Computing Platforms R&D

status and perspectives

Fabio HERNANDEZ *on behalf of*

KEK: T. Nakamura, T. Sasaki, W. Takase, H. Matsunaga, G. Iwai, K. Murakami, S. Suzuki, Y. Watase
CC-IN2P3: G. Rahal, R. Rumler, M. Puel, F. Wernli, R. Ferrand, V. Hamar, L. Caillat-Vallet, P-E. Macchi
CENBG: S. Incerti

Joint Workshop FKPPL & FJPPL — Seoul (Korea), May 18th, 2016

Background

- Guiding principle

  to explore technologies of potential interest for the data processing needs of HEP experiments

- Partners

  IN2P3 computing center, KEK computing center

- Method

  2-days-long annual meeting

  in-depth discussions and hands-on sessions about a broad set of topics [2015, 2016]

  remote joint work in-between

- Topics

  *platform for **data collection for operations***

  *continuous monitoring of **network connectivity** between CC-IN2P3 and KEK-CRC*

  *virtualisation (**cloud**) platforms*

  ***procurement** process and policies*

Motivation

- IN2P3 and KEK computing centres share the same mission

  provide data processing services to experiments supported by the research programme of their respective institutes

- They also share several challenges

  significant number of experiments supported, often with conflicting requirements

  increasing demand in computing resources within a context of flat budgets, at best

  manpower not increasing at all or not as fast as the demand for new, highly customised services

Ongoing activities

Data collection for operations

- CC-IN2P3 deployed a platform based on ElasticSearch (data collection) & Kibana (visualisation)

  100 million log entries and 1 billion metrics collected, centrally stored and indexed daily

  open source version lacks access control capabilities

  Kerberos-enabled authentication module deployed during hands-on session in Lyon, based on work performed by W. Takase (KEK-CRC)

  *this allows for **individual**, **personalised dashboard** and restricts control to the data repository*

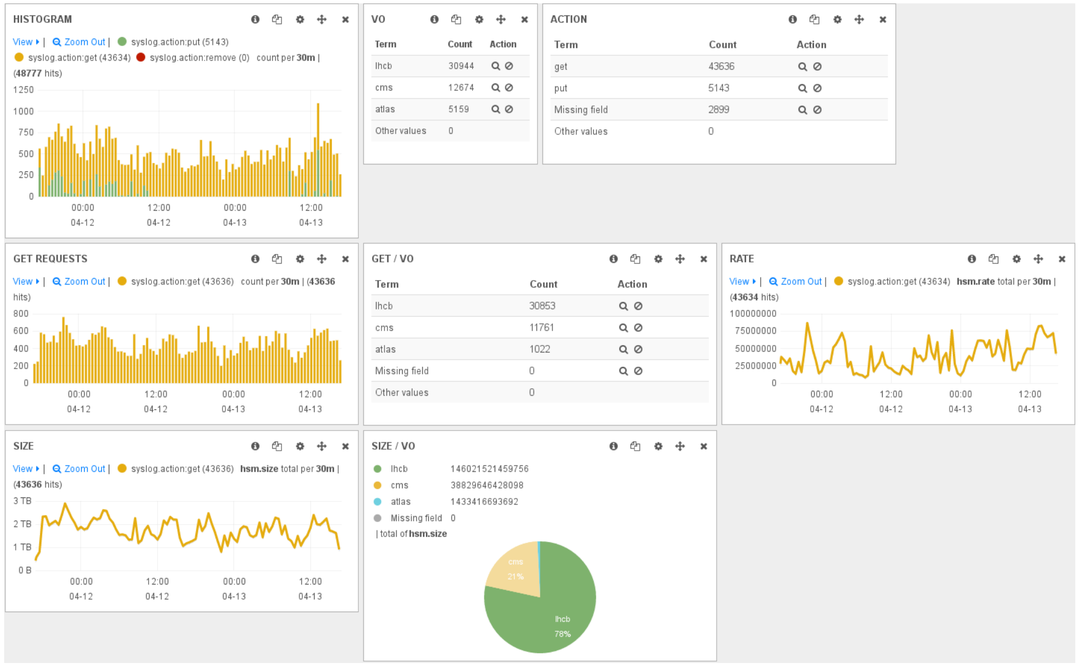

work presented at HEPiX workshop in April, 2016:

"Authorization extension for the secure use of ElasticSearch and Kibana" (W. Takase)

"Monitoring at scale: a needle in the haystack" (F. Wernli)

Network connectivity

- Continuous monitoring of the network connectivity between CC-IN2P3 and KEK-CRC

  upgraded CC-IN2P3's PerfSONAR host to use 10 Gbps network card

  observed low throughput which severely impacts data transfer

- KEK-CRC network configuration to be upgraded this year

  internal network will be reconfigured and new equipment deployed

  KEK connectivity to Europe will also be upgraded: two new 10 Gbps trans-Siberian links to connect KEK to Europe via London (GEANT)

  instabilities expected during the first stages of the deployment phase

  we intend to deploy additional equipment and tools for probing connectivity between CC-IN2P3 and KEK-CRC

Network connectivity (cont.)

Migration to SINET5 (Apr. 2016)

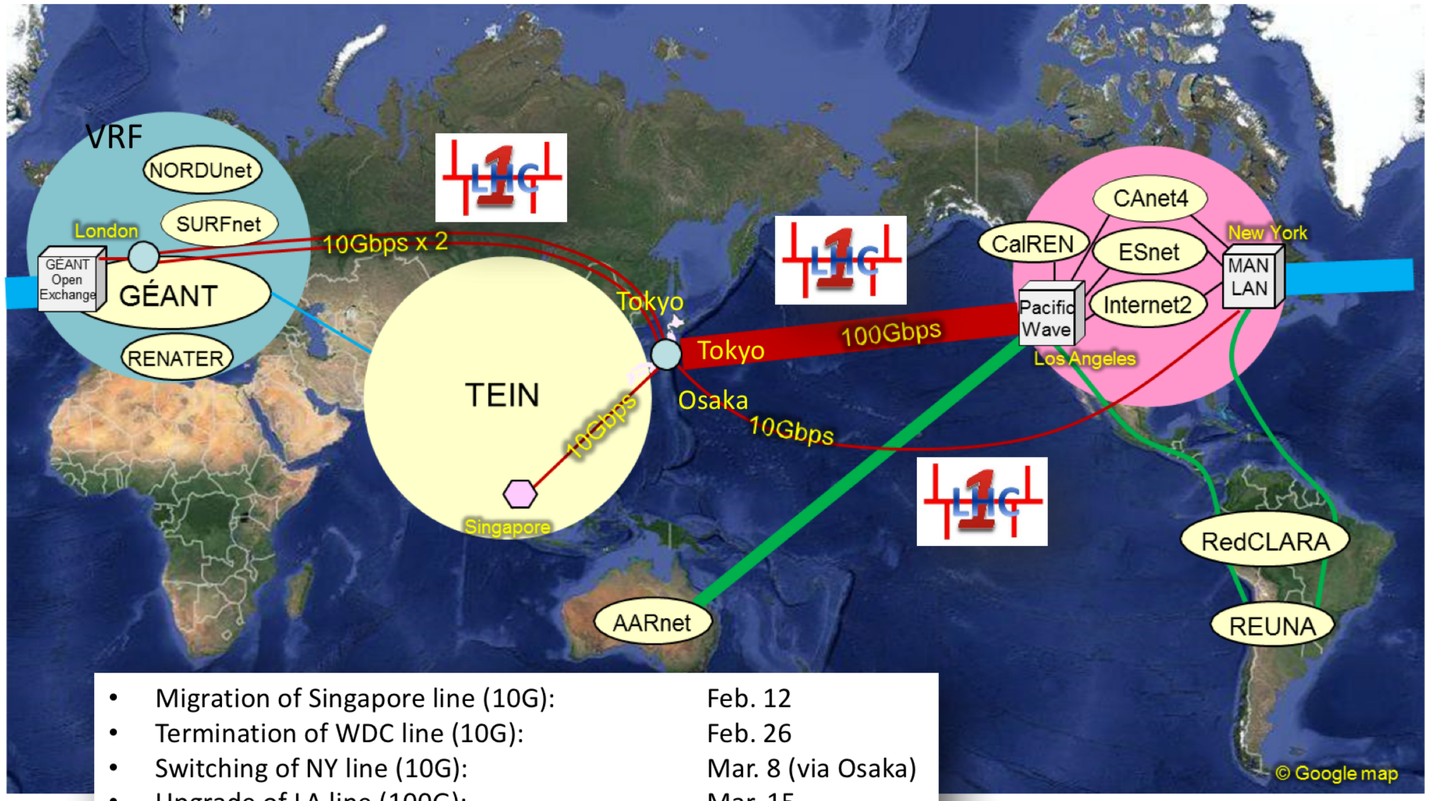

- Migration of Singapore line (10G): Feb. 12
- Termination of WDC line (10G): Feb. 26
- Switching of NY line (10G): Mar. 8 (via Osaka)
- Upgrade of LA line (100G): Mar. 15
- Connection from KEK (100G): Mar. 16
- New London Line (20G): Mar. 18

Y. Kubota (NII)

2016/04/18 KEK-CRC Tomoaki Nakamura 12

Source: T. Nakamura

Virtualisation technologies

- CC-IN2P3 currently operating a OpenStack-based cloud computing platform

  on-demand virtual machine management

  mainly for application services (web hosting, databases, development hosts, continuous integration, code repository, etc.)

  details: "IN2P3-CC cloud computing (IAAS) status"

- KEK-CRC considering the same technology for production

  conducted research for comparing efficiency of virtual machine vs Docker containers

  both for computation-heavy and for I/O-heavy workloads

  details: "Testbed for cloud computing and evaluation of Docker at KEK"

- Both sites regularly share experience on this topic

Equipment procurement and policies

- Detailed information exchange on KEK-CRC and CC-IN2P3 procurement process

including typical configuration, replacement and upgrade policies, costs, observed failure rates, data migration procedures, etc.

Upgrade of KEK Central Computer (KEKCC)

New KEKCC will start operation from Sep. 2016

Just in construction phase

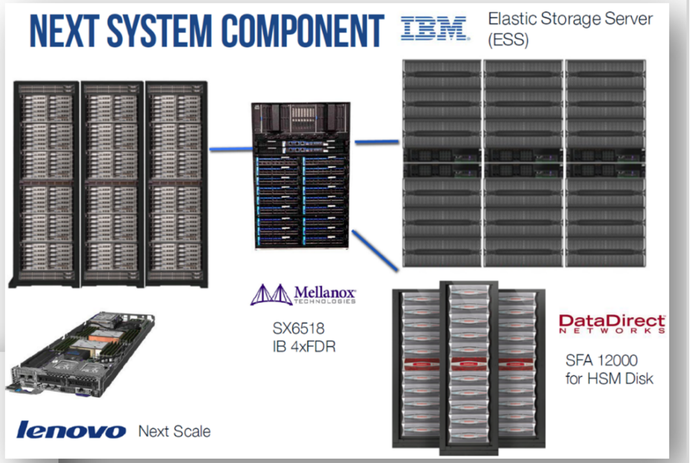

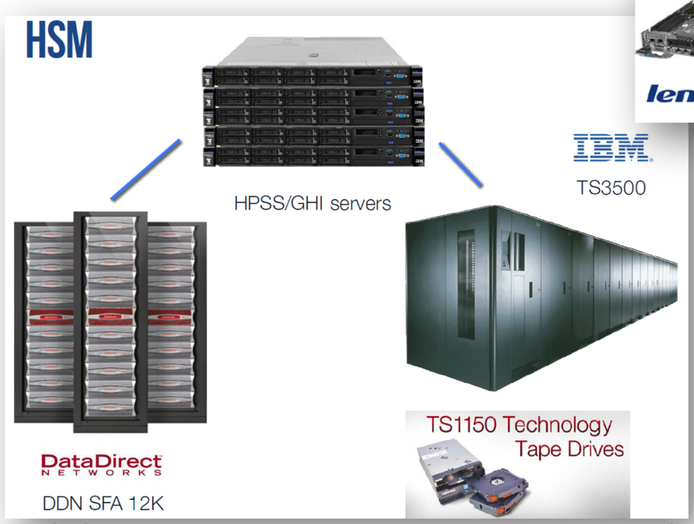

Very high Integration
- **10K CPU cores in 6 racks**
- **All comp. node have SSD**

High speed interconnect with storage
- **InfiniBand 4xFDR (56Gbps)**

Can use tape as disk transparently
- **GPSS-HPSS integration technology**

2016/04/18 KEK-CRC Tomoaki Nakamura 7

Source: T. Nakamura

Perspectives

Perspectives

- Submitted a project proposal to FJPPL 2016 call

  for continuing our joint work on the topics already covered

- New topic: to understand how to integrate a **GPU-based computing platform** to the daily operations of a computing center

  initial use case: simulation of physical and chemical processes at the molecular level by the Geant4-DNA project

  code is GPU-ready: demonstrated x250 execution speed-up compared to general purpose CPUs

  significant reduction in execution time directly benefits researcher as it shortens the experimentation cycle

  CC-IN2P3 platform to be composed of 20 Nvidia Tesla K80 GPU devices (40 GPUs in total), to enter production next summer

  CENBG to join CC-IN2P3 and KEK-CRC to work on this topic

Nvidia Tesla K80

Conclusions

Conclusions

- CC-IN2P3 and KEK-CRC maintain a long-standing effective collaboration around data processing for physics experiments

  in the last few years FJPPL funding has contributed to keep this relationship very active

- We intend to continue working together on the several challenges we share

References

References

FJPPL Computing Project Annual Workshop, Lyon, 10-11th March 2015, https://indico.in2p3.fr/event/11289

FJPPL Computing Project Annual Workshop, Lyon, February 10-11th 2016, https://indico.in2p3.fr/event/12701

Authorization extension for the secure use of ElasticSearch and Kibana, HEPiX Spring 2016 Workshop, Zeuthen, April 18th-22nd 2016, https://indico.cern.ch/event/466991/contributions/1143614

Monitoring at scale: a needle in the haystack, HEPiX Spring 2016 Workshop, Zeuthen, April 18th-22nd 2016, https://indico.cern.ch/event/466991/contributions/1143611

The next generation system of KEKCC, HEPiX Spring 2016 Workshop, Zeuthen, April 18th-22nd 2016, https://indico.cern.ch/event/466991/contributions/1143606