# Integrated Security Incident Management — Concepts and Real-World Experiences

Stefan Metzger, Wolfgang Hommel, Helmut Reiser
*Leibniz Supercomputing Centre*
*Munich, Germany*
*Email: metzger@lrz,de, hommel@lrz.de, reiser@lrz.de*

## Abstract

We present a holistic, process-oriented approach to ISO/IEC 27001 compliant security incident management that integrates multiple state-of-the-art security tools and has been applied to a real-world scenario very successfully for one year so far. The computer security incident response team, CSIRT, is enabled to correlate IT security related events across multiple communication channels and thus to classify any incidents consistently. Depending on an incident's classification, manual intervention or even fully automated reaction steps can be triggered; this starts with simple email notifications of system and network administrators, and scales up to quarantining compromised systems and subnetworks automatically. A formally specified security incident response (SIR) process serves as the basis that clearly defines responsibilities, workflows, and interfaces. It has been designed to enable quick reactions to IT security events in a very resource-conserving manner.

## Keywords

ISO/IEC 27001; computer security incident response team; IT service management; network abuse; intrusion detection

## I. Motivation

IT systems and services that are connected to the Internet are profitable targets for attackers for obvious reasons, such as money that can be made from selling personally identifiable information or leasing botnets to similarly malicious third parties. However, this does not only apply to e-commerce web sites and private end user PCs. Also public institutions, such as government agencies and higher education institutions (HEIs), are in the focus of attackers, as shown by the publicity of events as those described in [3]. In the case of HEIs, e. g., the personal data of employees (staff, faculty) and students including email addresses and personnel numbers are attractive targets for attackers interested in stealing data; similarly, the high Internet uplink bandwidth of many HEIs make remotely controlled PCs an interesting addition to botnets that are used, for example, for sending spam emails or participating in denial of service attacks. Besides the financial or otherwise direct damage caused by successful attacks, such as the results of data loss and the time needed to repair compromised systems, there often also is a huge damage to the organization's reputation, especially to the responsible computing center's one.

Network access from outside, i. e., originating from the Internet, is typically blocked using well-established tools like packet filter firewalls and intrusion detection/prevention systems (IDS/IPS). However, an attack's source may also be within the HEI's local area network. Attacker models include disgruntled employees, adventurous students, and WiFi users including guests and conference attendees, which are granted access only for a limited time. While the organization's servers and the employee's desktop PCs are usually centrally managed, and thus are being updated and patched in a reliable fashion, there typically is no such solution for the centrally triggered, contemporary update of students' and guests' hardware equipment (PCs, notebooks, netbooks, smartphones, tablet PCs, etc.).

For this very reason, solely preventive measures are not sufficient in such scenarios, as they can only be applied to a few selected parts of the overall IT infrastructure. As a consequence, additional measures for detecting and reacting to security incidents in a well-structured manner are required. In this paper, we suggest a process-oriented approach that is composed of the building blocks shown in Figure 1:

The overall coordination is one of the tasks within a precisely specified security incident response process that is based on the paradigm of process-oriented service management (cf. [14], [15], [16], [17]). Within this process, responsibilities, duties, and tasks are clearly assigned in order to control all the workflows that are triggered for the handling of security events and incidents. The process presented in this paper puts special emphasis on the integration of multiple ways to report security events that are complementary to each other:

- Incidents can be reported manually, e. g., by email, phone, facsimile, other institutions' computer security incident response teams (CSIRTs), or law enforcement agencies. In our experience, this manual approach is an essential and the most frequently used reporting method especially in Grid projects, i. e., distributed scientific computing and storage infrastructures provided multiple HEIs in a federated manner.
- Third party services like the security reporting service by the CERT of the DFN (”Deutsches Forschungsnetz”), which is Germany's national research and education network (NREN) provider. Especially DFN-CERT service *automated security warnings* is here of interest.

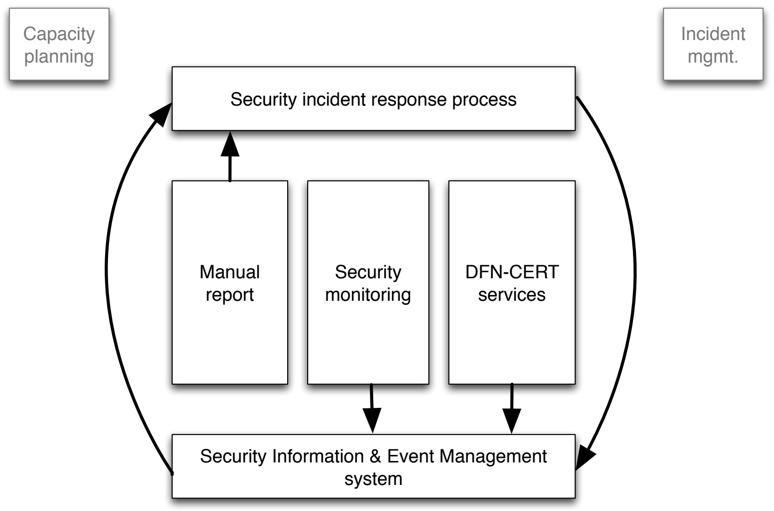
Capacity planning
Incident mgmt.
Security incident response process
Manual report
Security monitoring
DFN-CERT services
Security Information & Event Management system

Figure 1. Integrated management of security incidents in general

- Various local security monitoring mechanisms that are used to detect malware and unusual, typically disallowed communication behavior of clients that are connected to the LAN.
- The central analysis of security relevant events by means of a Security Information and Event Management (SIEM) system that offers correlation and automated reaction capabilities.

It is our explicit goal to combine all these reporting channels in a way that very quick reactions and responses can be achieved within a holistic, integrated security incident management process. We focus on the fully automated, next-to-realtime reaction for so called standad security incidents or minor security incidents in order to disburden the security staff and to gave them the time to cope with more complicated cases or major security incidents.

In this paper, we present our solution that has already been applied to a real-world scenario very successfully. The remainder is structured as follows: In the next section, we discuss the security information sources we use, including our NREN's CERT services. In Section III, we outline our scenario, the Munich Scientific Network (MWN), which spans more than 120.000 users and 80.000 devices connected per day in average, along with several challenges we met there. Furthermore, we give an overview of the security monitoring tools that are currently being used at the Leibniz Supercomputing Centre (LRZ), which operates MWN. In Section IV we summarize the ISO/IEC 27001 compliant security incident response process established at LRZ and go into detail about several of its key aspects. We then describe how the used mechanisms and communication channels can be combined with respect to the integrated management of security incidents and the integration into LRZ's overall IT service management strategy that is aligned to ISO/IEC 20000. The next section gives some advices and describes a few stumbling blocks for other HEIs if they decide to implement an integrated security incident response procedure. We conclude with a short summary of our real-world experiences and an outlook to the further research and enhancements we planned in this area.

## II. DFN-CERT Services: Automated security warnings

The wide range of services provided by the DFN-CERT consists of proactive (information on current vulnerabilities) and reactive services (automated security warnings, denoted as AW-Service) [4].

Proactive information on vulnerabilities of used operating systems and installed applications helps closing them before they are exploited successfully by an attacker. Additionally, a new preventive service is called *network checker*; based on the NMAP scan utility, HEIs can check whole networks and individual systems from an external point of view. Use of the network checker service provides a very simple method to determine, which systems and services are reachable and if essential security controls are in place (e.g., SSH access to web servers directly from the Internet should be prohibited). This service complements locally conducted scan activities because different firewall rulesets for internal and external source addresses typically lead to differences in the scan results. Thus, both views, internal and external, constitute an indispensable means for regular vulnerability assessments and risk management.

Closing potential vulnerabilities of centrally administered systems is a relatively simple task once a patch and deployment management process and supporting tools have been set up: New security patches provided on central update servers can be installed regularly and automatically on most systems. The implementation of these proactive controls reduces the probability of an occurrence of specific security incidents and therefore prevents harm of individual systems and services.

Despite using these preventive means, infections with a virus, malware, or in general the compromising of a system are still possible. Once such an incident happened, the DFN-CERT AW-Service comes into play. This reactive service is based on continually analyzing the detection results of different security sensors and other information sources all over the world, and in particular running their own security sensors, e.g., honeypots, at DFN-CERT. Honeypots act as a system providing vulnerable services. The intention and ambition of using honeypots is the analysis of malware installed by successful attackers and its behavior. Thus, HEIs using this DFN-CERT service receive a report about a conspicuous IP address, and additionally detailed information about the installed malware, e.g., its name, behavior, and how it can be removed. This information helps selecting the right, effective, and goal-oriented countermeasures to remove the malware, e.g., by installing dedicated removal tools. Each conspicuity is classifed, so that similar cases or behavior patterns belonging to the same category and thus the same reporting type can be identified and grouped. Currently, DFN-CERT distinguishes between about ten types, e.g., attack, botnet communication, portscanning, or spam. For certain types, e.g., botnet communication, a more detailed categorization takes place. For example, the name of the malware, e.g., Conficker or Mebroot, is additionally reported. Here is a short partial example of an email sent by the AW-Service:

```
System:       xxx.xxx.xxx.xxx
Type:         Bot
Timestamp:    2010-07-17 09:31:12 GMT+0200 (MEST)
Description:  Some botnet software seems to be operated on this
              system, which tries to contact a HTTP-based bot net
              control server. You can find further information about
              various malware types here:
              http://www.cert.dfn.de/index.php?id=bot

TCP source port   TCP dest port   malware type    HTTP Request
-----------------------------------------------------------------------
1667              80              Conficker       GET /search?q=0 HTTP/1.0
3812              80              Conficker       GET /search?q=0 HTTP/1.0
```

In addition, DFN-CERT reports serious security incidents and security-related events to certain contact persons of the affected HEI directly. Complementary, incidents that have been internally detected could be reported to DFN-CERT. This gives DFN-CERT the opportunity of a national overview of the current situation and to determine whether there is an attack against many different institutions running. If one of DFN's customers detects a security incident on an internal machine, DFN-CERT supports and helps resolving the incident quickly. Besides the competency and experience of each DFN-CERT member, it has a very beneficial effect on the solution of a current incident that DFN-CERT has a good overview on similar incidents at other DFN-related HEIs.

## III. Proactive network and system monitoring within the MWN at LRZ

The Munich Scientific Network (MWN) primarily consists of the locations of the Munich universities (Technische Universität München (TUM), Ludwig-Maximilians-University (LMU), Hochschule München (HM), and FH Weihenstephan-Triesdorf). Also several non-university institutions are associated. Currently more than 80.000 devices are connected each day. Because of only temporarily connected devices that are owned by students, guests, or, e.g., conference members, this figure varies and the total number is probably even higher. It can also be observed that this value increases continuously [5]. Because of its structure, the MWN could not be maintained like a typical enterprise network. The administrative responsibility of LRZ ends at the physical network outlet. Which system is connected to this network jack is in each department's and faculty's area of responsibility. They are also responsible for the maintenance and security of the connected end-systems. Despite of this distributed structure and the enormous number of connected devices, which are administered in a decentralized fashion, LRZ attempts to protect internal systems and systems operated for other institutions from attacks. LRZ also works on preventing attacks and mitigating already carried-out attack attempts that are coming from LRZ-internal or any MWN-connected devices. The primary objective here is that at least the impact and the potentially caused harm should be minimized.

In this section our currently used mechanisms are detailed:

- Network-based intrusion detection system (IDS)
- NAT gateway NAT-o-MAT
- Mail monitoring based on accounting data
- Analysis of netflow data

The detection of security-related events itself is not sufficient; rather an automatic processing, the notification of responsible adminstrators, and a near-real-time response are necessary. This is achieved by using a central security information and event management system.

### A. SNORT - the well-known intrusion detection system

We use the very well-known and worldwide popular open-source network intrusion detection system (IDS) SNORT [11] for the security monitoring of the inbound and outbound Internet traffic; it also serves for the detection of known attacks and compromised or worm-infected systems. The network traffic is mirrored using a SPAN-port, forwarded to the IDS machines, and analyzed. Analysis by SNORT is accomplished by comparing packets and their content with defined patterns.

Because of the worldwide usage of SNORT, the software itself and its rulesets are being continuously improved. Thus, an appropriate response to new attacks or attack methods is kind-of guaranteed. It is an essential criterion at LRZ that a software is widely used, maintained and improved.

Commonly available SNORT community rulesets are combined with self-written rules so that the detection of events that are considered especially important for the LRZ is achieved and an alarm is generated. For instance, an internal SSH attacker, which attempts to connect very often to (different) destination systems outside the MWN on destination port 22 (SSH scanning), is also detected as an internal system on which a FTP server is installed and provides its service somewhat obfuscated on a port different than the standard FTP port 21.

We observe a phletora of external SSH attacks targeted at systems inside the MWN on a daily basis. Combined with a dictionary attack to break weak passwords for default accounts, e.g., *root*, *admin*, or *guest*, those attacks are, if successful, a potential cause of a security incident. Internal SSH attackers, which scan systems outside the MWN are probably infected with a virus or some other kind of malware. The detection of such compromised systems by a SNORT-based IDS is a very efficient solution. However, the signature-based detection of SNORT is made difficult because of our 10 Gbit/s internet link. Therefore, up to eight instances of SNORT have to be started on the used multi-core machines in parallel, because the software itself doesn't support multi-threading yet. Additionally a few changes to the parameters of the network cards had to be made.

To reduce the network traffic, each IDS instance has to analyze (each instance should have to analyze the same volume) the traffic that is split based on Berkeley Packet Filters (BPF) [12]. These filters are parametrized empirically and checked monthly. The traffic is separated on the protocol level (e.g., TCP, UDP, ICMP, or IPv6) and by IP address ranges combined with individual ports (e.g., HTTP or SSH). Here is an example of a BPF-based filter for HTTP traffic:

```
ip and ( <subnets>) and tcp and ( port 80 or port 8080 or port 8180 )
```

Because of the filter-based traffic splitting, the usage of the common available SNORT community rulesets is not possible out-of-the-box anymore, since these rules are attack-based, i.e., for the detection of specific attacks, e.g., virus or trojan horses, and not protocol- or port-based. Thus, LRZ has to parse the community ruleset, disassemble the ruleset, and compose new rules so that each started SNORT-instance uses a minimal ruleset. This positively affects the performance of each instance and thus the time to analyze. However, this approach is not very elegant and requires a very high administrative, usually manual effort. Thus, LRZ plans to optimize this approach splitting the traffic at the router in the future.

### B. NAT-o-MAT - NAT gateway and security monitoring

Using private IP addresses is an effective protection mechanism for internal systems. A direct communication between external (Internet) and internal systems, and therefore a channel to attack them, is not possible. However, using externally provided services is not possible either.

Dedicated proxy servers, which are the often installed solution to this latter issue, had to be configured by each user, which led to a high number of support requests. Thus, this approach was inconvenient and error-prone. LRZ developed a NAT-gateway, which on the one hand translates private to public IP addresses, and on the other hand works as a security gateway, which automatically controls bandwidth usage and detects and prevents various kinds of attack methods (e.g., portscanning, (D)DoS activities and spam sending) [6], [13].

The core for NAT-o-MATs security-monitoring capabilities is a rate-based mechanism. For each event of exceeding a predefinied threshold value, e.g., 100 packets/sec, the source system is assigned one penalty point. Depending on the number of penalty points inside a sliding time-window, the internet connection will be temporarily suspended until this number decreases (below another treshold). The user will be informed by email and redirected to a landing page:

```
Blocked since 20.12.10 05:20
Number of hits     Protocol  Destination port and suspension reason
111                UDP       4600-4699 - Edonkey Filesharing
30                 UDP       7560
30                 UDP       5514
30                 UDP  54067
30                 UDP       4600-4699 - Edonkey Filesharing
30                 UDP       43793
30                 UDP       37283
```

At the moment, a new extended version of NAT-o-MAT called Secomat is being developed. The policy-based routing (PBR) load balancing is replaced. Additionally, the non-evolving classification of P2P protocols is going to be substituted by a SNORT-based method. In that case the NAT security gateway also provides intrusion detection capabilities, which will positively affect the flexible detection and lower the system and traffic load on the currently used IDS.

### *C. Mail monitoring based on accounting data*

Through analysis of accounting data, the identification of spam sending systems within the MWN is an effective task. The most common reason for such a behavior is that the offending system is compromised and infected with a specific kind of malware, usually some botnet (zombie) software. A rate-based monitoring mechanism counts each mail connection within a specific time window. If a pre-defined threshold value is exceeded, an alert is raised. Besides two different time intervals, five minutes and one hour, two threshold values are considered. If a soft limit event occurs, i.e., 20 mails within 5 minutes or 80 mails within one hour, an email is sent to the responsible system administrator. A hard event, i.e., 300 mails within 5 minutes or 1000 mails within one hour, triggers an automatic suspension of the Internet connection. For each time interval the administrator can optionally define an individual threshold value.

```
Monitoring details
==================

System:                [xxx.xxx.xxx.xxx]

Monitoring interval:     5 minutes  before  12.12.2010  12:30:28
Number of mail connections:      1616
```

A high (too high for a pre-defined threshold value), but only temporary occurrence of mail connections, e.g., for sending a monthly newsletter, could be excluded from a false positive alerting or suspension of a system, by using a whitelist. Dedicated mail servers, which are maintained by the individual institutions themselves and for which a higher number of mail connections should be considered as normal state, are also registered in this exception list. The maintenance of the exception list's entries is supported by each network or system administrator.

Also, the detection of spam sending systems through the intrusion detection system SNORT was taken into consideration and implemented prototypically. But this approach turned out to be unfeasible. The number of false-positives was much higher than using the rate-based mechanism. Additionally the maintenance effort (exception list, individual threshold values) had to be made by LRZ personnel and could not be delegated to the local network administrators.

### *D. NfSen - Analysis of netflow data*

Recently, the analysis of netflow data for security monitoring purposes and for detection of attacks has become a very important task at LRZ. The IDS detects only (signature based) know attacks. Rather than unknown, maybe zero-day attacks could be detected within netflow data using appropriate analyzing methods to determine traffic anomalies. At the moment the detection of botnet communication and the identification of spam sending systems using IPv6 or teredo technologies have priority [9].

NfSen provides a web-based interface for the NetFlow toolset *nfdump* and allows a convenient analysis of netflow data. Possible applications of NfSen and NFDUMP were presented by Haag at the 16th TF-CSIRT Meeting [8] in detail. Within various so-called profiles, sorted by flows, packets, or transmitted data volume and partitioned by used protocols the current situation can be displayed very clearly. A closer look into the netflow data can be taken by manual analysis using a special

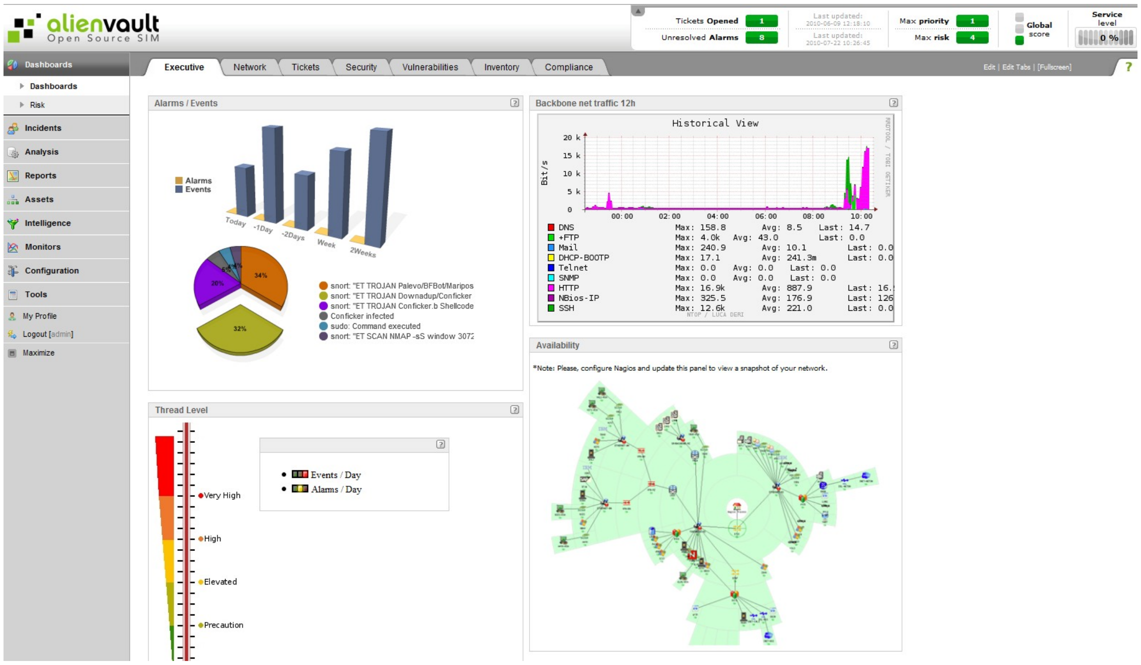
alienvault
Open Source SIM
Tickets Opened 1
Unresolved Alarms 8
Last updated: 2010-06-09 12:18:10
Last updated: 2010-07-22 10:26:45
Max priority 1
Max risk 4
Global score
Service level
0 %
Dashboards
Dashboards
Risk
Incidents
Analysis
Reports
Assets
Intelligence
Monitors
Configuration
Tools
My Profile
Logout [admin]
Maximize
Executive
Network
Tickets
Security
Vulnerabilities
Inventory
Compliance
Edit | Edit Tabs | [Fullscreen]
Alarms / Events
Alarms
Events
Today
-1Day
-2Days
Week
2Weeks
34%
32%
20%
5%
4%
4%
snort: "ET TROJAN Palevo/BFBot/Maripos
snort: "ET TROJAN Downadup/Conficker
snort: "ET TROJAN Conficker.b Shellcode
Conficker infected
sudo: Command executed
snort: "ET SCAN NMAP -sS window 3072
Backbone net traffic 12h
Historical View
Bit/s
20 k
15 k
10 k
5 k
0
00:00
02:00
04:00
06:00
08:00
10:00
RRDTOOL / TOBI OETIKER
DNS Max: 158.8 Avg: 8.5 Last: 14.7
+FTP Max: 4.0k Avg: 43.0 Last: 0.0
Mail Max: 240.9 Avg: 10.1 Last: 0.0
DHCP-BOOTP Max: 17.1 Avg: 241.3m Last: 0.0
Telnet Max: 0.0 Avg: 0.0 Last: 0.0
SNMP Max: 0.0 Avg: 0.0 Last: 0.0
HTTP Max: 16.9k Avg: 887.9 Last: 16.
NBios-IP Max: 325.5 Avg: 176.9 Last: 126
SSH Max: 12.6k Avg: 221.0 Last: 0.0
NTOP / LUCA DERI
Availability
*Note: Please, configure Nagios and update this panel to view a snapshot of your network.
Thread Level
Very High
High
Elevated
Precaution
Events / Day
Alarms / Day

Figure 2. Integrated, configurable dashboards of OSSIM

filter syntax. Keywords, like *ip*, *net*, or *dst port* can be combined in an intuitive manner. These filters are useful for manual analysis as well as for using the integrated, automated alerting mechanism.

### *E. OSSIM - Security Information & Event Management*

The Open Source Security Information Management (OSSIM) System by Alienvault [7] has become very popular [10]. OSSIM allows for the central collection, correlation, and analysis of security relevant event data. The analysis can be performed manually through efficient search functions. Events can be filtered and sorted by sensors, event-id, source IP, etc., and exported in various file formats. Within the integrated asset management, individual hosts, subnets, or networks can be consolidated to host or network groups, which improves clarity. Configurable dashboards, as shown in Figure 2, and scheduled reporting, e.g., to see the top 10 attackers of a particular day, complete the impressive feature set.

By means of event correlation, prioritizing security relevant events is a straight-forward task. Very interesting is the cross-correlation feature with which the events of different sensors can be correlated. Thus, it becomes possible to cross-correlate currently detected IDS events with vulnerabilities of the attacked system. IDS events could have a very high severity but could be irrelevant because the attacked service, the required operating system is not running, or an updated, non-vulnerable software version is already installed. OSSIM provides the facility to a flexible response in real-time, starting with simple email notification of the system administrator up to running dedicated scripts that perform more complex reactions.

OSSIM is the central instance for analyzing security events detected by internal sensors at LRZ. Some special correlation directives were implemented to respond suitably, based on defined polices, whereby the whole range of response facilities depending on the event severity and priority is used.

## IV. The Security Incident Response Process

Besides the internal monitoring systems and the information provided by DFN-CERT there is a third, highly important data source for reports about security events and incidents that we use: The manually created reports by local system and service administrators, which, for example, find out about security issues when auditing log files, or if they detect anomalies while working on their machines. In this section we outline the process for the handling of security incidents that has been established at LRZ during 2010. It was designed with compliance to the information security management system standard ISO/IEC 27001 [16] in mind.

Concerning the integration into LRZ's overall IT management tool landscape, it has to be mentioned that we are working on increasing the reliability and availability of all IT systems and services even further by applying the process-oriented IT service management (ITSM) paradigm, as it is defined in standards like ISO/IEC 20000 [14], [15] and best practices, such as ITILv3 [1], on a broad basis. One of our major strategic goals is to achieve an ISO/IEC 20000 certification, which in turn requires a solid security management process, of which the security incident response is an important part. A key prerequisite therefor is that processes and workflows are not only documented and instantiated in a reproducible manner, but that they are also subject to continuous improvement. The ITSM process Incident Management handles, among various other tasks, the receipt of reports about service incidents by customers and users as well as their quickest possible solution. Besides service requests by users that are being taken care of by the organization's first level support (e.g., help desk) immediately, two especially important types of incidents must be considered: Major Incidents, such as the total failure of an important service that impacts many customers and users, and Security Incidents, for whose handling security experts have to be consulted in addition to the local infrastructure and service operators. In theory, Security Incidents can quickly turn into Major Incidents, e. g., if an important system has been compromised and needs to be taken offline; in practice, however, Security Incidents happen in a much higher frequency than Major Incidents.

To enable the reporting of security related incidents by LRZ employees, we first established a hunt group telephone number and a mailing list as points of contact. Thus, the incident reporter can directly get into contact with a member of LRZ's CSIRT, i. e., with an expert for the handling of Security Incidents. Although the traditional first level support is bypassed this way, we made sure that the centrally operated trouble ticket system is being used to document Security Incidents, so that all incidents are being documented in a uniform manner regardless of their type and report channel.

For each new report that is made to the CSIRT, the security incident process is being instantiated; as shown in Figure 3, it primarily consists of the phases classification, escalation, analysis, diagnosis, solution, and closing. These phases, which are being discussed in more details below, are documented in the following ways:

- LRZ administrators and service managers have a deliberately short summary of the process ready to hand, which mentions contact details, the most important data that shall be reported about the security incident, and a few best practice recommendations for first aid and how to proceed in the case of emergency.
- LRZ CSIRT has a detailed description of the process at its disposal, which is rather voluminous given its more than 20 pages. It specifies the responsibilities, tasks, and workflows in detail and includes links to other documents and workflow descriptions, e. g., about basic IT forensic know-how that can be applied during the analysis of compromised systems. Forensic analysis is a very wide area and requires a team, which is up-to-date, knows current attacks and the behavior of specific malware and can also use special forensic tools. LRZ-CSIRT has currently insufficient personnel, which is aware of using forensic tools and thus we focus on basic, already installed, tools and commands to get an overview about the compromising.
- To avoid getting lost in the heat of the action, the LRZ CSIRT also uses a five pages checklist that summarizes the most important process steps and workflows in a comprehensive way.

During the initial analysis all the reported data, which, for example, shall include the contact details of the incident reporter and a description of the identified security issues, the sufficiency of the data is being evaluated and an initial classification of the security incident is being performed. This classification is based on aspects such as the number of affected systems and services, whether LRZ-internal or customer-operated machines are affected, which services and SLAs may be impacted, whether there are additional dependencies on other services, where the assumed attacker is located, and what type of attack has been observed. Additionally, a cross-correlation with other open and resolved security incidents is performed. On this basis, one of the four priorities low, medium, high, or very high is assigned. If the priority is very high, the incident is turned into a Major Incident.

So-called Standard Security Incidents (SSIs) are being specified in the course of time for specific types of security-related incidents as well as for those that occur very frequently; for each of these SSIs, certain simplifications of the process may be approved. For example, a brute force attack to guess a user's password via SSH login attempts is considered to be an SSI if the attack source is outside the Munich Scientific Network, i. e., if we assume that the attack originates somewhere from the Internet, and if the attack also does not follow the username naming conventions we use for our users' accounts. Based on such empirical criteria, the high number of attacks we observe in everyday's operations and whose success rate (and thus our risk) is considered very low, can be handled very efficiently. On the other hand, declaring such attacks as SSIs ensures that they are being documented and can be considered during the planning of additional or improved security measures. On the other hand, if any anomalies are detected during the handling of security incidents, their classification can also be changed in order to reflect any newly available information.

In the case that a security incident report has to be analyzed in detail, an additional rating scheme is applied to determine the impact and the urgency of the incident. The impact is derived using a weighting formula that also takes into account the priorization criteria mentioned above. However, opposed to the incidents that are handled by the first level support, security related incidents are always considered to be urgent. Based on these results, not only the internal queueing of incidents is being rearranged appropriately, but they also serve as a basis for the decision about which customer service level agreements

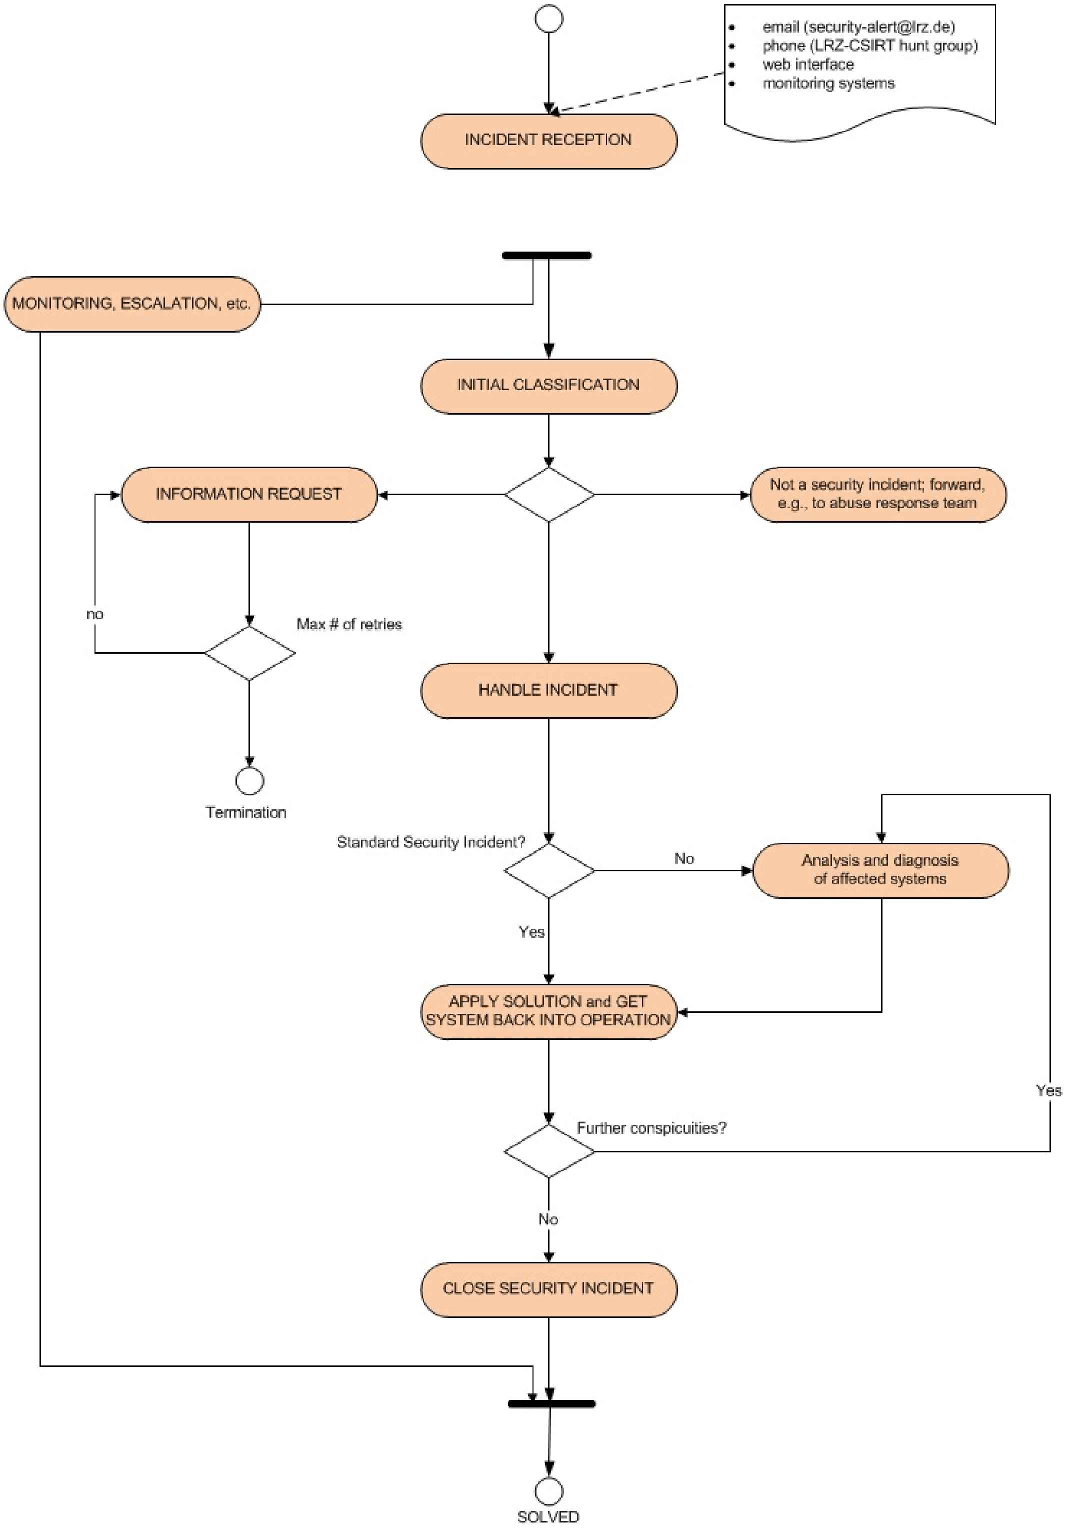
email (security-alert@lrz.de)
phone (LRZ-CSIRT hunt group)
web interface
monitoring systems
INCIDENT RECEPTION
MONITORING, ESCALATION, etc.
INITIAL CLASSIFICATION
INFORMATION REQUEST
Not a security incident; forward, e.g., to abuse response team
no
Max # of retries
HANDLE INCIDENT
Termination
Standard Security Incident?
No
Analysis and diagnosis of affected systems
Yes
APPLY SOLUTION and GET SYSTEM BACK INTO OPERATION
Yes
Further conspicuities?
No
CLOSE SECURITY INCIDENT
SOLVED

Figure 3. The workflow of our Security Incident Response Process

(SLAs) are affected, which maximum reaction time has been negotiated with the affected customers, and whether the incident is so severe that LRZ's management board has to be notified as soon as possible about the issue.

The further activities are then being coordinated by a Security Incident Coordinator (SIC) who is selected on a case-by-case basis. The SIC is granted extended decision-making power during the handling of a security incident and may call additional personnel in to consult about the necessary steps to find a solution to the incident. First countermeasures and duties may, for example, include the removal of the affected systems from the network, creating backups and disk images to preserve any evidence, performing basic IT forensics including the evaluation of any security, network, and system monitoring data that is available, escalating the incidents, and documenting all the steps that have been performed. The two primary objectives during this phase are to keep in touch with the incident reporter in order to ensure the sufficient flow of relevant information, and to find a way to restore the affected system or service as quickly as possible to avoid SLA violations. After a solution for a particular incident has been implemented, the system will be monitored intensively for up to 14 days in order to determine whether the incident actually happens again despite the implemented measures or if any other impact becomes evident.

Reviews are performed after each major incident as well as periodically. Their goal is, on the one hand, to discuss about possible improvements of the process as a whole and its workflows and activities. On the other hand, a trend analysis is performed to suggest consequences that should result from the previous security incidents; for example, the preventive as well as the proactive security measures might be enhanced to achieve a long-term improvement of the overall security level.

## V. Merging the various reports: integrated management of security incidents

As described in the previous sections there are three ways to report security relevant events and security incidents:

- Automatically generated warnings sent by DFN-CERT
- Various internal security monitoring mechanisms detect and forward events to a centralized SIEM (OSSIM) for correlation
- Manual reporting by sending an email or by phone

These three ways have to be combined to obtain the quickest possible response. To achieve this goal, all reports are merged in a centralized, ISO/IEC 20000 compliant trouble ticket system (TTS). Less severe events are processed fully automatically; i.e., there are dedicated, automatic countermeasures in place. In such a case, the created trouble ticket serves for tracking purposes only.

In this section our approach for an integrated management of security incidents will be detailed. First we describe the automated handling of the DFN-CERT warning messages. Section V-B explains the centralized analysis of events detected by internal monitoring tools and the automated response capabilities. In this section, also some issues will be mentioned, which will have to be improved at LRZ in future. After that, in Section V-C, the description of the manual, but structured processing of incidents and the used tools follows. Finally, in Section V-D, we give a short example of the functional interaction between our different reporting ways and the monitoring mechanisms.

### A. *Analysis of the DFN-CERT warnings*

LRZ-CSIRT will be notified of conspicuous behavior of an IT system within the MWN by DFN-CERT. A file attached to that email in an XML file format provides the opportunity for script-based, automated handling. Therefore, forwarding such information to the responsible on-site administrator, especially for systems within the MWN, is a very straight-forward task. This recipient selection is based on the MWN network documentation, in which information including the network administrator's email addresses or rough system's location is saved ([5]).

Administrators of selected networks (e.g., those with own maintenance teams) will be notified directly. The LRZ-CSIRT only gets a copy of that email and could thereby monitor the progress of the incident handling. Unfortunately the AW-service does not provide real-time alerting, and thereby the opportunity for response on short notice, but works digest-based. Private systems, especially those of students, guests, and sometimes of conference attendees, which are only connected via wireless LAN temporarily, are indeed detected by the monitoring tools of DFN-CERT, but a targeted launch of countermeasures is hardly possible because of the delayed notification.

The analysis of the AW-service warnings at LRZ serves on the one hand for review of its internal monitoring mechanisms and provides on the other hand a simple opportunity to improve the security tools we use.

### B. *Detection through internal monitoring mechanisms and automated response through SIEM*

As described in Section III, various monitoring tools are in use to detect security relevant events. After the detection of such an event, it will be forwarded to the centralized security information and event management system OSSIM. There it will be analyzed and correlated with other reported events. Based on various configured policies, an automated response takes place, starting with a notification of the responsible administrator or user right up to the suspension of the offending machine's internet access. The CSIRT team is notified and a trouble ticket is generated. Using the centralized trouble ticket system provides the opportunity to monitor the incident progress and to react in a professional way if an administrator or an user has questions regarding an alerting notification.

An appropriate escalation mechanism was developed to respond to certain events. At the first-time conspicuity of a system only an information or warning is sent to the administrator or user. If the system or user becomes conspicuous a second time, a friendly reminder email is sent under penalty of the suspension. On the third conspicuity, the blocking of the IP address or user account takes place. This three-level approach has proven itself in practice. Most administrators and concerned users answer already the first notification. However, if an internal SSH attacker is detected, the access is blocked immediately. OSSIM provides the needed flexibility and opportunity to start the appropriate reaction.

Blocking of external firewall network interfaces or gateway addresses is a severe sanction. In such a case often a whole building or institute is affected and using Internet services is not possible anymore for all the users therein. At LRZ, an exception list is maintained which can be used for individual IP addresses or whole subnets. The latter is necessary for transporting networks, NAT pools or dynamically assigned DHCP addresses. The CSIRT will be notified even if the conspicuous IP address is defined as an exception. This is urgently required to force an IP address blocking.

Translation of private addresses to public IP addresses using NAT (NAT-o-MAT) has to be taken into consideration. Blocking of an IP address dynamically assigned by NAT does not make sense. Hence a mechanism has to be developed, which analyzes connection tracking data to determine the corresponding private IP address. At the moment, about 8.800 private IP addresses are in use and controlled by NAT-o-MAT. Because of the current limitation of the NAT pool (only two class C networks), the correlation happens within seconds based on protocol, IP addresse and port information.

In larger buildings, locating an offending system only based on the IP address is not feasible. Because of that an automated query to *Nyx* [2] is made in the response scripts. Nyx — a tool developed by LRZ — provides the possibility to determine the corresponding switch port to which the conspicuous system is currently connected, based on an IP or MAC address information. This switch port information is supplied automatically to provide the quickest-possible physical localization of the system.

### *C. Manual process-oriented intervention*

Currently the manual reporting of a security-relevant event is done either by email or phone. Usually system administrators report a conspicuity on a server machine after performing a partially automated, mostly tool-based analysis of logging data, which has to be examined in detail. For instance, login activities to unusual times or from unknown source IP addresses need to be investigated. Sometimes users call attention to unusual behavior of a certain service to administrators, which could have a direct link to a security incident. It should be mentioned that for the handling of a security incident, also manual analysis tools have to be in place, like simple sorting or filtering functions integrated in the centralized SIEM system. At the touch of a button, the start and duration of a conspicuity can be determined. Because events of different monitoring tools could be analyzed in parallel and temporally sorted, it is possible to comprehend the attacking process in detail. The integrated export functionality can also be used for documentation of an incident.

Similarly, efficient means exist for the creation of SNORT signatures. If suspicious activities are reported by email, then the tool IP extractor can be used to create SNORT variables for the involved IP addresses, which can afterwards become part of a signature. Also, extending the creation of NfSen filtering rules is possible using this tool. The web-based NfSen user interface allows for the manual analysis of netflow data, and an integrated alerting mechanism enables the real-time notification of administrators if communication from or to the watched IP addresses occurs. A self-developed reporting functionality provides an automated analysis of filter-reduced data and the creation of clear reports, which can be sent to administrators by email.

During manual intervention, Nyx is an important tool for the identification of compromised systems. The CSIRT can pass information on to the on-site system administrators, letting them know the switch port to which the system is connected to. Based on this information and using the local system documentation, the machine and the responsible user can be identified in a timely manner.

Another important manual activity of CSIRT members is the (re-)activation of blocked IP addresses or user accounts. The administrators or users are encouraged to describe the measures which they have taken to clean the compromised system. However, further development in this area is required to achieve an automated activation procedure as well. Administrators should be able to activate the blocked IP addresses themselves. This would further reduce the manual overhead for LRZ personnel and would help to avoid delays from the customers' point of view.

### *D. Implementation example*

To illustrate the combination of the various reporting ways and the usage of different monitoring tools, a short example is provided. The handling process is shown in Figure 4:

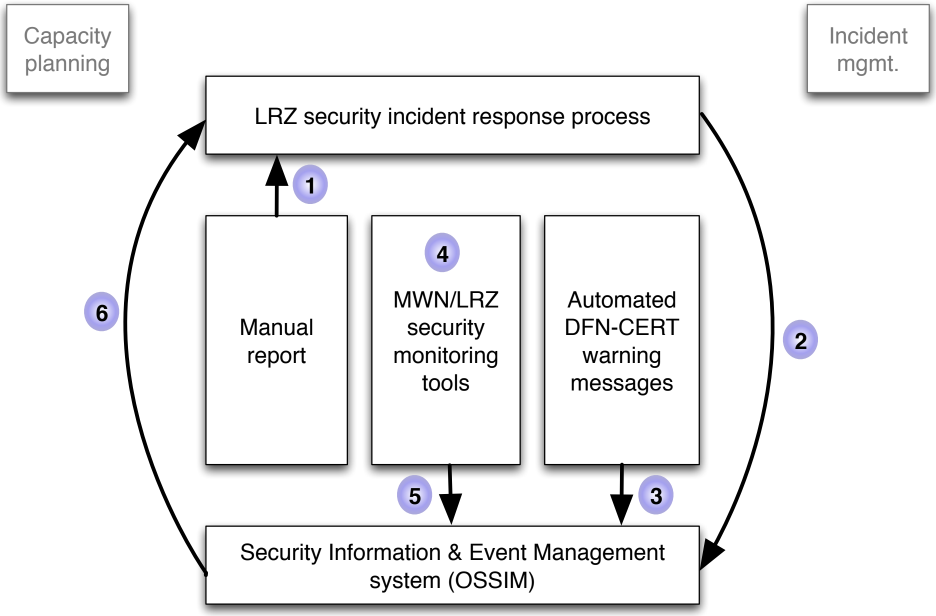
Capacity planning
Incident mgmt.
LRZ security incident response process
1
4
Manual report
MWN/LRZ security monitoring tools
Automated DFN-CERT warning messages
6
2
5
3
Security Information & Event Management system (OSSIM)

Figure 4. Integrated management of security incidents – LRZ example

In step 1, an email sent from an external Grid CSIRT team is forwarded to the LRZ-CSIRT, and a security incident ticket is created. In this email, some details about a security incident, in which different Grid nodes are involved, are summarized. All Grid sites (which received this email) are requested to look for network communication of suspicious IP addresses in their own part of the overall Grid infrastructure. On the Grid systems which are currently involved in the incident, the SSH daemon is replaced by a compromised version, which provides a backdoor-based administrative system access without any logging. Further conspicuities were not recognized so far.

The LRZ-CSIRT team starts the incident processing. At the initial classification a medium impact value is determined based on the classification scheme (outside the MWN; only a few Grid servers involved so far; root-compromised system through replacement of system files). Based on the provided information the LRZ-CSIRT starts to analyze the netflow data in step 2 (manual action). Communication is recognized between the reported suspicious IP addresses and two local systems. After that communication, these internal systems send a lot of traffic outbound-directed to various IP addresses. Hence the LRZ-CSIRT assumes that these local systems are compromised and inform the responsible administrators. Because these systems do not provide essential services, they shall be disconnected from the network to prevent further communication and a propagation of the incident. In parallel the integrated alerting mechanism provided by NfSen for analysis of the netflow data is activated to notify the LRZ-CSIRT if further communication between the suspicious IP addresses and local systems occurs.

At the same time the LRZ-CSIRT receives an email by the DFN-CERT's AW-Service (step 3); therein, the same two local IP addresses are reported. The event type is botnet communication, i.e., these IP addresses attempted to communicate with known command and control servers. This confirms the initial suspicion.

Now an analysis of the intrusion detection events is conducted. But filtering the IDS events by these IP addresses within the SIEM system shows no result although such command server communication should be monitored by some SNORT community rules. The LRZ-CSIRT decides to improve the available signatures with a self-developed one (step 4) based on the information provided by the AW-services email and the analysis of netflow data. Then, the CSIRT activates the integrated alerting and automated blocking mechanism for this special signature. Thus further communication between local systems and botnet command servers will be automatically detected. Additionally the system will be disconnected from the network and the responsible administrator and CSIRT will be notfified.

All currently available connection data will are exported and summarized in a report, which is attached to the security incident ticket. In parallel the first reporting CSIRT, i.e., the Grid CSIRT, will be informed about the compromised systems at LRZ.

The system administrators starts analyzing the systems. This encompasses the analysis of the provided log data as well as the inspection using selected forensic tools. Besides the replacement of the SSH daemon, as already mentioned within the first report, a bot software that has been installed by the attacker is found. This further information is summarized and forwarded to the Grid CSIRT to extend the first report. The Grid CSIRT can now forward this new information to all other Grid sites.

Two days later administrators receive an email informing them about an automated blocking of an IP address based on the special, self-developed SNORT signature (step 5). This system was shutdown few days ago and now restarted. The installed bot software attempts to contact the same C&C servers. Now the security incident has to be extended and this information is forwarded to the Grid CSIRT also. On the following day the email of DFN-CERT's AW-service confirms this conspicuity and thus the quality of the self-developed SNORT-signature.

Further communication between local systems and known command and control servers does not take place. Within the observation period of several days no more local systems become conspicuous and the security incident is closed successfully. Some insights based on the documented results lead to further improvement of the security incident response process and some further security measures are added on the local (Grid) systems.

## VI. Recommended proceedings for other higher education institutions

In this section a few recommendations and advices for other institutions concerning the implementation of an integrated security management approach are provided. Furthermore, some stumbling blocks and caveats are discussed. Information security should be considered as a holistic process divided into several tasks. Therein, endpoint and network security, but also, what is unfortunately often forgotten, employees or in general users of IT equipment and services are of equal importance. As is a well-known fact, absolute security cannot be reached and thus a level of security chosen based on required efforts and achieved benefits should be targeted. Following the basic need-to-know and least-privileges principles, the regular installation of currently available software patches for the operating system and all installed applications is not enough. Thus, other measures must be taken to minimize the resulting difference between the practically achievable and the theoretically best possible security; automation is a key issue in this regard.

The first step is the definition of a security incident response process, in which roles, responsibilities, and tasks have to be specified. Handling security incidents includes tasks of system administrators and users, but also network administrators, the public relations department, management, and optionally law enforcement authorities are involved. Thus, one person is needed who coordinates the whole process and delegates the various tasks, such as the communication with external (CSIRT) teams, forensic investigations, and the documentation of measures and results. These jobs must be assigned to trusted, security-experienced employees, such as members of an internal CSIRT team. Documentation starts with the reporting of the incident by a user or an administrator. The question what a security incident exactly is has to be answered by a clear definition, which makes a clear distinction to other issues, such as system misbehavior based on configuration errors. Possible reporting ways have to be defined and each user has to be aware of these ways. All information concerning a specific incident have to be gathered in a central location, such as a trouble ticket system. Stumbling blocks are, based on our experiences, the necessity to report a security incident and the support of different reporting ways. Some administrators solve the misbehavior recognized on their systems on their own and forget or avoid to inform the internal CSIRT team because they think they are fully responsible for the incident and expect worst-case consequences, or they simply do not know the reporting ways. Awareness trainings and official statements of the management could solve such issues. It should also be clear that first drafts and versions of the security incident response process have to be adapted sometimes. Regarding this aspect, embracing the ISO/IEC 27001 suggestion of continual improvement of the processes is a good and motivating solution.

After establishing the security incident response process (maybe one of the first versions), the monitoring tools have to be taken into consideration. Each security monitoring mechanism has its own event console in which events are displayed separately. Thus, a centralized security information and event management system should be installed and the security monitoring mechanism should forward their detected events (maybe additionally) to the SIEM system. As mentioned in the previous sections, SIEM provides the opportunity to (cross) correlate security events. Thus, a security analyst is able to get, supported by dashboards, a more detailed overview of the current security situation. Filtering and sorting functions allows to display potential relationships between individual events. Based on these events the analyst can inform the CSIRT team and provide details about the incident.

The next step is the definition of directives that (cross-) correlate security events automatically. No further manual intervention of an analyst is necessary. In such a case, a 24x7 event correlation is possible.

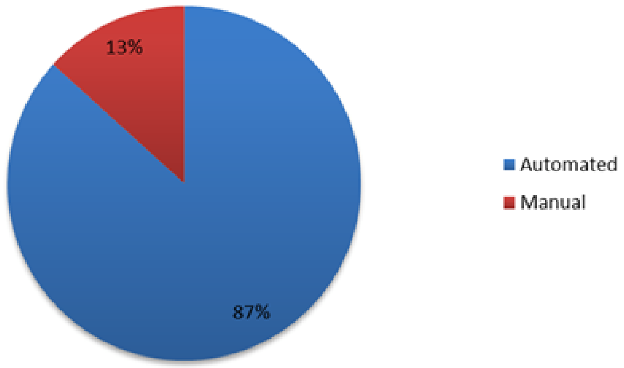

Figure 5. Percentage distribution - manual vs. automated processing

Based on the automated event correlation, actions can be triggered in the next step. Possible actions are the generation of a daily summary report, informing the analyst by email or pager, or creating a security incident ticket in real-time. The highest flexibility can be achieved if the action triggered by the SIEM is script-based, thus email notification, ticket generation, or even automated Internet or network access blocking are possible. The final step is triggering different actions based on a correlated event and context information. For instance, if an event occurs five times within normal business hours, the responsible system administrator is informed by email and a security incident ticket is generated. If the same event occurs five times outside normal business hours the network access of the compromised system is blocked, the administrator is informed by email, and a ticket is generated.

Besides, e.g., intrusion detection or firewall events (network security), also system logging events can be analyzed automatically and in a proactive manner. Through correlation of network and system events, combined with automatically triggered actions, a full 24x7 security incident response with limited personnel resources is possible.

## VII. Summary and outlook

The orchestration of various reporting capabilities, automatic analysis and response, and process-oriented intervention provides an effective and efficient approach for an integrated management of security incidents. Through a high level of automation, an adequate response in spite of limited personnel and outside the normal business hours is possible. As shown in Figure 6, this is urgently required at LRZ. Due to the proposed concepts and tools, more than 85 percent of all abuse cases can be (at least partially) automatically processed and handled as a *standard security incident*, as seen in Figure 5.

In sum it should be emphasized that for the quickest possible and structured response, the appropriate combination of external warnings, internal monitoring mechanisms, and internal notification by administrators is necessary. The AW-Service by DFN-CERT is an important part and supports our local security management. This service confirms the detection of compromised systems by internal tools and acts as an impulse to improve and extend the currently used sensors.

Additionally, various internal monitoring mechanisms can be combined suitably, with the analysis taking place centrally. Automated and flexible response capabilities are obviously essential. Automation also allows adequate response outside the business hours.

The main objective at LRZ is the enhancement and improvement of the currently used internal monitoring mechanisms and response capabilities. The number and impact of security incidents that occur have to be minimized. This can be achieved by the identification of compromised systems using internal detection methods as soon as possible. External warnings should never become necessary and only be used for confirmation purposes. Ideally, all security incidents are recognized locally and stemmed before external systems are affected.

We also plan an extension of the existing security mechanisms. One exemplary extension is the suspension of the switch port to which a compromised system is connected to, instead of only blocking the Internet access. Therefore, the spreading of malware within the MWN could be reduced. The implementation of a quarantine network in which only communication directed to specific servers (update servers of operating system and antivirus software) is allowed is also an ongoing project.

In the near future, a self service portal for system administrators will be taken into production. Using this portal, various tasks can be executed, e.g., determining the switch port based on Nyx or the re-activation of a suspended IP address. Automatic, scheduled reporting functions and the opportunity to configure thresholds for mail monitoring or maintenance

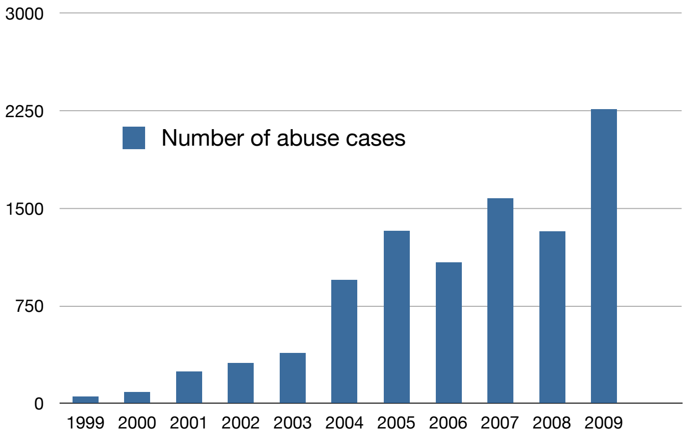

Figure 6. Number of abuse cases at LRZ per year [18]

of the exception lists complement the feature set. We are also planning on embedding the scan results of the DFN-CERT network checker in this integrated management approach.

## Acknowledgment

The authors wish to thank the members of the Munich Network Management Team (MNM-Team) for helpful discussions and valuable comments on previous versions of this paper. The MNM-Team, directed by Prof. Dr. Dieter Kranzlmüller and Prof. Dr. Heinz-Gerd Hegering, is a group of researchers at Ludwig-Maximilians-Universität München, Technische Universität München, the University of the Federal Armed Forces and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities. See http://www.mnm-team.org/.

## References

[1] Office of Government Commerce (OGC), *IT Infrastructure Library v3: Service Design, 2nd impression*, ISBN 978-0113310470, The Stationery Office (TSO), 2007

[2] R. Kornberger and H. Reiser, *Die Suche nach der Nadel im Heuhaufen — Nyx — Ein System zur Lokalisierung von Rechnern in grossen Netzwerken anhand IP- oder MAC-Adressen*, 21. DFN Arbeitstagung über Kommunikationsnetze, 2007

[3] E. Chickowski, *University Databases In the Bull's Eye*, http://www.darkreading.com/database_security/security/attacks/showArticle.jhtml?articleID=225702686&subSection=Attacks/breaches, 2010

[4] DFN-CERT, *Überblick über die Dienstleistungen des DFN-CERT*, https://www.cert.dfn.de/, 2010

[5] Leibniz-Rechenzentrum, *Das Münchner Wissenschaftsnetz (MWN) – Konzepte, Dienste, Infrastruktur, Management*, http://www.lrz.de/services/netz/mwn-netzkonzept/MWN-Netzkonzept-2010.pdf, 2010

[6] Leibniz-Rechenzentrum, *Nat-O-Mat: IP-Adressumsetzung (NAT) als Ersatz für Proxyserver*, http://www.lrz.de/services/netzdienste/nat-o-mat/, 2010

[7] Alienvault, *Alienvault OpenSource SIEM*, https://www.alienvault.com/products.php?section=OpenSourceSIM, 2010

[8] P. Haag,*NfSen and NFDUMP*, http://www.terena.org/activities/tf-csirt/meeting16/nfsen-haag.pdf, 2005

[9] C.Huitema, *Teredo: Tunneling IPv6 over UDP through Network Address Translations*, http://tools.ietf.org/search/rfc4380, 2006

[10] M. Hofherr and P. Wimmer, *Security (Information/Event) Management - Ein Praxisbericht*, 16. DFN-CERT Workshop Sicherheit in vernetzten Systemen, 2009

[11] Sourcefire, *Snort Official Documentation*, http://www.snort.org/docs, 2010

[12] S. McCanne and V. Jacobson, *The BSD Packet Filter: A New Architecture for User-level Packet Capture*, USENIX Winter, 1993

[13] D. Fliegl and T. Baur and B. Schmidt and H. Reiser, *Ein generisches Intrusion Prevention System mit dynamischer Bandbreitenbeschränkung*, 20. DFN–Arbeitstagung über Kommunikationsnetze, 2006

[14] ISO/IEC 20000-1:2005, *Information technology — Service management — Part 1: Specification*, 2005

[15] ISO/IEC 20000-2:2005, *Information technology — Service management — Part 2: Code of Practice*, 2005

[16] ISO/IEC 27001:2005, *Information technology — Security techniques — Information security management systems — Requirements*, 2005

[17] ISO/IEC 27002:2005, *Information technology — Security techniques — Code of practice for information security management*, 2005

[18] E. Bötsch, *Bearbeitung von Abuse Fällen*, http://www.lrz.de/services/security/abuse/, 2010