

StoRM evolution

Andrea Ceccanti, Enrico Vianello

GDB, June 10th 2020

andrea.ceccanti@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

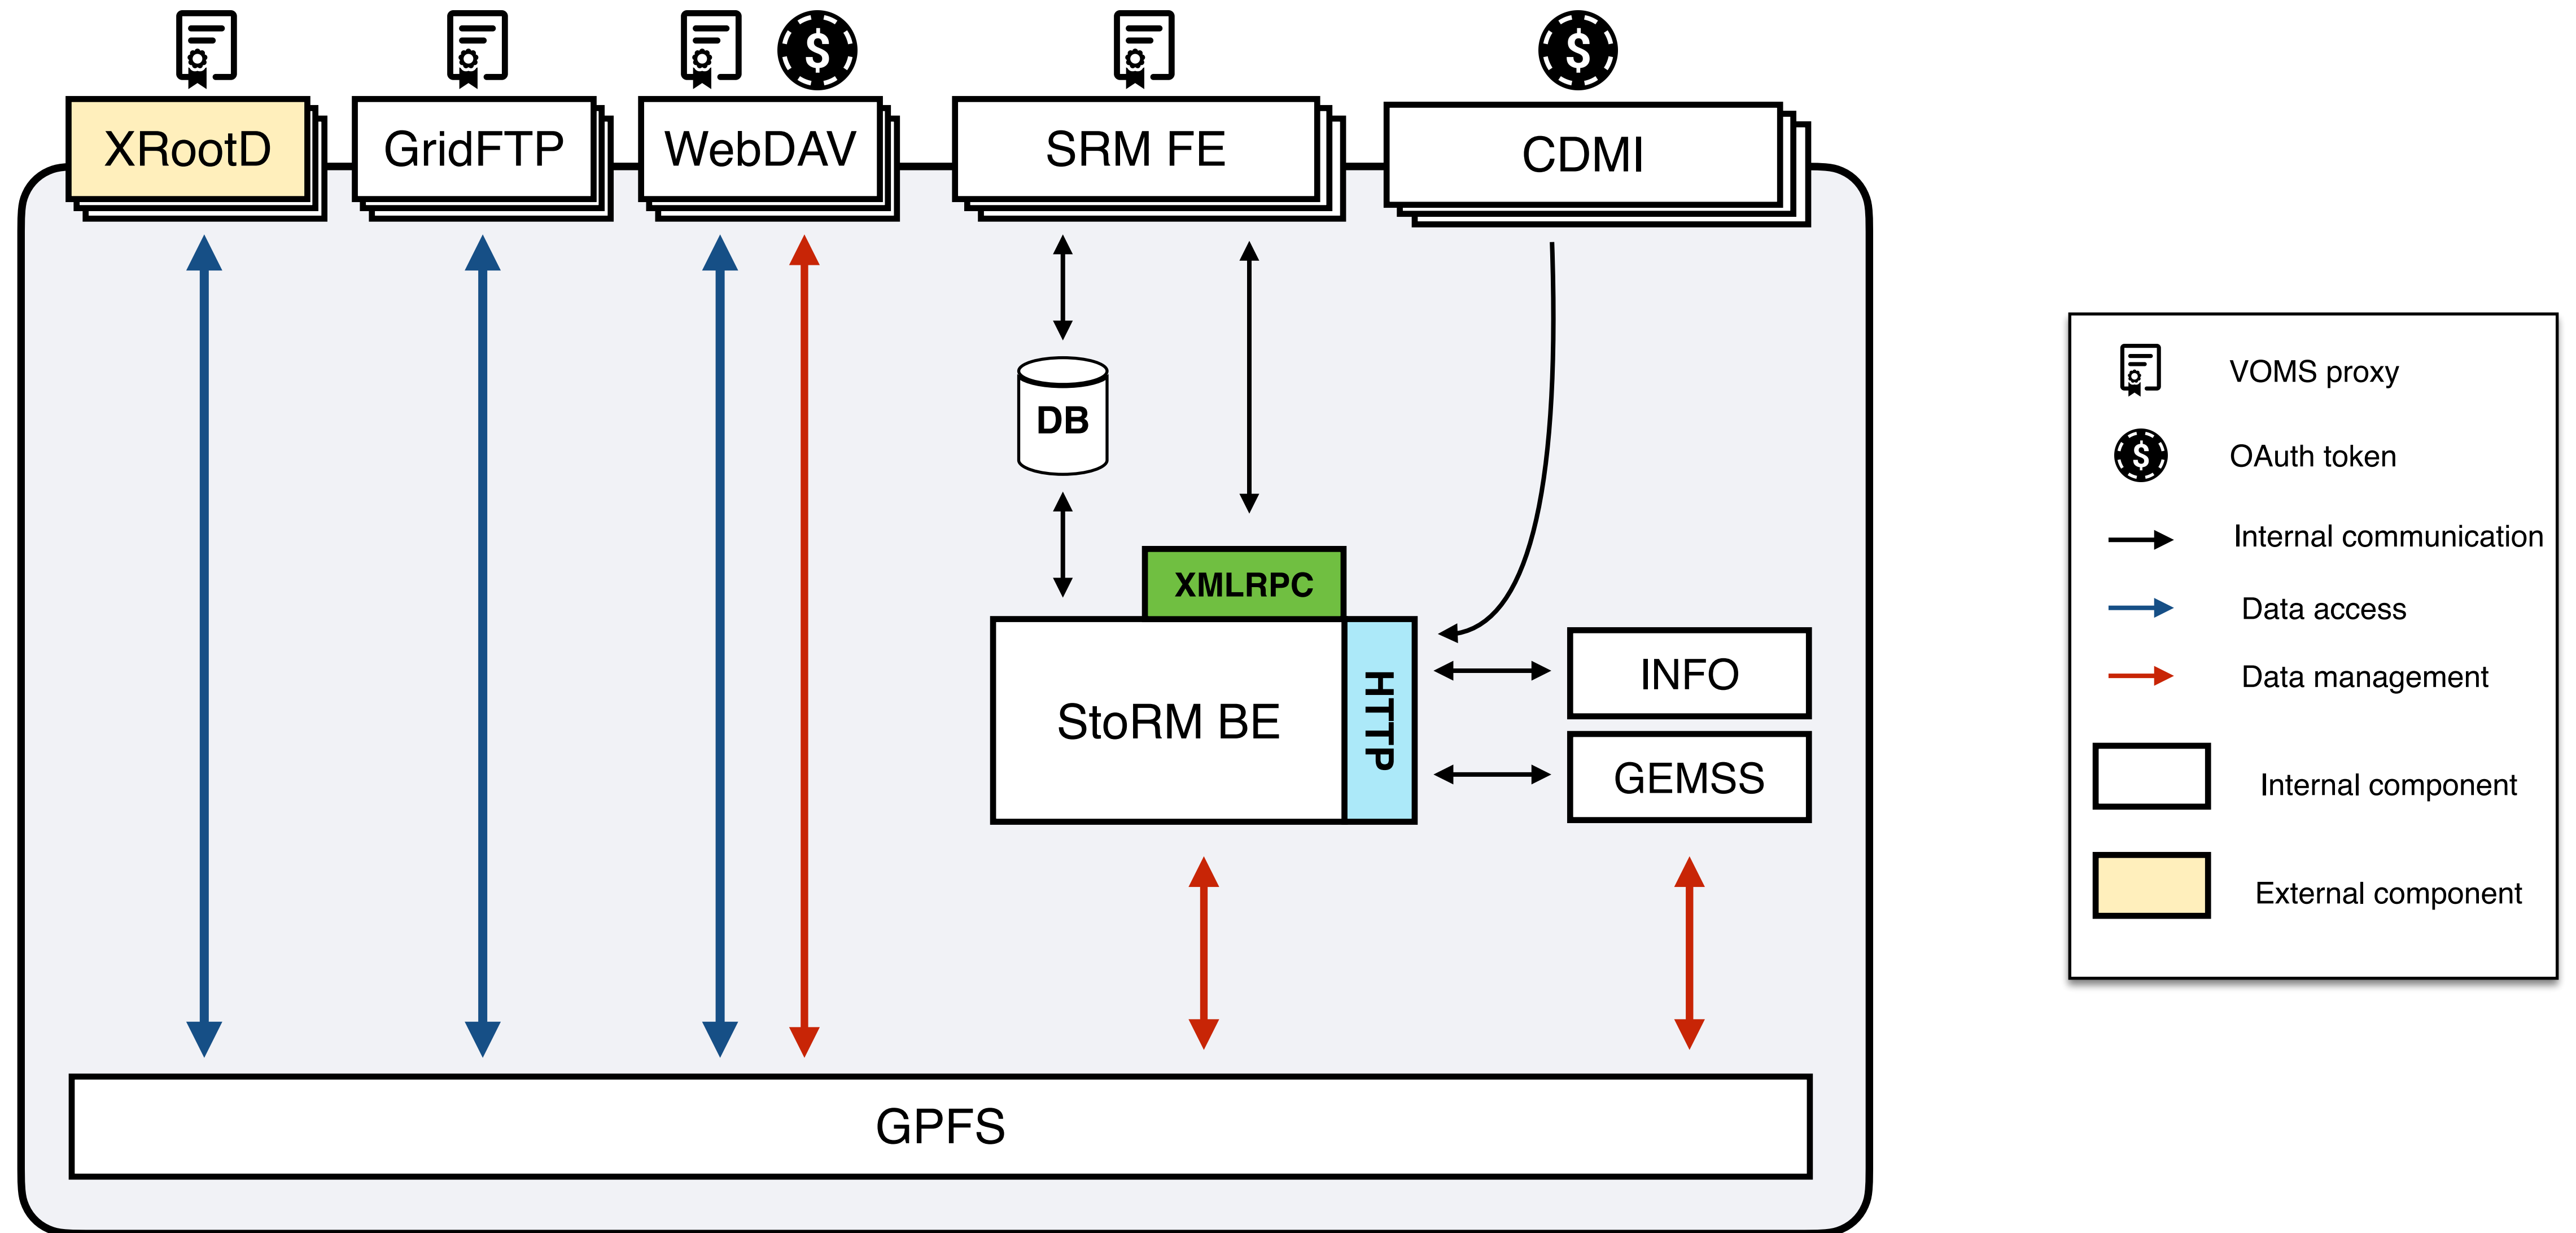
- Typically IBM GPFS or Lustre

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

Provides flexible AuthN/Z support:

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

StoRM: a typical deployment architecture



Main recent developments

CENTOS 7/SystemD porting

StoRM Puppet module

HTTP Third-party Copy support

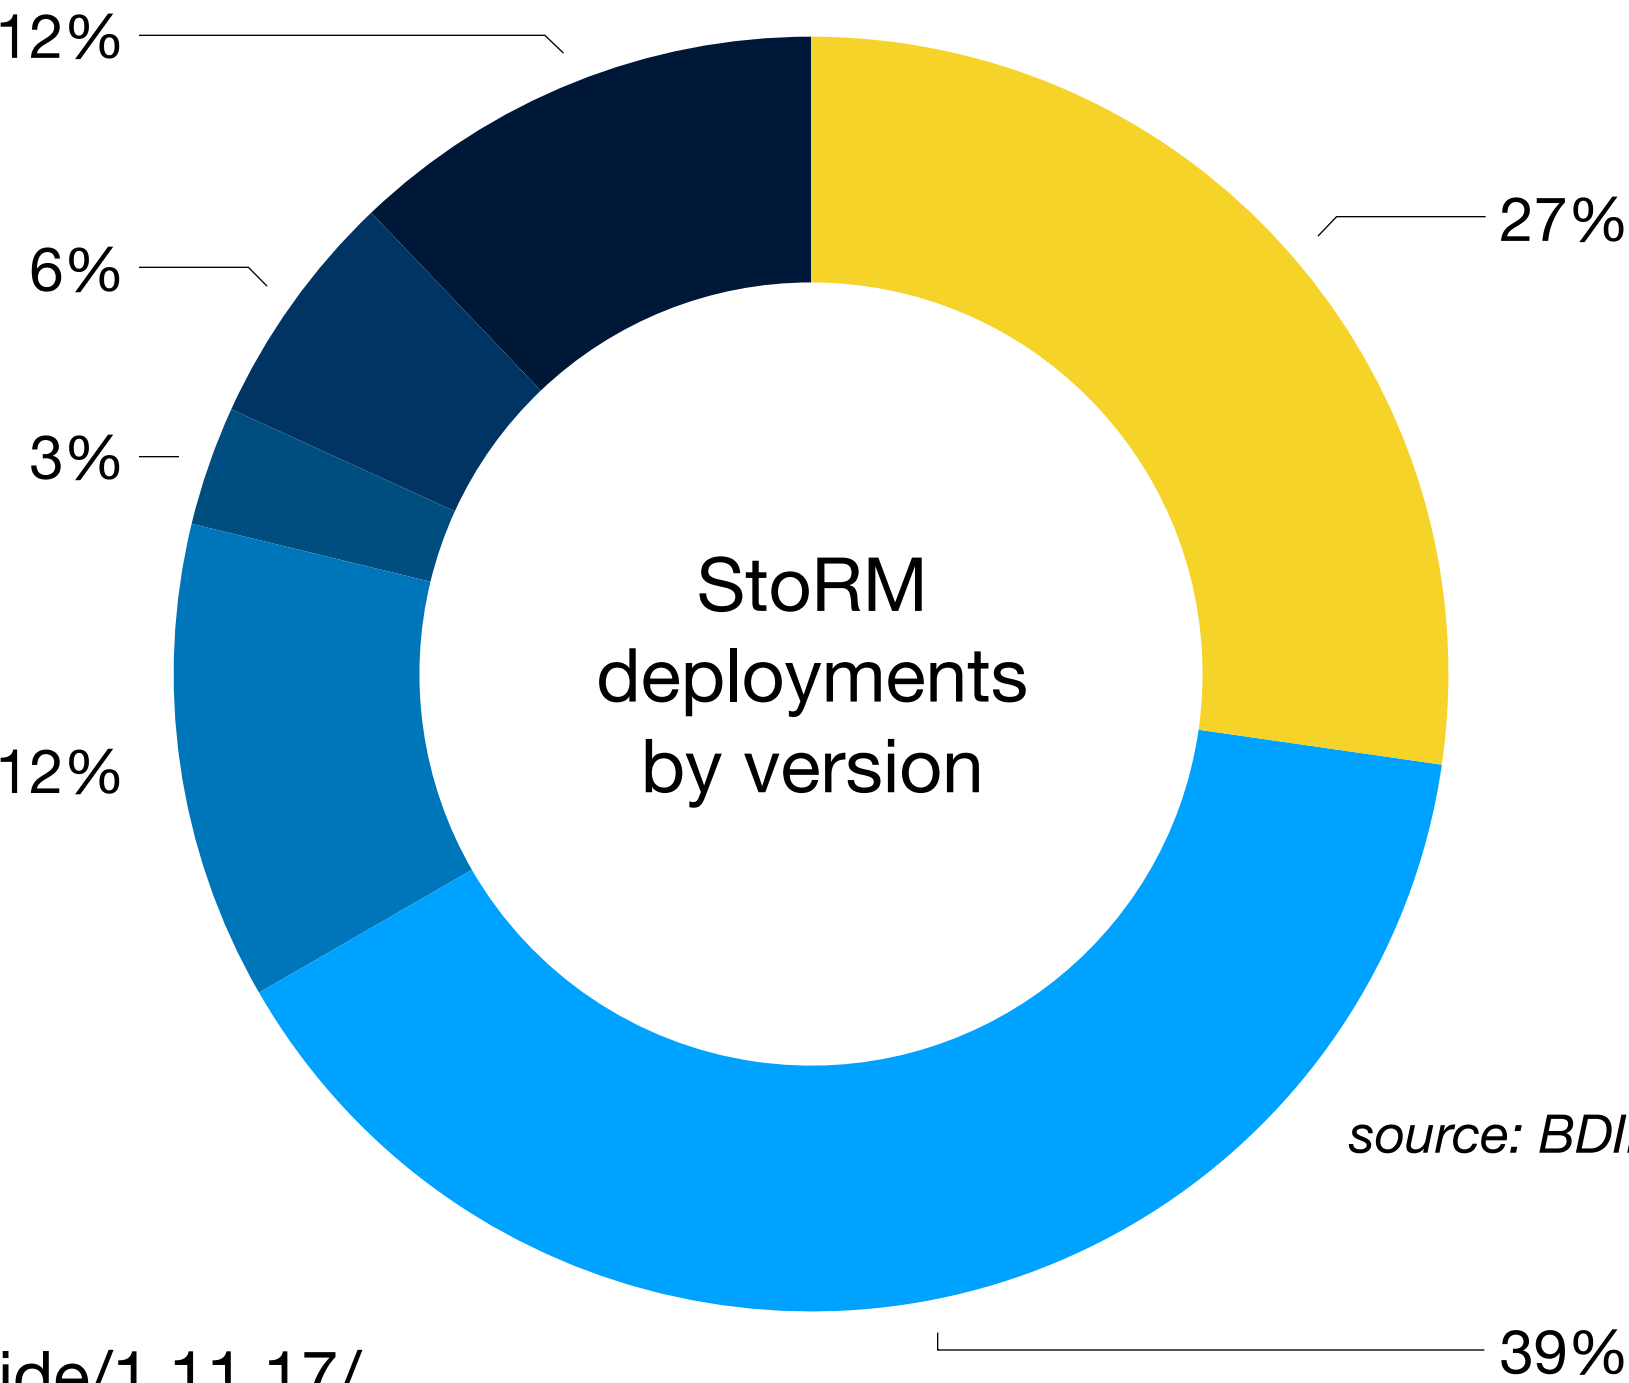

Token-based authentication & authorization

StoRM 1.11.17: latest release

Main changes:

- CENTOS 7 support for data transfer services: StoRM WebDAV and GridFTP

COMPONENT	VERSION	SUPPORTED PLATFORMS
StoRM WebDAV	1.2.1	CentOS6, CentOS7
StoRM GridFTP DSI module	1.2.3	CentOS6, CentOS7
StoRM Frontend	1.8.12	CentOS6
StoRM Backend	1.11.17	CentOS6
StoRM Info Provider	1.8.1	CentOS6
CDMI StoRM plugin	0.1.0	CentOS7



- References
 - StoRM downloads page <http://italiangrid.github.io/storm/download.html>
 - StoRM System Administration Guide <http://italiangrid.github.io/storm/documentation/sysadmin-guide/1.11.17/>

StoRM 1.11.18: release in preparation

Main new features:

- CENTOS 7 and puppet configuration support for frontend and backend services
- OpenID Connect authentication support on StoRM WebDAV

Main bug fixes:

- STOR-1207 StoRM WebDAV leaks file descriptors when Conscript is enabled
- STOR-932 Make Background DU configurable to run periodically in order to update used space info on db
- STOR-1197 StoRM Webdav should drop Authorization header in TPC redirects

ETA: end of June 2020

StoRM evolution: longer-term objectives

Drop unsupported dependencies

- i.e., anything Globus related

Simplify service operations and deployment

- Service containerisation, K8S, ...

Reduce maintenance and evolution costs

- Current complexity mostly due to unused SRM “features”

Support emerging AAI & QoS standards

- For AAI the direction it's clear, for QoS less so...

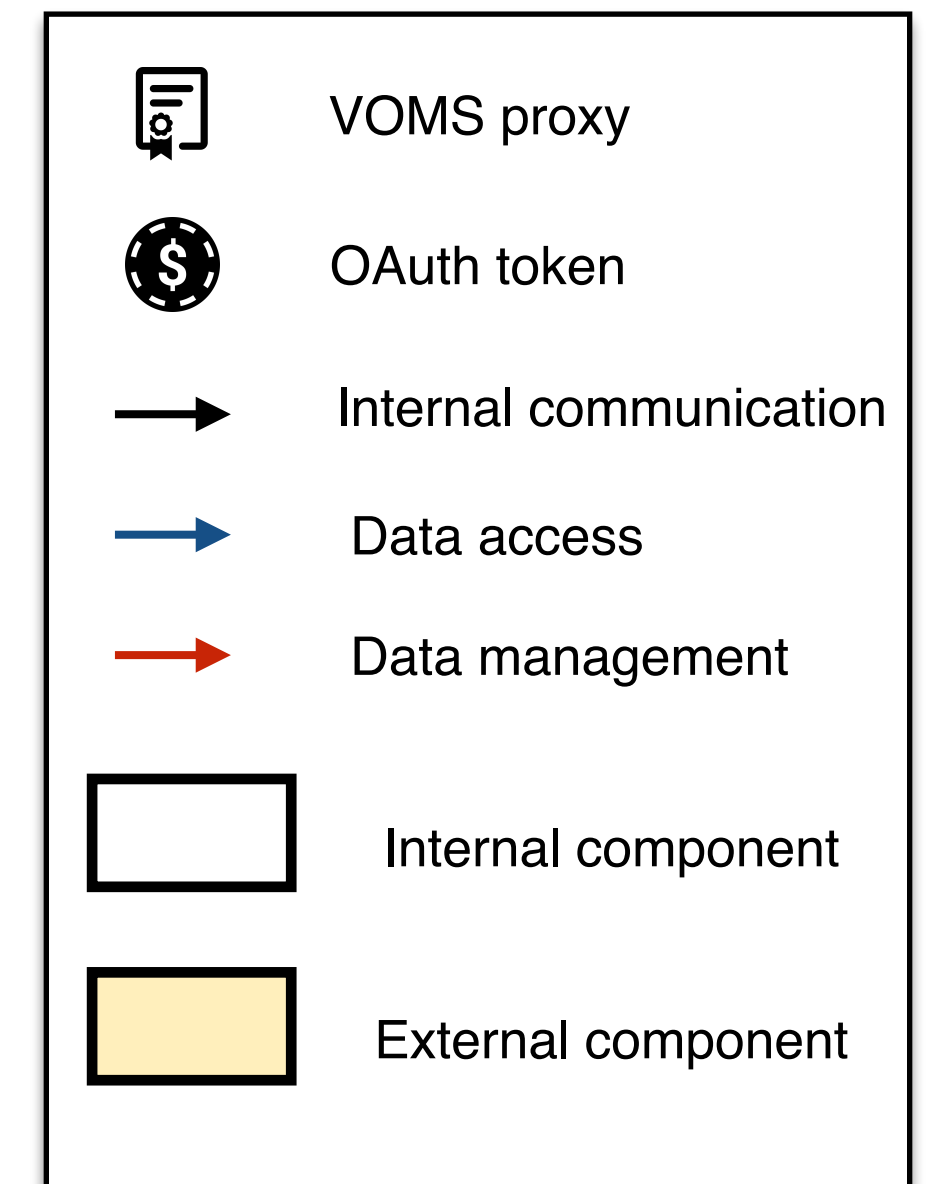
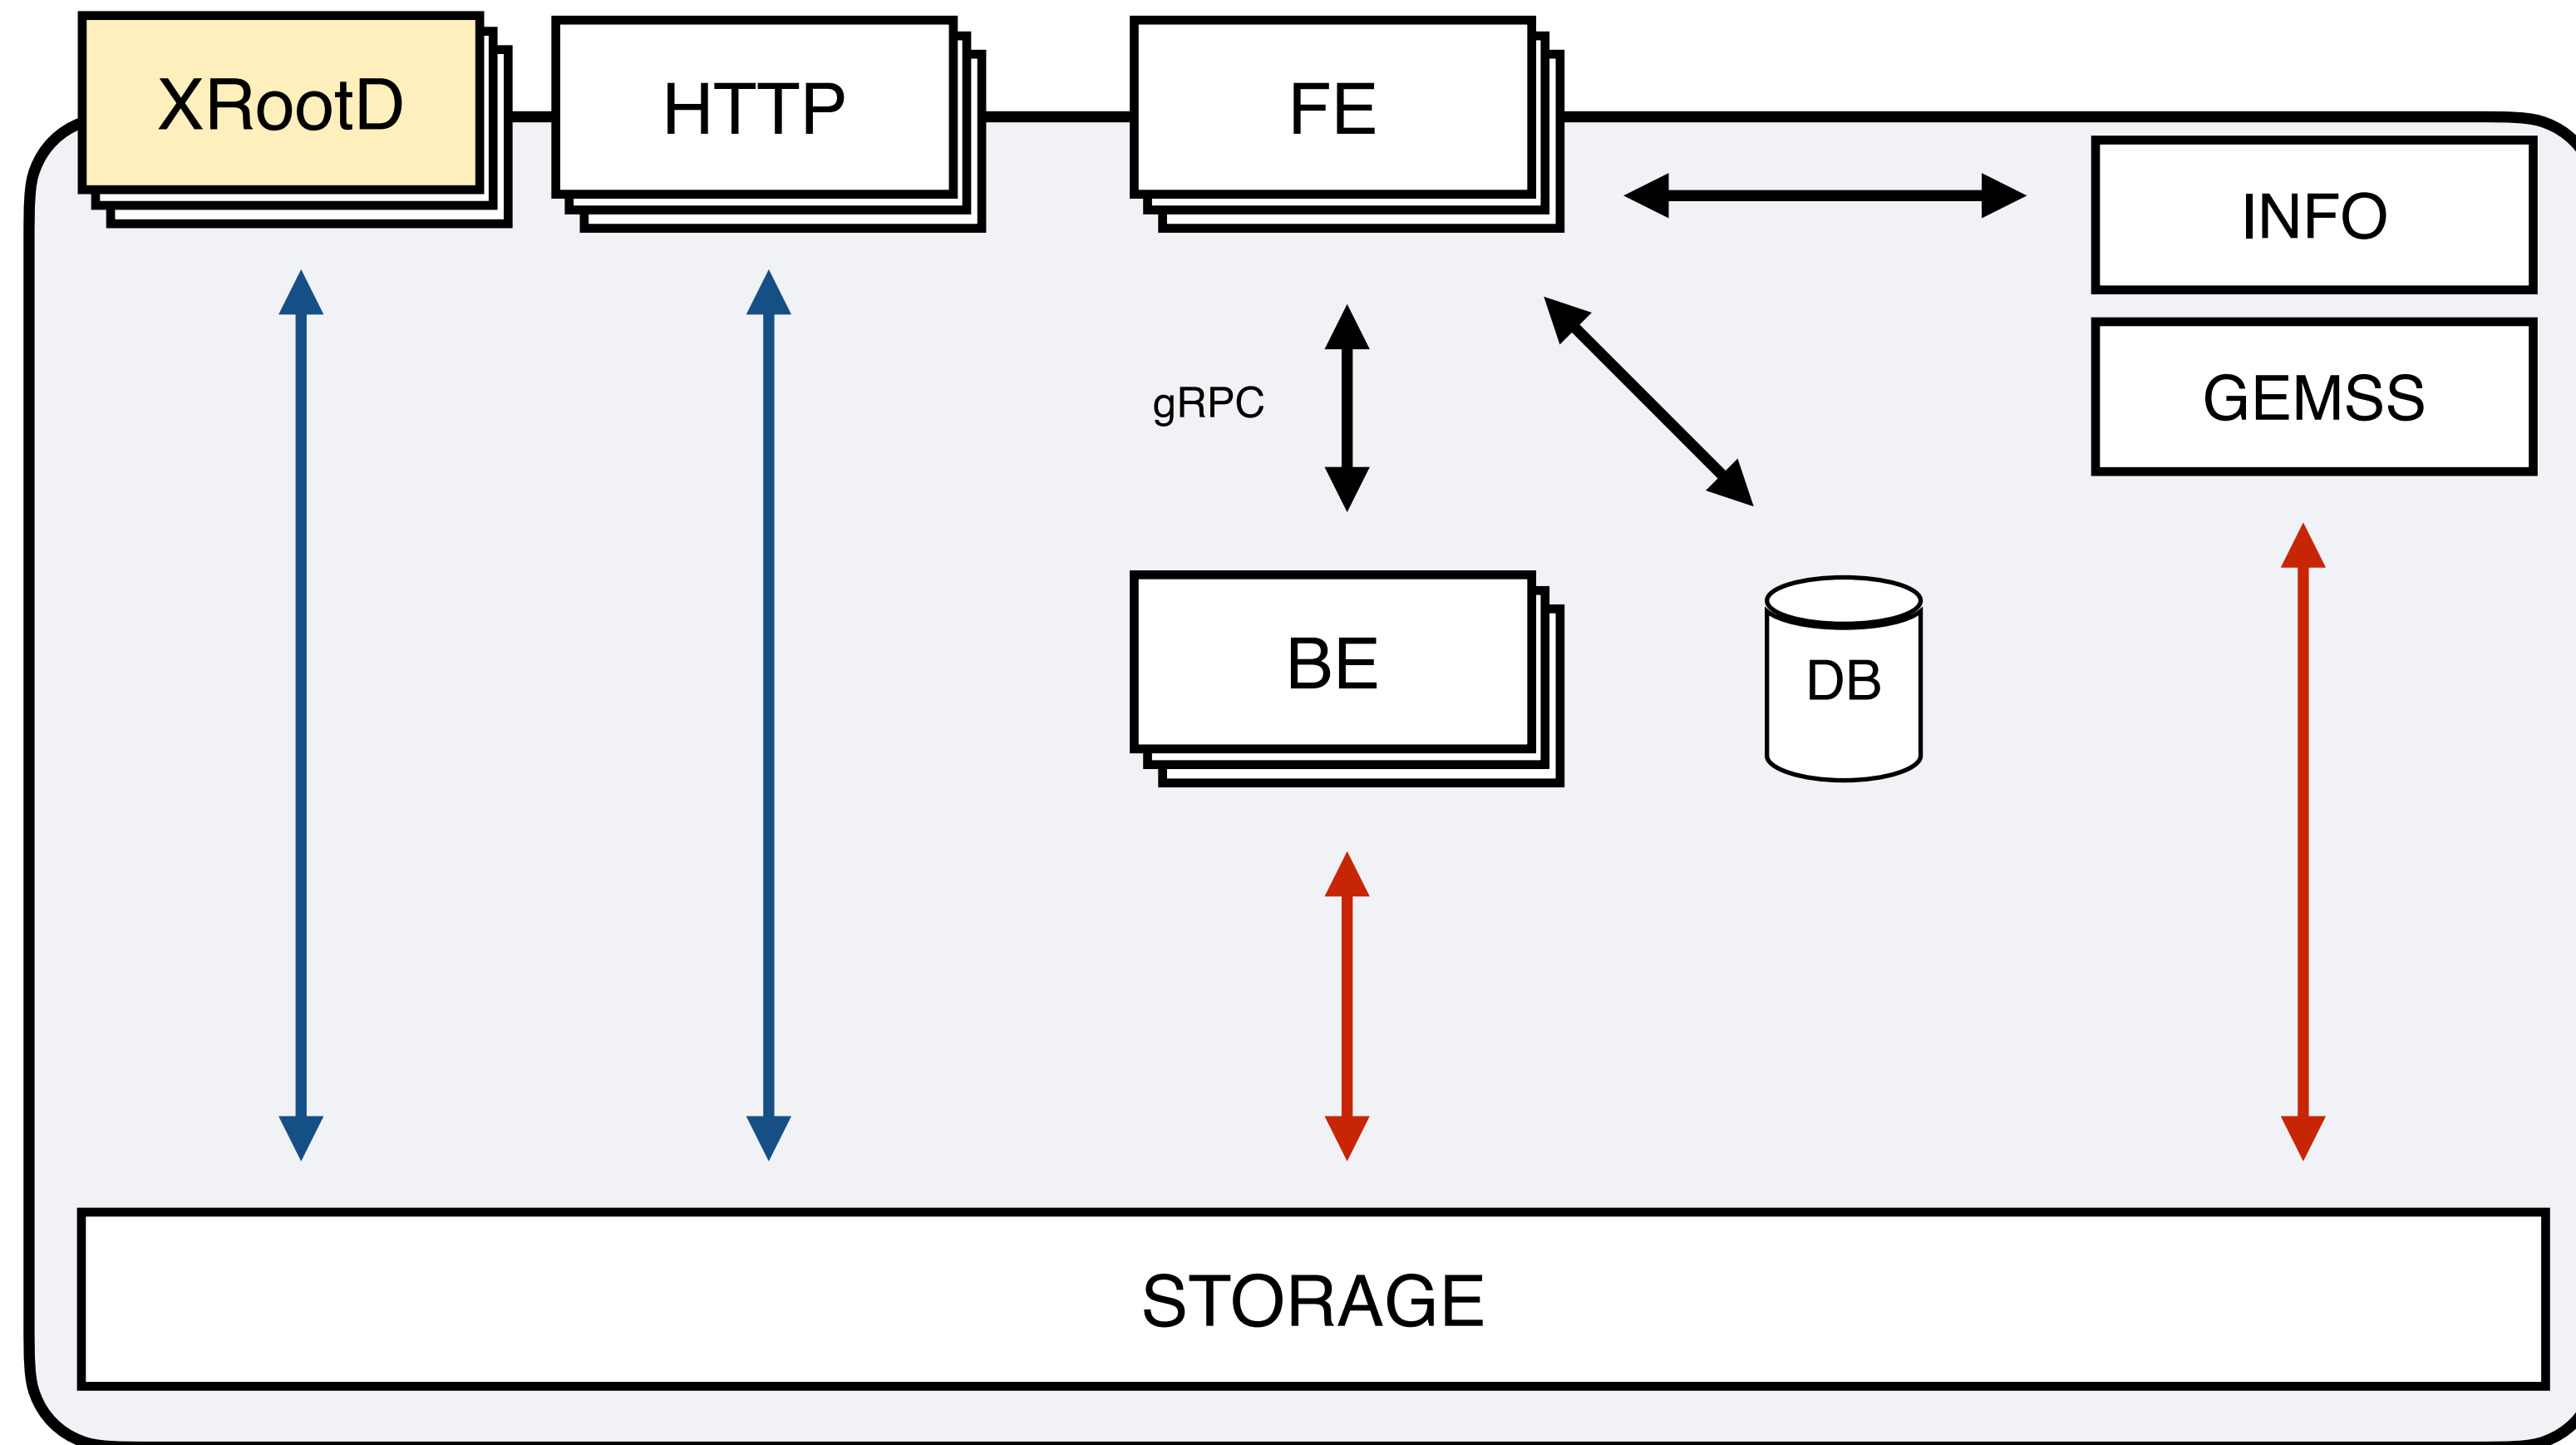
Provide horizontal scalability for all StoRM components

- Currently the StoRM Backend cannot be replicated (SPOF)

Improve observability

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

StoRM 2: high-level architecture



The new StoRM frontend



Implements management protocol endpoints

- SRM (the minimal subset needed to support current experiment requirements)
- WebDAV, CDMI, ?


Authentication, Authorization, Mapping

- X.509, VOMS, OAuth, ...

Translates incoming requests to a set of management operations on the storage that are executed by the backend

Stateless Spring Boot application

The new StoRM backend



Implements management operations on the storage

Stateless worker, completely decoupled from the frontend

- the two components can be scaled independently

Implemented in C++

Communication between fe and be implemented using gRPC

StoRM 2 and Globus/GSI

The new codebase does not have dependencies on Globus

- No support for GSI or Gridsite delegation

X.509 and VOMS are obviously supported

- Through VOMS libraries and a VOMS-aware NGINX module

StoRM 2 and QoS

StoRM 2 will keep a minimal SRM implementation for the management of tape recalls

- the only standard interface AFAIK that can handle bulk tape recall requests

StoRM 2 will also keep the CDMI interface developed during the INDIGO DataCloud project

- which, however, does not support bulk requests

We are also interested in developing a REST-ful alternative to SRM BringOnline

StoRM 2: when will it be ready?

Hard to say, since the development effort at INFN CNAF is limited and spread over many components

- IAM, StoRM 1, VOMS, Argus

Ideally, we want to be in good shape before the Globus end of support (end of 2021)

Thanks!
Questions?

Contacts and references

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

StoRM documentation: <http://italiangrid.github.io/storm/>

CDMI StoRM: <https://github.com/italiangrid/cdmi-storm>

Contacts:

andrea.ceccanti@cnaa.infn.it

enrico.vianello@cnaa.infn.it

storm-users@lists.infn.it