A RESTful approach to tape management in StoRM

Federica Agostini, Laura Cappelli, Tommaso Diotalevi, Angelo Galavotti, Francesco Giacomini,
Roberta Miccoli, Aksieniia Shtimmerman, Marcelo Vilaça Pinheiro Soares, **Enrico Vianello**

INFN-CNAF, Viale Berti Pichat 6/2,
40127 Bologna, Italy

Abstract

The **STOrage Resource Manager (StoRM)** [1] service relies on the SRM specification to recall files from tape. Although the SRM protocol has been successfully used for many years, its complexity has pushed the WLCG community to adopt a simpler approach, more in line with modern web technologies. The **WLCG tape REST API** offers a common HTTP interface allowing clients to manage disk residency of tape-stored files and observe the progress of file transfers to disk. In the context of the StoRM project developed at INFN-CNAF, the **StoRM Tape REST API** [2] implements this HTTP interface and it is deployed as a standalone component. It uses **NGINX** [3] reverse proxy as authentication engine and **Open Policy Agent (OPA)** [4] as authorisation policies enforcer. At the INFN Tier-1, this new service is required to coexist with the current StoRM deployments and to integrate smoothly within the existing infrastructure, in particular with the **Grid-Enabled Mass Storage System (GEMSS)** [5].

NGINX

NGINX is an open-source **HTTP server** and **reverse proxy**, known for its high performance, stability, rich feature set, easy configuration, and low resource consumption.

The service has been chosen as part of this deployment for **VOMS/TLS termination** and **authentication with JSON Web Tokens (JWT)** [6].

StoRM WebDAV integration

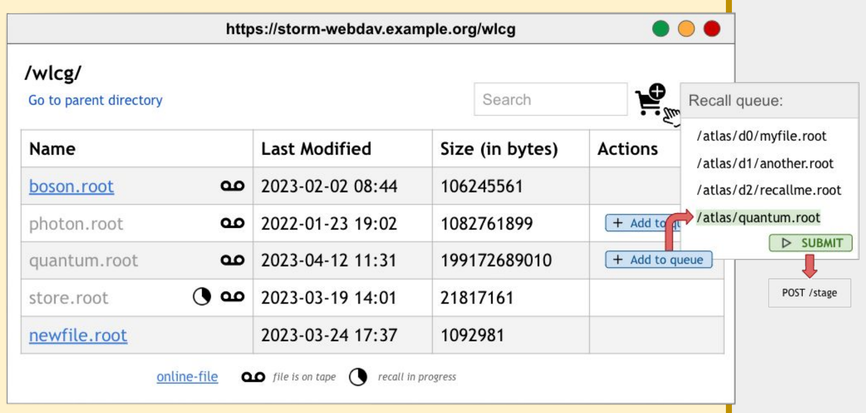

https://storm-webdav.example.org/wlcg

/wlcg/

Go to parent directory

Name	Last Modified	Size (in bytes)	Actions
boson.root	2023-02-02 08:44	106245561	
photon.root	2022-01-23 19:02	1082761899	+ Add to queue
quantum.root	2023-04-12 11:31	199172689010	+ Add to queue
store.root	2023-03-19 14:01	21817161	
newfile.root	2023-03-24 17:37	1092981	

Recall queue:
/atlas/d0/myfile.root
/atlas/d1/another.root
/atlas/d2/recallme.root
/atlas/quantum.root

StoRM WebDAV [1] allows users to navigate storage areas with a browser and some privileged users (identified by a JWT group membership or a DN for example) could also be able to trigger a bulk stage request through the folder view.

Monitoring

A monitoring dashboard will be available for some privileged users (identified by a JWT group membership or a DN for example) in order to monitor the ongoing stage requests.

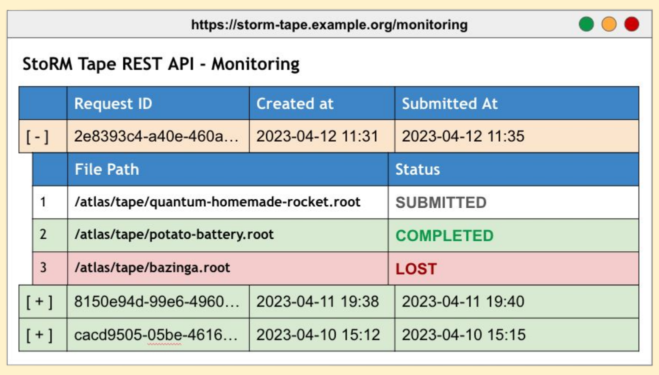

https://storm-tape.example.org/monitoring

StoRM Tape REST API - Monitoring

	Request ID	Created at	Submitted At
[-]	2e8393c4-a40e-460a...	2023-04-12 11:31	2023-04-12 11:35

	File Path	Status
1	/atlas/tape/quantum-homemade-rocket.root	SUBMITTED
2	/atlas/tape/potato-battery.root	COMPLETED
3	/atlas/tape/bazinga.root	LOST

	Request ID	Created at	Submitted At
[+]	8150e94d-99e6-4960...	2023-04-11 19:38	2023-04-11 19:40
[+]	cacd9505-05be-4616...	2023-04-10 15:12	2023-04-10 15:15

NGINX VOMS plugin

ngx_http_voms_module [11] is a module for NGINX developed at INFN-CNAF, that enables client-side authentication based on X.509 proxy certificates augmented with VOMS Attribute Certificates, typically obtained from a **Virtual Organization Membership Service (VOMS)** [12] server.

It defines a set of embedded variables (e.g. ***voms_fqans***), whose values are extracted from the Attribute Certificate.

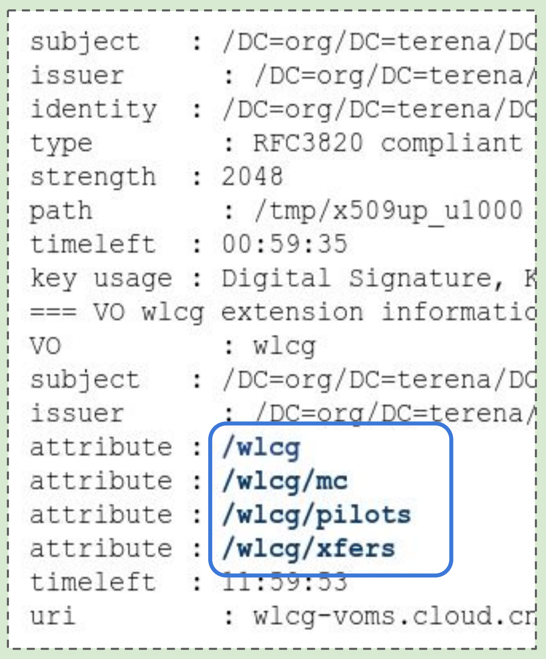

```
subject   : /DC=org/DC=terena/DC
issuer    : /DC=org/DC=terena/
identity  : /DC=org/DC=terena/DC
type      : RFC3820 compliant
strength  : 2048
path      : /tmp/x509up_u1000
timeleft  : 00:59:35
key usage : Digital Signature, K
=== VO wlcg extension informatio
VO        : wlcg
subject   : /DC=org/DC=terena/DC
issuer    : /DC=org/DC=terena/
attribute : /wlcg
attribute : /wlcg/mc
attribute : /wlcg/pilots
attribute : /wlcg/xfers
timeleft  : 11:59:53
uri       : wlcg-voms.cloud.cr
```

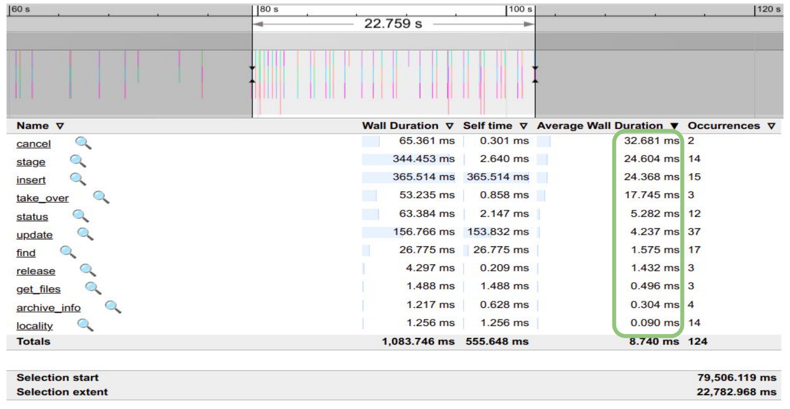

chrome://tracing

Name	Wall Duration	Self time	Average Wall Duration	Occurrences
cancel	65.361 ms	0.301 ms	32.681 ms	2
stage	344.453 ms	2.640 ms	24.604 ms	14
insert	365.514 ms	365.514 ms	24.368 ms	15
take_over	53.235 ms	0.858 ms	17.745 ms	3
status	63.384 ms	2.147 ms	5.282 ms	12
update	156.766 ms	153.832 ms	4.237 ms	37
find	26.775 ms	26.775 ms	1.575 ms	17
release	4.297 ms	0.209 ms	1.432 ms	3
get_files	1.488 ms	1.488 ms	0.496 ms	3
archive_info	1.217 ms	0.628 ms	0.304 ms	4
locality	1.256 ms	1.256 ms	0.090 ms	14
Totals	1,083.746 ms	555.648 ms	8.740 ms	124

Selection start	79,506.119 ms
Selection extent	22,782.968 ms

Basic profiling. During its execution, the service can produce a JSON file containing **tracing information** for selected instrumented functions.

Open Policy Agent (OPA)

Open Policy Agent is an open-source **authorization engine** that unifies policy enforcement across the stack, due to a high-level declarative language that allows the definition of policies as code [6].

OPA could be a valid alternative to Argus [7].

```
# GET /api/v1/stage/<id>
allow if {
  input.method == "GET"
  glob.match("/api/v1/stage/*", ["/"], input.path)

  any([scopes_allowed, fqans_allowed, dn_allowed])
}
```

has allowed WLCG scopes OR has allowed FQANs OR has allowed DN

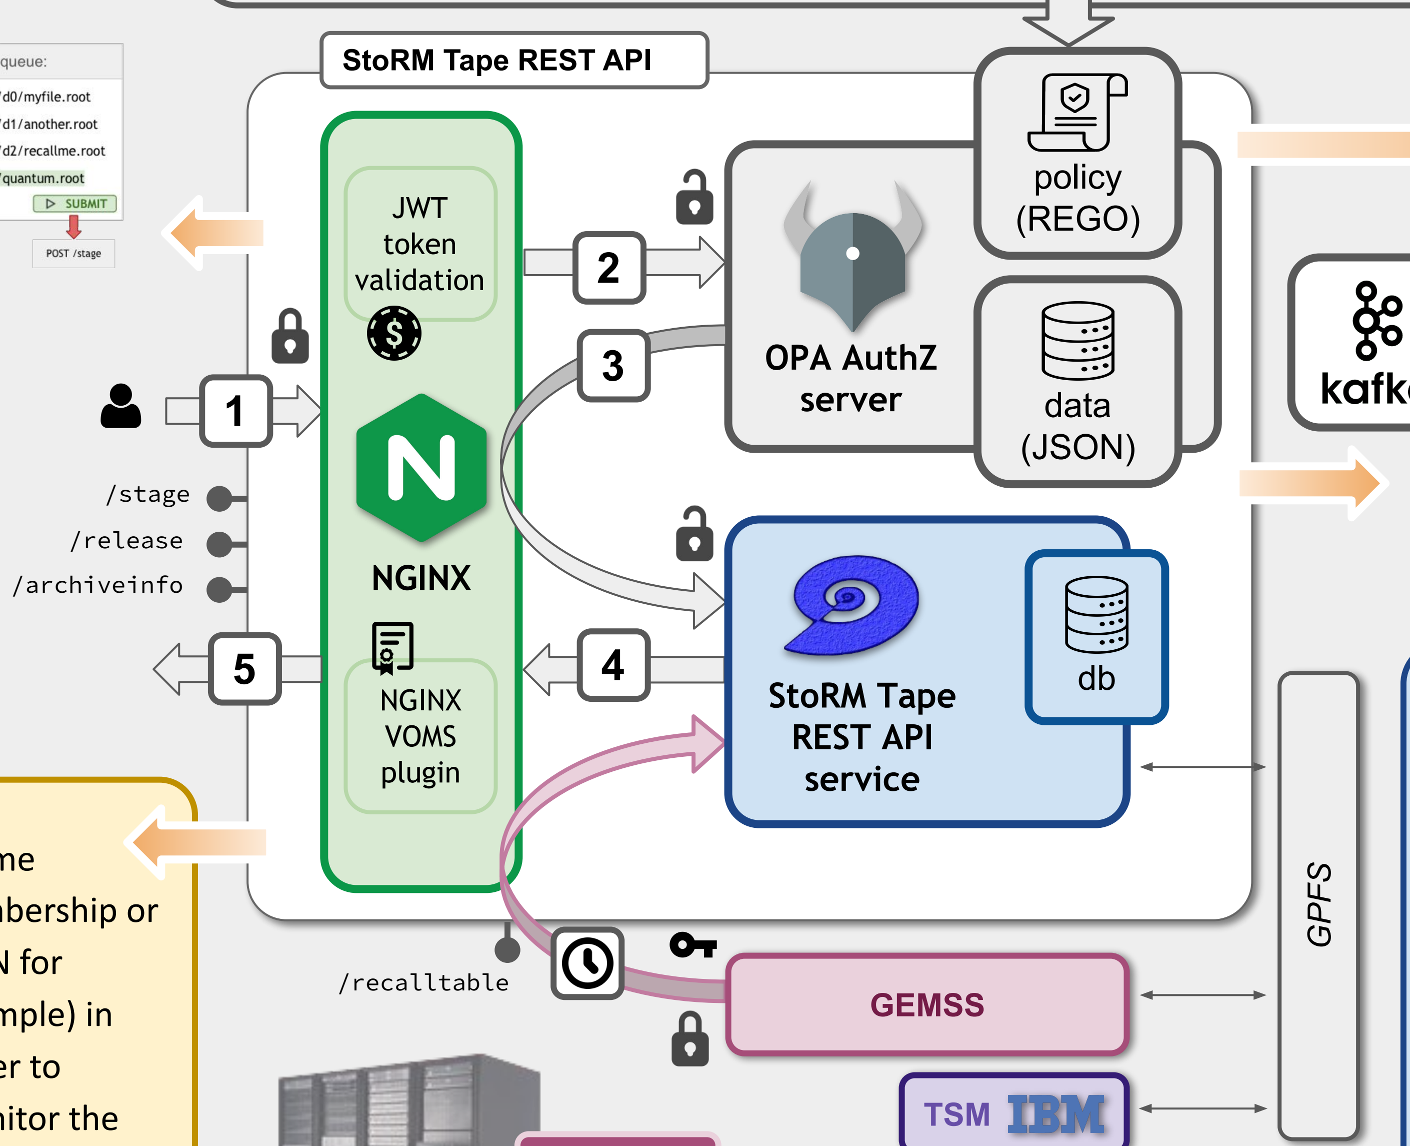

GEMSS

Periodically, **GEMSS retrieves the list of ready-to-recall files** from StoRM Tape REST API with an authenticated request. The access to this endpoint is also **restricted** by NGINX to a limited list of IP addresses.

Performance

Vegeta [13] is a versatile HTTP **load testing tool** built out of a need to drill HTTP services with a constant request rate.

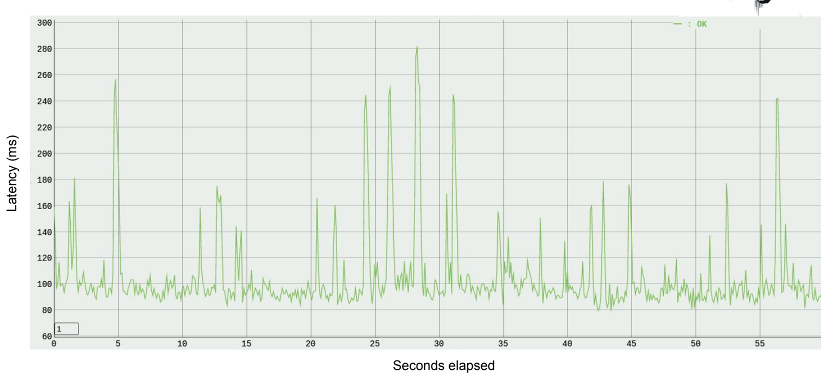

Deployment tests use the data-driven **Robot Framework** [6][14]. Fixes of failing tests are in progress.

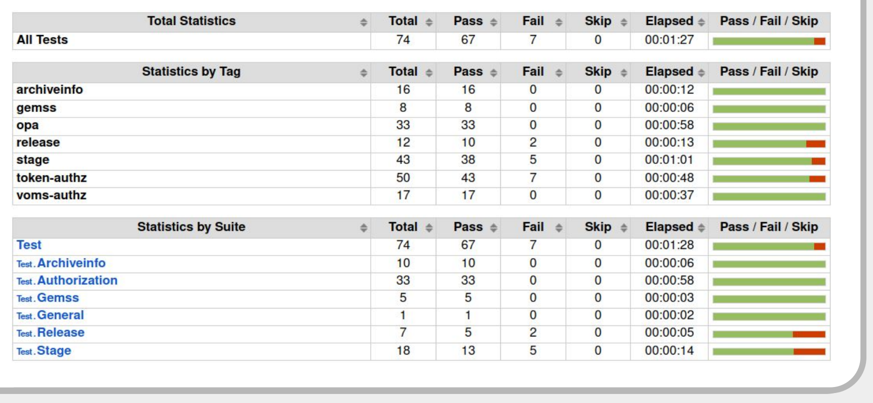

Total Statistics	Total	Pass	Fail	Skip	Elapsed	Pass / Fail / Skip
All Tests	74	67	7	0	00:01:27	

Statistics by Tag	Total	Pass	Fail	Skip	Elapsed	Pass / Fail / Skip
archiveinfo	16	16	0	0	00:00:12	
gemss	8	8	0	0	00:00:06	
opa	33	33	0	0	00:00:58	
release	12	10	2	0	00:00:13	
stage	43	38	5	0	00:01:01	
token-authz	50	43	7	0	00:00:48	
voms-authz	17	17	0	0	00:00:37	

Statistics by Suite	Total	Pass	Fail	Skip	Elapsed	Pass / Fail / Skip
Test	74	67	7	0	00:01:28	
Test . Archiveinfo	10	10	0	0	00:00:06	
Test . Authorization	33	33	0	0	00:00:58	
Test . Gemss	5	5	0	0	00:00:03	
Test . General	1	1	0	0	00:00:02	
Test . Release	7	5	2	0	00:00:05	
Test . Stage	18	13	5	0	00:00:14	

Legend

- JWT token issued by INDIGO IAM
- VOMS proxy
- Basic AuthN
- HTTPS communication
- HTTP communication
- Future developments

Support for remote OPA policies and data. OPA includes functionality for reaching out to external servers where to store Policies and Data. They may be cached locally.

Integration with Big Data Platform (BDP) at INFN-CNAF. It consists of

- a message broker: **Apache Kafka**
- a data processing instance: **Logstash**
- a data discovery and visualization engine: **OpenSearch**

Logstash

OpenSearch

StoRM Tape REST API service

StoRM Tape REST API service manages bulk stage requests of tape-stored files, making them available with disk latency. It also allows to: check the status of each submitted request; cancel a subset of staged files or delete a request completely; access to staged file metadata. The API service checks files locality directly from the underline storage system (GPFS).

It also provides an endpoint for GEMSS to replicate the current interaction with StoRM Backend.

The API service is written in C++, based on the Crow framework [8] and SQLite [9] as database engine. It is deployed as a standalone component, packed in a Docker image. It is also available as RPM on [10].

Ready as a preview!

NGINX + OPA Authorization flow:

1 Client submits an API request

2 NGINX sends the request to the OPA engine

3 OPA makes the AuthZ decision using its rules and data and sends it back to NGINX. In case of successful authZ, the request is forwarded to the StoRM Tape REST API service (or else Forbidden)

4 5 The response from the service is relayed to the client via NGINX

References

[1] https://italiangrid.github.io/storm/documentation.html
[2] https://baltig.infn.it/cnafsd/storm-tape
[3] https://nginx.org/en/docs
[4] https://www.openpolicyagent.org/docs/latest
[5] https://github.com/italiangrid/gemss
[6] https://baltig.infn.it/cnafsd/storm-tape-ts
[7] https://argus-documentation.readthedocs.io/en/stable
[8] https://github.com/CrowCpp/Crow
[9] https://sqlite.org
[10] https://repo.cloud.cnaf.infn.it
[11] https://baltig.infn.it/cnafsd/ngx_http_voms_module
[12] https://italiangrid.github.io/voms
[13] https://github.com/tsenart/vegeta
[14] https://robotframework.org