



## Red Hat Insights 2023

# Assessing and Monitoring Security Vulnerabilities on RHEL Systems

Understanding your Environmental Exposure to Potential Security Threats



# Red Hat Insights 2023 Assessing and Monitoring Security Vulnerabilities on RHEL Systems

---

Understanding your Environmental Exposure to Potential Security Threats

## Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

Use the vulnerability service to assess and monitor the status of security vulnerabilities on your RHEL systems, understand the level of exposure of your infrastructure, and plan a course of action. Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message .

## Table of Contents

<b>CHAPTER 1. OVERVIEW OF THE INSIGHTS FOR RHEL VULNERABILITY SERVICE</b> .....	<b>3</b>
1.1. HOW THE VULNERABILITY SERVICE WORKS	3
<b>CHAPTER 2. USER ACCESS FOR VULNERABILITY SERVICE USERS</b> .....	<b>4</b>
2.1. VULNERABILITY ADMINISTRATOR ROLE	4
2.2. VULNERABILITY VIEWER ROLE	4
<b>CHAPTER 3. COMMON VULNERABILITIES AND EXPOSURES (CVEs)</b> .....	<b>6</b>
3.1. RED HAT SECURITY ADVISORIES (RHSAS)	6
3.2. SECURITY RULES	7
3.2.1. Identifying security rules in the Insights for RHEL dashboard	8
3.3. KNOWN EXPLOITS	9
3.4. COMMON VULNERABILITIES AND EXPOSURES PROVIDE DEEP THREAT INTELLIGENCE WITH TRIAGE FEATURE	9
3.4.1. Identifying known-exploit CVEs in the Red Hat Insights for RHEL dashboard	10
<b>CHAPTER 4. REFINING VULNERABILITY-SERVICE RESULTS</b> .....	<b>11</b>
4.1. CVE-LIST AND SYSTEMS-LIST FILTERS	11
4.2. FILTERING SECURITY-RULE CVEs	13
4.3. FILTERING KNOWN-EXPLOIT CVEs	13
4.4. FILTERING LISTS OF SYSTEMS EXPOSED TO SECURITY RULES	14
4.5. INSIGHTS FOR RHEL GROUP FILTERS	14
4.5.1. Filtering Dashboard, CVEs, and Systems lists by group	14
4.6. DEFINING A BUSINESS RISK FOR A CVE	15
4.6.1. Setting a business risk for a single CVE	15
4.6.2. Setting a business risk for multiple CVEs	16
4.7. EXCLUDING SYSTEMS FROM VULNERABILITY SERVICE ANALYSIS	16
4.8. SHOWING PREVIOUSLY EXCLUDED SYSTEMS	17
4.9. RESUMING VULNERABILITY ANALYSIS FOR A SYSTEM	17
4.10. CVE STATUS	18
4.10.1. Setting a status for a CVE on all affected systems	18
4.10.2. Setting a status for a CVE and system pair	18
4.11. USING THE SEARCH BOX	19
4.12. SORTING CVE LIST DATA	19
<b>CHAPTER 5. SYSTEM TAGS AND GROUPS</b> .....	<b>21</b>
5.1. SAP WORKLOADS	21
5.2. SATELLITE HOST GROUPS	21
5.3. MICROSOFT SQL SERVER WORKLOADS	21
5.3.1. Setting up SQL Server assessments	22
5.3.1.1. Setting up the SQL Assessment on a timer	23
5.4. CUSTOM SYSTEM TAGGING	24
5.4.1. Tag structure	24
5.4.2. Creating a tags.yaml file and adding a custom group	24
5.4.3. Editing tags.yaml to add or change tags	25
5.4.4. Using predefined system tags to get more accurate Red Hat Insights advisor service recommendations and enhanced security	26
5.4.5. Configuring predefined tags	27
<b>CHAPTER 6. REFERENCES</b> .....	<b>30</b>
6.1. REFERENCE MATERIALS	30
<b>PROVIDING FEEDBACK ON RED HAT DOCUMENTATION</b> .....	<b>31</b>



# CHAPTER 1. OVERVIEW OF THE INSIGHTS FOR RHEL VULNERABILITY SERVICE

The vulnerability service enables quick assessment and comprehensive monitoring of the exposure of your RHEL infrastructure to Common Vulnerabilities and Exposures (CVEs) so you can better understand your most critical issues and systems and effectively manage remediations.

With your data uploaded to the vulnerability service, you can filter and sort groups of systems and CVEs to refine and optimize your views. You can also add context to individual CVEs when they pose an extraordinary risk to systems. After gaining an understanding of your risk exposure, report on the status of the CVEs to appropriate stakeholders, then create Ansible Playbooks to remediate issues to secure your organization.

## Prerequisites

The vulnerability service is available for all supported versions of RHEL 6, 7, 8 and 9. The following conditions must be met before you can use the vulnerability service:

- **Each system has the Insights client installed and registered to the Insights for Red Hat Enterprise Linux application.** Follow the [Red Hat Insights for Red Hat Enterprise Linux, Get Started instructions](#) to install the client and register your system(s).
- **The vulnerability service is fully supported for RHEL systems managed by Red Hat Subscription Management (RHSM) and Satellite 6 and later.** Using any other means to obtain package updates, other than Satellite 6 with RHSM or RHSM registered with [subscription.redhat.com](#) (Customer Portal), can lead to misleading results.
- **Vulnerability service remediations are not fully supported and may not work properly on Satellite 5 and Spacewalk-hosted RHEL systems.**
- **Some features require special privileges provided by your organization administrator.** Specifically, the ability to view Red Hat Security Advisories (RHSAs) associated with certain CVEs and systems, and to view and patch those vulnerabilities in the Red Hat Insights for Red Hat Enterprise Linux patch service, requires permissions granted through user access.

## Additional resources

- [Remediating Security Exposures using the Vulnerability Service and Ansible Playbooks](#)
- [Generating Vulnerability Service Reports](#)

## 1.1. HOW THE VULNERABILITY SERVICE WORKS

The vulnerability service uses the Insights client to gather information about your RHEL systems. The client gathers information about the systems and uploads it to the vulnerability service.

The vulnerability service then assesses the data against the Red Hat CVE database and security bulletins to determine if there are any outstanding CVEs that could affect the systems, and provides the results of those comparisons.

Once the data has been analyzed, you can view and sort the displayed results, assess the risks and priorities of the vulnerabilities, report their status, and create and deploy Ansible Playbooks to remediate them. The goal of the vulnerability service is to enable a repeatable process that protects against security weaknesses in your RHEL infrastructure.

## CHAPTER 2. USER ACCESS FOR VULNERABILITY SERVICE USERS

Before you can access certain features in the Red Hat Insights for Red Hat Enterprise Linux application, you must have the correct permissions, which are granted in [Red Hat Hybrid Cloud Console > User Access > Groups](#). An Organization Administrator or **User Access administrator** must add you as a member to a User Access group with the required roles.

By default, User Access on the Red Hat Hybrid Cloud Console has preconfigured **Vulnerability administrator** (all access) and **Vulnerability viewer** (read-only access) roles. If your organization determines that the predefined roles provide insufficient access, a **User Access administrator** can configure a custom role to provide the specific permissions required by a set of users.

The following sections in this chapter describe each of the predefined roles for vulnerability service users.



### IMPORTANT

Changes to User Access must be performed by an Organization Administrator on your Red Hat account, or by an account user who is a member of a User Access group with the **User Access administrator** role.

### Additional resources

- [User Access Configuration Guide for Role-based Access Control \(RBAC\)](#)

## 2.1. VULNERABILITY ADMINISTRATOR ROLE

The **Vulnerability administrator** role is a default role in the **Default access group**. All Red Hat Insights for Red Hat Enterprise Linux users on your account are members of the **Default access group** by default. In its default configuration, members of a group with the **Vulnerability administrator** role have access to all vulnerability service resources.

Your organization may decide that the default role is too limited, or too permissive. To limit access to some features, or to add additional permissions, a **User Access administrator** can customize the role and configure it with whatever permissions are required. By customizing the preconfigured role, the **Default access group** is replaced.

## 2.2. VULNERABILITY VIEWER ROLE

In its default configuration, the **Vulnerability viewer** role can read any vulnerability service resource. This is a preconfigured role but is not included in the **Default access group**. The **Vulnerability viewer** role includes the following permissions:

- View and triage all vulnerability service results, pages, and lists.
- View which systems have been opted out of reporting results to Red Hat Insights for Red Hat Enterprise Linux.
- Set filters and export data for .JSON and .CSV output.
- View and create advanced reports in [Red Hat Enterprise Linux > Vulnerability > Reports](#) .



If your organization determines that the default configuration of the Vulnerability viewer role is inadequate, a User Access administrator can create a custom role with the specific permissions required.

## CHAPTER 3. COMMON VULNERABILITIES AND EXPOSURES (CVEs)

Common Vulnerabilities and Exposures (CVEs) are security vulnerabilities identified in publicly released software packages. CVEs are identified and listed by the National Cybersecurity FFRDC (NCF), the federally funded research and development center operated by the Mitre Corporation, with funding from the National Cyber Security Division of the United States Department of Homeland Security. The complete list of CVEs is available at <https://cve.mitre.org>.

By highlighting CVEs with publicly known exploits and security rules associated with CVEs, the vulnerability service surfaces enhanced threat intelligence to aid in determining which CVEs pose the greatest potential risk to RHEL environments, enabling our users to effectively prioritize and address their most critical issues first.



### IMPORTANT

The vulnerability service does not contain every CVE included in the list of entries at <https://cve.mitre.org>. Only Red Hat CVEs, those CVEs for which Red Hat issues security advisories (RHSAs), are included in the vulnerability service.

The vulnerability service identifies CVEs impacting your RHEL systems, indicates the severity and enables you to efficiently triage the exposures that are most critical to resolve. The dashboard will alert you to the following types of CVEs:

- Known exploits
- Security rules
- Critical severity
- Important severity

The screenshot shows the Red Hat Hybrid Cloud Console interface. The main content area displays the 'CVEs' dashboard. At the top, there is a 'Filter by tags' dropdown. Below that, the title 'CVEs' is followed by four summary cards:

- CVEs with known exploits:** 4
- CVEs with security rules:** 13
- CVEs with critical severity:** 0
- CVEs with important severity:** 113

Below the cards, there is a search bar labeled 'Search ID or description' and a filter dropdown set to 'CVE'. There are also filter tags for 'Systems' (1 or more) and 'Security rules' (Has security rule). A 'Reset filters' link is visible. The sidebar on the left contains navigation options: Overview, Red Hat Insights, Inventory, Advisor, Patch, Vulnerability (expanded), CVEs (selected), Reports, Systems, and Compliance. The top navigation bar shows 'All apps and services' and user profile information.

:context: vuln-cves

### 3.1. RED HAT SECURITY ADVISORIES (RHSAS)

Red Hat Security Advisory (RHSA) errata document security vulnerabilities in Red Hat products for which there are remediations or mitigations available. The Red Hat Insights for Red Hat Enterprise Linux vulnerability service displays the advisory identifier tied to each system exposed to a CVE.

View this information by selecting a CVE and selecting the **Filter by affected systems** link in the security rule card. If an advisory exists for the system, the RHSA ID is visible as a link next to the system in the **Exposed systems** list, **Advisory** column. When there are no such advisories, the Advisory column is not visible, or will show “Not available.”

When an advisory exists for a system, users can view more information about the RHSA, including a list of affected systems. In the patch service, users can select systems to create an Ansible Playbook to apply the remediation.

The screenshot shows the Red Hat Insights interface for the advisory RHSA-2020:4183. The sidebar on the left contains navigation links for Dashboard, Advisor, Vulnerability, Compliance, Patch, Affected systems, Systems, Packages, Drift, Policies, Image Builder, Inventory, Remediations, Register Systems, Subscription Watch, and Documentation. The main content area displays the advisory details, including the title 'RHSA-2020:4183', a description of the vulnerability in the BIND DNS server, a severity indicator of 'Moderate', and a list of affected systems. The affected systems table has columns for Name, Packages, Applicable advisories, and Last seen. Three systems are selected, and a 'Remediate' button is visible.

Name	Packages	Applicable advisories	Last seen
<input checked="" type="checkbox"/> RHIQE.d602028f-25b3-43c6-87cb-6847d308a92d.iqe-insights-client-plugin	398	U 37 30 11	16 hours ago
<input checked="" type="checkbox"/> 4e6d5545-c506-4599-be95-3565a8815cd3	398	U 37 30 11	16 hours ago
<input checked="" type="checkbox"/> RHIQE.092a2477-ecb0-41dc-8677-d46019019597.iqe-insights-client-plugin	398	U 37 30 11	2 days ago
<input type="checkbox"/> 4500fd7-0b10-454f-b1ef-a69d7f6ead2d	398	U 37 30 11	2 days ago
<input type="checkbox"/> RHIQE.6b7500a8-6440-4190-b2c5-f2c2cba5f32c.iqe-insights-client-plugin	398	U 37 30 11	3 days ago

## 3.2. SECURITY RULES

Security rules are CVEs given additional visibility due to the elevated risk and exposure associated with them. These are security flaws that may receive significant media coverage and have been scrutinized by the Red Hat Product Security team, using the [Product Security Incident Response Plan](#) workflow to help determine your RHEL environment exposure. These security rules enable you to take the appropriate action to protect your organization.

Security rules provide deep threat intelligence, beyond analyzing the version of RHEL running on a system. Security rules are manually curated to determine whether you are susceptible to a security threat by analyzing system metadata collected by the Insights client. If the vulnerability service identifies a system as exposed to a security rule, there is the potential for elevated security risk and issues should be addressed with urgency.



### IMPORTANT

Addressing security rules on exposed systems should be your highest priority.


Finally, not all systems exposed to a CVE are also exposed to a security rule associated with that CVE. Even though you may be running a vulnerable version of software, other environmental conditions may mitigate the threat; for example, a specific port is closed or if you are running SELinux.

### 3.2.1. Identifying security rules in the Insights for RHEL dashboard

Use the following steps to view your infrastructure exposure to security rules.

#### Procedure

1. Navigate to the [Red Hat Insights for Red Hat Enterprise Linux dashboard](#).




#### NOTE

For simplicity, panels for services not related to security vulnerability assessment are minimized in the following screenshot.

The screenshot displays the Red Hat Insights dashboard with a sidebar on the left containing navigation options like Dashboard, Advisor, Drift, Inventory, Vulnerability, Compliance, Policies, Patch, Subscriptions, Resource Optimization, Register Systems, Remediations, and Product Materials. The main content area shows a summary of 7,648 systems registered with Insights, including 4,925 stale systems and 4,587 systems to be removed. A 'Latest critical notifications on your systems' panel highlights a newly released security rule for Linux-firmware. Below this, a 'Vulnerability' section provides a breakdown of CVEs: 18 with security rules and 4 with known exploits. A 'CVEs by CVSS score' pie chart and table are also shown, along with panels for Advisory recommendations, Remediations, and Subscription Watch utilization summary.

CVSS score	CVE totals	Known exploits
8.0 - 10	113	1
4.0 - 7.9	554	3
0.0 - 3.9	98	0

2. View the **Latest critical notifications** on your systems panel. These are security rules with an elevated severity rating of "Important" or "Critical." These are potentially your most critical issues and should be prioritized for remediation.
  - a. To the right of each notification, click the **Expand** button to see associated CVEs and the number of systems exposed in your infrastructure.



#### NOTE

You may see security rules in your critical notifications but have zero systems exposed. In this case, even though the CVE is present in your infrastructure, the security rule conditions may not exist.

- b. Below the name of the security rule, and under Associated CVEs, click the CVE ID link.
  - c. View which of your systems is impacted by the security rule CVE and optionally select exposed systems to create playbooks.
3. Next, view the information in the **vulnerability** card.
    - a. Note the number of “CVEs with **security rules** impacting systems.” This number includes security rules of any severity impacting at least one system.
      - i. Click **View CVEs**. Consider lesser-severity security rules your second highest priority for remediation, following high-severity security rules.

### 3.3. KNOWN EXPLOITS

Red Hat analyzes Metasploit data to determine whether code exists publicly to exploit a CVE, or a CVE has already been exploited publicly. The vulnerability service applies the “Known exploits” label to CVEs that meet that criteria.

This enhanced threat assessment can help users identify and address those CVEs that pose the most critical risks first. Red Hat recommends users review any CVEs with the “Known exploit” label with high priority and work towards remediating those issues.



#### IMPORTANT

The vulnerability service makes you aware that the known-exploit CVE exists on systems in your infrastructure. The “Known exploits” label does not mean that the vulnerability was exploited on your RHEL systems; the vulnerability service does not make that determination.

### 3.4. COMMON VULNERABILITIES AND EXPOSURES PROVIDE DEEP THREAT INTELLIGENCE WITH TRIAGE FEATURE

The vulnerability service provides you with data about individual Common Vulnerabilities and Exposures (CVEs) and their effect on your systems registered to Insights. CVEs are categorized as **vulnerable** or **affected but not vulnerable**. This level of threat intelligence is available for CVEs that have the **Security Rule** label or for those that have gone through Red Hat Product Security’s rigorous analysis.

This increased threat intelligence enables you to triage issues and address the most urgent ones first. When managing a large fleet of servers, this translates into expedited protection and significant efficiencies.

An **affected but not vulnerable** CVE status indicates that you are running software that has a vulnerability in it but is not currently exploitable. This system will need remediation but does not require immediate attention.

A **vulnerable** CVE status indicates flawed code with an open path to exploitation. An open path could be a port or an OS version that permits one of the following: confidential information to be leaked, the integrity of the system to be compromised or availability of the system to be hindered.

Let us look at an example of a **vulnerable** server versus an **affected but not vulnerable** server:

Suppose that **Server A** is running vulnerable software that allows root access to the system. **Server A** would be considered vulnerable and require immediate patching.

In contrast, suppose that **Server B's** current configuration prevents the vulnerability from manifesting, even when present in the affected code. **Server B** would be considered **affected but not vulnerable**. This would mean that **Server B** could be relegated to the to-do list, so that the more immediate threat, **Server A** could be remediated.

**IMPORTANT:** You should patch **Server B** once **Server A** has been addressed since it is running potentially vulnerable code. Version updates and other events could render it vulnerable in the future.

### 3.4.1. Identifying known-exploit CVEs in the Red Hat Insights for RHEL dashboard

Use the following steps to identify known-exploit CVEs in the Insights for Red Hat Enterprise Linux dashboard vulnerability card.

#### Procedure

1. Navigate to the [Red Hat Insights for Red Hat Enterprise Linux dashboard](#).



#### NOTE

For simplicity, panels for services not related to security vulnerability assessment are minimized in the following screenshot.

The screenshot shows the Red Hat Insights dashboard with the following data:

- Systems registered with Insights:** 7,648
- Stale systems:** 4,925
- Systems to be removed:** 4,587
- Latest critical notifications:** Newly released security rule: 24 Mar 2021. Linux-firmware: Denial of Service or Privilege Escalation in Bluetooth range.
- Vulnerability Card:**
  - 18 CVEs with security rules impacting 1 or more systems
  - 4 CVEs with known exploits impacting 1 or more systems
  - CVES by CVSS score:
 

CVSS score	CVE totals	Known exploits
8.0 - 10	113	1
4.0 - 7.9	554	3
0.0 - 3.9	98	0

2. On the **Vulnerability** card, note the **CVEs with Known exploits impacting 1 or more systems** and the number displayed.
3. Click **View Known exploits**.
4. View the filtered list of Known-exploit CVEs in the CVEs list.

## CHAPTER 4. REFINING VULNERABILITY-SERVICE RESULTS

Whether reporting results to stakeholders or prioritizing systems for remediation, the vulnerability service enables many ways to refine the views of your data, helping you and others focus on your most critical systems, workloads, or issues. The following sections describe the organization of your data and the sorting, filtering, and contextual features you can use to refine and enrich your results.

### 4.1. CVE-LIST AND SYSTEMS-LIST FILTERS

Filtering narrows the visible list of CVEs and associated systems, helping you focus on specific issues. Apply filters to the CVEs list to focus on CVEs by criticality or business risk, for example. After selecting an individual CVE, apply filters to the resulting list of affected systems to focus on those of a specific RHEL major or minor version, for example.

Filters are activated by selecting a primary filter from the drop down list of filters on the left, and then selecting a secondary subfilter from the filter options drop down list on the right. Selected filters are visible below the Filters menu and can be deactivated by clicking the X next to each one.

#### CVEs list filters

The screenshot shows the CVEs list interface. At the top, there is a "Filter by status" dropdown menu. Below it, the "CVEs" section is visible. A dropdown menu is open, showing the following options: CVE, Security rules, Known exploit, Severity, CVSS base score, Business risk, Systems exposed, Publish date, and Status. The "CVE" option is selected. Below the dropdown, a table of CVEs is displayed. The table has columns for CVE ID, Publish date, Severity, and CVSS base score. The table contains several rows of CVEs, each with a checkbox, a CVE ID, a publish date, a severity level, and a CVSS base score.

CVE ID	Publish date	Severity	CVSS base score
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Critical	8.8
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Moderate	6.8
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Important	8.1
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Important	9.0
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Important	9.0
> <input type="checkbox"/> CVE-2021-21887	04 Nov 2021	Important	9.0

The following primary filters are accessible from the CVEs page. Select the primary filter, then define a parameter in the subfilter:

- **CVE.** Search ID or description.
- **Security rules.** Show only CVEs with the "Security rule" label.
- **Known exploit.** Show only CVEs with the "Known exploit" label.
- **Severity.** Select one or more values: Critical, Important, Moderate, Low, or Unknown.
- **CVSS base score.** Select one or more ranges: All, 0.0-3.9, 4.0-7.9, 8.0-10.0, N/A (not applicable)
- **Business risk.** Select one or more values: High, Medium, Low, Not defined.
- **Systems exposed .** Select to only show CVEs with systems currently affected, or with no systems affected.
- **Publish date.** Select from All, Last 7 days, Last 30 days, Last 90 days, Last year, or More than 1 year ago.
- **Status.** Select one or more values: Not reviewed, In review, On-hold, Scheduled for patch, Resolved, No action - risk accepted, Resolved via mitigation.

### Systems list filters

The screenshot shows the 'Exposed systems' section with a dropdown menu open for 'Operating system'. The menu options are: Name, Security rules, Status, Advisory, Operating system, and Remediation. The background table lists systems with their respective tags and operating systems.

Name	Tags	OS
satellit	0	RHEL 7.9
idm8.r	9	RHEL 8.4
cap67.	6	RHEL 7.9
mhuth	0	RHEL 8.4
satellite.anzias.dns.regnat.com	0	RHEL 7.9

The following primary filters are accessible from the top of the list of systems on the CVE details page:

- **Name.** Find a specific CVE by entering the CVE ID.
- **Security rules.** If the CVE has a security rule associated with it, filter by other systems vulnerable to the same security rule, or show systems not affected by the security rule.
- **Status.** Show systems in specific status or workflow categories.
- **Advisory.** Show systems to which a Red Hat advisory applies for this CVE.
- **Operating system.** Show systems running specific RHEL (minor) versions.



- **Remediation.** Show systems included in an Ansible Playbook, a manual remediation, or that are not included in a current remediation plan.

## 4.2. FILTERING SECURITY-RULE CVEs

Security rules, especially high-severity security rules, pose the greatest potential threat to your infrastructure and should be considered the highest priority for identification and remediation. Use the following procedure to view only high-severity security-rule CVEs in the CVEs list and identify affected systems.



### NOTE


Not all systems exposed to a CVE are also exposed to a security rule associated with that CVE. Even though you may be running a vulnerable version of software, other environmental conditions may mitigate the threat; for example, a specific port is closed or SELinux is enabled.

### Procedure

1. Navigate to [Red Hat Enterprise Linux > Vulnerability > CVEs](#) in Red Hat Insights for Red Hat Enterprise Linux.
2. Click the filters dropdown list in the toolbar.
  - a. Apply the **Security rules** filter.
  - b. Apply the **Has security rule** subfilter.
3. Scroll down to view security-rule CVEs. CVEs with security rules display the security-rule label located immediately below the CVE ID.

## 4.3. FILTERING KNOWN-EXPLOIT CVEs

CVEs with the “Known exploit” label are determined by Red Hat to have exploits that exist in the wild; either the code exists publicly to exploit the CVE, or an exploit is known publicly to have already happened. For these reasons, known-exploit CVEs should be prioritized for identification and remediation.



### IMPORTANT

Red Hat does not determine whether any of your registered systems have been exploited. We are simply identifying CVEs that may pose an extraordinary risk.

Use the following steps to filter known-exploit CVEs in the CVEs list:

### Procedure

1. Navigate to [Red Hat Enterprise Linux > Vulnerability > CVEs](#) in Red Hat Insights for Red Hat Enterprise Linux.
2. Click the filters drop-down list in the toolbar.
  - a. Apply the **Known exploit** filter.

- b. Apply the **Has a known exploit** subfilter.
3. Scroll down to view the list of known-exploit CVEs.

## 4.4. FILTERING LISTS OF SYSTEMS EXPOSED TO SECURITY RULES

After filtering the list of CVEs to view only your most critical potential threats, select an individual CVE to view the list of exposed systems and apply a filter to the list.

### Procedure

1. After selecting a security-rule CVE, scroll down to the **Exposed systems** list. Not every system in the list has the security rule conditions present for the CVE to be a security rule. Apply the following filter to see only the systems with security rule conditions present.
2. Select the **Security rules** filter from the primary filter dropdown list.
3. Check the **Has security rule** box in the secondary subfilter dropdown list.
4. View the systems with exposure to that CVE that also have the conditions present for the security rules.

## 4.5. INSIGHTS FOR RHEL GROUP FILTERS

The ability to filter vulnerability service results by groups of systems or workloads enables users to view only those systems tagged as belonging to a specific group. These can be systems running SAP workloads (or by SAP ID), by Satellite host groups, or by custom tags added to the Insights client configuration file.

Group filtering can be set globally in Insights for Red Hat Enterprise Linux using the **Filter results** box located at the top of the page throughout the Insights for Red Hat Enterprise Linux application. Group selection persists when changing from service to service and page to page. However, the functionality varies within the different Insights for Red Hat Enterprise Linux services.

Group filtering works in the vulnerability Dashboard and vulnerability service CVEs and Systems lists.

Learn more about group tags and configuring custom tags in *Tags and system groups* section of this document.

### 4.5.1. Filtering Dashboard, CVEs, and Systems lists by group

Use the following procedure to filter vulnerability service CVE and Systems lists by group.

#### Procedure

1. Navigate to [Red Hat Hybrid Cloud Console](#) and log in.
2. Open the Red Hat Insights for Red Hat Enterprise Linux application.
3. Click the down arrow on the **Filter results** box located at the top of any page in the Insights application.
4. Select a group by which to filter your systems.  
Search or scroll to view available tags. To browse the full list of available tags, scroll to the bottom of the list and click **View more**.

Optionally,

- a. Select SAP workloads.
  - b. Select systems by specific SAP IDs.
  - c. Select Satellite host collections.
  - d. Select systems identified by custom group tags.  
To learn more about creating custom tags, see section, **Custom system tagging**, in this document.
5. Navigate to the service and view only systems or CVEs that belong to your selected group or groups.

## 4.6. DEFINING A BUSINESS RISK FOR A CVE


The vulnerability service allows you to define the business risk of a CVE with the following options: High, Medium, Low, or Not Defined (default).

While the list of CVEs shows the severity of each CVE, assigning a business risk lets you rank CVEs based on the impact they could have on your organization. This can give you more control in managing your risk efficiently in a large environment, and enable you to make better operational decisions.

By default, the business risk field for a specific CVE is set to **Not Defined**. After you set the business risk, it is visible in the [Red Hat Enterprise Linux > Vulnerability > CVEs](#) list, in the CVE row.

CVE ID	Publish date	Severity	CVSS base score	Systems exposed	Business risk	Status
<a href="#">CVE-2020-11008</a>	20 Apr 2020	Important	7.5	260	Medium	Resolved

Business risk is also visible on the details card for each CVE, which shows more information and lists affected systems.



### 4.6.1. Setting a business risk for a single CVE

Complete the following steps to set the business risk for a single CVE:



#### NOTE

The business risk for that CVE will be the same on *all* systems impacted by it.

1. Navigate to the [Red Hat Enterprise Linux > Vulnerability > CVEs](#) page and log in if necessary.
2. Identify a CVE for which to set a business risk.

- Click the **more-actions** icon (three vertical dots) on the right end of the CVE row and click **Edit business risk**.

>	<input type="checkbox"/>	<a href="#">CVE-2020-5260</a>	14 Apr 2020	Important	7.5	3	Not defined	Not reviewed	
>	<input type="checkbox"/>	<a href="#">CVE-2020-2754</a>	13 Apr 2020	Low	3.7	2	Not defined	Nc	

Edit business risk  
 Edit status

- Set a business risk value to the appropriate level and, optionally, add a justification for your risk assessment.
- Click **Save**.

#### 4.6.2. Setting a business risk for multiple CVEs

Complete the following steps to set the same business risk on multiple CVEs that you select:

- Navigate to [Red Hat Enterprise Linux > Vulnerability > CVEs](#) and log in if necessary.
- Check the boxes for the CVEs for which you want to set a business risk.
- Perform the following steps to set a business risk:
  - Click the **more-actions** icon (three vertical dots) to the right of the Filters dropdown menu in the toolbar and click **Edit business risk**
  - Set an appropriate business risk value and, optionally, add a justification for your risk assessment.
  - Click **Save**.

## 4.7. EXCLUDING SYSTEMS FROM VULNERABILITY SERVICE ANALYSIS

The vulnerability service allows you to exclude specific systems from vulnerability analysis. This can save you the time and attention required to review and re-review issues on systems that are not relevant to your organization's goals.

As an example, if you have the following category of servers: QA, Dev, and Production, you may not care to review the vulnerabilities for your QA servers and therefore want to exclude these systems from the analysis performed by the vulnerability service.

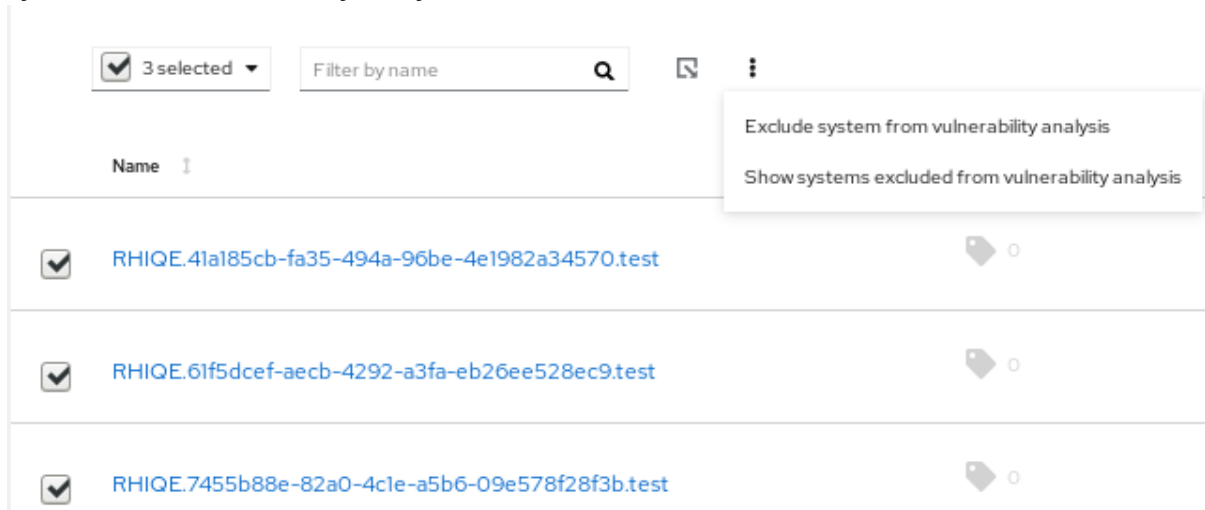
When you exclude systems from vulnerability analysis, the Insights client still runs per schedule on the system, but the results for the system are not visible in the vulnerability service. The continued operations of the client ensure that other Red Hat Insights for Red Hat Enterprise Linux services can still upload the data they need. It also means that you can still view results for those systems using filtering.

Complete the following steps to exclude selected RHEL systems from vulnerability service analysis:

#### Procedure

- Navigate to the [Red Hat Enterprise Linux > Vulnerability > Systems](#) tab and log in if necessary.
- Check the box for each system you want to exclude from vulnerability analysis.

- Click the **more-actions** icon in the toolbar, at the top of the list of systems, and select **Exclude systems from vulnerability analysis**.



- Optionally, you can exclude a *single* system by clicking the **more-actions** icon in the system row and selecting **Exclude system from vulnerability analysis**



## 4.8. SHOWING PREVIOUSLY EXCLUDED SYSTEMS

Complete the following steps to show a previously excluded system:

### Procedure

- Navigate to the [Red Hat Enterprise Linux > Vulnerability > Systems](#) tab and log in if necessary.
- Click the **more-actions** icon in the toolbar, at the top of the list of systems, and select **Show systems excluded from analysis**.
- See systems excluded from vulnerability analysis. This can be verified by the value of **Excluded** in the **Applicable CVEs** column.

## 4.9. RESUMING VULNERABILITY ANALYSIS FOR A SYSTEM

Complete the following steps to resume vulnerability analysis for a system:

### Procedure

- Navigate to the [Red Hat Enterprise Linux > Vulnerability > Systems](#) tab and log in if necessary.
- Click the **more-actions** icon in the toolbar, at the top of the list of systems, and select **Show systems excluded from analysis**.
- In the list of results, check the box for each system for which you want to resume vulnerability analysis.
- Click the **more-actions** icon again and select **Resume analysis for system**

## 4.10. CVE STATUS

Another method of managing CVEs impacting your systems is by setting a status for CVEs. The vulnerability service enables the following ways of setting a status for a CVE:

- Set a status for a CVE for *all* systems.
- Set a status for a *specific CVE + system pair*.

Status values are preset and include the following options:

- Not reviewed (default)
- In-review
- On-hold
- Scheduled for patch
- Resolved
- No action - risk accepted
- Resolved via mitigation

Setting a status for a CVE can facilitate better triaging through its life-cycle, from becoming aware of it to remediating it. Defining a status allows your organization to keep better tabs on where the most critical CVEs are in their life-cycle and where you should focus your efforts to address the most critical issues per your business need. The status for a CVE is visible in all CVE tables in the vulnerability service and in individual CVE views.

### 4.10.1. Setting a status for a CVE on all affected systems

Complete the following steps to set a status for a CVE and have that status apply to that CVE on all of the systems it impacts:

#### Procedure


1. Navigate to the [Red Hat Enterprise Linux > Vulnerability > CVEs](#) tab and log in if necessary.
2. Click the **more-actions** icon located on the right end of the CVE row and select **Edit status**.
3. Select the appropriate status and, optionally, enter a rationale for your decision in the **Justification** text box.
4. Check **Do not overwrite individual system status** if there are statuses set for this CVE on individual systems and that you want to preserve. Otherwise, leave the box unchecked to apply this status to all of the systems it is impacting.
5. Click **Save**.

### 4.10.2. Setting a status for a CVE and system pair

Complete the following steps to set a status on a CVE and system pair:

#### Procedure

1. Navigate to the [Red Hat Enterprise Linux > Vulnerability > Systems](#) tab and log in if necessary.
2. Identify the system and click the system name to open it.
3. Select a CVE from the list and check the box next to the CVE ID.
4. Click the **more-options** icon in the toolbar and select **Edit status**.



5. In the popup card, take the following actions:
  - a. Set a status for the CVE and system pair.



#### NOTE

If the box to **Use overall CVE status** is checked, you cannot set a status for the pair.

- b. Optionally, enter a justification for your status determination.
  - c. Click **Save**.
6. Locate the CVE in the list and verify the status is set.

## 4.11. USING THE SEARCH BOX

The search function in the vulnerability service works in the context of the page you are viewing.

- **CVEs page.** The search box is located in the toolbar at the top of the CVEs list. With the CVE filter set, search CVE IDs and descriptions.



- **Systems page.** The search box is located in the toolbar at the top of the list. Search for system name or UUID.



## 4.12. SORTING CVE LIST DATA

The sorting functions in the vulnerability service differ based on the context of the page you are viewing.

### Procedure

1. In the **CVEs tab**, you can apply sorting to the following columns:

- CVE ID
  - Publish date
  - Severity
  - CVSS base score
  - Systems exposed
  - Business risk
  - Status
2. In the **Systems tab**, the following column can be sorted:
- Name
  - Applicable CVEs
  - Last seen
3. After selecting a system in the Systems tab, the system-specific list of CVEs allows the following sorting options:
- CVE ID
  - Publish date
  - Impact
  - CVSS base score
  - Business risk
  - Status



## CHAPTER 5. SYSTEM TAGS AND GROUPS

Red Hat Insights for Red Hat Enterprise Linux enables administrators to filter groups of systems in inventory and in individual services using group tags. Groups are identified by the method of system data ingestion to Insights for Red Hat Enterprise Linux. Insights for Red Hat Enterprise Linux enables filtering groups of systems by those running SAP workloads, by Satellite host group, by Microsoft SQL Server workload, and by custom tags that are defined by system administrators with root access to configure the Insights client on the system.



### NOTE

As of Spring 2022, inventory, advisor, compliance, vulnerability, patch, drift, and policies enable filtering by groups and tags. Other services will follow.



### IMPORTANT

Unlike the other services that enable tagging, the compliance service sets tags within lists of systems in the compliance service UI. For more information, see [Group and tag filters in the compliance service](#).

Use the global, **Filter results** box to filter by SAP workloads, Satellite host groups, MS SQL Server workloads, or by custom tags added to the Insights client configuration file.

### Prerequisites

The following prerequisites and conditions must be met to use the tagging features in Red Hat Insights for Red Hat Enterprise Linux:

- The Red Hat Insights client is installed and registered on each system.
- You must have root permissions, or their equivalent, to create custom tags or change the `/etc/insights-client/tags.yaml` file.

## 5.1. SAP WORKLOADS

As Linux becomes the mandatory operating system for SAP ERP workloads in 2025, Red Hat Enterprise Linux and Red Hat Insights for Red Hat Enterprise Linux are working to make Insights for Red Hat Enterprise Linux the management tool of choice for SAP administrators.

As part of this ongoing effort, Insights for Red Hat Enterprise Linux automatically tags systems running SAP workloads and by SAP ID (SID), without any customization needed by administrators. Users can easily filter those workloads throughout the Insights for Red Hat Enterprise Linux application by using the global **Filter by tags** drop-down menu.

## 5.2. SATELLITE HOST GROUPS

Satellite host groups are configured in Satellite and recognized automatically by Insights for Red Hat Enterprise Linux.

## 5.3. MICROSOFT SQL SERVER WORKLOADS

Using the global **Filter by tags** feature, Red Hat Insights for Red Hat Enterprise Linux users can select groups of systems running Microsoft SQL Server workloads.

In May of 2019, the Red Hat Insights team introduced a new set of Insights for Red Hat Enterprise Linux recommendations for Microsoft SQL Server running on Red Hat Enterprise Linux (RHEL). These rules alert administrators to operating system level configurations that do not conform to the documented recommendations from Microsoft and Red Hat.

A limitation of these rules was that they primarily analyzed the operating system and not the database itself. The latest release of Insights for Red Hat Enterprise Linux and RHEL 8.5, introduces Microsoft SQL Assessment API. The SQL Assessment API provides a mechanism to evaluate the database configuration of MS SQL Server for best practices. The API is delivered with a rule set containing best practice rules suggested by the Microsoft SQL Server Team. While this rule set is enhanced with the release of new versions, the API is built with the intent to give a highly customizable and extensible solution, which enables users to tune the default rules and create their own.

The SQL Assessment API is supported by PowerShell for Linux (available from Microsoft), and Microsoft has developed a PowerShell script that can be used to call the API and store its results as a JSON formatted file. With RHEL 8.5, the Insights client now uploads this JSON file and presents the results in an easy-to-understand format in the Insights for Red Hat Enterprise Linux UI.

For more information about SQL Server assessment in Insights for Red Hat Enterprise Linux, see [SQL Server database best practices now available through Red Hat Insights](#).

### 5.3.1. Setting up SQL Server assessments

To configure the Microsoft SQL Assessment API to provide information to Red Hat Insights, the database administrator needs to take the following steps.

#### Procedure

1. In the database you wish to assess, create a login for SQL Server assessments using SQL Authentication. The following Transact-SQL creates a login. Replace `<*PASSWORD*>` with a strong password:

```
USE [master]
GO
CREATE LOGIN [assessmentLogin] with PASSWORD= N'<*PASSWORD*>'
ALTER SERVER ROLE [sysadmin] ADD MEMBER [assessmentLogin]
GO
```

2. Store the credentials for login on the system as follows, again replacing `<*PASSWORD*>` with the password you used in step 1.

```
# echo "assessmentLogin" > /var/opt/mssql/secrets/assessment
# echo "<*PASSWORD*>" >> /var/opt/mssql/secrets/assessment
```

3. Secure the credentials used by the assessment tool by ensuring that only the mssql user can access the credentials.

```
# chmod 0600 /var/opt/mssql/secrets/assessment
# chown mssql:mssql /var/opt/mssql/secrets/assessment
```

4. Download PowerShell from the microsoft-tools repository. This is the same repository you configured when you installed the **mssql-tools** and **mssqldb17** packages as part of SQL Server installation.

```
# yum -y install powershell
```

5. Install the SQLServer module for PowerShell. This module includes the assessment API.

```
# su mssql -c "/usr/bin/pwsh -Command Install-Module SqlServer"
```

6. Download the runassessment script from the Microsoft examples GitHub repository. Ensure it is owned and executable by mssql.

```
# /bin/curl -LJO -o /opt/mssql/bin/runassessment.ps1
https://raw.githubusercontent.com/microsoft/sql-server-samples/master/samples/manage/sql-
assessment-api/RHEL/runassessment.ps1
# chown mssql:mssql /opt/mssql/bin/runassessment.ps1
# chmod 0700 /opt/mssql/bin/runassessment.ps1
```

7. Create the directory that will store the log file used by Red Hat Insights. Again, make sure it is owned and executable by mssql.

```
# mkdir /var/opt/mssql/log/assessments/
# chown mssql:mssql /var/opt/mssql/log/assessments/
# chmod 0700 /var/opt/mssql/log/assessments/
```

8. You can now create your first assessment, but be sure to do so as the user mssql so that subsequent assessments can be run automatically via cron or systemd more securely as the mssql user.

```
# su mssql -c "pwsh -File /opt/mssql/bin/runassessment.ps1"
```

9. Insights for Red Hat Enterprise Linux will automatically include the assessment next time it runs, or you can initiate Insights client by running this command:

```
# insights-client
```

### 5.3.1.1. Setting up the SQL Assessment on a timer

Because SQL Server Assessments can take 10 minutes or more to complete, it may or may not make sense for you to run the assessment process automatically every day. If you would like to run them automatically, the Red Hat SQL Server community has created systemd service and timer files to use with the assessment tool.

#### Procedure

1. Download the following files from [Red Hat public SQL Server Community of Practice GitHub site](#).
  - **mssql-runassessment.service**
  - **mssql-runassessment.timer**
2. Install both files in the directory **/etc/systemd/system/**:

```
# cp mssql-runassessment.service /etc/systemd/system/
# cp mssql-runassessment.timer /etc/systemd/system/
# chmod 644 /etc/systemd/system/
```

3. Enable the timer with:

```
# systemctl enable --now mssql-runassessment.timer
```

## 5.4. CUSTOM SYSTEM TAGGING

By applying custom grouping and tagging to your systems, you can add contextual markers to individual systems, filter by those tags in the Insights for Red Hat Enterprise Linux application, and more easily focus on related systems. This functionality can be especially valuable when deploying Insights for Red Hat Enterprise Linux at scale, with many hundreds or thousands of systems under management.

In addition to the ability to add custom tags to several Insights for Red Hat Enterprise Linux services, you can add predefined tags. The advisor service can use those tags to create targeted recommendations for your systems that might require more attention, such as those systems that require a higher level of security.



### NOTE

To create custom and predefined tags, you must have root permissions, or their equivalent, to add to, or change the `/etc/insights-client/tags.yaml` file.

### 5.4.1. Tag structure

Tags use a **namespace/key=value** paired structure.

- **Namespace.** The namespace is the name of the ingestion point, *insights-client*, and cannot be changed. The **tags.yaml** file is abstracted from the namespace, which is injected by the Insights client before upload.
- **Key.** The key can be a user-chosen key or a predefined key from the system. You can use a mix of capitalization, letters, numbers, symbols and whitespace.
- **Value.** Define your own descriptive string value. You can use a mix of capitalization, letters, numbers, symbols and whitespace.



### NOTE

The advisor service includes Red Hat-supported predefined tags.

### 5.4.2. Creating a tags.yaml file and adding a custom group

Create and add tags to `/etc/insights-client/tags.yaml` simply by using **insights-client --group=<name-you-choose>**, which performs the following actions:

- Creates the **etc/insights-client/tags.yaml** file
- Adds the **group=** key and **<name-you-choose>** value to **tags.yaml**
- Uploads a fresh archive from the system to the Insights for Red Hat Enterprise Linux application so the new tag is immediately visible along with your latest results

After creating the initial **group** tag, add additional tags as needed by editing the `/etc/insights-client/tags.yaml` file.

The following procedure shows how to create the `/etc/insights-client/tags.yaml` file and the initial group, then verify the tag exists in the Insights for Red Hat Enterprise Linux inventory.

### Procedure to create new group

1. Run the following command as root, adding your custom group name after `--group=`:

```
[root@server ~]# insights-client --group=<name-you-choose>
```

### Example of tags.yaml format

The following example of a **tags.yaml** file shows an example of file format and additional tags added for the new group:

```
# tags
---
group: eastern-sap
name: Jane Example
contact: jexample@corporate.com
Zone: eastern time zone
Location:
- gray_rack
- basement
Application: SAP
```

### Procedure to verify your custom group was created

1. Navigate to [Red Hat Enterprise Linux > Inventory](#) and log in if necessary.
2. Click the **Filter results** dropdown menu.
3. Scroll through the list or use the search function to locate the tag.
4. Click the tag to filter by it.
5. Verify that your system is among the results on the advisor systems list.
6. Procedure to verify the system is tagged
7. Navigate to [Red Hat Enterprise Linux > Inventory](#) and log in if necessary.
8. Activate the **Name** filter and begin typing the system name until you see your system, then select it.
9. Verify that, next to the system name, the tag symbol is darkened and shows a number representing the correct number of tags applied.

### 5.4.3. Editing tags.yaml to add or change tags

After creating the group filter, edit the contents of `/etc/insights-client/tags.yaml` as needed to add or modify tags.


## Procedure

1. Using the command line, open the tag configuration file for editing.

```
[root@server ~]# vi /etc/insights-client/tags.yaml
```

2. Edit content or add additional values as needed. The following example shows how you can organize **tags.yaml** when adding multiple tags to a system.

```
# tags
---
group: eastern-sap
location: Boston
description:
- RHEL8
- SAP
key 4: value
```



### NOTE

Add as many key=value pairs as you need. Use a mix of capitalization, letters, numbers, symbols, and whitespace.

3. Save your changes and close the editor.
4. Optionally, generate an upload to Insights for Red Hat Enterprise Linux.

```
# insights-client
```

## 5.4.4. Using predefined system tags to get more accurate Red Hat Insights advisor service recommendations and enhanced security

Red Hat Insights advisor service recommendations treat every system equally. However, some systems may require a higher level of security than others, or require different networking performance levels. In addition to the ability to add custom tags, Red Hat Insights for Red Hat Enterprise Linux provides predefined tags which can be used by the advisor service to create targeted recommendations for your systems that might require more attention.

To opt in and get the extended security hardening and enhanced detection and remediation capabilities offered by predefined tags, you need to configure the tags. After configuration, the advisor service provides recommendations based on tailored severity levels, and preferred network performance that apply to your systems.

To configure the tags, use the **/etc/insights-client/tags.yaml** file to tag systems with predefined tags in a similar way that you might use it to tag systems in the inventory service. The predefined tags are configured using the same **key=value** structure used to create custom tags. Details about the Red Hat-predefined tags are in the following table.

**Table 5.1. List of Supported Predefined Tags**

Key	Value	Note
-----	-------	------

Key	Value	Note
security	<b>normal</b> (default) / <b>strict</b>	With <b>default</b> , the advisor service compares the system's risk profile to a baseline derived from the default configuration of the latest version of RHEL and from frequently-used usage patterns, keeping recommendations focused, actionable, and low in numbers. With the <b>strict</b> , value, the advisor service considers the system to be security-sensitive, causing specific recommendations to use a stricter baseline, potentially showing recommendations even on fresh up-to-date RHEL installations.
<b>network_performance`</b>	<b>null</b> (default) / <b>latency</b> / <b>throughput</b>	The preferred network performance (either latency or throughput according to your business requirement) would affect the severity of an advisor service recommendation to a system.



## NOTE

The predefined tag keys names are reserved. If you already use the key **security**, with a value that differs from one of the predefined values, you will not see a change in your recommendations. You will only see a change in recommendations if your existing **key=value** is the same as one of the predefined keys. For example, if you have a **key=value** of **security: high**, your recommendations will not change because of the Red Hat-predefined tags. If you currently have a **key=value** pair of **security: strict**, you will see a change in the recommendations for your systems.

### Additional resources

- [Using system tags to enable extended security hardening recommendations](#)
- [Leverage tags to make Red Hat Insights Advisor recommendations understand your environment better](#)
- [Custom system tagging](#)

### 5.4.5. Configuring predefined tags

You can use the Red Hat Insights for Red Hat Enterprise Linux advisor service's predefined tags to adjust the behavior of recommendations for your systems to gain extended security hardening and enhanced detection and remediation capabilities. This section describes how to configure the predefined tags.

## Prerequisites

- You have root-level access to your system
- You have Insights client installed
- You have systems registered within the Insights client
- You have already created the **tags.yaml** file. See [Creating a tags.yaml file and adding a custom group](#)

## Procedure

- Using the command line, open the **tags.yaml** configuration file located in **/etc/insights-client/** using your preferred editor. (The following example uses Vim.)

```
[root@server ~]# vi /etc/insights-client/tags.yaml
```

- Edit the **/etc/insights-client/tags.yaml** file to add the predefined **key=value** pair for the tags. This example shows how to add **security: strict** and **network\_performance: latency** tags.

```
# cat /etc/insights-client/tags.yaml
group: redhat
location: Brisbane/Australia
description:
- RHEL8
- SAP
security: strict
network_performance: latency
```

- Save your changes.
- Close the editor.
- **Optional:** Run the **insights-client** command to generate an upload to Red Hat Insights for Red Hat Enterprise Linux, or wait until the next scheduled Red Hat Insights upload.

```
[root@server ~]# insights-client
```

## Confirming that predefined tags are in your production area

After generating an upload to Red Hat Insights (or waiting for the next scheduled Insights upload), you can check whether the tags are in the production environment by accessing [Red Hat Enterprise Linux > Inventory](#). Find your system and look for the new tags. You should see something similar to what is shown in the following image.



Inventory > ruledev.jaylin.org

ruledev.jaylin.org

UUID: d2b5f00b-2691-4d61-8

Last seen: 03 Feb 2023 00:43

General information Adv

System properties

Host name ? ru

Display name ? ru

Name	Value	Tag source
group	redhat	insights-client
location	Brisbane/Australia	insights-client
security	strict	insights-client
description	RHEL8	insights-client
description	SAP	insights-client
network_performance	latency	insights-client

1 - 6 of 6 << < 1 of 1 > >>

system

RHEL 9.1

5.14.0

### Example of recommendations after applying a predefined tag

In the following image, the advisor service shows a system with the **network\_performance: latency** tag configured.

Recommendations Pathways ⓘ

> Name Filter by name Q

Systems impacted 1 or more x Status Enabled x Name NICs on Azure VMs x Reset filters

Name	Modified	Category	Total risk	risk of change	Syste...	Remediation
> NICs on Azure VMs encounter high network latency issue due to a known issue in the NETVSC driver	24 days ago	Performance	Important	Moderate	1	Playbook
> NICs on Azure VMs encounter network performance issue due to a known issue in the NETVSC driver	2 years ago	Performance	Moderate	Moderate	1	Playbook

Different "Total Risk"

The system shows a recommendation with a higher Total Risk that is categorized as Important. The system without the **network\_performance: latency** tag is categorized with a Total Risk of Moderate. You can make decisions about prioritizing the system with the higher Total Risk.

## CHAPTER 6. REFERENCES

See the following references materials to learn more.

### 6.1. REFERENCE MATERIALS

To learn more about the vulnerability service, see the following resources:

- [Remediating Security Exposures using the Vulnerability Service and Ansible Playbooks](#)
- [Generating Vulnerability Service Reports](#)
- [Red Hat Insights for Red Hat Enterprise Linux Documentation](#)
- [Red Hat Insights for Red Hat Enterprise Linux Product Support page](#)

# PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. To provide feedback, highlight text in a document and add comments.

## Prerequisites

- You are logged in to the Red Hat Customer Portal.
- In the Red Hat Customer Portal, the document is in the **Multi-page HTML** viewing format.

## Procedure

To provide your feedback, perform the following steps:

1. Click the **Feedback** button in the top-right corner of the document to see existing feedback.



### NOTE

The feedback feature is enabled only in the **Multi-page HTML** format.

2. Highlight the section of the document where you want to provide feedback.
3. Click the **Add Feedback** pop-up that appears near the highlighted text.  
A text box appears in the feedback section on the right side of the page.
4. Enter your feedback in the text box and click **Submit**.  
A documentation issue is created.
5. To view the issue, click the issue link in the feedback view.