

Getting the Most Out of OSA and HiperSockets with z/OS Communications Server

SHARE Orlando, Florida 2026 Winter Technical Conference
Session 183

Karthik Sundaresan – ksundar@ibm.com
Christopher Nyamful – cnyamfu@us.ibm.com

Agenda



- ❑ What is QDIO?
- ❑ OSA inbound “routing”
- ❑ Overview of selected OSA enhancements
- ❑ Best practices
- ❑ HiperSockets Converged Interface (HSCI)




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.



WHAT IS QDIO?

Queued Direct IO (QDIO) – How a zSystem accesses a LAN



OSA and QDIO work together


■ QDIO Layer 3 (z/OS)– for TCP/IP traffic only

- OSA IP Processing Assist (IPA) offloads many TCP/IP functions to the OSA adapter, of which some are:
 - Performs all ARP processing (IPv4 only)
 - Provides Multicast support
 - Builds MAC and LLC headers
 - Performs checksum processing
 - Performs outbound TCP segmentation
 - Inbound Workload Queuing (IWQ)
- To use for SNA/APPN/HPR traffic: done thru TN3270 and Enterprise Extender

■ OSA and QDIO are designed for high speed communication between zSystem and the Local Area Network


- Reduced TCP/IP path length (reduced z/OS CPU)
- QDIO IP Processing Assist
- LPAR-to-LPAR communication with port sharing
- Direct Memory Access (DMA) Protocol
 - Memory-to-memory communication
 - **DMA allows data to move directly from the OSA-Express microprocessor to the host memory**
- Dynamic customization

■ 25 Gigabit Ethernet, 10 Gigabit Ethernet, Gigabit Ethernet, 1000BASE-T Ethernet (1000, 100, and 10 Mbit)



A little more on QDIO and devices

- Control data between z/OS CS and the OSA port is exchanged over one read device and one write device (READ even address, WRITE odd address)
 - The READ and WRITE device pair is shared among all the IP interface definitions in an LPAR that use the same OSA port
- Data is exchanged over data devices
- One data device is used per:
 - A set of exactly one IPv4 and one IPv6 interface
 - IPv4: DEVICE+LINK, IPv6: INTERFACE
 - One IPv4 interface defined with INTERFACE
 - One IPv6 interface
- One data device is used by OSAENTA trace, if that trace is activated





OSA INBOUND “ROUTING”

OSA-Express VMACs



*No reasons not to use
VMACs. Use them!!!*


- **OSA MAC sharing problems do not exist if each stack has its own MAC**
 - A "virtual" MAC
 - To the network, each stack appears to have a dedicated OSA port (NIC)
- **MAC address selection**
 - Coded in the TCP/IP profile -OR-
 - Generated and assigned by the OSA adapter (Recommended!)
- **All IP addresses for a stack are advertised with the virtual MAC**
 - by OSA using ARP for IPv4
 - by the stack using Neighbor Discovery for IPv6
- **All external routers now forward frames to the virtual MAC**
 - OSA will "route" to an LPAR/Stack by virtual MAC instead of IP address
 - All stacks can be "routing" stacks instead of having to designate one primary router stack
- **Simplifies configuration greatly**
 - No PRIROUTER/SECROUTER!

What is a Virtual LAN (VLAN)?

Wikipedia:


- A **virtual LAN (VLAN)** is any broadcast domain that is partitioned and isolated in a computer network at the data link layer.
- Creates the appearance and functionality of network traffic that is physically on a single network but acts as if it is split between separate networks.

Network reconfiguration can be done through software instead of physically relocating devices.




z/OS and VLANs

Physical network diagram



Logical network diagram



- Each frame on the trunk mode connection carries a VLAN ID in the IEEE802.3 header that allows the network equipment to clearly identify which virtual LAN each frame belongs to.
- On an access mode connection, the switch will transport frames belonging to the configured VLAN ID for that access mode connection only.

VLAN is a LAN media virtualization technology that allows multiple independent IP networks (IP subnets) to share one physical media, such as a cable, an adapter, or a layer-2 switch. Connectivity between VLANs is under control of IP routers.

- Depending on switch configuration, the switch may interconnect the VLANs using a layer-3 IP router function.
- The subnets may belong to different routing domains or OSPF areas:
 - Test, production, demo
- The subnets may belong to different security zones:
 - Intranet, DMZ

OSA inbound routing

This is correct if all interfaces that share the OSA port use VLAN IDs or none of them do. If some do and others do not, this becomes utterly complex, rather ugly, and somewhat scary !!!!

VLAN ID

- If z/OS CS has set a global VLAN ID for an interface, OSA verifies the frame VLAN ID matches the global VLAN ID

Dest MAC

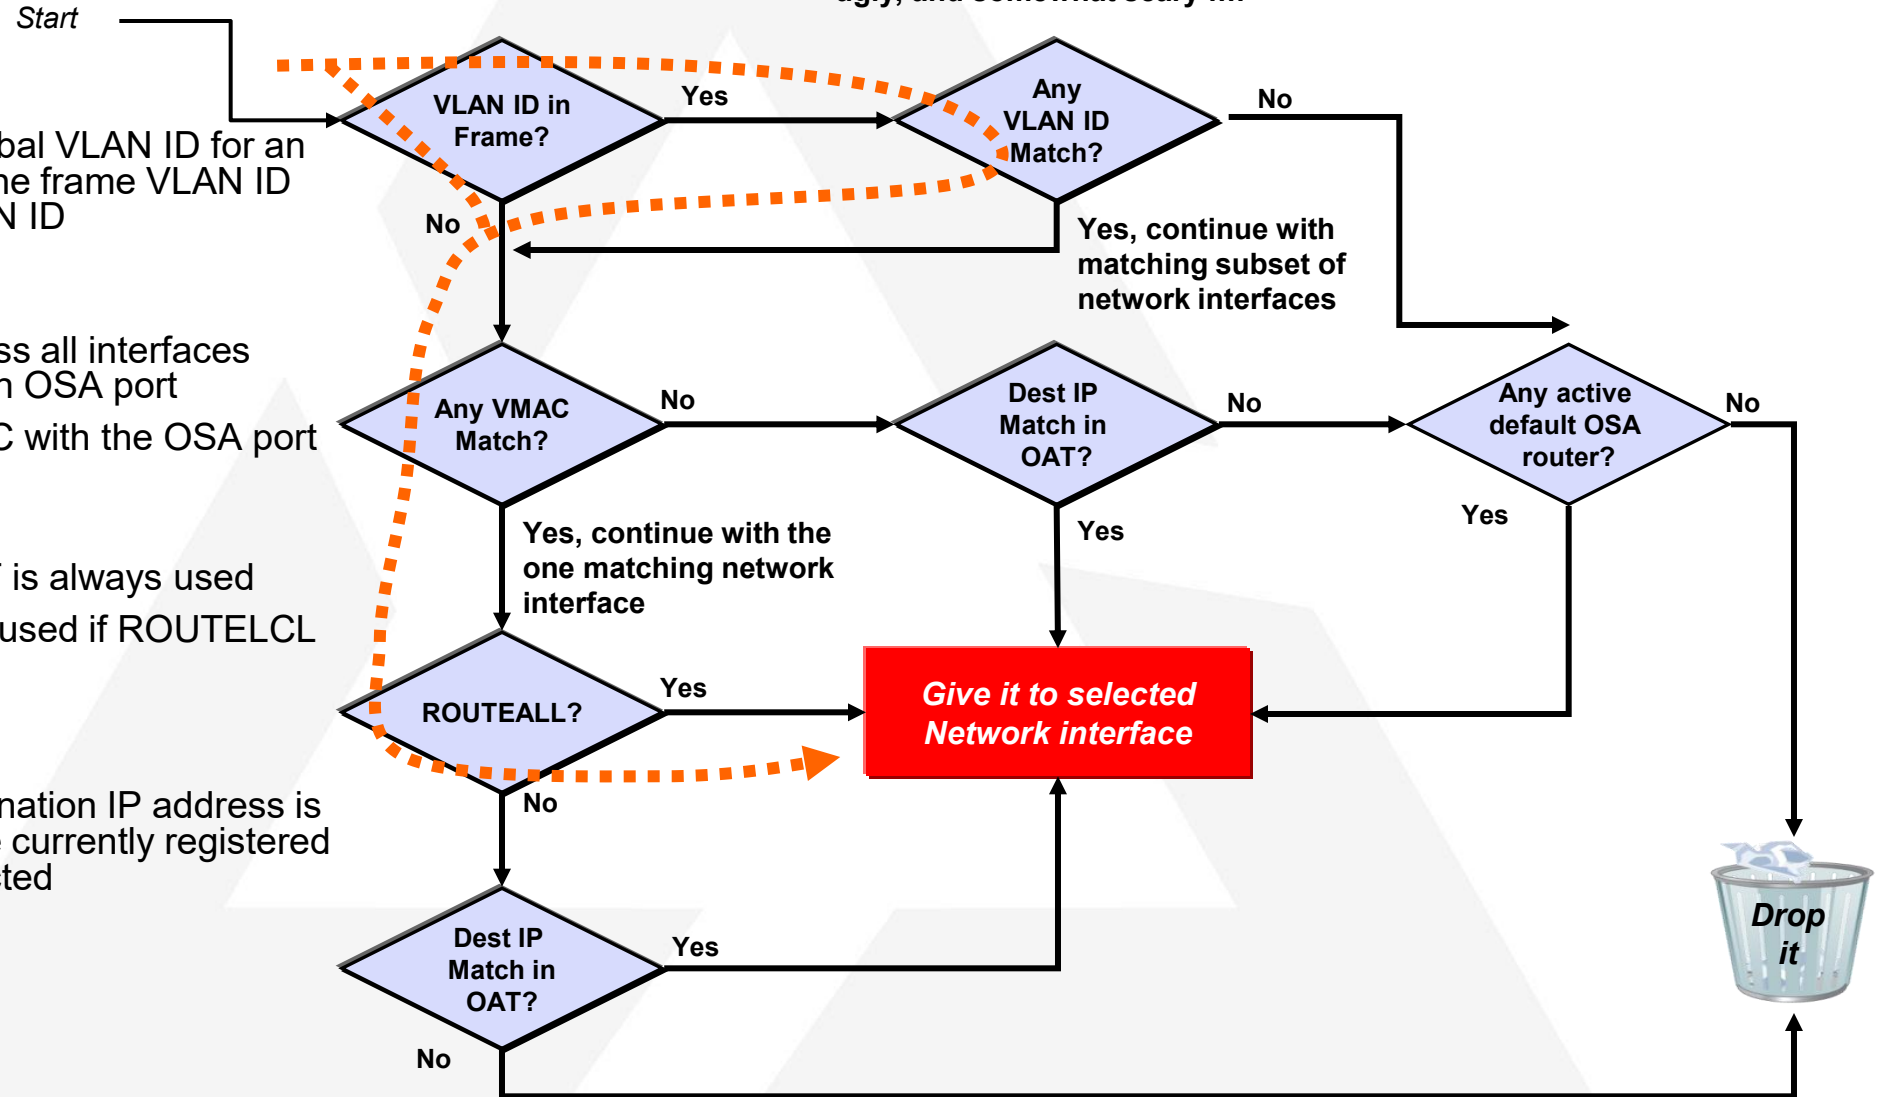
- VMACs are unique across all interfaces and VLANs that share an OSA port
- z/OS CS registers VMAC with the OSA port

Dest IP address

- Without VMAC, the OAT is always used
- With VMAC, the OAT is used if ROUTELCL is specified

Default OSA router

- If OAT is used and destination IP address is not in OAT, the interface currently registered as default router is selected




This elusive Maximum Transmission Unit – what is the size really?

 **PROFILE**

INTERFACE .. MTU 8992

 **OMROUTE**

OSPF_Interface .. MTU 1500



**IPv6 always uses
path MTU**

Interface MTU

- Configured on the INTERFACE statement or learned from the device
 - For IPv4, OSA reports 1492/8992 – even when you use Ethernet II frames
 - For IPv6, OSA reports 1500/9000 – because IPv6 always uses Ethernet II frames
- Reported as ActMtu in a DEVLINKS netstat report

Configured route MTU

- For static routes: the value coded in the BEGINROUTES section for the destination
- For dynamic routes: the MTU value from the OMROUTE interface over which the route was learned
 - If no MTU size defined in OMROUTE, a default of 576 will be used for IPv4

Actual route MTU

- Minimum of the interface MTU and the configured route MTU
- If path MTU is not enabled, this is the MTU that will be used for that route

Path MTU

- With path MTU enabled, TCP/IP starts out trying to use the actual route MTU and may then lower the MTU to what is discovered
 - With path MTU enabled, all frames are sent with the Do Not Fragment bit
 - Path MTU values are cached, but will timeout after a certain amount of time to accommodate topology changes that allows a higher MTU

MTU type	MTU value
Interface MTU	8992
Configured route MTU	1500
Actual route MTU	1500
Path MTU to this destination	896



OVERVIEW OF SELECTED OSA ENHANCEMENTS

SMC TCP/IP Configuration – OSA Interface

- IPAQENET INTERFACE statements

- SMC only valid for CHPIDTYPE OSD (also for OSH)
- SMC cannot be used with IPv4 OSD interfaces defined using DEVICE and LINK statements

**Must be non-zero subnet
For SMC**

OSD interface definition:

```
|__PORTNAME portname__IPADDR__ _ipv4_address/0_____ |__NONRouter__|_____>
|_ipv4_address_____ |__PRIRouter__| |__VLANID id_|
|_ipv4_address/num_mask_bits_| |__SECRouter__|
> |__INBPERF BALANCED_____ |_____>
|__INBPERF_____ |__DYNAMIC__| |__NOWORKLOADQ__| |__VMAC_____ |__ROUTEALL__|
|_____ |__WORKLOADQ__| |__macaddr__| |__ROUTECL__|
|__MINCPU_____ |
|__MINLATENCY_____ |
|__SMCR__| |__SMCD__|
> - - - - - |
|__NOSMCR__| |__NOSMCD__|
```

**Optional for SMC –
RoCE inherits VLANID
from associated OSD**


Inbound OSA/QDIO performance: Dynamic LAN idle timer

- Dynamic LAN idle timer is designed to reduce latency and improve network performance by dynamically adjusting the inbound blocking algorithm.
 - When enabled, the z/OS TCP/IP Stack is designed to adjust the inbound blocking algorithm to best match the application requirements.
- For latency sensitive applications, the blocking algorithm is modified to be "latency sensitive." For streaming (CPU sensitive) applications, the blocking algorithm is adjusted to minimize CPU utilization. In all cases, the z/OS TCP/IP stack dynamically detects the payload patterns, making the necessary adjustments to the blocking algorithm.
 - The monitoring of the application and the blocking algorithm adjustments are made in real-time, dynamically adjusting the application's LAN performance.
- System administrators can authorize the z/OS TCP/IP stack to enable the DYNAMIC setting, which was previously a static setting.
 - The z/OS TCP/IP stack is designed to dynamically determine the best setting for the current running application based on system configuration, inbound workload volume, CPU utilization, and traffic patterns.
- For an OSA Express port in OSD/QDIO mode the z/OS TCP/IP profile **INBPERF** parameter can be specified with one of the following options:
 - **MINCPU** - a static interrupt-timing value, selected to minimize host interrupts without regard to throughput
 - **MINLATENCY** - a static interrupt-timing value, selected to minimize latency
 - **BALANCED** (default) - a static interrupt-timing value, selected to achieve reasonably high throughput and reasonably low CPU
 - **DYNAMIC** - implements the dynamic LAN idle algorithm ← **recommended!**

Dynamic LAN Idle Timer: Performance Data


Dynamic LAN Idle increased a single RR connection's trans/second by 50% and ten concurrent RR connections' trans/second by 88%. Response time for these workloads is improved by 33% and 47%, respectively.

RR1 and RR10 Dynamic LAN Idle




- RR indicates request/response workload
- 1h/8h indicates 100 bytes in and 800 bytes out

Note: 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.



Which trace should I use?




OSA-Express Network Traffic Analyzer (OSAENTA) function

- Diagnosing OSA-Express QDIO problems can be somewhat difficult
 - TCP/IP stack (CTRACE and/or packet trace)
 - VTAM (VIT)
 - OSA (hardware trace)
 - Network (sniffer trace)
- Often not clear where the problem is and which trace(s) to collect
 - Offloaded functions and shared OSAs can complicate the diagnosis
 - ARP frames
 - Segmentation offload
- OSAENTA allows z/OS Communications Server to collect Ethernet data frames from the OSA adapter
 - Not a sniffer trace (but similar in some aspects)
 - Minimizes the need to collect and coordinate multiple traces for diagnosis
 - Minimizes the need for traces from the OSA Hardware Management Console (HMC)
 - Formatted with the standard z/OS CS packet trace formatter
 - Including the ability to convert to Sniffer/Wireshark format

QDIO outbound priority queue selection (based on WLM priority)



Basic principle is that if QoS policies are active, they will determine which priority queue to use.



Use workload objective of application address spaces to select QDIO outbound priority queue

An easy way to take advantage of QDIO's four outbound priority queues!

OSA Express (QDIO) WLM Outbound Priority Queuing




- Request-Response and Streaming mixed workloads
 - RR1/STR10: 1 RR session, 100 / 800 and 10 STR sessions, 1 / 20 MB
 - RR5/STR10: 5 RR sessions, 100 / 800 and 10 STR sessions, 1 / 20 MB
- WLM PRIORITYQ assigned importance level 2 to interactive workloads and level 3 to streaming workloads
 - The z/OS Workload Manager (WLM) system administrator assigns each job to its respective WLM service class
- z/OS with WLM I/O Priority provides 30 to 50% higher throughput for interactive workloads compared to without WLM I/O Priority
- z/OS with WLM I/O Priority provides 23 to 33% lower latency compared to without WLM I/O Priority

Note: 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.

OSA interface isolation



- Allow customers to disable shared OSA local routing functions
 - ISOLATE/NOISOLATE option on QDIO network INTERFACE definition
- OSA local routing can in some scenarios be seen as an issue as the packet may bypass packet inspection tools



If you enable ISOLATE, packets with a nexthop IP address of another stack that share the OSA port, will be discarded.

QDIO Accelerator

- 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
 - Allows for acceleration of Sysplex Distributed traffic when IP Security enabled



CPU / Transaction

A bar chart titled 'CPU / Transaction' comparing performance relative to a 'no QDIO Accelerator' baseline. The y-axis is labeled '% (Relative to no QDIO Accelerator)' and ranges from -75 to 75 in increments of 25. The x-axis has three categories: 'z/OS Client', 'z/OS Sysplex Distributor', and 'z/OS Target'. The 'z/OS Client' bar is green and labeled -0.9. The 'z/OS Sysplex Distributor' bar is blue and labeled -57.45. The 'z/OS Target' bar is green and labeled -0.33.


Category	% (Relative to no QDIO Accelerator)
z/OS Client	-0.9
z/OS Sysplex Distributor	-57.45
z/OS Target	-0.33

25




OSA Inbound Workload Queuing (IWQ) support

- Allow inbound QDIO traffic separation by supporting multiple read queues
 - “Register” with OSA which traffic goes to which queue
 - OSA Data Router function routes to the correct queue
- Each input queue can be serviced by a separate process
 - Primary input queue for general traffic
 - One or more ancillary input queues (AIQs) for specific traffic types
- Supported traffic types
 - Bulk data traffic queue
 - Serviced from a single process – reduces out of order packet delivery issues
 - Together PH44968 and PH43504 add enhancements to better detect streaming workloads - Network File System (NFS)
 - Sysplex distributor traffic queue
 - SD traffic efficiently accelerated or presented to target application
 - Enterprise Extender (EE) queue
 - All other traffic not backed up behind bulk data, EE, or SD traffic (Primary Queue)
- Dynamic LAN idle timer set per queue



TCP/IP defines and assigns traffic to queues dynamically based on local IP address and port

Inbound Workload Queuing: Mixed workload 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)



Note: 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.

Inbound Workload Queuing: EE performance Data

OSA 1GBe - mixed TCP and EE workloads



Transaction rate improved 33% for Enterprise Extender workload - Enterprise Extender queue provides internal optimizations such as avoiding a memory copy of data and no delays processing other workload traffic before EE

Note 1. 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.

Note 2. Enabling IWQ requires additional fixed CSM storage. Reference: IP Configuration Reference IWQ topic:

https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.3.0/com.ibm.zos.v2r3.halg001/rgvkrakn22.htm

z/OS IWQ IPsec


IPsec traffic serviced on its own processor / thread
(optionally on zIIP)

- Processing of IPsec queue is optimized and Q1 is no longer impeded (with encrypted traffic)

Value: Performance testing shows
throughput improvement! Example:

RR10 1k encrypted mixed with RR10 1k clear

	encrypted		/	clear	
Without IWQ:	59	mbps	/	53	mbps
With IWQ:	74	mbps	/	81	mbps
improvement!	25 %		/	52 %	



IWQ IPSec Enablement

- The following protocols are registered with OSA as new routing variables (RVs) to identify IPSec traffic
 - AH protocol (IPv4 and IPv6)
 - ESP protocol (IPv4 and IPv6)
 - UDP protocol + port 4500
 - UDP-encapsulated ESP packets (IPv4-only)
- IWQ IPSec enables with IWQ (**INBPERF DYNAMIC WORKLOADQ**)


z/OS IWQ zCX

Input queue for zCX traffic (zCX DVIPAs)

- zCX traffic serviced on its own processor / thread on zIIP
- Processing of zCX queue is optimized for packets destined for zCX server DVIPAs
- Enables with IWQ and when using zCX DVIPAs
- Available on V2R4 (PTFs UJ01511 and UI66733)

Value: Separation of zCX traffic from z/OS native traffic allows z/OS TCP/IP to optimize the zCX related processing including performing the zCX processing on a zIIP.

For some zCX workloads the GP reduction in z/OS TCP/IP processing can be significant.



Operator command to query and display OSA information


- OSA/SF (Support Facility) previously used to configure OSA and display the configuration prior to today's OSD devices (QDIO).
- OSD devices exclusively use IPA signals exchanged with the host to enable and configure features and register IP addresses to OSA.
- However, there was no mechanism to display the information directly from OSA without OSA/SF.
- Use the **DISPLAY TCPIP,,OSAINFO** command:
 - Base OSA information
 - OSA address table information
 - Information related to the new multiple inbound queues

```
D TCPIP,,OSAINFO,INTFN=V603ETHG0,MAX=100


EZD0031I TCP/IP CS V2R5    TCPIP Name: TCPSVT    12:16:29
Display OSAINFO results for IntfName: V607E25C0
PortName: 07E25C0P  PortNum: 00  Datapath: 2EE3    RealAddr: 00E3
PCHID: 0230          CHPID: 13    CHPID Type: OSD  OSA code level: 2222
Gen: OSA-E7S         Active speed/mode: 25 gigabit full duplex
Media: Multimode Fiber      Jumbo frames: Yes  Isolate: No
PhysicalMACAddr: AC620D0A2882  LocallyCfgMACAddr: 000000000000
Queues defined Out: 4  In: 6    Ancillary queues in use: 3
EIB Enabled: Yes          Port Packet Drops:          0
Connection Mode: Layer 3    IPv4: No    IPv6: Yes
...
```

TCP Segmentation Offload

- Segmentation consumes (high cost) host CPU cycles in the TCP/IP stack
- OSA feature called Segmentation Offload (also referred to as “Large Send”)
 - Offload most TCP segmentation (both IPv4 and IPv6) processing to OSA
 - Decrease host CPU utilization
 - Increase data transfer efficiency for packets
 - Configured using IPCONFIG/IPCONFIG6 statement (**SEGMENTATIONOFFLOAD**)



TCP segmentation offload performance measurements



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

Send buffer size: 180K for streaming workloads

Note: 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.



OSA BEST PRACTICES

Summary of OSA Best Practices

1. Configure an IP subnet mask on the OSA INTERFACE statement
2. Always use VMAC (OSA-generated)
3. When using VLAN for an OSA, configure your switch port in Trunk Mode
4. Define ReadStorage = Max (see white paper, next chart)
5. Enable segmentation offload
6. Use INBPERF setting of DYNAMIC
7. Enable inbound workload queues (IWQ) with WORKLOADQ (optimize bulk and interactive processing)
8. Utilize outbound priority queues with WLMRIORITYQ
9. When Sysplex Distributor host, use QDIO Accelerator with VIPAROUTE (or if DATAGRAMFWD configured)
10. Utilize jumbo frames with Path MTU Discovery (and ADJUSTDVIPAMSS for VIPAROUTE)

Note. See z/OS Communications Server and OSA Best Practices white paper (next chart)

White paper on OSA-Express best practices

IBM z/OS Communications Server and OSA-Express Best Practices

This white paper is provided for the purpose of aiding IBM z/OS customers by providing a general set of considerations (a checklist) for guidance focused on configuring OSA-Express for optimizing network performance.

<https://www.ibm.com/support/pages/node/6562173>



HIPERSOCKETS CONVERGED INTERFACE (HSCI)

HSCI Introduction

HSCI addresses two main issues:

1. z/OS incompatibility with the z/VM Bridge and Linux (Layer 2) environment:

The z/VM Bridge support provides a very competitive and seamless solution for the overall zSystem networking landscape for Linux on z clients. There are many advantages (e.g. allows single IP interface per Linux guest) to the z/VM bridge model. The z/VM Bridge is an integral part of the z/VM SSI and LGR (HA) architecture.

However, z/OS does not support (L2) connectivity to Linux in this configuration.

The missing z/OS support creates a gap in the overall value of the zSystem solution (Layer 2 Linux compatibility, HS, z/VM Bridge, VSwitch etc.), rendering the overall z solution as “partial/unfinished”.


2. z/OS HiperSockets usability:

When moving a z/OS instance to another CPC, it requires manual administrative actions to reconfigure the HS interfaces to match the new IP subnet(s) on the new CPC. This applies to all z/OS environments (even in z/OS only environments). Some customers regularly / frequently move their z/OS instances. Customers have expressed concerns about this issue (i.e. how this manual step impacts their day to day operations causing some customers to rethink or stop using HS).

Note. There is additional HSCI background information in the backup charts.


Issue: z/OS and z/VM VSwitch Bridge and Linux HS L2

Focus on Server D!




z/VM VSwitch Bridge: Workaround

Focus on Server D!




HiperSockets Converged Interface Concepts



Solution: z/OS HSCI


z/OS HiperSockets Converged Interface (HSCI)



z/OS HSCI and z/VM VSwitch Bridge L2 Compatibility

Focus on Server D!

CPC X




Very Important HSCI Concept!

z/OS HSCI is **bridge compatible** but is **not bridge eligible** (network traffic **does not** go thru (use) the IQD bridge)! z/OS communicates internally with hosts that are directly connected to IQD (L2 HS) and communicates externally using OSA.


z/OS HSCI Solution: Linux (Server D) is using a single IP interface (HS L2) and can now communicate internally with **z/OS (HSCI)** and externally (transparently) over the LAN using the z/VM VSwitch Bridge (using OSA). There are no changes to Linux, z/VM or HS.

z/OS only without HSCI: Each host sees multiple subnets




Each z/OS instance has two IP interfaces. Each IP interface has an IP address, IP Routes, and unique IP subnets (along with other interface related resources).

z/OS only with HSCI: Each host now sees single subnet




With HSCI:
Each z/OS instance has one IP interface. This removes all related resources (IP address, IP Route, etc.) for the HS interface). HS is no longer seen or managed by the administrator. Admin only sees the OSA interface.

Now each z/OS image has access to 1 network (via OSA). HS is not visible to the TCP/IP stack!



With HSCI: Moving a z/OS instance requires no HS configuration changes!



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

- **Online course:**
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>

Experience more with IBM

SHARE25 August 17-21, 2025
Cleveland, OH

Visit us at the IBM Booth #317

Connect with our experts, snap a photo with [the new z17 Plexi](#) or [the latest Telum II LEGO® model](#), and get an up-close look at our [Spyre AI Accelerator](#). Come back each day for fresh topics and rotating demos at our expert stations.

Register for the IBM Z Client Reception

Join us at IBM's Monday evening reception at Punch Bowl Social following the reception in the SHARE Technology Expo.
[Monday, 18 August | 7:00-10:00 PM](#)

Register



Don't miss the IBM Keynote in Tuesday's General Session featuring Elpida Tzortzatos, IBM Fellow - CTO AI, IBM Z & LinuxONE.

[Tuesday, 19 August | 7:45 –9:00 AM | Main Ballroom, Hilton Cleveland](#)

Join our Lunch and Learn sessions


Grab a bite as we dive into how agentic AI is reshaping the IBM Z platform.

[Monday, 18 August | 12:00–1:00 PM | Hope Ballroom B, Hilton Cleveland Downtown – Level 3](#)

Join the fun with our new AI-powered Security IBM z17 edition of our Wheel of Cybersecurity Game Show!

[Thursday, 21 August | 11:15 AM–12:15 PM | Hope Ballroom B, Hilton Cleveland Downtown – Level 3](#)

IBM TechXchange October 6 - 9




Explore, build and extend AI **where it matters most**. Join us in Orlando and experience the latest **IBM Z and LinuxONE** innovations with [150+ tech breakouts, labs, tech talks, demos and much more!](#)

Register now



IBM Z Day November 12



Go deeper at our largest tech event of the year, IBM Z Day - a [free virtual conference](#) featuring 6 tracks, 250+ expert speakers and participants from over 165 countries. **Earn up to 4 exclusive IBM digital skills badges.**

Register here




Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation

“Getting the most out of your
OSA and HS”
Session ID is # **628**





Trademarks, notices, and disclaimers

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

BigInsights	DFSMSdss	FICON*	IMS	RACF*	System z10*	zEnterprise*
BlueMix	DFSMSHsm	GDPS*	Language Environment*	Rational*	Tivoli*	z/OS*
CICS*	DFSORT	HyperSwap	MQSeries*	Redbooks*	UrbanCode	zSecure
COGNOS*	DS6000*	IBM*	Parallel Sysplex*	REXX	WebSphere*	z Systems
DB2*	DS8000*	IBM (logo)*	PartnerWorld*	SmartCloud*	z13	z/VM*
DFSMSdftp					z14	

* 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 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.

Trademarks, notices, and disclaimers

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

BigInsights	DFSMSdss	FICON*	IMS	RACF*	System z10*	zEnterprise*
BlueMix	DFSMSHsm	GDPS*	Language Environment*	Rational*	Tivoli*	z/OS*
CICS*	DFSORT	HyperSwap	MQSeries*	Redbooks*	UrbanCode	zSecure
COGNOS*	DS6000*	IBM*	Parallel Sysplex*	REXX	WebSphere*	z Systems
DB2*	DS8000*	IBM (logo)*	PartnerWorld*	SmartCloud*	z13	z/VM*
DFSMSdftp					z14	

* 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 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.

Backup Information: HiperSockets Converged Interface (HSCI)


Note. Screen cast of the entire z/OS CS HSCI Overview:

<https://www.youtube.com/watch?v=o0jK5walPoQ>

System z IQD Architecture: HS Mixing Layer 2 and Layer 3 Modes

IP interfaces can be defined in either L2 or L3 transport modes.


Not all Operating Systems support both modes.



On the same IQD CHPID, the green IQD connections using L2 mode are isolated from the yellow connections using L3. L2 and L3 IQD connections **can NOT communicate** over the same IQD CHPID!

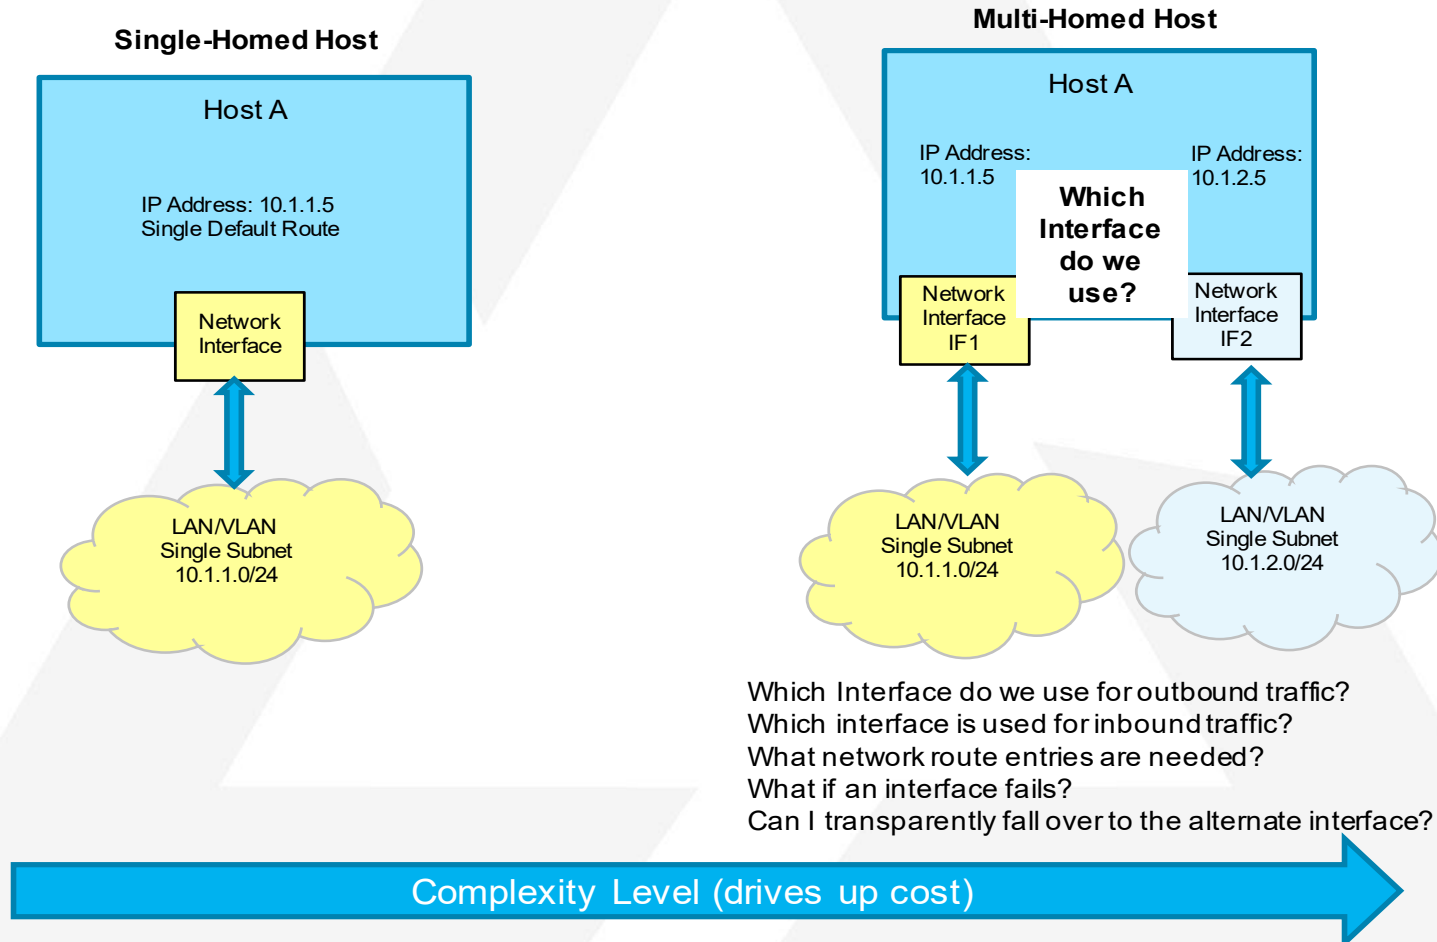
LPARs using HS in Layer 2 mode **can not** communicate with LPARs in Layer 3 mode.

Topic 1B. z/VM VSwitch HS Bridge Overview



- The z/VM VSwitch and HS **must be operating in the L2 transport mode (i.e. the Bridge requires L2)**
- No transport mode conversions
- Guests QA1, QA2, QA3 and QA4 have real (*dedicated*) QEBSM connections to HS PCHID.
 - Optimum performance configuration requiring almost no z/VM involvement
 - Bridged by default when using QDIO Assist (QEBSM, which provides connectivity to both HS and the external LAN segments)
- Guests VA1 and VA2 have virtual (NIC) connections through VSwitch A
 - Optimum performance config for guests that are not deployed with QEBSM on z/VM. Eliminates “shadow” queue overhead
 - Connectivity to HS and external LAN segments
- OSA uplink port BAU no changes in current support

Topic 1C. Single vs. Multi-Homed Hosts



z/OS 2.3 HSCI (AUTOIQDC) Highlights

1. The upper layers of the TCP/IP software stack only see an OSA Interface (i.e. the IP route(s), IP address/subnet, MTU, etc. all reflect the attributes of the OSA interface, IQDC is not visible)
2. The HSCI (L2) path is used by the lower layers of the stack when the peer host is found on the IQD network (i.e. when the next hop IP address is in the sending host's ARP/ND cache).
3. IPv4 and IPv6 are both supported.
4. The HSCI inherits the associated OSA INTERFACE attributes (IP address, VMAC, VLAN ID, MTU etc.)
5. All HSCI related artifacts are dynamically created with the following naming conventions:
 1. IQDC IPv4:
 1. Interface name = EZAIQCnn (where EZAIQC indicates IPv4 IQDC and nn = CHPID, e.g. EZAIQCE0).
 2. TRLE name = IUTIQCnn (where IUTIQC indicates IPv4 IQDC and nn = CHPID, e.g. IUTIQCEO).
 2. IQDC IPv6:
 1. Interface name = EZ6IQCnn (where EZ6IQC indicates IPv6 IQDC and nn = CHPID, e.g. EZ6IQCE0)
 2. TRLE name = IUTIC6nn (where IUTIC6 indicates IPv6 IQDC and nn = CHPID, e.g. IUTIC6EO).
6. Layer 2 architecture requires that software:
 1. transmits and receives "Ethernet frames".
 2. handles all ARP (IPv4) and Neighbor Discovery (IPv6) processing in software
7. z/OS QDIO Accelerator is supported (see backup for details)
8. AUTOIQDC can be changed with Vary Obey. The change does not impact current usage, but will impact any future attempts to use it.

z/OS 2.3 HSCI Enablement

APARs:

1. OA53198 (VTAM, FMID HVT6230)
2. PI83372 (TCP/IP, FMID HIP6230)

Notes.

1. z/OS 2.3 Announcement Letter (see HSCI on networking page 18):
[IBM z/OS 2.3 RFA URL](#)
2. The VTAM and TCP/IP PTFs are co-reqs (HSCI requires both)
3. HSCI is off by default, enabled with AUTOIQDC parameter
4. Terminology: HSCI (solution) is controlled by AUTOIQDC (parameter); HSCI = IQDC
5. There are no other software or System z hardware requirements (no other product dependencies or co-reqs)¹

Note 1. When Linux is executing as a z/VM guest virtual machine, then exploiting the z/VM VSwitch bridge is required. The IQD CHPID must be configured with External Bridge and the z/VM VSwitch Bridge support must be enabled for the same IQD CHPID that is being used by the Linux guest. Also refer to VM66135 for maintenance.