Design for XSEDE SP Deployment of Globus Connect Server version 5.4

Lee Liming

March 29, 2021

Globus Connect Server version 5.4 (GCS 5.4) is the next generation of the storage provider's interface to the Globus system. It connects a storage system to the Globus data transfer service, enabling researchers to use the Globus web application (https://app.globus.org), Globus CLI, Globus Transfer API, or a Web browser to access the storage system, subject to the administrator's data access policies.

GCS 5.4 replaces the X.509-based security mechanism used in GridFTP and Globus Connect Server version 4 with a new mechanism based on OpenID Connect 1.0 (OIDC) and OAuth 2.0 (OAuth2). OIDC/OAuth2-based security is widely used in the research community and in the mainstream Internet.

This document supplies a proposed design for XSEDE SPs who wish to use Globus Connect Server 5.4 for their XSEDE data transfer endpoints. The design specifies software configurations that implement specific data access policies. It also enables XSEDE system-wide conventions that improve the user experience for researchers who use XSEDE services.

TABLE OF CONTENTS

1. SCOPE

This design is focused on the requirements, behavior, and implementation of data access in the XSEDE context with Globus Connect Server version 5.4 (GCS 5.4). Specific XSEDE use cases are identified in Section 4.2.

Beyond GCS 5.4 itself, this design also describes related system components that are involved in data access, notably including the services—both within and external to XSEDE—that are involved in the authentication and authorization performed by GCS 5.4.

1. Data transfer nodes (host systems)
2. Network interfaces
3. Connected storage
4. XSEDE Mapfile API and CLI
5. Globus Auth
6. XSEDE Identity Provider
7. CILogon
8. OIDC, InCommon, and eduGAIN identity providers

2. ACRONYMS AND TERMINOLOGY

access control - a set of mechanisms, policies, and procedures that control access to a resource; includes but is not limited to authorization

allocation - in XSEDE, a time-limited authorization to use one or more specific resources in the XSEDE system; an allocation is granted by the XSEDE Resource Allocation Committee (XRAC) to a principal investigator who may, in turn, share the allocation with designated colleagues

authentication - the process that confirms a user's identity

authorization - a process that decides whether or not to provide access to a specific feature or data resource, usually involving the current user's identity/identities and related factors; one of several elements involved in access control

claim - an individual element (e.g., full name, username, email address, organizational affiliations) in an identity data structure

credentials - information (e.g., username and passwords) and artifacts (e.g., a mobile phone or one-time-password device) issued by or registered with an IdP to be used during user authentication

data transfer node - a Linux-based server specialized for performing high-performance data transfers to and from connected data storage

identity - a data structure by which an IdP uniquely identifies an individual; each identity includes a set of *claims* about the user's identity

identity set - in Globus, a set of identities from multiple IdPs that all correspond to the same individual

IdP (Identity Provider) - an organization that establishes and asserts identities for individuals; e.g.: the U.S. Social Security Administration; a place of employment; an educational institution

OAuth2 (OAuth 2.0) - an industry-standard protocol for authorization defined by the Internet Engineering Task Force (IETF); see [1]

OIDC (OpenID Connect Core 1.0) - a standard interface for user authentication and identity claims defined by the OpenID Foundation; see [2]

registration - 1. the process by which an IdP establishes an identity for an individual, typically involving recording information about the individual and issuing credentials; 2. the process by which an OIDC authentication service establishes an identity for an application or service that will use the OIDC service, typically involving recording information about the application or service and issuing a Client ID and credentials

resource - a computer system to which researchers are given access in order to run research applications; L1/L2 (Level 1 and Level 2) resources are high-value resources (>$1M) with significant security requirements

SP (Service Provider) - an organization that operates and offers resources for use in the XSEDE community

3. PRODUCT DESCRIPTION

Globus Connect Server version 5.4 (GCS 5.4) is software that provides remote access to data storage on a specific resource. GCS 5.4 is provided by Globus and is intended to be installed and configured on data transfer nodes associated with resources by the SP that operates each resource. GCS presents interfaces that allow remote data access and remote management, both within a rigorous access control system.

GCS 5.4 supports two data access interfaces. The first data access interface supports the Globus Transfer service, operated by Globus, which enables authorized researchers to access the storage system using the Globus web app (https://app.globus.org), the Globus CLI, and custom applications that use the Globus Transfer API. This interface is the GridFTP protocol, which has been modified to use OAuth2 tokens for authorization instead of X.509 credentials. The second data access interface is HTTPS, the encrypted protocol used by web browsers. The HTTPS interface allows researchers to directly upload or download data using their web browser. The HTTPS interface also uses OAuth2 tokens to authorize access, and all of these interactions take place within a rigorous access control system.

GCS 5.4 also provides an administrative interface that allows administrators to dynamically configure the service and keep the service's configuration consistent across multiple data transfer nodes. As an option, the administrative interface can also enable authorized resource users to configure guest collections: limited, secure guest access to specific folders on the system to support their research activities.

GCS 5.4 enables administrators to configure multiple distinct data access policies for their system. A single GCS 5.4 installation can present multiple user access points, each with its own access policy (authentication requirements, identity mapping algorithm) and access restrictions (accessible portions of the filesystem, accessible Globus service features). This feature is especially important for XSEDE SPs who support distinct user communities (XSEDE being one of several) and/or different modes of use (immediate use by individual researchers vs. programmatic or automated use by science gateways or other applications).

4. XSEDE REQUIREMENTS

This design for GCS 5.4 deployments on XSEDE resources is intended to satisfy a number of important needs within the XSEDE community. This section identifies these specific needs and is divided into two sections. The first section identifies the specific use cases in which GCS 5.4 deployment on SP resources plays a critical role. The second section identifies a set of access control requirements that are likely to be elements of specific XSEDE and XSEDE SP policies.

4.1 Use cases

Use cases describe things people can do in the context of the XSEDE system. In most cases, fully enabling a use case involves several parts of the system. Table 1 identifies the specific XSEDE use cases for which a GCS 5.4 deployment on an XSEDE SP resource—specifically as configured per Section 5.1 of this document—can enable the full realization of the use case.

In most cases, GCS 5.4 isn't the only software component that can satisfy these use cases. (XSEDE has been satisfying these use cases for years using earlier versions of Globus Connect software and its predecessors.) Table 1 explicitly indicates the few use cases that—to the best of our knowledge—can only be satisfied in XSEDE through use of GCS 5.4.

Use case	**Only with GCS 5.4?**
CAN-02: Managed file transfer	
CB-04: Access campus research data from a community resource	
CB-09: Access a community data collection from campus	
CB-10: Synchronize research data between campus and community resources	
DA-02: Prepare data for analysis	
DM-01: Create and share a data collection	
DM-03: Automate data ingestion from a set of sensors or instruments	
DM-04: Migrate data to a new resource	
DM-11: Post-allocation data access	
DM-12: Large-scale data transfer	

DM-15: Transfer data between a researcher's cloud storage and a community storage system	Yes[1]
HPC-01: Use a single HPC resource for a research project	
HPC-02: Use two or more HPC resources for a research project	
SGW-02: Transfer files to and from a community resource	Yes[2]
SGW-04: Enable gateway users to transfer files to or from a community resource	Yes[3]

TABLE 1. Use cases in which this design can be used

4.2 Access control requirements

Each XSEDE SP has its own unique set of access control requirements. In fact, it is common for a single SP to have several distinctly defined uses of their system, each having its own access control requirements.

This section enumerates specific access control requirements that may appear as elements of specific XSEDE and SP policies. Most policies will combine several of these requirements. It is unlikely that any specific policy will have all of these requirements. Section 5 ("High-level design") references these requirements in the recommended configuration.

1. Access is limited to individuals associated with an active XSEDE allocation.
2. Access requires multi-factor authentication (MFA).
3. Each access must be logged.
4. Each access must be traceable to an individual.
5. Access is limited to the smallest necessary portion of the system.
6. Resource administrators can disable access by a specific individual designated by an XSEDE identity or a resource-specific identity.
7. If credentials (see definition) are transmitted between systems, the transmission must utilize encryption and confidential communication.
8. XSEDE-issued credentials are not used by any service or application that is not operated by XSEDE Operations personnel.

[1] This use case requires the SP to have access to a Globus subscription that includes a cloud storage connector add-on for the type of cloud storage used. See https://www.globus.org/connectors.
[2] This is possible *but much more difficult and error-prone* using a GridFTP server, and it is also possible using Globus Connect Server version 4, but XSEDE has not documented a deployment design for GCS 4.
[3] This can also be accomplished using Globus Connect Server version 4, but XSEDE has not documented a deployment design for GCS 4.

5. HIGH-LEVEL DESIGN

This section describes the recommended configuration for GCS 5.4 on XSEDE SP resources (Section 5.1) and the resulting user experience when using this configuration (Section 5.2). Section 6 ("Detailed design") provides the details of how GCS 5.4 accomplishes the user experience described here and satisfies the access control requirements in Section 4.2.

5.1 Product configuration

The following sections describe the recommended configuration for GCS 5.4 on XSEDE SP resources. The configuration is divided into several sections, several of which are optional. This provides a menu of defined options from which SPs may select specific features for their resources. Each option references the specific access control requirements it satisfies.

5.1.1 Base configuration for XSEDE GCS 5.4 endpoints

Every GCS 5.4 installation begins with the creation of an ***endpoint***. The endpoint is the administrative interface for the GCS 5.4 installation. This administrative interface is used only by authorized administrators. *It does not enable data access.* In multi-DTN configurations, a single endpoint encompasses all of the DTNs.

Section 1 of the *XSEDE Globus Connect Server v5.4 Installation Guide* [15] describes the steps for setting up an endpoint. The following are the important design points.

- Endpoints that will be used to support XSEDE data access are "managed" by a Globus subscription. Whether using the XSEDE subscription or an SP subscription, configuring the endpoint with a subscription enables subscription features and also enables a visual cue in the Globus web interface's search results that assures people they're using the right collection.
- Globus's endpoint roles are used to share the administration of the endpoint with appropriate individuals in your organization. The administrator guide recommends creating one or more Globus groups that define the system administrators for the endpoint and assigning each endpoint role (Administrator, Activity Manager, Access Manager) to the appropriate group.
- Additional DTNs are added to the endpoint via the method described in Globus's *Globus Connect Server v5.4 Installation Guide* [16].

5.1.2 Data access configuration

In GCS 5.4, data access is configured by creating *storage gateways* and *collections*. A GCS 5.4 endpoint can have an unlimited number of storage gateways and collections. This enables

modular data access configuration: a single Globus endpoint can provide separate data access for a variety of distinct user communities and/or usage modes.

- The **storage gateway** defines a specific access policy. It specifies: *who* is authorized, *how* are identities managed and mapped, and *what* features are enabled.
- A **collection** is a specific set of storage areas (e.g., `/projects`, `/users`, and `/scratch`) made available **under the access policy defined by a storage gateway**. Each combination of (storage areas + access policy) is a distinct collection, and collections are what Globus users see when they search for storage locations in Globus. Collections may have overlapping storage areas.

For GCS 5.4 endpoints that provide access for XSEDE projects, **we strongly recommend separate storage gateways and collections for XSEDE use**, because use within the XSEDE context has different access requirements and identity mappings than uses outside the XSEDE context. (In the XSEDE context, an active XSEDE allocation is required and the XSEDE username is mapped to a local username.)

In a nutshell, the recommended XSEDE data access configuration is as follows.

- One collection for members of XSEDE-allocated projects.
- (optional) An additional collection that allows authorized individuals to create guest collections.
- (optional) An additional collection for use by XSEDE-allocated science gateways.
- (optional) Additional collections to support non-XSEDE uses.

The first two collections above can be combined if the local access policy permits it. Otherwise, each collection is distinct from the others.

5.1.2.1 Primary access - For members of XSEDE-allocated projects

The most common mode of data access on XSEDE SP resources is access by individual members of an allocated project. These individuals must be part of a project team defined in the XSEDE accounting service and the project must have an active allocation to use the resource. Each individual has a local account on the resource. Individuals may access data on the resource using the Globus web app (https://app.globus.org/), using the Globus CLI, or using an application that uses the Globus Transfer API.

Section 2.1 of the *XSEDE Globus Connect Server v5.4 Installation Guide* describes the steps for setting up this access mode. The following are the important design points.

For the storage gateway:

- The user must have an `xsede.org` identity.

- Authentication expires after 11 days.
- Local accounts are mapped from XSEDE identities using an external mapping script, which uses the XSEDE mapfile for the relevant XSEDE resource.
- Data access is restricted to filesystem paths designated for XSEDE use.

For the mapped collection:

- The display name and descriptive properties (description, contact information, web link) make the purpose of the collection clear in search results.
- Guest collections are disabled.

This configuration satisfies all of the access control requirements except #2: “Access requires multi-factor authentication.” This requirement is not satisfied because—although Globus requires an XSEDE identity for authorization—it allows linked identities for authentication, and a linked identity might not require MFA.

5.1.2.2 Option 1 - Enable project members to create guest collections

The access mode described in the previous section is consistent with traditional XSEDE access but is more restrictive than may be necessary. It limits access to individuals who have both an XSEDE account and a local account on the resource and who are part of an XSEDE project with an active allocation.

Globus can be configured to enable authorized individuals who meet the criteria above to share secure guest access to specific data with other Globus users. For example, a researcher who is a member of an XSEDE project might share read-only access to the results of a specific application run with a specific group of students, research assistants, or research partners who, themselves, are not part of the researcher’s XSEDE project.

Sharing is achieved in Globus via *guest collections.* An authorized individual first browses the system using the mapped collection. When the individual needs to share something, he or she creates a guest collection that addresses a specific directory. Once created, this guest collection allows the researcher to set permissions on directories within the guest collection that enable specific individuals or groups to access it. (Permissions can grant read-only or read-write access, assuming the creator of the guest collection has read-write access.) Individuals who are granted permission via the guest collection search Globus for the name of the guest collection (because they cannot use the mapped collection) and can only access folders within the guest collection to which they’ve been granted permission.

If the local access policy allows guest collections to be created in *all filesystem paths* accessible by Globus, then this access mode can be configured simply by enabling guest collections in the mapped collection described in Section 5.1.2.1. If guest collections are only permitted for a subset

of the accessible filesystem paths, then a separate mapped collection with more restricted access should be created.

Section 2.2 of the *XSEDE Globus Connect Server v5.4 Installation Guide* describes the steps for setting up an XSEDE mapped collection that also supports guest collections. The following are the important design points, and differences from the access mode in the previous section are highlighted.

For the storage gateway:

- The user must have an `xsede.org` identity.
- Authentication expires after 11 days.
- Local accounts are mapped from XSEDE identities using an external mapping script, **but the mapfile only contains mappings for individuals who are permitted to create guest collections.**[4]
- Data access is restricted to filesystem paths designated for XSEDE use **with guest collections**. (This may be a subset of the paths used in the previous section.)

For the mapped collection:

- The display name and descriptive properties (description, contact information, web link) make the purpose of the collection clear in search results.
- **Guest collections are enabled.**

This configuration satisfies access control requirements 3, 5, 6, 7, and 8.

- Requirement 1 ("Access is limited to individuals associated with an active XSEDE allocation") is relaxed to allow for sharing.
- Requirement 2 ("Access requires multi-factor authentication") is relaxed to allow sharing with individuals whose identity provider does not require MFA.
- Requirement 4 ("Each access must be traceable to an individual") is not satisfied unless the endpoint is managed by a High Assurance or HIPAA+BAA subscription, in which case it is satisfied by GCS 5.4's audit log.

5.1.2.3 Option 2 - Enable XSEDE-allocated science gateways

A science gateway is an application used by a community of researchers. The developer or operator of the application has been granted an allocation to use an XSEDE SP resource and is assigned a *community account* by XSEDE. The community account is a member of the project's team, so it has local accounts on resources allocated to the team. The application uses the

[4] An alternate configuration option is to use the complete XSEDE map file for mapping, but explicitly allow (or deny) guest collection creation for specific individuals. This is described in Section 2.2 of the *XSEDE Globus Connect Server v5.4 Installation Guide*.

community account to access the XSEDE SP resources when necessary. Researchers who use the application (usually) do not have accounts on the XSEDE SP resources. The science gateway application uses the Globus Transfer API—or in some cases the Globus CLI—to access data on the XSEDE SP resources.

Science Gateways represent a different risk profile for XSEDE SPs because access is granted to an application that has many different end users. XSEDE credentials are issued to the application developer or operator and are used to configure the application. Furthermore, XSEDE's science gateway documentation recommends creating a guest collection for use by the science gateway because it is simpler and less accident-prone to enable an automated application to use a Globus guest collection than a mapped collection. (See Section 5.1.2.2 for an explanation of guest collections.) But guest collections are not allowed by the primary access policy. For these reasons, we recommend a separate configuration for use by Science Gateways.

Section 2.3 of the *XSEDE Globus Connect Server v5.4 Installation Guide* describes the steps for setting up this access mode. The following are the important design points, and differences from the access mode in the first section are highlighted.

For the storage gateway:

- The user must have an `xsede.org` identity.
- Authentication expires after 11 days.
- Local accounts are mapped from XSEDE identities using an external mapping script, **but authorization is limited to community accounts.**
- Data access is restricted to filesystem paths designated for XSEDE use **by community accounts**. (This may be a subset of the paths used in the previous section.)

For the mapped collection:

- The display name and descriptive properties (description, contact information, web link) make the purpose of the collection clear in search results.
- **Guest collections are enabled.**

This configuration satisfies access control requirements 1, 3, 5, 6, 7, and 8.

- Requirement 2 ("Access requires multi-factor authentication") is relaxed to allow access by science gateways.
- Requirement 4 ("Each access must be traceable to an individual") is relaxed to allow access by science gateways.[5]

[5] Science gateways are required to report the identities of the gateway users who are served by XSEDE compute jobs, but this requirement does not apply to data access.

NOTE 1: This configuration enables community accounts to access the permitted filesystem space and to create guest collections. However, the purpose of enabling guest collections ***is not*** to enable gateways to share access beyond the community account, and it is expected that gateway operators will not set permissions in guest collections that allow identities other than the community account itself to access data. The purpose of enabling guest collections is to enable a simpler credential management workflow for the science gateway application as described by use case SGW-02 in Section 4.1 and in XSEDE's science gateway documentation.

NOTE 2: Some science gateways may also use guest collections to enable gateway users to transfer data between their personal systems and the gateway's XSEDE SP storage, as described by use case SGW-04 in Section 4.1 and in XSEDE's science gateway documentation. In these instances, **access control requirement 1 is also relaxed**.

5.1.2.4 For non-XSEDE uses of the endpoint

XSEDE SP resources often serve communities other than XSEDE. To support non-XSEDE use and users, this design recommends each SP create storage gateways and collections that are entirely separate from those used to support XSEDE uses as described in the previous sections.

A separate storage gateway can be configured to use an entirely different authentication policy.

- It can enable authentication using a different identity provider, such as an institutional login service or even local PAM authentication using Globus's built-in OpenID Connect option.
- It can enable a different identity mapping scheme.
- It can be configured to access separate filesystem spaces.
- It can be configured to separately enable or disable guest collection creation.

Any of these options might be needed to enable uses by non-XSEDE communities.

5.2 Access control behaviors

This section describes the behaviors enabled by the configuration described in Section 5.1.

5.2.1 Administrator access

The base configuration described in Section 5.1.1 enables administrator access for individuals who are granted specific administrative roles on the endpoint. Section 6.2.3 provides more detail about the features of administrative access.

Administrative access is via the GCS CLI or the Globus web app. (The GCS CLI is distinct from the Globus CLI, which is used for data access. It is installed on the server system when GCS 5.4 is installed.) The GCS CLI is used to configure storage gateways, collections, and other endpoint

features beyond the scope of this design. Administrators authenticate with the GCS CLI using the `globus-connect-server login` command.

The Globus web app (https://app.globus.org/) can also be used to change the descriptions of the endpoint itself and the mapped collections on the endpoint. (The endpoint description is seen only by administrators. Descriptive information for mapped collections is important for enabling XSEDE researchers to find and use the endpoint.) Administrators must login to the Globus web app using an identity that is (or is linked to) the identity with the administrative role on the endpoint.

5.2.2 Data access

Each of the four policies defined in Section 5.1.2 enables a distinct set of data access behaviors. These are summarized in the table below. The recommended configuration enables data access both using the Globus transfer service (Globus web app, CLI, and API) and using HTTPS (web browser upload/download). **Access control behavior is exactly the same for both protocols.**

Policy	Mapped Collection	Guest Collections
Primary (5.1.2.1)	• access requires an XSEDE identity and a local system account • restricted to accounts in XSEDE mapfile • tokens valid for 11 days • paths restricted to those specified by admin • local logs identify each access by local account	n/a
Guest collections (5.1.2.2)	• access requires an XSEDE identity and a local system account • restricted to accounts authorized to create guest collections • tokens valid for 11 days • paths restricted to those intended for guest collections • local logs identify each access by local account	• paths restricted to those configured to allow guest collections and (more specifically) the folder where each guest collection was created and its contents • access subject to limits on the account that created the guest collection • access requires an identity that has user or group permission in the folder's ACL • tokens auto-activate • local logs indicate each access using the account that created the guest collection
Science gateways (5.1.2.3)	• access requires an XSEDE community account mapped to a local system account • restricted to community accounts w/allocations • tokens valid for 11 days • paths restricted to those intended	• paths restricted to those intended for use by science gateways and (more specifically) the folder where each guest collection was created and its contents • access subject to limits on the account that created the guest collection • access requires an identity that has user

	for science gateways • local logs identify each access by local account	or group permission in the folder's ACL • tokens auto-activate • local logs indicate each access using the account that created the guest collection
Other uses (5.1.2.4)	• access requires identities and uses mappings defined by admin • tokens valid for period defined by admin • paths restricted to those specified by admin • local logs identity each access by local account	• available if allowed by admin • policies specified by admin

TABLE 2. Data access behaviors for the four data access policies in Section 5.1.2

5.2.3 Authentication

GCS 5.4 uses Globus Auth for user authentication and identity management. Details about Globus Auth are provided in Section 6.1.5. In brief, Globus Auth works as follows.

- Globus Auth provides an OpenID Connect (OIDC) [2] authentication service that enables applications and services to authenticate users and obtain user identity information.
- When authenticating, users select an organization to authenticate with from a list of thousands of OIDC, InCommon [4], and eduGAIN [5] IdPs. (XSEDE is a particularly significant option.)
- Globus Auth enables individuals to manage an *identity set* containing their own identities from multiple organizations. (The identities in the identity set are often referred to as *linked identities*.) This enables applications to learn about and use the individual's full set of identities. (See the detailed design section for more information.)
- In general, GCS 5.4 allows access to features and collections if any identity in the active identity set has been granted access.[6]

Specific user authentication behavior is determined by the collection being accessed and the IdP selected by the individual accessing it. Table 3 summarizes the behaviors and sources of identity data. See Section 6.1.6 for specific IdP behaviors.

[6] Stricter requirements are used when High Assurance Globus features are configured. High Assurance features are not used in this design, however.

Configuration	Authentication behavior is determined by...	Identity data comes from...
Administrative access (5.1.1)	IdP selected by the user for authentication	All IdPs appearing in the user's identity set
Primary data access (5.1.2.1)	IdP selected by the user for authentication	XSEDE IdP
Guest collections (5.1.2.2)	IdP selected by the user for authentication	• XSEDE IdP (for data access in mapped collections and for guest collection management) • IDPs appearing in guest collection ACLs (for data access in guest collections)
Science gateways (5.1.2.3)	IdP selected by the user for authentication	• XSEDE IdP (for data access in mapped collections and for guest collection management) • IDPs appearing in guest collection ACLs (for data access in guest collections)

TABLE 3. The IdP that determines behavior and that provides identity data in each product configuration

5.2.4 Authorization

In the **primary data access configuration** (Section 5.1.2.1), GCS 5.4 will authorize data access if all of the following are true:

- the path is within the path restrictions for the collection
- the individual has an XSEDE identity in his or her identity set
- the individual's XSEDE username appears in the local mapfile
- the mapped account has the required permissions

The behavior above applies for access both via the Globus Transfer service and via HTTPS.

In the **optional guest collection configuration** (Section 5.1.2.2), data access *in the mapped collection* is authorized by GCS 5.4 as it is in the primary configuration. (Note that both the mapfile contents and the path restrictions may be more restrictive than they are for the primary configuration.) Data access *in a guest collection* is authorized if the account that created the collection has access and if the current user's identity set contains an identity that is granted access by the guest collection's access control list. (These authorization behaviors apply for access both via the Globus Transfer service and via HTTPS.) GCS 5.4 will authorize *a mapped local account to create a guest collection* if all of the following conditions are met:

- the mapped collection is configured to allow guest collections
- the path for the guest collection is within the mapped collection's path restrictions
- the mapped local account is in the storage gateway's mapfile and is permitted to create guest collections by the mapped collection's guest collection restrictions
- the resource's internal access control mechanisms permit the mapped local account to access the path

In the **optional science gateway configuration** (Section 5.1.2.3), data access *in the mapped collection* is authorized by Globus Connect Server as it is in the primary configuration. (Note that both the mapfile contents and the path restrictions may be more restrictive than they are for the primary configuration.) Data access *in a guest collection* is authorized if the account that created the collection has access and if the current user's identity set contains an identity that is granted access by the guest collection's access control list. (These authorization behaviors apply for access both via the Globus Transfer service and via HTTPS.) GCS 5.4 will authorize *a mapped local account to create a guest collection* if all of the following conditions are met:

- the mapped collection is configured to allow guest collections
- the path for the guest collection is within the mapped collection's path restrictions
- the mapped local account is in the storage gateway's mapfile and is permitted to create guest collections by the mapped collection's guest collection restrictions
- the resource's internal access control mechanisms permit the mapped local account to access the path

6. DETAILED DESIGN

This section documents important details about how GCS 5.4 works, including its dependencies on external services.

6.1 Environment and dependencies

6.1.1 Data transfer nodes (host systems)

The technical specifications for the data transfer nodes associated with an XSEDE resource are usually determined as part of the resource's design, proposal, and acquisition. Key specs include as much volatile memory as possible and network interfaces that can fully utilize the throughput of the data storage system (storage interconnect) and the wide area network. If a single data transfer node is not able to fully utilize the lesser of the storage interconnect throughput or the wide area network throughput, multiple data transfer nodes can be used. GCS 5.4 can bind multiple data transfer nodes into a single logical "endpoint," using the combined resources of the nodes for transfers involving multiple files. An excellent resource for designing a state-of-the-art data transfer node, including a reference system, is provided by the Department of Energy's ESnet "Faster Data" website. [13]

6.1.2 Connected storage

This design assumes all connected storage is available via a standard POSIX filesystem interface. Use of non-POSIX storage connectors is not covered by this design.[7]

As with the design of data transfer nodes, SP storage and its connectivity is usually predetermined when a resource is acquired. The Department of Energy's ESnet "Faster Data" website referenced above includes recommendations for storage and storage connectivity.

6.1.3 XSEDE Mapfile API and CLI

All of the recommended data access configurations in Section 5.1.2 use a local mapfile to map OIDC identities to local accounts.

XSEDE's Central Database (XCDB) provides a RESTful API and a command-line tool (CLI) for generating these mapfiles. (The CLI uses the API.) [7] The API and the CLI both return a list of all users with active allocations on a given resource, including the local resource username and the corresponding XSEDE identity (*username*@xsede.org) for each user.

[7] GCS 5.4 supports a variety of storage connections, but non-POSIX storage requires subscription add-ons that are not included in the XSEDE subscription. Some XSEDE SPs have local subscriptions that include a storage connector add-on for use with their resource.

The API uses client authentication and HTTPS encryption to maintain the privacy of PII (specifically, XSEDE usernames) contained in the mappings. Each SP must obtain an API client key to access the API.

SPs are encouraged to use the CLI or API provided by XSEDE to generate their mapfiles, particularly the mapfile used for the primary data access configuration (Section 5.1.2.1). SPs are free to edit their mapfiles to add additional users or to ban specific individuals.

The two optional configurations recommend separate mapfiles containing (respectively) local accounts authorized to create guest collections (Section 5.1.2.2), and XSEDE community accounts used with science gateways (Section 5.1.2.3). These can be constructed by hand, either starting from the mapfile generated by the XSEDE tool or using a local configuration management tool. In both cases, the full mapfile can be used instead if permission to access data and create guest collections is explicitly permitted (or denied) to individual users or groups in the mapped collection configuration.

6.1.4 Globus Transfer service

As described in Section 6.2, the Globus Transfer service is one of two ways to access data using a Globus Connect Server collection. (The other way is direct web browser upload/download via HTTPS.) Access control for collections via the Globus Transfer service is described in detail in Section 6.2.1.

This section describes how the Globus Transfer service sets up transfers between two endpoints. All clients of the Transfer service—including the Globus CLI, Globus SDK, and applications that use the Globus Transfer API—work the same way. GCS 5.4 supports data transfers with other Globus Connect Server endpoints (versions 4.x and 5.4), with Globus Connect Personal endpoints, and with GridFTP endpoints.

All transfers managed by the Globus Transfer service happen in the context of a user session with the Transfer service. During that session, the user locates each of the transfer endpoints using the Transfer service's endpoint registry and *activates* each endpoint (separately) using the authentication mechanism supplied by the endpoint.

The result of activating an endpoint is that the Globus Transfer service has a credential allowing it to access the endpoint on behalf of the user. Depending on the type of endpoint, this credential could be an X.509 certificate, an X.509 proxy certificate, or an OAuth2 access token for a collection on the endpoint. The mechanisms supported by each type of endpoint are shown in Table 4. *The mechanisms supported by the two endpoints in a transfer need not be the same.*

Endpoint software	Supported activation mechanisms
GridFTP (Globus Toolkit 5.0+)	X.509 via MyProxy X.509 via MyProxy OAuth X.509 via CILogon
Globus Connect Server version 4.x	X.509 via MyProxy X.509 via MyProxy OAuth X.509 via CILogon
Globus Connect Server version 5.4	OAuth2 via OIDC IdP

TABLE 4. Endpoint activation mechanisms supported by various endpoint software

Once the Transfer service has the credentials necessary to create a control channel to each endpoint (two separate channels, one for each endpoint), it can use the control channels to establish data channels between the two endpoints. (Most transfers involve multiple parallel data channels, and new channels may need to be set up over time during a transfer.) When the Transfer service needs to establish a set of data channels, it uses the Data Channel Security Context (DCSC) feature in the GridFTP protocol to establish a unique security context with both endpoints. [8] This security context does not rely on the credentials used in endpoint activation, and instead uses a unique 2048-bit RSA key, which is communicated with each activated endpoint via the control channel. This security context is then used to mutually authenticate the data channels between the endpoints. The DCSC feature is supported by GridFTP 5.x and later versions and by all Globus Connect Server versions. If at least one of the two endpoints supports DCSC, this feature can be used to establish data channels between the endpoints.

6.1.5 Globus Auth

Globus Auth is an OAuth2-compliant authorization service that also provides an OIDC-compliant authentication service. Significant features of Globus Auth include the following.

- For user authentication, Globus Auth acts as a client to other OIDC services, consuming tokens for user identities and other claims. (The list of other OIDC services includes CILogon IDPs from InCommon [4] and eduGAIN [5] that report R&S attributes, the XSEDE IDP, Google, ORCID, and a variety of other government and research OIDC organizations.)
- Globus Auth allows individuals to create and manage identity sets by linking their own identities from multiple IdPs.
- Globus Auth provides self-service registration for client applications, enabling them to use Globus Auth for logins and to obtain access tokens for registered OAuth2 relying parties (services) that use Globus Auth for authorization.

- Globus Auth provides self-service registration for relying parties (services), enabling them to register scopes in Globus Auth so applications can request access tokens to use with them. All GCS 5.4 collections register unique scopes with Globus Auth.
- When an application uses Globus Auth for logins, individuals may login using any identity from the same identity set and Globus Auth will return consistent OIDC identity data to the application. (New identities can be linked as part of the login flow, if desired.)
- Globus Auth provides an authentication context that applications can use to enforce strict authentication policies. The authentication context supports device and browser isolation, limits for authentication lifetimes, authentication assurance requirements (e.g., use of MFA), and required IDP authentications.

6.1.5.1 Application registration

As is the case for all OIDC services, applications and services that use Globus Auth for authentication must be registered with Globus Auth. The registration process results in a unique identity for each application and service. Globus Auth asks users for permission to share their identity data with each application or service (and any other permissions requested by the application or service) and maintains a record of these user *consents*. Globus Auth provides a user interface that allows users to review the consents they've granted and revoke any consent at any time.

The installation process for GCS 5.4 includes registration with Globus Auth. This registration is performed by the endpoint's administrator. Globus provides a self-service web interface for registration. XSEDE provides an application registration continuity service [10] to help SPs retain access to the service registration in the event of staff turnover.

6.1.5.2 Identity sets

Globus maintains an internal database of *identity sets*. Each identity set is a set of identities from various IdPs, all of which correspond to a single individual. Each identity in the set was returned from successful authentication with an external IdP. The process by which the set is formed is described below.

When an individual successfully authenticates with an external IdP, Globus examines the identity data returned by the IdP and consults its identity set database to see if the identity has been used before with Globus. If it has, Globus retrieves the identity set containing the identity and uses this identity set when returning data to the client application or service. (See "Identity data" below.) If the identity doesn't appear in any identity set, Globus allows the individual to choose: (a) to begin a new identity set with this identity as its only member, or (b) to add the identity to an existing identity set. If the individual chooses option (b), the individual must prove that he or she is the owner of the identity set—by authenticating with an identity already in the set—before the

new identity is added. If the individual chooses option (a), a new identity set is created and the current identity becomes the *primary identity* in the set.

The individual's full identity set is included as an `identity_set` claim in the identity data returned from authentication or via token introspection. This allows applications and relying parties (services) to see all of the identities included in the user's identity set. This can be used in authorization decisions: for example, it can detect that the individual has been granted access to a resource via another identity in the identity set. The `identity_set` claim is an array of identity objects, each with its own set of claims.

GCS 5.4 uses identity sets in access control decisions. Specifically, for both data access and administrative access, access is permitted if *any identity in the identity set* has been granted access.[8]

6.1.5.3 Required identities

On successful authentication, Globus Auth returns identity data to the application per the OIDC specification. Per the OIDC specification, the application is expecting a single identity, not a set of identities. Globus Auth selects an identity from the identity set for the standard OIDC claims as follows. (As described in Section 6.1.5.2, the full identity set is available via the `identity_set` claim.)

When registering with Globus Auth, applications can specify a *required identity*. The "required identity" is any identity *from a specific IdP*. For an application registered in this manner, Globus Auth will always return identity data from the specified IdP, and authentication cannot succeed unless the user has an identity from this IdP in the user's identity set.

If the application's registration doesn't specify a required identity, Globus Auth will return the identity data from the primary identity in the identity set, regardless of which identity in the set was used to authenticate. The purpose of this behavior is to consistently return the same identity to an application when a given individual logs in, while allowing the user to choose—for whatever reason—to authenticate with different IdPs on subsequent visits. This simplifies the application's code and provides a more consistent user experience.

GCS 5.4 uses the required identity feature for data access with mapped collections. Every mapped collection specifies one or more required IdPs. (For mapped collections in XSEDE, the XSEDE identity is required.) Access is granted only if the individual has an identity from one of these IdPs in his/her identity set. The matching identity is used to map the access request to a local account.

[8] This is the basis for *authorization* (access control) decisions. However, some GCS 5.4 features use the authentication session feature (Section 6.1.5.4) to impose additional authentication requirements.

6.1.5.4 Authentication context (session data)

In order to support policies that place requirements on authentication methods or history, Globus Auth provides an *authentication context.*

The authentication context is associated with a *session*. Globus Auth opens a new session whenever a user logs into an application without an open session. The session is closed whenever one of the following occurs.

- The user quits the browser.
- The user explicitly logs out of Globus Auth.
- The user re-authenticates and the resulting identity set is different from the one associated with the current session.

Applications may access the authentication context for the current session via OAuth2 token introspection. The authentication context is included as a claim in the resulting identity data. The claim contains an object of authentication events. Each event records the user identity that was authenticated, the IdP used for the authentication, and the Unix epoch time when the authentication occurred. From this authentication context, an application can enforce a number of useful policies, including the following examples.

1. The application can require that the user must authenticate with a specific IdP (or any one of a list of acceptable IdPs).
2. The application can require that the user must re-authenticate if the time since the most recent authentication is beyond a limit.
3. The application can prohibit single sign-on (authentication that took place in a different web application or browser session).

GCS 5.4 uses the authentication context for some administrative features and for data access with mapped collections. Each mapped collection has a configured authentication lifetime. (For XSEDE mapped collections, the recommended lifetime is 11 days, which is the same as XSEDE's maximum lifetime for short-term X.509 certificates.) Individuals accessing the collection must authenticate if they haven't authenticated in the current session or within the required authentication lifetime. Administrative functions impose similar requirements.

6.1.6 Identity Providers (IdPs)

Globus Auth (see Section 6.1.5) relies on external Identity Providers (IdPs) for user authentication. IdPs are also the sources of the user identity data Globus Auth returns to applications, including GCS 5.4.

An IdP is a service provided by an organization for the purpose of identifying the organization's members. An IdP provides three closely related functions.

1. It enables an individual to register with the organization, establishing an identity and obtaining credentials for later use.
2. It enables an individual to authenticate using the credentials issued at registration time.
3. It enables an individual to release his/her identity data (whatever information the organization has and is configured to release) to an application or service so the application or service can customize the user experience for that individual.

The XSEDE IdP is the main focus of this section because it is the main identity data source used in L1/L2 resource deployments of GCS 5.4. However, as described in the high-level design and detailed below, there are circumstances when other IdPs will be used for authentication and as part of authorization decisions, which is why the following sections also describe other IdPs that are used by Globus.

6.1.6.1 XSEDE

The XSEDE IdP used by Globus plays a key role in access control for XSEDE's GCS 5.4 service because it is a required IDP for access to mapped collections used in XSEDE. (See Sections 5.1.2 and 5.2.) At the time this document is being written, Globus uses an XSEDE IdP operated by the University of Chicago. By the time this design is in use, Globus will have switched to using CILogon with the XSEDE InCommon IdP (idp.xsede.org). This section describes the latter configuration.

The XSEDE IdP used by Globus is provided by the CILogon service [12] using the XSEDE InCommon IdP (idp.xsede.org) [17]. When authenticating with the XSEDE IdP, an individual must enter an XSEDE username and password, which is validated with the XSEDE Kerberos service. The individual must also authenticate with XSEDE's Duo service. Both authentication factors (Kerberos and Duo) must succeed. If the individual has not previously enrolled with XSEDE's Duo service, the individual is prompted to enroll in the service.

Identity data provided by the XSEDE IdP comes from the XSEDE Central Database (XCDB), which combines original user registration data (entered by the user) with data from the user's XSEDE user profile. Specifically, the name, email, and organization claims in identity data from the XSEDE IdP are taken from the user's XSEDE user profile. The user profile is managed by the user via the XSEDE User Portal (XUP). No external data sources are used in XSEDE IdP identity data.

The XSEDE IdP reports the user's XSEDE username without anonymization. By policy, XSEDE never re-uses usernames, so a distinct username will always refer to the same individual. It is, however, possible for an individual to have more than one assigned XSEDE username.

6.1.6.2 OpenID Connect (OIDC)

Globus Auth allows users to authenticate to a large set of OIDC IdPs. The specific OIDC IdPs supported by Globus Auth is configured by Globus. Each OIDC IdP has its own user registration, credentialing, and user authentication mechanisms.

When a Globus Auth user selects one of these IdPs for authentication, Globus Auth directs the user to the OIDC IdP for authentication and uses the resulting OIDC identity.

6.1.6.3 InCommon and eduGAIN via CILogon

CILogon [12] provides a gateway from campus SAML authentication to OIDC. CILogon is an OIDC authentication service. For user authentication, it requires users to authenticate with a SAML IdP, offering a list of several thousand InCommon [4] and eduGAIN [5] IdPs as options. On successful authentication, CILogon translates the SAML identity claims to an OIDC identity. The user identity data that results from CILogon's OIDC authentication is thus an OIDC identity translated from a SAML identity.

Globus Auth regularly imports the list of SAML IdPs supported by CILogon and selects those that release Research & Scholarship (R&S) attributes [14]. These IdPs are listed in Globus Auth's user interface. When a Globus Auth user selects one of these IdPs, Globus Auth directs the user to CILogon to authenticate with the SAML IdP and returns the resulting OIDC identity to the calling application. The SAML IdP is presented in the OIDC claims as if it were an OIDC IdP.

6.1.6.4 Built-in OIDC identity provider

GCS 5.4 includes a built-in OpenID Connect (OIDC) server. This server is configured using the GCS CLI (see Section 6.2.3) and allows authentication using usernames and passwords via the host system's PAM stack. Identities returned by the built-in OIDC server are of the form `username@domain.name`, where the `domain.name` part is the domain name of the host system. (By default, GCS 5.4 uses unique domain names in the `data.globus.org` domain, but GCS 5.4 also allows endpoint administrators to use custom domain names provisioned by their own DNS services.)

The built-in OIDC server can be used to activate a GCS 5.4 endpoint, and identities issued by the built-in OIDC server can be used for setting permissions in guest collections. But the built-in OIDC service cannot be used to login to applications via Globus Auth.

6.2 Globus Connect Server version 5.4

Section 5.1 specifies the recommended GCS 5.4 configuration for use with XSEDE SP resources, and Section 5.2 describes the resulting behaviors, including authentication and authorization. This section provides additional detail regarding GCS 5.4's behaviors and implementation.

6.2.1 Accessing a collection - Globus transfer service

Access to a GCS 5.4 collection requires an access token. The access token is obtained via authentication with Globus Auth. (See Section 6.1.5 for details on Globus Auth.) All applications and services used to access data must be registered with Globus Auth and the user must grant permission for the application to use his/her identity data and access data on his/her behalf.

For the primary data access configuration (Section 5.1.2.1), authorization to the mapped collection requires all of the following to be true.

1. The access token must be validated by Globus Auth. (This cannot happen if the token is invalid or if the individual associated with the token has revoked his or her consent.)
2. The identity set associated with the token must include an XSEDE identity. (See “Identity sets” in the Globus Auth section below.)
3. The XSEDE identity must appear in the local mapfile, mapped to a local username.
4. Local resource permissions must allow the mapped local account to have the requested access.

For the optional “guest collections” configuration described in Section 5.1.2.2 and the optional “science gateways” configuration described in Section 5.1.2.3, access to the mapped collection (including creating guest collections) is authorized exactly as described above. However, these separate storage gateways and mapped collections have separate path restrictions, account mapping settings, and allowed users and groups. The administrator has likely further restricted the paths accessible by the collection and/or the local users authorized to access the mapped collection. The administrator may also have further limited guest creation collection to specific local accounts or specific local groups.

Access to guest collections requires all of the following conditions to be true.

1. The access token must be validated by Globus Auth. (This cannot happen if the token is invalid or has been revoked by the user.)
2. The local account under which the guest collection was created must still be valid and must still have access to the data. (If either of these are no longer true, no-one will be able to access the guest collection.)
3. The user’s identity set must include an identity that is granted access by one of the ACLs configured in the guest collection by the collection manager(s).

Note that these conditions do not require the user to have a local account or to have an XSEDE identity.

6.2.2 Accessing a collection - HTTPS upload and download

Access via HTTPS uses the same access control as access via the Globus Transfer service. All requirements for mapped collections and guest collections described in Sections 5.2 and 6.2.1 apply when using HTTPS.

GCS 5.4 deploys a standard Apache web server from the system's package repository. Each collection on the system (both mapped and guest collections) has its own VirtualHost with its own configuration settings.

The `globus` Apache module enables the following behaviors:

- GridFTP control channel using the HTTPS port
- Globus Auth integration (using a modified version of the `auth_openidc` module)
- HTTPS access to collections
- automatic reloads of Apache configuration when new collections are created

A modified version of the `auth_openidc` module is used to enable OIDC authentication.

6.2.3 Administrative access

In addition to its data access interfaces, GCS 5.4 also includes a RESTful web service called the GCS Manager service that supports administrative functions. The GCS Manager service allows dynamic configuration of storage gateways and collections and synchronizes configuration changes across all data transfer nodes in the endpoint. The GCS Manager is used by two clients: the GCS CLI (installed with GCS 5.4 on the GCS host system) and the Globus Transfer service.

- The GCS CLI (`globus-connect-server` command-line tool) is used by the endpoint's administrators to configure storage gateways and mapped collections.
- The Globus Transfer service and its clients (Globus web app, Globus CLI, Globus SDK) use the GCS Manager service to enable authorized individuals (see Sections 5.1.2 and 5.2) to create guest collections.

During the base endpoint setup (see Section 5.1.1), administrator roles are assigned. Endpoint configuration changes, including storage gateway and mapped collection creation, require authorization with an identity with an administrative role. This covers all of the configuration changes made by endpoint administrators using the GCS CLI.

Guest collections are not permitted by the primary data access configuration (Section 6.2), so unless other configurations are enabled, the GCS Manager service is only usable by system administrators. If the optional guest collection or science gateway configurations are enabled, or if other configurations (to support uses other than XSEDE's), the GCS Manager's guest collection functions are used to do that. In this case, permissions for creating guest collections are governed by the configurations for storage gateways and mapped collections as described in Section 5.2.4..

6.2.4 Documentation

The primary documentation for GCS 5.4 is provided by Globus [16].

And XSEDE-specific installation guide [15] is also available and recommended for use by XSEDE SPs. (This guide references the Globus documentation for details.)

6.2.5 Standards compliance

For data transfers between itself and another endpoint, GCS 5.4 supports the GridFTP v2 data channel. GCS 5.4's control channel is not compliant with GridFTP v2 control channel because: (a) it runs on port 443 rather than GridFTP's port 2811, and (b) it uses OAuth2 access tokens and TLS for authorization rather than GridFTP's X.509 TLS.

For HTTPS uploads and downloads, GCS 5.4 uses the Apache web server provided by the host system's configured package repository, so its standards compliance depends on the Apache server version.

The Globus Transfer service provides a RESTful API (HTTP with TLS) with access authorized via OAuth 2.0 access tokens issued by Globus Auth.

For authentication, both Globus Auth and the built-in OIDC service in GCS 5.4 are OpenID Connect 1.0 compliant.

Globus Auth and the access tokens it issues for use with GCS 5.4 are OAuth 2.0 compliant.

6.2.6 Interoperability

The use cases in Section 4.1 outline the scope of this design, but they do not explicitly state interoperability requirements. To enable these use cases, there are two critical interoperability requirements.

1. GCS 5.4 must be able to transfer data to and from other Globus endpoints, including Globus Connect Personal endpoints, Globus Connect Server version 4.x endpoints, and GridFTP endpoints.
2. GCS 5.4 must be able to support file uploads and downloads with web browsers and other HTTPS clients.

In regard to the first requirement, GCS 5.4 can transfer data to and from all of the Globus endpoints mentioned. This relies on the Globus Transfer service acting as a client and using the DCAU and DCSC GridFTP protocol features to negotiate and establish a shared security context for data channel connections between GCS 5.4 and the other endpoint. As long as one endpoint involved in a transfer supports DCSC (GCS 5.4 does), the security context can be established. (GridFTP 5.x and later, and all versions of Globus Connect Server also support DCSC.)

In regard to the second requirement, GCS 5.4 uses the Apache web server in the GCS host system's package repository as its HTTPS implementation. The `auth_openidc` module is used to implement OpenID Connect (OIDC) authentication with Globus Auth. HTTPS clients must be able to supply an HTTP authorization header containing the access token for the collection being accessed. (Each collection has its own OAuth2 scope.) HTTPS clients can obtain tokens from Globus Auth as described in Section 6.1.5.

A notable interoperability feature that *is not* satisfied by GCS 5.4 is that it does not support the GridFTP v2 protocol [18] defined in 2005. This means it cannot be accessed directly by GridFTP clients using that protocol. It is important to note that GridFTP v2 clients are not commonly used today. As noted in Section 6.2, GCS 5.4 is intended to be accessed either via HTTPS (for direct upload or download) or via the Globus Transfer service (for transfers between two servers).

6.2.7 Data collection, logging, and usage reporting

GCS 5.4 supports local logging for all data accesses in all of the recommended configurations. The GridFTP server logs accesses via the Globus Transfer service, and the Apache server logs accesses via HTTPS. For mapped collection access, the mapped account is logged. For guest collection access, the account that created the guest collection is logged.[9]

The Globus Transfer service (see Section 6.1.4) also logs all transfers it manages. Transfer request log records include the source and destination endpoints, the identity that requested the transfer, the number of files, folders, and bytes transferred, the start and end times for the transfer, and the completion status of the transfer. Globus operations personnel who have access to this data and the systems in which it is stored are trained to maintain confidentiality of this data.

Globus produces and delivers monthly summaries and detailed transfer logs for endpoints associated with each Globus subscription. (XSEDE's reports and transfer log cover transfers to and from all endpoints associated with the XSEDE subscription. They *do not* cover HTTPS access, as Globus is not involved in that type of access.) These reports and logs are gathered by XSEDE XCI personnel and archived on the XSEDE XCI Metrics server.

REFERENCES

[1] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <https://www.rfc-editor.org/info/rfc6749>.

[2] OpenID Connect Core 1.0 Specification. (https://openid.net/specs/openid-connect-core-1_0.html)

[9] With a High Assurance or HIPAA+BAA subscription configured on the endpoint, an additional audit log contains the details necessary to comply with data protection rules, including the identities used for each guest collection access.

[3] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Protocol Architecture", RFC 4251, DOI 10.17487/RFC4251, January 2006, <https://www.rfc-editor.org/info/rfc4251>.

[4] InCommon Federation. (https://www.incommon.org/federation/)

[5] eduGAIN. (https://edugain.org/)

[6] Globus Connect Server Version 5 Installation Guide. (https://docs.globus.org/globus-connect-server-v5-installation-guide/)

[7] XCI-196 Design: Deliver XSEDE user to OAuth identity mappings. (https://software.xsede.org/svn/xci/activities/xci-196/trunk/Deliverables/XCI-196-Design.pdf)

[8] See the “DCSC Specification” section in the Globus Toolkit 6.0 GridFTP Developer’s Guide. (http://toolkit.globus.org/toolkit/docs/6.0/gridftp/developer/index.html#gridftp-developer-dcsc-spec)

[9] Globus Connect Server v4 Authorization and Authentication. (https://docs.globus.org/security/authorization-authentication/)

[10] See the “Continuity service for application registrations” section in XSEDE Web SSO Service Overview. (https://software.xsede.org/development/xsede-web-sso/Web-SSO-Service-Overview.pdf)

[11] See the “Sessions for Native Apps” section in Globus Auth Session Guide. (https://docs.globus.org/api/auth/sessions/#sessions-for-native-apps)

[12] CILogon. (https://www.cilogon.org/faq)

[13] “Science DMZ: Data Transfer Nodes” on the ESnet website. (https://fasterdata.es.net/science-dmz/DTN/)

[14] “Research and Scholarship Entity Category” on the REFEDs website. (https://refeds.org/category/research-and-scholarship)

[15] “*XSEDE Globus Connect Server v5.4 Installation Guide.*” (Google Doc)

[16] Globus Connect Server v5.4 Installation Guide. (https://docs.globus.org/globus-connect-server/v5.4/)

[17] “XCI-30 Design Document: Provide InCommon Identity Provider for XSEDE Identities.” (https://software.xsede.org/svn/xci/activities/xci-030/trunk/Deliverables/XSEDE-InCommon-IdP-Design.pdf)

[18] I Mandrichenko, et all, GridFTP v2 Protocol Description, May 2005. (https://www.ogf.org/documents/GFD.47.pdf)