Incident Response

Basic guidelines for handling security-related incidents and network intrusions.

Allaire Security White Papers Series

(Version 1.0)

Abstract

Title	Incident Response
Date	January 8, 2001
Product	None
Target Audience	Network Administrators
Abstract	Maintaining a secure network can be a difficult process considering the vast number of potential security incidents that can occur. This document and other lockdown documents are Allaire's effort toward making this job a little easier.

Allaire Corporation • One Riverside Center • 275 Grove Street • Newton • MA • 02466
www.allaire.com • info@allaire.com • (617) 219-2000 •
security issues: secure@allaire.com
document feedback: lockdown@neohapsis.com

Table of Contents

Initial preparation and organization

Why preparation is essential

Staying on top of basic network and system security is hard enough without handling security incidents and intrusion attempts as well. Unfortunately, administrators are looking into the barrel of a daunting gun. Successful network penetrations are on the rise and any decline in this trend doesn't seem to be on the horizon. Fortunately there are a number of things you can do to prepare your organization, should the dreaded attack strike.

Preparation is the key to successful incident response. Technology only addresses perhaps half of the issues surrounding corporate security – if that. The other half is dependent on successful policies and procedures - efforts that are most effective when implemented *before* an attack occurs. You can employ a fleet of well-trained system administrators implementing the latest in ID (intrusion detection) and firewall technology and still be inviting disaster if your policies and procedures aren't nailed down. Regardless of the size of your organization, there are a few things to do up front.

Develop a security strategy and a set of baseline policies

First, define an AUP (acceptable usage policy) for your users. AUPs help employees understand their roles and responsibilities within the electronic workplace. AUPs also heighten awareness and serve as a preliminary education tool. Next, establish a set of baseline information security policies that delineate risk, data and infrastructure value -- and your organization's stance on security issues in general. Some organizations' policies are defined in great detail, while others take a more general "high-level" approach. Although resources vary, *every* organization should – at the least -- have a minimum set of baseline security policies. Finally, establish an approach to incident handling *before* you need one. A security breach is difficult to deal with in the best circumstances, so you certainly don't want to find yourself stumbling through one with no policies in place.

Some good resources and guidelines for initial policy definition:

- "Information Policies Made Easy" – a 500+ page guide by Charles Cresson Wood. Although a bit overwhelming and heavy-handed in some of its recommendations, it's a good starting point for anyone undertaking the challenges of policy writing. (see http://www.baselinesoft.com/).

- COBIT – "Control Objectives for Information and related Technology." While COBIT goes beyond basic policy definition, it's a good place to start. The

description from the ISACA web-site: “COBIT has been developed as a generally applicable and accepted standard for good IT (Information Technology) security and control practices. It provides a reference framework for management, users and IS audit, control and security practitioners.” (See http://www.isaca.org/cobit.htm)

- BS7799 – The British Standard for the “Code of Practice for Information Security Management.“ While more popular in Europe than in the US, BS7799 is a fine starting point for any organization. (See http://www.bsi-global.com/index.html)

Develop and define roles of responsibility

While defining responsibility can be a bit tricky for smaller organizations, addressing this early on is crucial. If a security breach or incident occurs without clear role definitions in place, chaos is almost always sure to follow. Suggested areas of definition:

- Technical – Who’s responsible for the network? Who handles the systems that run on the network? Who runs the telephone infrastructure? And who’s in charge of the satellite sites? If these roles aren’t clearly defined, you’ll double your trouble in the event of a breach.
- Managerial – Is your CIO concerned about corporate security? Unfortunately, if yours is like most, the answer is no. Upper-level management *must* understand the basic issues surrounding corporate security and defining who will be the decision-makers for “high-risk” security issues is essential.
- Users – Do your users know what they should and shouldn’t be doing? Make sure they’re aware of the company security policies and that they’re familiar with security staff.
- Law Enforcement – Have you made any contact with local law enforcement? Be smart: Make a quick phone call to introduce yourself and find out how things operate in your area. This will help immensely should you need to deal with something in a timely manner.

Prepare your systems and network

There are a number of things you can do that not only prevent security breaches, but aid in the response efforts should one occur.

- Have network diagrams handy and up-to date. These will help your incident response team and if you have to bring in outside help, these

maps will be essential. Regardless of who handles the incident, unless you run a *really* simple network, diagrams are invaluable.

- Stay on top of security holes. Whether you read BUGTRAQ (see http://www.securityfocus.com/), subscribe to the SANS Security Alert Consensus service (see http://www.sans.org/nwcnews/), or simply rely on direct vendor mailings, keeping up with system and service holes is *crucial*.
- Post warning or "banner" messages on all systems. This is important from a legal perspective. Login messages and prompts should warn potential users that their actions are logged, that the systems they use are the property of your organization and that unauthorized use is prohibited. Depending on what laws are enforced in your area, these warnings may heavily influence the outcome of any subsequent prosecution.
- Secure and check logging mechanisms. Auditing should not only be enabled, it should be handled in a secure manner. In addition to basic log enabling, system administrators should routinely review logs as part of their weekly (or daily) tasks. The worst time to find out your logging isn't working is after a security breach.
- Use write-once media for long-term log storage. Should any legal action be taken after an incident, the automated archiving of logs to secured write-once media such as WORM drives will add legal credibility to this potentially valuable form of evidence.

Discovery and initial response

Determine whether or not the incident is truly an intrusion attempt

Once a security-related issue has been identified, begin an initial level of investigation before kicking things into high gear. Many incidents simply appear to be hostile activities, when in reality, they aren't. These are commonly known as "false positives." Before alerting numerous people in the organization about a possible security breach, take the time to determine whether or not the incident in question is *truly* an intrusion attempt or break-in. Frequently, automated logins and network misconfigurations are behind the anomalous activity. Standardized and enforced change management policies can help avoid a fair amount of confusion when dealing with false positives. However, regardless of the cause, treat security-related incidents with care. Until the origin of an alert is determined, assume it is hostile.

Classify the severity of the attack

Once the incident has been identified as hostile, the first step in determining the scope of the intrusion is to classify the attack. Generally, security incidents can be classified as one or more of five types:

1. **Reconnaissance or probing attacks** – These are preliminary information gathering routines (see Allaire's 'Detecting Intrusions Security White Paper'). They usually consist of ping sweeps, port scanning and "banner grabbing." While they don't represent a threat themselves, they often are a sign that someone is "scoping you out" as a possible target.

2. **Denial of Service (DoS) attacks** – These attacks aren't designed to compromise systems, but rather cripple service availability. "Tear Drop," SYN flooding and the "Ping of Death" are all examples of DoS attacks. Denial of Service attacks are often considered a form of network maiming.

3. **User-level compromise** – A "user-level" (often referred to as "local") compromise results when an intruder gains access to a system as an unprivileged user. These compromises can be starting points for launching further, more complex attacks against internal machines. These are often the result of poor passwords, or poor password protection (users unknowingly supplying outsiders with information).

4. **Root or "administrative" compromise** – This type of attack is often the most devastating, because the intruder has full reign over the compromised system(s). In a "root compromise," the attacker gains super-user status, and can essentially access/modify anything on the compromised systems. Many

MS-specific attacks are usually considered “root compromises” since components of IIS and most other services run under the SYSTEM_LOCAL (essentially administrator) context.

5. **Internal data theft or modification** - These attacks are frequently unauthorized actions performed by someone within the organization, rather than technology-based attacks. While the technology side of the response to these activities can be similar to the other four, the procedures and legal ramifications can vary greatly. Because these attacks often fall under the category of industrial espionage/sabotage, they are outside of the scope of this document.

Determine the nature and scope of the incident

Intruders come in all shapes and sizes: disgruntled employees, competitors, malicious Internet users, teenage pranksters, etc. While their backgrounds may differ, the results are often the same: disruptive or unauthorized access to your systems and services. However, the “tell-tale signs” can vary widely (see Allaire’s ‘Detecting Intrusions Security White Paper’). Once you’ve determined that an intrusion has actually occurred, the following steps will help you learn the scope of the intrusion:

- Examine your log files. Have they been modified? Are you missing time chunks on your logs? How many servers have been affected or are missing log entries?
- Be aware of system dependencies. If you have an environment based on type of “trust” (single NT domain model, use of rhosts files, etc.) you must consider system dependencies. For example, if the attacker compromised your NT server at an administrative level (attack type 4) you must assume your passwords have been taken. If that machine is a domain controller, you can safely assume that the attacker has your entire domain password file.
- Look for unusual, hidden or changed files. Attackers will often leave behind exploit code, binary executables or snippets of information they gathered.
- Check the integrity of your system binaries (executables and services). Programs like TripWire (http://www.tripwire.com/) and AIDE (http://www.cs.tut.fi/~rammer/aide.html) are immensely useful for this task. However, they’re only useful if they’ve been run *before* a security breach.
- Check for strange or unknown processes or devices running on your systems. Attackers often place “back-door” programs (Back Orifice, netcat, netbus, Trojan inetds, etc.) as well as sniffers on your network.
- Investigate/check all scheduling services, such as CRON, AT, the NT Scheduler, etc. These are frequently tampered with, and back-door applications are installed.

- Examine user accounts and password files. Frequently attackers will insert new users for re-entry at a later date.
- Examine groups. Attackers often insert users into administrative groups which often go unnoticed.
- Examine all machines on the network that were accessed.
- Note that by initially limiting the use of the root or administrator account you'll reduce your amount of "tail-chasing" after an incident, as well as aiding in the identification process.

Understand that if a root compromise (attack type 4) was achieved, the integrity of *anything* on the compromised system is questionable, at best. Keep this in mind when examining the issues outlined above.

When an intrusion occurs, it's often best to assume the worst. Assume that the attacker has put a Trojan your machine. Assume the logs have been doctored. Assume the attacker has gone as far into the organization and systems as possible with the obtained level of knowledge. If the attacker got at it, assume he or she used it.

Here are a few other questions to consider: Was the attack targeted and specific, or did the attacker just stumble around? Did the attacker break in for a particular file or service? Was any content modified? If nothing seems modified, how confident are you that your data is intact?

Determine the risk in continuing operations

Once an intrusion has been detected and identified, there are primarily two possible courses of action: 1) disconnect the system from the network and begin containment procedures, or 2) continue the investigation with the compromised system(s) still running. There are no rules of thumb to help you in this decision, so weigh your options carefully. If you remove the affected system(s) from the network, you'll lose any services those systems offer along with alerting the attacker to the fact that you've detected the intrusion. You'll also relinquish many of the avenues by which you can track the attacker(s). If the origin or entry method hasn't been determined, weigh the possible loss of corporate assets against the problems of not knowing who the attacker(s) are or how they got in. However, if you have all the information you need about this intrusion (source of origin, method of entry, etc.) remove the system(s) for cleanup if possible.

In *most* cases the method of entry can be determined. Most attacks occurring on the Internet are based on *known* security holes. Also consider liability. While legal precedence has yet to be set in the area of information security best practices, be aware that your organization may someday be held liable for attacks launched from their own networks.

Start a documentation effort

A thorough documentation effort should be implemented the moment it has been determined that an intrusion occurred. Even if your organization has not determined whether they wish to prosecute, it's better to be safe then sorry. LOG YOUR ACTIVITY. This not only includes direct profit loss, loss of service, etc., it also includes employee TIME SPENT ON THE INCIDENT, as well as time spent coordinating with other organizations.

Notify Affected Parties

It's a good idea to notify the administrators of any networks that have been affected. This includes the administrators of the networks where attacks may have originated, as well as networks you might suspect that have been attacked from your site. While no one wants to hear about intruders in their site, most will be thankful in the long run, and cooperation is often one of the keys to successful incident handling.

Control any and all evidence

It is also very important that all evidence be controlled, preferably by a single person. If prosecution does occur, evidence is less credible if it can be proven that multiple people had access to it. Outlining legal considerations before an incident occurs is a very good idea.

Maintain a low profile

Tell only the minimum amount of people about the attack, and try to maintain a low profile. DO NOT, where possible, discuss the attack or investigative efforts over internal or insecure e-mail. Intruders are often able to target system and network administrators, and will often insert mechanisms for reading their mail. You should also note that public relations can be a killer for large organizations - especially those in the financial arenas. Keeping events contained is both a technical and very human challenge.

Containment and elimination

Remove affected systems from the network

Once the attack has been identified and classified, you'll probably want to remove the affected system(s) from your production network. While it's tempting to simply patch the suspected hole, change the passwords and re-deploy the system into the production environment, this can -- and has been -- a very costly mistake. Many attackers will leave behind back-doors, trojaned programs or "root-kits," and other methods of re-entry and further tampering. If any of your systems suffered a root-level compromise (type 4) you should assume that these types of activities occurred. See "Anatomy of a Network Intrusion" (http://www.nwc.com/1021/1021ws1.html) for further examples of post-attack activities. Removing the system for a real cleaning process is preferable over the "quick fix."

Backup the compromised systems and data

There are many reasons to perform a backup, but the primary one is for future analysis. Should you or law-enforcement officials wish to review the state of your system at a later date, you'll have a perfect snapshot of it. Many people skip this step, and later regret it.

Determine the entry point and eliminate the vulnerability

Perhaps one of the most important steps to effective incident handling is the determination and eradication of the point of entry. Without knowing how an attacker got in, it's often difficult to guarantee that they'll be kept out in the future. However, in most cases, effective and efficient patching and permissions checking will keep the majority of system attackers at bay. Attempt to find the exact point of entry – down to the service and system – but don't panic if you can't. There are still some basic techniques you can perform.

Most vulnerabilities and attacker entry points can be fixed by simply applying vendor patches and configuring services properly. Most attacks that occur on the Internet are based on known holes, or misconfigured applications. Rare are the custom, site-specific attacks that target a proprietary application – although these certainly occur. In any case, before re-introducing a system into your production environment you should make sure you've eliminated the vulnerability. Good starting points beyond simple patching are basic lock-down documents which we provide.

Perform vulnerability assessments of systems and networks

While you definitely want to seal up the hole the attacker used for entrance, once a network has been attacked there's a good chance it will be targeted in the future. All too frequently administrators recovering from their first known security breach find themselves fending off an onslaught of subsequent attacks from multiple sources. Word in the "underground" travels a lot faster then most people realize. Because of the threat of future probes and attacks, it's important that organizations perform additional audits on "neighboring" systems and networks – not just the systems or networks directly affected. If someone has attacked your network, chances are they'll be back. You want to preemptively patch any holes before they return.

If your organization does not perform routine vulnerability assessment drills, consider getting this effort up and running. There are a number of tools available to perform "automated" scanning for known security holes. Two of the better ones are ISS' Internet Security Scanner (http://www.iss.net/) and Network Associates' CyberCop Scanner (http://www.nai.com/). There are also many firms you can hire to perform these services for you. (See "Additional Resources" at the end of this guide.)

Rebuild and restore the system

Proceed cautiously when restoring systems that have been compromised. Make sure the system you think is clean isn't actually still contaminated when re-introducing it into your production environment. This mistake creates an even worse nightmare, since the attacker may now be able to gather new information, new password files, etc. You can also bet that most attackers who know they were discovered will be much more cautious and stealthy the second time in.

The only way to verify the integrity of your system and accompanying applications is to compare the compromised system to known and good binary signatures. If you did not deploy a binary integrity checker like TripWire or AIDE, you probably don't have a safe method for determining the integrity of your system files. Running a restore from backup media prior to the intrusion date is an acceptable method of re-building your system if -- and only if – you can definitively and absolutely determine the date of the initial intrusion and are 100 percent confident that there were no security breaches before that date. Of course, you must also patch the holes that created the problem in the first place after the restore procedure. However, if you have any doubts as to the integrity of your backups or services, a simple restore is out of the question. In most cases, you'll want to rebuild your system from vendor supplied installation media and only restore your data from backup media – not your system or applications.

While it may seem ridiculous, many administrators find that after an intrusion and the subsequent rebuild they're much more comfortable with their systems. Take advantage of the systems' downtime to shutdown unneeded services and subsystems, apply all recent and available patches and configure the system the right way. You don't want to do any of this more often then you have to!

Reintroduce the system to the production environment

When your system has been properly cleaned or rebuilt, locked-down, patched, updated and the services have been restored and configured securely then, and only then, is it ready to be reintroduced into your production environment. However, the incident shouldn't be closed just yet. Be sure to stay alert and read on for further recommendations.

Proactive measures

Make sure you change your passwords and keys

If your organization suffered a root-level compromise you should assume the attackers took everything to which they had access. This includes router configuration scripts, VPN keys, password and SAM files – everything. It's a good practice to not only change administrator passwords, but also change router login IDs, any static VPN or encryption keys and force users to change their passwords. Far too often administrators will lock out an intruder or close a back-door account, only to find out that the attacker continues to hop from valid-user account to valid-user account. It's often painful to force such changes, but they'll probably save you time and money in the long run. Remember, if the attackers had access to it, assume they will use it.

Enable advanced logging and auditing

While it's a good idea to actively monitor and log activity on your systems and services, it's especially important to do so after an intrusion or intrusion attempt. Most attackers return to the scene of the crime so make sure you can watch what they do.

While watching over a few NT event logs or UNIX syslogs can be a monotonous task, attempting to manage hundreds of them can prove impossible. Fortunately, there are products that can help. Network-based ID (Intrusion Detection) systems such as ISS' RealSecure (see http://www.iss.net/) and Network Security Wizards' Dragon (see http://www.securitywizards.com/) can help detect both hostile activity and network probes. In addition, host-based systems like Intrusion.com's products (see http://www.intrusion.com/) can not only help consolidate logs, but will help parse them and detect hostile activity as well. Regardless of how you watch your logs, make sure they're monitored to some degree.

Document and discuss the lessons learned

While it may sound morbid, sometimes a security breach is the best motivator. While few people enjoy cleaning up after a successful intrusion, the lessons learned can be invaluable. Take the time necessary to inform those who need informing. Go over the process and understand how events actually occurred. Focus on the positive side of the procedures and note where improvements were needed. Finally, make sure to implement and follow-up on any outstanding

items or fixes. While incident response can be extremely educational, it's usually equally exhausting – make sure you don't need to do it again anytime soon.