


SC23 Network Research Exhibitions NOTED, SciTags, AutoGOLE, SENSE, and P4

- * NOTED: Network-Optimized Transfer of Experiment Data
- * SciTags: packet marking and flow labelling
- * AutoGOLE/SENSE: deliver network services end-to-end in a fully automated way via the Network Service Interface Connection Service (NSI-CS)
- * P4: Programming Protocol-independent Packet Processors

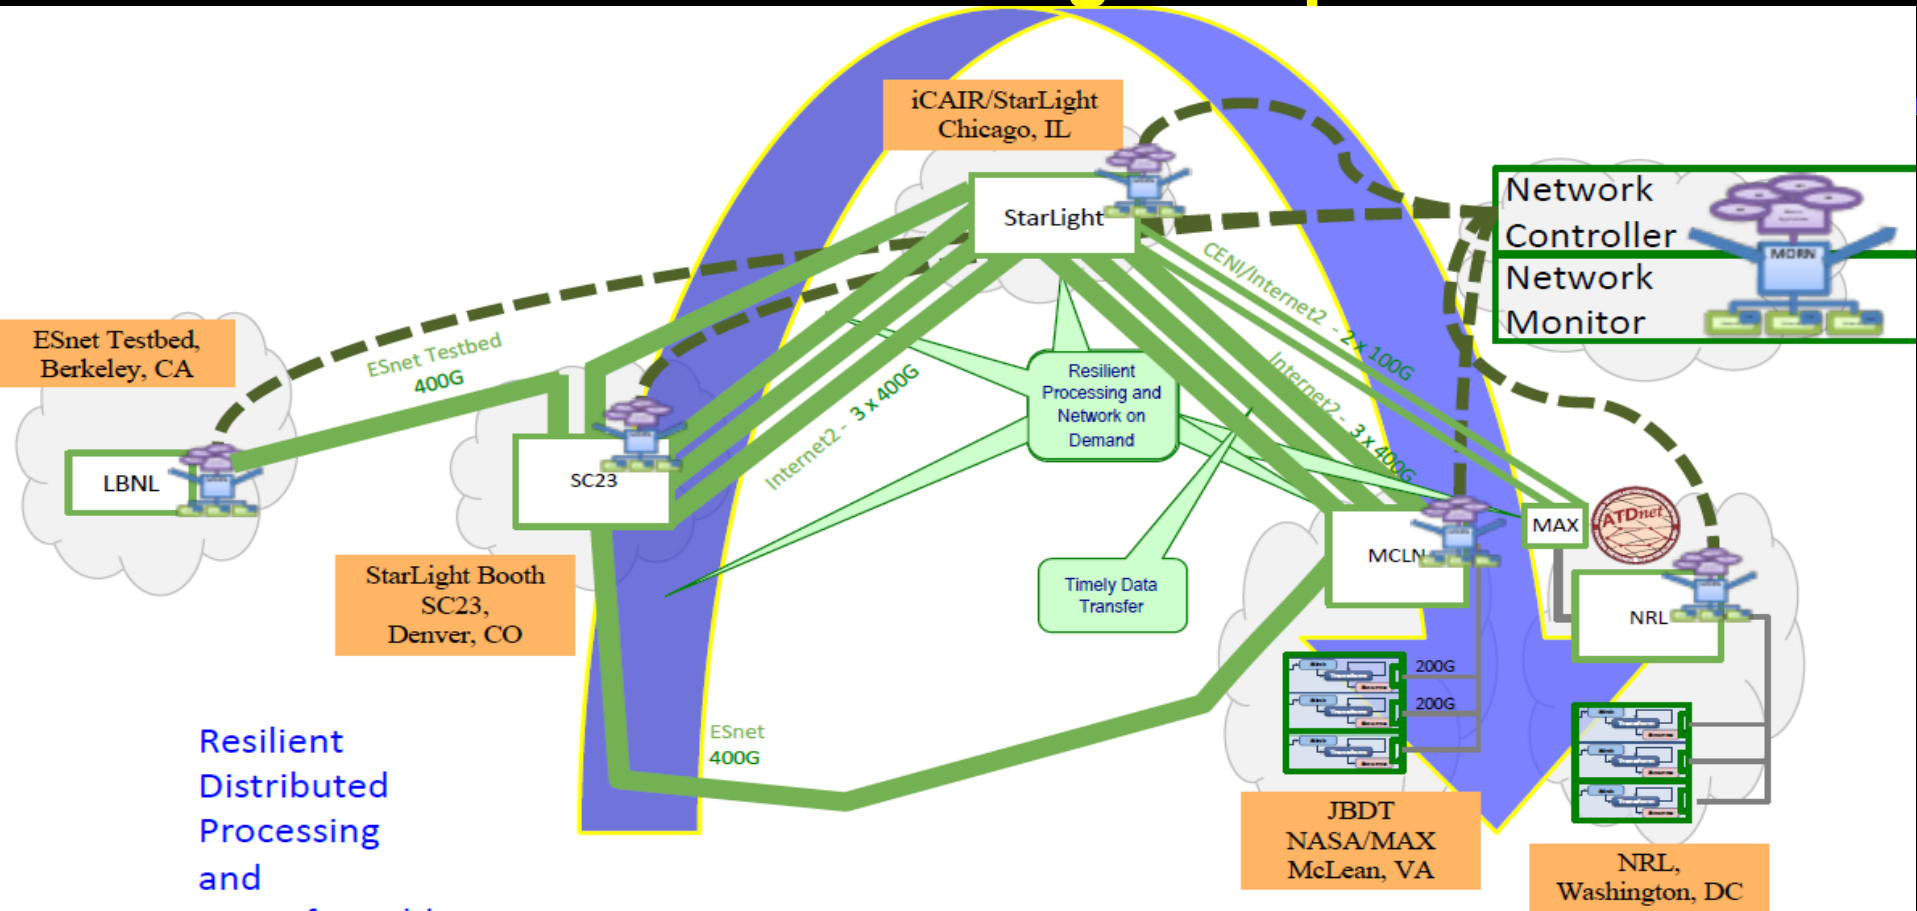


SC23 NRE map v.11, 2023-10-10 – WEJohnston, ESnet, wej@es.net

NOTED	SC21 NOTED infrastructure is in blue	ESnet	ESnet PoPs with High-Touch (line-rate, per packet) monitoring	McLean	Carrier hotels, etc
AG/SENSE	AutoGOLE / SENSE infrastructure is in red	Paris	Ovals are points of presence in regional infrastructure	ESnet	Rounded rectangles are individual switch/router
P4	P4 infrastructure is in green	Caltech	Sites	SOX	Exchange points (external or internal to a site)
general	Shared or general infrastructure is in black	SOX (Atlanta)		SC21 managed	
		Line Width			
		100G			
		200G			
		400G			
		800G			
		1 Tb/s			

- NOTES**
- 1) Within exchange points, etc., line width does not usually indicate bandwidth
 - 2) Map files (JPEGS, PDFs, and PPTX) are at <https://www.dropbox.com/sh/p2wcyppubsei7q/AAAMjF308xvUfospm3pRLLa1d?dl=0>
 - 3) P4 connections are only topological and are not associated with particular network link

Resilient Distributed Processing & Rapid Data Transfer



Resilient Distributed Processing and Reconfigurable Networks

Naval Research Laboratory
Center for Computational Science
SC23 Demonstration

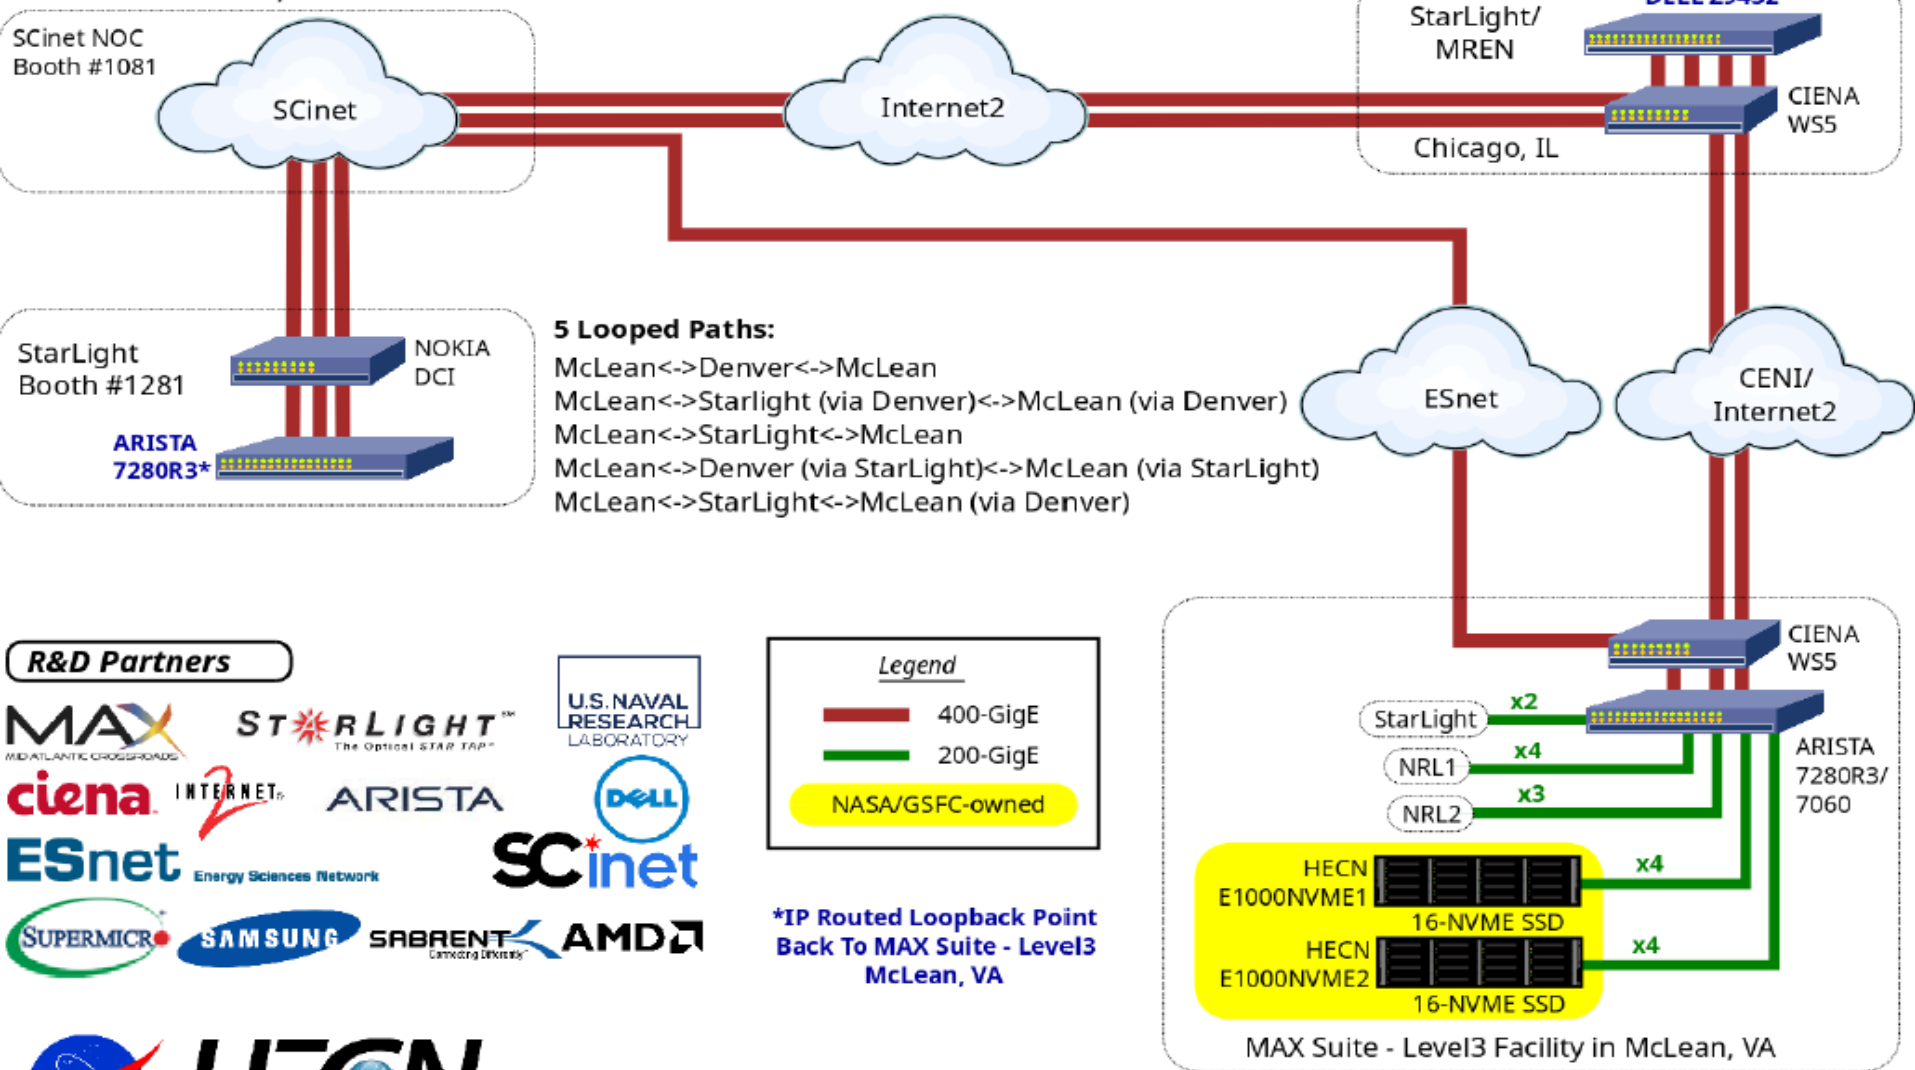


SC23 Joint Big Data Testbed

Demonstrations of 400 Gbps Disk-to-Disk WAN File Transfers using NVMe-oF/TCP

An SC23 Collaborative Initiative Among NASA and Several Partners

SC23 @ Denver, CO




R&D Partners



In memory of Paul Lang and Pat Gary

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink - 10/13/2023




Rucio

FTS

NOTED at KIT

NOTED at CERN

AutoGOLE SENSE



FTS
File Transfer Service

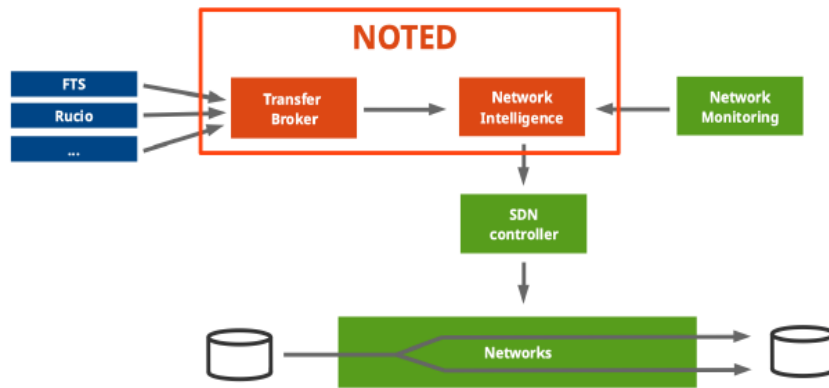
CRIC
Computing Resource Information Catalog

TRIUMF



elasticsearch

SKELETON AND ELEMENTS OF NOTED



FTS (File Transfer Service):

- ▶ Inspect and analyse data transfers to estimate if an action can be applied to optimise the network utilization → **get on-going and queued transfers**.

CRIC (Computing Resource Information Catalog):

- ▶ Enrichment to **get an overview** and knowledge of the **network topology** → get IPv4/IPv6 addresses, endpoints, rcsite and federation.

FLOWCHART AND DATASET STRUCTURE

- Input parameters: configuration given by the user
 - In noted/config/config.yaml → define a list of {src_rcsite, dst_rcsite}, maximum and minimum throughput threshold, SENSE/AutoGOLE VLANs UUID and user-defined email notification among others.
- Enrich NOTED with the topology of the network:
 - Query CRIC database → **get endpoints** that could be involved in the data transfers for the given {src_rcsite, dst_rcsite} pairs.
- Analyse on-going and upcoming data transfers:
 - Query FTS recursively → **get on-going** data transfers for each set of source and destination endpoints.
 - The total utilization of the network is the **sum** of on-going and upcoming **individual data transfers** for each source and destination endpoints for the given {src_rcsite, dst_rcsite} pairs.
- Network decision:
 - If NOTED interprets that the link **will be congested** → **provides a dynamic circuit** via SENSE/AutoGOLE.
 - If NOTED interprets that the link **will not be congested** anymore → **cancel the dynamic circuit** and the traffic is routed back.

DTN-UVIC-100g




1. Clients requesting data transfers from/to DTN-SC22-400g while passing science domain and activity fields via transfer protocols.

DTN-KIT-100g



canarie

4. High performance tests using eBPF-TC filters to test encoding of the science domains and activity fields in the IPv6 flow label at scale.



3. P4 programmable switch at CERN collecting the science domains and activity bits encoded in the packets.

CERN-LHCONE


P4 EdgeCore Wedge

CERN

DTN-CERN-100g
2.6GHz/32 cores
SSD, 100Gbps



5. Sampling of the low level TCP/IP metrics, which can be used by sites and R&Es to better understand the scientific flows.



DTN-SC22-400g
R7503 2.6 GHz
NVMe 2.0
2x200 Gbps

XRootD

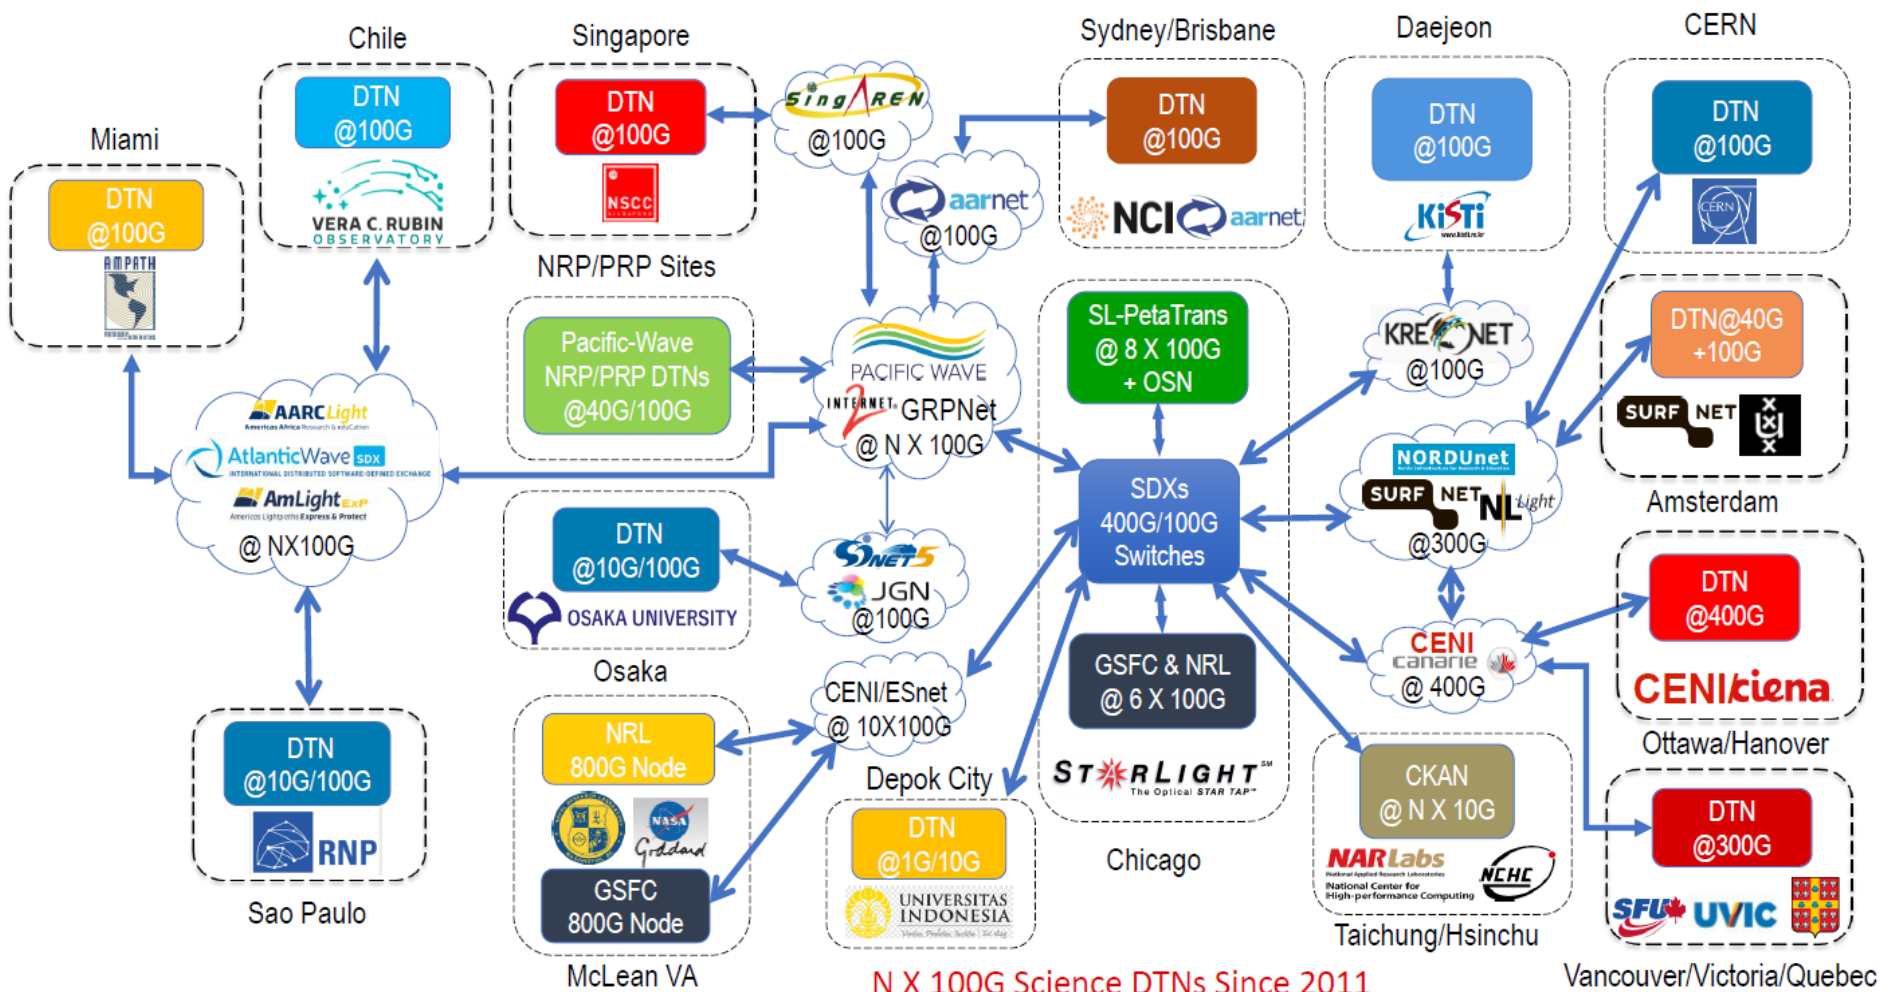
SC22 STARLIGHT™
Dallas, TX hpc accelerates.



2. XRootD storage responds to the client requests and marks the data transfer packets with the corresponding science domain and activity.

GRP DTNaaS For Petascale Science

GRP Service: DTNaaS for Petascale Sciences Data Movement



N X 100G Science DTNs Since 2011

Vancouver/Victoria/Quebec

DTN-as-a-Service – Demonstrated At SC23

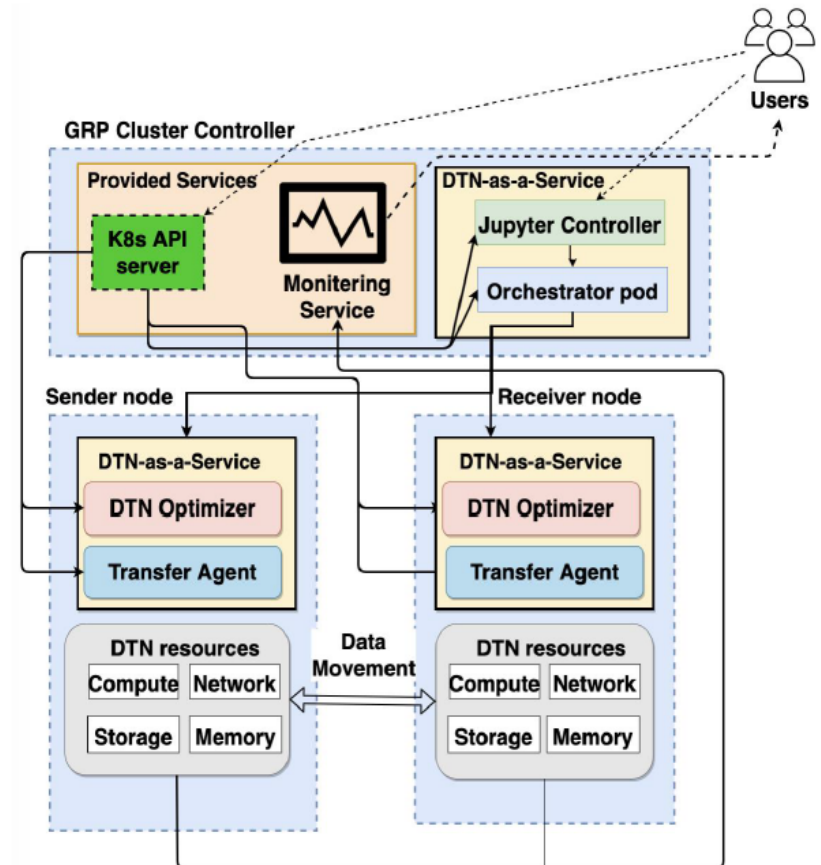
GRP Cluster with DTN-as-a-Service

DTN-as-a-Service(DTNaaS) provides a data movement workflow in GRP k8s cluster:

1. Deploy DTNaaS workloads via k8s API server
2. Use Jupyter to optimize and run transfers
3. Observe performance from monitoring service

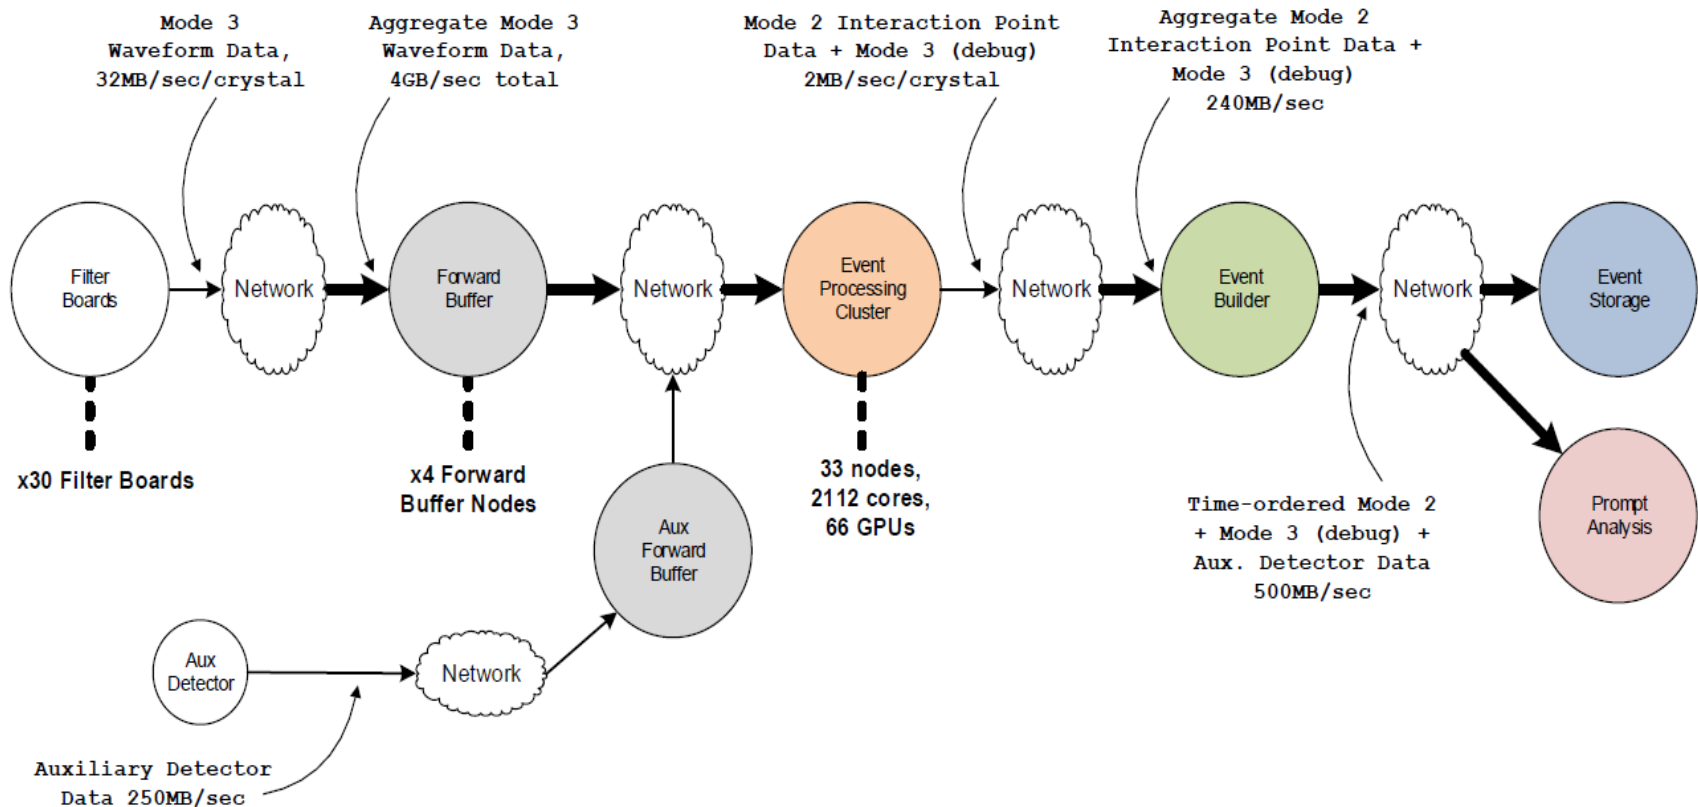
GRP DTNaaS Components:

- Orchestrator: controller of DTNaaS to manage agent and optimizer pods via REST API.
- Transfer Agent: run transfer jobs
- DTN Optimizer: optimize the DTN resources for workflow
- Jupyter: web interface to run DTNaaS interactively

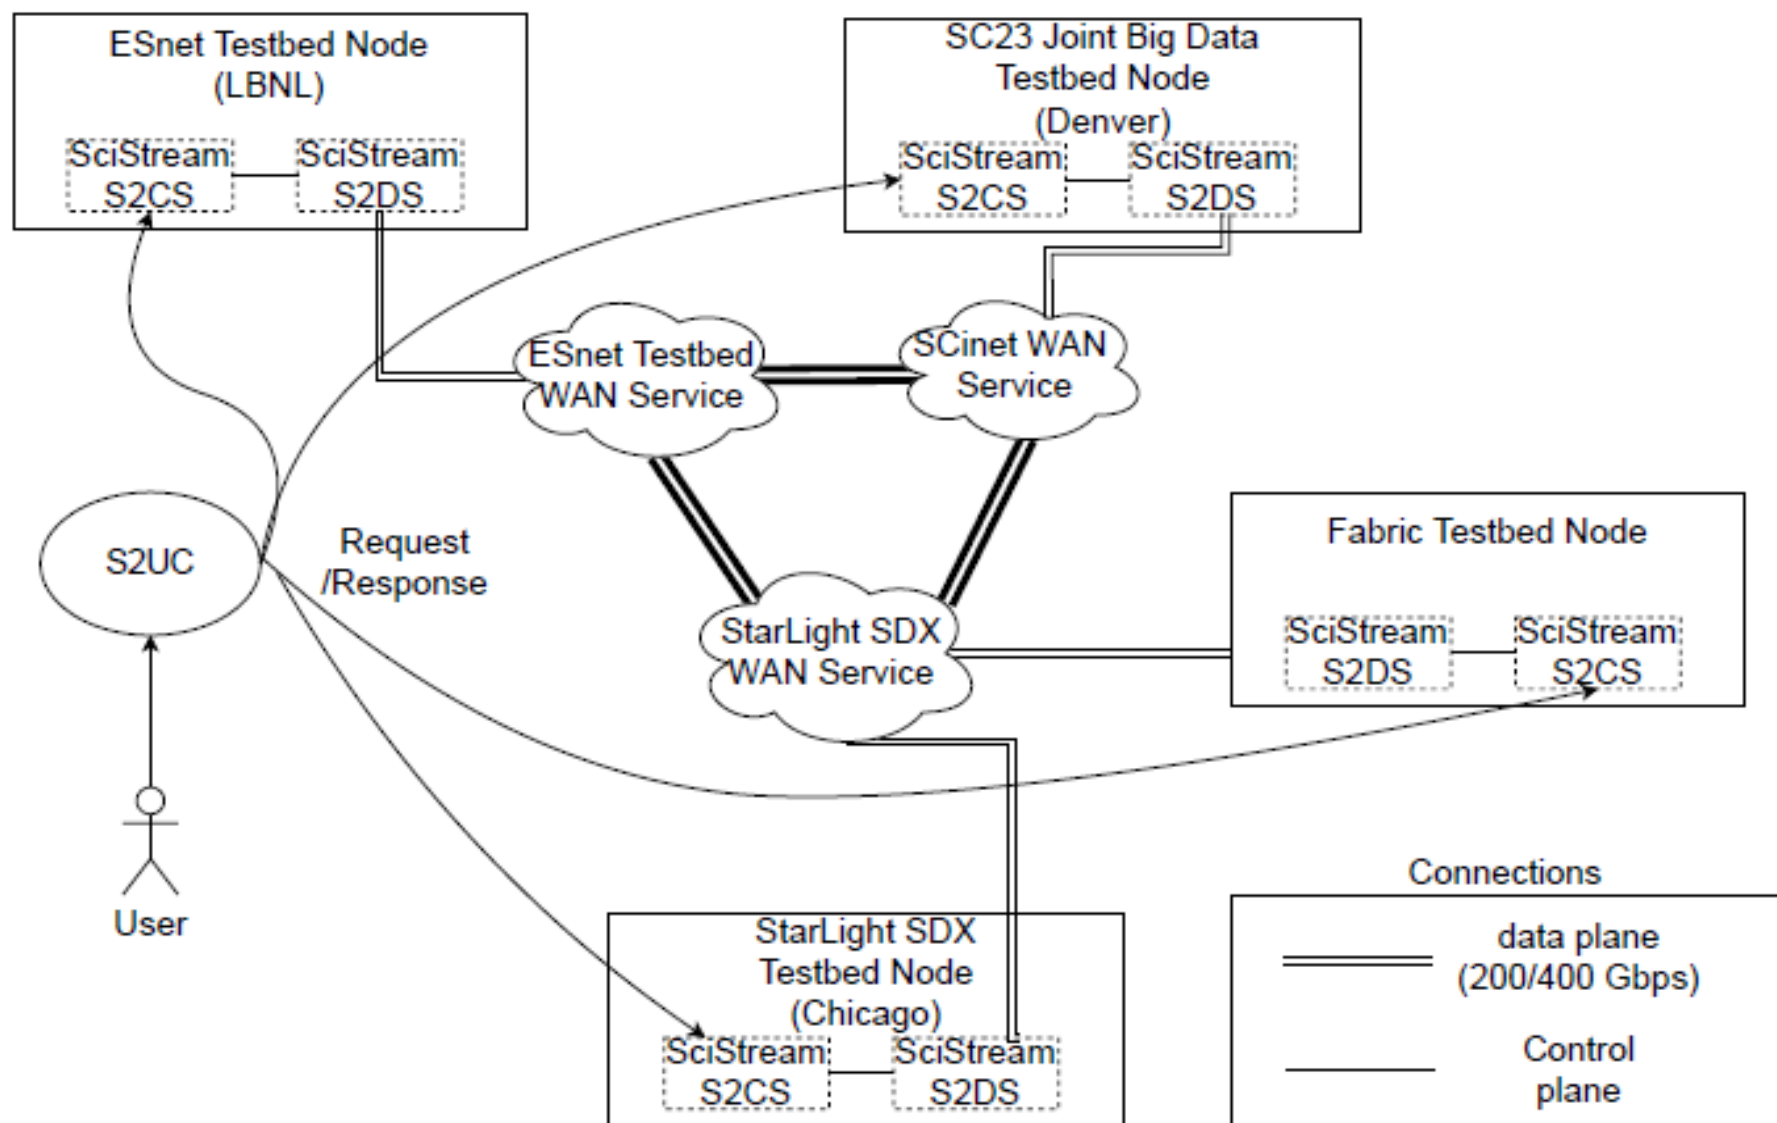


ESnet Gamma Ray Energy Tracking Array (GRETA)

GRETA Data Pipeline

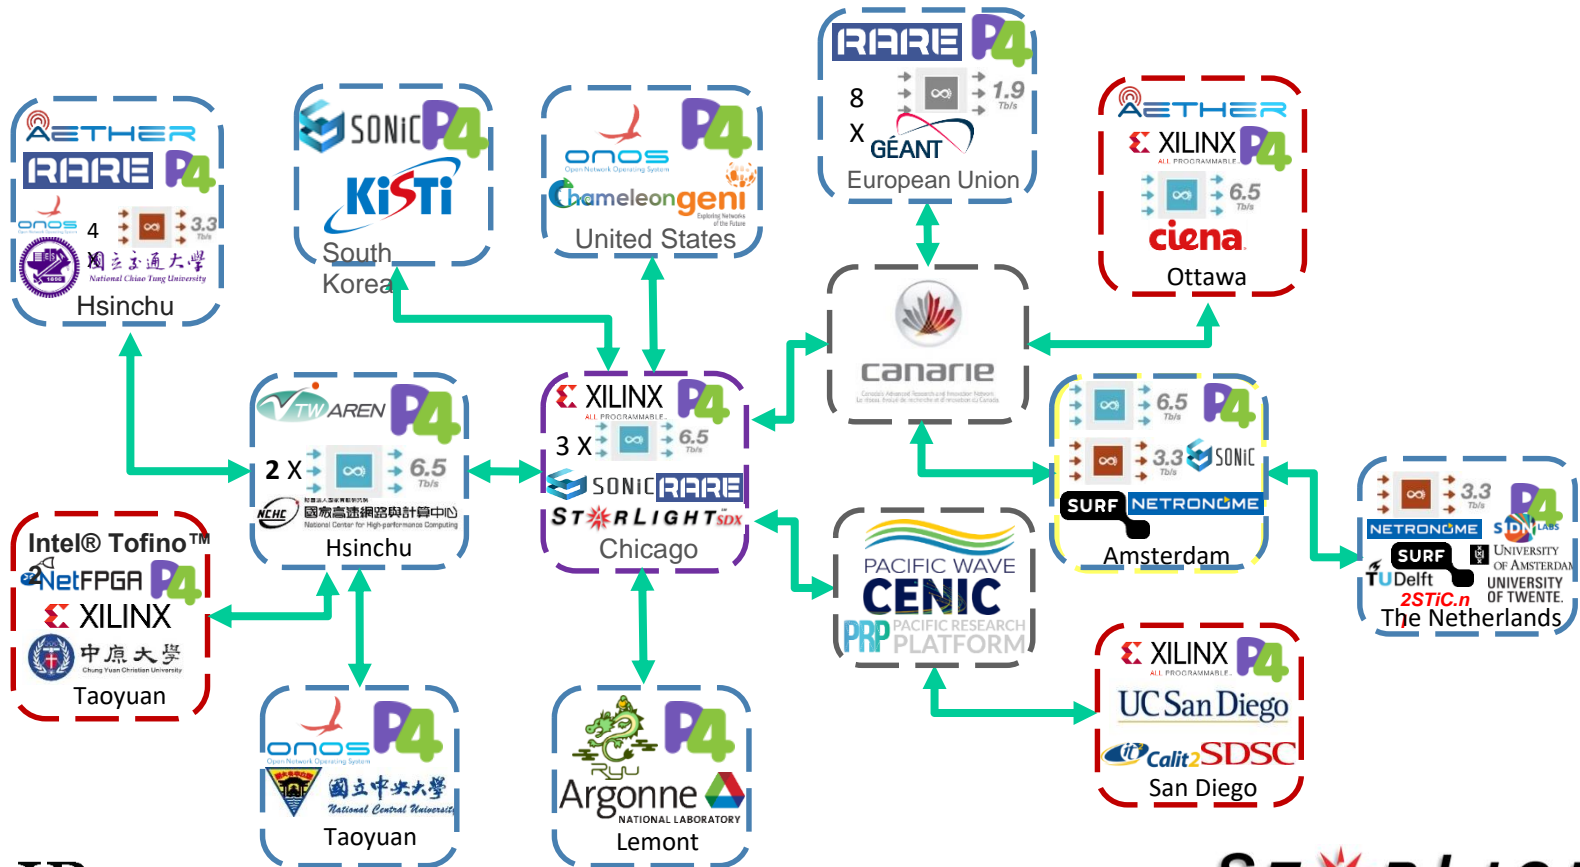


SC23 NRE-010: Multi-site data streaming orchestration with SciStream




International P4 Testbed Showcase at SC23

GRP Service: International P4 Experimental Networks (iP4EN)

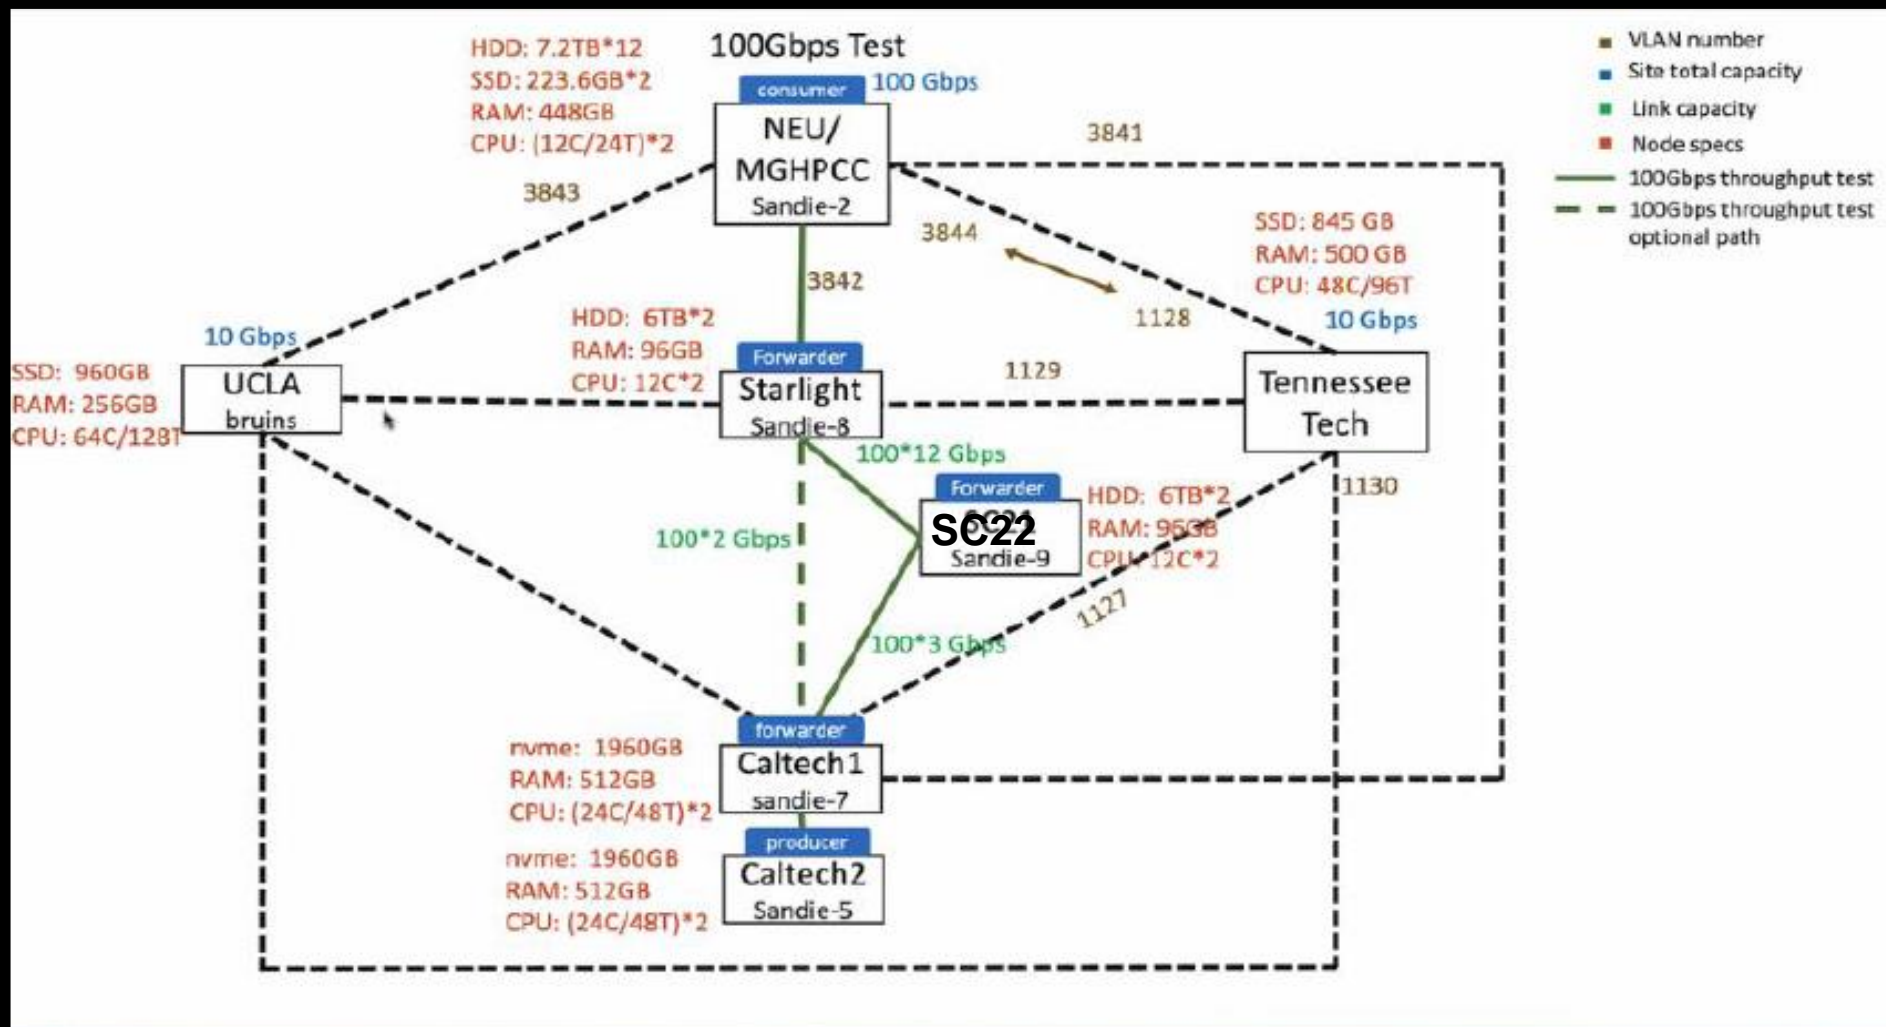


Integration With GEANT P4 Testbed

StarLight



Named Data Networking (NDN) for Data Intensive Science Experiments (N-DISE)





www.chameleoncloud.org

CHAMELEON: A LARGE SCALE, RECONFIGURABLE EXPERIMENTAL INSTRUMENT FOR COMPUTER SCIENCE

Kate Keahey

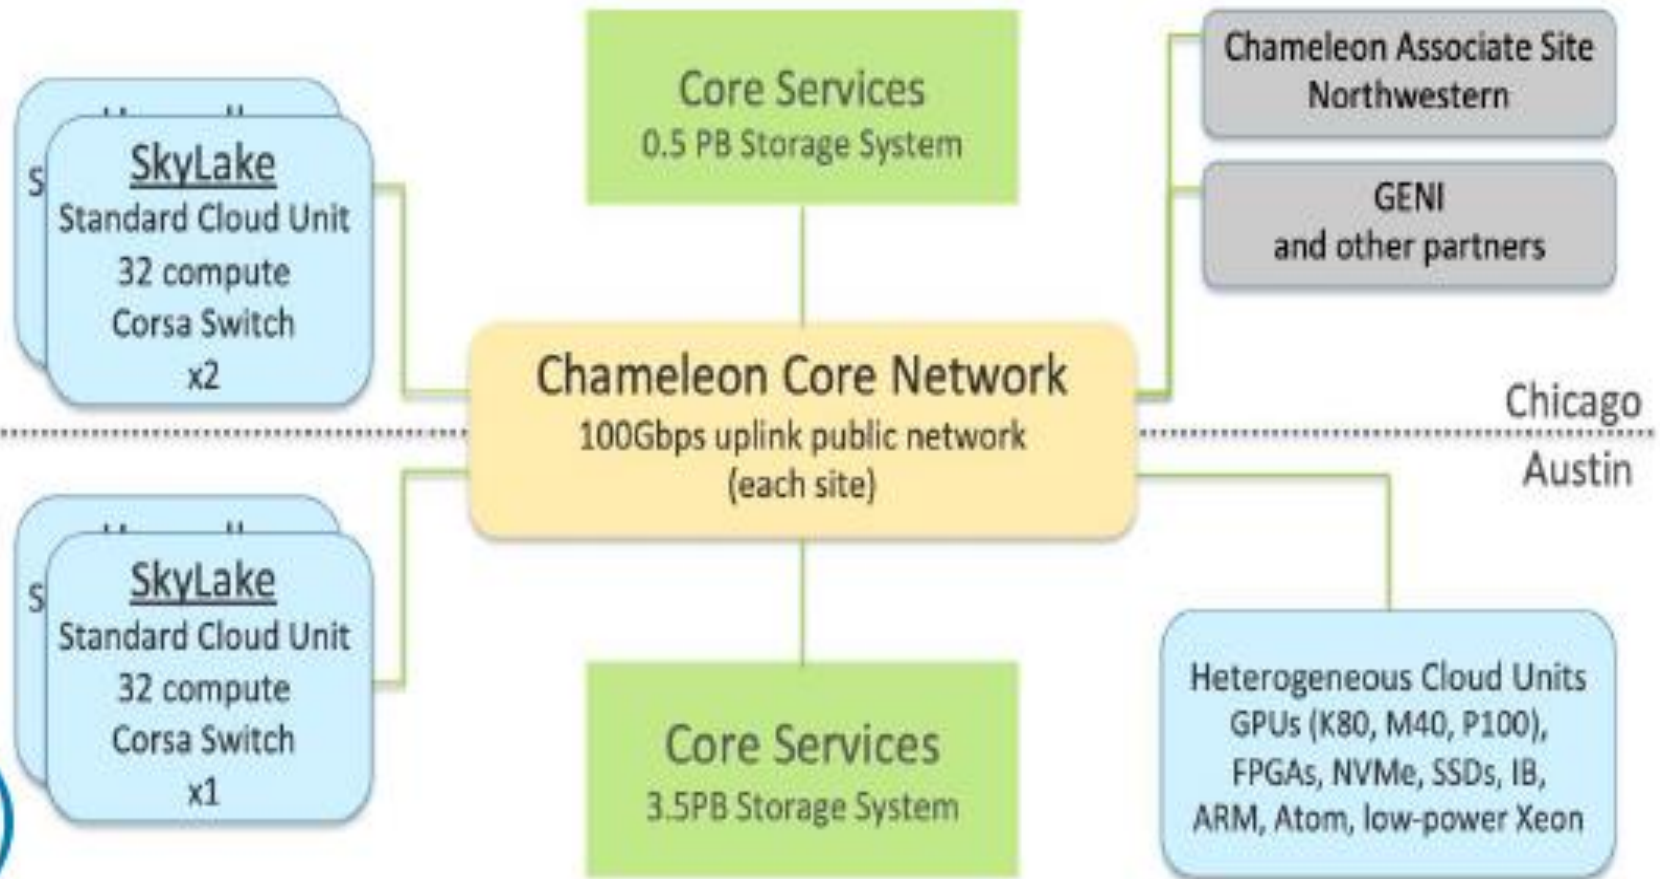
Joe Mambretti, Pierre Riteau, Paul Ruth, Dan Stanzione

SEPTEMBER 28, 2017


1



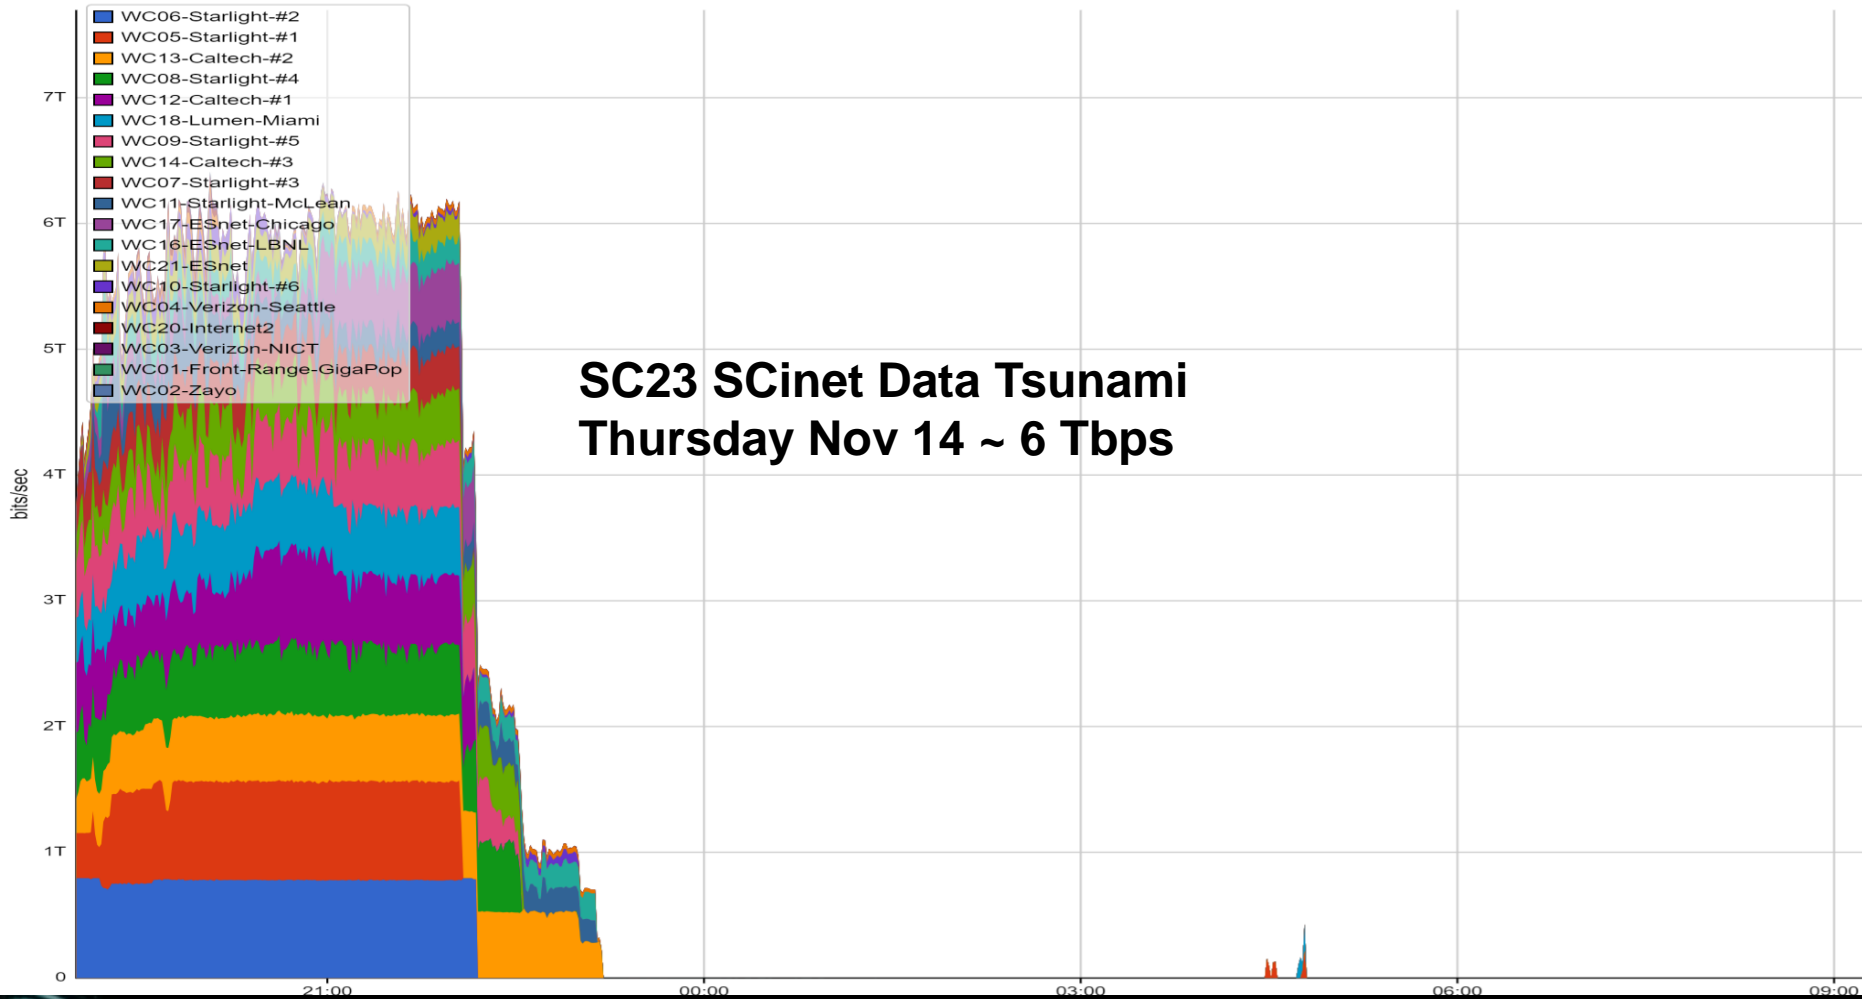
Chameleon CHI In A Box(CIAB) at StarLight



FABRIC Testbed (+FAB)




FABRIC Topology - with FAB Sites



Data Challenge

1





WLCG Site Network Input/Output



Row title (2 panels)


Transfer plots



National Research Platform


REGIONAL RESEARCH AND EDUCATION NETWORKS IN THE UNITED STATES

Proposed Extension of Nautilus 2021-2024



=> 5th NRP Workshop March 19-22, 2024 UCSD LIGHTSM

NRP+FABRIC




Source: Tom Lehman

STARLIGHTSM

February 18, 2024, a

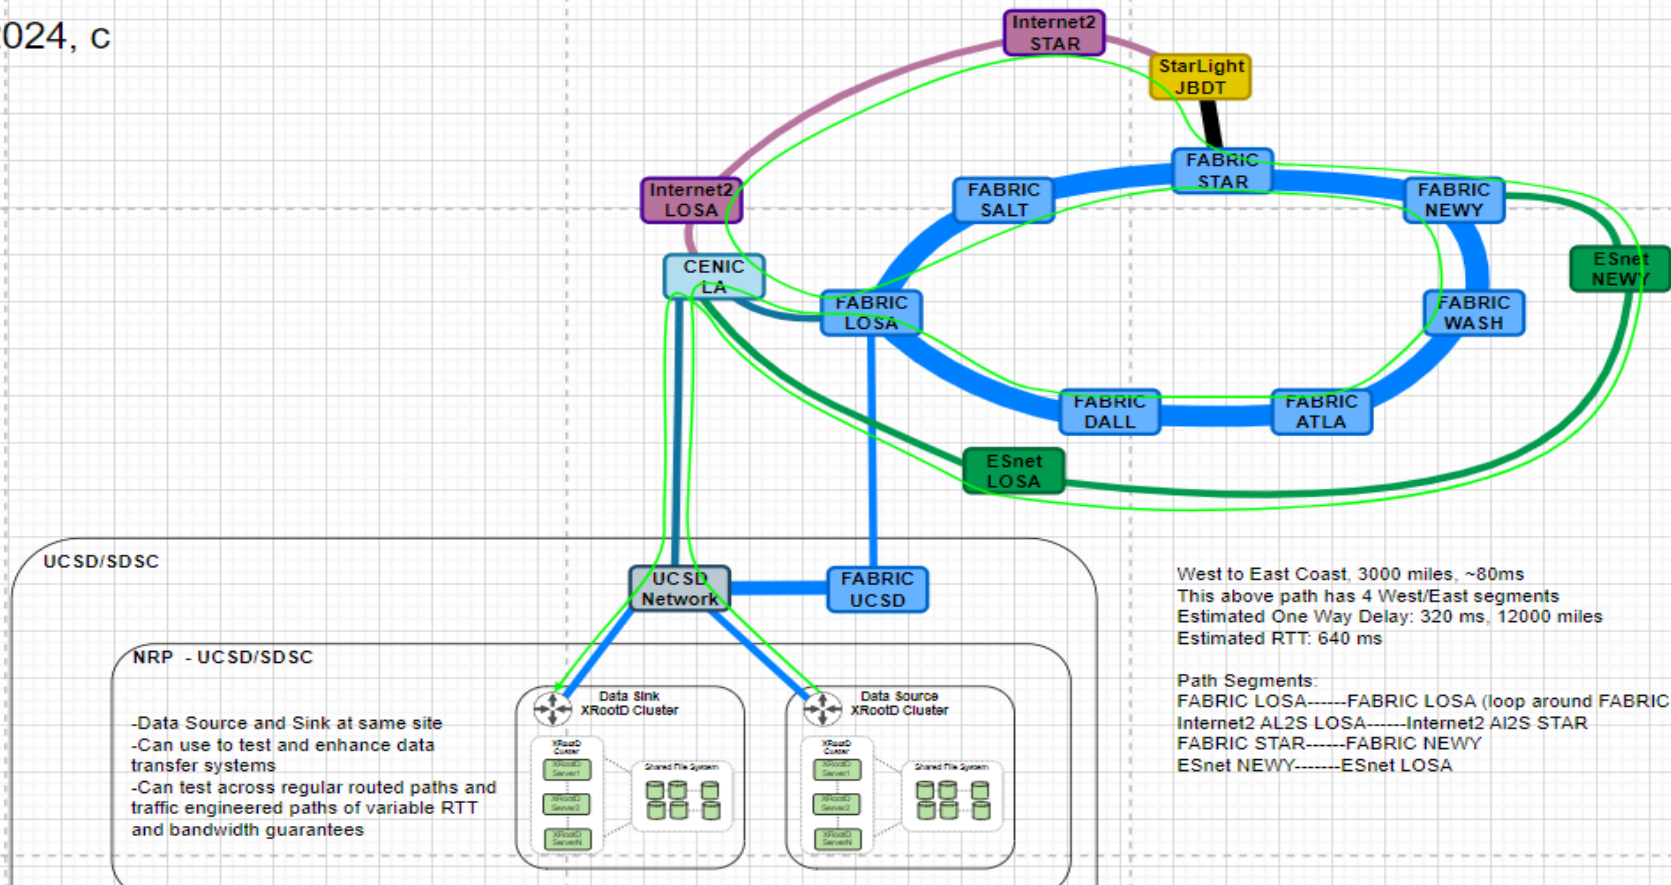
FABRIC SmartCity IP Plan:
-5 IP addresses for FABRIC SRX300
Headnode
Open Stack Virtual Router
2 IPs for VMs
-no firewalls are expected to be inline for these IPs.

VLAN Demo Assignments:
3100-3102: FABRIC Infrastructure
3103-3105: SDSC Gen5 Server
3106: Bristol

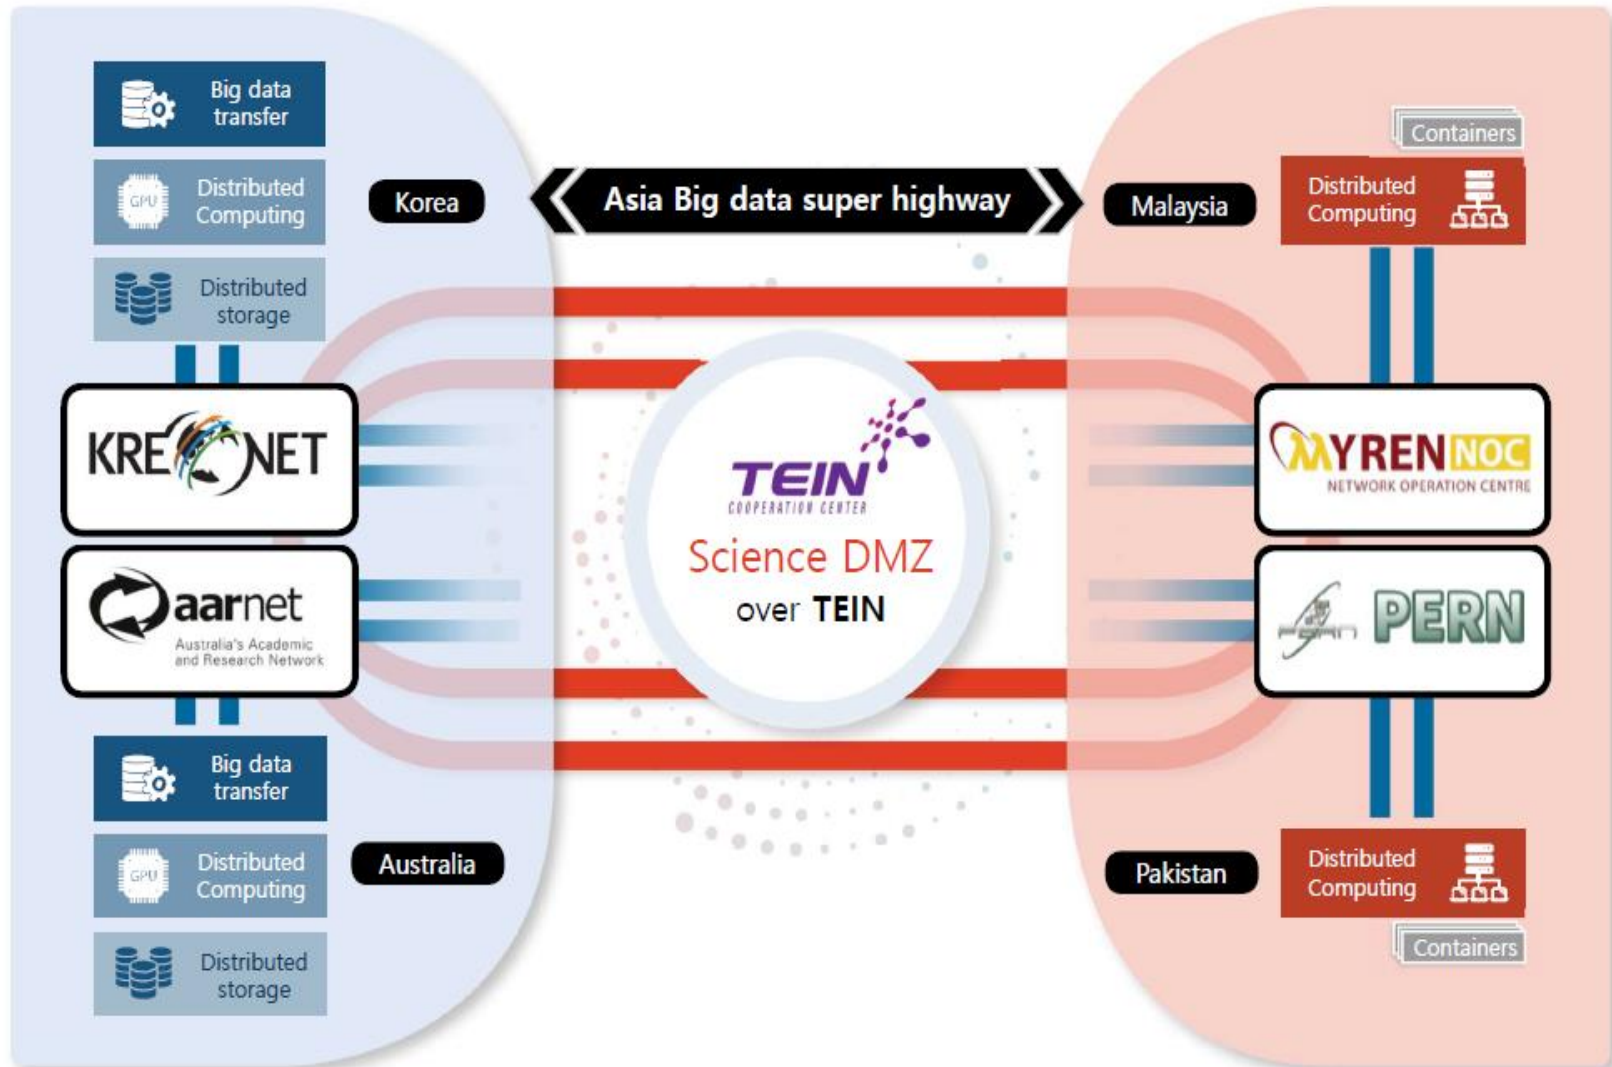


y 25, 2024, c

Example Topology for Testing, Research, Development Data Management and Movement Systems



Supports Asia Pacific Research Platform




Quantum Communications And Networks: Motivation

- **Quantum Enables Many New Applications**
 - Security – e.g., Quantum Key Distribution (QKD), Highly Secure Information Transmission, Quantum Encryption
 - Quantum Sensors
 - Quantum – e.g., Precise Clocks
 - New Applications Derived From Unique Properties (e.g., Superposition) And Novel Quantum Devices
 - Communications Among Quantum Computers, e.g., To Address Complex Computational Science Problems Through Distributed Quantum Environments




Emerging Chicago Quantum Exchange Testbed





Demo Lead Partner - NuCrypt (1) - Distribution of Quantum Entanglement Through Fiber With Co-Propagating Classical Data

(1) Spin Off From Northwestern University's Center for Photonic Communications and Computing, Which Was Also A Partner for the OFC 2023 Demonstrations (Prem Kumar, Director)





**Co-Propagation
And 400 Gbps WAN
Demonstrations
OFCnet Booth
March 2023**

STARLIGHTSM

www.startup.net/starlight

Thanks to the NSF, DOE, NASA,
NIH, DARPA
Universities, National Labs,
International & Industrial
Partners,
and Other Supporters

STARLIGHTSM