Comtrade Three-Sixty

Porting of XRootD to Windows as a part of EOS-wnc

Gregor Molan | **Branko Blagojević** | **Ivan Arizanović**

Comtrade Group / Comtrade 360

Ljubljana, 2023-03-23

Content

1. Introduction
2. Implementation of XRootD in EOS
3. XRootD and EOS client
4. Windows port of XRootD client
5. Resume

Introduction

1. Introduction
    1. Comtrade at CERN / EOS / XRootD
    2. EOS at CERN in numbers
    3. XRootD and EOS
    4. EOS client on Windows
2. Implementation of XRootD in EOS
3. XRootD and EOS client
4. Windows port of XRootD client
5. Resume

Comtrade at CERN / EOS / XRootD

Comtrade Three-Sixty

EOS at CERN in numbers

- Total space
  780 PB
- Total files
  ~8 Bil
- Total storage nodes
  ~1.300
- Total disks
  ~60.000

Timeline

- 2010 : 2 PB
- 2012 : 12 PB
- 2014 : 40 PB
- 2016 : 150 PB
- 2018 : 270 PB
- 2020 : 400 PB
- 2022 : 680 PB
  - 2023 : 780 PB
- 2024 : ???

EOS and XRootD

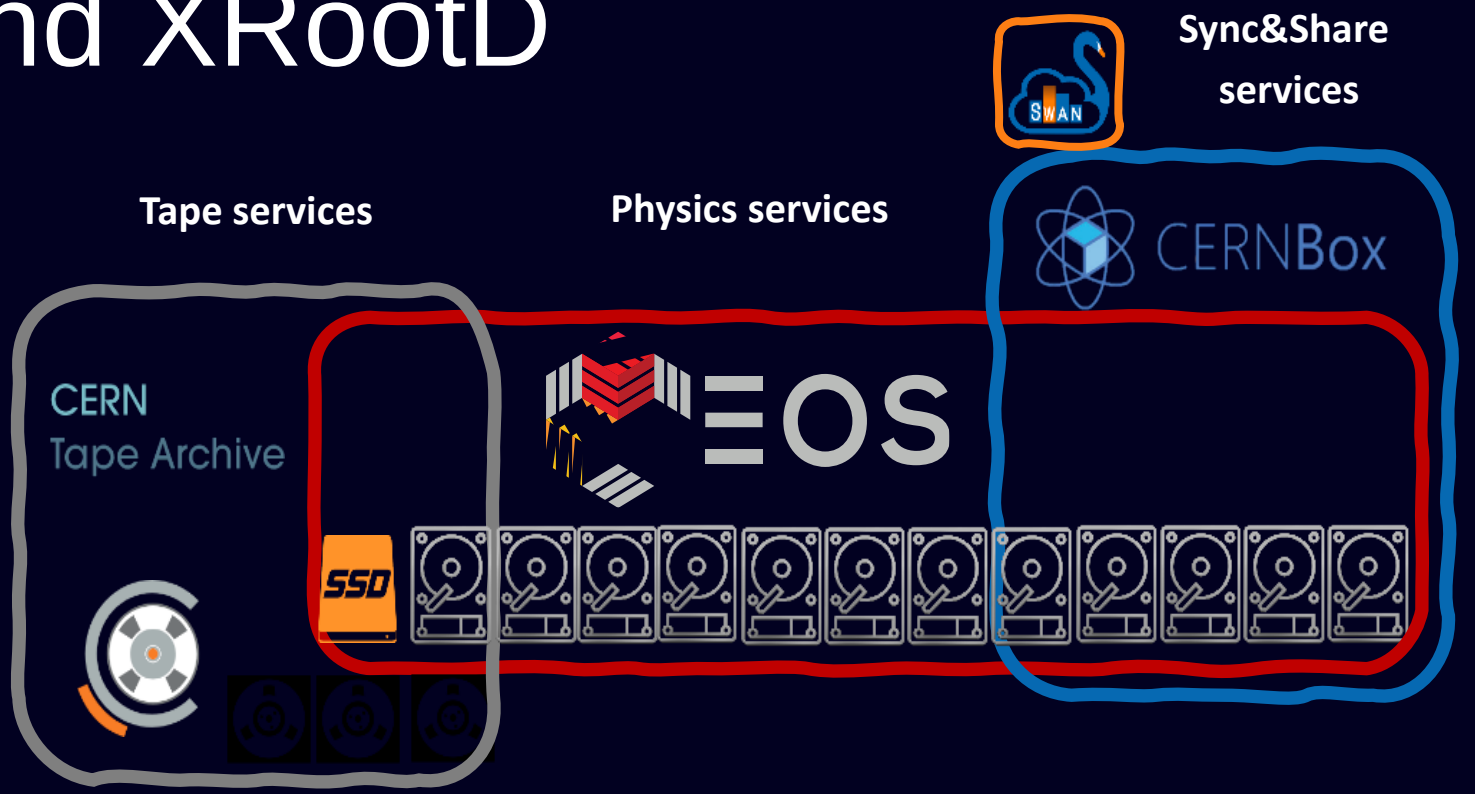

XRootD in EOS

- Feature-rich remote access protocol
- File access via native XRootD protocol

Old EOS client on Windows

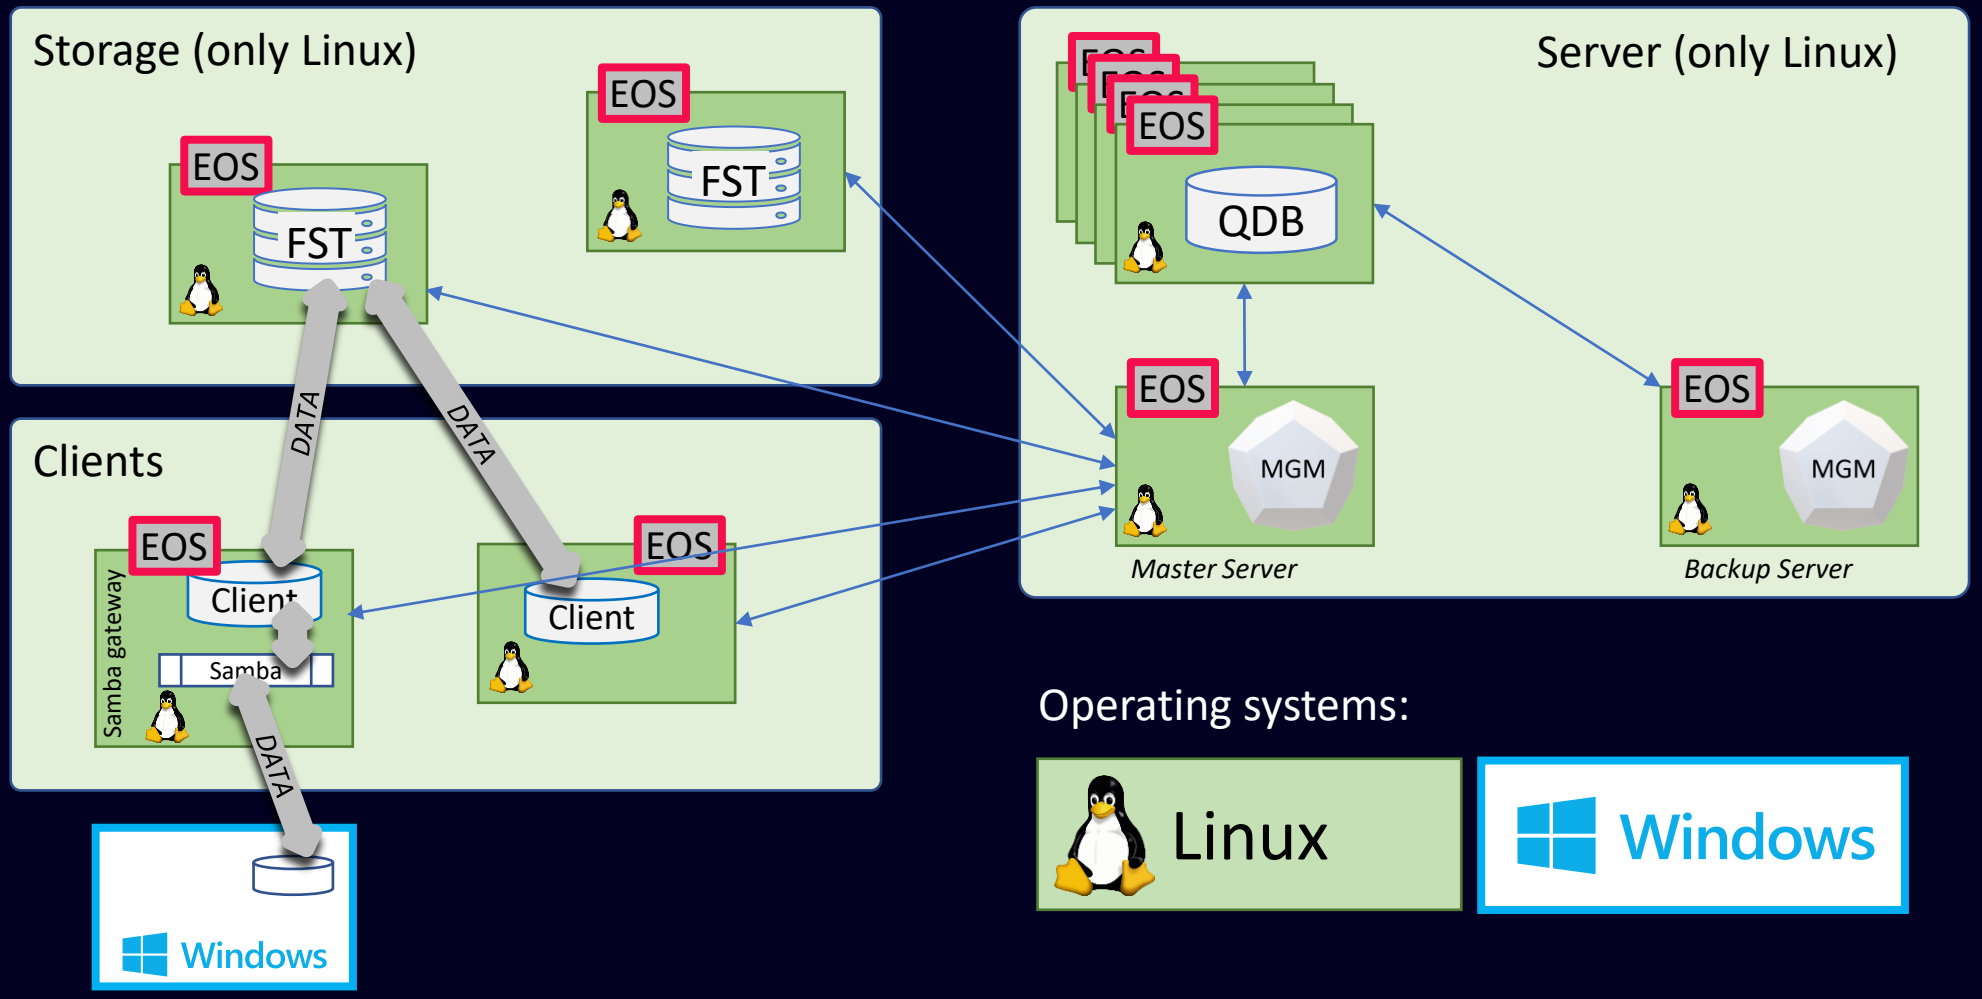

Operating systems:

Linux

Windows

Comtrade Three-Sixty

New EOS client on Windows

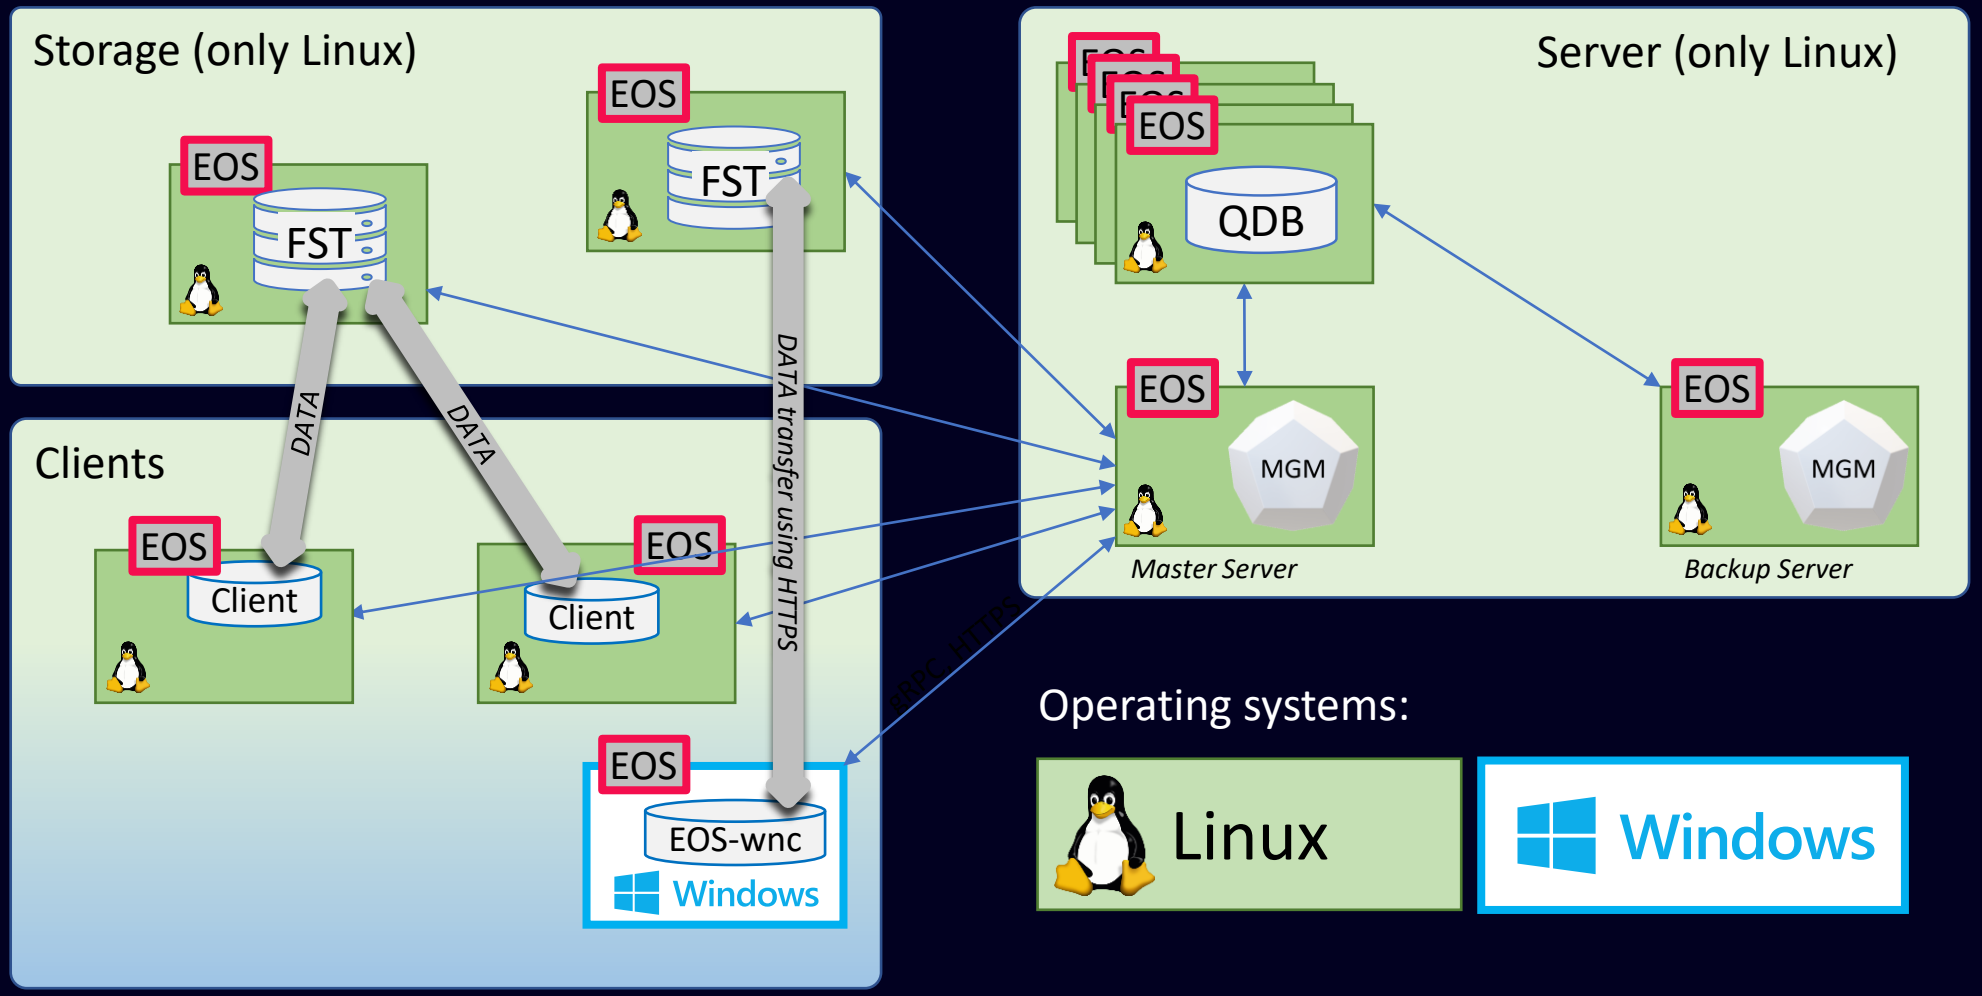

Implementation of XRootD in EOS

1. Introduction
2. Implementation of XRootD in EOS
    1. XRootD protocol for EOS
    2. FUSE client implementation (eosxd)
3. XRootD and EOS client
4. Windows port of XRootD client
5. Resume

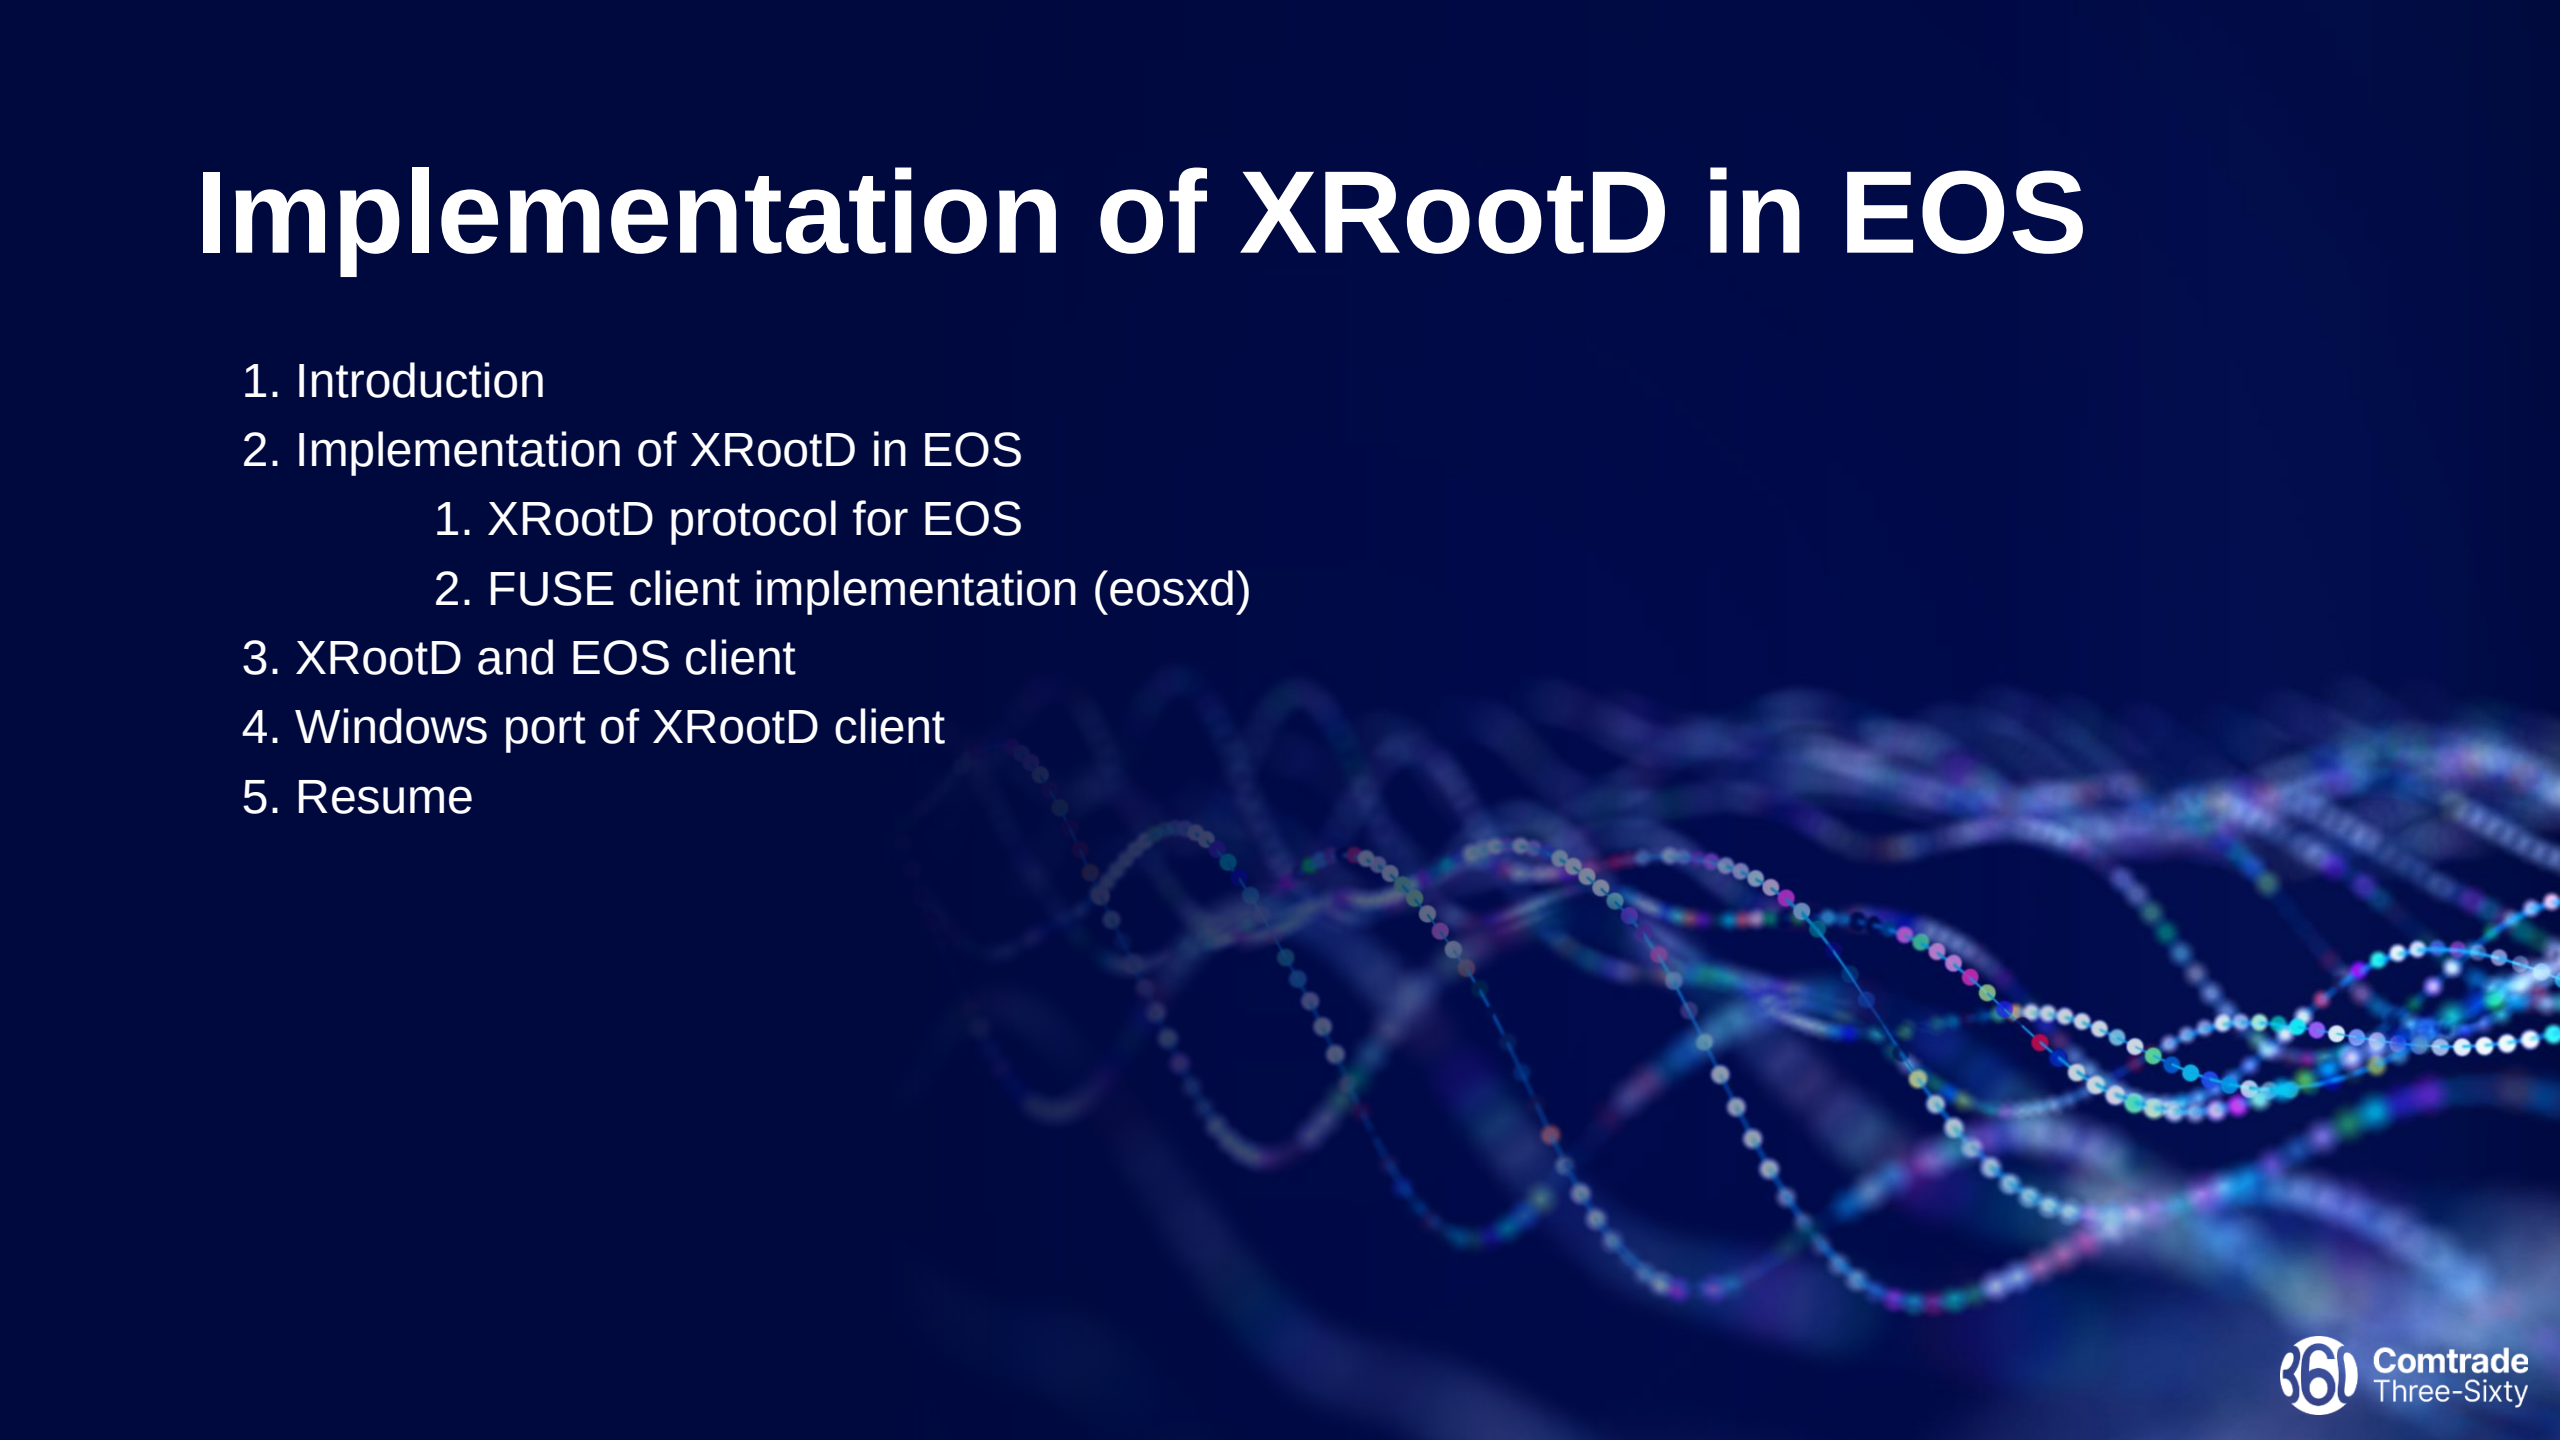

Comtrade Three-Sixty

XRootD protocol for EOS

XRootD protocol

- Strong authentication
- Redirection
- Remote access
- Third-party copy mechanism (TCP)

File system access with xrootdfs - native XRootD FUSE mount

FUSE client implementation (eosxd)

- FUSE
  - Allows to mount EOS as a filesystem

- FUSE client implementation (eosxd)

XRootD and EOS client

1. Introduction
2. Implementation of XRootD in EOS
3. XRootD and EOS client
    1. XRootD Clients and EOS
    2. Client component of XRootD
4. Windows port of XRootD client
5. Resume

Comtrade Three-Sixty

XRootD Clients and EOS

XRootD clients that EOS uses

- xrdcp
  - Copy application
- xrd
  - Filesystem application (listing, deletion ...)
- XrdCl
  - C++ client library

Client component of XRootD

- XrdCl - the new XRootD client
  - C++ implementation
  - multi-threaded implementation
- XrdCl - the basis for
  - xrdcp
  - Xrdfs

- XrdClient - the “old” XRootD client
  - available until XRootD v5

Windows port of XRootD client

1. Introduction
2. Implementation of XRootD in EOS
3. XRootD and EOS client
4. Windows port of XRootD client
    1. New XRootD client **XrdCl** on Windows
        1. Dependency packages for NEW xrdcp binary
        2. Windows porting issues: XRootD client XrdCl
        3. Networking needs pipes, sockets, flags, connections handling
        4. Porting of the security libraries
    2. New XRootD client XrdCl on Windows
        1. New XRootD client XrdCl on Windows
        2. Old XRootD client XrdClient on Windows
5. Resume

New XRootD client XrdCl on Windows

Investigation results

Dependency packages for NEW xrdcp binary

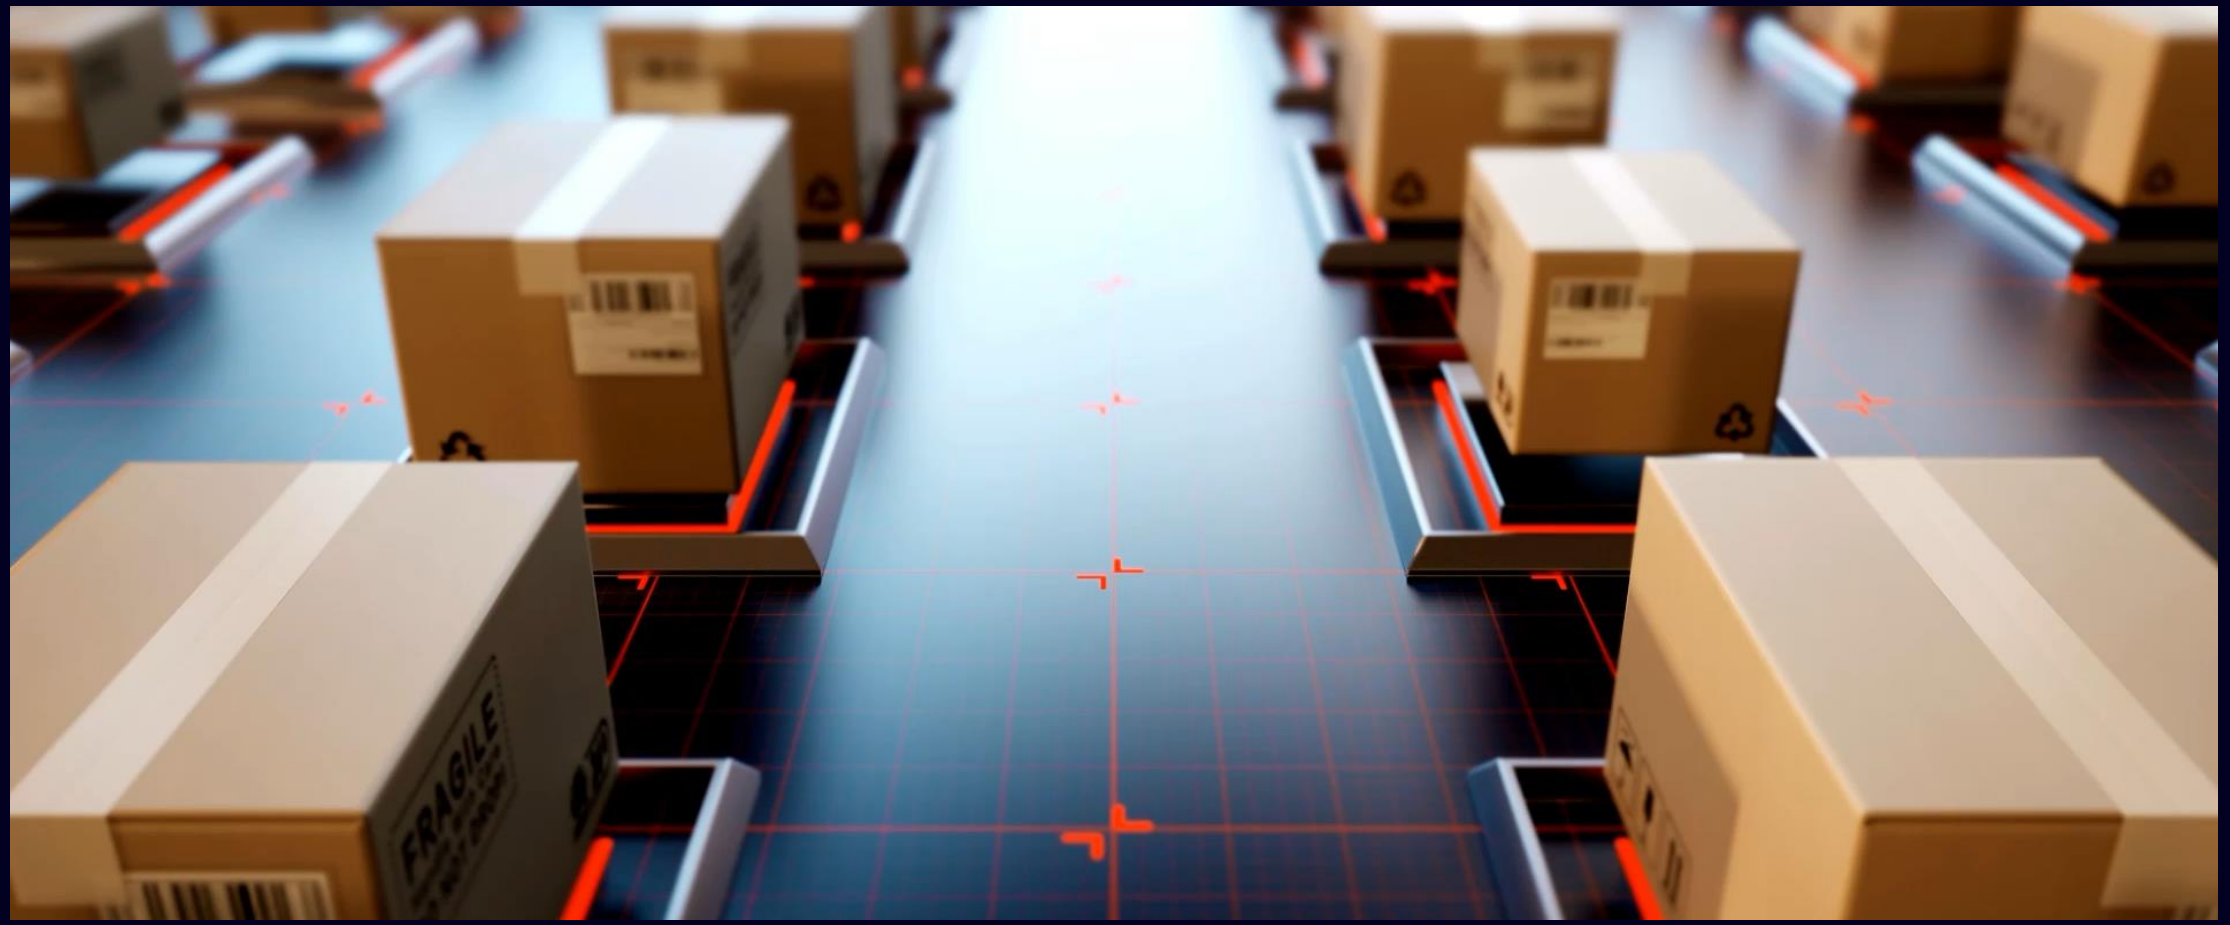

Dependency packages for NEW xrdcp binary

- pthread for Windows
- zlib (built with vcpkg)
- libxml2 (built with vcpkg)
  - requires iconv header file (used from vcpkg)
- dirent (used from https://github.com/tronkko/dirent)
- ibiconv (built with vcpkg)
- lzma (built with vcpkg)

Windows porting issues: XRootD client XrdCl

Security libraries:	libXrdSec-4.so
	libXrdSecgsi-4.so
	libXrdSecgsiAUTHZVO-4.so
	libXrdSecgsiGMAPDN-4.so
	libXrdSeckrb5-4.so
	libXrdSecProt-4.so
	libXrdSecpwd-4.so
	libXrdSecsss-4.so
	libXrdSecunix-4.so

Critical:

- libXrdXml.so

Networking needs pipes, sockets, flags, connections handling

Comtrade Three-Sixty

Networking needs pipes, sockets, flags, connections handling

Network libraries that are difficult to port to Windows

- fcntl()
- poll()
- select()
- pipe2()
- socketpair()

Porting of the security libraries

Security libraries that are difficult to port to Windows

- libXrdSec-4.so
- libXrdSecgsi-4.so
- libXrdSecgsiAUTHZVO-4.so
- libXrdSecgsiGMAPDN-4.so
- libXrdSeckrb5-4.so
- libXrdSecProt-4.so
- libXrdSecpwd-4.so
- libXrdSecsss-4.so
- libXrdSecunix-4.so

Decision

Porting of the old XRootD client XrdClient on Windows

New XRootD client XrdCl on Windows

Minimal requirement of libraries to build copy binary *xrdcp*:

- libXrdAppUtils.so
- libXrdUtils.so
- libXrdXml.so
- libXrdCl.so

Old XRootD client XrdClient on Windows

Minimal requirement of libraries to build copy binary *xrdcp-old*:

- libXrdAppUtils.so
- libXrdUtils.so
- libXrdClient.so

Successful porting!

Resume

1. Introduction
2. Implementation of XRootD in EOS
3. XRootD and EOS client
4. Windows port of XRootD client
5. Resume
    1. Out of Comtrade's investments
    2. Speed test results of EOS-wnc

Out of Comtrade's investments

Open-source community future work

- Using XrdHttp for EOS-wnc
- Support new Windows libraries
- Open source support for Windows libraries

Speed test results of EOS-wnc

Iterations (EOS)	26	(checksums OK)						
Iterations (IBM)	28	(checksums OK)						
Iterations (Ceph)	52	(checksums OK)	Number of files	30				
Iterations (Hadoop)	11	(checksums OK)	File size [MB]	1				

		Test [MB/s]	min	max	avg	trim25%	Avg time [ms]
Upload	Linux	EOS: xrdcp command	76,92	115,38	90,16	89,15	11,09
		EOS Fusex	26,95	59,64	45,10	45,86	22,17
		IBM Spectrum Scale	51,46	62,37	57,56	57,68	17,37
		Ceph on Linux	133,33	159,57	151,78	152,18	6,59
		Hadoop on Linux	2,51	2,76	2,62	2,61	382,21
	Windows	EOS-wnc	13,47	14,21	13,92	13,93	71,82
		EOS-drive ST	9,72	10,02	9,88	9,88	101,23
		EOS: Samba	8,45	25,04	22,58	24,32	44,28
		Ceph on Win	43,45	55,10	51,54	51,80	19,40
		Hadoop on Win	2,11	2,27	2,20	2,20	454,92
Download	Linux	EOS: xrdcp command	23,44	130,43	106,68	112,56	9,37
		EOS Fusex	28,82	41,44	36,07	36,36	27,73
		IBM Spectrum Scale	125,52	193,55	154,49	153,73	6,47
		Ceph on Linux	0,51	103,81	82,94	83,81	12,06
		Hadoop on Linux	2,69	2,97	2,86	2,86	350,10
	Windows	EOS-wnc	9,04	10,72	10,32	10,35	96,93
		EOS-drive ST	14,48	18,02	15,98	15,94	62,59
		EOS: Samba	7,32	15,19	12,63	13,06	79,17
		Ceph on Win	1,95	46,82	41,06	41,91	24,35
		Hadoop on Win	0,89	2,50	1,75	1,76	571,13

Legend:	[MB/s]		
Red	0	-	10
Orange	10	-	20
Yellow	20	-	30
Light green	30	-	40
Green	40	-	∞

Iterations (EOS)	28	(checksums OK)						
Iterations (IBM)	28	(checksums OK)						
Iterations (Ceph)	52	(checksums OK)	Number of files	10				
Iterations (Hadoop)	11	(checksums OK)	File size [MB]	100				

		Test [MB/s]	min	max	avg	trim25%	Avg time [ms]
Upload	Linux	EOS: xrdcp command	359,71	444,44	411,14	412,67	243,22
		EOS Fusex	134,72	192,64	160,16	160,07	624,38
		IBM Spectrum Scale	176,62	188,71	181,08	181,01	552,25
		Ceph on Linux	131,89	162,39	140,86	139,97	709,95
		Hadoop on Linux	9,43	10,17	9,94	9,97	10064,96
	Windows	EOS-wnc	174,21	204,60	186,28	185,66	536,82
		EOS-drive ST	186,54	210,24	197,67	197,40	505,89
		EOS: Samba	165,56	231,33	196,82	196,65	508,08
		Ceph on Win	102,68	141,96	136,28	137,07	733,77
		Hadoop on Win	4,50	5,22	4,61	4,56	21670,61
Download	Linux	EOS: xrdcp command	301,20	436,68	412,94	417,50	242,16
		EOS Fusex	186,67	217,11	206,36	207,14	484,59
		IBM Spectrum Scale	306,75	345,18	322,89	322,24	309,70
		Ceph on Linux	20,44	183,49	31,49	28,27	3175,40
		Hadoop on Linux	8,06	10,51	9,37	9,39	10668,22
	Windows	EOS-wnc	128,10	177,31	151,39	151,01	660,54
		EOS-drive ST	148,70	185,92	157,76	156,40	633,87
		EOS: Samba	72,97	97,50	81,26	80,39	1230,60
		Ceph on Win	17,63	81,00	25,54	23,66	3915,54
		Hadoop on Win	4,31	4,62	4,50	4,51	22217,73

Legend:	[MB/s]		
Red	0	-	100
Orange	100	-	150
Yellow	150	-	200
Light green	200	-	250
Green	250	-	∞

Iterations (EOS)	27	(checksums OK)						
Iterations (IBM)	28	(checksums OK)						
Iterations (Ceph)	52	(checksums OK)	Number of files	2				
Iterations (Hadoop)	11	(checksums OK)	File size [MB]	2000				

		Test [MB/s]	min	max	avg	trim25%	Avg time [s]
Upload	Linux	EOS: xrdcp command	329,49	405,27	371,03	371,17	5,39
		EOS Fusex	187,92	237,63	210,76	210,51	9,49
		IBM Spectrum Scale	283,61	318,22	294,47	293,28	6,79
		Ceph on Linux	141,00	163,37	157,56	158,17	12,69
		Hadoop on Linux	9,74	10,10	9,91	9,91	201,83
	Windows	EOS-wnc	158,40	331,09	231,25	227,75	8,65
		EOS-drive ST	212,22	294,47	237,44	234,72	8,42
		EOS: Samba	164,11	229,82	181,25	178,59	11,03
		Ceph on Win	128,18	158,19	153,32	154,04	13,04
		Hadoop on Win	4,61	4,72	4,66	4,66	428,85
Download	Linux	EOS: xrdcp command	328,68	365,97	353,00	354,47	5,67
		EOS Fusex	218,66	233,36	227,13	227,15	8,81
		IBM Spectrum Scale	328,95	364,96	342,54	341,65	5,84
		Ceph on Linux	188,80	355,49	265,08	264,04	7,54
		Hadoop on Linux	9,28	10,63	10,12	10,15	197,66
	Windows	EOS-wnc	119,92	213,86	170,17	169,49	11,75
		EOS-drive ST	179,86	210,49	190,24	189,72	10,51
		EOS: Samba	17,95	35,43	25,85	25,54	77,37
		Ceph on Win	105,38	141,66	122,82	122,90	16,28
		Hadoop on Win	4,30	4,73	4,55	4,55	440,00

Legend:	[MB/s]		
Red	0	-	100
Orange	100	-	150
Yellow	150	-	200
Light green	200	-	250
Green	250	-	∞

^ST - Single-thread
^MT - Multi-thread

Comtrade Three-Sixty

Speed test results of EOS-wnc

Iterations (EOS)	27	(checksums OK)	
Iterations (IBM)	28	(checksums OK)	
Iterations (Ceph)	52	(checksums OK)	Number of files: 2
Iterations (Hadoop)	11	(checksums OK)	File size [MB]: 2000

		Test [MB/s]	min	max	avg	trim25%	Avg time [s]
Upload	Linux	EOS: xrdcp command	329,49	405,27	371,03	371,17	5,39
		EOS Fusex	187,82	237,63	210,76	210,51	9,49
		IBM Spectrum Scale	283,61	318,22	294,47	293,28	6,79
		Ceph on Linux	141,00	163,37	157,56	158,17	12,69
		Hadoop on Linux	9,74	10,10	9,91	9,91	201,83
	Windows	EOS-wnc	158,40	331,09	231,25	227,75	8,65
		EOS-drive ST	212,22	294,47	237,44	234,72	8,42
		EOS: Samba	164,11	229,82	181,25	178,59	11,03
		Ceph on Win	128,18	158,19	153,32	154,04	13,04
		Hadoop on Win	4,61	4,72	4,66	4,66	428,85
Download	Linux	EOS: xrdcp command	328,68	365,97	353,00	354,47	5,67
		EOS Fusex	218,66	233,36	227,13	227,15	8,81
		IBM Spectrum Scale	328,95	364,96	342,54	341,65	5,84
		Ceph on Linux	188,80	355,49	265,08	264,04	7,54
		Hadoop on Linux	9,28	10,63	10,12	10,15	197,66
	Windows	EOS-wnc	119,92	213,86	170,17	169,49	11,75
		EOS-drive ST	179,86	210,49	190,24	189,72	10,51

Comtrade Three-Sixty

Comtrade Three-Sixty

Thank you

Porting of XRootD to Windows as a part of EOS-wnc

Gregor Molan
gregor.molan@comtrade.com

Branko Blagojević
branko.blagojevic@comtrade.com

Ivan Arizanović
ivan.arizanovic@comtrade.com

Comtrade Group / Comtrade 360