EOS Windows client productisation

Comtrade 360's AI Lab

Gregor Molan
Head of Research & AI

Branko Blagojević
Lead engineer

Ivan Arizanović
Lead engineer

Comtrade
Three-Sixty

A global technology solutions provider dedicated to excellence, performance and high-quality standards.

CERN EOS workshop 2022

The Journey of EOS productisation at Comtrade 360

- Comtrade 360 at CERN openlab
- Provide EOS for Windows clients
- Presentation of EOS in a Windows way
- Upgrades of CERN EOS
- CERN EOS productisation
- Directions of current and future work

Comtrade 360 at CERN openlab

Journey of Comtrade 360 through the CERN

1. **CERN openlab – Comtrade – Comtrade 360**
2. **Needs and characteristics of CERN EOS**
3. **EOS implementation**
4. **EOS productisation by Comtrade 360**

1. CERN openlab – Comtrade – Comtrade 360

Comtrade: 30 years of enterprise software development
7 years at CERN

- 1992: Comtrade
  A valuable business partner in an enterprise-level backup product development
- 2015: Comtrade
  Signed CERN openlab Project Agreement for EOS Productisation
- 2016: Comtrade
  Full installation kit for CERN EOS
- 2018: Comtrade
  CERN EOS User Manual
- 2022: Comtrade 360
  Provide EOS for Windows clients (EOS-wnc)

2. Needs and characteristics of EOS

What was the initial aim of EOS?

- Collecting data generated by sensors
- Allowing physicists insight into experiments at CERN
- Ability of a huge data inflow
- Secure data storage
- Extremely low latency
- Provide “unlimited” storage size extensions

3. EOS implementation

EOS: Software developed at CERN

- EOS fulfils the extreme requirements of CERN experiments
- EOS fulfils high-reliability (RAID-DP, RAID 6, Archive)
- EOS storage capabilities (set-up at CERN in 2022)
  - Storage volume: 500 PB
  - Number of hard disks: 40.000
  - Number of files: 7.000.000.000
  - Number of clients: 30.000

4. EOS productisation by Comtrade 360

Productisation of EOS at CERN in 2022

- **EOS-wnc** (EOS Windows native client)
- EOS-wnc maintenance
- Appropriate documentation for EOS-wnc (open)
- EOS-wnc to meet industry requirements for the installation process
- EOS-wnc to meet industry requirements for support

Provide EOS for Windows clients

5. **Disk access on Windows**
6. **Communication to EOS cluster**
7. **Interface to Windows operating system**

5. Disk access on Windows (c_1)

The technology used for Windows disk access

- Using Microsoft Visual Studio and Windows native APIs
- The fastest Windows-native low-level libraries
- No additional Linux-like functions
- No additional interface for EOS-wnc console to Windows disks

6. Communication to EOS cluster (c_2)

EOS client communication with EOS cluster

- cURL (HTTPS) for file transfers
- gRPC (HTTPS) for command calls
- Active Directory (used for EOS-drive)

7. Interface to Windows operating system (c_3)

EOS Windows client should be Windows native

- The same user experience
  - EOS client on Linux and EOS client on Windows
    - EOS commands
    - EOS console
- Improvements to EOS console on Windows
  - History
  - Search
  - Help

Presentation of EOS in a Windows way

8. **EOS drive using Dokan**
9. **EOS drive and Dokan 2**

8. EOS drive using Dokan (c_4)

EOS cluster is presented as Windows drive letter

- EOS as Windows drive
- EOS cluster access with Dokan
  - Virtual file system
- Avoid Windows Kernel programming by using Dokan driver
  - Avoiding of possible high-cost issues (bugs)
- Dokan: High-performance thin driver

9. EOS drive and Dokan v2

Continuous improvement of EOS drive

- Implementation of Dokan v2
  - New version of Dokan – first release 2021-12-31
- Re-implementation of cURL
  - Some additional requirements related to Dokan v2
- Dokan v2 advantages
  - Sequential request: 10-35 % faster related to Dokan v1
  - File search: 100-200 % faster related to Dokan v1

Upgrades of CERN EOS

10. **Assurance of performance**
11. **Continuous testing in CERN on EOS-PPS**

10. Assurance of performance

High performance – important EOS feature

- Performance testing environment at Comtrade
  - Allows local continuous performance testing
- Performance aware development
  - Performance testing for each newly added feature
- EOS commands implementation for EOS-wnc
  - Performance comparison with EOS commands on Linux

11. Continuous testing in CERN on EOS-PPS

Continuous testing on the testing cluster at CERN (EOS-PPS)

- Joined development
  - CERN openlab : Comtrade 360
- Special thanks to Luca Mascetti and Manuel Reis
  - Continuous performance testing at CERN by Manuel Reis
- Immediate warning according to possible performance issue
  - Immediate actions: check new code to improve performance

CERN EOS productisation

12. **Documentation of EOS-wnc**
13. **Business model for EOS-wnc (supporting)**
14. **Business model for EOS-wnc (selling)**
15. **Comparison: EOS and other DFS**

12. Documentation of EOS-wnc

Professional software needs professional documentation

- Documentation for the EOS system
  - Installation of EOS
  - Documentation of EOS commands
  - Setting up EOS cluster
  - Need to update according to current EOS version
- Documentation for EOS-wnc
  - Windows help for EOS-wnc commands
  - Documentation as a manual – need to provide

13. Business model for EOS-wnc (supporting)

How to support EOS-wnc

- EOS cluster is open source
  - Available open-source support
- EOS-wnc is planned as proprietary software
  - Freely available for CERN
  - Need to provide support as bug fixing
  - Need to provide support as upgrades to new Windows OS

14. Business model for EOS-wnc (selling)

How to sell EOS-wnc

- Selling open-source (EOS cluster)
  - Installation and set-up
  - Open-source support
- Selling proprietary software (EOS-wnc)
  - Selling new versions of EOS-wnc
- Convince customers
  - EOS is the most adequate solution

15. Comparison between EOS and other Distributed File Systems

Most DFS (Distributed File Systems) are the “best”

- Current action at Comtrade 360
- Comparison between following DFS:
  - CephFS
  - Hadoop
  - Lustre
  - CERN EOS
- Setting up DFS performance testing environment
- Planned
  - Consult with CERN about EOS performance testing

Directions of current and future work

Productisation of EOS & Productisation of EOS-wnc

To finalise:

- Missing commands
- Upgrade to Dokan v2
- Comparison testing
- Documentation
- Business model

Additional actions:

- Professional testing
- Marketing plan
- Support plan

Discussion: How to improve EOS-wnc productisation?

Thank you

Gratitude

Thanks to CERN:

1) Luca Mascetti
2) Manuel Reis

Thanks to Comtrade:

1) Veselin Jevrosimović
2) Alexis Lope-Bello

Gregor Molan | @Comtrade360 | www.comtrade360.com | phone: +386 81 605 200