

CSC and Pouta

Kalle Happonen

2.6.2015



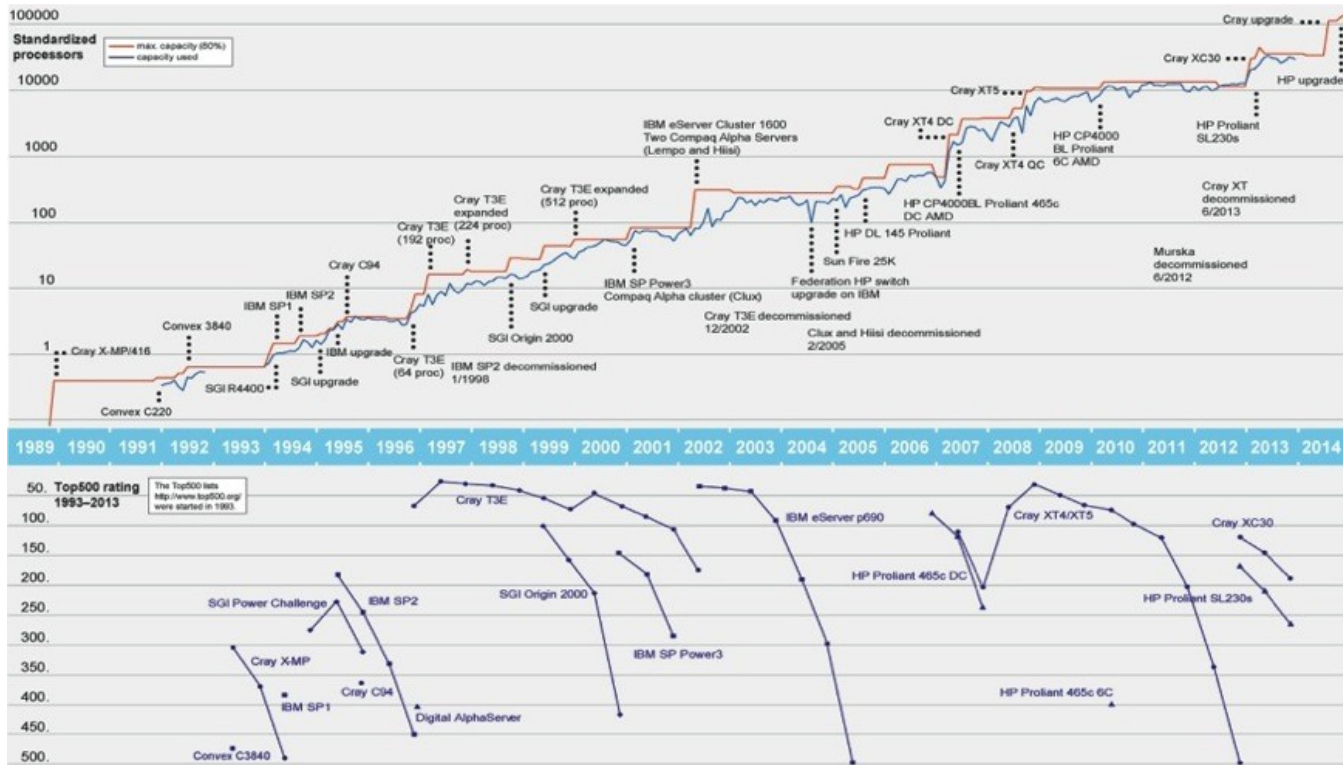
CSC-IT CENTER FOR SCIENCE

About CSC

- CSC offers IT -support and resources for research, education, culture and administration
- CSC provides Finland's widest selection of scientific software and databases and Finland's most powerful supercomputing environment that researchers can use via the Funet network
- CSC - IT Center for Science Ltd. is a state-owned, non-profit company administered by the Ministry of Education and Culture
- www.csc.fi

- Founded in 1971 as a technical support unit for Univac 1108
- CSC connected Finland to the Internet in 1988
- Reorganized as a company in 1993
- CSC's datacenter in Kajaani started in 2012
- Turnover 32,7 milj. euros in 2014

CSC Top 500





I'm a geek, give me numbers!

- 160 compute nodes (KVM)
 - Hpc, generic, io intensive
- 120TB + ~500 TB storage (iSCSI and Ceph)
- Prod, staging, devel environments ~ 25 management Vms (not in cloud)
- 10 GbE and 40 GbE networking

9:00am-9:40am	Austin Convention Center - Level 4 - MR 15	My calendar
Deploying OpenStack for The National Science Foundation's Newest Supercomputers HPC / Research		
PRESENTATION		
9:50am-10:30am	Austin Convention Center - Level 4 - MR 15	My calendar
High-Performance OpenStack for Science and Data Analytics in a Hybrid Cloud Environment HPC / Research		
PRESENTATION		
11:00am-11:40am	Austin Convention Center - Level 4 - Ballroom F	My calendar
Chasing the Rainbow – National Computational Infrastructure's Pursuit of High-Performance OpenStack Cloud HPC / Research		
PRESENTATION		
11:00am-11:40am	Austin Convention Center - Level 4 - MR 15	My calendar
Glance and SLURM: User-Defined Image Management on HPC Clusters HPC / Research		
PRESENTATION		
11:50am-12:30pm	Austin Convention Center - Level 4 - MR 15	My calendar
Bending Iron for Big Iron HPC / Research		
PRESENTATION		
2:40pm-3:20pm	Austin Convention Center - Level 4 - MR 15	My calendar
Cloud Infrastructure to Help Researchers Build 21st Century Microscopes HPC / Research		
PRESENTATION		
3:30pm-4:10pm	Austin Convention Center - Level 4 - MR 18 A/B	My calendar
Chameleon: An experimental Testbed for Computer Science as Application of Cloud Computing HPC / Research		
PRESENTATION		
3:30pm-4:10pm	Austin Convention Center - Level 4 - MR 15	My calendar
Building Efficient HPC Clouds with MVAPICH2 and OpenStack over SR-IOV Enabled Infiniband Clusters HPC / Research		
PRESENTATION		
4:30pm-5:10pm	Austin Convention Center - Level 4 - MR 16 A/B	My calendar
OpenStack for High-Performance Bioinformatics HPC / Research		
PRESENTATION		

Is cloud Ready for $H(P|T)C$?

But the overhead of virtualization...

*Yeah... you're **completely** missing the point*

<http://www.slideshare.net/moreirabelmiro/cpu-optimizations-in-the-cern-cloud>

The scheduling on the other hand...

Better, but still missing the point

Benefit is not measured in CPU cycles

Who uses the cloud?

- Hadoop/Spark users
- LHC Stuff
- BIO people
- Devs
- Courses
- Etc.




Jun 2, 2016

"So you've been running a cloud for 4 years?"

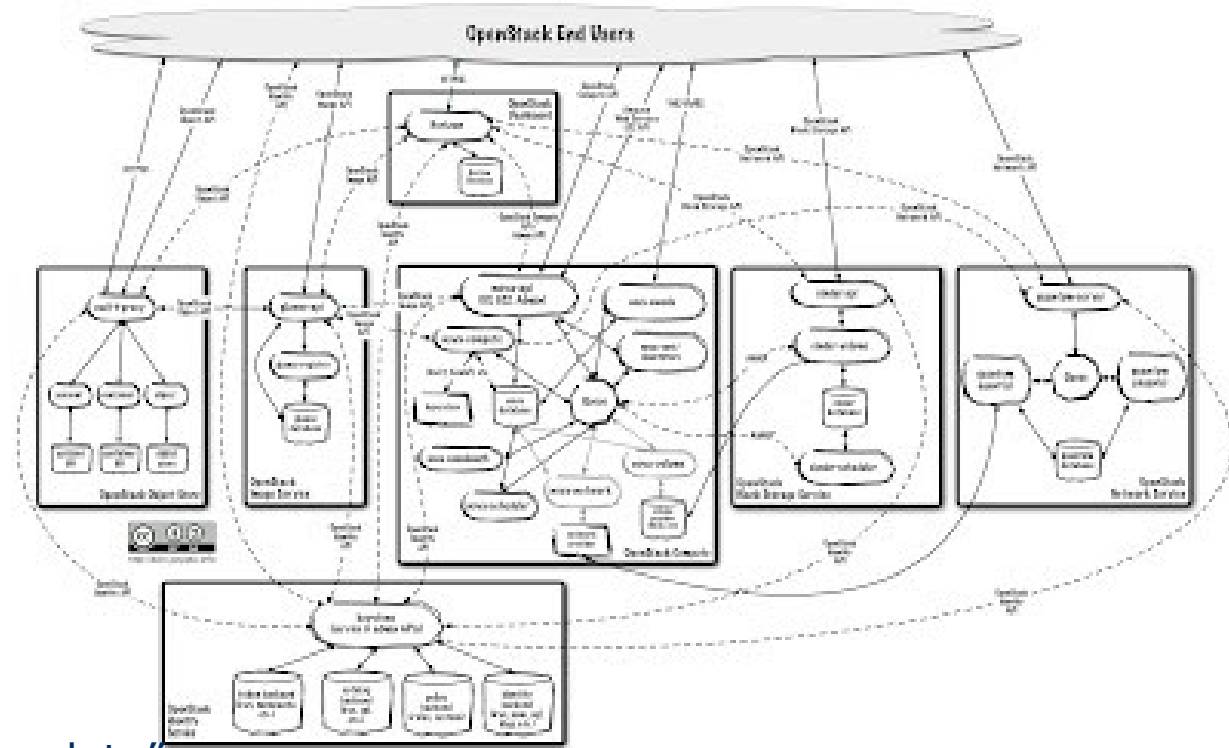
"Yeah"

"How's that been working out for you?"

"Uhm, do you have 2 hours?"



OpenStack



- 2 releases a year
- Upgrades aren't just "yum update"
- You ARE the IT department

Running OpenStack

- Configuration management
- Processes around configuration management
- Software project, not “traditional sysadmining”

Oh, and it can never be down, OK?

Questions?

@KalleHapponen