Metrics and Best Practices for Host-based Access Control to Ensure System Integrity and Availability

Urpo Kaila, Marco Passerini and Joni Virtanen, *CSC - IT Center for Science*

ABSTRACT: *Open access in academic research computing exposes the servers to many kind of brute-force attacks and vulnerability exploits. The system administrator has a delicate task to similarly ensure system integrity by proper access controls and by applying security patches but also to enable service availability and ease of use. This paper will present an analysis of aggregated log metrics for access history and service up-time processed with tools as Nagios and Splunk in conjunction of a set of cases of vulnerabilities, intrusions and faults. The paper will also compare and suggest improved best practices to be shared between sites.*

KEYWORDS: Information Security, Access Control, Availability, System Integrity, Monitoring

1. Introduction

CSC as provider of computing services

This public paper presents metrics and best practices for host based access control at CSC - IT Center for Science Ltd, an organization in Finland. As an introduction to the topic, some background facts are first presented about the hosting organization and its security requirements.

CSC - IT Center for Science Ltd. (CSC) is a non-profit limited company administered by the Ministry of Education, Science and Culture in Finland. Having core competences in modeling, computing and information services, CSC provides versatile IT services, support and resources for academia, research institutes, and companies. The Funet communication links provide research workers with Finland's widest selection of scientific software and databases and Finland's most powerful supercomputing environment [1].

CSC has provided computational and network services since 1971, when a Univac 1108 mainframe computer was installed on the premises of the Finnish State Computer Centre.

CSC employs over 200 people and had during year 2010 a turnover of 25.7 MEUR. CSC provides services to its customers – universities, polytechnics, research centres, public sector including Ministry of Education and industry – in five areas:

- Funet network services
- Computing services
- Application services
- Data services for science and culture
- Information management services.

Compared with other academic computing centres and classical National Research and Educations Networks (NRENs) [2] abroad CSC is a quite resourceful actor, especially considering the population base of Finland. Reasons for a steady growth and development over the years have been combining provisioning of both computing and NREN services to one agency and operating actively internationally, especially in EU projects. Also, CSC has lately found new areas for growth in data services for culture and in some selected information management services for public administration.

Despite its role as a limited company, CSC performs most major provisioning of IT systems with specific government funding and according to the Finnish Law of

public procurement. During 2011 the major CSC projects are the provisioning of the next advanced HPC system and a new data center in the city of Kajaani, in the north-eastern parts of Finland.

Security Requirements

To secure its computers, data and services against inappropriate risks CSC, has organized measures to ensure information security based on risk assessments, best practices and requirements for compliance and good governance. The principles for organizing information security at CSC have been defined in the information security policy [3] of CSC and related decision made by CSC senior management.

The external information security requirements for CSC consist of several sources. Some of the requirements affect CSC as an organization, while other requirements are more service specific. In addition to several laws related to information security and privacy, the sole owner of CSC, the Finnish Government, has developed guidelines for compliance with information security requirements. Currently, the most distinct of these guidelines, is the Government Manual for Information Security Levels [4], which is a wide management and maturity framework, partly based on the Information Security Management Maturity Mode ISM3 and the ISO/IEC 27001/27002 standards.

CSC has also been externally audited for consecutive years based on the requirements of the Government Information Security Level Manual.

Service specific security requirements include common security policies and guidelines for grid computing. CSC is involved in several interconnected grid infrastructures: DEISA (Distributed European Infrastructure for Supercomputing Applications), EGEE (The Enabling Grids for E-sciencE) and PRACE (The Partnership for Advanced Computing in Europe). Common criteria and operative guidelines are especially crucial for all peers in grid computing, as user from other grid sites can use CSC grid computing resources.

The technical and operational implementation of the security requirements is based on professional knowledge and skills but also best operational practices of information security. Without technical and operational implementations the security principles are just mere rhetoric.

Risk analysis

Except compliance, information security should be foremost based on risk management, by which the resources to be protected are identified, risks are identified and assessed, and finally risks are mitigated by security controls.

At CSC risk management is performed on both corporate level and on service level. The corporate level risk management programme has a wider approach which assesses strategic, operational and damage risks. The service based risks assessments are currently documented at CSC in the business continuity plan of each critical service.

In both corporate level and service level risk assessments CSC has identified a number of risks related to computing services, as presented in Table 1.

Selected Corporate level risks

- Chain reaction and service outage due internal service dependencies
- User Administration unavailable, not possible to log in
- CSC service platforms or networks unavailable or compromised
- Loss of data
- Failure to deliver or deliver with low quality and/or delays in external projects
- The provider fails to deliver the system, application or service in the agreed schedule
- The system, application or service has high fault frequency
- Fire or smoke in Data Centers
- Flood, Water damage in Data Centers
- Power failure in Data Centres
- Inoperability of Data Center cooling system
- Compromise of CSC Systems due vulnerability
- Compromise and / or exploit of an account (key, proxy, user account)
- Misuse or malicious use of admin accounts
- Turning off security measures for troubleshooting or other reasons

Risks related to Computing Services

- Stealing of, misuse, abuse of user accounts or keys, compromised passwords
- Infrastructure and data centre issues
- Physical or logical corruption of storage systems
- Compromise of administrative infrastructure

Table 1. Selected Corporate level risks

Confidentiality, Integrity and Availability

Amazingly, many persons outside the information security community, still thinks, that information security is primarily only about confidentiality and access control. That is a very narrow view which can isolate information security from the main business and service objectives of the organization delivering services. The current, well known main stream view [5] of information security is that should span in a balanced way the requirements for confidentiality, integrity and availability of the systems and services to be protected. A wider concept,

information assurance, is sometimes used to emphasize the wide ranges of information security.

To protect the confidentiality of computing systems and services without regard to the availability of the systems and the service provided by it would be trivial – just put the system off line. The real world challenge is naturally to optimize service and system availability in a flexible and user friendly way by simultaneously ensuring adequate security measures for the confidentiality and integrity of systems and data.

Hypothesis and Outline

This paper presents an introductory analysis of aggregated log metrics for CSC computing services for the time period from the beginning of year 2008 until the end of the first quarter of year 2011. We show access history and service up-time processed with tools as Nagios and Splunk in conjunction of a set of cases of vulnerabilities, intrusions and faults. The paper is also comparing and suggesting improved best practices to be shared between sites.

The discussion at the end of this paper will be based on following hypothesis, based on our professional experience:

1. Brute-force attacks to guess user passwords originate from a limited set of source addresses and are directed against a limited set of generic user names, directed attacks against actual user names are not common
2. Security incidents cause a considerable amount of downtime
3. Security risks related to intrusions, unauthorized access and system compromises can be mitigated with better access controls without degrading usability and user experience
4. An adequate intrusion detection system and an incident response plan will diminish downtime caused by intrusions
5. With adequate user management and operational security controls brute force attacks do not constitute major risks, it is merely noise
6. Implementing access controls in an optimal way requires sharing of best practices between peer sites

For simplicity this paper will limit the study to risks related to ssh based access, as ssh access is the most widely used mechanism to access academic computing services. There are also alternative access methods, for example grid jobs and Web based user access, as implemented by the Scientist User Interface (SUI) at CSC. Another limitation of this paper is, that we do not disclose such detailed information which might endanger the security of CSC or the privacy of its users, all information presented in this paper is considered safe to disclose.

2. Materials and methods

2.1 CSC Computing Environment

The CSC computing environment [6] consist currently (in spring 2011) of

- Massively parallel processing super computer Louhi (Cray XT4/XT5)
  - Front end servers for login and interactive work
  - High speed SeaStar interconnect
  - Shared 70TB lustre file system for binaries and scratch space
  - Interfaces for DEISA and PRACE
- Super cluster Murska (HP CP4000 BL ProLiant supercluster)
  - Front end servers for login and interactive work
  - 512 computing nodes, 2048 computing cores, 4608GB memory
  - High speed Infiniband interconnect
  - Shared 110TB lustre file system for binaries and scratch space
  - Interfaces for EGEE, MGRID and SUI
- Super cluster Vuori (HP CP4000 BL Proliant supercluster)
  - Front end servers for login and interactive work
  - 240 computing nodes, 2880 computing nodes, 5632GB memory
  - 32 dedicated computing nodes, 384 computing cores, 1024GB memory
  - 8 GPGPU nodes with 96 Intel cores, 16 Tesla 20x0 cards
  - High speed Infiniband interconnect
  - Shared 45TB lustre file system for binaries and scratch space
  - Interfaces for FGI and SUI
- Application Servers Hippu (HP ProLiant DL785 G5 server pair)
  - 2 Large memory “fat” nodes for interactive workload each with 32 computing cores and 512GB memory
  - Local FC scratch disk
  - Interface for SUI
- Other hosted computing systems

The basic network topology if CSC computing services is shown in Figure1.

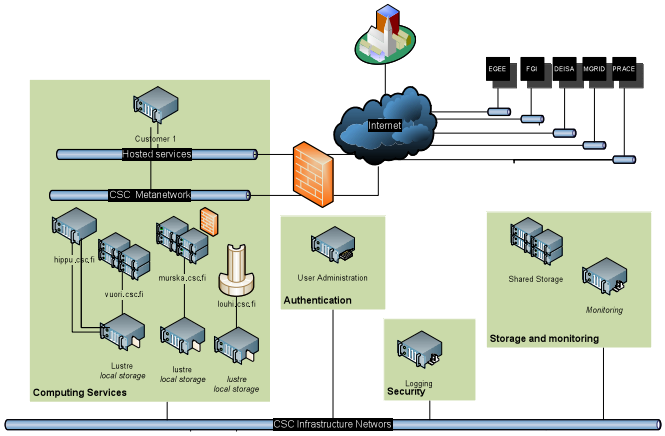

Figure 1. CSC Computing Services year 2011

2.2. Log Analysis

All successful and failed login attempts in CSC computing environment are logged both locally and centrally using the syslog protocol [7], the logs are stored on secure loghosts.

The logs can be analyzed by CSC Security team with scripts or currently also with tools as Splunk [8] . Plans have been made to implement a more automatic and intelligent distributed intrusion detection and prevention system.

Technically and operationally logging is trigged by the ssh daemon of the login nodes of that login node a user is accessing. A login attempt generates the following type of entries on the access log:

Nov 19 08:00:08 louhi-login3.csc.fi sshd[30676]: Failed password for invalid user xxxx from 172.16.0.1 port 51748 ssh2
Nov 19 08:00:09 louhi-login3.csc.fi sshd[30678]: Invalid user xxxx from 172.16.0.1
Nov 19 12:15:01 louhi-login3.csc.fi sshd[8435]: Accepted publickey for yyyy from 172.16.1.1 port 56954 ssh2
Jan 1 21:12:46 murska-login2.csc.fi sshd[10257]: Invalid user zzzz from 172.16.2.1

Table 2: Example of sshd log entries (user names and addresses are obscured for privacy and security reasons.)

The format and entries of the logs differs depending on sshd platform and software versions.

2.3 Availability Metrics

CSC monitors system availability both service based and centrally. The main tool for availability monitoring is Nagios [9]. During 2010, CSC begun to create a more structured view of its services and on the dependencies between them. CSC had already a long history of system monitoring with Nagios, but by introducing an add on to Nagios, the Nagios Business Process Intelligence (BPI) tool, could host availability status be grouped in aggregated views according to CSC services, for example computing services, university and AMK library systems, National Archives service.

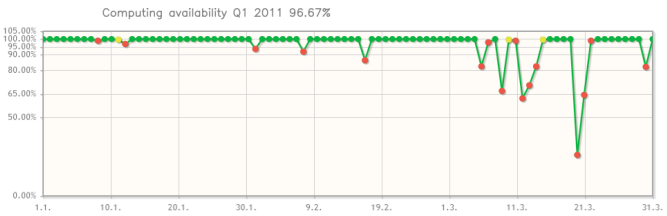

Figure 2. Example of availability monitoring of CSC computing services.

The aggregated availability monitoring with Nagios BPI is still in a pilot phase and the probes measuring availability still lack some depth to truly simulate user experience of availability. The current probes uses ping, http get command, ssh connect and disk space measurements to measure availability.

In some cases, as shown in Figure 2, measurement issues can also give false alarms on availability, despite the fact that the services are up and running.

The improved monitoring tools based on Nagios BPI have a current view, but can also be used to easily generate availability metrics for management and customers. CSC has created new roles and practices for monitoring and to react faster on service breaks. A cross –organizational Operations Manager and a Technical Operator, will have the joint primarily responsibility to ensure service availability and alert system experts in case of breaks. In most cases specific expertise is required to solve issues, a network engineer is not supposed to solve data center power issues and a data center engineer is not supposed to do core network router configuration. Weekly shifts for CSC Operations Manager and Technical Operator duties are organized by roster.

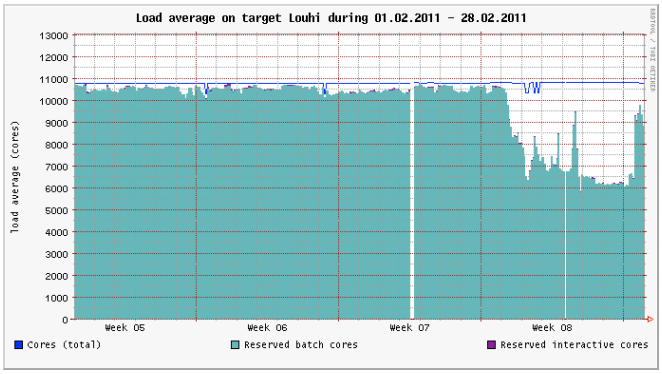

Figure 3. Example of monitoring load average of Louhi.

The service based monitoring for computing services is at CSC implemented with Cacti graphing [10].

2.4 Case Studies

In times of peace and quiet, security can sometimes be experienced as an unnecessary cost and burden which is implemented mainly because of external security requirements. During security incidents people often think, why one didn't prepare in advance to the threats, which often are well known and predictable.

With the case studies from two major system intrusions we want to improve our understanding to avoid and handle security incidents.

3. Results

3.1 Patterns detected from log analysis

Sanitizing scripts were first used to extract essential data from centralized ssh access logs, after that the logs were for convenience analyzed with Splunk.

Our first task was to check out, how many brute-force attempts were made from individual addresses.

The average for ssh password guessing attempts on CSC computing environment during 2008-2011/Q1 for the top 200 attacking addresses was 95 733, See Figure 4.

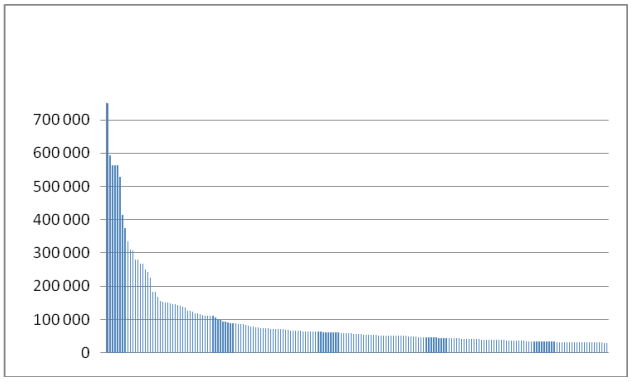

Figure 4: Number of password guessing attempts against CSC Computing Servers 2008-2011/Q1 with brute force ssh attacks per source IP address /Top 200

The geographic distribution of the top-10 attacking source addresses, were according to CSC logs as described in Table 3, but the validity and the reliability of the information can be questioned as the source addresses can contain any number of hijacked hosts (botnets).

Origin	Attacks	Share
China	6,469,719	23.1%
United States	3,337,965	11.9%
France	2,349,973	8.4%
Korea Republic of	1,521,382	5.4%
Russian Federation	978,433	3.5%
United Kingdom	821,596	2.9%
Japan	799,124	2.8%
Turkey	669,072	2.4%
Greece	627,732	2.2%
Italy	620,932	2.2%

Table 3: Geographical origin of the attackers addresses.

The hour distribution for the attacks within a day appears to be flat, each hour of the day scored between 4.451% (7 A.M) and 5.246% (4 A.M).

When studying the destination of the attacks, the following differences were found between the three most massive brute force attacks during the examined period.

Host	Attack#1	Attack#2	Attack#3
Louhi	99.992%	95.116%	99.989%
Hippu	0.003%	0 %	0.003%
Vuori	0.002%	4.88%	0.003%
Murska	0.002%	0.003%	0.004%
Accounts tried	15008	11173	11169
Length of attack:	5 days	12 hours	1 week

Table 4: Attack vectors/ Brute force attempts

The attempts on Louhi are high because on the other CSC computing hosts Denyhosts [11] will deny network access for a configurable period after a predefined amount of unsuccessful attempts. Denyhosts has not (yet) been installed on Louhi to ensure customer convenience and to avoid client lock outs due mistyped passwords in legitimate login attempts.

The massive brute-force attacks used a international set of presumed usernames, originating perhaps from an generic IT environment instead of a typical computing center.

User	#	User	#	User	#
Admin	9365	andrew	2086	Karl	1577

Test	6652	angela	2069	benny	1461
Oracle	3855	office	1919	web	1420
User	3452	amy	1915	www	1402
apache	3130	cvs	1862	kamal	1368
info	3056	student	1786		
guest	2802	fred	1766		
testing	2489	dummy	1726		
support	2411	webmaster	1717		
chris	2154	upload	1664		

Table 5. Top25 user names used in brute a force attack

3.2 Case Studies

3.2.1 Intrusion of Louhi 2009

To complement aggregated metrics we will present two cases of system intrusions at CSC.

On the Friday morning of April 17th, 2009, The CSC Security Manager received information about attacking hosts, but that information did not escalate to incident handling because unsuccessful brute force attacks are very common. After office hours, at 5:40 P.M. EET CSC received vendor information about a local exploit vulnerability related to the later to be announced udev vulnerabilities (CVE-2009-1185 and CVE-2009-11856 [12]).

On Sunday the 19th, administrators of CSC computing services, received a request from partners to investigate anomalies. On Monday morning of the 20th, an experienced CSC senior system administrator identified abnormal entries in the logs of Louhi. Various susceptible commands had been run with root privileges. Louhi had been compromised.

After the incident was identified, CSC Security Manager formed an incident handling group according to CSC guidelines practices. Affected and possibly affected customers and other external systems were taken off-line and system integrity of several CSC systems was analyzed. The incident handling group generated new passwords for all users of CSC Computing services, staff, partners and external users

The incident handling group informed CSC Customers and partners about the incident.

All existing staff and user ssh keys were removed, but a suggested blacklisting of possibly compromised keys could not be deployed due incompatibility issues. After the incident, storing unencrypted ssh private keys was banned.

After ensuring integrity of systems and user accounts, with fresh installations of some front end systems, all CSC external systems were made online and available for customer use on the afternoon of Thursday, April 23rd. CSC lost system uptime for four days for CSC external computing services, also a considerably amount of extra work was made because of the incident.

External access to Louhi was closed until the afternoon of April 24th, when all services were back online again.

The modus operandi of the attackers, who operated internationally and attacked many similar sites, seemed to be the following:

1. Gain shell access by a compromised grid multi site account
2. Gain root access by taking advantage of a fresh unpatched vulnerability
3. (Try to) install a keylogger and/or a rootkit to collect passwords
4. Scan user home directories for unencrypted ssh private keys (which should normally never be saved server-side, at least not unencrypted)
5. For found unencrypted keys, check history files where the user has previosly gone
6. Attack next host with found ssh keys and repeat from 1.

The vulnerability used during this attack, CVE-2009-1185, was a bug in previous versions of linux kernel udev device driver, allowing local users to gain root privileges. Local privilege escalation vulnerabilities are not uncommon. Most linux distributions and vendors issued warning about the vulnerability (For example Debian Security Advisory DSA-1772-1, Red Hat Security Advisory RHSA-2009:0427, the exploit was also available on the internet already on April the 17th.

3.2.1 Intrusion of murska 2010

Another system compromise was detected at CSC on Wednesday, September 29, 2010 when a password stealing program (a rogue sshd binary) was found in a login node of Murska.

Similar steps to react on the incident were taken as in during the Louhi incident which occurred the previous year.

By analyzing the centralized logs (local logs were again partly deleted) it was found out that the intrusion was made by using a legitimate user account from China. The intruder had gained root privileges by utilizing an unpathched vulnerability CVE 2010-3081 [12] and thereafter installed a rogue keylogger, to collect passwords from users logging in.

The intrusion also affected some other sites, which as CSC performed an enforced password change operations.

After the change of passwords and reinstallation of operating systems of the front ends to ensure system integrity, murska was opened back for customer use on the afternoon of October 6th.

3.3 Availability Metrics

According to CSC Computing Services availability metrics system availability has for the servers has been as described in table 6. and in table 7.

As we can see from Table 6, systems seem to mature and improve availability over time after the year of first deployment. Both vendors and administrators can learn from errors, bugs can be patched and faulty components can be replaced. There can still be big differences between systems from different vendors. Some systems never reach reasonable high availability. It must also be remembered, that true supercomputing involves considerable risks when reaching for maximum performance.

Host	2008	2009	2010	2011	Availa-bility
All breaks (scheduled and unscheduled)					
Louhi	1492h 51min	551h 22 min	469h 6min	29h 37min	91.29%
Murska	834h 30min	295h	237h	24h 35min	95.23%
Vuori	N/A	N/A	N/A	0h	100.0%
Hippu	N/A	N/A	N/A	26 min	99.98%
Unscheduled breaks (scheduled breaks removed)					
Louhi	341h 58min	350h 7min	348h	29h 37min	96.14 %
Murska	778h 30min	216h	237h	24h 35min	95.68 %
Vuori	N/A	N/A	N/A	0 h	100.0%
Hippu	N/A	N/A	N/A	26 min	99.98%

Table 6: System availability 2008-2011 (1.1.-30.4.2011).

As shown in Table 7. the amount of service breaks due security reasons (incidents, vulnerability patching has been done during scheduled breaks) have been quite high, although infrastructure issues (Power, HVAC, NFS) has also constituted the major cause for the breaks in addition to the main culprits for breaks, system specific hardware and software issues.

The hardware category contains also scratch and wok disks. The software category contains operating systems, cluster software and firmware. The group for It Infrastructure include NFS serves and external customer networks.

The big differences on the group IT Infra are because of issues on NFS Servers normally leads to a total service break on Louhi but on Murska the batch jobs can continue to run, users just can't log in.

The reason for differences on the group Data Center infra stems from the fact that Louhi and Murska are placed in separate Data Centre's. The unknown breaks are typically hardware or software issues.

	Louhi	Murska	Vuori	Hippu
Security	6.82%	12.98 %	N/A	0 %
Hardware	47.30%	61.35 %	N/A	0 %
Software	18.11%	11.48%	N/A	100 %
IT infra (NFS & al.)	7.75%	0.14 %	N/A	0 %
Data Center infra	0.35 %	14.01 %	N/A	0 %
Other	19.67%	0.04%	N/A	0%
TOTAL	**100 %**	**100 %**	**N/A**	**100 %**

Table 7. Reason for unscheduled breaks 2008-2011 (1.1.-30.4.2011)

Compared over time we can also follow an interesting availability competition between Louhi and Murska.

	2008	2009	2010	2011-
Louhi	83.00%	93.71%	94.64%	98.97 %
Murska	90.50%	96.63 %	97.29%	99.15%

Table 8. Availability including all breaks 2008-2011 (1.1.-30.4.2011)

4. Discussion

4.1 Analysis of the combined results

After looking at the data, intrusion metrics and cases combined with the availability metrics it is now time to ask how did the brute force attacks and the security affect availability of CSC Computing services. We will try to answer the hypothesis we presented in the beginning of this paper.

Brute-force attempts to guess user passwords originate from a limited set of source addresses and are directed against a limited set of generic user names, directed attacks against actual user names are not common

True, but it might be that many advanced attacks against actual usernames remain undetected. **Security incidents cause a considerable amount of downtime**

True. A substantial amount of the downtime was really caused by security incidents. Improving security

can advance efficiency and add value to customers. **Security risks related to intrusions, unauthorized access and system compromises can be mitigated much with better access controls without degrading usability and user experience.**

Partly true. The introduction of Denyhosts did filter out most of the brute force attacks but such tools cannot detect and prevent abuse of hijacked user accounts. Tighter access control, by for example requiring a cryptographic key or certificate to login or limit allowed networks might degrade at least somewhat usability. **An adequate intrusion detection system and an incident response plan will diminish downtime caused by intrusions.**

Most probably true. A faster reaction time and smoother incident response would at least make disaster recovery faster.

With adequate user management and operational security controls brute force attacks do not constitute major risks, it is merely noise

True. CSC has not experienced any break-ins to its computing systems due brute force ssh attacks

Implementing access controls in an optimal way requires sharing of best practices between peer sites.

True. Numerous open source and product based tools are available to improve access controls, but planning deployment for computing services requires investment in manpower, coordination, learning from results, improving configurations and setting up processes for adequate monitoring and reaction.

When the generic outline and main findings of this paper began to be ready, we interviewed a fellow senior System Administrator, Mr. **Esko Keränen** on the topic. Mr. Keränen is known for his long experience and very advanced skills on system administration of different generations or supercomputers

According to Mr. Keränen, the important factor to detect system intrusions is an understanding of normal system behavior and systematic monitoring, despite the massive amount of log data, of anomalies. Typical signs of intrusions are:

- Process and files with atypical user rights
- Process and files with atypical group rights
- Breaks in flows of logging

Automatic monitoring might improve detecting some anomalies, but can be difficult to detect directed advanced attacks with automatic tools.

In addition to these comments, most system administrators and security managers share the concern, about the common challenge for developers, vendors and systems administrators to race in patching the systems before security vulnerability is piggybacked by a abused user account. The known accounts we trust are typically a bigger risk than the unknown intruders.

Although script kiddies performing brute force attacks have not been the major threat for Computing Services, they might succeed to intrude if basic account security fails or by denying service for legitimate users by DDOS attacks (Distributed Denial of Services, by attacking a host from several source addresses), says Mr. **Tommi Tervo**, another experienced system administrator at CSC.

Another growing concerns adding risks and complexity to computing services is grid computing, where users are managed on peer sites and complex protocols are used. Grid computing requires a lot of trust between sites but on the other hand grid computing is due its flexibility also very strongly preferred by financiers and customer. The grid computing sites has also begun to share best security practices and the basic authentication method in grid computing is based on certificates instead of traditional passwords.

The finding of this paper also gives CSC management and system administrator's new facts for updating risks assessments and re-evaluate current access controls. Although it is still too early to definitely suggest new assessments and controls, the result shows that access controls could with intelligent tool and monitoring be made more strict without degrading user experience. Finally introducing required strong authentication methods, such as ssh keys, certificates or tokens, would also eliminate many of the risks stemming from brute force attacks.

4.2 Conclusions

Based on the analysis and discussions of CSC logs and security incidents we see a significant dependance of ensuring availability and protecting system integrity with access controls and other security tools. The tools are not enough, also skills and best information security practices are required to deploy tools and monitor security events and availability.

Automatic monitoring might improve detecting some anomalies, but it is much more challenging to detect directed advanced attacks.

To assure availability of computing services requires:

- Adequate operational security to ensure proper user authentication mechanism
- A method to ensure system integrity to detects intrusion and installed malware, such as rootkits and password loggers
- Smooth mechanism for incident response to minimize downtime
- Sharing of skills and information between peer sites to maintain capability to prevent and mitigate risks

4.3 Suggestions for further research

We authors are aware of some limitations of this paper, metrics and data analysis can in future research be more profound. Results of a single case study cannot be mechanically generalized to apply for all computing centers, the circumstances differs depending on stakeholders, service portfolio and size, to mention a few factors.

The consciously selected wide and explorative approach has limited the depth and perhaps also some of the reliability of the findings. A new in-depth analysis should be made on current log data to get a deeper understanding of the operational security.

4.4 Suggestions for sharing best practices

CSC utilizes many security tools to protect its computing services.

Except DenyHosts (many other computing centers use more feature rich and complex iptables firewall based Fail2Ban[13] instead) CSC is also ensuring account integrity and compliance by checking that user passwords meet minimum requirements by testing them with John the Ripper password cracker tool [14], which can also be run in parallel mode for greater speed.

Integrity of system binaries and configuration files are System integrity are tested on some host with AIDE [15].

To deploy security tools successfully requires skills and experience to not cause more harm than benefits. Configuring optimal threshold of Denyhosts, for example, must be done to avoid excessive false positives, locking out legitimate users mistyping their passwords. Professional system administrations also includes proactive monitoring, in case of Denyhosts, to see what kind of addresses has been blocked and to manually override settings when false positives has been detected.
As suggested previously [16], computing centers should be more active to systematically share experiences, skills, tools and best practices on both technical information security and security management – the same way as application specialist and systems administrators do.
An exchange and visiting program for junior and senior administrators should be initiated to develop best system administration and information security practices. Also service and security management should be more involved in sharing best practices, management issues can sometimes be the bottleneck for deploying better and more efficient security.
Information and computer security is not only about tools and technology; it also requires human networking, direct contacts between sites, and organized mailing lists for incident handling, vulnerability alerts and proactive security. Face to face meetings between sites are also an essential investment to improve site security. Proactive security cannot be created without allocating proper resources for it terms of tools, skills, peer networks and working hours.
Security is not only about tools and skills, a very important component of security is embedded already in hardware and software design, systems should be secure and reliable by design. Design flaws can be difficult or impossible to patch afterwards. This was also the reason for CSC to clearly highlight security and availability requirements in addition to capacity, features and usability requirements in its current acquisition of the next generation supercomputer and high performance cluster. Vendors are required to enter a security agreement with CSC to ensure proper security controls, vulnerability patching and incident response.

Security sharing tools and tips as presented on the web page of Mr. Kurt Carlson of ARSC [17] is a really valuable contribution to improve security of computing centers, more similar best practices should be made available for computing centers worldwide, both platform specific and as more generic advise.

Acknowledgments

The authors would like to thank colleagues, peers and vendor staff for good cooperation and for sharing and improving security together. Special thanks to members of Cray Extreme Group and colleagues on CSC Computing Services and Storage Services Groups, Jonas Dahlbom, Peter Jenkins, Pietari Hyvärinen, Esko Keränen, Samuli Saarinen, Dan Still, Tommi Tervo and Hannu Kivimäki.

About the Authors

Urpo Kaila is Head of Security at CSC, he is involved in a wide area of activities to promote and develop best practices of information security at National Research and Education networks and Computing services. Urpo and the other authors can be reached at CSC - IT Center for Science Ltd., P.O. Box 405, 02101 Espoo, Finland. Urpo's email address is urpo.kaila@csc.fi. Marco Passerini, marco.passerini@csc.fi, is a system administrator at CSC, Marco has specialized in information security. Joni Virtanen, joni.virtanen@csc.fi, is a senior system administrator and deputy group manager at CSC, Joni is a long-time CUG member and member the XTreme group.

When writing this article the authors divided their work in such way that Marco was responsible for log analysis with Splunk, Joni created and computed availability metrics and Urpo had the main responsibility for writing and editing the article.

References

[1] CSC Web site. http://www.csc.fi/english. Retrieved on 2011-05-07

[2]Cerf, V., "Thoughts on the National Research and Education Network", RFC 1167, July 1990.

[3] Information Security Policy of CSC ("Tietoturvapolitiikka"). 2004. Internal Document. CSC - IT Center for Science Ltd

[4] Ohje tietoturvallisuudesta valtionhallinnossa annetun asetuksen täytäntöönpanosta, VAHTI 2/2010. Valtiovarainministeriö. 2010. ISBN: 978-952-251-124-9.

[5]Tipton, Harold F. and Henry, Kevin. Official (ISC)2 Guide to the CISSP CBK. (Isc)2 Press Series. 2006. USA.

[6] Saarinen, Samuli. Business Continuity plan for CSC Computing Services. 2011. Internal Document. CSC - IT Center for Science Ltd.

[7] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.

[8] Splunk web site, http://www.splunk.com/. Retrieved on 2011-05-07.

[9] Nagios web site. http://www.nagios.org. Retrieved on 2011-05-07.

[10] Cacti web site, http://www.cacti.net/. Retrieved on 2011-05-07.

[11] Denyhosts web site. Retrieved on 2011-05-07. http://www.denyhosts.sourcforge.net/.

[12] Common Vulnerabilities and Exposures web site, http://cve.mitre.org/. Retrieved on 2011-05-07.

[13] Fail2ban web site. Retrieved on 2011-05-07. http://www.fail2ban.org/wiki/index.php/Main_Page

[14] John the Ripper password cracker web site. Retrieved on 2011-05-07. http://www.openwall.com/john/.

[15] http://aide.sourceforge.net/.Advanced Intrusion Detection Environment web site. Retrieved on 2011-05-07.

[16] Kaila, U., "Best Practices for Security Management in Supercomputing". CUG 2008 proceedings.

[17] Tools, Tips and Tricks for Managing Cray XT Systems Kurt Carlson. Arctic Region Supercomputing Center (ARSC). http://www.arsc.edu/~kcarlson/CUG2010/Carlson.html Retrieved on 2011-05-07