



Top 5 Essential Log Reports

Version 1.0

Contributors:

Chris Brenton - Independent Security Consultant - chris@chrisbrenton.org

Tina Bird, Security Architect, PGP Corporation

Marcus J Ranum, CSO, Tenable Network Security, Inc.

Introduction

In June of 2000, the "SANS/FBI Top 10 Critical Vulnerabilities" consensus list was created. This list identified the ten most frequently exploited vulnerabilities on the Internet. While the list was not intended to be a complete list of all possible threat models, it was an extremely useful action item list for network, system and security administrators alike. By securing the listed ten items, the administrator would receive the greatest increase in overall security and thus the greatest reduction in security risk from hostile attacks.

In the spirit of this original consensus, the SANS community has again banded together in order to create the "Top 5 Essential Log Reports" consensus. This list is not intended to be a complete review of all the potentially useful log reports. Rather, the focus is on identifying the five most critical log reports for a wide cross-section of the security community. These are the top reports which should be reviewed on a regular basis. The goal is to include reports that have the highest likelihood of identifying suspect activity, while generating the lowest number of false positive report entries. The log reports may not always clearly indicate the extent of an intrusion, but will at least give sufficient information to the appropriate administrator that suspect activity has been detected and requires further investigation.

The Top 5 Essential Log Reports

- 1) **Attempts to Gain Access through Existing Accounts**
- 2) **Failed File or Resource Access Attempts**
- 3) **Unauthorized Changes to Users, Groups and Services**
- 4) **Systems Most Vulnerable to Attack**
- 5) **Suspicious or Unauthorized Network Traffic Patterns**

#1 - Attempts to Gain Access through Existing Accounts

Failed authentication attempts can be an indication of a malicious user or process attempting to gain network access by performing password guessing. It can also be an indication that a local user account is attempting to gain a higher level of permissions to a system.

Why It's Important

Password guessing tools are extremely popular in the wild. There are tools available to assist attackers in breaking in through Windows, Linux/UNIX, IPSec, HTTP, SSH, and a host of other authentication mechanisms. If a brute force attempt can be detected in its early stages, you may be able to take corrective action to prevent an attacker from gaining system access. If the attacker does break in, recovery becomes far more difficult.

Report Description

A useful failed authentication report will give an indication of the source IP address of the attempts, as well as the login names used. A summary report is acceptable, so something similar to the following would be considered useful:

Failed Login Attempts:

jsmith from 1.2.3.4 against example-host performed 37 times

jdoe from 1.2.3.4 against example-host performed 16 times

Note that a useful bonus for this report would be an additional line item which indicates if the login name/source IP was eventually successful in gaining access to the system. It's acceptable however for the system to provide this information through a secondary report. For example the ability to report all authentication attempts sorted by login name and/or source IP in a time linear fashion.

Who Can Use This Report

This report is most useful to system, VPN and wireless administrators as they are typically charged with providing network access to resources. Webmasters may also find this report useful, as many times certain areas of a Web site will be secured using user level permissions. This report may also be useful for proxy administrators as failed authentication attempts may be an indication of external access, or local users attempting to gain access to restricted portions of the Internet.

False Positives

Clearly it is possible for a legitimate user to occasionally type in an incorrect password. If report information is provided in summary format, the amount of false positive "noise" should be minimal. The ability to pull additional authentication reports could also help an administrator clarify the situation. For example, multiple authentication failures from the same IP subnet a user commonly originates from are probably not a problem. If the source IP is completely foreign, further investigation is warranted. False positives could also be reduced by the ability to set thresholds within reports. For example it would be helpful if an administrator could define that they are only interested in seeing source IPs where "N" number of authentication failures occurred over "Y" period of time.

#2 - Failed File or Resource Access Attempts

Failed file or resource access attempts is a broad category which can impact many different job descriptions. In short, failed access attempts are an indication that someone is attempting to gain access to either a non-existent resource, or a resource to which they have not been granted the correct permissions.

Why It's important

Failed access attempts can be an early indication of an attacker probing a system. For example if you see an IP address on the Internet looking for multiple files on your Web server that do not exist, they are probably running some form of a vulnerability scanner against you in order to find your weak spots. Failed recursion attempts could be a sign that an attacker is attempting a cache poisoning attack against your network. Very rarely does an attacker get it right on their first try. By monitoring their failures you can get an early warning that your systems are being targeted.

Report Description

A proper failed access attempt report will give an indication of the resource to which access was attempted, the source IP performing the access attempt, along with account information where applicable. Again, a summary report format can be used to provide concise information as well as make it easier for an administrator to spot trends in the data. It is appropriate to consolidate information as required in order to reduce the line items generated.

This is a wide category, so clearly it is beyond this document to list all possible reporting options. The following is a sample broken down by discipline:

Name Server Failed Access Attempts:

Failed zone transfer from 1.2.3.4 for example.net against ns1 performed 3 times

Failed recursion attempt from 1.2.3.4 against ns1 performed 12 times

Web Server Failed Access Attempts:

Failed file access attempt for 1.2.3.4 performed 13 times

- /var/www/html/mambo
- /var/www/html/cvs
- /var/www/html/articles
- /var/www/html/cvs
- /var/www/html/xmlrpc.php
- /var/www/html/blog
- /var/www/html/blog
- /var/www/html/blogs
- /var/www/html/drupal
- /var/www/html/phpgroupware
- /var/www/html/wordpress
- /var/www/html/xmlrpc
- /var/www/html/xmlsrv

File Server Failed Access Attempts:

Failed write access on financial.xls for user jsmith from 1.2.3.4 performed 2 times

Who Can Use This Report

As mentioned, failed access attempts can clearly be useful for the greatest number of job descriptions. Some potential examples:

Name Server Administrator - Failed recursion and zone transfer attempts. Respectively these can be an indication of a cache poisoning attack or an attacker attempting to enumerate information about your network.

Web Server Administrator - Failed file and directory access attempts could be an indication of an attacker probing for possible vulnerabilities.

Mail Server Administrator - Failed mail relay attempts could be an indication that someone is attempting to use your mail system as a spam relay.

File Server Administrator - Failed file access attempts could be an indication that a user is attempting to access information to which they have not been granted permission.

Firewall Administrator - Failed inbound and outbound access attempts. Failed inbound attempts are an indication of port scans or probing. Failed outbound attempts could be an indication that an internal system has been compromised or an internal user is attempting to access non-authorized services.

False Positives

Unfortunately there are many automated processes that can cause false positive conditions when monitoring for failed access attempts. For example load balancers have been known to generate unsolicited zone transfer or recursion attempts. Web search engines typically look for a "robot.txt" file which can trigger a failed access attempt if the file does not exist. A reporting option which permits the administrator to define exception criteria for items they do not wish to see in a report can be extremely useful. For example, an administrator could choose to ignore zone transfer attempts from certain IPs or choose not to see robot.txt in the summary reports.

#3 - Unauthorized Changes to Users, Groups and Services

The modification of user and group accounts, as well as system services, can be an indication that a system has become compromised. While clearly, modifications to all three will occur legitimately in an evolving network, they warrant special attention because they can be a final indication that all other defenses have been breached and an intrusion has occurred.

Why It's Important

When a system becomes compromised, it's not uncommon for the attacker to create a user account with a high level of permissions so they can come back and access the system whenever they wish. Monitoring account modifications will give you an indication that this has occurred. Also, it's not uncommon for a compromised system to have changes made to running processes. For example the attacker may shutdown the systems firewall and AV software, or may launch new processes to help hide their system access. Monitoring when existing services are disabled or when new services are added will give you a warning indication that the system has become compromised.

Report Description

A useful report in this category will summarize results by authentication system or host, as applicable. A summary report which identifies the modification, the affected system, as well as the source of the change should be included when ever possible. Since users/groups and services are sometimes managed by different resources, it is acceptable to have their results summarized in different portions of the report.

Account changes for FS1:

New user: name=c0rt3z uid=1050

Group change: User c0rt3z added to group Administrator

Service changes for FS1:

antivirus.exe has been stopped

evilbackdoor started

sshd restarted

Who Can Use This Report

Tracking modifications to users, groups and services can clearly be useful information for a wide range of job descriptions. For example while system level administrators can clearly use this information, network administrators may find this information useful as well in managing the security of network devices.

False Positives

Obviously any legitimate modification to users, groups and services will show up in this report as well. There is simply no easy way to avoid these false positives. This means that distinguishing between legitimate and unauthorized changes will need to be performed through some other means. For example, if the network in question has a trouble ticket system, the changes could be matched against ticket system entries to ensure all changes are legitimate. If the installation of software is controlled, the appearance of new services can be checked against that system in order to ensure they are appropriate for that target host.

#4 - Systems Most Vulnerable to Attack

As indicated in the original SANS top 10 Critical Vulnerabilities list, as well as the current top 20, one of the most important steps you can take in securing your network is to stay up-to-date on patches. In an ideal world all systems would remain completely up-to-date on the latest patches; time management, legacy software, availability of resources, etc., can result in a less than ideal posture. A report that identifies the level of compliance of each network resource can be extremely helpful in setting priorities.

Why It's Important

A majority of worms on the Internet that have been responsible for wide scale damage have targeted known security problems to which a patch has existed. By ensuring that your systems are patched and up-to-date, you can protect yourself against a majority of the attacks experienced by networks attached to the Internet.

Report Description

As the item name suggests, a proper report will help an environment identify which systems are in the greatest need of patching. Obviously this report would typically be drawn from a vulnerability assessment tool rather than the results of any system level logs. Back end log analysis tools could provide additional levels of customization to these reports. For example, most environments would find it extremely beneficial to be able to customize the final weighting system. Examples would be customization based on a per vulnerability, as well as a per system basis.

Weight per vulnerability - Most vulnerability assessment tools include a weighting system based on the generally accepted risk level of a given exploit. Given that the needs of every environment are different, some level of customization control would be useful. A numeric system (say 1 is low risk while 10 is the highest) would be the most flexible.

Weight per system - Consider this to be a multiplying factor for the above vulnerability rating. In a given environment some systems are deemed to be more critical than others. An analysis system that reflects this relevance could be extremely useful in setting priorities. For example, an exposed e-commerce server might have a weight of "2" while an internal test server might have a weight of ".5". If both systems have identical patches missing, the exposed e-commerce server would score four times higher than the test server, thus indicating it is in greater need of being patched.

Who Can Use This Report

This report is useful to anyone responsible for managing system and network resources. It would also be useful to auditors as well as other who are responsible for identifying level of compliance. Trends of this data would also be useful to upper management. For example, the weighted values could be tracked over time in order to identify improvements or degradation in the current process of maintaining patch levels.

False Positives

False positives in the report would be more of a function of the software used to collect the data rather than being cause by the reports themselves. For example, if a vulnerability assessment tool checks the version of the software saved to the hard drive, rather than the version of the software actually running in memory; it's possible that a higher level of compliance could be reported than what actually exists for that system.

#5 - Suspicious or Unauthorized Network Traffic Patterns

Suspect traffic patterns can be described as unusual or unexpected traffic patterns on the local network. This not only includes traffic entering the local network, but leaving the network as well. This report option requires a certain level of familiarity of what is “normal” for the local network. With this in mind, administrators need to be knowledgeable of local traffic patterns in order to make the best use of these reports. With that said, there are some typical traffic patterns that can be considered to be highly suspect in nearly all environments.

Inbound ICMP Host Unreachable Errors (Type 3s)

Inbound ICMP unreachable errors are an indication that an internal host may have attempted to access a host on the Internet that is either off-line, does not exist, or is administratively prohibited. While it's not uncommon to see these errors occasionally generated, a large quantity of type 3 errors can be an indication that an internal host has been compromised and is currently scanning the Internet in an attempting to attack other systems. Type 3s have also been known to be used as a covert communication channel in order to send commands to a zombie system.

Outbound ICMP time Exceeded in Transit Errors (Type 11s)

ICMP time exceeded in transit errors are an indication that the TTL value in a packet has been decremented to one as it has traversed routers on the network. While it is possible routing errors can cause these packets, they are usually an indication that someone is probing the network with a tracing tool such as traceroute, Firewalk, TCPTraceroute, or similar. The latter two tools are the hardest to protect against as they typically target exposed and accessible services.

Unexpected Outbound DMZ Traffic

This traffic pattern requires a bit of an explanation before its ramifications can be fully identified. Consider the typical “third NIC off the firewall” where most environments locate their Internet accessible services. Typically this network will contain systems with known communication patterns like a Web server (inbound TCP/80 and possibly TCP/443), a name server (inbound and outbound port 53), as well as an SMTP server (inbound and outbound TCP/25). Since these traffic patterns are a known quantity, anything that falls outside of these traffic patterns should be highly suspect. For example, a Web server generating outbound TFTP sessions or a mail server generating outbound FTP sessions should be considered to be a critical situation and investigated. Both could be an indication of a compromised system where the attacker is attempting to pull down a tool kit.

Outbound TCP/25 from a non-SMTP server

It's not uncommon for attackers to turn a compromised system into Spam relays. With this in mind, monitoring outbound TCP/25 activity from all systems but your legitimate SMTP servers is an excellent way of catching these systems.

Outbound Internet Relay Chat (6660-6669, 7000, others)

Many attackers leverage the anonymity and functionality of IRC to control their zombies. This means that connection attempts to the above listed ports should be considered to be highly suspect. While IRC can be run on any port, monitoring ports 6660-6669 and 7000 will at least catch the zombies with a standard configuration.

Bandwidth Utilization

From a security standpoint, bandwidth utilization reports tend to be overlooked. They can, however, provide insight into suspect network activity. For example a desktop system which suddenly begins transmitting a high level of traffic to the Internet could be the result of a system compromise. A sudden spike in Web server access attempts could be the result of a DoS attack or a compromised system offering up Warez. Some potentially useful reports to consider:

- The top "N" transmitting/receiving system
- The top "N" busiest communication sessions
- The top "N" communication protocols

Why It's Important

Attackers have become far more skilled in compromising a system in such a way that the attack can not be detected on the system itself. With this in mind, your best bet for identifying hosts that have been compromised to this level is monitoring their network traffic for suspicious patterns.

Report Description

Report formats will vary based on the information being presented but as a general rule summary information is preferred. The source and destination IP addresses, internal host names, protocol, etc. should all be listed. For example:

Dropped Traffic From DMZ:

smtp1 outbound to 1.2.3.4 on TCP/80 performed 7 times

Suspicious Outbound Traffic:

accounting1 outbound to 1.2.3.4 on TCP/25 performed 1 time

accounting1 outbound to 1.2.3.5 on TCP/25 performed 1 time

accounting1 outbound to 1.2.3.6 on TCP/25 performed 1 time

Who Can Use This Report

While these reports would clearly be useful to firewall and IDS administrators, they can also be beneficial to system and network administrators as well. Administrators responsible for systems generating suspect traffic could use the reports to identify when a complete system audit is required.

False Positives

Because these reports require the administrator to be knowledgeable about normal traffic patterns on their network, the level of false positives will be a direct result of the tuning of the summary reports. Proper summary information can help reduce the amount of work required to distinguish between abnormal traffic patterns and patterns that are the result of an evolving network.

Contact: info@sans.org