


---

## z/OS Communications Server

### Performance enhancements related to sysplex distributor



© 2010 IBM Corporation

This presentation describes the new functions related to sysplex distributor performance in z/OS V1R12 Communications Server.

## Scalability, performance, constraint relief, and accelerators

- Performance improvements for sysplex distributor connection routing
- Performance improvements for streaming bulk data
- Extend sysplex distributor support for DataPower for IPv6
- Sysplex distributor support for hot-standby server

© 2010 IBM Corporation

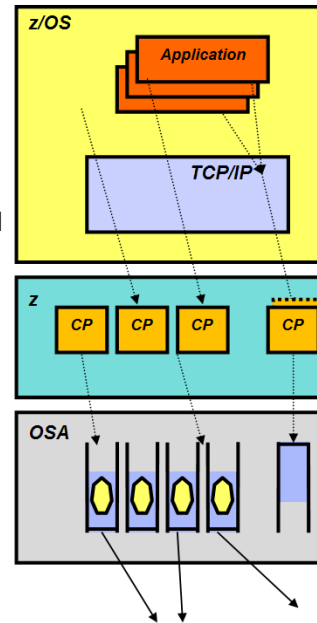
Inbound traffic separation is now being supported for OSA-Express3 using multiple read queues. TCP/IP will register with OSA which traffic is to be received on each read queue and OSA-Express Data Router function routes traffic to the correct queue. With sysplex distributor traffic separated onto its own read queue, it can be efficiently accelerated or presented to the target application. Also with bulk data traffic separated onto its own read queue, TCP/IP will service the bulk data queue from a single processor and this will assist in reducing the problems with out of order packet delivery.

Support is added for a sysplex distributor to load balance IPv6 connections to DataPower appliances. By deploying the sysplex distributor as the load balancing component for both tiers of distribution, z/OS administrators can use a single load balancing solution for the z/OS workload to the DataPower and z/OS application processing tiers. This simplifies the load balancing administration.

A new sysplex distribution method, HotStandby, is now supported. One target can be designated as a preferred target receiving all new incoming connection requests, and the hot-standby servers act as backup servers should the designated preferred target become unavailable.

## Pre V1R12 OSA inbound processing overview

- QDIO uses multiple write queues for traffic separation
  - Outbound traffic is separated by priority
  - Multiple CPs can manage the write queues
  
- QDIO uses only one read queue
  - All inbound traffic is received on the single read queue
  - Multiple CPs are used only when data is accumulating on the queue during bursts of inbound data
  - Single process for initial interrupt and read buffer packaging
    - TCP/IP stack performs inbound data separation



© 2010 IBM Corporation

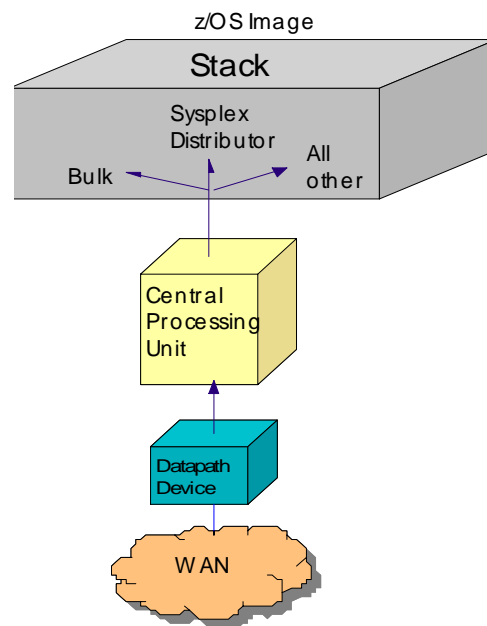
Outbound traffic separation (assignment to specific priority queue) on the multiple write queues can be accomplished by using Policy Agent and configuring a policy with the SetSubnetPrioToMask statement. Beginning in z/OS V1R11 Communications Server, outbound traffic separation can also be accomplished by using the WLM PRIORITYQ parameter on the GLOBALCONFIG statement. Each priority queue is processed independently of the others. For example, one processor can be building writes on priority queue one while another processor is building writes on priority queue four.

All inbound traffic is received on the single read queue. This includes both batch and interactive traffic and both traffic destined for this TCP/IP stack and traffic to be forwarded by this TCP/IP stack. Multiple processes only run for inbound traffic when data is accumulating on the read queue – typically during burst periods when z/OS Communications Server is not keeping up with the OSA. A single process is used to package the data, queue it, and schedule the TCP/IP stack to process it. This same process also performs acceleration functions, such as sysplex distributor connection routing accelerator. The TCP/IP stack separates the traffic types (sysplex distributor traffic, bulk inbound data such as FTP, and son on) to be forwarded to the appropriate stack component that will process them.

For these reasons, z/OS Communications Server is becoming the bottleneck as OSA-Express3 10GbE nears line speed. z/OS Communications Server is injecting latency and increasing processor utilization. This can impede scalability.

## Single read queue has negative impact for performance of bulk data

- Single read queue for all inbound QDIO traffic regardless of data type
- Single process for initial interrupt and read buffer processing
- TCP/IP stack must perform inbound data separation
- z/OS Communications Server is becoming the bottleneck as OSA nears 10GbE line speed
  - Injects latency
  - Increases processor utilization
  - Impedes scalability
- Multiple processes run when data is accumulating on the read queue
  - Out of order delivery of inbound traffic harms overall throughput



© 2010 IBM Corporation

All inbound QDIO traffic is received on a single read queue regardless of the data type.

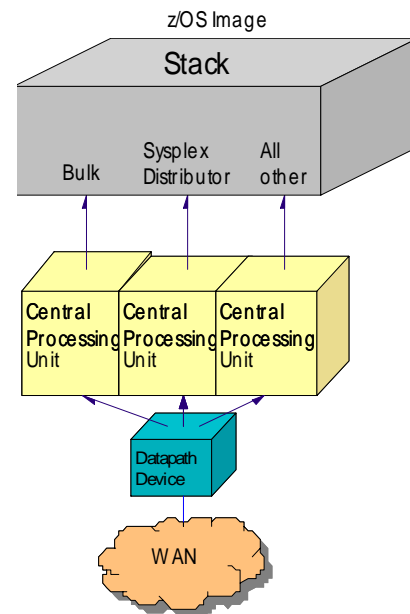
The maximum amount of storage available for inbound traffic is limited to the read buffer size (64K read SBALs) times the maximum number of read buffers (126).

Under the current z/OS Communications Server model, another QDIO input process will eventually be driven, and another TCP/IP stack thread, thus allowing multiple threads to process the one inbound read queue. However, this is only done when the OSA detects the host is now "falling behind" using the QDIO interrupt threshold algorithm.

Multiple processes are used for inbound traffic when data is accumulating on the read queue. This can cause bulk data packets for a single TCP connection to arrive at the TCP layer out of order. Each time the TCP layer on the receiving side sees out of order data, it transmits a duplicate ACK.

## OSA multiple inbound queue support

- Allow inbound QDIO traffic separation by supporting multiple read queues
  - “Register” with OSA which traffic goes to which queue
  - OSA-Express Data Router function routes to the correct queue
  
- Each input queue can be serviced by a separate process
  - Primary input queue for general traffic
  - One or more ancillary input queues (AIQs) for specific traffic types



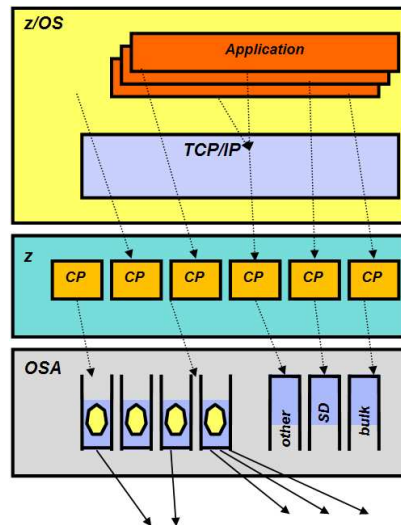
© 2010 IBM Corporation

With z/OS Communications Server V1R12, inbound traffic separation is supported using multiple read queues. TCP/IP will register with OSA which traffic to be received on each read queue. The OSA-Express Data Router function routes traffic to the correct queue.

Each read queue can be serviced by a separate process. The primary input queue is used for general traffic. One or more ancillary input queues (AIQs) are used for specific traffic types.

## Improved bulk transfer and sysplex distributor connection routing performance

- Supported traffic types
  - Bulk data traffic
  - Sysplex distributor traffic
  - All other traffic not backed up behind bulk data or SD traffic
  
- Dynamic LAN idle timer updated per queue



© 2010 IBM Corporation

The supported traffic types are streaming bulk data and sysplex distributor. Examples of bulk data traffic are FTP, TSM, NFS, and TDMF. Both IP versions are supported for all types of traffic.

With bulk data traffic separated onto its own read queue, TCP/IP will service the bulk data queue from a single processor. Using one processor solves the out of order delivery issue – there are no more race conditions. The determination of bulk data traffic is based in detecting streaming traffic. Bulk data traffic is registered per connection (5-tuple).

With sysplex distributor traffic separated onto its own read queue, it can be efficiently accelerated or presented to the target application. The determination of sysplex distributor traffic is based on active VIPADISTRIBUTE definitions. SD traffic is registered per DVIPA address.

All other traffic is processed simultaneous with the bulk data and sysplex distributor traffic. The dynamic LAN idle timer is updated independently for each read queue. This ensures the most efficient processing of inbound traffic based on the traffic type.

## TCP/IP profile changes

- INBPERF DYNAMIC WORKLOADQ enables QDIO inbound workload queuing
  - IPAQENET and IPAQENET6 INTERFACE statements only - no support for DEVICE/LINK definitions
  - INBPERF DYNAMIC WORKLOADQ requires VMAC
- Display TCPIP,,Netstat,DEVlinks to see whether QDIO inbound workload queuing is enabled for a QDIO

```
D TCPIP,TCPCS1,NETSTAT,DEVLINKS,INTFNAME=QDIO4101L
EZD0101I NETSTAT CS V1R12 TCPCS1
INTFNAME: QDIO4101L          INTFTYPE: IPAQENET   INTFSTATUS: READY
PORTNAME: QDIO4101  DATAPATH: 0E2A    DATAPATHSTATUS: READY
...
READSTORAGE: GLOBAL (4096K)
INBPERF: DYNAMIC
WORKLOADQUEUEING: YES
CHECKSUMOFFLOAD: YES
...
```

© 2010 IBM Corporation

The QDIO inbound workload queuing function is enabled with the INBPERF DYNAMIC WORKLOADQ setting on IPAQENET and IPAQENET6 INTERFACE statements. WORKLOADQ is not supported for INBPERF DYNAMIC on IPAQENET LINK statements. For steps to convert from IPv4 IPAQENET DEVICE, LINK, and HOME definitions to the IPv4 IPAQENET INTERFACE statement, see z/OS Communications Server: IP Configuration Guide.

The VMAC parameter can be specified with or without the MAC address.

For more information, see the IPAQENET INTERFACE and IPAQENET6 INTERFACE statements in z/OS Communication Server: IP Configuration Reference.

The Netstat DEVlinks/-d report has been updated for IPAQENET and IPAQENET6 interfaces to include a WorkloadQueueing field when InbPerf is Dynamic to indicate whether the QDIO inbound workload queuing function is enabled. The possible values are Yes indicating the function is enabled, No indicating the function is not enabled, and Unsupported indicating that the OSA-Express does not support the function. This field is not displayed when InbPerf is not Dynamic or for IPAQENET DEVICE/LINK definitions. This information can also be retrieved with a network management application that has been updated to use information returned by the Getlfs callable NMI.

In addition, the InbPerf field has been moved from the right column to the left column.

For more information, see the Netstat DEVlinks/-d report in z/OS Communications Server: IP System Administrator's Commands.

## Display QDIO inbound workload queues

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

```

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

```

© 2010 IBM Corporation

The output for the Display ID=trlename and Display TRL,TRLE=trlename commands has been updated to indicate whether the QDIO inbound workload queuing function is in use for the QDIO interface. For each input queue, it includes the queue ID and queue type in addition to the read storage. The queue type is PRIMARY for the primary input queue, BULKDATA for the bulk data AIQ, and SYSDIST for the sysplex distributor connection routing AIQ. The queue type value N/A indicates that the queue is initialized but is not currently in use by the TCP/IP stack. The queue ID and queue type can be used to correlate with VTAM tuning statistics, packet trace, and OSA-Express Network Traffic Analyzer (OSAENTA) trace output for the QDIO interface. For more information, see the DISPLAY ID command and DISPLAY TRL command in z/OS Communications Server: SNA Operation.



## Display new Bulk data session information (Netstat ALL/-A changes)

- Display TCPIP, Netstat, ALL to see whether QDIO inbound workload queuing is in use for bulk data

```

D TCPIP, TCPCS1, NETSTAT, ALL, CLIENT=USER1
EZD0101I NETSTAT CS V1R12 TCPCS1
CLIENT NAME: USER1                CLIENT ID: 00000046
LOCAL SOCKET:  ::FFFF:172.16.1.1..20
FOREIGN SOCKET:  ::FFFF:172.16.1.5..1030
  BYTESIN:          00000000000023316386
  BYTESOUT:         00000000000000000000
  SEGMENTSIN:       00000000000000016246
  SEGMENTSOUT:      0000000000000000922
  LAST TOUCHED:    21:38:53          STATE:          ESTABLSH
...
  ANCILLARY INPUT QUEUE: YES
  BULKDATAINTFNAME: LGBNS25G
...
APPLICATION DATA:  EZAFTPOS D USER1    C    PSSS
-----
1 OF 1 RECORDS DISPLAYED
END OF THE REPORT

```

The Netstat ALL/-A report has been updated to include the interface name for bulk data TCP connections that are using the QDIO inbound workload queuing function. For all other sockets, the ancillary input queue will display as N/A and the BULKDATAINTFNAME statement is not displayed. For more information, see the Netstat ALL/-A report in z/OS Communications Server: IP System Administrator's Commands. This information can also be retrieved with a network management application that has been updated to use information returned by the GetConnectionDetail callable NMI.

## Display number of Bulk data segments received (Netstat STATS/-S)

- Display TCPIP,,Netstat,STATS to see the total number of TCP segments received on bulk data queues

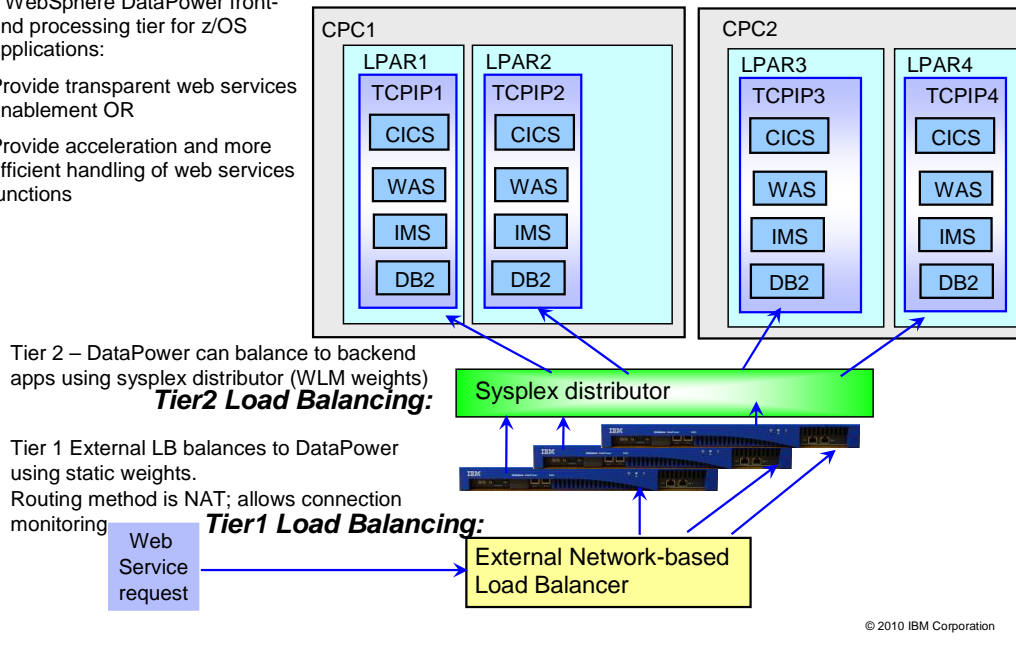
```
D TCPIP,TCPCS1,NETSTAT,STATS,PROTOCOL=TCP
EZD0101I NETSTAT CS V1R12 TCPCS1
TCP STATISTICS
  CURRENT ESTABLISHED CONNECTIONS      = 6
  ACTIVE CONNECTIONS OPENED            = 1
  PASSIVE CONNECTIONS OPENED           = 5
  CONNECTIONS CLOSED                   = 5
  ESTABLISHED CONNECTIONS DROPPED 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

The Netstat STATS/-S report has been updated to include the total number of segments received for all connections from the bulk data AIQ of the OSA-Express QDIO inbound workload queuing function. For more information, see the Netstat STATS/-S report in z/OS Communications Server: IP System Administrator's Commands. This information can also be retrieved with a network management application that has been updated to use information returned by the GetGlobalStats callable NMI.

## Background – DataPower in a sysplex using External Load Balancer

IBM WebSphere DataPower front-end processing tier for z/OS applications:

- Provide transparent web services enablement OR
- Provide acceleration and more efficient handling of web services functions



IBM WebSphere DataPower appliances are often used as a front-end processing tier to z/OS, providing transparent web services enablement or support of z/OS applications by providing acceleration and more efficient handling of web services functions.

In many of these environments, multiple DataPower instances are typically used in a cluster to allow for higher availability and scalability. This requires an external Tier 1 load balancer to distribute incoming web services connection requests to the DataPower appliance cluster.

Once DataPower has performed its portion of the processing of the incoming request, it will route a request to a z/OS application tier, such as CICS, IMS, WebSphere, or DB2. In a sysplex environment, the z/OS application tier will likely require a 2nd tier of load balancing to be performed.

Tier 2 load balancing can use DataPower's built-in load balancing support, which by default uses a round robin distribution. Another option, shown in the diagram, is to use the z/OS sysplex distributor which receives real time z/OS Workload Manager recommendations regarding current capacity and performance of each application and LPAR instance.

Sysplex distributor can be used as a load balancer for both tiers allowing one load balancing solution; this simplifies network management.

Support added in sysplex distributor and DataPower allows DataPower to communicate with sysplex distributor over an IPv4 TCP control connection which provides dynamic load balancing feedback:


Processor capacity information provided by the DataPower appliance allows sysplex distributor to optimize its load balancing decisions to avoid overloaded DataPower appliances by directing new requests to DataPower instances that are less used.

Connection state information provides server availability. It is used during a takeover. All connections are provided during takeover allowing Non-disruptive takeover by a backup sysplex distributor (planned or unplanned).

The return path from the DataPower appliance to a client can bypass sysplex distributor. Since Generic Resource Encapsulation (GRE) is used to forward inbound distributed packets to DataPower, the DataPower appliance can use the inner IP header to send outbound packets directly to client. Since DataPower sends connection termination information on a control connection, sysplex distributor does not need to monitor connection traffic to determine when a connection has ended.

## Support for IPv6 sysplex distribution to DataPower

- Support added to allow Tier 1 IPv6 sysplex distribution to DataPower
- Sysplex distributor communicates over a TCP Control Connection to a “sysplex distributor agent” running on DataPower



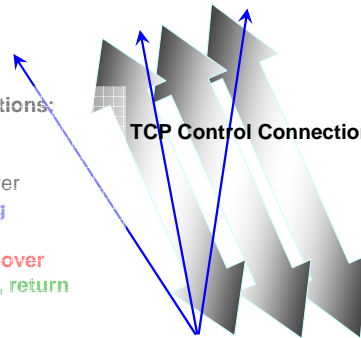
### IPv6 TCP Control Connections:

SD sends DVIPA/Ports  
DP sends

- DataPower weights & Server availability - **load balancing**
- Connection state --

- Non-disruptive takeover**
- IPv6 Encapsulation, return path bypasses SD**

### TCP Control Connections



Web  
Service  
request

In V1R12, support is added for a sysplex distributor to load balance IPv6 connections to DataPower appliances. By deploying the sysplex distributor as the load balancing component for both tiers of distribution, z/OS administrators can use a single load balancing solution for the z/OS workload to the DataPower and z/OS application processing tiers. This simplifies the load balancing administration.

Similar to IPv4 support, support added in sysplex distributor and DataPower allows DataPower to communicate with sysplex distributor over an IPv6 TCP control connection which provides dynamic load balancing feedback:

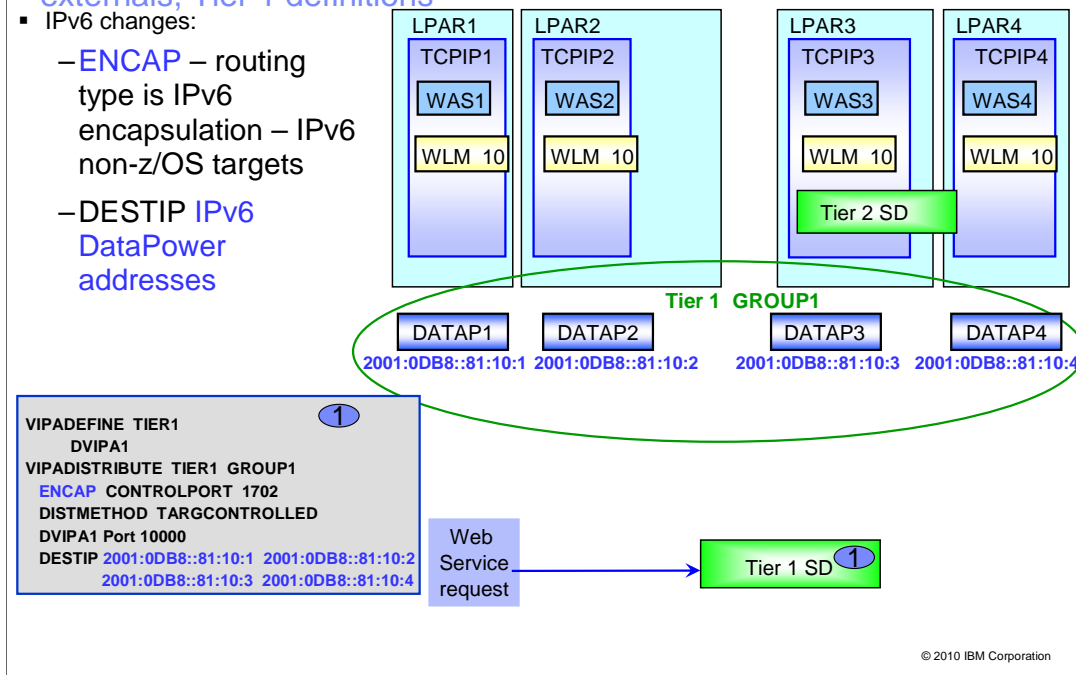
CPU capacity and Connection state information are received from the DataPower appliance on the TCP control connection.

Different from IPv4 support, IPv6 Encapsulation is used to forward inbound distributed packets to DataPower instead of GRE.

## Support for IPv6 sysplex distribution to DataPower, Configuration externals, Tier 1 definitions

- IPv6 changes:

- **ENCAP** – routing type is IPv6 encapsulation – IPv6 non-z/OS targets
- **DESTIP IPv6** DataPower addresses



© 2010 IBM Corporation

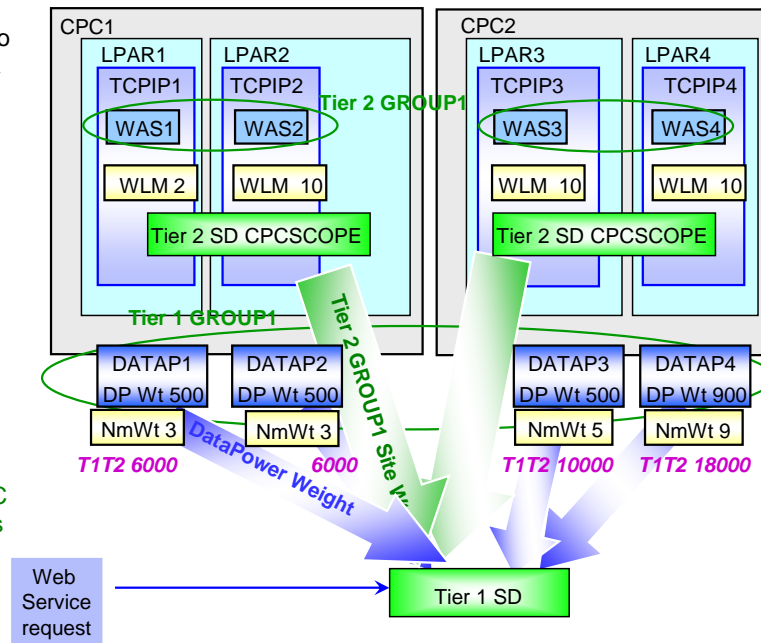
In V1R12, IPv6 support is added. The ENCAP parameter replaces the GRE parameter. ENCAP means that the Tier 1 targets are IPv6 non-z/OS targets, and IPv6 Encapsulation is used to send packets on a connection. IPv6 DVIPAs are supported to DataPower targets.

CONTROLPORT means that the sysplex distributor will bring up an IPv6 TCP control connection to the DataPower sysplex client listening on that port.

TARGCONTROLLED means the distribution is controlled by the target using the weights that it sends. The distributor uses a weighted round-robin distribution to the DataPower targets using the normalized DataPower weights received on the control connection. Weighted Active and Round Robin distribution can also be used.

## IPv6 Multi-Site support for DataPower

- Each DataPower is configured to connect to a group of servers on a site (CPC).
- Tier 1 load balancing uses normalized weights determined from combined Tier1/Tier2 weights:
- Combined Tier1/Tier2 weight =
  - DataPower weight \*
  - Tier 2 site weight - CPC weight of Tier 2 servers in matching group (GROUP1)



© 2010 IBM Corporation

IPv6 Multi-Site Support is added in V1R12. It works the same as IPv4 support.

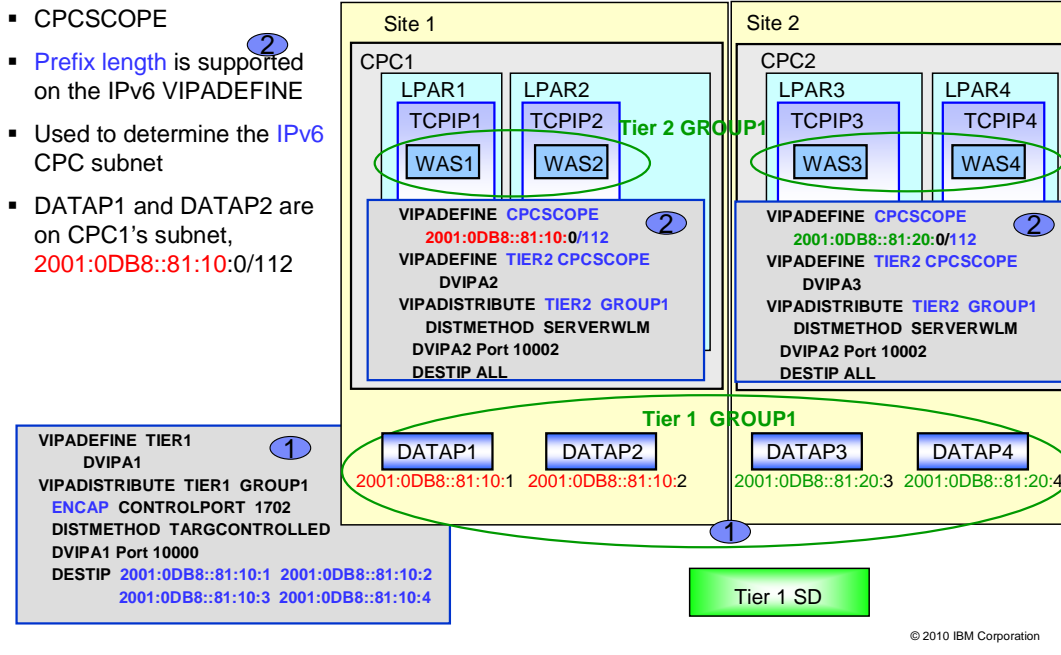
When DataPower is used in a multi-site environment, the load balancing configuration can be further enhanced by using multi-tier support of sysplex distributor in combination with load balancing to DataPower.

With this type of multi-CPC (multi-site) environment, each Tier 2 sysplex distributor is configured to load balance and route service requests to z/OS server application instances on its CPC. The Tier 2 distributors send a Site WLM weight of the group of target application instances in their CPC site to the Tier 1 distributor. The Tier 1 distributor makes its load balancing decision using a normalized weight based on the combined weight of each DataPower target along with the Site weight of its Tier 2 distribution targets. For DataPower DATAP1 and DATAP2, the weights of WAS1 and WAS2 are included in the combined weight. For DataPower DATAP3 and DATAP4, the weights of WAS3 and WAS4 are included in the combined weight.

For example, the combined Tier1/Tier2 weight of DATAP1 = (WAS1 weight 2 + WAS2 weight 10) \* DP Weight 500 = 6000.

The combined weights of all the DataPower appliances are then normalized against each other.

## IPv6 Multi-Site support for DataPower Configuration Example, Tier 2 definitions



For IPv6 Multi-Site support, each Tier 2 sysplex distributor is configured to load balance and route service requests to a set of IPv6 z/OS server application instances on its CPC. Since the Tier 2 DVIPA, DVIPA2, is defined with CPCSCOPE, the Tier 2 distributor will only load balance to targets on its CPC, CPC1 (even though DESTIP ALL is configured).

A CPCSCOPE DVIPA definition is used to define a CPC subnet. A CPCSCOPE DVIPA definition on one of the stacks in Site 1 (CPC1) is used to define the CPC1 subnet of 2001:0DB8::81:10:0/112. Another CPCSCOPE DVIPA definition on one of the stacks in Site 2 (CPC2) is used to define the CPC2 subnet of 2001:0DB8::81:20:0/112).

The Tier 1 distributor determines the CPC1 and CPC2 subnet as it receives the CPCSCOPE DVIPA definitions from the stacks that are on those CPCs. It determines that DATAP1 and DATAP2 have IP addresses that match the CPC1 subnet, 2001:0DB8::81:10:0/112. It determines that DATAP3 and DATAP4 have IP addresses that match the CPC2 subnet. It creates a combined Tier1/Tier2 weight for a DataPower appliance using the matching group weight that it receives from a distributor on that CPC with the appliance's weight.



## VIPADEFINE and VIPABACKUP statement syntax changes

- Addition of an IPv6 prefix length to each IPv6 address parameter

© 2010 IBM Corporation

The VIPADEFINE and VIPABACKUP statements have a new parameter – the IPv6 prefix length on each IPv6 address parameter.

This length is used with CPCSCOPE to define a site's DataPower subnet.



## IPv6 sysplex distribution to DataPower

### Netstat VIPADCFG/-F, VIPABACKUP and VIPADEFINE

#### – Show prefix length for IPv6 DVIPAs

```

NETSTAT VIPADCFG
MVS TCP/IP NETSTAT CS V1R11      TCPIP Name: TCPCS      19:49:12
Dynamic VIPA Information:
VIPA Define:
  IpAddr/PrefixLen: 201.2.10.17/26
  Moveable: Immediate SrvMgr: No Flg:
  IntfName: DVIPA6IF181I10
  IpAddr/PrefixLen: 2001:0DB8::81:10:1/112
  Moveable: Immediate SrvMgr: n/a Flg: 1

```

### Netstat VIPADCFG/-F DETAIL, VIPADistribute

#### – IPv6 external targets supported with routing type ENCAP

```

VIPA Distribute:
Dest:      201.1.181.10..10000
DestXCF:   9.42.105.53
SysPt:     No   TimAff: No   Flg: TargCtrl Tier1
OptLoc:    No
GrpName:   LOCALGROUP      RtgType: GRE CtrlPort: 01702

DestIntf:  DVIPA6IF181I10
Dest:      2001:0DB8::181:10..16000
DestXCF:   ALL
SysPt:     No   TimAff: No   Flg: TargCtrl Tier1
OptLoc:    0
GrpName:   LOCALGROUP      RtgType: ENCAP CtrlPort: 01702

```

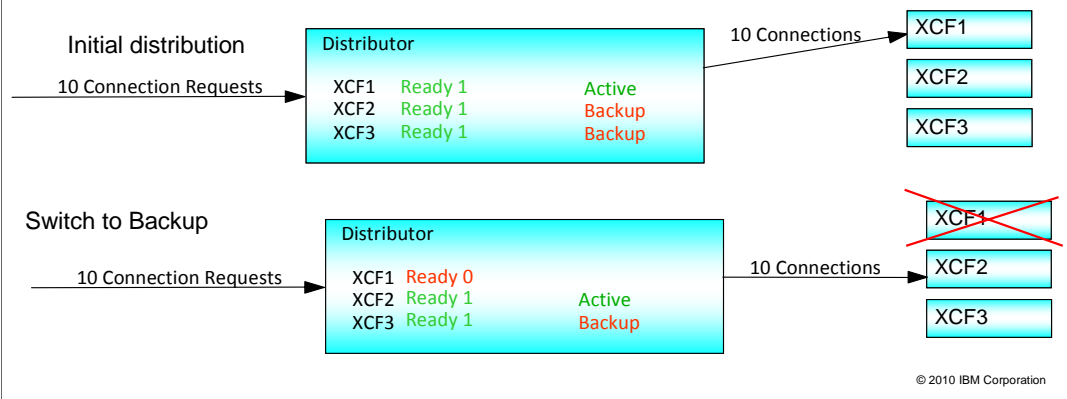
© 2010 IBM Corporation

The Netstat VIPADCFG/-F display will show the IPv6 prefix length when displaying information about the VIPADEFINE or VIPABACKUP statement.

The Netstat VIPADCFG/-F DETAIL display will show the IPv6 routing type of ENCAP when displaying information about a VIPADISTRIBUTE statement for IPv6 DataPower targets.

## Sysplex distributor support is needed for one active target with multiple backups

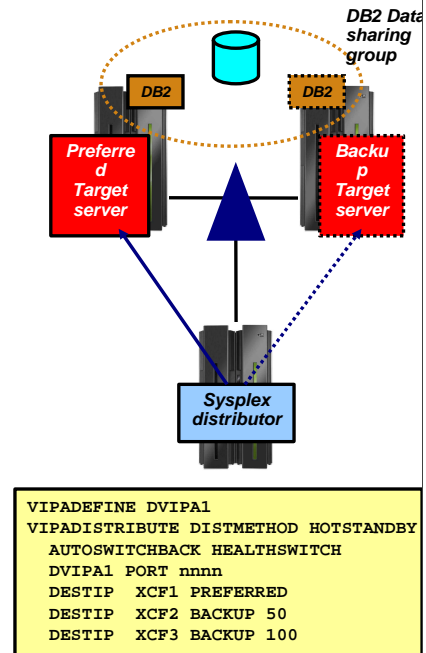
- Want only one active target to be used
  - Application runs more efficiently without database sharing
- Want multiple standby targets for availability
- 100% of workload distributed to active target if it is available
- Switch all workload to a backup, “hot-standby”, target if active target is not available



Customers have some applications that do not support database sharing or run more efficiently without database sharing, but they still want to have backup “hot-standby” targets for availability. The distributor keeps all of the workload on one active target as long as it is available. It switches to one of the backup targets if the active target is no longer available.

## Sysplex distributor hot standby support

- Have a single target server to receive all new connection requests
- Enable using a new HOTSTANDBY distribution method
  - One preferred target
    - AUTOSWITCHBACK option - switch to the preferred target if it becomes available
  - One or more backup targets ranked in order of preference
  - A target is not available when:
    - Not ready OR
    - Route to target is inactive OR
    - If HEALTHSWITCH option configured and target is not healthy



© 2010 IBM Corporation

A new distribution method, HotStandby, is supported. You have a single target server to receive all new connection requests. Other target servers are active but not receiving any new connection requests. The stack will automatically route traffic to a backup target server when the active target server is not available.

Using the new HOTSTANDBY distribution method, one target can be designated as a preferred target. If AUTOSWITCHBACK is configured, the distributor will switch to the preferred target when it is available unless the switch from the preferred target occurs because of health problems.

Designate one or more targets as backup targets. Use a backup target if the preferred target is not available. A target is not available if it is not ready, the route from distributor to target is inactive, or if HEALTHSWITCH is configured and the target is not healthy. A target is not healthy if the TSR is 0%, the abnormal termination rate is 1000 (out of 1000 transactions), or the server reported health is 0%.



## Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_PerformanceSD.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_PerformanceSD.ppt)

This module is also available in PDF format at: [../PerformanceSD.pdf](#)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, CICS, DataPower, DB2, IMS, TDMF, VTAM, WebSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2010. All rights reserved.