BiG Grid HPC Cloud^{Beta}

Floris Sluiter

- SARA Computing and Networking services
 - Amsterdam
 - www.cloud.sara.nl

Big Grid the dutch e-science grid

About BiG Grid

The BiG Grid project is a collaboration between NCF, Nikhef and NBIC, and enables access to grid infrastructures for scientific research in the Netherlands.

SARA is the primary operational partner of BiG Grid





About SARA

- A national High Performance Computing and e-Science Support Center, in Amsterdam
- Tier 1 site LHC Grid Computing
- SARA supports researchers with state-of-the-art integrated services, facilities and infrastructure:
 - High Performance Computing and Networking,
 - National HPC systems: Huygens, Lisa, Grid
 - Data storage
 - Visualization
 - E-Science services
 - Participation in National, European, Global projects as DEISA, PRACE, EGI, EGEE, NL-BiGGrid, and many others

BiG Grid the dutch e-science grid

HPC Cloud Team





"Our" definition of Cloud

Cloud Computing: Self Service Dynamically Scalable Computing Facilities

Cloud computing is not about new technology, it is about new uses of technology







Differences Grid vs HPC Cloud

- We could always run Grid Worker Nodes in our HPC Cloud...
- Return on investment
 - Grid: Cheap resources in bulk. Applications can be difficult to port -> Bulk computing
 - Cloud: more expensive hardware. But easy/no porting of applications -> Tailored Computing
- Time to solution shortens for many users Service Cost shifts from manpower to infrastructure
- Usage cost in HPC stays Pay per Use



Vision: Clone my laptop!

*Ou*r definition of Cloud Computing: Self Service Dynamically Scalable Computing Facilities





Virtual Private HPC Cluster

We plan to offer:

- Fully configurable HPC Cluster (a cluster from scratch)
 - Fast CPU
 - Large Memory (64GB/8 cores)
 - High Bandwidth (40Gbit/s Infiband)
- Users will be root inside their own cluster
- Free choice of OS, etc
- And/Or use existing VMs: Examples, Templates, Clones of Laptop, Downloaded VMs, etc
- Public IP possible (subject to security scan)
- Large and fast storage

Platform:

- Open Nebula
- Custom GUI (Open Source)



Roadmap

- 2009, Q3 Q4: Pilot Phase (finished)
 - Small testbed, 50 cores, 5 usergroups
- 2010, Q2, Q3: Pre-production Phase (almost finished)
 - Medium sized testbed, 128 cores, 100 Tbyte storage
- 2010, Q4,Q++: Production Phase
 - >=1024 cores *planned*, configuration pending



Pre-production Phase From POC to Pr.E...

Physical Architecture

- HPC Cloud needs High I/O capabilities
- Performance tuning: optimize hard- & software
- Scheduling
- Usability
 - Interfaces
 - Templates
 - Documentation & Education
 - Involve **users** in pre-production (!)
- Security
 - Protect user against self, fellow users, the world and vice versa!
 - Enable user to share private data and templates
 - Self Service Interface
 - User specifies "normal network traffic", ACLs & Firewall rules
 - Monitoring, Monitoring, Monitoring!
 - No control over contents of VM
 - monitor its ports, network and communication patterns



A bit of Hard Labour

























Being a pioneer is fun ...

Expert Administrators/developers to develop the infrastructure (and users do not notice the complexity)!!!





Self Service GUI

	😋 Applications: Maces: System 🔂 🖸 🔛 🏭						14 xL d 0 0 2 '	4 2. 4 0 2 'C C Tan Mar Re Dec 1, 19:18 - jerces		
	Openliebule Henspernert Corocle - Mccille Firefox									, W
		5 0	sere of https://ul.cla	autionaria nitrosta ph	p)			* * * 2	Geogle	and a
	No. of Alexand	en Ni Stashabar	B finance B Sustra	HINCY BOUNDA	SANA HIC B DIA	Phip B SAMA Data	transfers E Laco	tetus 🛢 Veuveerdinen aam	Trien Back Assess	
OpenHebule Hateportert Ceracle - Medile Firefox	Paciel 0 12	la Manageme	nt Console						Looped in an	onesdmin - insert
a ffe gat gen Highry Bookman's Sols (194)	um's L botts L	national agentie	nit console							
A A A A A A A A A A A A A A A A A A A	· · · · · · · · · · · · · · · · · · ·						Cr	eate host:		
New Adversers N	I Veneral recent	ing VM Total coul	nee cou Assigner	d cou Total mem	ory Free memory	Status		etrama:		
Openklahula Hanagement Concele	0 host011	800 8	00 800	3262916	3012120	on show #1	able - ak IN	in kan v		
OpenNebula Management Console	1. host02.0	800 8	008 00	3262916	3254028	on show de	able - ak yr	M: wmm kam w		
Create www.	2 host03.0	800 8	008 00	3262916	3206799	on show do	able ~ ak Th	tm_ssh v		
Id liter Nama Status/Du Mamon Host Daus Time	3 host04.6	On and I should	104 104	3363016	3451053	an abara da	able to 1 while the			
194 cloud12 ubuntu-9 shut 0 2097152 host05.22 05:11:30 show delete v nore v 04	4 host05 8	OpenNebul	a managen	nent Consol	e		Logged in a		101	
515 cloud14 ubuntu-9 runn 0 2097152 host06 07 06:19:16 show suspend v none v ok	5 host06.3	vin overview i	m configuration	disk image uplo	ad disk imagen	nanagement ho	ists networks	quotas		
701 oneadmin debian runn 0 2097152 host05 06 01:38:59 show suspend v none v ok		upicau mage.								
740 oneadminidebian runn 0 2097152 host05 05 05:15:21 show isopend v none v ek		Jumpl	cader							
743 cloud17 hadoop-1 runn 0 2097152 host05 05 04:26:33 show sespend v nove v 64	Applications Places System						-		st 1, 19:17 📄 jeros	es.
744 cloud17 hadoop-2 runn 0 2097152 host05 05 04:26:33 show suspend + none + ck	CONTRACTOR STRATEGICS AND ADDRESS	Paste	G Ado	G Flemov	🔋 Retry fa 💟	1 10	0 0	<u>n</u> o	0.000	×
745 cloud17 hadoop-3 runn 0 2097152 host04 05 04:26:33 show suspend + none + ok	the Est View History Bookma	(S) 1	8				Console - Moz	illa Firefox	00	
771 cloud09 ubuntu-9 stop 0 2097152 host0104 04:17:12 show restart v rose v sk	🔶 ih i 🗘 in 🎆 🔳			101:	ara.nl https://ui.claudi	ia.sara.nl/console.php	7id=161		 (q) 	<i>a</i> ,
774 cloud09 ubuntu-9 stop 0 2097152 host04 04 03:45:13 show restart v nere v et	🖪 NUril Algemeen Ni 🖪 Sieshdot			Pace of	rne Octors (lpa	overt Become Bence and 2,45	리카세 40al Balteri			
781 cloud09 ubuntu-9 stop 0 2097152 host04 04 02:56:21 show restart v none v sk	OpenNebula Manageme			/deu/	sda5: clean, 16; fean util-linux-	4/124496 Jiles.	20010/240076 8	locks	204	ut.
782 cloud01 ubunbu-9 runn 0 2097152 host04 04 02:43:47 show suspend v none v ck				ziliouz • Si	mapper/clundundi string_profiminar	e mult elean, " ny koymap	160002609600 41	lis, 226012/2436896 bl	rucks	
1045 cloud09 ubuntu-9 runn 0 2097152 host01 00 03:59:23 show suspend v none v ok	Cloud networks:	deb ar504	File Edit View I	History Bookmarks	s Tools Help					
1055 cloud12 hadoop-m runn 0 2097152 host06 00 03:42:32 show sespend * role * 64	Id User Name Type Br	2.23 34	CpenNebula N	anagement Cons				τ		
1056 doud12 hadoop-s runn 0 2097152 host05 00 03:41:31 show suspend v rate v at	0 onesdmin public Flored br		OpenNebula Ma	anagement Conse	le					
1057 doud12 hadoop-s runn 0 2097152 host05 00 03:41:31 show asspend v rese v ce	8 oneadmin private_01 Fixed br		un eveniew - un cen	ifiguration – disk imago up	load disk image manage	manes hases norway	Logged in an claude iss public forwall exe	epitions quarter		1
1058 (2010)2 hadron - runn 0 2097152 hostor 0 03:41:31 show septem - rese - of	9 oneadmin private_02 Fixed br	(Δ)	Create / Notify templa	ter:						
1060 churd12 hadrons mmn 0 2007152 host04.00 02:41:31 ahos septent v res v re	10 onesdmin private_03 Fixed br				•Tip: You can close a	clemplate by change	g the name during o	shing		
	11 onesdmin private_04 Flood br	3	Name: Permory:		Tobonial_VM_public	= _{MD}				
🗶 🔄 Hibos (3443 total) - Ex. 👘 Re: Opernebula Masa. 🛛 😝 OperNebula Manage. 🖉 jemenýtus	12 oneadmin private_05 Fixed br	Goto uploa STRI	Cou:		1 7					
			OS section:							
Doveloped at SARA			' Doot:		hd y					
			Disk section:							
-			Network section:		4 7					
Onon Source available at			Number of networks:		2 1					
open source, available at			Graphics section:		2 b					
- · · ·			VNC Password:		4				0	
www.opennepula.org			next						=	

BiG Grid

the dutch e-science grid

User participation 12 involved in Beta testing

nr.	Title	Core Hours	Storage	Objective	Group/instiute	
1	Cloud computing for sequence assembly	14 samples * 2 vms * 2-4 cores * 2 days = 5000	10-100GB / VM	Run a set of prepared vm's for different and specific sequence assembly tasks	Bacterial Genomics, CMBI Nijmegen	
2	Cloud computing for a multi- method perspective study of construction of (cyber)space and place	2000 (+)	75-100GB	Analyse 20 million Flickr Geocoded data points	Uva, GPIO institute	
3	Urban Flood Simulation	1500	1 GB	asses cloud technology potential and efficiency on ported Urban Flood simulation modules	UvA, Computational Science	
4	A user friendly cloud-based inverse modelling environment	testing	1GB / VM	running in the cloud supporting modelling, testing and large scale running of model.	Computational Geo-ecology, UvA	
5	Real life HPC cloud computing experiences for MicroArray analyses	8000	150GB	Test, development and acquire real life experiences using vm's for microarray analysis	Microarray Department, Integrative BioInformatics Unit, UvA	
6	Customized pipelines for the processing of MRI brain data	?	up to 1TB of data -> transferred out quickly.	Configure a customized virtual infrastructure for MRI image processing pipelines	Biomedical Imaging Group, Rotterdam, Erasmus MC	
7	Cloud computing for historical map collections: access and georeferencing	?	7VM's of 500 GB = 3.5 TB	Set up distributed, decentralized autonomous georeferencing data delivery system.	Department of Geography, UvA	
8	Parallellization of MT3DMS for modeling contaminant transport at large scale	64 cores, schaling experiments / * 80 hours = 5000 hours	1 TB	Goal, investigate massive parallell scaling for code speed-up	Deltares	
9	An imputation pipeline on Grid Gain		20TB	pipelines and, in particular, heavy imputation pipelines on a new HPC cloud	Groningen Bioinformatics Center, university of groningen	
10	Regional Atmospheric Soaring Prediction	320	20GB	Demonstrate how cloud computing eliminates porting problems.	Computational Geo-ecology, UvA	
11	Extraction of Social Signals from video	160	630GB	Video Feature extraction	Pattern Recognition Laboratory, TU Delft	
12	sequencing data from mouse tumors	?	150-300GB	Run analysis pipeline to create mouse model for genome analysis	Chris Klijn, NKI	



BiG Grid the dutch e-science grid

Usage statistics in beta phase

- Users liked it:
 - ~90.000 core-hours used in 10 weeks (~175.000 available)
 - 50% occupation *during beta testing*
 - Some pioneers paved the way for the rest ("Google" launch approach)
 - Evaluation meeting with users, outcome was very positive



User Experience

(slides from Han Rauwerda, transcriptomics UVA)

Microarray analysis: Calculation of F-values in a 36 * 135 k transcriptomics study using of 5000 permutations on 16 cores.

worked out of the box (including the standard cluster logic)

no indication of large overhead

Ageing study - conditional correlation

- dr. Martijs Jonker (MAD/IBU), prof. van Steeg (RIVM), prof. dr. v.d. Horst en prof.dr. Hoeymakers (EMC)
- 6 timepoints, 4 tissues, 3 replicates and 35 k measurements + pathological data
- Question: find per-gene correlation with pathological data (staining)
- Spearman Correlation conditional on chronological age (not normal)
- p-values through 10k permutations (4000 core hours / tissue)

Co-expression network analysis

- 6k * 6k correlation matrix (conditional on chronological age)
- calculation of this matrix parallellized. (5.000 core hours / tissue)

Development during testing period (real life!)

Conclusions

- Many ideas were tried (clusters with 32 64 cores)
- · Cloud cluster: like a real cluster
- Virtually no hick-ups of the system, no waiting times
- User: it is a very convenient system



Our Cloud What was, what is and what will be...



SIMPLY EXPLAINED - PART 17: CLOUD COMPLITING

- Pilot
- **Pre-production** (Now in Beta)
- Production system will take 3-4 months after goahead. And in the mean time we will continue to support and improve the beta system



What else is Cooking?



Extra features:

- AAA
 - Sharing resources
 - Accounting also on I/O & infra
 - Ldap / x509
- Finegrained firewall
- Scheduling also on memory and i/o bandwidth
- Selve Service Storage
 - CDMIFUSE
 - (prototype = working)
- Self service networking
 - Please supply use cases!
- More experiments!







Acknowledgements Our Sponsor: NL-BiGGrid Our Brave & Daring Beta Users And the HPC Cloud team: Tom Visser, Neil Mooney, Jeroen Nijhof, Jhon Masschelein, Dennis Blommesteijn, et. al.

http://www.cloud.sara.nl photo: http://cloudappreciationsociety.org/

