

PSI Site Report

(Spring HEPIX 2009)

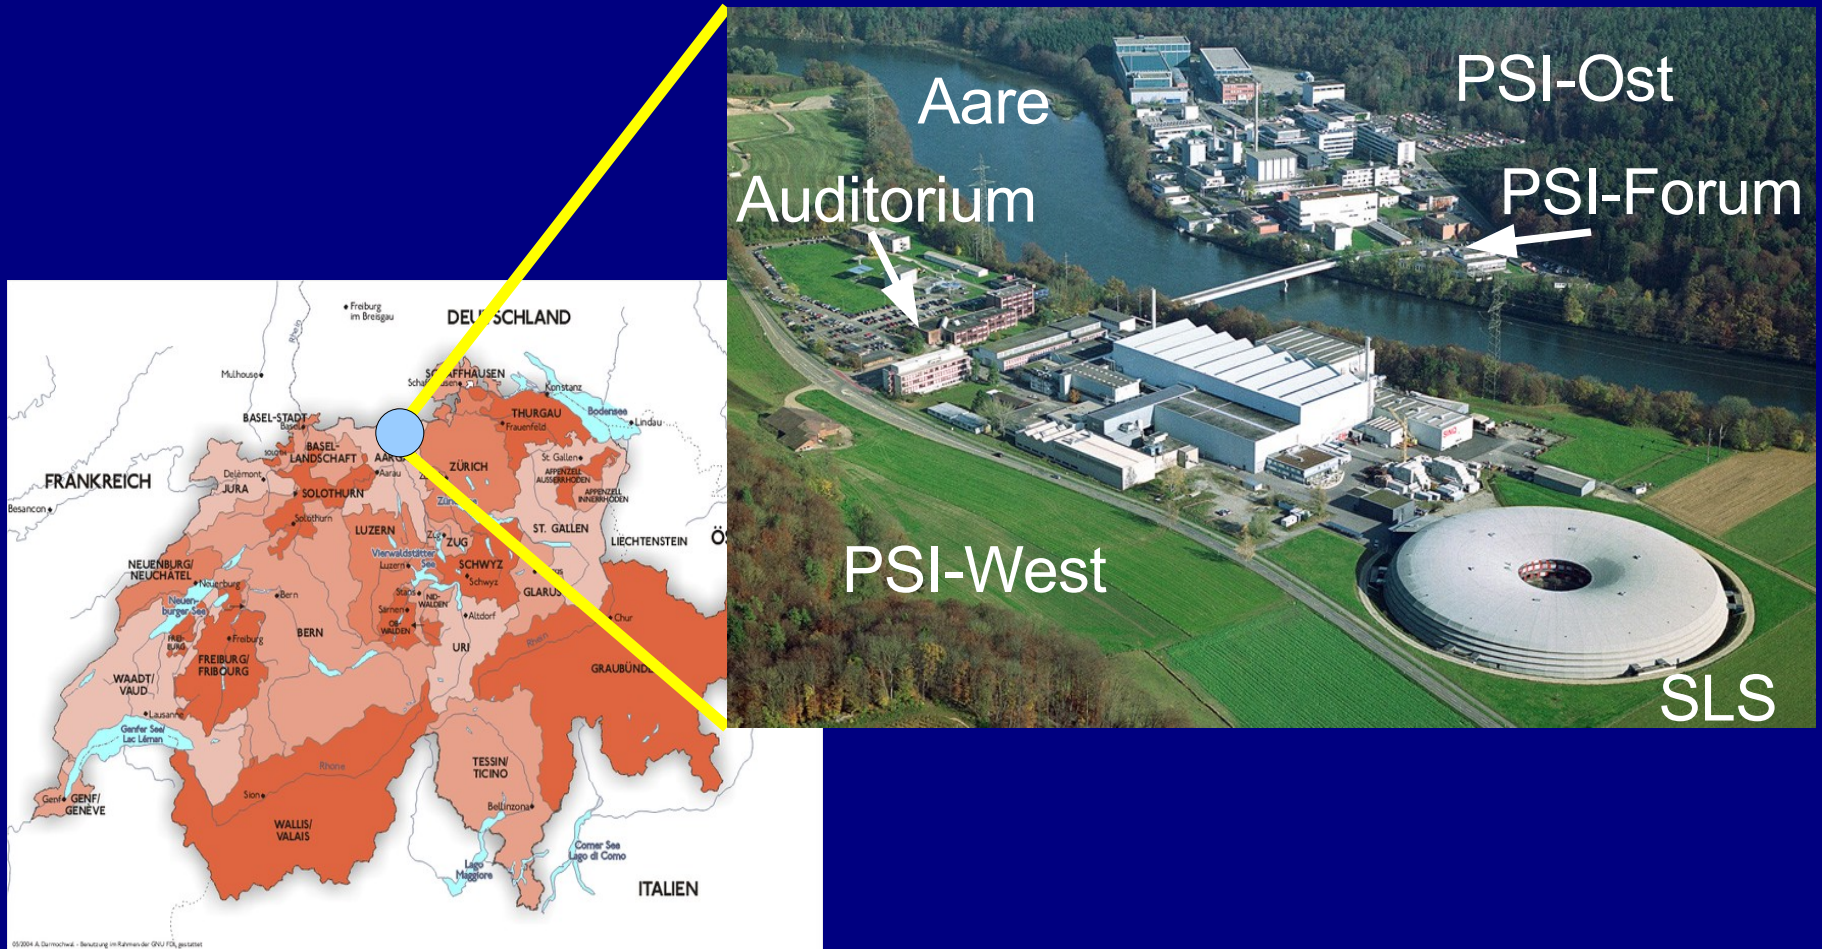


Gasser Marc

Paul Scherrer Institute, Switzerland

www.psi.ch

Paul Scherrer Institute, Villigen, CH



PSI: a brief Introduction

- **A leading national multidisciplinary research laboratory**
- **ca. 1670 Employees - 320 financed through 3rd party funding.
Ca. 2000 p/day on Campus**
 - Solid State Physics and Material Sciences (44%)
 - Life Sciences (17%)
 - Particle- and Astrophysics (15%)
 - Nuclear Energy Research and Safety (12%)
 - General Energy Research (12%)
- **User Lab (ca. 2300 long and short-time international guest scientists from ca. 50 nations)**
- **Education (Apprentices: 78, Doctoral Students: 300, Employees with Teaching Responsibilities: 75, Radiation School: 3000)**

Computing Infrastructure

- **Many suppliers/architectures:** Intel, HP, FS, SUN, IBM, SGI, Mac, Dell, ... (5500 IP's, 700 Wlan)
- **Many Operating Systems on Campus (not necessarily supported):**
 - WinXP Professional (2500)
 - Scientific Linux (700)
 - MacOS (300)
 - Win2000, WinNT, Win9*, VMS
 - Fedora, OpenSuse, Ubuntu
 - SunOS, Tru64Unix, AIX, HP/UX
- **Academic environment with many cultures & languages, balance between flexibility & security**

Computing Infrastructure

■ Numerous Applications:

- Email, WWW, FTP, CMS, DMS, Office, Graphics, Scientific, Visualization, Databases, High Performance Computing, Parallel Processing, Data Acquisition, Backup, Archiving, Monitoring
- High availability, 24h x 365d per year operation and usage

■ Printing:

- 250 network printers, 1000 desktop printers

■ IT Service Desk

- One single point of contact for customers
- First level support
- Forwarding requests to second level support

Computing Infrastructure: Examples

- **PSI operates an LCG Tier-3 cluster for the CMS groups of ETHZ, PSI and University of Zurich**
- **High Performance Computing Clusters:**
 - On site: small clusters, new HP enclosures with blades
 - Off site: joint ventures with Swiss National Supercomputing Centre (www.cscs.ch), Cray XT5
- **Filesystems:** EXT3, XFS, AFS (700 Lin., 300 Win., 100 Oth.), GPFS
- **Virtualization:** Vmserver2, ESX
- **Wiki:** Twiki
- **Telephony:** 2000 wired, 1500 Dect; migration to VoIP in 3 years

Linux

■ Scientific Linux PSI (SLP)

- SLP5 based on SL51, 550 systems
- SLP4 based on SL46, 150 systems
- Computer classes: Desktop, Server, ClusterNode

■ Customized Kickstart Installation

- Network installation via PXE boot and nfs or http
- Same framework for Desktop, Cluster and Server installation
- Kickstart file is modular. Customization scripts are used
- Desktop configuration done by puppet

■ Update of SLP5 based on SL53 in process

- New kickstart scripts
- Dedicated webserver for software (RPM) repositories

Linux

■ Puppet

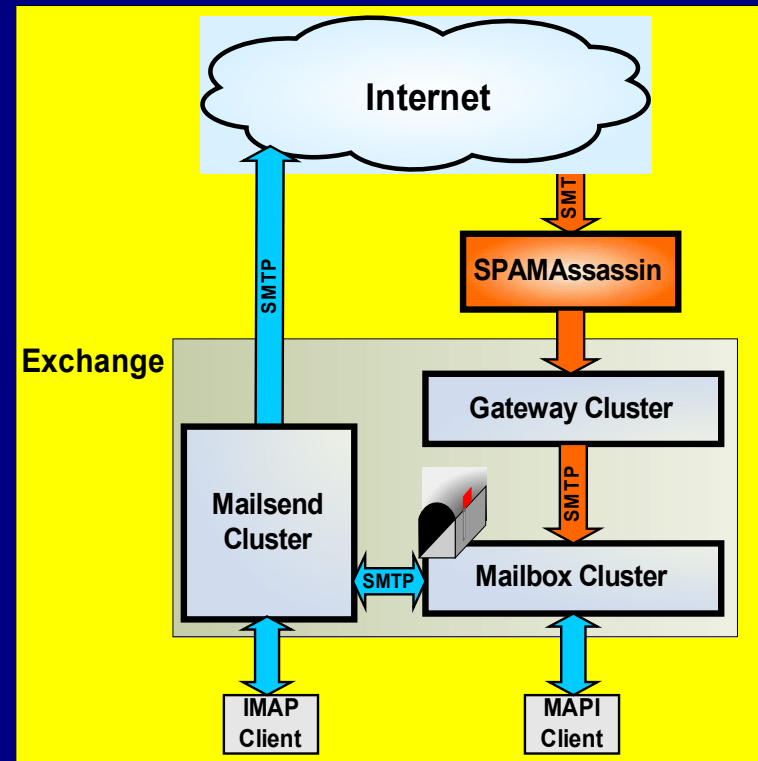
- One virtual puppet server in use for about 250 Desktop clients
- Problems:
 - Bad performance
 - Puppet server often crashes (0.24.1-1.el5)
 - Bad scalability and overview of current client configuration manifests
- Solution:
 - Update puppet server
 - Run several puppet servers, virtual or real hosts?
 - Reorganization of client configuration manifests

Spam Filter

■ Daily Spam mail volume @ psi: ca. 500 000

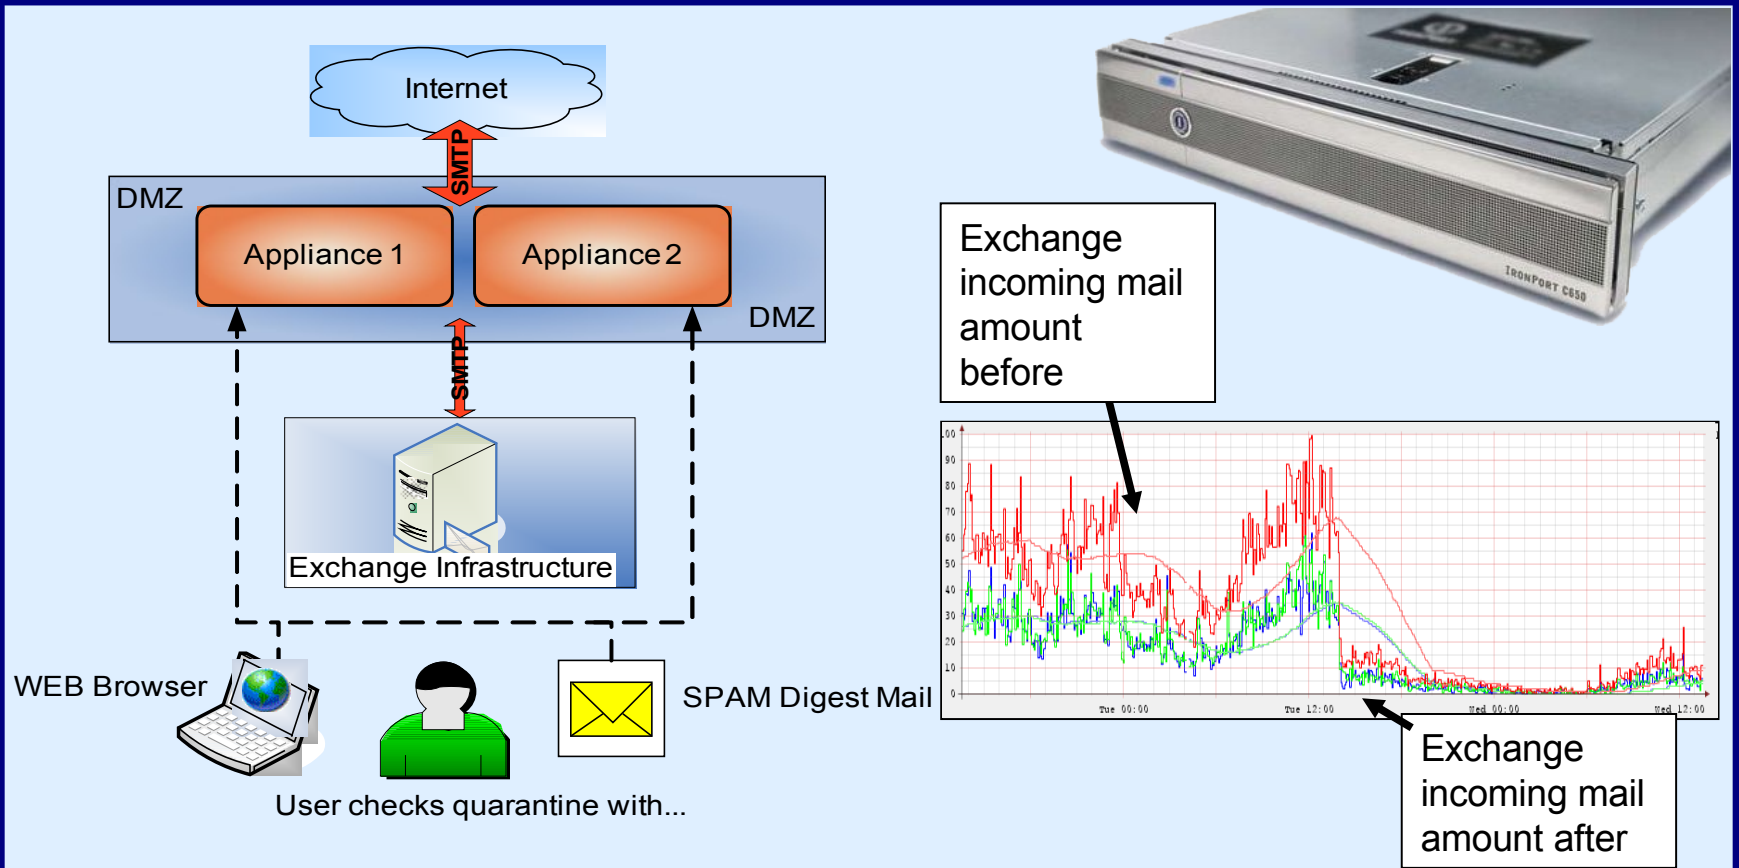
■ Currently, Spamassassin:

- Old Hardware
- Software update required
- No quarantine
- Too much administrative effort




Spam Filter

■ New proprietary spam filter:



Thank you for your attention



LCG Tier-3 cluster for CMS

- PSI operates a LCG Tier-3 cluster for the CMS groups of ETHZ, PSI and University of Zurich
- Storage Element: 6 Sun X4500 “Thumpers” (6 * 17.5 TB of RAID/Z) + 2 Linux head nodes
- Compute nodes: 8 * Sun X4150 (2*Xeon E5440, 16 cores)
- Storage available via Grid tools for stageout by external jobs
local jobs run via a local SGE batch farm (no grid enabled CE foreseen) System size

Year	Compute cores	Storage / TB
Q3 2008	64	105
Q4 2009	184	250

LCG Tier-3 cluster for CMS

CMS ist eines der vier grossen Experimente am CERN (CMS, ATLAS, LHcB und ALICE)

Für die Analyse der gigantischen Datenmenge des LHC-Experiments am CERN wurde über die letzten 8 Jahre das weltweites "LHC Computing Grid" (=LCG) ins Leben gerufen. Das Grid besteht aus einer hierarchisch organisierten Anzahl von Clustern, wobei

Tier-0: CERN

Tier-1: übernationale Zentren (ca 7, je nach Experiment)

Tier-2: nationale Zentren (in Manno am CSCS für die Schweiz)

Tier-3: Typischerweise Cluster einer Universität (unser PSI Zenter)

Benutzer können das Grid transparent als einen grossen Supercomputer bedienen. Sie können über sog. Grid User Interface Maschinen Jobs losschicken, die automatisch zu den Zentren laufen, die die zugehörigen Daten hosten. Der Benutzer erhält dann die Ergebnisse lokal auf seiner Maschine zurück.