XSEDE Science Gateway Security Policy & Guideline

January 16, 2014

Version 1.1

Table of Contents

A. Document History

Relevant Sections	Version	Date	Changes	Author
Entire Document	0.1	01/16/2014	Baseline	A. Slagell

B. Document Scope

"Community user" logins for Science Gateways present special accounting and security issues. This policy is designed to ensure that such users and logins are properly tracked and identified within the add-user process.

C. Document Body

C.1. Introduction

"Community user" logins for Science Gateways present special accounting and security issues. This policy is designed to ensure that such users and logins are properly tracked and identified within the add-user process.

C.2. Background: Science Gateways and Community Users

The XSEDE allocations process permits, and the XSEDE science gateways activities encourage, activities whereby entire research communities use XSEDE resources. The allocation for these communities is overseen by a project PI, but the precise number and identity of the investigators using the allocation may not be known in advance by the project PI, and in most cases, the project uses a portal or gateway interface to provide access to XSEDE resources by this community.

The allocations policy [1] states:

"Gateways or Community Services: Projects of this type provide services to a large community of users who are typically not directly collaborating with the project PI. An example of such a project would be an application portal service providing access to software and computer time to a community of biology researchers via a web-based interface. Requests to provide such a service must describe the details of the services provided, the methods used, the expected consumption of resources, and mechanisms for monitoring the users and usage of the service. Statistics of community usage should be reported quarterly and in renewal requests for resources, progress reports, and end-of-project reports."

This community user account policy document details the responsibilities and procedures for the XSEDE accounting system. It includes security recommendations for service providers. It also describes the responsibilities of and capabilities provided to allocated community services or science gateway projects.

C.3. Responsibilities and Procedures

C.3.1. XSEDE Accounting System

The XSEDE accounting system responsibilities include (a) permitting the creation of such community user logins within 10 business days and (b) identifying those users in a uniform manner so that security and accounting needs can be met.

When applying as a science gateway [2] the applicant may request a community account.

When a community user request is submitted, a community user login is created, which the requestor can access through normal XUP password reset procedures.

- To associate a certificate with the community account, the owner of this community account can submit a Distinguished Name (DN) for a valid host certificate (from a CA trusted by

XSEDE) with the original account request or later by logging into the XUP with this community account and associating DNs there.

- The contact information (email, phone, address) for the community user will be the gateway or project contact with primary oversight for the login (i.e. the PI). The gateway form will also allow entry of a primary administrative contact (if different from the PI).

Once created, the XSEDE accounting system will treat a community user as it would treat any other individual—assign the user to resources, permit usage uploading from assigned resources, and so on. The gateway or community project contact will be able to monitor usage by this "user" via the standard tools available to all users (e.g., xdusage or user portal).

C.3.2. Service Providers

Community accounts allow access to XSEDE resources through different means than command line logins by individuals. While community accounts can provide more secure access to XSEDE machines through a fully secured gateway server, the web interfaces often used and the large number of end users can introduce potential security risks. Service providers (SPs) are not required to place additional constraints on community accounts, but if they choose to do so, a uniform approach will make development easier for gateway programmers using multiple sites. For this reason the following recommendations are made.

The common set of recommended restrictions below has been developed jointly by XSEDE's security and gateway working groups.

C.3.2.1. Login Ability

If a service provider does not allow gateway developers to log in directly to a community account, the Unix "substitute user" or sudo access is recommended to allow a restricted set of developers to "sudo" into the community account for command line access and interactive development. Adding these developers to a distinct Unix group would be an additional benefit for gateway developers. In this way, developers could write to community software areas with the tools to be executed by the gateway account, and with properly set permissions the gateway account would be unable to modify the tools but just run them. Combined with a form of command restriction, such as, restricted shells [10] that could be setup by the SP, these restrictions can make it difficult for an exposed community account credential to be abused.

C.3.2.2. Remote Job Submission

SPs may support a number of different protocols for remote job submission. Examples include SSH (with passwords, key pairs and/or GSI authentication) and grid software such as Globus [3], Unicore [4] and Genesis II [5]. The time-limited nature and revocability of grid certificates used by GSI-SSH and grid software adds additional security beyond *unrestricted* key pairs and passwords used by SSH if only X.509 certificates are accepted. Furthermore, grid certificates allow use of GRAM and other tools, which can alleviate the need direct shell login with the community accounts. If on the other hand direct login must be utilized, it is preferable to do so with restricted ssh keys that only

allow certain commands to be executed. This is done using the *from* and *command* options in the SSH authorized_keys file [9].

C.3.2.3. Community Software Area

A community software area (CSA) gives developers access to disk space for the installation of executables and libraries that will be used by a community of users. CSAs are unpurged, backed up spaces accessible via a single path ($XD_COMMUNITY/my_directory_name) on all requested XSEDE systems. CSAs eliminate the need for developers to use home directories to install shared community applications, libraries, scripts, etc. Developers rather than XSEDE staff are responsible for software in these areas.

C.3.2.4. Gateway End-User Accounting Records

Science gateway end-user reporting is required of both science gateways and SPs. Some SPs require that gateways capture attributes and send them with job records to SPs. These SPs are then required to send locally-collected attribute information to the XSEDE Central Database through AMIE packets. Information on how to do this for gateway developers and SPs can be found at [11].

C.3.3. Community Activities and Science Gateways

Responsibilities of a PI, including those of a gateway PI, are described in the Acceptable Use Policy [6]. The gateway PI accepts responsibility for following the general XSEDE user policies as well as the additional policies in this document, and to appropriately make their use community aware of any relevant XSEDE restrictions, such as, storing personally identifiable information or personal health information on SP systems without prior authorization from the SPs in questions.

A key provision of these policies worth drawing attention to is that gateways must not directly consume XSEDE passwords, but instead should use the XSEDE MyProxy Delegation Service [12]. This follows directly from the XSEDE AUP [6].

It is recommended that researchers accessing science gateways can run executables provided by the gateway developer, but cannot upload arbitrary executables. Otherwise it may be difficult to for the gateway owner to attest to the function and validity of a particular job and thus has the potential to slow incident investigations and keep the community accounts disabled longer.

Community accounts can be requested at the XSEDE Portal and may be done as part of the development process. A PI must provide:

- IP address or DNS name of the gateway machine and any test servers.
- Current estimated long-term disk storage requirements for the community account. This number can be modified in the future to accommodate gateway growth.

Optionally, for each script or executable in the named directory, they may provide

- Estimated maximum number of processors/nodes

- Estimated maximum run time
- Estimated short-term storage requirements per user per job

This information is stored in the XSEDE Central Database and associated with the community account username, and it can be further edited as development proceeds. Some SPs may have extra requirements and request additional information.

Once a gateway is in operation, for all jobs run on XSEDE a PI must log and retain for at least 90 days the:

- End user's IP address
- Date and timestamp
- End user's username on the gateway

Authentication and access control is important for any account, and gateways differ from traditional single-user accounts in key ways. While gateway accounts are the responsibility of the requesting PI, these accounts are shared and do not map one-to-one with a single real person. Therefore, the XSEDE short-lived CA, following IGTF policies, does not issue end entity certificates for these accounts. Instead, any valid host certificate recognized by XSEDE will work for these accounts.

Related to this there are some guidelines regarding gateway and certificate usage.

- A gateway should be run on a dedicated host (bare metal or VM), not to be shared with other services. This is to protect the potentially long-lived private key material on the gateway.
- A gateway should have an identified party with root access to the gateway, and this person's contact information needs to be kept up-to-date to remain.
- A single gateway account should not to be used for multiple gateways. This is to reduce exposure of any private keys and limits the scope of potential compromise.
- Any host certificates used for authentication should match the submitted host names for the gateway server or one of the test servers.
- A gateway should never transfer private key material over a network.

Additional recommendations on secure practices for gateway developers are provided on the gateway link from http://www.xsede.org and in [13].

C.4. Additional Security Considerations

Community user logins present special security considerations and thus dictate the need for this policy.

The primary requirement from the security working group is to ensure that SPs are notified of the creation of such community user logins and that the information requested in the section above is provided.

Potential security risks have been identified in at least two Teragrid-funded studies. A 2008 Risk and Vulnerability Assessment conducted by the security-wg [7] identifies the following risks to the XSEDE (then TeraGrid) from Science Gateways:

- Because all jobs from the gateway are submitted as the same user, if there is a security problem, XSEDE cannot tell which Gateway User committed the offense (This is why gateways are required to maintain audit logs, mapping job submission to specific people.)
- If a gateway becomes compromised, the Gateway User could get a shell and attempt internal attacks on XSEDE. (This can be mitigated by using ssh keys with a restricted set of executables in the *authorized_keys* file options [9].)
- If gateways allow their users to upload unvetted code, it may be malicious or run poorly and consume unneeded resources. (This can be mitigated on the Gateway end by only allowing certain tasks to be performed.)
- A compromised user's workstation could be used by an attacker to log into the Gateway impersonating the user. (Damage is constrained by the types of jobs users can submit through the gateway.)

A 2010 paper co-authored by Teragrid security and gateway staff [8] also identifies compromise of the science gateway server, unauthorized use of the science gateway community account, unauthorized command and/or application execution on an SP resource by the community account and unauthorized replacement or update of gateway executables on the SP resource. In addition community account authentication methods (passwords, PKI credentials, SSH keys) may be compromised. The impact of credential compromise will vary based on the type of credential, its lifetime, and its corresponding capabilities.

C.4.1. Resources Required

Information entered on the form used to request a community account is also logged in the XCDB. If SPs choose to impose additional account restrictions, effort may be required to support sudo or GSI-SSH. Gateway developers may see an increased level of effort if there are restrictions where there were none previously, but if sites adopt the recommendations and provide a more standardized set of restrictions development efforts would be greatly reduced for developers using multiple SPs.

C.4.2. Scaling

This policy will provide a consistent means for identifying community users and thus should have a positive (though minor) impact on scaling. This policy provides recommendations for account restrictions at SP sites and so can improve the usage environment for gateway developers.

C.5. References

[1] https://www.xsede.org/web/guest/allocation-policies

[2] Science Gateways program description, http://www.xsede.org/

[3] The Globus Project, http://www.globus.org/

[4] UNICORE (Uniform Interface to Computing Resources), http://www.unicore.eu/

[5] Genesis II, http://www.cs.virginia.edu/~vcgr/wiki/index.php/The_Genesis_II_Project.

[6] Acceptable Use Policy, https://www.xsede.org/usage-policies

[7] TeraGrid Gateways Risk Assessment, James Rome, Internal TeraGrid publication, 2008.

[8] Hazlewood, V., Woitaszek, M. 2011. Securing Science Gateways. *TeraGrid 2011 Conference, Proceedings of the ACM/IEEE*. July 2011.

[9] OpenSSH SSHD man page, http://www.openbsd.org/cgi-bin/man.cgi?query=sshd&sektion=8

[10] Wikipedia: Restricted Shell, http://en.wikipedia.org/wiki/Restricted_shell

[11] Science Gateways Overview, http://www.xsede.org/gateways/

[12] XSEDE MyProxy Delegation Service, https://oa4mp.xsede.org/oauth/

[13] Science Gateway Security Recommendations, http://www.sciencegatewaysecurity.org/news/sgiw2013