Computing at MPP

Stefan Kluth

MPP Computing commission, chairman

MPP Project Review, 15.12.2009

Personnel C/N Group

Head of group: Leupold, U.

	Linux	MS	SW	LAN	HW	Pr	Grid
Leupold, U.	x	x	x	x	x		
Kriesel, A.	x(reg.)				x	x	
Krämer, M.				x			
Pan, Y.	(x)	x					
Salihagic, D.	x		x				x
Vidal, M.	x(AIX)		x		x		
Krebs, K.				x			

Hardware Overview

- Central servers
  - 2 IBM BladeCenter, 56 cores (ATLAS)
  - 28 fileservers with ~85 TB
  - mail, web, DNS, accounts, backup, printer, ...
  - > 500 cores for batch via condor (server+PCs)
- Experimental and engineering groups
  - ~300 PCs, 2/3 Linux, 1/3 MS
- Theory group
  - ~80 PCs (Linux), 1 DEC Alpha

Hardware News

- Homedir servers
  - RAID, failover (in prog.)
  - ATLAS/ILC/(s)BELLE
- WWW/indico
  - failover
- Mail service
  - 2 servers to handle load (~200k smtp req./d)
- New WLAN APs

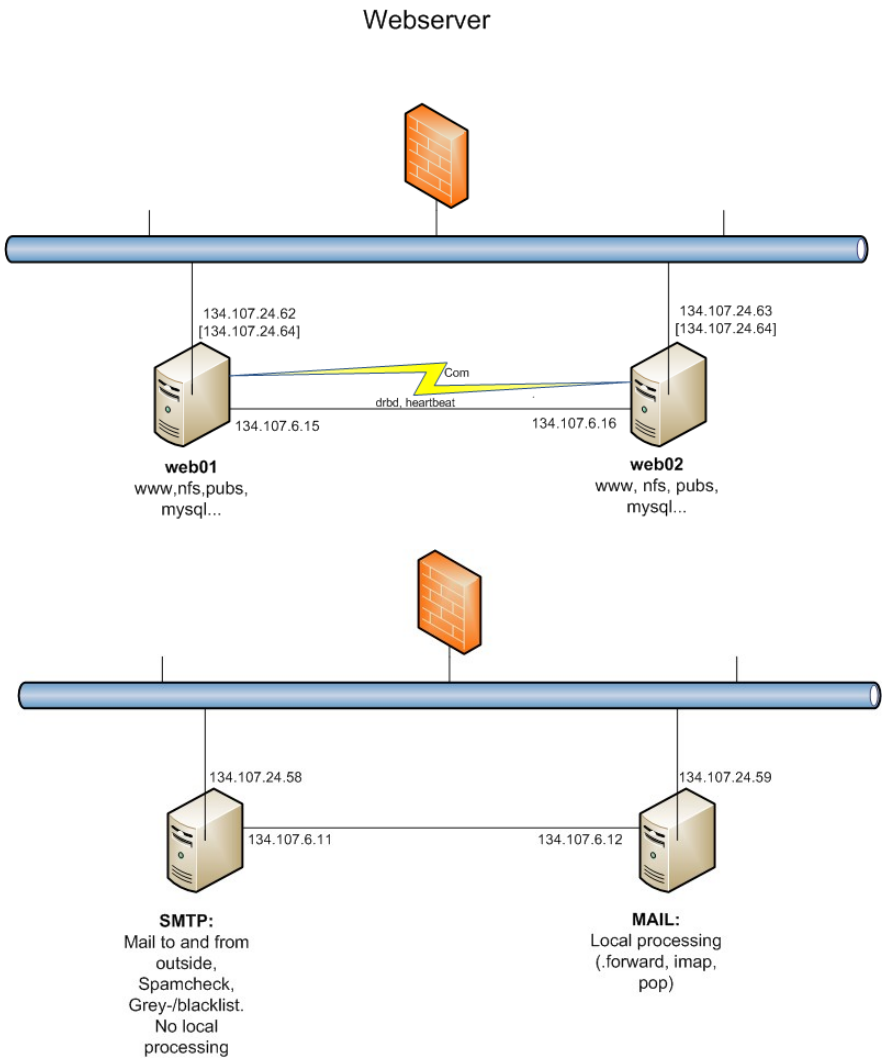

LAN

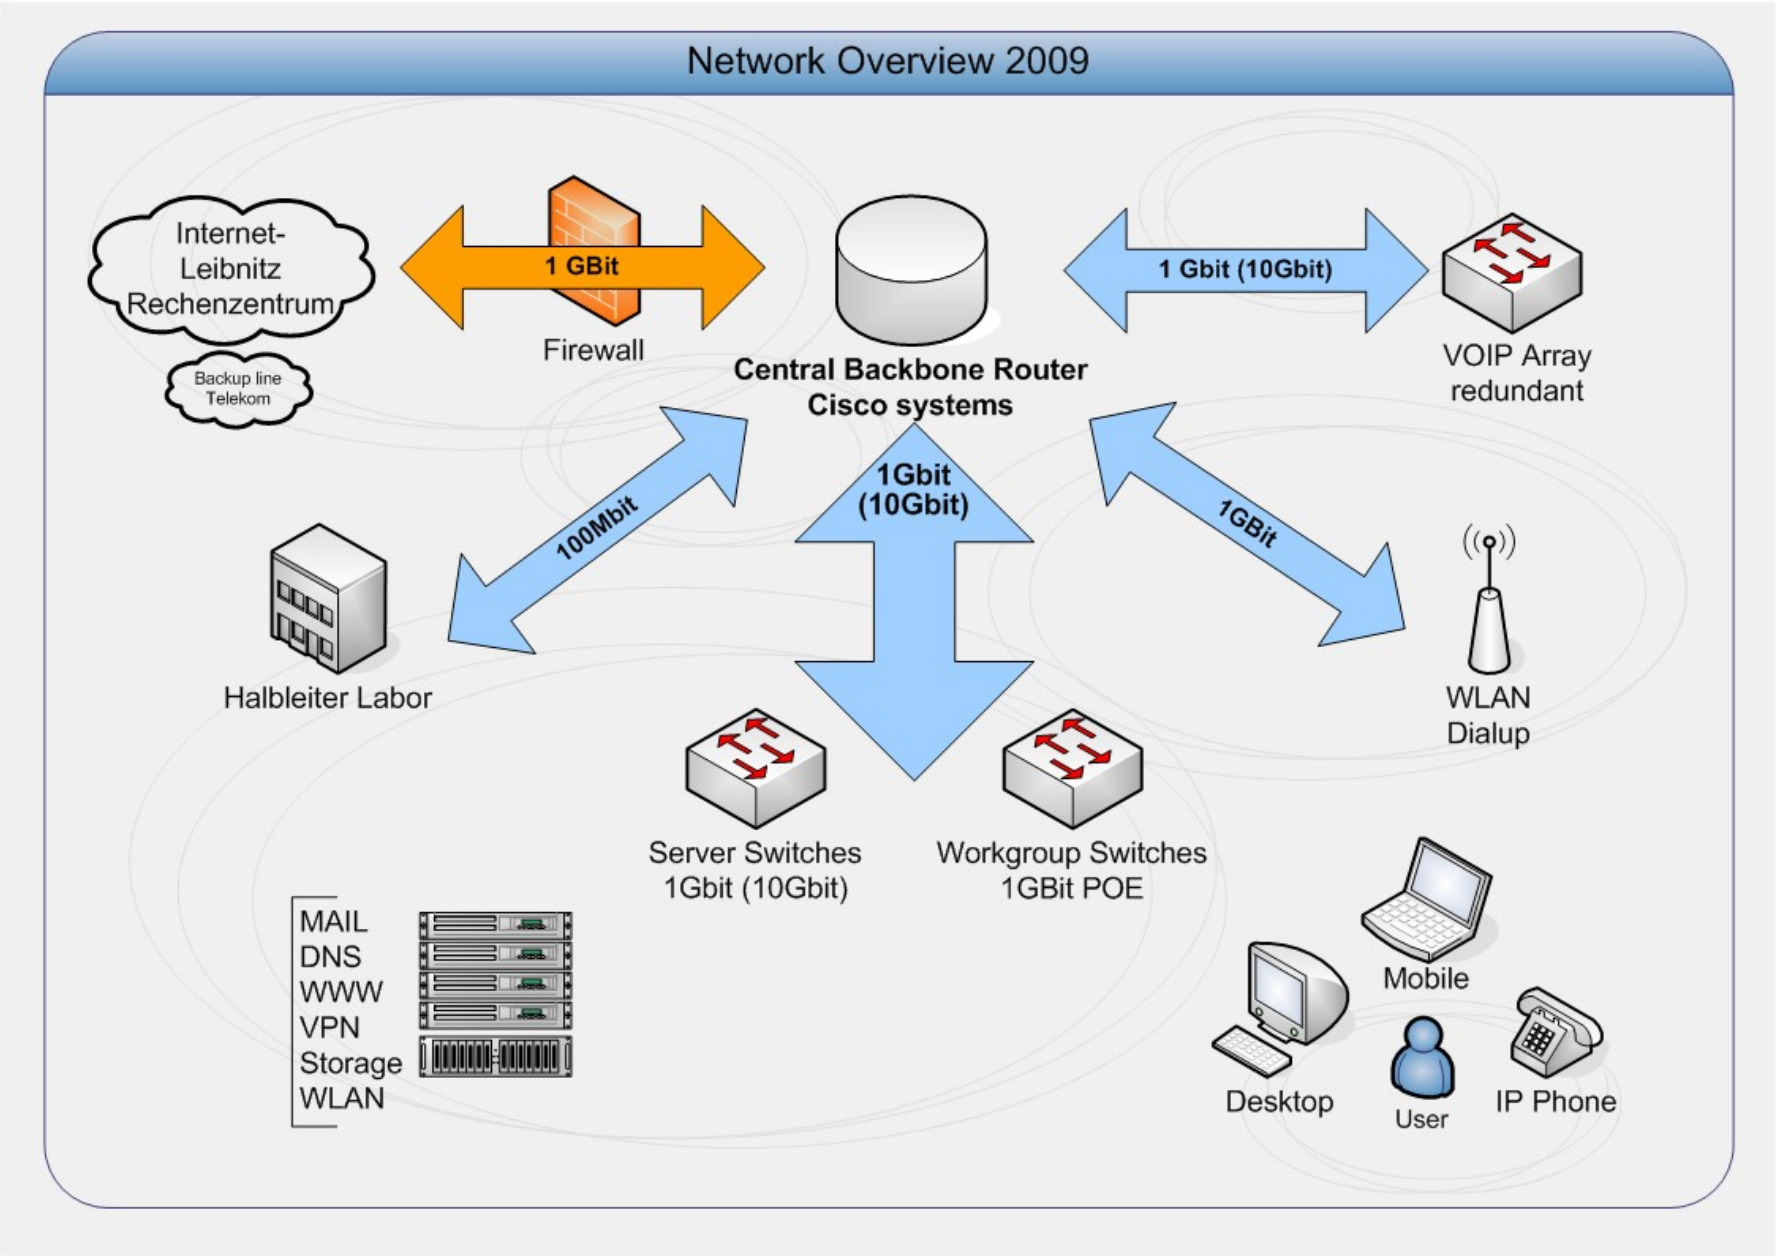
Network Overview 2009
Internet-
Leibnitz
Rechenzentrum
Backup line
Telekom
1 GBit
Firewall
Central Backbone Router
Cisco systems
1 Gbit (10Gbit)
VOIP Array
redundant
100Mbit
1Gbit
(10Gbit)
1GBit
Halbleiter Labor
WLAN
Dialup
Server Switches
1Gbit (10Gbit)
Workgroup Switches
1GBit POE
MAIL
DNS
WWW
VPN
Storage
WLAN
Mobile
Desktop
User
IP Phone

Software overview

- OS
  - Linux (MPI, Debian → Ubuntu (ATLAS), peb (theory) → Opensuse)
  - AIX (1 server left), Solaris (Elektronik)
    - MacOSX, VMS without C/N support
  - MS Windows (Admin, Labs, h1win, thwin)
- Applications/libraries
  - Mathematica, Maple, Portland pgf compiler, IDL, NAG, Matlab (Theory), ifortran
  - OSS in consultation with C/N group & mpicc

Software Applications

- Commercial
  - Oracle 10, Infoparc cms Fiona, Tivoli (backup), Wilma (WLAN manager), Gleitzeitserver
- OSS
  - Indico, mysql, root, CERNLIB, phpBB, twiki, Asterisk, eGroupware, ...

MPP WWW Pages

- Online since summer 08 (50th anniversary)
- Structure improvements
  - job offers, library, DE/EN pages, media archive pages, calendar
- PHP interfaces
  - calendar (indico), publications, MPP member search
- Ongoing
  - complete DE/EN, member search, complete indico integration, intranet

New Phonesystem

- MX-ONE plus Cisco/Asterisk VoIP system
  - consulting (BayCom) done
- In-house
  - new GBE LAN with PoE and vlan (done)
  - prototype service with O(10) phones (done)
  - install more phones (ongoing)
- Future (2010)
  - all users with new phones
  - phase out MX-ONE

Phonesystem Migration

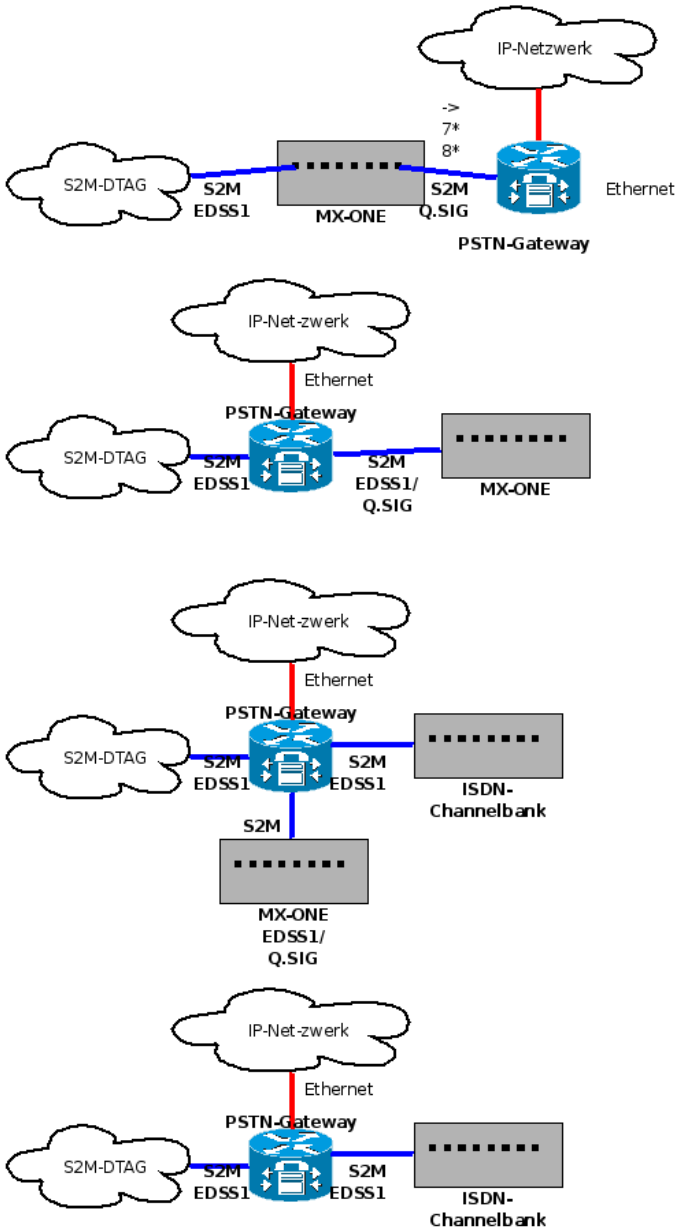

1. MX-ONE und VoIP parallel, replace phones, keep numbers
2. Phones replaced, MX-ONE serves old legacy phones
3. Connect ISDN channelbank, connect ISDN legacy phones
4. MX-ONE switched off

C/N Plans 2010

- ATLAS-ILC/Belle homedir servers
- Consolidate “RZF”
- New phones
- Replace central router and firewall
- LAN 1 → 10 GBE
- Linux migrations
- Request tracker (otrs)
- nfs3 → nfs4/afs?
- nis → Krb/LDAP?
- Condor → SGE?

MPP computing commission

- Subcommittee of IA
  - members: Bethke, Hahn, Kluth, Leupold, Wagner, Schieck, Reimann
- Mandate
  - oversight of C/N operations
  - medium- and longterm planning
- Please consult before buying hardware or requiring services

Rahmenverträge

- MPG has procurement contracts for IT
  - hardware, LAN, software, services
  - order without tendering procedure < 220k€
  - competitive prices (still check F&L or street prices)
  - FSC, HP, Acer, Dell, Lenovo, Apple, ...
- Experience
  - good for standard orders
  - companies respond well

Computing at RZG

- Operation
  - main users ATLAS (WLCG), MAGIC
  - open for all MPP groups (10% share)
  - connection via 1 GBE link
- Usage
  - need RZG account (via web-form)
  - direct access e.g. “ssh at01.t2.rzg.mpg.de”
  - direct access to dcache storage (dccp, dcap)
  - direct access to tape storage via afs

Computing at RZG

- Current status
  - 812 cores, 2 GB/core, > 300 TB disk
  - ATLAS Tier-2/3 and MDT calibration
  - MAGIC analysis centre
  - Others: theory, GERDA, ILC, BELLE(II)
- Plans for 2010
  - + > 336 cores, + > 500 TB disk, new LAN
- Other tasks possible at RZG
  - AIX (xlf), parallel computing

Computing at RZG

- Software
  - SLES10, WLCG, dCache, afs, nfs, SGE
  - access to tape storage (via afs)
  - Various gcc versions
  - SLES11 in 2010?
  - Experiment software
- Functions
  - WLCG: send/receive grid jobs and data
  - Local SGE batch jobs

Summary and Trends

- MPP IT landscape continuously changing
  - rapid changes in hard- and software
  - keeps C/N staff busy
- Mass data storage at MPP growing fast
  - > 100 TB soon? CPU power? Management?
  - consider using RZG cluster
- IT security
  - needs a well-managed setup
  - requires some protocol upgrades