AmLight ExP

Americas Africa Lightpaths **Express & Protect**

Vasilka Chergarova

Florida International University, IT Assistant Director

Resources supported by AMPATH/AmLight to enable ML/AI

November 16, 2024

AmLight-ExP (NSF #OAC-2029283)

Outline

- What is AMPATH and AmLight-ExP?
- Shared resources supporting Machine Learning and Artificial Intelligence research in Higher Education
  - National Research Platform (NRP)
  - Open Science Grid (OSG) & Open Science Data Federation (OSDF)
  - ACCESS
  - FABRIC
- Additional research: Leveraging the power of AI to detect DDoS attacks
- Conclusion

(AmLight ExP project is supported by NSF #2029283.)

Americas-Africa Lightpaths Express and Protect (ExP)

AmLight ExP is a hybrid network that uses Optical spectrum (Express) and Leased capacity (Protect) to build a reliable cutting-edge network infrastructure for research and education

AmLight ExP Network currently consists of:

- 600G of upstream bandwidth between the U.S., Latin America, Caribbean, and 100G to South Africa in Cape Town.
- Open Exchange Points (OXPs) with Points of Presence (PoPs) in: Florida (3), Brazil(2), Chile, Puerto Rico, Panama, and South Africa, New: Georgia (Atlanta), Argentina (Buenos Aires)
- Production Software-defined networking (SDN) Infrastructure since 2014
- Deeply Programmable R&E Network Infrastructure
- Highly instrumented: PerfSONAR, sFlow, Juniper Telemetry Interface (JTI), In-band Network Telemetry (INT)

Americas-Africa Lightpaths Express and Protect (ExP)

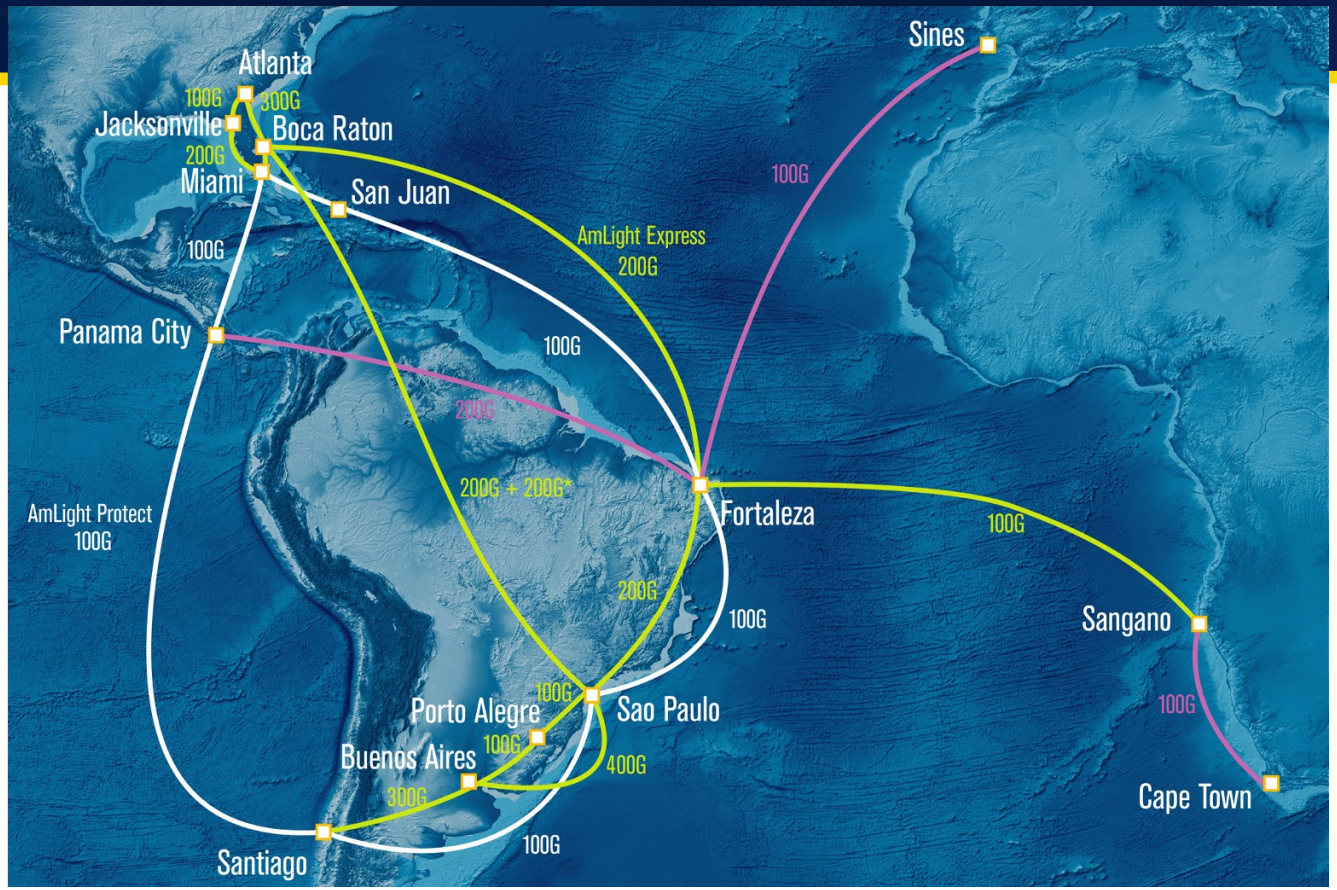

AmLight ExP
Americas Africa Lightpaths Express & Protect
(NSF Award # OAC-2029283)

AmLight ExP
Americas Africa Lightpaths Express & Protect

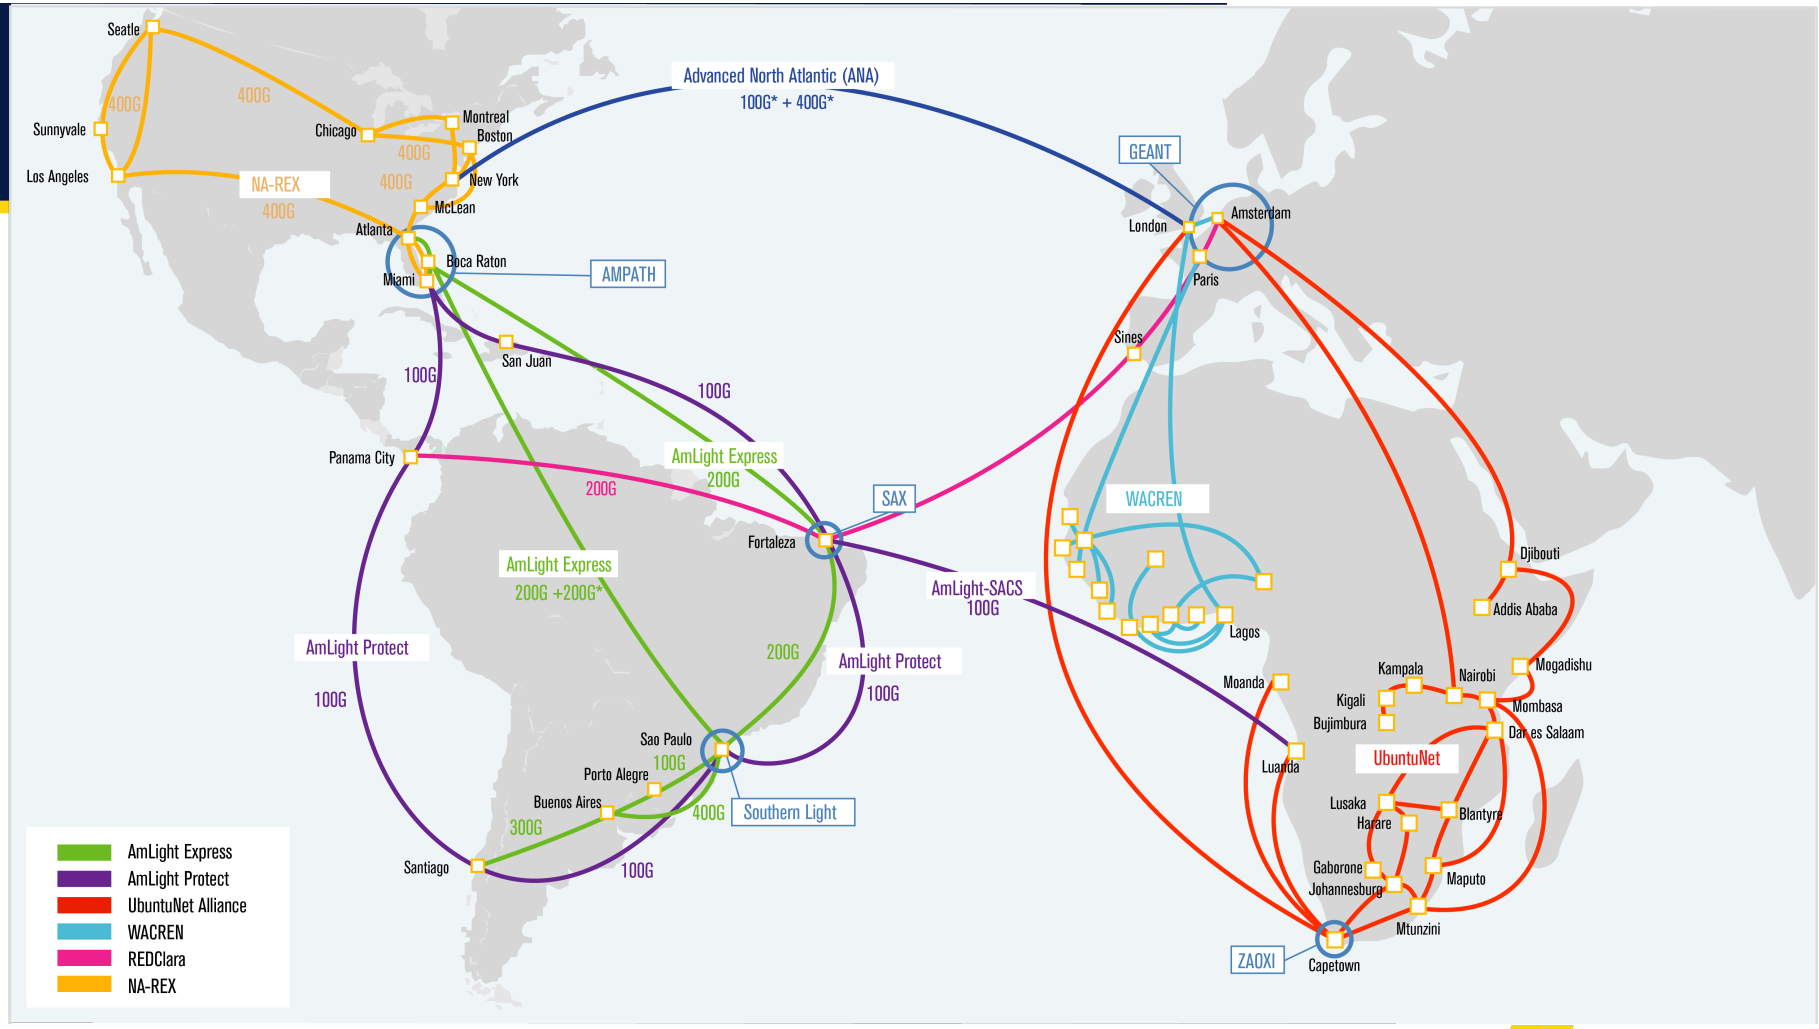

AmLight ExP
Americas Africa Lightpaths Express & Protect

Network performance during Supercomputing (SC23)

- Leveraged AmLight-ExP's 100Gbps infrastructure via South Atlantic Cable System (SACS)
- Demonstrated SANReN Data Transfer Nodes (DTNs)
- Achieved **2.5 Tbps data transmission** across the US, Brazil, and South Africa

Test Path: Miami to Cape Town

- Used 100Gbps SACS link (Miami - Fortaleza - Cape Town) - filled to +/- 98%.
- Achieved **70Gbps** throughput using iperf3, running 4 parallel 20Gbps streams with a 400MB TCP window

Collaboration Impact

- Highlighted **African NREN involvement** in global data transfer experiments
- Strengthened partnership between AmLight-ExP and SANReN

AmLight ExP
Americas Africa Lightpaths *Express & Protect*

Bridging Europe, Africa, and the Americas (BEAA) – New Collaboration for shared transoceanic R&E networking resilience

- Partnership between **seven regional/national R&E networks** across Europe, Africa, and Latin America
- Aimed at enhancing network resilience via **shared transoceanic links**

Memorandum of Understanding (MoU)

- Signed by AmLight/FIU, GÉANT, RedCLARA, RNP, CSIR, TENET, and the UbuntuNet Alliance
- **Initial 3-year agreement** for mutual backup and high-capacity resource sharing
- Ensures **back-up connectivity** for uninterrupted R&E collaborations

Impact

- Critical during outages, e.g., **March 2024 cable disruptions** off West Africa
- Enabled traffic migration to backup systems like **Google's Equiano** and **AmLight SACS academic link**

NRP NATIONAL RESEARCH PLATFORM

What is it?

- NRP is a partnership of more than 50 institutions, led by researchers and cyberinfrastructure professionals at UC San Diego

What does it do?

- Nautilus is a HyperCluster for running containerized Big Data Applications. It utilizes Kubernetes for managing and scaling containerized applications in conjunction with Rook for automating Ceph data services.
  https://docs.nationalresearchplatform.org/userdocs/jupyter/jupyterhub-service/

How is used?

- Login portal: https://portal.nrp-nautilus.io/
- Support:
  [matrix]: https://docs.nationalresearchplatform.org/userdocs/start/contact/
- Available Resources: https://portal.nrp-nautilus.io/resources

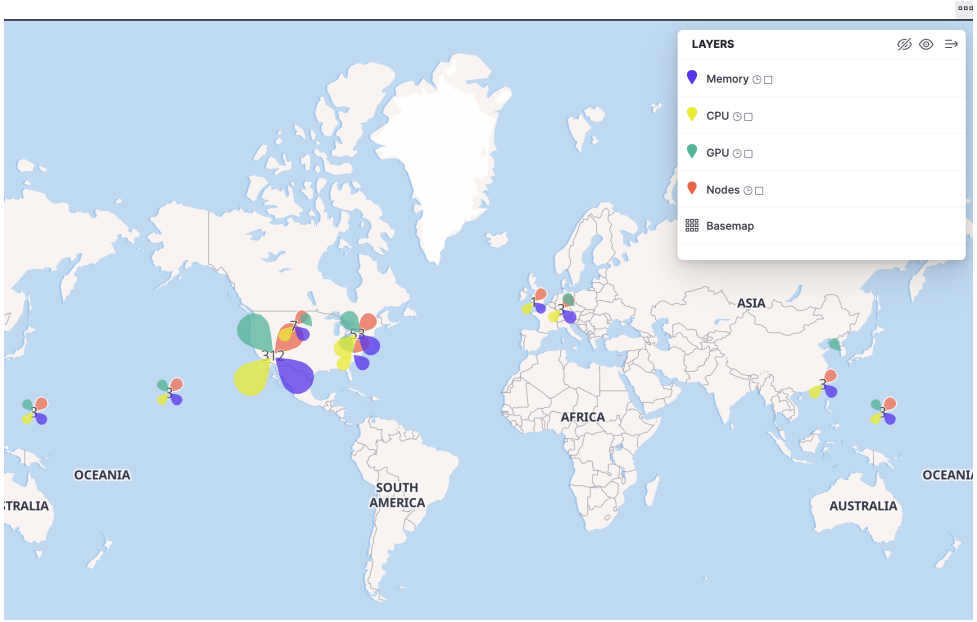

AmLight ExP
Americas Africa Lightpaths Express & Protect

- JupyterHub on Nautilus
- We provide the JupyterHub service running in our cluster, which is great if you need to quickly run your workflow and do not want to learn any kubernetes. Simply follow the above link (or https://jupyterhub-west.nrp-nautilus.io) and use your institutional credentials to login using CILogon. Choose the hardware specs to spawn your instance. Once authenticated you can run Jupyter notebooks as usual.
- Your persistent home folder initially will be limited to 5GB. If you need more, you can request it to be extended. You can also request for cephFS storage that is mounted to a shared disk space. Please use this to store all the data, code and results that you would need for long experiments.
- **NOTE:** your Jupyter container will shut down 1hr after your browser disconnects from it. If you need your job to keep running, don't close the browser window. You could either use a desktop with a persistent Internet connection or only use this for testing your code.
- **NOTE:** Available images are described in the scientific images section.
- If you need more, proceed to Step by Step Tensorflow with Jupyter.

NRP NATIONAL RESEARCH PLATFORM

# of Institutions	50+
# of Nodes	~300
# of GPUs	~1,100
# of CPU Cores	~10,000
Storage	~10PB
Processor Types	CPU, GPU, FPGA

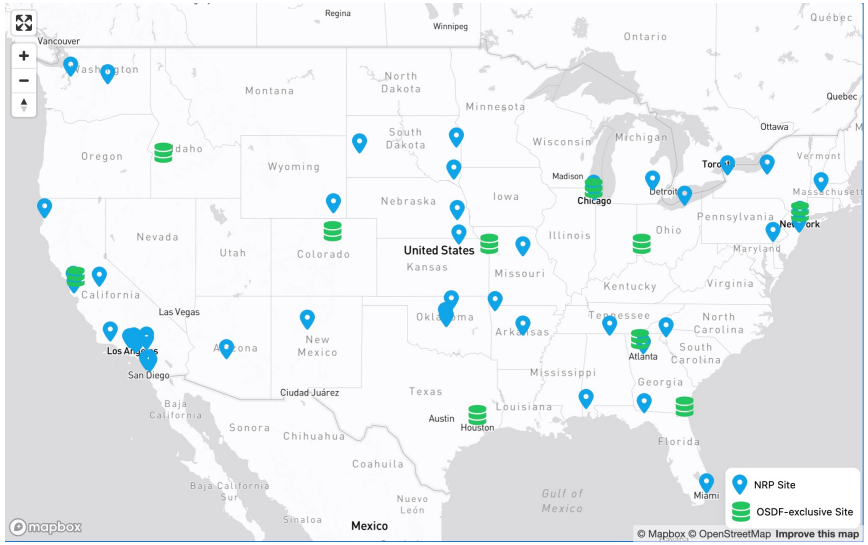

AmLight ExP
Americas Africa Lightpaths Express & Protect

Open Science Grid (OSG)

- What is it?
  - OSG Consortium operates a fabric of distributed High Throughput Computing (dHTC) services in support of the National Science & Engineering community
- What does it do?
  - **Open Science Pool**: Any researcher performing Open Science in the US can become an OSPool user (https://osg-htc.org/services/open_science_pool.html)
  - **Open Science Data Federation (OSDF)** (https://osg-htc.org/services/osdf.html) enables users and institutions to share data files and storage capacity, making them both accessible in dHTC environments such as the OSPool
- How is used?

  The OSPool provides its users with fair-share access (no allocation needed!) to processing and storage capacity contributed by university campuses, government-supported supercomputing institutions and research collaborations.

  Login: https://portal.osg-htc.org/application | Knowledge base: https://portal.osg-htc.org/documentation/

Open Science Grid (OSG)

OSG's Role

- Provides **software and services** to enable **opportunistic usage** of distributed computational resources
- Supports **35 institutions** across five continents through the **Open Science Data Federation (OSDF)**

Infrastructure

- **17 Origins** and **34 caches** globally contributing to the resource-sharing network

Challenges

- **Cache population delays** due to lack of data on **latency** and **geographical distance** between Caches and Origins
- Requires better insights into **round-trip time (RTT)** to improve efficiency

Open Science Data Federation (OSDF): Expanding Impact and Global Collaboration

- OSDF enables users and institutions to share data files and storage capacity
- Provides campuses and researchers with the ability to manage their data files, input and output, in support of running their dHTC workloads.
- Improves file access performance, resource consumption and reliability.
- OSG-Operated Access Points provide researchers with a default of 500GB of storage space on the OSDF.

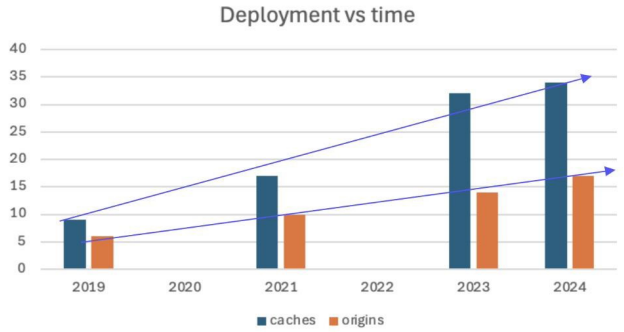

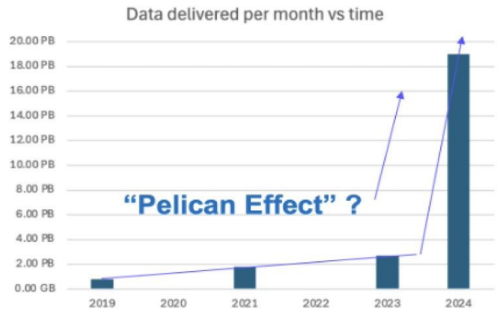

7x Increase in OSDF Usage

- Caching system saves **75% of 100G transnational network capacity**
- Now utilized by **one-third of OSPool users**, accounting for **10% of total reads**

Top Fields Using OSDF

- Biology, Physics, Math, Chemistry, Geological & Earth Sciences

Global Expansion: Discussions underway to deploy OSG caches in South America and Africa

AmLight ExP
Americas Africa Lightpaths *Express & Protect*

- What is it?
  - ACCESS is an advanced computing and data resource program
- What does it do?
  - Provides access to computing, data analysis, or storage resources
- How is used?
  - The user needs an ACCESS project and some resource units you can spend. The resource units are called Allocation (https://allocations.access-ci.org/get-your-first-project)
- Resource Type
  - Cloud, GPU Compute, Innovative / Novel Compute, CPU Compute, Storage and other services
  - Specialized Hardware: Composable hardware fabric and Large Memory Nodes
  - Specialized Support: Advance reservations, Preemption, Science Gateway support, ACCESS Pegasus

Allocation info https://allocations.access-ci.org/resources | Login: https://allocations.access-ci.org/login

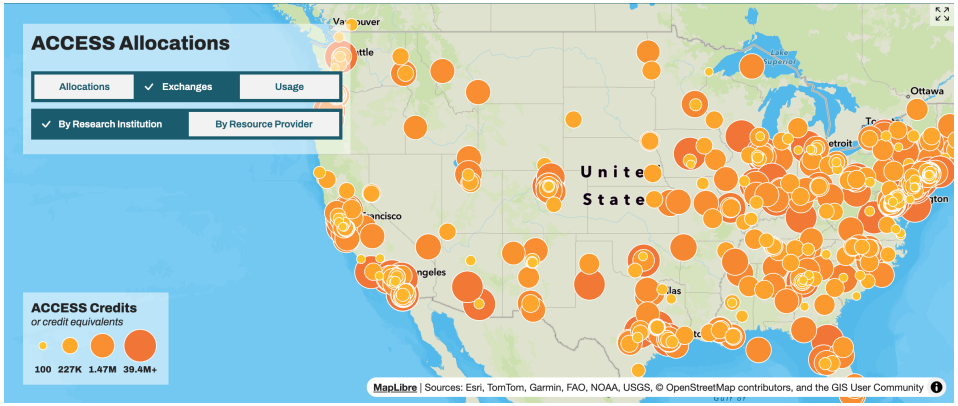

How to request allocation https://allocations.access-ci.org/how-to

Comparison Table

Project Type	Explore	Discover	Accelerate	Maximize
ACCESS Credits*	400,000	1,500,000	3,000,000	Awarded in resource units
Project duration	Supporting grant duration or 12 months	Supporting grant duration or 12 months	Supporting grant duration or 12 months	12 months
Requests accepted	Anytime	Anytime	Anytime	Every 6 months
	Multiple requests allowed	Multiple requests allowed	Multiple requests allowed	1 allowed (some exceptions)
Requirements and review process	Overview	1-page proposal	3-page proposal (max. length)	10-page proposal (max. length)
	Confirmation of eligibility and suitability of requested resources	Confirmation of eligibility and suitability of requested resources	Panel merit review	Panel merit review

*For more information see **ACCESS Credit Exchange & Calculator** page.

# of Institutions	15
# of Nodes	~8600
# of GPUs	~2300
# of CPU Cores	~850,000
Storage	~192PB
Processor Types	CPU, GPU, FPGA

FABRIC

Testbed Resource Summary							
Cores	Disk(GB)	RAM(GB)	GPU	NVME	SmartNIC	SharedNIC	FPGA
81799 / 88192	1944073 / 2038123	43466 / 62596	128 / 173	361 / 370	126 / 157	17621 / 18669	15 / 23

What is it?

- FABRIC is Adaptive ProgrammaBle Research Infrastructure for Computer Science and Science Applications

What does it do?

- FABRIC is an International infrastructure that enables experimentation and research at-scale in the areas of networking, cybersecurity, distributed computing, storage, virtual reality, 5G, machine learning, and science applications

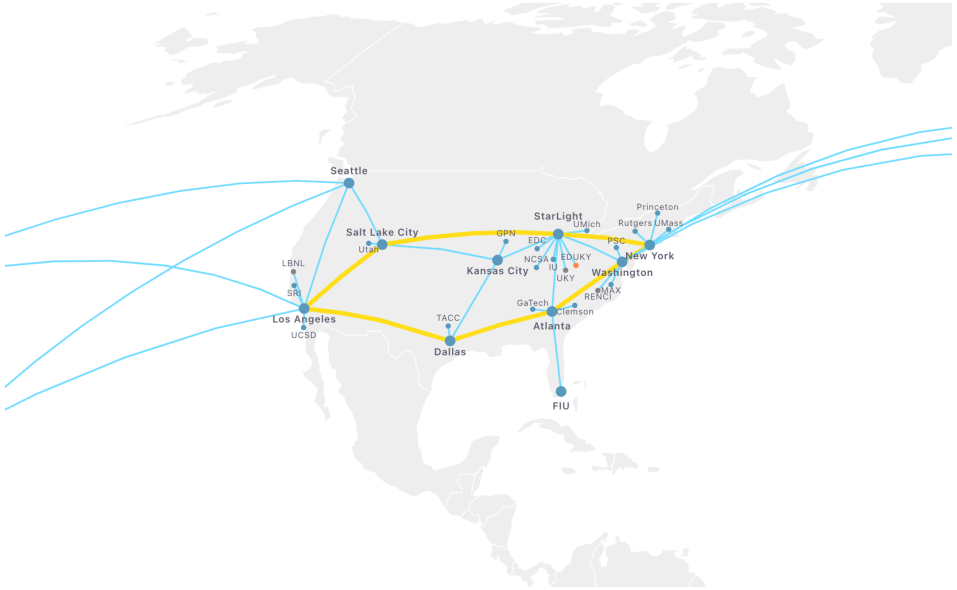

DALL (Dallas) Active

Resource	available / total
Cores	162/384
Disk (GB)	60731/64161
RAM (GB)	882/1434
GPU	5/5
NVME	10/10
SmartNIC	4/4
SharedNIC	362/381
FPGA	0/1

How is used?

- FABRIC is an International infrastructure that enables experimentation and research at-scale in the areas of networking, cybersecurity, distributed computing, storage, virtual reality, 5G, machine learning, and science applications

AmLight ExP
Americas Africa Lightpaths Express & Protect

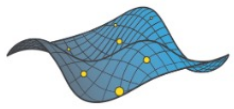

- Each of the 29 FABRIC sites features:
  - Large amounts of compute and storage
  - Interconnected by high-speed, dedicated optical links
- Additionally, it connects to:
  - Specialized testbeds (5G/IoT PAWR, NSF Clouds)
  - The Internet
  - High-performance computing facilities
- This setup creates a rich environment for a wide variety of experimental activities.
- FABRIC Across Borders (FAB) expands the network by adding 4 additional nodes in Asia and Europe and via AmLight to South America.

Login portal: https://portal.fabric-testbed.net/ | Knowledge base: https://learn.fabric-testbed.net/article-categories/technical-guides/ | Shared resources: https://portal.fabric-testbed.net/resources/all

# of Institutions	39
# of Nodes	144
# of GPUs	173
# of CPU Cores	88192
Storage	~2PB
Processor Types	CPU, GPU, FPGA, NIC

Getting started with FABRIC

- Prior to using FABRIC please familiarize yourself with the 'Getting Started' section of our documentation site: https://learn.fabric-testbed.net/article-categories/getting-started/
- Particularly pay attention to the following articles:
  - https://learn.fabric-testbed.net/knowledge-base/things-to-know-when-using-fabric-for-the-first-time/
  - https://learn.fabric-testbed.net/knowledge-base/fabric-user-roles-and-project-permissions/
- Please read the Project Lead policy and make sure the users you add are familiar with the Acceptable Use Policy: https://learn.fabric-testbed.net/design-documents/

- **Remember that you are responsible not just for your own conduct, but for the conduct of everyone you add to your projects.**
  Only add into your projects people you are personally familiar with.

Additional ML/AI research

INDIS SC24: Leveraging In-band Network Telemetry for Automated DDoS Detection in Production Programmable Networks: The AmLight Use Case

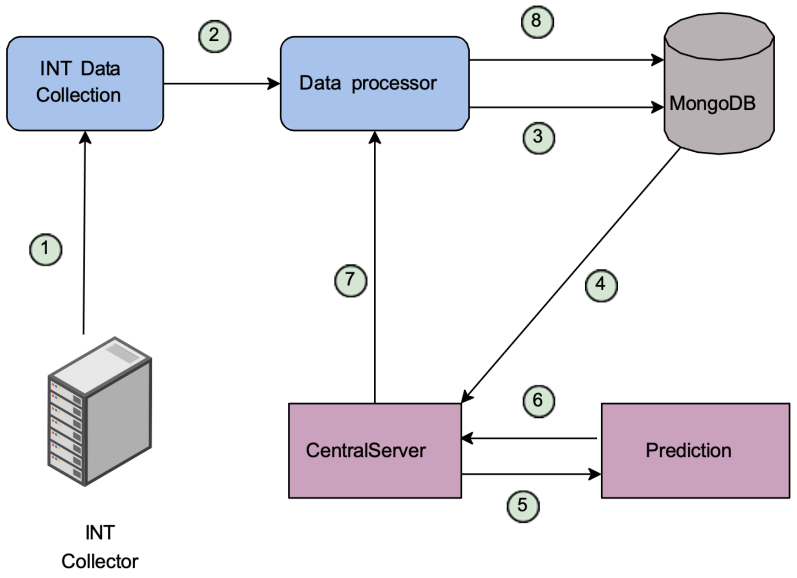

Automated DDoS Detection Proposed Mechanism

1. Gather INT data.
2. Send INT data to the Data processor:
   a. Flow ID: src/dst IP, src/dst ports, protocol.
   b. Flow-level features (e.g., Packets per second, Flows per second).
3. Save processed data to the database.
4. Retrieve processed data.
5. Send data to the prediction model.
6. Receive predictions.
7. Send predictions to the Data processor for aggregation.

AmLight ExP
Americas Africa Lightpaths Express & Protect

Experimental Results

- We achieved over 97% accuracy in predicting most attack types, with an average response time of under 2 seconds.
- The creation of new flows appears to introduce bottlenecks and increase prediction time.

Attack Type	Accuracy	Misclassified/ Number of Predicted Packets	Average Prediction Time (s)	Max Prediction Time (s)
UDP Scan	0.9947	14/2628	0.12	0.73
SYN Scan	0.9961	10/2542	0.44	1.81
SYN Flood	0.9984	27/2814	0.09	0.4
SlowLoris	0.9795	16/779	0.05	130.85
Benign	0.9417	136/2331	103.14	734.55*

- INT data proved effective in detecting DDoS attacks for both known and novel attack patterns.
- sFlow performs similarly but may miss data due to its sampling approach.
- Automated detection, addressing bottlenecks, can be achieved in under 2 seconds.
- Efficiently storing, processing, and analyzing INT data requires substantial computational resources and optimized techniques.
- Establishing precise timestamps remains challenging.
- With our network capacity of 100 Gbps, the simulated attack did not cause significant congestion, limiting our ability to observe the effects on queue occupancy.

Conclusion

Research and Educational Networks and Exchange Points can interconnect a wide range of resources to support Machine Learning and Artificial Intelligence research in higher education, as well as foster international collaborations.

THANK YOU

vchergar@fiu.edu

More about AmLight collaboration with Africa:
Come to hear about latest updates at the:
Caltech booth 845 – TBD
2.40PM - 3:00PM at the Illuminations Pavilion -Title: Ecosystems project - South Africa, Presenter: Lara Timm
4:40PM -5:00PM at the SCinet Theater 2049 -Title: High-Speed data transfers from South Africa to USA! Presenter: Kasandra Pillay