Mandiant

HOW MANDIANT RATES VULNERABILITIES

Introduction

Mandiant believes that effective vulnerability analysis is a combination of structured or algorithmic analysis and human analysis. Ideally, both should be used to capture and quantify the vulnerability's potential impact and the true threat that it poses to organizations. Use of structured algorithmic techniques, which are common in many models, allows for consistent and transparent rating levels, while the addition of human expert analysis allows experts to integrate factors that are difficult to quantify and tweak ratings based on real-world experience regarding the actual risk posed by various types of vulnerabilities. Vulnerability ratings start with the consequences of exploiting a vulnerability to provide a baseline risk level, then evolve based on factors to increase or decrease the risk the vulnerability poses (such as ease of exploitation, observed exploitation in the wild, and availability of exploit code). The final rating enumerates all considered factors to provide transparency, including any additional qualitative information provided by expert judgment from trained analysts.

Threat Detail

Mandiant's vulnerability ratings are primarily focused on increasing the ability of decision makers and on-the-ground security teams to better prioritize patching and mitigation based not only on the theoretical impacts of a given vulnerability, but how they could and are being exploited in the real world. These ratings are our expert assessment of the current threat a vulnerability poses to organizations, the impact an attacker could have on a targeted organization if they were to exploit a vulnerability, and the ease or reliability with which an actor could exploit the vulnerability. The goal of providing this information is to support resource prioritization for patch, mitigation, and version management within an organization. Our Exploitation State is intended to be taken as a primary factor in prioritization, since what an actor is realistically able to do and how they are able to do it in a real-world environment can drastically change the threat posed by a vulnerability. Our Risk Rating, which is primarily based on impact, should then be used to inform which other vulnerabilities are particularly noteworthy.

For more information about exploring vulnerabilities in the Mandiant Advantage Threat Intelligence (MATI) platform, see **How to Explore Vulnerabilities** (https://docs.mandiant.com/home/mati-how-to-explore-vulnerabilities).

	Mandiant Risk Ratings					NVD CVSSv3.x Ratings				
	2020	2021	2022	2023	%	2020	2021	2022	2023	%
Unrated	1	24	38	373	0.5%	85	251	307	986	1.8%
Low	11674	13181	15694	18799	63.7%	391	425	486	414	1.9%
Medium	5793	6570	7604	7703	29.8%	7428	8524	10134	12585	41.5%
High	652	674	1801	2417	6.0%	7570	8587	9984	10685	39.6%
Critical	2	2	0	0	< 0.1%	2648	2664	4226	4622	15.2%

Mandiant Vulnerability Ratings Defined

While Mandiant uses standard terminology to rate vulnerabilities, the resulting ratings of "low," "medium," "high," and "critical" vulnerabilities tends to differ compared to CVSSv2.0 (which does not contain a score range associated with the "critical" rating) and CVSSv3.1. Specifically, our ratings generally result in a lower proportion of high ratings and a much smaller proportion of critical ratings. We believe that reserving "critical" designations assists in the prioritization of mitigation efforts by identifying those vulnerabilities that require the most immediate remediation.

Our ratings are defined as the following:

- **Exploitation State**: Mandiant's Exploitation State indicates what is occurring in the wild in terms of exploitation-related activity. This may drive prioritization when operationalizing the vulnerability reporting structure into a Security Operations Center (SOC).

Rating	Description
No Known	No known exploitation activity, underground discussions, proof of concept (PoC) or exploit code, but has a low potential for exploitation.
Available	Exploit or PoC code is publicly available or underground discussions, alleged selling, or alleged privately held code observed.
Confirmed	Limited reported or confirmed exploitation activities.
Wide	Exploitation has been reported or confirmed to widely occur.

MANDIANT

- **Risk Rating**: Mandiant's Risk Rating is our expert assessment of the impact an attacker could have on a targeted organization if they were to exploit a vulnerability.

Rating	Description	Examples
LOW	Exploitation would have little to no security impact on targeted systems. Reliability of exploitation is likely low and unlikely to be performed on a wide scale.	• leak or modification of trivial information • temporary or local denial-of-service (DoS) • cross-site scripting (XSS)
MEDIUM	Exploitation would enable attackers to perform activities on the targeted device or network or could allow attackers to have a direct impact on the security of the targeted device or network but would require additional steps. Reliability of exploitation is questionable and may or may not be able to be performed on a wide scale.	• leak or modification of non-trivial information • persistent DoS • escalation of privileges • enablement of man-in-the-middle (MitM) attacks • security bypass • execution of arbitrary code with notable mitigating factors
HIGH	Exploitation would enable attackers to have a notable direct impact to the security of targeted devices and networks without needing to overcome any major mitigating factors. Reliability of exploitation is expected to be high and can typically be done on a wide scale.	• leak or modification of information obtained from key assets or systems • permanent DoS • path traversal • execution of arbitrary code with little to not mitigating factors
CRITICAL	Exploitation would fundamentally undermine the security of affected devices and networks, enable actors to perform significant attacks with minimal effort, impact a wide number of systems, often with little to no mitigating factors to overcome. Reliability of exploitation is most likely very high and can almost certainly be performed effectively at scale.	• CVE-2017-5638 • CVE-2019-19781

MANDIANT

We intentionally use the critical rating sparingly, most often in cases where exploitation has serious impact, exploitation is trivial with often no real mitigating factors, and the attack surface is large and remotely accessible. When Mandiant uses the critical rating, it is an indication that remediation should be a top priority for an organization due to the potential

impacts and ease of exploitation.

For example, Mandiant rated CVE-2019-19781 as critical due to the confluence of widespread exploitation, including by APT41, the public release of proof-of-concept (PoC) code that facilitated automated exploitation, the potentially acute outcomes of exploitation, and the ubiquity of the software in enterprise environments.

CVE-2019-19781 is a path traversal vulnerability of the Citrix Application Delivery Controller (ADC) 13.0 that, when exploited, allows an attacker to remotely execute arbitrary code. Due to the nature of the systems, successful exploitation could lead to further compromises of a victim's network through lateral movement or the discovery of Active Directory (AD) and/or LDAP credentials. Though the credentials are often stored in hashes, they have been proven to be vulnerable to password cracking. Depending on the environment, the potential second-order effects of exploitation of this vulnerability could be severe.

At the time, we described widespread exploitation of CVE-2019-19781 in our blog titled **Rough Patch: I Promise It'll Be 200 OK (https://www.mandiant.com/resources/blog/rough-patch-promise-it-will-be-200-ok)**, including a timeline from disclosure on Dec. 17, 2019 to the patch releases, which began a little more than a month later on Jan. 20, 2020. Significantly, within hours of the release of PoC code on Jan. 10, 2020, we detected reconnaissance for this vulnerability in Mandiant telemetry data. Within days, Mandiant observed weaponized exploits used to gain footholds in victim environments. On the same day the first patches were released, Jan. 20, 2020, Mandiant observed APT41, one of the most prolific Chinese groups we track, kick off an expansive campaign exploiting CVE-2019-19781 and other vulnerabilities against numerous targets (see **This Is Not a Test: APT41 Initiates Global Intrusion Campaign Using Multiple Exploits (https://www.mandiant.com/resources/blog/apt41-initiates-global-intrusion-campaign-using-multiple-exploits)**).

Mandiant doesn't make recommendations on specific sevice-level agreements (SLAs). It is the sole responsibility of the organization to determine their acceptable level of risk and to decide on timelines that suit their needs. Mandiant's patch prioritization is based on exploitation state, each of which is summarized in the following diagram.

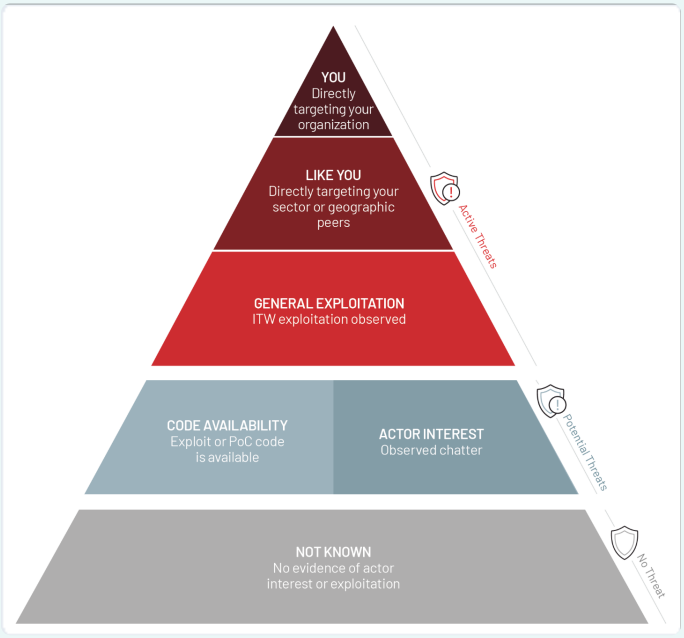

Pyramid chart that shows risk priority levels as suggested by Mandiant

Factors Considered in Ratings

Mandiant vulnerability analysts consider a wide variety of impact intensifying and mitigating factors when rating a vulnerability. While we do not use threat factors such as actor interest, availability of exploit or PoC code, or exploitation in the wild as a primary factor, in some cases it can inform our analysis or be a factor in later updates to a rating, particularly in instances where initial information is sparse.

For more details about how Mandiant applies additional context to exploring vulnerabilities within the Mandiant Advantage platform, see **How to Explore Vulnerabilities (https://docs.mandiant.com/home/how-to-explore-vulnerabilities)**.

Impact considerations: These factors help determine what impact exploitation can have on a targeted system.

- **Vulnerability Type**: The specific type of flaw that causes the vulnerability.
- **Exploitation Consequence**: The result of successful exploitation, such as privilege escalation or remote code execution.
- **Confidentiality Impact**: The extent to which exploitation of the vulnerability can compromise the confidentiality of data on the impacted system.
- **Integrity Impact**: The extent to which exploitation of the vulnerability allows attackers to alter information in impacted systems.
- **Availability Impact**: The extent to which exploitation of the vulnerability disrupts or restricts access to data or systems.

Mitigating factors: These factors help determine what an attacker needs to do to successfully exploit the vulnerability and how likely they are to be successful.

- **Exploitation Vector**: The known methods by which the vulnerability can be exploited.
- **Attacking Ease**: How difficult the exploit is to use in practice.
- **Exploit Reliability**: The degree of analysis performed by Mandiant (for example, unreviewed, reviewed, or tested).
- **Exploit Grade**: The state in which the code exists and its current capability (for example, proof-of-concept, weaponized, non-weaponized, scanner, fake, or unevaluated).
- **Access Vector**: What access is required to successfully exploit the vulnerability, such local, adjacent network, or network.
- **Access Complexity**: How difficult it is to gain access needed for the vulnerability.
- **Authentication Requirements**: Whether exploitation requires authentication, and if so, what type of authentication.
- **User-Interaction Requirements**: Whether exploitation requires user interaction to be successful, and, if so, whether that level of user interaction is likely to be achieved.
- **Vulnerable Product Ubiquity**: How commonly the vulnerable product is used in enterprise environments.
- **Product's Targeting Value**: How attractive a vulnerable software product or device would be for threat actors to target.
- **Vulnerable Configurations**: Whether exploitation requires specific configurations, either default or non-standard.

Mandiant Vulnerability Rating System Applied

The following are examples of cases in which Mandiant rated vulnerabilities differently than NVD by considering additional factors and incorporating information that either was not reported to NVD or is not easily quantified in an algorithm.

Vulnerability	Vulnerability Description	NVD Rating	MATI Rating	Explanation

CVE-2019-12650	A command injection vulnerability in the Web UI component of Cisco IOS XE versions 16.11.1 and earlier that, when exploited, allows a privileged attacker to remotely execute arbitrary commands with root privileges.	High	Low	This vulnerability was rated high by NVD, but MATI rated it as low risk because it requires the highest level of privileges—level 15 admin privileges—to exploit. Because this level of access should be quite limited in enterprise environments, we believe that it is unlikely attackers would be able to leverage this vulnerability as easily as others. There is no known exploitation of this activity.
CVE-2019-5786	A use-after-free vulnerability within the FileReader component in Google Chrome 72.0.3626.119 and prior that, when exploited, allows an attacker to remotely execute arbitrary code.	Medium	High	NVD rated CVE-2019-5786 as medium, while MATI rated it as high risk. The difference in ratings is likely due to NVD describing the consequences of exploitation as denial-of-service (DoS), while we know of exploitation in the wild which results in remote code execution in the context of the renderer, which is a more serious outcome.

As demonstrated previously, factors such as the assessed ease of exploitation and the observance of exploitation in the wild may result in a different priority rating than the one issued by NVD. In the case of CVE-2019-12650, Mandiant ultimately rated this vulnerability lower than NVD due to the required privileges needed to execute the vulnerability and the lack of observed exploitation. On the other hand, Mandiant rated the CVE-2019-5786 as high risk due to the assessed severity, ubiquity of the software, and confirmed exploitation.

In an early 2019 blog entitled **Disclosing vulnerabilities to protect users across platforms (https://security.googleblog.com/2019/03/disclosing-vulnerabilities-to-protect.html)**, Google reported two zero-day vulnerabilities were being used together in the wild: CVE-2019-5786 (a Chrome zero-day vulnerability) and CVE-2019-0808 (a Microsoft privilege escalation vulnerability). Google quickly released a patch for the Chrome vulnerability and pushed it to users through Chrome's auto-update feature on March 1. CVE-2019-5786 is significant because it can impact all major operating systems, including Windows, Mac OS, and Linux, and requires only minimal user interaction, such as navigating or following a link to a website hosting exploit code to achieve remote code execution. The severity is further compounded by a public blog post and PoC code that was released a few weeks later and subsequently incorporated into a **Metasploit (https://www.metasploit.com/)** module (see **CVE-2019-5786: Analysis & Exploitation of the recently patched Chrome vulnerability (https://blog.exodusintel.com/2019/03/20/cve-2019-5786-analysis-and-exploitation/)**).