



z/OS V2R1 Communications Server Performance Update

David Herr – dherr@us.ibm.com
IBM Raleigh, NC

Thursday, August 15th, 9:30am
Session: 13633



Trademarks

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

AIX*	DB2*	HiperSockets*	MQSeries*	PowerHA*	RMF	System z*	zEnterprise*	zVM*
BladeCenter*	DFSMS	HyperSwap	NetView*	PR/SM	Smarter Planet*	System z10*	z10	z/VSE*
CICS*	EASY Tier	IMS	OMEGAMON*	PureSystems	Storwize*	Tivoli*	z10 EC	
Cognos*	FICON*	InfiniBand*	Parallel Sysplex*	Rational*	System Storage*	WebSphere*	z/OS*	
DataPower*	GDPS*	Lotus*	POWER7*	RACF*	System x*	XIV*		

* 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.

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.

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.

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

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.

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

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

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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

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

* Other product and service names might be trademarks of IBM or other 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.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Agenda



- ❑ V2R1 Performance Enhancements
- ❑ Optimizing inbound communications using OSA-Express
- ❑ Optimizing outbound communications using OSA-Express
- ❑ OSA-Express4
- ❑ z/OS Communications Server Performance Summaries



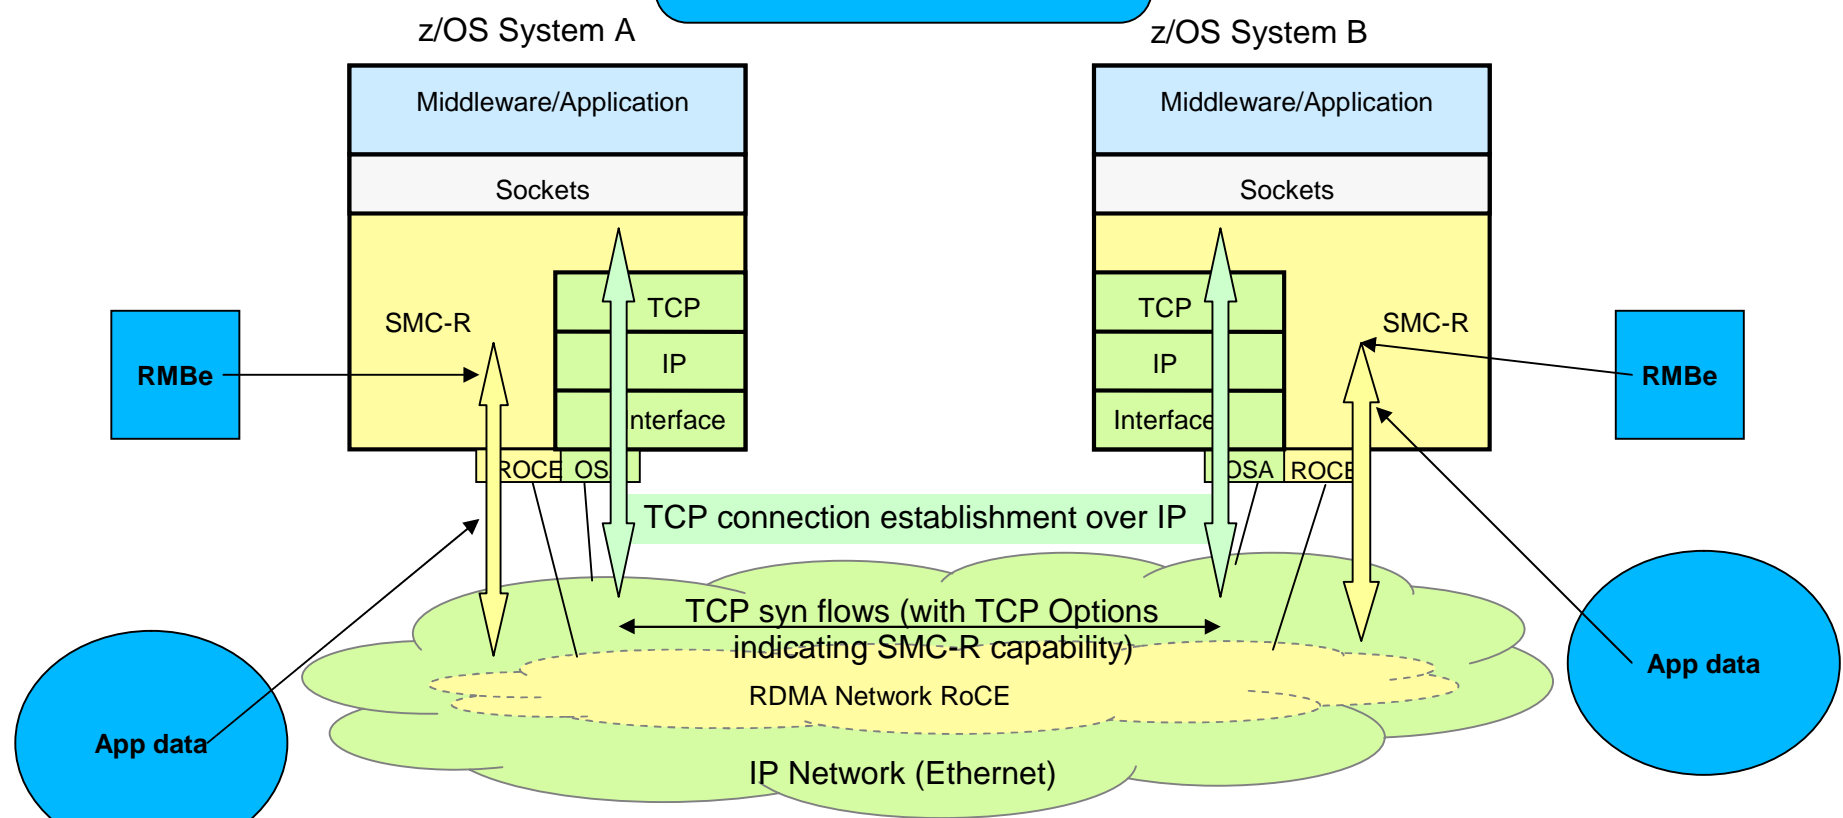
Disclaimer: All statements regarding IBM future direction or intent, including current product plans, are subject to change or withdrawal without notice and represent goals and objectives only. All information is provided for informational purposes only, on an “as is” basis, without warranty of any kind.

V2R1 Performance Enhancements

Shared Memory Communications – Remote (SMC-R)

SMC-R Background

Both TCP and SMC-R “connections” remain active



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

SMC-R - RDMA

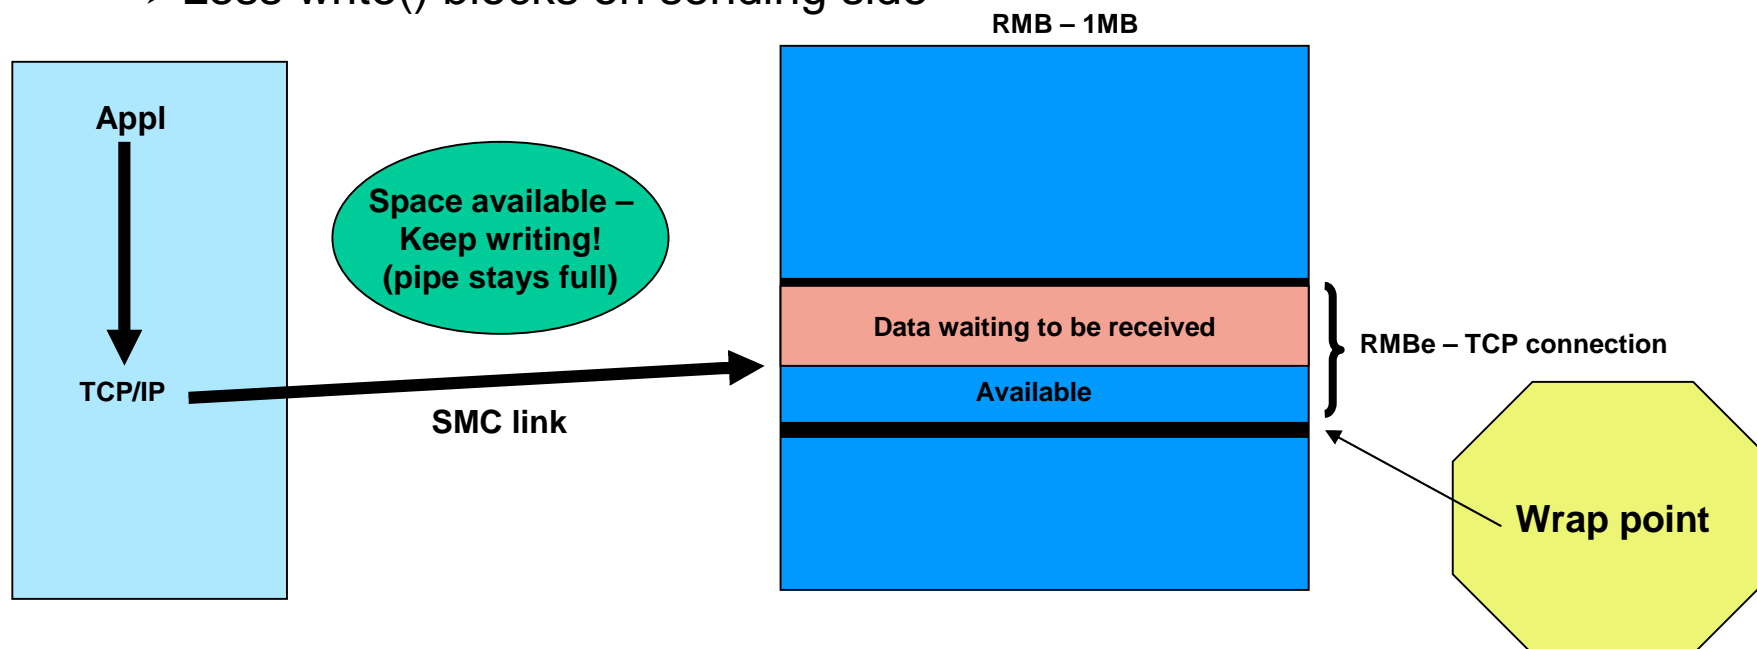
- Key attributes of RDMA
 - Enables a host to read or write directly from/to a remote host's memory *without* involving the remote host's CPU
 - By registering specific memory for RDMA partner use
 - **Interrupts still required for notification (i.e. CPU cycles are not completely eliminated)**
 - Reduced networking stack overhead by using streamlined, low level, RDMA interfaces
 - Key requirements:
 - A reliable “lossless” network fabric (LAN for layer 2 data center network distance)
 - An RDMA capable NIC (RNIC) and RDMA capable switched fabric (switches)

SMC-R - Solution

- Shared Memory Communications over RDMA (SMC-R) is a protocol that allows *TCP sockets* applications to transparently exploit RDMA (RoCE)
- SMC-R is a “hybrid” solution that:
 - Uses TCP connection (3-way handshake) to establish SMC-R connection
 - Each TCP end point exchanges TCP options that indicate whether it supports the SMC-R protocol
 - SMC-R “rendezvous” (RDMA attributes) information is then exchanged within the TCP data stream (similar to SSL handshake)
 - Socket application data is exchanged via RDMA (write operations)
 - TCP connection remains active (controls SMC-R connection)
 - This model preserves many critical existing operational and network management features of TCP/IP

SMC-R – Role of the RMBe (buffer size)


- The RMBe is a slot in the RMB buffer for a specific TCP connection
 - Based on TCPRCVBufsize – NOT equal to
 - Can be controlled by application using setsockopt() SO_RCVBUF
 - 5 sizes – 32K, 64K, 128K, 256K and 1024K (1MB)
 - Depending on the workload, a larger RMBe can improve performance
 - Streaming (bulk) workloads
 - Less wrapping of the RMBe = less RDMA writes
 - Less frequent “acknowledgement” interrupts to sending side
 - Less write() blocks on sending side



SMC-R – Micro benchmark performance results

- Response time/Throughput and CPU improvements
- Workload:
 - Using AWM (Application Workload Modeler) to model “socket to socket” performance using SMC-R
 - AWM very lightweight - contains no application/business logic
 - *Stresses and measures the networking infrastructure*
 - *Real workload benefits **will be smaller** than the improvements seen in AWM benchmarks!*
 - MTU: RoCE (1K and 2K) OSA (1500 and 8000)
 - Large Send enabled for some of the TCP/IP streaming runs
 - RR1(1/1): Single interactive session with 1 byte request and 1 byte reply
 - RR10: 10 concurrent connections with various message sizes
 - STR1(1/20M): Single Streaming session with 1 byte request (Client) and 20,000,000 bytes reply (Server)
 - Used large RMBs – 1MB

SMC-R – Micro benchmark performance results

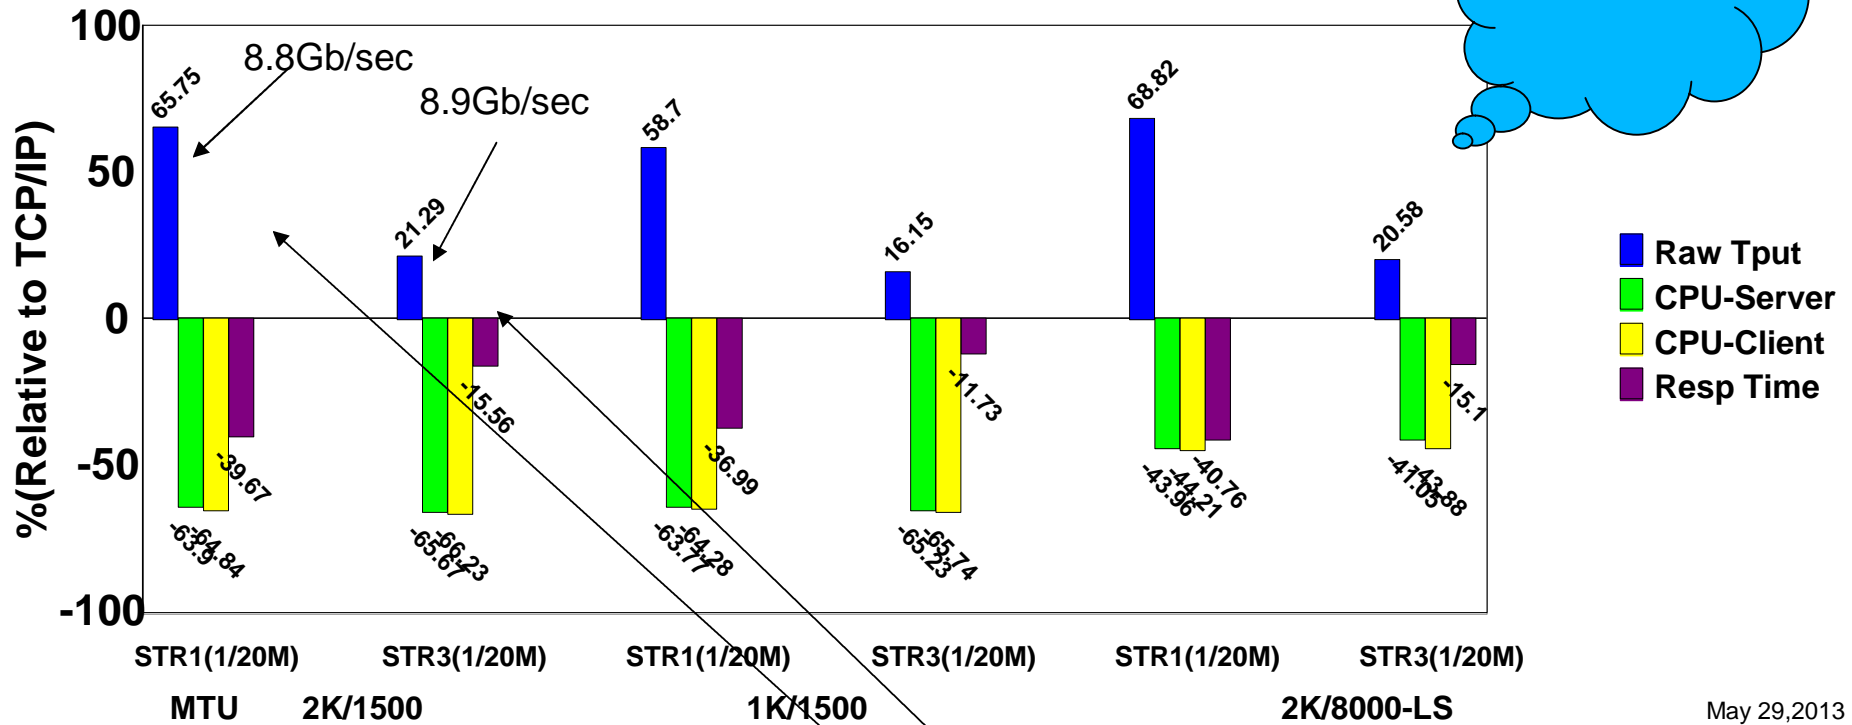



Significant Latency reduction across all data sizes (52-88%)
Reduced CPU cost as payload increases (up to 56% CPU savings)
Impressive throughput gains across all data sizes (Up to +717%)

Note: vs typical OSA customer configuration
 MTU (1500), Large Send disabled
 RoCE MTU: 1K

SMC-R – Micro benchmark performance results

z/OS V2R1 SMC-R vs TCP/IP
Streaming Data Performance Summary (AWM)



May 29, 2013
Client, Server: 2827-791 2CPs L
Interfaces: 10GbE RoCE Express and 10GbE OS

Notes:

- Significant throughput benefits and CPU reduction benefits
 - Up to 69% throughput improvement
 - Up to 66% reduction in CPU costs
- 2K RoCE MTU does yield throughput advantages
- LS – Large Send enabled (Segmentation offload)



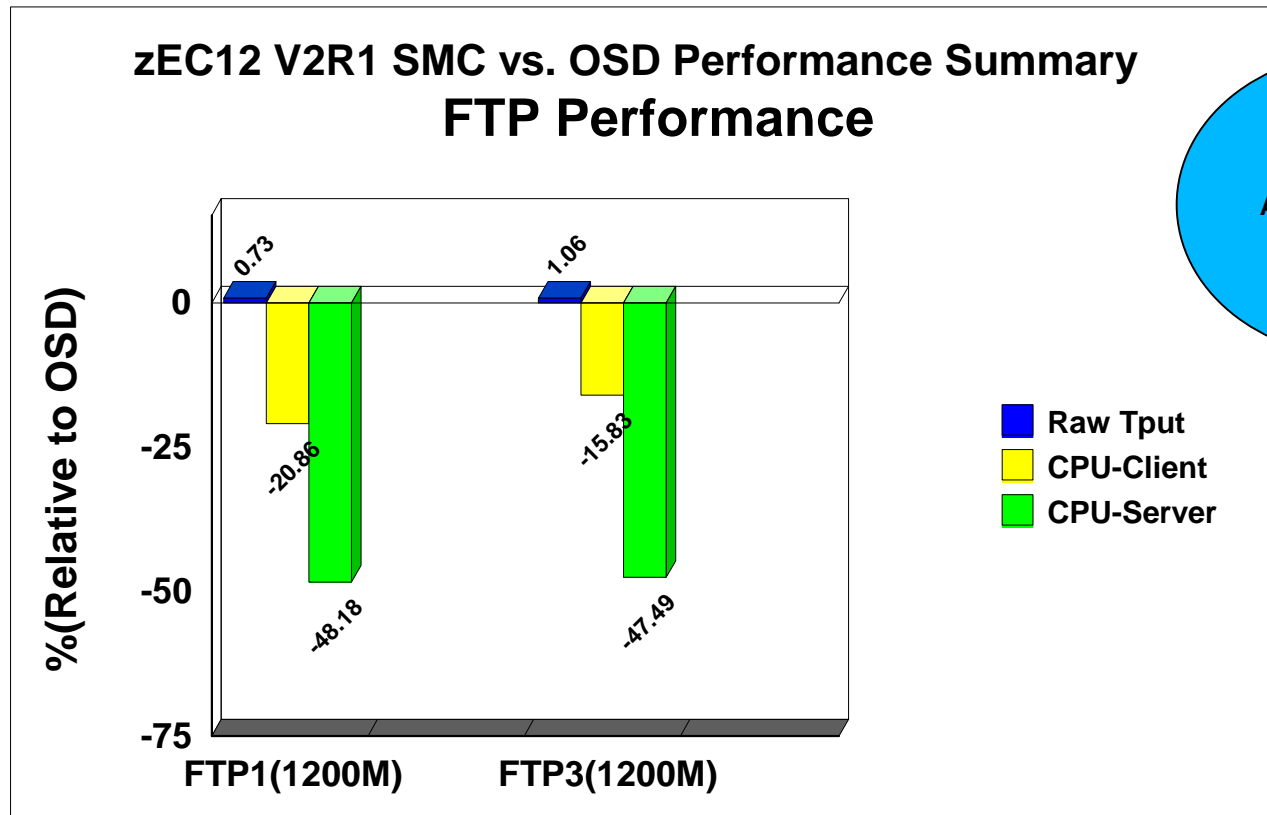
SMC-R – Micro benchmark performance results

- Summary –
 - Network latency for z/OS TCP/IP based OLTP (request/response) workloads reduced by up to 80%*
 - Networking related CPU consumption reduction for z/OS TCP/IP based OLTP (request/response) workloads increases as payload size increases
 - Networking related CPU consumption for z/OS TCP/IP based workloads with streaming data patterns reduced by up to 60% with a network throughput increase of up to 60%**
 - CPU consumption can be further optimized by using larger RMBE sizes
 - Less data consumed processing
 - Less data wrapping
 - Less data queuing

* Based on benchmarks of modeled z/OS TCP sockets based workloads with request/response traffic patterns using SMC-R vs. TCP/IP. The actual response times and CPU savings any user will experience will vary.

** Based on benchmarks of modeled z/OS TCP sockets based workloads with streaming data patterns using SMC-R vs. TCP/IP. The benefits any user will experience will vary

SMC-R – FTP performance summary



- FTP binary PUTs to z/OS FTP server, 1 and 3 sessions, transferring 1200 MB data
- OSD – OSA Express4 10Gb interface
- Reading from and writing to DASD datasets – Limits throughput

The performance measurements discussed in this document were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary significantly depending upon environments used.

SMC-R - WebSphere MQ for z/OS performance improvement

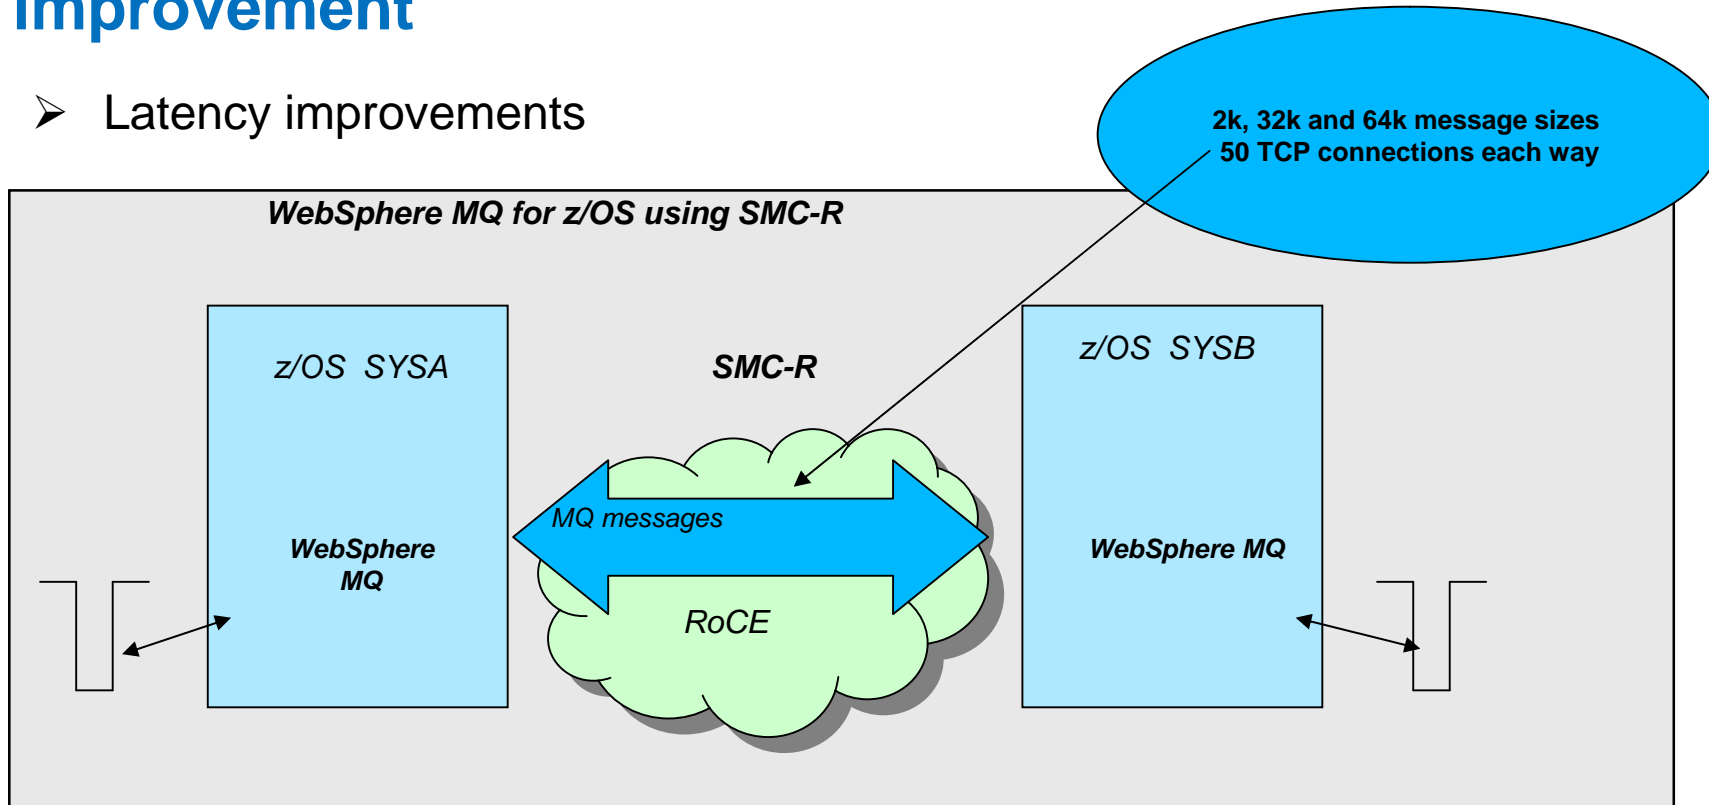
V2R1

- Latency improvements
- Workload
 - Measurements using WebSphere MQ V7.1.0
 - MQ between 2 LPARs on zEC12 machine (10 processors each)
 - On each LPAR, a queue manager was started and configured with 50 outbound sender channels and 50 inbound receiver channels, with default options for the channel definitions (100 TCP connections)
 - Each configuration was run with message sizes of 2KB, 32KB and 64KB where all messages were non-persistent
 - Results were consistent across all three message sizes

SMC-R - WebSphere MQ for z/OS performance improvement

V2R1

- Latency improvements



WebSphere MQ for z/OS realizes **up to a 3x increase** in messages per second it can deliver across z/OS systems when using SMC-R vs standard TCP/IP *

Based on internal IBM benchmarks using a modeled WebSphere MQ for z/OS workload driving non-persistent messages across z/OS systems in a request/response pattern. The benchmarks included various data sizes and number of channel pairs. The actual throughput and CPU savings users will experience may vary based on the user workload and configuration.

SMC-R – CICS performance improvement

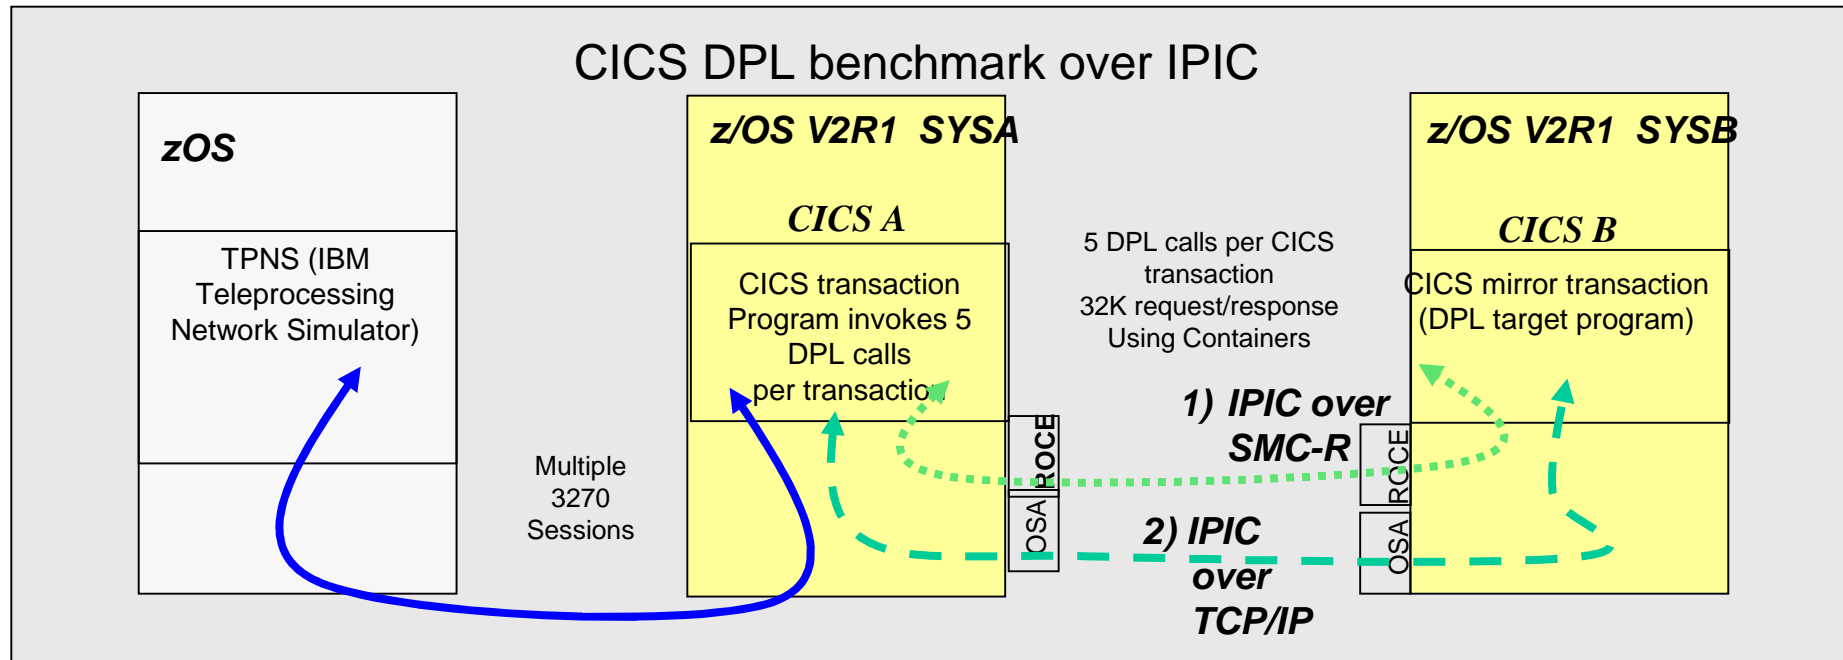
- Response time and CPU utilization improvements
- Workload - Each transaction
 - Makes 5 DPL (Distributed Program Link) requests over an IPIC connection
 - Sends 32K container on each request
 - Server program Receives the data and Send back 32K
 - Receives back a 32K container for each request

IPIC - IP Interconnectivity

- Introduced in CICS TS 3.2/TG 7.1
- TCP/IP based communications
- Alternative to LU6.2/SNA for Distributed program calls

Note: Results based on internal IBM benchmarks using a modeled CICS workload driving a CICS transaction that performs 5 DPL calls to a CICS region on a remote z/OS system, using 32K input/output containers. Response times and CPU savings measured on z/OS system initiating the DPL calls. The actual response times and CPU savings any user will experience will vary.

SMC-R – CICS performance improvement

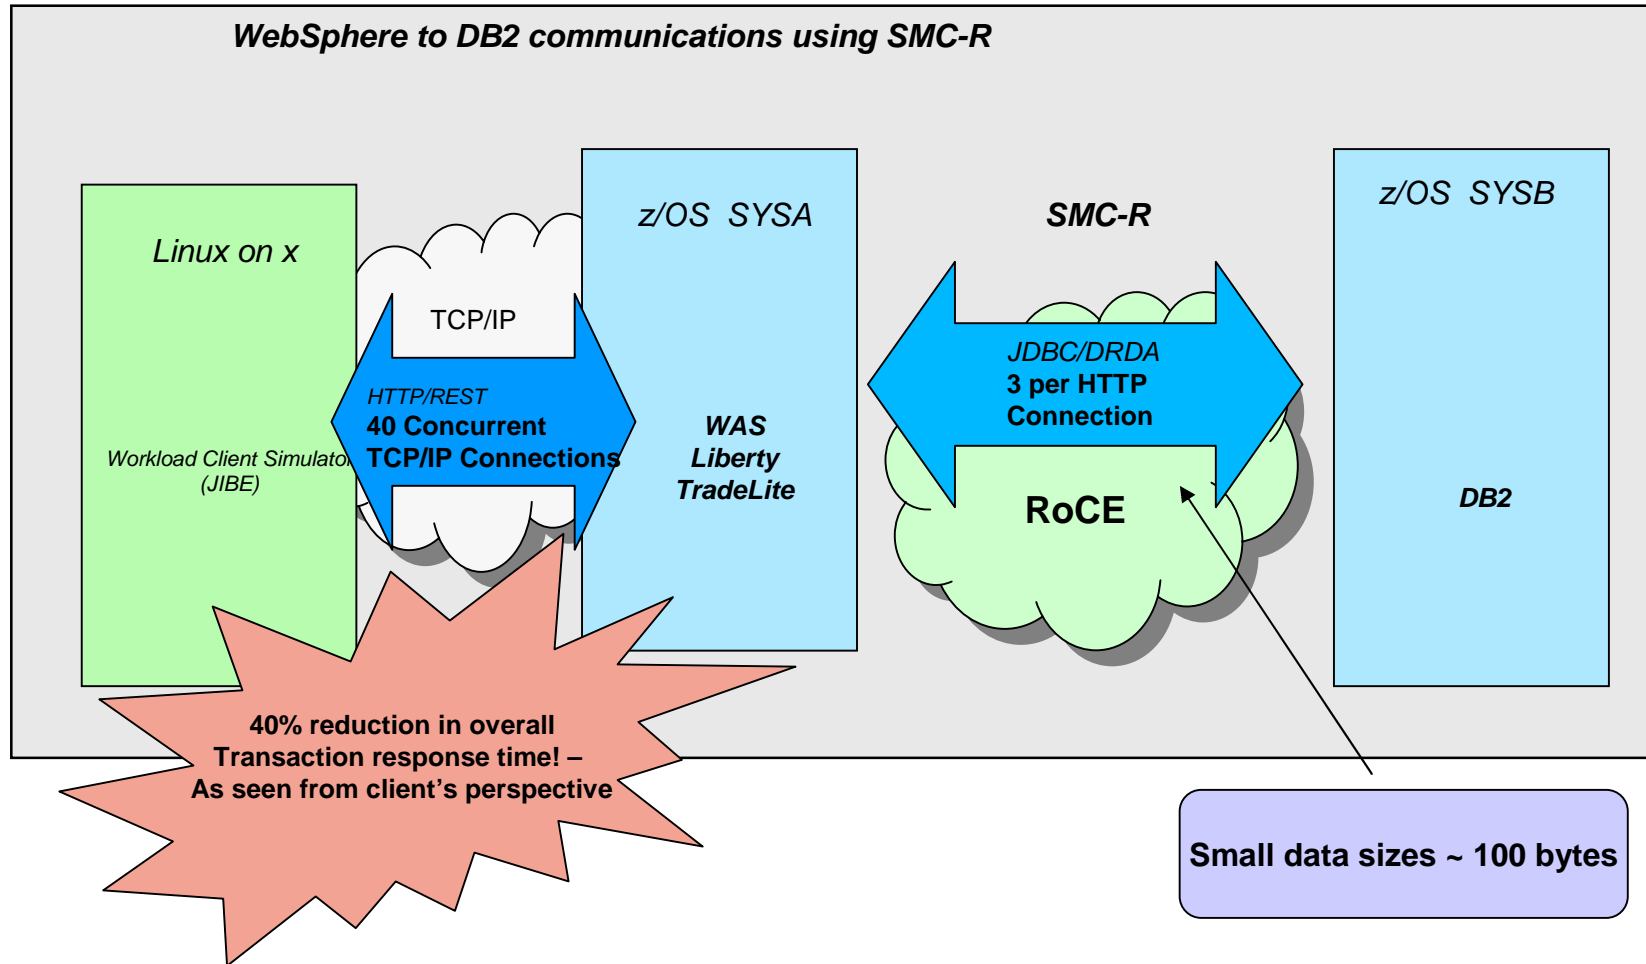


- Benchmarks run on z/OS V2R1 with latest zEC12 and new 10GbE RoCE Express feature
 - Compared use of SMC-R (10GbE RoCE Express) vs standard TCP/IP (10GbE OSA Express4S) with CICS IPIC communications for DPL (Distributed Program Link) processing
 - **Up to 48% improvement in CICS transaction response time** as measured on CICS system issuing the DPL calls (CICS A)
 - **Up to 10% decrease in overall z/OS CPU consumption** on CICS system issuing the DPL calls (SYSA)

SMC-R – Websphere to DB2 communications performance improvement

V2R1

- Response time improvements

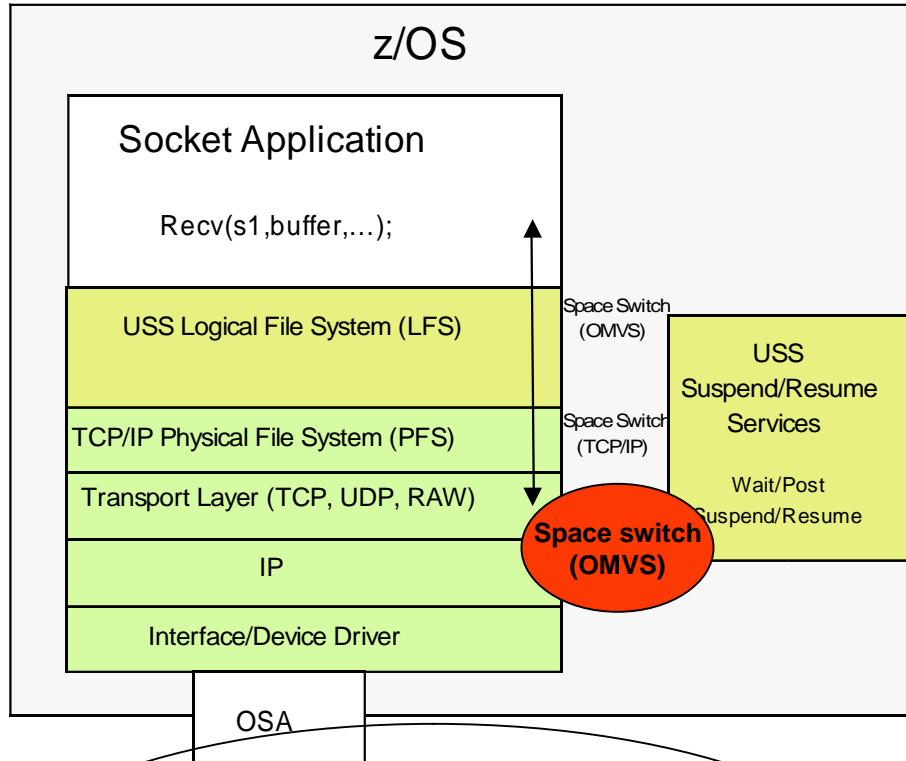


Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

TCP/IP Enhanced Fast Path Sockets

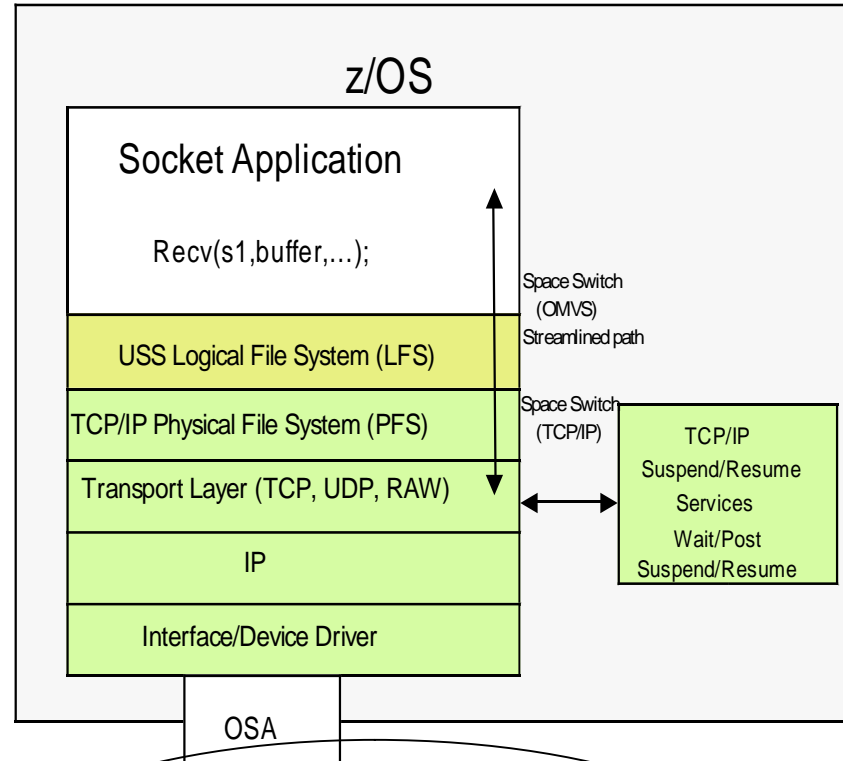
V2R1

TCP/IP sockets (normal path)



- Full function support for sockets, including support for Unix signals, POSIX compliance
- When TCP/IP needs to suspend a thread waiting for network flows, USS suspend/resume services are invoked

TCP/IP fast path sockets (Pre-V2R1)



- Streamlined path through USS LFS for selected socket APIs
- TCP/IP performs the wait/post or suspend/resume inline using its own services
- Significant reduction in path length

TCP/IP Enhanced Fast Path Sockets

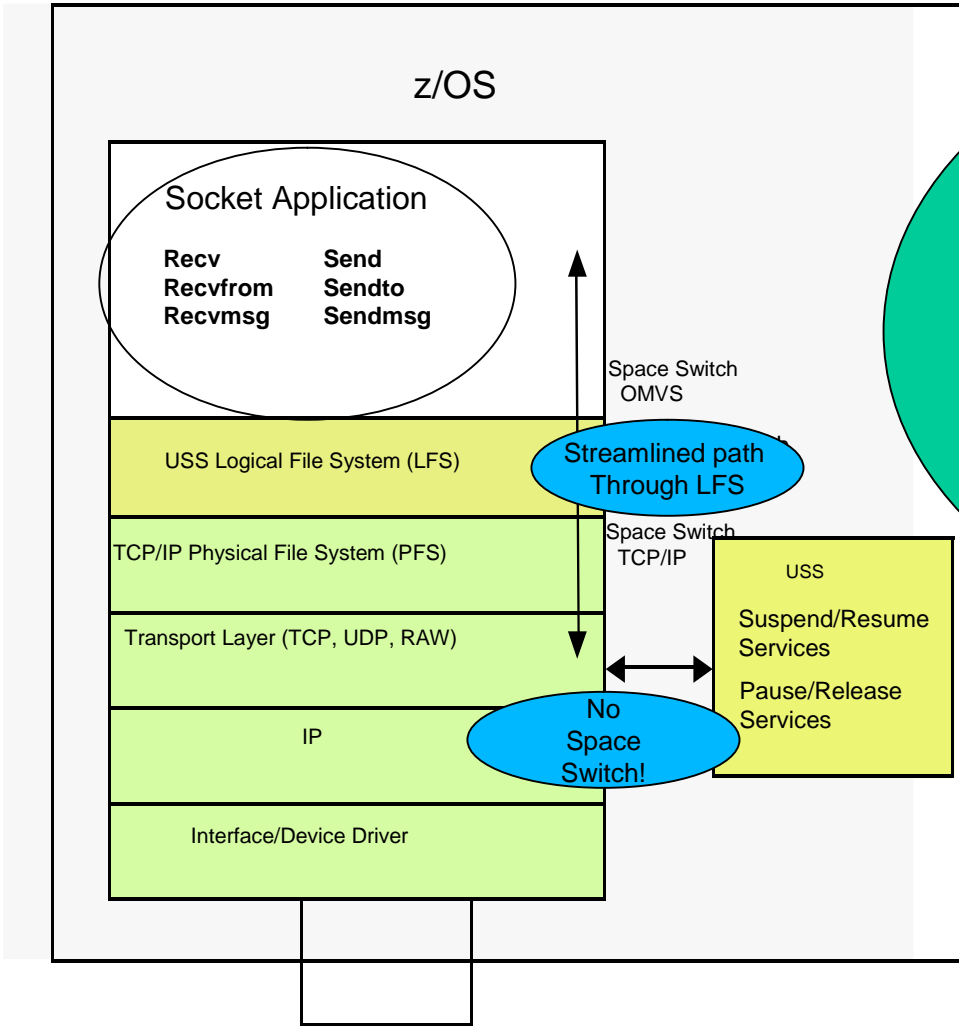
V2R1

Pre-V2R1 fast path provided CPU savings but not widely adopted:

- **No support for Unix signals (other than SIGTERM)**
 - **Only useful to applications that have no requirement for signal support**
- **No DBX support (debugger)**
- **Must be explicitly enabled!**
 - **BPXK_INET_FASTPATH environment variable**
 - **Ioctl#FastPath IOCTL**
- **Only supported for UNIX System Services socket API or the z/OS XL C/C++ Run-time Library functions**

TCP/IP Enhanced Fast Path Sockets

V2R1



Fast path sockets performance without all the conditions!

- Enabled by default
- Full POSIX compliance, signals support and DBX support
- Valid for **ALL** socket APIs (with the exception of the Pascal API)

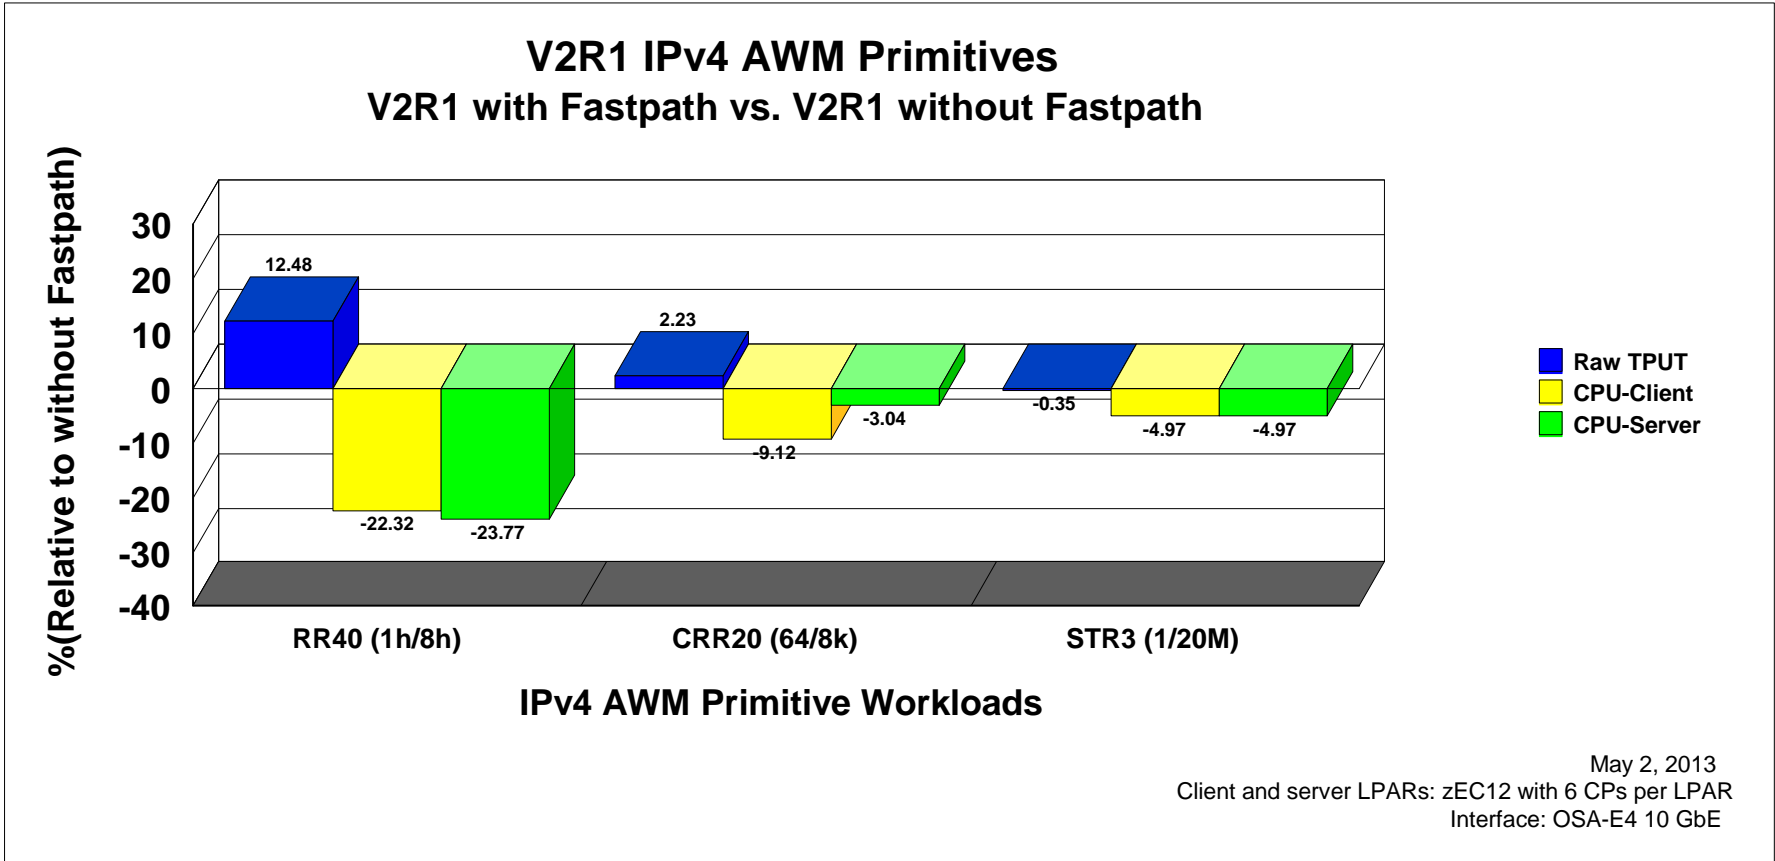
TCP/IP Enhanced Fast Path Sockets

V2R1

- No new externals
- Still supports “activating Fast path explicitly” to avoid migration issues
 - Provides performance benefits of enhanced Fast Path sockets
 - Keeps the following restrictions:
 - Does not support POSIX signals (blocked by z/OS UNIX)
 - Cannot use dbx debugger

TCP/IP Enhanced Fast Path Sockets

V2R1



Note: The performance measurements discussed in this presentation are z/OS V2R1 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

Optimizing inbound communications using OSA-Express

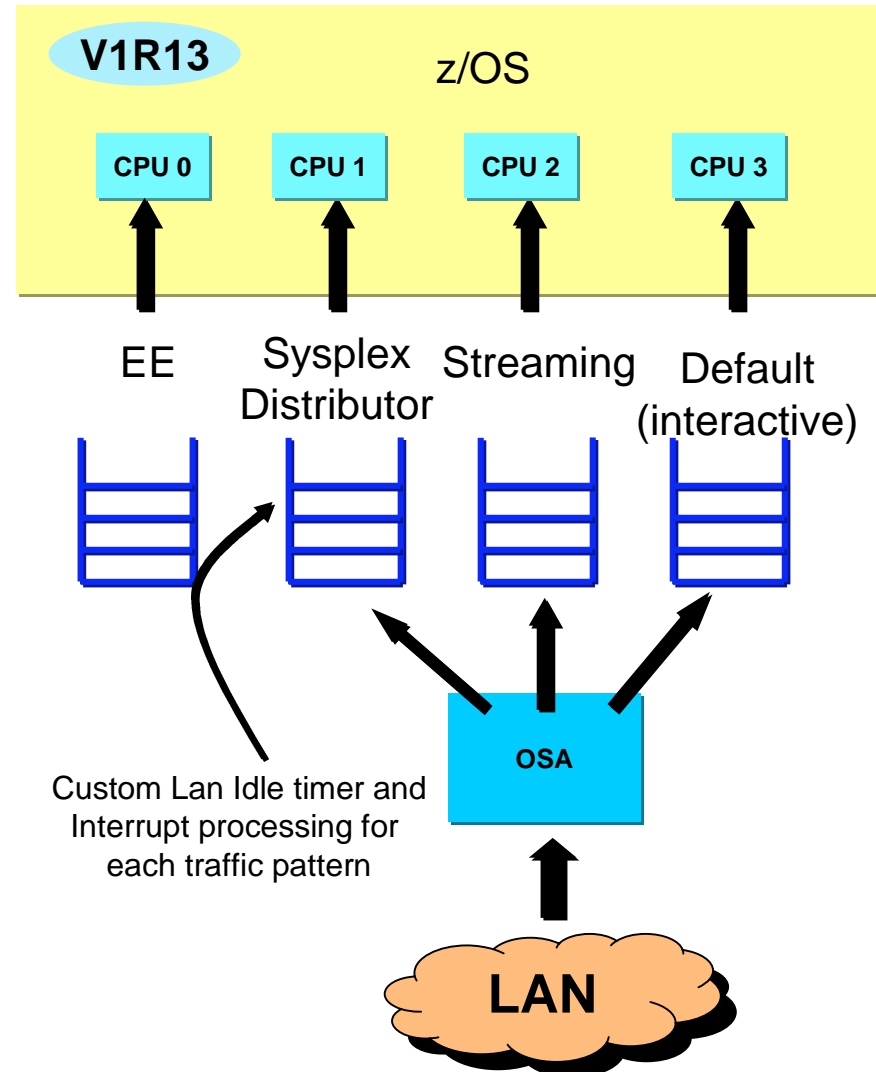
Inbound Workload Queuing

V1R12

With OSA-Express3/4S IWQ and z/OS V1R12, OSA now directs streaming traffic onto its own input queue – transparently separating the streaming traffic away from the more latency-sensitive interactive flows...

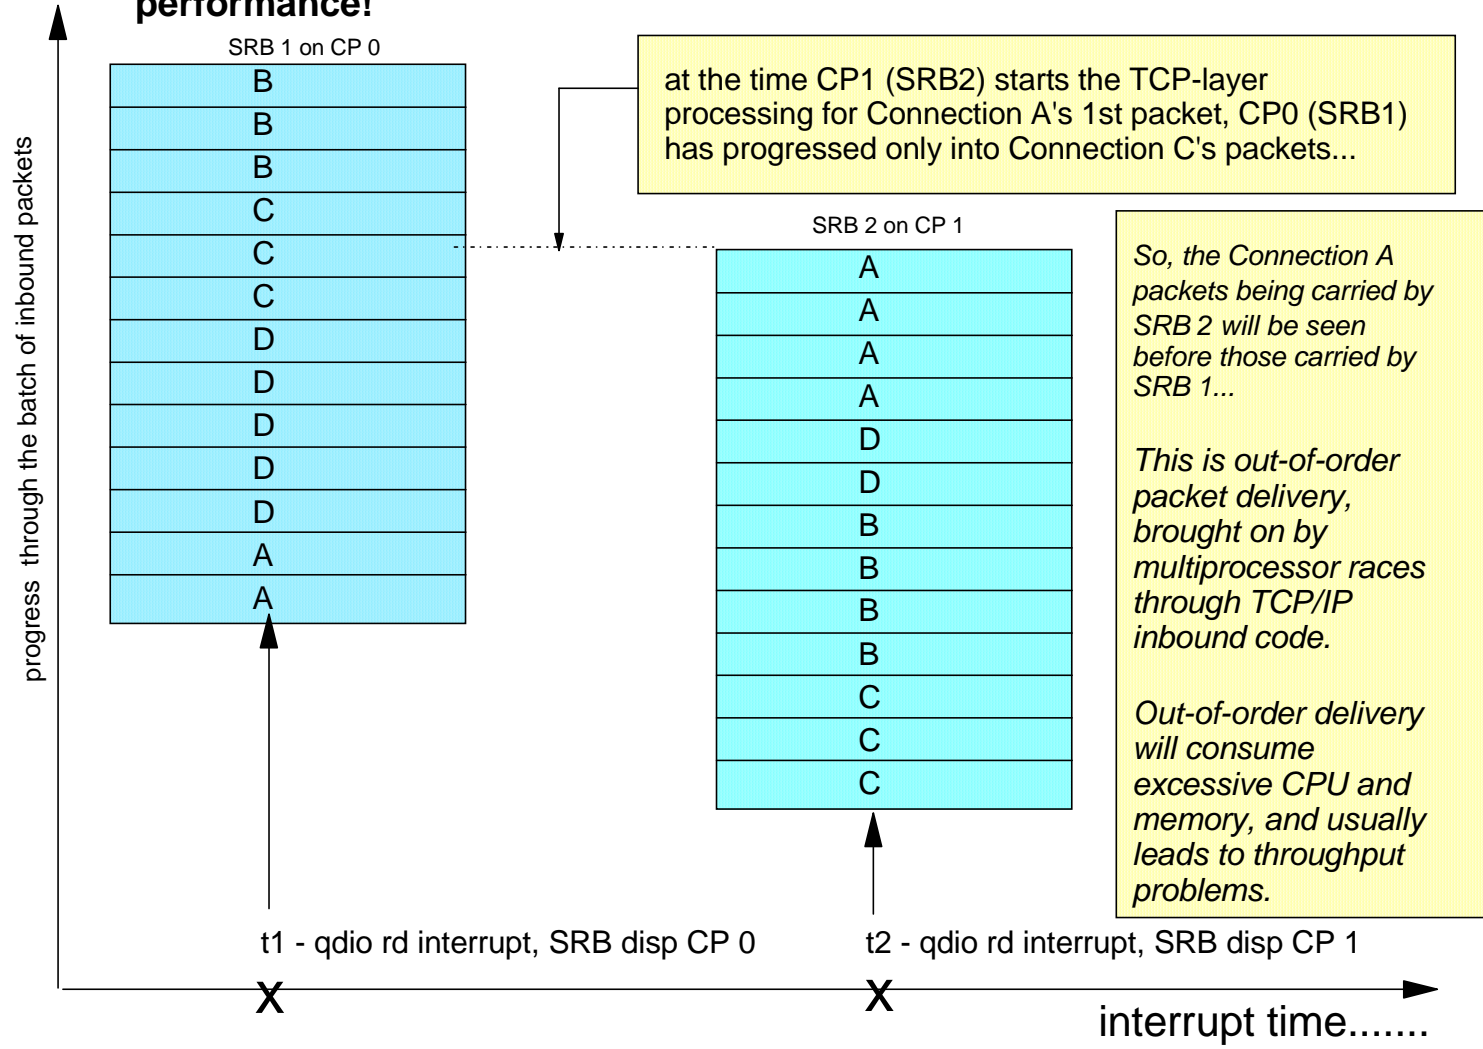
And each input queue has its own LAN-Idle timer, so the Dynamic LAN Idle function can now tune the streaming (bulk) queue to conserve CPU (high LAN-idle timer setting), while generally allowing the primary queue to operate with very low latency (minimizing its LAN-idle timer setting). So interactive traffic (on the primary input queue) may see significantly improved response time.

The separation of streaming traffic away from interactive also enables new streaming traffic efficiencies in Communications Server. This results in improved in-order delivery (better throughput and CPU consumption).



Improved Streaming Traffic Efficiency With IWQ

Before we had IWQ, Multiprocessor races would degrade streaming performance!



IWQ does away with MP-race-induced ordering problems!

With streaming traffic sorted onto its own queue, it is now convenient to service streaming traffic from a single CP (i.e., using a single SRB).

So with IWQ, we no longer have inbound SRB races for streaming data.

QDIO Inbound Workload Queuing

- Display OSAINFO command (V1R12) shows you what's registered in OSA

```
D TCPIP,,OSAINFO,INTFN=V6O3ETHG0
.
Ancillary Input Queue Routing Variables:
  Queue Type: BULKDATA Queue ID: 2 Protocol: TCP
  Src: 2000:197:11:201:0:1:0:1..221
  Dst: 100::101..257
  Src: 2000:197:11:201:0:2:0:1..290
  Dst: 200::202..514
  Total number of IPv6 connections:      2
  Queue Type: SYSDIST Queue ID: 3 Protocol: TCP
  Addr: 2000:197:11:201:0:1:0:1
  Addr: 2000:197:11:201:0:2:0:1
  Total number of IPv6 addresses:      2
36 of 36 Lines Displayed
End of report
```

5-Tuples

DVIPAs

- BULKDATA queue registers 5-tuples with OSA (streaming connections)
- SYSDIST queue registers Distributable DVIPAs with OSA

QDIO Inbound Workload Queuing: Netstat DEvlinks/-d

- Display TCPIP,,Netstat,DEvlinks to see whether QDIO inbound workload queueing is enabled for a QDIO interface

```
D TCPIP,,NETSTAT,DEVLINKS,INTFNAME=QDIO4101L
EZD0101I NETSTAT CS V1R12 TCPCS1
INTFNAME: QDIO4101L          INTFTYPE: IPAQENET    INTFSTATUS: READY
      PORTNAME: QDIO4101  DATAPATH: 0E2A      DATAPATHSTATUS: READY
      CHPIDTYPE: OSD
      SPEED: 0000001000
...
      READSTORAGE: GLOBAL (4096K)
      INBPERF: DYNAMIC
      WORKLOADQUEUEING: YES
      CHECKSUMOFFLOAD: YES
      SECCLASS: 255                      MONSYSPLEX: NO
      ISOLATE: NO                          OPTLATENCYMODE: NO
...
1 OF 1 RECORDS DISPLAYED
END OF THE REPORT
```

QDIO Inbound Workload Queuing: Display TRLE

- Display NET,TRL,TRLE=trlename to see whether QDIO inbound workload queuing is in use for a QDIO interface

```
D NET,TRL,TRLE=QDIO101
IST097I DISPLAY ACCEPTED
...
IST2263I PORTNAME = QDIO4101    PORTNUM =    0    OSA CODE LEVEL = ABCD
...
IST1221I DATA  DEV = 0E2A STATUS = ACTIVE        STATE = N/A
IST1724I I/O TRACE = OFF  TRACE LENGTH = *NA*
IST1717I ULPID = TCPCS1
IST2310I ACCELERATED ROUTING DISABLED
IST2331I QUEUE    QUEUE    READ
IST2332I ID       TYPE     STORAGE
IST2205I -----  -----  -----
IST2333I RD/1     PRIMARY   4.0M(64 SBALS)
IST2333I RD/2     BULKDATA 4.0M(64 SBALS)
IST2333I RD/3     SYSDIST  4.0M(64 SBALS)
...
IST924I -----
IST314I END
```

QDIO Inbound Workload Queuing: Netstat ALL/-A

- Display TCPIP,,Netstat,ALL to see whether QDIO inbound workload BULKDATA queueing is in use for a given connection

```
D TCPIP,,NETSTAT,ALL,CLIENT=USER1
EZD0101I NETSTAT CS V1R12 TCPCS1
CLIENT NAME: USER1                CLIENT ID: 00000046
LOCAL SOCKET:  ::FFFF:172.16.1.1..20
FOREIGN SOCKET:  ::FFFF:172.16.1.5..1030
  BYTESIN:                00000000000023316386
  BYTESOUT:               00000000000000000000
  SEGMENTSIN:            00000000000000016246
  SEGMENTSOUT:          00000000000000000922
  LAST TOUCHED:         21:38:53          STATE:          ESTABLISH
...
Ancillary Input Queue: Yes
  BulkDataIntfName: QDIO4101L
...
APPLICATION DATA:   EZAFTPOS D USER1      C      PSSS
-----
1 OF 1 RECORDS DISPLAYED
END OF THE REPORT
```

QDIO Inbound Workload Queuing: Netstat STATS/-S

- Display TCPIP,,NETSTAT,STATS to see the total number of TCP segments received on BULKDATA queues

```
D TCPIP,,NETSTAT,STATS,PROTOCOL=TCP
EZD0101I NETSTAT CS V1R12 TCPCS1
TCP STATISTICS
  CURRENT ESTABLISHED CONNECTIONS      = 6
  ACTIVE CONNECTIONS OPENED            = 1
  PASSIVE CONNECTIONS OPENED           = 5
  CONNECTIONS CLOSED                   = 5
  ESTABLISHED CONNECTIONS DROPPED      = 0
  CONNECTION ATTEMPTS DROPPED           = 0
  CONNECTION ATTEMPTS DISCARDED         = 0
  TIMEWAIT CONNECTIONS REUSED           = 0
  SEGMENTS RECEIVED                     = 38611
  ...
  SEGMENTS RECEIVED ON OSA BULK QUEUES= 2169
  SEGMENTS SENT                         = 2254
  ...
END OF THE REPORT
```



Quick INBPERF Review Before We Push On....

- The original static INBPERF settings (MINCPU, MINLATENCY, BALANCED) provide sub-optimal performance for workloads that tend to shift between request/response and streaming modes.
- We therefore **recommend customers specify INBPERF DYNAMIC**, since it self-tunes, to provide excellent performance even when inbound traffic patterns shift.
- Inbound Workload Queueing (IWQ) mode is an extension to the Dynamic LAN Idle function. IWQ improves upon the DYNAMIC setting, in part because it provides finer interrupt-timing control for mixed (interactive + streaming) workloads.

Dynamic LAN Idle Timer: Performance Data

Dynamic LAN Idle improved RR1 TPS 50% and RR10 TPS by 33%. Response Time for these workloads is improved 33% and 47%, respectively.


RR1 and RR10 Dynamic LAN Idle



1h/8h indicates 100 bytes in and 800 bytes out

Note: The performance measurements discussed in this presentation are z/OS V1R13 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

Inbound Workload Queuing: Performance Data

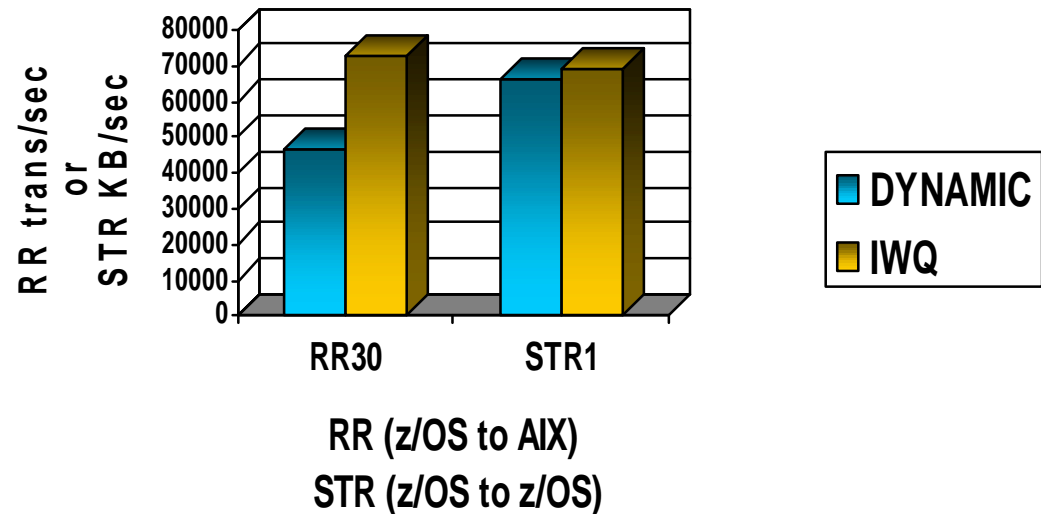


For z/OS outbound streaming to another platform, the degree of performance boost (due to IWQ) is relative to receiving platform's sensitivity to out-of-order packet delivery. For streaming INTO z/OS, IWQ will be especially beneficial for multi-CP configurations.


IWQ: Mixed Workload Results vs DYNAMIC:

- z/OS<->AIX R/R Throughput improved 55% (Response Time improved 36%)
- Streaming Throughput also improved in this test: +5%

Mixed Workload (IWQ vs Dynamic)




Inbound Workload Queuing: Performance Data



IWQ: Pure Streaming Results vs DYNAMIC:

- z/OS<->AIX Streaming Throughput improved 40%
- z/OS<->z/OS Streaming Throughput improved 24%

Pure Streaming (IWQ vs Dynamic)



For z/OS outbound streaming to another platform, the degree of performance boost (due to IWQ) is relative to receiving platform's sensitivity to out-of-order packet delivery. For streaming INTO z/OS, IWQ will be especially beneficial for multi-CP configurations.

IWQ Usage Considerations:

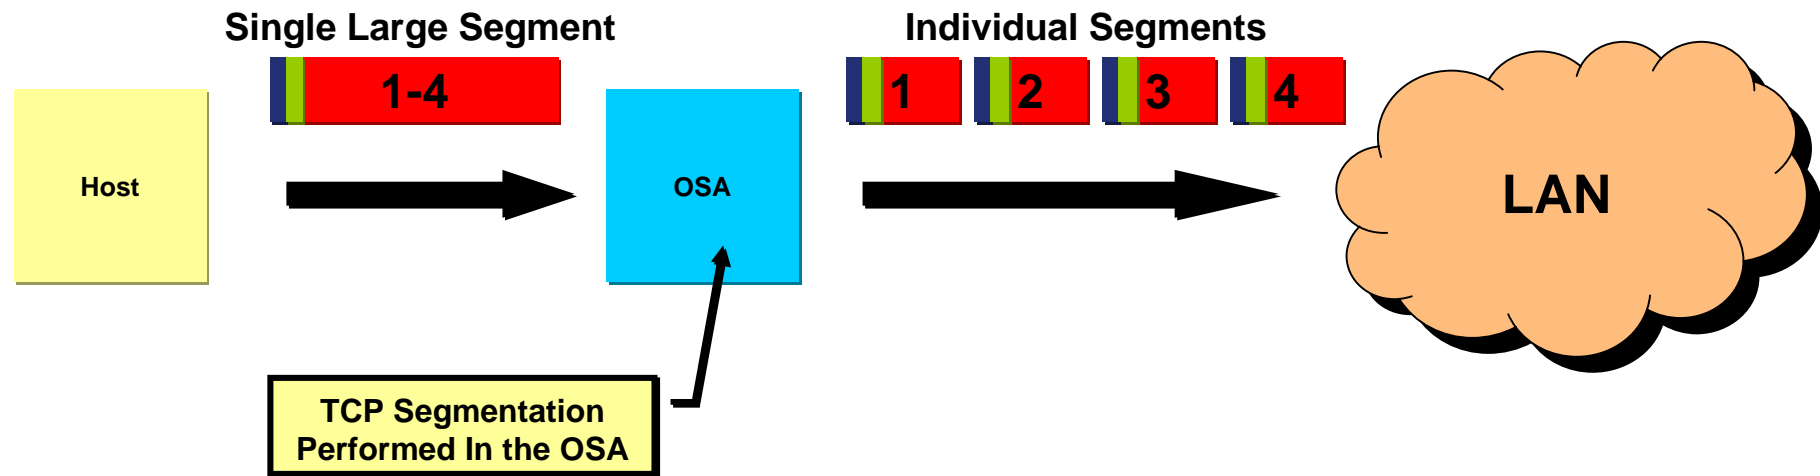
- Minor ECSA Usage increase: IWQ will grow ECSA usage by 72KBytes (per OSA interface) if Sysplex Distributor (SD) is in use; 36KBytes if SD is not in use
- IWQ requires OSA-Express3 in QDIO mode running on IBM System z10 or OSA-Express3/OSA-Express4 in QDIO mode running on zEnterprise 196/ zEC12.
- IWQ must be configured using the INTERFACE statement (not DEVICE/LINK)
- IWQ is not supported when z/OS is running as a z/VM guest with simulated devices (VSWITCH or guest LAN)
- Make sure to apply z/OS V1R12 PTF UK61028 (APAR PM20056) for added streaming throughput boost with IWQ

Optimizing outbound communications using OSA- Express

TCP Segmentation Offload


V1R13

- Segmentation consumes (high cost) host CPU cycles in the TCP stack
- Segmentation Offload (also referred to as “Large Send”)
 - Offload most IPv4 and/or IPv6 TCP segmentation processing to OSA
 - Decrease host CPU utilization
 - Increase data transfer efficiency
 - Checksum offload also added for IPv6

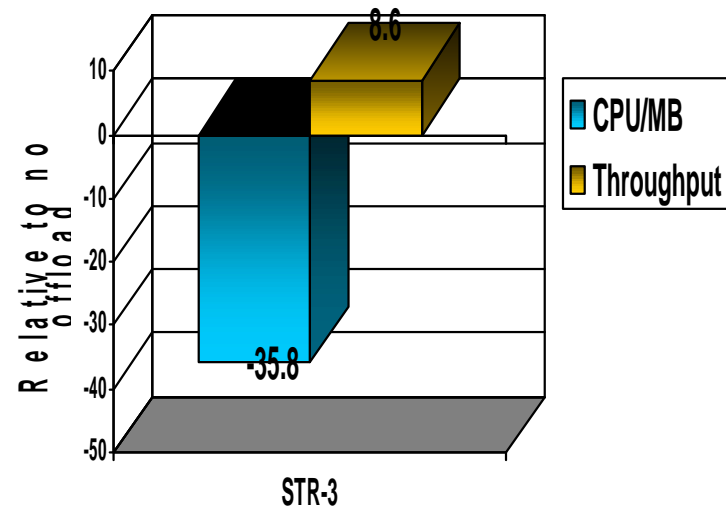


z/OS Segmentation Offload performance measurements

OSA-Express3 10Gb



OSA-Express4 10Gb



Send buffer size: 180K for streaming workloads

Segmentation offload may significantly reduce CPU cycles when sending bulk data from z/OS!

Note: The performance measurements discussed in this presentation are z/OS V1R13 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

TCP Segmentation Offload: Configuration

V1R13

- Enabled with IPCONFIG/IPCONFIG6 SEGMENTATIONOFFLOAD

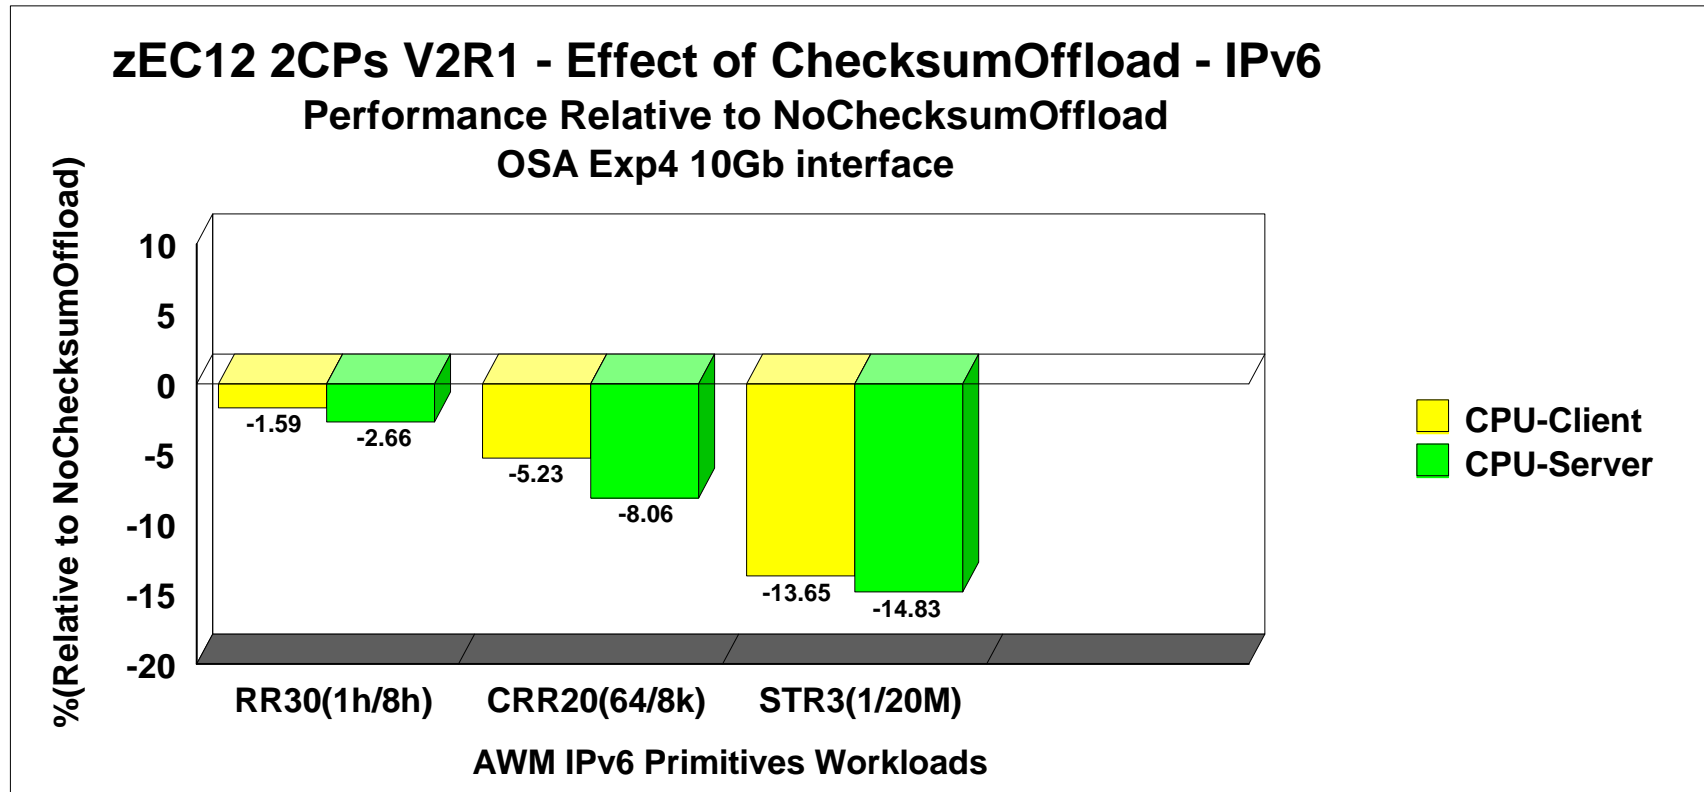
```
>>-IPCONFIG----->
.
.
>-----+-----+----->>
| .-NOSEGMENTATIONOFFLoad-. |
| +-----+-----+-----+ |
| '-SEGMENTATIONOFFLoad---' |
```

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

Reminder!
Checksum Offload
enabled by default

z/OS Checksum Offload performance measurements

V1R13




Note: The performance measurements discussed in this presentation are z/OS V2R1 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

OSA-Express4

OSA-Express4 Enhancements – 10GB improvements

- Improved on-card processor speed and memory bus provides better utilization of 10GB network

OSA 10GBe - Inbound Bulk traffic



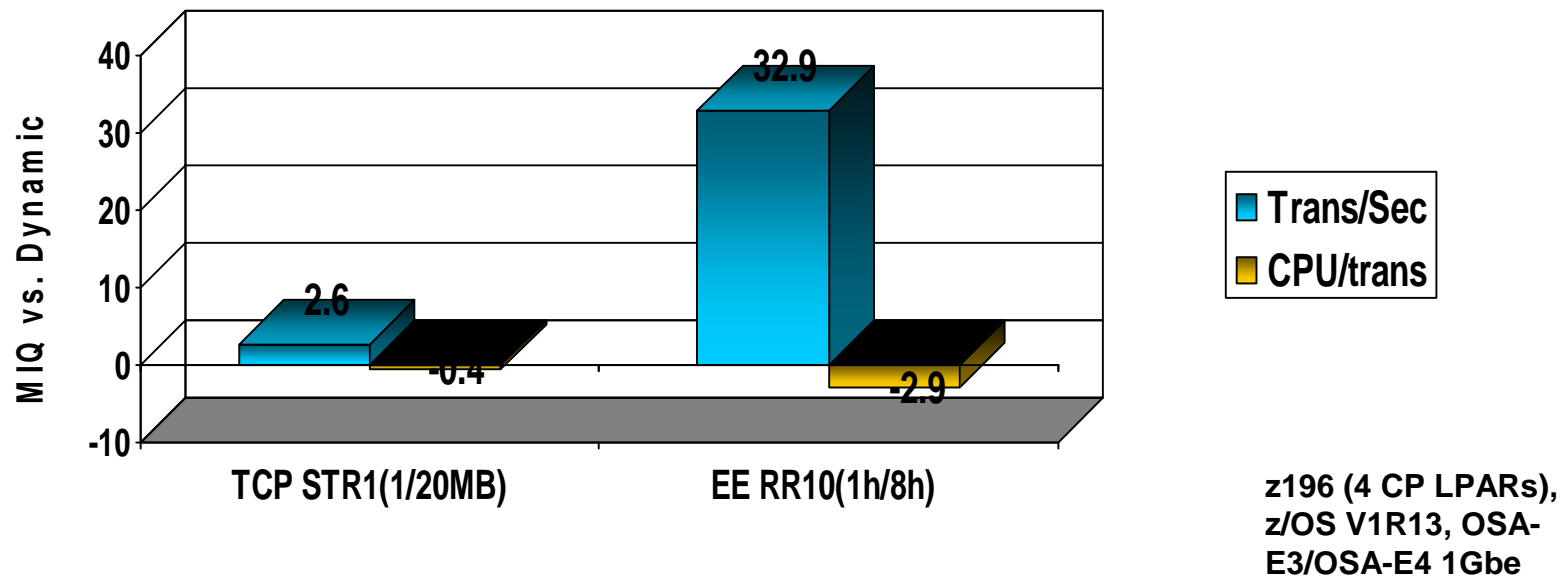
z196 (4 CP LPARs),
z/OS V1R13, OSA-
E3/OSA-E4 10Gbe

Note: The performance measurements discussed in this presentation are z/OS V1R13 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

OSA-Express4 Enhancements – EE Inbound Queue

- Enterprise Extender queue provides internal optimizations
 - EE traffic processed quicker
 - Avoids memory copy of data

OSA 1GBe - mixed TCP and EE workloads



Note: The performance measurements discussed in this presentation are z/OS V1R13 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments may vary.

OSA-Express4 Enhancements – Other improvements

- Checksum Offload support for IPv6 traffic
- Segmentation Offload support for IPv6 traffic

z/OS Communications Server Performance Summaries

z/OS Communications Server Performance Summaries

- Performance of each z/OS Communications Server release is studied by an internal performance team
- Summaries are created and published on line
 - <http://www-01.ibm.com/support/docview.wss?rs=852&uid=swg27005524>
- Ex: The z/OS V1R13 Communications Server Performance Summary includes:
 - The z/OS V1R13 Communications Server performance summary includes:
 - Performance of z/OS V1R13 Communications Server line items
 - Release to release performance comparisons (z/OS V1R13 Communications Server versus z/OS V1R12 Communications Server)
 - Capacity planning performance for:
 - *TN3270 (Clear Text, AT-TLS, and IPsec)*
 - *FTP (Clear Text, AT-TLS, and IPsec)*
 - *CICS Sockets performance*
 - CSM usage
 - VTAM buffer usage

z/OS Communications Server Performance Website

www-01.ibm.com/support/docview.wss?uid=swg27005524

The screenshot shows a web browser window displaying the IBM z/OS Communications Server performance index page. The browser's address bar shows the URL: <http://www-01.ibm.com/support/docview.wss?uid=swg27005524>. The page features the IBM logo and navigation menus for Home, Solutions, Services, Products, Support & downloads, and My IBM. The main content area is titled "z/OS Communications Server performance index" and includes a "White paper" section with an abstract and a list of content links. A right-hand sidebar contains a "Rate this page" section with a star rating, an "Add comments" button, and a "Document information" section listing software versions, operating systems, and reference numbers. The page is viewed in Mozilla Firefox, and the status bar at the bottom indicates "Done".

IBM z/OS Communications Server performance index - United States - Mozilla Firefox: IBM Edition

File Edit View History Bookmarks Tools Help

http://www-01.ibm.com/support/docview.wss?uid=swg27005524

United States [change]

Search

Home Solutions Services Products Support & downloads My IBM Welcome [IBM Sign in] [Register]

← IBM Support Portal

z/OS Communications Server performance index

Tags

Add a tag | Search all tags

Add a tag +

My tags | All tags

View as cloud | list

White paper

Abstract

z/OS Communications Server performance summary reports

Content

- [z/OS V1R13 Communications Server Performance Summary](#)
- [z/OS V1R12 Communications Server Performance Summary](#)
- [z/OS V1R12 Communications Server Performance Study: OSA Express3 Inbound Workload Queuing](#)
- [z/OS V1R11 Communications Server Performance Summary](#)
- [z/OS V1R10 Communications Server Large Send Performance Summary](#)
- [z/OS V1R10 Communications Server Performance Summary](#)
- [System z10 vs. System z9 Communications Server Performance](#)

Rate this page:

☆☆☆☆ Average rating 1 users

Add comments

Document information

[z/OS Communications Server](#)

All

Software version: 1.6, 1.7, 1.8, 1.9, 1.10, 1.11, 1.12, 1.13

Operating system(s): z/OS

Reference #: 7005524

Modified date: 2011-12-02

Translate my page

Done

Please fill out your session evaluation

- z/OS CS Performance Improvements
- Session # 13633
- QR Code:







Find us on Facebook at
<http://www.facebook.com/IBMCommserver>



Follow us on Twitter at
http://www.twitter.com/IBM_Commserver



Read the z/OS Communications Server blog at
<http://tinyurl.com/zoscsblog>



Visit the z/OS CS YouTube channel at
<http://www.youtube.com/user/zOSCommServer>