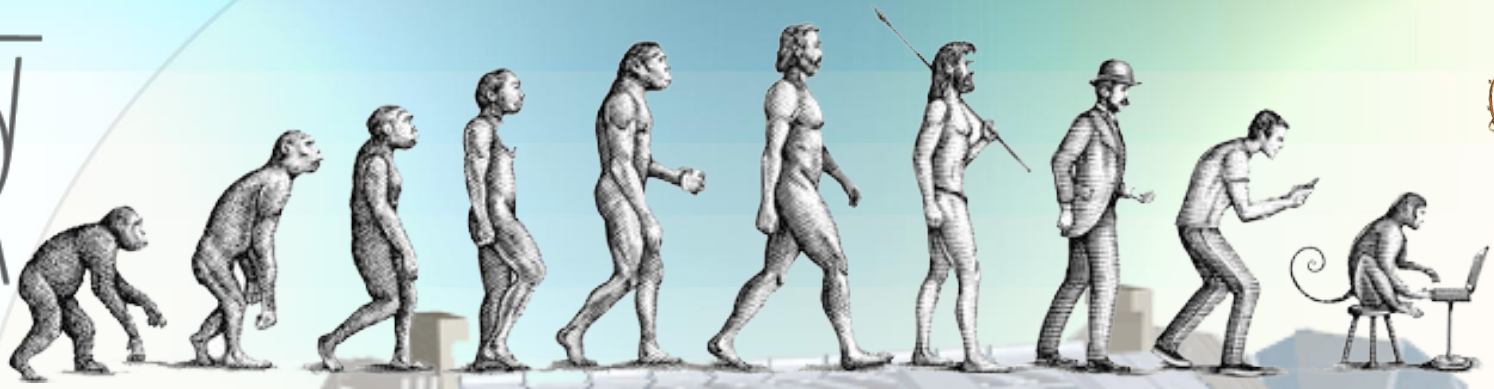

ATLAS TDAQ System Administration: an overview and evolution

ISGC 2013

Christopher Jon Lee

University of Johannesburg, South Africa
CERN

for and on behalf of the

ATLAS TDAQ SysAdmin team.

Introduction
LHC & ATLAS

- Large Hadron Collider, an accelerator, ~100 m underground
- 27 kilometres in circumference
- Hadrons are accelerated in opposite directions at 4 TeV
- Smashing together in the center of ATLAS, one of 7 experiments
- 600 million collisions per second
- Data from these collisions are recorded by the Trigger and Data Acquisition system
- After 3 years of LHC runs, Long Shutdown 1 (LS1) has started

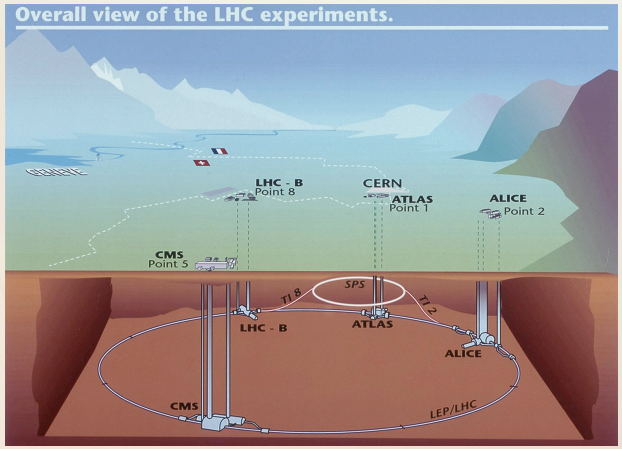

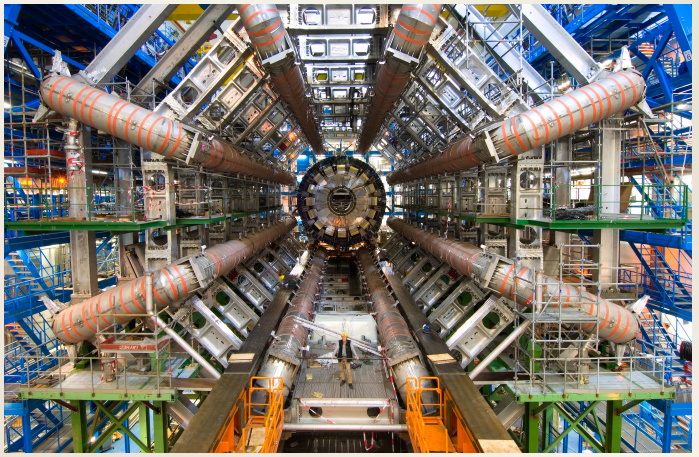

Introduction
Trigger and Data Acquisition

- Large online computing farm, used to process the data readout from ~100 million channels
- Ancillary functions (monitoring, control, etc.)
- ATLAS Point 1 – Counting Rooms
  - approximately 100 m underground, in close proximity to the detector (USA15)
  - on the surface near to the ATLAS Control Centre (P1, SDX1 & SCR)
- In General Public Network (GPN)
  - laboratory for software development, prior to implementation into P1, recently been commissioned (TestBed)

The Racks

- USA15:
  - 220 racks deployed over 2 floors
  - 2009, ~70% filled, 1MW
  - 2013, > 90% filled, 1.29MW*
  - 2.5MW of cooling can be provided
- SDX1:
  - 120 racks deployed over 2 floors
  - 2009, ~50% filled, 385 kW*
  - 2013, ~91% filled, 709 kW*
- TestBed:
  - 22 racks over 1 floor
  - 68% filled, 60kW of power. 11 Racks by TDAQ
  - 100 kW of Cooling can be provided
- Possibly an increase in hosts to deal with future requirments

*estimated at ~6.5 kW/rack

TDAQ Design Overview

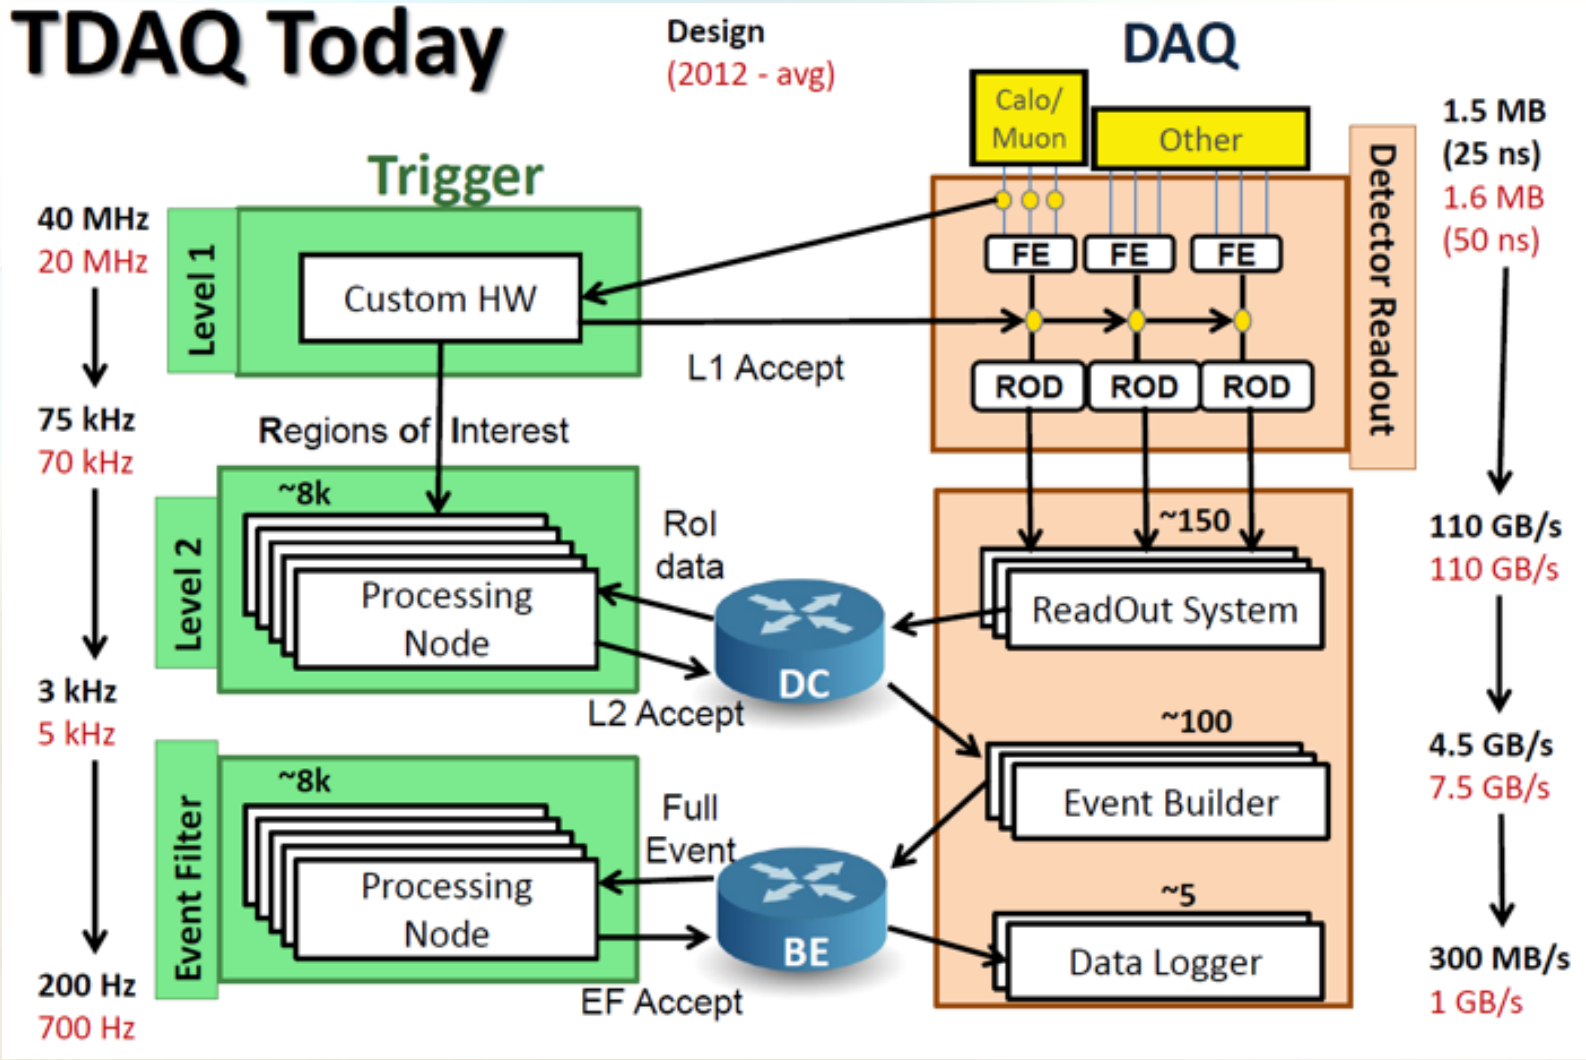

* Image made by Nicoletta Garelli - Athens Trigger Workshop

ATLAS Point1 Functional Layout

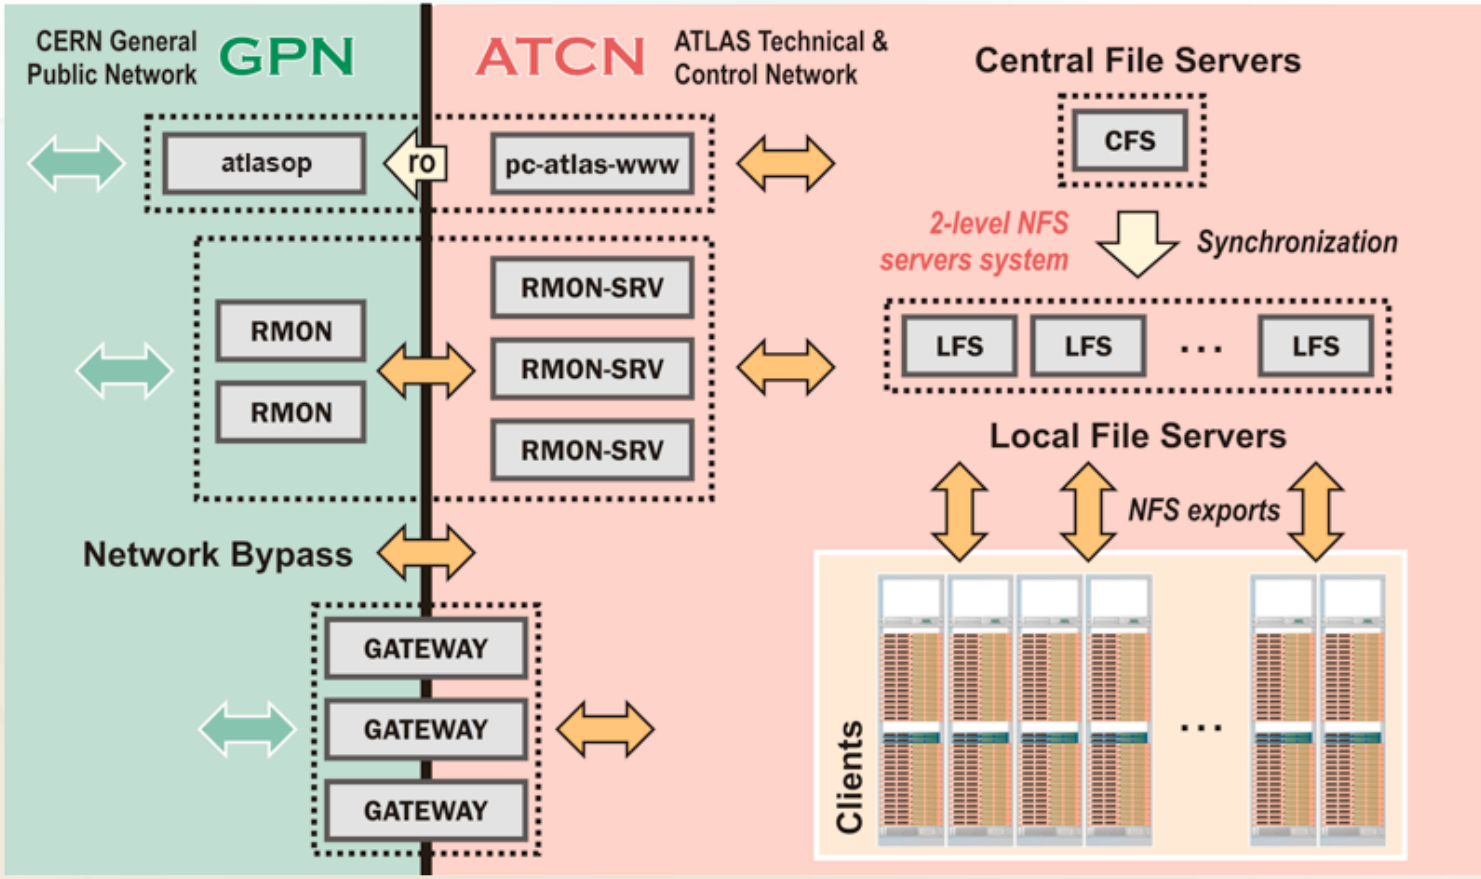

Redundancy

- Can't afford downtime & miss out on events provided by the LHC
- Two centralised UPS lines with diesel backup generators
- Independent UPS lines available to SDX1 for mission critical equipment
- Redundant configuration for critical services
  - DNS/DHCP/NTP
  - LDAP
  - Domain Controller
  - LFS'
- Network Attached Storage, serves the most critical NFS and CIFS areas
  - NetApp 3140, it has 12 TB over 84 HDD's in RAID DP, dual head system and redundant fibre channels
  - can survive the loss of one head, any one FC link, one FC interface on a shelf - just not of a whole disk shelf
- CERN's Tivoli for critical data
- Subversion for code and configurations

Centralised & Local Storage Systems

- CFS Linux node - no direct exports:
  - Trigger/DAQ and Offline Software installation
  - Coordination of synchronisations to LFS'
- NetApp Central Filer - NFS and selected CIFS exports :
  - user home directories, DAQ software distribution
  - Nagios RRD* files
  - node configuration area
  - DAQ configuration area
  - static and dynamic Web content
  - dedicated file exchange area for gateways
- LFS
  - provides boot services and NFS exports for clients
  - synced from NetApp

ATLAS Point1 Functional Layout

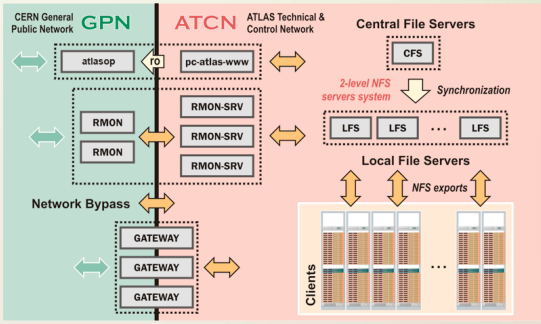

* round-robin database

Clients
Local Boot

- Provisioning by PXE + Kickstart
  - DHCP + PXE provided by an LFS from ConfDB* information
  - template-based kickstart files
- Quattor
  - CERN standard Configuration Management Tool
  - production system, managing 237 hosts in the Online Farm
  - tight control on installed packages
  - lack of flexibility for complex configuration/services dependencies
  - multiple languages for implementing modules

*See slide 14

Clients
The Puppet Master

- Puppet
  - widespread industry adoption, active development
  - full featured, highly flexible
  - focus on consistency and idempotency*
  - gentler learning curve
  - In production:
    - 25 complex servers in Point 1
    - entire Testbed
  - Complements Quattor
- Migrate to Puppet and SLC6
  - our manifest code base has grown to ~15000 LOC
- Puppet is being adopted by CERN IT and the other experiments

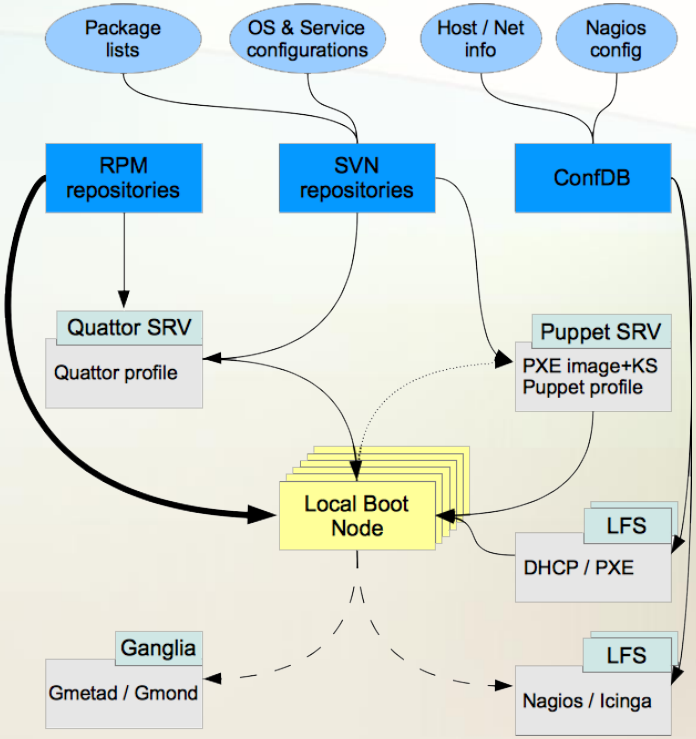

*unchanged in value following multiplication by itself;

Clients
Look Mom ... No disk

- The more components one has in a system, the greater the risk of failure. So… reduce any components that are not “needed”
- In ATLAS, extensive use of PCs with no operating system on disk
- “NetBooted” via PXE
- Advantages:
  - ease of maintenance
  - reproducibility on a large scale
  - reduced "installation" times
- Disadvantages:
  - requires ad-hoc development and support
  - not suitable for running servers

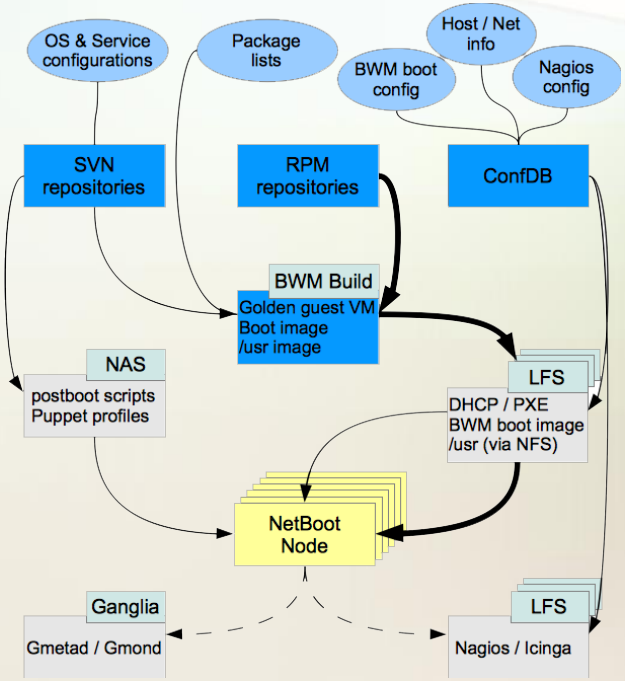

Clients
Netbooted Nodes

- ~ 2350 Nodes booted Scientific Linux CERN5 OS via PXE
- 75 Local File Servers (LFS), provide DHCP, PXE and TFTP for booting, `/usr` read-only directory via NFS
- Configuration of DHCP, PXE and boot parameters are provided by our in house built ConfDB*
- Boot With Me Tool (BWM)
  - generates PXE boot images (kernel + RAMdisk root) and `/usr`
  - uses reference SLC5 VM image as source
- BWM post-boot script system
  - hierarchy of shell scripts, configures servers, disk, and NFS mounts
  - store on a central Network-Attached Storage, executed by the client
- Subversion is used to keep track of image configurations and post-boot scripts

*See slide 14

Clients
Puppet again...

- Problems with BWM post-boot
  - no permanent changes to a working system
  - needs to be rebooted to apply configuration changes
  - bash scripts not easily maintainable
- Move net-booted to Puppet as well
  - manifests hosted on central NFS
  - **no** server
  - no daemon, run via cron job
  - no need to reboot node for changes to take affect
- Minor problems on low resource machines and TDAQ /ROS drivers

Tools
Configuration Database

- Management of large number of NetBooted nodes is far from trivial
- ConfDB was developed
  - python backend, using a MySQL database
  - web based GUI
- Keeps record of the base system configurations
- Greatly help to speed up routine tasks:
  - registering clients, with data extracted from CERN's network DataBase (LanDB)
  - configuring kernel boot options
  - client to LFS assignment and migration
  - deploying DHPC and NAGIOS configurations
  - issuing IPMI and system commands on multiple nodes in parallel

Tools
ConfDB UI

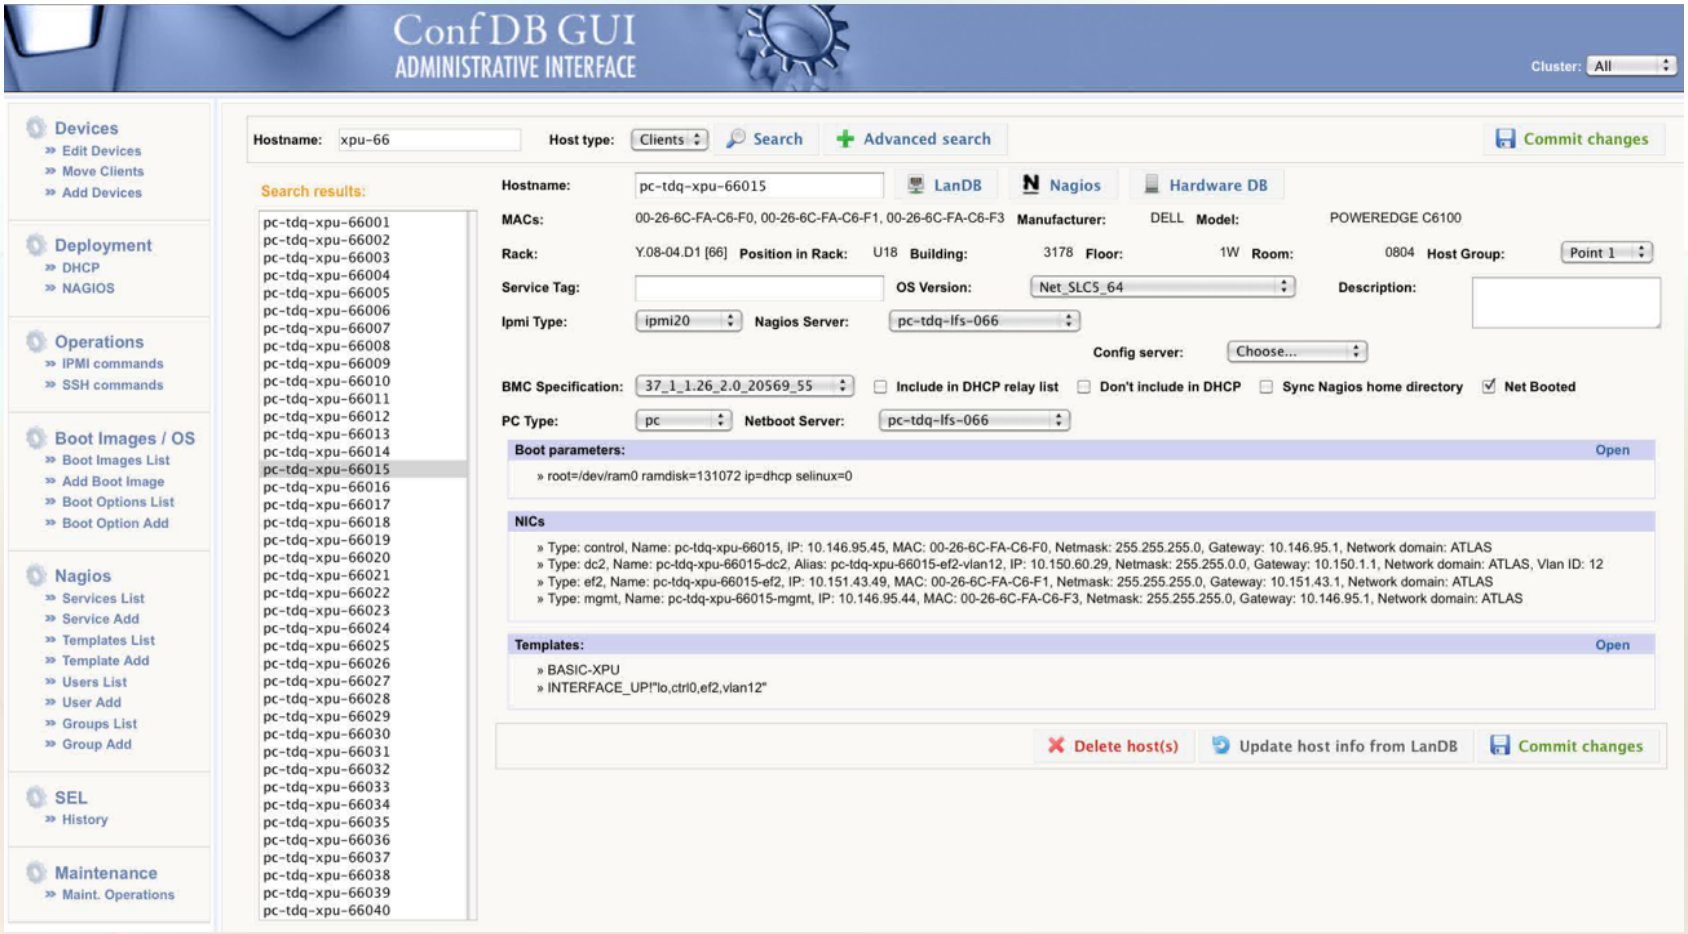

Tools
OS Repository Management

- Simple custom-made system for managing "time-frozen" snapshots of CERN package repositories
- Sufficient functionality through 2012, controlled upgrades, (theoretical) rollback capability
- We want more flexible functionality and easier management
  - partial upgrades, client status reporting etc.
  - may adopt a third-party open source tool, e.g. Pulp
- Rollback/versionlock are possible in principle but not easy: Puppet +yum does not offer the same detailed control as Quattor+SPMA

Monitoring
Nagios

- Nagios has been used since 2007
- Primarily monitors:
  - the health status of the OS
  - the hardware, selected services and network components
  - provides alerting for critical events
- Separate Nagios server instance on each of the LFS nodes
- Feeding data to a central RRD* storage and to a single MySQL Cluster
- A web-based interface was developed in-house
- Significant development effort into hardware status', via IPMI
- Configurations integrated in ConfDB
- Over 5 years, machines monitored by NAGIOS increased from ~1500 to ~3000 in Point 1

** round-robin database*

Monitoring

GROUPS	TOTAL	ONLINE	OFFLINE	BROKEN	RESERVED
⊞ Gateways	6	6	0	0	0
⊞ WebServers	2	2	0	0	0
⊞ FileServers	2	2	0	0	0
⊞ DNS	2	2	0	0	0
⊞ CFS	1	1	0	0	0
⊞ LDAP	4	4	0	0	0
⊞ MYSQL	5	5	0	0	0
⊞ VH	2	2	0	0	0
⊞ ACR	128	119	0	0	9
⊞ SCR	49	33	16	0	0
⊟ TDQ	2186	2168	2	9	7
⊞ LFS	74	73	0	0	1
⊞ ONL	33	33	0	0	0
⊞ AMS	7	7	0	0	0
⊞ MON	32	32	0	0	0
⊞ GMON	6	5	0	1	0
⊞ ROS	157	156	1	0	0
⊞ SFI	48	48	0	0	0
⊞ SFO	9	9	0	0	0
⊞ DC	12	12	0	0	0
⊞ L2SV	8	8	0	0	0
⊞ XPU	1208	1196	0	8	4
⊞ EF	448	448	0	0	0
⊞ PRESERIES	131	129	1	0	1
⊞ RMON SRVs	3	3	0	0	0
⊞ NET-MON	7	7	0	0	0
⊞ SYS	3	2	0	0	1
⊞ SBC	161	153	5	2	1
⊞ PUB	14	12	2	0	0
⊞ DCS	102	98	4	0	0
⊞ MU-CALSRV	2	2	0	0	0
⊞ SWITCH	100	100	0	0	0
⊞ OTHERS	156	132	21	0	3
TOTAL	2926	2841	50	11	20

Monitoring
Ganglia

- 2011: we introduced Ganglia for performance monitoring and trending
- We are currently evaluating on a smaller scale system
- Used primarily for special purpose nodes
  - local boot nodes
  - monitoring, Online and Output (SFO) sub-farms
- Single central server
- RRD* caching daemon used for I/O performance
- Ganglia is used by many Grid facilities and will be by LHCb as well
- Integration with Icinga

** round-robin database*

ICINGA and Gearman

Monitoring

- Icinga+Gearman provide active checks with distributed scheduling
- Alerting
- Icinga+Gearman adopted by CMS and LHCb
- Icinga can reuse Nagios plugins, and much of Nagios configuration

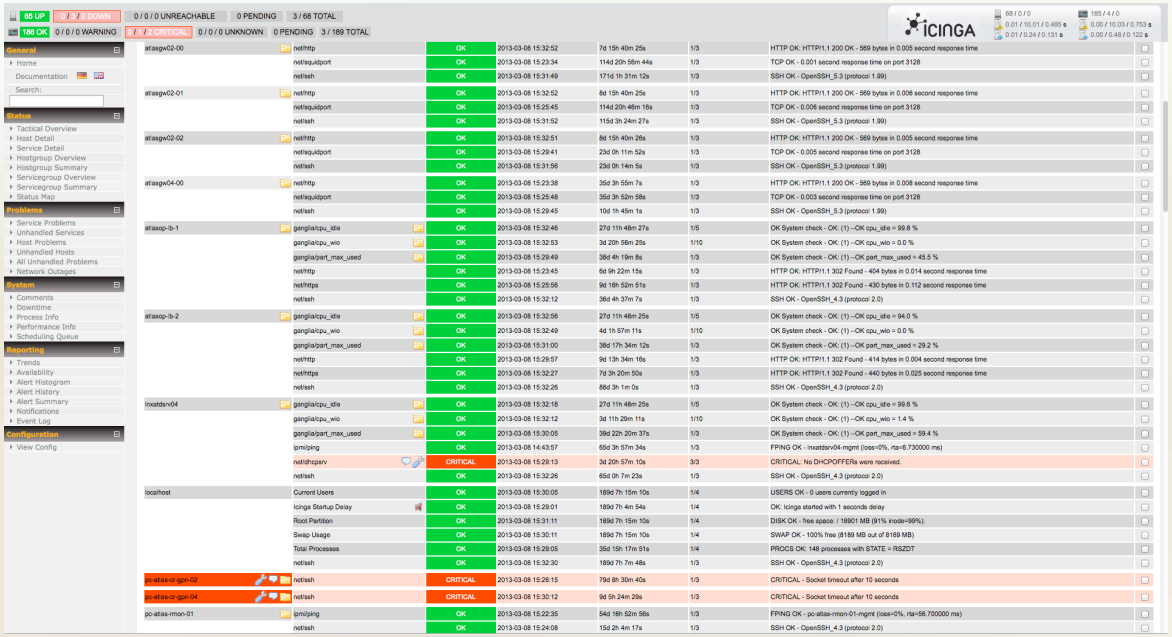

Security
Remote Access Subsystems

- Security at Point 1 is of utmost importance
- Access to the ATLAS Technical and Control Network (ATCN) is highly restricted
- Only allowed via one of the following gateway systems:
  - ATLAS Point 1, allowing expert users access to their restricted machines via SSH or SCP protocols
  - ATLAS Remote Monitoring System provides the graphical terminal services required for organising the remote participation in the ATLAS sub-detector monitoring shifts
  - ATLAS DCS Windows Terminal Servers, allowing experts access to the Detector Control System
- Host and network based accounting, security monitoring and intrusion prevention systems are configured on all the gateways

Security
Role Based Access Control

- Own user database in the form of an LDAP server based on OpenLDAP software
- Standalone but for consistency it is synchronised with CERN IT
- Slave Windows Domain Controller using the CERN NICE credentials.
- Local service accounts, authentication (passwords) in LDAP
- User based authentication, NOT group based authentication
- Role Based Access Control (RBAC) authorisation system
- ~360 unique roles in a hierarchical structure and ~4250 Users
- However..
  - shifter and expert roles automatically enabled/disabled for the person on shift / on call
  - time-limited remote access ONLY on case-by-case authorisation by Shift Leader

Virtualisation

- Adopted for specific services, not planning to use for the full scale farm
- Machines currently running as VM's
  - Gateways
  - Few Windows services
  - Core Nagios servers
  - Public Nodes
  - Domain Controllers
  - Development Web servers
  - Puppet & Quattor
- Xen on Gateways (Bastion Hosts) KVM with SLC6 for everything else
- Currently :
  - 35 Virtual machines in Point 1
  - 11 in GPN
- Expected:
  - Additional ~100 DCS
  - ~1500 transient VM instances for Sim@P1, on netbooted nodes

Sim @ P1

- HLT Farm would be mostly idle during Long shutdown 1 (LS1)
- Use it for ATLAS simulations jobs considerable fraction (>10%) of ATLAS Grid
- Support from BNL during the setup phases
- Support from NetAdmin for dedicated network
- Preserve security of ATCN by isolating VMs and VLANs
- Tests are ongoing in TDAQ Lab4 Test Bed

#	Type	CPU Cores: non-HT / HT	Memory	Local disk
341	Dell PE 1950	8 / 8 (*)	16 GB	80 GB
320	Dell PE 6100	8 / 16	24 GB	250 GB
832	Dell PE 6100	12 / 24	24 GB	250 GB
1493	**Total**	**15272 / 27816**	**33 TB**	**315 TB**

(*) HT not enabled

Sim @ P1
The Plan

- *OpenStack,* as an overlay infrastructure
  - provides necessary management of VM resources
  - support & control of physical hosts remain with TDAQ
  - delegate VM Farm support to Offline Operations
- Easy to quickly switch from HLT ↔ GRID
  - i.e.: during LS1: monthly full-scale test of TDAQ sw upgrade
- BNL, CERN IT, CMS
  - sharing experiences
  - support if needed
  - ATLAS is already, successfully using BNL cloud resources

Conclusion

- 3 years of LHC run, points to a good, linear scalability of the LFS-based architecture, OK overall, but to be kept under control
- Netboot limited functionality, great flexibility in OS migration
- Localboot management is improving with puppet
- The monitoring system will see significant improvements
- Overall, the current architecture is sound and performing well, and the possible changes would increase complexity and not give drastic improvements
- We have decided to focus on trying to simplify and streamline the system where possible, making it more robust and maintainable
- The experience in cloud and virtualisation from the Simulations in Point1 project can be useful for future evolutions

BACKUP

Sim @ P1

SDX1@P1: Full Scale System