



Contribution ID: 43

Type: 10 Minutes

EOS-drive for Windows: Architecture and file transferring system

Thursday 27 April 2023 11:00 (10 minutes)

EOS-drive is part of the EOS Windows Native Client package, it mounts the EOS filesystem as a Windows disk drive by which Windows applications interact with the EOS filesystem.

EOS-drive communicates with Windows applications through the user-mode Dokan library and kernel-mode Dokan driver. File operation requests from applications (e.g., CreateFile, ReadFile, WriteFile...) are sent to the Dokan driver, which is then forwarded to the Dokan library and subsequently to the EOS-drive. The results of this routine are sent back to the Windows application as a response to the operation request. The Dokan file system driver and library therefore acts as a proxy between Windows applications and the EOS-drive.

To collect data that should be sent to the Windows application as a response to an operation request, EOS-drive communicates with the EOS cluster through HTTPS protocol. EOS-drive uses the cURL library for file transferring and gRPC for metadata requests.

WriteFile and ReadFile operations are used for file transferring. Windows applications requests a process over chunks (data fragments) by providing offset and chunk size to the EOS-drive. Then, EOS-drive is made to serve Windows applications by performing WriteFile or ReadFile operations over specified chunks.

There are three mechanisms of file transferring in EOS-drive, these are: transfer chunk by chunk, transfer with the EOS-drive buffer, and single-session transfer.

Transferring chunk by chunk means that each requested chunk by the Windows applications will be uploaded/downloaded in a separate session. This method is the slowest but can be realized in any case.

Transferring with an EOS-drive buffer can be done in three different scenarios. In case of upload, EOS-drive buffers some of the chunks received from the Windows driver, sorts them if needed, and uploads them together per second. With regards to downloading "small" files, instead of downloading the requested chunk, EOS-drive will download the whole file and then use downloaded bytes for chunk requests. There is also a buffer between the upload and download process, i.e., the first part of an uploaded file will be buffered and then used for subsequent download requests.

Single-session transfer means that one session will be used for the whole file, not only for one chunk. Afterwards, if chunk requests from the driver match the incoming chunk from the EOS cluster, the chunk from the EOS cluster will then only be forwarded to the driver. Or in the opposite direction for the upload process. This method is the fastest, however, it has several restrictions.

Functionality and transfer speed are tested locally at Comtrade, as well as at CERN, therefore the best solution for file transfer in EOS-drive is single-session transfer, although this method is impossible in all situations due to many restrictions. The next best solution as per speed, is to implement transfer using the EOS-drive buffer, it is however limited per number of files to save memory resources. Subsequently, if none of the previous mechanisms can be realized, transfer chunk by chunk will be used.

Author: ARIZANOVIC, Ivan (Comtrade 360's AI Lab)

Co-authors: BLAGOJEVIC, Branko (Comtrade); MOLAN, Gregor (Comtrade 360's AI Lab)

Presenter: ARIZANOVIC, Ivan (Comtrade 360's AI Lab)

Session Classification: EOS Ecosystem

Track Classification: EOS, Data Management and Applications