

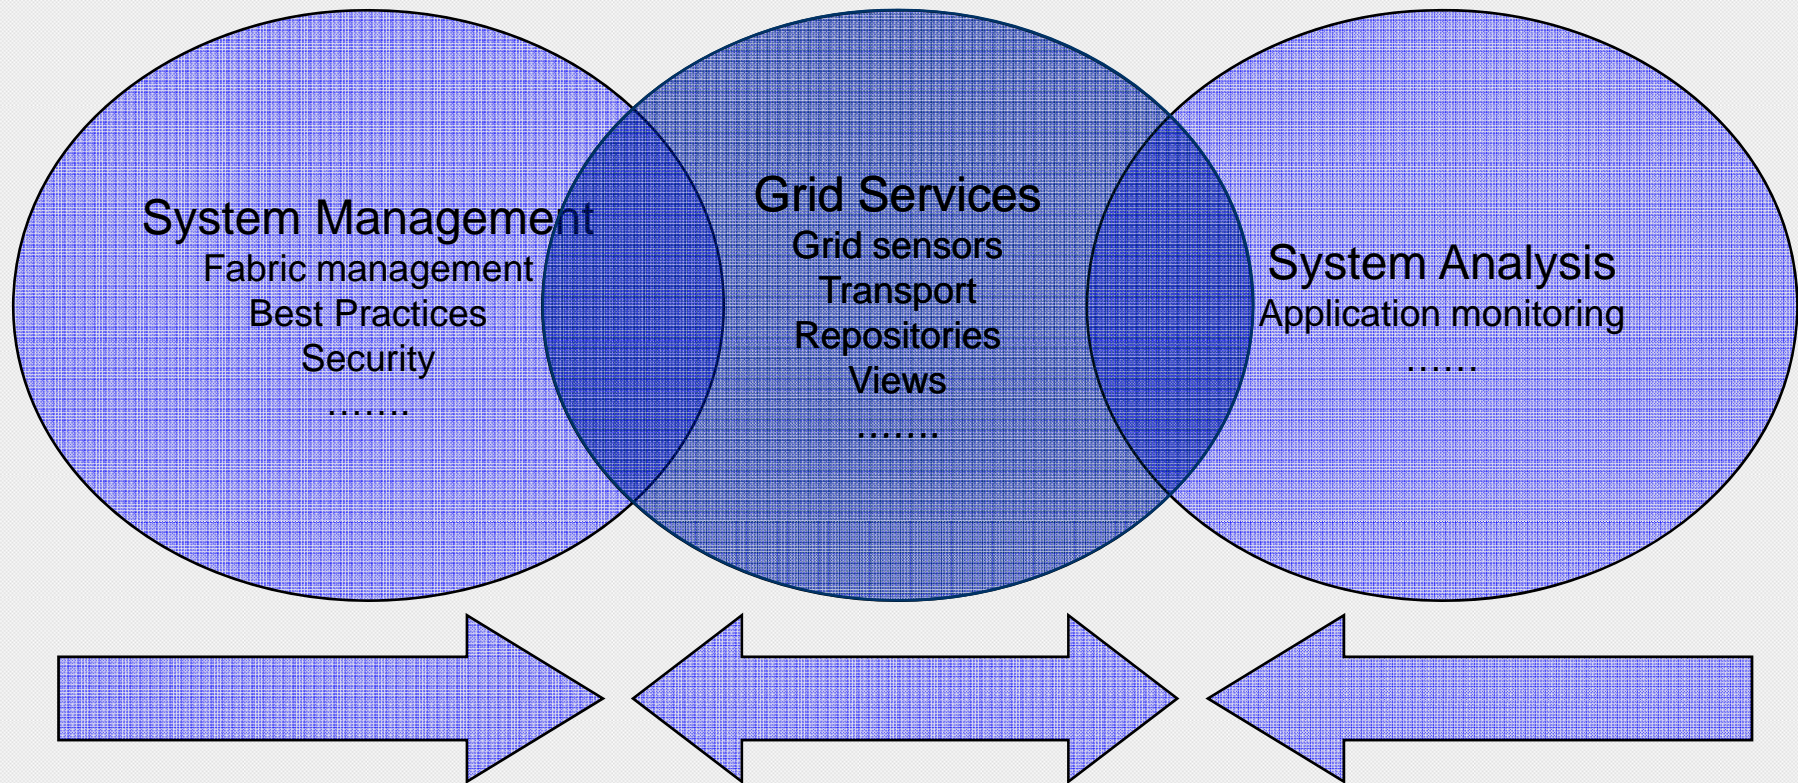
Grid Service Monitoring

James Casey, CERN IT-GD
WLCG/OSG Operations Meeting
14th June 2007



WLCG Monitoring Working Groups

- 3 groups proposed by Ian Bird to the LCG Management Board, Oct 06.
 - Goal to improve the reliability of the WLCG grid





Grid Services Monitoring WG

Mandate

- *"....to help improve the reliability of the grid infrastructure...."*
- *".... provide stakeholders with views of the infrastructure allowing them to understand the current and historical status of the service. ..."*
- *"... stakeholder are site administrators, grid service managers and operations, VOs, Grid Project management"*

<https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringWGMandate>



Monitoring

□ *You can't manage what you don't measure...*

appropriate metrics

- directly relevant to user experience
- clearly defined and understood

accuracy and credibility


measurement instrumentation

- active, passive, collection intervals, alarms

data collection points

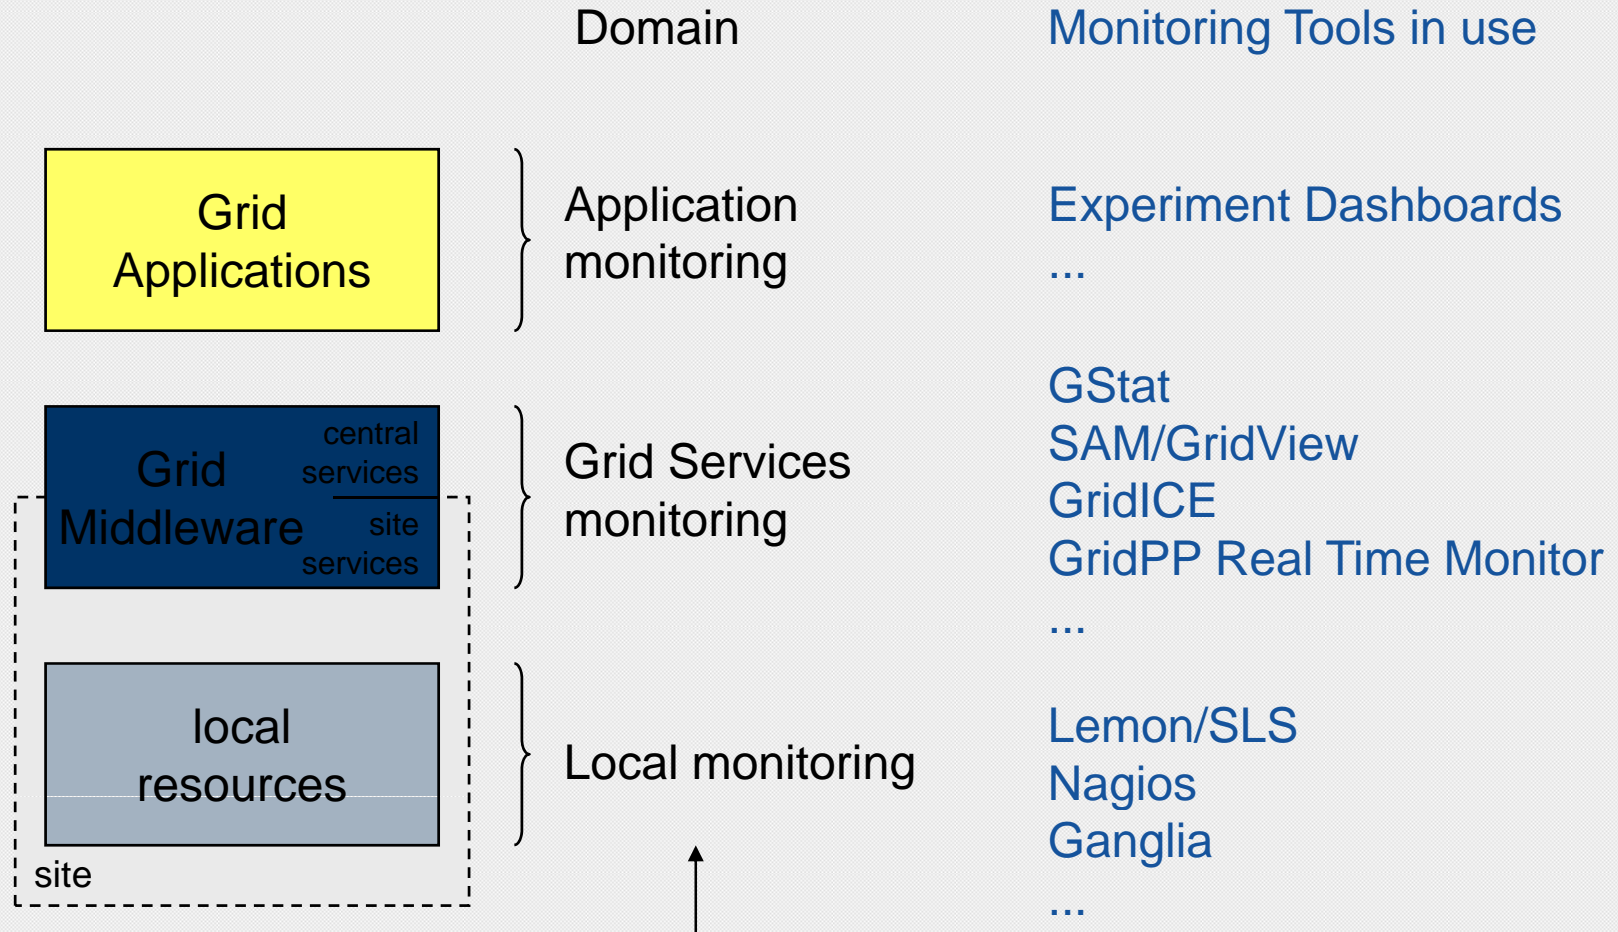
- system element ↔ service

real-time ↔ historical





WLCG Grid Monitoring Landscape



3 WLCG Monitoring Working Groups



Aims of grid services WG

- Not to provide yet another technical solution

But,

- Improve reliability of WLCG
- Consolidate existing solutions
 - Improve communication
 - Reduce overlap
 - Increase sharing



How?

- Engage with stakeholders
 - Operations meetings
 - WLCG Workshops
 - Questionnaires to site managers
 - Grid Service providers (EGEE, OSG)
 - Grid Middleware providers (gLite, VDT)
 - Monitoring software providers (SAM, GridIce, MonAmi, GridView, LEMON, Nagios, ...)
 - External experts (openlab EDS collaboration)
 - Other Working Groups
-



Tasks of grid services WG

- Collect descriptions of current grid services
 - So that probes can be written
 - Input from developers, deployment team, site admins
- Improve quality by providing technical guidance
 - Documenting best practices
 - Providing example components



Tasks of grid services WG

- Best practice notes
 - How to many grid proxies for monitoring
 - Message-level Security for monitoring
 - What information can/should be passed through site boundary
 - ...
- Create set of 'standard' WLCG probes
 - And how to calculate availability based on the metrics produced

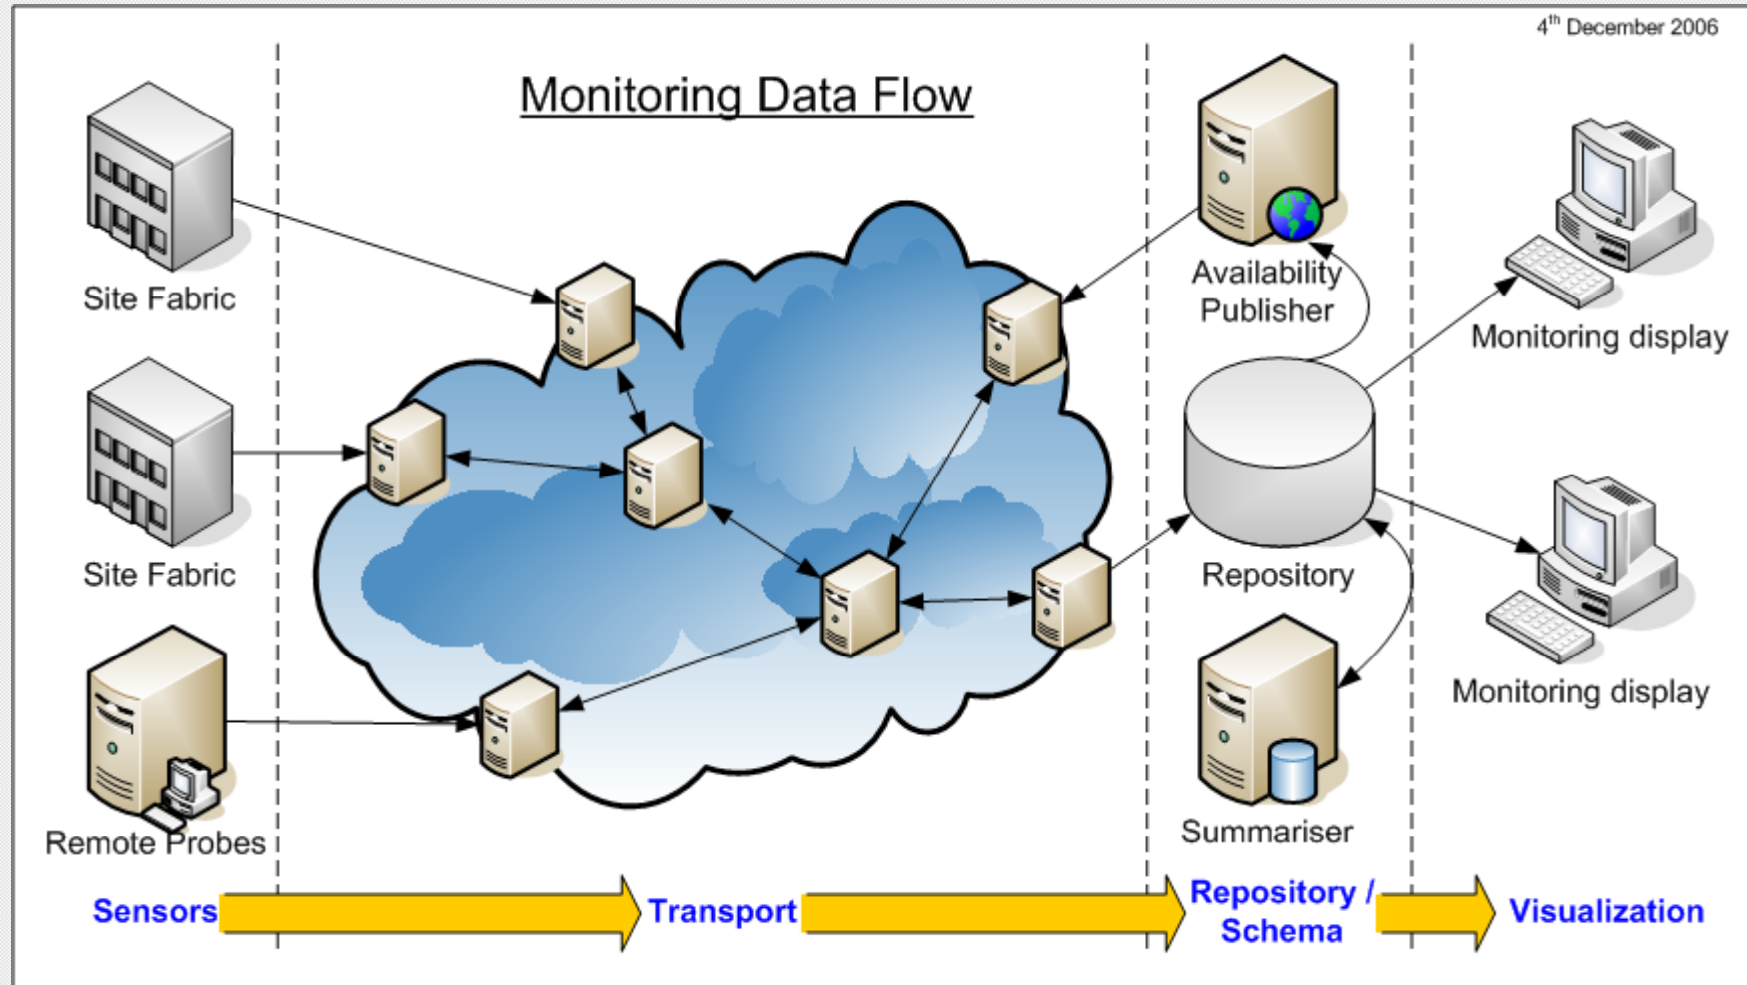


Direction

- Focus on the interaction points between the different systems
 - Allow for diversity across different grid infrastructures
- *“Specifications, not Standards”*
 - Timescales mean we can't get involved in long and heavyweight standards activities
 - Take best practices from existing systems, and document them
- Implement simple prototypes
 - And mature the bits that work !
- Get something out to the stakeholders
 - Close feedback loop is the key to adoption
 - Plan for a “standards based” solution in the future



High-level Model



See https://twiki.cern.ch/twiki/pub/LCG/GridServiceMonitoringInfo/0702-WLCG_Monitoring_for_Managers.pdf for details



Clearing up some terminology

- Metric
 - A data value gathered that tells us something about a service
- Probe
 - The actual code which gathers the metric/metrics
- Check & Sensor
 - A 'probe' in Nagios and LEMON respectively



Locality of Probes

- 'local' can mean two things ; (
- 'local' and 'remote' with respect to probing the interface of the service
 - **local** means on the site
 - **remote** means external to the site
- **(host-)local** probes
 - Gathering information from the operating system level
 - Traditional fabric management probes



3 Categories of Probe Locality

- So we have
 - host-local
 - is daemon running, free-space in size of /tmp
 - local
 - Can I probe the service interface from inside the site
 - remote
 - Can I probe the service interface from outside the site
- All give useful views of the service
 - And I believe all are required!



Service Descriptions

- We needed to understand what to monitor
 - Asked JRA1 developers to fill in a questionnaire
 - Then run past site admins to add their experience
- This should form the basis of the grid fabric monitoring

<https://twiki.cern.ch/twiki/bin/view/LCG/ServiceMonitoringDescriptions>



Node types gathered

gLite 3.1 Node Types

Node Type	Name (in yaim)	Node Description	Done?
gLite WMS and LB	WMSLB	Combined WMS LB node	WMS No WMPProxy L&B
glite CE	gliteCE	gLite Computing Element	Condor + BLAH
FTS	FTS	gLite File Transfer Server	FTS FTS-WS
FTA	FTA	gLite File Transfer Agent	FTS-agent
BDII	BDII	A top level BDII	BDII
Computing Element (middleware only)	CE		
Computing Element (with Torque) *	CE_torque		N/A
LCG File Catalog server *	LFC_mysql	LFC	LFC
R-GMA	MON	RGMA monitoring server	RGMA
e2emonit	E2EMONIT ?	MON box + E2EMONIT ?	N/A
Proxy	PX	Proxy Server	PX
Resource Broker	RB	Resource Broker	
Classic Storage Element	SE_classic	Storage Element on local disk	N/A
Disk Pool Manager (mysql) *	SE_dpm_mysql	DPM Head node	DPM
Disk Pool Manager disk *	SE_dpm_disk	DPM Disk server	as DPM above
dCache Storage Element	SE_dcach	dCache based Storage Element	
Re-locatable distribution *	TAR_UI/TAR_WN	Tarball based Worker Node or a UI	N/A
User Interface	UI	User Interface	
VO agent box	VOBOX	Machine to run VO agents	
Worker Node (middleware only)	WN		



Where should this go?

- This information needs to be curated in the long term
 - EGEE JRA1, SA1 and SA3 involvement is crucial
 - Along with providers of 'externals'
 - Proposal:
 - Simple web based structured repository
 - With database backend
 - Can generate fabric monitoring configuration for a release directly from the information
-



Specifications

Probe Specification

- Defines how a fabric monitoring system can interact with probes that test grid services
- Simple text-based protocol (lightweight)
- Decouples grid probes from the specifics of the fabric monitoring system
- Allows for currently existing probes to be re-used by any monitoring system
 - SAM Tests
 - EGEE CE ROC Nagios testing
 - OSG Tests



Example of Grid Probe

```
$ ./LFC-probe -u lfc://lfc101.cern.ch/ -m  
  ch.cern.LFC-Write -v dteam  
serviceType: glite-LFC  
gatheredAt: lxadm01.cern.ch  
metricStatus: OK  
timestamp: 2007-06-05T15:01:39.86Z  
voName: dteam  
summaryData: OK  
serviceURI: lfc://lfc101.cern.ch/  
metricName: ch.cern.LFC-Write  
EOT
```



Grid Data Exchange Format

- Query interface for repositories to provide stored information to clients
- HTTP message based
 - Query parameters encoded in URL
 - Response is single XML message
- Based on SAM Query format
 - Either current status or history
 - Structured data is returned
 - E.g. all metrics gathered for a site, a VO,
...



Example of exchange format

```
<?xml version="1.0"?>
<root xmlns="http://cern.ch/grid-mon/2007/05/mon-exchange-schema/">
  <Region name="CERN">
    <Site name="CERN-PROD">
      <type>Production</type>
      <status>Certified</status>
      <SiteMetric name="site-daily-avail">
        <measurement>
          <status>ok</status>
          <summary>0.3</summary>
          <timestamp>2007-02-25T00:00:00Z</timestamp>
        </measurement>
      </SiteMetric>
      <Service endpoint="https://ce101.cern.ch:2119/" type="CE">
        <isMonitored>true</isMonitored>
        <inMaintenance>>false</inMaintenance>
        ...
      </Service>
    </Site>
  </Region>
</root>
```

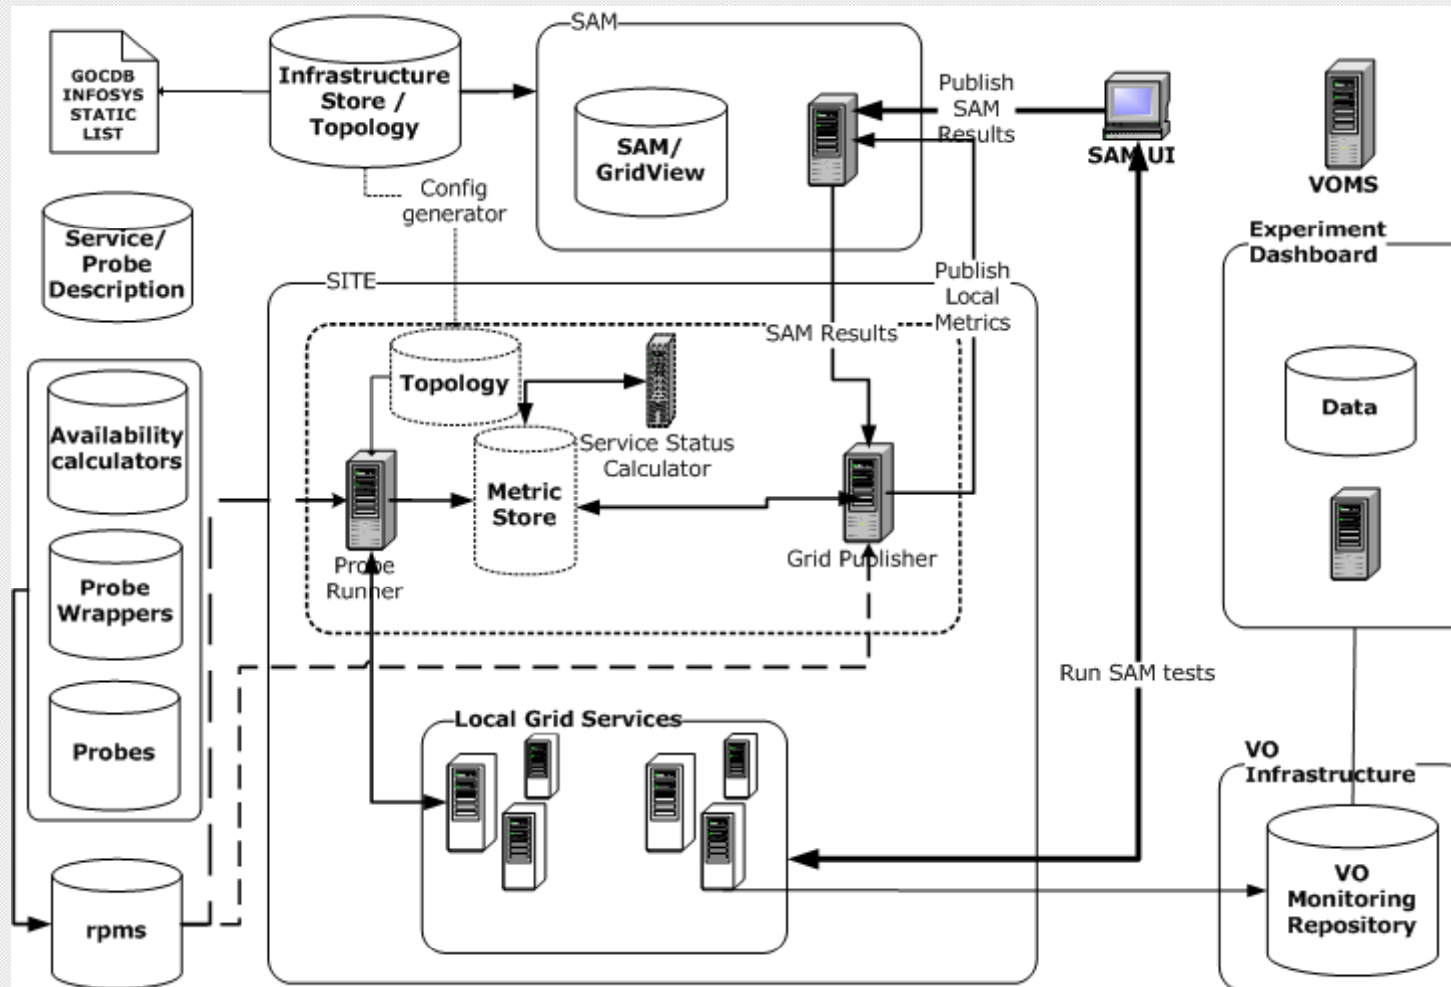


Site monitoring

- We can't/won't impose a solution on sites
 - They might/should have something already
- Specification based approach allows our probes fit into any fabric monitoring system
- Data Exchange format allows higher-level services consume the data regardless of fabric monitoring system



Component view



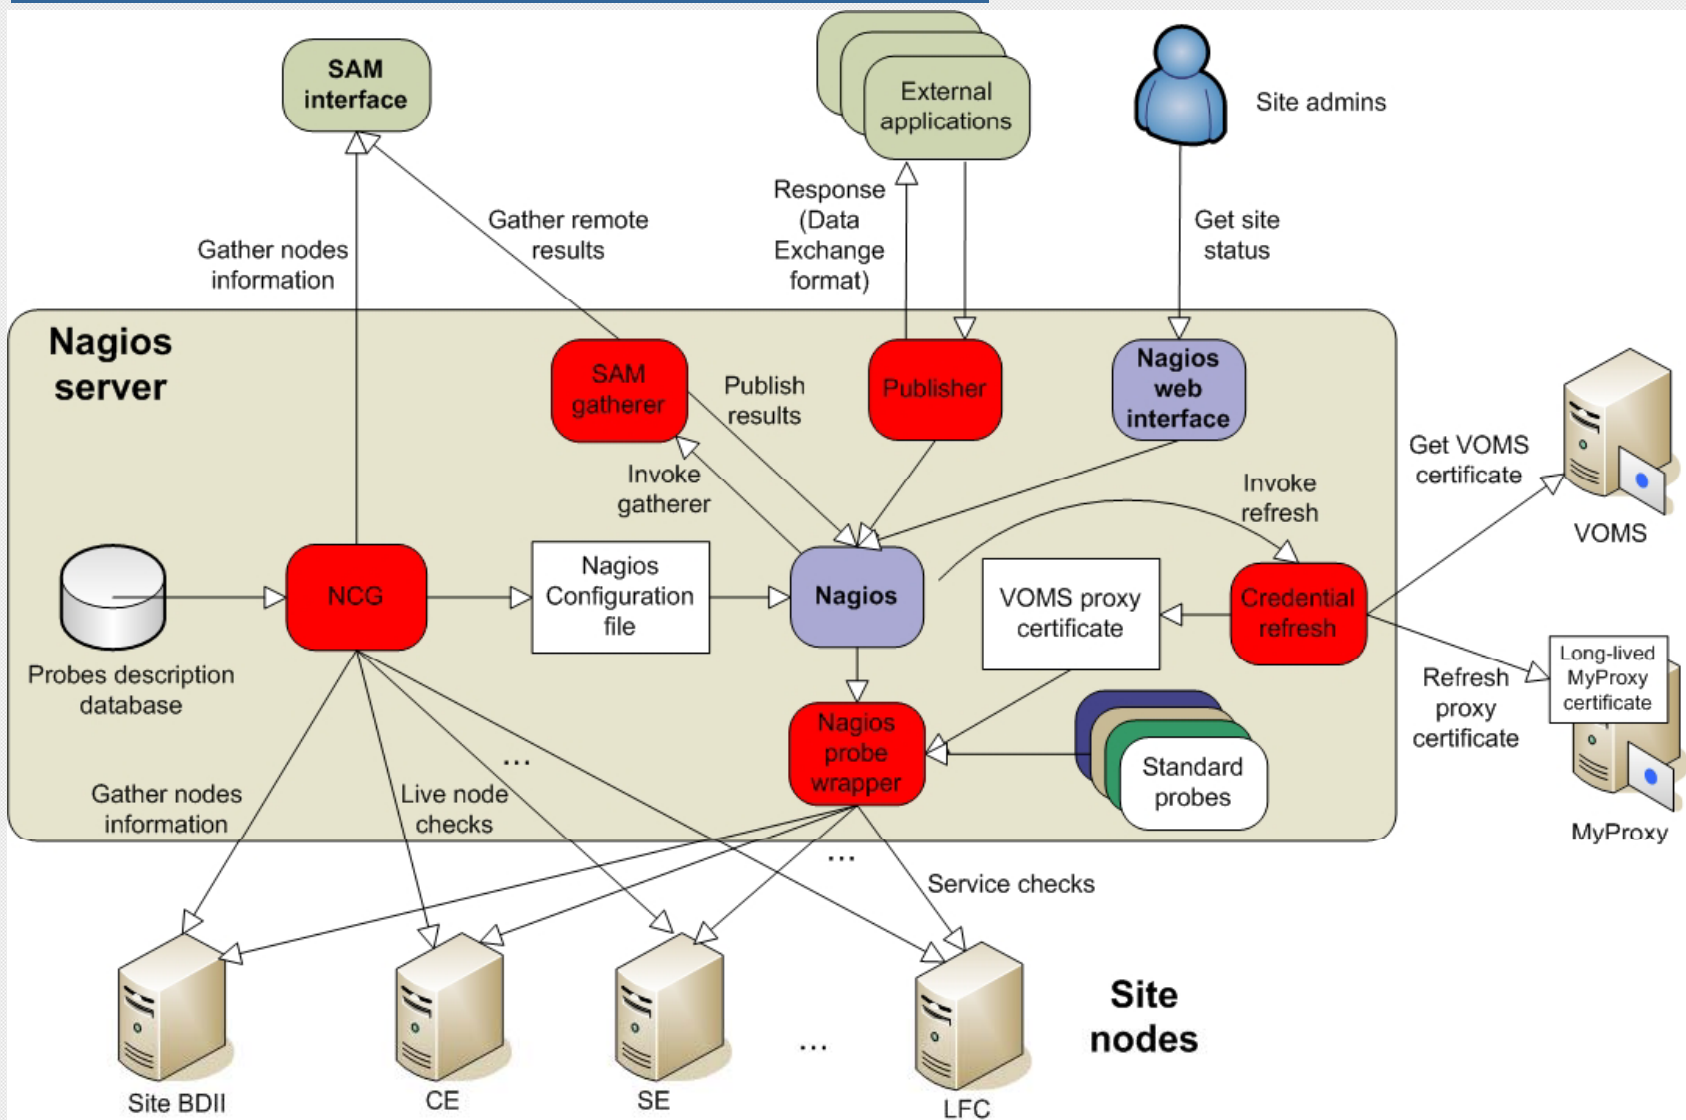


“The Nagios-based prototype”

- How to make this real?
- Initially implement using one fabric monitoring system - Nagios
 - ... but architecture checked with LEMON developers
- Implement some of the site components
 - Configuration Generation
 - Certificate handling
 - Service Status Calculator
- Staged approach for integrating sensors



Prototype site implementation





Configuration generation

- Nagios Specific Configuration Generation Script
 - Generates configuration files from SAM, BDII, live service checks
 - Generate checks for either 'remote' (SAM), or 'local (or 'both')
- Verbose mode that dumps the view of your site
 - Next version will allow you to add/remove services
- Produces a single Nagios .cfg file which can be integrated into an existing Nagios configuration



Stage I

- Feedback SAM info into site monitoring
- Single Nagios sensor – 'gather_sam'
 - Connects to SAM Web Service and gets SAM results for all nodes at the site



Stage II

- Run probes locally as well
 - Allows for local verification of site availability
- Another Nagios sensor – ‘check_wlcfg’
 - Heavily dependant on configuration
 - What probes are needed for what service?
- Start with a probe set consisting of:
 - Sample LFC probe
 - EGEE CE ROC probes (Emir Imamagic)



Stage III

- Run local fabric tests
 - gLite developers have described how to monitor logfiles, daemons etc. on the service nodes
- Integrate this information into the fabric monitoring
 - Where possible use existing sensors



Current status

- Prototype tested against
 - CERN PPS
 - egee.srce.hr site
- Installation and configuration instructions exist
- Packaging done

- Stage I, II are ready or testing



Nagios display

Nagios

General

- Home
- Documentation

Monitoring

- Tactical Overview
- Service Detail**
- Host Detail
- Hostgroup Overview
- Hostgroup Summary
- Hostgroup Grid
- Servicegroup Overview
- Servicegroup Summary
- Servicegroup Grid
- Status Map
- 3-D Status Map

Show Host:

- Service Problems
- Host Problems
- Network Outages

Comments

Downtime

- Process Info
- Performance Info
- Scheduling Queue

Reporting

- Trends
- Availability
- Alert Histogram
- Alert History
- Alert Summary
- Notifications
- Event Log

Configuration

- View Config

Current Network Status
 Last Updated: Wed Jun 6 11:53:57 CEST 2007
 Updated every 90 seconds
 Nagios@ - www.nagios.org
 Logged in as *nagiosadmin*
 - Notifications are disabled

[View History For all hosts](#)
[View Notifications For All Hosts](#)
[View Host Status Detail For All Hosts](#)

Host Status Totals

Up	Down	Unreachable	Pending
25	0	0	1

All Problems	All Types
0	26

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
158	0	137	24	0

All Problems	All Types
161	319

Service Status Details For All Hosts

Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information
castorgrid.cern.ch	GridFTP-Ping	OK	06-06-2007 11:53:33	11d 20h 18m 10s	1/4	FTP OK - 0.039 second response time on port 2811 [220 castorgrid04.cern.ch CASTOR GridFTP Server 1.12 GSS4 Globus/GSI wu-2.6.2(cern-2) (gcc32dbg, 1069715860-42
	GridFTP-Transfer	OK	06-06-2007 11:16:03	0d 0h 37m 54s	1/4	Upload to remote computer succeeded. Download from re computer succeeded. File successfully removed from ren computer. Received file is valid.
	SE-host-cert-valid-OPS-remote	OK	06-06-2007 11:38:08	0d 18h 29m 35s	1/1	SAM status: ok
	SE-lcq-cp-Atlas-remote	OK	06-06-2007 11:03:53	0d 18h 49m 54s	1/1	SAM status: ok
	SE-lcq-cp-CMS-remote	OK	06-06-2007 09:59:00	0d 1h 54m 57s	1/1	SAM status: ok
	SE-lcq-cp-DTeam-remote	OK	06-06-2007 11:47:54	0d 18h 21m 11s	1/1	SAM status: ok
	SE-lcq-cp-OPS-remote	OK	06-06-2007 11:00:03	0d 19h 2m 36s	1/1	SAM status: ok
	SE-lcq-cr-Atlas-remote	OK	06-06-2007 11:03:50	0d 18h 49m 59s	1/1	SAM status: ok
	SE-lcq-cr-CMS-remote	OK	06-06-2007 09:58:48	0d 1h 55m 9s	1/1	SAM status: ok
	SE-lcq-cr-DTeam-remote	OK	06-06-2007 11:47:51	0d 18h 21m 14s	1/1	SAM status: ok
	SE-lcq-cr-OPS-remote	OK	06-06-2007 11:00:00	0d 19h 2m 39s	1/1	SAM status: ok
	SE-lcq-del-Atlas-remote	OK	06-06-2007 11:03:56	0d 18h 49m 51s	1/1	SAM status: ok
	SE-lcq-del-CMS-remote	OK	06-06-2007 09:59:05	0d 1h 54m 52s	1/1	SAM status: ok
	SE-lcq-del-DTeam-remote	OK	06-06-2007 11:47:56	0d 18h 21m 8s	1/1	SAM status: ok
	SE-lcq-del-OPS-remote	OK	06-06-2007 11:00:05	0d 19h 2m 34s	1/1	SAM status: ok
	SE-seavail-OPS-remote	OK	06-06-2007 11:38:13	0d 18h 18m 53s	1/1	SAM status: ok
	SE-seused-OPS-remote	OK	06-06-2007 11:38:13	0d 18h 18m 53s	1/1	SAM status: ok



Prototype delivery timescale

- Stage I – ‘gather_sam’
 - Operations Workshop, mid-June 2007
- Stage II – ‘check_wlcfg’
 - End mid-July 2007
- Stage III – Local probes
 - CHEP, September 2007
- Expect to rapidly iterate, so perhaps only a few “early adopter” sites in June/July
 - Asking for Volunteers



RPM Packages

- `grid-monitoring-fm-nagios`
 - General nagios tools (including certificate handling for running local probes)
- `grid-monitoring-config-gen-nagios`
 - Configuration generator
- `grid-monitoring-probes-ch.cern`
 - Example probe for LFC
- `grid-monitoring-probes-hr.srce`
 - Full probe set for many services from EGEE CE Region



Who should try this?

- Site admins who already use nagios and want to integrate SAM
 - Simply use 'remote' generation
- Site admins who have no monitoring yet and are thinking of trying Nagios
 - Use both 'local' and 'remote' generation
 - Not for the faint hearted - you'll be a very early adoptor!
- RPMs will be available linked from twiki
- Mailing list set up
 - wlcg-monitoring-discuss@cern.ch



Futures and other work

- We focus here on the prototype
 - Since this is what we are delivering now
- Also working on
 - Specifications and example components
 - Security architecture
- Future work includes
 - Probe description database
 - Topology database
 - Messaging architecture for transport layer
- Closely involved with SAM team
 - Looking at how to use Nagios as a submission framework for SAM



Summary

- Effort invested to understand the current landscape
- Approach for improvement based on specifications of interfaces between components
- Prototype has been developed and tested on a small scale
- Now looking for early adopters to get feedback

Who wants to volunteer?

<https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringInfo>



Links to tools

- SAM/GridView Monitoring**
Portal: http://gridview.cern.ch/GRIDVIEW/job_index.php
TWiki: <https://twiki.cern.ch/twiki/bin/view/LCG/GridView>
- SAM (Service Availability Monitor)**
Test Page: <https://lcg-sam.cern.ch:8443/sam/sam.py>
TWiki: <https://twiki.cern.ch/twiki/bin/view/LCG/SamCern>
- GridICE Monitoring**
Portal: <http://gridice2.cnaf.infn.it:50080/gridice/>
Documentation: <http://gridice.forge.cnaf.infn.it/>
- Experiment Dashboard**
Portal: <http://dashboard.cern.ch/>
TWiki: <https://twiki.cern.ch/twiki/bin/view/CMS/Dashboard>
- GridPP Real Time Monitor**
Homepage: <http://gridportal.hep.ph.ic.ac.uk/rtm/> (2D map and 3D globe visualizations)
- GStat**
Portal: <http://goc.grid.sinica.edu.tw/gstat/>
TWiki: <http://goc.grid.sinica.edu.tw/gocwiki/GstatDocumentation>
- Lemon**
Portal (CERN Compute Center): <http://cern.ch/lemon-status/>
Documentation: <http://cern.ch/lemon/>
- Nagios**
Homepage: <http://nagios.org>