Sim@P1

- Many reports in the last weeks:
  - ✓ 4th July ATLAS DAQ/HLT Software and Operations (https://indico.cern.ch/getFile.py/access?contribId=3&resId=1&materialId=slides&confId=261227)
    - ! Today's slides are a subsetof the ones presented by Sasha Zaytsev (BNL) just few days ago + few more details

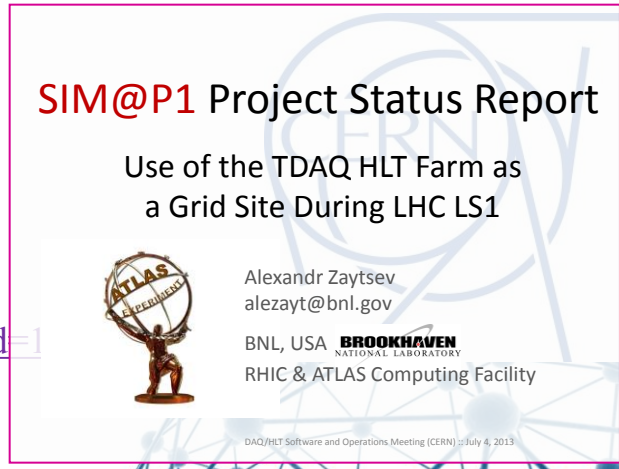

- ✓ 12 June ATLAS SW week (https://indico.cern.ch/getFile.py/access?contribId=15&sessionId=5&resId=1 onfId=210657)
- ✓ 24 May CERN IT Technical Forum (http://indico.cern.ch/getFile.py/access?contribId=1&resId=1&materialId=slides&confId=252610)
- ✓ …

Alessandro Di Girolamo

CERN IT-SDC-OL

People/Teams Closely Involved

Who	Team	Specialization
Sergio Ballestrero Franco Brasolin Cristian Contescu	**ATLAS TDAQ SysAdmins**	**Base Infrastructure**
Silvia Batraneanu M. E. Pozo Astigarraga	**ATLAS TDAQ NetAdmins**	**Network Infrastructure**
Alexandr Zaytsev	**RHIC & ATLAS Computing Facility (BNL)**	**Virtualization/Cloud Infrastructure (OpenStack)**
Alessandro Di Girolamo	**CERN IT SDC**	**Grid Infrastructure**

Two more teams are expected to join as the project enters the production phase:

- ✓ **SIM@P1 operators** in charge of VM contextualization and lifecycle maintenance (activity is already approved by ATLAS Computing management; organizational discussions are in progress)
- ✓ **SIM@P1 shifters** in charge of continuous workflow monitoring and providing the first level of support by supplying feedback for other teams

TDAQ HLT Point1 Infrastructure

Compute nodes

#	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

(*) No HT support

16.5k used by SIM@P1 already (+1.3k pending)

Network

56 Gbps available at the moment

Type	Before the upgrade	After the upgrade (current status)
P1 ↔ Castor	20 Gbps	80 Gbps (20 Gbps reserved for SIM@P1)

OpenStack Virtualization Platform

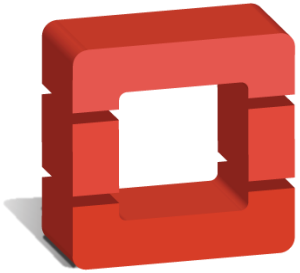

openstack™
CLOUD SOFTWARE

Keystone
Nova
Glance
Horizon

Swift
Ceilometer
Quantum
Heat

Series	Status	Releases	Date
Havana	Under development	Due	Oct 17, 2013
Grizzly	Current stable release, security-supported	2013.1	Apr 4, 2013
		2013.1.1	May 9, 2013
		2013.1.2	Jun 6, 2013
Folsom	Security-supported	2012.2	Sep 27, 2012
		2012.2.1	Nov 29, 2012
		2012.2.2	Dec 13, 2012
		2012.2.3	Jan 31, 2013
		2012.2.4	Apr 11, 2013
Essex	EOL	2012.1	Apr 5, 2012
		2012.1.1	Jun 22, 2012
		2012.1.2	Aug 10, 2012
		2012.1.3	Oct 12, 2012
Diablo	EOL	2011.3	Sep 22, 2011
		2011.3.1	Jan 19, 2012
Cactus	Deprecated	2011.2	Apr 15, 2011
Bexar	Deprecated	2011.1	Feb 3, 2011
Austin	Deprecated	2010.1	Oct 21, 2010

Currently used for SIM@P1

OpenStack Virtualization Platform
Components Used for SIM@P1

- Keystone
  - Central authentication and the service layout management
- Glance
  - VM base image storage and management
  - Central image distribution for Nova (must be assisted with image pre-distribution mechanism for SIM@P1)
- Nova
  - Central operations controller for hypervisors and VMs
  - CLI tools
  - VM scheduler
  - Compute node client
  - Virtual networking (limited use for SIM@P1)
- Horizon
  - WebUI for complete Openstack infrastructure/project/VM control (limited use for SIM@P1), additional CLI tools required for controlling large group of VMs for SIM@P1

Choosing Base VM Images: CERNVM Project

CernVM Downloads

[Software distribution method: CVMFS]

	VirtualBox	vmware	Xen	KVM	Hyper-V Server
2.6.0 Release notes	Desktop x86_64 x86 BOINC x86 Basic x86_64 x86	Desktop x86_64 x86 Basic x86_64 x86	Head Node x86_64 Batch Node x86_64 x86 Basic x86_64 x86	Batch Node x86 x86_64 Basic x86_64 x86 Desktop x86_64 x86 Head Node x86_64 x86	Desktop x86_64 x86 Basic x86_64 x86
2.5.3 Release notes	Desktop x86_64 x86 BOINC x86 Basic x86_64 x86	Desktop x86_64 x86 Basic x86_64 x86	Head Node x86 x86_64 Batch Node x86_64 x86 Basic x86_64 x86	Basic x86 x86_64 Head Node x86 x86_64 Batch Node x86_64 x86 Desktop x86 x86_64	Desktop x86_64 x86 Basic x86_64 x86

SLC 5.8 x86_64 based
SIM@P1 hypervisor: KVM

Post-boot contextualization method: script injected into the base image (Puppet in the near future)

New Network Infrastructure for SIM@P1

- New dedicated 1 Gbps physical network connecting the P1 rack data switches to the Castor router

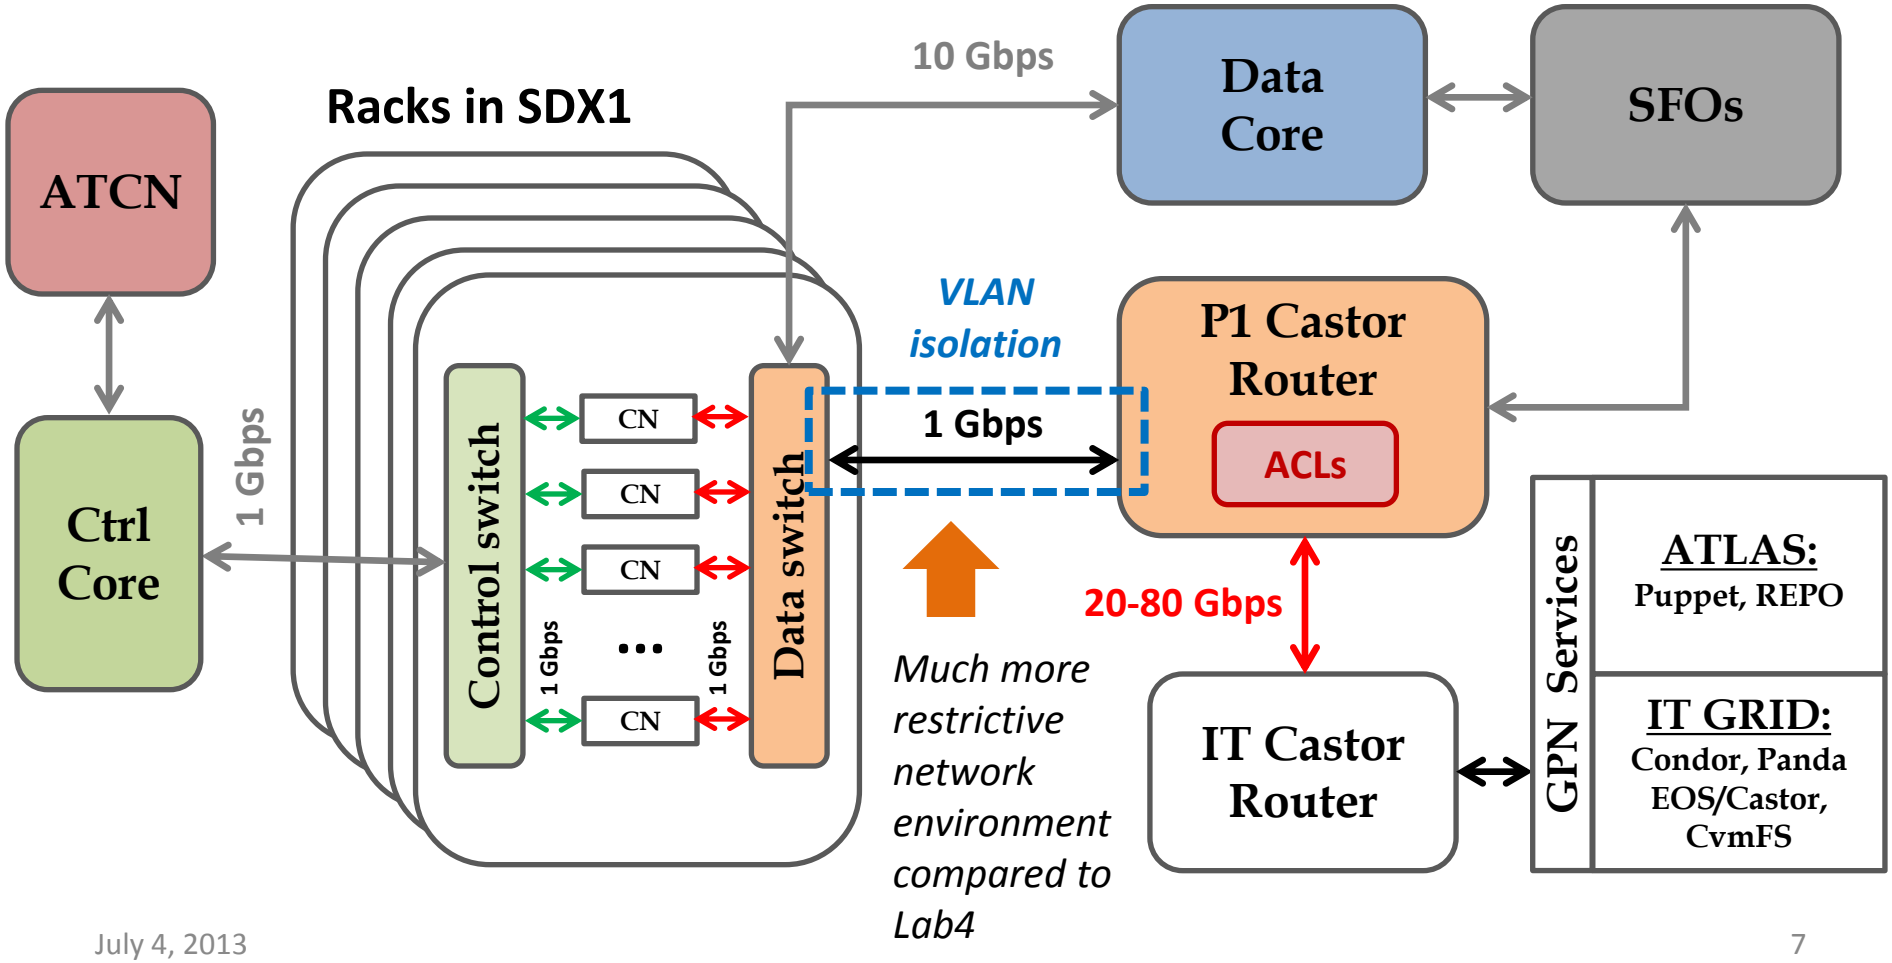

Full Scale Deployment in Point1 (SDX1)

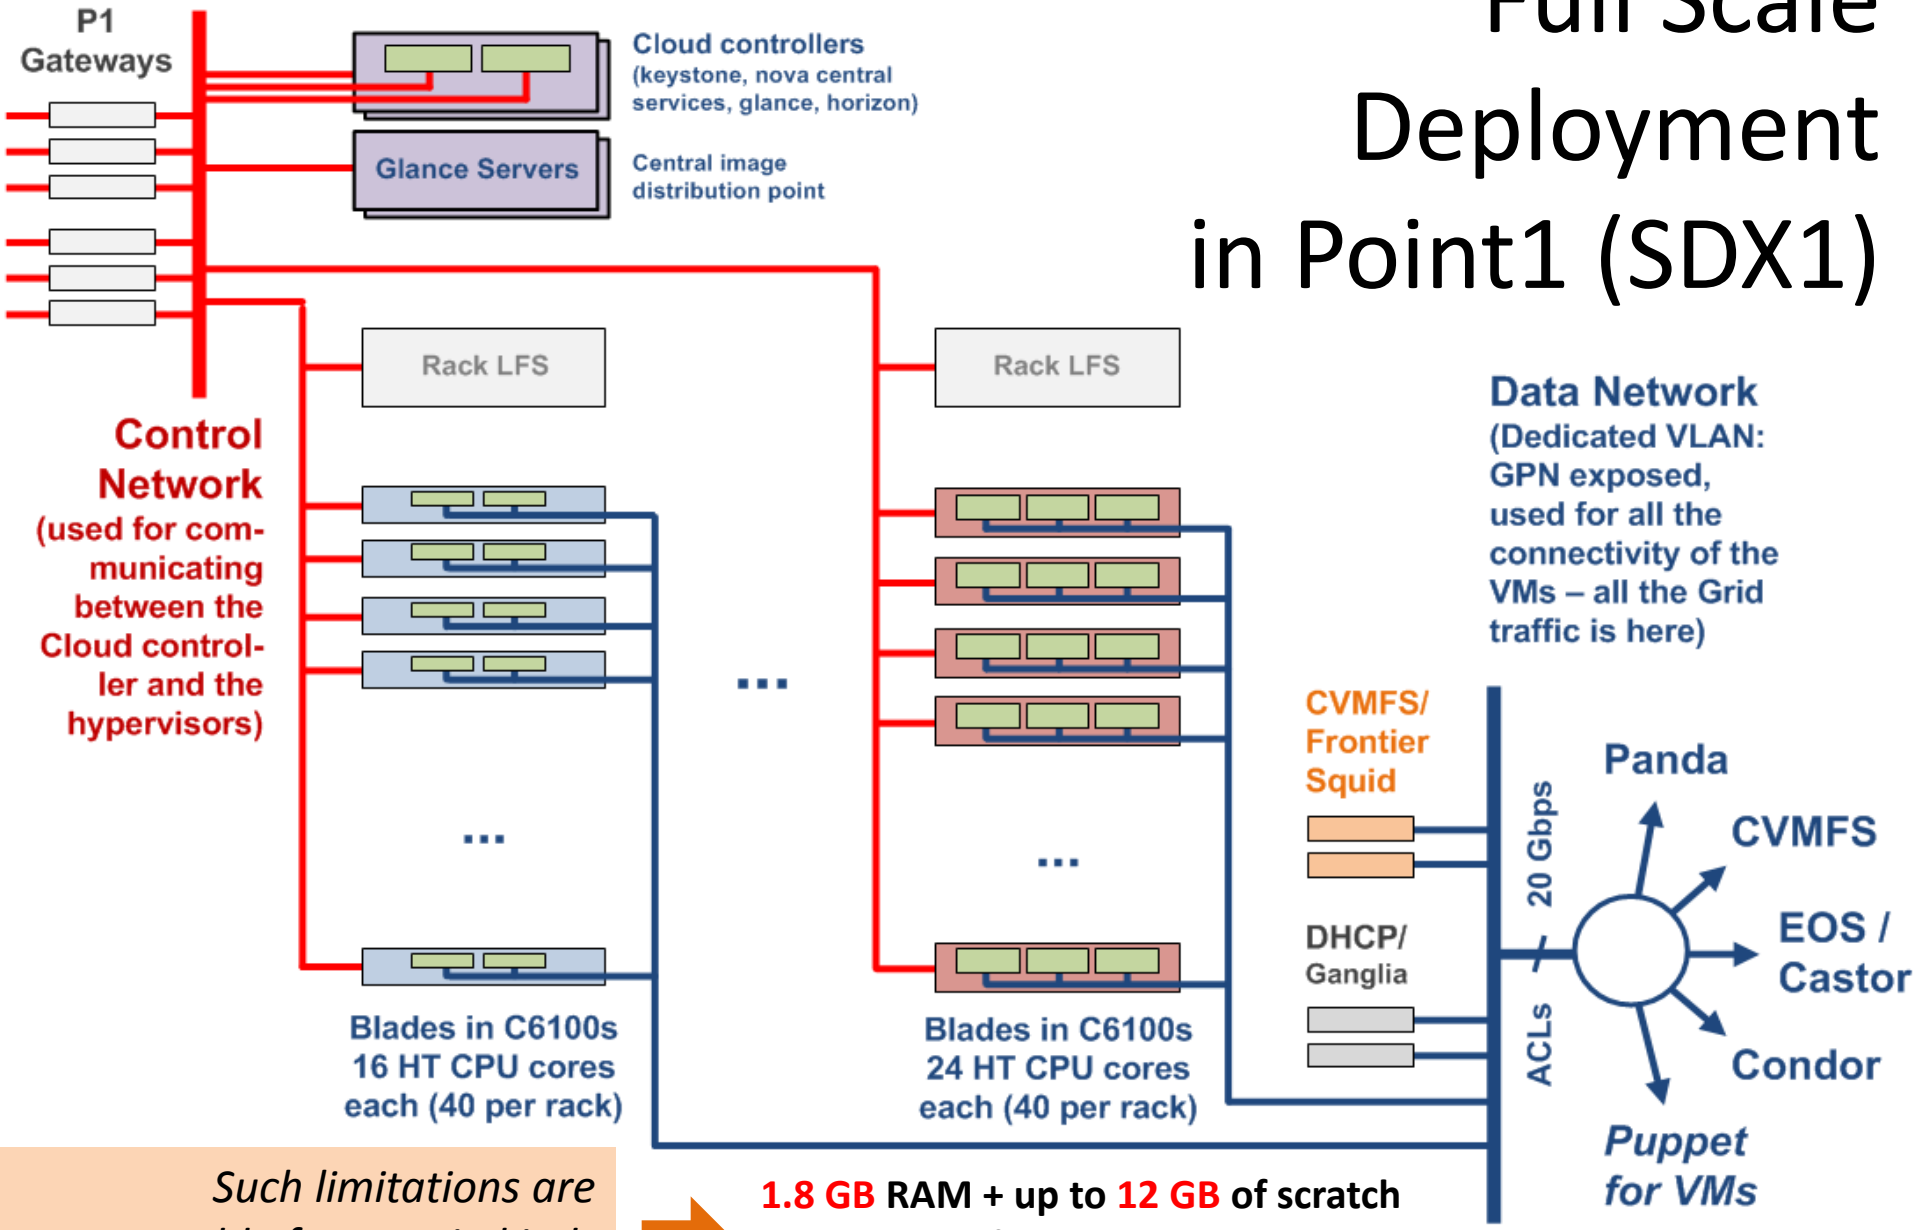

Such limitations are acceptable for certain kinds of ATLAS MC production jobs

1.8 GB RAM + up to 12 GB of scratch + 1.3 Mbps of ext. connectivity per non-HT CPU core (0.5x in full HT mode, 0.75x in half-HT mode)

Cloud Infrastructure of Point1 (SDX1)

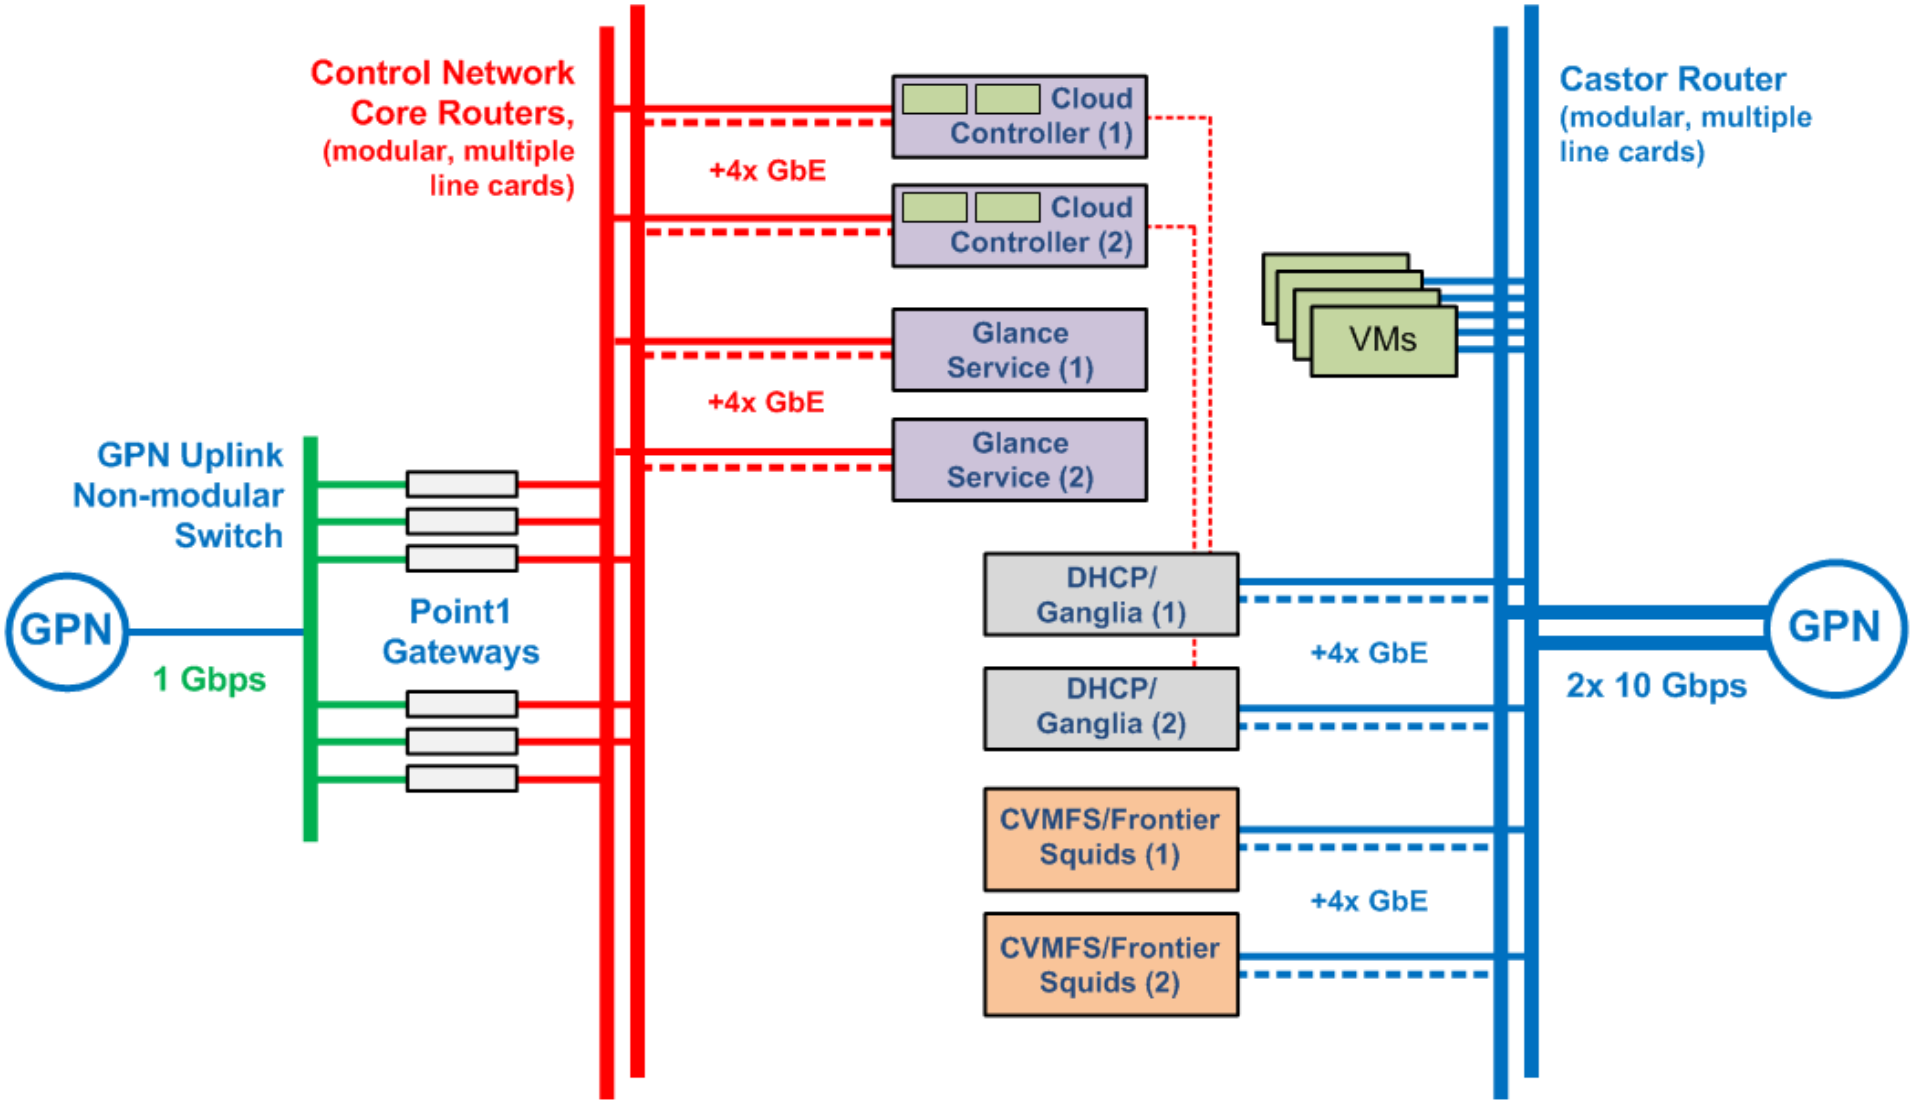

Getting to Production With the Entire Computing Capacity of SDX1

- Prospected communication layout between the teams involved in the operation and support of the production SIM@P1 infrastructure:

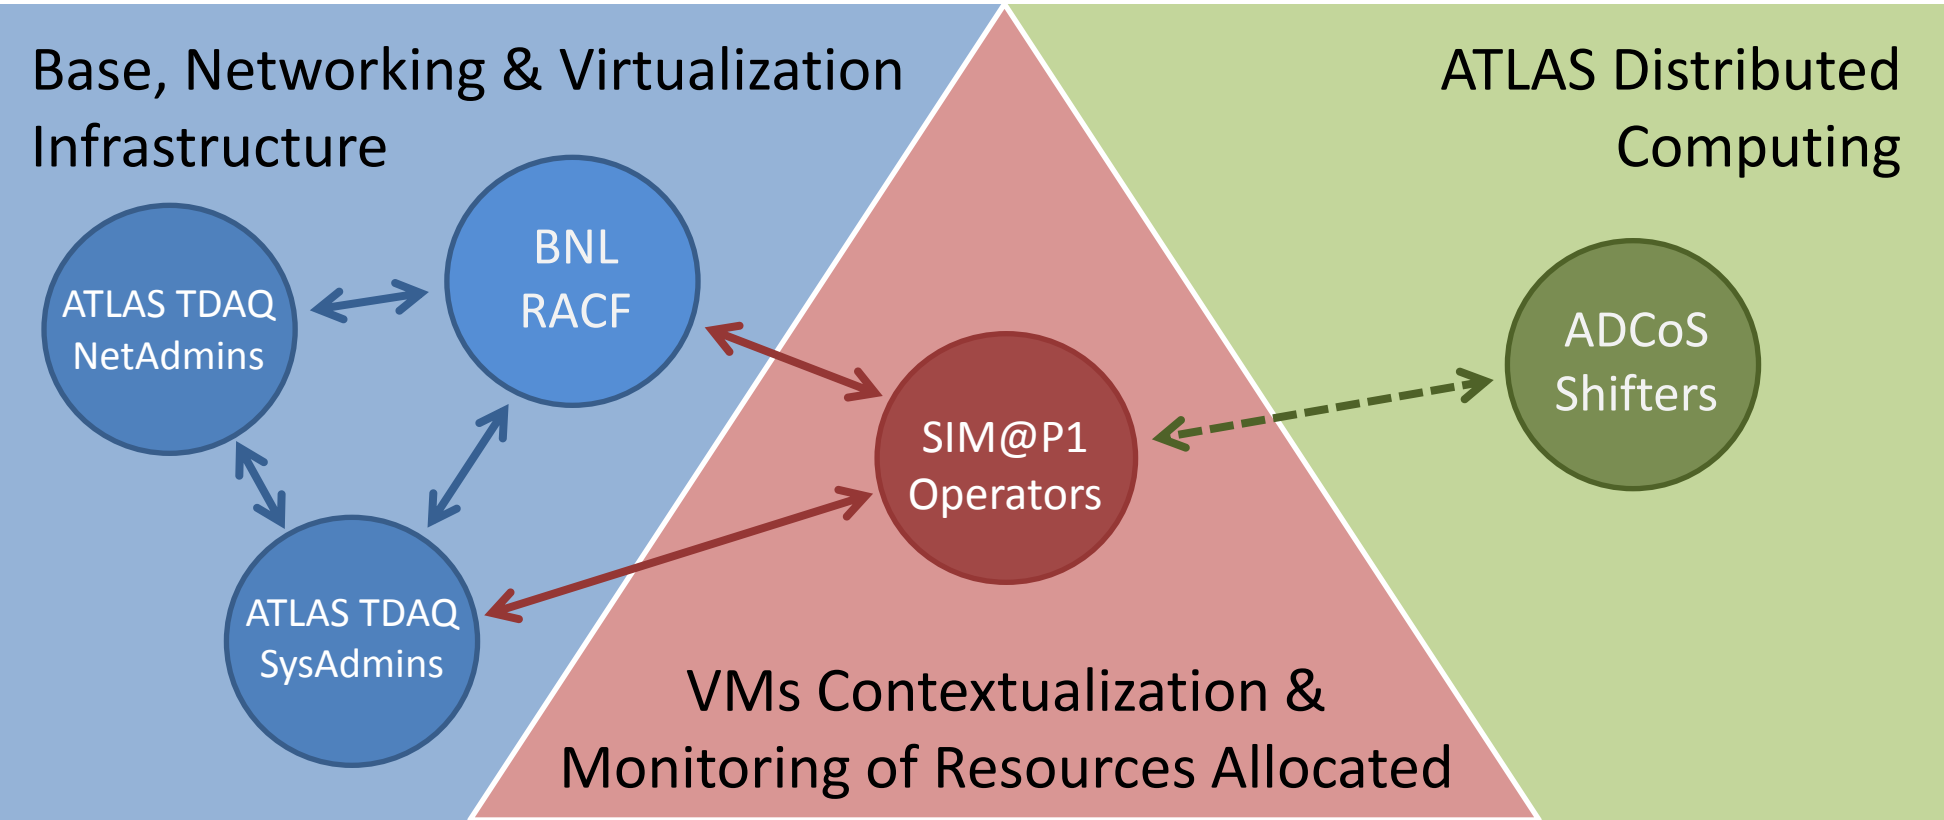

Full Scale Deployment in Point1 (SDX1) – 1/4

- Up to May 18, 2013: Physical deployment of Sim@P1 network is finished; LanDB sets/XPU interfaces are registered; ACLs control mechanisms are implemented for Castor routers
- June 05, 2013: Physical installation and cabling of the centralized Cloud infrastructure servers is finished in Point1
- June 07, 2013: The first rack of C6100 based XPUs (768 HT CPU cores) is configured for use as Nova Compute Nodes
- June 10, 2013: Configuration of the centralized Cloud infrastructure servers is finished (via Puppet)
- June 11, 2013: 96 fully contextualized VMs are running in Point1 under OpenStack control
- June 12, 2013: We have the first MC jobs running in Point1 under CERN-P1 site (HammerCloud stress test with input DC pattern mc12_8TeV*evgen.EVNT*)

Full Scale Deployment in Point1 (SDX1) – 2/4

- June 13, 2013: 2 racks of C6100s are under Sim@P1
  - 1.5k job slots with high density (0.9 GB RAM/job slots)
  - Running HammerCloud tests on the level of 1.1k
  - Observing RAM limitations on the VMs
- June 19, 2013: 12 racks of C6100s are under Sim@P1
  - 8k job slots with normal density: 1.4 GB RAM/job slots
  - Running HammerCloud tests and observing their instabilities on this new scale
  - 8 hour long EOS outage while trying to reach saturation on 8k jobs running:
    - internal EOS bug spotted by creating a flow of output file creation requests at 2.8 kHz rate
    - patch is developed by the EOS support team
    - no further problems of such a scale detected ever since
- June 20, 2013: Reaching 7.5k jobs running simultaneously under the HC test for the first time
  - Observing HammerCloud instability on the tests with lifetimes longer than 24 hours
  - Performance tuning for the Nova controllers needed
  - Limitation are observed for ramping up rate for the HC tests on the level of 1k jobs spawning per hour per submission host
  - start using multiple submission hosts in parallel for HC tests
  - Highest SIM@P1 generated spikes observed in the Castor router uplink: 15 Gbps for 5-10 min (width on 0.5 height)

Full Scale Deployment in Point1 (SDX1)

7.5k Job Slots Under HammerCloud Tests

June 19-20, 2013

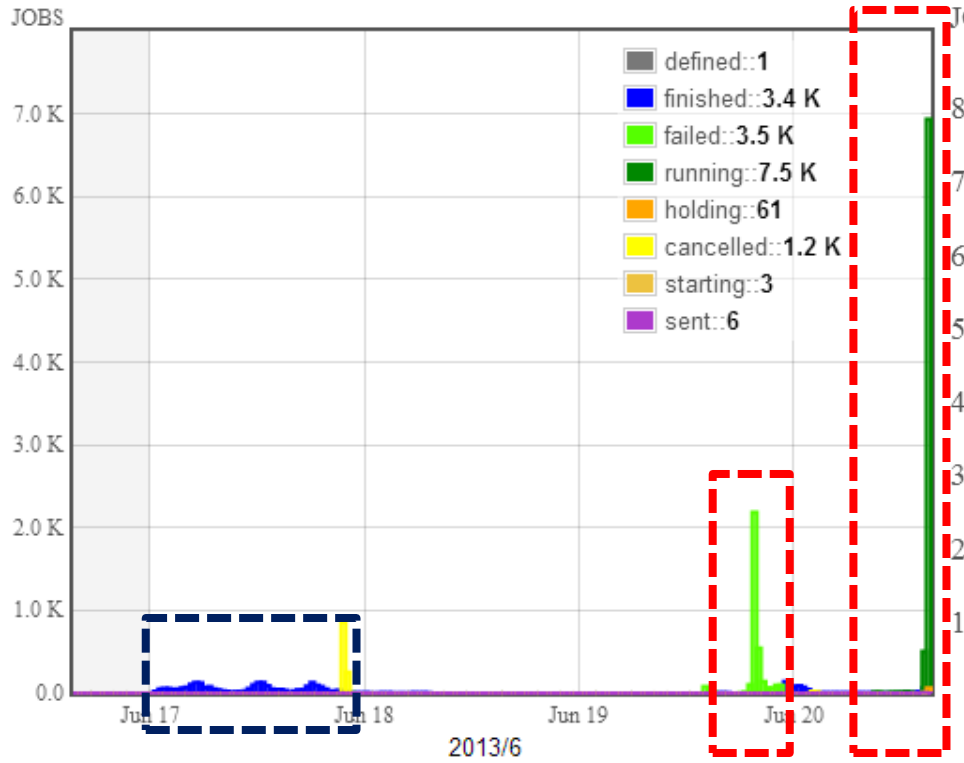

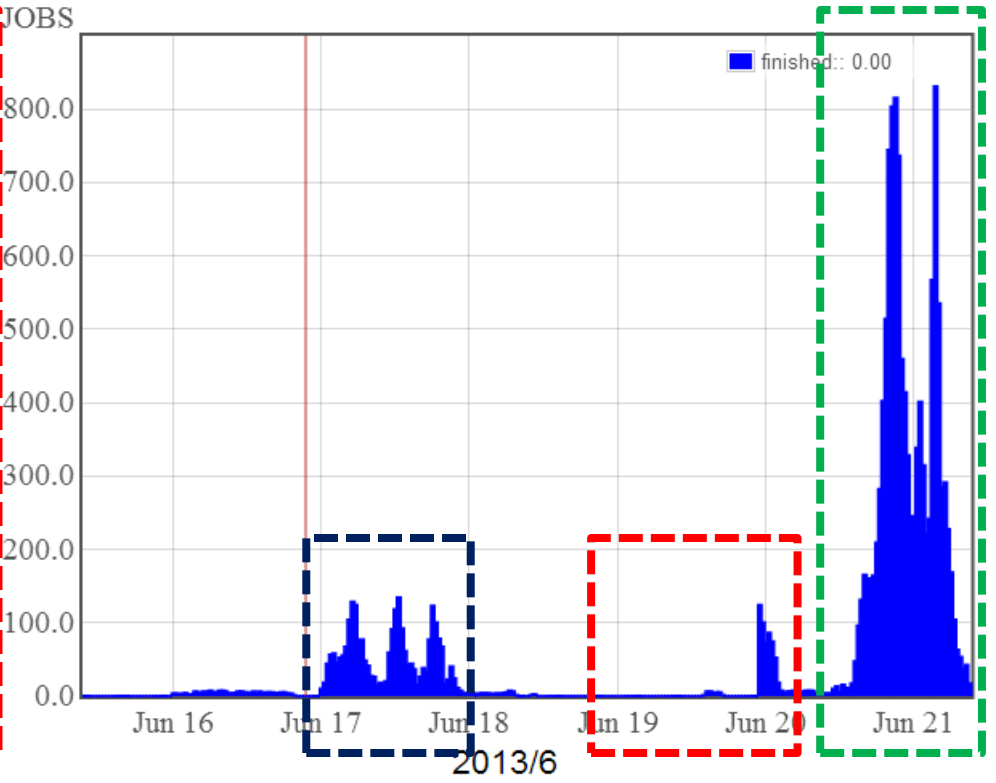

Full Scale Deployment in Point1 (SDX1) – 3/4

- June 22, 2013: Running the HC tests with 8k job slots occupied simultaneously
  - Observe the increase of the overall job failure up to 8% observed during the initial period of the test
  - Failure rate is dropping to 4% after the ramp-up is complete
  - Most of the failures are due to the EOS & LFC connectivity problems
  - Time correlated 1 GbE XPU rack uplink saturation is the primary suspect
- June 24, 2013: Adding compute nodes in SDX1 Level1 row 6
  - Reaching up to 1.2k hypervisors under Openstack control in Point1
  - Deploying the redundant group of CVMFS/Frontier servers in Point1
  - Attempting to scale up to the maximum of 2.1k VMs running (16.5k job slots), reaching the Openstack infrastructure saturation at the level of 1.3k VMs operating under single Nova controller
  - Further performance tuning is performed for the Openstack infrastructure
  - RabbitMQ is clusterized on the Nova controllers
  - The rates of internal messages between the Nova compute nodes and the Nova controller are suppressed
- June 30, 2013: The level of 1.2k hypervisors with 2.1k VMs running (16.5k job slots) is reached
  - The obtained configuration of the centralized Openstack infrastructure should scale up to up to 4-5k VMs running on the existing hardware
  - **Running the full scale production with real ATLAS MC jobs in SDX1**
  - Observing the 10k jobs slot limitation on the level of APF (Condor master)
  - Running real MC production on the level of 10.5k jobs executed in parallel

8k – 10.5k Job Slots Under Real ATLAS MC Production

June 22-30, 2013

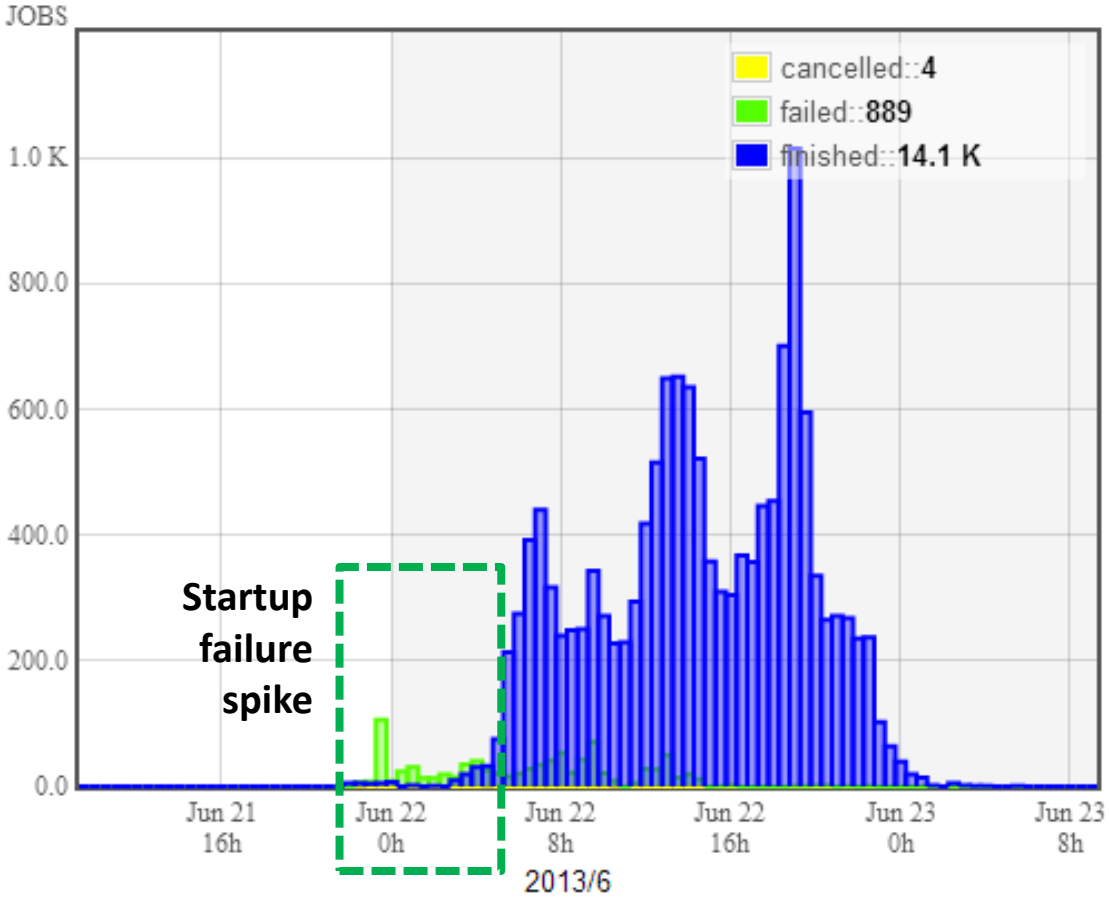

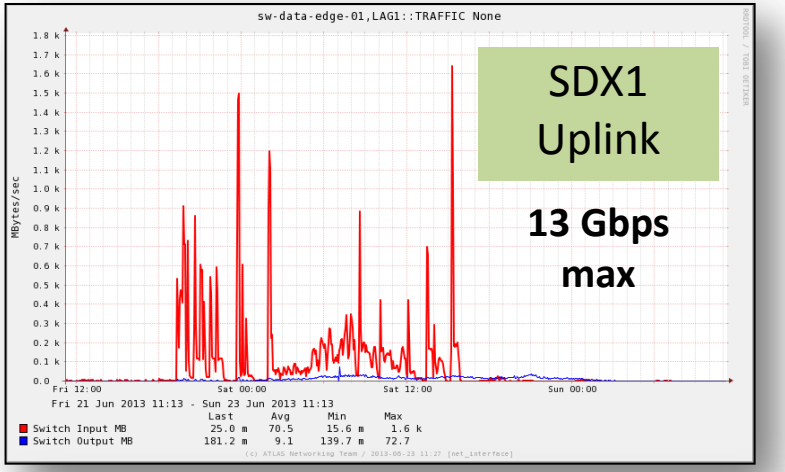

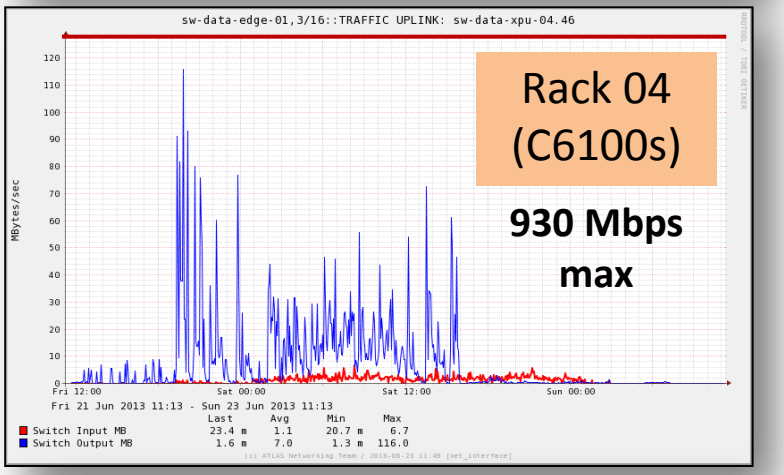

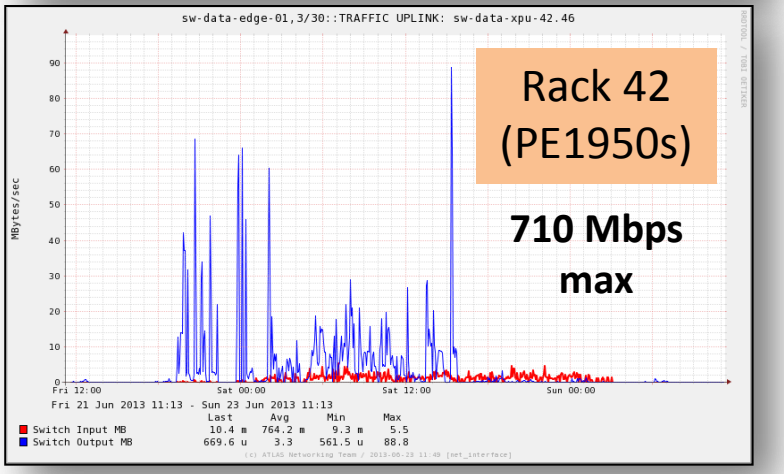

Full Scale Deployment in Point1 (SDX1) – 4/4

- July 1, 2013: Study of the 10k job slot limitation on APF side is on going; the first full scale ramp down test is performed
  - The limit is raised up to 12.5k by tuning Condor on the APF node
  - We need a dedicated Condor instance properly configured to operate on the scale of more than 10k job slots per Condor pool
  - Ramping down from 2.1k VMs to 0 VMs running under a single Nova controller takes 1 hour in a fully graceful mode; 10 minutes is expected for the ungraceful mode based on Puppet that is proposed for production (to be tested yet)
- June 2, 2013: The first full scale ramp up test with the new base VM image distribution is performed
  - A new Qcow2 image is pre-distributed on all the compute nodes in the compressed form prior to spawning the VMs (avoiding 1 GbE link saturation on the Openstack Glance servers)
  - The image unpacking/conversion to RAW format was being performed by Nova for all the instances (required only once the new image is performed)
  - Even with such overhead + 4 concurrent threads of spawning on the Nova controller (avoiding saturation of 1 GbE links on the CVMFS/Frontier servers) ramping up to 2.1k VMs running takes 1h 30 min (including VM boot and contextualization procedures) – still matching the highest rate of spawning of the MC jobs we've seen so far (10.5k jobs per 45 min)
  - Then the base VM image cached on the compute nodes is reused the total ramp up time should drop below 30 minutes (to be tested yet)

Full Scale Deployment in Point1 (SDX1)

12.5/16.5k Job Slots Under Real MC Load

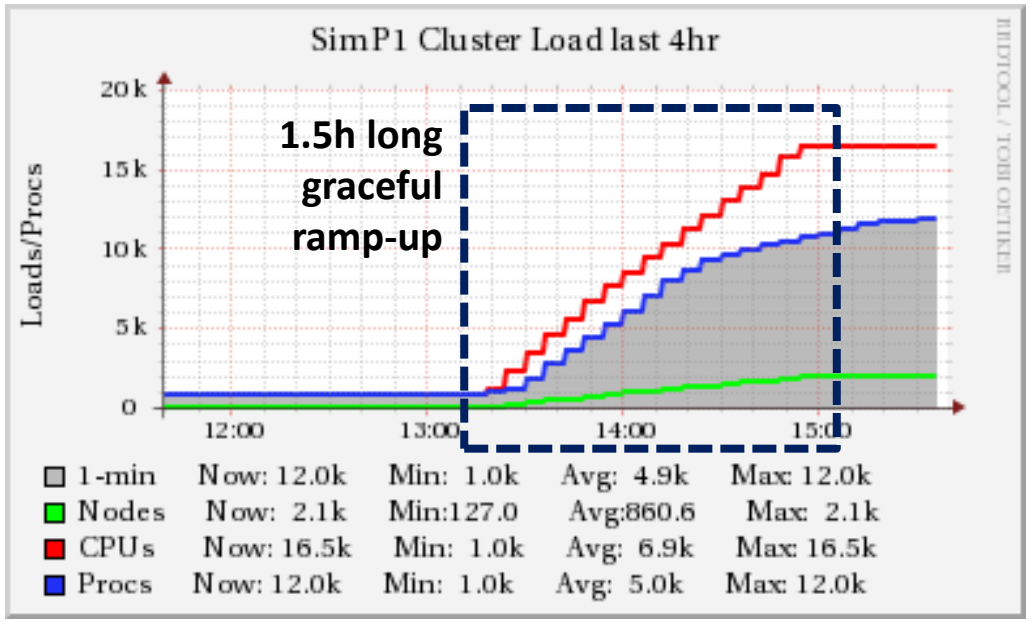

Since July 1, 2013

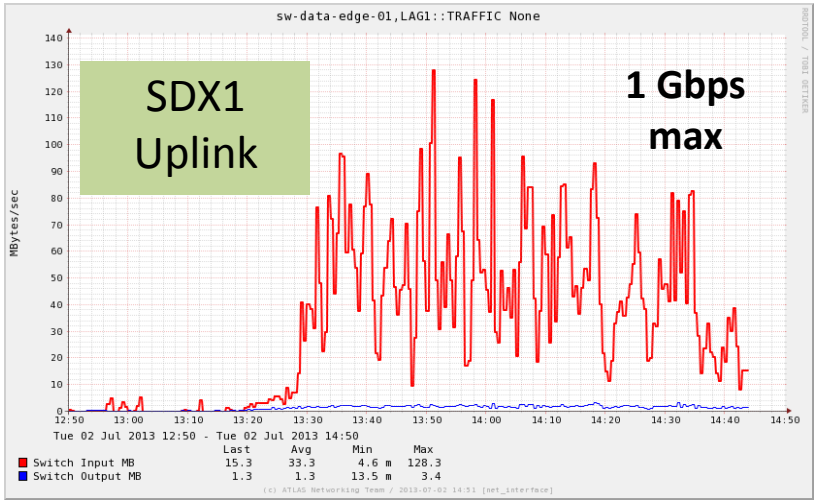

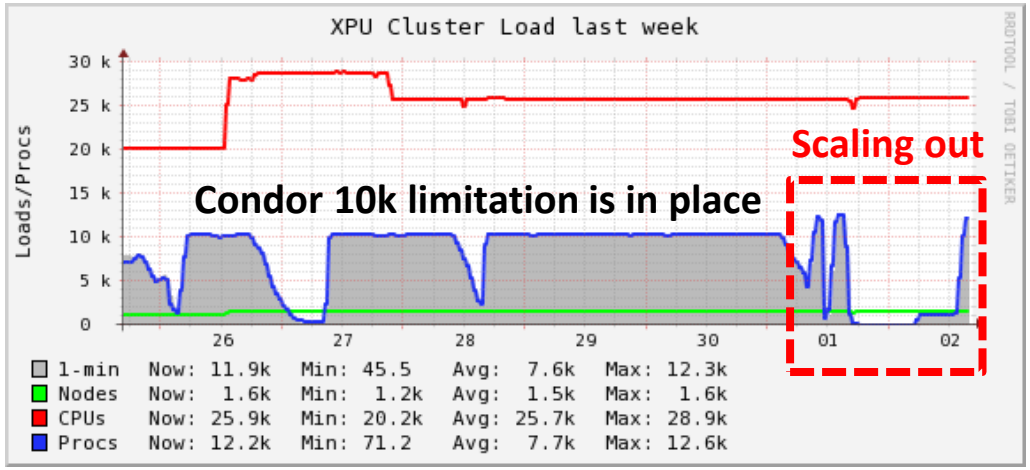

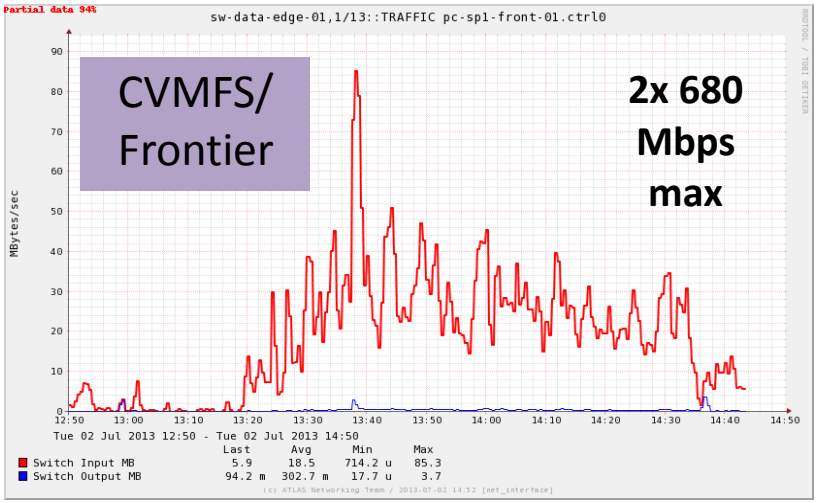

Full Scale Deployment in Point1 (SDX1)

12.5/16.5k Job Slots Under Real MC Load

Since July 1, 2013: Condor Scalability Issues

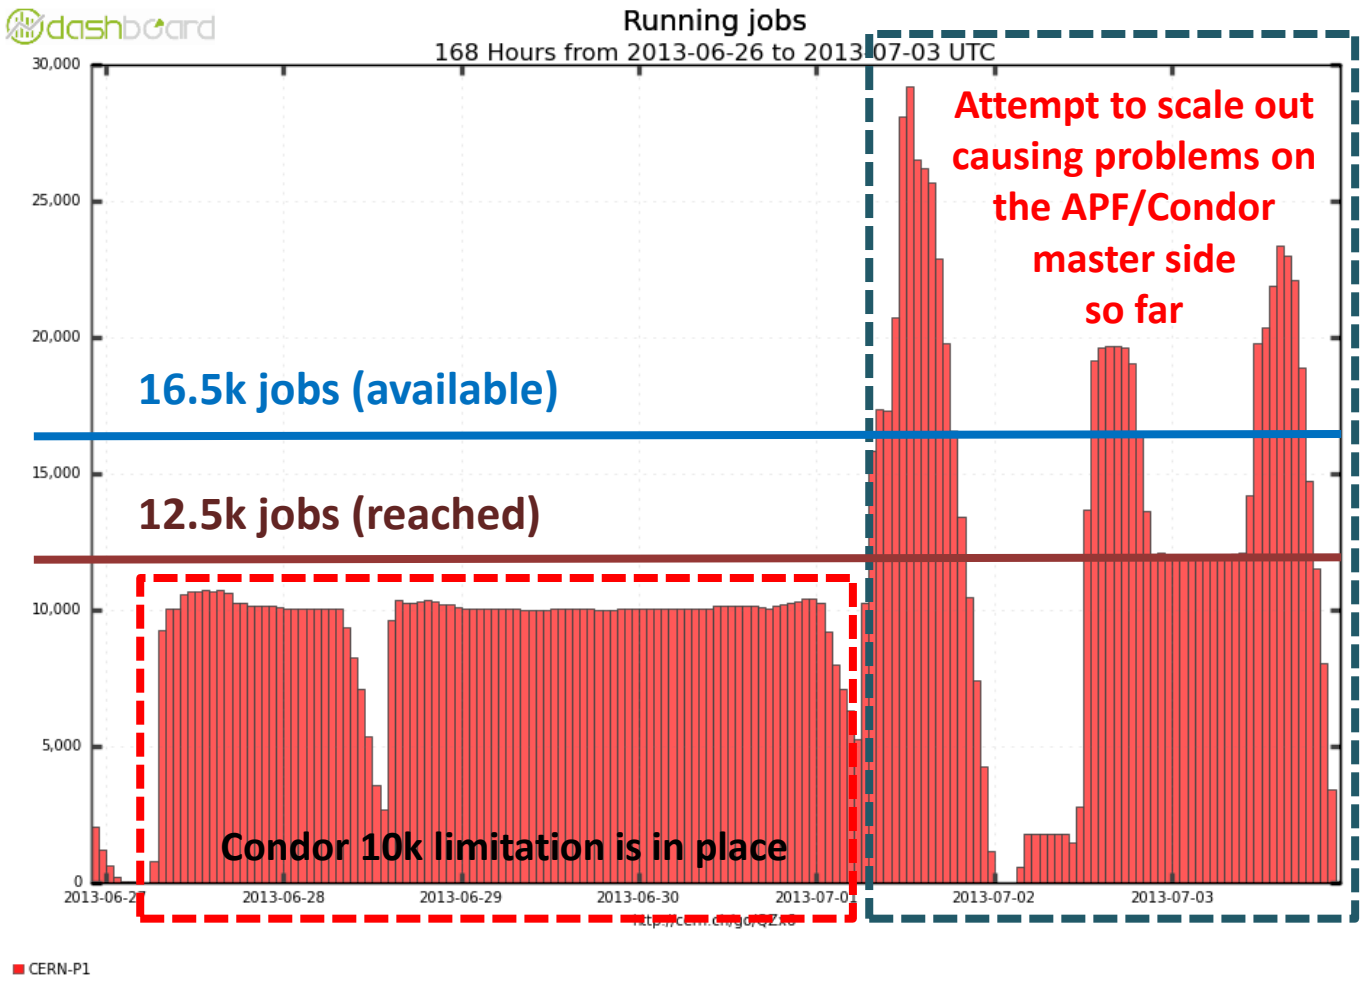

Overall SIM@P1 Project Status

- Functionality and scalability tests of the Cloud infrastructure with the testbed in Lab4 are finished successfully
- CERN-P1 Site is established and operational since Mar 2013
- Deployment of the physical server/networking infrastructure required for the full scale deployment in Point1 is finished
- The full scale deployment of the Openstack infrastructure in Point1 is finished (16.5k jobs slots are available at the moment, +1.3k may become available in the near future, 5.8k are reserved for TDAQ)
- The Sim@P1 operator team has been established and being provided with the necessary tools and access rights
- Sim@P1 successfully processed approximately 50k real MC production over the last week already, though many components are still being tested
- Currently running production on a scale of 12k jobs slots occupied simultaneously – hopefully until the beginning of the upcoming ATLAS TDAQ Run (July 15, 2013)
- Puppet based procedure of switching the racks between the TDAQ and Sim@P1 state is developed by ATLAS TDAQ SysAdmins
- Semi-automatic tools for control of large scale groups of VMs are developed but yet to be declared production ready/to be fully exposed to ATLAS TDAQ SysAdmins

Conclusion

- Sim@P1 project is proceeding very well
  - ✓ ATLAS is able to exploit, when available, the HLT resources
  - ! Few limitations observed and some of the solved, some of them being addressed now
    - E.g. Condor scaling: we do have other options to exploit the full HLT capacity, but for now we prefer to take this opportunity to address those.

? Questions still open

- Streamline VMs lifecycle between different Clouds
  - Ramon + myself wrote down a small doc to discuss this with ATLAS Cloud R&D task force
- Graceful termination of VMs
  - “etc machine features” could be a good approach for us

. . .

- . . .
  - . . .
  - . . .
- . . .
  - . .
  - . . .
  - . .
    - ! . . .