

Local Desktop Grid management experiences

Hannu Visti
University of Westminster

24/9/2014

UoW desktop grid



Cavendish Campus	570-580 nodes
Marylebone Campus	550-560 nodes
Harrow Campus	250-260 nodes
Regent Street Campus	390-400 nodes
Total	~1800 nodes

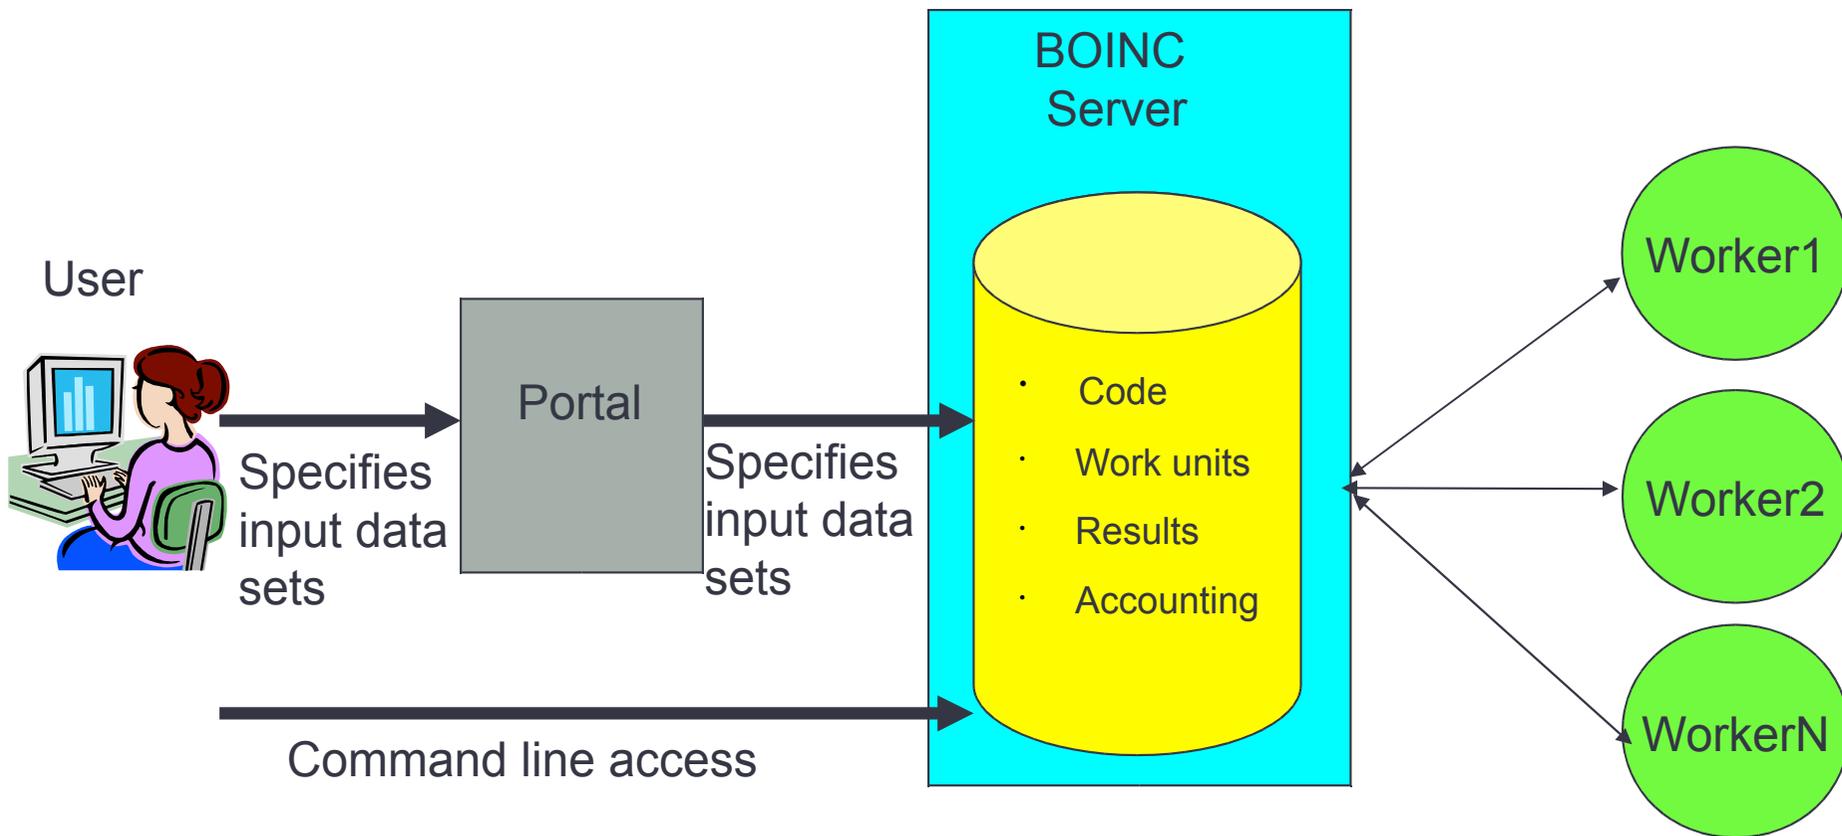
UoW grid

- “Closed” grid. Only university computers can become grid clients.
- Job submission either from command line (requires a local account) or through WS-PGRADE portal (works from outside university firewall)
- Mainly Windows 7 computers (32 bit and 64 bit)

UoW grid and WS-PGRADE portal

- Support local university researchers, teachers to use distributed computing
- Without writing a line of code
- Current applications:
 - Rendering (based on MentalRay 3D graphics software)
 - Docking (based on Autodock and Autodock Vina)
- Portal gives high level view for users
- Separate “view layout” for separate users/groups
- Hide the complex system from users

Job submission



System infrastructure

- Dg-portal powered by WS-PGRADE

dg-portal

- Regular BOINC server additional service

Boinc-server

- University Lab PCs

Boinc-clients

WESTMINSTER™ Desktop Grid Portal

Welcome Workflow Storage Settings Information Statistics Publications Help

Docking portlet Mental Ray

University of Westminster Desktop Grid Portal > End User

End User

Back Save on Server

Workflow name: PublicAutodock-NoAutogrid-Log-EndUser_2013-04-22-143759
Note: 2013-4-22

D, Best result number 5

A, Inputs file (*.zip) contain all map, molecules, etc. inputs.zip

B, Docking parameter file (*.dpf) docking.dpf

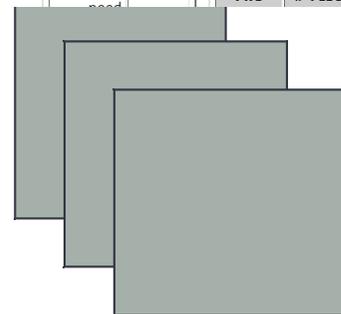
C, Number of docking 2000

Message:

7392 results 'Over' results 'Success' results 'Client error' results

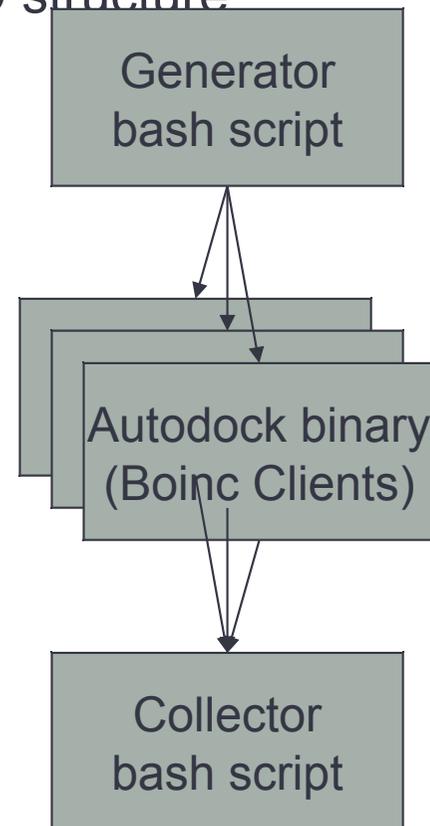
Server state	# results	Outcome	# results	Validate state	# results	Client state	# results
Inactive	0	Init	0	Initial	5	Downloading	0
Unsent	3	Success	6994	Valid	6989	Downloaded	0
Unsent (in work seq)	0	Couldn't send	0	Invalid	0	Compute error	221
In Progress	174	Client error	221	Skipped	0	Uploading	0
Over	7215	No reply	0	Inconclusive	0	Uploaded	0
		Didn't need	0	Too late	0	Aborted	0

File	# results



Portlets

- Autodock
 - Designed to predict how small molecules, such as substrates or drug candidates, bind to a receptor of known 3D structure
- Docking portlet functions:
 - Upload inputs, set parameters
 - Validate
 - Submit task
 - (Wait)
 - See the results
- Autodock Workflow



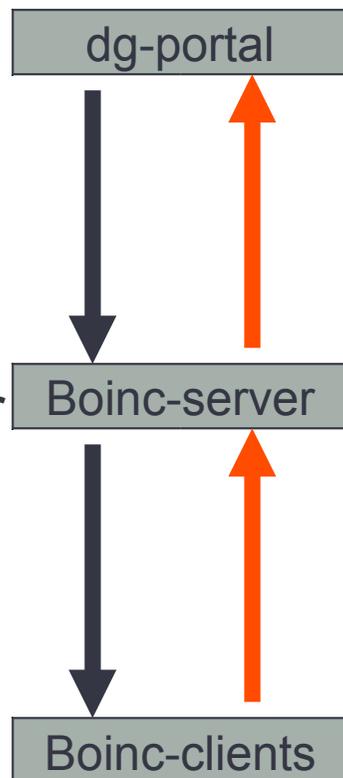
Dataflow

- Dg-portal

- BOINC server

- Lab PCs

- (workers)



UNIVERSITY OF WESTMINSTER University of Westminster Desktop Grid Portal SCI-BUS

Welcome Workflow Storage Settings Information Statistics Publications Help End User

University of Westminster Desktop Grid Portal End User

End User

Back Save on Server

Workflow name: PublicAutodock-NoAutogrid-Log-EndUser_2013-04-22_143759

Note: 2013-4-22

D. Best result number 5

A. Inputs file (*.zip) contain all map, molecules, etc. inputs.zip Browse...

B. Docking parameter file (*.dpf) docking.dpf Browse...

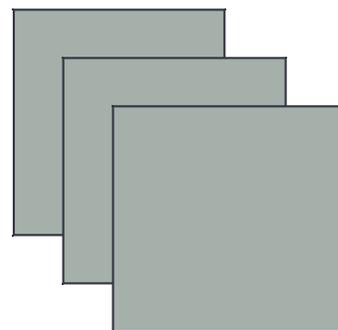
C. Number of docking 2000

Message:

7392 results 'Over' results 'Success' results 'Client error' results

Server state	# results	Outcome	# results	Validate state	# results	Client state	# results
Inactive	0	Init	0	Initial	5	Downloading	0
Unsent	3	Success	6994	Valid	6989	Downloaded	0
Unsent (in work seq)	0	Couldn't send	0	Invalid	0	Compute error	221
In Progress	174	Client error	221	Skipped	0	Uploading	0
Over	7215	No reply	0	Inconclusive	0	Uploaded	0
		Didn't	0	Too late	0	Aborted	0

File	# results



Applications

- Applications:
 - Autodock_vina
 - Public_autodock_4_2_3
 - Public_autodock_vina_1_1_2
 - Public_autodock_4_2_3_multiple_dockings
 - Autodock_4
 - Mentalray
 - Visivo_space_mission
 - Video_splitter
 - Video_converter
 - Video_merger

Closed grid management challenge

System administration tools

- “Project management” web page
- Shell and python scripts
 - Check host activity
 - Check application status
- Command line tools
 - Mostly related to disk space management

Security

- A malicious application could be deployed throughout university via the grid infrastructure
 - In practice this is mitigated in a closed grid by limiting administration access to grid server – users cannot deploy applications or upgrade existing ones.
- Nevertheless, UoW attempted running grid tasks inside Virtualbox
 - This attempt was a failure as Novell Zen was unable to assign required privileges.
 - Would work without Novell

Observed problems from Jan 2014

- Two major declines in host activity
- Mentalray 2014 stopped working, probably was never functional
- Obsolete applications
- Disk space problems
 - Out of storage space
 - Out of i-nodes
- Non-working application wrapper scripts

Declining host activity

- Measuring host activity takes time, as clients can sleep for 24 hours without contacting the server
- January 2014 only 930 hosts were active
 - A firewall change had isolated ECS building from the grid. This was fixed and activity returned to expected levels.
- Another decline in Spring/Summer:
 - 1/4/2014: 1794 active
 - 1/5/2014: 1709 active
 - 1/6/2014: 1571 active
 - 21/6/2014: 1262 active
 - 19/9/2014: 1133 active

Mentalray 2014

- For an unknown reason Maya 2014 does not work if Windows permissions of Boinc users are set from Novell Zen. With local permissions the application works.
- Tasks start but never finish
- Conclusion to wait for Maya 2015, which appears to be working based on initial analysis
- Problem not directly related to grid/Boinc but still interesting as there appears to be compatibility issues with centralised Windows management systems.

Obsolete applications

- Several obsolete and/or non-functional applications (mostly versions of Autodock) existed in the grid
- These do not cause problems as such, but removing an application from the grid is not trivial and to our best knowledge not covered by installation guides or documentation.

Disk space problems

- Grid server partition is susceptible of running out of disk space and inodes.
- Master/3g-bridge/input and master/3g-bridge/output are the main culprits.
- No automatic clean up exists. Can be automatized with cron jobs and scripting
- A large number of minuscule files in these directories ensure a standard inode allocation is not sufficient. It is not possible to increase number of inodes in ext4 file system without reformatting
 - UoW observation: Partition is 75% full when 100% of inodes occupied

Non-working application scripts

- Gitbox return values appear to be large random numbers, possibly due to an uninitialized variable or signed/unsigned conflict.
- All wrapper scripts that test gitbox return values, generally fail due to this.
- Quick fix to comment out all tests (done at UoW). Bug reported to SZTAKI and acknowledged by them in February but no fix issued.
 - Scripts cannot do any sanity checks or result checks
- Gitbox also only exists as a 32 bit version. No 64 bit version promised by SZTAKI.

Conclusion

- Boinc client problems – if any – are related to system image upgrades, centralised management systems and network connectivity. Clients and the server appear to be robust.
 - Clients sleeping for 24 hours make testing more complex
- Server infrastructure needs constant monitoring due to disk space and inode issues
- Obsolete applications may not be a problem as such but removing them is not a documented process.
- Gitbox appears not to be supported anymore

Differences between closed and public grids

- Closed grid
- Security is not a major concern
- Closed environment causes problems to grid computing due to way the infrastructure is managed
- Problems tend to be systematic: one computer seldom fails. They all work – or none do!

- Public grid
- Application security is a major concern
- Technical issues mainly related to server side administration
- Benefits from distributed administration as users are in charge of their computers, networks and firewalls

• Being end system

• Problems from

Recommendations

- Gitbox should either be supported by someone or removed from production. This far we have managed with workarounds and the 32 bit version only but this might not be enough in the future.
 - Another attempt with desktop virtualisation might solve the issue. This could also help solve problems caused by system image upgrades.
- A new grid installation should take the inode issue into account and format their grid data partition with 50% more inodes than default.
- Application removal process should be documented.
- A proper set of grid monitoring tools should be developed. UoW scripts are ad hoc solutions.

Questions?