



The Trusted Source for Computer Security
Training, Certification and Research

[SANS School Store](#)[Reading Room](#)[Internet Storm Center Certification](#)[GIAC](#)[S.C.O.R.E.](#)[Vendor Opportunities](#)

January 29, 2004: Tool Talk: PGP Education Series: Seamless Encryption in Government

[About SANS](#) | [Contact SANS](#) | [SANS Forum](#) | [What's New](#) | [F.A.Q.](#) | [PGP Key/Local Copy](#) | [Surveys](#) | [Webcasts](#)[Computer Security News](#) | [Research Projects](#) | [Resources](#) | [Press Room](#) | [Sample Policies](#) | [Top 20 List](#)

[Intrusion Detection FAQ](#)

What is the Role of Security Event Correlation in Intrusion Detection?

Introduction

Millions of dollars have been invested in security products such as firewalls, intrusion detection, and strong authentication over the past several years. However, system penetration attempts continue to occur and go unnoticed until it is too late. As a consequence financial losses continue to skyrocket for organizations. According to the 2002 CSI/FBI Computer Crime and Security Survey, average losses per respondent topped \$2,000,000 for the year! It is not that security countermeasures are ineffective against intrusive activity. Indeed, they can be very effective within an organization where security policies and procedures require analysis of security events and appropriate incident response. However, as pointed out by Steven Northcutt of SANS, deploying and analyzing a single device in an effort to maintain situational awareness with respect to the state of security within an organization is the "computerized version of tunnel vision". Security events must be analyzed from as many sources as possible in order to assess threat and formulate appropriate response. Extraordinary levels of security awareness can be attained in an organization's network by simply listening to what its devices are telling you.

This paper will demonstrate to intrusion analysts why correlative analysis must occur in order to understand the complete scope of a security incident.

Correlation Simplified


When law enforcement agents investigate a murder, they do more than examine the body for clues. The investigative process calls for searching the surrounding crime scene, interviewing individuals who know the victim, and soliciting requests to the public for anyone who might have information related to the crime.

A similar process should apply to intrusion analysis. If your web server is attacked, analyze more than the web server logs. Search the firewalls and intrusion detection systems protecting the web server for other activity from the source address. Share your log information about the activity with other analyst via websites such as incidents.org. Reviewing all of the information collectively provides a more complete picture of the

incident and assists in answering the who, what, when, where, and why's of an attack.

Correlation Demonstrated

Understanding the concepts of correlation can be dramatically simplified if the responses of various network devices are examined in the face of a probe or attack. The following scenario demonstrates how independently obscure security events can be correlated from multiple logs, and in doing so provide the higher level of vision necessary for accurate and expeditious intrusion analysis. We will conduct intrusion analysis of the log data independently and the collectively to show how a more complete picture can be attained through correlative analysis. The network below depicts a typical network layout where common countermeasures such as firewalls and intrusion detection systems are deployed.



In this network, we have established multiple log generators of interest (moving from the perimeter inward):

- Router: Access control lists (ACL's) can provide perimeter packet filtering with (typically) syslog-style alerting. If properly configured ACL's provide network perimeters a first line of defense by defining policy for traffic. When attackers conduct reconnaissance against target networks, ACL's will typically deny at least one element of the attacker's probes and generate a log entry documenting the action. For this exercise, this will be a Cisco 2600 series router.
- Firewall: Depending on the type of firewall and its configuration, extensive visibility can be gained from firewalls. Application proxy firewalls can provide extensive logging and access control capabilities that allow extensive visibility into network traffic passing through the perimeter. For this exercise, the firewall will be a Gauntlet (proxy) firewall.
- Network IDS: By inspecting network traffic, suspicious activity can be flagged and alerts generated. Depending on the type, network IDS typically monitors network traffic for suspicious activity against a database of well-known vulnerabilities and exposures. The alerts generated by network IDS provide valuable interpretative

analysis into network activity. However, because of the high rate of false positives associated with network IDS, correlation with other log sources is a must for alert validation. For this exercise, the network IDS will be Snort running on Linux.

- Application servers (ie. www, ftp, email): Application servers house the data of interest within organizations. As the targets of malicious activity, application servers are the reasons the rest of the security infrastructure is deployed. Common Internet services typically log both successful and failed transactions. These logs provide valuable insight into the overall intent of an attacker along with the success and/or failure of attacks. For this exercise, we will be probing an Apache web server running on Linux.

For this discussion, we will assume the devices have been configured with full logging capabilities such that maximum visibility is attained. For example, the firewall is configured to log both accepted and denied attempts.

We will analyze the log response of all of the network devices while Attacker1 is launching a series of probes searching for exploitable CGI scripts. This activity is being conducted by an attacker at 152.63.146.6 against an Apache web server (www) running on a typical Linux distribution. For this exercise, we will confine the probes to three well known exploits:

- CVE-1999-0067: CGI phf program allows remote command execution through shell metacharacters.
- CVE-1999-0172: FormMail CGI program allows remote execution of commands.
- CVE-1999-0936: BNBSurvey survey.cgi program allows remote attackers to execute commands via shell metacharacters.

Independent Analysis

First, let's review the log activity related to the probe activity for each device in the path of the probes. We will first analyze the information independently and later we will correlate all of the log data for a more complete picture of the incident.

Router Logs (Cisco):

```
May 31 09:27:44 router.company.com 1410875: May 31 09:27:43: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1459) -> xxx.yyy.zzz.1(80), 1 packet
```

```
May 31 09:27:50 router.company.com 1410880: May 31 09:27:50: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1673) -> xxx.yyy.zzz.2(80), 1 packet
```

```
May 31 09:27:54 router.company.com 1410883: May 31 09:27:53: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1750) -> xxx.yyy.zzz.3 (80), 1 packet
```

```
May 31 09:27:57 router.company.com 1410885: May 31 09:27:56: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1722) -> xxx.yyy.zzz.5(80), 1 packet
```

```
May 31 09:27:58 router.company.com 1410886: May 31 09:27:57: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1930) -> xxx.yyy.zzz.6(80), 1
```

packet

May 31 09:28:01 router.company.com 1410888: May 31 09:28:00: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(1976) -> xxx.yyy.zzz.7(80), 1 packet

May 31 09:28:05 router.company.com 1410891: May 31 09:28:04: %SEC-6-IPACCESSLOGP: list from-internet denied tcp 152.63.146.6(2167) -> xxx.yyy.zzz.8(80), 1 packet

.
 . <data pruned>
 .

For the source address 152.63.146.6, we observe that on May 31 at 09:27 a series of connect attempts occurred directed towards the xxx.yyy.zzz.0/24 network. This log data has been pruned but other entries show the activity directed towards the entire class C. By the destination TCP port 80 connection attempts, it appears as though 152.63.146.6 is conducting a broad scan searching for web servers. It is interesting to note that there was no denied log entry for an access attempt to xxx.yyy.zzz.4 because this is the web server in our network. Our router access control lists have been configured to allow inbound TCP port 80 traffic with ephemeral source ports to xxx.yyy.zzz.4 because this is our company web server.

By only looking at the router logs, the information presented to us suggests 152.63.146.6 swept the entire class C looking for web servers. Independently reviewed, we have no other insight into the intentions of the activity.

In summary for the router logs:

Who: 152.63.146.6
 What: Broad scanning of xxx.yyy.zzz.0/24 network for web servers.
 Likely found xxx.yyy.zzz.4.
 When: May 31 at 09:27-09:28
 Where: Company DMZ network
 Why: Likely reconnaissance.

Firewall Logs (Gauntlet):

**Jun 1 06:08:50 firewall.company.com http-gw[29142]: log
 host=nodnsquery/152.63.146.6 protocol=http cmd=get dest=xxx.yyy.zzz.4 path=/cgi-bin/phf ID=29142174970**

**Jun 1 06:08:54 firewall.company.com http-gw[29142]: log
 host=nodnsquery/152.63.146.6 protocol=http cmd=get dest=xxx.yyy.zzz.4 path=/cgi-bin/formmail ID=29142174971**

**Jun 1 06:08:58 firewall.company.com http-gw[29142]: log
 host=nodnsquery/152.63.146.6 protocol=http cmd=get dest=xxx.yyy.zzz.4 path=/cgi-bin/survey.cgi ID=29142174972**

For the source address 152.63.146.6, we observe that on June 1st, a series of http

connects were allowed to the corporate web server. We see that the URL path shows attempted access on three separate cgi scripts: phf,formmail, and survey.cgi. Reviewed independently, we have no way of knowing if the access attempt was successful. All we know is an attempt was allowed. And unless we are versed in known cgi vulnerabilities, we may simply overlook the activity as legitimate.

In summary for the firewall logs:

Who: 152.63.146.6
 What: Three http connects to xxx.yyy.zzz.4 with access attempts of cgi scripts: phf, formmail, and survey.cgi. Unknown if the scripts were accessed. There were no other http connections from this host so it appears as though this is malicious activity not associated with any other normal web traffic.
 When: June 1 at 06:08:50
 Where: Company DMZ network
 Why: Research shows that phf, formmail, and survey.cgi are all exploitable scripts. These access attempts in isolation suggest malicious activity because if accessed, these scripts could allow remote command execution.

IDS Logs (Snort):

```
[**] [1:886:3] WEB-CGI phf access [**]
[Classification: Attempted Information Leak] [Priority: 2]
06/01-06:08:50.764332 152.63.146.6:3308 -> xxx.yyy.zzz.4:80
TCP TTL:52 TOS:0x0 ID:61884 IpLen:20 DgmLen:280 DF
***AP*** Seq: 0x591AF831 Ack: 0x92D23FAF Win: 0x16D0 TcpLen: 32
TCP Options (3) => NOP NOP TS: 59902357 300726
[Xref => http://www.securityfocus.com/bid/629]
[Xref => http://www.whitehats.com/info/IDS128]
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0067]

[**] [1:884:2] WEB-CGI formmail access [**]
[Classification: Attempted Information Leak] [Priority: 2]
06/01-06:08:54.411065 152.63.146.6:3309 -> xxx.yyy.zzz.4:80
TCP TTL:52 TOS:0x0 ID:15383 IpLen:20 DgmLen:285 DF
***AP*** Seq: 0x85C51FDB Ack: 0xC0D4B803 Win: 0x16D0 TcpLen: 32
TCP Options (3) => NOP NOP TS: 59974615 372988
[Xref => http://www.securityfocus.com/bid/1187]
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0172]
[Xref => http://www.whitehats.com/info/IDS226]

[**] [1:871:2] WEB-CGI survey.cgi access [**]
[Classification: Attempted Information Leak] [Priority: 2]
06/01-06:08:58.609416 152.63.146.6:3310 -> xxx.yyy.zzz.4:80
TCP TTL:52 TOS:0x0 ID:32890 IpLen:20 DgmLen:295 DF
***AP*** Seq: 0x8B55C63C Ack: 0xC624745D Win: 0x16D0 TcpLen: 32
TCP Options (3) => NOP NOP TS: 59983434 381809
[Xref => http://www.securityfocus.com/bid/1817]
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0936]
```

For the source address 152.63.146.6, we observe that on June 1st, a series of cgi access

alerts occurred. The alerts point to vulnerabilities associated with these scripts that can be used for remote command execution.

These are the Snort rules that triggered these alerts:

```
alert tcp $EXTERNAL_NET any -> $HTTP_SERVERS 80 (msg:"WEB-CGI phf access";flags: A+; uricontent:"/phf"; nocase; reference:bugtraq,629; reference:arachnids,128; reference:cve,CVE-1999-0067; classtype:attempted-recon; sid:886; rev:3;)
```

```
alert tcp $EXTERNAL_NET any -> $HTTP_SERVERS 80 (msg:"WEB-CGI formmail access"; flags: A+; uricontent:"/formmail"; nocase; reference:bugtraq,1187; reference:cve,CVE-1999-0172; reference:arachnids,226; classtype:attempted-recon; sid:884; rev:2;)
```

```
alert tcp $EXTERNAL_NET any -> $HTTP_SERVERS 80 (msg:"WEB-CGI survey.cgi access"; flags: A+; uricontent:"/survey.cgi"; nocase; reference:bugtraq,1817; reference:cve,CVE-1999-0936; classtype:attempted-recon; sid:871; rev:2;)
```

These alerts only trigger when certain strings occur in a URL. These types of alerts are known for false positive alerts. Reviewed independently from other devices, we have no way of knowing if these access attempts were associated with other legitimate access and therefore false positives. However, we can infer that this is malicious because it would be highly irregular for these three scripts to be accessed by back-to-back-to-back connection attempts as indicated by the ephemeral source ports. Even so, we do not know if the attempts were successful or unsuccessful. All we know is that the attempts occurred.

In summary for the IDS logs:

Who: 152.63.146.6

What: Three http connects to xxx.yyy.zzz.4 with access attempts of cgi scripts: phf, formmail, and survey.cgi. Unknown if the scripts were accessed. There is no other log information available to evaluate if this is a false positive other than the fact these three scripts are unlikely to be accessed within this short of a time frame with incrementing ephemeral source ports.

When: June 1 at 06:08:50

Where: Company DMZ network

Why: If these vulnerable scripts are in operation on the web server, they would allow remote command execution by an attacker.

Web Server Logs (Apache):

access_log

```
152.63.146.6 - - [01/Jun/2002:06:08:50 -0400] "GET /cgi-bin/phf HTTP/1.0" 404 304 "-" "Lynx/2.8.5dev.2 libwww-FM/2.14 SSL-MM/1.4.1 OpenSSL/0.9.6a"
```

```
152.63.146.6 - - [01/Jun/2002:06:08:54 -0400] "GET /cgi-bin/formmail HTTP/1.0" 404 309 "-" "Lynx/2.8.5dev.2 libwww-FM/2.58 SSL-MM/1.4.1 OpenSSL/0.9.6a"
```

```
152.63.146.6 - - [01/Jun/2002:06:08:58 -0400] "GET /cgi-bin/survey.cgi HTTP/1.0" 404 311 "-" "Lynx/2.8.5dev.2 libwww-FM/2.14 SSL-MM/1.4.1 OpenSSL/0.9.6a"
```

error_log

```
[Sat Jun 1 06:08:50 2002] [error] [client 152.63.146.6] script not found or unable to
stat: /var/www/cgi-bin/phf
```

```
[Sat Jun 1 06:08:54 2002] [error] [client 152.63.146.6] script not found or unable to
stat: /var/www/cgi-bin/formmail
```

```
[Sat Jun 1 06:08:58 2002] [error] [client 152.63.146.6] script not found or unable to
stat: /var/www/cgi-bin/survey.cgi
```

For source address 152.63.146.6, we find log entries in both the access_log and error_log. The access log shows that on June 1st at 06:08, attempts to access cgi scripts phf, formmail, and survey.cgi in the cgi-bin subdirectory occurred. The error_log shows that this activity generated errors because these scripts were not in operation on this server. There were no other http connections from this host so it appears as though this is malicious activity not associated with any other normal web traffic.

In summary for the web server logs:

Who: 152.63.146.6. Likely a Unix/Linux host using Lynx v2.8.5dev.2 as the tool to conduct the activity.

What: Three http connects to xxx.yyy.zzz.4 with access attempts of cgi scripts: phf, formmail, and survey.cgi. The scripts were not accessed because they could not be found on the server. There were no other http connections from this host so it appears as though this is malicious activity not associated with any other normal web traffic. No other access log entries for 152.63.146.6 suggests this activity is malicious.

When: June 1 at 06:08:50

Where: Company DMZ network

Correlative Analysis

We demonstrated above that attempting to understand the full scope of a security incident is encumbered if logs and alerts from only a single device are analyzed. Each device has its own limits as to what it can tell us in the analysis process. However, collectively analyzed, the picture becomes much clearer. Let's take a look at what we can determine.

There are two separate episodes of activity that comprise the total picture of the security incident perpetrated by 152.63.146.6.

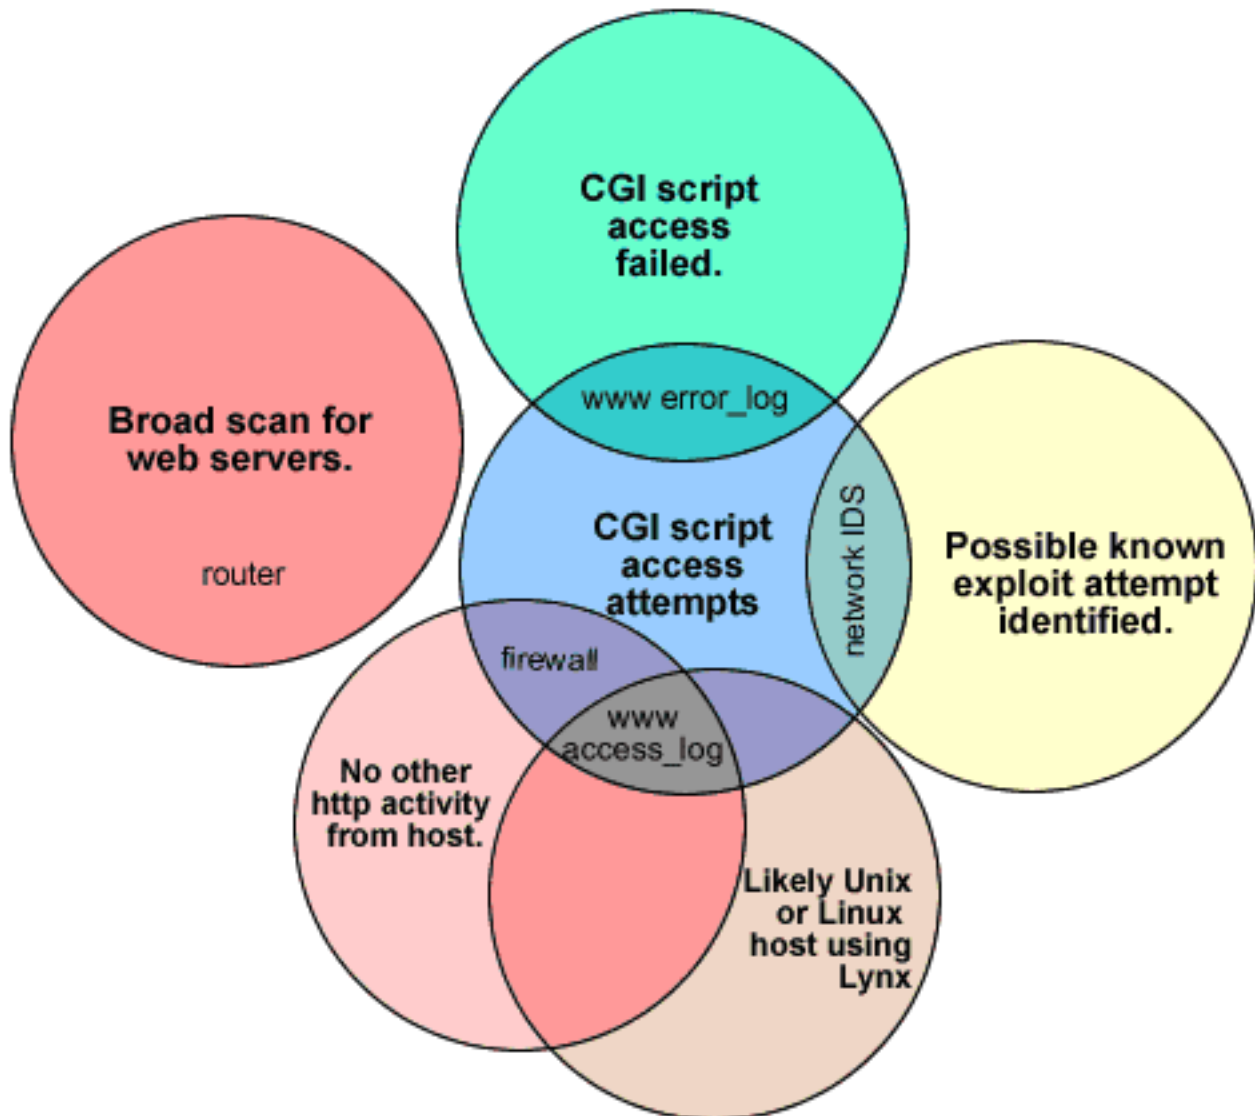
1. On May 31st, host 152.63.146.6 conducted a broad scan of the xxx.yyy.zzz.0/24 network likely in search of web servers (confirmed by router). The interior devices would not have seen this activity because of strong access control lists on the router. The xxx.yyy.zzz.0/24 network is the only public IP address space assigned to the company so it uncertain if this scanning activity is targeted at this company. The logs could be shared with third parties such as via the incidents.org mailing

list. This effort can reveal if this activity has been seen by others or if it is perhaps specifically targeted at the company.

2. On June 1st, host 152.63.146.6 attempted three distinct, and only three, http access attempts against the company web server xxx.yyy.zzz.4 (confirmed by the firewall and web server access_log).
 - a. These connection attempts requested the phf, formmail, and survey.cgi CGI scripts (confirmed by firewall, web server access_log, and network IDS).
 - b. These connection attempts failed (confirmed by web server error_log). It is therefore unlikely that a system compromise has occurred on the company web server.

The following Ven Diagram depicts how the individual network devices contributed to the overall situational awareness achieved through correlative analysis. It shows that removing the analysis of even just one of the device's log data, our understanding of the incident can drop dramatically.

Security Incident Ven Diagram



The diagram shows that removing the analysis of even just one of the device's log data, our understanding of the incident can drop dramatically. For example, if we remove the analysis of the web server error_log, we would not have known that the script access attempt failed. If we had not analyzed the router, we would not have known the probing host scanned the entire class C of addresses for web servers. If we had not analyzed the www access_log, we would not have known that the probing host was likely using Lynx as the web browser to check for the scripts. If we had not analyzed the network IDS logs, we may not have known that the activity was related to well known exploit attempts.

Conclusion

Analyzing a single device to in an attempt to conduct intrusion analysis is the "computerized version of tunnel vision" . Security events must be analyzed from as many sources as possible in order to assess threat and formulate appropriate response. Extraordinary levels of security awareness can be attained in an organization's network by simply listening to what its devices are telling you. This concept was demonstrated by examining how security events reviewed independently only paint part of the picture. However, when the correlation of event data across platforms occurs, a more clear understanding of the scope of security incidents is attained.

References

Curtin, Matt and Marcus Ranum. "Firewalls FAQ" URL: <http://www.faqs.org/faqs/firewalls-faq/> (May 31, 2002).

"Interpret Syslog and Console Messages Generated by Context-Based Access Control." The Cisco IOS Firewall Feature Set and Context-Based Access Control. URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113t/113t_3/firewall.htm (May 31, 2002).

"Log Files." Apache HTTP Server Version 2.0. URL: <http://httpd.apache.org/docs-2.0/logs.html> (June 1, 2002).

Northcutt, Steven. "Coordinated Attacks." IDS Signatures and Analysis-GCIA Courseware. 2001

Power, Richard. 2002 CSI/FBI Computer Crime and Security Survey. Vol. III, No. 1, Spring 2002.

Snort Users Manual - Snort Release: 1.9.x. URL: http://www.snort.org/docs/writing_rules/index.html (June 1, 2002).

Steven Drew



[< Previous Question](#) | [Back to Intrusion Detection FAQ Home](#) | [Next Question >](#)

© 2002-2004 The SANS Institute

SANS Web Privacy Policy: www.sans.org/privacy.php

Web Contact: webmaster@sans.org

SANS Press Room: www.sans.org/press

[Link to SANS](#)

[Feedback](#)

[< back to SANS Home](#) | [< Portal Home](#)

[printer friendly version >](#)