GridFTP and GSI Migration Plan

Team: Brian Lin[4], Brian Bockelman[1], Tim Theisen[4], Mátyás Selmeci[4], Carl Edquist[4], Edgar M Fajardo[2], Diego Davila[2], Aaron Moate[4], Derek Weitzel[3], Marian Zvada[3]
Institutions: [1]Morgridge Institute for Research, [2]UC San Diego, [3]University of Nebraska–Lincoln, [4]University of Wisconsin–Madison

Overview

The GridFTP data transfer protocol and GSI Authentication and Authorization Infrastructure (AAI) were selected as central components for the OSG ecosystem nearly 15 years ago. In both cases, approaches are becoming **increasingly niche** and the **support costs are increasingly shouldered by the OSG**. Thus, the OSG has the opportunity and motivation to evolve toward a data transfer protocol and security techniques that better fit our needs and allow us to connect to more vibrant software communities. For the data transfer, we are proposing **HTTP**; for the AAI, we are proposing the use of **bearer tokens, HTTPS, and OAuth2**.

August 2019
Beginning of OSG 3.5 (last release series depending on GridFTP/GSI)

October 2019
OSG no longer carries OSG-specific patches for GridFTP/GSI

January 2020
GSI-free site demonstration

July 2020
All GridFTP/GSI free components available in OSG repositories

January 2021
OSG series 3.6 released without GridFTP/GSI dependencies

January 2022
End of support for OSG 3.5 and GridFTP/GSI dependencies

Data Services

There are 70 active GridFTP services registered in the OSG that are used for data transfer. These installations will be replaced with **XRootD configured with HTTP/S support**. The OSG will release XRootD with support for **token-based and GSI authorization** to allow for a smooth transition, **third-party copy**, and **namespace translations** for the LHC experiments.

XRootD can be installed as a standalone entrypoint to a site's local storage, as a cache for a data federation, or as a storage-element to load-balance requests and storage.

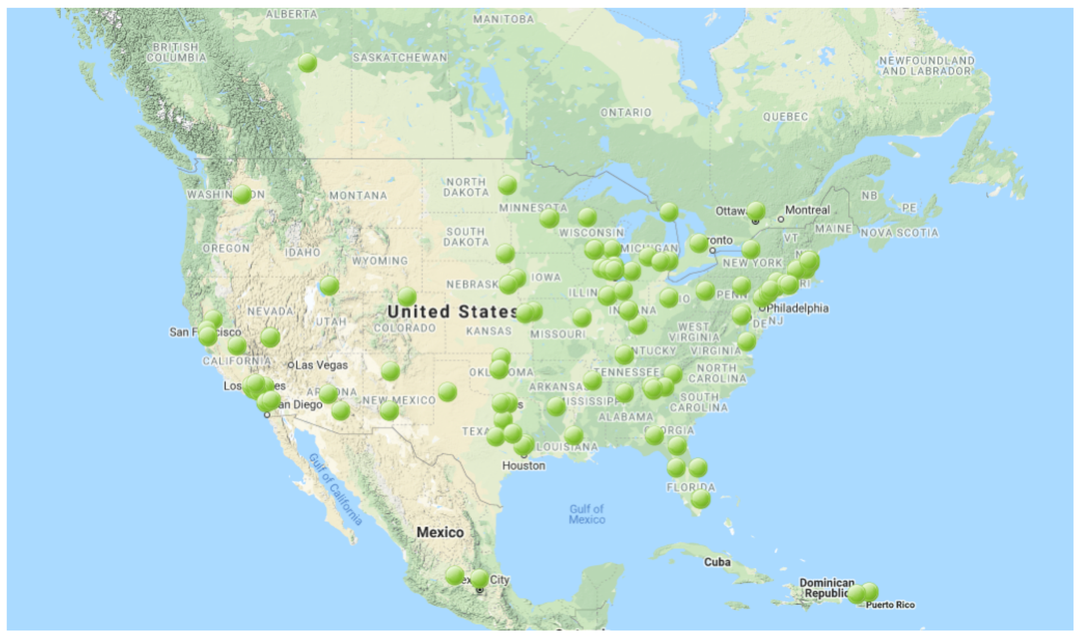

Above: Map of OSG sites affected by the migration
Right: Full migration plan (https://go.wisc.edu/5f3x6z)

Resource Provisioning Services

The software stack used for provisioning computing resources in the OSG relies heavily on GSI authentication for **communication between hosts** (VO pool ↔ pilot factory, pilot jobs ↔ VO pool) and **job submission** (pilot factory → site gateway).

All components in this software stack will be updated to use a combination of **HTCondor tokens** for communication between hosts and **SciTokens/WLCG tokens** for job submission.

Central OSG Services

Many centrally run OSG services use GSI for their AAI and need to be updated to use alternative methods for authentication:

- **Central Collector:** receives reports from active data federation caches and HTCondor-CEs, using GSI to authenticate each reporting host. This usage can be replaced with **SSL certificate authentication**.
- **GSI OpenSSH:** GSI is used to authenticate VO staff to manage their VO's software distribution areas in CVMFS. This can be replaced with **plain OpenSSH**.
- **Topology:** GSI authentication is required to access sensitive contact information, including emails and phone numbers. We intend to transition the contact data to **COManage**, which uses non-GSI methods for authentication.

This project is supported by National Science Foundation under Cooperative Agreement OAC-1836650. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.