


# Running Parallel, Distributed ROOT Analysis with PyRDF on Public Cloud

## AWS Lambda Case Study



26 Jan 2021

Jacek Kuśnierz  
Piotr Pasternak  
Maciej Malawski

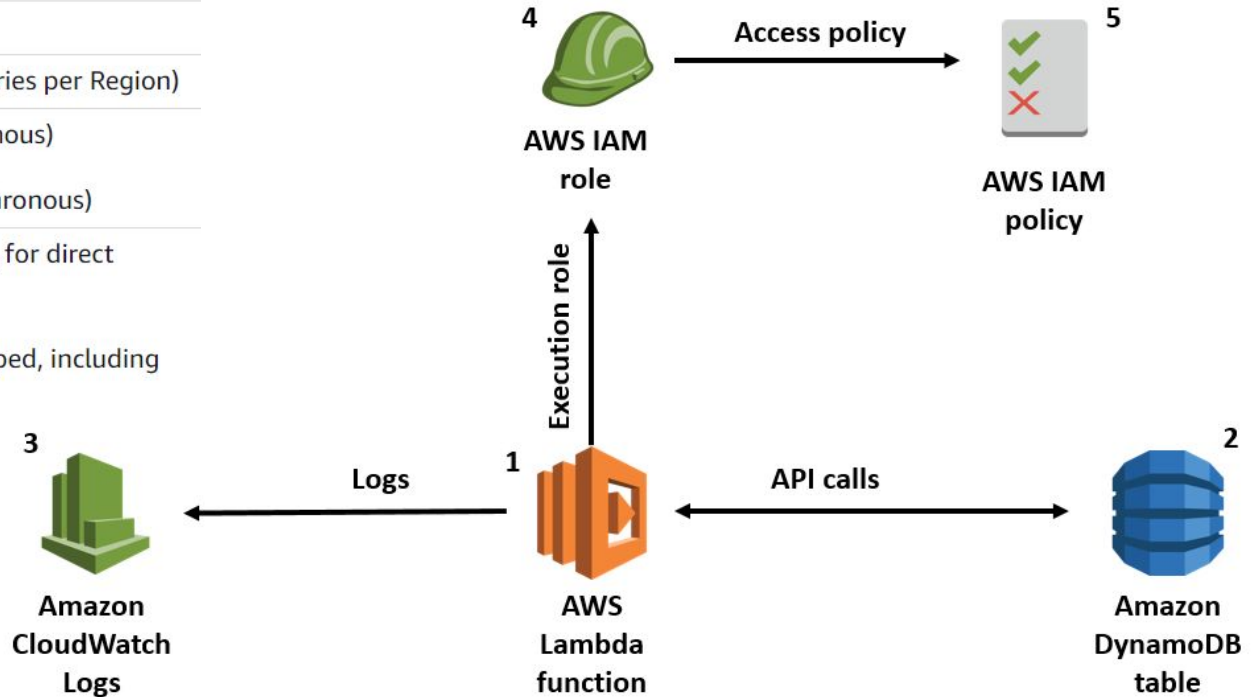

AGH University of Science and Technology  
Kraków, Poland 

# Presentation plan


1. Serverless, Lambda, ROOT and PyRDF
2. Project Structure Overview
3. Using the System
4. Technical Highlights

Function <b>memory allocation</b>	128 MB to 10,240 MB, in 1-MB increments.
Function <b>timeout</b>	900 seconds (15 minutes)
Function <b>environment variables</b>	4 KB
Function <b>resource-based policy</b>	20 KB
Function <b>layers</b>	5 layers
Function <b>burst concurrency</b>	500 - 3000 (varies per Region)
<b>Invocation payload</b> (request and response)	6 MB (synchronous) 256 KB (asynchronous)
<b>Deployment package</b> (.zip file archive) size	50 MB (zipped, for direct upload) 250 MB (unzipped, including layers)

# AWS Lambda



# Serverless



ROOT




ROOT

Data Analysis Framework

# PyRDF

```
PyRDF.use(backend)
d = RDataFrame(dataset)
f = d.Define(...)
    .Define(...)
    .Filter(...)


h1 = f.Histo1D(...)
h2 = f.Histo2D(...)
```



Local

APACHE  
Spark



...



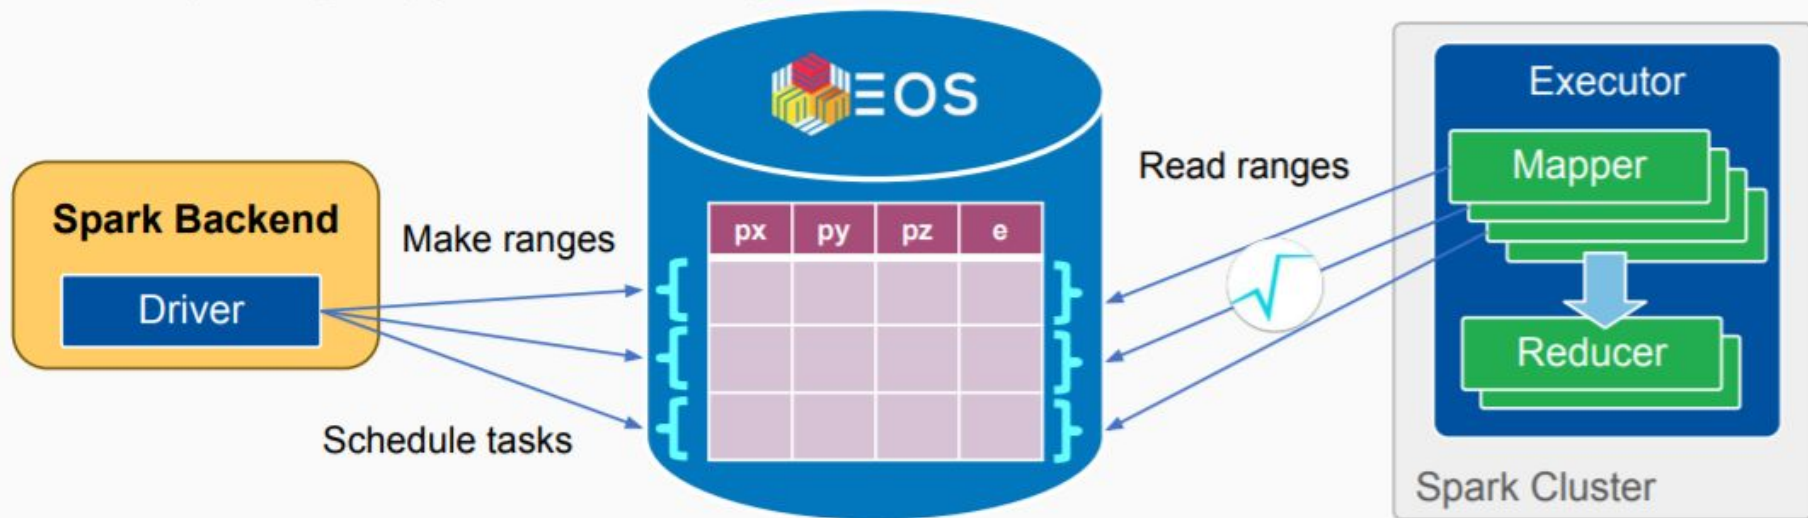
# PyRDF

```
PyRDF.use(backend)
d = RDataFrame(dataset)
f = d.Define(...)
    .Define(...)
    .Filter(...)

h1 = f.Histo1D(...)
h2 = f.Histo2D(...)
```

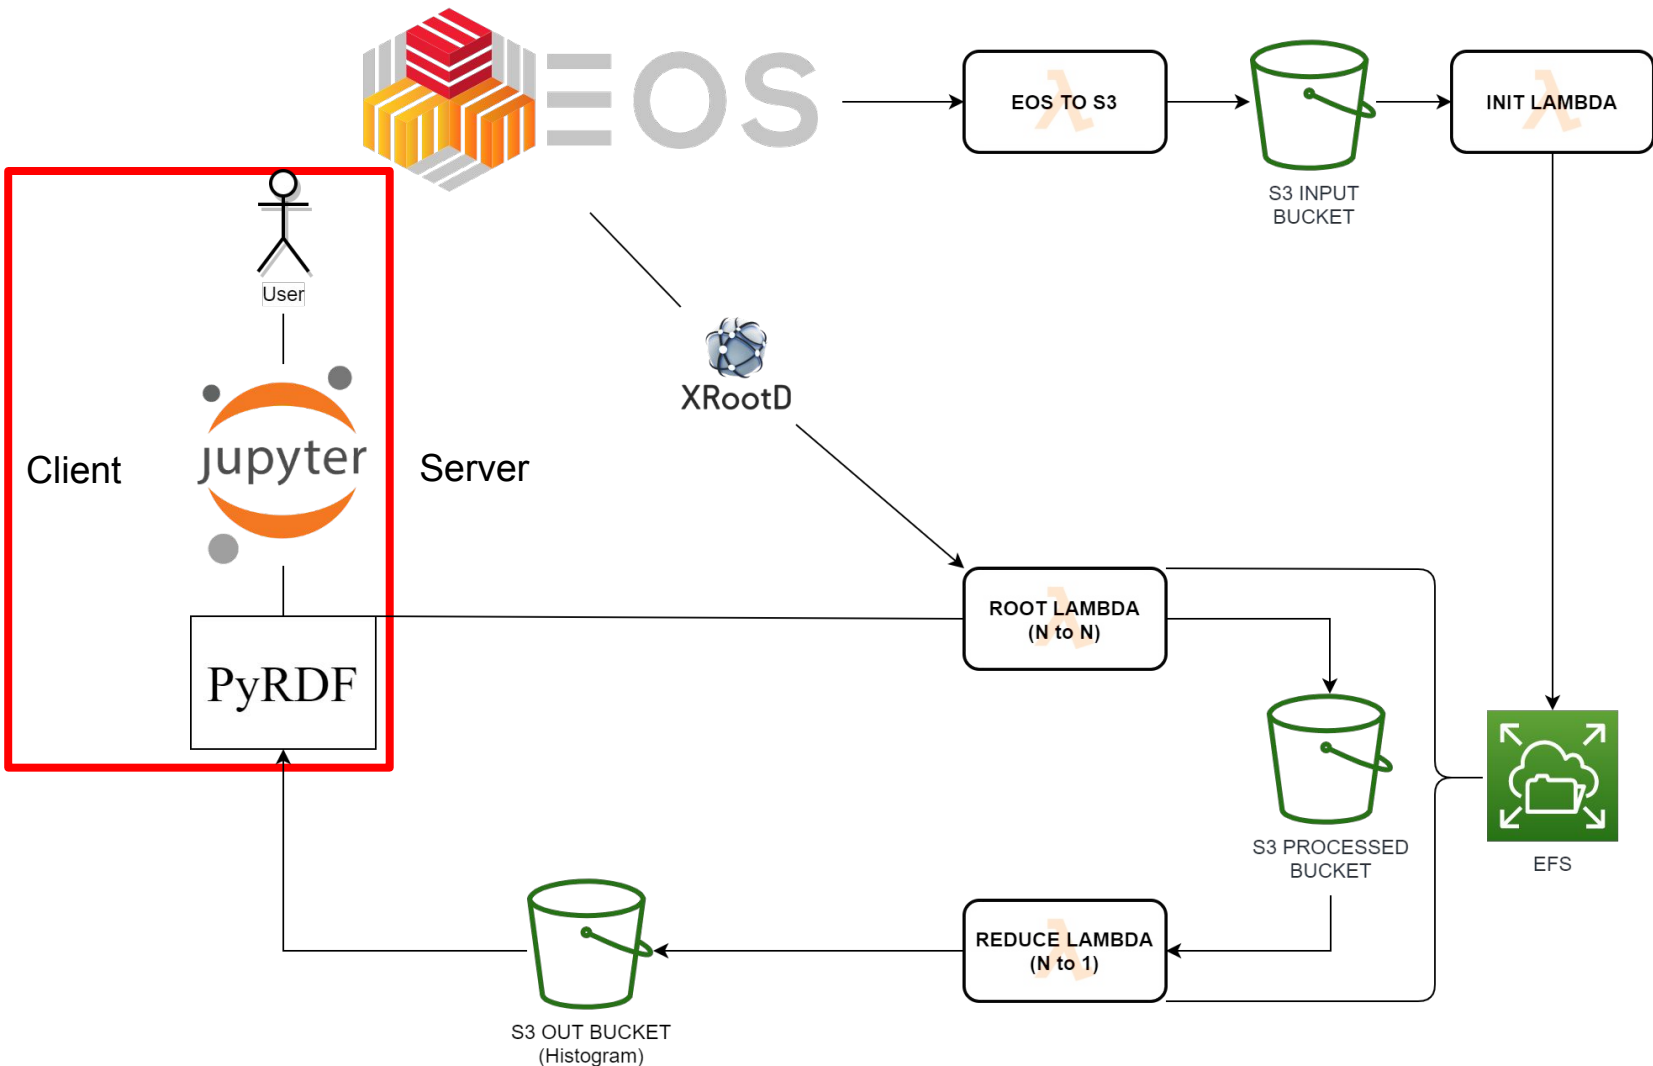


# PyRDF





# Structure Overview



# Using the System

# Using the system - create infrastructure

```
bash-4.2# cat ~/.aws/credentials
[default]
aws_access_key_id=ASIAUAHWLMWXWGET2WOG
aws_secret_access_key=TaN2q0ZoKrIJMxie9fBcd01yLheohb+q3Mhy9jf
aws_session_token=FwoGZXIvYXdzEE0aDN2AjoxtMCvWnIR6XyLLATNmfvVT/HvTydT/H4crFiFNZ4/7B7SGU91uksy7jwCSY5MBq1Ku9p3+tmH1VX+fv5
VuRwVdJl9tiT0kbwR3Db7c4AMFLimPm+D4ccg6Qd/mGsDzKRX8t/fyuhPEAcx6HueaxoRTs17N+WUczGA3UtjxeNvcaf6UsIWWZ1dBDSCqnqth45TbYNakJ
KKbaM1dYTBbR/IU7if5F402W4hUHsWw278JkZrMYEf41sAN1VfgZrRisoRWBxq1KvMEnIEd5y8cjsZ2GFyeigjKLmmvoAGMi2x5XJtGY8xrNsZ5o6BbwZwpJ
x7PkTpwdK1VPJky4xToNspZ4rSrFogS2KH9WE=
bash-4.2#
```

```
bash-4.2# terraform apply -var-file="logpass.tfvars" -auto-approve
null_resource.prepare_lambda_sources: Creating...
null_resource.prepare_lambda_sources: Provisioning with 'local-exec'...
null_resource.prepare_lambda_sources (local-exec): Executing: ["/bin/sh" "-c" "./build.sh"]
null_resource.prepare_lambda_sources (local-exec): Cloning into 'curl-to-s3'...
aws_eip.nat: Creating...
aws_vpc.main: Creating...
aws_efs_file_system.root_efs: Creating...
aws_s3_bucket.output_bucket: Creating...
aws_s3_bucket.input_bucket: Creating...
aws_s3_bucket.processing_bucket: Creating...
aws_eip.nat: Creation complete after 2s [id=eipalloc-0c58b9d9bb470172e]
```

# Using the system - analysis

```
- PyRDF.use("spark", {'npartitions': 32})
```

```
+ PyRDF.use("AWS", {'npartitions': 32})
```

localhost:8080/notebooks/df102\_NanoAODDimuonAnalysis.ipynb



df102\_NanoAODDimuonAnalysis Last Checkpoint: kilka sekund temu (autosaved)



Logout

Terminal

File Edit View Insert Cell Kernel Widgets Help

Trusted Python 3

Run Code

## NanoAOD files processed with Distristruted RDataFrame in Python

`df102_NanoAODDimuonAnalysis` ROOT tutorial running with PyRDF.

The NanoAOD-like input files are filled with 66 mio. events from CMS OpenData containing muon candidates part of 2012 dataset ([DOI: 10.7483/OPENDATA.CMS.YLIC.86ZZ](https://opendata.cern.ch/record/10.7483/OPENDATA.CMS.YLIC.86ZZ) and [DOI: 10.7483/OPENDATA.CMS.M5AD.Y3V3](https://opendata.cern.ch/record/10.7483/OPENDATA.CMS.M5AD.Y3V3)).

The macro matches muon pairs and produces an histogram of the dimuon mass spectrum showing resonances up to the Z mass. Note that the bump at 30 GeV is not a resonance but a trigger effect.

Some more details about the dataset:

- It contains about 66 millions events (muon and electron collections, plus some other information, e.g. about primary vertices)
- It spans two compressed ROOT files located on EOS for about a total size of 7.5 GB.

Date: April 2019

Author: Stefan Wunsch (KIT, CERN)

Adapted to PyRDF: Javier Cervantes Villanueva (CERN)

```
In [1]: import PyRDF
import ROOT

# Configure PyRDF to run on AWS splitting the dataset into 32 partitions
PyRDF.use("AWS", {'npartitions': '32'})

# Create dataframe from NanoAOD files
files = [
    "root://eospublic.cern.ch/eos/root-eos/cms_opendata_2012_nanoaod/Run2012B_DoubleMuParked.root",
    "root://eospublic.cern.ch/eos/root-eos/cms_opendata_2012_nanoaod/Run2012C_DoubleMuParked.root"
]

df = PyRDF.RDataFrame("Events", files);
```

Welcome to JupyROOT 6.22/07

# Using the system - analysis

```
In [6]: # Produce pdf of the plot
ROOT.gStyle.SetOptStat(0)
ROOT.gStyle.SetTextFont(42)

c = ROOT.TCanvas("c", "", 800, 700);
c.SetLogX(); c.SetLogY();
h.SetTitle("");
h.GetAxis().SetTitle("m_{#mu#mu} (GeV)"); h.GetAxis().SetTitleSize(0.04);
h.GetYaxis().SetTitle("N_{Events}"); h.GetYaxis().SetTitleSize(0.04);
h.Draw();

label = ROOT.TLatex()
label.SetNDC(True);

label.DrawLatex(0.175, 0.740, "#eta");
label.DrawLatex(0.205, 0.775, "#rho,#omega");
label.DrawLatex(0.270, 0.740, "#phi");
label.DrawLatex(0.400, 0.800, "J/#psi");
label.DrawLatex(0.415, 0.670, "#psi");
label.DrawLatex(0.485, 0.700, "Y(1,2,3S)");
label.DrawLatex(0.755, 0.680, "Z");
label.SetTextSize(0.040); label.DrawLatex(0.100, 0.920, "#bf{CMS Open Data}");
label.SetTextSize(0.030); label.DrawLatex(0.630, 0.920, "#sqrt{s} = 8 TeV, L_{int}");

c.SaveAs("dimuon_spectrum.pdf");|
Info in <TCanvas::Print>: pdf file dimuon_spectrum.pdf has been created
```

In [7]: %jsroot on

```
ROOT.gStyle.SetTextFont(42)


c2 = ROOT.TCanvas("c2", "", 800, 700);
c2.SetLogX(); c2.SetLogY();

h.SetTitle("");
h.GetAxis().SetTitle("m_{#mu#mu} (GeV)"); h.GetAxis().SetTitleSize(0.04);
h.GetYaxis().SetTitle("N_{Events}"); h.GetYaxis().SetTitleSize(0.04);
h.SetStats(False)
h.Draw();

label = ROOT.TLatex()
label.SetNDC(True);

label.DrawLatex(0.175, 0.740, "#eta");
label.DrawLatex(0.205, 0.775, "#rho,#omega");
label.DrawLatex(0.270, 0.740, "#phi");
label.DrawLatex(0.400, 0.800, "J/#psi");
label.DrawLatex(0.415, 0.670, "#psi");
label.DrawLatex(0.485, 0.700, "Y(1,2,3S)");
label.DrawLatex(0.755, 0.680, "Z");
label.SetTextSize(0.040); label.DrawLatex(0.100, 0.920, "#bf{CMS Open Data}");
label.SetTextSize(0.030); label.DrawLatex(0.630, 0.920, "#sqrt{s} = 8 TeV, L_{int} = 11.6 fb^{-1}");

label.DrawClone("Same")
c2.Draw()
```



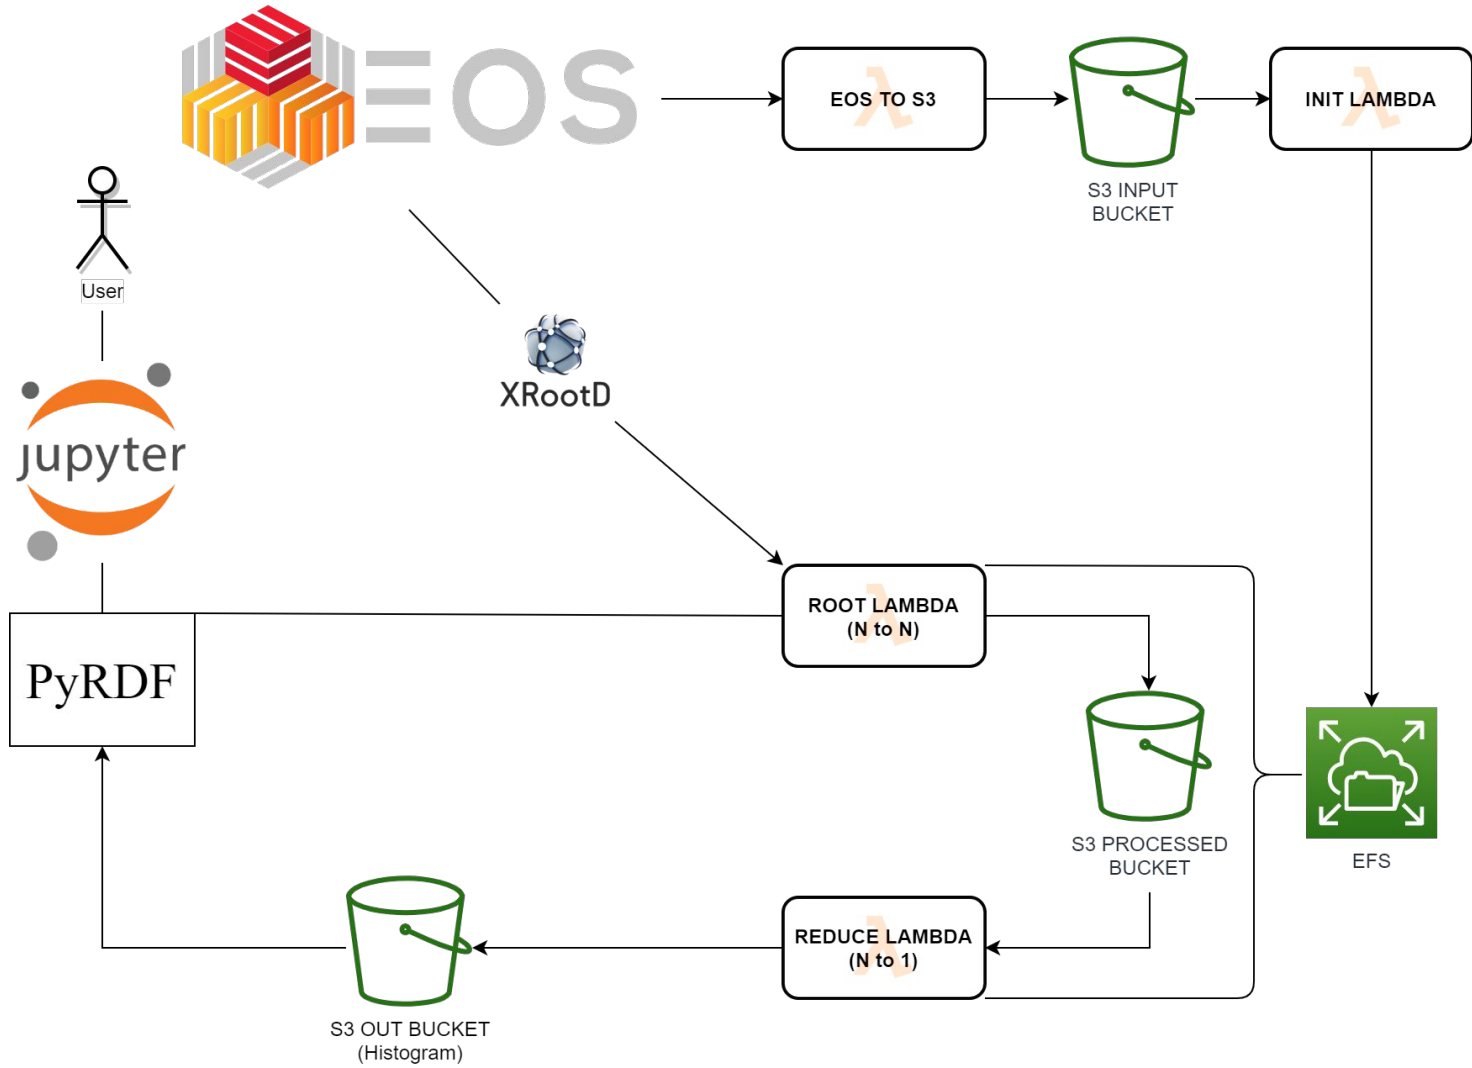
# Some Remarks

- Computation is quite expensive as it stands now
- There's not really much of exception handling at this point
- Remember to clean up after you're done:

```
bash-4.2# terraform destroy -auto-approve
```

# Technical Highlights





# CI/CD - Jenkins


```
bash-4.2# du -h --max-depth=1 /mnt/cern_root/
394M    /mnt/cern_root/chroot
368M    /mnt/cern_root/root_install
762M    /mnt/cern_root/
bash-4.2#
```

✓ [Main / RunWithSecond](#) < 308

Pipeline Changes Tests Artifacts Login

Branch: master 13m 29s Changes by noreply, jaku

Commit: 330f5da 19 days ago Branch indexing

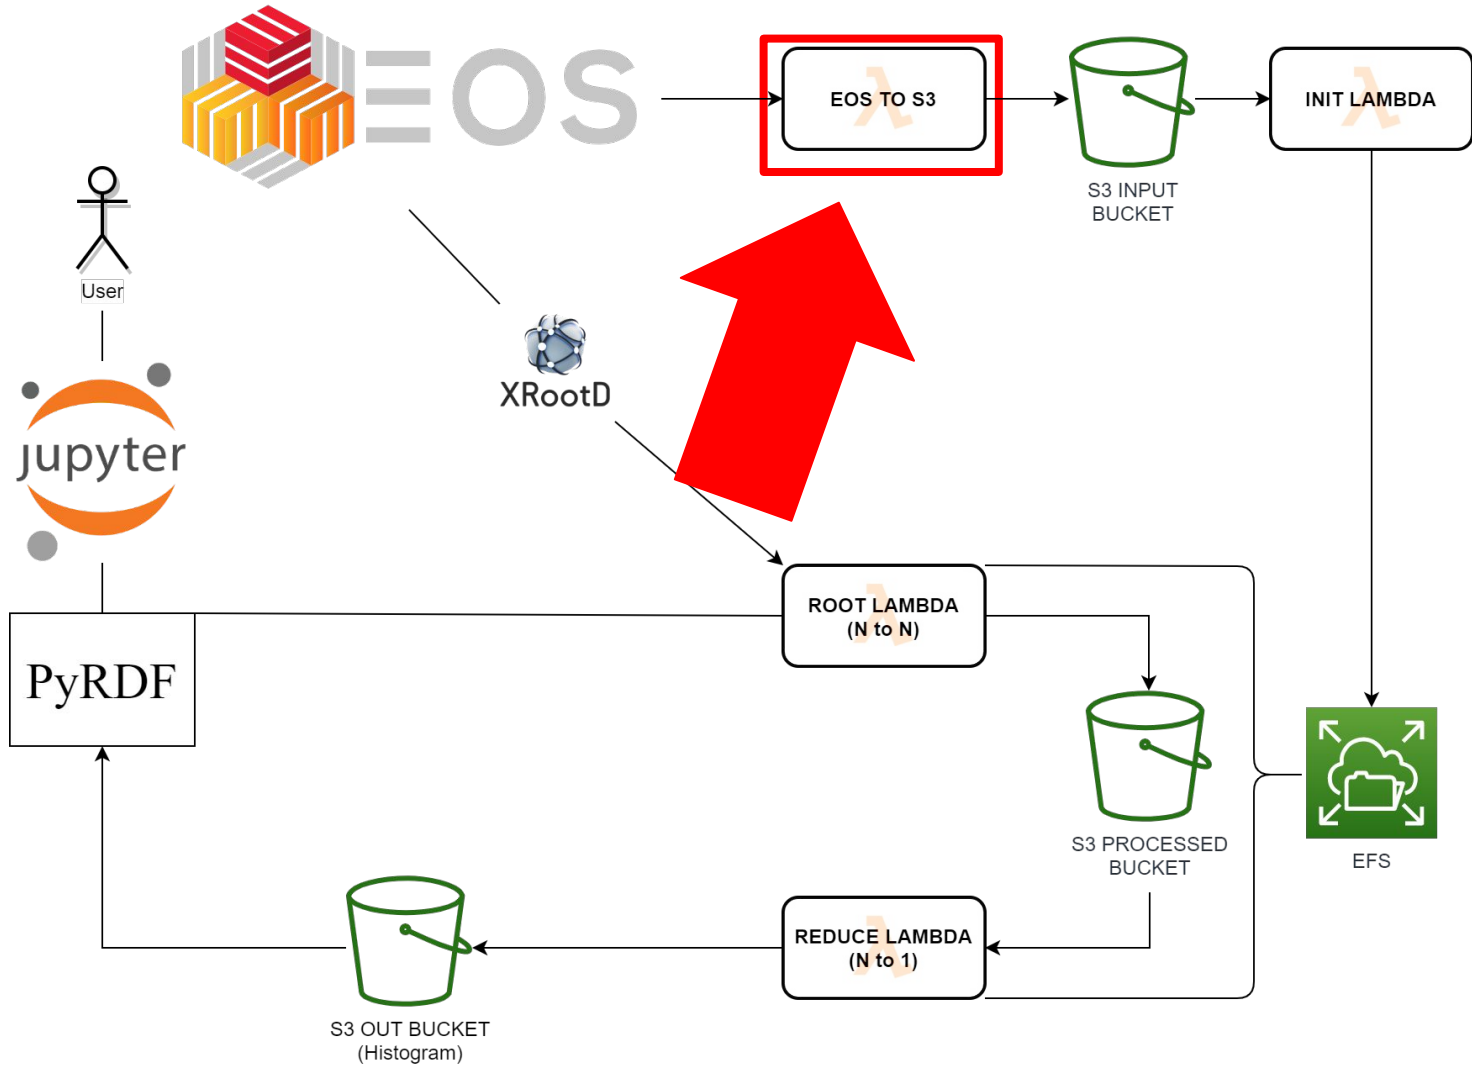


## Build images / command - <1s

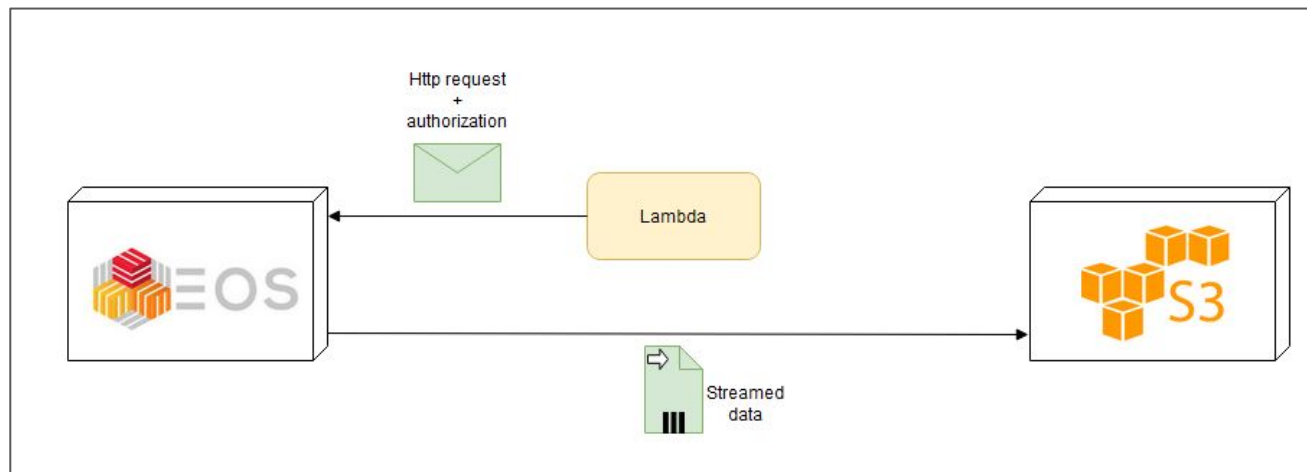
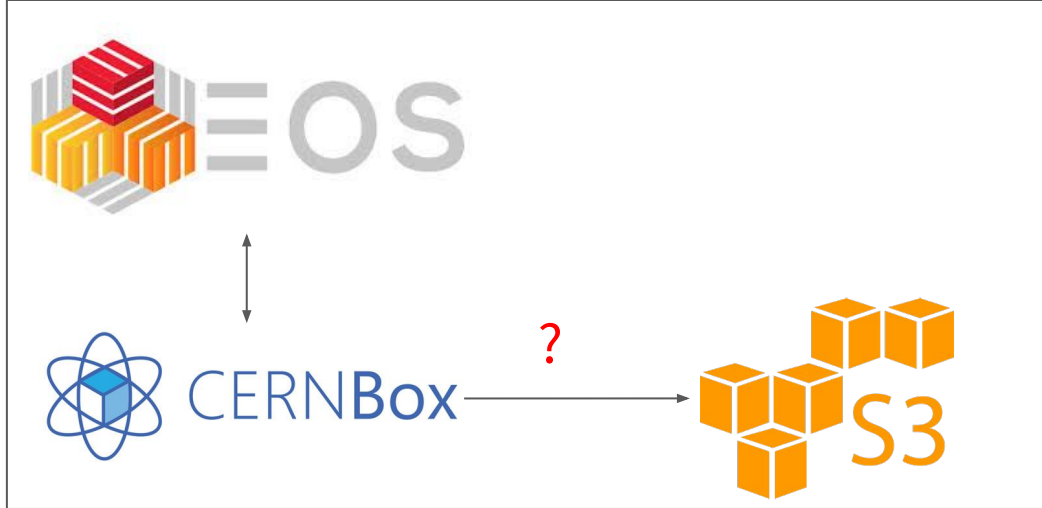
✓	> Checks if running on a Unix-like node	<1s
✓	> docker build -t pyrdf_terraform --network='host' commandRING -- Shell Script	7m 9s
✓	> Checks if running on a Unix-like node	<1s
✓	> docker inspect -f . pyrdf_terraform -- Shell Script	<1s
✗	> mkdir -p /mnt/dav/AWS_ROOT cp /*.pdf /mnt/dav/AWS_ROOT cd /mnt/cern_root zip -r /mnt/dav/AWS_ROOT/...	5m 59s

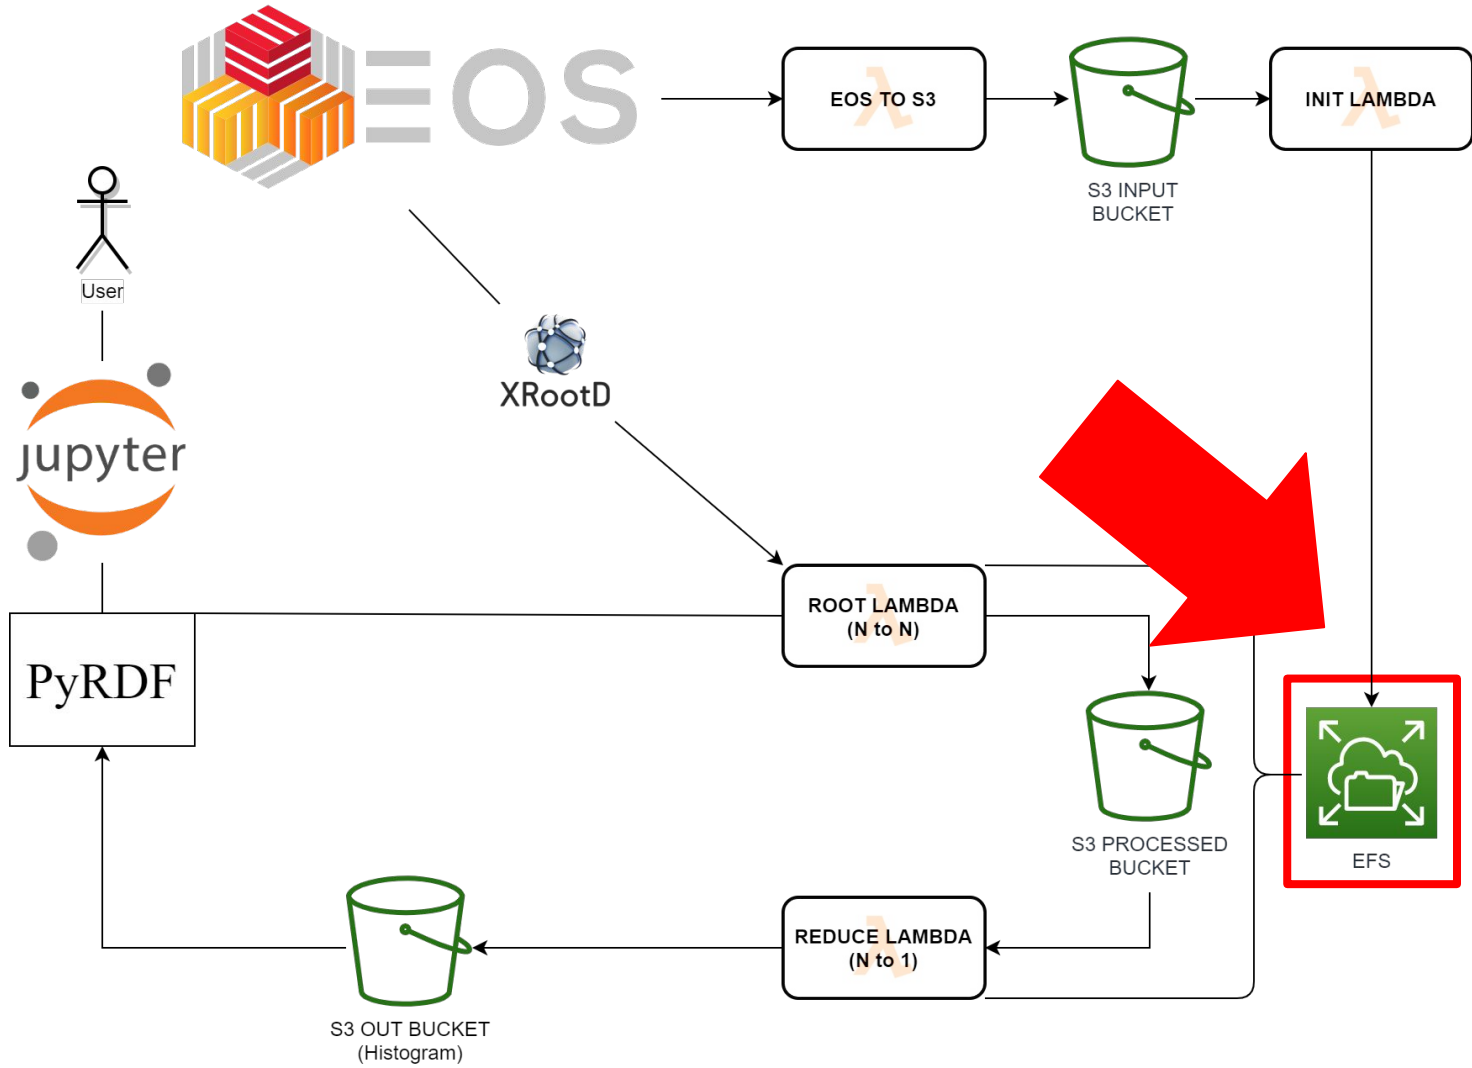



docker



# EOS Lambda







# ROOT Lambda



Lambda 1

Access Files

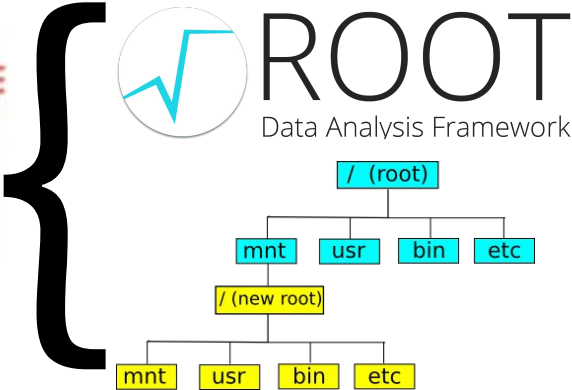


Lambda 2

Access Files



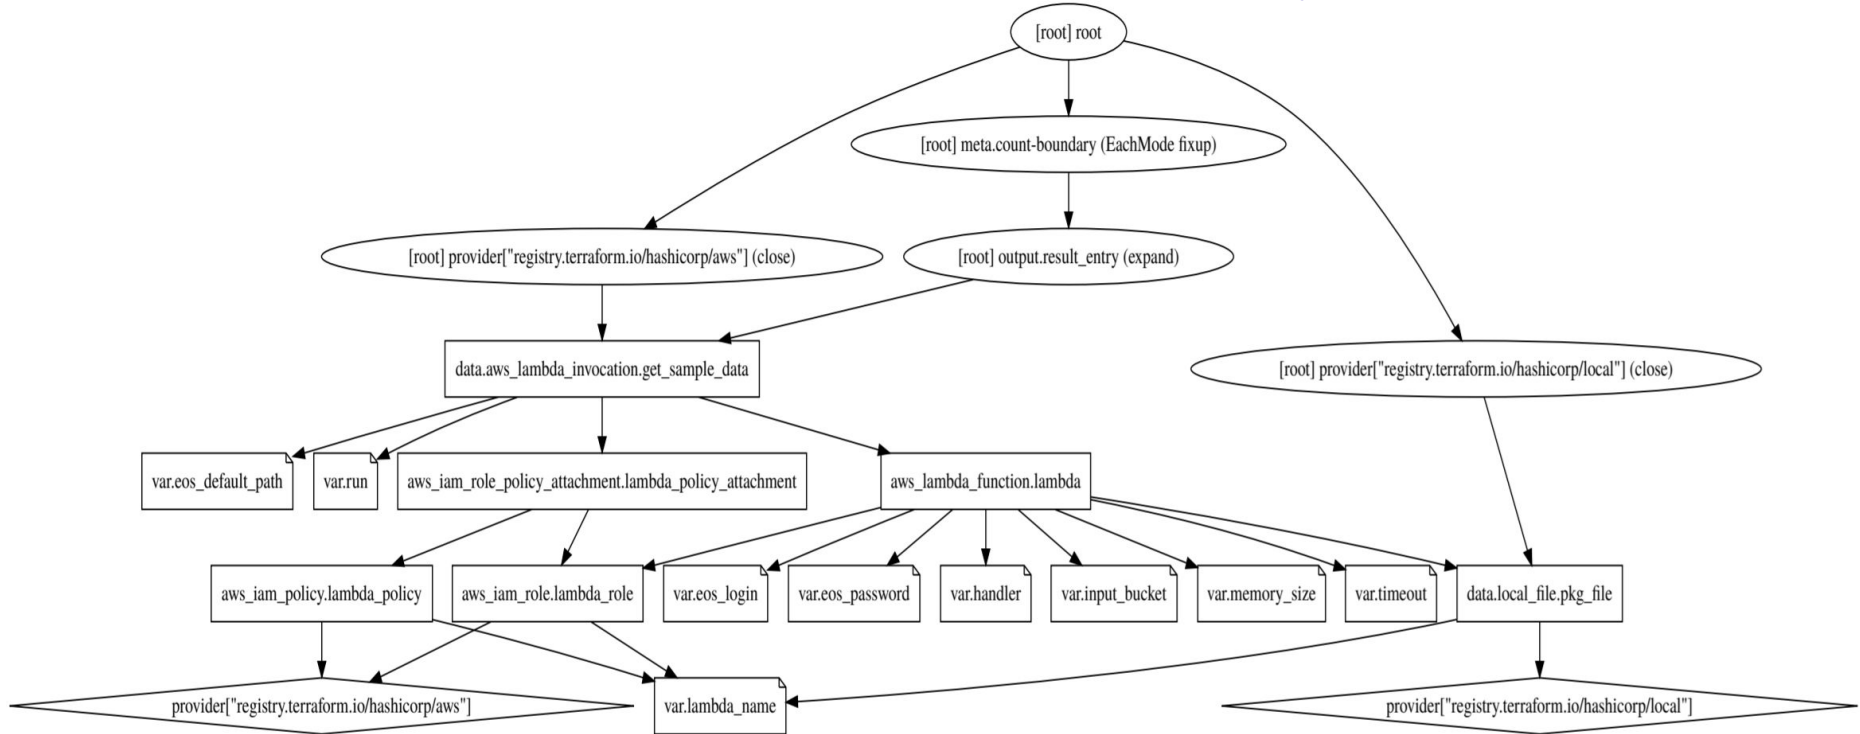
EFS



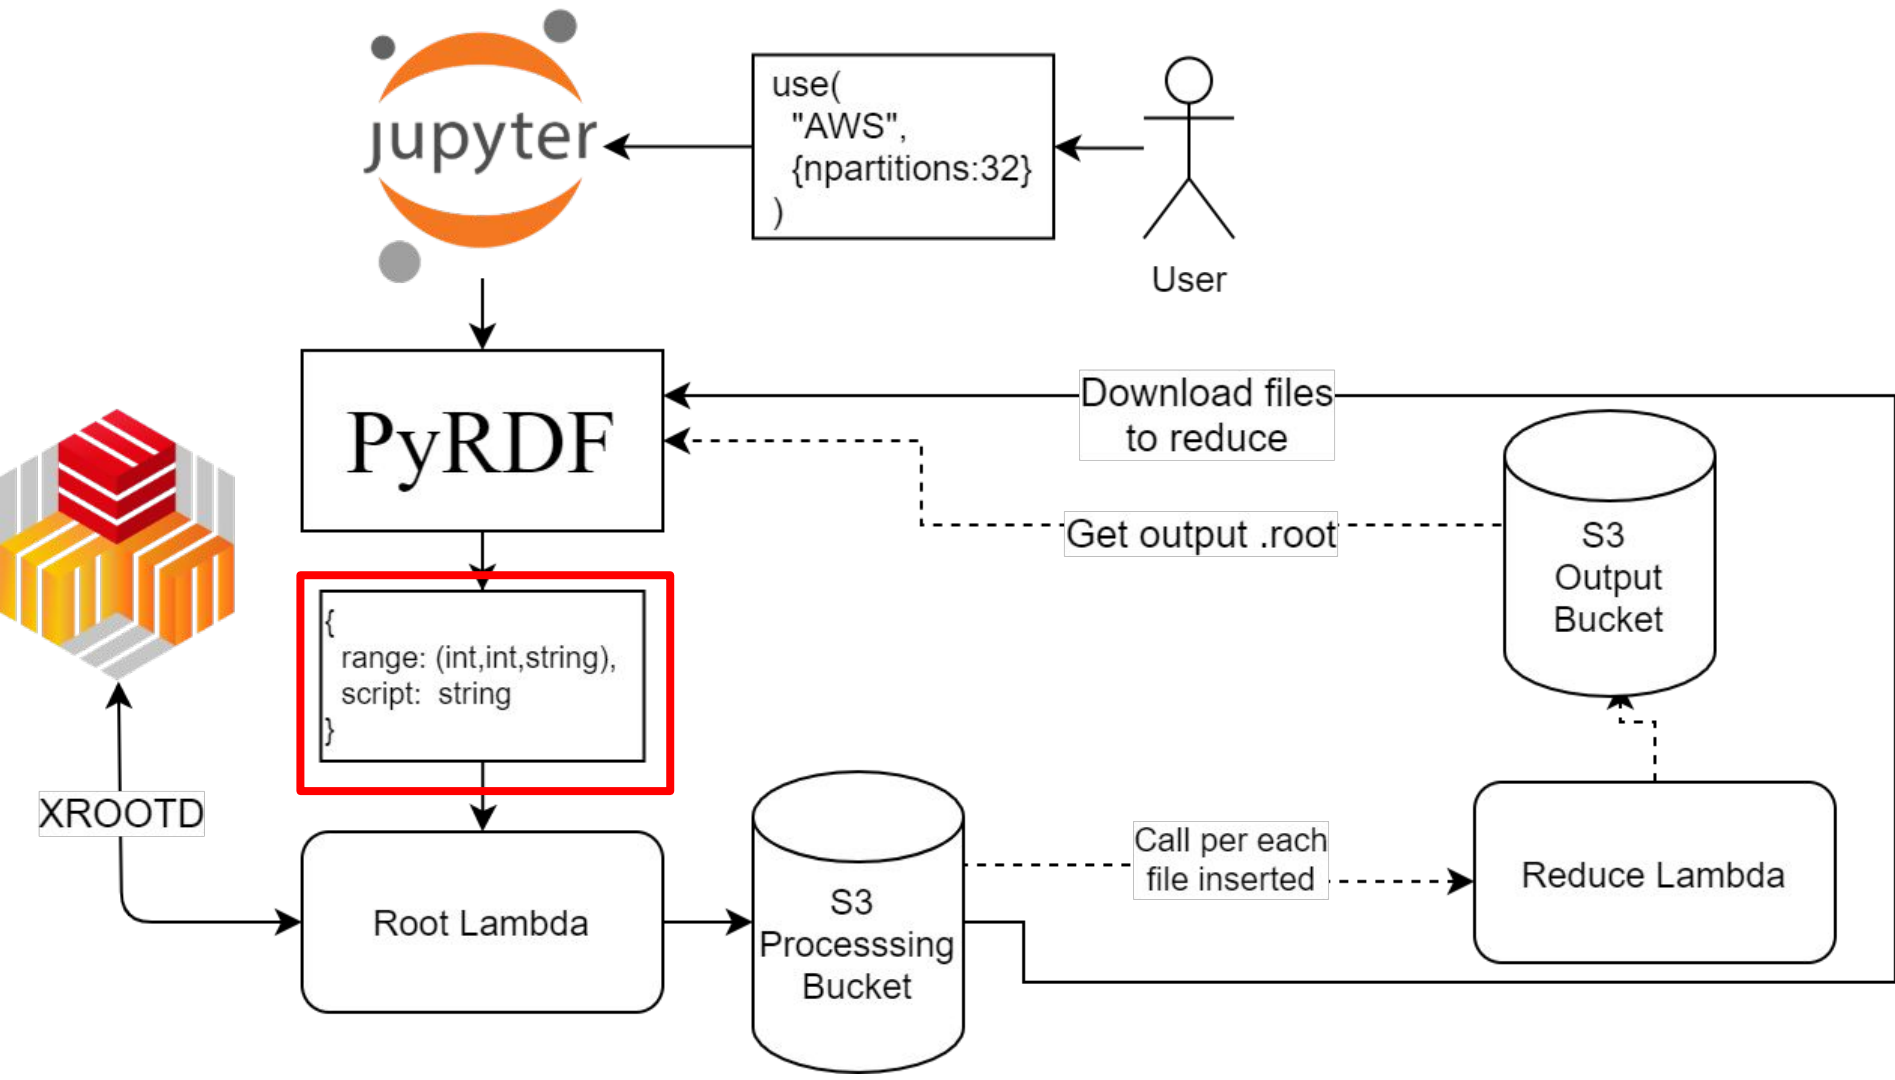
# Terraform



HashiCorp  
**Terraform**







# Results

- Proof of Concept works
- Looks promising despite 17s init time per lambda
- Shared file system on EFS - bottleneck
- High cost - new Docker Lambda implementation should solve it

Some complaints

# Summary and future work

- set up ROOT C++ JIT environment on AWS

# Summary and future work

- set up ROOT C++ JIT environment on AWS
- integrated solution for PyRDF

# Summary and future work

- set up ROOT C++ JIT environment on AWS
- integrated solution for PyRDF
- under certain conditions obtain not worse results than the current solution

# Summary and future work

- set up ROOT C++ JIT environment on AWS
- integrated solution for PyRDF
- under certain conditions obtain not worse results than the current solution
- run arbitrary software with own environment on Serverless Platform

# Summary and future work

- set up ROOT C++ JIT environment on AWS
- integrated solution for PyRDF
- under certain conditions obtain not worse results than the current solution
- run arbitrary software with own environment on Serverless Platform
- next step: move to Docker Lambda and test further



# Summary and future work

- set up ROOT C++ JIT environment on AWS
- integrated solution for PyRDF
- under certain conditions obtain not worse results than the current solution
- run arbitrary software with own environment on Serverless Platform
- next step: move to Docker Lambda and test further

## Thank you. Questions?

Special thanks to:

- Maciej Malawski, Associate Professor of AGH  
- Supervisor of the Project
- Vincenzo Eduardo Padulano - PyRDF expert

You can view code at <https://github.com/CloudPyRDF>

This work was supported by the Polish Ministry of Science and Higher Education, grant DIR/WK/2018/13

The project was realised as part of Bachelor Thesis at AGH University of Science and Technology in Kraków.

