WIND™

AN INTEL COMPANY

Securing Linux Systems in the Internet of Things

Four Essential Steps for Ongoing Threat Mitigation

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

Open source Linux is a popular choice for developers of embedded systems and devices in the Internet of Things (IoT). But with ever-increasing numbers of interconnected IoT devices being deployed, Linux software vulnerabilities have become more widespread than ever. Taking responsibility for identifying vulnerabilities and making the necessary updates to mitigate threats is often beyond the capacity of device developers and manufacturers. This paper outlines a proven four-step process for resolving Linux vulnerabilities: monitoring, assessment, notification, and remediation. It also explains the cost a company might incur for monitoring and fixing vulnerabilities in-house, and why it may make more sense to partner with an experienced security team to ensure ongoing protection of deployed devices and systems.

TABLE OF CONTENTS

OUR VULNERABLE WORLD

Open source Linux software has gained favor among IoT system developers for a variety of reasons. It gives developers more flexibility by freeing them from being locked into a proprietary vendor's standards. It also offers some practical benefits for IoT applications, notably support for the interoperability that IoT devices often require. Moreover, the cloud systems that run IoT solutions are increasingly built on open source, Linux-based operating systems.

In today's interconnected world, however, securing Linux-based systems and devices has become one of the most pressing and perplexing challenges facing developers and device manufacturers. Gone are the days of "fire and forget" device deployment. Virtually every device made these days is designed for interconnectivity with something, which makes them susceptible to security vulnerabilities. The reality is that connected devices are very likely becoming more vulnerable with every reported exploit.

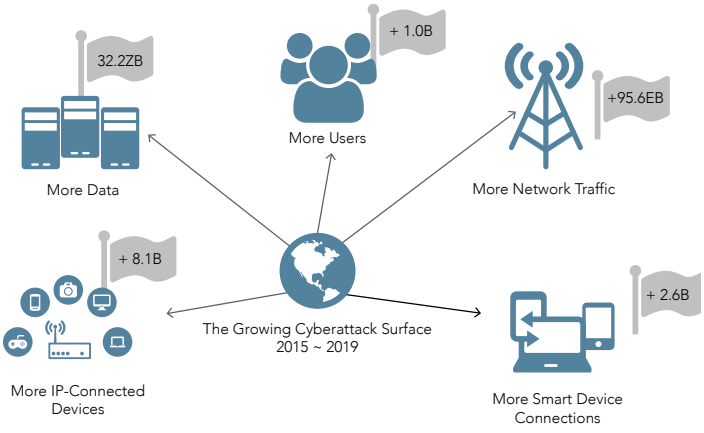

Figure 1. More connected devices means more data and more risk

With the rapid growth of IoT, interconnected devices are proliferating at exponential rates. This massive increase in devices, connections, data volume, network traffic, and users has brought a proportional increase in cyberthreats across a wider attack surface. In response, device manufacturers and developers of IoT applications are employing sophisticated methods to build in powerful security functionality at the earliest stages of design. And that's a good thing. In fact, it's essential. But it's not enough. Threats are constantly evolving. Operators of IoT systems need a mechanism to maintain security in devices over their entire useful life.

Manufacturers need to rethink their security strategies with an eye not only on system-level reinforcing, but also on agile integration of new vulnerability patches. Unless systems are constantly updated, they run the risk of being vulnerable to emerging threats, no matter how strong the built-in security may be.

Consider the example of your own laptop computer. There was a time when all you needed to secure it was a password, and the biggest external threat was an infected floppy disk. Once you connect it to the Internet, however, it becomes a target for attackers, typically via the applications that reside on it. Chances are you receive weekly or monthly update advisories or auto-updates from app providers intended to protect your computer from newly discovered software vulnerabilities.

Every IoT device running Linux needs that same level of ongoing protection. The question is how to accomplish it in a systematic, scalable, and cost-effective manner.

MIND THE GAPS

Before you can fix vulnerabilities in a system, you have to know what and where they are. That's becoming increasingly challenging, as security vulnerabilities are multiplying in parallel to the expansion of the IoT.

Common Vulnerabilities and Exposures (CVE) is the widely accepted, de facto industry standard for identifying, repairing, and reporting vulnerabilities. Using CVE identifiers, the information about a vulnerability can be correlated to appropriate security patches or protection technologies, which is especially vital in the open source software world. Disclosure of vulnerabilities can come from a variety of sources, including the software vendor, security vendors, independent researchers, community mailing lists, and government agencies such as the U.S. Computer Emergency Readiness Team (US-CERT). However, the CVE database has been challenged to keep pace with the volume and scale of vulnerabilities resulting from the IoT world.

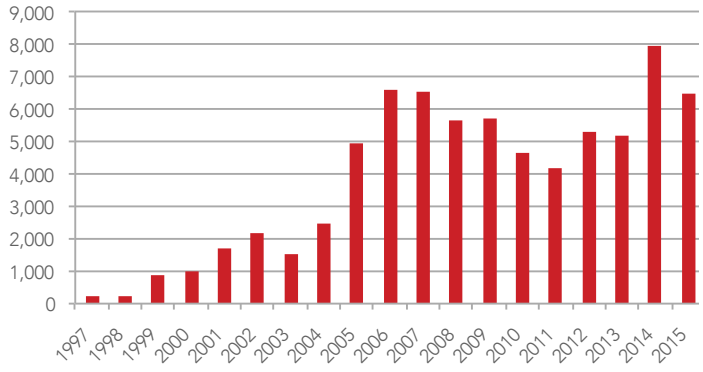

Figure 2. Growing total number of CVEs

The last decade saw an explosive growth in the volume of CVEs over the previous decade, with thousands reported each year. Moreover, they have tended to increase in severity. Based on the Common Vulnerability Severity Scoring system, the number of CVEs deemed high or critical in severity increased by around 25% from 2014 to 2015. According to the National Institute of Standards and Technology (NIST), 80% of all external attacks take advantage of known vulnerabilities in unpatched or misconfigured systems. Meanwhile, in its 2016 Threats Predictions, McAfee Labs reported that many recent "zero-day" attacks (those that exploit vulnerabilities before they become known to the vendor) specifically targeted vulnerabilities in open source software.

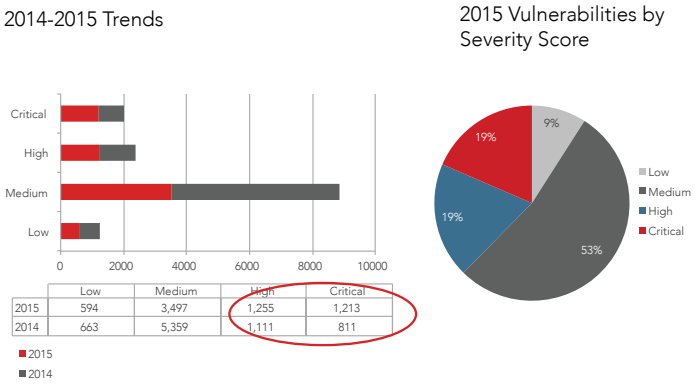

Figure 3. Severity Score trends, 2014–2015

Using open source software actually presents significant advantages from a security perspective. Ongoing threat mitigation requires the ability to update the software on a device as soon as a vulnerability is identified. Because of the large open source community, information about vulnerabilities surfaces quickly through legions of researchers, government agencies such as US-CERT, and dedicated mailing lists. As a result, users of open source in deployed devices can take fast action to lower a potential risk.

It is impractical to think that any system can be rendered 100% impervious to outside threats from persistent attackers given sufficient time and resources. But specific measures can be taken to make things extremely difficult for hackers and reduce the odds of breaches considerably.

THE FOUR ESSENTIAL STEPS

Ongoing threat mitigation in deployed systems requires a four-step approach: monitoring, assessment, notification, and remediation.

Monitoring

Think of monitoring as the "surveillance camera" in your security strategy. Assuming two houses have strong locks, the one with the surveillance camera is going to be better prepared against an intrusion. In this case, the cameras are operated by organizations that issue vulnerability alerts and advisories, such as US-CERT, NIST, the CVE database, various security vendors, private mailing lists, and communities focused on finding Linux vulnerabilities.

The challenge here is that, with dozens of organizations issuing advisories, there is bound to be a certain amount of speculation, making it critical to know which organizations can be relied upon for accurate and actionable information.

Assessment

Once an advisory or security report is received, the system operator or its software partners must make a determination as to whether its devices are vulnerable and to what extent. Vulnerability is typically ranked as high, medium, low, or not present, and prioritized based on likely severity, difficulty of attack, and likelihood of avoidance.

Assessment requires knowing exactly which packages and which versions are vulnerable, and also the exact configurations of your systems. For vulnerable products, the clock for finding a patch starts the minute the vulnerability is exposed.

Notification

Once the vulnerability has been assessed, affected users must be notified of the issue, the determination of vulnerability, and the action plan for remediation. This step requires the right tools and methodologies, so that notifications are sent to all affected parties in a timely and efficient manner.

Remediation

The timing and method of remediation is usually based on priority. Vulnerabilities deemed to be of high severity may require an immediate "hot fix," while others of lower severity may be covered in periodic software updates.

The challenge is having the capability to quickly deliver effective patches and distribute them to end users via a secure channel.

THE PRICE OF PROTECTION

If this four-step process sounds like a lot of work, it is. There's no denying it requires a substantial commitment of people, time, and effort. There are no shortcuts. Speed of response is of the essence. The ideal solution is a dedicated security response team to address every potential vulnerability.

What would it cost to assemble a dedicated security team in-house? Based on 8,000 to 10,000 CVEs uncovered each year, an organization would require a team of four or five highly skilled engineers to investigate and address each one. At an average annual salary of $100,000 for the requisite experience and skill set, the organization would need to budget as much as $500,000 annually for staff alone.

Most device manufacturers and operators of IoT systems would likely consider such specialized expertise outside of their core competency and beyond their budget. The more cost-effective alternative is to assign this responsibility to an experienced commercial Linux vendor with a dedicated security response team—a proven strategy for providing timely protection within hours of vulnerability publication, often weeks or months ahead of the upstream patching.

The right software partner would have the necessary connections within the Linux community and among advisory organizations—combined with its own monitoring and investigative capabilities—to stay on top of vulnerabilities as they are discovered. And because the provider is able to scale its security response services across multiple customers, outsourcing this critical responsibility costs far less than trying to manage it in-house.

THE WIND RIVER LINUX SECURITY RESPONSE PROCESS

As a leading provider of commercial-grade Linux software for embedded applications, Wind River® has devoted the resources necessary to help device manufacturers and their customers maintain ongoing threat mitigation over the life of their systems. The Wind River Linux Security Response Team identifies, monitors, resolves, and responds to Wind River Linux security vulnerabilities. The team follows the four-step process prescribed earlier and ensures adherence to the Wind River Security Response Policy, which establishes target response times based on the priority of the vulnerability.

The Wind River Security Response Team is constantly monitoring the CVE database at cve.mitre.org for potential issues affecting Wind River Linux and Wind River Pulsar™ Linux. This includes specific security notifications from U.S. government agencies and organizations such as NIST, US-CERT, and public and private security mailing lists. Wind River receives email alerts from each of these organizations whenever a new security threat arises. Alerts include both community-confirmed and potential vulnerabilities—the team looks into all of them.

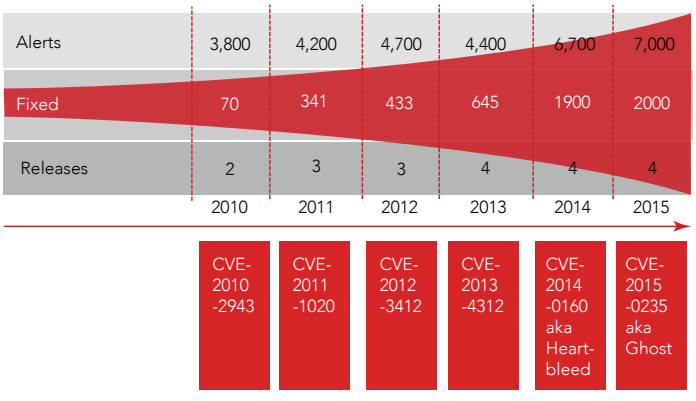

Figure 4. Product releases and integrated patches

Through its membership and participation in the appropriate forums, the security team is often privy to Linux vulnerabilities that have not yet been made public, allowing Wind River and the community to collectively close vulnerabilities and issue patches at a mutually agreed time that coincides with public announcement of the vulnerability. This results not only in a steady stream of security updates, but also in same-day closure of some of the most severe vulnerabilities.

The Security Response Team rolls all patches into future service packs and major releases of Wind River Linux, ensuring all releases contain no known security vulnerabilities.

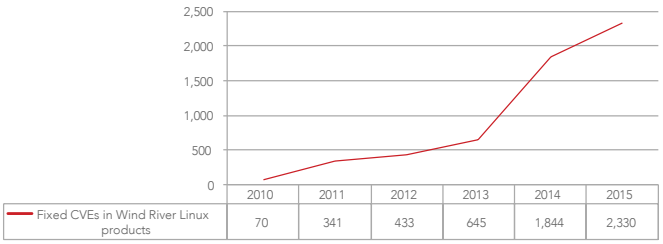

	2010	2011	2012	2013	2014	2015
Fixed CVEs in Wind River Linux products	70	341	433	645	1,844	2,330

Figure 5. Wind River security response timeline

CONCLUSION

Security vulnerabilities are simply a fact of life in today's interconnected world, and they are multiplying with the proliferation of embedded IoT applications. Managing them and mitigating threats is essential for the protection of end users, but requires a level of engagement that is beyond the scope of most IoT solution developers, device manufacturers, and system operators. Fortunately, the open source community is extremely vigilant in finding vulnerabilities that affect Linux software. By working with a software partner that is active in that community, with a proven process for monitoring, assessing, notifying customers, and fixing vulnerabilities, manufacturers and developers can help protect their customers against cyberthreats over the life of deployed IoT systems.

WIND
AN INTEL COMPANY

Wind River is a global leader in delivering software for the Internet of Things. The company's technology is found in more than 2 billion devices, backed by world-class professional services and customer support. Wind River delivers the software and expertise that enable the innovation and deployment of safe, secure, and reliable intelligent systems.