

StoRM for disk-based sites

Enrico Vianello
on behalf of the StoRM Product Team

EGI Conference 2021

enrico.vianello@cnafe.infn.it



StoRM

StoRM is a storage resource manager for disk-based storage systems

Provides a "thin" management layer (SRM, WebDAV) over a POSIX FS

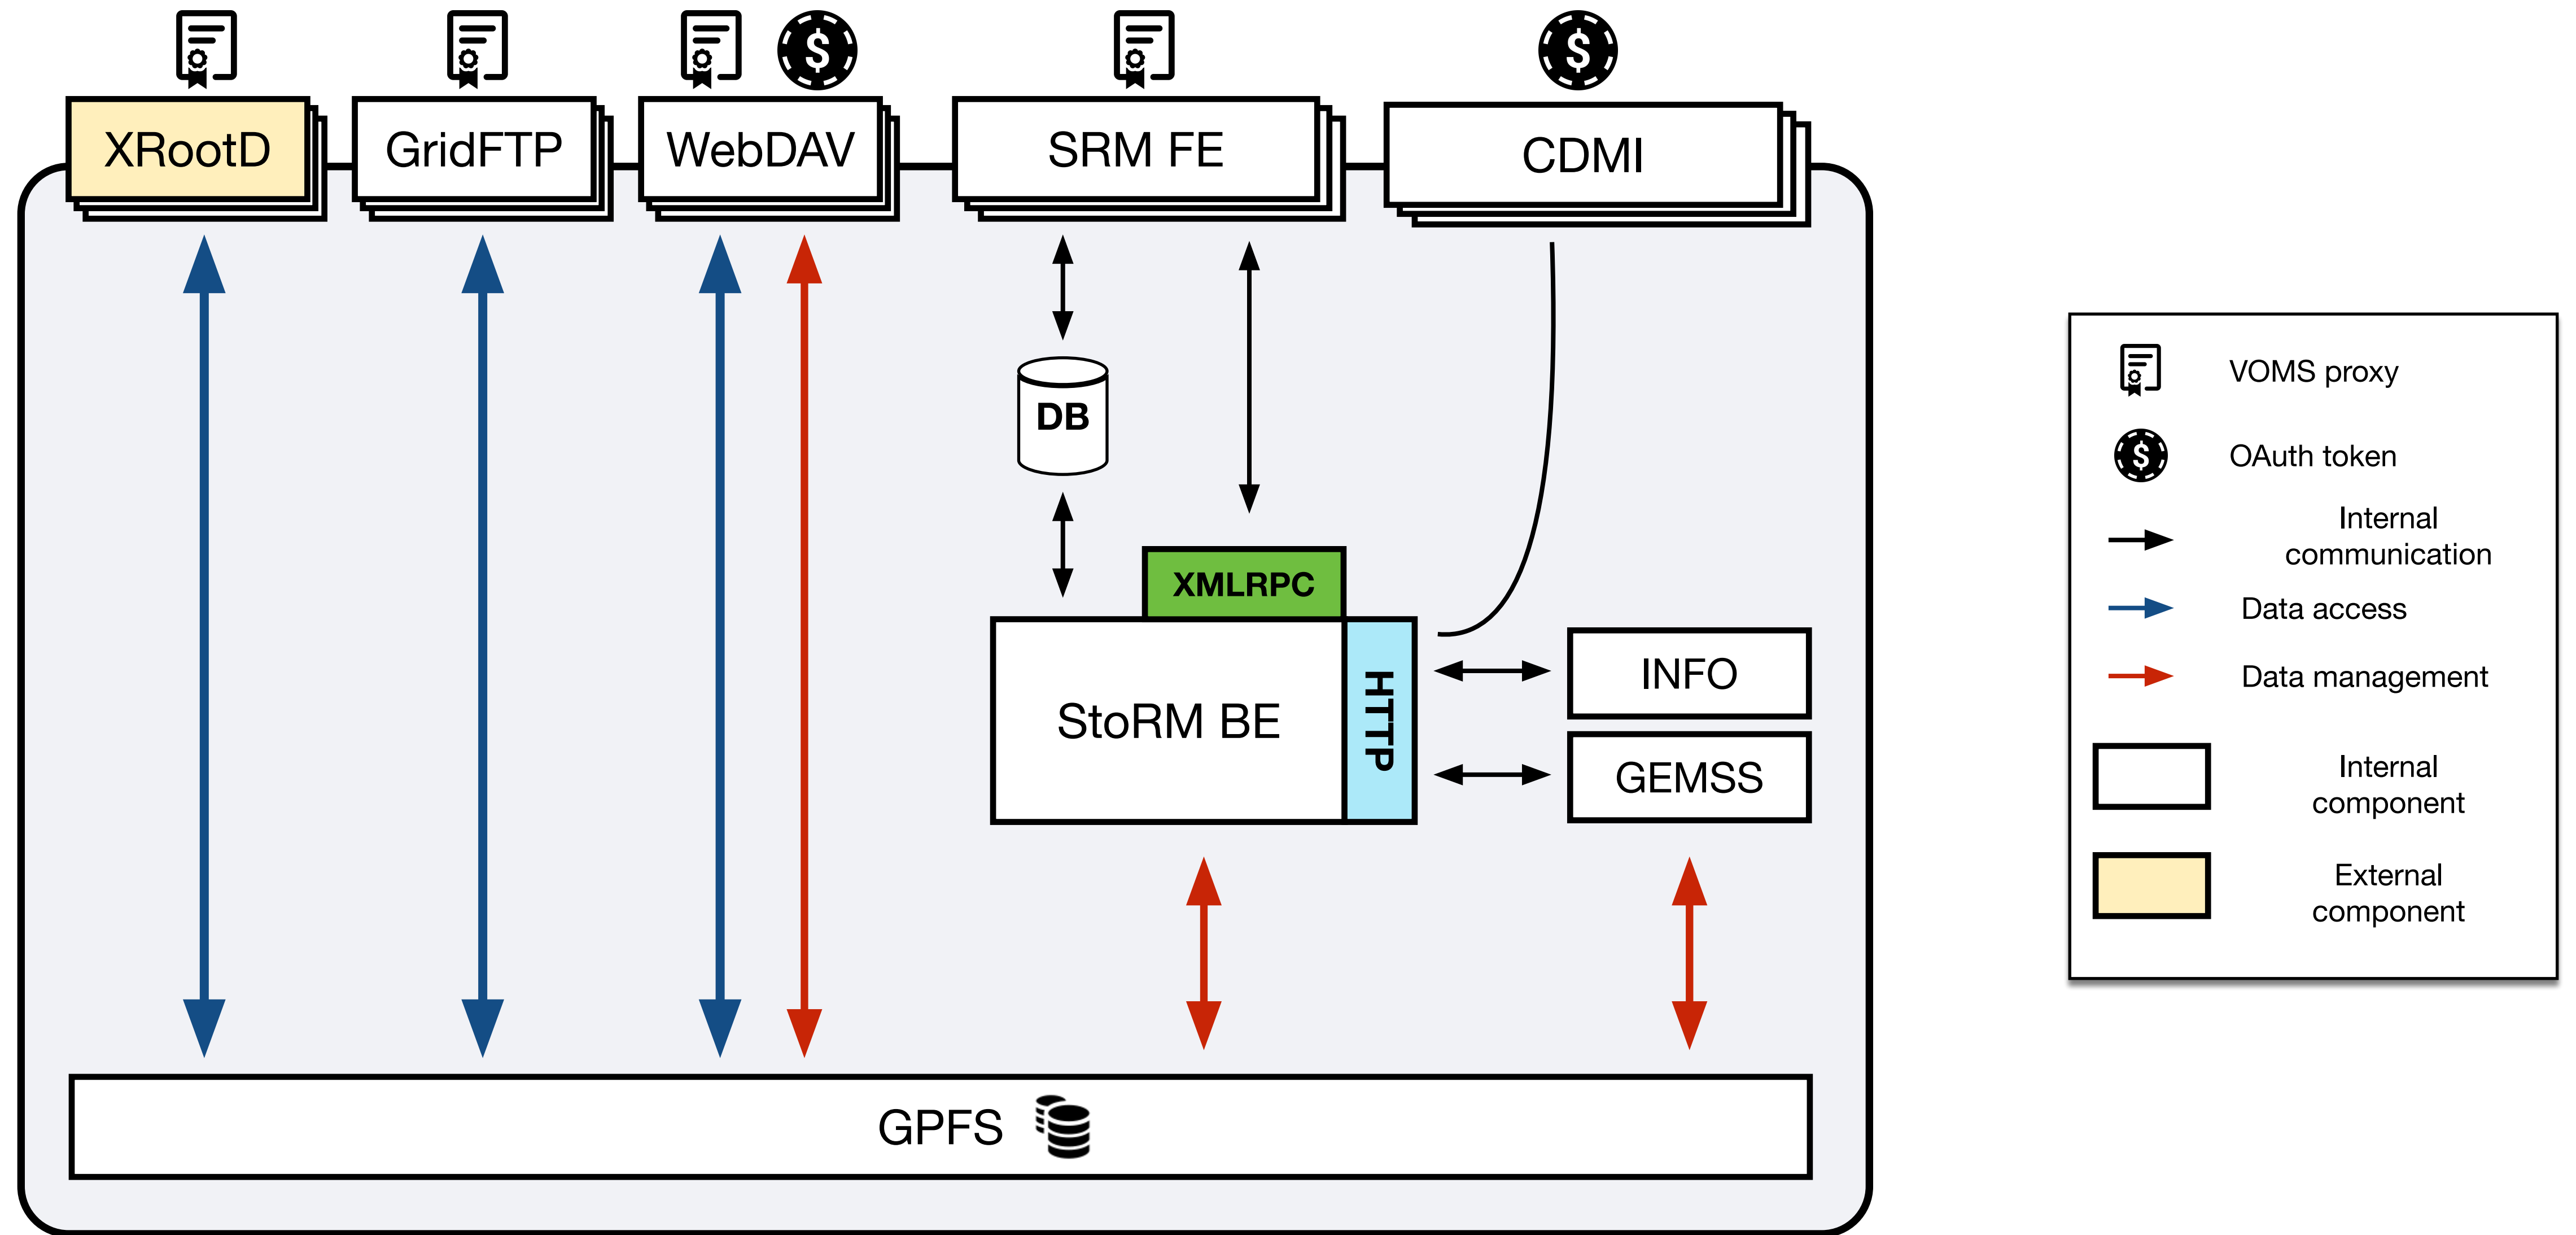
- Typically IBM GPFS or Lustre

Supports tape through integration with GEMSS, a tape library manager component also developed @ INFN-CNAF

Provides flexible AuthN/Z support:

- VOMS & OAuth/OIDC (WebDAV & CDMI interfaces)
- File access control is enforced via POSIX ACLs

StoRM: a typical deployment architecture



Main recent developments

CENTOS 7/SystemD porting & StoRM Puppet module

Token-based authentication & authorization

HTTP Third-party Copy support

Database connection pool refactoring
(available starting from v1.12.0 release)

Latest release: StoRM v1.11.21

<https://italiangrid.github.io/storm/2021/05/12/storm-v1.11.21-released.html>

Mainly a bugfix release:

- fixes the StoRM v1.11.20 known issue which could break connections with MariaDB
- ensures mariadb service is started before both Frontend and Backend;
- fixes the failed state shown on stop/restart of Java services;
- provides a set of scripts used to edit from command line the storage space info related to a storage area.


Already deployed at INFN-CNAF

- See INFN CNAF storage report [presentation](#)

Available in UMD

Next release: v1.12.0

- planned by the end of 2021



StoRM deployments by version
(source: egee-bdii.cnaf.infn.it)
(~20 sites)

StoRM evolution: objectives

Drop unsupported dependencies

- i.e., anything Globus related

Support emerging AAI & QoS standards

- i.e., new WLCG Tape management REST API

Simplify service operations and deployment

- Service containerisation, K8S, ...

Reduce maintenance and evolution costs

- Current complexity mostly due to unused SRM “features”

Provide horizontal scalability for all StoRM components

- Currently the StoRM Backend cannot be replicated (SPOF)

Improve observability

- Consistent logging across services, metrics, tracing, ...

Globus dependencies retirement plan

StoRM currently depends on Globus:

- in the SRM frontend component (via the CGSI/gSOAP dependency)
- in the GridFTP server implementation

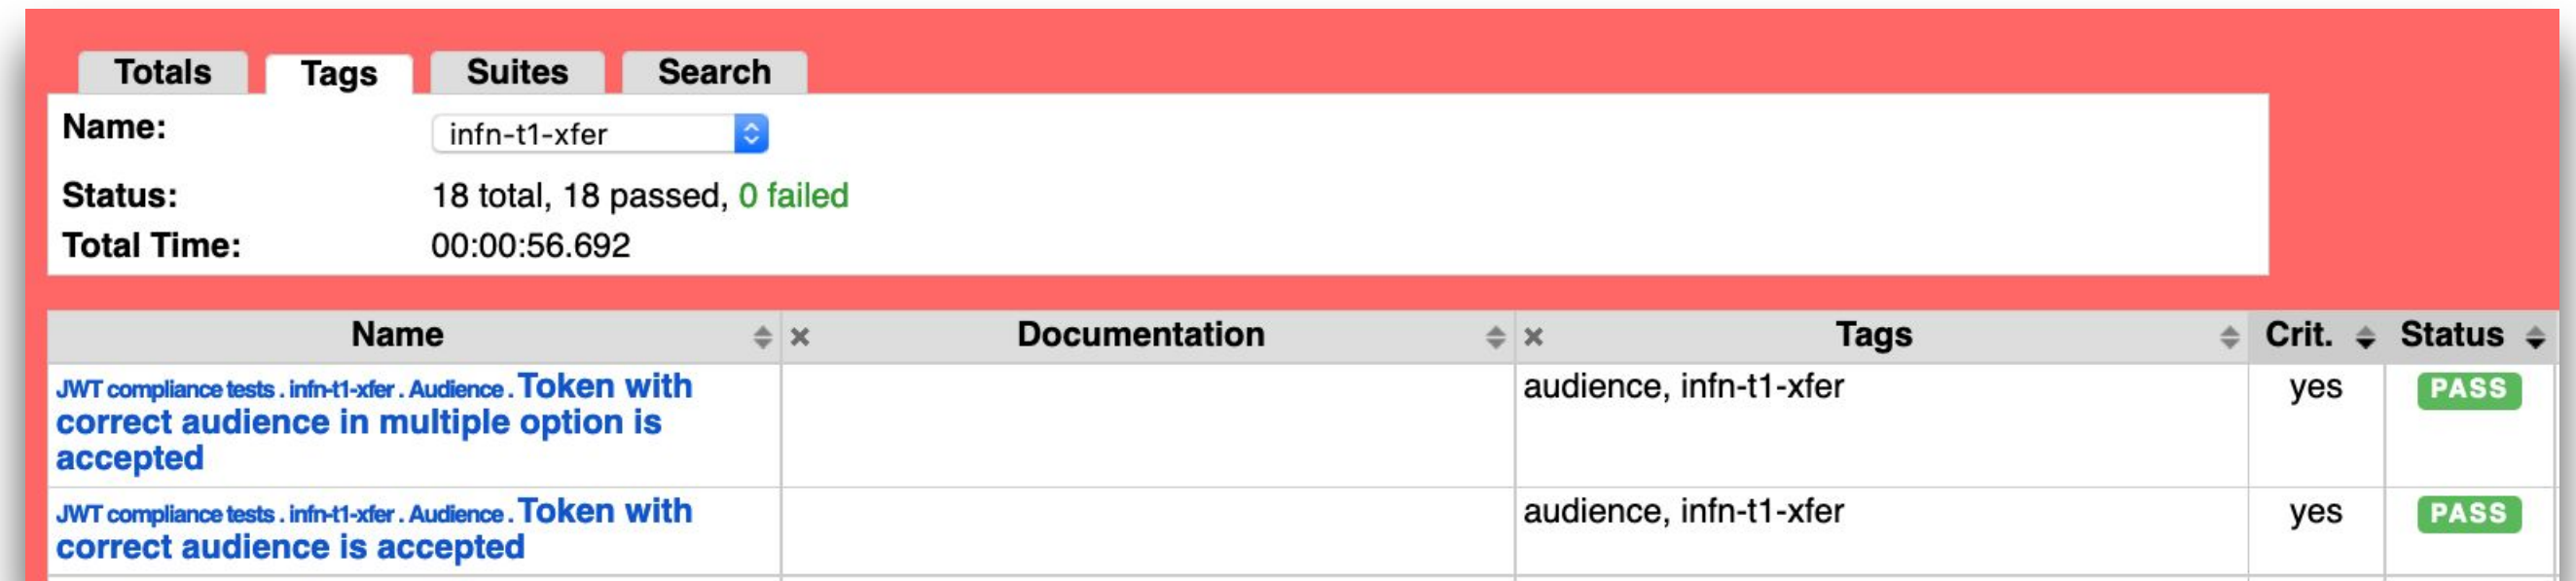
Plan:

- Develop Globus-free SRM frontend:
 - This means dropping support for the httpg protocol (i.e. HTTP over GSI)
 - **IN PROGRESS**, ETA: March 2022
- Provide GridFTP alternative supporting third-party copy → StoRM WebDAV
 - **DONE**

Token-based AuthN/Z support

StoRM WebDAV fully supports token-based authN/Z

- Storage-issued token support (aka "Macaroons") ✓
- OpenID Connect authentication support ✓
- OAuth access token support ✓
- WLCG JWT profile compliance:
 - audience enforcement support : ✓
 - capability-based authorization: ✓
 - group-based authorization: ✓



The screenshot shows a web interface for a test suite report. At the top, there are tabs for 'Totals', 'Tags', 'Suites', and 'Search'. Below these, a search bar contains 'inf-n-t1-xfer'. The status is '18 total, 18 passed, 0 failed' and the total time is '00:00:56.692'. Below this is a table with columns: Name, Documentation, Tags, Crit., and Status. Two rows are visible, both with a 'PASS' status.

Name	Documentation	Tags	Crit.	Status
JWT compliance tests . inf-n-t1-xfer . Audience . Token with correct audience in multiple option is accepted		audience, inf-n-t1-xfer	yes	PASS
JWT compliance tests . inf-n-t1-xfer . Audience . Token with correct audience is accepted		audience, inf-n-t1-xfer	yes	PASS

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

StoRM CDMI provides initial support for token-based AuthN/Z

Third-party copy support

StoRM supports HTTP-TPC in the StoRM WebDAV service

- GSI/Gridsite delegation is not supported
- RFC 3230 is supported (HTTP way of requesting checksum)
- Token-based authN/Z is supported (see previous slide)

There are no plans to support the XRootD TPC protocol in StoRM

StoRM SRM implementation already supports requesting HTTP as the transport protocol (with [appropriate configuration](#))

Storage Backends

StoRM services can run on top of any POSIX filesystem that:

- supports POSIX ACLs
- supports extended file attributes

StoRM is typically deployed on top of GPFS or Lustre

CEPH-FS is supported

- Tests are ongoing @CNAF to verify performance on top of CEPH-FS

StoRM WebDAV

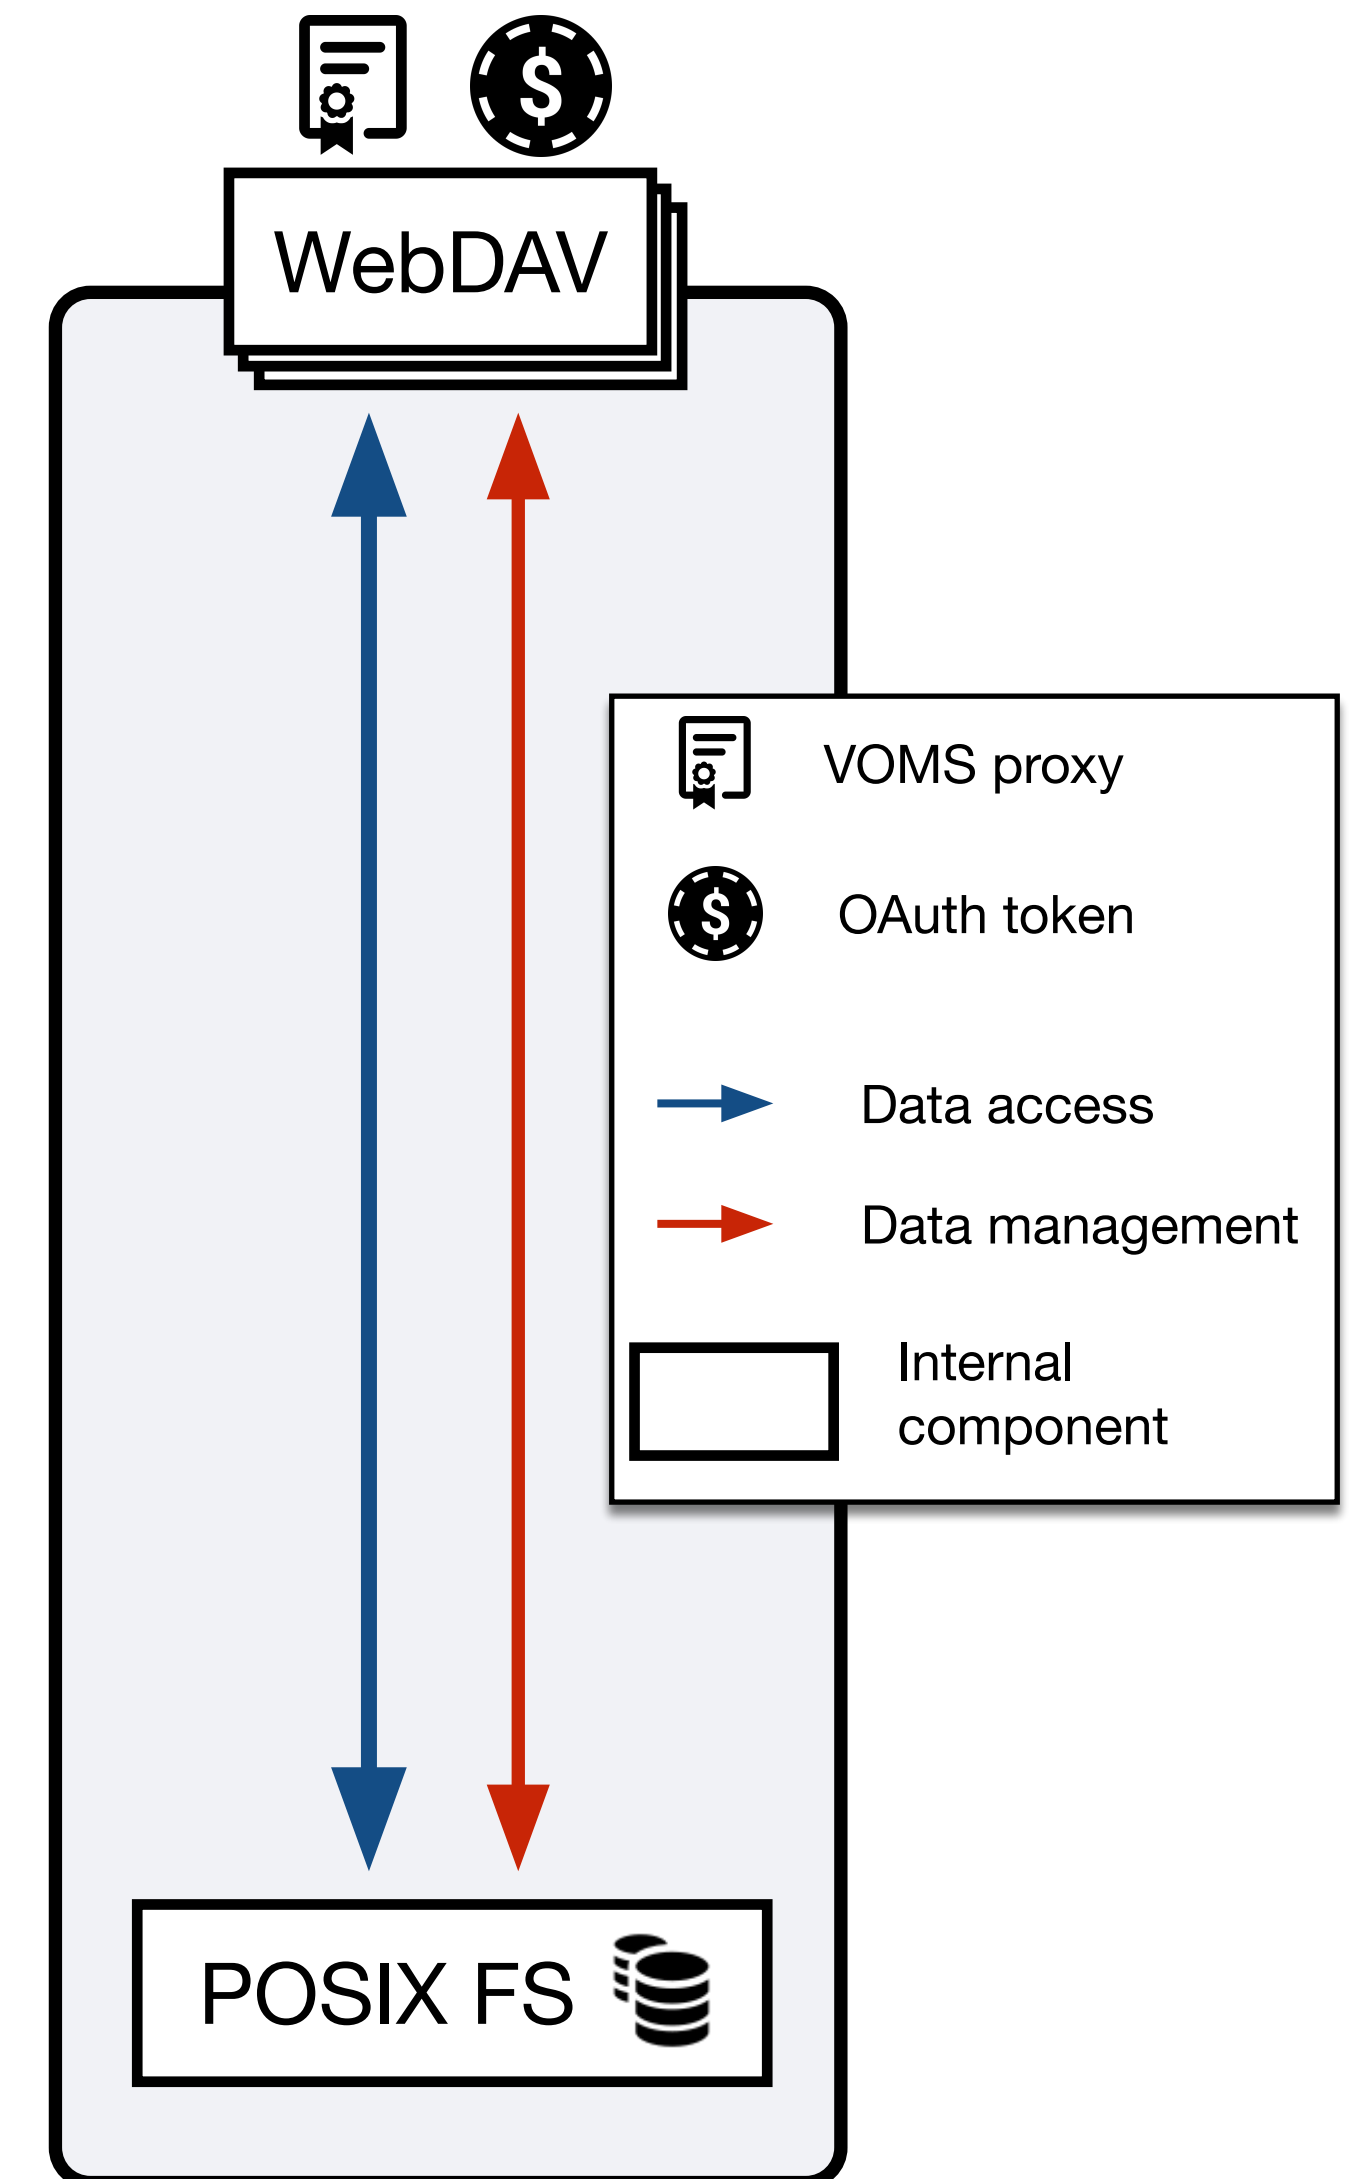
StoRM WebDAV service can run on top of any POSIX filesystem that supports extended file attributes

- GPFS, Lustre, CEPH-FS

StoRM WebDAV can be deployed as a standalone service, (no dependencies on the StoRM SRM backend)

- provides a browser-navigable interface
- supports WLCG JWT tokens and VOMS proxies
- supports HTTP Third Party Copy (HTTP-TPC)
- it can be easily configured using the [StoRM Puppet module](#)

```
$ puppet module install cnafsd-storm
```

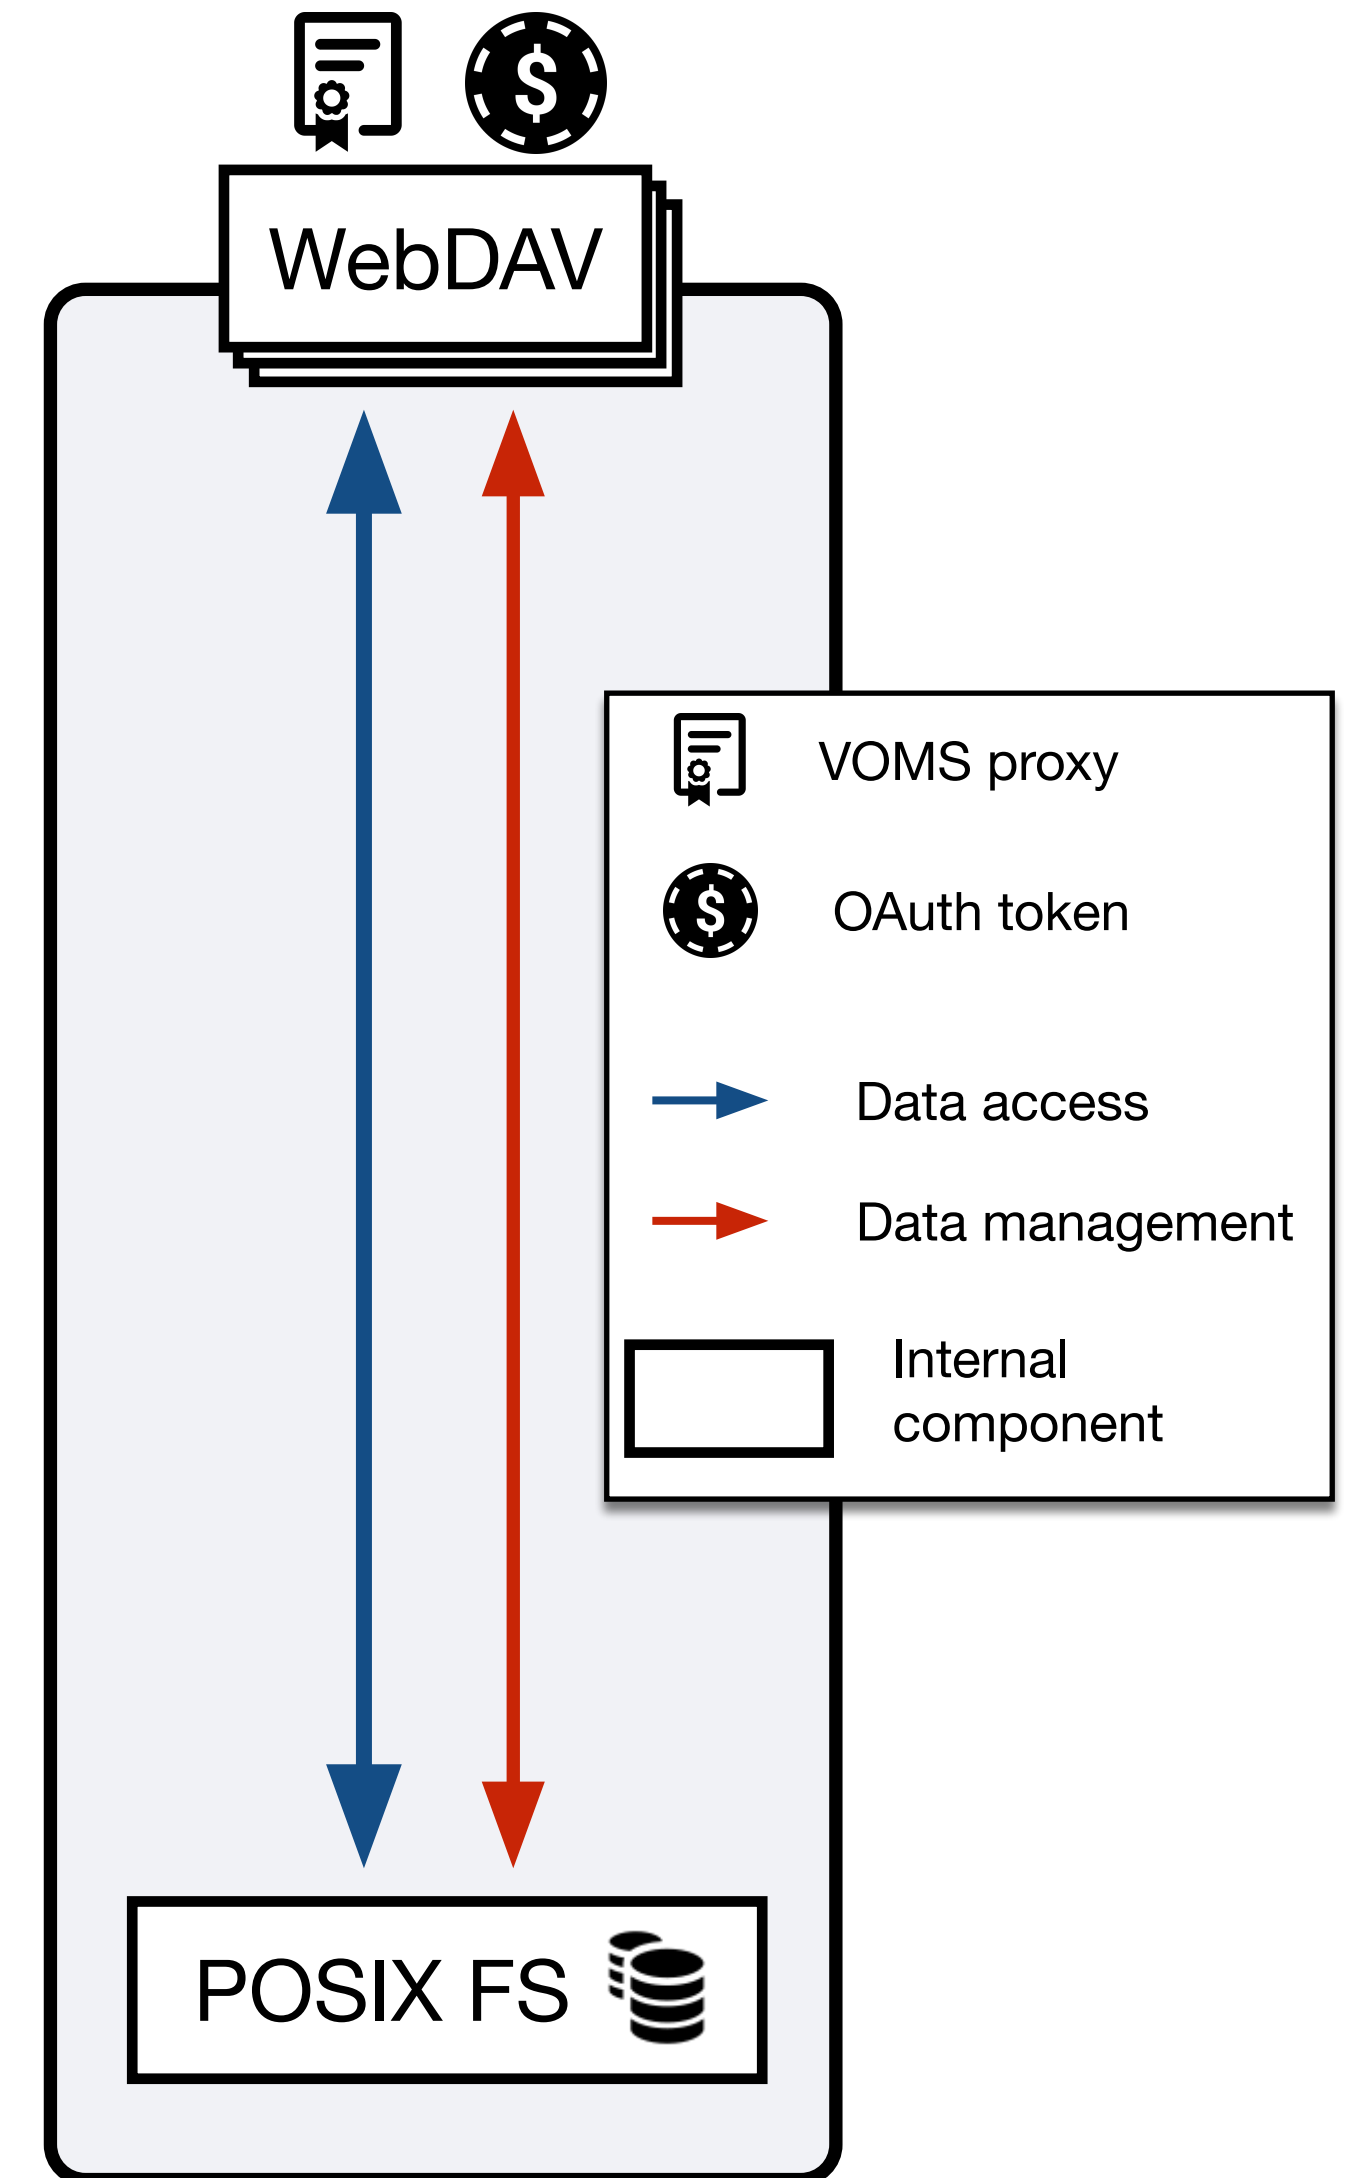


StoRM WebDAV

One disk storage area deployment example with VOMS proxy auth

```
class { 'storm::webdav':  
  storage_areas => [  
    {  
      name => 'escape',  
      root_path => '/storage/gpfs_escape',  
      access_points => ['/escape'],  
      vos => ['escape'],  
    },  
  ],  
  jvm_opts => '-Xms4192m -Xmx4192m',  
  use_conscrypt => true,  
  tpc_use_conscrypt => true,  
  enable_http2 => true,  
}  
storm::webdav::application_file { 'application-escape.yml':  
  source => '/path/to/your/application-escape.yml',  
}
```

```
tpc:  
  enable-tls-client-auth: false  
  max-connections: 500  
  use-conscrypt: true  
  max-connections-per-route: 330  
  progress-report-thread-pool-size: 16  
  timeout-in-secs: 30  
  
storm:  
  buffer:  
    file-buffer-size-bytes: 1048576
```



StoRM WebDAV

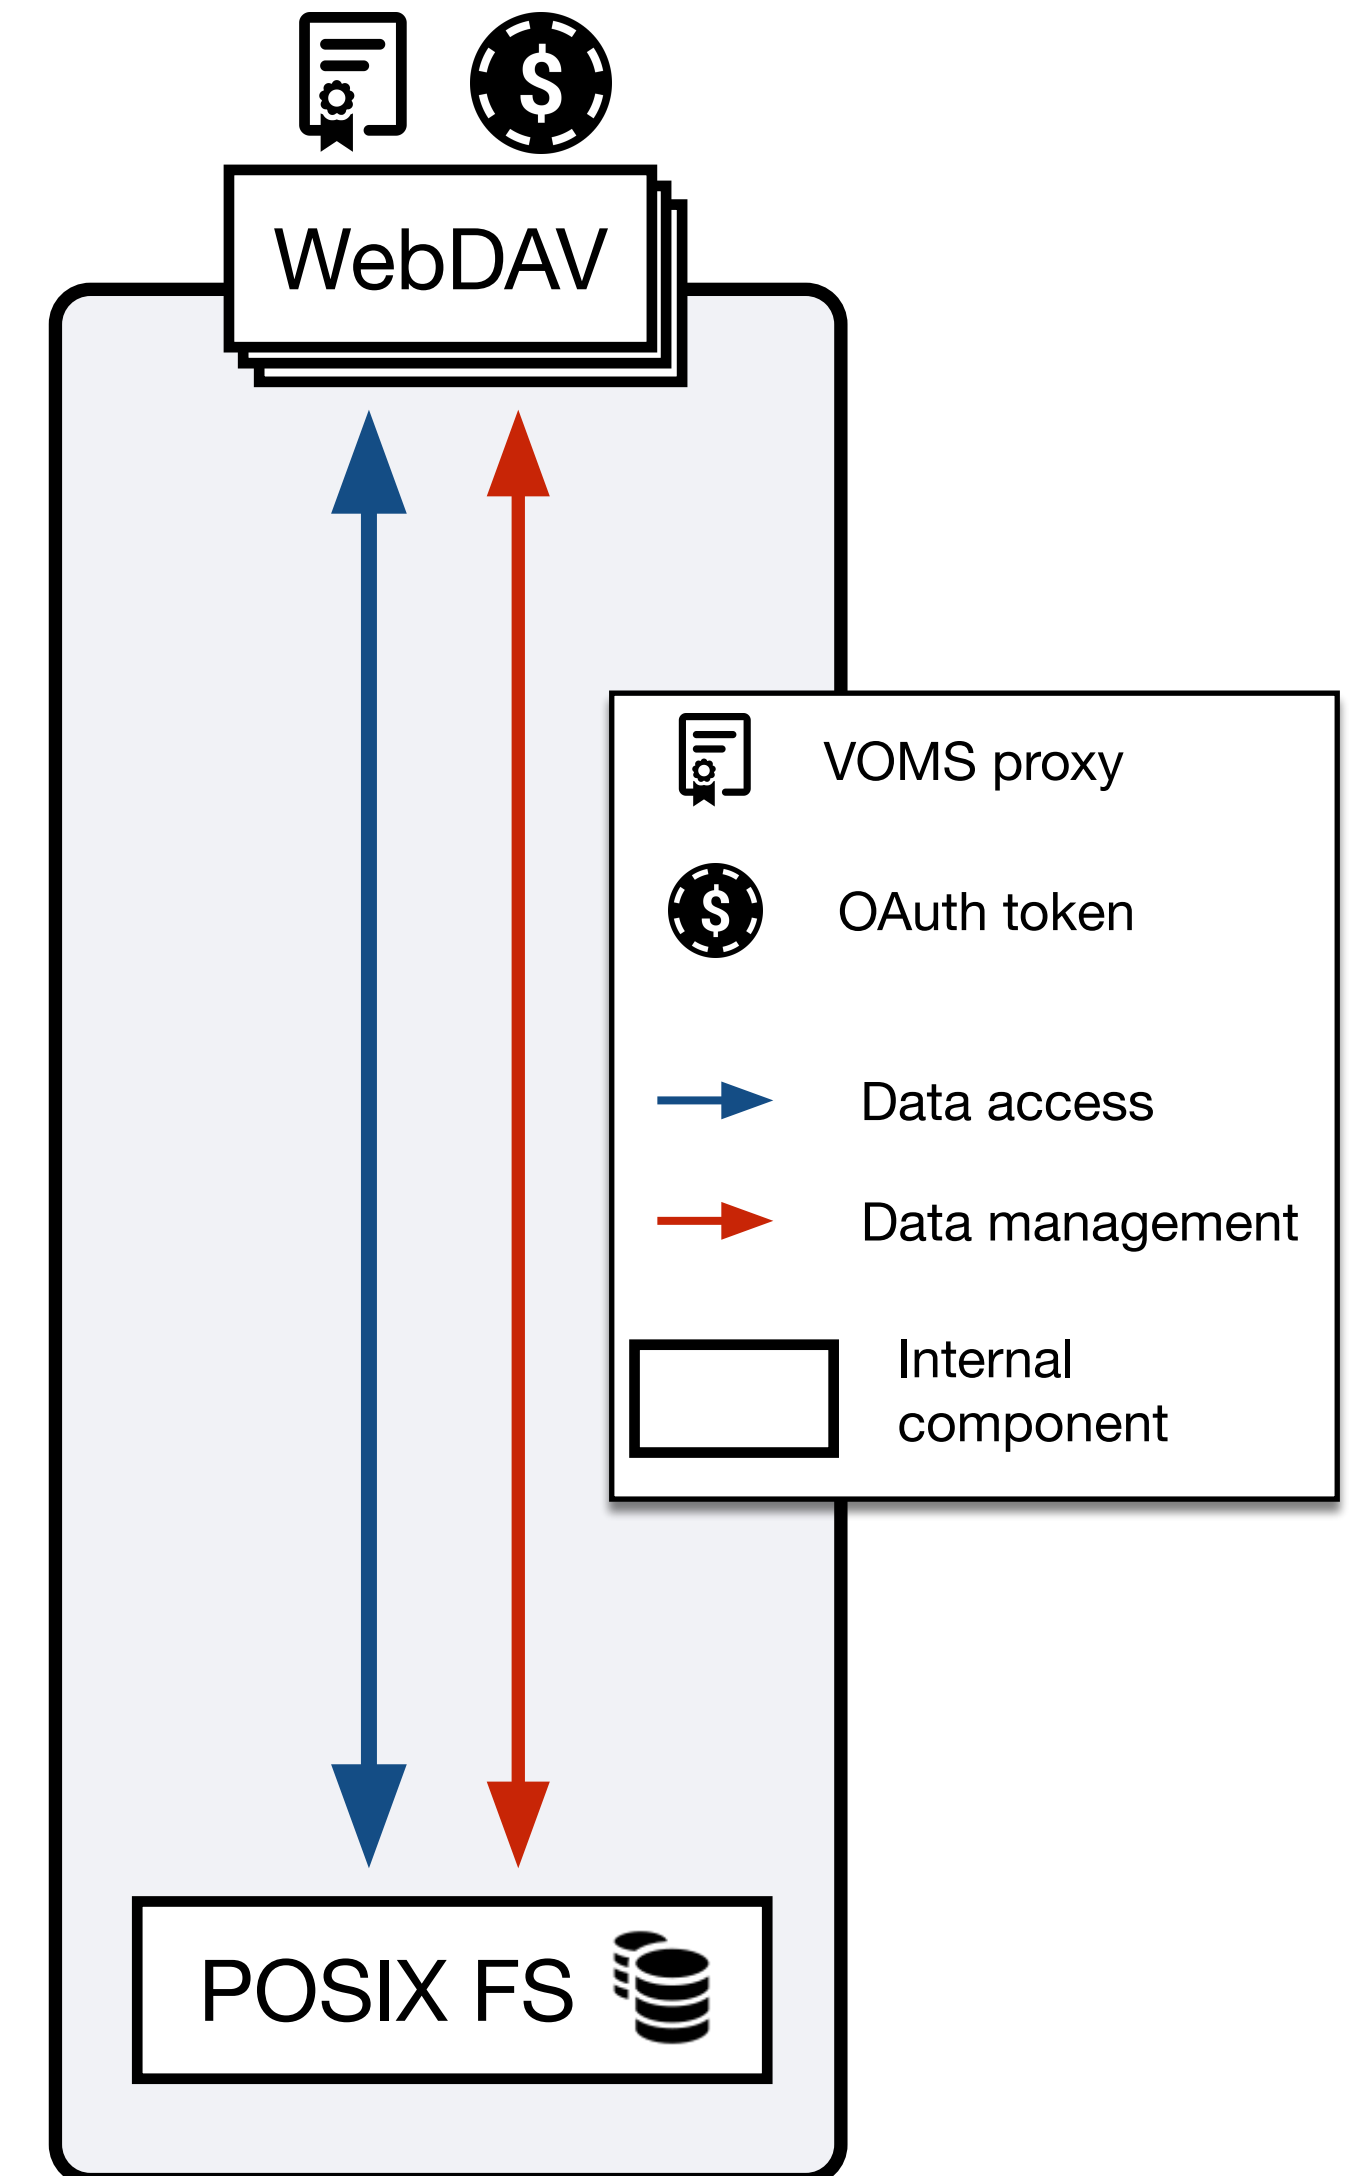
One disk storage area deployment example with VOMS proxy auth + OAuth token

```
class { 'storm::webdav':  
  storage_areas => [  
    {  
      name => 'escape',  
      root_path => '/storage/gpfs_escape',  
      access_points => ['/escape'],  
      vos => ['escape'],  
      orgs => ['https://iam-escape.cloud.cnaf.infn.it/'],  
    },  
  ],  
  jvm_opts => '-Xms4192m -Xmx4192m',  
  authz_server_enable => true,  
  authz_server_issuer  
    => 'https://xfer-escape.cr.cnaf.infn.it:8443',  
  authz_server_max_token_lifetime_sec => 14400,  
  authz_server_secret => '*****',  
  use_conscrypt => true,  
  tpc_use_conscrypt => true,  
  enable_http2 => true,  
}  
storm::webdav::application_file { 'application-escape.yml':  
  source => '/path/to/your/application-escape.yml',  
}
```

```
spring:  
  security:  
    oauth2:  
      client:  
        registration:  
          escape:  
            provider: escape  
            authorization-grant-type:  
              authorization_code  
            client-name: ESCAPE IAM  
            client-id: xfer.cr.cnaf.infn.it  
            client-secret: *****  
            scope:  
              - openid  
              - profile  
        oauth:  
          enable-oidc: true  
          issuers:  
            - name: escape  
              issuer: https://iam-escape.cloud.cnaf.infn.it/  
        tpc:  
          enable-tls-client-auth: false  
          max-connections: 500  
          use-conscrypt: true  
          max-connections-per-route: 330  
          progress-report-thread-pool-size: 16  
          timeout-in-secs: 30  
        storm:  
          buffer:  
            file-buffer-size-bytes: 1048576
```

StoRM Puppet module WebDAV class reference:

https://italiangrid.github.io/storm-puppet-module/puppet_classes/storm_3A_3Awebdav.html



Support, maintenance & evolution statement

StoRM is the SE adopted by the INFN CNAF data center, it will be maintained and evolved by INFN for the foreseeable future

- This includes providing support (through [GGUS](#)) to other StoRM-based sites



Ongoing work

- CENTOS 8 Stream porting
- SRM frontend refactoring
 - drop Globus dependencies
 - code cleanup
- New WLCG tape management REST API design & implementation
 - Github repo: <https://github.com/wlcg-storage/wlcg-tape-rest-api>
 - API Reference v1: https://docs.google.com/document/d/1Zx_H5dRkQRfju3xIYZ2WgjKoOvmLtsafP2pKGpHqcfY/

Thanks!
Questions?

Contacts and references

Sources:

StoRM Backend @ GitHub: <https://github.com/italiangrid/storm>

StoRM WebDAV @ GitHub: <https://github.com/italiangrid/storm-webdav>

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

StoRM Puppet module: [reference](#), [docs](#)

Contacts:

[StoRM users mailing list](#) storm-users@lists.infn.it

[StoRM developers mailing list](#) storm-devel@lists.infn.it