Cutting through the Acronyms: Finding a Path to Zero Trust

Technologies that can help implement Zero Trust

Ian Poynter
Research Fellow
Nemertes

John Burke
CTO
Nemertes

Q2 2021

Table of Contents

Executive Summary **3**

Zero Trust to the Rescue? **4**

Zero Trust Architecture **4**

Zero Trust Current State and Future Plans **5**

A Sea of Acronyms: CASB, SASE, SCAPE, SDP **6**

CASB (Cloud Access Security Broker) 6

Drilldown on CASB 7

SASE (Secure Access Service Edge) 7

Drilldown on SASE 8

SDP (Software Defined Perimeter) 8

Drilldown on SDP 9

SCAPE (Secure Cloud Access and Policy Enforcement) 9

A Bigger Picture 10

Getting from Here to There **10**

The Modern WAN 10

Thinking About Boundaries 12

Conclusion: Putting it All Together **12**

Zero Trust Checklist 13

Executive Summary

Zero Trust architecture envisages three major changes in the enterprise security environment:

- Shrinking and separating the security boundaries around entities;
- Constantly reconfirming entity identities and authorizations; and
- Constantly monitoring for anomalous behavior.

Although just 45.5% of organizations anticipate having begun to deploy ZT by the end of 2021, every organization should have a goal of implementing a Zero Trust architecture in the next few years. But there's a catch: in exploring ZT options, cybersecurity professionals will run into an alphabet soup of potential solutions or parts of solutions: Cloud Access Security Brokers (CASBs), Secure Access Service Edge solutions (SASEs), Secure Cloud Access and Policy Enforcement environments (SCAPEs), and Software Defined Perimeters (SDPs). These comprise an overlapping set of actual solutions and architectures or frameworks, none of which is a complete ZT solution, but any or all of which can contribute to one.

To select among the various technologies available to move an organization to ZT, cybersecurity professionals must base their choices on:

- The aspects of ZT they need to implement (e.g. cloud-only or hybrid).
- The technologies already in place that can participate.
- Integration with logging, monitoring, and the Security Operations Center (SOC).
- Budget, including the opportunity to consolidate solutions to reduce spending.

The following checklist of points to consider in evaluating ZT solutions is not in priority order; a cybersecurity organization should consider prioritizing it in the context of other projects and goals.

- **Identity management.** If the solution does not provide IDM itself, it must integrate with IDM solutions broadly.
- **Application authorization**. The cybersecurity organization needs a mechanism by which to provide and maintain the detailed and specific application access rights for authorized users.
- **Application use tracking.** Ideally, the solution should provide for more detailed application usage information tracking than network-level monitoring.
- **XDR/user behavior monitoring**. A ZT solution should have a solution for this, or integrate with behavioral threat analytics software more broadly. The system should provide a mechanism to adjust the ZT trust map based on what it sees.
- **Integration with ancillary security solutions**, especially for outbound and endpoint security. For example, secure web gateways (SWGs) and endpoint detection and response (EDR) solutions.

Zero Trust to the Rescue?

What does "zero trust" really mean? As a term, it has a formidable and forbidding sound. NIST provides a basic definition and high level architecture (please see Figure 1, below), but as is typical with standards bodies, not a functional architecture or blueprint for implementation. Fundamentally, ZT is a philosophy, not a technology or product. Adopting the basic tenets drives technology and product choices to implement them. ZT tenets include:

- **Never trust, always verify.** This is the core of the ZT approach. All users and devices are required to authenticate (and re-authenticate) each time they access enterprise systems and assets. There is no implicit trust: No perimeter exists within which all present are assumed to be "good guys." Authorized access to one system *never* implies authorized access to anything else. Authorization for one session does not imply authorization for the next. Authorization is not irrevocable within a session, and bad behavior could get a session terminated while in progress.
- **Secure all communication.** A corollary of no-implicit-trust, this maxim applies to traffic that in older architectures would have been assumed to be inherently secure, for example, on corporate networks in office settings or within enterprise data centers.
- **Assume breach.** All users, devices and systems are assumed to be compromised, until proven otherwise.
- **Log and monitor everything.** When, inevitably, a breach happens, this ensures there is forensic data for the investigation. Monitoring also allows security systems to build a profile of typical usage for all entities in the environment (users and systems), which in turn permits detection of anomalous behavior.

Zero Trust Architecture

Zero Trust architecture envisages three major changes in the enterprise security environment:

- shrinking and separating the security boundaries around entities,
- constantly ensuring the identity of entities (people, software, hardware) and confirming which are allowed to communicate with which, under what circumstances, and
- constantly monitoring for anomalous behavior.

To these ends, ZT architecture revolves around three major components (please see Figure 1, below):

- A policy engine (PE) decides whether a communication session will be allowed, based on a trust algorithm that looks at a trust map based both on policies—which will boil down to "no, it can't" or "it can, if conditions allow"—and on the context that determines whether "conditions allow" (e.g. what else the requesting entity is doing on the network, when and from where access is attempted, and more).
- A policy administrator (PA) manages the implementation of the decision.

- Policy enforcement points (PEPs) enforce the decision under the direction of the PA.

Bear in mind that these components are logical entities and cannot necessarily be bought directly off the shelf, and that solutions may combine functionality, e.g. be both PE and PA. Existing enterprise technology may also fill one or more roles, partially or fully, so rip-and-replace migration may not be required. For example, it is likely not necessary to replace an identity management system that could be integrated with any new technologies required; likewise an existing behavioral threat analytics solution could supply information to the PE regarding the ongoing trustworthiness of an entity.

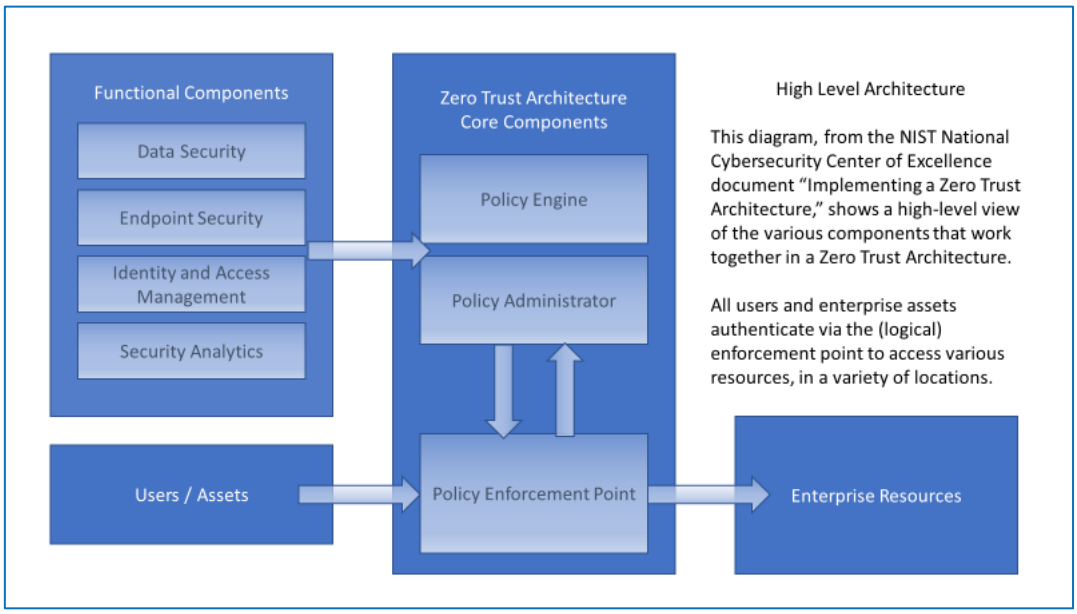

Figure 1: High Level Zero Trust Architecture[1]

Zero Trust Current State and Future Plans

As of May 2021, 16.1% of participants in the Nemertes *Beyond SASE 2021 Research Study* reported they had begun to implement Zero Trust, with 29.4% planning to begin in 2021, another 16% planning for 2022, and 24.5% evaluating.

But what does it mean to "implement" Zero Trust? In a nutshell, it means developing a Zero Trust architecture, then revising existing data, network, endpoint, and firewall architectures to align with the Zero Trust approach. Cybersecurity professionals must also pull together implementation teams, select tools and technologies, and finally review/revise cybersecurity policies. This entire effort typically takes two years or more, so it's clear that we're still in the early stages of most Zero Trust rollouts.

[1] The NIST National Cybersecurity Center of Excellence document "Implementing a Zero Trust Architecture" can be found here https://www.nccoe.nist.gov/sites/default/files/library/project-descriptions/zta-project-description-final.pdf and the ZTA project overview here https://www.nccoe.nist.gov/projects/building-blocks/zero-trust-architecture.

An added complexity when it comes to rolling out Zero Trust is that there's no single product or solution that "implements Zero Trust." Instead, there's a collection of often-overlapping terms and acronyms, each of which can contribute towards implementing Zero Trust. These terms and acronyms include or rely on each other, so we're including a set of definitions in Figure 2:

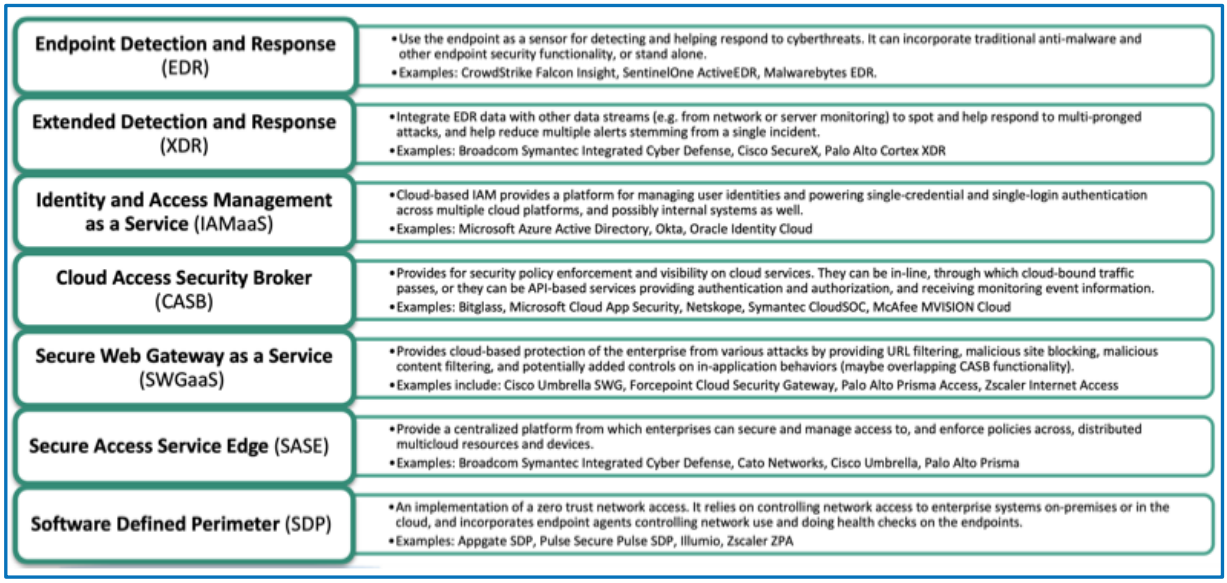

Figure 2: Key Definitions

A Sea of Acronyms: CASB, SASE, SCAPE, SDP

As is so often true in the technology industry, when it comes time to evaluate and select vendor tools to implement ZT, enterprises are faced with a sea of acronyms. When attempting to make sense out of these acronyms, it's important for cybersecurity professionals to take into consideration not just what the acronyms stand for, but where they come from.

CASB (Cloud Access Security Broker)

CASB is a marketing term created by Gartner. According to Gartner, a CASB is an *on-premises or cloud-based software that sits between cloud service users and cloud applications, and monitors all activity and enforces security policies. A CASB can offer a variety of services such as monitoring user activity, warning administrators about potentially hazardous actions, enforcing security policy compliance, and automatically preventing malware.*

The operative word here is "can"; CASB is essentially a grab bag of technologies designed to protect cloud-based apps (typically SaaS ones). There is no defined set of capabilities that must be included for a product or service to qualify as CASB; in fact, in 2015 the then-independent security company Blue Coat purchased two separate CASB companies,

Elastica and Perspecsys, which had virtually no functionality in common. Perspecsys included tokenization technology that protects corporate data moving to and from cloud applications, while Elastica focused on SaaS monitoring and AI-based analytics.

CASB is one of the most commonly deployed cloud solutions; as of May 2021, 51.3% of enterprises in the Nemertes *Beyond SASE 2021 Research Study* had implemented CASB.

Drilldown on CASB

CASBs enforce policy in one or both of two ways. The CASB can be interposed between users and cloud services (i.e. it acts as a proxy), or it can control and monitor access via application programming interface (API) integration to the cloud provider.

CASBs also often act as a funnel from internal enterprise networks to the cloud, with the goal of detecting unsanctioned uses of cloud services—"shadow IT"—including unsanctioned versions of sanctioned solutions. With remote or mobile users, the CASB can act as a focal point for access to approved resources. CASB solutions incorporate or interface with identity management systems to provide a level of single sign-on (SSO).

In a ZT environment, a CASB can inform the PE as to user behavior and user access rights in the cloud systems, and serve as a PEP by allowing the PA to modify those access rights.

CASB vendors include Netskope, Bitglass, Fortinet, and Microsoft (MCAS).

SASE (Secure Access Service Edge)

Like CASB, SASE is a marketing term developed by Gartner that encompasses a broad set of technologies including elements of edge computing, security, and wide-area networking (WAN).

The core, conceptually, is the marriage of security and SD-WAN functionality in a distributed cloud infrastructure. In an ideal SASE, users connect via the nearest point of presence (POP) for the solution, then ride a private and optimized middle mile network to the POP nearest the service they are accessing, or an egress to the larger Internet. Identity-centric security policies can be enforced at any point (ingress, transit, egress), and traditionally appliance-based services like firewalling and secured web access, are provided as part of the cloud service.

As with CASB, the precise capabilities delivered under the rubric of SASE vary significantly from one vendor to another: some incorporate secure web gateway (SWG) capabilities or CASB functionalities, for instance, while others don't. Moreover, many of the vendors jumping on the SASE bandwagon come from a network security background, and have acquired most of their cloud-based capabilities, from multiple sources. This means that many SASE solutions, although sold as integrated, are often not truly integrated under the covers, an issue that potential buyers must be aware of—and wary of, as it can lead to management headaches and misconfiguration problems. Other vendors, even if cloud

native, acquired capabilities outside their original core function—CASB, say, or SWGaaS—and can suffer from similar lack of management integration early on.

SASE is deployed by 37.3% of participants in the Nemertes *Beyond SASE 2021 Research Study*. They do not necessarily use all aspects of their SASE solution, though. More specifically, 45.2% of those using a SASE solution actually use its Secure Web Gateway functionality; 61.3% use its CASB functionality; and 41.3% use its Secure Access Edge capability. Again, this assumes that individual SASE implementations actually have these capabilities, which for the above reasons, is not a given.

Drilldown on SASE

SASE is a network-centric solution implementing identity-centric security. In a ZT architecture, a SASE solution, depending on what functions it *actually* performs from among the broad portfolio of possible features, can act as PE and/or PA and/or PEP for communications among entities not in the same physical location, but (in contrast to an SDP) does not extend into communications within sites.

Both security and network equipment vendors sell SASE solutions, as do cloud-native security and network as a service providers, so it is important to be aware of each vendor's core expertise, the extent to which their solution has grown through mergers and acquisitions as opposed to in-house expansion of functionality, and the extent to which it has truly unified management. The difference is important, because *ab initio* integration of features and management almost always translates into fewer failures, lower operational cost, and shorter learning curves.

Example SASE vendors include Cato Networks, Cisco, and Palo Alto Networks.

SDP (Software Defined Perimeter)

Unlike the previous two terms (three if you include Zero Trust) SDP is not a marketing buzzword but an actual architecture crafted by the Cloud Security Alliance. (In fairness, Zero Trust has an architecture defined ex post facto by NIST, but the original concept lacked a standardized architecture for several years.)

Because SDP is an architecture rather than a marketing buzzword, it is by definition more limited in scope than the other terms. That is, capabilities that aren't part of the SDP architecture aren't, technically speaking, part of SDP. By contrast, while some capabilities are typical to a SASE (such as SWG or CASB) and others are rare (like DLPaaS), few are specifically required by its definition, and nearly anything can and will be marketed and promoted by some vendor or other as "part of our SASE solution."

SDP focuses on the challenge of providing seamless access to enterprise resources regardless of where the resources or the users needing such access reside. A user who should not have access to an application will not even be able to transmit packets to that application's hosting infrastructure (be it server, VM, container, or cloud service).

SDP has a number of specific use cases including VPN replacement (in which it competes with SASE), enabling secure DevOps collaboration in the cloud, and migrating resources and applications to multicloud (since SDP implements a policy regardless of the location of the workload).

Of course, SDP also exists as a marketing term, and as such is seeing some stretch as well; for example, the introduction of "clientless" SDP. These solutions deliver functionality beyond what is encompassed by SDP proper, such as browser-based access to protected resources from unprotected devices, in order to provide a zero-trust access path for users on endpoints the enterprise does not control (and can't put an SDP client on).

Drilldown on SDP

An SDP is a flexible software-defined overlay network implementing Zero Trust. SDP is also sometimes known as Zero Trust Network Access (ZTNA).

An SDP enforces identity-centric access control at the network level. By default, an SDP throws out all network traffic that is not specifically authorized. Whenever any entity tries to communicate with any other, the SDP allows those entities to set up an encrypted tunnel through which to have the given conversation *only* if there is a policy in place specifically allowing for that communication. Permission depends not just on the authenticated identities, but also on *what specifically they are trying to do* (e.g. establish user portal access to a system vs. establish a system administration control panel connection), and where and when and from what platform, etc.

With SDP, the trust boundaries of networks shrink down to individual systems and users, and are made movable without the need to reconfigure hardware or deploy new networks. In essence, an SDP establishes a private, one-to-one microVPN for every conversation. Notably, unlike SASE, SDP can reach into the core of data center or cloud environments, rather than terminating at the boundary.

SDP is an implementation of Zero Trust, with one major caveat: the trust map in a ZT environment is by definition dynamic, adjustable on the fly in response to user behavior. Dynamic trust maps are not definitional to SDP, though they are in no way disallowed, and some SDP vendors do in fact implement them.

As noted earlier, SDP is an architecture rather than a marketing buzzword, so it's more narrowly focused on the specific challenge of granting users granular access to resources. Therefore, it's less likely to find SDP implementations that include everything but the kitchen sink (e.g. SWG, SD-WAN, EDR/XDR, and other capabilities). SDP vendors will typically integrate with point providers of such solutions, rather than offer them directly.

Example vendors include Appgate SDP, BeyondCorp, Certes, and Perimeter 81.

SCAPE (Secure Cloud Access and Policy Enforcement)

SCAPE is a marketing term coined by Nemertes (because why not?) to encompass the broader set of capabilities that are functionally required to deliver an end-to-end security

solution, including Endpoint Detection and Response (EDR), Extended Detection and Response (XDR) functionality, and SDP (Software Defined Perimeter).

Our intent in coining the term SCAPE is not to add to the acronym confusion, or market yet *another* buzzword to vendors, but rather to highlight the fact that no single solution or architecture is a cure-all. Enterprise security professionals we work with find SCAPE a useful concept when assessing the completeness of their cloud security architectures.

A Bigger Picture

As noted, each of these acronyms is defined to a greater or lesser extent (depending on whether it's a marketing buzzword or an actual architecture, or both). Their capabilities often overlap (see Figure 3). As a result, it's important to delve into the specific solutions that are delivered and evaluate how they relate to implementation goals.

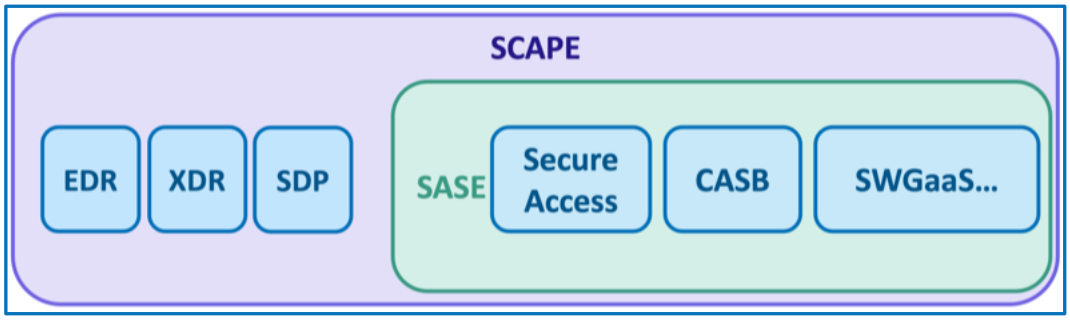

Figure 3: Relationship of Key Acronyms

Getting from Here to There

Having a clear vision of the target allows the comparison of vendor solutions to the goal, which means taking a close look at the security environment these solutions need to support and enable.

The Modern WAN

Most architects think of the corporate wide-area network as connecting "inside-to-inside": that is, premises-based users (e.g. workers in offices) to premises-based resources (e.g. applications running in data centers).

That's not only inaccurate in these post-Covid days, it's been inaccurate for a long time. Even in early 2020, when Nemertes published its *Next Generation Networks Research Study 2020-21*, we found that just 38% of all WAN traffic was "inside-to-inside." The remaining 62% was either "outside-to-inside" (e.g. WFH workers connecting to premises-based resources), "inside-to-outside" (e.g. on-premise workers connecting to the cloud), or "outside-to-outside" (e.g. WFH workers connecting to the cloud).

We have every reason to believe that the percentage of inside-to-inside traffic has declined

precipitously in 2020 and 2021, and will bounce back only slightly as some employees shift back to working in offices.

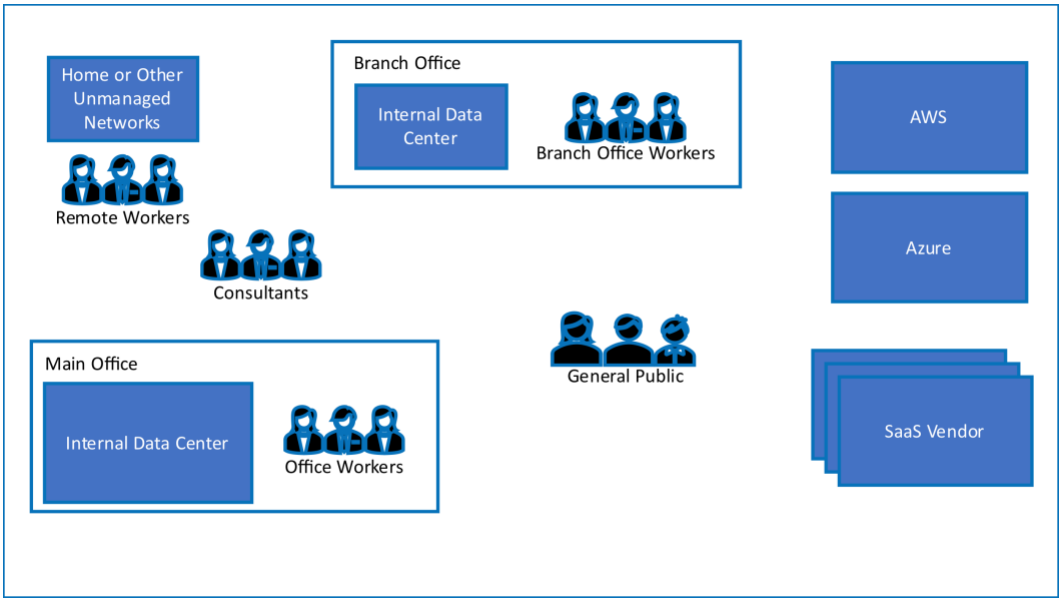

Figure 4: The (New) Corporate Network

This diagram shows the porous boundaries that constitute the new corporate WAN. While users may still exist within the perceived boundaries of the main or branch offices, the reality is that they are also remote, contractors or even sometimes members of the public. This also applies to the enterprise assets and equipment. While internal data centers are still found in many organizations, SaaS and other cloud-based solutions are much more commonplace, even for core services, such as email or CRM.

People are just as likely to use unmanaged home or coffee shop networks while accessing enterprise resources as networks under the control of corporate IT. This is precisely why ZT is so appealing. Rather than taking time and energy to either attempt to secure the networks or prohibit their use, ZT focuses on securing the access to protected resources.

The Covid-19 pandemic has accelerated many remote work requirements, even for enterprises at which it was not under consideration at all. With the changes in technology-driven organizations already in place, it will be exceedingly difficult to put the remote work toothpaste back in the tube.

In this context, the definition of the WAN changes. It is no longer the network tying together company facilities. The new corporate WAN is instead the set of technologies the enterprise uses to provide and control communications among entities *not* in the same physical location.

Thinking About Boundaries

In this version of the diagram, the squares overlaid on the diagram envision the narrowly defined ZT boundaries that the enterprise seeks to implement. These boundaries exist around users and their associated access devices, such as laptops, tablets, and smartphones. They also exist around enterprise resources individually, wherever located, not just around a data center or cloud environment, but around smaller services, components, and even pools of data. With the appropriate tools, we can specify detailed policies for access based on least privilege and need-to-know.

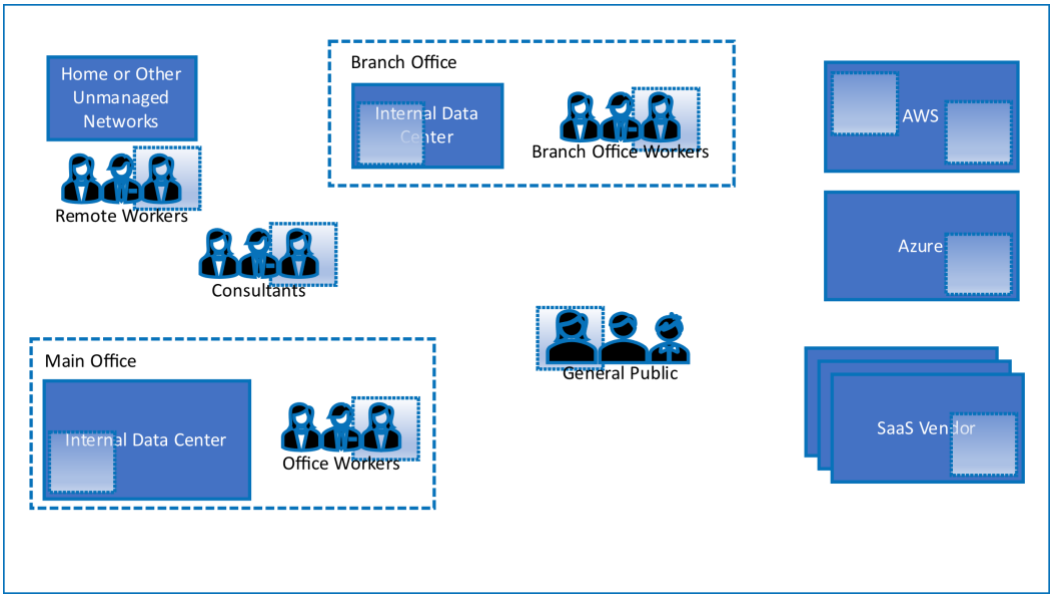

Figure 5: Overlaying ZT Boundaries

Any solutions that will provide ZT must have easily configured, easily audited, easily understood interfaces to allow for simple, error-free provisioning and management. Data classification plays a role in the definition of these policies and associated access rights, and should be considered as part of the provisioning process.

Conclusion: Putting it All Together

Every organization should have a goal of implementing a Zero Trust architecture in the next few years.

In pursuing zero trust, an enterprise has to shift how it thinks, what it expects, and what it does to secure systems and data. It has to let go of the idea that there is a perimeter defining the transition from a safe inside to a dangerous outside; everywhere is dangerous, nothing is to be implicitly trusted. It has to embrace the idea that access rights are temporary, conditional, and specific. It has to expect to make a (probably bumpy) transition to this new paradigm, as it sorts out who really needs to be able to communicate with whom, and under what circumstances. And it has to put in place systems and services that can support and implement ZT.

To select which of the various technologies to consider, cybersecurity professionals must base their choices on:

- The aspects of ZT they need to implement. If there are no data center workloads, then a hybrid product may be overkill. Conversely, if there are and will continue to be in-house systems, support for a hybrid models is essential.
- The technologies already in place. For example, if a currently deployed single sign-on (SSO) solution integrates easily with other ZT components from multiple vendors, it might obviate the need for a migration of that functionality.
- Logging and monitoring. This is a key aspect of any ZT implementation, and it is critical that all newly acquired technologies integrate with existing Security Incident and Event Management (SIEM), Security Orchestration, Automation, and Response (SOAR) and Security Operations Center (SOC) tools or services.
- Budget. This is obvious, but always worth considering before any new tools are deployed. Clearly, this also ties in with leveraging existing tools and vendors, as discussed above.
- Integration. As noted previously, not all "integrated" solutions are actually truly integrated. If a SASE solution comprises 6 different products from 5 different acquisitions, it's unlikely that the integration is solid.

Zero Trust Checklist

This checklist of points to consider in evaluating ZT solutions is not in priority order; an IT organization should consider prioritizing it in the context of other projects and goals.

- *Identity management.* If the solution does not provide IDM itself, it must integrate with IDM solutions broadly.
- *Application authorization.* IT, application owners, and cybersecurity professionals need a mechanism by which to provide and maintain the detailed and specific application access rights for authorized users for
  - on-premise resources
  - cloud resources.
- *Application use tracking.* Ideally, the solution should provide for more detailed application usage information tracking than network-level monitoring, for both
  - on-premise resources
  - cloud resources.
- *XDR/user behavior monitoring.* A ZT solution should have a solution for this, or integrate with behavioral threat analytics software more broadly. (Products in this space are sometimes also referred to as User and Entity Behavioral Analytics, or UEBA.) The system should provide a mechanism to adjust the ZT trust map based on what it sees.
- *Integration with ancillary security solutions*, especially for outbound and endpoint security. For example, a ZT architecture should include/provide, or integrate at policy and logging levels, with
  - Secure web gateway (SWG). These provide protection from malware and ensure enforcement of corporate web access policies. Some SASE solutions include an SWG.

- DLP. Controlling, monitoring and recording the transfer of enterprise data is an essential part of securing it. Some SASE solutions include DLP.
- Next Generation Firewalls (NGFW). These combine traditional firewall functionality with deep packet inspection, application firewalls, intrusion prevention and detection systems (IDS/IPS) and other capabilities. A complete and seamless ZT environment can in theory get by without one, since no unauthorized traffic will get to any system, but they are essential during the transition and should be integrated with the ZT solution if possible, for ease of policy management. In a ZT environment, they can act as "pre-filters" to reduce the amount of traffic the PAs and PEPs have to deal with.
- EDR and endpoint protections (EPP). These solutions expand the ability of endpoints to contribute usage data for use by the ZT environment's policy engine, and act as PEPs.

In sum, regardless of which acronyms you hang your hat on (SASE, CASB, SDP, SCAPE, etc.) an organization's Zero Trust security architecture should include strong endpoint protection and monitoring (EPP+EDR), behavioral threat analytics (XDR), application-level policies with network-level enforcement (SDP), secure web gateway, and identity management. Any solution that provides solid integration of these key functions is worth assessing.