



# Overview of PAtH

**Frank Würthwein**  
**OSG Executive Director**  
**Professor of Physics**  
**UCSD/SDSC**

**December 8<sup>th</sup> 2020**



# Partnership to Advance Throughput Computing (PATH)

Cooperative Agreement with NSF OAC CI

## Scope in a nutshell:

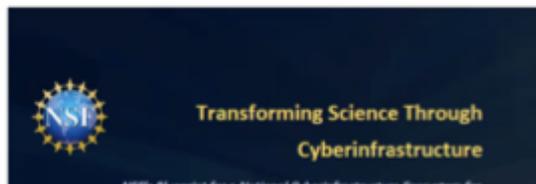
- HTCondor team
- Fabric of Services to support OSG Consortium beyond IRIS-HEP
  - Services instantiated via software
  - Services provided by people => Research Facilitation
- Global Infrastructure Laboratory
  - Ingest of ideas & software into PATH



# NSF Vision for PATH



## NSF's Vision for a National CI Ecosystem

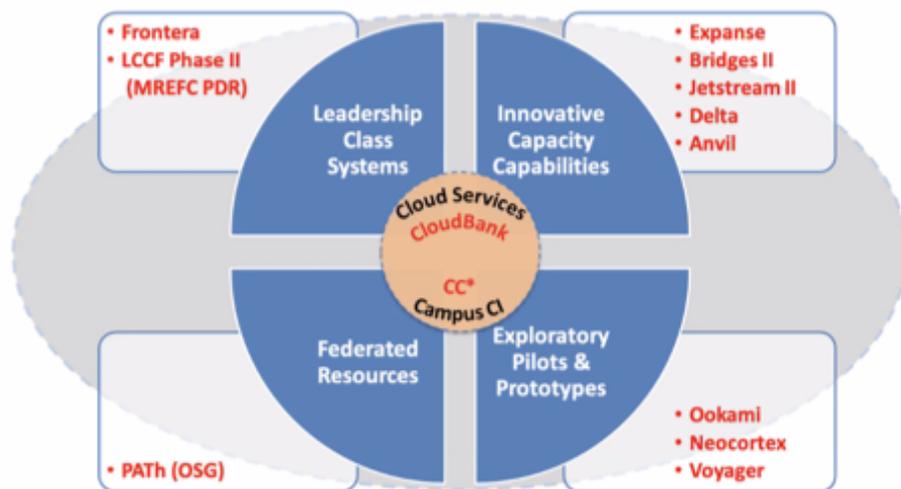


NSF's vision for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century

## for a National CI Ecosystem

*An agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of S&E research and education.*

### Computational Ecosystem: Elements



### Overarching principles:

- View CI more holistically
- Support translational research
- Balance innovation with stability
- Couple discovery and CI innovation cycle
- Improve usability

<https://www.nsf.gov/cise/oac/vision/blueprint-2019/>

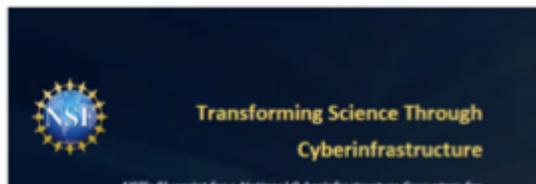
From Manish Parashar presentation to XSEDE SP last week.



# NSF Vision for PATH



## NSF's Vision for a National CI Ecosystem

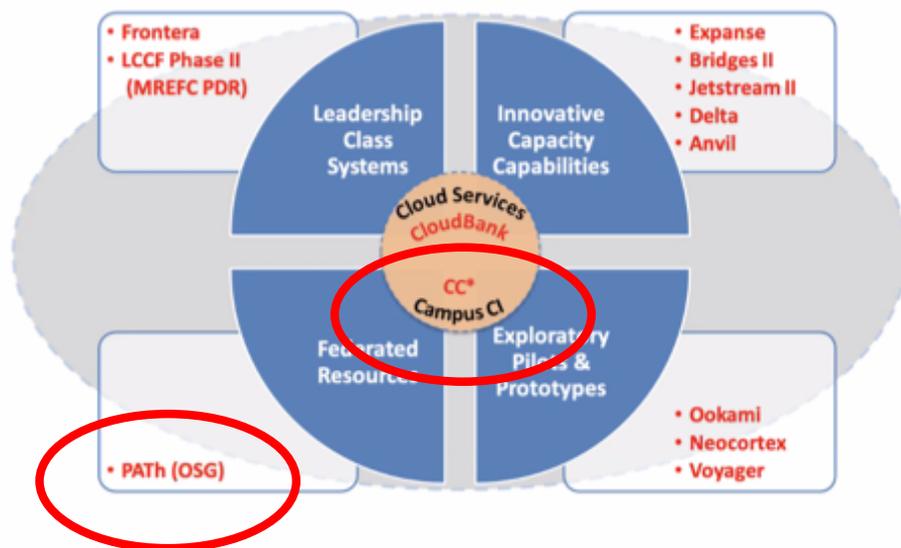


NSF's vision for a National Cyberinfrastructure Ecosystem for Science and Engineering in the 21st Century

*An agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of S&E research and education.*

## for a National CI Ecosystem

### Computational Ecosystem: Elements



### Overarching principles:

- View CI more holistically
- Support translational research
- Balance innovation with stability
- Couple discovery and CI innovation cycle
- Improve usability

<https://www.nsf.gov/cise/oac/vision/blueprint-2019/>

**PATH is strongly aligned with the CC\* program**

**From Manish Parashar presentation to XSEDE SP last week.**



# NSF Vision for PATH



## NSF's Vision for a National CI Ecosystem

PATH is expected to provide glue to connect CC\* awards.

Significant engagement role with campuses of all sizes.

PATH is also expected to coordinate with other "coordination services" ... XSED++, XDMoD++, ....

for a Nat

Con

- Fronte
- LCCF P
- (MRE

Resour

- PATH (OSG)

**PATH is strongly aligned  
With the CC\* program**

From

- The things relevant to the LHC that US ATLAS and US CMS Ops programs can agree on.
  - OSG software stack & support to US LHC T1, T2
  - Operational Security
  - Operations of Accounting et al.
  - Networking Performance Measurements

<b>Operations = UNL</b>	= 1FTE
<b>Security = Indiana University</b>	= 0.8FTE
<b>Software = U. Wisconsin – Madison</b>	= 3FTE
<b>Networking = U. Michigan</b>	= 0.4FTE

**5.2 FTE across ~12 people**



# **Excerpts from the PAtH proposal**

# Services Instantiated via Software (I)

**Compute Management.** The platform for delivering compute resources as effective capacity from existing raw capacity is based on the concept of resource provisioning and Annexes (Section 4.1.1). In the resource provisioning model, a researcher places jobs into a submit endpoint and each job describes the required execution environment. A resource provisioner determines and executes a strategy for acquiring resources into a pool of resources to be utilized by the workload. The services making up the OSCF include:

- **Submission endpoints** for providing researchers access to the shared compute power enabled by FoCaS. Each consists of an HTCondor submit host, scale-out storage, and identity management allowing login via SSH. Submission endpoints may be for the FoCaS-operated OSG Connect service or hosted by or for an individual campus, science collaboration, or gateway. FoCaS will initially have the responsibility for operating seven existing submission endpoints.
- **Resource pools** for aggregating resources on behalf of an organization, via an HTCondor central manager. A submission endpoint claims resources from a pool on behalf of jobs. PATH will initially operate four pools. The “Open Science Pool” [147] will continue with 35 attached submission endpoints: seven operated by FoCaS and 28 by campuses and science gateways.
- **Computing Entry points (CEs)** for contributing resources via the scheduler on a campus cluster into a resource pool, using the HTCondor-CE software (a special configuration of HTCSS). CEs can be operated by a campus or by FoCaS, on FoCaS hardware or on campus hardware.
- **Resource Provisioning and Annex components** for provisioning resources – from hosted CEs, XD machines, or cloud capacity – into the resource pools. The OSG Consortium hosts an instance of the provisioning software (*glideinWMS* [148]) for all communities with access to CEs worldwide.

## Supporting the OSG Compute Federation

# Services Instantiated via Software (II)

**Data & Software Management.** To support the execution of dHTC workloads, FoCaS will grow the Open Science Data Federation (OSDF), a content delivery network for all Open Science. Current software components include a *data endpoint* (a source or sink for data) and *data caches*. FoCaS will offer software tools for policy-based data placement and user software distribution capabilities. A **data endpoint** makes research data within a storage service (e.g., a cluster filesystem at a submit point) available to dHTC jobs over the wide-area network (WAN). The OSG data endpoint is based on **XRootD**. Currently, seven data endpoints currently serve as endpoints. **Data caches** provide on-demand data movement to jobs. A data cache is located near (in network terms) compute resources and manages the usage of WAN bandwidth to mitigate latency impact. For workloads with high locality of reference, caches allow for data reuse and reduction of total WAN bandwidth and load on data origins. The OSDF presently includes 15 data caches at network POPs and compute endpoints, distributed across the US, Europe, and Asia, mostly managed through distributed or federated k8s. The data federation can be accessed via HTTP using the data staging support in HTCSS for user jobs, using the command-line, or as a read-only POSIX filesystem [151]. The Software Development team will sustain Data Federation integration from its externally sourced components, and integration into HTCSS supported tools, enabling developers to work side-by-side with operators toward improving the management and scheduling of data transfers and job execution in HTCSS (Section 4.1.1). Larger collaborations have demonstrated the need for **policy-based data placement** to manage placement between passive archives, active storage, and across endpoints to make more data available or to replicate to protect against loss. To orchestrate transfers, PATH will operate a **Rucio** [152] instance, allowing an organization to maintain a catalog of its files, organize files into datasets, establish placement policies, and orchestrate transfers with an external service.

**Supporting the OSG Data Federation ... and then some ...**

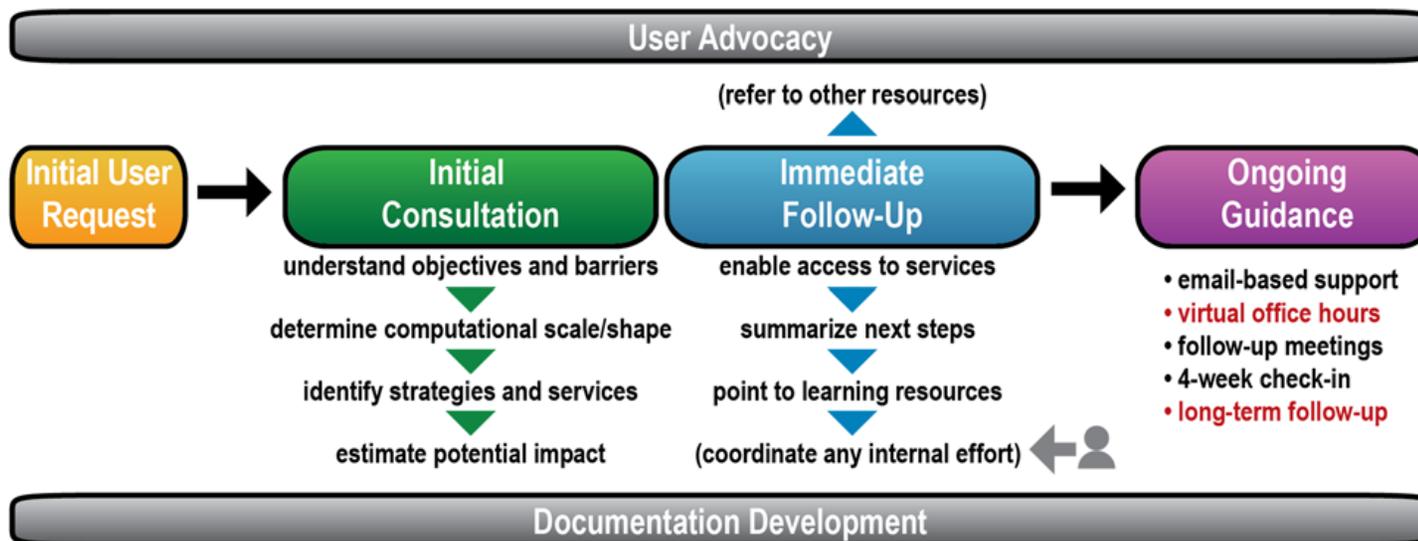


Figure 2: PATH Facilitator Activities. Newly proposed activities are highlighted in red.

**Specialized Workflow Consulting: Collaboration Support.** In-depth consulting on complex processing workflows will be offered to multi-campus collaborations and science gateways when they have a long-term need for automated data production or processing involving coordination between multiple campuses and implementation of the workload requires close integration with FoCaS. This activity will follow a standard pattern: consult to identify services, deploy and integrate the services into a federated platform, and jointly operate the platform. Collaborations will also receive ongoing facilitation and technical support at periodic check-ins, including mid-task operational assessments and post-mortem summaries. A single collaboration support activity may impact hundreds to thousands of researchers. This justifies the effort to perform infrastructure integration activities for a collaboration exceeding the effort that would be offered to a single group. However, long-term, the requesting collaboration is fully responsible for the workflow, as necessary for sustainable growth in the number of engaged collaborations.

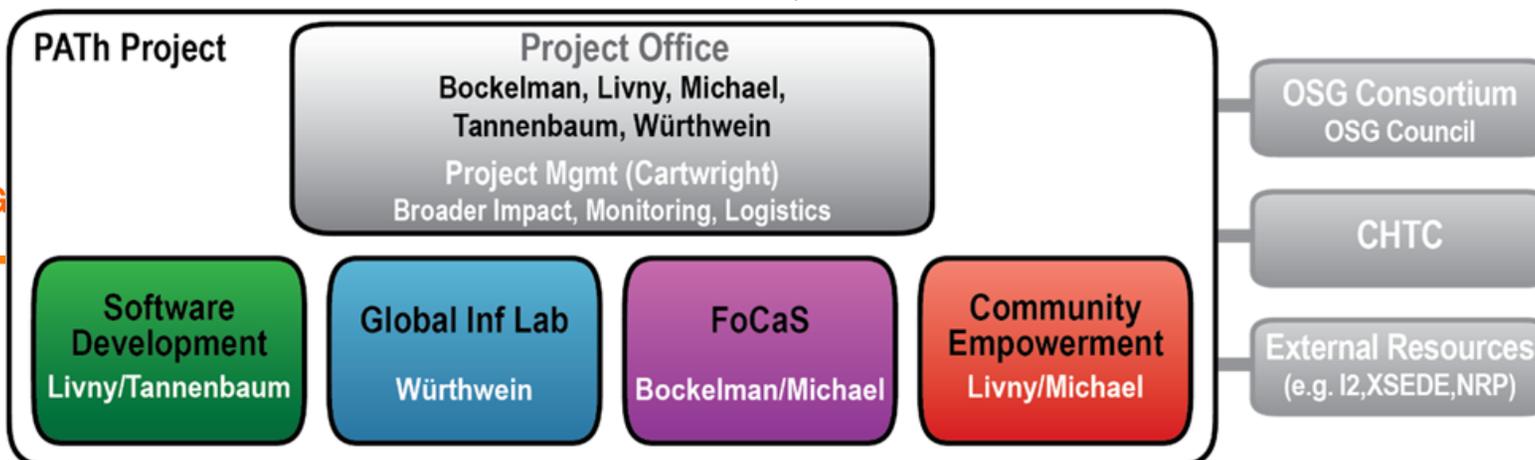


Figure 3: PATH Organization

Area	Activity	FTE Level	Area	Activity	FTE Level
<b>PATH Software Del.</b>	Core Development	6.45	<b>FoCaS</b>	Service Operations	2.40
	Integration & Release Mgmt	1.00		Infrastructure & HTCSS Beta Testing	1.30
	User Experience	0.60		Operational Security	0.50
	Leadership	0.35		Research Facilitation	2.76
	<b>Totals</b>	<b>8.40</b>		Collaboration Support	1.00
<b>GIL</b>	Investigations	0.83		Metrics, Measurement & Asset Management	0.66
	Leadership	0.10		Leadership	0.53
	<b>Totals</b>	<b>0.93</b>		<b>Totals</b>	<b>9.34</b>
<b>Project Office</b>	Logistics & BI Coordination	0.25		<b>Empowerment</b>	Event Coordination
	Project Manager	0.52	Communications		0.25
	Leadership Team	0.41	<b>Totals</b>		<b>0.55</b>
	<b>Totals</b>	<b>1.14</b>			
<b>Grand Total</b>					<b>20.36</b>

Table 1: A snapshot of effort within PATH by area and activity once the project is fully ramped.



**Comments & Questions ?**