paloalto NETWORKS®

the network security company™

Getting Started With a Zero Trust Approach to Network Security

Executive Summary

The continued, high frequency of successful cyberattacks against today's enterprises has made it abundantly clear that traditional, perimeter-centric security strategies are no longer effective. The failure of resulting architectures is a product not only of the outdated assumption that everything on the inside of an organization's network can be trusted, but also the inability of legacy countermeasures to provide adequate visibility, control, and protection of application traffic transiting associated network boundaries.

First introduced by Forrester Research, Zero Trust is an alternative security model that addresses the shortcomings of failing perimeter-centric strategies by removing the assumption of trust from the equation. With Zero Trust, essential security capabilities are deployed in a way that provides policy enforcement and protection for all users, devices, applications, data resources, and the communications traffic between them, regardless of location.

This paper discusses the need for and details of a Zero Trust approach to network security. It also itemizes the essential criteria and capabilities required of a Zero Trust solution, explains how the Palo Alto Networks® next-generation security platform delivers on these requirements, and provides guidance on how to progressively migrate to a Zero Trust design.

Benefits available to organizations that employ Palo Alto Networks solutions to implement a Zero Trust network include:

- Dramatically improved effectiveness in mitigating data loss via visibility and safe enablement of applications, and detection and prevention of advanced threats;
- greater efficiency for achieving compliance with security and privacy mandates;
- increased ability to securely enable transformative IT initiatives—such as user mobility and infrastructure virtualization; and,
- substantially reduced total cost of ownership (TCO) for IT security.

The Traditional Approach to Network Security is Failing

According to the 2014 Cyberthreat Defense Report, more than 60 percent of organizations fell victim to one or more successful cyberattacks in 2013.[1] Given the extent to which today's organizations continue to rely on perimeter-centric strategies, this finding should come as no surprise. The simple truth of the matter is that perimeter-based approaches to security are no longer effective.

Misplaced Trust

The primary issue with a perimeter-centric security strategy where countermeasures are deployed at a handful of well-defined ingress/egress points to the network is that it relies on the assumption that everything on the internal network can be trusted. However, this assumption is no longer a safe one to make given modern business conditions and computing environments where:

- Remote employees, mobile users, and cloud computing solutions blur the distinction between "internal" and "external;"
- wireless technologies, the proliferation of partner connections, and the need to support guest users introduce countless additional pathways into the network;
- branch offices may be located in untrusted "countries of interest;" and,
- insiders, whether intentionally malicious or just careless, may present a very real security threat.

Such strategies also fail to account for:

- The potential for sophisticated cyberthreats to penetrate perimeter defenses—in which case they would then have free rein over the internal network;
- scenarios where malicious users are able to gain access to the internal network and sensitive resources by using the stolen credentials of trusted users; and,
- the reality that internal networks are rarely homogeneous but instead include pockets of users and resources with inherently different levels of trust/sensitivity which should ideally be separated in any event (e.g., R&D and financial systems versus print/file servers).

Inadequate Capabilities

It is important to realize that a broken trust model is not the only item responsible for the diminishing effectiveness of perimeter-centric approaches to network security. Another contributing factor is that legacy devices and technologies commonly used to build network perimeters let too much unwanted traffic through. Typical shortcomings in this regard include the inability to:

- Definitively distinguish good applications from bad ones (which leads to overly permissive access control settings);
- adequately account for encrypted application traffic;
- accurately identify and control users (regardless of where they're located or what devices they're using); and,
- filter allowed traffic not only for known application-borne threats, but also unknown ones.

The net result is that merely re-architecting one's defenses in a way that delivers pervasive internal trust boundaries will not be sufficient. Care must be taken to also ensure that the devices and technologies used to implement these boundaries actually provide the visibility, control, and threat inspection capabilities needed to securely enable essential business applications while still thwarting modern malware, targeted attacks, and the unauthorized exfiltration of sensitive business data.

The Zero Trust Model—Providing Effective Security for Modern Networks

A promising alternative model for IT security, Zero Trust is intended to remedy the deficiencies with perimeter-centric strategies and the legacy devices and technologies used to implement them. It does this by promoting "never trust, always verify" as its guiding principle. This differs substantially from conventional security models which operate on the basis of "trust but verify."

In particular, with Zero Trust there is no default trust for any entity—including users, devices, applications, and packets—regardless of what it is and its location on or relative to the corporate network. In addition, verifying that authorized entities are always doing only what they're allowed to do is no longer optional; it's now mandatory.

The implications for these two changes are, respectively:

a) The need to establish trust boundaries that effectively compartmentalize different segments of the internal computing environment. The general idea is to move security functionality closer to the different pockets of resources that require protection. This way it can always be enforced regardless of the point of origin of associated communications traffic.

b) The need for trust boundaries to do more than just initial authorization and access control enforcement. To "always verify" also requires ongoing monitoring and inspection of associated communications traffic for subversive activities (i.e., threats).

The core Zero Trust principle and derivative implications are further reflected and refined in the three concepts that define the operational objectives of a Zero Trust implementation.[2]

Concept #1: Ensure that all resources are accessed securely regardless of location. This suggests not only the need for multiple trust boundaries but also increased use of secure access for communication to/from resources, even when sessions are confined to the "internal" network. It also means ensuring that only devices with the right status and settings (e.g, ones that are managed by corporate IT, have an approved VPN client and proper passcodes, and are not running malware) are allowed access to the network.

Concept #2: Adopt a least privilege strategy and strictly enforce access control. The goal in this case is to absolutely minimize allowed access to resources as a means to reduce the pathways available for malware and attackers to gain unauthorized access—and subsequently to spread laterally and/or exfiltrate sensitive data.

Concept #3: Inspect and log all traffic. This reiterates the need to "always verify" while also making it clear that adequate protection requires more than just strict enforcement of access control. Close and continuous attention must also be paid to exactly what is happening in "allowed" applications, and the only way to do this is to inspect the content for threats.

Zero Trust Conceptual Architecture

To help understand what Zero Trust looks like in practice, a conceptual architecture is shown in Figure 1.

The main components include the Zero Trust Segmentation Platform, trust zones, and associated management infrastructure.

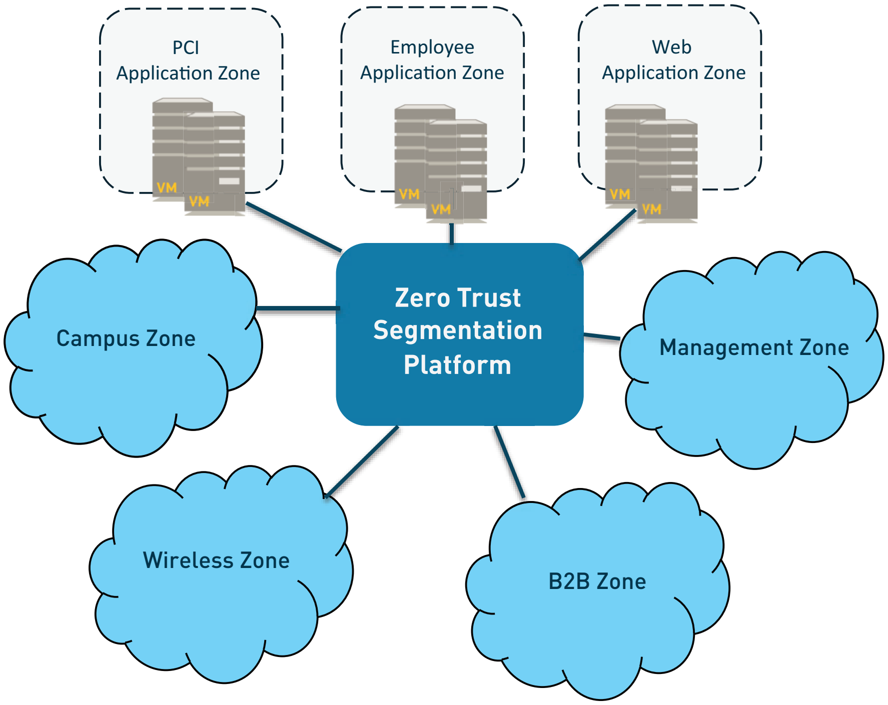

Figure 1 – Zero Trust Conceptual Architecture

Zero Trust Segmentation Platform. Referred to as a network segmentation gateway by Forrester Research[3], the Zero Trust Segmentation Platform is the component used to define internal trust boundaries. In other words, it is what provides the majority of the security functionality needed to deliver on the Zero Trust operational objectives—including the ability to enable secure network access, granularly control traffic flow to/from resources, and continuously monitor allowed sessions for signs of threat activity. Although Figure 1 depicts the Zero Trust Segmentation Platform as a single component in a single physical location, in practice—due to performance, scalability, and physical limitations—an effective implementation is more likely to entail multiple instances distributed throughout an organization's network. In addition, the solution is designated as a "platform" not only to reflect that it is an aggregation of multiple distinct (and potentially distributed) security technologies, but also that they operate as a wholistic threat protection framework to reduce the attack surface and correlate information about threats that are found.

Trust Zones. Referred to as a micro core and perimeter (MCAP) by Forrester Research[4], a trust zone is a distinct pocket of infrastructure where the member resources not only operate at the same trust level but also share similar functionality. Sharing functionality such as protocols and types of transactions is imperative in fact, because this is what is needed to actually minimize the number of allowed pathways into and out of a given zone and, in turn, minimize the potential for malicious insiders and other types of threats to gain unauthorized access to sensitive resources.

Example trust zones shown in Figure 1 include the user (or campus) zone, a wireless zone for guest access, a cardholder data zone, database and application zones for multi-tier services, and a zone for public-facing web applications.

It is important to note, too, that a trust zone is not intended to be a "pocket of trust" where systems (and therefore threats) within the zone are able to communicate freely/directly with each other. For a full zero trust implementation, the network would be configured to ensure that ALL communications traffic—including that between devices in the same zone—is intermediated by the corresponding Zero Trust Segmentation Platform.

Management Infrastructure. Centralized management capabilities are crucial to enabling efficient administration and ongoing monitoring, particularly for implementations involving multiple distributed Zero Trust Segmentation Platforms. In addition, a data acquisition network provides a convenient way to supplement the native monitoring and analysis capabilities for a Zero Trust Segmentation Platform. By forwarding all session logs to a data acquisition network, this data can then be processed by any number of out-of-band analysis tools and technologies intended, for example, to further enhance network visibility, detect unknown threats, or support compliance reporting.

Implementing Zero Trust with Palo Alto Networks

Because the heart of any Zero Trust network security architecture is the Zero Trust Segmentation Platform, it is imperative that organizations choose the right solution. Accordingly, this section identifies a set of key criteria and capabilities for IT security managers and architects to consider when making a selection. In each case, a brief synopsis is also provided of how the Palo Alto Networks next-generation security platform meets the corresponding requirements.

Comprehensive security functionality.

The Palo Alto Networks platform supports:

- **Secure access.** GlobalProtect™ delivers consistent secure IPsec and SSL VPN connectivity for all employees, partners, customers, and guests wherever they're located (e.g., at remote/branch offices, on the local network, or over the Internet). Policies to determine which users and devices can access sensitive applications and data can be defined based on application, user, content, device, and device state.

- **Inspection of ALL traffic.** App-ID™ accurately identifies and classifies all traffic, regardless of ports and protocols, evasive tactics such as port hopping, or encryption. This eliminates methods that malware may use to hide from detection and provides complete context into applications, associated content, and threats.
- **Least privileges access control.** The combination of App-ID, User-ID™, and Content-ID™ deliver a positive control model that allows organizations to control interactions with resources based on an extensive range of business-relevant attributes, including the specific application and individual functions being used, user and group identity, and the specific types or pieces of data being accessed (e.g., credit card or social security numbers) Compared to alternative solutions which let too much traffic through because they're limited to port and protocol level classification, the result is truly granular access control that safely enables the right applications for the right sets of users while automatically eliminating unwanted, unauthorized, and potentially harmful traffic from gaining access to the network. (see "Least Privileges … or Just Least Effective?" below)
- **Advanced threat protection.** A combination of anti-virus/malware, intrusion prevention, and advanced threat prevention technologies (Content-ID and WildFire™), provide comprehensive protection against both known and unknown threats, including threats on mobile devices. In addition, support for a closed-loop, highly integrated defense ensures that inline enforcement devices and other components in the threat protection framework are automatically updated with the findings from WildFire and other sources of threat intelligence.

Coverage for all IT domains. The Palo Alto Networks Zero Trust offering includes an extensive portfolio of virtual and hardware appliances that enables trust boundaries to consistently and cost-effectively be established throughout an organization's entire network, including in remote/branch offices, for mobile users, at the Internet perimeter, in the cloud, at the ingress to the datacenter, and for individual enclaves wherever they might exist.

High-performance design. By definition, a Zero Trust Segmentation Platform aggregates numerous security and networking capabilities. However, it must also be capable of delivering all of these features without becoming a performance bottleneck. The Palo Alto Networks solution achieves this objective first and foremost by utilizing a single-pass software architecture. Processing requirements and latency are minimized as, unlike with other solutions, there is no need for traffic streams to be processed multiple times (e.g., once for each security function). In addition, Palo Alto Networks hardware appliances feature separate control and data planes, plus function-specific, parallel processing hardware engines (i.e., custom chips) for core packet processing, acceleration of standard security functions, and dedicated content scanning. At the high-end, the result is 120 Gbps of Zero Trust throughput, with unmatched visibility and control of applications, users, and content.

Flexible, non-disruptive deployment. Ideally, it should be possible to implement a Zero Trust approach in a way that requires no modification to the existing network and is completely transparent to one's users. Opportunities to take advantage of major network overhauls are rare, and disrupting operations is not a good career choice. Thus, IT security

Least Privileges … or Just Least Effective?

Legacy security gateways and other devices that rely on stateful inspection technology are actually incapable of enforcing a least privileges policy (i.e., where only what's needed to support this business is allowed to pass). The issue with these devices is that their classification engines only understand IP addresses, ports, and protocols—and, therefore, can't distinguish the specific applications that reside behind/within these low-level "wrappers." With a stateful inspection device, for example, a rule permitting traffic using the HTTP protocol on TCP port 80 would allow the passage of not only a legitimate e-commerce application, but potentially numerous other web applications and utilities as well, such as those used for web mail, social networking, and countless other purposes. The net result is that such devices are, in fact, poor candidates for implementing a Zero Trust security model.

managers will need to make due as best they can, typically by converting to Zero Trust on the fly. The Palo Alto Networks next-generation security platform supports this requirement in numerous ways. For example:

- Virtual wire mode enables transparent, layer 1 insertion into the network, and does not require any configuration changes to surrounding or adjacent network devices. All next-generation security technologies are supported in this mode.
- A single hardware appliance can support multiple different connection modes (layer 1, layer 2, or layer 3), thereby maximizing its ability to accommodate trust zones with different needs.
- Support for a broad range of networking technologies (e.g., L2/L3 switching, dynamic routing, 802.1Q VLANS, trunked ports, and traffic shaping) guarantees the ability to integrate into practically any environment.
- Multiple management domains (see Figure 1) can be accommodated by taking advantage of a virtual systems capability that enables separate, isolated Zero Trust virtual instances on a physical appliance. Virtual systems allow you to segment the administration of all policies (security, NAT, QoS, etc.) as well as all reporting and visibility functions.

Centralized management. As discussed, the basis for this requirement is the need to be able to efficiently administer multiple, distributed Zero Trust Segmentation Platforms. In this case, the core need is met by Panorama, Palo Alto Networks optional centralized management offering. However, that is only the beginning of the management capabilities that Palo Alto Networks makes available to simplify the task of implementing and maintaining a Zero Trust security model. Other notable features include:

- A unified policy model and interface that avoids having to flip between multiple screens and/or consoles to view and configure access control and inspection rules
- A hierarchical policy and administration model that accommodates a combination of both global and local rules and configuration settings
- Advanced, graphical visualization tools for better understanding what applications and users are doing on your network at any given point in time
- A comprehensive RESTful API that enables quick-and-easy integration with third-party management, automation, and orchestration tools—for example, to ensure protection for newly-provisioned or relocated virtualized applications
- Integral reporting and logging, with real-time filtering for rapid forensic investigation into every session traversing the network

Further details on these and other capabilities that make the Palo Alto Networks next-generation security platform an ideal solution for implementing a Zero Trust security model can be obtained at www.paloaltonetworks.com/zerotrust.

A Progressive Approach for Implementing Zero Trust

In terms of moving forward with a Zero Trust re-design, it is important for IT security managers and architects to realize that it's not necessary to instigate or wait for the next comprehensive overhaul of their organization's network and security infrastructure. Indeed, one of the great advantages of a Zero Trust architecture featuring the Palo Alto Networks as the Zero Trust Segmentation Platform is that it is conducive to progressive implementation.

To get started, IT security teams can take advantage of the virtual wire feature to non-disruptively deploy Palo Alto Networks devices at one or more locations within the enterprise computing environment. Configured in listen-only mode, these units could then be used to obtain a detailed picture of transaction flows throughout the network, including where, when and to what extent specific users are using specific applications and data resources. Armed with these details, the security team would then be in an excellent position to incrementally (a) deploy devices in appropriate locations to establish internal trust boundaries for identified trust zones, and (b) configure the appropriate enforcement and inspection policies to effectively put each trust boundary "on line." Advantages of a progressive approach such as this include minimizing the potential impact on IT operations and being able to spread the required investment and work effort over time.

For those security teams that already have a good understanding of the transaction flows in their environment, an alternate approach is to map out trust zones and begin to establish corresponding trust boundaries based on relative risk and/or sensitivity of the data involved. A logical starting point in this case is to begin by identifying well-defined pockets of users and systems involving high concentrations of sensitive data—such as the 4Ps: payment card industry (PCI) or other financial data, personal healthcare information (PHI), other personally identifiable information (PII), and intellectual property (IP). From there, it then makes sense to consider progressively establishing trust zones/boundaries for other segments of the computing environment based on their relative degree of risk – for example:

- IT management systems/networks (where administrators often hold the proverbial “keys to the kingdom” and a successful attack could lead to compromise of the entire network)
- Partner resources and connections (B2B)
- High-profile, customer-facing resources and connections (B2C)
- Branch offices in “countries of interest” (followed by all other branch offices)
- Guest access networks (both wireless and wired)
- Campus networks

Adopting Zero Trust principles and concepts at major access points to the Internet also makes sense,. Doing so, however, will probably require replacing or augmenting entrenched, legacy security devices with a Zero Trust Segmentation Platform to obtain all of the requisite capabilities.

Benefits of Adopting Zero Trust Principles and Practices

There are several technical and business advantages associated with using the Palo Alto Networks platform to achieve a Zero Trust security architecture. These include being able to:

- Incrementally and non-disruptively make the transition to a Zero Trust model
- Obtain unparalleled situational awareness of enterprise computing activity, legitimate and otherwise
- Fully implement all Zero Trust principles and concepts, including strict enforcement of a least privileges access control policy (which is essential to reducing attack surface)
- Dramatically enhance the organization’s security posture and ability to prevent the exfiltration of sensitive data
- Simplify achieving and maintaining compliance with applicable standards and regulations (by using highly effective trust boundaries to segment off sensitive resources)
- Securely enable and easily adapt to accommodate business-driven IT initiatives—such as user mobility, social networking, infrastructure virtualization, and cloud computing
- Reduce total cost of ownership (by using a single consolidated security platform across the entire computing environment, instead of a disparate collection of disconnected point products)

Conclusion

Perimeter-centric security strategies continue to be sorely challenged. The issue is not only increasingly sophisticated cyberthreats, but also major changes to the technology and business landscape—such as user mobility, hyper inter-connectivity, and globalization—that invalidate the assumption that everything “on the inside” can be trusted. The bottom line is that such strategies—along with the legacy technologies used to implement them—are, for the most part, no longer effective.

Organizations looking to substantially improve their defensive posture against modern cyberthreats and more reliably prevent exfiltration of sensitive data should consider migrating to a Zero Trust security architecture. An alternative model for IT security, Zero Trust eliminates the faulty assumption of trust and rectifies the shortcomings of traditional perimeter-centric architectures by promoting the use of a

Zero Trust Segmentation Platform to establish secure "trust boundaries" throughout a computing environment and, in general, in closer proximity to sensitive resources.

Because the Zero Trust Segmentation Platform is the foundation of any Zero Trust initiative, the importance of selecting the right solution cannot be over-stated. In this regard, the Palo Alto Networks next-generation security platform represents an ideal candidate—one that combines unparalleled visibility, control, and threat protection capabilities with comprehensive coverage for all IT domains, from the datacenter and Internet gateway to branch offices, mobile users, and even the cloud.

Footnotes:

1. 2014 Cyberthreat Defense Report, CyberEdge Group, February 2014
2. No More Chewy Centers: Introducing the Zero Trust Model Of Information Security, Forrester Research, September 14, 2010
3. Build Security Into Your Network's DNA: The Zero Trust Network Architecture, Forrester Research, November 11, 2010
4. ditto

paloalto NETWORKS®
the network security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

Copyright ©2014, Palo Alto Networks, Inc. All rights reserved. Palo Alto Networks, the Palo Alto Networks Logo, PAN-OS, App-ID and Panorama are trademarks of Palo Alto Networks, Inc. All specifications are subject to change without notice. Palo Alto Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Palo Alto Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. PAN_WP_ZT_022414