Network Segmentation and Zero Trust Architectures

William R. Simpson and Kevin E. Foltz

***Abstract* — Network defense utilizes a comprehensive set of hardware and software tools to preclude malicious entities from conducting nefarious activities. Most current enterprises build their defenses upon a fortress approach. Network defense tools defend this fortress, which defines a clear boundary between the untrusted outside and the trusted inside. Network segmentation expands on the fortress idea to create a layered fortress model, where within a fortress there may be smaller fortresses with their own boundaries and protections. This provides more layers of defense, which limits threat mobility and helps to contain damage during exploits and intrusions. Zero trust starts with a different model, where the individual resources are protected and there is no reliance on the network for protection. This has the same goals of limiting threat mobility and containing damage. While network segmentation shares similar goals with zero trust architecture, it has fundamental incompatibilities that prevent it from being a useful security enhancement within a ZTA. This paper reviews the concepts of network segmentation and ZTA and illustrates why network segmentation is useful only for non-security purposes within a ZTA.**

***Index Terms* — Zero Trust, Network Defense, Network Segmentation, Networking, Security Architectures**

I. INTRODUCTION

Network defenses have traditionally been based upon the fortress approach. Computer Network Defense is defined as

> "Actions taken through the use of computer networks to protect, monitor, analyze, detect, and respond to unauthorized activity within the enterprise information systems and computer networks." [1]

The definition provides a very active defense seeking to find suspicious behavior and provide packet blocking, destruction or mis-direction, blocking of Internet Protocol addresses, and a range of other active measures. The current defense package assumes that the threat can be stopped at the front door. As shown in Figure 1, all traffic in the enterprise, both coming and going, is routed through this front door. However, despite our best efforts to restrict all traffic to this front door, exceptions are inevitably made to introduce multiple undocumented backdoors that compromise security.

Manuscript received 31 Dec 2020; revised 21 Jan 2021. This work was supported by the Institute for Defense Analyses. Such support does not constitute an endorsement by either the Institute for Defense Analyses or the U. S. Department of Defense.

William R. Simpson, corresponding author, Institute for Defense Analyses, 4850 Mark Center Drive, Alexandria, VA 22311 USA

Kevin E. Foltz is with the Institute for Defense Analyses. (email: kfoltz@ida.org)

The elements involved in implementing network and application defenses are numerous and complicated. Functionality is provided by a wide range of appliances [3-11]. This functionality may be for quality of service to the user or quality of protection to network resources and servers. These appliances are often placed in-line. Many operate at line speeds for all communications coming from or going to the enterprise, and some require access to content to provide their services. Figure 2 provides a representation of how these appliances come between the user and the application.

The fortress defense has failed to provide the promised boundary security, with breaches occurring almost daily. The appliances in the package do stop the current threats for a short period, but new threats materialize very shortly and once again defeat the fortress approach. Each time a new technology is put in place to counter discovered exploits, it makes the front door more complicated, more expensive, and more vulnerable. We must assume that threats may be present in the system at any time, and even with detection and mitigation, we must assume continued threat presence over long periods. The fortress approach has no answer for this paradigm, so we must use a different approach.

Alternatives to the fortress approach include network segmentation, distributed computing, end-point defenses, Zero Trust Architecture (ZTA), and Enterprise Level Security (ELS).

Network Segmentation [12] seeks to reduce the number of assets or resources in a segment of the network, separate the network segments, and require subnetwork security enforcement that limits lateral movement. Section II describes this in more detail.

Distributed computing [13, 14] is a model in which a capability of a system is implemented by multiple interacting components. These components are often identical in their hardware and software, and although they are independent nodes, they run as one system to provide the desired function. This is often done to improve reliability, due to the improved tolerance to individual node failures, and performance, due to increased parallelization.

With the improved reliability that distributed computing can provide, an attack on a single component of a distributed system need not compromise the system. It may be possible for the remaining components to isolate or repair the compromised node before the attack can spread.

ISBN: 978-988-14049-2-3
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

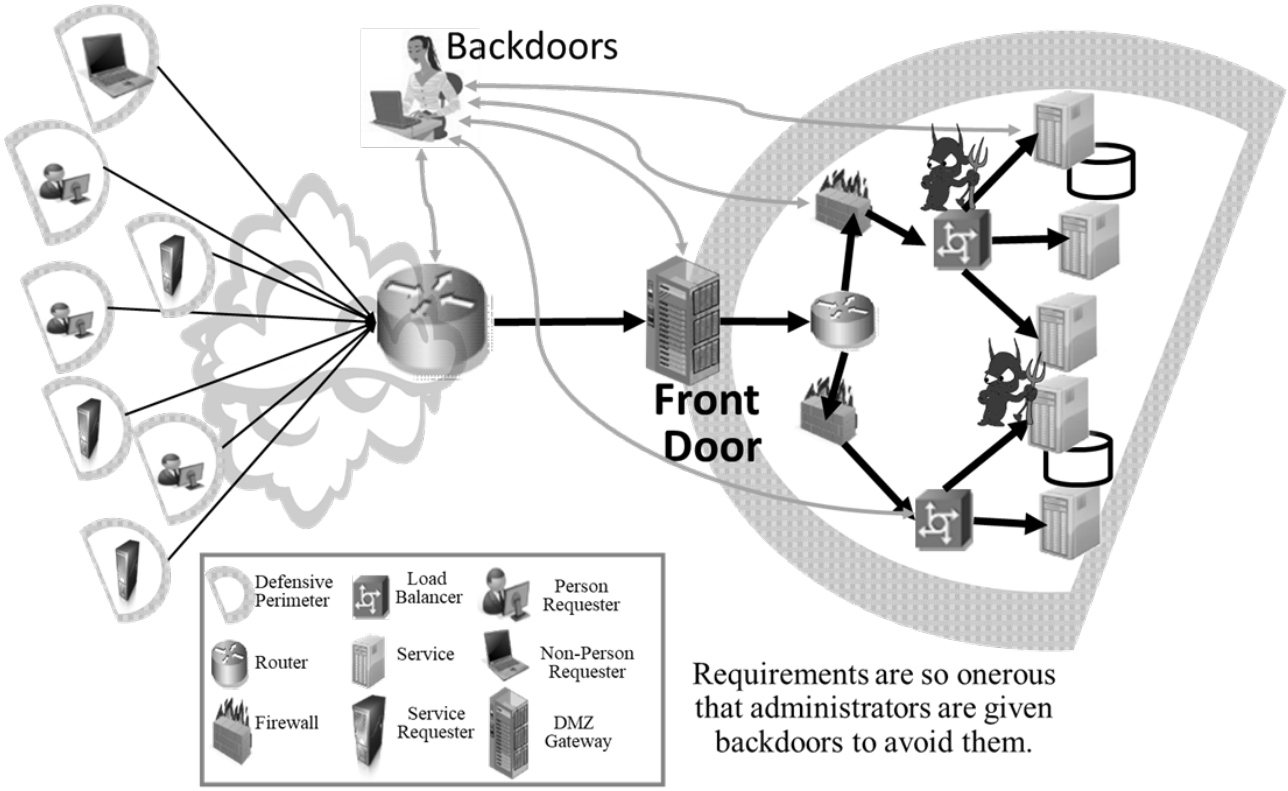

Figure 1. Fortress Protected Enterprises

Most distributed computing implementations provide parameters to adjust the degree of fault tolerance and parallelization, ensuring that different levels of security are possible through tuning of parameters. However, distributed computing alone provides no security. It still requires the addition of security functions that can run in a distributed environment.

End-point defenses [15, 16] define the requester and provider as the endpoints and put defense capabilities on these endpoints instead of on the network connecting them. Often they use endpoint health indicators and requester identity information to provide fine-grain access control to endpoint resources.

ZTA [17, 18] uses the principle of protecting individual resources within the enterprise, such as data and computing, instead of protecting network borders. Requests coming from the internal network are not inherently trusted and must verify their identities and access credentials at each resource. ZTA is designed to prevent data breaches and limit internal lateral movement in the enterprise. ZTA is described in detail in Section III.

ELS [19, 20] is a security architecture developed for the U.S. Air Force to overcome the assumptions inherent in fortress defenses. ELS encompasses many of the methods described above in an overall consistent security architecture.

For purposes of this paper we will concentrate on network segmentation and zero trust. The following sections describe each approach individually and then examine whether network segmentation is a useful approach within an existing ZTA.

II. NETWORK SEGMENTATION

Network segmentation is a term for dividing a network into multiple subnetworks, or segments, and managing access to these segments. Typically, it involves segregating traffic between the network segments and enforcing segment policies with firewalls or other security appliances. A typical segmentation is shown in Figure 3. Segmentation may involve the use of physical sub-neworks or Virtual Local Area Networks (VLANs). VLANs often rely on MAC address or incoming physical port numbers, and they provide, at best, machine-based security. They do not make distinctions based on requester identity credentials or resource access privileges.

The degree of network segmentation is determined by two opposing forces: the separation of resources into different segments, and the grouping together of resources within the same segment. The terms macro segmentation and micro-segmentation qualitatively describe different ends of this spectrum. With extreme macro segmentation, we arrive at the fortress approach for the entire enterprise. With extreme micro-segmentation we arrive at endpoint defense, where each endpoint is treated as its own fortress. Most real world implementations fall between these extremes and involve a number of segments that each contain a number of resources.

Micro-segmentation reduces lateral movement of threats and provides more granular access by allowing different rules at the policy decision points (PDP) for each of the segments. Resources are protected by appliances for both the segment and the whole network.

ISBN: 978-988-14049-2-3
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

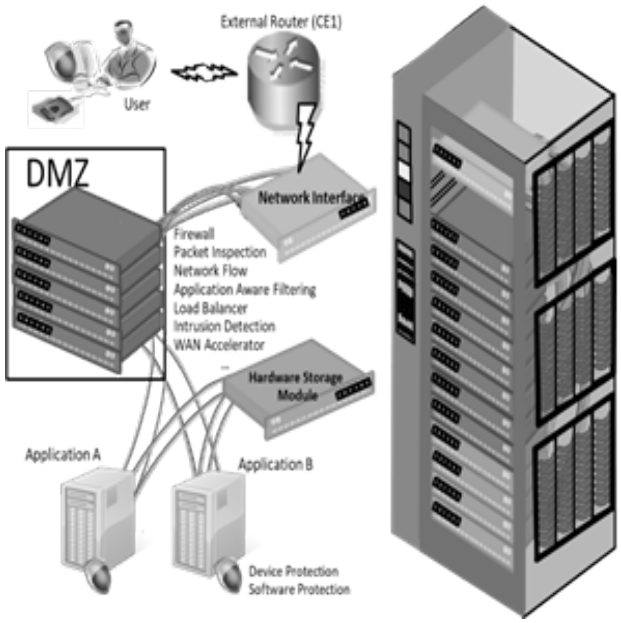

Figure 2. End-Point Access

Network segments do not have registered names and identity credentials. Segment security may be based on the user, the device, or both. Network access is inherently a one-way authentication process. The network authenticates the user, but the user does not (and need not) authenticate the network. The user does not care about the network as long as the requested resource is accessible.

One problem with segment access controls is that the segment policies must be based on the access policies for the resources within them. This is a security challenge because the network policies must be least as permissive as the most permissive policy for any resource within that segment. Otherwise, the network is blocking valid requesters from resources to which they are entitled. With complicated or diverse resource access policies, it may be very difficult to implement a meaningful segmentation security policy. Another challenge is that any change to access policies for resources must be propagated to the segments.

A fundamental problem for segmentation is based on its reliance on the fortress approach for security. It is still based on the flawed assumption that a robust front door can prevent attacks from outside. Adding more layers of a flawed approach leaves many of the same vulnerabilities that were previously present with the fortress approach. Attackers find ways through protections, and they will move laterally within segments, just as they did for the fortress. Segmentation increases the complexity of the fortress approach and must be carefully configured. Any misconfiguration is a new vulnerability.

III. ZERO TRUST ARCHITECTURE

ZTA was designed to address lateral threat movement within the network. ZTA embraces the principle of never trust, always verify. ZTA is a paradigm that moves defenses from network-based perimeters to focus on users, assets, and resources. More information on ZTA is provided by NIST SP 800-207 [21].

Each entity in a communication must have assurance that the party they are engaged with is a known entity and, specifically, the one to whom the communication is intended. Access and privilege should only be granted to an authenticated identity if credentials for access and privilege are presented, verified, and validated. Finally, all communications should be encrypted and provided with integrity protections that allow the recipient of communications to verify that what was received was actually sent. References [22] and [23] provide extensive descriptions of these processes.

Entities may be active or passive. Passive entities include storage elements, routers, wireless access points, some firewalls, and other entities that do not themselves initiate or respond to web service or web application requests. Passive entities do not view, create, or modify application layer content. Active entities are those entities that request or provide application layer services. Active entities include users, applications, and services. All active entities have identity credentials. Communication between active entities requires bi-lateral, end-to-end authentication using verifiable and trusted identity credentials.

A simple distinguishing feature of active versus passive entities is that active entities act as either sources or sinks of content, and passive entities act as pass-through elements. Note that there is no notion of entities being both sources and sinks. Such entities would be proxy or gateway elements, which break end-to-end security and hence are not allowed for ZTA. Moving to ZTA requires an assessment of the benefits versus the risks. Moving from a single boundary defense to multiple resource defenses allows increased flexibility of the defenses provided. Each resource can tailor its defenses to its own needs. However, this increased complexity, if not properly managed, can introduce its own vulnerabilities. Also, misconfiguration of any endpoint defense tool is a new vulnerability.

IV. COMBINING SEGMENTATION AND ZTA

Previous sections looked at segmentation and ZTA in isolation. Now we look at combining them. We first consider a full security implementation of both approaches. We then examine a hybrid solution that mixes parts of each. Finally, we consider non-security benefits.

A. *Full Security Combinatiton*

First, we consider implementing segmentation on an existing ZTA. ZTA requires seamless end-to-end encrypted communication for active entities as shown in Figure 4. Segmentation adds boundary security components that must break end-to-end security in order to view network traffic. These boundary components are passive entities in ZTA. As passive entities, they do not have the ability to decrypt network traffic, and they cannot perform their functions. So, in an existing ZTA, segmentation cannot help without breaking ZTA security.

ISBN: 978-988-14049-2-3
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

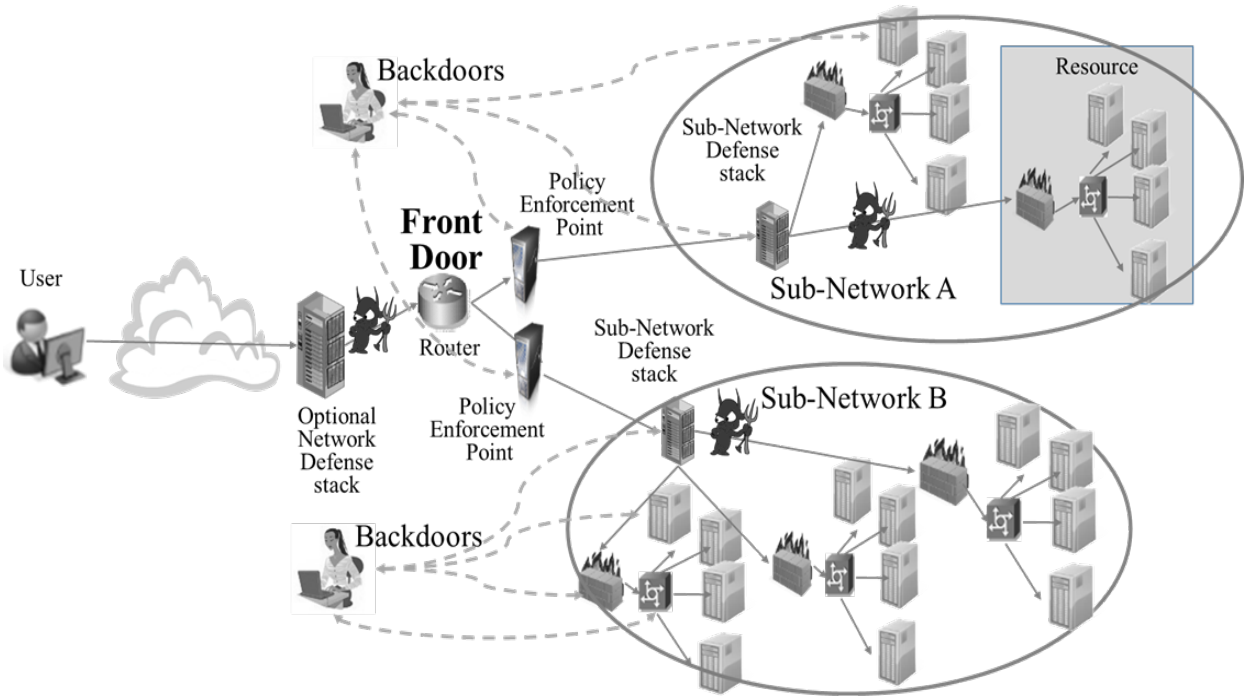

Figure 3. Segmented Network

Next, we look at adding ZTA to an existing segmentation. In this case, segment boundary protections break end-to-end active entity communication security at each segment boundary. This is done in the sub-network defense stacks in order to inspect content. A full ZTA implementation is not possible. However, ZTA is possible within each segment. ZTA delays content inspection until the content is at the server as described in [24], maintaining an unbroken encrypted communication path. Content inpection in the server only operates on that one server's traffic and does not require special hardware to handle full network traffic. Figure 5 shows individual micro-segments with ZTA enforcement applied. Note that segment defensive stacks are moved to the server, and administrator back doors are eliminated. ZTA limits lateral movement within and among each micro-segment, but unlike a full ZTA, it provides no security across macro segments. So ZTA can help in an existing segmentation, but only within individual segments that are fully ZTA.

Combining segmentation and ZTA results in problems from a security perspective. The key issue is how to handle secure communication at segment boundaries. Segmentation requires breaking it and ZTA requires preserving it. Because of this fundamental difference, it is not possible to fully implement both approaches in the same enterprise.

B. *Hybrid Approach*

A full implementation of both approaches does not work, but when segmentation is finely applied such that each segment is a micro-segmentation, conditions essentially match a full implementation of ZTA. Micro-segmentation to the individual resource together with an embedded network defense stack preserves the end-to-end communication path. This association of micro-segmentation and ZTA provides the basis for a hybrid solution.

Areas of micro-segmentation within an overall segmentation that includes both micro- and macro-segmentation can be converted to a local ZTA solution. This conversion of a single segment to ZTA can be applied to all regions of micro-segmentation, and neighboring ZTA segments can be combined into a single larger ZTA segment. Figure 5 also illustrates a hybrid enterprise segmentation using macro- and micro-segmentation. Using ZTA on the micro-segmentation paths and normal defense in depth on the macro segmentation provides the overall hybrid solution. Note that while the backdoors persist in the normal segmentations, the back doors are eliminated in the ZTA architecture, and administrators and other previous exceptions must go through the front door for connection. This is less onerous for administrators as they have an unbroken and direct encrypted connection to the end-point they seek. Coverting additionl parts of the macro segmentation into micro-segmentation results in a migration path from fortress to ZTA using segmentation.

C. *Other Considerations*

Although segmentation and ZTA cannot be fully combined for security, dividing network traffic between different segments may reduce the aggregate network traffic on each segment, which improves performance. Use of VLANs instead of hardware can offer cost savings and improved configurability. Software defined networks can improve network traffic performance. These segmentation benefits do not require breaking encryption at boundaries and show that although segmentation does not help ZTA security, it can provide other benefits.

ISBN: 978-988-14049-2-3
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

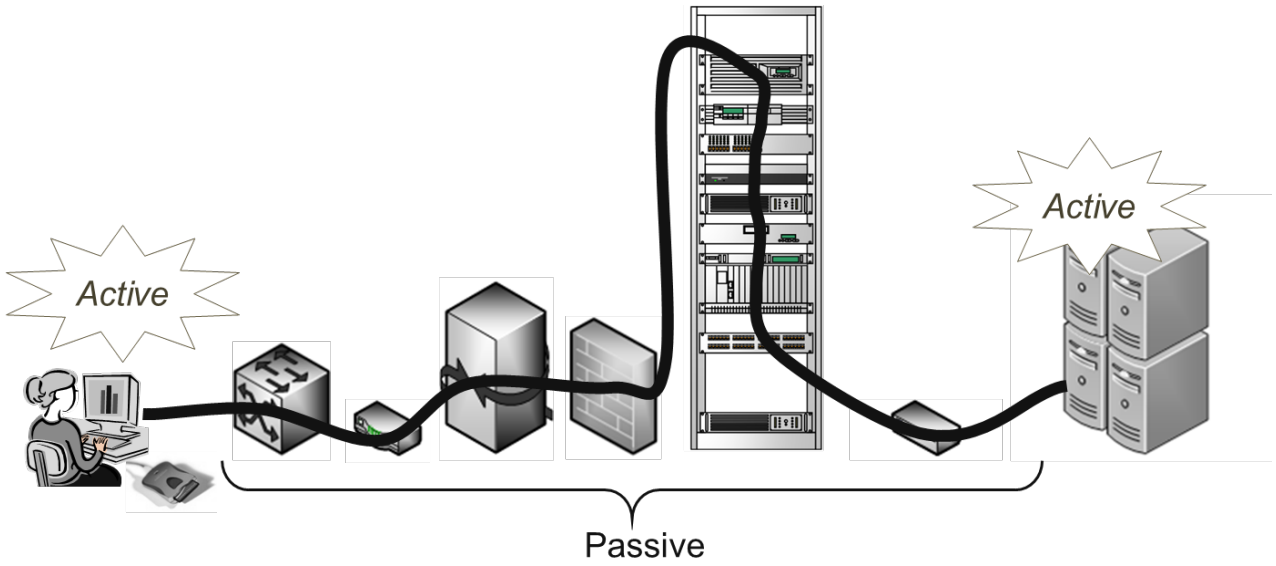

Figure 4. Typical Security Stack for Enterprise

V. CONCLUSION

Segmentation and ZTA are security approaches that improve on the fortress approach. Segmentation divides the network and repeatedly applies the fortress approach to each portion of the overall network, but ZTA explicitly does not trust the network and relies on the endpoints for security. They cannot be directly combined due to fundamental incompatibilities associated with end-to-end security. However, they can be used in a hybrid mode, where different parts of the enterprise use different approaches. ZTA can improve segmentation by providing security within individual segments, but segmentation cannot improve ZTA because it cannot break end-to-end encryption to perform security functions. Thus, improving the security of a hybrid approach naturally leads to more micro-segmentation and ZTA. Thus, the logical end state for security of a hybrid solution is ZTA. Segmentation provides the path to get from the fortress to ZTA, but this end state does not require segmentation for security. Network segmentation and ZTA can be combined in an enterprise, but ZTA does the security, and network segmentation provides other benefits, such as performance, broadcast traffic minimization, cost savings, and other efficiencies.

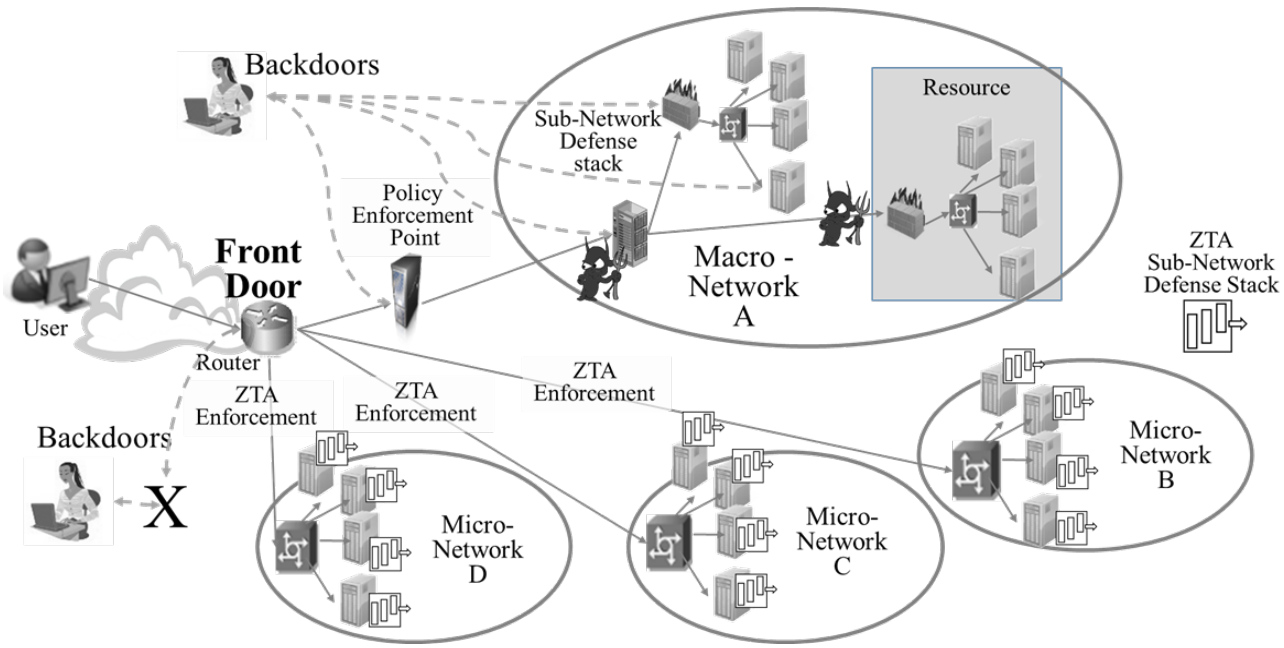

Figure 5. Combine Micro/Macro Segmentation for ZTA Transition

ISBN: 978-988-14049-2-3
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

REFERENCES

[1] Science Direct, Editors: Jason Andress, Steve Winterfeld, *Cyber Warfare*, ISBN 9781597496377, 2011, Jason Andress, Steve Winterfeld, Chapter 10 - Computer Network Defense, pp. 179–191, http://www.sciencedirect.com/science/article/pii/B9781597496377000101, last accessedon 24 November 2020.

[2] TechTarget.com, "backdoor (computing)," https://searchsecurity.techtarget.com/definition/back-door, last accessed 22 November 2019.

[3] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly, National Institute of Standards and Technology, NIST Special Publication 800-207, Zero Trust Architecture, August 2020, https://doi.org/10.6028/NIST.SP.800-207 , last accessed on 24 November 2020.

[4] Jack Wallen, "Five free, dead-easy IP traffic monitoring tools," Tech Republic, September 2011, https://www.techrepublic.com/blog/five-apps/five-free-dead-easy-ip-traffic-monitoring-tools/, last accessed 22 November 2019.

[5] Moskovitch R, Elovici Y. "Unknown malicious code detection – practical issues," In Proceedings of the 7th European Conference on Warfare and Security (ECIW'08), Plymouth, UK, 2008.

[6] A. Begel, S. McCanne and S. L. Graham, "BPF+: Exploiting global data-flow optimization in a generalized packet filter architecture," *Proc. of ACM SIGCOMM*, Cambridge, MA, USA, 1999, pp. 123–134.

[7] M. McDaniel and M.H. Heydari, "Content Based File Type Detection Algorithms," Proceedings of the 36th Annual Hawaii International Conference on System Sciences, IEEE, ISBN**:** 0-7695-1874-5, DOI: 10.1109/HICSS.2003.1174905, Jan 2003.

[8] Mike Fisk and George Varghese, "Fast Content-Based Packet Handling for Intrusion Detection," Los Alamos National Lab Computing Communications and Networking Division, May 2001, https://apps.dtic.mil/dtic/tr/fulltext/u2/a406413.pdf, last accessed 22 November 2019.

[9] Jian Song and Yanchun Zhang,. 2007, "Architecture of a Web Accelerator for Wireless Networks," in Proceedings of the thirtieth Australasian conference on Computer science – Volume 62 (ACSC '07), Gillian Dobbie (Ed.), Vol. 62. Australian Computer Society, Inc., Darlinghurst, Australia, Australia, pp. 125–129.

[10] Shin-ichi Kuribayashi, "Improving Quality of Service and Reducing Power Consumption with WAN Accelerator in Cloud Computing Environments," International Journal of Computer Networks & Communications (IJCNC) Vol.5, No.1, January 2013.

[11] Afzal, S., Kavitha, G. "Load balancing in cloud computing – A hierarchical taxonomical classification." J Cloud Comp 8, 22, December 23, 2019, https://doi.org/10.1186/s13677-019-0146-7

[12] N. Wagner et al., "Towards automated cyber decision support: A case study on network segmentation for security," 2016 IEEE Symposium Series on Computational Intelligence (SSCI), Athens, 2016, pp. 1–10, doi: 10.1109/SSCI.2016.7849908.

[13] Tanenbaum, Andrew S., Steen, Maarten van (2002). *Distributed Systems: Principles and Paradigms*, Upper Saddle River, NJ: Pearson Prentice Hall. ISBN 0-13-088893-1.

[14] Magnoni, L. (2015), "Modern Messaging for Distributed Sytems (sic)," Journal of Physics: Conference Series. 608 (1): 012038. doi:10.1088/1742-6596/608/1/012038. ISSN 1742-6596.

[15] Hassan N.A., (2019), "Endpoint Defense Strategies," in Ransomware Revealed, Apress, Berkeley, CA. https://doi.org/10.1007/978-1-4842-4255-1_4

[16] Mark Khai Shean Tan, Sigi Goode & Alex Richardson (2020), "Understanding negotiated anti-malware interruption effects on user decision quality in endpoint security," Behaviour & Information Technology, DOI: 10.1080/0144929X.2020.1734087

[17] Dayna Eidle, Si Ya Ni, Casimer DeCusatis, Anthony Sager, "Autonomic security for zero trust networks," Ubiquitous Computing Electronics and Mobile Communication Conference (UEMCON) 2017 IEEE 8th Annual, pp. 288–293, 2017.

[18] C. DeCusatis, P. Liengtiraphan, A. Sager and M. Pinelli, "Implementing Zero Trust Cloud Networks with Transport Access Control and First Packet Authentication," 2016 IEEE International Conference on Smart Cloud (SmartCloud), New York, NY, 2016, pp. 5–10, doi: 10.1109/SmartCloud.2016.22.

[19] William R. Simpson, CRC Press, "Enterprise Level Security – Securing Information Systems in an Uncertain World," by Auerbach Publications, ISBN 9781498764452, May 2016, 397 pp.

[20] Kevin E. Foltz, William R. Simpson, CRC Press, *Enterprise Level Security 2: Advanced Techniques for Information Technology in an Uncertain World*, by Taylor & Francis Group, September 2020, 338 pp., ISBN 9781003080787.

[21] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly, National Institute of Standards and Technology (NIST) SP 800-207, Zero Trust Architecture, August 2020.

[22] PKI Standards: PKCS#12 format PKCS #12 v1.0: Personal Information Exchange Syntax Standard, RSA Laboratories, June 1999 http://www.rsa.com/rsalabs/node.asp?id=2138 PKCS 12 Technical Corrigendum 1, RSA laboratories, February 2000.

[23] Organization for the Advancement of Structured Information Standards (OASIS) open set of Standards S. Cantor et al., "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0," OASIS Standard, March 2005

[24] William R. Simpson, Kevin E. Foltz, Proceedings of the 9th International Conference on Software Engineering and Applications (JSE 2020), Ed: David C. Wyld, Natarajan Meghanathan, ISBN: 978-1-925953-28-2, "Network Defense in an End-to-End Paradigm," pp. 177–187, virtual presentation, Zurich, Switzerland, November 21–22, 2020.