



# An Open Approach to Vulnerability Management

Red Hat's Methodology

Version 1.4

September 20, 2022



# Contents

Introduction	<b>3</b>
Defining a vulnerability	<b>4</b>
How Red Hat reports and evaluates vulnerabilities	<b>4</b>
Common Vulnerabilities and Exposures (CVE)	5
Common Weakness Enumeration (CWE)	6
Common Vulnerability Scoring System (CVSS)	6
Temporal & Environmental analysis	9
CVSS scoring differences	10
Red Hat Severity Ratings	11
Mitigating vulnerabilities	<b>12</b>
Life cycle considerations	13
Known exploited vulnerabilities	14
Backporting and rebasing	15
Content signing	17
Red Hat Security Advisories (RHSA)	17
Red Hat management tools	18
RPM package tools: yum/dnf	18
Learning about vulnerabilities	<b>19</b>
Red Hat security data	19
Open Vulnerability and Assessment Language (OVAL)	19
Common Vulnerability Reporting Format (CVRP)	20
Red Hat Security Data API	20
Red Hat Container Health Index (CHI)	21
Red Hat Incident Response Plan	21
Red Hat Product Security Risk Report	22
Third-party security scanners	<b>22</b>
Conclusion	<b>23</b>
<b>About Red Hat</b>	<b>24</b>



## Introduction

Over the years, Red Hat has published a large number of articles, blogs and other resources that describe different facets of how we handle security vulnerabilities in our products. This document builds on those efforts with the aim to bring it all together and help our customers and communities better understand how Red Hat categorizes, addresses and responds to security vulnerabilities.

This document describes the current state of Red Hat's vulnerability management process. This process evolves over time, and this document will be updated to reflect that as appropriate. Red Hat follows the open source philosophy of continuous improvement, and that includes our efforts to improve how it addresses vulnerabilities.

Red Hat welcomes feedback and comments from its customers, partners and open source communities. As custodians of this document, please direct any feedback or comments to [Red Hat Product Security](#).



## Defining a vulnerability

Sometimes software is flawed. This can happen regardless of the development model, be it open source or otherwise, and even despite meticulous reviews, bugs happen. When a defect or bug has security implications, it is referred to as a vulnerability. A vulnerability in software is a weakness or absence of a safeguard resulting in an unplanned adverse outcome.

Identifying and analyzing these vulnerabilities are vital to protecting Red Hat product offerings used by our customers. [Red Hat Product Security](#) and its Product Security Incident Response Team ([PSIRT](#)) have been serving Red Hat, our subscribers, communities, and partners since [September 2001](#). Red Hat Product Security oversees over 400,000 components/versions that are included within currently supported products and cloud services. Detailed information about coverage and support for products within the Red Hat portfolio can be found on the [Product Life Cycles](#) page.

### References

- [The Source of Vulnerabilities, How Red Hat finds out about vulnerabilities](#)

## How Red Hat reports and evaluates vulnerabilities

Vulnerabilities<sup>1</sup> are identified using an industry-standard called [Common Vulnerabilities and Exposures \(CVE\)](#). Every security defect that impacts a component within the Red Hat portfolio has an assigned CVE identifier. Red Hat is a [CVE Numbering Authority \(CNA\)](#) for all Red Hat-branded software and also supplies CNA services for many open source projects. As of September 2022, Red Hat has also become a [Root CNA for open source projects](#).

Red Hat Product Security is a member of the Forum of Incident Response and Security Teams (FIRST) and participates in the [FIRST CVSS SIG](#). Red Hat uses the [Common Vulnerability Scoring System \(CVSS\)](#) industry standard as an additional measurement on each vulnerability we address. All CVEs impacting Red Hat products are issued a CVSS base score.

---

<sup>1</sup> In this document, we use *vulnerabilities* and *CVEs* interchangeably. We report on all security issues that are applicable to Red Hat software.



Every fixed CVE has a public entry in the [Red Hat CVE database](#) on the Red Hat Customer Portal. Red Hat Product Security also collects and analyzes more detailed technical information in a publicly accessible bug (Bugzilla or Jira).

Every vulnerability reported to Red Hat Product Security is reviewed and analyzed by our team of open source software security specialists. These engineers understand how our offerings are composed, curated, hardened, packaged, delivered, and used by our customers. Their breadth of knowledge and experience on security-focused supply chain practices provide critical insights into potential impacts of these vulnerabilities on our products and services.

Additionally, Red Hat associates may find and report vulnerabilities in open source software to Red Hat Product Security, who then coordinate with other vendors as appropriate.

For every vulnerability impacting a component of our portfolio, Red Hat tracks the issue with a CVE identifier, issues a CVSS base score, and provides a Red Hat Severity Rating. We also identify and classify the type of vulnerability using an industry-standard called [CWE](#) (Common Weakness Enumeration).

## Common Vulnerabilities and Exposures (CVE)

The goal of CVEs is to assign a common identifier to a vulnerability that informs all hardware and software users of a unique problem that could impact one or more technical implementations of that affected component. It is common to see one CVE impact multiple vendors since they are potentially sourcing their components from the same upstream supplier. A CVE has the following format:

`CVE-XXXX-YYYY`

With XXXX being the year the CVE was issued, and the YYYY being a unique number issued by the relevant CVE Numbering Authority (CNA). This identifier, referenced by software suppliers, in a security advisory or bulletin notifies end-consumers of the vulnerability existing within a particular product or service requiring end-consumer action. Vulnerability aggregators and third-party security scanners leverage these CVE IDs as part of their processes.



## References

- [Red Hat CVE Database](#)
- [New and Improved CVE Pages](#)
- MITRE's [CVE site](#)
- [Red Hat extends CVE program expertise as newly-minted Root organization](#)

## Common Weakness Enumeration (CWE)

Common Weakness Enumeration (CWE) is a community-developed list of common software and hardware weakness types that have security ramifications. "Weaknesses" are flaws, faults, bugs, or other errors in software or hardware implementation, code, design or architecture that if left unaddressed could result in systems, networks or hardware being vulnerable to attack. Every vulnerability that impacts components provided by Red Hat will be evaluated using CWE for subsequent root cause analysis. CWE is described in the following format:

*CWE-ZZZ*

With ZZZ being a unique identifier for a type of weakness as described in the [CWE List](#).

## References

- [Red Hat is now CWE Compatible](#)
- MITRE's [CWE site](#)

## Common Vulnerability Scoring System (CVSS)

CVSS conveys how a particular vulnerability works and what aspects of the information security triad—confidentiality, integrity, and availability (CIA)—are impacted by the flaw. CVSS is [not a measurement of risk](#), Red Hat Product Security uses CVSS as a portion of our more holistic approach to vulnerability assessment and their impact on our software portfolio. Red Hat Product Security also conducts other assessments such as developing reproducers, analyzing the impact on layered products, and determining if the exploitability of the flaw is mitigated or reduced due to our build and compiling practices.

The CVSS standard provides the ability for an individual to analyze a security defect through the use of a series of metrics that help describe certain aspects of the flaw. It is divided up into a Base Metrics, and two optional analyses - [Temporal and Environmental](#).



CVSS:Q/AV:vv/AC:vx/PR:xx/UI:xy/S:yy/C:z/l:z/A:z

The CVSS v3 Base Metric group covers the constant aspects of a vulnerability:

- Attack Vector (AV) - Expresses the proximity to the vulnerable system required for an attack and how the vulnerability is exploited.
- Attack Complexity (AC) - Speaks to the difficulty of executing an attack and what factors are needed for it to be successful.
- User Interaction (UI) - Determines whether the attack requires an active human to participate or if the attack can be automated.
- Privileges Required (PR) - Documents the level of user authentication required for the attack to be successful.
- Scope (S) - Determines whether an attacker can impact a component beyond its security scope/authority.
- Confidentiality (C) - Determines whether unauthorized parties can access information resources and to what extent.
- Integrity (I) - Measures the impact to the trustworthiness and veracity of data.
- Availability (A) - Measures the impact on authorized user access to data or services .

A formula translates these measurements into a single, numerical base score, ranging from 0.0 (no impact) to 10.0 (highest base impact). Refer to [Common Vulnerability Scoring System v3.1: User Guide](#) for detailed descriptions of the Base Metrics. It is important to note that the CVSS Base Metrics were designed to be used with the other CVSS metric groups, notably the Temporal and Environmental Metrics, to provide an accurate representation of risk in customer environments. Alone, the Base Metrics offer a shallow view of the vulnerability itself without accounting for deployment or the environment.

As described in the [Common Vulnerability Scoring System v3.1: Specification Document](#) (emphasis added for clarity):

*“The Common Vulnerability Scoring System (CVSS) captures the **principal technical characteristics** of software, hardware and firmware vulnerabilities. Its outputs include numerical scores indicating the **severity of a vulnerability relative to other vulnerabilities.**”*



*“Base Scores are usually **produced by the organization maintaining the vulnerable product**, or a third party scoring on their behalf. It is typical for only the Base Metrics to be published as these do not change over time and are common to all environments. **Consumers of CVSS should supplement the Base Score with Temporal and Environmental Scores specific to their use of the vulnerable product to produce a severity more accurate for their organizational environment.** Consumers may use CVSS information as input to an organizational vulnerability management process **that also considers factors that are not part of CVSS** in order to **rank the threats** to their technology infrastructure and make informed remediation decisions. Such factors may include: number of customers on a product line, monetary losses due to a breach, life or property threatened, or public sentiment on highly publicized vulnerabilities. These are outside the scope of CVSS.”*

The use of CVSS, particularly the CVSS Base Score alone, in risk assessments has been misunderstood and misused in the industry for a long time. As a result, the CVSS v3.1 specification was updated to address this problem. As described in the [CVSS User Guide for changes in version 3.1](#) (emphasis added for clarity):

*“The CVSS Specification Document has been updated to emphasize and clarify the fact that **CVSS is designed to measure the severity of a vulnerability and should not be used alone to assess risk.**”*

*Concerns have been raised that the CVSS Base Score is being used in situations where a comprehensive assessment of risk is more appropriate. The CVSS v3.1 Specification Document now clearly states that the CVSS Base Score represents only the intrinsic characteristics of a vulnerability which are constant over time and across user environments. The CVSS Base Score should be supplemented with a contextual analysis of the environment, and with attributes that may change over time by leveraging CVSS Temporal and Environmental Metrics. More appropriately, **a comprehensive risk assessment system should be employed that considers more factors than simply the CVSS Base Score.** Such systems typically also consider factors outside the scope of CVSS such as exposure and threat.”*

Further information can be found on the [Understanding Red Hat security ratings](#) page.

It is important to note that Red Hat can only provide the Base Metrics which yield a Base score based upon default settings and most likely deployment scenarios. This becomes a starting





point for organizations to determine the real impact of a particular vulnerability to them, taking into account their own implementation, applications deployment, and risk environment.

Each organization is unique, with its own requirements and challenges, and all risks are not created equal, nor are they the same across companies. Consequently, during the passage of time knowledge of the characteristics of a vulnerability may change (for example, an automatable exploit may be available six months after disclosure rather than the day of disclosure). Thus, it is important for organizations to augment the Base scoring by using the Environmental and Temporal Metrics to better reflect the characteristics and risks a vulnerability may represent.

While commonly thought of as optional, the Environmental and Temporal Metrics are critical to using CVSS as any kind of risk-related score. Absent that, the Base score alone must only be used as a means of prioritizing which vulnerabilities to focus on mitigating, when mitigations are available. It must not be used as a measurement of risk without using these additional metrics that require user-input as to the environment and point-in-time potential exploit availability. Along with providing a Red Hat CVSS score, Red Hat provides a Red Hat Severity Rating (discussed below) for all vulnerabilities impacting our products. This is our primary guidance for our customers.

#### **References:**

- [Understanding Red Hat security ratings](#)
- [How Red Hat uses CVSSv3 to Assist in Rating Flaws](#)

## Temporal & Environmental analysis

The Temporal and Environmental reviews are important yet frequently overlooked areas of the CVSS analysis. These are methods that end-users should be aware of and use in their own Risk and Vulnerability Management programs. Temporal review allows for the Base Metric score provided by a vendor, such as Red Hat, to be modified based on details around current exploitation techniques, the existence of attacks leveraging the vulnerability, or the availability of patches or workarounds for the defect.

The other more important measurement is the Environmental Metrics. This is where the practitioner can add in organizational-specific details about mission-critical data, systems or controls that might exist in the end-consumer's environment that could alter the impact or probability of an attack being successfully executed.



CVSS Base Metric scores are generic and based on default or most-common configurations. They are *not* tailored to any one organization's configuration, sensitive data or systems, controls, regulatory or legal obligations, nor risk appetites. Consumers are always advised to conduct their own assessment of CVEs using all available data to inform their risk calculus.

## CVSS scoring differences

Red Hat Product Security is the authoritative source for vulnerability data and scoring information for Red Hat products, services and the components that comprise them. A qualified engineer with direct technological experience reviews each CVE. The scores Red Hat provides are based on how our software is selected, compiled, built and configured **at delivery**. The scoring reflects actual data and testing wherever possible. Each product component may be impacted differently by a specific vulnerability, so it is possible to see varying CVSS scores between different offerings and even between differing versions.

The Red Hat portfolio is based in large part on upstream open source software. As part of our productization processes, changes are made to make that code easier to digest for enterprises. Therefore, **the software Red Hat provides is not necessarily identical to what could be obtained or used from upstream**. Vulnerability analysis on those upstream projects and components are often not always directly applicable to a Red Hat offering.<sup>2</sup> As a result, the CVSS scores of an upstream package often differ from those for a Red Hat product.

While popular, third-party vulnerability aggregators, such as the [National Vulnerability Database](#) (NVD), are **not authoritative** around how a given issue can impact a component or product. Red Hat is a member of the CVE Board and CVSS Working Groups. We make every effort to work with our industry partners, peers and entities like NVD, but **a user should double-check any conclusion found in the NVD (or by any other aggregator) with the CVSS score and metrics determined by Red Hat** as shown in our [CVE database](#) for any particular issue.

Depending on how issues are discovered or reported, sometimes NVD and other aggregators may report different scoring information. Red Hat Product Security takes this seriously and actively works with organizations like MITRE (which maintains NVD) to provide the appropriate technical details. We actively collaborate with them when there are unique differences with a Red Hat implementation of a package, library or component.

---

<sup>2</sup> Conversely, vulnerability analysis on Red Hat products is not always applicable to upstream projects.



## References

- [Security flaws and CVSS rescore process with NVD](#)
- [Security flaws mitigated by compiler optimizations](#)

## Red Hat Severity Ratings

Red Hat Product Security uses a [four-point scale](#) to describe a particular bug’s severity based on rigorous analysis of the flaw. We designed this scale to align closely with similar scales used throughout the industry by other vendors and upstream open source communities. Our intent for the Red Hat Severity Rating is to help users determine which issues could pose more risk.

Ideally, this prioritized risk assessment helps customers understand how they may be exposed and enables them to better schedule updates to the systems they manage. We recognize that each business is unique, with its own requirements and challenges, and that all risks are not created equal, nor are they the same across companies.

The four-point scale rates vulnerabilities as Low, Moderate, Important or Critical. Critical vulnerabilities pose the most severe threat to an organization. As described in our rating methodology, a Critical vulnerability could be exploited remotely over a network (or the internet) or be automated in an attack, such as by a worm. Like many of our peers, we expand this definition to include flaws that affect web browsers or browser plug-ins that users might be susceptible to from malicious or compromised websites.

CRITICAL	IMPORTANT	MODERATE	LOW
A remote unauthenticated user can execute arbitrary code  Does not require user interaction  i.e. Worms	Allows local users to gain privileges  Unauthenticated remote users can view resources  Authenticated remote users can execute arbitrary code	Vulnerabilities are more difficult to exploit  Are exploitable via an unlikely configuration	Unlikely circumstances required to exploit  Impact is of minimal consequence



When Red Hat Product Security reviews a flaw, we look at how the software is sourced, built, packaged and deployed. A CVSS Base score for Red Hat software assumes our products are used as designed, with security-focused defaults and settings in place. If subsequent changes are made to system settings or security controls, system administrators must account for that as they evaluate the risk a vulnerability might pose inside their unique environments.

It is important to remember that no vendor can tell a business what is important to them nor dictate actions to take to protect their sensitive data. CVSS and the Red Hat Severity Rating are baselines of our software, a starting point for consumers to begin their own risk assessment. It is also important to note that CVSS Base scores do not influence the Red Hat Severity Rating; CVSS is used as a guide to assist with understanding a vulnerability. The Red Hat Severity Rating is Red Hat's standardized rating that speaks to the risk posed by a vulnerability.

## References

- [Red Hat Severity Ratings](#)
- [What does the severity rating in the security advisory mean?](#)

# Mitigating vulnerabilities

Knowing that there is a vulnerability is only a first step. It is essential to have a software lifecycle management process in place and stay alert to key documentation, fixes/patches and mitigations/workarounds. Red Hat provides a multitude of ways to not only stay up to date about what advisories have been issued, but also provides several simple paths to access signed and authorized updates. There are a few concepts that relate to all patches, regardless of the product, service or technology in scope.

Additionally, Red Hat's Customer Experience and Engagement (CEE) organization is available to help customers with questions about patched, or unpatched, vulnerabilities. If customers ever have questions about exposure, unfixed vulnerabilities they feel affect them, or other security-related questions, they are encouraged to reach out to their support representative.

## References

- [Red Hat CVE Database](#)
- [Red Hat Security Advisory Database](#)
- [Red Hat Product Security Center](#)
- Sign up for the [RHSA email list](#)



- [OVAL Data feed](#)
- [CVRF/CSAF Data feed](#)
- Red Hat Security Vulnerability [Data API](#)

## Life cycle considerations

Red Hat provides publicly available life cycle pages that describe how we support our software products during their lifetimes. There are three primary phases common to most offerings: Full Support, Maintenance Support, and then the Extended Life phase. The level of support subscribers can expect varies depending on which phase of the life cycle a product is in. However, there are a set of common expectations.

During all phases of a product's support life cycle, [Red Hat-rated](#) Critical and Important Severity issues will be addressed, typically asynchronously (outside of a scheduled major or minor release). Note that this may vary depending on the particular affected software and upon the complexity of the patch itself, taking into account whether the update is a backport, or a rebase, described in more detail in the following section. Refer to the [Red Hat Risk Report](#) for an understanding of how quickly vulnerabilities have been addressed.

The applicability of Critical and Important fixes is to all phases of a life cycle, be it Full Support, Maintenance Support or the Extended Life phase. What is generally out of scope are Low and Moderate fixes. Red Hat-rated Low Severity issues are those for which the impact of successful exploitation is minimal, and the odds of exploitation are deemed small. Low Severity issues may be addressed with other fixes for more severe issues at the next major or minor release, or at Product Engineering's discretion. Typically, Low Severity security fixes may be treated with the same urgency as bug fixes.

For example, perhaps a Low Severity flaw is only exposed to existing users with administrative privileges, meaning that *without* administrative privileges already in place, the flaw is not exploitable. This flaw could then only be exploited through a compromise of an actual administrative user account by an attacker, or if the user *were* the attacker (commonly referred to as "insider threat"). In a situation like this, the potential risk posed due to the presence of the flaw is itself insignificant. More significant is the potential exposure of an already-compromised administrative account, which would not need to use this flaw to obtain additional privilege, or that of a malicious insider who had somehow been granted administrative privileges.



When vulnerabilities of this kind are reported, they are assigned CVE names for completeness and tracking. Because they are unlikely to have a significant impact or pose any real risk, fixes are not often produced. **Implementing such a fix may actually *introduce more risk, including, e.g., incompatibilities with other software or functionality, instability of the running software, additional dependencies or libraries, and the possible introduction of bugs or new vulnerabilities.***

Vulnerabilities rated Moderate are treated on a case-by-case basis. Some Moderate vulnerabilities may be cause for concern and proactively addressed by Red Hat. Other Moderate vulnerabilities may be assessed to be less concerning and not be fixed.<sup>3</sup>

An example might be a vulnerability in Red Hat Enterprise Linux (RHEL) that has little bearing on its own but is used in such a way as having greater impact on Red Hat OpenShift . In these instances, a proactive approach is likely to be taken to mitigate the vulnerability on Red Hat OpenShift through the patching on RHEL. Typically, Moderate issues are corrected in an upcoming major or minor release and not asynchronously, as with Critical and Important fixes, at the discretion of Product Engineering teams.

## References

- [Red Hat Product Life Cycles page](#)
- [Red Hat Risk Report \(2020\)](#)

## Known exploited vulnerabilities

In addition to proactively addressing all Critical and Important vulnerabilities across the product life cycle, Red Hat Product Security tracks reports of vulnerabilities known to be actively exploited. Data shows the majority of exploited vulnerabilities rated Critical and Important are typically addressed, prior to being actively exploited.

As part of our risk-based approach to security, Product Security continues to monitor reports of exploitations in the wild. If a report is received that a vulnerability is actively being exploited and has not previously been fixed across the Red Hat portfolio, remediation is prioritized.

---

<sup>3</sup> Low and Moderate fixes are not guaranteed to be fixed in any phase of a product's life cycle.



Vulnerabilities with a Critical or Important severity, by the nature of the rating system, are the most cost effective targets for attackers, which is why addressing these vulnerabilities are prioritized. However when exploitation of a flaw is observed in the wild, a fix will be prioritized regardless of severity rating, reflecting our proactive risk-based approach.

In 2021, we began tracking exploit information, starting with the Cybersecurity and Infrastructure Security Agency (CISA) known exploited vulnerabilities catalog. Based on that list, out of 1596 vulnerabilities reported in 2021, only 26 were actively exploited (1.6%). Out of the 10 Critical vulnerabilities in 2021, only 1 was listed as actively exploited (10%) and of the 283 Important vulnerabilities in the same time period, only 7 were listed as actively exploited (2.5%)<sup>4</sup>.

## References

- [2021 Red Hat Product Security Risk Report](#)
- [CISA known exploited vulnerabilities catalog](#)

## Backporting and rebasing

As a commercial open source software provider, Red Hat derives most of its software from upstream open source repositories and communities. To provide enterprise-ready stability, we use two techniques to provide updates of software: backporting upstream code and updating to newer upstream versions (also known as rebasing).

“Backporting” is performed when a particular software feature, enhancement or fix is taken from a newer software version and applied to an older version of the same software. This is done to minimize additional code changes found in the newer upstream version. A common example is that Red Hat will backport features and enhancements from newer upstream kernel code into our stable Enterprise kernel.

There are two primary reasons for opting to backport distinct changes to current versions rather than use the new upstream version: the first is to ensure Application Programming Interface (API) and Application Binary Interface (ABI) compatibility with other software that depends on the component being updated; the second is to reduce the risk of introducing new potential vulnerabilities that may be present in other features of newer upstream versions and could adversely affect currently supported versions.

---

<sup>4</sup> Refer to the 2021 Red Hat Product Security Risk Report’s [“Known exploits of vulnerabilities in 2021”](#) section for further details.



The latter reason is important as backporting fixes have been demonstrably proven to reduce the risk of new, unknown vulnerabilities. And because the amount of code changed is limited to specific fixes or features, other bugs that may be present in upstream versions are **not** likely to appear in older versions. Often bugs of significance have their fixes backported to currently supported versions as well.

However, backporting often confuses some third-party scanners that perform simple version checks for vulnerable versions, a topic discussed more in the [Third-party scanners](#) section. Every updated package increments the *release* number, which is a mechanism to determine whether or not a package is newer or older than a vulnerable release. This is often referred to as the N-V-R (Name-Version-Release) of a package.<sup>5</sup>

As vulnerabilities are discovered in components that Red Hat provides and it is determined that the shipped software is not affected, that information is provided through our CVE pages and associated security metadata.

Often in major releases, when significant new features are included, or when backporting is not practical, Red Hat will "rebase" a package to a new upstream version. Rebasing is when the existing package is replaced with a newer version from upstream, thus future releases and backports are "based" on this new version. When a package is rebased, the full version number, rather than just the Red Hat release number, is updated to reflect the new upstream version.

---

<sup>5</sup> For example, `openssl-1.1.1d-4el8` would indicate the "openssl" package name, a "1.1.1d" version, and a "4el8" release, indicating the fourth release of this version for RHEL 8. A subsequent update to the package that does not change the version of the software would have a "5el8" release.





## References

- [What is Red Hat's security patch and backport practice?](#)
- [Security Backporting Practice](#)
- [What is backporting and how does it affect Red Hat Enterprise Linux \(RHEL\)?](#)
- [Is your software fixed?](#)
- [Security flaws on unsupported products or products with limited support](#)

## Content signing

It is critical for consumers to be able to verify that the software they are using is authentic and untampered with. All RPM-based<sup>6</sup> and container image content is signed using the authorized Red Hat signing server. End-users should always check to verify the software they have downloaded and are about to install is genuine and authentic, only using known trusted sources for any software installed within their environment.

## References

- [Product Signing Keys](#)
- [How to sign rpms with GPG](#)
- [Securing RPM signing keys](#)
- [Verifying image signing for Red Hat Container Registry](#)
- [How to test verifying image signatures?](#)
- [Red Hat Satellite Signing Packages Guide](#)

## Red Hat Security Advisories (RHSA)

Red Hat publishes several forms of advisories. Red Hat Security Advisories (RHSA) are published whenever an update to a product contains a security fix. Any RHSA can include fixes for multiple CVEs, and as such, always inherits the highest Red Hat Severity Rating of the CVEs being corrected.

Red Hat publishes our advisories over numerous channels directly to subscribers and the larger open source community.

---

<sup>6</sup> Content delivered using the RPM Package Manager (RPM) format, such as for Red Hat Enterprise Linux and other products.



## References

- [Red Hat Security Advisories database](#)
- [Explaining Red Hat Errata \(RHSA, RHBA, and RHEA\)](#)
- [The RHSA notifications you want, right in your Inbox](#)
- [Anatomy of a Red Hat Security Advisory](#)
- [Red Hat Product Errata Advisory Checker](#) (Customer Portal login required)

## Red Hat management tools

Red Hat's portfolio includes several products and services that help simplify the management of your Red Hat assets. These tools are tightly integrated with the data provided by Red Hat Product Security and offer different paths to understanding where vulnerabilities might lie in a customer's portfolio and how it can address them.

## References

- [Red Hat Satellite](#)
- [Red Hat Ansible](#)
- [Red Hat Insights](#)
- [Clair](#)
- [Red Hat closes acquisition of StackRox](#)
- [Creating a central patch management with ansible](#)
- [Managing the security of your Red Hat Enterprise Linux environment with Red Hat Insights](#)
- [Insights Security Hardening Rules](#)

## RPM package tools: yum/dnf

Many of Red Hat's offerings leverage RPM packages. This format and the associated utilities offer some unique capabilities in regards to vulnerabilities. For example, customers using RHEL may choose to only install security updates and ignore any non-security updates that may be available.



## References

- [Is it possible to limit yum so that it lists or installs only security updates?](#)
- [How do I check if a specific kernel is vulnerable to a specific CVE?](#)
- [How do I check the changes of a proposed package update?](#)
- [Can I install/run packages from different versions of RHEL?](#)
- [How to use yum to download a package without installing it](#)

# Learning about vulnerabilities

## Red Hat security data

All materials related to vulnerabilities impacting the Red Hat portfolio are publicly available after the vulnerability has been publicly disclosed. Our data is published on our [award-winning Customer Portal](#), and also in several industry-standard human- and machine-readable formats. Red Hat uses two well recognized methods (OVAL and CVRF, described below) to provide this data.

## References

- [Security Data](#)
- [Understanding Red Hat products' vulnerabilities](#)

## Open Vulnerability and Assessment Language (OVAL)

Red Hat has been heavily involved in providing our customers and open source communities with access to security data since 2002, when we became a founding board member of [Open Vulnerability and Assessment Language](#) (OVAL). Red Hat [announced OVAL compatibility in 2006](#). OVAL is an industry-recognized format for sharing vulnerability information in a consistent way. This data is freely published and available so any third-party scanning tool can leverage the data and have the ability to understand security flaws that exist on a system being scanned.



Red Hat provides two feeds for consumers to obtain information in the OVAL format. The first is a “classic” implementation, as originally intended, that describes all issues that are fixed within current versions of Red Hat Enterprise Linux. The second, more modern use case is the “v2” feed that supplies that same information and also augments it with information regarding issues that are known and unfixed, across the portfolio. We recommend that users consume the OVALv2 feed as, while OVALv1 is available for backwards compatibility to those already consuming it, the v2 feed is far more comprehensive.

### References

- [What is OVAL and how can I use it to learn about security issues?](#)
- [Evolving OVAL](#)

## Common Vulnerability Reporting Format (CVRF)

The Common Vulnerability Reporting Format is another industry-recognized standard that Red Hat Product Security helps to support and participates in its governance. It is a machine-readable data exchange format for supplying vulnerability advisory information.

### References

- [Red Hat Security Advisories in CVRF](#)

## Red Hat Security Data API

Red Hat has published information about vulnerabilities affecting our product portfolio since [1999](#). Since that time, delivery and formats have changed based on technological standards and consumer-demand. In 2016 we launched the Red Hat Security Data API that enables customers to interact with our security data through a modern API.

### References

- [Vulnerability API Blog](#)
- [Vulnerability API documentation](#)



## Red Hat Container Health Index (CHI)

Containers are an architectural format (gaining widespread usage) that enables cloud computing and agile software development. Red Hat OpenShift is a Kubernetes-based platform supported by the solid foundation of Red Hat Enterprise Linux and [SELinux](#).

Containers are a different deployment and delivery architecture than traditional Linux-based platforms and require new methods and tools to monitor, assess, and address security vulnerabilities with the ability to manage updates to that infrastructure.

The Red Hat Container Catalog is a platform that provides container images. The Container Health Index (CHI) is a means to provide security information and a “freshness” score to consumers to understand how up-to-date the containers they desire to use are. It is important to note that the CHI is based on unapplied, yet available, security fixes to the underlying products and does not account for vulnerabilities present for which there is no fix. At times it may seem that the CHI grade is at odds with what a third-party container scanner software provides; yet, in all likelihood they probably are not inconsistent. They both provide different views to container health -- the CHI tells you whether the container has the latest security fixes provided by Red Hat, whereas the container scanner will typically inform whether patches are available and unapplied *and* if there are vulnerabilities present for which there are no fixes.

### References

- [Container Health Index grades as used inside the Red Hat Container Catalog](#)
- [Security Scoring and Grading for Container Images](#)
- [Resources on Red Hat Container Security](#)
- [Scanning pods for vulnerabilities - OpenShift Container Platform 4.5 Security](#)
- [The OpenShift Security Guide Book Download](#)

## Red Hat Incident Response Plan

Managing vulnerabilities can be challenging for operators and administrators. Every year thousands of potential threats need to be sorted through, prioritized and corrected. Over recent years, researchers have taken to “branding” their flaws in an attempt to gain notoriety and attention to their findings. This dubious practice creates a great deal of anxiety and extra effort on behalf of organizations. Red Hat created the Incident Response Plan (IRP) containing the coordination process to help provide concise, clear advice and remove FUD, drama and hyperbole from the process of dealing with these “celebrity” vulnerabilities.



## References

- [Understanding Red Hat's Product Security Incident Response Plan](#)
- [Security Bulletins](#)

## Red Hat Product Security Risk Report

Red Hat publishes an annual Product Security Risk Report detailing the vulnerabilities and threats that arose during the calendar year that impacted the Red Hat portfolio. These reports detail a volume of issues and contextualizes them within the scope of events that occurred throughout the year.

## References

- [2021 Red Hat Product Security Risk Report](#)
- [The history of open source risk reporting](#)
- [Red Hat Risk Report: A tour of 2020's branded security flaws](#)

## Third-party security scanners

Understanding what vulnerabilities exist in any organization's environment is a critical, yet daunting task. Companies rarely run one technology alone throughout their enterprise, so they often rely upon third-party security scanners that can detect flaws across multiple technologies. If these tools use the appropriate vulnerability data, they can give a more accurate picture of actions that need to be taken by patch management staff. But when they do not, the scanning tools can create quagmires of false-positives and inaccuracies that distract administrative staff and slows down addressing the meaningful issues.

Red Hat has a new certification program for third-party security scanners. Using these certified tools assures our customers they are getting the most precise and accurate data possible from their vulnerability scanner. Red Hat welcomes additional participants in this certification program. The [Red Hat Vulnerability Scanner Certification](#) program was initiated in February 2021 to address inconsistencies in third-party scanners scanning Red Hat products.

Many scanning tool products in the market use questionable or oftentimes incomplete sources of data. By contrast, certified scanners use authoritative data produced by Red Hat Product Security which enables them to show users the Severity Ratings and CVSS scores produced by



Red Hat as experts on Red Hat software who understand backported software correctly. Thus, a user of a certified scanner that properly employs this data can see the right information in the right way -- few, if any, false positives and, more importantly, reduced false negatives. By using Red Hat's OVALv2 data, the certified scanner can know what software has fixes available and which vulnerabilities for which no patch remediation currently exists.

## References

- [Red Hat Vulnerability Scanner Certification](#)
- [Introducing Red Hat Vulnerability Scanner Certification](#)
- [Third-party Security Ratings and Backporting](#)
- [Determining your risk](#)
- ["To be, or not to be," vulnerable... How customers and partners can understand and track Red Hat security vulnerabilities](#)

## Conclusion

This paper has described the frameworks, standards, techniques and tools Red Hat uses around managing vulnerabilities discovered within components of our software portfolio.

For further information, please visit the [Red Hat Product Security Center](#).

Finally, if you have a question about a specific security vulnerability, or believe that you may have knowledge of such a vulnerability, please contact us at one of the methods listed on our [Security Contacts and Procedures](#) page.



## About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

Copyright © 2022 Red Hat, Inc. Red Hat and the Red Hat logo are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.