Expanding OpenStack Community in Academic Fields

Tom Fifield
tom@openstack.org
@TomFifield


We are not Special!
Challenge the must-have lists at project start
Are those requirements **really** justified ?
Accumulating technical debt stifles agility
There is no Moore's Law for people
Automation needs APIs, not documented procedures
Find open source communities and contribute
Understand ethos and architecture
Stay mainstream, stay up to date
04/11/2014
Tim Bell - OpenStack Paris
17


openstack
summit
BENVENUTO
TO THE OPENSTACK SUMMIT!

Industries

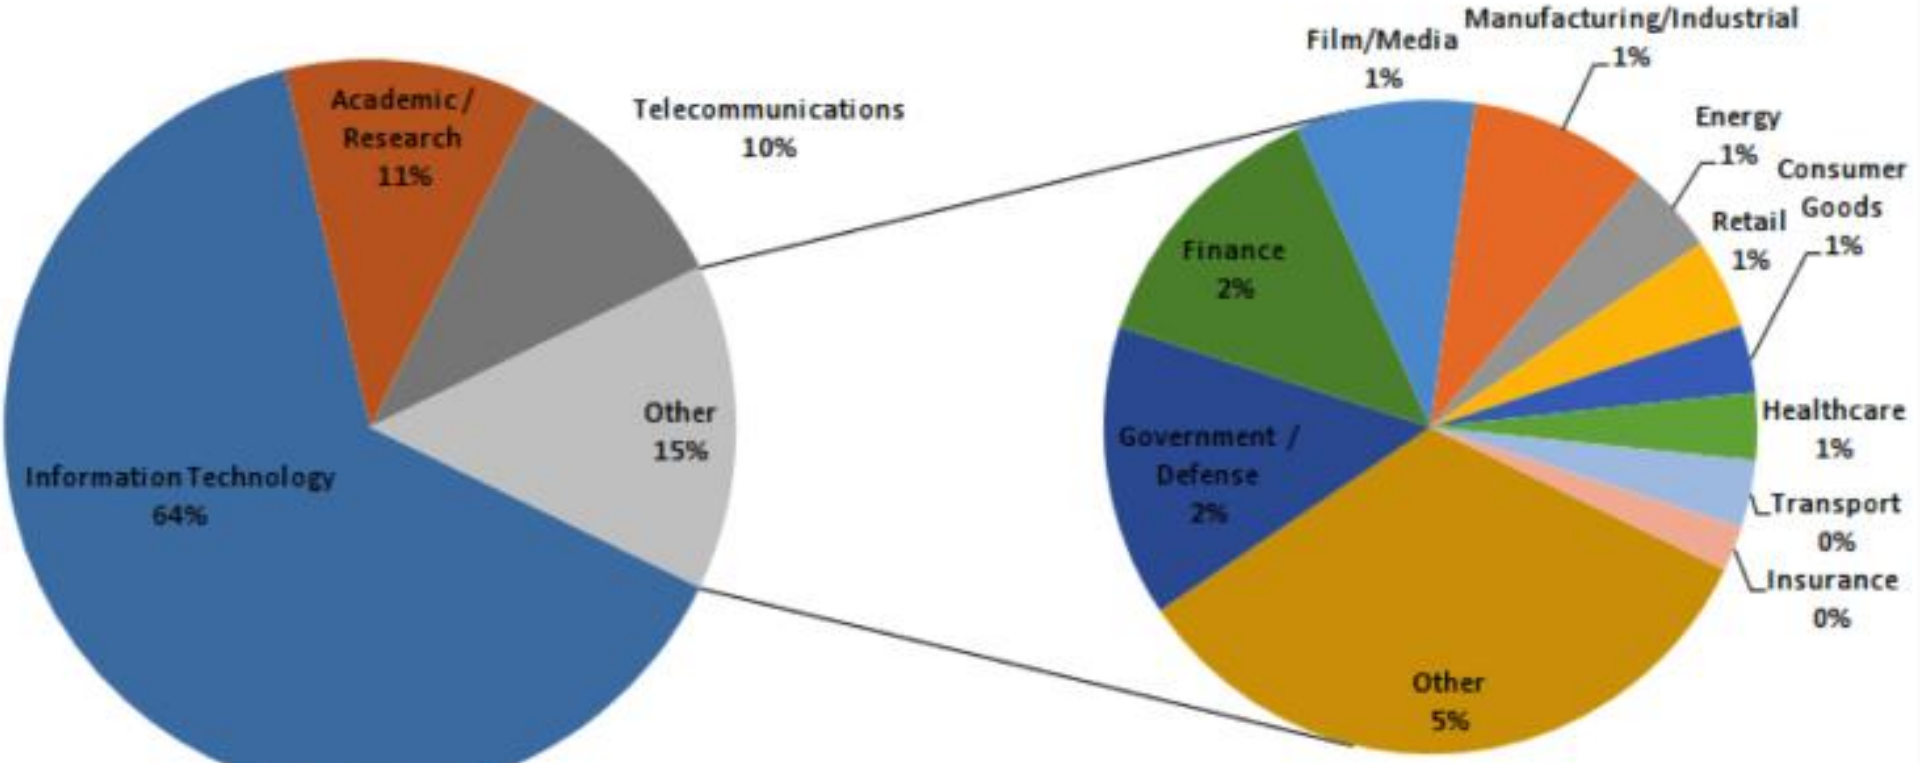

Survey of OpenStack users, Nov 2014 (N=740)

Code	Name	%	Cores
06	Biological Sciences	24	8336
08	Information And Computing Sciences	20	6769
10	Technology	11	3839
09	Engineering	7	2294
03	Chemical Sciences	6	1967
02	Physical Sciences	6	1912
04	Earth Sciences	5	1826
11	Medical And Health Sciences	5	1802
01	Mathematical Sciences	5	1752
14	Economics	5	1591
05	Environmental Sciences	3	1139
20	Language, Communication And Culture	1	274
12	Built Environment And Design	1	254
15	Commerce, Management, Tourism And Services	1	209
17	Psychology And Cognitive Sciences	1	189

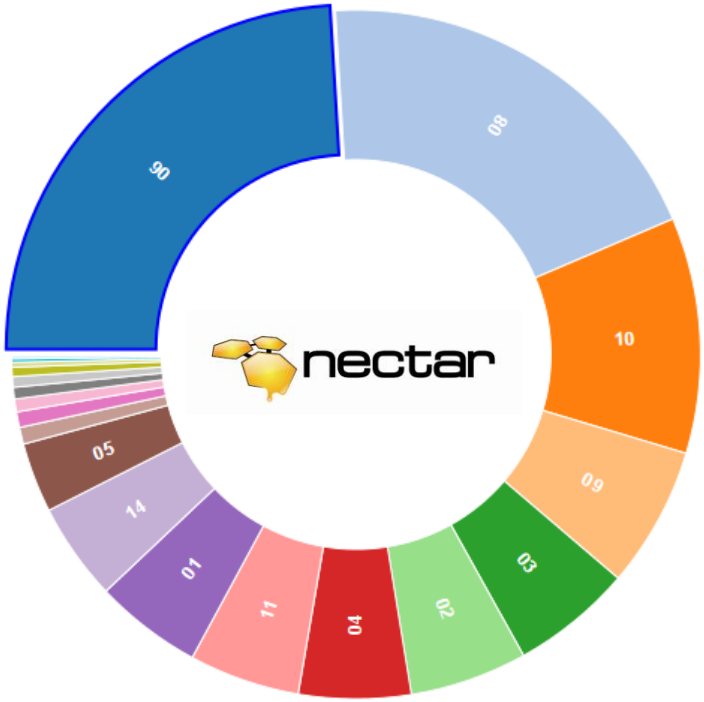

07	Agricultural And Veterinary Sciences	0	173
21	History And Archaeology	0	151
19	Studies In Creative Arts And Writing	0	79
16	Studies In Human Society	0	58

status.rc.nectar.org.au/allocations (N = 2193)

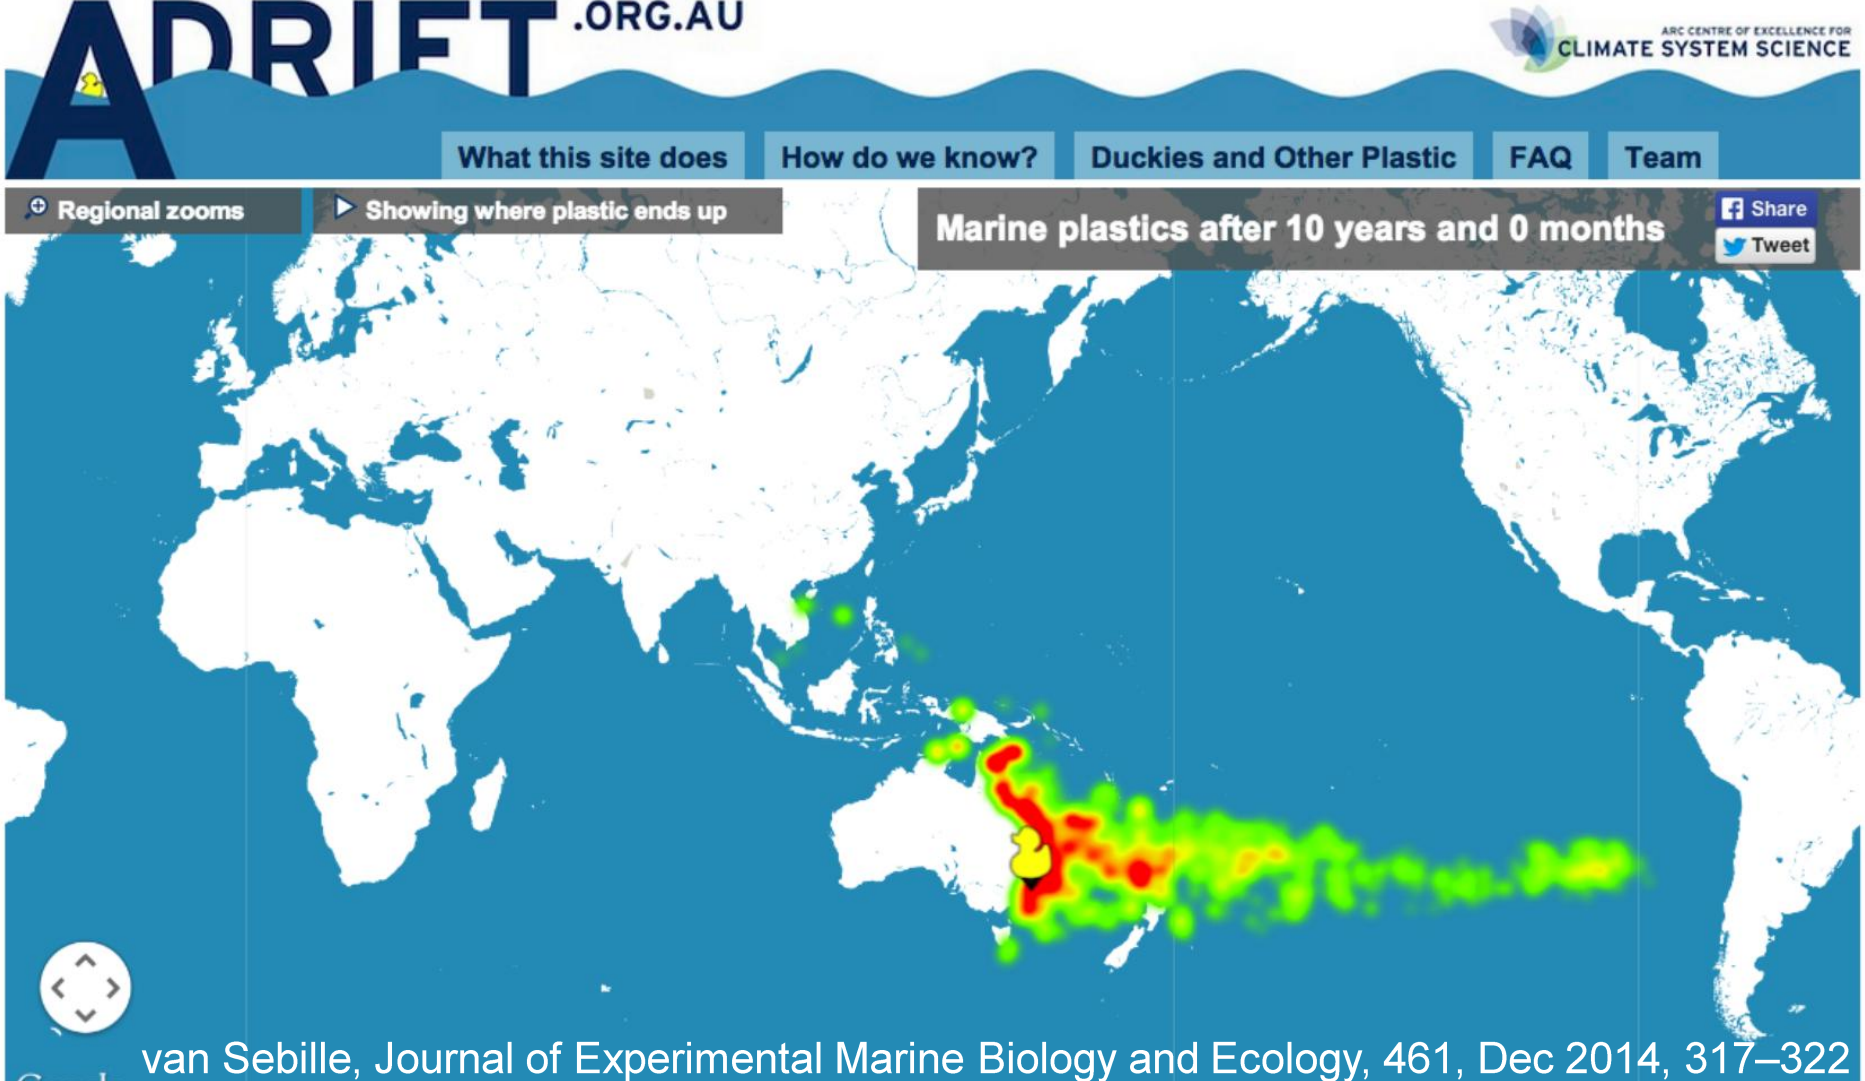

van Sebille, Journal of Experimental Marine Biology and Ecology, 461, Dec 2014, 317–322


COMCAST
openstack

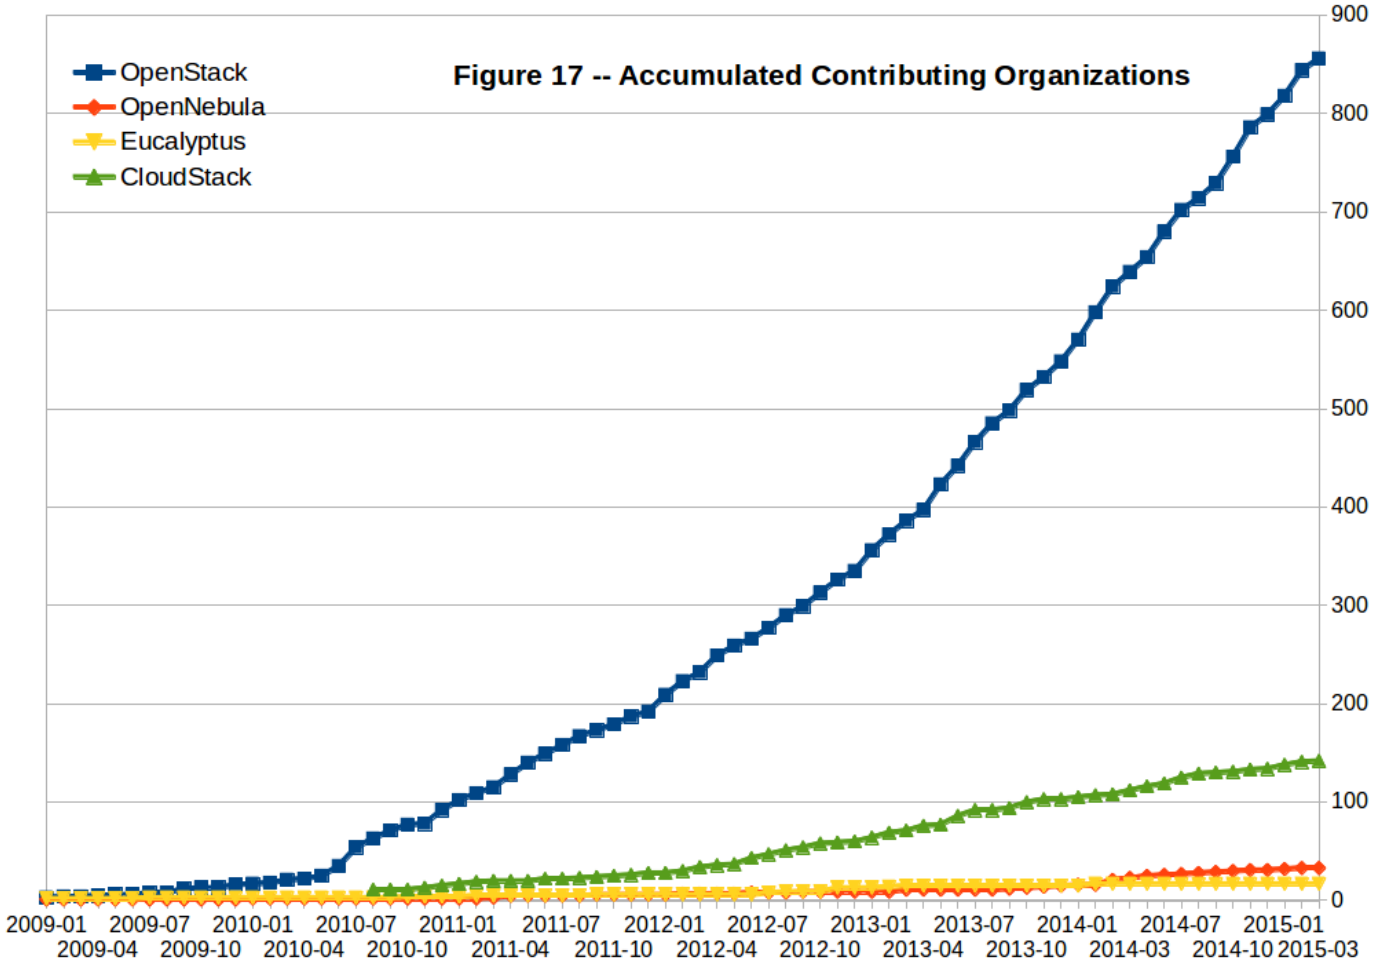

http://www.udpwork.com/item/14057.html - Qingye Jiang (John), University of Sydney, May 2015

Project is very active with many contributors

3,500+
Engineers

22,936

SUPPORTING ORGANISATIONS	AVERAGE CONTRIBUTORS MONTHLY	PATCHES MERGED MONTHLY	TEST CLOUDS BUILT DAILY
490+	660+	7,500+	700+

openstack
CLOUD SOFTWARE

Change-Id:	Ieffeb6cf34b225f0704321fa64fe6dfc227add8e
Owner	Steve Martinelli
Project	openstack/keystone
Branch	master
Topic	bug/1426128
Uploaded	2015-03-10 13:22
Updated	2015-03-28 16:18
Status	**Merged**

IBM

Commit Message *Permalink*

Add API to create ecp wrapped saml assertion

Create a new API that gives users the option to wrap a token based SAML assertion in an ECP envelope.

Co-Authored-By: Rodrigo Duarte Sousa <rodrigods@lsd.ufcg.edu.br>

Change-Id: Ieffeb6cf34b225f070432
Closes-Bug: 1426128
bp: ecp-wrapped-saml-assertions

UFCG.EDU.BR

Reviewer	*Code-Review*	*Verified*	*Workflow*
Steve Martinelli			
Haneef Ali			
John Dennis			
ayoung			
Morgan Fainberg			
guang-yee			
Boris Bobrov			
Marek Denis			
Rodrigo Duarte			
Lance Bragstad			
David Chadwick			

HP
Redhat
CERN
Mirantis
UFCG
Rackspace
Kent.ac.uk

Jenkins check	*2015-03-28 15:25*
gate-keystone-pep8	SUCCESS in 4m 02s
gate-keystone-docs	SUCCESS in 4m 23s
gate-keystone-python27	SUCCESS in 9m 38s
check-tempest-dsvm-full	SUCCESS in 44m 25s
check-tempest-dsvm-postgres-full	SUCCESS in 43m 26s
check-tempest-dsvm-neutron-full	SUCCESS in 52m 40s
check-grenade-dsvm	SUCCESS in 39m 18s
gate-tempest-dsvm-large-ops	SUCCESS in 24m 37s
gate-tempest-dsvm-neutron-large-ops	SUCCESS in 20m 13s
check-swift-dsvm-functional	SUCCESS in 22m 25s
gate-ra	

https://review.openstack.org/#/c/162866/

More than can fit in one presentation.

OpenStack Examples

Pioneer

ANL

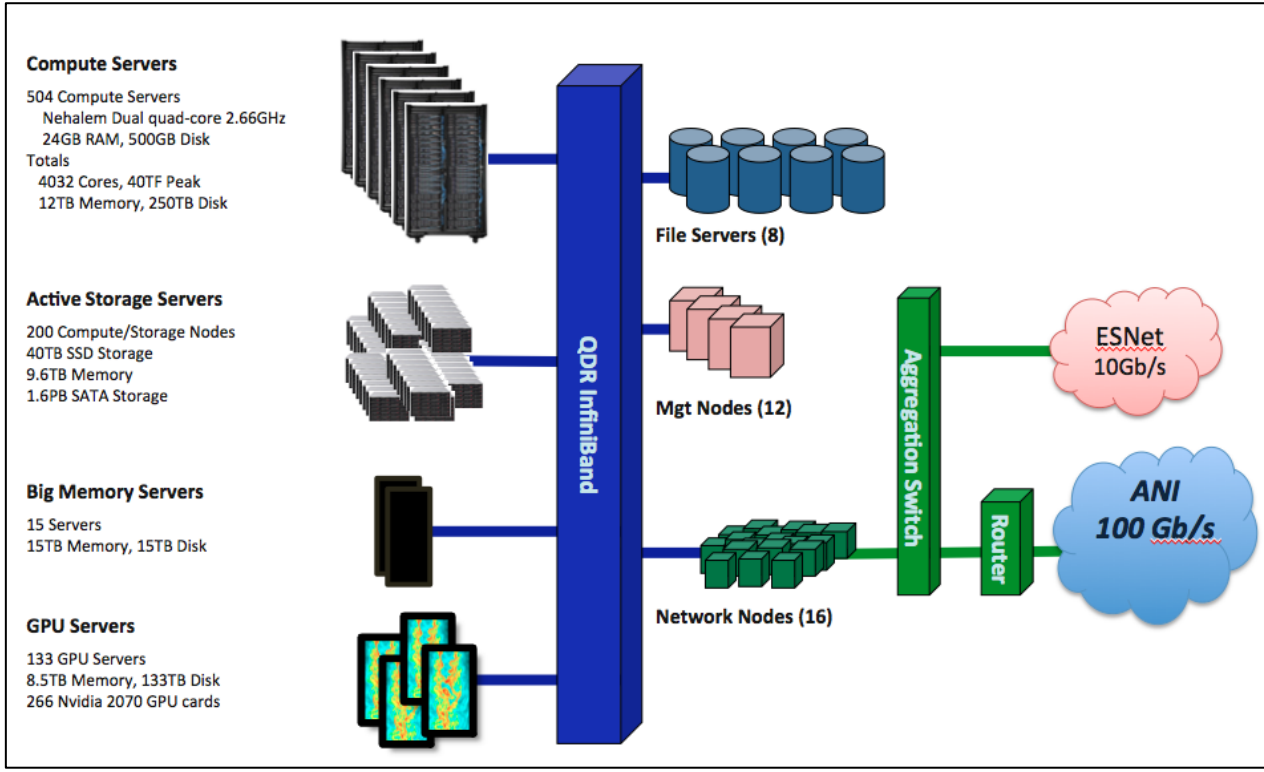
Compute Servers
504 Compute Servers
Nehalem Dual quad-core 2.66GHz
24GB RAM, 500GB Disk
Totals
4032 Cores, 40TF Peak
12TB Memory, 250TB Disk
Active Storage Servers
200 Compute/Storage Nodes
40TB SSD Storage
9.6TB Memory
1.6PB SATA Storage
Big Memory Servers
15 Servers
15TB Memory, 15TB Disk
GPU Servers
133 GPU Servers
8.5TB Memory, 133TB Disk
266 Nvidia 2070 GPU cards
QDR InfiniBand
File Servers (8)
Mgt Nodes (12)
Network Nodes (16)
Aggregation Switch
Router
ESNet
10Gb/s
ANI
100 Gb/s

Fast Switch!

ATLAS P1

Sim@P1

- The Simulation in Point1 Project, based on OpenStack, uses in an opportunistic way the resources of the TDAQ High Level Trigger (HLT) farm of the ATLAS experiment.
- More than 1300 compute nodes (CNs) running up to 2700 VMs are exploited for running event generation and Monte Carlo production jobs, mostly CPU and not I/O bound, for a total of up to 22K parallel running jobs

Mixed use

BNL

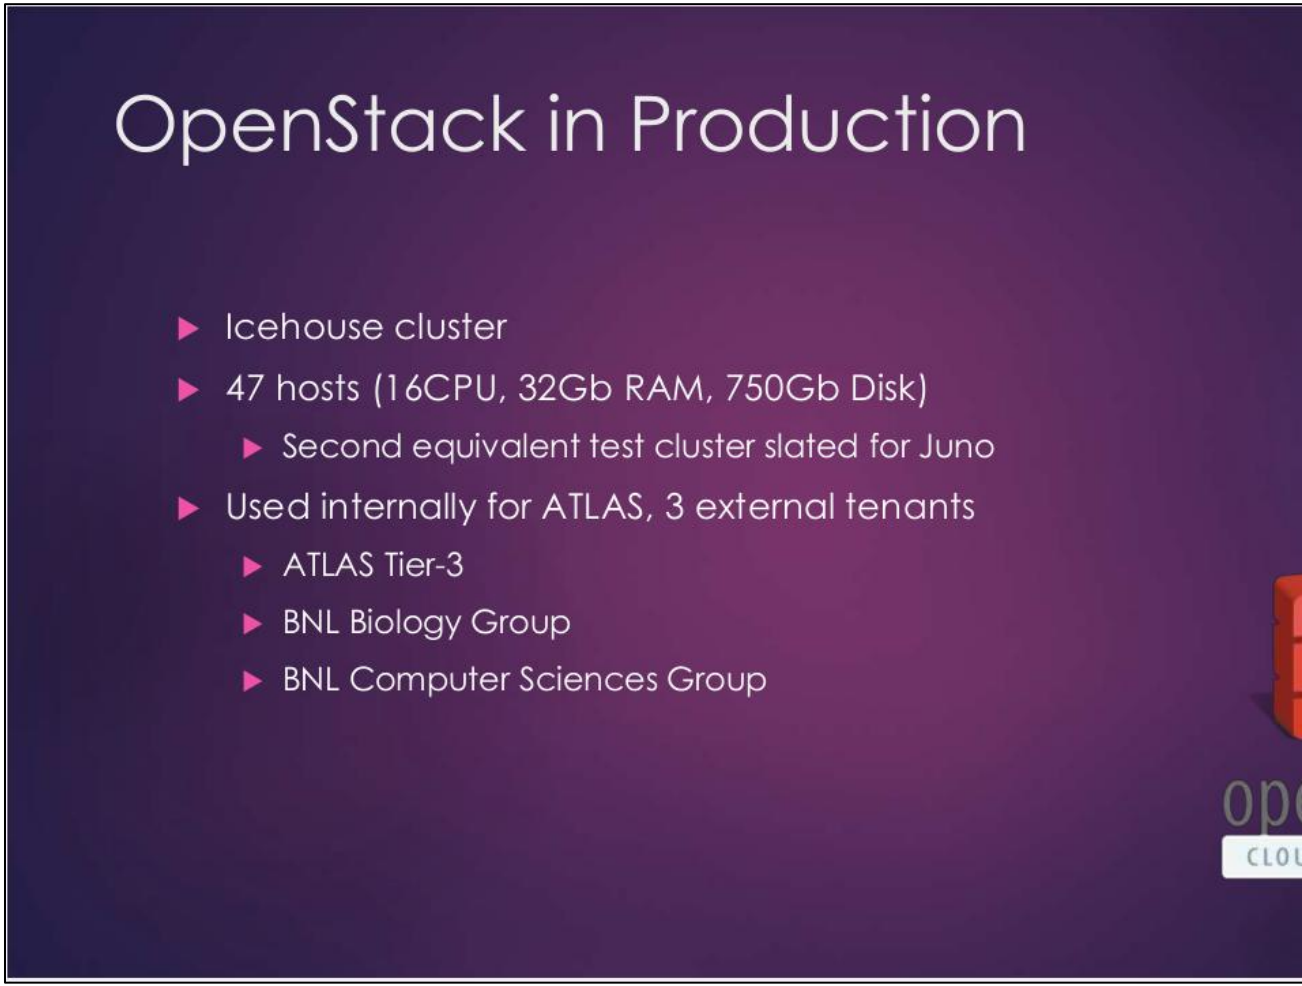

Enterprise+Research

Canada

WestGrid

compute • calcul CANADA

About | Resources | Support | Research

Home » About WestGrid » News & Media » Other HPC News » Cybera Offering Cloud Computing at No Cost to Alberta Researchers

Cybera Offering Cloud Computing at No Cost to Alberta Researchers

Up to one year of free access to cloud computing is now available to Alberta researchers and entrepreneurs through a pilot program designed to advance the use of technology in the province. Research and education organizations, and entrepreneurs, are the target users of the cloud services being offered by Alberta's not-for-profit advanced technology agency, Cybera. The service, called the Rapid Access Cloud, is available for prototyping and research.

"Most if not all of the major technology companies are using the cloud because it offers huge competitive advantages," says Robin Winsor, President and CEO of Cybera. "However, researchers and entrepreneurs – people trying out brand new ideas – might find the cost of this technology to be a barrier when they aren't yet sure what the return on investment will be. This program can provide that environment for discovery and invention."

Test, Host, Compute

CCIN2P3

Cloud services (Openstack Hardware)

Deployment :

- Scientific Linux 6 (requirement for 7 in Kilo release)
- Griddynamics, then EPEL and now RDO
- Configured with Puppet

Resources :

Bunch of C6xx, R6xx, M6xx DELL PowerEdge servers

Core services
- 150 CPUs
- 300GB RAM
- 5TB storage

Hosting
- 50 CPUs
- 200GB RAM
- 4TB storage

Computing
- 500 CPUs
- 3TB RAM
- 9TB storage

Preprod cluster :
- 300 CPUs
- 1TB RAM
- 28TB storage

Grand total:
- 1k CPUs
- 4.5TB Memory
- 100TB storage

+30 TB Cinder volumes

+24 TB Swift S3 storage

Non-HPC workloads

CSC

Pouta's use cases

- Enhanced security – isolated virtual machines
- Advanced users – able to manage servers
- Difficult workflows – can't run on Taito
- Complex software stacks
- Ready made virtual machine images
- Deploying tools with web interfaces
- "We need root access"

If you can run on Taito – run on Taito
If not – Pouta might be for you

- Pouta user guide: https://research.csc.fi/pouta-user-guide

All the things

CERN

CERN Private Cloud - Numbers

- Based on OpenStack Juno
- Spans between 2 datacentres
- 4700 hypervisors
  - 120000 cores
- 11000 VMs
- 1500 users
- 1800 projects

All the things

CERN

CERN Private Cloud - Numbers

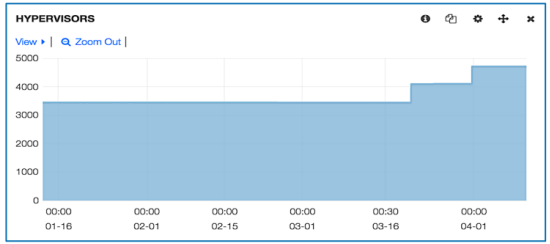

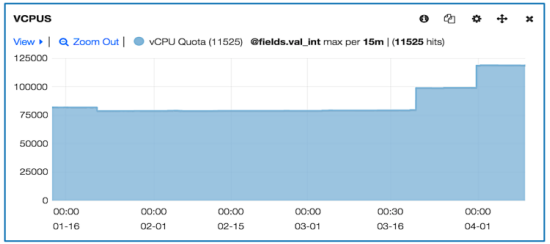

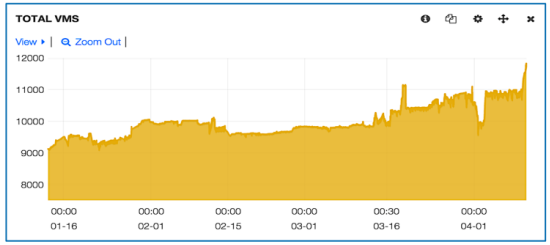

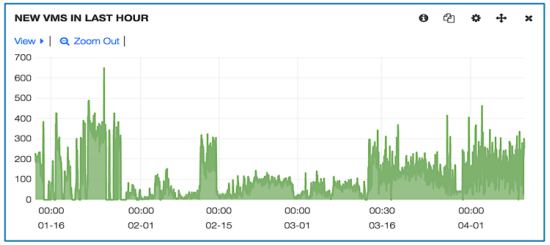

Look at the GUI

IHEP

IHEP Cloud established

- Released on 18th. November, 2014
- Built on openstack Icehouse
- 8 physical machines: 224 vm capacity
  - 1 control node, 7 computing nodes
  - User applies and gets the VM on line
- Three types of VM provided
  - Provide user VM that same as login node
    - DNS and IP management , Email and AFS account, puppet, NMS, ganglia …
  - Provide user VM with root right and no public IP
  - Provide administrator VM with root right and public IP
- Current Status
  - Active 172 VM, 628GB memory and 4.7TB disk

10

HTC

KIT

Enabling Shared Remote HPC Resources for HEP

- German funding agencies encourage sharing of computing resources among user groups
- Acquiring funding for dedicated HEP-only installations becomes more difficult

bwForCluster cluster installation case study

- Located at Freiburg University ∼150km south of Karlsruhe
- Shared by 3 diverse user groups: Particle Physics, Neuroscience, Microsystems technology
- State funding secured, full installation with ∼8000 cores is expected to be available in Fall 2015
- Virtualization is a key technology to allow for a efficient sharing among the user groups
- A 10Gbit dedicated link allows for an efficient data transfer (possible to upgrade to 100Gbit) to GridKa Tier-1 site
- Current prototype installation: system similar to the final bwForCluster installation
- Fully functional OpenStack setup with 800 cores

Beyond my research laboratory, this thesis heavily relied on a 1,500-core OpenStack computer cluster and I thank the patience of MIT CSAIL technical members Jonathan Proulx and Stephen Jahl for answering to my dozens of bug reports and other miscellaneous issues, and the generosity of Quanta Computer, who donated a large part of the cluster hardware. Many thanks as well to Garrett Wollman for his Unix expertise and NFS server skills.

In the thesis

MIT CSAIL

openstack
CLOUD SOFTWARE

Infiniband

NCI (AU)

NCI InfiniCloud: HPC InfiniBand performance in the Cloud

- OpenStack Cloud supported by a 56Gbs FDR InfiniBand fabric
  - High Performance IB MPI
  - High Performance Lustre File System access
- Built using Mellanox Neutron modules
- Flexible top performance computational resources ‘on demand’
- HPC quality, Curated and Managed OS and application stack supported by an I/O architecture which includes local Solid State Disk and high performance large file storage using Lustre

nci.org.au

Research Pipeline

OSDC

name / description	use	cloud software	compute nodes	Compute Cores	Compute RAM (GB)	RAW Storage (TB)	Usable Storage (TB)
Sullivan	Compute - General Projects + Selected Groups	OpenStack w/ GlusterFS	60	1056	4736	1096	435
Atwood	Protected Compute Cloud	OpenStack w/ GlusterFS	23	312	736	476	181
Bionimbus-PDC	Protected Compute Cloud for TCGA Data	OpenStack w/ Ceph and Swift	100	1208	9856	5736	276
Skidmore	Compute, Selected Projects	Hadoop	25	704	3200	756	295
OCC-Y	Compute, Selected Projects	Hadoop	61	928	3184	1464	1101
Public Data Commons	Public Data	GlusterFS	0	0	0	1440	967

openstack
CLOUD SOFTWARE

A national project

SWITCH

The Infrastructure

- Two locations (University DCs Lausanne/Zurich), each with
  - 32 2RU dual-Xeon (E5-2650v2) + 128 GB RAM + 2*10GE + 2*SSD
  - 16 servers also have 12*4TB 3.5” disks ➔ Ceph OSDs
  - 2*48-port 10GE (+6-port 40GE) switches + 1*48-port GigE
  - Uplink: 2*10GE w/BGP-4 (IPv4+IPv6) directly to backbone
- Currently two racks used per location
  - Each can grow up to ~20 racks
- Plus a staging setup with two (tiny) sites
  - in one of the two production locations

#1 SuperComputer

Tian He 2

Background

- Hardware for OpenStack deployment
  - 6400 nodes on TH-2
    - 50 cabinets
    - 128 nodes in each cabinet
  - Each node:
    - Intel Xeon Ivy Bridge (12 cores) * 2
    - 96 GB RAM
    - 1TB local disk
    - GE Nic * 2
    - TH-NI High-speed Network (160 Gbps)
    - Intel Xeon Phi * 3

Staff needed!

UNL

Future Plans

- OpenStack
  (because now we have to)

  Controllers (2x R710)
  Network (2x R710)
  Storage (5x Sun X4275 w/20TB)
  Compute (8x Sun X2200, 8 core, 32GB)

- $630k == Many, many unknowns
  (anyone looking for a job?)

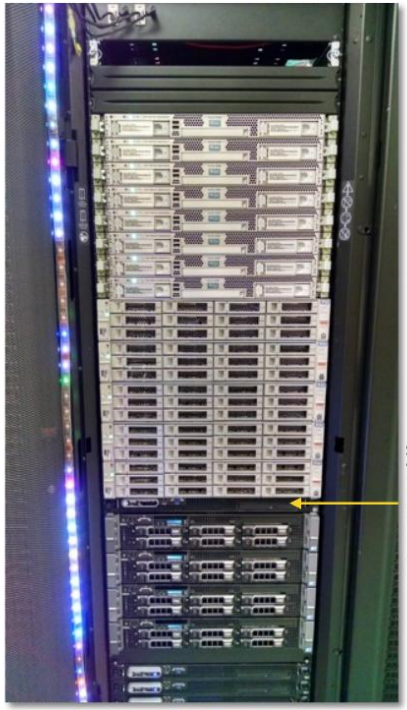

Researching OpenStack

ZHAW ICC

Making Openstack more Energy Efficient...

...a little story from your friends at ICCLab...

We've all got other things to do.

Operational Cost

Your OpenStack Journey

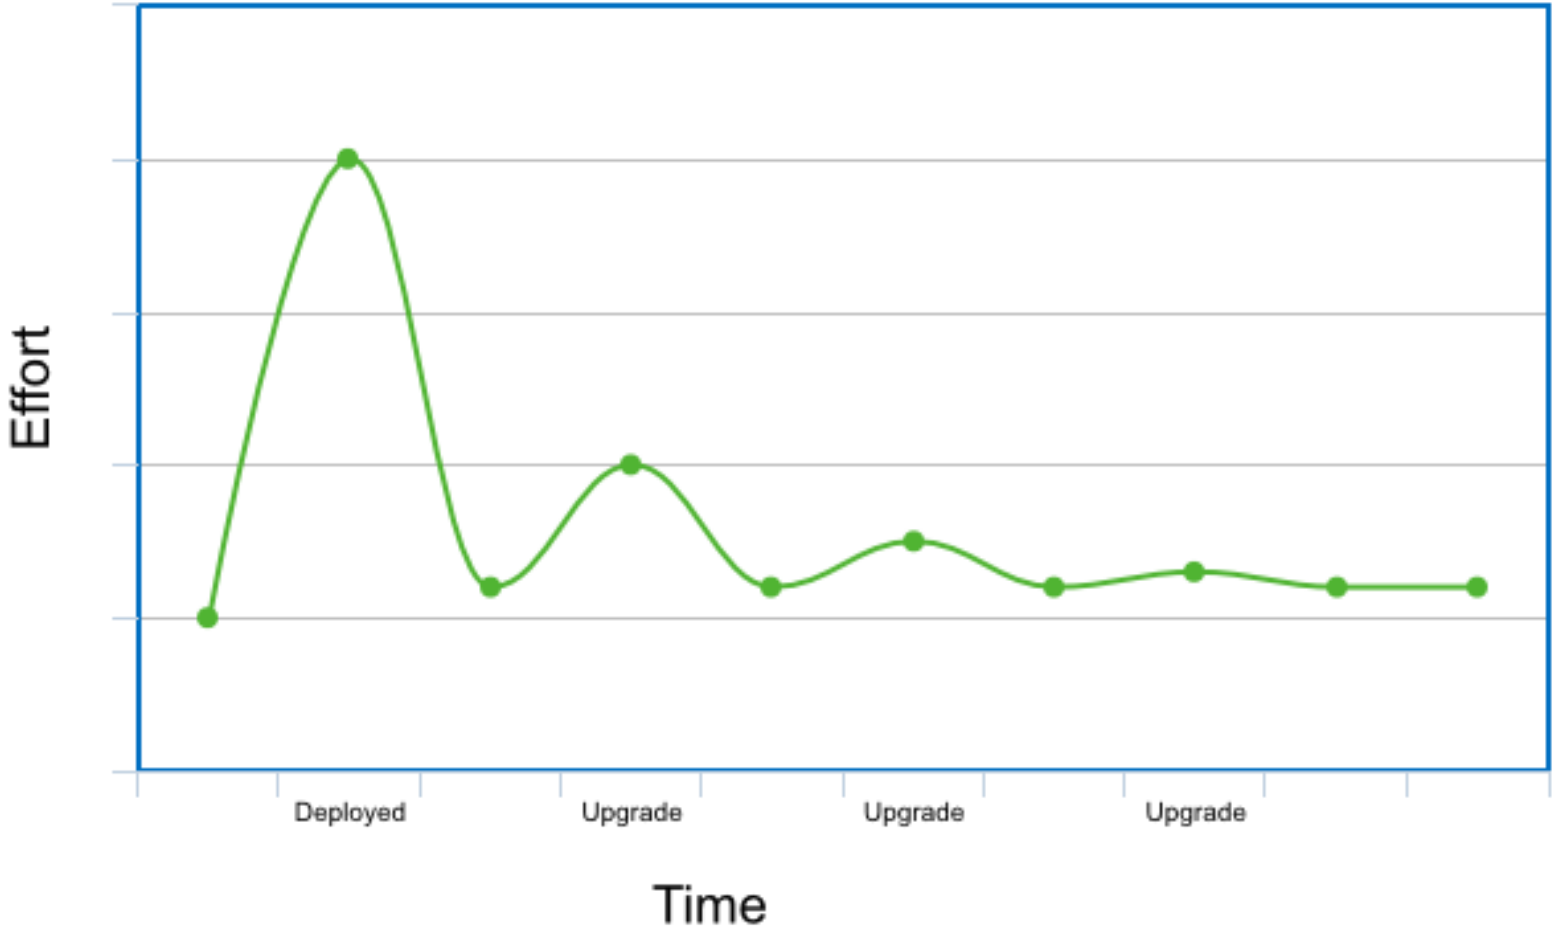

openstack
CLOUD SOFTWARE

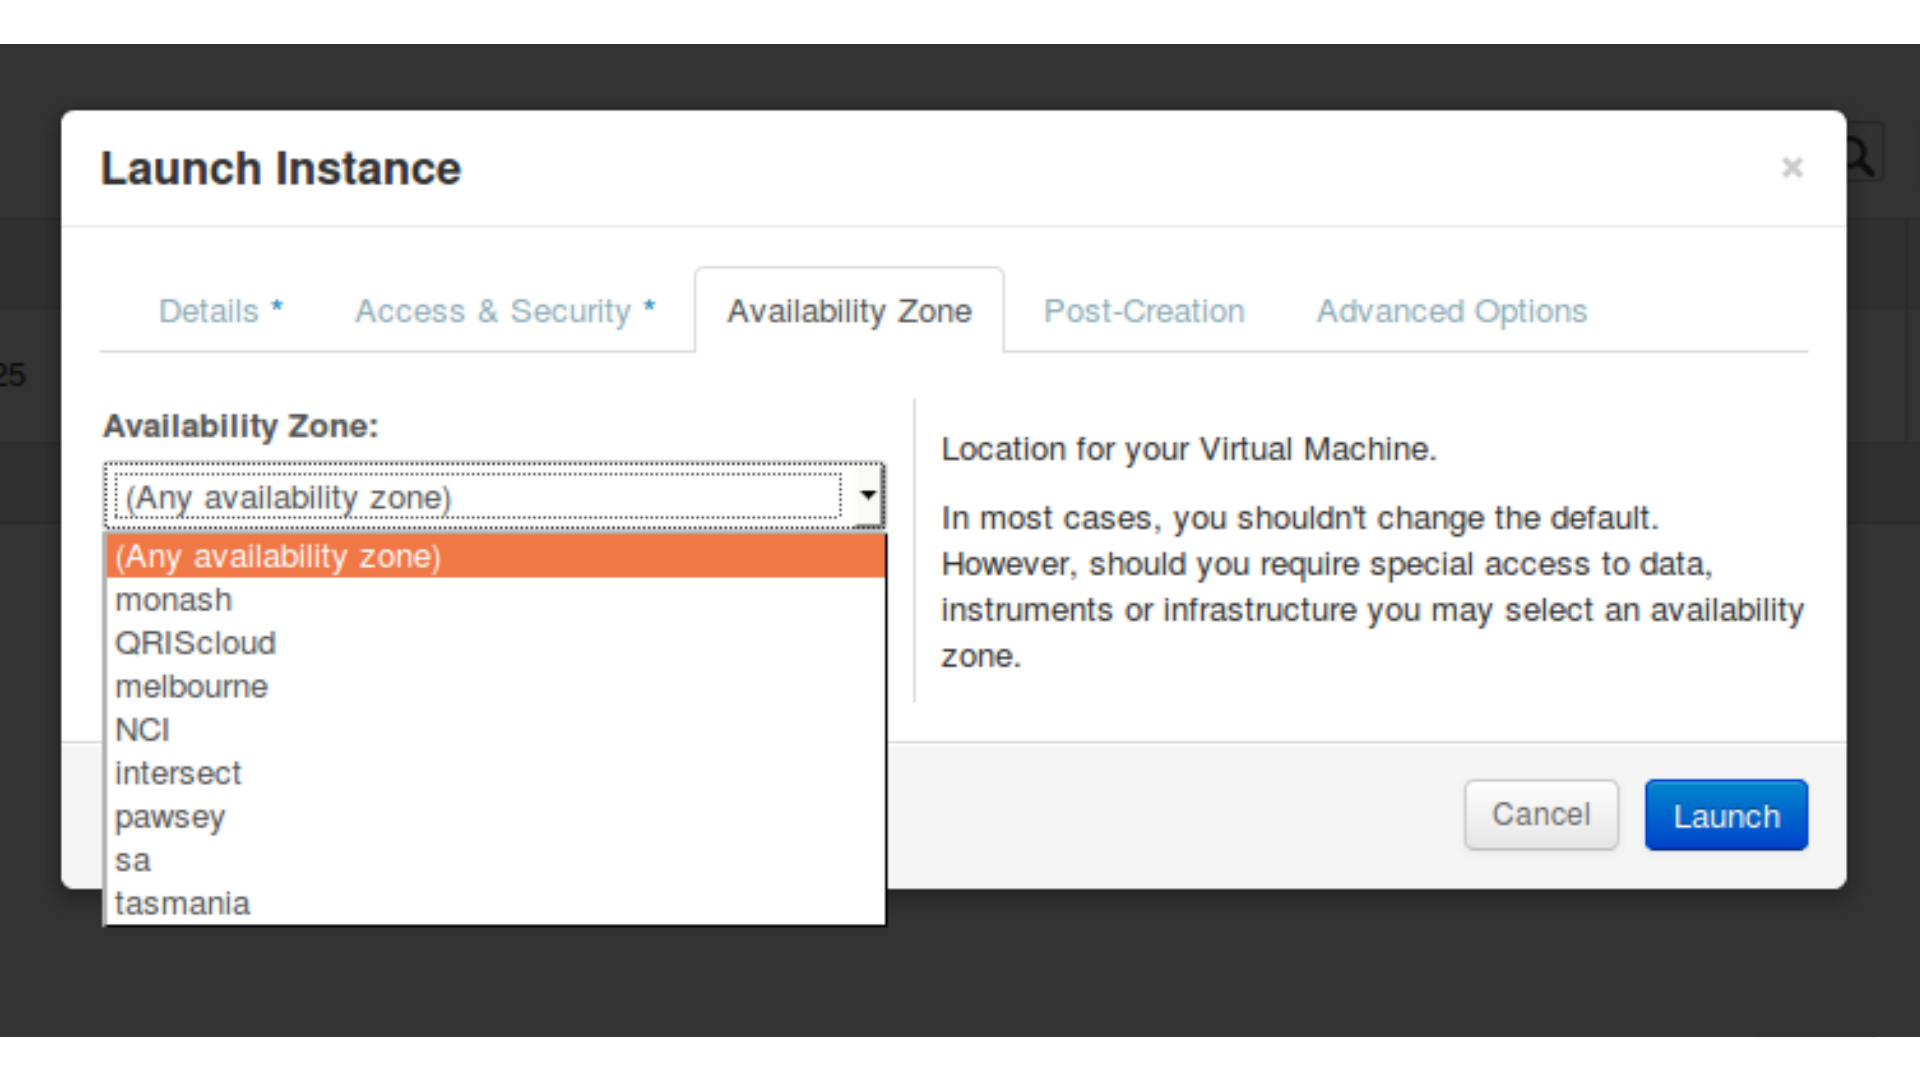
Launch Instance
Details *
Access & Security *
Availability Zone
Post-Creation
Advanced Options
Availability Zone:
(Any availability zone)
(Any availability zone)
monash
QRIScloud
melbourne
NCI
intersect
pawsey
sa
tasmania
Location for your Virtual Machine.
In most cases, you shouldn't change the default. However, should you require special access to data, instruments or infrastructure you may select an availability zone.
Cancel
Launch

Implementation/2

- Computing services:

Location	# of servers	CPU Model	# of CPUs	RAM (GB)	# of (HT) CPU-cores	OpenStack role
PD	4	E5-2609	2	32	8	Controller, Network nodes
PD	5	E5-2670v2	2	96	40	Compute nodes
PD	3	E5-2650v3	2	96	40	Compute nodes
LNL	6	E5-2650v2	2	96	32	Compute nodes
Total	**18**		**32**	**1472**	**544**	

- Storage:
  - ✓ 1 server iSCSI DELL MD3620i with 68TB at Padova for OpenStack services
  - ✓ 1 FibreChannel DELL PowerVault MD3600F with 48TB at LNL for Posix storage
  - ✓ Configured with GlusterFS for images (Glance service), VM ephemeral storage (Nova service) and block storage (Cinder service)

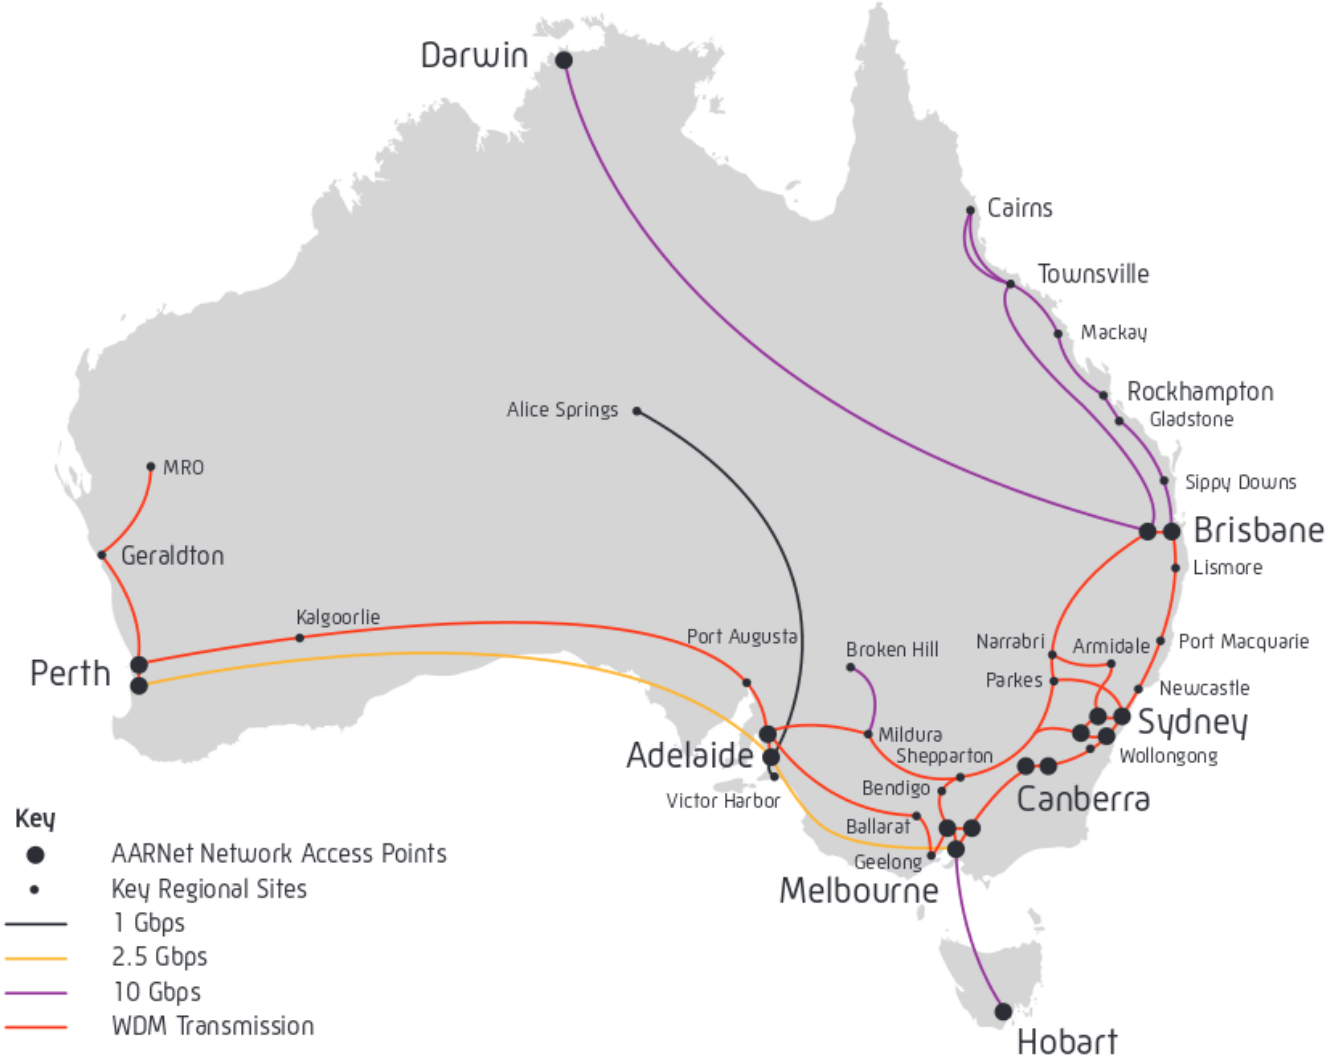
Darwin
Cairns
Townsville
Mackay
Rockhampton
Gladstone
Alice Springs
MRO
Sippy Downs
Brisbane
Geraldton
Lismore
Kalgoorlie
Port Augusta
Broken Hill
Narrabri
Armidale
Port Macquarie
Perth
Parkes
Newcastle
Sydney
Mildura
Shepparton
Wollongong
Adelaide
Bendigo
Canberra
Victor Harbor
Ballarat
Geelong
Melbourne
Hobart
Key
AARNet Network Access Points
Key Regional Sites
1 Gbps
2.5 Gbps
10 Gbps
WDM Transmission


Job Trends from Indeed.com
openstack
cloudstack
eucalyptus
Percentage of Matching Job Postings
0.04
0.02
0
Jan '06
Jan '07
Jan '08
Jan '09
Jan '10
Jan '11
Jan '12
Jan '13
Jan '14
Jan '1
indeed
one search. all jobs.


devs
openstack
summit
users
SAVE THE DATE! THE NEXT OPENSTACK SUMMIT
VANCOUVER, BC
May 18 - 22, 2015
STAY TUNED FOR DETAILS


Quijote
BA GU S
BiDi
CLUB ZOO

Thank you for supporting OpenStack

Ask Questions at ask.openstack.org

Tom Fifield
tom@openstack.org
@TomFifield