



Responding to a Customer's Security Incidents—Part 3: Following Up After an Incident

Vijay Masurkar—Sun Services

Sun BluePrints™ OnLine—September 2003



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
(650) 960-1300

Part No. 817-3733-11
Revision 06, 6/4/04
Edition: September 2003

Copyright 2003 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, California 95045 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun BluePrints, SunDocs, Sun Fire, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the US and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95045 Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd.

Sun, Sun Microsystems, le logo Sun, Sun BluePrints, SunDocs, Sun Fire, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Please
Recycle



Adobe PostScript

Responding to a Customer's Security Incidents—Part 3: Following Up After an Incident

The pressure of working and reacting in Internet time and protecting organizational or personal assets is constantly mounting on all of us. Adversaries are getting more sophisticated as the Internet is maturing. Both Nimda and CodeRed are recent examples of the advanced nature of threats combining software vulnerabilities and spreading of multiple vectors of infection. Security incidents have become widespread and difficult to contain in some situations. Yet, globally, enterprises working with various private and public security organizations, investigative agencies, IT vendors, governments, and academia must keep abreast of current incidents, plan for future incidents, cooperate in a concerted effort, and respond to incidents effectively.

The first article in this series discussed establishing a computer security incident response team (CSIRT) and a security policy. The second article discussed executing the policy. Security Incident Response (SIR) is the combination of resulting processes and actions an organization takes in responding to a security incident. It should be obvious that each and every security incident response program will contain unique elements that exist and make sense only for that organization. Before you read this third article, you should be familiar with the concepts described in the first two articles.

This third article focuses on following up after an incident. In this article, only the salient topics for best practices that can be executed in the follow-up phase are presented. These topics include acquiring incident data, resorting to legal actions when deemed necessary, and post-incident activities, such as taking inventory of the affected assets, assessing the damage, and capturing the lessons learned. The best practices presented in this article are generally preceded by a recovery phase and are only starting points for a more detailed analysis for building a policy with the associated processes and procedures.

This article is intended for computer security managers, security policy developers, system administrators, and other related staff, who are responsible for the creation or operation of a computer security incident response policy and service.

Understanding Key Points of the Follow-Up Phase

Why is *follow-up* so important to have an article dedicated to it? The follow-up, to a large extent, provides the closure on the incident by analyzing it, taking action against the cause of the incident, recording it in detail, and learning from it to improve the processes and procedures that are in place for the organization. Some unexpected or unusual questions might come up:

- If the incident information was received sooner, would the outcome be different?
- What would the staff do differently next time?
- Did management (at the customer site or at the organization providing the service) prove to be part of the problem and/or part of the solution? If yes, how?

On the other hand, a simple people communication issue might surface in the follow-up phase. For instance, a few years ago, in response to a local incident, a European CSIRT contacted the U.S. National level CSIRT (CERT/CC) before contacting its own national CSIRT. This was a procedural or execution-level lapse due to miscommunication among the teams in the same country. In another instance, at the U.S. federally funded agency, CIAC (<http://www.ciac.org/ctac/>), failures were reported in noting telephone numbers and email addresses of those who reported an incident. This procedural gap was spotted and fixed in the follow-up phase meeting.

Social Sciences

From a broad perspective, the integration of social sciences into incident response is critical for forming, applying, and reusing the skills of a CSIRT. The human aspects of social sciences can make or break a case. They involve the victims, the workers at the affected customer's site, the executives, and the perpetrator(s).

Take, for example, *insider attacks*, which are not uncommon. There are three important aspects that a CSIRT and the geo-based security officer must remember about these attacks:

- Everyone at the site is a suspect because the perpetrator is still part of the affected customer's enterprise.
- An insider attack could cause stress to many employees.
- The way the incident is processed will reflect on the reputation of the CSIRT and its parent organization and enterprise.

Peripheral Aspects

During the follow-up phase, several peripheral aspects that are beyond the security incident itself need attention. The investigators employed by the worldwide security team for an incident being handled by a virtual CSIRT (VCSIRT) must consider the possible media exposure, the customer's state of behavior, and reaction to the incident. In addition, attention must be paid to the visibility of the case within law enforcement, the interaction with external attorneys, such as those from the district attorney's office, and the interaction with the suspect's attorney. Lastly, the political aspects of the incident, particularly if it is a high-visibility case, cannot be ignored.

Documentation

Throughout the security incident response (SIR) follow-up phase, the responsible geo-based security officer must ensure that the answers to the following are properly captured in a document: what, when, who, why, and how. The documentation should include a chronology of events that can form a basis for prosecution, if needed, a postmortem analysis, and a lessons learned document so that security policies can be improved. The geo-based security officer should maintain this critical information for possible future use.

Incident Classes

Security incidents do not fall into a single, expected pattern. However, the analysis of an incident must address the scope of the incident very clearly. There are two general classes of incident analysis to consider:

- Intra-incident analysis
- Inter-incident analysis

The most common types of intra-incident analyses involve a specific incident. For instance, the analysis might involve items such as log files, artifacts left by the intruder (such as rootkits), software environments, and web-of-trust (that is, which person trusts which person, and which component trusts which component in a customer's infrastructure within an incident).

Inter-incident analysis involves relationships between incidents. This is aimed at finding symmetries between separate incidents that might indicate equivalent or related sources of intruder activity. For example, in the same week, if there are multiple attacks on the sites of an organization, it makes sense for the investigating VCSIRT to correlate log data from firewall and intrusion detection systems (IDSs) at these sites and to search for similarities between security events. A series of log entries qualifying as an event might contain several attack patterns.

Evidence

One general meaning of evidence that applies to security incident response is testimony—facts in support of something. The legal meaning that is more applicable in the context of this article is information given personally or drawn from documents, tending to establish fact. This explains what is required of the incident data if it needs to be used in a court of law. Without a doubt, security incident data that is gathered as evidence can make or break a case if a customer wants to prosecute the perpetrator. Note that there are options to prosecution. Throughout the investigation, vigilant collection of circumstantial evidence and use of a chain of custody is important. The evidence might be needed in a grand jury hearing or a trial, but remember that the definition of permissible evidence is not the same in every country.

Chain of Custody

The chain of custody is accomplished by having verifiable documentation indicating the sequence of individuals who have handled a piece of evidence and the sequence of locations where it was stored (including dates and times). For a proven chain of custody to occur, the evidence must be accounted for at all times, and the passage of evidence from one party to the next must be fully documented. If your organization has policy for preserving and proving chain of custody, ensure that your actions are in keeping with this policy. In reality, the concept of chain of custody in the context of a crime is well-known to law enforcement personnel, but not to field engineers or administrators of a VCSIRT or a servicing organization at the customer site. In addition, country-level laws differ for the handling of evidence. Thus, training of the legal implications of custody of the evidence collected during an incident must be

provided to the worldwide security team. The aim of a carefully crafted chain of custody is not only to protect the evidence, but also to make it difficult for a defense attorney to find a weakness in the custody process.

Scope

FIGURE 1 shows the scope of activities during the follow-up phase. The acquisition, authentication, preservation, analysis, response plan determination, and post-incident activities occur in almost every case. Legal and investigative activities occur only in specific cases. The VCSIRTs or security officers involved should determine the need for legal and investigative activities, based on the requirements of the customer affected by the incident.

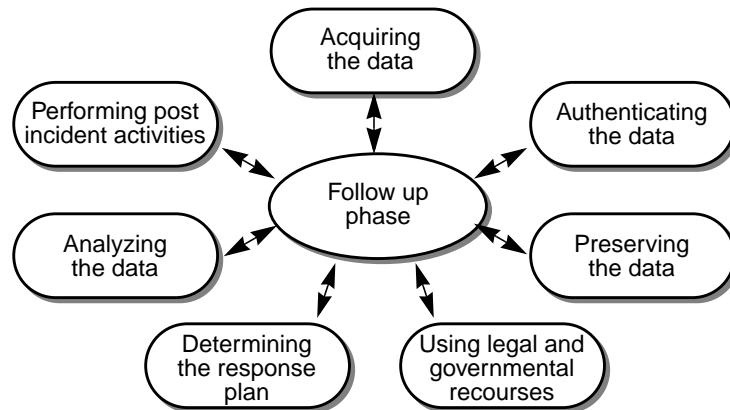


FIGURE 1 Follow Up Activities

The following sections highlight the salient points of the various activities of the follow-up phase, as shown in FIGURE 1.

Acquiring the Evidence

Handling of evidence data is critical when it comes to the first actions taken by the system administrator, support engineer, or field engineer before the geo-based security officer arrives at the affected customer site. Inadequacies in the data processing in storage or transportation could result in the case getting thrown out, if you plan to prosecute the perpetrator.

The organization's geo-based security officer who coordinates the VCSIRT should be jointly responsible with the organization's geo-based customer account manager for gathering the following evidential data:

- All chronological events pertaining to the systems and networks in question with what, when, and how it happened, who was involved from the VCSIRT and from the customer, how the staff received the information, and how it should receive it differently
- Names and contact information of all of the parties involved in every event
- All actions taken by the organization's VCSIRT and its members for the specific incident
- All internal and external conversations

Electronically stored data is a vital source for discovery. But, it is imperative that the VCSIRTs gather electronic data in a manner that ensures the admissibility of the evidence. Whenever necessary, legal advice must be sought by the geo-based security officer responsible for the customer site. While the technology continues to evolve, the basic technique for collecting evidence should remain consistent. The following conceptually lists the general steps for securing a system for data collection and for acquiring the data. The steps do not specify any particular operating system or machine type, so depending on the type of system, the procedure can vary.

1. Secure the customer's physical premises.

This first step is very important to prevent unauthorized personnel from tampering with the compromised system. Also, innocent bystanders could change the environment inadvertently. You should take photographs of any papers, disks, or peripheral devices in the area and collect all of the items in the vicinity that might contain evidence.

2. Confiscate and/or shut down the system or network, as necessary.

You must carefully evaluate the system, its network, and surrounding physical environment before shutting it down. Do not use the keyboard. Use the power button or other methods for graceful shutdown. The procedures you use must save programs that are running in memory onto the system's hard disk before you shut it down. Do not shut down the operating system because it might trigger logic bombs that are sometimes designed to destroy evidence in virtual memory.

Note – Logic bombs are programmed threats that are dormant in commonly used software for an extended period of time until they are triggered. At that point, they perform the function of the program in which they are contained.

3. Secure the system and network.

If you are seizing the computer or a network of computers, it should be sealed away from the rest of the environment before moving it into the examination area. Label all cables and connectors marking source and end points with evidence tape.

4. Examine the affected system and/or network.

The important action here is to check the current date and time and compare it to the known standard. Make a note of the difference, if any. This could be useful in correlating file timestamps to other incident data gathered as evidence.

5. Prepare the system for data acquisition.

The steps for preparing the system might involve changing the boot sequence, if necessary (for example, an alternative, forensic-clean drive and then the system's original hard drive).

6. Connect the target clean media with forensic tools.

Place a forensically clean drive into the system to use as a target drive with a known good operating system kernel and system level binaries and forensic tools.

7. Securely copy the media of the affected system to the clean system.

Boot the system again using the forensic disk. Then use the known good software on the forensic disk to copy the image of the original drive to the target drive. You must create and compare a MD5 cryptographic checksum of the source and the copied files.

8. Secure the evidence on the media.

Remove all of the drives from the system, and seal them with evidence tape in antistatic bags. Date and sign the evidence tape, and secure the drives in a locked container.

TABLE 1 summarizes the media list. The VCSIRT should take into account all possible data storage devices on the clients and servers as well as residual data in any hardware that could help in the follow-up. The following table does not address incident-related data that exists in various system, network, and manual logs.

TABLE 1 Making a Media List for Incident Related Data Collection

Data Storage Type	Possible Media Device Locations
Data files	Office desktop computer or workstation, home computer, notebook computer, and palmtop devices
Backup tapes and disks	System-wide backups (daily, incremental, weekly, and monthly), disaster recovery backups, offsite backups, and personal or <i>ad hoc</i> backups (for example, on diskettes and other portable media)
Other media types	Tape archives, floppy diskettes, CDs, Zip cartridges, and replaced, removed, or discarded hard drives

Making a complete and accurate copy of all of the data on the compromised system's drive is key in capturing the entire image (that is, an image copy or sector-by-sector copy). Otherwise, the acquired data could be deemed inadequate. For instance, in *Gates Rubber Company v. Bando Chemical Industries, Ltd.*, a U.S. court criticized a party's electronic expert for not making an image copy, concluding that when collecting evidence for judicial purposes, a party has "a duty to utilize the method which would yield the most complete and accurate results" (see *Gates Rubber Co. v. Bando Chemical Industries, Ltd.*, 167 F.R.D. 90, 112, D. Colo., 1996).

UNIX File System Imaging

After creating a list of all of the file systems on the suspect host, begin capturing them in a secure manner. There are a number of ways to capture the file systems, depending on the type of the suspect system and the device you have for storage.

Although `cat(1)` and `cp(1)` work well, the `dd(1M)` utility has several advantages on UNIX or Linux, in spite of the fact that it is slower. First, it can copy data using a specified block size that is suitable for specific hardware devices, and it reports how many blocks it processed. Second, it can keep an error log as a proof of your successful backup by using the following example command:

```
dd if=/dev/fd0 of=/home/does/test/ 2>/home/blow/error.txt
```

You can use the `dd(1M)` utility to clean up the analysis drive (that is, to ensure that no residue exists), and the `dd(1M)` utility can be used to copy chunks of data that will not fit on the analysis drive. For example, if you have an image that you need to

chop up into smaller pieces because the backup media is limited to four 1-gigabyte disks and the evidence contains 4 gigabytes of data, you could use `dd(1M)` with specific flags to create four images of the evidence, each 1 gigabyte in size, as in the following example:

```
dd if=/dev/st0 count=1000000 of=/dev/case10img1
dd if=/dev/st0 count=1000000 skip=1000000 of=/dev/case10img2
dd if=/dev/st0 count=1000000 skip=2000000 of=/dev/case10img3
dd if=/dev/st0 count=1000000 skip=3000000 of=/dev/case10img4
```

Finally, the `dd(1M)` utility can copy deleted blocks from the compromised drive. After copying, always compare an MD5 hash of the compromised device to that of the copied device. For further details on using this command, refer to the `dd(1M)` man page.

Copying a file system over a network can be slow, but it is convenient when you cannot remove the hard drive and have no tape drive. Set up Netcat on the listening (that is, target) host, then run `dd(1M)` on the suspect (that is, compromised) host, as in the following example:

```
Listening host: # nc -l -p <port#> > suspect.disk1.image
Suspect host: # dd bs=1024 < /dev/disk1 | nc <ip_address> <port#> -w 3
```

Note that using Netcat to transfer a disk image in the clear exposes the contents to anyone tracing the network. You should use the encryption option when necessary.

Data Discovery

For data discovery, you must follow a method. As a practical matter, finding the information stored on computers is becoming routine; however, it is an important part of the VCSIRT-employed attorney's discovery efforts. Typically, the data discovery items in TABLE 2 are required for an incident response team to gather data pertaining to legal proceedings. The following list should not be considered all

inclusive. For example, meta data, such as the data stored in MS Registry or other similar source (for example, preferences and history files), should also be taken into account.

TABLE 2 Required Data Discovery Items

Item	Description
Physical layout	This includes the layout of the compromised system, if known, and its associated networks, including the number and types of computers and the types of operating systems and application software packages used.
Electronic mail system	This includes the structure of the electronic mail system, the software used, the number of users, the location of mail files, and the password usage.
Applications and specialized utilities	This includes applications such as calendars, accounting programs, word processors, and encryption software. It is important to know when they were installed and when they were upgraded. Any routines for archiving and purging different types of data or databases must also be known.
Operating personnel	This includes the personnel who are responsible for the ongoing operation, upgrades, and maintenance of the software, hardware, and network.
Backups and restores	This includes the backup procedures used on all computer systems in the organization. It should also include descriptions of all of the devices (such as tape or floppy drives) and software used to create backups, the personnel responsible for backups, the backup schedules, and the tape rotation schedules. Finally, it should include details on any restoring that has been done.
Policies, processes, and procedures	This includes the process for archiving and retrieving backup media both onsite and offsite, the procedures used by system users, administrators, and superusers to log in to computers and in to the customer's network, and the use of passwords, audit trails, and other security measures used to identify the data created, modified, or otherwise accessed by particular users. It should also include whether access to particular files is controlled and how, along with access lists, if available. Finally, it should include how shared files are structured and named on the system.

Manual Logs and System-Generated Log Files

There are many kinds of logs at a customer site available for follow-up investigation. These include manual logs (for example, telephone logs) created by recording all of the events during a VCSIRT processing of an incident, systems logs, network logs,

database logs, web server logs, firewall logs, IDS event logs, and various other application and system level logs. All of these are important, and any of them could provide crucial evidence in a court of law.

Detailed written chronological logs must be maintained in a journal (with specific dates and times) for legal purposes with all of the details recorded during all of the phases of incident processing: discovery and reporting, containment, eradication, recovery, and follow-up. In addition, the logs should include all of the details of contacts, all of the actions taken (with when, where, whom, and why), all of the conversations, all of the audit records, and all of the logs from the affected applications, devices, systems, hosts, networks, and databases.

In practice, every hardware platform, every piece of equipment, all of the software, or other related facilities at the customer site could provide alarms and logs. Alarms are designed to ring when a specific log entry meets a pre-defined alarm criteria. These will be needed by the VCSIRT during the eradication and recovery phases of the incident response so that lessons can be learned for future deterrence or other countermeasures. However, these are also necessary to be gathered with legal advice in the electronic discovery and litigation support in the context of the follow-up phase.

The following two example cases involve logs. The first demonstrates the importance of the need for continuity of recorded information as evidence, and the second stresses the liability issue.

In the Netherlands (*State v. Ronald O.*, 1993-5), an alleged intruder was on trial. The evidence put forward by the prosecution included a set of logs. However, on court examination, several pages were found to be missing. They supposedly contained irrelevant information, so they were removed by the prosecution team. As a result, the defense argued that evidence was being withheld. A proper handling of log files with legal advice could have possibly averted this issue.

The following is a hypothetical case. You, as a member of a CSIRT for a customer or member of the worldwide security team, receive a log file from your organization's customer site. The file has entries that clearly show intruder activities, but you fail to follow up on the leads. If this fact is uncovered by your customer (who is within your SIR constituency), you may be liable for failing to act on the information.

When system created logs are used as evidence in court, their data must be complete from a legal point of view. TABLE 3 contains some notable characteristics of system-generated log files that contribute to the validity of the logged data as evidence.

TABLE 3 Important Log File Characteristics

Characteristic	Description in Support of Evidence
Timestamps in the log file	Timestamps must be present in the log file for every event to be recorded.
Origin of records in the log file	All of the details about the machine or computer-related equipment (for instance, the version, make, and manufacturer ID number) that produced the log file must be collected.
Authentication of the log file	The log file can be questioned in court if it was created before or after the incident, especially if there was no formal and/or built-in authentication process associated with the initial access of the log file.

Auditing logs produced by computer operating systems can provide useful user and kernel-level data. How often audits should be run depends on the criticality of the customer's environment and the customer's security policy. Refer to "Auditing System Security" (Sun BluePrints Online, May 2003) for best practices and facilities available in the Solaris™ Operating System (Solaris OS) at:

<http://www.sun.com/blueprints/browsesubject#security/>

The following list emphasizes key points regarding follow-up on all kinds of logs:

- Changes within the customer's DNS, DDNS, or DHCP setups might render the IP addresses in the log useless. So, the selected characteristics of the log must be valid for legal purposes.
- Full evidential value of the logs might not be realized until they are reviewed along with the configuration files, such as `/etc/syslog.conf`, that generated the logs.
- Response teams must ensure that the network time is maintained at the customer site using industry-standard protocols such as NTP and that system clocks are synchronized. Even the most accurate computer clocks are likely to vary due to manufacturing defects, changes in temperature, electric and magnetic interference, age of oscillator, or even computer load. If NTP is not used, the timestamps from the logs of different systems from the same site or different sites might not correlate and make sound conclusions about the incident impossible. Ideally, customers must be advised to reconcile time deltas between systems and devices with respect to logs. In practice, systems are often not configured to use NTP and therefore have significant deltas that could impact the admission of

evidence if prosecution is involved after an incident. For further insightful and useful technical details, see Sun's recommended NTP practices for the Solaris OS and for the Sun Fire™ servers:

- *Using NTP to Control and Synchronize System*
- *Using NTP to Control and Synchronize System Clocks - Part I: Introduction to NTP*
at: <http://www.sun.com/blueprints/browsesubject.html>
- *Using NTP to Control and Synchronize System Clocks - Part II: Basic NTP Administration and Architecture*
- *Using NTP to Control and Synchronize System Clocks - Part III: NTP Monitoring and Troubleshooting*
- *Using NTP on the Sun Fire 15K/12K Server* at:
<http://www.sun.com/solutions/blueprints/0603/817-2979.html>

Consult the `xntpd(1M)` man page in the Solaris OS for the full list of NTP configuration options. For public NTP servers, check the PUBLIC NTP Server list on <http://www.ntp.org>. Beginning in the Solaris 2.6 and 7 OE, Sun included an NTP version that was supported on all platforms. In the Solaris 8 OE, the included version was a slightly modified version of the open source version of NTP 3-5.93e.

The following lists important questions to consider:

- What category does the log file belong to?
There are various categories for log files depending on who the target consumer is. For example, you should not include a log file that contains specific network configuration information about your constituent customer. In addition, Ethernet sniffer logs have sensitive data including user names and passwords.
- Have you taken care to avoid illegal disclosure by sanitizing the log before sending it to someone?
For example, you should remove customer passwords. What about privacy related issues? Customer names, addresses, and related information is sensitive. For further information about the implications of privacy violations, see legislation pertaining to COPPA, HIPPA, and GLBA in the U.S. and EU privacy directives (TABLE 7 on page 30).
- Are you sending logs in a secure way?
You should use an encrypted channel to send logs between team members. Another way to secure data exchange is to use read-only copies of log material (for example, CD/DVD-ROMs).

- Is the log genuine when you receive it?

You should have agreed on authentication methods, such as digital signatures, between cooperating parties. MD5 one-way cryptographic hash and RSA digital signatures are popular methods. The latter establishes the identity of the sending party and authenticity of the sent data. After the log is received, storage security also needs attention.

- Is the relevant log information properly organized so that it follows proper disclosure rules, and at the same time, is meaningful enough in an incident follow-up?

Two competing Internet service providers might exchange logs, with a VCSIRT as the intermediate party. Good judgment is important for information exchange, without compromising competitive data.

- Do you have the tools to analyze logs in a timely manner?

Logs can contain an overwhelming amount of information. There are public domain tools that can help, such as Logsurfer at DFN-CERT web site (<http://www.cert.dfn.de/eng/logsurf/>).

The large amount of log information collected, such as the messages handled by the `syslog(1M)` daemon or the log files from information services, such as FTP and HTTP services, makes it nearly impossible to check logs manually to find any unusual activity. The Logsurfer program was designed to monitor any text-based log files on a system in real time.

Insider Attacks

It is well-known that vulnerability assessment tools only find known threats such as worms and unpatched systems rather than the attacks from insiders that can cause significant damage. For such attacks, the CSIRTs need full cooperation from the affected customer to acquire incident-related data from its employees and the physical premises, including the following:

- Legal and/or law-enforcement presence or support, as deemed necessary
- Names of the HR, physical security, or CSIRT members who will be conducting the interviews
- Names of the employees who need to be interviewed
- Appropriate place to conduct the interviews
- Times, dates, and order of the interviews
- Selection and presence of a third party as a witness

With approval from the customer, servicing organization's legal advisors, and the geo-based security officers involved, system-level real-time data should be acquired by auditing, tracing, or setting up honeypots on the customer's systems and networks.

Auditing of the operating system can provide detailed user and kernel-level data to uncover any unauthorized activities by an insider. See the Solaris OS auditing practices described in the articles at:

<http://www.sun.com/blueprints/browsesubject#security/>

Prior arrangements must be made with the customer by the VCSIRT to assign an appropriate level of audit without compromising system performance, particularly on business-critical servers.

UNIX systems provide sniffers, such as snoop in the Solaris OS, for watching the network traffic in detail. Sniffers can provide critical data as evidence in unauthorized network intrusions. In the Solaris OS, the presence of unauthorized sniffers can be detected with publicly available tools such as `ifstatus` that detect promiscuous mode on a system's interface. (Note that `ifstatus` is not supported by the Solaris 7 OE, and that it is supported only for the `hme` interface in the Solaris 8 OE. For the 64-bit Solaris OS, use `ifstatus64`.)

Honeypots are considered well-suited for detecting insider attacks when it comes to computer systems or network abuse or misuse. However, some experts argue against using them as a general tool for cybercrime and computer misuse. The following are a few tips for using honeypots:

- VCSIRTs at a customer site must consult with a legal representative before deploying a honeypot because the cost of intrusive monitoring might be great, even if there are no legal barriers in the country in which the incident occurred.
- Honeypots are potential victims. Therefore, they must be sufficiently secured, except for the functions that must be open to perpetrators.
- Dispersed honeypots work better in terms of the number of hits.
- Honeypots can be hardened for robustness before they are deployed. For information on hardening the Solaris OS, see:

<http://www.sun.com/blueprints/browsesubject#security/>

<http://www.sun.com/software/security/jass/>

Authenticating, Preserving, and Analyzing Incident Data

These activities for processing incident data will be discussed in detail in the next article in the series, “Responding to a Customer’s Security Incidents—Part 4: Processing Incident Data.” In this section, brief overviews are given.

As a general rule, all criminals leave evidence behind. Authentication means that evidence you have collected during and/or after an incident must be proven to be the same as what was left behind. In practice, both proof of integrity and time stamping are provided by calculating a value that represents an electronic footprint.

Any evidence that can be used in the court of law must be preserved with extra care and security. The assigned geo-based security officers and the worldwide security manager should consult the corporate security and legal departments for the incident servicing enterprise to review the evidence.

When the time comes for analyzing the data, there is a whole range of actions and processes a VCSIRT can execute depending on the time available to analyze events thoroughly and to disclose the outcomes to its constituent customers and to other teams. The decision as to the level of analysis to be conducted lies with the team, yet the affected customer must be consulted by the VCSIRT before making a decision.

An initial broad analysis (which might include cursory forensics) must precede detailed analysis because it helps to understand which response plan to follow. You can branch out on a Denial of Service or Unauthorized Access investigation and response plan, or others, depending on the conclusions reached from the broad analysis. The best practice should be to clearly demarcate the response plan so that team resources can be channeled in the right direction at appropriate times.

Conducting Post-Incident Activities

There are other activities to perform following the recovery of an incident that must be supervised by the geo-based security officer and tracked by the organization's worldwide security manager.

Inventory of System Assets

Some customers might not have equipment owned by the incident servicing organization at their sites. However, for those customers who do lease or maintain vendor-owned equipment, an inventory of the servicing vendor organization's system assets should be maintained per site by the customer so that a record is available to the incident response team's organization, if necessary, for examination during an incident investigation. In addition, the enterprise IT department must maintain an inventory of the pertinent LAN, VLAN, WAN, and WLAN systems that are connected with the organization's constituents. The organization's line management and worldwide security team should also monitor these systems.

Vulnerability Discovery and Removal

Some or all of the discovery and removal of vulnerabilities might occur in the eradication and/or recovery phases. However, the following items can also be identified during the follow-up phase:

- New unauthorized user accounts
- Processes owned by unfamiliar users
- Modified or deleted data and binaries of system or application executables or libraries
- Denial of services (for example, a customer's system suddenly goes into single user mode)
- Poor system performance
- Accounting discrepancies
- Suspicious login attempts

While network-level snooping tools passively observe and analyze network activity, network vulnerability scanners actively send packets over a network in search of vulnerabilities and malicious code on the hosts of the network.

With readily available tools and information on the Internet, finding specific vulnerabilities of an operating system is easy. A person with a malicious intent can find information about vulnerabilities from one of the several organizational sites that provide full disclosure, such as CVE (<http://cve.mitre.org>). The hacker can then trace a vulnerability number from CVE to a source code clip on the SecurityFocus web site (<http://www.securityfocus.com>). He or she can then find corresponding detailed instructions regarding the exploit on the SANS web site (<http://www.sans.org>) or some other public site. This can all be done in minutes.

Backdoors and Malicious Code

The VCSIRT should detect possible backdoors and malicious code introduced into the customer's network. An example is a rootkit, a set of software tools or scripts that enable a hacker to reenter the network for further misuse or damage. For more information on backdoors and malicious code, refer to the second article in this series, "Responding to a Customer's Security Incidents—Part 2: Executing a Policy."

Enabling Vulnerabilities

Enabling vulnerabilities are usually not caused by any malicious act. They tend to be genuine mistakes by system administrators and/or users. They also enable intruders to reenter the system or network using configuration weaknesses. Examples of enabling vulnerabilities on UNIX systems are poor or commonly used passwords, unused guest accounts, accounts with default passwords, misconfigured anonymous FTP, and inappropriate settings or entries in files such as `/etc/ttys/`, `/etc/ttytab`, or `/etc/aliases`. In addition, unpatched system software can also present unforeseen vulnerabilities.

Combination Threats

Recently seen combination or blended threats such as CodeRed, Nimda, and BugBear combined the characteristics of viruses, worms, trojan horses, and malicious code with server and Internet vulnerabilities to initiate, transmit, and spread an attack. By using distinct methods and techniques, these threats often spread rapidly causing widespread damage.

For example, the most recent BugBear.B mass mailer worm exploited a vulnerability in Internet Explorer. After it entered the network perimeter, it spread by using network shares. It automatically executed when the email message was previewed by an unpatched vulnerable host, so simply receiving the email triggered an infection because it emailed itself with a spoofed From address. The worm was polymorphic in nature. It also attempted to evade antivirus applications and to

deactivate them, along with host or desktop-based firewall applications. It also flooded printing devices. Another example of a blended threat is the ISS/sadmind worm that attacked the Windows and Solaris operating systems.

Vulnerability Best Practices

The primary concepts for best practices against vulnerabilities are as follows:

- A CSIRT must employ tools and procedures to verify and execute up-to-date vendor patches on all the customer systems in the response team's constituency.

For more information, refer to “A Patch Management Strategy for the Solaris Operating Environment” (Sun BluePrints OnLine, January 2003).

For example, for mission critical or business critical environments, the recommendation is to use three patch rollout schemes: regularly scheduled, rapid, and emergency. For an emergency rollout, testing is done in a few hours in the unit and integration test environment before rolling it out.

- A thorough scan with multiple vendor tools is effective.

The reason is that the customer should not fall into a trap by using a single vendor's tools that can succumb to some oversights by that vendor. The scans should include a thorough check of all possible vulnerabilities known within the customer's network involving hosts from multiple vendors.

Nessus (<http://www.nessus.org>) is well-maintained and widely popular. It scans systems and evaluates vulnerabilities present in services offered by that system. Although it has command-line and GUI modes, its GUI mode is more convenient. There are other related tools that can be useful (for example, nessQuick, a tool to manage Nessus reports using a database, and NessusWeb, a web interface to Nessus). It now includes better use of CERT advisories and CVE (common vulnerabilities and exposures) references.

There are some other excellent tools in the public domain. SATAN and Nmap are popular as network scanners. SATAN, which has undergone many changes since its initial release, is limited in its usage. A better version, SAINT, is aggressively maintained by the SAINT Corporation (www.saintcorporation.com) to keep up with the latest vulnerabilities. When vulnerabilities are detected, the SAINT vulnerability scanner categorizes the results in several ways, allowing customers to target the data they find most useful. SAINT can group vulnerabilities according to severity, type, or count. It can also provide information about a particular host or group of hosts. SAINT describes each of the vulnerabilities it locates, references CVEs, CERT advisories, and information assurance vulnerability alerts (IAVAs). It also describes ways to correct the vulnerabilities. In many cases, the SAINT vulnerability scanner provides links to sites where you can download patches or new versions of software that will eliminate the detected vulnerabilities.

Nmap is a superb low-level port scanner and a must-have for any CSIRT to scan different protocols (for example, TCP, UDP, ICMP, RPC, and Reverse-Ident). For more information, refer to the <http://www.insecure.org/nmap> site.

- Employees must be screened, and IT responsibilities must be segregated.
Insider attacks are usually a significant percentage of the total registered attacks in any thoroughly collected statistical security incidents data. This warrants cautious screening of employees without upsetting their rights to privacy. IT responsibilities should be segregated for clear accountability and internal auditing must take place routinely.
- Customers must be advised on careful categorization of event data from devices such as firewalls and intrusion detection systems. Of particular importance are:
 - Attacks and attack patterns
 - Events (realize that a single event might consist of multiple attacks)
 - Unique attackers (with single IP address or a set of repeatedly used addresses)
 - Segregation and classification process of events and attacks
 - Determination of vulnerabilities, their severity levels, and counts in a specified period
 - Automated references to industry-known vulnerability indices such as CVEs or IAVAs
 - At least two levels of summary reports—one technical for IT administrators and one for management are necessary.

In addition, reporting should contain information for the evaluation of a solution (for example, a software patch) against a threat and the potential impact on the stability of the monitored device. There should be a policy statement for actions to be taken, channel of communication, time period, and responsible contacts. TABLE 4 contains a brief report policy implementation example.

TABLE 4 Reporting Recommendations

Severity and Issue	Action	Contact	Communication
Critical: DNS00023	Immediately install of patches and/or other solutions to fix internal DNS server (in less than 24 hours)	Security policy owner and system administrator	Within eight hours
Moderately critical: WebServer00405	Install patches and/or other solutions on web server in the data center (in less than 48 hours)	Security policy owner and system administrator	Within 16 hours
Noncritical: Solaris10067	In the next maintenance window, update the Solaris 9 OS kernel on the directory server in the DMZ (in less than one week)	System administrator	Within 48 hours

- Host-based analysis tools help customers learn things about specific systems when following up on an incident. They are useful to detect malicious code or backdoors. In addition, they can assist in the detection of unauthorized changes to system files or applications. Tripwire is a useful tool in this regard.

Damage Assessment

As soon as the security breach has occurred, the entire system, its network, and all of its components should be considered suspect by the organization's VCSIRT. The analysis of the damage and extent can be time consuming, but the organization's customer account manager and the assigned security officer must drive the work with two main reasons in mind: it could lead to some insight into the motives and nature of the incident, and most prosecutors ask for an estimate of the loss when discussing a case or while considering sentencing guidelines.

A security incident of any kind has several cost components associated with it. The lead of the response team (VCSIRT) and the geo-based security officer, has to determine the extent of damage due to the break-in and notify the customer's site administrator, users, and servicing organization's worldwide security team. The damage should be taken into account for future precautions and risk assessment, clean up, and estimating the affect on users. Loss to the customer's enterprise in the form of decreased productivity must also be considered.

The proposed Senate Bill S2448, “The Internet Integrity and Critical Infrastructure Protection Act,” clarifies how loss should be calculated. It states that “the term ‘loss’ means any reasonable cost to any victim, including the cost of responding to an offense, conducting a damage assessment, and restoring the data, program, system, or information to its condition prior to the offense, and any revenue lost, cost incurred, or other consequential damages incurred because of interruption of service.” Thus, the costs that must be tallied include:

- Time spent by all the servicing organization's advisory personnel (such as the SAG), engaged VCSIRTs, the worldwide security team, and the customer's staff in cleaning up the damage and bringing systems back online (with tasks such as analyzing what has occurred, re-installing the operating system, restoring installed programs and data files).
- Lost productivity by system, network, and site administrators and end users who were prevented from using the systems during down time or during any DoS attacks associated with these individuals using compromised systems or networks at the affected site.
- Replacement of hardware, software, and/or other material or intellectual property that was damaged or stolen.

To assess the damage, you must determine the needed changes for a tainted system. That, in turn, will help estimate the amount of work involved. The following list contains some examples:

- Employing checksums of all associated media and usage of tools that provide before and after comparisons
- Looking at all centralized and decentralized logs for abnormalities
- Examining patterns of system usage in system accounting records for abnormalities

The following two tables contain examples of cost calculations for the VCSIRT and the users of the affected customer site. The cost and hour value must be based on estimated salaries plus overhead and any indirect cost. Cost calculation must take into account the variance based on known, determinable factors for cost to the servicing team members and the cost to the users. (The variance is shown in the table as x% and y%.)

TABLE 5 Security Incident Cost Analysis for CSIRT for Incident Tracking #100001

VCSIRT Worker	Hours	Cost per Hour	Total	-x% (with variance)	+x% (with variance)
Geo-based security officer					
Security engineer					
Field systems engineer					
Total labor cost					
Median cost +/- x%					

TABLE 6 Security Incident Cost Analysis for CSIRT for Incident Tracking #100001

Number of Users	Hours	Cost per Hour	Total	-y%	-y%
Web site users					
Application users					
System administrators					
Total cost (lost productivity)					
Median cost to users +/- y%					

In addition to the above costs, other relevant costs such as any materials used, travel, PR, legal, and investigative or government agency consulting must be taken into account in a similar way with variances.

Risk Analysis

A new risk analysis must be conducted by the organization's Security Advisory Group (SAG), working with the organization's worldwide security team. No matter what risk analysis process is used (for example, qualitative versus quantitative), the overall method should remain the same. In general, the process should include the following activities:

- Identifying the asset to be reviewed at the customer site
- Ascertaining the threats and associated risks

- Determining priorities on the risks
- Implementing corrective measures
- Monitoring the effectiveness of the controls or corrective measures implemented earlier

Fundamentally, risk analysis can be used to review any task, project, or idea. It could also relate to a recent event (that is, a lesson learned). As an arbitrary example, after a conversation with a Japanese JPCERT, one member of the Italian VCSIRT had inadvertently distributed information about a serious bug in the a Japanese vendor's operating system. Later, this turned out to be false information. The vendor was not pleased, and now the VCSIRT's parent organization and its enterprise are liable for wrongful disclosure.

Consider a hypothetical example. A VCSIRT member advised its customer to modify and reorder its boundary firewall rules to solve an IP-level filter performance problem. However, the fix silently opened up the customer's LAN to Internet intruders, which could subsequently result in a break-in. So, processes need to be fixed with regard to the team's internal and customer's review cycles, as appropriate. Information disclosure and configuration changes must reduce future liability risks arising from wrongful disclosures and advice.

In the context of the follow-up phase, two things need to happen:

- First, the servicing organization's SAG should use risk analysis to determine if a security architecture, design, development, process, or procedural project should be undertaken to improve the security of the entity where the root cause of the compromise (that triggered the incident response process and risk analysis) was determined.
- Second, the SAG's advice to the worldwide security team should follow up with a decision making process within the team, keeping in mind that the risk analysis process is required to support the business or mission of the customer's enterprise.

Lessons Learned

As depicted in FIGURE 2, lessons need to be captured throughout the incident response process and then fed back into the process at every step, as deemed necessary. A lesson learned in the Recovery phase could suggest improvements in the Evaluation and Containment phases.

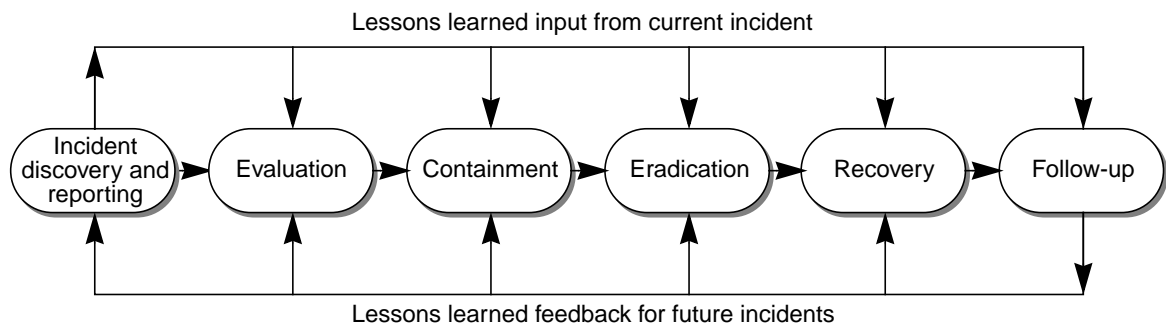


FIGURE 2 Lessons Learned in the Computer Security Incident Response Process

There is no apparent or predefined correlation as to which phase's input will affect improvements in other phases. The organization's geo-based security officers need to oversee the capturing of this vital information. The geo-based security officers should work with the geo-based customer account managers to understand the lesson as a result of the incident, and document it clearly in a standard, agreed upon format.

The content of the document must include:

- Incident description with date and time
- Location, site, network, and host system
- Names of all of the parties involved
- Recommended solutions
- Implemented solutions, with details about who, when, and how it was implemented, as well as which solutions succeeded and which failed
- Pending issues that need further investigation, if any
- Recommended changes in the security policy

You should include the Lessons Learned document in the revised security plan to prevent a similar incident from recurring and for review by the organization's SAG. You should also communicate any changes over the organization's security alias and ensure that all of the organization's security personnel are aware of changes in the policy, procedure, or administrative practice.

Upgrades of Policy, Processes, and Procedures

The organization's SAG is responsible for advising on upgrades to the policies, processes, and procedures. The procedure should be well-documented in the security incident response policy. The following are some specific examples of actions that need to be taken:

- Establishing mechanisms for updates of policies, procedures, and tools. The customer's enterprise corporate security principles must be considered by the servicing VCSIRT.
- Employing standardized security review specifically when introducing or upgrading to new applications and business partners in the organization's product or services delivery infrastructure.
- Establishing channels of communication to make all of the appropriate teams of the organization aware of the latest upgrades to the security incident response policy. At least a quarterly meeting must be held by the incident servicing geo-based security officer and monitored by the organization's worldwide security manager.

Using Legal, Investigative, and Government Recourses

Legal recourses could be needed when the recovery starts or after capturing evidential data. But, consider that every situation is different and that legal action might not be necessary in each case, even when it looks imminent. That is why an organization must consult its own legal department before proceeding any further. The following figure shows a typical prosecution decision to be taken by a customer working with the servicing VCSIRT during the recovery process.

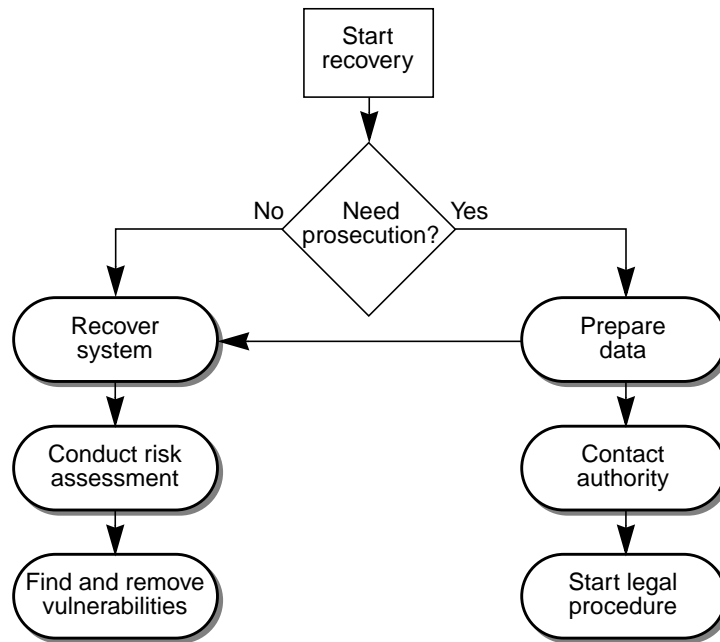


FIGURE 3 Starting the Legal Procedures After an Incident

Most countries have no computer crime laws. This makes it difficult to prosecute a suspect who is moving between countries and/or committing crimes in multiple countries. Therefore, VCSIRTs must seek legal advice to reduce their legal exposure by defining the purpose and boundaries of the VCSIRT and, in certain cases, the internal CSIRT.

The following are some examples of items that should be reviewed by the security officers with representatives from the enterprise's legal department:

- Written procedures that might cross company boundaries
- Negative publicity for the enterprise, its customers, or business partners that can result from non-containment, delayed resolution, or mishandling of an incident
- Downstream liability for the enterprise, its customers, vendors, or business partners
- Secure distribution of information regarding an incident or affected system
- Liabilities due to and during monitoring and follow-up of an ongoing incident
- Liabilities due to disclosure of court orders, leaks of information from CSIRTs, or leaks of information through intrusions
- Disclosure of information (for example, to the media, the competition, or third parties) that might be unnecessary or unauthorized

- Requirements for legal evidence
- What kinds of data can be collected under which circumstances (or what is allowed and what is prohibited under the local laws, such as the laws related to data protection and telephone communications)

Litigation Support and Cost

Suppose a customer has agreed to make a legal case out of the incident response investigation and costs are analyzed to see if the case makes economical sense. If it is deemed that it makes sense, the response team needs to find and organize a legal support team with the help of the enterprise's attorney or an external legal firm.

Litigation support is best described as the management of litigation cases through technology, primarily the organization of paper-based documents into an electronically organized database for review and reuse in production. The main functions of litigation support specialists are transcript management and trial preparation, along with managing the information database. Electronic discovery experts identify, collect, and analyze computer data from the compromised site.

Litigation and discovery support can be costly, so VCSIRTs must consider the cost because these services are usually outsourced. However, some of the litigation support specialists, IT specialists, and electronic discovery experts might be employees of the parent enterprise. In addition, some of their functions might overlap. However, in a high-profile case, you might need to find specialists in specific areas, which could result in higher costs. FIGURE 4 illustrates how the attorney should manage the investigation and interface with the incident response team.

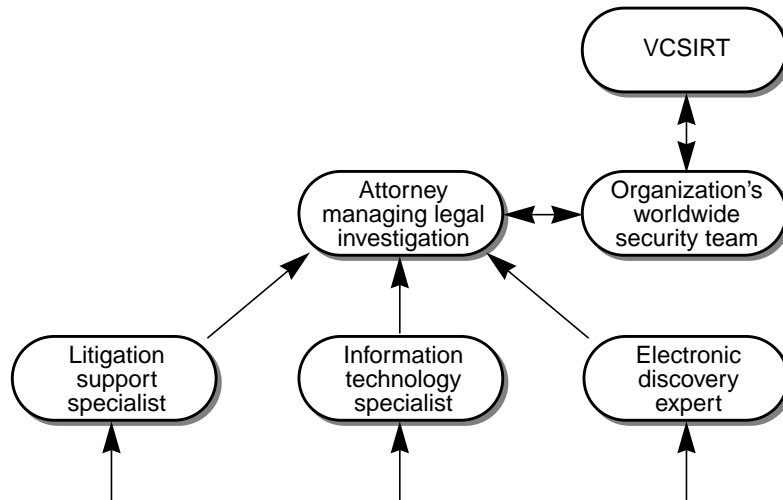


FIGURE 4 Management of Security Incident Legal Investigation

U.S. and International Laws of Interest

Legal advice can be sought at various points during a follow-up phase. For instance, if you intend to prosecute an intruder, you need to have complete, thorough, and convincing evidence that has been protected through a secure chain-of-custody procedure that tracks who has been involved in handling the evidence and where it has been stored. Law enforcement officials and your legal counsel are good sources for advice about how and when to collect and protect critical information. In addition, they can assist in interpreting application of current legislation (local to a particular country) that can be helpful at different junctures in the context of a specific incident.

TABLE 7 contains international and U.S. laws of significance to incident response and computer crime investigation personnel. To international readers, the U.S. laws are examples to pursue and compare to local laws for recourses you might have in the future. In addition, for transnational incident related issues, the laws provide guidance.

To U.S. residents, CSIRTs, and general readers, the laws can be very helpful when reporting security-related incidents to law enforcement, including violations of privacy legislation. You can find more information on these and other related laws at the <http://www.usdoj.gov/criminal/cybercrime> site.

TABLE 7 Security Incident Response Related Laws of Interest

Law	Location	Description
European Convention on Cybercrime by the Council of Europe (COE), 2001	Europe and international	The COE is a recent attempt to codify international standards of cooperation when investigating computer crimes.
European Commission's Regulation (EC) 45/2001 of the European Parliament and Council, December 18, 2000	Europe	This regulation contains data privacy regulations (http://europa.eu.int/comm/internal_market/en/dataprot/law/index.htm).
Electronic Communications Privacy Act (18 U.S.C. 2510), 1986	U.S.	The ECPA updated the Federal Wiretap Act to account for computers. This law protects against the unlawful interception of any wire communications (including email).
Economic Espionage Act (18 U.S.C 1831 and 1832), 1996	U.S.	The EEA outlaws foreign espionage and the theft of trade secrets.
Identity Theft and Assumption Deterrence Act (18 U.S.C. 2028), 1998	U.S.	This law makes it illegal to use interstate wire communication systems to commit a fraud to obtain money or property.
Computer Fraud and Abuse Act (18 U.S.C. 1030), 1984	U.S.	The CFAA, which has been updated several times, is the major federal law dealing with computer crimes (for example, fraud, extortion, and theft of financial information).
U.S. Department of Justice (DOJ) Guidelines for Warning Banners	U.S.	The DOJ released these guidelines for system administrators to help in the prosecution of computer crimes. It supports EEA and was primarily distributed by CERT/CC as Advisory 1992-19 "Keystroke Logging Banners."
HIPPA, Health Insurance Portability and Accountability Act of 1996	U.S.	HIPPA provides protection against illegal use of a personal medical information.
GLBA, Gramm-Leach-Bliley Act	U.S.	Title V of the GLBA generally prohibits any financial institution, directly or indirectly, from sharing personal information about an individual with a nonaffiliated party unless the financial institution's privacy notice states otherwise.
COPPA, Children's Online Privacy Protection Act	U.S.	COPPA, which became effective in April 2000, requires companies to implement safeguards when they collect marketing data online from children under 13.

Government Security and Investigative Resources

There is a huge community of incident responders belonging to various levels of government: regional, state, and national. The important points for best practices are as follows:

- Government resources can be particularly useful when the customer is dealing with classified information or when the customer is an agency of the government.
- Keep in mind that laws vary on every kind of crime from country to country. You need to determine the legal requirements of the country in which you are operating, especially the laws that pertain to collecting and protecting evidence, chain of custody, and sufficiency of evidence for prosecution. Then, you need to implement the necessary procedures to meet those requirements and weigh in all the alternatives to legal prosecution before proceeding.
- As an initial step, the enterprise security and/or legal department should be contacted by the organization's security officers when local government investigative and law enforcement help is needed to assist the customer in your constituency.
- Then the customer and/or you, as the servicing organization (or both together, as the predesigned SLA might dictate), will need to decide to whom to report the crime. Usually, it is better to deal with local or state authorities, if at all possible. For example, in the U.S., every state currently has laws against some sort of computer crime. If your local law enforcement believe the crime is more appropriately investigated by the federal government, they will suggest that you contact federal authorities. In some cases, local authorities might be more responsive, although they might not have complete jurisdiction for the crime your customer wants investigated.
- You do not need to determine the jurisdiction on your own. If you believe that a law has been violated in the U.S., you can call the nearest U.S. Attorney's office and ask them whom to contact. In some cases, when it is determined by your local team that you have no other choice, you will need to pursue legal means to protect your team, organization, and your constituent. But, do not encourage or develop a false sense of security. In some circumstances, cooperating with law enforcement is not a sufficient shield from liabilities. Expert legal advice from your CSIRT's or organization's lawyer might be of great help.
- Securing cyberspace and deterring, reducing, or eliminating security incidents is a difficult and strategic challenge. It requires a coordinated and focused effort from a global society, the federal governments, the state and local governments, the academic research organizations, and the private sector. Thus, maintaining contacts within the global community of incident response teams and keeping up with the policies, processes, and procedures of interaction with them is absolutely critical. FIRST, the Forum of Incident Response and Security Teams (<http://www.first.org>), is a good starting point for establishing outside contacts.

TABLE 8 contains a list of worldwide governmental resources.

TABLE 8 Worldwide Security Incident Resources

Country	Resource	Description
Australia and New Zealand	AusCERT (http://www.auscert.org.au/)	AusCERT is the national Computer Emergency Response Team for Australia and New Zealand and a leading CERT in the Asia and Pacific region. AusCERT maintains a worldwide recognized reputation and a trusted contact network of computer security experts. It also provides prevention, response, and mitigation strategies for members.
Brazil	Brazilian Computer Emergency Response Team (http://www.first.org/team-info/#NBSO)	The mission of the Brazilian CERT is to register and follow up on security problems in the RNP backbone and PoPs, as well as to help identify invasions and repair the damages caused by invaders. The Brazilian CERT is also involved in publicizing information on preventive action regarding network security.
Canada	Canadian Computer Incident Response Coordination Centre (http://www.ocipep.gc.ca/index.asp)	Canadian CERT is part of the Office of Critical Infrastructure Protection and Emergency Preparedness.
	DND CIRT	DND CIRT is part of the Department of National Defense.
	EWA-Canada/Canadian Computer Emergency Response Team (http://www.cancert.ca/Home/Default.php)	CANCERT was set up for the Canadian government, business, and academic organizations.
Croatia	CERT (http://www.cert.hr/)	Croatian CARNet is Croatia's computer emergency response team.
France	CERT Administration (http://www.certa.ssi.gouv.fr/)	CERT Administration is part of the French government.
	CERT-Renater	CERT-Renater is part of the Ministry of Research and Education.
Germany	CERT-Bund (http://www.bsi.bund.de/certbund/)	CERT-Bund is part of the German government.
India	Internet Security Center, New Delhi	Beginning in July 2003, this center was developed by India's central information technology ministry with assistance from U.S. CERT/CC. The center seeks to prevent cyber attacks on key defense, business, and government organizations.

TABLE 8 Worldwide Security Incident Resources (*Continued*)

Country	Resource	Description
Italy	CERT Italiano (http://www.dico.unimi.it/)	Italian Computer Emergency Response Team (CERT-IT) was founded in February 1994. The CERT-IT is a non-profit organization mainly supported by Dipartimento di Informatica e Comunicazione (DICO).
Japan	JPCERT (http://www.jpCERT.or.jp/)	JPCERT serves the Internet community in Japan.
Russia	Computer Security Incident Response Team RU-CERT (http://www.cert.ru/Eng/)	RU-CERT serves the Russian Federation (.ru and part of .su).
Sweden	SUNET CERT (http://www.cert.sunet.se/)	SUNET CERT serves Swedish universities with a direct SUNet connection.
	Swedish IT Center (http://www.sitic.se/)	Swedish IT Center is a Swedish government agency.
U.K.	Unified Incident Reporting and Alert Scheme (UNIRAS) (http://www.uniras.gov.uk/)	HM Government CERT/U.K. National Infrastructure Security Coordination Center (NISCC) serves the HM Government and specified sites within the .uk domain.
	Security Service (http://www.mi5.gov.uk/)	MI5 is the U.K. defensive security intelligence agency.
	Fight terrorism web site (http://www.homeoffice.gov.uk/terrorism/)	This information web site helps to fight terrorism (including cyber-terrorism) and contains useful information on fighting terrorism.
U.S.	Electronic Crimes Task Forces (http://www.ectaskforce.org)	Following the 9/11 attacks, the Patriot Act of October 2001 created Electronic Crimes Task Forces with the New York task force as the model. Regional task forces exist for the Bay Area, Chicago, Las Vegas, Los Angeles, Charlotte, Miami, New England, and New York (http://www.ecTaskForce.org/Regional_Locations.htm).
	Federal Bureau of Investigation (FBI)	The mission of the FBI is to investigate violations of federal laws and to protect the U.S. from foreign intelligence and terrorist activities, including cyber-terrorism.
	FedCIRC (http://www.fedcirc.gov)	FedCIRC is an incident response resource that is available to the entire federal government through the General Services Agency.

TABLE 8 Worldwide Security Incident Resources (*Continued*)

Country	Resource	Description
	Department of Energy Computer Incident Advisory Capability (http://www.ciac.org)	Although it principally serves the Department of Energy, the CIAC web site contains useful information for security and incident response, including one of the best virus hoax repositories.
	National Infrastructure Protection Center (NIPC) (http://www.nipc.gov)	NIPC is a Department of Justice clearing house and coordinator of computer crimes investigations.
	U.S. Department of Justice Computer Crimes Section (http://www.cybercrime.gov)	Computer Crimes Section of the Department of Justice represents U.S. prosecutors of computer crimes.
	CERT Coordination Center at Carnegie Mellon University (http://www.cert.org)	CERT Coordination Center is a federally funded research and development center. It handles computer security incidents and vulnerabilities, publishes security alerts, researches changes in networked systems, and develops information and training to help improve security in systems, networks, and web sites.
	Transportation Security Administration (TSA) (http://www.tsa.dot.gov)	TSA is responsible for protecting the U.S. transportation system and is charged with seeking excellence in transportation security through its people, processes, and technologies.
	U.S. Secret Service (http://www.secretservice.gov)	The Secret Service is responsible for the enforcement of laws relating to counterfeiting of obligations and securities of the U.S. It investigates financial crimes and computer-based attacks on U.S. financial, banking, and telecommunication infrastructures.
	Office of Homeland Security (http://www.whitehouse.gov/homeland)	This office coordinates the efforts of the Executive branch of the U.S. Government to detect, prepare for, prevent, protect against, respond to, and recover from terrorist attacks (physical and cyber). To locate Homeland Security directors at the state level, go to: http://www.whitehouse.gov/homeland/contactmap.html

TABLE 8 Worldwide Security Incident Resources (*Continued*)

Country	Resource	Description
	National Security Agency (http://www.nsa.gov)	N.S.A. coordinates, directs, and performs highly specialized activities to protect U.S. information systems and to produce foreign intelligence information.
	Critical Infrastructure Assurance Office (CIAO) (http://www.ciao.gov)	CIAO was started in response to a Presidential Directive (PDD-63).
	President's Critical Infrastructure Board (http://www.cybersecurity.gov)	This board consists of 25 agencies. Its scope covers information systems for critical infrastructure, including emergency preparedness communication and physical assets that support such systems.

Article Series

The “Responding to a Customer's Security Incidents” articles are an on-going series. The next article will cover best practices for processing incident data. It will discuss authenticating, preserving, and analyzing the data, determining the depth of the analysis and response plan, and understanding the essentials of forensics. The article also includes discussions on publicly available tools for forensic analysis.

References

The following lists the references that were used for this article:

- CERT. “Identifying Tools that Aid in Detecting Signs of Intrusion” (Tools for real time and forensic analysis), at: <http://www.cert.org/security-improvement/implementations/i042.07.html>
- Electronic Crimes Task Force and Helpful Links, at: <http://www.ectaskforce.org/> and http://www.ectaskforce.org/Helpful_Links.htm
- Feldman, J. “The Essentials of Computer Discovery,” 2002, at: <http://www.forensics.com>
- FIRST. Registered teams, at: <http://www.first.org/team-info/>
- Germany's DFN-CERT. Logsurfer home page, at: <http://www.cert.dfn.de/eng/logsurfer/>

- Masurkar, Vijay. “Responding to a Customer’s Security Incidents—Part 1: Establishing Teams and a Policy.” Sun BluePrints OnLine, March 2003, at: <http://www.sun.com/solutions/blueprints/>
- Masurkar, Vijay. “Responding to a Customer’s Security Incidents—Part 2: Executing a Policy.” Sun BluePrints OnLine, April 2003, at: <http://www.sun.com/solutions/blueprints/>
- Medford Police, at: <http://www.medfordpolice.org/>
- Noordergraaf, Alex. “Enterprise Security: Solaris Operating Environment.” Prentice Hall, 2002.
- Rude, T. “DD and Computer Forensics: Examples of Using DD within UNIX to Create Physical Backups,” at: <http://www.crazytrain.com/dd.html>
- Sun Microsystems, Inc. Solaris Security Toolkit, at: <http://www.sun.com/software/security/jass> and <http://www.sun.com/blueprints/tools>

Acknowledgments

The author would like to recognize the following individuals for their contributions: senior personnel from Sun Services, Engineering, and Sun Corporate and IT Security for reviewing this article and providing helpful comments (in particular, Joel Weise, Martin England, Glenn Brunette, Matthias Kussinger, Larry Dunn, Wayne Vincent, and Steve Gilliss). Sun BluePrints OnLine program contributors in revising and editing are very much appreciated (in particular, Dan Barnett and Billie Markim).

About the Author

Vijay Masurkar is a Principal Engineer and a leading services architect and technologist for Network and Security at Sun Microsystems. Currently in Sun Support Services, his research interests are best practices for enterprise and Internet level reliable and secure network architectures and customer security services. Vijay has been in the computer network and security industry for twenty-eight years. He has led research and development projects, consulting and support for VAX/VMS, Wang VS, Solaris, and TCP/IP based network and security products and services. He represents Sun in several industry forums. He is often invited to teach at Sun and in the industry. Vijay holds a B.S. in Electrical Engineering, an M.S. in Computer Systems Engineering, and an M.B.A.

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is `http://docs.sun.com/`

To reference Sun BluePrints OnLine articles, visit the Sun BluePrints OnLine Web site at: `http://www.sun.com/blueprints/online.html`

