Fall 12

XSEDE Security Risk Assessment

Prepared by Adam Slagell, NCSA

This is the executive summary of the security risk assessment performed during the first year of the XSEDE project as prepared by Adam Slagell with the review of the XSEDE risk assessment team and security operations group.

Introduction

Purpose

At a strategic level, this risk assessment forms the foundation on which to build a more formal, risk-based security program for XSEDE. This risk assessment will inform our security plans, i.e., security policies, procedures and blueprints.

Tactically, this risk assessment serves two purposes. First, it helps the security operations team prioritize its work for the next year or two by recommending specific countermeasures to implement to address the risks identified. It further prioritizes these actions based on the severity of the risks they address. Second, this risk assessment serves as a template for other XSEDE service providers (SPs) who are encouraged to follow a similar process at their specific institutions. This is recommended as a complete picture of risk for XSEDE depends both on the high-level federation issues, and the SPs on which much of its infrastructure and services are built.

The purpose and context into which this risk assessment fits is more fully described in *Developing a Risk-based Formal Security Program for XSEDE*.[1]

Scope

A project sizing[2] was completed to prior to beginning a risk assessment for XSEDE. This was critical to scope the risk assessment to the proper level for both the resources available to the team and time given for completion. It is always possible to drill deeper and deeper with a risk assessment, but at the risk of never finishing it without clear scope.

A risk assessment can be scoped both in terms of threats and assets considered. Congruent with our stated goals, we performed a risk assessment at the federation-level and are asking that individual Service Providers (SPs) follow a similar process for their local XSEDE resources. Therefore, threats that do not affect more than a single site, the fundamental shared resources of XSEDE or its underlying trust fabric, are not considered. For example, threats specific to the architecture of a specific HPC resource would not be considered at this level, but threats to an XSEDE shared authentication system (e.g., our Kerberos realm) would be.

The most general kind of risk assessment would consider all kinds of threats, including environmental, application error, physical failure and natural disasters. However, a primary goal for this risk assessment is to more wisely choose *cyber* security controls to give us the most return on our investment of resources. Therefore, we primarily considered application errors and insider/outsider threats (e.g., i.e. hacking, cracking and attacks). So while hackers stealing credentials were considered, fire damage at a particular site's data center were not. The latter would more appropriately be considered by a site-specific risk assessment.

---

[1] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Developing+a+risk-based+formal+security+program+for+XSEDE

[2] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Risk+Assessment+Project+Sizing

Risk Assessment Approach

The Charge

Senior management, including the project PI John Towns, approved the approach[3] and scope[4] of this risk assessment prior to any work beginning. While a risk assessment was promised in the XSEDE proposal, it was still necessary to get management buy-in to the process as it is ultimately up to management to decide how to act upon the recommendations made in this report. Furthermore, it was important to get the support of management as this risk assessment required resources and cooperation from many people across many groups in addition to the formal members of the risk assessment team.

The Team

Adam Slagell, a security analyst and the CISO at the NCSA, led the risk assessment team. This team was composed of a representative from each division of Operations: Security, Data Services, XSEDENet, Accounting, and System Operations Support. Jim Basney, Gary Rogers, Benny Sparks, Amy Schuele, Anjana Kar (and later Derek Simmel) fulfilled these roles, respectively. The level 3 Security manager, Randy Butler, and the Operations level 2 manager, Victor Hazlewood, provided additional support and advisement. Additionally, the XSEDE security operations team was invited to provide feedback at each stage of the process.

NIST 800-30

There is literally an alphabet soup of risk management approaches (e.g., NIST 800-30, OCTAVE, AS/NZS 4360:2004, ISO 31000 series, etc.). We chose NIST 800-30[5] for several reasons. First, Carnegie-Mellon's OCTAVE builds upon the NIST process, Australia and New Zealand's process upon OCTAVE, and ISO 31000 upon AS/NZ 4360:2004. Because of this, and the relative simplicity of the NIST process, it is often recommended that organizations start their first risk assessment with the NIST process. If it is later deemed insufficient, effort is not wasted as follow-up reassessments can build upon what was already done.

Secondly, these other approaches aren't as well suited to the geography and resources of XSEDE. For example, the OCTAVE approach relies heavily upon a series of self-directed workshops with management, operations, security and business heads walking through several scenarios, questionnaires and checklists. This is a significant commitment from the whole organization that XSEDE would not likely have made. And even if it would have, this approach is difficult to do with teams so geographically distributed.

For these reasons, we followed the NIST 800-30 process of a qualitative risk assessment which ranked risks in two dimensions according to a high, medium or low ranking (defined later). The first dimension was likelihood of a risk being realized, the second dimension the impact of a risk if it is realized. Combining the scores in those dimensions then allows one to create an overall risk rating to prioritize activities.

---

[3] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Developing+a+risk-based+formal+security+program+for+XSEDE

[4] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Risk+Assessment+Project+Sizing

[5] http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf

The over all process of a NIST 800-30 risk assessment is shown in the figure below.

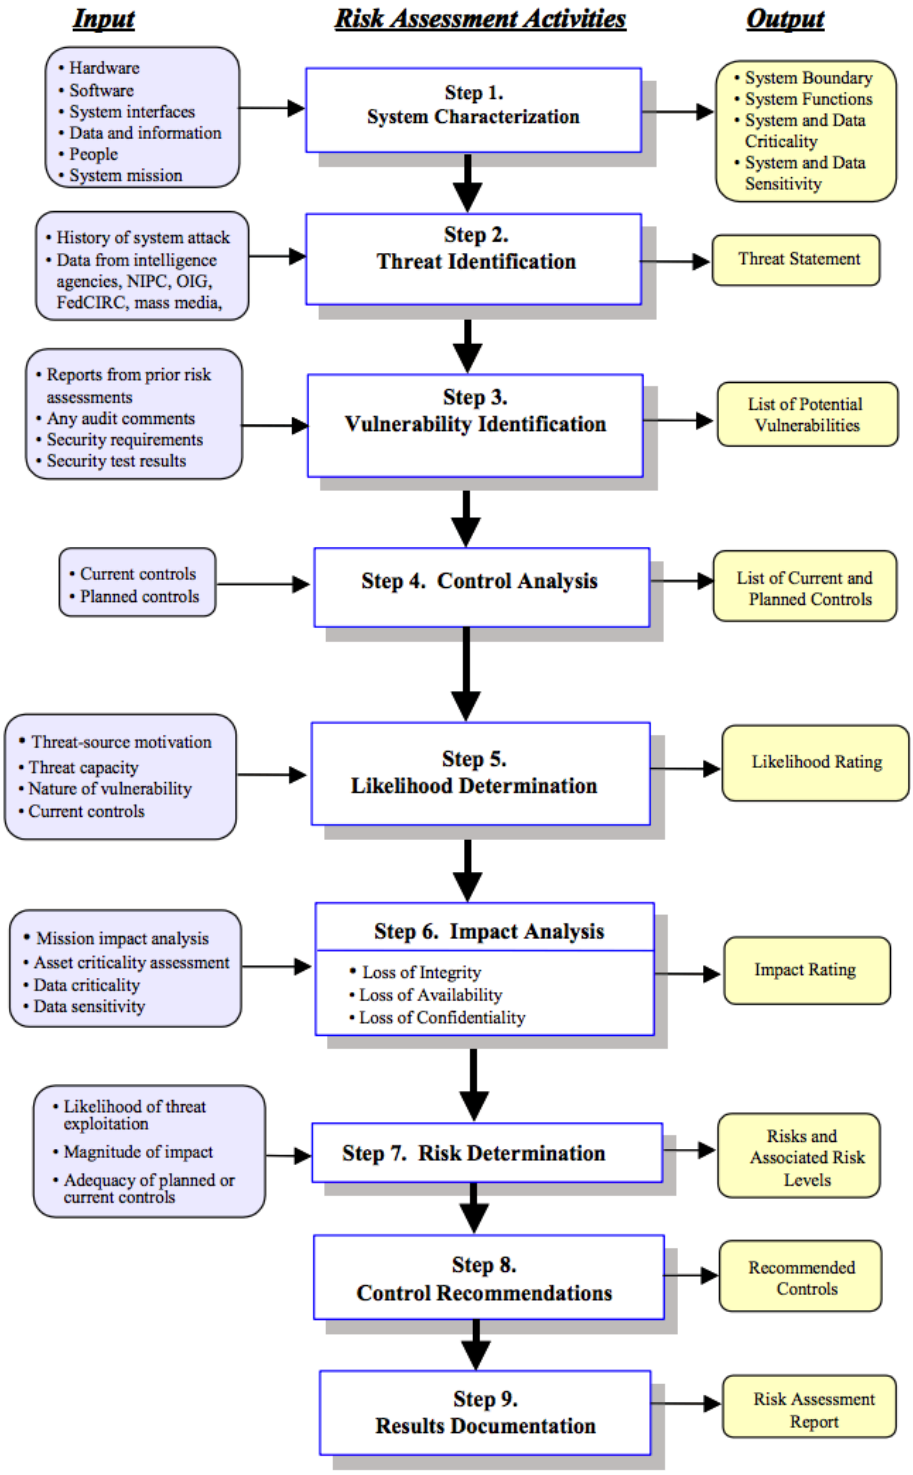

Figure 1: NIST 800-30 process

Vulnerability Analysis Process

A vulnerability is a flaw or weakness in a system's security procedures, design, implementation or internal controls that could potentially be exercised by a threat agent to result in a breach or violation of the system's security policy. In this phase of the risk assessment we identified as many federation-level vulnerabilities in XSEDE as we could, mapping each vulnerability to a threat action and threat source.

We developed this list by thoroughly investigating each asset from the System Characterization[6] (detailed in a later section), reviewing other risk assessments for generic vulnerabilities, and conducting interviews with people from each operations domain. For each vulnerability, we identified a potential threat source from the XSEDE Threat Profile[7] (detailed in a later section) and a threat action that could be taken to exploit the vulnerability. Finally, this was all compiled into a table[8] reviewed by the operations security and risk assessment teams for completeness and accuracy.

Control Analysis Process

Security controls are mechanisms in place to mitigate the risk of threats being realized and hence exploiting vulnerabilities in your infrastructure. Controls can be *administrative* (e.g., policies, standards, guidelines, training and other processes), *technical/logical* (e.g., authentication and authorization systems, file permissions, firewalls, intrusion detection systems, etc.), or *physical* (e.g., locked file cabinets, secured data centers, cameras, fences, etc.).

Since risks are addressed by controls, it is important to understand the security controls already in place or planned. Without this step, you don't know where there are gaps. One of the outputs of the risk assessment is to recommend changes and additions to controls to address the highest impact and most probable risks.

The Control Analysis Matrix[9] lists all the existing and planned controls with current status and notes. The sources for the list of these controls came from brainstorming, reference to the security controls catalog in NIST Special Publication 800-53[10], and a review of previously identified controls in the XSEDE System Characterization.

Risk Likelihood Evaluation Process

This phase of the assessment required the risk assessment team to rank the likelihood of each vulnerability being exploited as high, medium or low. In this ranking, three things are considered:

- Motivation and capabilities of the threat source
- Specifics of the vulnerability
- Effectiveness of current controls to mitigate associated risks

The following definitions for high, medium and low were used:

---

6 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+System+Characterization

7 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Threat+Profile

8 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Vulnerability+Identification

9 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Control+Analysis+Matrix

10 http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf

- **High:** The threat-source is highly motivated and sufficiently capable, and controls to prevent the vulnerability from being exercised are ineffective.
- **Medium:** The threat-source is motivated and capable, but controls are in place that may impede successful exercise of the vulnerability.
- **Low:** The threat-source lacks motivation or capability, or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised.

To come to consensus on the rank for each vulnerability we used the following process. As a first round, we used blinded (to all but the risk assessment lead) votes via email for each item. As long as 2/3 of the team voted on an item and the votes had a standard deviation less than 0.7, we took the average of the votes and chose the closest probability ranking. To compute a numerical average, each ranking received a numerical value of 0, 1 or 2, and the average was simply the arithmetic mean.

For other items, there was either widespread disagreement or a lack of understanding to gain quorum. These were resolved over a phone meeting. If quorum was the problem and the team lead could not get enough votes, then a threat scenario would need to have been written up on the issue by the most expert person of the team for that item (this did not happen as we always had quorum). Each contentious item was discussed on a phone meeting with quorum until a majority decision was made. The call started with a reminder of what the different rankings of high, medium and low mean, and the chair's synopsis of the disputed items and why he voted how he did. Detailed results of the individual votes and meeting notes can be found on the wiki[11].

Impact Analysis Process

After determining the likelihood of each identified vulnerability being exploited, we then focused on the impact if it were actually exploited. The impact of any exploit is going to depend upon (1) the mission of XSEDE, (2) the criticality of the vulnerable system or data to XSEDE, and (3) the sensitivity of the affected system or data. This information is usually found in a mission or business impact analysis, but XSEDE has no such documentation. However, there is a mission statement noted in the XSEDE System Characterization[12], and there is a categorization of XSEDE services into tiers in the XSEDE services master spreadsheet. In cases where it was not clear how critical a system is, we relied upon the specific system owner or maintainer's advice.

Impact from a security incident could affect the integrity, availability or confidentiality of a system or data. Depending on the subsystem affected, we could be concerned more with one kind of impact than another. Particularly, in light of XSEDE's mission, it would most often be that integrity and availability are of more concern than confidentiality. However, the impact on each of these three properties should be considered for any potential exploit. If the impact were unclear to the general risk assessment team, we would have asked the system/data owner to write a brief impact analysis for the particular vulnerability being exploited. However, that was not necessary for our disputed items.

The following definitions for high, medium and low were used:

---

11 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+Federation+Risk+Assessment#section-XSEDE+Federation+Risk+Assessment-LikelihoodDetermination

12 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+System+Characterization

- **High:** Exercise of the vulnerability (1) may result in the highly costly loss of major tangible assets or resources; (2) may significantly violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human death or serious injury.
- **Medium:** Exercise of the vulnerability (1) may result in the costly loss of tangible assets or resources; (2) may violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human injury.
- **Low:** Exercise of the vulnerability (1) may result in the loss of some tangible assets or resources or (2) may noticeably affect an organization's mission, reputation, or interest.

We followed the same process of semi-blinded voting on impact as in the previous phase of likelihood determination. The difference is that contentious items could potentially have been resolved with an impact analysis. The system/data owner would have written a short impact analysis, less than one page, and the risk assessment team would have revoted after reading it and applying the definitions above. However, since most of the vulnerabilities are discussed in generalities and we are doing a qualitative risk assessment, this was not required for any of our risks. Detailed results of individual votes can be found on the staff wiki[13].

Risk Determination Formula

With likelihood and impact analysis complete, we calculated numeric values for each risk to rank them. We used the table below, which gives a weight for both likelihood and impact.

	Low Impact (10)	**Medium Impact (50)**	**High Impact (100)**
High Likelihood (1.0)	Low 10 x 1.0 = 10	Medium 50 x 1.0 = 50	High 100 x 1.0 = 100
Med. Likelihood (0.5)	Low 10 x 0.5 = 5	Medium 50 x 0.5= 25	Medium 100 x 0.5 = 50
Low Likelihood (0.1)	Low 10 x 0.1 = 1	Low 50 x 0.1 = 5	Low 100 x 0.1 = 10

13 https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+Federation+Risk+Assessment#section-XSEDE+Federation+Risk+Assessment-ImpactAnalysis

System Characterization

The first official step of this risk assessment was performing a full system characterization, whose details can be found on the XSEDE staff wiki[14]. While the project sizing focused us on the kinds of threats and assets to be considered, the system characterization specifically enumerated these pieces and system boundaries. This phase further focused our risk assessment and helped to systematically breakup the larger system into logical areas that could be delegated to appropriate people for further evaluation and information gathering. It has also provided a nice single point of reference for anyone wanting to know more about a particular piece of XSEDE.

Our system characterization identified the boundaries of XSEDE (as opposed to site-specific resources suited to a lower-level risk assessment) along with the resources, organizations and information that constitute XSEDE. This included identifying the system's mission; major hardware components; key software and services; data and information with sensitivity assessment; user and support communities (includes XSEDE staffing groups as well as resource providers); logical network topology; system interfaces (both internal and external connectivity); flow of information (particularly if sensitive); references to existing security architecture and policy documentation; and management and monitoring controls (e.g., availability monitoring, intrusion detection systems).

Two subsections of the system characterization deserve further description as they required all members of the risk assessment team to gather information and interview domain experts beyond the simple review of existing XSEDE documentation: the Systems & Services section and the XSEDE Software Stack section.

Systems & Services

For each system or service we created a separate subpage of the System Characterization linked to from a common table. These subpages provided the following for each item:

- Summary description
- Location and administrator of the system or service
- Data types involved and whether or not any data was sensitive
- Flow of this data between systems
- System interfaces and protocols used to communicate with the system/service

We created such pages, interviewing the appropriate administrators, for the following services—a list which comes primarily from XSEDE Master Services spreadsheet maintained by Stephen McNally.

[14] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/XSEDE+System+Characterization

- Level 1 SP Compute & Storage Resources
- User portal, knowledgebase, web site & wiki
- Ticket system
- XSEDE central database
- Resource description repository
- AMIE
- XSEDE.org DNS
- XSEDE Kerberos realm
- Certificate authorities
- InCommon Portal
- OAuth MyProxy service
- Security wiki
- Security Jabber server
- Mail list server
- XUP file system services
- Conference call system
- POPS
- DOORS
- SharePoint
- CI Tutor
- Virtual Workshop
- Source repository
- XSEDE Bugzilla
- perfSONAR
- Integrated Information Services (IIS)
- User profile service
- Metrics for IIS
- INCA
- Speedpage
- Karnak Predictor
- Globus Listener
- Build & Test system
- GlobusOnline

XSEDE Software Stack

Here we list key pieces of software, such as toolkits and services, run at different SPs. Local customizations, particulars of base OSs and other software run at the individual sites were not included but belong in individual, site-specific risk assessments. We focused on software and configurations (e.g., the XSEDE trust store of certificates and Common User Environment) configured in XSEDE specific ways to support the services layer described in the XSEDE architecture. Much of this content is derived from the XSEDE Production Baseline for Service Provider Software and Services document.

- Registration Services
- INCA
- AMIE
- Globus Toolkit
- XSEDE CA Tarball
- UberFTP
- Globus-wsrf
- XSEDE Glue2 Publishing
- Local Resource Management System
- Gx-map
- Tgusage
- Modules
- Tgproxy
- Common User Environment
- Tginfo
- Tgresid
- SAGA
- MCP
- GUR
- CommSH

Threat Profile

A threat is the potential for a particular actor to exploit a particular vulnerability towards a malicious end. If there is no vulnerability, there is no threat. Furthermore, a threat is not the person or event that triggers a threat, that is the threat source/actor (e.g., a hacker, inept programmer, flood, etc.).

A threat profile is usually performed before a vulnerability assessment. Even though one may not know specifically if they have vulnerabilities for the threats to exploit, they can state what are the expected common threat actors as well as what they are not going to protect against or address within the scope of an assessment. For XSEDE, we are ignoring natural and environmental threats, which would need to be determined separately one level lower by the major service providers. Instead, we focused on human threats, both intentional (e.g., direct attack or theft) and unintentional (e.g., poorly written software). Furthermore, we only considered threats to shared XSEDE infrastructure, XSEDE's trust fabric and the overall stability of the federation as noted in the risk assessment's scope.

In our XSEDE threat profile we identified potential threat sources and made an estimation of their motivations, resources and capabilities. To gather this information, we created a large table with rows for potential threats, and columns for the corresponding threat sources, motivations and potential threat actions (e.g., credential harvesting, privilege escalation, social engineering, criminal activities, data theft, sabotage).

We used two methods to gather information. First, we used threat profile examples from other projects and systems to capture the generic sorts of threats that affect almost any system on the Internet. Then to get more specific we created a survey, which included questions on past security incidents, and had the security leads from each level 1 SP fill it out. This not only allowed us to come up with a list of threats, but to generally rank the likelihood of them. These surveys, the results, the full table, the external sources, and full threat statement can be found on full threat profile wiki page[15]. The summary follows.

Summary Threat Statement

The threats to XSEDE largely do not change from the threats to TeraGrid, besides some increased risk of credential harvesting with many additional level 3 service providers expected to join. Therefore, our historical data can intelligently inform our threat profile for XSEDE. We expect the most common problem to continue to be the cycle of credential harvesting, followed by privilege escalation, then Trojaning, and finally more credential harvesting. This threat is unlikely to change without major changes to authentication mechanisms and/or user interfaces. We will continue to see many, mostly unsuccessful dictionary-based attacks as well as bot herding for criminal activities. We anticipate a potentially new threat action of Bitcoin-mining and must be on the lookout for that. Other than that, our biggest threats are misconfigurations or software errors that must be addressed through configuration management and good software engineering practices. Fortunately we avoid some of the threats seen against other sites by the nature of the customers we serve. However, if that ever changes and we start serving more private industry with valuable intellectual property, governments with classified data or researchers working on politically hot topics, we will have to reevaluate our threat profile.

[15] https://www.xsede.org/web/staff/staff-wiki/-/wiki/Main/Threat+Profile

Risk Assessment Results

Risk #1 Total Risk: 5 Likelihood: M Impact: L

Vulnerability
Shared accounts make auditing difficult.

Potential Threat Source
Malicious user or cracker

Potential Threat Actions
Multiple XSEDE services (e.g., CVS, science gateways, and DOORS) all use shared accounts. Incident response involving any of those accounts either requires the assistance of a third party or cannot distinguish between actual people, thus allowing attackers to better obfuscate their identity.

Existing Controls
For science gateways there is policy and procedure before they are approved, and the gateways must retain records to allow actions to be mapped to a user. Also, the user actions are restricted, and they do not get a full shell. For other accounts there are no current controls.

Recommended Control Changes
Where possible, remove shared accounts (e.g., CVS, DOORS, etc.), and put into place mechanisms that don't require a shared password among several people. Additionally, extra auditing mechanisms should be applied in cases where this cannot be avoided.

Risk #2 Total Risk: 5 Likelihood: M Impact: L

Vulnerability
Inconsistent authentication standards make XSEDE authentication only as strong as the weakest SP.

Potential Threat Source
Crackers

Potential Threat Actions
Through use of x.509 certificates, users have single-sign-on capabilities to move between level 1 SP resources. However, the requirements for authentication mechanisms, identity vetting and operating a CA (certificate authority) between SPs are different, making the whole system only as strong as the weakest link. For example, some CAs aren't even IGTF accredited, and some sites accept public keys without restriction. Therefore, an attacker might exploit a weakness at one SP and leverage that to access other XSEDE resources to spread an attack.

Existing Controls
Future CAs added to the tarball must be IGTF accredited. Certificates are also short-lived, limiting the exposure window to some extent.

Recommended Control Changes
We recommend creating a standard for authentication mechanisms that could be used with a CA in the XSEDE tarball. Non-IGTF accredited CAs that are not used by many should also be removed from the tarball, though it is likely too disruptive to remove any critical non-accredited CAs. Hopefully, deployment of an XSEDE CA service will allow more of those CAs to retire.

Risk #3 Total Risk: 1 Likelihood: L Impact: L

Vulnerability
Some XSEDE services use their own authentication system, which is not subject to the same security requirements or account deactivation process.

Potential Threat Source
Disgruntled former employee/user or crackers

Potential Threat Actions
Several services (e.g., Sharepoint, DOORS, CI Tutor) use one-off authentication systems separate from the XSEDE Kerberos and CA systems. Users could exploit the fact that disabling an XSEDE Kerberos account does not deactivate these, and adversaries could exploit unknown vulnerabilities in these outside authentication systems if users synchronize their passwords for the portal and these other accounts.

Existing Controls
There are no existing controls to mitigate this risk.

Recommended Control Changes
We should identify and remove local or non-XSEDE authentication systems for XSEDE services and replace them with XSEDE Kerberos or GSI where possible.

Risk #4 Total Risk: 5 Likelihood: M Impact: L

Vulnerability
Anyone can create an XSEDE account, and identity vetting is dependent upon trusting email, delegations to PIs, and self-asserted attributes.

Potential Threat Source
Crackers or malicious users

Potential Threat Actions
Having a lower level of assurance in the user's identity, there is potential for a person to impersonate another or supply false contact information. This makes it difficult to always

hold a real person accountable for the malicious actions on the system or theft of intellectual property.

Existing Controls

The process has been formalized enough that TAGPMA is satisfied with it. Emails can only be used once, and it is expected that users would notice never getting an account activation email. And while the PI is not vetting to a high level of assurance, the PI is likely to notice a problem in account creation when her researcher cannot login later.

Recommended Control Changes

We recommend no changes at this time. While a more rigorous process could be developed for account creation and vetting, it would become much more manual and time intensive. Since this is a very low risk, such a decrease in usability is not justified.

Risk #5 Total Risk: 5 Likelihood: M Impact: L

Vulnerability

Plaintext authentication is used for mail list management.

Potential Threat Source

Crackers

Potential Threat Actions

Majordomo email lists are managed by plaintext passwords in the emails and trust the source address for authentication. This opens up the possibility for eavesdropping, list denial of service, and list hijacking for spam.

Existing Controls

There are spam filters to prevent some types of abuse of email lists. Security personnel use PGP encryption for confidential email to thwart eavesdropping. Finally, staff can use the wiki and just send URLs in the email for sensitive documents.

Recommended Control Changes

We recommend no changes at this time. The risk is quite low and does not justify the effort to change email list services or create a new interface for management of lists. Security training should remind users not to send sensitive documents over email, to a list or otherwise.

Risk #6 Total Risk: 1 Likelihood: L Impact: L

Vulnerability

The conference call system uses weak or loosely managed PINs.

Potential Threat Source

Crackers

Potential Threat Actions

The PINs used for AT Conference screen sharing are weak and highly predictable 4 digit PINs. The phone line PINs are longer and random, but they are widely shared on email lists and reused. They also are not rotated. This opens XSEDE up to the threat of eavesdropping on private meetings.

Existing Controls

The incident response team uses its own number that is a closely guarded secret. There are no controls other than a little security through obscurity for the other PINs.

Recommended Control Changes

We recommend no changes at this time. While a system could be developed to securely generate, delivery and rotate these PINs, a more onerous process could be very disruptive to meetings for some time. Since this risk is only ranked level 1, it is recommended to simply accept it at this time.

Risk #7 Total Risk: 5 Likelihood: M Impact: L

Vulnerability

Any adversary able to intercept or monitor emails can utilize self-service password resets to their advantage.

Potential Threat Source

Crackers

Potential Threat Actions

An attacker could try to reset a user's account from the portal, capturing the email with the reset code. Then they could log onto resources that they do not have access to normally and/or act as another person maliciously.

Existing Controls

Users would be unable to login themselves if their passwords were changed. An adversary simply able to read the message would not be able to prevent the user from seeing the reset email and being alerted to the issue.

Recommended Control Changes

Completely addressing the problem would make it more difficult to serve users who forget their passwords, a common occurrence. Instead, we should mitigate by simply requiring a secondary email address to be registered with the portal. Therefore, an attacker would have to stop 2 different emails to different accounts to prevent the user from seeing them.

Risk #8 Total Risk: 1 Likelihood: L Impact: L

Vulnerability
User credentials are not always encrypted on disk.

Potential Threat Source
Crackers or curious users

Potential Threat Actions
Kerberos tickets and proxy certificate keys are usually unencrypted on disk on the HPCs. Any one able to steal these could masquerade as another user.

Existing Controls
File permissions and short life spans are the primary mitigation against credential abuse in this case. On most systems, users are also alerted to their previous login details to help them identify someone using their credentials illicitly. Gateways such as Globus Online, which store many more credentials at a given time, have more protections in place.

Recommended Control Changes
We recommend no changes at this time. While shortening ticket or certificate lifetime could reduce this risk, it would decrease usability even more. It is currently balanced so that shortening it more provides diminishing returns for this very low risk.

Risk #9 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
Some services use unencrypted private keys to access XSEDE resources.

Potential Threat Source
Crackers or curious users

Potential Threat Actions
Several services run on the backend transferring data to/from the XDCDB, IIS, RDR, IIS metrics, etc. These often depend on permanent SSH keys or GSI certificates to run in automated scripts. Depending upon the lifetime of these credentials and the restrictions on the corresponding accounts, an attacker who compromises one of these systems could steal credentials to leverage elsewhere in an attack on the XSEDE infrastructure.

Existing Controls
Files system permissions and some security through obscurity are the main existing controls. Some services utilize a trusted proxy renewal service so that short-lived certificates can still be used in these cases, limiting the abuse of a one-time exposure.

Recommended Control Changes
Inventory how system accounts are used for various services and identify those using unencrypted certificates or SSH keys. Protect the credential as best as possible, limit the ca-

pabilities of the credential and corresponding account to only those needed, and monitor for any unexpected use of such accounts.

Risk #10 Total Risk: 5 Likelihood: M Impact: L

Vulnerability
Users control their keys and may not protect them adequately.

Potential Threat Source
Crackers

Potential Threat Actions
XSEDE does not control how users protect SSH private keys or keys corresponding to their X.509 certificates on their own systems. Keys may be unencrypted, or encrypted with poor passphrases. This means a compromise on a user's system could allow an attacker to steal their credentials to logon to an XSEDE resource and spread their attack.

Existing Controls
Automatically created proxy credentials are protected by file permissions, and default umasks, and home directory permissions protect user SSH keys installed onto XSEDE systems. However, nothing protects keys on user systems. Therefore, some SPs do not allow users to login with SSH keys. Short-lived certificates mitigate these problems somewhat by constraining the window of exposure.

Recommended Control Changes
Education can mitigate this risk some, and therefore we recommend putting something about protecting your credentials into the training. Really addressing the issue would mean disabling authentication mechanisms, or taking management of these credentials out of the users' hands. That is impractical for such a low risk.

Risk #11 Total Risk: 1 Likelihood: L Impact: L

Vulnerability
XSEDE lacks a centralized logging infrastructure.

Potential Threat Source
Crackers

Potential Threat Actions
There is no centralized logging for XSEDE services. This makes it easier for an attacker to erase his digital trail. It also makes it more difficult for incident response teams to investigate a complex, cross-site attack.

Existing Controls
Most individual sites do backup and logging, but this is not consistent across every service. XSEDE mitigates some of these challenges by trying to keep a tight collaboration

between the incident response teams at various SPs, actively funding security at these different sites.

Recommended Control Changes

If XSEDE provided a centralized syslog service, that would address this risk and potentially make it easier to detect attacks. However, that costs resources to run a new service, set it up, and provide for hardware. Also, it is difficult politically since institutional policy may prevent sharing of log data. The only recommended action for this very low risk is policy in the baseline documents that requires XSEDE services to send their logs to an log server as well.

Risk #12 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

DNS system lacks authentication for response and synchronization.

Potential Threat Source

Crackers

Potential Threat Actions

DNSSEC is not used for clients and inconsistently used for server synchronization. This lack of authentication could be exploited to give false responses or poison servers. Besides DoS attacks, this could allow man-in-the-middle attacks for non-certificate based services like regular SSHD and default Globus configurations that rely on DNS for hostname canonicalization.

Existing Controls

Syncs are done over XSEDENet, which is somewhat private, and there is a hidden master DNS server. Given that syncs are only allowed between certain IPs, that reduces the risk of some attacks that could corrupt the DNS tables and lead to denial of service. For services that require certificates, we control the trust store of XSEDE certificates and it would be difficult for a man-in-the-middle attack to succeed in those instances. However, some services, like SSHD, would give no indication of false DNS responses being used except for a change in the public key fingerprint.

Recommended Control Changes

Some subzones, such as xsede.psc.edu, already use DNSSEC. We recommend implementing this for all subdomains and the xsede.org root.

Risk #13 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

XSEDE hardening guidelines for SPs are optional and unaudited lending to the possibility weaker security at some SPs.

Potential Threat Source
Crackers

Potential Threat Actions
XSEDE has inherited the optional guidelines for system hardening from TeraGrid, but these are not enforced nor audited for compliance. Because of the shared trust fabric, XSEDE is only as strong as its weakest link, and an attacker could exploit a more lenient security posture at one SP to spread an attack on XSEDE.

Existing Controls
Some existing guidelines exist, but they are out-of-date and not widely followed or known.

Recommended Control Changes
Policies must be updated and baseline security plans developed for critical services. SPs must also be trained to be aware of these policies and their importance. Finally, regular auditing must be done to ensure compliance, and regular vulnerability scanning will help with that.

Risk #14 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
XSEDE has inconsistent or non-existent backup processes for key resources.

Potential Threat Source
Accident or incompetent staff

Potential Threat Actions
There are many services and systems distributed across XSEDE, but there is no centralized backup or backup policies for critical resources. The hosting SP determines what if anything is backed up. Equipment failure or a major security incident could make it difficult to bring these systems back online in a secure state (especially since not all sites test their restore processes).

Existing Controls
Most services are redundant across more than one SP. This limits the impact of any one failure.

Recommended Control Changes
It would be expensive and politically challenging to create a centralized backup service for XSEDE. Instead, we recommend that new policies and baselines require important services to be backed up regularly and restore procedures tested.

Risk #15 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
XSEDE does no centralized security monitoring.

Potential Threat Source
Crackers

Potential Threat Actions
XSEDE networks and many publicly facing systems do not utilize any sort of intrusion detection systems and could be compromised with delayed notice. Defacement or disruption of the portal, XSEDE's public face, would be potentially damaging, and lack of monitoring increases the exposure time during incidents. Given the large bandwidth and capabilities of HPCs and perfSonar nodes, a sizable amount of damage from a DoS attack using XSEDE resources could occur in even a short amount of time.

Existing Controls
Most SPs do some network monitoring, and there is considerable Bro expertise within the community that can be leveraged by other SPs. But not all are equal, and not all have a robust monitoring infrastructure. There is monitoring of intelligence channels that would help detect some attacks originating from XSEDE, but that is not a replacement for an IDS.

Recommended Control Changes
Due to the high bandwidth pipes of XSEDENet, it is prohibitively expensive to instrument it with a full IDS. Even filtering out gridFTP traffic, you still need very expensive networking hardware to passively tap this infrastructure. Add to that the cost of expertise, and we can see why SPs with smaller security groups do not do this. Even if all SPs had an IDS and we wanted to correlate alerts, many legal and political issues prevent easy, automated sharing of such data. Therefore, we recommend accepting this risk.

Risk #16 Total Risk: 50 Likelihood: H Impact: M

Vulnerability
Admins are not required to use strong authentication for management of XSEDE services.

Potential Threat Source
Cracker

Potential Threat Actions
Best practice is to use strong, i.e. two-factor, authentication for administrators given the ease with which passwords can be harvested. XSEDE has no central OTP service and different service providers have different rules about authentication for administrators. Attackers could take advantage of this and target sites with weaker policies to gain a foothold.

Existing Controls

Since there are no enforced baselines for securing XSEDE resources, there is really no consistent way administration is handled. Some sites do use two-factor as well as additional controls to protect these administrative interfaces, but not consistently. Therefore, this is a largely an unmitigated risk.

Recommended Control Changes

We recommend a new policy, also part of the baseline security requirements, to require two factor authentication before escalation to privileged accounts or logging into administrative interfaces. This is one of the first changes we would like to make and verify. It is recognized that there may be edge cases where this is not possible, and so the policy must have in place a way to handle exceptions.

Risk #17 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

A critical service is vulnerable to a denial of service attack.

Potential Threat Source

Crackers

Potential Threat Actions

Any Internet facing system could be vulnerable to a DoS attack, given sufficient adversarial resources. Even though XSEDE tries to mitigate this by having replication across multiple sites and redundant network paths, it is still possible for an adversary to mount such an attack, especially against non-replicated services like the XUP and POPS.

Existing Controls

Replication in most cases mitigates the risk of a service going down because a particular SP is targeted in a DoS, but it does much less if someone is purposefully targeting XSEDE as they could target all replicas. Large network pipes help to keep floods from bringing down the network, but one does not have to saturate the network to bring down something like a web server. Virtualization and good backups, where used, help with a quicker recovery. Finally, some services use Amazon's EC2 for hosting which has its own DoS protections.

Recommended Control Changes

Services without offsite replication should be identified. Then those should be replicated in order of criticality. This helps protect against more threats than just targeted DoS.

Risk #18 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

There is no consistent patch management process for all XSEDE services and systems.

Potential Threat Source
Crackers

Potential Threat Actions
Patch management is ad hoc, usually handled by individual SPs hosting services. There is no one monitoring for vulnerabilities in the various services or their dependencies, and McNally's group relies upon notifications about new software vulnerabilities from the security operations team. Some sites have practices and regular patch schedules like SDSC, which are applied to AMIE and the XDCDB servers quarterly. There are no security baselines detailed for each of these services. All of this taken together increases the probability of a vulnerability exposure, which an adversary could take advantage of. This brings home the need for some sort of vulnerability management solution like Qualys.

Existing Controls
Members of the XSEDE incident response team do monitor Bugtraq and other channels for news of new vulnerabilities and active exploits. When they are aware that XSEDE is using such software, they do alert XSEDE system administrators and recommend actions.

Recommended Control Changes
We recommend regular scanning and auditing of XSEDE resources and support systems to ensure systems are patched against vulnerabilities deemed significant threats to XSEDE. Accordingly, we are providing a resource towards this end, a service through the commercial provider Qualys.

This should also come with the development of a patch management policy and standard procedures for all SPs to apply to the XSEDE software stack. For critical services, we also recommend developing specific security baselines that we can audit against.

Risk #19 Total Risk: 1 Likelihood: L Impact: L

Vulnerability
User data has weak isolation guarantees on most resources.

Potential Threat Source
Crackers and curious users

Potential Threat Actions
There is little besides file system permissions or ACLs that isolate users and their data on most XSEDE systems. There is potential for such basic mechanisms to be overcome and allow data snooping by adversaries. When networked file systems are used without encryption, this threat is increased.

Existing Controls
Users are warned that they should not put highly confidential materials on XSEDE, or if they do to use encryption appropriately. However, the onus is completely on the user.

Recommended Control Changes

This is a very low risk considering the types of data used on XSEDE. While more complicated file systems that use encryption could be utilized, the performance cost would be significant for such a low risk. Therefore we recommend accepting this risk and just reminding users during regular security training.

Risk #20 Total Risk: 5 Likelihood: M Impact: L

Vulnerability

Helpdesk tickets are emailed in plaintext.

Potential Threat Source

Crackers

Potential Threat Actions

Much of the ticket system communication is done over plaintext emails. Since some of the tickets are sensitive and contain security relevant information, attackers snooping those emails could gain an advantage.

Existing Controls

We have alternative, secure communication channels for security issues. Staff can also link to content on the staff wiki rather than email it directly.

Recommended Control Changes

We could stop putting the contents of the tickets in the email, but instead just send a URL to the ticket. However, this does not address the first email that creates a ticket, and it decreases usability measurably. Being such a low risk, we do not recommend making any technical changes. However, it would be prudent to remind users in training not to put very sensitive information in these emails. Sensitive details could instead be revealed over the phone or through other means.

Risk #21 Total Risk: 1 Likelihood: L Impact: L

Vulnerability

Incident response team members may log sensitive IM chats.

Potential Threat Source

Crackers

Potential Threat Actions

The incident response Jabber server uses SSL and does not log conversations. However, there is no control over the endpoints when using the Jabber server for incident response. So messages could be logged on client hosts and accessed more easily or exposed if a laptop is lost. This inside information could be used for gain by an attacker.

Existing Controls

There are no controls other than the fact that these are very security conscious people who likely harden their systems and don't log sensitive chats without encryption. However, there is no way to be sure the person on the other end is not logging.

Recommended Control Changes

We could create a policy requiring people on the incident response team to use full disk encryption on their workstations or turn logging off. However, there would be no way to verify this, and it is somewhat onerous to require for a very low risk.

Risk #22 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

Deployed software could be out of date, especially without coordinated patch management and stale CTSS registrations info.

Potential Threat Source

Crackers

Potential Threat Actions

While there is a common software stack for XSEDE compatibility, there is inconsistency across sites on versions deployed, with some sites using very out of date software. There is a threat that an attacker could leverage an exploit at one site to gain a foothold at another, or an attacker could exploit an old vulnerability in outdated software on a support service.

Existing Controls

There is a list of current software for the XSEDE software stack, but this is not audited nor enforced for patch-levels.

Recommended Control Changes

The most effective control is to update policy so that XSEDE software must be current or without currently exploitable vulnerabilities that could adversely affect XSEDE. Security operations would also need to regularly audit for compliance with this policy, and SEI would have to work with SPs to update as appropriate. Regular vulnerability scanning may help a little, but often this requires one to log onto a system to test. Furthermore, much of it is custom software not recognized by products like Qualys.

Risk #23 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

XSEDE relies heavily on in-house software that has not had code audits.

Potential Threat Source

Crackers

Potential Threat Actions

Many applets and pieces of software for XSEDE have been developed in-house without code reviews or expertise in security. There are likely unknown security flaws that could be exploited in a targeted attack. This is especially true of something as complex as the user portal, whose compromise would be harmful to XSEDE's reputation.

Existing Controls

There are no existing controls, but XSEDE benefits from the obscurity of much of this software, especially the pieces installed on HPCs.

Recommended Control Changes

It is expensive to hire out for code audits, and XSEDE lacks the internal experience and resources to perform full source code audits from a security perspective. New configuration items are getting lower-level architectural security reviews, but still no code audits. Besides requiring new configuration items to use some freely available fuzzing tools to check for common buffer overflow problems, little can be done without significant resources.

There are some new projects that may begin to fill this need for NSF communities, and XSEDE might make an ideal first customer. Until such services are available, we recommend accepting this risk.

Risk #24 Total Risk: 1 Likelihood: L Impact: L

Vulnerability

Some XSEDE resources depend upon proprietary, unvetted protocols.

Potential Threat Source

Crackers

Potential Threat Actions

Some proprietary protocols have been created for services (e.g., Globus listener over UDP). Developing secure protocols in nefariously hard, and in some cases there is no indication that any encryption or signing has been done. A very targeted attacker could exploit protocol vulnerabilities in ways that are very difficult to detect or deter.

Existing Controls

Obscurity, lack of a likely threat source, and the fact that there are many simpler methods of penetration make this unlikely. However, there are few controls deployed to protect against this threat other than generic methods to detect compromises. The only effective control we see used is that some of these protocols are tunneled through SSH or over SSL.

Recommended Control Changes

New protocols get reviewed as part of the design review in SD&I, and here we try to identify protocol weaknesses. Also, we push hard for communicating over secure tunnels to mitigate these types of risks. While we could try and force all existing protocols through the review process, this is infeasible as SD&I struggles to keep up with its cur-

rent load. Furthermore, this is a very time intensive process. Therefore, given the low risk and the difficulty of running every grandfathered-in service through a new review, we recommend accepting this risk.

Risk #25 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
XSEDE depends upon software that is no longer actively supported.

Potential Threat Source
Crackers

Potential Threat Actions
XSEDE depends on some software that no longer has active development, such as, the Java GSISSH terminal. This means there is no one there to fix security or reliability bugs, which could be exploited maliciously.

Existing Controls
There are no controls currently applied to this general problem. In fact, not even all of the currently unsupported software is identified in one place.

Recommended Control Changes
Unsupported software in the XSEDE software stack must first be identified. Then for each item, a decision must be made either replace it or commit resources to adopt and maintain it.

Risk #26 Total Risk: N/A Likelihood: N/A Impact: N/A

This risk was retired part way through the risk assessment process but left as a placeholder so not to confuse readers. These numbers are referenced all throughout the risk assessment documents and emails, and renumbering would lead to confusion unless it was done retroactively in 100% of the places used.

Risk #27 Total Risk: 50 Likelihood: H Impact: M

Vulnerability
There is a zero-day root escalation exploit in the wild for Linux or some common piece of the XSEDE software stack.

Potential Threat Source
Crackers

Potential Threat Actions

Software vulnerabilities are commonly found, and there are often crackers sitting on harvested user credentials waiting for the next Linux zero-day that could allow them to escalate their privileges to obtain root on an XSEDE resource.

Existing Controls

Short-lived certificates reduce the usefulness of harvesting those credentials. Some sites ban public key authentication for this reason. Portals like GlobusOnline are provided an OAuth MyProxy service so that users do not have to expose their XSEDE credentials to more places where they could potentially be harvested. The security operations team also monitors Bugtraq and other intelligence channels looking out for new exploits in the wild and recommend reactions based on the estimated threat to XSEDE. Finally, users are notified of the time and hostname from where they last logged in, helping them to notice if their account is compromised.

Recommended Control Changes

This is one of our highest risks because it has traditionally been one of our largest problems, and we still rely on basic password authentication for XSEDE. While we do not recommend forcing OTP for all users as it is costly for our user base and decreases usability, there are additional mitigations we recommend.

Automated vulnerability scanning and a successful patch management strategy will help us close the gap on these types of exploits more quickly, before the corresponding vulnerabilities might be exploited. However, there will always be cases where we cannot patch quickly enough, and in those cases it is important to detect these compromised accounts before they are exploited towards malicious ends. Usually an attacker will at least test an account once, even if they plan on sitting on it for a while. Therefore, we recommend that XSEDE security operations provide simple, lightweight user profiling tools for SPs to detect compromised accounts before they are used in conjunction with a zero-day exploit.

Risk #28 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

There is a common XSEDE service with a remote exploit.

Potential Threat Source

Crackers

Potential Threat Actions

By virtue of having services online, there is always a risk that a new vulnerability is discovered that allows remote exploitation that could either be combined with a local root escalation or that gives root itself. If such an exploit is in the wild and XSEDE is vulnerable, it is only a matter of time before scans by crackers discover it and exploit it.

Existing Controls

The security operations team monitors Bugtraq and other intelligence channels looking out for new exploits in the wild and recommend reactions based on the estimated threat to XSEDE.

Recommended Control Changes

We recommend identifying remotely accessible services listening on XSEDE systems and disabling unnecessary and insecure services. As part of the new security baselines and standards it should be policy not to run services as root if possible, and in other cases jails or other containment mechanisms should be utilized. Finally, our recommendation to employ regular vulnerability scanning will help mitigate this risk as well.

Risk #29 Total Risk: 50 Likelihood: H Impact: M

Vulnerability

Security policies and procedures from TeraGrid are out-of-date.

Potential Threat Source

Crackers

Potential Threat Actions

All the security policies and procedures are inherited from TeraGrid and fairly out-of-date and/or difficult to find. There is a risk that they are no longer completely applicable or adequate for XSEDE's operations, and new procedures will need to be made on-the-fly during an incident.

Existing Controls

This is risk currently unmitigated.

Recommended Control Changes

Review policies and procedures from TeraGrid and update and supplement as needed. Some additional policies to consider are a legal, law enforcement and regulatory plan, security responsibilities policy for SPs (includes policy regarding 2-factor authentication), a password policy, a separate AUP for PIs who are responsible for vetting users, and policy for campus gateways.

Risk #30 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

Services and software grand-fathered into XSEDE have not gone through the formal review processes that are now a part of XSEDE.

Potential Threat Source

Crackers

Potential Threat Actions

With inconsistent standards, attackers could target grandfathered-in services that might not pass security reviews today.

Existing Controls

Nothing currently mitigates the specific risk of targeting grandfathered-in services. These services are protected by the many generic security controls in place already, but nothing extra is done for specifically for them.

Recommended Control Changes

No changes are currently recommended. While running these through the SD&I process and performing code-audits for existing software would mitigate the risk, it has already been discussed why those controls are impractical at this time. Therefore, we recommend accepting this risk.

Risk #31 Total Risk: 50 Likelihood: H Impact: M

Vulnerability

There is no regular security training for users and service providers, and security policies lack visibility.

Potential Threat Source

Staff or users making mistakes

Potential Threat Actions

If users do not know what they should and should not do with respect to security, it is more likely that they can be taken advantage of by social engineers, will not protect their credentials well enough, or practice poor security hygiene. Because we do not audit SPs, this makes it more likely that they are not following security baselines either. Failures to implement policy and procedure are therefore more likely and could result in more down-time and slower response to an incident.

Existing Controls

There are some existing policies and a user MOU for new accounts inherited from TeraGrid. There is also some institutional memory of policies and procedures, though a lot has been lost. Some training materials have been created, but not delivered.

Recommended Control Changes

We should develop targeted training for different levels of SPs, users & XSEDE staff. The goal is to familiarize them with important policy and procedure. Some of this should be interactive (whether online or at meetings) and perhaps mandatory for staff. Just as important as developing materials is planning for adequate and regular dissemination (in-cluding at the annual XSEDE conference).

Risk #32 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
Incident response resources at different SPs vary, and the XSEDE incident response (IR) team is geographically distributed.

Potential Threat Source
Crackers

Potential Threat Actions
Different sites are more or less prepared to detect and respond to incidents, giving adversaries strategic options when deciding where to attack XSEDE. Being geographically and institutionally separated, the XSEDE IR team members must work harder to coordinate activities and overcome hurdles with information sharing that could slow response.

Existing Controls
We have bi-weekly security meetings and weekly incident response calls, both of which are rarely cancelled. We also have an incident response shared wiki space, a shared GPG email key, and a secure Jabber server for communication. Finally, there is an XSO (XSEDE Security Officer) responsible for security in XSEDE, a change from the less centralized TeraGrid.

Recommended Control Changes
In addition to all that we regular do with the existing controls, XSEDE should have regular security incident drills, requiring participation from all level 1 SPs.

Risk #33 Total Risk: 25 Likelihood: M Impact: M

Vulnerability
There are no security baselines for most services, and there is no regular auditing with respect to security.

Potential Threat Source
Crackers

Potential Threat Actions
As time goes by, systems likely drift from more secure and up-to-date configurations to less secure states. Crackers can target such infrastructure and find vulnerabilities to exploit more easily.

Existing Controls
The new System, Software & Engineering group does have a change control and testing process, which should help keep XSEDE-specific software more up-to-date. However, this does nothing for the base OS. Also, the SD&I security reviews and collaboration with the deployment team helps to ensure that when new services initially go online they are up-to-date. This does not help to **keep** things up-to-date though.

Recommended Control Changes

For this risk, as for others, we recommend regular scanning and auditing of XSEDE resources and support systems to ensure systems are patched against vulnerabilities deemed significant threats to XSEDE. Also, as recommended elsewhere, we should develop per service security baselines for critical XSEDE support services. Finally, none of this is effective without any auditing. Therefore, we recommend policy and implementation of a new XSEDE auditing process.

Risk #34 Total Risk: 25 Likelihood: M Impact: M

Vulnerability

XSEDE has very complex organizational and procedural structures.

Potential Threat Source

Mistakes or oversights by staff

Potential Threat Actions

With a large complex organizational structure with many geographically separated employees only part-time committed to XSEDE, there is increased risk of inaction or slow action at critical moments. This is exacerbated when roles and process are unclear and has an impact beyond just security—though agility is particularly important for security.

Existing Controls

In many ways this is a generalization of risk #32. So some of those controls are applicable to at least the security team, which is more tightly knit than many groups. But where security interacts with other groups, these organizational issues can become a security risk.

Recommended Control Changes

We should maintain strong leadership and regular meetings for XSEDE security operations. Participation of level 1 SPs should not be optional. Security personnel and resource allocations should also be evaluated and redistributed at least annually to be proportional to where the work is done or needs to be done. While we cannot do anything about XSEDE's overall organizational structure or geography, we can keep the operations security group tight knit with good lines of communication.

Risk #35 Total Risk: 5 Likelihood: M Impact: L

Vulnerability

Resources for security in XSEDE could be inadequate.

Potential Threat Source

Mistakes or oversights by staff

Potential Threat Actions

With the merger of XSEDE and XROADS there is much more work but not more resources. If resources are contentious, it forces choices about what to focus on and how much vetting to do for security approval. There is also the risk that focus is on the wrong thing when choices are made.

Existing Controls

There are annual reviews in XSEDE, and resources have already been shifted and reallocated after the first year.

Recommended Control Changes

The same recommendation to risk #34 applies here. Also, one of the major outcomes of doing this risk assessment is to help us optimize our impact with the resources we have.

Summary

There were 35 total risks identified, none of them considered high. Their distribution by severity based on likelihood and impact can be seen in the table below. The four highest risks were ranked 50 out of 100 possible points and are our primary risks to address in the next year. There are a further 15 medium risks of rank 25, 10 of which we recommend addressing after the 4 primary risks. The remaining 6 medium risks would be more costly to address than to simply accept. Finally, there are 15 low risks ranked no more than 5 of 100 potential points, most of which we recommend accepting. The few we do recommend controls for would come as the lowest priority, after all the medium risks are addressed.

	Low Impact	Medium Impact	High Impact
High Likelihood	Low (10)	Medium (50)	High (100)
	No risks	**4 risks**	No risks
Med. Likelihood	Low (5)	Medium (25)	Medium (50)
	8 risks	**15 risks**	No risks
Low Likelihood	Low (1)	Low (5)	Low (10)
	7 risks	No risks	No risks

Recommended Controls

High Priority Controls

- Automate regular scanning and auditing of XSEDE resources and support systems to ensure systems are patched against vulnerabilities deemed significant threats to XSEDE. This should also come with the development of a patch management policy and standard procedures for all SPs to apply to the XSEDE software stack.
- Set policy to require two factor authentication for administrative interfaces and functions on XSEDE resources where possible. The XSEDE security operations team must also audit this regularly for all tier 1 & 2 services (to be defined by the new policy).
- Provide simple, lightweight user profiling tools across SPs to detect compromised accounts before they are used in conjunction with a zero-day exploit.
- Review policies and procedures from TeraGrid and update and supplement as needed. Some additional policies to consider are a legal, law enforcement and regulatory plan, security responsibilities policy for SPs (includes policy regarding 2-factor authentication), a password policy, a separate AUP for PIs who are responsible for vetting users, and policies for campus gateways.
- Develop targeted training for different levels of SPs, users and XSEDE staff. The goal is to familiarize them with important policy and procedure. Some of this should be interactive (whether online or at meetings) and perhaps mandatory for staff. Just as important as developing materials is planning for adequate and regu-

lar dissemination, and we recommend at least one security training session at XSEDE annual meetings.

Medium Priority Controls

- Maintain strong leadership and regular meetings for XSEDE security operations. Participation of level 1 SPs should not be optional. Security personnel and resource allocations should also be evaluated and redistributed at least annually to be proportional to where the work is done or needs to be done. While we cannot do anything about XSEDE's overall organizational structure or geography, we can keep the operations security group tight knit with good lines of communication.
- Develop per service security baselines for critical XSEDE services.
- After updating policies and baseline security requirements for SPs, perform audits to make sure XSEDE systems and services stay true to these baselines. This could also mitigate multiple risks if the baseline includes requirements for offsite backups. Audits need not be invasive, and may consist in part of questionnaires to be filled out by SPs or checklists for them to complete.
- Identify unsupported software in the XSEDE software stack and either (a) retire and eradicate it, (b) update or replace it, or (c) commit resources to adopt and maintain it.
- Have regular security incident drills and require participation from all level 1 SPs.
- Identify services not replicated across at least two XSEDE SPs and either replicate them or move to a service provider with DoS protections. Focus on critical services first.
- Identify remotely accessible services listening on XSEDE systems, and disable unnecessary and insecure services.
- Identify and minimize services running as root on XSEDE systems; reconfigure to run as non-root wherever possible, and recommend jails or other containment mechanisms where this is not possible.
- Inventory how system accounts are used for various services and identify those using unencrypted certificates or SSH keys. Protect the credential as best as possible, limit the capabilities of the credential and corresponding account to only those needed, and monitor for any unexpected use of such accounts.
- Deploy DNSSEC on all XSEDE.org subzones and the root domain.

Low Priority Controls

- Where possible, remove shared accounts (e.g., CVS, DOORS, etc.) and put into place mechanisms that don't require a shared password among several people. Extra auditing mechanisms should be applied, in cases where this cannot be avoided.
- Create a standard for authentication to acquire an XSEDE acceptable certificate, and audit for compliance with the standard.
- Regularly review and update the XSEDE risk assessment to stay focused on the most relevant security issues.
- Identify and remove local or non-XSEDE authentication systems for XSEDE services and replace with XSEDE Kerberos or GSI where possible.
- Require a second email address to be registered for account creation.