



# Services & Tools

Vincent Garonne, <[vincent.garonne@bnl.gov](mailto:vincent.garonne@bnl.gov)>  
*Scientific Data and Computing Center (SDCC)*

ATLAS Pre-Scrubbing Review - June 27, 2022



@BrookhavenLab

# Services & Tools — Overview

- 9 staff members
  - 7 of whom are partially funded by ATLAS \*
  - FY23: 2.95 FTEs in WBS 2.3.1
- Responsible for these ATLAS services:
  - **Central Storage Software**
  - **Tape Archive Software interface**
  - **Data transfer**
  - **Web front-ends** — USATLAS.org website
  - **Analysis Facilities Operation & Continuous Integration and Operations** (WBS 2.3.4 & 2.3.5) — Cf. Ofer's Talks

## Services & Tools

Vincent Garonne\* - Group Lead

### **Data/Storage management**

Carlos Gamboa \*

Qiulan Huang \*

Shigeki Misawa \*

Hironori Ito \*

Matt Snyder \*

Ofer Rind \*

### **General Web Services**

Christian Lepore \*

Dmitry Arkhipkin

**New hire**




# Central Storage at SDCC

C. Gamboa, Q. Huang, V. Garonne

- Biggest ATLAS Tier-1 storage service
  - Disk: 27PB
  - Tape: 50PB
  - 30 GB/s aggregated throughput
  - Peak at 300Hz of deletion rate


- A good and well-understood service consumer ↔ service provider relationship between ATLAS and BNL

- 24|7 service with storage operator on call
- WLCG Availability & Reliability > 99%
- Pledge 2022



- Continuous delivery & improvements

- Latest dCache golden release (7.2.\*)
- Regular interactions with dCache team with releases addressing our issues (e.g., deletion)



ATLAS data at BNL over time



BNL-ATLAS Avail: 100% Unkn: 0%

# Network+Disk data challenge status

**2021 target:** 10% of HL-LHC in 2027 and the highest rate for special periods of Run-3

During the last period:

**In: 50 Gb/s**


With **25 Gb/s from CERN**

**Out: 60 Gb/s**

with **20 Gb/s to T1s**

**40 Gb/s to T2s**

Transfers throughput for BNL (all activities)



The target has been reached with no special storage and network issue

# Tape related activities for ATLAS

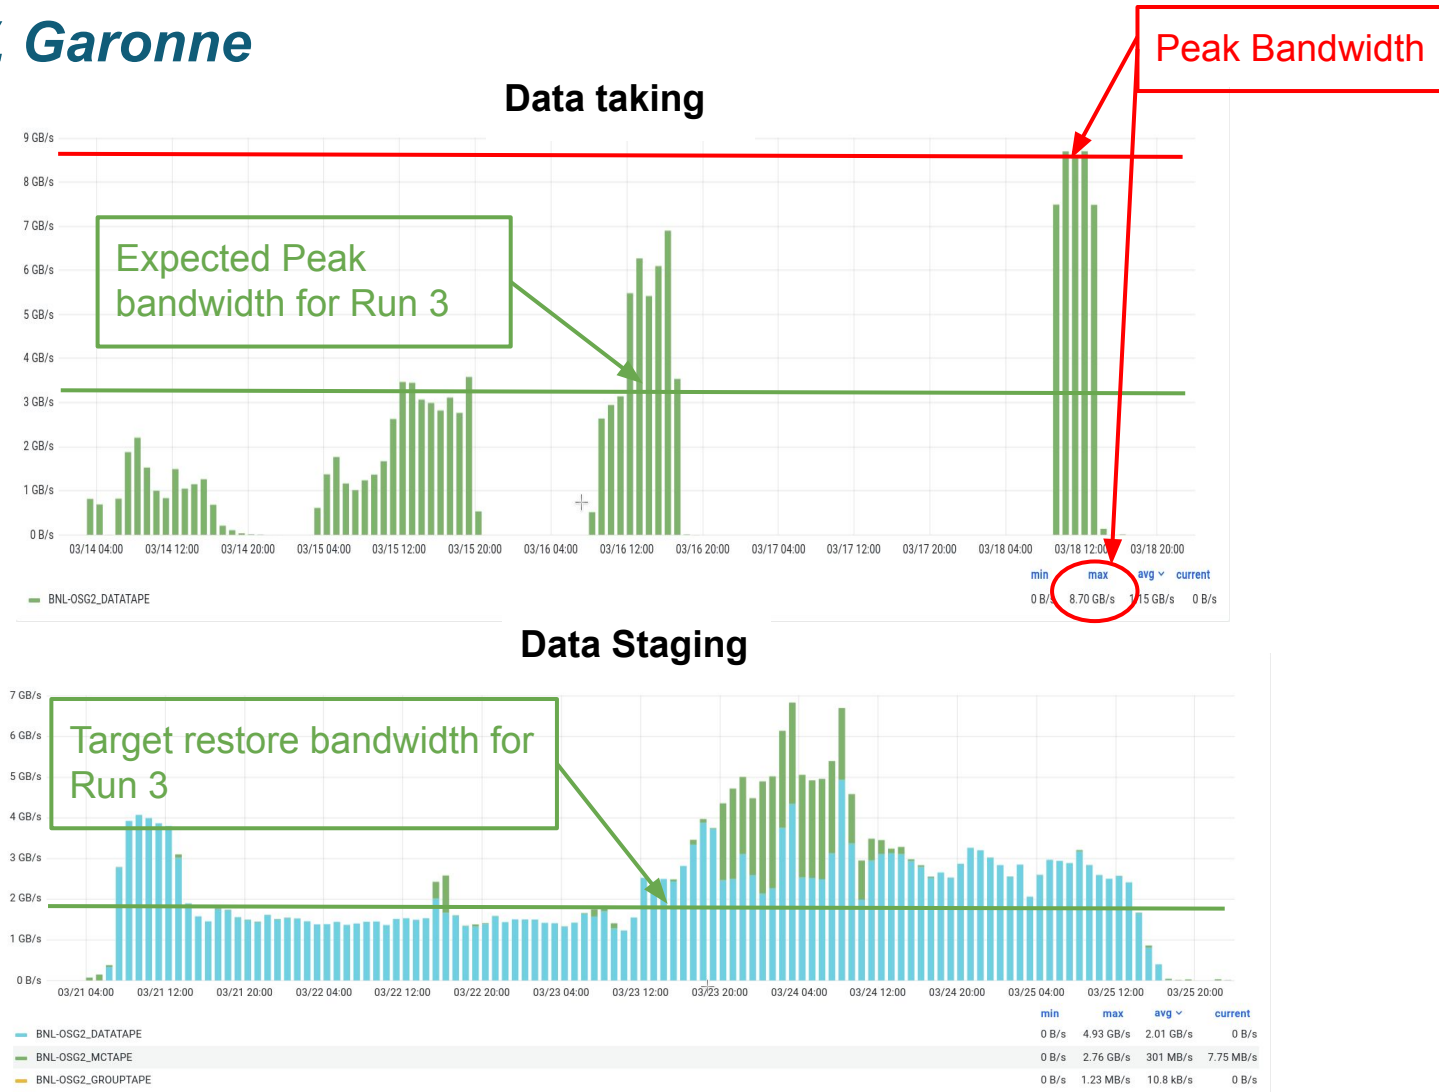
*C. Gamboa, Q. Huang, S. Misawa, V. Garonne*

Continuous improvements in dCache tape interface

- Development/Adaptation of new “glue code” connecting dCache to HPSS
- Commissioned and deployed for the tape challenge
- Great improvement of system scalability
  - 80k → 200K concurrent staging requests

Active member of the WLCG and ATLAS tape working groups

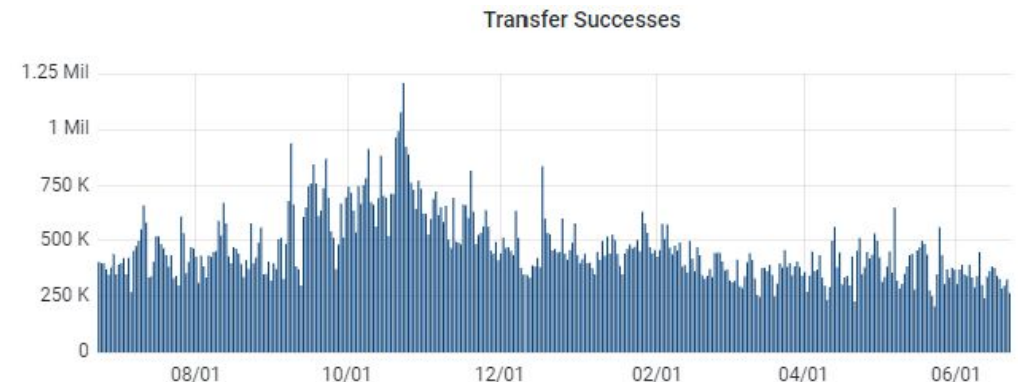

- 4 talks given in the last 10 months



# Data transfer service

*H. Ito, M. Snyder*


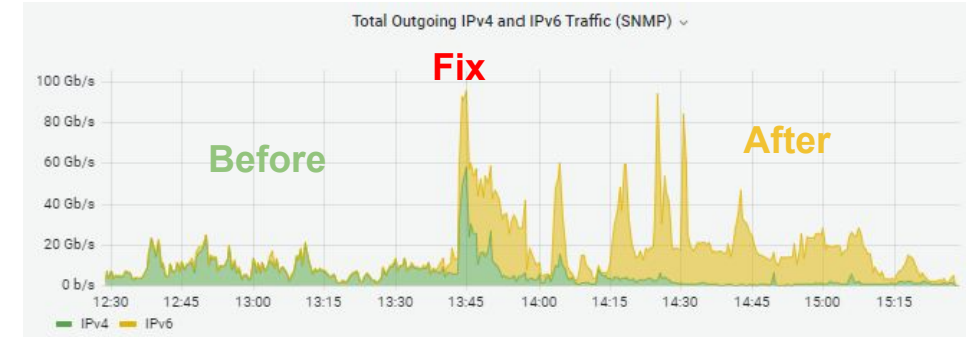
- BNL has an operational responsibility for the ATLAS data transfer service (FTS)
  - FTS@BNL is upgraded to each major release (3.11)
- The BNL instance serves one third of the total ATLAS traffic  
10GB/s per day, 1M transfer per day
  - Service availability > 99 %
- BNL has been used for entire ATLAS when CERN FTS went down for extended period



# IPv6 and Network Activity

H. Ito

- IPv6 Issue
  - ATLAS's IPv6 data volume started to decrease since the last summer.
  - BNL with the help of the dCache developers identified the cause of the issue.
    - Cause: Java's default network priority setting
  - BNL has tested the suggested fix provided by the developers.
  - BNL has broadcasted the fix to WLCG and HEPIX community.
- PerfSonar at 100 Gbps
  - One of handful of sites with deployment
- Packet Marking
  - Flowd was installed at all doors



# RUCIO QoS


*H. Ito, M. Snyder*

- RUCIO Quality of Service (QoS) provides ability to stage files from the tape storage to disk cache and pin them for desired durations according to RUCIO Replica Rules
- RUCIO will issue bring-online for all QoS files whenever accessing them regardless of the status of the files
- With the guaranteed pin times, the disk cache can act as a regular storage endpoint
- QoS storage will be more efficient than the current copy-to-data disk model.
  - No need to copy to data disk.
  - No need to delete from data disk.
- RUCIO QoS will/should work with PANDA and Data Carousel without any major changes
- Contributions to Rucio code source and LTS releases
- [Talk at the last ATLAS software week](#)



# USATLAS Website

*Ch. Lepore, V. Garonne*



SDCC has a strong expertise in supporting websites (~6 drupal websites) and is responsible for [USATLAS.org](https://www.usatlas.org)

- Providing support, guidance and cybersecurity expertise to the US ATLAS collaboration
- Following up closely the activity (gap analysis) with DataArt, an IT consultancy company, responsible for enhancing the capabilities of the website to prepare the long term service support
- **USATLAS.org Drupal website deployment is completed**

# Future Activities & Synergies with ADC

Increasing need for analytics et instrumentation of data storage, E.g., Tape recall optimisation

- Timeline: Q4 2023

Storage evolution — Cf. Future storage dedicated talk

- Timeline: Q4 2022

Inclusion within [ATLAS/ADC](#)

- Conduct an ["ATLAS Grid Sites" survey](#) for ADC
  - 28 questions / 31 questionnaire answers
  - [Summary and analysis](#) delivered at the last software week
- [ATLAS ADC activities](#) of interests:
  - "DC-2: Storage evolution" — Timeline: Q4 2025
  - "DC-8: Storage optimisation" — Timeline: Q4 2025
  - "DC-9: Disk management: secondary (cached) datasets" — Timeline: Q2 2023
  - "DC-11: Optimising the user analysis experience" — Timeline: Q4 2023
  - Analytics working group
  - Dynamic Data Handling Task Force