ZEDEDA

Industry Brief

Securing Intelligent, Connected Systems at the Distributed Edge

Key Considerations for Protecting Critical Operations and Turning Security into a Profit Center

TABLE OF CONTENTS

Introduction

As organizations execute on their digital transformation strategies, they are incorporating more edge computing technology into their operations to help improve business outcomes, lower costs and ensure data security and privacy. Edge computing workloads are being driven by IoT, AI, networking and security use cases across all industry verticals to deliver outcomes such as remote condition monitoring, predictive maintenance, logistics optimization, safety and security and improved customer experiences.

The edge is a continuum with inherent technical tradeoffs that drive a need for specific tools to address different parts of the paradigm. The well-established security and management solutions and practices from the data center solutions are not suitable for scaling computing deployments at the distributed edge.

This white paper is targeted at technology professionals looking to better understand the unique security challenges for distributed edge computing, along with key considerations for dealing with these challenges. It addresses these topics in the real-world context of widely-publicized security breaches from the past several years.

This paper also highlights the importance of leveraging an open foundation and prioritizing usability to meet the needs of the inherent mix of skill sets in the field. Finally, it provides insights on how the emergence of trust fabrics will foster entirely new business models and customer experiences, thereby turning security from a cost center to a profit center.

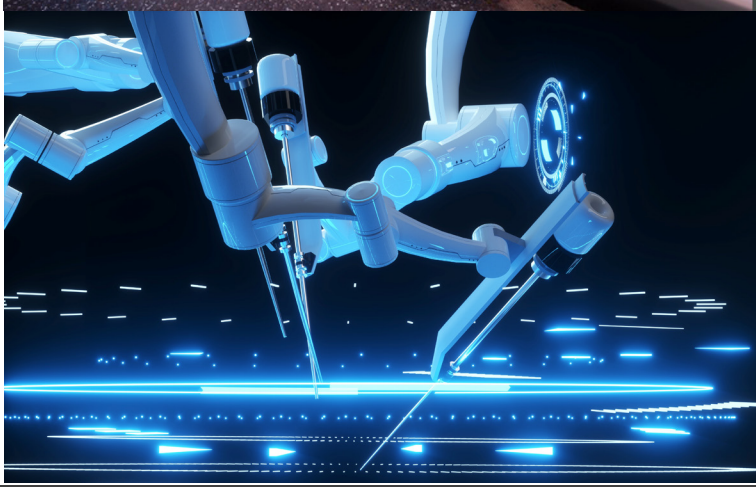

Overview of Real-World Edge Security Threats

Attacks on mission-critical networks at the edge are increasing in volume and velocity, costing organizations billions of dollars each year. In the Spring of 2021, news of the Verkada hack that breached 150,000 surveillance cameras sent shockwaves through the security world. As the attack surface of intelligent systems increasingly expands in the field, we'll see more and more attacks of this sort if we continue to leverage the same security stance and tools that we've used for decades within perimetered locations like data centers and secure telecommunication facilities.

In a hack reminiscent of Verkada, the 2016 Mirai virus infected millions of cameras and turned them into bots that together launched a massive DDOS attack on upstream networks, briefly taking down the internet in the Northeastern US. Something on the order of under twenty combinations of username and password got into all those cameras, because the developers made it too easy to change, or not even possible to change, these security credentials. Often this was due to prioritizing usability and instant gratification for users over security.

Another commonly-cited example is the massive Target data breach in 2014 that was a result of attackers accessing millions of customers' credit card information by way of a networked HVAC system. The hackers stole credentials from a HVAC contractor and were able to access the payment system because the operations network the HVAC was on wasn't properly segmented from the IT network.

The 2010 Stuxnet breach involved malware that was loaded into process control systems by using a USB flash drive to bypass the network air gap. The worm then propagated across the internal process control network, scanning for Siemens S7 software on industrial PCs. When successful, the virus would send unexpected commands to PLCs controlling industrial processes while giving the operators a view of normal operation.

Viruses like Stuxnet that focus on compromising industrial systems are especially concerning because attacks can lead to immediate loss of production, or worse life. This is compared to breaches of IT systems which typically play out over long periods of time, with compromises to privacy, financial data and IP.

All of these breaches could have been avoided had usable security tools and procedures been implemented as part of the core solution architecture.

As connected systems become more intelligent through the adoption of AI, so will the sophistication of hackers using AI to automate attacks. Gartner estimates that 30% of workloads will be deployed at the edge by 2022 and this will only increase over time. As such, it's critical to architect security into your edge solution from the start and in a flexible manner so your data and infrastructure can be augmented to address new threat vectors as fast as both your attack surface and the sophistication of hackers continues to grow.

Defining the Distributed Edge

The edge isn't a single location, rather a continuum of locations that can be categorized into three distinct paradigms based on inherent technical tradeoffs.

These paradigms are driven by:

1) Whether an edge computing node has the resources available to support abstraction in the form of virtualization or containerization

2) Whether it is on a LAN or a WAN relative to the users/processes it serves (compute for time-critical vs. sensitive workloads will always be deployed on a LAN), and,

3) Whether the compute infrastructure and applications are deployed in a centralized data center or is physically-accessible in the field. The 2020 LF Edge taxonomy white paper is a great resource to better understand the details behind these tradeoffs.

The focus of this paper is on security related to hardware and applications deployed at the distributed edge, which correlates to a blurring of the Smart Device and On-Prem Edges outlined in the taxonomy white paper. The lower limit of the distributed edge spectrum is driven by available compute resources and the upper limit by whether or not the compute node is deployed in a physically-secure data center.

Examples of distributed edge computing nodes include an industrial PC embedded in a connected machine to an IoT gateway on a factory floor to a small cluster of servers at the fringes of the data center.

These deployments can be in locations such as a drone, truck, wind turbine, oil rig, factory floor, retail store and beyond.

Location of the Distributed Edge Within the Edge Continuum

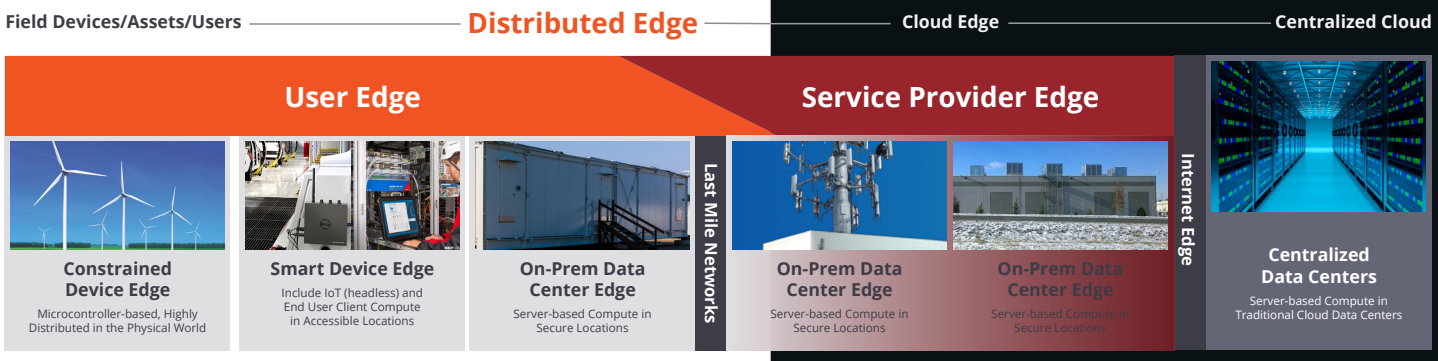

Five Key Edge Security Challenges

The distributed edge presents five unique security challenges when compared to the typical concerns in traditional, centralized data centers.

Lack of Physical and Network Perimeters: Distributed edge computing resources rarely have the defense of four physical walls or a traditional network perimeter. This requires a security approach that assumes that these resources can be physically tampered with and doesn't depend upon an owned network for defense.

Heterogeneity: The edge is at the convergence of the physical and digital worlds. In addition to a highly heterogeneous landscape of technologies, we also have to account for diverse skill sets spanning Operations Technology (OT) and IT (e.g., network and security admins, DevOps, production, quality and maintenance engineers, data scientists).

Scale: Part of the value of IoT and edge computing comes from having devices connected across the organization, providing a holistic view of operations. Over time we will see edge device deployments grow into the trillions, and traditional data center security and management solutions for security and management are not designed for this kind of scale.

Varying Priorities: As OT and IT converge at the edge, each organization's often conflicting priorities must be considered. While OT typically cares about uptime and safety, IT prioritizes data security, privacy and governance. Security solutions must balance these priorities in order to be successful.

Constrained Devices and Legacy Systems: Many IoT devices are too resource-constrained to host security measures such as encryption, plus the broad install base of legacy systems in the field was never intended to be connected to broader networks, let alone the internet. Because of these limitations, these devices and systems need to rely on more capable compute nodes immediately upstream to serve as the first line of defense, providing functions such as root of trust and encryption.

Securing Distributed Edge Computing

Addressing these unique challenges requires an orchestration solution featuring a robust Zero Trust security model. Interactions with edge hardware and applications should be performed through a remote interface featuring Role-based Access Control with multi-factor authentication, however many of the foundational security features must be built into the distributed edge hardware itself. The following five elements form the basis of a proper Zero Trust security model for distributed edge computing.

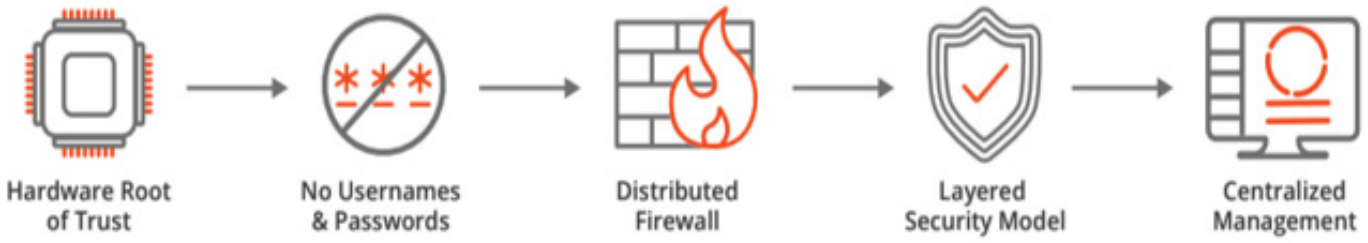

Hardware Root of Trust: Deployed edge nodes that leverage a cryptographic identity created in the factory or supply chain in the form of a private key generated in a hardware security model (e.g., TPM chip). This identity never leaves that chip and the root of trust is also used to store additional keys (e.g., for an application stack such as Azure IoT Edge). In turn, the public key is stored in the remote console.

No Edge Node Usernames and Passwords: Each edge compute node leverages its silicon-based trust anchor (e.g., TPM) for identity and communicates directly with the remote controller to verify itself. This eliminates having a username and password for each edge node in the field, instead, all access is governed through the centralized console. Hackers with physical access to an edge computing node have no way of logging into the device locally.

Distributed Firewall: Built-in granular, software-defined networking controls, enabling admins to govern traffic between applications, compute resources, and other network resources based on policy. The distributed firewall can be used to govern communication between applications on an edge node and on-prem and cloud systems and detect any abnormal patterns in network traffic. Furthermore, the solution must provide admins with the ability to remotely block unused I/O ports on edge devices such as USB, Ethernet and serial. Combined with there being no local login credentials, this physical port blocking provides an effective measure against insider attacks leveraging USB sticks.

Layered Security Model: All of these tools must be implemented in a curated, layered fashion to establish defense in depth with considerations for people, process, and technology.

Centralized Management: All security features built into edge nodes should be exposed through an API that is accessed through a remote orchestration console. Edge nodes must be designed to block unsolicited inbound instruction, instead reaching out to their centralized management console at scheduled intervals and establishing a secure connection before implementing any updates.

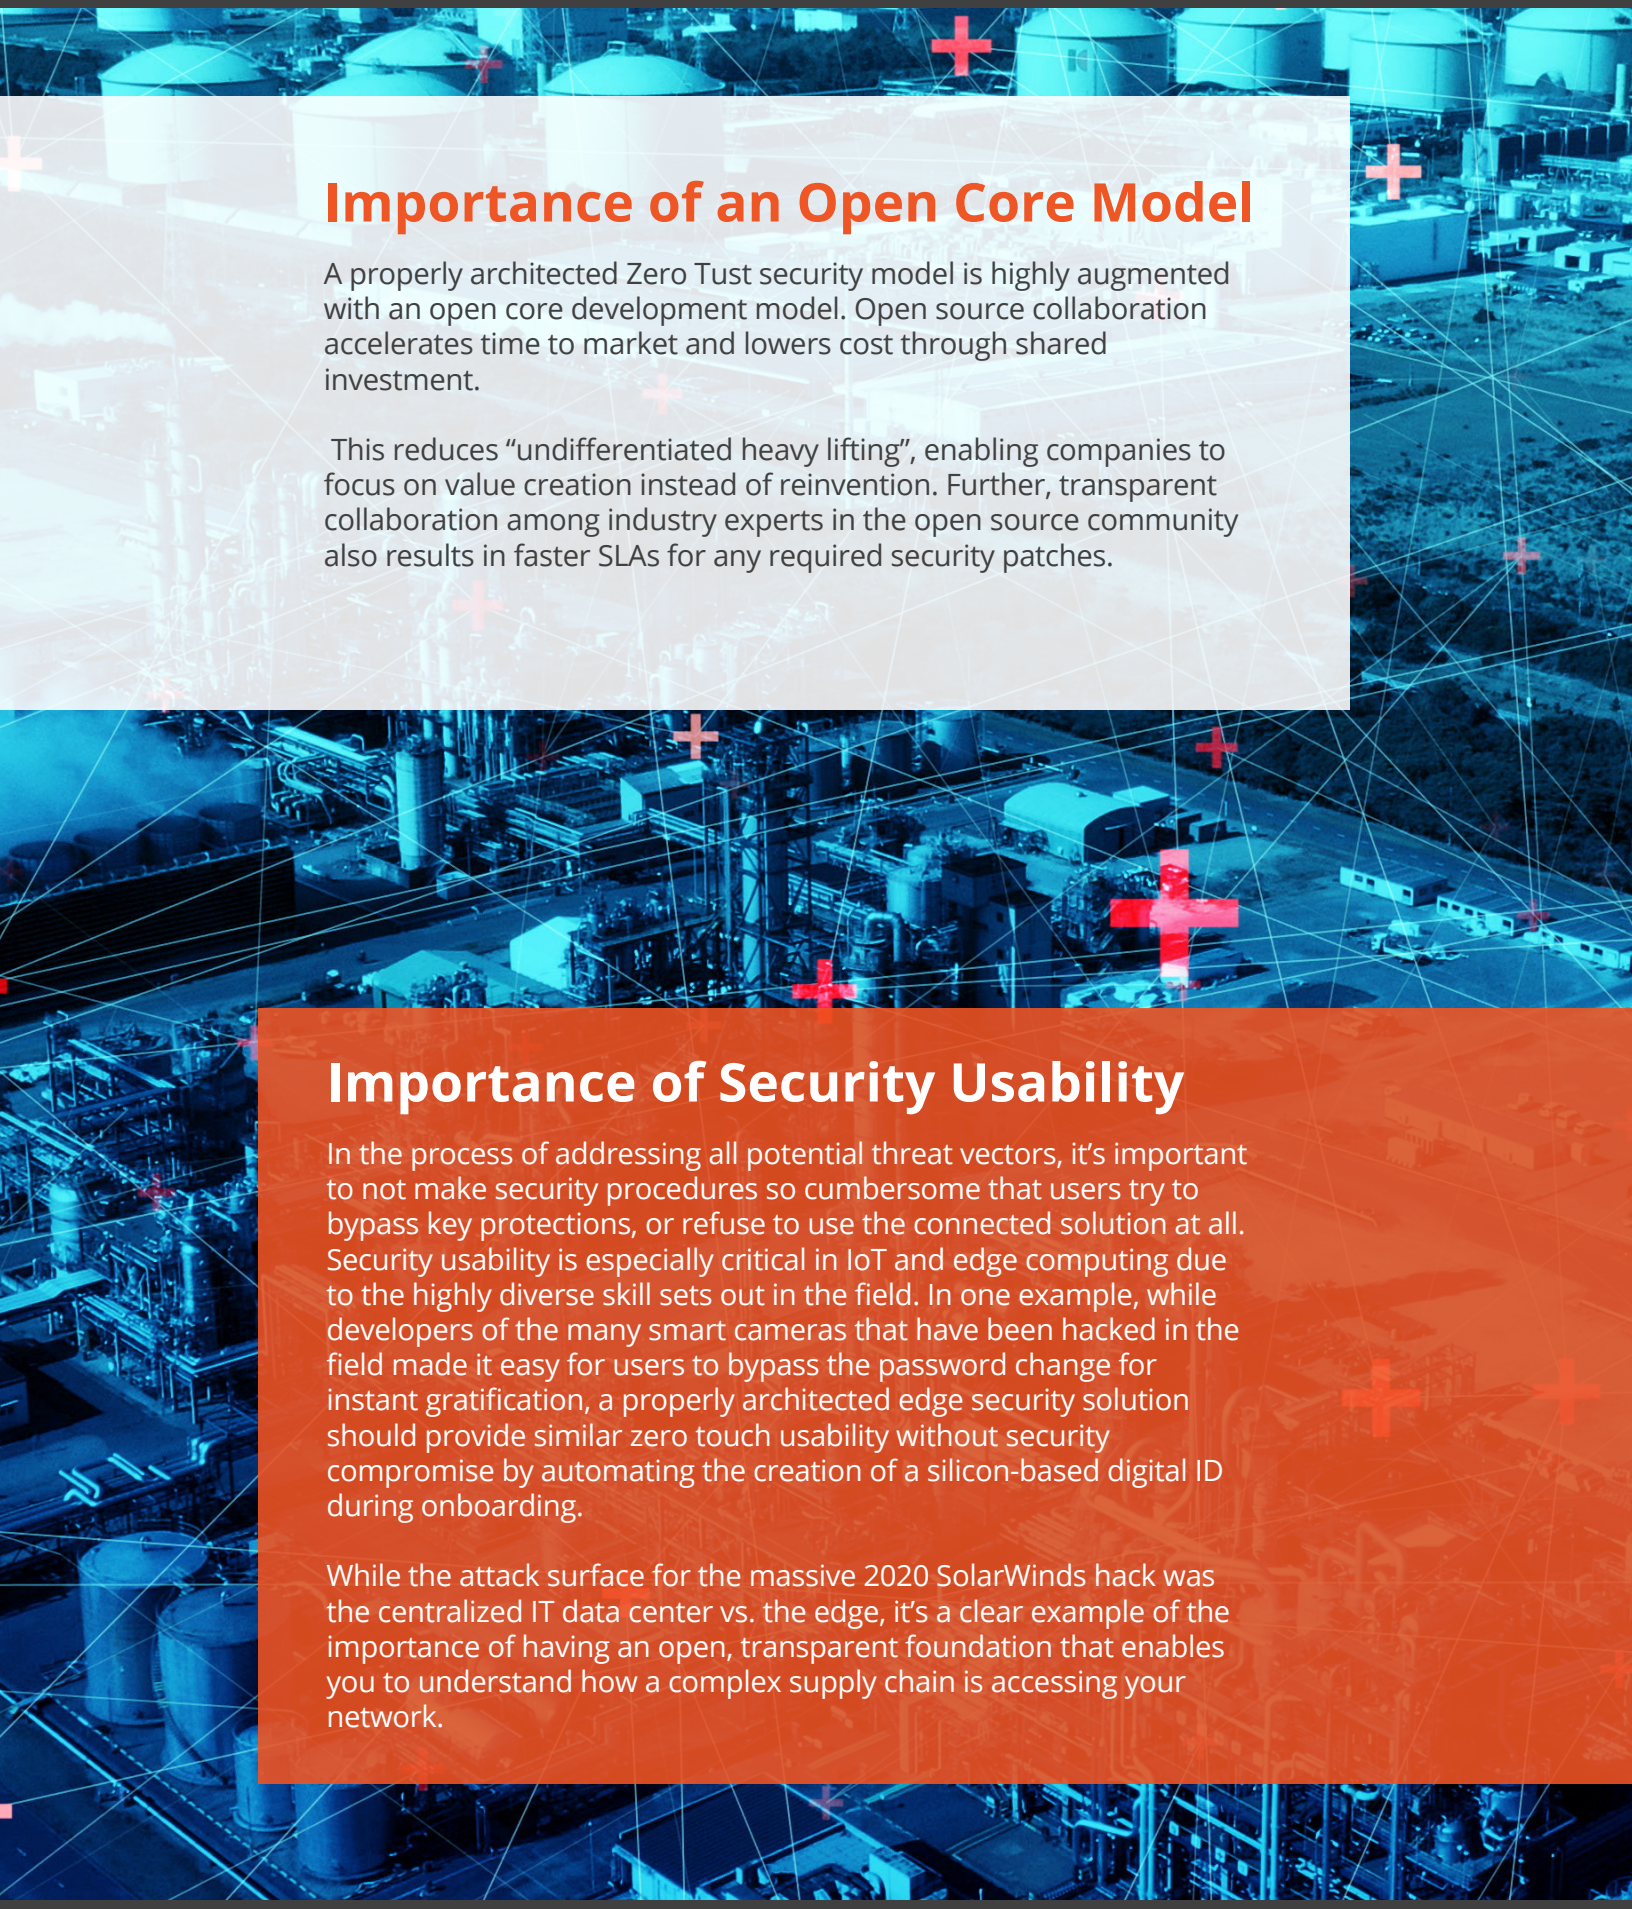

Importance of an Open Core Model

A properly architected Zero Tust security model is highly augmented with an open core development model. Open source collaboration accelerates time to market and lowers cost through shared investment.

This reduces “undifferentiated heavy lifting”, enabling companies to focus on value creation instead of reinvention. Further, transparent collaboration among industry experts in the open source community also results in faster SLAs for any required security patches.

Importance of Security Usability

In the process of addressing all potential threat vectors, it’s important to not make security procedures so cumbersome that users try to bypass key protections, or refuse to use the connected solution at all. Security usability is especially critical in IoT and edge computing due to the highly diverse skill sets out in the field. In one example, while developers of the many smart cameras that have been hacked in the field made it easy for users to bypass the password change for instant gratification, a properly architected edge security solution should provide similar zero touch usability without security compromise by automating the creation of a silicon-based digital ID during onboarding.

While the attack surface for the massive 2020 SolarWinds hack was the centralized IT data center vs. the edge, it’s a clear example of the importance of having an open, transparent foundation that enables you to understand how a complex supply chain is accessing your network.

Overview of ZEDEDA's Orchestration Solution

ZEDEDA has architected its SaaS-based orchestration solution specifically for distributed edge computing to meet the unique security, safety, uptime and usability needs of both OT and IT organizations, enabling them to focus on driving business outcomes. Access to the ZedCloud orchestration interface is sold as a subscription and the solution is built with an open core model in that it leverages the open source EVE-OS from the Linux Foundation's LF Edge organization. EVE-OS provides a robust security foundation while abstracting the complexity of the diverse hardware, connectivity and software landscape at the distributed edge.

The shared technology investment of developing EVE-OS through vendor-neutral open source collaboration is important to accommodate the diversity of hardware and software found at the edge, and because the open, vendor-neutral orchestration API establishes a de-facto standard that prevents lock-in (even to ZEDEDA).

EVE-OS can be thought of as the "Android of the Edge" in that it eliminates vendor lock-in and serves as an open anchor point for an ecosystem of hardware, software and services providers that can deliver full edge solutions. An additional benefit of EVE-OS being open source is that a growing number of hardware OEMs are supporting it from their factory, including both support for both necessary drivers and silicon-based root of trust.

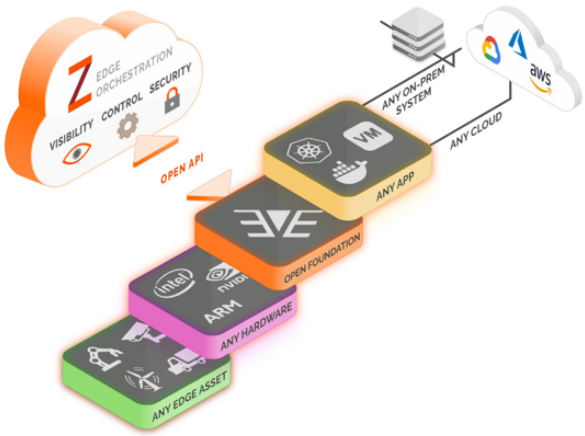

ZEDEDA Distributed Edge Orchestration Solution

ZEDEDA's solution is designed to streamline usability for both OT and IT organizations throughout the lifecycle of deploying and orchestrating distributed edge computing solution. This provides users with the confidence to connect diverse edge assets and applications to choice of on-prem systems and the cloud to solve business challenges, regardless of their technical training.

ZEDEDA has also achieved SOC-2 compliance as an additional assurance that we have the proper protections and processes in place for our cloud-based solution. For more detail, please refer to our in-depth technical whitepapers on security and overall solution architecture.

ZEDEDA's Zero Trust Security Model in Practice

Returning to the examples of real-world security breaches in the introduction, what would the impact of these attacks have looked like if these edge systems were secured and managed by ZEDEDA's orchestration solution? In short, there would have been multiple opportunities for the breaches to be prevented, or at least discovered and mitigated immediately.

No Physical Breach

In the Verkada and Mirai examples, the entry point would have had to be the camera operating system itself, running in isolation on top of top EVE-OS. However, this would not have been possible because EVE-OS itself has no direct login capabilities, rather the device is coupled with ZEDEDA's cloud (ZEDCloud) and all actions must be performed through Role-based Access Control (RBAC) in the console with multi-factor authentication and logging of all activities performed. The same benefit would have applied in the Target example, and in the case of Stuxnet, admins could have remotely locked down USB ports on the industrial PCs at the edge to prevent a physical insider attack.

Network Communications Intercepted

In all of these example attacks, the distributed firewall within EVE-OS would have limited the communications of applications, intercepting any attempts of compromised devices to communicate with any systems not explicitly allowed. Further, edge computing nodes running EVE-OS deployed immediately upstream of the target devices would have provided additional segmentation and protection. The ZEDCloud interface provides powerful, granular control and visualization of these network flows.

Activities Monitored

ZEDCloud would have provided detailed logs of all of the hackers' attempts to hack edge hardware running EVE-OS with chosen applications. It's unlikely that the hackers would have realized that the operating system or application they breached was actually virtualized on top of EVE-OS.

Automatic Quarantine

Security policies established within ZEDCloud and enforced locally by EVE-OS would have detected unusual behavior by each of these devices at the source and immediately cordoned them off from the rest of the network, preventing hackers from inflicting further damage.

One-click Mitigation

Centralized management from ZEDCloud means that updates to applications and their operating systems (e.g. a camera OS) could have been deployed to an entire fleet of edge hardware running EVE-OS with a single click. Meanwhile, any hacked application operating system running above EVE-OS would be preserved for subsequent forensic analysis by the developer.

Reputation Preserved

ZEDEDA's Zero Trust approach, comprehensive security policies, and immediate notifications would have drastically limited the scope and damage of each of these breaches, preserving the company's brand in addition to mitigating direct effects of the attack.

Augmented with an Open Ecosystem

Securing distributed edge computing solutions requires a defense in depth strategy and ZEDEDA is building an Open Edge Ecosystem to augment its robust Zero Trust security foundation. Partner offers for security within our ecosystem include solutions in areas such as active threat detection, virtual firewall and SD-WAN. ZEDEDA's ecosystem also includes software and services experts in areas such as analytics, data management and networking. Partner applications are easily deployed to edge hardware fleets running EVE-OS with a single click from within the marketplace in ZEDCloud.

Formed through a partnership between Dell, the IOTA Foundation, Intel and ZEDEDA, Project Alvarium leverages a framework that binds trust insertion technologies such as hardware root of trust, secure operating system, confidential computing, open data APIs, immutable storage and distributed ledger.

The framework includes an algorithm that scores data confidence as it flows through a given fabric, thereby enabling trusted interactions between different stakeholders.

Turning Security Into a Profit Center

The ultimate value of digital transformation is interconnecting ecosystems to foster new business models, revenue streams and customer experiences. Examples include data monetization, more sophisticated supply chains, cities offering trusted services to their citizens and service providers enabling B2B2C crossover for retailers, insurance providers, utilities into the home. In all cases, policy-based control of the balance between privacy and value must be maintained, but organizations and consumers will typically give up a little privacy in exchange for value - as long as they trust that the organization providing that value will not compromise sensitive information, IP or safety in the process.

As part of the collaboration, ZEDEDA is working towards establishing EVE-OS as one of the baseline trust insertion technologies in the overall reference stack.

In addition to the opportunity for business growth driven by trust fabrics, these fabrics will also enable organizations to address growing challenges with privacy regulations such as GDPR and combat the growing problem of fake data generated by AI. In the case of GDPR, a trust fabric would enable an organization to delete all of a customer's data in place globally with the click of a button, should they take up the right to be forgotten.

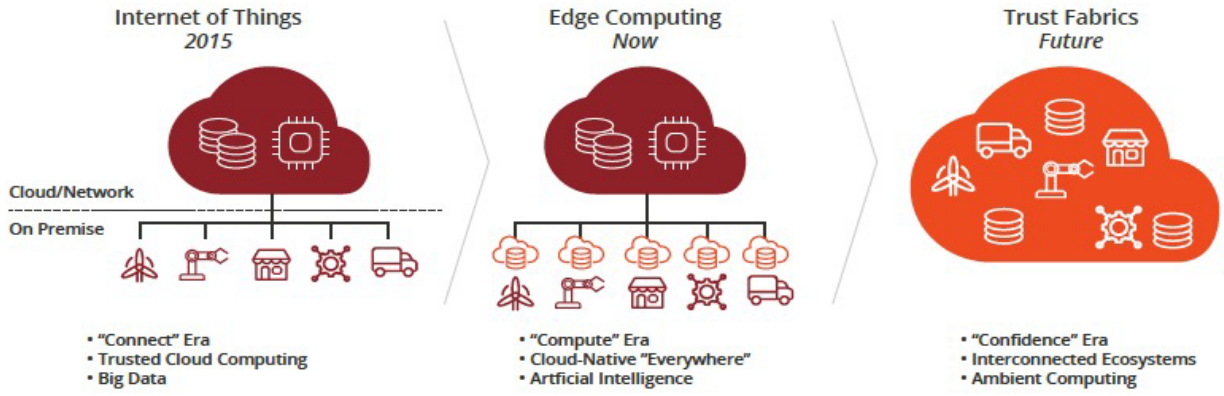

Evolution in a Connected World Requires an Open Edge

Conclusion

Security at the distributed edge begins with a Zero Trust foundation, balancing stakeholder needs, a high degree of usability, and an open core model and ecosystem. ZEDEDA has architected its orchestration solution from the ground up to include these considerations and is working with the industry to take the guesswork out so customers can securely scale their edge computing deployments with a choice of hardware, applications, and clouds, with limited IT knowledge required. ZEDEDA's goal is to enable the adoption of distributed edge computing to address key business challenges today while also getting our customers on a path to creating entirely new revenue streams over time, all without taking on unnecessary risk.

Zero Trust Foundation

Balanced for Stakeholder Needs

Usability

Open Core Model

Open Ecosystem

Architecting your edge infrastructure with an open core model enables you to smart small, leverage the expertise of the broader community, and scale over time into creating new business models and experiences for both internal stakeholders and end customers. In this regard, investments in security tools aligned with the principles in this paper will both protect your organization today and turn security from a cost center to a profit center over time.

Reach out to us today to learn more about how we can help you scale your edge solution today while being on the path to new possibilities. A world of interconnected ecosystems enabled by trust fabrics will take time to realize, but it's only achievable with investment in an open and highly secure foundation.

You can learn more about our comprehensive security approach in this technical white paper.

Contact us at sales@zededa.com to learn more about this white paper and how we can help you with your digital transformation.