

Campus Bridging Made Easy via Globus Services

Ian Foster
The University of Chicago and
Argonne National Laboratory
Argonne, IL 60439
+1 630-252-4619

foster@anl.gov

Steve Tuecke
The University of Chicago and
Argonne National Laboratory
Argonne, IL 60439
+1 630-252-8711

tuecke@uchicago.edu

Thomas Hauser
Research Computing
University of Colorado
Boulder, CO 80309
+1 303-735-3400

Thomas.Hauser@colorado.edu

Rajkumar Kettimuthu
The University of Chicago and
Argonne National Laboratory
Argonne, IL 60439
+1 630-252-0915

kettimut@mcs.anl.gov

Daniel Milroy
Research Computing
University of Colorado
Boulder, CO 80309
+1 303-735-4307

daniel.milroy@colorado.edu

Stuart Martin
The University of Chicago and
Argonne National Laboratory
Argonne, IL 60439
+1 630-252-0904

smartin@mcs.anl.gov

Brock Palen
University of Michigan
2281 Bonisteel Blvd
Ann Arbor, MI 48109
+1 723-936-1985

brockp@umich.edu

Jazcek Braden
Research Computing
University of Colorado
Boulder, CO 80309
+1 303-735-3886

Jazcek.Braden@colorado.edu

ABSTRACT

As science becomes more computation and data intensive, computing needs often exceed campus capacity. Thus we see a desire to scale from the local environment to other campuses, to national cyberinfrastructure providers such as XSEDE, and/or to cloud providers—in other words, to “bridge” to the wider world. But given the realities of limited resources, time, and expertise, campus bridging methods must be exceedingly easy to use: as easy, for example, as are Netflix and Amazon movie streaming services. We report here on experiences with a service called Globus Online, which seeks to do for campus bridging what Netflix and Amazon do for movies: that is, use powerful cloud-hosted services and simple, intuitive web interfaces to make it “so easy that your grandparent can do it.” Specifically, we describe Globus Transfer, which addresses the important campus bridging use case of moving or synchronizing data across institutional boundaries. We describe how Globus Transfer achieves both ease of use for researchers and ease of administration for campus IT staff. We provide technical details on the Globus solution; quantitative data on usage by more than 25 early adopter campuses; and experience reports from two early adopters, the University of Michigan and the University of Colorado Boulder.

Keywords

Globus, campus bridging, XSEDE, computer networks, data-intensive computing

(c) 2012 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

XSEDE12, July 16 – 20, 2012, Chicago, IL, USA

Copyright 2012 ACM 978-1-4503-1602-6/12/07...\$15.00.

1. INTRODUCTION

Increasingly computational and data intensive science means that even smaller projects often need sophisticated cyberinfrastructure. The ability to engage local campus, peer-institution, and national resources in a convenient and ideally seamless manner is thus increasingly critical to successful research. In other words, researchers at US universities and other institutions want tools that let them *bridge* [11] between their campuses and the national cyberinfrastructure—while IT staff at those same campuses want tools that let them enable that bridging in a manner that is easy, effective, and manageable.

Others have provided detailed analyses of campus bridging use cases [11, 14]. The following are two examples. 1) A researcher runs a code on a local cluster. As requirements grow, she wants to run on remote clusters, large-scale cyberinfrastructure providers such as XSEDE, and/or commercial or academic cloud computing resources such as Amazon Web Services and Microsoft Azure. Thus, she needs to be able to access those systems, move data to them, run computations, and move results back—ideally without having to learn an entirely new way of doing things. 2) A user runs an experiment at a specialized experimental facility such as a light source or genome sequencing facility, and then wants to return experimental data to campus for analysis.

In these and other use cases, the need is not simply to run software remotely, as is enabled by systems such as nanoHUB [8], or to download data from a remote site, as is supported by for example Protein Data Bank and Earth System Grid. Rather, the goal is:

“the seamlessly integrated use of cyberinfrastructure operated by a scientist or engineer with other cyberinfrastructure on the scientist’s campus, at other campuses, and at the regional, national, and international levels as if they were proximate to the scientist” [11].

This objective has implications for two distinct groups of people, with different skill sets, goals, and motivations. *End users* want seamless integration of local and remote resources without

learning about complex software. *Campus information technology* (IT) teams want to deliver seamless integration to their users—but because they often lack the resources to deploy and maintain sophisticated software infrastructures, they need to be able to achieve that goal with modest effort.

Researchers and practitioners have experimented with many approaches to the seamless integration that campus bridging aspires to provide: for example, distributed computing middleware (e.g., Distributed Computing Environment), grid middleware (e.g., Open Science Grid [12]), and distributed file systems (e.g., Andrew File System). While such systems have been operated successfully at scale, none is ubiquitous across campuses, and all seem to require excessive effort by campus IT.

We describe here a new approach to these problems based on the use of software-as-a-service (SaaS) methods, akin to those used today to deliver consumer and commercial IT services such as home movies, email, and accounting. With SaaS, the web browser is the computer: a user points it at the right URL and can immediately start interacting with a local or remote resource. Because the application is hosted elsewhere (“in the cloud”), complexity can be reduced enormously for the end user—and for the user’s IT staff, if indeed they have any. Users may not need to download any software at all, and if they do, it is typically just an auto-updateable agent.

We describe, more specifically, experiences with a SaaS system called Globus Online [4, 6] for campus bridging, and in particular the use of the Globus Transfer service for data movement and synchronization among sites. As the examples above show, data movement, while not always a researcher’s *only* campus bridging need, is often an *important* need, particularly as data volumes grow and science becomes more data-driven. We describe how the Globus Transfer service bridges researchers’ personal computers and campus systems with XSEDE and other resources for seamless data movement.

The rest of the paper is as follows. In Section 2, we present the Globus services that enable campus bridging. In Section 3, we discuss how this solution is used at various campuses and provide some statistics. We describe how Globus services are used for campus bridging at two campuses, University of Michigan and University of Colorado, in Sections 4 and 5. In Section 6, we describe how the consolidated support and troubleshooting enabled by Globus Online’s SaaS model helps campuses. We discuss future directions in Section 7 and summarize in Section 8.

2. GLOBUS AND CAMPUS BRIDGING

Researchers use a hierarchy of computational resources ranging from desktops to lab machines to campus clusters to national supercomputers. As they compute on these resources, they inevitably want to move data among them. There are a number of obstacles to moving data among distributed resources. Each resource has its own security domain and the researcher may often require a separate identity at each resource. Researcher desktops and campus clusters often lack sophisticated data movement tools. Transient network and system failures have to be dealt with. For these and other reasons, data movement is a frequently cited hindrance to the use of national and other off-campus resources. Researchers may even face challenges in moving large data between their desktop and campus clusters.

Globus Online provides two SaaS services to address these challenges: Globus Transfer and Globus Nexus. *Globus Transfer* is a fast, reliable file transfer service that simplifies the process of

secure data movement. It uses SaaS methods to provide easy fire-and-forget transfers, high performance data movement, and automatic fault recovery. That is, a researcher uses a convenient Web interface (or, for scripting, a command line interface) to request transfers—handing off to Globus Transfer responsibility for ensuring that the transfer completes correctly and efficiently. Globus Transfer uses Grid Security Infrastructure (GSI), which is based on X.509 certificates, to authenticate with data sources and sinks. *Globus Nexus* provides identity and group management, and supports signing on to the Globus Online ecosystem from widely adopted federated identity systems such as InCommon [5] and from OpenID providers such as Google. It also manages the multiple X.509 user credentials that are often required when accessing remote resources.

Data sources and sinks are termed “endpoints” in Globus Transfer. Globus Transfer acts as a third party agent and moves data between such endpoints. Since Globus Transfer is a cloud-hosted service, users do not have to install Globus Transfer on their machine. Globus Transfer uses GridFTP for data movement. Many national computing centers and scientific facilities have GridFTP servers installed and most are available as Globus Transfer endpoints. Most endpoints are associated with a MyProxy Online Certificate Authority (CA) [14] that issues the X.509 certificates required to access the GridFTP server. MyProxy Online CA is typically tied to the institutions’ local authentication system and issues X.509 certificates by validating the username and password (or one time password) against the local authentication system. Thus, any authorized user can immediately start moving data among a wide range of locations, via both Web and SSH interfaces—simply by registering with Globus Online and using the appropriate username and password to access the site. Globus Transfer simply passes the username and password used to access the site to the site’s MyProxy Online CA—it does not store them. For sites that do not have MyProxy Online CA and that require the use of X.509 certificates issued by widely trusted certificate authority such as DOEGrids CA, a proxy certificate can be delegated to Globus Transfer via `gssissh`.

The ability to invoke data movement and synchronization functions from a Web or SSH interface provides end users with a seamlessly integrated view of national cyberinfrastructure. However, in order for that view to extend to the campus, we must also enable creation of endpoints on personal computers for individual access and on campus clusters for multi-user access. Here we encounter the concerns of the second set of people noted above: the campus IT professionals responsible for installing software, who need that installation process to be as easy as possible. In the rest of this section, we describe Globus Transfer and describe how Globus Connect and Globus Connect Multi User address installation concerns.

2.1 Globus Transfer interfaces

Globus Transfer provides researchers with REST, Web, and command line interfaces for requesting, monitoring, and managing transfers, and for configuring the transfer environment.

Globus Transfer’s **REST interface** uses HTTP GET, PUT, POST and DELETE operations against a defined set of URLs that represent various Globus Transfer resources. Documents passed to and from HTTP requests can be formatted using either JSON or XML. Several security mechanisms are supported, including HTTPS mutual authentication with an X.509 client certificate, and HTTPS server authentication with cookie-based client authentication. The Globus Transfer **Web interface** is built entirely on top of the REST interface, using AJAX programming

techniques. A **command line interface (CLI)** is valuable for client-side scripting, but typically requires installation of client-side libraries. Globus Online avoids the need for client software installation by providing all GO users with a restricted shell, to which any user can connect via ssh to execute commands.

2.2 Globus Transfer implementation

The Globus Transfer implementation consists of a set of *user gateways* that enable users to access the system via the Web, command line and REST interfaces; a set of *workers* that orchestrate data transfers and perform other tasks, such as notifying users of changes in state; and a *profiles and state database* that maintain user profiles, request state, and endpoint information. These components are all hosted on Amazon Web Services infrastructure (see Figure 1).

Currently, Globus Transfer uses one Amazon c1.xlarge instance to run all worker processes. Based on scalability tests, we estimate that this instance size can support a total of ~200 concurrent worker processes. To prevent denial of service attacks, we enforce limits on per-user concurrent transfer activity and the total number of active transfers. The daily peak number of concurrent transfers is currently ~25; thus, we estimate we can handle roughly eight times current load under our current configuration. We use Ganglia to monitor worker CPU load and memory in order to know if we are nearing capacity. Once we reach capacity, a first step will be to migrate the Globus Transfer system to a more powerful Amazon instance type. Beyond that, we will leverage features of our design that allow additional worker instances to be added. A small amount of development remains to enable the automated expansion (and contraction) of worker instances.

Globus Transfer performs automated performance tuning for transfers, applying heuristics to set parameters based on the number and size of files in a request. It sorts files by size and performs the transfer in chunks, setting parameters for each chunk according to its average file size. If a chunk has more than 100 files and an average file size smaller than 50 MB, it uses more concurrency (transfer multiple files concurrently) and less parallelism (transfer multiple chunks of a same file concurrently). If all files in a chunk are larger than 250 MB, it uses more parallelism and moderate concurrency. In its current instantiation, Globus Transfer does not prevent many users from driving large transfers to/from the same GridFTP endpoint and/or over the same network link(s). We plan to develop tools that will allow endpoint administrators to manage bandwidth usage.




Figure 1: Globus Transfer Architecture

2.3 Globus Connect

Globus Connect solves the “last mile problem” of transferring data between a user’s personal computer and a campus cluster or a regional or national computing facility. Globus Connect is a specialized packaging of the GridFTP server binaries for

Windows, Mac OS X, and Linux that turns a personal computer into a Globus Transfer endpoint. Globus Connect can be easily installed and used by anyone (without administrative privileges) on their own computer, even if they are behind a firewall or Network Address Translation device that only allows out-bound connections. It just takes one click and one copy/paste to install Globus Connect, with no manual security configuration required.

Figure 2 shows a data transfer flow with Globus Connect as one endpoint. (Support for Globus Connect-to-Globus Connect transfers is in the works.) During November to April 2012, an average of 80 TB was moved to/from Globus Connect endpoints per month. Here are two quotes from Globus Connect users:

“I just used Globus Connect to transfer 28 GB from the Trestles cluster with a single click in a web browser at a speed of 183.3 MBits/sec. 20 minutes instead of 61 hours!”—XSEDE user, who was moving data to a poorly connected computer on which scp performed particularly badly.

“Fantastic! I have already started using Globus Connect to transfer data, and it only took me five minutes to set up. Thank you!”—NERSC user.

The fact that people are so excited about Globus Connect is due in part to Globus GridFTP’s reputation as a powerful but hard-to-use tool. The GridFTP protocol specification [3] extends the FTP protocol to provide secure, reliable, and efficient transfer of huge volumes of data across wide-area distributed networks. It can deliver orders-of-magnitude higher throughput over inefficient data transfer methods such as secure copy (scp). Many science communities rely upon Globus GridFTP [2] for their production operations. However, while GridFTP is powerful, installation and configuration complexities—particularly from a security perspective—hinder its use. Simplicity of endpoint instantiation is critical for campus bridging, as individual researchers and campuses typically lack resources and expertise to setup and maintain complex software. By addressing this concern, Globus Connect makes Globus Transfer’s high-performance, easy-to-use file transfer capabilities available to many more users and uses.




Figure 2: Globus Connect

2.4 Globus Connect Multi-User

Globus Connect Multi User (GCMU) [8] is a version of Globus Connect that creates a Globus Transfer endpoint in multi-user environments such as campus clusters, enabling easy integration of such resources into campus and national cyberinfrastructures.

As shown in Figure 3, GCMU combines a MyProxy Online Certificate Authority (CA) server [10] plus a GridFTP server configured with a custom authorization callout. When GCMU is installed, MyProxy Online CA ties to a local authentication system such as LDAP [9] via a Pluggable Authentication Module (PAM) [13] API. An interactive install script prompts the system administrator for ten simple inputs and then sets up a GSI-enabled secure Globus Transfer endpoint.




Figure 3: Globus Connect Multi-User

Once the endpoint is set up, a user can access it from Globus Transfer using the credential (username/password, OTP, etc.) that they normally use to access that cluster—providing, once again, seamless integration. Behind the scenes, an X.509 credential is obtained from the MyProxy online CA at the site and used to secure transfers. Neither endpoint administrators nor end users have to go through the process of obtaining and installing the X.509 credentials themselves. The custom authorization callout takes care of mapping an X.509 credential to a unique user at the endpoint; endpoint administrators do not have to create and maintain this mapping explicitly. Figure 4 illustrates the steps involved in the authentication process.

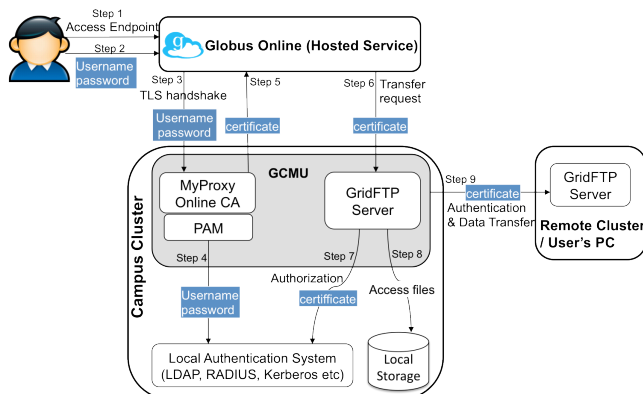


Figure 4: Interaction between Globus Online and GCMU

Here are two quotes from campus system administrators:

“With GCMU and Globus Online, GridFTP has become a critical service for us practically overnight. Globus Online is an awesome tool that is really helping our user community.”

“We have been extremely impressed with Globus Online and how easy it is to use. Now with Globus Connect Multi-User, setting up a GridFTP server and handling authentication for multiple users is equally easy. The way it ties in seamlessly with Globus Online and allows for simple user administration is fantastic.”



Figure 5: GCMU deployments as of April 2012

As can be seen from Figure 4, users provide their username/password to their GCMU endpoint via Globus Transfer. To mitigate the security concerns associated with passing credentials through a third-party site, Globus Transfer supports OAuth [7]. With an OAuth server set up on a GCMU endpoint, users do not enter a username or password on Globus Transfer. Instead, when they access a GCMU endpoint, they are redirected to a web page running on the endpoint; when they enter the username/password on that site, Globus Transfer is provided with a short-term certificate from that endpoint via the OAuth protocol.

Figure 5 shows GCMU deployments as of April 2012, which include more than 25 U.S. university campuses and have moved over 250 TB since first release in October 2011.

2.5 Support for External Identity Providers

Globus Online allows users to associate their Globus Online account with their own institutions’ identity via InCommon or external identity providers such as Google OpenID or a MyProxy server. A user can associate multiple identities with an account, but each external identity can only be associated with a single Globus account. Once a Globus Online account is associated with an external identity, users can login to Globus Online using that external identity.

InCommon support means that a user can also authenticate to Globus with a campus credential. In InCommon parlance, Globus Online is a service provider, which relies on existing InCommon identity providers. Thus, Globus Online can be integrated with other systems, such as science gateways, with single sign-on. If a user has already logged into a third party web site (e.g., a science gateway) via InCommon, the user will typically be able to log into Globus Online (e.g., to check transfer status) without retyping the password for their InCommon identity provider account.

InCommon and OpenID can be used to authenticate to Globus Online, but not for authentication between Globus Transfer and endpoints. An X.509 credential is required in order for Globus Transfer to authenticate with the endpoint and to transfer data as requested by a user. At present, this X.509 credential is handed off to Globus Transfer through the MyProxy Online CA at the endpoint or via gsissh delegation. To enable the use of InCommon or other credentials in this context, Globus Online plans to use technologies developed within the CILogon project [1]. CILogon translates an authenticated InCommon identity into a GSI-compatible X.509 certificate. If a user logs into the Globus Online using their InCommon identity, Globus Online can then use CILogon to convert that identity into an X.509 certificate that Globus Transfer can use to access endpoints on the user’s behalf.

3. USAGE STATISTICS

Globus Connect was first made available in April 2011 and GCMU in October 2011. Figure 6 shows the number of unique GC users (users that completed at least one transfer to/from a GC endpoint) per month during the period October 2011 to April 2012, and also the number of those users that completed one or more transfers to/from an XSEDE endpoint. Figure 7 shows both the total bytes transferred per month to/from GC endpoints over the same period and, again, the same data subsetted to XSEDE endpoints. Figure 8 shows the same data as Figures 6 for GCMU—and, in addition, the number of unique GCMU endpoints and the number of unique GCMU users that completed one or more transfers to/from an XSEDE endpoint. Figure 9 shows the same data as Figure 7 for GCMU.




Figure 6: Globus Connect endpoints and user counts




Figure 7: Bytes transferred with Globus Connect

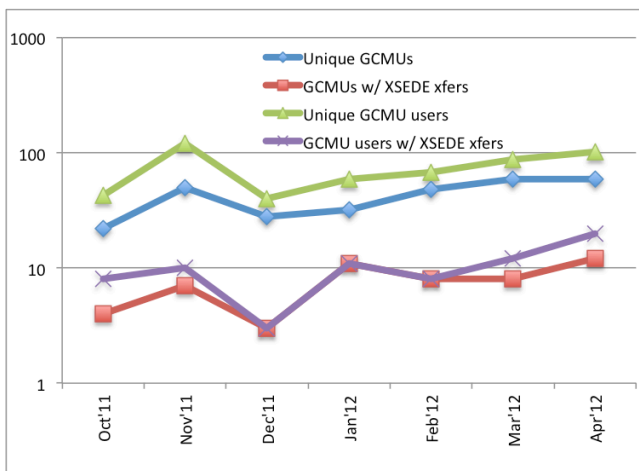



Figure 8: GCMU endpoints and users statistics

We see that both Globus Connect and GCMU play a vital role in linking end users and campuses with XSEDE and other resources. Over the past six months, ~120 TB have been moved per month using Globus Connect and GCMU. As of April 2012, an average of more than 1 TB a day is moved between GC/GCMU and XSEDE. This is only a small fraction of the total data moved via Globus Transfer, but it shows campus users are already benefiting from these tools. We see a clear trend of growing campus adoption (number of GCMU endpoints in Figure 8) and number of users (number of GC users in Figure 6 and number of

GCMU users in Figure 8).



4. Figure 9: Bytes transferred with GCMUGLOBUS BRIDGING AT U. MICHIGAN

The Computer Aided Engineering Network (CAEN) HPC Group at the University of Michigan Ann Arbor, located in the College of Engineering, operates two primary HPC systems. These systems provide traditional HPC cycles for a large group of diverse researchers as well as undergraduate education. Currently a modest staff provides resources to 133 research groups and over 500 users from all major schools and colleges at Michigan.

CAEN staff report that Globus Transfer and the existence of packaged, preconfigured Globus tools in the form of Globus GCMU and Globus Connect has enabled a modest-sized staff to bring easy-to-use, reliable, high performance data transfers to a large user base. The simple Web interface makes user education easy and scalable, an important factor for such a large population. In the opinion of CAEN staff, a similar training effort for the traditional Globus Toolkit would have required significantly larger time commitment from CAEN HPC staff.

The GCMU package is used to enable the campus resource so it may be used as a Globus endpoint. To facilitate deployments at campuses that want to set up multiple endpoints, GCMU has an option to set up the GridFTP server alone (in addition to the default option of setting up both GridFTP server and MyProxy server) and associate it with an existing MyProxy server. This is the configuration that CAEN uses: It runs a campus-wide MyProxy server (set up from GCMU) that is tied to the campus-wide Kerberos authentication service. Building on this configuration, other clusters on campus that use Michigan user IDs install just the GridFTP from the GCMU package; there is no need for those individual users or groups to set up their own MyProxy server. This simplicity has enabled the establishment of five additional GridFTP servers at Michigan, allowing transfers of files among on-campus resources as well as to and from national resources such as XSEDE and DOE supercomputers.

Globus support for third-party transfers has also allowed CAEN staff to keep data concentrated to major network pathways between end-points: with GCMU endpoints installed on all of the major data storage centers, users naturally end up moving data *directly* between remote resources (e.g., XSEDE) and campus centers, rather than moving it via their lab as used to be commonplace. At the same time, CAEN reports that researchers

who need to move data to and from personal machines find the Globus Connect package effective and easy to use.

Due to its reliability, Globus Transfer has been particularly useful for users who experience poor network conditions such as when travelling abroad or when working with unstable resources that may crash or drop off the network. Prior to the availability of Globus Transfer, these users would not have access to their data or would have little certainty about the state of their transfer after a disruption. To be confident that data integrity was maintained, users often restarted the entire transfer following a failure. This strategy added significant time to the transfer as well as unneeded stress on disk and network systems. With Globus Transfer and its reliable service, such restarts are no longer required.

The University of Michigan CAEN HPC group and its partners in the campus community such as Collage of Engineering, Literature Science and the Arts (LSA), Medical School, and School of Public Health have found the Globus Online software-as-a-service to be simple, reliable, maintainable, and useful to their efforts in the computational mission of the University. Since setting up the GCMU endpoints in October 2011, their users have moved more than 40 TB to and from these endpoints. More than 30 unique users used those endpoints during April 2012.

5. GLOBUS BRIDGING AT CU-BOULDER

The integration of computing resources, software, and networking, along with data storage, information management, and human resources to advance scholarship and research is a fundamental goal of cyberinfrastructure. The University of Colorado Boulder (CU-Boulder)'s Research Computing (RC) unit was established via a partnership between the Vice Chancellor of Research and the Associate Vice Chancellor for Information Technology and Chief Information Officer. The mission of RC is to provide leadership in developing, deploying, and operating such an integrated cyberinfrastructure.




Figure 10: Overview of the RCN on the CU-Boulder campus with connections to JILA and NSIDC.

5.1 Janus and Networking Infrastructure

RC operates compute and storage resources at three locations. The Janus supercomputer, with 1368 compute nodes and about 800 TB of high performance storage, is located in a containerized data center [15]. A research computing network (RCN) connects Janus

to storage, brings individual dedicated 10 Gbps circuits to various campus locations, and provides a 10 Gbps link between Janus and storage at the National Center for Atmospheric Research (NCAR). CU-Boulder participates in Internet 2 and is an active member of the Front Range GigaPoP (FRGP). Figure 10 provides an overview of the RCN, showing the connections between the Janus supercomputer, two data centers on the CU-Boulder campus, and the connectivity to the outside world. Research Computing offers private VLANs to institutes participating in the RCN.

5.2 Demand for Globus Connect Multi-User

The deployment of Janus and related cyberinfrastructure at CU-Boulder introduced the challenge of moving terabyte-scale datasets between Janus and large-scale remote computing facilities such as NCAR and TACC. Prior to the development of GCMU, traditional data transfer mechanisms were considered as solutions. Legacy options such as rsync and scp are viable for small transfers, but the probability of lost connections or timeouts increases with connection time, becoming unsupportable at terabyte order. An unsophisticated but effective alternative is to ship disks between facilities. The speed of GridFTP obviates physical disk transfer, but the learning curve is high for users not thoroughly versed in the syntax and semantics of globus-url-copy. From an administrative standpoint, installation and configuration of GridFTP can be a demanding process.

RC staff report that Globus Transfer and GCMU address these issues, offering drag-and-drop simplicity for all users coupled with the speed of GridFTP and comparative ease of installation. Finally, RC does not support SSH key authentication, necessitating multiple One Time Password uses to complete transfers. Globus Transfer's credential-based authentication allows users to activate endpoints for long periods of time, enabling automation of file transfers.

5.3 RC GCMU Configuration and Attributes

RC offers its users the resources to move data across public networks as well as through private VLANs within the RCN. We summarize here the modifications that were provided to enable this capability with GCMU and report endpoint usage statistics. We also relate RC's findings regarding the relative performance of Globus Transfer automated protocol configuration and manually configured parameters via the CLI "transfer" command.

5.3.1 Public VLAN Endpoint - RC Environment

RC designated four Dell PowerEdge R710 rack servers with 10 Gb NICs as GridFTP hosts. Each node has two Intel Xeon X5660 CPUs and 48 GB of memory. To facilitate transfers to Janus, RC created a public logical endpoint on Globus Transfer through the CLI server. Colorado#gridftp is a composite of four physical nodes running Globus Connect Multi User, with one primary node functioning as a MyProxy CA server. All that is necessary to aggregate the nodes into the logical endpoint is to add them sequentially by FQDN, specifying the subject Distinguished Name of each node's certificate. The endpoint must be modified to map MyProxy authentication requests to the physical MyProxy CA host. Globus Transfer then performs server load balancing and modulates transfer parallelism, concurrency, and pipeline depth. User authentication through Globus Transfer integrates automatically with the CU-Boulder OTP protocol.

5.3.2 Private VLAN Endpoints - JILA, NSIDC

Two institutes at CU-Boulder need to execute transfers entirely within the campus research network. This scenario presents the difficulty of distributing the control and data channel connections

across different network interfaces: the Globus Transfer control channel connection must be established through the host's gateway device, while the data channel is opened on a private VLAN. This organization can be accomplished by editing GCMU's configuration.

Since GCMU runs under xinetd, RC staff created a new gridftp-go configuration file for each instance. By default, Globus Transfer uses port 2811 for control channel connections, and 7512 for MyProxy CA connections. If several instances of GCMU are running on a node, each must have a unique control port. RC can use one instance of MyProxy CA on the primary node to authenticate other GCMU nodes; the MyProxy CA port is left unaltered. To make a new GCMU instance, the existing gridftp-go file in /etc/xinetd.d/ is copied and the control port ("port") altered. To direct data to a different VLAN, --data-interface <interface IP address> is appended to server_args in the configuration file. Adding physical hosts to the logical endpoint introduces the additional step of specifying the control channel port. This can be accomplished via the CLI, or through a web browser with Globus Transfer's endpoint management functionality.

5.3.3 Statistics and Performance Testing

Table 1 presents cumulative statistics for service usage and maximum observed transfer rates from September 2011 until April 2012. This data is obtained from Globus Transfer logs.

Table 1: Transfer data for Colorado endpoints

Data transferred from colorado#gridftp	102.8 TB
Data transferred to colorado#gridftp	17.6 TB
Peak transfer rate between distinct endpoints	2.9 Gb/s
Peak transfer rate to/from Janus (disk)	5.9 Gb/s
Peak transfer rate to/from Janus (memory)	9.5 Gb/s

Table 2: Comparison of Globus Transfer performance when using either auto tuning or the CLI "transfer" command with manually selected parameters (all "cc 4 p 4 -pp 4" except for that labeled +, "cc 4 p 4")—with and without jumbo frames

Test	Protocol	MTU	Destination	Rate (Mb/s)
1	auto	9000	colorado#jila	4720
	manual	9000	colorado#jila	5924
2	auto	9000	colorado#gridftp	3025
	manual	9000	colorado#gridftp	4090
3	auto	9000	colorado#nsidc	2736
	manual	9000	colorado#nsidc	3014
4	auto	1500	colorado#jila	4799
	manual	1500	colorado#jila	6225
5	auto	9000	colorado#nsidc	4190
	manual	9000	colorado#nsidc	7881
6	auto	1500	colorado#jila	2315
	manual	1500	colorado#jila	5804
7	auto	1500	JILA hyperion	1307
	manual	1500	JILA hyperion	1758

8	auto	9000	JILA hyperion	1503
	manual	9000	JILA hyperion	1836
	manual+	9000	JILA hyperion	1951
9	auto	1500	NCAR	1269
	manual	1500	NCAR	1933

In an effort to understand transfer rates between Janus and other computing facilities, RC staff compared performance of Globus Transfer's Web API, which invokes transfers with automatically selected parameters, and that of the "transfer" command executed via SSH on cli.globusonline.org with manually selected parameters. Table 2 lists characteristic transfers. Test files are 100 x 1 GB blocks of zeros for tests 1-6 and 9 and 28 x 1 GB blocks of zeros for tests 7 and 8, with all files on the Janus Lustre file system. Each trial was performed successively within a test, meaning that in each case the GO API was tested, immediately followed by "transfer." Transfer rates are those reported by Globus Transfer. The observed average increase in transfer rate for all tests between the Globus Transfer API (auto tuning) and "transfer" command via cli.globusonline.org (manual tuning) is 66% for MTU 1500, and 27% for MTU 9000. Excluding internal testing (e.g., Janus to Janus) the observed average increase is 44% for default packets and 26% for jumbo frames. These results demonstrate the value of jumbo frames. They also show that there can be significant benefit to manually configuring transfer parameters—suggesting that the Globus Online team has more work to do on their automated configurations.

5.4 Future Directions at CU-Boulder

Jumbo frames were tested and placed into production in April, resulting in a maximal increase in test transfer rate of 70%. The tests transferred 100 x 1 GB files from colorado#gridftp to colorado#gridftp using the CLI "transfer" command. While the test differs significantly from user transfers between distinct endpoints, taken with results reported by NCAR it indicates that enabling jumbo frames will result in increased transfer rates between CU-Boulder endpoints and those supporting MTU 9000. With jumbo frames enabled, NCAR experienced transfer rates of greater than 2.6 Gbps from Janus to NCAR's Boulder cluster. This test preceded a series of transfer timeouts reported by users in late April. The problem was localized to the data channel connection over RC's public VLAN, making efforts to detect packet loss due to attempted fragmentation prohibitively time consuming for devices beyond the purview of the CU-Boulder network. Each appliance connecting our GridFTP nodes to the outside world was tested to confirm support for jumbo frames. Data channels for colorado#jila and colorado#nsidc are under the aegis of the Research Computing; both have MTU 9000 enabled. Research Computing wishes to thank the Globus Online support team for his assistance in identifying this problem.

RC intends to create a new endpoint to transfer data across a public VLAN with jumbo frames enabled. Users will be advised that connections may time out if appliances external to CU-Boulder do not support jumbo frames. In this case a user may revert to colorado#gridftp for default packet size.

6. TROUBLESHOOTING

In addition to providing easy-to-use tools for bridging campuses and end users to the Globus Online ecosystem, Globus Online's SaaS model incorporates consolidated support and troubleshooting. Proactive monitoring of transfers by Globus team

members helps identify and resolve problems rapidly. When transfers fail due to endpoint errors, the Globus support team notifies the endpoint administrators of the problem and often helps resolve the problem before the end user notices the failures.

For example, in one recent event, transfers from an XSEDE endpoint to a Globus Connect endpoint on a campus system and from a supercomputing facility to the same XSEDE endpoint failed with an “end of file” error every time. After verifying that it was not a firewall problem and that transfers between these three endpoints and other endpoints worked fine, it took a team of people involving Globus developers, endpoint administrators, and network engineers a few days to identify the problem. It turned out to be a jumbo frame fragmentation problem at a single router. The team identified the problem as a jumbo issue by sending a smaller size data and a bigger size data with netcat. Then they used traceroute and piecewise ping to identify the router that was causing this problem. This problem would surely have taken longer to detect and correct if the user had been using a standalone data transfer tool. Furthermore, this experience then allowed the Globus Online team to identify quickly the same problem when it occurred in other situations. The team helped identify the jumbo frame issue at CU-Boulder (see Section 5.2.4) and in another instance—soon after the team noticed transfer failures.

7. FUTURE DIRECTIONS

The Globus Online team plans to include an OAuth server in the GCMU package and make it trivial for campus administrators to set up the OAuth server. This feature will allow any site to ensure that its users provide their credentials only on a Web page run by the site and not to any third party agents. As mentioned in Section 2.3, the team is also adding InCommon support in Globus Online, via CILogon, for authentication with GridFTP servers. This will enable campuses embracing InCommon to provide a single identity for their users to transfer data to/from their campus.

Even though CILogon provides a GSI-compatible X.509 credential to authenticate with the endpoint, there are challenges in binding an X.509 identity to a campus local identity, especially for campus clusters that use campus identities. The challenges are not due to inherent limitations in Globus Transfer or GCMU but to the fact that in many campuses, not all campus clusters use a (single) campus identity for providing access to users. The Globus Online team is enhancing GCMU to allow users to link their InCommon identity to their local account on a campus cluster.

8. SUMMARY

We have presented experiences implementing campus bridging with Globus tools. Specifically, we described Globus Transfer, which uses software-as-a-service methods to address the important use case of moving or synchronizing data across institutions. We also described how Globus Connect and Globus Connect Multi-User provide both ease of use for researchers and ease of administration for campus IT staff, and let researchers bridge between campuses and national cyberinfrastructure and other external resources. We presented quantitative data on usage by the more than 25 early adopter campuses. We also presented experience reports from two early adopters, the University of Michigan and the University of Colorado Boulder.

ACKNOWLEDGMENTS

We thank the Globus Online team for their work, Craig Stewart and others involved in the XSEDE Architecture and Design activity for their insights into campus bridging, and John Hanks for help setting up the infrastructure at University of Colorado at Boulder. This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-06CH11357; the National Science Foundation, under contract OCI-534113; and the National Institutes of Health, under NCRR grant “the Biomedical Informatics Research Network” (1 U24 RR025736-01).

REFERENCES

1. CILogon Service. [Accessed April 1, 2012]; Available from: www.cilogon.org/service.
2. Allcock, B., Bresnahan, J., Kettimuthu, R., Link, M., Dumitrescu, C., Raicu, I. and Foster, I., The Globus Striped GridFTP Framework and Server. SC'2005, 2005.
3. Allcock, W. GridFTP: Protocol Extensions to FTP for the Grid. GFD-R-P.020, Global Grid Forum, 2003.
4. Allen, B., Bresnahan, J., Childers, L., Foster, I., Kandaswamy, G., Kettimuthu, R., Kordas, J., Link, M., Martin, S., Pickett, K. and Tuecke, S. Software as a Service for Data Scientists. *Communications of the ACM*, 55(2):81-88, 2012.
5. Barnett, W., Welch, V., Walsh, A. and Stewart, C.A. A Roadmap for Using NSF Cyberinfrastructure with InCommon, <http://hdl.handle.net/2022/13024>, 2011.
6. Foster, I. Globus Online: Accelerating and democratizing science through cloud-based services. *IEEE Internet Computing*(May/June):70-73, 2011.
7. Hammer-Lahav, E. The OAuth 1.0 Protocol. Internet Engineering Task Force (IETF) RFC 5849, 2010.
8. Kettimuthu, R., Lacinski, L., Link, M., Pickett, K., Tuecke, S. and Foster, I., Instant GridFTP. 9th Workshop on High Performance Grid and Cloud Computing, 2012.
9. Koutsonikola, V. and Vakali, A. LDAP: Framework, Practices, and Trends. *IEEE Internet Computing*, 8(5):66-72, 2004.
10. Novotny, J., Tuecke, S. and Welch, V., An Online Credential Repository for the Grid: MyProxy. 10th IEEE International Symposium on High Performance Distributed Computing, San Francisco, 2001, IEEE Computer Society Press.
11. NSF Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging Final Report, http://www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf or <http://pti.iu.edu/campusbridging/> March 2011.
12. Pordes, R., Petravick, D., Kramer, B., Olson, D., Livny, M., Roy, A., Avery, P., Blackburn, K., Wenaus, T., Würthwein, F., Foster, I., Gardner, R., Wilde, M., Blatecky, A., McGee, J. and Quick, R., The Open Science Grid. Scientific Discovery through Advanced Computing (SciDAC) Conference, 2007.
13. Samar, V. and Schemers, R. Unified Login with Pluggable Authentication Modules (PAM). OSF RFC 86.0, 1995.
14. Stewart, C., Knepper, R., Ferguson, J., Bachmann, F., Foster, I., Grimshaw, A., Hazlewood, V. and Lifka, D., What is Campus Bridging and what is XSEDE doing about it? XSEDE, Chicago, IL, USA, 2012, ACM Press.
15. Tufo, H.M., Patterson, M.K., Oberg, M., Woitaszek, M., Cobb, G., Strong, R. and Gutowski, J., Janus: Co-designing HPC systems and facilities. SC'11, Seattle, WA, USA, 2011, ACM.