GRID SECURITY

David P. Kelsey
CCLRC, Rutherford Appleton Laboratory, Chilton, Didcot, Oxon. OX11 0QX, UK

Abstract

This paper describes the main Grid security issues, in terms of both technology and policy, that have been tackled over recent years in the LHC Computing Grid (LCG) and related Grid projects. Achievements to date are described and opportunities for future improvements are addressed.

INTRODUCTION

The aim of Grid computing is to enable the easy and open sharing of resources between large and highly distributed communities of scientists and institutes across many independent administrative domains. Convincing site security officers and computer centre managers to allow this to happen in view of today's ever-increasing Internet security problems is a major challenge. Convincing users and application developers to take security seriously is equally difficult.

This paper concentrates on security issues from the point of view of HEP Grid computing. The LHC Computing Grid (LCG) [1] and the Enabling eScience for Europe (EGEE) [2] projects are both referred to as a way of illustrating what are in fact more general issues. Other HEP Grid projects are involved, as are other application communities, some of which have more stringent security needs, but it is not possible to give them all adequate coverage here.

SECURITY REQUIREMENTS

Many Grid projects have collected and documented large lists of detailed security requirements. One of the first to do this was the EU DataGrid project (EDG) [3]. EDG enumerated [4] 112 requirements across all areas of security, including Authentication, Authorization, Confidentiality, Integrity, Non-repudiation and various management and performance areas.

It would not be appropriate to consider these details here. Instead, just some high-level requirements are mentioned.

User Requirements

Grid users require

- Easy and **open access** to all of the compute and data resources available to them as members of an HEP experiment, across the Grid
- The need to **register just once** per Virtual Organization (VO), i.e. per HEP experiment. The scaling problems of users having to register at each and every LCG site, today a total of approximately 80 sites worldwide, would be totally prohibitive
- The need to **login just once** per session. This concept of single "sign-on" is another strong requirement. The identity and/or rights of the user are then delegated to other Grid services allowing them to act on the user's behalf without further requirements to "login"
- Not to be **bothered by security**. Users often confuse this with the incorrect statement of "I don't need security". Users always require, often without realising it, that the various services are available, i.e. not affected by compromise or denial of service attacks, and that the integrity of their data is assured. Having said that, they do still want security to be simple and not to get in the way of doing their work

Site Requirements

Computer centre managers and site security officers require:

- Full **local control** of access to their resources. No site is willing to fully delegate all control to outsiders. The decision to open up a site to a large community of remote users and giving them the ability to run programmes of their own choice is not one which is made lightly.
- **Knowledge** of user identities. Legal responsibility and the need to prove due diligence means that sites have to know who is using their resources. The registration information of each user which is gathered by the VO has to be made available to duly authorized managers in the sites.
- The ability to perform adequate **audits**. This has to be at the individual user level, giving the site the ability to determine who did what and when in terms of Grid access to their site.
- **Secure middleware**, applications and services; an extremely important requirement. Sites need to be convinced that the services they are asked to deploy have been well designed and implemented, with security taken fully into account right from the start. All failure modes must have been fully tested and the deployment and configuration of the service must also have been shown to be secure. The integrity of their whole site depends critically on the quality of the Grid services they run.

THE PLAYERS

Many HEP Grid projects are working on Grid Security. EDG, LCG and EGEE have already been mentioned. Other EU and national projects such as DataTAG (EDT)

[5], UK GridPP [6] and the Italian INFN Grid [7] are also very much involved.

In the USA, Grid3 [8] and its constituent projects, and now the Open Science Grid consortium (OSG) [9] are all very active. There are strong links, particularly in the area of Grid Authentication, with ApGrid [10] and ASCC [11] in the Asia/Pacific region.

Security, being one of the basic Grid infrastructures, is an area where interoperability is extremely important. For this reason many projects collaborate on common technology, interfaces and policies, either via bi-lateral links, e.g. members of OSG are active members of EGEE security groups, or via participation in multi-lateral activities such as the Global Grid Forum [12]

THE GRID SECURITY MODEL

A much simplified overview of the Grid security model, as used in HEP, can be presented in terms of four essentially independent areas; namely Authentication, Global Authorization, Local Authorization and Policy.

Authentication

Authentication is concerned with the proof of identity of people, machines and services. Given that so many of the early adopters of Grid were using the Globus toolkit [13] and for reasons of interoperability between these projects, an early decision was taken to base Grid authentication on the Globus Grid Security Infrastructure (GSI) [14]. This uses a Public Key Infrastructure (PKI) [15] based on X.509 certificates. A GSI extension to standard X.509, namely proxy-certificates, allows for short-lived, typically 12 hours, delegation of identity to other services enabling single sign-on. Authentication, however, only proves identity and has nothing to say about the user's right to access resources; hence the need for Authorization mechanisms.

Global Authorization

Global Authorization is the mechanism by which users are given the right to access the compute and data resources by their VO. The VO manages the list of their registered members, allocates them to groups, grants them the ability to acquire certain privileged roles, such as "production manager", and defines and manages global VO policy and resource allocation.

Local Authorization

Local Authorization is the mechanism by which sites and resource managers control access to their resources. This may be achieved by the mapping of global identities into local security mechanisms, such as UNIX users and file permissions, or via Grid-aware technology which is able to make decisions directly on the basis of the user's global identity (X.509 distinguished name) and/or the global authorization attributes given to the user by the VO.

Policy

Grid sites and/or projects define and maintain the various security policy documents. In terms of more dynamic policy, i.e. those attributes that can and need to be checked by Grid services before deciding whether or not to grant access, there are many different policy stakeholders. These include the VO and the site who contribute to building the many parts of the policy which are finally evaluated in a policy decision point and then used for the final enforcement, i.e. a "yes" or "no", at the policy enforcement point, thereby granting or denying access to the Grid service.

RISK ANALYSIS

The LCG Security Group [16], now called the Joint Security Policy Group (JSPG), as it works on policy and procedures for both LCG and EGEE, has performed a security risk analysis [17].

This considered risks arising from intentional or malicious incidents and also from accidental events. The malicious incidents were further split into sub-classes of misuse of resources, problems with confidentiality and/or integrity and disruption to Grid services.

For each problem identified, the risk, being the product of estimated likelihood and impact, was calculated. This analysis is useful in identifying those areas where the projects should concentrate their efforts.

The top four risks were identified to be:

- Misuse of the Grid resources to launch attacks on external sites
- The distribution, sharing or storage of illegal or inappropriate data
- Disruption to Grid services because of the exploitation of vulnerabilities in the Grid middleware or operating systems
- Disruption caused by damage from Trojans, worms and/or viruses

PROGRESS AND ACHIEVEMENTS

This section presents the three areas of Grid Security in which there has been the most significant progress over recent years. This is not to say that there have not been other achievements but these are the ones that have either already had a significant impact or are likely do so in the near future

Authentication

The EDG project made a very early decision to keep Authentication and Authorization separate. This followed the observation that employing institutes are best placed to assert the identity of an employee and confirm their membership of the academic research community. HEP experiments are in turn best placed to confirm the individual's membership of the VO and hence authorize them to access the VO's resources.

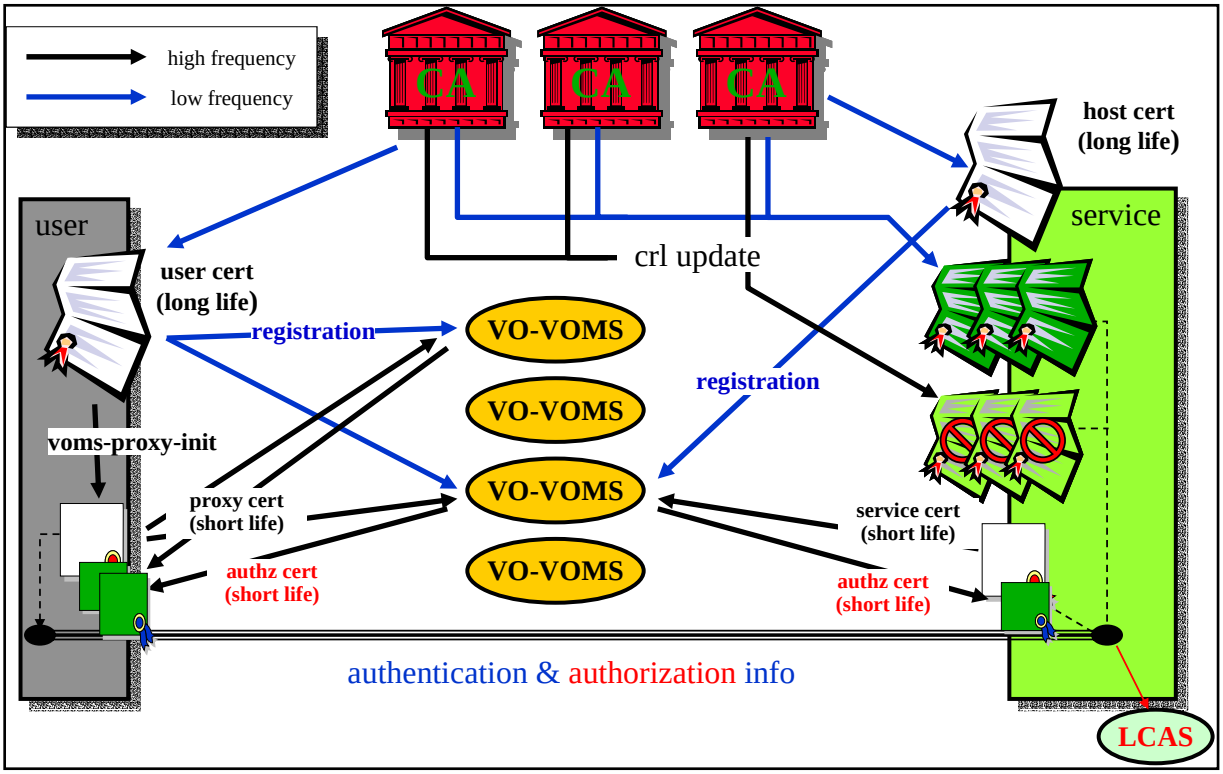

Figure 1: Mutual Authentication and Authorization using VOMS and LCAS

A second important aim was to provide the user with one, and only one, electronic identity for use in many different Grid projects. X.509 certificates and associated private keys are complex enough without requiring the user to manage many of them.

Participants from many national and international Grid projects collaborated together under the auspices of the EDG Certification Authority (CA) Coordination Group [18] to build a large global Public Key Infrastructure (X.509) for use in GSI-based mutual authentication of people and services. The CA is the trusted third party which cryptographically signs the certificate, thereby binding the name and public key to the identity of the person or service.

It was decided that the most appropriate scale was to aim for at most one Certification Authority per country, with the additional aim of providing Grid certificates for a broad community of Grid users and services in that country.

Today, this large global PKI is coordinated by the EU Policy Management Authority for Grid Authentication [19]. The members of the body determine best practice and minimum requirements for the accreditation of new CAs following a process of peer review. There are 27 accredited CAs today, covering not only the majority of Europe, but also North America, Taiwan, Israel, Pakistan, Russia and other places. Similar bodies are now being created in Asia/Pacific [20] and the Americas.

This infrastructure allows users from many parts of the world to participate in global Grid applications by the use of their locally obtained certificate. This is a major achievement.

Authorization

Many projects have been active in the important area of Grid Authorization. A few of the developed components are mentioned here. There are several others including VOX, SAZ, LRAS, PERMIS, and CAS [21]. There is not enough space to do a full review and comparison of all technologies.

Figure 1 shows the components and flow of information used in the mutual authentication of a user and a Grid service and of the authorization to use the service based on signed attribute certificates from several VOMS [22] servers and the enforcement of local policy via the use of LCAS [22].

The Virtual Organization Membership Service (VOMS) is a joint development of EDG and EDT. It allows the VO to create and define a database of the VO members and assign groups and roles to them. Users authenticate against the VOMS service using their Grid identity credentials and receive a digitally signed authorization attribute certificate from VOMS confirming membership of the VO and granting of requested groups and roles. This authorization information is then included inside the short-lived proxy certificate of the user as a non-critical extension of the certificate. This means that services which are capable of decoding the authorization information can do so while those which are not VOMS-

aware ignore the additional information and treat the authentication as a normal proxy certificate.

EDG also developed two local site authorization components: LCAS and LCMAPS.

The Local Centre Authorization Service (LCAS) is a framework for handling local policy decisions. Plug-ins are provided for a number of standard requirements, including the support of lists of approved or banned users.

The Local Credential Mapping Service (LCMAPS) provides the mapping of global Grid identities into credentials required in the local fabric, such as UNIX accounts or Kerberos tokens for access to AFS files.

These Authorization components, once fully deployed, will represent a major step forward in the ability to manage flexible policy and access control across a large distributed VO.

Policy

The LCG project started working on its security policy during 2003 with the aim of having this in place for the startup of the LCG-1 service in Autumn 2003. This policy, which actually consists of a top-level Security and Availability Policy and a number of associated documents, was prepared by the Joint Security Policy Group already mentioned above. The approval of the policy takes place in the LCG Grid Deployment Board and Project Executive Board.

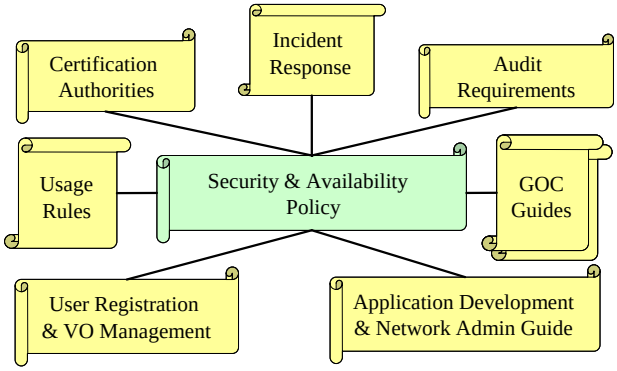

Figure 2: LCG Security Policy documents

It was recognised that a large-scale production Grid project, such as LCG, needed an agreed common security policy to:

- Define the attitude of the project towards security and availability
- Give authority for certain defined actions
- Put responsibilities on individuals and bodies within the project.

Figure 2 shows the set of documents currently forming the LCG policy. There are made available [23] as required reading to all new sites joining the LCG.

The actual negotiation and agreement of these policy documents took some considerable time and the achievement of what seemed like a daunting task at the outset should not be underestimated.

These policy documents are now in use in EGEE and various national Grid projects, although there is a need to make them less LCG-specific as mentioned below.

FUTURE OPPORTUNITIES

There are many issues in Grid security which are current topics of research and development. These span the whole range of technical, operational and policy concerns. Many of these are worthy of future development work and are already foreseen, for example, as future work for EGEE [24].

Authentication

There are many concerns about the problems of the current user-managed identity credentials in the Grid PKI. It is essential that the user's private key, encrypted on disk, is properly protected both in terms of file protections and in the quality of the pass-phrase used to encrypt it. As is often the case in the management of user-held ssh private keys, this does not always happen. The fact that the management of certificates and private keys inside web browsers is also very complex does not improve the situation.

Several solutions are being actively pursued to improve the situation. These are all aimed at removing the need for long-lived private keys held on disk:

- Credential repositories, such as the MyProxy service [25] will be used to hold long-lived user credentials. The service will issue short-lived proxy certificates on demand
- Site Integrated Proxy Services, such as the Kerberos CA, can issue short-lived X.509 credentials following authentication of the user by a local site mechanism, in this example Kerberos. A One Time Password mechanism for the initial authentication is also possible
- A security token, such as a SmartCard, where the private key is held securely inside a hardware device in the possession of the user. The operational costs and issues with this solution are considerable, but there are many potential benefits

The concern with all of these approaches is that the trust model is now different. Rather than having to rely on one CA per country and user-held credentials, we will have to agree how to build trust in site-based credential services with the related potential scaling problems given the large number of sites.

Authorization

Much progress has been made already in this area, but to date little of the exciting new technology has actually been deployed. It is important that the additional functionality in both global and local authorization, made possible by components such as VOMS and LCAS/LCMAPS is made available soon to solve the problems caused by the current inflexible mechanisms.

Delegation

Today the delegation technology used is not able to restrict the delegation to a subset of the user's rights. The ability to do this is an important requirement for future attention.

Firewalls

There are many issues related to the desire of many sites to restrict unnecessary incoming and outgoing network traffic. Today's Grid services all too often require ports to be opened in site firewalls which conflict with existing site policies. Work is needed in the area of proxy services, or application level firewalls, which can make the required links to the outside world in a more controlled way

Secure web services

The Grid middleware future is most definitely based on web services. There are many components in the web services security arena, WS-Security, only some of which have been standardised to date. Even though GSI proxy-certificates have now been standardised in IETF [26], there is little sign that the GSI delegation mechanisms will be supported in commercial web services. Much work is also required on taking today's Authorization components forward into a standards-based web services.

Middleware development

There are many potential issues with current implementations of Grid middleware. All middleware and application development must follow best practice for developing secure software. LCG has recently published a guide on this as part of the policy set [27]. Topics to be considered are ease of use, minimised need for off-site network connectivity, the use of least privilege, testing for failures as well as functionality, the existence of appropriate logging and well tested configuration and deployment. There is room for improvement in many of these areas.

Securing the end host

Several projects are exploring the benefits of controlling the environment on an end Grid system, or worker node, very carefully, either via sandboxing or some type of virtual machine. Via these mechanisms, the interaction with the base operating system is minimised and new virtual environments can be reinstalled for each use, reducing the risk of any long-lived security problem.

Policy and procedures

Next versions of some of the LCG Security Policy and procedures documents are now being prepared not only by LCG and EGEE but also in collaboration with OSG. The aim is to make the scope as general as possible, meeting the needs of many Grids and application communities. The November 2004 meeting of the EU eInfrastructure Reflection Group [28] plans to tackle common Acceptable Use and Authorization policies to encourage easy sharing of resources across Europe. The LCG/EGEE/OSG documents will be used as input to their reflections.

Operational security

There is a need to invest much more in this whole area than has been to date. The OSG project is working on the important area of Incident Handling and Response. LCG/EGEE is looking at using this as the basis for a revision to their procedures.

EGEE is in the process of creating a new Operational Security Coordination Team. One of their first tasks will be to organise and run several tests of the operational security procedures as one of a series of LCG service challenges.

Trust, responsibility and liability

Building "trust" between independent administrative domains is far from easy. Experience in the current Grid projects shows that an important ingredient in this is detailed consideration of what happens when incidents take place. Discussions often boil down to what responsibility, and perhaps even liability, is one Grid entity willing to take for problems caused to others in the Grid because of their or their users actions.

FINAL WORDS

There has already been much progress in the field of security for Grids, both in terms of technology and policy, over the last 4 years, as described in this paper.

There is, however, much more to be done:

- The security policies need to broadened in scope and made more generally applicable to many Grid projects and the whole Grid community, not just HEP
- The various components of Authorization technology which have been developed need to be deployed and used
- The various middleware services need to be made more secure, both in terms of functionality and quality, as part of the development and re-engineering process for production quality middleware
- The deployment and operational security teams need to ensure that the services are always configured and used in a fully secure manner

It is generally true that users of a fully secure, well designed and professionally run system tend not to notice security. When the situation is less perfect, there is a much greater chance of disruption due to compromise of services or data and/or lack of availability caused by attacks or denial of service problems. Under these circumstances the users most definitely will notice!

While not particularly wishing to end on a negative note, it is an unfortunate consequence of human nature that the major effort required for fully securing Grids will only be possible once we start to experience the inevitable security problems of the current implementations of Grid services. This has, after all, been the way that security has

gradually improved in all operating systems and other non-Grid internet applications. I see no reason to believe that Grids will be any different.

In the meantime, many dedicated and talented people will continue to work hard in this challenging area. They require and deserve the full support and understanding of the user and developer communities.

It will be interesting to review the situation at the next CHEP conference in 2006.

ACKNOWLEDGEMENTS

I acknowledge the support of the EU and the various national funding agencies, especially PPARC in the UK, which has allowed me and others to work on Grid security within the GridPP, EDG, LCG and EGEE projects. I am grateful to Ákos Frohner and Ian Neilson, both from CERN, for the preparation of figures 1 and 2, respectively. I particularly thank all of my colleagues, too numerous to mention by name, from the various projects for their major contributions to Grid security and for many illuminating discussions. I thank the organisers of this conference for inviting me to present this paper.

REFERENCES

[1] http://lcg.web.cern.ch/
[2] http://www.eu-egee.org/
[3] http://www.edg.org/
[4] https://edms.cern.ch/document/340234/
[5] http://datatag.web.cern.ch/
[6] http://www.gridpp.ac.uk/
[7] http://grid.infn.it/
[8] http://www.ivdgl.org/grid2003/
[9] http://www.opensciencegrid.org/
[10] http://www.apgrid.org/
[11] http://www.sinica.edu.tw/
[12] http://www.gridforum.org/
[13] http://www.globus.org/
[14] A Security Architecture for Computational Grids. I. Foster, C. Kesselman, G. Tsudik, S. Tuecke. Proc. 5th ACM Conference on Computer and Communications Security Conference, pp. 83-92, 1998.
[15] There are many good overviews of PKI, such as: Understanding PKI: Concepts, Standards, and Deployment Considerations, Second Edition. Carlisle Adams and Steve Lloyd. Addison Wesley. November, 2002. ISBN: 0-672-32391-5
[16] http://proj-lcg-security.web.cern.ch/
[17] http://proj-lcg-security.web.cern.ch/proj-lcg-security/RiskAnalysis/risk.html
[18] http://marianne.in2p3.fr/datagrid/ca/
[19] http://www.eugridpma.org/
[20] http://www.apgridpma.org/
[21] VOX, SAZ & LRAS:http://www.uscms.org/s&c/VO/
PERMIS : http://sec.isi.salford.ac.uk/permis/
CAS : http://www-fp.globus.org/security/CAS/GT3/
[22] https://edms.cern.ch/document/344562/
[23] http://proj-lcg-security.web.cern.ch/proj-lcg-security/sites/for_sites.htm
[24] https://edms.cern.ch/document/487004/
[25] http://grid.ncsa.uiuc.edu/myproxy/
[26] http://www.ietf.org/rfc/rfc3820.txt
[27] https://edms.cern.ch/document/452128/
[28] http://www.e-irg.org/