Infrastructure for testing accelerators and new technologies @CNAF

Daniele Cesini, Gaetano Maron
INFN-CNAF

INFN
Istituto Nazionale
di Fisica Nucleare

Accelerators at CNAF

- Small cluster Infiniband interconnected independent from T1
  - Funded by INFN-CNAF, INFN-BO, UniBo, former COKA project

- 27 Worker Nodes
  - CPU: 872 HT cores
    - 608 HT cores E5-2640v2
    - 48 HT cores X5650
    - 48 HT cores E5-2620v2
    - 168 HT core E5-3683v3
  - **15 GPUs:**
    - 8 Tesla K40
    - 7 Tesla K20
    - 2x(4GRID K1)
  - **3 MICs:**
    - 2 x Xeon Phi 5100
    - 1 x Xeon Phi 3100

- Dedicated STORAGE
  - 2 disks server
  - 60 TB shared disk space
  - 4 TB shared home

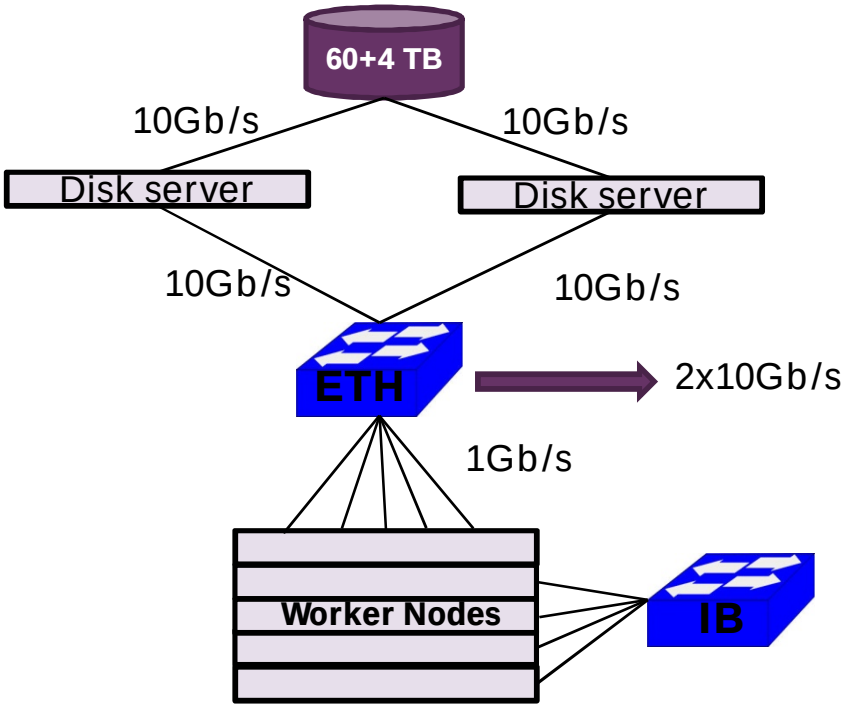

	CPU	GPU	MIC	TOT
TFLOPS (DP - PEAK)	**6**	**19**	**3**	**28**

Clusters Users

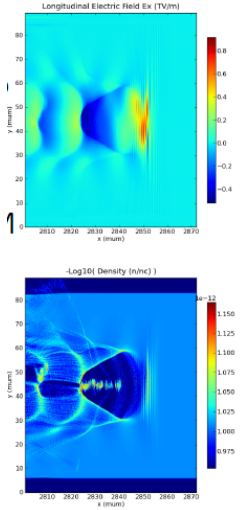

- Simulation of acceleration and plasma physics
  - 80% of the usage
  - Bologna and CERN groups
    - Preparatory runs before launching in supercomputers
- Astrophysics simulations
  - Tests by COSMO_WNEXT (PLANK in particular)
  - Collaboration between CNR, UniBo, INFN
- Radio-protection simulations at particle accelerators
  - INFN-UniBO collaboration
- **Currently 90% of the production usage is on the CPUs, not accelerators**
- Testing and performance evaluation of hw and sw

Testing hw and sw

- Recent testing activities
  - GPU and Phi performance
  - Application porting to GPU
  - Intel vs gcc compilers evaluation
  - GPU in the cloud
  - GPU and DOCKER
  - LHCb event building sw
- Cannot be done in production
  - Interference with other jobs
  - The batch system add complexity

- A constant number of nodes is out of the cluster, reducing the total power
- Frequent administrator interventions to install needed libraries or operating system and to restore the nodes back to production

Low power cluster

- The COSA project clusters (see next talk…)

16xARMv7
2xARMv8

4xINTEL AVOTON C-2750
4xINTEL XEOND-1540

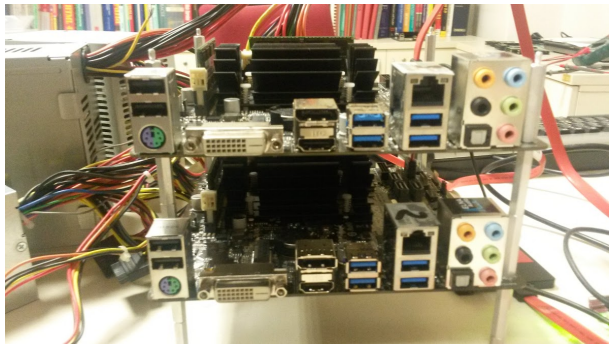

4xINTEL N3700

Other testbeds

- Testbed for event building sw
  - Few nodes low latency interconnected
    - Intel and Mellanox fabrics
    - Traditional and low power architectures
    - Event building for LHCb tested

- Testbed for the new Intel KNL Intel XEON-PHI with Omnipath

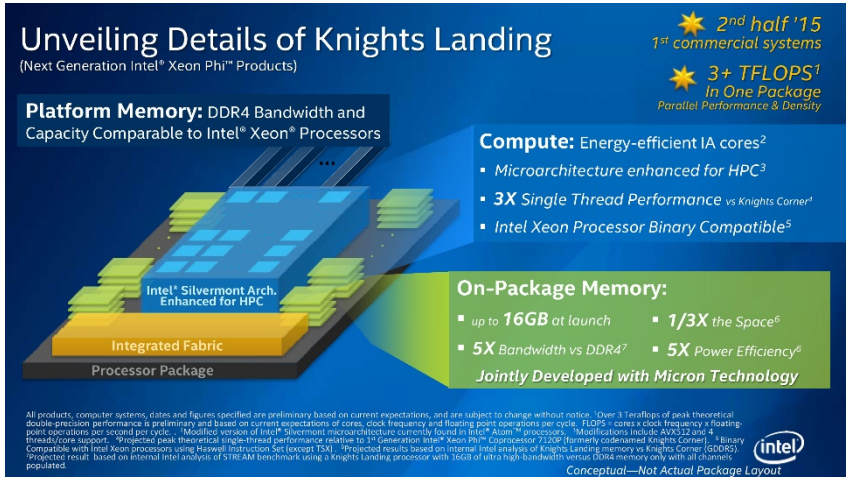

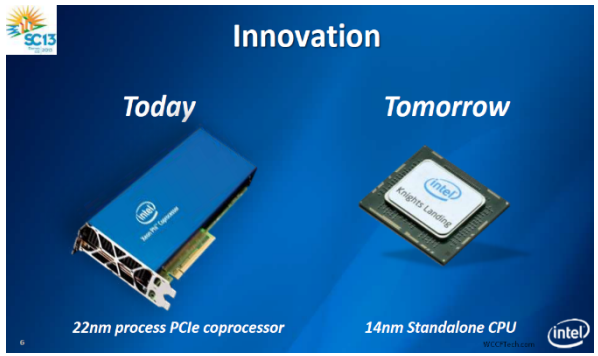

(We hope)

An infrastructure for testing new technologies @CNAF

- Reconfiguration of the current testbeds to create a single infrastructure dedicated only to testing purposes
  - Accelerators (GPU, KNC, KNL) – How many?
  - Low Power CPUs
  - Low latency interconnection for event builder sw
    - how many nodes?
  - New CPU and storage systems
- Services
  - Access to hw and sw to interested INFN users
  - Tools to book in advance and access the systems
  - System administration support

Research Agreement with the CERN Open Lab

- We are signing a Research Agreement with the CERN Open Lab
  - Tests at CERN OpenLab will be possible
  - Participation to Open Lab projects
    - Intel "Code Modernization" project as soon as the RA will be signed
      http://openlab.web.cern.ch/technical-area/computing-platforms-offline
      - Access to Intel sw and hw products located at CERN
      - 30k euro per year for training (@CERN)

Collaboration with the ISSS infrastructure

Infrastruttura di Supporto allo Sviluppo Software

Il progetto ISSS si rivolge agli sviluppatori software dell'INFN, soprattutto a quelli appartenenti a esperimenti di dimensione medio-piccola, con l'obiettivo di mettere a loro disposizione una infrastruttura di strumenti e servizi che li aiuti a produrre codice di qualità crescente, a costi ridotti e rispettando le scadenze.

Il progetto vuole inoltre rappresentare un punto d'incontro tra gli sviluppatori, per favorire la costruzione di una comunità all'interno della quale ci possa essere condivisione di conoscenza e di esperienze.

Grazie alla collaborazione offerta dai Servizi Nazionali, al momento sono disponibili i seguenti strumenti, accessibili attraverso un'interfaccia web:

- Project tracker, basato su JIRA
- Continuous Integration, basato su Jenkins
- Code repository, basato su git

L'accesso ai servizi è automatico ed è basato sulle credenziali dell'AAI dell'INFN. Per il project tracker è possibile usare il certificato X.509 personale.

The “Nuove tecnologie” infrastructure would be the hardware counterpart of the ISSS project

Funding

- Manpower from CNAF
  - i.e. fraction of the user support team
  - open to any possible collaboration
- HW & SW
  - CSN V
  - Any other CSN interested (i.e requests from CSN2/VIRGO)
  - Direct funding from interested experiments
  - Letter of Intent to the GE from interested experiments on specific topics
  - CCR on specific topics
    - Already contributing to sw license (i.e. Intel compiler)

Summary

- At CNAF about 20 accelerators are available in a cluster, low latency interconnected
  - Interested users can request access for testing purposes and/or production
- We are restructuring all our testbeds in an infrastructure open to projects and collaborations with other departments and experiments
- We are signing a Research Agreement with CERN OpenLab
  - Participation to OpenLab project
  - Intel Code Modernization