

GridPP

UK Computing for Particle Physics

Verify scheduling at the Tier1

James Thorne
STFC RAL



Science & Technology
Facilities Council



- Exercises idle disk servers
- Earlier detection of bad disks
- Can correct errors or force rebuilds

- Not disk “scrubbing”
- Performance hit
 - 6% on writes
 - 14% on reads (3ware, lowest priority)
- Takes hours or days to run
- Different implementation on each controller type

- Reduces performance hit on a disk pool
 - Only verify 1/5 of a disk pool simultaneously
 - Use disk server database to determine which servers to verify
- Ensures that every array is verified regularly
 - Every 5 weeks
- Record the date of verify for reference
 - In Overwatch



PSCS Overwatch Deployment

[Home Page](#) [Add New System](#)

What, Where and Why

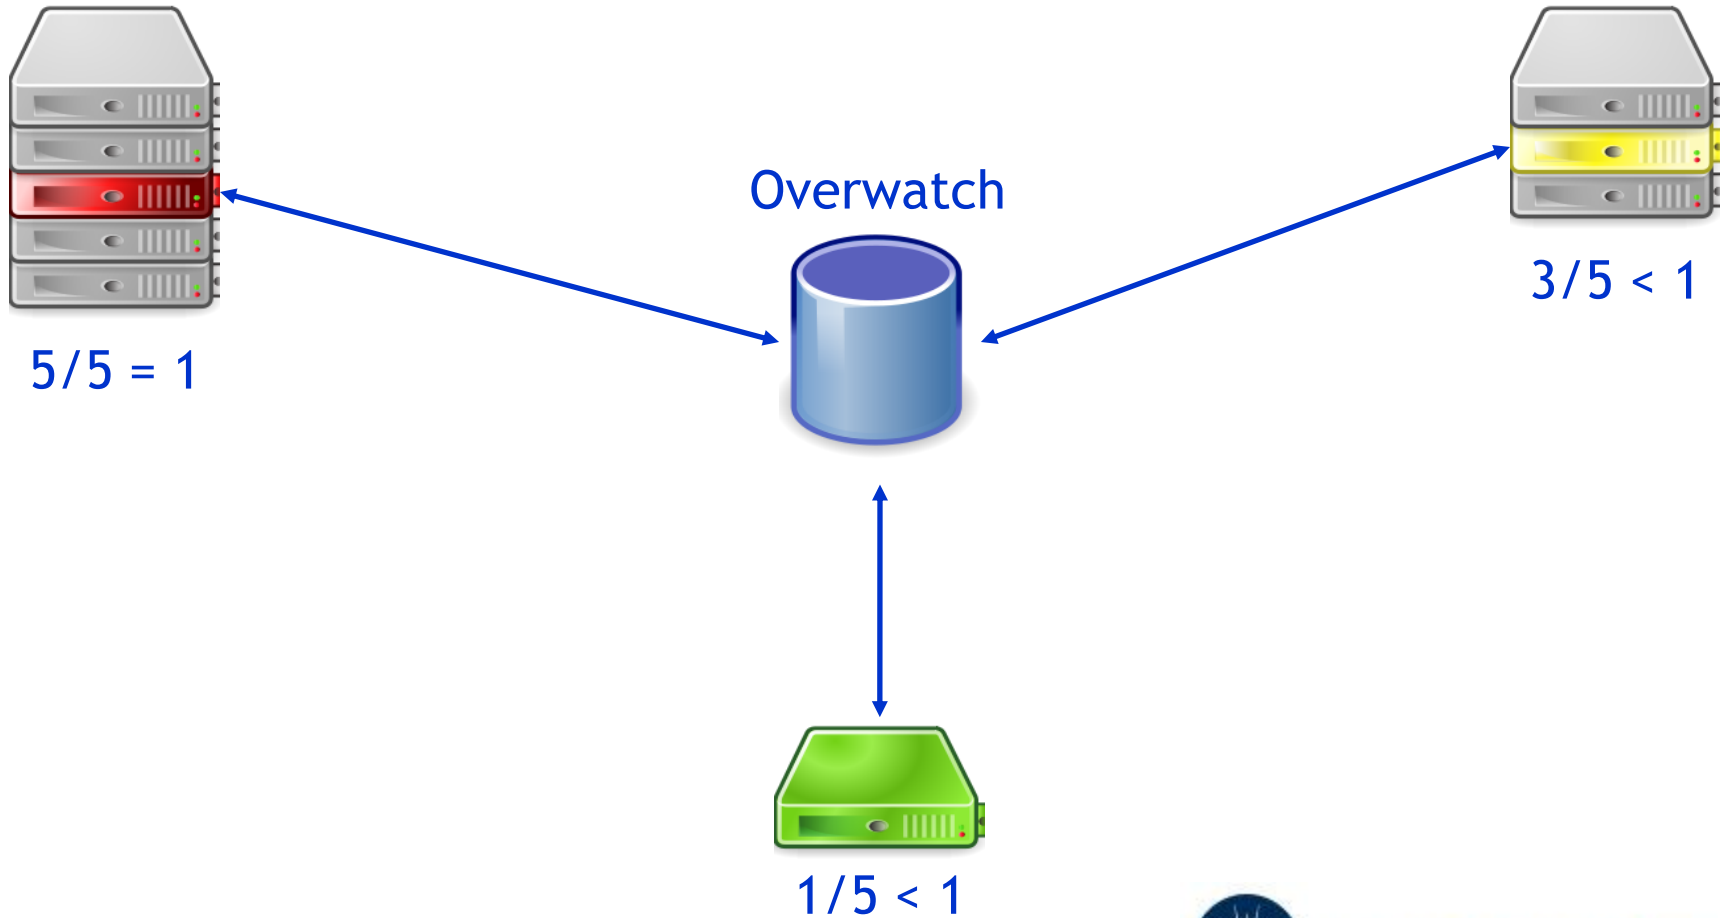
View a System: View by Instance: View by Disk Pool: View by Status: View by Nagios Grouping: Update a System:

Delete a System:

System Name	Current Status	Normal Status	Current Service Team	Service Type	Castor Instance	Virtual Organisation	Disk Pool	Size (Tb)	File Systems	Hardware Type	Managed by Puppet	Comments	Last Verified	NORC	Update	Delete
qdss211	Ready	Deployment	fabric	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES	(22/12/08) fsprobe errors #39374 (6/11/09) back to ready	2009-06-09	dit0	Update	Delete
qdss218	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES	RT# 40131 Moved into ready state	2009-06-09	dit0	Update	Delete
qdss223	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-09	dit0	Update	Delete
qdss229	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-09	dit0	Update	Delete
qdss230	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-09	dit0	Update	Delete
qdss231	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-09	dit0	Update	Delete
qdss232	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss233	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss234	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss235	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss236	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss237	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-15	dit0	Update	Delete
qdss238	Ready	Ready	castor	castor	Atlas	Atlas	atlasNonProd	8.90	3	V07a	YES		2009-06-22	dit0	Update	Delete



- TAVS (Tier1 Array Verify Scheduler)
 - Pointless acronym
 - Weekly Perl cron job
 - Runs on each disk server
 - Runs verify if server is in oldest fifth and not verified in last 28 days
 - CLI wrapper script
 - Detects the installed CLI and runs it.
 - On 3ware, verifies only run 11 AM Monday to noon Thursday
 - Provided by RPM



- Supported RAID card (currently most 3ware and Areca cards)
- Database to link host to verify date
- Appropriate CLI
 - tw_cli
 - areca_cli

- Went live just before STEP so only deployed on non-production disk servers
- Haven't noticed more failures yet
- To be installed on the rest of the disk servers at the end of July