

Reva and EOS:

A lovely story



Hugo G. Labrador

04/02/2020


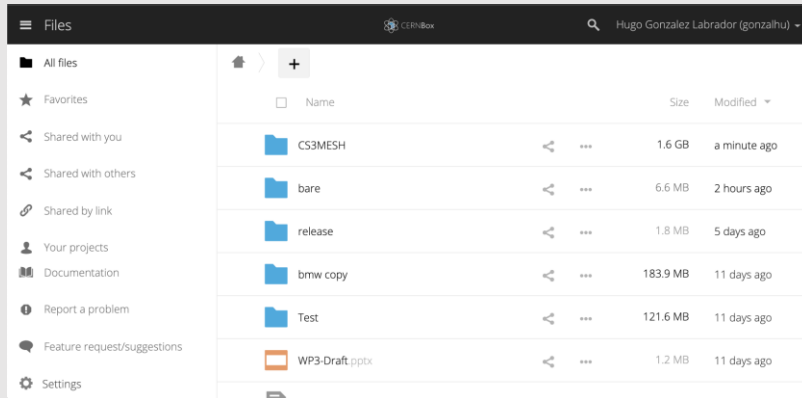
Eos Workshop - CERN



Reva is a nutshell

[More info at: https://reva.link/](https://reva.link/)

Your sync and share box


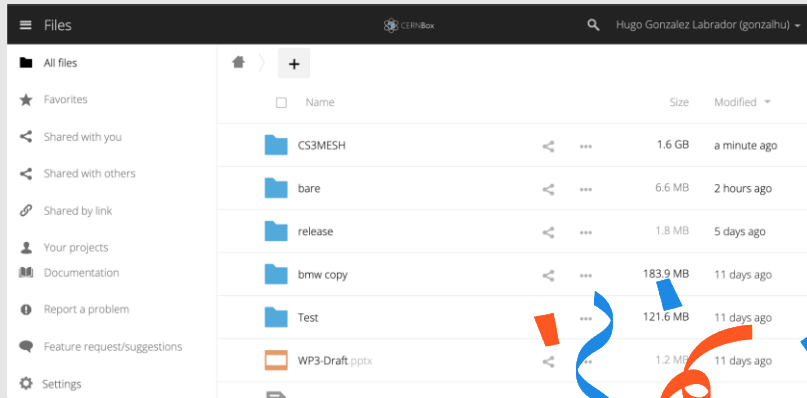


Other storages as well
(cephfs, gluster, USB, ...)


Access through the back
not feasible



Our box



CERNBox



Reva



MySQL™




OS

OS

OS



The next box

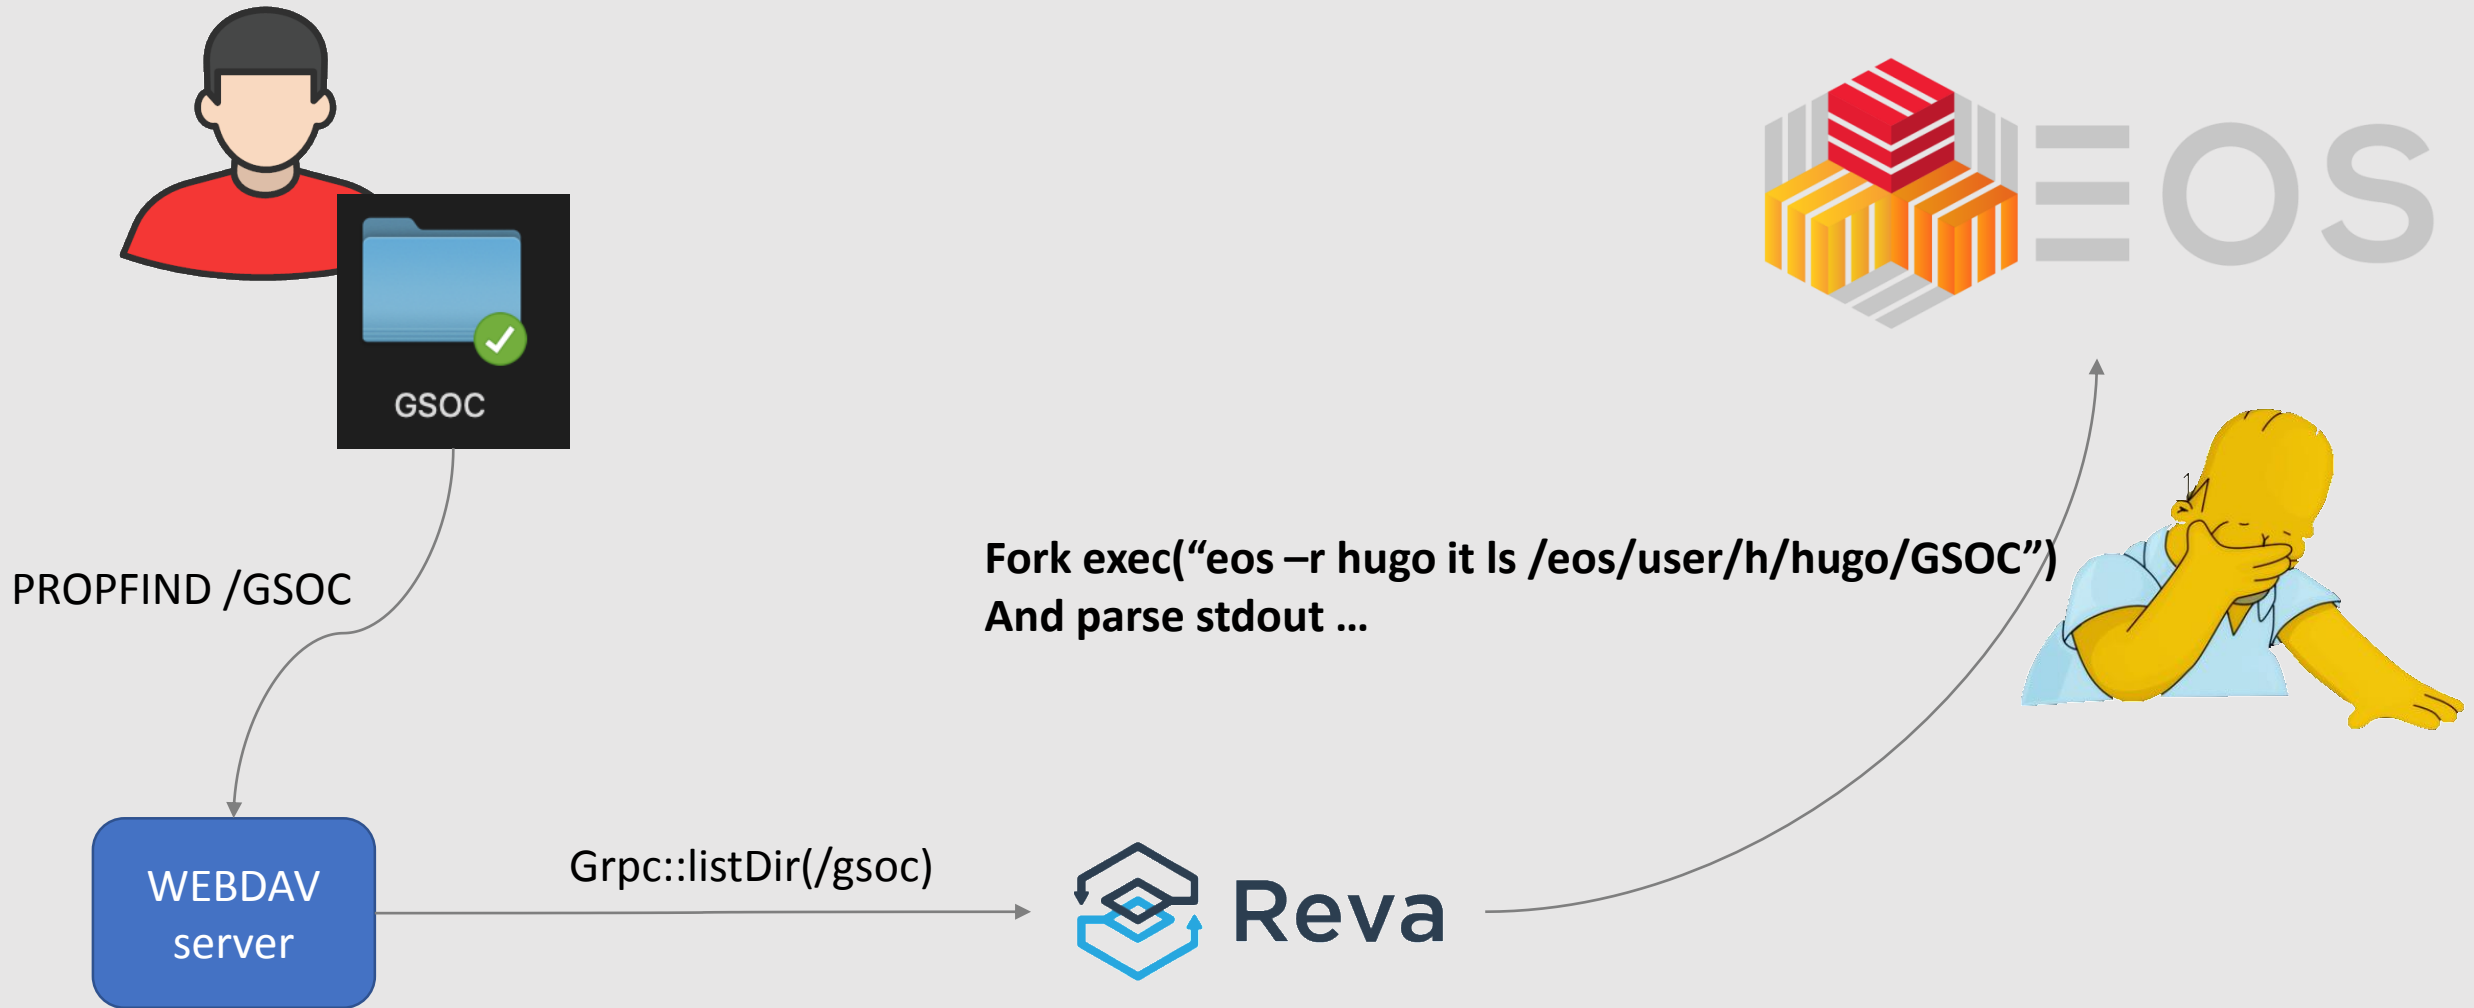


MENU

DEMO

Protocols, protocolos and more protocolos

List my files ...








```
26 service Eos {  
27   // Replies to a ping  
28   rpc Ping (PingRequest) returns (PingReply) {}  
29   //  
30 }
```



protobuf
Protocol Buffers

Andreas is the wizard of this, ask him ...

gRPC



```
1 package main
2
3 import "github.com/cernbox/eosgrpc-go/eos"
4 import "google.golang.org/grpc"
5 import "log"
6 import "context"
7
8 func main() {
9     host := "myeos.cern.ch:50051"
10    con, err := grpc.Dial(host, grpc.WithInsecure())
11    if err != nil {
12        log.Fatal(err)
13    }
14
15    client := eos.NewEosClient(con)
16
17    log.Println("Sending Ping message to " + host)
18    ctx := context.Background()
19    req := &eos.PingRequest{Authkey: "hugo-test-key", Message: []byte("hola")}
20    res, err := client.Ping(ctx, req)
21    if err != nil {
22        log.Fatal(err)
23    }
24    log.Printf("Server replied with %+v\n", res)
25 }
```

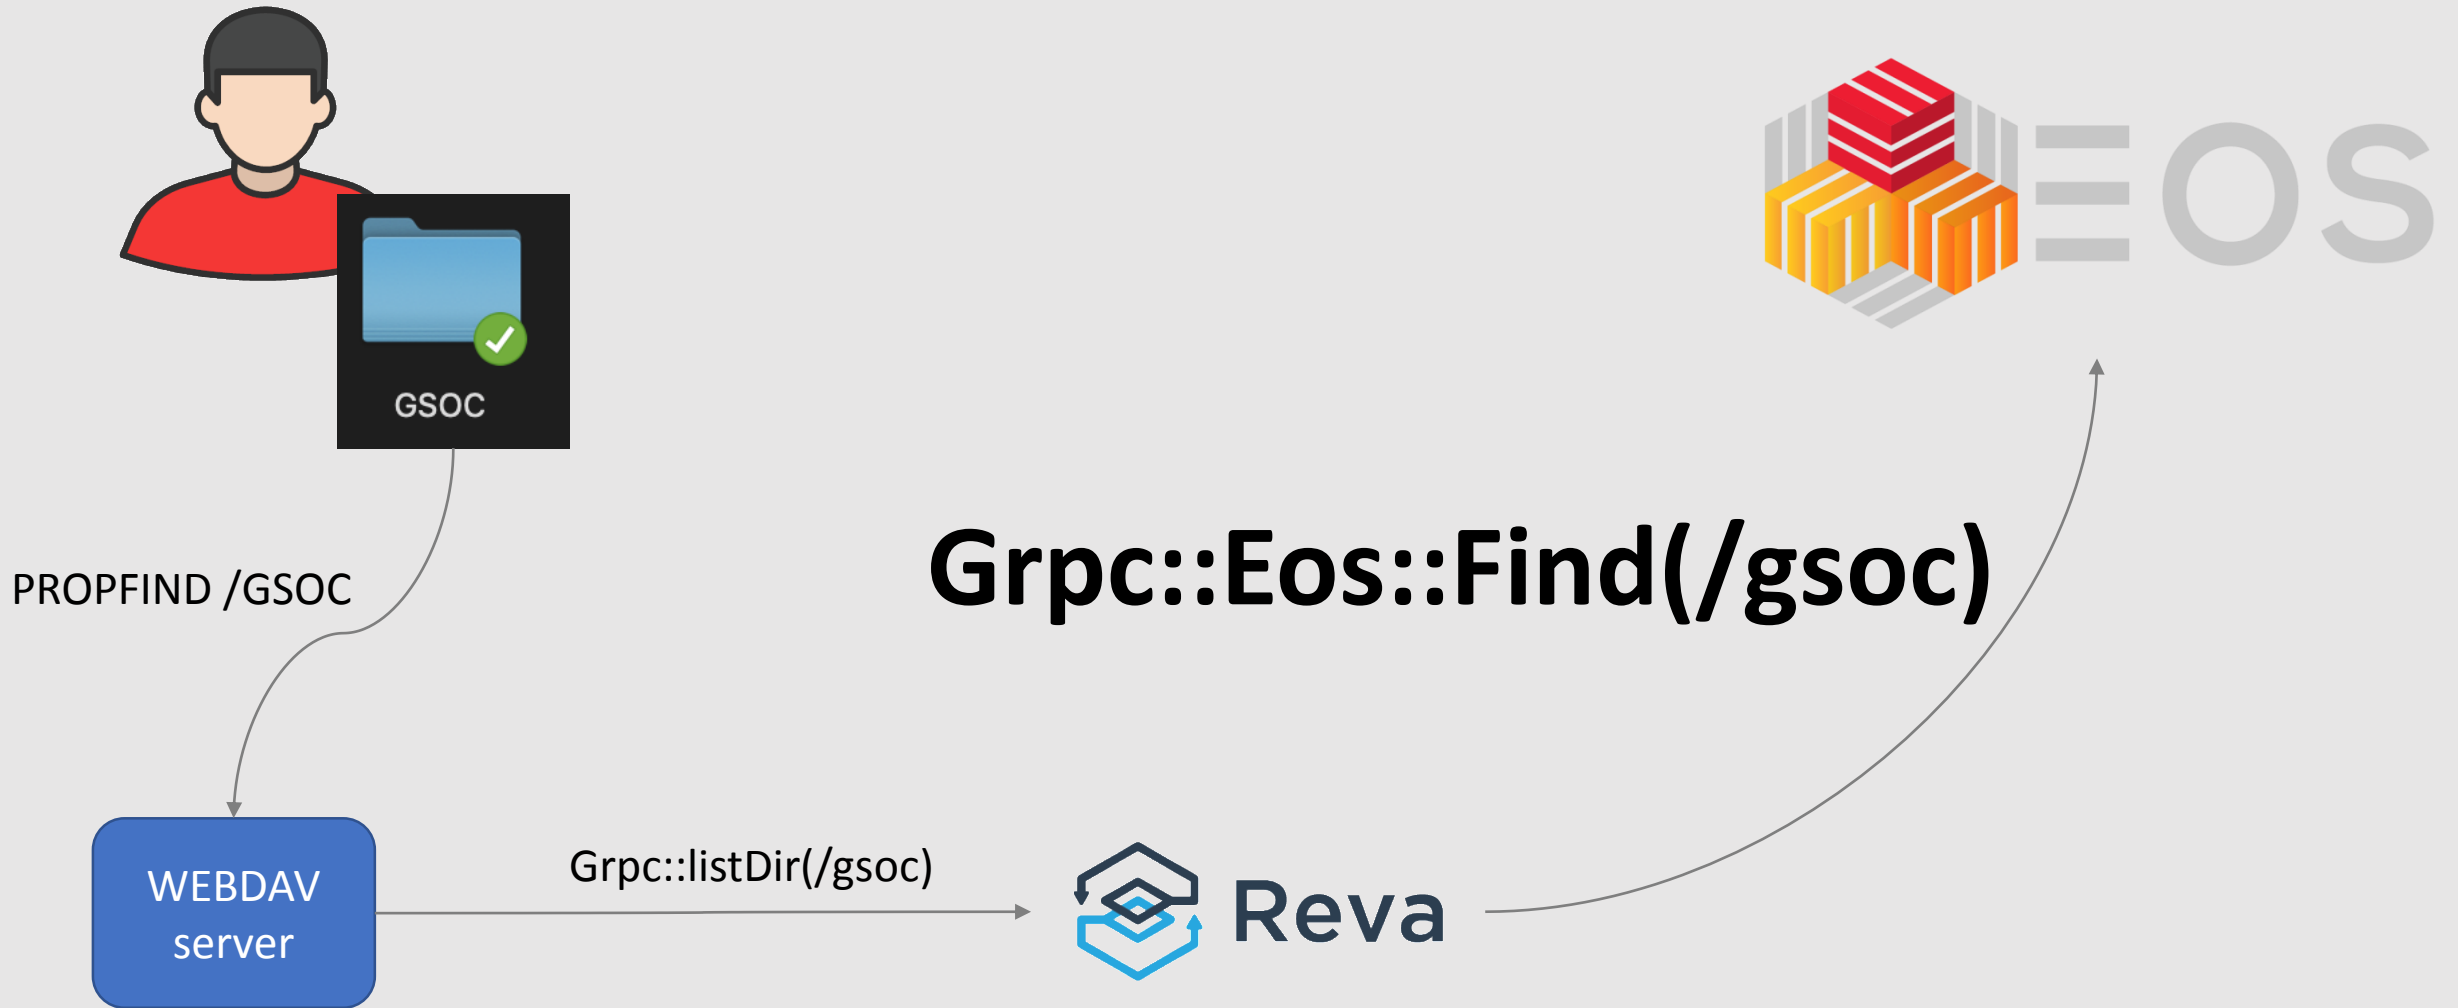


```
[root@eospps-fe1 (mgm:master mq:master) tmp]$ ./hey  
2020/02/03 07:33:32 Sending Ping message to eospps.cern.ch:50051  
2020/02/03 07:33:32 Server replied with message:"hola"
```


GRPC brings you many goodies that you don't care about
BUT your users will:

- Load balancing with default or custom algorithms
- Error retries before making them user visible ...
- Universally routable (HTTP/2), no ISP would dare to block ports 80, 443
- Proved to scale (Dropbox, Netflix, Google, ...)

GRPC as metadata control plane for EOS



Connect directly to EOS from a browser?



Would be nice to have ...

Other storages



```
ceph.dir.rbytes="88890171665717"  
ceph.dir.rctime="1430393746.09144842250"  
ceph.dir.rentries="10335945"  
ceph.dir.rfiles="9853255"  
ceph.dir.rsubdirs="482690"
```

CEPH FS

A POSIX-compliant distributed file system, with a Linux kernel client and support for FUSE

Only available on FUSE layer



Introduced in 2019

Watch folder API is an IBM Spectrum Scale feature that allows you to monitor folders, filesets, and inode spaces for file accesses.



Once upon a time was EOS



```
Directory: '/eos/user/g/gonzalhu/GSOC' Treesize: 119663
Container: 2 Files: 0 Flags: 42700
Modify: Mon Apr 8 04:09:19 2019 Timestamp: 1554689359.586192517
Change: Mon Jul 15 09:26:11 2019 Timestamp: 1563175571.48705357
Sync : Mon Apr 8 04:52:02 2019 Timestamp: 1554691922.294225849
Birth : Thu Jan 1 01:00:00 1970 Timestamp: 0.0
CUid: 95491 CGid: 2763 Fxid: 00af889a Fid: 11503770 Pid: 2463024 Pxid: 00259530
ETAG: af889a:1554691922.294
```

WFE Engine

The **workflow** engine is a versatile event triggered storage process chain. Currently all events are created by file operations. The policy to emit events is described as extended attributes of a parent directory. Each **workflow** is named. The default workflow is named 'default' and used if no workflow name is provided in an URL as ?

`eos.workflow=default.`

Recursive
Accounting



Arbitrary flexibility

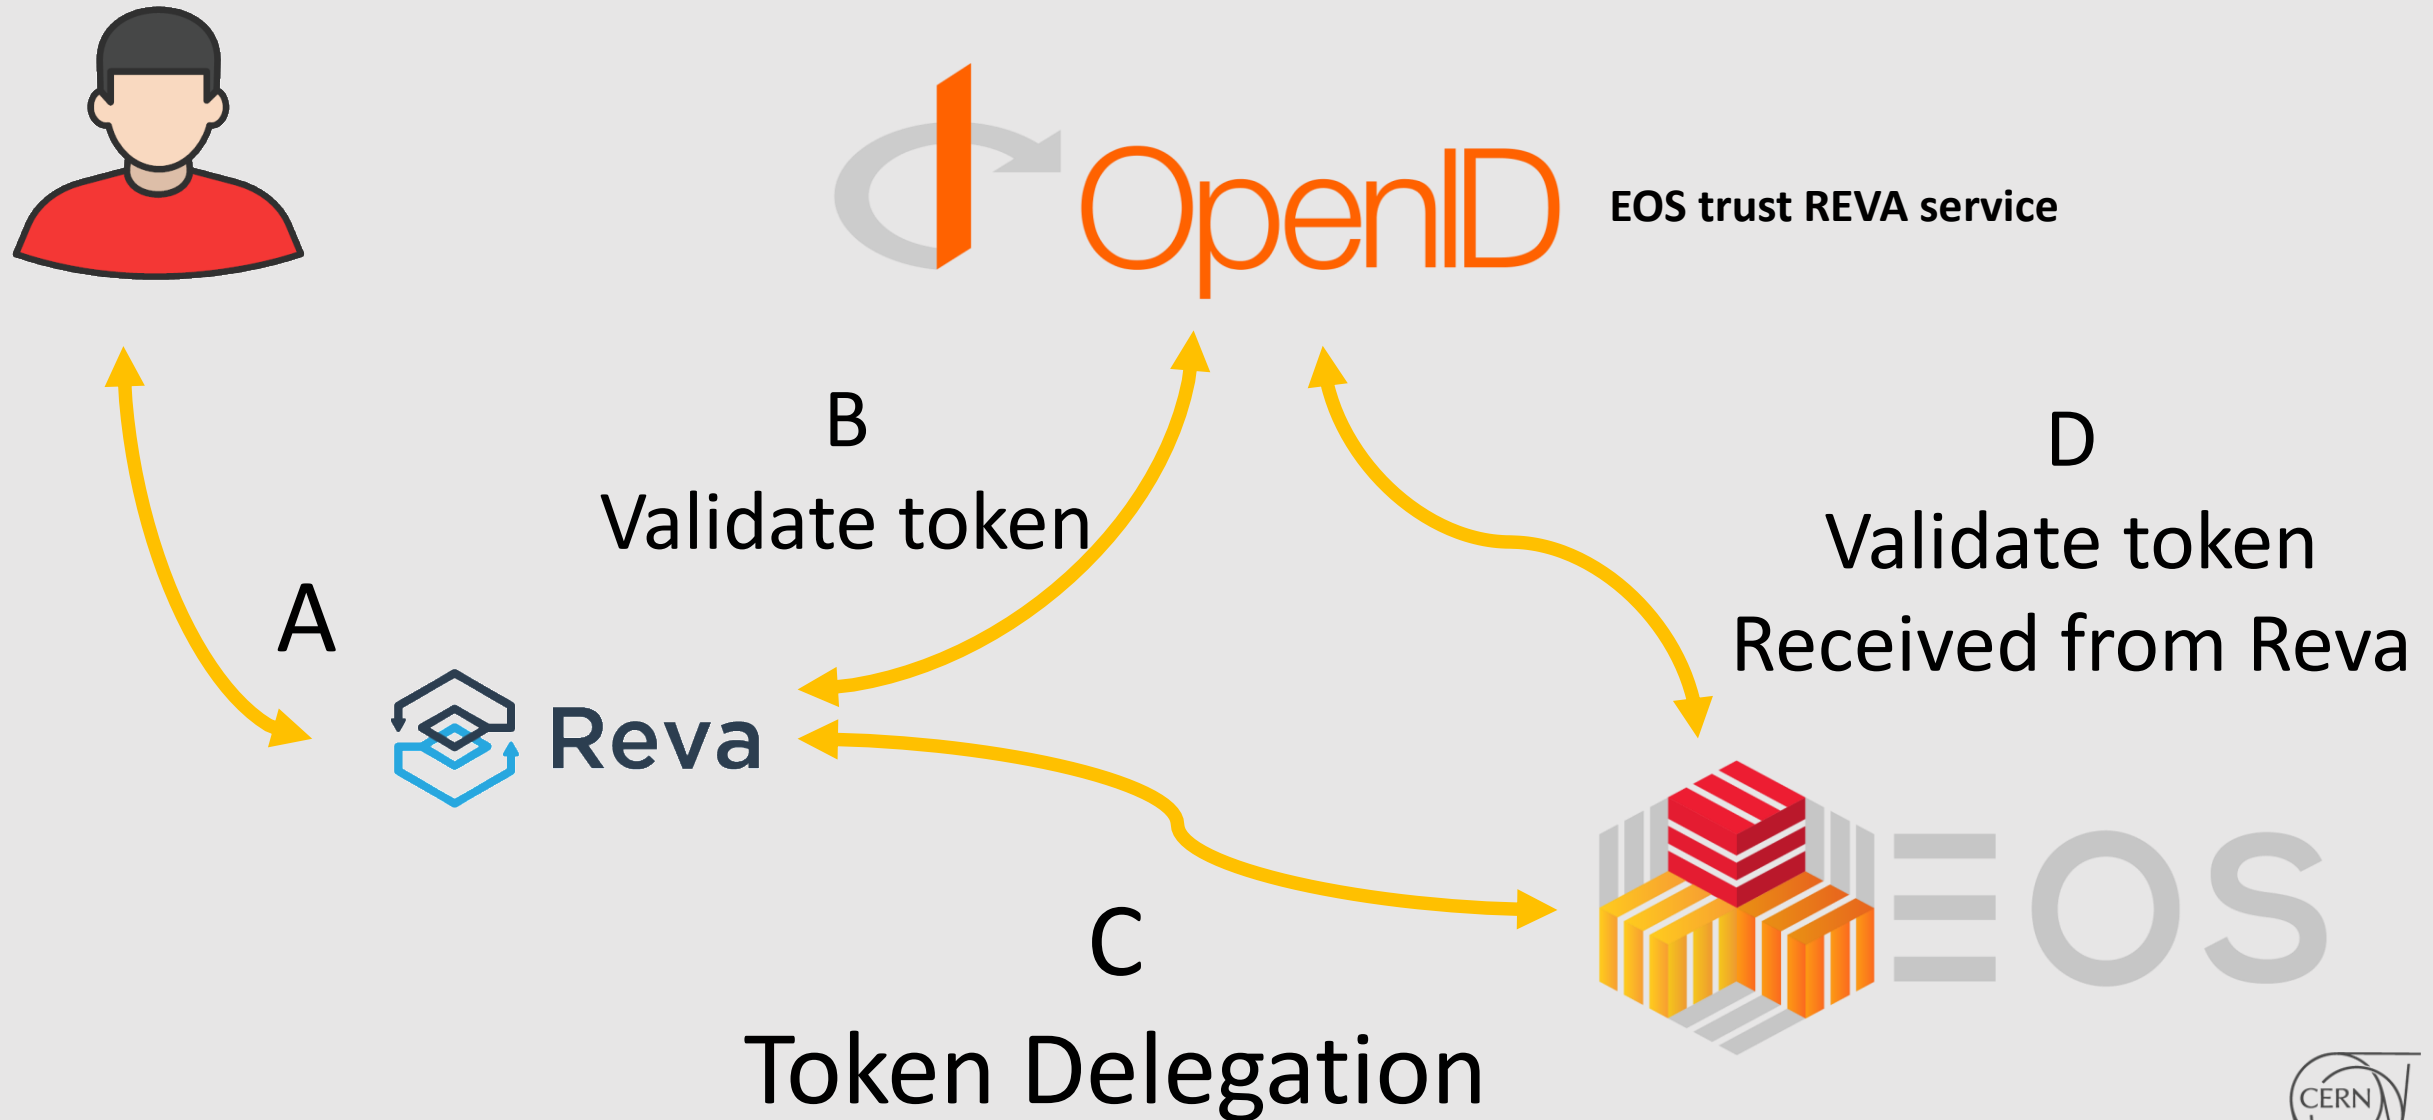
```
EOS Console [root://localhost] |/eos/user/g/gonzalhu/MyShares/> attr ls Australia-Wildlife.mov  
sys.eos.btime="1580716250.24115127"  
sys.reva.fav="1"  
sys.reva.ref='cs3://home-a:1234'"  
sys.reva.tags.Important="1"
```

MyShares +

<input type="checkbox"/>	Name	Size	Modified
<input checked="" type="checkbox"/>	Australia-Wildlife.mov Important!	0 KB	6 hours ago

1 file 0 B

Modern Auth: ~~Kerberos, X509~~, OpenIDConnect



The future European Federated Cloud

H2020 project: CS3MESH4EOSC

- EU-funded project (coordinated by CERN)
 - 6M EUR, 12 partners, 2020-2022
- Goal: Global collaborative environment for research
 - Share documents, files, projects, data, ...
 - Connected Application Hubs
 - Data Science Environments
- Federation of existing CS3 sites
 - 30+ sites (e.g. CERNBox, DesyBox, Universities, ...)
 - 300K+ users
 - cs3community.org
- Reva is the core inter-operability platform for the project
- The project will enable the federation of extremely large data sets and to connect EOS instances to have global delivery fabric for science (from small to big organizations)

