



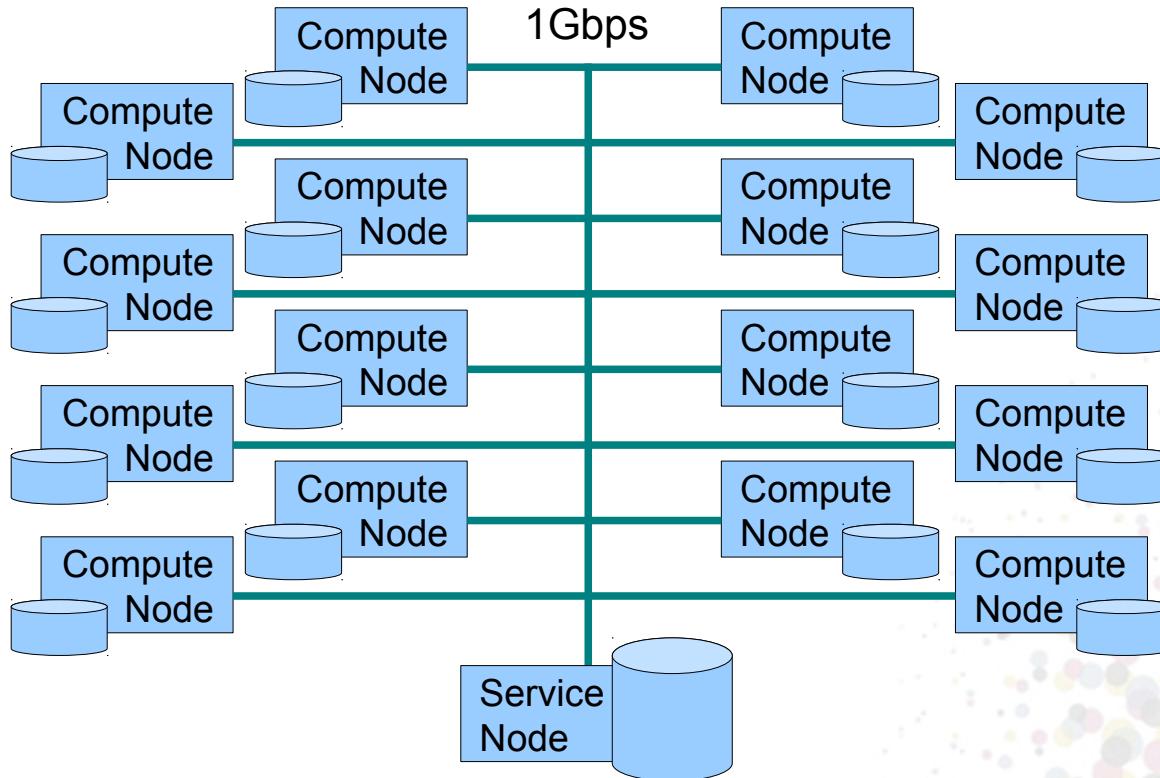
The SARA/BigGrid HPC Cloud



“Calligo” From Beta to Production

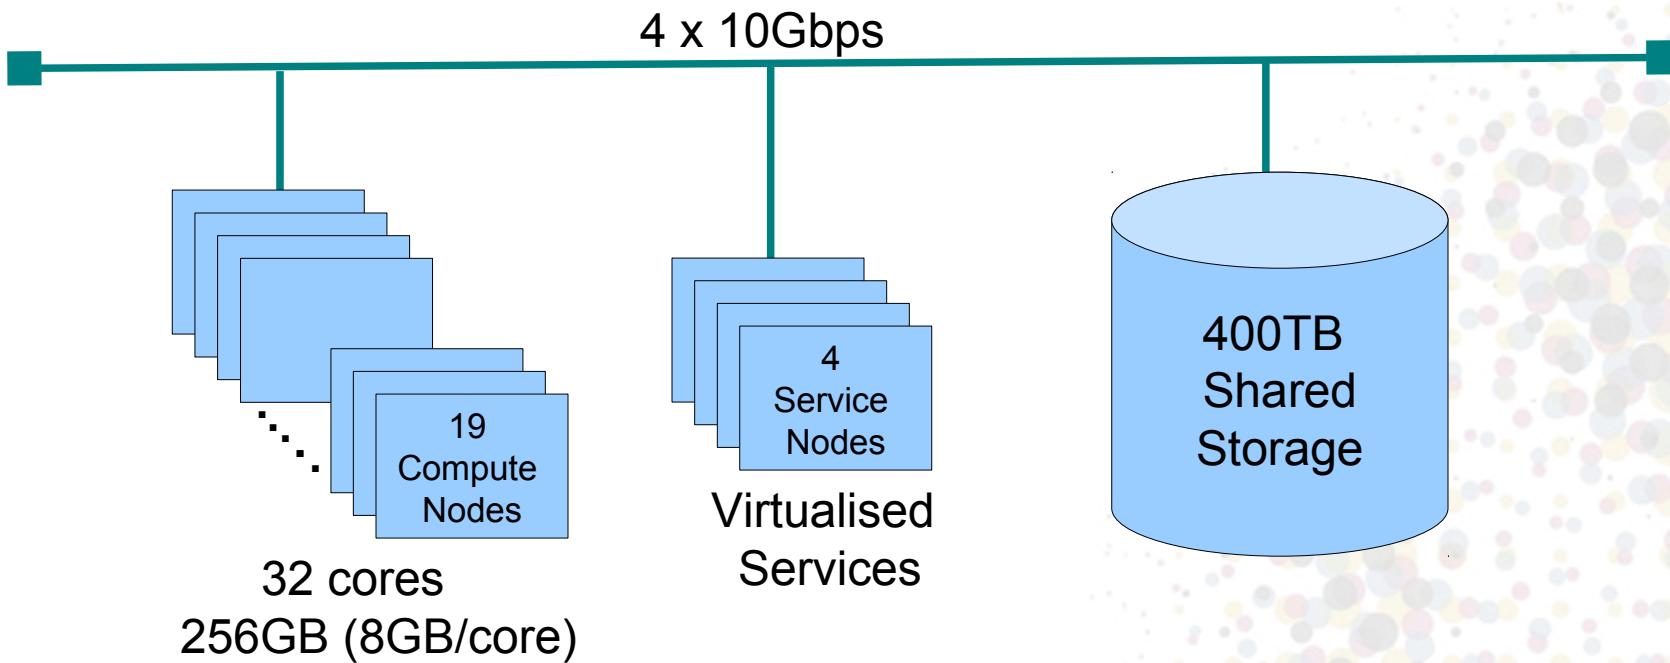
Jhon Masschelein
Jhon.Masschelein@SARA.nl

HPC Cloud – Beta Cloud

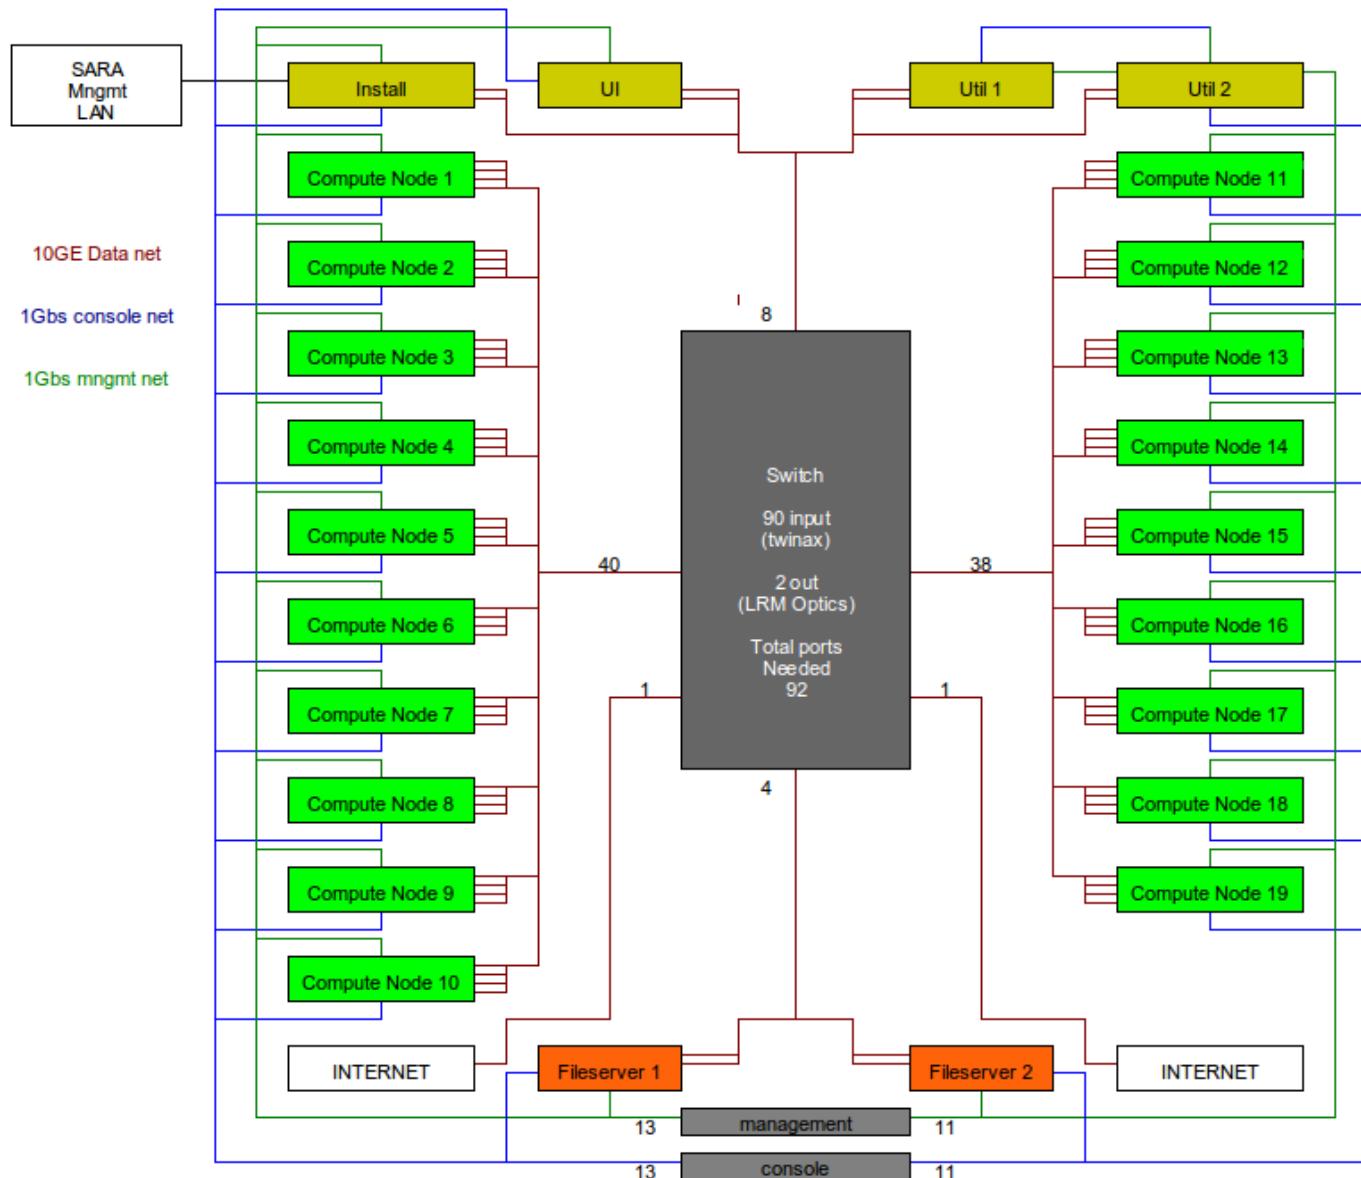


HPC Cloud – Design Choices

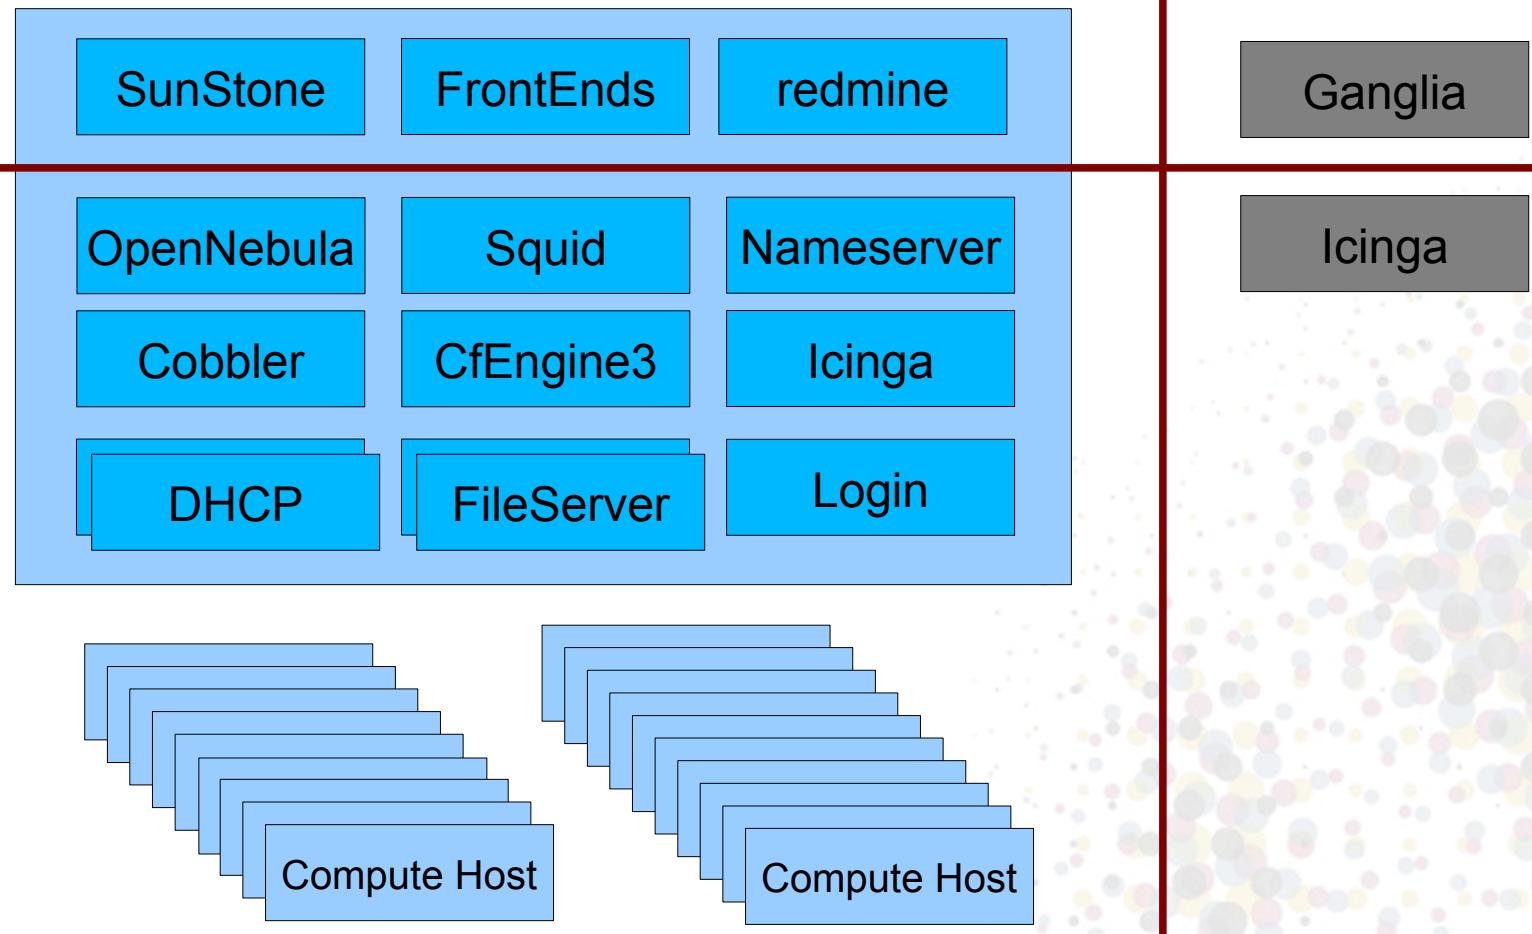
- Driven by beta user comments:
 - More cores and memory per VM
 - Faster start up and shutdown of VMs
 - ▶ make persistent images usefull
 - MORE! MORE! MORE!



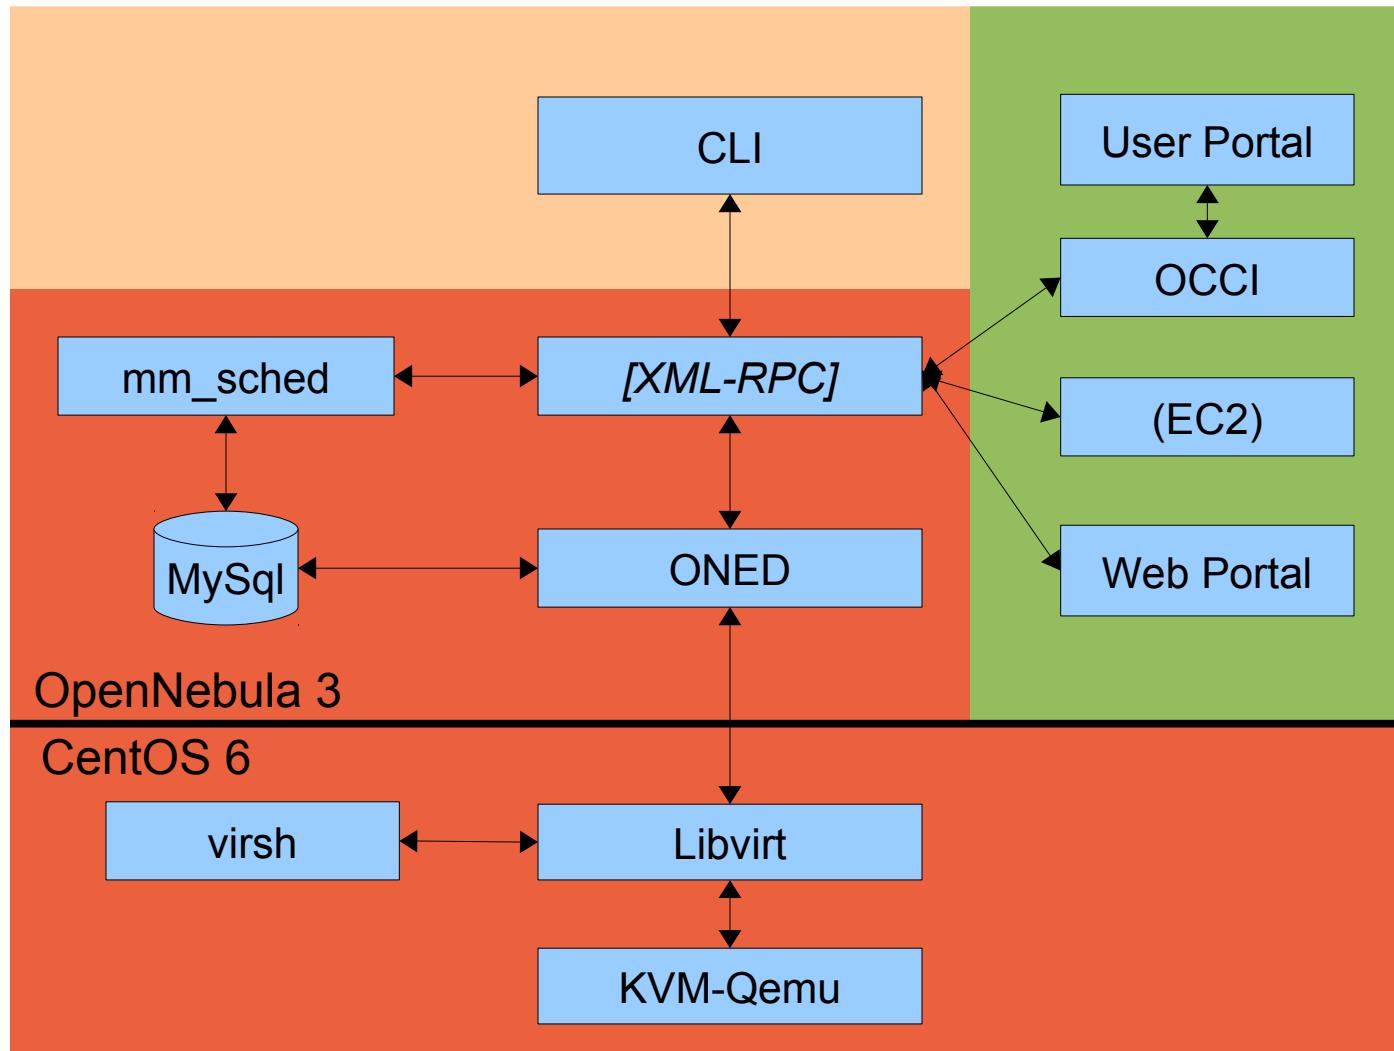
HPC Cloud – Network Detail



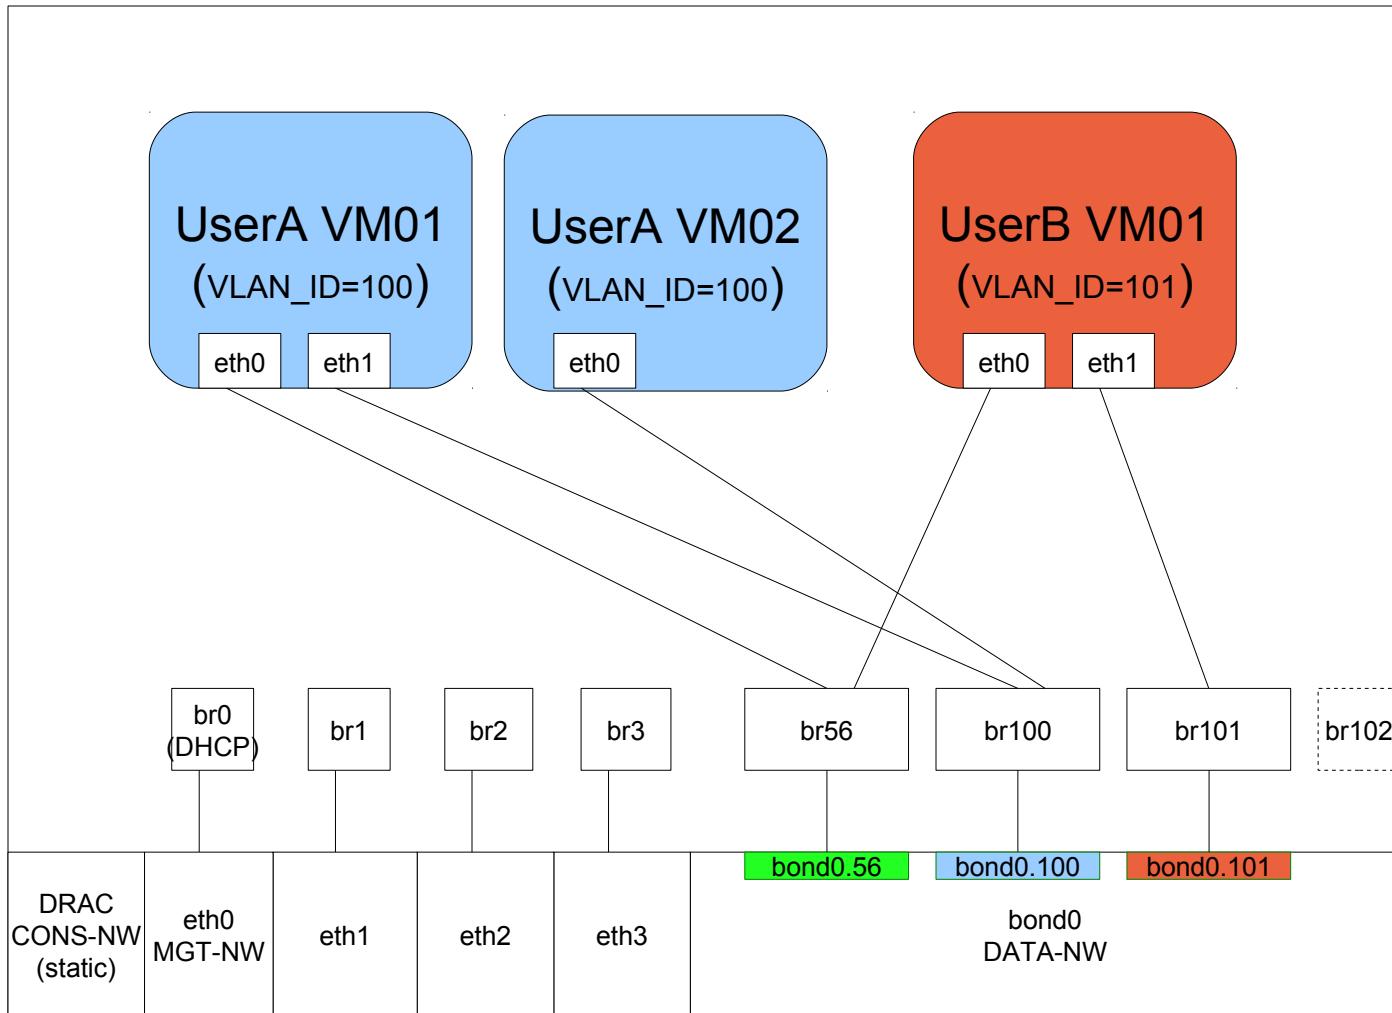
HPC Cloud - Layout



HPC Cloud – Software



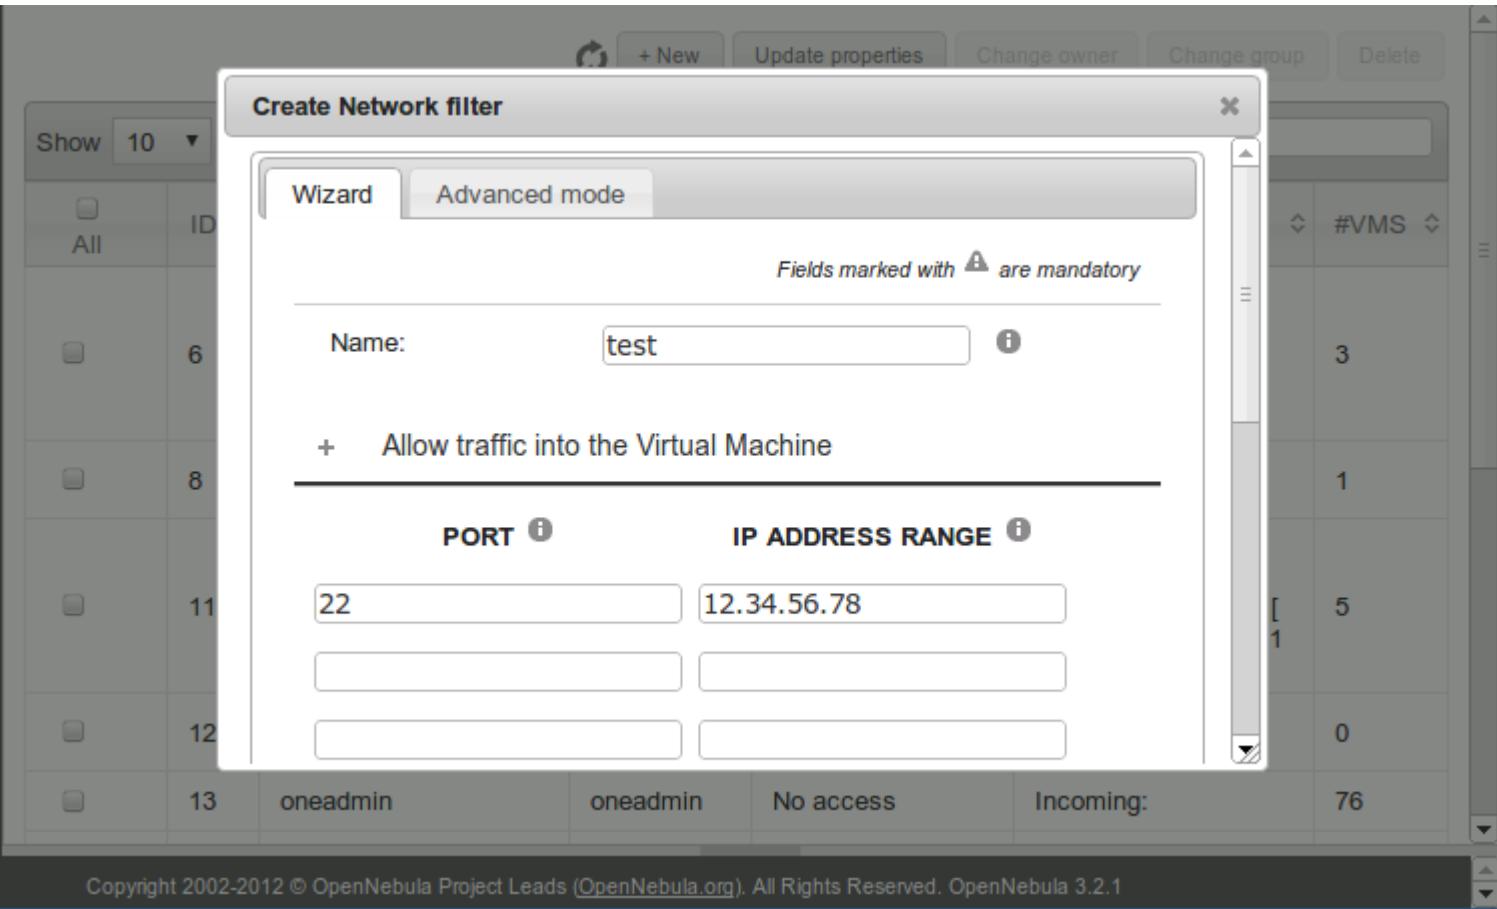
HPC Cloud – VM network



Slide by courtesy of Alain van Hoof (@SARA.nl)

HPC Cloud – Firewall Issue

- OpenNebula Virtual Network template attribute
 - ▶ LIBVIRT_FILTER=calligo-public-network

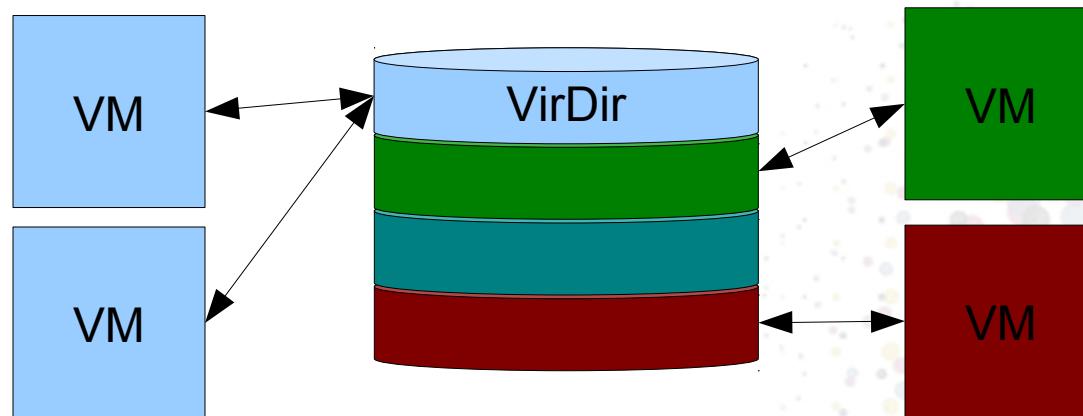


The screenshot shows the OpenNebula interface for creating a network filter. The main window displays a table of existing filters, with columns for ID, Name, and various status metrics like #VMS, #VMs, and various counters (e.g., 3, 1, 5, 0, 76). A modal dialog titled "Create Network filter" is open, showing the "Advanced mode" tab selected. It prompts for a "Name" (set to "test") and a rule description "Allow traffic into the Virtual Machine". Below this, there are fields for defining a port (22) and an IP address range (12.34.56.78). At the bottom of the dialog, there are fields for "oneadmin" and "No access" under "Incoming". The footer of the interface includes the copyright notice "Copyright 2002-2012 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved. OpenNebula 3.2.1".

HPC Cloud – Storage Issue

- Choice for shared storage makes disk IO sub-optimal
- Users want to access project data from all VMs
- We want to avoid a lot of copying and duplication of data

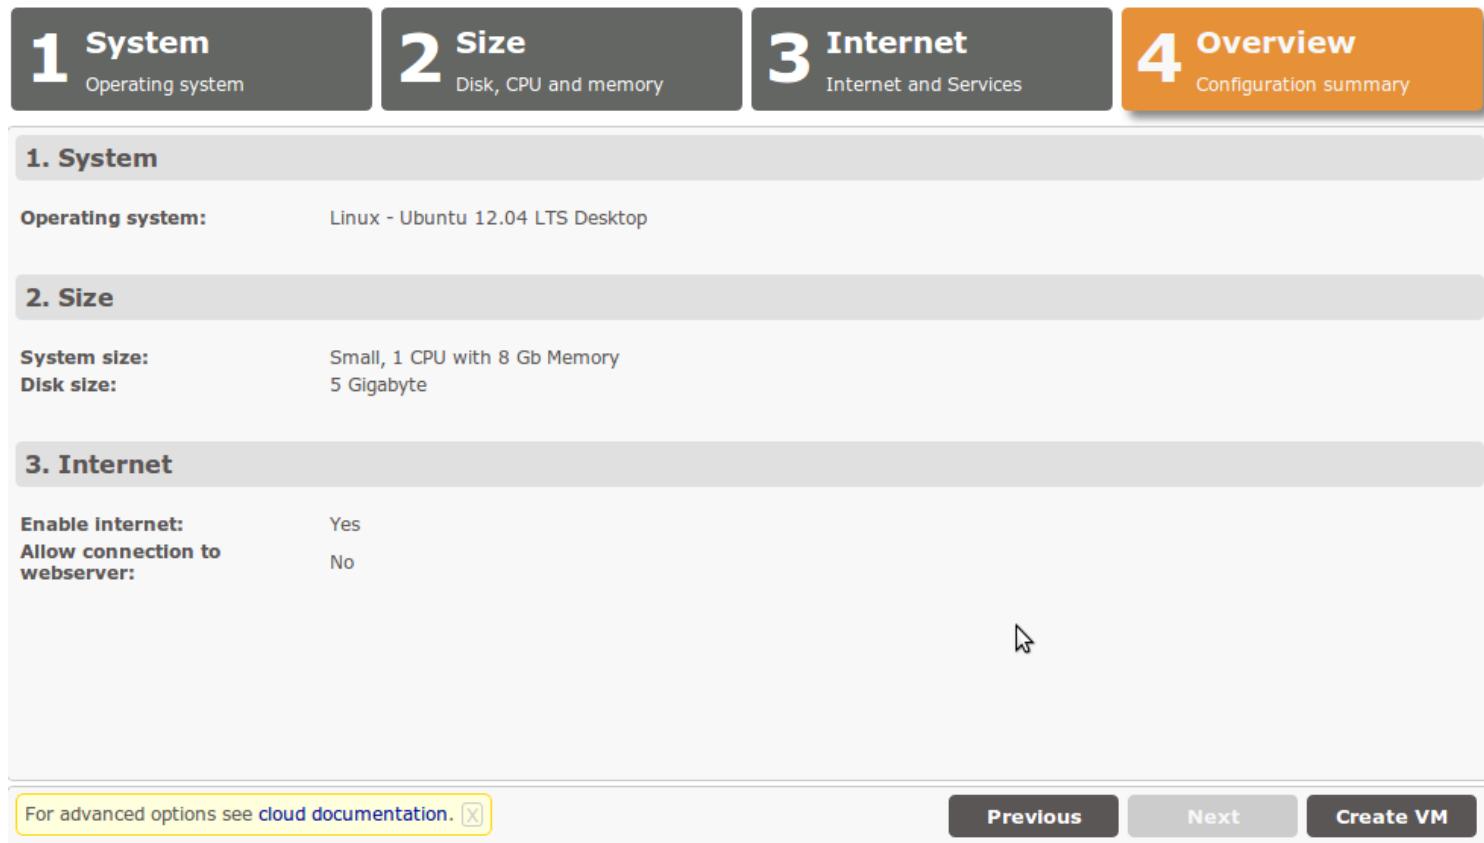
- VirDir (Virtual Directory):
 - NFSv4 shares (dedicated server)
 - “ALL_SQUACH” to project user
 - XFS Quota enabled



HPC Cloud – Ease of use Issue

VM creation Wizard in Sunstone Web Interface

- 4 easy steps to create template & disk images and launch VM
- Prepared Images for CentOS and Ubuntu



The screenshot shows the Sunstone VM creation wizard interface. It consists of four numbered steps: 1. System, 2. Size, 3. Internet, and 4. Overview. Step 4 is highlighted with an orange background.

1. System
Operating system: Linux - Ubuntu 12.04 LTS Desktop

2. Size
System size: Small, 1 CPU with 8 Gb Memory
Disk size: 5 Gigabyte

3. Internet
Enable internet: Yes
Allow connection to webserver: No

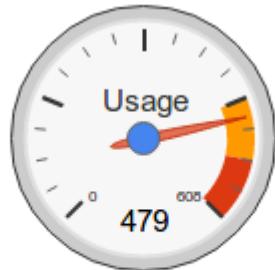
For advanced options see [cloud documentation](#).

Previous Next Create VM

HPC Cloud – Availability Issue

<https://www.cloud.sara.nl>

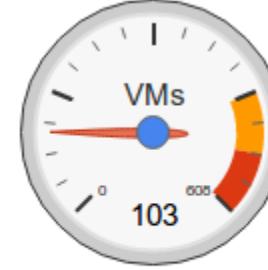
node	used	free	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
d-node01	26	6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node02	31	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node03	22	10	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node04	22	10	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node05	32	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node06	31	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node07	17	15	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node08	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node09	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node10	15	17	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node11	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node12	16	16	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node13	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node14	28	4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node15	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node16	29	3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node17	26	6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node18	27	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
d-node19	22	10	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	



A total of 479 cores out of a maximum of 608 cpu cores are in use at this time. (78%).

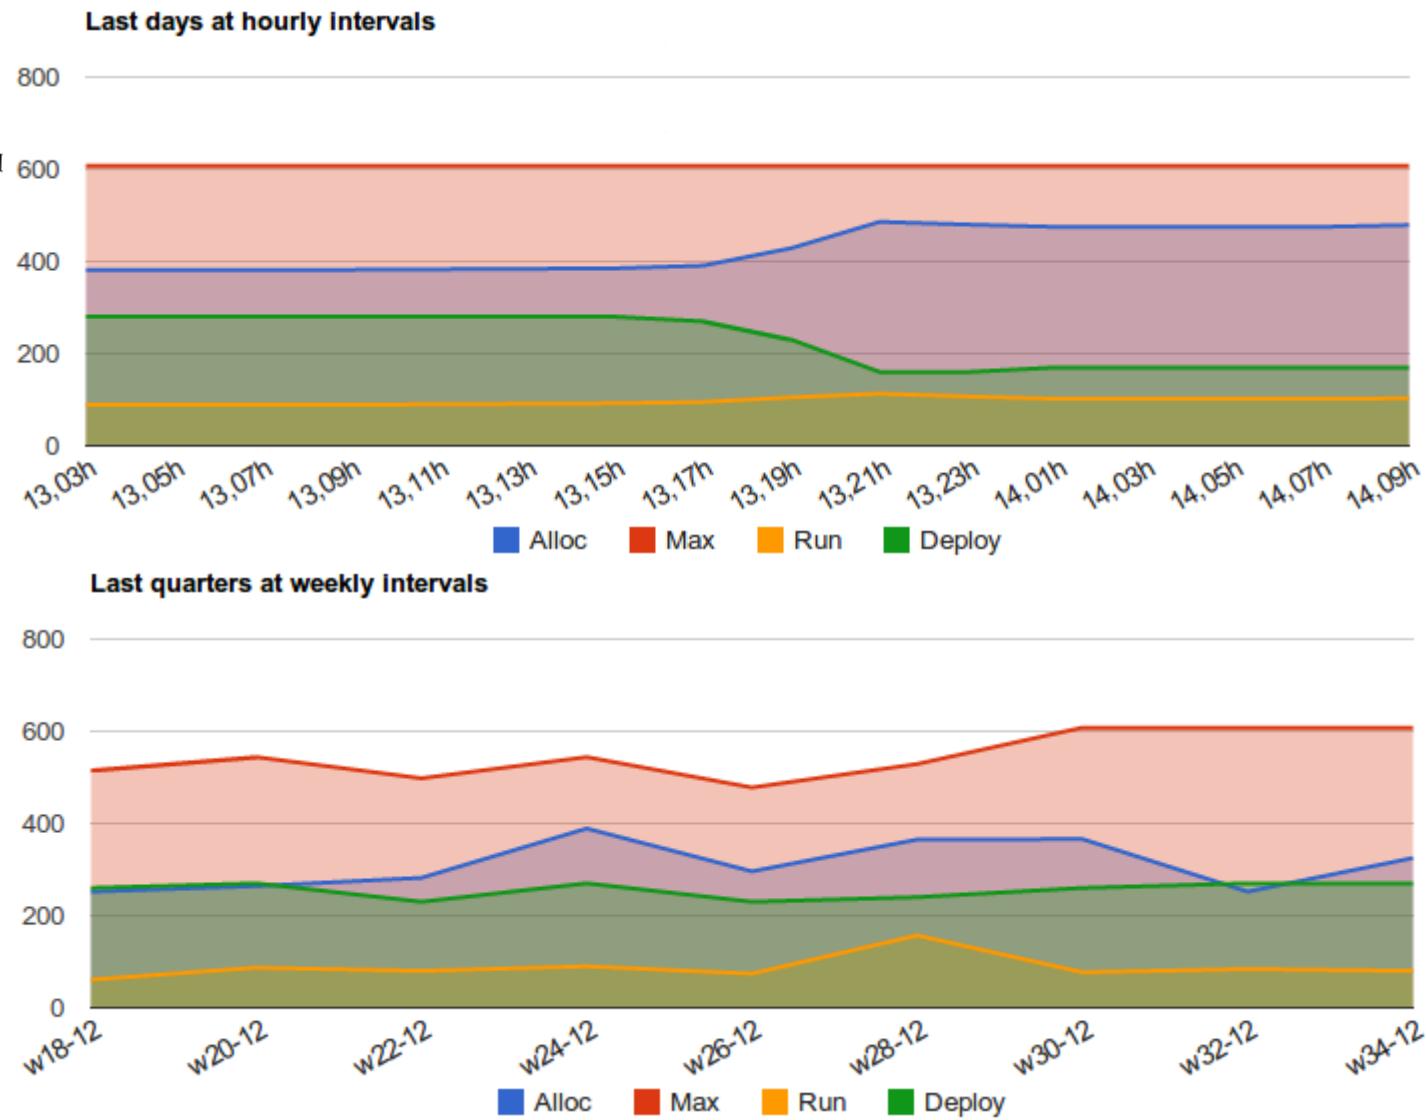


The largest VM that can currently be deployed can use at most 17 cores.



The number of Virtual Machines currently running.

HPC Cloud – Availability Issue

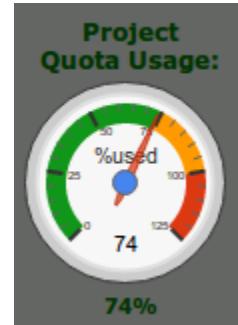


HPC Cloud - Accounting Issue

	Resource usage		
	Hours	Days	Years
2011-12	42,718	1,780	5
2012-01	58,654	2,444	7
2012-02	102,562	4,273	12
2012-03	131,664	5,486	15
2012-04	238,816	9,951	27
2012-05	229,389	9,558	26
2012-06	246,380	10,266	28
2012-07	262,580	10,941	30
2012-08	276,448	11,519	32
2012-09	301,853	12,577	34
Total	1,717,270	71,553	196

Data on 14-09-2012 at 13:30.

Data for September extrapolated.



HPC Cloud – Accounting issues

Field	Type
uuid	binary(16)
datestamp	int(16)
month	varchar(12)
vmid	int(16)
vmname	varchar(128)
templateid	int(16)
templatename	varchar(128)
userid	int(16)
username	varchar(128)
groupid	int(16)
groupname	varchar(128)
start	int(16)
end	int(16)
stateid	int(16)
state	varchar(128)
lcmstateid	int(16)
lcmstate	varchar(128)
usedcpu	int(16)
usedmem	int(16)
usednetrx	int(16)
usednettx	int(16)
usedstorage	int(16)
imageids	varchar(128)
alloccpu	int(16)
allocmem	int(16)
durationrunning	int(16)
durationstopped	int(16)
calculatedduration	int(16)
tokens	int(16)
calculatedtokens	int(16)

30 rows in set (0.00 sec)

HPC Cloud – Remaining Issues

- Shared filesystem presents an I/O bottleneck for certain workloads.
 - ▶ VirDir should be able to solve most of these issues.
- LTS Linux distributions are evolving slowly
 - ▶ Not using LTS would be worse.
- Hardware unreliability
 - ▶ Everything we can do to limit downtime is in place
- Fast Ethernet through SR-IOV not (easily) available in OpenNebula
 - ▶ Waiting for OpenNebula

HPC Cloud - Roadmap

- Improvement of OCCI interface
 - ▶ Interest is rising
 - ▶ EGI Federated Cloud integration
- Better quota enforcement
 - ▶ More feedback to users
- Dynamic modification of running Virtual Machines in new OpenNebula
 - ▶ Disk hotplug
 - ▶ PCI devices
 - ▶ Networkfilters
 - ▶ VLANs

HPC Cloud – Questions?

- CPU (19*32) 608 (Intel Intel Xeon-E7 "Westmere-EX")
- MEM (19*256GB) 4.864GB = 4,75TB
- 10GE, 1-hop, non-blocking interconnect (Arista DCS-7504)
- 400TB shared storage (DDN F10K)



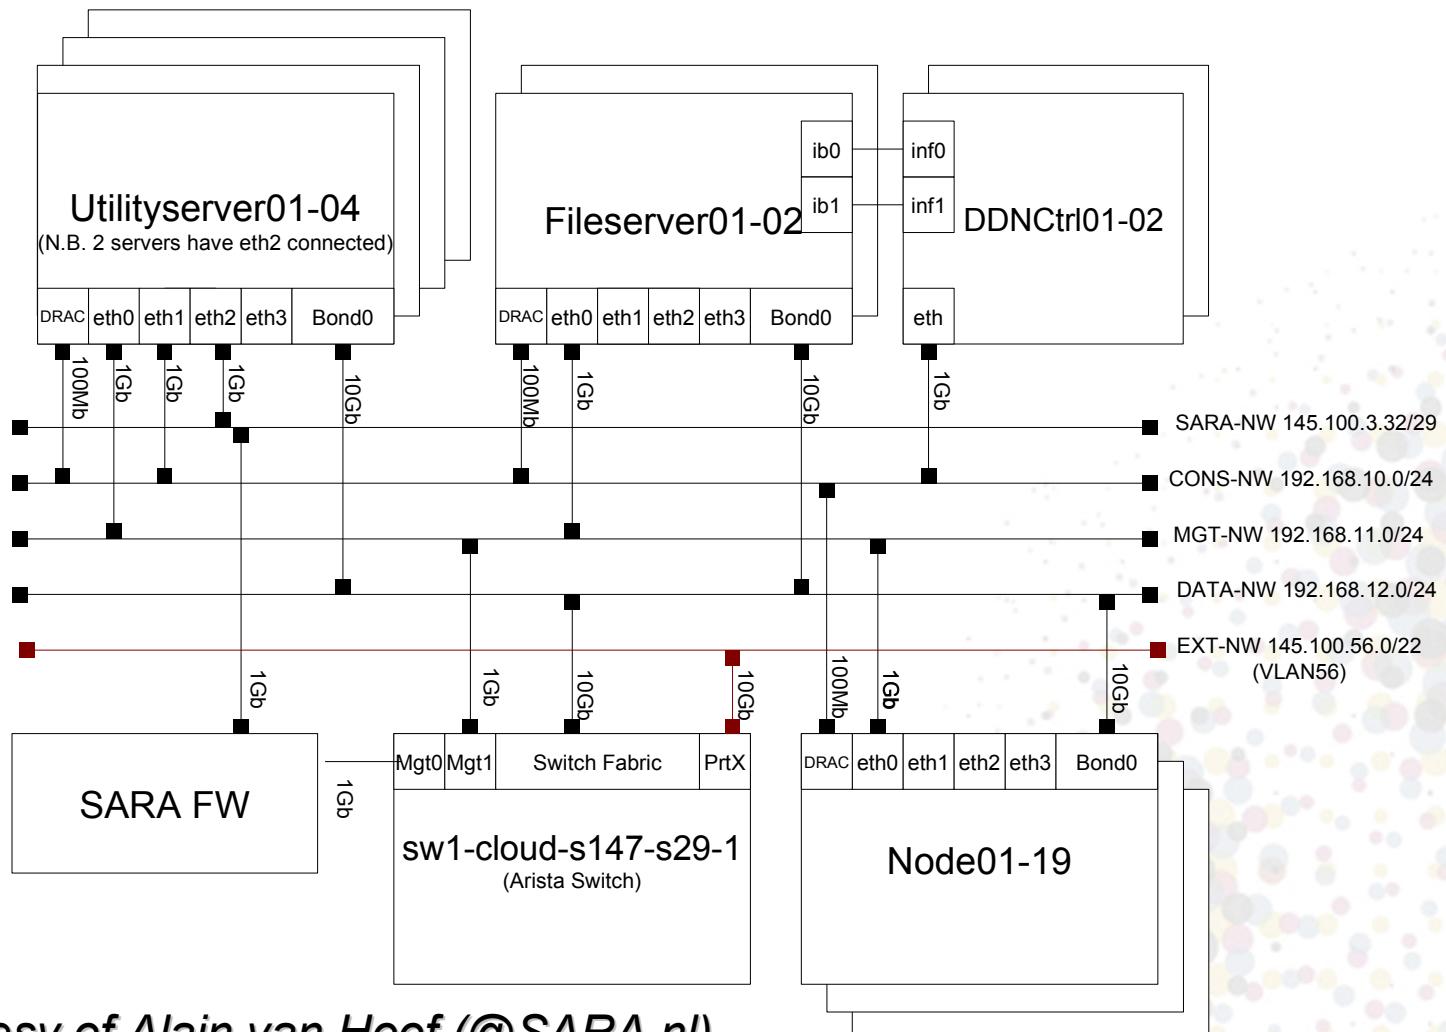


BiG Grid
the dutch e-science grid

Questions?

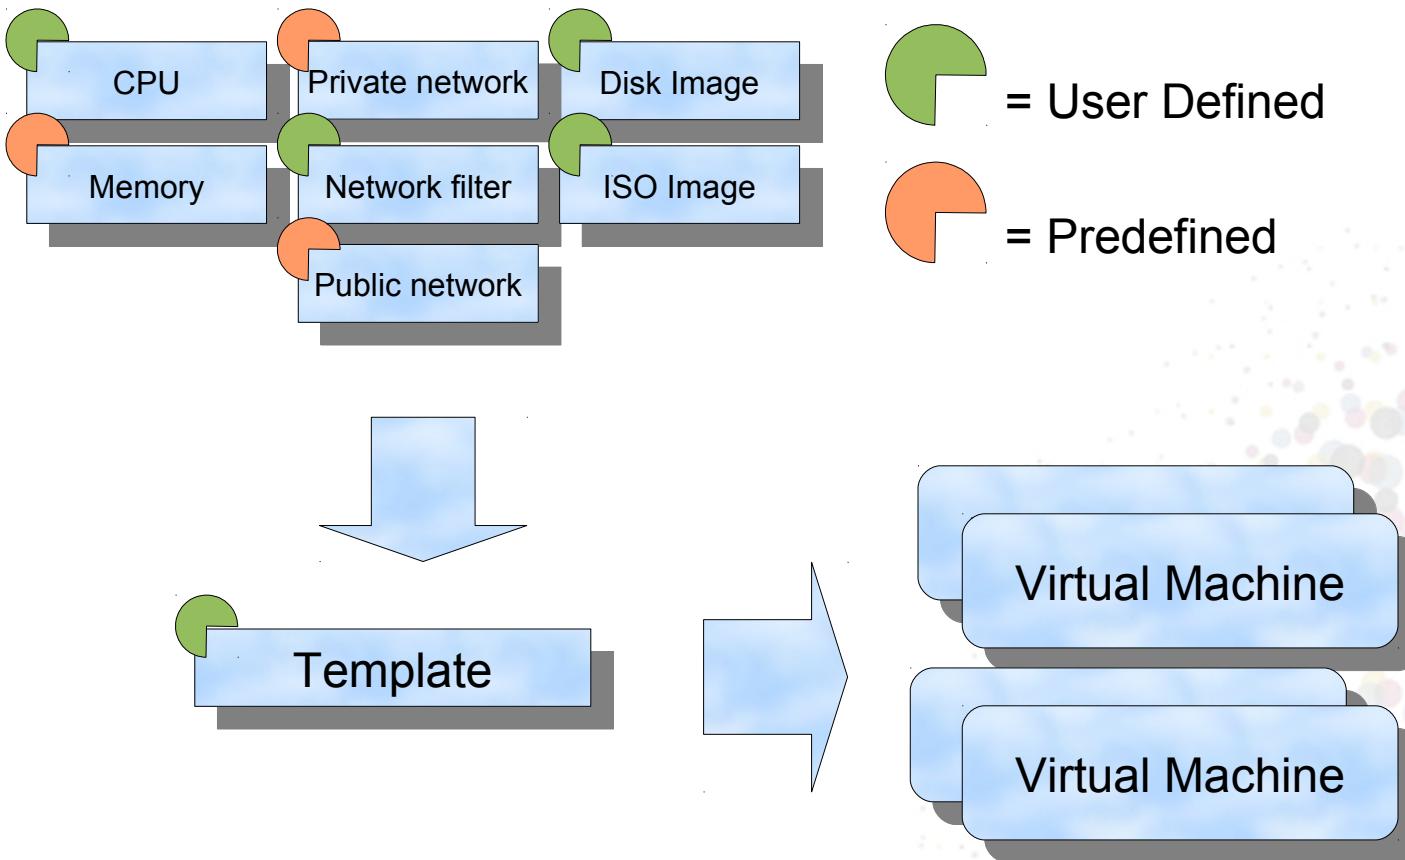
jhon.masschelein@sara.nl

HPC Cloud – Networks



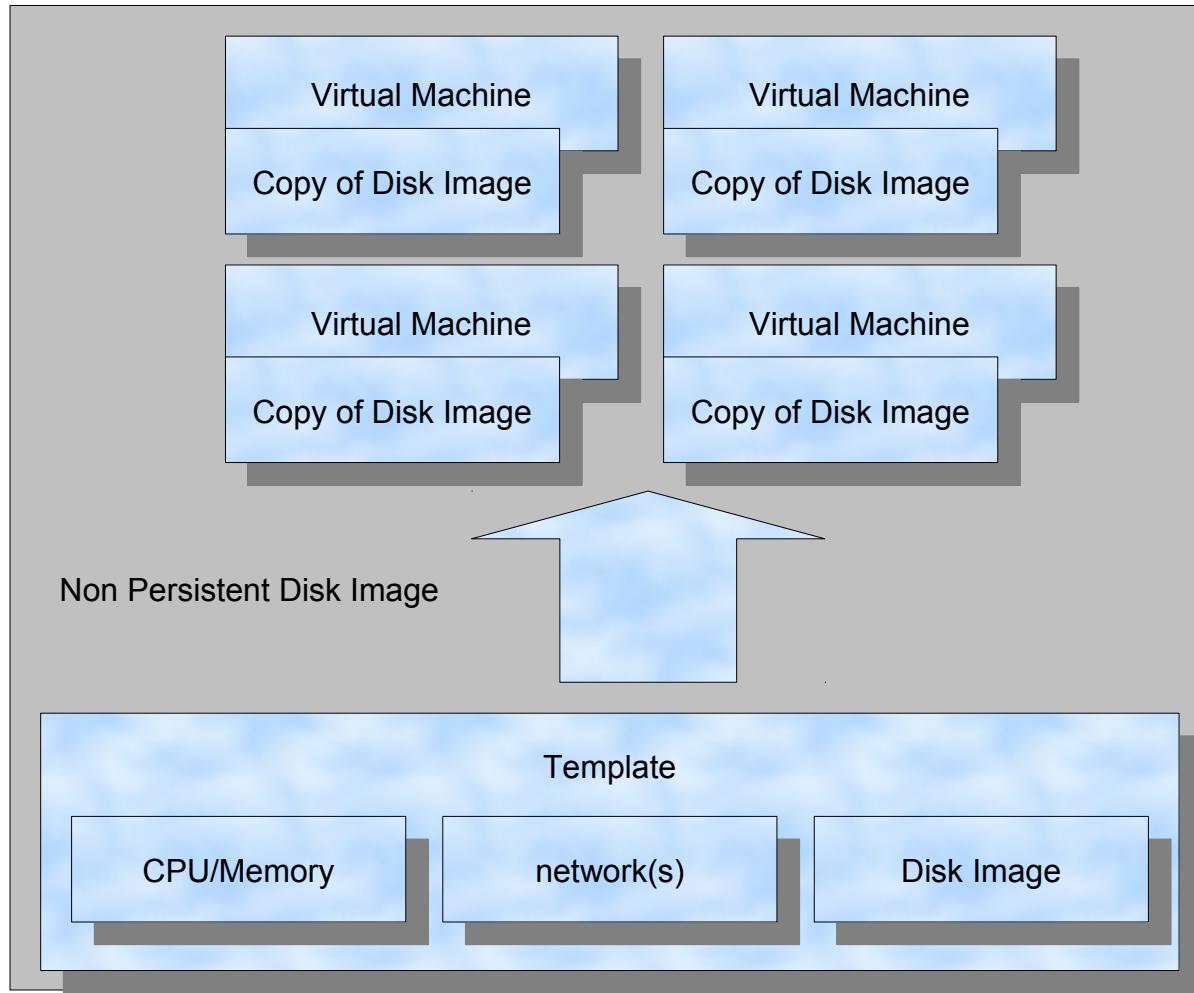
Slide by courtesy of Alain van Hoof (@SARA.nl)

HPC Cloud – Calligo



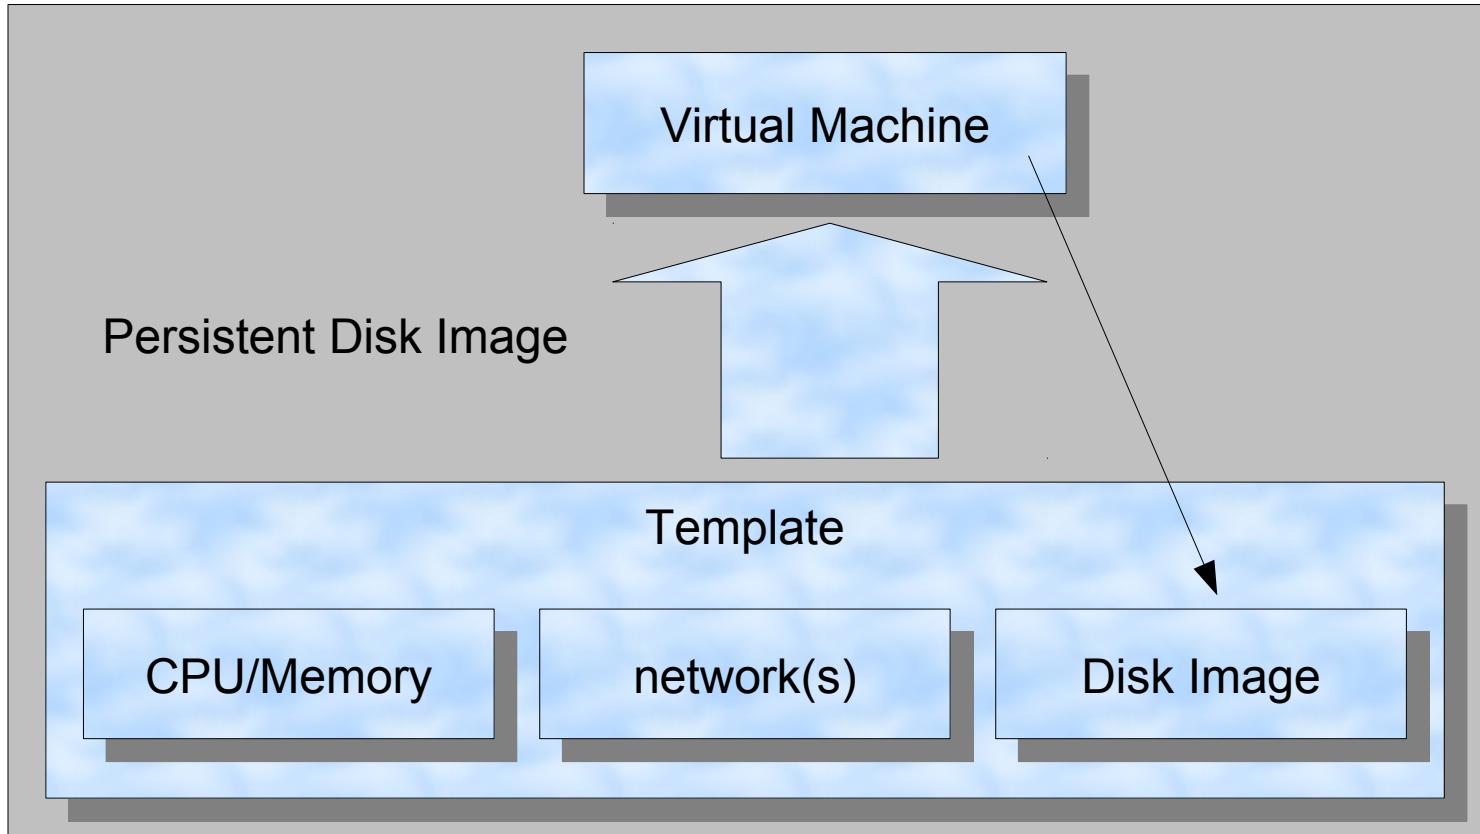
Slide by courtesy of Alain van Hoof (@SARA.nl)

HPC Cloud – Calligo



Slide by courtesy of Alain van Hoof (@SARA.nl)

HPC Cloud – Calligo



Slide by courtesy of Alain van Hoof (@SARA.nl)