UFNet2: A vision for the next generation UF network

Executive Summary

High-speed networks have enabled real-time collaboration, research and teaching at a pace unimaginable only a decade ago. At the same time, they create a significant risk of sensitive data exposure and disruption of University operations. The need to provide both an open environment for research and teaching, as well as a closed environment for data protection and operational effectiveness are seemingly at odds with each other. A new way of organizing our people and resources at the network level is required to accomplish these goals. One that is both fined grained and ubiquitous. An architecture, which is configurable based on the users or systems role, not where the user or system is on campus, and dynamic, providing a consistent user experience enterprise-wide.

Virtualizing the Network Backbone

The first step to a new architecture is to virtualize the network backbone. Using similar paradigms to server virtualization, the network backbone may also be virtualized to increase flexibility and efficiency. The first step is to break up the existing network into a low number of distinct virtual network environments (vNEs). The guidelines for environment and environment classification would be:

- Have the fewest number of environments possible.
- Create the environments such that people are not moved around.
- Be able to oversee data movement between environments in some controlled fashion.
- Provide open access to academics/researchers. Move them out from behind network impediments.
- Protect sensitive data and the "business of the university" IT infrastructure.

Given these guidelines, the most logical approach seems to be to define the following environments:

- Academic: Faculty/staff who participate in teaching and general research and the systems that provide for them.
- Administrative: "Business of the university" staff and resources. This environment contains PII data related to the operation of the institution.
- Health: PHI/HIPAA type data and the individuals/systems for which their primary purpose is working on this data.
- ScienceDMZ: Researchers/Instruments which require high speed external connectivity or direct, unfettered access to Campus Research Network resources.
- External: Visitors to the university, or those who are on campus, but not employees, faculty, staff, students, or contractors. This nVE is outside of the standard UF security parameter.

Each vNE has its own distinct set of operational criteria. Exact guidelines for each environment are beyond the scope of this document. Future environments may be defined; however they should meet the vNE guidelines presented above. Environments create a broad and generalized approach

to network architecture that is systems and user based as opposed to location based. It is acceptable to have areas that are slightly more restrictive in a given environment; however, the mission of and data within the environment should be respected when devising these restrictions. In addition, data migration between environments should be controlled via official University policy and procedures and may include systems designed to detect such movement. To provide more fine-grained control, tiers may be established within each environment. These tiers should be classified into user tiers and services tiers. Services tiers provide a generalized service to users, and are roughly equivalent campus wide. Examples may include voip phones, building automation, audio/visual systems, vending, etc. User tiers are the networks on which users reside. Security policy for access to these networks and restricted data is done on a case-by-case basis. General recommendations will be made and templates will be established for common access control scenarios. The important point to this design is that decisions on open vs. closed access within tiers is made much closer to the user (i.e. subnet manager or equivalent), so changes to these controls may be made in a rapid fashion without a large amount of administrative overhead and **without risking sensitive systems or data**.

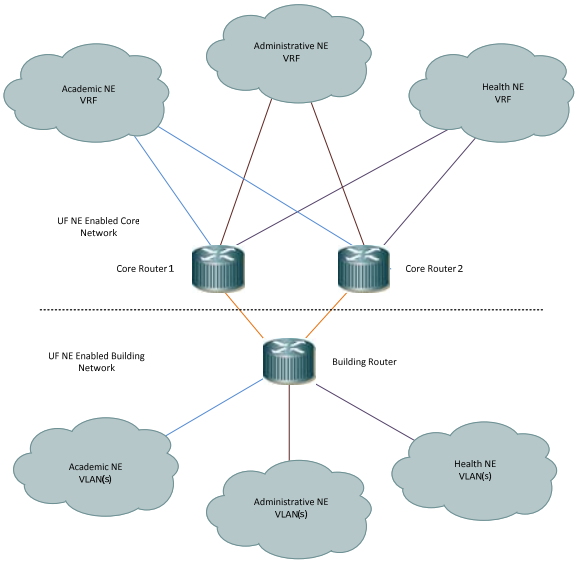

Figure 1: Generalized Core and Building Network Architecture

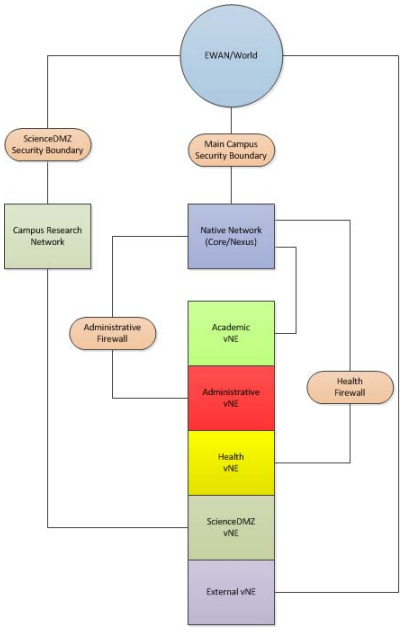

Figure 2: Core VRF Interaction

Core Network Architecture

To provide the network environments in a fine grained and ubiquitous manner, each environment would be implemented as a Virtual Routing and Forwarding instance (VRF) on top of a Multi-Protocol Label Switching (MPLS) core network. Each network environment would be put into a VRF defined on the core network. These VRFs can be seen as distinct "virtual core networks," or slices, which isolate traffic to a given environment until it reaches a well-controlled point in the network where traffic between environments or external networks may flow. This common core network would be responsible for providing network services to all vNE enabled buildings everywhere on campus.

Points where traffic flows between environments provides for the broad security posture in each vNE. It is expected that the current AHC boundary would form the basis for the Health vNE. The current core network would form the basis of the new Academic NE, and the Administrative NE would be developed from scratch. The exact final layout

of the NEs would be determined at a later date and may be modified to best fit traffic patterns and NE use campus-wide. It's also worthy of note that the new UF data center networks on the main campus and eastside campus can provide the same environments, as can the FLR state wide network.

Building Network Architecture

The current UF and AHC building network architecture would be expanded to handle NEs. Currently the standard building design typically calls for two distinct connections to the UF core network, exchanging routes over a single Virtual LAN (VLAN) per connection using the OSPF routing protocol. The extension would be to use a technique called VRF Lite. Rather than a pair of VLANs going to the core network, you would have a pair of VLANs per NE. The building layer 3 devices would "peer" with each NE, and using VLANS, keep these NEs distinct within the building. In that way, a device that is on the Health NE can sit right next to a device that is on the Academic NE. Moving between NEs involves a VLAN change on a port within the building. It's important to note that many VLANs in a given building may be associated with an NE, and traffic between VLANs in the same NE would follow the standard traffic rules we use today. Because we use VLANs in the building and VRFs in the core, traffic in different NEs **may not touch** until it reaches a NE crossing point in the network (at the core level). More discussion surrounding the location and management of cross NE traffic will need to take place. It is envisioned that a "custodial group" will be named for each NE. That group will be responsible for implementing and maintaining the security posture of the NE.

Campus Research Network and vNEs

In this new design, security posture would be determined not by what physical network you are in, but rather what NE you are in. The CRN is therefore initially part of the ScienceDMZ vNE with a couple of important distinctions:

- It is typically 4x10 times faster than the current core network, thus it is geared towards **very** high-speed sources and sinks of data. These may include large scale storage and compute (such as in the HPC center), gene sequencers, etc.
- It has slightly different "production ethics" to allow for rapid changes in support of research and research which may involve the network itself (i.e. SDN).

Thus, the CRN itself would no longer be an environment in and of itself. In its current capacity, it would provide the ScienceDMZ vNE at a much higher speed. It could best be thought of as an area within the ScienceDMZ vNE. It could also provide, for instance, the Health vNE without putting PHI outside of the Health NE security boundary. It is important for the CRN to continue to maintain different production rules and "flexibility" to service the needs for which it was designed.

For data requirements at 10 Gbps and below, the UF core network would be used to satisfy those demands. CRN connectivity would only be necessary for needs at the 40 G and 100 G levels.