

# Practical approach to SMB Multichannel in IBM® Storage Scale

Samba XP 2026

April 19/20 , 2026 | Göttingen, Germany

Michael Diederich

[mailto: diederich@de.ibm.com](mailto:diederich@de.ibm.com)



# SMB Multichannel

## **Introduction:**

Windows Server 2012 includes a new feature called SMB Multichannel, part of the SMB 3.0 protocol, which increases the network performance and availability for File Servers

## **Benefits:**

SMB Multichannel allows file servers to use multiple network connections simultaneously and provides the following capabilities:

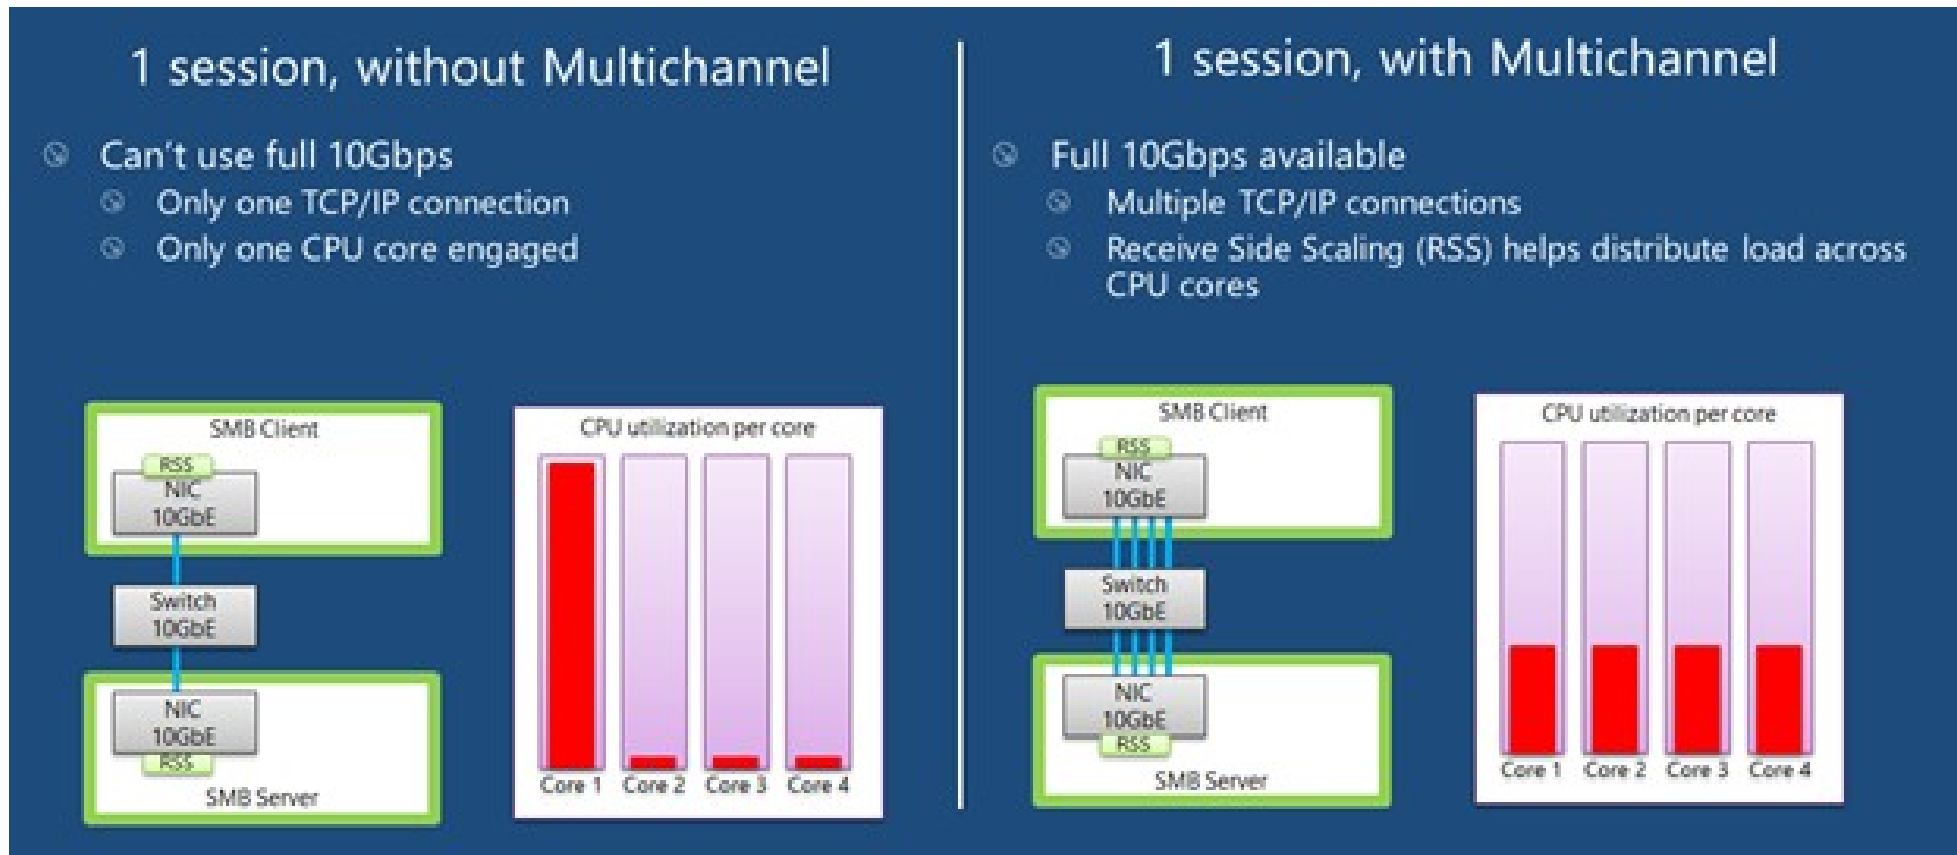
- Increased throughput:
- Network Fault Tolerance:
- Automatic Configuration:

**Requirements:** SMB Multichannel (SMB 3.0 and above)

At least one of the configurations below:

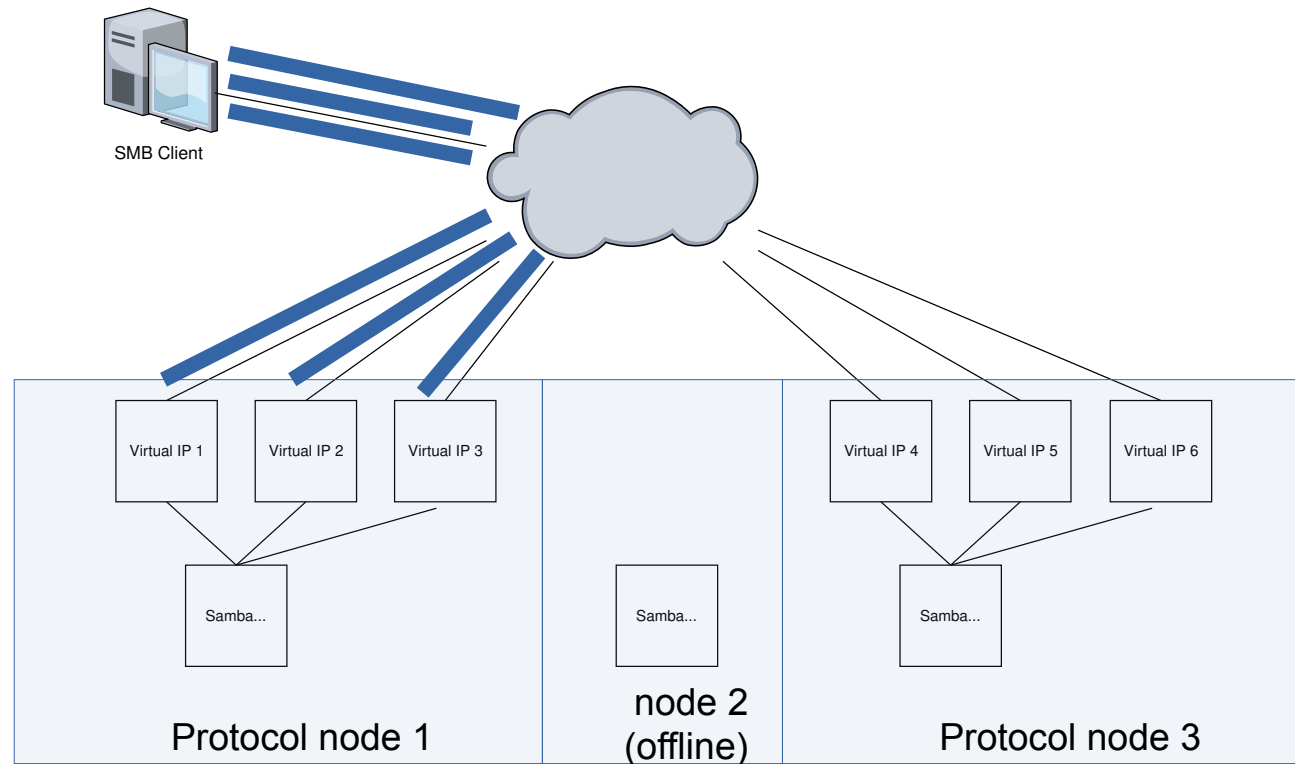
- Multiple network adapters
- One or more network adapters that support RSS (Receive Side Scaling)
- One of more network adapters configured with NIC Teaming
- One or more network adapters that support RDMA (Remote Direct Memory Access)

With SMB Multichannel, SMB will create multiple TCP/IP connections for a single session, avoiding a potential bottleneck on a single CPU core when lots IOs are required.



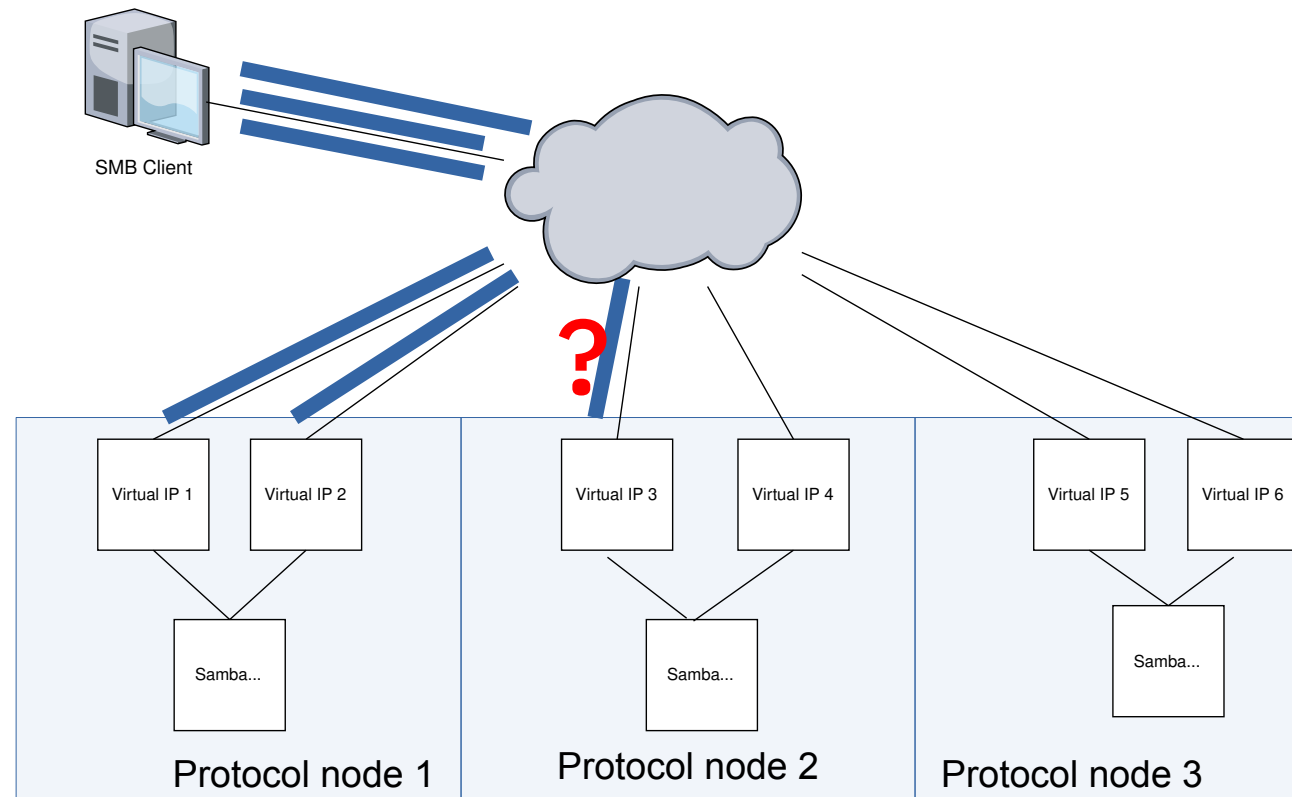
Source : <https://learn.microsoft.com/de-de/archive/blogs/josebda/the-basics-of-smb-multichannel-a-feature-of-windows-server-2012-and-smb-3-0>

Multichannel in a cluster with virtual IP addresses :  
Multiple channels (IP1, IP2, IP3) bound to same SMB3 session :

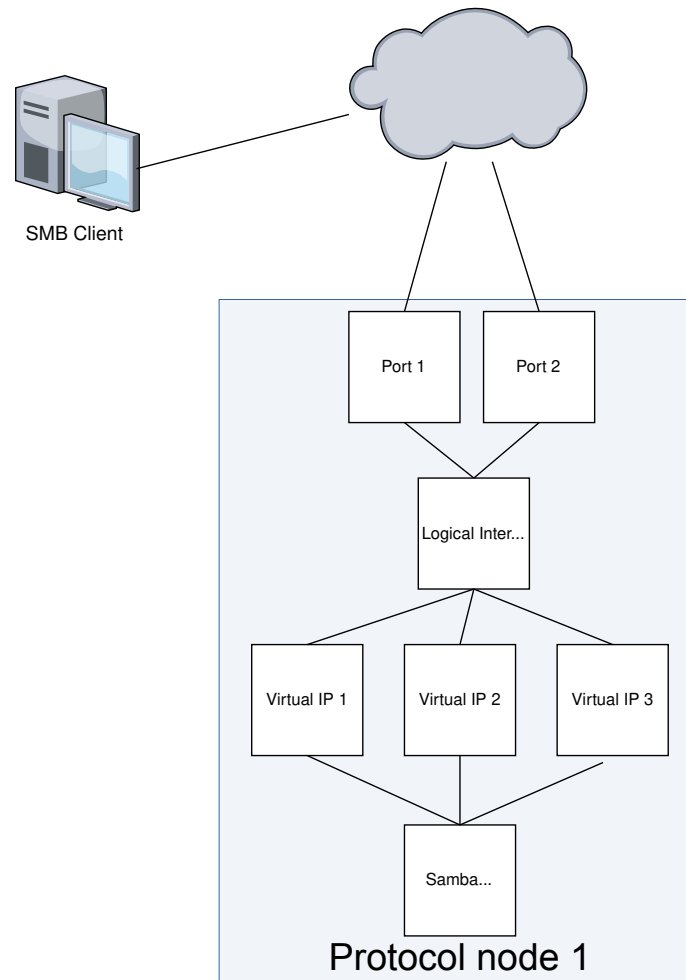


Multichannel in a cluster with virtual IP addresses :

Multiple channels (IP1, IP2) bound to same SMB3 session – IP3 has moved to other node :



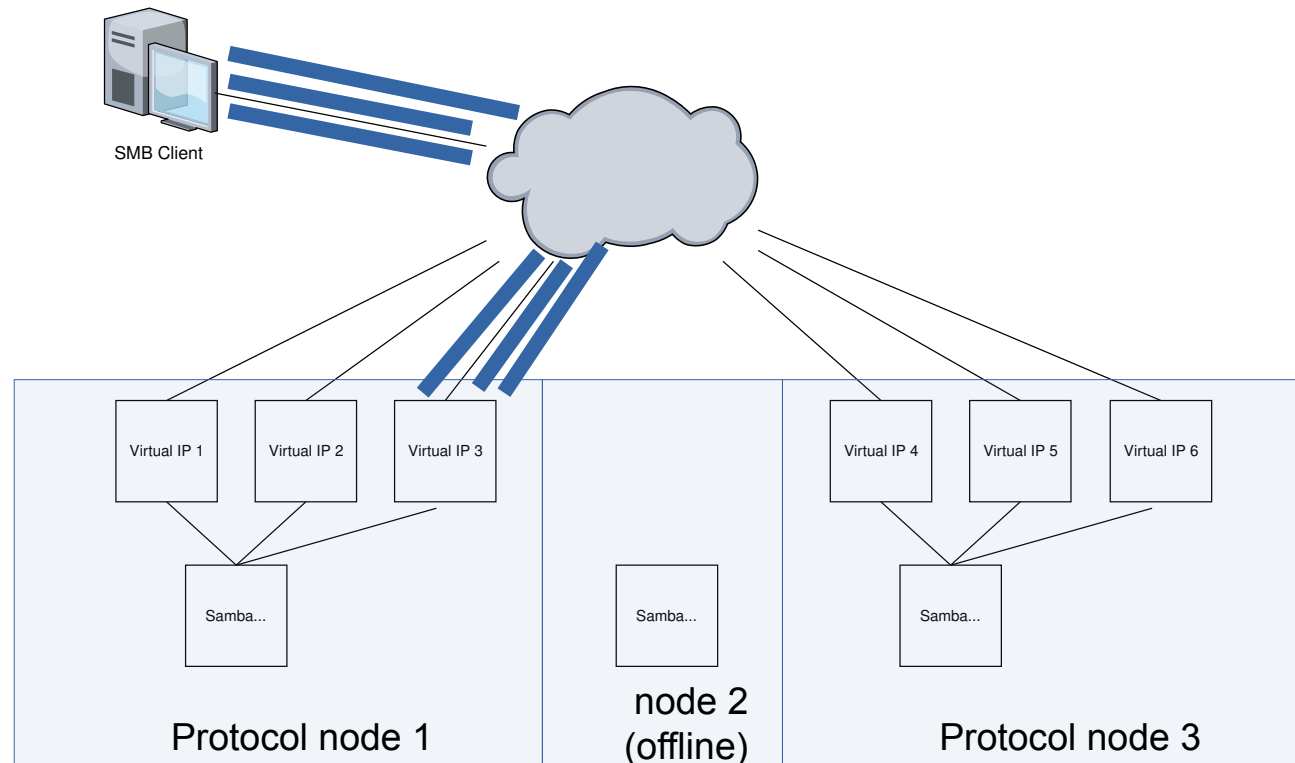
## Virtual Ips and multiple Interfaces – in reality



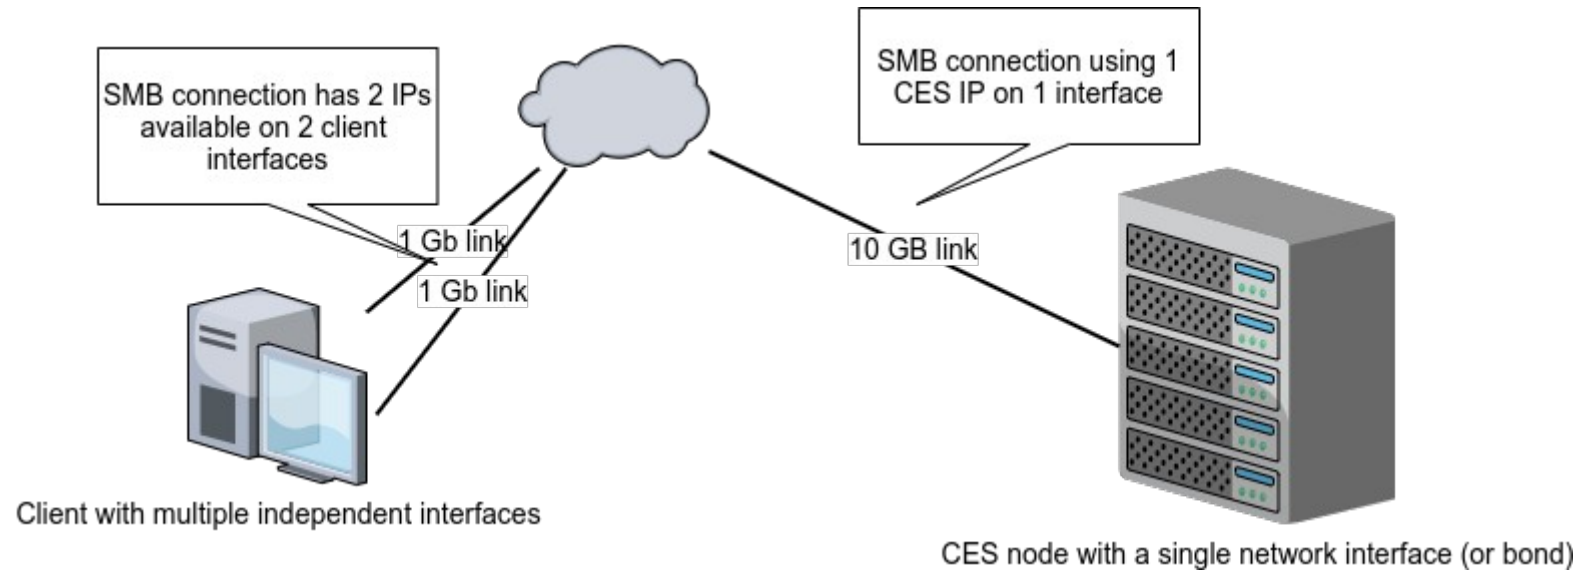
In most installs - Protocol nodes use one set of network ports for SMB IO

- Usually 2 ports are combined using the Linux “bond” mechanism for HA (could be more than 2, can include link aggregation)
- Virtual IP addresses are started on a single logical interface
- Available bandwidth is usually not a concern
- Additional interfaces connect nodes in the Samba cluster and/or Storage backend

Multichannel in a cluster with virtual IP addresses :  
Multiple channels (1, 2, 3) bound to same SMB3 session and IP :



# SMB Multichannel with multiple interfaces on client (single interface and single IP on CES node)



# Scale Protocol node typical Networks

0000:01:00.0 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]  
0000:01:00.1 Infiniband controller: Mellanox Technologies MT28908 Family [ConnectX-6]  
0001:08:00.0 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
0001:08:00.1 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
0001:08:00.2 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
0001:08:00.3 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
0033:01:00.0 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]  
0033:01:00.1 Ethernet controller: Mellanox Technologies MT28800 Family [ConnectX-5 Ex]

2\*100Gbit/s

01:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
01:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
25:00.0 Ethernet controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
25:00.1 Ethernet controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
61:00.0 Ethernet controller: Intel Corporation Ethernet Controller X550 (rev 01)  
61:00.1 Ethernet controller: Intel Corporation Ethernet Controller X550 (rev 01)  
81:00.0 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
81:00.1 Infiniband controller: Mellanox Technologies MT2910 Family [ConnectX-7]  
eth0,eth1 – 40G, ib0,1,2,3 – 200G, X550 1 G

2\*40Gbit/s

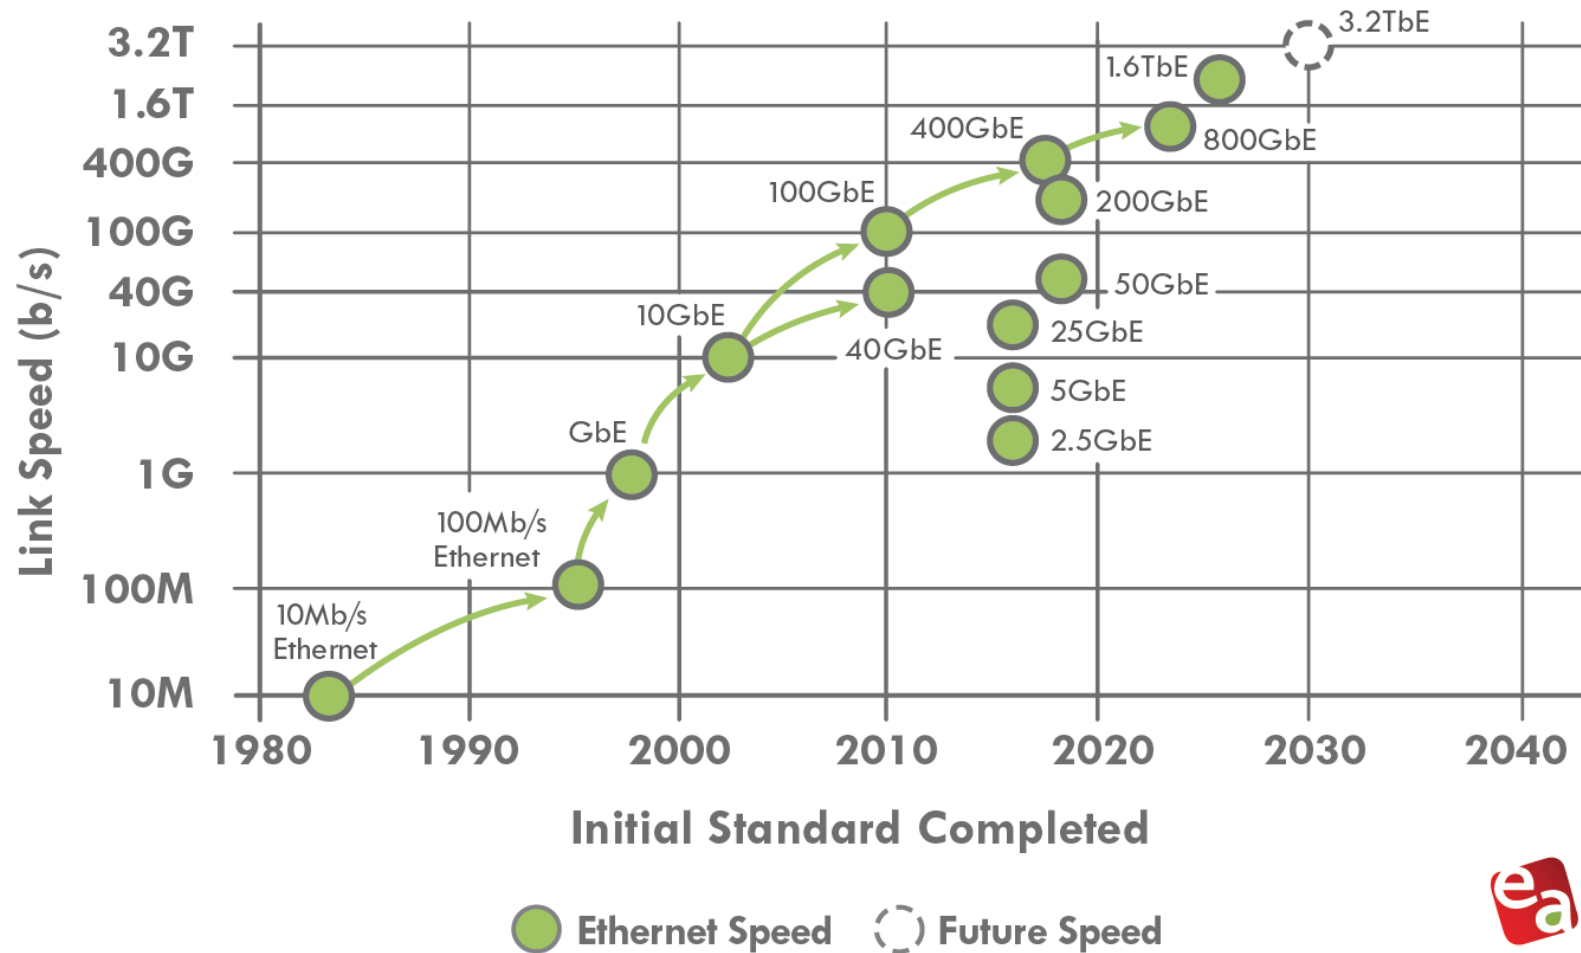
02:00.0 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
02:00.1 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
02:00.2 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
02:00.3 Ethernet controller: Broadcom Inc. and subsidiaries NetXtreme BCM5719 Gigabit Ethernet PCIe (rev 01)  
5d:00.0 Ethernet controller: Broadcom Inc. and subsidiaries BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller (rev 01)  
5d:00.1 Ethernet controller: Broadcom Inc. and subsidiaries BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller (rev 01)  
b0:00.0 Ethernet controller: Broadcom Inc. and subsidiaries BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller (rev 01)  
b0:00.1 Ethernet controller: Broadcom Inc. and subsidiaries BCM57414 NetXtreme-E 10Gb/25Gb RDMA Ethernet Controller (rev 01)  
2 bonds

2\*25Gbit/s

08:00.0 Ethernet controller: Intel Corporation Ethernet Connection X722 for 1GbE (rev 09)  
08:00.1 Ethernet controller: Intel Corporation Ethernet Connection X722 for 1GbE (rev 09)  
86:00.0 Ethernet controller: Intel Corporation Ethernet Controller XXV710 for 25GbE SFP28 (rev 02)  
86:00.1 Ethernet controller: Intel Corporation Ethernet Controller XXV710 for 25GbE SFP28 (rev 02)  
2f:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]  
2f:00.1 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]  
1 bonds 50GB aggregated + infiniband 2x100G

2\*25Gbit/s

# Ethernet roadmap



# Sample TCP capture

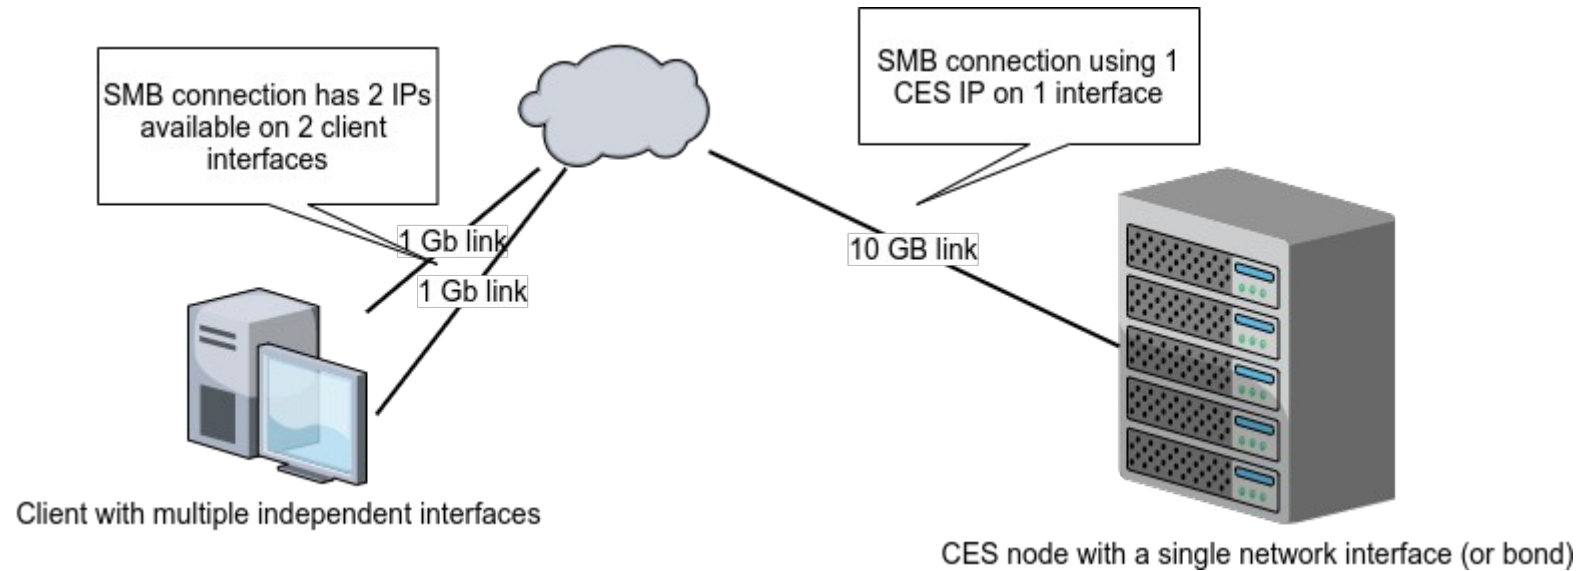
- Using network packet:

2025-08-28 23:14:14.930253	10.0.100.101	10.0.100.201	SMB2	179	Negotiate Protocol Response
2025-08-28 23:14:14.931299	10.0.100.201	10.0.100.101	SMB2	179	Session Setup Request, NTLMSSP_NEGOTIATE
2025-08-28 23:14:14.935795	10.0.100.101	10.0.100.201	SMB2	179	Session Setup Response, Error: STATUS_MORE_P
2025-08-28 23:14:14.936739	10.0.100.201	10.0.100.101	SMB2	179	Session Setup Request, NTLMSSP_AUTH, User: W
2025-08-28 23:14:14.956831	10.0.100.101	10.0.100.201	SMB2	179	Session Setup Response
2025-08-28 23:14:14.957915	10.0.100.201	10.0.100.101	SMB2	179	Tree Connect Request Tree: \\10.0.100.101\sm
2025-08-28 23:14:14.975156	10.0.100.101	10.0.100.201	SMB2	179	Tree Connect Response
2025-08-28 23:14:14.976019	10.0.100.201	10.0.100.101	SMB2	179	Ioctl Request FSCTL_QUERY_NETWORK_INTERFACE_I
2025-08-28 23:14:14.977099	10.0.100.101	10.0.100.201	SMB2	179	Ioctl Response FSCTL_QUERY_NETWORK_INTERFACE_I
2025-08-28 23:14:15.066834	10.0.100.201	10.0.100.101	SMB2	179	Create Request File: Desktop.ini
2025-08-28 23:14:15.067207	10.0.100.201	10.0.100.101	SMB2	179	Create Request File:
2025-08-28 23:14:15.074287	10.0.100.101	10.0.100.201	SMB2	179	Create Response, Error: STATUS_OBJECT_NAME_N
2025-08-28 23:14:15.075274	10.0.100.201	10.0.100.101	SMB2	179	Create Request File: AutoRun.inf
2025-08-28 23:14:15.082174	10.0.100.101	10.0.100.201	SMB2	179	Create Response File:
2025-08-28 23:14:15.084428	10.0.100.101	10.0.100.201	SMB2	179	Create Response, Error: STATUS_OBJECT_NAME_N
2025-08-28 23:14:15.170096	10.0.100.201	10.0.100.101	SMB2	179	Create Request File:
2025-08-28 23:14:15.182287	10.0.100.101	10.0.100.201	SMB2	179	Create Response File:
2025-08-28 23:14:15.185781	10.0.100.201	10.0.100.101	SMB2	179	Create Request File: ;Find Request (Level:0x5
2025-08-28 23:14:15.192063	10.0.100.101	10.0.100.201	SMB2	179	Create Response File: ;Find Response, Error:
2025-08-28 23:14:15.192504	10.0.100.201	10.0.100.101	SMB2	179	Find Request File: SMB2_FIND_ID_BOTH_DIRECT
2025-08-28 23:14:15.194840	10.0.100.101	10.0.100.201	SMB2	179	Find Response
2025-08-28 23:14:15.195221	10.0.100.201	10.0.100.101	SMB2	179	Close Request File:
2025-08-28 23:14:15.196663	10.0.100.101	10.0.100.201	SMB2	179	Close Response
2025-08-28 23:14:15.197494	10.0.100.201	10.0.100.101	SMB2	179	Create Request File: ;Notify Request
2025-08-28 23:14:15.197748	10.0.100.201	10.0.100.101	SMB2	179	Find Request File: SMB2_FIND_ID_BOTH_DIRECT
2025-08-28 23:14:15.202058	10.0.100.101	10.0.100.201	SMB2	179	Create Response File:
2025-08-28 23:14:15.202590	10.0.100.101	10.0.100.201	SMB2	179	Find Response, Error: STATUS_NO_MORE_FILES
2025-08-28 23:14:15.202665	10.0.100.101	10.0.100.201	SMB2	179	Notify Response, Error: STATUS_PENDING
2025-08-28 23:14:15.202973	10.0.100.201	10.0.100.101	SMB2	179	Create Request File: ;Notify Request
2025-08-28 23:14:15.203174	10.0.100.201	10.0.100.101	SMB2	179	Close Request File:
2025-08-28 23:14:15.207558	10.0.100.101	10.0.100.201	SMB2	179	Create Response File:
2025-08-28 23:14:15.209027	10.0.100.101	10.0.100.201	SMB2	179	Close Response
2025-08-28 23:14:15.209123	10.0.100.101	10.0.100.201	SMB2	179	Notify Response, Error: STATUS_PENDING
2025-08-28 23:14:15.251674	10.0.100.201	10.0.100.101	SMB2	179	GetInfo Request FS_INFO/FileFsVolumeInformat
2025-08-28 23:14:15.253147	10.0.100.101	10.0.100.201	SMB2	179	GetInfo Response;GetInfo Response
2025-08-28 23:14:15.775266	10.0.100.201	10.0.100.101	SMB2	179	GetInfo Request FS_INFO/FileFsFullSizeInforma

```

> SMB2 Header
v Ioctl Response (0x00b)
  > StructureSize: 0x0031
    Reserved: 0000
  > Function: FSCTL_QUERY_NETWORK_INTERFACE_INFO (0x001401fc)
  > GUID handle
    Flags: 0x00000000
    Reserved: 00000000
    Blob Offset: 0x00000070
    Blob Length: 0
    In Data: NO DATA
    Blob Offset: 0x00000070
    Blob Length: 152
  v Out Data, RSS
    v Network Interface, RSS, 4294967.0 GBits/s, IPv4: 10.0.100.101
      Next Offset: 0x00000000
      Interface Index: 2
      v Interface Capabilities: 0x00000001, RSS
        .... = RDMA: This
        .... = RSS: This
      Reserved: 0
      Link Speed: 4294967295000000, 4294967.0 GBits/s
    v Socket Address, IPv4: 10.0.100.101
      Socket Family: 2
      Socket Port: 0
      Socket IPv4: 10.0.100.101
  
```

# SMB Multichannel with multiple interfaces on client (single interface and single IP on CES node)



Windows client powershell shows the available interfaces and channels :  
Server is "scale-1" with CES IP 10.0.100.101,  
Client has 10.0.100.201 on Interface Index 10 and 10.0.100.202 on Interface Index 24  
showing output of : "Get-SmbMultichannelConnection | Format-List"

```
ServerName      : scale-1
Selected       : True
Failed         : False
FailureCount   : 0
ClientInterfaceIndex : 24
ClientRSSCapable : True
ClientRdmaCapable : False
ClientLinkSpeed : 1 Gbps
ClientIpAddress : 10.0.100.202
ServerInterfaceIndex : 2
ServerRSSCapable : True
ServerRdmaCapable : False
ServerLinkSpeed : 4294967.295 Gbps
ServerIpAddress : 10.0.100.101
MaxChannels    : 4
CurrentChannels : 0

ServerName      : scale-1
Selected       : True
Failed         : False
FailureCount   : 0
ClientInterfaceIndex : 10
ClientRSSCapable : True
ClientRdmaCapable : False
ClientLinkSpeed : 1 Gbps
ClientIpAddress : 10.0.100.201
ServerInterfaceIndex : 2
ServerRSSCapable : True
ServerRdmaCapable : False
ServerLinkSpeed : 4294967.295 Gbps
ServerIpAddress : 10.0.100.101
MaxChannels    : 4
CurrentChannels : 1
```

Before transfer of large file → current channels 0,1

After transfer of large file → current channels 4,4

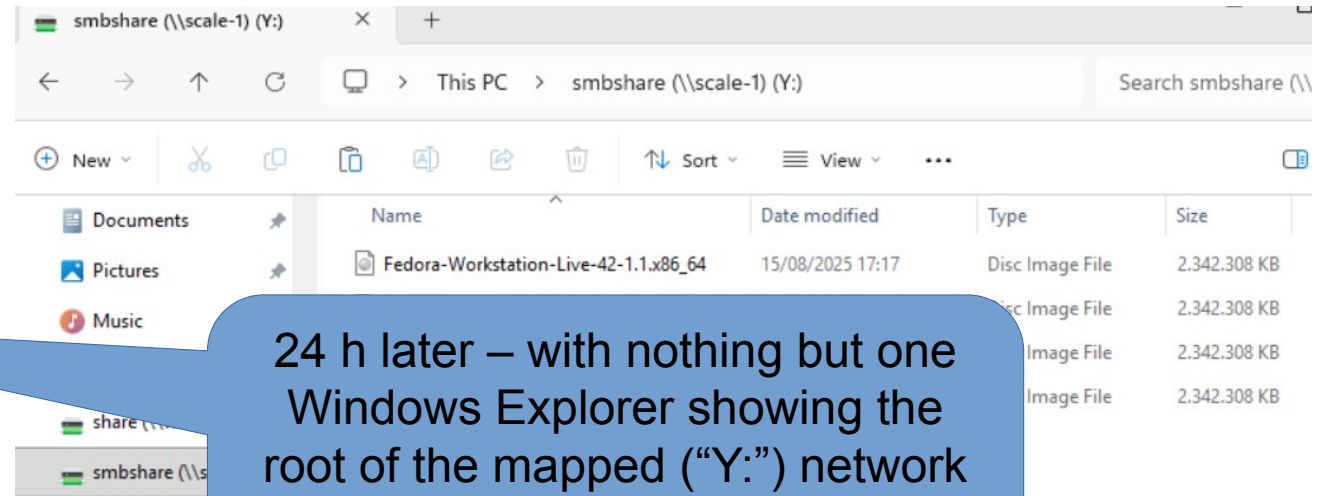
```
ServerName      : scale-1
Selected       : True
Failed         : False
FailureCount   : 0
ClientInterfaceIndex : 24
ClientRSSCapable : True
ClientRdmaCapable : False
ClientLinkSpeed : 1 Gbps
ClientIpAddress : 10.0.100.202
ServerInterfaceIndex : 2
ServerRSSCapable : True
ServerRdmaCapable : False
ServerLinkSpeed : 4294967.295 Gbps
ServerIpAddress : 10.0.100.101
MaxChannels    : 4
CurrentChannels : 4

ServerName      : scale-1
Selected       : True
Failed         : False
FailureCount   : 0
ClientInterfaceIndex : 10
ClientRSSCapable : True
ClientRdmaCapable : False
ClientLinkSpeed : 1 Gbps
ClientIpAddress : 10.0.100.201
ServerInterfaceIndex : 2
ServerRSSCapable : True
ServerRdmaCapable : False
ServerLinkSpeed : 4294967.295 Gbps
ServerIpAddress : 10.0.100.101
MaxChannels    : 4
CurrentChannels : 4
```

On CES node -  
we can see the channels in netstat / ss output

```
[root@scale-13 ~]# ss -an | grep 10.0.100.20
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.202:65090
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.202:65089
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.201:49253
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.201:65092
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.201:65091
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.202:65088
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.202:65087
tcp ESTAB 0 0 10.0.100.101:445 10.0.100.201:65093
```

4 connections from IP 10.0.100.201  
and  
4 connections from IP 10.0.100.202



24 h later – with nothing but one  
Windows Explorer showing the  
root of the mapped (“Y:”) network  
drive and no other workload – all 8  
channels are still active !

```
# smbstatus -p
```

```
Samba version 4.22.4-gpfs-5
```

```
PID Username Group Machine
```

```
0:3102325 WIN2K25\winuser WIN2K25\domain users 10.0.100.201 (ipv4:10.0.100.201:49253)
```

```
Protocol Version Encryption
```

```
SMB3_11
```

```
- partial(AES-128-GMAC)
```

```
Signing
```

# Some performance numbers

Num Threads	IO Size (KB)	No-multichannel			Multichannel = 4		
		Avg MB/s	Std Dev	CV	Avg MB/s	Std Dev	CV
1	32	970.08	41.56	4.28	942.89	51.08	5.42
1	64	1068.44	21.39	2.00	982.19	5.56	0.57
1	256	1083.42	21.74	2.01	1047.73	22.65	2.16
1	1024	1093.71	30.40	2.78	1012.56	19.86	1.96
1	2048	1188.77	35.03	2.95	1119.89	14.21	1.27
1	4096	1147.95	5.80	0.51	1127.73	21.18	1.88
1	8192	1125.91	22.24	1.98	1139.74	66.76	5.86
2	32	1451.82	4.58	0.32	1472.75	25.01	1.70
2	64	1502.50	9.31	0.62	1530.91	6.50	0.42
2	256	1470.18	59.48	4.05	1557.17	19.47	1.25
2	1024	1493.82	24.27	1.62	1524.07	19.73	1.29
2	2048	1565.87	38.71	2.47	1695.09	8.12	0.48
2	4096	1536.59	33.23	2.16	1672.61	13.91	0.83
2	8192	1582.44	12.76	0.81	1702.15	31.48	1.85
4	32	1769.82	20.43	1.15	2050.39	3.81	0.19
4	64	1823.30	7.71	0.42	2035.30	39.98	1.96
4	256	1767.44	24.79	1.40	2055.97	36.64	1.78
4	1024	1794.66	51.62	2.88	2033.41	51.09	2.51
4	2048	1832.12	32.14	1.75	2125.26	3.77	0.18
4	4096	1889.57	20.57	1.09	2092.30	5.34	0.26
4	8192	1837.50	50.98	2.77	2141.98	21.51	1.00
8	32	1989.83	11.37	0.57	2144.92	3.60	0.17
8	64	1981.09	17.77	0.90	2149.37	4.17	0.19
8	256	1997.33	4.12	0.21	2133.32	6.98	0.33
8	1024	1951.02	9.05	0.46	2070.06	8.48	0.41
8	2048	1988.85	9.61	0.48	2119.74	14.08	0.66
8	4096	2010.72	10.24	0.51	2111.37	2.12	0.10
8	8192	2016.45	6.68	0.33	2105.21	1.22	0.06

Reads

# Some performance numbers

**Writes**

Num Threads	IO Size (KB)	No-multichannel			Multichannel = 4		
		Avg MB/s	Std Dev	CV	Avg MB/s	Std Dev	CV
1	32	453.50	5.74	1.27	547.19	20.77	3.80
1	64	476.08	10.66	2.24	555.89	21.18	3.81
1	256	520.37	32.14	6.18	532.20	4.79	0.90
1	1024	552.11	10.51	1.90	539.97	12.56	2.33
1	2048	535.76	5.21	0.97	528.79	3.23	0.61
1	4096	539.92	3.37	0.62	533.12	0.79	0.15
1	8192	541.28	13.45	2.48	521.15	27.17	5.21
2	32	995.86	22.71	2.28	977.55	45.91	4.70
2	64	992.44	28.28	2.85	951.34	33.26	3.50
2	256	989.21	24.54	2.48	975.43	24.59	2.52
2	1024	938.12	16.24	1.73	1004.51	8.51	0.85
2	2048	959.98	22.21	2.31	971.92	9.63	0.99
2	4096	978.91	31.07	3.17	992.20	37.43	3.77
2	8192	957.00	6.18	0.65	960.69	20.27	2.11
4	32	1698.11	30.75	1.81	1753.71	34.17	1.95
4	64	1750.42	49.29	2.82	1733.50	7.25	0.42
4	256	1820.37	41.57	2.28	1778.91	27.76	1.56
4	1024	1963.40	38.77	1.97	1796.44	58.89	3.28
4	2048	1939.98	38.61	1.99	1739.58	27.50	1.58
4	4096	1917.56	35.34	1.84	1695.35	12.76	0.75
4	8192	1903.66	29.01	1.52	1696.06	26.27	1.55
8	32	2222.25	58.42	2.63	2105.55	27.03	1.28
8	64	2188.05	18.57	0.85	2097.74	9.64	0.46
8	256	2127.88	16.71	0.79	2076.59	7.25	0.35
8	1024	2148.77	24.73	1.15	2123.76	32.74	1.54
8	2048	2169.01	2.92	0.13	2122.40	13.29	0.63
8	4096	2238.97	14.06	0.63	2103.80	9.18	0.44
8	8192	2186.33	28.10	1.29	2145.46	13.93	0.65

# IBM Storage Scale Developer Edition

Learn, develop, and test with IBM Storage Scale for free for non-production use. Create a single global namespace of high-performance or tiered data using new or current storage up to 12TB with no time limits.

<https://www.ibm.com/account/reg/us-en/signup?formid=urx-41728>

Michael Diederich

[mailto: diederich@de.ibm.com](mailto:diederich@de.ibm.com)