

**Franklin & Marshall College
Campus Cyberinfrastructure (CI) Plan
January 2019**

A. Executive Summary

F&M has, year over year, matured its approach and available resources in support of campus researchers. This work is highlighted as a priority in the College's most recent strategic plan. Examples of this maturation include:

- creation of F&M's first-ever Faculty Center in 2014 with support from the Mellon Foundation
- development of a professional College Grants Office, which received a prestigious NIH Biomedical/Biobehavioral Research Administration Development (BRAD) Award to enhance its support for sponsored research and faculty development in health-related fields. Ultimately, all researchers will benefit from this work.
- receipt of a grant from the NSF in excess of \$400K to acquire a confocal microscope for F&M.

F&M's researchers expect access to world-class research instrumentation and infrastructure. They expect the same from the College's cyberinfrastructure. More than a utility, this infrastructure plays a crucial role in virtually all of the scientific research happening at F&M. To better support this, F&M applied for and received a CC* award in 2015 which allowed us to improve dramatically upon our inter- and intra-campus connectivity. Our Science DMZ, data transfer node, access to Globus and associated upgrades to a variety of campus fiber runs, including a move to a 10 Gb/s connection to PennREN and joining Internet2 have been transformative for our researchers.

As we look forward, we recognize that the work of recent years leaves us with one significant shortcoming in our ability to provide quality cyberinfrastructure to our researchers. Our current campus compute resources are significantly lagging and at the same time, demand is increasing particularly given developments such as those mentioned above. With this in mind, F&M's CI plan has four significant priorities.

Plan Goals

- Continue to leverage Science DMZ and DTN to provide Franklin & Marshall College with premium access to the world-wide research community.
- Enhance campus compute resources to support researchers on campus and improve their access to compute resources off-site.
- Provide student researchers, the researchers of tomorrow, access to world-class infrastructure to prepare them for further study and careers.
- Continue to assure F&M's entire cyberinfrastructure is secure, well maintained, contemporary and meeting the needs of the entire community.

B. Review of Relevant Existing Cyberinfrastructure and Practices

Participation in the InCommon Federation Identity Management

F&M has been a member of InCommon since 2012. All of our services are available via single-sign-on. F&M has been connected to Internet2 for many years and became a full member organization in 2015. F&M staff regularly participate in Internet2 events and programs and take advantage of the Net+ collection of services available to members. F&M is registered with InCommon and fully meets the InCommon Baseline Expectations for Trust in Federation. As well, F&M fully meets the requirements of the Research and Scholarship Entity category specification. F&M releases attributes to all research and scholarship service providers. F&M also makes extensive use of the SSL certificate service that is offered through InCommon.

Campus network backbone and external connections

F&M's network core was completely replaced in 2014. The core nodes are connected via 4 - 10 Gb/s links. The core nodes have support for both IPv6 and Software Defined Networks (SDN) in hardware.

F&M currently has 2 connections to the Internet. We have a 10 Gb/s connection to PennREN. The unused portion of the PennREN connection is available for connections to other PennREN-connected organizations via the KINBER Member Exchange. As a secondary back-up connection, F&M has a 1 Gb/s connection. Both connections come to our campus via fiber optic cable and offer the technical capability of being upgraded to 100 Gb/s or beyond.

F&M was the first organization to connect to PennREN after the initial build out. We worked with KINBER to design and install the fiber necessary to connect to the PennREN backbone as an end-node. There has been significant interest from other organizations in Lancaster to join PennREN and now numerous entities connect to PennREN via F&M. This initial investment has yielded other benefits such as allowing us to get connected to Internet2 many years ago and became a full member of Internet2 in 2015.

F&M has increased the bandwidth of uplinks from the buildings to the core, in most cases to 10 Gb/s with the implementation of new network edge switches in 2015-16. We also added redundant connections from the edge switches to the core, which provides a more reliable network connection and also doubles the available bandwidth.

F&M's current wireless network was installed in 2016 and fully supports the 802.11ac protocol. As well, F&M's primary SSID on campus is eduroam and we routinely see evidence that our community is taking advantage of eduroam access as they travel around the world.

F&M's Science DMZ, DTN and PerfSonar deployment

In 2015, F&M received an award from the NSF to build a Science DMZ and DTN. The build was completed in early 2018. As part of that work we also deployed a number of perfSONAR nodes. All of this infrastructure is currently in production and researchers are actively using the DTN and Globus to leverage our Science DMZ. This has been immediately useful to our researchers who are moving data to various xsede clusters. As was noted in our final report to the NSF, the impact of this infrastructure, even in the first few months, has been dramatic. As an example, in one of our more recent data transfers, we transferred half a Terabyte from Sunnyvale, California in just over 14 minutes (Effective Speed 5 Gb/s) which we can compare to before the Science DMZ when we transferred 500GB at 2 Mb/s which took over 78 hours from West Virginia. We have consistently attained speeds at or above 1 Gb/s to every transfer location.

Networking Best Practices and Information Security

F&M is a member of MANRS.org, the Mutually Agreed Norms for Routing Security and currently meets all four best practices MANRS recommends.

Under direction of the College's CISO, F&M has a number of information security initiatives both completed and underway. These include:

[REDACTED]

[REDACTED] Most relevant to our researchers, an entire restructuring of F&M's vlans, completed over the past few years, helped to make our Science DMZ all the more secure.

Data Center and Data Storage Infrastructure

F&M currently has two data centers on campus. The main data center contains F&M's

[REDACTED]

F&M plans to relocate the current secondary data center, which contains several additional HPC clusters, to a new space that is being built in summer 2019. This improved data center will provide greater space, advanced environmental controls and more robust network connections than the current secondary data center. This will also allow F&M to more uniformly spread administrative and HPC clusters between the two data centers for greater redundancy and reliability.

F&M makes available significant local data storage capacity for scientific research, when needed with high performance SAS/SSD drive arrays, at no cost to the researcher. We anticipate expanding the available resources as part of this proposal as it is an area of need.

High-Performance Computing at F&M

F&M is home to diverse fields of faculty and student research requiring access to a variety of HPC resources. Our researchers use cluster and grid technology, both on campus, and across the country. From our home base in Lancaster, Pa., to Puerto Rico, Singapore, and Australia, F&M faculty share data with collaborators across the globe.

One of our main goals for HPC at F&M is to further centralize computing resources to simplify management, make more efficient use of available resources and meet the growing need for these resources. We extensively leverage open source tools to manage the existing resources on campus such as ROCKS cluster management, HTCondor, Torque, and SGE. While F&M houses cluster computing, grid systems, and a variety of single high-performance computers that are driving innovative research in disciplines such as Astronomy, Biology, Computer Science and Mathematics, the resources are inadequate and demand continues to grow.

Achieving IPv6 Compatibility

F&M has applied for and received IPv6 address space and has end-to-end IPv6 capabilities. We have conducted tests of using IPv6 protocol but have not configured it permanently on our equipment. Our Internet routers, our DNS servers, and various other equipment are also IPv6 compliant. In addition, most of our servers also support IPv6.

Fiber Infrastructure

F&M has a robust fiber infrastructure that extends to every building on the main campus and almost every off-campus, College-owned building. The first fiber was installed in 1987, and the fiber plant has been regularly expanded since then, particularly in adding single-mode fiber runs to most buildings. Almost every building has multiple fiber cables entering at diverse locations to increase redundancy and resiliency.

Performance Monitoring

F&M makes extensive use of [REDACTED] to monitor the status of every piece of network electronics and the connections between them, other devices, our services, and our connections to the Internet. [REDACTED] is configured to maintain graphs of link utilization over time so that we can identify issues and identify trends in usage. This is particularly important for our core connections and our Internet connections. Our server provides notifications via email and text messages in the event of a service disruption, and it provides considerable flexibility in determining when a notification should be sent—and to whom.

F&M uses Splunk to collect and analyze log files from all network electronics and most of our servers. Splunk allows us to correlate data from many sources and to identify the cause of issues promptly. This works well in tandem with our Intermapper server.

F&M has three perfSONAR nodes. These include a 10 Gb/s tester funded by our 2015 NSF CC* award, a 1 Gb/s tester received from KINBER as part of our research and education network's effort to deploy more perfSonar nodes in Pennsylvania and a node F&M built that runs in our central campus

infrastructure environment. Our perfSONAR nodes are used to benchmark our Science DMZ as well as our entire network and identify bottlenecks both on campus and as we move data inter-institutionally.

C. Sustainability

F&M's CI strategy attempts to leverage federal research funding opportunities to enhance infrastructure and services available to our community. In addition, per College policy, researchers who receive grants or start up funds that require computing resources are required to contribute resources to a centralized campus compute cluster and not attempt to fund and/or manage their own. This policy was established in 2015. However, we recognize that these are highly competitive programs and the institution has a responsibility to maintain these assets and use institutional resources to do so as necessary, particularly as they are often also utilized in teaching our students. F&M consistently allocates ██████ annually for cyberinfrastructure renewal and at times more for significantly more expensive initiatives. These resources would be used in the future to renew the requested resources as needed.

D. Future Initiatives Planned In Support of Plan Goals

- Move, in March 2019, to a new single sign-on service. We continue to closely monitor the work being done by InCommon to increase security and provide additional features. In addition, we have had personnel participate on some of the planning teams associated with Internet2's TIER program, which we imagine making use of in the future.
- Complete evaluation of edge switch options and begin a phased migration to new hardware beginning summer 2019.
- Build a storage array dedicated to F&M's faculty researcher needs and optimized with those needs in mind.
- Build a centralized robust campus compute resource to be shared by all researchers in need of compute cycles and take advantage of collaborations such as Open Science Grid.
- Construct a larger more appropriate second data center to create greater resiliency in support of the College's research and administrative needs.
- Re-evaluate current cloud vendor DR strategy and consider alternatives.

Conceptual Functional Network Diagram - 2019

