



AN AGENT-BASED APPROACH TO GRID SERVICE MONITORING

KEITH ROCHFORD, BRIAN COGHLAN AND JOHN WALSH*

Abstract. The centralised management of distributed computing infrastructures presents a number of considerable challenges, not least of which is the effective monitoring of physical resources and middleware components to provide an accurate operational picture for use by administrative or management staff. The detection and presentation of real-time information pertaining to the performance and availability of computing resources is a difficult yet critical activity.

This architecture is intended to enhance the service monitoring experience of a Grid operations team. We have designed and implemented an extensible agent-based architecture capable of detecting and aggregating status information using low-level sensors, functionality tests and existing information systems. To date it has been successfully deployed across eighteen Grid-Ireland sites.

1. Introduction. Effective monitoring of complex distributed infrastructures, such as electricity supply and traffic management systems, is essential if expected levels of service are to be achieved. Operations Centres are frequently established to satisfy this requirement, providing a central authoritative point of operational information. They are, however, only as good as the underlying information systems they depend on.

Distributed computing infrastructures such as Grids have similar monitoring requirements. For those responsible for the management and administration of the individual resources or the infrastructure as a whole it is desirable that some form of monitoring system is put in place to highlight problems and abnormalities. Centrally managed infrastructures, such as Grid-Ireland[17] or EGEE[3], present additional challenges. Large quantities of status information from disparate resources must be detected, aggregated and stored.

We describe the principles of and justifications for the monitoring system currently deployed on the Grid-Ireland infrastructure and provide a technical overview of the sensors, communication and presentation systems which make the aggregate real-time status information available to the members of the operations team. The monitoring information is presented in such a manner as to facilitate the management of the infrastructure in response to status and security events, and provide a tool to maximise the impact of public relations activities and demonstrations.

2. Grid Monitoring. Grid monitoring may be described as the activity of measuring significant grid related resource parameters in order to analyse usage, availability, behaviour and performance[12]. Its objectives include:

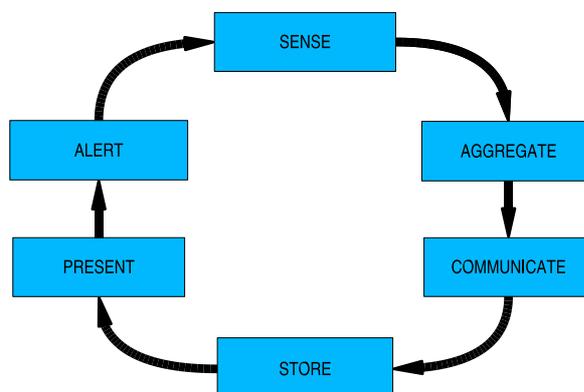
- Locating performance problems
- Tuning for better performance
- Fault detection and recovery
- Input to prediction services

The size and nature of grid systems means that monitoring systems must be capable of measuring a wide range of metrics relating to a substantial number of entities distributed across a large geographical area. Monitoring activities present those responsible for the maintenance of the infrastructures with considerable challenges and has become an active area of research in its own right[22]. Traditional network and host monitoring tools often lack the scalability, dynamicity, or extensibility required for effective deployment within computational or data grids.

3. Context. In order to better understand the motivation for this work we will briefly discuss the architecture of the underlying infrastructure. Grid-Ireland is a managed layer providing grid services above the Irish research network. It allows researchers to share computing and storage resources using a common interface, and facilitates international collaboration by linking Irish sites to into European grid infrastructures being developed under such EU projects as EGEE[3], LCG[9], CrossGrid[2], and int.eu.grid[8].

The Grid-Ireland operations centre is based at Trinity College Dublin and is staffed by members of our research group. Among its responsibilities is the deployment and maintenance of an homogeneous core infrastructure[17] comprised of *grid gateways*. Installed at each site, these gateways provide a point of access both for external grid users to gain access to the site resources and local site users to gain access to remote grid resources.

*Department of Computer Science, Trinity College Dublin, Ireland. keith.rochford, {coghlan, john.walsh}@cs.tcd.ie

FIG. 2.1. *The Monitoring Cycle*

The gateway[15] provides the essential functionality of a grid site, thus freeing site administrators to concentrate on the management of their own resources. The gateways are remotely managed by the operations centre. Local cluster administrators need not concern themselves with the provision of grid services as these are an integral part of the deployed gateways. Resources at a site remain under the control of the site administrators who are free to manage as they see fit, without having to concern themselves with the finer details of grid integration.

Thus the core infrastructure of grid gateways remain the responsibility of the operations team. The ensuing possibility of homogeneity of this core infrastructure presents a number of opportunities to the users, grid operators and the site administrators. These include:

- Minimizing the proportion of software components that need to be ported to non-reference platforms.
- Maximizing the availability of the infrastructure while reducing the effort required for software deployment. The common software and hardware components facilitate centralised management and push-button transactional deployment of middleware components[18], guaranteeing uniform responses to management actions.
- Decoupling of grid infrastructure and site management. The fact that the site resources and the grid infrastructure are independent of each other allows for variation in design, deployment, and management.
- The installation of heterogeneous site resources based on non-reference architectures is also supported. In order to support a wide range of connected resources, the operations team carryout porting of the necessary middleware components[20].

Each grid gateway is composed of a set of seven machines: a firewall, an install server, a compute element (CE), a storage element (SE), a user interface (UI), a test worker node that is used for gateway testing, and optionally a network monitor. In general these are implemented as virtual machines on one physical host. All sites are identically configured with grid software based on gLite3. It is important to recognise that despite being deployed within remote networks, the management and availability of these hosts and services remains the responsibility of the centrally located Grid Operations Team.

4. Centralised vs. Agent-based Monitoring. Remote entities are typically monitored using one of two approaches; a centralised approach where one or more processes executing at a single location are responsible for the monitoring of all entities or a federated approach involving multiple distributed processes and sensors. In a centralised system, all monitoring processes are executed from a single location and attempt to determine the status of monitored entities either through the establishment of some form of connection with the entity or by performing some operation involving it. Distributed monitoring architectures employ monitoring processes on or closer to the monitored entity. These processes, often referred to as agents, perform the required tests and report the outcomes, through push or pull mechanisms, to the interested parties. The area of software agents is one of active research and many types of agents have been identified or proposed. These range from simple processes to complex intelligent autonomous mobile elements of software. It is argued that the popularity of

the term has led to its misuse and it might therefore be appropriate to present an explanation of what should be inferred by the term in the context of this work. We will use the term 'agent' to refer to an independent software component that performs operations on behalf of a user or other piece of software and exhibits some degree of autonomic behaviour.

Both methods of monitoring have advantages and the choice of which is best suited to a given situation is dependent on a number of factors including the size and nature of the systems, the required information and the degree of control that is required. It is often claimed that one approach is superior to the other, however the merits of each must be evaluated with respect to the individual systems and monitoring requirements. While it is true that the deployment of distributed monitoring processes can incur additional effort it is not necessarily true that the configuration need be any more complex. Unfortunately, centralised monitoring typically implies a more limited depth of data gathering than is possible with remote agents. In addition the potential for the management and control of remote entities through interaction with the distributed agents makes them an attractive choice for complex infrastructures. Also the scale of distribution can vary. The requirement for the installation of a software agent on each monitored host is dependent on the motivation of the monitoring activity and the amount of information pertaining to each host that is required. If the motivation for the adoption of an agent-based approach is merely to overcome the limitations imposed by network access policies, it may be sufficient to deploy a single agent to each network segment, domain, or site.

A combination of both approaches might be considered the best solution in terms of grid monitoring. Remote agents provide a level of access and control that would otherwise be unfeasible, yet there are a number of advantages offered by a centralised monitoring approach. It is of limited value to know the internal status information pertaining to a particular resource if we cannot determine whether or not the resource is accessible and usable through the normal operational mechanisms. Furthermore, in the interest of consistency and conceptual simplicity a central monitoring process might be implemented as a monitoring agent deployed at the Operations Centre.

5. Motivations. Early grid service monitoring systems, such as MapCentre[13] and Ganglia[4], offered valuable but limited functionality, required complex configuration, and were employed with limited degrees of success on our particular infrastructure[17]. Monitoring traffic is often denied access into the remote networks hosting the monitored grid resources, greatly reducing the usefulness of these tools. The unsuitability of these systems prompted the investigation and assessment of alternatives.

Several existing host and network monitoring tools, including some supporting distributed monitoring, were assessed but did not offer a suitable solution. Many were intended for deployment within a single administrative domain. While changes to network settings, such as firewall rules, can be quickly authorised and actioned from within such environments, this is not true from outside the domain. From the viewpoint of an operations team the majority of the monitored resources reside within remote administrative domains which are beyond their control. In cases where the resources are remotely managed the remote network still remains beyond control; this makes efficient centralised monitoring of the resources impossible. This is the case for Grid-Ireland, which has a notably integrated remotely managed dedicated grid infrastructure, and in consequence a decision was taken to undertake research into this rather specialist monitoring field.

The monitoring solution described below employs a combination of distributed and centralised monitoring sensors along with the aggregation of information from existing monitoring tools and is specifically designed to speed the response to status and security events.

The immediate aims were to develop and deploy an extensible monitoring framework which would overcome the problems traditionally associated with the monitoring of distributed resources spanning multiple administrative domains to satisfy the information requirements of the operations team. The framework would take advantage of the considerable existing body of work carried out in the area of grid monitoring by aggregating existing systems and tools where provided, and allowing for the development of new sensors where none exist. A key concern was the administrative overhead involved in the deployment and management of the monitoring solution. While there were initially only six sites to monitor, the introduction of an additional twelve sites highlighted the need for an efficient solution.

6. Architecture. The architecture obeys the general principle of ensuring scalability by devolving work to the monitored nodes where possible. This is achieved through the combined use of monitoring agents local to each site, plus a central monitoring process. Traditional systems like Nagios[10] are configured at the site for the site, but this becomes an excessive burden. Here a more efficient integrated approach is taken.

6.1. Monitoring Agents. A single Monitoring Agent, implemented in Java, is deployed to each of the sites with the responsibility of executing service checks for hosts and services at that site. In addition to the distributed agents a central agent running at the Operations Centre determines metrics, such as site connectivity, and aggregates information from additional sources such as Site Functionality Tests[11], Grid Intrusion Detection Systems[21] and existing grid information systems, so that it may be included in the analysis and presentation process.

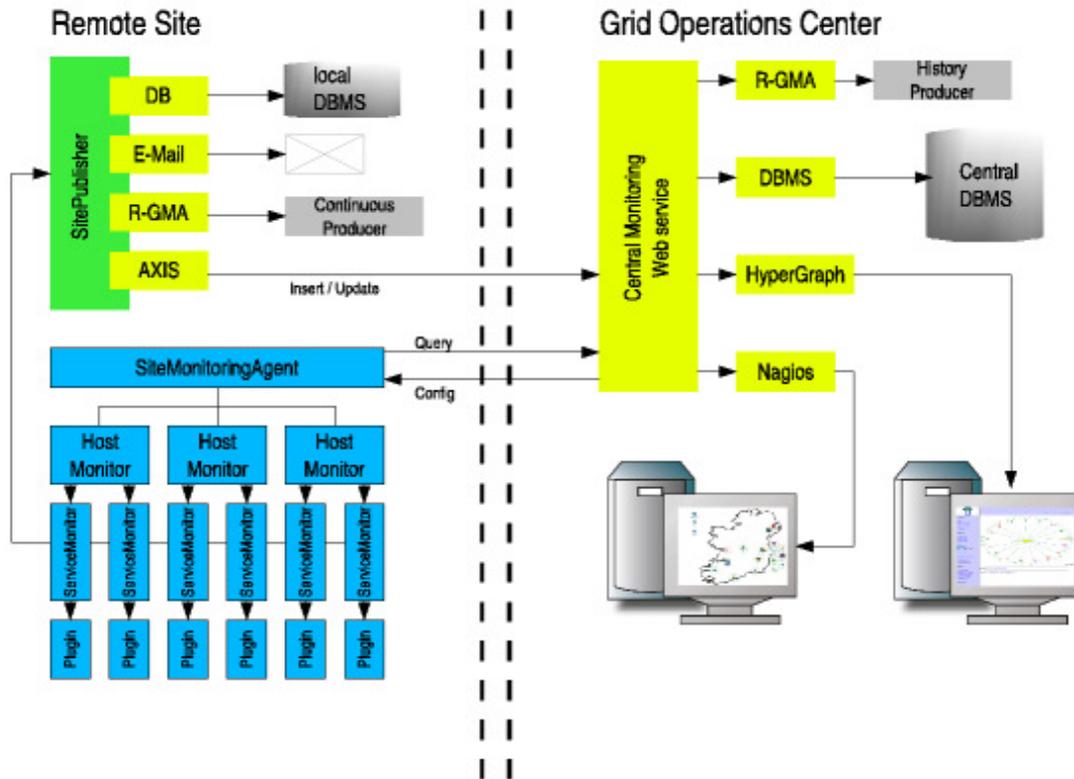


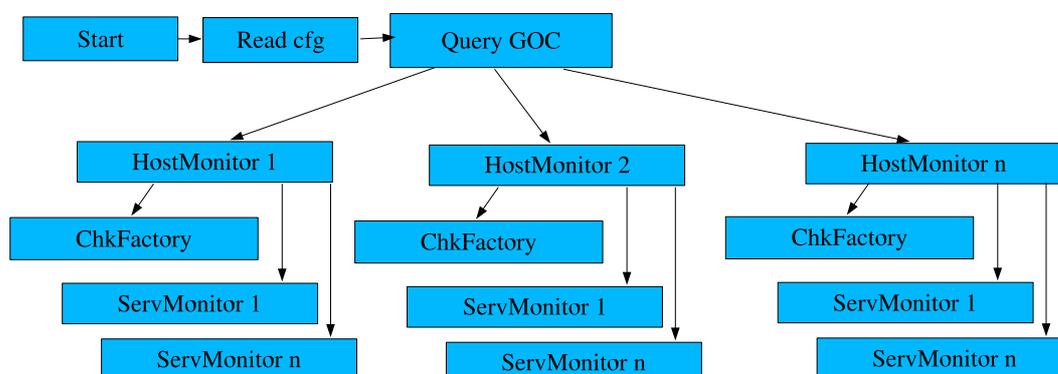
FIG. 6.1. Agent-based monitoring architecture

Configuration information specific to each site, host, or service is stored in a central database and made available to the agents via a web service, thereby minimising the amount of local site-specific configuration required. As described in Figure 6.2, upon start-up a remote agent queries the central configuration server for details of the hosts and services that it should monitor at its site. It then creates local monitoring objects that determine the availability of the specific service through the use of existing or custom service checking objects. At intervals specific to each service, the monitoring objects report their status to interested subscribers via a local publisher. One such subscriber is the central monitoring service, where the results are archived and, if necessary, brought to the attention of an operator. The use of remote sensors communicating over established ports (HTTPS) not only distributes the overhead associated with the monitoring operations but also eliminates many of the connectivity problems commonly experienced with centralised network monitoring.

In order to satisfy the requirement for extensibility, the status of each monitored entity is determined by the agent using plugin modules implemented as Java objects or as Nagios-compliant plugins written in either Perl, C, or as a shell script. This allows the agent to take advantage of a wide range of existing service checks developed for the Nagios project[10]. Examples of such custom plugins include log parsing mechanisms, information system consumers and mechanisms to query a Ganglia[4] monitoring daemon for system load information.

The collected status information includes but is by no means limited to:

- The tcp connectivity to grid service ports such as GRAM, GRIS etc.
- Network reachability of the machines comprising the grid gateway
- Queue / Job activity on the Compute Elements

FIG. 6.2. *Monitoring Agent start-up sequence*

- The output of the Site Functionality Tests
- Grid Portal availability
- SSH reachability of the managed hosts
- Load information obtained via the Ganglia gmon daemon

6.2. Configuration Database. The configuration information defining what hosts and services should be monitored at each of the sites can be represented in text files or database tables, etc. In this case, advantage is taken of the database schema for version 1 of the LCG GOCDB[5] maintained at Rutherford Appleton Laboratory, UK, with several extensions. Hence the configuration information is stored in a MySQL database at the Operations Centre. A sub-set of the configuration database schema is illustrated in Figure 6.3. This shows the data stored for each host, service, and the data required to relate services and hosts for monitoring purposes.

6.3. Central Monitoring Webservice. The central monitoring process, for entirely pragmatic reasons, needs to perform the following functions:

- provision of configuration information to each remote agent
- aggregation of information from remote agents and existing monitoring tools
- archiving and publishing of monitoring/status information

It is implemented as an AXIS[1] web service hosted at the Grid Operations Centre. Configuration information is made available to the monitoring agents via a web service employing the Jakarta Database Connection Pooling mechanisms to query the MySQL database. The use of this pooling mechanism to manage the database connections is an important part of ensuring the scalability of the central Webservice. Since the Web Services are invoked frequently to transfer small amounts of information, establishing a database connection per session would result in slower response times and greater server load.

Using publication mechanisms similar to those embedded within the monitoring agents, status information reported back to the central server is made available to consumer APIs, via re-publisher mechanisms, or passed on to registered listeners. Currently implemented listeners include the Nagios NSCA client, JDBC, XML-RPC, and R-GMA[14] mechanisms. In addition to the central publisher mechanism, the status reports are cached in memory providing rapid access to the latest metrics for all monitored entities. Further optimizations are being explored.

7. Presentation and Alerting. Just as in many traditional Operations Centres, monitoring information is brought to the attention of the operations team in real-time by means of wall-mounted TFT displays, web-based tools and email/messaging alerting systems. Currently four 18 inch LCD display panels are used at the Operations Centre, but soon a number of 40 inch displays will replace these. In the near future, monitoring information and alerts will also be made available to the members of the operations team remotely through the use of mobile devices running web browsers and thin clients.

Nodes Table

Field	Type	Attributes	Null	Default
nodeID	mediumint(8)	UNSIGNED	No	
siteID	mediumint(8)	UNSIGNED	No	0
nodetype	varchar(50)		No	
nodetype2	varchar(50)		Yes	NULL
hostname	varchar(50)		No	
domain	varchar(200)		Yes	NULL
nagios_name	varchar(255)		Yes	NULL
ip	varchar(15)		Yes	NULL
grp	varchar(20)		Yes	LOG-1
hostdn	varchar(255)		Yes	NULL
system	varchar(50)		Yes	NULL
middleware	varchar(50)		Yes	NULL
install	varchar(50)		Yes	NULL
status	varchar(50)		Yes	NULL
operation	varchar(50)		Yes	NULL
comment	varchar(250)		Yes	NULL
monitor	enum('Y', 'N')		No	Y

Services Table

Field	Type	Attributes	Null	Default
serviceID	mediumint(9)		No	
serviceName	varchar(30)		No	
protocol	enum('TCP', 'UDP', 'PING', 'HTTP')		Yes	NULL
port	smallint(6)	UNSIGNED	No	0

Node-Services Mapping

Field	Type	Attributes	Null	Default
hostID	mediumint(9)		No	0
serviceID	mediumint(9)		No	2
checkPeriod	varchar(128)		No	24x7
maxAttempts	int(11)		No	3
checkInterval	int(11)		No	5
retryInterval	int(11)		No	3
contactGroups	varchar(255)		No	

FIG. 6.3. A section of the Configuration Database Schema

7.1. Reports/Displays/Alerts. The Nagios host and network monitoring system is employed at the operations centre for presentation and alerting, solely in order to take advantage of its web based display and reporting functionality in addition to its comprehensive alerting mechanisms. The configuration is purely passive in that Nagios is not responsible for any service monitoring directly but rather the information gathered from the remote sensors is fed into Nagios by means of the Nagios Service Check Acceptor server daemon, allowing the status of defined hosts and services to be updated.

7.2. Navigation of Information. For many existing host and network monitoring tools, navigation of the information is rather neglected. Not only is it desirable that this be easy, it should also be fast, so that an experienced operator can exhibit a “musicians touch”, particularly in emergency situations. The most promising approach thus far assumes tree structures.

HyperGraph[6] is an open source project which provides Java code to work with hyperbolic geometry and in particular hyperbolic trees. Its extensive API facilitates the visualisation of graphs and hyperbolic trees which are extremely useful when dealing with large volumes of data in a hierarchical structure.

The HyperGraph API is used as an example of an alternative methodology for the display of status information, allowing the identification of problems ‘at a glance’ and manipulating the graph for an improved view of the necessary information. Navigation of the infrastructure tree is very fast, which is useful when attempting to find the location of an urgent problem.

In this case (see Figure 7.3), the colour of the edges between the Operations Centre node and the individual site nodes is determined by the network reachability of the gateway machine at that site. The status of the site nodes themselves is determined by the latest results of the Site Functionality Tests for that site. The colour of a host node is determined by the maximum alert value of the services monitored on that host. Nodes representing services are coloured based on the output of the plugin for that service. The graph is constructed in such a way that hovering the pointer over a node causes further information pertaining to that node to be displayed. Table 7.1 defines the colour code used in the graph.

Although very useful, HyperGraph is not the panacea. It’s assumption of tree structures is quite restrictive. Fast hierarchical navigation, particularly relational, is also very desirable and even necessary given our reliance on R-GMA. This is under investigation.

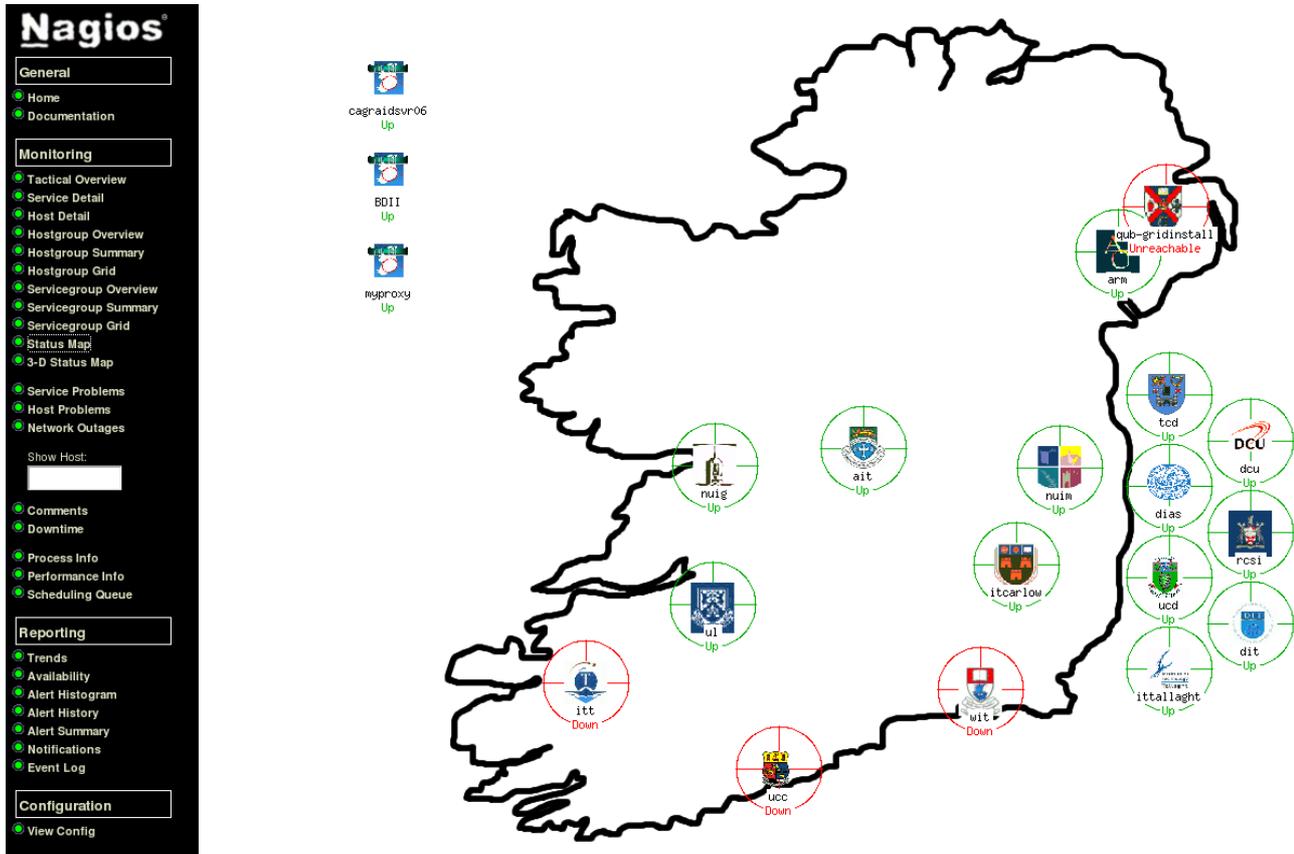


FIG. 7.1. Example Nagios Display

Current Network Status
 Last Updated: Wed Jun 22 11:19:09 IST 2005
 Updated every 90 seconds
 Nagios® - www.nagios.org
 Logged in as guest

[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
4	0	0	0
All Problems		All Types	
0		4	

Service Status Totals				
Ok	Warning	Unknown	Critical	Pending
11	0	0	0	0
All Problems		All Types		
0		11		

Service Status Details For Host Group 'tcd'

Host	Service	Status	Last Check	Duration	Attempt	Status Information
tcd-gridgate	GFTP	OK	06-08-2005 21:20:33	40d 18h 39m 26s	1/3	TCP OK - 0.459 second response time on port 2811
	GRAM	OK	06-08-2005 21:20:49	40d 18h 39m 26s	1/3	TCP OK - 0.387 second response time on port 2119
	GRIS	OK	06-08-2005 21:20:58	40d 18h 39m 26s	1/3	TCP OK - 0.345 second response time on port 2135
	SSH	OK	06-22-2005 11:15:38	16d 7h 7m 10s	1/3	SSH OK - OpenSSH_3.6.1p2-CERN20030917 (protocol 1.99)
tcd-gridinstall	APC	OK	06-22-2005 11:16:27	0d 7h 38m 9s	1/3	IN 228.9v (208.0 - 253.0) OUT 228.9v 50.0 Hz, Load 38.3 Percent, Charge 100.0 Percent, Batt 27.4v, Time left 33.0 mins, Internal Temp 31.5 C
	HTTP	OK	06-22-2005 11:17:09	49d 20h 40m 43s	1/3	HTTP ok: HTTP/1.1 200 OK - 0.003 second response time
	SSH	OK	06-22-2005 11:15:38	22d 16h 30m 58s	1/3	SSH OK - OpenSSH_3.1p1 (protocol 1.99)
tcd-gridstore	GFTP	OK	06-08-2005 21:22:57	49d 20h 41m 0s	1/3	TCP OK - 0.236 second response time on port 2811
	SSH	OK	06-22-2005 11:15:38	54d 17h 29m 15s	1/3	SSH OK - OpenSSH_3.6.1p2-CERN20030917 (protocol 1.99)
tcd-gridui	HTTPS	OK	06-22-2005 11:17:37	0d 21h 6m 42s	1/3	HTTP ok: HTTP/1.1 200 OK - 0.071 second response time
	SSH	OK	06-22-2005 11:15:38	54d 17h 31m 56s	1/3	SSH OK - OpenSSH_3.6.1p2-CERN20030917 (protocol 1.99)

FIG. 7.2. Detailed service monitoring information

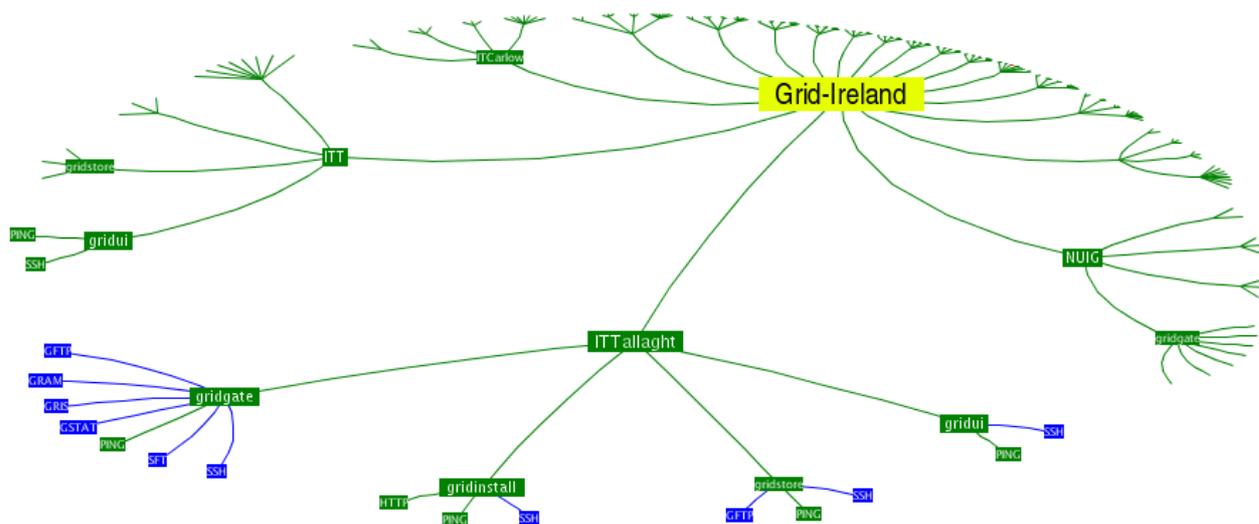


FIG. 7.3. A snapshot of the Hypergraph Display of Infrastructure Status in mid-navigation

TABLE 7.1
HyperGraph Infrastructure Monitoring Key

Colour	Connotation
Green	No problems or warnings
Amber	Warnings exist for this host or service
Red	Errors or critical warnings exist
Blue	The information for this host is considered stale
Grey	The status could not be determined

7.3. Alert Analysers. Of course, alerts are themselves an interesting data set, amenable to extraction of useful information regarding stability, behaviour, etc. More in-depth analysis and intelligent alerting is made possible through the publication and aggregation of the monitoring information into the R-GMA system [19] and the use of R-GMA consumers in the form of custom alert analysers [16]. These analysers can trigger event handlers and notification mechanisms based on queries made on the R-GMA producers.

It is important to recognise that the power of the SQL queries within individual analysers may be extended by building trees of or a series of dependent analysers, allowing complex analysis of the monitoring information. Analysers may contain advanced processing logic and include historical analysis, providing alert escalation mechanisms or the investigation of correlations among service outages and higher level alerts or security events etc.

Examples of event handlers and notification mechanisms currently under development are outlined in table 7.2.

TABLE 7.2
Example Analyser event handlers

Handler Name	Description
Console Handler	Prints the alert to an open console
JDBC Handler	invokes an insert statement on a JDBC resource
RGMA Handler	invokes an insert statement on an RGMA producer
MAILTO Handler	Sends an email to a user/operator detailing alert
SMS Handler	Sends a Short Message Service (text) to a user/operator
XML-RPC Handler	Invokes a method on an XML-RPC server

8. Deployment. The system described has been deployed across 18 Grid-Ireland sites with the task of monitoring almost 200 services on over 70 hosts. Up to 15,000 service/host check results are reported to the central monitoring service each day. It has achieved its objectives and proves to be a valuable tool for the operations team. Its extensible nature is tested regularly as additional checks are requested.

9. Future Work. While the current system is production-ready and has a stable history, a number of future avenues of research are planned. These include further extension to the agent architecture, improved agent and resource management, on-demand monitoring, and the aggregation of additional information sources, particularly with respect to security information [21].

The area of real-time and historical status information analysis, with relevant warning and alerting systems, is where the majority of our future research lies. Customised presentation systems employing 3D visualisation and virtual instrumentation will also be investigated.

A Sourceforge project has been setup [7] to promote the deployment and further development of the architecture. It is hoped that it will attract interest of other users and developers, particularly those interested in tailoring the solution to their own specific infrastructure requirements.

10. Conclusions. This paper has outlined some of the difficulties associated with the effective monitoring of distributed computing infrastructures and has presented our efforts in the development of a new approach to the execution and management of grid service monitoring. Following a brief discussion of distributed versus centralised monitoring, we described some of the difficulties encountered in the deployment of monitoring tools within our own infrastructure and introduced our motivations for an agent-based solution. It has demonstrated the combined use of remote and centralised monitoring mechanisms along with the aggregation of existing information systems in order to satisfy the information requirements of a Grid Operations Centre. A number of presentation and alerting systems currently in use were also described.

Our distributed monitoring solution boasts a number of advantages. The lightweight Java agents are easily deployed at remote sites requiring minimal local configuration. Configuration management is achieved through web interfaces to the configuration database. The use of standard communication protocols has resulted in a reliable system, capable of operating within the constraints of tightly managed network security due to its use of standard ports. Site firewall rules for these ports are not typically subject to change following security audits. Where appropriate, existing tools have been used in a flexible and extensible manner, improving the efficiency of the development effort and the overall usefulness of the system. The archiving and republishing of the monitoring information makes this a valuable component on which to base future work.

REFERENCES

- [1] Axis - Apache WebServices Project. <http://ws.apache.org/axis/>
- [2] CrossGrid. <http://www.crossgrid.org/>
- [3] Enabling Grids for E-science (EGEE). <http://www.eu-egee.org/>
- [4] Ganglia project documentation. <http://ganglia.sourceforge.net/docs/>
- [5] Grid Operations Centre Database. <http://goc.grid-support.ac.uk/gridsite/gocdb/>
- [6] Hypergraph project documentation. <http://hypergraph.sourceforge.net/docs.html>
- [7] I4c - Wide Area Network Monitoring. <http://i4c.sourceforge.net/>
- [8] Interactive European Grid Project. <http://www.interactive-grid.eu/>
- [9] Large Hadron Collider Computing Grid project (LCG). <http://lcg.web.cern.ch/LCG/>
- [10] Nagios project documentation. <http://www.nagios.org/docs/>
- [11] Site Functionality Tests. http://goc.grid.sinica.edu.tw/gocwiki/Site_Functional_Tests
- [12] S. ANDREOZZI, N. DE BORTOLI, S. FANTINEL, A. GHISELLI, G. TORTONE, AND C. VISTOLI. Gridice: a monitoring service for the grid. Proc. Cracow Grid Workshop, Poland, December December, 2003.
- [13] P. P. BONNASSIEUX F., HAKAKALY R. Mapcenter: an open grid status visualization tool. In *ISCA 15th International Conference on parallel and distributed computing systems*, Louisville, Kentucky, USA, September September 19-21, 2002.
- [14] R. BYROM, B. COGHLAN, A. COOKE, R. CORDENONSI, L. CORNWALL, A. DATTA, A. DJAOU, L. FIELD, S. FISHER, S. HICKS, S. KENNY, J. MAGOWAN, W. NUTT, D. O'CALLAGHAN, M. OEVER, N. PODHORSZKI, J. RYAN, M. SONI, P. TAYLOR, A. WILSON, AND X. ZHU. R-gma: A relational grid information and monitoring system. Proc. Cracow Grid Workshop, Poland, December December, 2002.
- [15] S. CHILDS, B. COGHLAN, D. O'CALLAGHAN, G. QUIGLEY, AND J. WALSH. A single-computer grid gateway using virtual machines. In *Proc. AINA 2005*, pages 761-770, Taiwan, March 2005. IEEE Computer Society.
- [16] B. COGHLAN AND S. KENNY. Grid-wide intrusion detection: First results after deployment. *Submitted to special issue of the Journal of Parallel and Distributed Computing on Security in Grid and Distributed Systems*.
- [17] B. COGHLAN, J. WALSH, AND D. O'CALLAGHAN. Grid-ireland deployment architecture. In P. M. Sloot, A. G. Hoekstra,

- T. Priol, A. Reinefeld, and M. Bubak, editors, *Advances in Grid Computing - EGC 2005*, LNCS3470, Amsterdam, The Netherlands, February February, 2005. Springer.
- [18] B. COGHLAN, J. WALSH, G. QUIGLEY, D. O'CALLAGHAN, S. CHILDS, , AND E. KENNY. Transactional grid deployment. In M. Bubak, M. Turala, and K. Wiatr, editors, *Proc. Cracow Grid Workshop (CGW'04)*, pages 363–370, Cracow, Poland, December 2004. Academic Computer Centre CYFRONET AGH.
- [19] A. COOKE, A. GRAY, L. MA, W. NUTT, J. MAGOWAN, M. OEVERS, P. TAYLOR, R. BYROM, L. FIELD, S. HICKS, J. LEAKE, M. SONI, A. WILSON, R. CORDENONSI, L. CORNWALL, A. DJAOUI, S. FISHER, N. PODHORSZKI, B. COGHLAN, S. KENNY, AND D. O'CALLAGHAN. R-gma: An information integration system for grid monitoring. In *Proc.Int.Conf. Cooperative Information Systems (CoopIS'03)*, Catania, Sicily, November November, 2003.
- [20] E. KENNY, B. COGHLAN, J. WALSH, S. CHILDS, D. O'CALLAGHAN, AND G. QUIGLEY. Autobuilding multiple ports of computing nodes for grid computing. In *Cracow Grid Workshop (CGW'05)*, Cracow, Poland, November 2005.
- [21] S. KENNY AND B. COGHLAN. Towards a grid-wide intrusion detection system. In P. M. Sloot, A. G. Hoekstra, T. Priol, A. Reinefeld, and M. Bubak, editors, *Advances in Grid Computing - EGC 2005*, LNCS3470, Amsterdam, The Netherlands, February February, 2005. Springer.
- [22] S. ZANIKOLAS AND R. SAKELLARIOU. A taxonomy of grid monitoring systems. *Future Gener. Comput. Syst.*, 21(1):163–188, 2005.

Edited by: Dana Petcu

Received: May 31, 2007

Accepted: June 18, 2007