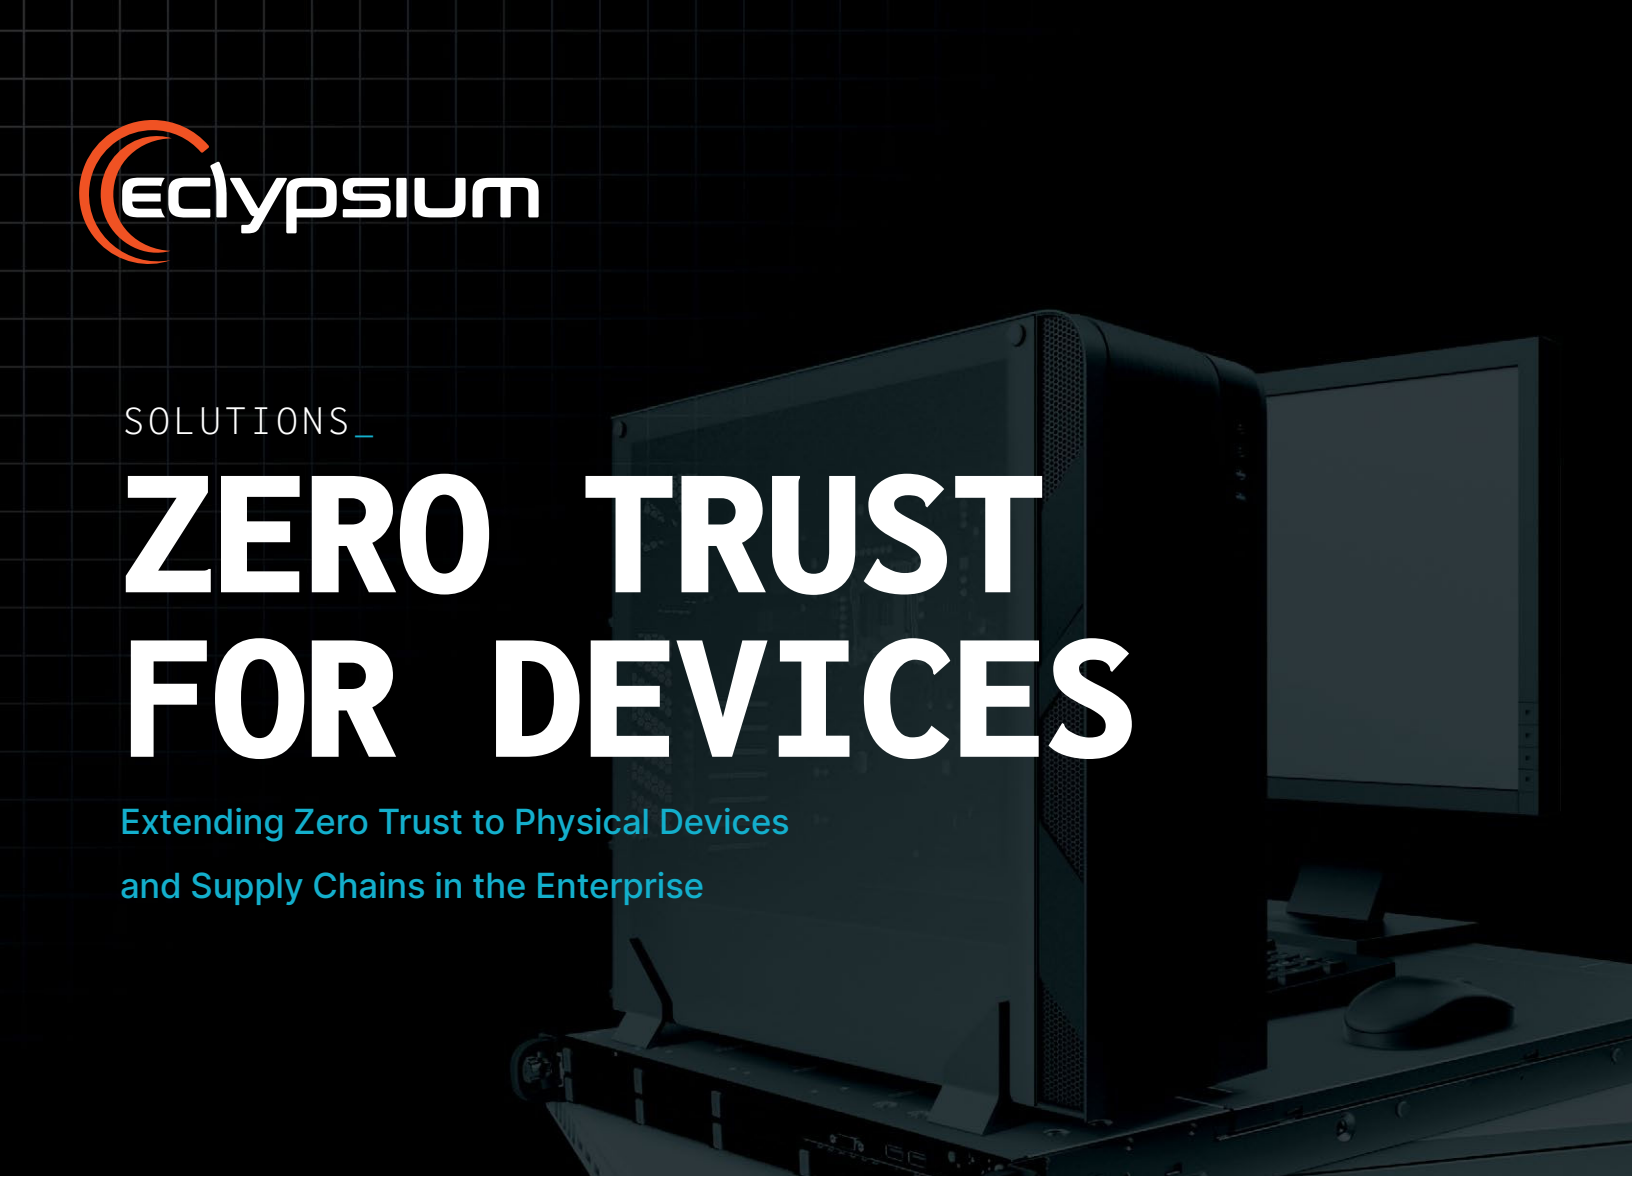

SOLUTIONS_

ZERO TRUST FOR DEVICES

Extending Zero Trust to Physical Devices and Supply Chains in the Enterprise

WHO SHOULD READ THIS_

Organizational cybersecurity leaders, including CISOs and CIOs; endpoint security managers, security and network architects; teams responsible for data center security, infrastructure security, or network security who are planning and executing Zero Trust projects.

WHAT THEY WILL LEARN_

How and why Zero Trust principles apply to devices and their supply chains, and best practices for extending these principles down to physical hardware including chips, processors, and system components.

FURTHER READING_

- TAG Cyber white paper, “Making the Case for Firmware in the Context of Zero Trust Security”
- Eclypsium white paper on Executive Order 14028
- Teledyne Lecroy & Eclypsium joint white paper, “Applying Zero Trust in the Supply Chain to Prevent DMA Attacks”

Zero Trust is a foundational principle of enterprise cybersecurity strategies across virtually every industry. Likewise, Zero Trust designs, architectures, and practices are increasingly codified in government mandates and regulations. In the U.S. examples include Presidential Executive Order (EO) 14028, *Memorandum For the Heads of Executive Departments and Agencies* M-22-09, and NIST's SP 1800-35 2022 draft, *Implementing a Zero Trust Architecture*, while international standards and frameworks like the Singapore Cybersecurity Strategy of 2021, the Monetary Authority of Singapore, and ISO-27001 have all adopted Zero Trust principles into their frameworks and regulatory language.

But how should organizations apply these principles to the devices and assets that are delivered by information and communication technology (ICT) supply chains? How will they stand up to the multitude of embedded components that exist within these devices? This solution brief answers those questions and gives strategists and practitioners guidance on achieving a Zero Trust posture at the asset and hardware layer. .

UNDERSTANDING ZERO TRUST_

The O'Reilly textbook, *Zero Trust Networks: Building Secure Systems in Untrusted Networks* by Evan Gilman and Doug Barth, sets out five basic principles at work in a Zero Trust network design:

1. The network is always assumed to be hostile.
2. External and internal threats exist on the network at all times.
3. Network locality is not sufficient for deciding trust in a network.
4. **Every device, user, and network flow is authenticated and authorized.**
5. **Policies must be dynamic and calculated from as many sources of data as possible.**

This brief focuses on Principles 4 and 5 from the Wiley book, where the Zero Trust "rubber" hits the real-world "road." It's here we start thinking about the natural extension of Zero Trust principles down into the devices and hardware they use, while also calculating the unique, highly dynamic risks not only of our devices, but of the countless hardware, firmware, and software components that make them whole.

ZERO TRUST FOR EVERY DEVICE_

Enterprise networks consist of two major entities: people and devices. "People" are generally covered by identity and access management (IAM) programs. Devices, likewise, need to be uniquely authenticated, typically either through user interaction, embedded X.509 certificates, SSH keys, or other methods. Principle 4 as defined by O'Reilly states that "every device, user, and network flow is authenticated and authorized."

> **"Every device, user, and network flow is authenticated and authorized"**

However, Zero Trust requires more than simply checking the identity of devices. Organizations also need to know that their assets are free from threats and vulnerabilities down to the underlying code, hardware, and components. For example, a device could pass identity checks by having the appropriate certificate yet still be compromised by a backdoor or UEFI/BIOS implant.

In order to verify the integrity of physical devices, security teams will need new types of visibility that understand the inner workings of devices and the supply chains that produce them. We need to recognize that each individual asset is actually an amalgamation of components and code, often from dozens of disparate supply chain suppliers including motherboards, central processing units (CPUs), memory, PCI cards, solid state drives, system management modules (SMM), baseboard management controllers (BMCs) in servers. Each component is a potential source of risk, and each supplier is a potential point of compromise in the supply chain.

This potential risk has rapidly turned into a documented reality as threat actors have aggressively pivoted to exploiting technology supply chains. While enterprise networks are often well-defended, supply chain vendors often are not. For every product, attackers can target dozens of supply chain suppliers to compromise devices

before they are ever delivered, or can similarly steal keys and source code in order to deliver low-level implants within product updates. Within the past several months, threat actors have compromised major supply chain vendors affecting chipsets, motherboards, laptops, and server baseboard management controllers. **Thus, in the same way that organizations must assume their networks are hostile, they must also assume their devices and supply chains are hostile**.

To address this risk, enterprise security teams must actively validate all their critical devices and the components within those devices. How do we verify device components? Some use signed certificates and keys, but for the majority, we can verify them through the firmware and microcode embedded in them by their manufacturer.

According to research from analyst firm Gartner:

- Every endpoint is delivered, on average, with 15-20 firmware components
- Every server is delivered with around 30 components, and sometimes more than 50
- Every network device is now shipped with dedicated firmware

This firmware and low-level code represents potential risk, but it also provides a way to uniquely establish the "integrity" of the devices at the most fundamental levels. Unlike user applications that are constantly changing, firmware and other critical forms of code remain very predictable. This makes it possible to take cryptographically assess the code and recognize if it has been altered or modified in any way.

ESTABLISHING ZERO TRUST THROUGH FIRMWARE_

Firmware is simply software code that's been embedded directly onto hardware components by that hardware's manufacturer, or one of their suppliers. It doesn't reside in common storage locations for files and data, but in specialized chips.

Firmware is increasingly used as an initial attack vector. This table lists common hardware or firmware-based vulnerabilities and the recent exploits that leverage them.

info@eclypsium.com
eclypsium.com

Component	Role	Vuln?	Exploited?
Central Processing Unit (CPU)	Often called microcode, CPU-level firmware is powerful and privileged. Microcode firmware typically resides in special high-speed memory and translates machine instructions, state machine data, or other input into sequences of detailed circuit-level operations.	Yes	Dirty Cow Spectre Meltdown
Unified Extensible Firmware Interface (UEFI)	A specification that defines a software interface between an operating system and platform firmware. UEFI replaces the legacy Basic Input/Output System (BIOS) boot firmware.	Yes	BlackLotus Moon Bounce Cosmic Strand TrickBoot
Trusted Platform Module (TPM)	An international standard for a dedicated microcontroller designed to secure hardware through integrated cryptographic keys. The term can also refer to a chip conforming to the standard.	Yes	Various probing, side- channel, interposer attacks
Management Engines (ME)	Intel's autonomous subsystem that has been incorporated in virtually all of Intel's processor chipsets since 2008.	Unk	Conti is focused here
Baseboard Management Controller (BMC)	Provides the intelligence in an Intelligent Platform Management Interface (IPMI). It is a specialized microcontroller embedded on the motherboard of a server computer. The BMC manages the interface between system-management software and platform hardware through dedicated firmware and RAM.	Yes	iLOBleed USBAnywhere
Network Card (NIC)	A network interface controller (NIC, also known as a network interface card, network adapter, LAN adapter or physical network interface, and by similar terms is a computer hardware component that connects a computer to a computer network.		EtherLED NetSpectre
Direct Memory Access (DMA)	A feature of computers that allows certain hardware subsystems to access main system memory independently of the central processing unit (CPU) and run commands through onboard RAM.	Yes	Various DMA and side channel attacks
Embedded Controller	A microcontroller in computers that handles various system tasks. Usually merged with Super I/O, especially on mobile platforms.	Unk	Multitude • Firmware-based • Network-based • Side-based
System Management Module (SMM)	Chips that enable System Management Mode, which when active provides an alternate firmware-based software system with higher privileges.	Yes	HPE devices AMD chips
Solid State Drive (SSD)	A solid-state storage device that uses integrated circuits to store data persistently, typically using flash memory, and functioning as secondary storage.	Soon	SSD attacks Micron Flex

In addition to these common components, most computers have additional chips for video processing, sound processing, digital signal processing, and other application-specific integrated circuits (ASICs). Each of these components represents a source for compromise or vulnerability. Each must be included in the universe of data we use in the decision-making processes of Zero Trust systems. By assessing the attributes of this low-level code – its version, source date, binary signature, and provenance – practitioners can build trust in these underlying (and invisible) components.

How do we make those critical Zero Trust decisions? We assess asset risk down to the chip- level.

ASSESSING CHIP-LEVEL RISK_

The 5th and final point from the Zero Trust Networks book cited earlier is both specific and almost impossibly broad: "Policies must be dynamic and calculated from as many sources of data as possible."

> **"Policies must be dynamic and calculated from as many sources of data as possible"**

"Policy" refers to the output from a trust engine. A trust engine, in turn, calculates risk based on system inputs. Given the data on hand, do we trust this component or not? Can we allow this device or this user on the network, or not?

We can generate a Zero Trust test case using an example from the "BMC" row in the table on the previous page:

- A subject system on the network is an HPE Gen9 server using iLO4, HP's "integrated lights out BMC" module.
- A process (or person) wants to store critical or sensitive data on this server.
- But as this post explains, this BMC module has been actively exploited in ransomware attacks known as iLOBLeed, and it is difficult to tell whether an implant is at work on this HPE server without doing a firmware-level scan.
- Do we trust it? Or do we explicitly distrust it?

To answer this question, we need to reliably verify the code within the BMC itself. However, this can be easier said than done. Not only does it require firmware and BMC expertise, the iLOBleed threat actively tries to prevent the device from being updated and then reports false information in order to appear that the update was successful. So in reality the device is vulnerable and compromised, yet traditional security and vulnerability scans see it as safe. To close this gap, teams need independent, purpose-built visibility into low-level code, firmware, and components within their assets.

HOW ECLYPSIUM HELPS ACROSS THE ENTERPRISE_

At its core, Zero Trust is about rooting out areas where trust is assumed and replacing that assumption with *active verification* of a trust state. For many organizations, the firmware, hardware, and supply chain code within devices remains the most glaring blind spot where trust remains "by default". This undercuts the principles of Zero Trust and puts the very foundation of an organization's technology at risk.

To close the gap, organizations must take steps to actively and continually assess all their assets - laptops, servers, networking infrastructure, and all their components and code - to validate Zero Trust posture and integrity. Assets should be assessed before being allowed on the network or given to end users. Assets should likewise be assessed before being granted access to sensitive resources, and then, continuously monitored to detect any changes to the integrity or posture of the asset over time.

Eclypsium's platform makes it possible for organizations to easily and automatically extend their Zero Trust program to their supply chains and down to the fundamental code and hardware in their endpoints, servers, and network equipment.

Eclypsium has built a new type of security platform designed for the needs of the supply chain. The solution is designed to support all types of critical assets (devices, software, and cloud) and can be used in any layer of the supply chain – allowing vendors to easily validate the components they receive from their suppliers and to document the security of their products, and for enterprises to validate the security and integrity of the products and services they use.

This includes the following key capabilities.

- **Detection of supply chain-specific threats and vulnerabilities** - Supply chain threats and vulnerabilities are often buried deep within products (or early in the development process) where they are beyond the view of traditional scanners. Eclypsium combines industry-leading research and proprietary, OS-independent detection mechanisms to reliably see into the deepest levels of an asset to detect vulnerabilities, misconfigurations, and threats including unknown threats or threats that have been inserted into "valid" vendor code.

- **Integrity validation** - While most security tools look for what is "bad", a supply chain solution must also verify what is "good". Eclypsium maintains what is by far the industry's largest database of critical code and components, down to the level of firmware packages, open-source branches, and dependencies. This spans many different types of assets (laptops, servers, apps, cloud), many different vendors, many different suppliers, many different software stacks, and open-source projects. This base of information allows organizations to quickly detect if any of their assets have been altered either in the supply chain or after deployment

- **Product-level validation** - Products aren't just a collection of parts - everything has to fit together and dozens of low-level systems and settings all have to work together. If a single bit gets flipped or the right setting isn't enabled, then the collective protections can fall apart. Eclypsium takes the all-important step back to validate the product or service as a whole, ensuring that all available protections and systems are properly configured and working together.

Eclypsium makes it possible to apply these core capabilities across the full lifecycle of an asset. Most importantly, Eclypsium automates the analysis and provides simple results, so that teams can quickly understand their assets without the need for supply chain or firmware expertise.

- **Procurement teams** can assess prospective vendors to identify exactly what components and code are used within a product and if they contain any vulnerabilities or misconfigurations.

- **IT teams** can quickly validate that all received assets meet the manufacturer's software bill of materials (SBOM) down to the lowest levels and verify that the integrity of the product has not been compromised.

- **Security and vulnerability management teams** can continuously monitor assets for newly identified vulnerabilities, threats, or integrity changes.

- **IR teams** can assess devices involved in incidents to ensure that the asset is free of backdoors and other low-level persistence mechanisms.

Together, this gives a cohesive and consistent approach that ensures that Zero Trust principles can be applied to the supply chain and the lowest levels of a device. To learn more about this end-to-end approach to achieving a Zero Trust network posture, visit the Eclypsium website or schedule a demo through email.