

# StoRM towards DC 2024

Enrico Vianello  
INFN-CNAF

Data Challenge 2024 Workshop  
November 9-10th, 2023  
[enrico.vianello@cnafe.infn.it](mailto:enrico.vianello@cnafe.infn.it)



# Topics

- StoRM status
  - Components and latest updates
  - Token-based AuthN/Z support
  - Monitoring logs & metrics
- StoRM deployments and monitoring @ CNAF-T1
  - Deployments summary
  - Token support and open points
  - Monitoring alarms & metrics
  - Monitoring from log files
- StoRM evolution
  - Objectives on our roadmap
  - Plans for DC 24

# StoRM status

# StoRM - STOrage Resource Manager

StoRM main components are:

- StoRM Backend + StoRM Frontend
  - SRM disk and tape file management
- StoRM WebDAV
  - HTTP/HTTPS file management and transfer
- StoRM Tape REST API
  - HTTP/HTTPS tape file management
- StoRM GridFTP
  - GridFTP file transfer

Platforms:

- CentOS 7
- RedHat Enterprise 9 (next)
  - Using UBI docker images to build rpms

Provides flexible AuthN/Z support:

- VOMS & OAuth/OIDC
- File access control is enforced via POSIX ACLs

**Latest release:**

StoRM v1.11.22

(2023-06-21)

[release-notes](#)

available on

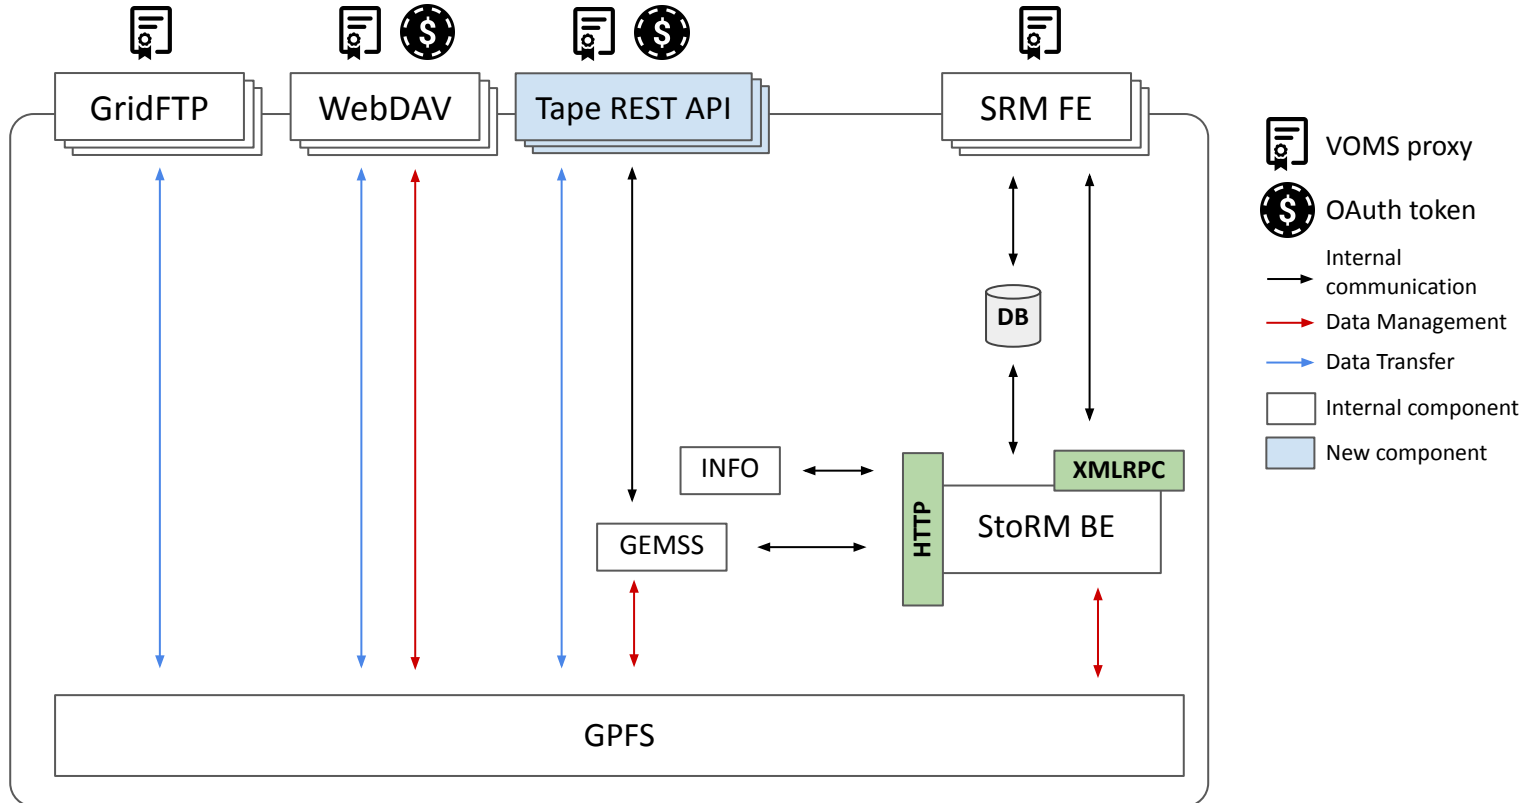
StoRM [stable repo](#)

or UMD repositories

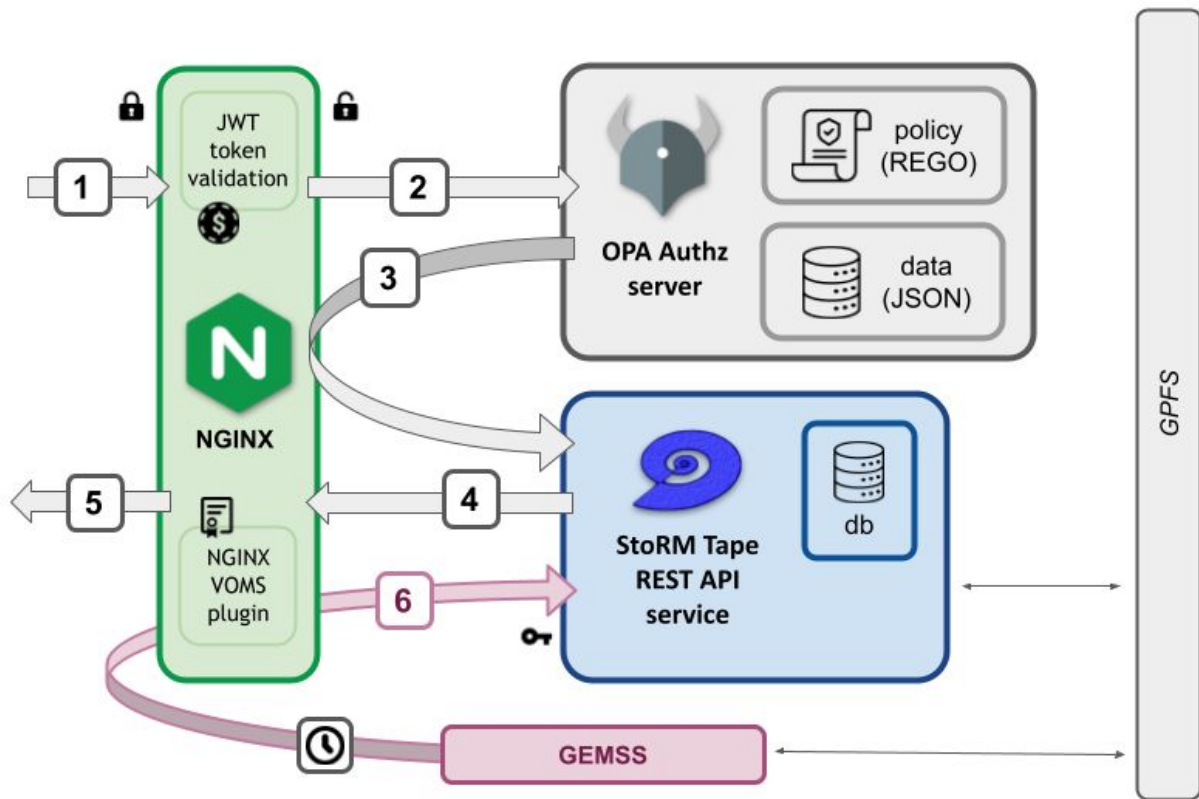
Name	Latest version
<a href="#">StoRM Backend</a>	<a href="#">1.11.22</a>
<a href="#">StoRM Frontend</a>	<a href="#">1.8.15</a>
<a href="#">StoRM WebDAV</a>	<a href="#">1.4.2</a>
<a href="#">StoRM Native Libs</a>	<a href="#">1.0.7</a>
<a href="#">StoRM Info Provider</a>	<a href="#">1.8.4</a>
<a href="#">StoRM SRM Client</a>	<a href="#">1.6.1</a>
<a href="#">StoRM GridFTP</a>	<a href="#">1.2.4</a>
<a href="#">StoRM XMLRPC-C</a>	<a href="#">1.39.12</a>
<a href="#">StoRM Utils</a>	<a href="#">1.0.0</a>
<a href="#">StoRM Puppet module</a>	<a href="#">4.1.0</a>
<a href="#">StoRM Tape REST API</a>	<a href="#">0.4.0</a>

recently released

# StoRM: current deployment architecture



# StoRM Tape REST API



The StoRM implementation of [WLCG tape REST API spec](#)

Supports bulk-request of tape-stored files

It can coexist with a typical SRM deployment

Accessed via both X509 + VOMS (proxy-based) or token based (JWT) authentication mechanisms

More details into [pre-GDB Tape Evolution talk](#)

# WLCG JWT scope-based authorization

Since version 1.3.1, StoRM WebDAV introduces:

- support for [WLCG JWT profile scope-based authorization](#)

## Test Statistics

Screenshot from the [WLCG JWT Compliance testsuite](#) report

Total Statistics	Total	Pass	Fail	Skip	Elapsed	Pass / Fail / Skip
All Tests	26	24	2	0	00:00:25	

StoRM WebDAV fully supports token-based authN/Z

- Storage-issued token support (aka “Macarons”) ✓
- OpenID Connect authentication support ✓
- OAuth access token support ✓

WLCG JWT profile compliance:

- audience enforcement support ✓
- capability-based authorization ✓
- group-based authorization ✓
- path-enforced-authz-checks ✗ (tag “not-critical”)

- a new authorization engine which allows for finer-grained authorization
  - allows admins to define access [policies](#) for individual files/folders in the storage area, and targeting specific authenticated users or group of users

## Finer grained authorization: example

- **sa:** `wlcg`  
**description:** Deny write access to `/wlcg/protected` area  
**actions:**
  - `write`
  - `delete`**paths:**
  - `/protected/**`**effect:** `deny`  
**principals:**
  - **type:** `anyone`
- **sa:** `wlcg`  
**description:** Allow read access to `/wlcg/protected` area to `wlcg` members  
**actions:**
  - `read`
  - `list`**paths:**
  - `/protected/**`**effect:** `permit`  
**principals:**
  - **type:** `vo`  
**params:**
    - `vo:` `wlcg`
  - **type:** `jwt-group`  
**params:**
    - `iss:` `https://wlcg.cloud.cnaf.infn.it/`
    - `group:` `/wlcg`



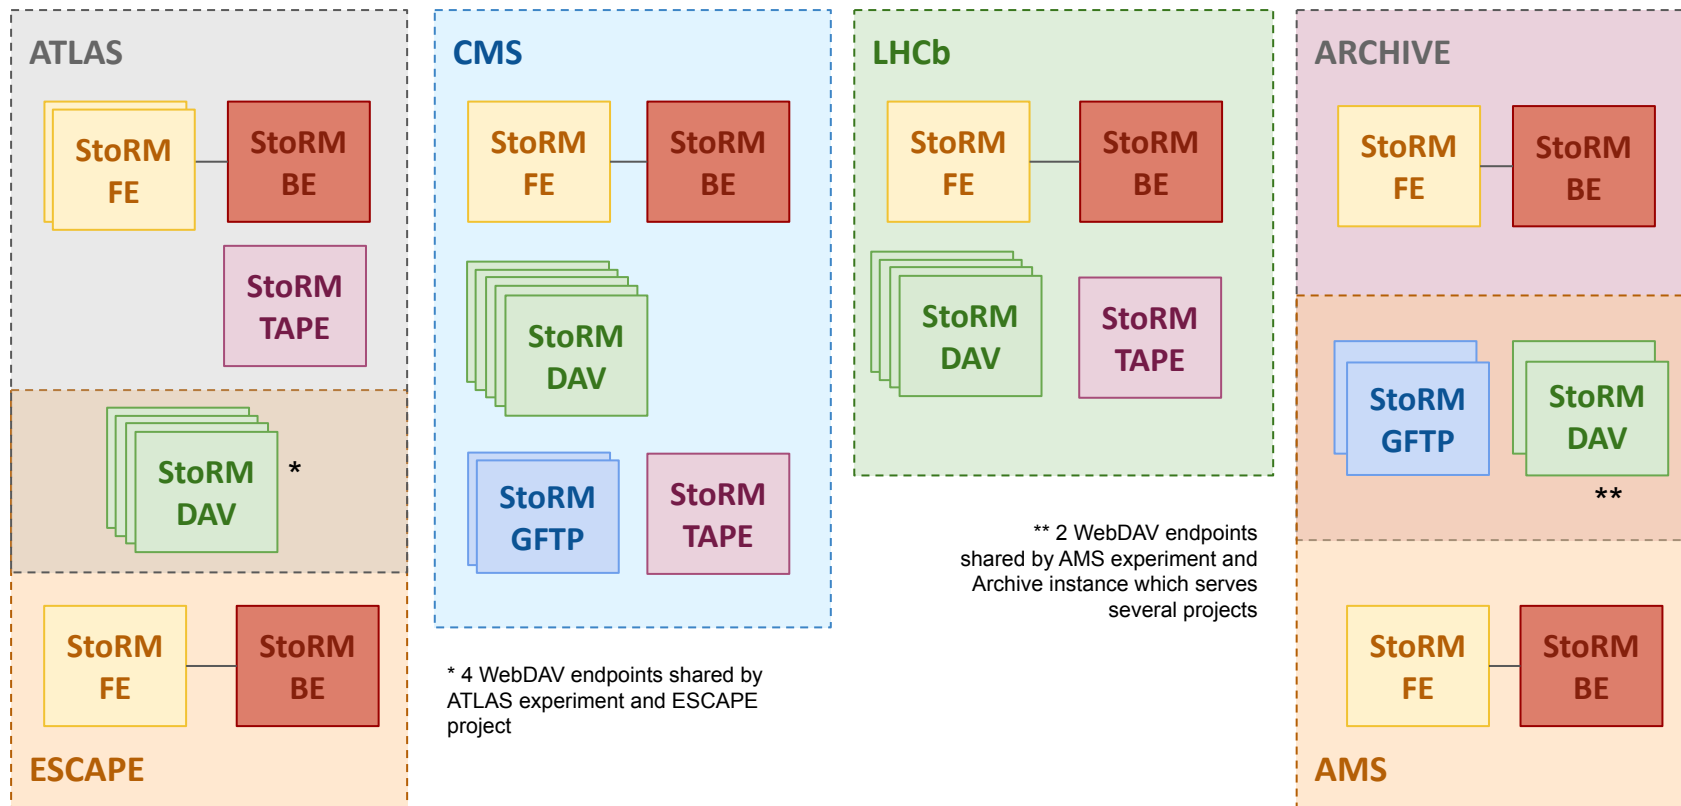
# StoRM Monitoring files & endpoints

- StoRM Backend log files
  - **heartbeat.log** - *This is an aggregated log that shows the number of synch and asynch requests occurred from startup and on last minute.*
  - **storm-backend-metrics.log** - *A finer grained monitoring of incoming synchronous requests, contains metrics for individual types of synchronous requests (synch.mkdir, synch.mv, synch.rm, synch.rmDir, fs.aclOp, fs.fileAttributeOp, ea, etc.)*
- StoRM Frontend log files
  - **monitoring.log** - *reports both information about requests that have been performed in the last Monitoring Round (default 1 minute) and information considering the whole Frontend execution time (aggregate monitoring)*
- StoRM WebDAV
  - **/status/metrics** endpoint - *A finer grained monitoring for jvm, embedded http server, TPC, requests per storage-area, http responses, fs operations, etc.*
- StoRM Info Provider storage report exported through WebDAV
  - **/info/report.json** storage report - *Contains information more oriented to storage areas: used/free space status, supported transfer protocols, transfer endpoints, etc.*

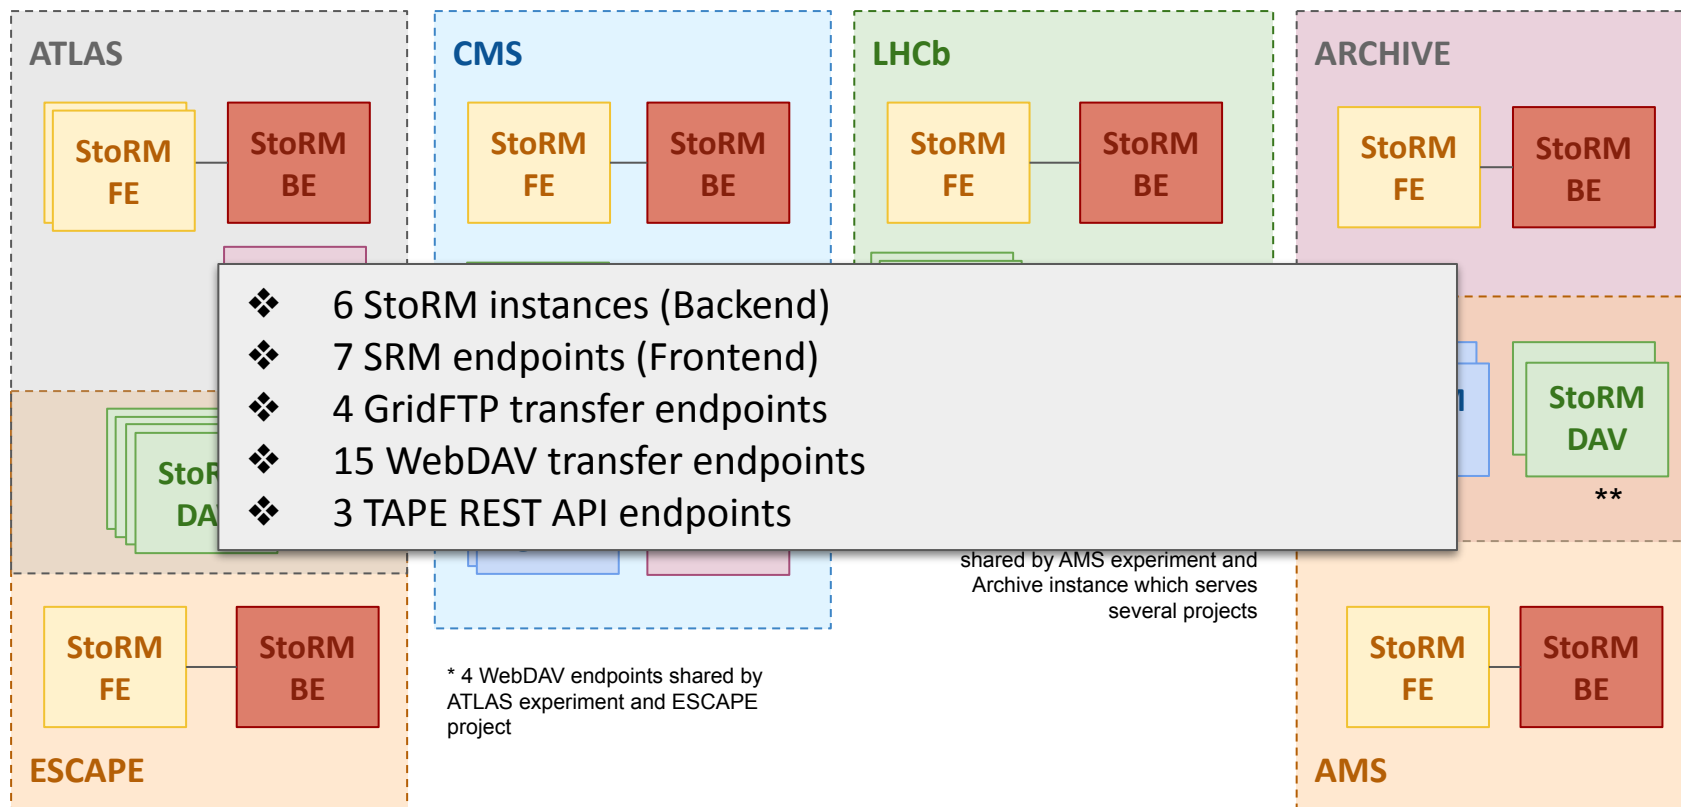
# StoRM deployments and monitoring @ CNAF-T1

Thanks to [storage@inf-n1](mailto:storage@inf-n1)

# StoRM Deployments @ CNAF-T1 - Summary



# StoRM Deployments @ CNAF-T1 - Summary



## StoRM Tape REST API @ CNAF-T1

- Currently v.0.4.0 *beta* version is available in production at INFN-CNAF
  - Installed via RPM and deployed as a standalone component
  - Deployed three separated instances, one for each WLCG experiment
    - <https://tape-atlas.cr.cnaf.infn.it:8443>
    - <https://tape-cms.cr.cnaf.infn.it:8443>
    - <https://tape-lhcb.cr.cnaf.infn.it:8443>
  - Direct access from remote users
    - not deployed within StoRM WebDAV
  - Configured to authorize users via VOMS proxies (ATLAS, CMS, LHCb)
    - ATLAS and CMS authZ via WLCG tokens is also available
- not much used yet ⇒ ready for tests and production

## StoRM Deployments @ CNAF-T1 - Token support

- WLCG JWT profile scope-based authorization is enabled on:
  - ATLAS, CMS, DUNE disk storage areas
  - CMS tape storage area
  - Looking forward to LHCb input
- More than other 30 experiments have (non WLCG) token-based authorization enabled
  - BELLE, JUNO, CTA-LST, HERD, etc.
  - Mixed AuthZ based on capabilities and/or identities/groups

# StoRM Deployments @ CNAF-T1 - Token support open point


WLCG Profile seems to not contain a clear statement on how a scope's path relates to the requested path

- There are several WLCG Profile's open discussions (issues?) about scope path definition:
  - [No discussion on how to map the scope path to request path](#)
  - [Limitation of the storage claim](#)
  - [Improve storage.\\*:/path description](#)
- With related StoRM issues into our issue-tracker:
  - [WLCG-scope-based authZ doesn't allow the access point inside the scope path](#)

## Monitoring with Sensu + InfluxDB + Grafana

Alarms & metrics are monitored with a Sensu + InfluxDB + Grafana infrastructure

- Server health metrics
- Network metrics
- StoRM services metrics
  - e.g. StoRM WebDAV metrics  
<https://xfer.cr.cnaf.infn.it:8443/status/metrics>
- StoRM Backend metrics from log files
  - heartbeat.log
  - storm-backend-metrics.log



Internal and public views can be easily accessed through [t1-metria dashboards](#)



retention one\_week

General info

GPFS servers 71

StoRM servers 8

HSM servers 5

Gateway servers 54

GridFTP 10

Xrootd 36

Webdav 18


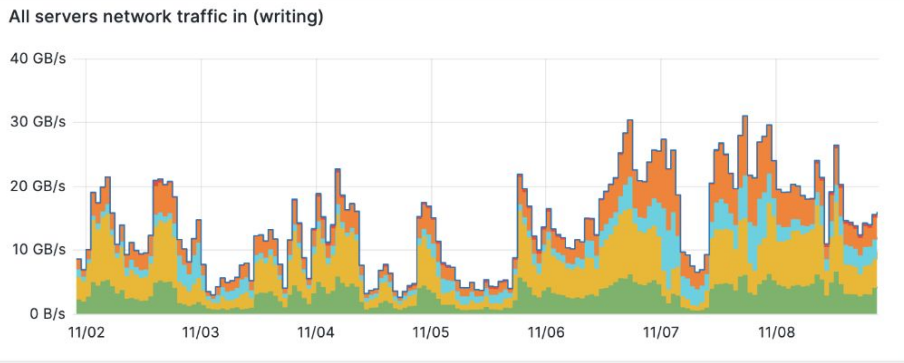
Disk net capacity 69.7 PB

Disk used space 48.4 PB

Tape net capacity 132 PB

Tape used space 110 PB


Data traffic



# Plots from t1metria

from BE's heartbeat.log →


Number of async PTP / min (sum over hosts)



	min	max	avg	current
■ async_ptp_tot	0	474	8.16	0.458
■ async_ptp_ok	0	474	8.16	0.458

from WebDAV metrics endpoint  
↓

HTTP requests rate



■ xs-604.cr.cnaf.infn.it	0.00309	76.5	4.53
■ xs-605.cr.cnaf.infn.it	0.00290	75.4	4.50
■ xs-704.cr.cnaf.infn.it	0.483	75.7	4.58

# Monitoring from log files: Big Data Platform (new)

## Producer



SSL encrypted

Ship

one tag for each  
CNAF internal team

Producer custom

## Message broker



Configured with 3 partitions, 3 replicas and one week retention. Authorization through ACLs.

Consumer custom



## Consumer

Log files parsed using grok filters based on Regular Expressions.



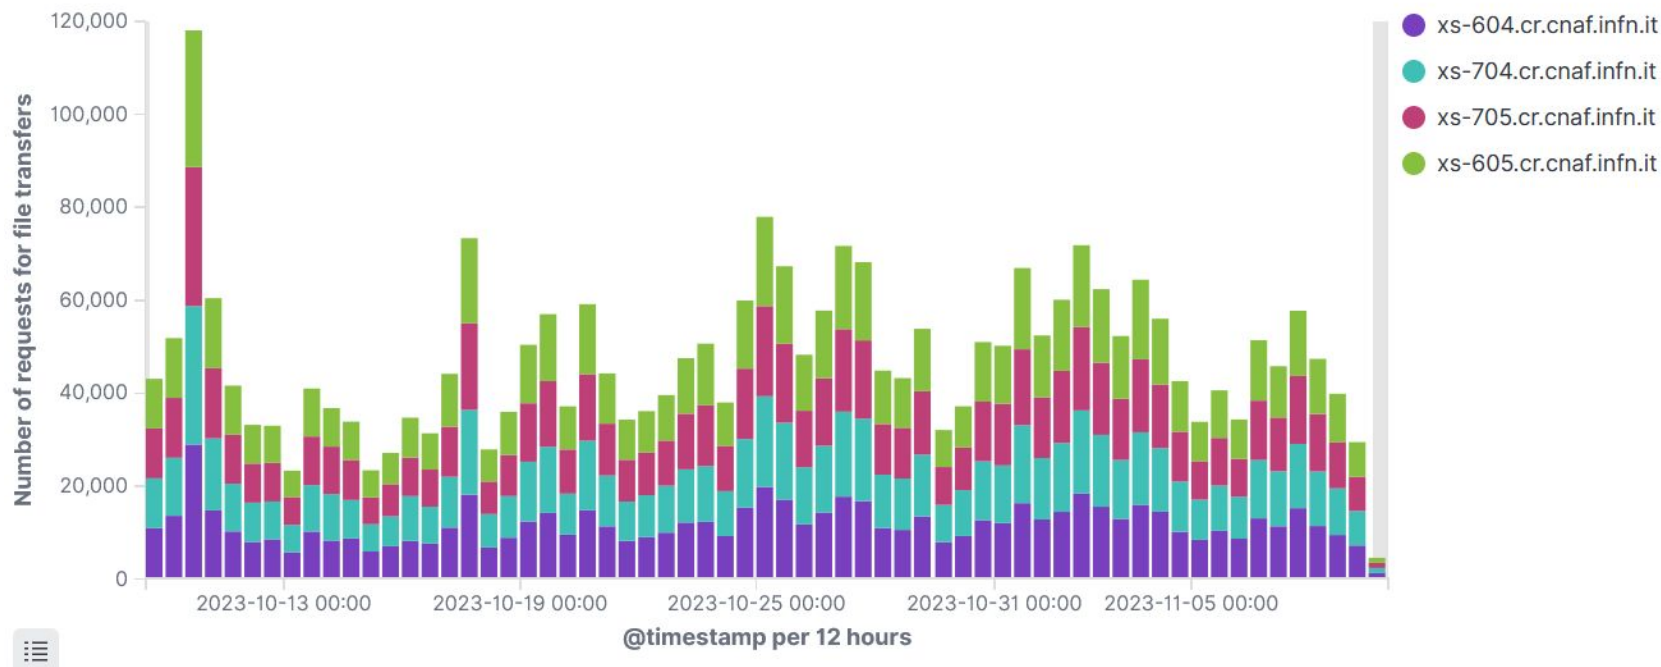
user

AuthN/Z

Allows to inspect and visualize log files. Integrated with INDIGO IAM: Opensearch backend roles are mapped to IAM groups



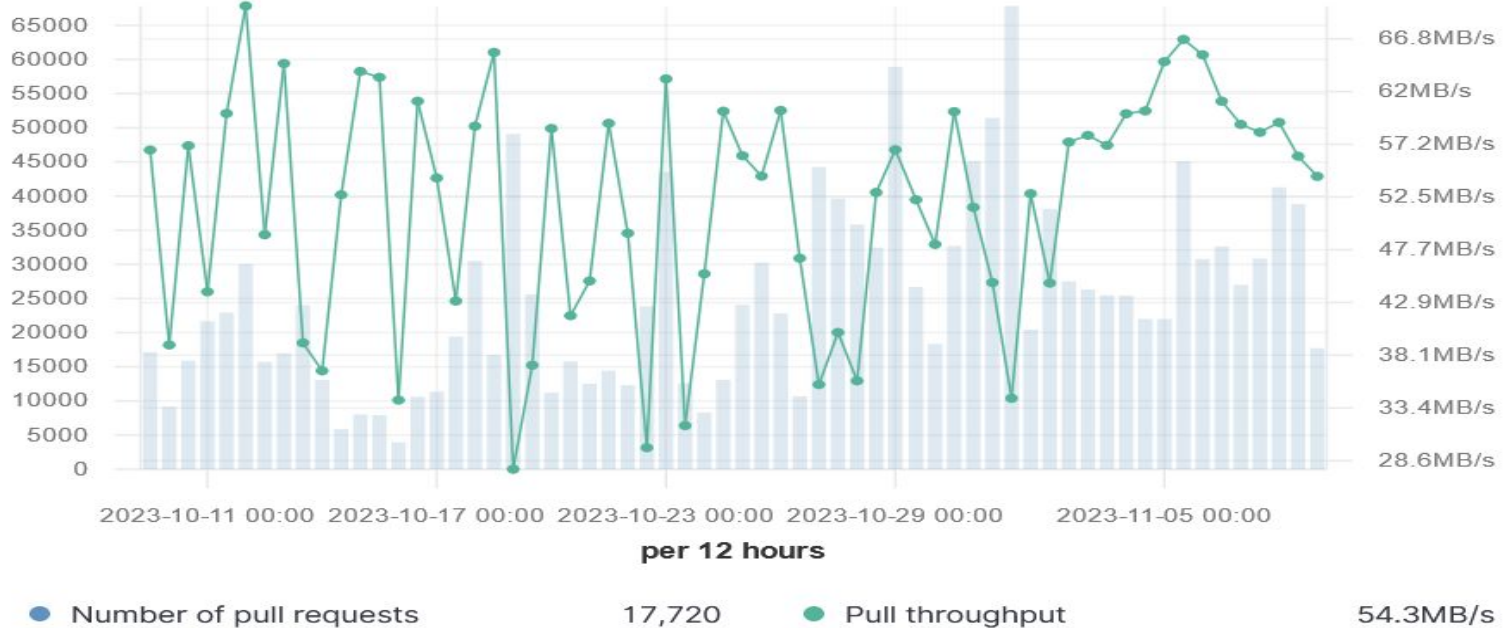
## Big Data Platform: examples of plots



*Distribution of file transfer requests among the 4 StoRM WebDAV endpoints dedicated to ATLAS*

# Big Data Platform: examples of plots

StoRM WebDAV TPC Pull throughput with requests



*Last month ATLAS TPC view computed from log files: histogram shows the number of TPC requests, the pointed line is the total throughput*

# StoRM evolution

## Objectives on our roadmap

- Finalize new component StoRM Tape REST API
  - Improve AuthZ path specific policies ... *in progress*
- Go beyond CentOS 7
  - build and release packages for EL 9 platform ... *in progress*
- Simplify service operations and deployment
  - Service containerisation, K8S, ... *started with StoRM Taper REST API*
- Improve JVM profiling with tools
  - i.e. Java Flight Recorder, YourKit, ... *started*
- Improve support for Cloud Storage providers by StoRM WebDAV
  - i.e. Google Cloud Storage, Amazon S3, ...
- Improve observability
  - Consistent logging across services, metrics, tracing, ...
- Drop Globus dependency on StoRM Frontend, if/when needed

## StoRM plans for DC 24

- Solve StoRM WebDAV [thread contention](#) which manifests as load unbalancing between deployed instances
- Support packet marking
- Provide requested metrics (if any)



**Thanks! Questions?**

# Contacts and references

GitHub: <https://github.com/italiangrid/storm>

Documentation: <http://italiangrid.github.io/storm/>

Contacts:

[storm-support@lists.infn.it](mailto:storm-support@lists.infn.it) and [storm-users@lists.infn.it](mailto:storm-users@lists.infn.it) for users support

[storm-devel@lists.infn.it](mailto:storm-devel@lists.infn.it) for developers