

Sysplex Networking Technologies

SHARE 2026 Winter (session 204)

Grant Mericle- gmericle@us.ibm.com
Paul Gartman - paul.gartman@ibm.com

Agenda

- ❑ Sysplex overview
- ❑ Workload balancing considerations
- ❑ Intra-Sysplex connectivity
- ❑ Networking Sysplex availability
- ❑ Network subplexing

Sysplex Networking Technologies and Considerations

Sysplex Overview

High Availability extends to Communications Server

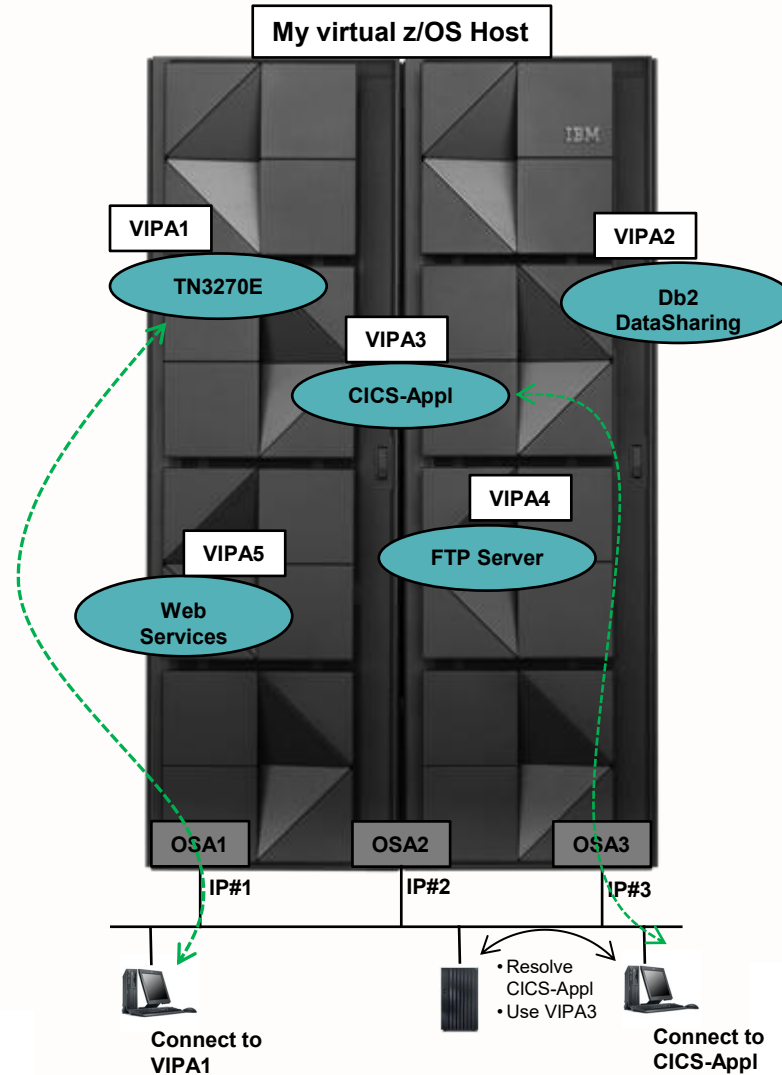
Hide the complexity of a Sysplex by providing a single system image view

□ **The goals of the Parallel Sysplex cluster environment are:**

- Application location independence
- Ability to shift application workload between LPARs
- Application single system image view from the network
- Application capacity on-demand
- Component failures do not lead to application failure

□ **To gain these benefits:**

- Carefully design redundancy of all key hardware and software components in symmetric configurations
- Supporting functions in z/OS and middleware
- Cooperation by applications
- Operational procedures



The objective is to make the Sysplex look like one large server that has a number of physical network interfaces for performance and resiliency – and that provides a set of highly available and scalable services.

Types of virtual IP addresses on z/OS

❑ Static VIPAs

- Assigned to a single TCP/IP stack
- Protects against a network interface failure as VIPA can be reached via a redundant network interface
- Used by Enterprise Extender

❑ Dynamic VIPAs

- Assigned to a single TCP/IP stack
- Unlike static VIPAs, can be dynamically moved to another TCP/IP stack within the Sysplex
- Created using VIPADEFINE statement

❑ Sysplex-specific dynamic VIPAs (i.e. Sysplex Distributor)

- Known as distributed dynamic VIPAs
- Intra-Sysplex load balancing performed across the multiple server application instances that are using the dynamic VIPA
- Created using the VIPADISTRIBUTE statement

Types of virtual IP addresses on z/OS...

❑ Application-specific dynamic VIPAs

- Dynamically created/deleted when a server application's socket binds or closes (can also use MODVIPA utility)
- Dynamic VIPA moves with server application if restarted on another TCP/IP stack within the Sysplex

```
VIPARANGE 255.255.255.0 10.10.1.0
```

❑ zCX dynamic VIPAs

- Dynamically created/deleted when the z/OS Container Extensions (zCX) address space is started/stopped
- Dynamic VIPA moves with zCX if recovered on another TCP/IP stack within the Sysplex

```
VIPARANGE 255.255.255.248 10.10.2.0 ZCX
```

And new for z/OS Container Platform!

❑ zCPA dynamic VIPAs

- Dynamically created/deleted when a z/OS Control Plane Appliance (zCPA) address spaces is started/stopped
- Subnet mask must indicate a single host IP address
- Dynamic VIPA moves with zCPA server if recovered on another TCP/IP stack within the Sysplex

```
VIPARANGE 255.255.255.255 10.20.1.100 ZCPA
```

❑ zOSCP container dynamic VIPAs

- Dynamically created/deleted when a z/OS Container Platform (zOSCP) container is started/stopped
- No control over which dynamic VIPA is assigned to a running container (next available in the subnet)

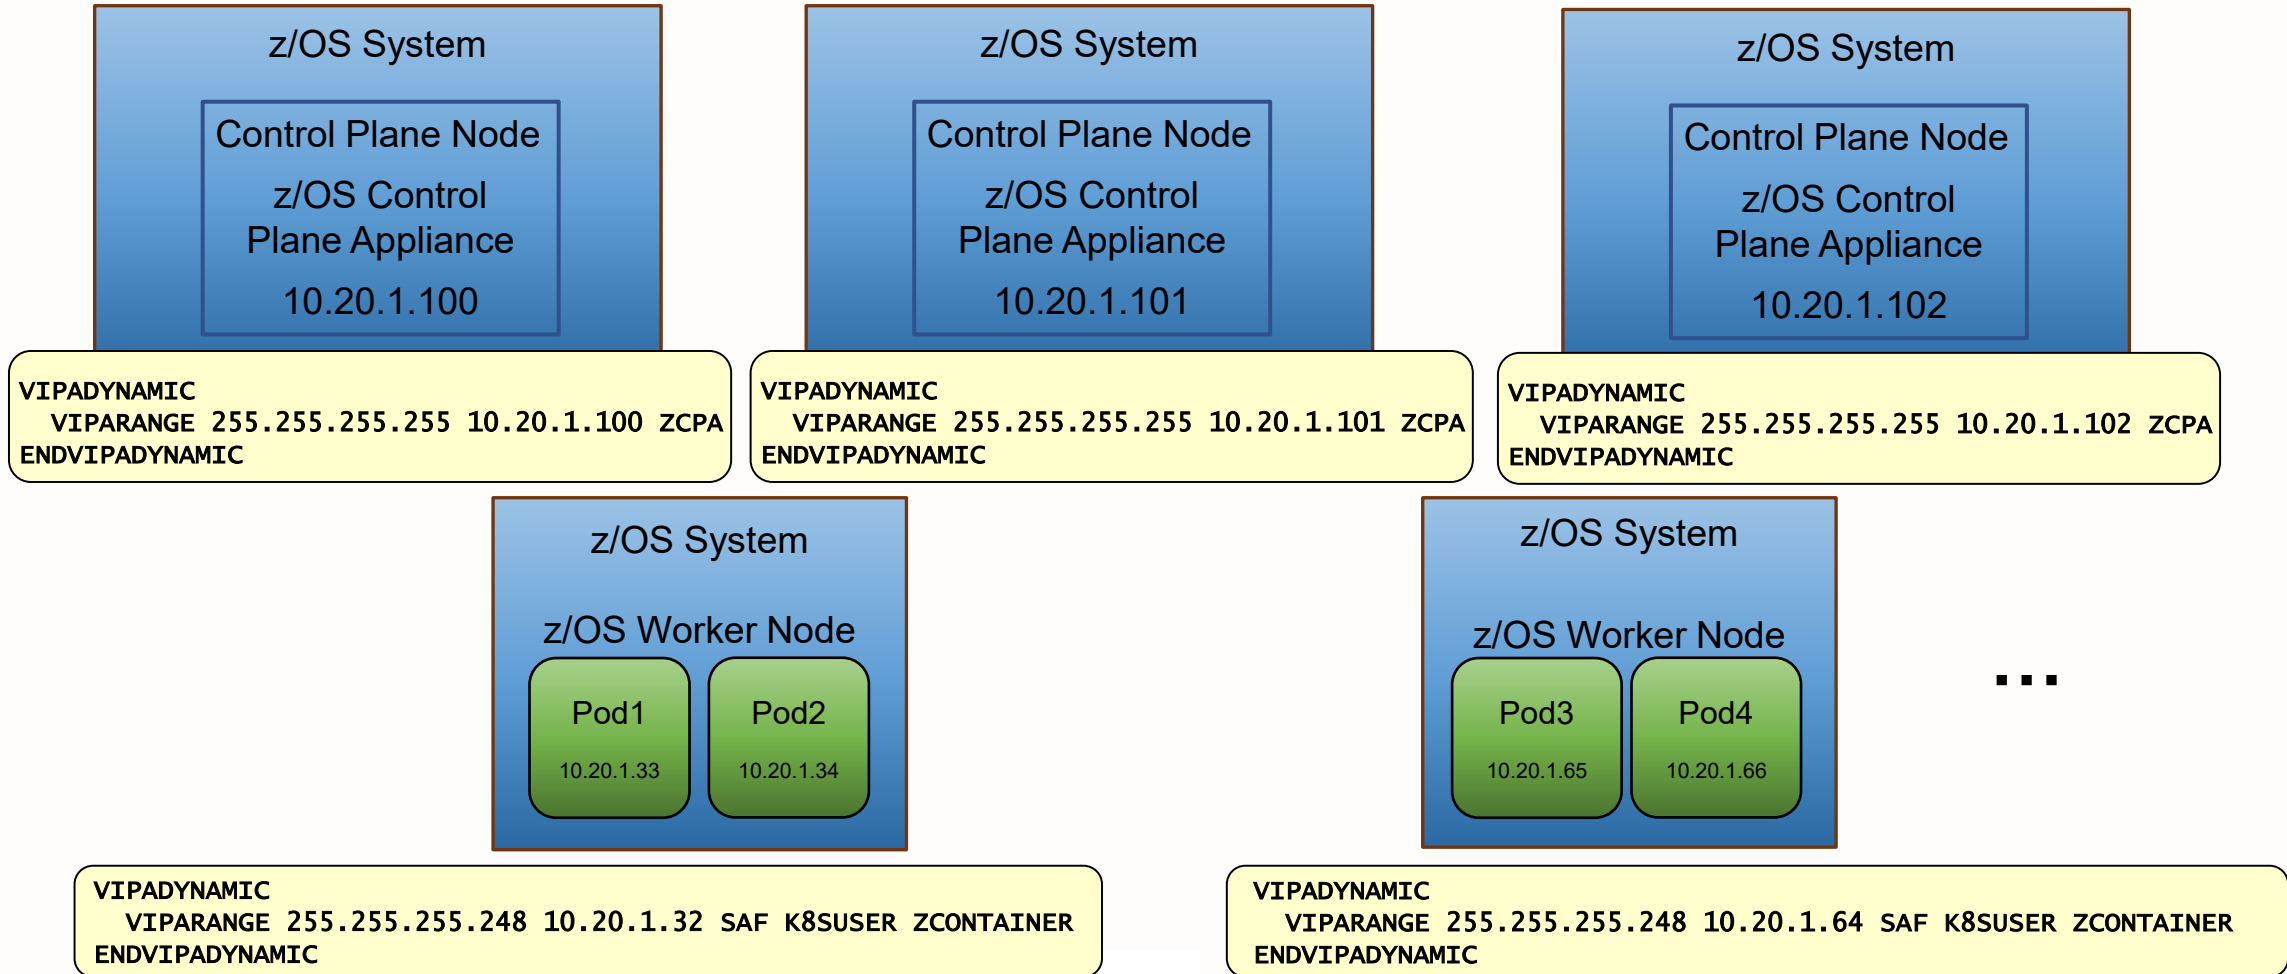
```
VIPARANGE 255.255.255.248 10.20.1.32 SAF PODMANUSER ZCONTAINER
:
VIPARANGE 255.255.255.248 10.20.1.64 SAF K8SUSER ZCONTAINER
:
```

❑ Sysplex-specific dynamic VIPAs for external targets

- Sysplex Distributor dynamic VIPAs for z/OS Control Plane Appliances (zCPA) and zCX
- Provides high availability for a Kubernetes cluster
 - Multiple zCPAs acting as Kubernetes control plane nodes

```
VIPADISTRIBUTE DISTMETHOD ROUNDROBIN EXTTARG 10.20.1.1
DESTIP 10.20.1.100 10.20.1.101 10.20.1.102
```

And new for z/OS Container Platform ...



Introducing IBM z/OS Container Platform (zOSCP)

Leverage the strengths of IBM Z and the agility of container technology



*Run containerized
z/OS UNIX
applications
natively on z/OS*

Application Modernization & Co-location

Target z/OS UNIX applications to modernize existing workloads and build new applications

Enable performance improvements through data affinity

Industry Standard Technologies

Based on industry standard open-source technology to provide a common look-and-feel and behavior

Developer Experience

Build common development skills as developers interact directly with the open-source container technology without requiring specialized Z skills

Embrace a DevOps model utilizing common industry standard tooling and automated CI/CD pipeline to improve development productivity

Leverage an isolated self-service environment for development and test

Traditional Management

Containerized workloads run as an address space in z/OS alongside traditional z/OS workloads

Traditional z/OS security, resource management and networking support is being utilized

Container Images

IBM provided images available through IBM Container Registry (ICR)

Customize and utilize a common image across the enterprise

Sysplex Networking Technologies and Considerations

Workload Balancing Considerations

What are the main objectives of network workload balancing?

□ Performance

- Workload management across a cluster of server application instances
- One server application instance on a single hardware node may not be sufficient to handle all the workload requests

□ Availability

- As long as one server application instance is up-and-running, the “service” is available
- Individual server application instances and associated hardware components may fail without impacting overall “service” availability


□ Capacity management & horizontal growth

- Transparently add/remove server application instances and/or hardware nodes to/from the pool of server applications in the cluster

□ Single System Image

- Give users one target hostname to direct requests to
- Number of and location of server instances is transparent to the user

All server application instances must be able to provide the same basic service. In a z/OS Sysplex that means the applications must be Sysplex-enabled and be able to share data across all LPARs in the Sysplex.



In order for the load balancing decision maker to meet those objectives, it must be capable of obtaining feedback dynamically, such as server application instance availability, capacity, performance, and overall health.

z/OS IP network workload balancing overview

□ Two main technologies:

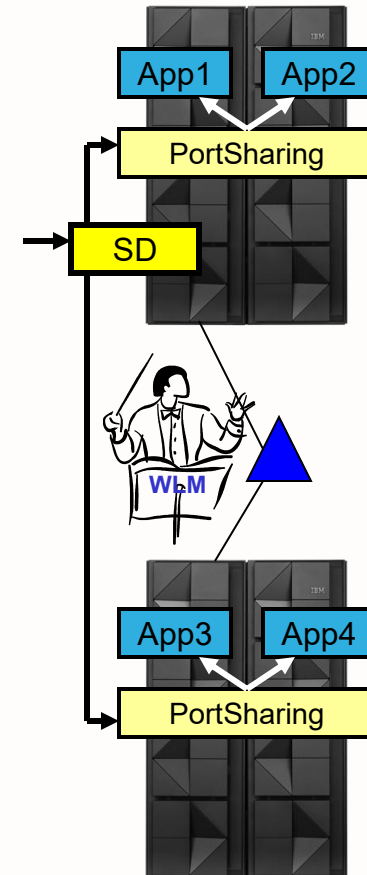
- Sysplex Distributor
- Port sharing

□ Sysplex Distributor

- Sysplex Distributor performs layer-4 load balancing
 - Makes routing decision when it sees an inbound TCP SYN segment destined for one of the distributed dynamic VIPA IP address & port combinations it load balances
- Sysplex Distributor uses MAC-level forwarding when connection routing takes place over XCF or HiperSockets links
- Sysplex Distributor uses GRE headers when connection routing takes place over any other network between the z/OS images
 - Based on definition of VIPAROUTE
- All inbound packets for a distributed connection must be routed through the Sysplex Distributor LPAR
 - Only the Sysplex Distributor TCP/IP stack advertises routing ownership for the distributed dynamic VIPA
- All outbound packets from the server application instances can take whatever route is most optimal from the server application instance node back to the client

□ Port sharing

- PORTSHARING can be used within a TCP/IP stack to distribute connections among multiple server application address spaces within that TCP/IP stack
 - SHAREPORT – TCP/IP Server Efficiency Factor (SEF) value used to perform a weighted round robin distribution to server application instances for new connections
 - SHAREPORTWLM – WLM input is used to select a server application instance for a new connection



Sysplex Distributor distribution method overview

❑ z/OS target applications without WLM recommendations

- ROUNDROBIN
 - Static distribution of incoming connections, does not account for target system capacity to absorb new workload
- WEIGHTEDACTIVE
 - Incoming connections are distributed so the available server application instances' percentage of active connections match specified weights


❑ z/OS target applications using WLM recommendations

- BASEWLM
 - Based on LPAR level CPU capacity/availability and workload importance levels
- SERVERWLM
 - Similar to BASEWLM but takes into account WLM service class and how well individual server applications are performing (i.e. meeting specified WLM goals) and how much CPU capacity is available for the specific workload being load balanced
 - Enhanced to account for WLM provided server health
 - **Generally, the recommended distribution method for Sysplex Distributor**

Sysplex Distributor built-in awareness of abnormal conditions

❑ TSR – Target Server Responsiveness

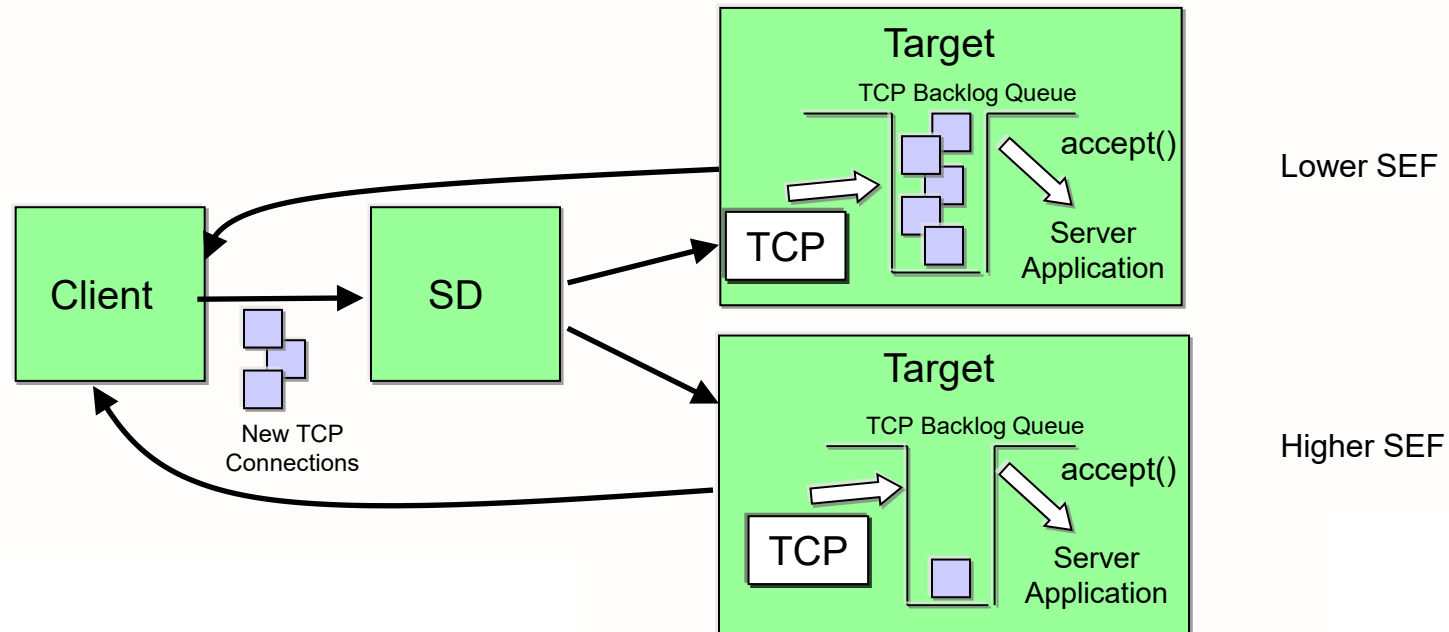
- How healthy is the target system and application from an SD perspective? A percentage, 0-100%
- Comprised of several individual health metrics:
 - **TCSR** – Target Connectivity Success Rate (percentage: 100 is good, 0 is bad)
 - Are connections being sent to the Target System making it there?



- **CER** – Connection Establishment Rate (percentage: 100 is good, 0 is bad)
 - Is connectivity between the target system and the client ok?
 - By monitoring TCP connection establishment (completed 3 way handshake between client and server), can detect whether a connectivity issue exists
 - Note: CER no longer part of TSR directly but is included directly in SEF and continues to be calculated and reported separately

❑ TSR – Target Server Responsiveness (continued)

- **SEF** – Server Efficiency Fraction (percentage: 100 is good, 0 is bad)
 - Is the target server application keeping up with new connections in its backlog queue?
 - Is the new connection arrival rate higher than the application accept rate? (Is backlog growing over time?)
 - How many connections in the TCP backlog queue? How close to maximum backlog queue depth? Were any new connections dropped because the backlog queue depth was exceeded?
 - Is the server application hung? (Is server application not accepting any connections?)
 - Is the number of half-open connections on the backlog queue growing? (Similar to CER)



Middleware/Application issues (“Storm Drain Problem”)

- ❑ TCP/IP and WLM are not aware of all problems experienced by load balancing targets (middleware/applications)
 - The server application needs a resource such as a database, but the resource is unavailable
 - The server application is failing most of the transactions routed to it because of internal processing problems
 - The server application acts as a transaction router for other back-end applications on other system(s), but the path to the back-end application is unavailable
- ❑ In each of these scenarios, the server may appear to be completing the transactions quickly (using little CPU capacity) when they are actually not completing successfully
- ❑ This is sometimes referred to as the *Storm Drain Problem*
 - The server application is favored by WLM since it is using very little CPU capacity
 - As workloads increase, the server application is favored more and more over other server applications
 - All this work goes "down the drain"

Improving WLM awareness of Application Health

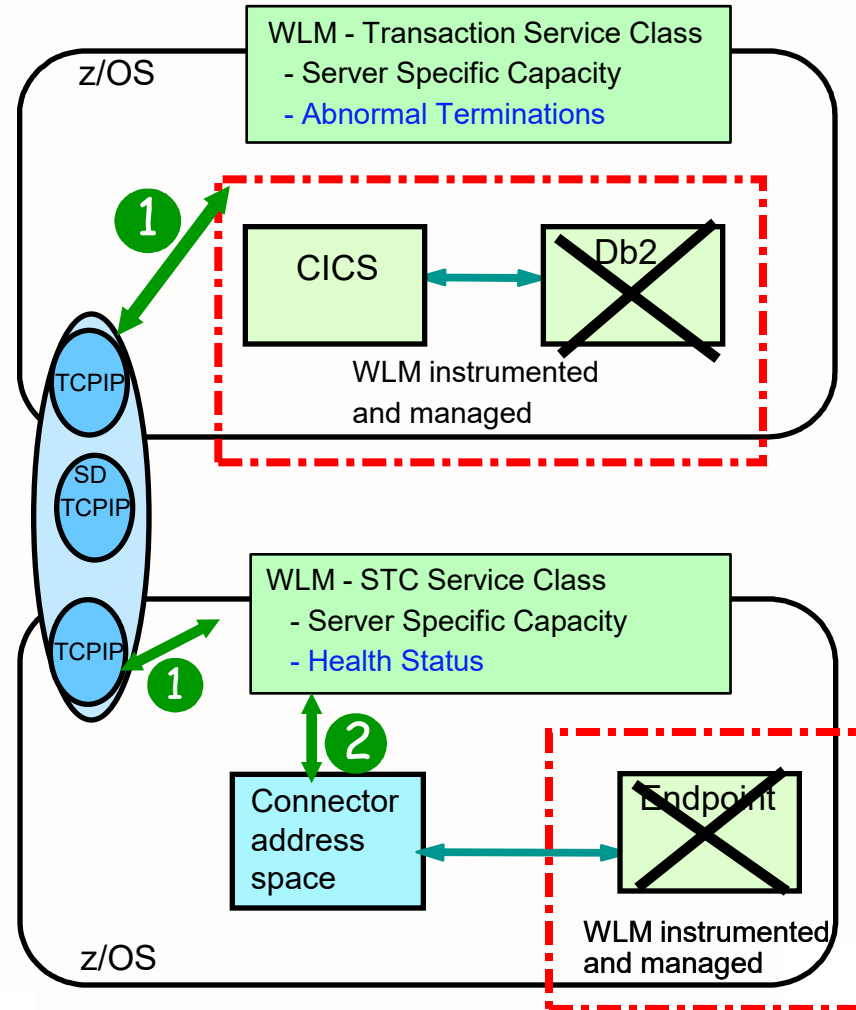
Server Scenarios

1 IWM4SRSC WLM Service

- Used by Sysplex Distributor to obtain WLM recommendations
- Abnormal Termination information: Reported by 1st-tier server application when transactions can not complete because back-end resources are not available
 - WLM uses this information to reduce the recommendation for unhealthy server applications

2 IWM4HLTH WLM Service

- Allows address spaces which are not instrumented with WLM to set a health status which is also returned by IWM4SRSC
- The ServerWLM recommendations are reduced when the health is <100%
- Exploited by CICS Transaction Gateway, Db2, LDAP, IMS, and WebSphere Application Server



Using Netstat VDPT Display to monitor Sysplex Distributor

Target Server Responsiveness (TSR) and subcomponents (applied to WLM weight)

ActConn: Active number of connections to this target at this time. Note connections in Timewait or Finwait states also show up here.

WLM Information: Raw Weights, Proportional Weights, Abnormal Transaction Rate and Middleware reported health

```

NETSTAT VDPT DETAIL
MVS TCP/IP NETSTAT CS V2R3          TCPIP Name: TCPCS          15:35:26
Dynamic VIPA Distribution Port Table for TCP/IP Stacks:

Dest IPaddr      DPort DestXCF Addr      Rdy TotalConn  WLM  TSR  Flg
-----
192.168.10.14    00244 192.168.10.16     001 0002304546 12  080  1
  DistMethod: ServerWLM
  TCSR: 100 CER: 095 SEF: 080
  Weight: 58
  Raw          CP: 58 zAAP: 00 zIIP: 5
  Proportional CP: 04 zAAP: 00 zIIP: 5
  Abnorm: 0000 Health: 100
  ActConn: 0000000101
  QosPlcAct: *DEFAULT*
192.168.10.14    00244 192.168.10.17     001 0001543454 10  100  1
  DistMethod: ServerWLM
  TCSR: 100 CER: 100 SEF: 100
  Weight: 40
  Raw          CP: 40 zAAP: 00 zIIP: 40
  Proportional CP: 06 zAAP: 00 zIIP: 34
  Abnorm: 0000 Health: 100
  ActConn: 0000000030
  QosPlcAct: *DEFAULT*
  
```

WLM Weight after all adjustments TSR, Subsystem Health, Abnormal Connection Rate. Final value divided by 4 to end up with 0-16 value range

TotalConn: Total number of connections since DVIPA was activated – ever increasing value

SERVERWLM method: What is displaceable LPAR capacity?

- ❑ LPAR capacity that is currently being used for less important work than what Sysplex Distributor wants to send to the LPAR
- ❑ Controlled with ILWEIGHTING on VIPADISTRIBUTE statement
- ❑ An example:
 - New work will run at Importance level 2
 - Which LPAR is best (LPAR 1 or LPAR 2)?
 - They both have 500 service units of displaceable work
 - Typically, both would be considered equally good targets
 - Sysplex Distributor can take importance level of displaceable workload into consideration
 - LPAR2 will be preferred since the importance level of the work being displaced is lower than the work that would be displaced on LPAR1

New workload at IL=2 (can displace IL=3 to IL=7 workload)

IL 0: High importance
IL 7: Low importance

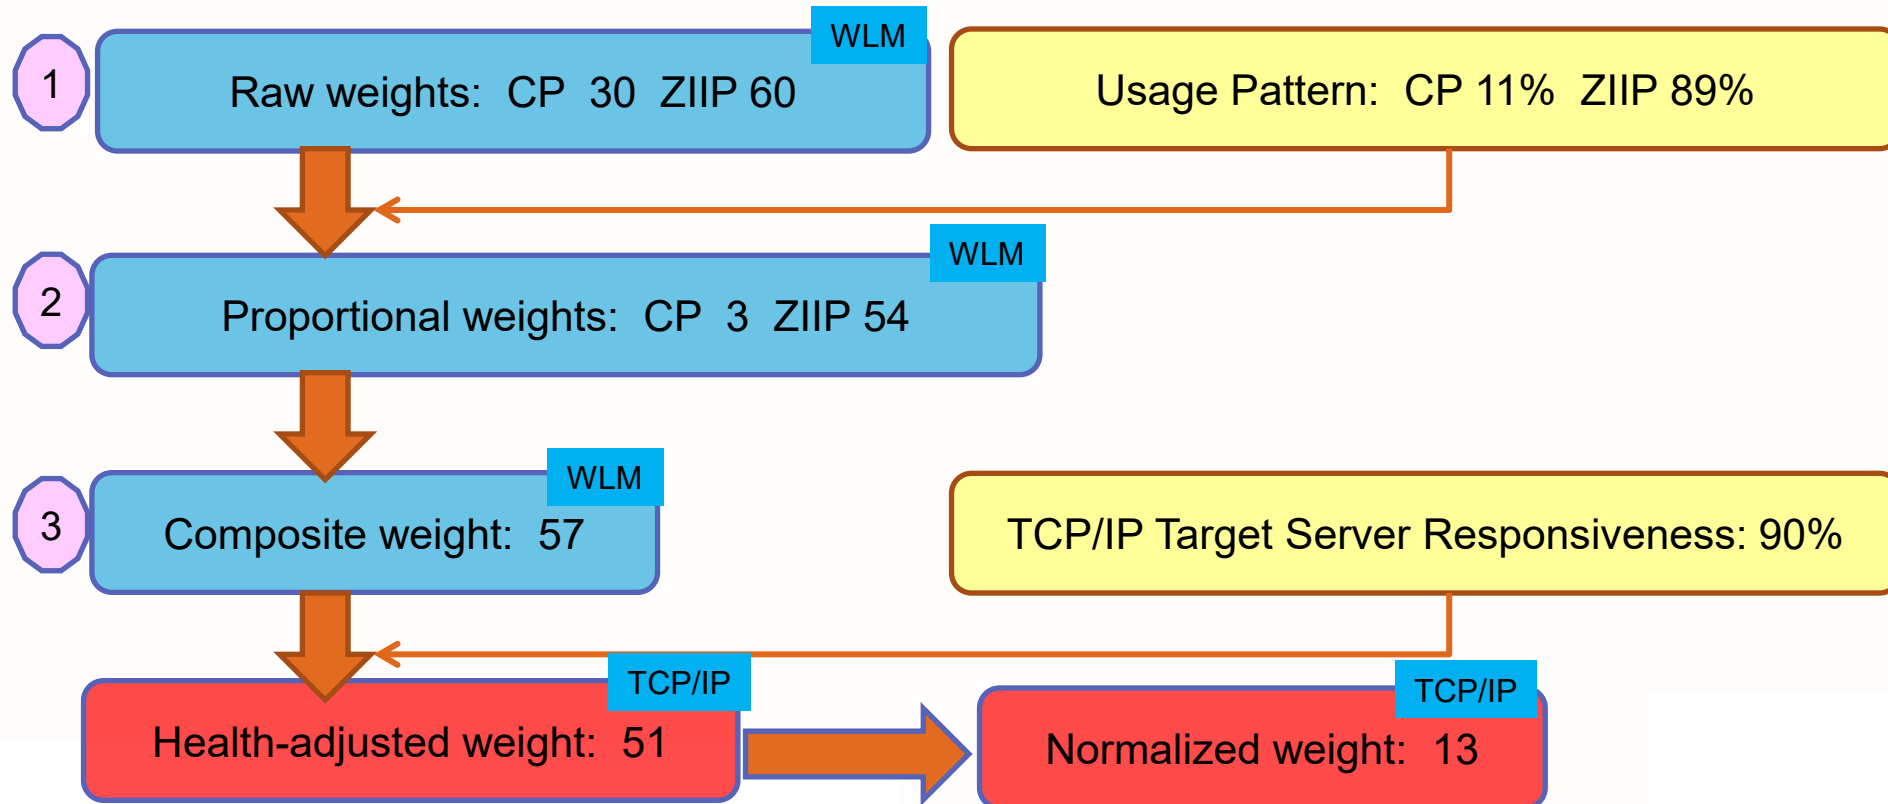
LPAR1		LPAR2	
I	SUs	I	SUs
L		L	
0	0	0	0
1	0	1	0
2	0	2	0
3	500	3	0
4	0	4	0
5	0	5	500
6	0	6	0
7	0	7	0

SERVERWLM method: specialty processors - overview

When using WLM server-specific weights.

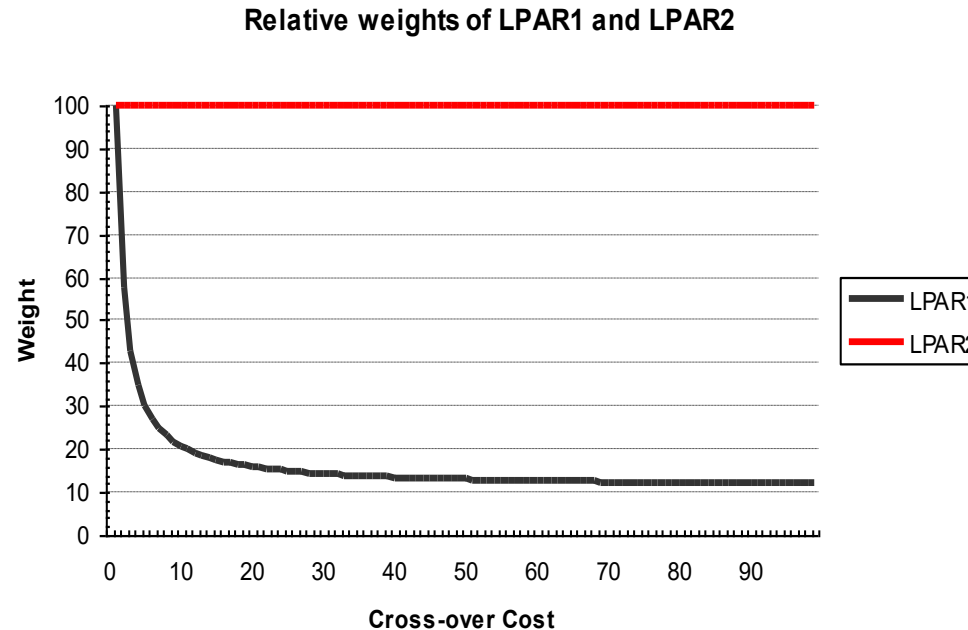
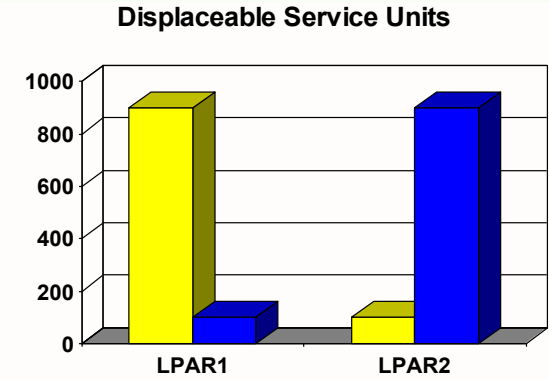

– WLM returns

- The raw CP and zIIP system weights
- Proportional weights – raw weight modified by actual server usage pattern as observed by WLM
- Composite weight



SERVERWLM method: zIIP cross-over to CP

- ❑ Controlled with PROCXCOST
- ❑ Application designed to use 10% CP and 90% zIIP
- ❑ LPAR1 and LPAR2 are targets
 - LPAR1 has 900 CP SUs and 100 zIIP SUs that can be displaced
 - LPAR2 has 100 CP SUs and 900 zIIP SUs that can be displaced
- ❑ Without a cross-over cost, the two targets are equally good to receive new workload
- ❑ As a cross-over cost is applied, LPAR1 is less attractive than LPAR2
- ❑ Cross-over cost can be set to a value between 1 and 100
 - 1: no penalty for cross-over
 - 100: maximum penalty for cross-over




Sysplex Distributor Polling Intervals

- ❑ By default, Sysplex Distributor queries WLM every 60 seconds
 - During this same interval, Sysplex Distributor also calculates health metrics such as SEF, TSR, etc.
 - In environments where changes in workload conditions can occur very quickly, a smaller polling interval is often more desirable
 - Allows Sysplex Distributor to have more current WLM recommendations and health metrics
 - The polling interval can be set via the ***SYSPLEXWLMPOLL*** keyword on ***GLOBALCONFIG*** statement
 - Specified in seconds (1-180)
 - A value of 10 seconds is recommended for obtaining the most recent WLM recommendations
 - ***Should be configured on all*** TCP/IP stacks in the Sysplex environment (distributor and targets)
 - Each TCP/IP stack implements its own timer

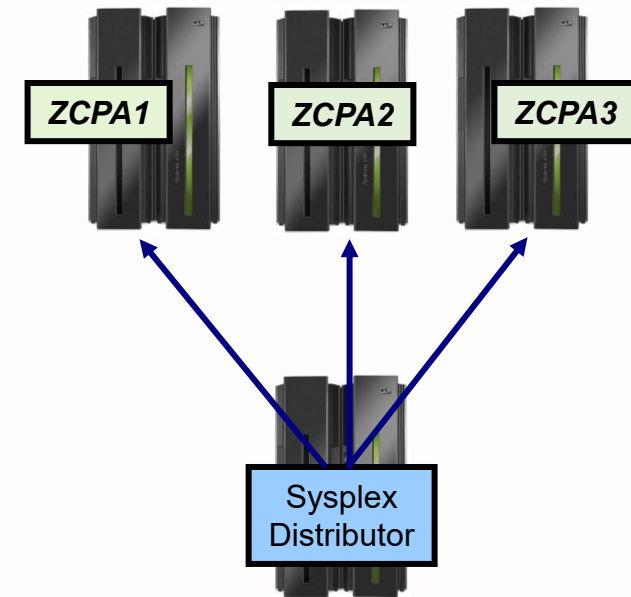
Sysplex Distributor: Hot Standby support

- Data sharing provides for highly scalable and highly available configuration
 - Additional server application instances and LPARs can be cloned and added to increase capacity and improve performance
 - But what if the work can comfortably fit within a single LPAR?
 - Data sharing becomes primarily an availability feature
 - A Hot Standby configuration would allow all work to be routed to a single LPAR
 - Minimizing data sharing overhead!
 - While retaining high availability!



Sysplex Distributor: External Target support

- Utilize Sysplex Distributor technology to load balance across non-z/OS targets
 - DESTIP lists the set of external targets to monitor and distribute connections
 - z/OS Control Plane Appliances (zCPA)
 - z/OS Container Extensions (zCX)
 - Ports are dynamically discovered as server applications or Docker containers are started within a zCPA or zCX
 - Dynamically discover Kubernetes NodePort services as they are created on control plane nodes for workload deployments

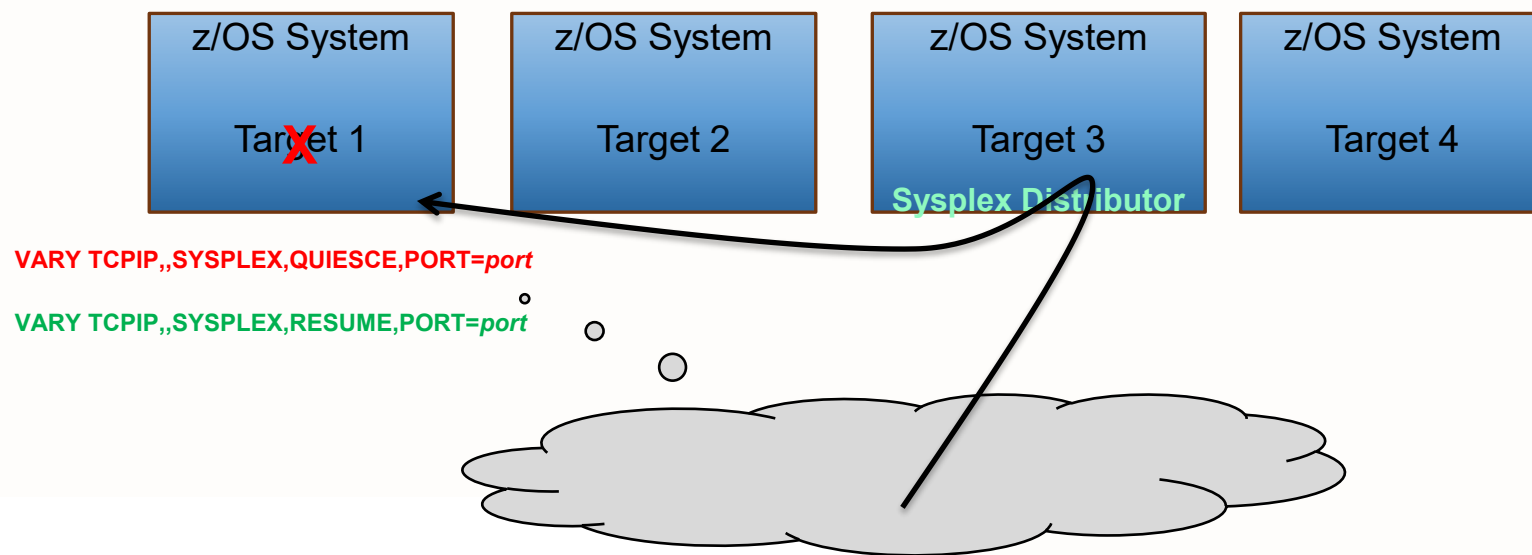


```
VIPADISTRIBUTE DISTMETHOD ROUNDROBIN EXTTARG ipaddress  
DESTIP ZCPA1 ZCPA2 ZCPA3
```

Quiescing a Sysplex Distributor target

On a target stack:

- A quiesce command to stop sysplex distribution of new connections to a specified port or application
 - Informs distributing stack that this application on this target stack should not receive new connections – existing connections are unaffected
 - Useful for draining workload away from target before taking it down
- A resume command to reactivate the port or application for sysplex distribution

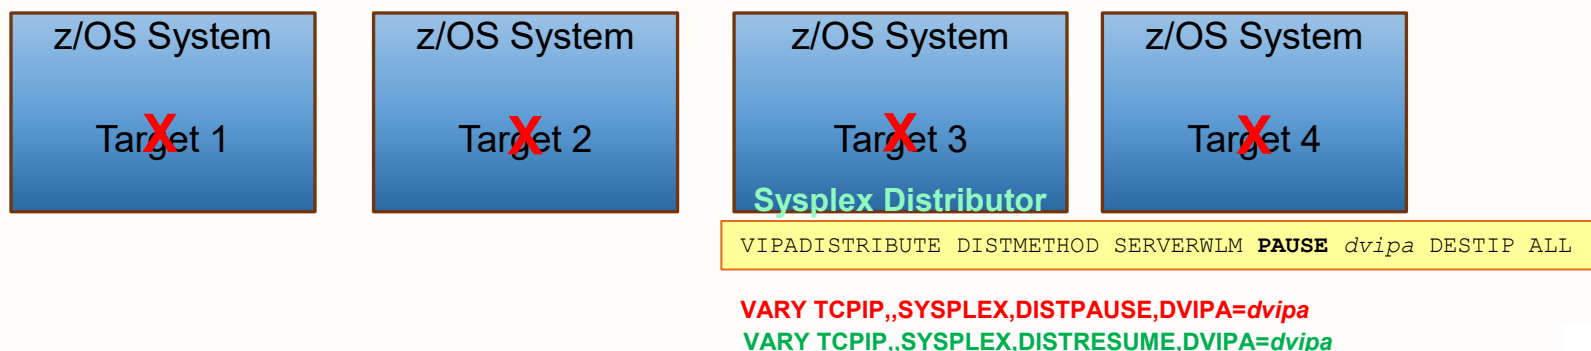


Persistent pausing all Sysplex Distributor targets

On Sysplex Distributor stack:

New in 3.1!

- A VIPADISTRIBUTE statement to define the dynamic VIPA in a paused state
- A pause command to stop sysplex distribution of new connections to a specified port or all ports across all targets
 - Informs distributing stack that this dynamic VIPA is not available to receive new connections for a specific port or all ports for any targets – existing connections are unaffected
 - Persists across application and target stack restarts
 - Useful for draining all work away from a production sysplex prior to restarting workload on recovery site – persists across a restart of production sysplex
- A resume command to reactivate the dynamic VIPA and optional port for sysplex distribution across all targets



Increasing PORT limit on VIPADISTRIBUTE statement

- Used to specify one or more individual ports, ranges of ports, or a combination of individual ports and ranges
 - Servers that bind to the specified DVIPA, 0.0.0.0, or :: (IPv6 unspecified address) and one of the specified ports, cause the target stack to become eligible to receive connection requests
- The maximum number of ports that is specified, including all individual ports and all ports within ranges, has been raised to 256
 - Available on z/OS 3.1 via APAR PH63320
- Care should be taken to not combine large numbers of both ports (via the PORT keyword) and target TCP/IP stacks (via DESTIP statement)
 - Could result in very large XCF messages that could impact the performance of Sysplex Distributor
 - Recommendation is to use DESTIP ALL when configuring more than 64 ports on a VIPADISTRIBUTE statement

Sysplex Networking Technologies and Considerations

Intra-Sysplex Connectivity

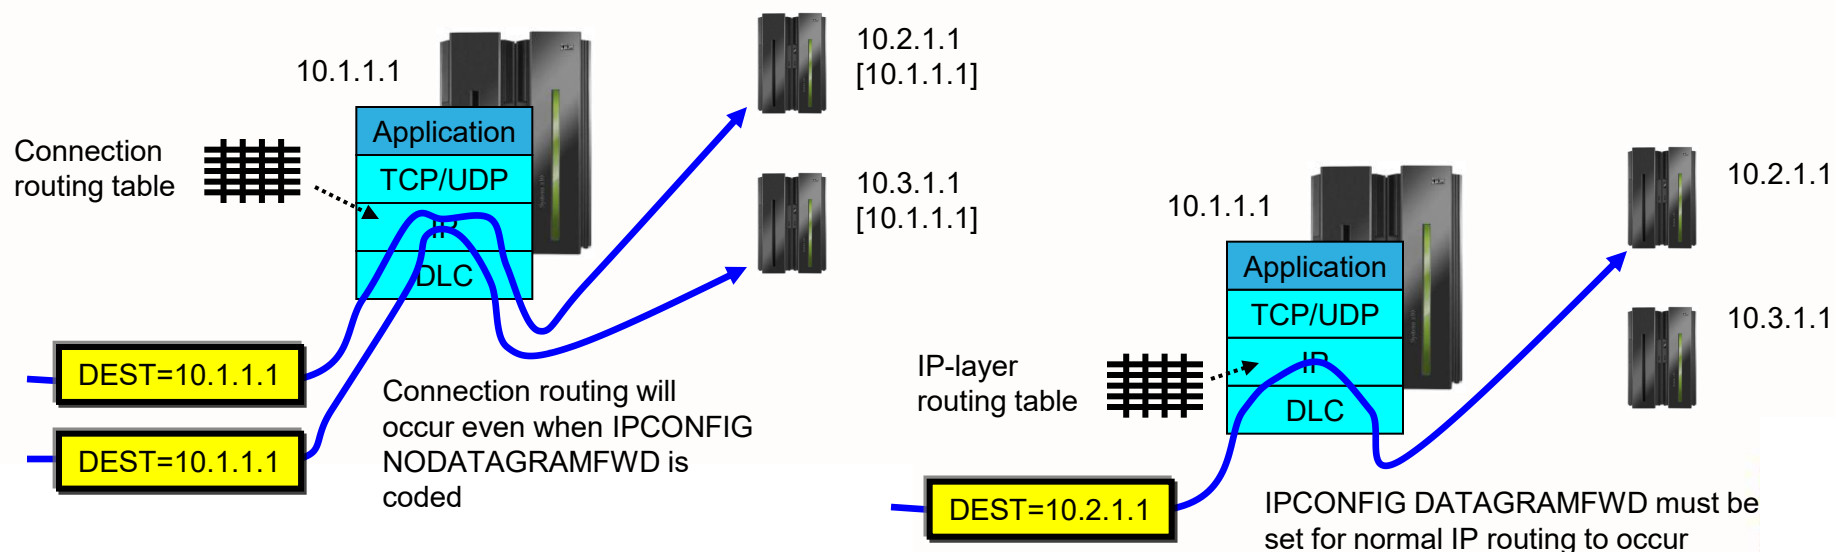
Two types of intra-Sysplex routing

❑ Connection routing

- IP routing decision based upon connection routing table (CRT), destination IP address, and specific connection (4-tuple)
 - Packets to the same IP address, but belonging to two different connections, may go to two different targets
- Used by Sysplex Distributor and movable Dynamic VIPA support
- Not subject to the setting of IPCONFIG DATAGRAMFWD/NODATAGRAMFWD

❑ Normal IP routing

- IP routing decision based upon IP-layer routing table and destination IP address
 - All packets to the same IP address are treated the same
- Forwarding to z/OS TCP/IP stacks through another z/OS TCP/IP stack
- Subject to the setting of the IPCONFIG DATAGRAMFWD/NODATAGRAMFWD option



Role of XCF, HiperSockets, and external interfaces in Sysplex

❑ XCF

- All XCF control messaging between stacks in a Sysplex always go via XCF messages
- DynamicXCF Sysplex Distributor connection routing (but only if no VIPAROUTE defined)
- If configured for static or dynamic routing, normal IP routing between LPARs – not recommended

❑ HiperSockets

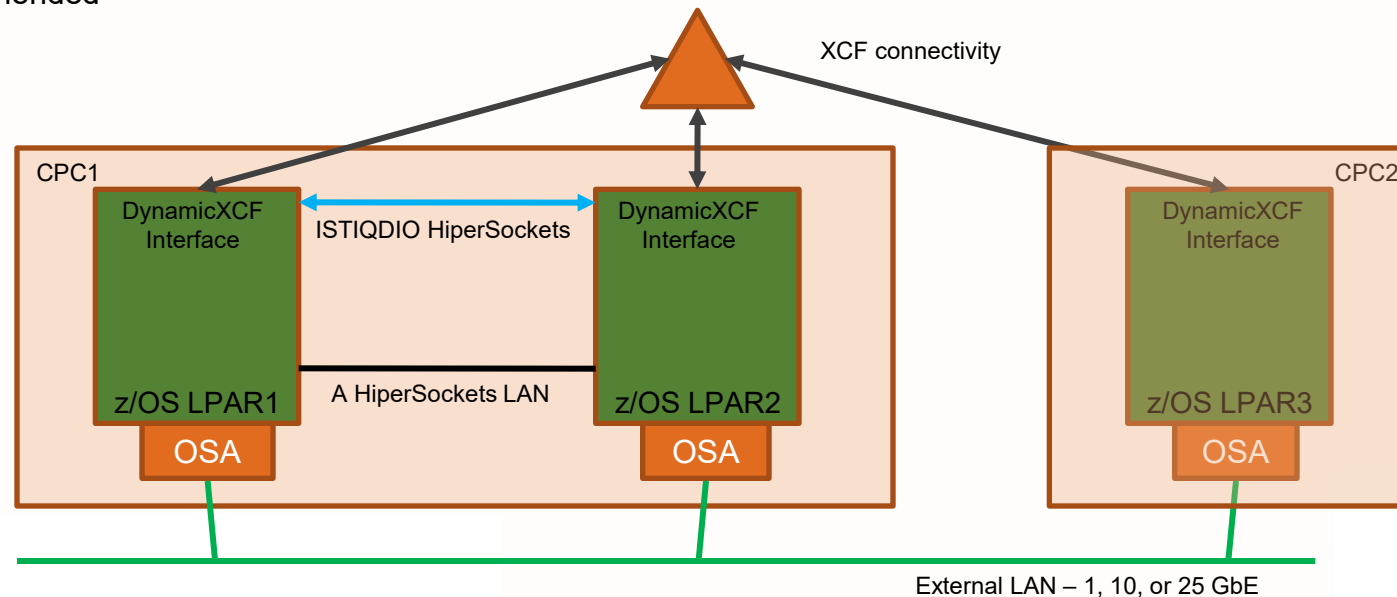
- If ISTIQDIO is defined (in VTAM), DynamicXCF Sysplex Distributor connection routing between LPARs on same CPC uses this interface instead of XCF
- Considered part of the DynamicXCF network interface – no separate DEVICE/LINK or INTERFACE definitions required

❑ External interface or a manually defined HiperSockets


- If VIPAROUTE defined, then used to determine Sysplex Distributor connection routing between LPARs
 - VIPAROUTE is generally recommended
- Normal IP routing

Only define DynamicXCF interfaces as OSPF interfaces if you want to be able to use XCF as a last-resort connectivity between z/OS TCP/IP stacks.

If you have “enough” redundancy built into your OSA adapters and data center switches, you may not need to ever use XCF for normal IP routing.



VIPARROUTE vs XCF – Performance Comparison

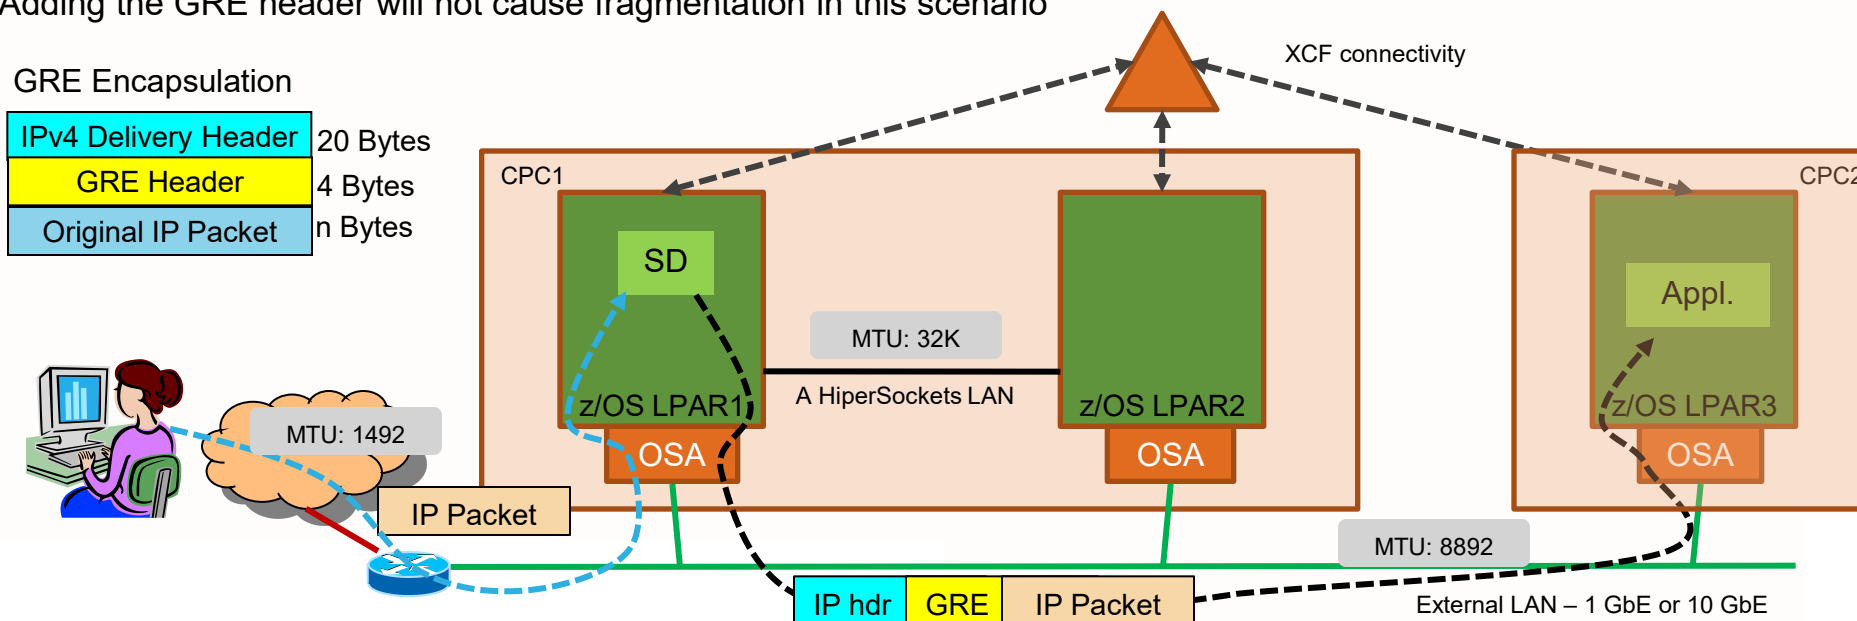


Significant throughput benefits for both request/response (RR) and streaming data (STR) patterns
Significant reduction in networking CPU overhead on both Sysplex Distributor and Target LPARs

Note: The performance measurements were collected in IBM internal tests using a dedicated system environment. The results obtained in other configurations or operating system environment significantly depending upon environments used. Therefore, no assurance can be given, and there is no guarantee that an individual user will achieve performance or throughput improvements to the results stated here.

VIPAROUTE and MTU size considerations

- ❑ When VIPAROUTE is used, the distributing TCP/IP stack adds a GRE header to the original IP packet before forwarding to the target TCP/IP stack
- ❑ Two ways to avoid fragmentation between distributing and target stacks:
 - Have clients use path MTU discovery
 - z/OS will factor in the GRE header size (24 bytes) when responding with next-hop MTU size
 - Not always possible to control distributed node's settings from the data center
 - Use jumbo-frames on the data center network
 - The access network will typically be limited to Ethernet MTU size (1492 bytes), while the data center network will be able to use jumbo frame MTU size (8892 bytes)
 - Adding the GRE header will not cause fragmentation in this scenario



VIPARROUTE fragmentation avoidance

- ❑ VIPARROUTE is used by many to offload Sysplex Distributor forwarded traffic from XCF links
 - When used in combination with QDIO Accelerator for Sysplex Distributor, can result in dramatically reduced overhead for Sysplex Distributor forwarding



- ❑ Fragmentation is still a concern
 - Resulting from the extra 24 bytes that are needed for the GRE header
 - Path MTU Discovery helps but doesn't solve the issue in some environments (where ICMP messages cannot flow across firewalls)
 - Jumbo frames can cause fragmentation for outbound traffic routed through standard network (1492 MTU)
 - Fragmentation can cause significant performance degradation

VIPAROUTE fragmentation avoidance

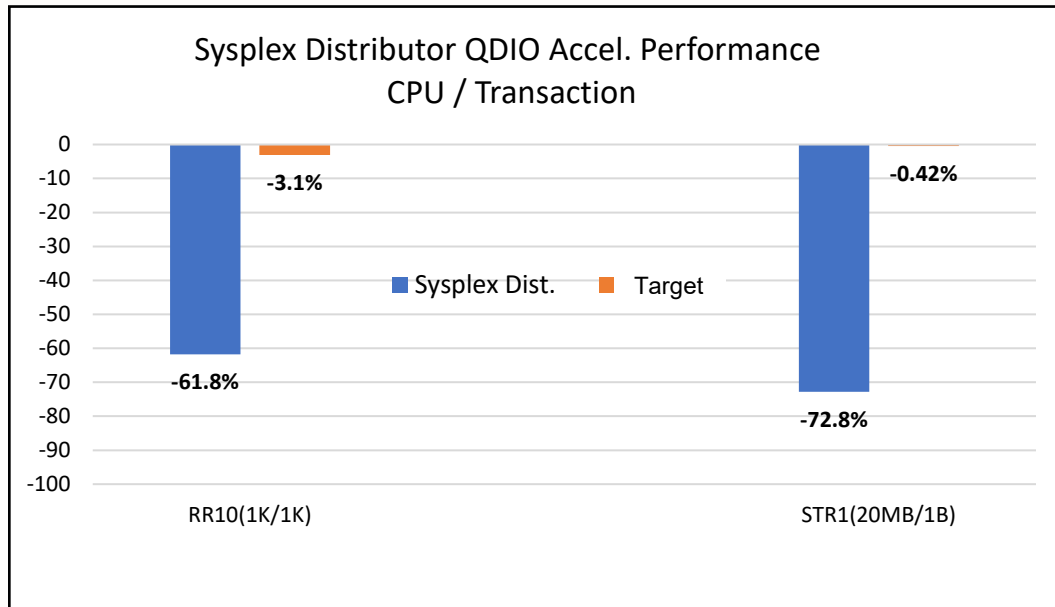
- ❑ An autonomic option that will automatically reduce the MSS (Maximum Segment Size) of a distributed connection by the length of the GRE header
 - This allows client TCP/IP stacks to build packets that account for the 24 bytes of the GRE header to be added without any fragmentation being required
 - **ADJUSTVIPAMSS** on GLOBALCONFIG
 - Defaults to *AUTO* - enables adjusted MSS
 - On target TCP/IP stacks when VIPAROUTE is being used
 - On Sysplex Distributor TCP/IP stack if it is also a target and VIPAROUTE is defined
 - Option *ALL* – Enables adjusted MSS for all connections using a dynamic VIPA (distributed or not)
 - Option *NONE* - If you are already exploiting VIPAROUTE and know that there's no fragmentation possible in your environment you can disable this function
 - *Note:* This is a TCP/IP stack specific option - if the default is not taken, then it must be configured on all TCP/IP stacks in the sysplex

QDIO routing accelerator

- ❑ **Provides fast path IP forwarding for:**
 - Inbound OSA → Outbound OSA or HiperSockets
 - Inbound HiperSockets → Outbound OSA or HiperSockets
- ❑ **Adds Sysplex Distributor acceleration**
 - Inbound packets over HiperSockets or OSA
 - When Sysplex Distributor forwards to the target stack using either:
 - Dynamic XCF connectivity over HiperSockets
 - VIPAROUTE over OSA
- ❑ **Improves performance and reduces CP utilization for such workloads**
- ❑ **Restrictions:**
 - QDIO routing accelerator is IPv4 only
 - Requires IP Forwarding to be enabled (for non-Sysplex Distributor acceleration)
 - No acceleration for:
 - Traffic which requires fragmentation in order to be forwarded
 - VIPAROUTE over HiperSockets
 - Incoming fragments for a Sysplex Distributor connection
 - Interfaces using optimized latency mode (OLM)



Sysplex Distributor Accelerator Performance



✓ Intended to benefit all existing Sysplex Distributor users

✓ Request/Response data pattern (1K request, 1K response, 10 concurrent sessions) – RR10

✓ Streaming data pattern (1/20M – 1 byte in, 20MB response, 20M/1 – 20MB in, 1 byte response)

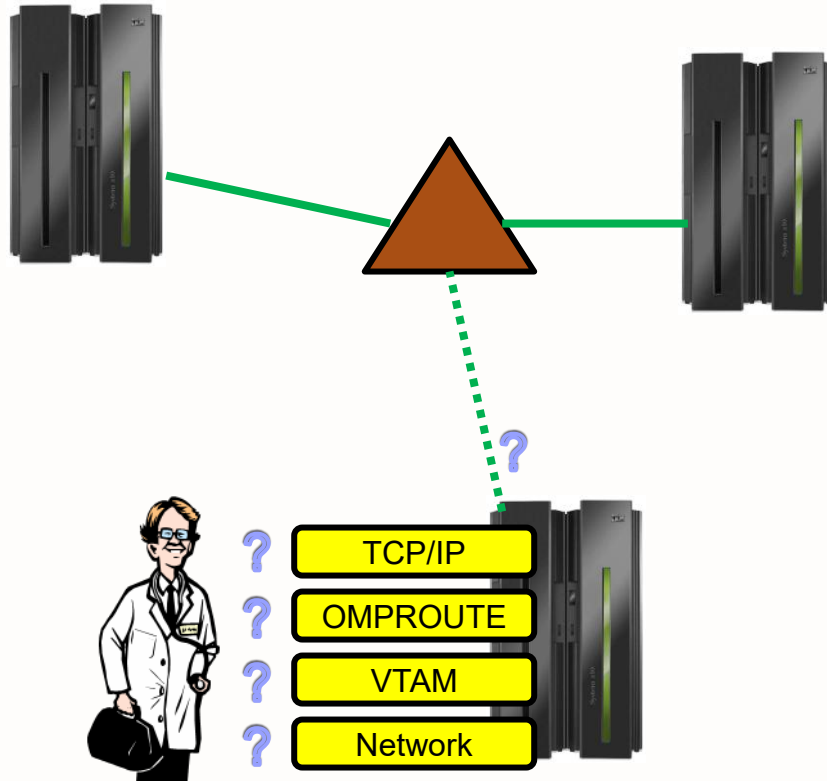
✓ Percentages relative to no acceleration (both using VIPAROUTE and 10GbE OSA Express 7)

Note: The performance measurements were collected in IBM internal tests using a dedicated system environment. The results obtained in other configurations or operating system environments may vary significantly depending upon environments used. Therefore, no assurance can be given, and there is no guarantee that an individual user will achieve performance or throughput improvements equivalent to the results stated here.

Sysplex Networking Technologies and Considerations

Network Sysplex Availability

Sysplex problem detection (and recovery!)



Sick? Better remove myself from the IP Sysplex!



Feeling better? Maybe it's time to rejoin the IP Sysplex

Monitoring:

- Monitor Communications Server health indicators
 - Storage usage critical condition (>90%) - CSM, TCP/IP Private, and ECSA
- Monitor dependent networking functions
 - OMPROUTE availability
 - VTAM availability
 - XCF links available
- Monitor for abends in Sysplex-related TCP/IP stack components
 - Selected internal components that are vital to Sysplex processing
- Monitor for repetitive internal abends in non-Sysplex related TCP/IP stack components
 - 5 times in less than 1 minute
- Selected network interface availability and routing (MONSYSPLEX)
- Detect when CSM FIXED, CSM ECSA, or CSM HVCOMMON has been constrained (>80% utilization) for multiple monitoring intervals
- IPsec infrastructure active and operational

Actions:

- Remove the stack from the TCP/IP Sysplex (manual or automatic)
 - Retain the current Sysplex configuration data in an inactive state when a TCP/IP stack leaves the Sysplex
- Reactivate the currently inactive Sysplex configuration when a TCP/IP stack rejoins the Sysplex (manual or automatic)

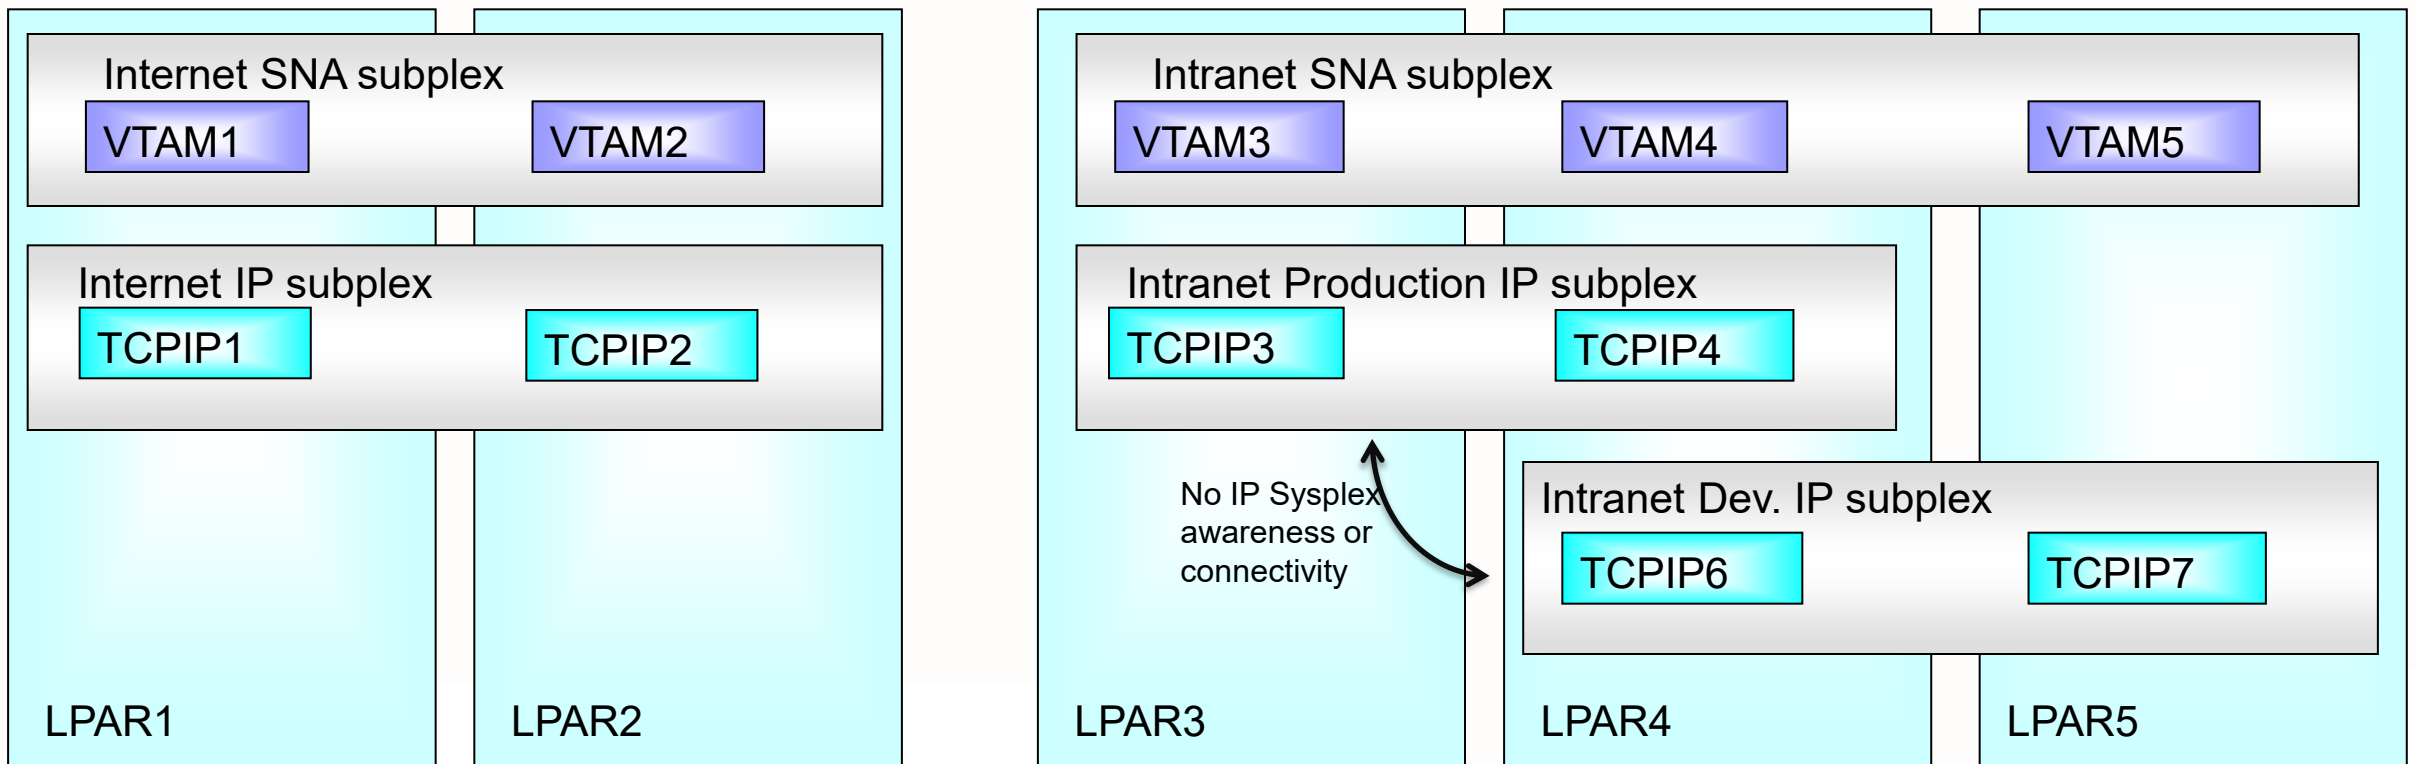
Sysplex Networking Technologies and Considerations

Network subplexing

Networking sub-plexing within a z/OS Sysplex

Subplexing allows network isolation of multiple secure areas

- TCP/IP Sysplex networking:
 - Establish IP Connectivity only to stacks in subplex; dynamic IP address discovery only within a subplex
- Coupling Facility resources:
 - VTAM generic resources, Sysplex ports resources, SWSA resources all defined at subplex level



Sysplex Networking Technologies and Considerations

Questions?

Experience more with IBM

SHARE25 August 17-21, 2025 CLEVELAND, OH

Visit us at the IBM Booth #317

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

Register for the IBM Z Client Reception

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

Don't miss the IBM Keynote in

Tuesday's General Session featuring Elpida Tzortzatos, IBM Fellow - CTO AI, IBM Z & LinuxONE.
[Tuesday, 19 August | 7:45 -9:00 AM | Main Ballroom, Hilton Cleveland](#)

Join our Lunch and Learn sessions


Grab a bite as we dive into how agentic AI is reshaping the IBM Z platform.
[Monday, 18 August | 12:00-1:00 PM | Hope Ballroom B, Hilton Cleveland Downtown - Level 3](#)

Join the fun with our new AI-powered Security IBM z17 edition of our Wheel of Cybersecurity Game Show!
[Thursday, 21 August | 11:15 AM-12:15 PM | Hope Ballroom B, Hilton Cleveland Downtown - Level 3](#)

Register →



IBM TechXchange October 6 - 9

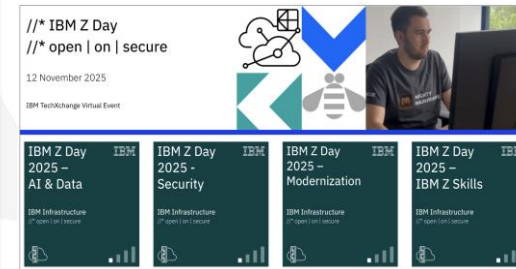


Explore, build and extend AI **where it matters most.** Join us in Orlando and experience the latest **IBM Z and LinuxONE** innovations with **150+ tech breakouts, labs, tech talks, demos and much more!**

Register now →



IBM Z Day November 12



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


Register here →



Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation



Notices and Disclaimers



- © 2025 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- 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.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and Disclaimers

- 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 about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml

Trademarks

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

IBM*
IBM Logo*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

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

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

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Docker and the Docker logo are trademarks or registered trademarks of Docker, Inc. in the United States and/or other countries. Docker, Inc. and other parties may also have trademark rights in other terms used herein.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.