NETWORK DEFENSE IN AN END-TO-END PARADIGM

William R. Simpson and Kevin E. Foltz

The Institute for Defense Analyses (IDA), Alexandria, Virginia, USA

ABSTRACT

Network defense implies a comprehensive set of software tools to preclude malicious entities from conducting nefarious activities. For most enterprises at this time, that defense builds upon a clear concept of the fortress approach. Many of the requirements are based on inspection and reporting prior to delivery of the communication to the intended target. These inspections require decryption of packets when encrypted. This decryption implies that the defensive suite has access to the private keys of the servers that are the target of communication. This is in contrast to an end-to-end paradigm where known good entities can communicate directly with each other. In an end-to-end paradigm, maintaining confidentiality through unbroken end-to-end encryption, the private key resides only with the holder-of-key in the communication and on a distributed computation of inspection and reporting. This paper examines a formulation that is pertinent to the Enterprise Level Security (ELS) framework.

KEYWORDS

Appliance, end-to-end security model, ELS, network defenses, web server handlers

1. INTRODUCTION

The Enterprise Level Security (ELS) framework has evolved from a fortress approach, in which the assumption that the threat is stopped at the front door, to a distributed security system that eliminates or mitigates many of the primary vulnerability points inherent with that system, as shown in Figure 1.The basic process of identification involves a two-way contract between two entities that are initiating a communication. Each entity needs to have some assurance that the party they are engaged with is a known entity and, specifically, the one to whom the communication should be allowed. The presentation of claims by each party that can be verified and validated accomplishes this. These claims are often in the form of credentials. [1] provides an extensive description 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. Active entities are those entities that request or provide services according to ELS. Active entities include users, applications, and services. All active entities have PKI certificates, and their private keys are stored in tamper-proof, threat-mitigating storage. Communication between active entities requires bi-lateral, PKI, end-to-end authentication. A verifiable identity claims-based process provides authentication.

David C. Wyld et al. (Eds): NLP, JSE, MLTEC, DMS, NeTIOT, ITCS, SIP, CST, ARIA - 2020
pp. 177-187, 2020. CS & IT - CSCP 2020 DOI: 10.5121/csit.2020.101414

2. LITERATURE REVIEW FOR CURRENT DEFENSE PACKAGES

The elements involved in implementing network and application defense are numerous and complicated. Functionality is provided by a wide ranges of appliances (and by other means).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, and some require access to content to provide their service. Figure 2 provides a representation of how these appliances come between the user and the application.

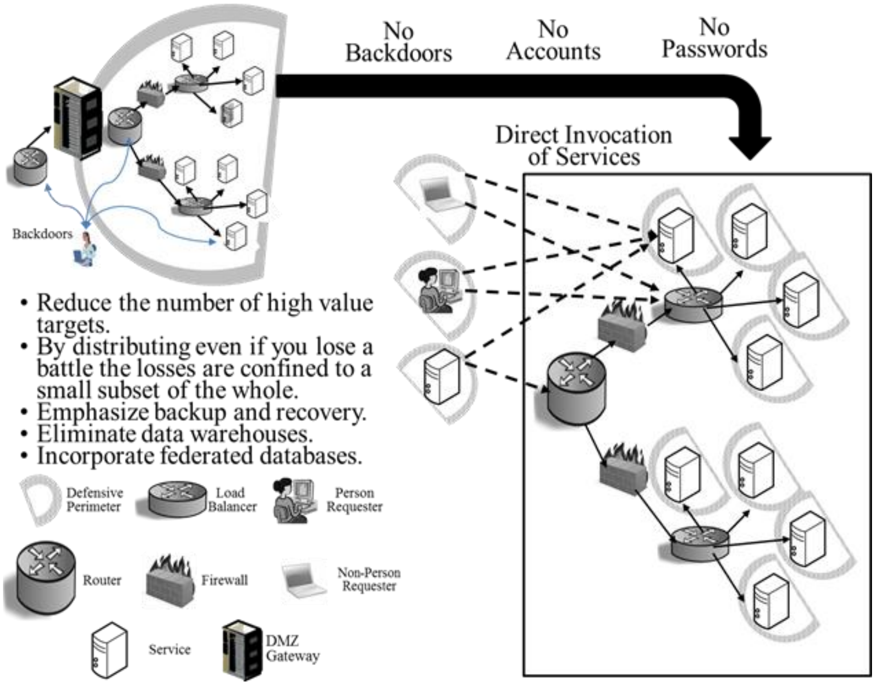
No Backdoors
No Accounts
No Passwords
Direct Invocation of Services
Backdoors
• Reduce the number of high value targets.
• By distributing even if you lose a battle the losses are confined to a small subset of the whole.
• Emphasize backup and recovery.
• Eliminate data warehouses.
• Incorporate federated databases.
Defensive Perimeter
Load Balancer
Person Requester
Router
Firewall
Non-Person Requester
Service
DMZ Gateway

Figure 1 Distributed Security Architecture

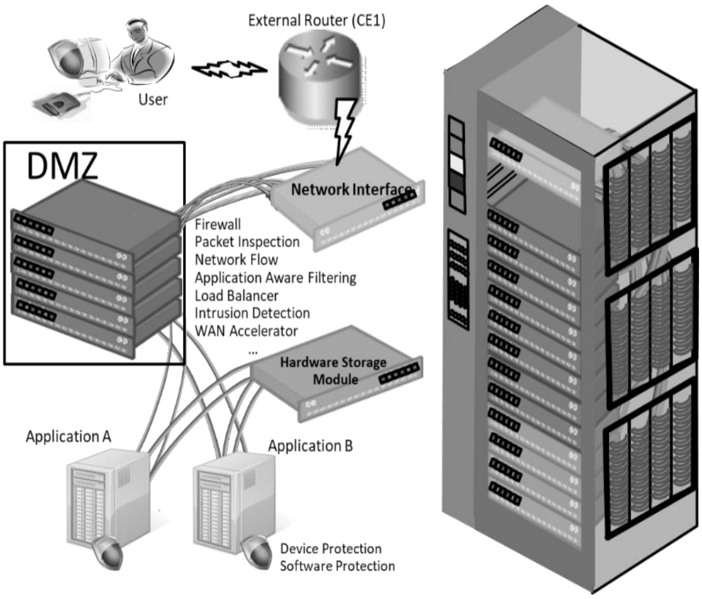

Figure 2 End-Point Access

The number and types of appliances can be quite large. Below is a partial list of functional types as provided in the current literature:

- Header-based scanner/logger [2]
  - Views only unencrypted portion of traffic
  - Synchronous or asynchronous operation
  - Scans for defined behavior, logs traffic
- Content-based scanner/logger [3]
  - Views decrypted transport layer security (TLS) content
  - Synchronous or asynchronous operation
  - Scans for defined behavior, and logs traffic/content
- Header-based firewall [4]
  - Views only unencrypted portion of traffic
  - Synchronous operation
  - Scans for and blocks defined behavior
- Content-based firewall – block only [5]
  - Views decrypted TLS content
  - Synchronous operation
  - Scans for defined behavior and blocks (terminates) connection
- Content-based firewall – modify malicious content [6]
  - Views decrypted TLS content
  - Synchronous operation
  - Scans for defined content, and blocks connection or removes content without blocking the connection
- Web accelerator [7]
  - Views decrypted TLS content
  - Synchronous operation
  - Modifies content for performance
- WAN accelerator [8]
  - Views decrypted TLS content
  - Multi-party system
  - Synchronous operation
  - Modifies content representation between parties, but no end-to-end modification
- Load Balancers [9]
  - Distributes load among destination end-points to improve throughput and reduce latency
  - May decrypt content:
    - May combine encrypted flows through a “secure sockets layer (SSL) accelerator”
    - May distribute content by request to different servers based on load
    - These load balancers are active entities
  - May not decrypt content:
    - Using “sticky” or end-point balances may route all requests from an entity to the same server
    - These load balancers are passive entities

3. SHORTCOMINGS OF THE CURRENT APPROACHES

Each of the appliances above offers some functionality and increases the threat exposure. None of these are free from vulnerabilities from a security standpoint, and they do increase the threat surface and the vulnerability space. For example, default passwords or other improperly secured access methods allow an attacker access to any data that the appliance can access. For detailed scans, this could include all decrypted network traffic to and from a server. With a large number of independent appliances, this represents a significant security risk. Use of any appliance must be balanced by the increased functionality and the increased vulnerability. The situation is further complicated by vendor offerings of load balancers with firewall capability, “smart” accelerators that scan content, and software-only offerings that will provide most of these functionalities in a modular fashion.

This work is part of a larger body of work termed “Consolidate Enterprise IT Baseline (CEITB).”In this paper, we review the communication models for current network defenses. We then review the inspection processes and its basic architecture. Next, we show how network inspections and reporting are available while maintaining end-to-end communications. Finally, we provide the unique factors that arise with end-to-end approaches and network defenses.

4. THE REAL DE-MILITARIZED ZONE (DMZ)

Figure 3 provides a real-world defense package. Although it may look like a network defense package you have seen, it is not and it is only for illustration purposes. The first thing you see is that it is very complex and has many elements requiring proper configuration to function correctly. In reality, it occupies several racks of equipment. Secondly, the first stop after initial entry from the external router is a load balancer that will decrypt the encrypted packets. This is accomplished by either providing the private keys of all servers or allowing the load balancers (LB1 or LB2) to access the hardware storage module (HSM) of the server as if it were the server. Both break the end-to-end paradigm. Additionally, in most instances, forwarded packets are unencrypted as the appliances are assumed trusted. Each appliance has its own set of vulnerabilities, and this complicates the network defense appreciably.

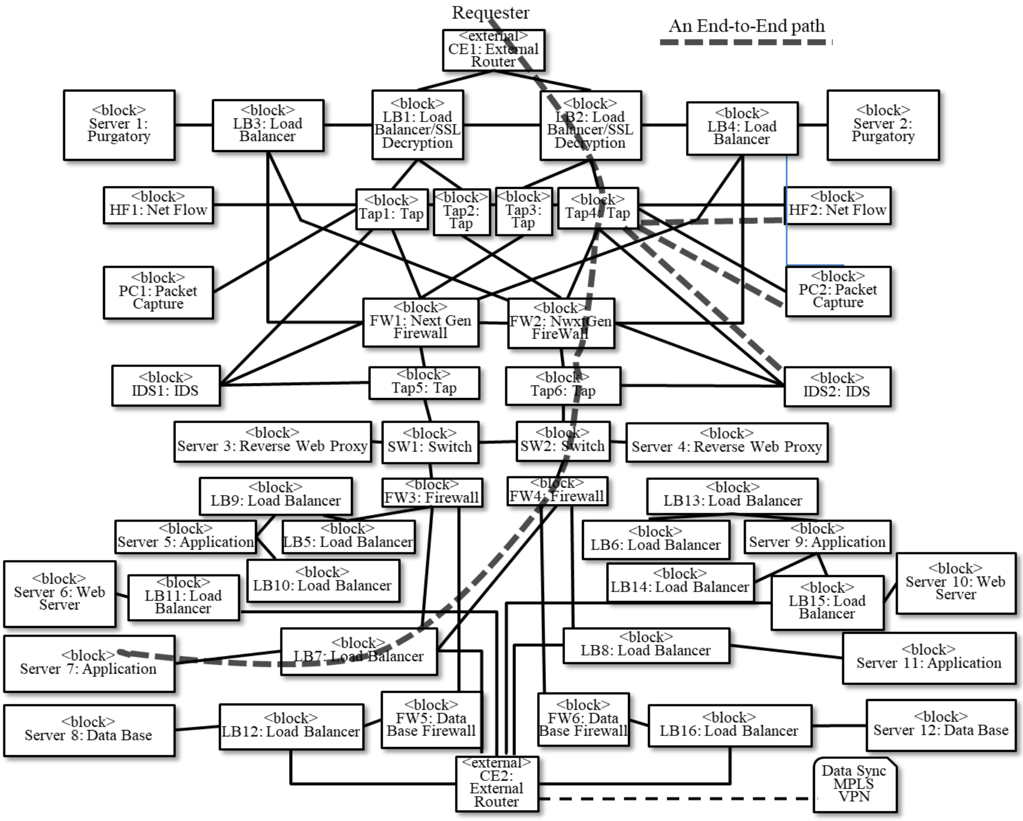
Requester
An End-to-End path
<external> CE1: External Router
<block> Server 1: Purgatory
<block> LB3: Load Balancer
<block> LB1: Load Balancer/SSL Decryption
<block> LB2: Load Balancer/SSL Decryption
<block> LB4: Load Balancer
<block> Server 2: Purgatory
<block> HF1: Net Flow
<block> Tap1: Tap
<block> Tap2: Tap
<block> Tap3: Tap
<block> Tap4: Tap
<block> HF2: Net Flow
<block> PC1: Packet Capture
<block> PC2: Packet Capture
<block> FW1: Next Gen Firewall
<block> FW2: Nwxt Gen FireWall
<block> IDS1: IDS
<block> Tap5: Tap
<block> Tap6: Tap
<block> IDS2: IDS
<block> Server 3: Reverse Web Proxy
<block> SW1: Switch
<block> SW2: Switch
<block> Server 4: Reverse Web Proxy
<block> LB9: Load Balancer
<block> FW3: Firewall
<block> FW4: Firewall
<block> LB13: Load Balancer
<block> Server 5: Application
<block> LB5: Load Balancer
<block> LB6: Load Balancer
<block> Server 9: Application
<block> Server 6: Web Server
<block> LB11: Load Balancer
<block> LB10: Load Balancer
<block> LB14: Load Balancer
<block> LB15: Load Balancer
<block> Server 10: Web Server
<block> Server 7: Application
<block> LB7: Load Balancer
<block> LB8: Load Balancer
<block> Server 11: Application
<block> Server 8: Data Base
<block> LB12: Load Balancer
<block> FW5: Data Base Firewall
<block> FW6: Data Base Firewall
<block> LB16: Load Balancer
<block> Server 12: Data Base
<external> CE2: External Router
Data Sync MPLS VPN

Figure 2 A Real DMZ

5. A NEW APPROACH –CREATING THE PSEUDO-APPLIANCE

The main contribution and unique approach to network defense in a distributed system is in maintaining the inspection process without breaking the end-to-end encryption of communications. The pseudo-appliance captures all of the inspection processes and places them into one software process that resides in the application for processing. This is the first step in realigning the priorities between the current approach and the end-to-end approach, as shown in Figure 4.The path from the user to the application in the top half of the figure shows the processes needed for inspection. Note that the private key for server 7 has been hand passed to the initial load balancer so that the exchange of information is visible. Next, the load balancer decrypts packets for inspection. This includes not only the inspection, but also the necessary reporting.

In the second half of Figure 4, we show the user directly communicating with the load balancer in front of the application (which now contains the inspection process).We have reduced the bandwidth necessary to handle the traffic at the network interface and distributed the computing burden. Tagging the communications between the requester and provider bypasses the DMZ stack. The initial handshake (which is unencrypted) includes the exchange of two white-listed PKI certificates. This exchange in ELS is the bi-lateral authentication of entities and is the initial setup for TLS encryption of all communications. This exchange allows for this tagging. As the decryption now occurs in server 7 prior to inspection, key passing is no longer required, and the end-to-end confidentiality is maintained. Untagged traffic will go through the normal DMZ processing. The reduction in traffic bandwidth at the front door may reduce the need for several of the downstream load balancers. Figure 5 shows the handler makeup in the server.

ELS enhances protection of the application server and provides additional security protections as discussed in the following section.

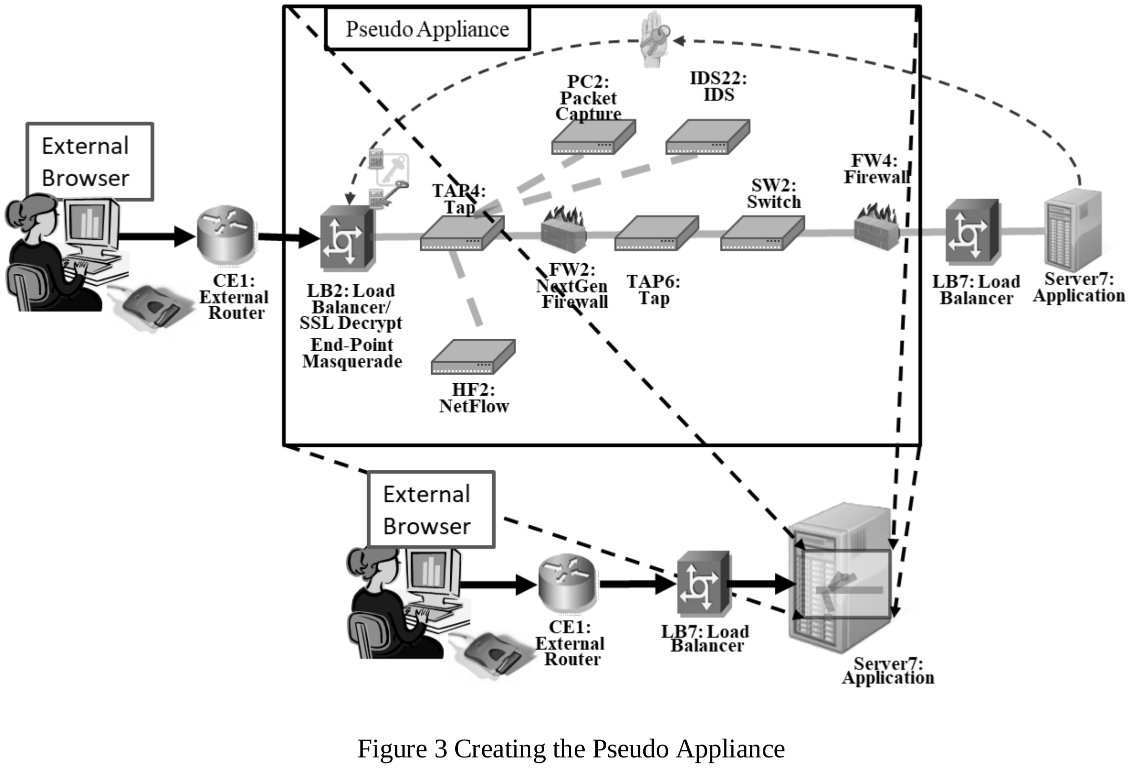

Figure 3 Creating the Pseudo Appliance

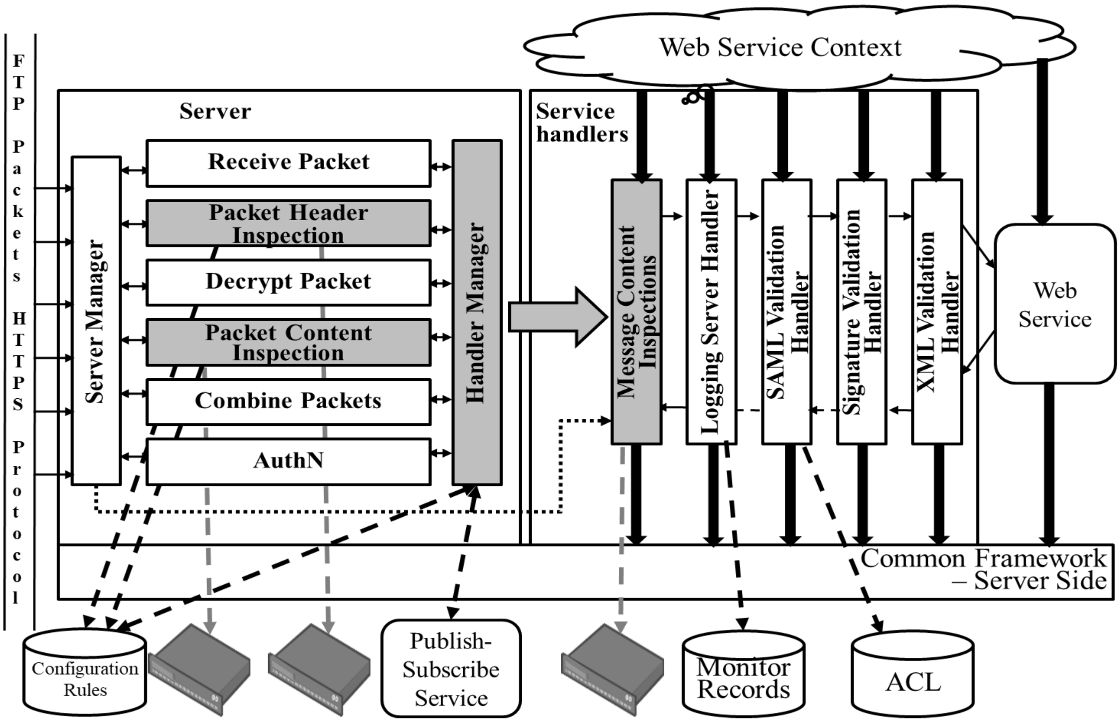

Figure 4 ELS End-point Network Security Functions

6. END-POINT PROTECTION SYSTEMS

The end-point protection system must provide firewall functionality under certain circumstances (as shown in Figure 6) based on end-point, claimed identity, requested action, and other factors.

- Black list – The only functionality enforced is block or drop packets. The black list is centralized, managed, and "pushed" to the protection system (ELS compliant)
- White – Varying degree of firewall enforcement based upon device and criticality. White membership includes The S3ecurity Token Server (STS), for example.
- Gray – Full firewall functionality is enforced. Functionality includes virus scan, malware scan, and other deep packet techniques.

The protection system has the capability to monitor, filter, or shut down traffic to given ports. It scans for malicious code. It examines incoming and outgoing traffic for anomalies or known exploits. It acts in the security context of the end-point for both requester and provider and examines not only the encrypted traffic but also the clear text traffic for malicious behavior or code. This requires access to the unencrypted traffic as well as the encrypted traffic. The protection system provides most but not all of the checks. Figure 6 walks through checks in an ELS enclave provided by the protection system, the server handlers, the service handlers, and the service itself, minimizing the need for in-line appliances.

This capability of the protection system is defined in terms of functional elements, some of which are listed below:

- Maintaining an inventory of assets on all hosts with situational awareness
- Detecting and removing of viruses, Trojans, worms, bots, and root kits in incoming and outgoing email

- Identifying unsafe websites during searches
- Detecting and repairing computer problems
- Enforcing policies on local machine
- Monitoring asset configurations and compare against baseline to detect changes
- Preventing use of unauthorized USB and flash media
- Blocking known and unknown buffer overflow exploits
- Preventing malicious code installation/execution
- Identifying activities that deviate from DoD or organizational policy
- Ensuring firewall functionality
- Monitoring DHCP requests on the network
- Marking any system that does not check in as rogue
- Scanning for compliance with policies
- Identifying host vulnerabilities on the network
- Making data available to the consumer, using ELS security
- Providing situational awareness
- And others as indicated by threat modelling

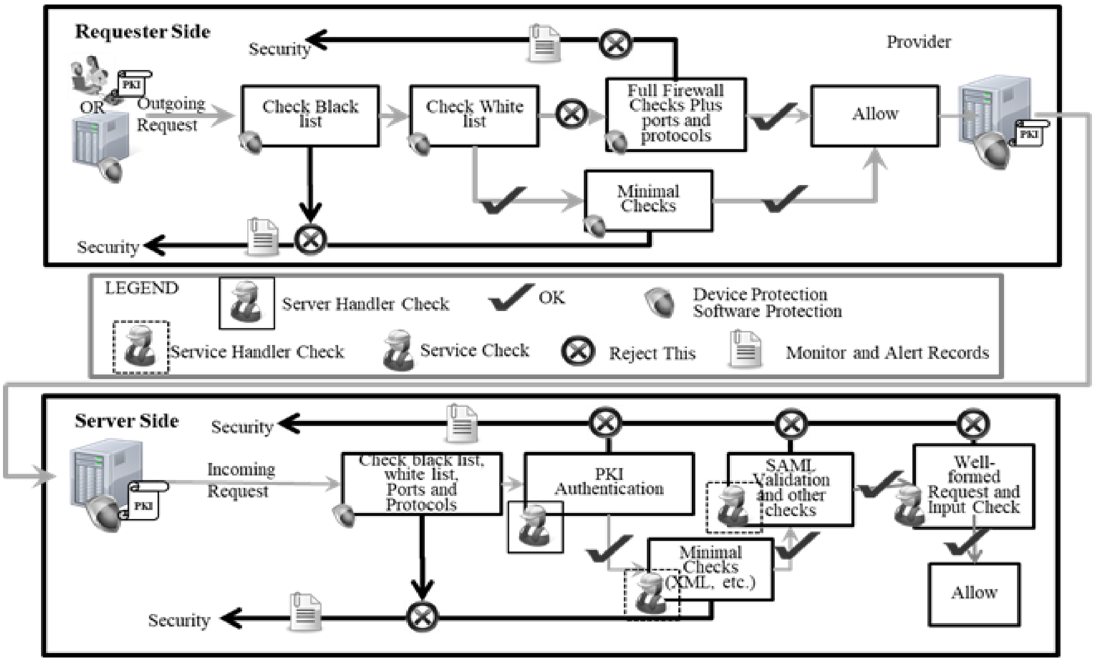

Figure 5 Protection Provided Without In-Line Appliances

The end-point protection system maintains an inventory of what is present (virtual and real) on all devices in the enterprise. Regular updates to this list ensures timely measures can be taken when an incident occurs. The protection system scans applications, configurations, permissions, services, registry entries, and other attributes to ensure that any changes from the baseline configuration have proper authorization. Any unauthorized or questionable differences from an approved baseline are reported to a central monitoring facility.

The protection system detects and removes malicious software from email by extracting, sandboxing and executing attachments to email in the user’s security context before the user can do this. The execution is monitored and if malicious the attachment is removed from the email and forwarded to the security team for further analysis. Phishing can overcome people’s mistrust of such attachments; this is an important part of device protection.

To prevent web-based attacks, the protection system flags potentially malicious sites to warn users. The protection system uses both heuristics and historical data to determine whether a site is safe or not. As search accesses many new sites, this is the ideal time for performance of such protection functions.

The protection system provides mechanisms to fix problems. Of course, a fully compromised system might be unresponsive to commands to fix certain issues, so this is not always possible. However, for most situations, remotely fixing the problem instead of requiring on-site manual intervention is the best course of action.

The protection system enforces policy on the local machine and enforcement of group policy or other methods for setting policy for compliance. Policies that are not enforced by the device itself must be monitored explicitly by the protection system.

The protection system keeps an accurate record of what the approved baseline configuration is for a given device [10].After a scan of the device, any differences are recorded and made available to the central monitor.

With new threats evolving through non-standard interfaces, such as USB, printers, and other attached devices, the protection system provides a way to manage these interfaces, either by monitoring or filtering traffic on them, disabling them, or using other methods to prevent attacks from these sources.

By closely monitoring code execution, the protection system prevents buffer overflows. Low-level system calls are monitored to track any attempts at writing to unallocated memory spaces, stopping both known and unknown buffer overflows from being exploited. This type of monitoring and prevention requires elevated privilege, as it requires access to system level resources, not just user data.

The protection system stops a user from installing new executable code, unless they are explicitly authorized. This prevents a user from compiling and running code downloaded from, or modified by, a malicious entity. It also provides a generic catch-all for any executables that may have bypassed the email or web monitoring functions. By stopping the user from installing executables, the protection system also stops malicious entities from using hijacked user accounts or sessions to run malicious code.

Enterprise enforcement of rules that govern behavior on their networks and devices is partially achieved by the protection system [11].Although many of these rules will already be handled through group policy or device Security Technical Implementation Guides (STIG), some activity can only be monitored dynamically through the protection system. For example, use of TLS with appropriate version, ciphers, two-way authentication using PKI, and use of appropriate extensions is not typically monitored using existing tools and must be implemented by the protection system.

7. END-POINT PROTECTION IN ELS

In ELS, an agent-type model is preferred. In this model, the packet header filtering and other security functions reside at the web server in the handler chain of the web service. The basic configuration of end-point protection in ELS shown in Figure 6; it provides a complete set of security functions for packet, message, and application layer security tailored for the specific web service being protected. The new functions added in the server are packet header inspection, packet content inspection, message content inspection, and application protection. These functions implement the ports and protocols protection, as well as other security functions normally provided by network devices such as intrusion detection/protection, packet and message content filtering, deep packet inspection, and application/web content filtering such as included in an application firewall.

A service requestor uses HTTPS to establish communication with the server hosting the target web service according to the ELS practice. The packet received by the destination server and the packet header are immediately inspected to perform the ports and protocols blocking, source whitelist/blacklist checking, and other filtering based on only the header, including stateful tracking of client addresses and ports. Until an HTTPS session is established, only packets addressed to the server's IP address and port 443 are allowed. Other ports may be opened as needed as part of the web service following establishment of the HTTPS session.

On the return path, the messages follow a similar process. In effect, the "packet header inspection" module performs the required network-layer filtering and can block traffic based on ports, protocols, and IP address. This makes the personal firewall essentially two-way in its filtering capacity.

In the ELS end-point protection architecture, the end-point protection modules can be configured to communicate with additional security monitoring appliances, such as a NetScout (or other traffic monitoring products), that can compile and track statistics about the security status of the server and the web service. The security appliances should be active entities and communicate with the server via TLS with mutual authentication. If required, the server could send the decrypted message traffic to other security appliances through this interface for additional security functions.

The end-point protection functions are configured through the server configuration management interface, which communicates with the server by TLS with mutual authentication. The ports and protocols, whitelist information, and any software updates are provided through this interface.

It is recommended that the initial configuration of the packet header deny all ports and protocols, both incoming and outgoing (as opposed to the traditional incoming only), and that permissions be configured in as they are identified as needed.

8. HANDLING AND INSPECTION OF TRAFFIC

Handling and inspection is done in software-only modules in the server. The handlers are embedded in the server handler chain at the point when and where the communication is prepared for their use and when and where the functionality has been distributed to packet-header inspection, packet content inspection, and message content inspection. Each of these may perform inspection related to intrusion detection or blacklist blocking, etc.

This is the preferred embodiment for enterprise applications. It moves the inspections to the point of the application itself by inserting handlers within the server and service to do the inspections at the point it makes most sense. The inspections that can be done without decrypting the packets may be done at the front of the web server because they are passive entities. Moving inspections of decrypted traffic inside the server not only preserves the end-to-end paradigm, it encapsulates the security and allows tailoring for the application itself. The encapsulated security with the application is virtualization ready.

9. CONCLUSIONS

We have reviewed the ELS security model and the end-to-end requirement within the enterprise. We have also reviewed the “normal” network defense process, and described the issues that the current network defenses raise and the vulnerabilities that may be introduced. Finally, we have provided an end-to-end approach that allows for both network inspection and reporting and the maintaining of unbroken encryption to the final destination, including enhanced defensive protections afforded by ELS. This approach is based on identifying the instances of official business and deferring the initial inspection until arrival at the target server. For enterprise operations, defining a clear end-to-end approach means a reduced attack space. The approach also reduces bandwidth requirements at the front door of the enterprise and may reduce the need for some load balancing. We have also reviewed the specific requirements for an enterprise level security that is bi-laterally authenticated and encrypted end-to-end. This paper is part of a body of work for high-assurance enterprise computing using web services. Elements of this work are described in [12-22].

REFERENCES

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

[2] 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.

[3] 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.

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

[5] 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.

[6] 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.

[7] 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, 125-129.

[8] 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.

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

[10] William R. Simpson and Kevin E. Foltz, Lecture Notes in Engineering and Computer Science, Proceedings of the World Congress on Engineering (WCE) 2018, “Enterprise End-point Device Management”, pp. 331-336, Imperial College, London, 4-6 July 2018, IBSN: 978-988-14047-9-4, ISSN: 2078-0958.

[11] William R. Simpson and Kevin E. Foltz, Lecture Notes in Engineering and Computer Science, Proceedings World Congress on Engineering and Computer Science(WCECS) 2017, Volume 1, “Enterprise Level Security: Insider Threat Counter-Claims”, pp112-117, Berkeley, CA. October 2017.

[12] William R. Simpson and Kevin E. Foltz, Proceedings of the Information Security Solutions Europe (ISSE) 2016, ISBN: 9781541211445, “The Virtual Application Data Center”, pp. 43-59,

https://www.amazon.com/isse2016-3-Information-Security-Solutions-Europe/dp/1541211448, Paris, France, November 2016.

[13] William R. Simpson and Kevin E. Foltz,Haeng Kon Kim • Mahyar A. Amouzegar (eds.), Transactions on Engineering Technologies, Special Issue of the World Congress on Engineering 2015, Chapter 15, pp. 205-220, "High Assurance Asynchronous Messaging Methods", 15 pp., DOI 10.1007/978-981-10-2717-8, Springer Dordrecht 2017.

[14] William R. Simpson and Kevin E. Foltz, Lecture Notes in Engineering and Computer Science, Proceedings of the World Congress on Engineering (WCE) 2017, "Assured Identity for Enterprise Level Security", pp. 440-445, Imperial College, London, July 2017, IBSN: 978-988-14047-4-9.

[15] William R. Simpson and Kevin E. Foltz, Proceedings of The 21th World Multi-Conference on Systemics, Cybernetics and Informatics: WMSCI, "Data Mediation with Enterprise Level Security",WMSCI 2017, Orlando, Florida, 8-11 July 2017, 6 pages.

[16] William R. Simpson and Kevin E. Foltz, Proceedings of the 22nd International Command and Control Research and Technology Symposium (ICCRTS), "Escalation of Access and Privilege with Enterprise Level Security", ISBN: 978-0-9997246-0-6, Los Angeles, CA. September 2017.

[17] William R. Simpson and Kevin E. Foltz, Sio-Long Ao, et. al. (eds.), IAENG Transactions on Engineering Sciences, Special Issue of the Association of Engineers Conferences 2016, Volume II, pp. 475-488, "Electronic Record Key Management for Digital Rights Management", 14 pp., World Scientific Publishing, Singapore, ISBN 978-981-3230-76-7, 2018.

[18] William R. Simpson and Kevin E. Foltz, "Secure Identity for Enterprises," IAENG International Journal of Computer Science, vol. 45, no. 1, pp 142-152, ISSN: 1819-656X, February 2018.

[19] William R. Simpson and Kevin E. Foltz, Proceedings of the 8th International Conference on Electronics, Communications and Networks (CECNet 2018), Volume 1, "Cloud Security and Scalability", pp 27, Bangkok, Thailand, November 2018.

[20] William R. Simpson and Kevin E. Foltz, "Insider Threat Metrics in Enterprise Level security," IAENG International Journal of Computer Science, vol. 45, no. 4, pp 610-622, ISSN: 1819-656X, December 2018.

[21] Simpson W. and Foltz K., Lecture Notes in Engineering and Computer Science, Proceedings World Congress on Engineering and Computer Science 2015, Volume 1, "Maintaining High Assurance in Asynchronous Messaging," pp. 178–183, Berkeley, CA, October 2015.

[22] William R Simpson, and Kevin E. Foltz, "Mobile Ad-hoc for Enterprise Level Security," Lecture Notes in Engineering and Computer Science: Proceedings of The World Congress on Engineering and Computer Science 2018, 23-25 October, 2018, San Francisco, USA, pp172-177.

Dr. Simpson has over two decades of experience working to improve systems security. He has degrees in Aeronautical Engineering and Business Administration. He also attended several schools for military and government training. He spent many years as an expert in aeronautics before delving into the field of electronic and system test, and he has spent the last 20years on IT-related themes (mostly security, including processes, damage assessments of cyber intrusions, IT security standards, IT security evaluation, and IT architecture).

Dr. Foltz has over a decade of experience working to improve security in information systems. He has degrees in Mathematics, Computer Science, Electrical Engineering, and Strategic Security Studies. He has presented and published research on different aspects of enterprise security, security modelling, and high assurance systems.