Building on a Root of Trust

As cybersecurity threats evolve, DoD servers are the new ground-zero of attacks

CHALLENGE

Like the burglar who ignores the well-lit front entrance of a marked house and instead enters through an unlocked backdoor, cyber hackers looking to penetrate defense systems are choosing the path of least resistance. They tiptoe around well-fortified perimeter securities to attack a less conspicuous — and often less secure — point of entry: server firmware.

For years, defense agencies have emphasized the value of a defense-in-depth approach to security. But often such efforts overlook the vulnerabilities at the firmware level. Firmware is software embedded in hardware, usually in read-only memory, that directs key hardware functions. Unlike other hardware and software on the network, firmware generally is not covered by traditional security solutions.

"Defense agencies have done a great job protecting the perimeter. That's where all the focus has traditionally been," said Sam Ceccola, Account CTO of U.S. Department of Defense, CTO of Public Sector, at Hewlett Packard Enterprise (HPE). "So why would anyone attack the strong side of an offensive line? They wouldn't. That's why there's an increase in attacks on firmware."

For defense agencies, the surge in hostile acts against servers comes at a time when forward-deployed personnel are more reliant than ever on edge computing and servers at the edge. Unlike mainframes or cloud computing, edge computing relies on a distributed IT architecture and devices with decentralized processing power to avoid bottlenecks, reduce latency, and shrink distances between users and IT resources.

Compliance with NIST guidance requires a level of server security that protects, detects and recovers: protecting hardware, firmware and data; detecting intrusions; and in the event of a breach, recovering data and system operability. Not following NIST guidance is "the biggest risk," Ceccola said.

Consider a Navy vessel deployed to Asia that comes under cyberattack.

"It is great if I can protect the server, and good if I can detect that it's been breached, but what if I can't recover what was lost or damaged in the attack? I can't send a technician out to the South China Sea. That server better be able to recover itself back to a last known good state to preserve that mission," Ceccola said. "There's a lot of the industry that does a great job protecting part of the firmware, but that's not good enough, because that still leaves something exposed, and that's what gets attacked."

Here's the challenge. A **cryptographic hash** function is an algorithm that generates enciphered text — a **hash value** — that can be used as a password. On less-secure servers, cryptographic hash keys reside in firmware. When booting up, those servers run millions of lines of code before boot security kicks in, creating a vulnerability gap. To a would-be hacker, that gap looks like an invitation — much the way a backdoor that's ajar beckons to a burglar.

At the root of the challenge is the base management controller (BMC) firmware. Vendors who do not manufacture proprietary BMC silicon and firmware — and who don't have supply chain visibility and traceability — create server vulnerability that has far-reaching ramifications, Ceccola said. "If I can infect BMC, I become the basis on which everything else in that server — applications and data — is built on, and I can affect everything," he said.

SOLUTION

The world's most secure industry-standard servers use advanced "silicon root of trust" technology that establishes security before a single line of firmware code runs on a server. This solution relies on cryptographic hash keys that open servers' algorithmic locks before authorizing access. "There is a huge difference between a root of trust being in firmware and only protecting half of the server, and a root of trust being in silicon and protecting the whole server," Ceccola said. "It's baked into the silicon as an immutable fingerprint. That's why it's so effective."

HPE's exclusive silicon root of trust solution is based on a hardware-validated boot process that ensures a computer system can only be started using code from an unalterable source. Combined with a cryptographically secured signature, the zero trust-based solution leaves no accessible gaps for hackers to exploit. The technology is available with all HPE Gen10 servers including ProLiant, Apollo, Synergy, EdgeLine 8000 and hyperconverged systems.

Root of trust aligns with DoD interest in zero trust security models. Zero trust involves authenticating the identity and permission levels of a user and the security status of a device every time a user or system accesses a network resource. In this environment, bolting security solutions on top of the BMC isn't good enough as over 4 million lines of code run before the server boots up.

"If you are going to have technology that is making authoritative decisions for you, it better be at the lowest possible point" of the architecture, Ceccola said. "You cannot establish a zero trust architecture if the foundation is not trusted."

By definition, the **root of trust** is a set of functions that is always trusted by the server's operating systems. Operating as a separate computer engine, the root of trust controls the trusted computing platform cryptographic processor.

Most servers are configured to engage the root of trust security only at boot time, further creating exposure to potential intruders, such as attacks on servers where malicious code has been introduced. The most secure servers allow agencies to re-authenticate all firmware at runtime, for example through an out-of-band management capability. "It executes the authentication and doesn't degrade the performance of the server," Ceccola said. "It's protecting the whole boot cycle and protecting it in runtime, not just at boot time."

For example, HPE's silicon root of trust solution can conduct daily runtime firmware validation. If compromised code or malware is inserted in any of the critical firmware, an HPE audit log alert is created to notify the customer that a compromise has occurred.

Root of trust security is also resilient. In the event of a security breach, the most secure servers are able to recover the last known good state or "golden image" of all server firmware. Captured on encrypted manned flash, operators access recovered versions through the BMC silicon. Agencies have the option of configuring servers with the capability to automatically recover known good versions of firmware anytime anomalous firmware is detected.

Hewlett Packard Labs is also working on innovative research programs focused on new methods to identify Operating System (O/S) Kernel Intrusion Detection as a future next line of defense beyond the silicon root of trust. This is a critical area of focus that is needed to prevent difficult to detect, persistent and sophisticated cybersecurity attacks on the O/S. Today, a compromised O/S cannot identify that it has been compromised and this represents a growing and very sinister cyber threat to both governments and enterprises today.

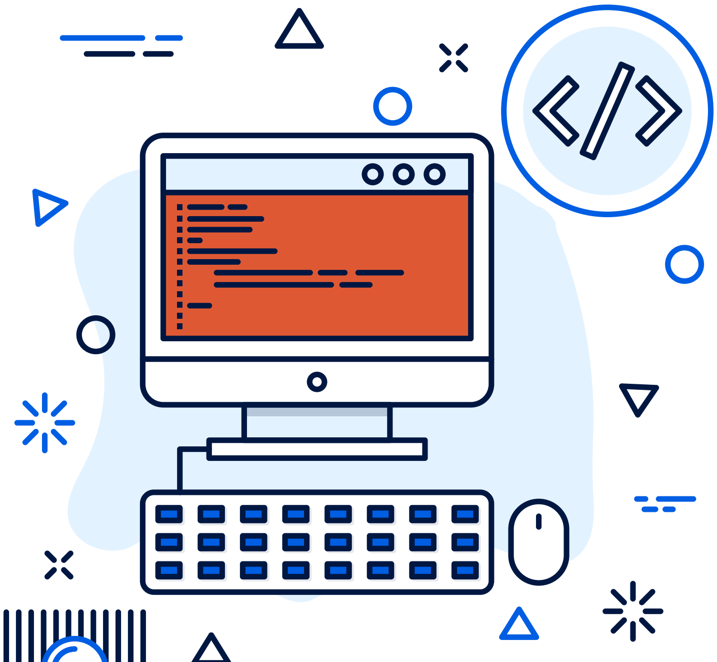

STATS

70%

Percent of organizations without a firmware upgrade plan that will be breached by 2022.

$18.8 billion

Projected spending by the U.S. government for cyber security in 2021.

2.86 million

Users of DoD IT networks (active duty service members, civilian personnel and National Guard/Reserve service members).

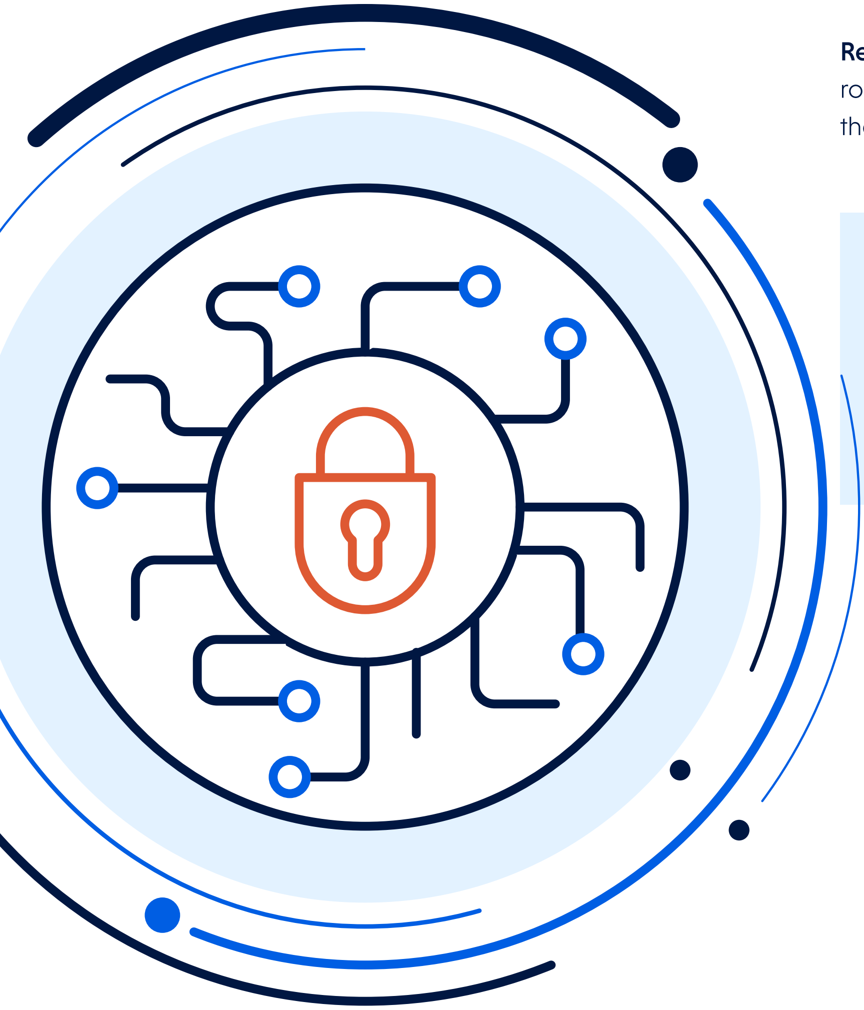

TIPS FOR SUCCESS

The NIST Cybersecurity Framework, a roadmap for securing IT enterprises, rests on a foundation of five functions that are enhanced by silicon root of trust technology:

Identify - Root of trust security helps organizations develop expand awareness needed to manage cybersecurity risk — to systems, people, assets, data and capabilities.

Protect - Secure servers should protect firmware at the lowest level possible, establishing a root level of trust that secures firmware before code runs during the booting process.

Detect - Root of trust provides maximum protection when it runs in boot and runtime. Less-secure servers apply root of trust security only to boot up, making it more difficult to detect anomalies.

Respond – Improved threat-detection capabilities of silicon root of trust technology lead to faster, more-effective responses to security events.

Recover - The most secure servers, equipped with silicon root of trust technology, can always restore firmware to the last known good version.

Learn more about how Affigent and HPE can ensure a reliable and secure IT environment at affigent.com/products/hpe.

Affigent
AN AKIMA COMPANY

Hewlett Packard Enterprise

govloop