ATLAS Journey Integrating into the SENSE/Rucio Priority Data Transfer Service

CHEP 2026

Aashay Arora on behalf of the Rucio/SENSE team





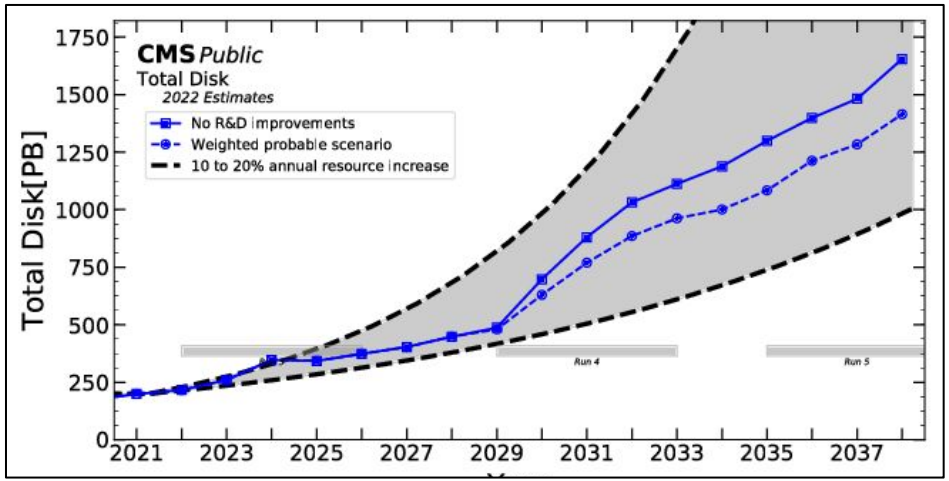









Motivation: HL-LHC



CMS **2022** Estimated Disk requirements

2025 estimations suggest we will have less disk space than expected

Less Disk => More Tape
More Tape => **More Data Movement**

Less Disk => Less Replication
Less Replication => **More Data Movement**

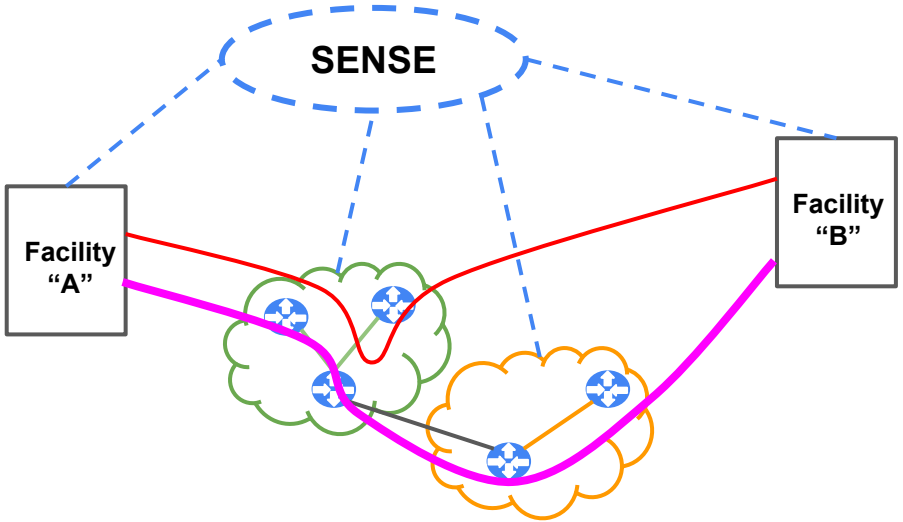
The network has never been a problem in the past. The HL era might change that.

SENSE

A Network Services Orchestration System developed by ESnet.

It allows for experiments to communicate intentions and needs to the R&D Networks in real time.

Can be used to build traffic engineered paths between Computing Facilities and treat the network as a schedulable resource.



Regular (red) vs Traffic engineered path (pink) across two Computing Facilities

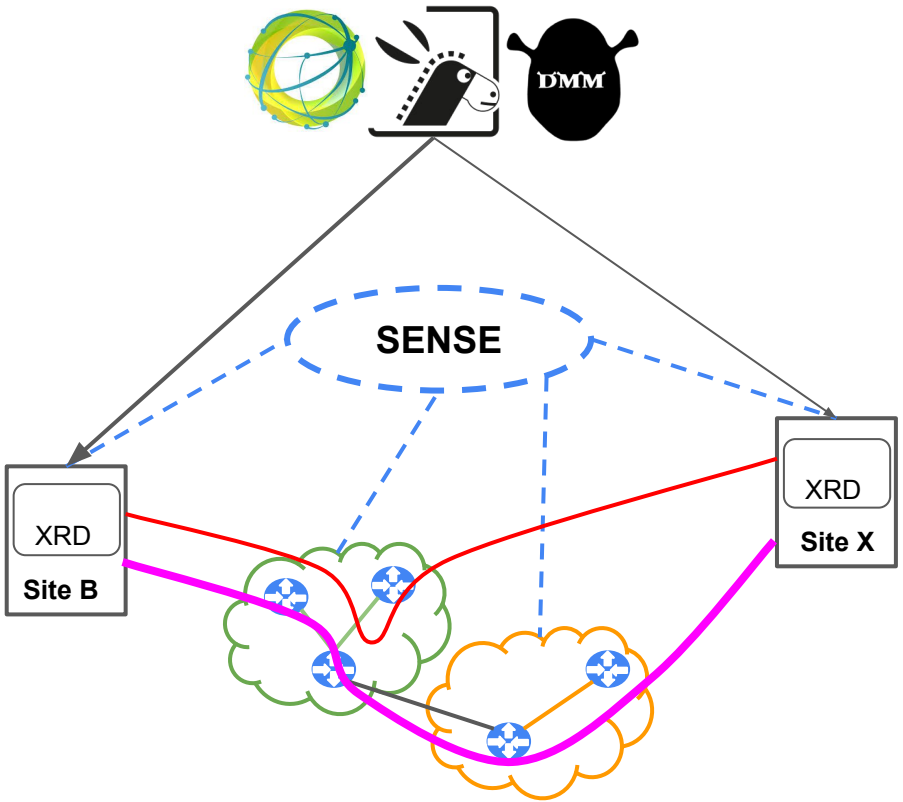
"The better the Network understand its user's needs the smarter decisions they can take"

The Rucio Use-Case

Within the WLCG world, Rucio was the first use-case selected to make use of these SENSE capabilities.

The objective of the Rucio/SENSE integration is to make Rucio capable of requesting **bandwidth-guaranteed** network circuits for its more **important and large data flows**

Bandwidth Guarantee =
Predictability &
Resilience to Traffic Variability



Network Control - Requirements

From the point of view of a site and in order to achieve the above, SENSE requires access to the following network capabilities:

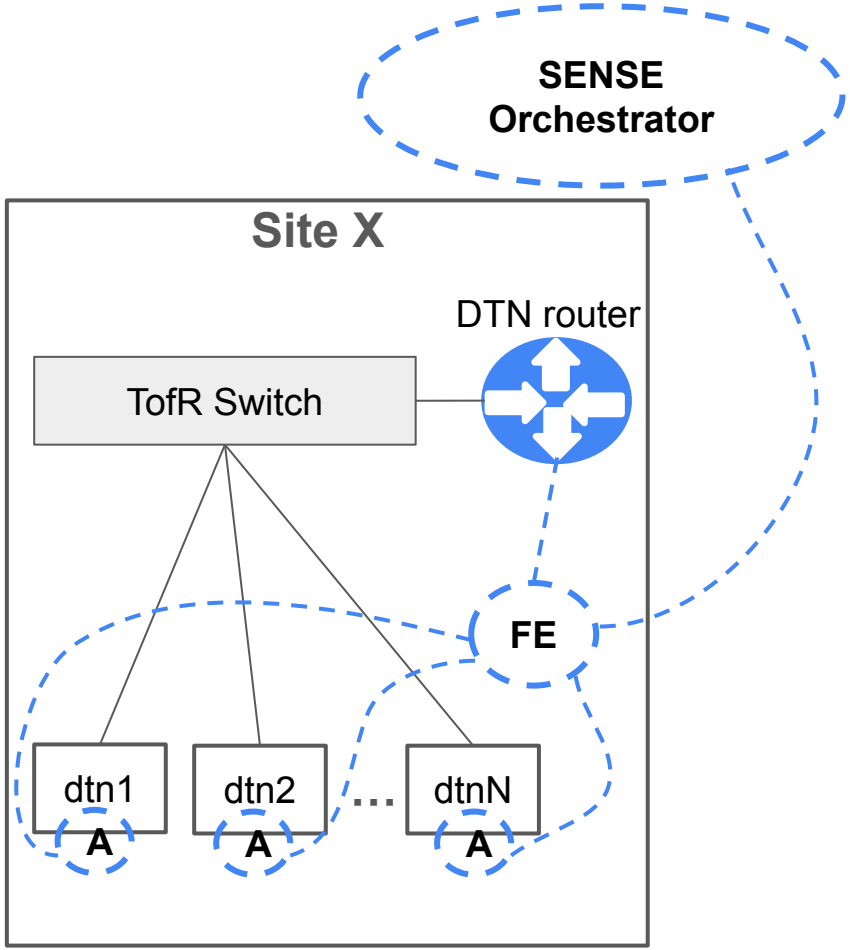
- Modify the routing of the DTNs data flows
  - Manage VLANS to create Point-2-Point L2 network paths
  - Manage BGP announcements to route desired traffic over the above L2
- Configure QoS policies on the DTNs
  - Used to provide bandwidth allocations
  - Done using the Linux Traffic Controller

These requirements are **restricted to a defined set of subnets and VLANS** which are used only for SENSE-managed data flows.

Network Control - SiteRM

At the site level, SENSE requires the deployment of 2 main Site Resource Managers (SiteRM):

1. SiteRM Frontend.
    a. Receives commands from the Orchestrator
    b. Applies changes in the DTN router
    c. Sends commands to the Agents
2. SiteRM Agent.
    a. Runs in the DTNs
    b. Execute commands on the DTN
    c. Enforces QoS policies



More info about SiteRM: https://sdn-sense.github.io/#siterm-resource-manager

Software-Router

The SENSE system is based on traffic engineering IPv6/64 subnets for the purpose of providing end-to-end priority data transfer service. SENSE is typically allowed to apply a **limited set of configurations** on the DTN router to effect these features at the site level.

When the above network access is not possible, SENSE provides an option for a software based traffic engineering approach .This solution is based on the following 3 software suites:

1. Free Range Routing (FRR)
2. Vector Packet Processing (VPP)
3. Data Plane Development Kit (DPDK)

Software-Router

Generally speaking, network routing is composed of 2 elements: control-plane and the data-plane. The former is in charge of the logic and the later of the execution.

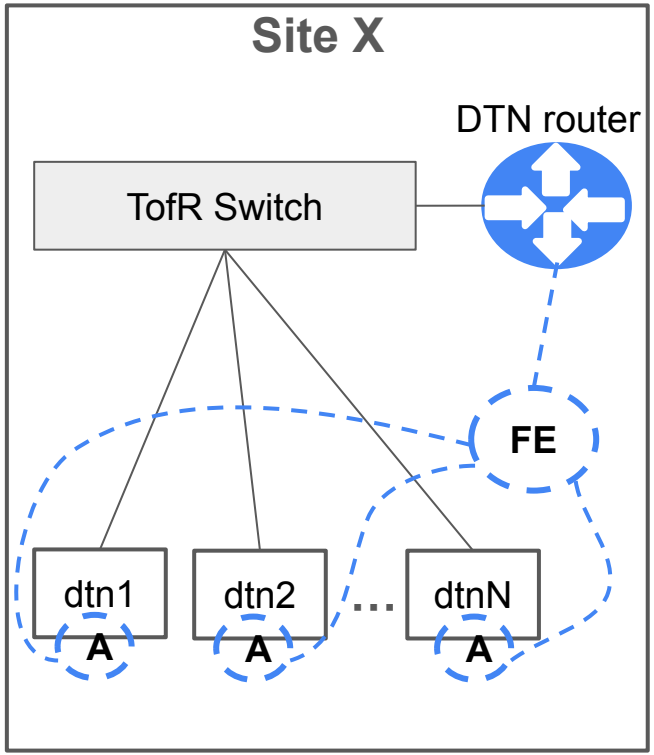
FRR implements the control-plane logic. When used alone, it offloads the data-plane logic to the Linux Kernel which process packets one-by-one.

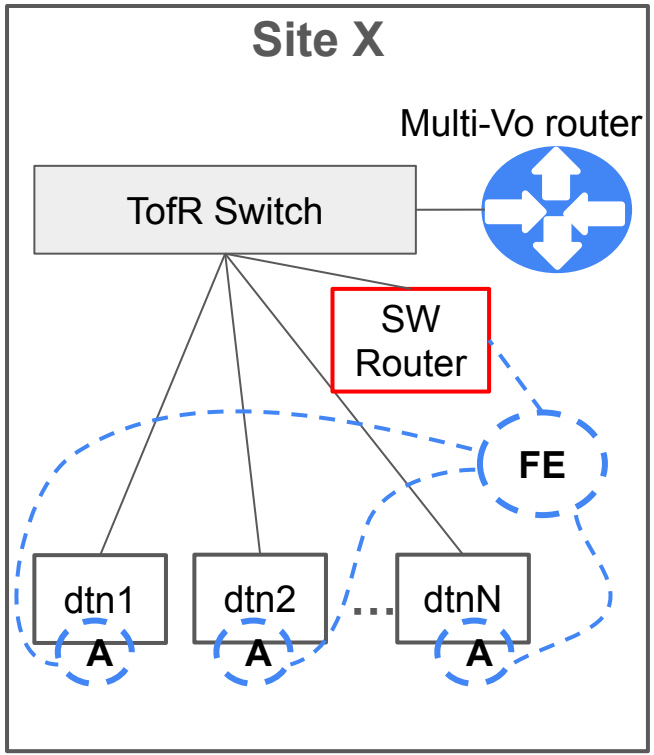
When dealing with high-throughput the above becomes a bottleneck. This is where **VPP/DPDK** comes into play.

VPP process packets in vectors which makes the process faster. Moreover it can be integrated with **DPDK** allowing for directly control of the NIC, bypassing the Linux kernel network stack hence eliminating the interrupt overhead

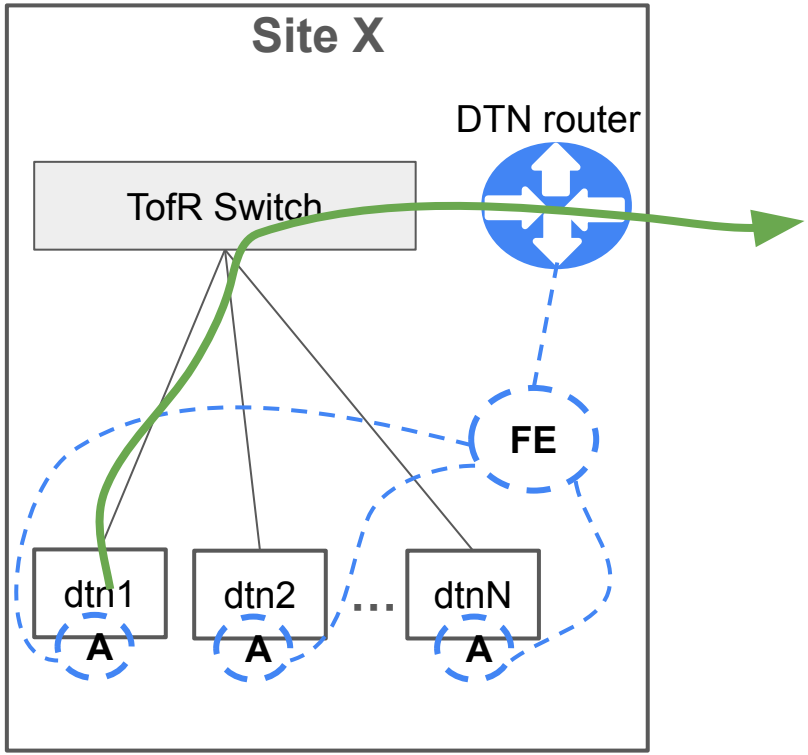
More information here: https://sdn-sense.github.io/optional-install/software-based-router-dpdk/

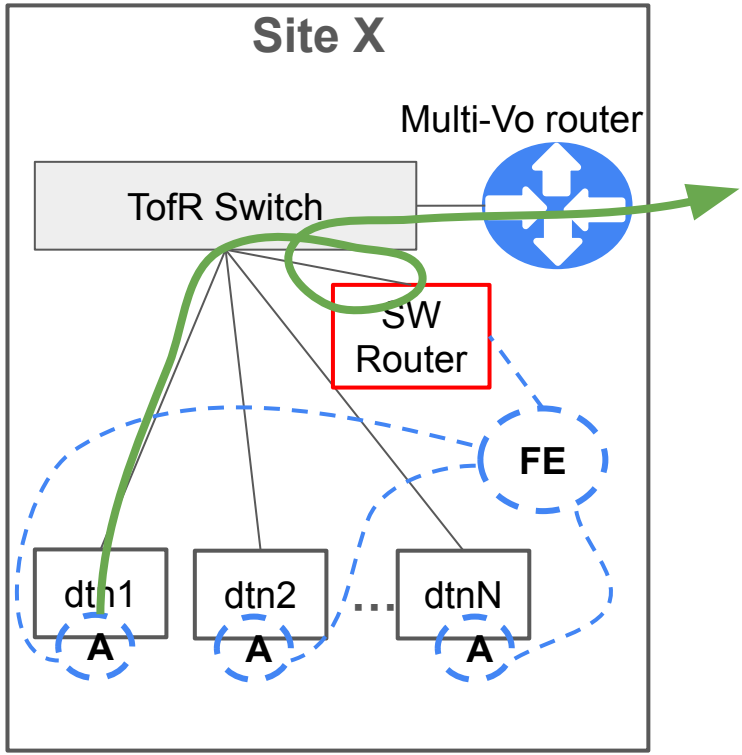
Regular Router vs Software-Router





Routing Difference





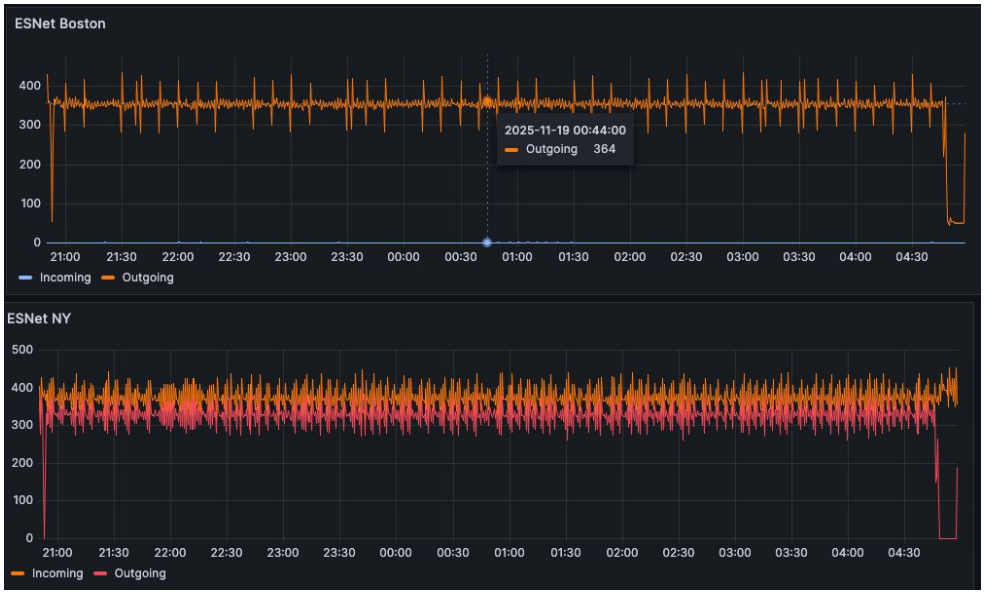
UMass Story

- UMass was the first ATLAS deployment in the SENSE testbed.
- A 400Gbps-capable node was deployed at the site and connected to the nautilus k8s cluster of the NRP project.
- SENSE Network control:
  - Access to the DTN router was granted
  - It was restricted to a **limited set of commands** on a limited set of **subnets** managed by SENSE
- Deployment Development:
  - Started testing with a self-contained XRootD deployment within the NRP node
  - Later on we swapped XRootD for a dcache proxy to UMass prod-storage
  - Finally we added 2 more proxies an added an HAproxy load balancer on top

UChicago Story

- UChicago joined the project in via IRIS-HEP/SSL effort
- They had their own k8s cluster to which we got access
- SENSE Network control:
  - We did not get access to the DTN router
  - Agreed on testing the **software-router approach**
  - Supplied us with a powerful (ConnectX7) NIC 200Gbps to be used for FRR/VPP
- Deployment Development
  - All contained in K8s (no access to productions storage)
  - XRootD-based cluster with an HAproxy load balancer
  - Storage was limited so we used the XrootD Blackhole plugin

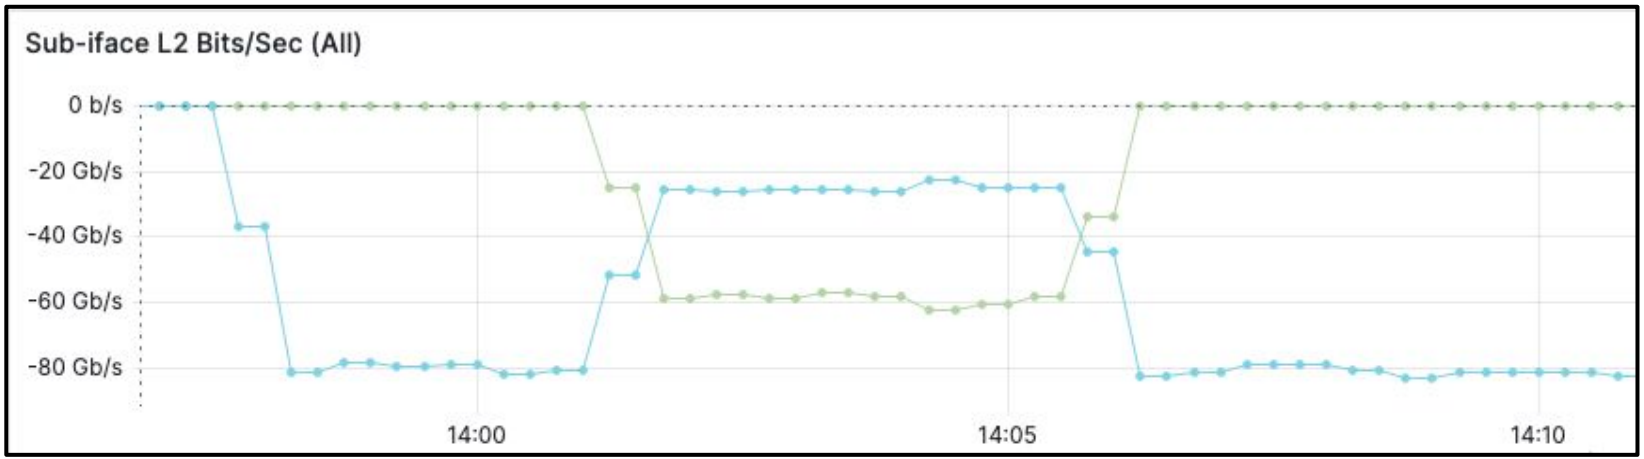
Story of Success: UMass



UMass performing 1.1Tbps at SC25.
Top: UMass/Boston => SC25 400 Gbps
Bottom: UMass/NY <=> SC25 400 + 300 Gbps

Using SENSE we were able to configure UMass to temporarily make use of their backup network path to increase their network capacity from 400 to 800 Gbps (symmetric) for the duration of the demonstration.

Story of Success: UChicago



During the WLGC 2026 Capability Mini-challenge we were able to demonstrate a 100Gbps SENSE managed data flow between UChicago and UMass

Lessons Learnt

- Deploying SENSE in a site is a lot easier when access to the DTN's router is available but this will not always be the case.
- We need to document, in detail, the minimal access needed by SENSE and best practices around it so that sites are more comfortable deploying this.
- The software-router works great at lower bandwidth levels. For higher bandwidth scaling and ease of deployment, SENSE access to local network element is recommended

Current Status

- Now that we have 2 ATLAS sites we have split the testbed by experiment
- The above required the deployment of a Rucio/DMM instance for ATLAS
  - Currently under testing
- Looking for joint-Demos where SENSE can help
  - Currently working with the ServiceX to investigate if SENSE can help reducing jitter
- On the lookout for more ATLAS sites to join the testbed

Acknowledgments

This ongoing work is partially supported by the US National Science Foundation (NSF) Grants:

- IRIS-HEP: OAC-1836650, PHY-2323298
- NRP: OAC-2112167.

Finally, this work would not be possible without the significant contributions of collaborators at ESNet, UMass, UChicago and SDSC.

Thank You!
Questions?











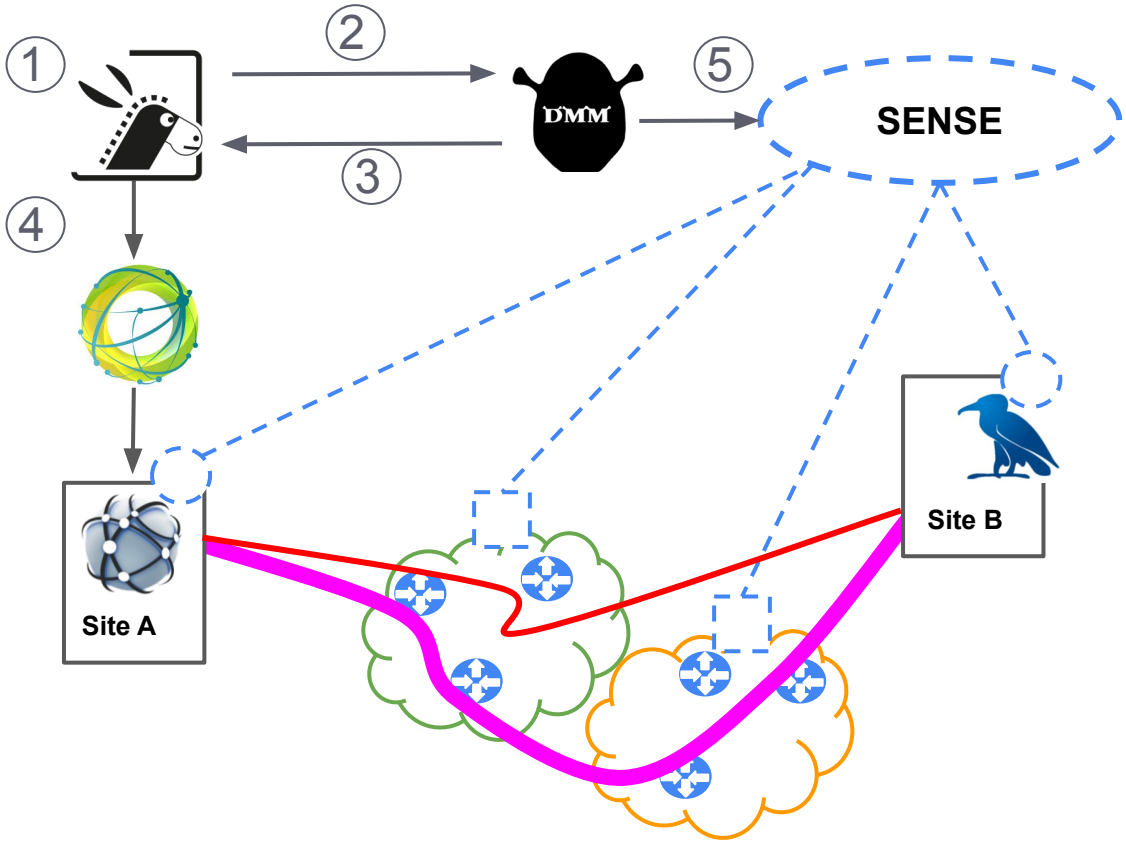



References

- SENSE: https://sdn-sense.github.io/
- Software-Router documentation: https://sdn-sense.github.io/optional-install/software-based-router/
- FRR: https://frrouting.org/
- VPP: https://s3-docs.fd.io/vpp/26.06/
- DPDK: https://www.dpdk.org/

Backup slides

What Makes SENSE? End-to-End Orchestration



1. A user creates a Rucio rule with the ***SENSE*** flag (“I want priority”).

2. Rucio contacts *Data Movement Manager* to get new pair of **endpoints**.

3. These are the **SENSE network touchpoints** on the **same sites** but in **special IPv6 subnets**.

4. Rucio-FTS transfers now **bind to** the SENSE priority path by replacing the old endpoints with the new ones.

5. Behind the curtain, *DMM* requests the creation of the SENSE path.

** SENSE error? Transfer remains on or fails back to default path **transparently.***