

# Accelerating Data Movement To Support Lattice QCD Computations

## CHALLENGE

File transfer is both a critical and frustrating aspect of compute-intensive research such as Lattice quantum chromodynamics. For a relatively mundane task, moving terabytes of data reliably and efficiently can be surprisingly complicated (one must deal with determining available protocols, negotiating firewalls and authentication, pre-staging files for transfer, detecting and responding to failures, diagnosing network misconfigurations, configuring software, etc.). Often, Lattice scientists use supercomputers where time has been allocated (such as TeraGrid or DOE machines) – but use of these facilities is costly and can be time-consuming, since scientists must deal with managing file transfers to and from these machines based on wherever computing time happens to be available.

*“Globus Online frees up my time to do more creative work rather than typing scp commands or devising scripts to initiate and monitor progress.”*

If file transfer could be accelerated and simplified, scientists would benefit not only from time and cost savings but also from the elimination of tedious, repetitive tasks that take their attention away from their core research.

## SYSTEM OVERVIEW

Globus Online is a hosted file transfer system that provides simple, reliable data movement for virtually any researcher and facility.

Globus Online makes high performance file transfer capabilities – traditionally available only on expensive, special-purpose software systems – accessible to any researcher with an Internet connection and a laptop. Users simply sign up for the service, log in, access their data using existing credentials, and then click or issue a simple command to transfer files.

### Key features:

- Reliability through automatic fault recovery and integrity checking
- Integrated management of transfers across multiple security domains with multiple user identities
- Optimized performance by auto-tuning transfers based on number of files and file sizes
- Performance monitoring, fine-grained activity logging, and status reporting
- Conditional file synchronization based on attributes such as size, timestamp, or checksum mismatch
- Simplified transfers from mass storage using pre-staging of data and automatic retries
- Transfer from machines behind firewalls and NATs using a lightweight client
- Simple creation and maintenance of personal endpoints using Globus Connect
- New features immediately available to scientists as a result of Software from Service approach

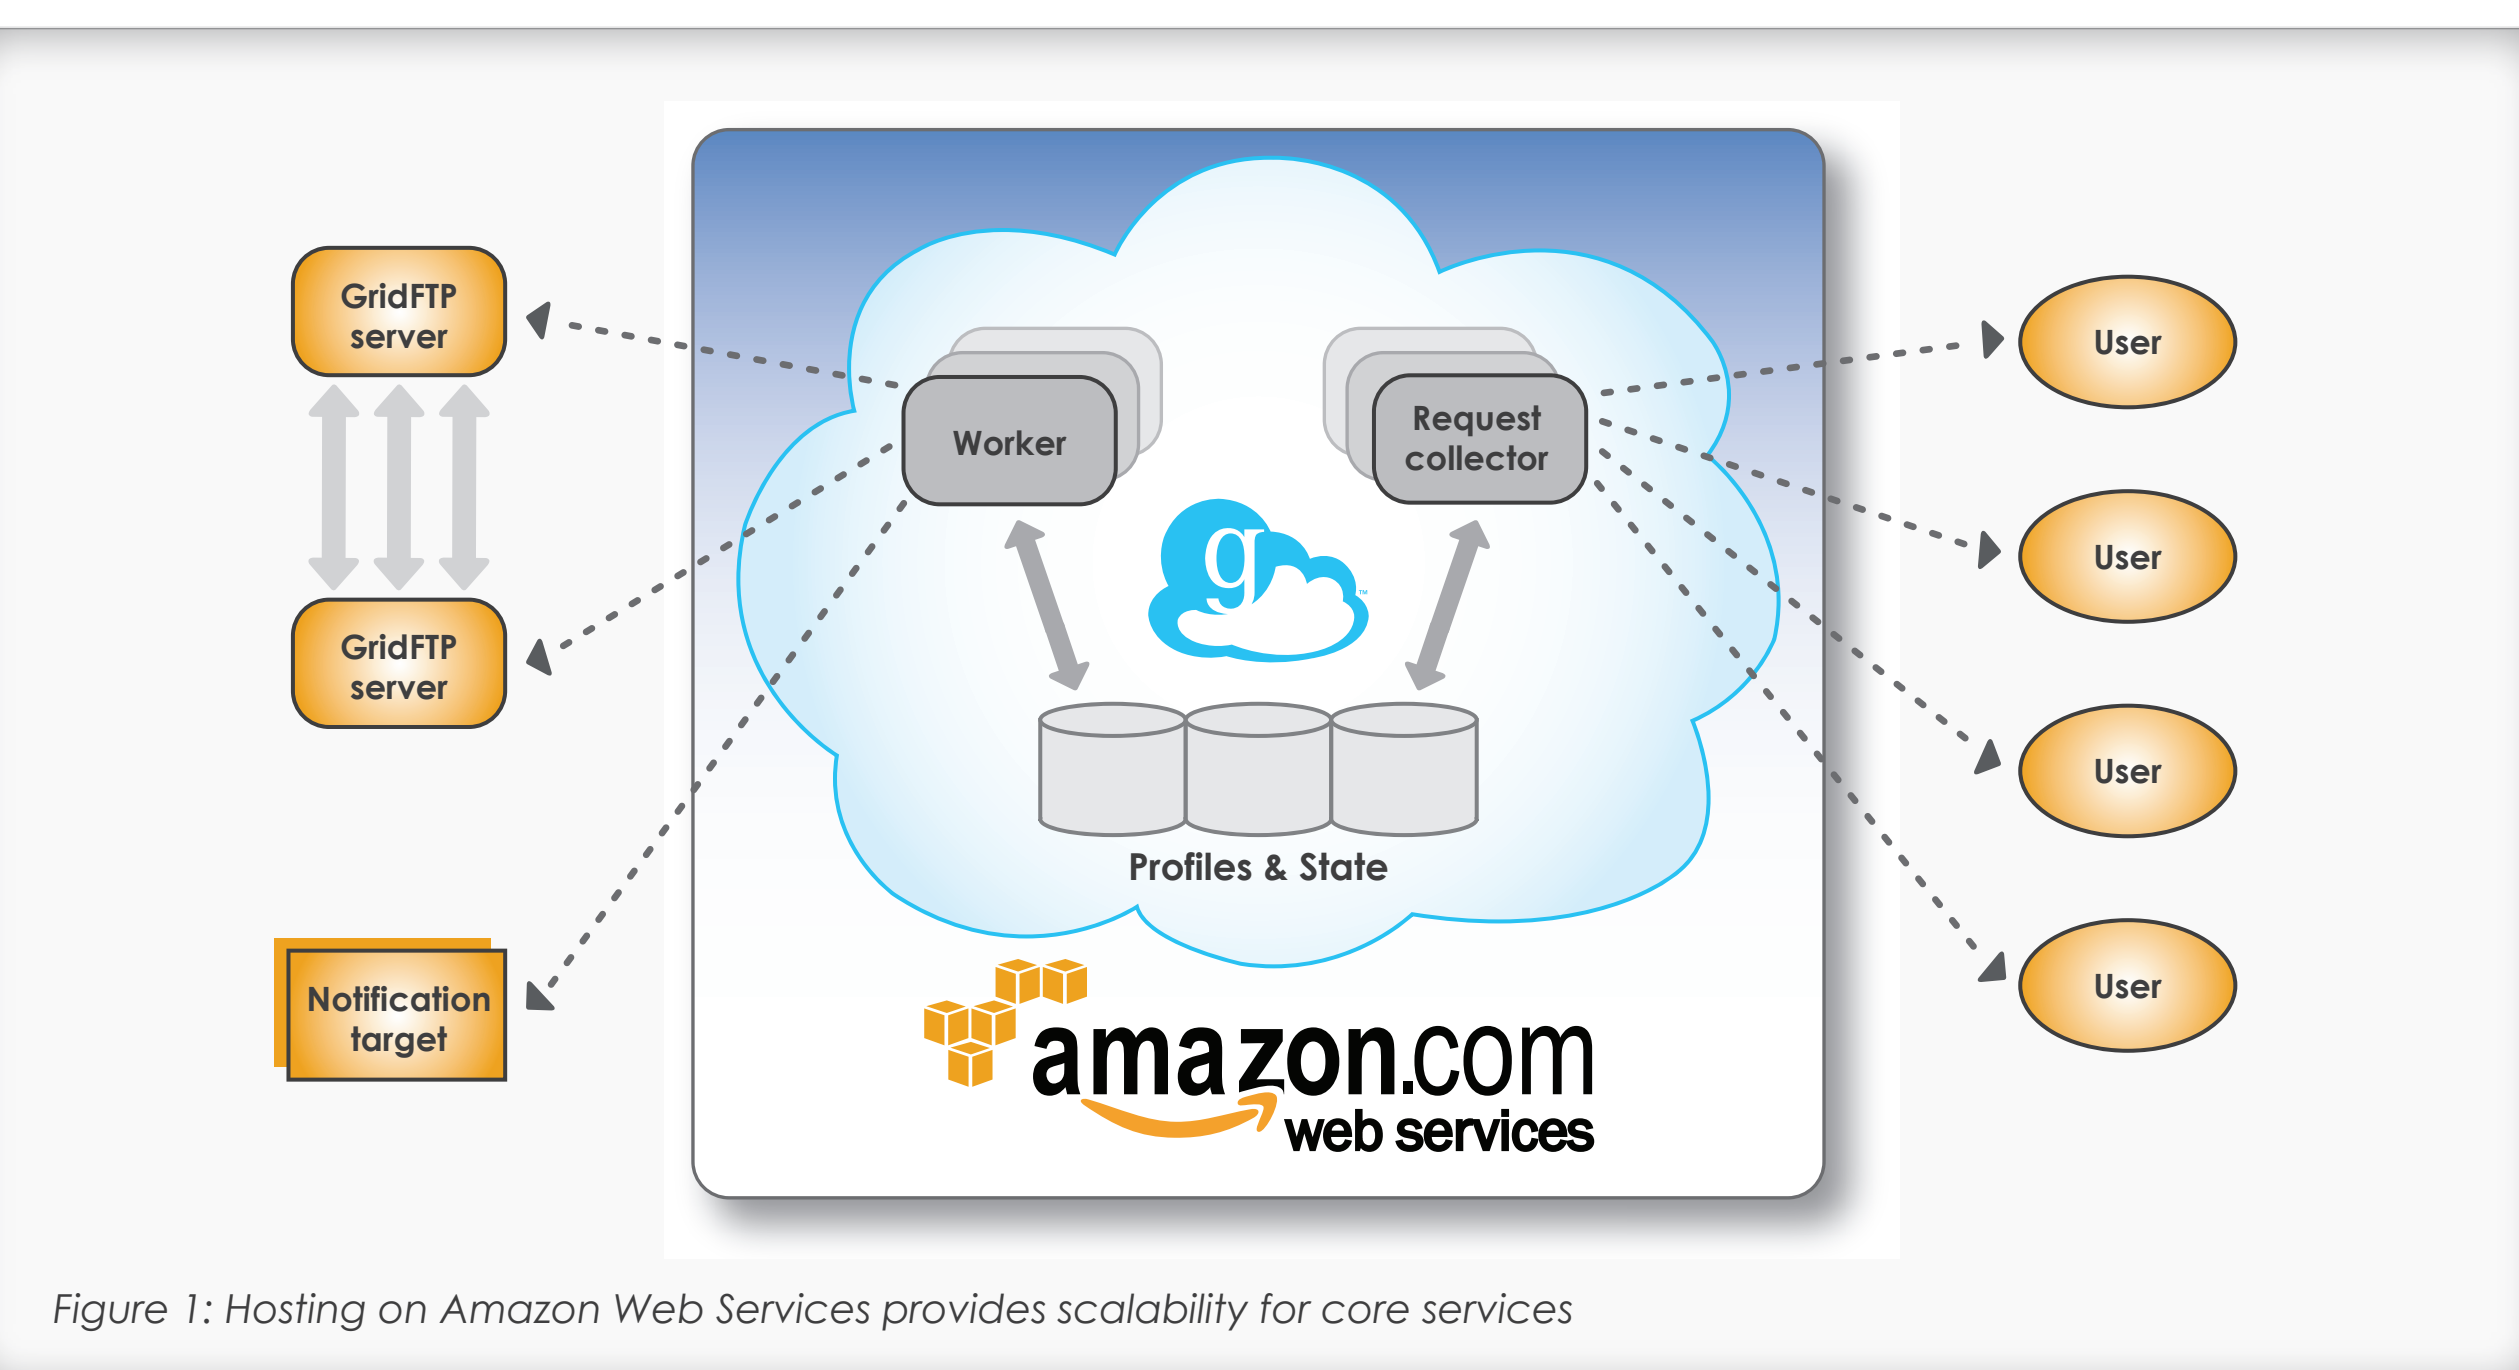


Figure 1: Hosting on Amazon Web Services provides scalability for core services

## USER SCENARIO: LATTICE QCD DATA MOVEMENT

Scientists at the MIMD Lattice Computation (MLC) Research Collaboration are using Globus Online to simplify this process of moving large files around to different supercomputing centers (including TeraGrid machines like Kraken, Athena, Lincoln and Longhorn; or DOE computers like Franklin, Intrepid and Hopper) on a frequent basis. According to Steven Gottlieb, Distinguished Professor, Indiana University and a senior member of the MLC Collaboration, the resulting speed up has been “very impressive,” with (on one occasion) one hundred 7-gigabyte files transferred in just 90 minutes (a process that would have take over 3 days with scp). The reduction in delays and manual intervention has “made a big difference” in the convenience of moving projects between centers.

*“Globus Online is the most beneficial grid technology I have even seen.”*

## ARCHITECTURE

Globus Online comprises:

- A set of user gateways, which exist to provide interaction between users and the system via the Web, CLI and REST interfaces
- A set of one or more workers, which exist to orchestrate data transfers and perform other tasks, such as notify users of changes in state
- A profiles and state database used to maintain user profiles and file transfer state information.

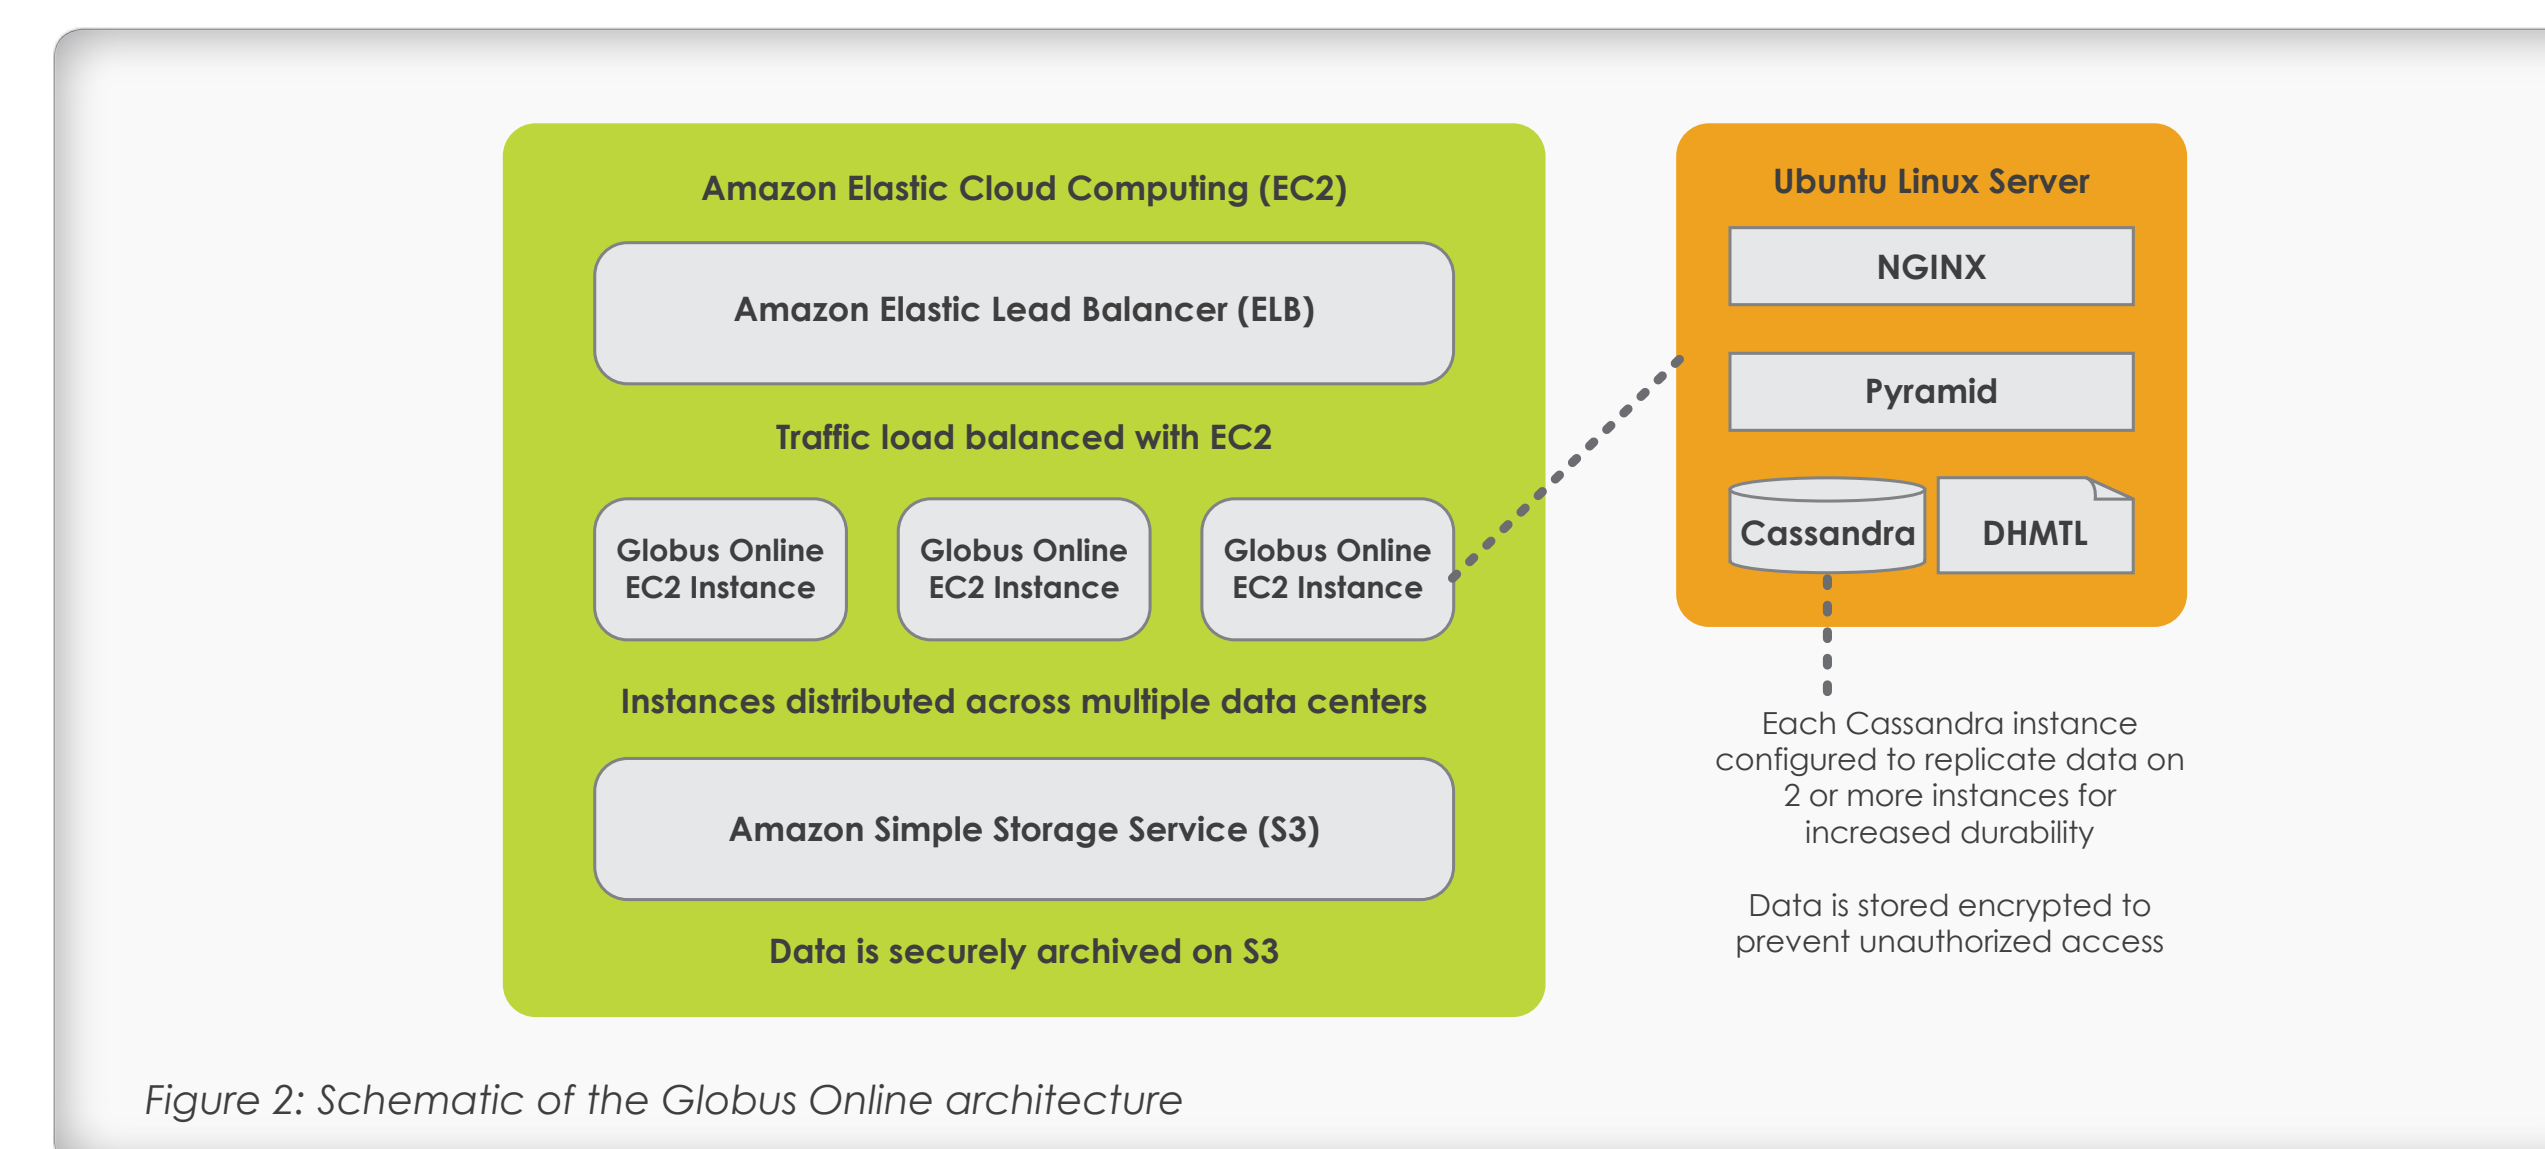


Figure 2: Schematic of the Globus Online architecture

Personal (local) machines are connected to Globus Online using Globus Connect, a one-click download-and-install application for MacOS, Linux, and Windows. Globus Connect comprises a GridFTP server that runs as a normal user (rather than as root from inetd like a typical GridFTP server), and a gsi-openssh client configured to establish an authenticated connection to a Globus Online relay server in order to tunnel the GridFTP control channel requests from Globus Online. Globus Connect only establishes outbound connections and thus can work behind a firewall or other network interface device that does not allow for inbound connections.

### • Front-end architecture and performance

- Web rendering engine and supporting core services are developed in Python using the Pyramid web service stack
- Globus Online services are lightweight, highly scalable services that present RESTful interfaces, leveraging HTTP semantics for ease-of-use, performance and scalability
- User data is stored in Cassandra, a highly scalable, second generation distributed database originally developed by Facebook and now supported as an Apache project
- URL rules and proxy logic are handled by NGINX, a high-performance HTTP server and reverse proxy
- Web pages are defined with a simple, yet powerful declarative model that enables assembly to be performed client side or server side
- Web widgets are developed in pure DHTML. No presentation logic is performed server side

### • Security and robustness

- All website traffic is handled over HTTPS to ensure security of sensitive data and cookies
- Support for fine-grained access control to user information
- All user profile information, including personal data and security credentials, is strongly encrypted in storage to further protect against unauthorized access

### • Reliability and scalability

- Globus Online is hosted on Amazon’s Elastic Compute Cloud (EC2) for optimum performance and scalability
- Core services are hosted on three separate Amazon data centers for redundancy
- Globus Online easily scales with the addition of new EC2 instances
- Traffic to Website is load-balanced by Amazon Elastic Load Balancer (ELB) for scalable reliability

## PERFORMANCE

Below are results from one Lattice researcher’s use of Globus Online over the past few months:

- 1.20.11: Moved 100 7.3-GB files in 88 minutes
  - Improved transfer rate from 20 Mb/sec to 1100 Mb/sec
  - Same transfer would have taken over 3 days with scp
- 2.8.11: Moved 108 GB in 35 minutes
- 4.23.11: Moved ~1.5 TB in 6 hours
- 5.19.11: Moved 109 GB in 22 minutes

To demonstrate how well Globus Online’s performance optimization logic performs in practical situations, we compared Globus Online performance with that achieved when using scp and the globus-url-copy (GUC) client. Figure 3 shows the performance of various transfer mechanisms for transfers over a high-speed wide area network (ESNet) between two high-performance parallel storage systems at ALCF and NERSC. It has to be noted that the performance of Globus Online is better than or similar to the performance of tuned-GUC (tuned by a data movement expert). Actually, the comparison with the untuned GUC and SCP shows the performance gains that many users can expect to gain from Globus Online.

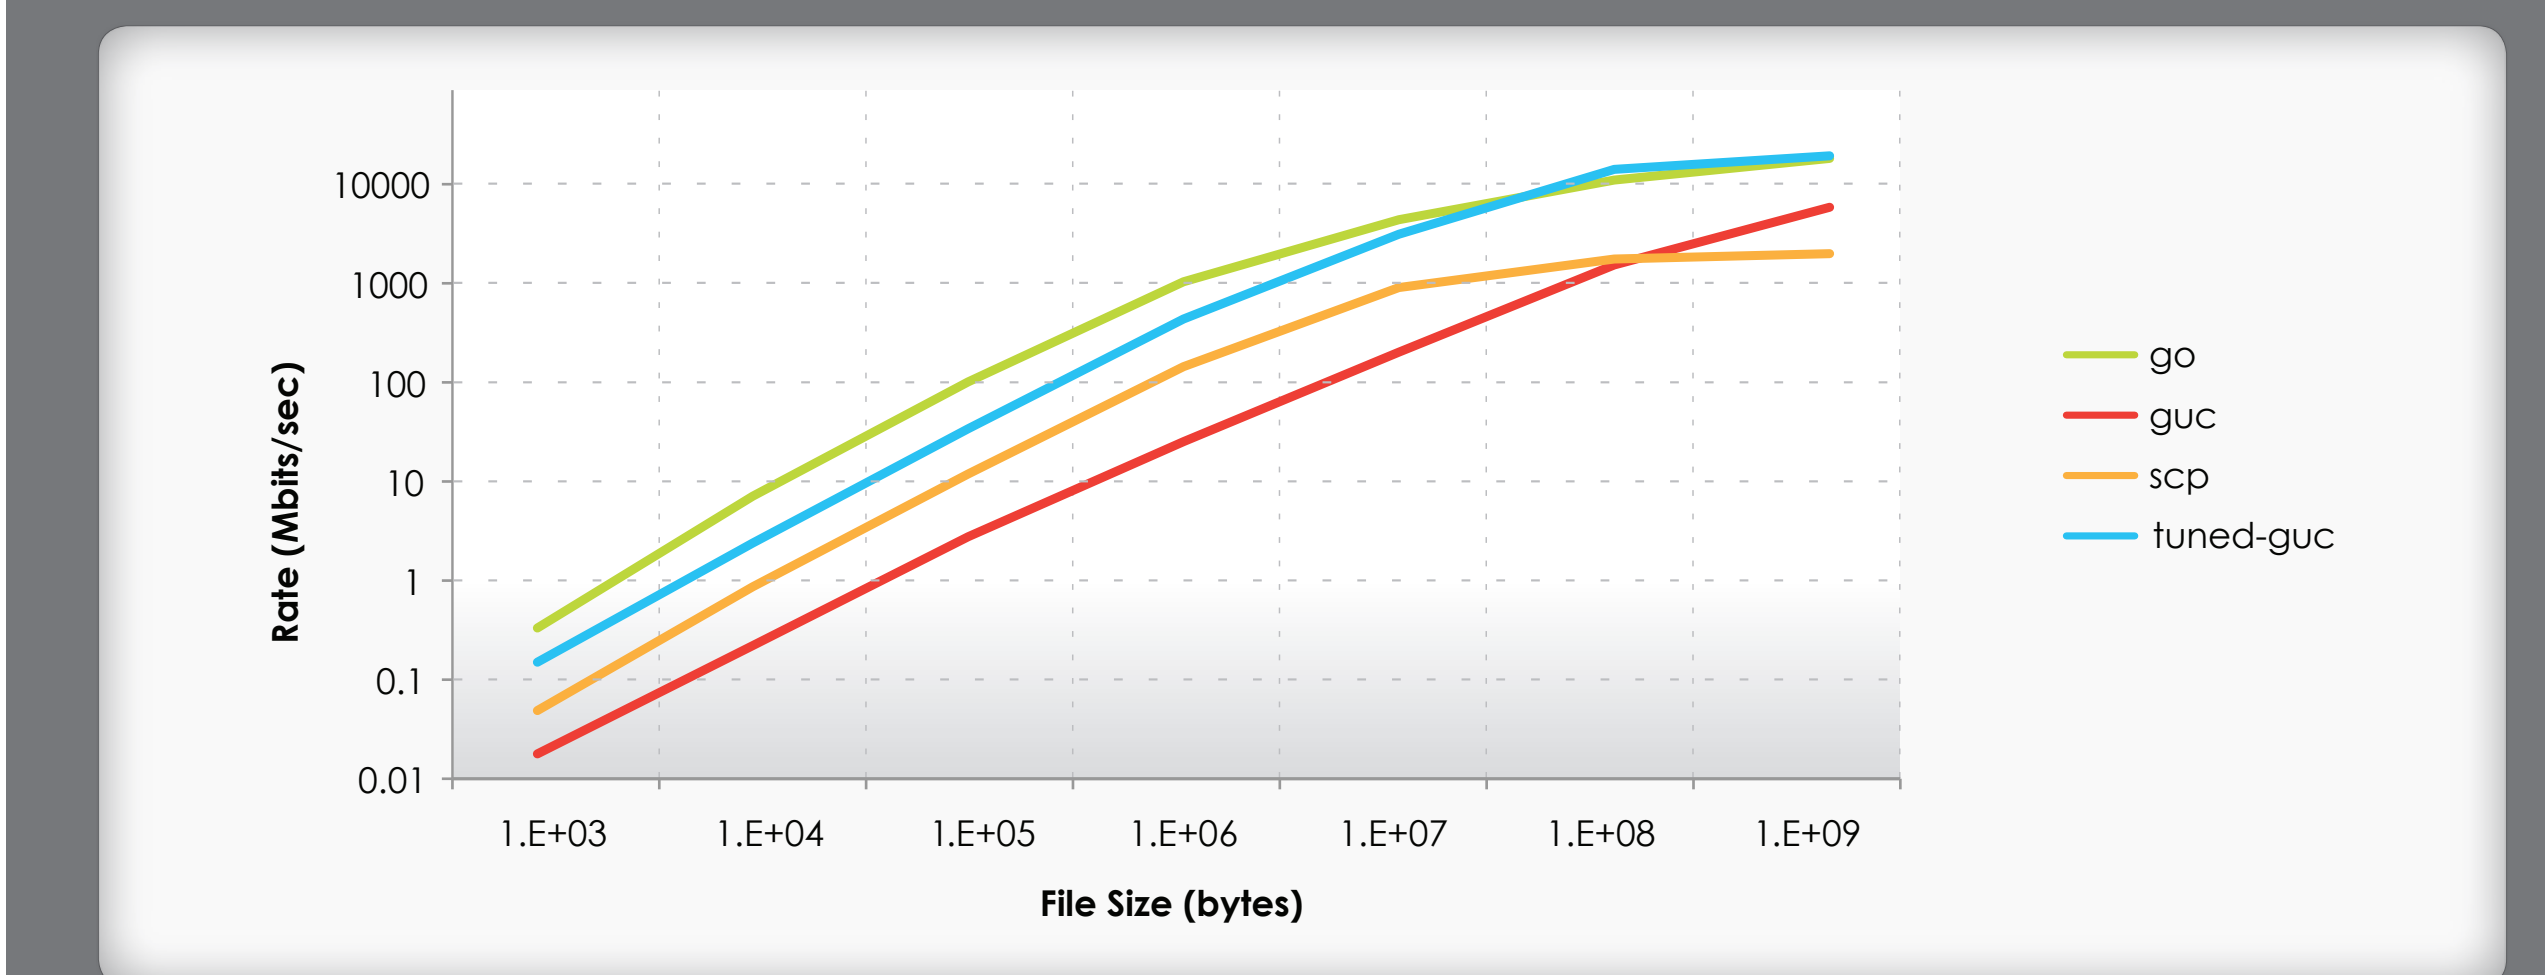


Figure 3: Data transfer performance between ALCF and NERSC for various file sizes and transfer technologies

## INTERFACES

Globus Online features a web interface, a command-line interface, and a REST API for moving files securely and reliably:

**Web UI:** Using just a web browser, files or entire directories from one endpoint can be selected and transferred to another endpoint.

**CLI Commands:** For those more comfortable at a command prompt, the ‘scp’ and ‘transfer’ commands provide additional options for creating and controlling file transfer requests. You may use a familiar scp-like interface for moving files between two endpoints, or take advantage of the GlobusOnline transfer command for increased flexibility in specifying a list of individual files or directories for transfer between endpoints.

**Transfer REST API:** This “RESTful” Globus Online interface is designed for developers who need to deliver solutions to a community of users or integrate to an external website or application. The Transfer REST API makes it possible to integrate reliable data movement into users’ HPC workflows, such as integrating with Java clients or Web-based portals, with no requirement for specialized software.

*“I moved 100 7.3 GB files in about 1.5 hours. The same transfer would have taken over 3 days with scp.”*

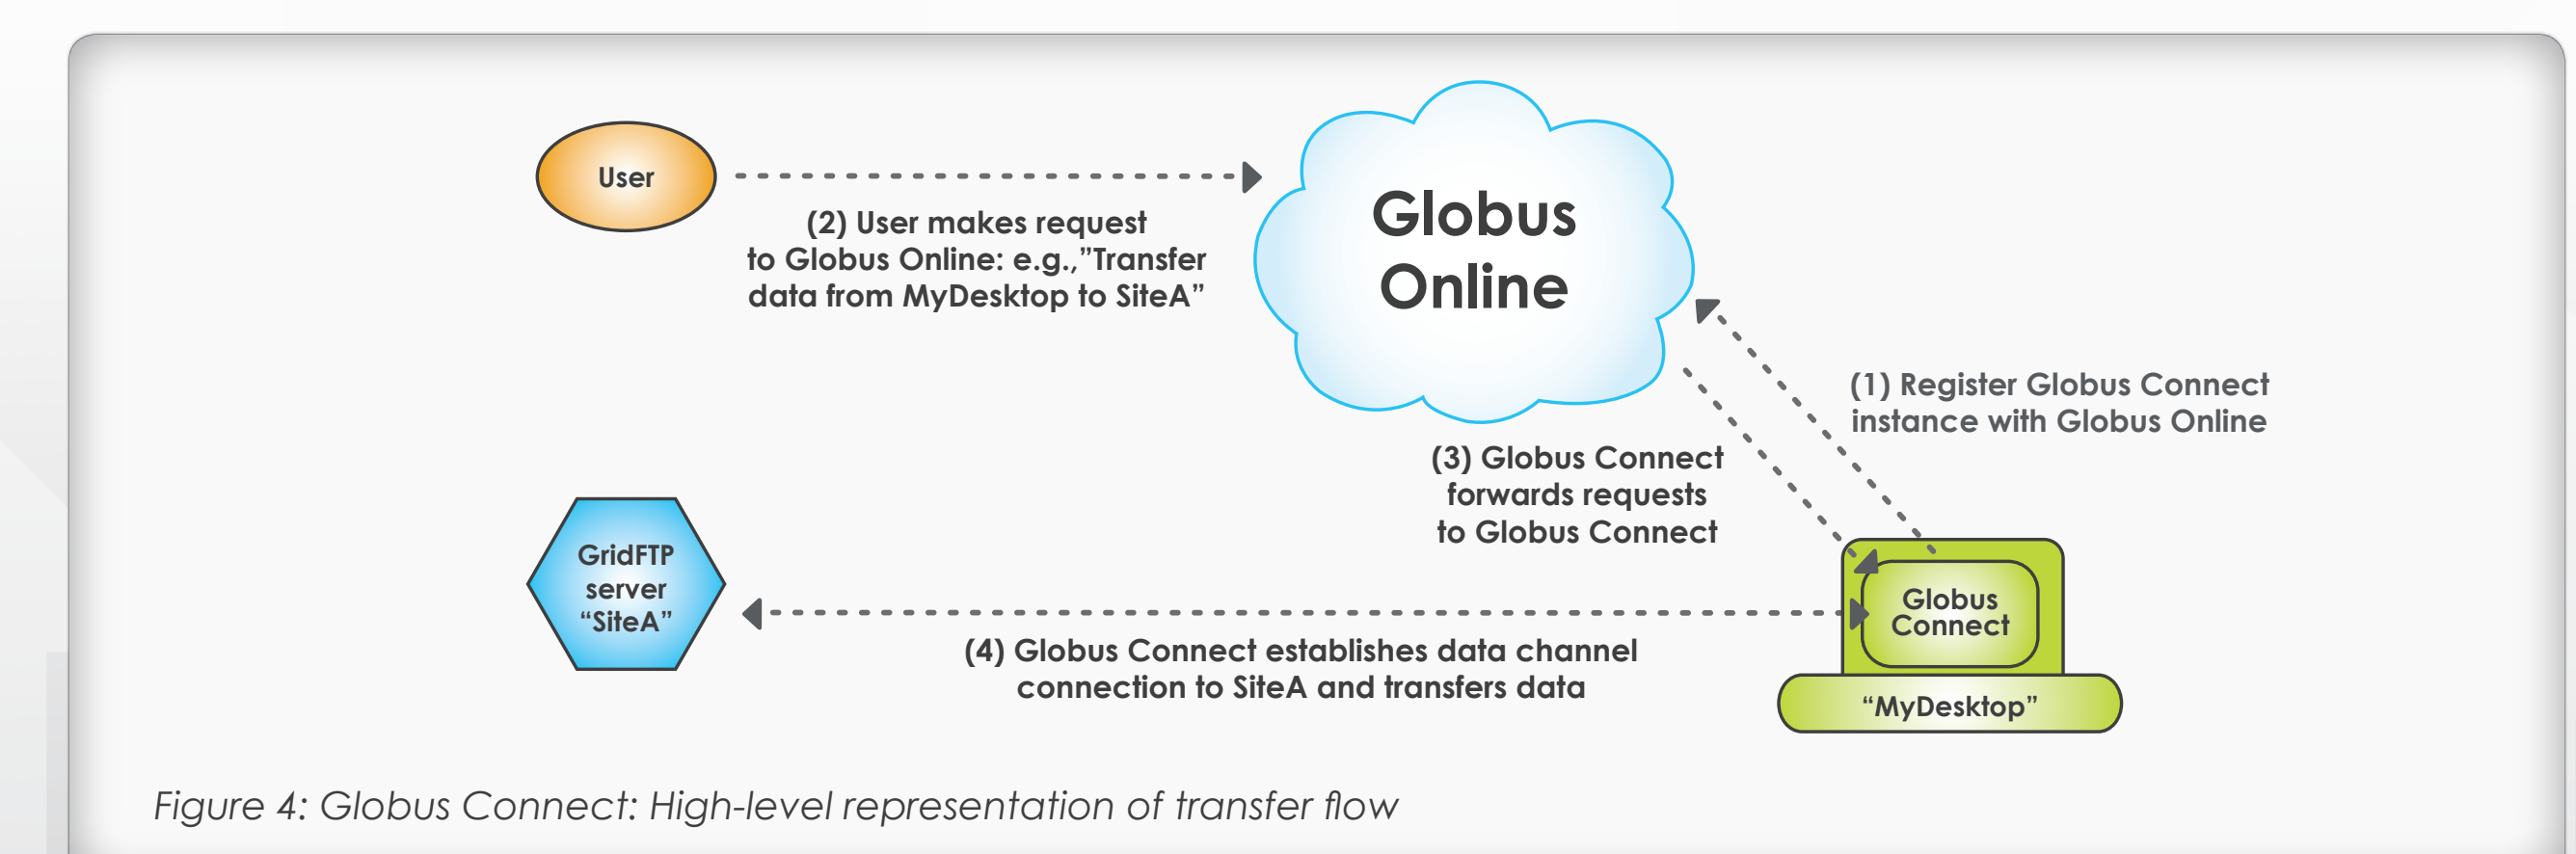


Figure 4: Globus Connect: High-level representation of transfer flow