now available from the Inventory reports menu.
- Assets that respond to SNMP enumeration of the CAM/MAC table are now automatically classified as switches.

v1.1.4

2019-11-18

- A bug in the email notifier has been resolved that prevented delivery in rare circumstances.

v1.1.3

2019-11-14

- The Network Bridges report has been optimized for improved readability.

v1.1.2

2019-11-11

- The agent inventory search term has been added. This accepts a UUID as well as the agent name.

v1.1.1

2019-11-08

- Email notifications for expiring, deactivating, and unused trial accounts.

v1.1.0

2019-11-05

- Version 1.1.0 is a rollup of post-1.0.0 point release work.

v1.0.20

2019-11-04

- OS detection now prefers the most granular fingerprint in the case of multiple matches.
- Stopped tasks now show in the Failed list and not the Completed list.

- Scans now support the Max Group Size option to limit the number of concurrent scan targets.
- The site inventory search term now accepts a UUID as well as the site name.
- The task inventory search term now shows all assets that were last updated by a given task ID.
- The lowest_ttl field is no longer considered as a significant change in asset change notifications.
- A bug that led to some assets being stored with only IPv6 link-local addresses has been resolved.
- A bug that prevented OS version numbers from being tracked properly has been resolved.
- A bug that allowed hostname-based OS matching to override more reliable fingerprints has been resolved.

v1.0.19

2019-11-01

- Trial accounts and associated data are now automatically removed two weeks after the license period expires.

v1.0.18

2019-10-31

- The Network Bridges report now supports filtering using the same syntax as the Asset Inventory search.
- The Network Bridges report is now much faster, with the caveat that nodes are no longer draggable.
- The Network Bridges report now supports up to 1,000 multi-homed nodes at a time.
- The Completed Tasks view no longer shows tasks in a Cancelled or Error state (these are now listed under Failed Tasks).
- The Asset Inventory view now allows merging multiple assets into a single asset.

v1.0.17

2019-10-26

- Failed tasks can now be cleared via the Clear action button.

v1.0.16

2019-10-25

- The Export API now treats the Organization ID in the URL as optional.
- The Asset and Services inventory now support the id search term to find assets by ID.

v1.0.15

2019-10-25

- The Export API now offers a list of sites in three formats (CSV, JSON, JSONL).
- The change summary titles in the task view now link to their respective sections in the body.

v1.0.14

2019-10-24

- A small number of typos and cosmetic improvements resolved with the web console.

v1.0.13

2019-10-23

- The SYN probe now retries twice if no RST is received. This improves reliability at the cost of a small increase in scan times. This can be changed by the `syn-max-retries` advanced parameter for the SYN probe in the new scan configuration screen.

v1.0.8

2019-10-18

- Single Sign On support is now available for all users (SAML2). This is available the SS0 Settings button on the My Team page.

v1.0.7

2019-10-07

- A handful of small cosmetic issues have been resolved in the web interface.

v1.0.6

2019-10-08

- The asset correlation algorithm now accepts looser matches for TCP/IP stack fingerprints when matching an asset by IP address. This reduces the chance of asset churn when one or more services change between scan runs.

v1.0.5

2019-10-07

- Scan tasks now have Name and Description fields. These can be set during task creation and updated in Recurring and Scheduled tasks. These fields are shown in the main task view right after the task type.
- Existing scan tasks can now be copied to a new scan configuration (Copy action in the task list). This simplifies the process of running a one-off scan from an existing recurring scan definition.
- Less common scan options have been moved to the Advanced Scan Options section, which is now visible by default in the scan configuration page.
- The Team Manage menu now include an option to reset the security tokens of a team member (available to superusers only).
- The Sites listing now links the site name to the inventory search query. To edit the Site definition, use the Update Site action from the Modify menu.

v1.0.3

2019-10-06

- A race condition was fixed that could lead to some scan tasks failing with the error agent failed to queue task.
- Minor cosmetic improvements were applied to the console.

v1.0.2

2019-10-03

- A bug was fixed that caused some Scanner downloads to have a license expiration set in the past.
- The task scheduler is now more tolerant of temporary network errors when queueing scans.

v1.0.0

2019-10-01

- Rumble Network Discovery is [out of Beta with version 1.0.0!](#)

Older Explorer release notes

Explorer release notes (agents)

Starting with version 1.7.9 all release notes have been [consolidated](#) into one page.

v1.7.8

2020-05-23

- Fingerprint updates.

v1.7.7

2020-05-22

- Fingerprint updates.

v1.7.6

2020-05-14

- Corrects inconsistent use of the new service attributes when processing the dynamic MAC address filter.

v1.7.5

2020-05-14

- Asset and Service attributes have been normalized. All keys are now camelCase and most service attributes are now prefixed by the protocol name.

v1.7.4

2020-05-13

- Support has been improved for the following database protocols: Memcached (TCP), CouchDB, Cassandra, Redis, ElasticSearch, Riak (TCP/HTTP), MySQL, PostgreSQL, MongoDB, MSSQL, and Oracle.

v1.7.3

2020-05-07

- Cisco HSRP MAC addresses are now ignored for the purposes of asset correlation.
- Updated Ethernet fingerprints.

v1.7.2

2020-05-06

- A bug in the mDNS probe that could lead to a hung scan on certain platforms has been resolved.

v1.7.1

2020-05-06

- Diagnostic information is now returned in the “systeminfo” response.
- Updated Ethernet fingerprints.

v1.7.0

2020-05-04

- The bundled npcap driver has been upgraded to version 0.9991.
- The TLS probe now reports `tls.notBeforeTS` and `tls.notAfterTS` fields as unix timestamps.
- Updated Ethernet fingerprints.

v1.6.10

2020-05-03

- Updated Ethernet fingerprints.

v1.6.9

2020-05-01

- The runZero Agent will verify its own binary and exit on startup if corrupted.
- Updated Ethernet fingerprints.

v1.6.8

2020-04-23

- Updated Ethernet fingerprints.

v1.6.7

2020-04-23

- The scan engine can now identify TCP services on the scanning system across all platforms.

v1.6.6

2020-04-22

- An issue that could lead to the scan engine hanging with misbehaving HTTP services has been resolved.

v1.6.5

2020-04-22

- Updated Ethernet and BACnet fingerprints.

v1.6.4

2020-04-17

- Devices that relay mDNS from other networks (ex: Ubiquiti USG) are no longer associated with the relayed asset information.
- Additional Google Chrome paths are considered for screenshot collection. Snap packages of Chromium are no longer used.

v1.6.3

2020-04-14

- An issue that could lead to scans hanging while processing HTTP services has been resolved.

v1.6.2

2020-04-13

- The Linux agent now restarts properly on Ubuntu 14.04, CentOS 6.10, and other Upstart-based platforms.
- Fingerprint updates for Crestron, ELAN, MAC addresses, and BACnet.

v1.6.1

2020-04-08

- The Windows agent will now try harder to work around temporary issues while installing an update.
- The Windows agent will now clear any stale chrome.exe processes running as LocalSystem during the update process.

v1.6.0

2020-04-06

- Screenshots will now limit the number of concurrent Chrome processes based on core count, available RAM, and architecture.
- The bundled npcap build has been updated to version 0.9990.

v1.5.6

2020-04-04

- The RDP probe now collects the full NTLMSSP response for more platforms.
- The HTTP probe now collects information about web forms and their inputs.

v1.5.5

2020-03-27

- The SNMP probe no longer reports invalid MAC addresses found in ARP caches or MAC tables.

v1.5.4

2020-03-26

- The TCP probe now handles a wider variety of RDP responses.

v1.5.3

2020-03-26

- The SMB probe now reports subprotocols (smb1, smb2, and smb3) consistently.
- The SMB probe now collects hashing, encryption, and compression methods from SMBv3 servers.
- The SMB probe now reports the server-allocated Session ID for smb2 and smb3.
- The TCP probe now collects NTLM information from Remote Desktop endpoints and reports the protocol as rdp.
- The HTTP probe now collects additional information from VMware SOAP endpoints.

v1.5.2

2020-03-20

- A bug that could lead to the HTTP/2 probe stalling during TLS negotiation has been resolved.

v1.5.1

2020-03-14

- A full system disk no longer results in a deadlocked scan.
- Fingerprints have been updated for Ethernet MAC addresses, BACnet vendors, and Enterprise IDs.
- HTTP/1 probes now explicitly disable HTTP/2 upgrades even when advertised. HTTP/2 is handled separately.
- Generic protocol negotiation is no longer attempted on NDMP ports (10,000/30,000). A future release will support improved NDMP detection and negotiation.
- A potential deadlock in the runtime library has been resolved by reverting to an older runtime version.

v1.5.0

2020-03-04

- A NTP probe has been added that reports the clock skew compared to the scanning instance.
- A TFTP discovery probe has been added that requests a non-existent file and stores the response. The TFTP probe supports port ranges.
- An OpenVPN probe has been added that can detect remote instances across multiple ports.
- A dTLS discovery probe has been added that handles both bare dTLS and CAPWAP-encoded variants.

- Microsoft Remote Desktop Gateway instances are now fingerprinted through dTLS and HTTP, reporting the rdg.Transport service key.
- The protocol handlers for NATPMP, WS-Discovery, and UPnP Device XML now parse out specific subfields for easier matches and future fingerprinting efforts.
- The UPnP Device XML parser now triggers a request to download and report the device icon.
- The SYN scanner has been updated to improve reliability and report more accurate progress.
- The HTTP probe now identifies and reports web site icons as base64-encoded images along with their MD5 hashes.
- The HTTP probe now extracts the generator meta tag from HTML responses.
- The HTTP probe now extracts splunkd versions from HTML responses.
- The RPCBind probe now sends a null call to every UDP service and probes the NFS daemon directly.
- VMware ESXi detection has improved and will be used as a fallback in more cases.
- TCP protocol fingerprinting will retry more often on temporary network errors.
- Empty fields in the result structure within the JSON output are now omitted.
- Linux on ARM 64-bit (aarch64) is now a supported platform.
- Improved detection and early rejection of invalid CIDRs.

v1.4.5

2020-02-19

- The SMB probe now records the NTLMSSP response from a wider range of operating systems.
- The HTTP probe now stores the response to GET / and the response after any redirects are followed. Key names for the redirect responses are prefixed by last, such as last.http.code.
- The HTTP probe now handles compression and chunked transfer encoding properly, storing the normalized HTTP body.
- The HTTP probe now reports a banner consisting of the raw HTTP response.
- The HTTP probe now supports collecting environment data from LANDesk Management Agents.
- HTTP screenshots are now only collected when a 2XX HTTP response code is seen.
- HTTP screenshot processing is now more reliable.

v1.4.4

2020-02-16

- The SMB Server GUID attribute is now used to correlate results to assets.
- The SNMP sysName and sysObjectID attributes are now used to unmatch assets that have changed IPs or were mistakenly matched through another attribute (shared bogus MAC addresses or similar).
- Interfaces with no global unicast addresses (including RFC1918) are no longer considered by the ARP and SYN scanners.

- VLAN-tagged frames are now ignored by the SYN scanner resolving an issue where packets could be sent on the wrong interface by mistake.
- SYN scans now have a mandatory delay between retry attempts, which improves reliability and decreases change churn when small network ranges are scanned.

v1.4.3

2020-02-13

- Concurrent scans now use less resources and provide more accurate results.
- A bug that caused some HTTP requests to be sent with an empty Host header has been fixed.
- Version 1.4.1 and 1.4.2 were internal test releases and not deployed.

v1.4.0

2020-02-04

- Version 1.4.0 is a rollup of post-1.3.0 point release work.

v1.3.2

2020-02-02

- Agents now support concurrent scans, configured via the cloud console.
- Agents now write log files and rotate these log files automatically.
- Agents now scan faster on local segments within AWS VPCs.
- Agents now support SNMP v3 Context values in the scan configuration.
- Agents now try harder to recover from error cases during installation.

v1.3.1

2020-01-26

- A race condition was resolved that could leave abandoned chrome.exe processes after a scan.

v1.3.0

2020-01-07

- Version 1.3.0 is a rollup of post-1.2.0 point release work.

v1.2.3

2019-12-19

- The agent now tracks virtual hosts better for reported HTTP and TLS services.

v1.2.2

2019-12-19

- Support for the BACnet protocol has been added.

v1.2.1

2019-12-13

- The protocol detection engine has received a number of small improvements (mongod recognition among others).

v1.2.0

2019-12-01

- Version 1.2.0 is a rollup of post-1.1.0 point release work.

v1.1.15

2019-12-01

- Improved normalization of wireless network fields for the wlan-list probe.

v1.1.14

2019-11-27

- Additional bug fixes for SNMP processing.
- Initial support for the wlan-list probe module.

v1.1.13

2019-11-26

- Better support for truncated HTTP responses.

v1.1.12

2019-11-24

- Invalid SNMP responses are now handled more efficiently.

v1.1.11

2019-11-24

- A bug that could lead to memory exhaustion when Max Group Size was set to zero has been resolved.

v1.1.10

2019-11-23

- A bug in the SNMP probe that could result in the scan missing the last round of enumeration results has been fixed.
- Scans that result in excessive memory usage will now automatically upload a heap profile to enable support diagnostics.

v1.1.9

2019-11-22

- Improved error handling and logging, minor performance increase.

v1.1.8

2019-11-22

- Reduced memory usage on scan reply deduplication.

v1.1.7

2019-11-19

- Cisco-specific MIBs are now enumerated for CAM/MAC table enumeration.
- SNMP v2 is now queried two ways by the SNMP probe to improve device compatibility.
- SNMP v3 authenticated enumeration is now available.

v1.1.6

2019-11-19

- A number of small bugs in the SNMP probe have been resolved.

v1.1.5

2019-11-18

- The SNMP probe will now try to obtain the full interface and MAC address list from each asset.

v1.1.4

2019-11-14

- Miscellaneous fingerprinting improvements.

v1.1.3

2019-11-07

- Improved protocol detection for the Click Modular Router daemon.

v1.1.2

2019-11-05

- A bug that could lead to a agent being stuck in “stopping” status when a scan is stopped has been resolved.

v1.0.15

2019-11-04

- Scans now support the Max Group Size option to limit the number of concurrent scan targets.
- The MAC address prefix database and various other dependencies and fingerprints were updated.

v1.0.9

2019-10-24

- The SYN probe now sends retries using the same source port and sequence number to minimize duplicate responses. MAC address fingerprints have been updated.

v1.0.8

2019-10-23

- The SYN probe now retries twice if no RST is received. This improves reliability at the cost of a small increase in scan times. This can be changed by the `syn-max-retries` parameter in the console.

v1.0.7

2019-10-21

- Scanner performance is no longer reduced when the ARP probe is enabled for non-local scan targets.

v1.0.3

2019-10-06

- The macOS agent now supports additional interface types (loopback and tunnel adapters).

v1.0.2

2019-10-02

- A race condition was fixed that could cause an agent to crash mid-scan. Affected scans would have an error status of `task lost` to `agent restart`.

v1.0.1

2019-10-01

- Rumble Network Discovery is [out of Beta with version 1.0.0!](#)

Older scanner release notes

Scanner release notes

Starting with version 1.7.9 all release notes have been [consolidated](#) into one page.

v1.7.8

2020-05-23

- Fingerprint updates.

v1.7.7

2020-05-22

- Fingerprint updates.

v1.7.6

2020-05-14

- Corrects inconsistent use of the new service attributes when processing the dynamic MAC address filter.

v1.7.5

2020-05-14

- Asset and Service attributes have been normalized. All keys are now camelCase and most service attributes are now prefixed by the protocol name.

v1.7.4

2020-05-13

- Support has been improved for the following database protocols: Memcached (TCP), CouchDB, Cassandra, Redis, ElasticSearch, Riak (TCP/HTTP), MySQL, PostgreSQL, MongoDB, MSSQL, and Oracle.

v1.7.3

2020-05-07

- Cisco HSRP MAC addresses are now ignored for the purposes of asset correlation.
- Updated Ethernet fingerprints.

v1.7.2

2020-05-06

- A bug in the mDNS probe that could lead to a hung scan on certain platforms has been resolved.

v1.7.1

2020-05-06

- Updated Ethernet fingerprints.

v1.7.0

2020-05-04

- The bundled npcap driver has been upgraded to version 0.9991.
- The TLS probe now reports `tls.notBeforeTS` and `tls.notAfterTS` fields as unix timestamps.
- Updated Ethernet fingerprints.

v1.6.10

2020-05-03

- Updated Ethernet fingerprints.

v1.6.9

2020-05-01

- Support for Recog development with `--fingerprints` and `--fingerprints-debug` options.
- The Switch Topology and Network Bridges reports are now available for the CLI scanner.
- Updated Ethernet fingerprints.

v1.6.8

2020-04-23

- Updated Ethernet fingerprints.

v1.6.7

2020-04-23

- The scan engine can now identify TCP services on the scanning system across all platforms.

v1.6.6

2020-04-22

- An issue that could lead to the scan engine hanging with misbehaving HTTP services has been resolved.

v1.6.5

2020-04-22

- Support for the runZero Starter Edition.
- Updated Ethernet and BACnet fingerprints.

v1.6.4

2020-04-17

- Devices that relay mDNS from other networks (ex: Ubiquiti USG) are no longer associated with the relayed asset information.
- Additional Google Chrome paths are considered for screenshot collection. Snap packages of Chromium are no longer used.

v1.6.3

2020-04-14

- An issue that could lead to scans hanging while processing HTTP services has been resolved.

v1.6.2

2020-04-13

- Fingerprint updates for Crestron, ELAN, MAC addresses, and BACnet.

v1.6.1

2020-04-08

- The MAC address database fingerprints have been updated.

v1.6.0

2020-04-06

- Screenshots will now limit the number of concurrent Chrome processes based on core count, available RAM, and architecture.
- The bundled npcap build has been updated to version 0.9990.

v1.5.6

2020-04-04

- The RDP probe now collects the full NTLMSSP response for more platforms.
- The HTTP probe now collects information about web forms and their inputs.

v1.5.5

2020-03-27

- The SNMP probe no longer reports invalid MAC addresses found in ARP caches or MAC tables.

v1.5.4

2020-03-26

- The TCP probe now handles a wider variety of RDP responses.

v1.5.3

2020-03-26

- The SMB probe now reports subprotocols (smb1, smb2, and smb3) consistently.
- The SMB probe now collects hashing, encryption, and compression methods from SMBv3 servers.
- The SMB probe now reports the server-allocated Session ID for smb2 and smb3.
- The TCP probe now collects NTLM information from Remote Desktop endpoints and reports the protocol as rdp.
- The HTTP probe now collects additional information from VMware SOAP endpoints.
- A race condition in --nopcap mode that led to inconsistent results has been resolved.

v1.5.2

2020-03-20

- A bug that could lead to the HTTP/2 probe stalling during TLS negotiation has been resolved.

v1.5.1

2020-03-14

- Fingerprints have been updated for Ethernet MAC addresses, BACnet vendors, and Enterprise IDs.
- HTTP/1 probes now explicitly disable HTTP/2 upgrades even when advertised. HTTP/2 is handled separately.
- Generic protocol negotiation is no longer attempted on NDMP ports (10,000/30,000). A future release will support improved NDMP detection and negotiation.
- A potential deadlock in the runtime library has been resolved by reverting to an older runtime version.

v1.5.0

2020-03-04

- A NTP probe has been added that reports the clock skew compared to the scanning instance.
- A TFTP discovery probe has been added that requests a non-existent file and stores the response. The TFTP probe supports port ranges.
- An OpenVPN probe has been added that can detect remote instances across multiple ports.
- A dTLS discovery probe has been added that handles both bare dTLS and CAPWAP-encoded variants.
- Microsoft Remote Desktop Gateway instances are now fingerprinted through dTLS and HTTP, reporting the rdg.Transport service key.
- The protocol handlers for NATPMP, WS-Discovery, and UPnP Device XML now parse out specific subfields for easier matches and future fingerprinting efforts.

- The UPnP Device XML parser now triggers a request to download and report the device icon.
- The SYN scanner has been updated to improve reliability and report more accurate progress.
- The HTTP probe now identifies and reports web site icons as base64-encoded images along with their MD5 hashes.
- The HTTP probe now extracts the generator meta tag from HTML responses.
- The HTTP probe now extracts splunkd versions from HTML responses.
- The RPCBind probe now sends a null call to every UDP service and probes the NFS daemon directly.
- VMware ESXi detection has improved and will be used as a fallback in more cases.
- TCP protocol fingerprinting will retry more often on temporary network errors.
- Empty fields in the `result` structure within the JSON output are now omitted.
- Linux on ARM 64-bit (aarch64) is now a supported platform.
- Improved detection and early rejection of invalid CIDRs.

v1.4.5

2020-02-19

- The SMB probe now records the NTLMSSP response from a wider range of operating systems.
- The HTTP probe now stores the response to GET / and the response after any redirects are followed. Key names for the redirect responses are prefixed by `last`, such as `last.http.code`.
- The HTTP probe now handles compression and chunked transfer encoding properly, storing the normalized HTTP body.
- The HTTP probe now reports a banner consisting of the raw HTTP response.
- The HTTP probe now supports collecting environment data from LANDesk Management Agents.
- HTTP screenshots are now only collected when a 2XX HTTP response code is seen.
- HTTP screenshot processing is now more reliable.

v1.4.4

2020-02-16

- The SMB Server GUID attribute is now used to correlate results to assets.
- The SNMP `sysName` and `sysObjectID` attributes are now used to unmatch assets that have changed IPs or were mistakenly matched through another attribute (shared bogus MAC addresses or similar).
- Interfaces with no global unicast addresses (including RFC1918) are no longer considered by the ARP and SYN scanners.
- VLAN-tagged frames are now ignored by the SYN scanner resolving an issue where packets could be sent on the wrong interface by mistake.
- SYN scans now have a mandatory delay between retry attempts, which improves reliability and decreases change churn when small network ranges are scanned.

v1.4.3

2020-02-13

- A bug that caused some HTTP requests to be sent with an empty Host header has been fixed.

v1.4.2

2020-02-05

- Version 1.4.2 improves fingerprinting and type classification of Windows operating systems.

v1.4.1

2020-02-04

- Version 1.4.1 resolves a regression where `--input-targets` was being ignored.

v1.4.0

2020-02-04

- Version 1.4.0 is a rollup of post-1.3.0 point release work.

v1.3.2

2020-02-02

- Support for the new `--arp-fast` option to send ARP probes at the configured `--rate`, without additional delays. This option is on by default in AWS VPC environments.
- Support for the new `--snmp-v3-context` option to specify the SNMP v3 Context for queries.

v1.3.1

2020-01-26

- Support for the `--baseline` option (`-b`) to load a previous `assets.jsonl` and use existing asset IDs in the new `assets.jsonl` output.
- Support for multiple `--import` files. This allows multiple scans to be recombined into a single output.

- Support for the --upload option, which uses --api-key, --upload-site, and other API-related options to upload scan data to the runZero Console.
- A race condition was resolved that could leave abandoned chrome.exe processes after a scan.
- Prevent use or display of a blank virtual host for HTTP servers

v1.3.0

2020-01-07

- Version 1.3.0 is a rollup of post-1.2.0 point release work.

v1.2.3

2019-12-19

- The CLI scanner now generates a protocols.csv file with a simplified list of services and their URLs.

v1.2.2

2019-12-19

- Support for the BACnet protocol has been added.

v1.2.1

2019-12-13

- Many new fingerprints were added for HTTP and SIP endpoints.
- Asset correlation now occurs after the scan completes, not as it runs, for more consistent results.
- The protocol detection engine has received a number of small improvements (mongod recognition among others).
- Network scans of segments where a device responds to all ARP requests with the same MAC address will now be handled appropriately.

v1.2.0

2019-12-01

- Version 1.2.0 is a rollup of post-1.1.0 point release work.

v1.1.15

2019-12-01

- Automatic generation of wireless.jsonl and wireless.csv reports when the wlan-list probe returns results.
- Improved normalization of wireless network fields for the wlan-list probe.

v1.1.14

2019-11-27

- Additional bug fixes for SNMP processing.
- Initial support for the wlan-list probe module.

v1.1.13

2019-11-26

- Better support for truncated HTTP responses.

v1.1.12

2019-11-24

- Invalid SNMP responses are now handled more efficiently.

v1.1.11

2019-11-24

- A bug that could lead to memory exhaustion when Max Group Size was set to zero has been resolved.

v1.1.10

2019-11-23

- A bug in the SNMP probe that could result in the scan missing the last round of enumeration results has been fixed.

v1.1.9

2019-11-22

- Improved error handling and logging, minor performance increase.

v1.1.8

2019-11-22

- Reduced memory usage on scan reply deduplication.

v1.1.7

2019-11-19

- Cisco-specific MIBs are now enumerated for CAM/MAC table enumeration.
- SNMP v2 is now queried two ways by the SNMP probe to improve device compatibility.
- SNMP v3 authenticated enumeration is now available.

v1.1.6

2019-11-19

- A number of small bugs in the SNMP probe have been resolved.

v1.1.5

2019-11-18

- The SNMP probe will now try to obtain the full interface and MAC address list from each asset.

v1.1.4

2019-11-14

- Network topology links are now reported in the `_links.*` asset attributes.
- Miscellaneous fingerprinting improvements.

v1.1.3

2019-11-07

- Additional SSH fingerprints, covering Debian, Ubuntu, Raspbian, and FreeBSD.
- Improved protocol detection for the Click Modular Router daemon.

v1.1.2

2019-11-05

- A bug that could lead to a scan engine hang when the scan is interrupted has been resolved.

v1.0.15

2019-11-04

- The scanner now supports the `--max-group-size` option to limit the host working set (default is 4096).
- A race condition in the `--text` mode output that could rarely lead to a crash has been resolved.
- The MAC address prefix database and various other dependencies and fingerprints were updated.

v1.0.10

2019-10-25

- The scanner now supports the `--overwrite` option to allow reuse of an existing output directory

v1.0.9

2019-10-24

- The SYN probe now sends retries using the same source port and sequence number to minimize duplicate responses. MAC address fingerprints have been updated.

v1.0.8

2019-10-23

- The SYN probe now retries twice if no RST is received. This improves reliability at the cost of a small increase in scan times. This can be changed by the `-syn-max-retries` parameter.

v1.0.7

2019-10-21

- Scanner performance is no longer reduced when the ARP probe is enabled for non-local scan targets.

v1.0.3

2019-10-06

- The macOS scanner now supports additional interface types (loopback and tunnel adapters).
- The macOS scanner no longer prints warnings about unusable interfaces unless the verbose flag is set.
- A bug was fixed that led to the wrong result count being reported when using the --text interface of the scanner.

v1.0.2

2019-10-02

- A race condition was fixed that could the scanner to crash mid-scan.
- The macOS scanner no longer crashes if an unusable interface is found and the scanner is not run as root.

v1.0.1

2019-10-01

- Rumble Network Discovery is [out of Beta with version 1.0.0!](#)

Index

2

2FA: 84

A

Account API: 85 399
 Action: 300
 Activating your account: 3
 Active scanning: 106
 Adaptive Cards: 303
 Add multiple users: 89
 Address ranges: 111
 Administrators: 82
 Alert messages: 302
 Alerts: 7
 Alerts: 299 299 380
 Alive: 286
 Amazon Web Services: 146
 Analysis reports: 268
 Annotators: 83
 API: 399
 API client credentials: 399
 API key / token: 399
 ARP: 286
 Asset data: 7
 Asset risk report: 289
 Asset route pathing: 298
 Asset route pathing report: 383
 Assets: 282 285
 Assets back online: 264
 Assets changed: 264
 Assets ignored: 264
 Assets marked offline: 264
 Assets unchanged: 264
 Assets updated by task: 264
 Atlassian Insight: 228
 Authenticated SNMPv3: 119
 Authentication passphrase: 119
 Authentication protocol: 119
 AuthNoPriv: 119
 AuthPriv: 119
 Automated installers: 36
 Automatic queries: 382
 AWS: 146 153
 Azure AD: 99 155

B

Billing:	82
Billing information:	104
Binary signature:	34
Block Kit:	302
Bridge report:	297
Bulk asset update:	283
Bulk import users:	89

C

Caret:	313
CDE:	297
CEE:	64
Censys Search:	160
Censys Universal Internet Dataset:	161
Change report:	277
Channel:	300
Channels:	299
CIDR allow list:	119
Cisco ASA:	450
CMDB:	234
Colors:	384
Community strings:	118
Comparison:	388
Connect:	283
Connector task:	144
Container:	25
Coverage reports:	386
Credentials:	140
Criticality:	288
CrowdStrike:	164
CSV:	396
Curly brackets:	303

D

Dashboard:	7
Dashboard:	279
Data formats:	269
Data retention:	78
Debug SNMPv3:	121
Defender:	176
Department:	47
Deploy:	25
Deployment plan:	39
Detected by:	286
Device type:	286
DHCP:	282
Diagnostics collection:	76

Differences:	388
Discovery scan:	110
DNS:	285
Docker:	25
Domain membership:	268
Download:	25
Duplicate assets:	49

E

EC2:	153
ElasticSearch:	64
Email:	299
Embedded devices:	25
ESSID:	356
ESXi:	222
Event records:	4
Events:	299
Example queries:	324
Explorer:	25 110 131
Export API:	399
Export asset data:	396
External Asset Report:	394
External network discovery:	25
External users:	91

F

Falcon:	168
Filter:	322
Fingerprint results:	282
Firewall:	450
First scans:	116
Flat networks:	49
FQL:	168

G

GCP:	170
Global credential:	140
Goal tracking:	294
Goals:	294
Google Chrome:	452
Google Cloud Platform:	170
Google Workspace:	173
Grace period:	5
Graphviz:	384
Group mapping search:	375
Group search:	373
Groups changed:	265
Groups unchanged:	265
Groups updated by task:	265

H

Hardware:	286
High-availability:	71
Historical changes:	388
Hop:	383
Host ping:	115
Hosted zone:	110
Hostnames:	111
Hostnames:	282 285 285
HP iLO:	397

I

IAM:	153
ICMP:	286
ICMP ping:	115
Insights:	7
Insights:	280
InsightVM:	199 200
Installation:	3
IntegrationHub ETL:	234
Inventory:	282
Inventory views:	7
Invoices:	105
IoT:	455
IP addresses:	282 285
IPv4:	386
IRE:	234
ITOM:	235

J

Jira Service Management:	228
Journal:	63
Journalctl:	64
JSON:	396
JSONL:	269

K

Knowledgebase:	450
----------------	-----

L

LDAP:	180
Licensing:	104
Logging:	63

M

MAC addresses:	282 285
----------------	---------

Magnifying glass:	387
Mattermost:	299
Max group size:	113
Max group size:	451
MFA:	4
Microsoft 365:	176
Microsoft 365 Defender:	176
Microsoft Azure:	183
Microsoft Intune:	188
Microsoft Teams:	303
MID Server:	234
Miradore:	192
Missing subnets:	386
MSI package:	36
Multi-factor authentication:	84
Mustache:	303

N

NameID:	93
NAT:	451
Nessus:	218
Nessus Professional:	215
Network bridges:	267
Network examples:	42
Network paths:	383
Network segmentation:	297
Newly discovered assets:	264
Newly discovered groups:	264
Newly discovered users:	264
Nexpose:	199 203
No access:	83
NoAuthNoPriv:	118
Npcap:	133

O

Offline:	285
Offline asset expiration:	4
Offline installation:	69
Okta:	101
On-premises:	51
Operating system:	282 286
Organization:	274
Organization API:	399
Organization Overview Report:	392
Organization search:	377
Organizations:	110
OT:	455
Outlier score:	287
Outliers:	267

Ownership: 290

P

Packet rate: 112
 Panther: 233
 Passive: 106
 Password rules: 24
 PCI DSS: 297
 Planning: 3
 Ports: 112 454
 PostgreSQL: 70
 Prebuilt searches: 382
 Prescan modes: 115
 Primary address: 285
 Privacy passphrase: 119
 Privacy protocol: 119
 Private IP addresses: 386
 Probes: 112 124
 Profile images: 24
 Project: 48
 Project assets: 104
 Protocols: 453

Q

Qualys VMDR: 195
 Queries: 321
 Queries: 382
 Query language: 382
 Query library: 382

R

Rapid7 InsightVM: 200
 Rapid7 Nexpose: 203
 Raspberry Pi: 37
 Recent assets: 104
 Recurring scan: 293
 Red border: 387
 Reinstall: 132
 Remote logging: 64
 Report: 388
 Reports: 7
 REST: 399 400
 RFC 1918: 116 138 267
 RFC1918: 49 386
 Risk: 288
 Roles: 80
 Router: 450
 RPMs: 70
 RTT: 286

Rules:	299 299
Rules:	302
Rules submenu:	302

S

SAML:	96
Sample Queries:	265
Sample rate:	115
Scan completed:	308
Scan data:	7
Scan data expiration:	4
Scan probe:	144
Scan result:	272
Scan scope:	285
Scan speed:	5
Scan speed:	451
Scan stats:	271
Scan template:	135
Scanner:	141 401
Scans:	292
Scheduled scans:	5
Scope:	111
Screenshot:	131
Screenshots:	27 7 9
Search:	282
Search assets:	331
Search credentials:	377
Search explorers:	369
Search groups:	361
Search queries:	379
Search query rule:	308
Search reports:	368
Search scan templates:	363
Search services:	346
Search software:	351
Search syntax:	322
Search users:	359
Search vulnerabilities:	351
Search WiFi networks:	356
Searching tasks:	364
Secondary addresses:	285
SecurityGate.io:	233
Segmentation:	297
Self-hosted:	51 141
Sendgrid:	299
SentinelOne:	205
Service Graph connector:	234
ServiceNow:	234
Services:	282 7
Shodan:	208

Side-by-side:	388
Sign up:	21 24
SIP:	450
Site:	274 282 285
Site search:	376
Sites:	110
Slack:	299 302 320
SNMP:	118
SNMP credentials:	119
SNMP v2 Communities:	119
Software:	7
Software groups:	350
Splunk:	237 237
SSID:	356
SSO:	4 84 93 101
SSO group mapping:	96
Stale asset expiration:	4
Stateful firewalls:	110
Statistics:	387
Subnet analysis:	387
Subnet sampling:	115
Subnet size:	115
Subnet utilization:	138
Subnet utilization:	267
Subnets:	282 386
Sumo Logic:	238 239 245
Superuser:	82
Switch topology:	139
Switch topology:	266 267
Syslog:	64
System events:	380
System requirements:	3

T

Task details:	264
Tasks:	292
Tasks overview page:	292
TCP ping:	115
TCP ports:	124
TCP services:	286
Team member:	80
Templates:	302
Tenable Nessus:	211 218
Tenable Nessus Professional:	215
Tenable Security Center:	211
Tenable Vulnerability Management:	211
Tenable.io:	211 212
Tenable.sc:	211 220
Thinkst Canary:	262
Time To Live:	113

Tines:	252
Topology:	266
Traceroute:	383
Traffic sampling:	106
Two-factor authentication:	84
Type of Service:	113 113

U

UAC:	26
UDP ping:	116
UDP ports:	124
UDP services:	286
Unmapped:	266
Unmapped MACs:	139
Upgrade:	132
User groups:	86
User search:	371
Users changed:	264
Users unchanged:	264
Users updated by task:	264
UUID:	274

V

VCenter:	222
Verifier:	34
Viewers:	83
Virtual machine:	222
Virtual machines:	451
Virtualized system:	25
VMware:	25 222 451
VPN gateways:	110
Vulnerabilities:	7

W

Web screenshots:	3
Webhook:	299
Webhooks:	302
Widget library:	279
WiFi:	284
Wireless networks:	7
Wiz:	224