

# EOS-drive for Windows: Architecture and file transferring system

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



# Content

- Introduction
- Architecture of EOS-drive
- File transferring system in EOS-drive
- File transferring system testing
- Summary

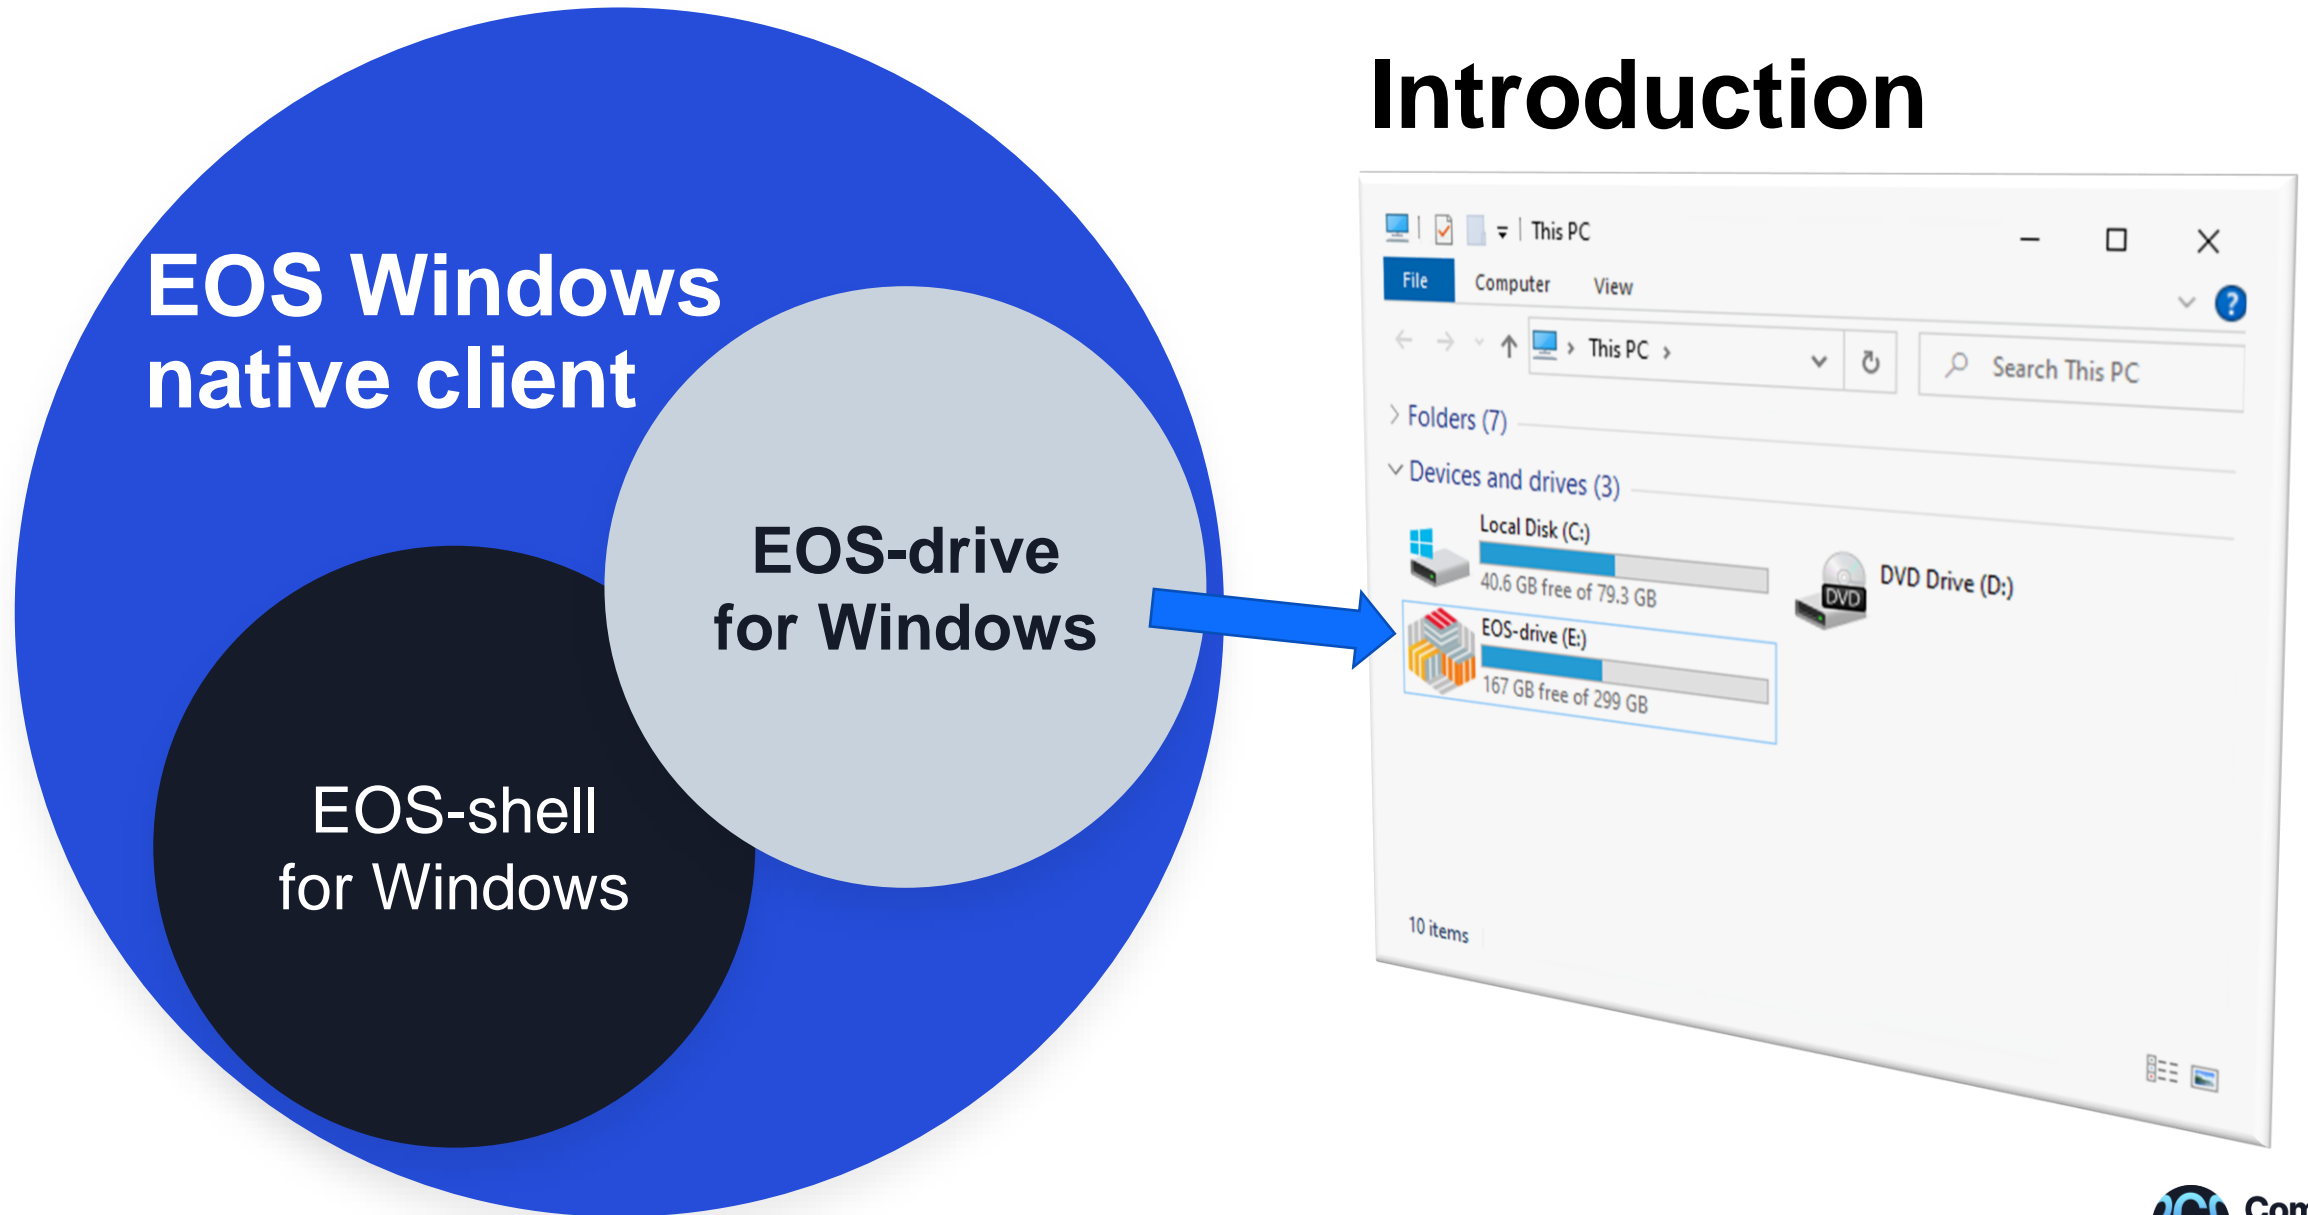


# Introduction

EOS Windows native client:

- EOS-shell
  - EOS via command line interface
- EOS-drive
  - EOS filesystem as a Windows disk drive

# Introduction



# Architecture




**Windows  
applications**



**Dokan:**  
kernel-mode driver  
user-mode library



**EOS-drive  
for Windows**



**HTTPS protocol:**  
cURL library for file transfer  
gRPC library for metadata

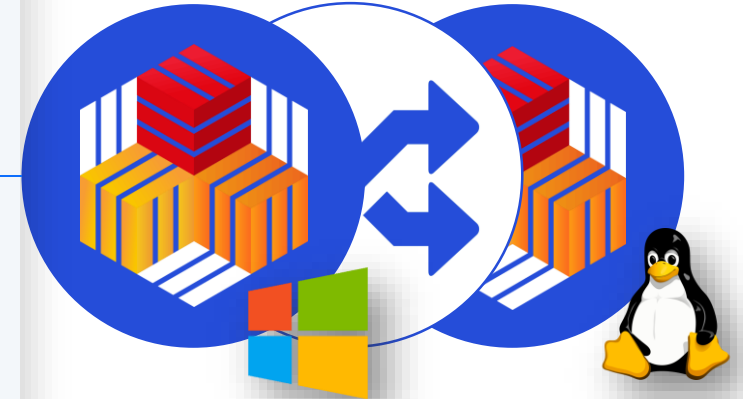
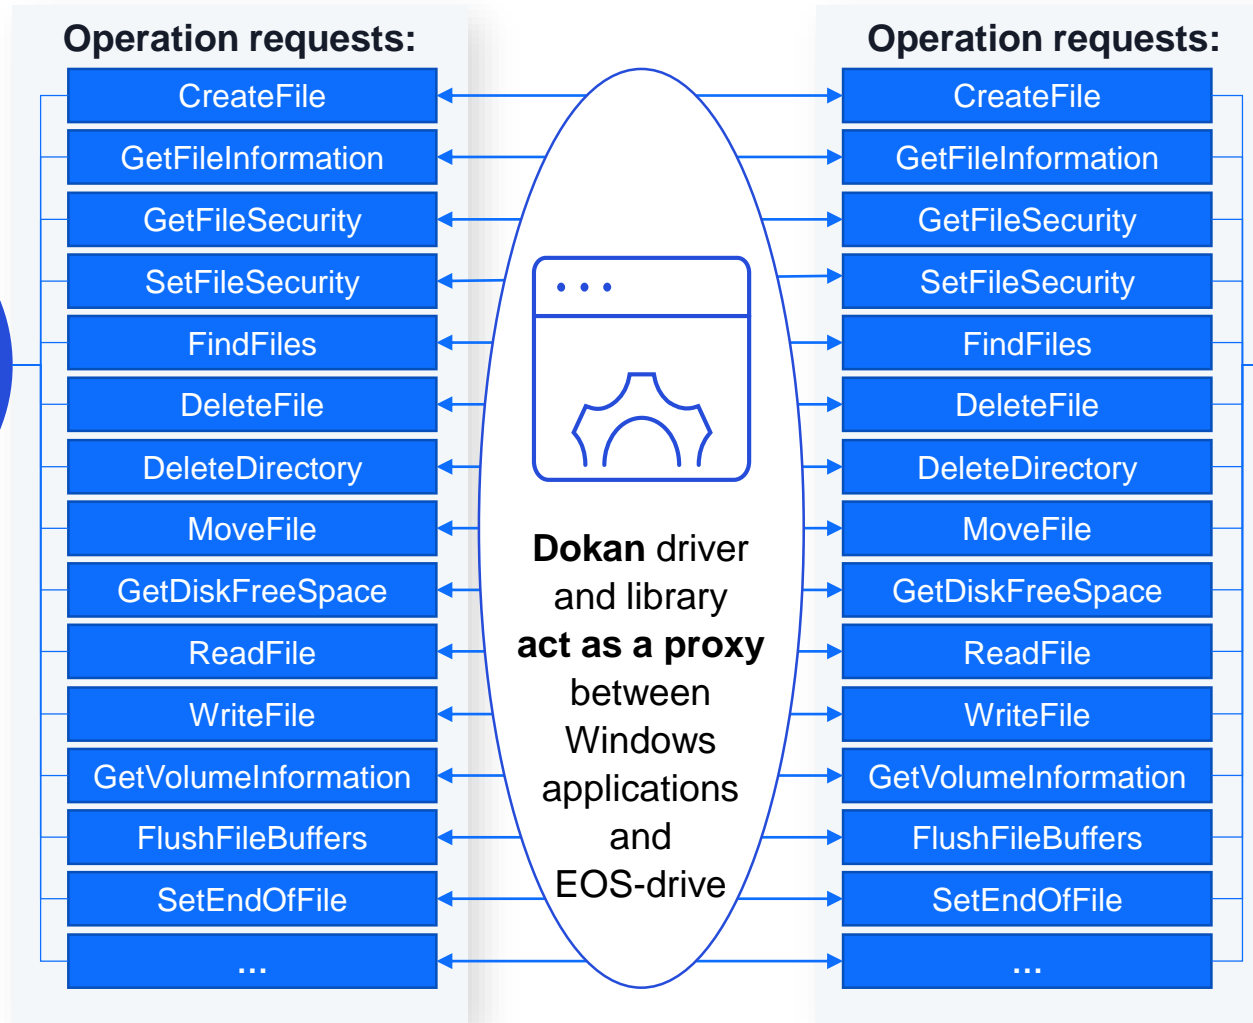


**EOS cluster  
on Linux**

# Architecture

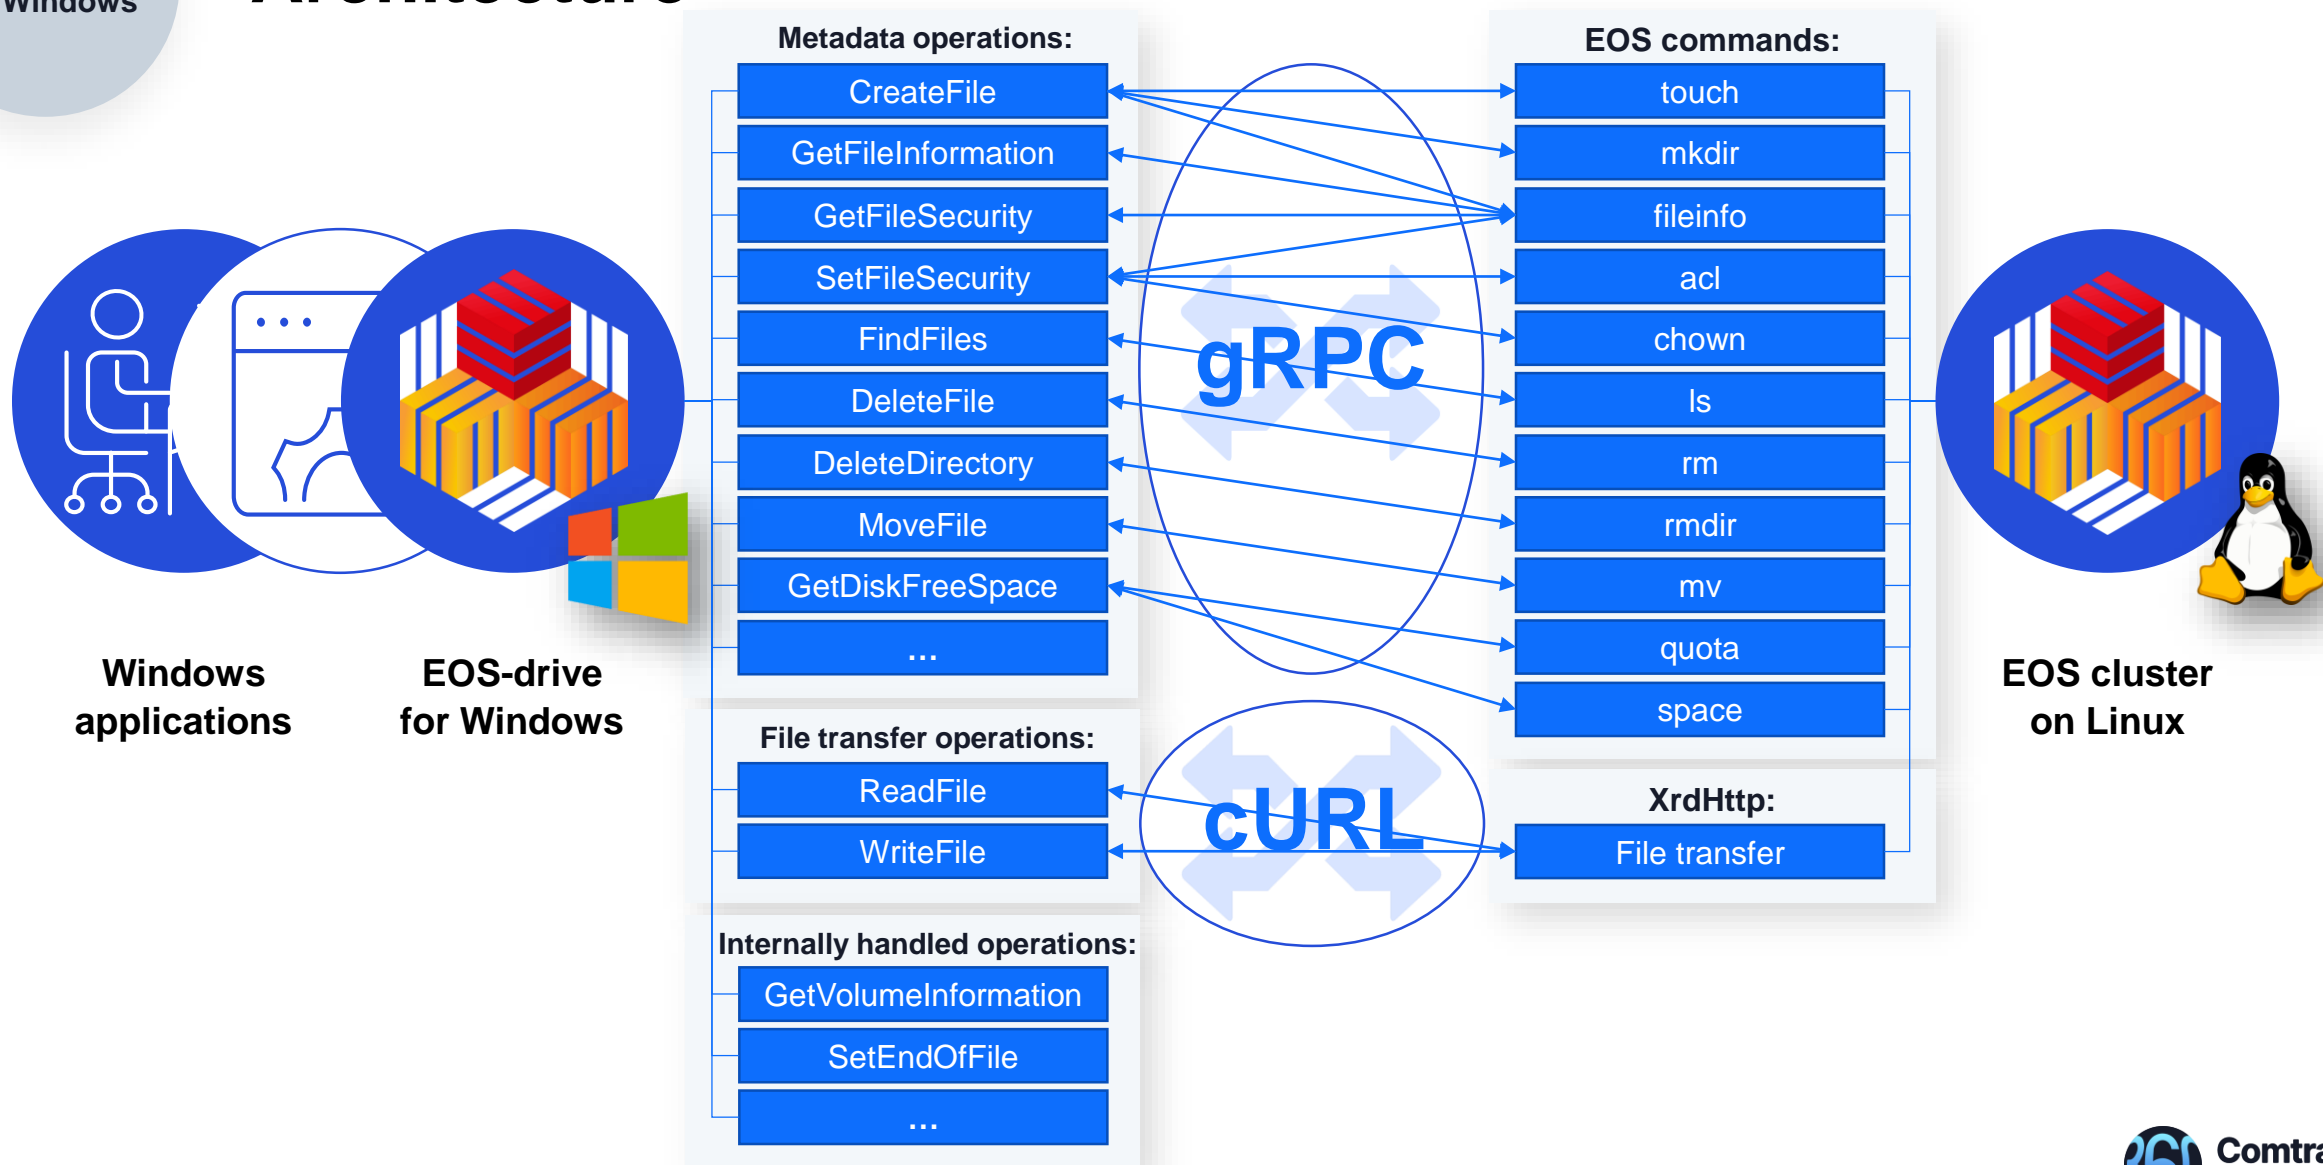


Windows  
applications

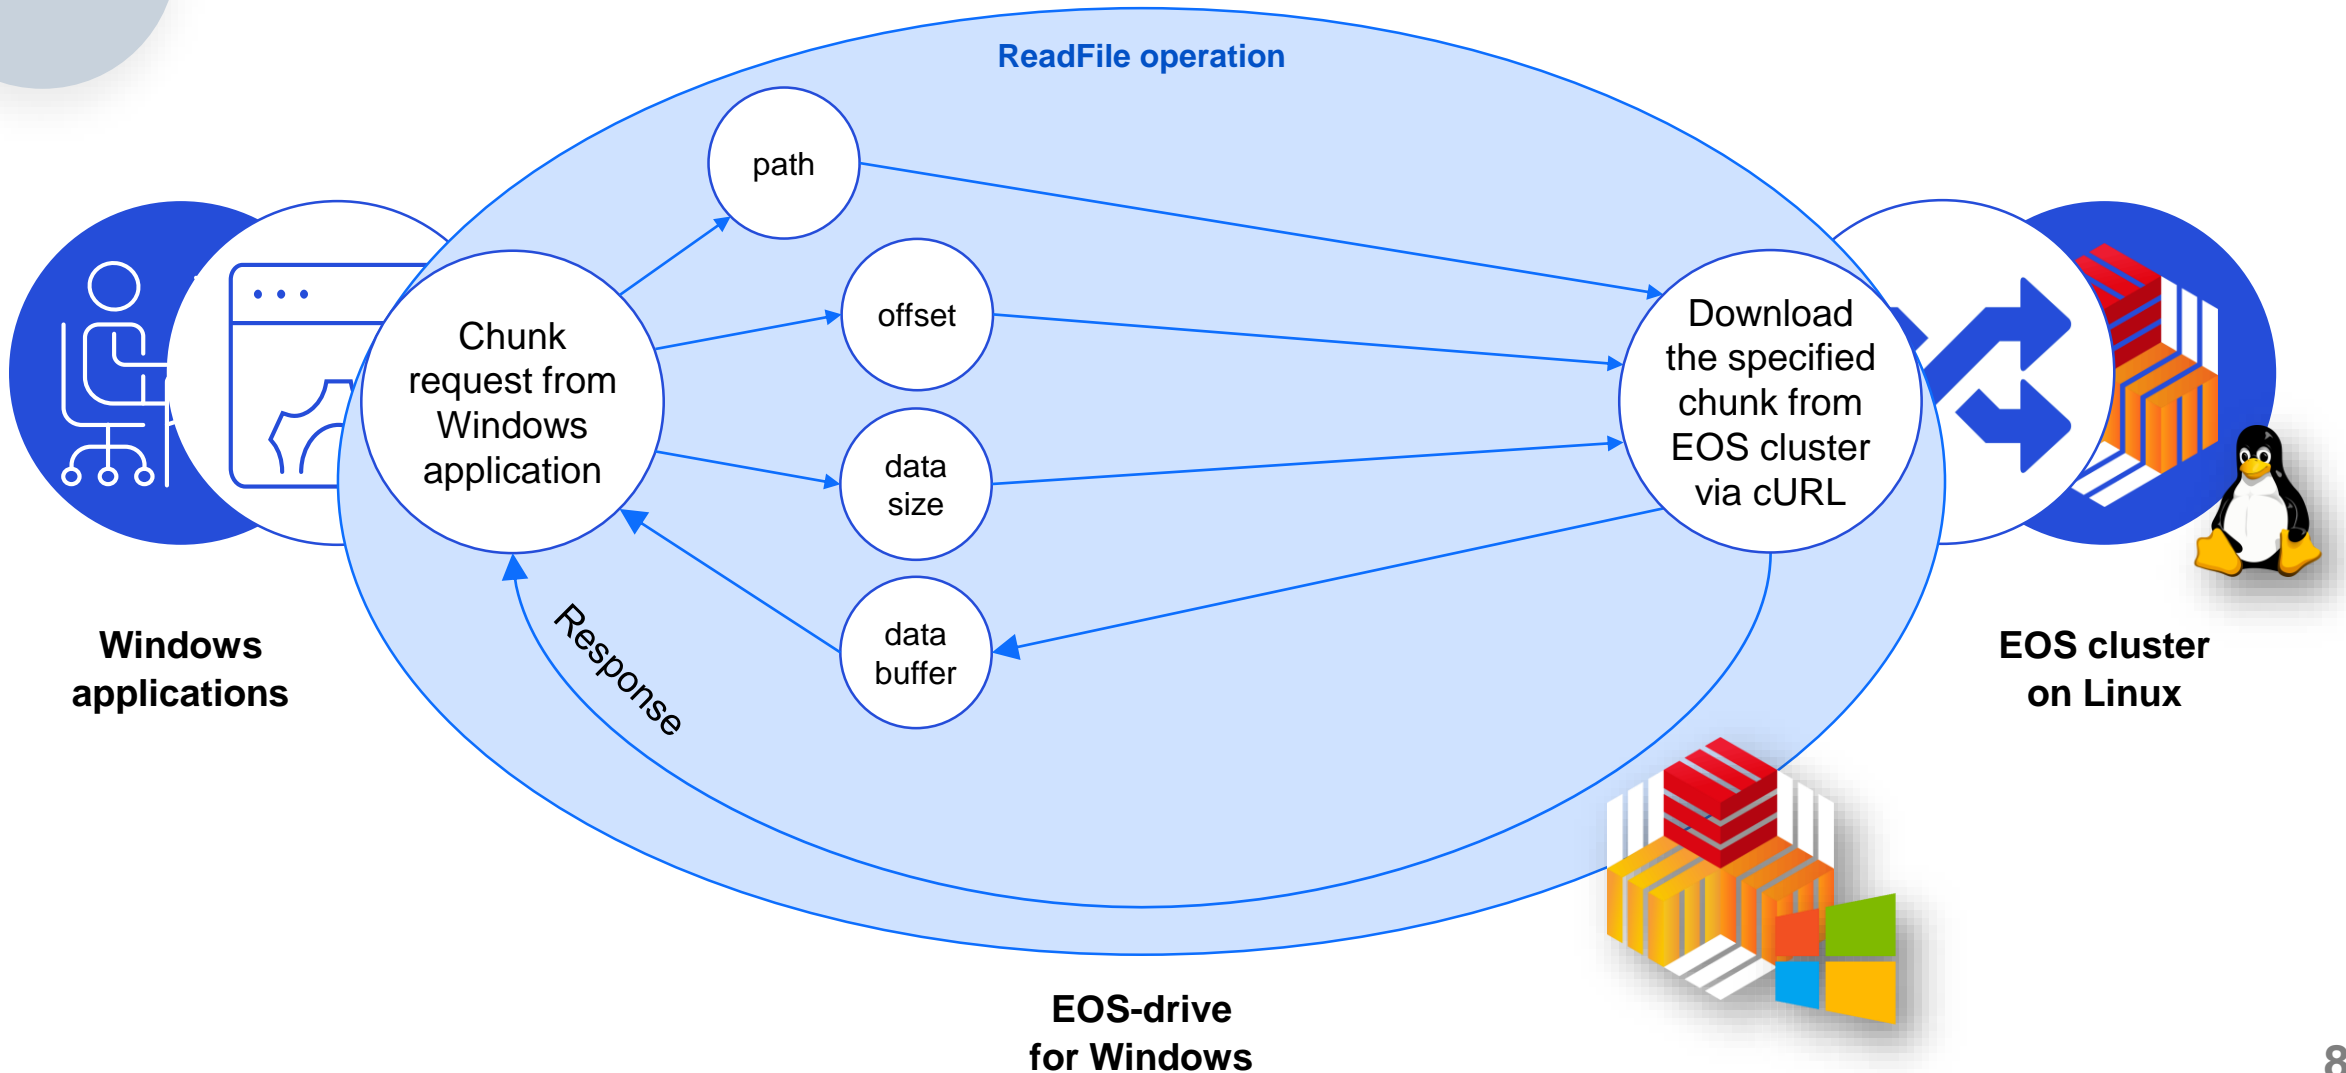


EOS-drive  
for Windows

# Architecture

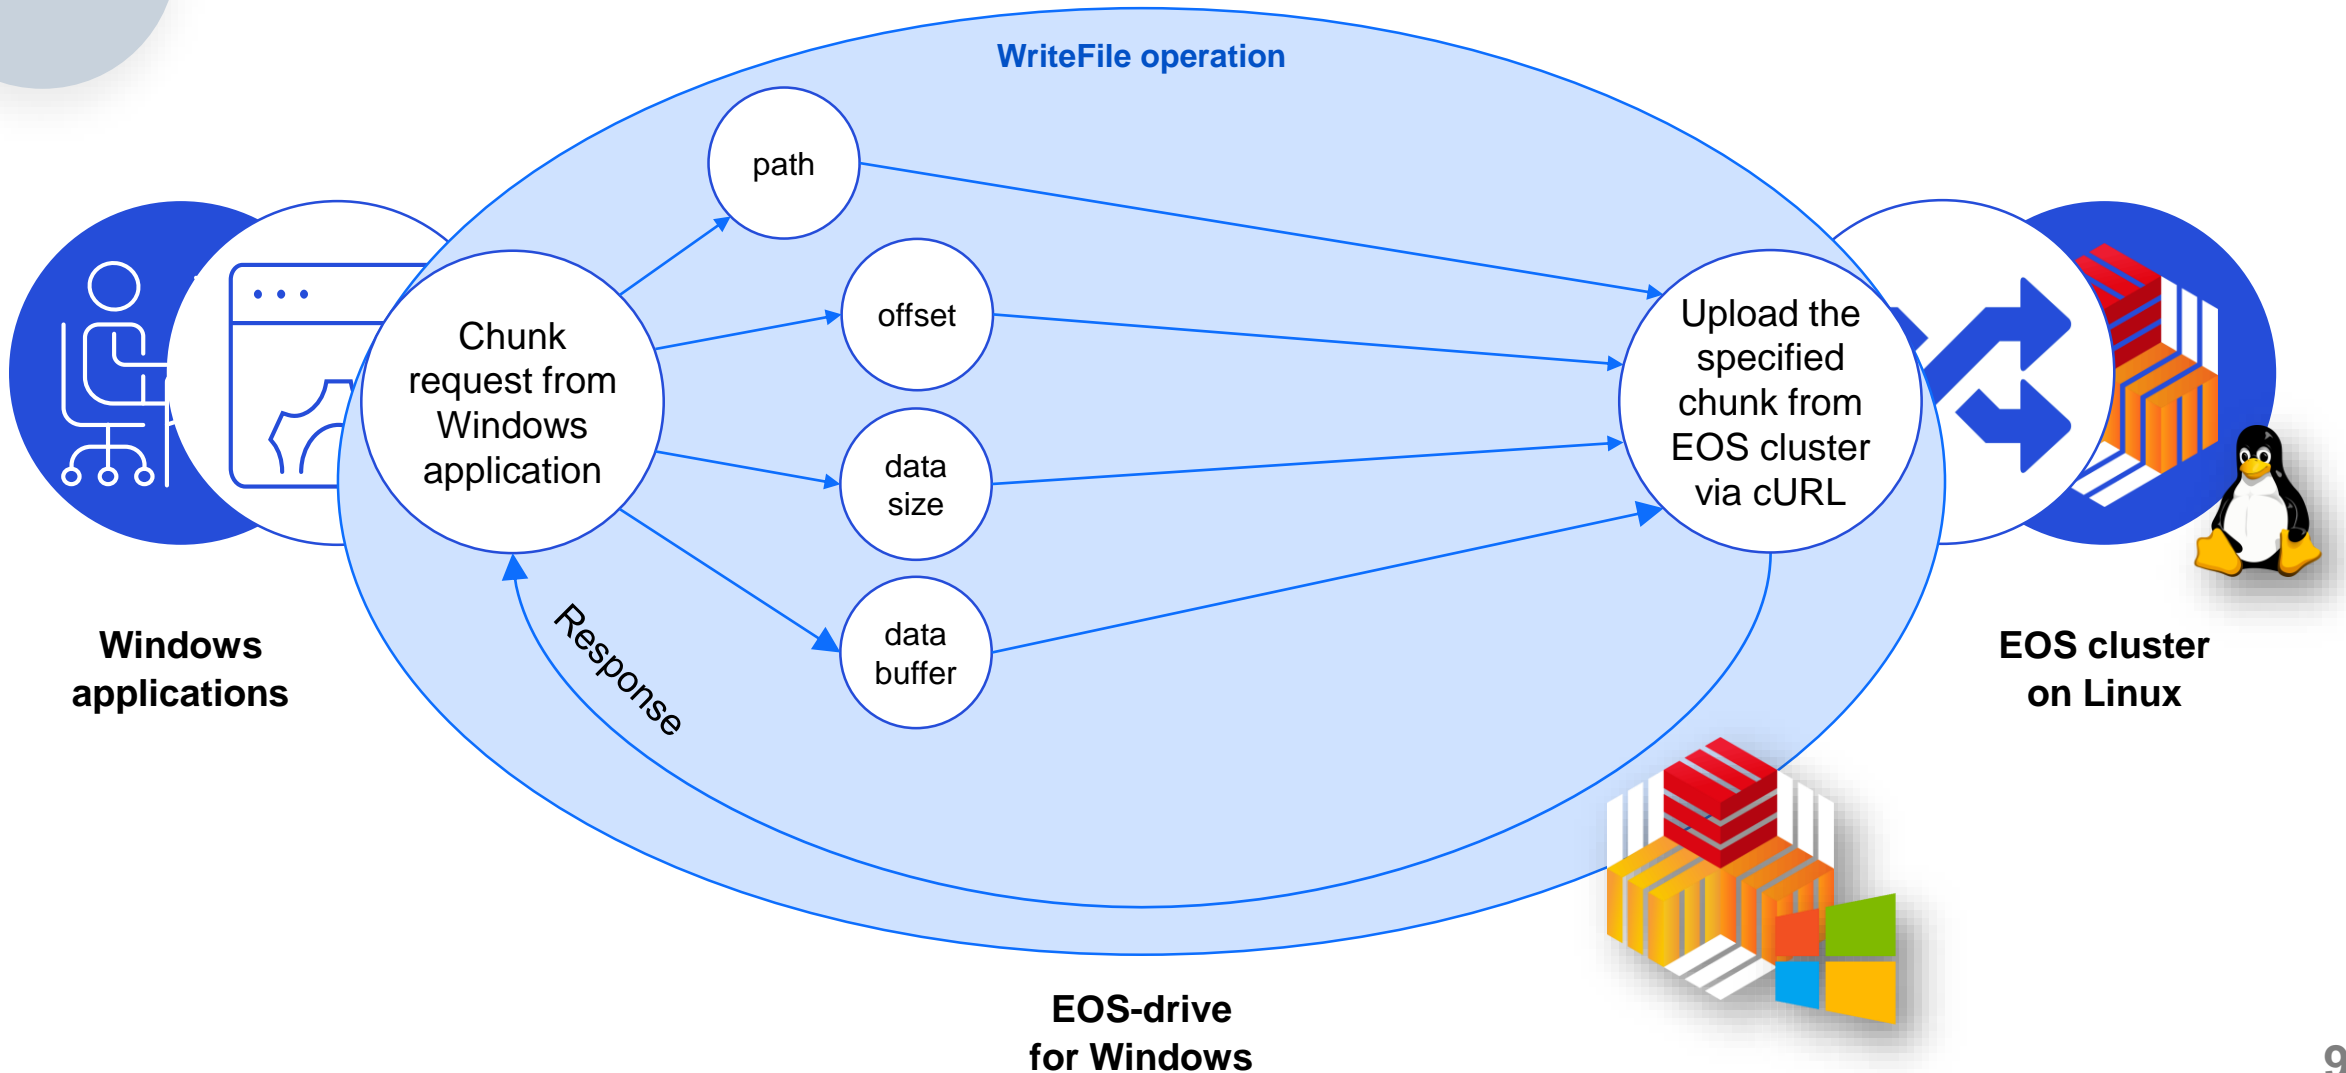


# File transferring chunk by chunk for downloading

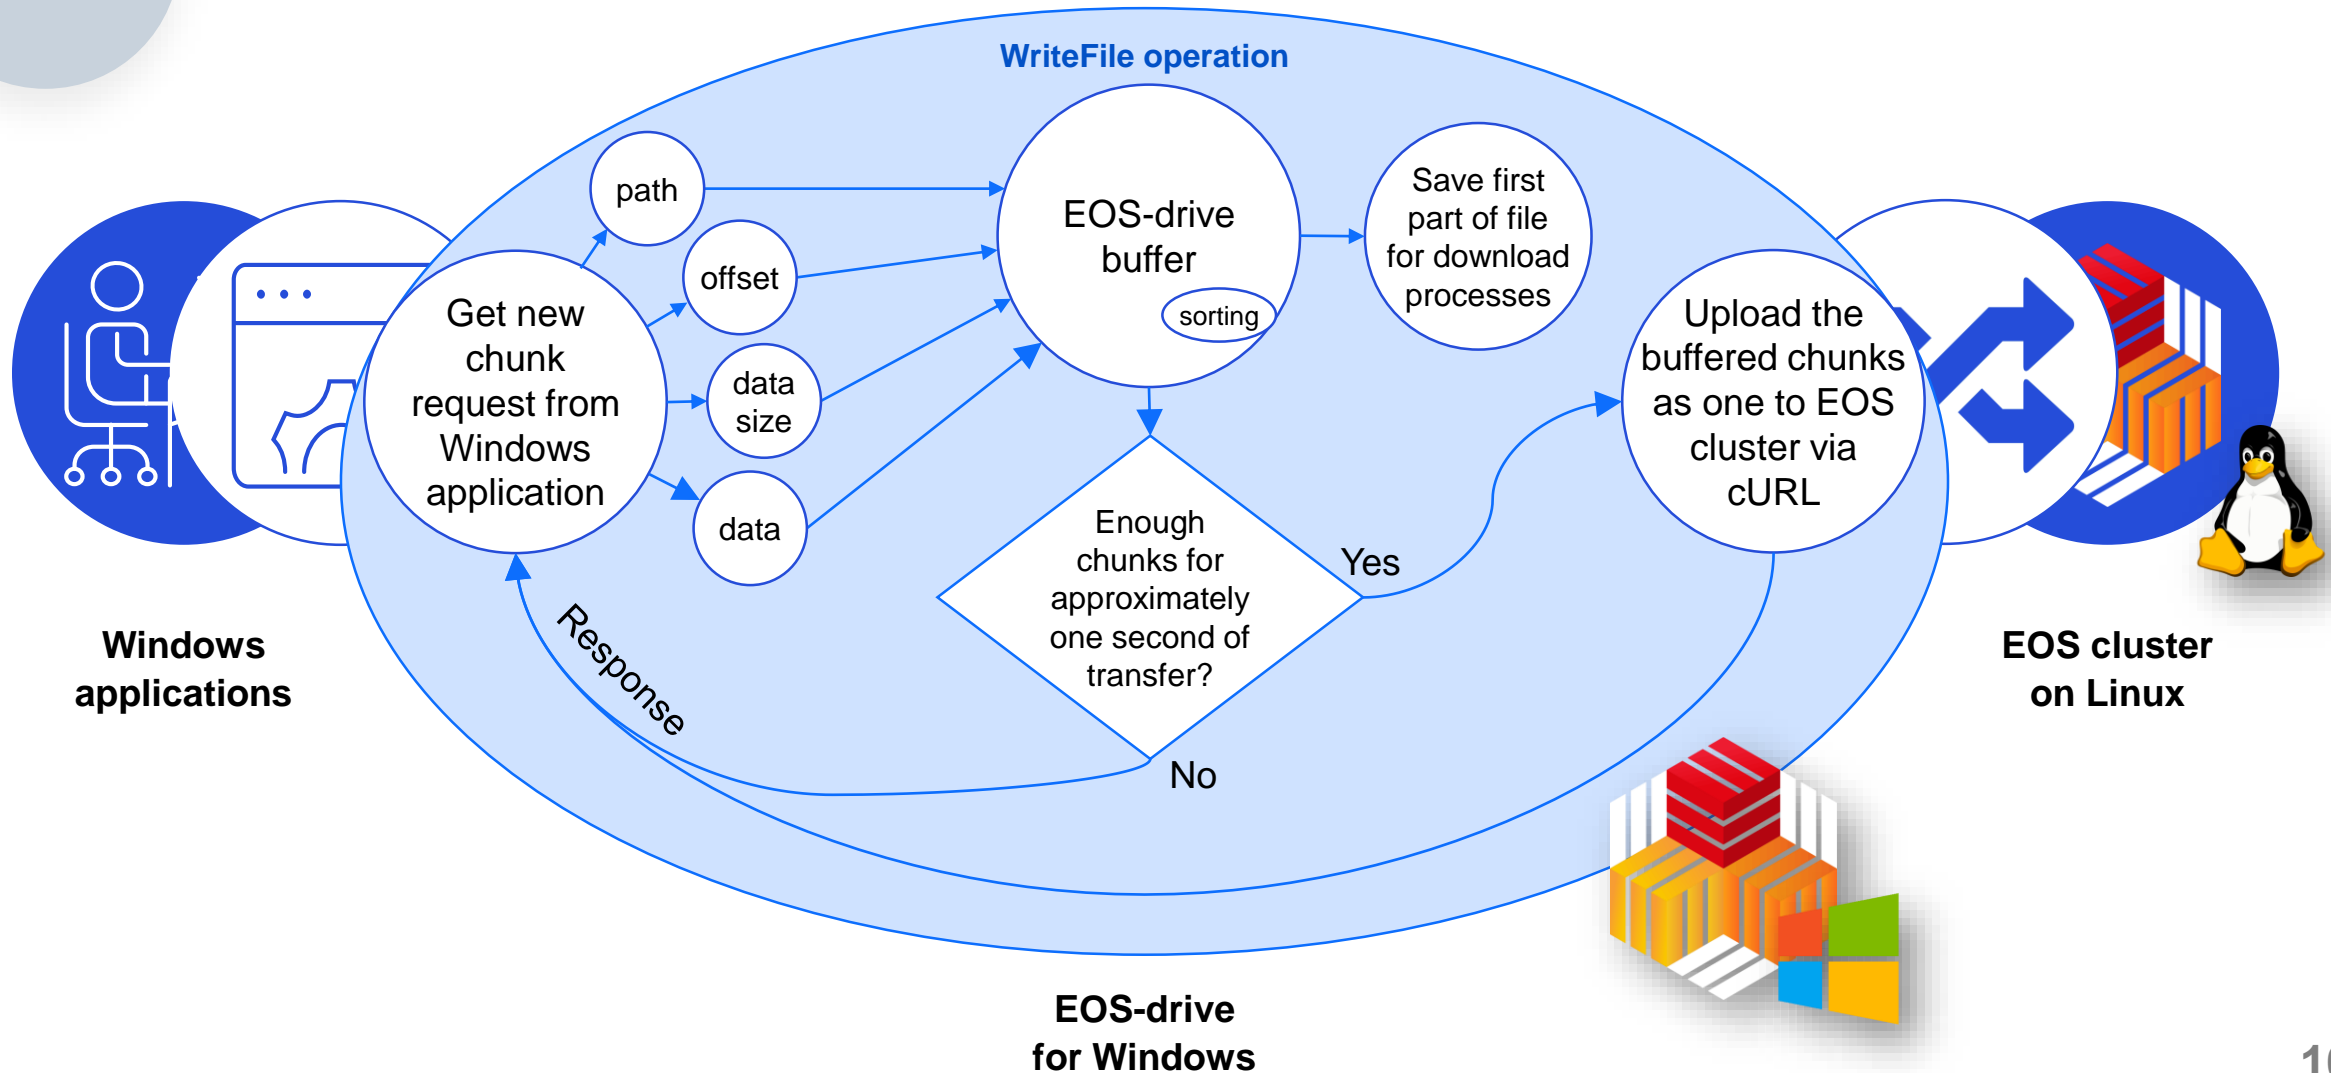




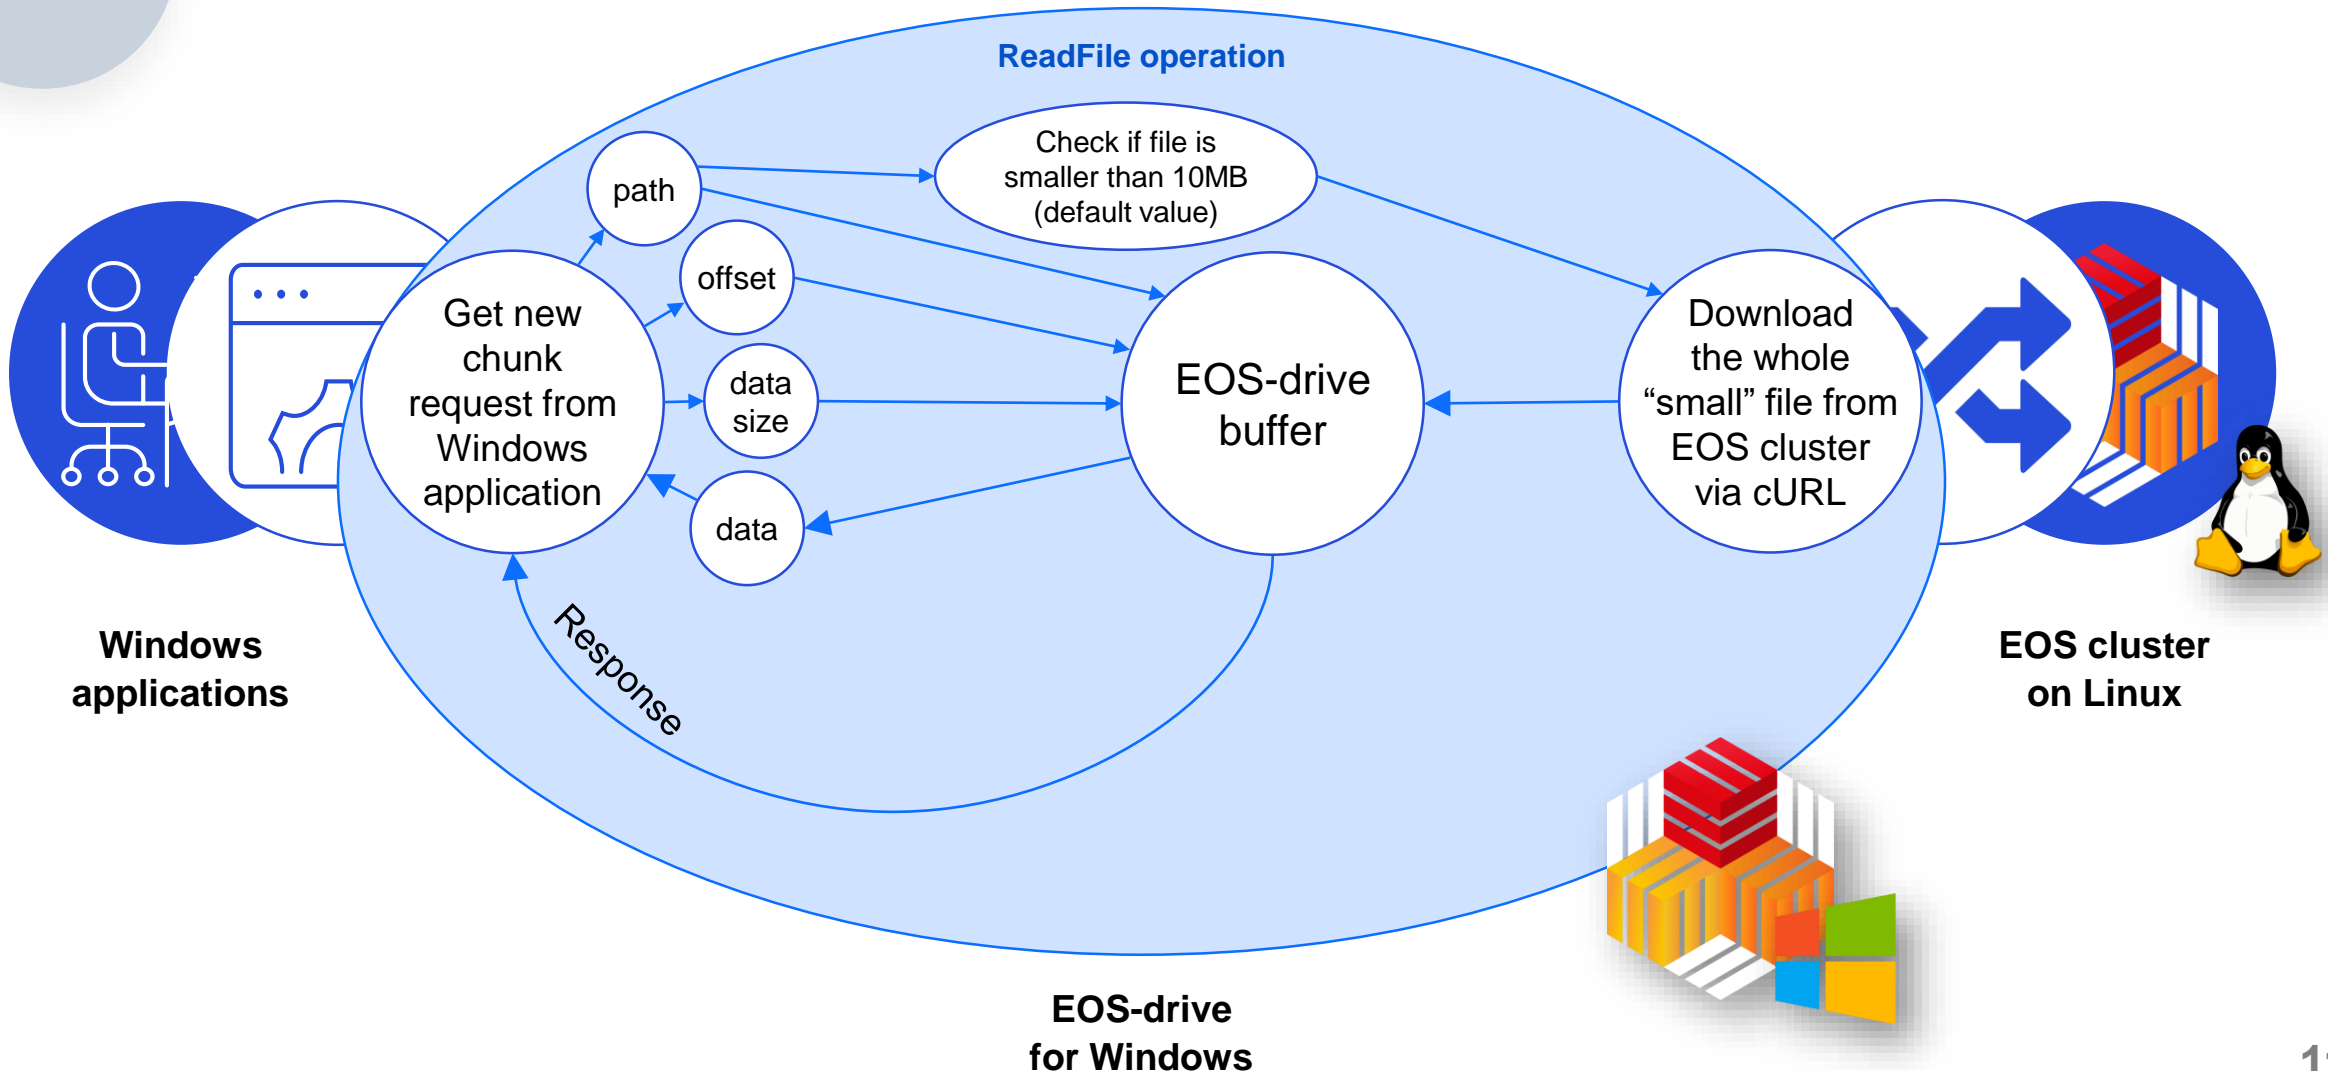
# File transferring chunk by chunk for uploading



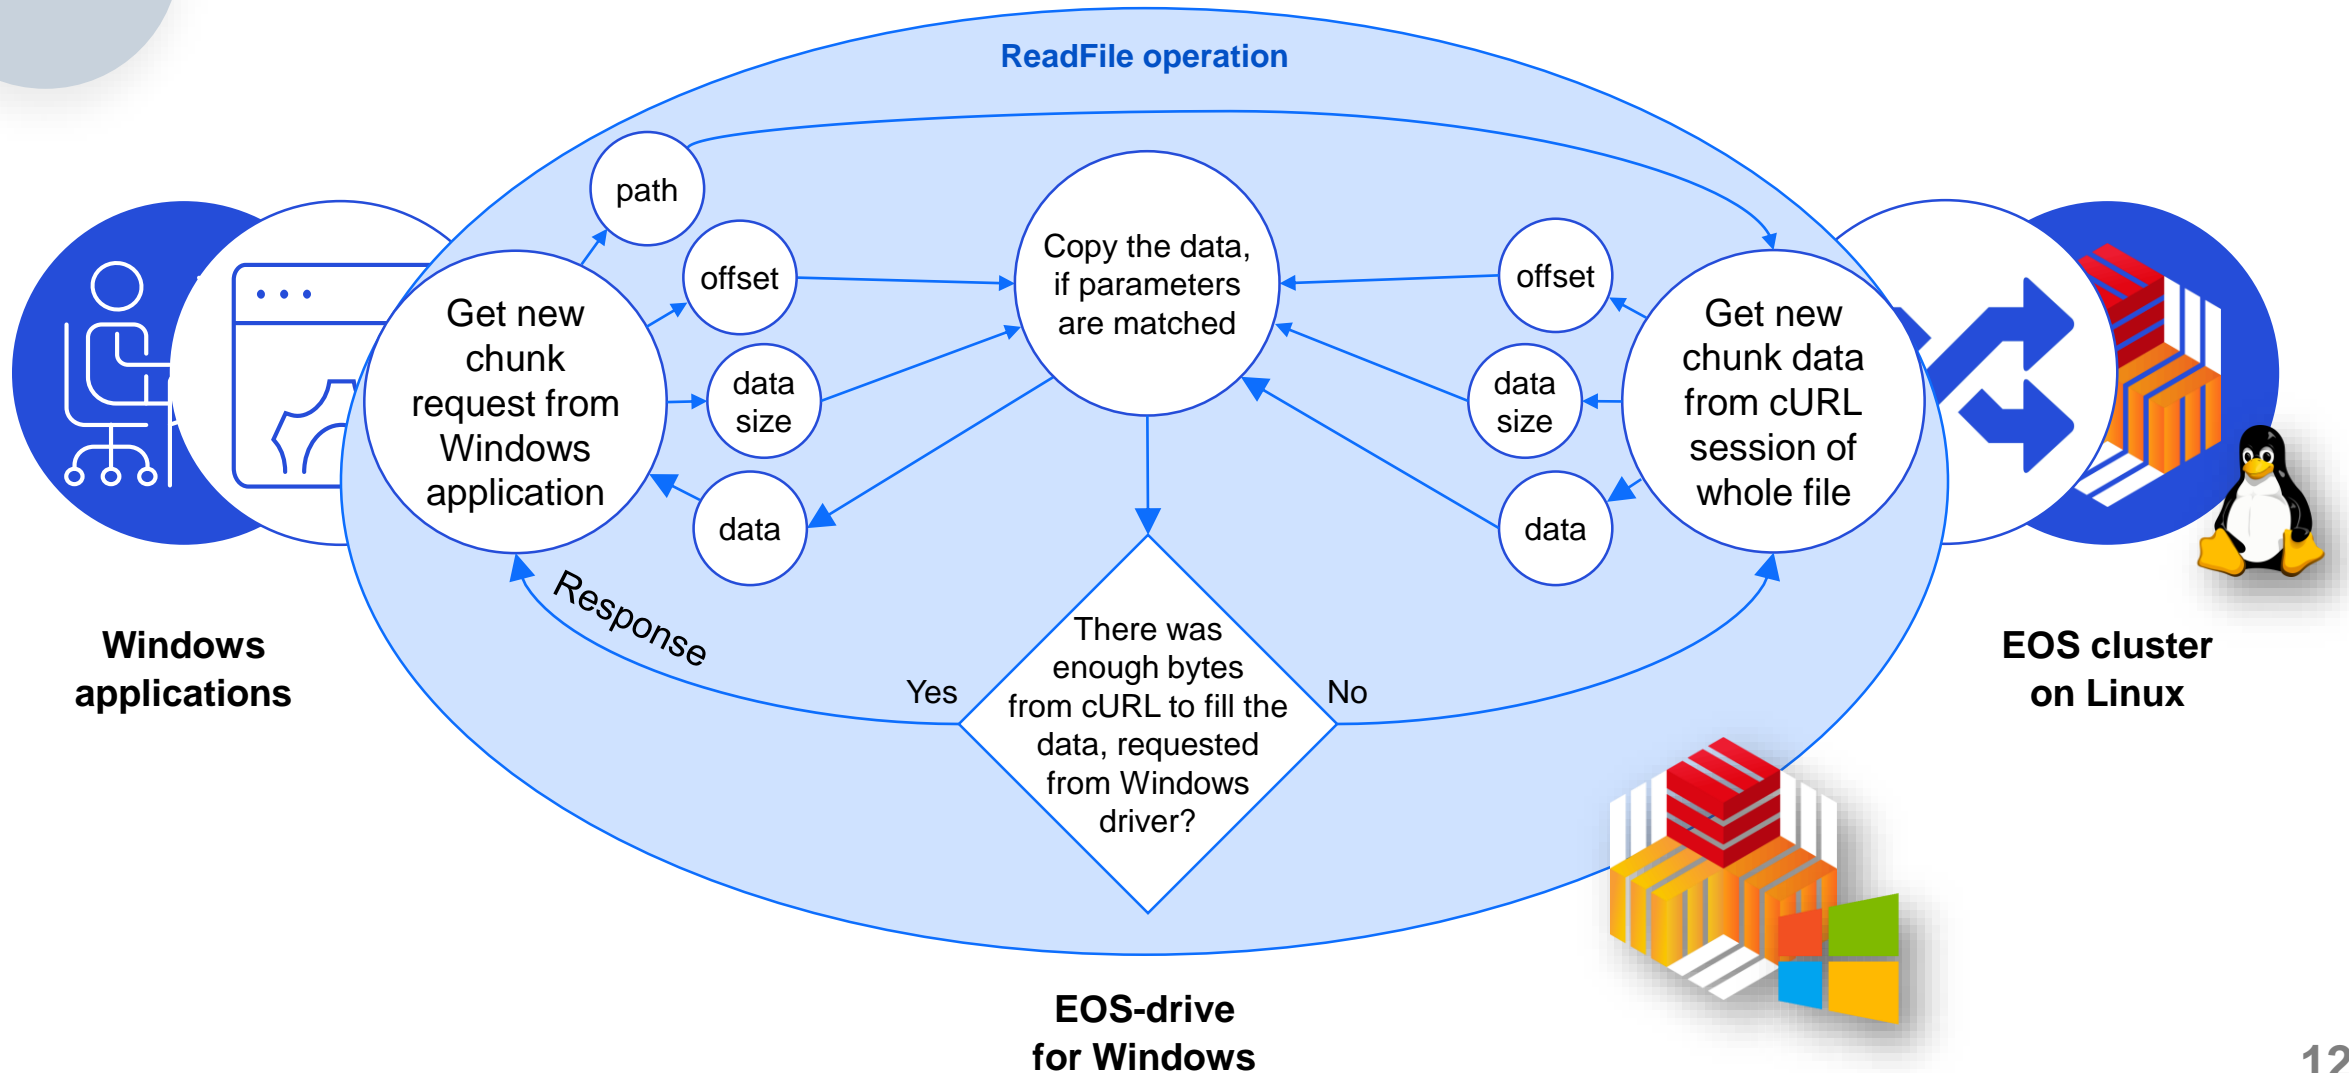
# File transferring with EOS-drive buffer for uploading



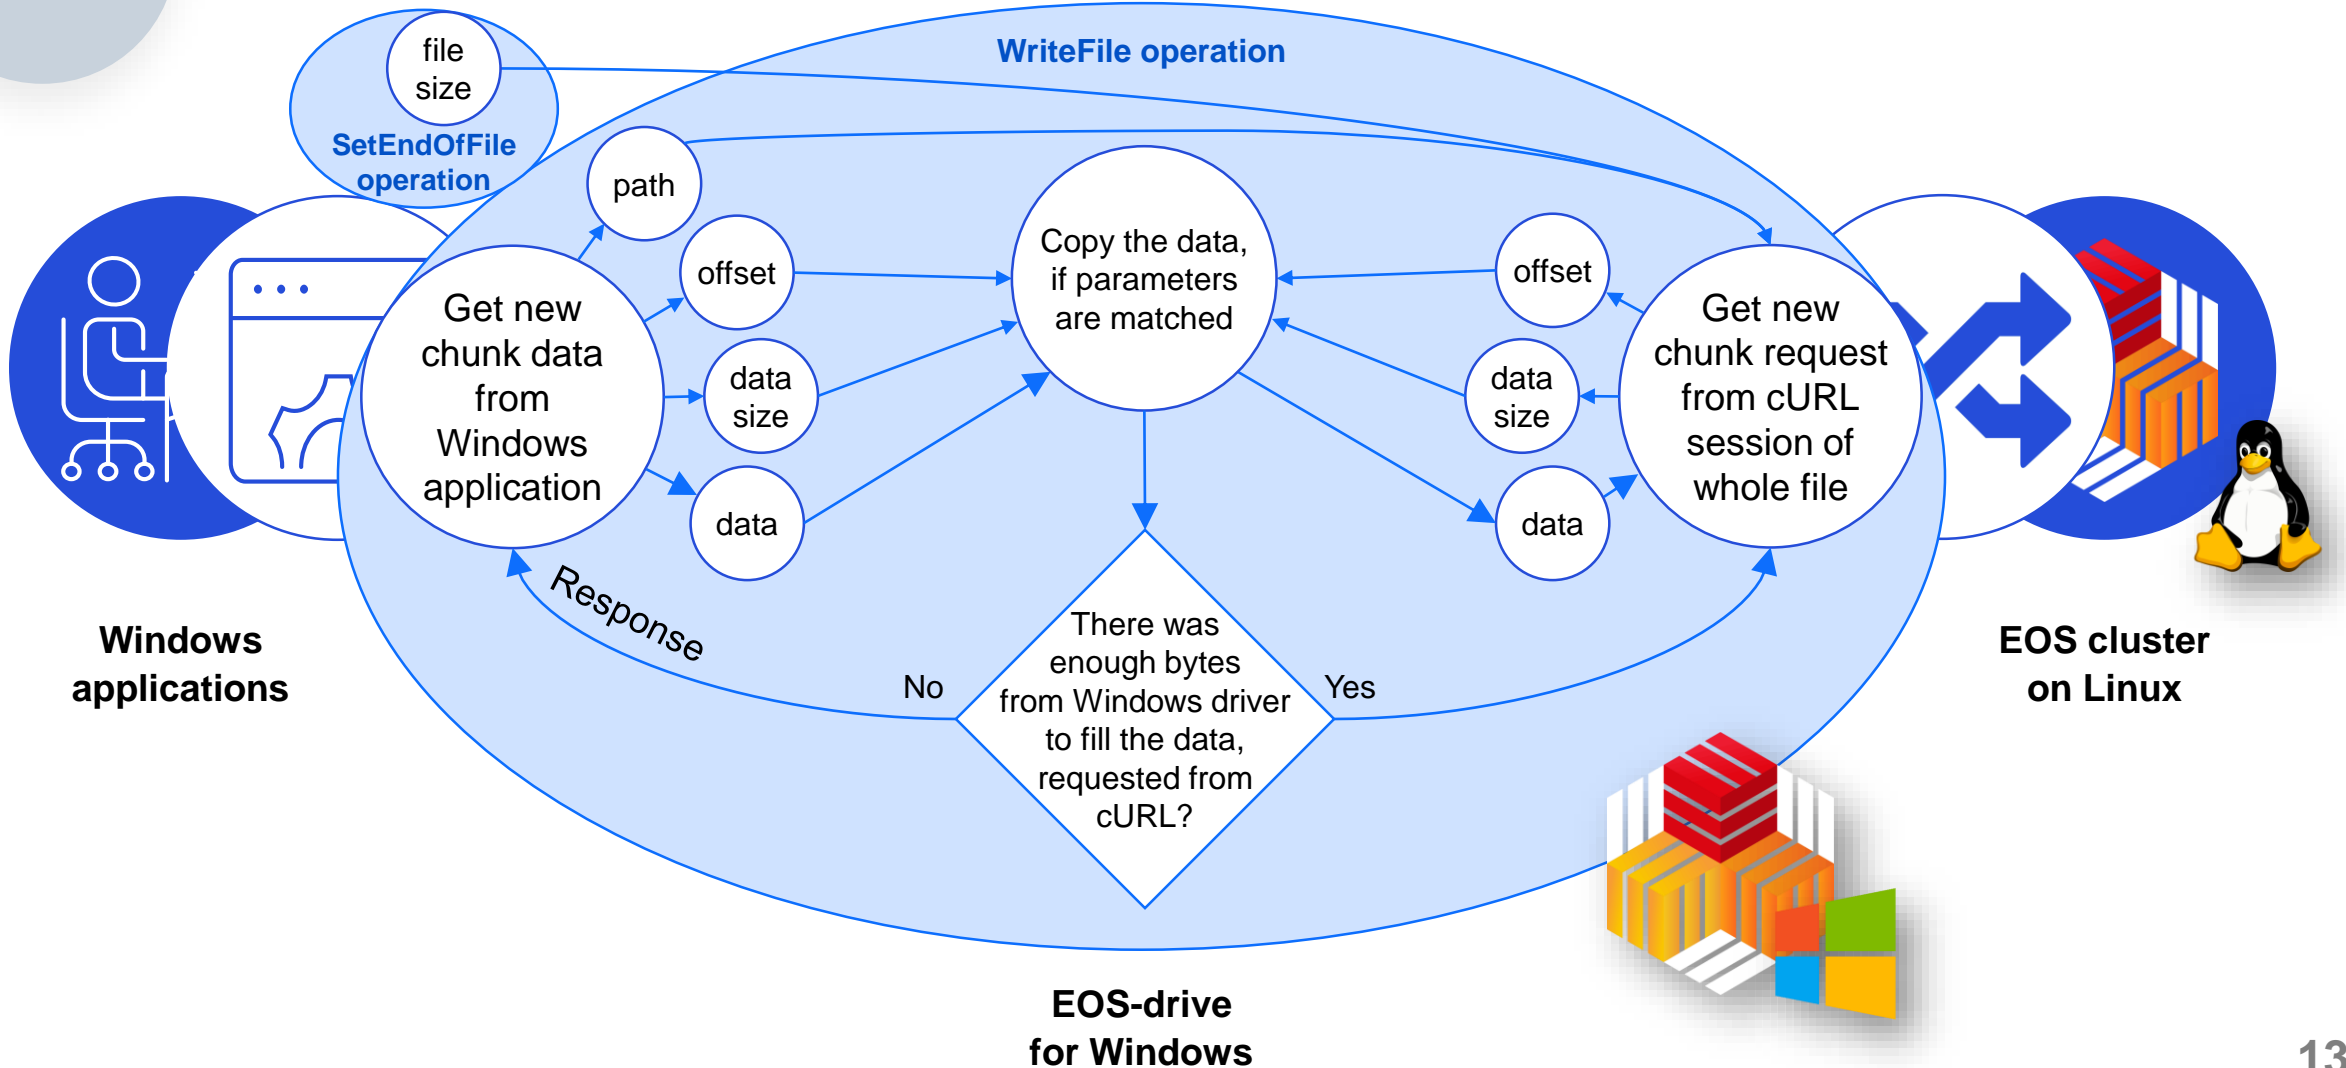
# File transferring with EOS-drive buffer for downloading



# File transferring with single-session for downloading



# File transferring with single-session for uploading





# File transferring system testing

## Testing of EOS-wnc at Comtrade:

Iterations (EOS)		390 (checksums OK)				Number of files	100
Iterations (IBM)		407 (checksums OK)					
Iterations (Ceph)		23 (checksums OK)				File size [MB]	1
Iterations (Hadoop)		14 (checksums OK)					
Test (MB/s)		min	max	avg	trim25%	Avg time (ms)	
Upload	EOS: xrdcp command	18.08	208.33	163.02	163.98	★	6.13
	EOS Fusex	21.05	52.03	25.38	23.15	★	39.40
	IBM Spectrum Scale	26.16	59.84	55.23	55.31	★	18.10
	Ceph on Linux	145.99	156.25	150.62	150.50	★	6.64
	Hadoop on Linux	3.46	3.64	3.55	3.55	★	281.92
	EOS-wnc	11.06	17.27	14.82	14.82	★	67.46
Download	EOS-drive ST	9.70	11.51	9.97	9.96	★	100.30
	EOS: Samba	7.73	24.51	9.86	8.53	★	101.38
	Ceph on Win	50.28	56.50	53.52	53.51	★	18.68
	Hadoop on Win	3.13	3.21	3.18	3.18	★	314.61
	EOS: xrdcp command	22.68	208.33	170.81	171.51	★	5.85
	EOS Fusex	33.82	53.28	37.24	36.60	★	26.85
Linux	IBM Spectrum Scale	143.47	205.76	170.60	170.64	★	5.86
	Ceph on Linux	7.47	110.13	94.19	97.40	★	10.62
	Hadoop on Linux	3.68	4.30	3.97	3.97	★	251.71
	EOS-wnc	8.55	13.83	10.59	10.58	★	94.42
	EOS-drive ST	7.08	19.03	11.74	11.43	★	85.21
	EOS: Samba	5.12	15.82	7.48	6.89	★	133.61
Windows	Ceph on Win	1.93	47.25	35.38	37.60	★	28.26
	Hadoop on Win	1.69	2.60	2.20	2.21	★	454.84

Legend:	[MB/s]
Red	0 - 10
Orange	10 - 20
Yellow	20 - 30
Light green	30 - 40
Green	40 - 50

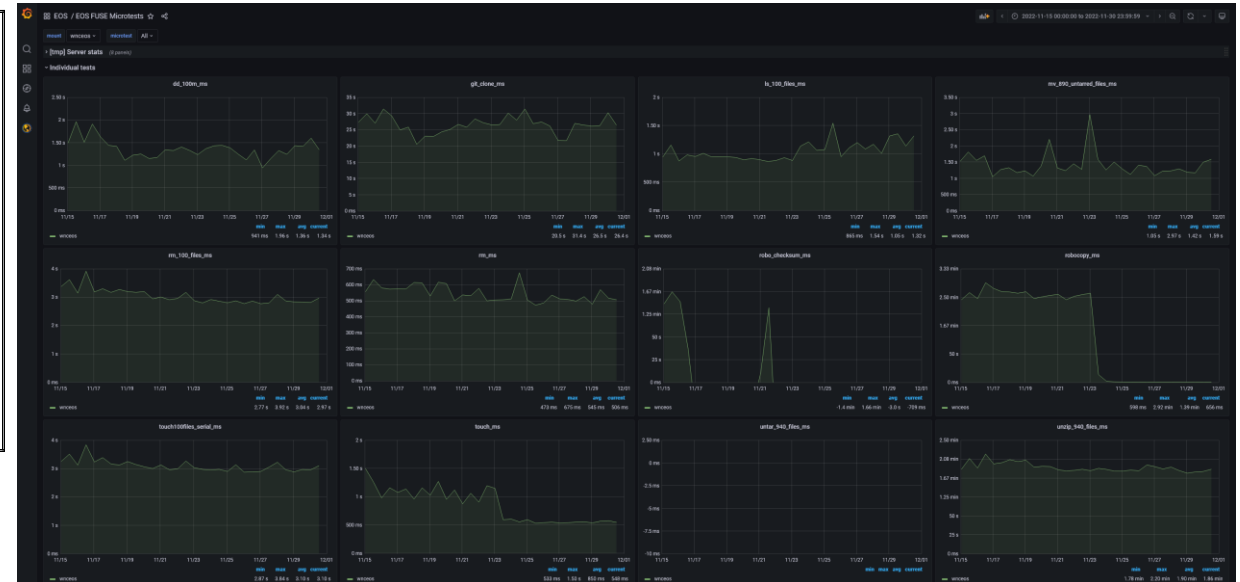
Iterations (EOS)		414 (checks OK)				Number of files	10
Iterations (IBM)		414 (checks OK)					
Iterations (Ceph)		52 (checks OK)				File size [MB]	100
Iterations (Hadoop)		11 (checks OK)					
Upload	Test [MB/s]		min	max	avg	trim25%	
	EOS: xrdcp command		172.71	452.49	410.38	412.37	★ 243.68
	EOS Fusex		127.34	192.64	153.62	153.63	★ 650.96
	IBM Spectrum Scale		153.68	188.71	177.58	177.71	★ 563.12
	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
	EOS-wnc		129.32	206.05	180.04	179.90	★ 555.43
	EOS-drive ST		160.36	250.43	189.69	189.08	★ 527.17
	EOS: Samba		142.01	245.42	178.27	176.63	★ 560.96
	Ceph on Win		102.68	141.96	136.28	137.07	★ 733.77
Hadoop on Win		4.50	5.22	4.61	4.54	★ 21670.61	
Linux	EOS: xrdcp command		155.52	442.48	383.77	390.45	★ 260.57
	EOS Fusex		104.96	227.58	151.87	148.75	★ 658.45
	IBM Spectrum Scale		291.63	375.66	325.87	325.22	★ 306.87
	Ceph on Linux		20.49	183.49	31.48	28.77	★ 3175.40
	Hadoop on Linux		8.06	10.51	9.37	9.38	★ 10668.22
	EOS-wnc		98.70	267.68	161.52	160.95	★ 619.12
Windows	EOS-drive ST		74.73	221.85	117.37	115.93	★ 852.00
	EOS: Samba		51.38	97.50	66.38	65.48	★ 1506.49
	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 - 300

Iterations (EOS)		399 (checks OK)				Number of files	2	
Iterations (IBM)		414 (checks OK)						
Iterations (Ceph)		52 (checks OK)				File size [MB]	2000	
Iterations (Hadoop)		11 (checks OK)						
Download	Test [MB/s]		min	max	avg	trim25%	Avg time [s]	
	Linux	EOS: xrdcp command	265.60	440.04	381.54	382.53	★	5.24
		EOS Fusex	164.71	237.63	201.69	201.73	★	9.92
		IBM Spectrum Scale	256.46	318.22	285.54	285.53	★	7.00
		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	150.60	332.40	230.26	227.29	★	8.69
		EOS-drive ST	177.71	367.24	223.57	222.85	★	8.95
		EOS: Samba	151.15	247.39	183.37	178.84	★	10.91
		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
	Linux	EOS: xrdcp command	242.28	378.43	342.92	344.30	★	5.83
EOS Fusex		85.17	428.91	196.16	203.19	★	10.20	
IBM Spectrum Scale		290.30	364.96	325.84	326.47	★	6.14	
Ceph on Linux		188.80	355.49	265.08	265.44	★	7.54	
Hadoop on Linux		9.28	10.63	10.12	10.15	★	197.66	
Windows	EOS-wnc	119.14	248.14	184.88	185.42	★	10.82	
	EOS-drive ST	133.97	238.34	170.04	169.85	★	11.76	
	EOS: Samba	14.72	47.71	25.40	24.98	★	78.74	
	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 - 300

## Testing of EOS-wnc at CERN:



The number of transfer sessions per file, plays a leading role in performance of EOS-drive.

# Summary

01.

**Single-session  
transfer**



**The fastest solution,**  
but it needs sorted chunk  
requests, it cannot be used  
for parallel transfer and it  
needs filesize parameter  
for uploading processes.

02.

**Transfer with the  
EOS-drive buffer**



**Middle solution per  
speed,** but it is limited per  
number of parallel files  
to save memory resources.

03.

**Transfer  
chunk by chunk**



**The slowest solution,**  
but can be realized  
in any case.

# Thank you

EOS-drive for Windows:  
Architecture and file transferring system

**Ivan Arizanović**

ivan.arizanovic@comtrade.com  
ivan.arizanovic@cern.ch

Gregor Molan | Branko Blagojević