Sysplex distributor enhancements

z/OS Communications Server Development, Raleigh, North Carolina

This presentation describes the enhancements in z/OS V1R11 Communications Server for scalability, performance, constraint relief, and acceleration. This theme is a major area of enhancements in z/OS V1R11 Communications Server.

Sysplex distributor enhancements in z/OS V1R11

- ✓ Sysplex autonomics improvements for FRCA
- ✓ QDIO accelerator
- ✓ Sysplex distributor connection routing accelerator
- ✓ Sysplex distributor optimization for multi-tier z/OS workload
- ✓ Sysplex distributor support for DataPower®

Sysplex distributor has several performance-related enhancements in z/OS V1R11. These enhancements are described in the next slides.

Health of a sysplex distributor target stack - overview

- The distributor uses several health metrics to monitor target stacks and servers
  - **Target Server Responsiveness fraction** (TSR) is a compound measure
    - **Target Connectivity Success Rate** (TCSR)
      - Connectivity between distributing and target stacks - are new connection requests reaching the target?
    - **Connection Establishment Rate** (CER)
      - Network connectivity between server and client - are new connections being established?
    - **Server accept Efficiency Fraction** (SEF)
      - Is the server accepting new work? Are there connections on the backlog queue?

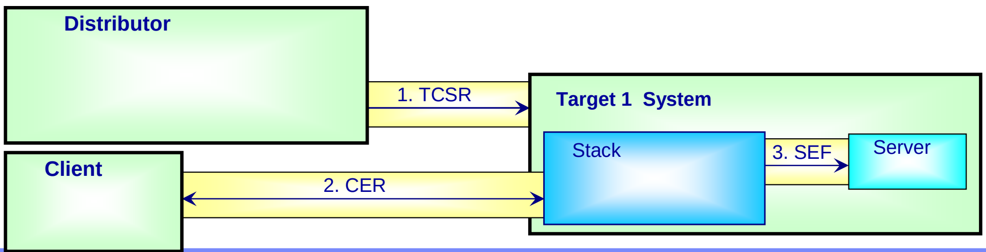

Let’s start by going over some background and terms.

The target server responsiveness (TSR) describes values calculated from measurements that sysplex distributor makes to measure the responsiveness of target servers in accepting new TCP connections approximately every 60 seconds. The TSR values are used to modify the weight used to favor servers that are more successfully accepting new TCP connection setup requests and are displayed as a percentage. A value of 100 indicates full responsiveness and zero indicates no responsiveness.

The target connectivity success rate (TCSR) measures the success rate of connections reaching the target stack from the distributing stack.

The connection establishment rate (CER) is an indication of how many TCP setup requests received by a target stack have become established over an interval.

The server efficiency factor (SEF) is one of the measures of the health of a particular server on a particular stack. The higher the value, 100 being the max, the healthier the application is. It generally measures how well a server application is processing, or accepting, established connections on its backlog queue. A reduced SEF can cause sysplex distributor or Load Balancing Advisor to reduce the number of new connections that it routes to a target stack server.

The backlog queue contains half-open connections, established connections that have not been accepted by the server, and FRCA connections that might never be accepted.

Half-open connections are connections which have not yet transitioned to established state (completed the three way handshake). They exist in the server’s backlog but have not yet been presented to the server to be accepted.

FRCA connections might never be accepted by the server application as requests can be handled by the cache. This is a normal FRCA connection and should not be counted against the health of this server.

Fast response cache accelerator - FRCA

- FRCA provides a hybrid Web server environment
  - Partly Web server
  - Partly TCP/IP stack
- The Web server loads a set of TCP/IP stack exits to parse received data
  - Allows FRCA to work with any protocol, not just HTTP
- Web pages are cached within the TCP/IP stack
  - Requests are handled without traversing the full protocol stack up to to the Web server
  - Significant performance improvements when compared to the Web server handling all requests

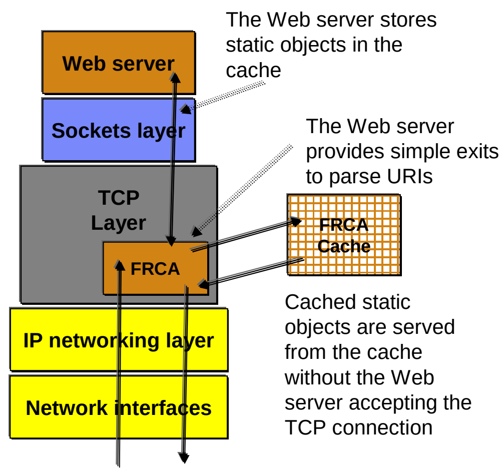

Fast Response Cache Accelerator (FRCA) is a Communications Server function that can significantly improve the performance of the z/OS HTTP Server and the WebSphere® Application Server on z/OS. Web pages are cached within the operating system kernel. Established FRCA connections are handled without traversing the entire kernel or entering the application space.

Established FRCA connections remain in the server's backlog queue unless they are presented to the application to be accepted. If a FRCA connection is serviced by the cache, it remains in the server backlog until the connection is closed.

FRCA use by WebSphere Application Server - overview

- **The FRCA cache is an HTTP cache that is maintained by TCP/IP**
- **Cached responses can be served with high performance using minimal processor cycles**
  - Serve static requests from the FRCA cache
    - Provide equivalent performance on WebSphere Application Server as is possible with the FRCA cache on the Web server
  - Serve dynamic content from the FRCA cache
    - Serve the same content that the Dynamic Cache serves
    - Record HTTP access log entries for requests served from the FRCA cache
- **WebSphere Application Server V7 exploits the FRCA cache**

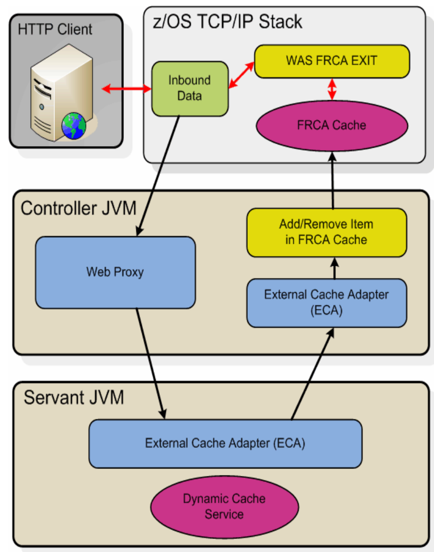

FRCA is used with the IBM HTTP Servers on some platforms. The FRCA cache is managed by TCP/IP. Requests can be served from the cache at the TCP/IP level without the need for the request to reach the application level. WebSphere Application Server currently has support to push static and dynamic content to the FRCA cache of an IBM HTTP Server on Windows NT® and Windows 2000 operating systems. This is done with the Dynamic Cache service's external cache group support. With this new support the WebSphere Application Server on z/OS activates the FRCA cache and pushes static and dynamic content to it directly. There is no involvement of the IBM HTTP Server in this scenario.

The WebSphere Application Server exploits TCP/IP interfaces to create the FRCA Cache and to add and remove objects from the cache. Once a listener socket is associated with the FRCA Cache, then all connections accepted for that listener are handled uniquely. When data arrives on one of the connections, TCP/IP drives a WebSphere Application Server exit with that data. That exit parses the HTTP request to determine the request and checks to see if a response for that request is in the FRCA Cache. If the response is in the cache then the exit builds the response headers and passes them back to TCP/IP. TCP/IP then sends the response headers and the response data from the cache. If the response is not in the cache, then TCP/IP passes the data up to WebSphere Application Server.

When a response is served from the FRCA cache, it is done very efficiently without WebSphere Application Server touching the response data. The performance of this is outstanding, especially considering the normal z/OS flow. In the normal z/OS flow, the request is read into the Controller region and then passed to the Servant region where the response is generated and passed back to the Controller before being written out to the socket. Also depending on the size of the request/response, there can be multiple trips back and forth between the Controller and the Servant regions. All of the cross address space calling is completely eliminated if the response is served from the FRCA cache where all the work is done in the TCP/IP address space.

The FRCA cache is configured as an external cache group under the Dynamic Cache services. If Dynamic Cache is active with the com.ibm.ws.cache.servlet.Afpa. ECA bean and servlet caching is enabled, then objects are added and removed from the FRCA cache in accordance to the caching policy (cachespec.xml) of the application.

Processor cost savings using FRCA compared to dyna cache

- **The amount of processor time needed to process a request is dramatically reduced using FRCA as compared to Dyna Caching**
  - **Dyna Cache is 37 times more costly than FRCA caching for 1k file sizes**
  - **Larger file sizes, 5k to 100k, Dyna Cache is 27 to 11 times more costly**

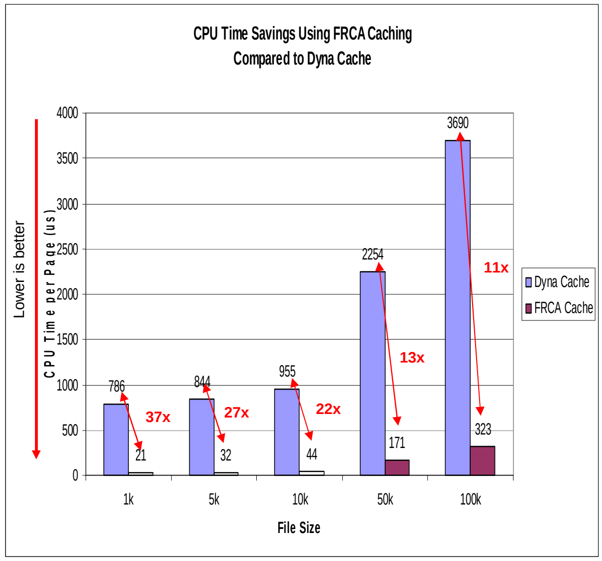

System Configuration
Workload: Simple File Server App
SUT: IBM z10 Processor (model 2097 – 720) 4 x 4.4 GHz, 32 GB Real
Driver: x/335 model 8676-21X, 4x3.06 GHz, 2 GB RAM

This chart compares the processor cost of serving static files of different sizes using both Dyna caching and FRCA caching. The application used is a simple file serving application with little or no application logic. The purpose is to demonstrate the performance benefits using the FRCA cache to serve static files as compared to existing caching techniques such as Dyna Caching. Five different file sizes were used to demonstrate the effectiveness of FRCA for both small and large files. For the 1k file size the improvement in the amount of processor consumed per request was dramatic. Dyna Cache took about 37 times more processor cycles to process a request than FRCA. As the file size got bigger the advantage over Dyna Cache was smaller but still 27 to 11 times faster. This is because as the file size increased, the processing required just to handle the data dominates the total cost for both Dyna Cache and FRCA cache.

Sysplex distributor enhancements

- FRCA connections that are serviced by the cache are not be measured against a server's health
- Half-open connections are not weighed as heavily against the server's health
  - Half-open connections reflect more of a routing issue than a problem with the server
- The CER is now a diagnostic aid only

- TCPSTACKSOURCEVIPA no longer uses MOVING DVIPAs (Dynamic VIPA)
  - Allow DVIPA to transition to original state
  - Use normal source IP address selection process

The solution to the first problem is that FRCA connections which are serviced by the FRCA cache are no longer counted against the SEF. These connections do still remain in the server's backlog but the count is subtracted before the calculation of the SEF is performed.

Half-open connections are not factored as heavily against a server's health now. They are more a measure of the route between the server and client. They are still reflected in the server's SEF but in a smaller measure.

Now that half-open connections are measured in the SEF, the CER is not used to gauge a target server's health. This field is now used for diagnostic purposes only.

TCPSTACKSOURCEVIPA uses the configured IP address for all outbound TCP connections. A DVIPA in MOVING state is currently a valid IP address for TCPSTACKSOURCEVIPA. Therefore, it is possible for a DVIPA in MOVING state to remain that way for an indefinite period of time. That is, as long as there are new outbound TCP connections. This violates the intent of the MOVING state. A DVIPA should remain in MOVING state only as long as it takes to complete processing on any connections that existed before the DVIPA moved to another stack.

In z/OS V1R11, DVIPAs in MOVING state are no longer valid for TCPSTACKSOURCEVIPA. The normal source IP address selection process is used if a DVIPA configured for TCPSTACKSOURCEVIPA has transitioned to MOVING state.

Display command example

- Display TCPIP,,NETSTAT,ALL to see the new ServerBacklog and FRCABacklog fields

```
MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPCS          16:49:28
Client Name: BPXOINIT                 Client Id: 00000021
  Local Socket: 0.0.0.0..10007
  Foreign Socket: 0.0.0.0..0
…
…
 MaximumBacklog:      0000000010
 CurrentBacklog:      0000000008
    ServerBacklog:    0000000003      FRCABAcklog: 0000000002
 CurrentConnections: 0000000000       SEF:                100
```

Established: Full three-way handshake has completed.

Half-open: You have received SYN and responded with a SYN+ACK, but are waiting for the final ACK from the client

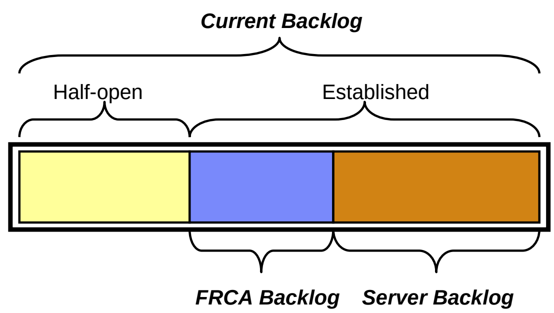

This is an example of the D TCPIP,,NETSTAT,ALL command that shows what the new ServerBacklog and FRCABacklog fields look like. The new fields are in bold blue.

The CurrentBacklog count is made up of half-open and established connections. The established connections are the sum of the connections in the ServerBacklog field plus the connections in the FRCABacklog field.

HiperSockets accelerator scope is limited

- Provides fast path IP forwarding for these DLC combinations:
  - Inbound OSA-Express™ QDIO → Outbound HiperSockets™
  - Inbound HiperSockets → Outbound OSA-Express QDIO
- No support for sysplex distributor acceleration
- No acceleration between two OSA-Express QDIO interfaces
- No acceleration between two HiperSockets interfaces

In V1R2, z/OS Communications Server added support for HiperSockets accelerator. This function provides accelerated routing (or redirecting of packets) which come inbound over OSA-Express QDIO and go outbound over HiperSockets (or inbound HiperSockets and outbound OSA-Express QDIO). This function occurs at the DLC layer such that the forwarding stack never processes the accelerated packets. This function is sometimes referred to as iQDIO routing.

While HiperSockets accelerator works for normal IP forwarding, it does not provide any benefits for sysplex distributor forwarding.

A new, more complete QDIO routing accelerator

- Provides fast path IP forwarding for these DLC combinations:
  - Inbound OSA-E QDIO → Outbound OSA-E QDIO or HiperSockets
  - Inbound HiperSockets → Outbound OSA-E QDIO or HiperSockets
- Adds sysplex distributor acceleration
  - Inbound packets over HiperSockets or OSA-E QDIO
  - When SD gets to the target stack using either:
    - Dynamic XCF connectivity over HiperSockets
    - VIPAROUTE over OSA-E QDIO
- Improves performance and reduces processor usage for such workloads
- Configure QDIOACCELERATOR on the IPCONFIG statement in the TCPIP profile

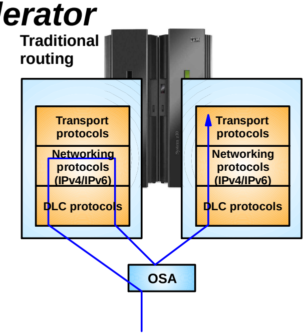

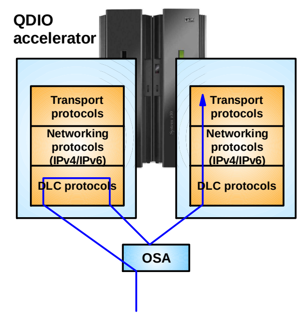

Page 10

QDIO Accelerator extends the HiperSockets Accelerator to provide fast path forwarding between two OSA-Express QDIO interfaces (or two HiperSockets interfaces). It also provides acceleration for sysplex distributor forwarding in any of the four supported inbound/outbound DLC combinations. The outbound interface can either be the dynamic XCF HiperSockets interface or can be a VIPAROUTE over an OSA-Express QDIO interface.

To enable QDIO Accelerator, configure QDIOACCELERATOR on the IPCONFIG statement in the TCPIP profile.

A subsection of the IP routing table is pushed down into the DLC layer so routing decisions can be performed at the DLC layer. Therefore, QDIO accelerator uses fewer processor resources and has lower latency.

QDIO accelerator cannot be used in combination with IP Security.

The IP layer is required to inspect packets when IP Security is enabled and cannot be bypassed in that case.

The QDIO accelerator in z/OS V1R11 supports IPv4 forwarding only.

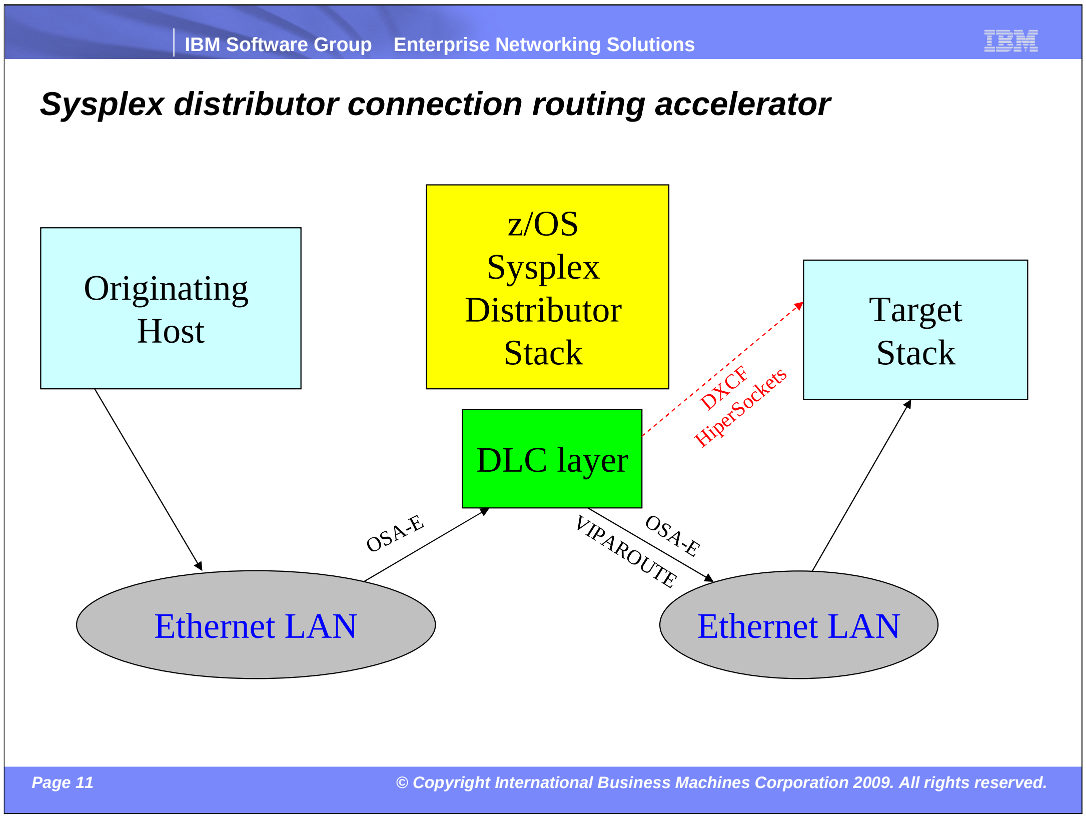

With sysplex distributor (SD), the distributor stack receives packets and forwards them to a target stack. This picture shows SD forwarding with QDIO Accelerator. In this example, the SD is using a VIPAROUTE over an OSA-Express QDIO interface to get to the target stack (as shown by the solid black arrows). This function also provides SD acceleration when the SD reaches the target stack using Dynamic XCF connectivity over HiperSockets (as shown by the dotted red arrow). And, while the example shows inbound OSA-Express QDIO, the SD acceleration function also applies to data received inbound over HiperSockets.

Sysplex distributor accelerator performance

- ✓Intended to benefit all existing sysplex distributor users
- ✓Measurements with Interactive workload
- ✓Small data sizes (100 in, 800 out)
- ✓Percentages relative to no acceleration

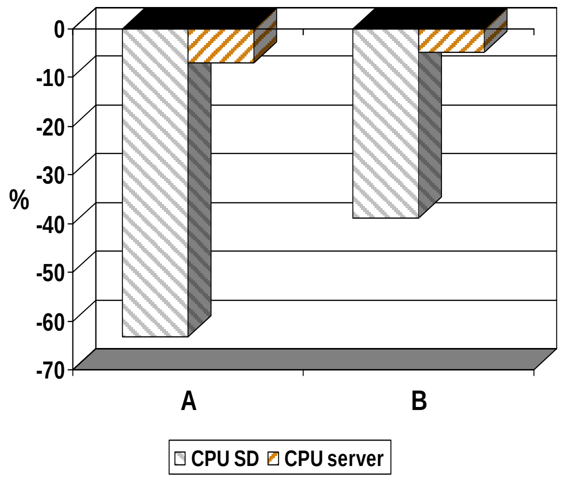

Configuration A – Three z10 LPARs with OSA Express 3 cards and HiperSockets between SD and server LPARs

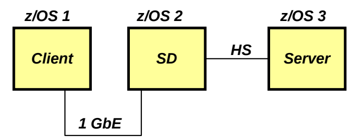

Configuration B – Three z10 LPARs with OSA Express 3 cards

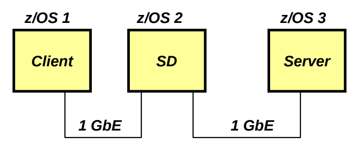

Note: The performance measurements discussed in this presentation are preliminary z/OS V1R11 Communications Server numbers and were collected using a dedicated system environment. The results obtained in other configurations or operating system environments might vary

Use of sysplex distributor accelerator is expected to provide a noticeable reduction in processor usage on the distributing stack that does the connection routing of inbound IP packets.

Netstat VCRT/-V example

```
NETSTAT VCRT DETAIL

MVS TCP/IP NETSTAT CS V1R11       TCPIP Name: TCPCS          14:16:16
Dynamic VIPA Connection Routing Table:
Dest IPaddr      DPort  Src IPaddr       SPort  DestXCF Addr
-----------      -----  ----------       -----  ------------
201.2.10.11      00021  201.1.10.85      01027  201.1.10.10
Intf:   OSAQDIOLINK
    VipaRoute: Yes      Gw: 199.100.1.1
  Accelerator: Yes
```

Netstat ROUTe/-r example

```
NETSTAT ROUTE QDIOACCEL

MVS TCP/IP NETSTAT CS V1R11      TCPIP NAME: TCPCS       09:51:02
Destination       Gateway          Interface
-----------       -------          ---------
9.67.4.1/32       0.0.0.0          OSAQDIO4
9.67.5.2/32       0.0.0.0          OSAQDIO5
9.67.20.3/32      0.0.0.0          HIPERSOCK2
```


First is an example Netstat VCRT/-V report showing the new Accelerator: Yes/No field which indicates whether the connection is accelerated.

Second is an example report from a Netstat ROUTe/-r command with the QDIOACCEL parameter. This shows each route in the accelerator routing table along with the corresponding outbound interface. These routes are a subset of the stack IP routing table and do not include sysplex distributor routes. This report is very similar to that of the existing Netstat ROUTE-/r command with the IQDIO parameter which was added for HiperSockets Accelerator.

Use VTAM tuning statistics to verify packets are accelerated

- Shows packets/bytes which are:
  - Received over an interface and then accelerated (read direction)
  - Accelerated over an interface (write direction)

```
IST1233I DEV      = 0E2A       DIR      = READ
IST1236I BYTECNTO =          0 BYTECNT  =       9628
IST1810I PKTIQDO  =          0 PKTIQD   =         14
IST1811I BYTIQDO  =          0 BYTIQD   =       9368
IST924I -------------------------------------------------------------…

IST1233I DEV      = 0E2F       DIR      = WR/1
…
IST1236I BYTECNTO =          0 BYTECNT  =       7424
IST1810I PKTIQDO  =          0 PKTIQD   =         10
IST1811I BYTIQDO  =          0 BYTIQD   =       6840
IST924I -------------------------------------------------------------
```


This is an example of VTAM® tuning stats output. PKTIQD and BYTIQD are existing fields for HiperSockets Accelerator which also apply to QDIO Accelerator.

For the read direction, PKTIQD shows the number of packets received over this interface and accelerated using QDIO Accelerator or HiperSockets Accelerator. BYTIQD shows the number of bytes received over this interface and accelerated using QDIO Accelerator or HiperSockets Accelerator. This is shown by device 0E2A in the example.

In the tuning statistics for the outbound direction (device 0E2F in the example), these counters show the number of packets and bytes which were accelerated outbound over the device.

Before you enable QDIO accelerator

- QDIO routing accelerator is IPv4 only
- Mutually exclusive with IPSECURITY
- Requires IP Forwarding to be enabled (for non-SD acceleration)
- No acceleration for:
  - Traffic which requires fragmentation in order to be forwarded
  - VIPAROUTE over HiperSockets
  - Incoming fragments for an SD connection
  - Interfaces using optimized latency mode (OLM)

These are the restrictions for QDIO Accelerator. Like HiperSockets Accelerator, the function only applies to IPv4 and cannot be enabled if IPSECURITY is enabled. If IP Forwarding is not enabled, then QDIO Accelerator only applies to sysplex distributor acceleration. If the traffic being forwarded requires fragmentation based on the MTU of the outbound route, then those packets are not accelerated but instead forwarded by the stack. For sysplex distributor, this function provides acceleration for VIPAROUTE over OSA, but not VIPAROUTE over HiperSockets. This is because of differences in how acceleration works over HiperSockets compared to OSA-Express QDIO. Also, incoming fragments for a sysplex distributor connection are not accelerated but instead sent to the stack because the SD requires the fragments to be reassembled before forwarding can occur. Similarly, the distributor stack needs to process SYN packets in order to select a target stack so SYN packets for SD are not accelerated. Finally, acceleration is not allowed to or from interfaces which are using optimized latency mode (OLM) function which is new in V1R11.

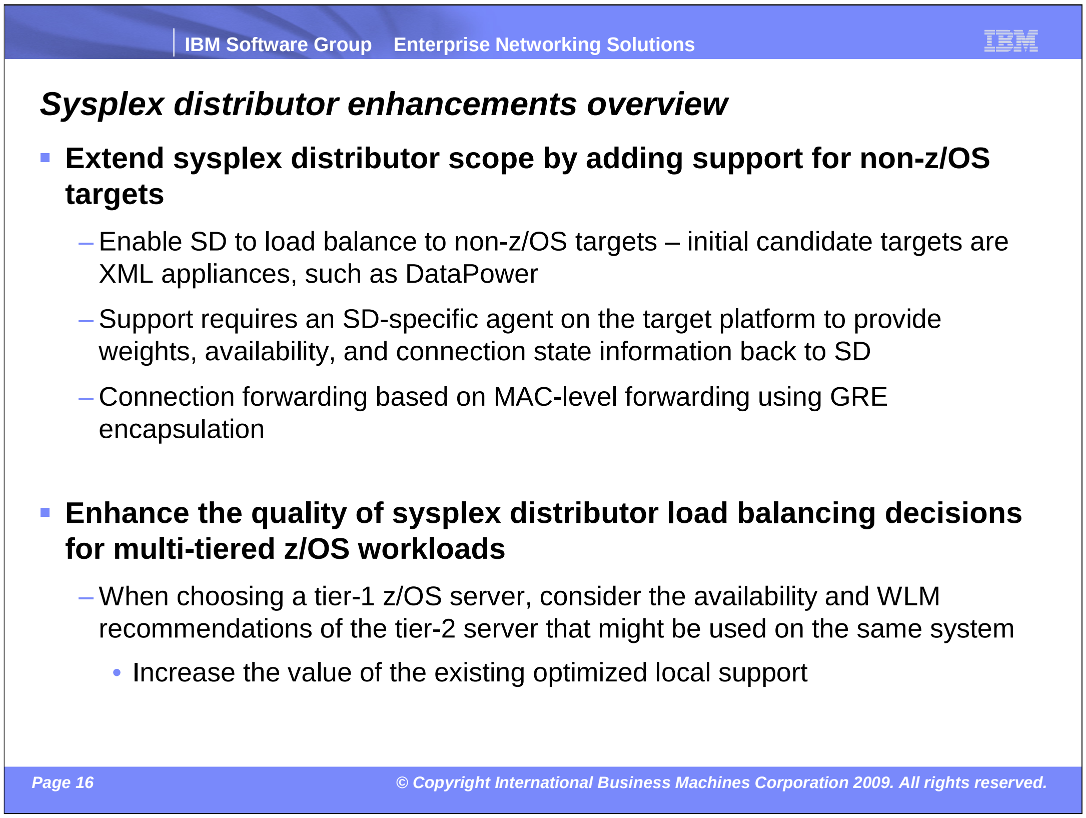

There are many enhancements to the sysplex distributor (SD) function in this release. This slide summarizes those enhancements. The next slides describe some of them in more detail.

SD in this release adds support for distributing workload (TCP connections) to target systems other than z/OS. Doing so also means that SD cannot rely on XCF communication with the target systems and it cannot rely on WLM to provide server weights for the targets. SD solves that problem by implementing support for an SD-specific agent on the target systems that uses an SD-specific protocol to communicate with SD. The agent provides metrics back to SD, which SD uses to determine availability and capacity of the target servers. In addition, the agent also provides SD with connection state information so SD can maintain its normal connection routing table information, and allow non-disruptive takeover of non-z/OS targets by a backup SD. Incoming IP packets to connections that have been distributed to a non-z/OS target are forwarded to the target using generic routing encapsulation (GRE). The initial target platform is IBM DataPower.

SD also implements accelerated connection routing which is aimed at reducing processor overhead and latency for inbound routing through the distributing node.

Finally, SD improves optimized local processing. This is done in two ways. First, SD enables the distributing stack to factor in metrics for the tier-1 servers on each target TCP/IP stack. SD also factors in metrics for the tier-2 server that is used as a target by that tier-1 server on the same TCP/IP stack.

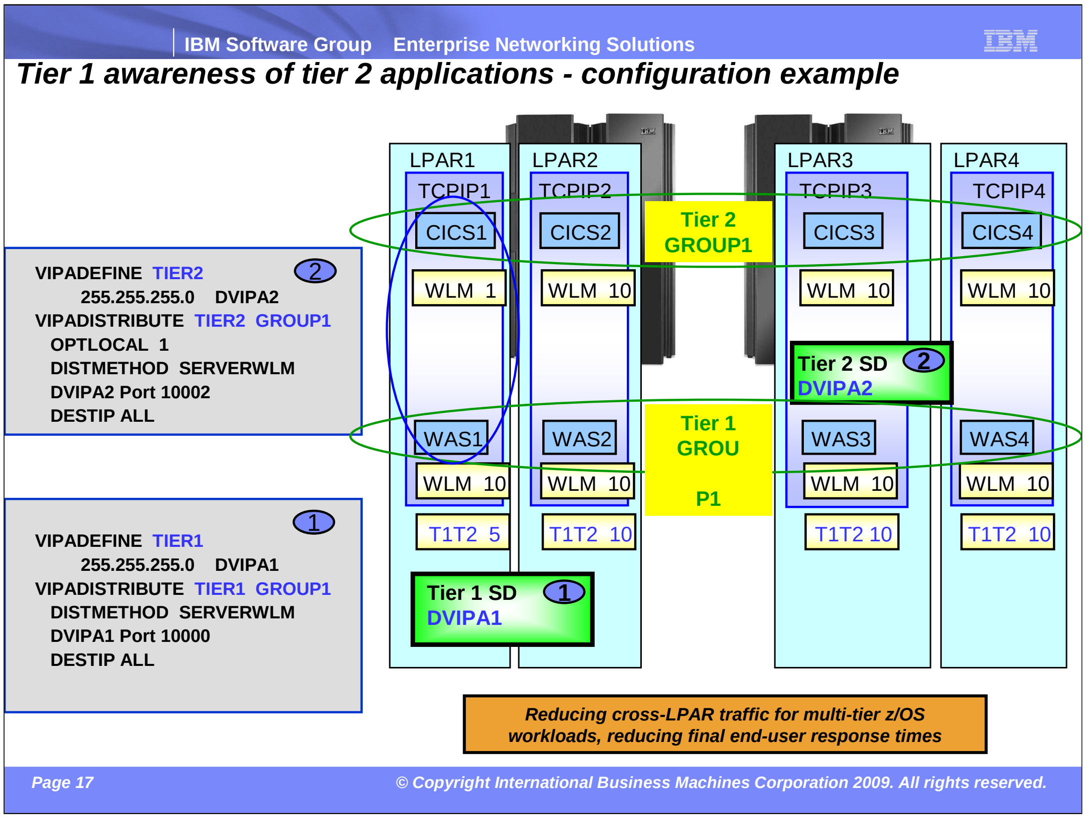

This is an example of the changes to configure a Tier 1 and Tier 2 group.

The VIPADEFINE statement has a new parameter indicating that the DVIPA is for a TIER1 or TIER2 group.

The VIPADISTRIBUTE statement has the same new tier parameter followed by the name of the group (in this case GROUP1).

Tier 1 and Tier 2 combined weights are calculated ONLY when both distribution methods use WLM weights; either BASEWLM or SERVERWLM. SERVERWLM should be used when possible since this is a recommendation based on a server's capacity for new work instead of the LPAR's capacity for new work.

Multi-tier support for z/OS targets provides an enhancement to the existing z/OS Communications Server optimized local (OPTLOCAL) support.

Sysplex distributor includes availability, health, and performance metrics of both the tier-1 z/OS server and the tier-2 z/OS server on a target LPAR when determining which LPAR the tier-1 connection should go to. This increases the likelihood of such connections gaining the performance benefits of optimized local support, such as use of fast local sockets.

In the example on this slide, a connection request comes into SD on SYS1. SD consults with WLM and the TCP/IP stacks in the z/OS Sysplex to determine availability, health, performance, and capacity of the target systems. This is done for both the HTTP tier server instances and the EIS tier server instances on each LPAR. When the chosen HTTP server connects to the tier-2 server and optimized local support is in effect, that second connection stays on the chosen LPAR. When optimized local is configured with a default value of 1, the second connection stays local if the WLM weight is greater than 0 and the server is healthy.

The support increases the likelihood that the optimized local option for the 2nd tier connection is effective.

CPCSCOPE of dynamic VIPA addresses

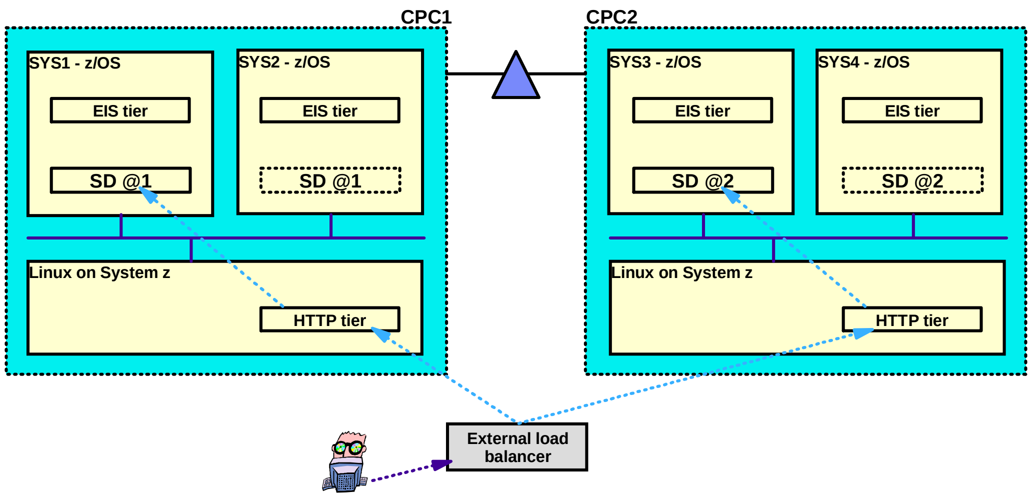

You can define dynamic VIPA addresses with a scope of a CPC in z/OS V1R11 Communications Server.

Such DVIPAs are never activated on LPARs that reside in other CPCs than where the address is currently defined.

In scenarios where Linux® on System z® acts as the first tier server, CPCSCOPE DVIPAs can be used to keep tier-2 connections and traffic on the same CPC where Linux is running. The HTTP tier on CPC1 can be configured to connect to SD@1 for the EIS tier, while the HTTP tier on CPC2 can be configured to connect to SD@2 for the EIS tier.

This topology keeps traffic between the Linux systems and the z/OS systems on HiperSockets interfaces. The goal is to improve overall response times due to reduced cross-CPC traffic in mixed Linux on System z and z/OS workload scenarios.

What is DataPower?

- DataPower can perform advanced Web services operations, contents-based routing, and transformation of requests to traditional z/OS applications
  - Security: XML/SOAP firewall capability, Web services security processing
  - XML offload: XML parsing on specialty device
  - Message transformation
    - Transform XML/SOAP to traditional z/OS application data formats
    - Interface with existing z/OS applications such as HTTP, MQ, JMS, IMS-Connect, or DB2® DRDA®
  - Contents-based routing: based on data in request (including data protected by Web services security)
    - Select proper target server, request type, format, and mode
  - Monitoring and SOA governance

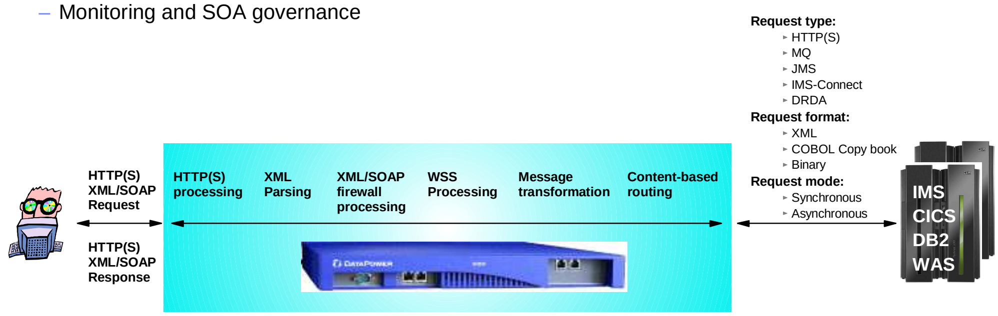

DataPower can be used to provide Web services access to other platforms, such as z/OS.

DataPower provides a full Web services protocol stack, including support for Web services security. DataPower can be customized to act as a Web services gateway to z/OS using traditional transaction interfaces to existing z/OS applications such as MQ, IMS-Connect, and DB2-DRDA. In such a setup, DataPower can provide the ability to integrate existing z/OS transactions into a Web services environment.

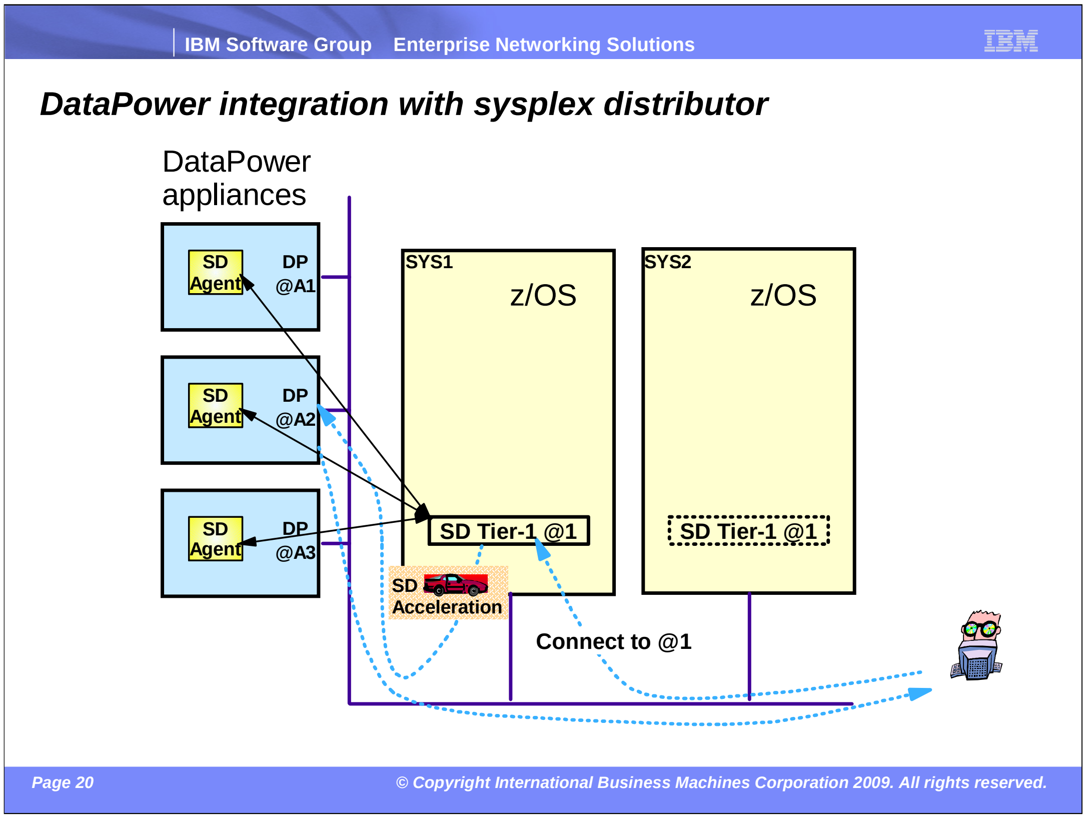

z/OS V1R11 Communications Server continues to focus on logical integration between DataPower and z/OS by supporting a DataPower feedback technology to z/OS sysplex distributor.

This support allows sysplex distributor to include DataPower appliance availability, health metrics, and performance metrics when load-balancing connections to a set of DataPower appliances.

This logical integration allows sysplex distributor to make much higher quality load balancing decisions than any other existing load balancing technology, when load balancing connections to DataPower appliances.

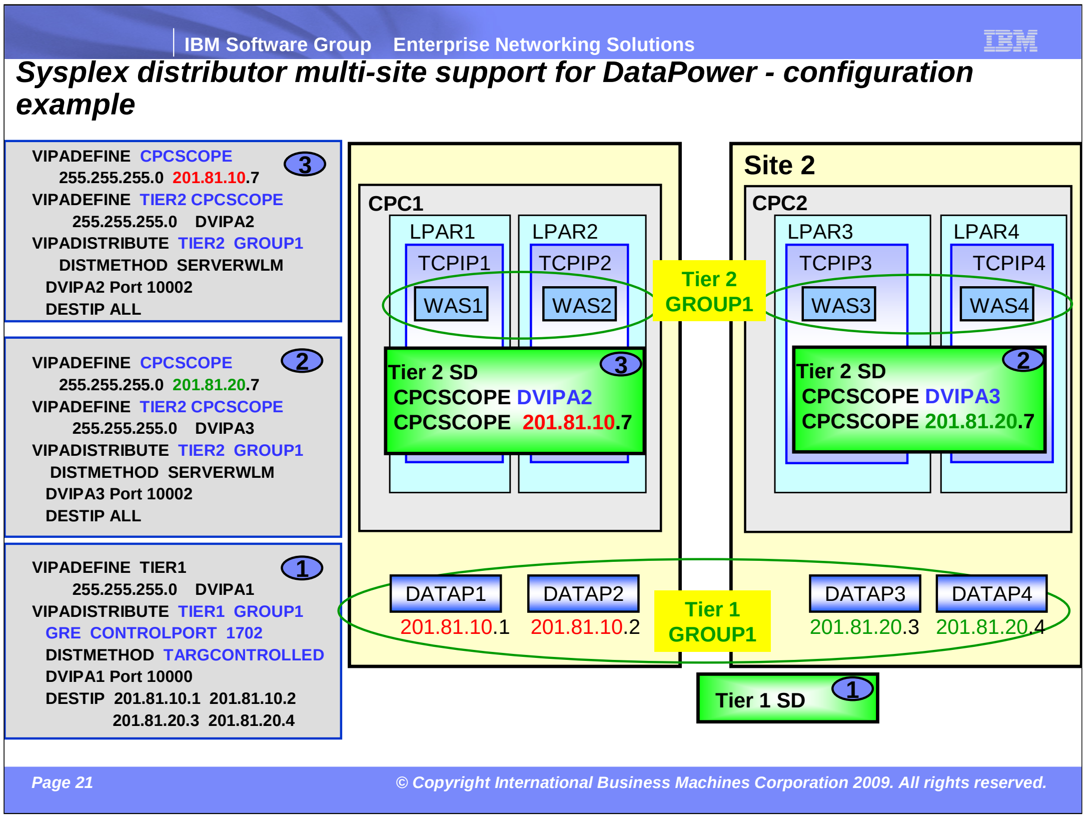

This diagram shows the full set of VIPADEFINE and VIPADISTRIBUTE statements needed on the Tier 1 distributor and the two Tier 2 distributors.

This support applies to installations with a multi-site (multi-CPC) sysplex with DataPower appliances and tier-2 z/OS applications in both sites. In such cases sysplex distributor is able to tie DataPower targets as tier-1 servers together with z/OS tier-2 servers. This can be done in such a way that cross-site communication for tier-2 connections can be eliminated.

Sysplex distributor selects a DataPower appliance based on the availability, health, and performance of both tiers, the tier-1 DataPower appliances and the tier-2 z/OS servers in that site.

If the network between the z/OS systems and DataPower appliances is implemented as a secured network, the overhead of doing connection-based encryption/decryption can be eliminated. This can further enhance the overall performance of the solution.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DataPower	DB2	DRDA	Express	HiperSockets	System z	VTAM	WebSphere	z/OS

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

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

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

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

