

## **Campus Cyberinfrastructure (CI) Plan**

### **Background**

The NCShareDMZ project is developed and designed to provide new and improved Cyberinfrastructure resources for lesser funded Institutions of Higher Education (IHE) in North Carolina. A unique partnership, consisting of Duke University, MCNC, NC Central University (NCCU) and Davidson College (Davidson), was formed with a planning collaboration funded through the NSF funded “CC\* Planning: NC Regional Science DMZ” (NSF OAC-2018401 PI: Tracy Futhey) to understand the science driven cyberinfrastructure requirements at under-resourced institutions in North Carolina. The NCShareDMZ infrastructure, a virtualized regional Science DMZ, is the result of this planning process with the expectation it will lead to a regional service model that can be cost effectively extended to other IHE’s in North Carolina while also facilitating access to network-based cyberinfrastructure resources including scientific computing, remote instrumentation, and data storage to name a few.

Each member of the NCShareDMZ project brings a unique perspective and skills to the project:

- Duke University is a private R1 university that has developed, deployed and benefited from significant cybersecurity programs and resources on its campus and brings substantial technical expertise already in place at Duke’s Office of Information Technology, whose staff has architected and maintains an innovative Science DMZ that was created with support from NSF (ACI-1246042, CNS-1243315, ACI-1440588, OAC-1925550);
- MCNC is the state Research and Education Network provider supporting regional and wide area network connectivity; it leverages MCNC’s excellent reputation, expansive (4000-fiber-mile) network footprint, extensive and growing set of network-based services, and near-universal IHE customer base;
- NC Central University (NCCU) is a publicly funded university and one of 11 HBCUs in the state and is an important link to the greater North Carolina System institutions;
- Davidson College is a small private college with a strong desire to increase scientific computational support centrally within a limited budget, and is part of the NC Independent Colleges & Universities Consortium.

Through the Planning Grant process the team also worked with each campus to develop a Campus Cyberinfrastructure plan identifying that would be addressed through the NCShareDMZ infrastructure and beyond.

### **NCCU**

- The on-premise storage arrays are not adequately sized to serve as both primary and secondary storage of research data. At times, researchers have relied upon high-capacity, consumer grade hard drives for storage of their data.
- There is a lack of access to dedicated HPC resources to support research.
- There is not currently a resource to provide both access and expertise to open-source tools such as R Studio.

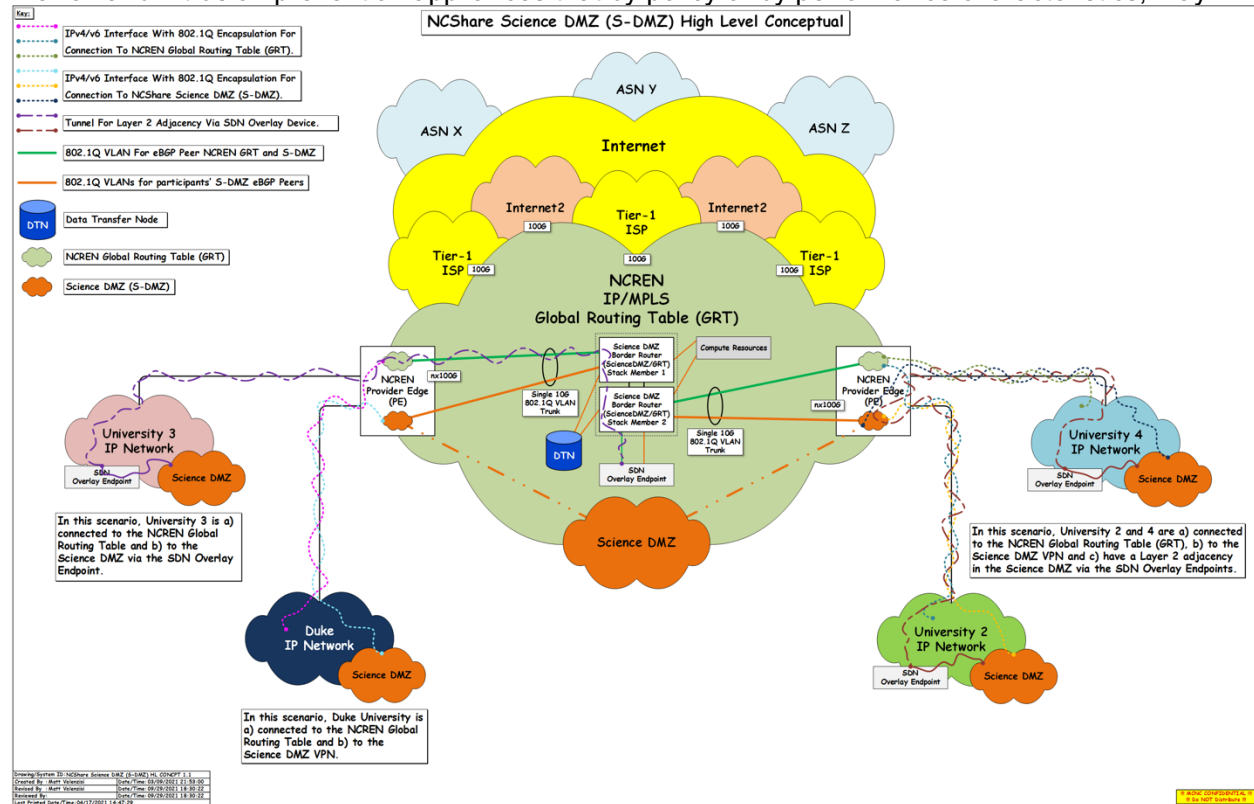
### **Davidson**

- The initial deployment of Jupyter Notebooks has received very positive feedback among faculty. Davidson seeks to continue adoption of that deployment as well as provide RStudio capabilities via web browser within the Jupyter framework for the many students and faculty using R for teaching, learning and research.
- Research computing uses currently share space with enterprise computing equipment in the campus’ single data center. Work and planning is needed to create a secondary location to host research computing and departmental servers. The establishment of the NCShareDMZ would allow for some of the hosting of these services through the MCNC co-location services.

- Review middleware platforms (such as Cirrus Identity or possible future inter-institutional collaborations) as an option for metadata compatibility within our going-forward Azure AD SSO platform. We will also review the requirements for the Research & Scholarship Adopters and SIRTFI communities in the InCommon federation.

### NCShareDMZ Design Details

Figure 1 below illustrates the conceptualized NCShareDMZ architecture developed through the CC\* Planning grant process. As indicated in the diagram, existing connectivity to NCREN will be utilized by connecting campuses, such as NCCU and Davidson, by virtualizing existing last-mile circuits between university networks and NCREN. A parallel "friction-free" path (PFFP) to the NCShareDMZ will be established, thereby creating an optimized path between a participating IHE and the NCShareDMZ that is engineered to bypass traditional campus network edge infrastructure, such as firewall and intrusion prevention appliances that by policy or by performance characteristics, may



impede the movement of S-DMZ flows. They can be provisioned with minimal complexity and without the participating IHEs incurring additional circuit fees. The virtualization of the circuits is achieved via the use of 802.1Q VLAN tags which facilitate the establishment of Layer 2 PFFPs paths from the participating IHEs to the NCShareDMZ. Over these Layer 2 PFFPs IPv4/IPv6 peerings can be established between the participating IHEs and the NCShareDMZ providing low latency access to remote compute infrastructure and other cyberinfrastructure connected resources within the NCShareDMZ at MCNC or nationally via NCREN/Internet2.

The NCShareDMZ fully **supports Mutually Agreed Norms for Routing Security (MANRS) and IPv6**. MCNC and Duke will discuss the benefits of MANRS and IPv6 with participating institutions as part of connecting the to the NCShareDMZ.

## **Regional Network**

The NCShareDMZ will utilize MCNCs regional network. MCNC delivers broadband services, through the North Carolina Research and Education Network (NCREN), to all public K-20 institutions in NC and 70% of the independent and private colleges and universities. In addition to servicing anchor institutions throughout North Carolina such as the Renaissance Computing Institute (RENCI), National Oceanic and Atmospheric Administration (NOAA), and Pisgah Astronomical Research Institute (PARI), MCNC works to support digital inclusion efforts across the state and prioritize North Carolina communities that are considered underserved in terms of broadband infrastructure, economic stress, or other conditions. Duke, NCCU, Davidson College all receive broadband services through MCNC. NCREN aggregates connections throughout the state, to various Internet providers, and to Internet2 for all participants. One of these sites includes the North Carolina Research Campus (NCRC) located in Kannapolis, NC which houses collaborators for Duke, NCCU, and Davidson.

NCCU has diverse 10Gbps links which connect to each of the campus's network cores. Davidson College has redundant 10GBps links with a 3Gb/s subscription for services. See the Facilities, Equipment, and Other Resources document for current utilization.

## **Federated Identity**

Duke, NCCU, and Davidson are ongoing members of InCommon.

**Duke** was early participant in both Internet2 and the InCommon Federation and has a multi-decade history of national collaboration. Duke users regularly rely upon the InCommon federation for access to more than two dozen external services; likewise, Duke-sponsored services are accessible through the federation to users throughout the InCommon community, and Duke is registered with InCommon as supporting the Research and Scholarship (R&S) Entity Category to streamline integration with research applications. Within the wider Internet2 community, Duke has a long history of involvement in both nationwide networking initiatives and Internet2- and NMI-sponsored middleware initiatives, including the MACE Paccman and MACE Grouper initiatives, the latter of which actually employs a senior Duke IT analyst as one of its primary developers.

Updates to **NCCU's** IdP have been executed to meet the InCommon Baseline Expectations for Trust in Federation. The process to register with Incommon as supporting the Research and Scholarship (R&S) Entity Category is presently being reviewed.

**Davidson** utilizes Microsoft Active Directory Federated Services (ADFS) as our identity provider when InCommon federated metadata exchange is required and Microsoft Azure AD is the primary SAML-compatible identity provider.

## **NCShareDMZ Participating Campuses**

### **NC Central University**

#### **Campus Data Network**

NCCU's campus network architecture is built upon the three-tiered hierarchical internetworking model having distinct layers for the core, distribution and access. This model was first implemented in 2018 when the University's leadership authorized a comprehensive network refresh. This refresh included the replacement of existing network equipment as well as the

installation of a new single-mode fiber optic backbone to support the hierarchical design. The core devices are split across two separate buildings and have a 40Gbps connection between them. Virtual Routing and Forwarding-Lite (VRF-Lite) is used to create logical security zones across the network. Next-Generation Firewall devices with 10Gbps interfaces are able to manage traffic flows across the VRFs without becoming a bottleneck.

Each campus building has a 10Gbps connection to its upstream distribution point. At the access layer, there are 1Gbps links between the floors as well as to each endpoint. To support wireless connectivity, a controller-based deployment of 802.11ac exists. Both the 2.4GHz and 5GHz spectrums are made available, accepting minimum data rates of 11Mbps to support a variety of legacy and IoT devices. Although not currently deployed, all existing networking and server equipment are capable of supporting IPv6. The implementation of IPv6 will be reviewed on an annual basis as part of an overall network efficiency initiative. Both 3<sup>rd</sup> party and network vendor tools are deployed to monitor the performance of the data network and identify any saturation or latency. The current design of the network allows for upgrading the 10Gbps links to 40Gbps, through funding as part of the NCShareDMZ grant, to the research centers.

### **Compute and Storage**

Centralized ITS provides academic, research, and administrative areas with hosting services for research systems within two modern data centers and a variety of storage services through enterprise SAN and HCI storage arrays as well as SaaS solutions (Microsoft and Google). There are currently no dedicated HPC services and researcher specific compute and storage solutions are generally still deployed within local labs and departments. Researchers, at times, struggle to store large data sets as well as transfer data with collaborators.

### **Security and Privacy**

NCCU's information security program includes policies, regulations and rules that document the expectations for maintaining the confidentiality, integrity and availability of university data. A defense-in-depth approach is applied across the institution. Security awareness training is provided on an ongoing basis. NCCU supports Mutually Agreed Norms for Routing Security (MANRS) and is reviewing the process of joining the initiative.

## **Davidson College**

### **Networking**

Davidson's core network, refreshed in 2020, is based on a Cisco architecture and includes 40 gigabit transit capabilities, including 40 gigabit capacity to each of multiple network segments. Most academic/research buildings are connected to the distribution and core layer via 10 gigabit links; most other campus buildings have 1 gigabit links. The campus is served by a robust fiber plant first developed in the mid-1990s. A Palo Alto Networks next-generation firewall (refreshed 2021) provides both north-south protection as well as east-west inspection of data across different segments. Employees, students, guests, IoT/industrial devices, campuswide servers, and enterprise data center servers are all segmented on the campus network. The network also contains a Frontier zone intended for faculty web sites and related projects. Both 3<sup>rd</sup> party and network vendor tools are deployed to monitor the performance of the data network and identify any saturation or latency.

### **Computing and Storage**

Over the past two years, Davidson has been able to repurpose some enterprise servers as well as purchase servers and a 50TB storage array to dedicate to scientific computing for research and education. Encouraged by the success of these programs, Davidson would like to continue

to entice researchers to build into central IT offerings through enhanced hosting services and a secondary location.

### **Security and Privacy**

Historically, the types of research data in use at Davidson did not carry with it significant requirements for secure storage; however, T&I is working on a range of services related to research data that include mechanisms for secure storage as well as partnerships across campus to support research data management.

Security reviews are performed in conjunction with the College's access to electronic communications policy, as approved by the faculty executive council in 2017.

### **Sustainability**

Financial sustainability: Central to sustainability is MCNC's 40+ years of operation as a trusted connector for NC. Their service and financial model establishes even smaller IHEs with a 10Gb link, where they pay for a tier of service corresponding to their typical traffic level (e.g., 1.5Gb or 5Gb). MCNC's commitment that this headroom capacity can be utilized at no direct cost in support of Science DMZ flows establishes our Virtual, Shared Science DMZ model as one that does not carry significant direct ongoing hardware maintenance and replacement cost, separate from that which the institution would already incur for its commodity Internet services. Further, MCNC's commitment to support the one-time and annual recurring costs for the network-related licensing, hardware and service costs (e.g., MPLS VPN licenses, port connections, IP address space, access to 10Gb lambda service) minimizes the ongoing financial burden on participants. This represents several thousand dollars per year as opportunity costs savings.

Technical sustainability: MCNC and Duke will drive the technological evolution of the network, which both are committed to continue to do as part of our service missions to the community. NCCU and Davidson will commit technical staff to partner with Duke and MCNC to fully utilize these resources. NCCU has a goal of becoming an R2 university; technologies such as the NCShareDMZ offer a critical component of providing campus researchers (including those in and outside STEM) with critical infrastructure to meet this goal. For its part, Davidson continues to emphasize undergraduate research experiences – the college ranked tenth in a recent national rating of colleges and universities based on Undergraduate Research/Creative Projects – which will be aided by the resources NCShareDMZ provides.

Any participating IHEs connected in year 2 of the project will be required to review the best practices established as part of the NCShareDMZ and make the same commitments as NCCU and Davidson College.