APS Scientific Computing Strategy: APS-U Feature Beamline Interfaces

July 10, 2020

Introduction

The Advanced Photon Source (APS) at Argonne National Laboratory (ANL) is a U.S. Department of Energy (DOE) Office of Science-Basic Energy Sciences (BES) scientific user facility. The core mission of the APS is to serve the scientific community by providing experiment facilities utilizing x-rays (beamlines) to allow users to address the most important basic and applied research challenges facing our nation.

All aspects of APS operation depend on computation, but data analysis software and computing infrastructure are of particular importance for facility productivity. Demands for increased computing at the APS are driven by new scientific opportunities, which are enabled by new measurement techniques, technological advances in detectors, multi-modal data utilization, and advances in data analysis algorithms. The priority for the APS is to further improve our world-class programs that benefit most from high-energy, high-brightness and coherent x-rays. All of these require advanced computing. The revolutionized high-energy synchrotron facility proposed in the APS-U project will increase brightness and coherence, leading to further increases in data rates and experiment complexity, creating further demands for advanced scientific computation.

The APS and ANL are poised well to employ advanced computing to maintain a world-leading position in the synchrotron community. The APS has a world-class photon science program with a large and diverse user base, and ANL is home to world-leading supercomputing infrastructure and computer science expertise in the Computing, Environment, and Life Sciences directorate (CELS). This colocation provides an unprecedented opportunity for collaboration.

The APS has organized the core groups required to achieve these goals under the X-ray Science Technologies (XST) umbrella within the X-ray Science Division (XSD). The XSD Beamline Controls (BC) group is responsible for beamline data acquisition, through control and operations systems and software. The XSD Computational X-ray Science (CXS) group is mainly responsible for the development of theory, mathematical models, algorithms and software for interpreting x-ray measurements. The XSD Scientific Software Engineering & Data Management (SDM) group is responsible for software engineering for data analysis applications and data management tools, enabling high-performance computing (HPC). The management and support of information technology resources within the APS is handled by the APS Engineering Support (AES) division Information Technology (IT) and Information Solutions (IS) groups.

This document identifies the core challenges facing the APS in this area, specifically for the APS-U feature beamlines, and presents a plan for addressing these challenges over the coming years within the anticipated resource environment at the APS, and by leveraging synergies with other parts of ANL, the Center for Advanced Mathematics for Energy Research Applications (CAMERA), across the BES and DOE complex, and with the outside community. Each group maintains strategies and work plans listed at the end of this document that describe activities in greater detail.

Data Acquisition and Experiment Control

A powerful and up-to-date data acquisition and experiment control system is key to realizing the potential of the APS. In order for the next generation of instruments in the APS-U era to help answer the most pressing scientific challenges, the facility must support the most advanced and flexible control and beamline operation systems. Despite the need for flexibility, employment of standardization and pre-existing solutions are required to take advantage of economies of scale.

Currently, data acquisition and experiment control tools at the APS vary depending on particular technique and experiment requirements. SPEC is by far the most widely used experiment control software. Other tools, such as APSPy, JBlueIce, TXM, and uProbeX, serve a specific user base.

For the feature beamlines planned as a part of the APS-U project, and for new beamline improvements, the APS will adopt the bluesky data acquisition and experiment control system from the National Synchrotron Light Source II (NSLS-II) at Brookhaven National Laboratory (BNL). The bluesky system provides a flexible data acquisition and experiment control environment to coordinate detectors, motors, shutters, timing signals, and other beamline instrumentation, and it supports fast fly scans and data acquisition rates and provides experimental feedback.

The APS collaborates closely with bluesky's development team at NSLS-II. The majority of core development effort for bluesky was provided by the NSLS-II. APS operations funding will provide resources for the deployment and customization of bluesky for the APS-U feature beamlines and for new beamline improvements.

This effort will fulfill the needs described in the following APS-U beamline functional requirements and interface control documents: ASL Requirement 5.1, ATOMIC/3DMN Requirement 5.1, CHEX Requirement 5.1.4, CSSI Requirement 5.1.5, HEXM Requirement 5.1.6, ISN Requirement 5.1.9, PtychoProbe Requirement 5.1.7, POLAR Requirement 5.1.10, and XPCS Requirement 5.1.7.

Data Analysis Algorithms and Software

The APS is focusing data analysis algorithm and software development in the areas needed to answer the novel scientific inquiries enabled by the future APS and the planned APS-U project, which are the techniques driven by coherence, imaging, and high-energy, as well as multi-modal techniques. The table that follows lists the areas targeted for data analysis algorithm and software advancements critical for operation of the APS-U feature beamlines. For each algorithm and software requirement listed in the APS-U beamline functional requirements and interface control documents, the software efforts that fulfill the respective requirement is listed followed by the funding source.

APS-U Feature Beamline Requirement	Software Requirement	Software Package	Funding
ASL 5.1	X-ray fluorescence microscopy elemental mapping	XRF-Maps	APS Operations
		PyMCA	ESRF
		GeoPIXE	CSIRO
	XAFS, EXAFS, XANES, XES, and RIXS analysis	Demeter	NIST
		RIXS Plot	APS Operations
		MiniXES	APS Operations Univ. of Washington
ATOMIC/3DMN 5.1	3D Bragg coherent diffraction imaging	Bragg CDI	APS Operations ANL LDRDs
	Space-mapping	RSMap3D	APS Operations
	Laue depth reconstruction	LaueGo	APS Operations
CHEX 5.1.5	X-ray photon correlation spectroscopy software that processes multi-tau, two-time, higher-order, and spatial-temporal correlations	XPCS-Eigen	APS Operations LBNL CAMERA
	Ptychographic reconstruction software that processes time-resolved and polychromatic data	PtychoLib	APS Operations ASCR/BES Grant Northwestern Univ. IARPA
		SHARP	LBNL CAMERA
		tike	APS Operations
		xscale	LBNL CAMERA
	3D Bragg coherent diffraction imaging	Bragg CDI	APS Operations ANL LDRDs
	Grazing-incidence small angle x-ray scattering and coherent surface scattering reconstruction	GI-SAXS	APS Operations ANL LDRDs
		HipGISAXS	LBNL CAMERA
CSSI 5.1.6	Spatial-temporal correlation	XPCS-Eigen	APS Operations
	Grazing-incidence small angle x-ray scattering and coherent surface scattering reconstruction software	GI-SAXS	APS Operations ANL LDRDs
		HipGISAXS	LBNL CAMERA
HEXM 5.1.7	Far- and near-field diffraction	MIDAS	APS Operations ANL LDRDs Industrial Partners
	Imaging and diffraction tomography	TomoPy	APS Operations ANL LDRDs LBNL CAMERA
	3D Bragg coherent diffraction imaging	Bragg CDI	APS Operations ANL LDRDs
ISN 5.1.10	Ptychographic reconstruction	PtychoLib	APS Operations ASCR/BES Grant Northwestern Univ. IARPA

		SHARP	LBNL CAMERA
		tike	APS Operations
		xscale	LBNL CAMERA
	X-ray fluorescence microscopy elemental mapping	XRF-Maps	APS Operations
PtychoProbe 5.1.8	Ptychographic reconstruction	PtychoLib	APS Operations ASCR/BES Grant Northwestern Univ. IARPA
		SHARP	LBNL CAMERA
		tike	APS Operations
		xscale	LBNL CAMERA
	X-ray fluorescence microscopy elemental mapping	XRF-Maps	APS Operations
POLAR 5.1.11	Scattering and diffraction data processing	RSMap3D	APS Operations
	X-ray fluorescence microscopy elemental mapping	XRF-Maps	APS Operations
	3D Bragg coherent diffraction imaging	Bragg CDI	APS Operations ANL LDRDs
	Dichroic data processing	XMCD	APS Operations
XPCS 5.1.8	X-ray photon correlation spectroscopy software that processes multi-tau, two-time, higher-order, and spatial-temporal correlations	XPCS-Eigen	Operations LBNL CAMERA

Table 1 - Requirements for data analysis algorithms and software for the APS-U era. Table describes APS-U beamlines that will require particular algorithms and software, current efforts including the names of tools that meet the requirements, and how the efforts are funded.

These techniques already number amongst the most computationally intensive techniques performed at the APS and throughput demands are expected to grow by as much as multiple orders of magnitude, due to improved detectors and the upgraded source. Data reduction and analysis will rely heavily on the use of high-performance computing (HPC), utilizing appropriate technologies such as multi-threading, MMPs and GPUs, and distributed computing environments to obtain results with near real-time completion, so that results allow user-driven or even automated steering of experiments.

Most software will largely be developed as open source and will be made available with user community code contributions encouraged. A graded approach according to impact and priority will be applied to development. Packaging and active support either as distributable applications or as Software-as-a-Service (SaaS) will be provided for software systems that have been deemed to be most important for the success of APS users. Beamlines not directly driving the APS-U will also benefit from the reuse of tools developed for priority applications.
Key accomplishments have been made in this area over the past years. The TomoPy reconstruction toolkit, the XPCS-Eigen correlation software for x-ray photon correlation spectroscopy, the HPC-enabled Microstructural Imaging using Diffraction Analysis Software

(MIDAS) for high-energy diffraction microscopy, the multi-threaded XRF-Maps software for elemental mapping, the distributed GPU-based PtychoLib software for ptychography, the Bragg CDI package for 3D Bragg coherent diffraction imaging, the RSMap3D application for space-mapping scattering and diffraction data, the LaueGo software for Laue depth reconstructions, the GI-SAXS package for grazing-incidence small-angle scattering analysis, and the XMCD analysis tool for x-ray magnetic circular dichroism data are some examples that form the staple applications critical to the APS-U era facility. These programs will continue to be developed for improved performance and algorithms. A more complete list of software produced at the APS can be found at https://www1.aps.anl.gov/Science/Scientific-Software.

New efforts are currently underway to address the development of new algorithms and HPC software for multi-modal analysis, including fluorescence tomography, fluorescence ptychography, tomography diffraction, and Bragg CDI and ptychography, and for a new approach to Laue diffraction reconstructions. The MONA (Monitoring, Optimization, Navigation, Adaptation) project will prototype data streaming coupled with real-time data analysis and automated feedback.

One way the Laboratory supports the computational efforts at the APS is via LDRD funding. Beginning in FY11, the *Tao of Fusion* LDRD helped seed the TomoPy application; likewise, the FY13 *Next Generation Data Exploration: Intelligence in Data Analysis, Visualization and Mining* LDRD was aimed at multi-modal analysis. Other previously funded LDRDs include *Visualization and Mining, Modeling, Analysis, and Ultrafast Imaging (MAUI)*, *Multimodal Imaging of Materials for Energy Storage (MIMES)*, *Enabling Nanometer-scale X-ray Fluorescence Tomography*, and *Coherent Surface Scattering Imaging*.

Currently funded LDRDs of direct benefit to the APS in the computing space include FY17 *A Universal Data Analytics Platform for Science,* FY17 *COHED: Coherence for High-Energy Diffraction*, FY17 *Developing Advanced Coherent Surface Scattering Reconstruction Method Incorporating Dynamical Scattering Theory*, FY17 *Enabling Multidimensional X-ray Nano-Tomography*, FY18 *A.I C.D.I: Atomistically Informed Coherent Diffraction Imaging*, FY18 *Integrated Approach to Unravel Four Dimensional Spatiotemporal Correlation in Highly Transient Phenomena: Ultrafast X-ray Imaging and High-Performance Computing*, FY18 *Novel Capabilities for Ultra-fast and Ultra-low-dose 3D Scanning Hard X-ray Microscopy*, FY19 *Enabling Automatic Learning of Atmospheric Particles through APS-U*, FY19 *Finding Critical Processes of Deformation in Structural Materials with Artificial Intelligence*, FY19 *Learning and Differentiating: Using Artificial Intelligence to Image Beyond the X-ray Depth of Focus Limit*, FY19 *Machine Learning Enabled Advanced X-ray Spectroscopy in the APS-U Era*, FY20 *Machine Learning Methods for Spectral Data from X-ray Transition Edge Sensor Arrays*, FY20 *Coded Apertures for Depth Resolved Diffraction*, and FY20 *Intelligent Ptychography Scan via Diffraction-Based Machine Learning*.

In addition to LDRDs, the APS receives funding and personnel support from the ALCF Data Sciences Program (ADSP) for *Large-Scale Computing and Visualization on the Connectomes of the Brain*, *Developing High-Fidelity Dynamic and Ultrafast X-ray Imaging Tools for APS-Upgrade*, *X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future,* and *Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction*, and from the NERSC Exascale Science Applications Program for *Optimization of data-intensive tomography workflows at light sources*. Effort for CDI ptychography was initially funded by ASCR and now via the Intelligence Advanced Research Projects Activity (IARPA) and Northwestern University; early efforts for the MIDAS software were funded from APS industrial partners.

The APS has actively collaborated with other facilities and organizations, and members of the user community to develop data analysis algorithms and software. As examples, most ANL-funded LDRDs in these areas involve collaborators from ANL's Mathematics and Computer Science Division, Computational Science Division, or Data Science and Learning Division. Select user groups have contributed greatly to analysis algorithms and software. The Center for Advanced Mathematics for Energy Research Applications (CAMERA) aids in the development of software for modeling and GISAXS, the SHARP ptychographic reconstruction package, and most recently has been involved in the development of the XPCS-Eigen correlation application, in addition to running annual workshops for tomography, ptychography, and XPCS. The APS has been involved in the NOBUGS conference community, and maintains active participation in the series of hack-a-thons organized by the Experimental Facilities Computing (ExFaC) Working Group. The APS and the NSLS-II at BNL have developed a comprehensive computing collaboration plan so as to best utilize our scare resources. Since 2017, the directors of the five BES funded light sources chartered the Light Source Data & Computing Steering Committee (previously the Data Working Group) tasked with finding avenues for collaboration amongst the facilities. In 2019, the directors of the five BES funded light sources and the directors of the 4 ASCR computing and networking facilities charted the BES-ASCR Working Group tasked identifying how the ASCR facilities can help meet the needs of the BES facilities with regard to data and computing. These activities are directly aligned with facility priorities and the core requirements listed in Table 1.

The efforts described above will fulfill the needs described in the following APS-U beamline functional requirements and interface control documents: ASL Requirement 5.1, ATOMIC/3DMN Requirement 5.1, CHEX Requirement 5.1.5, CSSI Requirement 5.1.6, HEXM Requirement 5.1.7, ISN Requirement 5.1.10, PtychoProbe Requirement 5.1.8, POLAR Requirement 5.1.11, and XPCS Requirement 5.1.8.

Data Management

The APS Data Management System, the facility-wide software and hardware system for managing data, provides a data management and storage system sufficient to retain experimental data in accordance with sponsor requirements for all current and future APS beamlines.

The need for data management and distribution tools and data storage resources continues to grow. Currently the APS X-ray Science Division operated beamlines collect on the order of 5 PB

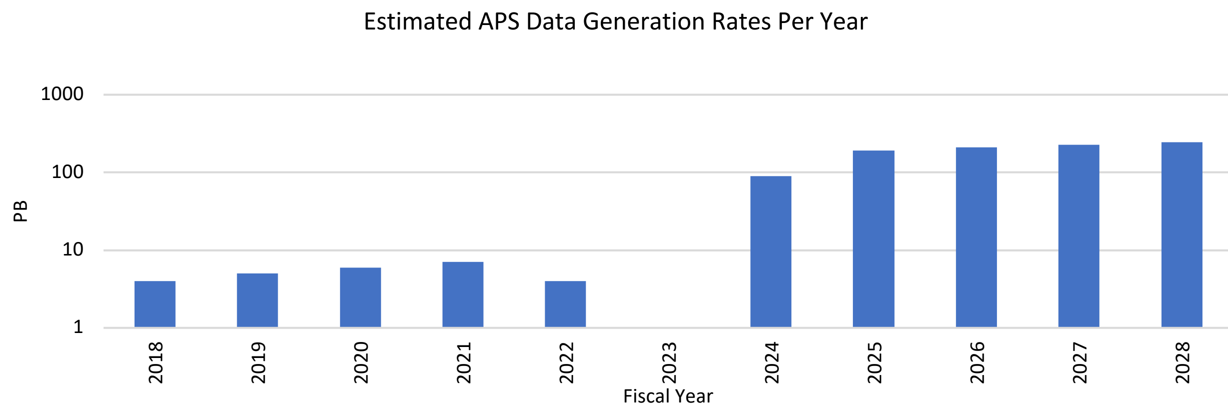

Figure 1 - Anticipated aggregate APS X-ray Science Division data generation per year. Data generation during 2022 and 2024 are estimated to be at 50% of the peak due to reduced user operations before and after the APS storage ring replacement. Data generation will stop during the 2023 long shutdown period.

of raw data per year. Over the next decade, it is estimated that the data storage needs of the APS are anticipated to increase by at least two orders of magnitude to over 100 PB of raw data per year (see Figure 1). Great strides have been made in this area over the past years. The APS will continue to deliver a multi-tiered data management and distribution system for all current and future APS beamlines.

During FY13 - FY15, the APS piloted facility-wide data management and distribution tools and resources with effort and funding from LDRDs. These activities provided R&D effort and seeded ongoing connections between the APS and the Argonne Leadership Computing Facility (ALCF), the Mathematics and Computer Science (MCS) division, and the Globus Services team at the University of Chicago.

Through APS operations effort, the APS Data Management System integrates with beamline data workflows, and large data storage systems. These tools help automate the transfer of data between acquisition devices, computing resources, and data storage systems. Ownership and access permissions are granted to the users signed-up to perform the experiment. A metadata catalog allows beamline staff to populate experiment conditions and information for access via a web portal. Users can download data at their home institutions using Globus Online (globus.org).

Medium-term data storage is available within the APS; longer-term storage systems are provided by the Argonne Leadership Computing Facility (see Figure 2). Currently, the APS provides approximately 3.6 PB of central disk storage (easily expandable to 15 PB) for medium-term data retention. The Argonne Leadership Computing Facility currently provides approximately 10 PB of tape storage (easily expandable to meet future APS needs) for longer-term data retention. In the near future, the ALCF will deploy a 100 PB community file system along with additional tape storage that will be available for APS use. Both the National Energy Research Scientific Computing (NERSC) Center and the Oak Ridge Leadership Computing Facility (OLCF) are deploying similar systems. These resources are currently funded, and will continue

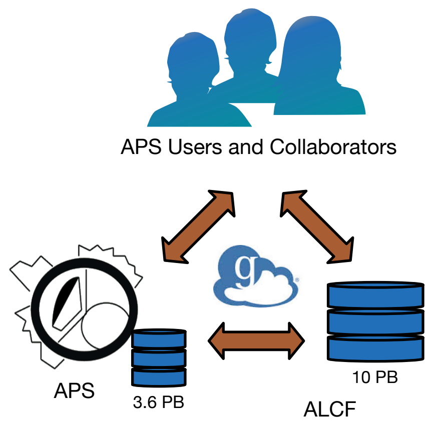

Figure 2 - Current storage available for APS beamlines. A 3.6 PB data storage system located at the APS serves medium-term needs. The Argonne Leadership Computing Facility (ALCF) provides 10 PB of tape storage for the APS. Capacity will be expanded as needed to meet sponsor requirements.

to be funded, by APS and ALCF operations budgets, respectively, for all current and future APS beamlines, including the APS-U feature beamlines and APS-U enhanced beamlines.

The current APS Data Management Policy can be found at https://www1.aps.anl.gov/users-information/help-reference/data-management-retrieval-practices.
This effort will fulfill the needs described in the following APS-U beamline functional requirements and interface control documents: ASL Requirement 5.1, ATOMIC/3DMN Requirement 5.1, CHEX Requirement 5.1.6, CSSI Requirement 5.1.7, HEXM Requirement 5.1.8, ISN Requirement 5.1.11, PtychoProbe Requirement 5.1.9, POLAR Requirement 5.1.12, and XPCS Requirement 5.1.9.

The APS is working with the Globus team to develop a computational data fabric for multiple techniques and instruments in the facility. This fabric will connect and automate all stages of the data lifecycle from acquisition to processing to publication. A science web portal will allow APS users to view and download their data and reprocess analysis jobs on ALCF and other large-scale computing resources. The Materials Data Facility (MDF) will serve as a DOI generating service for APS datasets. The APS and Globus team have prototyped a computational fabric for XPCS (see Figure 3) and serial crystallography, and plan to develop such workflows for ptychography, High-Energy Diffraction Microscopy (HEDM), and Bragg Coherent Diffraction Imaging (BCDI) over the next two years.

Computing and Network Infrastructure

As with data storage, the computing resources required by the APS are anticipated to grow by at least two orders of magnitude. Prior to the APS-U, most data processing can be performed with a few PFLOPS of computing resources. In the APS-U era, data processing at the APS will require on-demand access to over 50 PFLOPS. To satisfy these needs, the APS has adopted a graded approach to resource utilization. Small-scale resources, such as Intel Phi processors and GPUs, local to beamlines will be used when sufficient. For moderate computational needs, the APS maintains an on-site computing cluster, and ANL maintains computing resources as a part

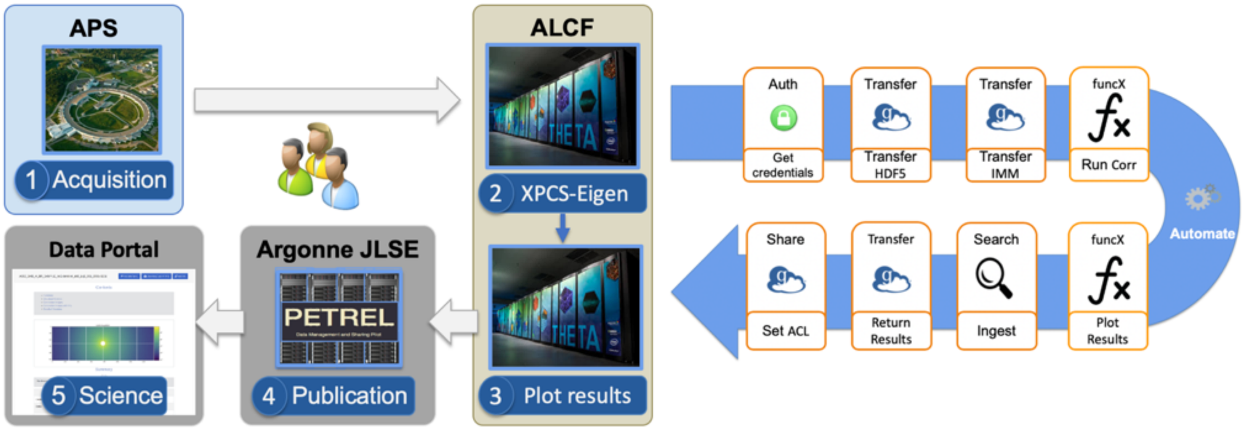

Figure 3 - Prototype automation used to perform on-demand analysis of XPCS data using Theta. Data are transferred to ALCF where compute nodes are provisioned to perform analysis, extract metadata, and plot results, which are then published to a Globus data portal for user consumption.

of the Laboratory Computing Resource Center (LCRC). For the most demanding computational problems, large-scale computing facilities must be used, including the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing (NERSC) Center, and the Oak Ridge Leadership Computing Facility (OLCF).

The ALCF has modified its queue structure and policies to better suit the size and frequency of APS processing jobs. An umbrella allocation at the ALCF managed by the APS will be utilized to provide compute access to APS users.

The APS has been involved in many activities aimed at using centralized and large-scale computing resources. Notable activities include:

- Efforts are currently underway to utilize Argonne Leadership Computing Facility (ALCF) super-computers, such as Theta, for routine beamline use. Software for ptychographic and tomographic reconstructions, elemental mapping, and XPCS correlation processing have been demonstrated to run on Theta.
- Researchers at the ALCF are currently developing workflow tools to easily enable transfer and on-demand job submission of beamline data at the APS. ALCF researchers demonstrated the use of the Balsam resource and queue scheduler to run XPCS processing jobs in an on-demand fashion.
- During SC'19, ALCF researchers demonstrated real-time reconstruction and visualization of tomography data. Data acquisition was simulated at the convention center and streamed to Theta at the ALCF for processing and denoising. Data was visualized back on the show floor.
- Efforts are currently underway at the APS to routinely utilize NERSC for high-energy diffraction microscopy and tomography reconstructions. Applications such as MIDAS (high-energy diffraction microscopy) and TomoPy have been run at scale with up to 30,000 cores, allowing data analysis during beam time.

Networking is another key consideration when planning both computing and storage resources. Sufficient bandwidth and latency constraints are required both within the APS, and within and outside of ANL. These demands will evolve with the architecture and utilization of computing and storage needs. Over the previous few years that APS completed a redundant 100 Gbps network connection to the central laboratory network and a redundant 100 Gbps connection to the ALCF computing center in building 240. Working with ESnet and ANL, the APS implemented a Science DMZ for particular outside user groups in order to reduce latency and better enable remote access and experiment control, and a Science DMZ between the APS and NERSC.

The APS is currently in the process of upgrading its fiber networking complex. Each LOM is being upgraded with an additional 4 x 48 strands of single mode backbone fiber to the APS main computer room. From each LOM fiber is connected to each sector. Single mode fiber can handle network speeds of 10G, 40G, 100G and beyond. This work is planned to be completed during 2020. Redundant core switches are planned for CAT network connectivity that enable multiple 10, 40, and 100 Gbps connectivity from CAT networks similar to the XSD beamline network.

In coming years, costs for use of internal vs. externally managed computing resources will be reviewed to ensure that the APS is most wisely using laboratory funding.

Group Strategies and Work Plans

XSD-BC Strategy:
https://www.aps.anl.gov/sites/www.aps.anl.gov/files/APS-Uploads/XSD/XSD-Strategic-Plans/BCDA_Strategy.pdf

XSD-CXS Strategy:
https://www.aps.anl.gov/sites/www.aps.anl.gov/files/APS-Uploads/XSD/XSD-Strategic-Plans/CXS_Strategy.pdf

XSD-SDM Strategy:
https://www.aps.anl.gov/sites/www.aps.anl.gov/files/APS-Uploads/XSD/XSD-Strategic-Plans/SDM-Strategy.pdf