

Triton and Science IT

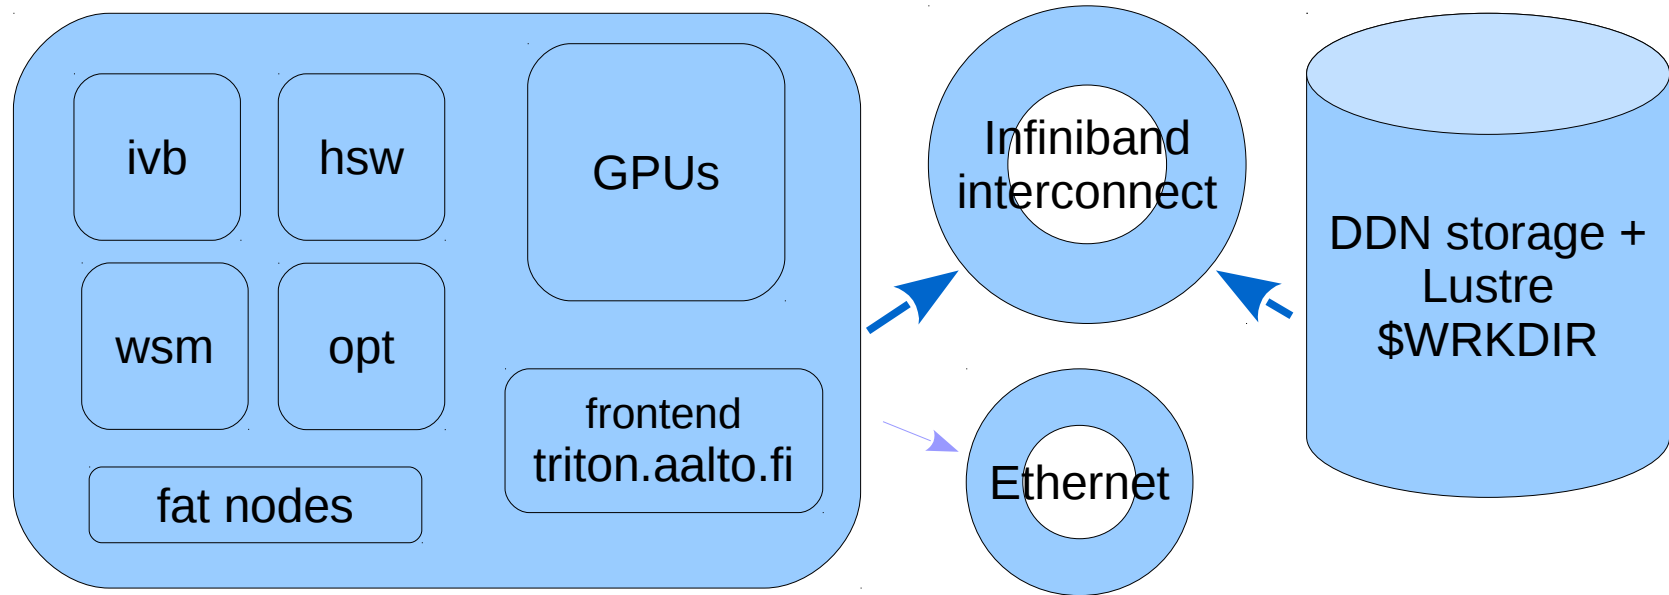
SCIP Kickstart 2017

*Ivan Degtyarenko, Janne Blomqvist, Mikko Hakala,
Simo Tuomisto, Richard Darst*

School of Science, Aalto University

June 5-7, 2017

Triton cluster



Heterogeneous: several generations of hardware since 2009

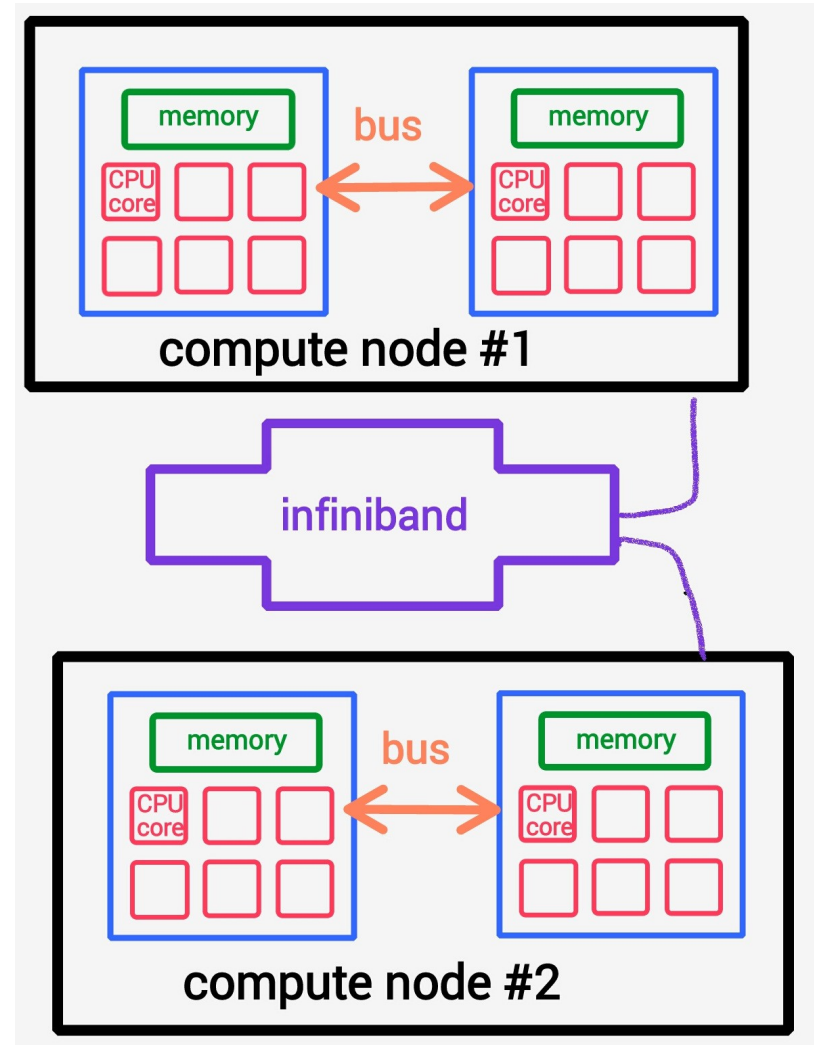
Triton as a parallel computer

In principle, all the computers today are parallel from a hardware perspective (!)

Triton's compute nodes are multi-processor parallel computers in itself. On top of that they are connected through network to make larger computer: a cluster.


Hybrid architecture: it is both and shared memory and distributed memory

- **Shared:** within one node
- **Distributed:** over the cluster



Triton in nutshell

- Heterogeneous **compute nodes** with different Xeon and Opteron nodes:
- Memory **up to 256GB** on “normal nodes”, plus several fat (large memory) nodes with **1TB of RAM**
- **InfiniBand** as interconnect
- GPU computing with **Tesla**
- DDN + **Lustre filesystem** as a storage
- Software: **CentOS 7, SLURM** as a batch system, additional apps provided through “**module**”



Reference: Triton tech specs

- **Dell PowerEdge C6320 (pe[1-82])**
 - 2x Xeon E5-2680 v3 @ 2.50GHz (12core)
 - 128 and 256 GB RAM
 - SATA disk (7.2k) 1TB of local storage
- **DELL PowerEdge C4130 (gpu[20-22])**
 - 2x Xeon E5-2620 v3 @ 2.40GHz (12core)
 - 128GB RAM
 - 4x Tesla K80
- **Upcoming soon**
 - Tesla P100
 - more C6320 with E5-2680 v4
- **HP SL230s G8 (ivy[1-48])**
 - 2x E5 2680 v2 @ 2.8 GHz (10 cores)
 - 1.8TB local storage
 - 64/256 GB RAM
- **HP SL390s G7 (wsm[1-142],gpu[1-19])**
 - 2x Xeon X5650 @ 2.67GHz (6core) / 48 GB RAM
 - 2x local SATA disks (7.2k) -> 800GB of local storage
 - 2x Tesla M20[5|7|9]0 GPU (Fermi, 3/6 GB mem) on 19 servers / 24 GB RAM
- **HP DL580 G7 (fn[1,2])**
 - 4x X7542 @ 2.67 GHz (6core)
 - 1.3TB local storage
 - 1024 GB RAM
- **HP BL465c G6 (opt[1-112])**
 - 2x AMD Opteron 2435 @ 2.6GHz (6core)
 - 200GB local storage
 - 16/32/64 GB RAM

Triton storage



- **Parallel filesystem **Lustre**:**
 - v2.5
 - mounted under /scratch
- **DDN SFA12K**
 - 1.8PB of raw space
 - 100 000 IOPS, 30 GB/s read/write
 - Round Robin and IML
- **\$HOME**
 - Cheap/slow 1GbE

Triton interconnect



- **Infiniband** fabrics on Triton: fast & expensive, very high throughput and very low latency
 - for **MPI**, **Lustre** [aka \$WRKDIR] on Triton
 - Xeon E5-2680 v3: PowerEdge C6320, C4130: **4xFDR Infiniband** (56GT/s)
 - Xeon Westmere nodes: SL390, DL580: **4xQDR Infiniband** (40GT/s)
 - Xeon IvyBridge nodes: SL230: **4xFDR Infiniband** (56GT/s)
 - Opteron nodes: BL465: **4xDDR Infiniband** (20GT/s)
 - 4x spine switches
 - Mostly fat-tree configuration, subscription differs
 - See full Infiniband map at <https://wiki.aalto.fi/display/Triton/Cluster+overview#Clusteroverview-Networking>
- **Ethernet**: cheap & slow
 - **1G** internal net for SSH, NFS [aka /home] on Triton
 - **10G uplink** towards Aalto net

GPUs on Triton

As of spring 2017 the main driven force are [Tesla K80](#) cards. P100 cards are upcoming fall 2017. The full list includes



- 12x [Tesla K80](#) @ gpu[20-22]
- 22x [Tesla M2090](#) @ gpu[1-11]
- 6x [Tesla M2070](#) @ gpu[17-19]
- 10x [Tesla M2050](#) @ gpu[12-16]

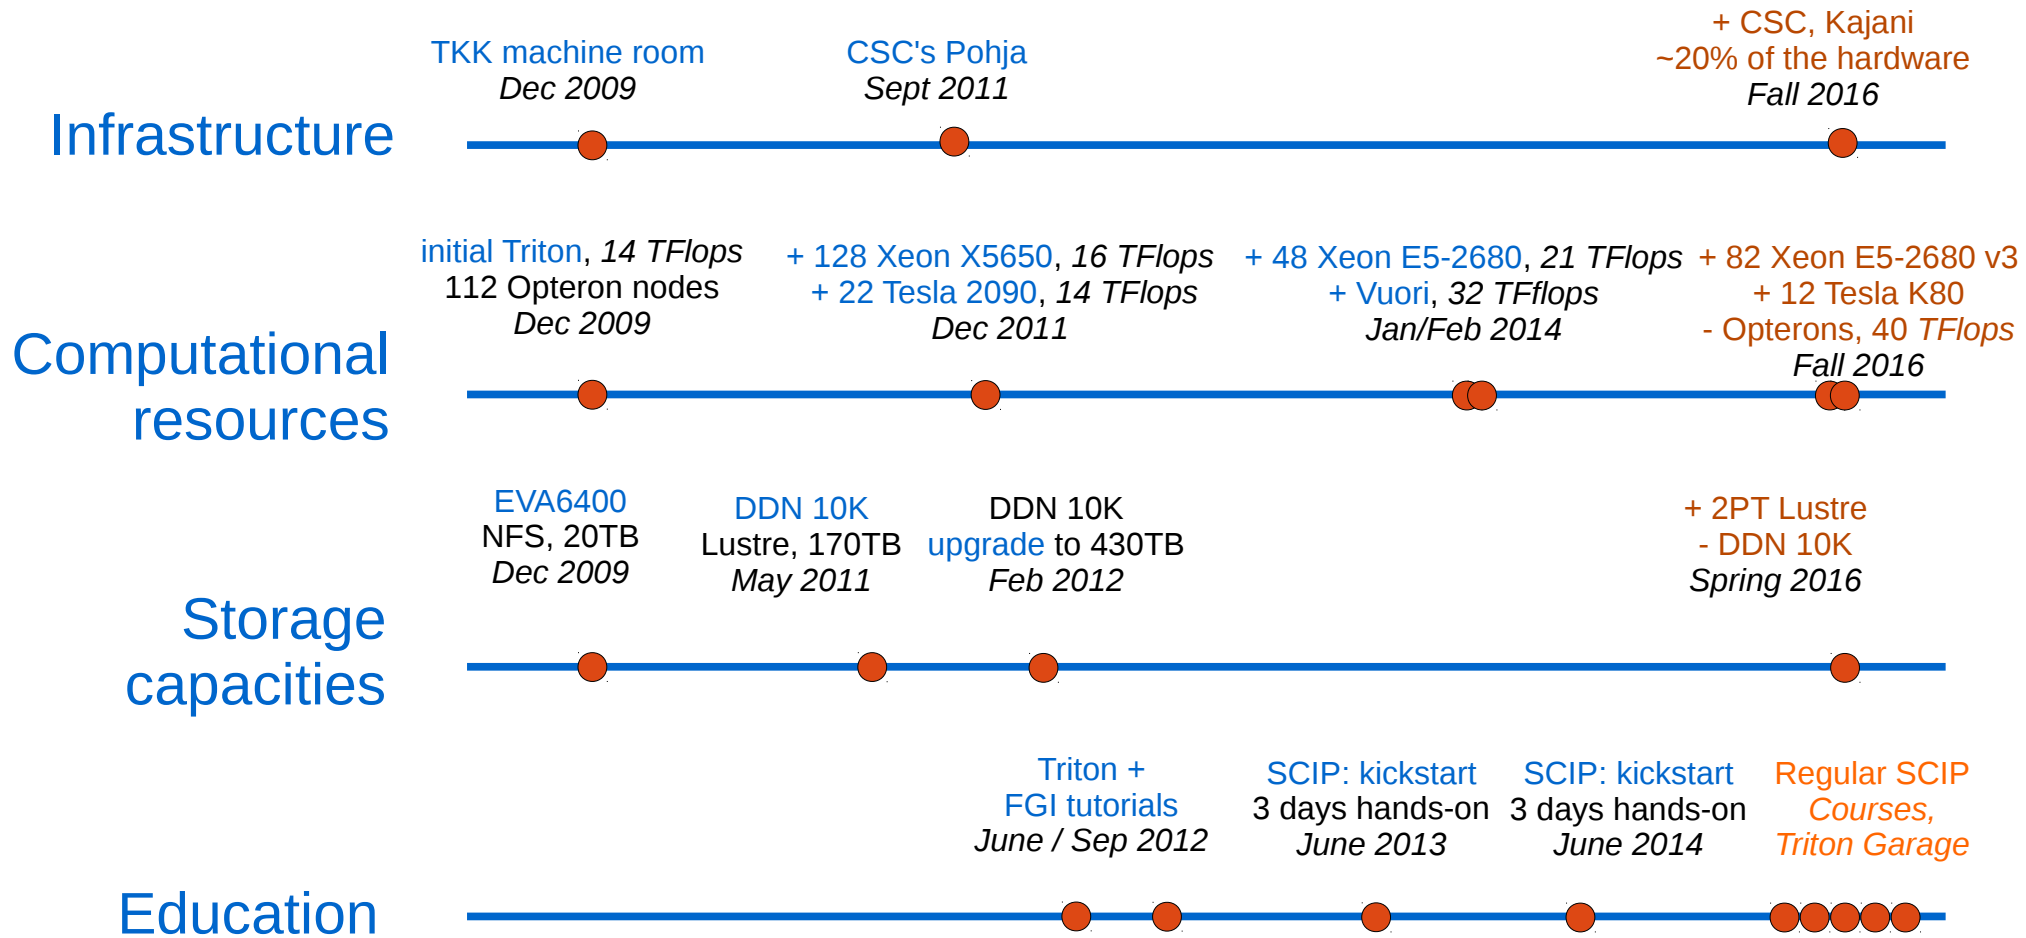
Scientific Computing @ Aalto:

Science IT (<http://science-it.aalto.fi>)

- Computational Studies are within top-5 highest priorities at Aalto
- midrange Scientific Computing resources for Aalto researchers: Triton-cluster
- wide collaboration between Finnish Universities and CSC, part of European grid infrastructure
- 2005 collaboration started under M-grid; 2010 -- FGI, a wider consortium; 2014 – FGCI, Finnish Grid and Cloud Infrastructure; 2017 – FGCI phase 2 (*in progress*)
- joint procurement and administration collaboration, focus on serving (special) needs of Aalto's computational science, i.e. raw CPU and GPU power, big data sets, large memory servers



Science IT timeline



Science IT management

Management board

- Prof. Martti Puska (SCI/PHYS) , Prof. Keijo Heljanko (SCI/CS, chair), Prof. Mikko Kurimo (ELEC/SPA), Prof. Lauri Parkkonen (SCI/NBE), Kenrick Bingham (SCI, Technology Manager), Prof. Petteri Kaski (SCI/CS)

Stakeholder model

- All costs are distributed to schools (departments) based on the agreed Science-IT share
- Key departments (PHYS, CS, NBE) contribute 50% to administrative personnel responsible for day-to-day tasks. Local support person named at each participating unit.
- 20% of resources for Grid => Freely available for Finnish research community via grid-interface
- Small share 5% for users outside stakeholder model

Science IT support team

The core team

- [Mikko Hakala](#), D. Sc. (Tech) (CS+NBE/SCI)
- [Janne Blomqvist](#), D. Sc. (Tech) (PHYS+NBE/SCI)
- [Simo Tuomisto](#), (CS/SCI)
- [Richard Darst](#), D. Sc. (Tech) (CS/SCI)
- [Ivan Degtyarenko](#), D. Sc. (Tech) (PHYS/SCI)
- responsible for "daily cluster maintenance"

Support team members

- Jussi Hynninen (ELEC)
- named by Departments to be [local Science-IT support team members](#)
 - provide [department user support](#), being close to researchers
- being [contact person](#) between schools (departments) and Science-IT support

Getting help

- points of search for help:
 - at wiki.aalto.fi look for [Triton User Guide](#)
 - access requires valid Aalto account
 - see also [FAQ](#) over there
 - follow triton-users@list.aalto.fi list: all users MUST BE there
 - <https://version.aalto.fi/gitlab/AaltoScienceIT/triton/issues> for any issue; see whether it has been published already, if not, then shoot
 - local [support team](#) member: for granting access, quota