PROJECT SUMMARY

Overview:

The management of security credentials for scientific workflows is a burden for scientists and information security officers. Problems with security credentials (e.g., expiration, privilege mismatch) cause scientific workflows to fail to fetch needed input data or store valuable scientific results, distracting scientists from their research by requiring them to diagnose the problems, re-run their computations, and wait longer for their results. In an effort to avoid these problems, scientists often use long-lived, highly-privileged credentials (e.g., enabling the workflow to fully impersonate their identity), increasing risks to their accounts and to the underlying computational infrastructure, causing headaches for information security officers. The SciTokens project delivers an open source software infrastructure to address these challenges, using IETF-standard OAuth tokens for capability-based secure access to remote scientific data. SciTokens uses OAuth refresh tokens, maintained securely on the submission node, to delegate short-lived, least-privilege OAuth access tokens to scientific workflows, to enable their remote data access. The access tokens convey the specific authorizations needed by the workflows, rather than general-purpose authentication impersonation credentials. These least-privilege authorization tokens help to address the risks of scientific workflows running on distributed infrastructure including NSF resources (e.g., LIGO Data Grid, Open Science Grid, XSEDE) and public clouds (e.g., Amazon Web Services, Google Cloud, Microsoft Azure).

Intellectual Merit:

The proposed activity will advance knowledge through use of distributed, least-privilege authorization in our nation's scientific computing infrastructures. Adoption of current Internet standards (i.e., OAuth) enables knowledge transfer of Internet security methods across sectors.

Broader Impacts:

The SciTokens software will enable scientists to perform widely distributed computational science more reliably and securely. The project includes participants from the Laser Interferometer Gravitational-Wave Observatory (LIGO) Scientific Collaboration and the Large Synoptic Survey Telescope (LSST) project to provide target use cases and facilitate adoption of the SciTokens software. Integration with the widely-used HTCondor software and collaboration with Open Science Grid and the Extreme Science and Engineering Discovery Environment (XSEDE) facilitates adoption by the wider scientific community.

Project Description

Access control on remote data is an essential capability for computational science. Maintaining the privacy of scientific data prior to publication helps avoid premature science claims and facilitates a healthy competition between research groups. Maintaining the integrity of scientific data (i.e., preventing intentional or unintentional data alteration) helps avoid erroneous scientific claims and preserves the provenance of scientific results. Computational science is often geographically distributed, with scientists collaborating across organizations and using computational resources spread across distributed facilities such as the LIGO Data Grid, the Open Science Grid, XSEDE, and public cloud providers (Amazon, Google, Microsoft, etc.). Secure remote access to scientific data enables these collaborations.

Access control for distributed scientific computing across administrative domains is fundamentally different from local computing environments. When computing on the desktop or a local high performance cluster, the local operating system provides filesystem access control using local accounts, groups, and file permissions. In contrast, when computing across distributed facilities, access control is managed by the collaboration (often called a *virtual organization* or VO), and access to computing resources and scientific data is determined by membership in that collaboration. A remote compute job, which is often part of a larger computational science workflow, will read input data and write output data to/from remote file servers. These network-accessible file servers may be operated on a long-term basis, serving multiple VOs, or they may be instantiated on-the-fly by the VO (often called a *glide-in*).

The remote compute job needs security credentials to access those file servers. In common practice today, those credentials are *identity tokens*, carrying the identity of the individual researcher or the VO, enabling the job to act on behalf of that researcher or VO when accessing remote resources like file servers. Using identity tokens in this way creates significant risk of abuse, since they are used by jobs that are running on remote, less trusted systems and if stolen, these tokens provide wide access to the attacker. The continuous stream of news reports about compromised passwords at sites like LinkedIn and Yahoo! teach us how using the same credential across multiple sites enables attacks to spread widely. The same concern applies to the credentials we use for academic and scientific computing.

To address this risk for current distributed scientific workflows and to enable use of additional distributed computational resources for more scientific domains, our SciTokens project applies the well-known principle of capability-based security to remote data access. Rather than sending unconstrained identity tokens with compute jobs, we send capability-based access tokens. These access tokens grant the specific data access rights needed by the jobs, limiting exposure to abuse. These tokens comply with the IETF OAuth standard [RFC6749], enabling interoperability with the many public cloud storage and computing services that have adopted this standard. By improving the interoperability and security of scientific workflows, we 1) enable use of distributed computing for scientific domains that require greater data protection and 2) enable use of more widely distributed computing resources by reducing the risk of credential abuse on remote systems.

As illustrated in Figure 1, our SciTokens model applies capability-based security to three common domains in the computational science environment: Submit (where the researcher submits and manages scientific workflows), Execute (where the computational jobs run), and Data (where remote read/write access to scientific data is provided). The Submit domain obtains the needed access tokens for the researcher's jobs and forwards the tokens to the jobs when they run, so the jobs can perform the needed remote data access. The Scheduler and Token Manager work together in the Submit domain to ensure

that running jobs have the tokens they need (e.g., by refreshing tokens when they expire) and handle any errors (e.g., by putting jobs on hold until needed access tokens are acquired). The Data domain contains Token Servers that issue access tokens for access to Data Servers. Thus, there is a strong policy and trust relationship between Token Servers and Data Servers. In the Execute domain, the job Launcher delivers access tokens to the job's environment, enabling it to access remote data.

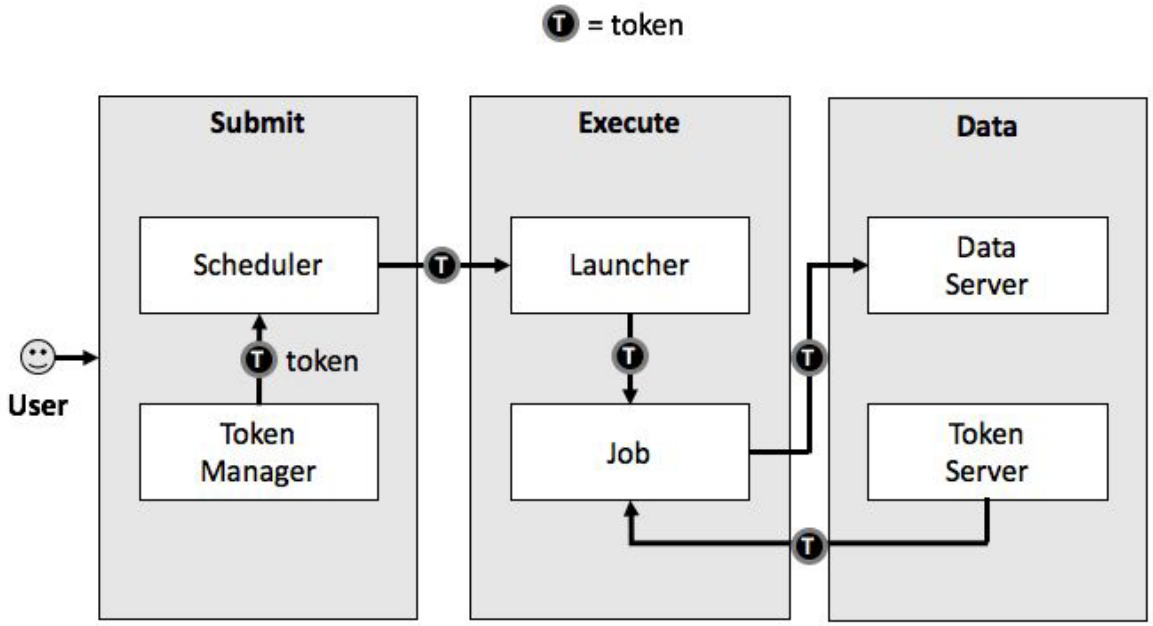

Figure 1: The SciTokens Model

The SciTokens model adopts token types from OAuth (see Figure 2). Users authenticate with identity tokens to submit jobs/workflows, but identity tokens do not travel along with the jobs. Instead, at job submission time the Token Manager obtains OAuth *refresh tokens* with needed data access privileges from Token Servers. The Token Manager securely stores these relatively long-lived refresh tokens locally, then uses them to obtain short-lived *access tokens* from the Token Server when needed (e.g., when jobs start or when access tokens for running jobs near expiration). The Scheduler then sends the short-lived access tokens to the jobs, which the jobs use to access remote data.

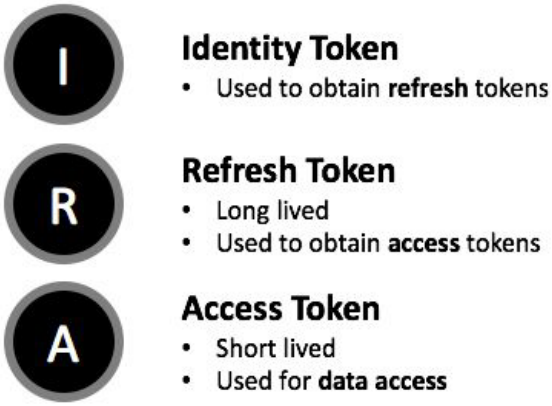

Figure 2: Different Token Types

The remainder of our Project Description is organized as follows: We first discuss the target use cases and partnerships that will ensure broad impact of our work. Then we discuss existing capabilities in the software used by LIGO and other science projects that we will be leveraging and existing. Then we detail our technical approach for implementing the SciTokens model. Lastly, we discuss related work, our open source licenses, and our team qualifications including our results from prior NSF support.

Broader Impacts

To ensure that our work has broad impact on our nation's infrastructure for computational research: 1) we form a team that includes representatives from multiple NSF science projects, 2) we build on the software components currently in use across NSF science projects today, and 3) we establish partnerships with NSF science projects to evaluate and adopt the SciTokens approach for their scientific workflows. In this section we highlight our partnerships with four projects: LIGO, LSST, OSG, and XSEDE. Additional information is provided in attached Letters of Commitment from Randy Trudeau (LIGO Chief Information Security Officer), Steve Kahn (LSST Project Director), Frank Würthwein (Open Science Grid Executive Director), and John Towns (XSEDE PI and Project Director).

The proposed activities also directly impact training for students, scientists, and computer specialists in the use of emerging cyberinfrastructure. For example, LIGO's PyCBC search workflows are often run by undergraduate and graduate students. By simplifying access to computational resources, our proposed work will lower the barriers to entry for students and scientists to make meaningful contributions to LIGO science. This will have disproportionate benefits to scientists with fewer resources from smaller institutions within the LIGO Scientific Collaboration, who typically rely on the national cyberinfrastructure to pursue their science.

LIGO

The observation of binary black hole mergers [LIGO1][LIGO2][LIGO3] by the Advanced Laser Interferometer Gravitational-wave Observatory (LIGO) [LIGO4] marks a transformative moment for physics and astronomy. Gravitational waves are ripples in the fabric of spacetime produced by the coherent relativistic motion of masses. Gravitational-wave observatories can peer into the cores of exploding stars, study the interiors of neutron stars, and explore the physics of colliding black holes. The challenge of gravitational-wave detection is to separate the tiny mirror motions caused by gravitational waves from the motion caused by all the other noise sources in the detector. LIGO's searches are most sensitive when we have prior knowledge of the shape of the gravitational waves. In this case, we can use matched filtering to extract the signals from the noisy data. Since we do not know the physical parameters of any given source in advance, we must search for many different sources by matched filtering the data against a large "bank" of gravitational-waveform templates. The noise in the LIGO detectors consists of a stationary, Gaussian component from fundamental sources and non-Gaussian "glitches" of both environmental and instrumental origin. To eliminate glitches from the data, we record a large amount of instrumental health and status information that must be folded into the search. To make confident detections, we demand that a gravitational-wave signal is present in two or more detectors in the network with consistent signal parameters. To measure the significance of a particular candidate signal, we must compare its amplitude to that of the noise-induced background in the network. We measure this background by repeating the search many times with the detector data offset

by time intervals larger than the gravitational-wave travel time between the detectors. The statistical significance of candidates is then computed and any significant events are followed-up by additional analysis of auxiliary and environmental detector status channels.

Successfully executing all of the steps described above requires the execution of a search workflow with components that span scientific algorithms, cyberinfrastructure, data management, and distributed computational hardware [Brown06]. A typical gravitational-wave search may generate a workflow of hundreds of thousands of discrete computational tasks, with job dependencies and a large numbers of intermediate data products. The structure of a workflow will vary depending on the input parameters for a specific search. Advanced LIGO uses the "PyCBC" search [PyCBC16] to detect and study gravitational-waves from binary black holes, and to search for binary neutron stars and neutron-star–black-hole mergers. PyCBC-generated workflows are written in an abstract workflow format that can be planned by the Pegasus Workflow Management System [Deelman15] and executed by HTCondor [Thain05] on LIGO Data Grid, Open Science Grid, and XSEDE HTC resources (Comet and Stampede).

Authentication tokens are needed at multiple stages during the creation, planning, and execution of the PyCBC search workflow. A program pycbc_make_coinc_search_workflow is used to create a workflow to analyze a specific block of data; this is typically two calendar weeks, an interval chosen to allow measurement of the search's noise background to a level where detections can be made. The workflow creation script must make an authenticated query to two metadata servers: the first reports the on the availability and quality of LIGO data, and the second is used to locate the input data files needed. These files may be stored on a local filesystem, on a XrootD server, or available via GridFTP. At execution time, the jobs use an authentication token to access the LIGO data via CVMFS/XrootD or via GridFTP and to fetch any additional workflow data from the submission site via either the GridFTP or scp protocols. The token is then used to push the data from a job back to the submission site for publication (if it is a final data product) or use by other jobs in the workflow (in the case of an intermediate data product).

Analysis of two weeks of data by a PyCBC workflow typically requires approximately 2500 CPU days and takes several wall-clock days to complete. It is not uncommon for intervals of poor data quality (which trigger additional processing in the workflow) to extend the runtime of a workflow to approximately one wall-clock week. At present all authentication during the workflow is performed using a user's X.509 grid proxy credential, which is created with a lifetime that nominally exceeds the execution time of the workflow. This single credential must exist on the submission site and be copied to all of the job execution sites, which can be of order 10,000 individual nodes for a production OSG/XSEDE run. This credential is the same token that can be used to log in to LIGO clusters and authenticate to many LIGO services. If it is compromised, then an attacker can masquerade as the compromised user for the duration of the credential, gaining access to all of the user's files and to all LIGO data.

LSST

LSST’s Batch Production service for computing jobs and data processing is still currently being implemented and tested. While the overall design of the system has been architected, there are many details that have to be decided upon [Kowalik16]. LSST has yet to determine authentication and authorization access methods at remote sites. One key stated requirement of LSST’s Batch Production service is to provides credentials and endpoint information for any needed LSST services. It is anticipated that access will be in part governed and managed through X.509 certificate infrastructure. It is clear that

with LSST's remote computing and data access needs, OAuth token support would greatly simplify access.

The LSST Batch Production service executes campaigns on computing resources to produce the desired LSST data products. A campaign is composed of a set of pipelines, a set of inputs to run the pipelines against, and a method of handling the outputs of the pipelines. A pipeline is a body of code. A campaign is the set of all pipeline executions needed to achieve an LSST objective and pipelines within a campaign can have a dependency chain since outputs from one pipeline might be required as inputs for an following pipeline.

LSST's pipeline orchestration follows a fairly typical pattern. Initially, the best site and computing resources are scheduled or reserved for the pipeline. Pipeline preparation is then run on the remote site and monitored. The pipeline interfaces with LSST's Data Backbone--which serves and manages LSST data products--for required data which then stages the input data and other supporting files into the staging area. It is expected that the Data Backbone will make use of Pegasus [Deelman15] to manage a pipeline's needed data products. The Data Backbone does not include the handling of authorization or authentication of users or services. This functionality is expected to be provided by layers on top of the Data Backbone.

Using HTCondor DAGman [Couvares07], LSST's pipeline workflow creates a DAG representing pipeline execution. A pipeline would generate multiple condor jobs, comprising the tasks of the pipeline. During execution of the pipeline, a task may require additional data to be staged-in and may require expansion of initial computing resources using "pilot job" functionality. As each task ends, any output might be staged from the compute nodes to the Data Backbone or this might not occur until the pipeline concludes. At completion of the pipeline output data is staged out, various custodial tasks concerning the pipeline are executed, a clean up process is executed and the pipeline releases nodes it may have reserved.

OSG

The Open Science Grid is national, distributed computing partnership for data-intensive research. It allows participating universities and labs to perform distributed, high-throughput computing across heterogeneous resources and at large scale: over 1.3 billion CPU hours were reported in the last year. While the OSG was historically dominated by HEP experiments, it has been rapidly diversifying in the last few years. An important enabler of this diversification is migrating from the legacy X.509-based authentication and authorization layer to more user-friendly mechanisms. Unfortunately, this migration has only been performed for computing jobs: supported storage services still require X.509 authentication. SciTokens will allow OSG to make necessary progress in storage and data access, and represent an important new capability for the broad set of OSG users.

XSEDE

NSF's Extreme Science and Engineering Discovery Environment (XSEDE) provides a single virtual system that scientists can use to interactively share computing resources, data, and expertise to enhance the productivity of scholars, researchers and engineers. Its integrated, comprehensive suite of advanced digital services is designed to federate with other high-end facilities and with campus-based resources, serving as the foundation for a national e-science infrastructure ecosystem.

Secure distributed access to scientific data is an essential function of the XSEDE infrastructure. Today access control in XSEDE is primarily identity-based, via user accounts, certificates, Duo, and InCommon. The SciToken approach, using OAuth for capability-based (rather than identity-based) access to remote scientific data, can provide new options for XSEDE integration with remote data services (on campus, in the cloud, operated by virtual organizations, etc.). XSEDE has experience using OAuth for web single sign-on to components including CILogon, Globus, MyProxy, and the XSEDE User Portal.

Building Blocks

Our SciTokens project builds on the software currently in use for LIGO scientific workflows and in many other NSF science projects, which will ease deployment of our work in these environments by updating existing software rather than forcing use of new software. In this section we describe the existing authentication, authorization, and credential management capabilities of HTCondor, CVMFS, and CILogon OAuth that provide the foundational building blocks for our project.

HTCondor

The general architecture of HTCondor [Thain05] is best explained at a high level by understanding three major components. The first major component is known as the "Submit Machine" and is where users typically interface with HTCondor. They submit their jobs and workflows, monitor progress, and see the results of their jobs from this point. The second component is a pool of "Execute Machines" that can vary in size. These are the compute resources which actually execute the users' jobs, and the user does generally not interface with these machines directly. In some environments, users may have access to multiple pools of execute machines that are controlled by different organizations. The third component is the "Matchmaker", which handles the scheduling of jobs and matches them to compute resources according to their requirements, priority of the user, and several other factors. This process runs continuously so that as new jobs enter the queue, they are matched to available resources. Also, if the pool of resources grows (for example, in a cloud computing environment) then again jobs are matched to the newly available resources. When a job is matched, HTCondor spawns a new process on the Submit Machine called the job "Shadow", as well as a new process on the Execute Machine called the job "Starter". These two processes exist for the lifetime of the job and are responsible for moving data, monitoring the job, reporting the status of the job back to the submit machine while the jobs is running, and moving data again when the job has completed.

The HTCondor components communicate with each other over a network using TCP/IP. Sometimes all components of the cluster belong to a single institution and are on a private network behind a firewall and additional security measures are not needed. In other scenarios, the HTCondor components are widely across the insecure Internet. HTCondor has a wide array of possible security configurations depending on the requirements of the site and system administrator. In the recommended configuration, communication channels from one HTCondor process to another are authenticated using one of HTCondor's supported mechanisms such as a Shared Secret, Kerberos, or X.509 certificates. Thus, all components of an HTCondor system can trust the other components they interface with. HTCondor also supports encryption of data, both for communication between HTCondor components and also for transferring data as part of running users' jobs. Because the authentication process can be relatively computationally intensive when thousands of nodes are involved, it is desirable not to authenticate every single time a new network connection is made. HTCondor accomplishes this by setting up secure and reusable sessions between

components which greatly increases the scalability of the system. For our proposed program of work, we will leverage the secure communication channel between a job's Shadow (on the Submit Machine) and its Starter (on the Execute Machine). It is through this channel that all of a job's executable files, credentials, input, and output data are transmitted. When a large HTCondor system is running, it may be starting dozens of jobs per second, all of which need to establish secure connections between their Shadow and Starter. In this case, the use of sessions does not help. Instead, we rely on a mechanism in which the Matchmaker delegates a trust relationship at the time a job is matched by providing both the Submit and Execute machines with a unique secret key, or “match password”, they can use only with each other. [Miller10].

HTCondor currently has a sub-component on the Submit Machine which manages credentials, called the "CredD". The CredD is responsible for securely managing all credentials. The CredD has a plug-in architecture for storing and managing credentials of different types, which we will leverage to add support for OAuth tokens as described below. The plug-in architecture includes hooks for refreshing credentials and performing other credential transformations. For example, CERN uses a plug-in to manage Kerberos tickets for jobs and uses also transforms the Kerberos ticket into an AFS token for the running job.

CVMFS

CVMFS (originally, “CernVM File System”) is a highly-scalable, read-optimized, global distributed filesystem. End-to-end data integrity is achieved through the use of a public key signing and a Merkle-tree-based integrity scheme. CVMFS repositories can consist of millions of files, yet the global system has far more scalable namespaces than traditional cluster filesystem such as Lustre.

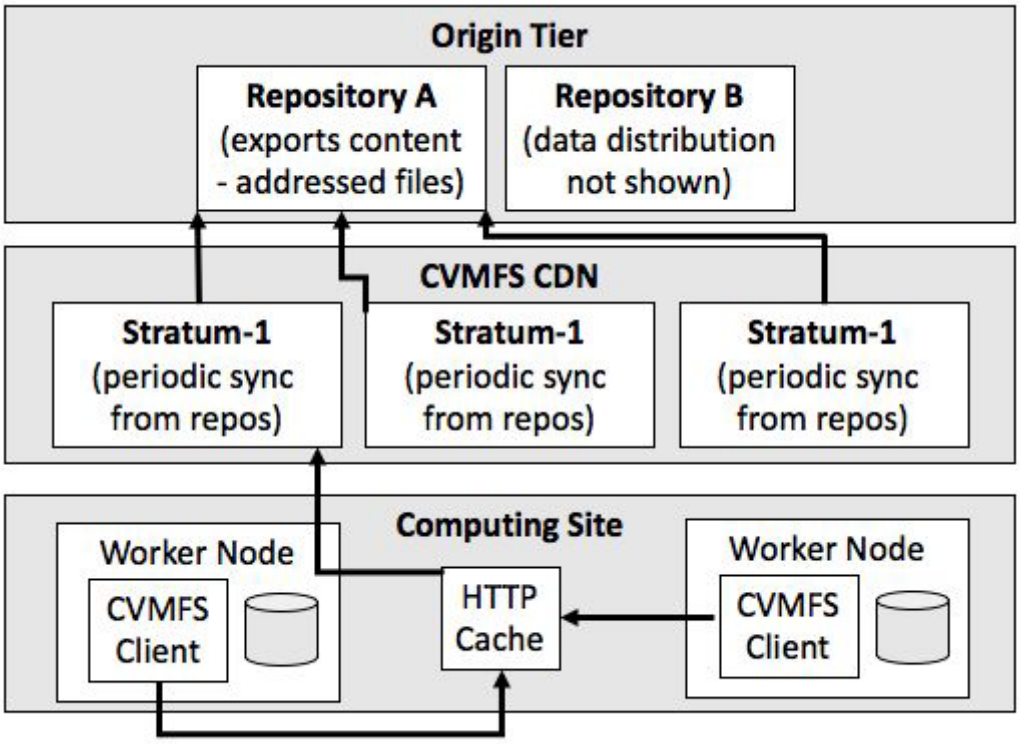

Figure 3: CVMFS Architecture

This scalability is achieved by performing all writes at one location (the “repository”) and introducing transaction-based semantics namespace updates. All data and namespace metadata is then content-addressed. The resulting filesystem is extremely amenable to caching techniques; cache hit rates at each layer in the deployed system are often greater than 99%. Files are distributed using a multi-layer, HTTP-based content distribution network (CDN), site-local HTTP caches, and worker node disk. The traditional CVMFS architecture is outlined in Figure 3.

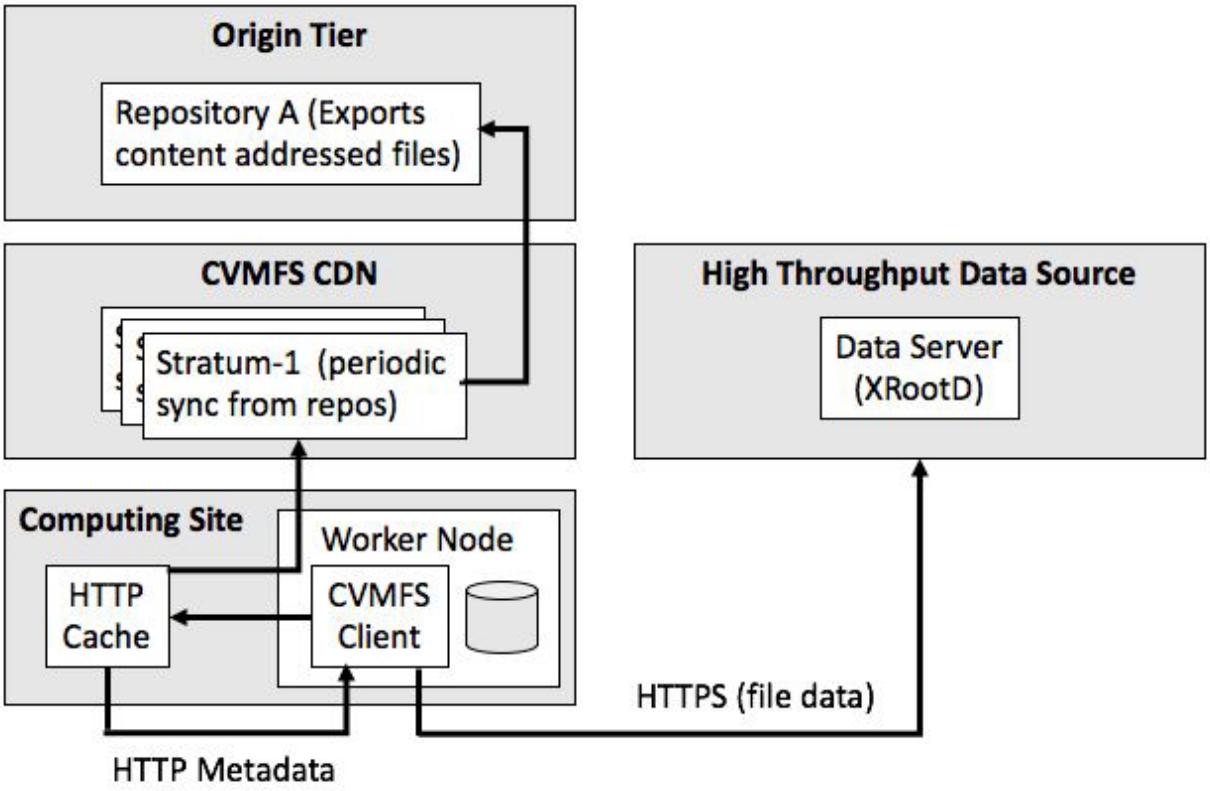

Figure 4: Accessing LIGO files in OSG

The traditional CVMFS infrastructure works extremely well for distribution of software environments or containers where the data is *public* and the working set size is one to ten gigabytes (typically, working set size limitations are limited by the HTTP cache and worker node disk size). In 2016, the proposal team at Nebraska implemented a set of changes to the CVMFS client to efficiently handle user authentication / authorization and larger working set sizes. This was achieved by allowing CVMFS to access a separate data access infrastructure based upon the XRootD software, utilizing both a high-performance storage system at Nebraska and a series of very large caches run by OSG. This allowed the Nebraska team, in collaboration with the OSG, to configure a CVMFS repository containing the LIGO frame files. Figure 4 outlines the architecture used to access LIGO files on the OSG.

When a user process accesses a restricted-access repository, the CVMFS process will request an external process to authenticate and authorize the client (the list of authorizations is distributed as an extended attribute in the CVMFS namespace). If the external process can successfully authorize the user process, it will return success and any user credentials back to CVMFS. This authorizes the user to

access the local disk cache; if the requested file is not found locally, the user credentials are utilized by the CVMFS client to authenticate the file download from the CVMFS CDN. This is illustrated in Figure 5.

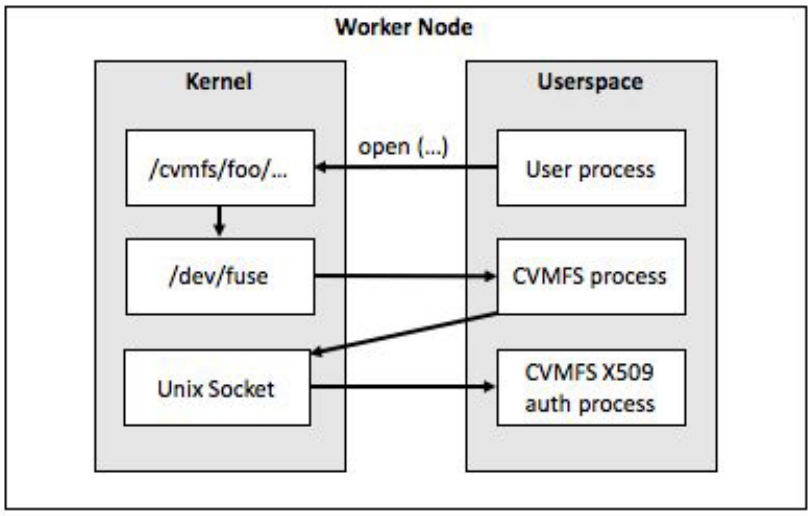

Figure 5: Accessing Restricted Data

Currently, this plugin API utilizes Globus GSI / X.509 certificates for authentication and a list of authorized DNs or VOMS attributes [Alfieri04] for authorization. As described below, the SciTokens project will implement a new plugin for OAuth tokens.

CILogon OAuth

CILogon [Basney14] provides open source software and an operational service for using federated identities in science projects. CILogon's support for SAML [Cantor05] enables interoperability with the US InCommon federation and other federations worldwide via the eduGAIN interfederation service. CILogon's support for X.509 certificates, compliant with standards from the Interoperable Global Trust Federation (IGTF), enables distributed scientific computing in the Grid Security Infrastructure [Welch03]. Over 3,000 scientists regularly use CILogon for authentication, including over 200 LIGO scientists. CILogon includes support for OAuth [RFC6749] and the OpenID Connect [OIDC] standards, using open source software originally developed for NSF science gateways [Basney11]. This OAuth software contains lightweight Java OAuth client/server libraries, with support for JSON Web Tokens [RFC7519], which we will use for our SciTokens implementation.

Technical Approach

In this section we present our technical plans for implementing the capability-based approach for remote data sharing, for LIGO, LSST, and other scientific workflows. Our technical approach modifies HTCondor to manage the tokens for the workflow, modifies CVMFS to accept the tokens for authorizing remote data access, and implements an OAuth Token Server that issues the capability tokens. To validate our approach, we will demonstrate LIGO and LSST workflows executing with SciTokens. In the following subsections, we discuss our technical implementation plans for each system component.

HTCondor

As the component that actually executes a scientific workflow, HTCondor serves as the linchpin that ties together all the SciToken components. To best communicate our planned HTCondor program of work, we first present a walk-thru of how HTCondor will orchestrate the component interactions upon submission of a job, followed by a discussion of integration research opportunities.

As illustrated in Figure 6, the process begins when the researcher submits the computational job using the condor_submit command (or more likely using Pegasus or similar workflow front-end that then runs condor_submit). As part of the submission, the researcher specifies required scientific input data and locations for output data storage in the condor_submit input file. For example, in a LIGO PyCBC [PyCBC16] submission, the researcher will specify a set of data "frames" from the LIGO instrument that are the subject of the analysis. Then condor_submit authenticates the researcher to the token_server(s) to obtain the tokens needed for the job's data access; as an optimization, condor_submit may first check for any locally cached tokens from the researcher's prior job submissions. The token_server determines if the researcher is authorized for the requested data access, based on the researcher's identity and/or group memberships or other researcher attributes. If the authorization check succeeds, the Token Server issues an OAuth refresh token back to condor_submit, which stores the refresh token securely in the condor_credd, and sends the job information to the condor_schedd. Since condor_submit gathers all the needed data access tokens, there is no need to store any identity credentials (e.g., passwords, X.509 certificates, etc.) with the job submission, thereby achieving our goal of a capability-based approach.

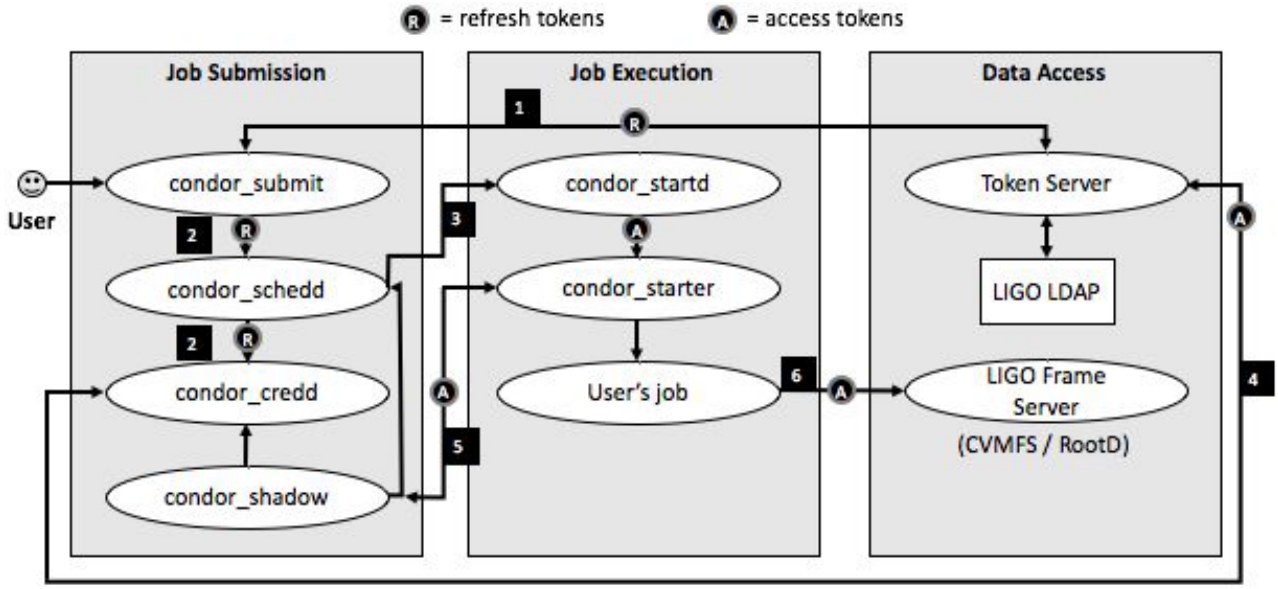

Figure 6: The SciTokens System Architecture

The next phase of the process begins when the condor_schedd has scheduled the job on a remote execution site. The condor_schedd communicates with the condor_startd to launch the job, establishing a secure communication channel between the condor_shadow on the submission side and the condor_starter on the remote execution side. The condor_starter then requests access tokens from the condor_shadow for the job's input data. The condor_shadow forwards the access token requests to the

condor_credd, which uses its stored refresh tokens to obtain fresh access tokens from the token_server. The condor_credd returns the access tokens to the condor_shadow which forwards them to the securely condor_starter which provides them to the researcher's job. Note that only access tokens are sent to the remote execution environment; the longer-lived refresh tokens remain secured in the submission environment which typically resides at the researcher's home institution. Lastly, the job uses the access tokens to mount CVMFS filesystem(s) to access scientific data. CVMFS verifies each access token to confirm that the token was issued by its trusted token_server and that the token's scope includes access to the scientific data being requested. If verification succeeds, CVMFS grants the requested data access. If the access token needs to be refreshed, the condor_starter makes another request back to the condor_shadow.

Note that SciTokens can leverage additional aspects of HTCondor, such as the fact that the condor_shadow can be made explicitly aware if the job is staging input data, accessing data online while the job is running, or staging output data. We will investigate allowing the job submission to state three different sets of access tokens, which will only be instantiated at file stage-in, execution, and file stage-out, respectively. This will enable long running jobs, for instance, to fetch a very short-lived write token for output that will only be instantiated once processing has completed. We will also investigate adjusting the granularity of access token restrictions; for instance, the condor_shadow could request fresh access tokens for each job instance, allowing the token to be restricted in origin to a specific execution node. Alternatively, for greater scalability, access tokens could be cached at the credd and shared across all condor_shadow processes serving jobs that need the same data sets. Finally, we will investigate scenarios in which the data service and its accompanying token service is not fixed infrastructure, but instead is dynamically deployed upon execute nodes, perhaps by the workflow itself. In this scenario, the token service could be instantiated with a set of recognized refresh tokens a priori.

CVMFS / XRootD

The CVMFS (client) and XRootD (server) stack will be updated to understand the SciTokens authorization model. At the minimum, both must be extended to determine whether the bearer of the token is considered within the VO (and able to access the data) or outside. However, as we research the more fine-grained models afforded by Macaroons [Birgisson14], we will be able to prototype more restrictive access control policies -- allowing reads or writes at individual directory level for groups inside the VO. In both cases, this can be done without either CVMFS client or data server needing to know a global identity of the user (as is the case today).

We will implement a new authentication and authorization callout process for CVMFS, based on the experience with our X.509 implementation. This will validate the user's authentication token and, as appropriate, authorize it for use with the CVMFS repository. As the authorization is handled by a separate process, we will be able to prototype the implementation quickly in a scripting language such as Python. The callout API currently only allows callout processes to return X.509 certificate-based client authorization; we will extend it to handle HTTPS bearer authorization (the existing CVMFS authorization plugin API allows for backward-compatible changes such as this).

We will conduct end-to-end performance tests, addressing any deficiencies in the setup. For example, we plan to improve the libcurl-based HTTPS connection cache to handle authenticated connections. This will allow session reuse across file downloads. We will work with the CVMFS upstream to integrate this into a software release, designing and implementing necessary the integration tests and documentation.

While any HTTPS server implementation can likely be made to work with the OAuth model, we will implement changes to the XRootD server suite to have its HTTPS protocol implementation support bearer authentication and the tokens issued by the SciToken service. This server implementation was selected in order to integrate cleanly with the existing service at Nebraska and provide continuity with the existing X.509-secured LIGO repository. We will integrate our token authorization format with XRootD's authorization plugin framework, allowing token-based reading and writing for appropriately enabled XRootD servers. The improved plugin will be deployed and operated at the load-balanced XRootD endpoint at Nebraska; it will be the origin point for LIGO's CVMFS repository. After successfully integrated with the Nebraska service, we will work to deploy it across other OSG-run XRootD endpoints.

OAuth Tokens

The distributed, large-scale architecture of CVMFS presents a special challenge for the OAuth access tokens. Tokens are usually validated via the token server's introspection endpoint [RFC7662]; however, the CVMFS client on each worker node must perform access control. This presents a scalability challenge for the LIGO use case. Ideally, the token's verification should be decentralized as opposed to relying on a token service callback. We will investigate both JSON Web Tokens [RFC7519] and Macaroons [Birgisson14] as a token format. Both formats can be verified knowing only the public key of the signing service; CVMFS contains built-in mechanisms for secure public key distribution. Macaroon-based access tokens have the additional advantage in that they can be delegated or limited by HTCondor without interacting with the token service. If desired, HTCondor could assign unique access permissions for each job in a workflow based on a single access token or make the token only valid for that particular job.

We will evaluate multiple deployment options for the open source token servers, including statically deployed token servers, a token server per VO, with CVMFS supporting multiple VOs, and token servers deployed on-the-fly as part of a data server glide-in.

LIGO and LSST Workflows

To evaluate our end-to-end approach, we will demonstrate use of the SciTokens software for LIGO and LSST workflows. We will work with the LIGO PyCBC workflow [PyCBC16], which uses the Pegasus workflow management system [Deelman15] to manage HTCondor job submissions. (See attached letter of commitment from Ewa Deelman of the Pegasus project.) In addition to demonstrating the use of OAuth tokens for fetching input frame data from CVMFS as part of the LIGO workflows, we will also demonstrate storing PyCBC output data to XrootD and OAuth-capable cloud storage (e.g., Google Cloud Storage), features desired by the PyCBC group.

Related Work

The most significant qualities of the SciTokens model that differentiate it from related work are:

- **highly distributed**: The SciTokens approach does not require a centralized service. Multiple instances of submit, execute, and data domains, operated by different organizations, can interoperate in a highly distributed manner.
- **open source**: All software components in the SciTokens approach are open source and can be deployed by science projects and other operators independent of the project team.

- **standards compliant**: Strict compliance with the OAuth [RFC6749] and JSON Web Token [RFC7519] specifications ensures that our work will be usable with the growing number of services that are adopting these standards.

In the following subsections we compare our approach with ALICE and Globus.

ALICE XrootD Tokens

One of the experiments on the LHC with a comparatively smaller footprint within the US, ALICE never adopted the X.509-credential-based authentication system used throughout the OSG. Instead of using common interfaces (e.g., Globus GridFTP) for data management, ALICE runs private XRootD hosts on top of site storage systems. Having control over an end-to-end VO-specific system allowed ALICE to develop an innovative authorization system. Users would request data access from the central ALICE file catalog; if access was granted, the central service would return an encrypted, base64-encoded XML document describing a list of read / write permissions the bearer is granted. The ALICE-managed XRootD servers would decrypt the token and allow the bearer appropriate authorizations. The data servers would not need the user identity to allow read / write access - effectively, they delegated the management of the ALICE VO's storage allocation to the central ALICE service.

ALICE demonstrated the viability of many of the concepts within SciTokens; we further the approach by:

- Utilizing the standardized OAuth framework instead of a homegrown format (admittedly, ALICE's work predates these standards by more than 5 years).
- Investigating token formats that allow decentralized validation: each CVMFS client will need to perform validation, not just the data service.
- Integrate token generation with federated identity.

Globus Auth

Globus Auth [Tuecke16] provides an OAuth-based cloud service for remote scientific data access using GridFTP [Allcock05], with extensions for delegated access tokens. In contrast to our SciTokens model, Globus Auth has a strong dependency on a centralized, closed-source token service operated by the Globus organization. Tokens in Globus Auth are opaque and require token introspection callbacks to the central service for validation.

Open Source Licenses

SciTokens is an open source project. For existing software components, we use their current open source licenses. CILogon and HTCondor use the Apache 2.0 license. CVMFS uses a BSD 3-clause license. For new software developed in the project, we use the Apache 2.0 license.

About the Project Team

The SciTokens project team consists of leading experts in the relevant project areas. PI Basney has over 15 years experience in federated identity and access management for distributed scientific computing, starting with TeraGrid in 2001. He is also the security technical lead for XSEDE's Community Infrastructure (XCI) area. As OSG's Technology Area Coordinator, co-PI Bockelman has contributed to a

wide range of distributed computing projects, including implementing access control in CVMFS. Co-PI Brown has 17 years experience in LIGO workflow design and is a developer of the PyCBC search for gravitational waves. Co-PI Tannenbaum has over 18 years of experience creating effective distributed high-throughput computing solutions and is the HTCondor Project Technical lead. Co-PI Withers is the Information Security Officer for LSST and has 15 years experience in scientific computing and cybersecurity operations.

Results from Prior NSF Support

Basney is PI of **NSF award #1547268**. **Budget:** $499,973. **Period:** January 2016 through December 2018. **Title**: "CICI: Secure Data Architecture: CILogon 2.0 - An Integrated Identity and Access Management Platform for Science." **Intellectual Merit:** The project advances knowledge and understanding in the field of identity and access management by demonstrating international interfederation, integrating multiple identity protocols into a common platform, and disseminating results to the community. **Broader Impacts:** The project enables science projects to meet their identity and access management needs more effectively so they can allocate more time and effort to their core mission of scientific research. **Publications**: [Basney14] describes the original CILogon system. **Research Products**: CILogon open source software is available on GitHub at https://github.com/cilogon.

Bockelman is senior personnel on **NSF award #1148698**. **Budget**: $18,750,000.00. **Period**: June 2012 through May 2017. **Title**: “THE OPEN SCIENCE GRID The Next Five Years: Distributed High Throughput Computing for the Nation's Scientists, Researchers, Educators, and Students”. **Intellectual Merit**: The OSG is a national, distributed computing partnership for data-intensive research. It provides a fabric of services and a framework for sharing and utilizing a heterogeneous set of computational resources. **Broader Impacts**: The OSG underlies the US portion of the computing strategy of several major NSF investments, including the ATLAS and CMS experiments at the Large Hadron Collider (LHC). Approximately one-third of the OSG usage is US LHC experiments, one-third are other High Energy Physics (HEP) experiments and one-third are other fields of science. **Publications**: For a broad overview, see [OSG]. For a technical paper describing the earliest OSG work on CVMFS, see [OASIS]. **Research Products**: See http://www.opensciencegrid.org for links to further publication, software, and software packaging produced by this proposal.

Brown is PI on **NSF award #1404395**. **Title**: “Gravitational Wave Astrophysics With Advanced LIGO.” **Budget**: $360,000. **Period**: July 2014 through June 2017. **Intellectual Merit**: The direct detection of gravitational waves from binary black hole mergers with Advanced LIGO; measuring the rate of binary black hole mergers in the universe with Advanced LIGO; searching for gravitational waves from binary neutron stars and neutron star–black hole binaries with Advanced LIGO; developing, testing, and deploying searches for compact binary coalescence for use with Advanced LIGO; improving the astrophysical sensitivity of Advanced LIGO data by assessing the impact of non-Gaussian detector noise transients on searches for compact-binary mergers; validation and review of the Advanced LIGO detector’s calibration and assessing the impact of calibration accuracy on the search for compact-binary mergers; exploring the effect of matter in binary neutron star waveforms by comparing numerical simulations to analytic models, and assessing the accuracy of the inspiral phase of neutron star–black hole waveforms using numerical relativity; exploring the ability of Advanced LIGO and next-generation gravitational-wave observatories to test the no-hair conjecture using black hole ringdown signals. **Broader Impacts**: Advanced LIGO’s detection of gravitational waves from colliding black holes provided an exceptional opportunity to increase public scientific literacy and public engagement with science.

Brown helped raise funds to restore Holden Observatory on the Syracuse campus, which houses an eight inch Alvin Clark telescope dating from 1887. Brown and his students regularly give tours of the observatory to the public, hosting star parties and discussions on current events in physics and astronomy. **Publications**: [PyCBC16], [LIGO1]-[LIGO21]. **Research Products**: PyCBC open source software is available on GitHub at https://github.com/ligo-cbc/pycbc.

Tannenbaum is Co-PI on **NSF award #1321762**. **Budget**: $7,269,960.00. **Period**: July 2013 through June 2018. **Title**: “Accomplishment Based Renewal (ABR) to the award Flight-Worthy Condor: Enabling Scientific Discovery”. **Intellectual Merit**: The end-to-end process of identifying challenges and opportunities for high-throughput computing (HTC), developing frameworks to address them, translating novel distributed computing concepts into dependable software tools, maintaining nearly one million lines of HTCondor source code and supporting a broad and diverse user community. **Broader Impacts**: HTCondor’s impact are interdisciplinary and span both academia and industry. Researchers who use HTCondor are able to increase their computing throughput, and consequently increase the size and complexity of the problems they study. HTCondor fuels commercial offerings, such as Red Hat's MRG and Cycle Computing's CycleServer. **Publications**: [Bockelman15], [Fajardo15]. **Research products**: See http://htcondor.org for evidence of ~12 significant HTCondor releases per year and materials from past HTCondor Week workshops; http://htcondor-git.cs.wisc.edu/ for source repository; http://htcondor.org/manual/ for 1000+ page HTCondor Manual; http://wiki.htcondor.org for Wiki including FAQs and development tickets; http://htcondor.org/new.html for HTCondor news feed.

Withers is PI on **NSF award #1547249**, “CICI: Secure Data Architecture: Shared Intelligence Platform for Protecting our National Cyberinfrastructure”, $499,206, 12/1/2015-12/1/2018. **Intellectual Merit**: The SDAIA project actively promotes sharing of intelligence among science DMZ participants as well as with the national scale XSEDE cyberinfrastructure and other open science projects. The SDAIA project lays the foundation for an intelligence sharing infrastructure that will provide a significant benefit to the cybersecurity research community. **Broader Impacts**: The SDAIA project has brought awareness and opened avenues to increased threat intelligence sharing. Several institutions in the open science networks have volunteered to make available their resources for the SDAIA project and thus help to increase threat awareness. **Publications**: No publications were produced yet under this award. **Research products:** SDAIA software is available in GitHub at https://git.ncsa.illinois.edu/awithers/sdmz.

References Cited

[Alfieri04] R. Alfieri, R. Cecchini, V. Ciaschini, L. dell'Agnello, Á. Frohner, A. Gianoli, K. Lõrentey, and F. Spataro, "VOMS, an Authorization System for Virtual Organizations." In: Fernández Rivera F., Bubak M., Gómez Tato A., Doallo R. (eds) Grid Computing. Lecture Notes in Computer Science, vol 2970. Springer, Berlin, Heidelberg. https://dx.doi.org/10.1007/978-3-540-24689-3_5

[Allcock05] William Allcock, John Bresnahan, Rajkumar Kettimuthu, Michael Link, Catalin Dumitrescu, Ioan Raicu, and Ian Foster. 2005. The Globus Striped GridFTP Framework and Server. In Proceedings of the 2005 ACM/IEEE conference on Supercomputing (SC '05). IEEE Computer Society, Washington, DC, USA. https://doi.org/10.1109/SC.2005.72

[Basney11] Jim Basney and Jeff Gaynor, "An OAuth Service for Issuing Certificates to Science Gateways for TeraGrid Users," TeraGrid Conference, July 18-21, 2011, Salt Lake City, UT. https://dx.doi.org/10.1145/2016741.2016776

[Basney14] Jim Basney, Terry Fleury, and Jeff Gaynor, "CILogon: A Federated X.509 Certification Authority for CyberInfrastructure Logon," Concurrency and Computation: Practice and Experience, Volume 26, Issue 13, pages 2225-2239, September 2014. https://dx.doi.org/10.1002/cpe.3265

[Birgisson14] Arnar Birgisson, Joe Gibbs Politz, Ulfar Erlingsson, Ankur Taly, Michael Vrable and Mark Lentczner, "Macaroons: Cookies with Contextual Caveats for Decentralized Authorization in the Cloud," NDSS Symposium, February 2014. https://dx.doi.org/10.14722/ndss.2014.23212

[Bockelman15] B Bockelman, T Cartwright, J Frey, E M Fajardo, B Lin, M Selmeci, T Tannenbaum and M Zvada "Commissioning the HTCondor-CE for the Open Science Grid", *Journal of Physics: Conference Series*, Vol. 664, 2015

[Brown06] D. A. Brown, P. R. Brady, A. Dietz, J Cao, B. Johnson, and J. McNabb. A case study on the use of workflow technologies for scientific analysis: Gravitational wave data analysis. In Ian J. Taylor, Ewa Deelman, Dennis Gannon, and Matthew S. Shields, editors, Workflows for e-Science, chapter 5, pages 41–61. Springer-Verlag, 2006.

[Cantor05] Scott Cantor, John Kemp, Rob Philpott, and Eve Maler (eds.), "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0," OASIS Standard, March 2005. http://docs.oasis-open.org/security/saml/v2.0/

[Couvares07] Peter Couvares, Tevik Kosar, Alain Roy, Jeff Weber and Kent Wenger, "Workflow in Condor", in In Workflows for e-Science, Editors: I.Taylor, E.Deelman, D.Gannon, M.Shields, Springer Press, January 2007 (ISBN: 1-84628-519-4)

[Deelman15] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger, "Pegasus: a Workflow Management System for Science Automation," Future Generation Computer Systems, vol. 46, pp. 17-35, 2015.

[Einstein16] A. Einstein. Approximative Integration of the Field Equations of Gravitation. Preuss. Akad. Weiss. Berlin, page 688, 1916.

[Fajardo15] E M Fajardo, J M Dost, B Holzman, T Tannenbaum, J Letts, A Tiradani, B Bockelman, J Frey and D Mason, "How much higher can HTCondor fly?", *Journal of Physics: Conference Series*, Vol. 664, 2015

[Kowalik16] M. Kowalik, H. Chiang, G. Daues and R. Kooper, “A Survey of Workflow Management Systems”, LSST Technical Note, DMTN-025, 2016.

[LIGO1] The LIGO Scientific Collaboration and the Virgo Collaboration. Observation of Gravitational Waves from a Binary Black Hole Merger. Phys. Rev. Lett., 116(6):061102, 2016.

[LIGO2] The LIGO Scientific Collaboration and the Virgo Collaboration. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. Phys. Rev. Lett., 116(24):241103, 2016.

[LIGO3] The LIGO Scientific Collaboration and the Virgo Collaboration. Binary Black Hole Mergers in the first Advanced LIGO Observing Run. Phys. Rev., X6:041015, 2016.

[LIGO4] The LIGO Scientific Collaboration and the Virgo Collaboration). GW150914: The Advanced LIGO Detectors in the Era of First Discoveries. Phys. Rev. Lett., 116(13):131103, 2016.

[LIGO5] The LIGO Scientific Collaboration and the Virgo Collaboration. The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. 2016. To appear in Astrophysical Journal Letters.

[LIGO6] The LIGO Scientific Collaboration and the Virgo Collaboration. Upper limits on the rates of binary neutron star and neutron-star–black-hole mergers from Advanced LIGO’s first observing run. 2016.

[LIGO7] LIGO Scientific Collaboration and the Virgo Collaboration. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. Class. Quant. Grav., 33(13):134001, 2016.

[LIGO8] LIGO Scientific Collaboration and the Virgo Collaboration. Properties of the Binary Black Hole Merger GW150914. Phys. Rev. Lett., 116(24):241102, 2016.

[LIGO9] LIGO Scientific Collaboration and the Virgo Collaboration. Tests of general relativity with GW150914. Phys. Rev. Lett., 116(22):221101, 2016.

[LIGO10] LIGO Scientific Collaboration and the Virgo Collaboration. Observing gravitational-wave transient GW150914 with minimal assumptions. Phys. Rev., D93(12):122004, 2016.

[LIGO11] LIGO Scientific Collaboration and the Virgo Collaboration. Directly compar- ing GW150914 with numerical solutions of Einstein’s equations for binary black hole coalescence.
Phys. Rev., D94:064035, 2016.

[LIGO12] LIGO Scientific Collaboration and the Virgo Collaboration. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. Phys. Rev., D93(12):122003, 2016.

[LIGO13] LIGO Scientific Collaboration and the Virgo Collaboration. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. 2016.

[LIGO14] Tito Dal Canton, Alexander H. Nitz, Andrew P. Lundgren, Alex B. Nielsen, Duncan A. Brown, Thomas Dent, Ian W. Harry, Badri Krishnan, Andrew J. Miller, Karl Wette, Karsten Wiesner, and Joshua L. Willis. Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors. Phys. Rev., D90(8):082004, 2014.

[LIGO15] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. Phys.Rev., D89:024010, 2014.

[LIGO16] L. K. Nuttall, T. J. Massinger, J. Areeda, J. Betzwieser, S. Dwyer, A. Effler, R. P. Fisher, P. Fritschel, J. S. Kissel, A. P. Lundgren, D. M. Macleod, D. Martynov, J. McIver, A. Mullavey, D. Sigg, J. R. Smith, G. Vajente, A. R. Williamson, and C. C. Wipf. Improving the Data Quality of Advanced LIGO Based on Early Engineering Run Results. Class. Quant. Grav., 32(24):245005, 2015.

[LIGO17] Kevin Barkett, Mark A. Scheel, Roland Haas, Christian D. Ott, Sebastiano Bernuzzi, Duncan A. Brown, Be ́la Szila ́gyi, Jeffrey D. Kaplan, Jonas Lippuner, Curran D. Muhlberger, Francois Foucart, and Matthew D. Duez. Gravitational waveforms for neutron star binaries from binary black hole simulations. Phys. Rev., D93(4):044064, 2016.

[LIGO18] A. H. Nitz, A. Lundgren, D. A. Brown, E. Ochsner, D. Keppel, and I. W. Harry. Accuracy of gravitational waveform models for observing neutron-star–black-hole binaries in Advanced LIGO. Phys.Rev., D88:124039, 2013.

[LIGO19] Prayush Kumar, Kevin Barkett, Swetha Bhagwat, Nousha Afshari, Duncan A. Brown, Geoffrey Lovelace, Mark A. Scheel, and Bela Szilagyi. Accuracy and precision of gravitational-wave mod- els of inspiraling neutron star-black hole binaries with spin: Comparison with matter-free numerical relativity in the low-frequency regime. Phys. Rev., D92(10):102001, 2015.

[LIGO20] Benjamin D. Lackey, Sebastiano Bernuzzi, Chad R. Galley, Jeroen Meidam, and Chris Van Den Broeck. Effective-one-body waveforms for binary neutron stars using surrogate models. 2016.

[LIGO21] Swetha Bhagwat, Duncan A. Brown, and Stefan W. Ballmer. Spectroscopic analysis of stellar mass black-hole mergers in our local universe with ground-based gravitational wave detectors. Phys. Rev., D94(8):084024, 2016.

[Miller10] Zach Miller, Dan Bradley, Todd Tannenbaum, Igor Sfiligoi, "Flexible Session Management in a Distributed Environment", Journal of Physics: Conference Series Volume 219, Issue 4, Year 2010.

[OASIS] B Bockelman, J Caballero Bejar, J De Stefano, J Hover, R Quick, and S Teige3=, "OASIS: a data and software distribution service for Open Science

Grid," Journal of Physics: Conference Series, Volume 513, Track 3, Year 2014. https://doi.org/10.1088/1742-6596/513/3/032013

[OIDC] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros, and C. Mortimore, “OpenID Connect Core 1.0,” November 2014. http://openid.net/connect/

[OSG] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Würthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick, "The Open Science Grid," Journal of Physics: Conference Series, Volume 78, Volume 78, Year 2007. https://doi.org/10.1088/1742-6596/78/1/012057

[PyCBC16] Samantha A. Usman, Alexander H. Nitz, Ian W. Harry, Christopher M. Biwer, Duncan A. Brown, Miriam Cabero, Collin D. Capano, Tito Dal Canton, Thomas Dent, Stephen Fairhurst, Marcel S. Kehl, Drew Keppel, Badri Krishnan, Amber Lenon, Andrew Lundgren, Alex B. Nielsen, Larne P. Pekowsky, Peter R. Pfeiffer, Harald P.and Saulson, Matthew West, and Joshua L. Willis. The PyCBC search for gravitational waves from compact binary coalescence. Class. Quant. Grav., 33(21):215004, 2016.

[RFC6749] D. Hardt (ed.), “The OAuth 2.0 Authorization Framework,” IETF RFC 5849 (Standards Track), October 2012.

[RFC7519] M. Jones, J. Bradley, and N. Sakimura, "JSON Web Token (JTW)," IETF RFC 7797 (Standards Track), May 2015.

[RFC7662] Richer, J., Ed., "OAuth 2.0 Token Introspection", IETF RFC 7662 (Standards Track), October 2015.

[Thain05] Douglas Thain, Todd Tannenbaum, and Miron Livny, "Distributed Computing in Practice: The Condor Experience" Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005.

[Tuecke16] S. Tuecke, R. Ananthakrishnan, K. Chard, M. Lidman, B. McCollam, S. Rosen, I. Foster, “Globus Auth: A Research Identity and Access Management Platform,” 12th IEEE International Conference on eScience, October 25, 2016.

[Welch03] V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kesselman, S. Meder, L. Pearlman, and S. Tuecke, "Security for Grid services," High Performance Distributed Computing, 2003. Proceedings. 12th IEEE International Symposium on, 2003, pp. 48-57. https://dx.doi.org/10.1109/HPDC.2003.1210015

James Basney, Ph.D.
Senior Research Scientist
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

a. Professional Preparation

Oberlin College	Oberlin, OH	Computer Science and English	B.A. 1995
University of Wisconsin-Madison	Madison, WI	Computer Sciences	M.S. 1997
University of Wisconsin-Madison	Madison, WI	Computer Sciences	Ph.D. 2001

b. Appointments

Senior Research Scientist, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, 2001 - present.

c. Products

(i) Five most closely related to proposal project

1. Jim Basney, Terry Fleury, and Jeff Gaynor, "CILogon: A Federated X.509 Certification Authority for CyberInfrastructure Logon," Concurrency and Computation: Practice and Experience, Volume 26, Issue 13, pages 2225-2239, September 2014. http://dx.doi.org/10.1002/cpe.3265
2. Rion Dooley, Joe Stubbs, and Jim Basney, "The MyProxy Gateway," International Workshop on Science Gateways, June 2014, Dublin, Ireland. http://dx.doi.org/10.1109/IWSG.2014.8
3. Jim Basney, Terry Fleury, and Jeff Gaynor, "CILogon: A Federated X.509 Certification Authority for CyberInfrastructure Logon," XSEDE Conference, July 2013, San Diego, CA. http://dx.doi.org/10.1145/2484762.2484791
4. Jim Basney and Jeff Gaynor, "An OAuth Service for Issuing Certificates to Science Gateways for TeraGrid Users," TeraGrid Conference, July 18-21, 2011, Salt Lake City, UT. http://dx.doi.org/10.1145/2016741.2016776
5. Jim Basney, Terry Fleury, and Von Welch, "Federated Login to TeraGrid," 9th Symposium on Identity and Trust on the Internet (IDtrust 2010), Gaithersburg, MD, April 2010. http://dx.doi.org/10.1145/1750389.1750391

(ii) Five other significant products

1. Jim Basney, Rion Dooley, Jeff Gaynor, Thejaka Amila Kanewala, Suresh Marru, Marlon Pierce, and Joe Stubbs, "Integrating Science Gateways with XSEDE Security: A Survey of Credential Management Approaches," XSEDE Conference, July 2014, Atlanta, GA. http://dx.doi.org/10.1145/2616498.2616559
2. Thejaka Amila Kanewala, Suresh Marru, Jim Basney, and Marlon Pierce, "A Credential Store for Multi-Tenant Science Gateways," International Symposium on Cluster, Cloud and Grid Computing (CCGrid), May 2014, Chicago, IL. http://hdl.handle.net/2022/17379
3. Jim Basney, Rion Dooley, Jeff Gaynor, Suresh Marru, and Marlon Pierce, "Distributed Web Security for Science Gateways," Gateway Computing Environments Workshop (GCE11), November 17, 2011, Seattle, WA. http://dx.doi.org/10.1145/2110486.2110489
4. Jim Basney, Von Welch, and Nancy Wilkins-Diehr, "TeraGrid Science Gateway AAAA Model: Implementation and Lessons Learned," TeraGrid Conference, August 2-5, 2010, Pittsburgh, PA. http://dx.doi.org/10.1145/1838574.1838576

5. J. Basney, M. Humphrey, and V. Welch, "The MyProxy Online Credential Repository", Software: Practice and Experience, Volume 35, Issue 9, July 2005. http://dx.doi.org/10.1002/spe.688

d. Synergistic Activities

The following activities highlight service, scholarship, and teaching contributions:

1. MyProxy project lead, providing credential management software distributed through the Globus Toolkit, the Fedora Project, GitHub, and SourceForge.
2. CILogon project lead, providing an operational service that issues digital credentials to the research community.
3. Interoperable Global Trust Federation member, representing CILogon, InCommon, NCSA, and OSG.
4. Program Committee member: PEARC17, XSEDE16, CCGrid 2015, XSEDE15, XSEDE14, SC14, XSEDE13, SC13, SC12, MGC 2011, HPDC 2011, SC10, Grid 2010, MGC 2010, MGC 2009, SC10, Grid 2010, CCGrid 2009, IDtrust 2008, 2008 IEEE TCSC Doctoral Symposium, SC08, Grid 2008, MGC 2008, CCGrid 2008, AINA 2008, CCGrid 2008, IDtrust 2008, 2008 IEEE TCSC Doctoral Symposium, SC08, Grid 2008, MGC 2008, AINA 2007, WWW 2007, 2007 IEEE TCSC Doctoral Symposium, Grid 2007, SC07, SECOVAL 2007, MGC 2007, CCGrid 2007, AINA 2007, WWW 2007, Grid 2006, MGC 2006, GCE 2006, WASR 2006, HPDC-15, MGC 2005, MGC 2004, DIVO 2004, ISWC 2004 Trust Workshop

Dr. Brian Bockelman
Computer Science & Engineering
118C Schorr Center
(402) 472-5029/Fax (402) 472-3892
bbockelm@cse.unl.edu

Professional Preparation

Institution	*Location*	*Major*	*Degree*
University of West Georgia	Carrollton, GA	Math	B.S.
University of Nebraska-Lincoln	Lincoln, NE	Math, Computer Science	M.S., Ph.D.

Appointments

Year(s)	*Title*
6/2013-present	Asst. Research Professor, University of Nebraska-Lincoln
2013-present	USCMS Computing Project Execution Team member
10/2011-present.	OSG Technology Area Coordinator
10/2011-10/2016	Any Data, Anytime, Anywhere project technical lead
2008-6/2013	Postdoc Researcher, University of Nebraska-Lincoln
1/2008-2013	USCMS Grid Services team member
1/2008-10/2011.	OSG Metrics and Measurements Area coordinator
5/2006-12/2008	Graduate Research Assistant for Nebraska CMS Tier-2 Compute center

Publications

Publications most closely related to the proposed project

1. Zhang, Z., Bockelman, B., Tannenbaum, T. and Carder, D. 2015. Lark: Bringing Network Awareness to High Throughput Computing. *Proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2015)*, Shenzhen, Guangdong, China, May 2015.
2. Weitzel, D.; Bockelman, B.; Swanson, D. Distributed Caching Using the HTCondor CacheD. *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*: 341-346. (2015).
3. Bauerdick, L., Benjamin, D., Bloom, K., Bockelman, B., Bradley, D., Dasu, S., Ernst, M., Gardner, R., Hanushevsky, A., Ito, H., Lesny, D., McGuigan, P., McKee, S., Rind, O., Severini, H., Sfiligoi, I., Tadel, M., Vukotic, I., Williams, S., Würthwein, F., Yagil, A., and Yang, W., "Using Xrootd to Federate Regional Storage", 2012 *J. Phys.: Conf. Ser*. **396** 042009.
4. B. Bockelman, "Using Hadoop as a Grid Storage Element", 2009 J. Phys.: Conf. Ser. 180 01204.
5. Altunay, M., Avery, P., Blackburn, K., Bockelman, B., Ernst, M., Fraser, D., Quick, R., Gardner, R., Goasguen, S., Levshina, T., Livny, M., McGee, J., Olson, D., Pordes, R., Potekhin, M., Rana, A., Roy, A., Sehgal, C., Sfiligoi, I., and Würthwein, F., "A Science Driven Production Cyberinfrastructure—the Open Science Grid", 2011 *J. of Grid Computing,* doi:10.1007/s10723-010-9176-6.

Other significant publications and presentations

1. Bockelman, B. "LIGO on OSG". Presentation at the OSG All-Hands Meeting. https://indico.fnal.gov/getFile.py/access?contribId=9&sessionId=3&resId=0&materialId=slides&confId=10571 (2016)
2. B. Bockelman, "Big Data – Flexible Data – For HEP", Computing in HEP (keynote speaker), October 2013.
3. B. Bockelman, "Putting Condor in a container: Adapting virtualization techniques to batch systems", BiG Grid e-Infrastructure colloquium, January 2012.
4. B. Bockelman, "The Worldwide LHC Computing Grid: Data Processing on a Global Scale", 2009 Hadoop Summit
5. B. Bockelman, "Improving Data Access for the LHC using Xrootd", Conference on Federated Data Stores, November 2011.

Synergistic Activities

- Lead the usage of the Hadoop Distributed File System as a storage system within USCMS.
- Started the CMS initiative for optimizing wide-area, remote I/O in USCMS using Xrootd.
- Leads the OSG Technology Investigations team, starting 2011, for managing the present OSG software stack and its evolution.
- Built, scale, monitor, and operate a nation-wide data access infrastructure for CMS based on Xrootd consisting of USCMS sites.
- Implemented X509-based authorization and scalability features for CVMFS.

BIOGRAPHICAL SKETCH: DUNCAN A BROWN

Professional Preparation:

M.Math.	Mathematics. (University of Newcastle Upon Tyne. 1999)
Ph.D.	Physics. (University of Wisconsin–Milwaukee. 2004)

Appointments:

2016-	Charles Brightman Professor of Physics, Syracuse University
2011-2016	Associate Professor of Physics, Syracuse University
2007-2011	Assistant Professor of Physics, Syracuse University
2004-2007	Postdoctoral Scholar in Physics, California Institute of Technology

Five most closely related products:

1. B. P. Abbott, *et al.*, "*Binary black hole mergers in the first advanced LIGO observing run,*", Phys. Rev. X **6**, 041015 (2016).

2. B. P. Abbott, *et al.*, "*GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence,*", Phys. Rev. Lett. **116**, 241103 (2016).

3. B. P. Abbott, *et al.*, "*Observation of Gravitational Waves from a Binary Black Hole Merger,*", Phys. Rev. Lett. **116**, 061102 (2016).

4. B. P. Abbott, *et al.*, "*GW150914: First results from the search for binary black hole coalescence with Advanced LIGO,*", Phys. Rev. **D93**, 122003 (2016).

5. S. A. Usman, *et al.*, "*The PyCBC search for gravitational waves from compact binary coalescence,*", Class. Quant. Grav. **33**, 215004 (2016).

Five other products:

1. Swetha Bhagwat, Duncan A. Brown, and Stefan W. Ballmer, "*Spectroscopic analysis of stellar mass black-hole mergers in our local universe with ground-based gravitational wave detectors*", Phys. Rev. **D94**, 084024 (2016).

2. Kevin Barkett *et al.*, "*Gravitational waveforms for neutron star binaries from binary black hole simulations*", Phys. Rev. **D93**, 044064 (2016).

3. B. P. Abbott, *et al.*, "*Calibration of the Advanced LIGO detectors for the discovery of the binary black hole merger GW150914*", arXix:1602.03845 (2016).

4. B. P. Abbott, *et al.*, "*Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914*", Class. Quant. Grav. **33**, 134001 (2016).

5. Prayush Kumar, *et al.*, "*Accuracy and precision of gravitational-wave models of inspiraling neutron star-black hole binaries with spin: Comparison with matter-free numerical relativity in the low-frequency regime*", Phys. Rev. **D92**, 102001 (2015).

Synergistic Activities:

- Co-chair of the 2016 Gordon Research Conference on Physics Research and Education (Relativity and Gravitation: Contemporary Teaching of Einstein's Physics).
- Fellow of the American Physical Society. Member at large of the executive committee of the APS Division of Gravitational Physics (2015–2017). Member at large of the executive committee of the APS Division of Computational Physics (2015–2017). Chair of the APS Division of Gravity Nominations committee (2016). Member of the APS Division of Computational Physics Rahman Prize committee (2016).
- Awarded the university-wide 2010 Meredith Outstanding Teaching Award for teaching excellence in the AST101 introductory astronomy class and for involving undergraduate students in gravitational-wave research. Awarded 2013 and 2014 Syracuse University Physics Teaching Excellence Prizes.
- Awarded a Research Corporation for Science Advancement Cottrell Scholar award for faculty members who are committed to excel at both teaching and research. 2015, 2016 Research Corporation Scialog Fellow.
- Program Coordinator for 2016 Kavli Institute for Theoretical Physics Rapid Response Program: Astrophysics from LIGOs First Black Holes.

Todd Tannenbaum

Department of Computer Sciences, Center for High Throughput Computing
University of Wisconsin-Madison
1210 West Dayton Street
Phone: +1 (608) 263-7132
Email: tannenba@cs.wisc.edu

Professional Preparation

Institution	*Location*	*Major*	*Degree & Year*
University of Wisconsin-Madison	Madison, WI, USA	Computer Sciences	B.S., 1990
University of Wisconsin-Madison	Madison, WI, USA	Computer Sciences	M.S., 2013

Appointments

2004-present	Technical Lead, HTCondor Project, Center for High Throughput Computing, University of Wisconsin-Madison
2004-present	Researcher, Department of Computer Sciences, University of Wisconsin-Madison
2006-2007	Interim Deputy Facility Coordinator, Open Science Grid
1997-2004	Associate Researcher, Department of Computer Sciences, University of Wisconsin-Madison
1996-1999	President, Coffee Computing Corporation
1995-1997	Director, Model Advanced Facility (MAF), UW Computer Aided Engineering Center, University of Wisconsin-Madison
1994-1998	Technology/Contributing Editor for Network Computing Magazine, CMP Media
1990-1994	UNIX Systems Manager, UW Computer Aided Engineering Center, University of Wisconsin-Madison

Products

Five Products Most Closely Related to the Proposed Project:

[1] Zhe Zhang, Brian Bockelman, Dale Carder, and Todd Tannenbaum, "Lark: Bringing Network Awareness to High Throughput Computing", *Proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2015)*, Shenzhen, Guangdong, China, May 2015.

[2] James Frey, Todd Tannenbaum, Ian Foster, Miron Livny, and Steven Tuecke, "Condor-G: A Computation Management Agent for Multi-Institutional Grids", *Proceedings of the Tenth IEEE Symposium on High Performance Distributed Computing (HPDC10)* San Francisco, California, August 7-9, 2001.
Note: In 2012, this paper was recognized as the third best paper submitted inthe 20 year history of the HPDC conference; see http://hpdc.org/best.php

[3] Douglas Thain, Todd Tannenbaum, and Miron Livny, "Distributed Computing in Practice: The Condor Experience" *Concurrency and Computation: Practice and Experience*, Vol. 17, No. 2-4, pages 323-356, February-April, 2005.
[4] Zach Miller, Dan Bradley, Todd Tannenbaum, Igor Sfiligoi, "Flexible Session Management in a Distributed Environment", *Journal of Physics: Conference Series Volume 219, Issue 4, Year 2010.*
[5] Alexandru Iosup, Dick H.J. Epema, Todd Tannenbaum, Matthew Farrellee, Miron Livny, "Inter-Operating Grids through Delegated MatchMaking", in proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC07), Reno, Nevada, November 2007. http://www.cs.wisc.edu/condor/doc/SC07-Iosup.pdf

Five Other Significant Products:

[1] Zach Miller, Todd Tannenbaum, and Ben Liblit, "Enforcing Murphy's Law for Advance Identification of Run-time Failures", *Proceedings of the 2012 USENIX Annual Technical Conference (USENIX ATC '12)*, pages 203-208, Boston, Massachusetts, USA, June 13-15, 2012
[2] Todd Tannenbaum, Derek Wright, Karen Miller, and Miron Livny, "Condor - A Distributed Job Scheduler", in Thomas Sterling, editor, *Beowulf Cluster Computing with Linux*, The MIT Press, 2002. ISBN: 0-262-69274-0. http://www.cs.wisc.edu/condor/doc/beowulf-chapter-rev1.pdf
[3] Todd Tannenbaum and Michael Litzkow, "Checkpointing and Migration of UNIX Processes in the Condor Distributed Processing System", *Dr Dobbs Journal*, Feb 1995. http://www.cs.wisc.edu/condor/doc/dobbs_95.ps
[4] Andrew Baranovski, Gabriele Garzoglio, Igor Terekhov, Alain Roy and Todd Tannenbaum, "Management of Grid Jobs and Data within SAMGrid", *Proceedings of the 2004 IEEE International Conference on Cluster Computing*, pages 353-360, San Diego, CA, September 2004.
[5] How to Measure a Large Open Source Distributed System, Douglas Thain, Todd Tannenbaum, and Miron Livny, Concurrency and Computation: Practice and Experience, volume 18, issue 15, page 1989-2019, 2006. http://www.cse.nd.edu/~dthain/papers/measure-ccpe.pdf

Synergistic Activities

- Besides research publications, authored over 25 articles that were published in several of the nation's mainstream trade software development and information technology periodicals, such as Dr. Dobbs Journal, Network Computing Magazine, and Information Week.
- Contributed as a proposal reviewer of SBIR proposals for DOE, and as a paper reviewer on the SC09 Papers Committee (Grids Area)
- Have contributed to the development of the HTCondor distributed high throughput computing software system for 24 years.
- On the steering committee for the UW-Madison Center for High Throughput Computing

Alexander Withers

Cyber Security Division, National Center for Supercomputing Applications, University of Illinois Urbana-Champaign, Urbana, IL 61802 email: alexw1@illinois.edu

A. Professional Preparation

Gonzaga University, Spokane, WA	Computer Science	B.S. 2002
Stony Brook University, Stony Brook, NY	Computer Science	M.S. 2004

B. Appointments

2014– Senior Security Engineer, *National Center for Supercomputing Applications*
Leading and supporting grant funded activities related to security log analysis. Working with and directing the activities of graduate students, as well as collaborating with fellow security researchers and engineers. Developing log analysis solutions and supporting intrusion detection efforts in emerging areas like SCADA.

2013-2014 Technology Architect, *Brookhaven National Laboratory*
Member of the Cyber Security Incident Response Team and Cyber Security Operations (unclassified) group. Tasked with the protection of computer and network based assets at the laboratory by architecting enterprise security solutions. Helped to ensure regulatory compliance and risk mitigation.

2008-2013 Senior Technology Engineer, *Brookhaven National Laboratory*
Member of the Cyber Security Incident Response Team and Cyber Security Operations (unclassified) group. Responded to incidents and performed forensic investigations. Managed and mined log data to detect intrusion, computer use violations, and other anomalous behavior. Developed software to assess security controls, prevent reconnaissance, and correlate data. Deployed and maintained intrusion detection systems running Snort and Bro.

2004–2008 Technology Engineer, *Brookhaven National Laboratory*
Member of RHIC/USATLAS Computing Facility Linux Farm group responsible for administering a 4000+ CPU cluster and 20+ infrastructure support servers. Maintaining and developing complex resources sharing policies using the HTCondor workflow management system. Worked to integrate Grid middle-ware with HTCondor. Developed and maintained database, virtualization and monitoring software for high performance computing applications and infrastructure.

2002–2004 Teaching and Research Assistant, Department of Computer Science, *Stony Brook University*
Worked under Prof. Barbara Jacak to develop and explore pervasive computing for the PHENIX experiment at RHIC using common grid middle-ware technology. Investigated authentication, distributed data storage, and job management. Started project to develop framework for scheduling jobs over the grid which developed into my masters project.

2001 ERULF Fellowship, *Idaho National Laboratory*
Summer intern working with cyber security group. Implemented backdoor detection algorithms as a snort preprocessor.

C. Products

[1] Phuong Cao, Eric C. Badger, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer, Alexander Withers, and Adam J. Slagell. "Towards an unified security testbed and security analytics framework." *In Proceedings of the 2015 Symposium and Bootcamp on the Science of Security* (HotSoS 2015). ACM, New York, NY, USA, , Article 24 , 2 pages. DOI=10.1145/2746194.2746218 http://doi.acm.org/10.1145/2746194.2746218

[2] Greeling, K., Bashir, M., Withers, A., "Factors for Differentiating Human from Automated Attacks," *Digital Forensic Research Workshop*, Nov., 2017, *Submitted*

D. Synergistic Activities

- Member of the NSF Large Facility CISO Working Group, 2016-Present.
- Information Security Officer for the Large Synoptic Survey Telescope.
- Reviewer for the SBIR and STTR Programs at the U.S. Department of Energy for 2011-2013.
- On program committee for the Second Workshop on the Changing Landscape in HPC Security, 2015

Facilities, Equipment and Other Resources

Unfunded Collaborations

The following unfunded collaborators have provided letters that document their commitments to the project: John Towns (XSEDE), Randy Trudeau (LIGO), Steven Kahn (LSST), Gerardo Genis and Jakob Blomer (CVMFS), Frank Würthwein (OSG), and Ewa Deelman (Pegasus).

Facilities - NCSA

Facilities, Equipment, and Other Resources

NCSA continues to support user communities by offering the resources that are the foundations of advanced cyberinfrastructure. NCSA's resources are located in two buildings on the University of Illinois' Champaign-Urbana campus. The NCSA building provides office space for most NCSA staff, advanced visualization laboratories, and 2,500 square feet of climate-controlled raised access floor machine room space for small-scale production and experimental system development. NCSA's primary production computing facility is housed in the National Petascale Compute Facility.

The National Petascale Computing Facility

The state-of-the-art 88,000-square-foot National Petascale Computing Facility (NPCF) is located at the corner of Oak Street and St. Mary's Road on the University of Illinois' south campus and houses Blue Waters and other NCSA major infrastructure, as well providing office space for NCSA and vendor staff members. NPCF began operations June 1, 2010.

The facility includes 30,000 square feet of 6' raised floor, with 20,000 square feet of contiguous, un-obstructed space for computing equipment, free of air handling equipment and support structures. NPCF combines top-flight physical and cyber protection with the open, collaborative research attitudes of a public institution. It is a highly efficient, environmentally-friendly facility and one of the very few large scale data centers to achieve LEED Gold certification. Energy efficiency is an important driver in the operation of the facility and all equipment will be as energy efficient as possible. Efficient space use is to be achieved with tall high-density equipment if possible. Full power metering and environmental monitoring are provided.

NPCF:

- Takes advantage of three on-site water economization cooling towers to provide essentially energy free cooling water about 50% of the year, depending on weather conditions.
- Takes advantage of the campus' highly reliable electrical supply providing 24 MWs of usable power with three independent feeds and a fourth feed for redundancy.
- Utilizes the campus chilled water distribution system and thermal storage facilities.
- Has a 0.5 MW UPS system for selected equipment such as metadata servers, network routers and cyber protection systems.

- Provides a 480 volt AC power distribution infrastructure to compute systems thus reducing power conversion losses. Storage and communication equipment is strongly encouraged to use 480V, but may use 120/208V if the 480V alternative does not exist
- Has multiple cooling loops. All high heat load equipment is required to be water-cooled. The Blue Waters system, for example, has all its compute, storage and server racks using state of the practice liquid cooling.
- Will operate continually at the high end of the ASHRAE standards. Equipment must be able to operate with a 60-65°F inlet water temperature and a 78°F inlet air temperature.
- Uses state of the art monitoring and control systems.

NCSA Computational and Data Resources

Blue Waters

The Blue Waters system provides unprecedented, highly productive resources and services for computational and data intensive science. It is designed for maximum throughput on very large-scale, complex applications using Cray's Gemini interconnect architecture and is a landmark petascale system with a peak speed of 13.1 PF, 1.66 PB of memory, 26 PB of user accessible on-line storage, 300 PB of usable near-line storage and 400 Gbps in external networking capability. More important than the peak speed is that Blue Waters achieves real, sustained petascale performance on multiple real science problems averaging 1.3 PF/s. Blue Waters combines into one, single, fully open system, all connected with a single, uniform, best of class interconnect fabric, the one of the world's largest general purpose computational resources, one of the largest accelerator based resources, and the most intense storage system to enable future discoveries that are simply beyond reach today. The system characteristics are summarized in the below.

- Cray XE6/XK7 system with:
    - 22,752 XE6 nodes each with 2 AMD Interlagos processors (16 FP cores) and 64 GB of RAM.
    - 4,224 XK7 nodes each with 1 AMD Interlagos processor (8 FP cores), 1 NVIDIA Kepler K20x GPU, and 32 GB of RAM.
- Cray Gemini 3D torus Interconnect with dimensions of 24x24x24 providing a peak per node injection bandwidth of 9.6 GB/s and minimum bisection bandwidth of 10.4 TB/s.
- Online storage provided by 212 Cray Sonexion Lustre appliances combining 17,280 disks for 26 PBs of usable storage out of the 35 PBs of raw space. Usable online storage bandwidth to the Cray compute engine in excess of 1.1 TB/s.
- Nearline storage provided in 4 (to be expanded to 6 in 2014) dual-arm libraries, each with 16,000 slots and an aggregate throughput of 100 GB/s. The environment currently provides usable storage up to 199 PB with an upgrade to 300 PB later in 2015 with further expansion to more than 500 raw PB limited only by media budget.
- 78 Dell servers provide high-speed file transfer via Globus Online to both online and near line storage at up to 400 Gbps to the NCSA WAN infrastructure.
- The Illinois developed Integrated System Console that collects, analyzes and takes action on millions of events per day from the system to assist in management and monitoring as well as fault detection and prediction.

iForge

iForge is NCSA's premiere and dedicated HPC resource for Private Sector Partners. iForge features two distinct hardware platforms, each configured for a different set of computational needs.

- 144 Dell Power Edge servers utilizing
    - dual Intel Haswell processors with 128 GB of RAM
    - or dual AMD Abu Dhabi processors with 256 GB of RAM.
- Interconnect: QDR InfiniBand
- 3456 cores in total
- 700 TB GPFS filesystem

ROGER – Resourcing Open Geospatial Education and Research

ROGER is NCSA's newest computational resource dedicated to research in Geographic Information Science. ROGER combines multiple computational capabilities in a modest sized system with a large, fast file system to enable data intensive work. The computational capabilities include a traditional HPC system including a portion of the nodes with GPU accelerators, a set of nodes dedicated to HADOOP and an OpenStack capability.

- 36 Dell Power Edge servers utilizing
    - dual Intel Haswell processors with 128 GB of RAM and a 500GB HD
    - 12 nodes include an NVIDIA K40 GPU
- 18 Dell Power Edge servers utilizing
    - dual Intel Haswell processors with 256 GB of RAM
    - 800 GB SSD
- Interconnect: 10 and 40 Gb Ethernet
- 1000 cores in total
- 5 PB GPFS file system

Storage Condominium

The NCSA Storage Condominium provides a reliable mid-scale data storage resource to projects that is not tied to any specific compute resource. The storage condo supports projects needing as little a TB of storage all the up to a PB and more. Multiple access methods are supported including GridFTP, NFS and native GPFS clients. The condo also offers a virtual machine capability to support various utility services for projects. The current storage condo provides about 2 PBs of total usable storage.

Ice House

The Ice House is a cold data storage service providing reliable long-term data storage for infrequently used data from a TB to many PBs. The service utilizes a Spectra Logic T950 library with space for 8,000 LTO tapes and a Spectra Logic Black Pearl provide the data transfer node and interface to users. Data is stored using the LTFS format to allow tapes to be exported and read by others. The system currently has 4.5 PBs of available storage.

High-Performance Network

All computing platforms are interconnected to a multi-10 gigabit network core. The high performance Ethernet backbone provides 1 or 10 gigabit connections as required with up to 300 gigabit external network capacity, eventually moving to 40 and/or 100 Gbps links. Currently 120 Gbps of external

aggregate bandwidth is configured via 10 Gb links to multiple national networks including Internet2, XSEDE, and more via links to major research network hubs in Chicago such as MREN and OmniPOP. Upgrades to multiple 100 Gbps connections are in progress

CyberProtection

All systems at NPCF must comply with cyberprotection and operational standards to enable highly effective and efficient operations. NCSA's cyberprotection includes monitoring all exit bandwidth, deep packet inspection and tracking of network flows with state-of-the-the art tools including an 80-node Bro cluster.

Applications Software

NCSA offers a variety of third-party applications and community codes that are installed on the high-performance systems at NCSA. These applications cover a wide range of science and engineering domains, data analytics and visualization, mathematics and statistics. Some additional software available via University of Illinois campus licensing programs (e.g. Oracle, ABAQUS, ANSYS, etc.) is also available for use by University of Illinois researchers.

Facilities - UW-Madison

This project will reside in the University of Wisconsin-Madison (UW-Madison) Center for High Throughput Computing (CHTC) within the Department of Computer Sciences. The department has sufficient centralized office space, professional information technology support (including software and hardware support and maintenance, data backups, managed web and file servers), and administrative support available to all personnel. The CHTC is a campus-wide organization dedicated to supercharging research on campus by working side-by-side with domain scientists on infusing high throughput computing techniques into their routine. The CHTC consists of approximately 20 full-time staff with significant scientific distributed computing operational and development experience. UW-Madison is the lead institution for the NSF/DOE Open Science Grid (OSG), and runs a Tier-2 computing center for the international CMS LHC experiment. Beyond just being familiar with deployment and use of such systems, UW-Madison staff has also been intimately involved in the design, implementation and deployment of the software. Existing UW technology infrastructure that can be leveraged includes CPU capacity, network connectivity, middleware connectivity, and storage availability including the UW-Madison Digital Collections Center's MINDS@UW digital archive to indefinitely maintain final-form research materials.

High Throughput Computing Compute Clusters. For high-throughput computing (HTC) capability, UW-Madison maintains many compute clusters across campus, which are managed via the HTCondor software developed by the and maintained by the CHTC; therefore, these clusters are linked together to share resources via widely adopted distributed computing technologies. Together these clusters represent roughly 30,000 CPU cores in support of

research. Between 1/1/2015 and 12/31/2015, the CHTC provided over 300 million CPU hours of computing work. Temporary file space for large individual files can support up to hundreds of terabytes of total working data. For single computing runs needing significant memory on a single server, beyond the typical value of 128 GB, the CHTC maintains two multi-core servers, one with 1 TB of memory and one with 2 TB of memory. The CHTC can also engage computing resources from the Open Science Grid (OSG). Individual users of our high-throughput computing (HTC) system, including Open Science Grid capacity, can frequently obtain in excess of 200,000 CPU hours per day.

Build and Test Cluster. The UW NMI Build and Test Lab (BaTLab) facility enables managed, automated builds and tests which lead to the creation and hardening of production-quality software. BaTLab infrastructure consists of an HTCondor pool comprised of 2 submit machines, a central manager, a database server and 20 execute machines covering 17 distinct 32- and 64-bit platforms. Several machines host platforms as virtual machines using the KVM virtualization software. Each front-end machine for job submission is an 8 (2x4) core Intel Xeon 2.40GHz server with 4TB disk (~2TB usable as RAID 1) and 24GB RAM. Each execute machine/infrastructure server is a 12 (2x6) core AMD Opteron server with 4x300GB disk (~1TB usable as RAID 0) and 32GB RAM. Additional resources include a file server (10TB SAN) and 2 Cisco UCS controllers with 7 nodes. Each node is a 12 (2x6) core Intel server with 2x146GB 15Krpm SAS drives and 64GB RAM. Each platform receives scheduled operating system updates along with installs of common system packages (OpenSSL, Perl, XML, etc) and is subject to regular quality control builds to ensure that all software is working properly. The lab's network is a 1GB switched LAN using Cisco 3750x switches with a 10GB uplink. Each machine has 2 x 1GB interfaces and is IPv4 and IPv6 ready. The SAN has a 1GB interface. The entire Lab is on a VLAN connected to the UW-Madison network described below.

Network. The UW-Madison network is currently comprised of a 10GB backbone with 10GB connections to heavy-use buildings and departments, and 1GB connections to the rest. Redundancy is built into every network node. An equitable funding model assures that network resources are kept current; upgrades to 20GB connections and beyond are already being planned. WiscWaves, the high-speed optical network connection to Chicago, provides researchers with 10GB dedicated research networks (lamdas). For example, the UW Department of Physics uses a dedicated lambda to the High Energy Physics Large Hadron Collider project. UW has been fundamental to the establishment of the Broadband Optical Research Education And Science network (BOREAS). This Regional Optical Network (RON) connects to the CIC OmniPoP in Chicago, providing a high-speed gateway to various research networks, including Internet2, National Lambda Rail (NLR), ESNet and other global research networks.

Facilities - UNL

This document details the equipment resident in the Holland Computing Center (HCC) as of

August 2016.

HCC has two primary locations directly interconnected by a pair of 10 Gbps fiber optic links (20 Gbps total). The 1800 sq. ft. HCC machine room at the Peter Kiewit Institute (PKI) in Omaha can provide up to 500 kVA in UPS and genset protected power, and 160 ton cooling. A 2200 sq. ft. second machine room in the Schorr Center at the University of Nebraska-Lincoln (UNL) can currently provide up to 100 ton cooling with up to 400 kVA of power. Brocade MLXe routers, one in each location, provide both high bandwidth and Software Defined Networking (SDN) capability. The Schorr machine room connects to campus and Internet2/ESnet at 100 Gbps while the PKI machine room connects at 10 Gbps.

HCC's resources at UNL include two distinct offerings: Sandhills and Red. Sandhills is a linux cluster dedicated to general campus usage with 5,408 compute cores interconnected by low-latency Infiniband networking. 175 TB of Lustre storage is complemented by 50 TB of NFS storage and 3 TB of local scratch per node.

The largest machine on the Lincoln campus is Red, with 6,960 job slots interconnected by a mixture of 1 Gb and 10 Gb ethernet. More importantly, Red serves up over 4.6 PB of storage using HDFS (Hadoop Distributed File System), an open source version of the file system Google uses. Red is integrated with the Open Science Grid (OSG), and serves as a major site for storage and analysis in the international high energy physics project known as CMS (Compact Muon Solenoid).

The largest HCC clusters, named Tusker and Crane, are located at PKI (Peter Kiewitt Institute) in Omaha. Tusker offers 6,784 cores interconnected with Mellanox QDR Infiniband along with 523TB of Lustre storage. Each compute node is an R815 server with at least 256 GB RAM and 4 Opteron 6272 (2.1 GHz) processors.

Crane debuted at 474 on the Top500 list with an HPL benchmark or 121.8 TeraFLOPS. Intel Xeon chips (8-core, 2.6 GHz) provide the processing with 4 GB RAM available per core and a total of 7,232 cores. The cluster shares 1.5 PetaBytes of lustre storage.

Attic and Silo form a near line archive with 1 PB of usable storage. Attic is located at PKI in Omaha, while Silo acts as an online backup located in Lincoln. Both Attic and Silo are connected with 10Gbps network connections.

Anvil is an OpenStack cloud environment consisting of 1,520 cores and 400TB of CEPH storage all connected by 10 Gb networking. The Anvil cloud exists to address needs of NU researchers that cannot be served by traditional scheduler-based HPC environments such as GUI applications, Windows based software, test environments, and persistent services.

These resources are detailed further below.

1. **HCC at UNL Resources**:

1.1 Sandhills

- 62 4-socket Opteron 6376 (2.3 Ghz/192GB RAM/64 bit) (64 cores per node)
- 42 4-socket Opteron 6128 (2.0 Ghz/128GB RAM/64 bit) (32 cores per node)
- 2 4-socket Opteron 6168 (1.9 Ghz/256GB RAM/64 bit) (48 cores per node)
- QLogic Infiniband (QDR)
- 1 and 10 GbE networking
    - 5x Dell N3048 switches

- 50 TB shared storage (NFS) -> /home
- 175TB shared scratch storage (Lustre) -> /work
- 3TB local scratch

1.2 Red
- (USCMS Tier-2 resource, available opportunistically via the Open Science Grid)
- 60 2-socket Xeon E5530 (2.4GHz/64 bit) (16 cores per node)
- 20 2-socket Xeon E5520 (2.27 GHz/64 bit) (16 cores per node)
- 36 2-socket Xeon X5650 (2.67GHz/64 bit) (24 cores per node)
- 16 2-socket Xeon E5-2640 v3 (2.6GHz/64 bit) (32 cores per node)
- 40 2-socket Xeon E5-2650 v3 (2.3GHz/64 bit) (40 cores per node)
- 30 2-socket Opteron 2354 (2.2 GHz/64 bit) (8 cores per node)
- 28 2-socket Xeon E5-2650 v2 (2.6GHz/64 bit) (32 cores per node)
- 48 2-socket Xeon E5-2660 (2.2GHz/64 bit) (32 cores per node)
- 2 2-socket Xeon E5-1660 v3 (3.0GHz/64 bit) (16 cores per node)
- 4,600 TB HDFS storage (2,300 TB usable)
- Mix of 1GbE and 10GbE networking
    - 1x Dell S6000-ON switch
    - 1x Dell S4048-ON switch
    - 5x Dell S3048-ON switches
    - 2x Dell S4810 switches

1.3 Silo (backup mirror for Attic)
- 1 Mercury RM216 2U Rackmount Server 2 Xeon E5-2630 (2.6GHz/64 bit) (12 cores)
- 10 Mercury RM445J 4U Rackmount JBOD with 45x 4TB NL SAS Hard Disks

2. **HCC at PKI Resources**:

2.1 Tusker
- 106 PowerEdge R815 systems
    - 102x with 256 GB RAM, 2x with 512GB RAM, 2x with 1024GB RAM
    - 4-socket Opteron 6272 Interlagos (64-core, 2.1GHz / 64Bit)
- Mellanox QDR Infiniband
- 1 GbE networking
    - 3x Dell Powerconnect 6248 switches
- 523TB Lustre storage over Infiniband

2.2 Crane
- 452 Relion 2840e systems from Penguin
    - 452x with 64 GB RAM
    - 2-socket Intel Xeon E5-2670 (8-core, 2.6GHz / 64Bit)
- Intel QDR Infiniband
- 1 and 10 GbE networking
    - 1x QuantaMesh 10 GbE switch
    - 11x QuantaMesh 1 GbE switches
- 1500 TB Lustre storage over Infiniband
- 3 Supermicro SYS-6016GT systems

    - 3x with 48GB RAM
    - 2-socket Intel Xeon E5620 (4-core 2.4GHz / 64Bit)
    - 2 Nvidia M2070 GPUs
- 3 Supermicro SYS-1027GR-TSF systems
    - 3x with 128GB RAM
    - 2-socket Intel Xeon E5-2630 (6-core 2.3GHz / 64Bit)
    - 3 Nvidia K20M GPUs
- 1 Supermicro SYS-5017GR-TF systems
    - 1x with 32GB RAM
    - 1-socket Intel Xeon E5-2650 v2 (8-core 2.6GHz / 64Bit)
    - 2 Nvidia K40C GPUs
- 5 Supermicro SYS-2027GR-TRF systems
    - 5x with 64GB RAM
    - 2-socket Intel Xeon E5-2650 v2 (8-core 2.6GHz / 64Bit)
    - 4 Nvidia K40M GPUs

2.3 Attic

- 1 Mercury RM216 2U Rackmount Server 2 Xeon E5-2630 (2.6GHz/64 bit) (12 cores)
- 10 Mercury RM445J 4U Rackmount JBOD with 45x 4TB NL SAS Hard Disks

2.4 Anvil

- 76 PowerEdge R630 systems
    - 76x with 256GB RAM
    - 2-socket Intel Xeon E5-2650 v3 (10-core, 2.3GHz / 64Bit)
    - Dual 10Gb Ethernet
- 12 PowerEdge R730xd systems
    - 12x with 128GB RAM
    - 2-socket Intel Xeon E5-2630L v3 (8-core 1.8GHz / 64Bit)
    - 12x 4TB NL SAS Hard Disks and 2x200GB SSD
    - Dual 10Gb Ethernet
- 2 PowerEdge R320 systems
    - 2x with 48GB RAM
    - 1-socket Intel E5-2403 v3 (4-core 1.8GHz / 64Bit)
    - Quad 10Gb Ethernet
- 10 GbE networking
    - 6x Dell S4048-ON switches

Facilities - Syracuse

The Physics Department at Syracuse University will provide office space for students and faculty supported by this proposal. In addition, Syracuse University provides rooms dedicated to gradu- ate students who work on sponsored projects. The Syracuse University Gravitational Wave Group (SUGWG) has a dedicated meeting room, equipped with a projector, dedicated computer, and videoconferencing facilities for effective face-to-face or teleconferenced meetings. These facilities are available for the participants in the proposed project.

Syracuse University Information and Technology Services (ITS) manages the computing infrastructure for the SUGWG and the ITS computing facilities will be available to the participants in this proposal. These facilities include:

1. Access to the SU ITS Academic Virtual Hosting Environment (AVHE) which can provision virtual servers using a VMWare cloud. AVHE servers are used for compute nodes HTCondor job submission, analysis, and post-processing. AVHE services are also used to provision web services (including wikis and Git repositories) and for integration with the Open Science Grid (OSG).
2. Access to the CRUSH Virtual Compute cloud and the provision of dedicated compute nodes for LIGO data analysis. As described in the Memorandum of Understanding Between the Syracuse University and The Laser Interferometer Gravitational Wave Observatory (LIGO) Laboratory[1], Syracuse University will dedicate computing resources at the level of 4,000 equivalent Intel X5650 cores for LSC data analysis activities on a continuous basis for the duration of this proposal. Priority access to these computing resources will be given to members of the SUGWG group and collaborators.
3. Access to the SU ITS Graphics Processing Unit (GPU) compute farms. SU ITS operates a GPU cluster containing 224 NVIDIA GTX 750Ti cards. This large-scale compute resource uses cards based on the the GM107 Kepler microarchitecture. Each 750 Ti has 2 Gb of GDDR5 RAM and a single-precision floating point performance of 1306 GFLOPS. Together, this resource provides over 292 TFLOPS of computing performance that can be used for LIGO data analysis. Additionally four GTX 1080 cards are available for developing code for the current Pascal microarchitecture.
4. Access to OrangeGrid; a high-throughput computing resource available free of charge to the Syracuse campus community. OrangeGrid uses Syracuse's desktop computers, HTCondor, and a custom virtual machine manager to allow idle computers to perform research- computing tasks. Support and maintenance of OrangeGrid is provided by Syracuse Univer- sity Information and Technology Services and deployment has now reached over 20,000 CPU cores. OrangeGrid is available for SUGWG faculty, postdocs, and students to use for their research computing needs in support of this proposal.
5. SU ITS has a full-time cyberinfrastructure engineer to help with the use and deployment of scientific codes on both Syracuse and OSG resources.

Servers and computing services are located in the Syracuse University Green Data Center, a 12,000 square-foot facility with 6,000 square feet of infrastructure space and 6,000 square feet of raised-floor data center space.

[1] https://dcc.ligo.org/LIGO-M1500083

Data Management Plan

Introduction

This document outlines the data management plan for the proposal “CICI: CE: SciTokens: Capability-based Secure Access to Remote Scientific Data” (“SciTokens”), PI Basney.

The primary outputs of the SciTokens project are expected to be software implementing the SciTokens authentication scheme (as part of the larger CILogon software suite), improvements and add-ons to CVMFS, plugins for the XRootD software suite, measurements of the corresponding system’s performance, and technical reports and papers describing the outcomes of the project. The raw data may be network measurements, log files, transfer history, and/or production job histories; this raw data will be collected from existing monitoring sources at UNL, NCSA, and Syracuse. We will be making the raw data public when the corresponding technical reports are written.

Roles and responsibilities

Dr. Basney (PI) has overall responsibility for management of project data. Dr. Basney will ensure that system administrators at the University of Illinois properly implement this data management plan. Dr. Basney delegates responsibility to each subaward PI for management of project data stored at subawardee institutions. The co-PIs will be responsible to ensure the primary outputs collected at their institutions are transferred to NCSA by the end of the SciTokens project. Should Dr. Basney leave the University of Illinois, his direct supervisor (Adam Slagell, Chief Information Security Officer, National Center for Supercomputing Applications at the University of Illinois) will assume responsibility for project data.

Types of data

The project manages the following types of data:

- **Software** will be publicly accessible under Open Source licenses from commercial software hosting providers (SourceForge, GitHub).
- The final **configurations** (except sensitive data such as passwords) of the production hosts will be documented and stored alongside the corresponding software documentation, where practical. The configurations are essential outcome of the integration work, as they demonstrate the environments where the work was done.
- **Technical reports** or papers and their raw data will be made available through their respective journals and corresponding institutional libraries. Preference will be given to open access journals.

Policies for access/sharing and appropriate protection/privacy

The software for SciTokens will be open source code and made freely available. Raw data (such as log files and performance traces) will be made available alongside the corresponding journal papers and upon request.

Policies for re-use, re-distribution, and production of derivatives

All software components will be licensed under an Open Source license. Where practical, we will use the Apache Software License 2.0[1]; where possible, the technical documents will be licensed under the Creative Commons 3.0 license (Attribution / Non-Commercial / Share-Alike)[2]. Contributions to larger code bases that are *incompatible* with these licenses will default to the preferred license of that code base. Any papers resulting from this work will be copyrighted and licensed according to the journal where they are published.

Data storage and preservation of access

All reports, presentations, manuscripts, and other documents that record research outputs generated under this project will be deposited in IDEALS (https://www.ideals.illinois.edu/), the Illinois Digital Environment for Access to Learning and Scholarship. All datasets and accompanying documentation generated under this project will be deposited in the Illinois Data Bank (https://databank.illinois.edu/), the file-based repository for research data at the University of Illinois at Urbana-Champaign. Both repositories are optimized for their respective content types and support robust indexing and stable access.

Other Scientific Results

Scientific results and other data in repositories protected by SciTokens will *not* be considered outputs of this project and are *not* covered by this data management plan.

For example, LIGO frame files stored in the LIGO CVMFS repository are not covered by this data management plan and are covered by the existing cooperative agreement between the LIGO Scientific Collaboration and NSF.

[1] http://www.apache.org/licenses/LICENSE-2.0.html
[2] http://creativecommons.org/licenses/by-nc-sa/3.0/us/

Postdoc Mentoring Plan

The Syracuse University postdoc's professional development will be enhanced by a structured program of mentoring activities. The goal of this mentoring program is to ensure that the postdoc gains the scientific, technical, and professional skills necessary to excel in their career. To accomplish this goal, the mentoring plan will follow the guidance of the National Academies of Science and Engineering on how to enhance the postdoctoral experience, by providing a structured mentoring plan, career planning assistance, and opportunities to learn a number of career skills such as writing grant proposals, teaching students, writing articles for publication and communication skills [1]. Specific elements of the mentoring plan will include:

- Working with the postdoctoral researcher to develop an **individual development plan** which will define expectations and goals for their professional development. Through a combination of **self-assessment and formal evaluations** (at least annually) the postdoc will ensure that their goals are appropriate and are being achieved.
- Attending seminars and workshops organized by the Syracuse University Office of Sponsored Programs (OSP) on **funding opportunities, writing competitive grant applications, the grant submission process, crafting budgets and budget narratives and best practices for expending sponsored project funds.**
- Participation in seminars and workshops on **teaching and learning** and advice on preparation of the teaching statement required for a faculty application.
- Attending the active **journal club** run by the SU physics postdocs and graduate students. This club meets every Monday afternoon. The participants discuss and critique recent journal articles in the field and to discuss how to write and submit journal articles.
- **Present the results of their research to the broader collaboration:** The postdoc will present their results at LIGO Scientific Collaboration meetings and collaboration teleconferences, as appropriate (travel funds are included in the budget).
- Participation in the weekly research group meetings of the Syracuse University Gravitational-Wave Group. All group members will be expected to present their research regularly, and feedback and coaching will be given to help all members to **develop their communication and presentation skills.**
- Participation in a monthly brown bag lunch series for postdoctoral fellows and graduate students in the Physics Department, in which speakers will be invited to discuss subjects related to **career development** such as how to apply for a faculty position, career paths outside of academia, tips for negotiating salary and start-up funds, how to plan and independent research agenda, etc.

Success of this mentoring plan will be assessed by tracking the progress of the postdoc through their individual development plan, with interviews to assess their satisfaction with the mentoring program, and tracking of their progress toward their career goals after finishing the postdoc.

[1] National Academy of Science, National Academy of Engineering, Institute of Medicine, Enhancing the Postdoctoral Experience for Scientists and Engineers: A Guide for Postdoctoral Scholars, Advisers, Institutions, Funding Organizations, and Disciplinary Societies, National Academies Press, 2000.

February 15, 2017

James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Jim:

I write to express the commitment of LIGO to collaborate with the SciToken project that you are proposing in response to the NSF Cybersecurity Innovation for Cyber infrastructure (NSF 17-528) solicitation.

The SciToken project will address a real security risk inherent in LIGO data analysis pipelines. Currently, these pipelines require the use of an X.509 certificate or an RFC 3820 certificate proxy to perform routine tasks such as data movement. However, these certificates and proxies are authentication tokens that can be used for purposes other than the intended task, including for login access to various systems operated by LIGO and associated institutions around the world. While every effort is made to secure these certificates and proxies, there is always residual risk that must be accepted associated with this usage.

SciToken will mitigate this risk. The LIGO Security Team commits to reviewing the SciToken infrastructure and providing security feedback to the development team should this proposal be funded. If the expected reduction in residual risk is realized, the LIGO Security Team will require the use of SciToken authorization for LIGO data analysis pipelines where possible in place of the current authentication tokens.

I look forward to collaborating with you and wish you success with your proposal.

Sincerely,

Randy J Trudeau
LIGO Chief Information Security Officer

European Organization for Nuclear Research
Organisation européenne pour la recherche nucléaire

European Laboratory for Particle Physics
Laboratoire européen pour la physique des particules

Geneva, February 17th, 2017

Attention: James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Dr Basney,

We write to express the commitment of the CernVM project to collaborate with the SciToken project that you are proposing in response to the NSF Cybersecurity Innovation for Cyberinfrastructure (NSF 17-528) solicitation.

CernVM-FS is an integral part of most high energy physics experiments' software distribution strategy. CernVM-FS has demonstrated the ability to meet the massive scale of the LHC while providing an intuitive POSIX interface for physicists. Within the last year, we have been collaborating with your co-PI, Brian Bockelman of the University of Nebraska on CernVM-FS extensions to make it more applicable to other fields of science. For example, we have added X509-based authentication and authorization to meet the data security needs of distributing LIGO's physics data. We look forward to the addition of the more standard OAuth-based authorization mechanism proposed by the SciToken project.

In particular, we intend to continue the CernVM-FS project for the lifetime of the proposed new project and and it is our intent to commit adequate resources to provide code reviews, maintain release plans, integrate the relevant code in the CernVM-FS main repository. We commit to working with the SciToken team to further refine the existing authorization plugin mechanisms and shephard their research outputs into a CernVM-FS release.

We look forward to collaborating with you and wish you success with your proposal.

Sincerely,

Gerardo Ganis
Staff Physicist
Leader of the CernVM Project
Gerardo.Ganis@cern.ch

Jakob Blomer
Staff Computer Engineer
CernVM–FS author and main developer
Jakob.Blomer@cern.ch

Software Group in the CERN Physics Department
EP-SFT: http://ep-dep-sft.web.cern.ch

February 16, 2017

James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Jim:

It is with great enthusiasm that I write as PI and Project Director of the NSF XSEDE program to express XSEDE's intent to collaborate with the *SciTokens* project that you are proposing in response to the NSF Cybersecurity Innovation for Cyberinfrastructure (NSF 17-528) solicitation. Supported by National Science Foundation grant number ACI-1548562, XSEDE is a single virtual system that scientists can use to interactively share computing resources, data, and expertise to enhance the productivity of scholars, researchers and engineers. Its integrated, comprehensive suite of advanced digital services is designed to federate with other high-end facilities and with campus-based resources, serving as the foundation for a national e-science infrastructure ecosystem.

Secure distributed access to scientific data is an essential function of the XSEDE infrastructure. Today access control in XSEDE is primarily identity-based, via user accounts, certificates, Duo, and InCommon. The *SciToken* approach, using OAuth for capability-based (rather than identity-based) access to remote scientific data, can provide new options for XSEDE integration with remote data services (on campus, in the cloud, operated by virtual organizations, etc.). XSEDE has experience using OAuth for web single sign-on to components including CILogon, Globus, MyProxy, and the XSEDE User Portal.

Given the alignment of objectives, I am pleased to commit XSEDE to work closely with your team via XSEDE's Community Infrastructure (XCI) activity, to further develop the nation's cyberinfrastructure ecosystem and to further our goal of improving the productivity of the broad range of research communities XSEDE supports. We look forward to this collaboration.

Sincerely,

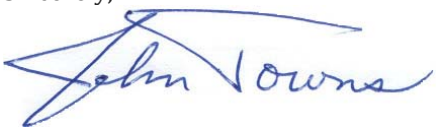

John Towns
PI and Project Director, XSEDE
Executive Director for Science & Technology, NCSA
Deputy CIO for Research IT, Office of the CIO
University of Illinois
jtowns@ncsa.illinois.edu

UNIVERSITY OF CALIFORNIA, SAN DIEGO UCSD

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO SANTA BARBARA • SANTA CRUZ

Frank Würthwein
DEPARTMENT OF PHYSICS
Mayer Hall, 5515
9500 Gilman Drive
La Jolla, CA 92093-0319

fkw@ucsd.edu
TEL: (858) 822-3219
FAX: (858) 534-0173

February 15, 2017

James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Jim:

I write to express the commitment of the Open Science Grid to collaborate with the SciToken project that you are proposing in response to the NSF Cybersecurity Innovation for Cyberinfrastructure (NSF 17-528) solicitation.

Open Science Grid is the premier platform for Distributed High Throughput Computing within the United States, supporting over 1.3 billion CPU hours of research computing a year. Historically, our authorization and authentication infrastructure has been based on Globus GSI and user X509 certificates. We have eliminated their need for many user workflows - but not those involving storage or data access. We are looking forward to the SciToken's modernized approach to authorization.

The Open Science Grid is a significant user and operator of CVMFS. We commit to continue running this infrastructure, packaging / distributing the CVMFS client for OSG sites, and upgrading the CVMFS infrastructure as needed to support this project. Additionally, we will deploy any new CVMFS components related to supporting LIGO's data file repository.

I look forward to collaborating with you and wish you success with your proposal.

Sincerely,

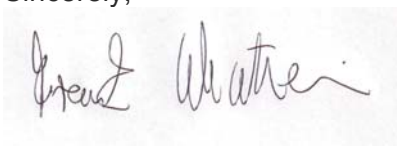

Frank Würthwein
Open Science Grid Executive Director
Professor of Physics, UCSD

USC
UNIVERSITY
OF SOUTHERN
CALIFORNIA

February 15, 2017

James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Jim:

I write to express the commitment of the Pegasus project (https://pegasus.isi.edu/) to collaborate with the SciToken project that you are proposing in response to the NSF Cybersecurity Innovation for Cyberinfrastructure (NSF 17-528) solicitation.

Many NSF research projects rely on the Pegasus workflow management system, including LIGO, which executes the PyCBC analysis pipeline using Pegasus. Pegasus manages PyCBC's cross-site data transfers and computations in a reliable, scalable, and efficient manner, across the LIGO Data Grid, Open Science Grid, and XSEDE, using HTCondor as the underlying scheduler. Secure and reliable management of credentials is critical for these workflows. Pegasus launches workflows with the needed credentials and detects when problems arise, due to credential expiration, authorization mismatch, or other runtime errors. The SciToken project's support for OAuth token refresh and least-privilege delegation will be a valuable improvement over current methods.

The Pegasus project will provide input to the SciToken design to ensure it provides the interfaces and capabilities needed by the advanced scientific workflows, such as LIGO's PyCBC, that Pegasus supports. We will plan to use the SciToken capabilities in future Pegasus releases.

I look forward to collaborating with you and wish you success with your proposal.

Sincerely,

Ewa Deelman, PhD
Research Professor, USC Computer Science Department
Research Director, Science Automation Technologies, USC Information Sciences Institute
4676 Admiralty Way, suite 1001
Marina del Rey, CA 90292
(310) 448-8408
deelman@isi.edu http://www.isi.edu/~deelman

Information Sciences Institute

Divisions

Advanced Electronics

Computational Systems and Technology

Informatics

Intelligent Systems

University of Southern California
USC Viterbi School of Engineering
4676 Admiralty Way
Suite 1001
Marina del Rey, California 90292-6601
Tel: 310 822 1511
Fax: 310 823 6714

Large Synoptic Survey Telescope

Dr. Steven Kahn, LSST Director | 950 N. Cherry Avenue, Tucson, AZ 85719 SKahn@lsst.org | www.lsst.org

February 15, 2017

James Basney
Senior Research Scientist, Cybersecurity Directorate
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 West Clark Street
Urbana, Illinois 61801

Dear Jim:

I write to express our commitment to collaborate with the SciToken project that you are proposing in response to the NSF Cybersecurity Innovation for Cyberinfrastructure (NSF 17-528) solicitation.

This project proposes a useful integration and is in line with our expected evolution of the ecosystem, which has been the basis of our plan. NSF support of scientific infrastructure software is vital for the Large Synoptic Survey Telescope (LSST) to achieve its scientific mission. HTCondor is the primary component in LSST's workflow management system and Pegasus is being seriously considered in addition to HTCondor. The proposal to extend these tools to support modern and flexible authentication and authorization technology provides an obvious advantage to LSST. LSST must provide access to computing resources and data to a very large and diverse user base. Giving more flexibility beyond X.509 certificates for managing user access is vital in providing data in a secure manner. To that end, we expect to establish dialogue with the project, present our use cases to the project, and use the newly-produced features consistent with our own plans.

LSST benefits from a strong partnership with NCSA's Cyber Security Division on topics including identity and access management. I look forward to our ongoing collaboration and wish you success with your proposal.

Sincerely,

Steven M. Kahn
LSST Director

List of Project Personnel and Partner Institutions

1. Basney, James; University of Illinois; PI
2. Blomer, Jakob; CERN; Collaborator
3. Bockelman, Brian; University of Nebraska-Lincoln; Co-PI
4. Brown, Duncan; Syracuse University; Co-PI
5. Deelman, Ewa; University of Southern California; Collaborator
6. Genis, Gerardo; CERN; Collaborator
7. Kahn, Steven; LSST; Collaborator
8. Tannenbaum, Todd; University of Wisconsin-Madison; Co-PI
9. Towns, John; University of Illinois; Collaborator
10. Trudeau, Randy; California Institute of Technology; Collaborator
11. Withers, Alex; University of Illinois; Co-PI
12. Würthwein, Frank; University of California at San Diego; Collaborator

Project Plan

The overall goals of the project are to:

- extend existing open source software infrastructure--CVMFS, XrootD, HTCondor, PyCBC, Pegasus--to support OAuth tokens for capability-based secure access to remote scientific data,
- engage with scientific research projects (LIGO and LSST) on evaluation and adoption of the platform,
- disseminate our results to the wider community.

PI Basney and co-PI Withers, together with the subawardees, will collaborate to achieve the above goals according to the following schedule of project milestones.

Project Period	**Project Milestones**
Y1Q1	OAuth: Issue and verify signed JWTs with filesystem scopes. CVMFS: Initial OAuth Bearer Token Support in CVMFS; prototype callout process.
Y1Q2	OAuth: LDAP plug-in for LIGO authorization. CVMFS: Access token verification in XrootD.
Y1Q3	CVMFS: changes submitted upstream.
Y1Q4	HTCondor: End-to-end HTCondor support for OAuth tokens. PyCBC: Output to XrootD and use Chirp for data movement (replacing GridFTP). CVMFS: support for "cvmfscp" and HTCondor file transfer plugins.
Y2Q1	CVMFS: Implement authenticated session caching in CVMFS client. PyCBC: Initial PyCBC workflow demonstration.
Y2Q2	HTCondor: OAuth support improvements: stage-in/stage-out. OAuth: Issue and verify Macaroons. CVMFS: Roll out new CVMFS client to pilot sites on OSG. Update callout process and XRootD auth for decentralized verification.
Y2Q3	HTCondor: OAuth support at all phases. Documentation on system drafted and disseminated. Present at HTCondor Week. CVMFS: Demonstrate PyCBC workflow using oauth tokens on production infrastructure (LDG, OSG). Integrate session caching into a CVMFS release.
Y2Q4	CVMFS: Roll out CVMFS updated client to official OSG release. Prototype on commercial cloud services (e.g., Dropbox). Journal article documenting design, implementation, and results.

Metrics

We identify the following success metrics for our project:

1. Adoption metric: Number of science projects adopting the SciTokens software (goal: 3 or more, including LIGO and LSST)
2. Adoption metric: Number of submission sites used for SciToken-enabled workflows (goal: 5 or more, including LIGO, LSST, and OSG sites)
3. Adoption metric: Number of execution sites where SciTokens enabled jobs have run (goal: 20 or more, including LIGO, LSST, and OSG sites)
4. Adoption metric: Number of external packages with contributions from the SciTokens project team (goal: 2 or more, including CVMFS)
5. Performance metric: Number of job submissions in 1 minute (goal: over 10,000)
6. Performance metric: Number of data source connections on job startup in 10 seconds or less (goal: 10 or more data sources, i.e., less than 1 second/authorization)
7. Performance metric: Sustainable data delivery rate to SciTokens users (goal: 20Gbps)
8. Software development metric: Number of SciTokens software releases (goal: 4 releases corresponding to above project milestones)

In addition to an outcomes assessment covering the above success metrics, we will report the results of micro- and macro-benchmarks for the SciTokens system, covering both JSON Web Tokens and Macaroons, in a journal article produced at the conclusion of our project.