

Monitoring of Castor at RAL

Overview of current monitoring

- Nagios
- LSF monitoring – see Brian's talk later
- Statistical – TSBN / ganglia and cacti
- Diagnostic – checkreplicas
- Of servers running castor

NAGIOS

Capabilities of Nagios

- Browser frontend
- API access – Mimic created by James Adams RAL
- Email alerting
- Structured around groups and services
- Scheduled downtimes for hosts / services / groups
- Acknowledgements of alerts
- Comments to alerts and scheduled downtimes
- Triggers to execute external scripts
- Dependencies between alerts
- User authentication
- Supports snmp, plugins to nrpe, nsca forwarding
- Supports pager and email alerts

Overnight callout

- Triggers
 - Detected errors will execute a script that beeps a pager and sends a message
- Support person
 - refers to callout documentation on how to fix problems
 - Fixes castor from home
 - logs actions in RT ticketing (helpdesk) system
- This works well !

Current Network Status - Microsoft Internet Explorer

Address: https://nagios.gridpp.rl.ac.uk/nagios/cgi-bin/status.cgi?hostgroup=castor-linux&style=grid

Current Network Status
 Last Updated: Wed Feb 18 09:41:07 GMT 2009
 Updated every 90 seconds
 nagios@2.11 - www.nagios.org
 Installed in
 /C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney.ketley
 Service checks are disabled

Up	Down	Unreachable	Pending
325	5	0	0
All Problems		All Types	
5		330	

Ok	Warning	Unknown	Critical	Pending
10114	172	6	184	12
All Problems		All Types		
362		10488		

[View Status Grid For All Host Groups](#)
[View Service Status Detail For This Host Group](#)
[View Host Status Detail For This Host Group](#)
[View Status Overview For This Host Group](#)
[View Status Summary For This Host Group](#)

Status Grid For Host Group 'castor-linux'

[Castor linux servers \(castor-linux\)](#)

Host	Services	Actions
castor100	SWAP c linux cpu c linux load c linux proc atd c linux proc cron c linux proc omond c linux proc klogd c linux proc rntd c linux proc portmap c linux proc ssh c linux proc syslogd c linux space root c linux space swap c linux space tmp c linux space usr c linux space var c linux stale-ntfs c linux vsz c ns ora-conn c ns proc Cupvdaemon c ns proc msqdaemon c ns proc nsdaemon c ns proc vdmnserv c ns proc vmoraemon c ns showq all down c ns showq down c ns showq queued c ns showq unkn	
castor101	SWAP c linux cpu c linux load c linux proc atd c linux proc cron c linux proc omond c linux proc klogd c linux proc rntd c linux proc portmap c linux proc ssh c linux proc syslogd c linux space root c linux space swap c linux space tmp c linux space usr c linux space var c linux stale-ntfs c linux vsz c midhunter missing c stager carbemiar c stager free-tapes free c stager loq ora32102 c stager proc MidHunter new c stager proc rtsrver new c stager proc rtsrver priv new c stager proc rtcplntd c stager proc stager new	
castor102	SWAP c dft proc dlsrver c dft proc expertd c dft proc jobmgr c linux cpu c linux load c linux proc atd c linux proc cron c linux proc omond c linux proc klogd c linux proc rntd c linux proc portmap c linux proc ssh c linux proc syslogd c linux space lst c linux space root c linux space swap c linux space tmp c linux space usr c linux space var c linux stale-ntfs c linux vsz c proc httpd	
	SWAP c linux cpu c linux load c linux proc atd c linux proc cron c linux proc omond c linux proc klogd c linux proc rntd c linux proc portmap c linux proc puppet noc c linux proc ssh c linux proc syslogd c linux space lst c linux space root c linux space swap c linux space tmp c linux space usr c linux space var c linux stale-ntfs c linux vsz c lst bhsts ok c lst bhsts unavail c lst bhsts done c lst bhsts exit	

https://nagios.gridpp.rl.ac.uk/nagios/cgi-bin/extinfo.cgi?type=2&host=castor103&service=c+linux+space+root

Note the easy-to-see errors.

Overview – grid format

15 servers monitored in this view

File Edit View Favorites Tools Help

Address <https://nagios.gridpp.rl.ac.uk/nagios/cgi-bin/extinfo.cgi?type=6#HOSTDOWNT> Go

Links G Y! Clst YIM GM YIBT BBC RT Foot Incident NGS inc C Plan OvW Reds

All Host and Service Scheduled Downtime
 Last Updated: Wed Feb 18 09:44:34 GMT 2009
 Updated every 90 seconds
 Nagios® 2.11 - www.nagios.org
 Logged in as /C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney ketley
 - Service checks are disabled

[[Host Downtime](#) | [Service Downtime](#)]

Scheduled Host Downtime
[Schedule host downtime](#)

Host Name	Entry Time	Author	Comment	Start Time	End Time	Type	Duration	Downtime ID	Trig ID
ccsc06	03-10-2008 17:38:17	/C=UK/O=eScience/OU=CLRC/L=RAL/CN=bonny strong	Server converted to test repack to move to gen stager host	03-10-2008 17:37:25	03-10-2009 17:37:25	Fixed	365d 0h 0m 0s	5026	N/A
srm0375	17-10-2008 16:01:32	/C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney ketley	no longer in use	17-10-2008 16:01:16	17-10-2009 18:01:16	Fixed	365d 2h 0m 0s	5029	N/A
drno	18-12-2008 09:57:55	/C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney ketley	damaged sata cable problem	18-12-2008 09:57:24	01-03-2009 11:57:24	Fixed	73d 2h 0m 0s	7176	N/A
	18-12-	/C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney ketley	being rebuilt	18-12-	01-03-		73d 2h		

So alerts are not sent during downtime.

Note comments.

Scheduling of downtimes

Multiple servers put into downtime via hostgroup.

Links >>

Current Network Status

Last Updated: Wed Feb 18 09:40:31 GMT 2009
 Updated every 90 seconds
 Nagios® 2.11 - www.nagios.org
 Logged in as /C=UK/O=eScience/OU=CLRC/L=RAL/CN=cheney ketley
 - Service checks are disabled

[View Service Status Detail For All Host Groups](#)
[View Host Status Detail For This Host Group](#)
[View Status Overview For This Host Group](#)
[View Status Summary For This Host Group](#)
[View Status Grid For This Host Group](#)

Host Status Totals

Up	Down	Unreachable	Pending
325	5	0	0

All Problems | **All Types**

5 | 330

Service Status Totals

Ok	Warning	Unknown	Critical	Paused
9797	170	6	182	

All Problems | **All Types**

358 | 10164

Service Status Details For Host Group 'castor-linux'

Display Filters:
 Host Status Types: All
 Host Properties: Any
 Service Status Types: Critical
 Service Properties: Any

Host	Service	Status	Last Check	Duration	Attempt	Status Information
castor103	c lsf bhosts unavail	CRITICAL	18-02-2009 09:15:15	0d 16h 14m 52s	3/3	-1- processes in state - unavail- levels=0,1,1
ccsc05	c lsf bhosts unavail	CRITICAL	18-02-2009 08:39:04	15d 15h 32m 10s	3/3	-2- processes in state - unavail- levels=0,1,1
	c lsf closed or disabled	CRITICAL	18-02-2009 09:09:07	70d 10h 14m 21s	3/3	29 Nodes are closed
ccsc06	c mighunter missing	CRITICAL	18-02-2009 09:33:00	76d 9h 23m 1s	3/3	at least one service class has not enough mighunters running.

Red alerts only

Nagios can show only the errors if desired
 Useful for easy monitoring

Address <https://nagios.gridpp.rl.ac.uk/nagios/cgi-bin/extinfo.cgi?type=2&host=castor103&se> Go

Links [G](#) [YI](#) [Clist](#) [YIM](#) [GM](#) [YIBT](#) [BBC](#) [RT](#) [Foot](#) [Incident](#) [NGS inc](#) [C Plan](#) [OvW](#) [Reds](#)

[View Availability Report For This Service](#)
[View Notifications For This Service](#)











130.246.183.219

Service State Information

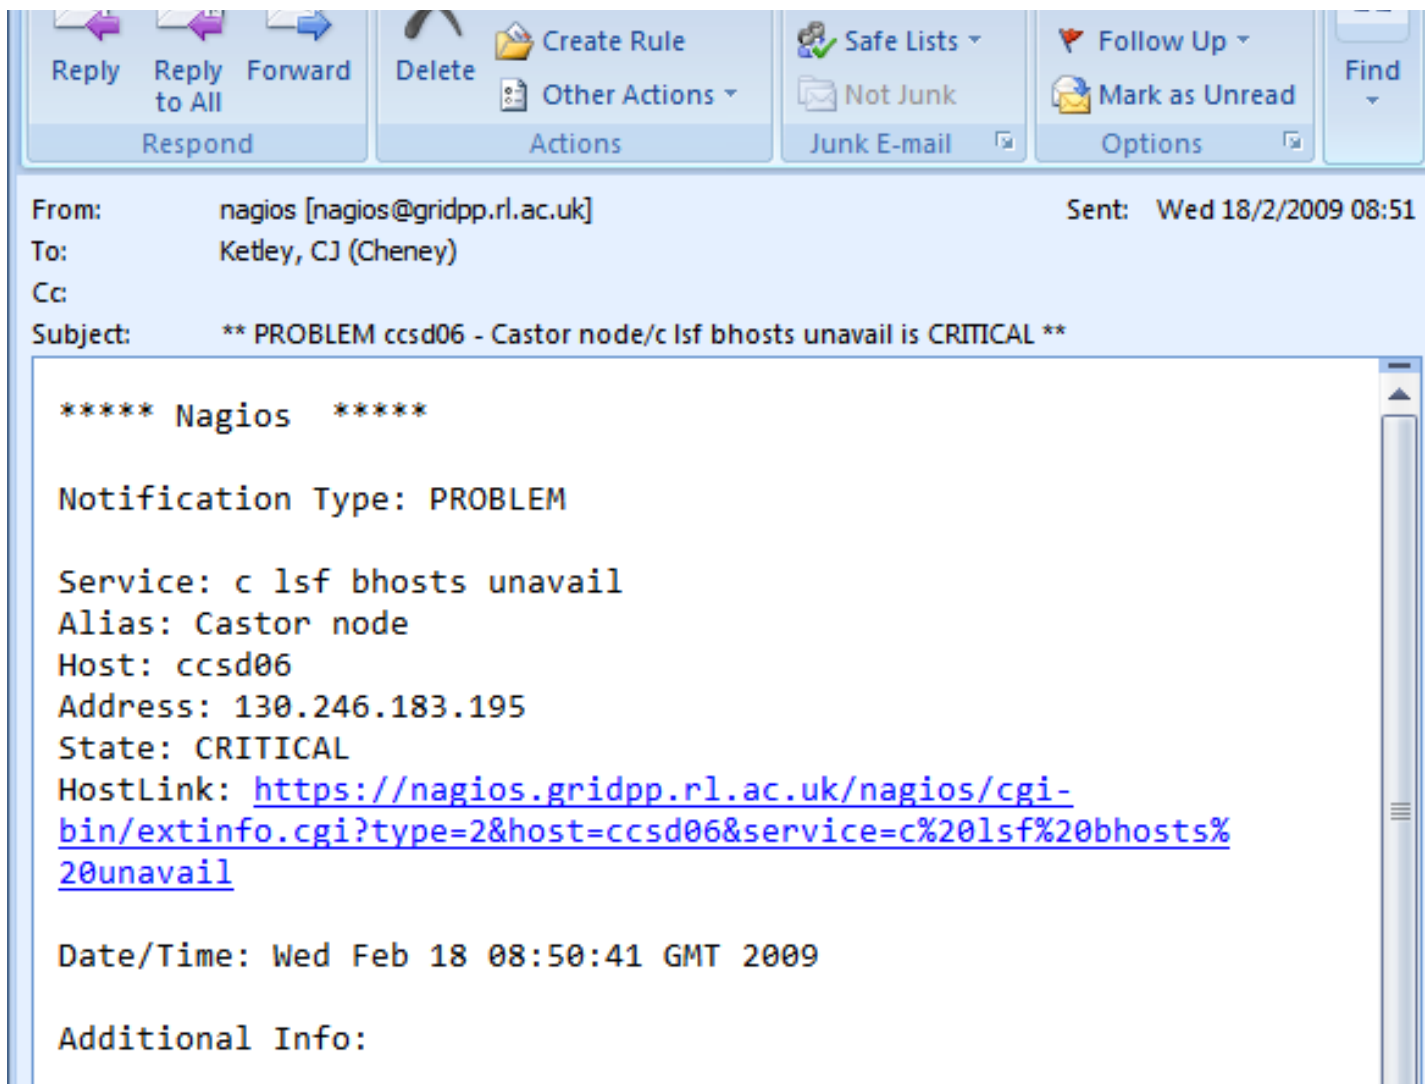
Current Status:	CRITICAL
Status Information:	-1- processes in state -unavail- levels=0,1,1
Performance Data:	
Current Attempt:	3/3
State Type:	HARD
Last Check Type:	ACTIVE
Last Check Time:	18-02-2009 09:15:15
Status Data Age:	0d 0h 27m 44s
Next Scheduled Active Check:	N/A
Latency:	1957.377 seconds
Check Duration:	0.282 seconds
Last State Change:	17-02-2009 17:25:36
Current State Duration:	0d 16h 17m 23s
Last Service Notification:	18-02-2009 08:09:34
Current Notification Number:	2
Is This Service Flapping?	N/A
Percent State Change:	N/A
In Scheduled Downtime?	NO
Last Update:	18-02-2009 09:42:39

Active Checks:	DISABLED
Passive Checks:	ENABLED
Obsessing:	ENABLED
Notifications:	ENABLED
Event Handler:	ENABLED

Service Commands

-  [Enable active checks of this service](#)
-  [Submit passive check result for this service](#)
-  [Stop accepting passive checks for this service](#)
-  [Stop obsessing over this service](#)
-  [Acknowledge this service problem](#)
-  [Disable notifications for this service](#)
-  [Delay next service notification](#)
-  [Schedule downtime for this service](#)
-  [Disable event handler for this service](#)
-  [Disable flap detection for this service](#)

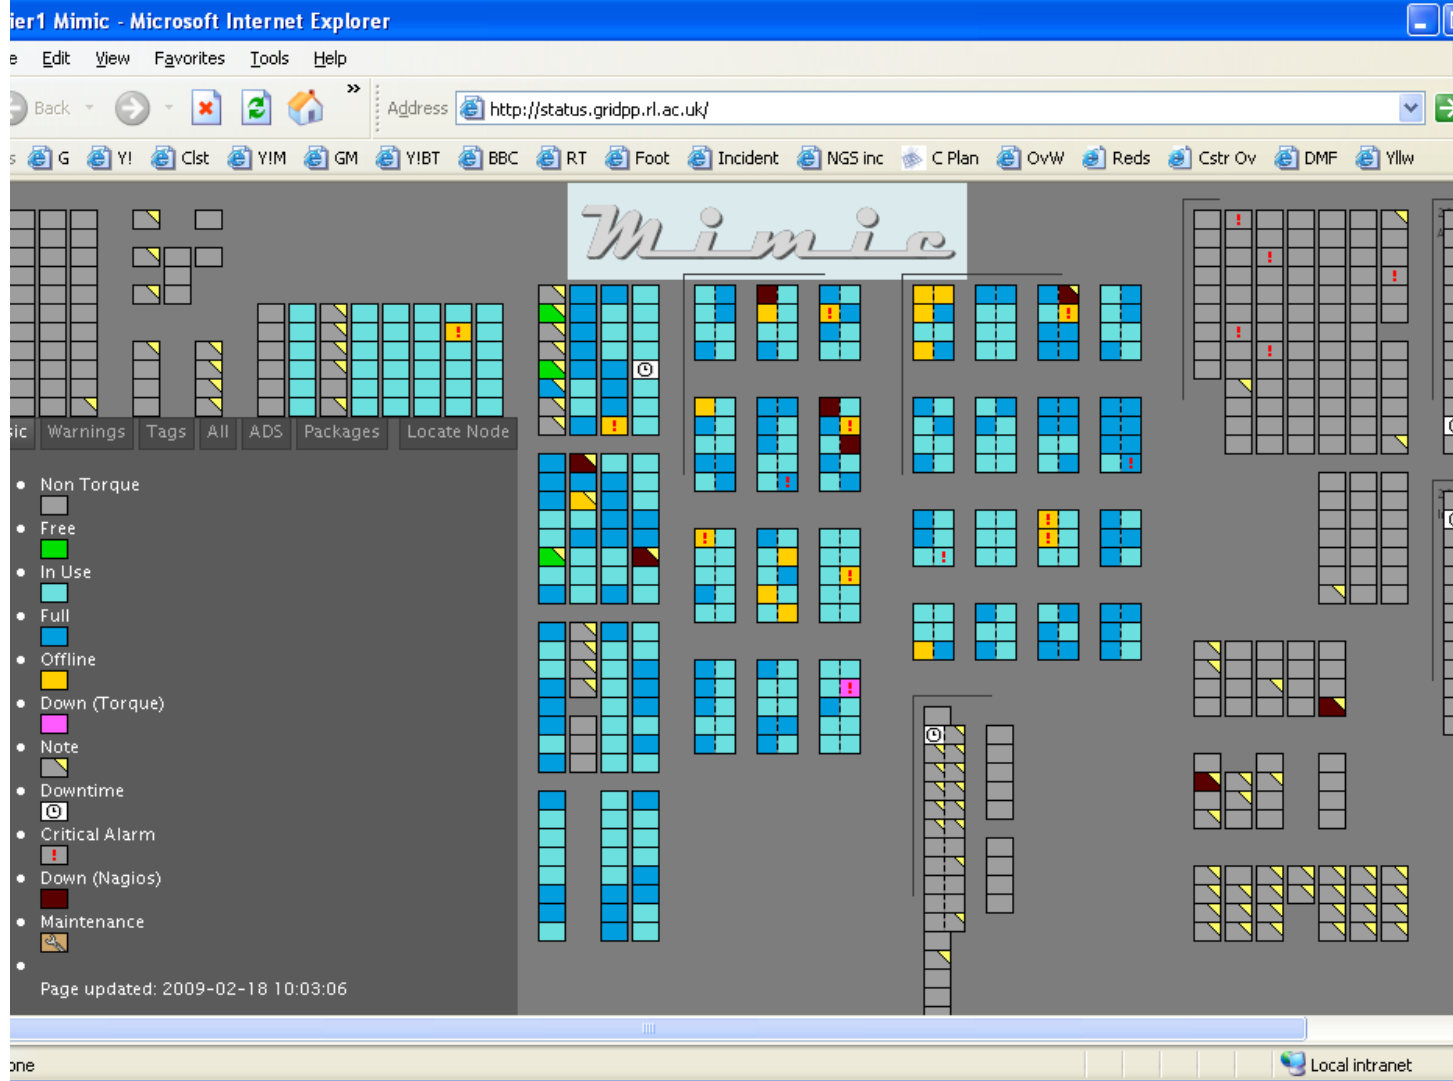
Dealing with an alert Note that acknowledgement is not the only way
 Note a comment is also displayed but not shown here



Example of an email alert Note that extra text can be added which is useful for explaining the problem.

Methods of detecting problems

- NRPE
 - Custom plugins can be written
 - Standard plugins cover very many situations
- SNMP
- NSCA
 - Send a message from a user-written script with an alert code
 - Useful where nrpe cannot be installed



Mimic – created by RAL Server room layout at a glance

Summary of Nagios experiences

- Overall – very positive
- Plus points
 - Supports overnight callout via pager alerts
 - A powerful tool
 - Scales well – into hundreds of servers so far
 - Supports many architectures / OS's – besides castor
- Negative points
 - Re-configuration to add server requires restart
 - Dependency checks not fully functional

OTHER MONITORING TOOLS

TSBN

- Extract from tape server logs with analysis by MS Excel pivot tables
- Perl extract written by Tim Bell (Cern)
- MS Excel analysis written by Cheney Ketley (RAL)
cheney.ketley@stfc.ac.uk

TSBN-analysis.v10-8.xls [Compatibility Mode] - Micr... PivotTable Tools

Home Insert Page Layout Formulas Data Review View Options Design

Font: Arial, 10, Bold, Italic, Underline, Text Color, Background Color, Font Color

Alignment: Left, Center, Right, Top, Bottom, Merge, Unmerge

Number: General, Percentage, Decimals, Thousands Separator

Styles: Conditional Formatting, Format as Table, Cell Styles

Cells: Insert, Delete, Format

Editing: Sort & Filter, Find & Select

Security Warning: Data connections have been disabled

A14 fx ctsd11

Thrashing by tape server

tapeserver	KB transferred	File count	Mounts	Position count	Releases	Tape mounts per file	Tape positionings per file	Tape releases per file	Gi
castor201	988,142,245	8807	640	8807	640	0.1	1.0	0.1	
castor208	11,210,872,968	4864	1448	4834	1417	0.3	1.0	0.3	
ctsc12	3,180,355,712	1644	282	1638	282	0.2	1.0	0.2	
ctsc16	294,843,395	468	786	467	932	1.7	1.0	2.0	
ctsc17	353,144,454	502	824	502	976	1.6	1.0	1.9	
ctsc19	2,607,096,523	1135	303	1134	304	0.3	1.0	0.3	
castor200	3,831,343,014	1666	329	1669	329	0.2	1.0	0.2	
ctsc18	3,517,252,944	1313	322	1308	321	0.2	1.0	0.2	
ctsc20	1,170,644,049	418	131	416	132	0.3	1.0	0.3	
ctsd11	2,871,467,932	1140	309	1136	308	0.3	1.0	0.3	
ctsd12	1,284,752,785	463	122	463	120	0.3	1.0	0.3	
ctsd13	3,770,719,570	1727	309	1715	310	0.2	1.0	0.2	
ctsd14	1,306,535,924	450	130	450	130	0.3	1.0	0.3	
ctsd15	11,603,032,877	5061	1413	5032	1382	0.3	1.0	0.3	
ctsd16	3,711,904,536	1864	312	1857	311	0.2	1.0	0.2	
Grand Total	51,702,108,928	31522	7660	31428	7894	0.2	1.0	0.3	

Library stats | Tapeserver thrashing per month | **Tapeserver thrashing** | VO thrashing | Top 3 | Worst 20 mounts | Trends | Raw data

100%

TSBN - analysis by tape server File screen can be used to select particular data and ranges

TSBN-analysis.v10-8.xls [Compatibility Mode] - Microsoft Excel

Layout Formulas Data Review View

Font Font Color Background Color

Alignment

Number

Styles

Options...

C D E F G H

Tape library stats

	KB transferred	Tape Xfer MB/sec	Tape MB/sec elapsed	Mounts	Position count	Releases
	10,602,008,828	58.4	21.6	2,752	5,359	3,041
	41,100,100,100	90.4	30.2	4,908	26,069	4,841
	51,702,108,928	81.3	27.9	7,660	31,428	7,891
	51,702,108,928	81.3	27.9	7,660	31,428	7,891

preserver thrashing per month Tapeserver thrashing VO thrashing Top 3 Worst 20 mounts

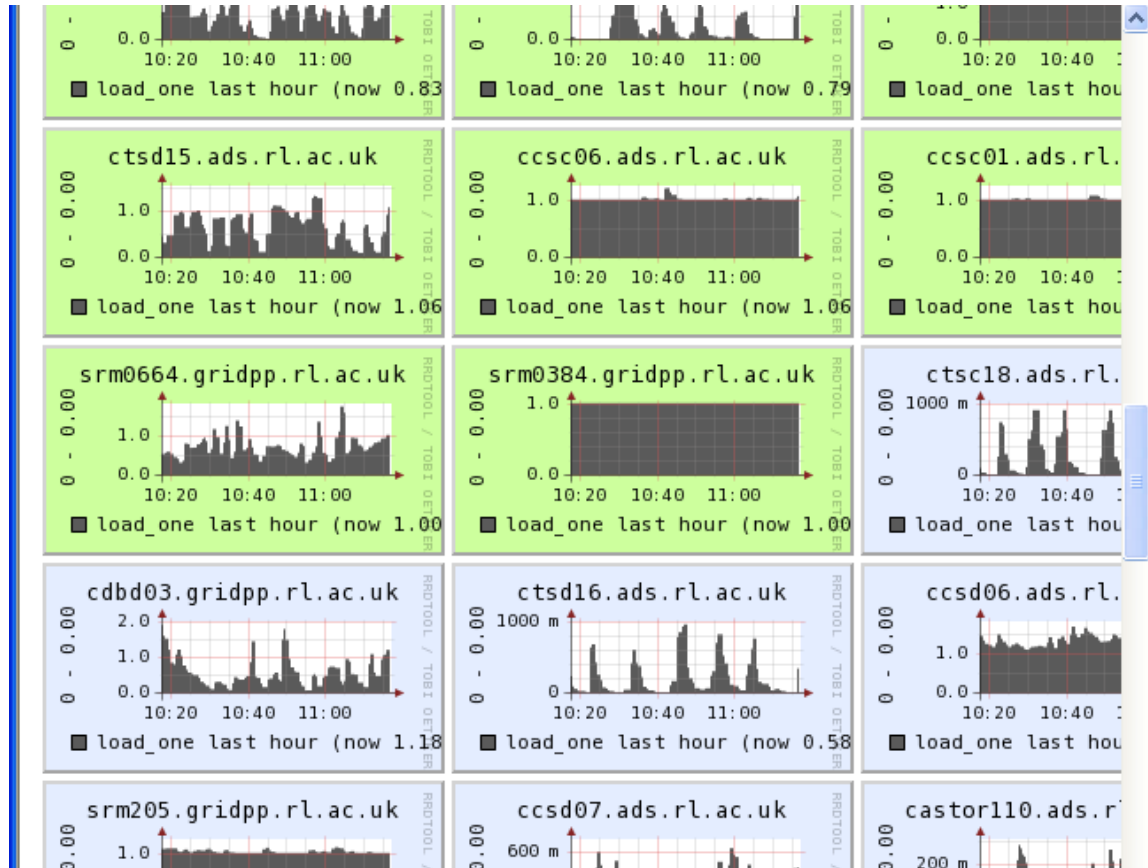
TSBN - overview

Note there are many tabs at page bottom for different views on data

Provided by Excel pivot tables

Ganglia, Cacti, Gridview

- Ganglia monitors server performance, Cacti is for the network.
- Ganglia also shows feeds into Castor e.g. FTS
- Gridview gives a broad view



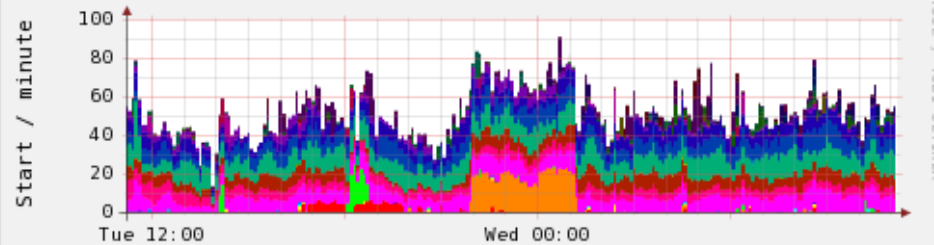
Ganglia plots

This shows performance graphs by server or by service

RALLCG2 as Source or Destination

Started Transfers (per minute):

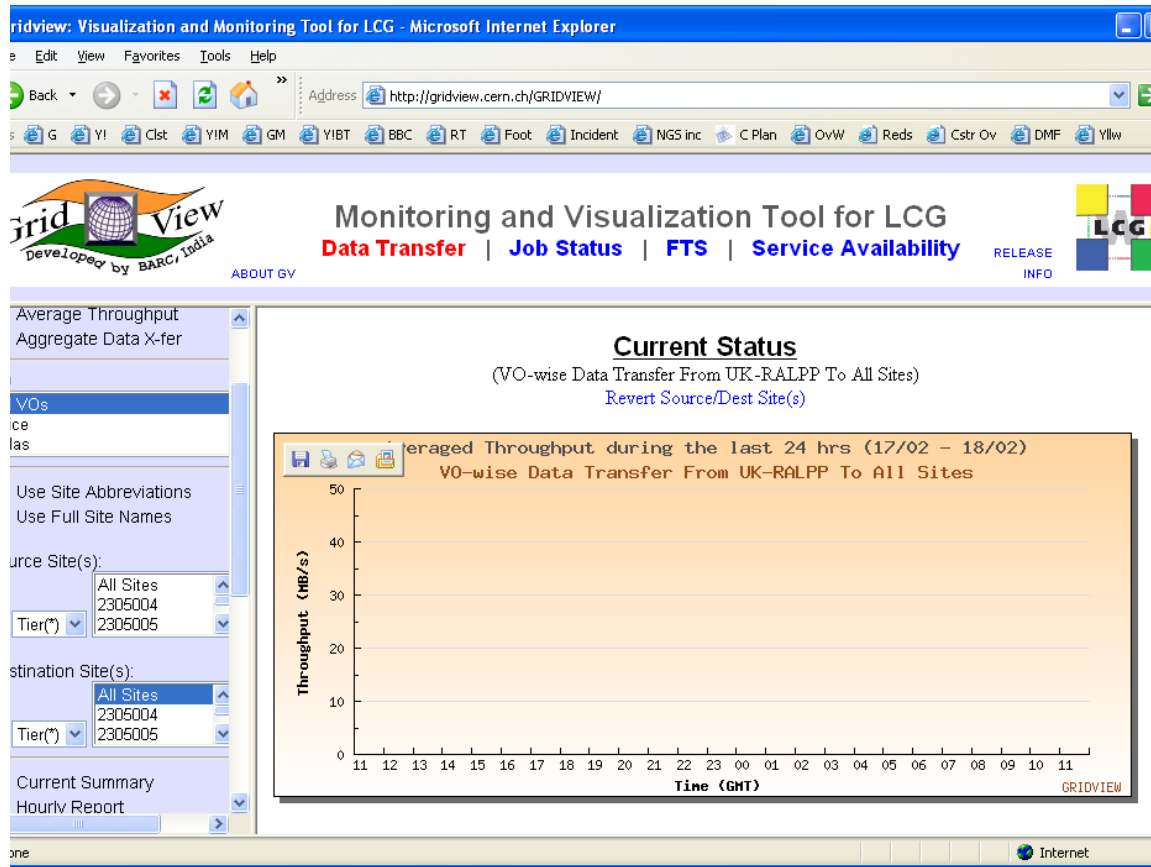
-LCG2 Tier-1 FTS Started Transfers (RALLCG2 as source/destination)



■ BNLLCG2-RALLCG2 (0.15; average = 2.21)
■ FZKLCG2-RALLCG2 (0.55; average = 9.40)
■ IN2P3CC-RALLCG2 (1.40; average = 0.76)
■ INFNT1-RALLCG2 (2.00; average = 5.94)
■ NDGFT1-RALLCG2 (1.78; average = 0.39)
■ PIC-RALLCG2 (0.55; average = 0.58)
■ RALLCG2-UKILT2BRUNEL (0.24; average = 0.22)
■ RALLCG2-UKILT2QMUL (0.20; average = 0.26)
■ RALLCG2-UKILT2RHUL (6.28; average = 7.44)
■ RALLCG2-UKINORTHGRIDLANCSHEP (4.36; average = 4.55)
■ RALLCG2-UKINORTHGRIDLIVHEP (6.30; average = 5.72)
■ RALLCG2-UKINORTHGRIDMANHEP (0.20; average = 0.34)
■ RALLCG2-UKINORTHGRIDSHEFHEP (0.40; average = 0.35)
■ RALLCG2-UKISCOTGRIDDURHAM (0.20; average = 0.19)
■ RALLCG2-UKISCOTGRIDECDF (0.40; average = 0.38)
■ RALLCG2-UKISCOTGRIDGLASGOW (12.57; average = 9.70)

Another ganglia plot

Useful for analysis by feed e.g. FTS



Gridview

This gives an overview of site performance

DIAGNOSTICS

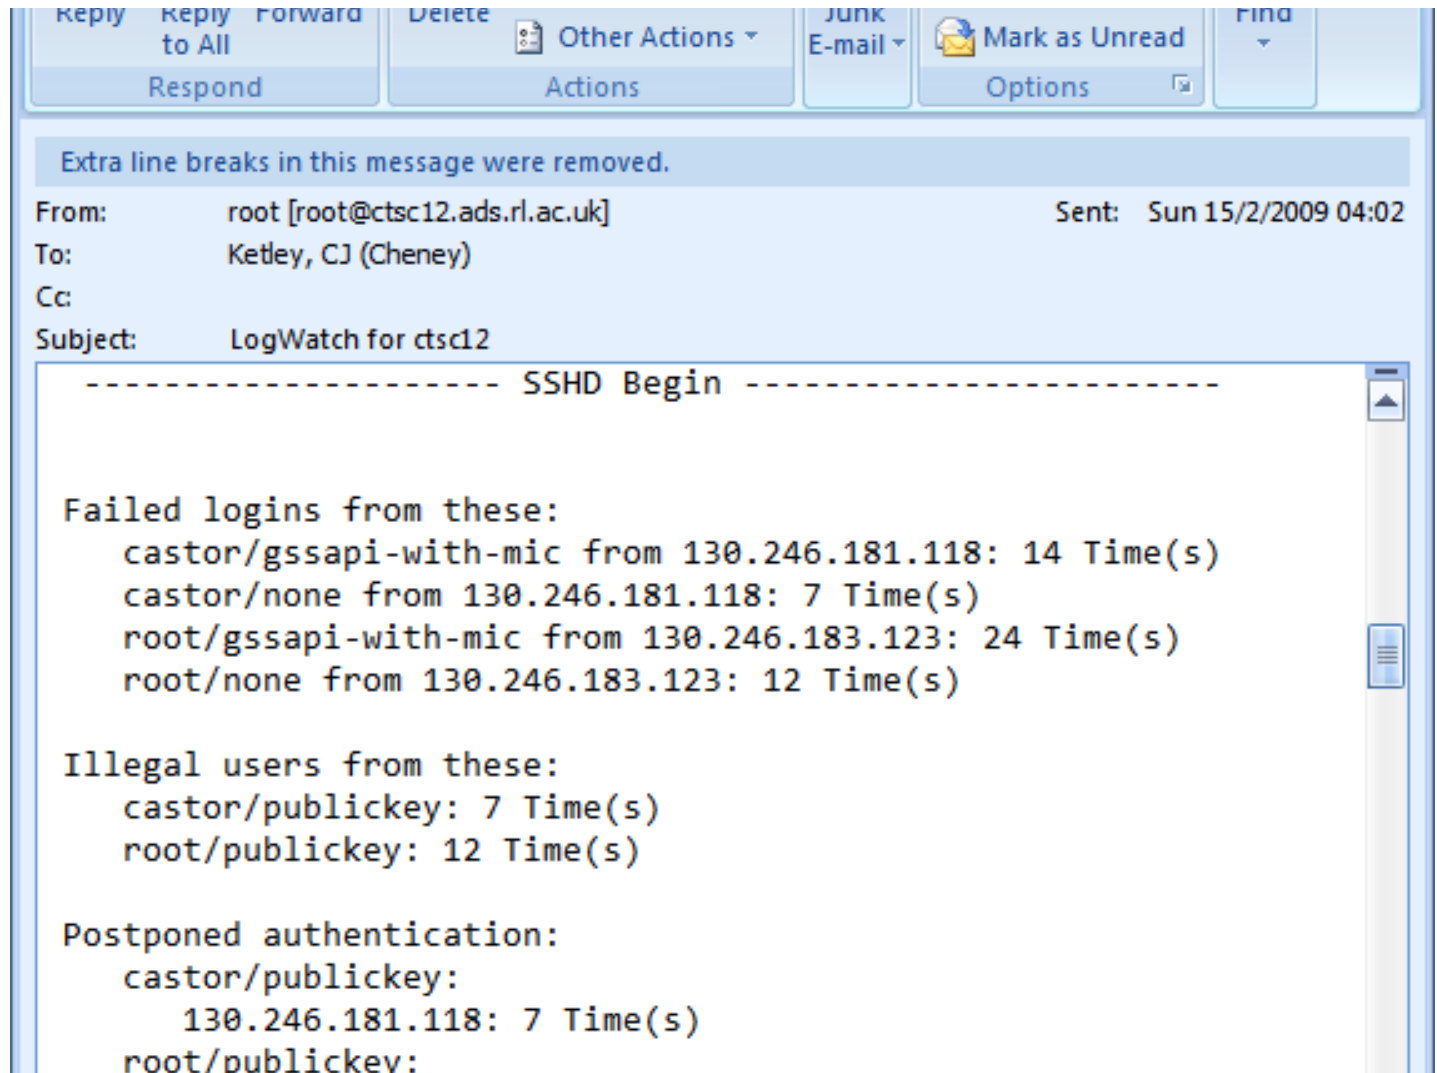
Diagnostics

- Cern checkreplicas.pl (perl script)
 - <https://cern.ch/castor/DIST/checkreplicas>
- Webmin wrapper to checkreplicas with hyperlinks to drill down into data
 - Broken at the moment !
 - Contact is cheney.ketley@stfc.ac.uk
- Other ad hoc tools exist – sql scripts written by Bonny
 - Contact is chris.kruk@stfc.ac.uk

SERVER-SIDE

Operating system and environment

- Logwatch – covers diskspace, logins, email usage, cron usage, packages updated, and other messages not recognised by logwatch
- IPMI being rolled out (for server physical alerts and remote control)
- Nagios also monitors the operating system components
- Backups – email messages from Amanda



Logwatch message

Shows key information regarding operating system
But causes problems with email overload...

FUTURE

Future direction

- More Nagios plugins
- Performance metrics
- More diagnostic tools

Contacts

- Jonathan Wheeler is main Nagios administrator but not responsible for Castor
 - Jonathan.wheeler@stfc.ac.uk
- Cheney Ketley is secondary Nagios administrator and supports Castor
 - Cheney.ketley@stfc.ac.uk
- James Adams wrote Mimic
 - James.adams@stfc.ac.uk