



HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Automatic Discovery of Diverse and Changing Network Services

AMICT 2009 Workshop

Petrozavodsk State University

19th May 2009

Mikko Pervilä, prof. Jussi Kangasharju (instructor)

Department of Computer Science, University of Helsinki



Presentation Outline

- Goal: the ratio of **common-mode** (CMF) to normal failures
 - Most common causes for CMFs
- Describe a work in progress measurement framework
 - Some self-healing also a possibility
 - Data suitable for Bayesian analysis
- Main problem: the environment keeps changing
- Fixes: automatic discovery, distributed monitoring



CMFs – The Basics

- From *Fault Tolerance by Design Diversity: Concepts and Experiments* by A. Avižienis and J. Kelly, 1978:
- N -fold computation in time, hardware, and software
 - **Repetitions** from (1T / 1H / 1S) to (XT / YDH / ZDS)
 - D is for **diversity**
- M -plex faults affect M out of the N computations
 - The faults may either be independent or related
 - Their cause may either be operational or **by design**



CMFs – Well known in early CS

- Dionysius Lardner, Babbage's Calculating Engine, in the *Edinburgh Review*, July 1834:
 - “The most certain and effectual check upon errors which arise in the process of computation, is to cause the same computations to be made by separate and independent computers; and this check is rendered still more decisive if they make their computations by different methods.”



CMFs – What is in a name?

- “Common-mode failures” more common than “*M*-plex”
- First occurrence from 1930 (?) in the *Journal of American Ceramic Society* by J. Otis Everhart
 - “The common mode of failure in the autoclave is by crazing [...]
The common mode of failure during freezing is by spalling, [...]”
- Physical stress and temperatures seem to be reoccurring themes



CMFs – How common are they today?

- Nvidia GPU Failures Caused By Material Problem, Sources Claim. [Tom's Hardware](#), Aug. 26th, 2008
 - \$200 million for repairs

- Microsoft Zune 30 GB [meltdown](#), Dec. 31st, 2008
 - Bad leap year parsing code causes device lockups

- Enter the Poorly Designed MLC, [AnandTech](#), Sep. 8th 2008
 - Some SSD controllers cause random 1 second writes

- Seagate firmware fix bricks [Barracudas](#), Jan. 21st, 2009
 - Firmware fix for 1 TB drives causes 500 GB drive failures



CMFs – User reports are problematic

- The problem with these reports is their **credibility**
 - Reported by home users, enthusiasts, and hardware sites

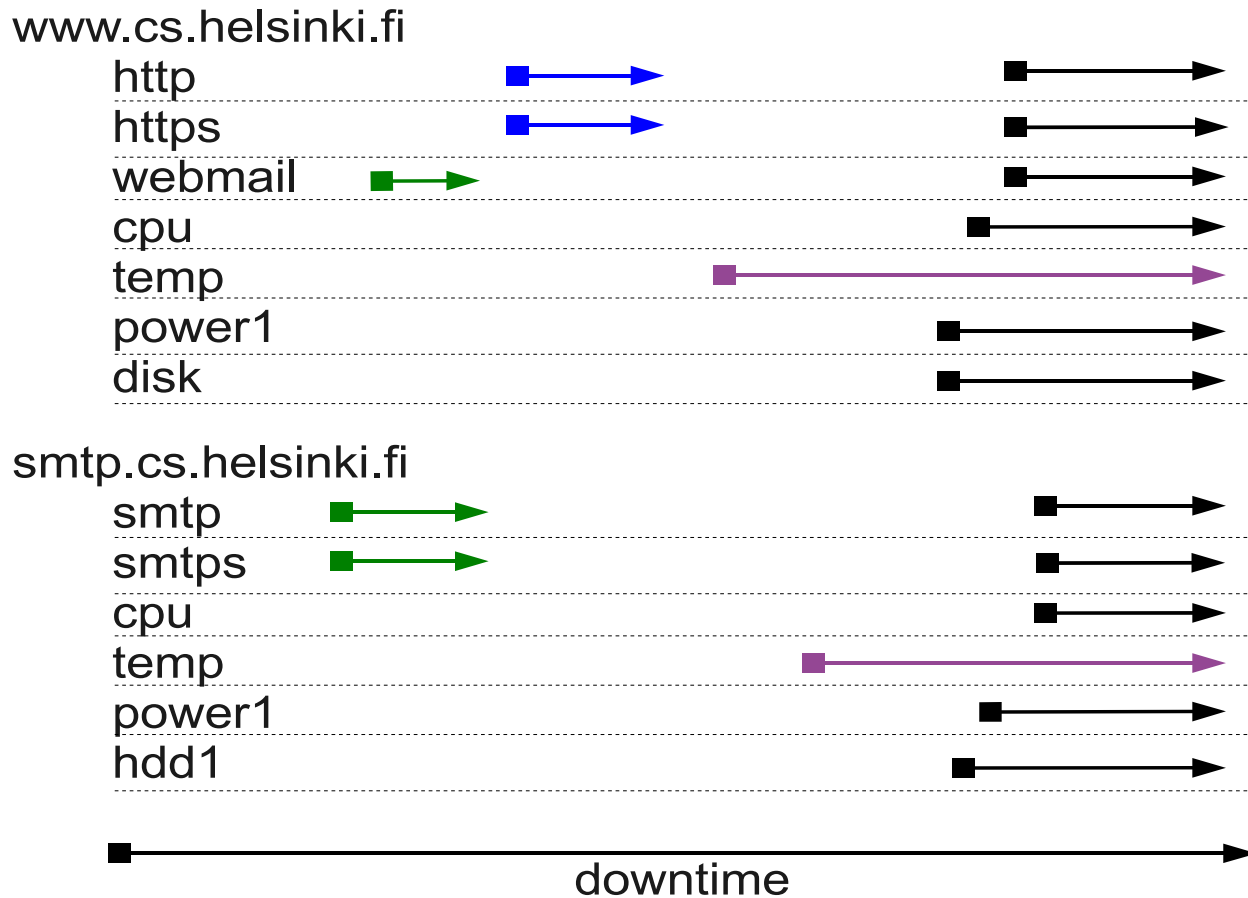
- Scientific background of the reporters a question
 - Methodology?
 - Bias?
 - Repeatability?

- Product failure rates are business secrets
 - Data sets seldom available



CMFs – Measurement goal

- Study **related** downtime; Data mining, Bayesian models





Nagios – The sentinel service

- Basic idea: run input / output checks against services
 - Versatility: checks run by plug-ins; any program code
 - Nagios handles scheduling and interleaving checks
 - Output outside given parameters causes a notification
- Primary focus: **network** services
- Distributed monitoring catches local services
 - Fan speeds, temperatures, SMART attributes for storage, ...



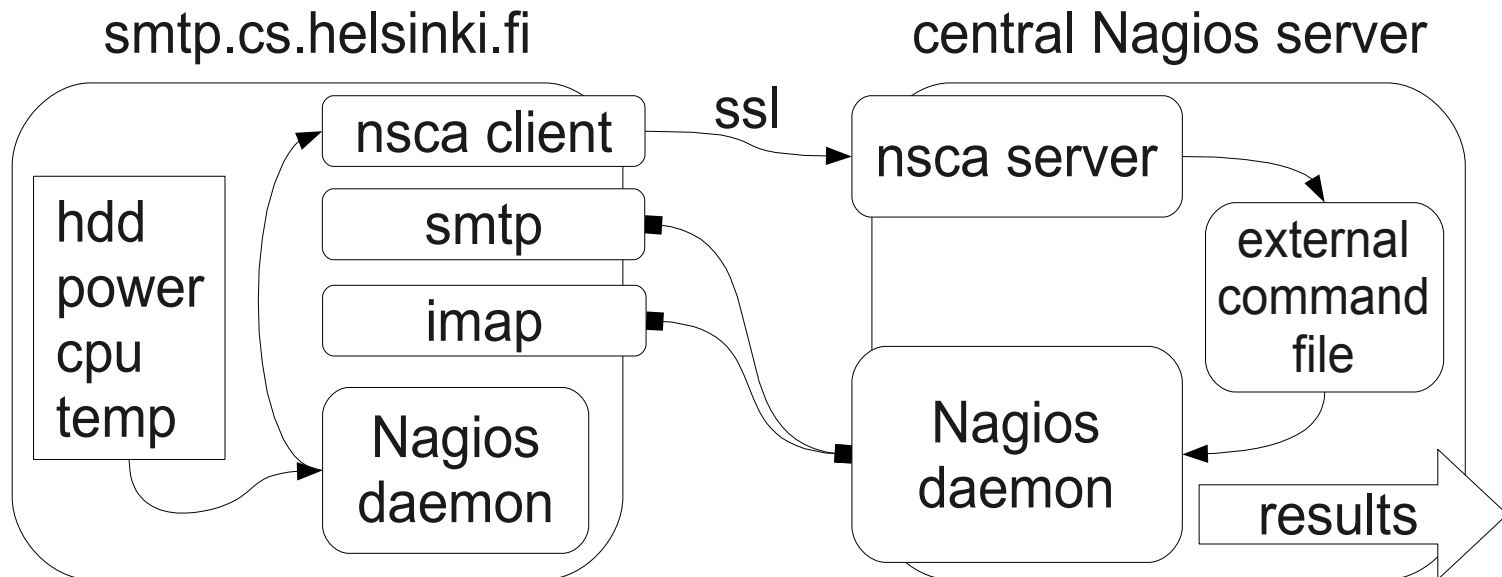
Nagios – Network services

- Monitoring the CS Dept. network is challenging
 - New hosts and services come and go
 - Research groups administer their own hosts
- Partial solution: [Nmap Security Scanner](#)
 - Scans IP blocks, discovers services
 - Nmap produces XML output
- Nmap → [Nmap3Nagios](#) → Nagios
 - Our open source tool for configuring Nagios



Nagios – Local services

- Distribute **local** Nagios daemons
 - Run checks against local services
 - Nagios' client-server tunnel NSCA reports back
 - Results may be stale if workstation is shut down





Nagios – Self-healing

- When a service malfunctions
 - Plugin notices abnormal output
 - Nagios notifies administrators with mail, SMS, ...
- Nagios can also call external event handlers
 - Event handlers perform scripted actions
 - E.g., restart services, analyze log files
- Requires special privileges
 - But very flexible



Nagios – Problems

- For administrators, fixing problems is a priority
 - Acknowledging Nagios secondary
 - Planning downtime tertiary, or even less
- Nagios' GUI very old school
 - Administrators can not redefine hosts or services
 - Not integrated with local issue trackers (yet)
 - Many alternative GUIs, none really good for us



Nagios – Problems cont'd

- Nagios is a delicate instrument
 - It detects failures usually invisible for human users
 - Scheduled backup runs
 - Automatic software upgrades

- Service dependencies complex
 - Manual work still necessary
 - Where should dependencies be stored?
 - NACE tool uses SNMP fields for this

- Dual-booting between Windows and Linux



Conclusions

- Common-mode failures seem very common
- Monitoring failures can be done, requires work
- Keeping up with administrators very difficult
- Working on a toolkit, will publish data



Questions, Comments?

- Nmap3Nagios tool available from
 - <http://www.cs.helsinki.fi/u/pervila/Nmap3Nagios/>
 - Other tools will follow

- pervila@cs.helsinki.fi

Thanks - Спасибо!

DO NOT ENTER

**IF THE FALL DOES NOT KILL YOU
THE CROCODILE WILL**

