

USCMS Pre-DC24 Challenge Fall 2024

Garhan Attebury, Diego Davila, Asif Shah, Preeti Bhat

Data Challenge 2024 Workshop at CERN
November 9-10, 2023

USCMS Pre-DC24 Challenge Overview

Purpose: Conduct mini-data challenge to understand tooling and current functionality in preparation for February 2024 data challenge. Not a scale test but more of a functionality test.

Dates: Oct 30th through Nov 13th 2023, weekends and holidays excluded

Target data rate: 25Gbps for FNAL (as opposed to 250Gbps), 10Gbps for T2 sites - limited by request of ESnet to not saturate interchange points

Miscellaneous: No testing of AAA, no correlation with link quality

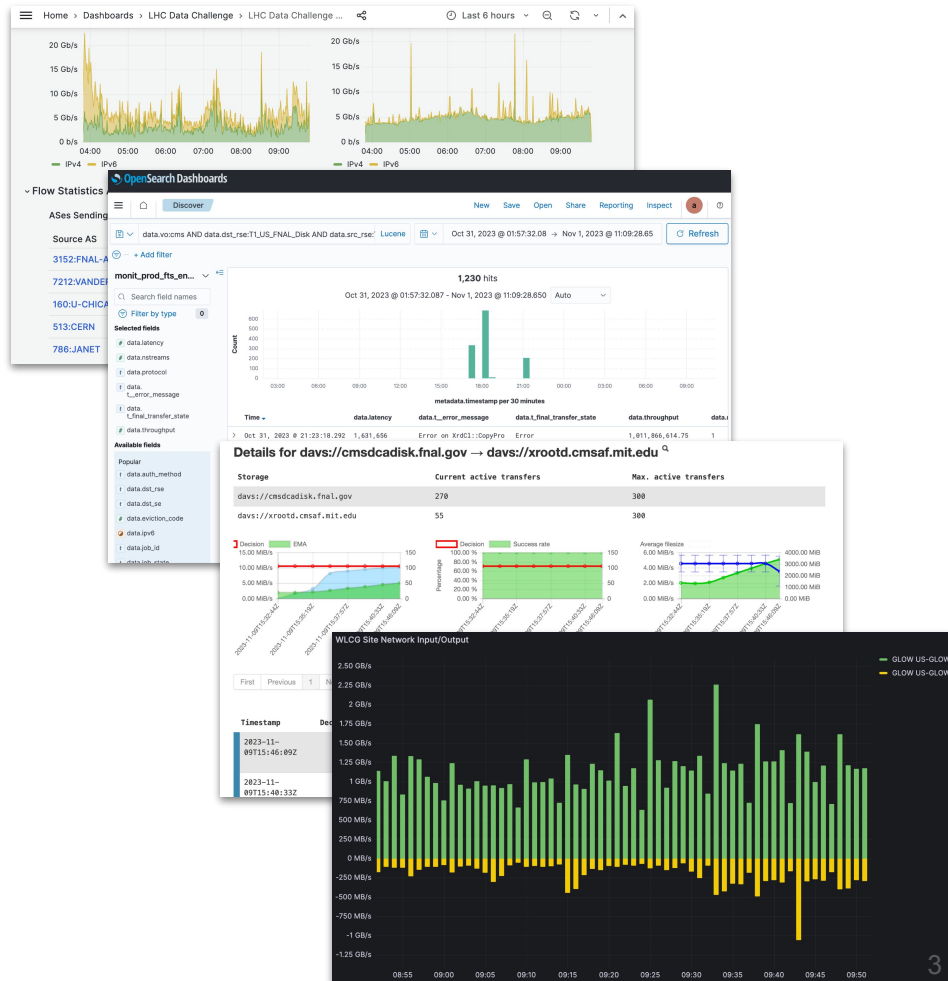
More details in [USCMS DC24 Rehearsal report - v1](#)

Tooling

- [Rucio injection tool](https://gitlab.cern.ch/atlas-adc-ddm/dc_inject)
https://gitlab.cern.ch/atlas-adc-ddm/dc_inject

Monitoring

- [Stardust](#) - ESnet peering SNMP
- [MONIT](#)
- [FTS \(CERN\)](#) and [FTS \(FNAL\)](#)
- [WLCG Site Network Monitoring](#)
- [DC FTS Dashboard](#)



Schedule

Day	AM	PM
Oct 30th	FNAL → Caltech	Caltech → FNAL
Oct 31st	FNAL → Florida	Florida → FNAL
Nov 1st	FNAL → MIT	MIT → FNAL
Nov 2nd	FNAL → Nebraska	Nebraska → FNAL
Nov 3rd	FNAL → Purdue	Purdue → FNAL
Nov 6th	FNAL → All Tier2	All Tier2 → FNAL
Nov 7th	FNAL → Vanderbilt	Vanderbilt → FNAL
Nov 8th	FNAL → Wisconsin	Wisconsin → FNAL
Nov 9th	FNAL → SPRACE	SPRACE → FNAL
Nov 13th	FNAL → UCSD	UCSD → FNAL


Results

Rates are Gbps

Day	AM	PM	Write Rate (Avg/Max)	Read Rate (Avg/Max)	Observations
Oct 30th	FNAL → Caltech	Caltech → FNAL	45/60	70/100	Higher data rate by mistake
Oct 31st	FNAL → Florida	Florida → FNAL	12/20	45/50	Increased rate to 20Gbps, no affect due to concurrent FTS transfer limits
Nov 1st	FNAL → MIT	MIT → FNAL	17/22	12/37	Site loaded before test, no real impact
Nov 2nd	FNAL → Nebraska	Nebraska → FNAL	50/100	40/120	Increased rate to 25Gbps, transfers always ended before next injection
Nov 3rd	FNAL → Purdue	Purdue → FNAL	45/75	60/130	Transfers ended before next injection
Nov 6th	FNAL → All Tier2	All Tier2 → FNAL	90/210	50/205	Transfers ended before next injection
Nov 7th	FNAL → Vanderbilt	Vanderbilt → FNAL	15/40	25*/27	FTS limit and tool issue affected reads
Nov 8th	FNAL → Wisconsin	Wisconsin → FNAL	/70	/58	
Nov 9th	FNAL → SPRACE	SPRACE → FNAL			
Nov 13th	FNAL → UCSD	UCSD → FNAL			

FNAL → AII


> 200Gbps outgoing from FNAL
with ~50Gbps background



[Stardust FNAL monitoring link](#)

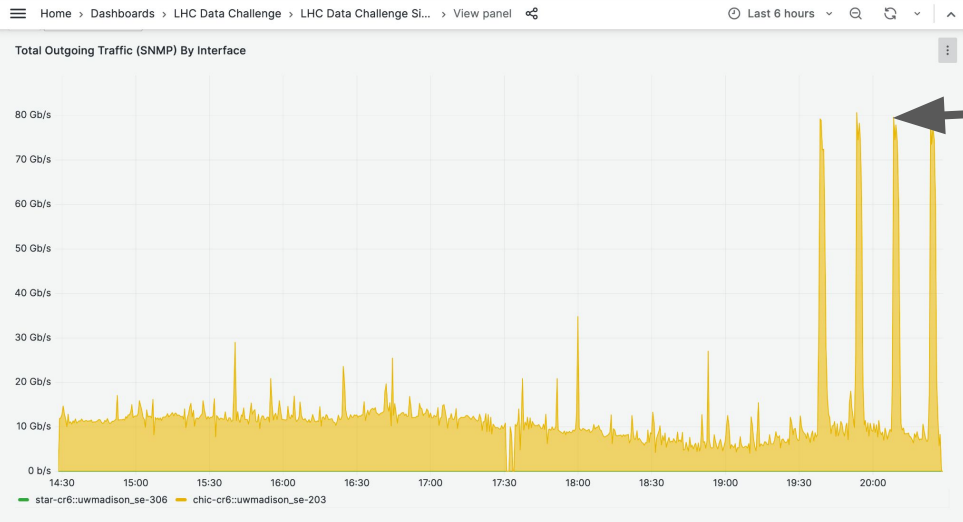

FNAL ← AII

> 200Gbps into FNAL from T2s



T1_US_FNAL → T2_US_Wisconsin

800 GB / 5 minute window
equates to ~21Gbps




80Gbps seen on ESnet interfaces

What question are we asking?

- *Network operator view*
- *Experiment (disk to disk) view*

We care about both!

T1_US_FNAL ↔ T2_US_Wisconsin



[WLCG Site Network Link](#)

[MONIT FTS Link](#)

FTS / Rucio Optimizations

- Limit too low to strain sites (200 transfers)
- Scaling
 - Running with `-big-first` results in low number of rules @ 25Gbps
 - Running without can result in too many rules generated
 - Practical limits of FTS when doing All $\leftarrow \rightarrow$ All testing?
- Utilize REST API to modify limits on the fly

Steps between now and DC24 Challenge in February

- Network quality visibility and improvements
 - PerfSONAR upgrade campaign
 - New perfSONAR and DTNs for most sites, 100Gb capability improvement
- FTS improvements
 - Reliability of transfers (reduce failure rates, easier diagnosing of failures)
 - Wishlist: Summarization of failure reasons per link by day/hour/etc
 - Scale / efficiency testing of FTS
 - Automation with REST API (responsiveness during challenge)
- Monitoring improvements
 - Instrument all sites with WLCG SNMP monitoring
 - Hard to see site-to-site traffic, even harder for individual domain flows
 - SENSE could enable this