


Performance Testing of High-Speed Scientific Networks in Preparation for the Exascale Era



Fermilab

Lael Verace¹, Syed Asif Shah², Oliver Gutsche², Philip DeMar², Kevin Black¹

¹University of Wisconsin-Madison, WI, USA; ²Fermi National Accelerator Laboratory, IL, USA



Background

As the High-Luminosity Large Hadron Collider (HL-LHC) era approaches, scientific experiments are expected to generate data at unprecedented scales. To support this, terabit-scale high-speed networks will be critical for transferring data efficiently between global computing facilities. For instance, Fermilab—a CMS Tier-1 site—is currently connected to a high-performance scientific network with a capacity of 1.6 Tbps. Ensuring the reliability and readiness of such infrastructure requires robust performance testing under peak load conditions.




Figure 1: LHC luminosity is due for significant increases after LS3 and LS4

Research Questions



Are data transfers between CMS computing sites performing at expected speeds under full load?



How can we effectively saturate network links to test throughput and identify bottlenecks in a simulated HL-LHC environment?

Methodology

We designed a **distributed load generation framework** that deploys lightweight software containers across computing sites. These containers generate and monitor network traffic in parallel, simulating realistic data transfer scenarios. A **centralized orchestrator**, currently hosted at Fermilab, dynamically manages container deployment and test coordination. Network traffic is aggregated from multiple sites to saturate links and measure throughput accurately.

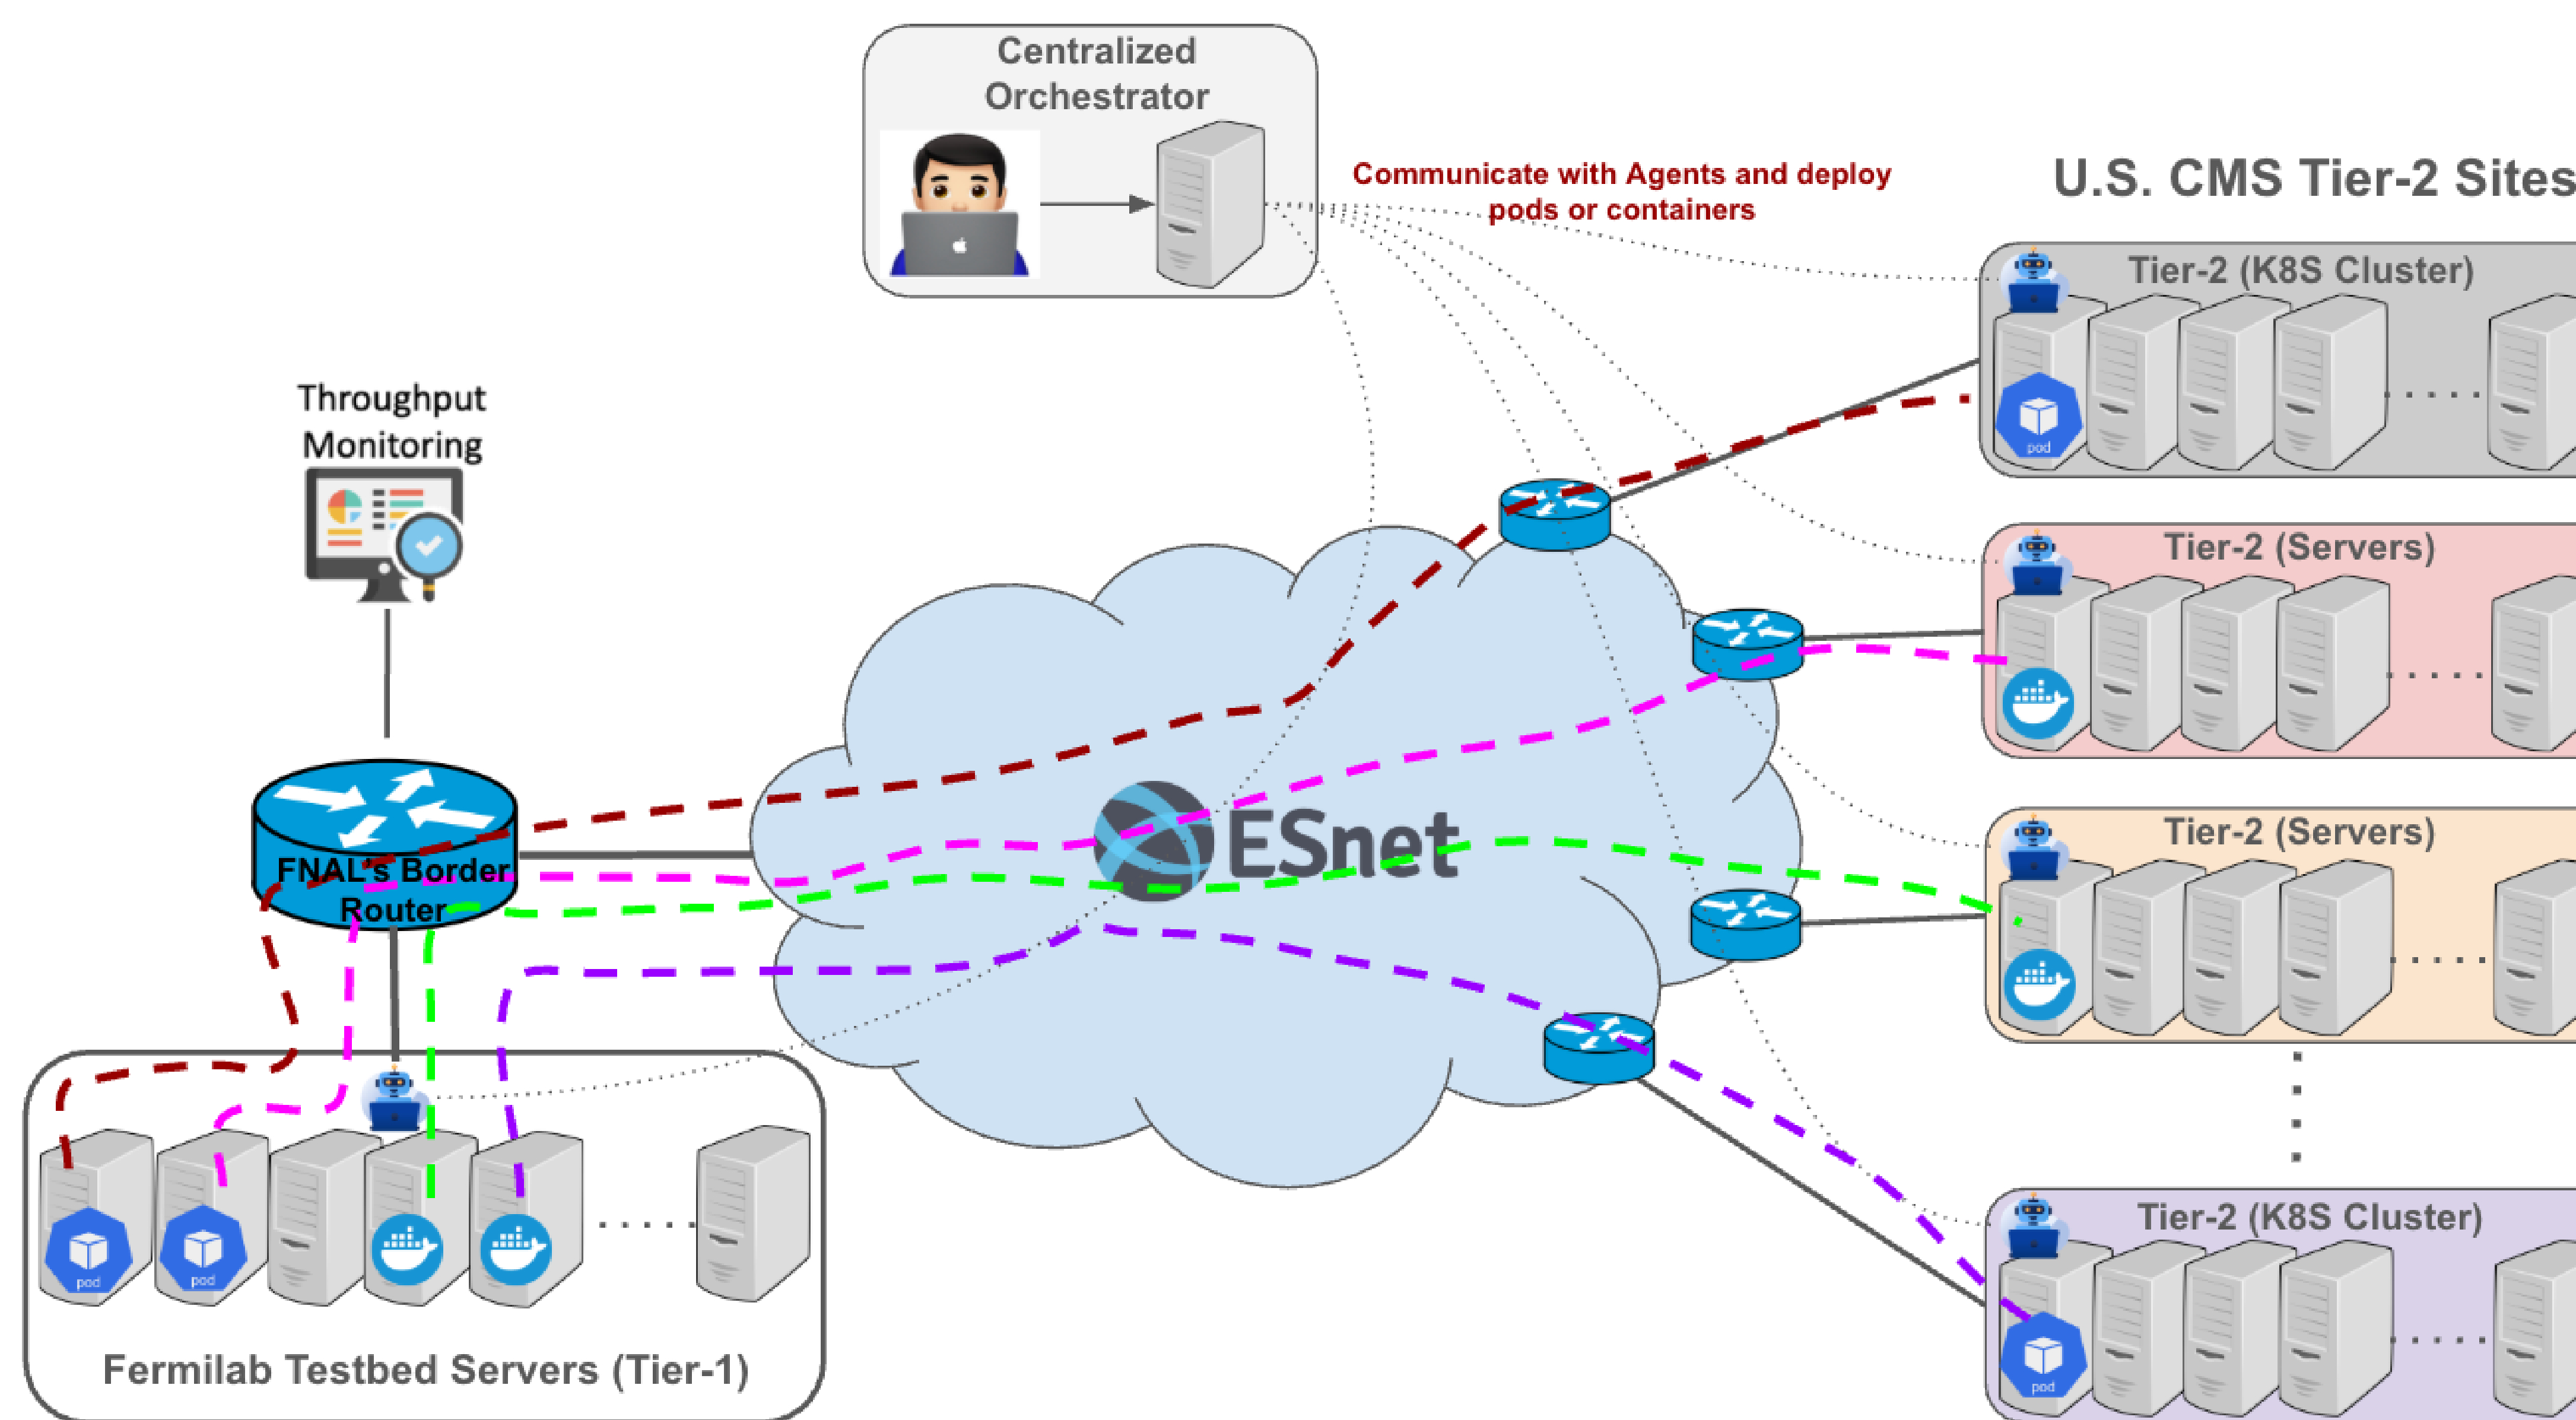


Figure 2: Diagram of the infrastructure. Commands containing test parameters are sent to the centralized orchestrator, which then interacts with server agents (🤖 icons) via an API to generate traffic and complete tests. The agents are specialized to run on Kubernetes clusters (📦 icons) or standard servers running Docker/Podman (🐳 icons)

Preliminary Results

The infrastructure has successfully completed its initial round of tests

We tested the methodology of our concept in a controlled environment at Fermilab over a 300 Gb/s connection, treating one server cluster as a Tier-2 site and the other as the Tier-1 site.




Figure 3: Aggregated traffic from 3 server pairs allows the 300 Gb/s connection between sites to be tested at full load.

This test showcases the ability of our framework to successfully saturate the network and monitor performance

Conclusion & Future Work

We have demonstrated a scalable and dynamic approach for high-volume network testing across CMS computing infrastructure.

Next steps include:

- Expanding support for Tier-2 site integration
- Enhancing the orchestrator's user interface
- Automating diagnostics for real-time anomaly detection

Acknowledgement

This project is supported by the Department of Energy award number DE-SC0023528