

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

**Urpo Kaila, Marco Passerini  
and Joni Virtanen**

CSC – Tieteen tietotekniikan keskus Oy  
CSC – IT Center for Science Ltd.

# Outline



- Introduction
  - About CSC
  - About Security
  - The objectives of the paper
- CSC environment
- CSC procedures and metrics
- Results
  - Data and discussion
- Q & A

# CSC at a glance




- Founded in 1971 as a technical support unit for Univac 1108
- Connected Finland to the Internet in 1988
- Reorganized as a company, CSC – Scientific Computing Ltd. in 1993
- All shares to the Ministry of Education and Culture of Finland in 1997
- Operates on a non-profit principle
- Facilities in Espoo, close to Otaniemi campus (of 15,000 students and 16,000 technology professionals) and Kajaani
- Staff 200
- Turnover 2009 21,9 million euros



# Mission




- CSC, as part of the Finnish national research structure, develops and offers high quality information technology services



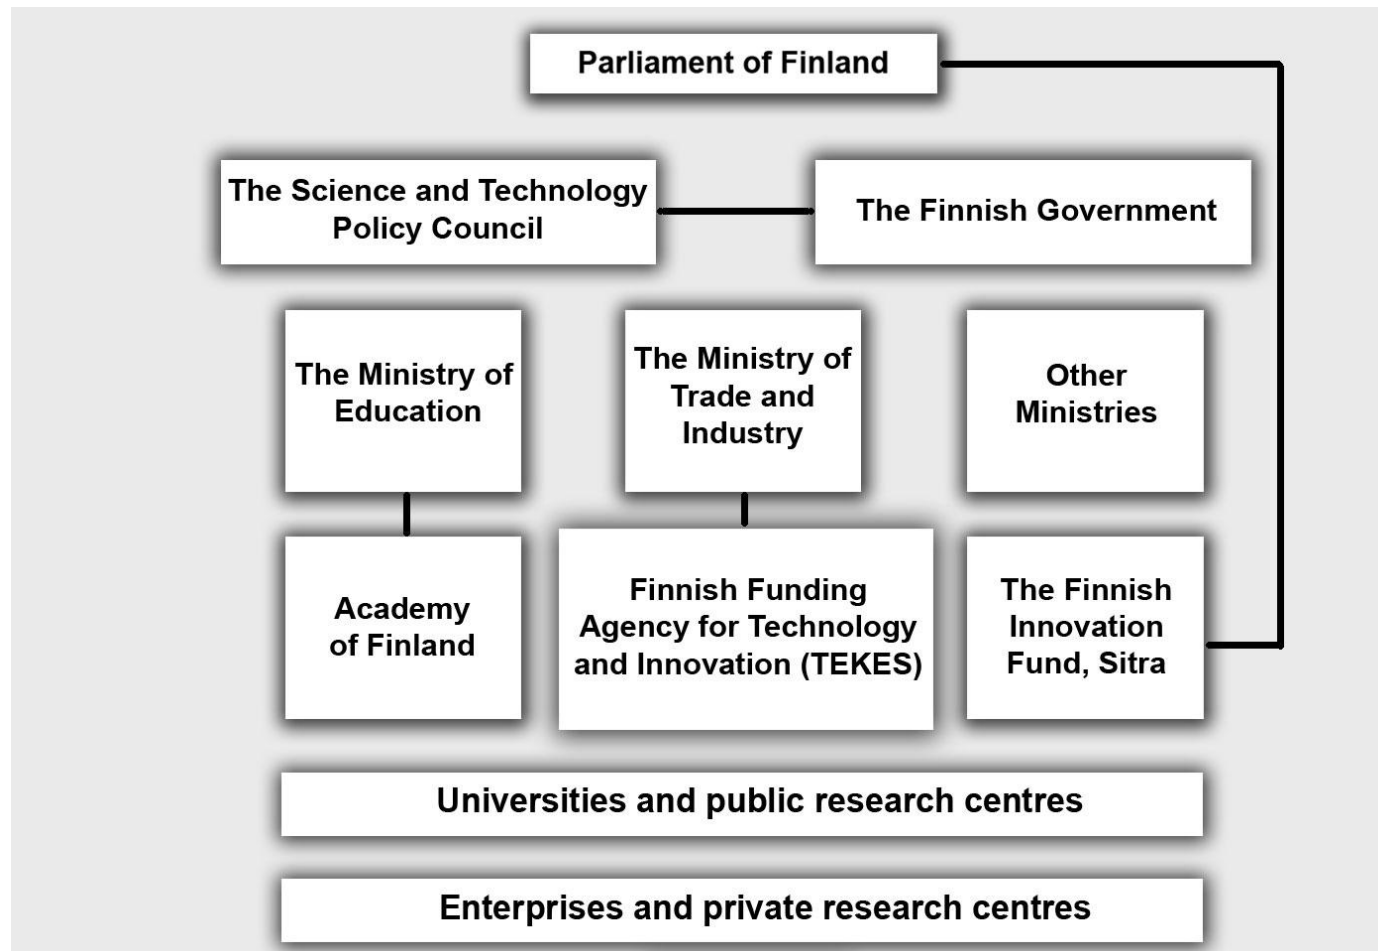
# Operational goals



- Improves conditions for research and product development
- Provides national level, centralized services in fields that would be impracticable to support at university level
- Promotes collaboration
- Provides internationally competitive supercomputing and data communication services
- Serves as a pioneer and information provider



# CSC supports the national research structure



# Important national and international actor



- Offers high-level expert services for the usage of softwares, databases and methods
- Participates actively on European high performance computing development projects



# Customers



- 3000 researchers use CSC's computing capacity
- Funet connects about 80 organizations to the global research networking infrastructure
  - universities
  - polytechnics
  - 35 industrial clients and research institutions
  - Total of 350 000 end users





# CSC's services




Funet Services



Computing Services



Application Services



Data Services for Science and Culture



Information Management Services



# Funet Backbone Network



- Funet backbone provides reliable and high-capacity connections for all Funet member organizations in Finland. Funet is connected to international academic networks via NORDUnet.
- Funet backbone supports advanced services like IPv6 and IP multicast. Link speeds range up to 10 Gbps.
- Since spring 2009, light paths, dedicated high-capacity links for special applications and users, have been available in many locations.

# International collaboration

- Computing centers
- International research network organizations:
  - NORDUnet, TERENA, Internet2, Dante (Géant2)
- International science network organizations:
  - European Molecular Biology Network (EMBnet), EMBRACE
- Nordic and European HPC projects and GRID-organizations:
  - Nordic Data Grid Facility, NorduGrid, DEISA2, EGEE-III, NEG, ESO, Sirene, PRACE, EGI
- CSC chairing: TERENA, E-IRG, EGI, NORDUnet, PRACE (vice-chair)





# ABOUT SECURITY


# What is information security all about?



- Information security is about protecting assets (systems, data services and reputation) against risks.
- Assets can be protected to prevail their
  - Confidentiality
    - To prevent intentional or unintentional disclosure
  - Integrity
    - To prevent unauthorized modification and protect consistency
  - Availability
    - To protect reliability and timely access
- **Information Security is**
  - a building block of quality
  - a management responsibility
  - implemented by security controls
  - the responsibility of each and everyone of all staff

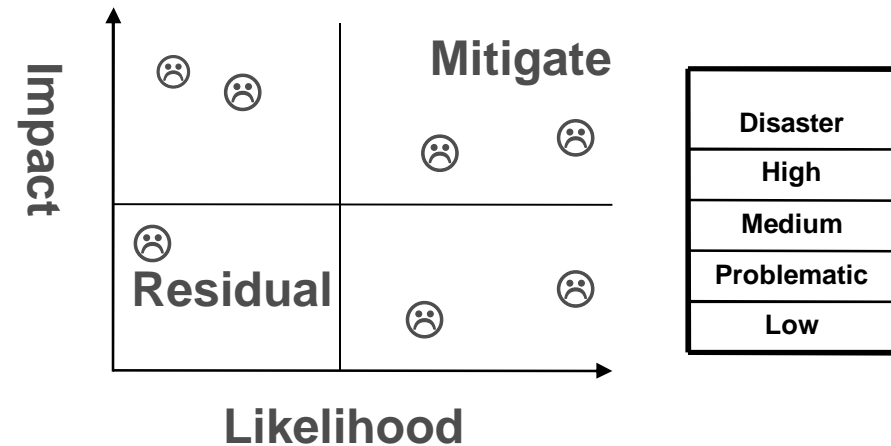


# What kind of risks do we have?



- Stealing of account because of account sharing or weak passwords
- Misuse or malicious use
- Infrastructure problems (fire, power supply, cooling, flood, malfunction, ...)
- Hard to meet regulatory requirements on privacy
- Loss of data because lack of skills
- Break in through scans and queries because problems with change management
- Privacy issues due phishing
- Services down because of DDOS
- Theft
- Vulnerability exploit due insecure configurations
- Unnoticed backdoors due lack of time for proper system administration

**Internal – Intentional\***  
**Internal - Accidental**  
**External - Intentional**  
**External - Accidental**



# Compliance and Best Practices

CSC owner, partners and peers require to comply with

- Privacy and security laws
  - Act on the Openness on Government Activities
  - Act on the Protection of Privacy in Electronic Communication
  - Criminal Act
  - Decree on Information Security
- Government and industry regulation
  - The Government Information Security Level Manual (GISLM): COMPLIANCE WITH RAISED LEVEL
  - (FICORA regulation)
- Contracts
  - The MinEdu contract
  - Several customer and peer contracts
- Best practices




Several interrelated best practices for IS and ISM

- ISO27001/27002 (ISF SOGP)
- ISM3
- ITIL
- EFQM/ EA
- COBIT

# The objectives of our paper



- Analysis of aggregated log metrics for CSC computing 2008-
- Access history and up-times with Nagios and Splunk
- Cases of intrusions
- Suggest best practices to be shared between sites
- We do disclose a lot, but not everything, not security through obscurity (or through babbling :)





# Hypothesis



1. Brute-force attacks comes from a limited set of sources and are directed against a limited set of user names, directed attacks against actual user names not common
2. Security incidents cause a considerable amount of downtime
3. Unauthorized access can be mitigated with better access controls without degrading usability
4. Intrusion detection system will diminish downtime
5. With adequate user management brute force attacks do not constitute major risks
6. Implementing optimal access controls requires sharing of best practices between peers



# CSC ENVIRONMENT

# CSC's CrayXT4/XT5



## CRAY XT4/XT5 alias **Louhi**

- 2356 AMD Quad Opteron  
2,3 GHz CPUs
- 10864 cores
- Memory ~ 11,7 TB
- Theoretical computing  
power 100 teraflop/s



# CSC's CrayXT4/XT5




- Louhi has been upgraded with two Cray XT5 – cabinets (alias Loviatar)
- 360 AMD Opteron Quad Core 2,3GHz CPUs (was updated with AMD Shanghai Quad Core processors early 2009)
- 1440 cores
- Theoretical computing power 13,24 teraflop/s
- The new Cray XT5 -cabinets are part of PRACE-project (Partnership for Advanced Computing in Europe). PRACE selected a broad coverage of promising architectures for petaflop/s-class systems to be deployed in 2009/2010. PRACE is a project funded in part by the EU's 7th Framework Programme.



# Other CSC Computing resources (1/2)




- 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 inter-connect
  - Shared 110TB lustre file system for binaries and scratch space
  - Interfaces for EGEE, MGRID and SUI

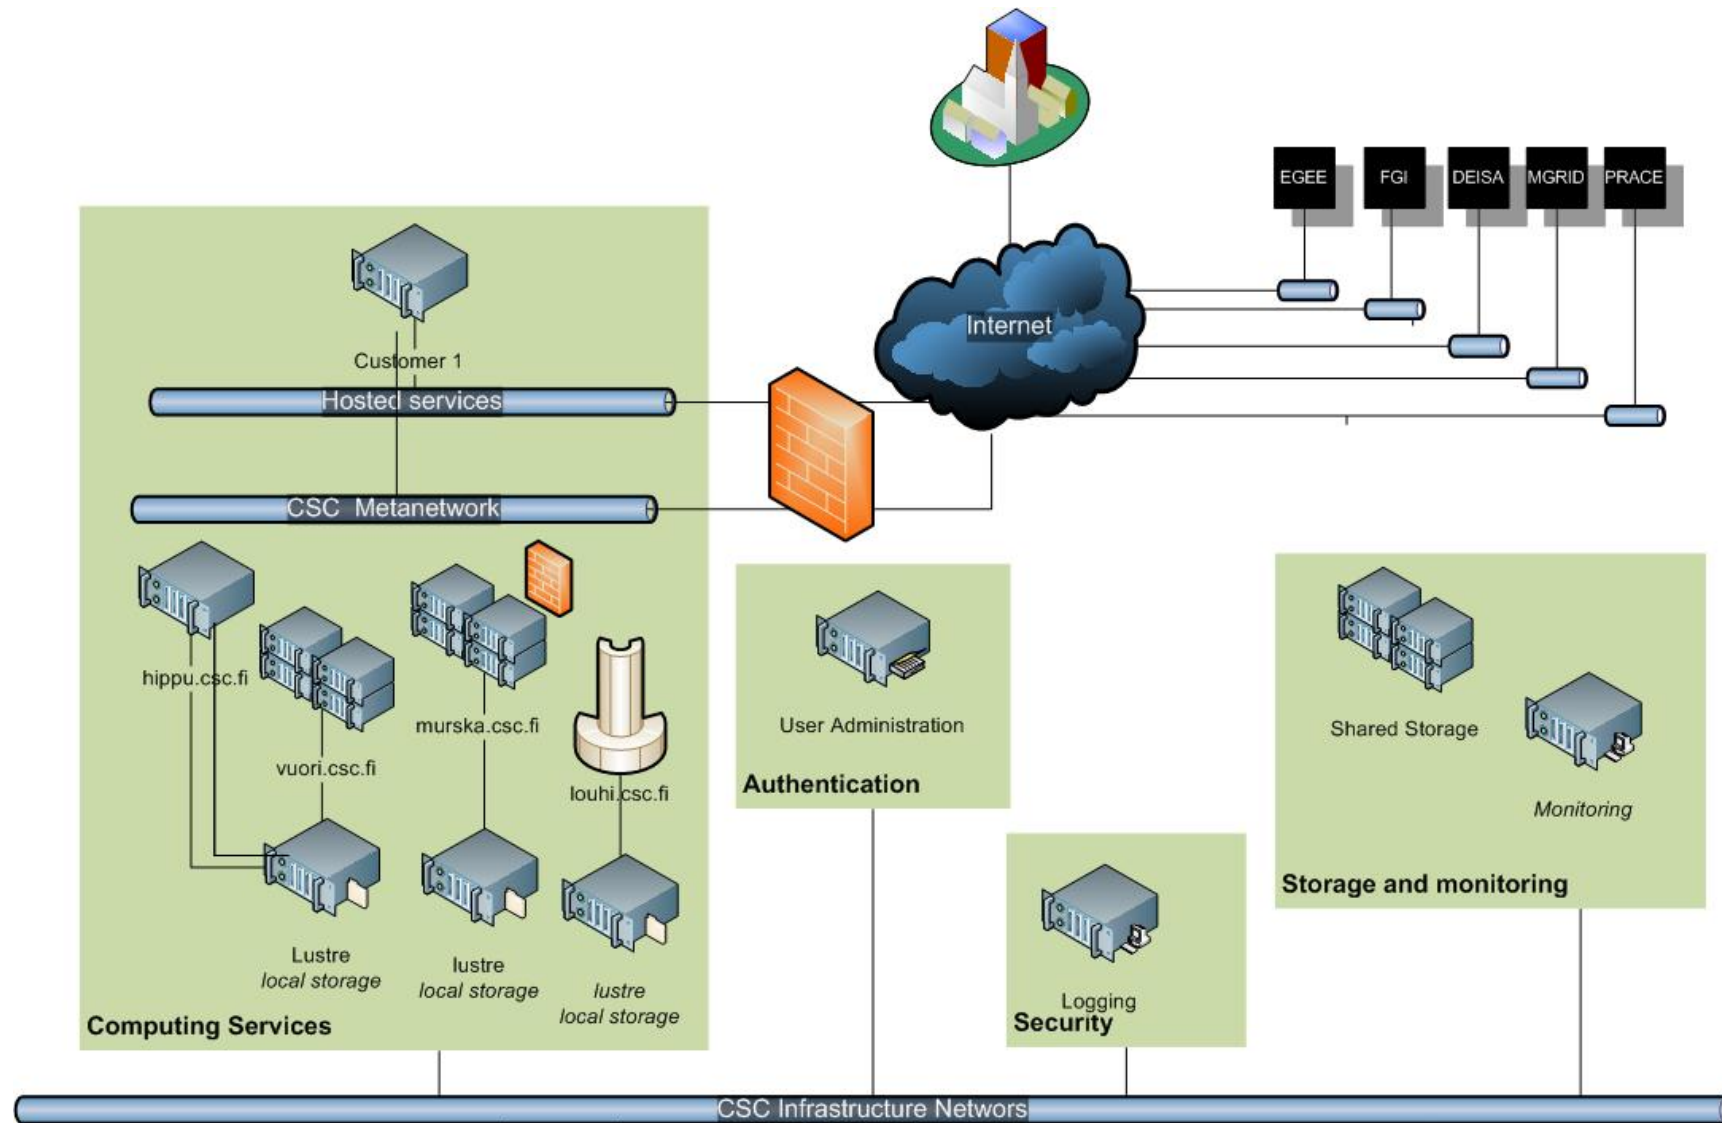


# Other CSC Computing resources (2/2)

- 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 CSC Computing Environment



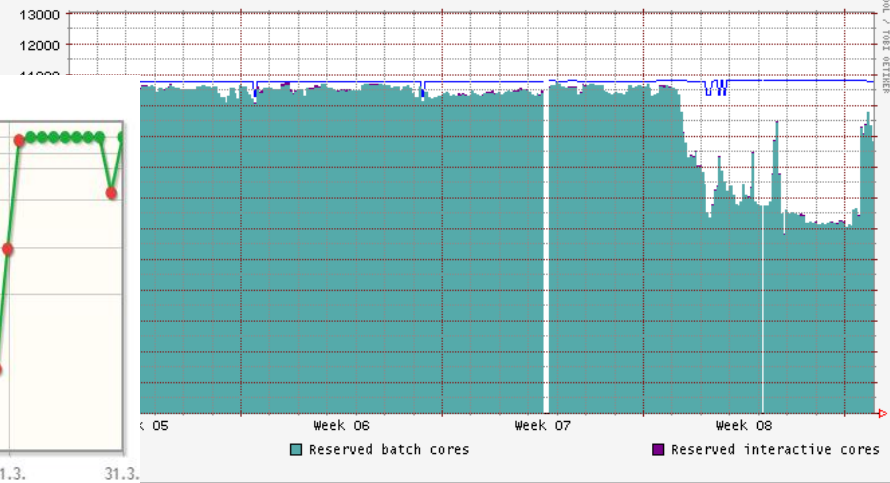
# CSC PROCEDURES AND METRICS




Computing availability Q1 2011 96.67%



Load average on target Louhi during 01.02.2011 - 28.02.2011





**CSC Critical Services**  
under Operations Management monitoring

18.05.2011 17:34



Availability graphs

Ok	Services BPI	0 problem(s)			
Ok	▶ Hippi	0 problem(s)	Hippi application servers (CORE)	100.00%	100.00%
Ok	▶ KDK	0 problem(s)	National digital library (DATU)	100.00%	100.00%
Ok	▶ Libraries	0 problem(s)	University and AMK library systems (DATU)	98.35%	98.35%
Ok	▶ License	0 problem(s)	License services (SE)	99.98%	99.99%
Ok	▶ Louhi	0 problem(s)	Louhi supercomputer (CORE)	97.55%	97.55%
Ok	▶ Murska	0 problem(s)	Murska supercluster (CORE)	99.98%	99.98%
Ok	▶ SUI	0 problem(s)	Scientist's User Interface (SE)	99.76%	99.76%
Ok	▶ Vuori	0 problem(s)	Vuori supercluster (CORE)	100.00%	100.00%
Ok	▶ VAPA	0 problem(s)	National archives VAPA service (DATU)	99.94%	99.81%
State	Service	# of problems	Description (responsible group)	05/2011	2011 total


Autorefresh is on In case of production or quality issues, contact Operation [Waiting for csc.fi...](#)






# RESULTS

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



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	15008	11173	11169
Length of attack:	5 days	12 hours	1 week

**Brute force attacks per origin country**



**Attack vectors/ Brute force attempts**

# Case intrusion of Louhi 2009



- April 17th, udev vulnerabilities CVE-2009-1185 and CVE-2009-11856
- April 19th, warning from partners
- April 20th, an admin identified abnormal entries in the logs of Louhi, Louhi compromised. Incident Handling Group formed, Systems were off-line, integrity checks, new passwords and ssh keys for all users, Customer info, Ban on storing unencrypted ssh private keys

## The modus operandi of the attackers

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 previously gone
6. Attack next host with found ssh keys and repeat from 1.



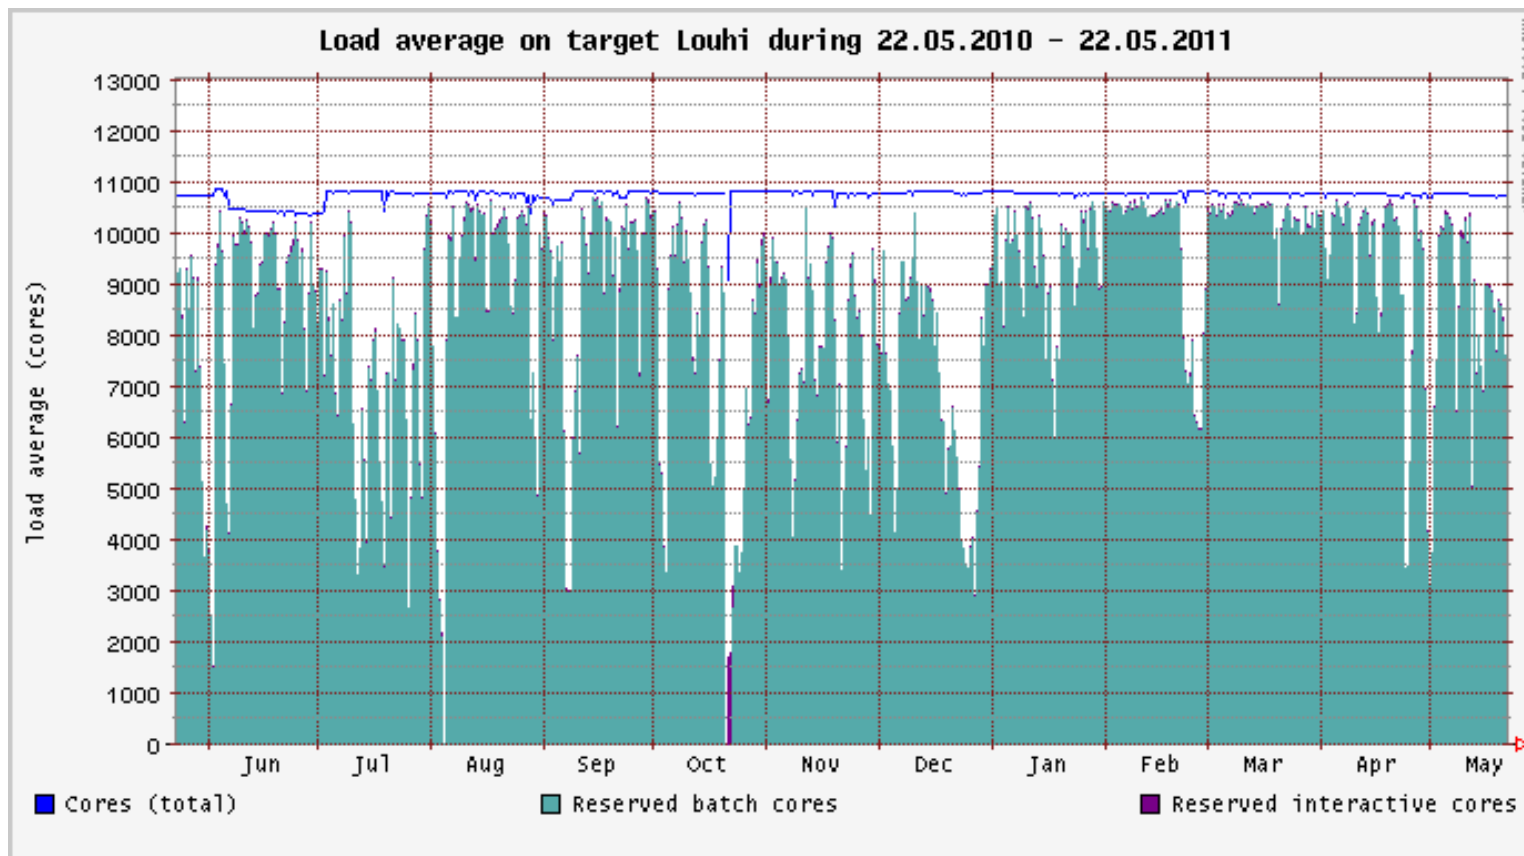
# Case intrusion of Murska 2010



- September 29: a password stealing rogue sshd binary was found in a login node of Murska
- Similar steps to react as in the Louhi case
- intrusion was made with a legitimate user account
- intruder used unpatched vulnerability CVE 2010-3081
- installed a rogue keylogger to collect passwords
- The intrusion affected some other sites
- October 6th: back on line



# Louhi load average




# Availability metrics



Host	2008	2009	2010	2011	Availability
<b>All breaks (scheduled and unscheduled)</b>					
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	100.00%
Hippu	N/A	N/A	N/A	26 min	99.98%
<b>Unscheduled breaks (scheduled breaks removed)</b>					
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	100.00%
Hippu	N/A	N/A	N/A	26 min	99.98%


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



# Availability metrics Contd.



	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%
<b>TOTAL</b>	<b>100.00 %</b>	<b>100.00 %</b>	<b>N/A</b>	<b>100 %</b>



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

	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)**

# Hypothesis



1. Brute-force attempts from a limited set of sources..
  - True, but many advanced attacks might remain undetected.
2. Plenty of downtime from security incidents
  - True. Improving security can also advance efficiency
3. Less intrusion with better access controls
  - Partly true. The introduction of **Denyhosts** did filter out most of the brute force attacks
4. IDS will diminish downtime
  - Most probably true. At least faster recovery
5. Adequate user management protects against brute force attacks
  - Mostly True
6. Implementing access controls in an optimal way requires sharing of best practices
  - 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.

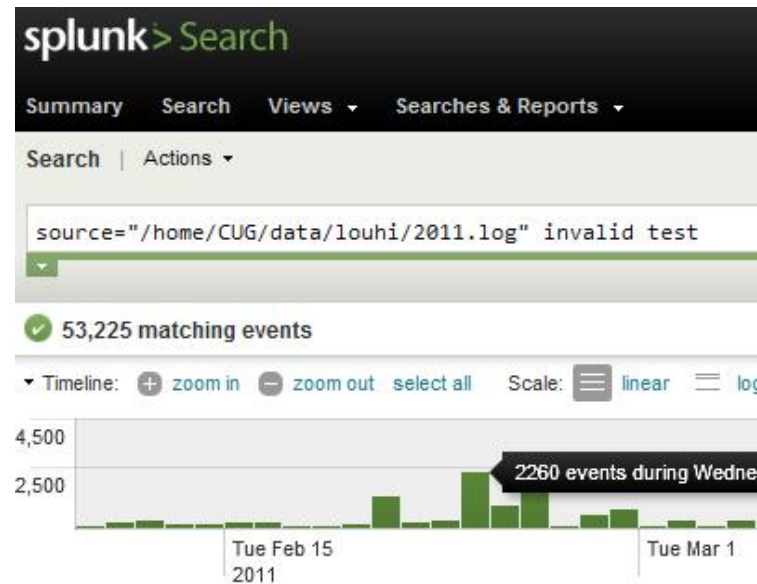




# Security tools and best practices



- Splunk
  - Log analysis
- Denyhosts (as an easy to use alternative to fail2ban)
  - Block brute force attacks
- Nagios BPI
  - Availability metrics
- AIDE
  - Integrity testing
- John the Ripper
  - Password quality testing
- Nessus
  - Network scans (/w credentials)
- Rkhunter
  - Search for rootkits and insecure configurations
- Security policies, agreements and guidelines (BCP's and DRP's)
  - BOF's, Training, professional certifications, site visits
- Metrics and audits!




# Conclusions and suggestions



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



Suggestions:

- let's share between sites best practices for deploying security tools , such as DenyHosts, Aide, John the Ripper, IDP,..
- A great example: <http://www.arsc.edu/~kcarlson/CUG2010/Carlson.html>
- Let's start an exchange and visiting program for junior and senior administrators for best system administration and information security practices.
- Also service and security management should be involved



# Thank you for your attention!

Now we would like to have some feedback and discuss our findings and suggestions with you!



**Urpo Kaila**



**Marco Passerini**



**Joni Virtanen**

**Email: `Firstname.Lastname@csc.fi`**