



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Disconnect from the Internet – Whale’s e-Gap In-Depth

Kevin Gennuso

GSEC Practical Assignment Version 1.2f

September 13, 2001

Introduction


While firewalls are a critical part of today’s externally-connected networks, their weaknesses have been revealed time and time again. Some of the world’s most widely implemented firewall systems, including Check Point’s FireWall-1, Cisco’s PIX, NAI’s Gauntlet, and Axent’s Raptor, have had serious vulnerabilities exposed in recent history, and all of these could be exploited remotely by a malicious party in order to gain access to the backend systems. These vulnerabilities were able to exist because of three fundamental design flaws that all firewalls have: a) they all speak TCP/IP, a protocol fraught with inherent vulnerabilities; b) they all connect both the DMZ and internal network in the same way that a router does; and c) holes must be created to allow network traffic to flow through to the inside. These three flaws make firewalls a less-than-optimal solution for protecting critical systems from intruders, especially when those systems house the extremely sensitive data found in today’s e-business systems.

The arrival of air gap technology is a welcome addition to any security engineer’s arsenal of perimeter defenses. Air gap technology allows the engineer to disconnect his or her application and database servers from the Internet at the physical layer of the OSI model. While certainly not a replacement for a firewall, air gap systems can certainly add another layer of protection to your network and the data residing on it. As Michael Hurley points out in his SANS paper “Network Air Gaps – Drawbridge to the Backend Office”, air gap systems strip out all TCP/IP and operating system commands. This prevents a malicious party from launching nearly every remote exploit available today. By doing so, air gaps are a solution to the three design flaws that are at the foundation of all firewalls. While there are a number of variations on the air gap concept, the focus of this paper will be on one implementation of this technology: Whale Communications’ e-Gap.

e-Gap to the Rescue

The e-Gap product is a unique combination of hardware and software that allows for this “disconnection from the Internet” to take place. Instead of physically connecting both the DMZ and internal networks, the e-Gap system allows communication to happen by way of a solid-state SCSI memory bank connected to a host in the DMZ and a host on the inside. The SCSI appliance has a high-speed analog switch that allows only one of the two hosts to communicate with it at a time. When a request comes in from the external network, the host in the DMZ running Whale’s external software accepts the request, strips out all TCP headers, and dumps only the data payload into the SCSI e-disk. The analog switch then disconnects the external host and connects the internal host to the appliance. The host, running Whale’s internal software, then processes the request, adds its own TCP headers, and forwards it along to the real application server. This process happens in reverse when the response is sent back to the user. The system can also be set to “one-way mode” using a physical key that locks the analog switch

in place, allowing either inbound or outbound communication only.



This description and diagram by themselves do not accurately depict what happens behind the scenes. In order to paint the complete picture of how this technology really works, we must take a closer look at the e-Gap software's "shuttle" functionality. We will delve further into the four most common shuttles provided by Whale: the Web Shuttle, File Shuttle, Mail Shuttle, and Dynamic Services Shuttle. A fifth shuttle exists in the form of an SDK for non-TCP-based messaging, but will not be discussed here.

Web Shuttle

The Web Shuttle allows users who wish to make an HTTP or HTTPS request to a server behind the e-Gap to do so, even though the server is not connected to the DMZ. The external e-Gap host acts as a virtual web server, which handles the initial request and passes the data payload into the SCSI appliance. The analog switch then disconnects the external server, connects the internal host to the memory bank, and the internal host retrieves the request. Now that the internal server has the request and is untouchable by a potential cracker, the security engineer has a number of filtering options at their disposal.

The first of these filters allow the engineer to permit only the specific URLs that are valid for each web site. A Recorder utility provided in the software allows the engineer to "record" which

URLs are valid for the site, and then implicitly deny everything else. There is also an Optimizer utility that looks for duplicate requests and combines those with similar parameters, making the overall set of rules smaller, and therefore, faster. The advantages of URL filtering are obvious. Many Web servers, such as Microsoft's IIS, are susceptible to malformed URL attacks, most recently demonstrated by the outbreak of the Code Red worm. This type of attack is stopped in its tracks since the e-Gap will not allow invalid URLs to pass through to the real Web server. The latest version of Whale's e-Gap software even takes into account the extreme URL malformations produced by Rain Forest Puppy's whisker code such as URL encoding and HTTP mis-formatting.

The Web Shuttle also allows any application behind the e-Gap to have SSL encryption and user authentication retrofitted, even if the app was never written to use them. The e-Gap supports third-party authentication products such as SecurID, RADIUS, LDAP, and TACACS+, and allows for easy configuration of all four. The engineer can also define their own authentication scheme with relative ease. The e-Gap software comes with canned HTML and ASP pages which handle the user validation process, but these can be changed if needed.

Additional filters are also included to extend the flexibility and control of the protected Web site. Server Name Translation ensures that any outgoing HTML containing the name of the real Web server is replaced with that of the virtual Web server and vice versa. Finally, the Web Shuttle can utilize the IP Blacklist feature, where the engineer can build a list of undesirable or troublesome hosts or networks and implicitly deny them.

File Shuttle

Whale describes file shuttling over the e-Gap system as "ultra-fast, secure 'sneakernet'", since this is essentially how all data transfers function. The File Shuttle can use FTP, RCP, or any other file transfer protocol. Once a file is written to the e-Gap host (be it the external host for incoming transfers, or the internal host for outgoing transfers), the File Shuttle uses the SCSI memory bank to mirror the data on the opposite server at 100Mbit/sec, assuming no other traffic is passing through the e-Gap. If file transfers are the main reason why the e-Gap is being implemented, the administrator can assign a higher priority to this type of traffic using built-in Quality of Service (QoS) settings.

Filters can be configured on the internal e-Gap host (away from potential attackers on the untrusted network) in the same way that they were used in the Web Shuttle. These rules can be extremely granular if need be. For example, one could configure the File Shuttle to only allow files to be uploaded with a JPG extension, a particular string in the file name, and a creation date falling within a one-week period. Filters can also trigger supported anti-virus software, such as Norton, McAfee, or eSafe. Depending on how the filters are configured, the rule sets will tag the file as either "accepted" (file is copied to opposite server), "deferred" (file remains is held until approved by a security engineer), or "rejected" (file is discarded and processing begins on next request).

Mail Shuttle

The Mail Shuttle works very similarly to the File Shuttle with its ability to filter and process SMTP messages based on user-defined rule sets. When a message is received by the virtual SMTP server (i.e. the external e-Gap host), it strips out all protocol headers and shuttles the message to the internal server via the SCSI appliance. Once it exists on the trusted host, away from the prying eyes of would-be attackers, the message is subject to strict content inspection and virus scanning.

The built-in SMTP filters allow for very close inspection of the data within the message. They can be configured to trigger due to matches on attachment extension, message size, and date range. They can also be set to trigger on regular expression string matches, either within the message, its attachments, and/or the sender, recipient, and subject fields. There are also SMTP processors that allow the administrator to replace a matching character string with one that they define. The filters and processors can also call external batch files and executables, should the need to use a third-party content inspection or virus scanning package arise (currently, the only virus-scanning package natively supported by the e-Gap for email is eSafe). Depending on how the rules are defined, the filters will flag messages as being “accepted” (sent for delivery), “delayed” (held for a specified period of time), “deferred” (sent to the engineer for review), or “rejected” (thrown out due to a rule violation).

Dynamic Services Shuttle

Whale defines “dynamic services” as user-defined protocols that operate over TCP. This basically allows the security engineer to use the e-Gap system to permit traffic other than HTTP, HTTPS, and SMTP to talk across the two disconnected networks using the shuttling method similar to the ones described for the other protocols. The amount of filtering and content inspection that can be done here out of the box is extremely limited, but can be tailored using the SDK. Dynamic services can also be made subject to the IP Blacklist feature used in the Web Shuttle.

High Availability Array

In a configuration where network load is very heavy, or where uptime is of critical importance, Whale offers a High Availability option that allows multiple internal and external e-Gap hosts to function as one. Whale allows you to use a hardware-based solution such as an F5 or Cisco device, or you can use the Resonate Central Dispatch software that comes with the e-Gap Array. The High Availability option also comes with a Central Configuration utility that deploys all configuration settings across all of the machines belonging to the Array, which reduces administrative overhead and allows changes to be made across all systems much easier. Finally, there is a KeepAlive feature that works in conjunction with the Web Shuttle that ensures that shuttling is functioning properly. The external e-Gap host will periodically poll a defined Web site to confirm that the request makes it through to the real server on the other side of the air gap, allowing for immediate detection of failed sites or e-Gap hosts.

Monitoring the System

For monitoring the performance and system events on the system, Whale ships their e-Gap Monitor application with the product. It can be run on either the internal host or on a remote system, although you must configure the server you want to monitor to allow remote monitoring connections. The interface is broken down into three monitoring modes: System Info (general system performance), Application Messages (communications along each memory channel or “trunk”), and Security Messages (violations of filters on all shuttles). You then have the ability to either view individual messages or categorized statistics about each trunk. You can also create real-time graphs that display information from any of the three monitoring modes. Unfortunately, there is no built-in reporting mechanism that allows for the creation of monthly utilization reports, for example, but Whale does supply a MIB file for those who wish to use a third-party SNMP monitoring tool to collect statistics and create reports.

For monitoring and making decisions on email and files that were inspected using filters and rule sets, Whale provides the e-Gap Security Console. This application allows you to view logs of every file and email that passed through the system, and it gives you extremely detailed information about each item and its status. It also allows you to approve or reject any file or email that was tagged as “deferred” by a filter once the security engineer has inspected the content by hand.

Potential Vulnerabilities

One cannot assess a security technology without taking a look at the potential weaknesses associated with it. Since the Whale software functions only on Microsoft platforms (Windows NT and Windows 2000), that should automatically raise a warning flag. Whale strongly suggests that the security engineer “harden” the external machine as much as possible, although they do not go into specifics about those hardening procedures. One of the first things that any engineer should do to harden a box is to remove any services that are not essential to the machine’s operation. For example, shutting down the Server service and setting it to Manual startup is one of the easier tasks to complete. Disabling File and Print sharing is also relatively simple and well-documented. Administrators should follow the “Securing Windows” guides from SANS for the relevant version of Windows being used with the e-Gap, but even after the machine is tighter than a drum, there will still be ports listening that shouldn’t be. It doesn’t hurt to put the external e-Gap host behind a firewall for an added layer of security.

So now that the boxes are hardened, what attacks could be launched against the external host? There aren’t many. The external host only listens for connections coming in on a particular TCP port and then dumps that request into the SCSI appliance, making it relatively dumb. Absolutely no processing happens on the external machine aside from the removal of TCP headers. In addition, the e-Gap software can only receive configuration commands and rule changes via its SCSI controller, making it impossible for someone to send it spoofed commands over a remote network. There may be a way to flood the external host with enough traffic so as to create a denial of service scenario, but so far, I am not aware of a way to do so.

Attacking the e-Gap system from within the trusted network is a little easier, but it would still be

extremely difficult to modify any of the filters or rule sets configured for the protected site. This is due to the multiple passwords and levels of encryption that are used to protect the console and configuration files. In order to make changes to the rules, the attacker would first need to gain administrator access on the inside host. They would then need to launch and enter the password for the e-Gap Configuration application. Once this is done, the attacker could make changes, but those changes are not pushed to the e-Gap hosts until the passphrase is entered. This passphrase unlocks the Blowfish-encrypted configuration files, which are virtually impossible to crack without it.

Potentially successful attacks could be made against IIS, which must run on the internal host for the Web Shuttle to function, if it is not properly patched. There are a number of holes in a default IIS installation, not the least of which could be exploited to gain administrator access to the machine. Also, the remote monitoring port (TCP 50001) is extremely sensitive to connection flooding, pegging the CPU at 100% and stopping all communication through the e-Gap after about 60 connections. This allows for an easy denial of service attack that will take the e-Gap system offline, although Whale is promising a fix for this issue in an upcoming software release. The lessons here have been stated often: keep your patch levels current and, as with all machines with access to the internal trusted network, you should take precautions to prevent unauthorized access from the inside.

Conclusion

Although today's publicly-connected networks could not survive without firewalls, their vulnerabilities are dangerous and can threaten the integrity of today's e-business systems. Holes in Web, FTP, and mail server software are even more menacing and pervasive, and although software vendors are claiming to become more "security-centric", this has been slow in the coming. Even when fixes to known vulnerabilities are released, it is sometimes unrealistic to take every production server offline to apply the patch. When your databases house detailed customer information or thousands of credit card numbers, you must utilize every tool at your disposal to mitigate the risk of that sensitive data being tampered with or stolen. An old saying the security realm states: "The best way to secure a box is to unplug it from the network." Whale Communications has provided a way to do exactly that with their e-Gap product.

References

Whale Communications – e-Gap Products
<http://www.whalecommunications.com/0200.htm> (30 August 2001)

Whale Communications – Air Gap Technology
<http://www.whalecommunications.com/0300.htm> (30 August 2001)

Bobbitt, Michael. "(Un)Bridging the Gap" Information Security. July 2000 (2000): 35 – 47
<http://www.infosecuritymag.com/articles/july00/cover.shtml> (30 August 2001)

Whale Communications. "e-Gap User Guide Version 2.2", May 2001. 4 – 11

Chambers, Chris, Dolske, Justin, Iyer, Jayaraman. "TCP/IP Security"
http://www.linuxsecurity.com/resource_files/documentation/tcpip-security.html (30 August 2001)

Bellovin, S.M.. "Security Problems in the TCP/IP Protocol Suite" April 1989
http://www.ja.net/CERT/Bellovin/TCP-IP_Security_Problems.html (2 September 2001)

SecuriTeam.com – Security News Archive, March 2000 thru September 2001
<http://www.securiteam.com/securitynews/archive.html> (2 September 2001)

Curtin, Matt & Ranum, Marcus J. "Internet Firewalls: FAQ rev. 10.0" 1 December 2000.
<http://www.faqs.org/faqs/firewalls-faq/> (2 September 2001)

Edwards, John. "Unplugging Cybercrime" 1 May 2000
http://www2.cio.com/archive/050100_development_content.html (2 September 2001)

McClure, Stuart, Scambray, Joel, Kurtz, George. "Hacking Exposed: Network Security Secrets & Solutions", McGraw-Hill, 1999. 314

Hurley, Michael. "Network Air Gaps – Drawbridge to the Backend Office" 4 April 2001.
<http://www.sans.org/infosecFAQ/firewall/gaps.htm> (2 September 2001)

Rain Forest Puppy. "A look at whisker's anti-IDS tactics" 30 December 1999
<http://www.wiretrip.net/rfp/pages/whitepapers/whiskerids.html> (4 September 2001)

Whale Communications. "e-Gap Version 2.2 Release Notes", April 2001. 34