



High Throughput, Low Impedance e-Science on Microsoft Azure

David Abramson & Blair Bethwaite

Monash e-Science and Grid Engineering Lab (MeSsAGE Lab)
Faculty of Information Technology

Monash e-Research Centre

Outline




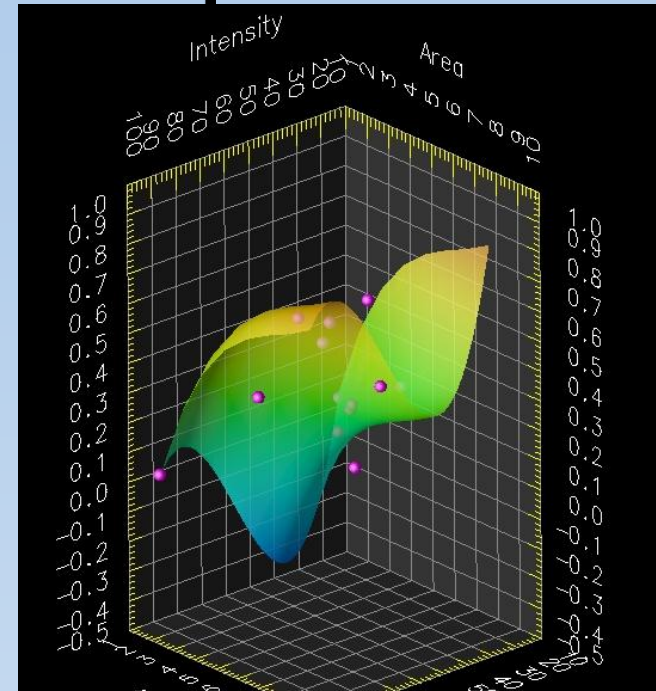
- A new kind of science
 - When you get the tools right magic happens
 - Robust Science with Nimrod
 - Nimrod/K workflows
- From Clusters to Grids to Clouds
- New Azure application drivers
- Conclusions




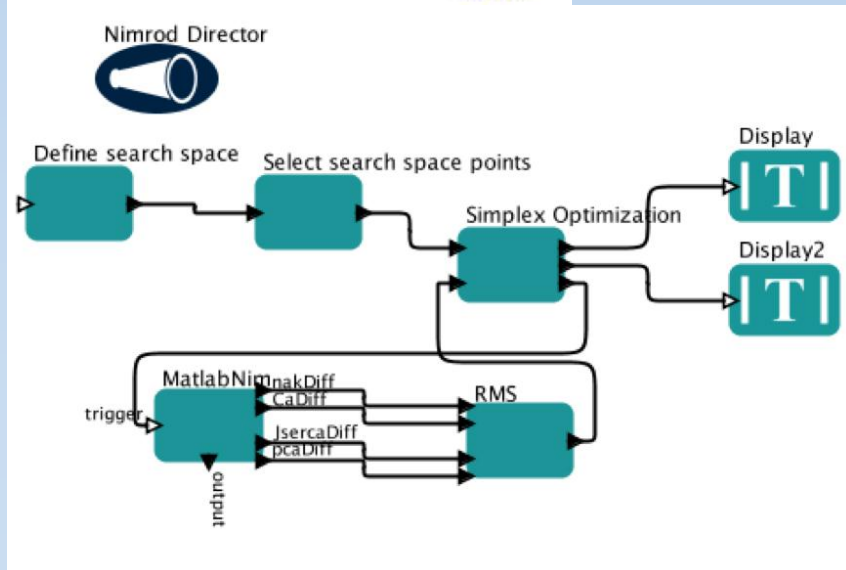


When you get the tools right magic happens

A new kind of science


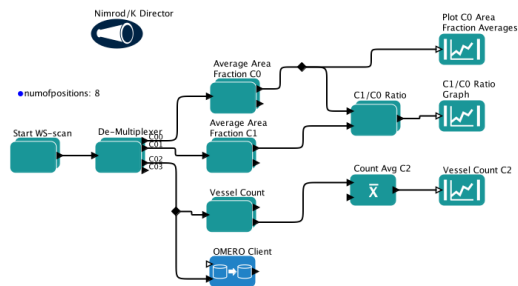

Science that was impossible



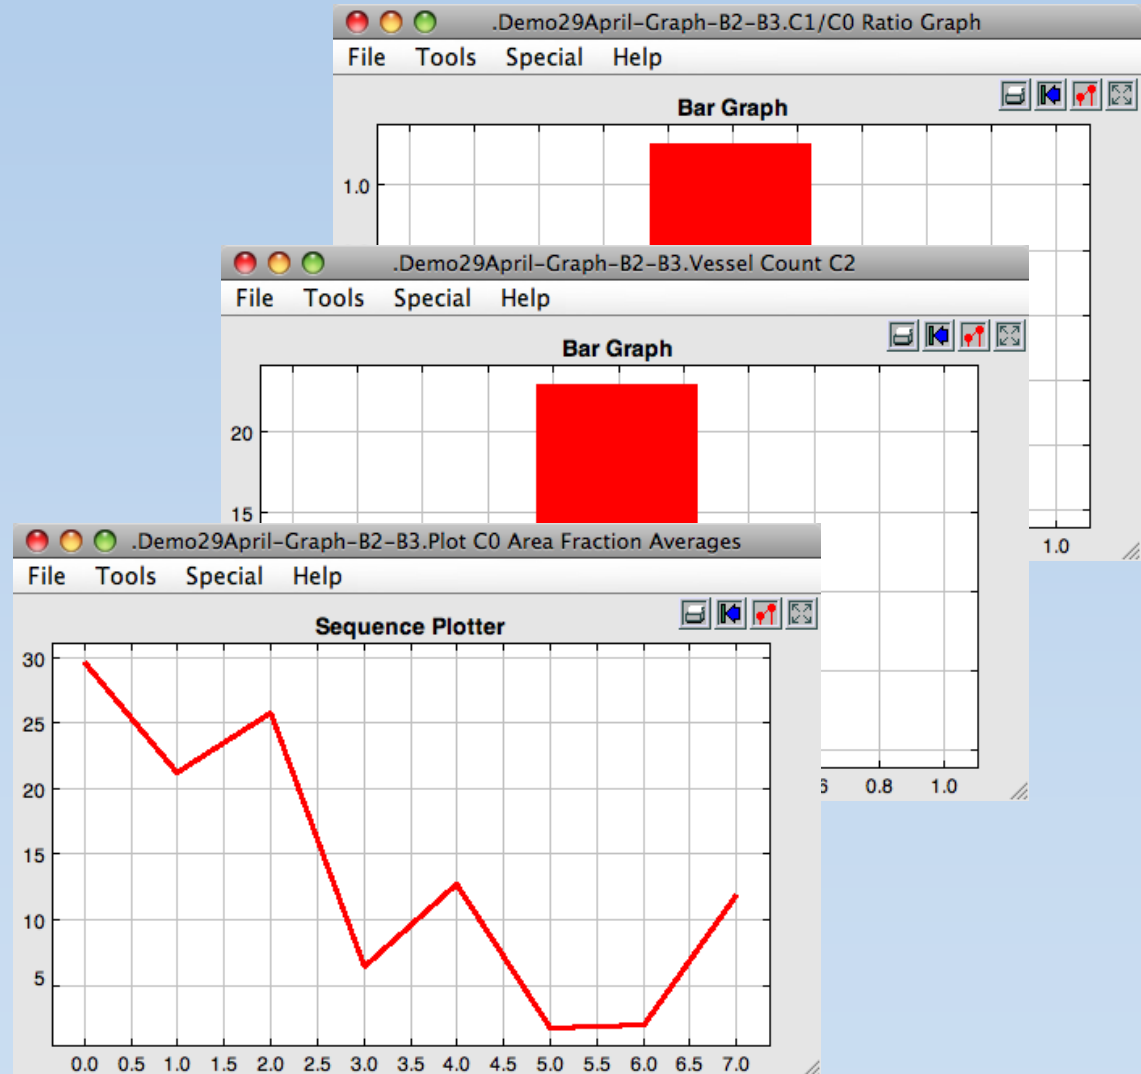
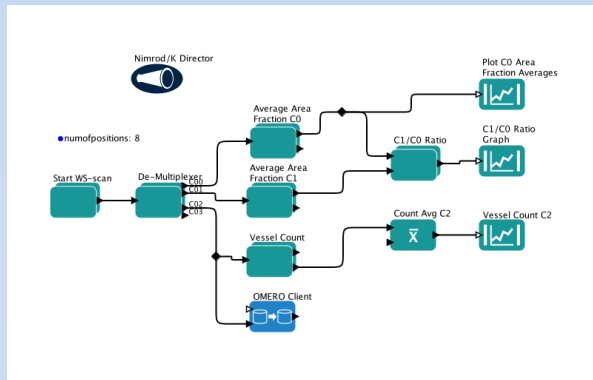

New ways of thinking



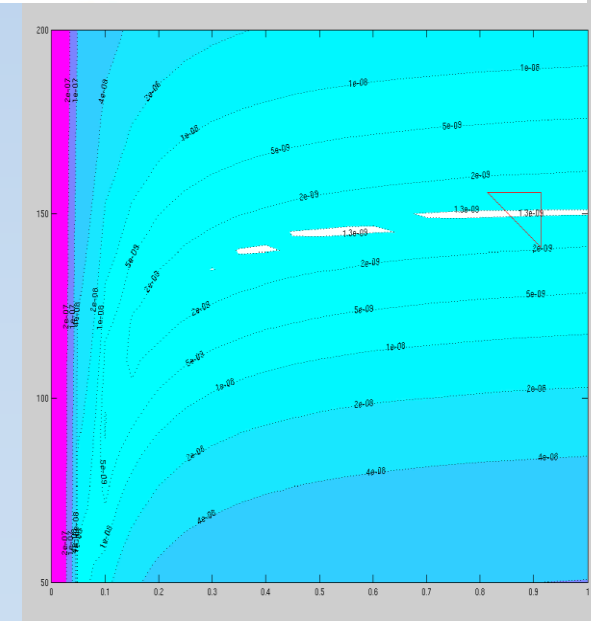

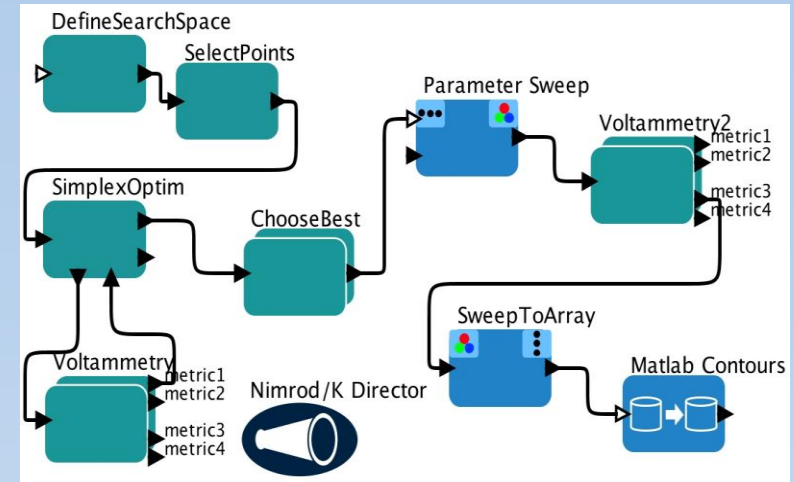

New ways of thinking



New ways of thinking



That work on other problems





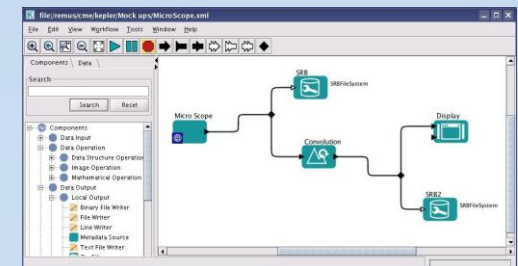



Robust science with Nimrod

A new kind of science

Robust science with Nimrod



- A full parameter sweep is the cross product of all the parameters (Nimrod/G)
- An optimization run minimizes some output metric and returns parameter combinations that do this (Nimrod/O)
- Design of experiments limits number of combinations (Nimrod/E)
- Workflows (Nimrod/K)



The Nimrod family



Plan File

Nimrod Portal

Nimrod/O

Nimrod/E

Nimrod/G

Actuators

Grid Middleware

Nimrod Portal

Parameter Section

<input type="checkbox"/>	<input type="checkbox"/>	temp	label	<input type="text"/>	float	range	from	<input type="text" value="200"/>	to	<input type="text" value="300"/>	Points	<input type="text" value="2"/>
<input type="checkbox"/>	<input type="checkbox"/>	pressure	label	<input type="text"/>	float	range	from	<input type="text" value="5000"/>	to	<input type="text" value="6000"/>	Points	<input type="text" value="4"/>
<input type="checkbox"/>	<input type="checkbox"/>	concent	label	<input type="text"/>	float	range	from	<input type="text" value="0.002"/>	to	<input type="text" value="0.005"/>	Points	<input type="text" value="2"/>
<input type="checkbox"/>	<input type="checkbox"/>	material	label	<input type="text"/>	text	select	<input type="text"/>	<input type="button" value="Add"/>	<input type="button" value="Fe"/>	<input type="button" value="Al"/>	<input type="button" value="Remove"/>	

Add a new parameter

Add a comment

Tasks Section


- ☐ task rootstart
- ☐ task nodestart
- ☐ task main
- ☐ task nodefinish
- ☐ task rootfinish



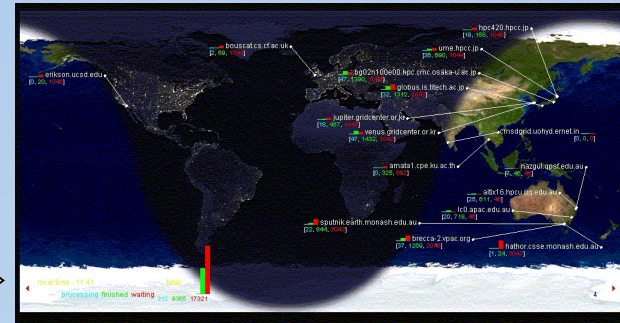
Add a new task

Save

Cancel and Reload

Text mode





Jobs Scheduled Executed Dynamically


```
root@east-
Terminal Tabs Help
8 0/ 0 18/ 18
0 1/ 8 1/ 8
3 0/ 0 53/ 53
[screen 8: test] blair@nimrod1:~
Terminal Tabs Help
checked command 30 ready:False
checked command 70 ready:False
:1
e internal queue:
ing:0
81707500 pausing for 17.2020218372
main thread
:1
main thread
e internal queue:
ing:0
parallel
81707500 pausing for 18.2319529057
main thread
checked command 72 ready:True
main thread
unchecked command 72 ident:1236
checked command 70 ready:True
main thread
unchecked command 70 ident:1241
checked command 30 ready:True
tj$ ssh-keyscan -t rsa ec2-174-129-155
204.compute-1.amazonaws.com SSH-2.0-Ope
4.compute-1.amazonaws.com ssh-rsa AAAAB
xQD75xcN6nRoedkR0TnJSLMurlmrCnwvoTA17kq
fCm32YJmFTX0AamtaqI1IAw//SnaEz0cFIIGfQ
emljXRqo40VEIhdWf9i9YAtlG7ikgRe8LJY2MJq
9oKQ==
tj$
tj$
tj$ tailf ~/.nimrod/ec2.log
unchecked command 31 ident:1251
:1
e internal queue:
ing:0
:1
e internal queue:
ing:0
81707500 pausing for 18.2695720196
main thread
main thread
main thread
:1
e internal queue:
ing:0
81707500 pausing for 18.2608129978
main thread
```

Nimrod portal - Swiftfox

File Edit View History Bookmarks Tools Help

https://messagelab.monash.edu.au/NimrodARCS/cgi-bin/np-ml.cgi?page=experiments?window=

Gmail Latest BBC Headlin... Zazz Cricket Aus Chaser foobar2000 Inbox - blair.bethwai... Most Visited Latest Headlines

Disable Cookies CSS Forms Images Information Miscellaneous Outline Resize Tools View Source Options


Inbox - ... ITS Mes... Nimr... pyar.py... Build a ... Amazo... Cloud E... A - Z El... a.http...ges a.http...0& boto - P... http:...txt

Boto Amazo... The Pyt... 11.1. os... Amazo... Amazo... 36.16. c... 18.5. po... 18.1. su... /trunk/jni... Eucalyp... Eucalyp...

17.2. th... Amazo... MeSsA... 7. Built-i... Testing ... ec2 bot... Amazo... Tip: Get...

Nimrod Portal

Welcome Blair Bethwaite (you are an administrator)



Status of test


What's this?

The experiment has executed for 6mins and 0secs.

The experiment has started. 29 jobs have been completed, there are 23 jobs waiting, 38 jobs executing and 0 jobs failed.

[Check for grid errors](#)

[Pause Experiment](#)



Plan file

What's this?

```
parameter x integer range from 7 to -1 step 1;
parameter y float range from 0.125 to 1.25 step 0.125;

task main
  copy work.py node:.
  node:execute /bin/chmod +x work.py
  node:execute /bin/hostname >> output
  node:execute /bin/date >> output
  node:execute /bin/pwd >> output
  node:execute /bin/echo "Working..." >> output
  node:execute ./work.py ${x} ${y} >> output
  copy node:output output.${jobname}.x${x}.y${y}
endtask
```

You cannot change the plan after execution has started.

Find:


Previous Next Highlight all Match case

Done

38 rows

blair-dev=#

Experiments

 test

[New experiment](#)

[Resource management](#)


[Administrative tools](#)

[Online Manual](#)

[Grid World](#)

[About](#)

[Logout](#)



Applications Places System

New Settings Show Discard

Hardy
Saved

Intrepid
Powered Off

NG Dev
Powered Off

XPL (Pre SFU)
Running

Details

Gene

1 21 41 61 81

21 41 61 81

Hard

CD/D

Flopp

Audic

WaspFactory.txt

gt4-test.sh

pcm9_linux

linux-phc-0.3.0-pre1

EVOPlayer.jnlp

e169g-switch_0.3_all.deb

temp

pslinuxv180

Nimrod Viewer for experiment test - Swiftfox

https://messagelab.monash.edu.au/NimrodARCS/cgi-bin/np-ml.cgi?page=nvapplet?param=test?window=main

1

21

41

61

81

none

nimrodtest.q@https://east-globus.enterprisegrid.edu.au:8443/SGE@

cloud1//home1/blair/ec2/david-id/home1/blair/ec2/david-secret2:5:m1.small:ami-0d729464

Jobs	W	E	D
0 (0%)	0	-	0
67 (74%)	7	32	28
23 (25%)	5	13	5
Total: 90	12	45	33

Job: 1 Status: Executing ReQue... Mark as Do... Mark as Fail...

Time: 0:00:00 Parameters: x "7" y "0.125"

Applet nimrodviewer.Applet started messagelab.monash.edu.au Open Notebook zotero


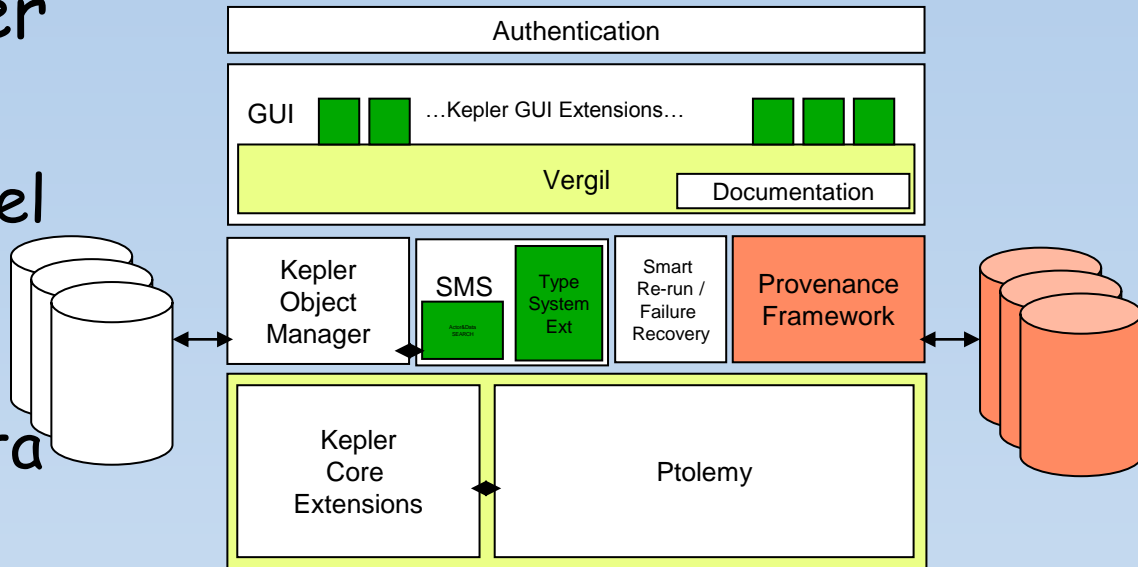
docs.google.com Open Notebook zotero

Nimrod/K Workflows

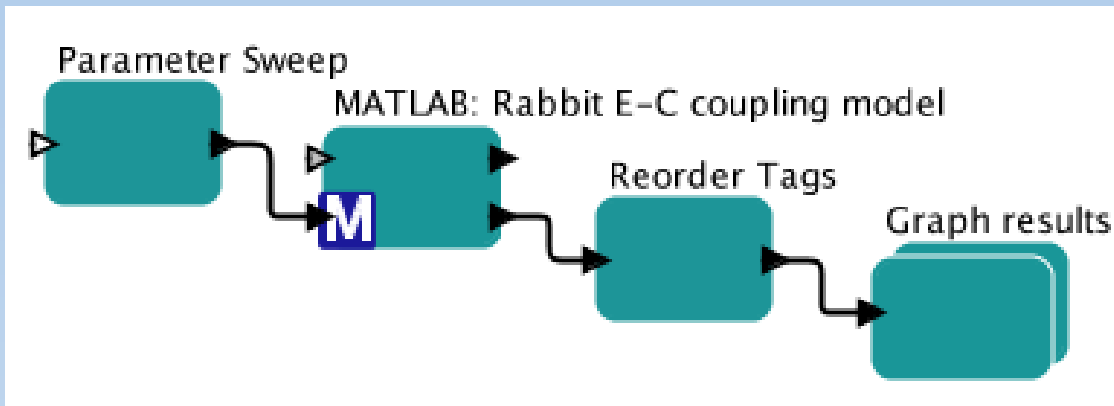


- Nimrod/K integrates Kepler with


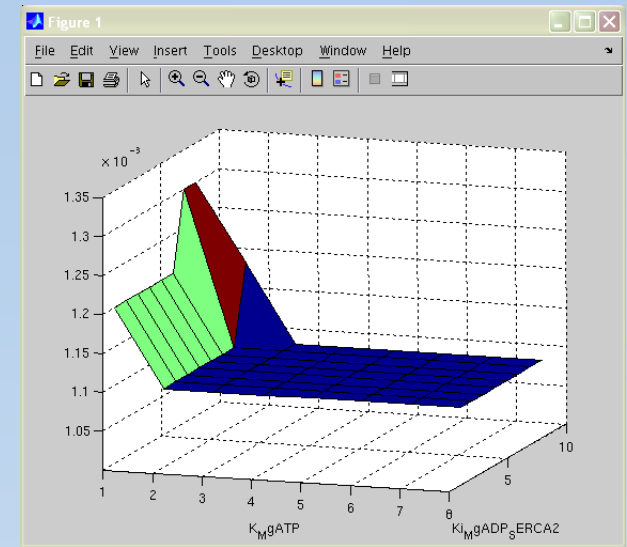
- Massively parallel execution mechanism
- Transparent data transport
- Robust design from Nimrod/G/O/E



Complete Parameter Sweep



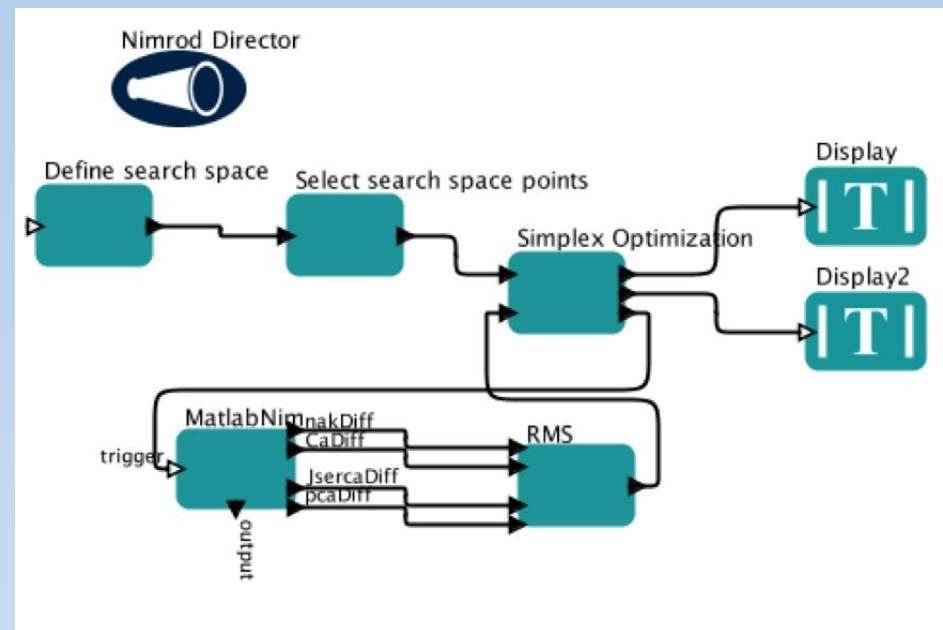
- Using a MATLAB actor provided by Kepler
- Local spawn
 - Multiple thread ran concurrently on a computer with 8 cores (2 x quads)
 - Workflow execution was just under 8 times faster
- Remote Spawn
 - 100's – 1000's of remote processes



Nimrod/OK Workflows



- Nimrod/K supports parallel execution
- General template for search
 - Built from key components
- Can mix and match optimization algorithms





From Clusters, to Grids, to Clouds

Nimrod over Clusters



Jobs / Nimrod experiment



Nimrod

Actuator, e.g., SGE, PBS, LSF, Condor



Local Batch System

axceleon™
high-performance distributed computing software

HOME | CORPORATE | PRODUCTS | SOLUTIONS | PARTNERS | DOWNLOADS | SUPPORT | NEWS | CONTACT

EnFuzion CLOUDFUZION

- Why EnFuzion
- EnFuzion Features
- EnFuzion Benefits
- Technical Specs

- User Interface
- EnFuzion Data Sheet
- EnFuzion3D Data Sheet
- CloudFuzion Data Sheet
- Applications Notes
- Demo Applications

DOWNLOAD EnFuzion

DOWNLOAD CloudFuzion

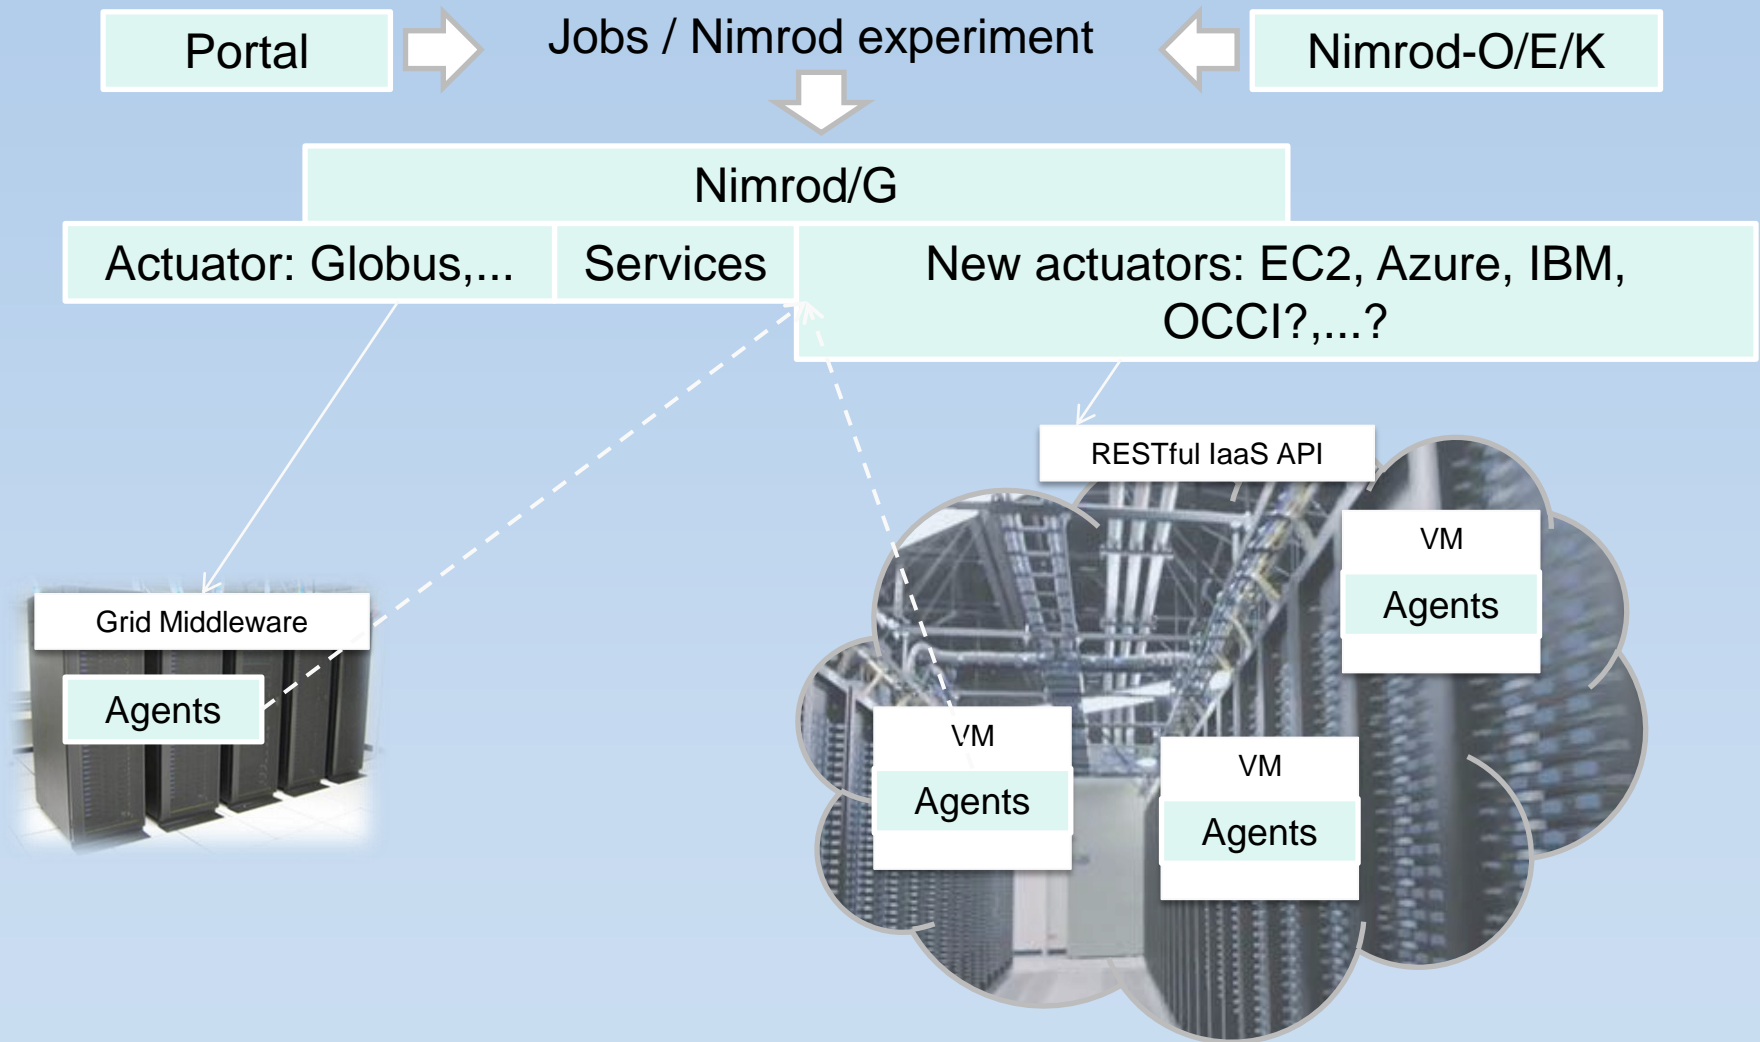
- [illegible]

Leveraging Cloud Infrastructure

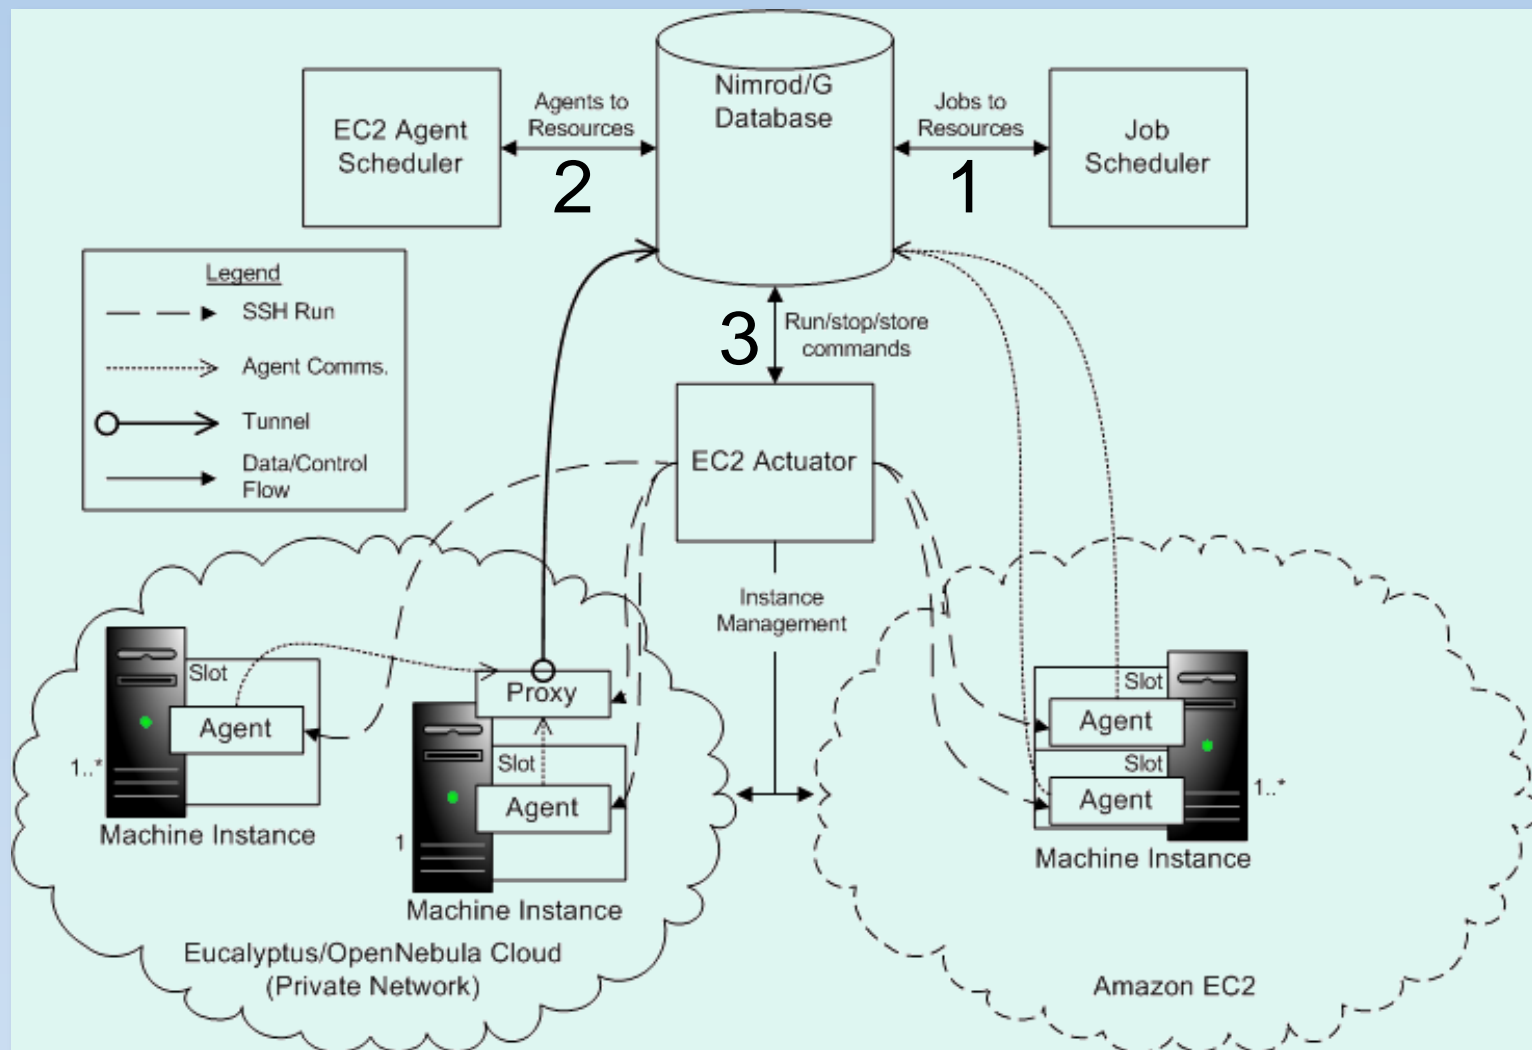


- Centralisation is easier
 - (Clusters vs Grid)
- Virtualisation improves interoperability and scalability
 - Build once, run everywhere
- Computational economy, for real
 - Deadline driven
 - "I need this finished by Monday morning!"
 - Budget driven
 - "Here's my credit card, do this as quickly and cheaply as possible."
- Cloud bursting
 - Scale-out to supplement locally and nationally available resources

Integrating Nimrod with IaaS



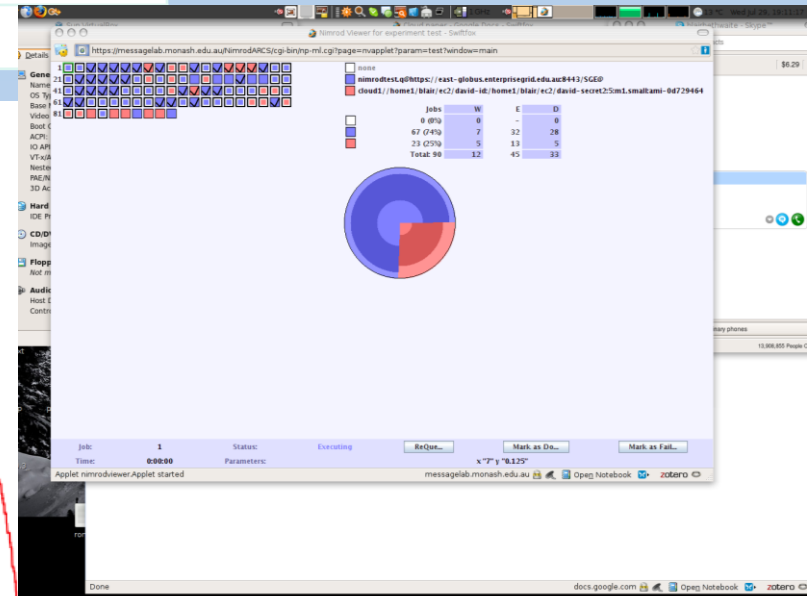

Integrating Nimrod with IaaS



Cloud bursting



Resource	#jobs completed	Total job time (h:m:s)	μ / σ Job runtime (mins)
East	818	1245:37:23	91/5.7
EC2	613	683:34:05	67/14.2





Progress to date

Azure implementation

PaaS is trickier...




- More variety
 - Azure vs AppEngine
- Designed for web-app hosting
 - Nimrod provides a generic execution framework
- Higher level PaaS too prescriptive
 - AppEngine: Python and Java only

Nimrod-Azure Mk.1

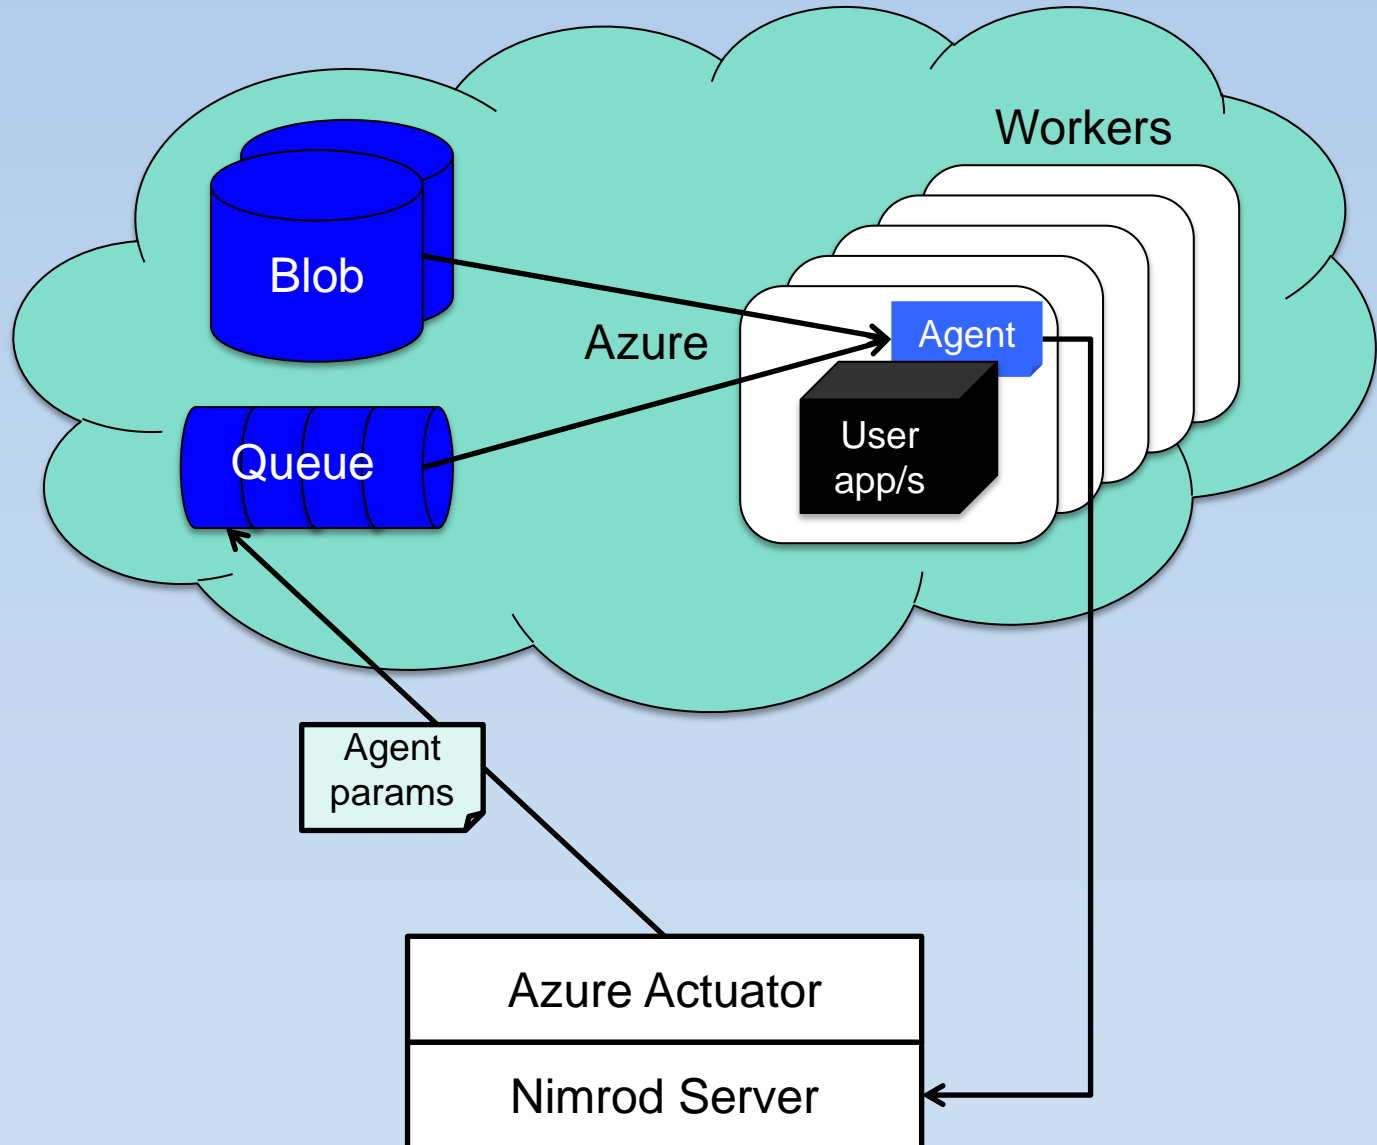


- Nimrod server runs on a Linux box external to Azure
- Nimrod-Azure actuator module contains the code for managing Nimrod agents on Azure
 - pre-defined minimal NimrodWorkerService cspkg;
 - library for speaking XML over HTTP with the Azure Storage and Management REST APIs

Integrating Nimrod with Azure




Integrating Nimrod with Azure



New Application Drivers



- Volcanic ash dispersion modelling
 - Pt.1 (NG-Tephra) completed with EC2
 - Pt.2 (Ceniza) to run on Azure
- Energy economics of distributed generation technology
 - to run on Azure




Acknowledgements

MeSsAGE Lab



- Faculty Members
 - Jeff Tan
 - Maria Indrawan
- Research Fellows
 - Blair Bethwaite
 - Slavisa Garic
 - Jin Chao
- Admin
 - Rob Gray
- Current PhD Students
 - Shahaan Ayyub
 - Philip Chan
 - Colin Enticott
 - ABM Russell
 - Steve Quinette
 - Ngoc Dinh (Minh)
- Completed PhD Students
 - Greg Watson
 - Rajkumar Buyya
 - Andrew Lewis
 - Nam Tran
 - Wojtek Goscinski
 - Aaron Searle
 - Tim Ho
 - Donny Kurniawan
 - Tirath Ramdas
- Funding & Support
 - Axceleon
 - Australian Partnership for Advanced Computing (APAC)
 - Australian Research Council
 - Cray Inc
 - CRC for Enterprise Distributed Systems (DSTC)
 - GrangeNet (DCITA)
 - Hewlett Packard
 - IBM
 - Microsoft
 - Sun Microsystems
 - US Department of Energy



Questions?



More information:

<http://messagelab.monash.edu.au>