Report from WLCG Workshop 2017: WLCG Network Requirements

GDB - CERN
12th of July 2017

edoardo.martelli@cern.ch

WLCG Network Requirements

The 4 major LHC experiments were asked to provide their network requirements for the coming years, in terms of:

- bandwidth
- special capabilities
- monitoring

ALICE

Recommendations for Tier2s:

- WAN: 100Mbps of WAN per 1000 cores
- LAN to local storage: 20Gbps per 1000 cores
- Better read data locally: CPU efficiency get a 15% penalty per 20ms RTT

ALICE depends on and strongly encourages:

- full implementation of LHCONE at all sites
- plus the associated tools (like PerfSONAR), properly instrumented to be used in the individual Grid frameworks

ATLAS

ATLAS moving to a Nucleus and Satellite Model

Nucleus will store primary data and act as a source for data distribution:

- Storage capacity > 1PB
- Good Network throughput
- Site availability: > 95%

Nucleus	Now	5 year (2022)	10 year (2027)
Storage Capacity (PB)	2	5	12.5
Total CPU (kHS06)	40	100	250
LAN (Gb/s)	40	200	1000
WAN (Gb/s)	20	60	200

Disk-less	Now	5 year (2022)	10 year (2027)
Total CPU (kHS06)	20	50	125
WAN (Gb/s)	4	20	100

ATLAS would like

- better visibility into networks to improve ability to resolve network issues
- to see network traffic data from R&E Networks

CMS

Site types and requirements:
- **Full Tier-2**, many CPUs and large disk capacity: Some 10Gbit/s or 100Gbit/s for both LAN and WAN are advisable for sites with several 1000 cores
- **Disk-rich Tier-2**, more storage than average, perhaps hosting disk for co-located CPU only site: Good WAN connection even more important

CMS schedules typically 10-15% of jobs with remote data access
- Penalty in CPU efficiency: drop of ~10% for remote access
- But large gain in flexibility

Computing model likely to evolve towards scenarios that require fast interconnects

Need to improve CMS transfer system and scheduling system to better exploit network metrics

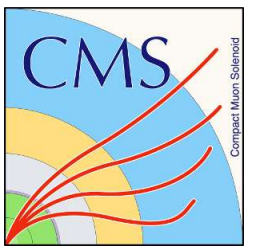

LHCb

- Majority of workflows are simulation jobs without input data
- A certain fraction are user jobs with local access to input data
- Currently small usage of working group productions. Will increase
- In case of input data, always read by default from local site
- Output data always goes to different storage areas on T1 sites

	Job Length	Input Data	Output Data
Monte Carlo Simulation	6 hours	None (start from random seed)	O(500MB), close T1
Monte Carlo Reconstruction	1 – 2 hours	Download O(5 GB), close T1	O(5GB), same T1
"Helper" for Data Processing	12 – 24 hrs	Download O(5-10 GB), close T1	O(5GB), same T1
User Analysis @ T2 ("toy MC")	< 2 hours	None	Close T1 user area
User Analysis @ T2-D	< 2 hours	Protocol access, local storage	Close T1 user area
Working Group Production	< 3 hours	Protocol access, local storage	Close T1 data area

Monitoring, including network monitoring, available from within Dirac:
- for "helper data processing" sites to check the WAN connectivity to a T1 storage
- for data management to check WAN quality

Network throughput WG

WG has established an infrastructure to monitor and measure networks:

- proven record on fixing existing network problems and improving transfer efficiency
- stable production infrastructure

Mid-term evolution topics:

- Network capacities planning
- Network utilization monitoring, both site-level and WAN
- Evolving and integration of monitoring data: new sources, dashboards and network stream
- Network Analytics: Alerting/Notifications and anomaly detection
- SDN Networking demonstrators and testbeds

GEANT

GEANT is investigating ways to increase ability to differentiate, break dependence on vendors, and provide capabilities to automate/orchestrate multi-domain services

Asked the Experiments if interested in participating

From the discussions

Need to maintain, monitor and develop network monitoring;
effort is required to:

- respond to GGUS tickets
- track service and metric status and follow up on down services, mis-configuration, etc.
- evolve the system, implementing new user interfaces, analytics, alerting and new measurements
- coordinate network activities with others

Some interest has been expressed about **network programmability**, especially in light of foreseen commercially driven network developments.

References

Presentations available here:
https://indico.cern.ch/event/609911/sessions/238341/#20170620

Questions?

edoardo.martelli@cern.ch