

Networking Aspects in the DPASA Survivability Architecture: An Experience Report*

Michael Atighetchi, Paul Rubel, Partha Pal, Jennifer Chong, Lyle Sudin
BBN Technologies
{matighet, prubel, ppal, jchong, lsudin}@bbn.com

Abstract

The same network infrastructure that is essential for the operation of today's high valued distributed systems can also be misused by malicious attackers. Experience shows that implementing absolute security or preventing cyber attacks completely is infeasible when systems must be highly interconnected and comprise of COTS components with unknown security characteristics. Therefore, focus is shifting towards making high valued distributed systems survivable, enabling them to operate through attacks. This paper focuses on the networking aspects of the DPASA¹ survivability architecture, which was used recently to defense-enable a DoD information system.

1. Introduction

The DPASA survivability architecture was recently tested in the context of a Joint Battlespace Infosphere (JBI) [1] system under sustained attacks from a sophisticated adversary (class A red team). The JBI is a distributed command and control information system (developed by the US Air Force) consisting of clients that communicate with each other using the publish-subscribe paradigm. The clients, organized in LANs, are connected to the JBI core, which itself is a LAN that hosts the services that mediate inter-client publish-subscribe operations over a public IP network.

Network communication is not only fundamental for the functional aspect (i.e., the client's ability to publish and subscribe); it is also the primary facilitator of attacks mounted by intruders. In order to defend against this threat, the design of the survivability

architecture introduces multiple mechanisms (representing both COTS and research grade technologies) to protect the confidentiality, integrity, and availability of network communication. In addition, the design makes extensive use of early detection and reporting of network incidents, and supports recovery and graceful degradation to mitigate compromises in the network.

The defense enabled JBI aims to establish a new high watermark in survivable system design and implementation. The architecture is designed not only to place a very high barrier to unauthorized entry from the outside (intrusion), but also to place similar resistance against an attacker who is attempting to expand his initial privilege or presence in the network from the inside. The architecture uses the principles of defense-in-depth and least privilege in order to achieve this goal. Redundancy and modularization are also used to facilitate tolerance and containment of attack effects.

Apart from the significant amount of internal evaluation of the design and implementation, the weeklong red team exercise in March 2005 was the most visible external evaluation of the defense-enabled JBI. Detailed analysis of the exercise results are not yet complete, but it is clear that the survivability architecture significantly improved the system's ability to deflect, detect and tolerate cyber attacks.

In this short paper we summarize the networking aspects of the DPASA architecture, our experience at the red team exercise, and lessons learned.

2. DPASA Overview

Many of today's security mechanisms, such as firewalls or intrusion detection systems, offer point solutions. However, an approach to combine the three major aspects of defense, namely protection, detection, and reaction, is needed to build a system that can withstand a wide range of threats for a considerable amount of time. The DPASA survivability solution

* This work is supported by DARPA under contract No. F30602-02-C-0134.

¹ DPASA stands for Designing Protection and Adaptation into a Survivability Architecture.

takes the form of a survivability architecture, which can be defined as the well-defined organization, placement, and interaction of a diverse set of defense mechanisms amongst the components of the undefended system. This allows for a careful organization of multi-dimensional layers of defense with each barrier backing up or managing the gaps of another.

The DPASA survivability architecture (see [2] for more details) rests upon a foundation of a robust network infrastructure. This foundation consists of networking elements that support redundancy in the architecture and provide security services such as packet filtering, source authentication, link-level encryption, and network anomaly detection. Middleware-based components within the architecture react to detected violations with defensive responses that change the configuration and usage of the networking fabric.

The DPASA architecture introduces redundancy in the JBI core in the form of four core quadrants (quads) where each quad runs on a dedicated LAN implemented as a Virtual LAN (VLAN). Each LAN is fronted by a VPN firewall, and the LANS are connected together over an emulated WAN. Hosts in the four quads run three different operating systems, i.e. SELinux, Windows, and Solaris. All client hosts run SELinux except for legacy clients, which run on Solaris hosts. Each SELinux and Windows host is equipped with an Autonomic Distributed Firewall Network Interface Card (ADF NIC) (see Section 3 for more details) that performs packet filtering and enforces encryption policies. The same functions are performed by an ADF equipped SELinux host configured as a bump-in-the-wire for each Solaris host². The core quads are organized into three zones. The executive (innermost) zone contains the overall management and control functions of the system. The operations (middle) zone contains hosts that are responsible for the main functional operations of the system, including the publish-subscribe-query service and supporting repositories (PSQ henceforth). The crumple (outer) zone acts as the region of first impact and proxies the operations zone functions for the clients. The zones concept physically impacts network wiring, and both inter- and intra-zone communications are strictly controlled via ADF policies and managed switches to limit attack propagation. The managed switches are powered via a custom device called the Quadrant Isolation Switch.

² The driver for the NIC card was not available for the Solaris platform.

Hosts within the zones are assigned specific functions. Access Proxies (AP) in the crumple zone perform checks on the dataflow being proxied, and establish a separation between the inside and outside of the core network. System Managers (SMs) in the executive zone gather system information, control other components via adaptive algorithms, and suggest appropriate defensive actions to a human operator. The operations zone in each quad consists of a Network Intrusion Detection System (NIDS³) as a network sensor, a Correlator for aggregating alert information, a Downstream Controller (DC) that acts as an intermediary between the AP and the SM, a PSQ server hosting the PSQ function and associated repositories, and an ADF Policy Server (PS) which is a 3rd party control server for the ADF NICs.

3. Network Design and Implementation

This section explains the role of the network-centric mechanisms in the architecture.

VPN firewalls form the first line of defense against attackers coming from the untrusted public IP network (WAN) into LANs of the defended system. All communication between client LANs and the four core LANs is sent over encrypted IPSEC tunnels. These VPNs effectively hide internal network addresses and payload content from packet sniffers on the WAN and deny opportunities for fine-grained traffic analysis. Any invalid or replay VPN traffic is dropped by the firewall before it enters a client or a core LAN.

Each host contains an **Autonomic Distributed Firewall Card** [3]. These firewalls are built into the network interface firmware, are separate from the operating system, and are controlled via encrypted messages by the Policy Server. This separation means that even if an attacker compromises a host, the attacker cannot corrupt the firewall policy on the compromised host. ADFs enforce separation between zones by only allowing network paths and protocols required by the system. Whereas a traditional firewall at a DMZ boundary needs to pass any traffic that any host behind it may possibly need, DPASA enforces a least-privilege policy for network traffic on a per-host basis. In addition to filtering out unwanted traffic both incoming and outgoing, ADFs also provide confidentiality and authentication using Virtual Private Groups (VPGs) [4]. VPGs encrypt messages sent between ADF NICs in the same group with a shared key. Non-members cannot send messages to the group without the key and ADFs prohibit a host from sending

³ Client LANs also have a NIDS.

spoofed packets. The VPN firewalls and the ADF NICs are an example of two layers of defense backing each other up. The coarse grained, but industry proven VPN firewalls stop unauthorized traffic before it gets a chance to potentially overwhelm the ADF cards, allowing the ADF cards to deal solely with traffic within the VPNs.

Attempts to compromise the JBI core from within the VPN layers are confronted by the Access Proxy hosts. APs are equipped with 2 ADF cards and contain dedicated proxy processes for all servers in the operation zone. The proxies receive incoming traffic from the exterior ADF, inspect the content, enforce rate-limiting, and then forward the traffic to the corresponding server via the interior ADF. Proxy processes, like most other DPASA processes, are started and executed in their own process protection domain, which limits their privilege and isolates them from each other via technologies such as SELinux, Cisco Security Agent, and Java Security Policies.

Four Managed Switches, one per quad, connect the individual quad LANs into a core network. These switches can be configured to block any connection not permitted by the architecture at the link level via source-port filtering.

Power to the managed switches is controlled via a set of four Quadrant Isolation Switches (QISs), allowing a human operator to isolate an entire quad if necessary. Usually, such a drastic measure is taken when other adaptive attempts to recover the quad have failed. Each QIS, a custom piece of hardware, is cross-connected to the other 3 QISs and its local SM via serial ports. Each SM can vote whether or not to isolate any of the four quads. In order to cut off a quad from the network, three SMs need to agree, in which case the DPASA protocols will continue to work in degraded mode with less than 4 quads.

During network hardening, unnecessary services were disabled/uninstalled from hosts, switches, and other network-enabled devices. This ensures that any alarms or traffic involving the removed services are immediately noted as a possible attack. In order to deny ARP attack points to adversaries (see [5,6,7] for more details), static ARP tables are used wherever possible. Furthermore, TCP/IP stack settings were hardened on all three operating systems to increase resistance to attacks such as SYN flooding.

The defense-enabled system utilizes NIDS appliances on all LANs for network detection. These systems analyze network traffic and report anomalies to the Correlator. Through the use of both ADF and non-ADF NICs, the NIDS can perform signature-based checks on

VPG traffic and detect non-VPG traffic indicative of attack activity.

In addition to the NIDSes, DPASA uses violations from various policy enforcement mechanisms for intrusion detection. At the application level, illegitimate network actions (such as opening a non-specified server port) are detected. At the network level, each time an ADF equipped host sends an outgoing packet that violates the host's ADF policy, the NIC drops the packet and generates an audit alert to the PS.

To deal with components abusing allowed communication paths, the DPASA components and protocols are designed to be robust against malformed data and report possibly malicious actions using custom alerts. These alerts include flooding, replay, and traditional access control violation reports. Furthermore, most DPASA components send heartbeat messages to the core. If heartbeat messages are failing to reach the core on account of network problems, there is a good chance that scenario traffic would fail similarly. During the red team exercise, missing heartbeats often indicated the first sign of an attack.

Automatic adaptive network defenses enable the DPASA system to continue to function even for cases in which the attacker has managed to break through protection layers and started to affect availability or integrity of the system. The overall goal is to engage the attacker through both local and more coordinated dynamic changes in the system, and thus slowing down attack propagation and keeping the system operational, albeit at a reduced level (i.e., graceful degradation). For example, the SM can isolate a host by putting all its ADF NICs into block all mode, which causes the NICs to drop all incoming and outgoing traffic. Furthermore, application-level queues with threshold schemes are used for sending out alerts in a throttled manner, and the clients run greedy algorithms for selecting the best (fastest non-corrupted) AP in the current environment.

4. Evaluation

Validation activities took place throughout the development of the defense-enabled JBI. We identified ADF NICs as the main target for network attacks through internal white boarding sessions and stochastic model-based simulation [8]. Further empirical studies explored the behavior of ADF NICs under stress [9]. The remainder of this section focuses on experiences and results from the March 2005 red team exercise.

The main objective of this exercise was to determine whether the defense-enabled JBI could survive 12 hours of sustained attacks and complete its mission.

Two different red teams participated on separate days; each had complete knowledge of the system, and was fairly unrestricted in terms of rules of engagement. In addition, both red teams had high-privilege access to the WAN switch, enabling them to sniff all traffic and send management traffic to the switch. Both red teams aimed their attacks mainly at the VPN firewalls in order to decrease network availability. The first red team was successful in flooding the firewalls with massive amounts of seemingly valid replay traffic⁵, causing lost heartbeat messages from clients and delay in PSQ requests. However, the red team had no insight into the effectiveness of this attack, and moved on to different unsuccessful attacks. The system recovered and the mission was completed. The second red team exploited a zero day attack involving the Dynamic Trunking Protocol (DTP) on the WAN switch. This attack was highly effective, disrupting all legitimate WAN traffic and stopping the mission. Using the same attack over a port without DTP, however, did not have any impact on the defended system.

In order to test defense in depth, we tested a number of system configurations with escalated levels of red team privilege. The defense-enabled system was able to provide mission functionality in all such cases.

5. Conclusion

The red team evaluation of the defense-enabled JBI demonstrated that it is possible to develop a well-configured adaptive system that can defend critical system functionality against sustained attacks from a sophisticated adversary over a reasonably long time. While some attacks caused disruptions during the mission, they required high-privilege local access on the emulated network, and potentially a flaw in the COTS VPN routers. This level of access would be difficult to obtain in a fielded system, but not impossible. It should be noted that the attacks did not impact integrity and confidentiality, and did not compromise or damage any of the host OSes.

This experience shows that, in the context of defending the network, careful combination of appropriate technologies can make it very hard for attackers to find a successful attack. This is especially true when dealing with attacks aimed to breach integrity or confidentiality.

The DPASA experience is an incremental milestone in the ongoing fight against cyber threat. Examples where we expect to find the next generation of

survivability tools and technologies include exploration of 1) regenerative ideas to replace lost capabilities with newer and improved ones, perhaps with different security policies and configuration settings, 2) how to combine adaptive security with protection measures that are grounded in hardware or sound cryptographic techniques, and 3) how to use policy violations and IDS alerts to compute trustworthiness of system components.

6. References

- [1] AFRL JBI homepage - <http://www.infospherics.org>
- [2] "Case Study: The Intrusion Tolerant JBI", extended version of this paper, <http://www.dist-systems.bbn.com/papers/2005/DpasaCaseStudy/index.shtml>
- [3] Payne, C., and Markham, T. "Architecture and applications for a distributed embedded Firewall". In 17th Annual Computer Security Applications Conference (December 2001).
- [4] Tom Markham, Lynn Meredith, and Charlie Payne. "Distributed embedded firewalls with virtual private groups". In DARPA Information Survivability Conference and Exposition -Volume II, Washington, D.C., April 2003. DARPA, IEEE.
- [5] Michael Atighetchi, Partha Pal, Franklin Webber, Richard Schantz, Christopher Jones, and Joseph Loyall. "Adaptive Cyberdefense for Survival and Intrusion Tolerance". IEEE Internet Computing, Vol. 8, No. 6, November/December 2004, pp. 25-33.
- [6] W. Nelson, W. Farrell, M. Atighetchi, S. Kaufman, L. Sudin, M. Shepard, and K. Theriault. "APOD Experiment 1: Final Report", BBN Technologies LLC, Technical Memorandum 1311, May, 2002.
- [7] W. Nelson, W. Farrell, M. Atighetchi, J. Clem, L. Sudin, M. Shepard, and K. Theriault, "APOD Experiment 2: Final Report" BBN Technologies LLC, Technical Memorandum 1326, Sep, 2002.
- [8] F. Stevens, T. Courtney, S. Singh, A. Agbaria, J. F. Meyer, W. H. Sanders, P. Pal, "Model-Based Validation of an Intrusion-Tolerant Information System", from the 23rd Symposium on Reliable Distributed Systems (SRDS 2004), posted July 16, 2004.
- [9] M. Ihde and W. H. Sanders, "Barbarians in the Gate: Packet Flooding NIC-based Distributed Firewalls", submitted for publication

⁵ Mostly Internet Security Association and Key Management (ISAKMP) traffic.