SC25 Network Research Exhibition: Demonstration Preliminary Abstract

StarLight AIDTN-as-a-Service for High-Performance Data Transport with Research Platforms

Jim Chen, Fei Yeh, Joe Mambretti jim-chen, fyeh, j-mambretti@northwestern.edu
International Center for Advanced Internet Research - Northwestern University,
Se-Young Yu, youf3@es.net, ESnet. Ting-Yu Lin, NCKU, p78101530@gs.ncku.edu.tw
Chun-Hao Yang, NYCU, chyang.cs12@nycu.edu.tw

Abstract

Building on prior StarLight DTN-as-a-Service (DTNaaS) initiatives, including those in the published paper, 'AIDTN: Towards a Real-Time AI Optimized DTN System With NVMeoF,' the StarLight consortium and its partners will showcase an advanced AI-optimized DTNaaS framework at SC25. This project aims to prototype AI-assisted network data movement services for end-to-end WAN infrastructure operating at speeds up to N x 400 Gbps.

The AIDTN-as-a-Service framework is designed to significantly enhance the performance of massive data capacity transport, especially for data intensive science, across wide-area networks (WANs) and local-area networks (LANs) within cloud environments, particularly those leveraging orchestrators like Kubernetes. By employing sophisticated AI techniques, the system dynamically sets up, optimizes, and monitors underlying network and system resources.

This SC25 demonstration will highlight cloud-native data transport services both within and between Kubernetes clusters, streamlining complex big data movement workflows among widely distributed sites. A key function of AIDTN-as-a-Service is its Jupyter controller, which empowers data scientists to intuitively identify, examine, and fine-tune Data Transfer Nodes (DTNs) for optimal performance in diverse research workflows. This integration of AI techniques enables efficient, high-performance data movement over long-distance WANs utilizing various networking fabrics.

Specifically, the demonstration will feature Memory-to-Memory, Disk-to-Disk, and NVMe-over-Fabrics (NVMeoF) with ROCE and TCP data transfers over multiple SC25 SCinet N x 400G WAN networks. This builds directly upon our successful StarLight SC24 demonstration, where we achieved approximately 2 x 387 Gbps end-to-end throughput over 2 x 400G WAN paths connecting StarLight, the Joint Big Data Testbed at McLean, Virginia, and the SC24 show floor. For SC25, we will further explore StarLight AIDTNaaS performance over N x 400G end-to-end WAN infrastructure and investigate additonal capabilitues, such as the potential integration of AIDTN with StarLight DTNaaS for even more enhanced data movement network services.

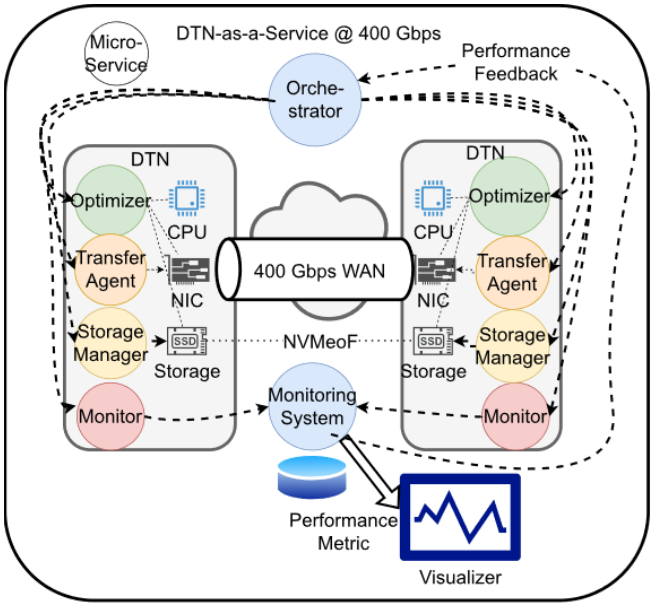

Figure 1. DTN-as-a Service Architecture

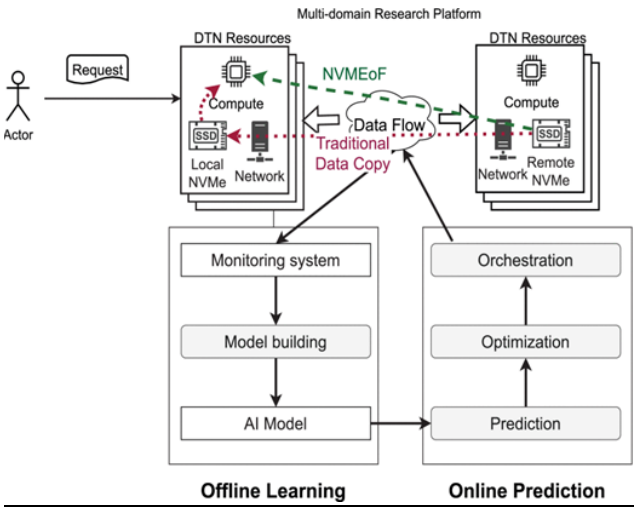

Figure 2. AIDTN* Architecture

Goals

1.Optimize Large Data Transport: The StarLight Exchange AIDTN-as-a-Service aims to significantly improve data transmission performance for large datasets among data-intensive science sites, including cloud environments like Kubernetes. For SC25, this capability will leverage high-performance optical transport switches, 400 Gbps L2 switches, and 2x400 Gbps servers utilizing a microservice architecture. The prototype will include Memory-to-Memory, Disk-to-Disk, NVMeoF with RoCE, and TCP transfers over N x 400G networks.

2.Implement Cloud-Native Services: We will implement cloud-native services for data movement within and among Kubernetes clusters, utilizing the DTN-as-a-Service framework to set up, optimize, and monitor underlying systems and networks via Jupyter notebooks.

3.Demonstrate AI-Assisted Data Movement: We will demonstrate the AI-controlled movement of data among nodes in Kubernetes using cloud-native services implemented in AIDTN-as-a-Service to improve and analyze performance over high-performance networks.

4.Scalable Cloud Resource Optimization: The StarLight AIDTN-as-a-Service architecture is specifically designed to work with cloud environments to optimize underlying Kubernetes resources (compute, storage, memory, and network) for scalable service deployment and orchestration.

5.Enhanced Resource Control and Monitoring: Enhancements include extending the Kubernetes capabilities within the AIDTN-as-a-Service software stack. This will enable precise control over DTN resources for data movement using cloud-native services and Jupyter notebooks, coupled with real-time monitoring for comprehensive data movement analysis.

6.Support for Diverse Hardware Platforms: For SC25, StarLight AIDTNaaS will demonstrate support for both AMD and Intel PCI-e Gen5 with 2 x 400G systems.

Impacts

The scientific research and Generative AI communities are addressing a critical challenge: large scale data movement requirements that are increasing by nearly 250%+ annually, exceeding the capacity of most existing infrastructures.

This project is developing a prototype AI-driven optimization workflow to manage and automate these massive data transfers over high-performance WAN networks.

Our work builds directly upon the published paper, 'AIDTN: Towards a Real-Time AI Optimized DTN System With NVMeoF,' and leverages prior StarLight DTN-as-a-Service (DTNaaS) initiatives.

SCinet provides an important testbed that enables demonstrating these techniques at scale.

In the future, we will continue to build on these techniques. We also intend to provide this prototype to SCinet participants, network exchanges, and the broader science community, enabling them to manage escalating data transport demands, especially over WANs. Target use cases include optimizing data center-to-data center transfers, enabling system-to-system AI system meshes, and powering automation processes involving big data movements.

Resources Required

SC25 SCinet WAN services totaling 1.2 Tbps end-to-end are required. This includes connectivity from the StarLight International/National Communications Exchange Facility in Chicago to the StarLight booth at the SC25 venue, between StarLight and the Joint Big Data Testbed (JBDT) Facility in McLean, Virginia, between the JBDT Facility and the StarLight booth at the SC25 venue, and among all sites. Additionally, the 400 Gbps ESnet testbed at Berkeley, connected to the StarLight Facility, will be utilized.

Involved Parties

- Jim Chen, iCAIR, jim-chen@northwestern.edu
- Se-Young Yu, ESnet, youf3@es.net
- Ting-Yu Lin, NCKU, p78101530@gs.ncku.edu.tw
- Chun-Hao Yang, NYCU, chyang.cs12@nycu.edu.tw
- Fei Yeh, iCAIR, fyeh@northwestern.edu
- Joe Mambretti, iCAIR, j-mambretti@northwestern.edu
- Marc Lyonnais, Ciena, mlyonnai@ciena.com
- Gaurauvdeep Shami, Ciena, gshami@ciena.com