Certificates for Automated Clients

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.

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 it's task, 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 it's distinguished name. 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 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 it's distinguished name. 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 is used to generate proxy certs, and is used to authenticate the client side of a connection. You could reasonably expect to see these certs in a grid-mapfile. This is similar to a personal identity cert.
3) The accountable party is cannot be clearly identified from the DN on the cert, but might nominally by 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 cert
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. Reuse of certificates makes these certs similar to personal identity certs.

Because these certs are neither fish nor fowl in the current model, there needs to be clarification on how they can be handled in a standard fashion. Another set of use cases that haven'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.

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 affected 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.