

Supporting Production Science at Major Facilities by Leveraging Research and Education Networks

Julio Ibarra Research Professor Principal Investigator

June 25, 2024



Motivation for this webinar





Observations

- NSF Research Infrastructure Workshops: <u>2023</u> and <u>2024</u>
- CI4MF 2024 Collaboration in Action
- Observations about MFs and their Challenges:
 - Designing, building and operating CI for MFs is increasingly challenging
 - Lifecycle changes: Planning & Design, Construction, Operations, periodic enhancements
 - Data archival is an increasingly important issue
 - Storage needs for MFs is increasing
 - Preservation of scientific data for future generations is critical

Clouds may be a solution for operations and data archival for some MFs



Challenges transitioning to the Cloud

- Participants at RIW 2023, CI4MF 2024, and RIW 2024 reported challenges transitioning to the Cloud
 - Managing costs, vendor lock-in, and storage were the topics most discussed as being challenging for the MFs to navigate
- <u>Managing Costs</u>: Costs are incurred for data storage, egress, compute resources, serverless functions, data movement, API requests, etc.
- Vendor Lock-in: Transitioning between Cloud vendors can become costly, because of differing APIs, tools, and differing services
- The national CI ecosystem offers a plethora of options for MFs to achieve their goals by <u>collaboratively</u> building effective solutions to serve the goals of their mission

CI Compass: NSF Major Facilities Cloud Use Cases and Considerations Report



Leveraging R&E Networks to build effective solutions

- How can MFs leverage R&E networks to build effective solutions to serve the goals of their mission?
- To answer this question, I will present the following topics:
 - A brief review of the Fundamentals of R&E Networks
 - A sampling of the R&E Network ecosystem
 - Measurement Instrumentation and Resources of R&E Networks
 - Major Facilities leveraging R&E Networks
 - An SLA-driven Science Use Case: the Vera Rubin Observatory
- Propositions after each section to stimulate discussion in the Q&A
- Conclusion





Fundamentals of R&E Networks





Fundamentals of R&E Networks

- R&E Networks are sponsored and operated by R&E organizations
 - Universities (FIU, USC), Regional R&E Networks (Florida LambdaRail, Cenic), National Research and Education Networks (NRENs: Internet2, Esnet)
 - NSF International Research Network Connections (IRNC) program
- R&E Networks span a territory/region to serve their Community of Practice
 - E.g., FLR serves the R&E community of Florida, Cenic of California, etc.
- R&E Networks are mission-oriented, sponsored by a government or a Community of Practice
 - NOIRLab operates a production R&E network to support 5 programs: CTIO, CSDC, Gemini, KPNO, and Vera Rubin Observatory
 - The <u>ESnet network</u> is a high-performance, unclassified network built to support scientific research, providing advanced networking services to more than 50 DOE research sites



Fundamentals of R&E Networks [2]

- R&E Networks conduct experimental research to support continued technical evolution of the R&E networks
 - Software-Defined Networking, Programmable Data Planes, In-band Network Telemetry, Innetwork computation
- Collaboration and Cooperation are core tenets of R&E Networks to support research and education, and the mission of their stakeholders
 - The Vera Rubin Observatory Long-haul Network is built from a collaboration and cooperation of 10 R&E networks

V. Cerf, Thoughts on the National Research and Education Network, <u>RFC 1167</u> NSF <u>International Research and education Network Connections (IRNC)</u> program announcement





Sampling of the R&E Network Ecosystem



Internet2 (NREN) and the Quilt Regional RENs





ESnet6 National R&E Network



11 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks





International R&E Networks



Advanced North Atlantic (ANA) Collaboration



13 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks



International Networks @ Indiana University (IN@IU)

- 25 year history of operating international networks and building partnerships
- Currently support 4 100G links with 2 US NSF awards (TransPAC and NEA3R)
- Coordinate globally to ensure US researchers have worldwide access to networking services
- Work with research groups to improve international data flows via engagement and training
- Projects carry traffic data from researchers in 84% of the world





Pacific Islands R&E Networking (PIREN)





Pacific Wave International Peering Exchange





AmLight International R&E Network: FIU, PI: Julio Ibarra

- 600G of upstream capacity between the U.S., Latin America, Caribbean and 100G to Africa
 - Increasing to 1.1 Tbps between Florida and Brazil
- OXPs: Florida(3), Atlanta, Brazil(2), Chile, Argentina, Puerto Rico, Panama, and South Africa
- Production SDN Infrastructure since 2014
- Deeply programmable across the network stack
 - Programmable P4 Data Plane
 - Open Source SDN Controller
 - Fine-grained telemetry
 - Run-time network verification
 - Closed-Loop Orchestration
- Highly instrumented
 - PerfSonar, sFlow, Juniper Telemetry Interface (JTI), Inband Network Telemetry (INT)





Network Startup Resource Center (NSRC), U of Oregon, PI: Steven Huter







NSRC Engage Project

1100+ Field Activities & 60,000+ Trained









Propositions

- R&E networks have an abundance of bandwidth capacity
- R&E networks span local, regional, national and international geographies to support their Communities of Practice
- R&E networks are well connected, forming a global fabric of network connectivity
- MFs in remote locations can leverage the R&E global fabric of network connectivity to accomplish their science





Monitoring, Measurement Instrumentation and Resources of R&E Networks



Network Performance and Troubleshooting

Network Performance and Troubleshooting (perfSONAR)



perfSONAR: Comprehensive monitoring and measuring of network resources

Introduction

perfSONAR is a tool for end-to-end monitoring and troubleshooting of multi-domain network performance. perfSONAR provides network engineers with the ability to test and measure network performance, as well as to archive data in order to pinpoint and solve service problems that may span multiple networks and international boundaries. perfSONAR is currently deployed at over 2000 locations around the world, and is extremely valuable when doing <u>network troubleshooting</u>. perfSONAR has been developed through an international collaboration led by Internet2, ESnet, Indiana University, and GEANT.

perfSONAR Overview

Current list of global perfSONAR services Download the latest version of the <u>perfSONAR Toolkit</u> ESnet's <u>perfSONAR Dashboards</u>

https://www.es.net/network-r-and-d/perfsonar/

Table of Contents

Introduction
Why use perfSONAR?
How to effectively use

perfSONAR 4. Future directions

5. Contact perfSONAR

6. Additional Information

The Engagement and Performance Operations Center (EPOC)

- EPOC supports 5 main activities:
 - Roadside assistance and consultation
 - Application Deep Dive
 - Network Analysis
 - Data Mobility Exhibition
 - Coordinated Training

https://epoc.global/



Vera Rubin Observatory PerfSonar Network



23 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks



Network Performance Resources for MFs

For Engineers:

- TCP Tuning
- 100G Tuning
- Data Transfer Nodes

For Scientists:

- Performance Expectations
- Data Transfer Tools
- Science DMZ

24

ESnet Fasterdata Knowledge Base

ESnet Fasterdata Knowledge Base

An Expert Guide for End-to-End Performance Tuning, Tools and Techniques

The Fasterdata Knowledge Base provides proven, operationally sound methods for troubleshooting and solving performance issues. Since 1986 ESnet has operated an advanced research network with the goal of enabling the highest levels of performance for the Department of Energy (DOE) scientific community. During this time, our engineers have identified a common set of issues that hinder performance. We share our experiences and findings in this knowledge base.

Our solutions fall into five categories:

- Network Architecture, including the Science DMZ model
- Host Tuning

CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks

- Network Tuning
- Data Transfer Tools
- Network Performance Testing
- Wireless and 5G networking for science

How long should it take to transfer a TeraByte of data across your network? It is probably less time than you think. Check out our <u>performance expectations guide</u>.

Want to contribute material or know of something that should be on the site? Send an email to fasterdata@es.net.



Research and Education Network



Network Troubleshooting

100G Host Tuning

Linux Tuning

TCP Tuning

Tools

Techniques

Science DMZ

Architecture

Data Transfer Nodes

Performance Expectations for a 100G Host

Data Transfer Tools

Data Transfer Tools Overview

NSF International

Globus

NSF Info

Globus Science DMZ and DTN application

- Globus engages storage systems via DTNs located in a Science DMZ
- The Globus Endpoint (Globus Connect Server agents) within each Science DMZ can be requested to perform a transfer
- The request can then proceed without passing through a firewall

Slide by Ian Foster



* Not actively listening; only used when transfer is in progress; may be restricted to private network Please see TCP ports reference: https://docs.globus.org/resource-provider-guide/#open-tcp-ports_section





Propositions

- R&E Networks have very experienced and talented network architects and engineers to design and build networks for MFs
- R&E Networks have a comprehensive set of tools and resources for monitoring, measuring, tuning, and troubleshooting networks to assist MFs achieve their goals





Major Facilities leveraging R&E Networks



NOIRLab

- Cerro Tololo Inter-American Observatory (CTIO)
- Gemini Observatory
- Kitt Peak National Observatory
- Vera Rubin Observatory



28 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks



NOIRLab R&E Network



LHCONE

- Private network connecting HEP Tier1 and Tier2 sites
- Sharing network resources
- Collaboration among 30 national and international R&E networks
- 110 connected sites

E. Martelli, Networks for High Energy Physics





NSF Campus Cyberinfrastructure(CC*) Program

- Funds networking, compute and storage at the campus level
- Provides support for Science DMZs
- Requires a minimum of 20% resource sharing for compute and storage awards by leveraging PATh or ACCESS (resource providers)
- The <u>CC* program</u> has made over 500 awards



31 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks





Propositions

- MFs are leveraging R&E networks and RPs to achieve their goals
- Collaborations between MFs, the R&E networks at the campus, regional, national and international levels, and the RPs, leveraging shared CI, can provide cost effective solutions





SLA-driven science use case: Vera Rubin Observatory



Vera Rubin Observatory operation use case

- Vera Rubin is a large-aperture, wide-field, groundbased optical telescope under construction in northern Chile
- The telescope will take a picture of the southern sky every 27 seconds (transfer window), and produce a 13 Gigabyte data set
- Each data set must be transferred to the U.S. Data Facility at SLAC, in Menlo Park, CA, within 5 seconds, inside the 27 second transfer window
- Challenges
 - High propagation delay in the end-to-end path
 - RTT from the Base Station to the USDF is approximately 180+ ms
 - 0.001% of packet loss will compromise the Rubin Observatory application



34 CI Compass Webinar: Supporting Production Science at Major Facilities by leveraging R&E Networks



Vera Rubin Observatory Long-Haul Network (LHN)

- The LHN was designed and built to respond to the Vera Rubin SLA
- 10 R&E networks collaboratively and cooperatively have built and are operating the LHN
- The LHN leverages the R&E networks
- Operations is scheduled to start July 2025







Propositions

- R&E Network operators can design and build R&E networks in response to requirements (or SLAs) from MFs
- R&E Network operators cooperate and collaborate to design and build solutions that leverage their networks to support MFs





Conclusion



Can MFs leverage R&E networks to build effective solutions?

- I conclude that the answer is <u>Yes</u>, by
 - Collaborating and cooperating to design and build solutions for the MFs that leverage R&E networks
 - Leveraging the broader CI ecosystem on the campuses, at the regional, national, and global levels
 - Leveraging the expertise, talent and relationships in the R&E ecosystem, in collaboration with the MFs, to build effective solutions



THANK YOU













