

CERN IT procurement update

HEPiX Autumn/Fall 2018 at Port d'Informació Científica



Oct 2018

Olof Barring, CERN IT-CF


Outline

- Brief intro to team & activities
- Hardware tech: trends & wishes
- Inventory & outlook
- Repairs
- Some recent hardware issues
- Workflow automation
- Conclusions

Brief intro


- Team

- New section leader (as of ~now): Eric Bonfillou
- 4 staff (2 engineers + 2 technicians)
- 4 temporary (Fellows, project associates) for Openlab and Techlab






- Activities

- Procurement: requirement, specification, evaluation, order, delivery and acceptance
- Hardware support
 - Train and assist repair contractor
 - Follow up systematic issues and plan fw upgrades etc.
- Hosting Openlab & Techlab clusters
 - GPU, FPGA, different CPU architectures (x86 and non-x86)
 - End-user access and support



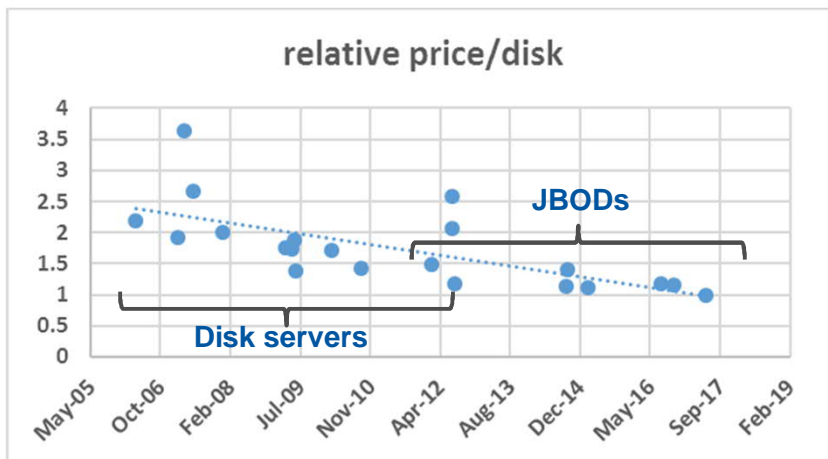


HW Technology Trends (+ wishes)

Compute performance



- Unit = server, with memory, storage, PSUs etc.
- Price comparison not corrected for currency exchange
- Memory and Flash market suffered important demand growth in 2017
 - \$/DRAM ~x2-x3 over 12-18 months

Storage capacity



- The relative price for JBOD hosted disk does not include hosting server
 - Adds 8-40% to price depending on nb. JBODs per server
 - Can attach 1-8 JBODs per server:
 - 24 – 192 disks
 - 144 – 1152 TB (6TB disks)
 - Max today 2.3 PB (192x 12TB disks)

Growing use of SSDs

- Since 2016 we buy majority of servers with SSDs instead of mechanical disks
- Hypervisors mostly don't need much capacity but can benefit a lot from the better IOPs
 - Typically ~10-20% of local 6TB was used
- In principle (much) better failure rate
 - Initial concern about wear-out turned out groundless
 - We rarely (~never) need >1 DWPD
 - But SSDs can have their own (fw) issues (see later)
- Bulk (JBOD) storage still of course with 'normal' disks (12TB today)

SSD trends / maybes

- Interesting form-factors
 - M.2 (no hot-swap), U.2
 - Intel EDSFF (Rulers) short (11cm) and long (30cm)
 - Samsung's NGSFF (11cm)
- Access protocols
 - SATA → NVMe
 - Scales better and a bit more future proof (Gen3/4/5)
- NVMe over Fabric?
 - Rack-level fabric with ToR storage boxes
 - May enable disaggregation of local storage...
 - Zoning granularity?
- NVDIMM...
 - DRAM+Flash + supercap
 - 3D Xpoint

Accelerators

- GPGPUs
 - Exist in small scale in Techlab since 2014
 - Nvidia: V100, P100, GTX1080(Ti)
 - AMD: Radeon Vega
 - Purchasing another 4 nodes, each with 4x V100, for production use
 - Cloud-provisioned in Openstack work fine (both PCI Passthrough and vGPU)
 - Consumer cards: beware of the **EULA restrictions** for using the drivers etc.!

Accelerators

- GPGPUs

- Exist in small form factors
 - Nvidia: V100, T400
 - AMD: Radeon Instinct
- Purchasing and operating costs are high
- Cloud-provisioning (e.g. AWS, Azure, GCP)
- Consumer cards (e.g. GeForce RTX) are also used as accelerators

These operating systems, or other operating systems derived from the source code to these operating systems, may be copied and redistributed, provided that the binary files thereof are not modified in any way (except for unzipping of compressed files).

2.1.3 Limitations.

No Modification or Reverse Engineering. Customer may not modify (except as provided in Section 2.1.2), reverse engineer, decompile, or disassemble the SOFTWARE, nor attempt in any other manner to obtain the source code.

No Separation of Components. The SOFTWARE is licensed as a single product. Its component parts may not be separated for use on more than one computer, nor otherwise used separately from the other parts.

No Sublicensing or Distribution. Customer may not sell, rent, sublicense, distribute or transfer the SOFTWARE; or use the SOFTWARE for public performance or broadcast; or provide commercial hosting services with the SOFTWARE.

No Datacenter Deployment. The SOFTWARE is not licensed for datacenter deployment, except that blockchain processing in a datacenter is permitted.

TERMINATION

use
(vGPU)
drivers

Accelerators

- GPGPUs

- Exist in small form factors
 - Nvidia: V100, T400
 - AMD: Radeon Instinct
- Purchasing and operating costs are high
- Cloud-provisioning (e.g. AWS, Azure, Google Cloud)
- Consumer cards (e.g. GeForce RTX) are also used as accelerators

These operating systems, or other operating systems derived from the source code to these operating systems, may be copied and redistributed, provided that the binary files thereof are not modified in any way (except for unzipping of compressed files).

2.1.3 Limitations.

No Modification or Reverse Engineering. Customer may not modify (except as provided in Section 2.1.2), reverse engineer, decompile, or disassemble the SOFTWARE, nor attempt in any other manner to obtain the source code.

No Separation of Components. The SOFTWARE is licensed as a single product. Its component parts may not be separated for use on more than one computer, nor otherwise used separately from the other parts.

No Sublicensing or Distribution. Customer may not sell, rent, sublicense, distribute or transfer the SOFTWARE; or use the SOFTWARE for public performance or broadcast; ~~or provide commercial hosting services with the SOFTWARE.~~

No Datacenter Deployment. The SOFTWARE is not licensed for datacenter deployment, except that blockchain processing in a datacenter is permitted.

TERMINATION

use
(vGPU)
drivers

Accelerators

- GPGPUs
 - Exist in small scale in Techlab since 2014
 - Nvidia: V100, P100, GTX1080(Ti)
 - AMD: Radeon Vega
 - Purchasing another 4 nodes, each with 4x V100, for production use
 - Cloud-provisioned in Openstack work fine (both PCI Passthrough and vGPU)
 - Consumer cards: beware of the **EULA restrictions** for using the drivers etc.!
- FPGA
 - One (Intel Arria 10) in Techlab
 - Licenses add to procurement complexity (OpenCL + BSP)

Open Compute

- Started by Facebook in 2011
 - “...most efficient server, storage and data centre hardware designs for scalable computing”
- In 2015 we procured 6 racks with 108 servers + 12 JBODs (~2PB)
- In production since ~3 years
- Experience:
 - Difficult to procure in a public RFP
 - See CHEP'15 paper
 - Painful commissioning
 - Running ok in production:
 - Not much hw issues
 - Easy to repair



Other trends

- Software Defined Infrastructure:
 - Rack-scale architecture
 - Standardization efforts around mgmt. protocols, orchestration
 - Redfish, Swordfish
 - Flexible (disaggregated) infrastructure
- Open19.org
 - A bit like OCP but different objectives
 - Mechanical and connectivity standard
 - Future perhaps uncertain?
 - Started by LinkedIn, which were acquired by Microsoft
 - Microsoft contributed Olympus to OCP...

Inventory



Machine Inventory

- Current inventory (Meyrin + Wigner):
 - 2'850 (JBOD) disk arrays
 - 768 single board servers
 - 3'028 2U4N (Quads) } → 12,880 servers
 - 1'200 (Meyrin) intelligent PDUs
- Receiving ~now:
 - 167 JBODs: 6TB and 12TB disks
 - 377 servers (some with 40GbE)
 - 85 intelligent PDUs
- Outlook (LS2 & ramp up for Run3):
 - 2019: ~like 2018
 - 2020 – 2021: ~6 000 servers + ~1 000 Disk arrays

Hardware repairs

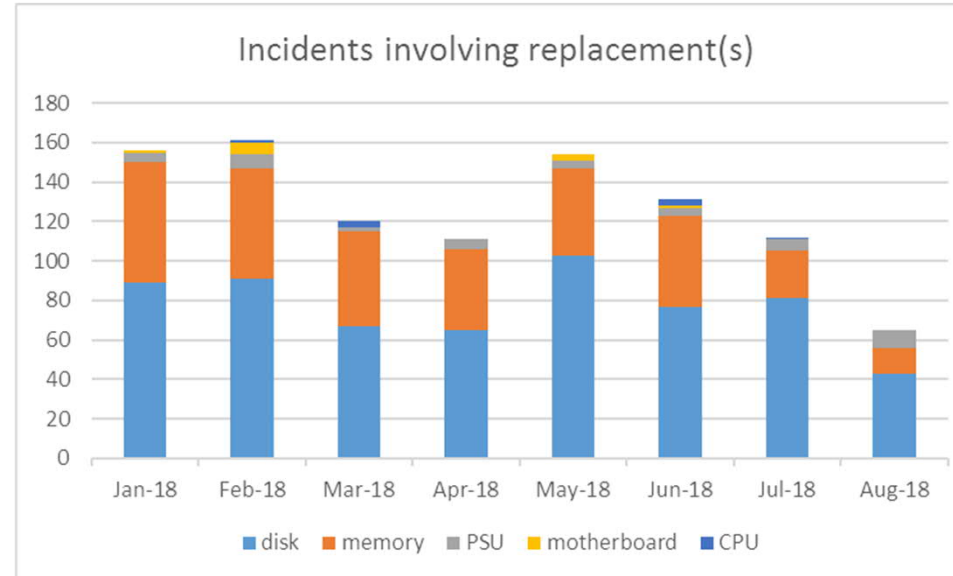


Hardware repairs

- Since 2013: single contractor for all repairs
 - Works well
 - Require some initial training on our hardware + regular refresh
 - New contract started 1/10/2018
- Automatically triggered via Service NOW (SNOW) tickets
 - Automation: Monitoring exception → ticket
 - Technician: incident analysis → repair if needed

Incidents with replacements

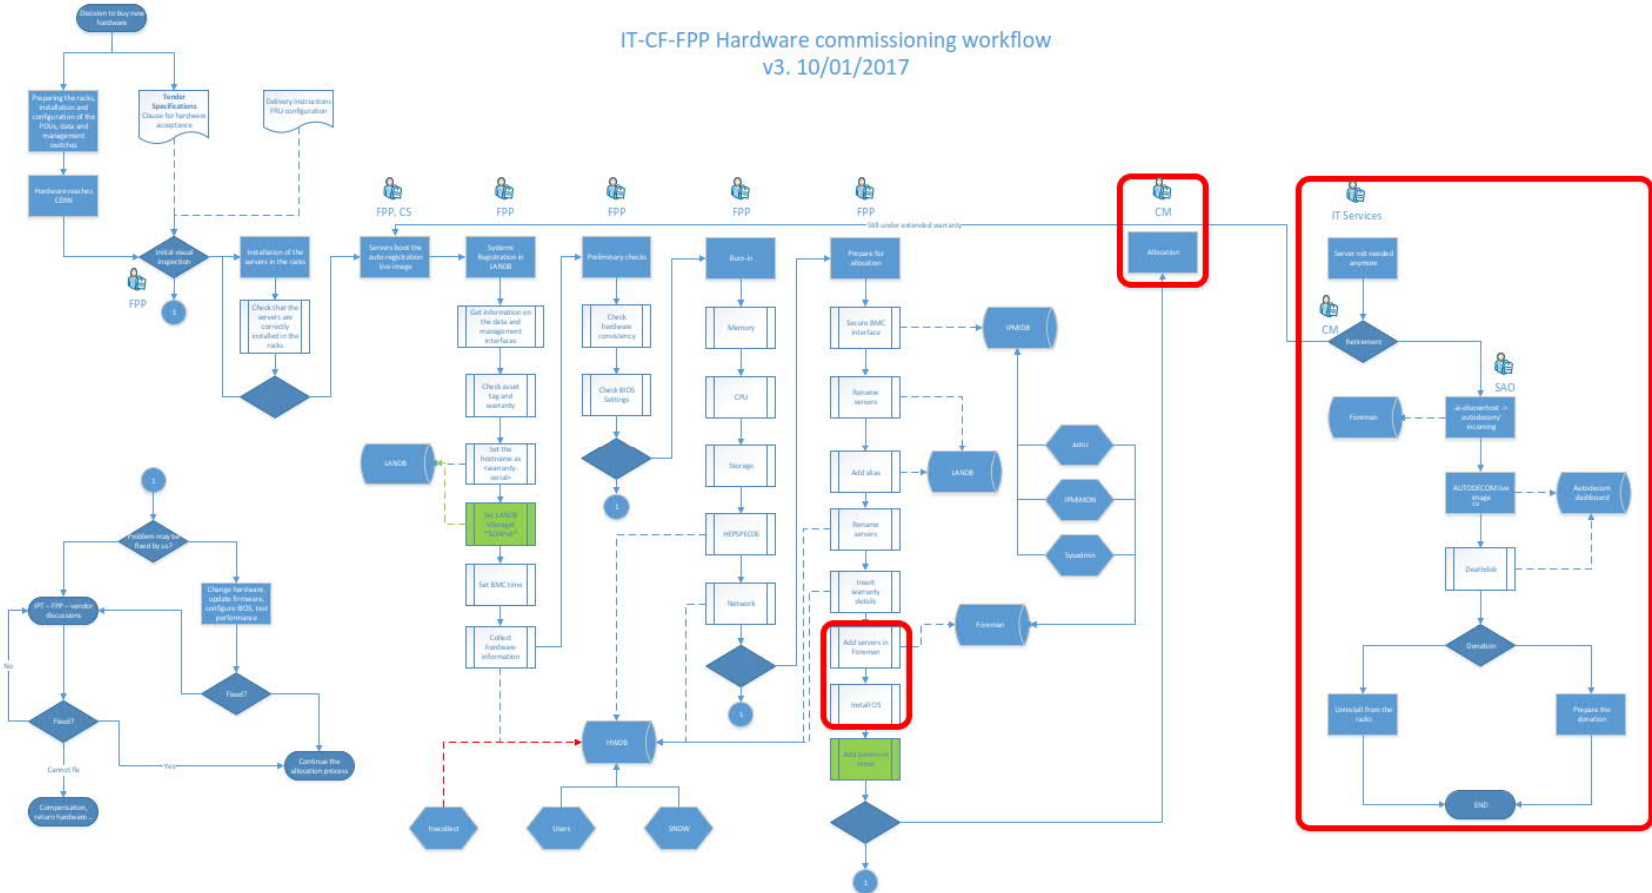
- Hardware monitoring is sensitive
 - Only ~10% events lead to component replacements
- Disk replacements dominate..
- ..but also memory has become significant
 - Mostly repeated correctable ECC. Un-correctable errors are rare



Recent (ongoing) hw issues

- Intel 2130W PSU issue → severe CPU throttling:
 - **Intrusive**
 - Requires BIOS/BMC/ME and PSU firmware update
 - Affected systems: 178
- Intel DC3510 SSD firmware issue → risk aborts on write
 - **Intrusive**
 - Requires reboot after SSD firmware update
 - Affected systems: 1528
- SMASH protocol enabled by defaults on Quanta S2L/S2S
 - **Transparent**
 - Only requires BMC reboot after reconfiguration
 - Affected systems (BMC): 3212

Workflow automation



Workflow automation

- Project with Openstack@CERN
 1. Feasibility of using Ironic for existing workflows for (*)
 - a) Registration & Burn-in
 - b) Commissioning,
 - c) Repurposing, and
 - d) Decommissioning
 2. Hardware (component) inventory (**)
 - a) Agent-free (OOB discovery, e.g. Redfish)
 - b) Track and reconcile changes (e.g. parts replacement)
- Good progress on workflow migration
 - In process of enrolling all servers in Ironic
- Inventory part just started (possibly with Oath)

(*) <http://openstack-in-production.blogspot.ch/2018/03/hardware-burn-in-in-cern-datacenter.html>

(**) https://docs.google.com/document/d/144g8E_fzGD4WZzMvswkeowzclLL4hxg4QDS-46KYCcQ/edit?usp=sharing

Summary

- Some team changes
- Performance and capacity continue to improve
 - But not anywhere near the rate needed to meet Run4 timescale
- Accelerators
 - GPGPUs in Techlab and soon also in production
 - FPGAs tamed to data centre use
- Contractor for hardware repairs has worked well
 - But the alarming could be improved (only 10% lead to a repair)

Links

HEPiX Fall 2010 : CERN IT Facility Planning and Procurement. <https://indico.cern.ch/event/92498/contributions/1278919/>
HEPiX Spring 2012 : New burn-in test. <https://indico.cern.ch/event/160737/contributions/1407865/>
HEPiX Fall 2013 : Automatic server registration and burn-in framework. <https://indico.cern.ch/event/247864/contributions/1570349/>
HEPiX Spring 2014 : Open Compute at CERN. <https://indico.cern.ch/event/274555/contributions/615732/>
HEPiX Spring 2015 : A look beyond x86: OpenPOWER8 & AArch64. <https://indico.cern.ch/event/346931/contributions/817833/>
HEPiX Fall 2015 : Asset management in CERN data centres. <https://indico.cern.ch/event/384358/contributions/909190/>
HEPiX Fall 2016 : Deploying Open Compute hardware at CERN. <https://indico.cern.ch/event/531810/contributions/2298924/>
HEPiX Spring 2018 : Techlab benchmarking web portal. <https://indico.cern.ch/event/676324/contributions/2981840/>
CHEP'13 : Experience with procuring, deploying and maintaining hardware... <https://indico.cern.ch/event/214784/contributions/1512499/>,
CHEP'15 : Experience of public procurement of Open Compute servers. <https://indico.cern.ch/event/304944/contributions/1672389/>

