



**TELEDYNE
LECROY**



APPLYING ZERO TRUST IN THE SUPPLY CHAIN TO PREVENT DMA ATTACKS

Sophisticated state-based adversaries are increasingly targeting the technology supply chain. These adversaries can surreptitiously and deeply embed threats within trusted components long before a final product is delivered to the end customer. The firmware within modern SSD drives represents a key potential target, allowing malicious actors to launch devastating DMA attacks that gain full control over a device and subvert the traditional security controls running in the operating system. This makes it critical for SSD drive suppliers and drive qualification teams to employ appropriate practices and security controls to ensure the integrity of their components throughout all phases of the supply chain.

In this paper, we will look at how organizations can apply the core concepts of Zero Trust to mitigate their risk and prevent firmware-based attacks such as DMA attacks. Specifically, we will:

- Introduce DMA and firmware-level attacks and their security impact on a device.
- Provide an overview of Zero Trust principles and how they relate to DMA attacks and supply chain security.
- Propose a set of specific best practices and tests that can be applied to ensure the integrity of SSDs through all phases of the supply chain.

While this is not intended to be a comprehensive analysis of Zero Trust or DMA attacks, the goal is to provide vendors, validation teams, and IT and Security teams with a practical approach to improve their security going forward.

FIRMWARE-BASED THREATS AND DMA ATTACKS

The firmware within systems and their components is some of the first code to execute when a device boots up and is some of the most privileged code on a machine. Any compromise of this code can allow an attacker to gain control of how the system boots, subvert other security controls on the system, and ultimately, maintain persistence even if the operating system is completely re-installed. This is a critically important concept as it allows an attacker to compromise a device below the level of the operating system (and before the OS even runs), where most traditional security controls are run. When that compromise happens on some of your most core systems where storage resides it can become especially problematic.

This risk has translated to real attacks in the wild including some of the most damaging malware and ransomware attacks. The Trickbot malware recently added a new module dubbed "TrickBoot" to check devices for well-known vulnerabilities that can allow attackers to read, write, or erase the UEFI/BIOS firmware of a device. This is a significant

development given Trickbot's role in maintaining persistence for a variety of malware campaigns, including the Ryuk family of ransomware. Likewise a wide range of nation-state threat actors have targeted firmware within network devices and VPNs to gain access to target networks.

DMA Attacks

DMA or Direct Memory Access attacks are another critical example of firmware threats. Direct Memory Access is a normal and necessary part of system operation that allows components such as SSD drives to quickly read and write to system memory directly without the need to be processed by the main CPU and OS. This enables the system to move data at high speeds that would otherwise be impossible.

However, attackers can abuse this same functionality to steal sensitive data from memory or to overwrite that memory and gain control over kernel execution of the device. This is an extremely powerful attack as it can provide the adversary with complete control over the device and the freedom to perform virtually any malicious activity.

This is a particular concern in the context of today's SSD drives. In recent years, many SSD drives have evolved from using traditional SATA interfaces to much faster PCIe and NVMe technologies. However, these drives have more freedom to read and write to arbitrary memory addresses due to their use of the PCIe interface. This is even true of SATA Host Bus Adapters (HBAs). This access to memory means that any vulnerable or compromised firmware in a PCIe/NVMe or SATA HBA could enable an attacker to remotely execute code in the pre-boot environment. Such code may alter the initial state of an operating system, violating common assumptions on the hardware/firmware layers and breaking OS-level security controls. This means that as developers chase faster data transfers from their SSDs, they are also carrying a greater risk of DMA attacks.

ZERO TRUST, DMA, AND THE SUPPLY CHAIN

"Zero trust" (ZT) is the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. "Zero trust" assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned).

The National Institute of Standards and Technology (NIST) special publication SP 800-207 provides an in-depth

description of Zero Trust concepts and how they can be applied to security operations. One of the core concepts is as follows:

No resource is inherently trusted. Every asset must have its security posture evaluated via a PEP (Policy Enforcement Point) before a request is granted to an enterprise-owned resource.

The 2021 Executive Order on Improving the Nation's Cybersecurity likewise calls out the importance of Zero Trust. The executive order is a watershed document that introduces new perspectives and directions for the prevention of cyber-attacks. Two sections of the Executive Order stand out as clear mandates for Federal Agency cybersecurity teams, but also as innovative, new approaches for civilian teams who need to improve their strategies to counter new adversary tactics:

1. Section 3, which calls for "Modernizing Federal Government Cybersecurity," focusing especially on the design and implementation of Zero Trust architectures in government networks, and;
2. Section 4, which concentrates on strengthening and securing the software supply chain. While all ten sections of the Executive Order serve as clear instructions for federal agencies and forward-thinking guidance for CISOs in the commercial sector, these two sections mark significant departures from previous best practices.

ZERO TRUST IN THE CONTEXT OF DMA ATTACKS

Both of these mandates have important implications in the context of DMA attacks. By design, DMA provides an implicit level of trust to a component such as an SSD drive in exchange for better performance. As stated previously, this is particularly true in the case of PCIe/NVMe SSD drives. DMA trusts a device to directly access one of the most sensitive resources on a system.

This is a fundamental break from Zero Trust principles.

Some may try to justify such a trade-off by assuming that other security controls would prevent such a compromise from happening in the first place. However, firmware-level and boot protections can vary wildly from device to device and can be vulnerable even in the best of circumstances. And protections against insider attacks are either non-existent or just beginning to be contemplated.

If the firmware in a device or component is compromised in the supply chain then it often becomes inherently trusted as a valid part of the system. For example, many firmware checks will simply take measurements of firmware to verify that the firmware hasn't unexpectedly changed. If a component is already compromised in the supply chain before such UEFI measurements take place, the system will inherently trust the compromised component firmware.

This forces teams to think about Zero Trust both within a device as well as in the context of the larger supply chain. Internal components and firmware can't be inherently trusted. Likewise the components and systems delivered from partners and suppliers can't be inherently trusted.

For SSD device manufacturers and SSD consumers, this means your firmware can no longer be assumed to be trust-worthy. Drive validation and qualification processes must be employed to verify an SSD's trustworthiness. As CISO's look at drive firmware with increased skepticism, vendors must be able to show a set of practices to verify the integrity of their firmware development and supply chains to ensure their firmware has not been compromised.

Gartner puts an emphasis on this argument in their 2020 report "Roadmap for Improving Endpoint Security" when they say, "Firmware may well be the next endpoint battleground for advanced adversaries as script controls tighten."

This triggers a chain of logic that calls for attention from CISOs, security architects and practitioners alike:

- For cybersecurity programs to be successful they must rely on Zero Trust strategies, tactics, and postures


- A successful Zero Trust program must have an active, expanded understanding of Device Integrity
- Device Integrity in turn requires deep, firmware and hardware level discovery, evaluation, and remediation capabilities

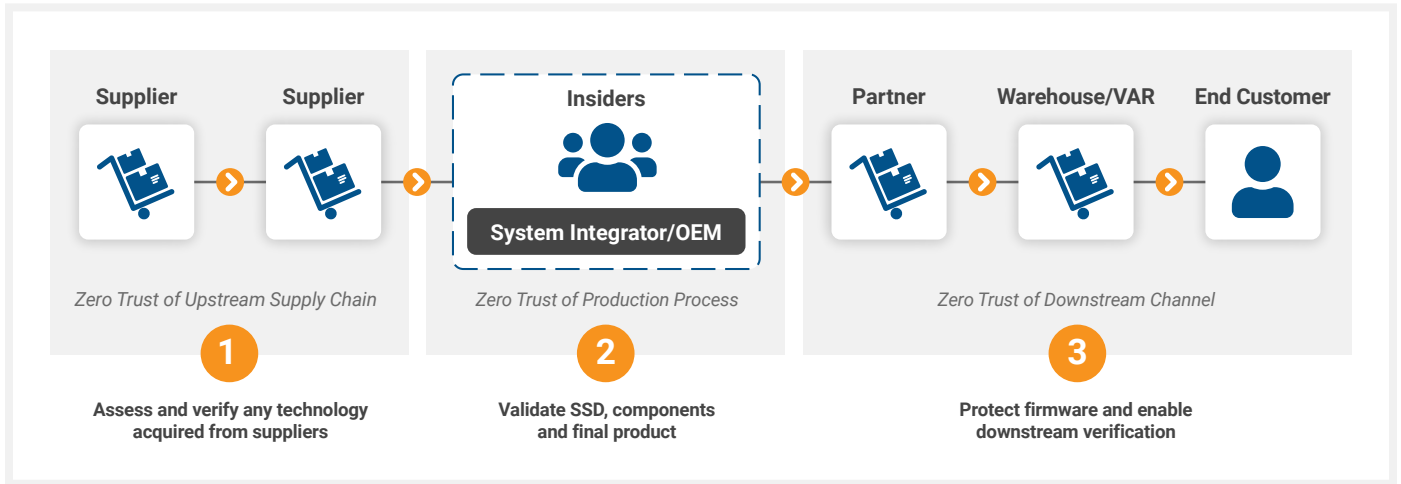
Additionally, malicious firmware such as the Equation Group's HDD implant allowed attackers to hide malicious code on a drive that remained invisible to the host operating system.

BEST PRACTICES TO ESTABLISH ZERO TRUST AND PREVENT DMA ATTACKS

Technology supply chains and development processes are naturally complex, and there is the potential for risk at virtually every step. System integrators and OEMs will need to apply Zero Trust principles in several ways throughout the process. As shown in the diagram below, we have broken these into three high-level phases from the perspective of an SI or OEM organization – the upstream supply chain, the integration or manufacturing of the complete product, and delivery through the downstream channel. Next, we have provided example best practices that can be applied to ensure Zero Trust during each phase.

Of note, an SI/OEM will need to work closely with their upstream and downstream partners to define what is expected from each organization in terms of security. However, Zero Trust means that an SI/OEM cannot trust that partners will meet their obligations. As a result, each section below is heavily focused on what an SI/OEM can do to make the process and product as safe and auditable as possible for the ultimate end customer.






1 APPLYING ZERO TRUST TO THE UPSTREAM SUPPLY CHAIN

As a part of Zero Trust, the OEM/SI organization should never implicitly trust any technology delivered by their vendors. All incoming firmware should be scanned for known vulnerabilities, misconfigurations, and known threats. The assessment process can begin even before a supplier is selected by scanning any supplied firmware and components for known vulnerabilities and threats. This same assessment should also be performed for goods that are received to validate the integrity of the delivered components. Firmware security platforms such as Eclipsium are examples of tools that can be used to assess vendors. Key steps can include:

- Define the security requirements and expectations of all suppliers** - SI/OEMs should work with their suppliers to ensure that each supplier recognizes that they are responsible for the security of all code they deliver, including the code of any sub-system suppliers that they may contract with. An SI/OEM will naturally verify all code in subsequent steps. However, this step will help to ensure that suppliers work with reputable sub-contractors and that security checks are introduced into the supply chain as early as possible.
- Scan components and firmware for known vulnerabilities and misconfigurations** - These are weaknesses in the code that would allow an attacker to insert a threat (see below). Many vendors will reuse the same libraries, which can allow the same well-known vulnerabilities to show up in a variety of products. Any vulnerabilities in SSD firmware could allow an adversary to implant malicious code within firmware that would then drive a DMA attack. An SI/OEM supplier should scan all firmware for known weaknesses and misconfigurations. Likewise, firmware protections should be properly implemented to help prevent firmware from easily being tampered with. No firmware updates should be allowed without the firmware being properly signed. Teams can refer to established open compute project (OCP) tests (PDF) for recommended configurations as well as firmware security platforms to automate these types of assessments.
- Verify integrity of received firmware** - SI/OEMs should work with their suppliers to identify the latest approved firmware from the supplier and develop a firmware bill of materials (SBOM). The supplier should be verifying this BOM in each device before delivery of the component components to ensure that they match the valid version and expected SBOM. This can be done by cryptographically comparing the hashes of the observed firmware against the expected version for each piece of firmware and can look to the new OCP 1.0 standard for process checks to ensure the firmware if the intended firmware.
- Scan firmware for known threats and implants** - Firmware threats such as implants and back doors are malicious code embedded into the firmware. Organizations should also assess firmware for known implants. Even a well-meaning supplier may be unknowingly compromised by an attacker. This is also an important step as attackers have repeatedly reused code from earlier malicious implants. For example, the MosaicRegressor implant discovered in 2020, reused much of the same code from the Hacking Team implant released five years before.

- **Implement robust validation processes** - These processes should ensure the delivered firmware has not been compromised by an insider. This is a critical step for identifying a SolarWinds style of attack in which attackers compromise the valid development process. Expected actions should include independent review of the code, and lifecycle qualification tests that look for aberrant behavior that might occur within a drive's normal life. These qualification steps should be implemented by a separate team removed and unknown to the development team creating the firmware. The testing should monitor the

behavior of firmware to identify any anomalous actions that could be indicative of an unknown threat. To accomplish this monitoring, suppliers should implement a firmware security platform or service that scans for any anomalous actions that could be indicative of an unknown threat. The testing or service should include a means of validating PRP and SGL memory access locations such as can be found with Teledyne LeCroy Oakgate's memory fencing. Being able to see and stop any unauthorized memory addresses will prevent users from placing unauthorized code in hidden places in memory.



2 APPLYING ZERO TRUST TO OEM/SI PRODUCT DEVELOPMENT AND PRODUCTION

As a part of Zero Trust, an acquiring organization should never implicitly trust any technology delivered by or validation steps performed by their vendors. All firmware should be scanned for known vulnerabilities, misconfigurations, and known threats. While the previous assessment phase helps to identify any obvious weaknesses or problems from a supplier and their supply chain, the SI validation process provides another in-depth and active assessment of the security both of individual components and the final product.

For example, teams will want to incorporate tests that verify how a drive actually behaves by modeling their environment. Some of the following steps can require a level of firmware expertise not present in the SI/OEM, and organizations may want to consider a firmware security platform or a firmware security service to automate the process.

- **Perform supplier Audits** - SI/OEM's should implement as part of their component supplier review process a checklist to confirm they meet the acceptable security practices outlined above for critical components like SSD's. The burden for ensuring safe firmware should be spread back into the value chain to improve the level of security robustness.
- **Repeat scan to verify the integrity of firmware and to identify known vulnerabilities and threats** - Qualification teams should apply the same scans for known vulnerabilities, misconfigurations, and threats as done by suppliers. This testing is to ensure that vulnerabilities or threats were not introduced during the manufacturing or assembly of the final product. Teams can assess individual components as well as the final assembled system. No firmware updates should be allowed without the firmware being properly signed.
- **Observe firmware-level behavior for anomalies** - Most firmware components should exhibit very predictable behaviors. A firmware security platform or service will be able to identify anomalous behaviors that could indicate the presence of an unknown or custom threat. One such anomaly is the SSD is supposed to honor PRP and SGL memory descriptors. However, this trusted state allows internal bad actors on a platform to launch DMA attacks. SI/OEM's should implement their own proprietary set of test leveraging tools such as OakGate's Memory Fencing technology and Eclipsium's firmware security platform or service. This second level of tests provides additional protection against hard to see insider attacks. SI/OEM's should not assume that just because the supplier builds exclusively in country X that an insider cannot be compromised.
- **Incorporate Firmware BOM into a System Wide BOM** - Firmware from your SSD will be a subset of firmware on the system. All firmware should be included in a proprietary firmware SBOM which is checked on the system before it is packaged and shipped to the final customer or channel.
- **Perform 3rd party code reviews** - If possible, organizations may want to have internal teams or firmware specialists perform analysis of the supplied firmware source code to proactively identify any unknown potential weaknesses or vulnerabilities.

3 APPLYING ZERO TRUST TO THE DOWNSTREAM CHANNEL

SI/OEM teams will need to take steps to ensure that products remain safe and can be easily audited by downstream partners as well as the ultimate customer. It will naturally be harder for an SI/OEM to have direct control over the security of a product once it leaves their control. As a result, this phase focuses on ensuring that any firmware is protected from unauthorized modifications, working with partners to establish additional verifications that may be necessary, and making it easy for partners and end customers to verify the integrity of the products that they receive.

- **Verify firmware-level protections and security configurations** - SI/OEMs will also need to verify that the SSD and other component firmware is properly configured to ensure that it can't easily be altered by 3rd parties. For example, teams should ensure that all firmware updates are required to be cryptographically signed prior to being committed. Without this basic protection, the firmware could easily be altered at any point in the supply chain. Once again OCP tests or a dedicated firmware security platform can be used to identify these issues.
- **Establish SBOMs for critical firmware** - Section 4 of the previously referenced Executive Order heavily focuses on the requirement for "critical software" to have a Software Bill of Materials (SBOM) to assure it maintains its intended integrity. The order specifically encourages organizations to press their vendors for complete SBOMs to verify their equipment throughout the acquisition and deployment processes. By confidentially providing SBOMs for their firmware, SI/OEMs can make it much easier for downstream

partners and consumers to verify the integrity of their products. Any channel partner opening the box and adding components should be required by the SI/OEM to verify the SBOM as part of their assembly process. If any added components have firmware, the SBOM should be updated before sending on to the final-end customer. This is particularly important because it provides a way for the end consumer to verify that the product was not altered between the SI/OEM, the channel partner and the final consumer.

- **Set security requirements for downstream channel partners** - Firmware integrity and vulnerability checks, should be applied anytime that a product is opened or modified. If partners need to unbox a product to make any changes, the SI/OEM should require that channel partners have the correct controls in place to ensure the product has not been altered. These requirements are particularly true for product destined for the defense industry or other critical infrastructure targets.

SUMMARY/CONCLUSIONS

Firmware threats such as SSD-enabled DMA attacks can be incredibly damaging yet hard to detect. By introducing such threats in the supply chain, adversaries and threat actors can deliver their threats to a wide range of customers, without the need to infiltrate each individual target. The complexity of modern supply chains requires system integrators and OEMs to think about this risk in depth and to apply Zero Trust principles throughout their process. At a low-level, these organizations must be sure that individual components within a system are never trusted. Special attention should be paid to DMA-capable components due to the inherent trust associated with direct memory access. Likewise, organizations will need to approach Zero Trust at the level of the relationships with their partners, suppliers, and customers. System integrators and OEMs must implement processes to continually verify the posture and integrity of the components they receive from suppliers, and likewise to enable downstream partners and customers to verify the final products that they receive.

While this will require some planning and effort, organizations have the tools and services needed to automate this process and ensure the highest levels of security for their products.

REFERENCES

Ross Stenfort (Facebook), Ta-Yu Wu (Facebook), Lee Prewitt (Microsoft), Paul Kaler (HPE), David Derosa (HPE), William Lynn (Dell EMC), Austin Bolen (Dell EMC), (2021) *Open Compute Project: Data Center NVMe SSD Specification*

<https://www.opencompute.org/documents/datacenter-nvme-ssd-specification-v2-0r21-pdf>

Scott Rose (NIST), Oliver Borchert (NIST), Stu Mitchell (Stu2Labs), Sean Connelly (DHS), (2020) Zero Trust Architecture, SP 800-207

<https://csrc.nist.gov/publications/detail/sp/800-207/final>

Peter Firstbrook (Gartner) (2018) Roadmap for Improving Endpoint Security, G00343353

<https://www.gartner.com/en/documents/3879573/roadmap-for-improving-endpoint-security>

The White House (2021) Executive Order on Improving the Nation's Cyber Security

<https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>