Building Your Zero Trust Strategy with **NIST 800-207** and **Arista NDR**

Zero Trust has moved from a buzzword to reality very quickly as organizations have accelerated their digital transformation efforts in light of the COVID-19 pandemic. To use a cliché, however, zero trust is a journey - not a destination. Unfortunately, for many companies, this journey doesn't have a well-defined end in sight. As a result, NIST put together a framework for the Zero Trust journey that enables organizations to gauge the effort and path to improve their cybersecurity posture[1]. As NIST documents go, this isn't a page-turner, but it aligns nicely with what Arista has heard from many of our customers as they set out on their own zero trust endeavors.

Before we discuss our key takeaways from this document, one helpful suggestion for those looking to read the NIST 800-207 guidance is to approach it in a non-linear fashion. While section 2 of the NIST document provides a basic primer on zero trust architecture for the uninitiated, reading section 7 next will help you understand what you are signing up for. Section 5 then provides a good perspective on how your threat model will evolve post-migration. While this is not covered explicitly in the NIST document, in working with our customers, it is useful to benchmark how this new threat model will compare to your current threat model.

Spoiler alert: Zero trust does not mean the threats go away! Section 5 clarifies that you must put in place a set of controls to ensure your zero trust architecture is secure. Finally, if you can't get through it all, section 3 is the other required reading section, in our opinion. What is this going to look like in your environment?

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

Migrating to a Zero Trust Architecture (ZTA)

We often talk to teams that believe Zero Trust is a project that has a well-defined and relatively short timeline. Unless you are rebuilding the network from scratch or have a relatively simple network, section 7 will help you understand the ongoing processes needed to build and maintain a zero trust architecture (ZTA).

Figure 1: NIST Guidance on Migrating to a Zero Trust Architecture

Based on that table of contents (Figure 1), Section 7 goes into a fair amount of detail about the considerations and steps in rolling out ZTA. As you might expect, an endeavor of this type is rarely "one and done" since architects do not have the luxury of starting from scratch or dramatically changing network topology in one fell swoop.

As Figure 2 shows, the process is an ongoing effort with a feedback loop.

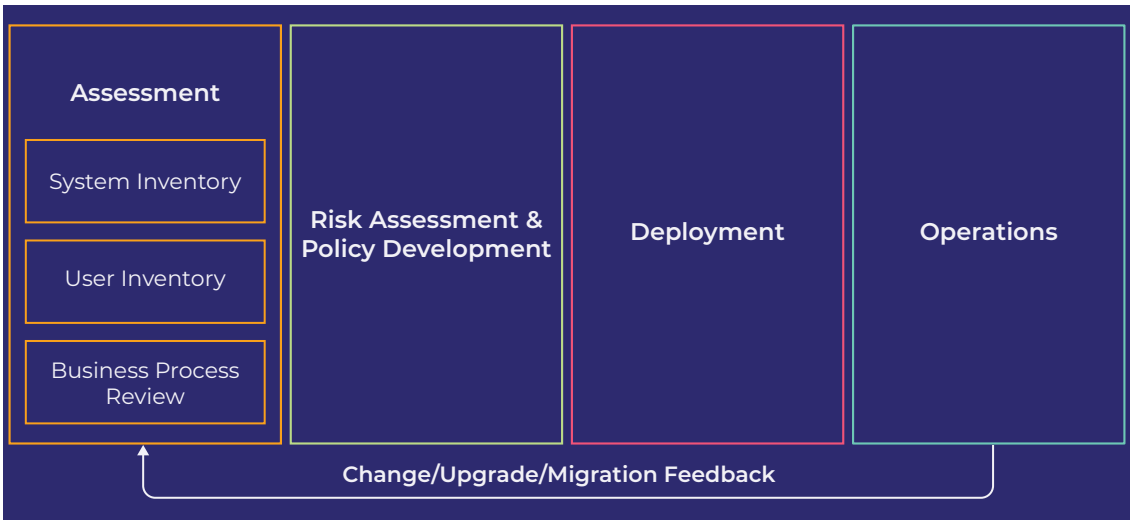

Figure 2: Zero Trust Architecture Migration Process[2]

Before migrating to ZTA, you must understand what resources/assets need to be protected and who is accessing them. Starting with Active Directory or a CMDB is certainly a good option, but how confident are you in their accuracy? In Arista NDR deployments, we find that the average customer "manages" less than 50% of the actual resources and devices on the network and thus lacks visibility into much of the unmanaged infrastructure. As sections 7.3.1 and 7.3.2 point out, getting a handle on that is foundational to an effective ZTA strategy.

[2] Adapted from NIST 800-207

A good starting point in building your ZTA strategy is to address visibility and control at the network level. An increase in shadow IT, cloud IoT, BYOD and third-party devices on the enterprise network require an agentless approach in addition to traditional agent-based solutions.

Arista NDR deploys with EntityIQ™, an AI-based security knowledge graph that identifies, profiles and tracks all devices, users and applications on an enterprise network with just a network connection so that your analysts immediately have context on what is on their network without requiring new integrations or agents.

In addition, Arista NDR's adversarial modeling language (AML) can be used to surface and monitor business processes and detect when those processes are being undermined. For instance, an Arista customer in the retail space maintains ZTA with respect to their PCI[3] enclave. Arista NDR monitors this environment for attempts by unauthorized devices or users to access the environment. This capability was used pre-ZTA to inventory all the devices and users that needed access versus those that could be consolidated, allowing the customer to define a narrow allow list. Post migration, Arista NDR's capabilities now provide oversight of any attempts to subvert this allow list—whether malicious or not.

Finally, as SP 800-207 notes, the ZTA journey is not without bumps. Section 7.3.6 stresses the need for ongoing monitoring and tuning. Doing that requires the right telemetry from the network and the ability to analyze that telemetry at scale without requiring humans to crunch large amounts of data or manually build effective rules. The AML provides the flexibility to do just that! It enables the comparison of access patterns post migration with historical patterns, allowing administrators to quickly identify and rectify network interruptions.

The Pre- and Post-ZTA Threat Model

Given today's aggressive threat landscape, security teams must always have a clear understanding of the organizational threat model. This threat model serves another purpose during a zero trust migration: it helps identify which processes bear the most risk to the organization and, therefore, should move to ZTA first. Arista NDR automatically scores the risk of every entity, whether internal devices, users, or applications, as well as external destinations such as domains, ASNs, and IP addresses. This capability has enabled organizations to prioritize business processes with higher-risk entities earlier for migration to ZTA.

What about a post-ZTA threat model? Zero Trust doesn't make all threats disappear, but it improves security and resiliency. SP 800-207 makes the case that organizations cannot have a secure network even if it was architected according to zero trust principles.

Figure 3: NIST Guidance on Threats to Zero Trust Architectures

So, what does "doing it right" mean? Organizations would be well-served to start with the threats they are concerned about today. For instance, is there a decrease in traditional "inside the firewall" applications, with a corresponding increase in SaaS and other cloud-based platforms? Do monitored security metrics show an evolution of threats toward more non-malware (a.k.a. living-off-the-land) threats? Industry statistics show most breaches today show no traces of malware. Instead, attackers use tools within the

[3] https://www.pcisecuritystandards.org/

environment, abuse insider credentials, or use popular sites like Twitter and Google Drive for command and control. Security teams must now detect malicious intent within an activity that looks identical to normal business activity. Once the threat model is built, the organization can look for foundational controls to implement secure zero trust.

Examples include:

Section 5.3 emphasizes the importance of context when making runtime zero trust decisions. For example, if user "Bob" from the finance department is logging in from a corporate device and accessing a resource he usually does, that is acceptable. But if that access is coming from a personal device, mobile phone, or workload in your cloud, perhaps it isn't. Similarly, if "Bob" attempts to access the source code control system, even if that access has not been explicitly disallowed, that might be an indication of malicious intent. Today, uncovering threats like these often require significant investments in manual threat hunting, making the process both onerous and not repeatable. Arista NDR's autonomous hunting flags these behavioral threats and triggers automated response actions on analysts' behalf via integration with other parts of your infrastructure, including orchestration and ticketing systems, network enforcement solutions, and endpoint agents.

Section 5.4 highlights another challenging network security trend: more than 90% of web traffic is now encrypted. Arista's own data shows that more than 50% of so-called east-west traffic is encrypted even inside the network perimeter. Moreover, organizations are increasingly hesitant to decrypt this traffic due to potential policy and privacy violations. From a technical perspective, applications using end-to-end encryption and TLS 1.3 present further roadblocks. Attackers are not shy about this trend and increasingly use encrypted traffic to evade network detection. Accordingly, Gartner recently reported that over 70% of malicious traffic is now encrypted. As prescribed in NIST SP 800-207, Arista NDR uses data science to perform encrypted traffic analysis, thus working within technology, privacy, and policy constraints and ensuring security teams do not fly blind. Such use cases include the identification of application-specific protocols and communications and the nature of the encrypted traffic that is going over the wire. To illustrate, it might not be of concern when an encrypted file is transferred over a Zoom meeting session, but if that file is transferred from a Power- Shell script to a Dropbox account might merit investigation.

ZTA in Practice

Section 3 of the NIST document does a great job describing the components of ZTA (Figure 4). Specifically, section 3.4 outlines some of the key network requirements for implementing ZTA effectively. This includes the ability to effectively distinguish between managed and unmanaged devices and the ability to capture and analyze all network traffic, including encrypted traffic (as described above).

[4] https://www.crowdstrike.com/blog/global-threat-report-foreword-2020/
[5] https://www.bondcap.com/report/itr19/

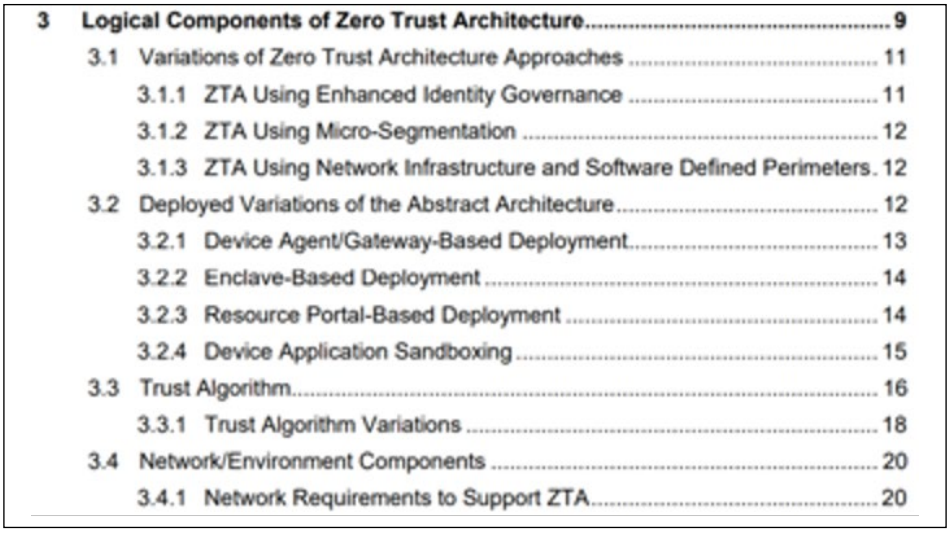

Figure 4: NIST Guidance on Components of Zero Trust Architectures

The concepts underlying ZTA sound simple, but it is essential to plan for practical implementation challenges while planning your journey. As an illustrative example, the Trust Algorithm described in section 3.3 must make decisions by accounting for information including the user or identity requesting access, the type of access being requested, the minimum security requirements to allow access (e.g., multi-factor authentication, updated patch levels, etc.), and data from both internal and external threat intelligence. Putting these controls in place might allow us to ask highly specific questions of our network, such as "is this access being made from a device that is an outlier from its peer group in requesting this specific access" or "has the user or device in question shown weak signals of compromise, like repeated connections to a relatively rare domain?"

Arista NDR automates this analysis using a combination of machine learning approaches. Legacy solutions rely primarily on unsupervised learning to spot anomalies from "normal" baselines. Often, anomalies are not malicious (resulting in false positives) and conversely, pre-existing compromises are missed because they are assumed to be "normal." These false negatives are even more dangerous because security teams unknowingly live with the risk. The Arista NDR platform compares entity behaviors not just to what was observed in the past but what is seen from the peer group and the rest of the organization. The lack of reliance on baselining also speeds up both time to value and ongoing operational costs, especially when compared to solutions that require 1-3 months of "training" and then constant retraining every time the environment changes.

Summary: Arista Enables Your Zero Trust Journey

Figure 5: Arista Enables Your Zero Trust Architecture Journey

Building a zero trust architecture is a noble goal, especially in times like these where IT teams are being asked to make the organization productive irrespective of the location of users and needed resources. However, as NIST 800-207 points out, this is not a shiny red button that can magically transform the organization. Embarking on this journey needs careful thought to ensure organizational productivity and enhanced security.

As Figure 5 illustrates, Arista provides you with a number of the tools necessary to migrate and operationalize a zero trust architecture effectively.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office
1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062