

Certificates for Automated Clients

Status of This Memo

This memo provides information to the Grid community regarding new class of certificates for automated Grid Clients . It does not define any standards or technical recommendations. Distribution is unlimited.

Copyright Notice

Copyright © Global Grid Forum (2003). All Rights Reserved.

Abstract

In evaluating Grid technology for production use, several members of the DOESG community feel there is a need for another class of certificates to be used for authenticating automated clients that connect to Grid servers. The usage of these new certificates is different enough so that current CP policies will need to be updated.

Contents

Abstract	1
Introduction	2
Motivation	2
Discussion.....	2
Security Considerations	4
Author Information	4
Glossary	5
Intellectual Property Statement	5
Full Copyright Notice.....	5
References.....	5

Introduction

Current DOESG and EDG certificate policies cover issuing x509 certificates for individual identities and for authenticating hosts and services. In evaluating Grid technology for production use, several members of the DOESG community feel there is a need for another class of certificates to be used for authenticating automated clients that connect to Grid servers. The usage of these new certificates is different enough so that current CP policies will need to be updated.

We propose that a new namespace, differentiated by the "OU=" attribute be created that can be used to distinguished between certificates used for automated services, from the certificates used to identify hosts and individuals.

Motivation

One of the major differences between a production environment and a research environment is in procedure: production environments typically have a large collection of procedures that are used to maintain "production quality", examples are service monitoring systems, backup systems and various housecleaning utilities. Production environments also typically support repetitive, ongoing processes - either internal system processes, or processes relating to the applications being run at the site, for example, processes that move datasets from one site to another.

These procedures and repetitive processes are typically automated, run across multiple machines, and generally run using an identity with the necessary privileges to perform their tasks, and little else. As Grid technology is put into production, we are finding that we need to perform the same kind of repetitive, automated tasks, but using Grid credentials over Grid services. However, there does not seem to be any allowance within the DOESG and EDG certificate policies for issuing certificates to support these kinds of tasks.

Discussion

Currently, 2 classes of certificates are issued:

Personal Identity certificates - these are used to identify individuals, and have the "OU=people" attribute in their subject distinguished names. Typically, these certificates are used to authenticate the client side of a client/server request, and the certificate must be on any machine from which a client request could originate. The private key is supposed to be while stored encrypted, and user action is required to generate a proxy certificate based on the key. The accountable party is the individual identified on the certificate.

Host/Service certificates - these are used to authenticate a host, or a server running on the host, and generally have the "OU=services" attribute in their subject distinguished names. These certs are generally used to authenticate servers to clients, and are uniquely located on that machine. There is usually no reason for a client connecting to a server to present an "OU=services" certificate (for example, you would not expect to see such a cert in a grid-mapfile). The private key is stored unencrypted, and is used "as-is, without generating a proxy certificate. The accountable party is nominally the individual that submitted the certificate signing request, but in fact it is the IT group that maintains the machine or service identified on the certificate.

A certificate for an automated client has attributes that span both these types of certificates:

- 1) The private key needs to be stored unencrypted so that automated tools can use it. This is similar to a Services cert.

- 2) The certificate may be used to generate proxy certs, and is used to authenticate the client side of a connection. You could reasonably expect to see the subject names of these certs in a grid-mapfile, as you would a personal identity certificate.
- 3) The accountable party may not be clearly identified in the DN on the cert, but might nominally be the person submitting the certificate signing request. However, once again, the true accountable party would be the IT group that operates the automated client. This is similar to the Services certificate.
- 4) Many of these automated clients do the same task from multiple machines. For example, automated backup clients all do essentially the same task, and usually run as the same restricted identity across multiple machines. For data replication clients, the replication service may operate from multiple machines. Getting and managing multiple certificates for a small number of machines is annoying, getting and managing certificates for hundreds of machines (such as the typical HENP cluster) becomes problematic. IT staff may be inclined to reuse the same certificate on multiple hosts. If this is not acceptable practice, there needs to be some mechanism to prevent or detect it.

Because these certificates are neither fish nor fowl in the current model, there needs to be clarification on how they can acceptably be handled. Another set of use cases that hasn't been discussed, but are related, is authentication of peer to peer services.

The discussion on the DOESG mailing list also seemed to identify these main issues:

Identity/accountability - who exactly would be accountable for these certs? The correct answer from the perspective of an organization is that an IT group would be accountable for the certificate. Individuals within an IT group may share or change responsibilities, and turnover is not unheard of - at which point, the certificate will no longer accurately reflect the accountable party (arguably, it may never have in the first place). In a managed, production environment, an *organizational abstraction* is the actual accountable party, and not an individual. However it might be argued that assigning accountability to an *abstraction* dilutes accountability.

Shared certificates - the notion of sharing the certificate across multiple machines increases the potential damage from a compromised private key. However, this same form of risk is managed on a daily basis by IT groups that have standard passwords for system accounts (such as root). The exposure is contained by the fact that there is a de facto trust boundary at the site perimeter - no machine managed by another group shares the same "root" password. If shared certificates were issued, it would be wise to ensure that all machines that trust that certificate are under the same administrative domain.

Analog from existing practice

Kerberos is already used by automated clients at many sites. To explore the problem space, we examine the practice at one such site, Fermilab.

Service principals are used directly as clients when it is the service itself that is acting as the client of another service. For example, when a mass storage systems initiates a transfer to or from another mass storage system, it identifies itself as the same name by which it's known to clients.

When the automated process initiating a connection isn't a known service itself, it uses what's known at this site as a "project principal." The name of such a principal has three components (besides the realm) which identify its function, the responsible organizational unit, and the single host from which it is expected to authenticate itself. For example, "level3/cdf/b0dax42.fnal.gov@FNAL.GOV" acts as part of the Level-3 trigger for the CDF experiment, and authenticates using a keytab file (analogous to the private key corresponding to a certificate) stored on the host b0dax42.fnal.gov. Since a Kerberos initiator must first communicate with the KDC to obtain credentials, the KDC has a record of where the

authentication requests for this principal come from and, to the extent that DNS can be relied upon, may reject requests from elsewhere than that one host or raise an audit flag for them.

A large class of automated client requests are performing actions for a specific user, generally through the unix "cron" scheduled execution facility. Fermilab users have the ability to create and destroy per-user, per-host principals for themselves with names of the form *username/cron/hostname@FNAL.GOV* and to create a keytab file on the corresponding host. Tying each such principal to a host and a user facilitates recovery from any security incident involving the user or the host.

Proposals

One of the most basic, and probably least controversial proposals is to create a separate namespace for this new class of certificates. Reusing service certificates sets up a precedent of having "server side" certificates suddenly become client side proxy certificates, and makes it harder to clearly identify misuse of certificates. Using personal certificates for these services requires that individuals expose their personal certificates to theft because the private key is not encrypted. This is especially dangerous for administrators, whose personal accounts may contain sensitive information and or privileges. In addition, these client service accounts typically map onto severely restricted local accounts, instead of the normal user account.

The more controversial proposals would include somehow altering the policies to make allowance for the notion that "permanent" IT groups (as opposed to ad hoc virtual organizations) can be accountable parties. Possibly the most controversial is how to handle the notion of shared certificates - this is a convenience for IT organizations, but needs to be mitigated by security concerns.

None of the proposed changes can be effected unilaterally, and we also believe that the issues are of concern to European Data Grid users as well as the DOE Science Grid users. In fact, it will likely be a concern for all large Certificate Authorities, and the situation provides an opportunity for the DOESG and EDG to exercise leadership in setting standards for how these certificates can be handled.

Security Considerations

This is a REQUIRED section.

Author Information

Stephen Chan
NERSC/LBNL
One Cyclotron Road
Berkeley, CA USA 94720

Matt Crawford
Fermilab MS-369
PO Box 500
Batavia IL 60510
USA

Glossary

TBD

Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director.

Full Copyright Notice

Copyright (C) Global Grid Forum (date). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

References