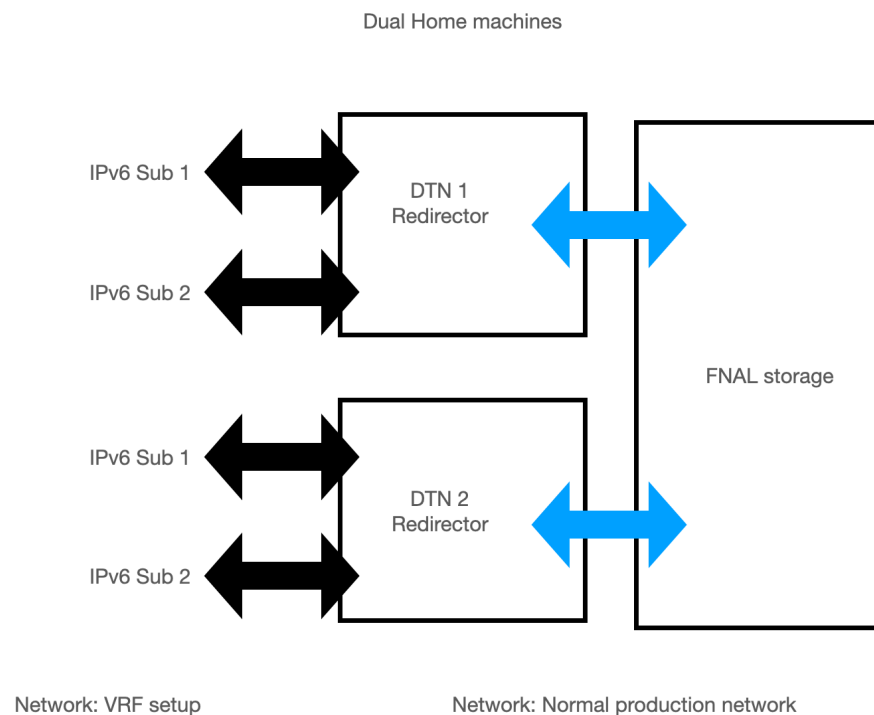


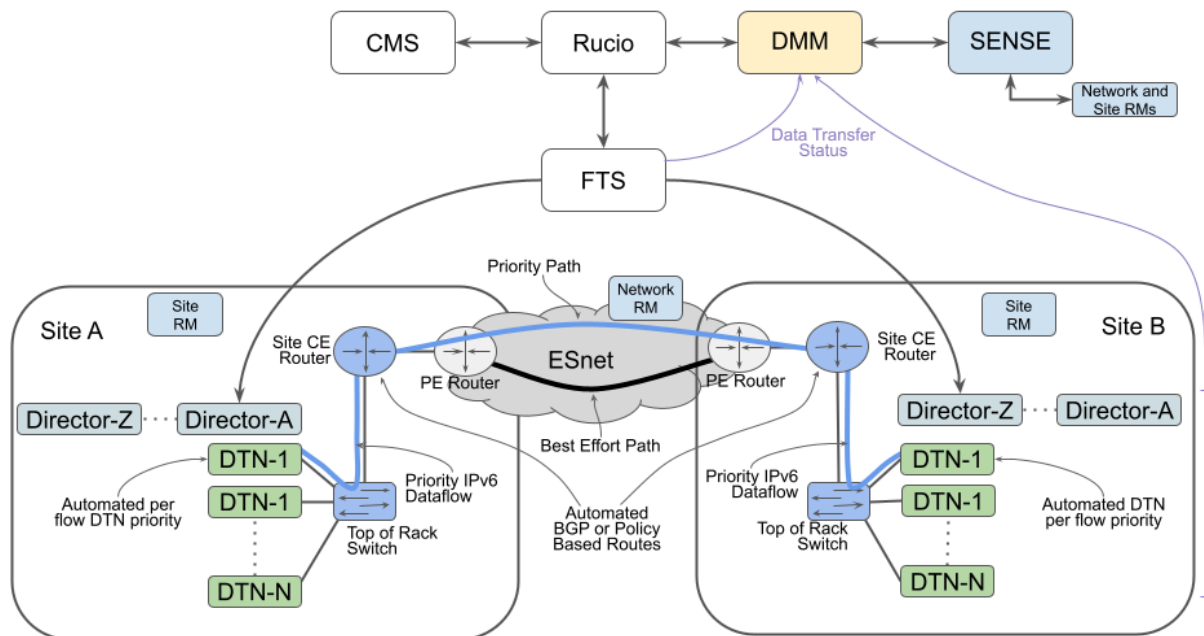
FNAL Rucio/SENSE testbed

220603



- Goal: support 100 Gbps aggregate through the IPv6 subnets
- How many DTN machines do we need?
 - Start with lower number of machines
 - Then scale up to reach 100 Gbps
- Switching should already be setup for 100 Gbps
 - We can support 100 Gbps NICs in the DTNs if needed
 - We can add a 100 Gbps node if needed
- **Decision:**
 - **we start with 2 DTN nodes**
 - Dual-homed 10 Gbps
 - 1 for IPv6 subnets
 - 1 for the FNAL storage network

- Wait if additional control plane connection is needed (FNAL storage network is overwhelmed)
- **Additional server for the directors**
 - Needs to have access to all the IPv6 subnets
 - Control plane in the FNAL network
 - Dual-homed as well
- Important:
 - IPv6 subnets are not FNAL subnets, they are SENSE subnets, only used for data transfers
- Find 3 nodes in FCC → Dave (and Fagan)
- ↗Later, we need dedicated network switching infrastructure
 - FCC is better for that reason as well



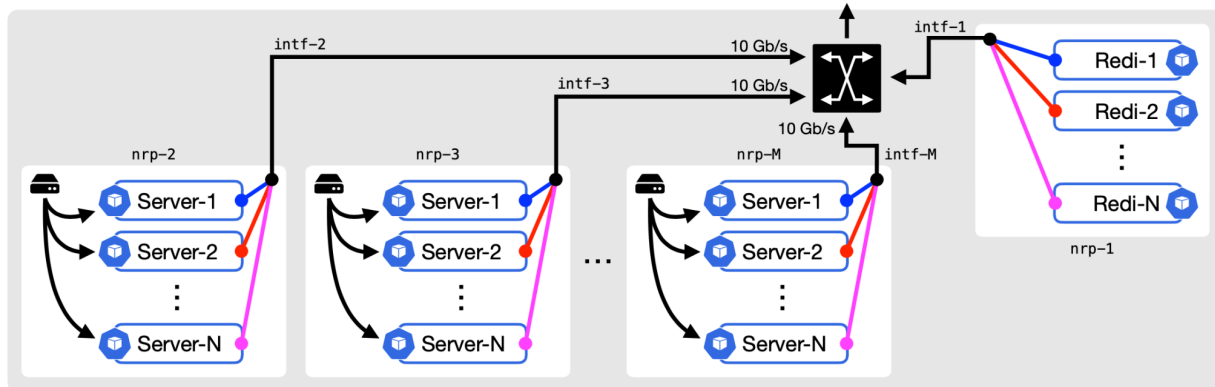
- **In general, we want to manage the software on the nodes and the network, not externally**
- Time schedule
 - Provision 3 nodes → 1 week
- Need to get more detailed information/documentation how the DTNs are configured at Caltech/UCSD
- Slack channel → Dave **DONE**
- Listserv → Oli

220617

- Hardware is progressing
 - Most nodes have dual network cards
 - Request for 3 nodes
 - Where would we like them to be? → want to start figuring out how to shape the virtual network
 - Dave will add networking people on the ticket
- Long term goals for Fermilab
 - Explore WAN network management within the Fermilab storage architecture (Is the DTN model the right one for large installations like FNAL?)
 - We will run our own services, no remote access to run services
 - We want to perform the tests within the same infrastructure and the same tools as the production infrastructure
 - Example: we want to create virtual networks within the production network
 - Example: operation model: puppet vs. machine in the corner manually administered
- Joining the testbed
 - Stage 1: just to see how this works
 - We can put virtual networks into our production network (VRFs)
 - Need to put in a SENSE control plane, this is a server (we need to manage it)
 - We want to use one NVME pool (100 TB) as storage backend
 - Stage 2:
 - If need special networking hardware, then we need to integrate that and have special routing (example: P4 switch)
 - Longer term: proxy'ing on xrootd doors → this will not take a long time (it is already implemented for gridFTP and WebDAV)
- Slides giving more details of the UCSD setup:
https://www.dropbox.com/s/1x9kezshvtwc6ef/220615%20-%20rucio-sense_site-config.pdf?dl=0



R&D Rucio-SENSE Prototype



- For N priorities, M available nodes/interfaces:
 - Each node has 1 **physical interface** with N internal virtual interfaces (i.e. **blue**, **red**, and **magenta** above)
- Servers write to disk (🗄️) on each node

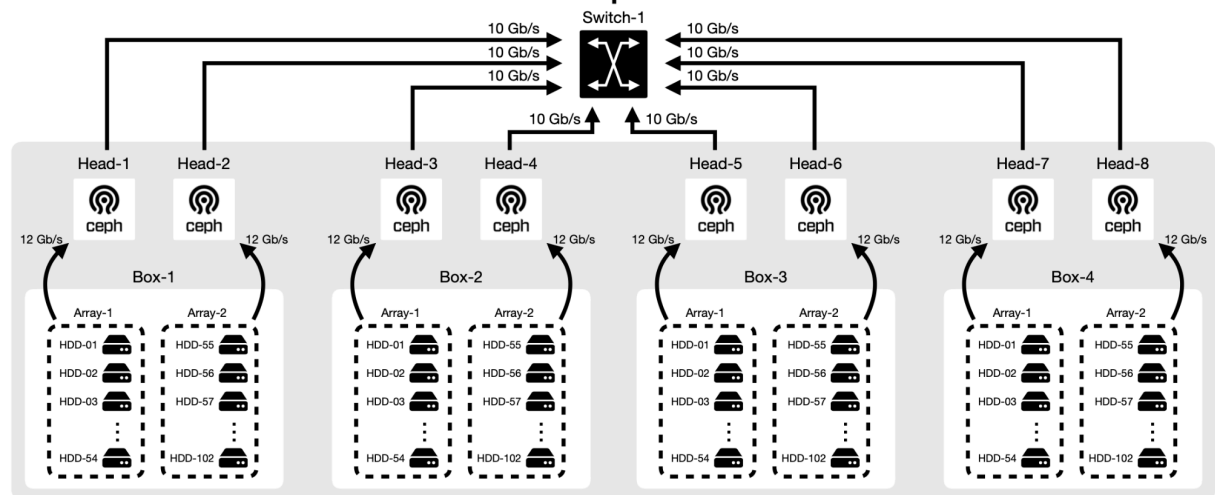
2

UC San Diego



Ceph @ UCSD

Switch-1 Max Capacity: $2 \times 40 \text{ Gb/s} + 100 \text{ Gb/s}$ **Current Max Output from Ceph: 80 Gb/s**



Only Box-1 & Box-2 (effectively) will be mounted until more hardware arrives; ETA ~1 mo. (end of Feb.)

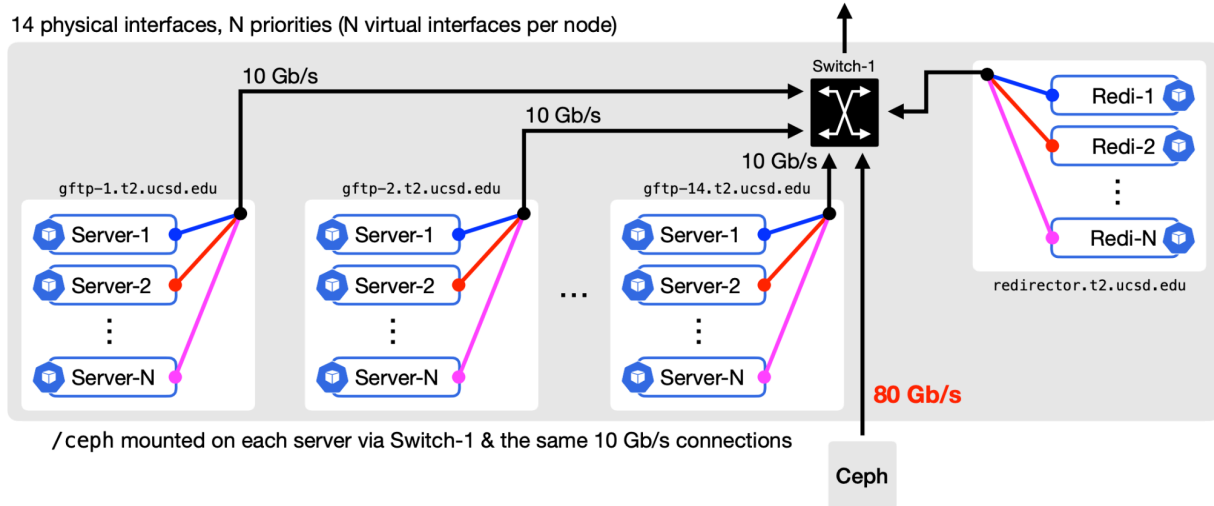
3

UC San Diego



Ceph + Rucio-SENSE Prototype @ UCSD

14 physical interfaces, N priorities (N virtual interfaces per node)



4

UC San Diego

- How do we proceed
 - We get the 3 machines in hand → Dave
 - We get the NVME array (somewhere in FCC) → Oli
 - We get the SENSE control plane machine → Andrey
 - Create layout of the network (where are the machines most likely) → Andrey, Ray, Phil
 - For now, we are looking at 10 Gbps! We are not looking into putting in 100 Gbps at this time.

220701

- Rucio/SENSE integration news from meeting with Tom Lehman
 - Talked about integration of the systems from the high level: function of DTN, FTS, XRootD redirector, Rucio
 - Rucio → SENSE
 - Rucio → DTN
 - FTS is doing the transfer job assignment
 - Not clear what the relation with xrootd is
 - Checking back with UCSD developers → not clear to them who is launching the data transfers: DTN, XRootD or FTS?
 - They are documenting:
 - Who is launching the data transfer
 - Some high level design of the architecture seems not to be fully understood
 - “Bottom design that does not pass a top-down review”

- Eduardo (WAN person at CERN) asks if FNAL wants to be a part of a SC22 demo of NOTED (detects WAN transfers and chooses optimal paths)
 - Last year they only got Triumph and wiped out the network in British Columbia
 - Only FTS based, not at the level of complexity of SENSE/Rucio test
 - SENSE is used to arrange the WAN for this
 - We answered: if it is not too much effort and also they don't expect to run at scale, we could consider to be part of it
 - Arrange a fixed path within the testbed
 - Deadline: October 1st to get something setup
 - Question:
 - Do we need to provision a circuit? OSCAR like circuit, switch on and off
 - Will sit down with Tom Lehman to figure out details
 - October/November: IERC and PIP2, SURF will come online at the same time as this demo, might not have effort
 - From a networking point of view: stitching together the network for the testbed
- Hardware:
 - 3 DTN machines
 - Machines have been identified: FCC-2-1585, FCC-2-1586
 - Dave had not yet getting hands on them
 - Dave is setting this up, but also on vacation this week
 - Interfaces: we have 2 interfaces
 - What network connections need to be established? Physical connections.
 - One end is the CMS network
 - Other and is the VRF network
 - → both network connections will go to CMS
 - Phil: need to take step back and look at what we want to do: both for Caltech-UCSD test but also for the NOTED test
 - Ask Eduardo what we will need for NOTED → Phil
 - Then compare with the Caltech test environment, discuss with Andrey from network point of view, then Dave for system level point of view
 - SENSE control plane machine (SENSE RM)
 - Somebody needs to provide machine
 - Somebody needs to install software
 - We are planning that the new hire would install it, the person is not there yet
 - Software we get from SENSE people
 - Worry that we don't have people to do that right now
 - Let's find out what it takes (repository, instructions, etc.) → Oli
 - We take one of the testbed systems for this
 - NVME: Oli needs to follow up
 - Think about it: if we get an NVME, do we get a server in front of the NVME as well?

- Backup: After August 9, we can use an NVME array that Dmitry is using for Spark studies with an intern
- Proxying in xrootd
 - Dmitry is implementing this right now
 - If you have pools in the DMZ and pools in the private network, the xrootd door will be able to serve files through proxying (this is the BNL setup)
 - Xrootd door on DTN, pool in private network
 - Data will flow from private pool through DTN to customer

220715

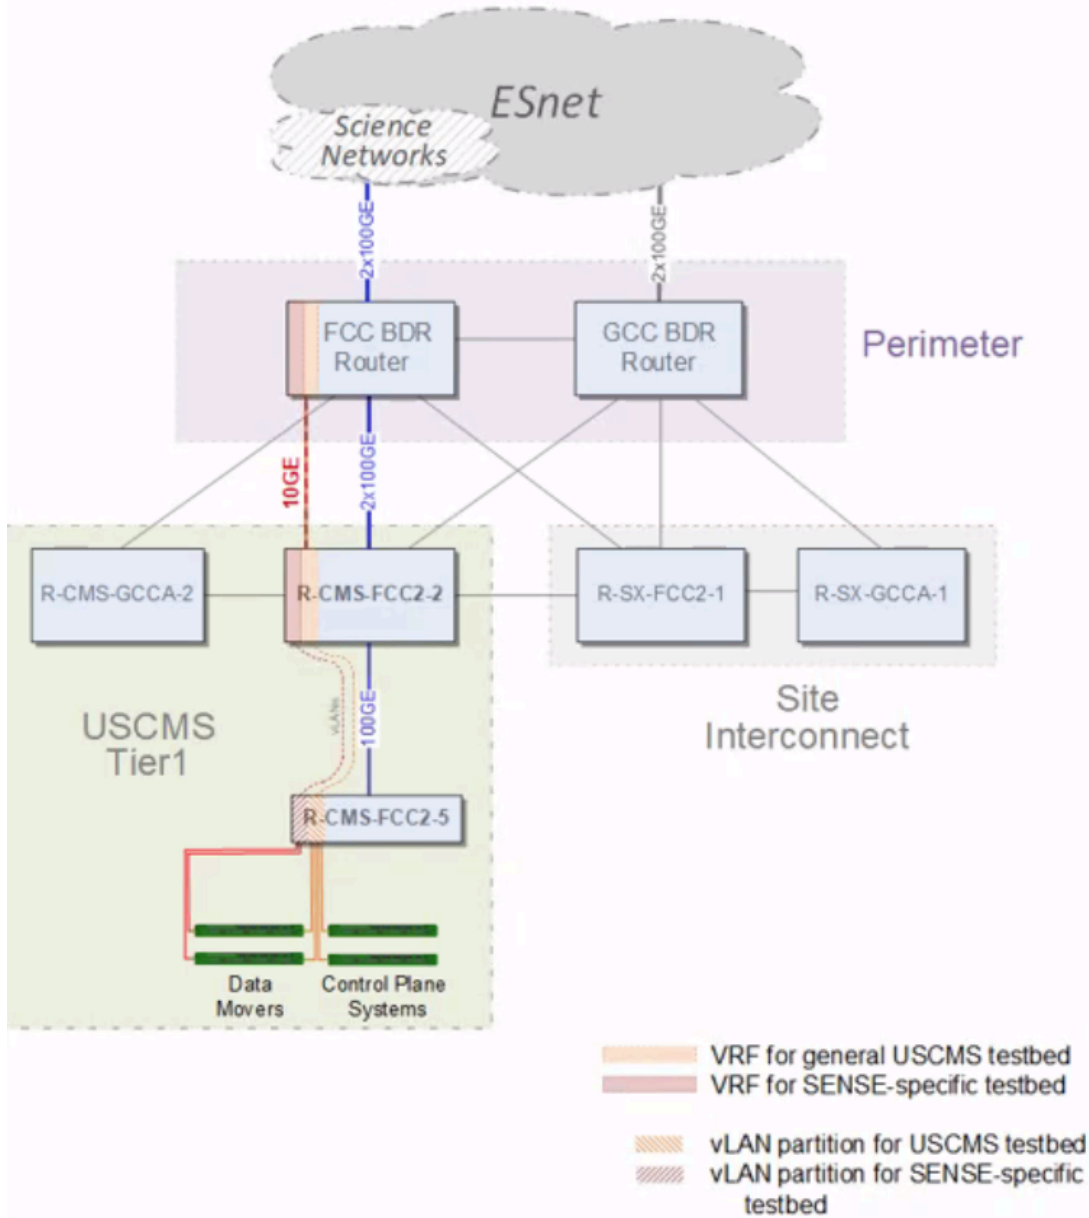
- New listserv: fnal-networking-r-n-d-testbed@listserv.fnal.gov
- 3 DTN notes
 - Waiting for Dave for an update
 - Systems:
 - Fcc-2-1585
 - cmsspare1300 -- cmssense1
 - cmsspare1306 -- cmssense2
 - fcc-2-1586
 - cmsspare1303 -- cmssense3
 - Dcache testbed is the source of storage
 - Planned is to use a NVMW node
 - Each of the DTN will act as an xrootd door
 - Could either mount nvme through PNSF
 - Good first method to try it
 - Could make them part of the test dcache and could communicate directly
 - Would require proxying in xrootd, currently being worked on but needs more development work
 - Every DTN node has access to all files
- Network setup
 - Andrey very busy, next week sitting down with Phil to think about details of networking setup of the testbed
 - VRF: virtual routing facility (similar to IPv6 testing setup)
 - Maybe some dedicated physical infrastructure will be added in the mix
 - An extension would be a P4 switch
 - Dual-NIC configuration
 - NIC1: WAN movement through SENSE
 - NIC2: general connectivity → have not yet talked about the subnet, but would look like a production subnet
- Local network controller: NRM (network resource manager)
 - Phil has boxes from old testbed
 - Need to define who is setting it up and define details
- Will layout time schedule next week

- Andrew: Vanderbilt joining testbed
 - Hardware is available
 - Question: configuration at Vanderbilt border → requires process
 - DTN boxes with storage mounted (LStore), posix mount, standard xrootd server setup
- Software
 - Prototypes for testbed software components available, but still under active development

220812

- New people joined the meeting
 - Alan
 - Ryan
- Dmitry - Storage
 - Currently adding proxying data in the xrootd door → this is done, available in the master branch, available in the next golden release in September
 - DTN has two network interfaces
 - 1: SENSE VPN IPv6 interface
 - 2: different network to dCache server
 - New xrootd server will be installed on DTN, can receive request on SENSE VPN interface and can server files through the DTN from the dCache servers, although the dCache servers are on a different network that the SENSE VPN interface
 - DTN setup
 - One general interface for all traffic
 - One special interface for all VPN IPv6 data movement
 - NVME:
 - Dmitry has one
- DTN nodes:
 - DTN nodes are available
 - Nodes have QLogic 10 Gbps cards
 - Should find/buy Intel/Mellanox cards → DaveM is telling DaveF to go ahead
 - Should be provisioned and ready to go
- Next steps
 - Dmitry: install xrootd on DTNs and dCache on NVME
 - Network setup

USCMS Testbed Design



- Talk to Frank and his SDSC people to talk about which parts of the network need to be exposed to public and which not → implementation is changing
- Next question: who is learning about the SENSE/Rucio testbed team software and installs the software on data movers and control plane systems
 - Dave will talk to SSI for help to deploy the software

220826

- Oli unable to attend; Phil running the meeting
- Need to have timeline estimates on all testbed components to determine level of possible participation in SC demo's
- Network Update:
 - Standard change expected in order to implement VRFs within the T1 for the testbed
 - Question on whether testbed systems are completely dedicated to testbed, as well as any other dependencies on production systems
 - Will get new IPv4 & IPv6 address blocks assigned
 - Access model needs to be developed
 - Access to T1 only, or beyond
 - Model for security privileges on testbed systems needs to be defined (need to understand what access will be required on DTNs)
- Storage update:
 - Chih-Hao will check with Dmitry on whether changing to new subnets will impact the storage deployment for the testbed
- DTNs/systems update:
 - Phil will follow up with Dave on whether SSI support for deploying SENSE-specific software on DTNs & control systems

220909

- Oli on vacation; Phil running the meeting
- Network Update:
 - Networking is establishing a spreadsheet of tasks to get the network component of the test bed deployed
 - Currently targeting completion for very early October
 - Will have the spreadsheet with target completion dates for the next meeting
- System/DTN update:
 - 3 DTN/control systems should be ready to go; just need new addresses & network connections
- Storage update:
 - NVME system is available:
 - It has only a single (40GE?) NIC
- Access model for test bed subnets:
 - Offsite access will initially be default-deny with exceptions for collaborating remote systems

- T1 subnets will have access to the test bed for administrative/management purposes:
 - Should also satisfy DNS and authentication services needs
- 100GE capability:
 - Initial deployment will be based on 10GE
 - Looking toward 100GE networking as part of initial rollout, but probably this is a second step
 - DTNs would require only a single 2x100GE NIC
 - Talk to Dave F about procurement when 100GE becomes plausible

220923

- Networking side
 - Working on deployment
 - Broken into several steps
 - Plan is mature, implementation is progressing, competing with other priorities at the lab
- 4 machines are available and software can be installed
 - 3 DTN's
 - 1 NVME
 - Login to machines?
 - NVME login available
- Order of deployment
 - First network has to be configured
 - Then we can talk about installing software and talking about architecture
- SuperComputing
 - Week of November 14
 - Two potential projects
 - SENSE/Rucio → UCSD/Caltech/ESnet
 - Noted → CERN
 - Need a month for both teams to implement things to include us in the demos
 - Oct. 10
 - We are on target for network deployment

221007

- NVME machine
 - storagedev202
 - How many interfaces?
 - One interface with 100 Gpbs
- Network deployment
 - Ryan:
 - most of the networking is configured on the Fermilab side
 - Monday: move the servers so that they have 100 Gbps in the future

- Dave Fagan's suggestion to move them
 - Can configure IPs on these machines now (today or Monday)
 - Connectivity offsite: request to ESnet to bring up the connectivity in the discussed isolated fashion
 - On track (as long as we can interface with ESnet) to finish deployment next week
 - Interface names and information on hand to configure them
 - Phil has a spreadsheet with names, will circulate
 - Can change if other consistency in naming is needed
 - Can send out go for people to start installing software when done
 - Discussion about storagedev202
 - Which network should it go through?
 - Actually would need another interface for control network
- Machines
 - Dave Fagan might have 100 Gbps cards, otherwise we might order some (supply chain issues?)
- After network setup complete, need to deploy software
 - Box that controls SENSE
 - Other middleware that needs to be deployed to tie in Rucio
 - Configure the DTNs
 - Who does what?
 - Dmitry: dCache on storagedev202
 - Maybe Chih-Hao: deploy the DTNs
 - Justas or Chih from ESnet: deploy control box (to get started, will figure out general setup for the future)
 - Communication: add Justas to Slack channel
 - Also add the networking people (Ryan, Alan) to it
 - First initial meeting:
 - In two weeks
 - Overview followed by discussion → from the software deployment point of view

221021

- Network deployment:
 - Physical server connections are now in place, but not plugged into servers
 - Access exists between T1 and the testbed subnets
 - Connectivity to the border in place and is working
 - Off-site access via general IP routed network
 - IPv6 access from rest of site still needs resolution of DNS anycast in the testbed subnets
 - Addresses have been assigned:
 - NVME addresses need to be put in spreadsheet
- NVME server:

- Currently has 1 100GE & 4 10GE interfaces
 - Will investigate replacing Qlogic 100GE NIC with Mellanox 2x100GE NIC to provide 100GE access for both testbed networks
- Ready to be configured
- Machines:
 - All the initial testbed machines are in place
- After network setup complete, need to deploy software:
 - This is a Dave M. & Oli issue
- Action items:
 - Distribute address spreadsheet
 - Contact Dave F. about address configuration on systems

221104

- Network deployment:
 - Two testbeds should be functional
 - General purpose testbed
 - SENSE IPv6 testbed
 - Traffic offsite and internally should be ok
 - IP information for 4 servers sent to Dave Fagan
 - Access to the testbed from Fermilab
 - First only T1 (includes FCC and GCC)
 - Added all hardwired connections of Computing Division (both FCC and WH)
 - Spreadsheet with IP addresses and MAC addresses available, will publish in appropriate place
 - Plan: now that the testbed is functional, looking into upgrading to 100 Gbps
 - Next step: talk to Tom Lehman for IPv6 off-site routing, right now it is part of the FNAL IPv6 block (works, ESnet intervention was necessary)
 - Old PerfSonar box was reconfigured and will be moved into the testbed (one of the two, need to make decision)
- Machines and software
 - Waiting to get a hold of some people for documentation/discussion
 - Should try to log in to machines
 - Ryan: all nodes are accessible
 - repurposed , accessible from their T1 names
 - Put in task for Dave Fagan: assign testbed IPs → this is coming
 - 3 systems are intended DTNS and/or control systems
 - 4th is the NVME system → this is 100 Gbps connected
 - Dmitry seems to have lost access to the NVME machine → Dave will follow up on this, .k5login needs to be updated
 - Name is cmsense4
- Oli:
 - Mail Justas and Tom Lehman to get going (caveat is SC22)

221118

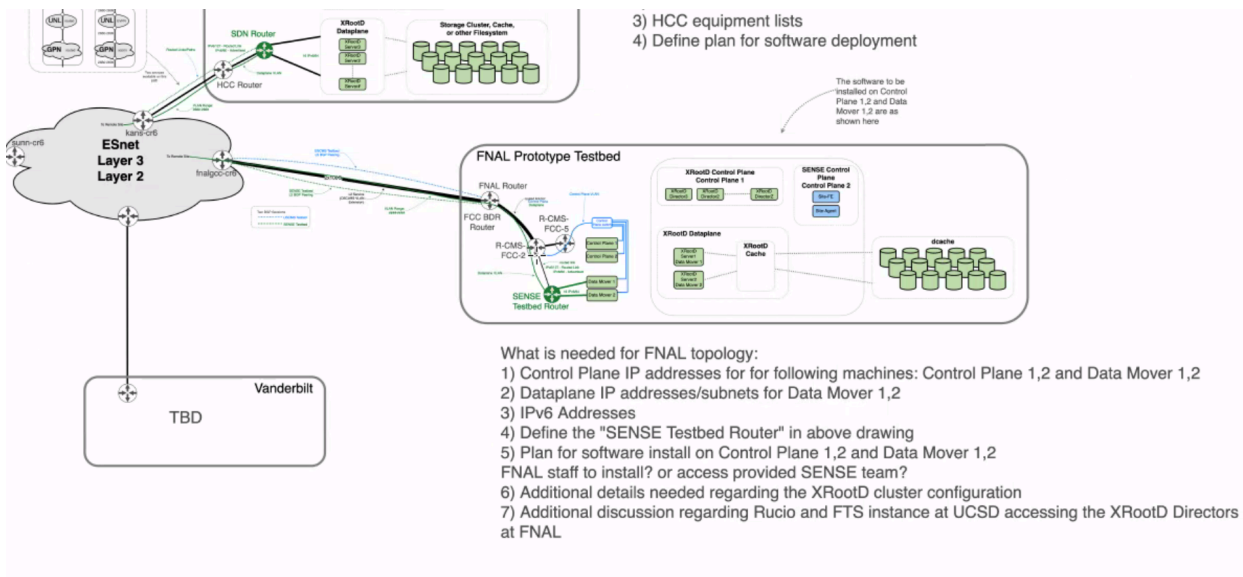
- Networking update
 - 10 Gpbs is established
 - Talking about 100 Gpbs
 - Port availability
 - LR optics: long range optics: 10 km
 - SR optics: short range optics: 10m
 - Co-location with SR, anywhere on the floor with LR
 - Working on PerfSONAR
 - On general part of the testbed
 - Perfsonartest1 → repurposed node, dedicated for testbed
 - 10 Gbps
 - ESnet setup all done
- Software installation update
 - Combination of Chih-Hao and Dmitry will do the software installation
- Justas:
 - How do we proceed with the installation? Are you ok with giving me access - or Fermilab requires all installation to be done by Fermilab employees? I see a huge benefit if I am not installing it - so we can prepare a list of questions for other sites and improve documentation. → yes, he does not get access to install software
 - For Start - for SiteRM installation, we will need to know the following:
 - 1. What is the router NOS (We might not support that yet, and we need to develop support). It will be easier if I have access (I hope it is possible). Is the Router separated - or do you plan to separate it even more (with Tofino switch you have)? What is the endpoint (port/switch) facing ESNet?
 - Not in separate devices, part of the infrastructure
 - Clarify what is meant with Tofino switch. We have an EdgeCore Tofino switch envisioned to go into the testbed, but not at the access layer. Justas, why do you ask about this?
 - Purpose (from Andrew): source-based routing using this kind of switch
 - 2. For each server (agent or any place xrootd will be running), we need to know: the hostname, IP, macvlans, controlled interface max speed, and split IPv6 ranges assigned to each macvlans.
 - Lot of that information in the spreadsheet
 - Single IPv6 subnet deployed so far
 - Close to have that info, need to work with Justas
- questions/clarifications
 - Need to clarify with Justas what we (FNAL) allow on the network side as well
 - Need to clarify the question about Tofino switch
- Time schedule
 - Waiting for SC to be over

- Then we have the holidays

221202

- Network Update

- Collecting IP information for Justas, completing the spreadsheet today (Phil), includes all IP information and VLAN information
-



- 10 Gbps testbed at FNAL
 - Two VRFs and corresponding VLANs
 - For more than the SENSE/Rucio tests, for example also for NOTED
- SENSE/Rucio testbed RM would like to be able to control a switch in our testbed network
 - Because our testbed is part of the production infrastructure at FNAL, we are adding an additional switch for the additional 100 Gbps interfaces on the DTNs
- 100 Gbps testbed at FNAL
 - In parallel to 10 Gbps testbed
- DTNs will have both
 - Dual-10-Gbps
 - Dual-100-Gbps
- But we should proceed with the 10 Gbps testbed, while in parallel we get the 100 Gbps setup
- Software to be installed on
 - DTNs → Chih-Hao coordinating with Dmitry
 - Connect DTNs to storage
 - Need to ask what to install → to Chih-Hao and Justas (Dave will ask next Monday)
 - Storage (NVME) → Dmitry coordinating with Chih-Hao

- Resource Manager (RM)
 - Expect Justas to help us, give instructions
 - Not resolved who on the FNAL side will help with this → Chih-Hao and Justas
 - Only custom software, configuring the DTNs and chose the VLAN

221216

- Network setup
 - 10 Gbps
 - No loose ends from network perspective
 - Provisioned networks as planned, tested connectivity to T1
 - Maybe changing the node names, they are generic right now, to something that indicates DTNs, controllers, etc. → **Dave need to say if this is changes in DNS if this has negative effects → Dave: no, nothing negative**
 - 100 Gbps
 - Need to investigate the programmable router, how would the testbed communicate with this router
 - What is going to program the router, what network its going to be on, what language is going to be used, how much access needs to be granted → to be discussed with Justas
 - Have not closed the loop about hardware
 - Instead of going with most performance cards, go with what the T1 is usually using
 - NVME has 100, but not sure if 2x100
 - Dmitry needs back his login to the NVME node!!!
 - Depends on a separate task outside of the testbed
 - Software update has to happen on networking side, the border has to be updated and another 100 Gbps card to be added
 - Dave (Fagan and Mason): old 10 Gpbs nodes cannot handle 100 Gbps, we're pulling additional newer server nodes to do this
- Software
 - Wednesday, Justas promised to reply to Dave's email about software installation
 - Has not happen yet 😞
 - Justas thinks easiest to use containers (in a setup that does not need remote administration)
 - Question: is this something we do in practice for the T1?
 - **Follow up with Dave**
 - Discussion about NVME node
 - Was moved to SSI
 - Is this correct, or reassign the zone to DMS
 - Chih-Hao is following up

- Cancel meeting on the 30th.
 - Happy Holidays!

230106

- Network
 - 10 Gbps systems in place
 - Need to increase the number of IPv6 subnets for SENSE from 1 to 6
 - Need to decide on the address block
 - 100 Gbps
 - DaveFagan is adding two systems capable of 100 Gbps, ordered NICs for them
 - Both will be kept in parallel
 - 100 Gbps Network path needs some reconfiguration at the border
- Software
 - Waiting for meeting with Justas to get going
 - Filling the doodle poll
 - Alan and Ryan need doodle poll link, Andrew as well

230120

- Network
 - 1st stage: Working on IPv6 subnets following requirements list
 - Investigating SENSE IP addresses and defaults one that they use, checking that this is ok for our setup
 - Many IPv4 and IPv6 addresses were added
 - *Want to make sure that there is no conflict between private addresses SENSE is using and we are using on site*
 - 2nd stage: Router that SENSE can manipulate
 - 100 Gbps hardware available
 - Licensing? Nexus hardware routing does not come by default, about \$4-5k, then figure out where it comes from
 - Another option: Arista testing hardware is arriving, not sure if it arrives in time to use it in the testbed → something to look at in the future (licensing would not be an issue)
 - Andrey and team are looking beyond Cisco for the T1
 - Deliverable: addressing table for SENSE
 - Still figuring out the details before publishing the table
 - 100 Gbps systems for testbed
 - Fagan got the network cards, will install the machines
- Software
 - No progress, still waiting for the hardware
 - 10 Gbps systems are not ready, Fagan still doing work

- Doubts that 10 Gbps on some of them is not going to work
- Dave and Chih-Hao are following up
- Instructions of what to install are available
- Weekly meeting now with Justas


230203

- Network
 - 1st stage: Working on IPv6 subnets following requirements list
 - Moving things around, also to the new switch with controls the routing manipulation from Justas (NEXUS)
 - Justas wants to look at how to interface with the NEXUS switch for his control logic
 - Need to figure out what kind of authentication Justas uses (VPN, directly over SENSE network) → ask on Wednesday
 - 2nd stage: Router that SENSE can manipulate
 - NEXUS
 - ESnet asking for IPv6 addresses, waiting for
 - Control is still in the CMS address space, needs to be moved to the testbed address space, then we send the list
- Software/Hardware
 - Not a lot of progress
 - Changed hardware to new platforms with 100 Gbps
 - We ran out of 100 Gbps slots, needed to find a new home
 - Not complete yet
 - Move started yesterday
 - 3 new machines are somewhere, older machines are somewhere else, we will move everything into the test rack
 - Go ahead to move them, Dave and Chih-Hao are waiting
 - NB. CMSSENSE1 will be part of the testbed, but not as a data mover
 - Installation instructions available, but not have not yet done anything on the machines
 - Before installation, make sure that the new interfaces are properly configured (control and SENSE networks) and working → Alan
 - CMSSENSE4 (nvme node)
 - Unconfigured state in puppet
 - Date partitions are not there, needs to be fixed → Fagan or TimS have to do that, Dave will ask
 - Identical node (storagedev201) has the right configuration
 - General
 - Hardware is now all 100 Gbps (except control)
 - But the routing to the border is currently 10 Gbps, that needs to be updated later
 - Question: the connection between the CMS network and the testbed network

- For puppet, do we need to configure the machines before we change the IPs to the test network → Phil and Dave take the lead here to discuss and finalize
 - Testbed in separate routing domain that the production T1 network
 - For puppet, we need routing and firewall rules to go from one domain to the other

230217

- Network



-
- Moved hardware into new rack
- Links are in place except the 100 Gbps connecting from border to fcc2-2, which should be in place shortly, its 10 Gbps for now
- And the 100 Gbps from Nexus to data movers are either in place or will be in place shortly
- Address change still needs to be done (in blue)
 - Blue: test network, will be also used for other tests (like Noted)
 - Green: IPv6 subnets used for testbed
- ESnet asked for IPv6 address ranges → has to wait for the blue reconfiguration
- Did we figure out how Justas will authenticate to the nexus? No, has to be coordinated between Ryan/Alan and Justas at a Wednesday meeting
- Software/Hardware
 - Data Mover 1 and 2 are new systems and have mellanox 100 Gbps cards
 - Control-Plane: cmssense1
 - NVME: cmssense4
 - Data mover 1: cmssense2
 - Data mover 2: cmssense3
 - All changes have been implemented to have the nodes in place
 - No further changes yet: next step is to install software together with Justas
 - Puppet question: need change so that systems can see FNAL puppet through current networking: Ryan/Alan have to do this
 - Routing between different VRFs need to be verified by Ryan/Alan (network services)

- Phil will contact Ryan/Alan and then one of Alan/Ryan will contact Chih-Hao

230303

- Networking making progress on getting reconfigured. Updating addresses and as of this morning getting names (accessible on FNAL network until now only by address)
- Ryan got in ticket to rename nodes, cmssense01, 04 done.
- Chih-Hao, Dmitry able to log in.
- Longer term from Ryan to Justas on node addressing (Phil will follow up/bring up in Weds meeting)
- Waiting on Dave F for config change on cmssense2,3
- cmssense4 (nvme) has wrong partitioning to be useful. Dave M following up

230317

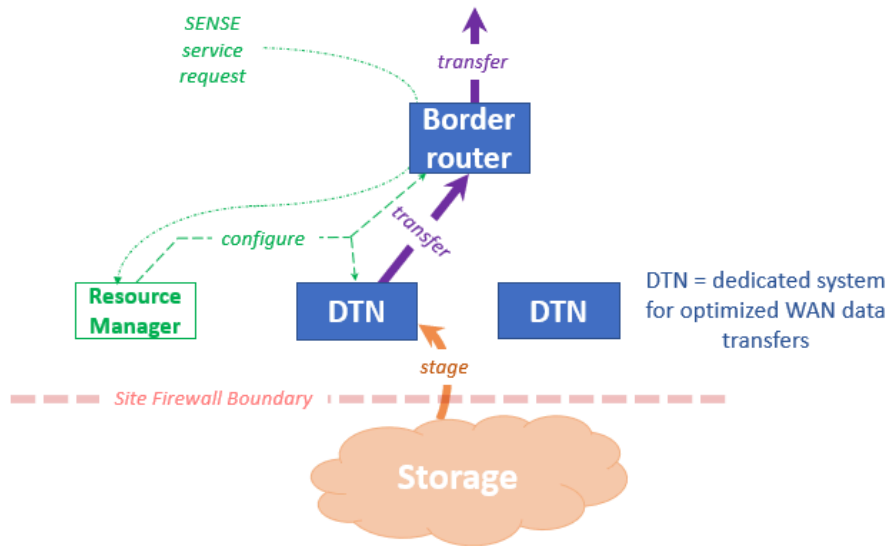
- Justas was able to access to the local devices, needs local credentials on the router, etc.
 - Ryan will work with Justas offline
- Tom: need to advertise the routes
- General setup is ready
- 10 Gbps link still in the path → fine for testing
- Current status of cmssense1-4
 - Ssh access to 4 nodes works
 - Chih-Hao does not see the configuration in Justas' github → asked Justas via email
 - Justas said that docker was not installed → Chih-Hao: we put podman on it, need to explain to Justas by Chih-Hao
 - Machines are managed the same way as the cmsdms nodes are managed
 - Glenn's group does underlying system administration
 - Chih-Hao and Dmitry will handle the application part of puppet
 - Chih-Hao can access puppet but it is locked for unknown reasons, will check with Fagan
 - Cmssense4 has wrong partitioning → ticket is open
 - Partition table from Slack was a different machine
 - Dmitry needs to define how he wants the machine → Dave is checking with Dmitry
- Glenn's meeting on Wednesday was used to clarify who manages what
 - Clarified also how much Glenn's group should watch monitoring → currently not at all
- Network
 - Work on adapting the networking setup to the way Tom/ESnet needs it
 - From our side, everything is working

- Justas was able to login to the router
- Still have the 10 Gbps in the path, no target date yet to upgrade to 100 Gbps
- Software
 - Justas and Chih-Hao finished the configuration files for FNAL (github)
 - Configured cmssense1 as controlplane/frontend
 - Firewall is setup
 - Started installation process and started service
 - Something is running!
 - Will continue with cmssense2 and cmssense3
 - Need to figure out how to configure and run the xrootd service with Justas
 - Do we have a plan how to deploy cmssense4
 - Install dcache
 - Oli learns something
 - Cmssense2 and cmssense3 essentially are redirectors
 - And cmssense4 is the xrootd server with attached storage
 - Secret implementation? Xrootd redirect through different networks
 - Data has to flow through cmssense2 and cmssense3 because they only see the IPv6 networks, cmssense4 does not see the IPv6 networks
 - DTN model!
 - Dmitry will install cmssense4
 - Then Chih-Hao and Dmitry will have to make the whole chain work

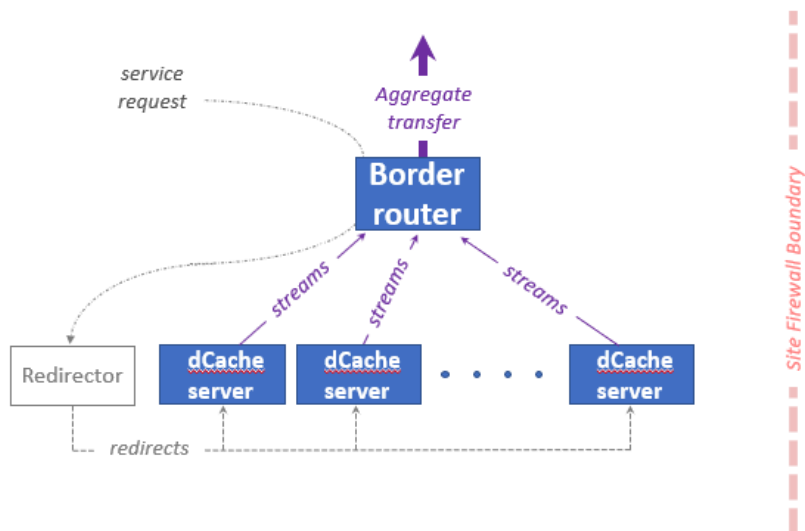
230414

- Network
 - No change in network status; 10GE bottleneck upgrade still needs to be done
 - Discussion on differences between ScienceDMZ/SENSE model vs dCache model

ScienceDMZ Model



dCache Transfer Model



- Software
 - CMSSENSE2/3 are configured as systems, need dCache namespace installed
 - DCache installed on CMSSENSE4
 - DCache doors to be installed on cmssense2/3
 - Need to check on whether cmssense4 has dual 100GE NIC
- Action items:
 - We are “close” to putting some data on cmssense4 and attempting to move it around

230428


- Network
 - Nothing missing for testing
 - 10 Gbps part no eta yet, because production border
 - Cmssense2
 - Network interface
 - One with dual stack configuration
 - One for sense data traffic (IPv6-only)
 - Default gateway should be defaulting to the dual stack
 - netstat -rn
 - Chih-Hao tested default interface
 - When testing xrootd on cmssense2 and cmssense3
 - cmssense2 every interface was slow
 - Restart fixed that
 - Didn't test the sense interface(s) on cmssense2(3), this is going to be controlled and tested by Justas' software
- Software
 - Installed everything according to instructions
 - Justas didn't reply yet
 - Now Justas and Chih-Hao have to test
- Action items:
 - Chih-Hao waiting for Justas to start testing
 - Dave will put some data on cmssense4, will create some FTS jobs

230512

- <Meeting canceled due to CHEP>

230526

- Network
 - Working on two new nodes for Noted
 - Next step: found the devices, working on the network
 - Ryan working on 10 Gbps path
 - Got complicated because of the delays in delivery of the new border routers
 - Doing inventory of 100 Gbps on old border, maybe we have to ask ESnet to add a couple more 100 Gbps because the 400 Gbps border router is delayed
 - Should be coming soon
 - Will troubleshoot VLAN connectivity issues that might appear during Chih-Hao's and Justas' testing

- Software
 - Diego from UCSD wanted to use the FNAL FTS to do tests
 - Justas asked Chih-Hao and was told cmsftssrv3
 - Have todo list from shared document:
 -  Questions for a Site (this applies to each endhost you want to add t...
 - Working through task list
 - Probably earliest 2 weeks from making first tests (because Justas on vacation the next two weeks)
 - Dave todo: find datasets and put them into cmssense4 (nvme)
 - First order: SAM test dataset
 - Cmssense1
 - Old enough that it would not do Linux 8 and later
 - We might come up with another machine to come up with a new one
 - Cmssense1 is the orchestrator, not a lot of load
- Will add Noted to this meeting!

230609

SENSE:

- Network
 - Upgrade to 100GE on the current 10GE bottleneck is in progress and should be completed next week
 - Networking is looking into implementing virtual interfaces on the border router & the CMS T1 core router to facilitate troubleshooting at layer-2
 - Would prefer to do this in coordination with ESnet and other SENSE sites so that network troubleshooting can be done across network boundaries
- Software
 - Chih_Hao is working on the SENSE site task list & end system configuration with Justas
 - Chih-Hao is requesting DNS entries for the dynamic IPv6 addresses on the SENSE-side interfaces of CMSSENSE2 & CMSSENSE3 (the two SENSE DTNs), using the following convention:
 - Need DNS Entries for all these new IPs (as example):
 - 2620:6a:0:2841::2 cmssense2-origin-2841-1.fnal.gov
 - 2620:6a:0:2841::3 cmssense3-origin-2841-2.fnal.gov

 - 2620:6a:0:2842::2 cmssense2-origin-2842-1.fnal.gov
 - 2620:6a:0:2842::3 cmssense3-origin-2842-2.fnal.gov

 - 2620:6a:0:2843::2 cmssense2-origin-2843-1.fnal.gov
 - 2620:6a:0:2843::3 cmssense3-origin-2843-2.fnal.gov

- Chih-Hao can not reach default gateway using these addresses from CMSSENSE2 & CMSSENSE3 respectively
 - Network services will check to see if the default gateway can be reached using the statically assigned IPv6 address on those interfaces
- Remaining system configuration items:
 - DNS entries (noted above...)
 - Host certificates with alternative names
 - Xrootd configuration
 - Debugging
- Other remaining items from previous meetings that we neglected to cover in this meeting:
 - Datasets on cmsense4 (nvme)
 - SAM test dataset?
 - CMSSENSE1 can't be upgraded to Linux 8; need to decide whether to replace with newer system

NOTED:

- Need to do a NOTED overview presentation (next meeting...)
 - Will be deploying a separate virtual network for NOTED, with separate end systems and a separate VRF on the testbed Nexus switch
- Expect to keep effort focused on SENSE until data movement is working before putting effort into configuring NOTED system(s):
 - But networking is moving forward on getting NOTED systems connected now

MISC:

- From CMS Management Meeting at CERN (Andrew) - there will be pre-DC24 testing in the fall
 - On a CMS-wide scale
 - Challenge will be a network trial
- Andrew is requesting that routing configuration requirements/needs of SENSE be detailed to the Vanderbilt networking people:
 - Suggestion is to use the weekly Wednesday afternoon (3:00pm CT) SENSE deployment meeting to seek this information from Justas or Tom

230623

SENSE

- Network
 - Upgrade to 100 Gbps was completed on Tuesday
 - Networking is looking into implementing virtual interfaces on the border router & the CMS T1 core router to facilitate troubleshooting at layer-2
 - Would prefer to do this in coordination with ESnet and other SENSE sites so that network troubleshooting can be done across network boundaries
 - Alan will look into this, need to coordinate with Ryan

- Default offsite path for control network (still trying to verify where it is going)
 - Asif: checked, goes through FNAL-FCC router
 - Still need to check the reverse path
 - Explanation: default path for control network should use the default path of the T1
 - Want it to go through GCC, it seems to go through FCC, need to change some things
- Software
 - Still need DNS names for every host-system/vlan combination on the SENSE side
 - Alan is looking into it
 - Question: why is 2840 block not used
 - Chih-Hao can not reach default gateway using these addresses from CMSSENSE2 & CMSSENSE3 respectively
 - This is actually a SENSE RM issue
 - Layer 2 is more complicated
 - That's the reason why we want the virtual interface on the border router to check
 - Need to clarify with Chih-Hao, this is a little bit unclear where the problem is
 - Otherwise Chih-Hao is on vacation

Noted

- Planned presentation is postponed, we need to figure out how to handle this on the FNAL side
- We start on this after SENSE is established
- Action item for Phil to talk with Eduardo Martelli @ CERN

230706

SENSE/Rucio

- Network
 - Need to setup the gateway for the 3 VLANs on cmssense1 and 2
 - Discussing with Phil and Alan how to setup the gateways
 - Tom: need to setup on all FCC routers the 3600 VLANs → some testing?
 - FCC 1 and 2, and the ESnet router
 - We have missing VLANs from 3616-3619
 - Will be added
 - Then complete: 3601-3619 → 20 VLANs
- Software
 - Almost ready to test
 - Need to config FTS3 → will be done today
 - Asif: please include him in the testing and explain software install
 - No requests for the moment

Noted

-

- Plan: next week start discussing with Justas how to start the tests

230721

SENSE/Rucio

- Network (Asif and Alan)
 - Path testing is going on
 - Finished successful testing of SENSE path from the testbed router to outside on vlan 3600
 - Will be validating the paths for vlans 3601-3619 once additional virtual interfaces for testing are configured
 - One issue remaining is default gateway for MACvlans on DTNs; we are working on it
 - More technical detail need to be discuss with Justus about default gateway
- Software (Chih Hao)
 - Waiting on resolution of the DTN default gateway configuration issue by networking
 - Backup FTS service is configured for use by SENSE/RUCIO
 - Not clear whether actual data has been loaded on to the NVME server yet; need to check on this

230804

SENSE/Rucio

- Network (Alan, Asif)
 - Working on tasks from Tom/Justas
 - Deleted VLANs
 - Adding networks for testing
 - Nothing outstanding
 - Testing end-to-end path creation for the VLANs
 - Solved the default gateway problem
 - Previously tested only 2 VLANs, now all 20 VLANs are added to the border routers
 - VLANs have been deleted from the SENSE testbed router → will be controlled and configured by the site rm
 - Tom/Justas are working on ansible module for CISCO to control the SENSE testbed router
 - Preeti/Asif are testing all 20 VLANs end-to-end
- Software (Chih-Hao, Asif)
 - Not much has been changed

- Configured FTS3 to allow maximum 2000 transfers between sites
 - Test instance → FNAL's standby server
- Waiting for Justas' network work to complete, then continue configure the software
 - Want to test FTS as well
- Question: is xrootd installation completed?
 - Yes, it's installed and running, might need minor configuration changes
 - Data is not on the system yet
- Also need to figure out how to interact with the Rucio test instance for SENSE

Notes:

- Ask about data and Rucio, how to get Rucio data on the test system and use it with Rucio → Oli will follow up offline

230818

- SENSE/Rucio
 - Network
 - Testing all vlans
 - Problem on orchestrator, when deleting the vlan, OSCAR circuit is not removed
 - Need ESnet people to remove circuit, not yet deleted
 - Software
 - Nothing changed since last meeting
 - How can be place data on our systems
 - Discussion with Diego → Asif is following up
 - We need to create directory structure following CMS namespace
 - Diego need to set up an endpoint to populate the test Rucio instance
 - PFN piece, LFN piece
 - Endpoint URL/name
 - LFN Logical File Name
 - PFN Physical File Name
 - Authentication/Authorization: are we using the CMS production certificate?
 - Next step: Diego copies a dataset
 - Also: investigate the Rucio load test functionality
 - Load test: creates random files, transfers them, deletes them again
 - Old load test: dial to produce a specified bandwidth
 - Diego needs to know
- Noted
 - Discussion with Eduardo, Phil sending clarification

- Eduardo replied: needs layer 2 part from CERN to FNAL, request to Tom is out, Tom is working on the circuit
- Also talking about FTS details
 - Noted is using the CERN FTS and is using the FTS logs to decide where to shift traffic
 - They need log access to whatever FTS we are using
- Creating the testbed, we are figuring out what we have to do on our (FNAL) end
 - Plan is to create between CERN, KIT and FNAL
- Eduardo asks that we are not using different paths than LHCOPN and LHCONE
- Following up with Phil after he is back from vacation

230901

- SENSE/Rucio
 - DC24 blueprint to define US goals for data challenge
 - Testbed is almost complete
 - Diego is working with Chih-Hao on Rucio on the FNAL testbed
 - Dmitry is coordinating now because of WEBdav access to dCache
 - ESnet/OSCAR circuit still has a problem, ESnet will fix in 2 weeks → then we can test VLANs
 - Possibly a software issue, need to be fixed and deployed
 - Nexus plugin for SiteRM → Justas is working on it, will be completed soon
- Noted
 - Separate system (3rd system)
 - Ran out of 100 Gbps ports, talking to networking services
 - Replacing switch with new switch, same nexus switch family, no impact on Justas → switch is on site, just needs to be deployed
 - Maybe switching dtn2 to noted while we are waiting for the new switch
 - Tom (ESnet) is setting up a LHCONE like circuit for SC23
 - Seems ok for Eduardo, opposite to what we said last meeting
 - Noted requires FTS logs from testbed
 - Oli does not think they need access to the FNAL FTS logs, CERN will use CERN FTS instance(s)
- Network in general
 - Traffic between the buildings → crossing 1 Tbps
 - Cross checking monitoring information

230915

- SENSE/Rucio
 - Network part of testbed is functional for SENSE
 - Dimitry working with UCSD on actual data movement via XrootD

- 3 XrootD door configured; XrootD is working; HTTP not yet
 - Certificates have been configured on CMSSENSE1
 - Will need Rucio-cataloged data; intra-FNAL Rucio transfers to test systems seems to be preferred means of accomplishing that
 - After that, they are expected to work on load testing
 - Asif/Preete testing site Resource Manager for deployment at other US T2s
- NOTED:
 - NOTED DTN system in place:
 - Checking on O/S version for being up-to-date:
 - Network Services is configuring testbed switch to support NOTED DTN data transfer
 - Also getting address assignments for DTN; will try to get by with IPv6 only
 - DTN will start out using 10GE connections, but migrate to 100GE as soon as new testbed switch is in place
 - Replacement testbed switch is in hand; Network Services planning started on forklift upgrade of existing switch:
 - Don't have a timeline yet
 - Detailed technical discussions with CERN planned for next week on what specifically will be required on our end
- Network in general:
 - 2x100GE has been added to FCC border router WAN connectivity (LHCONE & LHCOPN paths), now at 4x100GE (400GE total)
 - Should be OK to run testbed WAN traffic at full 100GE without impacting production traffic

230929

- SENSE/Rucio
 - Between UCSD and Caltech being tested
 - FNAL side
 - Configuring the siteRM, making changes to configs to dynamically configure our router
 - When done, it can dynamically configure the testbed router and configure the end-to-end path
 - Next week deploy the container
 - Nebraska
 - Deployed SiteRM already
 - ESnet (Xi) was testing on the OSCAR side, still some issues, OSCAR people are still upgrading the software

- Next Wednesday meeting will ask for update
 - Status of FTS transfer transfers
 - HTTP transfers were needed, deployed the doors, did the proxy configuration on cmssense1 (single endpoint to be used to HTTP)
 - Diego reported transfers are working
 - Asif will talk to Diego to transfer a larger dataset
 - Asif needs to learn how to use and monitor Rucio
 - Oli will find some training for everyone to learn Rucio
 - For FTS, Chih-Hao will train Asif
 - Diego requested OSG certificate update in cmssense2, 3 from Chih-Hao, this should be in place, Asif will check
- NOTED
 - System is in place, OS is done, system is accessible to Chih-Hao and Asif
 - Working on configuring the remaining services
 - Got the IPv4 and IPv6 addresses, requested the network team to configure the SENSE part
 - Need to test the routing from FNAL to CERN (end-to-end), figure out if we use LHCOPN or LHCONE
 - Requested another DNS entry from NOTED side
 - Currently working on 10 Gbps, when switch is switched we go to 100 Gbps
 - Next is also to test FTS transfer into the NOTED endpoint
 - Dmitry will configure an additional webdav endpoint on cmssense1 for cmsnoted1
 - Also need to configure xrootd on cmsnoted1 → Chih-Hao and Asif
 - Cmsnoted1 needs a certificate
- Network in general
 - Switch replacement is being worked on
 - Next couple of weeks it will be replaced
 - Border at 4x100GE, testbed is at 100GE
 - Will not change till the border moves to 400GE
 - Will stay with this for 6 months

231013

- SENSE/Rucio
 - Testing is still in progress – Sense/Rucio team is testing the DMM functionality which is being developed at UCSD
 - FNAL side
 - SiteRM locally tested and working with Cisco switch
 - But one problem when running container with host network and ports configuration and accessible issue from outside
 - Decided to not run with host network and use IPv4 address to manage the testbed switch

- Asked the network team to assign IPv4 address on switch for smooth working of siteRM container.
 - Nebraska
 - SiteRM deployed and DTNs also part of SENSE orchestrator now
 - New ESnet OSCAR production version still in progress to deploy – not deployed completely
 - Asif and Diego tested the larger dataset transfer from UCSD to FNAL
 - Tested by using cern-pilot FTS and FNAL's FTS instances.
 - Used Rucio instance deployed at UCSD
 - Learn about Rucio monitoring on MonIT
- NOTED
 - Cmsnoted1 system is ready and accessible
 - Chih-Hao & Dmitry configured the endpoints to test with CERN
 - Routing configuration is in progress to configure the LHCONE path and working on a quick database to advertise noted system IP addresses.
 - Dave Fagon sent an email about the port issue on NOTED server: Chih-Hao will take care of that because it is allowed as an endpoint for NOTED.
 - Chih-Hao will perform a data transfer test between FNAL–CERN, asif share the details of endpoints.
- General Networking:
 - Alan set up the IPv4 address needed for the testbed switch.

231027

- SENSE/Rucio
 - Testing between T1 and UCSD. Small issue regarding site RM; Justas is working on this.
 - New switch being tested with DTNs; things seem to be working well, same as the old switch. Switch is now being managed with IPv4, a current requirement with the site RM
 - Agents on DTN still need to be installed. Asif is doing this. Minor issue with certificates being worked on.
 - ESnet expects to complete deployment of new version of OSCARs this week. Once this is completed, we should be able to run end-to-end SENSE service testing.
- NOTED
 - NOTED server is ready with the new switch, but using the 10GE NIC. There is an issue with the 100GE Mellanox-6 NIC. Dave F. is working on this.
 - End-to-end path testing between NOTED server and CERN us working, but using the default routed path. Yan is working on making the LHCONE path the default path.
 - Tom L. (ESnet) is working on getting the OSCARs circuit to CERN in place. This is tied into the ESnet OSCARs transition.

- Carmen (CERN) has asked Asif to set up continuous transfers between FNAL & CERN.
- NOTED vLAN need to be changed from 1323 to 1335. Network Services will take care of this.
- Asif having an issue with initiating transfers from CERN (EOS). Works for Chih-Hao, but not for him. Error is authentication error on CERN's end. Dave suggests the next step is to try the same transfer from production T1 server (LPC node?). That test should help determine whether it's a system issue or an Asif credentials issue. Chih-Hao will also try a transfer test from the testbed NOTED server, using his credentials, to see if it works for him.
- General Networking
 -

231110

- SENSE/Rucio
 - Circuit between Caltech and FNAL is available, BGP peering is done
 - Need to re-architect the storage for SC23
 - Looks different than what the SENSE/Rucio testbed expects
 - Since Wednesday working on that
 - Physically reconfigure
 - Previously:
 - 2 DTNs in front of NVME system
 - But was configured as one of the DTN was configured as HA proxy, and all 3 machines were accessible
 - For SENSE/Rucio SC23 test (a.k.a Panic Mode)
 - Need to access the NVME system directly
 - Serve data from NVME system, needed a 2nd interface card to make it work
 - Yesterday completed setup
 - Today: meeting with Tom, Justas, Chih-Hao, Dmitry → install SiteRM and test at 12 PM (if everything is fine, fingers crossed 🙏)
 - After SC23: need to step back and evaluate what we actually want to test and then setup the testbed accordingly
 - Fallback-fallback, use local storage on DTNs
 - Alan needs to reserve an IP on the DNS, IP is already in use and working
- NOTED for SC23
 - Created SENSE circuit (few minutes ago, were waiting for ESnet), trying to ping now
 - Then we setup BGP peering
 - The we can test data transfer

- Once we have the circuit, we have to test if the storage setup still works while we reconfigure the SENSE/Rucio setup
- General networking
- DC24
 - After Thanksgiving and SC23, get together and discuss
 - Outcome pre-challenge
 - Expectations for DC24
 - Talk about FNAL: overhead needed on top of DC24 CMS traffic and DC24 Dune traffic and SENSE/Rucio traffic
 - Invite Dune to this meeting

231208

- SENSE/Rucio:
 - SC23 SENSE demos went well:
 - Needed to run SENSE demos directly to/from storage (cmssense4)
 - This is considered a temporary configuration, not a long term configuration
 - SENSE demos worked showing prioritized SENSE data movement
 - Post-SC23 discussions on how to configure storage in a way the mimics production
 - Planning to add a data server to the testbed; will not be an NVME system, but will be used to create testbed dCache cluster
 - New server installation is in progress (Dave F.)
 - Will need help from Dimitry & Chih-Hao to do this
 - SC23 NOTED demo had issues with BGP configuration/peering:
 - Some tests were run successfully, but not quite what we had intended
 - Still working on post-SC23 tests; this remains a work-in-progress
 - BGP configuration remains an issue
 - Asif has added additional subnets (up to 16) to the testbed systems
 - DNS entries requests are submitted; not completed yet
 - An SSI task
 - Dimitry needs to submit certificates request
 - Alan will look into deploying a breakout cable for the new switch to provide 10GE support
 - This will require replacement of the 10GE UTP NIC with a fiber 10GE NIC on cmssense1
- Nothing new to report or to be done on the testbed system end at this point
 - Chih-Hao will test data transfer capability from cmsnoted1 to CERN EOS end point
- Next Friday testbed meeting will be Jan 5

240112

- Today's agenda and DC24 plan (SENSE/Rucio)
 - <https://docs.google.com/presentation/d/1LzcZGFpgpFod-hdI0FRCx7ipMW8uvYr3/edit?usp=sharing&oid=108943909987414331647&rtpof=true&sd=true>
- SENSE/Rucio:
 - Discuss the DC24 SENSE/Rucio Plan and activities
 - Overview of current status SENSE/Rucio architecture and changes required
 - Due to issues at green path for transfer, the separation of green path VLAN and SENSE subnet VLAN is required
 - Ryan already separated the VLANs on the testbed router and Asif/Justas will test it this/next week.
 - Storage change required in testbed and Dmitry is working on it.
 - One webdavs door (cmssense2) and it will redirect to storage NVMe server (cmssense4)
 - We are trying to replicate testbed storage and software where it should be look like production (Asif)
 - In production, we have multiple xrootd (sire/regional) redirectors contact with WebDavs doors (Chih-Hao)
 - Dmitry configured new certificate with new DNS entries on all servers
 - Additional Storage nodes requested from Dave for testbed and we will deploy in parallel with cmssense4 in future (Asif)
 - Need to plan DUNE in parallel with the DC24 activity of CMS (Oli)
 - Dave will form a coordination meeting for DC24 between CMS and DUNE (Oli)
 - Also for SENSE/Rucio and DUNE in parallel.
 - CMS/USCMS DC24 plan is to achieve at least 250 Gbps for CMS and at least 80~100 Gbps for USCMS
 - Asif will use NetworkLoadGen Tool to stress the FNAL network before DC24.
 - **Oli → suggested that our upcoming meeting should be in-person for an in-depth discussion of DC24 activity and architectural discussion on future testbeds.**
 - **Asif will schedule it and send an email to FNAL's members coordinating on the SENSE/Rucio testbed.**
- NOTED
 - It was tested using FTS before and working (Asif), now coordinated with Carmen and Katy for further testing using SENSE path and LHCONe paths
 - Might be FNAL's NOTED setup will participate in DC24 (Need more discussion if required)
- General networking
 - Alan will connect server with breakout cable on new switch and remove old switch

- IP addresses and DNS entries configured on all CMSSENSE servers ([Link](#))

240126

- Today's agenda and DC24 plan (SENSE/Rucio)
 - https://docs.google.com/presentation/d/112WHmTZQFT6MctiHVVNOjdzJuukj-jOQ/edit#slide=id.g265a05fb36_0_0
- SENSE/Rucio:
 - An overview of the original SENSE/Rucio testbed plan and few changes accommodate since the start of the testbed.
 - changes highlighted in network and storage architecture as per the requirement of the SENSE/Rucio team.
 - Reconfigured the network new VLAN from sense testbed router to servers (16 IPs) (light green network)
 - SENSE VRF ended at testbed router from FNAL's border routers
 - NOTED path was also configured with a totally separate IPv6 subnet in the testbed.
 - SENSE is not using cmssense2 and cmssense3 as data movers as originally planned.
 - SENSE using storage node (cmssense4) directly for data transfer and added that node to the light green path.
 - Current status SENSE/Rucio testbed at FNAL:
 - Testbed is working as expected for DC24, but few testing are on-going related to the network.
 - Planned to use cmssense2 as webdav door and cmssense4 as the storage nodes
 - DC24 proposed plan for SENSE/Rucio:
 - Mainly on day 4 SENSE/Rucio activity is scheduled.
 - All set for this DC24 exercise, just need to check with DUNE because parallel activities are planned.
 - Post-DC24 open discussion/plans for SENSE/Rucio Testbed @FNAL:
 - Need to plan the testbed by considering the scalability of network, storage, and sense itself.
 - All agreed to replicate production like setup in testbed environment and discussion on couple of storage configuration options:
 - Multiple IPs (macvlans) on production storage servers:
 - Maybe a scalability limitation, because how to orchestrate it on thousands of storage nodes in future? (Dave)

- Webdav doors serve the data at the front-end and we can configure the multiple IPs (macvlans) as required for SENSE.
 - Need to further explore this option in the testbed. Looks feasible for production in future.
 - Need to add multiple storage nodes in testbed after DC24
 - Disk based storage server may not compatible with NVMe – I/O speed
 - Need to adjust all disk nodes or nvme nodes in future.
- Next meeting will be on 9th Feb. 2024 to discuss the testbed readiness and execution plan for DC24.

240209

- Presented:
 - https://docs.google.com/presentation/d/1gIGxczMF4u3JmpWHk5Meas6dqtXPrHMJ/edit#slide=id.g2b3665a3b81_0_0
- SENSE/Rucio:
 - DC24 preparation and current status is ready for testing
 - NVMe server is not exceeding the throughput more than 40G
 - Additional options discussed with SENSE/Rucio team:
 - Instead of two flows run the experiment with three flows:
 - Flow #1: FNAL → Caltech (~30G)
 - Flow #2: UCSD → Caltech (~50G)
 - Flow #3: UCSD → Caltech (~20G)
 - Swapping the better performing NIC (100G) from cmsense2 to cmsense4
 - Discuss the previous test results from FNAL → Caltech
- General Networking:
 - 200G cap issue testing at FNAL is still in troubleshooting stage
 - ESnet's router (Nokia) having issue with QoS setup
 - LHCOPN path was not working

240308

- Presentation:
 - <https://docs.google.com/presentation/d/1oiPIr1Kz6nVYK98UCxmHa8ZN8zdkDdhr/edit?usp=sharing&oid=108943909987414331647&rtpof=true&sd=true>
- SENSE/Rucio:
 - Consolidation of green (SENSE data movement) and blue (control plane) networks is proposed:
 - Migrates away from dual NIC complexities of current configuration

- Aligns design of testbed network with current deployment model in production T1
 - Question of which IPv6 address block (2620:6a:0:4840::/SENSE data movement or 2620:6a:0:4421/control network) to use:
 - 4421 seems preferred
 - Question on whether IPv4 support is still necessary in the testbed at all?
 - Consensus is to keep IPv4 support for now to mimic dual stack configuration of the T1
 - However, storage system itself should be capable of IPv6-only
 - Upgrade of 100GE uplink from testbed router to r-cms-fcc2-2 to 400GE should be doable
 - Being upgraded now in preparation for deployment of new site border routers to 400GE capability
 - Need to obtain 400GE optics to do this
- Misc system changes/upgrades
 - Existing systems need to be upgraded to Alma9, and more current kernel to improve throughput
 - NIC card on current NVME system (cmsense4) probably needs to be upgraded from QLogic to Mellanox
 - New NVME system will be available only on a temporary basis, and will be replaced with a new procurement, when the latter arrives
 - Probably don't want to physically move the temp NVME system over to the testbed rack
- Asif will provide a schedule for the testbed reconfiguration as soon as one is developed

240322

- Presentation:
 - <https://docs.google.com/presentation/d/1JEtQ9VEUUdOoTzPfnwNwrKU3CgruP9fQ/edit?usp=sharing&ouid=108943909987414331647&rtpof=true&sd=true>
 - https://docs.google.com/document/d/19k1lq4iUtq7BYmLV8AaSWKHRYXwjR_Ng6q5Qa0VhX88/edit
- SENSE/Rucio:
 - Recap of last meeting and discuss the major points:
 - Updating the servers, re-architecting network topology, and re-architecting storage servers
 - Upgrade activities schedule will start from 1st April because currently Tom will have SENSE/Rucio demos in OFC conference.
 - Shared the schedule of all activities need to work out, some activities can be started in parallel and some are interdependent.
 - Next week we can start work on taking backups of network configurations and servers:

- Alan will take the backup of testbed router's current configurations
 - Chih-Hao: Firewall rules need to be listed down and need to configure after updating the OS (Alma9), because Alma9 deals differently with security configurations
 - Macvlans configurations and SENSE software (SiteRM, Agent, exporter, etc.) backups will be taken by Asif/Justus.
 - Dmitry will take care of storage related backups (if required)
- Asif will discuss with Dave mason about moving new storage node to testbed rack
- Discussed the network configuration changes required on testbed (NEXUS) router
 - Making blue VRF as default and removing the green VRF from switch
 - Alan is agreed with new configuration changes but he will discuss with Ryan
 - Alan will also discuss with ryan, if any changes required on border router

240419

- SENSE/Rucio:
 - Testbed servers (cmssense[2-4]) upgraded with new OS Alma9 and ready to use
 - Cmssense1 server is old and cannot be upgraded, so we will keep it as it is for siteRM.
 - New storage server deployed at testbed rack (Cmssense5) and connected with network
 - Need to request host CERT for this server
 - Do we need to map all 16 subnets/hostnames in the CERT for this storage server? Discussion needed with Dmitry.
 - Initial testing of all newly upgraded servers done with the existing network.
 - Chih-Hao: completed the nftables setup on cmssense[2-3], however, cmssense[4-5] will be taken care of by Dmitry.
 - Networking re-architecting is in progress: Changing from multiple VRFs to single VRF
 - Ryan/Alan need more clarification about this new upgrade before start changes on the testbed router.
 - Asif will discuss with Ryan in-person today
 - Phil: New border router in placed but not in production, Ryan may wants to connect testbed with new router (But need more discussion)
 - Phil: Upgrading the testbed's 100G to 400G is also discussed with Ryan.

240503

- SENSE/Rucio:

- Did couple of meetings with Ryan about new network configuration
- He completed the Network configuration of TB router and Border routers
 - Created only single VRF (Blue side)
 - Remove the Green VRF
- Initially we had issue of testing servers to TB Router
 - Ryan updated the OS on router (NXOS 10)
- We configured 16-IPs on servers (cmssense2-4) top of Blue VRF (VLAN1312)
- Created the Macvlan on servers and tested from servers to TB router connectivity – all the macvlan on servers working fine
- **Some pending items:**
 - Cmssense4 is still connected with 10G port, we need to move this on 100G port (need swap the port on TB router) – Alan/Asif will take care of that
 - Upgrading link capacity 100G → 400G in progress – Ryan/Alan will do that
 - Generated ticket of assignment of 16-IPs to cmssense5 (new storage node)
 - Once done then I will generate new ticket for its Grid CERT
 - Need to configure the siteRM and agents config files for FNAL
 - Interfaces information changed – Need to add – Asif will do that
 - Updated the NXOS 9 → 10: I discussed with Justas and tested with the current module of SiteRM but it is not working.
 - Justas need to develop a new Module (NXOS 10) that will be compatible with TB router.
 - Justas asked about some details of TB Router NX and I will provide him then he will work on new module for siteRM

240517

- SENSE/Rucio:
 - Testbed almost ready after upgradation and completed the initial testing
 - Cmssense4 connected with 100G ports, by changing the IPs addresses
 - 16-IPs (MACVLANS) on cmssense5 assigned and configured and Grid CERT also generated.
 - New siteRM module for Cisco NEXUS version is configured successfully.
 - Node exporters and agents also configured on servers, cmssense5 is still need to configure
 - Discuss with Alan about old switch removal, but still cmssense1 is connected – need to plan the removal in future
 - NVMe disk on cmssense4 need replacement – contact with system team about it
 - Discussion on two interfaces connectivity as SENSE/Rucio group suggesting:
 - Dmitry suggested to first testing with original plan (single link)

- Phil had discussion with the storage team about future plans and production storage might need dual home servers.
- Justus wanna to add his DN for xrootd permission – Dmitry will take care of that
- Phil: 400G link upgradation at border routers in progress then testbed will also be upgraded.

240712

- Updates of SENSE/Rucio testbed:
 - All testbed machines updated to Alma9 (EL9)
 - Chih-Hoa configured the similar firewall/nftables as we had before
 - Cmssense1 machine retired — SiteRM moved to cmssense2 — waiting for new machines as cmssense1, already talked with David Fagon (Provided the cmssense1 but with 10GE copper ports)
 - Phil have 10G optical NIC which can be used for cmssense1 machine
 - Cmssense2 SiteRM having problem to access its Frontend outside, issue of nftable and docker compatibility.
 - Need to talk with Tim to check any security related solution possible (Asif will talk with tim)
 - In parallel working is Justus to figure out the hack to resolve the issue on cmssense2 for SiteRM.
 - Cmssense4 was restarted and it was noticed the IPs were not permanently configured on NIC (single interface)
 - Discussed with David Fagon to fix this issue permanently. Add in Foreman
 - Supercomputing 2024 participation — Need to ready our testbed as per production environment.
 - Networking:
 - Ryan said 400G from sense to CMS is up
 - He will migrate the TB CMS to the border connection to the new border router next week.
 - Before the start of SC24 we will have a testbed with 400G link.
 - Storage setup configuration as proxy is required now to test whatever we decided before.
 - Asif will discuss with dmitry to configure the storage nodes and proxy setup in the testbed and need a plan to complete it before start of SC24.
 - Need extensive testing of scalability and performance.
 - Just for information, a Madison student (Lael) is also using these testbed machines for network throughput testing using Perfsonar.
 - Everything will be inside docker containers.

240726

- Updates of SENSE/Rucio testbed:
 - New cmssense1 machine replace with old and ready for use
 - SiteRM software was running on cmssense2 temporary and now move to cmssense1
 - Had a problem of port forwarding with nftable and docker, it is a compatibility issue but found a hack to run the siteRM inside the docker with port 8443.
 - Cmssense4 permanent IP configuration issue fixed.
 - Preparation is ongoing for SC24 with the SENSE/Rucio team.
- 400G Network Transition status:
 - Two border routers installed at GCC & FCC but still science traffic routed with old paths including SENSE testbed traffic.
 - Alan will check and update the current status of network transition
 - New network diagram need to be shared with SENSE/Rucio team for clear understanding for routing and VLANs (3600-3619) configuration
 - Phil suggested configuring the VLANs on the new border router for data traffic.
- Storage status & proxy-based storage plan
 - First testing with direct storage nodes (cmssense [4-5]) need to be done
 - Dmitry will check the configuration of cmssense5 configuration at the storage node.
 - Next week SENSE/Rucio team will test the storage nodes and then Dmitry will setup proxy mode
 - Need extensive testing of scalability and performance with proxy mode before SC24
- System/Software/Security updates:
 - Chih-Hao will look the nftable solution for docker

240809

- Updates of SENSE/Rucio testbed:
 - Network link 400G updated
 - FCC-cr6 ESNet router configured ports (Lag-2) — 4x400 (1.6Tbps)
 - Old router is still in placed but our testbed routing is managed via New Arista router
 - OSCARS VLANs 3600-3619 currently in-placed via new Arista router
 - All VLANs are being in testing, few configurations required at software (orchestrator & SiteRM) and network (ESnet's WAN links) levels
 - Links throughput testing will be the first step after end-to-end VLANs configured
- Storage nodes testing
 - Cmssense4 & cmssense5 are our dCache machines, we have to ensure all dCache services are running on these machines
 - We have to make sure that the Cmssense5 server is working fine and dCache is properly configured? Data is available? – Asif will discuss with Dmitry
 - Once network testing done then we will start direct storage testing via SENSE/Rucio/FTS

- Dmitry will then configure the proxy setup.
- System/software/security updates – by Chih-Hao
 - Need to check the nftables compatibility with docker solution – Chih-Hao will take care of it.
 - Cmsense4 one NVMe disk was faulted, it was in warranty and Chih-Hao/Dmitry will get updates from Edwards.