

Evolving StoRM WebDAV: delegation of file transfers to NGINX and support for SciTags

Luca Bassi^{1,2}, Federica Agostini¹, Jacopo Gasparetto¹, Francesco Giacomini¹,
Roberta Miccoli¹, Enrico Vianello¹

¹INFN-CNAF, Bologna, Italy

²GARR, Rome, Italy

CHEP 2024, 22 October 2024

StoRM is a storage resource manager for disk-based storage systems that provides a “thin” management layer (SRM, WebDAV) over a POSIX file system: typically a distributed file system such as IBM GPFS or Lustre.

StoRM is a suite of components.

StoRM provides flexible authentication and authorization mechanisms:

- ▶ VOMS proxies and OAuth/OIDC JWT tokens
- ▶ File access control is enforced via POSIX ACLs

StoRM supports a tape system through integration with GEMSS, a full Hierarchical Storage Management (HSM) system that integrates:

- ▶ IBM General Parallel File System (GPFS)
- ▶ IBM Tivoli Storage Manager (TSM)
- ▶ StoRM Backend

StoRM main components are:

StoRM Backend + StoRM Frontend SRM disk and tape file management

StoRM Tape HTTP/HTTPS tape file management

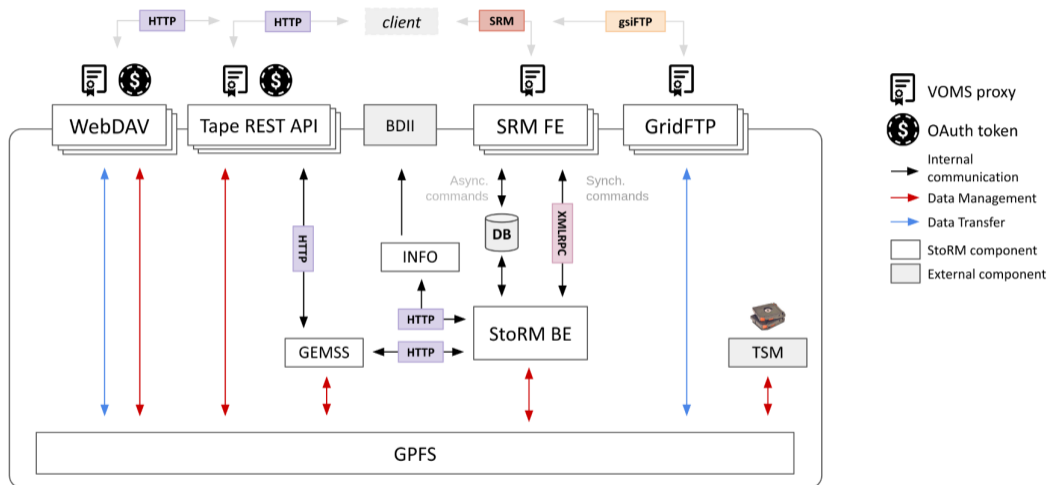
StoRM GridFTP GridFTP file transfer

StoRM WebDAV HTTP/HTTPS file management and transfer

StoRM is the Grid/Cloud storage solution adopted by the INFN Tier 1 at CNAF and some other Tier 2.

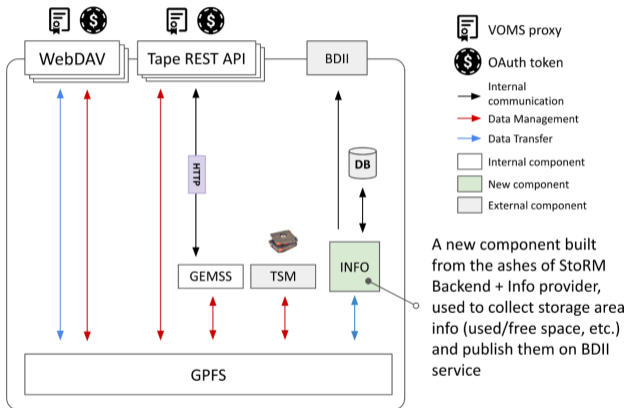
It will be maintained and evolved by INFN for the foreseeable future: including support (through GGUS tickets or mailing-list) to StoRM-based sites.

StoRM current typical deployment



Future architecture: no-SRM deployment

- ▶ All Globus GridFTP will be turned off soon
- ▶ StoRM WebDAV + Tape components provide the necessary data transfer/management functionalities on a tape-enabled storage system, without the need of legacy SRM components



StoRM WebDAV is a StoRM component which provides data transfer functionality through the WebDAV (Web Distributed Authoring and Versioning) protocol.

WebDAV is an extension of the HTTP protocol that allows users to create, change and move resources on a web server.

StoRM WebDAV also provides a browser-based data management interface.

It supports authorization based on JWT tokens, X.509 certificates and VOMS proxies. It supports Third Party Copies (TPC), relying on an extension of the WebDAV COPY verb, which consists in bulk transfer requests between two remote storage endpoints.

StoRM WebDAV is evolving and improving.

In particular:

- ▶ Packet and flow marking, thanks to SciTags support
- ▶ Improved efficiency, delegating file transfers to NGINX

Scientific network tags (scitags) is an initiative promoting identification of the science domains and their high-level activities at the network level.¹

Two ways to mark research traffic:

- ▶ Flow Marking (UDP firefly)
- ▶ Packet Marking (Flow Label)

¹<https://www.scitags.org/>

Each flow identifier has two fields:

- ▶ Experiment/virtual organization (required)
- ▶ Experiment activity (optional)

These fields are encoded in the SciTag HTTP header: $\langle \text{expID} \rangle \ll 6 \mid \langle \text{actID} \rangle$

So it's a 16-bit positive integer > 64 and < 65536 , for example for ATLAS (expID: 2) doing a Data Consolidation (actID: 4) the header will be SciTag: 132


The IDs are statically mapped in a flow registry².

²<https://www.scitags.org/api.json>

SciTags: Flow Marking (UDP firefly)

Send a UDP packet at the start and end of each transfer to a specific destination port with information about the experiment and the activity that generated the traffic.

Example: HTTP-TPC push mode A \rightarrow B




SciTags: Flow Marking (UDP firefly)

Send a UDP packet at the start and end of each transfer to a specific destination port with information about the experiment and the activity that generated the traffic.

Example: HTTP-TPC push mode A \rightarrow B

Client



Leverage the flow label of the IPv6 header to add information about the experiment and the activity that generated the traffic.

- ▶ Activity identifier encoded in bits 24-29
- ▶ Experiment identifier encoded in bits 14-22
- ▶ 5 entropy bits (12-13, 23, 30-31)

The next release of StoRM WebDAV (1.4.3) will have opt-in support for SciTags.

In particular, it uses *flowd*³, a daemon developed by M. Babik and T. Sullivan, that has different plugins to retrieve flow identifier and a set of backend to mark the traffic.

StoRM WebDAV uses:

- ▶ The *np_api* plugin that use a pipe to receive the identifiers
- ▶ The *udp-firefly* backend to send the UDP firefly packets

StoRM WebDAV writes to `/var/run/flowd`:

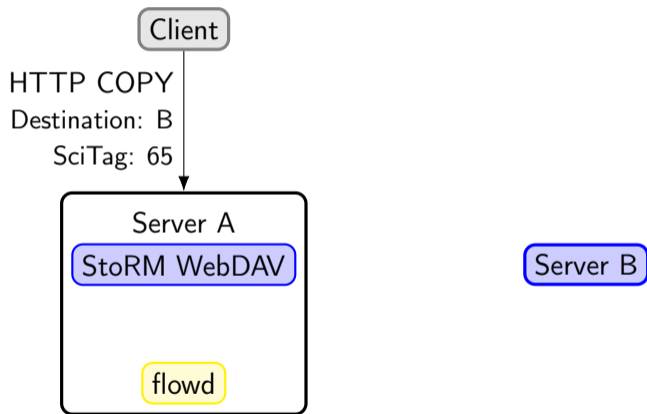
```
state prot src_ip src_port dst_ip dst_port exp act
```

where `state` is `start/end` and `exp/act` the information extracted from the SciTag header.

³<https://github.com/scitags/flowd>

SciTags support in StoRM WebDAV


Example: HTTP-TPC push mode A \rightarrow B



SciTags support in StoRM WebDAV

Example: HTTP-TPC push mode A \rightarrow B


Client



SciTags support in StoRM WebDAV

Example: HTTP-TPC push mode A \rightarrow B


Client



SciTags support in StoRM WebDAV

Example: HTTP-TPC push mode A \rightarrow B

Client



StoRM WebDAV at the moment uses Jetty and implements Milton framework to transfer data.

We are working to provide a deployment option behind NGINX.

NGINX would do TLS termination and authentication with VOMS proxies⁴.

NGINX would also actually perform at least some of the data transfers on behalf of StoRM WebDAV: we leverage the X-Accel-Redirect header to issue an internal redirect.

We expect measurable improvements in terms of resource usage (CPU and memory), performance and codebase maintainability.

⁴https://baltig.infn.it/cnafsd/nginx_http_voms_module

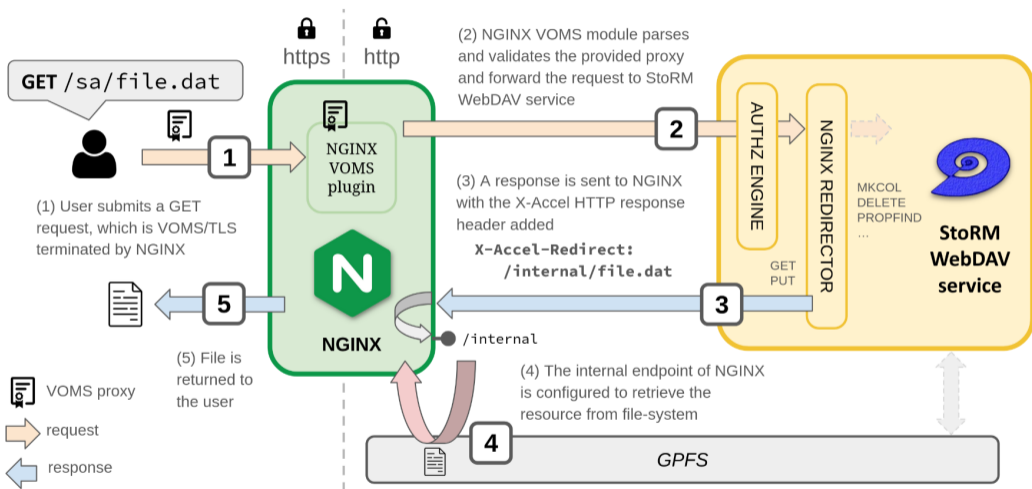
NGINX permits to implement internal redirects.

The backend can send the header `X-Accel-Redirect` with the URI.

This is leveraged to do authorization with StoRM WebDAV, but delegate the file serving to NGINX.


```
location /internal-get {
    internal;
    alias /;
}
location / {
    proxy_pass http://127.0.0.1:8086;
}
```

Delegating transfer requests to NGINX



Improved efficiency delegating file transfers to NGINX

- ▶ An average of +10% in terms of transfer performance
- ▶ Half usage in terms of memory: WebDAV 820 MB vs NGINX 491 MB
- ▶ A huge difference in terms of CPU consuming: WebDAV 10% vs NGINX 0,1%



StoRM WebDAV provides a data transfer functionality through the WebDAV protocol.

The next release of StoRM WebDAV will have opt-in support for SciTags, an initiative for flow/packet marking.

We are working to support a StoRM WebDAV deployment behind NGINX to improve transfer performance.

StoRM WebDAV source code: <https://github.com/italiangrid/storm-webdav>

Thank you for your attention

Questions?

