

May 2020

Kacper Abramczyk, Mateusz Rejkowicz, Radek Rowicki, Marcin Stolarek

Data management in autonomous driving projects

Who we are?

Aptiv - **Aptiv PLC** (formerly known as **Delphi Automotive PLC**) is a Jersey-domiciled auto parts company headquartered in [Dublin, Ireland](https://en.wikipedia.org/wiki/Aptiv). (<https://en.wikipedia.org/wiki/Aptiv>)

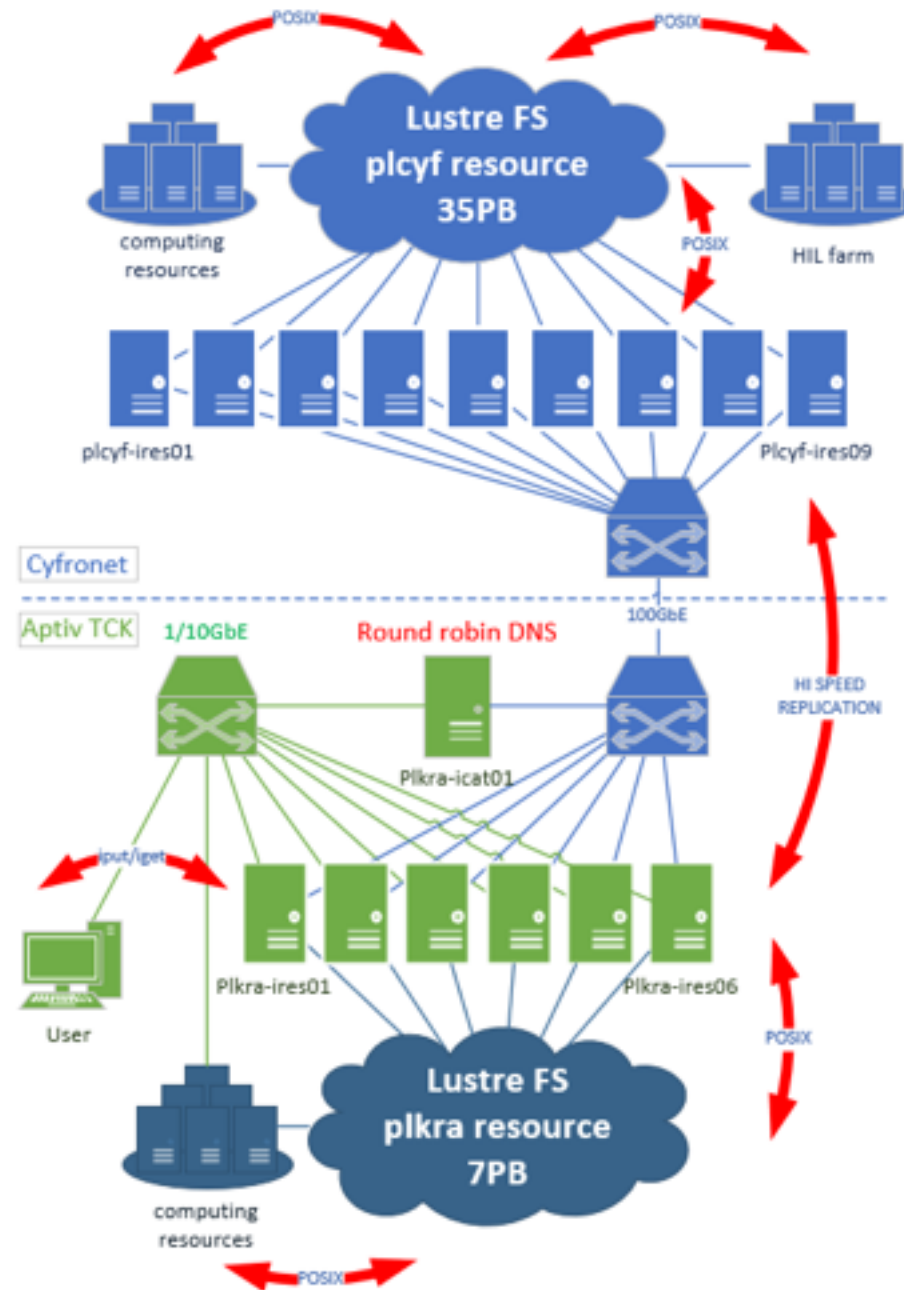
iRODS is part of Aptiv HPC infrastructure managed by the core team located in Kraków, Poland. The team acts globally with the help of local employees in US and “engineering groups” champions all around the world. Supporting “full stack” - from bare metal, through OS and services layers up to end-user support.



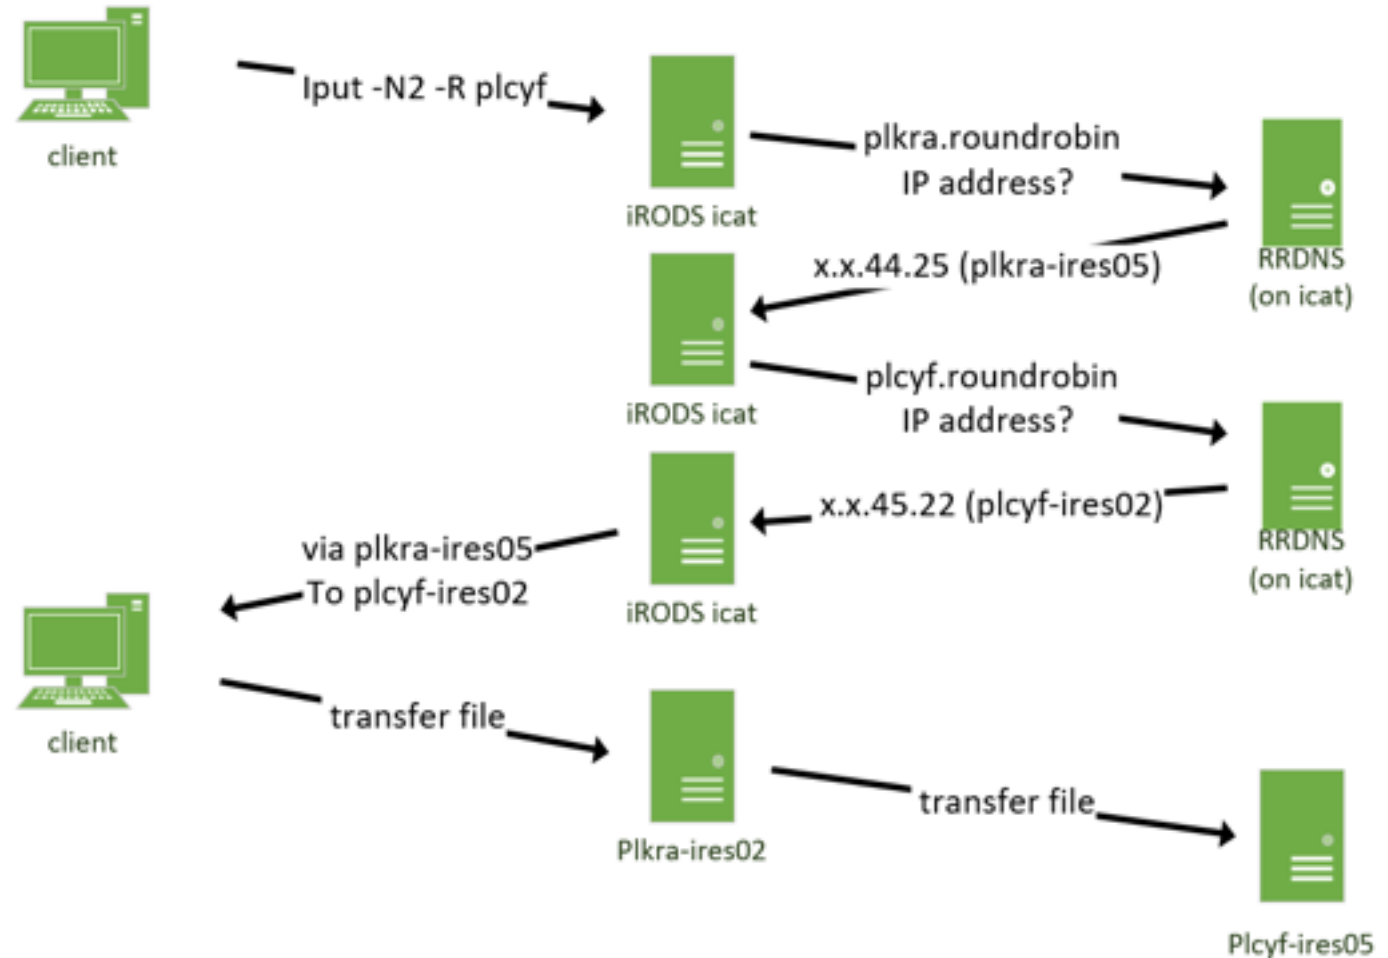

**Cracow,
Poland**

iRODS DNS round-robin for resource

- achieved ~90Gbps continuous upload
- active-active HA configuration
- iRods random/round-robin resource is not appropriate for a shared file system accesible from a group of hosts (like Lustre, BeeGFS) with single namespace
- need to implement your own DNS server not complaint with RFC3484
- gives an option for health check - implemented in DNS server.



iRODS DNS round-robin request explained



WIP: Preparing to enable auto registration on Lustre (and BeeGFS in our 2nd deployment in the future)

- **Currently running with a workaround script**
 - the user gives a file name he expects to exist, the script calls ireg over sudo.
- **Difficulties - need to compile.**
- **Terse instructions, lack kind of bulk registration.**
- **ZeroMQ - not very common.**
- **Risk - slow read -> full lustre log = FS outage**

Failed attempt to use audit rule engine

Full disk on rabbitMQ server effectively took iRODS down - probably because of very high number of connections at the time.

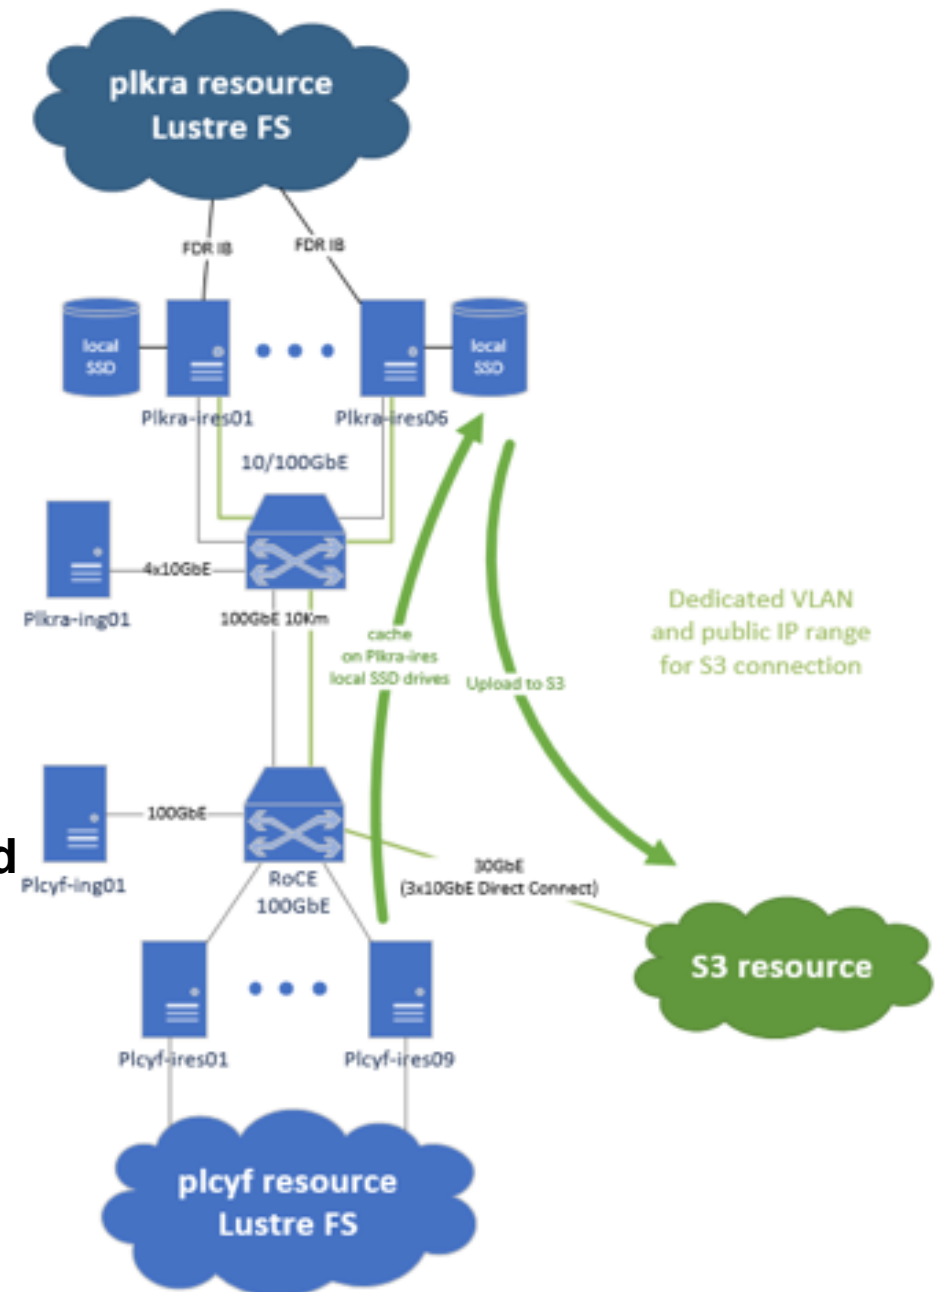
It will be great to have better set of examples and example filter on the rules that should go to audit rule engine with plots we can get out of those rules.

We'll probably re-enable it, however, production and efficiency is the major concern.

Issues with s3 integration

- Lack of really cacheless connector - efficiency of local drives/SAS controller is limiting our throughput with just four regular SSD drives
- Need to upload many files at once to get full throughput:
"Too many open files"
- S3 authorization issues we didn't fully understand. Sometimes it didn't work for a day after configuration and started to work without any changes.

`S3_DEFAULT_HOSTNAME=s3.<bucket-region>.amazonaws.com`



Smaller issues we've seen

- Reached limit for I1 desc
- end-users difficulties:
 - **imv - only within resource, it will be nice to behave like mv:**
 - rename within resource
 - irepl/itrim with different destination resource specified.
 - **Need to use -N2 to enforce redirect (don't transfer over icat)**
 - We don't know how to trigger it from rule engine (not as per file size definition)
 - It will be nice to have an option on user commands meaning "go directly".
- **How to limit reply size and send nice message to users generating slow queries?**
- **Probably python client with PAM authorization leading to (serious or not):**
remote addresses: 127.0.0.1 ERROR: readWorkerTask - readStartupPack failed. -4000

Summary

- **very flexible tool, not very easy to understand error messages.**

```
Jun  2 04:14:21 pid:38141 remote addresses: 10.214.44.23, 10.234.56.125 ERROR: _rcConnect:  
connectToRhost error, server on plcyf-ires.roundrobin:1247 is probably down status = -115000  
SYS_SOCKET_READ_TIMEOUT
```

- **educating users (need to prepare better tutorial, define workflow)**
- **not easy to introduce without full understanding of all capabilities and process definition**
- **difficulties of Windows centric organization**