


# Transforming a Campus Network for Data Intensive Science

Eric Boyd, Amy Liebowitz  
University of Michigan






# University of Michigan Network Services




ITS network services connect campus to WiFi, university networks, and the Internet. We also design and maintain the university's communication backbone, connecting the university to external education and research networks nationally and around the globe.



# U-M Network Goals



- Support academic and administrative network needs
- Enable 100G connectivity for every building on campus
- Enable WiFi 6E for every client
- Enable data intensive science
- Re-architect network design for advanced services and operational efficiency
- Enabled security without compromising aforementioned goals




# U-M Network Design Goals



Adopt design paradigms that support high bandwidth flows by maximizing throughput / minimizing loss and latency

- Prioritize network connections of key computation and storage resources
- Data intensive science collaborations
  - within U-M (e.g. HITS  $\longleftrightarrow$  ARC-TS)
  - with external cloud providers (e.g. AWS, Google Cloud, MS Azure)
  - reliant on external data sources and instruments (e.g. AGLT2, Cryo-EM)



# U-M Network Design Principles

Adopt designs that ensure higher levels of performance and stability of the network

- Core and Backbone Interconnect Network (BIN) - full mesh topology between core routers
- Diversely routing cabling between sites - each campus has a pair of fully redundant core routers
- N+1 power supplies - UPS and generator power at all sites

# Engineering New Solutions



1. Core network upgrade - Core deployed 9/22
2. Campus fiber upgrade - Largely complete
3. DL replacement - In progress (3 years)
4. MACC network refresh - In progress (3 years)
5. AL replacement - Unit driven

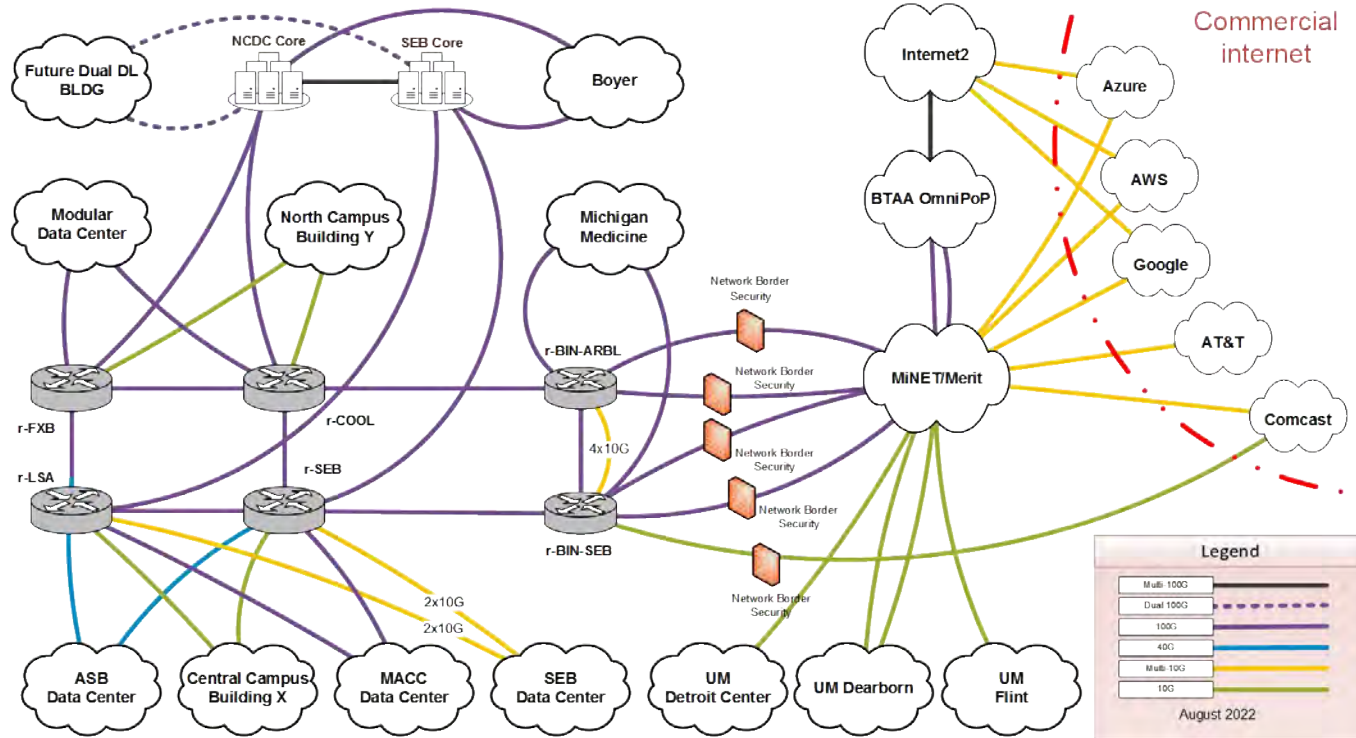
## Why Replace Core Network?

- Cannot provide many 100 Gbps connections at a reasonable cost
- Routers are nearing 10 years old
- Current network design is spread across too many (4) sites

## Why Replace DLs

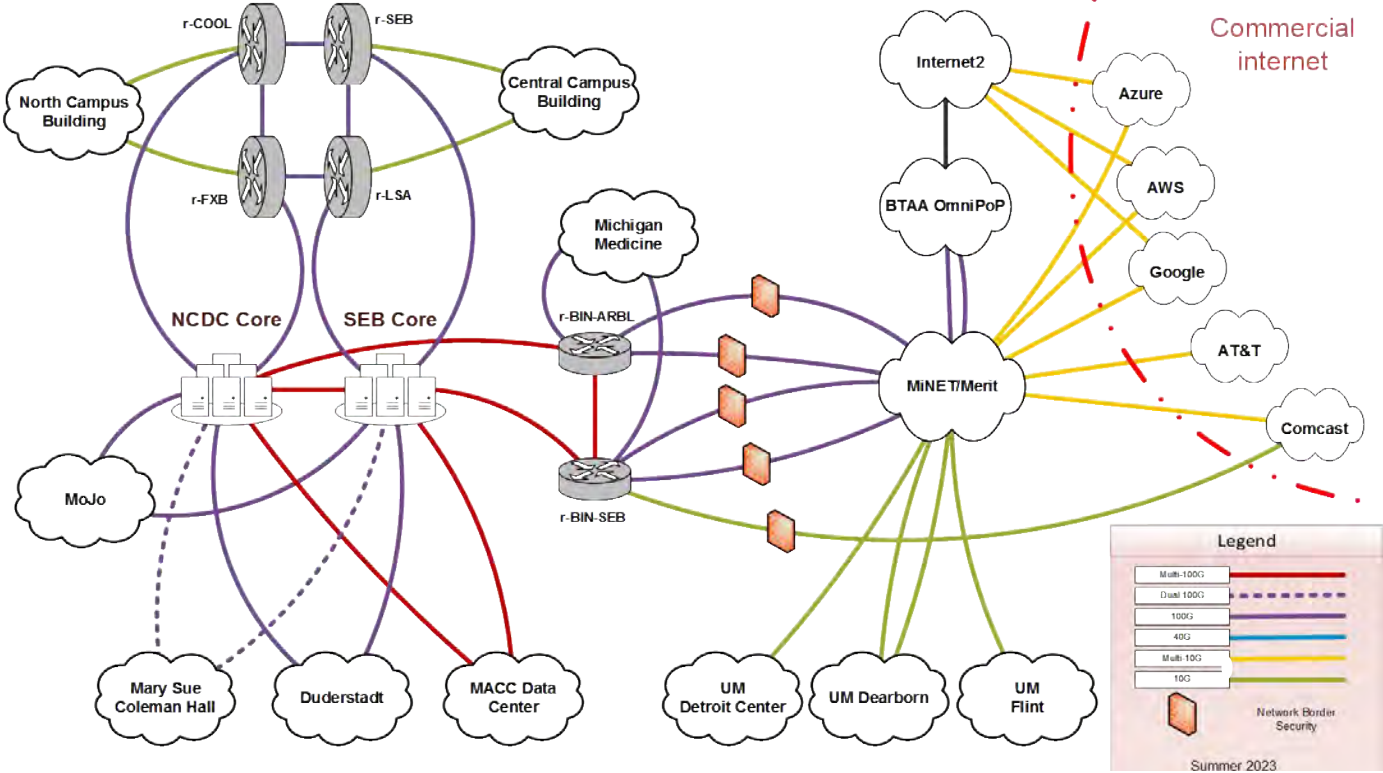
- Most current DLs are well beyond end of life and cannot support higher speed connections

# Core Network Transformation: Parallel Cores - Current

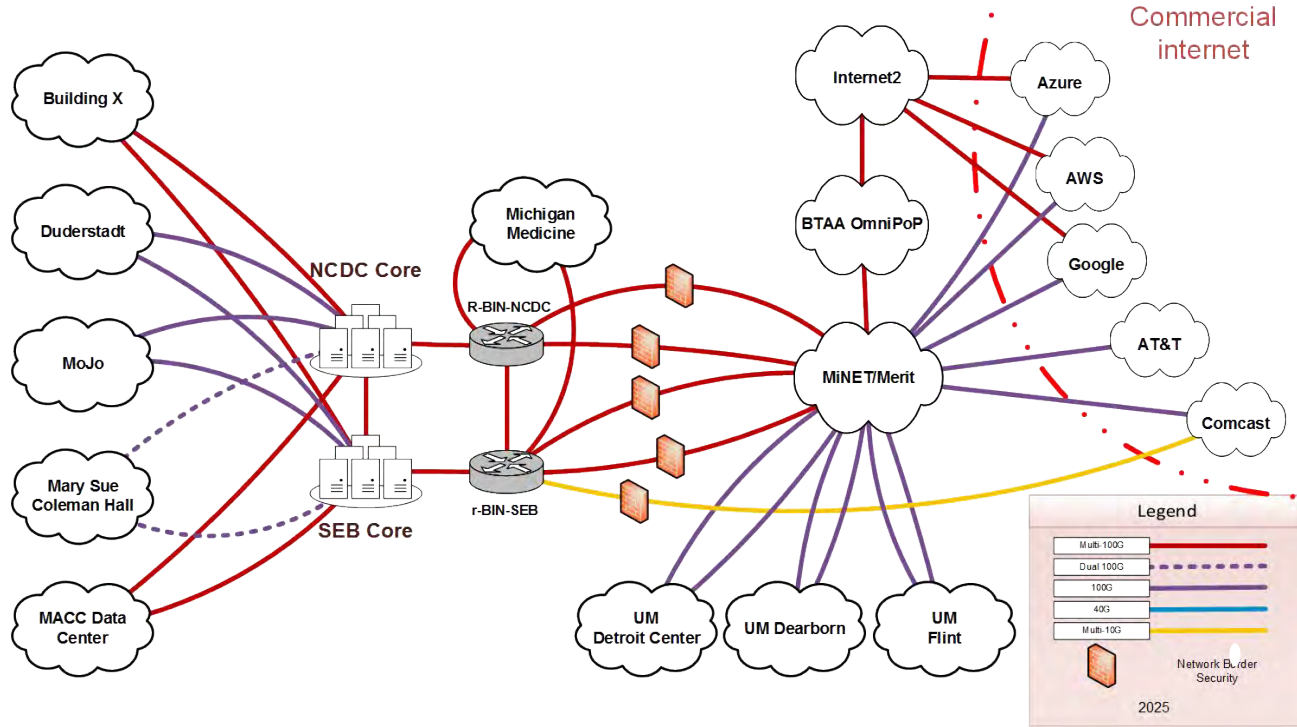




# Core Network Transformation: Flip the Cores - Stage 2



# Core Network Transformation: DL Project Complete - Final



# Backbone Network Requirements

- Capacity for 100G to all campus buildings
- Deep buffer switches for buildings supporting HPC (no dedicated Science DMZ)
- Macro segmentation (L3VPN) for steering traffic to centralized security devices
- Limited layer 2 stretch between buildings
  - For legacy implementations where no other solution is available

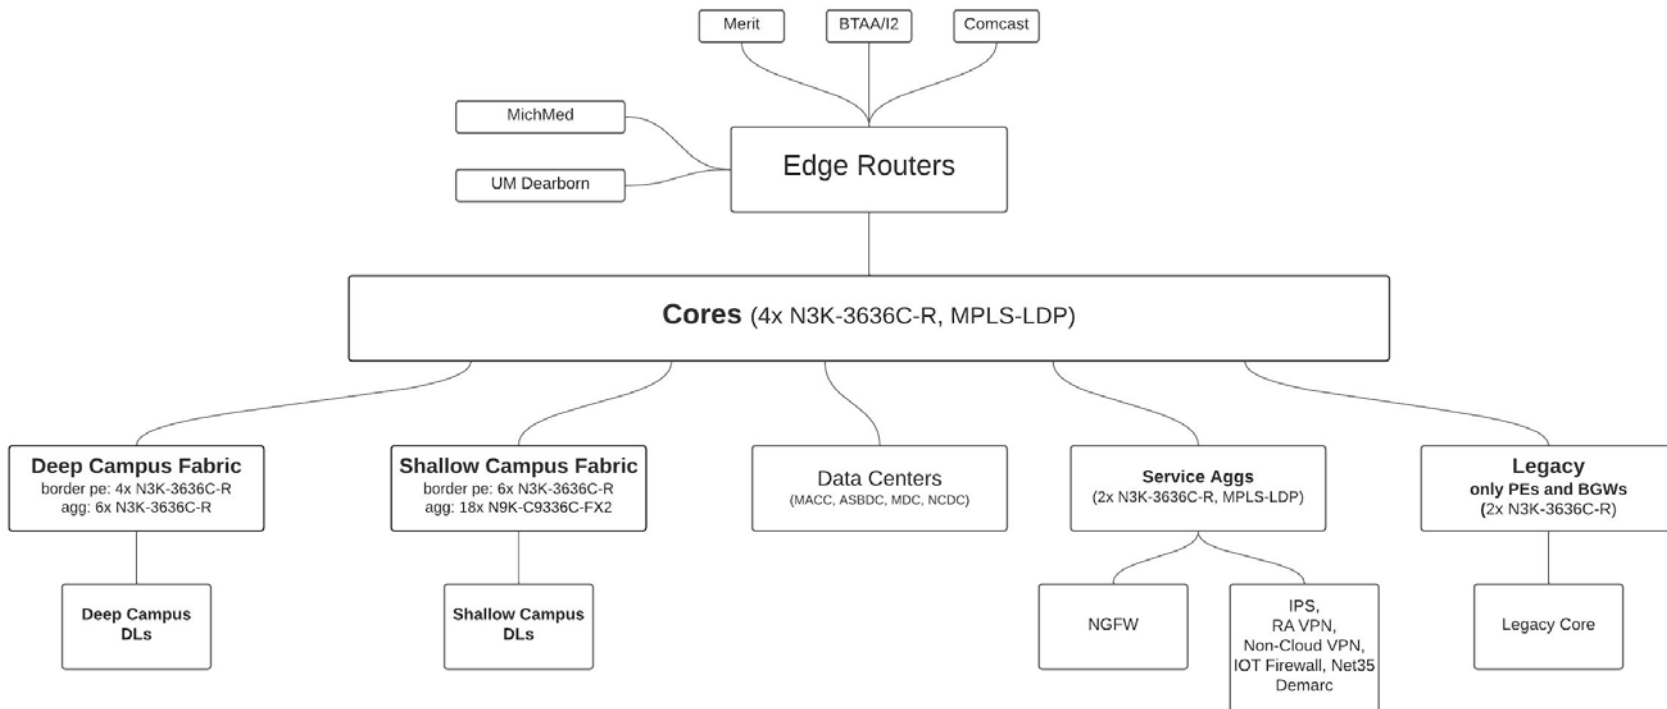


# Backbone Design Overview

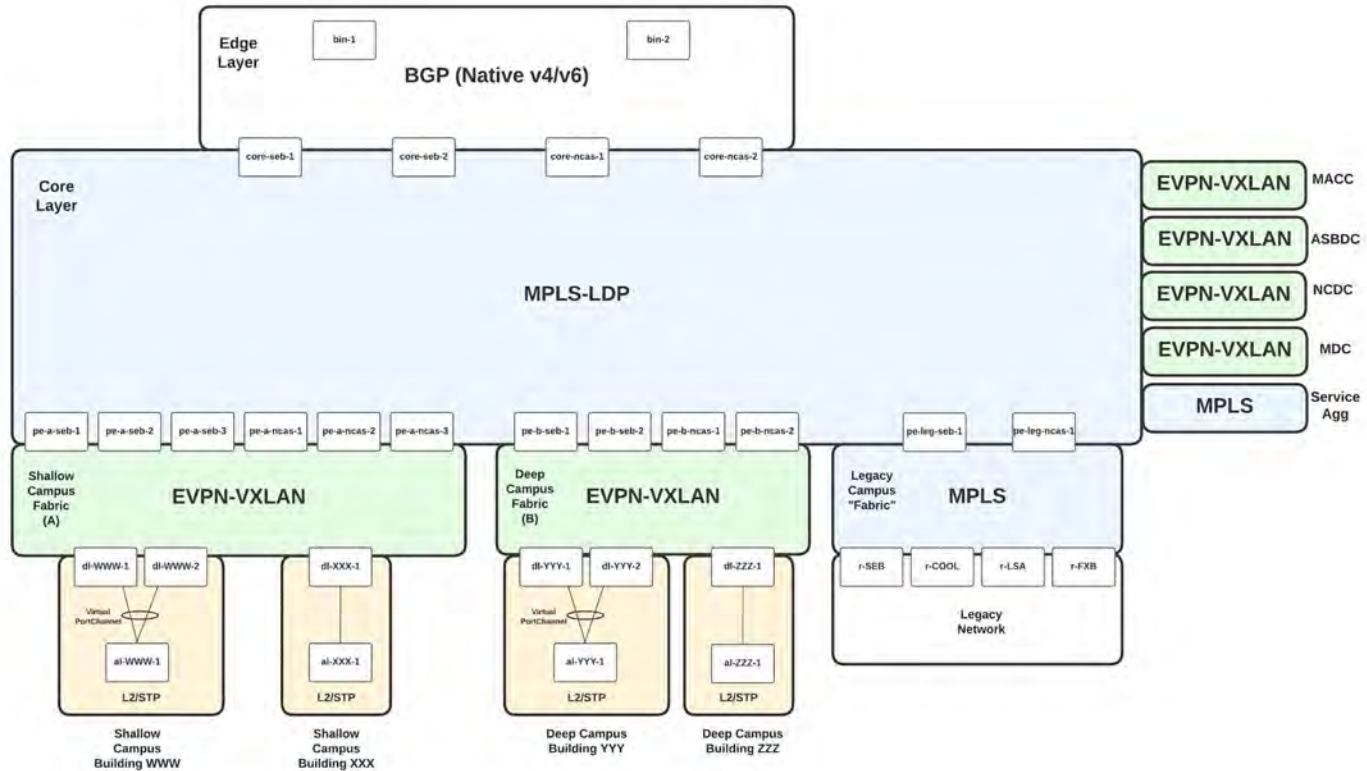
- Adapted version of Cisco's VXLAN EVPN Multi-Site Design
- Heavily leverages 100G fixed config Nexus platforms: N3K-3636C-R and N9K-9336C-FX2
- Campus broken up into EVPN/VXLAN fabrics
  - Deep buffer fabric for HPC buildings
  - Shallow buffer fabric for other buildings
  - Each DC is it's own fabric
- Fabrics are connected via an MPLS L3VPN Core
- Core layer connects fabrics to edge routers, service aggregation, and legacy network.



# Backbone High Level Overview



# Backbone High Level Overview

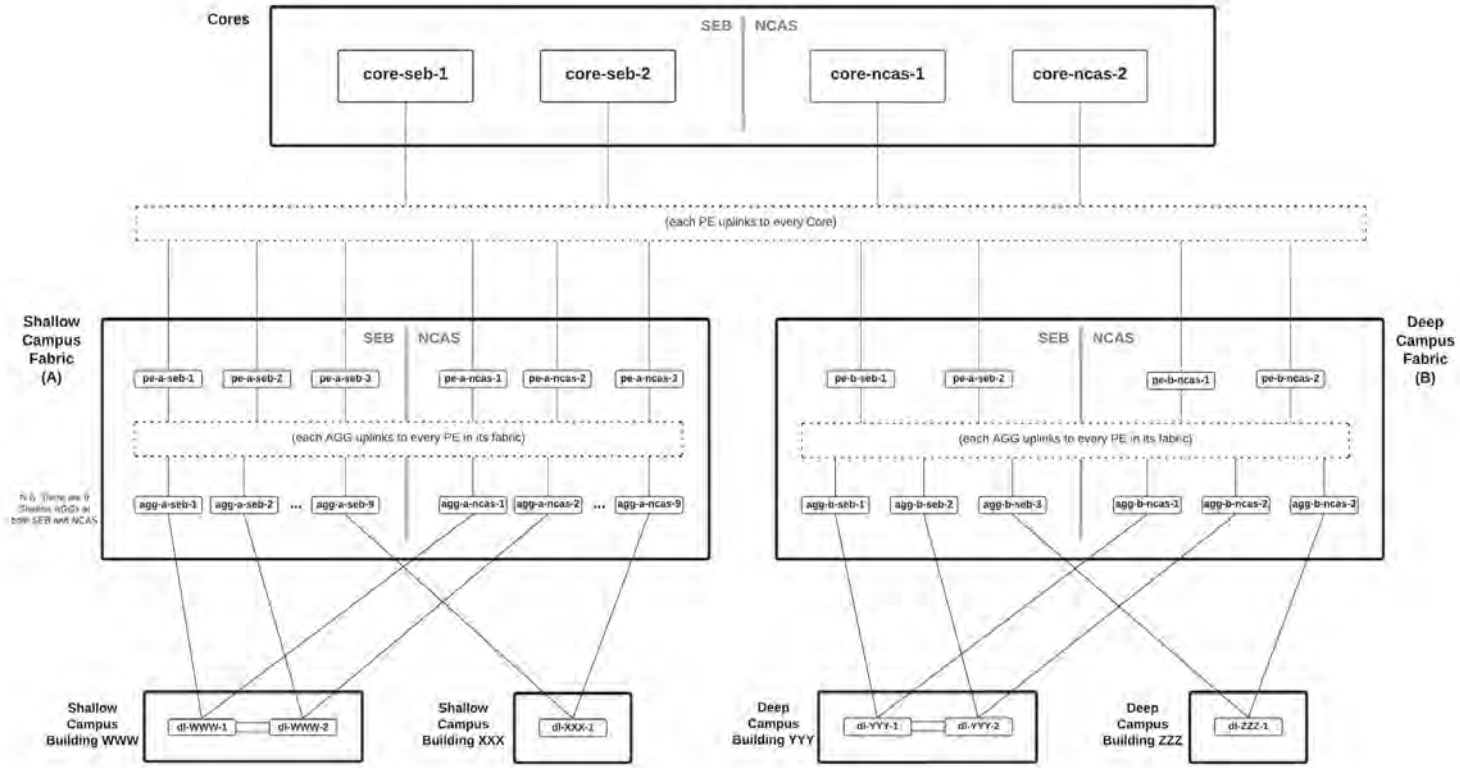


# Campus Fabrics

- Two core locations host redundant connections to all campus buildings
- Both fabrics exist in both core locations and are heavily interconnected via dedicated fiber between the two sites
- Fabrics consist of multiple 100G 1U Nexus devices: N3K-3636C-R and N9K-9336C-FX2
- Each fabric has two layers:
  - Aggregation layer hosts connections to buildings
  - “PE” layer aggregates aggregation and converts between EVPN/VXLAN and MPLS/L3VPN



# Campus Fabrics



N.B. There are 9 Shallow AGGs at both SEB and NCAS



# Securing the Network



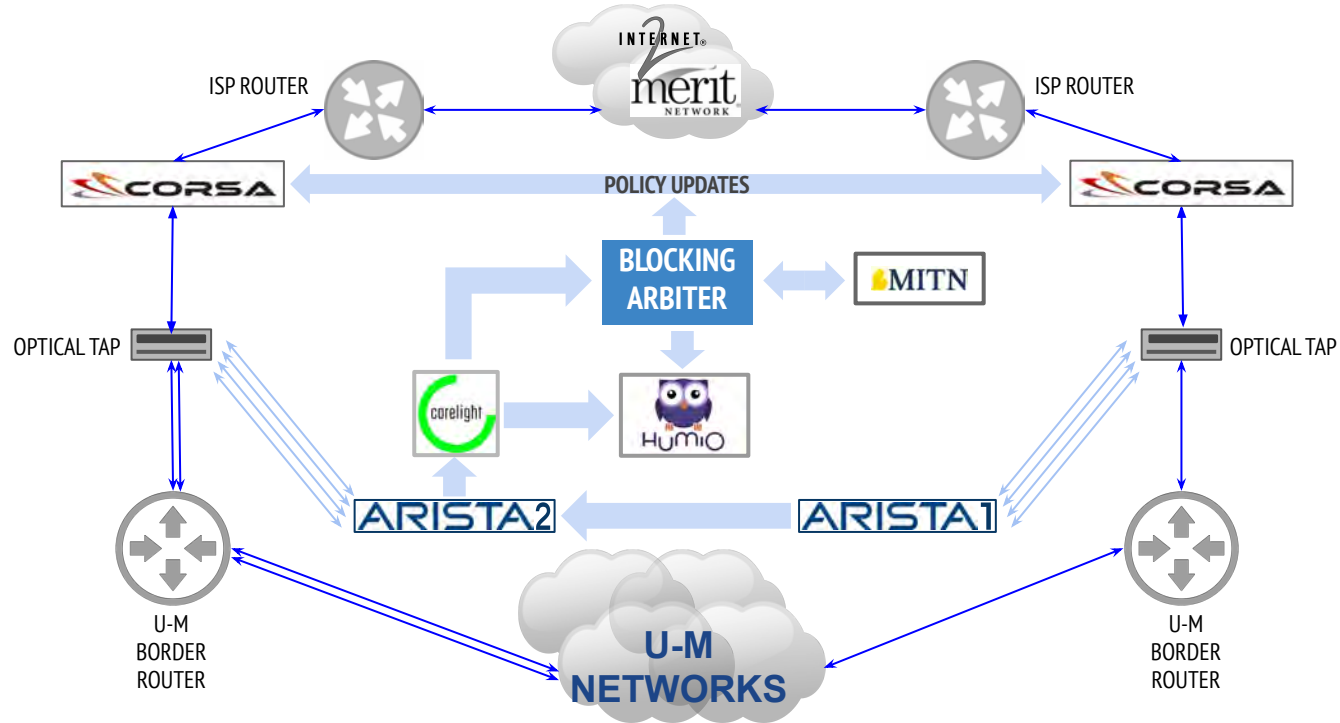
# Network Border Infrastructure System (NBIS) Goals

*Devise and deploy a high-speed, large-scale, cost-effective network border security service that does not negatively impact the university's research traffic.*









- Visibility of and security enforcement on network traffic to and from external entities
- High-speed: Line-rate Nx100G visibility and filtering
- Large-scale: Up to 10M filtering entries (problematic ...)
- Cost-effective: An order of magnitude less expensive than an Nx100G stateful firewall cluster



# Current NBIS Architecture



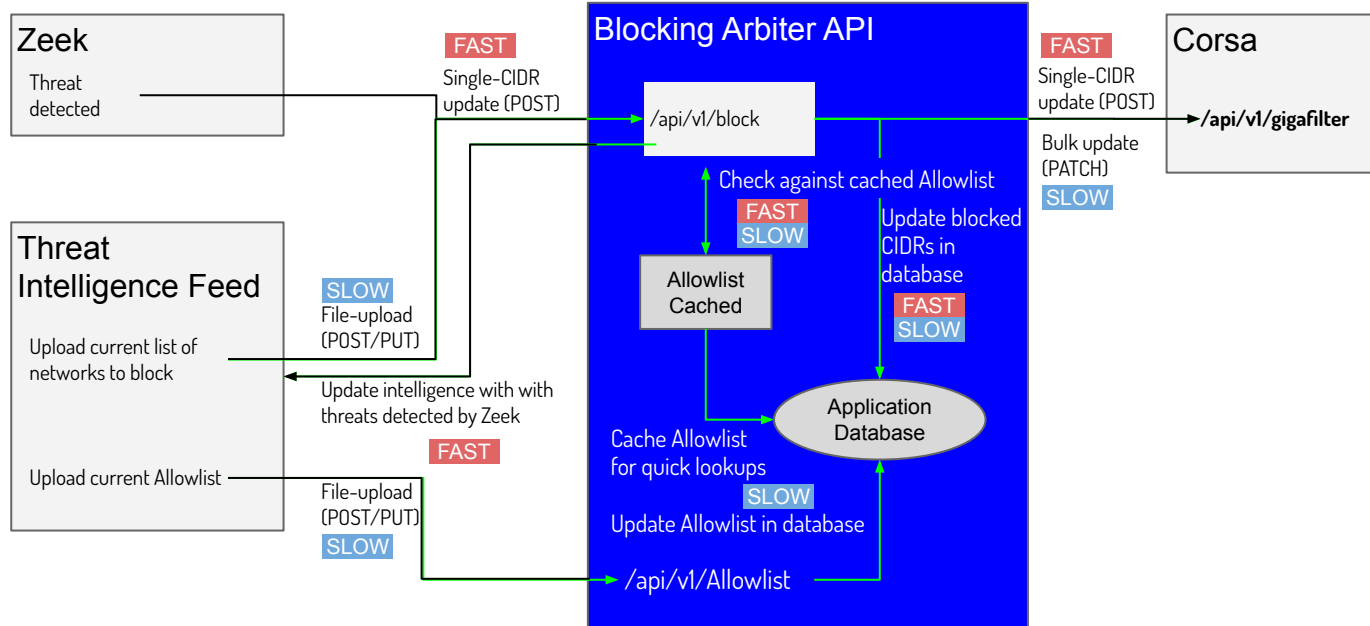
# Current NBIS Components

	Filtering/policy enforcement	Corsa NSE 7200	Filter network traffic based on programmed policy (using the "Gigafilter") Certain networks can bypass enforcement (e.g. research scans)
	Traffic mirroring	Gigamon Optical taps	Copy live network traffic for further processing
	Tap aggregation	Arista 7504E	Aggregate tapped traffic from multiple taps/sites Shunt low-value network traffic to reduce processing load
	Traffic analysis	Corelight AP3000 Zeek	Inspect traffic, analyze data, correlate events, detect attacks, record logs Scalable / flexible / programmable
	Log aggregation / analysis	Humio (li'l bit of Splunk)	Enable large scale data ingestion and export Provide instantaneous data views
	Threat intelligence repository	MITN	Collect threat data from multiple trusted sources (e.g. REN-ISAC, Spamhaus...), participating universities, and locally-generated data (e.g. Zeek)
	Policy / configuration updates	Custom in-house software	Glue together the various components of the system Update Corsa enforcement policy, perform data sanity checking Plans for open sourcing this software package
	Network traffic metrics collection / analysis	U-M perfSONAR infrastructure	Evaluate impact of security configuration changes

# Middleware (Blocking Arbiter) and Test Data Repository

## Blocking Arbiter

### Application Architecture



# NBIS Flexible Components → Upcoming Improvements

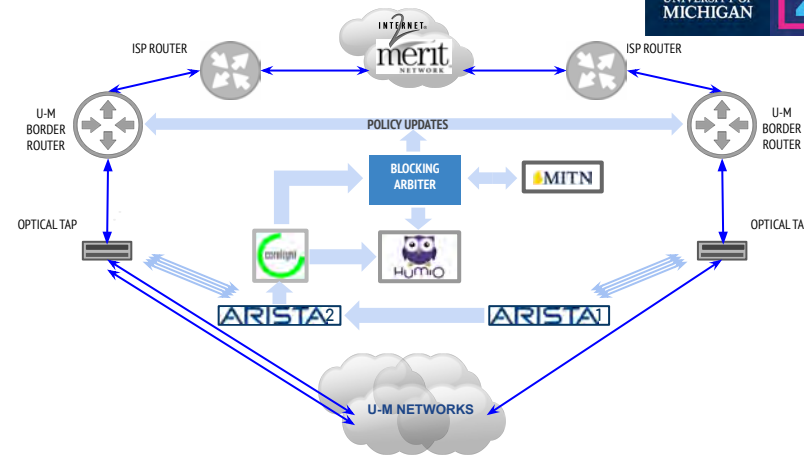

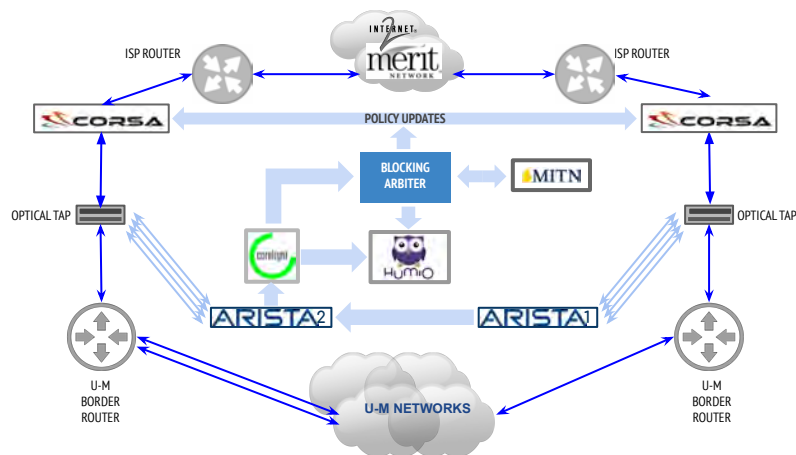


- Replace Corsas with Null Routings solution
  - Vendor does not fully support the product any more;
  - Null routing\* identified as best alternative;
  - Modified Blocking Arbiter code to support null routing.

\*A **null route** or black hole route is a network route (routing table entry) that goes nowhere.

Matching packets are dropped (ignored) rather than forwarded. The act of using null routes is often called blackhole filtering.

# Implementation Details: Current vs. New Architecture



## Current “North” to “South” architecture

- Merit (i.e. Internet)
- **Corsa (i.e. enforcement mechanism, including bypass)**
- Optical Tap (i.e. copy of traffic to send to Corelights / MiTN)
- **Border Router**
- Various university services

## New “North” to “South” architecture

- Merit (i.e. internet)
- **Border Router (including enforcement mechanism and bypass)**
- Optical tap (i.e. copy traffic to send to Corelights / MiTN)
- Various university services

# WiFi Everywhere





# About the WiFi Upgrade Project



## Built a world-class, transformational network

- Enables innovation
- Supports demand for fast and secure WiFi
- Concurrently supports high collaborations volumes
- Improves service for densely populated environments

In less than a year, replaced/added ~16,300 WiFi Access Points (APs)

First university in the nation to upgrade campus network to WiFi 6E



# WiFi Design Standards



## Redundant Architecture

- Three diverse data node locations for WiFi controllers
- Diverse connections to core infrastructure
- Redundant power supplies

## Standardized Hardware and Software

- Same (controller) version of code and type of hardware across campus
- Same AP equipment model whenever possible

## Authentication: Aruba's ClearPass

- Registration and authentication monitoring
- 11+2+1 dedicated servers located in multiple locations

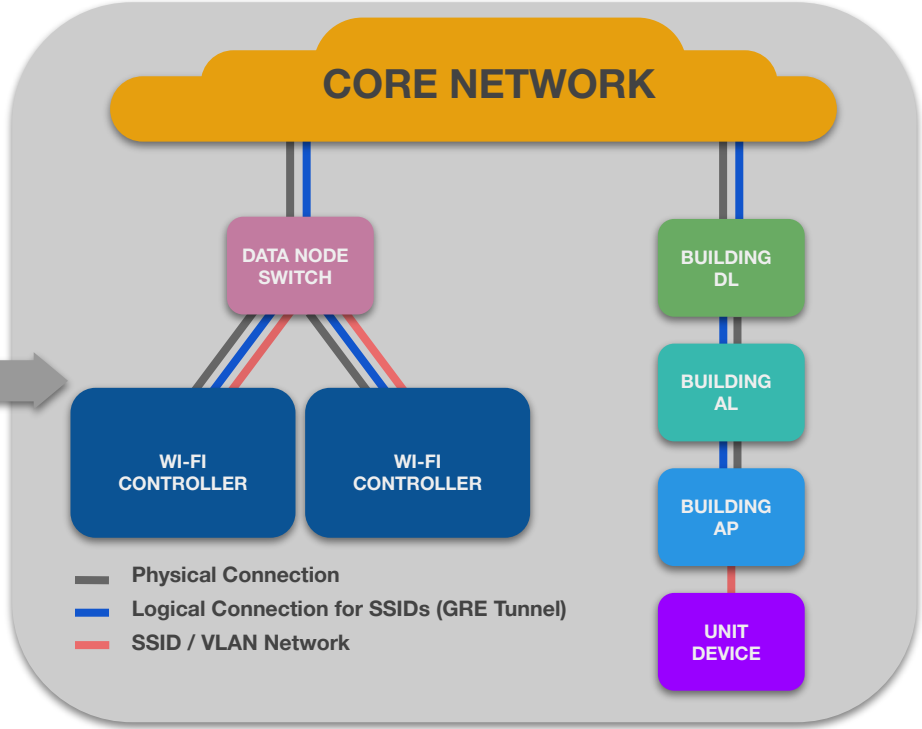

## Central Management Platform

- Monitoring and reporting for entire campus



# Typical Wi-Fi Architecture Diagram

## U-M Wi-Fi Architecture – Wi-Fi Controllers



# New WiFi In Production



## Noticeable Demand Increase: U-M WiFi September 2022 Metrics



42 Wireless Controllers



16,261 Access Points (APs)

TOTAL DATA TRANSFER  
**2.18 PB**


TOTAL DATA DOWNLOADED  
**1.73 PB**

TOTAL DATA UPLOADED  
**467.58 TB**

### Monthly WiFi Usage Totals (August difference)

Unique Clients Total	336,763 (+46,678)
Sessions Total (including connection attempts)	228.134 M
Average Session Duration	11 min
Traffic Total	2.18 PB (+1.3 PB)
Average Traffic / Client	1.58 GB
Average Traffic / Session	10.57 MB

### Monthly WiFi Usage / SSID



- MWireless (2.06 PB)
- eduroam (207.71 TB)
- MGuest (174.96 TB)
- MSetup (84.74 TB)
- Others (9.78 TB)

# New WiFi 6E Technology

*6 GHz spectrum approved by the FCC in April 2020*


## Benefits\*

- > doubled number of available channels
- HD video streams in densely populated settings
- Groundwork to provide next generation content sharing technology

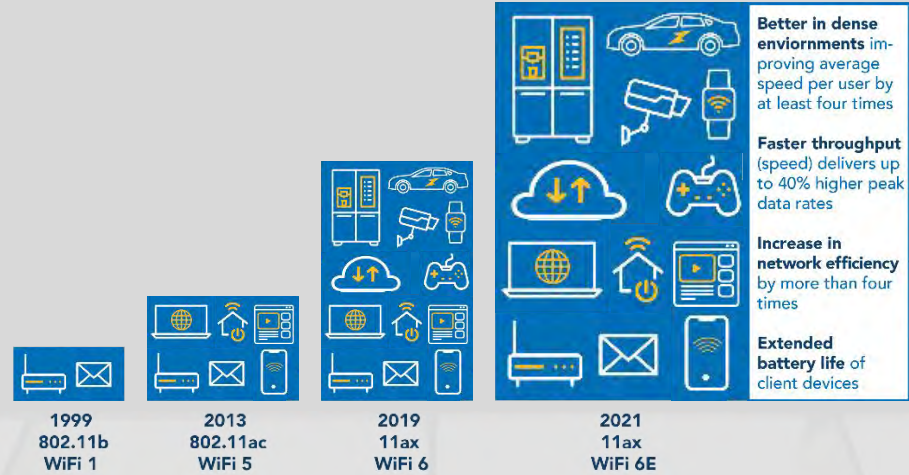
*\* Only WiFi 6E-enabled devices provide full benefits of the new technology*



# WiFi 6GHz Improvements vs. Previous Generation



Data provided by Broadcom



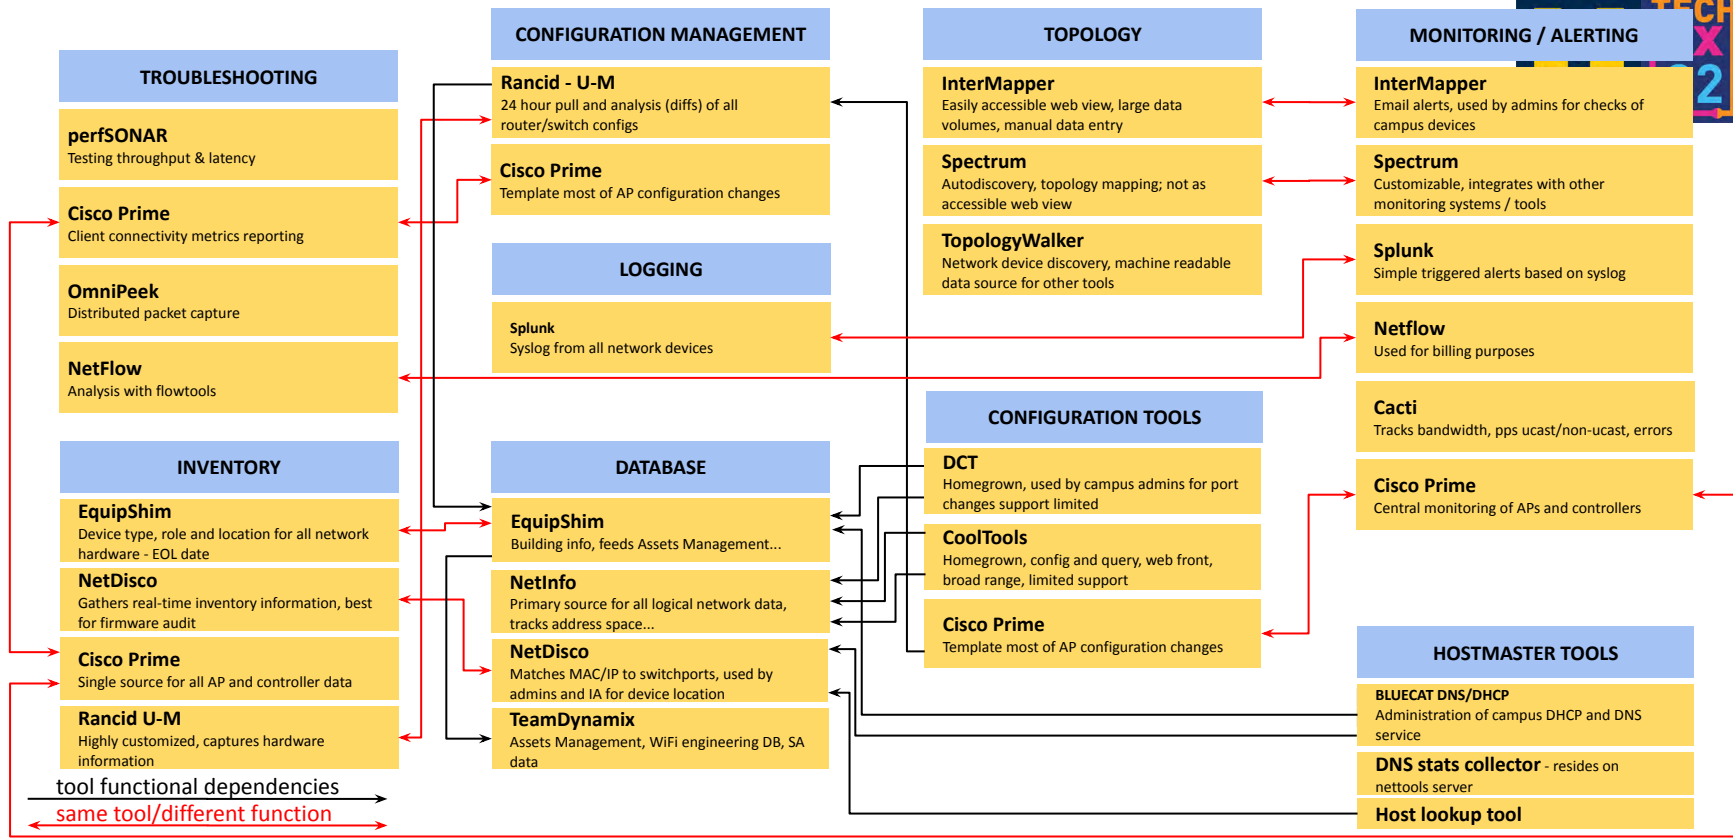


# Network Automation

*Technology empowers the leaders and best. We enrich the University of Michigan experience with technology that inspires people to do amazing things.*

*– University of Michigan IT Vision –*

# Current State of Network Tools



Significant overlaps in functionality, tool interdependencies, and staff turnover have made our current tool chain difficult to maintain.



# Legacy Tool Replacement Strategy

1. Implement “service abstraction” to
  - Better organize and structure data
  - Increase reliability and maintainability of software
2. Cisco’s NSO product meets the needs of this generation of network



# Establish Sources of Truth

## NetBox: DCIM/IPAM - in development

- Device inventory and categorization
- Asset tracking
- Prefix, VLAN, and VRF assignments

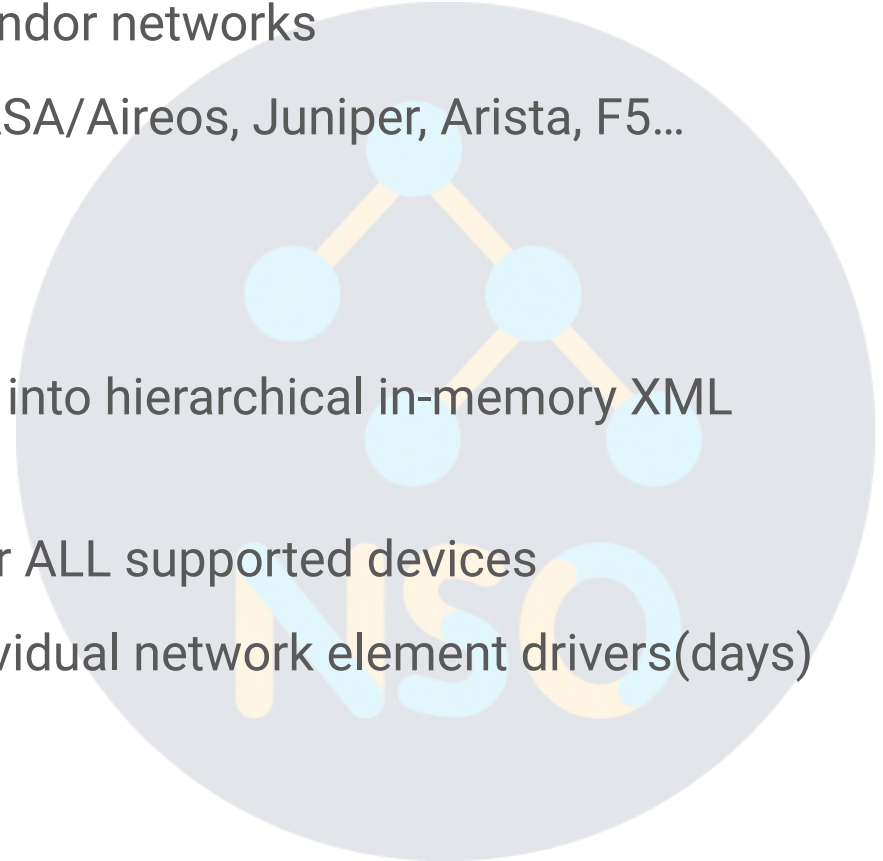
## NSO: Configuration

- Device configuration management
- Source of Truth (SoT) for service data
- Drift check
- Operational state snapshot



# NSO - Multi-vendor Automation

- NSO - purpose-built to manage multi-vendor networks
  - support for Cisco IOS/XR/XE/NX/ASA/Aireos, Juniper, Arista, F5...
  - Based on NETCONF+YANG
  - Not an appliance
- Abstracts network device configuration into hierarchical in-memory XML database
- Transactional configuration changes for ALL supported devices
- Fast feature development cycle for individual network element drivers(days)



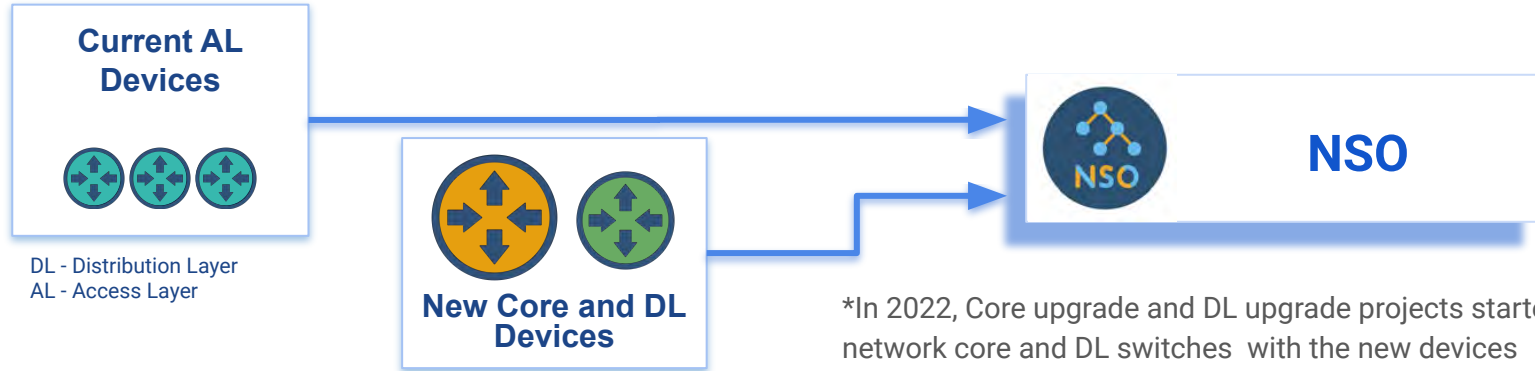
# NSO Data Migration

## Current

- Configuration backups - RANCID
- Configuration automation
  - Perl scripts
  - Ansible

## Upcoming

- NSO - configuration management
- NSO to represent campus network
  - Add existing AL devices
  - Only new Core and DL devices to NSO\*



# Replace Legacy Configuration Tool / DCT

## Current

*DCT: "Device Network Configuration Tool"*

Retrieves configuration and operational state via SNMP and displays

Configures AL devices directly via SSH/expect

Monolithic perl app encompasses web front-end and back end communications.



## Upcoming

*NetDash*

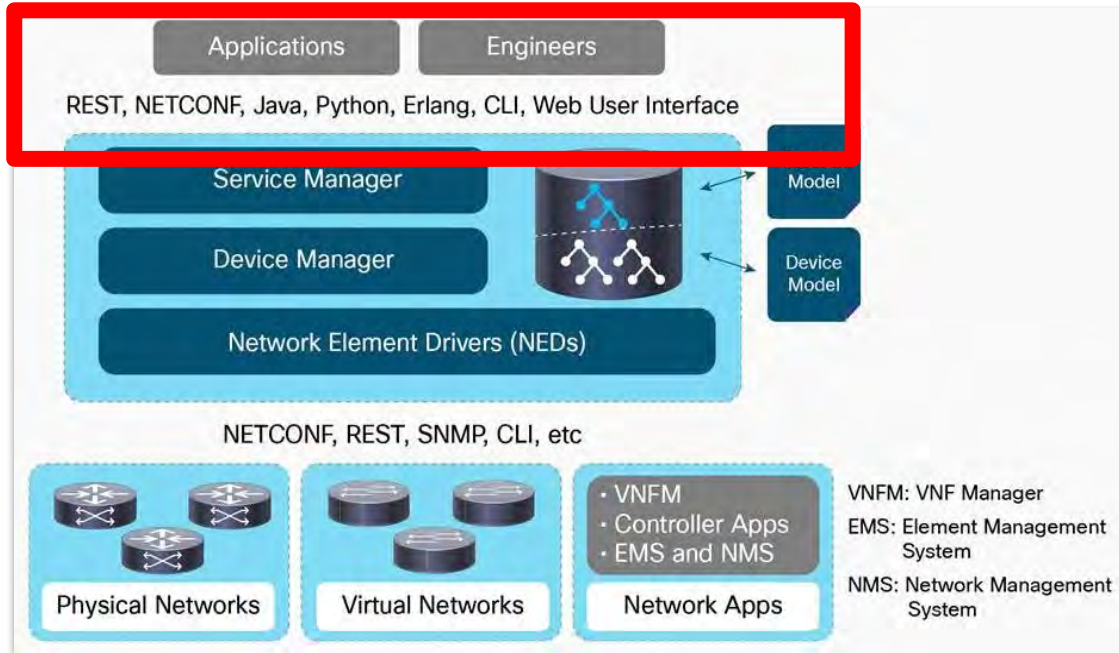
Retrieves configuration and operational state via custom NSO actions

NSO handles communication with devices

Front and back end separation makes app support easier

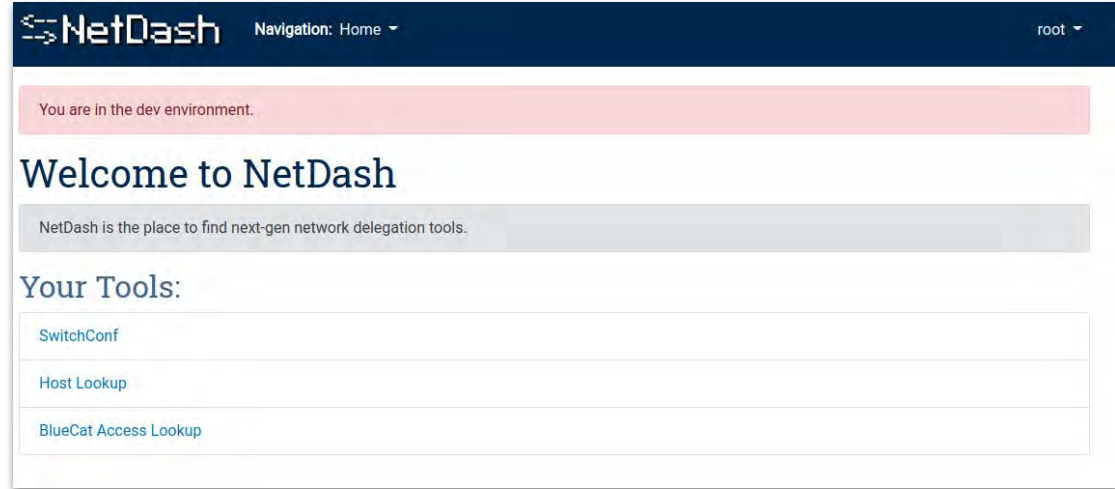
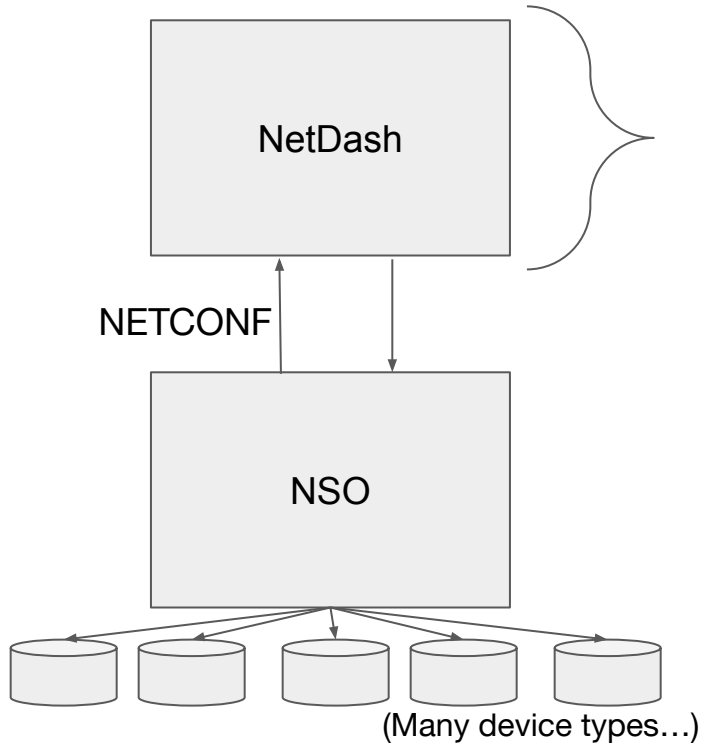


# Building Software on NSO - Interfaces



- REST
- **NETCONF**
- Java
- Python
- Erlang
- CLI
- Web UI

# Building Software - “NetDash”



- Uses Django (Python web framework)
- Much simpler codebase, targeting only “NSO” instead of dozens of device types. (Fewer code paths)

# NetDash Web Interface



INFORMATION AND TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN

Login

Back

IP: 10.224.36.106

Location:

Manufacturer:

Model: C9300-48T

Name: al-ilab-a1-5

OS:

Edit

Port	Descri...	VLAN ...	VLAN Id	VoIP	Speed	Duplex	Admin...	Operat...	MAC ...
Gi1/0/1	unused	default	1	✗	auto	auto	✓	✗	None
Gi1/0/2	unused	default	1	✗	auto	auto	✓	✗	None
Gi1/0/3	hello world	V-BLDGA-...	13	✓	auto	auto	✓	✓	[00:d0:c9:...



UNIVERSITY OF MICHIGAN





# NetDash Web Interface

Back

Edit

Port	Descri...	VLAN ...
Gi1/0/1	unused	default
Gi1/0/2	unused	default
Gi1/0/3	hello world	V-BLDGA-...

## Edit Data

Description

VLAN 

- ✓ -Select VLAN
- 10: NGFW-LSA-GADGETS-BLDGA
- 13: V-BLDGA-USER

Speed

Duplex

Admin Status

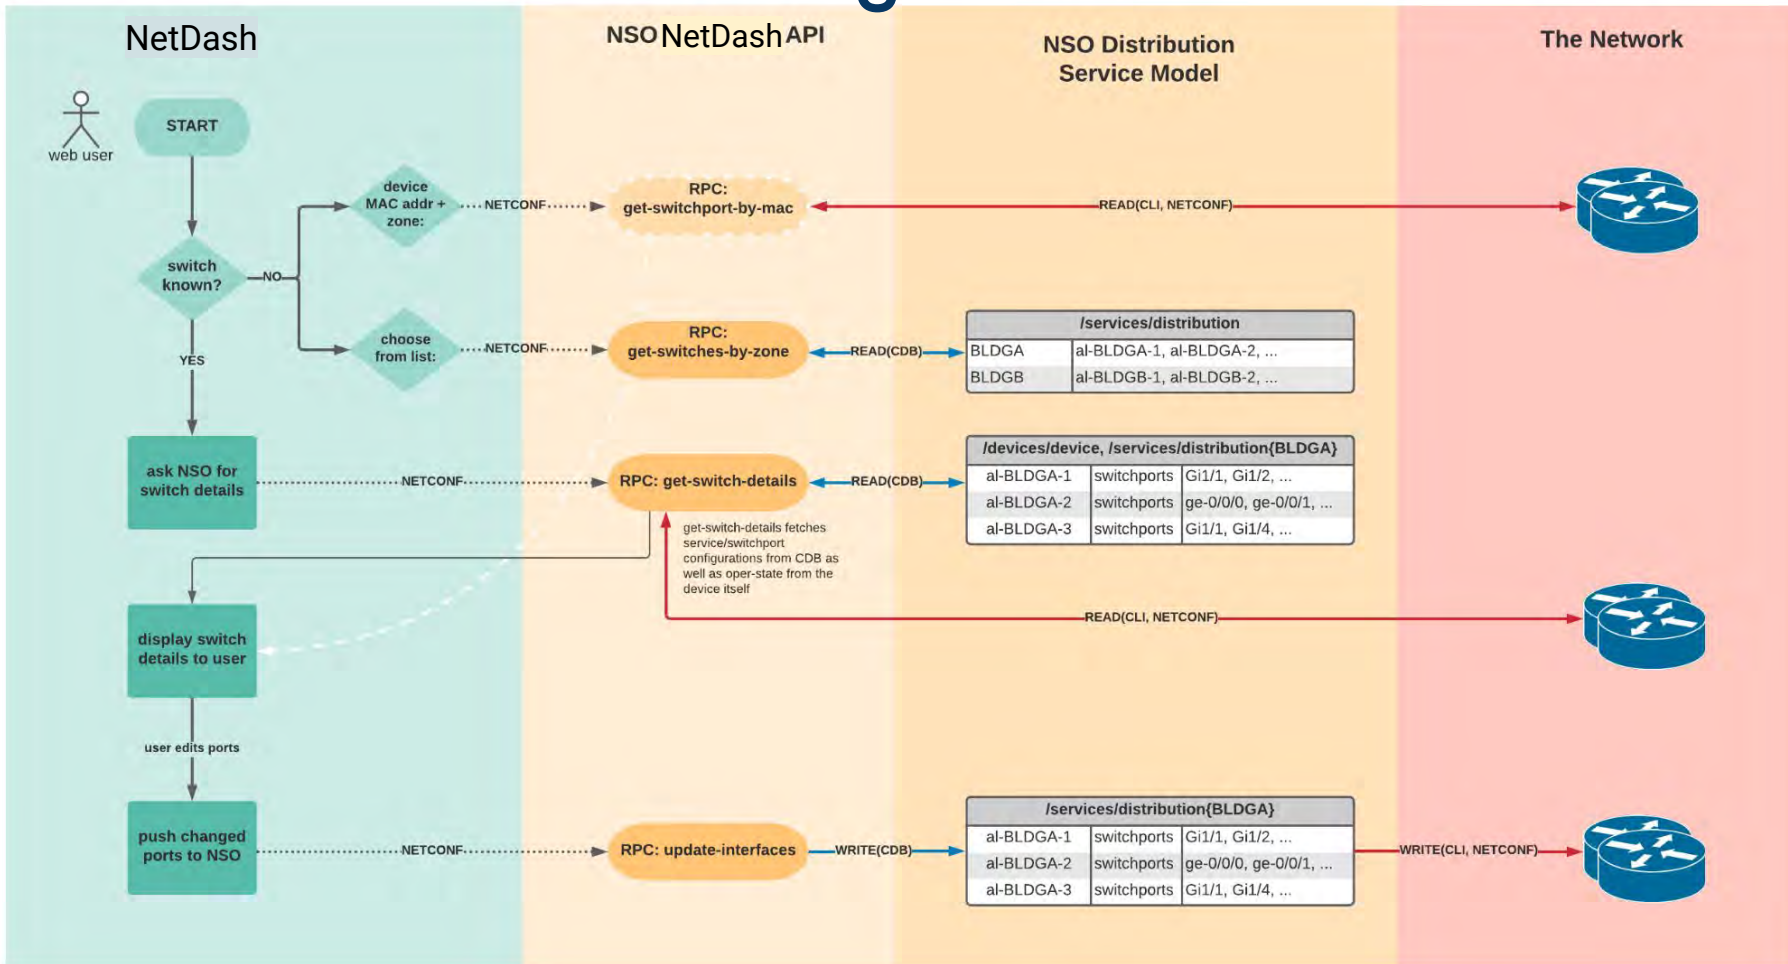
VoIP

Close

Save changes

Admin...	Operat...	MAC ...
✓	✗	None
✓	✗	None
✓	✓	[00:d0:c9:...

# NetDash / NSO Processing Flowchart



# Building Software - NetDash SwitchConf (3)


```
<action xmlns="urn:ietf:params:xml:ns:yang:1">
  <services xmlns="http://tail-f.com/ns/ncs">
    <netsplash xmlns="http://umnet.umich.edu/netsplash">
      <update-interfaces xmlns="http://umnet.umich.edu/netsplash">

        <config xmlns="http://tail-f.com/ns/config/1.0">
          <services xmlns="http://tail-f.com/ns/ncs">
            <distribution xmlns="http://umnet.umich.edu/distribution">
              <name>ARBL</name>
              <switch>
                <name>s-ARBL2-1496-1</name>
                <switchport>
                  <name>ge-0/0/0</name>
                  <description>1410-03C</description>
                  <mode>
                    <access>
                      <vlan>NGFW-ITS-COMM-AL</vlan>
                    </access>
                  </mode>
                </switchport>
              </switch>
            </distribution>
          </services>
        </config>

      </update-interfaces>
    </netsplash>
  </services>
</action>
```

Envelope - NSO Action

Configuration Payload



# Skills, Training, and Technical Standards

- Git / GitLab / GitHub
- Ansible / NSO / YANG
- NetBox
- Python [py4e.com](https://py4e.com)
- Software development practices



# Questions

Eric Boyd

<[ericboyd@umich.edu](mailto:ericboyd@umich.edu)>



# Thank You

