Computers & Security (2006) **25**, 351—370

Computers & Security

www.elsevier.com/locate/cose

On Incident Handling and Response: A state-of-the-art approach

Sarandis Mitropoulos*, Dimitrios Patsos, Christos Douligeris

Department of Informatics, University of Piraeus, 80, Karaoli and Dimitriou Street, Piraeus 185 34, Greece

Received 24 February 2005; revised 18 July 2005; accepted 23 September 2005

KEYWORDS
Incident Handling;
Incident Response;
Computer forensics;
Internet forensics;
Software forensics;
Trace-back mechanisms

Abstract Incident Response has always been an important aspect of Information Security but it is often overlooked by security administrators. Responding to an incident is not solely a technical issue but has many management, legal, technical and social aspects that are presented in this paper. We propose a detailed management framework along with a complete structured methodology that contains best practices and recommendations for appropriately handling a security incident. We also present the state-of-the art technology in computer, network and software forensics as well as automated trace-back artifacts, schemas and protocols. Finally, we propose a generic Incident Response process within a corporate environment.

Introduction

We stand today in front of the most important advances of Information Technology. Considering that nearly one billion computing systems in our planet are already connected through the Internet (Global Reach, 2005), as well as the convergence of mobile telephony services and e-commerce, huge amounts of information flow from one network to another with a single command, request or click.

To this extent many technologies, platforms and infrastructures are thriving to provide services to the end user, who becomes the target point: it is the user who requests services; it is the user who accesses networks and resources; it is the who that requires security and privacy. Users are 'carrying' their digital identities (usernames, passwords or PINs, digital certificates or biometric features) over several different platforms and applications to request access: on their corporate workstations, on their home computers, on their mobile phones or PDAs, leaving important traces of their selections, habits and personal data.

In this context, there are many system vulnerabilities that, if not addressed properly, could potentially lead to a set of unauthorized actions varying from a denial-of-service attack to identity theft. Therefore it is of utmost importance to treat

* Corresponding author.
E-mail address: sarandis@unipi.gr (S. Mitropoulos).

0167-4048/$ - see front matter © 2005 Published by Elsevier Ltd.
doi:10.1016/j.cose.2005.09.006

every security incident (i.e. any related activity with negative security implications (CERT/CC, 2005)) to its full extent, by applying appropriate response methods, mechanisms and/or policies in order to minimize its effect(s). On the other hand, wherever applicable, these actions should aim to reach the actual source of an incident. These countermeasures can vary from simple fixes (e.g. a software update) to extremely complex Incident Response Policies that have to be enforced within organizations. Incident Response can be defined as the process that aims to minimize the damage from security incidents and malfunctions, and monitors and learns from such incidents (BSI, 1999).

This paper proposes a management framework of Incident Response by examining proposed techniques (both research and applied), best practices and technology implementations. We first introduce the concept of Incident Response and present the management issues that arise when an organization is asked to handle a specific security incident. The aim is not only to present the complexity of the Incident Response process but also to identify and categorize the relations arising with various aspects of the Information Security and the Information Technology (IT) Management sectors of an organization. We later introduce a structured methodology to isolate and handle a security incident and present the various phases of this methodology by proposing best practices that could be enforced in every phase. A conceptual process that can be used to examine the system log files is also proposed. We refer to all these procedures as passive Incident Response (also known as Incident Handling).

We continue the paper by examining the active ways of responding to an incident, the techniques used to find the source of the incident, i.e. the computer domain, the network, the system, the user within a system or the actual person(s) who started the incident in the best case. We then present the most well-known automated trace-back mechanisms that have been proposed in the context of software forensics concepts at the application level, computer forensics at the system level and Network Forensics at the network level. We refer to all these procedures as active Incident Response (also known as trace-back). A generic Incident Response scenario within a corporate environment is also provided to show the applicability of these techniques.

We finally examine the issues that remain open from a technical, management, legal and social perspective by emphasizing upon identity theft that reflects the changing focus of modern attacks: from the corporate environment to the ordinary Internet user. The paper concludes with ideas for further research and development in this area.

The problem space

Organizations and individuals invest on preventive security measures to protect their assets, such as safes, secure doors, windows, etc. Their main operation is to prevent someone from acting in a way that could pose harm on these assets. Security by prevention, though, is not enough. Most of those who implement such measures also invest on detective and monitoring mechanisms (e.g. CCTV, alarms, lights, etc.) to supplement the effectiveness of the previous measures. Nevertheless, even the combination of these measures does not guarantee total security. Appropriate response mechanisms are necessary when a security-related incident occurs. For example, when the alarm sounds when a burglar enters a house, all the previously taken preventive countermeasures are worthless if the police do not show up in a short time interval.

A similar situation appears in the Information Security area. Many products that are able to protect a company's information assets and resources from unauthorized network access have been developed, in the last decade, firewalls being perhaps the most popular and most effective. Intrusion Prevention/Detection technology is offering monitoring services thus supplementing the effectiveness of the firewalls as well as a wide range of other preventive countermeasures. If a security incident does happen though, Incident Response is necessary to mitigate the immediate damage, eliminate any possible consequential losses and prevent any possible future recurrence.

The management perspective of Incident Response

The Internet is being the network that links the entire planet so responding to security incidents often requires the coordination of international efforts. The Computer Emergency Response Team/ Coordination Center (CERT/CC) at Carnegie Melon University has set the first milestone by providing a central location for the reporting of such incidents as well as for providing the appropriate solutions. Similar efforts can be found in Europe (Forum of Incident Response and Security Teams — FIRST), Australia (Australian Computer Emergency Response Team — AusCERT) and elsewhere. Computer Security Incident Response Teams (CSIRTs) are becoming an essential part of modern Information Security Standards (like ISO/IEC 17799) (International Standards Organization, 2000). The work of CSIRTs is fully described in RFC 2350 (Internet Engineering Task Force, 1998).

CSIRTs can be categorized by the specific constituency (internal, external, commercial, vendor, governmental, academic, etc.) they serve or by the type of their existence and formation (formal, ad hoc, etc.) (Van Wyk and Forno, 2001). CSIRTs can be members of international cooperation efforts or can even operate under transnational agreements (Council of Europe, 2001).

A detailed analysis of Computer Security Incident Response Teams Organizational Models can be found in Killcrece et al. (2003). Van Wyk and Forno (2001) summarize the strengths and weaknesses of the various types of Computer Incident Response Teams as in Table 1.

In addition to Table 1, there is a growing debate regarding the role of CSIRTs that can be best summarized in Dr Schultz's (2004) statement that CSIRTs need to change, since most of the analysis centers provide the same type of information without a really thorough examination of the specific incidents, offering to a large extent the same amount of information regarding a vulnerability, threat or attack.

When an organization decides to implement a corporate Incident Response Policy (also known as Incident Response Capability), it is of major importance that the organization clarifies specific roles and responsibilities in its jurisdiction. It is important to stress that responding to a security incident is not solely a technical issue. Therefore, people from various departments outside the IT department have to actively participate in this effort. A clearly-defined, easy to implement and execute management structure is necessary. Fig. 1 presents a management schema for Incident Response that contains the necessary parties (also known as Contacts) that should participate in a corporate Incident Response Capability along with their primary relationships and interactions.

The key-person in such an organizational-wide program is the *Incident Response Capability*

Table 1 Advantages and disadvantages of the different types of CSIRTs (Van Wyk and Forno, 2001)

Team type	Strengths	Weaknesses
Public resource (coordination/analysis center)	Low cost per organization.	Problematic to provide one-to-one support to each site.
	Good source of statistics.	Community funding model can be overly bureaucratic.
	Wide distribution of alerts.	
Internal	Dedicated service for parent organization.	Difficult to get big picture for widespread security incidents.
	Hands-on support.	Very difficult to get adequate funding and resources.
	Leverage for quantity discounts on related products.	
	Direct feedback loop for improving overall security.	
Commercial	Staffing expertise should be world class.	Cannot know the organization as well as employees do.
	Only engaged (and paid for) on demand.	Second tier support means still having to handle day-to-day incidents internally.
	Contractual leverage to get desired level of service.	Timeliness to bring team on-site when seconds matter.
Vendor	Focus on rapidly responding to product vulnerabilities.	Narrow focus.
	Availability of product design engineers to fix bugs.	Can have spin control orientation.
Ad hoc	Better than nothing.	No permanent resources.
	Can at least help get key players involved during crises.	No management buy-in.

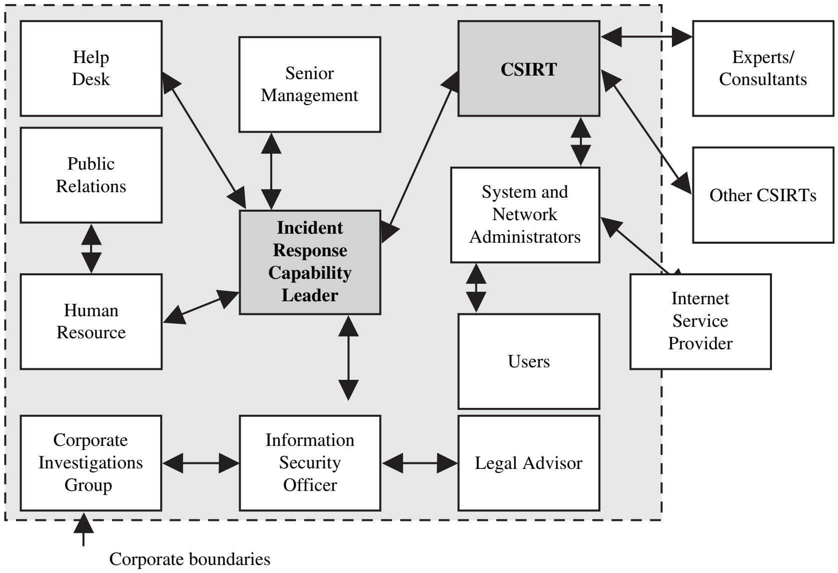
Help Desk
Senior Management
CSIRT
Experts/ Consultants
Public Relations
System and Network Administrators
Other CSIRTs
Incident Response Capability Leader
Human Resource
Users
Internet Service Provider
Corporate Investigations Group
Information Security Officer
Legal Advisor
Corporate boundaries

Figure 1 The Incident Response Capability Contacts.

Leader of the organization who should refer directly to the *Senior Management* of the organization. He/she should cooperate effectively and efficiently with all other managers in order to make them aware of an incident occurrence and to facilitate management decisions to be made. This manager can, in some cases, direct the *Computer Security Incident Response Team* (*CSIRT*), which is responsible for designing, implementing and updating the technical solution, the procedures and all the necessary guidelines for maintaining the Incident Response Capability of an organization, as described earlier. The *System and Network Administrator*(*s*) are the people with the most technical knowledge within the organization, since their everyday jobs include design, installation and fine-tuning of systems and networks and they can provide a very useful feedback in case an incident happens. They could either belong to the CSIRT or they could be alarmed only when there is a significant need. However, they need to be aware of any instance regarding a security breach. Last, but not least, the *Help Desk* personnel should participate in the Incident Response Capability, since there could be cases, such as a denial-of-service attack on the company's public servers (e.g. Web Server, FTP or Email), where the organization should answer relevant enquiries.

Apart from IT-related sections of an organization, there is an emerging need for cooperation with other departments as well. For example, the *Public Relations* department could be responsible for handling the corporate public image after an incident occurs. This task involves communication with third parties, like contractors, media agencies, etc. Furthermore, the *Human Resources* department could participate in the Capability in order to take the appropriate actions when an internal incident (caused by an employee) happens. The role of the *Corporate Investigations Group* is also important since it has to keep the fact secret for as long as it is needed, facilitate the tracing of the incident to the responsible party and not allow information to flow outside the organization that could affect its public image. This Group is often in close cooperation with the *Information Security Officer* as with all other security-related personnel who investigate security incidents and communicate with the *Law Enforcement Agencies* in case a serious incident happens (e.g. theft of IT equipment, unauthorized copying of proprietary software and data or whenever required by the law). The *Legal Advisor* of the organization is also essential in the Incident Response Capability development, being the person in charge of compliance and provision of legal advice during the various phases of Incident Response. Finally, the corporate *Users* should be trained to react according to the policies, procedures and guidelines. In many occasions, the security incidents are discovered by ordinary users; a correct and rapid communication is essential in a timely and effective response, especially in cases of internally-caused incidents. Therefore, the users have to know how to react in the first place and with whom to communicate when such an incident occurs.

Finally, according to the severity of a security incident there may be a need to cooperate with external parties that can play an important role in the discovery and response phases of a security incident. For example, the *Internet Service Provider (ISP)* can provide useful information when trying to trace a network connection, since the ISPs are providing the link of the corporate systems and networks to the outside world (be it the Internet, Extranets, etc.). Finally, other Computer Security Incident Response Teams (CSIRTs) and/or *external security experts* can provide extremely useful services when an incident is beyond the capabilities, knowledge and training of an organization's CSIRT.

A methodology on Incident Response

Apart from a well-defined and communicated management framework, both the literature and the commercial world provide well-structured methodologies consisting of several distinct phases to isolate an incident and appropriately respond to it. The most important of these methodologies are the ''Framework for Incident Response'' by the Information Security Team of DePaul University (Information Security Team, 2002), the ''Handbook For Computer Security Incident Response Teams'' by the Carnegie Mellon Software Engineering Institute (West-Brown et al., 1998) and the ''Computer Security Incident Handling Guide'' by NIST (National Institute of Standards and Technology, 2004). The aim is to 'patch a hole' and minimize the magnitude of incidents' effects while trying to trace their source(s).

However, it is necessary to comment on the fact that, until now, there is no complete de jure worldwide Incident Response standard, either as a dedicated document or as part of larger, enterprise-wide Information Security Standard. One can only find specifically designed techniques.

For example, all US Federal Agencies are obliged to provide an Incident Response Capability to comply with the OMB's Circular No. A-130, Appendix III (OMB's, 2005) as well as the Federal Information Security Management Act (FISMA) of 2002 (United States Code, 2002). On the other hand, ISO/IEC 17799:2000 (section Escalation level 2), denotes that Incident Reporting and Handling are essential for an organization but does not provide specific information on this (International Standards Organization, 2000).

Most of the Incident Response methodologies are strongly combined with the science of digital forensics, i.e. the processes of unearthing data of probative value from computer and information systems (Mandia and Procise, 2002). Forensics include the necessary actions to trace-back a security incident to its actual source and, in most cases, the physical person(s) who caused the incident. Forensics require a strong understanding of the network protocols and operating systems and demand not only a high degree of patience but also an ability to follow law-related rules and procedures.

Even though the field of forensics has been developed and has matured in Law Enforcement Agencies and not in the academic research world (Yasincac and Manzano, 2001), it seems that best practices and methodologies can be successfully transferred to the Incident Response domain.

In order to combine efficiently and effectively appropriate actions when an incident occurs, most well-acknowledged methodologies include several distinct parts. A typical sketch of such a methodology is shown in Fig. 2, based upon the phases proposed in National Institute of Standards and Technology (2004) and Patsos (2002). Example recommendations and best practices are proposed throughout the following sections based upon the existing literatures National Institute of Standards and Technology (2004), Patsos (2002), Allen (2001), Kossakowski et al. (1999). We limit ourselves only to the absolutely necessary actions that should be taken in every phase, without exhausting all possible options and alternative solutions.

Preparation phase

Considering that necessary security mechanisms are already in place both at the corporate gateway (firewalls, antivirus software, strong authentication mechanisms, etc.) and at critical internal parts (Host and Network-Based Intrusion Detection Systems (IDS)), additional specific software and hardware have to be installed as well (e.g. sniffers, audit log consolidation software, backup software, etc.). In the preparation phase, useful actions include the following.

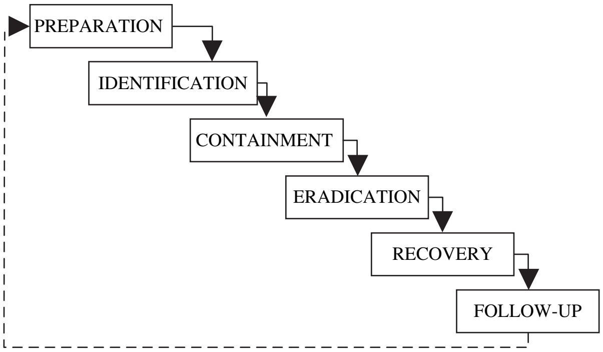

Figure 2 Incident Response Methodology model.

Operating systems and boot disks archive

It is important to maintain a complete archive of all different versions of the operating systems installed in the corporate systems. Moreover, original or trusted distribution media of any third-party software and applications that have been installed should also be archived to be available when needed. Furthermore, boot disks of all operating systems installed should also be archived. Boot disks allow system restart from a well-known (pre-existing) configuration and provide assurance that compromised files or malicious software has not been not loaded.

Security patches archive

All operating systems and applications have some (software) vulnerabilities. When vulnerabilities are discovered, the vendor is obliged to issue a security-related patch to address it. The common practice to download patches is through the vendor's website. It is very important to download and keep an archive of all security patches for every piece of software installed on the systems, along with their description and their cryptographic checksums. Apart from ensuring that no malicious code is included in the patch, a correct cryptographic checksum also verifies its source (see section Cryptographic checksums of critical operating system and application files).

Reinstallation tools

In case system reinstallation is needed it should be performed from a trusted source. When a number of hosts (usually workstations) need simultaneous reinstallation this could be performed with the use of a network server that contains generic versions of the whole system (probably images) and tools that are able to retrieve, unpack, verify and install software patches as well.

Backup procedures

Data backups are needed so that when an incident occurs damage to data is minimized. A very important issue when backing up or restoring information systems' configuration is data integrity; the system administrator(s) has to be sure that no malicious code is contained in the data backed-up or the data restored. An off-site backup copy that is held and protected away from the corporate premises should also be maintained.

Cryptographic checksums of critical operating system and application files

When installing a new system it is essential to record the cryptographic checksum of critical operating system and application files. This will provide assurance when rebuilding a compromised system. These checksums should be stored in special media (e.g. CD-ROMs) to make modifications impossible.

Backup media

During the response process backups are used to restore from the latest available configuration of the system and the last version of the data files. Backups also preserve the evidence (in case the organization decides to investigate the incident) and are useful when an isolated (or test) system is analyzed. Recommended media are optical disks (CD-R or DVR ± R) that provide relatively high capacity (650—700 MB and 4.7—9.4 GB, respectively) and prevent from accidental or intentional modification. Other options include data tapes that can store huge volumes of data when capacity is a major concern (120 GB or more). Separated volumes of a Storage Area Network (SAN) can also be very helpful in this situation.

Resource kit

The resource kit should contain all the necessary tools that may be needed by the CSIRT during the response process, e.g. disk imaging, file comparison, configuration editors, cryptographic checksum builders, etc. The importance of the resource kit should not be underestimated as it saves valuable time when handling an incident. Write-protected media (or optical disks) are preferred in this case as well. When handling an incident, special hardware may be needed as well. For example, a laptop computer can be used to monitor the network traffic or perform a virus scan in certain systems. A Resource Kit Form may prove useful information when creating a resource kit of all software and hardware devices. It also provides assistance when questions regarding a specific host's configuration arise. These forms can be archived in printed media or stored in electronic format (e.g. on a CD-ROM, an internal Web Server or a Database).

Test systems and networks

If there are sufficient hardware resources available, isolated (both physically and logically) systems and networks similar to the original ones should be maintained. In case of a system compromise these test systems can provide a smooth transition and allow for continuous operation while the compromised systems can be used for evidence collection and analysis. If this is not the case, the CSIRT must have the capability to create an ''ad hoc'' test environment. It is important that all entities of the test environment must have the same security level as the original (production) ones.

Audit trail

An audit trail of all systems should be kept and protected. The complexity of today's systems and networks leads to numerous audit logs. A recommended practice is a central System Log Server (often called Syslog) that keeps a copy of all the audit information that is saved locally in every system: firewalls, Intrusion Detection Systems, routers, hosts, etc. This system does not need to be of the latest technology; therefore it can be inexpensive but it should be isolated and protected. On the other hand, all traffic directed to this server should be encrypted if possible, to prevent an attacker from reading information in transit. Moreover, this system should be hardened to eliminate access to an attacker. The operating systems and application hardening procedures limit the system configuration only to the necessary services a system provides. Finally, the use of time synchronization protocols, like NTP (Network Time Protocol), is of major importance (Internet Engineering Task Force, 1992). All information of this kind should be synchronized with a trusted time source making it possible to use the audit trail as evidence later.

Recent technological advances provide unified event log management systems, known as Security Information Management Systems. These systems are configured to collect, normalize, aggregate, correlate, report and archive security events written in various audit log files (e.g. firewalls, routers, systems). The most important benefit of a Security Information Management System deployment within an organization is the forensic information they can provide, by allowing specific queries regarding the source and destination as well as the time and type of an incident.

Apart from Security Information Management Systems, ordinary printers can provide great help in certain cases. For example, a printer can be locally attached to a system under attack in a way that every interaction with the system is immediately printed. This method prevents attackers from altering their recorded actions (like they could possibly do with audit log files) and, moreover, printers cannot be detected by attackers (Stoll, 1990).

Identification phase

The identification stage is of crucial significance: this stage can identify the starting point of an event and this is the stage when critical decisions have to be made to categorize an event and respond accordingly. If procedures fail in this stage, the whole methodology may collapse and be of no use.

The beginning of evidence collection should start immediately after the identification or even the suspicion of an incident. The decision, though, whether abnormal activity corresponds to an actual attack or an attack pattern is quite tricky. Technology offers assistance through a variety of methods via Intrusion Detection and (Near) Real Time Threat Management Systems that require a wide deployment into corporate networks. It should be mentioned though that in most cases it is the human factor that has the knowledge on what consists abnormal activity in a specific corporate environment. Two main approaches regarding network incidents depending on their severity can be identified:

- immediately close the attacker's point of entry and eliminate all possible access means, or;
- remain 'open', as long as possible, and gather as much information as possible to be used later as evidence.

Audit log collection, examination and analysis

Perhaps the most tedious and tiring tasks are audit log collection and analysis. Since information about an incident can be found at various sources (e.g. firewall(s), IDS(s), router(s), etc.) a great amount of effort and time is required to correlate them before reaching trustworthy conclusions. Even worse, if the working environment is heterogeneous with respect to systems and network technologies, audit log information comes in various formats and sizes.

The importance of having a central System Log Server that performs a log analysis is based on the observation that using one central system and applying filters can provide useful conclusions in a relatively short period of time. However, if this system does not exist, an alternative manual approach is needed.

A high level process and the respective possible actions are proposed in Fig. 3, while the necessary steps are briefly described in the following paragraphs.

Both Host and Network-Based Intrusion Detection Systems maintain a large database with patterns of well-known attacks (known as ''signatures''). This database can provide useful information regarding the type as well as the origin of an incident. If the incident is a known one, i.e. belongs to the IDS database, the affected system

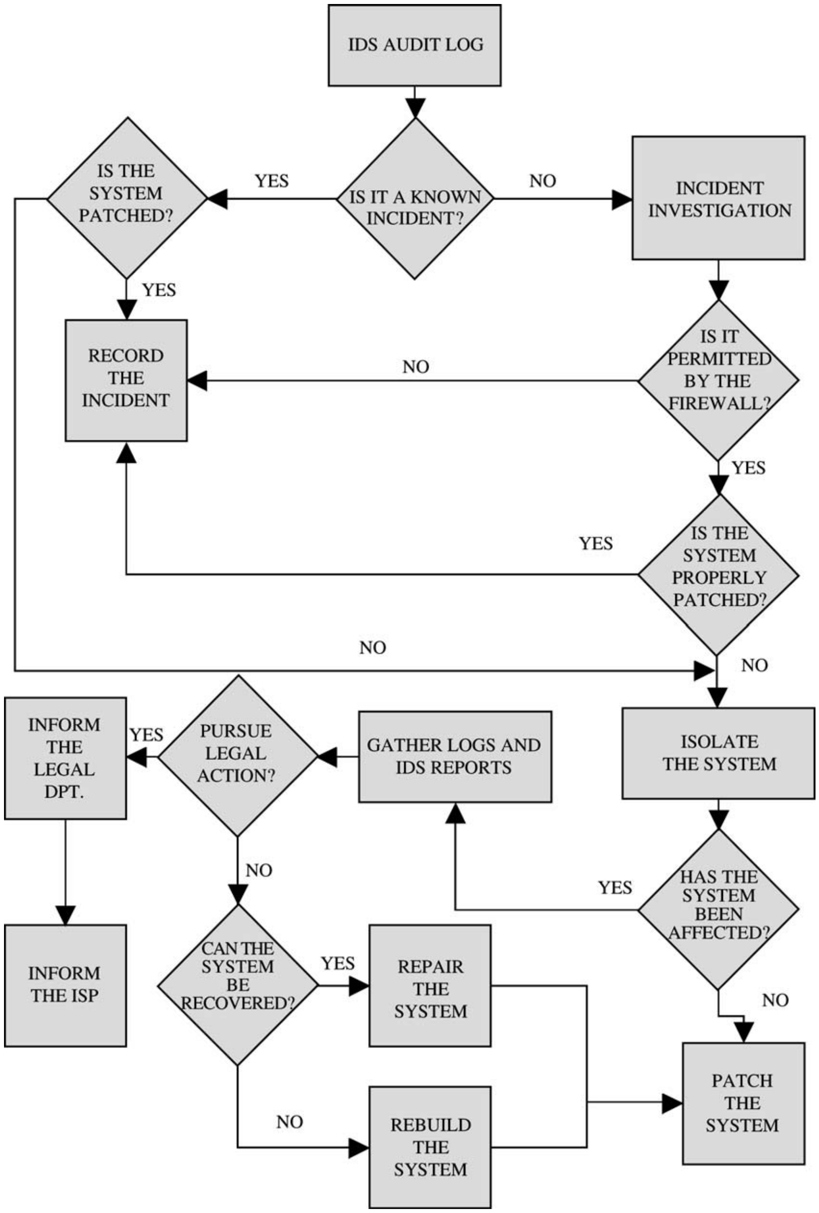

Figure 3 IDS Log Analysis and possible actions.

should be checked against the vulnerability that caused the incident. If the system has the appropriate countermeasure against this vulnerability (usually a software patch) then the incident is recorded for archival purposes.

If the incident is a new one, then the security gateway audit logs (usually the firewall logs) should be correlated in order to provide more useful information. If the firewall prohibits the connection then the IDS alarm is probably false. In this case, the incident has to be recorded and the whole process stops. On the other hand, if the firewall permits this connection, then the related system(s) should be checked against the corresponding vulnerability, as said before, using the information provided by the IDS. If the system is vulnerable, it should be isolated in some way in order to minimize the possibility that the incident propagates to other systems and networks.

Two major issues arise when a system is found vulnerable. At first, all possible electronic evidence should be collected in a forensically sound manner to be handed for analysis at a later phase. Secondly, the affected system should be recovered either by applying the appropriate patch available (see section Containment phase), or by the recovery and backup media (see section Eradication phase). If no such option is possible, the system should be rebuilt from scratch (see sections Eradication phase and Recovery phase).

Incident reporting and assessment

As soon as an incident is identified it should be categorized in one of the three possible security levels based on the impact of a specific incident. The impact is expressed in terms of financial impact, impact to manufacturing, sales, corporate image or impact to trust by the company's customers. The scope of the impact will probably determine the initial response. A draft listing of these security levels as well as a definition/description of each one of them is proposed in Table 2.

Classification of the incident in one of these categories is part of what is called the Incident Reporting Form. In this form all relevant to the incident information is documented. The importance of this form is very significant, since it contains useful information that is reviewed later during a forensics analysis or during the follow-up phase (see section Follow-up phase). Example information includes:

- date and time of reporting;
- date and time of incident discovery;
- the system in which the incident was first identified;
- possible systems and networks that may be affected;
- system configuration, host applications and criticality;
- name and other credentials of the person that completes the form.

System information

All network connections, running processes and open files as well as the active users that are logged onto the network to give an overview of the state of the system should be recorded. Use of the IDS information in parallel to system and disk imaging utilities could possibly achieve this.

Table 2 Classification of incidents based upon severity and impact

Security level	Description
High	The impact is severe. Examples include: disruption of the corporate network, compromise of confidential information, a virus that has affected a large number of systems, etc.
Medium	The impact is significant. Examples include: delayed delivery of electronic mail and other services, exploited vulnerabilities that do not affect the overall system operation, etc.
Low	The impact is minimal. Examples include: harmless SPAM email, isolated virus infections, etc.

Disk imaging of the affected system(s)

A disk image (also referred to as a bit stream image) is a bit-by-bit copy of the system affected by a security incident. This imaging records the active state of the system at the time this operation is performed (snapshot). It is recommended that at least two images of every compromised system should be taken and saved in hardware-write-protected or optical media.

Common practices indicate that the imaging process is performed after a system is compromised to assist during the computer forensics analysis but, in many cases, ''live'' system images can be found in servers that incorporate mirroring techniques (e.g. RAID arrays). These (image) copies can be later reinstalled on test systems for further analysis. One of these copies should be used as evidence. Therefore, it should be sealed and stored in a secure location and not used for any operational task.

A very important issue when an organization decides to perform a forensics analysis in order to find and prosecute an attacker is to create and carefully maintain a chain of custody. Furthermore, image copying should be done with great care, as unusual disk activity could alert an intruder who has gained access to the system(s).

Other systems analysis

In a networked environment usually more than one system is affected by an incident, therefore various other systems should be examined thoroughly as well, including:

- systems at the same IP address range;
- systems at the same network segment;
- systems in the same network domain;
- other critical systems.

In a remote attack, for example, the usual point of entry of the attacker is the HTTP (Hyper Text Transfer Protocol) server or the FTP (File Transfer Protocol) server. In the case of a computer virus, worm or malicious software in general, it is recommended that the SMTP (Simple Mail Transfer Protocol) server should be examined in the first place.

Containment phase

The next step is to apply immediate solutions, thus limiting the extent of the incident and letting the attack function only up to the desirable extent. Not all attacks should be immediately stopped, since — for certain reasons computer forensics analysis may be needed afterwards. Common techniques include patch installations and configuration changes into critical perimeter, public and internal systems.

Disabling of specific system services

The results of the log analysis in combination with the Incident Reporting Form should provide some information about the scope of the incident as well as the possible vulnerability that was exploited.

Therefore, system services related to this vulnerability should be disabled until a patch is issued, downloaded, verified and installed.

By disabling specific services the problem is isolated, there is enough time to address it properly and the system continues to operate providing a significant part of its services. However, if this vulnerability is not a known one, a software fix may take a long time before it is issued. If this is the case, either the service must remain disabled until the 'patch' is issued, or ad hoc methods should be followed (e.g. allow certain commands upon a service like FTP-get but no FTP-put) according to the nature of the problem.

Changing of passwords and account disabling

This is the usual task all the administrators perform when a network incident occurs, since compromised accounts are the usual point of entry of an attacker. If this is the case and the intruder has not used some other means to gain access (e.g. a vulnerability that bypasses user authentication) this practice will terminate its access path.

On the other hand this change provides the attackers with a notification that someone is watching, a situation that may not be desirable if the organization has decided to remain 'open' to trace the incident. When there is no other alternative, advanced security techniques are used (e.g. use of honey pots).

Disconnection of the compromised system from the network

In serious network incidents the need to disconnect the compromised system from the network may arise in order to prevent the attacker from gaining access to other systems (escalation attack) or from taking complete control over this particular system.

Before doing this, the scope of the incident has to be carefully estimated in contrast with the acceptable level of risk faced by business processes. It needs to be stressed here that the Incident Response Capability should clearly define who is empowered or authorized to take this decision, especially when an incident takes place in non-working hours.

Temporary shut down of the compromised system

If there is no alternative option the compromised system should be shut down. This shut down will prevent further damages to the specific system (as well as to other systems) and will also provide some time for an initial analysis but such a shut down will deny access to corporate users that need these connections. Moreover, all available memory contents will vanish, thus destroying any possible evidence that may reside therein (e.g. running programs of the attacker). This practice is not recommended and when performed it should be done with great care and only for a limited time period. If there are no alternative solutions, a memory dump that contains memory contents should be written to the disk prior to the shut down process in order to preserve any information that could be used later as evidence.

Restoration of the compromised system

When there are enough hardware resources available a clean backup copy of the system should replace the compromised one. However, it has to be made sure that the incident's source (i.e. vulnerability, access path, etc.) is identified and the system is patched. This will limit the systems' downtime and will provide enough time to perform a thorough analysis of the compromised host without causing operational problems.

Eradication phase

This stage refers to the mid and long-term solutions that have to be applied on the affected systems in order to eliminate any possible means for recurrence of the specific attack. Possible countermeasures at this stage include policy compliance checks, independent security audits, policy updates, etc.

Changing of all passwords in all compromised systems

Even if the comparison with the cryptographic checksums that were recorded during the preparation phase does not indicate anomalies, it is not uncommon to have all password changed in all systems affected. This is a time-consuming task performed by the administrators but this practice provides an additional assurance both to the systems administrator and to the management as well.

Complete elimination of intruder access and identification of possible changes

Attackers usually leave backdoor programs or Trojan horses behind to be able to gain access at a later time. Initially, a comparison of the critical system files to the cryptographic checksums should be performed. This comparison guarantees that no modification has been made during the incident. If time allows and if cryptographic checksums were recorded during the preparation phase, the same comparison should be made regarding the data. Although this is not a common practice, it may provide useful information.

Complete reinstallation of the compromised systems

A complete reinstallation of all the compromised systems is among the best possible practices. It may be performed faster when restoring from bit stream images (taken during the preparation phase) rather than from ordinary system backups. It should not be forgotten though that the compromised systems need the necessary patches and fixes before they are put back to production.

System rebuilding

If backups and bit stream images have not been taken during the preparation phase, the system administrator can restore the compromised systems from the original distribution media. This is probably one of the most discouraging and time-consuming tasks, but also one that provides the highest possible level of assurance. This system rebuilding should be definitely done in the ''background'', in a backup system, or could be transferred to the next phase of the Incident Response Capability if the effects of the incident allow for this. A fast system restoration from backups (or images) is preferable in order to guarantee business continuity and a good level of assurance, giving the system administrator the opportunity to build the compromised system from scratch at a later time.

Recovery phase

After all the previous steps have been successfully followed, system restoration and security mechanisms enhancements should begin in order to bring the whole system back to production without any security holes open. Examples include complete system rebuilding, data recovery from backup media, installation of extra security mechanisms, etc. Before putting compromised systems into production again, it is advised to run a vulnerability assessment or a penetration test in order to disclose possible existing vulnerabilities.

System rebuilding from scratch

As it was discussed in the previous phase, a complete rebuilding of the system should be done during the recovery phase, where a larger time frame permits the smooth execution of this heavy task. Furthermore, incident analysis may provide feedback to the technical staff of the CSIRT enabling them to determine the desired security level.

Restoration of user data from trusted backups

Acting proactively can save precious time when restoring from backups. Original Cryptographic checksums guarantee data integrity and authentication. However, the recording of cryptographic checksums for every available system turns out to be a time-consuming task especially in cases where data content is frequently updated (e.g. a financial database) or is dynamic (e.g. fetching web pages from another source). Backups made closest to the time of the incident should be carefully used to provide both a high-level of confidence and a quick restoration time.

System configuration review — auditing

A convenient way to provide the maximum level of confidence is to perform a system audit at the compromised hosts. Apart from the maximum level of assurance, auditing procedures will also discover any possible mistakes or omissions during system rebuilding (residual risk). If the systems' restoration were performed during the eradication phase, it should not be forgotten that most of the operations were performed under heavy stress; possible omissions or mistakes may still exist.

Review of the protective and detective mechanisms

In case a network incident has happened, the firewall's configuration and rule base should be reviewed and updated. Security is not a static process; attack methods are evolving and security administrators have to constantly keep up with them. The same should be done with the IDS configuration in order to enable better monitoring and reporting of security-related activities. The intelligence of the IDSs relies on their proper configuration and the continuous updating of their signatures' database.

Follow-up phase

Last but not least, all actions and information concerning the incident should be documented and electronic evidence should be disseminated for analysis in a forensically sound manner to experts. Furthermore, a post-mortem meeting with Senior Management should take place in order to assess the damage done, the policies' strengths and weaknesses and the procedures need to be followed. The aftermath of an incident may indicate or even require that security policies, procedures and guidelines have to be updated in order to be able to address future attacks of the same type. After the complete analysis of the incident is performed, changes in system configurations must be documented and the inventory of systems and network assets has to be updated to reflect these changes.

Tracing back a security incident

Handling a security incident may only be the beginning in the Incident Response process, since, there could exist some occasions where the source of the incident has to be identified (i.e. the actual attacker(s) should be found in order to be held accountable for their actions). The task though is not easy since the attackers use impersonation techniques to cover their traces and hide their identities. Impersonation attacks (also known as spoofing or masquerade attacks) can be reproduced in various different ways. The most well-documented impersonation procedures are:

- MAC address spoofing (Whalen), i.e. impersonating of a user's MAC address;
- IP spoofing (Bellovin, 1989), i.e. impersonating of a user's source IP address;
- application layer spoofing, i.e. impersonating of a user's identity.

As the results of several Distributed DoS attacks indicate (Lemos, 2001), the attackers often use intermediate hosts and networks (in order to obtain different routing paths) before launching an attack. Besides that, they may also use some intermediate compromised hosts (known as stepping-stones) that act as conduits and change the essence of the attack process by using, for example, encryption (Zhang and Paxson, 2000). The reconstruction of the attack path back to the attacker, who has used one of the previously mentioned impersonation techniques, is not straightforward. There are many ways that the correct reconstruction of the attack path (from the victim machine back to the attacker machine) is prevented. In general, let $C = h_1h_2...h_i\ h_{i+1}...h_n$ be the connection path between hosts h_i ($i = 1,...n$), then the trace-back problem is to recursively identify the actual IP address of h_{n-1}, ... h_1 (attacker), when the actual IP address of host h_n (victim) is given.

Before proceeding to the essence of automated trace-back mechanisms, it is important to define the issue of network tracing, that — until now — has been used by network engineers to troubleshoot routing functions and protocols. As the relevant RFC denotes (Postel, 1981), the IP protocol provides the *Record Route* option in the protocol header that supports network tracing. When this option is used the routing devices along a path are mandated to append their IP addresses to IP options portion of the IP header. The protocol header has a fixed part of 20 bytes and a variable part (corresponding to the IP Options Field). The total length of every IP packet in 32-bit words is mandated by the value of the 4-bit IHL (IP Header Length), which is 15 (1111 in binary), resulting in a maximum of 60 bytes, of which only 40 bytes are available for IP Options. With the current size of the Internet and the heavy routing information used in modern networks the *Record Route* field appears rather limited to withstand recording for every hop a packet traverses.

Apart from the tremendous amount of processing overhead in routing devices since — at least — 32-bits of information have to be appended to data in flight in every routing device, a wily attacker can use another option in the IP header options field (e.g. the *Loose source routing* that mandatory defines a list of routers that should not be missed during routing), ''invent'' additional hops in the path and fill the 40 bytes available for IP options with false or misleading information.

IP marking trace-back

In order to counter the limitations mentioned in the previous paragraph many automated trace-back mechanisms techniques have been developed in order to enable routers to probabilistically mark packets and therefore make the reconstruction of the complete path easier (Savage et al., 2000; Song and Perrig, 2001; Park and Lee, 2001). Mechanisms that rely on IP marking use very complicated mathematical algorithms in order to identify the origins of IP packets, especially of the spoofed ones. These mechanisms have proved robustness, effective probability rates and — relatively — scalable deployment. On the other hand, they require that all traffic is in clear text, so an obvious issue arising is the compatibility with IPSec (Kent and Atkinson, 1998) or other network encryption methods that prevent routing devices from appending marking information and achieve trace-back.

ICMP-based trace-back

Apart from IP marking techniques, the research community and the IETF are evaluating a series of ICMP-based trace-back mechanisms. The current IETF standard is the *iTrace* scheme proposed by Bellovin (2003), which is based on the capability of routing devices to generate a ''trace'' packet for every packet they forward (and marked for tracing). Both the original and the ''trace'' packet are collected at the destination host and the route is reconstructed. The *iTrace* scheme supports HMAC (US Department of Commerce, 2002) and the X.509 protocol for authenticating and evaluating the *iTrace* messages (Adams). In the current revision of the standard, the number of *iTrace* packets generated by a router is relatively small, resulting in a low overhead (statistically, around 0.005%) (Mankin et al., 2001). However, this scheme requires a significant amount of traffic from a small number of sources, due to the lower probability of generating *iTrace* packets (Savage et al., 2000). An enhancement to Bellovin's approach, known as the Intention-driven *iTrace* (Mankin et al., 2001), is based upon the addition of one extra bit (called intention-bit) in the routing and forwarding process as well upon the functionalities provided by the Border Gateway Protocol (BGP) (Rekhter and Watson).

IP tunneling trace-back

The *CenterTrack* schema introduces the concept of some special types of routers (called tracking routers (Stone, 2000)) that create an overlay network (Fig. 4). These routers have a physical or virtual adjacency with the edge routers of an autonomous system. In turn, all edge routers are linked to a central tracking router (or a simple network of tracking routers) via IP tunnels and create an overlay network. *CenterTrack* requires that all edge and tracking routers are able to perform input debugging functions. If this is not the case the use of network ''sniffers'' for traffic analysis and attack pattern recognition is allowed. The malicious traffic that is originally destined for the victim is routed through the overlay network via dynamic routing protocols, thus initiating a

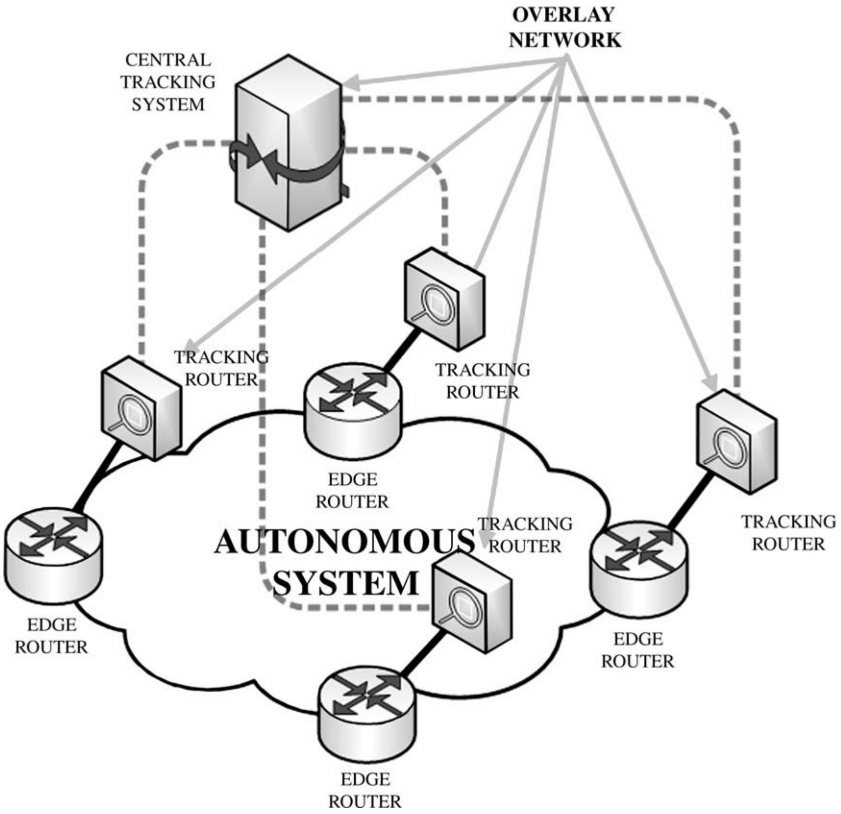

Figure 4 CenterTrack simplified.

hop-by-hop tracking, starting from the router that is closest to the victim.

In the entire process, a significant number of static routes (both in the egress and the ingress routers closest to the victim) must be configured so that attack traffic flows explicitly through the overlay network, while — at the same time — legitimate traffic is not intercepted. The last fact is quite complicated since in general it is very difficult to filter and reroute certain volumes of traffic that match specific attack patterns. A drawback of this model is that an attacker can detect the presence of tracking systems by statistically measuring the observed latency via fragmented packets sent to the victim during the information gathering phase of an attack (McClure et al., 2001), exactly the same way sniffing programs or IDSs are detected. Besides this, an attacker can cause a DoS to the *CenterTrack* (which is a single-point-of-failure). Last but not least, if the attack target is the edge router itself then the system would try to reroute traffic destined to the edge router through this specific edge router, resulting either to tunnel collapsing or to endless routing loops. An alternative suggestion using the concept of an overlay network has been proposed by Baba and Matsuda (2002).

Host-based trace-back approaches

Historically, the first trace-back mechanisms proposed aimed to identify the hosts that formed the connection path. The most accredited of these mechanisms are the *Caller Identification System (CIS)* and the *Caller ID* system that are briefly explained in the following. *CIS* is a trace-back system that intends to identify the attacker through the login process (Jung et al., 1993). Its operation relies on the information exchanged during the login process within the systems involved in a connection chain. So, when a user from host h_1 connects to system h_n ($n > 2$) through some intermediate hosts ($h_2,\ldots h_{n-1}$) the h_n-th system recursively queries the h_{n-1}-st host about the login information. Therefore, for every system where a user logs-in, all previous login information is checked before access is granted (or denied). Important drawbacks of this concept are the rather outdated method and the native vulnerabilities introduced by authentication techniques. Apart from that an important overhead in the login process is added, which can detected by the attackers. *Caller ID*, proposed by Staniford-Chen and Heberlein (1995) introduces a manual trace-back in every intermediate host of the connection chain. When an attacker connects to system h_n through h_1, h_2, $h_3,\ldots$, h_{n-1}, the system owner or security personnel breaks into h_{n-1} to verify the origin of the connection, most often by using techniques used by attackers. He then breaks into h_{n-2} until reaching h_1, which is the attacker's machine. Apart from the ethics and legal complications, this method does not introduce important overheads like *CIS* and could be possibly scalable to cross-administration domains or even the Internet. However, considering the current high-speed networks the manual processes that have to be performed for every host traced make this approach non-applicable. However, it has been reported that the US Air Force uses this method (Wang et al., 2001).

Application-based trace-back

A very recent and promising research effort in automated intrusion response has resulted in the *Intruder Detection and Isolation Protocol (IDIP)* (Schnackenberg et al., 2002; Feiertag et al., 1999) (Fig. 5). The *IDIP* is currently being scaled to multiple administration domains across the Internet, since it is featuring low cost integration with current intrusion detection techniques by adding new response mechanisms and algorithms. IDIP is based on the *Common Intrusion Specification Language (CISL)*, developed by the Common Intrusion Detection Framework (CIDF), providing a unified explanation of security incidents (Feiertag et al., 1999).

Recent results have shown that the protocol is performing well and is integrating smoothly with IDS systems within the DARPA research community (Schnackenberg et al., 2002). Although *IDIP* is currently supporting a minimal set (''block'' and ''allow'') through the *CISL*, these actions can be performed against a variety of different objects (e.g. users, applications, processes, connections, states, systems, etc.) so that the overall combination of responses provides a significant number, therefore facilitating granular policy development.

Performing a forensics analysis

Further to the trace-back mechanisms that allow the discovery of the host(s) that participated in the attack, the overall goal of a forensics analysis is to develop comprehensive techniques to mirror a computer host to an actual entity.

Computer forensics

Computer forensics is the science that is dealing with finding appropriate evidence in end systems

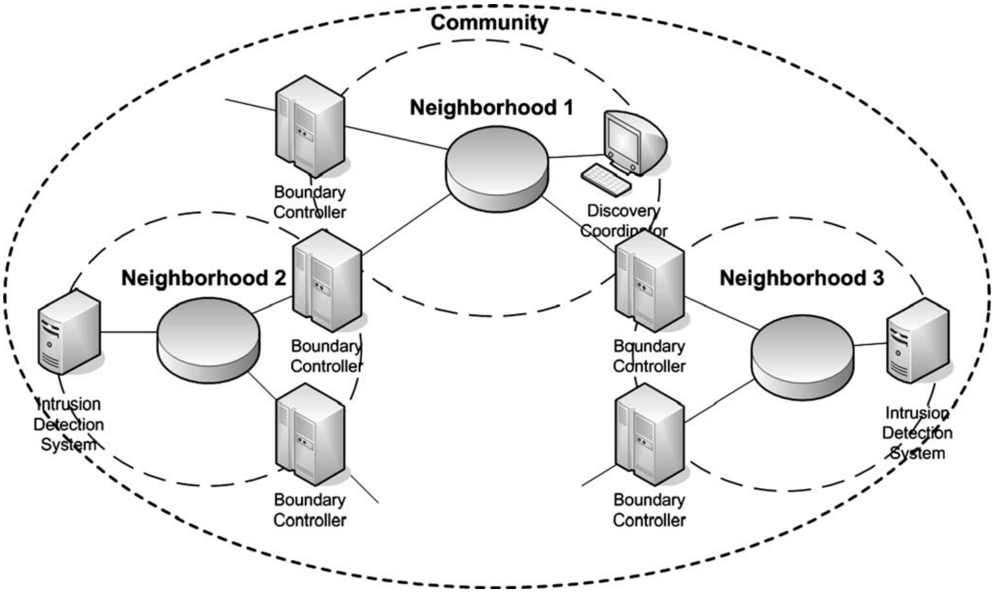

Figure 5 The *IDIP* architecture (Schnackenberg et al., 2002).

(probably both in the attacker's and the victim's machines). Computer forensics is about the ''preservation, identification, extraction, documentation and interpretation of computer data'' (Kruse and Heiser, 2002). The objective is to identify the actual programs that were used during an attack as well as to simulate the entire attack process. Apart from being an extremely difficult task, since there are millions of files that should be analyzed, this process requires a deep knowledge of the underlying legal framework. In many cases attackers are using spoofing techniques to hide their original identity, thus misleading a forensics expert. Last but not least, since international law is not harmonized the whole process can be driven to an endless loop. Computer forensics techniques require specialized software to analyze and correlate miscellaneous files and logs, so as to provide the big picture of the security puzzle. Possible actions to identify sources of evidence include (Berghel, 2003):

- decrypting files;
- decompressing data;
- cracking passwords;
- bit stream imaging of hard disk volumes;
- analyzing Free & Slack Space of one or more hard disks;
- examining actual applications files;
- reconstructing the 'swap-file', etc.

We have to note that even if such an analysis may lead to the actual attacker, it is often difficult to monitor these computer actions in such a way as to be able to stand in the Court of Law. Moreover, it is even more difficult to explain to the Court of Law how the whole attack has been carried out. The latter makes the importance of using forensically sound techniques and tools extremely important, as a minor aberration can destroy all the previous efforts. Apart from computer forensics that have been extensively analyzed in scientific literature other important forensics techniques include network and software forensics.

Network Forensics

Network Forensics (also known as Internet forensics (Berghel, 2003)) primarily deal with the data found across a network connection (mostly ingress and egress traffic from one host to another). We have to state that the major difference with the automated trace-back mechanisms that were presented in the previous section is the fact that Network Forensics deal with the reconstruction of the attack path long after an incident has been occurred, mostly for accountability reasons and legally-driven purposes. Furthermore, one major difference with computer forensics is the actual data that both examine.

While, on one hand, computer forensics deal with data that are to be found on a system after an incident has occurred (data that is easily cloned and therefore examined), Network Forensics analyses ephemeral and volatile data logged through specific security countermeasures (e.g. packet filters, firewalls, intrusion detection, etc.). The tools and skills used for such an analysis are almost the same as the ones used by Internet hackers. It is the ethics and purpose that make the difference. Not surprisingly, Network Forensics evolved as a response to the hacker community (Berghel, 2003). In such an analysis, possible patterns of data for analysis can be found among others in:

- the firewall log files;
- the Host and Network Intrusion Detection Systems' log file;
- other monitoring devices or software;
- the router log files;
- the users' directory.

Artificial Intelligence and Fusion techniques that are used in order to speed up the whole process as well as assisting in reaching useful conclusions have also been proposed (Nong et al., 1998). Last but not least, a Network Forensics analysis should progress in parallel to computer forensics, if such a separation is desirable when investigating a security incident.

Software forensics

An idea introduced in the early-90s by Spafford and Weeber is software forensics, i.e. tracing code back to their author(s) (Spafford and Weeber, 1992). Although this may seem extraordinary at a first glance, the concept relies upon special individual characteristics that can be found in any computer program segment, like Language, Formatting, Special Features, Comment Styles, Variable Names, Spelling and Grammar, Use of Language Features, Execution paths, Metrics, etc. Obvious limitations include code reuse, application changes and the amount of code to be examined. However, one does not intend to solely identify an attacker through using only software forensics but expects to use them in collaboration with other forensic techniques. It is worth noting that the problem of code authenticity is also related to financial crime, copyright of digital media, patent protection and computer viruses. Extensions on this concept include, Microsoft's Authenticode, a technique also relying upon the benefits of Public Key Cryptography (Microsoft Corporation), and code signing extensions introduced by Sun Microsystems's Java Programming Language version 1.1 (McGraw and Felten, 1999). However, we have to comment on the fact that both these two techniques have been developed under the prism of providing assurance on the authenticity of a piece of program rather than on providing accountability attributes.

A generic Incident Response process in a corporate environment

In this section a combination of the management framework proposed in Fig. 1 along with the methodology presented in Fig. 2 is attempted. We briefly propose the necessary actions that have to be taken in order to minimize the effects of a security incident along with the efforts that will lead to the source of the attack.

In this context, it is very important not to overlook the fact that an incident's significance is increasing as time passes (which is true for the most of known security incidents). This parameter is called ''escalation level'' and usually falls within one of the following — broad — categories:

- Level 0, where the operations are normal and there is no evidence that a security incident is occurring.
- Level 1, where a threat is discovered and the initial responses are taken.
- Level 2, where the threat is spreading and containment actions are taken.
- Level 3, where the threat has become significant and containment along with recovery actions are taken.

The proposed process along with the necessary actions of every Incident Response Contact is depicted in Fig. 6, while it is briefly explained in the following.

Escalation level 0

The CSIRT usually monitors all known sources for possible security alerts. These sources can be logs of any proactive or detective security countermeasure (e.g. firewalls, Intrusion Detection Systems, Syslogs), various security mailing lists, web-sites of analysis centers, etc.

Escalation level 1

When a threat is identified and categorized then the corporate CSIRT documents the incident (usually through the Incident Reporting Form), determines possible initial actions and notifies the Incident Response Capability Leader (IRCL). The IRCL has to approve the initial actions and alert the Help Desk if user actions are required.

Escalation level 2

When the threat has become significant then the IRCL notifies the CSIRT about the threat impact

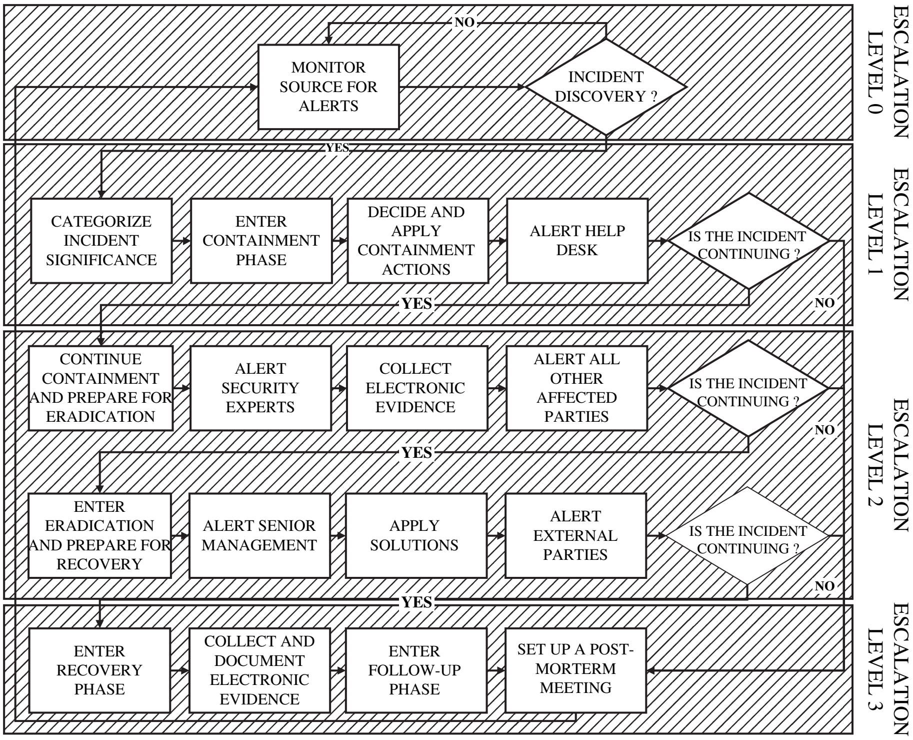
MONITOR SOURCE FOR ALERTS
INCIDENT DISCOVERY ?
NO
YES
ESCALATION LEVEL 0
CATEGORIZE INCIDENT SIGNIFICANCE
ENTER CONTAINMENT PHASE
DECIDE AND APPLY CONTAINMENT ACTIONS
ALERT HELP DESK
IS THE INCIDENT CONTINUING ?
YES
NO
ESCALATION LEVEL 1
CONTINUE CONTAINMENT AND PREPARE FOR ERADICATION
ALERT SECURITY EXPERTS
COLLECT ELECTRONIC EVIDENCE
ALERT ALL OTHER AFFECTED PARTIES
IS THE INCIDENT CONTINUING ?
NO
YES
ENTER ERADICATION AND PREPARE FOR RECOVERY
ALERT SENIOR MANAGEMENT
APPLY SOLUTIONS
ALERT EXTERNAL PARTIES
IS THE INCIDENT CONTINUING ?
NO
YES
ESCALATION LEVEL 2
ENTER RECOVERY PHASE
COLLECT AND DOCUMENT ELECTRONIC EVIDENCE
ENTER FOLLOW-UP PHASE
SET UP A POST-MORTERM MEETING
ESCALATION LEVEL 3

Figure 6 An example Incident Response process.

and alerts the System and Network Administrators. Then, in turn, the IRCL collects any reports from them and starts to collect evidence. The CSIRT determines the required containment actions (possibly in cooperation with external Security Experts/Consultants) and reports all actions taken to the IRCL. According to the incident's effects, other Contacts can be possibly alarmed at this stage (e.g. Corporate Investigations Group, Help Desk, Legal Advisor, etc.).

Escalation level 3

When the threat is severe or widely spread then the IRCL has to communicate with the Senior Management to gain support for further actions that need to follow (e.g. shut down of a specific system or a network subnet). The CSIRT and administrators decide the containment and/or eradication actions that will follow. In the entire process electronic evidence is collected and documented since the management may decide to further investigate an incident through a forensics analysis.

Postincident analysis

When the incident has been treated properly then the IRCL has to set a post-mortem meeting with representatives from all involved parties and from the Senior Management where the following issues will be addressed:

- estimation of the damage/impact;
- all actions taken during the incident;
- follow-up actions needed for the complete elimination of the vulnerability;
- policies and procedures that need updating;
- further actions that are possibly needed;
- possible electronic evidence that will be handed to the responsible senior manager;
- update of the Incident Response Capability procedures.

Open issues

The methodologies, mechanisms and measures presented in the previous sections provide the ''rich picture'' on the current state of Incident Handling and Response. However, we have to state that all those practices are closely related to corporate IT Systems and Networks, since these were the most usual attack targets until nowadays. In the last few years, experiencing the internetworking capabilities provided by 3G and broadband networks along with advances of E-Government and E-Commerce applications, there is a major shift on the target of Internet attacks from corporate systems to the ordinary Internet user. In this situation, the consequences vary from a single system restore to identity theft implications. The latter often leads to an individual inexistence for many daily operations, as described above. Identity theft, as recent surveys show, is the scourge of our time (Harris Interactive, 2003), with nearly 33.4 million individual victims in the US only. Although the term is rather legal than technical (since if someone else is forging our personal details they still belong to us), it refers to the unauthorized manipulation of personal data found in public or private databases as the product of inadequate security. For instance, personal data include name and surname, financial details, home address, penal records, military obligations, medical history, etc. These data can be found in various public or private databases. Identity theft is not the same as a regular credit card fraud for when a credit card's validity is cancelled the scenario stops. In identity theft, an attacker can use a person's details and manipulate them in detrimental ways. This type of attack is not a new one to Information Security. It has evolved throughout the years from the simple impersonation and has reached this spread due to our heavy reliance on automated systems for financial, military, governmental and miscellaneous other applications as well as due to the existence of infrastructures that support these operations (e.g. the Internet, 3rd Generation Mobile Phones, etc.). Quoting D. Solove, ''Identity Theft does not just happen. It has been constructed'' (Solove, 2004).

Although technology seems to be the obvious culprit in identity theft, since it is a reincarnation of a typical 'masquerade attack' widely reviewed in the literature, it is mostly the (international) laws and policies that cause the difficulties in an efficient response. The legal framework concerning security in the storing and processing of personal data in IT Systems is not clear (if defined) in many occasions.

The issue of how an entity should respond when a digital identity-related security breach occurs still remains open. To our knowledge, no formal methodology has been proposed so far to underline this specific issue in any Information Security Standard, RFC or best practice guide and we hope this paper will encourage heavy research in this area. Classic Incident Response methodologies are of no use in this situation, since most of their internal distinct phases cannot be directly — or indirectly — applied. A classic dilemma that lies on this is that incidents have to be prevented while privacy should be safeguarded and guaranteed. This could mean that many security products have to be completely redesigned for functional, security and individual privacy purposes.

Discussion and conclusions

In this paper we proposed a detailed management framework along with a complete structured methodology for appropriately handling a security incident, we presented the state-of-the-art technology in computer, network and software forensics as well as automated trace-back artifacts, schemas and protocols and we proposed a generic Incident Response process within a corporate environment.

Responding to security incidents has always been an important part in the science of Information Security. With our heavy reliance on computers, networks and applications for nearly all sorts of daily operations a major concern arises: privacy. Formal methodologies for Incident Response as well as forensics need to take into serious consideration the attribute of privacy. Furthermore, since new types of attacks arise containing human intelligence and exploiting intangibles, it seems that the time has come for the Incident Response to move towards the human factor. As we mentioned earlier, it is the human that lies in the heart of computer-based applications: technology merely supports our interaction. Once again, we found ourselves in classic dilemmas concerning how to implant certain aspects of our conceptual world, like trust and privacy, into technology and how to legally treat these aspects when transformed to digital. We simply stress the fact that a person's rights for or against a computer forensics analysis are not clarified yet, while they must be protected.

The art and science of Information Security experiences tremendous advantages worldwide, many of the times due to the oxymoron of failures or poor security design in several widely used modern systems. Open issues, apart from responding to identity theft, include the responses to critical infrastructure security incidents that form a critical part in most of western societies, since late or false responses could have fatal results in case of information warfare.

Incidents should be both proactively and reactively addressed to counter against a launched attack. The road of response includes advances in automated trace-back mechanisms and all types of forensics is still in their early infancy. There seems to be a fertile ground for research. Industry, on one hand, is providing good tools and resources for analyzing, determining the source and technically responding to an incident. Academia, on the other hand, focuses upon the advances of Incident Response providing analysis, coordination and further research. Finally, governmental institutions are obliged to treat all incidents with respect to privacy and issue appropriate laws and directives.

The need for international cooperation of all key players is emerging. Advances in security mechanisms can be used in order to provide solutions but they cannot be used alone. For example, embedded hardware authentication mechanisms, like the use of digital certificates hard-coded into Network Interface Cards could provide a tactical solution to the trace-back problem as they can provide both authentication and accounting for users' activities. On the contrary, they cannot provide privacy when a user wants to visit various Internet sites for all of their actions are still logged usually against their wish. In addition to this, smart cards can provide the utmost level of security but technology cannot decide whether DNA data should be included in a chip or if a smart card should be used as a national identity to counter identity theft (Hiltz et al., 2003).

The IPSec protocol was a milestone in the Information Security history. The engineers and cryptographers who cooperated in this effort managed to implant fundamental cryptographic techniques into a protocol that was designed with no security whatsoever. Perhaps the next big thing is to provide auditing and accountability attributes to the IP protocol itself, since it comprises the heart and mind of our networking world. However, at this point of time, little information or research is to be found in this extremely interesting effort.

Acknowledgements

This paper has been partially supported by GSRT under a PENED grant and by the IST FET Coordination Action ACCA (006475).

References

Adams C. Request for comments (RFC) 2510 — Internet X.509 public key infrastructure certificate management protocols. Available from: http://www.ietf.org.

Allen J. CERT guide to system and network security practices. Addison-Wesley 2001.

Baba T, Matsuda S. Tracing network attacks to their sources. IEEE Internet Computing 2002;6(3).

Bellovin SM. Security problems in the TCP/IP protocol suite. Computer Communication Review April 1989;19(2):32—48.

Bellovin SM. ICMP traceback messages, Internet draft (work in progress); February 2003.

Berghel H. The discipline of Internet forensics. Communications of the ACM August 2003;46(8).

BSI. Information security management, BS7799, part 1: code of practice for information security management; 1999.

CERT/CC. Security of the Internet [online]. Available from: http://www.cert.org/encyc_article/tocencyc.html; August 2003 [3/01/2005].

Council of Europe. Convention on cyber crime. In: European treaty series — no. 185, Budapest; 2001.

Feiertag R, Kahn C, Porras P, Schnackenberg D, Staniford-Chen S, Tung B. A common intrusion specification language. Available from: http://people.emich.edu/pstephen/other_papers/CISL-Original.PDF; June 1999.

Global Reach. Global Internet statistics. Available from: http://www.glreach.com/globstats/; June 1999 [3/01/2005].

Harris Interactive. Identity theft new survey & trend report. Commissioned by Privacy & American Business; August 2003.

Hiltz SR, Han HJ, Briller V. Public attitudes towards a national identity "Smart Card:" privacy and security concerns. In: Proceedings of the 36th Hawaii international conference on system sciences (HICSS'03). Hilton Waikoloa Village, Island of Hawaii, January 6—9; 2003.

Information Security Team, DePaul University. A framework for incident response (draft), <http://security.depaul.edu>; 13th December 2002.

International Standards Organization. Code of practice for information security management 2000. ISO/IEC 17799.

Internet Engineering Task Force, Request for Comments (RFC) 1305. Network time protocol (version 3) — specification, implementation and analysis. Available from: http://www.ietf.org; March 1992.

Internet Engineering Task Force, Request for Comments (RFC) 2350. Expectations for computer security incident response. Available from: http://www.ietf.org; June 1998.

Jung HT, Kim HL, Seo YM, Choe G, Min SL, Kim CS, et al. Caller identification system in the Internet environment. In: Proceedings of fourth USENIX security symposium; 1993.

Kent S, Atkinson R. Request for comments (RFC) 2401 — security architecture for the Internet protocol. Available from: http://www.ietf.org; November 1998.

Killcrece G, Kossakowski KP, Ruefle R, Zajicek M. Organizational models for computer incident response teams (CSIRTs). Report: CMU/SEI-2003-HB-001. Carnegie Melon University/Software Engineering Institute; December 2003.

Kossakowski KP, Allen J, Alberts C, Cohen C, Ford G, Fraser B, et al. Responding to intrusions. Report: CMU/SEI-SIM-006. Carnegie Melon University/Software Engineering Institute; February 1999.

Kruse W, Heiser J. Computer forensics. Canada: Addison-Wesley; 2002.

Lemos R. Study: sites attacked 4,000 times a week. CNET News. Available from: http://news.com.com/2100-1001-258093.html?legacy=cnet; 22 May 2001.

Mandia K, Procise C. Incident response: investigating computer crime. NY: Osborne/McGraw-Hill; 2002.

Mankin A, Massey D, Wu CL, Zhang L. On design and evaluation of intention-driven ICMP traceback. In: IEEE international conference on computer communications and networks (ICCCN); October 2001.

McClure S, Scambray J, Kurtz G. Hacking exposed. McGraw-Hill; 2001.

McGraw, Felten. Securing Java: getting down to business with mobile code. NY: Wiley; 1999.

Microsoft Corporation. Introduction to code signing, <http://msdn.microsoft.com/library/default.asp?url=/workshop/security/authcode/intro_authenticode.asp>; 1999.

National Institute of Standards and Technology. Computer security incident handling guide January 2004. NIST Special Publication 800-61.

Nong Y, Giordano J, Feldman J, Zhong Q. Information fusion techniques for network intrusion detection. In: IEEE information technology conference, information environment for the future, Syracuse, NY, USA; September 1998

OMB's Circular No. A-130. Appendix III online. Available from: http://www.whitehouse.gov/omb/circulars/a130/a130trans4.html; January 2004 [05/01/2005].

Park K, Lee H. On the effectiveness of probabilistic packet marking for IP traceback. In: Proceedings of 2001 conference on applications, technologies, architectures and protocols for computer communication, ACM SIGCOMM'01. San Francisco, US; August 2001.

Patsos D. A strategic approach to incident response, M.Sc. thesis. London: Department of Mathematics/Information Security Group, Royal Holloway University of London; 2002.

Postel J. Request for comments (RFC) 791 — Internet protocol. Available from: http://www.ietf.org; September, 1981.

Rekhter Y, Watson TJ. A border gateway protocol 4 (BGP-4). Available from: http://www.ietf.org; September, 1981.

Savage S, Wetherall D, Karlin A, Anderson T. Practical network support for IP traceback. In: Proceedings of SIGCOMM'00. Stockholm, Sweden; August 2000.

Schnackenberg D, Djahandari K, Reid T, Wilson B. Cooperative intrusion traceback and response architecture (CITRA), Boeing Phantom Works and NAI Labs, prepared under contract N66001-01-C-8048 for Space and Naval Warfare System Center (SSC), San Diego; February 2002.

Schultz E. Incident response teams need to change. Computers and Security Journal January 2004;23:87—8.

Solove DJ. The legal construction of identity theft. In: Symposium: digital cops in a virtual environment Yale law school; March 26—28, 2004.

Song DX, Perrig A. Advanced and authenticated marking schemes for IP traceback. In: Proceeding of the IEEE INFOCOM01. Anchorage, Alaska; April 2001.

Spafford EH, Weeber SA. Software forensics: can we track code to its authors? Purdue Technical Report CSD—TR 92—010; February 1992.

Staniford-Chen S, Heberlein LT. Holding intruders accountable on the Internet. In: Proceedings of IEEE symposium on security and privacy; 1995.

Stoll C. The cuckoo's egg, pocket; reprint edition; November 1, 1990.

Stone R. CenterTrack: an IP overlay network for tracking DoS floods. In: Proceedings of 9th Usenix security symposium; August 2000.

United States Code, Chapter 35 of Title 44, Subchapter III — Information Security, Federal Information Security Management Act (FISMA) of 2002.

US Department of Commerce. Federal Information Processing Standards Publication 198, The Keyed-Hash Message Authentication Code (HMAC); March 6, 2002.

Van Wyk K, Forno R. Incident response. NY: O'Reilly; 2001.

Wang XY, Reeves DS, Wu SF, Yuill J. Sleepy watermark tracing: an active intrusion response framework. In: Proceedings of the 16th international information security conference (IFIP/Sec'01); June 2001.

West-Brown MJ, Stikvoort D, Kossakowski KP. Handbook for computer security incident response teams (CSIRTs). Report: CMU/SEI-98-HB-001. Carnegie Melon University/Software Engineering Institute; December 1998.

Whalen S. An introduction to ARP spoofing, (white paper). Available from: http://www.gmx.net; December 1998.

Yasincac, Manzano Y. Policies to enhance computer and network forensics. In: Proceedings of the 2001 IEEE workshop on information assurance and security. West Point, NY: United States Military Academy; 5—6 June 2001.

Zhang Y, Paxson V. Detecting stepping stones. In: Proceedings of the 9th USENIX security symposium. Denver, Colorado, August 14—17; 2000.

Sarandis Mitropoulos is a visiting lecturer at the Department of Informatics of the University of Piraeus, Greece, and a System Analyst at a Bank supervised by the Ministry of Economics of Greece. He completed his Ph.D. at the Department of Electrical and Computing Engineering of the National Technical University of Athens (NTUA) in 1994. His Ph.D. dissertation focused on Distributed System and Network Management. He received his degree in Informatics and Computing Engineering from the same department of NTUA in 1990. He has been working in technical and project management and business development in European R&TD and integrated solution projects, in the areas of system and network management, of advanced telecommunication/telematic services, and of management information systems, as well as of system and network security. He taught in NTUA from 1991 to 1994. He is a member of IEEE, CNOM, and Technical Chamber of Greece.

Dimitrios Patsos is a Ph.D. candidate at the Department of Informatics of the University of Piraeus, Greece. He received his B.Sc. in Informatics from the Department of Informatics of the Athens University of Economics and Business (AUEB) in 2001 and his M.Sc. in Information Security from the Department of Mathematics of the Royal Holloway University of London, UK, in 2002. He is co-author of the book ''Security of Information: in computers, the Internet and the ordinary life'', one of the first information security-related books to be published in Greek. He has been serving as an IT Security Consultant for a major Greek Systems Integrator since 1999, dealing with various aspects of Information Security with a strong emphasis on the Financial Institutions Sector and especially to Banks. His main research interests are security management, cryptography, network security, Incident Response and electronic crime.

Christos Douligeris received the Diploma in the Electrical Engineering from the National Technical University of Athens in 1984 and his M.S., M.Phil., and Ph.D. degrees from the Columbia University in 1985, 1987, 1990, respectively. He has held positions with the Department of Electrical and Computer Engineering at the University of Miami, where he reached the rank of associate professor and was the associate director for engineering of the Ocean Pollution Research Center. He is currently teaching at the Department of Informatics of the University of Piraeus, Greece. He has served in technical program committees of several conferences. His main technical interests lie in the areas of performance evaluation and security of high-speed networks, neurocomputing in networking, resource allocation in wireless networks and information management, risk assessment and evaluation for emergency response operations. He was the guest editor of a special issue of the IEEE Communications Magazine on ''Security for Telecommunications Networks'' and he is preparing a book on ''Network Security'' to be published by IEEE Press/Wiley. He is an editor of the IEEE Communications Letters, a technical editor of IEEE Network and a technical editor of Computer Networks (Elsevier).