DOI: 10.7250/eunis.2013.006

IT security concept documentation in higher education data centers: A template-based approach

Wolfgang Hommel[1], Stefan Metzger[2], Helmut Reiser[3], Felix von Eye[4]

[1-4] Leibniz Supercomputing Centre, 85748 Garching, Germany, {hommel,metzger,reiser,voneye}@lrz.de

Keywords
Information security, ISO/IEC 27001, document management, compliance.

1. ABSTRACT

Data centers of universities and IT departments of smaller higher education institutions provide dozens of IT services such as email, web hosting, e-learning, and file storage. The number of server machines and appliances that need to be operated often reach three-digit numbers depending on the number of services, users, and high-availability setups. Many services can be used via the Internet to improve usability. As one of the consequences, many servers are subject to Internet-based attacks. Typically, security mechanisms such as firewalls and intrusion prevention systems are used to counter these attacks. However, in practice still a lot of server machines get compromised, e.g., due to vulnerabilities in server software that is not patched fast enough, or due to improper configuration of the software running on these machines.

In an ideal world, there would be enough IT personnel to operate all these IT services, and each IT administrator would also be an IT security specialist who knows exactly how to make his or her own servers almost perfectly secure. In reality, however, often a very small IT staff needs to run more servers than can easily be handled, and IT services such as groupware or e-learning systems have become extremely complex regarding their core functionality. Consequently, administrators only have diminishing resources, i.e., time and know-how, to properly secure their IT services. Specialization then typically leads to the foundation of dedicated security teams, such as CERTs (computer emergency response teams) and CSIRTs (computer security incident response teams). While those security teams consist of security experts, their primary problem is a lack of in-depth-knowledge about all those IT services and their specific configuration. In order to facilitate the security team's efficient handling of, for example, security incidents in an e-learning service, knowledge transfer from the e-learning administrator to the security team about the specific setup must be fostered, and accurate knowledge must be available on-demand, for example, if a security incident happens while the service administrator is on holidays.

In theory, each IT service should be properly documented along with all of its operational and security-specific properties, and this documentation should always be kept up-to-date. In reality, most IT administrators have no time to write documentation, dislike this task, and often do not even know what should be documented in a service-specific "IT security concept". Therefore, many security incidents are handled in a patch-on-demand manner: Once a service has been compromised by an attacker, it is set up again, e.g., from a clean backup, and minimum configuration changes are applied to prevent the same type of attack from being successful again. While this approach is somewhat pragmatic, it obviously cannot be considered as a good and sustainable solution.

We present a template-based approach towards the documentation and management of IT security concepts tailored to the demands of real-world IT service operation in higher education institutions. Our documentation template is intended to be filled in easily, provides a uniform document structure across many types of IT services, encourages IT administrators to think about IT security target-oriented, and supplies security teams with the information they require for security incident handling. Its contents are based on security standards and good practices, such as ISO/IEC 27001, ITIL v3, and the "IT base protection catalogues" by the German Federal Office for Information Security. We are working on a web-based management frontend that makes it easy to initially write, update, access, and utilize the security concept documents, which are stored in a repository that also serves as a foundation for an inter-organizational exchange of IT security concepts.

DOI: 10.7250/eunis.2013.006

2. MOTIVATION FOR TEMPLATE-BASED SECURITY CONCEPT DOCUMENTATION

As it may not be obvious why a template-based approach towards documenting security concepts is the solution we chose, we first discuss our motivation for creating the template presented below.

For almost each IT service there is generic literature about how to make it more secure, for example, as a part of the server software product documentation or available as "security how-to" from the Internet. However, most IT services are configured specifically for the environment they are operated in, for example by activating additional features that are not turned on by default, or by coupling the service with the university's LDAP server or identity management system for user authentication and authorization. Therefore, usually a lot of security-relevant settings are unique to the specific instance of the IT services.

Documenting these specifics requires input from those people who are most familiar with the local installation and configuration, i.e., the responsible IT administrators. Given the number of IT services, their dynamics (e.g., regarding updates and configuration changes), and personnel constraints, we assume that there is *no* dedicated technical writer available who documents security concepts based on, e.g., interviews made with these IT administrators. Instead, the IT administrators shall be enabled to document their services' security concepts on their own. Of course, security experts may assist them with the details and implementation of service-specific security mechanisms, but the documentation task shall be easily enough handled by each IT service's administrator.

Unfortunately, writing a security concept is not an intuitive task for most people. Given this duty, many IT administrators will not know where to start and what to put into such a document. Internet search engines do not help much either, because almost nobody seems to be willing to publically share detailed security concepts for various comprehensible reasons, and there hardly are any useful generic exemplars or copy-and-paste-boilerplates either. The consequences are disenchanting:

- People tasked with documenting a security concept often become frustrated and give up or procrastinate.
- If they struggle through it, they have to start from scratch, which consumes a lot of time and causes several wheels to be re-invented.
- Each security concept will look different in terms of, e.g., structure, contents, and depth of details. This makes it harder, e.g., for security teams to get an overview of all IT services and to quickly extract required information, e.g., during the handling of a security incident.

Providing IT administrators with a documentation template remedies these issues and provides additional benefits:

- The security concept template, even when not filled in but only read, presents security-related topics and ideas that IT administrators should consider when operating their IT services. This gives thought-provoking impulses even if the IT administrator has focused on the core functionality of the service only so far and not yet dedicated much time to security-specific aspects.
- Given his or her knowledge about the IT service, the IT administrator can start documenting the security mechanisms right away without having to worry first about what content should be put into it. Of course there will be details for each service that need to be amended based on service- or product-specific knowledge, but many sections will be contained in each and every security concept.
- The resulting documents follow a uniform structure, which makes reading and extracting specific information easier, e.g., for the members of a local IT security team that has access to the security concepts of all IT services.

Of course, template-based security concepts inherit the vantages and disadvantages of any type of service documentation: On the one hand, they make it easier, for example, to bring holiday replacements up to speed, but on the other hand, they become useless or even dangerous if they are not updated appropriately, and there should be a review-and-approval process to ensure the correctness and quality of the created documents. In the following sections, we first outline selected contents of our security concept template and then discuss our ongoing work on a web-based management frontend and the processes it facilitates.

DOI: 10.7250/eunis.2013.006

3. SECURITY CONCEPT TEMPLATE: CONTENT AND INFILLING

Providing a documentation template inherently is a trade-off between the discussed demand for uniformity and a desired degree of individuality in order to map the IT service specifics to its documentation. The template and its exemplary content, which is presented here, is the result of several refining iterations based on its practical application to several IT services operated by the Leibniz Supercomputing Centre (LRZ) in Munich, Germany. Several input sources have been used to elicit the requirements of the template's content and structure, including the following:

- Interviews with the security team members: Which security-related information is needed to get a good overview of the documented service and which information is most relevant to handle acute security incidents?
- Interviews with IT service administrators: What is the best layout for each section of the template, i.e., when should check boxes, select lists, or free text forms be used in order to enable quick infilling?
- ISO/IEC 27000: Within this family of standards, ISO/IEC 27001 defines requirements for information security management systems (ISMS). Its normative appendix A defines more than 130 security controls, grouped by so-called control objectives, which organizations shall consider for implementation. Most of them are organizational measures, such as assigning responsibilities and creating as well as enforcing security policies, but several of them are quite technical, such as ensuring that timestamps in log entries are based on synchronized server time settings. Therefore, ISO/IEC 27001 can be refined to create a useful checklist about basic, but very useful and highly recommended security measures for each IT service. Furthermore, appendix E of ISO/IEC 27005 provides an overview of typical threats that need to be considered when securing an IT service.
- ITIL v3 and ISO/IEC 20000: ITIL is a very comprehensive good practice documentation for IT service management (ITSM), whereas ISO/IEC 20000-1 is quite a compact standard for this topic. In practice, security management always needs to be closely intertwined with many other ITSM processes: For example, server machines and software installations are managed by the configuration management process, which is supported by a configuration management database (CMDB, or configuration management system (CMS)) that can be used to look up, e.g., on which physical machines a server software is running. ITSM also covers change management, a process that, among many other aspects, covers software update handling; when considering the fact that a lot of software updates become necessary due to security patches, it becomes obvious that change management must also be considered when a security concept is developed.
- IT base protection catalogues: The German Federal Office for Information Security provides an extensive collection of information security best practices. With their focus on technical security measures, they complement ISO/IEC 27001 very well and can be used, for example, as a checklist for hardening Linux and Microsoft Windows servers.
- Reviews of existing security concepts and security incident records: We analyzed several security concepts that had been created at LRZ in the past without any templates in order to determine what typically has already been documented in the past. We also reviewed notes about how security incidents have been handled at LRZ in the past in order to figure out which pieces of important information were typically necessary and eventually hard to retrieve without a documented security concept.
- Third party documents: We requested security concept documentation from other university data centers and used Internet search engines to skim over countless security documents and security recommendations for various IT services in order to identify topics that may be useful to include in security concept documentation.

For obvious reasons, the aggregation of topics from these inputs results in a very bulky and highly redundant collection of security-related material. Putting all information in each security concept would result in very large documents that take a lot of time to write, and only a fraction of the information therein would be required for most practical use cases, such as handling urgent security incidents. Therefore, we made several iterations to set priorities, marked certain topics as optional, omitted topics that we determined to be of less importance, grouped topics, and brought everything remaining in a specific order, which we consider a trade-off between an intuitive writing order and

DOI: 10.7250/eunis.2013.006

the desire to have the most important information for handling security incidents in a compact format at the very beginning of the security concept documentation.

We present selected content of our security concept documentation template below, which we consider the most interesting information for other organizations. It quickly becomes obvious that certain topics and content parts are not purely security-specific. For example, a short description of the IT service whose security concept is documented is typically already available on the data center's website, in a service catalog, or in other internal service documentation. We consider this information relevant for a security concept, but suggest the use of hyperlinks to avoid redundant documentation. For example, our security concept document management web-frontend, which we outline below, can fetch the names and IP addresses of servers, on which a certain IT service runs, from a CMDB at run-time, so it is not necessary to document and update this information manually. This means that on the one hand, we want to make the template to be filled in as easily and quickly as possible, but on the other hand the writer is responsible for ensuring that readers know where to find the information that has not explicitly been documented in the security concept. It must also be kept in mind that a ready-made template may deter some administrators from thinking about additional, specific security aspects of their service, which may lead to the undesired state that only the standard security measures suggested by the template will be implemented for new services. However, since the documentation template is only one small building block of the overall information security management strategy, its benefits outweigh this disadvantage in our opinion, and along with other measures such as trainings to improve security awareness, we are confident that the overall level of security can be increased by using our template.

3.1. Security concept documentation: Table of Contents

Our security concept template is available from http://git.lrz.de/secdoc and structured as follows:

- **Metadata**, including the name of the documented IT service, author and reviewer information, date and version of the document, and authoritative storage location (where to find the most recent version).
- **Security overview**, including contact information, hardware and software overview, classification of the data processed by the service, service dependencies, service criticality, and service-specific risks.
- **Security measures description**, including software update management, the application of security-specific software such as anti-malware programs, identity & access management for administrators and regular users, data protection and privacy, network security measures such as service-specific firewall configuration, logfile management, and backup / restore procedures.
- **References and notes**, including links to other documentation, such as vendor security recommendations, instructions and hints for the security team in case of emergencies, and an outlook to planned changes and improvements that have not yet been implemented.

The content of each document chapter is discussed in more details in the following sections.

3.2. Metadata

The metadata section of our security concepts is quite similar to other, not security-related documentation and mostly intended to support document management. On its title page, a security concept documentation based on our template includes the following information:

- Document name, i.e., security concept for <*service name*>.
- Author name(s): Who contributed to the documentation?
- Information classification: Security policies at LRZ require each document to be classified, e.g., as *internal*, along with a list of authorized individuals or groups. Typically, security concepts are made available to the organizational group operating the IT service, the security team, and executive management.
- Names of individuals who have reviewed and approved this version of the document: Our document management workflow specifies that an individual who is not the author of the document verifies its technical correctness; this is usually done by another IT administrator of the same service, eventually the designated holiday replacement. The document release

then needs to be approved, e.g., either by the head of the responsible department or by a designated security team member.
- Authoritative storage location: Unless a central security concept repository is used, where can the most recent version of this document be found?
- Version number of the template that has been used for the document. Based on feedback from those who use it, we continue to improve our template, but do not expect that all security concepts are re-written or updated when we release a new template version. The information which template version has been used is useful to determine, for example, whether a certain section is missing because it is too new or because it has explicitly been left out for the specific service.
- Date of the document's last modification, and deadline until which the document has to be reviewed and a new version has to be released. In order to avoid outdated documentation, we have a policy to review and update documents at least every six or twelve months, depending on the type of service.

The beginning of the template also includes a generic introduction about how the template is intended to be used, typographic conventions, suggestions on where and how extensions or deliberate omissions should be made, and contact information about the parties that can assist the security team or answer questions about the template.

3.3. Security Overview

The first chapter that needs to be filled in by the authors of a security concept assembles basic information about the documented IT service and contains the information that typically is most relevant for handling acute security incidents. Our security incident response team is made up of seasoned security practitioners across all data center departments, which means that the security officer on duty who has to handle a security incident is not necessarily familiar with the compromised service or machine. Therefore, the security overview starts with generic information about the service and then turns more and more security-specific:

- Short service description, e.g., "Dovecot IMAP server for student email accounts". Instead of repeating a more detailed description from other documents or websites, links to a more detailed should be provided.
- Contact information: Who should be contacted with questions about this service regarding, among others, the server hardware, the operating system, the service software and its configuration, and reporting security issues? Typically, colleagues will be named, but depending on service contracts also external help desks or important "customer" data may be mentioned here.
- Server information: Which physical and virtual machines (VM) provide the documented service? Details such as DNS names, IPv4/IPv6 addresses, room names and rack locations or VM hosts must be given. Redundant hardware, e.g., for high-availability purposes, and different instances of the service, e.g., separate testing and production environments, must also be specified. Ideally, this information can be retrieved from a CMDB and does not have to be entered manually.
- Software information: Which operating system and service software are used in which version? Again, this may be retrieved from a CMDB, or otherwise has to be updated manually when, e.g., major software upgrades are performed.
- Data classification: Which classes of information does the service process? For example, if a machine which processes personal data, such as a campus management system, is compromised, security incident handling differs from cases when, e.g., a web server containing only public information is attacked.
- Service dependencies: This section details which other services the documented service depends on and vice versa. On the one hand, this includes, for example, NAS volumes on a file server, LDAP servers for user authentication, and service load balancers that must be up and running in order to make this service work. On the other hand, other services may rely on the availability and integrity of this service, so if one of its machines gets compromised, it might affect other services as well. Again, this information can ideally be extracted from a CMDB, which stores each service as so-called configuration item (CI) along with dependency

links between the CIs, or otherwise should be specified manually in an appropriate granularity. For example, if a central Microsoft Active Directory service is documented on which dozens of other Microsoft-Windows-based services rely, it may not be feasible to know and name them all, so highlighting only the most important ones and providing a more generic description of which other services use this one may suffice.

- Service criticality: This section describes service usage in more detail and includes information about
    - whether the service is open to public, available on campus only, or data center internal;
    - when the service is typically used the most, for example, during office hours (e.g., mail server) or during the night (e.g., backup and archival services);
    - approximately how many registered users the service has and how many concurrent users there are during peak hours;
    - which other services or machines must be considered compromised if this service has been compromised (consider, for example, the potentially fatal consequences of hacked LDAP servers which authenticate and authorize all users, or manipulated software installation repositories for local Linux or Windows servers).
- Data criticality: A rating of the service's confidentiality, integrity, and availability requirements. For example, a public web server has lower confidentiality requirements than a student and exam management system, and the availability of the campus-wide email system may be more important than the availability of a file server operated for a small student lab.
- Service-specific risks: A description of potential attacks against this service, which are either very typical or even unique for this service. More generic risks, such as fire in the data center or long electrical power outages, should not be discussed in each single service security concept, but are the subject of organization-wide risk management and business continuity management. However, the administrator of a web server might point out vulnerabilities in the used content management system as likely attack vectors, and the administrator of a learning management system may think of insider threats, such as students who try to cheat at electronic exams. Ideally, the descriptions made here put the security team on the right track about what kind of attacker is most likely behind the attack and how the service got compromised once a security incident has occurred.

Except for the service-specific risk analysis, our template consists of tables, checklists, and cloze texts that can easily and quickly be filled in. For example, we provide a list of more than ten central IT services, such as LDAP or syslog servers, that are typically used, so the document author only needs to tick the appropriate checkboxes and eventually provide some additional information, such as the name of a NAS volume that is mounted on a Linux server.

3.4. Description of security measures

The second chapter is the core of each security concept documentation and provides insight into the security-specific configuration of the IT service. To get the authors hooked on writing, this chapter again starts with more generic information, such as how operating system and other software updates are performed. Those procedures will typically be the same for most of the machines which the security concept author takes care of, and therefore are easy to fill in. The template then moves on to more security-specific tools and settings on both the machines that provide the IT service as well as network-based protection such as firewalls. It is structured as follows:

- Operating system and software updates:
    - Who is responsible for OS and service software updates? This may be, e.g., a dedicated Linux server team, the service administrator, or an external third party with a managed service contract.
    - Are updates done manually or automatically? When will necessary reboots be performed, e.g., after a Linux kernel update? Are certain software packages explicitly excluded from the regular update mechanism?
    - Are updates tested (on a separated machine) before rollout?
    - Which channels and media are used to receive update and security news for the deployed software? For example, vendor newsletters or security mailing lists.

DOI: 10.7250/eunis.2013.006

- Dedicated security software:
  - Is anti-virus/anti-malware software installed? What happens if malware is found, i.e., who is notified, is the file deleted or put into quarantine, etc.?
  - Is host intrusion detection or other attack detection software installed? For example, is a host IDS system like *samhain*, a system compromise detection software like *chkrootkit*, or a denial-of-service protection software like *fail2ban* used and how is it configured?
  - Does the service software use server certificates, e.g., to authenticate SSL/TLS connections? If so, are self-signed certificates used, or is the PKI operated by the national research and education network provider used? How is it made sure that expiring certificates are renewed in time?
- Identity and access management:
  - How do administrators connect to the service? For example, is one of the dedicated management SSH gateways or a management Windows terminal server used? Is a dedicated management software necessary or are there web-based management frontends?
  - Which individuals or groups do know the administrator authentication credentials, such as passwords? Are those stored in the organization-wide emergency password vault? Do external third parties need privileged access, e.g., for maintenance?
  - Are there any default accounts, such as *root* or *admin*, and have their passwords been changed?
  - How is user authentication performed? For example, are password hashes stored locally, is a central LDAP or Active Directory server used, or are SSH keypairs or user certificates in use? Is single sign-on, e.g., based on Shibboleth or Kerberos, supported?
  - How are new user accounts set up and how are old accounts deleted?
  - How are permissions and authorizations assigned to user accounts? Are, for example, role-based access details stored locally within the service, or are they fetched, e.g., via LDAP or Shibboleth?
- Communication network security:
  - Which of the local network zones (VLANs) is the service assigned to?
  - Have dedicated rules for this service been set up in the organization-wide firewall or router access control lists?
  - Is a "personal firewall" in use, such as the Windows firewall or *netfilter*/*iptables* on Linux?
  - Is management access, e.g., via SSH, restricted to dedicated management gateways or selected data-center-internal networks?
- Data availability and privacy:
  - Are hardware or operation system level data availability mechanisms used, e.g., multiple machines, RAID for physical drives, live migration techniques for virtual machines, or NAS file system snapshots?
  - Are backups performed regularly? Is the organization-wide default backup-to-tape software used or is a service-specific backup software required to ensure the backup's data consistency?
  - Does a recovery plan exist, which details how restoring the service from backup works exactly? When has this procedure been verified to work?
  - According to German data protection and privacy laws, a so-called process description must be created for the service. It documents, among other topics, which personal data (such as name and email address) is processed by the service for which purposes. Where is this document stored and when has it been approved by the organization's chief privacy officer?
  - Logfile management:
    - Where are system and service logs stored, e.g., locally on the server machine or on a central *Syslog* server?
    - Is the log data retention in compliance with the organization's security policy?

- Is the system time used for timestamping log entries synchronized, e.g., using the organization's NTP server? Is the correct system time verified regularly?
- When and how are logfiles evaluated? For example, are they analyzed manually after incidents have been reported, e.g., by users? Are automatisms, such as scripts or reporting engines, used to create, e.g., usage statistics?

It is noteworthy that most of the topics in this chapter of the security concept template are verbalized in interrogative form. If the documentation author answers *yes* to a question, some more details about the specific security measure have to be filled in. If an answer is *no*, because, for example, no anti-malware software is used, we ask for a short free text explanation why the mentioned security measure is not deemed necessary for this service.

The answer sections in our template consist mostly of checkboxes and selection lists, and they also mention typically used security software by name. While we have no intentions of making all IT administrators use exactly these tools, we want to point out different categories of security mechanisms, such as denial-of-service protection and system compromise detection, in case the IT administrator has not yet thought about these security aspects for the documented service yet. Naming a few tools that are successfully in operation for other services has worked out well as a trigger for further improvement of a service's security mechanism landscape.

3.5. References and notes

The final chapter of the security concept template assembles any additional security-specific material and annotations that do not fit into the overview and security measures documentation. Document authors are requested to elaborate on the following topics:

- Related documents and references:
  - Are there other organization-specific documents, such as a service operations concept or documentation for users, from which additional information can be gathered?
  - Where is vendor-created hardware and software documentation stored?
  - Is the local service configuration aligned to security recommendations that can be found in vendor documentation or what are the URLs of websites with such information?
- Annotations for the security team:
  - Is there anything else that the security team should know in case of emergency?
  - Are there any "do's and don'ts" to react appropriately when none of the contacts specified in the overview chapter can be reached?
- Planned changes to the security configuration: Have security issues been identified or are there any plans for, e.g., the application of additional security tools, which have not yet been resolved / implemented due to time and effort constraints?
- Have there been security incidents in the past that have influenced the current service design and configuration security-wise?
- Who is responsible for implementing the documented security concept and how much time is required for initial implementation and continuous improvement?

Although this last topic is marked optional, we consider it a good idea to document how much time and effort is required to secure one's IT services. More often than not "customers", users, and even some superiors assume that IT services are secure by default and expect the IT administrators to focus on other aspects, such as functionality and performance. However, for personnel planning and security management decisions it is indispensable to know how time-consuming proper security operations management is in practice.

This concludes the overview of our security concept template's content. In the next section, we discuss how documented security concepts can be managed on an organization-wide scale.

DOI: 10.7250/eunis.2013.006

4. SECURITY CONCEPT DOCUMENTATION MANAGEMENT WORKFLOW

In our experience, higher education data centers have, not unlike many other organizations and enterprises, very heterogeneous approaches towards document management. We have seen security concepts stored on file servers, uploaded to internal websites via content management systems or Wiki servers, stored in versioning systems like SVN or git, and distributed solely via email.

In addition to the proper selection of a storage location along with the management of appropriate access permissions, a template-based approach requires the template to be delivered in a format that makes it easy to work with. Many of our colleagues work with text processing software, such as Microsoft Word, while some more open-source- and research-oriented colleagues prefer LaTeX for writing documentation. And yet others prefer web-based forms and want to fill in their texts using a web browser.

In the beginning, we used a LaTeX document and applied the *latex2rtf* and *latex2html* tools to create versions for Microsoft Word and a Wiki template. However, the resulting templates are static, which means that depending on whether *yes* or *no* is answered to one of the questions asked during filling in the template, the follow-up questions of the other answer either have to be deleted manually or remain as half-empty text fragments in the final documentation. Meanwhile, we have started working on a web-based management frontend for security concepts based on our template. This approach delivers the following benefits:

- A central repository, such as a file system, relational database, or versioning system can be used to store all security concepts.
- Along with the central storage, access and permission management becomes easier.
- Web-based text entry can make use of dynamic forms: Depending on the answer chosen to one question, only the relevant follow-up questions will be asked.
- Workflow support can be implemented: For example, the review, approval, and release process can be triggered when an updated version of a security concept has been submitted. The management tool can send email reminders when the deadline for the semi-annual document review is approaching.
- The security concept authors implicitly always work with the most recent version of the template. Authors can be informed whenever a new template version is released, which may, for example, introduce additional topics and sections that should also be filled in for existing security concepts.
- Security concepts can easily be searched and evaluated based on arbitrary criteria. For example, statistics about services that still store user passwords locally instead of using the central LDAP server can be compiled. This helps risk management and the security team to determine, which new and additional security measures should be taken.
- Security concepts can be exported in various formats, including HTML, plain text, and PDF.
- The overhead for releasing a new version of the template or any of the security concept documents created with it becomes significantly lower because, for example, creating a PDF, emailing it for approval, and publishing the final version via a content management system are no longer necessary.

We also plan to make different views on the same security concept possible. For example, shortened versions of a security concept could be exported from the management tool in order to make it available to other academic data centers without revealing information that must be kept strictly internal. If multiple data centers operate services in a similar manner, they could share those common parts of their security concepts in order to further reduce the effort required for documenting them.

In the long run, such functionality could also be used to create (optionally anonymized) comparisons, statistics, and benchmarks of the security state-of-the-art across higher education data centers.

The web-based management frontend is currently in development and not yet fully usable. We plan to make it available as open source via http://git.lrz.de/secdoc once it has matured.

DOI: 10.7250/eunis.2013.006

5. SUMMARY

Documenting the security concepts and security-specific configuration aspects of IT services is a tedious, often only reluctantly performed task that is, however, essential for organization-wide knowledge management and the very basis for a security team when handling security incidents. We are convinced that a template-based approach simplifies this matter for both the authors and the readers of security concept documentation.

However, a lack of suitable security concept blueprints has motivated us to create a template of our own, which is aligned with the specific needs of higher education data centers and the services they typically provide. We compiled relevant topics and material from various sources, including the ISO/IEC 27000 standard series, IT service management best practice documentation, our own previous experience, and various documents published on the web. We grouped and prioritized the results and created a new security concept documentation template, which since then is used in our production environment and continually improved based on feedback from our colleagues who use it as their first choice when it comes to security documentation.

In this article, we first presented our motivation and course of action. We believe that our document template is also useful for other higher education data centers and described its structure and selected content in some details. We also outlined how we make use of tables, checkboxes and selection lists to make filling the template in as easy and efficient as possible. Finally, we sketched the functionality and benefits of a web-based security concept management frontend that we currently develop; it will be released as open source software.

6. AUTHORS' BIOGRAPHIES

Wolfgang Hommel is the Chief Information Security Officer of the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities, where he is also the head of the communication networks planning group.

He studied computer science at Technische Universität München and has a Ph.D. as well as a postdoctoral lecture qualification from Ludwig-Maximilians-Universität in Munich, Germany, where he teaches information security lectures and labs. His research, for which he was granted the Karl Thiemig foundation's young academics award in 2011, focuses on information security and IT service management in complex large-scale and inter-organizational scenarios.

Stefan Metzger is a member of the communication networks planning group at the Leibniz Supercomputing Centre. He holds an international CISSP certification. As head of the LRZ security working group his main focus lays on ISO/IEC 27000-based security management.

He attained in 2005 a Diploma degree in Computer Science from Technical University Munich. Currently he's doing his PhD in security management in large-scaled, inter-organizational infrastructures and offers lab courses in security at Ludwig Maximilians University (LMU) Munich.

Helmut Reiser is the head of the department of communication networks at the Leibniz Supercomputing Centre.

In 1997 he attained a Diploma degree in Computer Sciences from Technical University Munich, in 2001 a PhD degree and in 2008 the lecture qualification (venia docendi) from Ludwig Maximilians University (LMU) Munich. He is a member of the faculty and doctoral thesis committee of the department of Mathematics, Statistics and Computer Science at LMU and offers lectures and lab courses in security. His research comprises security and IT-management in large-scaled open and inter-organizational infrastructures.

DOI: 10.7250/eunis.2013.006

Felix von Eye is a member of the communication networks planning group at the Leibniz Supercomputing Centre and works in different research projects in the area of network management, IT security, and intrusion detection.

He studied computer science and mathematics and attained a diploma degree in mathematics from Augsburg University in 2009. Currently he is working on his PhD in dynamic and automatic threshold methods for intrusion detection systems. At the Ludwig Maximilians University Munich and the Technical University Munich he supervises different theses and is involved in the security lectures.

Background information

The Leibniz Supercomputing Centre (LRZ) is the common IT service provider of the universities and higher education institutions in the greater Munich area, Germany. It offers IT services for more than 130,000 regional users and operates the Munich Scientific Network, which is part of the German national research and education network. LRZ is also one of Germany's national supercomputing centres and offers high performance computing services to scientific communities across Europe.

Acknowledgment

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful 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.