Global Information Assurance Certification Paper

SANS Security Essentials Course (GSEC)
GSEC Practical Requirement (V.1.3)

The Increase in Vulnerability/Incident Alerts and its Effect on System Support Group Resources

Name: **Neale Rankin**

Course Certification **GSEC**

Assignment Version **V.1.3**

Submission: **Resend**

Course Location **Online**

ABSTRACT

Since the beginning of the 1990s with the explosion in the use of the Internet , there has also been a explosion in the number of vulnerability's and incident alerts. This has resulted in managers have to find additional resources in already stretched support teams to review and manage these alerts.

This paper discusses the increase in the number of alerts, a process for handling these alerts using the recent SNMP vulnerability as an example, and looks briefly at one effort to reduce this burden on support manager's resources.

Table of Contents

Table of Tables

Introduction

I have been in the computing field since the mid 1980's. Since then I have seen a trickle of vulnerability's and alerts in the early 1990's change to what can only be described as a flood in 2001/2002. Almost daily we are notified of another vulnerability in Operating Systems or applications. On top this, there are also the incident reports. Have our software programmers become that bad?

In this paper, I will discuss this trend and the impact it can have. Topics in this discussion will include:

- Why is this a Problem,
- Sources of Vulnerability Alerts,
- Sources of Incident Alerts,
- Growth Statistics,
- What do you do when you get an alert, and
- What can be done about the problem.

Why Is This A Problem

Each alert inevitably requires some form of remedial action such as patching of the operating system. This volume can also lead to complacency that can be more dangerous to the integrity of the network in the long term, than the original vulnerability. A manager who is seeing a considerable amount of their work-force hours being used to continually patch systems from potential threats, may decide that it is a waste of time, and decide not to allocate resources to review alerts. This could lead to a major compromise of the systems.

The volume of alerts can easily overwhelm a support team that looks after systems. The following questions need to be asked:

1. Which ones do you action?
2. Who investigates it?
3. Do you patch each time or not?
4. Who is responsible to make that decision?

In the UNIX environment, the number of types of operating systems and the number of vendors involved also complicates this. On top of that, there are also the applications and tools, many of which are Open Source, which run on these servers. For an outsourcing company with 300 to 400 servers maintained by 25-30 Unix Administrators this could fast become a nightmare.

In many instances there is not a dedicated person to handle the Unix security issues. This is left to the company information security team, which may not have experienced UNIX administrators as members of the team, or it is assumed to be done by the UNIX support team.

Sources of Vulnerability Alerts

For a person new to the security side of system administration, the hardest thing is to determine where to get their security information such as vulnerability alerts or incident reports from. There are many WEB sites offering "*expert information*", but which ones do you trust?

The WEB site, which is generally accepted as a good starting point is the CERT Coordination Center (CERT/CC) site operated by Carnegie Mellon University.

The purpose of the CERT/CC is:

> "We study Internet security vulnerabilities, handle computer security incidents, publish security alerts, research long-term changes in networked systems, and develop information and training to help you improve security at your site."[1]

This site is considered an authoritative source of information by security and system administrative staff.

In Australia, we have the AUSCert (Australian Computer Emergency Response Team). The purpose of AUSTCert is to:

> "provide a single, trusted point of contact in Australia for the Internet community to deal with computer security incidents and their prevention. AusCERT's aims are to reduce the probability of successful attack, to reduce the direct costs of security to organisations and lower the risk of consequential damage."[2]

These 2 groups provide a good starting point for understanding what vulnerabilities exist. Their WWW sites offers a good range of information to assist in understanding the vulnerability as well as other useful information.

The draw back of relying on these sites is that you do not get the detailed vendor specific information. For this, you need to go to the vendor WWW sites. This is where it starts to get more complex. For a UNIX team, this may mean another 4 or more sites that have to be monitored. Refer to table 1 below.

Vendor	Operating System	WWW Site
Sun	Solaris	http://sunsolve.sun.com
HP	HP-UX	http://www.itresourcecenter.hp.com
IBM	AIX	www.ibm.com
Compaq	Tru-64	www.compaq.com

Table 1 - Common UNIX Vendor WWW Sites

The vendors, in many cases, will not advise of a vulnerability unless it has been exploited or is likely to be exploited. It is not in their best interests to. None of the vendors listed above have direct links from their home page to any security related problems with

[1] CERT/CC WEB Site www.cert.org
[2] The Australian Computer Emergency Response Team WEB Site www.auscert.org.au

their software.

Now if you add in the products that run on these servers, such as databases (eg Oracle), WEB servers, Fire Walls and the problem gets larger. Then there are the Open Source products such as SAMBA, OpenSSH or Apache.

Why do you need to worry about them? The answer to this is simple, the security of the server is only as good as its weakest link. Privileged access can be gained via the Operating System or via an application. The attacker will exploit the weakest link.

At this stage we have not looked at incidents and already there are about 10 WWW sites.

Sources of Incident Information

This is where you have to be very selective. You should find a number of sources you trust and stay with them. I recommend using the following resources:

- www.cert.org
- www.auscert.org.au
- www.incidents.org which is part of the SANS institute,
- SANS Security Digest, and
- www.iss.net/security_center/alerts/alerts.php

I have found that these sites are more reliable and trustworthy.

One source of information that should not be trusted is e-mail from a non-trusted source or friend. An e-mail alert from an external source, should only be used for informational purposes only. Never act on e-mail alone, unless it is coming from your company's internal security team..

Growth Statistics

To see the growth in the number alerts we need only to look at the statistics from the CERT/CC. Table 2 shows the growth in the number of incidents report:

Year	Incidents	Vulnerability's	Security Alerts	Security Notes
1988	6		1	
1989	132		7	
1990	252		12	
1991	406		23	
1992	773		21	
1993	1334		19	
1994	2340		17	
1995	2412	171	31	
1996	2573	345	53	
1997	2134	311	50	

1998	3734	262	34	15
1999	9859	417	22	11
2000	21756	1090	26	57
2001	52658	2437	41	341

Table 2 - Statistic from the CERT/CC

(Source: http://www.cert.org/stats/cert_stats.html)

From the statistics from the CERT/CC, it can be seen that since 1998 there has been greater than a doubling of incidents each year and a corresponding approximate doubling of vulnerability's reported. Each vulnerability has to be treated as if it could potentially lead to an incident.

So far this year the CERT/CC has issued 8 advisories to the end of March, compared to 5 by the same time last year. Similar statistics can also be seen from the AUSCERT site as well. Refer to Table 3.

Year	AusCERT Advisories	AusCERT Alerts	External Security Bulletins
1997	29		167
1998	4	3	198
1999	2	6	204
2000	3	11	404
2001	5	20	547

Table 3 - AusCERT Statistics

(Source: http://www.auscert.org.au)

Note here the steady increase in AusCERT Alerts and External Security Bulletins since 1997. If this trend continues, there will be excess of 600 external security bulletins issued this year.

What does this mean to a manager?

To a manager this means lost resource time. For example, assume it takes 5 minutes to review each external security bulletin to decide if they are relevant to the organisation or not and there are 600 this year. That is 50 hours lost. Lets assume that only a quarter (150) required more detailed investigation and that this takes 1 hour for each one. That takes it to 200 hours lost already.

Out of these 150, lets assume that only 20% (or 30 bulletins) need actioning. Limiting the time spent reviewing and developing an action plan to a total of 10 resource hours for each one. This adds another 300 hours to the total, bringing it to 500 hours. At this stage, the manager has lost the equivalent of a quarter of a person, and this is before any remedial action has been taken.

From this it is easy to see how some managers can become complacent and ignore

the risk. It would be hard to justify this amount of resource hours if there was not a high level company policy directing managers to perform this task with a well defined process for doing it. This process would also need to have a mechanism to audit results.

The Alert Handling Process

The Policy and Procedure

The biggest hurdle that an organisation has is defining the policy and procedures that are going to be followed and allocating who is responsible to carry out these procedures. This is where management needs to be involved and accept responsibility for the policy/process. With-out management's signoff there is no policy or procedure.

Getting management to accept that there is a need for a policy and a procedure is often the hardest part of this task. In many cases this will only occur after an incident. The manager who signs off the policy must have sufficient authority to enforce it and accept responsibility for enforcing it. The higher the management authority the better.

The policy must define the sources of vulnerability/incident alerts that the organisation is going to investigate. This should include CERT/CC, AUSCERT (of Australian organisations), and the vendors.

The procedure must identify the steps to be carried out and the positions in the organisation that will carry out each step. Try not to identify people as they have a tendency to change, resulting in the procedure being out of date because someone has moved or left the company.

Once the procedure has been approved it must be circulated to all staff. The individuals responsible for each step of the procedure must be made aware of their individual responsibilities and duties. There must also be a number of checks and balances to ensure that the procedure is working correctly.

Framework of a Procedure

The following is a basic framework of a procedure to deal a CERT or a vulnerability alert.

Step 1:	The vulnerability Alert is received
Step 2:	The Alert is reviewed by the person(s) designated to handle them. If the alert requires investigation and is relevant to the organisation, then it is passed on to the relevant Subject Matter Expert (SME). If it is not relevant, a recommendation of no action required is made to the manager responsible.
Step 3:	The SME than investigates the alert using vendor information as well as information from the security SME and other sources.
Step 4:	SME reports findings and makes a recommendation back to the manager. In many cases this may mean patching of the servers.
Step 5:	The manager than reviews the findings and recommendations. If the recommendation is accepted, the manger approves the work to

be carried out.

Step 6: Any remedial work is carried out, verified and recorded as being done. This has to be done with in the organisations configuration control procedures.

Step 7: All records are kept for review in the future if required.

This process will require several resource hours to complete. Initially, the time required will be high until all involved get experienced in the process and get use to the different sources of data available. After a few times through the procedures, the SME's will get better at it, thus reducing the time required. It is also important not to rely on the same people each time. There should always be a replacement trained up ready to go for each SME.

The Procedure In Practice

The best way to see how a procedure works is to apply it to a real situation. The recent SNMP vulnerability alert is a good example to use. This was a very public alert with the news media telecasting it within hours of it been issued as a CERT Advisory. What made this vulnerability worse was that it affected potentially all network devices.

With in hours of it being released, both client and mangers wanted answers. They all wanted to know what affect it was going to have on their servers/equipment. They all wanted to know what action we where going to take, and when it was to be done. I will now apply the above procedure to this vulnerability.

Step 1: Alert is Received

The alert was received as an e-mail from the CERT/CC automatic mailing system. It was issued as CERT Advisory CA-2002-03 on the 12th February 2002. Refer to http://www.cert.org/advisories/CA-2002-03.html for the text of the advisory.

This was also followed by a similar e-mail from may companies on internal security advisory system. Before the end of the day, we had already received inquiries from our clients requesting information on what action we were going to take.

Step 2: Initial Review Of The Alert

I reviewed the CERT advisory and found that it was relevant to our organisation and had the potential to affect all of our 400+ servers that our group looked after. Due to the nature of the alert, its public profile, and the number servers involved, my manager was notified.

Each of the SMEs for the relevant operating systems requested to investigate the impact on their operating system. Due to the profile of the alert, a team consisting of a

senior team leader, the security SME and the operating systems SMEs was set up to coordinate the response.

Step 3: SME Investigation Of The Alert

The investigation was carried out in by 5 SMEs. Four were for the different operating system we looked after, these were SOLARIS, HP-UX, AIX and TRU-64. The security SME investigated/monitored the security sites to determine if the vulnerability was being exploited and if so try and determine the risks to our servers. The operating system SMEs based their recommendations on information from vendor sites.

Investigations also found that the perimeter Firewalls did not have the SNMP package loaded, and they were not allowing any SNMP traffic to pass. Note, the routers and switches where being maintained by a separate network group.

While this was going on, the different security news groups were being monitored to see if there was any sizable increase in SNMP activity that would indicate that the vulnerability was being exploited. It was found that there may be a powerful SNMP attack tool in the computer underground. This information was found on the Internet Security Systems Inc, advisories WEB page, refer to http://www.iss.net/security_center/alerts/advise110.php

At the same time a check of the SecurityFocus mailing list of incidents showed what appeared to be the use of a scanning tool to target the SNMP. Refer to the Security Focus incidents mailing list for the week ending 21/02/2002, at http://online.securityfocus.com/cgi-bin/archive.pl?id=75&threads=0&start=2002-02-15&end=2002-02-21. A check of the port report for port 161 on the DShield WEB site, indicated some increase in activity, but none to the levels that would rate as a concerted exploitation of the vulnerability. This report can be found at the following URL, http://www.dshield.org/port_report.php?port=161. Note, it will only show the last 30 days.

By the end of this stage we had what we considered to be a good picture of our situation. We where confident of the following:

1. Our Firewalls where safe and did not have the SNMP package loaded.
2. Firewall rules were in place not to pass SNMP traffic.
3. There appeared to be a tool that could be used to exploit the vulnerability.
4. Information from Dshield indicated that port 161 was being targetted but not to a alarming level.
5. A number of vendors had released patches, and
6. We had to take remedial action to protect our servers.

Step 4: SME reports findings and makes a recommendation back to the

manager

This information was passed on to the manager with the following recommendations:

1. Where possible all SNMP services were to be disabled.
2. Vendor patches where to be installed on all servers where SNMP was required. Note the servers where SNMP was deactivated would have the patches installed during the next patching cycle, so if the service was reactivated it would be patched.
3. The outer perimeter Firewalls were safe because the SNMP package was not loaded.
4. The Firewalls where configured not to pass SNMP traffic, and
5. The patching of the servers needed to be done but the timing was not critical, e.g. it did not need to be done in the next few days.

Step 5: The manager then reviews the findings and recommendations

The management review agreed with the SME's findings and accepted that patching or disabling of the SNMP services should occur. By this time, a considerable amount of resource time has been spent on the vulnerability alert. An estimate of the time spent to this stage is approximately 30 hours, and that is before remedial action has started.

Step 6: Remedial Work

Remedial work of checking, patching or disabling SNMP started. There were 400+ servers involved, with a large number of clients involved. Each client had to be notified that the server was going to be patched and the reason for doing it. Configuration Control Change requests had to be raised and the servers then had to be patched. Approximate time required was 25 to 30 minutes per server. This results in approximately 150 to 200 hours of work.

Step 7: Recording Of Actions

This is an important step, and needs to be carried each time. The records may be needed by auditors when they are reviewing processes, or may be required as proof that a server was patched if there is a compromise. Worst case, they may be needed to show that you did perform all reasonable actions to protect a clients server, if that server is comprised and the client seeks damages.

Summary of the Alert Handling Procedure

As can be seen above, a vulnerability alert, which involves different operating systems, can result in up to 200 resource hours or more just to handle it. If only 5 vulnerabilities require this amount of resources, that is up to 1000 resource hours, before you look at that other vulnerabilities and incidents that occur. It is no wonder that managers are reluctant to act on these alerts, and it also explains why there are so many servers out there that are not patched.

What can be done about the problem?

There has to be some way of determining what servers need to be patched and what patches need to be applied. All the system administrators in the support groups need is a report that states which server needs to be patched with what patches. Currently this does not exist.

The US Government General Services Administration was expected to award a contract to Science Applications International Corp. to set up a system that will provide notification and available fix to different The US government agencies. (Refer to http://www.fcw.com/fcw/articles/2002/0325/news-patch-03-25-02.asp for the full story.)

This type of system will flow on to the commercial world, thus reducing the work load for the different support groups in a large organisation. It will also reduce the requirement for different groups to evaluate each vulnerability and determine the action to be taken. It will become a corporate directive which manager's will have to enforce.

Summary

Since the early 1990's the spread of computers has also lead to an increase in the number of vulnerabilities and incident alerts. This started out as a trickle and now is a flood. Unfortunately the increase has not, as of yet, resulted in a procedure/technology for large organisations and governments to handle these vulnerabilities and alerts in an efficient cost effective way. The result of this is that more resource hours have to be spent each year on handling them, and managers being reluctant to "waste" resources on the problem.

The number is only going to increase, we need to get smarter on how we handle them.

References

CERT Coordination Center
http://www.cert.org/

Australian Computer Emergency Response Team

http://www.auscert.org.au/

Sun Microsystems Pty Ltd Support WEB Site
http://sunsolve.sun.com

Hewlett-Packard Company Support WEB Site
http://www.itresourcecenter.hp.com

IBM Corporation WEB Site
www.ibm.com

Compaq Corporation WEB Site
www.compaq.com

The Incidents.org (part of the SANS Institute) WEB Site
www.incidents.org

Internet Security Systems, Inc alerts WEB page
www.iss.net/security_center/alerts/alerts.php

CERT Coordination Center Statistics WEB page
http://www.cert.org/stats/cert_stats.html

CERT® Advisory CA-2002-03 Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP)
http://www.cert.org/advisories/CA-2002-03.html

Internet Security Systems, Inc alert advisory
http://www.iss.net/security_center/alerts/advise110.php

Security Focus incidents mailing list for the week ending 21/02/2002,
http://online.securityfocus.com/cgi-bin/archive.pl?id=75&threads=0&start=2002-02-15&end=2002-02-21

Dsheild.org Report on Port 161 activity
http://www.dshield.org/port_report.php?port=161

News Article on US Government General Services Administration contract to Science Applications International Corp for a notification and availability fix system.
http://www.fcw.com/fcw/articles/2002/0325/news-patch-03-25-02.asp