z/OS V2R4 Communications Server
IWQ support for z/OS Container Extensions

Contents

List of Tables.....v

Chapter 1: New Function Summary.....7
IWQ support for z/OS Container Extensions..... 8

Chapter 2: IP Configuration Guide.....11
QDIO inbound workload queueing..... 12

Chapter 3: IP Configuration Reference.....15
TCP/IP profile (PROFILE.TCPIP) and configuration statements..... 16
INTERFACE - IPAQENET OSA-Express QDIO interfaces statement..... 16
INTERFACE - IPAQENET6 OSA-Express QDIO interfaces statement..... 28

Chapter 4: IP System Administrator's Commands.....45
DISPLAY TCPIP,,OSAINFO..... 46
Format..... 46
Parameters..... 47

Chapter 5: SNA Messages.....55
IST1221I.....56
IST1230I.....67
IST1233I.....76

Chapter 6: z/OS Upgrade Workflow.....77
Upgrade to an IBM z14 server..... 78
IP services: Ensure storage availability for IWQ zCX traffic..... 84

List of Tables

Table 1: Task topics to enable IWQ support for z/OS Container Extensions.....8

Table 2: All related topics about IWQ support for z/OS Container Extensions.....9

Table 3: z14 server functions included in the base support for z/OS V2R2, z/OS V2R3, and z/OS V2R4.....78

Table 4: Exploitation of z14 server functions supported by z/OS V2R2, z/OS V2R3, and z/OS V2R4.....79

Table 5: Information about this upgrade action.....82

Table 6: Information about this upgrade action.....85

Chapter 1

New Function Summary

Topics:

- IWQ support for z/OS Container Extensions

IWQ support for z/OS Container Extensions

z/OS® V2R4 Communications Server, with VTAM APAR OA58300 and TCP/IP APAR PH16581, is enhanced to support inbound workload queueing for IBM z/OS Container Extensions (zCX) workloads for OSA-Express® in QDIO mode.

Inbound workload queueing uses multiple input queues for each QDIO data device (subchannel device) to improve TCP/IP stack scalability and general network optimization. To implement the performance improvements for zCX workloads, enable inbound workload queueing to process zCX, IPSec, EE, sysplex distributor, and streaming bulk data traffic all concurrently with other types of inbound QDIO traffic. When you enable these improvements for a QDIO interface, zCX, inbound IPSec, EE, sysplex distributor, and streaming bulk data traffic are each processed on their own ancillary input queue (AIQ). All other inbound traffic is processed on the primary input queue.

Incompatibilities: This function does not support IPAQENET interfaces that are defined by using the DEVICE, LINK, and HOME statements. Convert your IPAQENET definitions to use the INTERFACE statement to enable this support.

Dependencies:

- This function is limited to OSA-Express6S Ethernet features or later in QDIO mode running on IBM z14™. For more information about the QDIO inbound workload queueing function and the OSA-Express features that support it, see QDIO inbound workload queueing in *z/OS Communications Server: IP Configuration Guide*. See the 3906DEVICE or 3907DEVICE Preventive Service Planning (PSP) bucket for more information.
- This function is supported only for interfaces that are configured to use a virtual MAC (VMAC) address.

To enable IWQ support for z/OS Container Extensions, complete the appropriate tasks in Table 1: Task topics to enable IWQ support for z/OS Container Extensions on page 8.

Table 1: Task topics to enable IWQ support for z/OS Container Extensions

Task	Reference

Enable inbound workload queueing for a specific QDIO interface by specifying the WORKLOADQ parameter on the IPAQENET or IPAQENET6 INTERFACE statement (if not already configured).	• See the following statements in *z/OS Communications Server: IP Configuration Reference*:
	  • INTERFACE-IPAQENET OSA-Express QDIO interfaces
	  • INTERFACE-IPAQENET6 OSA-Express QDIO interfaces
	• Steps for enabling QDIO inbound workload queueing in *z/OS Communications Server: IP Configuration Guide*

Display whether inbound workload queueing is in effect for the QDIO interface by issuing the `Netstat DEvlinks/-d` command.	Netstat DEvlinks/-d report in *z/OS Communications Server: IP System Administrator's Commands*
Display whether inbound workload queueing is in effect for the QDIO interface and display the workload queueing functions and queue IDs for that interface by issuing the `DISPLAY NET,ID=trle` command or the `DISPLAY NET,TRL,TRLE=trle` command.	See the following topics in *z/OS Communications Server: SNA Operation*: • DISPLAY ID command • DISPLAY TRL command
Monitor whether inbound traffic is using inbound workload queueing and display statistics for each queue by initiating VTAM® tuning statistics for the QDIO interface.	MODIFY TNSTAT command in *z/OS Communications Server: SNA Operation*

Task	Reference
Monitor whether inbound traffic is using inbound workload queueing and display statistics for each queue by using the TCP/IP callable NMI GetIfStatsExtended request.	TCP/IP callable NMI (EZBNMIFR) in *z/OS Communications Server: IP Programmer's Guide and Reference*
Determine the QID on which a specific packet was received, and the associated workload queueing function, from a packet trace.	Formatting packet traces using IPCS in *z/OS Communications Server: IP Diagnosis Guide*
Determine the QID on which a specific packet was received from an OSAENTA trace.	Formatting OSA traces using IPCS in *z/OS Communications Server: IP Diagnosis Guide*

To find all related topics about IWQ support for z/OS Container Extensions, see Table 2: All related topics about IWQ support for z/OS Container Extensions on page 9.

Table 2: All related topics about IWQ support for z/OS Container Extensions

Book name	Topics
z/OS Communications Server: IP Configuration Guide	• QDIO inbound workload queueing
z/OS Communications Server: IP Configuration Reference	• INTERFACE-IPAQENET OSA-Express QDIO interfaces • INTERFACE-IPAQENET6 OSA-Express QDIO interfaces
z/OS Communications Server: IP Programmer's Guide and Reference	• TCP/IP callable NMI (EZBNMIFR)
z/OS Communications Server: IP Diagnosis Guide	• Formatting packet traces using IPCS • Formatting OSA traces using IPCS
z/OS Communications Server: IP System Administrator's Commands	• DISPLAY TCPIP,,OSAINFO • Netstat DEvlinks/-d report
z/OS Communications Server: SNA Operation	• DISPLAY ID command • DISPLAY TRL command • MODIFY TNSTAT command
z/OS Communications Server: SNA Messages	• IST1221I • IST1230I • IST1233I

Chapter 2

IP Configuration Guide

Topics:

- QDIO inbound workload queueing

QDIO inbound workload queueing

OSA-Express3 or later features can perform a degree of traffic sorting by placing inbound packets for differing workload types on separate processing queues. This function is called QDIO inbound workload queueing (IWQ). With the inbound traffic stream already sorted by the OSA-Express feature, z/OS Communications Server provides the following performance optimizations:

- Finer tuning of read-side interrupt frequency to match the latency demands of the various workloads that are serviced
- Improved multiprocessor scalability, because the multiple OSA-Express input queues are now efficiently serviced in parallel

When QDIO IWQ is enabled, z/OS Communications Server and the OSA-Express feature establish a primary input queue and one or more ancillary input queues (AIQs), each with a unique read queue identifier (QID) for inbound traffic. z/OS Communications Server and the OSA-Express feature cooperatively use the multiple queues in the following way:

- The OSA-Express feature directs an inbound packet (received on this interface) that is to be forwarded by the sysplex distributor to the sysplex distributor AIQ. z/OS Communications Server then tailors its processing for the sysplex distributor queue, notably by using the multiprocessor to service sysplex distributor traffic in parallel with traffic on the other queues.
- The TCP layer automatically detects connections operating in a bulk-data fashion (such as the FTP data connection), and these connections are registered to the receiving OSA-Express feature as bulk-mode connections. The OSA-Express feature then directs an inbound packet (received on this interface) for any registered bulk-mode connection to the TCP bulk-data AIQ. z/OS Communications Server tailors its processing for the bulk queue, notably by improving in-order packet delivery on multiprocessors, which likely results in improvements to CPU consumption and throughput. Like other AIQs, processing for data on the bulk queue can be in parallel with traffic on the other queues.
- The OSA-Express feature directs an inbound Enterprise Extender packet (received on this interface) to the Enterprise Extender AIQ. This allows z/OS Communications Server to process inbound traffic on the Enterprise Extender queue in parallel with inbound traffic on the other queues for this interface.
- The OSA-Express feature directs inbound AH packets, ESP packets, and UDP-encapsulated ESP packets that are received on this interface to the IPSec AIQ. This allows z/OS Communications Server to process inbound traffic on the IPSec queue in parallel with inbound traffic on the other queues for this interface.
- The OSA-Express feature directs inbound packets (received on this interface) that are destined to a zCX server to the zCX AIQ. This allows z/OS Communications Server to process inbound traffic on the zCX queue in parallel with inbound traffic on the other queues for this interface. Additionally, it also allows for this inbound traffic to be processed on a System z Integrated Information Processor (zIIP).
- If a packet is not directed to an AIQ, the OSA-Express feature directs the packet to the primary input queue.

Requirement: QDIO IWQ is limited to OSA-Express3 or later Ethernet features in QDIO mode that are running on an IBM® System z10 GA3 or later server. For more information, see the 2097DEVICE, 2098DEVICE, or 2817DEVICE Preventive Service Planning (PSP) buckets as appropriate.

Requirement: QDIO IWQ for IPSec and zCX is limited to OSA-Express6S or later Ethernet features in QDIO mode that are running on an IBM z14 or later server. For more information, see the 3906DEVICE or 3907DEVICE Preventive Service Planning (PSP) bucket.

Results:

- If you already enabled IWQ on an OSA-Express6S, installed the IWQ for IPSec enablement PTF (TCP/IP APAR PI77649), and have IPSec enabled, z/OS Communications Server will automatically define the IPSec input queue. This new queue will not be backed by CSM storage until the first IPSec tunnel is activated on that stack. Therefore, if you activate any IPSec tunnels, each OSA-Express6S interface on that stack that enables IWQ will consume an additional 36 K of fixed ECSA and an additional amount of fixed 64-bit CSM equal to the READSTORAGE setting for the interface.
- If you already enabled IWQ on an OSA-Express6S and installed the IWQ for zCX enablement PTFs (VTAM APAR OA58300 and TCP/IP APAR PH16581), z/OS Communications Server will automatically define the zCX

input queue. This new queue will not be backed by CSM storage until the first zCX dynamic VIPA is activated on that stack. Therefore, if you activate any zCX dynamic VIPA, each OSA-Express6S interface on that stack that enables IWQ will consume an additional 36 K of fixed ECSA and an additional amount of fixed 64-bit CSM equal to the READSTORAGE setting for the interface.

Restrictions:

- QDIO IWQ is not supported for IPAQENET interfaces defined with the DEVICE, LINK, and HOME statements. You must convert your IPAQENET definitions to use the INTERFACE statement to enable this support. For more information, see Steps for converting from IPv4 IPAQENET DEVICE, LINK, and HOME definitions to the IPv4 IPAQENET INTERFACE statement.
- QDIO IWQ is not supported for a z/OS guest on z/VM® using simulated (virtual) devices, such as virtual switch (VSWITCH) or guest LAN.
- Bulk-mode TCP connection registration is supported only in configurations in which a single inbound interface is servicing the bulk-mode TCP connection. If a bulk-mode TCP connection detects that it is receiving data over multiple interfaces, QDIO IWQ is disabled for the TCP connection and inbound data from that point forward is delivered to the primary input queue.
- QDIO IWQ does not apply for traffic that is sent over an OSA port that is shared by the receiving TCP/IP stack when an indirect route (where the next hop and destination IP address are different) is being used; this traffic is placed on the primary input queue. QDIO IWQ does apply when traffic on the shared OSA path uses a direct route (where the next hop and destination IP address are the same).

Guideline: The WORKLOADQ parameter requires an additional amount of fixed 4K CSM HVCOMMON storage per AIQ. The amount of storage consumed per AIQ is based on the amount of storage defined for READSTORAGE for this interface. The bulk AIQ is always backed with this additional CSM storage. The remaining AIQs are not backed by the additional CSM storage until the specific function (EE, SD, IPSec, or zCX) is used. The EE AIQ is backed by fixed 4K CSM DSPACE64 storage (instead of HVCOMMON). To verify the amount of CSM storage that is being used for each input queue, display the VTAM TRLE name that is associated with the interface. The WORKLOADQ parameter also requires an additional 36K of ECSA per AIQ.

Tip: The additional CSM storage consumed by each OSA interface using WORKLOADQ consumes fixed (real) storage. It is recommended that you verify that the additional fixed storage required by enabling WORKLOADQ (per OSA interface) will not approach any of the following limits:

- The CSM FIXED MAXIMUM value used by Communications Server (use the D NET, CSM command and see the CSM FIXED MAXIMUM value defined in IVTPRM00)
- The actual real storage available to this z/OS system (see D M=STOR or D M=HIGH)

Guideline: CPU might increase significantly on a busy Sysplex Distributor if QDIO Accelerator is not installed when using IWQ.

Chapter 3

IP Configuration Reference

Topics:

- TCP/IP profile (PROFILE.TCPIP) and configuration statements

TCP/IP profile (PROFILE.TCPIP) and configuration statements

This topic contains the following information:

- Summary of TCP/IP address space configuration statements
- PROFILE.TCPIP search order
- Statement syntax for configuration statements
- Statements and descriptions

Configuring the stack for IPv6 is done in the BPXPRMxx member of SYS1.PARMLIB. For more information about configuring the stack to support IPv6, see *z/OS Communications Server: IP Configuration Guide* or *z/OS Communications Server: IPv6 Network and Appl Design Guide*.

INTERFACE - IPAQENET OSA-Express QDIO interfaces statement

Use the INTERFACE statement to specify an OSA-Express QDIO Ethernet interface for IPv4.

Restriction: This statement applies to IPv4 IP addresses only.

To determine the OSA-Express microcode level, use the DISPLAY TRL command. If a specific OSA-Express function is documented with a minimum microcode level, you can use this command to determine whether that function is supported. IBM service might request the microcode level for problem diagnosis. For more information about the DISPLAY TRL command, see *z/OS Communications Server: SNA Operation*.

For more information about OSA-Express features that support QDIO mode, see *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

When you start an IPAQENET interface (and you did not specify VMAC with ROUTEALL), TCP/IP registers all non-loopback local (home) IPv4 addresses for this TCP/IP instance to the OSA-Express feature. If you subsequently add, delete, or change any home IPv4 addresses on this TCP/IP instance, TCP/IP dynamically registers the changes to the OSA-Express feature. The OSA adapter routes datagrams destined for those IPv4 addresses to this TCP/IP instance.

If a datagram is received at the OSA adapter for an unregistered IPv4 address, then the OSA-Express feature routes the datagram to the TCP/IP instance, depending on the setting of a virtual MAC (VMAC) address or definition of an instance as PRIROUTER or SECROUTER. If the datagram is not destined for a virtual MAC address and no active TCP/IP instance using this interface is defined as PRIROUTER or SECROUTER, then the OSA-Express feature discards the datagram. See the router information in *z/OS Communications Server: IP Configuration Guide* for more details and primary and secondary routing in *z/OS Communications Server: SNA Network Implementation Guide*.

For detailed instructions on setting up an OSA-Express feature, see *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

For more information about missing interrupt handler (MIH) considerations with TCP/IP interfaces, see Missing interrupt handler factors.

Syntax

Rule: Specify the required parameters and the CHPIDTYPE parameter in the order shown here. The OSD Interface Definition and OSX Interface Definition parameters can be specified in any order.

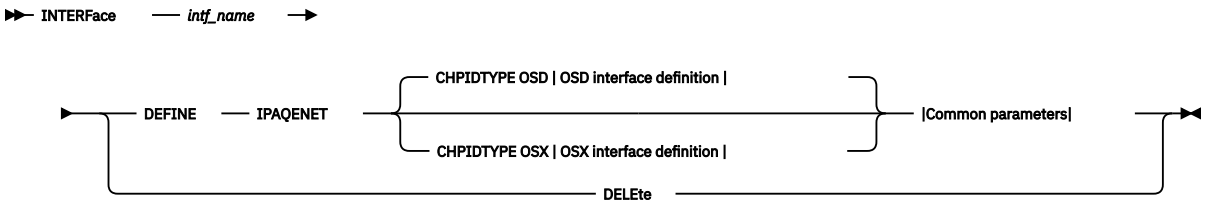

Figure 1:

OSX Interface Definition

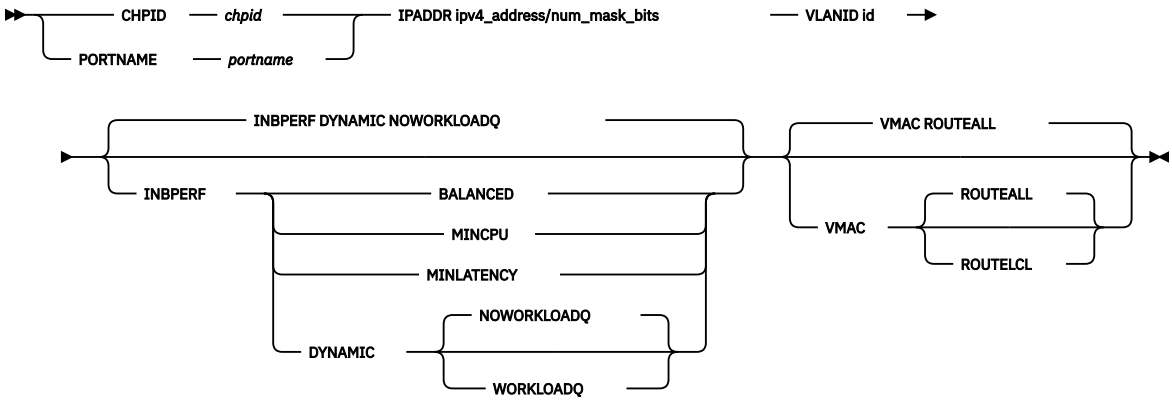

Common parameters for OSD and OSX interface definitions

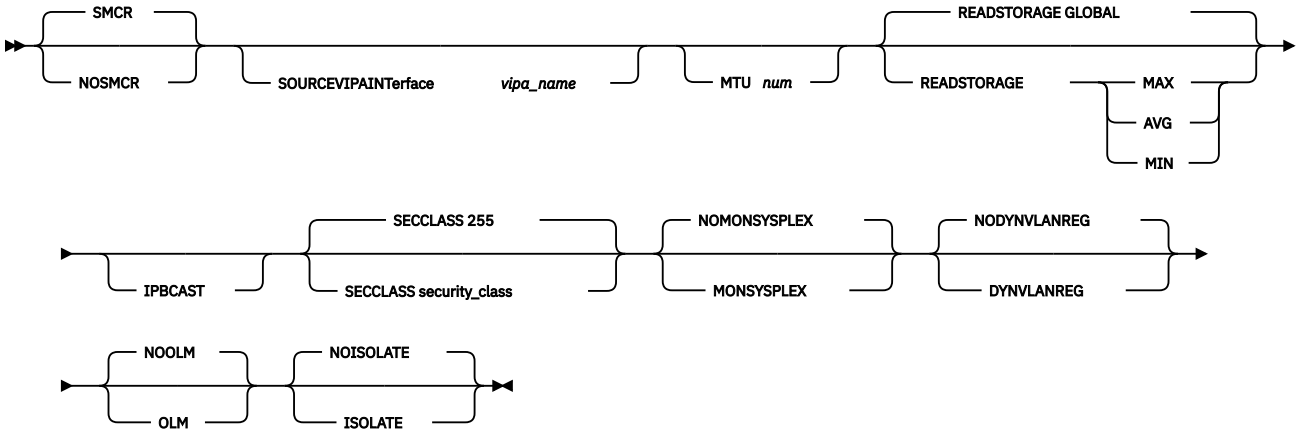

Parameters

intf_name

The name of the interface. The maximum length is 16 characters.

Requirement: This name must be different than the name specified for the PORTNAME parameter.

DEFINE

Specifies that this definition is to be added to the list of defined interfaces.

DELETE

Specifies that this definition is to be deleted from the list of defined interfaces. The *intf_name* value must be the name of an interface previously defined by an INTERFACE statement. Specifying INTERFACE DELETE deletes the home IP address for the interface.

IPAQENET

Indicates that the interface uses the interface based on IP assist, which belongs to the QDIO family of interfaces, and uses the Ethernet protocol.

CHPIDTYPE

An optional parameter indicating the CHPID type of the OSA-Express QDIO interface.

OSD	Indicates an external data network type. This is the default value.
OSX	The intraensemble data network. See *z/OS Communications Server: IP Configuration Guide* for information about requirements necessary to make an OSX work.

CHPID ***chpid***

This parameter applies only to interfaces of CHPIDTYPE OSX and is used to specify the CHPID for the interface. This value is a 2-character hexadecimal value (00 - FF).

PORTNAME ***portname***

Use this parameter to specify the PORT name that is in the TRLE definition for the QDIO interface. The TRLE must be defined as MPCLEVEL=QDIO. For details about defining a TRLE, see *z/OS Communications Server: SNA Resource Definition Reference*.

Requirement: The *portname* value must be different than the name specified for *intf_name*.

IPADDR

ipv4_address	The home IP address for this interface. **Requirement:** The IP address must be specified in dotted decimal form.
num_mask_bits	An integer value in the range 0 - 32 that represents the number of leftmost significant bits for the subnet mask of the interface. This value also controls how ARP processing for VIPAs is handled for this interface. When you specify a nonzero value, the TCP/IP stack informs OSA to perform ARP processing for a VIPA only if the VIPA is configured in the same subnet as the OSA (as defined by this subnet mask). The default is 0 for CHPIDTYPE OSD. This parameter is required for CHPIDTYPE OSX.. **Requirement:** If you are configuring multiple IPv4 VLAN interfaces to the same OSA-Express feature, then you must specify a nonzero value for the *num_mask_bits* variable for each of these interfaces and the resulting subnet must be unique for each of these interfaces. **Rule:** If you are using OMPROUTE and OMPROUTE is not configured to ignore this interface, ensure that the subnet mask value that you define on this parameter matches the subnet mask used by OMPROUTE for this interface. The subnet mask used by OMPROUTE is the subnet mask value defined on the corresponding OMPROUTE statement (OSPF_INTERFACE,

RIP_INTERFACE, or INTERFACE) for this interface. If no OMPROUTE statement is specified for this interface, the subnet mask used by OMPROUTE is the class mask for the interface IP address.

TEMPIP

Specifies that the interface starts with an IP address of 0.0.0.0. The interface can be used for broadcast traffic. This parameter applies only to interfaces that are defined with CHPIDTYPE OSD.

Guideline: Use TEMPIP interfaces in a unit test environment to support an application that provides a DHCP client, such as IBM Rational Developer for System z® Unit Test feature (Rdz-UT). For more information about configuring a TEMPIP interface, see Using TEMPIP interfaces in *z/OS Communications Server: IP Configuration Guide*.

NONROUTER

If a datagram is received at this interface for an unknown IP address, the datagram is not routed to this TCP/IP instance. This is the default value.

The PRIROUTER and SECROUTER parameters interact with the VLANID parameter. See the VLANID parameter definition to understand this relationship.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

PRIROUTER

If a datagram is received at this interface for an unknown IP address and is not destined for a virtual MAC, the datagram is routed to this TCP/IP instance. This parameter interacts with the VLANID parameter. See the VLANID parameter definition to understand this relationship.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

SECROUTER

If a datagram is received at this interface for an unknown IP address and is not destined for a virtual MAC, and there is no active TCP/IP instance defined as PRIROUTER, then the datagram is routed to this TCP/IP instance. This parameter interacts with the VLANID parameter. See the VLANID parameter definition to understand this relationship.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

VLANID id

Specifies the decimal virtual LAN identifier to be assigned to the OSA-Express interface. This field should be a virtual LAN identifier recognized by the switch for the LAN that is connected to this OSA-Express interface. The valid range is 1 - 4094. This parameter is optional for CHPIDTYPE OSD and required for CHPIDTYPE OSX.

Guideline: Installation configuration on other platforms or related to Ensemble networking can limit the maximum VLANID of 4096.

The VLANID parameter interacts with the PRIROUTER and SECROUTER parameters. If you configure both the VLANID parameter and either PRIROUTER or SECROUTER parameter, then this TCP/IP instance acts as a router for this VLAN (ID) only. Datagrams that are received at this device instance for an unknown IP address and are not destined for a virtual MAC are routed only to this TCP/IP instance if it is VLAN tagged with this VLAN ID. For more information about VLANID parameter interactions, see OSA VLAN in *z/OS Communications Server: IP Configuration Guide*.

Rule: If you are configuring multiple VLAN interfaces to the same OSA-Express feature, then you must specify the VMAC parameter (with the default ROUTEALL attribute) on the INTERFACE statement for each of these interfaces.

Restriction: The stack supports a maximum of 32 IPv4 VLAN interfaces to the same OSA-Express port. Additional VLANID limitations might exist if this interface can be used with Shared Memory Communications over Remote Direct Memory Access (SMC-R). See VLANID considerations in *z/OS Communications Server: IP Configuration Guide* for details.

INBPERF

An optional parameter that indicates how processing of inbound traffic for the QDIO interface is performed.

There are three supported static settings that indicate how frequently the adapter should interrupt the host for inbound traffic: BALANCED, MINCPU, and MINLATENCY. The static settings use static interrupt-timing values. The static values are not always optimal for all workload types or traffic patterns, and the static values cannot account for changes in traffic patterns.

There is also one supported dynamic setting (DYNAMIC). This setting causes the host (stack) to dynamically adjust the timer-interrupt value while the device is active and in use. This function exploits an OSA hardware function called Dynamic LAN Idle. Unlike the static settings, the DYNAMIC setting reacts to changes in traffic patterns and sets the interrupt-timing values to maximize throughput. The dynamic setting does not incur additional CPU consumption that might be produced when you specify any of the static settings. In addition, the DYNAMIC setting uses the OSA Dynamic Router Architecture function to enable QDIO inbound workload queues for specific inbound traffic types.

Result: When you specify OLM on the INTERFACE statement, the INBPERF parameter is ignored and the statement takes the value DYNAMIC.

Valid values for INBPERF are:

BALANCED

This setting uses a static interrupt-timing value, which is selected to achieve reasonably high throughput and reasonably low CPU consumption. This is the default value for CHPIDTYPE OSD.

DYNAMIC

This setting causes the host to dynamically signal the OSA-Express feature to change the timer-interrupt value, based on current inbound workload conditions. The DYNAMIC setting is effective only for OSA-Express2 or later features on at least an IBM System z9® that supports the corresponding Dynamic LAN Idle function. See the 2097DEVICE Preventive Service Planning (PSP) bucket for more information about the OSA-Express3 adapter that supports this function. The DYNAMIC setting should outperform the other three static settings for most workload combinations. This is the default value for CHPIDTYPE OSX.

If the DYNAMIC setting is specified for an OSA-Express adapter that does not support the dynamic LAN Idle function, the stack reverts to using the BALANCED setting.

WORKLOADQ | NOWORKLOADQ

This subparameter controls the QDIO inbound workload queueing function for the interface. QDIO inbound workload queueing is effective only for OSA-Express features in QDIO mode that support the corresponding Data Router Architecture. OSA-Express features that support workload queueing do not necessarily support workload queueing for all possible traffic types. For more information about the QDIO inbound workload queueing function and the OSA-Express features that support it, see QDIO inbound workload queueing in *z/OS Communications Server: IP Configuration Guide*.

NOWORKLOADQ	Specifies that QDIO inbound workload queueing is not enabled for inbound traffic. All

inbound traffic for this interface uses a single input queue. This is the default value.

WORKLOADQ

Specifies that QDIO inbound workload queueing (IWQ) is enabled for inbound traffic.

If the WORKLOADQ subparameter is specified, QDIO inbound workload queueing is enabled for specific inbound traffic types. A primary input queue is reserved for all other traffic types.

Ancillary input queues (AIQs) are created for the following inbound traffic types when supported by the OSA-Express feature:

- Sysplex distributor
- Streaming workloads (for example FTP)
- Enterprise Extender (EE)
- IP Security (IPSec)
- IBM z/OS Container Extensions (zCX)

Requirement: You must specify the VMAC parameter with WORKLOADQ to enable QDIO inbound workload queueing.

Restrictions:

- Bulk-mode TCP connection registration is supported only in configurations in which a single inbound interface is servicing the bulk-mode TCP connection. If a bulk-mode TCP connection detects that it is receiving data over multiple interfaces, QDIO IWQ is disabled for the TCP connection and inbound data from that point forward is delivered to the primary input queue.
- QDIO IWQ does not apply for traffic that is sent over an OSA port that is shared by the receiving TCP/IP stack when an indirect route (where the next hop and destination IP address are different) is being used; this traffic is placed on the primary input queue. QDIO IWQ does apply when traffic on the shared OSA path uses a direct route (where the next hop and destination IP address are the same).

Guideline: The WORKLOADQ parameter requires an additional amount of fixed 4K CSM HVCOMMON storage per AIQ. The amount of storage consumed per AIQ is based on the amount of storage defined for READSTORAGE for this interface. The bulk AIQ is always backed with this additional CSM storage. The remaining AIQs are not backed by the additional CSM storage until the specific function (EE, SD, IPSec, or zCX) is used. The EE AIQ is backed by fixed 4K CSM DSPACE64 storage

(instead of HVCOMMON). To verify the amount of CSM storage that is being used for each input queue, display the VTAM TRLE name that is associated with the interface. The WORKLOADQ parameter also requires an additional 36K of ECSA per AIQ.

Tip: The additional CSM storage consumed by each OSA interface using WORKLOADQ consumes fixed (real) storage. It is recommended that you verify that the additional fixed storage required by enabling WORKLOADQ (per OSA interface) will not approach any of the following limits:

- The CSM FIXED MAXIMUM value used by Communications Server (use the D NET, CSM command and see the CSM FIXED MAXIMUM value defined in IVTPRM00)
- The actual real storage available to this z/OS system (see D M=STOR or D M=HIGH)

If the WORKLOADQ setting is specified for an OSA-Express adapter that does not support the Data Router Architecture function, the stack reverts to using a single input queue.

MINCPU

This setting uses a static interrupt-timing value, which is selected to minimize host interrupts without regard to throughput. This mode of operation might result in minor queueing delays (latency) for packets flowing into the host, which is not optimal for workloads with demanding latency requirements.

MINLATENCY

This setting uses a static interrupt-timing value, which is selected to minimize latency (delay), by more aggressively sending received packets to the host. This mode of operation generally results in higher CPU consumption than the other three settings. Use this setting only if host CPU consumption is not an issue.

VMAC *macaddr*

Specifies the virtual MAC address, which can be represented by 12 hexadecimal characters. The OSA-Express device uses this address rather than the physical MAC address of the device for all IPv4 packets sent to and received from this TCP/IP stack. For CHPIDTYPE OSD, using a virtual MAC address is optional. For CHPIDTYPE OSX, using a virtual MAC address is required, so the VMAC parameter is the default

The *macaddr* value is optional. The *macaddr* value is optional for CHPIDTYPE OSD and cannot be specified for CHPIDTYPE OSX. If you do not code the *macaddr* value, then the OSA-Express device generates a virtual MAC address. If you do code the *macaddr* value, it must be defined as a locally administered individual MAC address. This means the MAC address must have bit 6 (the universal or local flag U bit) of the first byte set to 1 and bit 7 (the group or individual flag G bit) of the first byte set to 0. The second hexadecimal character must be 2, 6, A, or E. The bit positions within the 12 hexadecimal characters are indicated as follows:

```
|               1|1               3|3               4|
|0              5|6               1|2               7|
+----------------+----------------+----------------+
|xxxxxxUGxxxxxxxx|xxxxxxxxxxxxxxxx|xxxxxxxxxxxxxxxx|
```

```
+----------------+----------------+----------------+
```

Rules:

- The same virtual MAC address generated by the OSA-Express device during interface activation remains in effect for this OSA-Express for this TCP/IP stack, even if the interface is stopped or becomes inoperative (INOPs). A new virtual MAC address is generated only if the INTERFACE statement is deleted and redefined or if the TCP/IP stack is recycled.
- The NONROUTER, PRIROUTER, and SECROUTER parameters are ignored for an OSA-Express interface if the VMAC parameter is configured on the INTERFACE statement.

Guideline: Unless the virtual MAC address representing this OSA-Express device must remain the same even after TCP/IP termination and restart, configure VMAC without a *macaddr* value and allow the OSA-Express device to generate it. This guarantees that the VMAC address is unique from all other physical MAC addresses and from all other VMAC addresses generated by any OSA-Express feature.

ROUTEALL

Specifies that all IP traffic destined to the virtual MAC is forwarded by the OSA-Express device to the TCP/IP stack. This is the default value. See the router information in *z/OS Communications Server: IP Configuration Guide* for more details.

ROUTELCL

Specifies that only traffic destined to the virtual MAC and whose destination IP address is registered with the OSA-Express device by this TCP/IP stack is forwarded by the OSA-Express. See the router information in *z/OS Communications Server: IP Configuration Guide* for more details.

SMCR | NOSMCR

Specifies whether this interface can be used with Shared Memory Communications over Remote Direct Memory Access (SMC-R) for IEDN or external data network communications.

NOSMCR	Specifies that this interface cannot be used for new TCP connections with SMC-R for IEDN or external data network communications.
SMCR	Specifies that this interface can be used for new TCP connections with SMC-R for IEDN or external data network communications. This is the default setting.

Rules:

- SMCR has no effect unless at least one Peripheral Component Interconnect Express (PCIe) function ID (PFID) value is specified by using the PFID subparameter of the SMCR parameter on the GLOBALCONFIG statement.
- SMCR has no effect unless a nonzero subnet mask is configured on the INTERFACE statement.

Guideline: If you enable Multipath and equal-cost interfaces are associated with different IP subnets, enabling SMC for some of, but not all, the interfaces can cause unpredictable SMC usage. You must specify either SMCR or NOSMCR on all equal-cost interfaces.

SMCD | NOSMCD

Specifies whether this interface can be used with Shared Memory Communications - Direct Memory Access (SMC-D).

NOSMCD	Specifies that this interface cannot be used for new TCP connections with SMC-D.
SMCD	Specifies that this interface can be used for new TCP connections with SMC-D. This is the default setting.

Rules:

- SMCD and NOSMCD are valid with CHPIDTYPE OSD definitions only.
- SMCD has no effect unless a nonzero subnet mask is configured on the INTERFACE statement.

Guideline: If you enable Multipath and equal-cost interfaces are associated with different IP subnets, enabling SMC for some of, but not all, the interfaces can cause unpredictable SMC usage. You must specify either SMCD or NOSMCD on all equal-cost interfaces.

SOURCEVIPAINTERFACE ***vipa_name***

Specifies which previously-defined static VIPA interface is used for SOURCEVIPA (when IPCONFIG SOURCEVIPA is in effect). The *vipa_name* value is the interface name (or link name) for a VIRTUAL interface. This parameter is optional.

Requirement: The VIRTUAL interface must be defined prior to specifying this INTERFACE statement to the TCP/IP stack. It must either already be defined, or the INTERFACE statement (or DEVICE and LINK statements) that define the static VIPA must precede this INTERFACE statement in the profile data set.

Tip: The SOURCEVIPAINTERFACE setting can be overridden. See the information about Source IP address selection in *z/OS Communications Server: IP Configuration Guide* for the hierarchy of ways that the source IP address of an outbound packet is determined.

MTU ***num***

The maximum transmission unit (MTU), in bytes. This value can be in the range 576 - 8992. The minimum MTU for IPv4 is 576. The stack takes the minimum of the configured value and the value supported by the device (returned by OSA).

The MTU default, which depends on the value that is supported by the device, is the following value:

- Gigabit Ethernet default MTU = 8992
- Fast Ethernet default MTU = 1492

The MTU default is 1492 for Fast Ethernet; otherwise, it is 8992.

Rule: If you are using OMPROUTE and OMPROUTE is not configured to ignore this interface, ensure that the MTU that you define on this parameter matches the MTU used by OMPROUTE for this interface. The MTU used by OMPROUTE is the MTU value defined on the corresponding OMPROUTE statement (OSPF_INTERFACE, RIP_INTERFACE, or INTERFACE) for this interface. If an MTU value is not defined on the corresponding OMPROUTE statement for this interface or if no OMPROUTE statement is specified for this interface, the MTU used by OMPROUTE is the minimum MTU for IPv4 (576).

Tip: See *z/OS Communications Server: IP Configuration Guide*, in section Maximum transmission unit considerations, for additional information about how TCP/IP uses the MTU to determine the largest size frame to send.

READSTORAGE

An optional parameter indicating the amount of fixed storage that z/OS Communications Server should keep available for read processing for each input queue for this interface. Multiple input queues are used when WORKLOADQ is enabled. See description of the WORKLOADQ parameter for more details.

Use the QDIOSTG VTAM start option to specify a value that applies to all OSA-Express adapters in QDIO mode. You can use the READSTORAGE keyword to override the global QDIOSTG value for this adapter based on the inbound workload that you expect over this interface on this stack. The valid values for READSTORAGE are:

GLOBAL	The amount of storage is determined by the QDIOSTG VTAM start option. This is the default value.

MAX	Use this value if you expect a heavy inbound workload over this interface.
AVG	Use this value if you expect a medium inbound workload over this interface.
MIN	Use this value if you expect a light inbound workload over this interface.

Tip: See the description of the QDIOSTG VTAM start option in the *z/OS Communications Server: SNA Resource Definition Reference* for details about exactly how much storage is allocated by z/OS Communications Server for each of these values.

Tip: The amount of storage which corresponds to each of these values is different for an OSA with at least 25 GbE of bandwidth than for an OSA with less than 25 GbE of bandwidth. See QDIOSTG start option in *z/OS Communications Server: SNA Resource Definition Reference* for details about exactly how much storage is allocated by z/OS Communications Server for each of these values.

IPBCAST

Specifies that the interface both sends and receives IP broadcast packets. If this parameter is not specified, no IP broadcast packets are sent or received on this interface.

SECCLASS *security_class*

Use this parameter to associate a security class for IP filtering with this interface. For traffic over the interface to match a filter rule, the filter rule must have the same security class value as the interface or a value of 0. You can specify filter rules in the TCP/IP profile or in an IP security policy file that is read by the Policy Agent. Filter rules can include a security class specification on the IpService statement in an IP Security policy file or on the SECCLASS parameter on the IPSECRULE statement in the TCP/IP profile.

Valid security classes are identified as a number in the range 1 - 255. The default value is 255. For more information about security class values, see *z/OS Communications Server: IP Configuration Guide*.

The TCP/IP stack ignores this value if IPSECURITY is not specified on the IPCONFIG statement.

MONSYSPLEX | NOMONSYSPLEX

Specifies whether sysplex autonomics should monitor the interface's status.

NOMONSYSPLEX	Specifies that sysplex autonomics should not monitor the interface's status. This is the default value.
MONSYSPLEX	Specifies that sysplex autonomics should monitor the interface's status. **Restriction:** The MONSYSPLEX attribute is not in effect unless the MONINTERFACE keyword is specified on the GLOBALCONFIG SYSPLEXMONITOR profile statement. The presence of dynamic routes over this interface is monitored if the DYNROUTE keyword is also specified on the GLOBALCONFIG SYSPLEXMONITOR profile statement.

DYNVLANREG | NODYNVLANREG

This parameter controls whether or not the VLAN ID for this interface is dynamically or statically registered with the physical switch on the LAN.

Restriction: This parameter is applicable only if a VLAN ID is specified on the statement.

Dynamic registration of VLAN IDs is handled by the OSA-Express feature and the physical switch on your LAN. Therefore, in order for the DYNVLANREG parameter to be effective, both must be at a level that provides the necessary hardware support for dynamic VLAN ID registration. After the interface is active, you can view the Netstat DEvlinks/`-d` report output to determine whether your OSA-Express feature can support VLAN dynamic registration. This Netstat report also displays whether dynamic VLAN ID registration has been configured for the interface.

NODYNVLANREG	Specifies that if a VLAN ID is configured for this interface, it must be manually registered with the physical switches on the corresponding LAN. This is the default value. If this parameter is specified without a VLAN ID, then it is ignored.
DYNVLANREG	Specifies that if a VLAN ID is configured for this interface, it is dynamically registered with the physical switches on the corresponding LAN. If this parameter is specified without a VLAN ID, then warning message EZZ0056I is issued and the NODYNVLANREG setting is used instead.

OLM | NOOLM

An optional parameter indicating whether an OSA-Express adapter operates in optimized latency mode.

NOOLM	Specifies that the OSA-Express adapter should not operate in optimized latency mode. This is the default value.
OLM	Specifies that the OSA-Express adapter operates in optimized latency mode (OLM). Optimized latency mode optimizes interrupt processing for both inbound and outbound data. Use this mode for workloads that have demanding latency requirements. Because this mode can provide significant increases of throughput, particularly for interactive, non-streaming workloads. For more information about optimized latency mode, see the optimized latency mode topic in *z/OS Communications Server: IP Configuration Guide*.

Guidelines:

- Because of the operating characteristics of optimized latency mode, you might need to change your configuration to direct traffic to particular OSA-Express write priority queues and to limit the number of concurrent users sharing an OSA-Express configured for optimized latency mode. For more information about OLM, see the optimized latency mode topic in *z/OS Communications Server: IP Configuration Guide*.
- The optimized latency mode function targets a z/OS environment with a high-volume, interactive workloads. Although optimized latency mode can compensate for some mixing of workloads, an excessive amount of high-volume streaming workloads, such as bulk data or file transfer, can result in higher CPU consumption.

Restrictions:

- This function is limited to OSA-Express3 or later Ethernet features in QDIO mode that are running with an IBM System z10 or later. See the 2097 DEVICE Preventive Service Planning (PSP) bucket for more information.
- Traffic that is either inbound over or being forwarded to an OSA-Express configured to use optimized latency mode is not eligible for the accelerated routing function provided by HiperSockets Accelerator and QDIO Accelerator.
- For an OSA-Express configured to use optimized latency mode, the stack ignores the configured or default INBPERF setting and uses the value DYNAMIC.

ISOLATE | NOISOLATE

Specifies whether packets should be directly routed between TCP/IP stacks that share the OSA adapter.

NOISOLATE

Route packets directly between TCP/IP stacks sharing the OSA adapter. In this mode, if the next hop address was registered by another stack that is sharing the OSA adapter, then the OSA-Express adapter routes the packet directly to the sharing stack without putting the packet on the external LAN.

ISOLATE

Prevent OSA-Express from routing packets directly to another TCP/IP stack that is sharing the OSA adapter. In this mode, OSA-Express adapter discards any packets when the next hop address was registered by another stack that is sharing the OSA adapter. Packets can flow between two stacks that share the OSA only by first going through a router on the LAN. For more details, see the OSA-Express connection isolation information in *z/OS Communications Server: IP Configuration Guide*.

Tips:

- If you isolate an interface, there might be an adverse effect on latency.
- You can selectively apply OSA-Express connection isolation to individual virtual LANs.
- The OSA-Express adapter requires that both stacks sharing the port be non-isolated for direct routing to occur. Therefore, for traffic between two stacks sharing the OSA adapter, as long as at least one of the stacks is isolated, connection isolation is in effect for traffic in both directions between these stacks.

Restriction: This function is limited to OSA-Express2 or later Ethernet features in QDIO mode and running at least an IBM System z9 Enterprise Class (EC) or z9 Business Class (BC). See the 2094DEVICE, 2096DEVICE, 2097DEVICE, or 2098DEVICE Preventive Service Planning (PSP) bucket for more information.

Steps for modifying

See Summary of INTERFACE statements for modification information.

Examples

```
INTERFACE OSAQDIO24
DEFINE IPAQENET
PORTNAME OSAQDIO2
SOURCEVIPAINT VIPAV4
IPADDR 100.1.1.1/24
```

Related topics

- BEGINROUTES statement
- BSDROUTINGPARMS statement
- DEVICE and LINK - MPCIPA OSA-Express QDIO devices statement
- GLOBALCONFIG statement
- INTERFACE - IPAQENET6 OSA-Express QDIO interfaces statement on page 28
- SACONFIG statement
- START statement
- STOP statement

INTERFACE - IPAQENET6 OSA-Express QDIO interfaces statement

Use the INTERFACE statement to specify an OSA-Express QDIO Ethernet interface for IPv6.

To determine the OSA-Express microcode level, use the DISPLAY TRL command. If a specific OSA-Express function is documented with a minimum microcode level, you can use this command to determine whether that function is supported. IBM service might request the microcode level for problem diagnosis. For more information about the DISPLAY TRL command, see *z/OS Communications Server: SNA Operation*.

For more information about OSA-Express features that support QDIO mode, see *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

When you start an IPAQENET6 interface (and you do not specify VMAC with ROUTEALL), TCP/IP registers all non-loopback local (home) IPv6 addresses for this TCP/IP instance to the OSA-Express feature. If you subsequently add, delete, or change any home IPv6 addresses on this TCP/IP instance, TCP/IP dynamically registers the changes to the OSA-Express feature. If stateless address autoconfiguration is enabled for this interface, TCP/IP dynamically registers autoconfigured addresses to the OSA-Express feature. This includes both public and temporary autoconfigured addresses. The OSA-Express feature routes datagrams destined to those IPv6 addresses to this TCP/IP instance.

If a datagram is received by the OSA adapter for an unregistered IPv6 address, then the OSA-Express feature routes the datagram to the TCP/IP instance, depending on the setting of a virtual MAC (VMAC) address or whether the definition of an instance is PRIROUTER or SECROUTER. If the datagram is not destined for a virtual MAC address and no active TCP/IP instance using this interface is defined as PRIROUTER or SECROUTER, then the OSA-Express feature discards the datagram. For more details about the OSA-Express feature routing considerations, see the router information in *z/OS Communications Server: IP Configuration Guide* and primary and secondary routing in *z/OS Communications Server: SNA Network Implementation Guide*.

For detailed instructions on setting up an OSA-Express feature, see *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

For more information about missing interrupt handler (MIH) considerations with TCP/IP interfaces, see Missing interrupt handler factors.

Restriction: This statement applies to IPv6 IP addresses only.

Syntax

Rule: Specify the required parameters and the CHPIDTYPE parameter in the order shown here. The OSD Interface Definition and OSX Interface Definition parameters can be specified in any order.

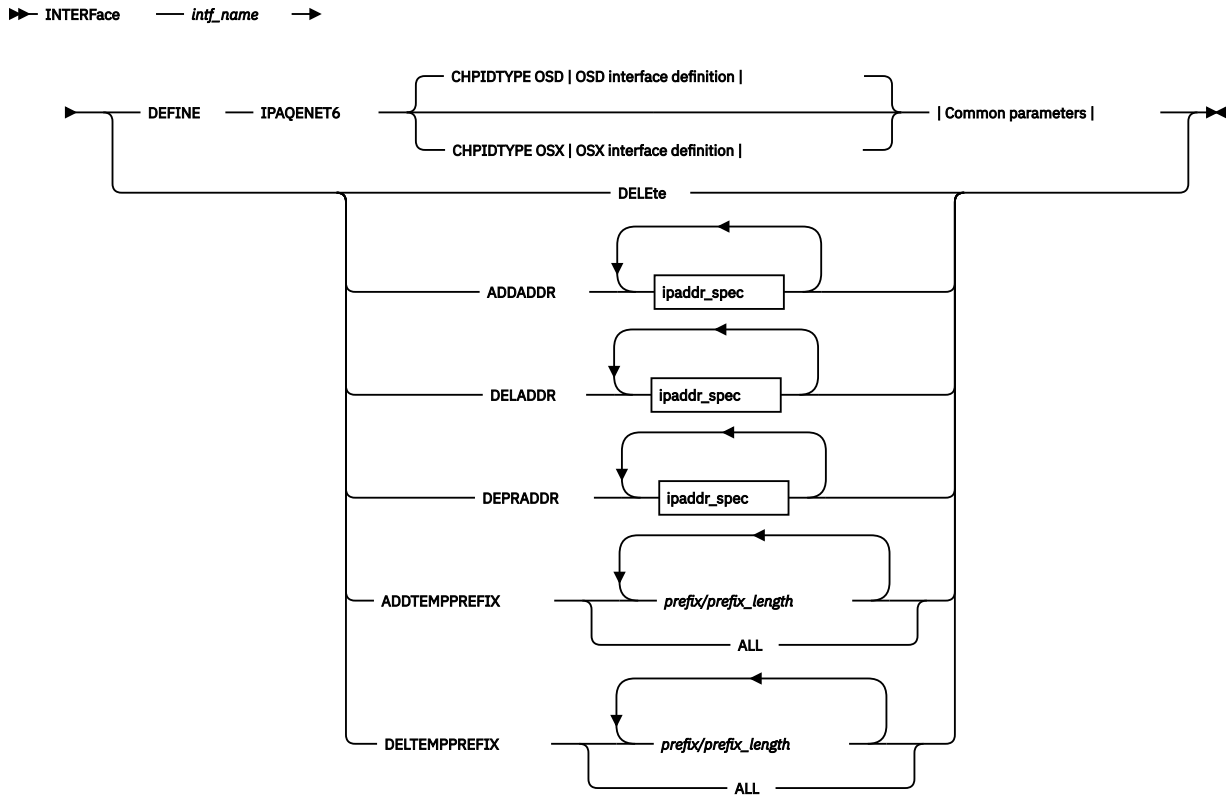

Figure 2:

OSX Interface definition

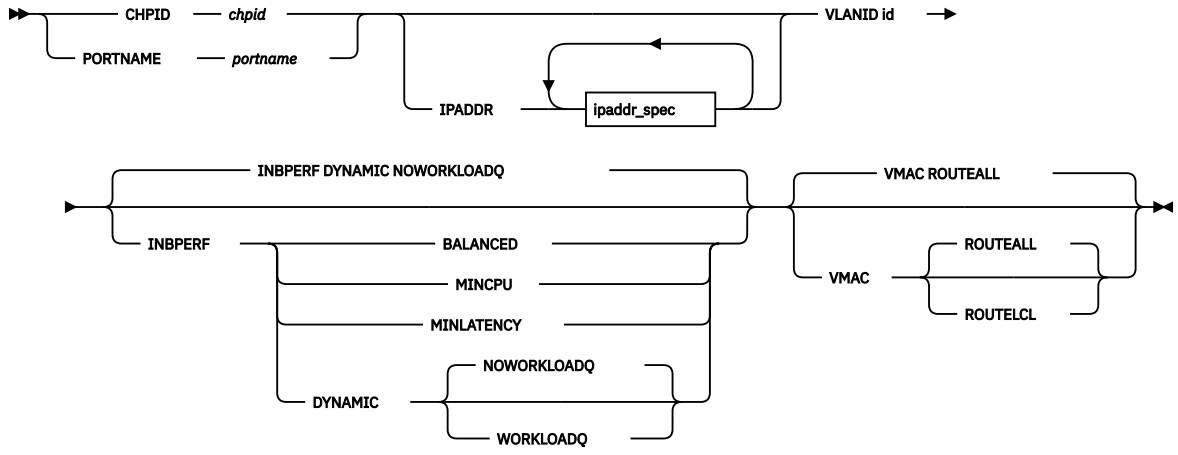

Common parameters for OSD and OSX interface definitions

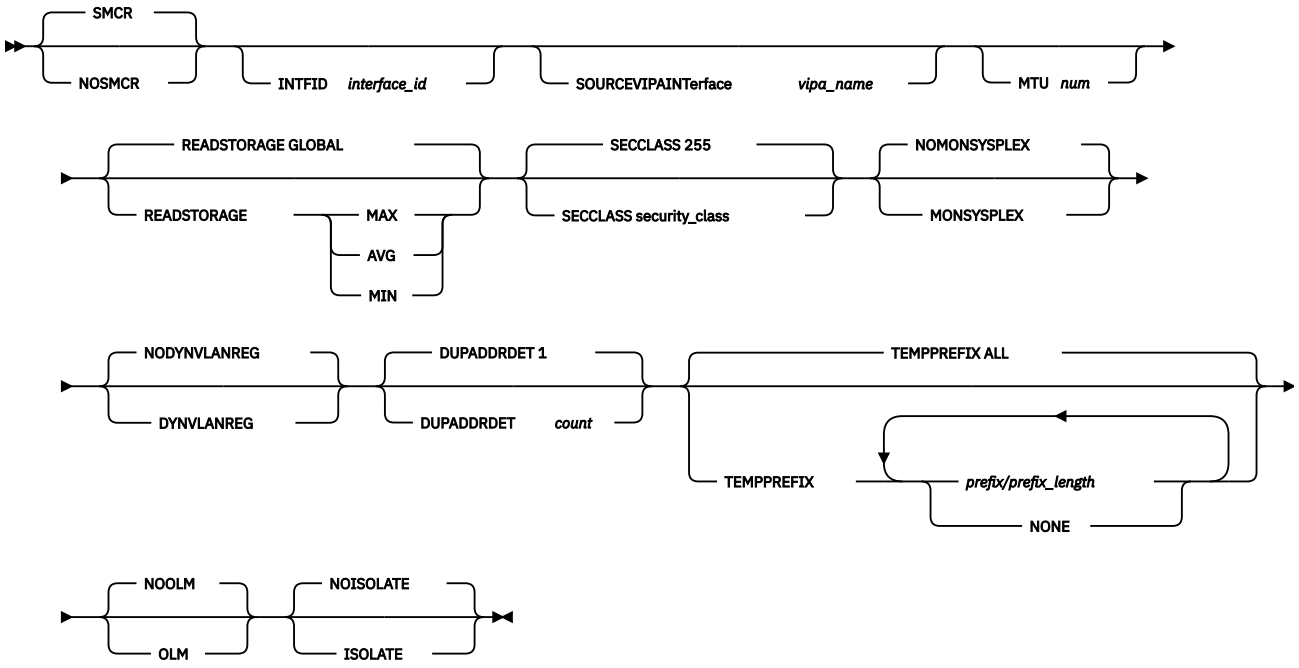

ipaddr_spec

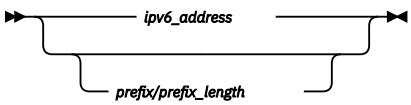

Parameters

intf_name

The name of the interface. The maximum length is 16 characters.

Requirement: This name must be different than the name specified for the PORTNAME parameter.

Restriction: Do not specify the value PUBLICADDRS or TEMPADDRS for the interface name. The values PUBLICADDRS and TEMPADDRS are keywords on the SRCIP statement. These values are not recognized if they are specified as an IPv6 interface name on an SRCIP entry.

DEFINE

Specifies that this definition is to be added to the list of defined interfaces.

DELETE

Specifies that this definition is to be deleted from the list of defined interfaces. The *intf_name* must be the name of an interface previously defined by an INTERFACE statement. INTERFACE DELETE deletes all home IP addresses for the interface.

CHPIDTYPE

An optional parameter indicating the CHPID type of the OSA-Express QDIO interface.

OSD	The external data network. This is the default value.
OSX	The intraensemble data network. See *z/OS Communications Server: IP Configuration Guide* for information about requirements necessary to make an OSX work.

IPADDR ***ipaddr_spec***

The IPADDR parameter is optional, and is used to configure the interface's IPv6 addresses other than the link-local address, which is generated internally by TCP/IP. If IPADDR is not specified, TCP/IP enables IPv6 stateless autoconfiguration for the interface. For more information, see Stateless address autoconfiguration in *z/OS Communications Server: IPv6 Network and Appl Design Guide*.

Tip: Autoconfiguration is enabled if a router or some other device provides a router advertisement.

The following values can be specified for *ipaddr_spec*:

- ipv6_addr (A fully qualified IPv6 address is in colon-hexadecimal format.)
- *prefix/prefix_length*

  The digits (in colon-hexadecimal format) before the / represent the prefix. The *prefix_length* represents the length of the prefix in bits. If a prefix length is coded, it must be equal to 64. When a prefix is specified, TCP/IP forms the IPv6 address by appending the interface ID to the specified prefix. The interface ID is either the value specified by way of the INTFID keyword, or the value formed by the stack using information returned by the device when the interface was started.

  Restriction: If you code a prefix that is longer than 64 bits, it is truncated to 64 bits, and no error messages are issued.

For information about the IPv6 address restrictions, see Restrictions on IPv6 addresses configured in the TCP/IP profile

ADDADDR *ipaddr_spec*

Allows the addition of IP addresses to an existing INTERFACE definition (similar to updating the HOME list with the VARY TCPIP,,OBEYFILE command) without having to delete and redefine the INTERFACE. This can be used to change the autoconfiguration state of an interface. If ADDADDR is coded and this is the first manually configured IP address for the interface, then TCP/IP disables autoconfiguration for the interface. The *intf_name* coded with ADDADDR must be the name of an interface previously defined by an INTERFACE statement.

Any public or temporary addresses that had previously been autoconfigured for the interface are deleted.

DELADDR *ipaddr_spec*

Allows you to delete IP addresses from an existing INTERFACE definition. If DELADDR is coded for the last or only manually configured IP address for an interface, then TCP/IP enables autoconfiguration for the interface. DELADDR is valid only for an IP address or prefix configured manually. The *intf_name* coded with DELADDR must be the name of an interface previously defined by an INTERFACE statement. DELADDR is valid only in a data set specified on a VARY TCPIP,,OBEYFILE command.

Guideline: If you specify a prefix for DELADDR, then the only IP addresses affected are those defined by way of the same prefix specified on IPADDR or ADDADDR.

DEPRADDR *ipaddr_spec*

The DEPRADDR keyword allows you to deprecate an IP address. This can assist with site renumbering. DEPRADDR is valid only for an IP address or prefix configured manually. If you use DEPRADDR to deprecate an IP address, you can subsequently use ADDADDR again to make that IP address preferred. For DEPRADDR, the *interface_name* must be the name of an interface previously defined by an INTERFACE statement. DEPRADDR is valid only in a data set specified on a VARY TCPIP,,OBEYFILE command.

Guideline: If you specify a prefix for DEPRADDR, then the only IP addresses affected are those defined by way of the same prefix specified on IPADDR or ADDADDR.

ADDTEMPPREFIX

Use the ADDTEMPPREFIX keyword to add prefixes to the temporary prefixes list of an existing INTERFACE definition without having to delete and redefine the INTERFACE statement. The temporary prefixes list limits the set of prefixes for which temporary IPv6 addresses can be generated. A temporary IPv6 address is generated when a router advertisement containing the prefix is processed, and the prefix is included in one of the prefixes in the temporary prefixes list. For example, if the temporary prefixes list for an interface contains a single prefix 2001:0db8:58cd::/48, a temporary address is generated for advertised prefix 2001:0db8:58cd:0001/64; however, a temporary address is not generated on this interface for advertised prefix 2001:0db8:5555:0001/64. The *intf_name* variable coded with ADDTEMPPREFIX must be the name of an interface that was previously defined by an INTERFACE statement.

prefix/prefix_length
: The digits (in colon-hexadecimal format) before the slash (/) represent the prefix. The *prefix_length* value represents the length of the prefix in bits. Valid values for *prefix_length* parameter are in the range 1 - 64.

ALL
: Causes temporary addresses to be generated for all prefixes that are learned over this interface by way of router advertisements.

DELTEMPPREFIX

Use the DELTEMPPREFIX keyword to delete prefixes from the temporary prefixes list of an existing INTERFACE definition. The temporary prefixes list limits the set of prefixes for which temporary IPv6 addresses can be generated. A temporary IPv6 address is generated when a router advertisement containing the prefix is processed and the prefix is included in one of the prefixes in the temporary prefixes list. The *intf_name* variable coded with the DELTEMPPREFIX keyword must be the name of an interface that was previously defined by an INTERFACE statement.

prefix/prefix_length
: The digits (in colon-hexadecimal format) before the slash (/) represent the prefix. The *prefix_length* value represents the length of the prefix in bits. Valid values for the *prefix_length* are in the range 1 - 64. All temporary addresses for this interface whose prefix is not included in the updated temporary prefixes list are deleted.

ALL
: Delete all prefixes from the temporary prefixes list, which sets the temporary prefixes list to NONE. All temporary addresses for this interface are deleted, and no more temporary addresses are generated for this interface.

IPAQENET6

Indicates that the interface uses the interface based on IP assist, belongs to the QDIO family of interfaces, and uses the Gigabit Ethernet or Fast Ethernet protocol.

INTFID *interface_id*

An optional 64-bit interface identifier in colon-hexadecimal format. IPv6 shorthand is not allowed when specifying the interface ID. If specified, this interface ID is used to form the link-local address for the interface, and is also appended to any manually configured prefixes for the interface, to form complete IPv6 addresses on the interface. If you do not configure manual IP addresses on the interface, the INTFID value is appended to any prefixes that are learned over this interface by way of router advertisements to form public IPv6 addresses on the interface. The INTFID value is not used to form temporary IPv6 addresses. A randomly generated interface ID is appended to any learned prefixes to form temporary IPv6 addresses on the interface (if temporary addresses are enabled).

If INTFID is not coded, TCP/IP builds the Interface ID using information returned from the OSA-Express Adapter (during Interface activation). The built Interface ID value is then used to form the link-local address. This value is also used to complete the formation of other IPv6 addresses on the interface, if you choose to configure only the prefix portion of the addresses (by way of IPADDR or ADDADDR). Also, if you do not configure manual IP addresses on the interface, the built interface ID value is appended to any prefixes learned over this interface by way of router advertisements to form public IPv6 addresses on the interface. The built interface ID value is not used to form temporary IPv6 addresses. A randomly generated interface ID is appended to any learned prefixes to form temporary IPv6 addresses on the interface (if temporary addresses are enabled).

When defining the interface ID, the local/universal flag (the U bit, bit 6 shown in the following example) must be set to 0. The group/individual flag (the G bit, bit 7 shown in the following example) must also be set to 0. If either flag is set incorrectly, interface definition fails. Additionally, an interface ID value correlating to an ISATAP address or a Reserved Anycast address is not allowed. (An ISATAP Interface ID has '00005EFE'x in bits 0 - 31, and a Reserved Anycast Interface ID has 'FCFFFFFFFFFFFF8' in bits 0 - 56.)

```
|                1|1              3|3              4|4
  6|
|0               5|6              1|2              7|8
  3|
+----------------+----------------+----------------
+----------------+
|xxxxxxUGxxxxxxxx|xxxxxxxxxxxxxxxx|xxxxxxxxxxxxxxxx|
xxxxxxxxxxxxxxxx|
+----------------+----------------+----------------
+----------------+
```

SOURCEVIPAINTERFACE ***vipa_name***

SOURCEVIPAINTERFACE is optional. Use this parameter to specify which previously defined static VIPA interface is to be used for SOURCEVIPA (when IPCONFIG6 SOURCEVIPA is in effect).

Tip: The use of the SOURCEVIPAINTERFACE parameter can be overridden. See the information about source IP address selection in *z/OS Communications Server: IP Configuration Guide* for the hierarchy of ways that the source IP address of an outbound packet is determined.

The *vipa_name* is the interface name for a VIRTUAL6 interface. If the VIPA has multiple IP addresses, then the sourcevipa address for outbound packets is selected from among these addresses according to the default source address selection algorithm. For more information, see*z/OS Communications Server: IPv6 Network and Appl Design Guide*.

Requirement: The VIRTUAL6 interface must be defined prior to specifying this INTERFACE statement to the TCP/IP stack. It must either already be defined or, the INTERFACE statement that defines it must precede this INTERFACE statement in the profile data set.

CHPID ***chpid***

This parameter applies only to interfaces of CHPIDTYPE OSX and is used to specify the CHPID for the interface. This value is a 2-character hexadecimal value (00 - FF).

PORTNAME ***portname***

Use this parameter to specify the PORT name contained in the TRLE definition for the QDIO interface. The TRLE must be defined as MPCLEVEL=QDIO. For details about defining a TRLE, see *z/OS Communications Server: SNA Resource Definition Reference*.

Requirement: The *portname* value must be different from the name that is specified for the *intf_name* value.

NONROUTER

If a datagram is received at this interface for an unknown IP address, the datagram is not routed to this TCP/IP instance. This is the default value.

PRIRouter and SECRouter parameters interact with the VLANID parameter. See the VLANID parameter to understand this relationship.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

PRIROUTER

If a datagram is received at this interface for an unknown IP address and is not destined for a virtual MAC, the datagram is routed to this TCP/IP instance.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

SECROUTER

If a datagram is received at this interface for an unknown IP address and is not destined for a virtual MAC, and there is no active TCP/IP instance defined as PRIROUTER, then the datagram is routed to this TCP/IP instance.

Rule: This keyword applies only to interfaces of CHPIDTYPE OSD and is ignored if the VMAC parameter is configured on the INTERFACE statement.

DUPADDRDET ***count***

Use this parameter to specify the number of times to attempt duplicate address detection. The minimum value is 0, maximum is 2 and default is 1. This is an optional parameter.

Guideline: A value of 0 means that TCP/IP does not perform duplicate address detection for this interface.

MTU ***num***

The maximum transmission unit (MTU) in bytes. This value can be up to 9000. The minimum MTU for IPv6 is 1280. The stack takes the minimum of the configured value and the value supported by the device (returned by the OSA adapter).

The MTU default, which depends on value supported by device, is the following value:

- Gigabit Ethernet default MTU = 9000
- Fast Ethernet default MTU = 1500

Tip: See *z/OS Communications Server: IP Configuration Guide*, in section Maximum transmission unit considerations, for additional information about how TCP/IP uses the MTU to determine the largest size frame to send.

VLANID id

Specifies the decimal virtual LAN identifier to be assigned to the OSA-Express INTERFACE. This field should be a virtual LAN identifier recognized by the switch for the LAN connected to this OSA-Express. The valid range is 1 - 4094. This parameter is optional for CHPIDTYPE OSD and required for CHPIDTYPE OSX.

Guideline: Installation configuration on other platforms or related to Ensemble networking can limit the maximum VLANID of 4096.

The VLANID parameter interacts with the PRIRouter and SECRouter parameters. If you configure both the VLANID parameter and either PRIROUTER or SECROUTER parameter, then this TCP/IP instance acts as a router for this VLAN (ID) only. Datagrams that are received at this device instance for an unknown IP address and are not destined for a virtual MAC are routed only to this TCP/IP instance if it is VLAN tagged with this VLAN ID. For more information about VLANID parameter interactions, see OSA VLAN in *z/OS Communications Server: IP Configuration Guide*.

Rule: If you are configuring multiple VLAN interfaces to the same OSA-Express feature, then you must specify the VMAC parameter (with the default ROUTEALL attribute) on the INTERFACE statement for each of these interfaces.

Restriction: The stack supports a maximum of 32 IPv6 VLAN interfaces to the same OSA-Express port. Additional VLANID limitations might exist if this interface can be used with Shared Memory Communications over Remote Direct Memory Access (SMC-R). See VLANID considerations in *z/OS Communications Server: IP Configuration Guide* for details.

READSTORAGE

An optional parameter indicating the amount of fixed storage that z/OS Communications Server should keep available for read processing for each input queue for this interface. Multiple input queues are used when WORKLOADQ is enabled. See description of the WORKLOADQ parameter for more details.

The QDIOSTG VTAM start option allows you to specify a value which applies to all OSA-Express adapters in QDIO mode. You can use the READSTORAGE keyword to override the global QDIOSTG value for this adapter based on the inbound workload you expect over this interface on this stack. The valid values are:

GLOBAL	The amount of storage is determined by the QDIOSTG VTAM start option. This is the default value.
MAX	Use this value if you expect a heavy inbound workload over this interface.
AVG	Use this value if you expect a medium inbound workload over this interface.
MIN	Use this value if you expect a light inbound workload over this interface.

Tip: See the description of the QDIOSTG VTAM start option in the *z/OS Communications Server: SNA Resource Definition Reference* for details about exactly how much storage is allocated by z/OS Communications Server for each of these values.

Tip: The amount of storage which corresponds to each of these values is different for an OSA with at least 25 GbE of bandwidth than for an OSA with less than 25 GbE of bandwidth. See QDIOSTG start option in *z/OS Communications Server: SNA Resource Definition Reference* for details about exactly how much storage is allocated by z/OS Communications Server for each of these values.

Rule: If you define both a LINK and INTERFACE statement for the same adapter, then the READSTORAGE value on the LINK statement must match the READSTORAGE value on the corresponding INTERFACE statement. If you define an INTERFACE statement that contains a value for READSTORAGE that conflicts with the READSTORAGE value for a previous LINK statement for the same adapter, then TCP/IP rejects the INTERFACE statement.

INBPERF

An optional parameter that indicates how processing of inbound traffic for the QDIO interface is performed.

There are three supported static settings (MINCPU, MINLATENCY, and BALANCED).that indicate how frequently the adapter should interrupt the host for inbound traffic: BALANCED, MINCPU, and MINLATENCY. The static settings use static interrupt-timing values. The static values are not always optimal for all workload types or traffic patterns, and cannot account for changes in traffic patterns.

There is also one supported dynamic setting (DYNAMIC). This setting causes the host (stack) to dynamically adjust the timer-interrupt value while the device is active and in use. This function exploits an OSA hardware function called Dynamic LAN Idle. Unlike the static settings, the DYNAMIC setting reacts to changes in traffic patterns, and sets the interrupt-timing values at the

point where throughput is maximized. In addition, the DYNAMIC setting uses the OSA Dynamic Router Architecture function to enable QDIO inbound workload queues for specific inbound traffic types.

Result: When you specify OLM on the INTERFACE statement, the INBPERF parameter is ignored and the statement defaults to the value DYNAMIC.

Valid values are:

BALANCED

This setting uses a static interrupt-timing value, which is selected to achieve reasonably high throughput and reasonably low CPU consumption. This is the default value for CHPIDTYPE OSD.

DYNAMIC

This setting causes the host to dynamically signal the OSA-Express feature to change the timer-interrupt value, based on current inbound workload conditions. The DYNAMIC setting is effective only for OSA-Express2 or later features on at least an IBM System z9 that supports the corresponding Dynamic LAN Idle function. See the 2094DEVICE Preventive Service Planning (PSP) bucket and the 2096DEVICE Preventive Service Planning (PSP) bucket for more information about the level of OSA-Express2 adapter that supports this function. See the 2097DEVICE Preventive Service Planning (PSP) bucket for more information about the OSA-Express3 adapter that supports this function. The DYNAMIC setting can decrease latency and provide increases in throughput for many interactive workloads. For all other workload combinations, this setting provides performance similar to the three static settings. This is the default value for CHPIDTYPE OSX.

If the DYNAMIC setting is specified for an OSA-Express adapter that does not support the dynamic LAN Idle function, the stack reverts to using the BALANCED setting.

WORKLOADQ | NOWORKLOADQ

This subparameter controls the QDIO inbound workload queueing function for the interface. QDIO inbound workload queueing is effective only for OSA-Express features in QDIO mode that support the corresponding Data Router Architecture. OSA-Express features that support workload queueing do not necessarily support workload queueing for all possible traffic types. For more information about the QDIO inbound workload queueing function and the OSA-Express features that support it, see QDIO inbound workload queueing in *z/OS Communications Server: IP Configuration Guide*.

NOWORKLOADQ	Specifies that QDIO inbound workload queueing is not enabled for inbound traffic. All inbound traffic for this interface uses a single input queue. This is the default.
WORKLOADQ	Specifies that QDIO inbound workload queueing (IWQ) is enabled for inbound traffic. If the WORKLOADQ subparameter is specified, QDIO inbound workload queueing is enabled for specific inbound traffic types. A primary input queue is reserved for all other traffic types. Ancillary input queues (AIQs) are created for the following inbound traffic types when supported by the OSA-Express feature: • Sysplex distributor • Streaming workloads (for example FTP) • Enterprise Extender (EE) • IP Security (IPSec)

- IBM z/OS Container Extensions (zCX)

Requirement: You must specify the VMAC parameter with WORKLOADQ to enable QDIO inbound workload queueing.

Restrictions:

- Bulk-mode TCP connection registration is supported only in configurations in which a single inbound interface is servicing the bulk-mode TCP connection. If a bulk-mode TCP connection detects that it is receiving data over multiple interfaces, QDIO IWQ is disabled for the TCP connection and inbound data from that point forward is delivered to the primary input queue.
- QDIO IWQ does not apply for traffic that is sent over an OSA port that is shared by the receiving TCP/IP stack when an indirect route (where the next hop and destination IP address are different) is being used; this traffic is placed on the primary input queue. QDIO IWQ does apply when traffic on the shared OSA path uses a direct route (where the next hop and destination IP address are the same).

Guideline: The WORKLOADQ parameter requires an additional amount of fixed 4K CSM HVCOMMON storage per AIQ. The amount of storage consumed per AIQ is based on the amount of storage defined for READSTORAGE for this interface. The bulk AIQ is always backed with this additional CSM storage. The remaining AIQs are not backed by the additional CSM storage until the specific function (EE, SD, IPSec, or zCX) is used. The EE AIQ is backed by fixed 4K CSM DSPACE64 storage (instead of HVCOMMON). To verify the amount of CSM storage that is being used for each input queue, display the VTAM TRLE name that is associated with the interface. The WORKLOADQ parameter also requires an additional 36K of ECSA per AIQ.

Tip: The additional CSM storage consumed by each OSA interface using WORKLOADQ consumes fixed (real) storage. It is recommended that you verify that the additional fixed storage required by enabling WORKLOADQ (per OSA interface) will not approach any of the following limits:

- The CSM FIXED MAXIMUM value used by Communications Server (use the D NET, CSM command and see the CSM FIXED MAXIMUM value defined in IVTPRM00)

- The actual real storage available to this z/OS system (see D M=STOR or D M=HIGH)

If the WORKLOADQ setting is specified for an OSA-Express adapter that does not support the Data Router Architecture function, the stack reverts to using a single input queue.

MINCPU

This setting uses a static interrupt-timing value, which is selected to minimize host interrupts without regard to throughput. This mode of operation might result in minor queueing delays (latency) for packets into the host, which is not optimal for workloads with demanding latency requirements.

MINLATENCY

This setting uses a static interrupt-timing value, which is selected to minimize latency (delay), by more aggressively presenting received packets to the host. This mode of operation generally results in higher CPU consumption than the other three settings. Use this setting only if host CPU consumption is not an issue.

Rule: If you define both a LINK IPAQENET and an INTERFACE IPAQENET6 statement for the same adapter, then the following rules apply for the INBPERF parameter on these statements:

- The value on the LINK statement must match the INBPERF value on the corresponding INTERFACE statement.
- The INTERFACE statement supports the subparameters WORKLOADQ and NOWORKLOADQ for the INBPERF DYNAMIC parameter. These subparameters are associated with QDIO inbound workload queueing support and are not supported on the LINK IPAQENET statement. So, if you specify the INBPERF DYNAMIC parameter for both the LINK and the INTERFACE statements, then you must use the default or specify the NOWORKLOADQ subparameter for the INBPERF DYNAMIC parameter on the INTERFACE statement. This ensures that the INBPERF DYNAMIC setting for both statements is the same.
- If you define an INTERFACE IPAQENET6 statement that contains a value for INBPERF that conflicts with the INBPERF value for a previous LINK IPAQENET statement for the same adapter, then TCP/IP rejects the INTERFACE statement.

SECCLASS *security_class*

Use this parameter to associate a security class for IP filtering with this interface. In order for traffic over the interface to match a filter rule, the filter rule must have the same security class value as the interface or a value of 0. Filter rules can be specified in the TCP/IP profile or in an IP Security policy file read by the Policy Agent. Filter rules can include a security class specification on the IpService statement in an IP Security policy file or on the SECCLASS parameter on the IPSEC6RULE statement in the TCP/IP profile.

Valid security classes are identified as a number in the range 1 - 255. The default value is 255. For more information about security class values, see *z/OS Communications Server: IP Configuration Guide*.

The TCP/IP stack ignores this value if IPSECURITY is not specified on the IPCONFIG6 statement.

MONSYSPLEX | NOMONSYSPLEX

Specifies whether or not sysplex autonomics should monitor the interface's status.

NOMONSYSPLEX	Specifies that sysplex autonomics should not monitor the interfaces's status. This is the default value.
MONSYSPLEX	Specifies that sysplex autonomics should monitor the interface's status.

Restriction: The MONSYSPLEX attribute is not in effect unless the MONINTERFACE keyword is specified on the GLOBALCONFIG SYSPLEXMONITOR profile statement. The presence of dynamic routes over this interface is monitored if the DYNROUTE keyword is also specified on the GLOBALCONFIG SYSPLEXMONITOR profile statement.

DYNVLANREG | NODYNVLANREG

This parameter controls whether or not the VLAN ID for this interface is dynamically or statically registered with the physical switch on the LAN.

Restriction: This parameter is applicable only if a VLAN ID is specified on the statement.

Dynamic registration of VLAN IDs is handled by the OSA-Express feature and the physical switch on your LAN. Therefore, in order for the DYNVLANREG parameter to be effective, both must be at a level which provides the necessary hardware support for dynamic VLAN ID registration. After the interface is active, you can view the Netstat DEvlinks/-d report output to determine if your OSA-Express feature can support VLAN dynamic registration. This Netstat report also displays whether or not dynamic VLAN ID registration has been configured for the interface.

Rule: If you define both a LINK and INTERFACE statement for the same adapter, then the dynamic VLAN ID registration parameter value on the LINK statement must match the value of this same parameter on the corresponding INTERFACE statement. If you define an INTERFACE statement that contains a dynamic VLAN ID registration parameter value that conflicts with the same parameter value for a previous INTERFACE statement for the same OSA-Express feature, then TCP/IP rejects the INTERFACE statement.

NODYNVLANREG
Specifies that if a VLAN ID is configured for this interface, it must be manually registered with the physical switches on the corresponding LAN. This is the default value. If this parameter is specified without a VLAN ID, then it is ignored.

DYNVLANREG
Specifies that if a VLAN ID is configured for this interface, it is dynamically registered with the physical switches on the corresponding LAN. If this parameter is specified without a VLAN ID, then warning message EZZ0056I is issued and the NODYNVLANREG setting is used instead.

VMAC *macaddr*

Specifies the virtual MAC address, which can be represented by 12 hexadecimal characters. The OSA-Express device uses this address rather than the physical MAC address of the device for all IPv6 packets to and from this TCP/IP stack. For CHPIDTYPE OSD, using a virtual MAC address is optional. For CHPIDTYPE OSX, using a virtual MAC address is required, so the VMAC parameter is the default.

The *macaddr* value is optional for CHPIDTYPE OSD and cannot be specified for CHPIDTYPE OSX. If the *macaddr* value is not coded, then the OSA-Express device generates a virtual MAC address. If the *macaddr* is coded, it must be defined as a locally administered individual MAC address. This means the MAC address must have bit 6 (the universal or local flag U bit) of the first byte set to 1 and bit 7 (the group or individual flag G bit) of the first byte set to 0. The second

hexadecimal character must be 2, 6, A or E. The bit positions within the 12 hexadecimal characters are indicated as follows:

```
|               1|1              3|3               4|
|0              5|6              1|2               7|
+----------------+----------------+----------------+
|xxxxxxUGxxxxxxxx|xxxxxxxxxxxxxxxx|xxxxxxxxxxxxxxxx|
+----------------+----------------+----------------+
```

Rules:

- The same virtual MAC address generated by the OSA-Express device at interface activation remains in effect for this OSA-Express for this TCP/IP stack, even if the interface is stopped or becomes inoperative (INOPs). A new Virtual MAC address is generated only if the INTERFACE statement is deleted and redefined, or if the TCP/IP stack is recycled.
- The NONROUTER, PRIROUTER, and SECROUTER parameters are ignored for an OSA-Express interface if the VMAC parameter is configured on the INTERFACE statement.

Guideline: Unless the virtual MAC address representing this OSA-Express device must remain the same even after TCP/IP termination and restart, configure VMAC without a *macaddr* value and allow the OSA-Express device to generate it. This guarantees that the VMAC address is unique from all other physical burned-in MAC addresses and from all other VMAC addresses generated by any OSA-Express feature.

ROUTEALL

Specifies that all IP traffic destined to the virtual MAC is forwarded by the OSA-Express device to the TCP/IP stack. This is the default value. See the router information in *z/OS Communications Server: IP Configuration Guide* for more details.

ROUTELCL

This specifies that only traffic destined to the virtual MAC and whose destination IP address is registered with the OSA-Express device by this TCP/IP stack is forwarded by the OSA-Express. See the router information in *z/OS Communications Server: IP Configuration Guide* for more details.

SMCR | NOSMCR

Specifies whether this interface can be used with Shared Memory Communications over Remote Direct Memory Access (SMC-R) for IEDN or external data network communications.

NOSMCR	Specifies that this interface cannot be used for new TCP connections with SMC-R for IEDN or external data network communications.
SMCR	Specifies that this interface can be used for new TCP connections with SMC-R for IEDN or external data network communications. This is the default setting.

Rule: SMCR has no effect unless at least one Peripheral Component Interconnect Express (PCIe) function ID (PFID) value is specified by using the PFID subparameter of the SMCR parameter on the GLOBALCONFIG statement.

Guideline: If you enable Multipath and equal-cost interfaces are associated with different IP subnets, enabling SMC for some of, but not all, the interfaces can cause unpredictable SMC usage. You must specify either SMCR or NOSMCR on all equal-cost interfaces.

SMCD | NOSMCD

Specifies whether this interface can be used with Shared Memory Communications - Direct Memory Access (SMC-D).

NOSMCD	Specifies that this interface cannot be used for new TCP connections with SMC-D.
SMCD	Specifies that this interface can be used for new TCP connections with SMC-D. This is the default setting.

Rule: SMCD and NOSMCD are valid with CHPIDTYPE OSD definitions only.

Guideline: If you enable Multipath and equal-cost interfaces are associated with different IP subnets, enabling SMC for some of, but not all, the interfaces can cause unpredictable SMC usage. You must specify either SMCD or NOSMCD on all equal-cost interfaces.

OLM| NOOLM

An optional parameter indicating whether an OSA-Express adapter operates in optimized latency mode.

OLM	Specifies that the OSA-Express adapter operates in optimized latency mode (OLM). Optimized latency mode optimizes interrupt processing for both inbound and outbound data. Use this mode for workloads that have demanding latency requirements. Because this mode can provide significant increases of throughput, this mode is particularly suited for interactive, non-streaming workloads. For more information about OLM, see the optimized latency mode topic in *z/OS Communications Server: IP Configuration Guide*.
NOOLM	Specifies that the OSA-Express adapter should not operate in optimized latency mode. This is the default value.

Guidelines:

- Because of the operating characteristics of optimized latency mode, you might need to change configuration to direct traffic to particular OSA-Express write priority queues and to limit the number of concurrent users sharing an OSA-Express adapter configured for OLM. See the optimized latency mode topic in *z/OS Communications Server: IP Configuration Guide* for more information.
- The optimized latency mode function targets a z/OS environment with high-volume interactive workloads. Although optimized latency mode can compensate for some mixing of workloads, an excessive amount of high-volume streaming workloads, such as bulk data or file transfer, can result in higher CPU consumption.

Restrictions:

- This function is limited to OSA-Express3 or later Ethernet features in QDIO mode that are running with an IBM System z10 or later. See the 2097 DEVICE Preventive Service Planning (PSP) bucket for more information.
- For an OSA-Express configured to use optimized latency mode, the stack ignores the configured or default INBPERF setting and uses the value DYNAMIC.

NOISOLATE | ISOLATE

Specifies whether packets should be directly routed between TCP/IP stacks that share the OSA adapter.

NOISOLATE	Route packets directly between TCP/IP stacks that share the OSA adapter. In this mode, if the

next hop address was registered by another stack that is sharing the OSA, then OSA-Express routes the packet directly to the sharing stack without putting the packet on the external LAN.

ISOLATE

Prevent OSA-Express from routing packets directly to another TCP/IP stack that is sharing the OSA adapter. In this mode, OSA-Express discards any packets when the next hop address was registered by another stack that is sharing the OSA adapter. In this mode, packets can flow between two stacks that share the OSA adapter only by first going through a router on the LAN. For more details, see OSA-Express connection isolation information in *z/OS Communications Server: IP Configuration Guide*.

Tips:

- If you isolate an INTERFACE, that action might have an adverse effect on latency.
- You can selectively apply OSA-Express connection isolation to individual virtual LANs.
- OSA-Express requires that both stacks sharing the port be non-isolated for direct routing to occur. Therefore, for traffic between two stacks sharing the OSA adapter, as long as at least one of the stacks is isolated, connection isolation is in effect for traffic in both directions between these stacks.

Restriction: This function is limited to OSA-Express2 or later Ethernet features in QDIO mode and running at least an IBM System z9 Enterprise Class (EC) or z9 Business Class (BC). See the 2094, 2096, 2097, or 2098 DEVICE Preventive Service Planning (PSP) and the 2096DEVICE Preventive Service Planning (PSP) buckets for more information.

TEMPPREFIX

TEMPPREFIX specifies the set of prefixes for which temporary IPv6 addresses can be generated. A temporary IPv6 address is generated when a router advertisement containing a prefix is processed and the prefix is included in one of the prefixes in the temporary prefix list. For example, if TEMPPREFIX 2001:0db8:58cd::/48 is specified for an interface, a temporary address is generated for advertised prefix 2001:0db8:58cd:0001/64; however, a temporary address is not generated for advertised prefix 2001:0db8:5555:0001/64.

ALL

Generate temporary addresses for all prefixes that are learned over this interface by way of router advertisements. ALL is the default.

NONE

No IPv6 temporary addresses are generated for this interface.

prefix/prefix_length

The digits (in colon-hexadecimal format) before the slash (/) represent the prefix. The

prefix_length value represents the length of the prefix, in bits. Valid values for *prefix_length* are in the range 1 - 64.

Rules:

- Temporary addresses are generated only on an interface that is enabled for stateless address autoconfiguration.
- Temporary addresses are generated only when the TEMPADDRS keyword is specified on the IPCONFIG6 statement.

Requirement: You must specify the job name of an application in the SRCIP statement block with a value of TEMPADDRS to cause a temporary IPv6 address to be preferred over a public IPv6 address as the source IP address for the application; otherwise, the default source address selection algorithm prefers public IPv6 addresses over temporary addresses. For more information, see the information about the default source address selection algorithm in *z/OS Communications Server: IPv6 Network and Appl Design Guide*.

Steps for modifying

See Summary of INTERFACE statements for modification information.

Examples

```
INTERFACE OSAQDIO26 ; OSA QDIO (Fast Ethernet)
DEFINE IPAQENET6
PORTNAME OSAQDIO2
SOURCEVIPAINT VIPAV6
IPADDR 2001:0DB8:1:9:67:115:66     ; (Global Address)
```

Usage notes

Restriction: For each interface, the PRIROUTER and SECROUTER attributes can be in effect for only one TCP/IP instance within a central processor complex (CPC). If PRIROUTER is specified for an IPAQENET6 interface, but the IPv6 primary router attribute is already in effect on another TCP/IP instance for the same OSA-Express, then TCP/IP issues a warning message during interface activation and ignores the PRIROUTER parameter. Therefore, only one TCP/IP instance can be the primary router for the OSA-Express. Depending on the level of OSA-Express being started, either only one or multiple TCP/IP instances can be allowed to have SECROUTER specified. If OSA-Express allows only one secondary router, any TCP/IP instance subsequently starting that interface with SECROUTER receives a warning message during START processing for the interface. If OSA-Express allows multiple secondary routers, then OSA-Express can select any TCP/IP instance which specifies SECROUTER as the secondary router. There is no requirement that the same TCP/IP instance be specified PRIROUTER or SECROUTER for all OSA-Express adapters attached to the CPC.

Rule: To configure a single OSA port for both IPv4 and IPv6 traffic, consider the following conditions:

- If you use DEVICE/LINK/HOME for the IPv4 definition and INTERFACE for the IPv6 definition, the PORTNAME value on the INTERFACE statement must match the device_name on the DEVICE statement. This combination shares a single DATAPATH device.
- If you use INTERFACE for both IPv4 and IPv6 definitions, the PORTNAME value on the IPv4 INTERFACE statement must match the PORTNAME value on the IPv6 INTERFACE statement. This combination results in separate DATAPATH devices.

Related topics

- BEGINROUTES statement
- DEVICE and LINK - MPCIPA OSA-Express QDIO devices statement
- GLOBALCONFIG statement
- INTERFACE - IPAQENET OSA-Express QDIO interfaces statement on page 16

- START statement
- STOP statement

Chapter 4

IP System Administrator's Commands

Topics:

- DISPLAY TCPIP,,OSAINFO

DISPLAY TCPIP,,OSAINFO

Use the DISPLAY TCPIP,,OSAINFO command to retrieve information for active IPAQENET and IPAQENET6 interfaces. An interface represents a single datapath device of an OSA-Express feature. The information is retrieved directly from the OSA-Express feature.

Tips:

- If you have an INTERFACE and a DEVICE or LINK definition with the same port name and both are active, specifying either the INTERFACE or link name on the command will generate a report with both IPv4 and IPv6 information.
- You can use the DISPLAY TCPIP,,OSAINFO command to retrieve information for an active OSAENTA interface although only the base portion of the report is pertinent.

The command output provides the following sections of information:

Base	Contains physical characteristics and attributes for the interface and OSA-Express feature.
Registered addresses	Contains the Layer 2 MAC addresses or Layer 3 unicast and multicast IPv4 and IPv6 addresses registered to the OSA-Express feature.
QDIO inbound workload queueing routing variables	If QDIO inbound workload queueing is in effect for the interface, this section contains the routing variables for the ancillary input queues. Routing variables identify which inbound packets are to be presented on an ancillary input queue. For more information about ancillary input queues, see QDIO inbound workload queueing in *z/OS Communications Server: IP Configuration Guide*.

Both the modifiers and the MAX parameters can be used to limit the number of output lines.

Format

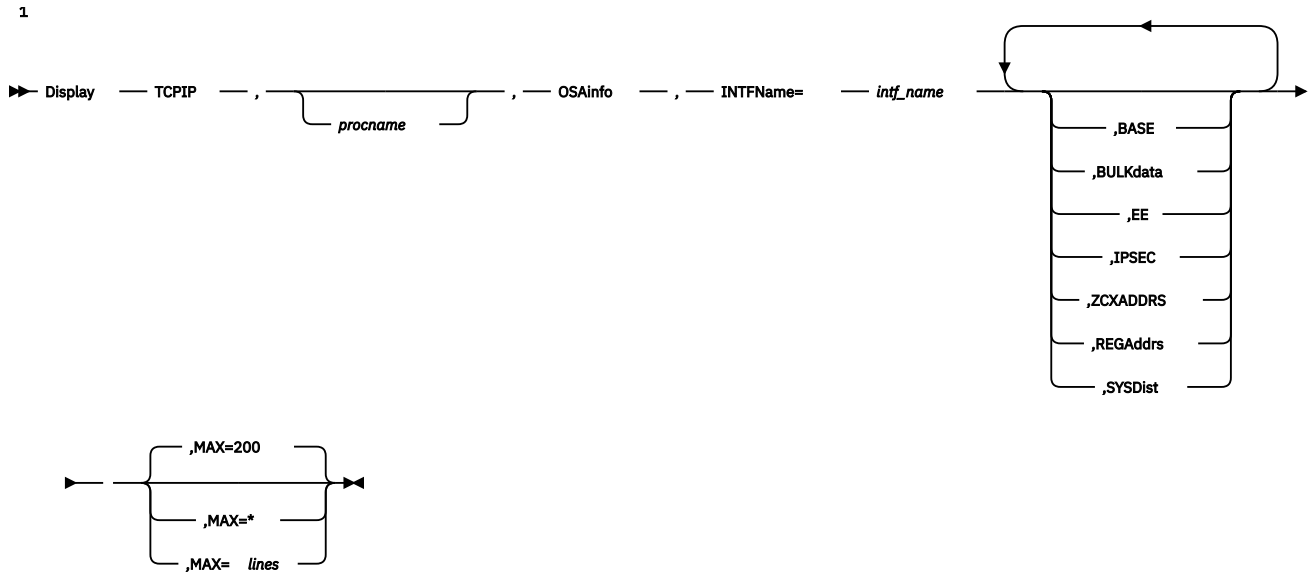

Notes:

[1] If no modifiers are specified, all sections for which information exists are displayed.

Rule: You must specify the parameters in the order that the syntax diagram shows.

Parameters

OSAINFO

Requests OSA information.

INTFNAME = *intf_name*

Specifies the name of the OSA-Express QDIO interface whose datapath device information is requested. The *intf_name* value can be one of the following names:

- The name that was configured on a LINK IPAQENET profile statement.
- The name that was configured on an INTERFACE IPAQENET or IPAQENET6 profile statement.
- The name of an OSAENTA trace interface, which is EZANTA*portname*, where the *portname* value is the name that is specified on the PORTNAME keyword in the TRLE for the OSA-Express port that is being traced.

Tip: To obtain a list of names to use as the value of the INTFNAME parameter, use the Netstat DEvlinks/-d command.

BASE

Indicates that the physical characteristics and attributes of the interface and OSA-Express feature are to be included in the report.

BULKDATA

Indicates that QDIO inbound workload queueing routing variables for the BULKDATA ancillary queue are to be included in the report. The BULKDATA routing variables are comprised of source and destination IP addresses, source and destination ports, and protocol. That combination uniquely identifies those packets that the OSA-Express will route to the BULKDATA ancillary queue.

EE

Indicates that QDIO inbound workload queueing routing variables for the Enterprise Extender (EE) ancillary queue are to be included in the report. The EE routing variables are comprised of destination IP addresses, destination ports, and protocol. That combination uniquely identifies those packets that the OSA-Express will route to the EE ancillary queue.

IPSEC

Indicates that QDIO inbound workload queueing routing variables for the IPSec ancillary queue are to be included in the report. The IPSec routing variables include IPSec protocols AH and ESP and a combination of protocol and port for UDP encapsulated ESP packets. These routing variables uniquely identify those packets that the OSA-Express will route to the IPSec ancillary queue.

REGADDRS

Indicates that registered Layer 2 MAC addresses or Layer 3 unicast and multicast addresses are to be included in the report.

SYSDIST

Indicates that QDIO inbound workload queueing routing variables for the SYSDIST ancillary queue are to be included in the report. The SYSDIST routing variables are comprised of destination IP addresses and protocol. That combination uniquely identifies those packets that the OSA-Express will route to the SYSDIST ancillary queue.

ZCXADDRS

Indicates that QDIO inbound workload queueing routing variables for the zCX ancillary queue are to be included in the report. The zCX routing variables are comprised of the zCX dynamic VIPA address and the protocol of TCP or UDP. These routing variables uniquely identify those packets that the OSA-Express will route to the zCX ancillary queue.

MAX = *lines*| *

The maximum number of lines to be displayed on the console. Valid *lines* values are in the range 4 - 65533. Specify an asterisk (*) to allow up to 65533 lines to be displayed.

- If MAX=* is specified and the report is truncated as the result of exceeding the multi-line WTO maximum, the following message is displayed:

```
Report truncated due to greater than 65533 lines of output
```

- In all other cases, the total number of lines that is displayed and the total number of lines that could have been displayed are shown in the following output line, where *n* is the number of lines displayed and *m* is the total number of lines that could have been displayed.

```
n of m lines displayed
```

Reply field descriptions

Interface

Interface name from the display command.

Base section:

The Base section of the report is displayed if the BASE modifier is specified or none of the modifiers are specified.

PortName
Portname specified on the INTERFACE definition, specified as the device name, or both when the datapath device is shared by both definitions. This name also matches *portname* on the VTAM TRLE definition.

PortNum
Physical port on the OSA-Express that is used for the interface.

Datapath
Hexadecimal datapath device address on the OSA-Express that is used for the interface.

RealAddr
Hexadecimal logical address and unit address of the interface.

PCHID
Physically installed channel path that is used by this QDIO datapath device.

CHPID
Channel path identifier that is used by this QDIO datapath device.

CHPID Type
The CHPID type of the interface, which can have the following values:

OSD	External network
OSM	Intra node management network
OSX	Intra ensemble data network

OSA code level
OSA-Express processor code level of the QDIO datapath device.

Gen
Generation of the OSA-Express feature. The following values are supported:

OSA-E3	OSA-Express3
OSA-E4S	OSA-Express4S

	OSA-E5S	OSA-Express5S
	OSA-E6S	OSA-Express6S
	OSA-E7S	OSA-Express7S

Active speed/mode

Switch speed and duplex mode of the interface. The following values are supported:

- 10 mb/sec half duplex
- 10 mb/sec full duplex
- 100 mb/sec half duplex
- 100 mb/sec full duplex
- 1000 mb/sec half duplex
- 1000 mb/sec full duplex
- 10 gigabit full duplex
- 10 gigabit/sec full duplex
- 25 gigabit/sec full duplex
- Unknown

Media

Transmission media (copper or fiber). If fiber is the transmission media, it can be single-mode fiber (LR/LX) or multimode (SR/SX). The following values are supported:

- Copper
- Multimode Fiber
- Single-mode Fiber

Jumbo frames

Indicates whether jumbo frames are supported.

Isolate

Indicates whether this TCP/IP stack is prohibited from communicating directly through the interface with other TCP/IP stacks that are sharing the OSA-Express feature.

PhysicalMACAddr

Physical Medium Access Control (MAC) LAN address for the interface.

LocallyCfgMACAddr

Local Medium Access Control (MAC) LAN address for the interface.

Queues defined

Out	Number of output priority queues that are defined for this interface.
In	Number of input queues that are defined for this interface. A value of 1 indicates only the primary queue is defined. A value larger than 1 indicates that QDIO inbound workload queueing ancillary queues are defined and the number of ancillary queues is 1 less than the value reported.

Ancillary queues in use
Number of QDIO inbound workload queueing ancillary input queues (SYSDIST, BULKDATA, and so on) in use by this interface.

Connection Mode
Connection mode of the interface. The following values are supported:

- Layer 2
- Layer 3

IPv4
Indicates whether an IPv4 interface is active for the datapath device.

IPv6
Indicates whether an IPv6 interface is active for the datapath device.

SAPSup
Information used for problem analysis by IBM support.

SAPEna
Information used for problem analysis by IBM support.

IPv4, IPv6, or Layer 2 attributes
This section displays the attributes for the interface.

VLAN ID
Decimal virtual LAN identification number that is defined on this interface.

VMAC Active
Indicates whether the interface is using a virtual MAC address.

Defined Router
The defined router attribute. This field is displayed for Layer 3 only when VMAC is not active. The following values are supported:

Pri
The interface is a primary router.

Sec
The interface is a secondary router.

Non
The interface is not a router.

Active Router
Indicates whether this interface is the active router for the OSA-Express feature. This field is displayed for Layer 3 only when VMAC is not active and is applicable only for PRIROUTER and SECROUTER interface configurations.

VMAC Addr
Displays the virtual MAC address in use for this interface. This field is displayed only when VMAC is active.

VMAC Origin
Indicates the origin of the virtual MAC address. This field is displayed only when VMAC is active. The following values are supported:

Cfg
The virtual MAC address was configured in the TCP/IP stack PROFILE

OSA
The virtual MAC address was assigned by the OSA-Express

VMAC Router
This field is displayed for Layer 3 only when VMAC is active. The following values are supported:

All
Indicates that the OSA-Express is routing everything that was received on the virtual MAC address to the interface without regard for registered addresses.

Local
Indicates that the OSA-Express is routing everything that is received on the virtual

MAC address, and to a registered IP address, to the interface.

AsstParms This field is displayed only for Layer 3. It is information used for problem analysis by IBM support.

OutCkSumEna This field is displayed only for Layer 3. It is information used for problem analysis by IBM support.

InCkSumEna This field is displayed only for Layer 3. It is information used for problem analysis by IBM support.

Registered Addresses

This is the registered addresses section of the report and is displayed if the REGADDRS modifier is specified or none of the modifiers are specified and only if there are registered addresses.

For Layer 3, there are four subsections that are included only if there are addresses to report:

- IPv4 Unicast Addresses
- IPv4 Multicast Addresses
- IPv6 Unicast Addresses
- IPv6 Multicast Addresses

For Layer 2, there is only one subsection which is included only if there are addresses to report:

- Layer 2 Multicast MAC Addresses

ARP

Indicates whether the OSA-Express is providing address resolution for the corresponding registered IPv4 address.

Addr

IPv4 or IPv6 address.

Total number of IPv4 addresses or Total number of IPv6 addresses

Shows the cumulative number of IPv4 or IPv6 addresses immediately preceding this message.

MAC

The Medium Access Control (MAC) LAN address corresponding to the Layer 2 or registered multicast IP address.

Total number of Layer 2 MAC addresses

Shows the cumulative number of MAC addresses immediately preceding this message.

Ancillary Input Queue Routing Variables

The Ancillary Input Queue Routing Variables section of the report is displayed if any of the following modifiers were specified, or none were specified:

- BULKDATA
- SYSDIST

Queue Type

Displays the workload name for an ancillary queue. The following values are supported:

BULKDATA

Specifies that the input queue is used for streaming workloads.

EE

Specifies that the input queue is used for enterprise extender workloads.

IPSEC

Specifies that the input queue is used for IP security workloads.

ZCX

Specifies that the input queue is used for IBM z/OS Container Extensions (zCX) workloads.

SYSDIST

Specifies that the input queue is used for sysplex distributor workloads.

Queue ID

Ancillary queue number.

Protocol

TCP

Src

Source address and port. This information is displayed only for the BULKDATA queue.

Dst

Destination address and port. This information is displayed only for the BULKDATA queue.

Total number of IPv4 connections or Total number of IPv6 connections

Displays the cumulative number of BULKDATA IPv4 or IPv6 Src/Dst combinations immediately preceding this message.

Example

These are examples of the OSAINFO display when zCX dynamic VIPAs are registered.

```
Display OSAINFO results for IntfName: QDIO4103L
PortName: QDIO4103  PortNum: 00  Datapath: 0E36   RealAddr: 0E36
PCHID: 1234         CHPID: C4    CHPID Type: OSD  OSA code level: ABCD
Gen: OSA-E6         Active speed/mode: 1000 Mb/sec half duplex
Media: Multimode Fiber         Jumbo frames: Yes  Isolate: No
PhysicalMACAddr: 40000E340300  LocallyCfgMACAddr: 000000000000
Queues defined Out: 4  In: 5   Ancillary queues in use: 2
Connection Mode: Layer 3       IPv4: Yes  IPv6: No
SAPSup: 000FB001               SAPEna: 000FB001
IPv4 attributes:
  VLAN ID:  501            VMAC Active: Yes
  VMAC Addr: 420001AA0E36  VMAC Origin: OSA       VMAC Router: All
  AsstParmsEna: 00206C57   OutCkSumEna: 0000003A  InCkSumEna: 0000001A
Registered Addresses:
  IPv4 Unicast Addresses:
    ARP: Yes  Addr: 172.16.1.3
    Total number of IPv4 addresses:      1
  IPv4 Multicast Addresses:
    MAC: 01005E000001  Addr: 224.0.0.1
    Total number of IPv4 addresses:      1
```

```
Ancillary Input Queue Routing Variables:
  Queue Type: ZCX       Queue ID:  5  Protocol: TCP
    Addr: 9.42.105.1
    Total number of IPv4 addresses:      1
  Queue Type: ZCX       Queue ID:  5  Protocol: UDP
    Addr: 9.42.105.1
    Total number of IPv4 addresses:      1
30 of 30 lines displayed
End of report
```

This is an example of when the new ZCXADDRS keyword is specified.

```
Display OSAINFO results for IntfName: QDIO4103L
Ancillary Input Queue Routing Variables:
  Queue Type: ZCX       Queue ID:  5  Protocol: TCP
    Addr: 9.42.105.1
    Total number of IPv4 addresses:      1
  Queue Type: ZCX       Queue ID:  5  Protocol: UDP
    Addr: 9.42.105.1
    Total number of IPv4 addresses:      1
11 of 11 lines displayed
End of report
```

Chapter 5

SNA Messages

Topics:

- IST1221I
- IST1230I
- IST1233I

IST1221I

chtyp DEV = *device_address* STATUS = *status* STATE = *system_state*

Explanation

VTAM issues this message as part of a message group in response to:

- A DISPLAY ID command to identify the operational status of all **READ** and **WRITE** subchannels.
- A DISPLAY ID command for an MPC line or a transport resource list entry (TRLE).
- A DISPLAY TRL command for an active TRLE.

The message group varies if the TRLE is using the Queued Direct I/O (QDIO) interface to either an IBM OSA-Express Adapter or a HiperSockets device. If the message group is for a TRLE that is not using QDIO, the message group will look as follows:

This message group displays a TRLE that does not represent an OSA-Express adapter or HiperSockets interface.

```
 IST075I  NAME = nodename, TYPE = LINE
 IST486I  STATUS = currentstatus, DESIRED STATE = desiredstate
 IST087I  TYPE = line_type, CONTROL = line_control, HPDT = hpdtvalue
 IST1954I TRL MAJOR NODE = trl_major_node_name
 IST1715I MPCLEVEL = mpc_level MPCUSAGE = mpc_usage
 IST1717I ULPID = ulp_id
[IST2219I resource ACTIVATION WAITING FOR MINIMUM NUMBER OF DEVICES]
[IST1801I UNITS OF WORK FOR NCB AT ADDRESS stor_addr]
[IST1802I pn  CURRENT = cur  AVERAGE = avg  MAXIMUM = max]
[IST1577I HEADER SIZE = hpsize DATA SIZE = dsize STORAGE = storage]
 IST1221I chtyp DEV = device_address STATUS = status STATE = system_state
#
[IST1577I  HEADER SIZE = hpsize DATA SIZE = dsize STORAGE = storage]
[IST1221I chtyp DEV = device_address STATUS = status STATE = system_state]
#
 IST314I  END
```

Note: VTAM displays all **WRITE** subchannel addresses for *nodename* value, followed by all **READ** subchannel addresses for *nodename*. For TCP/IP channel DLC connections, there is only one **R/W** subchannel.

The following is an example of the message group if it is for a TRLE that uses QDIO:

```
 IST075I  NAME = nodename, TYPE = LINE
 IST486I  STATUS = currentstatus, DESIRED STATE = desiredstate
 IST087I  TYPE = line_type, CONTROL = line_control, HPDT = hpdtvalue
 IST1954I TRL MAJOR NODE = trl_major_node_name
 IST1715I MPCLEVEL = mpc_level MPCUSAGE = mpc_usage
[IST1716I PORTNAME = port_name LINKNUM = link_num OSA CODE LEVEL =
 code_level]
[IST2263I PORTNAME = port_name PORTNUM = port_num OSA CODE LEVEL =
 code_level]
[IST2337I CHPID TYPE = ch_type   CHPID = chpid_num  PNETID = network_id]
[IST2184I QDIOSYNC = ALLINOUT - SYNCID = TRAP01 - SAVED = YES]
[IST1577I HEADER SIZE = hpsize DATA SIZE = dsize STORAGE = storage]
 IST1221I chtyp DEV = device_address STATUS = status STATE = system_state
 #
[IST1577I HEADER SIZE = hpsize DATA SIZE = dsize STORAGE = storage]
[IST1221I chtyp DEV = device_address STATUS = status STATE = system_state]
#
 IST924I  -------------------------------------------------------------
[IST1221I DATA  DEV = device_address STATUS = status STATE = system_state]
[IST1724I I/O TRACE = iotrc TRACE LENGTH = length]
[IST1717I ULPID = ulp_id ULP INTERFACE = ulp_interface]
[IST2309I ACCELERATED ROUTING ENABLED]
```

```
[IST2310I ACCELERATED ROUTING DISABLED]
[IST2331I QUEUE   QUEUE     READ             QUEUE                 ]
[IST2332I ID      TYPE      STORAGE          STATUS                ]
[IST2205I ------  --------  ---------------  ----------------------]
[IST2333I qid     qtype     storage_amount   qstat                 ]
[IST924I -------------------------------------------------------------]
[IST2457I OSA DETECTED INBOUND ERRORS ON queue_name QUEUE:]

[IST2458I   TOTAL INVALID INBOUND PACKETS = num_packets]

[IST2459I   LAST INVALID INBOUND PACKET AT time ON date]
[IST2460I   TOTAL INBOUND SBAL ERRORS     = error_count]

[IST2461I   LAST INBOUND SBAL ERROR AT error_time ON error_date]

[IST2462I   CONSECUTIVE SBAL ERRORS: LAST = last_count MAX = max_count]

[IST2463I   MAX CONSECUTIVE SBAL ERRORS AT error_time ON error_date]

[IST924I -------------------------------------------------------------]
[IST2305I NUMBER OF DISCARDED INBOUND READ BUFFERS = sbalcnt]
[IST2386I NUMBER OF DISCARDED OUTBOUND WRITE BUFFERS = wbufcnt]
[IST1757I PRIORITYx: congstate  PRIORITYx: congstate]
[IST1757I PRIORITYx: congstate  PRIORITYx: congstate]
[IST2190I DEVICEID PARAMETER FOR OSAENTA TRACE COMMAND = deviceid]
[IST1801I UNITS OF WORK FOR NCB AT ADDRESS stor_addr]
[IST1802I pn  CURRENT = cur  AVERAGE = avg  MAXIMUM = max]
IST924I  -------------------------------------------------------------
#
[IST1221I TRACE DEV = device_address STATUS = status STATE = system_state]
[IST1724I I/O TRACE = iotrc TRACE LENGTH = length]
[IST1717I ULPID = ulp_id ULP INTERFACE = ulp_interface]

[IST2319I PHYSICAL CHANNEL ID = pchid]
[IST2309I ACCELERATED ROUTING ENABLED]
[IST2310I ACCELERATED ROUTING DISABLED]
[IST2331I QUEUE   QUEUE     READ             QUEUE                 ]
[IST2332I ID      TYPE      STORAGE          STATUS                ]
[IST2205I ------  --------  ---------------  ----------------------]
[IST2333I qid     qtype     storage_amount   qstat                 ]
[IST2305I NUMBER OF DISCARDED INBOUND READ BUFFERS = sbalcnt]
[IST2386I NUMBER OF DISCARDED OUTBOUND WRITE BUFFERS = wbufcnt]
[IST1757I PRIORITY1: congstate  PRIORITY2: ****NA****]
[IST1757I PRIORITY3: ****NA****  PRIORITY4: ****NA****]
 IST1801I UNITS OF WORK FOR NCB AT ADDRESS stor_addr
IST1802I pn  CURRENT = cur  AVERAGE = avg  MAXIMUM = max
[IST924I  -------------------------------------------------------------]
#
 IST314I  END
```

Note: VTAM displays all **WRITE** subchannel addresses for the node specified by the *nodename* value, followed by all **READ** subchannel addresses for that node, followed by all **DATA** subchannel addresses, followed by all **TRACE** subchannel addresses. DATA subchannel addresses and TRACE subchannel addresses are displayed only for an OSA-Express TRLE. For each DATA and TRACE subchannel address that is currently being used by an upper-layer protocol (ULP), the name of the z/OS Communications Server ULP (for example, the TCP/IP procedure name) using that data subchannel is displayed.

IST075I

In the message text:

nodename

The name of the resource that was entered on the DISPLAY ID command.

nodetype

The resource type of the major or minor node. The *nodetype* value is always **LINE** for this message group.

IST087I

In the message text:

line_type

The type of line. The *line_type* value is always **LEASED** for this message group.

line_control

The *line_control* value is always **MPC** (multipath channel) for this message group.

hpdtvalue

The *hpdtvalue* can be one of the following:

YES

Indicates the connection is capable of performing channel I/O directly to or from communications storage manager (CSM) buffers.

NO

Indicates the connection is not capable of performing channel I/O directly to or from communications storage manager (CSM) buffers.

NA

Is displayed when the connection is not active.

IST486I

In the message text:

currentstatus

The current status of the node. See the *z/OS Communications Server: IP and SNA Codes* for status information.

desiredstate

The node state that is desired. See the *z/OS Communications Server: IP and SNA Codes* for status information. If VTAM cannot determine the desired state, *desiredstate* is *****NA*****.

IST1221I

In the message text:

chtyp

The type of subchannel. Possible values are READ, WRITE, R/W, DATA, or TRACE

device_address

The hexadecimal address of the subchannel that is displayed.

status

The condition or state of the subchannel that is displayed. Possible values are:

ACTIVE	Subchannel is active.
INOP	Subchannel path is inoperative.
RESET	Subchannel path is not ready.
SLOWDN	Subchannel path is in slowdown mode.
ACTPEND	DLC is in the process of activating.
OPEN.PEND	DLC is in the process of opening a connection.

IDX.PEND	DLC is in the process of IDXINIT transmission for a DATA channel.
START.PEND	DLC is in the process of starting data flow for a connection.
INACT.PEND	DLC is in the process of deactivating.

system_state

The *system_state* value can be one of the following:

ONLINE	An MVS™ VARY ONLINE command for the subchannel has completed successfully and the channel is now available for use.
OFFLINE	An MVS VARY OFFLINE command has been issued for the subchannel and the command has completed successfully. The subchannel is no longer available for use.
PEND_OFFLINE	An MVS VARY OFFLINE command has been issued for the subchannel and the subchannel is in the process of completing the command.
N/A	The system state cannot be determined for DATA subchannel addresses. This state is also displayed in cases where VTAM has not allocated or could not allocate the UCB for the subchannel. **Tip:** If the MVS status of the subchannel is required, you can use the D U,xxxx command, where xxxx is the subchannel address.

IST1577I

This message is displayed only when HPDT=YES in message IST087I. This message is not displayed if the TRLE is IUTSAMEH, which is the TRLE for same-host communication.

In the message text:

hpsize

The MPC header segment size, in bytes.

dsize

The maximum MPC data segment size, in kilobytes.

storage

The storage medium that is used for inbound data (on READ subchannels). Possible values are:

ECSA

An extended common service area buffer provided by the communications storage manager (CSM).

DATASPACE

A data space buffer provided by the communications storage manager (CSM).

*****NA*****

Not applicable. This value is issued for WRITE subchannels.

IST1715I

In the message text:

mpc_level

The level of MPC connection. Possible values are:

HPDT

Indicates that the connection is capable of performing channel I/O directly to or from communications storage manager (CSM) buffers.

NOHPDT

Indicates that the connection is not capable of performing channel I/O directly to or from communications storage manager (CSM) buffers.

QDIO

(Queued Direct I/O) Indicates that the connection performs channel I/O operations using direct IO instead of CCW channel operations. The connection is also HPDT capable, and can therefore perform the direct IO to or from communications storage manager (CSM) buffers.

mpc_usage

Indicates whether the MPC connection can be used exclusively by only one ULP, or shared by multiple ULPs. Possible values are:

SHARE

Indicates that the connection can be shared by multiple ULPs.

EXCLUSIVE

Indicates that the connection can only be used by the first ULP that requests usage of the MPC connection.

IST1716I

This message is displayed only for TRLEs representing an IBM OSA-Express Adapter or an IBM Open Systems Adapter used for native access to an ATM network.

In the message text:

port_name

The port name to be assigned to the port on the IBM Open Systems Adapter. Each IBM Open Systems Adapter has one *port_name* that is represented by one TRLE.

link_num

The relative adapter number of the OSA-Express Adapter port represented by this TRLE. For an IBM Open Systems Adapter used for native access to an ATM network, the *link_num* value is N/A.

code_level

The OSA processor code level of the OSA-Express. For some versions of OSA-Express, the *code_level* value is N/A. For detailed instructions about setting up an OSA-Express feature, see the *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

IST1717I

This message is displayed for all TRLEs that are currently being used by at least one ULP. A separate IST1717I message will be displayed for each ULP using this TRLE. For a dynamic TCP TRLE, or an exclusively owned TRLE, only one message with ULPID will be issued, because there can only be one ULP using each of these TRLEs. For an OSA-Express Adapter, one message with ULPID will be issued for each Datapath channel address in use by a ULP. For other TRLEs, more than one ULPID message might be issued, depending on how many upper-layer protocols are using the TRLE.

In the message text:

ulp_id

The name of a z/OS Communications Server upper-layer protocol (ULP) that is using the TRLE or using one of the datapath channels of an OSA-Express TRLE.

- For TCP/IP ULPs, the *ulp_id* value is the TCPIP job name.
- For ANNC ULPs, the *ulp_id* value is the SNA PU name.

- For ATM or EE ULPs, the *ulp_id* value is the XCA major node name.

For all TRLEs with MPCLEVEL = QDIO, IST1717I will also display the interface dedicated to this datapath channel address. For all TRLEs whose MPCLEVEL is not QDIO, the ulp_interface will be *NA*.

ulp_interface

The name of either the interface or the device that is using one of the datapath channels of an OSA-Express TRLE.

IST1724I

This message is issued in response to DISPLAY ID or DISPLAY TRL commands. This message appears for a TRLE representing an OSA-Express adapter.

In the message text:

iotrc

Specifies whether I/O Trace is active for this OSA-Express data device (ON or OFF).

length

Specifies the number of bytes being recorded for I/O Trace for this OSA-Express data device.

For information about setting up an OSA-Express feature, see *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

IST1757I

This message is issued in response to DISPLAY ID or DISPLAY TRL commands. This message will appear for a TRLE representing an OSA-Express Adapter.

In the message text:

x

The write priority level.

congstate

The congestion state of that priority level. The *congstate* value is **CONGESTED** when, at least once in the last congestion reporting window, all 128 writes for the priority level were unavailable. Otherwise *congstate* will be **UNCONGESTED**.

IST1954I

In the message text:

trl_major_node_name

The name of the TRL major node defining the TRLE.

IST2184I

This message is displayed for only a TRLE that represents an OSA-Express2 or later adapter and only when the OSA-Express2 or later adapter is armed for QDIOSYNC. See QDIOSYNC trace in *z/OS Communications Server: SNA Diagnosis Vol 1, Techniques and Procedures* for a description of the QDIOSYNC trace function.

In the message text:

armstate

The OPTION operand value from the MODIFY TRACE command or TRACE start option.

Tip: The OSA might be collecting more than what is specified by the *armstate* value while OSA merges the options for all Armed data devices.

Possible values are:

ALLIN

OSA is collecting inbound diagnostic data for all devices.

ALLINOUT

OSA is collecting inbound and outbound diagnostic data for all devices.

ALLOUT

OSA is collecting outbound diagnostic data for all devices.

IN

OSA is collecting inbound diagnostic data for devices defined to this VTAM.

INOUT

OSA is collecting inbound and outbound diagnostic data for devices defined to this VTAM.

OUT

OSA is collecting outbound diagnostic data for devices defined to this VTAM.

syncid

The SYNCID operand value from the MODIFY TRACE command or TRACE start option. This value is to be used as part of a correlator when the OSA-Express2 or later diagnostic data is captured.

saved_state

The SAVE operand value from the MODIFY TRACE command or TRACE start option. Valid values are YES or NO.

IST2190I

This message is issued in response to DISPLAY ID or DISPLAY TRL commands for a TRLE configured with an MPCLEVEL parameter value of QDIO representing an OSA-Express adapter. This message appears for each active datapath channel if the OSA supports the OSA-Express network traffic analyzer (OSAENTA) trace function. The message displays the DEVICEID parameter, a number that uniquely identifies this datapath channel to the OSA-Express adapter. When a TCP/IP stack is performing the OSAENTA trace function for this OSA, this DEVICEID parameter can be specified on a TCP/IP OSAENTA profile configuration statement or a VARY TCPIP,,OSAENTA command to limit the tracing to just the user of this data device. See OSA-Express network traffic analyzer trace in *z/OS Communications Server: IP Configuration Guide* for more information about the OSAENTA trace function.

In the message text:

deviceid

The form *cs-mf-cl-us*, where

cs is the Channel subsystem ID for this data path device.
mf is the LPAR multiple image facility ID for the LPAR using this datapath device.
cl is the control unit logical identifier for this datapath device.
ua is the unit address for this data path device.
Each identifier is a 2 digit hexadecimal value in the range 00-FF.

IST2219I

This message is issued if the *resource* value defines an MPC channel-to-channel group, the activation of which is presently suspended waiting for the minimum required number of read and write devices to become available.

In the message text:

resource

The name of the TRLE or MPC subarea line that defines the MPC group.

IST2263I

This message is displayed if either of the following scenarios is true:

- PORTNUM is specified on the QDIO TRLE definition statement.
- VTAM detected it is connected to an IBM OSA-Express3 or later feature in QDIO mode.

In the message text:

port_name

The port name to be assigned to the port on the IBM Open Systems Adapter. Each IBM Open Systems Adapter has one port name that is represented by one TRLE.

port_num

The OSA-Express3 or later physical port number to be used for this QDIO MPC group. For OSA-Express2 and earlier or later adapters, only one physical port is available, so the *port_num* value will be 0.

code_level

The OSA processor code level of the OSA-Express. For some versions of OSA-Express, the *code_level* value will be N/A.

IST2305I

This is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents HiperSockets or an OSA-Express adapter.

In the message text:

sbalcnt

The number of storage block address lists (SBAL) that have been discarded since the activation of the device.

IST2309I

This message indicates that the upper-layer protocol (ULP) that is using the datapath channel of the OSA-Express or HiperSockets TRLE is using accelerated routing. If the ULP is a TCP/IP stack, then you can display the accelerator routing table by issuing the Netstat ROUTe/-r command with the QDIOACCEL modifier for a particular TCP/IP stack. For details about how to display the accelerator routing table, see the Netstat ROUTe/-r report in *z/OS Communications Server: IP System Administrator's Commands*.

IST2310I

This message indicates that the upper-layer protocol (ULP) that is using the datapath channel of the OSA-Express or HiperSockets TRLE is not using accelerated routing.

IST2319I

This message is issued if the TRLE that is displayed represents an IBM iQDIO Adapter (CHPID).

In the message text:

pchid

The internal QDIO (IQD) physical channel identifier is an internal system generated identifier that represents the internal logical network. The *phcid* is associated with the IQD CHPID and can span the entire central processor complex (CPC), based on the system configuration of the IQD CHPID. Operating systems that are running on this CPC, which are connected to the same IQD physical channel identifier, are using the same internal logical network and therefore have network connectivity. The IQD *pchid* is subject to change during a power-on reset of the CPC, or with dynamic I/O updates for the IQD CHPID.

IST2331I

This message is the first of two header messages for the information displayed in message IST2333I.

IST2332I

This message is the second of two header messages for the information displayed in message IST2333I.

IST2333I

When OSA Express supports QDIO inbound workload queueing, z/OS Communications Server can initialize multiple input queues. IST2333I is displayed once for each initialized read queue.

In the message text:

qid

The queue identifier of the read queue. The *qid* value is in the form RD/*qid*. RD/1 represents the primary read queue and RD/2 through RD/8 represent any initialized ancillary read queues.

qtype

The queue type for this read queue. Possible values are PRIMARY, BULKDATA, EE, IPSEC, SYSDIST, or ZCX.

storage_amount

The amount of storage defined by the VTAM start option QDIOSTG (or IQDIOSTG for iQDIO data devices). The VTAM start option value can be overridden on an individual device basis when READSTORAGE is configured on the LINK or INTERFACE statement in the TCP/IP profile.

A *storage_amount* value of `***NA***` appears if the *qstat* value is not ACTIVE. The queue has no read buffers and cannot be used by OSA Express to present inbound data.

The *storage_amount* value is displayed both in megabytes and in the number of QDIO read buffers that are storage block access lists (SBALs) that VTAM will use for this data device for inbound (read) processing. The *storage_amount* value is in the following format:

```
n.nM(nnn SBALS)
```

where *n.n* is the amount of storage in megabytes and *nnn* is the number of SBALs.

For an OSA-Express in QDIO mode, the size of an SBAL is fixed at 64 KB. For an iQDIO (HiperSockets) device, the SBAL size is variable. The iQDIO SBAL size is configured in a hardware configuration definition (HCD) when the maximum frame size (MFS) is specified. The default MFS is 16 KB, and the values 24 KB, 40 KB, and 64 KB are also supported. For an iQDIO device, both the VTAM start option IQDIOSTG and TCP/IP profile LINK or INTERFACE statement parameter READSTORAGE have an effect only when an MFS of 64 KB was configured.

qstat

The status of this read queue. Possible values are:

ACTIVE	The queue type is initialized and currently in use by the TCP/IP stack.
INITIALIZATION FAILURE	The queue type failed to initialize and will not be used by the TCP/IP stack.
NOT IN USE	The queue type is not currently in use by the TCP/IP stack.
NOT SUPPORTED BY OSA	The queue type is not supported by the OSA-Express adapter and will not be used by the TCP/IP stack.

IST2337I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents HiperSockets or an OSA-Express adapter.

In the message text:

chpid_type

The type of channel path identifier (CHPID) used by this TRLE:

OSD

Channel type for an OSA-Express CHPID configured in QDIO mode.

OSM

Channel type for an OSA-Express CHPID configured for attachment to the intranode management network.

OSX

Channel type for an OSA-Express CHPID configured for attachment to the intraensemble data network.

IQD

Channel type for HiperSockets (Internal Queued Direct I/O) communications.

chpid_num

The hexadecimal channel path identifier (CHPID) for the OSA adapter or HiperSockets device.

network_id

The physical network identifier.

- When *chpid_type* is OSX, *network_id* is always **IEDN**.
- When *chpid_type* is OSD or IQD, *network_id* is either the configured network identifier of the adapter, or ****NA**** if no network identifier was configured for the adapter.
- For all other *chpid_type* values, *network_id* is ****NA****.

IST2386I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents HiperSockets or an OSA-Express adapter.

In the message text:

wbufcnt

The number of outbound write buffers that have been discarded since the activation of the device.

IST2457I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors have been detected since the activation of the device. This message is followed by a group of messages that display QDIO error information (can include IST2458I, IST2459I, IST2460I, IST2461I, IST2462I, and IST2463I, depending on the type of QDIO errors). This message and the group of messages that follow it can be displayed once for each read queue.

In the message text:

queue_name

The name of the inbound read queue.

IST2458I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

num_packets

The number of inbound packets marked with invalid frame errors since the activation of the device.

IST2459I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

time

The time of the last inbound invalid frame error.

date

The date of the last inbound invalid frame error.

IST2460I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

error_count

The count of inbound SBAL errors since the activation of the device.

IST2461I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

error_time

The time of last inbound SBAL error.

error_date

The date of last inbound SBAL error.

IST2462I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

last_count

The count of the last consecutive batch of SBAL errors.

max_count

The largest count for consecutive SBAL errors.

IST2463I

This message is issued in response to DISPLAY NET,ID=*trlename* or DISPLAY NET,TRL,TRLE=*trlename* commands when the TRLE represents an OSA-Express adapter, and OSA QDIO errors were detected since the activation of the device.

In the message text:

error_time

The time of occurrence of the largest count of consecutive SBAL errors.

error_date

The date of occurrence of the largest count of consecutive SBAL errors.

System action

Processing continues.

Operator response

For MPC or TRLE configurations defined with multiple READ and multiple WRITE devices, MPC dynamics enables an operator to dynamically add and remove subchannels to and from the MPC/TRLE group.

- If a READ or WRITE MPC or TRLE subchannel displays as OFFLINE, issue an MVS ONLINE command (for example, VARY cua,ONLINE) to dynamically add the device back to the MPC or TRLE group.
- If a READ or WRITE MPC or TRLE subchannel displays as ONLINE, and you want to remove the subchannel from the group, issue the MVS OFFLINE command (for example, VARY cua,OFFLINE) to dynamically remove the device from the MPC or TRLE group.

Restriction: For subarea MPC connections, the MPCDYN=YES operand must be coded on the MPC GROUP or LINE definition to enable MPC dynamics.

System programmer response

For message IST2333I, use the *storage_amount* value to confirm the system storage use and to tune the performance of a specific data device.

Routing code

2

Descriptor code

5

IST1230I

TIME = *time* DATE = *date* ID = *id*

Explanation

This message is the first in a group of messages that displays tuning statistics.

If message IST1505I appears in the display, see that message for a complete description of the message group for XCF connections.

If message IST1613I appears in the display, see that message for a complete description of the message group for TCP/IP resources.

If message IST2366I appears in the display, see that message for a complete description of the message group for "RoCE Express" resources.

A complete description of the message groups for other connections follows.

This message group displays tuning statistics for a multipath channel-to-channel adapter. Messages IST924I - IST1570I are repeated for each subchannel address.

```
IST1230I TIME     = time     DATE     = date     ID  = id
IST1231I IPDU     = ipdu     OPDU     = opdu
IST1569I INLP     = inlp     ONLP     = onlp
IST1232I TSWEEP   = tsweep   QSWEEP   = qsweep
IST924I  ------------------------------------------------------------
IST1233I DEV      = dev      DIR      = dir
IST1234I BSIZE    = bsize    MAXBYTES = maxbytes
IST1235I SIO      = sio      SLOWDOWN = slowdown
IST1236I BYTECNTO = bytecnto BYTECNT  = bytecnt
IST1570I NBYTECTO = nbytecto NBYTECT  = nbytect
IST314I END
```

This message group displays tuning statistics for a queued direct I/O (QDIO) or a HiperSockets (iQDIO) adapter. Message subgroup IST924I - IST1570I is repeated for each subchannel address. Message subgroup IST924I - IST1811I is repeated for each direct I/O READ or WRITE queue.

The IST924I - IST1811I message subgroup for a direct I/O READ queue contains different messages than the message subgroup that is displayed for a direct I/O WRITE queue. In this example, each direct I/O queue subgroup contains an actual value in the *dir* field in message IST1233I, so each subgroup can be identified clearly.

```
IST1230I TIME     = time     DATE     = date     ID  = id
IST1231I IPDU     = ipdu     OPDU     = opdu
IST1569I INLP     = inlp     ONLP     = onlp
IST1232I TSWEEP   = tsweep   QSWEEP   = qsweep
IST924I  -------------------------------------------------------
IST1233I DEV      = dev      DIR      = dir
IST1234I BSIZE    = bsize    MAXBYTES = maxbytes
IST1235I SIO      = sio      SLOWDOWN = slowdown
IST1236I BYTECNTO = bytecnto BYTECNT  = bytecnt
IST1570I NBYTECTO = nbytecto NBYTECT  = nbytect
IST924I  -------------------------------------------------------
IST1233I DEV      = dev      DIR      = RD/x (queue_type)
[IST1719I PCIREALO = pcirealo PCIREAL  = pcireal]
[IST1720I PCIVIRTO = pcivirto PCIVIRT  = pcivirt]
[IST1750I PCITHRSO = pcithrso PCITHRSH = pcithrsh]
[IST1751I PCIUNPRO = pciunpro PCIUNPRD = pciunprd]
[IST2316I EARLYINO = earlyino EARLYINT = earlyint]
[IST2317I ULPRETUO = ulpretuo ULPRETU  = ulpretu]
[IST1752I RPROCDEO = rprocdeo RPROCDEF = rprocdef]
[IST1753I RREPLDEO = rrepldeo RREPLDEF = rrepldef]
IST1754I NOREADSO = noreadso NOREADS  = noreads
IST1721I SBALCNTO = sbalcnto SBALCNT  = sbalcnt
IST1722I PACKCNTO = packcnto PACKCNT  = packcnt
IST2185I FRINVCTO = frinvcto FRINVCT  = frinvct
IST1236I BYTECNTO = bytecnto BYTECNT  = bytecnt
IST1810I PKTIQDO  = pktiqdo  PKTIQD   = pktiqd
IST1811I BYTIQDO  = bytiqdo  BYTIQD   = bytiqd
IST924I  -------------------------------------------------------
IST1233I DEV      = dev      DIR      = WR/1
IST1755I SBALMAX  = sbalmax  SBALAVG  = sbalavg
IST1756I QDPTHMAX = qdpthmax QDPTHAVG = qdpthavg
IST1723I SIGACNTO = sigacnto SIGACNT  = sigacnt
IST1721I SBALCNTO = sbalcnto SBALCNT  = sbalcnt
IST1722I PACKCNTO = packcnto PACKCNT  = packcnt
IST2242I SIGMCNTO = sigmcnto SIGMCNT  = sigmcnt
IST1236I BYTECNTO = bytecnto BYTECNT  = bytecnt
IST1810I PKTIQDO  = pktiqdo  PKTIQD   = pktiqd
IST1811I BYTIQDO  = bytiqdo  BYTIQD   = bytiqd
IST314I END
```

Many of the messages in this group contain both a count and an overflow. Both the count and overflow are maintained in unsigned 32-bit variables (unless otherwise indicated). Since an unsigned 32-bit variable can only contain a value up to and including 4294967295 ('FFFFFFFF'X), the variable will wrap through 0 if an increment results in exceeding this value. When such a wrap occurs, the overflow is incremented by 1. Therefore, the total count is determined as follows:

```
            Total = (overflow * 4294967296) + count
```

Note: in the context of the QDIO (Queued Direct IO) interface, a single DATAPATH subchannel address represents both read and write DMA (Direct Memory Access) based data transfer points. In addition, the write transfer point is comprised of multiple priority levels (sometimes referred to as queues). For QDIO DATAPATH subchannel addresses, this sequence of messages is repeated for each data transfer point or priority level.

IST1230I

In the message text:

date* and *time

The *date* and *time* values specify when the record was reported. See DATE and TIME formats for information about the *date* and *time* values.

id

The name of the link for which tuning statistics are being recorded. The link name is from one of the following sources:

- The name specified on the LINE definition statement in the associated channel-attached major node.
- The name specified on the TRLE definition statement in the TRL definition deck.
- An internally generated TRLE name.

IST1231I

In the message text:

ipdu

The total number of inbound PDUs received.

opdu

The total number of outbound PDUs transmitted.

IST1232I

In the message text:

tsweep

The number of sweeps initiated as a result of a timeout. A sweep is a special data block that is exchanged with the adjacent host to verify that data has not been lost. Only the host with the highest subarea number will initiate timer sweeps.

qsweep

The number of sweeps initiated as a result of excessive receive queue depth.

- Receive queue depth represents the number of bytes of data waiting to be passed to the next layer.
- Excessive receive queue depth usually indicates a problem with reserialization of the data segment. Possible causes are a mismatch in the speed of the read subchannels or a lack of available I/O buffers.

IST1233I

In the message text:

dev

The hexadecimal subchannel address of the device for which tuning statistics are being recorded. It corresponds to one of the addresses coded on the READ, WRITE, or DATAPATH statement on the LINE or TRLE definition statement.

dir

The direction of this resource. Possible values are:

- READ
- WRITE
- RD/*x* (*queue_type*)
  - The *x* value is the QDIO read input queue number for this QDIO data device.
  - The *queue_type* value is the queue type. The *queue_type* value can be a maximum of 8 characters. Possible values include PRIMARY, BULKDATA, SYSDIST, EE, IPSEC, ZCX, or N/A. The N/A value indicates that the queue is initialized but is not currently in use by the TCP/IP stack.
- WR/*x*
  - The *x* value is the QDIO write priority level.

IST1234I

In the message text:

bsize

The maximum buffer size supported by this device.

maxbytes

The number of bytes used in the largest channel program. This field provides information about the utilization or packing of data into the transmit or receive buffer. As this number approaches *bsize*, this indicates that maximum instantaneous utilization of the device's buffer has occurred.

IST1235I

In the message text:

sio

The number of start I/O operations counted for the subchannel. This number is reset each time VTAM reports tuning statistics and is expressed in decimal. The value of *sio* is never larger than 65535. If *sio* is 65535, its value is reset to 0 when the next start I/O operation takes place.

slowdown

The number of times slowdown mode has been entered. If *slowdown* is incrementing, this indicates a lack of available I/O buffers.

- If **DIR = READ** in message IST1233I, *slowdown* is incremented every time the channel program cannot be reinitiated immediately because of a lack of I/O buffers to unpack the inbound data.
- If **DIR = WRITE** in message IST1233I, *slowdown* is ****N/A****.

IST1236I

In the message text:

bytecnto

The PDU byte count overflow.

bytecnt

The byte count. This value represents the accumulated number of bytes of PDUs transmitted on the **WRITE** subchannel or received on the **READ** subchannel.

IST1569I

In the message text:

inlp

The total number of inbound NLPs received.

onlp

The total number of outbound NLPs transmitted.

IST1570I

In the message text:

nbytecto

The NLP byte count overflow.

nbytect

The NLP byte count. This value represents the accumulated number of bytes of NLPs transmitted on the **WRITE** subchannel or received on the **READ** subchannel.

IST1719I

In the message text:

pcirealo

The real PCI (Program Controlled Interrupt) overflow.

pcireal

The real PCI count. This value represents the accumulated number of real PCI interrupts fielded by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

A real PCI is an execution of the QDIO Program-Controlled Interrupt Exit as a result of a call from the system interrupt handler. The higher the ratio of real PCI to virtual PCI, the less successful QDIO is at avoiding the overhead of the system interrupt handler.

IST1720I

In the message text:

pcivirto

The virtual PCI overflow.

pcivirt

The virtual PCI count. This value represents the accumulated number of virtual PCI interrupts fielded by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

A virtual PCI is an execution of the QDIO Program-Controlled Interrupt Exit as a result of a call from the QDIO device driver. The higher the ratio of virtual PCI to real PCI, the more successful QDIO is at avoiding the overhead of the system interrupt handler.

IST1721I

- *sbalcnto* is the storage block address list (SBAL) count overflow.
- *sbalcnt* is the SBAL count. This value represents the accumulated number of SBALs used for I/O on the QDIO **WRITE** priority level or the QDIO **READ** data transfer point.

  An SBAL serves a similar function in QDIO that an indirect-data-address word (IDAW) list serves in traditional I/O.

IST1722I

- *packcnto* is packet count overflow.
- *packcnt* is packet count. This value represents the accumulated number of packets transmitted on the QDIO **WRITE** priority level or received on the QDIO **READ** data transfer point.

  A packet is a single unit of data as presented to the QDIO device driver (for example, a datagram or TCP packet).

IST1723I

- *sigacnto* is SIGA (Signal Adapter) count overflow.
- *sigacnt* is SIGA count. This value represents the accumulated number of SIGA instructions issued for the QDIO **WRITE** priority level.

  SIGA tells the QDIO adapter that data is ready to be written. This count will not directly correlate with the data rate as QDIO employs a tack-in mechanism similar to Seldom Ending Channel Program.

IST1750I

- *pcithrso* is threshold PCI overflow.
- *pcithrsh* is threshold PCI count. This value represents the accumulated number of threshold PCI interrupts fielded by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

  A threshold PCI is a real PCI generated by the adapter because one of the threshold conditions which controls the PCI processing was met.

  A count of zero indicates that the QDIO device driver is providing sufficient resources to keep pace with the inbound data stream from the adapter.

IST1751I

- *pciunpro* is unproductive PCI overflow.

- *pciunprd* is unproductive PCI count. This value represents the accumulated number of unproductive PCI interrupts fielded by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

  An unproductive PCI is a real PCI where the QDIO Program-Controlled Interrupt Exit failed to find any completed reads. An unproductive PCI will occur when a virtual PCI causes the processing of read completions for which a real PCI is pending. An unproductive PCI indicates that the system interrupt handler overhead was incurred unnecessarily.

IST1752I

- *rprocdeo* is read processing deferral overflow.
- *rprocdef* is read processing deferral count. This value represents the accumulated number of read processing deferrals by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

  A read processing deferral occurs when the QDIO Program-Controlled Interrupt Exit must defer a read completion because a control block cannot be obtained to represent the inbound data. Read processing deferrals (along with read replenishment deferrals) might cause the NOREADS count to be nonzero.

IST1753I

- *rrepldeo* is read replenishment deferral overflow.
- *rrepldef* is read replenishment deferral count. This value represents the accumulated number of read replenishment deferrals by the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

  A read replenishment deferral occurs when the QDIO Program-Controlled Interrupt Exit does not have enough available read buffers to tack-in a new read. Read replenishment deferrals (along with read processing deferrals) might cause the NOREADS count to be nonzero.

IST1754I

- *noreadso* is reads exhausted overflow.
- *noreads* is reads exhausted count. This value represents the accumulated number of times reads were exhausted on entry to the QDIO Program-Controlled Interrupt Exit for the QDIO **READ** data transfer point.

  This value is incremented by 1 each time the QDIO Program-Controlled Interrupt Exit is invoked and it detects that all the read buffers are full (the adapter has no place to move additional inbound data). A zero total is preferred because lack of read buffers might result in the adapter discarding inbound data. Examination of the read processing and read replenishment deferral counts might indicate the reason the QDIO device driver is not providing sufficient resources to accept the inbound data from the adapter.

IST1755I

sbalmax is the maximum number of active SBALs at the completion of the write initiation process for the QDIO **WRITE** priority level. This value will be in the range 0-128 (0 meaning the priority level had no outbound activity in the interval and 128 meaning that at one point in the interval ALL the SBALs for the priority level were active).
sbalavg is the average number of active SBALs at the completion of the write initiation process for the QDIO **WRITE** priority level. This value will be in the range 0-128 (0 meaning the priority level had no outbound activity in the interval and 128 meaning that every time the QDIO write initiator completed, all 128 SBALs were active).

IST1756I

- *qdpthmax* is the maximum number of work elements left on the outbound work queue at the completion of the write initiation process for the QDIO **WRITE** priority level. This value will be a number in the range 0–54 or a character constant **>254**.
- *qdpthavg* is the average number of work elements left on the outbound work queue at the completion of the write initiation process of the QDIO **WRITE** priority level. This value will be a number in the range 0–254 or a character constant **>254**.

  A nonzero value for either of these fields suggests the adapter is not accepting outbound data as fast as the device driver is presenting it.

IST1810I

pktiqdo is the IQDIO packet count overflow.

pktiqd is the IQDIO packet count. If this message is for a QDIO or IQDIO DATAPATH device with a direction of read, the value represents the accumulated number of packets received on this particular QDIO or IQDIO READ data transfer point and subsequently routed outbound using IQDIO routing or QDIO Accelerator. If this message is for a QDIO or IQDIO DATAPATH device with a direction of WR, the value represents the accumulated number of packets transmitted on the write priority level for this device after being received inbound and subsequently routed outbound using IQDIO routing or QDIO Accelerator.

A packet is a single unit of data as presented to the QDIO device driver (for example, a datagram or TCP packet).

IST1811I

bytiqdo is the IQDIO byte count overflow.

bytiqd is the IQDIO byte count. If this message is for a QDIO or IQDIO DATAPATH device with a direction of read, the value represents the accumulated number of bytes received on this particular QDIO or IQDIO READ data transfer point and subsequently routed outbound using IQDIO routing or QDIO Accelerator. If this message is for a QDIO or IQDIO DATAPATH device with a direction of WR, the value represents the accumulated number of bytes transmitted on the write priority level for this device after being received inbound and subsequently routed outbound using IQDIO routing or QDIO Accelerator.

IST2185I

frinvcto is the frame invalidation error count overflow.

frinvct is the frame invalidation error count. This message applies to inbound data received on a QDIO DATAPATH device. The value represents the accumulated number of frame invalidation packets that were received and discarded. A frame invalidation packet is an inbound packet marked as invalid by OSA-Express3 or later.

Result: These fields are 0 for OSA-Express, OSA-Express2 or later, and iQDIO ports.

For detailed instructions about setting up an OSA-Express feature, see the *IBM Z: Open Systems Adapter-Express Customer's Guide and Reference*.

IST2242I

sigmcnto is the iQDIO (HiperSockets) multiple write SIGA count overflow.

sigmcnt is the iQDIO (HiperSockets) multiple write SIGA count. This value is the accumulated number of signal adapter (SIGA) multiple write instructions issued for the IQDIO WRITE priority level.

SIGA multiple write tells the IQDIO device that multiple storage block address lists (SBALs) with data are ready to be written. This count does not directly correlate with the data rate because the number of SBALs represented by each SIGA multiple write might vary.

IST2316I

earlyino is the early interrupt count overflow.

earlyint is the early interrupt count. This message applies to inbound data that is received on a QDIO DATAPATH device. The value represents the number of times that an OSA-Express feature that is operating in optimized latency mode signaled the host with an interrupt that indicates that no data was currently available to process, but that data that is destined to this host was received.

Result: These fields are both 0 for any QDIO DATAPATH device that is not operating in optimized latency mode.

IST2317I

ulpretuo is the Upper-Layer Protocol (ULP) return with no data available count overflow.

ulpretu is the ULP return with no data available count. This message applies to inbound data received on a QDIO DATAPATH device. The value represents the number of times that VTAM presented data to the ULP from an OSA-Express feature that is operating in optimized latency mode, and when VTAM received the ULP return, it found that there was no more data to process.

Result: These fields are both 0 for any QDIO DATAPATH device that is not operating in optimized latency mode.

System action

Processing continues.

Operator response

To discontinue statistics recording, enter the MODIFY NOTNSTAT command.

System programmer response

For additional information on tuning and analyzing tuning statistics, see the *z/OS Communications Server: SNA Network Implementation Guide*.

Routing code

2

Descriptor code

4

Example

This message group displays tuning statistics for a Multi Path Channel Adapter with the following TRLE definitions:

```
 TRLE1C   TRLE  LNCTL=MPC,
              READ=(0408,02F0),
              WRITE=(0508,03F0)

IST1230I TIME     = 07570625   DATE     = 08102      ID  = TRLE1C
IST1231I IPDU     =      53987 OPDU     =      52031
IST1569I INLP     =          0 ONLP     =          0
IST1232I TSWEEP   =          0 QSWEEP   =          0
IST924I ------------------------------------------------------------
IST1233I DEV      = 0508       DIR      = WRITE
IST1234I BSIZE    =       4095 MAXBYTES =       3930
IST1235I SIO      =       8203 SLOWDOWN = ****NA****
IST1236I BYTECNTO =          0 BYTECNT  =  157031499
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST924I ------------------------------------------------------------
IST1233I DEV      = 03F0       DIR      = WRITE
IST1234I BSIZE    =       4095 MAXBYTES =       3930
IST1235I SIO      =      21988 SLOWDOWN = ****NA****
IST1236I BYTECNTO =          0 BYTECNT  =   77055685
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST924I ------------------------------------------------------------
IST1233I DEV      = 0408       DIR      = READ
IST1234I BSIZE    =       4095 MAXBYTES =        341
IST1235I SIO      =      53032 SLOWDOWN =          0
IST1236I BYTECNTO =          0 BYTECNT  =    8199621
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST924I ------------------------------------------------------------
IST1233I DEV      = 02F0       DIR      = READ
IST1234I BSIZE    =       4095 MAXBYTES =        301
IST1235I SIO      =         35 SLOWDOWN =          0
IST1236I BYTECNTO =          0 BYTECNT  =       7719
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST314I END
```

This message group displays tuning statistics for a QDIO or an iQDIO Adapter with the following TRLE definitions:

```
 NSQDIO11 TRLE  LNCTL=MPC,
```

```
                MPCLEVEL=QDIO,
                READ=(0E28),
                WRITE=(0E29),
                DATAPATH=(0E2A,0E2B),
                PORTNAME=(NSQDIO1,0)

IST1230I TIME     = 18051835   DATE     = 09182      ID  = QDIO101
IST1231I IPDU     =          0 OPDU     =          0
IST1569I INLP     =          0 ONLP     =          0
IST1232I TSWEEP   =          0 QSWEEP   =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E29       DIR      = WRITE
IST1234I BSIZE    =       4096 MAXBYTES =          0
IST1235I SIO      =          0 SLOWDOWN = ****NA****
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E28       DIR      = READ
IST1234I BSIZE    =       4092 MAXBYTES =          0
IST1235I SIO      =          0 SLOWDOWN =          0
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1570I NBYTECTO =          0 NBYTECT  =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = RD/1 (PRIMARY)
IST1719I PCIREALO =          0 PCIREAL  =         12
IST1720I PCIVIRTO =          0 PCIVIRT  =          0
IST1750I PCITHRSO =          0 PCITHRSH =          0
IST1751I PCIUNPRO =          0 PCIUNPRD =          0
IST2316I EARLYINO =          0 EARLYINT =          0
IST2317I ULPRETUO =          0 ULPRETU  =          0
IST1752I RPROCDEO =          0 RPROCDEF =          0
IST1753I RREPLDEO =          0 RREPLDEF =          0
IST1754I NOREADSO =          0 NOREADS  =          0
IST1721I SBALCNTO =          0 SBALCNT  =          6
IST1722I PACKCNTO =          0 PACKCNT  =          6
IST2185I FRINVCTO =          0 FRINVCT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =        528
IST1810I PKTIQDO  =          0 PKTIQD   =          5
IST1811I BYTIQDO  =          0 BYTIQD   =        444
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = RD/2 (SYSDIST)
IST1754I NOREADSO =          0 NOREADS  =          0
IST1721I SBALCNTO =          0 SBALCNT  =          6
IST1722I PACKCNTO =          0 PACKCNT  =          6
IST2185I FRINVCTO =          0 FRINVCT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =        528
IST1810I PKTIQDO  =          0 PKTIQD   =          6
IST1811I BYTIQDO  =          0 BYTIQD   =        528
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = WR/1
IST1755I SBALMAX  =          0 SBALAVG  =          0
IST1756I QDPTHMAX =          0 QDPTHAVG =          0
IST1723I SIGACNTO =          0 SIGACNT  =          0
IST1721I SBALCNTO =          0 SBALCNT  =          0
IST1722I PACKCNTO =          0 PACKCNT  =          0
IST2242I SIGMCNTO =          0 SIGMCNT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1810I PKTIQDO  =          0 PKTIQD   =          0
IST1811I BYTIQDO  =          0 BYTIQD   =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = WR/2
IST1755I SBALMAX  =          0 SBALAVG  =          0
IST1756I QDPTHMAX =          0 QDPTHAVG =          0
IST1723I SIGACNTO =          0 SIGACNT  =          0
```

```
IST1721I SBALCNTO =          0 SBALCNT  =          0
IST1722I PACKCNTO =          0 PACKCNT  =          0
IST2242I SIGMCNTO =          0 SIGMCNT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1810I PKTIQDO  =          0 PKTIQD   =          0
IST1811I BYTIQDO  =          0 BYTIQD   =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = WR/3
IST1755I SBALMAX  =          0 SBALAVG  =          0
IST1756I QDPTHMAX =          0 QDPTHAVG =          0
IST1723I SIGACNTO =          0 SIGACNT  =          0
IST1721I SBALCNTO =          0 SBALCNT  =          0
IST1722I PACKCNTO =          0 PACKCNT  =          0
IST2242I SIGMCNTO =          0 SIGMCNT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1810I PKTIQDO  =          0 PKTIQD   =          0
IST1811I BYTIQDO  =          0 BYTIQD   =          0
IST924I -------------------------------------------------------------
IST1233I DEV      = 0E2A       DIR      = WR/4
IST1755I SBALMAX  =          0 SBALAVG  =          0
IST1756I QDPTHMAX =          0 QDPTHAVG =          0
IST1723I SIGACNTO =          0 SIGACNT  =          0
IST1721I SBALCNTO =          0 SBALCNT  =          0
IST1722I PACKCNTO =          0 PACKCNT  =          0
IST2242I SIGMCNTO =          0 SIGMCNT  =          0
IST1236I BYTECNTO =          0 BYTECNT  =          0
IST1810I PKTIQDO  =          0 PKTIQD   =          0
IST1811I BYTIQDO  =          0 BYTIQD   =          0
```

IST1233I

DEV = *dev* DIR = *dir*

Explanation

VTAM issues this message as part of a group of messages that displays tuning statistics for multipath channel (MPC) attached resources. The first message in the group is IST1230I. See that message for a complete description.

Routing code

2

Descriptor code

4

Chapter 6

z/OS Upgrade Workflow

Topics:

- Upgrade to an IBM z14 server
- IP services: Ensure storage availability for IWQ zCX traffic

Upgrade to an IBM z14 server

Description

The IBM z14™ (z14™) is designed to be the cornerstone for a trust economy. The z14 provides a scalable, highly available platform that delivers differentiated value to enable business growth, reduce cost, and protect existing investments. z/OS V2R3 enhances the role of the z14 by supporting solutions for a trusted digital economy. Capabilities that are designed to optimize high availability, performance, security, and operational flexibility can help organizations grow and secure their most critical transaction environments.

The IBM z14 family of mainframes includes the following models, which are dual frame configurations: M01, M02, M03, M04, M05. The newest addition to the IBM z14 family, the IBM z14™ Model ZR1 (z14 ZR1), is a single frame configuration. The z14 ZR1 is designed for clients who do not need the capacity of the larger M01-M05 models.

In this document, references to "z14" pertain to all models of the IBM z14, including the z14 Model ZR1, unless otherwise noted.

z/OS base support for specific z14 functions depends on the z/OS release. Table 3: z14 server functions included in the base support for z/OS V2R2, z/OS V2R3, and z/OS V2R4. on page 78 lists the z14 functions and indicates whether each function is included in a specific z/OS release (Yes or No), or whether the function requires a PTF (PTF) or a web deliverable (Web). For information about finding the appropriate PTFs, see Actions you can take before you order an IBM z14 server. For more information about V1R13 support, see the *Table Notes*.

Table 3: z14 server functions included in the base support for z/OS V2R2, z/OS V2R3, and z/OS V2R4.

This table shows a summary view of which z14 server functions are included in the base support for z/OS V2R2, z/OS V2R3, and z/OS V2R4.

z14 function included in base z/OS support (Y/N)	V2R2 [4]	V2R3 [4]	V2R4 [4]
Base support [1]	PTF	PTF	Yes, and in PTFs
CHPID type CS5 for coupling	Yes	Yes	Yes
CHPID type OSE supporting 4 or 2 ports per feature	Yes	Yes	Yes
CHPID type OSM for intranode management network (INMN)	Yes	Yes	Yes
CPU measurement facility	PTF	Yes	Yes
Crypto Express6S Toleration [2]	One of the following: • z/OS V2.2 with Cryptographic Support for z/OS V1R13 - z/OS V2R2 (HCR77B1) with PTFs • z/OS V2.2 with Cryptographic Support for z/OS V2R1 - z/OS V2R2 (HCR77CO) with PTFs	PTF	Yes
FICON® Express8S (CHPID type FC)	Yes	Yes	Yes
FICON Express16S (CHPID type FC)	Yes	Yes	Yes

z14 function included in base z/OS support (Y/N)	V2R2 [4]	V2R3 [4]	V2R4 [4]
FICON Express16S+ (CHPID type FC) when using FICON or channel-to-channel (CTC) [3] • **FICON Express16S+ LX** • **FICON Express16S+ SX**	PTF	PTF	Yes
IBM z BladeCenter Extension (zBX) Model 004	Yes	Yes	Yes
IBM z Unified Resource Manager	Yes	Yes	Yes
New z14 machine instruction (assembler support) [3]	PTF	PTF	Yes
OSA-Express6S (CHPID type OSD) [3] • **OSA-Express6S GbE LX** • **OSA-Express6S GbE SX** • **OSA-Express6S 10 GbE LR** • **OSA-Express6S 10 GbE SR** • **OSA-Express6S 1000BASE-T**	Yes	Yes	Yes
Parallel Sysplex Infiniband (PSIFB) Coupling Link [5]	Yes	Yes	Yes
z/OS global resource serialization (GRS) support for FICON CTCs	Yes	Yes	Yes

Table Notes:

1. **z/OS V1R13 (compatibility only). Extended support contract for IBM Software Support Services for z/OS is required with PTFs.**
2. **z/OS V1R13 with PTFs, the Enhanced Cryptographic Support for z/OS V1R13-z/OS V2R1 (HCR77B0) Web deliverable installed, and an extended support contract for IBM Software Support Services.**
3. **z/OS V1R13 with an extended support contract for IBM Software Support Services and PTFs.**
4. **PTFs for base support have the following fix categories:**
   - **For IBM z14 models M01 through M05, use FIXCAT value "IBM.Device.Server.z14-3906.RequiredService" and the FIXCATs for earlier processors.**
   - **For IBM z14 Model ZR1, use FIXCAT value "IBM.Device.Server.z14ZR1-3907.RequiredService" and the FIXCATs for earlier processors.**
5. **InfiniBand coupling links are not supported on the z14 Model ZR1.**

The exploitation of specific z14 server functions depends on the z/OS release. Table 4: Exploitation of z14 server functions supported by z/OS V2R2, z/OS V2R3, and z/OS V2R4. on page 79 lists z14 functions and indicates whether each function is exploited by a specific z/OS release (Yes or No), or whether the function requires a PTF (PTF) or a web deliverable (Web). For information about finding the appropriate PTFs, see Actions you can take before you order an IBM z14 server. For more information about V1R13 support, see the *Table Notes*.

Table 4: Exploitation of z14 server functions supported by z/OS V2R2, z/OS V2R3, and z/OS V2R4.

This table shows a summary view of which z14 server functions can be exploited in z/OS V2R2, z/OS V2R3, and z/OS V2R4.

z14 function included in base z/OS support (Y/N)	**V2R2 [2]**	**V2R3 [2]**	**V2R4[2]**

Processors: The maximum number of processors that can be configured per server.
For IBM z14 models M01 through M05:
• **Up to 170 processors configurable as CPs, zIIPs, IFLs, ICFs, or optional SAPs.**
• **The sum of CPs and zIIPs configured in a single z/OS LPAR cannot exceed:**
  • **170 on z/OS V2R1 or later in non-SMT mode**
  • **128 on z/OS V2R1 or later in SMT mode.**
For IBM z14 Model ZR1:
• **Up to 30 processors configurable as CPs, zIIPs, IFLs, ICFs, or optional SAPs.**
• **The sum of CPs and zIIPs configured in a single z/OS LPAR cannot exceed 30 on z/OS V2R1 or later (in either SMT or non-SMT mode).**
	Yes	Yes	Yes

Two-way simultaneous multithreading (SMT)	Yes	Yes	Yes

Channel subsystems. The maximum number of channel subsystems (CSS) that can be configured per server:
• **For IBM z14 models M01 through M05:**
  • **Up to six channel subsystems**
  • **Up to four subchannel sets per CSS.**
• **For IBM z14 Model ZR1:**
  • **Up to three channel subsystems**
  • **Up to three subchannel sets per CSS.**
	Yes	Yes	Yes

Up to 4 TB of Redundant Array of Independent Memory (RAIM) per z/OS LPAR with z/OS V2.1 or later.	Yes	Yes	Yes
Transactional memory	Yes	Yes	Yes
2 GB Large Pages	Yes	Yes	Yes

z14 function included in base z/OS support (Y/N)	V2R2 [2]	V2R3 [2]	V2R4[2]
Logical partitions. The maximum number of logical partitions (LPARs) that can be configured per server: • **For IBM z14 models M01 through M05, up to 85 LPARs can be configured.** • **For IBM z14 Model ZR1, up to 40 LPARs can be configured.**	Yes	Yes	Yes
Coupling Facility Control Code Level 22 (CFLEVEL 22)	PTF	Yes [3]	Yes
Coupling Express LR (CHPID type CL5)	PTF	Yes	Yes
Support for 256 Coupling CHPIDs	Yes	Yes	Yes
Support of High Performance FICON (zHPF) single-track or multitrack operations (CHPID type FC)	PTF	Yes	Yes
IBM Integrated Coupling Adapter Fanout (ICA SR) (CHPID type CS5)	Yes	Yes	Yes
CHPID type OSX for access control to the intra-ensemble data network (IEDN) from IBM z14 to Unified Resource Manager functions	Yes	Yes	Yes
CHPID type OSC supporting TN3270E and non-SNA DFT	Yes	Yes	Yes
OSA-Express6S (CHPID type OSD) without maximum port exploitation	Yes	Yes	Yes
OSA-Express6S IWQ IPSec	PTF	PTF	Yes
OSA-Express6S IWQ zCX	No	No	PTF
CHPID type OSD with exploitation of two ports per CHPID	Yes	Yes	Yes
Inbound workload queuing for Enterprise Extender (CHPID type OSD)	Yes	Yes	Yes
Checksum offload for IPv6 packets (CHPID type OSD)	Yes	Yes	Yes
Checksum offload for LPAR-to-LPAR traffic for IPv4 and IPv6 packets (CHPID type OSD)	Yes	Yes	Yes
Large Send for IPv6 packets (CHPID type OSD)	Yes	Yes	Yes
Crypto Express6S support of VISA Format Preserving Encryption (VFPE)	Yes	Yes	Yes
Crypto Express6S support of greater than 16 Domains	Yes	Yes	Yes

z14 function included in base z/OS support (Y/N)	V2R2 [2]	V2R3 [2]	V2R4[2]
Crypto Express6S Exploitation	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)
Crypto Express6S support of PCI-HSM compliance	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)	Web: Cryptographic Support for z/OS V2R1 - z/OS V2R3 (HCR77C1)
10 GbE RoCE Express2 for Shared Memory Communications-Remote Direct Memory Access (SMC-R)	PTF	Yes	Yes
Single Instruction Multiple Data (SIMD) instruction set [4]	Yes	Yes	Yes
IBM Virtual Flash Memory [1]	Yes	Yes	Yes
zHyperLink Express	PTF	PTF	Yes
HiperDispatch enhancements	PTF	Yes	Yes
Data set encryption	PTF	Yes	Yes
Guarded Storage Facility (GSF)	PTF	Yes	Yes
IBM z/OS XL C/C++ compiler support of the ARCH(12) and TUNE(12) parameters.	No	Yes	Yes
IBM z/OS XL C/C++ runtime support of ARCH(12) and TUNE(12) parameters.	PTF[4]	Yes	Yes
Instruction Execution Protection (IEP)	PTF	Yes	Yes
zEDC capability using zEDC Express	Y (requires the z/OS zEDC software feature to be enabled)	Y (requires the z/OS zEDC software feature to be enabled)	Y (requires the z/OS zEDC software feature to be enabled)

Table Notes:

1. **z/OS V1R13 with PTFs, the z/OS V1.13 RSM Enablement Offering web deliverable installed, and an extended support contract for IBM Software Support Services. The web deliverable is available at .**
2. **PTFs for exploitation have the following fix categories:**
   - **For IBM z14 models M01 through M05, use FIXCAT value "IBM.Device.Server.z14-3906.Exploitation" and the FIXCATs for earlier processors.**
   - **For IBM z14 Model ZR1, use FIXCAT value "IBM.Device.Server.z14ZR1-3907.Exploitation" and the FIXCATs for earlier processors.**
3. **Full support (duplexing with master/slave for CFLEVEL 22 on a z14) requires a V2R3 PTF.**
4. **Running the code on z/OS V2R1 or z/OS V2R2 requires the PTF for APAR PI12281.**

Web deliverables are available from the z/OS downloads page at .

Table 5: Information about this upgrade action on page 82 provides more details about this upgrade action. Use this information to plan your changes to the system.

Table 5: Information about this upgrade action

This table provides more details about the upgrade action. Use this information to plan your changes to the system.

Element or feature:	Multiple.

When change was introduced:	• IBM z14 models M01 through M05, which became available September, 2017. • IBM z14 Model ZR1, which became available May, 2018.
Applies to migration from:	z/OS V2R3 and z/OS V2R2.
Timing:	Anytime before you introduce a z14 server into your environment.
Is the upgrade action required?	Yes, if you want to run z/OS V2R4, z/OS V2R3, or z/OS V2R2 on a z14 server, or if you want to run a coupling facility on a z14 server. If you plan to run only a coupling facility on a z14 system, only the sysplex-related actions are relevant. However, you must install the appropriate z/OS service for systems that are running on other servers that use the z14 coupling facilities.

Target system hardware requirements:	• z/OS V2R4 with IBM z14 (z14) requires a minimum of 8 GB of memory to IPL. When running as a z/VM guest or on an IBM Z® Personal Development Tool (zPDT), a minimum of 2 GB is required for z/OS V2R4. If the minimum is not met, a warning WTOR is issued at IPL.
	• A z14 with Hardware Management Console (HMC) Version 2.14.0 plus microcode change levels (MCLs) and the Support Element Version 2.14.0.
	• Additional hardware required for specific functions.
	  • IBM devices that are attached to z13, z13s, zEC12, and zBC12 servers are supported for attachment to z14 channels, unless otherwise noted. The subject I/O devices must meet FICON architecture requirements.
	  • IBM Virtual Flash Memory requires FC #0604
	  • zEDC Express requirements:
	  • IBM z14 models M01 through M05 require FC #0420.
	  • IBM z14 Model ZR1 requires FC #0420 and PCIe I/O drawer (FC #4001).
	  • 10 GbE RoCE Express2 requires FC #0412
	  • New OSA-Express6S cards require the following feature codes:
	  • OSA-Express6S GbE LX requires FC #0422
	  • OSA-Express6S GbE SX requires FC #0423
	  • OSA-Express6S 10 GbE LR requires FC #0424
	  • OSA-Express6S 10 GbE SR requires FC #0425
	  • OSA-Express6S 1000BASE-T requires FC #0426
	  • New FICON Express16S cards require the following feature codes:
	  • FICON Express16S+ LX requires FC #0427
	  • FICON Express16S+ SX requires FC #0428
	  • zHyperLink Express requires FC #0431
	  • Coupling Express LR requires FC #0433
	  • Crypto Express6S requires FC #0893
	  • Use of hardware acceleration for cryptographic operations, including the use of Visa, Inc. Data Secure Platform (DSP) functions requires a CPACF (FC #3863) or a CEX5S (FC #0890) feature, or both.
	  See Table Note 1.
	  • Use of a new Trusted Key Entry (TKE) workstation requires TKE Rack Mount w/4768 (FC #0085) or TKE w/4768 (FC #0086). The TKE workstation is a combination of hardware and software, with other features required, depending on which functions you plan to use. Different feature codes are required if you are upgrading an existing TKE workstation from an earlier server.
	  TKE 9.0 (FC #0879) is also required if you plan to use the TKE workstation to manage any of the following:
	  • A Crypto Express6S
	  • The new CCA mode PCI-HSM settings that are available on the Crypto Express6S
	  • Any Crypto Express feature that is running in IBM Enterprise PKCS #11 (EP11) mode.

Target system software requirements:	See the topic "Support is delivered by service, z/OS features, and web deliverables" in General recommendations and considerations for a z14 server and "Install the necessary z/OS service", as indicated in the PSP buckets that are described in Actions you can take before you order an IBM z14 server.

Other system (coexistence or fallback) requirements:	• The suggested practice is to install and run the required service on your existing server. This action enables you to fall back from a hardware perspective, yet maintain your software level.
	• If you have not installed the preconditioning and exploitation PTFs for CFLEVEL 17-21, note that those PTFs are required to be installed throughout your sysplex before implementing CFLEVEL 22. Use the appropriate fix category, as follows:
	  • For IBM z14 models M01 through M05, use FIXCAT value "IBM.Device.Server.z14-3906.RequiredService"
	  • For IBM z14 Model ZR1, use FIXCAT value "IBM.Device.Server.z14ZR1-3907.RequiredService"
	• If you plan to implement zEDC, you should install the z/OS toleration PTFs on systems that access data that is compressed using zEDC Express on a z14 server.
	• If you are using ICSF and plan to share keys with other z/OS systems that have an earlier level of ICSF, you should install the ICSF coexistence PTFs on those earlier levels of ICSF. For ICSF coexistence PTFs, see Actions you can take before you order an IBM z14 server for the appropriate SMP/E FIXCAT values to use.

Restrictions:	See Restrictions for a z14 server.
System impacts:	None.
Related IBM Health Checker for z/OS check:	None.

Table Note:

1. **IBM Z® cryptography features include VISA Format Preserving Encryption (VFPE) technology, which forms part of the Visa, Inc. Data Secure Platform (DSP). The use of VFPE technology requires a service agreement with Visa, Inc. For details, contact your Visa account manager or Visa at P2PE@visa.com. Clients who use IBM Z® cryptography features, but do not make use of the FPE functionality, are not required to enter into any such agreement with Visa.**

Steps to take

Follow the General recommendations and considerations for a z14 server, adhere to the Restrictions for a z14 server, and perform the tasks that are described in the following topics:

- Actions you can take before you order an IBM z14 server
- Actions you can take after installing an IBM z14 server
- Accommodate functions for the z14 to be discontinued on future servers

Reference information

For more information about new functions to exploit, see the IBM z14 Library.

IP services: Ensure storage availability for IWQ zCX traffic

As of z/OS V2R4 with VTAM APAR OA58300 and TCP/IP APAR PH16581, the processing of IPAQENET and IPAQENET6 INTERFACE statements is enhanced when you use OSA-Express6S. If you enabled QDIO inbound workload queuing (WORKLOADQ) and you have defined a dynamic VIPA for IBM z/OS Container Extensions (zCX), an additional ancillary input queue (AIQ) is established for zCX inbound traffic. Additional storage is allocated for this input queue.

Each AIQ increases storage utilization in the following two areas:

- Approximately 36 KB of fixed ECSA
- 64-bit CSM HVCOMMON for READSTORAGE

  If you are using zCX, when the first zCX tunnel is activated (protocols ESP or AH), the new AIQ is backed with 64-bit CSM HVCOMMON fixed storage. The amount of HVCOMMON storage that is used is based on the specification of the INTERFACE READSTORAGE parameter.

If you configured QDIO inbound workload queuing (WORKLOADQ), ensure that sufficient fixed ECSA and fixed (real) 4 KB CSM HVCOMMON storage is available for the AIQ for zCX traffic.

The following table provides more details about this upgrade action. Use this information to plan your changes to the system.

Table 6: Information about this upgrade action

Element or feature:	z/OS Communications Server.
When change was introduced:	z/OS V2R4 (with VTAM APAR OA58300 and TCP/IP APAR PH16581).
Applies to upgrade from:	z/OS V2R4 (without VTAM APAR OA58300 and TCP/IP APAR PH16581).
Timing:	Before installing z/OS V2R4.
Is the upgrade action required?	No, but it is recommended if you have the WORKLOADQ parameter that is specified on the OSA IPAQENET and IPAQENET6 INTERFACE statements, you have zCX traffic and you have concerns about using additional ECSA or real (fixed) storage.
Target system hardware requirements:	OSA-Express6S Ethernet features or later in QDIO mode running on IBM z14 or z15.
Target system software requirements:	None.
Other system (coexistence or fallback) requirements:	None.
Restrictions:	None.
System impacts:	None.
Related IBM Health Checker for z/OS check:	To determine whether sufficient fixed storage is available for IWQ zCX enabling, use the following IBM health checks: • IBMCS,ZOSMIGV2R4PREV_CS_IWQZCX_*tcpipstackname* issues a message if the TCP/IP stack has inbound workload queuing (IWQ) and zCX enabled, but the OSA does not support IWQ for zCX. Only OSA-Express6S and later support IWQ for zCX. This health check is shipped INACTIVE and set to run once. • IBMCS,CSTCP_IWQ_ZCX_*tcpipstackname* issues a message if the TCP/IP stack has IWQ and zCX enabled, and the OSA does support IWQ for zCX. This health check is shipped ACTIVE and is set to run once. These health checks are available with VTAM APAR OA58300 and TCP/IP APAR PH16581 for V2R4.

Steps to take

1. If you are using or plan to use OSA-Express6S or later, verify that the following conditions are true:
   - WORKLOADQ is specified on the IPAQENET and IPAQENET6 INTERFACE statements.
   - Have zCX traffic; protocols ESP or AH.
2. If step 1 is applicable, IWQ zCX uses additional storage. Continue with step 3 - 6. Otherwise, there is no increase in storage usage, and no further action is required.

3. To calculate the total storage increase, count the total number of IPAQENET and IPAQENET6 INTERFACE statements that are coded with the WORKLOADQ parameter that are associated with OSA-Express6S or later. Make a note of the number.
4. Verify that sufficient ECSA is available. To calculate this, multiply the total INTERFACE statements that are counted in step 3 by 36 KB. The resulting number indicates how much additional ECSA is required.

   To determine whether sufficient ECSA is available to enable this function, verify the following ECSA definitions:

   - D CSM usage (PARMLIB member IVTPRM00, ECSA MAX value)
   - VTAM (Start Options CSALIMIT and CSA24)
   - TCPIP (GLOBALCONFIG ECSALIMIT statement in the TCPIP PROFILE)
5. Verify that sufficient real (fixed) storage is available. 64-bit fixed (CSM HVCOMMON) storage is used for the zCX AIQ read buffers. To calculate this value, multiply the total number of INTERFACE statements that are counted in step 3 by your configured READSTORAGE value (for example, 4 MB). You can verify how much storage is being used for READSTORAGE by using the D NET, TRLE command. The resulting number indicates how much additional fixed (64-bit) storage is required. To verify that the additional fixed storage is not a constraint for your system, take the following actions:

   - Use the DISPLAY CSM command to verify that sufficient fixed storage is available (CSM FIXED MAXIMUM defined in IVTPRM00).
   - Verify the actual amount of real storage available to this z/OS system by using D M=STOR or D M=HIGH.
6. If sufficient ECSA or real storage is not available, increase the available real storage or consider defining some of the OSA-Express6S (or later) INTERFACE statements with the NOWORKLOADQ parameter. If your CSM FIXED MAXIMUM is too low, increase this value in IVTPRM00.

Reference information

None.