

Publication and Protection of Sensitive Site Information in a Grid Infrastructure

Shreyas Cholia

Lawrence Berkeley National Laboratory
scholia@lbl.gov

R. Jefferson Porter

Lawrence Berkeley National Laboratory
rjporter@lbl.gov

Abstract

In order to create a successful grid infrastructure, sites and resource providers must be able to publish information about their underlying resources and services. This information makes it easier for users and virtual organizations to make intelligent decisions about resource selection and scheduling, and can be used by the grid infrastructure for accounting and troubleshooting services. However, such an outbound stream may include data deemed sensitive by a resource-providing site, exposing potential security vulnerabilities or private user information to the world at large, including malicious entities. This study analyzes the various vectors of information being published from sites to grid infrastructures. In particular, it examines the data being published to, and collected by the Open Science Grid, including resource selection, monitoring, accounting, troubleshooting, logging and site verification data. We analyze the risks and potential threat models posed by the publication and collection of such data. We also offer some recommendations and best practices for sites and grid infrastructures to manage and protect sensitive data.

1. Introduction

Grid computing has become a very successful model for scientific collaborations and projects to leverage distributed compute and data resources. It has also offered the research and academic institutions that host these resources an effective means to reach a much larger community. As grid computing grows in scope, and as an increasing number of users and resources are plugged into the grid, there is an increasing need for metadata services that can provide useful information about the activities on that grid. These services allow for more sophisticated models of computing, and are fundamental components of scalable grid infrastructures. The scope of these services is fairly

broad and covers a variety of uses including resource selection, monitoring, accounting, troubleshooting, logging, site availability and site validation. This list could grow, as grids evolve and other types of metadata become interesting to users and administrators. This means that it becomes important for a grid infrastructure to provide central collection and distribution points that can collate information gathered from multiple sources.

The typical publication model involves pushing data from site based informational end points to central collectors, using streaming feeds or periodic send operations. The central collectors then make this data available to interested parties using standard interfaces and protocols in the form of web services and database query engines. The usability of the grid depends on the widespread availability of this information. Given the increasingly open nature of grid computing these collectors and information services generally present publicly accessible front-ends.

Now consider the implications of this model for a site providing grid resources. Being included in a grid infrastructure means that a large amount of site information suddenly enters the public domain. This could include information deemed as sensitive or private from the perspective of the site, the user or the grid collaboration as a whole. It becomes very important then, to have controls on the access and flow of this data, so that the information sources can decide what data they want published and what data they want restricted. Since these models of informational flow are still evolving in today's grids, these controls are still in the process of being designed into the software infrastructure. As such, there isn't a standard way to control this flow of information. We think there is an urgent need to study the various vectors of information being provided by sites to grid infrastructures. This includes an analysis of the nature of the information

itself, as well as the software publishing this information.

In our work, we use the Open Science Grid (OSG) as a case study for this model of information flow, looking at the five major information collection systems within the OSG, and analyzing the security implications of this infrastructure. We also provide some recommendations on improving the current infrastructure to preserve the privacy and security of sensitive information.

2. The Open Science Grid

The OSG offers a shared infrastructure of distributed computing and storage resources, independently owned and managed by its members. OSG members provide a virtual facility available to individual research communities, who can add services according to their scientists' needs.

It includes a wide selection of resource providers, ranging from small universities to large national laboratories. This broad range of sites results in a diverse set of security requirements. Reconciling these diverse security priorities is a challenge, and requires close interaction between the sites and the OSG managers. One approach to addressing this issue is to provide the necessary tools in the grid middleware stack, so that sites can configure security policies directly into the software. The OSG provides a software distribution called the Virtual Data Toolkit (VDT). This includes a packaged, tested and supported collection of middleware for participating compute and storage nodes, as well as a client package for end-user researchers.

The OSG also provides support and infrastructure services to collect and publish information from participating sites, and to monitor their resources. These services are provided by the OSG Grid Operations Center (GOC). The GOC provides a single point of operational support for the OSG. The GOC performs real time grid monitoring and problem tracking, offers support to users, developers and systems administrators, maintains grid services, and provides security incident responses. It manages information repositories for Virtual Organizations (VOs) and grid resources.

3. Information Collection in OSG

There are currently five major information collection systems in the OSG, which rely on information feeds from sites to centralized servers. The following is a description of each of these services, and

an analysis of the information being published by them from a site security perspective.

3.1 Resource Selection Information

In the OSG framework, the Generic Information Provider (GIP) gathers site resource information. GIP aggregates static and dynamic resource information for use with LDAP-based information systems. Information published is based on Glue Schema. The CEMon (Compute Element Monitor) service is responsible for publishing this information to a central OSG information collector service called the CEMon Consumer. CEMon connections are authenticated and encrypted (using GSI). This information is then made public in two ways:

1. Class-ads are published to a Condor matchmaker service called the Resource Selection Service (ReSS), which allows Condor clients to select appropriate resources for job submission.
2. The Berkeley Database Information Index (BDII) collects this information for resource brokering. It tracks status of each participating cluster in terms of available CPUs, free CPUs, supported VOs, etc.

The Glue Schema provides a more detailed list of attributes supported in this scheme. For the purposes of this study, we concentrate on those attributes published by GIP that may be deemed sensitive by certain sites. This includes:

- Operating System version/patch information
- Authentication method (grid-mapfile, GUMS)
- Underlying job-manager and batch system information
- Internal system paths

In some sense, publication of this information is essential to a site's successful participation in the grid. However, a site must understand the implications of making this information public. Prior to joining the grid, much of this information was inherently under the control of the site, and limited to people under its own administrative domain. As such, administrators must be aware of any conflicts with the current site security policy and requirements that may have been drafted prior to participation in the grid.

Additionally, a site may only want to provide this information up to a desired level of detail. Since the GIP software will publish all available information in its default mode, a site may want to consider limiting, or overriding some of the attributes being published.

Another consideration is the public nature of this information, once it has been sent to the CEMon Consumers. Given that this information is only useful to actual users of the grid, it might be useful to provide

some minimal restrictions so that the information is only accessible to current members of the OSG (or collaborating grids).

3.2 Accounting

The Gratia software provides the accounting framework for the OSG. Gratia consists of two components:

1. The Gratia probes that run on the site resource and interface with the site-specific accounting and batch systems. These probes extract resource usage information from the underlying infrastructure and convert it into a common Usage Record-XML based format. This is then sent to a central collector.
2. The Gratia collector is a central server operated by the OSG GOC that gathers information from the various probes, and internally stores this in a relational database. It makes this information publicly available through a web interface, in certain pre-defined views. The web interface also allows viewers to create their own reports and custom SQL queries against the usage data.

The Gratia records include information that might be considered sensitive by both the sites and the grid users. Specifically, we identified the following information as potentially sensitive:

- User account names
- User DN information
- Job file and application binary names

Given that this information can be accessed through a public SQL interface, all user activity on the OSG can be traced and analyzed in fairly sophisticated ways, by anyone with a web browser.

User account and DN information could be used by an attacker that has compromised an account on one site to query a list of sites with the same user account/DN, thus increasing the scope of the attack. It is not being suggested that masking this information will protect a site from a compromised account on another system. Certainly, once an account has been compromised, any other site that uses a common set of login credentials should be considered vulnerable. However, making this information less accessible to an attacker could mitigate the scope of the attack.

Job file or application names would be less useful to attackers, but could reveal information about the nature of the jobs being run. There is the potential for a rival project to gain valuable clues about the research being done from this information. A researcher may want to restrict this information to a limited set of people. On the other hand, from an accounting standpoint, the

underlying file descriptions may not be as interesting as the actual resource consumption being measured. In most cases, the accounting software only needs to be able to uniquely identify a job, and doesn't care about the specifics of underlying job or application names.

For these reasons, it is recommended that access to this data be restricted along user and VO lines using grid certificates as the mechanism for controlling this. Sites can also mask sensitive information by modifying the probe software to apply filters to the records.

3.3 Logging

The OSG uses Syslog-ng to provide centralized logging of user activity on the Grid. Syslog-ng is an extension to the Syslog protocol that provides more flexible support for distributed logging and richer content filtering options.

Currently OSG resources optionally log all information related to Grid processes using syslog-ng, and send this to a central collector managed by the GOC. The primary uses for this information are:

1. Troubleshooting – Being able to trace the workflow of a distributed job is very useful as a debugging tool for failures. It makes it significantly easier to detect how and why a job might be failing, especially when multiple sites are involved. The OSG GOC has a troubleshooting team to deal with such cases.
2. Security Incident Response – Having centralized logs available to the OSG security team, makes it very useful to be able to analyze the scope and extent of a security compromise. It allows the GOC to identify compromised sites or users, and to judge the nature of the compromise. Affected sites can then be notified for rapid incident response.

In the troubleshooting case, there is the need to protect failure modes from becoming publicly available, as this could reveal possible avenues for attack. For example, a poorly configured site may have vulnerabilities in the execution path. While not apparent through the standard client software, these may be exposed through syslog information. In general, logging information should only be available to authorized personnel within the OSG administrative domain, or to specific users when debugging problems. Another approach to this issue involves the level of logging performed by the site, so that only a minimal amount of information is logged by default. This translates to logging only the start and stop times for jobs and data transfers for a given user. In the event of a failure, the site can increase the level of logging, and

work in conjunction with the troubleshooting team and the user to diagnose the specific problem.

Security incident information is perhaps even more sensitive, and syslog information revealing incident details must have tight access controls. Once again, this points to restricting the information to an authorized set of security personnel.

Syslog-ng allows for collectors on a per site basis, that can then filter out the information getting passed to the OSG wide collector. This would allow sites to collect detailed information internally, while filtering the information sent to the OSG. Any information sent to the OSG GOC should be encrypted. As long as there is enough information being sent to identify a failure or compromise at a central level, the relevant sites can be notified of this. The sites can then address the specifics of the problem, and provide more information to the OSG GOC and security team, as necessary. This is the model that is expected to go into production for future OSG deployments.

3.4 Site Availability and Validation Data

The OSG GOC performs site availability and validity tests on participating compute and storage elements, and publishes these results online. These tests are run at regular intervals, either using a Perl script (*site_verify.pl*) or using a customizable set of probes called RSV (Resource and Service Validation). The basic aim is to validate the services being advertised through the resource selection and monitoring modules (CEMon). Much of the information being collected here is analogous to CEMon information, and subject to the same issues. The RSV probes use a push model, similar to the Gratia service. The *site_verify.pl* script takes the form of a remote grid job run by the GOC at individual sites, relaying information back using the standard Globus data movement protocols (GASS, GridFTP). Possibly sensitive information being reported includes:

- Account Names
- Historical system availability information
- Currently running software information
- Internal System Paths

Given that site validation data is both being collected at regular intervals, and being archived, it offers the ability to track the state of a system over time. This may provide information about regular system downtimes, when a system may be in a transitional state and particularly susceptible to an attack.

Moreover, the archived nature of this information suggests that the site is subject to a “Google Hack”,

even if system data is no longer been published. An attacker can use standard search-engine technology to scan the Internet for systems that match certain keywords. This can be used to scope out systems with known vulnerabilities based on advertised software levels. This is compounded by the fact that modern search-engines like Google do their own external caching and archiving of information, creating a situation where anything that is published on the web has the chance of persisting, despite a site no longer wishing to make that information publicly available. There are known methods to prevent a site from being listed in a search engine, and it is recommended to use these for this kind of data.

3.5 Monitoring

The OSG uses the CEMon software for monitoring sites. An analysis of this has already been included in the “Resource Selection Information” section.

The OSG also supports an optional package called MonALISA (MONitoring Agents using a Large Integrated Services Architecture) to monitor system availability and load. Sites using MonALISA send system information to a central MonALISA service, which allows general users to query site information from a web-based clickable map interface. It monitors the following information:

- System information for computer nodes and clusters.
- Network information (traffic, flows, connectivity, topology) for WAN and LAN.
- Performance of applications, jobs and services.
- End user systems, and end-to-end performance measurements.

Since this includes performance and load information for systems and networks, it could be used to determine whether a machine is susceptible to a Denial-Of-Service attack. In other words, it could be used to target systems that are running close to their maximum capacity.

This type of information is, however, extremely useful to legitimate users of a grid - it helps them determine the optimal locations for their workloads. If possible, it should only be made available to grid users, without exposing it to the outside world.

4. Summary of Security Risks

So far we have identified the following pieces of information, that are published to the OSG, as being potentially sensitive to a site:

1. Operating system and software level information
2. Local account names

3. Supported grid user DNs
4. Underlying authentication methods
5. Job-manager / batch-system information
6. Internal system paths
7. Job names
8. Error and failure information
9. System load and performance information
10. User activity at the site
11. Historical system availability data

While much of this data is very important to users and VOs on the grid, and essential in creating a robust and flexible grid architecture, it is important to design the systems that publish this information such that they can support the desired level of protection for the data. In other words, information should be restricted to legitimate users of the grid, and sites should have ultimate control over what information they wish to publish, and at what level of detail.

5. Recommended Grid Middleware Configuration

While software may evolve, and the specific methods for configuring software may change, the general goals for proper middleware configuration remain the same. The following recommendations will help provide some amount of control to sites that wish to protect sensitive data:

1. Override attributes that are considered sensitive with alternate values that can convey the equivalent information. For example the GIP allows named attributes to be overwritten by specifying them in a special file (*alter-attributes.txt*). This could allow a site to replace detailed software levels with more generic information.
2. Use site level collectors for multi-resource sites. This will allow the site to filter sensitive data at this level before forwarding it to OSG. Syslog-ng is designed with this sort of architecture in mind.
3. Turn down level of detail for the published information to the minimum required – during troubleshooting efforts, this can be turned up for more diagnostic information. This limits the overall exposure of the site.
4. Always use encrypted data streams and secure protocols to send information, instead of using clear text. Many OSG services, such as Gratia or Syslog-ng, offer both SSL and clear-text options to send data to their respective collectors. Sites should always use the former, when given a choice.

6. Recommendations for Data Protection

Additionally, it is in the best interest of the grid provider (OSG), to provide methods for protecting this data. This protection must happen in multiple ways:

1. All grid infrastructure software that transmits or collects data from public networks should support secure and encrypted communication protocols.
2. The software design should allow sites to override arbitrary attributes being published.
3. Information collectors should endeavor to authenticate the machines that publish site data – only machines whose identities can be verified should be allowed to publish their information. This prevents third parties from publishing fake or invalid data for a given site. GSI host certificates are an effective way to achieve this kind of authentication. CEMon already uses this, and the model could easily be extended to other OSG collection services.
4. Use of grid certificates to restrict access to data where possible. Web servers should attempt to verify the identity of the user before allowing access to grid resource information. Current technologies, (e.g. *mod_gridsite* for Apache based web servers) provide the ability to control access based on the user certificates. Additionally, this information could be restricted along VO lines, so that a VO is only authorized to access its own data.
5. Prevent indexing or caching of dynamic site information on web servers by search engines. This can be done by using files like *robots.txt* to prevent search engines from storing this information.
6. In the long run, there should be a concerted effort to consolidate software systems collecting similar information, so that site administrators and security officers have a single point of control for protecting such information. For example the Teragrid's Inca monitoring system consolidates resource validation, troubleshooting and monitoring functionality under a single engine.

Some of these features already exist in the OSG software, but there also needs to be a comprehensive effort to integrate these types of features across the middleware and collector infrastructure.

7. Applicability to Other Grids

While our work has largely been a case study on the OSG, the general principles of securing site information are applicable to any major grid infrastructure. Collection and publication of resource

information is a common feature across grids, and results in similar requirements and goals with respect to protection of such information.

Indeed, many of the discussed software systems are currently deployed in other grid infrastructures as well (e.g. CEMon and MonALISA at various EGEE sites). Other grids have their own information services providing equivalent functionality. The Teragrid uses the Inca monitoring system for resource availability, validation and monitoring purposes, collecting and publishing similar site information as that discussed in section 3. These systems face similar risks with respect to sensitive site information, and we expect the general techniques for protecting this information to be applicable as well.

There is an increasing trend towards interoperability among grids, with international collaborations and VOs driving usage and infrastructure requirements. There is a shift away from centralized grid providers, towards integrated VO architectures, where a given VO frames its own usage model. This points to cross-grid collection services that operate on a per-VO basis. Since VOs work in close collaboration with the major grid providers, many of the current technologies discussed have uses cases for such VO based services. For example, the ALICE VO uses MonALISA to provide integrated monitoring of its supporting resources. This means that VOs must also take site security requirements into consideration as they build their grid information frameworks.

8. Future Work

The focus of this work has been on the OSG, and its tools, infrastructure and metadata. It would be useful to extend this analysis to other major grid infrastructures such as the Teragrid or EGEE, to understand how they approach issues pertaining to sensitive site-related information. This would highlight common problems and solutions, and provide alternative approaches towards protecting site data.

Also, given that scientific collaborations are increasingly adopting the VO model of grid computing, where a VO maintains a certain amount of control over its own users and metadata, it would be interesting to analyze how VOs manage sensitive information, and how they publish and integrate this data across one or more grid infrastructures.

9. Conclusions

While a bulk of this paper has been devoted to the importance of protecting information that might reveal

weaknesses in a site's security infrastructure, this should not be taken as an endorsement of the "security by obfuscation" philosophy. We recognize that there is no substitute for hard security – regular fixing and patching of software, intelligent system monitoring, and strong security policies and practices are essential for a truly secure platform. However, practical security considerations demand that administrators account for the fact that not all vulnerabilities may be known at a given time. There may also be delays between the discovery and the patching of a vulnerability. Thus, it is prudent to minimize the amount of information available to a malicious entity and limit the extent of a compromise. While it is necessary to make certain kinds of information public for the success of open grid computing, it is also in the resource provider's best interest to understand the risks involved in doing so. Since grid architectures tend to be as generic as possible, some of the published information may be extraneous. The site must find a balance between how much information it seeks to publish about itself, and how much information it wishes to protect. It may also want to limit the consumers of this information to a controlled set of persons.

We believe that this paper would serve as a useful tool for sites that wish to identify these channels of information, so that they can determine the appropriate level of protection they wish to apply to their published data. We also hope to motivate further study and discussion on the protection of site information across various grid infrastructure and middleware providers.

References

- [1] The Open Science Grid Consortium, <http://www.opensciencegrid.org/>.
- [2] The Virtual Data Toolkit (VDT), <http://www.cs.wisc.edu/vdt/>.
- [3] OSG Grid Operations Center, <http://www.grid.iu.edu/>.
- [4] CEMon Service Guide, <https://edms.cern.ch/document/585040>.
- [5] OSG information services - discussion. A. Padmanabhan, OSG Site Administrators Meeting, Dec 2007.
- [6] Generic Information Provider. L Field. <http://twiki.cern.ch/twiki/bin/view/EGEE/GIP>.
- [7] GLUE Schema Specification version 1.3 Draft 3.
- [8] GRATIA, a resource accounting system for OSG. P. Canal, P. Constanta, C. Green, J. Mack. CHEP'07, Victoria, British Columbia, Canada. Sep 2007.
- [9] Usage Record – XML Format. Global Grid Forum. 2003
- [10] OSG Resource and Service Validation Project. <http://rsv.grid.iu.edu/documentation/>.
- [11] The CEDPS Troubleshooting Architecture and Deployment on the Open Science Grid. B. L. Tierney, D.

Gunter, J. M. Schopf. J. Phys.: Conf. Ser. 78 012075, SciDAC 2007.

[12] Syslog-ng Logging System.

<http://www.balabit.com/network-security/syslog-ng/>.

[13] MonALISA: An Agent Based, Dynamic Service System to Monitor, Control and Optimize Distributed Systems. I.

Legrand. CHEP'07, Victoria, British Columbia, Canada. Sep 2007.

[14] Globus Toolkit. <http://globus.org>.

[15] Gridsite. <http://www.gridsite.org/>.

[16] Google Hacking.

<http://www.acunetix.com/websecurity/google-hacking.htm>.

[17] Inca: User Level Grid Monitoring.

<http://inca.sdsc.edu/drupal/>.

[18] Teragrid. <http://www.teragrid.org/>.

[19] Enabling Grids for E-Science. <http://www.eu-egee.org/>.

[20] MonALISA Repository for Alice.

<http://pcalimonitor.cern.ch/map.jsp>.