SANS Institute
Information Security Reading Room

How to Create a Comprehensive Zero Trust Strategy

Dave Shackleford

SANS

A SANS Whitepaper

How to Create a Comprehensive Zero Trust Strategy

Written by **Dave Shackleford**

September 2020

Sponsored by:

Cisco

Introduction: Microsegmentation Evolution

Security professionals are starting to rethink their approach to network segmentation and access control. Some of the things they realize they need to address include:

- Focusing on trust relationships and system-to-system relationships in general within all parts of the environment. Most of the communications in enterprise networks today are either wholly unnecessary or irrelevant to the systems or applications really needed for business.
- Looking at the entire environment as potentially untrusted or compromised versus thinking in terms of "outside-in" attack vectors. Increasingly, the most damaging attack scenarios are almost entirely internal due to advanced malware and phishing exercises that compromise end users.
- Gaining a better understanding of application behavior on the network and looking at what types of network communication-approved applications really should be transmitting.
- Emphasizing approaches to protection that align with data classification types and processing environments.

Today, many traditional static network segmentation and access controls are not capable of accomplishing these objectives. Compounding this challenge is the increased use of highly virtualized and converged workloads, as well as public cloud workloads that are very dynamic in nature. There are many tools and controls available to monitor internal systems and data moving between increasingly hybrid software-based environments. However, enterprises are starting to adopt one overarching theme when designing a dynamic security architecture model: microsegmentation of specific workloads and application traffic patterns, which may facilitate a "zero trust" access control model.

Zero trust is a model where data and application behavior are the central focus points of all isolation and segmentation tactics, and all assets in an IT operating environment are considered potentially untrusted by default until network traffic and behavior is validated and approved. In its early iteration, the concept of zero trust referred to segmenting and securing the network across locations and hosting models. Today, though, there's more integration into individual services and workloads to inspect application components, binaries and the behavior of systems communicating in an application architecture. The zero trust approach does not involve eliminating the perimeter. Instead, it leverages network microsegmentation to move the perimeter in as close as possible to privileged apps and protected surface areas, along with continuous assessment of identity relationships and privileges in use. To date, there have been a wide variety of approaches taken to achieve this overarching concept of zero trust.

Enterprises are starting to adopt one overarching theme when designing a dynamic security architecture model: microsegmentation of specific workloads and application traffic patterns, which may facilitate a "zero trust" access control model.

To implement a zero trust model, security and operations teams have traditionally focused on two key concepts. First, security controls were usually integrated into workloads themselves, moving with the instances and data anywhere they might run (in data centers, virtual infrastructure, cloud environments, and so on). By creating a layer of policy enforcement that travels with workloads wherever they go, organizations have a much stronger chance of protecting data regardless of where the system runs. In some ways, this shifts security policy and access control back to the individual instances versus solely residing within the network itself. But this host-based orientation has often been associated with zero trust products and tools.

Second, the actual behavior of the applications and services running on each system need to be much better understood, and the relationships between systems and applications need more intense scrutiny than ever to facilitate a highly restricted, zero trust operations model that doesn't impact network connectivity adversely. Dynamic assets, such as virtual instances (running on technology like VMware internally or in Amazon Web Services and Microsoft Azure) and containers, are difficult to position behind "fixed" network enforcement points. That has led organizations to seek to adopt a network isolation strategy that only allows traffic to flow between approved systems and connections, regardless of the environment they are in.

Microsegmentation Technology

The National Institute of Standards and Technology (NIST) recently released SP 800-207, "Zero Trust Architecture,"[1] which focuses on a zero trust architecture model. This model includes the following elements in a comprehensive data/network security strategy that meets zero trust principles:

- **Identity—**Role and privilege definitions for user/account access
- **Credentials—**Authentication controls, such as passwords and keys
- **Access management—**Controls and policies that govern what assets and services can be accessed, and from where
- **Operations—**The overarching tools and processes needed to define, implement, maintain and monitor zero trust architectures
- **Endpoints—**Distinct systems and workloads that are part of a zero trust environment
- **Hosting environments—**The environment where a zero trust architecture is implemented (for example, a data center or cloud provider infrastructure)
- **Interconnecting infrastructure—**Tools and platforms that facilitate connectivity to and from assets both within a zero trust architecture and external to it

With these components and principles as a foundation, zero trust architecture has evolved to include a variety of controls that can facilitate a comprehensive approach to protecting systems and data.

Challenges of Microsegmentation

It hasn't always been simple to implement a microsegmentation access control model. Many organizations, when looking to implement microsegmentation, have encountered numerous challenges, including:

- **Silos of technology—**In most organizations, there is a wide variety of technology in place that likely includes legacy operating systems and applications, development tools and platforms, third-party applications and services along with "homegrown" applications, and many more. Commonly, some of these technology types don't work well with others, leading to silos of technologies. Examples might include a single primary network vendor that doesn't interoperate with other network tools, hardware limitations or a specific OS version that doesn't support updates to critical applications without performance issues. Zero trust architecture and technology usually require a broad degree of access to infrastructure and platforms, and a large number of technology silos can diminish effectiveness.

[1] https://csrc.nist.gov/publications/detail/sp/800-207/final

- **Lack of technology integration—**For zero trust technology to provide the greatest benefit, some degree of technology integration is necessary. For example, endpoint agents will need to be installed on systems and mobile devices, and any issues with this integration can easily derail a zero trust implementation. Similarly, the policy engine for zero trust should be able to integrate with user directory stores in use to continually assess accounts, roles and permissions. One of the most pressing technology integration hurdles can emerge when third-party solutions and platforms aren't available or won't run on specific cloud provider infrastructure. As more organizations move to a hybrid or public cloud deployment model, this issue can be crippling.
- **Rapidly changing threat surface and threat landscape—**The threat landscape is constantly changing, of course, which can potentially lead to challenges with some types of zero trust technology that only focus on one type of environment or are limited in deployment modality. For example, many attackers are now targeting end users as the primary ingress vector to a network and can functionally "assume" that user and device identity. A zero trust technology that doesn't incorporate strong machine learning with behavioral modeling and dynamic updates will likely miss some of the most critical types of issues we see today with attacks. Some of these include:
    - **Lateral movement between systems—**This is a common scenario in today's attack campaigns where initial ingress into a network environment is usually followed by probes and attempts to compromise additional systems nearby. Detection and prevention of lateral movement requires an understanding of trust relationships within the environment, as well as real-time monitoring and response capabilities.
    - **Insider threats—**These types of threats are notoriously difficult to detect because insiders usually already have access granted, thus potentially limiting the efficacy of identity or endpoint-specific controls in a zero trust design. Only deep understanding of behaviors that are expected or unexpected in specific user-oriented interaction with data and applications can help detect insider threats.
- **Lack of simplicity and ease of use—**The more complicated a security technology is, the less likely it will be rolled out effectively or consistently applied. One of the biggest challenges with zero trust and microsegmentation tools today is the complexity and lack of simple administration and policy design/implementation. An intuitive console and policy engine, well-designed reporting and metrics, easy installation and monitoring for endpoint and workload agents, and strong support for numerous technologies and platforms (especially cloud environments) are critical for any enterprise-class technology that seeks to help organizations implement zero trust.

- **New infrastructure requirements—**To support newer generation firewalls and software-defined networking infrastructure, many organizations will need to upgrade and modify their infrastructure. This might include network systems like switches and routers, virtual appliances in converged and cloud-based environments, and monitoring tools and services.

Critical Capabilities

To implement microsegmentation and zero trust effectively, solutions should provide several critical capabilities including:

- Policy discovery
- Policy validation and simulation
- Host-based enforcement
- Capability to address several operational challenges

We describe each of these capabilities in more detail in the following sections.

Policy Discovery

To get started, a policy discovery effort is critical, especially in brownfield environments where numerous existing technologies are already running. Most microsegmentation and zero trust technologies include some form of scanning and discovery tools to find identity use and privilege allocation, application components in use, traffic sent between systems, device types, and behavioral trends and patterns in the environment. Security teams should work with identity and access management (IAM) teams or those responsible for key IT operations roles and functions, such as user provisioning and management, to understand the different groups and users within the environment, as well as the types of access they need to perform job functions. The same should be done for all types of user devices, primarily laptops and desktops in use by privileged users.

Policy Validation and Simulation

Once some basic discovery has been done, any mature access control (microsegmentation) policy engine should be able to start linking detected and stated identities (user, groups, devices, privilege sets, and so on) with network traffic generated by specific services and application components across systems. To get the most benefit from a zero trust strategy, this stage of planning and project implementation needs to carefully accommodate business- and application-centric use cases. Security teams should plan to evaluate what types of behavior are actually necessary in the environment versus those that may be simply allowed or “not denied” explicitly. This takes time and careful analysis of systems that are running.

This planning and discussion should lead security and network teams to develop policies that isolate and compartmentalize applications and specific traffic in an environment based on traffic and identity requirements. A microsegmentation solution should allow for simulated network segmentation that closely aligns with a specific type of system or workload once these have been discovered. Zero trust's traditional concept of network microsegmentation strives to prevent attackers from using unapproved network connections to attack systems, move laterally from a compromised application or system, or perform any illicit network activity regardless of environment. Essentially, zero trust facilitates the creation of affinity policies, where systems have relationships and permitted applications and traffic, and any attempted communications are evaluated and compared against these policies to determine whether the actions should be permitted. This happens continuously, and effective zero trust control technology will also include some sort of machine learning capabilities to perform analytics processing of attempted behaviors, adapting dynamically over time to changes in the workloads and application environments. These capabilities should be readily tested during the policy validation and simulation phase of any deployment. In essence, the chosen platform should be capable of policy generation based on detected traffic and communication patterns, and simulation of attacks or disallowed traffic patterns that demonstrate the policy's effectiveness.

By potentially eliminating lateral movement, a zero trust microsegmentation model also reduces the post-compromise risk when an attacker illicitly gains access to an asset within a data center or cloud environment. Security architecture and operations teams (and often DevOps and cloud engineering teams) refer to this as "limiting the blast radius" of an attack because any damage is contained to the smallest possible surface area and attackers are prevented from leveraging one compromised asset to access another. This works not only by controlling asset-to-asset communication, but also by evaluating the actual applications running and assessing what these applications are trying to do. For example, if an application workload (web services such as NGINX or Apache) is legitimately permitted to communicate with a database server, an attacker would have to compromise the system and then perfectly emulate the web services in trying to laterally move to the database server (even issuing traffic directly from the local binaries and services installed).

Host-Based Enforcement and Benefits

To manage access across a diverse range of workloads and network environments, some type of host-based enforcement capabilities should be a part of any solution. Most microsegmentation solutions accomplish this by deploying a custom agent to any workloads that incorporate custom firewall controls. In some cases, microsegmentation agents include (or can integrate with) anti-malware and exploit protection capabilities and tools. Endpoint monitoring for both malicious signatures and behaviors can significantly improve the security posture of trusted workloads, with automated quarantine and elimination of detected threats adding even more value.

By integrating host-based security capabilities with a monitoring and enforcement engine that incorporates network traffic monitoring and identity roles and privileges, organizations can create and maintain a more complete model of zero trust over time.

Operational Challenges

A number of operational concerns can accompany the shift to microsegmentation and zero trust. Figure 1 shows the key stages of implementation.

Figure 1. Key Stages of Implementation

A breakdown of key stages in any zero trust implementation should include:

1. **Discovery: Discover, catalogue, and classify data and assets—**The discovery phase is one of the more important phases in a zero trust architecture model because the different types of assets and data in any environment will need to be continually found and assessed against defined policies. In most traditional data centers, discovery has proven to be challenging at scale due to a lack of cohesive visibility into all network segments. In a zero trust environment, discovery should focus on network monitoring that discovers, catalogues, and classifies data in all storage and application deployments. Mapping data flows for sensitive data scenarios is another important function that discovery tools should facilitate.

   Operational challenges in this phase include implementing agents and network monitoring in all network segments, as well as classifying applications and assets appropriately.

2. **Deployment: Microperimeters and architecture—**In a zero trust deployment, some sort of microsegmentation engine must be in place to enact access control policies defined by a central policy engine. This engine may include cloud-native microsegmentation tools such as Amazon EC2 Security Groups, as well as internal identity-aware policy engines that can restrict and limit access between assets running in on-premises data centers and cloud provider environments. Once identities are confirmed and validated—through integration with directory services and other identity stores—a least privilege access model should be enforced through policy.

   Operational challenges in this phase include integration with identity stores; categorization of users, groups and roles; and determination of appropriate privileges for a particular application scenario or use case.

3. **Detection: Network and application traffic monitoring—**Monitoring will be in place to continuously detect and track network traffic and local system processes in the environment, mapping usage of applications, services exposed, user interaction models and patterns of network behavior in expected and unexpected application usage scenarios.

   Operational challenges include aligning security operations teams and incident investigation workflows with any detected anomalies or unusual workload activity.

4. **Response: Policies and automation actions—**In today's dynamic environments, automated response actions are becoming more commonplace for specific use cases and playbooks. Response actions can be "triggered" through continuous monitoring and detection of events and behaviors that likely indicate compromise or attempted compromise, and may include quarantine of assets, suspension or deletion of systems and workloads, suspension or removal of user accounts and identities, and more. For this to happen at scale, most enterprises likely need powerful analytics and environment integration.

   Operational challenges are primarily policy tuning and adjustment.

In addition to these specific tactical challenges, many organizations will also find that larger organizational governance challenges arise. One of the key decisions teams should make early in the adoption cycle of microsegmentation and zero trust is policy ownership and input. Network teams, identity teams and security teams all have a definitive stake in the development of policies, and they should adopt a collaborative model of policy development and updates if possible. It is common for a single team to "own" microsegmentation platforms and tools, but policies need to include a variety of stakeholders for input and discussion.

Another key operational challenge is the ongoing "Day 2 compliance"—in other words, the maintenance of configuration and compliance state once a solution and policies have been implemented. Because microsegmentation and zero trust are major technology shifts for any organization, it's important to enact a routine evaluation of policy effectiveness and issues that arise and develop an escalation path for troubleshooting policies as applications and environments change. This is especially true for highly dynamic cloud infrastructure and workloads.

Compliance Examples

There are a number of ways that microsegmentation and zero trust initiatives can positively impact compliance requirements and control choices. In the following sections, we demonstrate several examples of industry compliance requirements that can be effectively met with microsegmentation technology.

SWIFT Customer Security Program (CSP)

The Society for Worldwide Interbank Financial Telecommunication (SWIFT) is a global provider of financial messaging services. Banks and other financial institutions use SWIFT to send secure messages and orchestrate financial transactions globally, and SWIFT users must adhere to a number of strict security requirements. Table 1 shows how SWIFT requirements can be met with zero trust technology.

Table 1. Zero Trust Meets SWIFT Requirements[2]

SWIFT Requirement	How Zero Trust Applies
Securing the SWIFT environment	Under SWIFT requirements, all SWIFT-related assets and communications must be entirely segmented from any other IT infrastructure. Microsegmentation and zero trust allow for software-defined segmentation policies to isolate assets and control data flow.
Knowing and limiting environment access	Because identity and privilege allocation is a core element of any microsegmentation and zero trust deployment, it is easy to create a least privilege model using segmentation policies aligned with user and group membership.
Detecting and responding to threats	With microsegmentation, network policies and monitoring at the workload level can be used to monitor and respond to suspicious and malicious behavioral patterns and detected signatures of intrusion and attacks.

[2] www.guardicore.com/swift-compliance/

HITRUST

The Health Information Trust Alliance (HITRUST) Common Security Framework (CSF) is a robust privacy and security controls framework that has been used extensively to evaluate and certify healthcare-related organizations that must comply with security and privacy regulations such as the Health Insurance Portability and Accountability Act (HIPAA). Zero trust deployments may help to meet the HITRUST CSF control requirements[3] shown in Table 2.

Table 2. Zero Trust Meets HITRUST Requirements

HITRUST Requirement	How Zero Trust Applies
Protection of remote diagnostic and configuration ports	Logical access to diagnostic and management ports for connected devices and systems can be controlled with microsegmentation policies.
Network segregation	Segregation of distinct network zones and segments can be easily defined through software-based microsegmentation policies that limit access to only approved users, groups and workloads.
Control of network connections	Zero trust policies can limit access to the organization's shared networks and applications to only the appropriate users and groups through identity integration.
Isolation of sensitive systems	Microsegmentation and zero trust allow for software-defined segmentation policies to isolate sensitive assets and control data flow to and from them.
Asset inventory	Microsegmentation and zero trust discovery tools can collect information about connected assets in a network environment, providing a comprehensive inventory of both assets and running processes and services.
Control of technical vulnerabilities	With a focus on continuous monitoring of all assets at both the network and workload levels, microsegmentation and zero trust can facilitate rapid detection of possible vulnerabilities being attacked or exploited. These same tools also can provide remediation and response actions and controls.

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) outlines a set of security controls and practices that organizations accepting or processing credit card payments must have in place. The PCI DSS applies to merchants and service providers of all sizes involved with credit card payment processing globally. There are a number of areas where microsegmentation can help meet PCI DSS compliance requirements, including those in Table 3.

Table 3. Zero Trust Meets PCI DSS Requirements

PCI DSS Security Focus	How Zero Trust Applies
Focusing/narrowing PCI DSS compliance scope	Limiting the scope of PCI DSS to the smallest necessary systems and transactions is essential to control operational overhead and threat surface, as well as costly audits. Software-defined segmentation policies can enforce access and communications restrictions that meet the PCI DSS requirements for scope limitation.
Improving network and system visibility	PCI DSS includes several requirements that call for restricting communication and blocking all insecure services, protocols and processes. Workload and network monitoring and control policies for microsegmentation can effectively define and maintain all access controls in a PCI DSS network.

A Zero Trust Roadmap for Executives

Executives should carefully evaluate technology in the space and determine which solutions most adequately meet the needs of the organization. They also should look for products that work in both internal and public cloud environments where possible—this will almost always require an agent-based solution. When implementing a zero trust microsegmentation strategy, follow this basic roadmap:

1. **Start with passive application discovery, usually implemented with network traffic monitoring.** Allow for several weeks of discovery to find the relationships in place and coordinate with stakeholders who are knowledgeable about what "normal" traffic patterns and intersystem communications look like. Enforcement policies should be enacted later, after confirming the appropriate relationships that should be in place, along with application behavior.

[3] For more information about HITRUST CSF control requirements, see "HITRUST CSF Version 9.4," June 2020, https://hitrustalliance.net/hitrust-csf/ [Registration required.]

2. **Design zero trust architecture based on how data moves across the network, and how users and apps access sensitive information.** This will assist in determining how the network should be segmented, and where protection and access controls should be positioned using virtual mechanisms and/or physical devices between the borders of different network segments.
3. **Take the time to categorize systems and applications, which will help in building application traffic baselines and behaviors.** More advanced zero trust tools integrate with asset "identities," which may be part of an application architecture, aligned with a business unit or group, or representative of a specific system type.

At its heart, zero trust is really a modern take on least privilege access that is dynamically updated and tied to both network- and identity-based behaviors and components. For any zero trust project, it's critical to help build a least privilege business case that looks at the various ways a robust zero trust enablement technology can improve access control and organizational security across the board.

Examples of common business drivers include:

- **Diverse endpoints and users—**The number and types of endpoints and users/groups functioning within an organization is growing, in some cases rapidly. Especially for large organizations with a massive and diverse set of technologies and user variations in place, choosing a technology that can accommodate all of these could vastly simplify the implementation and maintenance of access control operations for the foreseeable future.
- **Cloud and new service layers—**With the drive to hybrid and public cloud deployment models, the need to find technologies that support a wide range of hosting and infrastructure deployment grows rapidly. Today, few zero trust access control platforms have strong support across numerous cloud provider environments, as well as on-premises data centers. Consolidation and integration and deployment support could easily help shift least privilege strategy to a unified technology solution that works in all environments.
- **Business continuity and contingency planning—**Many organizations are now realizing that business continuity and contingency planning needs to better embrace the unexpected and unknown scenarios that could easily occur in the future. There is no way to plan for all scenarios, but embracing more flexible approaches to endpoint technology and rapidly changing business use cases could drive access control models toward a ubiquitous zero trust strategy and technology implementation.
- **Machine learning/AI enhancements—**In a nutshell, machine learning is training machines to solve problems. This is most often applied to problems that can be solved by repeated training and pattern recognition development. Machine learning and AI techniques can help security professionals and technologies recognize patterns in data, and this can be extremely useful at scale. For example, collected threat intelligence data provides perspective on attacker sources, indicators of compromise and attack behavioral trends. Threat intelligence data can be aggregated, analyzed at scale using machine learning, and processed for likelihood/predictability models that are then fed back to zero trust access control policies and platforms to help dynamically update detection and response capabilities.

As a zero trust deployment matures and becomes more stable, organizations will want to determine how effective the approach has been in their environments. Industry-wide zero trust metrics are currently few and far between, as most organizations are still in early phases of deployment and have not yet developed mature, long-term strategies for this type of architecture. Potential metrics for a zero trust strategy might include:

- Number of applications identified and mapped (often tracked per network range, business unit or geographic locale)
- Number of tracked and labeled identities catalogued
- Percent reduction in network access control alerts (after implementing enforcement policies)
- Percent reduction in compromised systems and application workloads after zero trust implementation

Conclusion

A zero trust architecture should include authentication and authorization controls, network access and inspection controls, and monitoring/enforcement controls for both the network and endpoints. No single technology currently will provide a full zero trust design and implementation. Organizations need a combination of tools and services to provide the full degree of necessary coverage. For most, a hybrid approach of both zero trust and existing infrastructure will need to coexist for some period of time, and emphasis should be placed on the common components and control categories that could suitably enable both, such as identity and access management through directory service integration, endpoint security and policy enforcement, and network monitoring and traffic inspection. As zero trust frameworks mature and evolve, so will standards and platform interoperability, likely facilitating more streamlined and effective approaches overall.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Sponsor

SANS would like to thank this paper's sponsor:

Last Updated: September 22nd, 2020

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS October Singapore 2020	**Singapore, SG**	**Oct 12, 2020 - Oct 24, 2020**	**Live Event**
SANS Community CTF	**,**	**Oct 15, 2020 - Oct 16, 2020**	**Self Paced**
SANS SEC504 Rennes 2020 (In French)	**Rennes, FR**	**Oct 19, 2020 - Oct 24, 2020**	**Live Event**
SANS SEC560 Lille 2020 (In French)	**Lille, FR**	**Oct 26, 2020 - Oct 31, 2020**	**Live Event**
SANS Tel Aviv November 2020	**Tel Aviv, IL**	**Nov 01, 2020 - Nov 06, 2020**	**Live Event**
SANS Sydney 2020	**Sydney, AU**	**Nov 02, 2020 - Nov 14, 2020**	**Live Event**
SANS Secure Thailand	**Bangkok, TH**	**Nov 09, 2020 - Nov 14, 2020**	**Live Event**
APAC ICS Summit & Training 2020	**Singapore, SG**	**Nov 13, 2020 - Nov 21, 2020**	**Live Event**
SANS FOR508 Rome 2020 (in Italian)	**Rome, IT**	**Nov 16, 2020 - Nov 21, 2020**	**Live Event**
SANS Community CTF	**,**	**Nov 19, 2020 - Nov 20, 2020**	**Self Paced**
SANS Local: Oslo November 2020	**Oslo, NO**	**Nov 23, 2020 - Nov 28, 2020**	**Live Event**
SANS Wellington 2020	**Wellington, NZ**	**Nov 30, 2020 - Dec 12, 2020**	**Live Event**
SANS OnDemand	**OnlineUS**	**Anytime**	**Self Paced**
SANS SelfStudy	**Books & MP3s OnlyUS**	**Anytime**	**Self Paced**