GPLAZMA 'GRID-AWARE PLUGGABLE AUTHORIZATION MANAGEMENT SYSTEM': INTRODUCING RBAC (ROLE BASED ACCESS CONTROL) SECURITY IN DCACHE/SRM

Abhishek Singh Rana[†], Frank Würthwein[‡]
University of California at San Diego, CA 92093, USA

Timur Perelmutov, Robert Kennedy, Jon Bakken, Dane Skow, Ian Fisk
Fermi National Accelerator Laboratory, IL, USA

Patrick Fuhrmann, Michael Ernst
Deutsches Elektronen Synchrotron (DESY), Hamburg, Germany

Markus Lorch, IBM, Germany

Abstract

We introduce gPLAZMA (grid-aware PLuggable AuthoriZation MAnagement) for dCache/SRM in this paper. Our work is motivated by a need for fine-grained security (Role Based Access Control or RBAC) in Storage Systems, and utilizes VOMS extended X.509 certificate specification for defining extra attributes (FQANs), based on RFC3281. Our implementation, the gPLAZMA module for dCache, introduces Storage Authorization Callouts for SRM and GridFTP. It allows using different authorization mechanisms simultaneously, fine-tuned with switches and priorities of mechanisms. Of the four mechanisms currently supported, one is an integration with RBAC services in the Open Science Grid (OSG) USCMS/USATLAS Privilege Project, others are built-in as a lightweight suite of services (gPLAZMALite Authorization Services Suite) including the legacy dcache.kpwd file, as well as the popular grid-mapfile, augmented with a gPLAZMALite specific RBAC mechanism. Based on our current work, we also outline a list of future tasks.

This work was undertaken as collaboration between PPDG Common project, OSG Privilege project, and the dCache/SRM groups at DESY, FNAL and UCSD.

INTRODUCTION

Data Grids, or Storage Grids, are projected to be the most cost-effective and efficient technology solution for management of PetaByte-scale distributed data and metadata. In 2007, Large Hardon Collider (LHC) at CERN, will produce a sustained stream of data in the order of 300MB/sec, equivalent to a stack of CDs as high as the Eiffel Tower once per week. LHC thus serves as the motivation behind many global projects and efforts aimed at designing and developing next-generation compute, storage, and network systems.

A major player in the set of these storage systems is dCache/SRM. dCache/SRM has proven to be capable of managing the storage and exchange of several hundreds of TeraBytes of data, transparently distributed among dozens of disk storage nodes.

Motivation

Such Data Grids to support future LHC needs, and underlying storage systems, are being built around the concept of Virtual Organizations (VOs). A Virtual Organization (VO) is a logical collection of resources and personnel derived from subsets or even supersets of real enterprises. A VO may support a complex set of relationships defining which users are part of which projects within the VO and which users are designated to perform higher-privileged roles within the VO at various times. Examples of these relationships include:

- A single user may be a member of several projects. Not only are there different resource allocations for these projects, but also the resource usage must be properly charged to the correct project and the correct VO.
- A single user may have multiple roles in a VO. At times the user may act as a project administrator and at other times the user acts a regular VO member.
- A group of individuals may alternate the administration of the VO with only one individual at a time acting as the administrator. To insure non-interfering administration of the VO, the administrator function may be permitted to a given individual only during a pre-determined period of time.

[†]Corresponding Author. Abhishek Singh Rana is Senior Software Architect at UC San Diego, (Email: rana@fnal.gov, abhisheksinghrana@gmail.com).
[‡]Frank Würthwein is Associate Professor of Physics at UC San Diego, (Email: fkw@ucsd.edu).

Drawbacks of Current Approaches

Security mechanisms exhibit several limitations when attempting to cope with the complex VO structure illustrated above. This is true for data grid middleware in general and dCache/SRM in particular.

Enforcement mechanisms at the storage system level may not be aware of VO groups or roles. Storage systems therefore have no basis for differentiating between users from a given VO. A more flexible authorization mechanism is required that can distinguish between individual users and between the roles an individual user can hold.

Multiple users are frequently mapped to the same POSIX user account (e.g., all users from a given VO are mapped to a single, shared user account). But in this scheme, every access is granted with the full set of access privileges that the VO as a whole is authorized to assume. Moreover, user operations are not well insulated from each other. Such many-to-one mappings provide limited support for management of authorization policies and have low system security.

Site-level policies are maintained on each host. This leads to administrative and maintenance burden. A site-centralized identity management system for such policy information can improve the maintenance and promote the consistent enforcement of site access control policies.

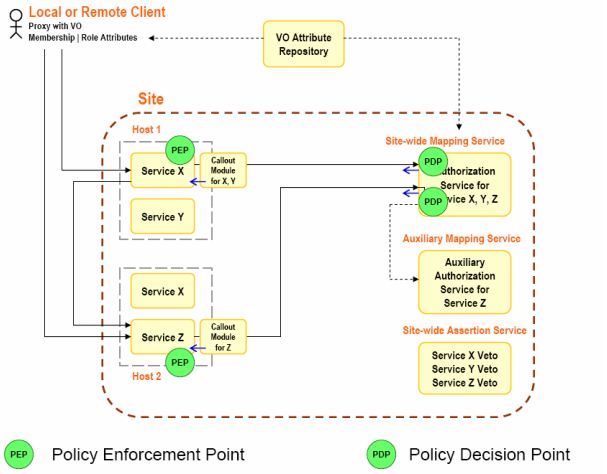

Fig.1. OSG Approach to Authorization

The OSG (Open Science Grid) Approach

As outlined in Fig.1, the approach consists of:

- VO-global specification of privilege per role.
- Site-central mapping of all roles to site's local implementation of privilege attributes.
- Local enforcement of privilege attributes.
- Use of VOMS extended X.509 Attribute Certificate specification for defining extra attributes FQANs.
- Based on RFC-3281. FQANs contain role and VO membership information for a User.
- VO defines roles and associated privileges by specifying expected functionality. For eg., *cmssoft* may install software in area that is read-only by all *cmsuser* jobs running on a site, *cmsphedex* may have special access to dCache/SRM system.
- Site maps VO-scoped grid identities to local-scoped resource-level identities.
- Site-wide management of mapping.
- Service-level granularity of mapping.
- Site enforces VO privilege policies within local scoped identities.
- Authorization is granted to an access request if the user is allowed by VO, but not vetoed by site.

GPLAZMA IN DCACHE/SRM - ROLE BASED ACCESS CONTROL, SITE-WIDE SECURITY, PLUGGABLE MODES OF AUTHORIZATION, QUASI-FIREWALL

The gPLAZMA implementation in dCache/SRM is a dynamically loadable module fully integrated with SRM server and GridFTP servers within standard dCache/SRM software packages. Accompanying components are GUMS (Grid User Management System), PRIMA (Privilege Management System) java classes to implement OGSA-Authz, and a Storage Authorization Service (SAS) that works as an intermediary to GUMS server. GUMS itself serves as the interfacing service with VOMS (VO Management System) and VOMRS (VO Management and Registration Service) servers. VOMS proxy management clients are used to obtain extended X.509 proxy certificates with extra membership and role attributes embedded by a VOMS server as X.509 Attributes Certificates (ACs). We have also attempted a custom implementation of OGSA-Authz that may serve as a more feature-rich and lightweight substitute to PRIMA java classes if needed.

Role Based Access Control/Authorization

Users obtain extended X.509 proxy certificates from VOMS servers using voms-proxy-init in VOMS client suite. In addition to user identification X.500 Distinguished Name (DN), these short-lived proxy certificates also have VO membership and role information. Our implementation also supports backward compatibility with plain proxy certificates obtained using grid-proxy-init in Globus GSI client suite. Such certificates only contain the user identification DN.

A VO takes the responsibility to regulate the types of proxy certificates all member users must obtain, and thus make use of role based access control (or not).

When the SRM server (or GridFTP servers) in dCache/SRM receives a data transfer request from a user, the server makes a callout to the gPLAZMA module and passes the associated security handshake context (GSSContext) to the module. gPLAZMA extracts user identification, trust and authorization information from

this context. This includes DN of user credentials, DN of the SRM server certificate, FQANs containing the user's VO membership or a subgroup membership with the associated role, DNs of the Certificate Authorities (CAs) and Attribute Authority (AA). It is to be noted that VOMS server, which had issued the Attribute Certificate (AC) to a requesting user, is an AA. gPLAZMA module repackages this information and loads its various *plug-ins*. We designed a common interface and decided to build plug-ins to interact with different modes of authorization as discussed below.

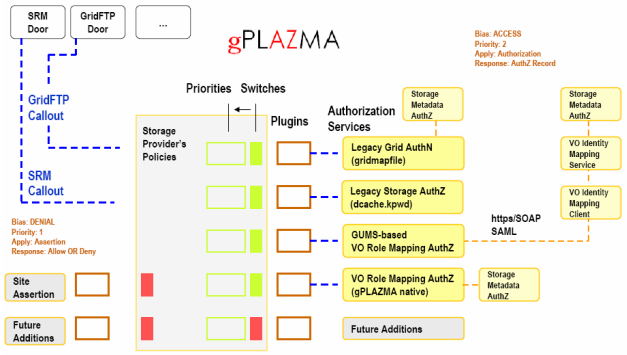

Fig.2. Architecture of gPLAZMA in dCache/SRM

These modes can either be from the built-in gPLAZMALite Authorization Services Suite or can be one among the pre-integrated modes from our partners.

Pluggable Modes of Authorization

gPLAZMA module, upon loading the configuration file dcachesrmgplazma.policy, decides which mode is to be used to perform authorization, or to use multiple modes if one fails, and the *priorities* (ie., order) with which to attempt multiple modes. Modes can be enabled or disabled using *switches*. The order of attempting different modes is fully customizable and provides for pluggability and overriding of authorization services, based on a given site's needs at a given time. There are four different modes supported in the first implementation:

- dcache.kpwd: This is original legacy dCache/SRM security mechanism and is supported for backward compatibility. Provides one-to-one or many-to-one mapping of a user identity to local storage resource privileges.
- grid-mapfile (built-in gPLAZMALite suite): This is based on Globus GSI and has been built-in gPLAZMA to be used by many small VOs and sites that do not use role-based authorization. It was one of the original file formats in Globus software, and is almost the common default of most Globus services. Provides one-to-one or many-to-one mapping of a user identity to local storage resource privileges.
- grid-vorolemapfile (built-in gPLAZMALite suite): This is a new file format introduced by gPLAZMA, and can be considered as the next generation evolution of grid-mapfile. This service makes full use of membership and role attributes, and thus provides many-to-many mapping of a user identity to local storage resource privileges.
- SAS web service interfacing with GUMS service and gPLAZMA metadata service. SAS provides the site-central mapping using GUMS. This service makes full use of membership and role attributes, and thus provides many-to-many mapping of a user identity to local storage resource privileges.

Site-wide Security

GUMS (Grid User Management System) has been extended as part of the OSG Privilege project to an online identity mapping service. GUMS maps a user's grid entity to a site-local username at the requested resource based on the entity's X.500 DN and FQAN. It facilitates site-centralized management of authorized users, and possibly site-consistent allocation of local user accounts. Within GUMS, a variety of allocation alternatives are available including dynamic allocation from a pool of user accounts, mapping to role-specific shared accounts, and mapping of individual (statically allocated) accounts.

SAS serves as an intermediary to GUMS, and uses an internal gPLAZMA Metadata service to add extra authorization decision factors to the username received from GUMS. These include UID, GID, authorized user's home path, root path, and access privileges such as 'read-only' or 'read-write'. This entire privilege set is returned to components at dCache/SRM. gPLAZMA returns this authorization decision and privileges to the SRM server or GridFTP servers. The data transfer request can then be processed accordingly by the underlying filesystem components (currently, PNFS) in dCache/SRM.

GUMS and SAS communicate using SOAP over an HTTPS connection. Emerging specification sets, SAML and XACML, are employed in this communication as agreed upon in the OGSA-Authz interface specification.

Quasi-Firewall

A site can turn all storage authorization 'off' immediately, without a need to shut individual services in dCache/SRM. This can prove useful in provisioning timely incident response to security breaches, identity thefts, and similar emergencies. This quasi-firewall functionality comes into action when all switches are turned off. Thereafter, the SRM server and GridFTP servers, although still up and running, refuse to entertain all access requests.

A site can also take refuge by only enabling a mode exposing less risk, and authorizing only well-known selected users until a security incident is resolved.

None of the dCache/SRM services need to be restarted for gPLAZMA's configuration related changes to take effect.

Current Deployment and Usage

gPLAZMA and associated software components are deployed and used at USCMS Tier1 Center at Fermi National Accelerator Laboratory, and at USCMS Tier2 Center at UC San Diego. Production usage on all USCMS sites is expected in Summer 2006.

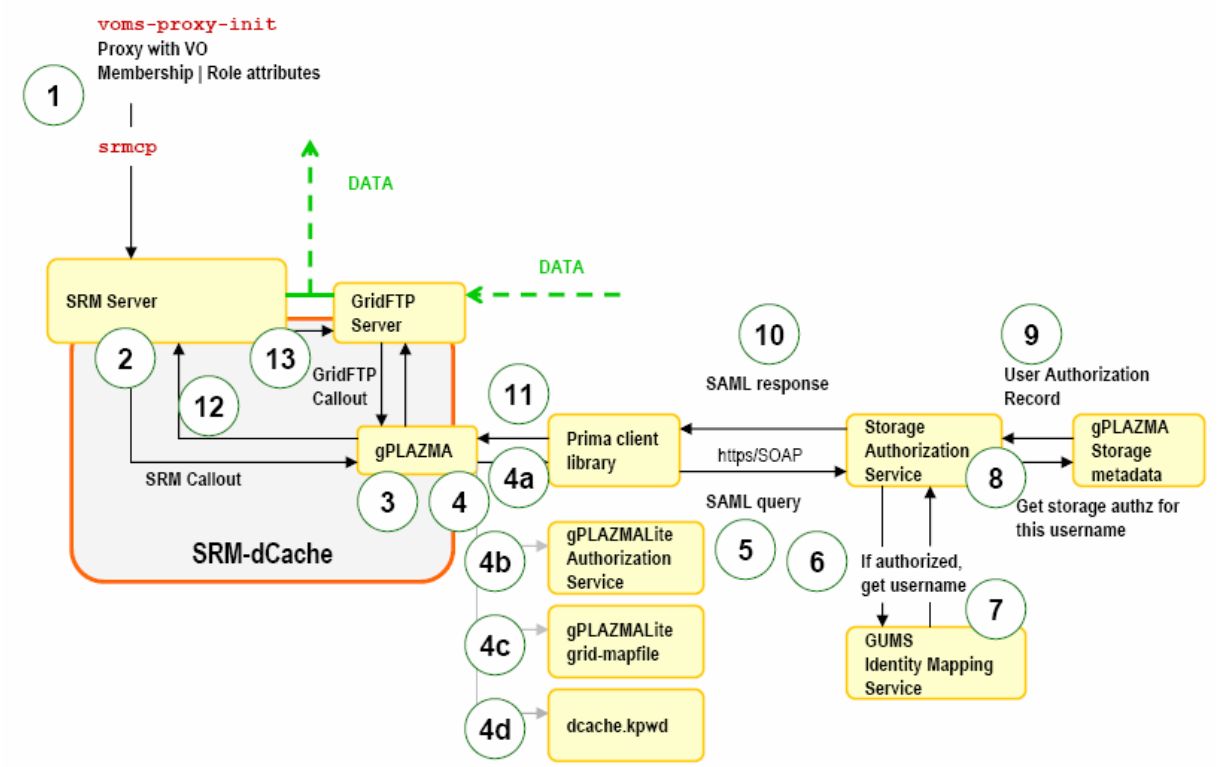

Fig. 3. Current gPLAZMA Design and Implementation

FUTURE WORK

A few near-term tasks that need to be undertaken to enhance current capability and add new features in the current gPLAZMA implementation in dCache/SRM are:

- Integration with DCap server/door to provide role-based functionality to users of this important native protocol.
- Replacement of current ASCII file formats with a database (eg., MySQL) backend.
- Evaluation of Chimera, the next-generation filesystem in dCache, followed by gPLAZMA-Chimera integration. This may lead to design of new authorization decision factors. For e.g., per VO or user or role priority, per VO or user or role quota, ACLs, space reservation attributes, etc.
- Scalability of GUMS-based SAS mode needs to be evaluated. For high-throughput processing of large number of requests, overall effect of network latencies, along with CPU and memory load on SAS and GUMS services needs to be determined.
- Optional usage of site-wide /etc/passwd and /etc/group system files if needed.
- Addition of a 'Site Assertion Service' (as part of built-in gPLAZMALite suite) with a bias on denial, such that selected users, VOs, or roles are always denied access. This will provide the important security feature known as *blacklisting*.
- Adaptation to evolve with future policy, security, and storage management standards.

CONCLUSIONS

We discussed general limitations of previous authorization approaches, listed the current OSG approach associated with OSG Privilege Project design, and summarized our introduction of gPLAZMA architecture and implementation in dCache/SRM.

ACKNOWLEDGEMENTS

The authors would like to express sincere gratitude to all members of OSG Privilege Project, of Particle Physics Data Grid (PPDG), and of dCache/SRM teams at DESY, FNAL, and UCSD. This work was supported in part by the National Science Foundation under Award No. 0533280, and by the U.S. Department of Energy under Award No. DE-FC02-01ER41201.

REFERENCES

[1] http://www.dcache.org
[2] http://srm.fnal.gov
[3] http://sdm.lbl.gov/srm-wg
[4] http://home.fnal.gov/~rana/gplazma