OSG Workplan for Years 3-5

Overview

The workplan for the two out-years of the current phase of the Open Science Grid (OSG) project will continue to implement the commitment of the OSG consortium to advance and promote the adoption of Distributed High Throughput Computing (DHTC) technologies. The workplan assures the continuous evolution of the OSG production environment to meet the growing and changing computing needs of the U.S. LHC program, while facilitating the increasing demand for HTC capacity by the U.S. research communities at DOE labs and across university campuses. In addition, the workplan supports continuous engagement in a number of forward-looking initiatives to sustain and extend the impact of the services provided by the OSG on the broader research community.

Today, the Open Science Grid offers a diverse fabric of services that advances scientific discovery through improvements in computational throughput. The OSG is both a virtual facility distributed across more than 100 sites and an ecosystem. At its core is the concept of DHTC, the shared utilization of large ensembles of autonomous resources toward a common goal, where all of the elements are optimized for maximizing computational throughput.

At its core, the OSG provides a powerful and dependable platform and an ecosystem of DHTC capabilities, that enables Virtual Organizations (VOs) to run workflows and data storage systems across all OSG sites. This platform is essential for OSG's main stakeholders, the LHC experiments and other large experiments and VOs.

During the first two years of the current phase of the project we used the concept of "overlay schedulers" to establish the *OSG Direct* facility. This unique and powerful DHTC facility uses the capabilities of HTCondor and GlideinWMS to effectively harvest CPU resources opportunistically from all OSG sites, that stakeholder VOs would otherwise leave idle (for example, between production campaigns for major releases of experiment software). The *OSG Direct* service is used by a large and diverse community of researchers and science platforms. It is the basis for OSG being an XD Level 2 Service Provider to which science groups and individual principal investigators (PIs) acquire access through XRAC allocations. Finally it is used for specialized science gateways e.g. for biology and medical applications using web portals such as Galaxy.

The most recent development, extending but leveraging the components of *OSG Direct*, is the *OSG Connect* platform that provides a "retail" login and group management service for VOs and individuals with HTC workloads. Once signed-up and authorized, new users from small research communities gain access to services for job submission (HTCondor), storage, data transfer (Globus) and access (http, Posix) services so as to quickly on-board their workloads across the OSG sites, usually within minutes for simple pipelines. Additionally, using HTCondor-based flocking mechanisms with simple ssh public key security protocol, the platform allows small campuses with minimal IT staff to bring their clusters to the OSG ecosystem without installing and maintaining any new services. Finally, the concept has been generalized (*CI Connect*) to deliver campus grids as a service, allowing campuses to overflow jobs from their home clusters onto the OSG, to bridge with campuses having a sharing agreement, or for HPC resource targets using XRAC allocations. Campuses configured in this way thus become fully-fledged OSG campus grids without having to deploy any services locally. A small initial set of campuses are setup this way, and this approach provides an exciting path forward to extend the reach of OSG and DHTC for researchers on campuses.

OSG DHTC services also include user and host certificates (through the OSG CA, the follow-up of DOEgrids CA), software distribution services (the OASIS service, based on CVMFS), a distributed software environment modification service based on the popular Environment

Modules tool found on most campus and all XSEDE clusters, a growing repository of common scientific applications uniformly distributed by OASIS (including apps found in the XSEDE campus bridging yum repository), network monitoring and a network performance dashboard that provides an archive of network performance information, and several other services.

To deliver these broad and dependable fabric of services the OSG project is organized into the following major areas: Operations, Technology Evaluations and Software, and User Support. Additional areas or sub-areas include Campus Grids, Software Releases, Cybersecurity, Network Monitoring and Project Coordination. The OSG status report for year 1 and 2 (together with two annual reports) provide a detailed account of achievements and status. The following section describes the ongoing work in each area and current effort allocation that we consider to be the baseline workplan at the midpoint of the current phase of the OSG project.

OSG Baseline Workplan at the Midpoint of the Project

OSG Operations provides operations and maintenance for the evolving set of infrastructure services, operations support at the Grid Operations Center hosted by Indiana University's GRNOC, operations of the GlideinWMS systems at UCSD, as well as support for operations deliverables from other areas, such as from the cybersecurity and incident response team, network monitoring, software releases, running the OSG CA, the software distribution services, etc. The effort allocation for year 3 is 8.8 FTE.

OSG User Support provides consulting on technologies, architectures, and general assistance to users in adapting applications for the DHTC environment. It helps PIs and science groups to succeed using the *OSG Direct* service, either through XRAC allocations, through *OSG Connect* or through direct use of the facility. A general deliverable is to spread knowledge on high-throughput computing as a science problem solver. The effort allocation for year 3 is 3.0 FTE.

OSG Technologies and Software is to develop concepts and blueprints, and to deliver an evolving software stack, focusing on a small number of areas at a time, such as the HTCondor-CE, scalability, authorization, etc. The majority of the effort goes into the software "factory": patching, integration, packaging, and nightly testing of the components of the OSG software stack, thus making it ready for release, distribution, and deployment by sites. The effort allocation for year 3 is 7.8 FTE.

OSG Release Management is the ongoing effort to validate the updates made by the OSG Software team on test sites ("integration testing") and perform the actual release in conjunction with the OSG Operations team as well as participating sites in the Integration Testbed. The effort allocation for year 3 is 1.5 FTE.

OSG Campus Grids provides the new *OSG Connect* service platform described above, and runs the Campus Infrastructure Community, the forum to share best practices for connecting campus resources into the national cyberinfrastructure. The effort allocation for year 3 is 2.8 FTE.

OSG Networking provides readily available information about network performance, bottlenecks and problems for all OSG stakeholders and provides network metrics to higher-level services so they can make informed decisions about bandwidth usage. The effort allocation for year 3 is 0.3 FTE, pulling additional effort from the software and operations teams.

OSG Security Team guides and assists the OSG community in operating a secure fabric of DHTC services while maintaining the right balance between dependability and usability. It provides the following services: risk assessment, incident response, security testing and evaluation, identity and access management, security policies, and education and training. The effort allocation for year 3 is 2.1 FTE.

OSG Project Office and Communications provide all aspects of project and budget management support for OSG, as well as publishing monthly research highlights, periodic newsletters, the public-facing OSG web site, and updates to general documentation resources. The effort allocation for year 3 is 1.2 FTE.

Evolving the OSG Workplan for the 2nd Half of the Project

The OSG workplan for the 2nd half of the project continues to be driven mainly by the large VOs and resources owners. The LHC is preparing for Run2, and with it comes a large increase in computing demands. During the 2nd half of the OSG project we expect to see a factor of 2 increase in resource use and maybe more, and OSG is actively working with the LHC experiments to help with this transition. During this same period, the LHC collaborations will start to plan for the luminosity upgrade of the LHC (expected in ~ 2022) that will require computing innovations to help the transition from petabytes to exabytes of data volumes, and from 100,000 cores to 100M cores within the global DHTC systems over the next ten years. This will require new development and deployment approaches to DHTC workflows and distributed data management. As before, the OSG will play an important role in working with the experiments on their upgrade plans while ensuring the production infrastructure remains usable and robust.

We are pro-actively working with Intensity Frontier experiments as a new primary stakeholder, in collaboration with Fermilab, defining the future computing architecture for these experiments. OSG User Support is working to get Fermilab-related Intensity Frontier groups running on OSG one by one: NOvA, g–2, MicroBoone, mu2e, and LBNE (35Ton). OSG also works with Belle–II and with SNO+ to allow opportunistic use of resources. Belle-II and NOvA are routinely running monte-carlo (MC) production, SNO+ has run validation tests, MicroBoone is finalizing the code and testing for its MC Data Challenge in mid-August, and LBNE has used OSG for fast-simulation ahead of the DOE review last May. OSG is determined and committed to have all these experiments run efficiently on the OSG DHTC infrastructure.

Balance of benefits to the LHC experiments and to other parts of particle physics

OSG is deeply embedded into the U.S. LHC collaborations software and computing infrastructure. OSG services and software benefit from the LHC advances in needs and technology pushes, as well as joint activities and contributions. Other particle physics communities – including the emerging neutrino community – are increasingly active in deploying and operating a distributed computing model to benefit researchers at universities and laboratories that are part of their organizations. While their overall data and computing needs are not at the same scale as the LHC, there are specifics such as patterns of data analysis, use of ancillary databases etc., that will require changes or in some cases new services, all of which will need to be integrated, deployed and operated. Both OSG and collaboration staff are increasingly working together on these, with current examples being the deployment and evolution of the OASIS software distribution service, deployment of a common data distribution service for ancillary input data, and potentially services for dynamically provisioning cloud based resources. This model of collaboration allows for across-the-board prioritization and planning to take account of both the LHC and the other particle physics needs in a coordinated way, and to ensure that the key milestones and critical path items can be delivered to specification and in time.

Extending OSG for the "Long Tail of Science"

Since the very beginning of OSG, it has been our declared vision and goal to extend the benefits of DHTC to the broadest, most diverse set of scientific communities possible. In the early years, we approached this primarily via the concept of "Virtual Organizations" (VOs). We had expected that scientific communities would self-organize in similar ways to high energy physics, and support themselves by forming community organizations. OSG's role would then be to support those organizations. Within the last three years, we accepted the fact that, while successful in some communities, this was overall too limiting a goal. We understand now that to reach a broad spectrum of scientists at a variety of scales, from individual users at a single campus to multi-institutional experiments, we needed to diversify our support and access models. This led us to become an XSEDE service provider, to develop the concept of *OSG Direct* (login nodes at campuses and institutes that feed our provisioning infrastructure), to create the concept of *OSG Connect*, i.e. operate a service we offer to campuses to facilitate their research communities' access to DHTC, and to support flocking arrangements from campus and institute resources onto the OSG facility via the OSG operated GlideinWMS provisioning system. We have thus developed a varied portfolio of solutions based on a small set of core technologies in order to achieve our original vision and goals.

For the second half of the project, in addition to keeping up the ongoing work described above, we will focus in particular on enabling scientists and campuses for easy access to the OSG DHTC facility.

This will allow us to make OSG more available for the "Long Tail of Science", namely the large number of researchers that do not have access to sufficient in-house computing resources and/or the right level of skills, or that would profit from access to larger computing resources. There are an ever increasing number and diversity of small to mid-size collaborations of domain scientists that struggle with crossing order of magnitude boundaries in their computing, transitioning from gigabytes to terabytes of data volumes at scales of 1–10 million CPU hours. Over the next 10 years, we expect the number and diversity of groups at this end of the spectrum to increase as data and compute intensive science becomes the norm rather than the exception.

Already today the OSG successfully serves both the large experiments supporting them to utilize effectively their huge distributed set of resources, and the smaller groups at the campuses that do not own large resources but profit from DHTC. In our assessment, it is crucial to prevent the capability gap between the two extremes from growing. While "exascale" problems need to be solved for the large collaborations, it is equally necessary to ensure solutions are available for the many scientists challenged at the terascale and petascale on their passage to the exascale. To help our growing spectrum of users, we need to keep up with increasingly dynamic and heterogeneous environments while ensuring that domain scientists with limited computing expertise can use those environments .

We believe that the campuses will have to play a role in filling this gap and supporting scientists and groups to cross over to DHTC. Thus for our user community, including for the LHC, it is very important that we continue and increase to include the campuses. It will be of strategic importance to partner with projects that support campuses, like the ACI-REF project.

In the following description of workplans we highlight those elements that focus on providing opportunities for the "Long Tail of Science" explicitly.

OSG Operations

The primary important OSG deliverable is ongoing operation of the DHTC environment for the U.S. LHC and other U.S. researchers. This includes running the Grid Operations Center, a set of

Grid computing services including the GlideinWMS services at UCSD, providing the OSG PKI Certificate Authority (CA), the *OSG Direct* opportunistic access facility, the cyber security response team, etc. OSG has ongoing effort to support and maintain the OSG software distribution through a series of releases, incorporating changes required to develop functionality and to address any security issue.

In the first half of the project we established the *OSG Direct* facility to provide campus researchers with access to opportunistic resources, and PIs with access to XD allocations as allocated by the XRAC process, and this will be significantly extended in the coming years. A focus for the second half of the OSG project will be to extend and solidify the provisioning of opportunistic non-owned resources to science users. We expect a large peak-demand for additional resources from all large science users, and in particular from the LHC experiments, increasingly so in 2015 and beyond when LHC Run2 gains momentum and data sizes are rapidly increasing. In preparation for Run2, Atlas and CMS are starting to use OSG-provided technologies and approaches to enable "overflow" into non-owned resources for peak demands. As we grow the community of U.S. researchers who want to use OSG for their science, the project is working with resource providers to improve our ability to harvest opportunistic cycles and to make them available to all OSG users in a timely fashion.

Workplan Elements Operations

1) Continue to provide service desk support and ticket exchange

2) Operate OSG PKI service

3) Operate OSG Common services

4) Plan and implement needed upgrades and new Services

a) Upgrade Gratia Hardware

b) XRootD Monitoring

c) MaDDash/perfSONAR Datastore and Visualizations

5) Serve as interface for U.S. LHC experiments to WLCG & EGI

Added focus to provide opportunities for the "Long Tail of Science"

1) Plan and implement needed upgrades and new Services

a) Upgrade OSG-XD submission Node

b) Host iRODS service for distributed data storage and access

c) Transition from OASIS v1 to v2 and beyond

Technology Evaluations and Software Factory

The Technology area provides the OSG with the mechanism for long-term technology planning, and provides the software factory where software components that power the OSG fabric of DHTC services are evaluated, built, tested, configured, documented and packaged into a coherent and dependable software distribution.

Technology evaluation and further developing the OSG blueprint is conducted in the context of the domain sciences it will enable, and it is driven by specific challenges. OSG's existing capabilities are effective but basic and primitive compared to those necessary to support collaborative extreme-scale science in the 21st century. OSG needs to achieve new levels of

usability for portable, transparent services across an increasingly diverse and complex set of resource types serving a broader range of science domains with growing diversity of scientific computing skills. Statically federated resources (our classical view of the Grid) need to be integrated with dynamically allocated resources (like clouds) causing new challenges for resource planning, acquisition and provisioning. The ongoing increases in size, complexity and diversity of both the science applications and the computer technologies present significant research and technical challenges.

The OSG cyber-infrastructure is growing more and more diverse with the integration of new resource types, including public and commercial clouds, high-performance computers, etc. To manage such diversity, the project is exploring opportunities for improved resource provisioning mechanisms, with the goal to enable communities to organize OSG resources in tiers and define complex provisioning policies. OSG is evaluating methods for providing on-demand and allocation-based resources to its users, like access to high-performance computing centers including NERSC and XSEDE, and to community or commercial clouds like Amazon.

Moving forward OSG faces foundational changes in technologies, scales of use and collaborative heterogeneity. Computing hardware is undergoing a foundational change in the number of compute cores available on a single node; The emergence and prevalence of mobile computing interfaces causes foundational changes in the interaction of individual researchers with the computing environment; The size and reach of data transport and access are undergoing foundational changes in throughput, connectivity and heterogeneous sharing; More scientists use high level tools to develop complex applications out of community specific building blocks. The nature and diversity of the collaborative research groups are undergoing foundational changes in the breadth and inter-dependence of the research questions being addressed and computational tools employed.

Therefore, for OSG to address these challenges it needs an influx of innovative frameworks and technologies in the areas of data, security, systems, workflows, tools and collaborative environments, working on issues like transparent usage, provisioning and management of resources, maximizing throughput and total benefit, improving robustness, managing identity information and trust, improving usability and integration. Innovation in these broad but inter-related areas can be only accomplished through a coordinated and collaborative CS research effort developing new methods and capabilities to become part of the OSG ecosystem.

The OSG Technology area itself can target only a narrow set of high-priority items as given in the workplan elements below. Other continuing activities include running the OSG blueprint forum to develop the technology foundations and DHTC principles, and to evaluate new technologies as required by stakeholders and the developing OSG software stack.

Workplan Elements Technology/Software

1. Finish HTCondor-CE transition.
2. Address “common input data” problem.
3. More flexible software delivery (OASIS) to sites.
4. Support the LHC computing upgrades program and validate operation at 2X scale.
5. Support access to dynamically provisioned resources XSEDE, commercial and community clouds.
6. Maintain and update software platforms based on stakeholder requests.
7. Quarterly Blueprint meetings to address emerging DHTC topics.
8. Reduce the number of supported software components by 25%.

Added focus to provide opportunities for the "Long Tail of Science"

1) Integrate BOSCO / HTCondor-G+SSH into the production grid.
2) Support of workloads that require significant data on input, or produce data on output.
3) Data movement service based on Globus Online (at the OSG edge) and potentially other backhaul content distribution networks that transparently stream or cache user data near CPU resources.

User Support

The User Support area is to enable new communities to quickly adopt the OSG DHTC model and to improve productivity for all VOs as OSG service capabilities evolve. The area supports our aspirations to be "Open" towards supporting a diversity of research and science, taking into account requirements and modalities intrinsic to a domain. Leveraging efforts from all OSG areas, the team works to understand and facilitate troubleshooting of problematic or systemic failure modes that might be encountered by a particular VO's workflow. The project provides user support to individual scientists, to VOs and new communities interested in joining, helps with new applications and requirements, helps to integrate new sites, and runs the OSG VO environment providing opportunistic resources to campus users and PIs with XD XRAC allocations on OSG.

Workplan Elements User Support

1) Support DOE programs other than the High Energy Frontier in achieving maximum leverage of OSG platforms, services, and opportunistic resources. Special attention in the near term is on the Intensity Program, Cosmology and the Hall D program at JLab
    a. OASIS adoption
    b. Support IFF Data delivery plan
    c. Provisioning - preferential access to collaborating sites

Campus Grids and *OSG Connect*

One of the lessons learned during the first 5 years of OSG is that our reach into campuses is limited if all we provide is software that IT professionals then need to install and operate. The threshold for learning about and then maintaining services based on such software is too big a hurdle for many campus IT organizations with limited budgets and manpower.

In 2013/14 we thus started offering "OSG as a Service", which we refer to as *OSG Connect*. *OSG Connect* allows campuses to connect to OSG through a set of services hosted and operated by OSG staff. This allows campuses to start using OSG without having to deploy any special service or configuration.

OSG Connect "encapsulates" basic infrastructure services, from the network services (Globus Online etc), to identity management services allowing existing campus identities of researchers to be used across the full set of OSG services. Access to compute resources is based on HTCondor technologies, including the new BOSCO services developed by OSG. BOSCO is a self-installable and self-configuring robust tool that allows to submit from your local machine, including laptops, to compute globally across the OSG. *OSG Connect* also allows research communities to select services or tools that they require, without additional complexity.

This enables to setup easy access to a variety of facilities and resources. It allows campuses to create Campus Grids, provides campuses with the services to connect to the OSG and a useful set of services that help setting up a campus grid or extending an application into a campus grid: campuses don't have to build their own — OSG is running these services for them.

One aspect of this approach is that it lends itself to extending the OSG eco system to the HPC community. First interfaces to HPC installations using the *OSG Connect* approach are emerging: ATLAS is pioneering this with XSEDE (the Stampede cluster at TACC), and CMS is using BOSCO to enable their job execution environment to access the NERSC DOE facility.

The *OSG Connect* approach is gaining significant momentum now. The *OSG Connect* set of services will be further developed to connect researchers and campuses. Also the LHC experiments are interested in this approach to connect the campuses of collaborating institutions, and both ATLAS and CMS have started to base Tier–3 computing on the *OSG Connect* technology (ATLAS Connect and CMS Connect). Now that we have established the usefulness of this approach we will further develop it during the second half of the project.

Workplan Elements Campus Grids & User Support, focus to provide opportunities for the "Long Tail of Science"

1. Provide seamless access/usability for researchers based on *OSG Connect*; make it easy for campus researcher to join OSG, run jobs and move data.
    a. Mature and extend the *OSG Connect* services
    b. Target particular campus opportunities: e.g. 3/6 ACI-REF campuses
    c. Target partner communities: LHC, others
2. Grow the OSG opportunistic facility
    a. Provide the *OSG Direct* facility as an easy to use OSG "On-ramp" for sending opportunistic jobs to the OSG production fabric
    b. Increase throughput, manageability, and observability of this facility
    c. Research opportunistic eco-system and develop and implement recommendations for growing the accessible opportunistic pool toward supplying up to 20% of total OSG usage
3. Ongoing support of OSG as a high-quality Level 2 Service Provider in XD and grow number of XRAC requests against OSG
    a. Use XSEDE Campus Champions as a channel to more U.S. campus researchers
    b. Provide tutorials and documentation on how to structure jobs for use in OSG and how to submit jobs
4. Simplify software and execution environments
    a. Extend the supported use/services/science applications
    b. Support for special processing needs
    c. Large RAM support, long runtime support, etc. i.e. apps that are slightly out of our main sweet spot

Security

Providing the security team services in the context of a DHTC environment requires a framework that captures voluntary trust relationships between the participants sites, users, and software providers that may operate under different security models. The inherent heterogeneity of a shared environment that is composed of autonomous principles requires means to codify mutual expectations when one party relies on another party to perform specific actions in reaction to certain stimuli and to audit these interactions.

Work on extending trust relationships will continue. The recent successes of the security team to introduce job traceability capabilities to replace certificates will be extended. The project is seeking mechanisms to improve identity management to enable resource access as seamlessly as possible. These efforts range from certificate-less access to resources, to the reductions of host

certificates, to better integration of central user identity information with the campus grid (*OSG Connect*) submission mechanisms.

Workplan Elements Security

1. Provide operational security
2. Decrease dependence on PKI certificates, both user and host certs; investigate certificate-free operational methods for data delivery services
3. Plan and deploy evolution for OSG identity management services beyond the expiration of the current OSG PKI service (using DigiCert) in June 2016
4. Enable use of campus identities for access to OSG resources.

Network Monitoring

Work on network monitoring will continue, providing important network performance information to applications including the LHC, and providing an invaluable resource of network related performance parameters for networking researchers, networking providers and users. We are working on integrating tools from the community to increase network awareness. We plan to collect and store metrics associated with network connectivity in OSG. Using network management probes at OSG sites through the PerfSONAR-PS, we are planning to release a dashboard with the goal of gathering, storing, and publishing current and historical network metrics for alarming, troubleshooting, and higher-level service integration.

Workplan Elements Network Monitoring

1. Develop a fast, robust, maintainable datastore for OSG networking
2. Support higher-level services by providing API for users and applications to access the network information they need from the datastore
3. Develop alarming and alerting for network problems

Promoting use of DHTC for Science

The OSG is the exemplar of a DHTC facility and HTC is a paradigm used and explored broadly throughout DOE and NSF science, including in genomics, structural biology, high energy physics, nuclear physics and other fields. Effective DHTC requires making use of a diverse set of resources to achieve maximum capacity and cost effectiveness, and establishing mutual trust among the participants while still respecting their autonomy. Achieving these DHTC attributes presents significant research and technical challenges as the applications and technologies increase in size and diversity.

OSG has limited resources to provide the intellectual resources, expertise, and solutions needed to advance the computing throughput of DOE and NSF science, but has a number of important activities in this area, that are crucial to further develop the adoption and use of DHTC across the OSG stake holders.

Workplan Elements to Promote use of DHTC for Science

1. Focus Groups, provide technical leadership/guidance LHC, IF, others
2. Campus Infrastructure Communities, workshops etc
3. Evolve partnership XD/XSEDE
4. Establish partnership with ACI-REFs to provide 2nd layer support for them as they include OSG in their solutions for local campus researchers
5. Tutorials and documentation

6. OSG User School
7. Publish OSG Research Highlight monthly

Workplan Summary

The base-line effort distribution was shown in the table, with a total FTE count of about 27.5. Following the workplan described above the effort allocation will be adjusted incrementally to address the focus areas, at the level of fractions of FTEs between the areas, while we expect the overall effort envelope to stay about constant.

1. Provide value to VOs and resource owners by enabling DHTC and resource sharing for their set of applications

Technology/Software

1. Finish HTCondor-CE transition
2. Address "common input data" problem
3. More flexible software delivery (OASIS) to sites
4. Support the LHC computing upgrades program and validate operation at twice the current scale
5. Support access to dynamically provisioned resources XSEDE, commercial and community clouds
6. Maintain and update software platforms based on stakeholder requests
7. Quarterly Blueprint meetings to address emerging DHTC topics
8. Reduce the number of supported software components by 25%

Operations

1. Continue to provide service desk support and ticket exchange
2. Operate OSG PKI service
3. Operate OSG Common services
4. Plan and implement needed upgrades and new Services
    a. Upgrade Gratia Hardware
    b. XRootD Monitoring
    c. MaDDash/perfSONAR Datastore and Visualizations
5. Serve as interface for U.S. LHC experiments to WLCG & EGI

Network Monitoring

1. Develop a fast, robust, maintainable datastore for OSG networking
2. Support higher-level services by providing API for users and applications to access the network information they need from the datastore
3. Develop alarming and alerting for network problems

Security

1. Provide operational security
2. Decrease dependence on PKI certificates, both user and host certs; investigate certificate-free operational methods for data delivery services

3. Plan and deploy evolution for OSG identity management services beyond the expiration of the current OSG PKI service (using DigiCert) in June 2016
4. Enable use of campus identities for access to OSG resources.

User Support

1. Support DOE Intensity Frontier programs in achieving maximum leverage of OSG platforms, services, and opportunistic resources
    a. OASIS adoption
    b. Support IFF Data delivery plan
    c. Provisioning - preferential access to collaborating sites

2. Provide opportunities for the long tail of science

Campus Grids & User Support

1. Provide seamless access/usability for researchers based on *OSG Connect*; make it easy for campus researcher to join OSG, run jobs and move data.
    a. mature and extend the *OSG Connect* services
    b. Target particular campus opportunities: e.g. 3/6 ACI-REF campuses
    c. Target partner communities: LHC, others
2. Grow the OSG opportunistic facility
    a. Provide the OSG Open Facility as an easy to use OSG "On-ramp" for sending opportunistic jobs to the OSG production fabric
    b. Increase throughput, manageability, and observability of this facility
    c. Research opportunistic eco-system and develop and implement recommendations for growing the accessible opportunistic pool toward supplying up to 20% of total OSG usage
3. Ongoing support of OSG as a high-quality level 2SP in XD; grow number of XRAC requests against OSG
    a. Use XSEDE Campus Champions as a channel to more U.S. campus researchers
    b. Provide tutorials and documentation on how to structure jobs for use in OSG and how to submit jobs
4. Simplify software and execution environments
    a. Extend the supported use/services/science applications
    b. Support for special processing needs
    c. Large RAM support, long runtime support, etc. i.e. apps that are slightly out of our main sweet spot

Technology & Software

1. Integrate BOSCO / HTCondor-G+SSH into the production grid. Year 3
2. Support of workloads that require data on input, or produce data on output
3. Data movement service based on Globus Online (at the OSG edge)

Operations

1. Plan and implement needed upgrades and new Services
    a. Upgrade OSG-XD submission Node
    b. Host iRODS service for data delivery
    c. Transition from OASIS v1 to v2 and beyond

3. Promote use of DHTC for science

1. Focus Groups, provide technical leadership/guidance LHC, IF, others?
2. Campus Infrastructure Communities, workshops etc
3. Evolve partnership XD/XSEDE
4. Establish partnership with ACI-REFs to provide 2nd layer support for them as they include OSG in their solutions for local campus researchers
5. Tutorials and documentation
6. OSG User School
7. Publish OSG Research Highlight monthly

Milestones and Deliverables

Goals and deliverables are separated into sustaining and extending efforts. We make more detailed plans annually taking into account lessons learned and the evolving needs of the OSG stakeholders.

Ongoing Operations:

- Operate service desk support, ticket exchange, and accounting interface to WLCG for U.S. LHC community
- Operate OSG Common services at SLA levels to enable the stakeholder's production plans
- Operate OSG PKI service
- Provide operational security
- Maintain and update software platforms based on stakeholder requests
- Ongoing support of OSG as a high-quality level 2 Service Provider in XD
- Conduct Annual OSG User School for ~25 students and researchers

Year 3:

- Finish HTCondor-CE transition and Integrate BOSCO / HTCondor-G+SSH into the production grid. (Technology, Software)
- Provide for needs of LHC in 2015 at 2X scale (all)
- Deploy first production release of Monitoring Datastore (Operations, Networking)
- Support access to dynamically provisioned resources on one XSEDE HPC system and one DOE HPC system, commercial and community clouds (all)
- Have all running IF experiments use OSG as needed for simulation (all)
- Establish partnership with ACI-REFs to enable broader access to DHTC to for U.S. campus researchers (User Support)
- Add 3 campuses using OSG-Connect technology to access the OSG production fabric for DHTC access (Campus Grids)

Year 4:

- Provide production access to full network monitoring data and visualizations (Operations, Technology/Software, Networking)
- Multi-community support of workloads that require large data on input, or produce large data on output (Operations, User Support)

- Provide production access to dynamically provisioned resources on commercial and community clouds as well as multiple HPC systems. (all)
- Provide services to access the network information by users and applications (Networking, Operations)
- Provide Identity Management services beyond the expiration of the current OSG PKI service (using DigiCert) in June 2016 (Security, Operations)
- Enable use of campus identities for access to OSG resources. (Security, Operations)
- Production support in OSG Connect for 1-2 users with specific (out of current OSG sweet spot) application and processing needs. (Campus Infrastructure)

Year 5:

- Production support for new hardware architectures requester by stakeholders (all)
- Implement recommendations (based on research) to grow the opportunistic pool towards 20% of OSG usage (User Support)
- Initial deployment of alarming and alerting for network problems (Networking, Operations)