

# z/OS Communications Server Performance Update

*SHARE 2026 Winter Technical Conference*  
*Session 140 – February 22, 2026*

*Christopher Nyamful - [cnyamfu@us.ibm.com](mailto:cnyamfu@us.ibm.com)*

All performance information was determined on dedicated hardware in a controlled environment. Actual results may vary.

Performance information is provided “AS IS” and no warranties or guarantees are expressed or implied by IBM.

IBM statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

# Agenda

- z/OS Enhanced Queued Direct I/O (EQDIO) Performance on z17
- z/OS HiperSockets and SMCDv2 Performance on z17
- z/OS Communications Server Best Practices
- z/OS 3.2 AT-TLS Performance

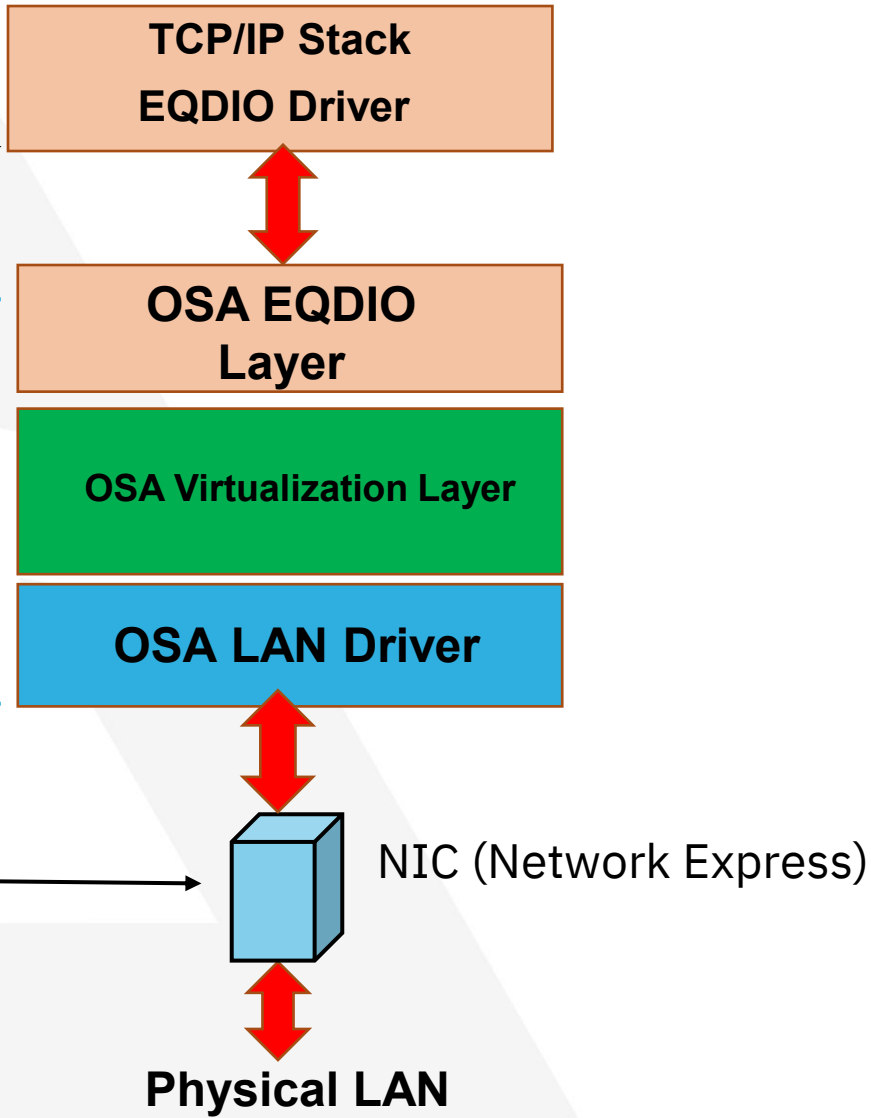


# z/OS Enhanced Queued Direct I/O (EQDIO) Performance on z17

# Enhanced QDIO (EQDIO) Architecture

Enhanced QDIO (EQDIO) architecture is introduced to allow software to interoperate with the OSA firmware and adapter. There are 3 significant system changes for networking:

- 1. **Software:** The z/OS Communications Server support for “OSA” is a complete rewrite. Key changes are the TCP/IP stack’s INTERFACE statement and device driver.
- 2. **OSA firmware/zCP:** The OSA firmware is updated to support OSH and EQDIO
- 3. **Adapter Hardware:** The physical NIC / RNIC hardware adapter is new

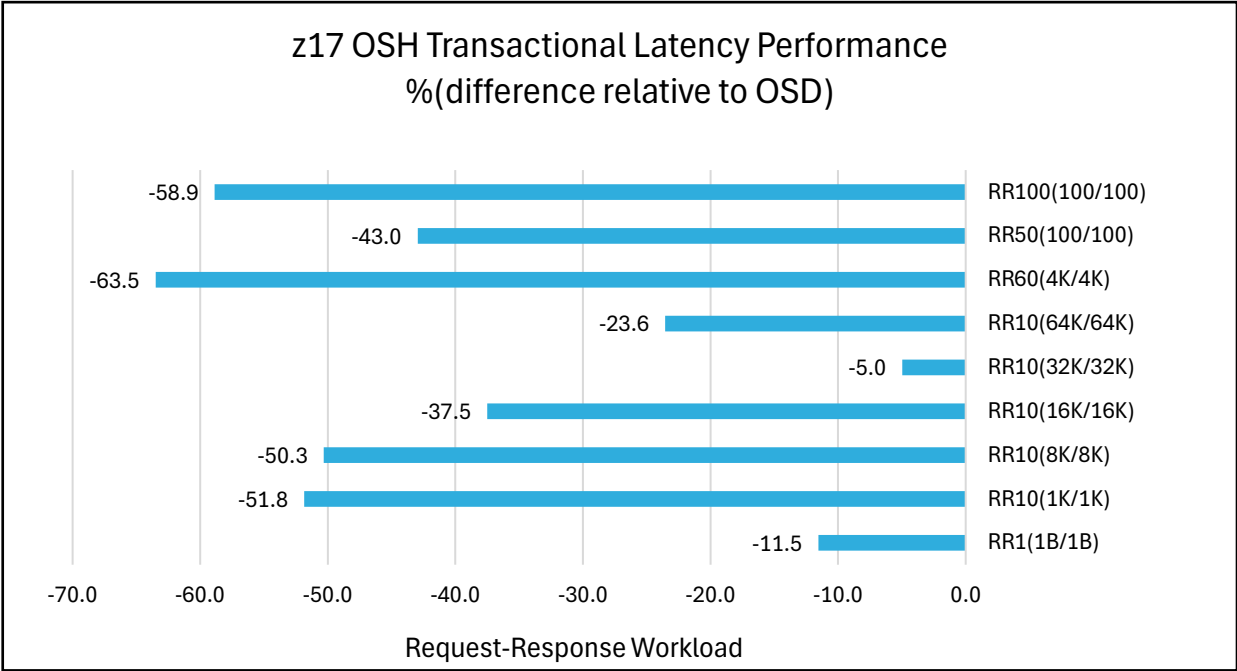


# New terminology for z/OS Networking on z17

1. New z17 “OSA” network feature name is called “**Network Express**”
2. Each physical card / slot becomes two PCHIDs / CHPIDs, each PCHID supports a single physical port
3. Each PCHID / port supports “two personalities”:
  - **OSA** for standard Ethernet (via EQDIO) defined as an **OSH CHPID** defined with channel devices
  - **RoCE** for RDMA for SMC-R (via native PCI architecture) defined with FIDs using a new **FID Type NETH**. NETH supports native PCI functions such as RoCE or standard Ethernet used by Linux.
4. EQDIO architecture:
  - Each OSA interface is defined with a single OSH device number. **Control read/write devices are replaced with control queues.**
  - **The z/OS user defined VTAM TRLE is eliminated!** A dynamic TRLE is created for each **EQENET INTERFACE.**
  - Layer 2 only (with IP assists/offloads)
  - Default MTU is 9000 (jumbo frames)

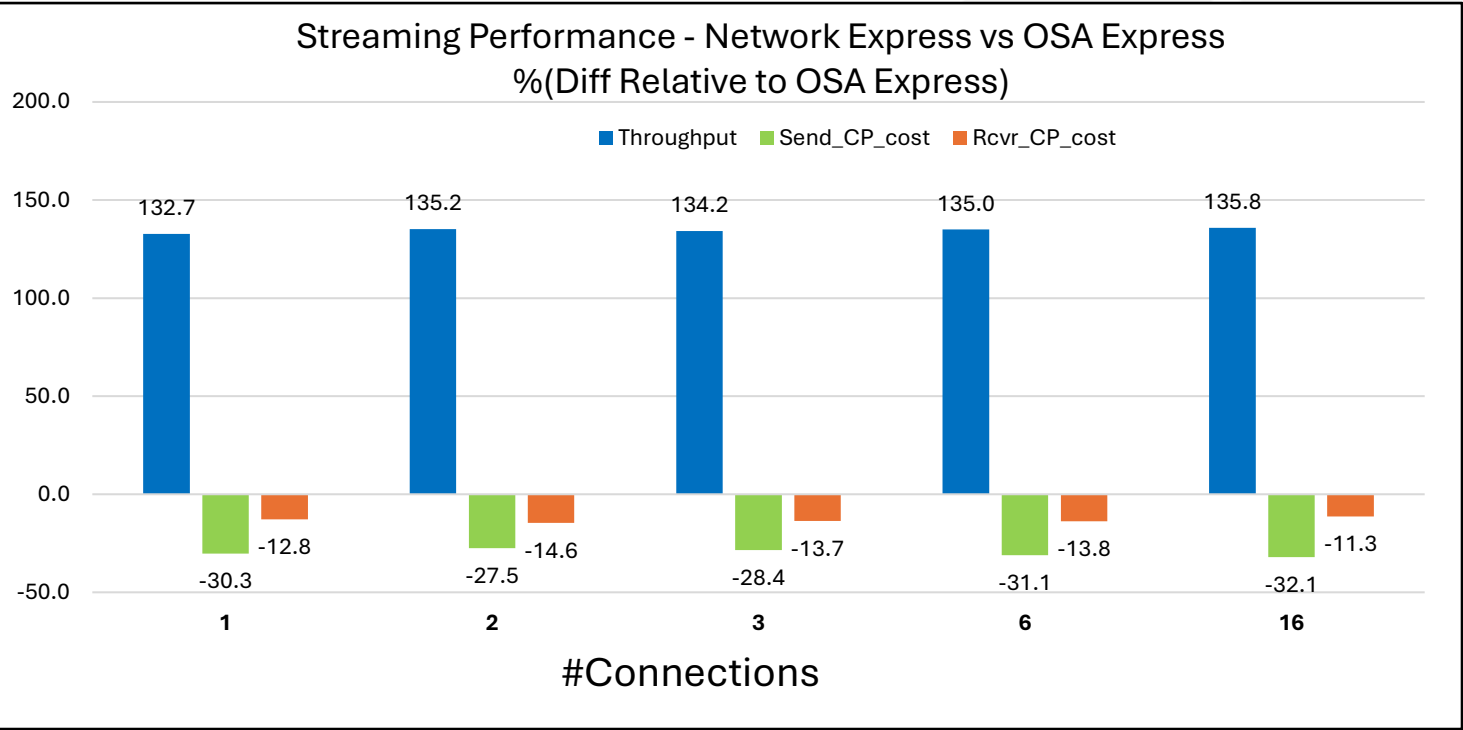
# Network Express OSH Performance

# Interactive (RR) Performance



- OSH (EQDIO) compared to OSD (QDIO)
- 2CPs + 25Gb + 1500 MTU + LSO
- Significant reduction in transactional latency
  - Up to 60% improvement in network response times
- Upcoming hardware bundle updates could further improve performance

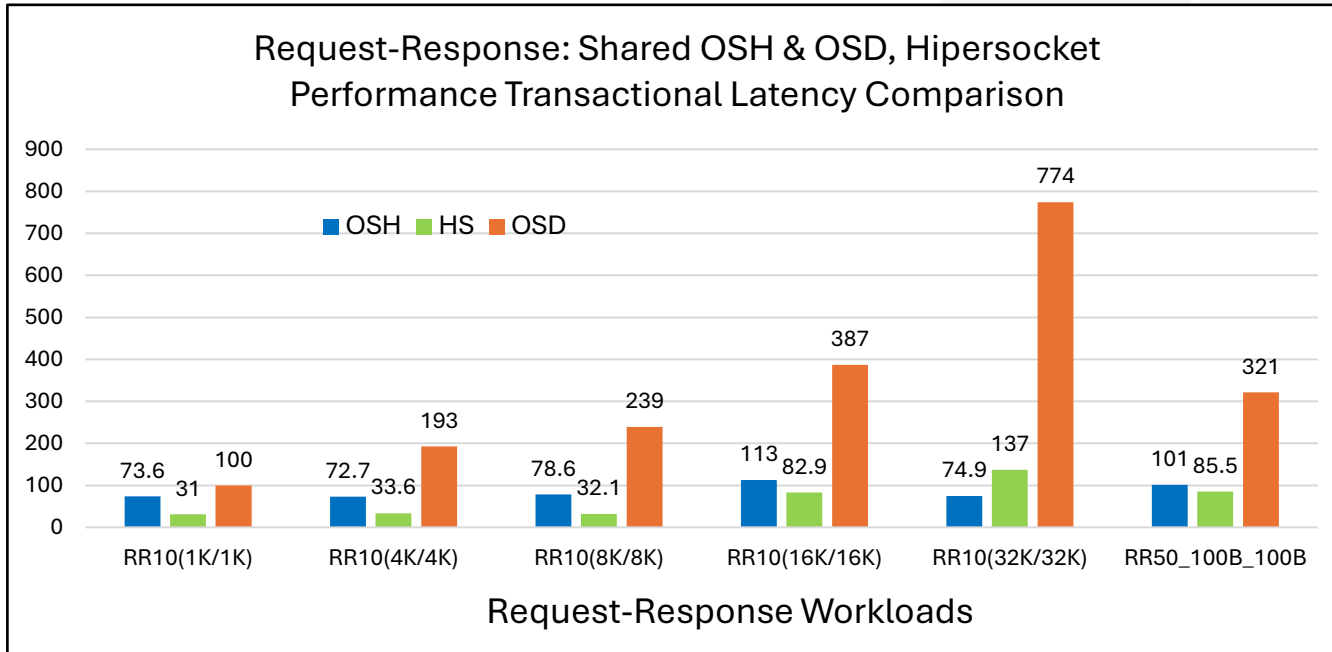
# Streaming (STR) Performance



- OSH (EQDIO) compared to OSD (QDIO)
- 2CPs + 25Gb + 1500 MTU + LSO
- Streaming performance (throughput, sender & receiver CPU cost)
- Significant performance improvement
  - Up to 135% improvement in throughput compared to OSD
  - 30% average savings in CPU cost

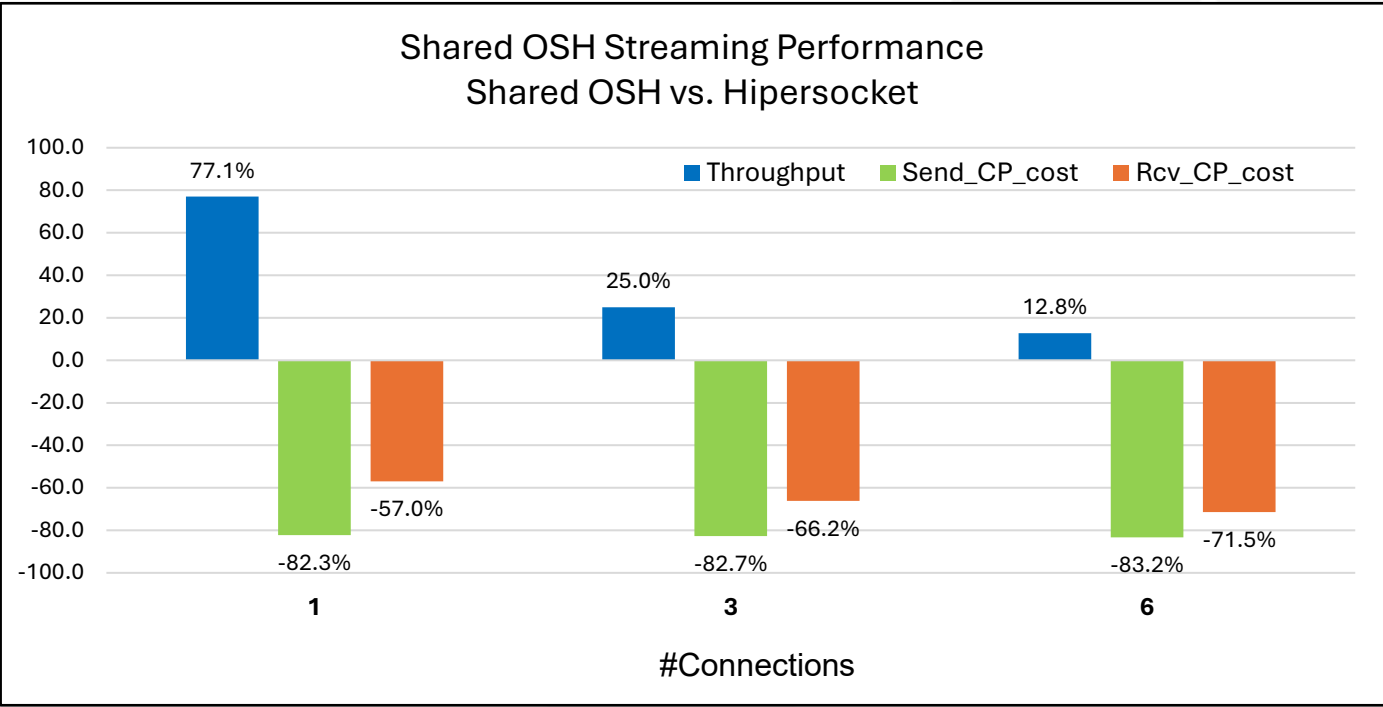
# Shared OSH Performance

# Shared OSH Performance – networking CPU cost



- Response time comparisons between shared Network Express OSH, shared OSA Express and Hipersocket on z17
- OSD vs Hipersocket
  - Small payload : 474%
  - Medium payload : 366%
  - Large payload : 464%
- OSH vs Hipersocket
  - Small payload : 116%
  - Medium payload : 36%
  - Large payload : -45%
- Shared OSH shows an average of **30%** reduction in CPU cost compared to Hipersocket

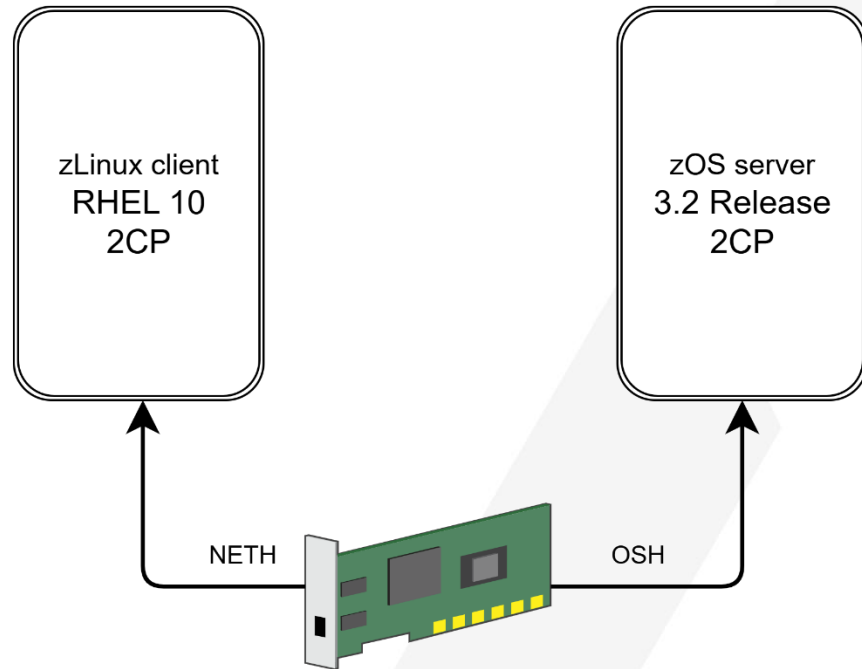
# Shared OSH Performance – Throughput improvements



- Sharing (internal loopback) over Network Express OSH compared with Hipersockets (HS)
- Shared OSH shows significant performance improvement over Hipersocket
- The sender and receiver shows an average of 70% reduction in network related CPU cost compared to Hipersocket

# Shared Network Express Linux on Z, z/OS

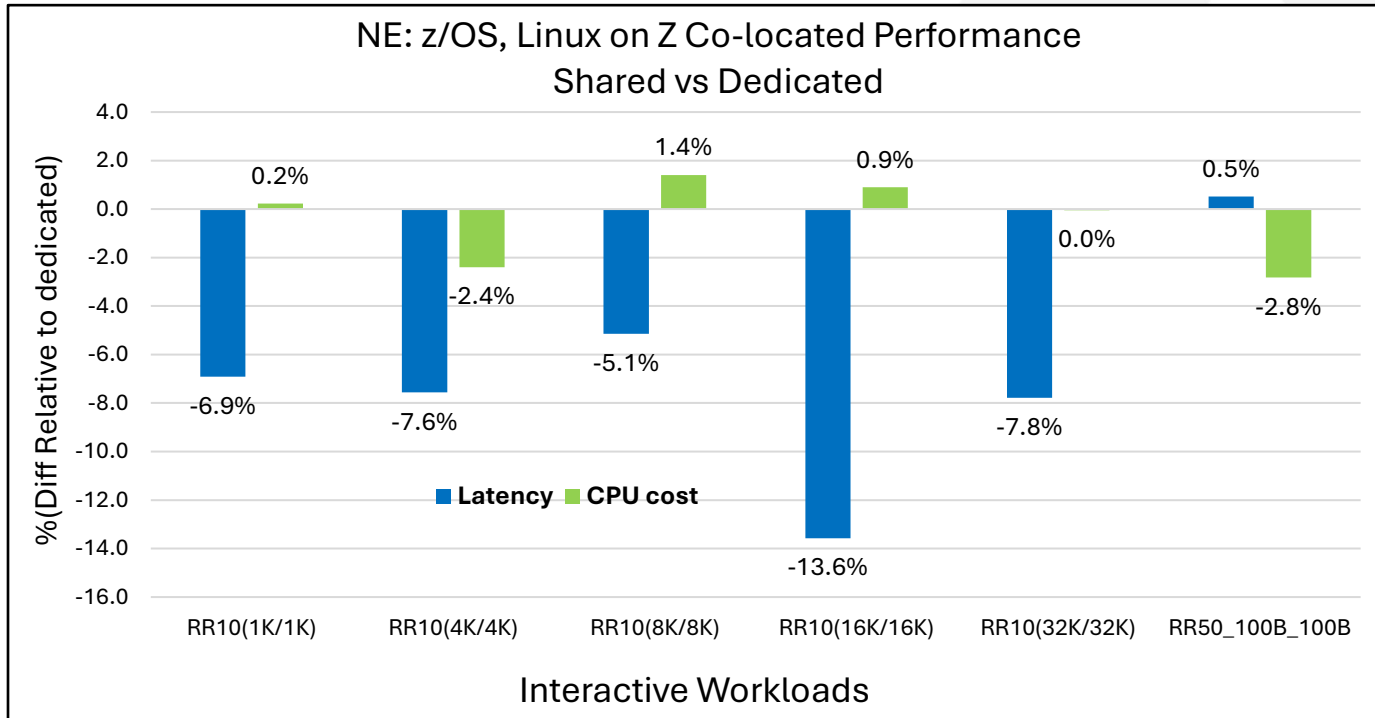
# Shared Network Express Configuration



- Network Express (NE) port can be configured as type OSH and NETH at the same time on z/OS
- Linux uses native PCI support for NETH to support IP and RDMA traffic
- A port on the NE can be shared between zLinux client as type NETH and z/OS server as type OSH
- z/OS and Linux applications running co-located can benefit from the card

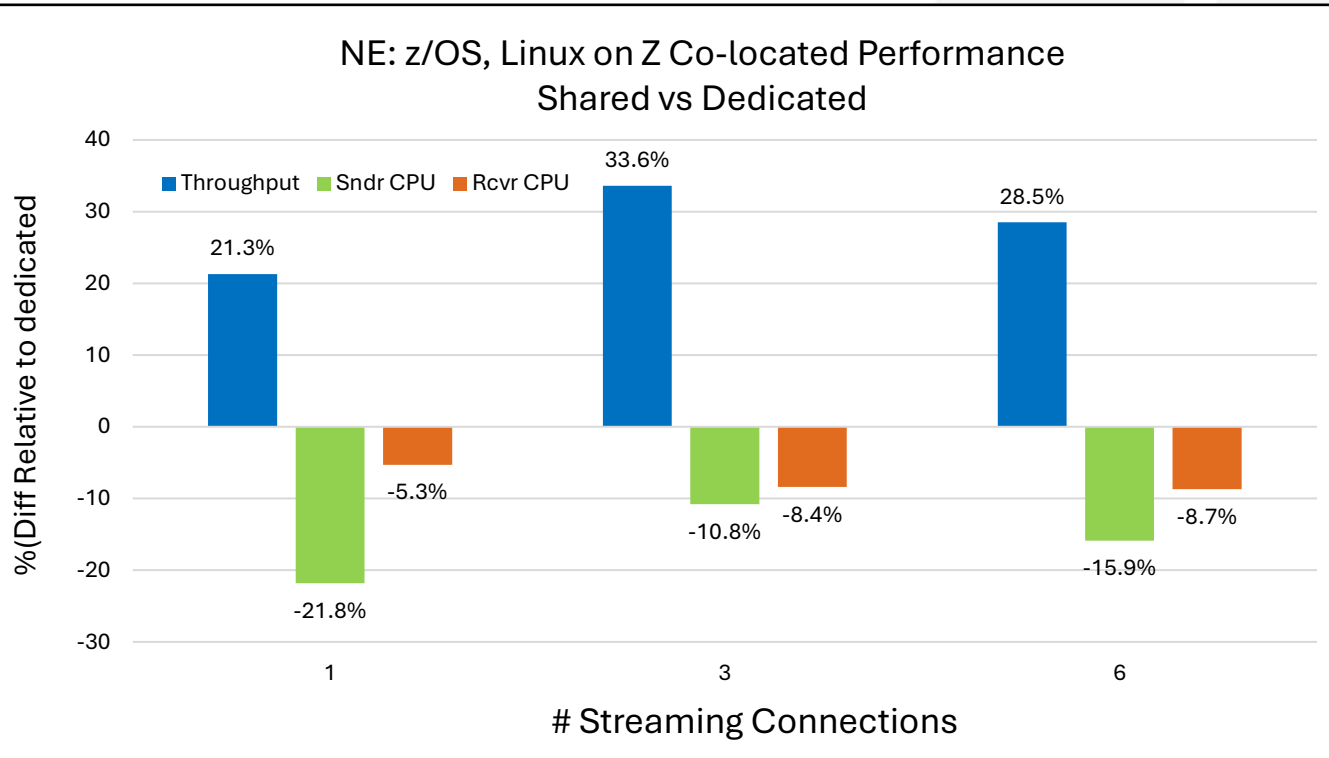
*Note: NETH on Linux on Z doing IP traffic*

# Shared NE: Interactive Workload Performance



- Chart shows how shared NE (same port) compares to dedicated adapters between z/OS and Linux on Z
- Aside saving on one card, shared NE shows a moderate reduction on transactional latency up to **13%**
- There is no significant change for network related CPU cost

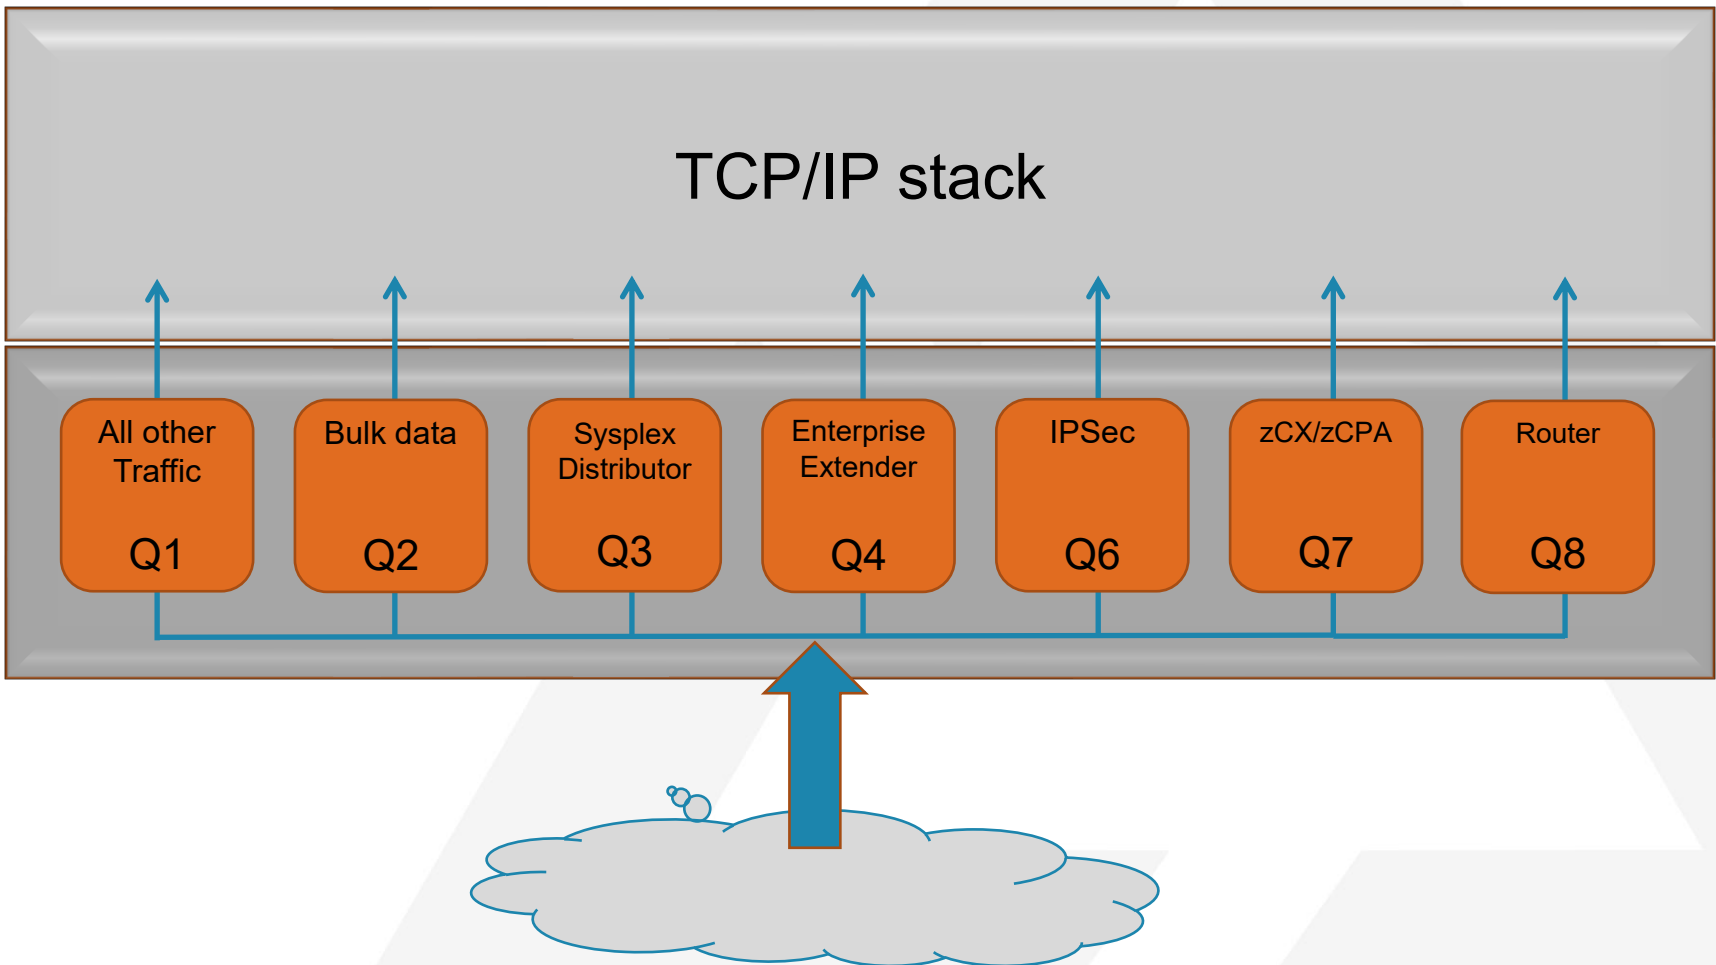
# Shared NE: Streaming Workload Performance



- Streaming from z/OS server to Linux on Z client.
- Shared NE shines with streaming bulk workloads compared to dedicated
- Significant throughput increase, up to **33%** improvement.
- Good CPU savings, up to **21%** reduction on the sending side, and **8%** on the receiving side


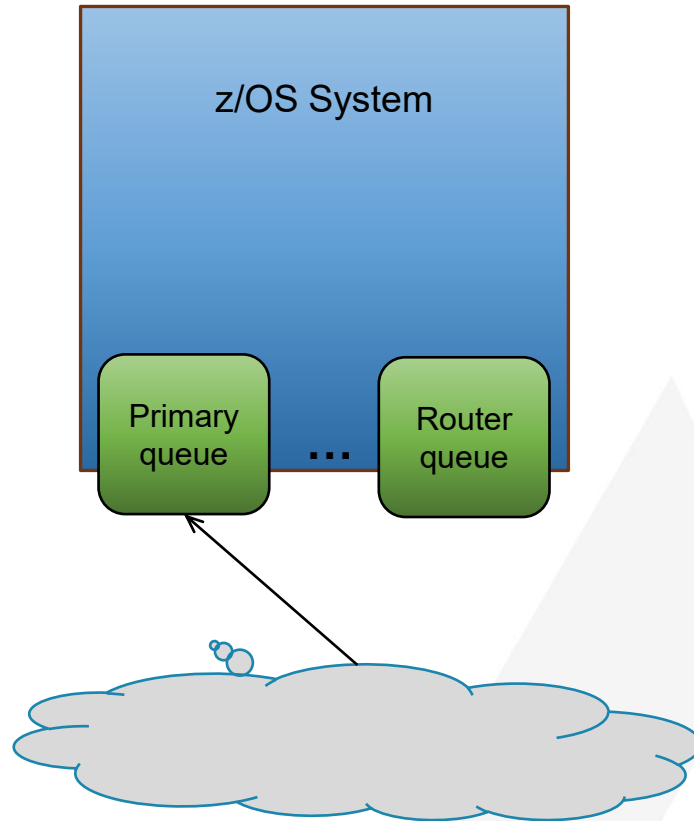
# IP Router Input Queue

# Inbound Workload Queuing



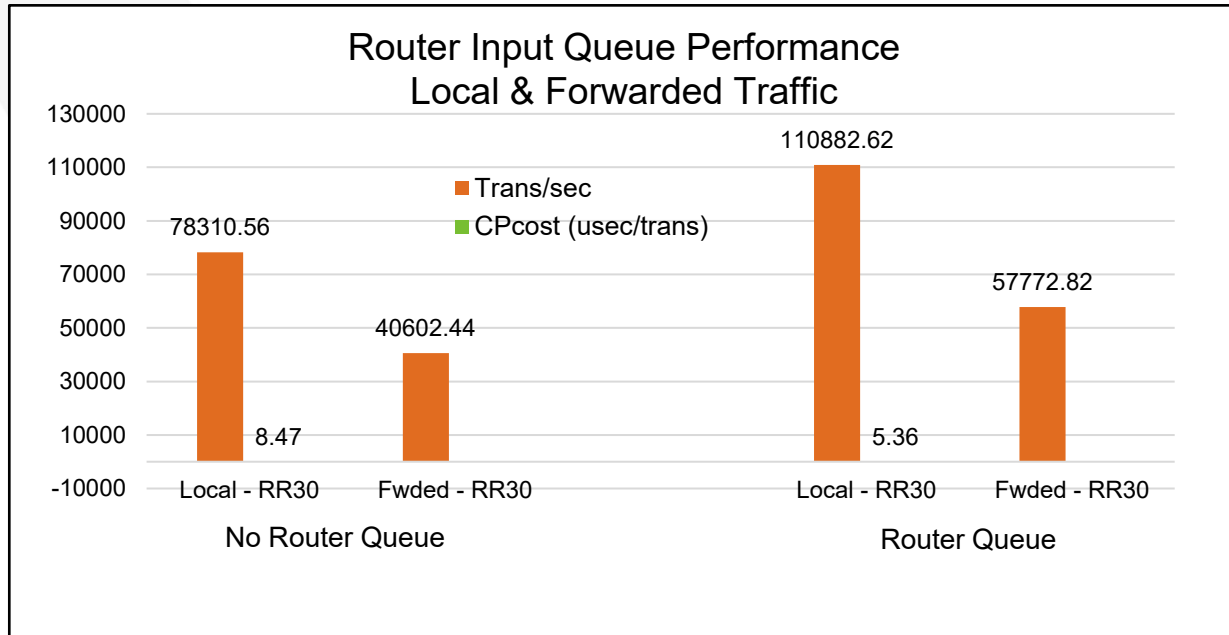
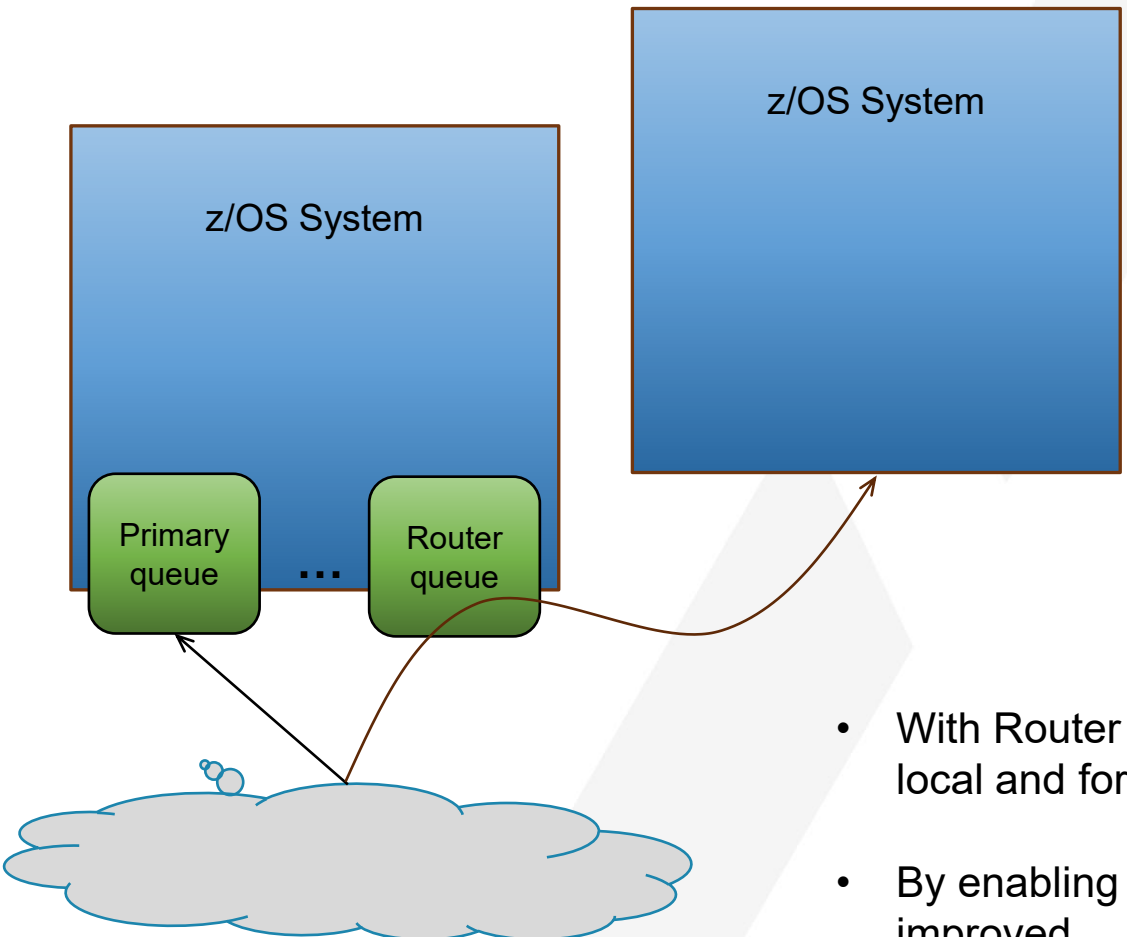
- Splitting 'routed' traffic away from Q1 (primary queue) improves processing of primary queue
- Enabled automatically for OSH interfaces if:
  - DATAGRAMFWD configured on IPCONFIG statement
  - ROUTEALL is configured on INTERFACE statement
  - QDIOACCELERATOR configured on IPCONFIG statement

# Router Input Queue Performance



- Avoiding look ups for every packet should translate into reduction in latency and CPU savings
- By enabling Router input queue, network response times for interactive traffic reduced by 18% and networking CPU reductions of 21%

# Router Input Queue Performance...

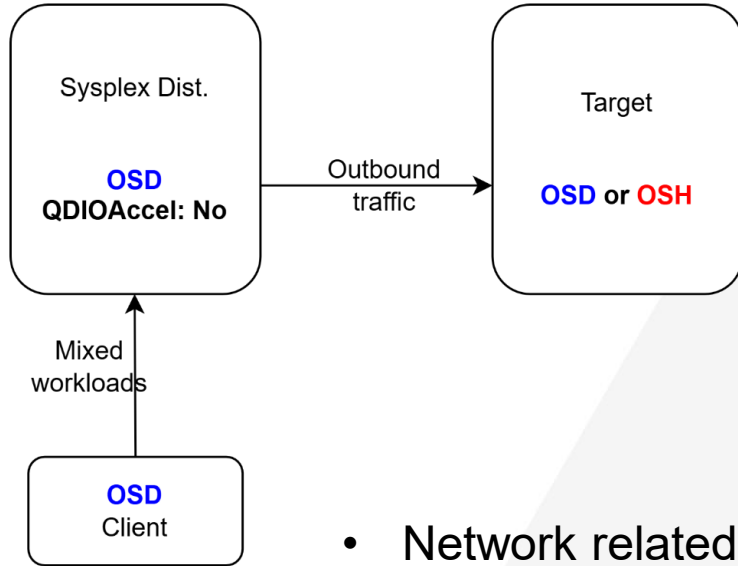


- With Router input queue, transaction rates should increase due to separation of local and forwarded traffic
- By enabling Router input queue, transaction rates for interactive traffic significantly improved
  - Local traffic improved by 41%
  - Forwarded traffic shows a 42% improvement
  - Networking CPU reduced by 36% on “routing” system

# GRE IWQ Performance

# GRE IWQ Performance

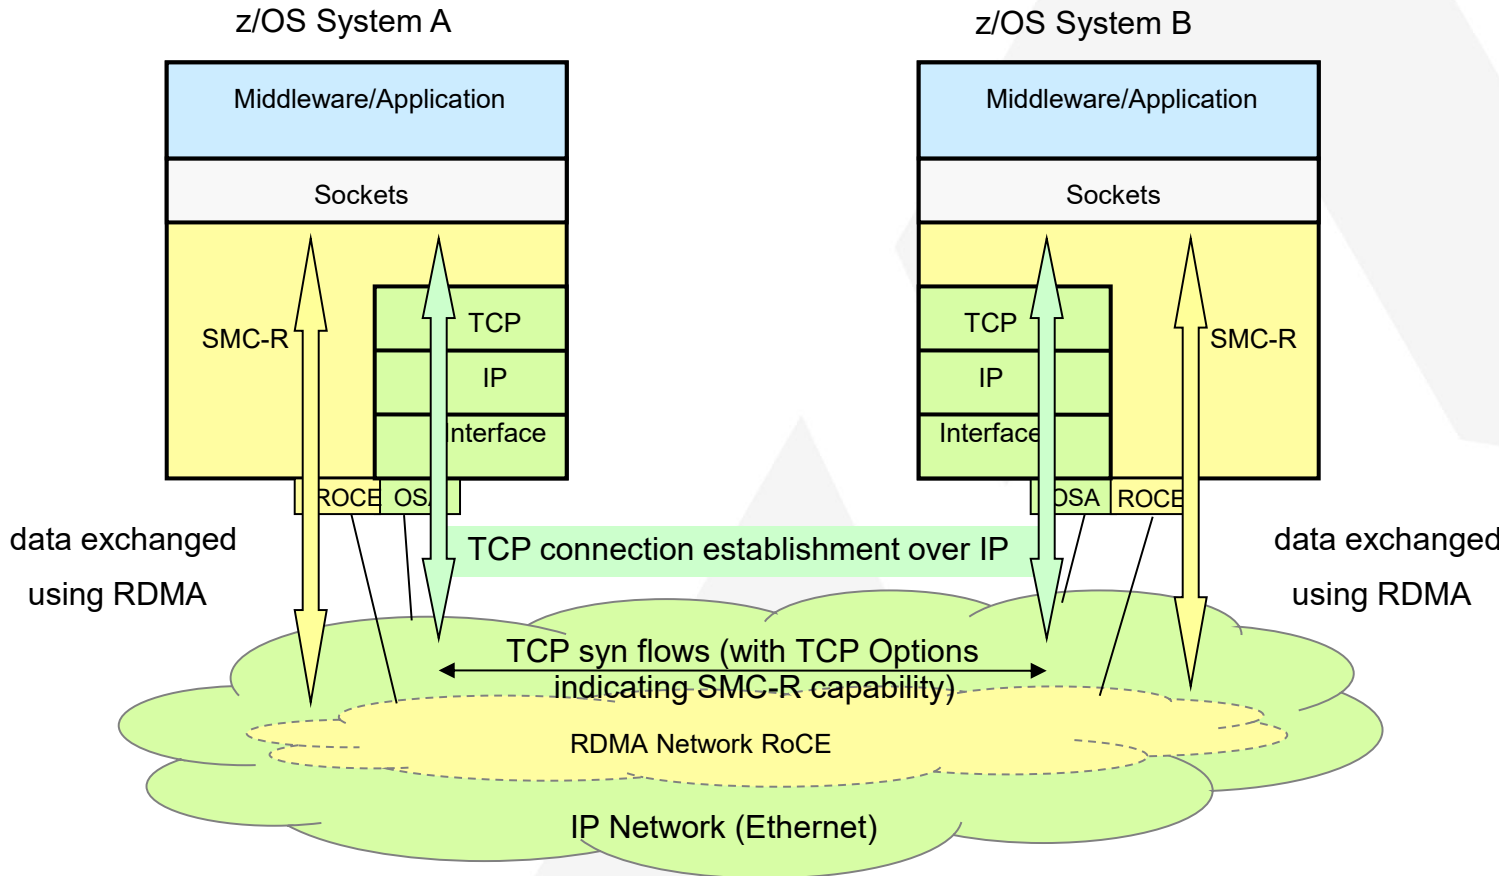
- Generic Routing Encapsulation (GRE) allows you to override the routing behavior of a packet by putting a different IP header in front of it
- Viparoute Sysplex distributor stack uses GRE to forward traffic to target stack
- New enhancement allows GRE encapsulated bulk traffic to be put on its queue



- To test GRE IWQ enhancement, we use OSD to distribute traffic to target OSH or OSD stack
  - We disable QDIO accelerator on the distributing stack for fair comparisons
  - Mixed workloads (interactive and stream traffic) running at the same time to target stack
- 
- Network related CPU cost on target stack reduce by **9%** due reduced out-of-order packets
  - Distributor stack shows **4%** reduction in CPU cost
  - About **1%** reduction in transactional latency

# NETH (SMC-R) performance

# SMC-R Overview



- ✓ **Optimized Network Performance (leveraging RDMA technology)**
- ✓ Transparent to (TCP socket based) application software
- ✓ Leverages existing Ethernet infrastructure (RoCE)
- ✓ Preserves existing network security model
- ✓ Resiliency (dynamic failover to redundant hardware)
- ✓ Preserves existing IP topology


Dynamic (in-line) negotiation for SMC-R is initiated by presence of TCP Options

TCP connection transitions to SMC-R allowing application data to be exchanged using RDMA

# NETH SMC-Rv2 Performance

- Network Express (NE) feature configured Functional ID (FID) type NETH
- SMC-Rv2 protocol using NETH on z17 versus OSD + RoCE on z16
- For interactive workload, SMC-Rv2 over NETH shows **10%** reduction in transactional latency compared to same protocol on z16
- For streaming workload, there is no net throughput difference but SMC-R over NETH averages **10%** CPU cost reduction compared to OSD + RoCE on z16

# SMC-Rv2 Performance: One Card Versus Two Cards

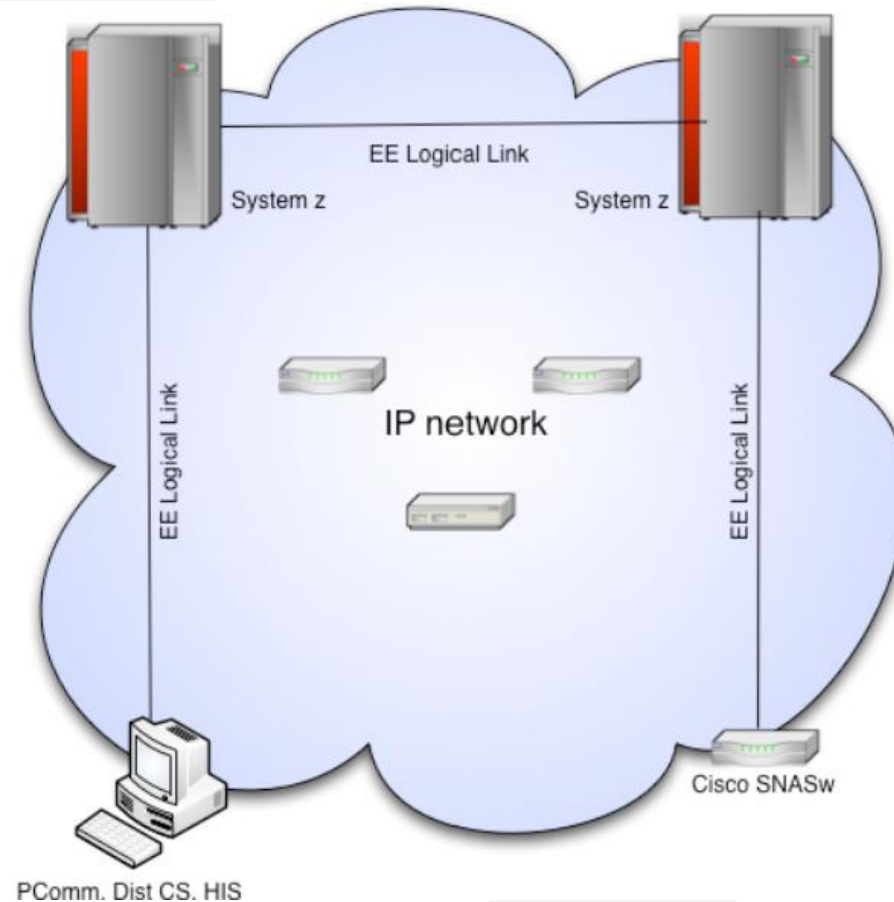


- TCP traffic over OSH running on port 1
- SMC-R traffic over NETH on port 2
- Mixed Workloads – interactive over TCP and streaming over SMC-R
- Chart shows there is no significant difference in transactional volume
- One card running two different protocol is as good as using two separate cards for the same work


# Enterprise Extender (EE) performance

# Enterprise Extender Overview

- Allows use of IP network for SNA sessions
- EE allows enablement of IP applications and convergence on a single network transport while preserving SNA application and endpoint investment.
- EE is APPN HPR routing over an IP network
  - To the IP network, EE looks like a UDP application



# Network Express OSH: EE Performance

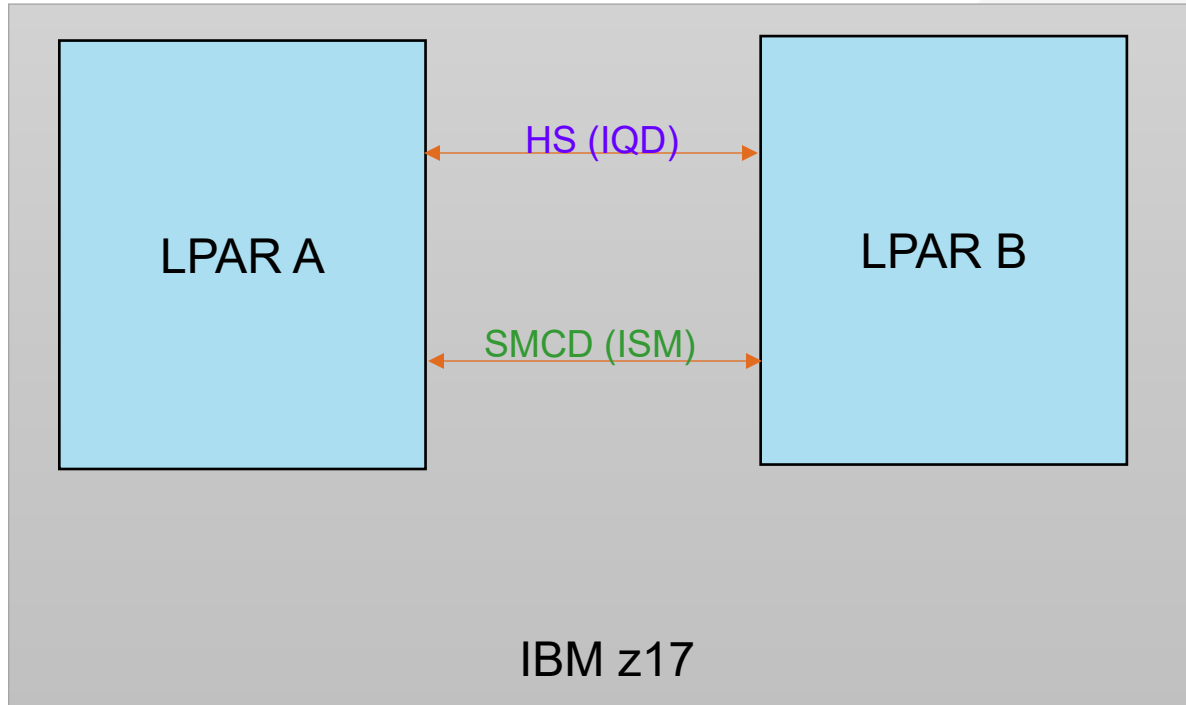


- EE protocol over OSH compared to OSD
- Small, medium, and large size RR workloads
- Significant reduction in response time across all payload sizes
  - Up to **55%** reduction compared to OSD
- No significant network related CPU difference



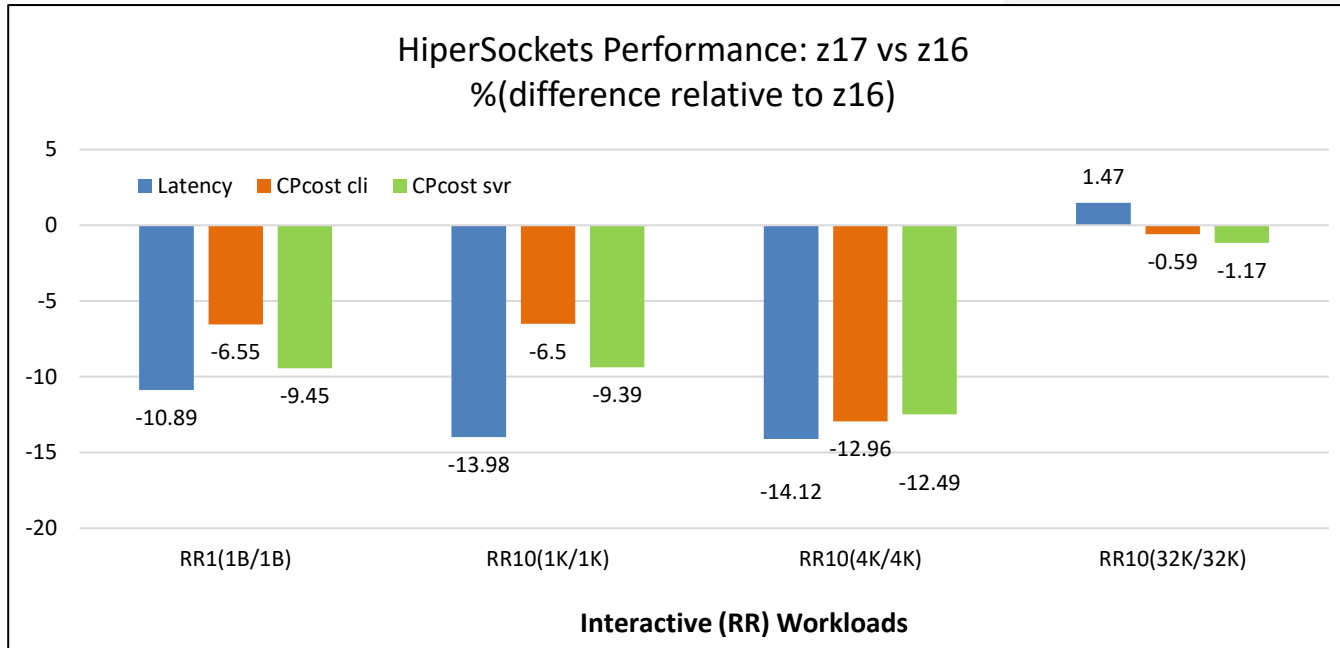
# z/OS HiperSockets and SMCDv2 Performance on z17

# HiperSockets and SMCD Overview



- HiperSockets (HS) technology provides high-speed TCP/IP connectivity within a CPC
- HS communication is through system memory of the processor, so systems are connected to form an “internal LAN”
- HS implementation based on Queued Direct I/O (QDIO) protocol, hence HS is also called internal QDIO or IQDIO
  
- Shared Memory Communications – Direct Memory Access (SMCD) provides highly optimized communication within a CPC
- Uses an internal Shared Memory (ISM) virtual PCI that enables direct access to shared virtual memory
- SMCD can use OSA or HiperSockets associated with the ISM device
- ISMv2 allows SMCD over different IP subnets on the same CPC

# HiperSockets Performance – z17 vs z16




- Interactive workloads (Request-Response) over HiperSockets
- Transactional latency and CPU cost per transaction comparison
- Low to medium payloads
  - Up to 14% reduction in latency
  - Up to 12% networking CPU savings
- High payloads
  - No significant change in latency and networking CPU

*Reduced response times and CPU savings on z17 is a result of faster CPU and larger processor cache*

# SMCDv2 Performance – z17 vs z16

SMCDv2 Performance z17 vs z16  
%(difference relative to z16)




Interactive (RR) Workloads

- Interactive workloads (Request-Response) over SMCD
- Transactional latency and CPU cost per transaction comparison
- Low to medium payloads
  - Up to 23% reduction in latency and 10% networking CPU savings
- High payloads
  - About 15% reduction in latency and networking CPU cost

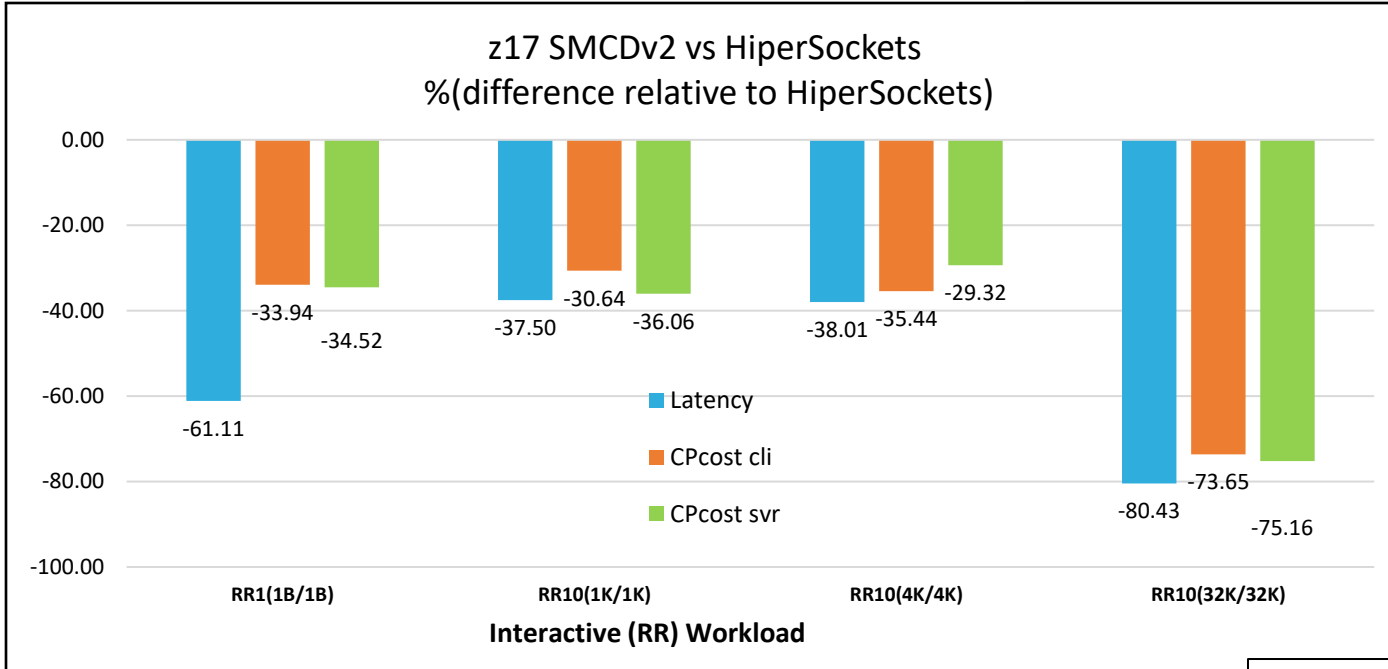
- Streaming workloads (20MB “puts”)
- About 16% average improvement in throughput
- Sender networking CPU cost reduced by up to 17%
- Receiver networking CPU cost reduced up to 13%

SMCDv2 Performance: z17 vs z16  
%(difference relative to z16)



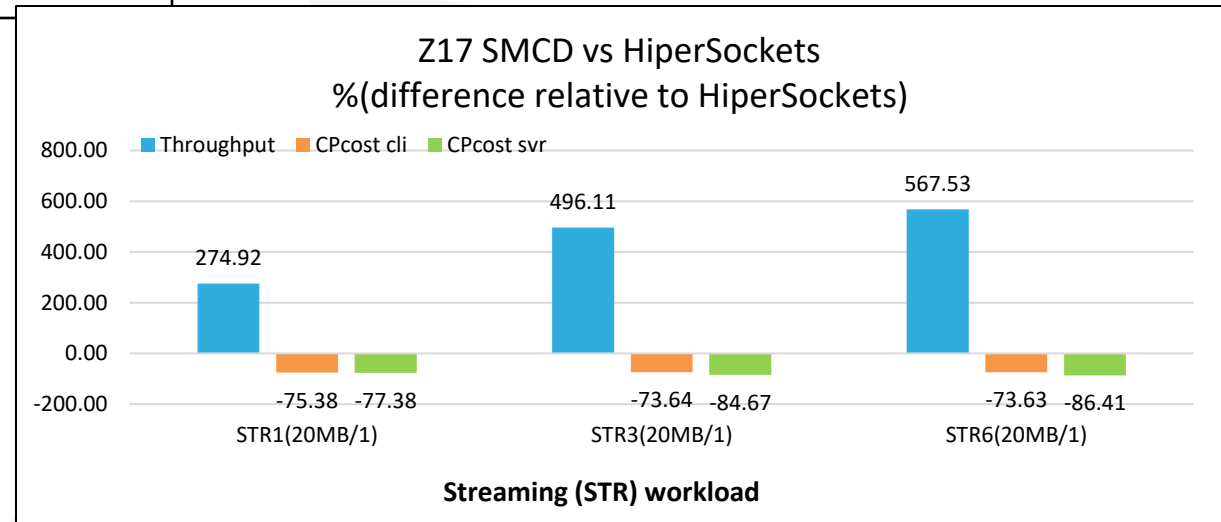
Streaming (STR) Workloads

# HiperSockets vs SMCDv2 Performance



- Interactive workloads with 1k/1k & 4k/4k payloads
  - Latency: up to 38% reduction in latency
  - CPU cost: up to 36% reduction in networking CPU cost
- Interactive workloads with 32k/32k payload
  - Latency: up to 80% reduction in latency
  - CPU cost: up to 75% reduction in networking CPU cost

- Streaming workloads (20MB “puts”)
- Throughput: up to 567% (~6x) increase in throughput
- CPU cost: up to 86% reduction in networking CPU cost





# z/OS Communications Server Best Practices

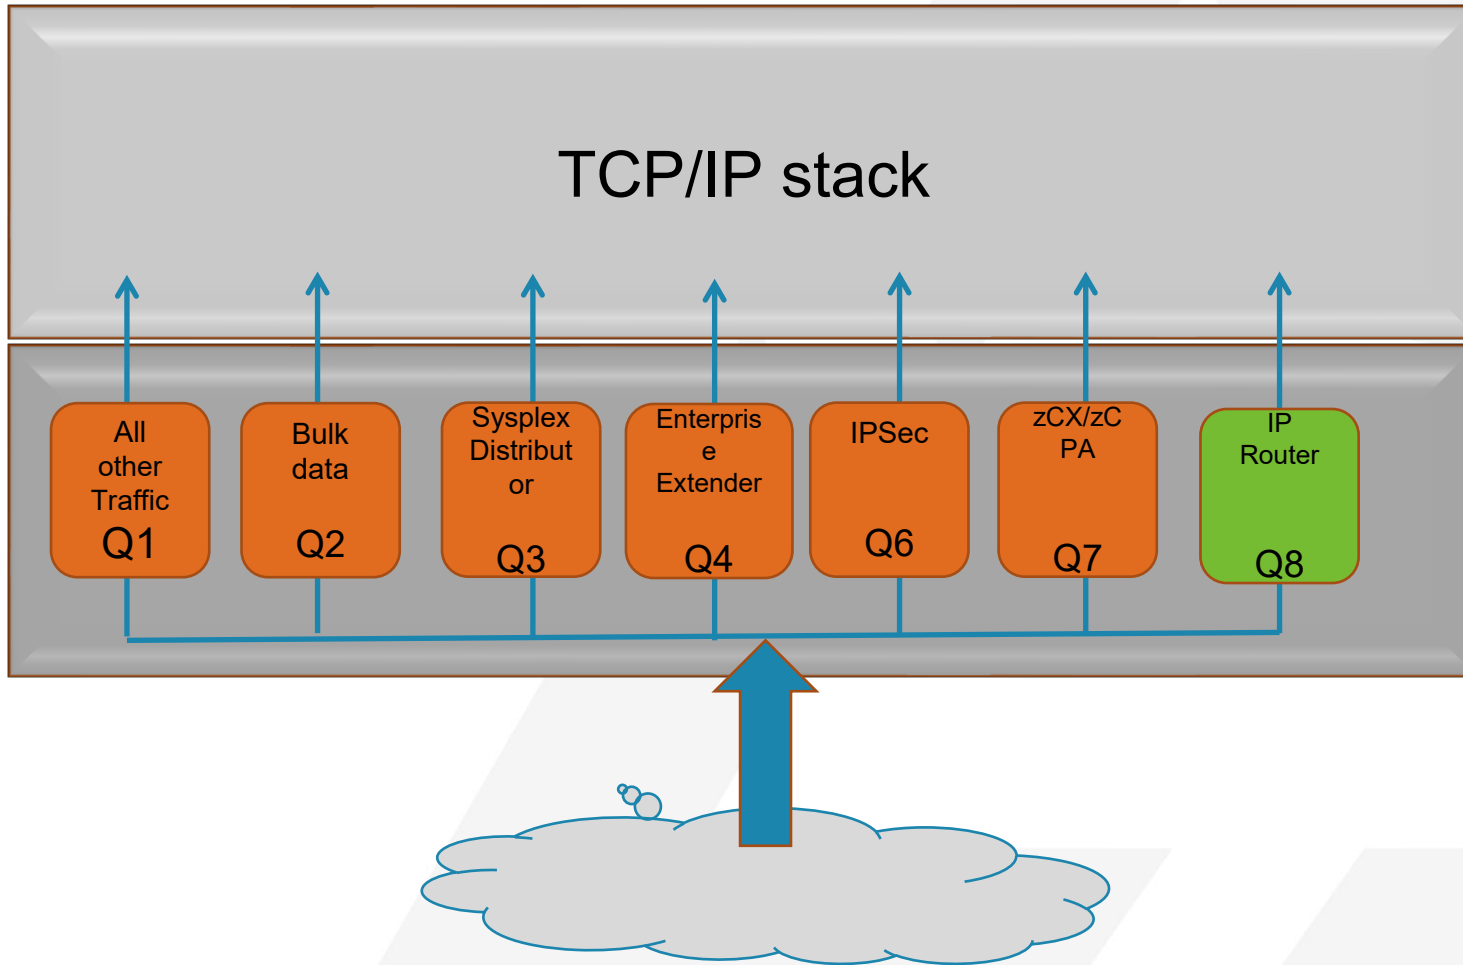
# Inbound Workload Queuing configuration

- INBPERF DYNAMIC WORKLOADQ enables OSA Inbound Workload Queuing (IWQ)

```
>>-INTERFace--intf_name----->
.
.-INBPERF BALANCED-----
>+-----+>
|                .-NOWORKLOADQ-. |
\ -INBPERF+-DYNAMIC+-----+ +-'
|                \-WORKLOADQ---' |
+-MINCPU-----+
\ -MINLATENCY-----'
```

- For Streaming workloads, enable on z/OS endpoints (sender and/or receiver)
- **ONLY APPLIES TO OSD OSA's – NETWORK EXPRESS (OSH) ALWAYS USES IWQ!**

# Inbound Workload Queuing ...



## Benefits:

- Prevents bulk traffic from starving out OLTP traffic
- Reduces out-of-order packets for streaming traffic
- Enables SD and EE traffic to take optimized path directly to SD targets or VTAM
- Enables IPSec and zCX/zCPA traffic to immediately be scheduled onto zIIP processors

## TCP Segmentation Offload configuration


- Enabled with IPCONFIG/IPCONFIG6 SEGMENTATIONOFFLOAD

```

>>-IPCONFIG----->
.
.
>-----+-----+-----<
| .-NOSEGMENTATIONOFFLoad- |
+--+-----+-----+
| '-SEGMENTATIONOFFLoad---' |
  
```

- Disabled** by default (previously enabled via GLOBALCONFIG)
- Segmentation cannot be offloaded for
  - Packets to another TCP/IP stack sharing OSD OSA port
    - Segmentation offload is enabled for shared OSH OSA port
  - IPSec encapsulated packets
  - When multipath is in effect (unless all interfaces in the multipath group support segmentation offload)

# TCP Segmentation Offload performance



➤ **Segmentation offload can significantly reduce network CPU cycles when sending bulk data from z/OS!**

## z/OS Communications Server - A few more suggestions/tips

- Sysplex Distributor:
  - Configure VIPADynamic VIPAROUTE with GLOBALCONFig ADJUSTDVipamss
    - Improved networking performance both in latency and CPU usage
    - Eliminates fragmentation
  
- TCP configuration:
  - TCPCONFIG AUTODELAYAcks
    - Optimizes ACK transmission based on workload patterns
  - TCPCONFIG SELECTIVEACK
    - Can reduce unnecessary packet retransmissions

## z/OS Communications Server - A few more suggestions/tips...

- Streaming data over shared OSA Express:
  - Do not stream when source/destination reside in same subnet with IWQ-enabled
    - Contention with outbound/inbound paths within OSA can negatively impact throughput
  - Recommendations
    - Use SMC-D
    - Use SHARED OSH
    - Use HiperSockets
    - Move source and destinations into different subnets

**RECOMMENDATION APPLIES ONLY TO OSD OSA'S!**

# z/OS Communications Server Performance Summaries

- Performance characteristics of each z/OS Communications Server release is studied by an internal performance team
- Summaries are created and published online
  - <http://www-01.ibm.com/support/docview.wss?rs=852&uid=swg27005524>
- z/OS 3.2 Communications Server Performance Summary
  - Will be available in few days
- White paper on OSA-Express best practices
  - Provides a general set of considerations (a checklist) for guidance focused on configuring OSA-Express for optimizing network performance
    - <http://ibm.biz/OSACSBP>



# z/OS 3.2 AT-TLS Performance



# AT-TLS Performance – Tips for non-persistent connections

- Utilize Crypto Express cards
  - Offload CPU overhead of performing asymmetric encryption processing for TLS/SSL handshakes
    - Networking latency for non-persistent connections reduced by up to **42x**
    - Networking CPU costs reduced by over **94x**
- Configure Session ID (SID) or Session Ticket Caching
  - Allows repeated connections from same clients to perform an optimized TLS/SSL handshake
  - Reduces number of round-trip network flows needed for TLS/SSL handshakes to 1 (from 2 round-trips required for a full handshake)
    - Networking latency for connections using session caching reduced by an additional **25%**
    - Networking CPU costs reduced by up to **25%**
- Enable LE Runtime Option HEAPPOOLS64
  - Removes contention for accessing user heap storage in an AT-TLS environment
  - HEAPPOOLS64 becomes a factor with TLSv1.3 under very low handshake volumes
  - z/OS Comm. Server **APAR (PH59425)** ensures HEAPPOOLS64 is always enabled
    - Networking latency for connections using session caching reduced by an additional **63%**
    - Networking CPU costs reduced by up to **40%**


# TLS v1.3 versus TLSv1.2 Performance

```
ConnID: 00000037
JobName: APF1
LocalSocket: 10.67.170.128..5001
RemoteSocket: 10.67.170.126..1034
SecLevel: TLS Version 1.2
Cipher: C030 TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
KeyShare: N/A
CertUserID: N/A
MapType: Primary
FIPS140: Off
SessionID: 00000094 0943AA7E 040A0000 00000000
           00000000 00000000 65AAA842 00000009
SignaturePairs: 0804 TLS_SIGALG_SHA256_WITH_RSASSA_PSS
                0401 TLS_SIGALG_SHA256_WITH_RSA
```

```
ConnID: 000000C6
JobName: APF7
LocalSocket: 10.67.170.128..5001
RemoteSocket: 10.67.170.126..1065
SecLevel: TLS Version 1.3
Cipher: 1302 TLS_AES_256_GCM_SHA384
KeyShare: 0023 secp256r1
CertUserID: N/A
MapType: Primary
FIPS140: Off
SessionID: 00000035 0943AA7E 04290000 00000000
           00000000 00000000 65AA9742 00000009
ClientKeyShareGroups: 0023 secp256r1
ServerKeyShareGroups: 0023 secp256r1
SignaturePairs: 0804 TLS_SIGALG_SHA256_WITH_RSASSA_PSS
```

- We impose much similarities in algorithm being used between TLSv1.2 and TLSv1.3 protocol (both using ECDHE, RSASSA-PSS, AES\_256\_GCM, SHA384)
- Key type RSA and key size 2048
- HEAPPOOLS64 LE option enabled for AT-TLS by default

# TLS v1.3 versus TLSv1.2 Performance chart




- Short-lived connection benchmark – CRR40(64B/8K)
- All the increased CPU cost of the TLSv1.3 handshake is seen on the client side
- There is on net increase in the server-side CPU
- Little to no performance impact for persistent connections

*When using similar cryptographic algorithms, TLSv1.3 performance is competitive with that of TLSv1.2*

# DH Group Size, ECDH Curve Performance

## ECDHE, DHE Key Exchange Performance Comparison



- TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (C030) Curve –
- TLS\_DHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (009F)
- DH Group (1024 vs 2048) vs ECDH curve 256
- ECDHE key exchange accelerated in hardware. DHE key exchange is done software
- Conclusion: ECDHE key exchange operations are performed in CPACF providing great performance as opposed to DHE key exchange, which is done completely in software. Best to prioritize ECDHE cipher suites higher than DHE.

*ECDHE key exchange achieves 878% (~9x) more handshakes per second, and 98% lower CPU cost compared to DHE key exchange. Best to prioritize DHE cipher suites towards the bottom of your cipher suite list*

# Performance Impact of Session Object Audit

- ICSF options AUDITKEYLIFETKDS and AUDITPKCS11USG
- Controls auditing lifecycle and events related to PKCS #11 services
- SESSIONOBJ 'YES' creates, use, or delete ephemeral key
- Recommend setting SESSIONOBJ to 'NO'

OPTIONS	Handshakes per sec	TLS client CPU cost per tran	TLS server CPU cost per tran
TOKOBJ(YES) SESSOBJ( <b>NO</b> )	3067	282	261
TOKOBJ(YES) SESSOBJ( <b>YES</b> )	645	5790	3260

- Cipher suite – 1301 (TLSv1.3)
- 40 concurrent short-lived connections
- Setting SESSOBJ to YES for PKCS #11 significantly impact performance:
  - **78%** reduction in handshake volumes
  - Up to **20x** increase in network CPU cost

- *Session objects are intended to be used for ephemeral keys*
- *Actions for a single TLS session is expensive and only gives like debug data*

# z/OS Communications Server TLS Performance Paper

- TLS Performance White Paper Published in February
- Contains good information on the network cost associated with:
  - ECDHE, DHE, TLS\_RSA based cipher suites
  - TLSv1.3
  - Session reuse
  - HEAPPOOLS benefits
  - FIPS 140-2 and more
- Link to paper - <https://www.ibm.com/support/pages/node/317829>

Questions?

# Experience more with IBM



## [Visit us at the IBM Booth #113](#)

After a full day of technical sessions, take a break with us!

Connect with our experts, snap a photo with the z17 Plexi or the latest Telum II, and get an up-close look at our Spyre Accelerator.

Come back each day for fresh topics and demos at our expert stations.

## [Think 2026](#)

Join 5000+ senior business and technology leaders who are seizing the AI revolution to unlock unprecedented growth and productivity at **Think 2026**.


Find out more information using the QR code below.



## [IBM Digital Asset Haven](#)

IBM Digital Asset Haven is the operational backbone for financial institutions and regulated enterprises entering the digital asset economy.

Find out more information using the QR code below.



# Your feedback is important!

## Submit a session evaluation for each session you attend:

[www.share.org/evaluation](http://www.share.org/evaluation)



Thank you

# IBM z/OS Communications Server

A high availability enterprise transaction and data server built for optimum productivity



Comm Server Community  
(<https://ibm.biz/cscommunity>)  
Rich and up-to-date technical content, including blogs, videos, and events.

**High-speed connectivity**



**Network Management**

**High availability**



**IPv6 enablement**

**Network Security**



**Simplification and usability**

# Digital Badges & Online Courses

Start your **free IBM Z online learning** and earn IBM open badges!



## Networking on z/OS - Foundations

Foundational understanding of networking on z/OS.

- **IBM Open Badge:** <https://ibm.biz/zosnetworkingbadge>

- **Online course:** <https://ibm.biz/zosnetworkingcourse>



## z/OS TCP/IP Configuration with NCA

Use the IBM Configuration Assistant for z/OS Communications Server (NCA) to create and manage TCP/IP profiles.

- **IBM Open Badge:** <http://ibm.biz/NCAbadge>

- **Online course:** <http://ibm.biz/NCATCPIPcourse>




## z/OS Network Security - Foundations

Knowledge and foundational understanding of z/OS network security.

- **IBM Open Badge:** <http://ibm.biz/zosnetsecuritybadge>

- **Online course:** <http://ibm.biz/zosnetsecuritycourse>




## zERT Policy Enforcement Configuration with NCA

Configure zERT Policy Enforcement using the IBM Configuration Assistant for z/OS Communications Server (NCA)

- **IBM Open Badge:** [http://ibm.biz/NCA\\_zERTbadge](http://ibm.biz/NCA_zERTbadge)

- **Online course:** [http://ibm.biz/NCA\\_zERTcourse](http://ibm.biz/NCA_zERTcourse)



## TCP/IP on z/OS Essentials - Level 1

General knowledge and understanding of TCP/IP on z/OS, including network layers, protocols at each layer, and the hardware that facilitates the transport of data.

- **IBM Open Badge:** <http://ibm.biz/tcpipl1badge>

- **Online course:** <https://ibm.biz/tcpipl1course>

# Notices and Disclaimers



- © 2025 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

# Notices and Disclaimers

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**

**IBM\***  
**IBM Logo\***

\* Registered trademarks of IBM Corporation

**The following are trademarks or registered trademarks of other companies.**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Docker and the Docker logo are trademarks or registered trademarks of Docker, Inc. in the United States and/or other countries. Docker, Inc. and other parties may also have trademark rights in other terms used herein.

\* All other products may be trademarks or registered trademarks of their respective companies.

**Notes:**

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.



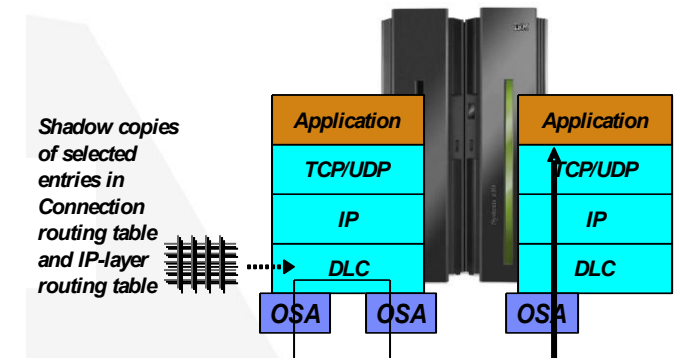
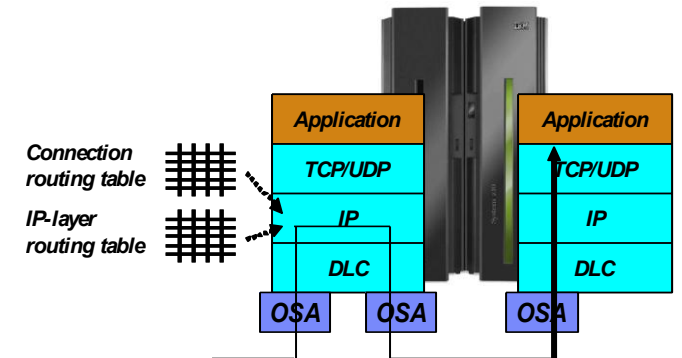
# Backup Slides



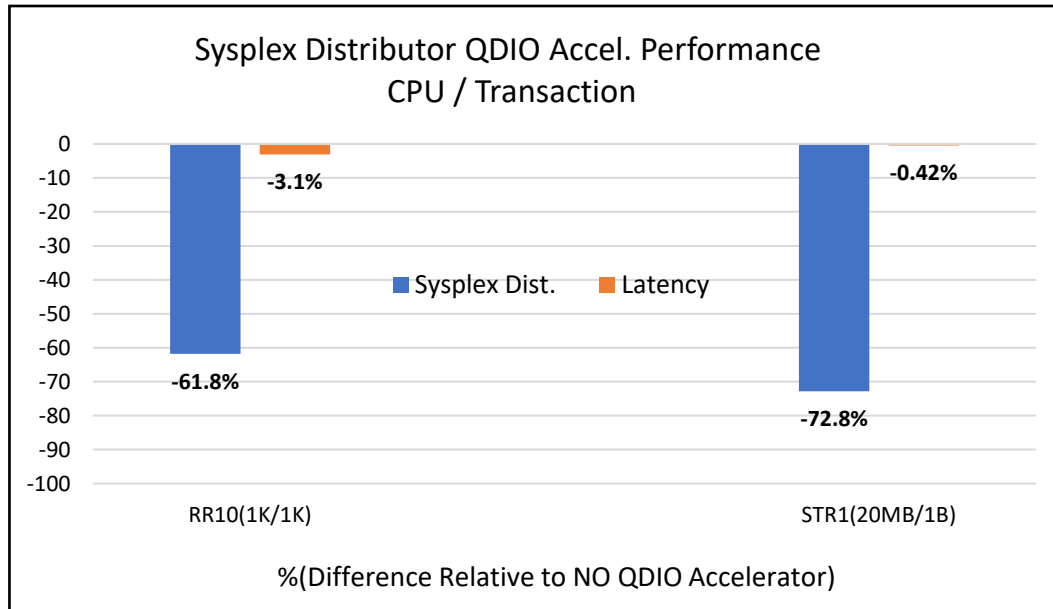
# SMC – Sysplex Distributor Performance

# QDIO Accelerator – Background

- Provides fast path IP forwarding for these interface combinations:
  - Inbound OSA, outbound OSA or HiperSockets
  - Inbound HiperSockets, outbound OSA or HiperSockets
- Sysplex Distributor (SD) acceleration
  - Inbound HiperSockets or OSA
  - When Sysplex Distributor gets to target stack using either
    - Dynamic XCF connectivity over HiperSockets
    - VIPAROUTE connectivity over OSA
  - Improves performance and reduces network CPU usage for such workloads

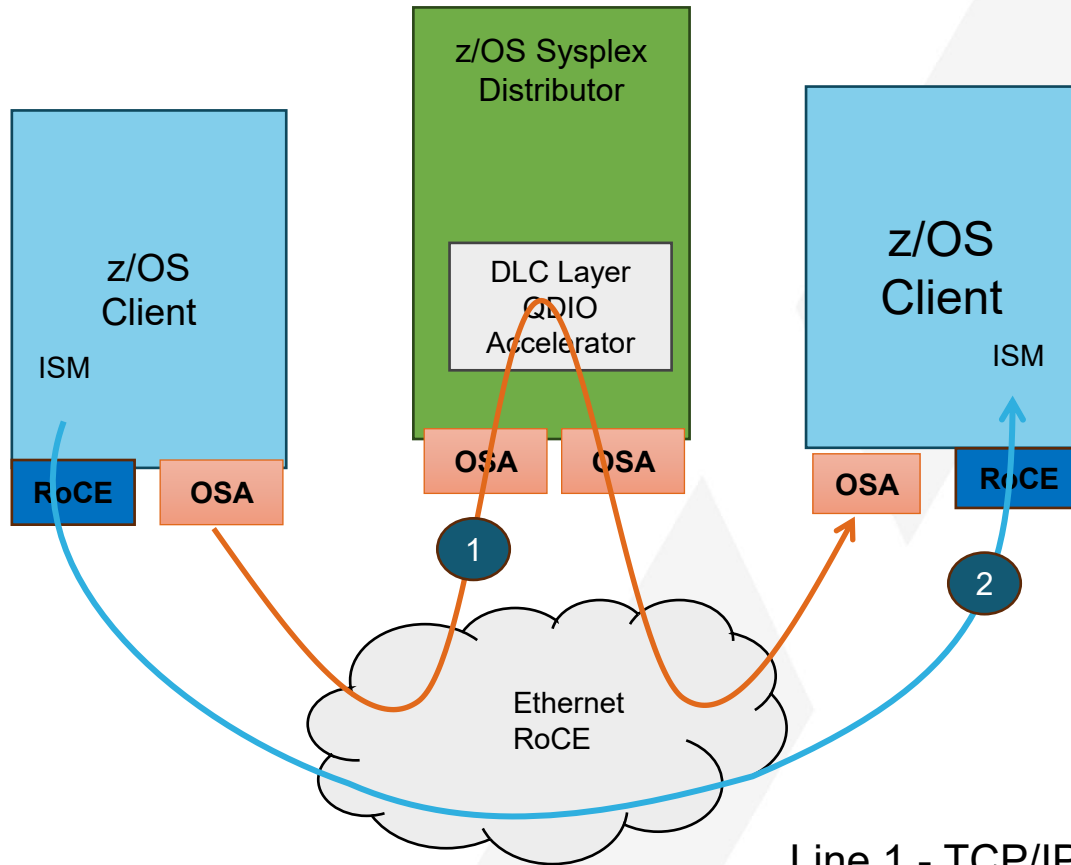


# QDIO Accelerator – performance results



- Request-Response workload – RR10(1K/1K)
- Streaming workload – STR1(20M/1B)
- Significant CPU savings on z/OS Sysplex distributor stack with QDIO acceleration enabled

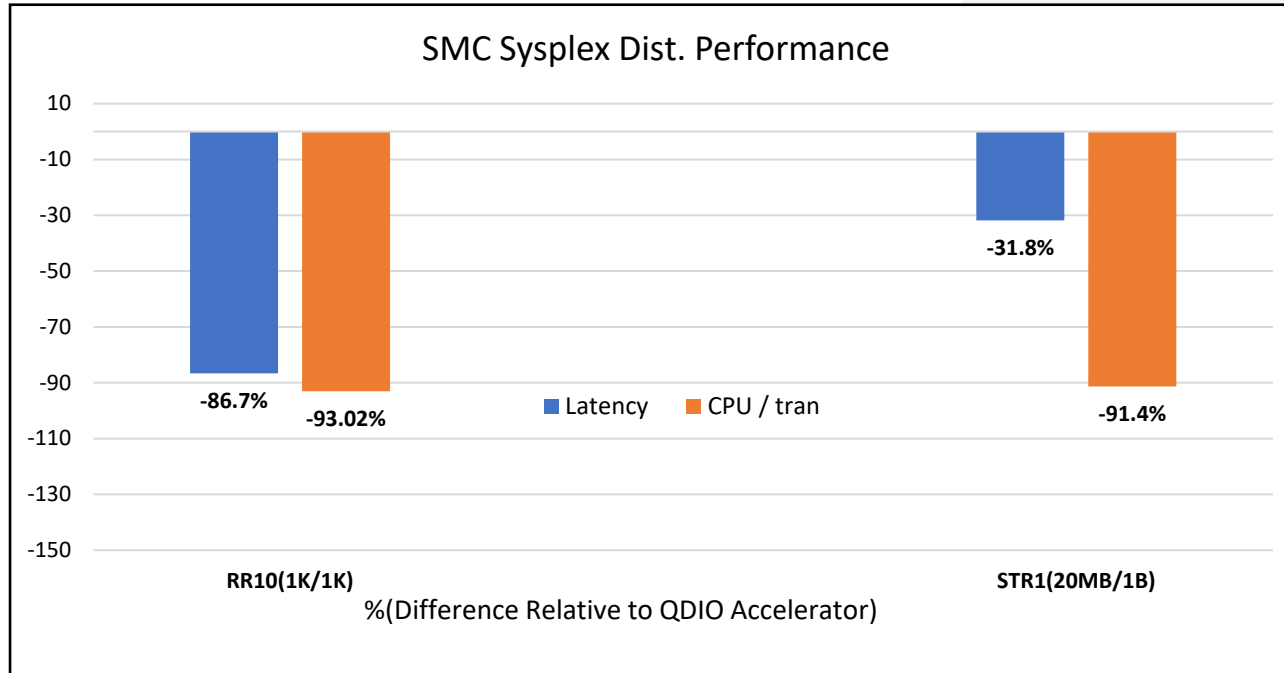
# SMC with Sysplex distributor function



- With SMC the distributing stack is bypassed for inbound data.
- Connection setup and SMC-R rendezvous packets will be the only inbound traffic going through the distributing stack.
- Remember that all outbound traffic bypasses the distributing stack for all scenarios.

Line 1 - TCP/IP distributed connections utilizing QDIO Accelerator  
Line 2 - SMC-R distributed connections

# SMC Sysplex distributor performance



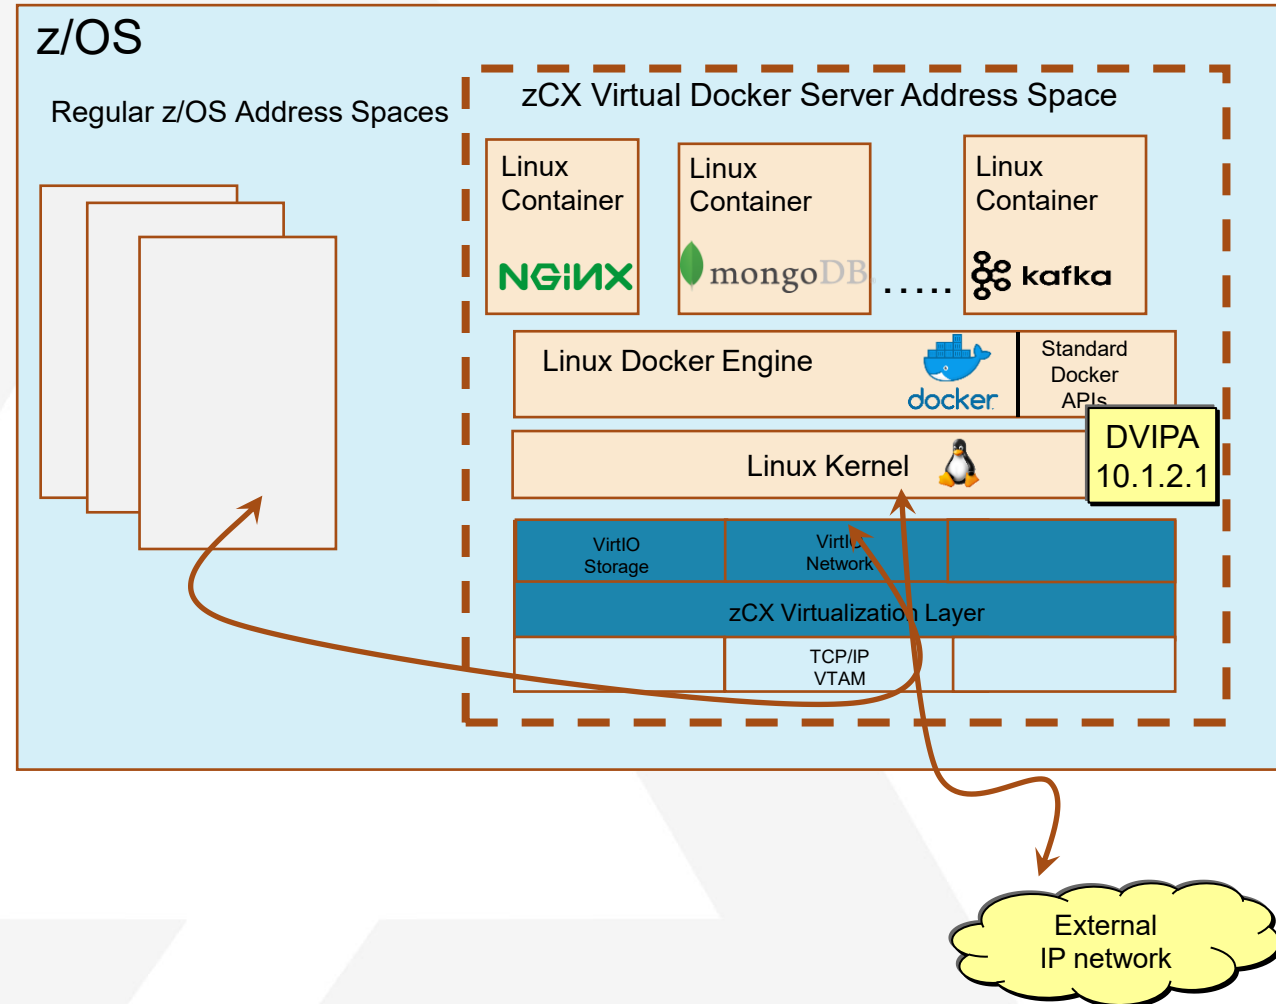
- Chart summarizes how SMC performance compares with QDIO Accelerator
- Results from Sysplex distributing stack perspective
- SMC removes virtually all CP processing on distributing stack – up to 93% savings
- SMC also improves latency, 86% reduction for RR and 31% for streaming workload.



# z/OS Containers Extensions (zCX) Network Recommendations

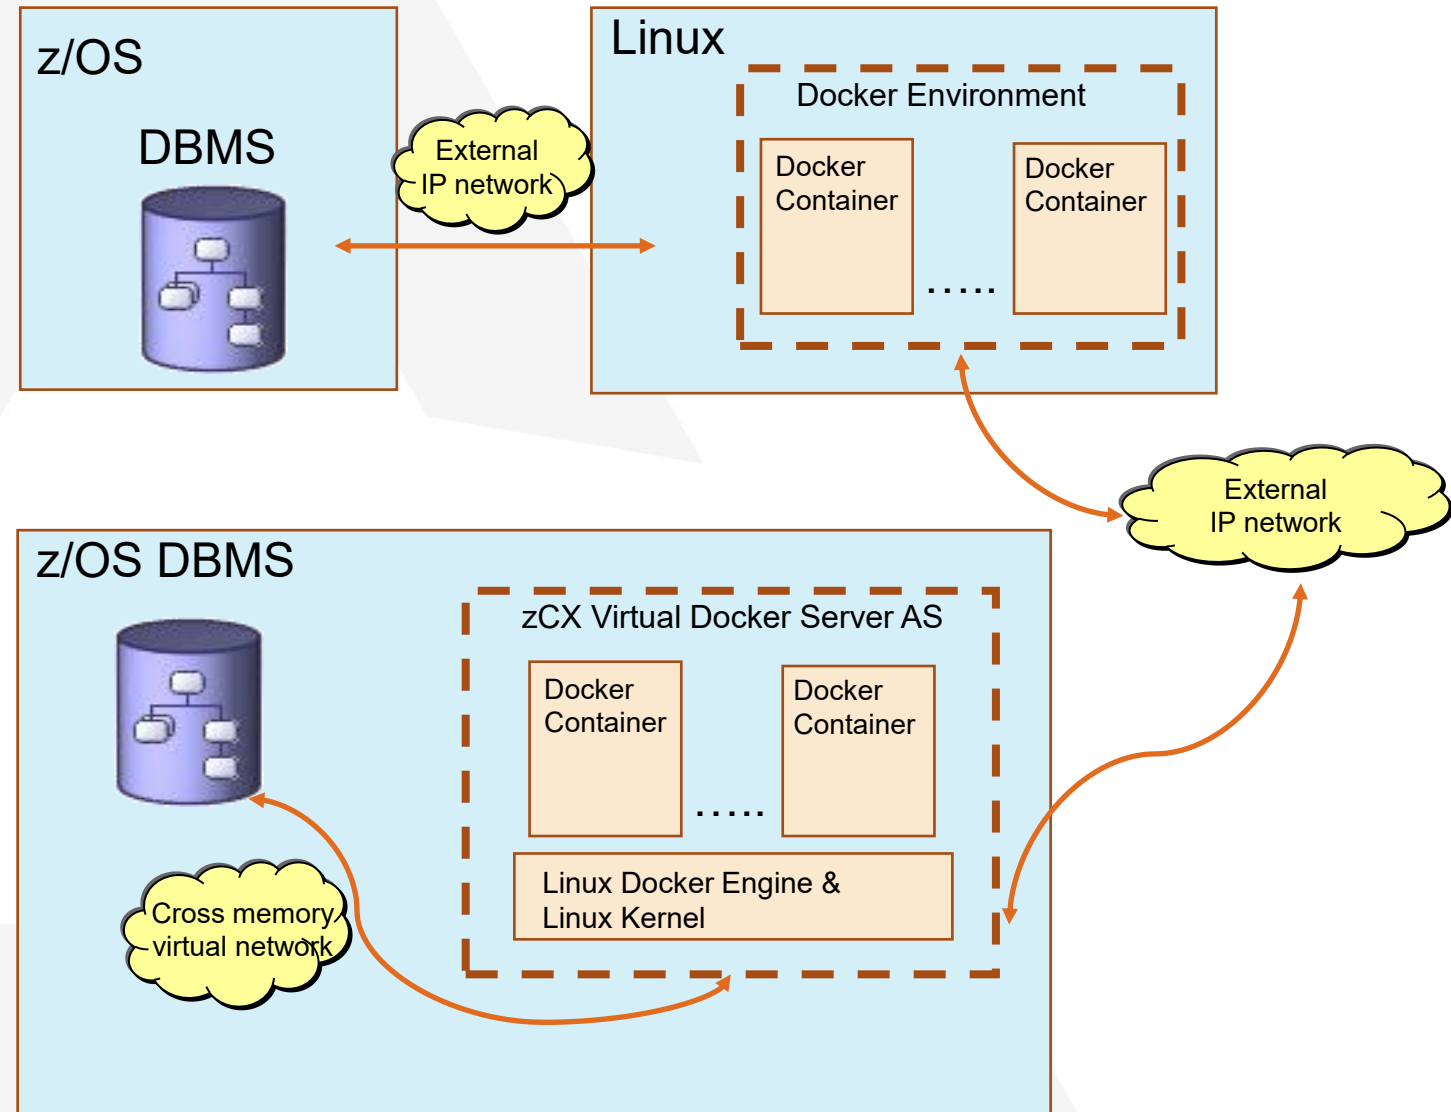
# IBM zCX – z/OS Network Integration

- z/OS Linux Virtualization Layer
  - Allows virtual access to z/OS Network
  - Using virtio Linux interfaces
    - Stable, well defined interfaces used to virtualize Linux
- Linux network access via high speed virtual *SAMEHOST* link to z/OS TCP/IP protocol stack
  - Each Linux Docker Server represented by a z/OS owned, managed and advertised Dynamic VIPA (DVIPA)
    - Allows restart of a CX instance in another system in the sysplex
  - Provide high performance network access across z/OS applications and Linux Docker containers – leveraging cross memory
    - All communications between zCX containers and z/OS applications over TCP/IP
  - External network access via z/OS TCP/IP
    - z/OS IP filters to restrict external access



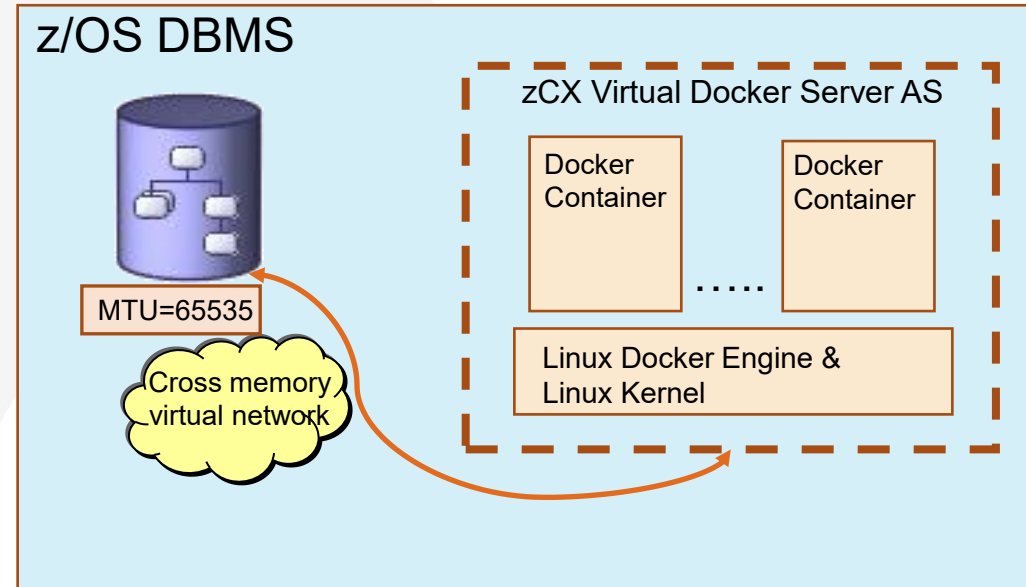
# IBM zCX – Moving Docker containers to zCX

- Application tier running in Docker Container on Linux server
  - All communication with Data tier must traverse external network
- Application tier running in Docker container within zCX
  - Co-locating Application tier with Data tier can significantly reduce network latency
  - Reduced network latency for interactive workloads by **45%** while increasing network transaction rates by **81%**
  - Reduced network latency for streaming workloads by **67%** while increasing throughput by over **200%**



# IBM zCX – Optimizing cross memory virtual network

- Virtual network not constrained to packet size limits imposed on physical networks
- When streaming data between the Application and Data tiers, using a larger MTU can provide significant benefits
  - Reduced network latency by **44%** while increasing throughput by **80%**
  - Reduced network related costs on GCPs by **34%** and by **60%** on zIIPs



# IBM zCX – Considerations for non co-located zCX

- Application tier and Data tier running in different z/OS LPARs
  - All communication with Data tier must traverse external network
- Configure Inbound Workload Queuing (IWQ) on OSA-Express
  - Better preserve order of packets delivered to zCX and utilize zIIPs for more network processing
  - Reduced network latency for interactive workloads by **26%** while improving network transaction rates by **34%**
  - Move nearly **40%** of network processing for interactive workloads to zIIPs

