SCTE CABLE-TEC EXPO 22
SEPTEMBER 19-22 • PHILADELPHIA, PA

2022 Fall Technical Forum
SCTE • CABLELABS • NCTA

Zero Trust Security Architecture for The Enterprise

A Technical Paper prepared for SCTE by

Sarah Weinstein
Vice President, Engineering
Comcast
Philadelphia, PA
(215) 286-3412
sarah_weinstein@cable.comcast.com

Christopher Zarcone
Distinguished Engineer
Comcast
Moorestown, NJ
(856) 638-4116
christopher_zarcone@cable.comcast.com

Stephen Zevan
Director, Product Management
Comcast
Philadelphia, PA
(215) 286-8275
stephen_zevan@cable.comcast.com

Table of Contents

List of Figures

Title **Page Number**

1. Introduction

1.1. Zero Trust Architecture

Zero Trust Architecture (ZTA) is an information security model that de-emphasizes computer networks as trust factors, focusing instead on strong user & device authentication and contextual, risk-based authorization. In this paper, we:

- Review the background and history of ZTA (Section 2).
- Summarize the history of our organization's Zero Trust journey (Section 3).
- Discuss highlights of our organization's ZTA program, including its guiding principles, security focus areas, and governance (Sections 4 – 6).
- Reflect on lessons learned (Section 7).
- Consider future directions for our ZTA program (Section 8).
- Evaluate ZTA as it relates to service providers (Section 9).

2. Background

2.1. History of Zero Trust

Zero Trust Architecture is an information security model based on the principle of **Never Trust, Always Verify**. The term "zero trust" can be traced to a 1994 doctoral dissertation by Stephen Paul Marsh of the University of Stirling, and the concept was later popularized in a 2010 whitepaper by John Kindervag of Forrester Research. ZTA argues that the traditional perimeter network security model is obsolete, and that modern security programs should instead focus on authentication of security principals, security policy compliance, and continuous risk assessment.

ZTA assumes that all computer networks are untrusted by default, and that enterprise networks are no different – and hence no more secure – than non-enterprise or public networks. Viewed in this light, Never Trust, Always Verify requires:

- Strongly authenticating, authorizing, and auditing all access to resources and services.
- Measuring user and device compliance against organizational security policy.
- Continuously assessing the security posture of all principals (users, devices, resources).

2.2. Elements of ZTA

2.2.1. Foundational

The foundational elements of a ZTA are:

- **Users**. Users are human actors and include the organization's employees, contractors, business partners, and other third parties.
- **Devices**. A device is a physical computing resource designed to be used interactively by a user. Desktop computers, laptop computers, smartphones and tablets are all examples of user devices.

- **Resources**. Resources are repositories of organizational data, and as such, primary targets of attack. They include both physical and virtual machines, as well as the server operating systems, containers, services, and applications running on such machines.

At the most basic level, Users use Devices to access Resources. All three elements are security principals, in that they have identities and can attempt to perform actions. A ZTA authenticates identities and authorizes actions as required.

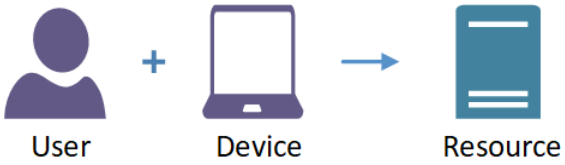

Figure 1 – Users, Devices & Resources

Additional dimensions to each element must be considered as well:

- For any given Resource, User audiences might vary (employee, business partner, guest, etc.). Additionally, the Users may be local to one of the organization's offices, or they may be remote.
- Resources might be owned and administered organizationally at one or more premises locations. Alternatively, the Resources might be hosted by a service provider (e.g., **Cloud** services).
- Devices might be owned and administered organizationally. Alternatively, Devices might be owned and administered by a third party, or personally owned by their Users.

2.2.2. Security Support Structure

The Security Support Structure (SSS) is an all-encompassing term used to describe the set of services that enable ZTA. To varying degrees, the foundational elements depend on each of these services. Common SSS components include:

- Application Access Frameworks – on-premise, off-premise
- Configuration Management Databases (CMDB)
- Directory Services (LDAP)
- Endpoint Detection & Response (EDR)
- Identity & Access Management (IAM) & Identity Providers (IdP)
- Multi-Factor Authentication (MFA)
- Mobile Device Management (MDM)
- Public Key Infrastructure (PKI)
- Security Information and Event Management (SIEM)
- User Behavior Analytics (UBA)

Collectively, the SSS bolsters the Zero Trust posture of Users, Devices, and Resources. An example SSS can be depicted as follows (significant integrations illustrated):

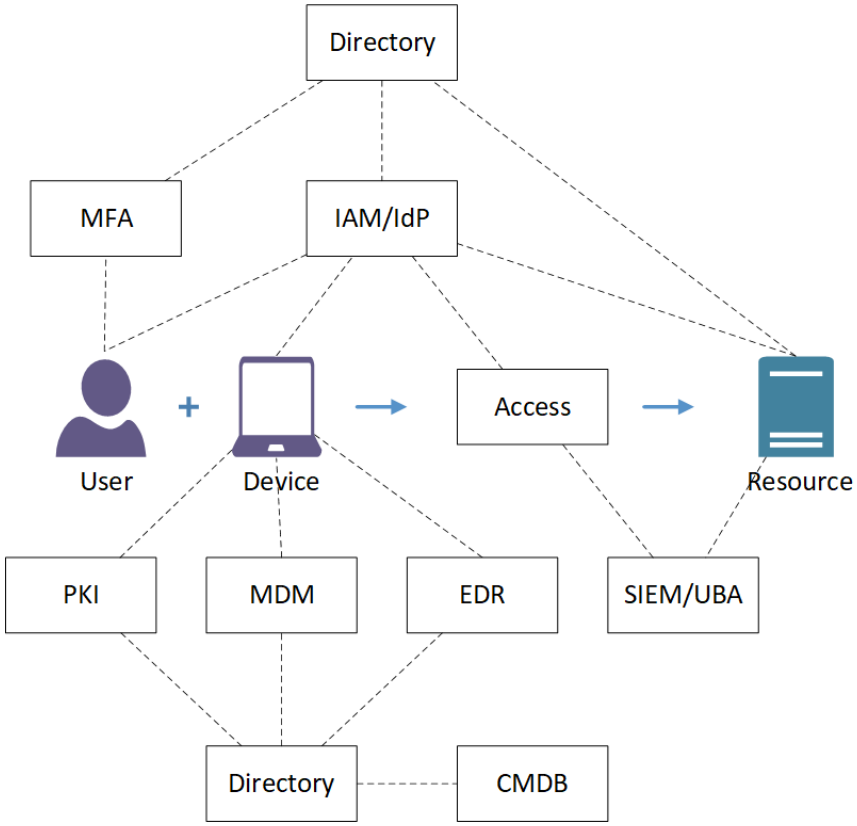

Figure 2 – Security Support Structure

2.3. ZTA Fundamentals

All Users and Devices begin a work session in a state of Zero Trust. That is, absolutely no security relevant attributes are assumed or inferred about a User or a Device. Examples:

- Assertions of Identity
  - "I'm Alice!"
  - "I'm Bob's computer!"
- Assertions of Network Location
  - "I'm on the Internal network!"
- Assertions of Compliance
  - "I'm fully patched!"
  - "My storage is encrypted!"

Users and Devices must therefore establish trust before obtaining access to Resources. Trust is established in three primary ways:

- **Authentication**. User and Device identities must be confirmed with high confidence. This typically requires strong authentication, which can be defined as:
    - Multi Factor Authentication (MFA) for interactive User access. Examples include knowledge factors (passwords, PINs) combined with physical or logical possession-based factors (time- or event-based One-Time Password (OTP) generators, digital certificates, challenge/response mechanisms, FIDO 2 tokens).
    - Cryptographically sound authentication for Device-level, non-interactive access. Examples include shared symmetric keys, asymmetric key pairs, digital certificates, refresh tokens, and so forth.

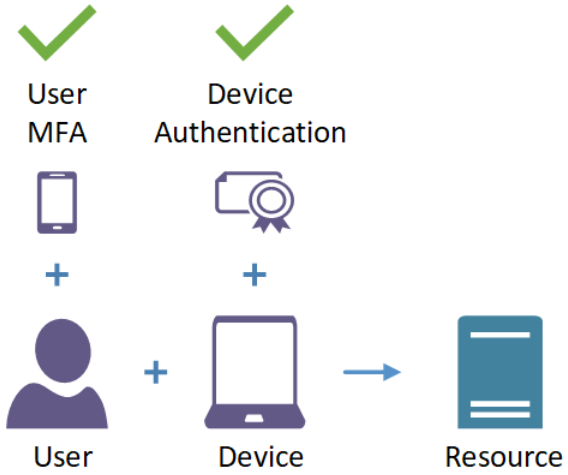

Figure 3 – Establishment of Trust Via Authentication

- **Security Policy Compliance**. Users and Devices must be assessed for compliance with organizational security policy. This could include factors such as:
    - User Authorization Checks
        - There are many possible criteria for user authorization, including directory group memberships and/or other directory attributes, requested Resource, work schedule (time of day, day of week, etc.) and so forth.
    - Device Authorization Checks
        - Determination that device properties and configuration are acceptable (e.g., idle UI timeouts are enabled, storage media is encrypted)
        - Verification that all expected security tooling is installed and operational
        - Confirmation that OS and/or security tool patch levels are suitably recent

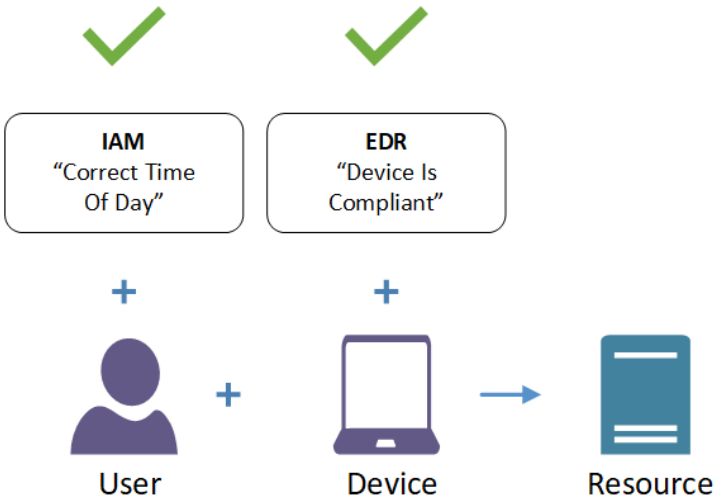

Figure 4 – Establishment of Trust Via Risk & Compliance Assessment

- **Continuous Risk Assessment (Context)**. Risk is not a static property to be evaluated only once, at session initiation. Rather, risk is a dynamic property, and the level of risk posed by a User or Device can change mid-session. For example, a Device could move from a trusted state to an untrusted state (e.g., EDR detects a malware infection, or Device IP address reputation changes from benign to malicious). For this reason, User and Device risk should be continuously assessed throughout the duration of a work session. This continuous risk assessment is referred to as Context.

Once a sufficient level of trust in a User and a Device has been established, access may be granted to one of more Resources as specified by policy.

2.4. ZTA Frameworks

The descriptions above provide a generic, high-level understanding of basic ZTA principles. In practice, however, ZTA cuts across many different policy and technology domains, and implementing a full ZTA program can require considerable effort and attention to detail.

Industry, academia, and standards bodies have responded to this challenge by producing several reference frameworks for ZTA, to guide organizations in their approach to Zero Trust. Here, we review a few examples.

2.4.1. Google – BeyondCorp

Starting in 2014, Google published a series of influential whitepapers documenting their own multi-year Zero Trust implementation. Collectively referred to as "BeyondCorp," the six whitepapers propose Google's vision of a ZTA, with special emphasis on migrating to a network architecture without traditional perimeter controls.

Google's ZTA features five design objectives:

1. Securely Identifying the Device
2. Securely Identifying the User
3. Removing Trust from the Network
4. Externalizing Applications and Workflows
5. Implementing Inventory-Based Access Control

In general, BeyondCorp shifts access controls from network edges and perimeters to individual User and Device identity, embracing the oft-repeated maxim that "identity is the new perimeter." This focus allows an organization's employees, contractors, and business partners to work more securely and remotely from any location without the need for a traditional remote access technologies like Virtual Private Networks (VPN). Key to this architecture is a next-generation access model for resources, an Internet-Facing Access Proxy (alternatively called an "Identity-Aware Proxy" or simply an "Access Proxy"). The Access Proxy enforces the security requirements of traditional VPN – strong authentication, strong encryption, etc. – in a centralized manner, without the protocol overhead of IPSec and similar tunnelling approaches.

2.4.2. NIST – Special Publication 800-207

In 2020, National Institute of Standards and Technology (NIST) developed its own framework for ZTA (NIST Special Publication 800-207). The standard does not propose a definitive architecture for Zero Trust; instead, it proposes several architectural variations, emphasizing the advantages and disadvantages of each. Each of these variations, however, reflect a common set of design goals (or "tenets" as they are described in the standard). They are as follows:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture.

2.4.3. Other Frameworks

Other Zero Trust frameworks of note:

- Although not exclusively billed as a "Zero Trust" architecture, Gartner's Continuous Adaptive Risk and Trust Assessment (CARTA) framework embodies many basic ZTA principles, such as continuous risk assessment.

- The US Cybersecurity and Infrastructure Security Agency (CISA) has published a "Zero Trust Maturity Model," intended primarily for consumption by federal agencies. CISA framework focuses on five domains, each with maturity levels ranging from "traditional" to "optimal."
- The United Kingdom's National Cyber Security Centre (NCSC) outlines a six-pronged Zero Trust network architecture in its broader published guidance for device security.

2.4.4. Using the ZTA Frameworks

There is no uniform, "one size fits all" approach to ZTA. Organizations will have to evaluate their current IT strengths, capability gaps, and the overall threat landscape to develop their own unique approach to ZTA. To the extent that ZTA frameworks (such as those mentioned above) can guide and assist with this effort, so much the better.

2.5. Scope of ZTA

It should be noted that the scope of a ZTA typically applies to:

- An organization's enterprise Users, Devices, and Resources.
- Third parties (contractors, business partners, vendors, etc.) performing enterprise duties on behalf of the organization.

Although the same fundamental concepts apply, information security management of an organization's customers or subscribers should be treated separately from its enterprise needs.

2.6. Embarking on ZTA

It is often stated that Zero Trust is not a specific product or technology. Rather, it is an assemblage of multiple components, all operating in concert, to realize a specific security posture. As such, any approach to ZTA is most fundamentally an effort in systems integration. In terms of project management, organizations would be well advised to approach ZTA from this perspective.

It is also often stated that Zero Trust is not a specific project deliverable or operational practice, but rather a "journey." This choice of language reflects the reality that – for all but the smallest of organizations – migrating to a Zero Trust posture will require sustained effort over a significant period of time. For example:

- In a 2022 executive memorandum from the US Office of Management and Budget titled "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," OMB noted that "Transitioning to a zero trust architecture will not be a quick or easy task for an enterprise as complex and technologically diverse as the Federal Government." The memo further requires that specific Zero Trust goals be met within a three-year timeframe (end of Fiscal Year 2024).

Accordingly, claims of "overnight" transformations from legacy networks to Zero Trust should be regarded with skepticism. The Zero Trust journey is one of incremental successes – it is a marathon, not a sprint.

3. Brief History of ZTA at Comcast

3.1. 2019 – The Start of Our ZTA Program

Comcast has been on its ZTA journey since 2019. The journey began as most large initiatives do – by building the business case for Zero Trust and why it was important to the organization.

At the time, the need to consider new security strategies was self-evident – the rate of publicly disclosed security incidents was increasing dramatically. To cite just one example of many, in the third quarter of 2019, the RiskBased Data Breach QuickView Report noted that there had been 5,183 breaches during the year, exposing 7.9 billion records. Compared to the 2018 Q3 report, the total number of breaches had increased 33% year over year, and the total number of records exposed had increased 112%.

Zero Trust was quickly identified as the security strategy most worth pursuing. In terms of socializing this new approach to information security, the "Castle and Moat" metaphor really helped paint the picture whenever we needed to explain the concept of Zero Trust to internal stakeholders. For example, a traditional medieval castle may feature high walls made of brick or stone, perhaps further protected by a moat and a drawbridge-style gate. The walls and moat form the "perimeter" by which the residents of the castle are protected from external threats. Entry and egress to the protected castle is governed by the gate. However, if an intruder can get past the gate, they have unfettered access to what is inside the castle. There may be some impedance once past the walls and moat, but the reality is that access can be quite open.

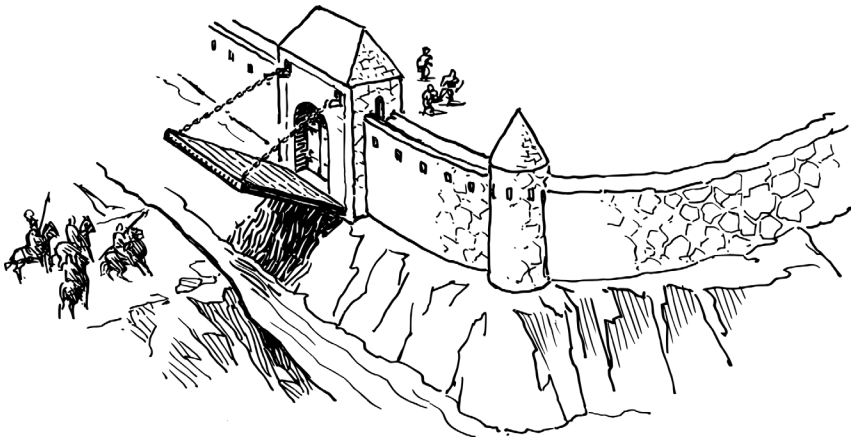

Figure 5 – Classic Physical Perimeter Defense Architecture

"Drawbridge" by j4p4n. Licensed under Creative Commons Zero 1.0 License. openclipart.org

Replace "castle" with "internal and/or enterprise network" and "gate" with "firewall" and one arrives at the traditional network security perimeter model, which was developed in the late 1980s and served the Internet community well for the better part of two decades.

In our new reality, however, with the rise of business partner integrations, public cloud services, and application deployment at network's edge, it is not enough to rely on the traditional network perimeter model. We need to verify ever User, every Device, and every attempt to access Resources (systems, applications and information), every time.

The "Castle and Moat" analogy served to ground technical and non-technical professionals and executives, which in turn was important to gain their support. A tremendous amount of stakeholder engagement was required on how we were approaching Zero Trust because it required ongoing partnership across the organization, with entities including our network engineering team, information technology, and our geographic divisions. Another important partner in this endeavor was Comcast's Cybersecurity Guild. With over 200 members, the Cybersecurity Guild is a rich and interactive community of employees who are cyber-minded or have day-to-day responsibilities for cybersecurity in their business units.

Although developing the business case involved significant effort, it found a receptive audience, as our executive leadership has long regarded cybersecurity as a critical partner to the business. We didn't have to convince everyone that it was important to embark on the Zero Trust journey – everyone knew that we just had to do this. The challenge was figuring out how we were going to do this and devising a strategy in a relatively short period of time.

We soon began developing the overall structure of our Zero Trust program. Some of the earliest influences included the BeyondCorp framework (discussed in Section 2.4.1) and Microsoft, which was underway with enterprise solutions and SSS to enable Zero Trust at scale.

Our leadership, architects, engineers, and other stakeholders eventually coalesced around a three-pronged program structure, consisting of:

- **Guiding Principles**
- **Security Domains**
- **Program Governance**

We discuss these in more detail in Sections 4, 5, and 6, respectively.

3.2. 2020 – COVID-19

In early 2020, the COVID-19 pandemic emerged as a worldwide health concern. The pandemic created a seismic shift in how our employees and business partners worked, and everyone was living with the reality of remote work and further diminished network perimeters. Companies like ours could no longer rely on the safeguards that we took for granted before the pandemic – working together, in common facilities, on an enterprise network that provided a basic degree of assurance. Our new way of working made it much more clear why Zero Trust is so important.

Arguably, because of the COVID-19 pandemic, Zero Trust is even more relevant today than it was in 2019.

The threat landscape also changed significantly during COVID-19. The Harvard Business Review highlighted the challenge that with the migration to remote work during COVID-19, cyberattacks increased exponentially. Businesses saw more attacks of every kind, but the headline for 2020 was ransomware attacks, which increased 150% over the previous year. Equally worrisome, the amount of money paid by victims of these attacks increased more than 300% in 2020.

COVID-19, however, did not deter our ZTA efforts. Quite to the contrary, COVID-19 accelerated them. Moving a significant portion of our workforce to remote work re-emphasized the need to move beyond the traditional perimeter security model.

3.3. 2021 to Present – Acceleration

Recently (and particularly within the last one to two years) there has been more industrial awareness of security as a business enabler. Conversations have shifted, with business clients increasingly asking about their service provider's adherence to Zero Trust practices.

In May 2021, we saw the issuance of the White House's "Executive Order on Improving the Nation's Cybersecurity." The order states (emphasis added):

> The United States faces persistent and increasingly sophisticated malicious cyber campaigns that threaten the public sector, the private sector, and ultimately the American people's security and privacy. The Federal Government must improve its efforts to identify, deter, protect against, detect, and respond to these actions and actors. The Federal Government must also carefully examine what occurred during any major cyber incident and apply lessons learned. But cybersecurity requires more than government action. Protecting our Nation from malicious cyber actors requires the Federal Government to partner with the private sector. ***The private sector must adapt to the continuously changing threat environment, ensure its products are built and operate securely, and partner with the Federal Government to foster a more secure cyberspace.***

4. Guiding Principles

As previously discussed, the Zero Trust model assumes that every security principal (every User and every Device) is a potential threat until verification. This assumption holds regardless of whether those principals' network location appears to be internal or external to the organization.

We also assume that despite our best efforts to adopt sound information security practices and secure our enterprise, incidents can and will occur. Many factors give rise to this reality, most notably Zero-Day exploits (exploits directed at previously unknown and/or undisclosed vulnerabilities).

These assumptions give rise to the two guiding principles of our ZTA:

1. **Never Trust, Always Verify.** Users and Devices must establish trust before obtaining access to Resources.
2. **Prepare To Be Breached.** Our networks, systems, and applications must be designed to limit the impact of incidents.

Principle #1 was discussed at length in Section 2 but bears repeating. Users/Devices begin a work session in a state of Zero Trust, and Users/Devices must therefore establish trust before obtaining access to Resources. Trust is established through authentication, policy compliance, and risk assessment. Importantly, Trust must ***not*** be established by the perceived physical or logical network location of a User or Device.

Principle #2 asserts that if incidents cannot be prevented with certainty, the next best strategy is to plan for their eventuality. The goal is to minimize the degree to which Resources are affected by incidents in the broader technical environment. This is often referred to as minimizing the "blast radius" of an incident, preventing spillover effects and lateral network movement.

These two guiding principles informed our work across our entire ZTA program. For every significant architectural decision, we would ask the same two questions:

1. How does this decision improve the trustworthiness of Users and/or Devices?
    a. Or conversely, how does this decision further reduce our reliance on networks as trust factors?
2. How does this decision help mitigate the effects of an incident?

5. Security Domains

Next, we applied the guiding principles across three broad security domains:

- **Security Hygiene**. Hygiene concerns the security posture of our main ZTA principals – Users, Devices, and Resources.
- **Network Microsegmentation**. Microsegmentation focuses on partitioning our network environments into smaller, workload-specific enclaves.
- **Application Access**. Access governs how our Resources are exposed to Users and Devices, regardless of their physical or logical network locations.

Each of these domains is discussed in turn.

5.1. Security Hygiene

Like many large companies, Comcast has a complex IT environment that was built through years of product and platform growth and acquisitions. There was some foundational work that needed to be addressed to ensure that as a company, we were ready for our Zero Trust journey. This foundational work was called **Security Hygiene** (or simply "Hygiene"), defined as a focus on the security posture of our security principals: User, Devices, and Resources.

We defined the scope of this work across several parallel workstreams, termed **Pillars**. For each Pillar, we defined the problem space that we needed to resolve, and definitions of what success looked like. We also aligned on quantified targets for success wherever possible.

The Pillars were defined as follows:

1. **User Identity & Access Management**
2. **Resource Identity & Access Management**
3. **Asset Ownership**
4. **Device Identity & Management**
5. **Visibility & Hardening**

5.1.1. User Identity & Access Management

Section 2.3 notes that strong authentication for Users plays a pivotal role in ZTA. This more broadly requires an examination of how an organization performs Identity & Access Management (IAM) for its Users. IAM refers to an entire suite of tools and services for provisioning, managing, and deprovisioning User identities. IAM is also used to manage User privileges and entitlements, as determined by a User's title, position, and/or organizational roles.

A core component of IAM is the Identity Provider (IdP). The IdP's primary responsibilities are:

- **Authenticating Users**. This can be accomplished using one or more of the mechanisms described in Section 2.3.
- **Asserting User Identities to Resources**. This can be accomplished using protocols such as Kerberos, Security Assertion Markup Language (SAML), and Open Authorization & OpenID Connect (OAuth, OIDC). Proprietary protocols also exist.

Ideally, an organization should implement one, centralized IAM platform with as few IdPs as possible (again, ideally one).

The alternative is multiple IAM platforms with competing/duplicative identity stores that complicates:

- User provisioning and deprovisioning
- Assignment of user authorizations and entitlements
- Compliance
- Reporting
- Far from least important, the User experience (e.g., access to some systems works whereas access to other systems does not, confusion over which credentials or MFA/SSO mechanisms to use for a given system, and so forth).

5.1.2. Resource Identity & Access Management

Resource IAM is essentially the complement to User IAM. If ZTA requires Users to produce strongly authenticated assertions of identity, it becomes necessary for Resources to consume those assertions. Otherwise, the identity assertions serve no purpose.

This generally requires Resources to align with:

- The same Identity Provider (IdP) employed for Users
- The same authentication protocols (SAML, OAuth/OIDC) supported by the IdP

The amount of effort required to integrate a Resource with a modern IdP varies considerably. It can range from the trivial (e.g., simple changes to a configuration file) to the substantial (e.g., installation and integration of third-party authentication libraries, custom code, and extensive testing). The latter is often the case with older, "legacy" Resources, for which modern protocols like SAML/OIDC did not exist at the time they were deployed.

5.1.3. Asset Ownership

Information systems cannot be protected if they are unknown to the organization. This requires an accurate inventory of:

- Devices
- Resources
- The organizational owners of such Devices and Resources

In most organizations, this requires the creation and maintenance of a CMDB or similar inventory/asset management system. As with user identities, the ideal is a single CMDB that provides complete coverage for all Devices and Resources in the organization. Multiple, competing CMDBs exhibit many of the same challenges experienced with multiple, competing IAM platforms – inconsistent views of asset ownership, impaired ability to protect unregistered assets, and so forth.

5.1.4. Device Identity & Management

As with User identity, Device identity is extremely important for purposes of ZTA. It allows organizations to assign specific roles and entitlements to the Device (for example, an administrative workstation, compared to workstations intended for normal workforce use), establish unique accountability for Device activity, and so forth.

Device identity should be established through cryptographically strong mechanisms. Examples include digital certificates, public key pairs (independent of certificates), bearer tokens such as OAuth refresh tokens, and so forth. These credentials are typically created at the time a Device is activated or joined to an administrative framework (such as an Active Directory domain, a Kerberos realm, and so forth).

Device management is equally important as establishment of identity. A Device is said to be "managed" if an organization has administrative control of the Device and can apply security policy to it; otherwise, the Device is considered unmanaged. An insufficiently managed or unmanaged Device poses risk to a ZTA. The device may lack appropriate security tooling like EDR, or may not be configured to require commonplace security controls, such as idle UI session timeouts.

A subset of device management, Mobile Device Management (MDM) refers to the administration of mobile devices, such as laptops, smartphones and tablets. MDM caters to the unique characteristics of mobile devices, in that such devices are typically more limited in their capabilities than server, desktop or workstation computers. Mobile devices are also more easily lost or stolen than their desktop counterparts; as such, MDM is frequently used to enforce confidentiality-preserving security requirements, such as storage encryption.

5.1.5. Visibility & Hardening

Visibility (often also called Telemetry) refers to the monitoring of systems and networks in real time.

Resource Visibility refers to the degree to which system-level properties and activity can be observed on a Resource. Examples include:

- Resource utilization (CPU, memory, storage, network & other input/output)
- Process and thread activity
- File- or filesystem-level artifacts (existence of absence of specific files, etc.)
- (Windows) Existence and/or values of specific registry settings
- Collection and/or examination of system logs

Depending on the level of visibility required, some management frameworks (such as membership in an administrative domain or MDM) may provide sufficient data. In other cases, the deployment and management of specialized agent software may be required.

Resource Visibility also assists with maintaining accurate asset management (Pillar 3).

By contrast, **Network Visibility** refers to the monitoring of network communications in real time. Ideally, such monitoring can be summarized into data flows, or sequences of network packets from sources to destinations.

Communications may be monitored:

- Via the capabilities of standard network infrastructure (both physical and virtual), such as routers and switches.
- Via specialized, media-specific collection devices, often referred to as network taps.

The purpose of network visibility is to detect signals that possibly indicate:

- Lateral movement within an organization
- Data exfiltration

Hardening applies to Devices, Resources, and network infrastructure (collectively, "nodes") and primarily consists of:

- The disablement of all unnecessary network services on a network-attached node.
- Enabling all practical security controls on the node's remaining network services.

The main objective is to reduce the **attack surface** of network-attached nodes.

In security taxonomy, "attack surface" is the set of all possible entry points into a Device, a Resource or other information system. A system with a large amount of attack surface is thought to be more vulnerable than a system with less attack surface. The rationale is simple – every entry point represents a potential avenue of attack, either in terms of vulnerability, incorrect configuration, or both. As such, the more entry points an attacker has to work with, the greater the likelihood that one of them can be successfully exploited (and lead to system compromise).

Fortunately, attack surface can be reduced through appropriate system hardening. However, hardening does have its practical limits. Some services must inevitably remain operational for a Resource to perform its intended function – a web server must respond to requests for HTTP content, for example, and a DNS server must respond to name service queries. Other services might need to remain for purposes of monitoring or management, such as SNMP or SSH.

5.2. Microsegmentation

5.2.1. Definition

Even with proper hardening, the overall attack surface of many environments continues to increase, in many cases simply due to the sheer number of Resources and other nodes that are internetworked. The inevitable result is:

- More humans interacting with machines.
- More machines interacting with machines.

In terms of ZTA, one way to address this challenge is through the adoption of **segmented** network architectures. A network is said to be segmented if it is physically or logically separated into smaller, isolated subnetworks. The goals of this approach are twofold:

1. Separate public-facing components from private (e.g., non-public-facing) components.
2. Separate unrelated components from each other.

Taken to its logical end state, this approach results in **microsegmentation**, or the subdivision of networks down to the level of discrete, application-specific workloads. In a microsegmented environment, each distinct application receives its own dedicated, logically independent subnetwork.

Many technologies exist to facilitate microsegmentation. Virtual Local Area Networks (VLANs) are a classic example, and variations of this approach (e.g., Private VLANs) exist as well. More contemporary approaches include:

- Software Defined Networks, which use a centralized component to make forwarding decisions for individual packets. OpenFlow is a classic implementation of SDN.
- Overlay networks (e.g., VxLAN), which are networks layered on top of other networks.
- Hypervisors, which can enforce network isolation between guest virtual machines.
- Host-based firewalls, in some cases complemented with agent-based implementations to facilitate firewall rule management.

It is important to consider the various different approaches to microsegmentation when building and deploying new applications and application components.

5.2.2. Example

A good motivation for microsegmentation comes from the shipbuilding industry. The watertight body of a maritime ship is referred to as the hull. Modern shipbuilding techniques separate the hull into "compartments," or individual watertight subdivisions. The rationale for this design is that damage to any portion of the hull can (hopefully) be limited to a specific compartment, instead of flooding the entire hull and sinking the ship.

An example of compartmentalization is depicted in Figure 6. Here, one shipboard compartment has sustained damage and subsequently floods with water. Adjacent compartments, however, are sufficiently isolated from the damaged compartment, and hence do not flood. In effect, the severity and extent of the incident has been minimized.

Figure 6 – Shipboard Compartmentalization

This design can be extended into the worlds of computer networks and distributed systems. For example, consider Figure 7:

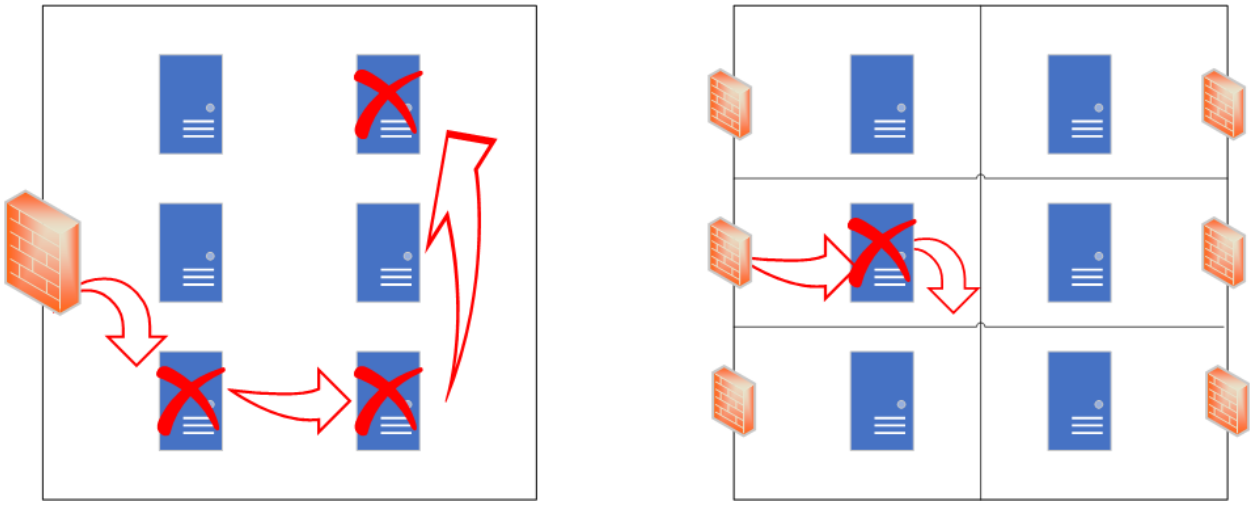

Figure 7 – An Unsegmented Network and a Microsegmented Network

Here we have an example of an unsegmented network (left). In typical "Castle and Moat" architecture, it has a substantial exterior perimeter but no interior perimeters. Should an adversary make it past the perimeter firewall and compromise a Resource, they can leverage that access as a foothold to further move laterally within the network. In the depicted example, an attacker manages to compromise one system (lower left), then uses that system to launch additional attack against other reachable systems in the same network (lower right, upper right).

Compare this to the microsegmented network depicted on the right, where each system occupies its own individual subnetwork. Each of these "micro" subnetworks is isolated from one another, and each features its own individual perimeter. Breaching the perimeter of one subnetwork (and the systems contained within) does not gain an attacker access to other, unrelated subnetworks. In this manner, and as with the example of the shipboard compartment, breaches and incidents can be isolated to the fewest possible number of systems. Not only does this architecture reduce the "blast radius" of an incident, but it also gives organizations more time to detect the incident and mount a response, before too much additional damage is done.

5.3. Application Access

Before delving into our ZTA access model for application Resources, it would be helpful to review the traditional access model that has accompanied "Castle and Moat" perimeter architectures.

In the traditional model, organizations deploy applications to their internal, trusted enterprise networks. Users access these applications from the safety of those internal, trusted networks, shielded from external abuse and attack by firewalls and other perimeter controls.

Now consider the use case where the Users are remote – other organizational locations, business partner locations, or perhaps the Users' residences – and hence in untrusted network locations. In such cases, technologies like VPNs are often employed to provide secure connectivity over public networks, like the Internet. In the traditional model, VPNs provide paths to trusted networks over untrusted networks.

However, ZTA asserts that computer networks should be regarded as untrusted. Viewed this way, VPNs provide paths to untrusted networks over *other untrusted* networks, which is nonsensical. ZTA essentially puts all users into a perpetual remote access security posture, but traditional remote access mechanisms (like VPN) are not philosophically aligned with ZTA.

This situation gives rise to the Identity-Aware Proxy/Access Proxy models discussed in Section 2.4. Comcast decided to pursue deployment of Access Proxy infrastructure as the third domain of its broader Zero Trust program.

5.3.1. Architecture

With ZTA, the Access Proxy becomes the primary model by which Resources are exposed to Users and Devices.

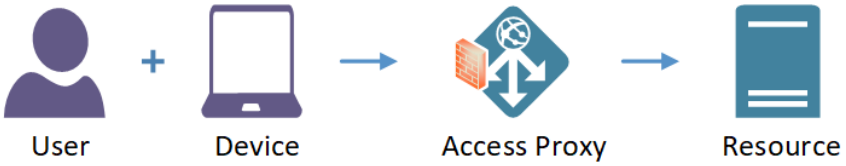

Figure 8 – Access Proxy Arbiting Access to Resources

We decided to decompose the Access Proxy into three functional components:

- Policy Enforcement Point (PEP)
- Policy Decision Point (PDP)
- Enterprise Risk Engine (ERE)

This was done primarily so that we could:

- Develop each component independently, on its own release schedule.
- Scale each component horizontally within a given datacenter environment.
- Geographically load balance the components across our enterprise footprint.
- Facilitate an active/active configuration for high availability and redundancy.

An exploded view of these components and their integration can be depicted as follows:

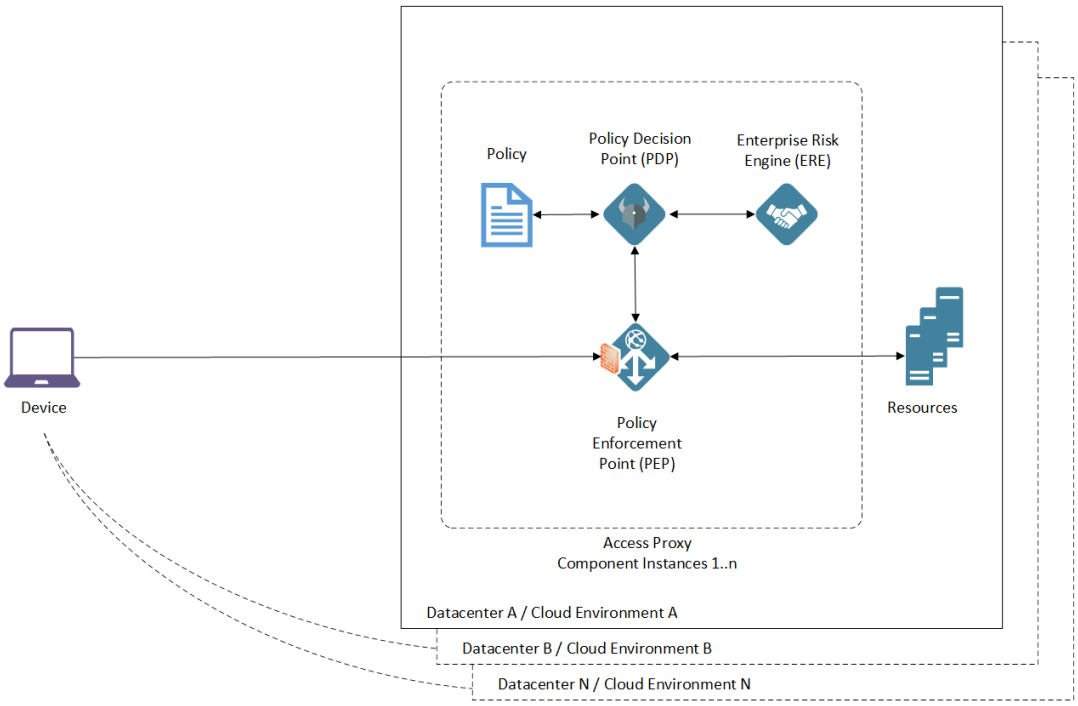

Figure 9 – Exploded View of Access Proxy Components

5.3.2. Policy Enforcement Point

The PEP is essentially an identity-aware reverse proxy. It is responsible for directly mediating connectivity between security principals (Users and Devices) and enterprise Resources. The PEP serves as the ZTA boundary between Resources and their microsegmented networks (which are managed by Resource owners) and all other networks.

5.3.3. Policy Decision Point

The PDP is primarily responsible for authorizing access to enterprise Resources. The PDP uses a combination of static policy (in the form of predefined security rules) and context (in the form of dynamic, near-real-time risk data and other signaling) when making access control decisions. Access control decisions are then enforced by PEPs as necessary.

5.3.4. Enterprise Risk Engine

The ERE is the primary aggregator and provider of context, primarily in the form of risk signals. A risk signal is any kind of data or metadata that provides contextual insight into the security posture of a User, a Device, or a Resource the User/Device is attempting to access. Examples of risk signals include geolocation, velocity, IP address reputation, and Device compliance.

The ERE continuously receives risk signals from sources of truth, including MDM, EDR and IP Reputation. It can then evaluate and share the risk levels of Users and Devices across the security ecosystem.

At of this writing, the PDP is the primary consumer of ERE services, which in turn considers risk signals provided by the ERE when making access control decisions. The ERE is essentially an adjunct used by the PDP to make context-aware access control decisions.

6. Program Governance

We believe that success is tied to how you engage with your enterprise partners. Any effort as large as a ZTA journey requires significant governance and project management. To this end:

- Our ZTA program was championed by our Executive Vice President, who is also our Chief Information Security and Product Privacy Officer.
- Product, Program and Architecture leaders were nominated to guide the overall effort.
- A team was defined to support the overall ZTA program. Each hygiene Pillar was assigned an Executive Leader, a Product Lead and a Program Lead. In most cases, this team was comprised by cybersecurity professionals; in other cases, it included partners from other departments within our organization (IT, Reliability Engineering).
- An Executive Committee was established to provide guidance on a quarterly basis that included leaders from our organization, our business units, Human Resources, Procurement, Legal, Compliance and Finance. We work for a very financially disciplined company. The meetings are a force and function to regularly remind us of the goals that we set, our progress, what's next and whether we need to shift course.
- A Steering Committee was established to provide stakeholder direction between the various organizations contributing development, engineering, and support resources to the program. This committee meets twice a quarter.

We engaged our Executive and Steering Committee members directly very early on in the program to ensure we were properly framing our strategy. They in turn connected us with the critical influencers in our organization who were very supportive and conveyed their commitment to our ZTA journey.

As ZTA started to gain traction, we expanded the formal means of communication to include:

- Cybersecurity Stakeholders – This includes engagement with our Business Information Security Officers (BISOs), Portfolio Leads and Cybersecurity Leads across the company.
- Employee Resource Groups – We closely partnered with technology resource groups to promote, inform and engage partners in our ZTA program:
    - Cybersecurity Guild – Our primary partner is our Cybersecurity Guild with hundreds of members. The Guild hosts a ZTA-themed event at least every quarter that is open to all of our Comcast technical resource groups.
    - We partner with other technical resource groups where they can provide necessary subject matter expertise. For example, we presented our architectural strategy as early concepts to Comcast's Software Ambassador's Guild, and we gained invaluable feedback that influenced our onboarding approach and commitment to automation wherever possible.
- Our ZTA scope is included in annual and quarterly roadmap planning across our technology and business units.

- We are very intentional about how we approach wider employee communications. We have a very tech-savvy workforce and recognized that ZTA would have an impact on how we work. We design our employee experiences and workforce communications around the north star of making it easy for our employees to do the right thing to protect our organization. We partner with our Communications Team to produce campaigns that make the concepts relatable. Cybersecurity Awareness Month is always observed in October, during which time we socialize important security trends. Last year, we introduced ZTA in a game called "Cyber Splash," a popular smartphone application that is available to all our employees. Our "Trust Titan Superhero" made his big debut in October 2021.

As our ZTA program gained further traction, our BISOs soon became critical partners. The BISOs facilitate a streamlined engagement model for delivering security services to our business units. The BISOs worked with the leaders across Comcast to manage ongoing priority setting for ZTA programs and to continually assess feedback on how those programs were going. The feedback wasn't always easy to hear. There are times when we had to address technical issues, process issues, and team issues. In addition, our technical and business partners also felt the strain of the convergence of several compliance programs. We view feedback as a gift and used the opportunity to improve our approach and for corporate leadership to stay better aligned.

7. Lessons Learned

7.1. Engagement

It is crucially important to get key stakeholders involved early. An organization's IT function is the most obvious candidate for collaboration, and our security organization spent much time meeting with IT and aligning efforts. Other key partners included our Office of the CIO, Network Engineering, Procurement, Finance, Human Resources, and Corporate Communications. The governance framework described above ensured ongoing executive engagement and alignment with our key stakeholders.

7.2. Cloud First

Applying elements of ZTA to existing systems and infrastructure – particularly microsegmentation – can be a challenge. Early experimentation and prototypes confirmed this reality. As such, we initiated our ZTA journey with a priority on our new cloud deployments. We focused on enabling microsegmentation and other controls through automated policy. Over the course of one year, we trained our development teams in a "cabin crew" structure that allowed for imbedded expertise within each development organization. Later, we expanded the focus to our legacy applications that require a more hands on, consultative approach.

7.1. Buy vs. Build

In one sense, ZTA is like any other information technology initiative, in that solutions can be realized in many ways, each with its advantages and disadvantages:

- Purchase commercial solutions (licenses, subscriptions, etc.)
- Leverage open-source solutions
- Develop in-house, proprietary solutions

This is often referred to as the "Buy vs. Build" scenario.

More often than anticipated, we had to make buy vs. build decisions throughout our program. These buy/build decisions were driven by our requirements and speed. We had to say "no" to some prototype development efforts in the beginning of our journey, and that's hard to do when you have talented architecture and engineering resources at your disposal. Like any strategic initiative, ZTA needed to deliver value for the company. In some cases, we built our own ZTA components and SSS; in others, we purchased commercial solutions and services. In the latter case, vendor partnerships and strategic integrations have helped us get to where we need to be.

In our experience, the key to successful ZTA is exploring and finding the optimal mix of all three approaches (commercial off-the-shelf offerings, open-source, custom development) and spending the time and effort to integrate these solutions.

7.2. Application Owner Collaboration

For Application Access, we developed an onboarding process to identify strategic applications and register them for access via our Access Proxy infrastructure. Our initial approach entailed identifying the application owners for such applications. We soon discovered that our CMDBs used for identifying and cataloging application and asset information can become outdated over time. To accommodate this, we employed an intensive "white glove" engagement strategy to help identify application owners. This impeded our initial progress, but we learned and adapted the application ownership step as part of the onboarding process if necessary.

The next step in the onboarding process was to identify application authentication mechanisms and determine if they met the minimum thresholds established by our IAM program. In our case, that minimum threshold was based on the OpenID Connect standard (built on top of the OAuth 2.0 protocol). This proved to be a challenge for some application owners, for two reasons:

1. The longer an application has been in existence, the higher the probability that its organizational ownership has changed, possibly multiple times; and
2. When these transitions occur, knowledge transfer for the application's business logic often overshadows IAM knowledge transfer, due to time constraints and other factors.

The result is that the basic application security knowledge is sometimes lost. In such cases, the ZTA team works collaboratively with application owners to understand their authentication frameworks.

And finally, it goes without saying that implementing a comprehensive ZTA often requires application owners to undertake multiple security workstreams, sometimes in parallel. For example, an application team may be asked to simultaneously:

1. Modernize their authentication mechanisms
2. Migrate to a microsegmented network architecture
3. Onboard their application for Access Proxy integration

This can be a tall order for even the most well-resourced application owners. Multiply this single example by the total number of candidate applications in the environment, and one can see that this can seem like an overwhelming set of asks for an organization's application community.

It is imperative that the ZTA team convey and communicate to application owners "security is a team sport" and that everyone is in the ZTA journey together. Defining ZTA priorities at a program level will help build trust between the security and application teams and prevent the perception that ZTA is a series of individual asks, but rather a set of unified strategies to improve the overall security posture of the enterprise.

7.3. Business Partner Collaboration

We spent considerable time defining the quantitative goals for ZTA and the impacts across our internal business units. We had big goals (e.g., "harden X thousand Resources by Q4") and we had to define the intermediate milestones to get to those goals. The challenge was that we didn't always know exactly how we were going to get there. We persistently took stock of what we needed to address and defined risk-based priorities (e.g., "harden Internet-exposed Resources first"). Our business partners appreciated that we put so much thought into the strategy.

Still, this required a high degree of collaboration between Cybersecurity and our business and technical partners. Successful ZTA requires partnership across an organization, not competition. It was important to us that the collaboration meant that the outcomes were better because of teams working together. There were a lot of synergies across the Pillars and meeting security needs also addressed other needs for our business.

At the core of a team is solving those unsolved problems and making the most of our ecosystem. We had to become a provider not just of standards but of solutions for our enterprise. Our business units expect us to look out for them and to anticipate the needs to secure the business.

8. Future Work

In this section, we briefly examine opportunities for enhancement and improvement to our ZTA.

All estimates of this nature are subject to change due to the ever-dynamic needs of our business, as well as unforeseeable events and trends in the information security landscape. As the entire Internet community learned with the Log4j vulnerability of 2021, even the best plans can be upended by a single ill-timed Zero-Day.

8.1. The Multi-Year Journey

In Section 2.5, we noted that for all but the smallest organizations, ZTA is a multi-year endeavor. Overall, Comcast is making good progress on its ZTA journey.

We expect much (if not all) work in the Security Hygiene domain to be complete by the end of 2022. We have also been microsegmenting in cloud/virtual environments for the past three years, and our Access Proxy infrastructure entered production in 2021. These were all big mountains to climb. We are lucky to be a part of an organization where our employees appreciate leaning into a challenge.

Continued expansion of microsegmentation and Access Proxy infrastructure will be primary focus areas of 2023 and beyond. Most net-new deployments already align with our ZTA; it is the uplift of existing and/or legacy Resources that will command most of the effort.

8.2. Centralized and Continuous Risk Evaluation

One of the key tenets of the 2022 OMB memorandum (Section 2.6) is the need for a Continuous Diagnostic and Mitigation (CDM) process. A CDM process ensures that continuous policy evaluation and risk assessment can be conducted on Users and Devices during their active sessions. The memorandum further states that:

> A necessary foundation for any enterprise-wide zero trust architecture is a complete understanding of the devices, users, and systems interacting within an organization.

The generally accepted practice in introducing and incorporating an enterprise wide CDM system into a company's security ecosystem is to:

1. Enable risk signal collection for all Users, Devices and Resources
2. Evaluate the signals in a policy driven manner
3. Gradually and incrementally enable and enforce the policies to improve the overall enterprise security posture.

Comcast has developed its ERE component, currently part of the Access Proxy infrastructure. This essentially serves as a CDM engine. Our future plans are to enhance the ERE with more risk signaling from more risk sources, ultimately including Resources and other end-user applications. The latter possibility is especially intriguing, as it would allow for suspicious application-level activity (e.g., Mallory transferring unusually large amounts of money in a finance application) to serve as a source of risk signals.

8.3. Microsegmentation

Our initial forays into microsegmentation have largely centered around cloud deployments (both on-premise cloud environments as well as hosted/public cloud environments). It is generally easier to microsegment in a virtual environment, as individual application workloads can be allocated their own virtual microperimeter environments (virtual switches, virtual networks, virtual firewalls, etc.)

In 2022 and beyond, we seek to retrofit the microsegmentation concept into "bare metal" and other non-virtualized environments. Several strategies have been considered (see Section 5.2 for examples) and our ultimate strategy may yet employ multiple approaches.

8.4. Converged Application Access

As an organization with a significant mobile/remote workforce (made even further mobile/remote due to COVID-19) we deployed our Access Proxy model initially to support remote access use cases, as an alternative to VPN. Over the long term, however, we wish to route all application access through Access Proxies, both on-premises as well as off-premises.

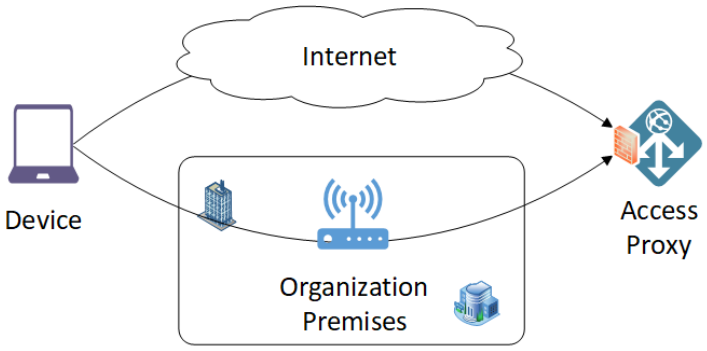

Figure 10 – Converged Access Proxy Architecture

Depicted in this figure are two network paths – one from an organization premises network, the other from an untrusted network such as the Internet. Ultimately, both paths converge on an Access Proxy – one network path is not favored or preferred over another. This is the ultimate end-state architecture for application access in ZTA.

9. Service Provider Considerations

In many ways, ZTA for Multi-System Operators (MSOs) and Internet Service Providers (ISPs) is not fundamentally different than other organizations.

9.1. Infrastructure Hardening

For many years, it has been common practice for operators to harden their access networks to Zero Trust security standards, long before the term "Zero Trust" even entered commonplace use. The Data-Over-Cable Service Interface Specifications (DOCSIS) 3.0 Security Specification, which dates to 2006, states its goals as follows:

1. To provide cable modem (CM) users with data privacy across the cable network;
2. To prevent unauthorized users from gaining access to the network's RF MAC services.

These goals align neatly with ZTA, in that:

1. Assume that the network provides no assurances of privacy. If privacy is desired, it must be established by higher-layer communications protocols.
2. No assumptions are made regarding the identity or authenticity of devices (CMs) on the network. Devices must establish identity through strong authentication before services are provisioned.

DOCSIS then specifies solutions that achieve these goals:

1. For privacy, the Baseline Privacy Interface (BPI) protocol.
2. For authentication, X.509 digital certificates.

Similarly, providers have long hardened the network infrastructure – routers, switches, Cable Modem Termination Systems (CMTSes), etc. – that are used to provision DOCSIS services. Examples include:

- Centralized Authentication, Authorization & Accounting (AAA) frameworks, such as RADIUS, Diameter, TACACS+, LDAP, Kerberos, etc.
- Strong administrative authentication (OTP, public key, etc.)
- Role-based authorization (read only, read/write, CLI command authorization, etc.)
- Encrypted administrative interfaces (HTTPS, SSH, etc.)
- Access Control Lists (ACLs)

In general:

- Service providers should begin regarding their enterprise networks as akin to their access networks, in terms of default trust levels and security guarantees.
- Service providers can and should use their experience hardening access network infrastructure to similarly secure enterprise infrastructure to the same degree.

9.2. Network Visibility

Service providers often deploy significant network monitoring infrastructures, to as to monitor their access networks for performance, faults, and other useful diagnostic information. Simple Network Management Protocol (SNMP) – now in its third version and with significant security enhancements – is the most well-known monitoring protocol.

Existing investments in access network monitoring can and should be deployed to monitor enterprise networks as well, for the reasons discussed in Section 5.1.5.

10. Conclusions

ZTA represents a fundamental shift away from traditional approaches to information and network security. In the past, "internal" networks were safe places, walled off from external threats by "the firewall" and other perimeter controls. Now, however, we are in a state of constant vigilance, where networks do not provide the security assurance they once did.

ZTA recognizes this reality and proposes an alternative approach, one where Users and Devices must establish trust before gaining access to Resources. Trust must be established with confidence (e.g., strong authentication) and network location is not to be used as an authenticator or a credential.

Finally, ZTA is not an experiment – it is a practical, pragmatic approach to security that can be realized with modern IT support structure. If not already on their journey, organizations should strongly consider the merits of ZTA and act accordingly.

Abbreviations

AAA	Authentication, Authorization & Accounting
ACL	Access Control List
BISO	Business Information Security Officer
BPI	Baseline Privacy Interface
CARTA	Continuous Adaptive Risk and Trust Assessment
CDM	Continuous Diagnostic and Mitigation
CISA	Cybersecurity and Infrastructure Security Agency
CM	Cable Modem
CMDB	Configuration Management Database
CMTS	Cable Modem Termination System
DOCSIS	Data-Over-Cable Service Interface Specifications
DNS	Domain Name System
EDR	Endpoint Detection & Response
ERE	Enterprise Risk Engine
FIDO	Fast IDentity Online
IAM	Identity & Access Management
IdP	Identity Provider
ISP	Internet Service Provider
MDM	Mobile Device Management
MFA	Multi-Factor Authentication
MSO	Multi-System Operators
NCSC	National Cyber Security Centre
NIST	National Institute of Standards and Technology
OAuth	Open Authorization
OIDC	OpenID Connect
OMB	Office of Management and Budget
OTP	One-Time Password
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PKI	Public Key Infrastructure
RADIUS	Remote Authentication Dial-In User Service
SAML	Security Assertion Markup Language
SCTE	Society of Cable Telecommunications Engineers
SIEM	Security Information and Event Management
SNMP	Simple Network Management Protocol
SSS	Security Support Structure
TACACS	Terminal Access Controller Access-Control System
UBA	User Behavior Analytics
VPN	Virtual Private Network
ZTA	Zero Trust Architecture

Bibliography & References

"Compartment (ship)." Wikipedia. Available: https://en.wikipedia.org/wiki/Compartment_(ship)

Executive Office of the President, Executive Order on Improving the Nation's Zero Trust Cybersecurity Posture, May 12, 2021. Available: https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/

Cybersecurity and Infrastructure Security Agency (2021) Zero Trust Maturity Model. (Department of Homeland Security, Washington, D. C.), June 2021. Available: https://www.cisa.gov/zero-trust-maturity-model

"Data Breach QuickView Report – 2019 Q3 Trends," RiskBased Security, Richmond, VA, 2019. Available: https://pages.riskbasedsecurity.com/data-breach-quickview-report-2019-q3-trends

Data-Over-Cable Service Interface Specifications (DOCSIS) 3.0 – Security Specification, CM-SP-SEC3.0-C01-171207, Cable Television Laboratories, Inc., 2006. Available: https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=9b3e83b9-daf0-4434-9258-7b6b4e8f2c2e

Executive Office of the President, Office of Management and Budget (2022, January 26). *M-22-09, Moving the U.S. Government Toward Zero Trust Cybersecurity Principles*. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf

L. Cittadini, B. Spear, B. Beyer and M. Saltonstall, "BeyondCorp Part III: The Access Proxy," *;login:,* vol. 41, no. 4, pp. 28-33, 2016. Available: https://research.google/pubs/pub45728/

J. Kindervag, "No More Chewy Centers: Introducing The Zero Trust Model of Information Security," Forrester Research, Cambridge, MA, Sep. 2010.

S. P. Marsh, "Formalising Trust as a Computational Concept," Ph.D. dissertation, Department of Computing Science and Mathematics, University of Stirling, Scotland, UK, 1994.

National Cyber Security Centre (2021) Device Security Guidance. (Government Communications Headquarters, London), June 2021. Available: https://www.ncsc.gov.uk/collection/device-security-guidance/infrastructure/network-architectures

National Institute of Standards and Technology (2020) Zero Trust Architecture. (Department of Commerce, Washington, D.C.), Special Publication 800-207, August 2020. Available: https://doi.org/10.6028/NIST.SP.800-207

S. Potti. "BeyondCorp Enterprise: Introducing a safer era of computing." Google. Available: https://cloud.google.com/blog/products/identity-security/introducing-beyondcorp-enterprise

R. Ward and B. Beyer, "BeyondCorp: A new approach to enterprise security," *;login:,* vol. 39, no. 6, pp. 6-11, Dec. 2014. Available: https://research.google/pubs/pub43231/

"Zero trust security model." Wikipedia. Available: https://en.wikipedia.org/wiki/Zero_trust_security_model

