IBM Spectrum Scale Strategy Days 2021

March 4-5, 2020
Online Event

Improving Spectrum Scale performance using RDMA

Ralph Würthner, Olaf Weiser, Ingo Meents
IBM Germany Research & Development GmbH
Kelsterbach

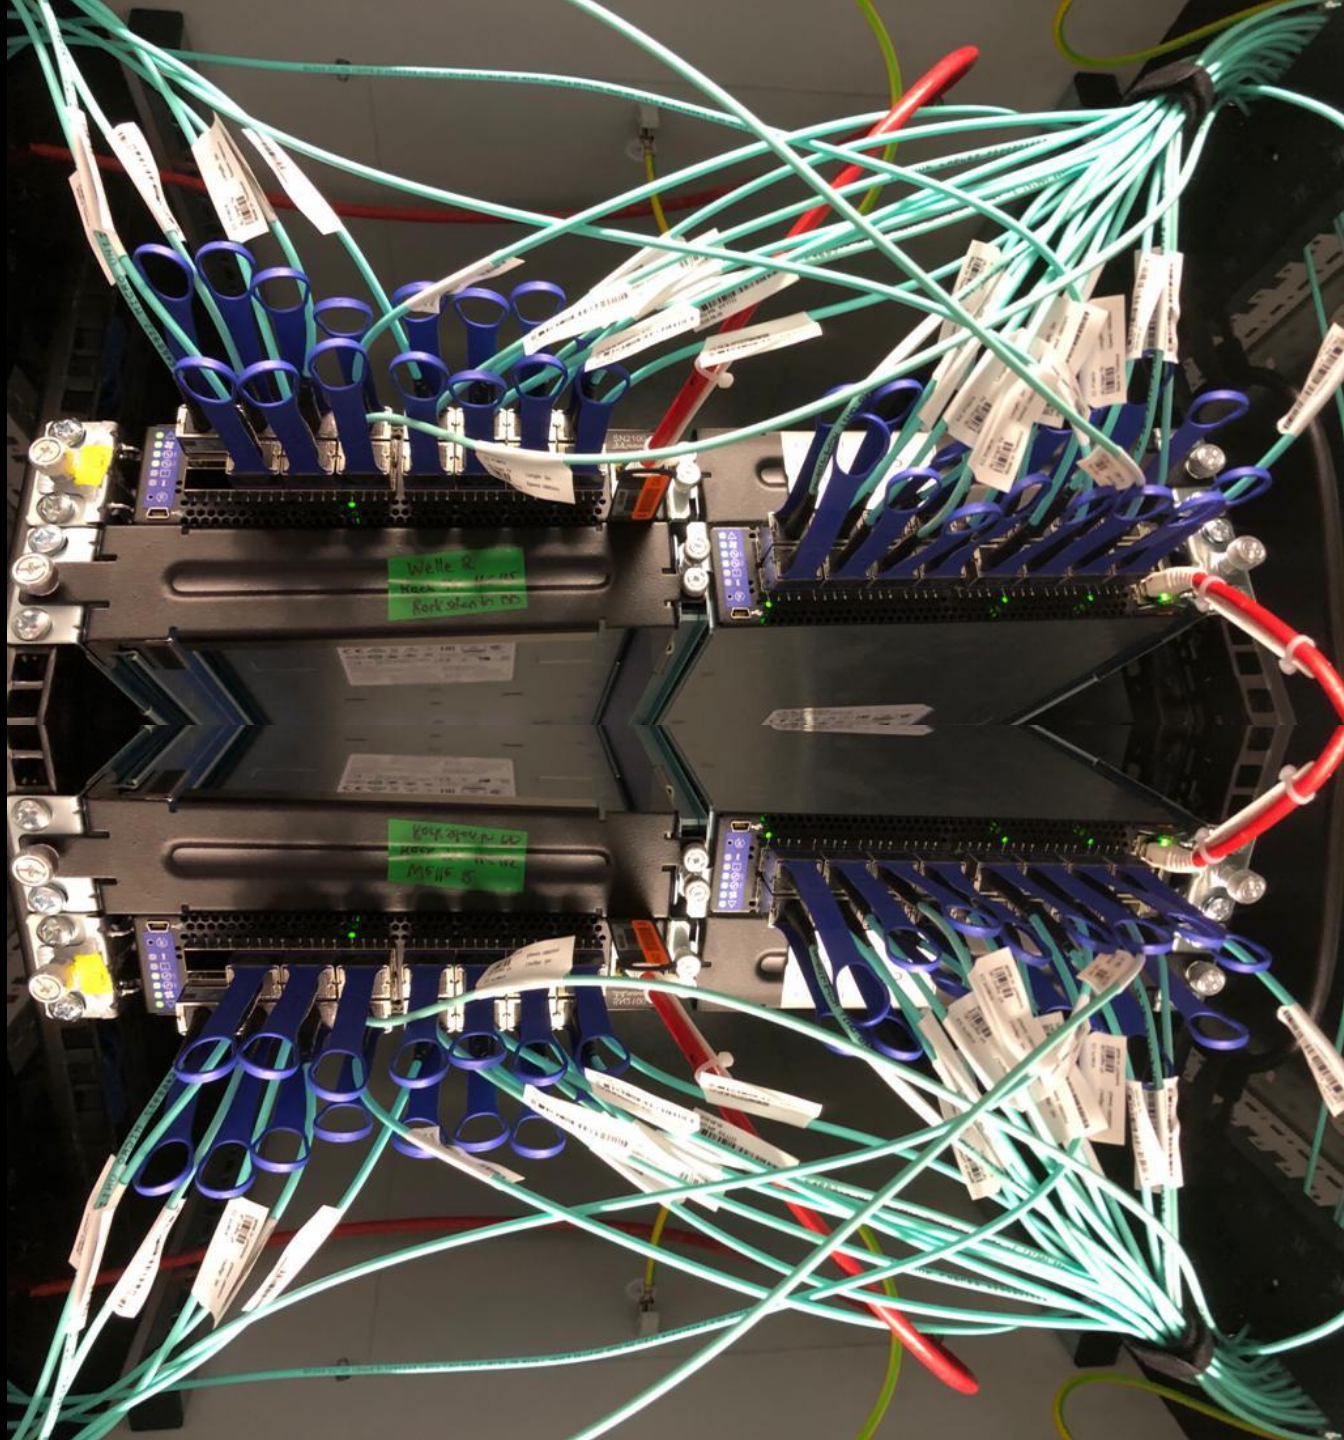

RDMA Agenda

Spectrum Scale
Strategy Days 2021

03. - 04. März 2021

Improving Spectrum Scale performance using RDMA

Part1
Introduction -
Why and how?

Ingo Meents
IT Architect
Kelsterbach, Germany

Part 2
Options Deep Dive -
How to configure?

Ralph Würthner
IT Specialist
Kelsterbach, Germany

Part 3
Performance results –
What does it gain?

Olaf Weiser
IT Specialist
Ehningen, Germany

Why RDMA ?

Computational Power rises
faster CPUs, more Cores, GPUs

Storage gets faster
HDD, SSD, NVMe, Storage class memory

Data-driven use cases keep rising
Big Data, AI workloads, Cognitive computing, Autonomous driving, etc.

The interconnect between storage and computation must not become a bottleneck.

- Move data directly, avoid data copying and context switches
- Save CPU cycles, let DMA engines do the work, bypass kernel
- Convergence: Single fabric to support Storage and Compute

- **RDMA = Remote Direct Memory Access**

- Goal: High performance networking
  - Low latency
  - High throughput

Traditional TCPIP vs RDMA Stack

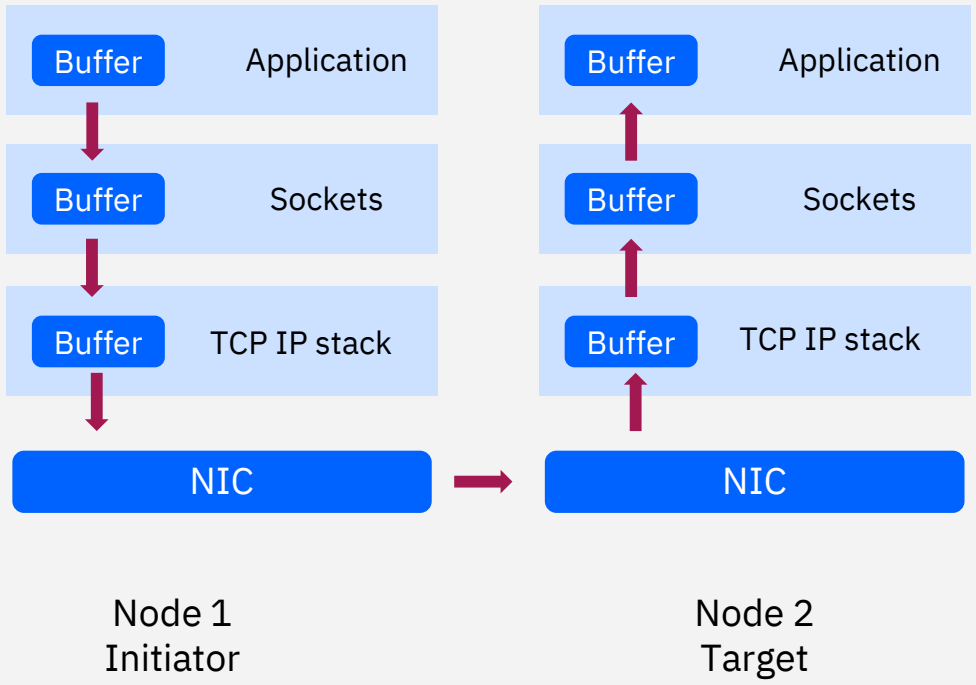

Data Flow without RDMA
context switches, data copy, blocking threads

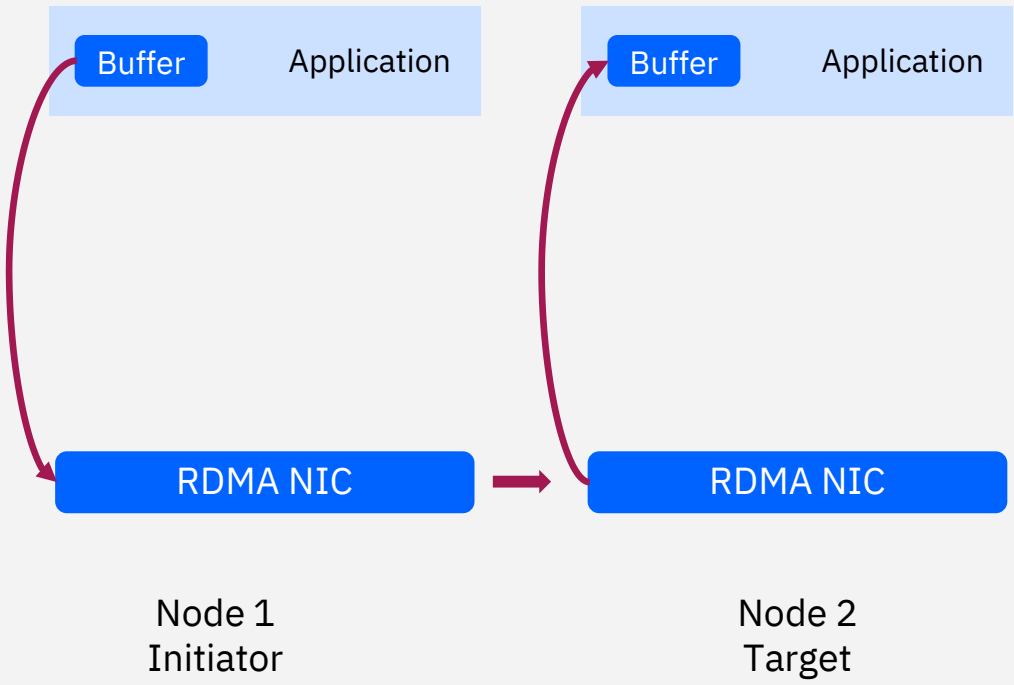

Data Flow with RDMA
OS is bypassed, network stack offload,
Direct placement of data, user thread not blocked

RDMA - What does it gain?

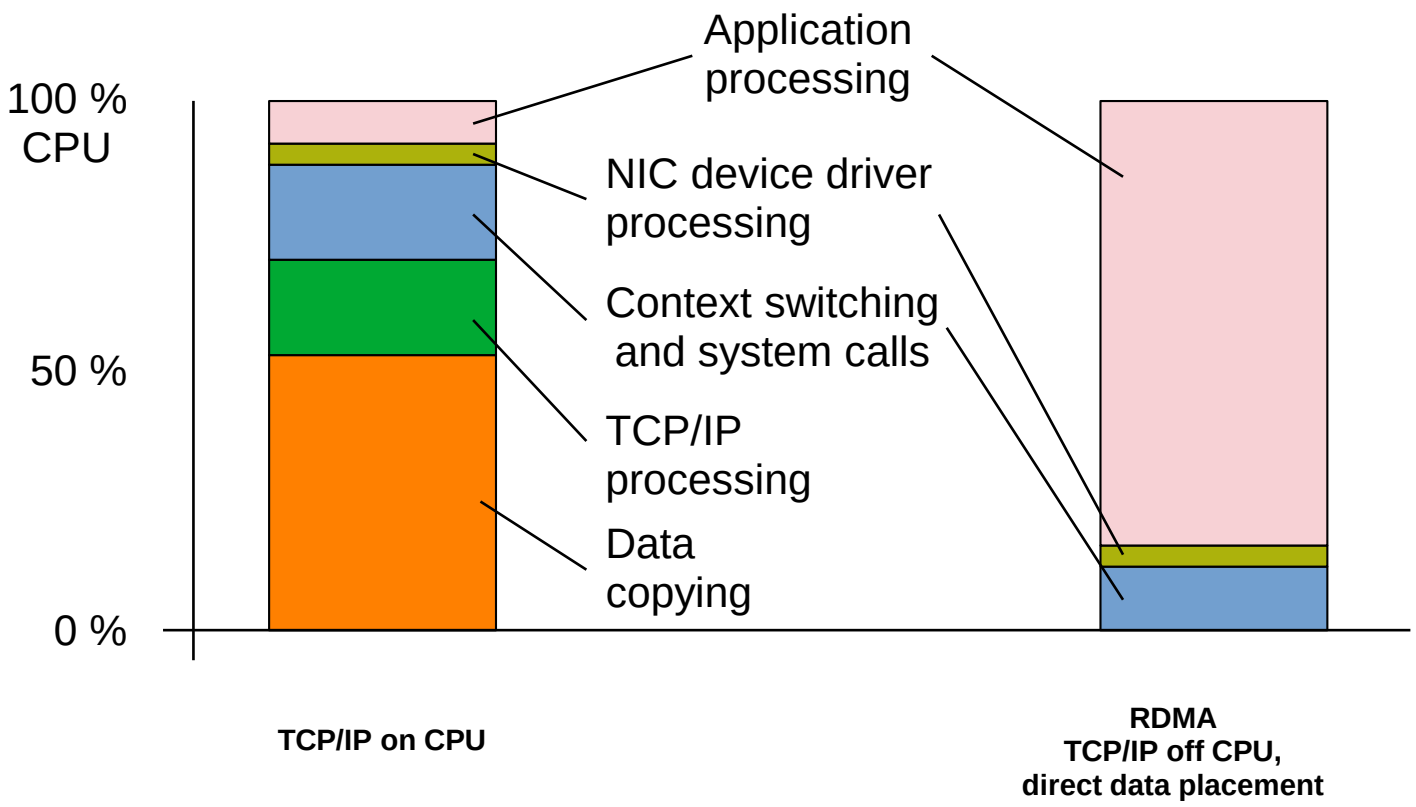

Sample data, may vary by use case

RDMA Support in Spectrum Scale

Features

- Support of both fabrics, Infiniband and RoCE (*)
- Introduced to support bulk-data (verbsRdma)
- Can also send daemon to daemon RPCs via RDMA (verbsRdmaSend)
- RDMA is from/to pagepool between gpfs cluster nodes
- TCP/IP is always required for daemon network (may be IPoIB)
- Robustness:
  - Fallback possible from RDMA to TCP/IP possible
  - detection of hung RDMA requests
- Health monitoring
- Performance counters (also in GUI)
- Speeds supported: Ethernet: up to 200 Gbps, IB: up to 4 links HDR (200Gbps)

Pre-requisites

- Linux on ppc64(le), x86_64 (no AIX or Windows)
- Mellanox HCAs
- cables for high-speed networking
- Fabric with switches
- RDMA software stack: OFED or MOFED

(*) Omnipath also still in support list

RDMA – HW Prerequisites

NIC

- Network Cards (HCA, host channel adapters)
- Number of ports (1 or 2)
- QSFP (Quad Small Form-Factor Pluggable)
  - e.g. QSFP28: 100Gbps, QSFP56: 200Gbps, QSFP-DD:400 Gbps
- PCIe3.0 or PCIe4.0
- May be configurable between IB and Ethernet

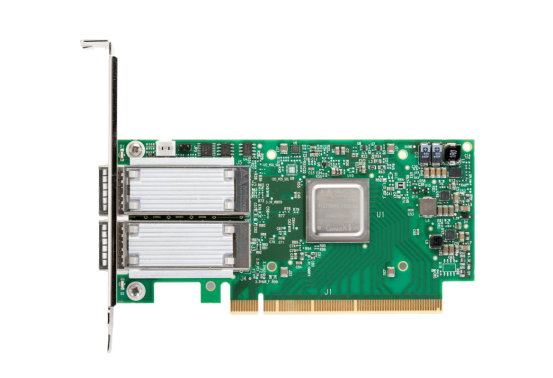

Cable

- Optical vs copper (speed, length, $$, flexibility)
- Transceiver with multiple channels
  - Eg. QSFP28 can have 4 channels with 28Gbps each
- Can do 100Gbps Ethernet as well as IB HDR Infiniband
- Diff. speeds: 25, 50, 100, 200 GbE/HDR

Switch

- Dedicated IB or Ethernet
- Number of ports (1 up to 800)
- Max speed, connector types
- Managed / unmanaged

Images: https://www.mellanox.com

RDMA – SW Prerequisites

- RDMA stack
  - Provides a.o.
    - libibverbs (user-space part of RDMA interface, rdma-core)
    - IB resource allocation (kernel space modules, driver/infiniband)
    - Vendor specific libraries and kernel modules
  - Distributions
    - Ofed (Open Fabric Enterprise Distribution) → not supported
    - Linux distro
      - E.g. packages like rdma, libibverbs, opensm + drivers
      - Optionally: librdmacm, librdmacm-utils, ibacm, libibverbs-utils, infiniband-diags, ibutils, perftest, qperf
    - Mellanox (MOFED)
      - Mellanox implementation of OFED libs and kernel modules
      - Has optimizations for mlx HW (mlx4 and mlx5 providers)
      - https://www.mellanox.com/products/infiniband-drivers/linux/mlnx_ofed

Which stack to use?

Spectrum Scale With **IB**

Spectrum Scale With **IB**

Spectrum Scale With **RoCE**

RDMA Agenda

Spectrum Scale
Strategy Days 2021

03. - 04. März 2021

Improving Spectrum Scale performance using RDMA

Part1 Introduction	Ingo Meents IT Architect Kelsterbach, Germany
Part 2 Deep Dive	Ralph Würthner IT Specialist Kelsterbach, Germany
Part 3 Performance results	Olaf Weiser IT Specialist Ehningen, Germany

GPFS RDMA Configuration (Basic settings)

- verbsRdma = [yes|no]
  - Enable general RDMA support in Spectrum Scale
  - Default is disabled
- verbsPorts = *<device>/<port>/<fabric>/<service level>*
  - List of RDMA ports to be used
  - *<device>* : RDMA device, required, e.g. mlx5_0, mlx5_1
  - *<port>*: RDMA port on device, default 1, valid values are 1 or 2
  - *<fabric>*: virtual fabric number, default 0, valid values are >= 0
    - Only verbsPorts using a common *<fabric>* are connected
  - *<service level>*: QoS level, default is 0, valid values defined in SM configuration
- verbsRdmaSend = [yes|no]
  - Enable RDMA Send for sending and receiving RPC and using inline sends of small write I/O
  - Default is disabled
  - verbsRdma must be enable as well

Spectrum Scale I/O Model – No RDMA (verbsRdma = no)

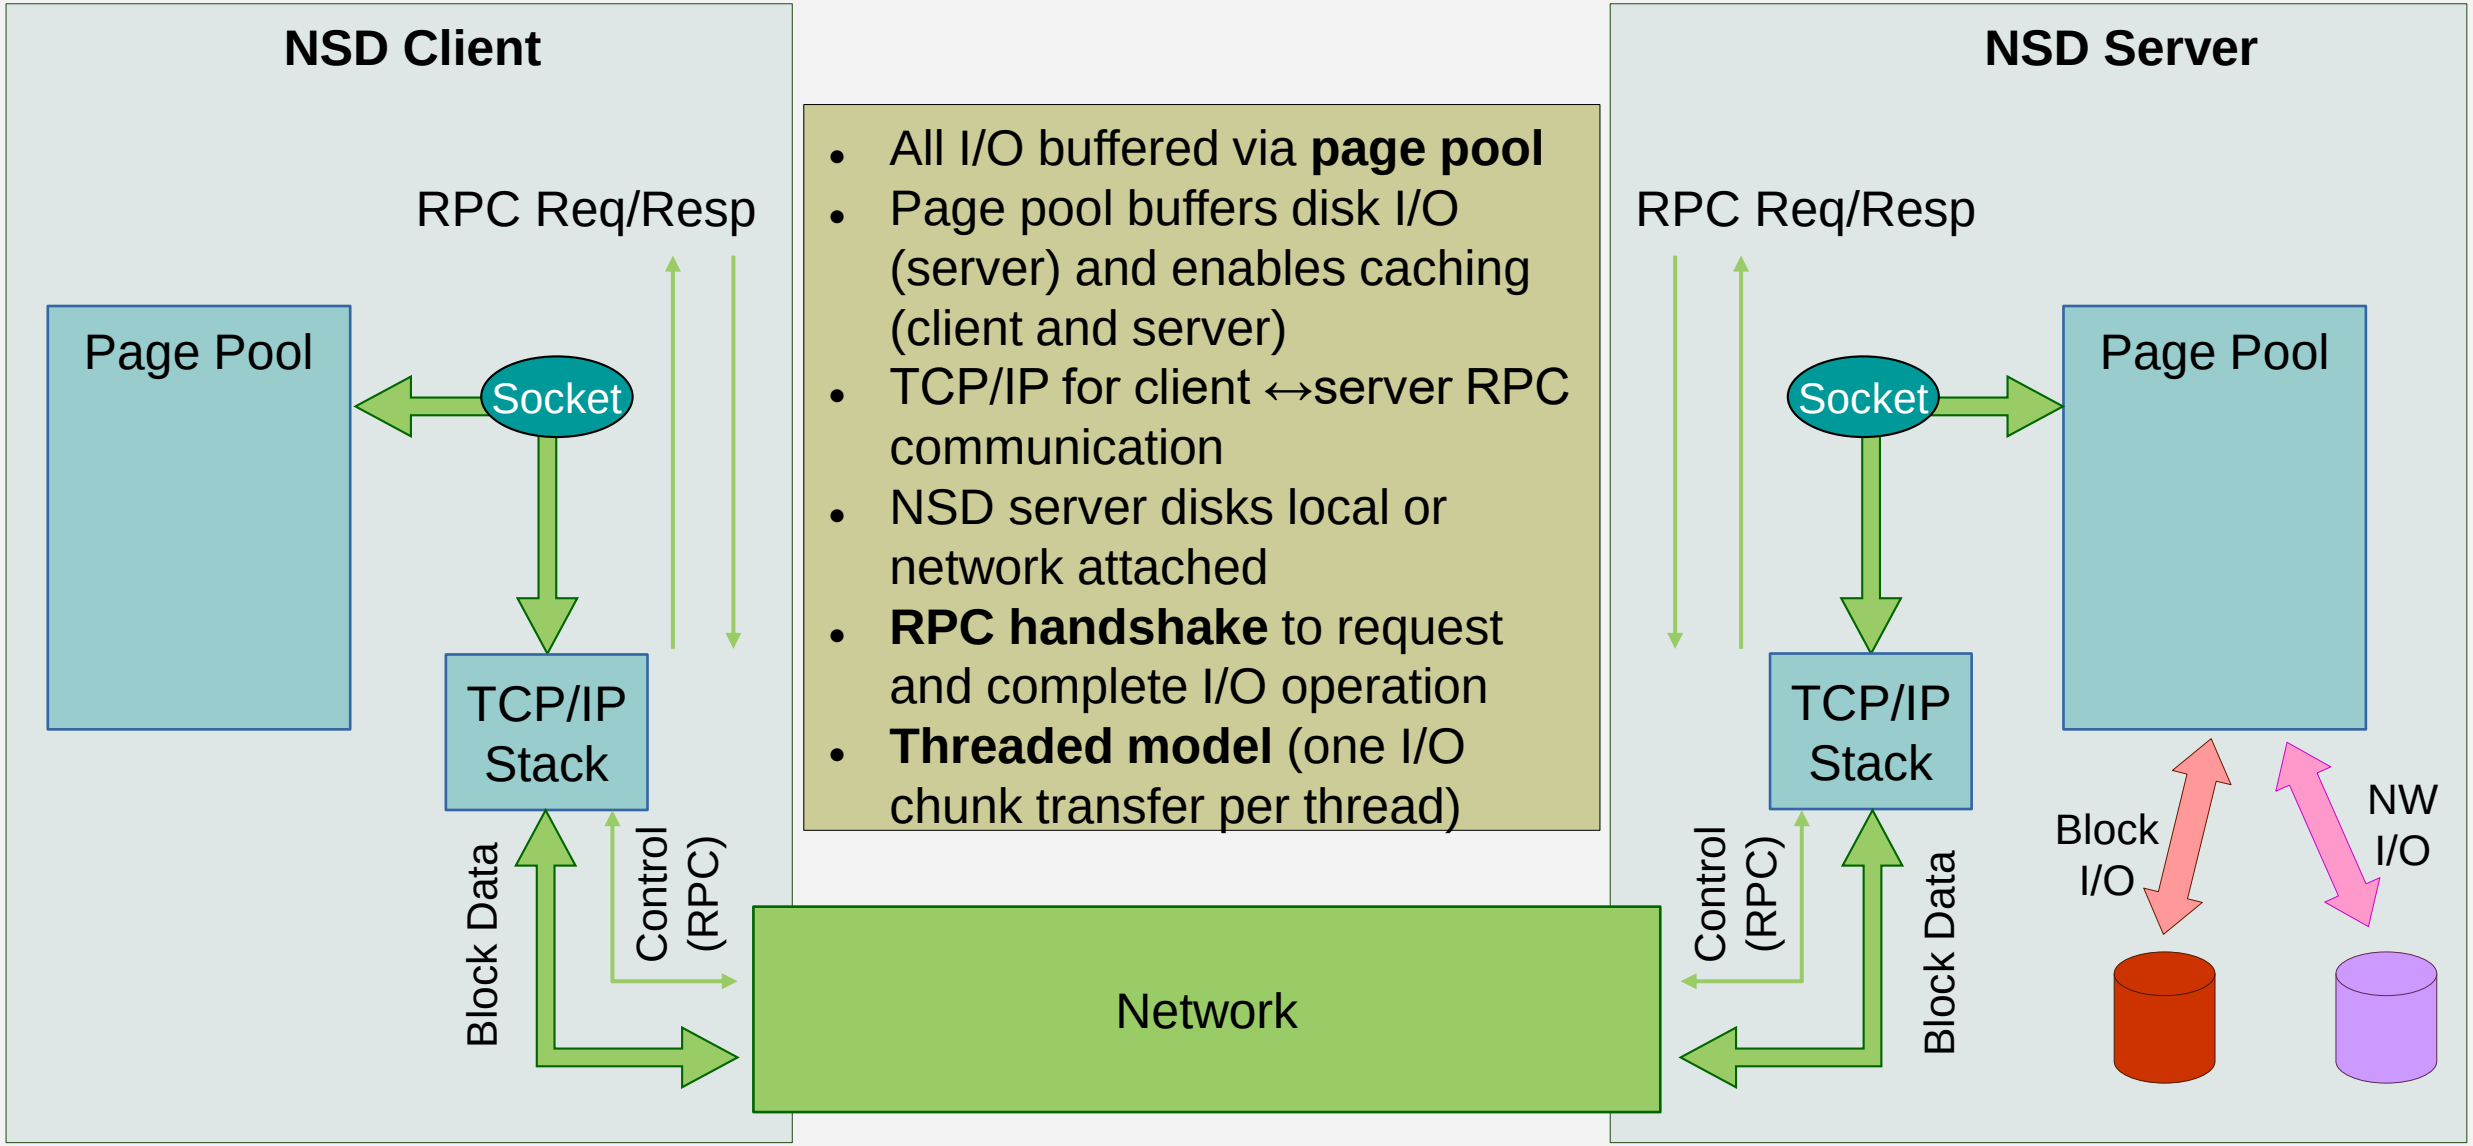

- All I/O buffered via **page pool**
- Page pool buffers disk I/O (server) and enables caching (client and server)
- TCP/IP for client ↔server RPC communication
- NSD server disks local or network attached
- **RPC handshake** to request and complete I/O operation
- **Threaded model** (one I/O chunk transfer per thread)

Spectrum Scale I/O Model – No RDMA (verbsRdma = no)

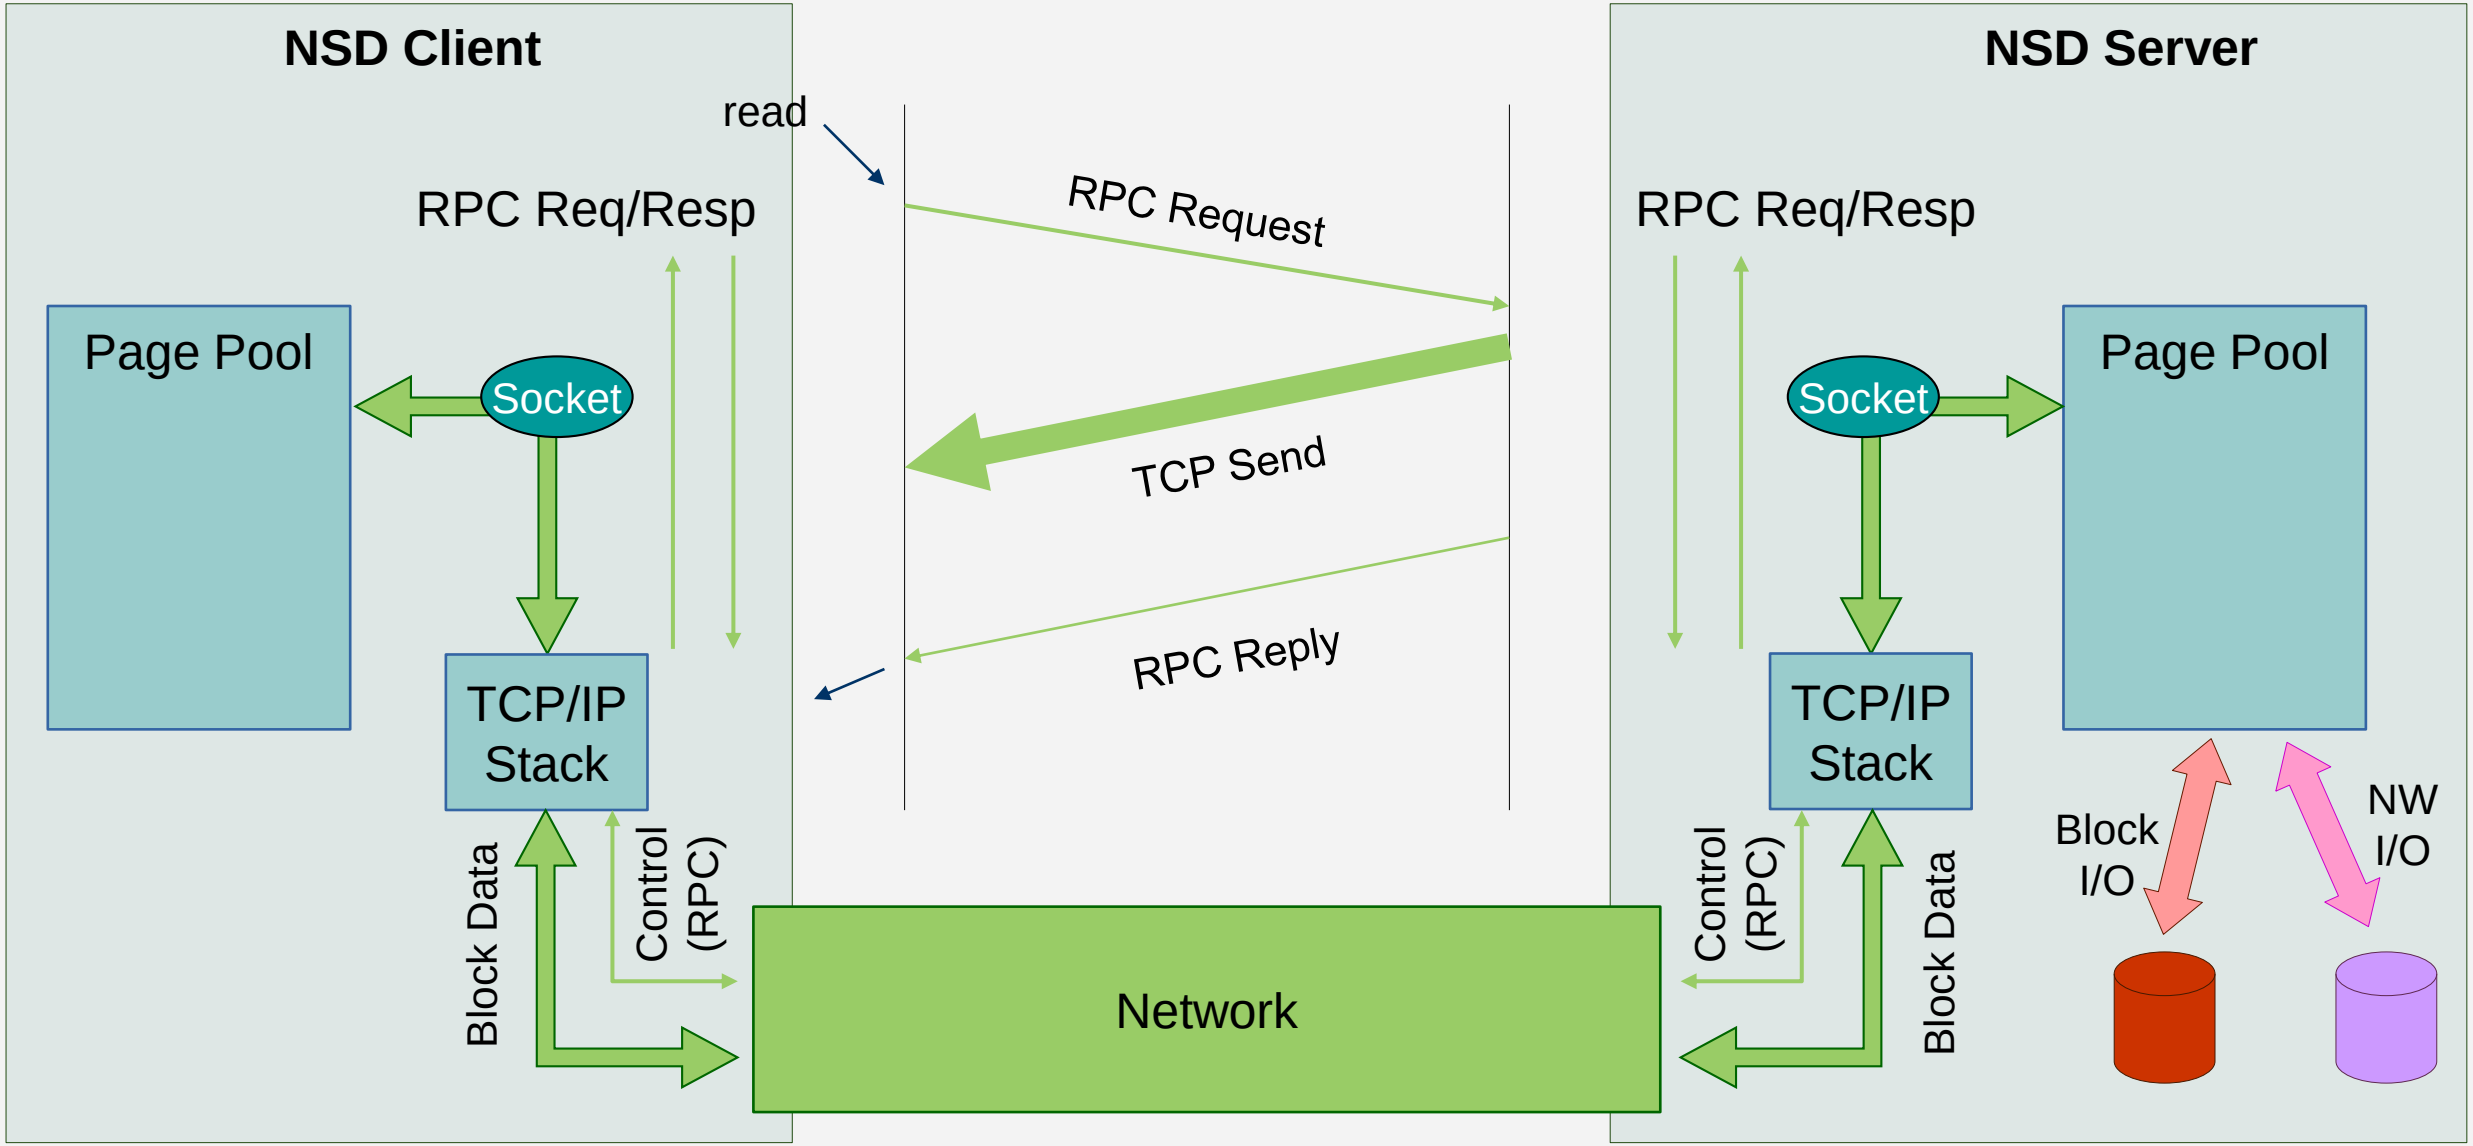

Spectrum Scale I/O Model with RDMA I/O (verbsRdma = yes)

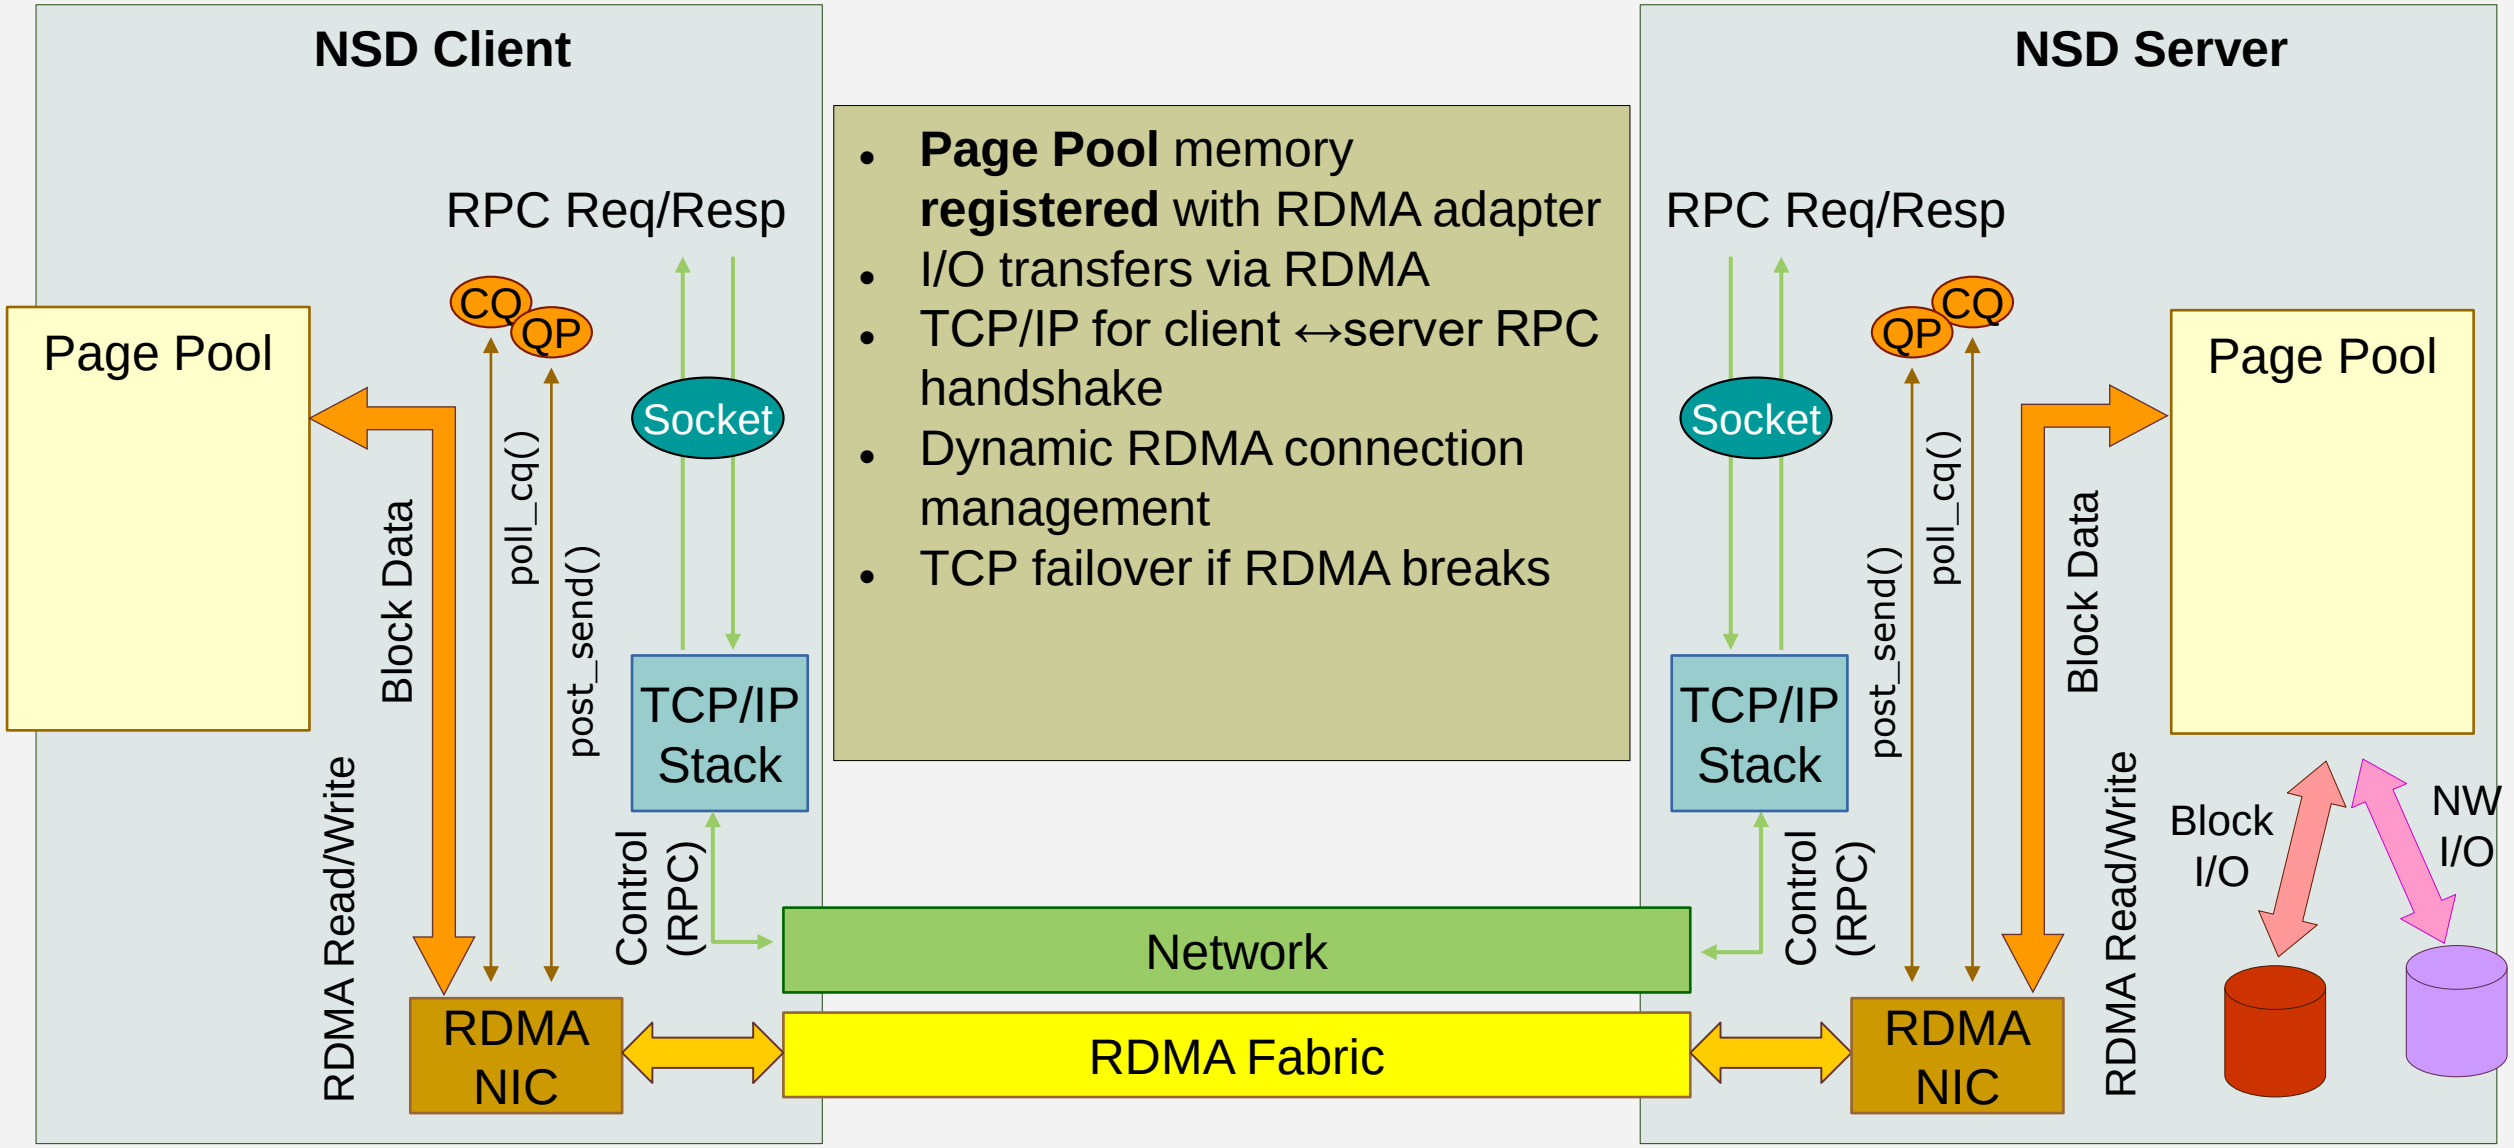

- **Page Pool** memory **registered** with RDMA adapter
- I/O transfers via RDMA
- TCP/IP for client ↔server RPC handshake
- Dynamic RDMA connection management
- TCP failover if RDMA breaks

Spectrum Scale I/O Model with RDMA I/O (verbsRdma = yes)

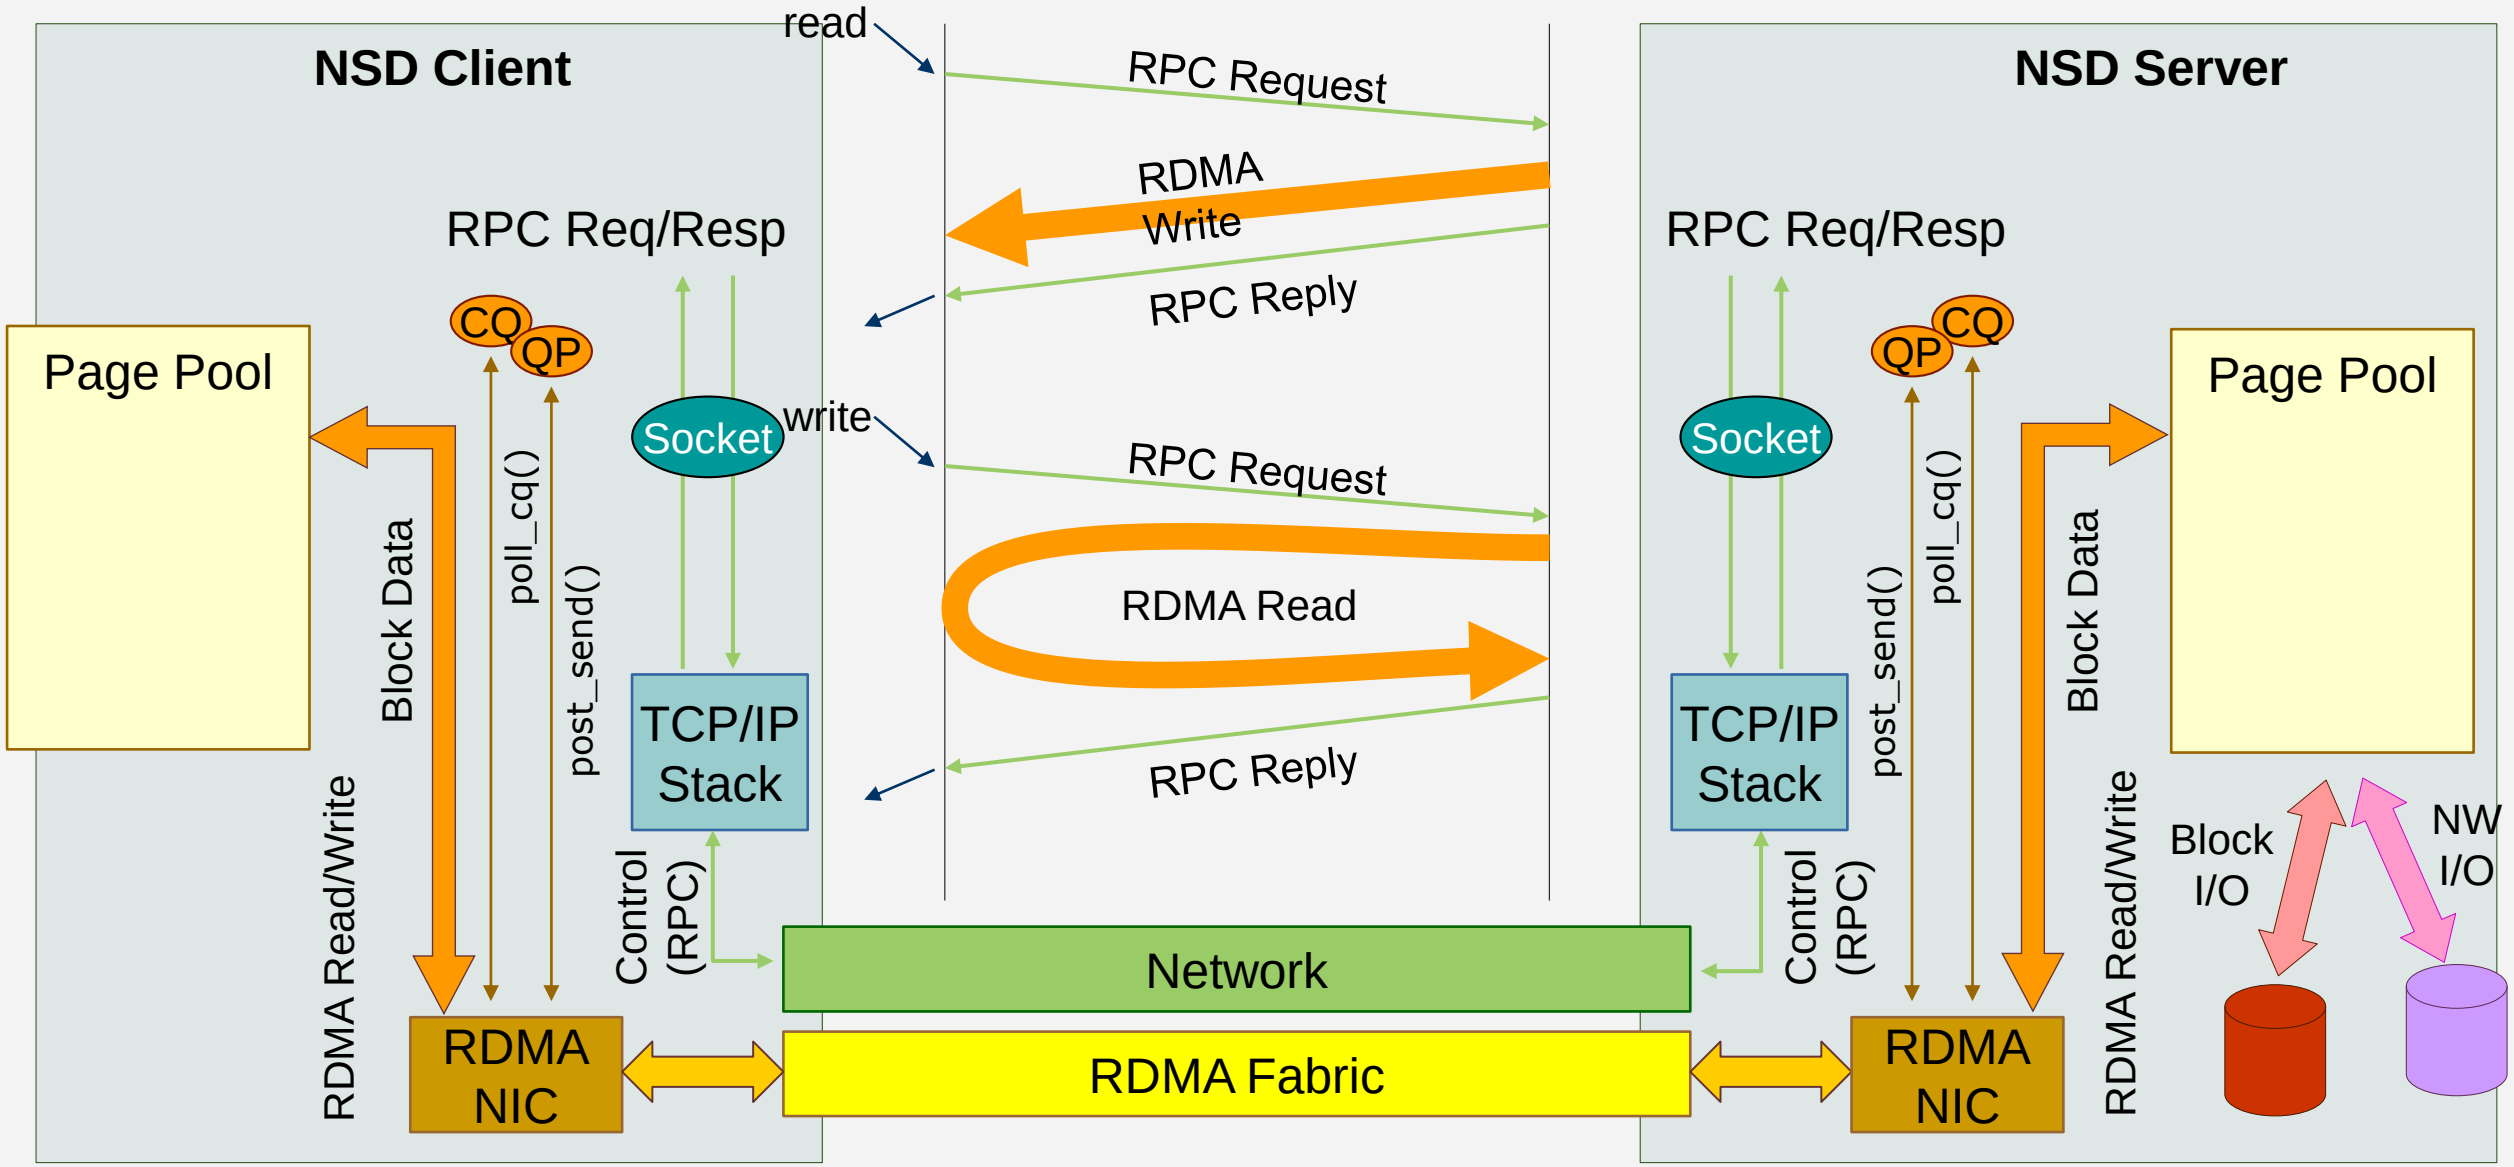

Spectrum Scale I/O Model with full RDMA I/O (verbsRdma = yes, verbsRdmaSend = yes)

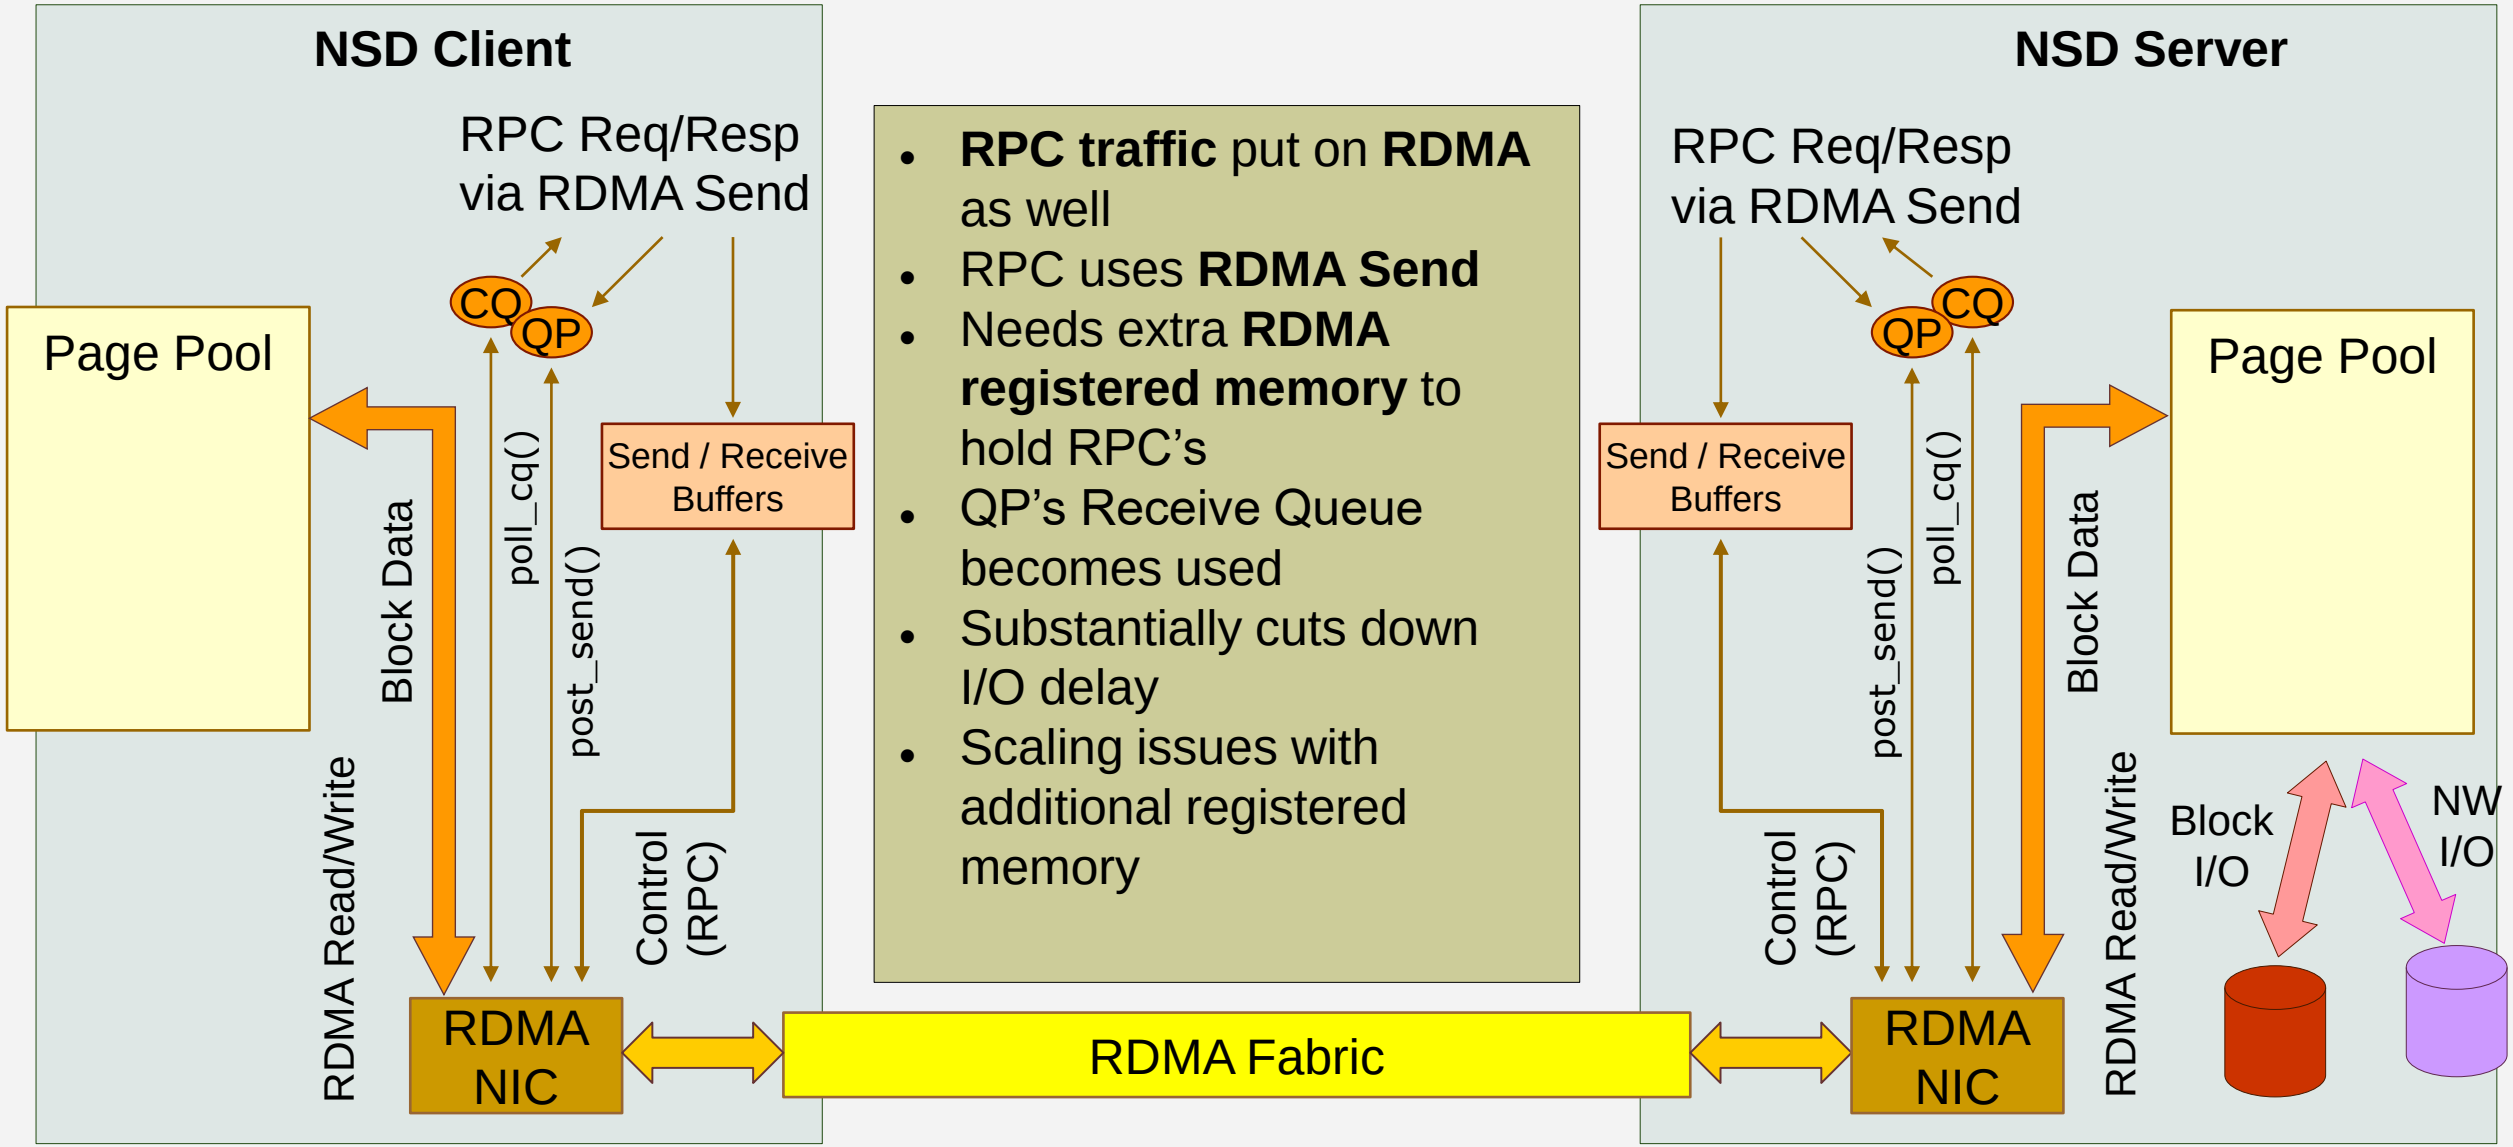

- **RPC traffic** put on **RDMA** as well
- RPC uses **RDMA Send**
- Needs extra **RDMA registered memory** to hold RPC's
- QP's Receive Queue becomes used
- Substantially cuts down I/O delay
- Scaling issues with additional registered memory

Spectrum Scale I/O Model with full RDMA I/O (verbsRdma = yes, verbsRdmaSend = yes)

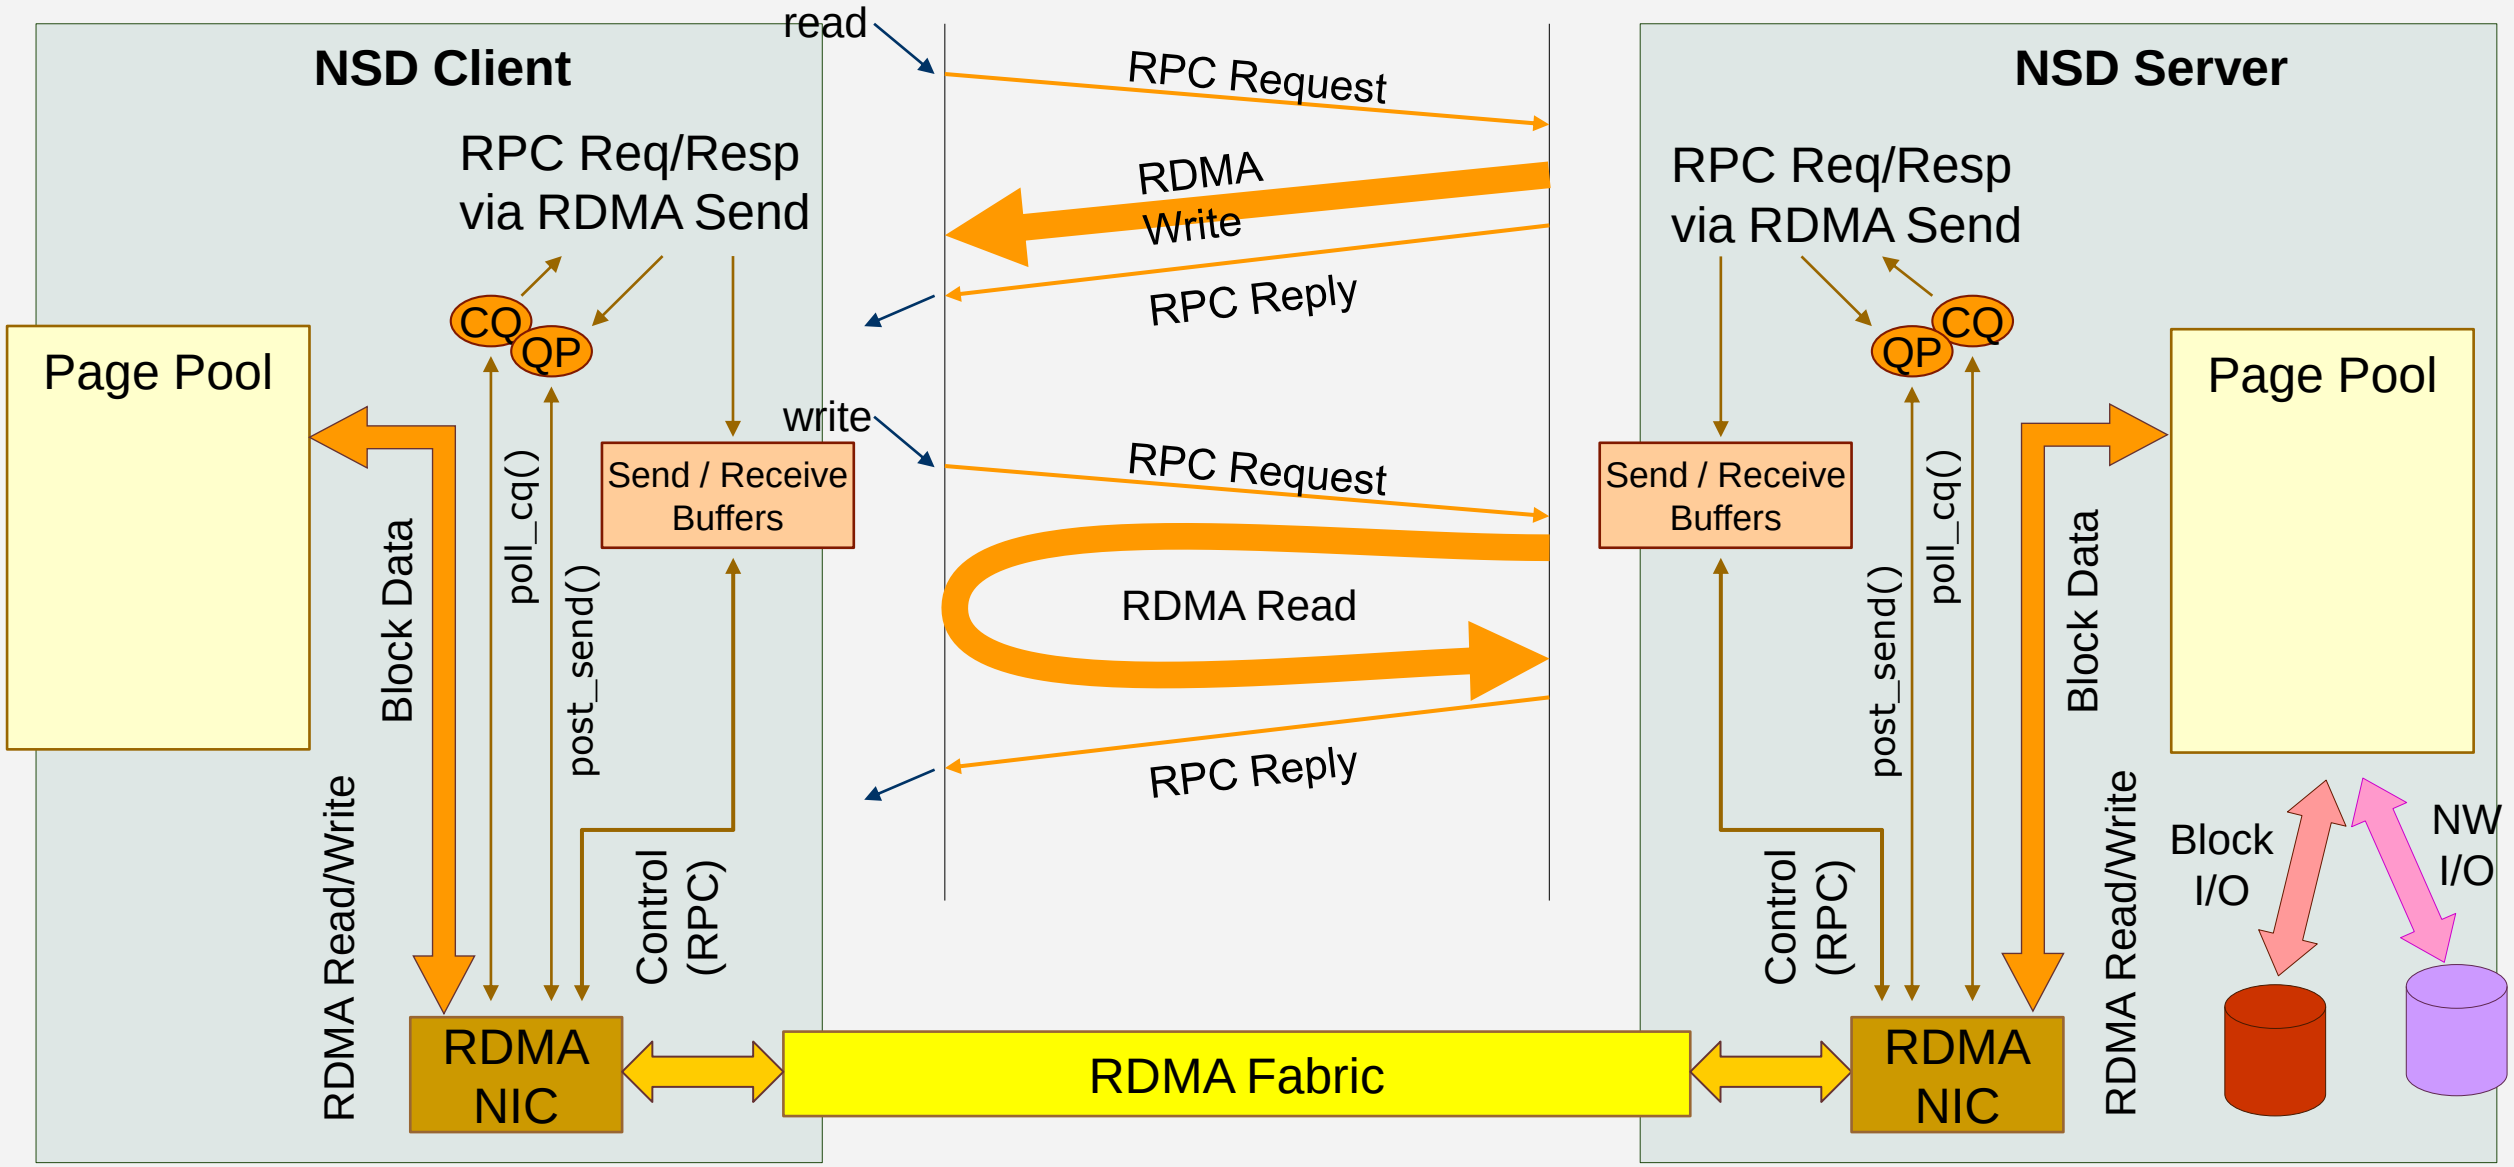

RDMA Agenda

Spectrum Scale
Strategy Days 2021

03. - 04. März 2021

Improving Spectrum Scale performance using RDMA

Part1 Introduction	Ingo Meents IT Architect Kelsterbach, Germany
Part 2 Deep Dive	Ralph Würthner IT Specialist Kelsterbach, Germany
Part 3 Performance results	Olaf Weiser IT Specialist Ehningen, Germany

Motivation

```
[root@speedyL ~] lspci -tv |sed -n '/ +-\[0000:\(85\|3a\)\]/,/8546/p'
+-[0000:85]-+-00.0-[86-93]--+-00.0-[87-93]--+-00.0-[88]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-01.0-[89]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-02.0-[8a]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-03.0-[8b]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-04.0-[8c]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-05.0-[8d]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-06.0-[8e]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-07.0-[8f]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-08.0-[90]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-09.0-[91]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-0a.0-[92]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 \-0b.0-[93]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               \-00.1  PMC-Sierra Inc. PM8546 B-FEIP PSX 96xG3 PCIe Storage Switch
+-[0000:3a]-+-00.0-[3b-48]--+-00.0-[3c-48]--+-00.0-[3d]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-01.0-[3e]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-02.0-[3f]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-03.0-[40]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-04.0-[41]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-05.0-[42]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-06.0-[43]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-07.0-[44]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-08.0-[45]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-09.0-[46]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 +-0a.0-[47]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               |                 \-0b.0-[48]----00.0  Samsung Electronics Co Ltd NVMe SSD Controller 172Xa/172Xb
|           |               \-00.1  PMC-Sierra Inc. PM8546 B-FEIP PSX 96xG3 PCIe Storage Switch
[root@speedyL ~]
```

```
[root@fscc-sr650-13 ~]# fio --filename=/dev/nvme1n1 --direct=1 --rw=read  --bs=256k --nur
bw="
  READ: bw=1942MiB/s (2036MB/s), 1942MiB/s-1942MiB/s (2036MB/s-2036MB/s)
[root@fscc-sr650-13 ~]#
```

Bei 24 drives == 48 GB/s

EDR:	at least 5 ports
HDR (100)	at least 5 ports
HDR (200)	at least 3 ports

… use RoCE , when …

– high bandwidth

– keep system overhead as small as possible

– small latency

– scale out over multiple NICs

RoCE - requirements

– RoCE enabled network cards

– special network features / support of
  – join TechTalk about RoCE config*

– Linux OS + [M]OFED
  – join TechTalk about RoCE config*

– each interface needs to have at least one IP

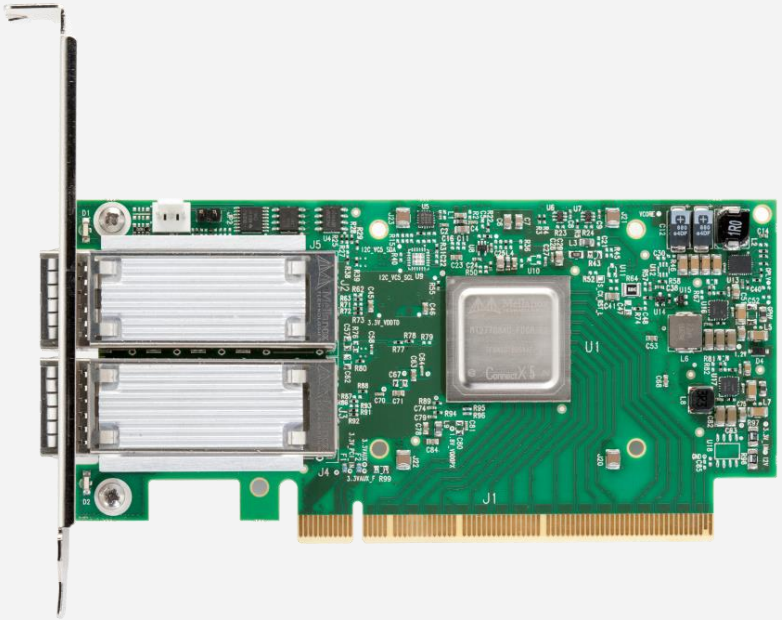

** not scheduled yet

RoCE – simplest topology

MLX Onyx

[standalone: master] (config) # show roce
RoCE mode : lossless

cumulus@switch:~$ net add roce lossless
cumulus@switch:~$ net pending
cumulus@switch:~$ net commit

nodeA

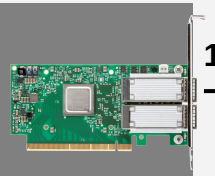

192.168.12.102/24

192.168.12.104/24

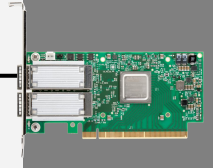

nodeB

```
#Set DSCP (L3) as trust mode
#  for the NIC

for INT in `ibdev2netdev | grep Up | awk '{print $5}' | xargs`
do
mlnx_qos -i $INT --trust dscp
mlnx_qos -i $INT --pfc 0,0,0,1,0,0,0,0
done
for MLX in `ibdev2netdev | grep Up | awk '{print $1}' | xargs`
do
echo 106 > /sys/class/infiniband/${MLX}/tc/1/traffic_class
 cma_roce_tos -d $MLX -t 106
 done
```

SpectrumScale's – little demo (1 client / 1 port , one ESS 8 ports)

```
[root@fscc-sr650-13 ~]# /usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/$(hostname)/file500g -n 500g -r 16m
-th 16 -fsync
/usr/lpp/mmfs/samples/perf/gpfsperf create seq /gpfs/ess5k16m/fscc-sr650-13/file500g
 recSize 16M nBytes 500G fileSize 500G
 nProcesses 1 nThreadsPerProcess 16
 file cache flushed before test
 not using direct I/O
 offsets accessed will cycle through the same file segment
 not using shared memory buffer
 not releasing byte-range token after open
 fsync at end of test
   Data rate was 7923914.14 Kbytes/sec, Op Rate was 472.30 Ops/sec, Avg Latency was 30.465 milliseconds, thread utilization
0.899, bytesTransferred 536870912000
[root@fscc-sr650-13 ~]#
[root@fscc-sr650-13 ~]#
```

TCP

SpectrumScale's – little demo

```
[root@fscc-sr650-13 ~]# /usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/$(hostname)/file500g -n 500g -r 16m
-th 16 -fsync
/usr/lpp/mmfs/samples/perf/gpfsperf create seq /gpfs/ess5k16m/fscc-sr650-13/file500g
 recSize 16M nBytes 500G fileSize 500G
 nProcesses 1 nThreadsPerProcess 16
 file cache flushed before test
 not using direct I/O
 offsets accessed will cycle through the same file segment
 not using shared memory buffer
 not releasing byte-range token after open
 fsync at end of test
   Data rate was 7923914.14 Kbytes/sec, Op Rate was 472.30 Ops/sec, Avg Latency was 30.465 milliseconds, thread
utilization 0.899, bytesTransferred 536870912000
[root@fscc-sr650-13 ~]#
[root@fscc-sr650-13 ~]# mmchconfig verbsRdma=enable,verbsRdmaSend=yes,verbsPorts mlx5_0/1,verbsRdmaCm enable
[root@fscc-sr650-13 ~]# mmshutdown -a ; mmstartup -a
```

TCP

SpectrumScale's – little demo

```
[root@fscc-sr650-13 ~]# /usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/$(hostname)/file500g -n 500g -r 16m -th 16 -
fsync
/usr/lpp/mmfs/samples/perf/gpfsperf create seq /gpfs/ess5k16m/fscc-sr650-13/file500g
 recSize 16M nBytes 500G fileSize 500G
 nProcesses 1 nThreadsPerProcess 16
 file cache flushed before test
 not using direct I/O
 offsets accessed will cycle through the same file segment
 not using shared memory buffer
 not releasing byte-range token after open
 fsync at end of test
   Data rate was 7923914.14 Kbytes/sec, Op Rate was 472.30 Ops/sec, Avg Latency was 30.465 milliseconds, thread utilization 0.899,
bytesTransferred 536870912000
[root@fscc-sr650-13 ~]#
[root@fscc-sr650-13 ~]# mmchconfig verbsRdma=enable,verbsRdmaSend=yes,verbsPorts mlx5_0/1,verbsRdmaCm enable
[root@fscc-sr650-13 ~]# mmshutdown -a ; mmstartup -a
```

TCP

```
[root@fscc-sr650-13 net]# /usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/$(hostname)/file500g -n 500g -r 16m -th 16 -
fsync
/usr/lpp/mmfs/samples/perf/gpfsperf create seq /gpfs/ess5k16m/fscc-sr650-13/file500g
  recSize 16M nBytes 500G fileSize 500G
  nProcesses 1 nThreadsPerProcess 16
  file cache flushed before test
  not using direct I/O
  offsets accessed will cycle through the same file segment
  not using shared memory buffer
  not releasing byte-range token after open
  fsync at end of test
    Data rate was 11936976.73 Kbytes/sec, Op Rate was 711.50 Ops/sec, Avg Latency was 20.544 milliseconds, thread utilization 0.914,
bytesTransferred 536870912000
[root@fscc-sr650-13 net]#
```

RDMA / RoCE

RoCE – complex, but powerful & flexible topology

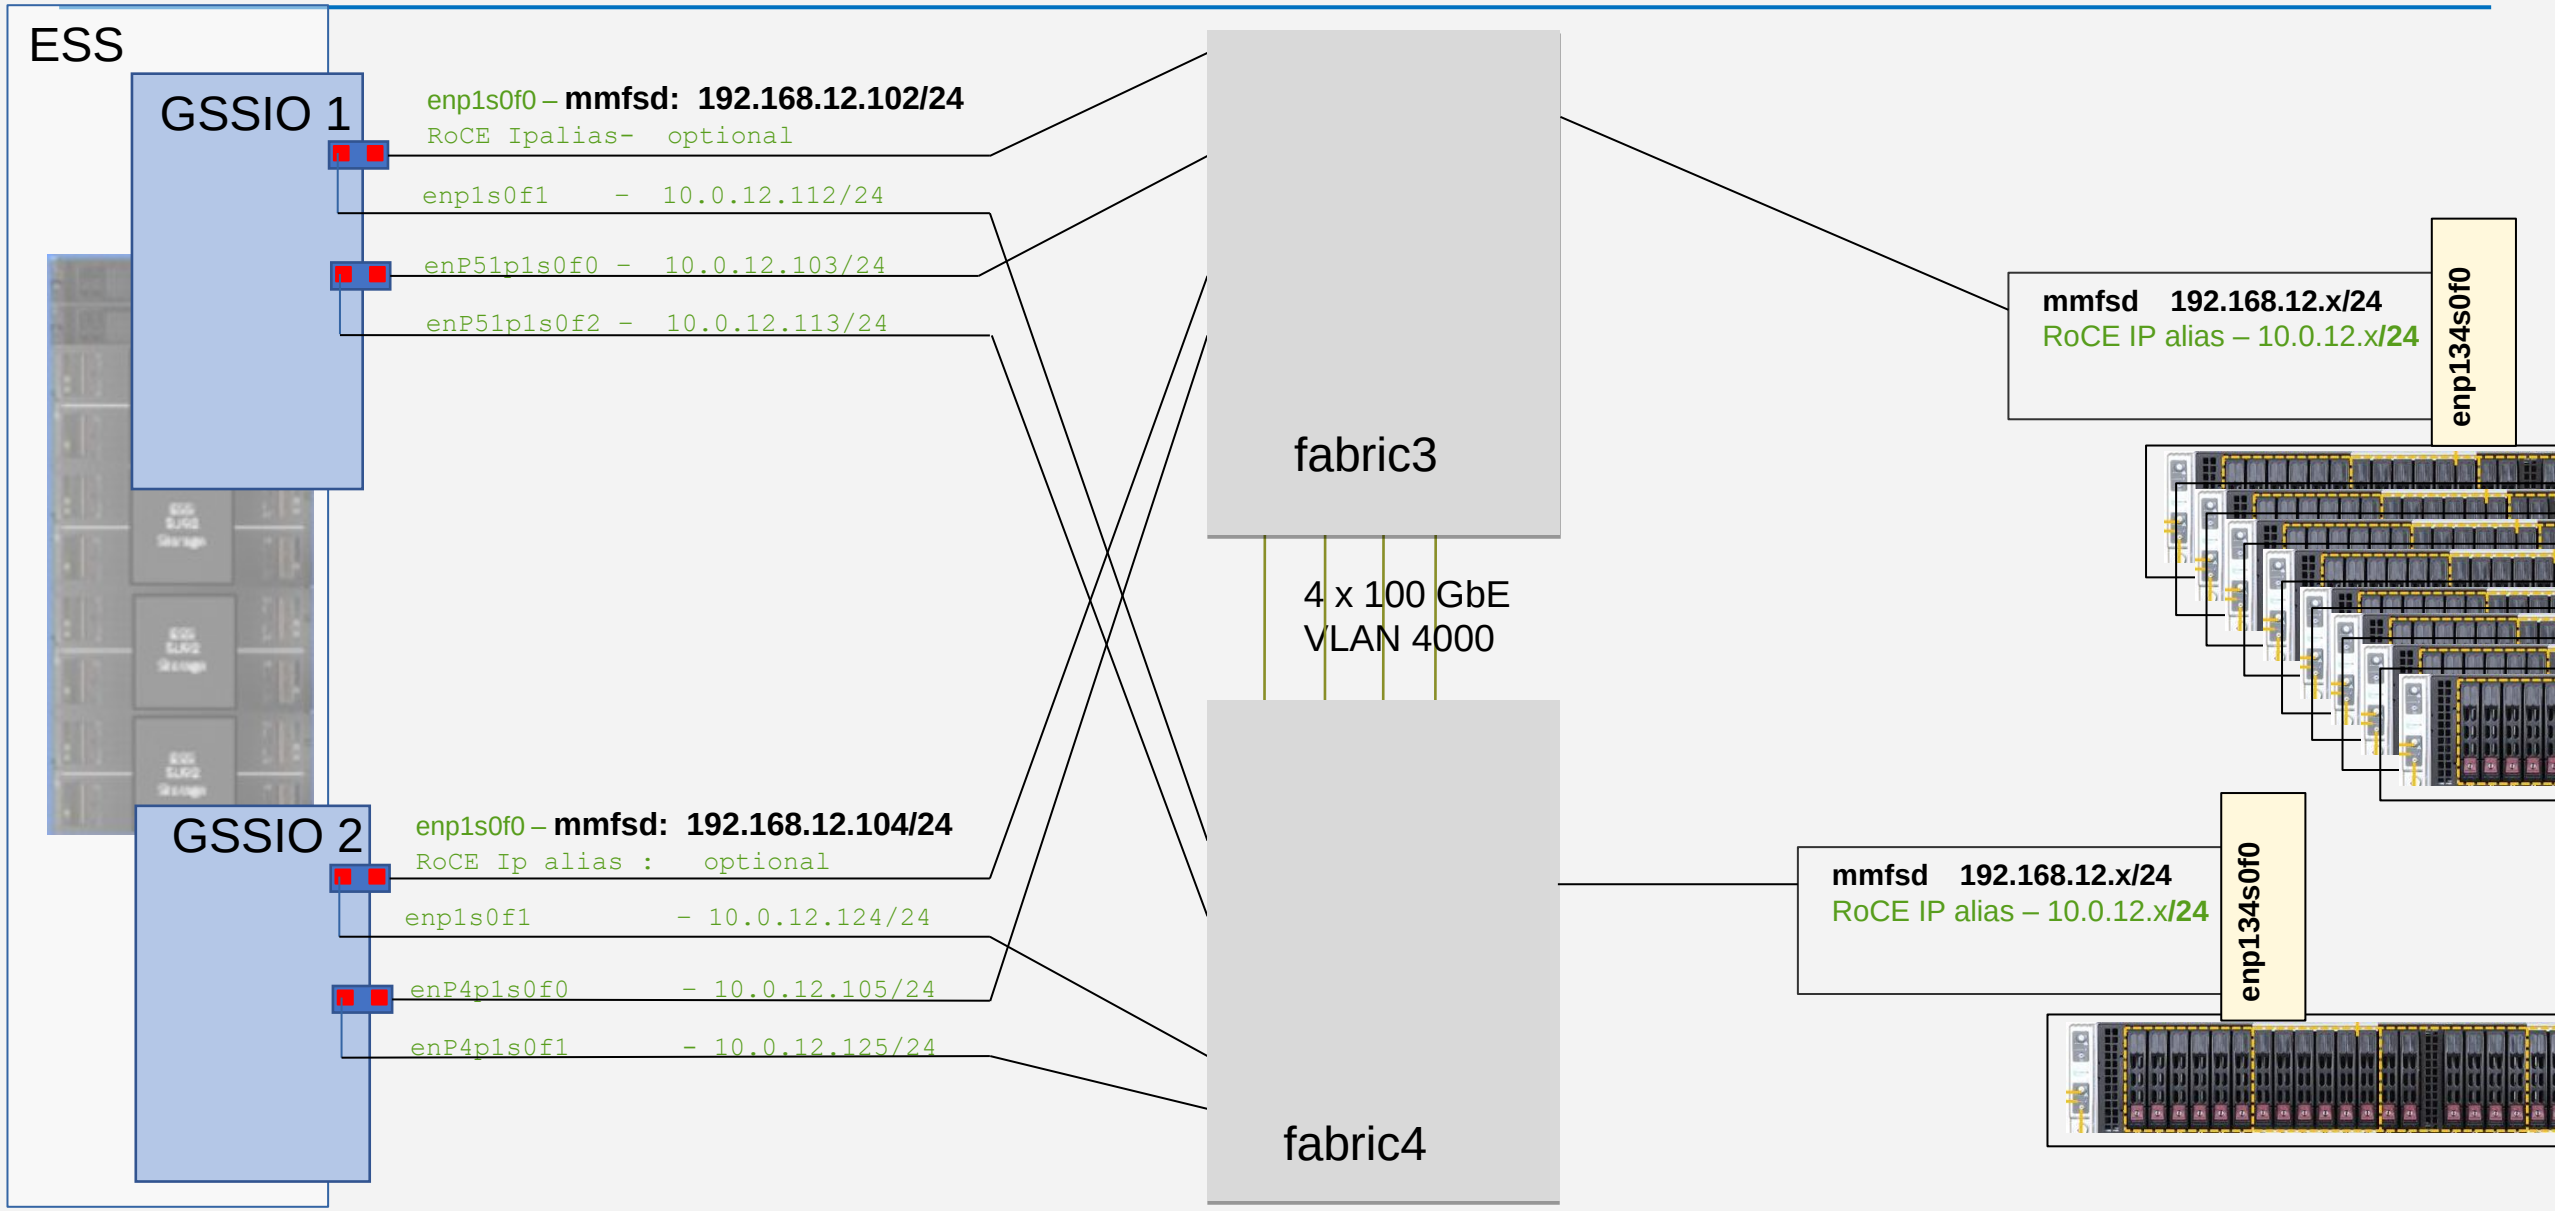

SpectrumScale's enhancement – RoCE over multiple NICs

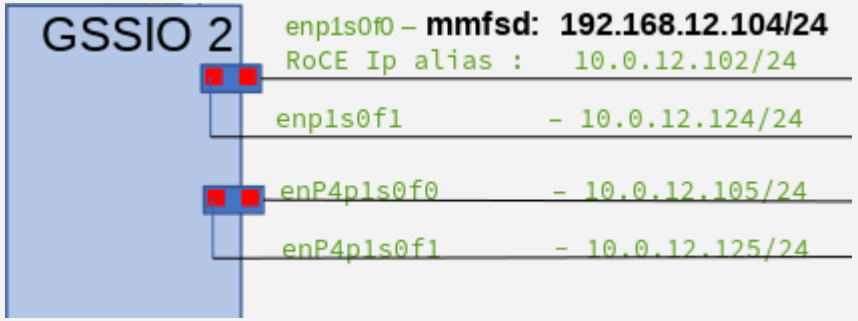

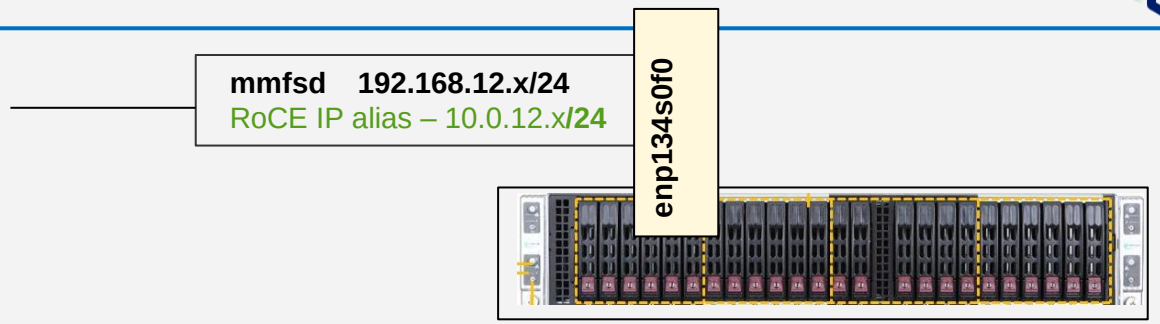

TCP :	1 socket	- one cable	100 GbE	one cable
RDMA:	4 QP	- four cables	400 GbE	one cable

```
RDMA Connections between nodes:
  Fabric 0 - Device mlx5_0 Port 1 Width 4x Speed EDR lid 0
    hostname                              idx CM state VS buff RDMA_CT(
    [...]
    ece-13                                0   Y  RTS   (Y)256  1
  Fabric 0 - Device mlx5_1 Port 1 Width 4x Speed EDR lid 0
    [...]
    ece-13                                1   Y  RTS   (Y)256  64
  Fabric 0 - Device mlx5_2 Port 1 Width 4x Speed EDR lid 0
    [...]
    ece-13                                2   Y  RTS   (Y)256  1
  Fabric 0 - Device mlx5_3 Port 1 Width 4x Speed EDR lid 0
    [...]
    ece-13                                3   Y  RTS   (Y)256  1
```

```
Device details:
  devicename     speed     mtu       duplex    rx_dropped rx_errors tx_
  enp134s0f0     100000    9000      full      0          0         0
RDMA Connections between nodes:
 Fabric 0 - Device mlx5_0 Port 1 Width 4x Speed EDR lid 0
   hostname                              idx CM state VS buff RDMA_CT(ERR
 […]
   ess5kio1                              0   Y  RTS   (Y)256  0      (0
     ess5kio2                              0   Y  RTS   (Y)256  0      (0
   ess5kio1                              1   Y  RTS   (Y)256  0      (0
     ess5kio2                              1   Y  RTS   (Y)256  0      (0
   ess5kio1                              2   Y  RTS   (Y)256  0      (0
     ess5kio2                              2   Y  RTS   (Y)256  0      (0
   ess5kio1                              3   Y  RTS   (Y)256  0      (0
     ess5kio2                              3   Y  RTS   (Y)256  0      (0
[…]
```

SpectrumScale's enhancement – RoCE over multiple NICs

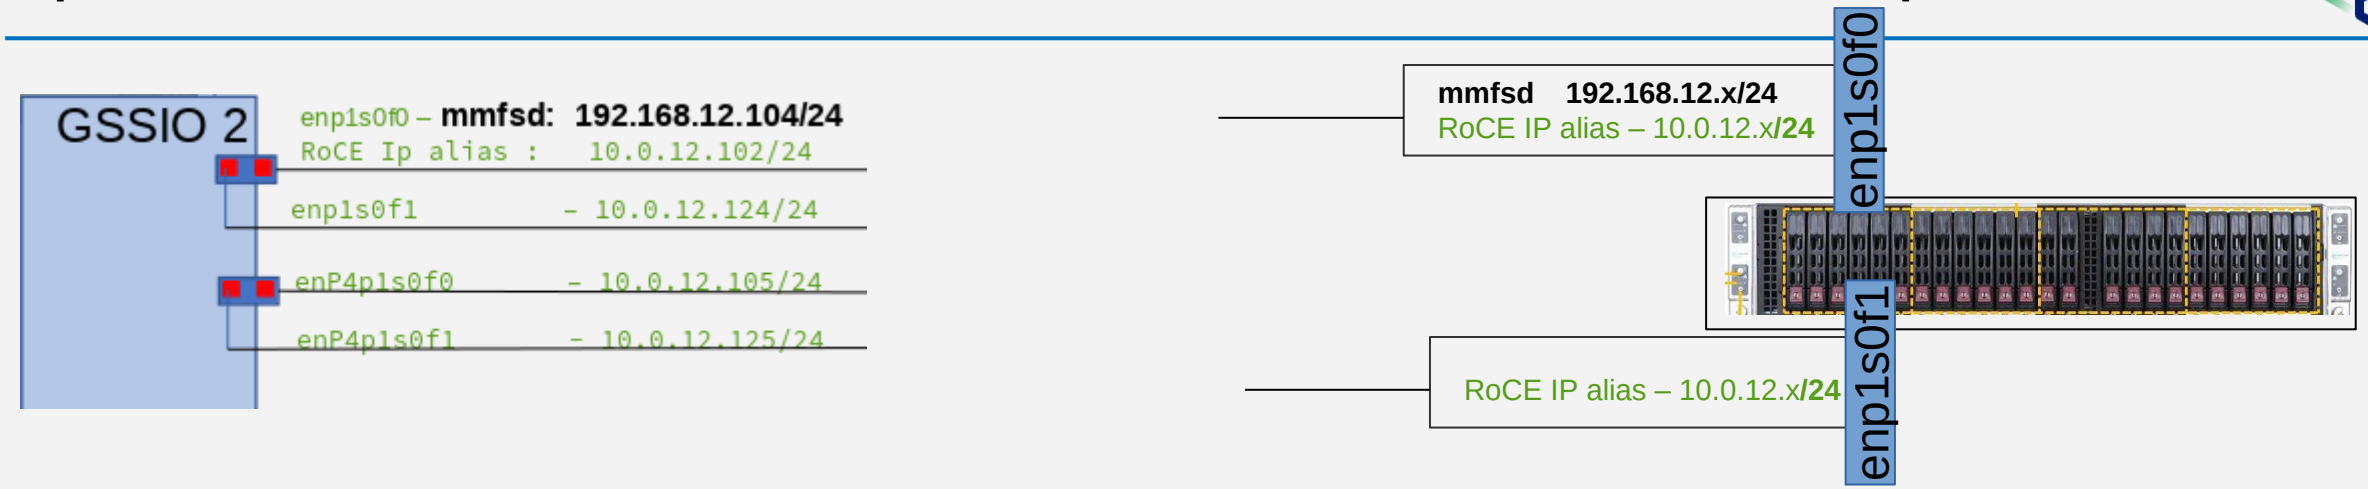


```
[root@ems5000 ~]# /usr/lpp/mmfs/samples/perf/gpfsperf read seq  /gpfs/ess5k16M/$(hostname)/file500g -n 100g -r 16m -th 32 -fsync
/usr/lpp/mmfs/samples/perf/gpfsperf read seq /gpfs/ess5k16M/ems5000.localdomain.local/file500g
 recSize 16M nBytes 100G fileSize 500G
 nProcesses 1 nThreadsPerProcess 32
 file cache flushed before test
 not using direct I/O
 offsets accessed will cycle through the same file segment
 not using shared memory buffer
 not releasing byte-range token after open
   Data rate was 22435014.38 Kbytes/sec, Op Rate was 1337.23 Ops/sec, Avg Latency was 22.973 milliseconds, thread utilization 0.960,
bytesTransferred 107374182400
[root@ems5000 ~]# ^C
```

Performance – TCP vs RDMA

RDMA:

```
[root@fscc-sr650-13 ~]# mmdsh -N ece-13,ece-14,ece-15 "/usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/\$(hostname)
ece-13.mmfsd.net:      Data rate was 10325041.64 Kbytes/sec, Op Rate was 615.42 Ops/sec, Avg Latency was 22.211 milliseconds, threa
ece-14.mmfsd.net:      Data rate was 10219641.73 Kbytes/sec, Op Rate was 609.14 Ops/sec, Avg Latency was 24.257 milliseconds, threa
ece-15.mmfsd.net:      Data rate was 10192525.19 Kbytes/sec, Op Rate was 607.52 Ops/sec, Avg Latency was 24.322 milliseconds, threa

[root@fscc-sr650-13 ~]# mmchconfig verbsRdmaCm=disable
[...]
[root@fscc-sr650-13 ~]# mmshutdown -a ; mmstartup -a
[...]
[root@fscc-sr650-13 ~]#
```

TCP:

```
[root@fscc-sr650-13 net]# mmdsh -N ece-13,ece-14,ece-15 "/usr/lpp/mmfs/samples/perf/gpfsperf create seq  /gpfs/ess5k16m/\$(hostname
ece-15.mmfsd.net:      Data rate was 6915593.66 Kbytes/sec, Op Rate was 412.20 Ops/sec, Avg Latency was 35.266 milliseconds, threa
ece-14.mmfsd.net:      Data rate was 6895936.18 Kbytes/sec, Op Rate was 411.03 Ops/sec, Avg Latency was 35.587 milliseconds, threa
ece-13.mmfsd.net:      Data rate was 6809766.16 Kbytes/sec, Op Rate was 405.89 Ops/sec, Avg Latency was 35.441 milliseconds, threa
[root@fscc-sr650-13 net]#
```

System utilization – TCP vs RDMA

TCP

```
[root@fscc-sr650-13 ~]# vmstat 1
procs -----------memory---------- ---swap-- -----io---- -system-- ------cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs  us sy id wa st

10  0      0 38184948 438956 14603296    0    0     0     0 137523 29996  8 18 74  0  0
 5  0      0 38182236 438956 14603300    0    0     0     0 147294 31296  7 19 74  0  0
12  0      0 38186504 438956 14603300    0    0     0   964 179346 30488  8 18 74  0  0
12  0      0 38170932 438956 14603852    0    0     0     8 247307 30226  9 17 74  0  0
12  0      0 38186996 438956 14603856    0    0     0     0 238873 31551  9 19 73  0  0
16  0      0 38177540 438956 14603856    0    0     0     0 189341 37039  9 18 73  0  0
10  0      0 38176680 438956 14603872    0    0     0     8 177513 35386  8 18 75  0  0
11  0      0 38168844 438956 14603872    0    0     0     0 212547 31395  9 19 72  0  0
18  0      0 38175756 438956 14603864    0    0     0     0 245655 35327  8 18 74  0  0
 9  0      0 38180520 438956 14603140    0    0     0    16 319003 35797  9 19 71  0  0
14  0      0 38171788 438956 14603864    0    0     0     0 329079 30652  8 18 75  0  0
 6  0      0 38170956 438956 14603872    0    0     0     8 165768 29923 10 20 70  0  0
 9  0      0 38182612 438956 14603316    0    0     0     0 155475 27275  8 17 75  0  0
14  0      0 38183884 438956 14603316    0    0     0     0 146383 35045  6 16 79  0  0
16  0      0 38184716 438956 14603320    0    0     0    60 202398 24787  8 21 70  0  0
 8  0      0 38186912 438956 14603324    0    0     0     8 157181 35263  7 18 75  0  0
12  0      0 38183064 438956 14603324    0    0     0     0 261088 26387  7 17 76  0  0
13  0      0 38184844 438956 14603324    0    0     0     0 281521 29047  7 18 75  0  0
 8  0      0 38182928 438956 14603328    0    0     0     0 308903 33511  7 19 74  0  0
```

RDMA

```
[root@fscc-sr650-13 ~]# vmstat 1
procs -----------memory---------- ---swap-- -----io---- -system-- ------cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs  us sy id wa st

15  0      0 36855756 438956 14614296    0    0     0     0 14743 14913 11 14 75  0  0
10  0      0 36856228 438956 14614300    0    0     0     0 13738 14639 11 11 78  0  0
12  0      0 36854956 438956 14614340    0    0     0     4 20024 19665 13 17 70  0  0
12  0      0 36856436 438956 14614260    0    0     0    44 13683 13272 11 11 78  0  0
11  0      0 36857056 438956 14614264    0    0     0     0 13810 18681 10 12 77  0  0
12  0      0 36856952 438956 14614264    0    0     0     0 17028 13837 11 15 74  0  0
11  0      0 36857172 438956 14614272    0    0     0     0 14731 13109 12 13 75  0  0
17  0      0 36857000 438956 14614272    0    0     0    28 15525 13813 12 15 73  0  0
 3  0      0 36856764 438956 14614272    0    0     0   188 16106 13691 13 14 73  0  0
13  0      0 36855412 438956 14614280    0    0     0     0 14351 13213 10 12 78  0  0
10  0      0 36857768 438956 14614288    0    0     0     0 15371 13317 12 14 74  0  0
21  0      0 36858212 438956 14614292    0    0     0     0 14593 11774 13 12 75  0  0
 7  0      0 36856176 438956 14614236    0    0     0     0 18269 16061 13 18 70  0  0
 3  0      0 36856516 438956 14614180    0    0     0    20 13565 20803 11 10 79  0  0
 8  0      0 36856904 438956 14614184    0    0     0     0 14220 13212 11 14 75  0  0
12  0      0 36857896 438956 14614184    0    0     0     0 11278 12052  9  9 82  0  0
 7  0      0 36858148 438956 14614188    0    0     0     0 18394 16525 15 19 66  0  0
13  0      0 36859016 438956 14614188    0    0     0     0 13509 19105 11 11 78  0  0
 5  0      0 36859916 438956 14614192    0    0     0     0 12517 17178 10 10 80  0  0
```

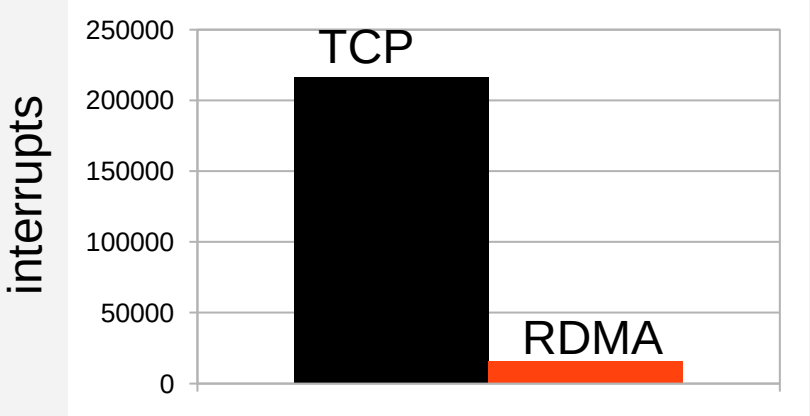

IBM