

Sysplex Distributor Overview

—

Chelsea Jean-Mary
Client Technical Specialist

Linda Harrison
Client Technical Specialist



Contents

Introduction	03
Background on Virtual IP Addressing (VIPA)	04
What is Sysplex Distributor?	11
How it Works	18
Workload Distribution Methods	23
Implementation Steps	26
Performance Improvements	29
Troubleshooting	34

Introduction

Based on four key components:

Virtual IP Addressing (VIPA)

Dynamic VIPA Takeover support

Sysplex Technologies

Workload Management (WLM)



Background on Virtual IP Addressing

Creates a virtual IP address that is not associated with a physical network interface and thus cannot fail

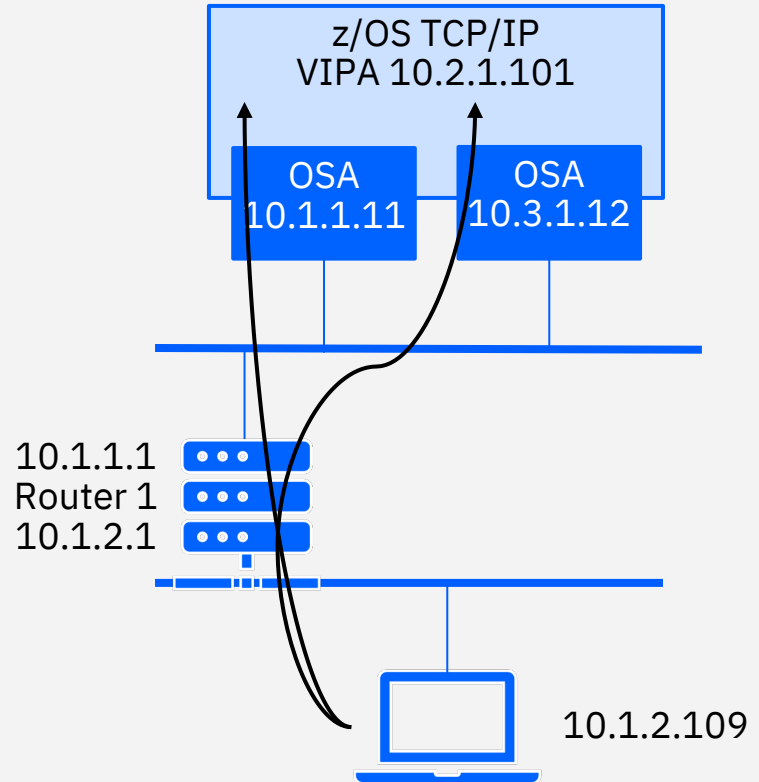
VIPA

Virtual IP Addressing (VIPA)

How this works on z/OS

- Remote servers connect to the VIPA address through one of the physical interfaces for z/OS system
- If the physical interface fails:
 - Dynamic route updates will be sent over the other physical interfaces
 - Downstream routers or servers will update their IP routing tables to use an alternate path to the VIPA

Benefit: TCP connections will not be broken and will recover non-disruptively through remaining interfaces



Has the capability to have the VIPA of a failed stack or failed application taken over automatically by an alternate stack or by restarting the failed application on a different LPAR

DVIPA

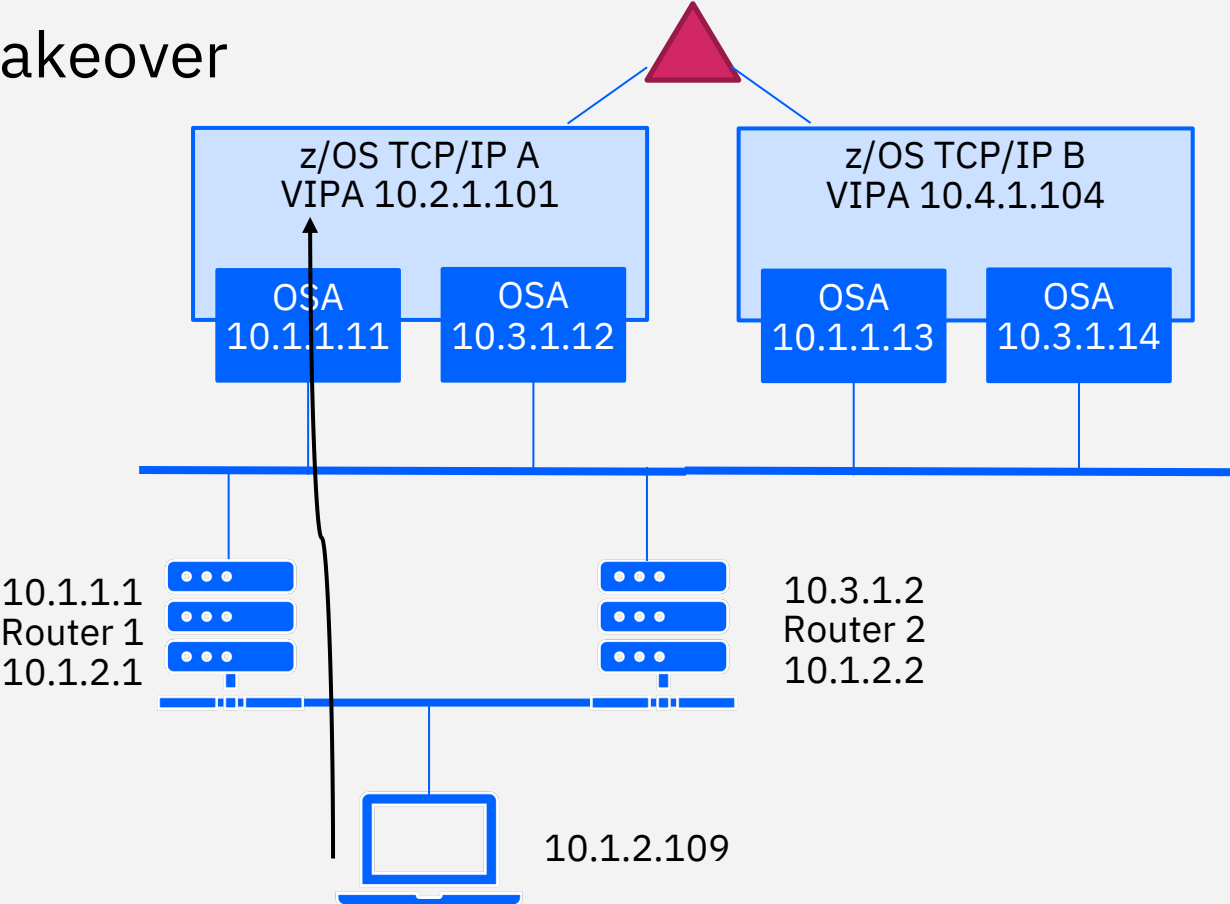
Dynamic VIPA Takeover

How this works on z/OS

- If the TCP/IP stack fails:
 - Connections with clients will be broken.
 - A standby stack will be dynamically modified to take over the VIPA of the failed stack.
 - Afterwards, remote clients can reconnect to the VIPA of the failed system.
- If the application fails:
 - When app is restarted on a different LPAR, it will connect to the same VIPA when it is coming up on a new LPAR
 - All requests will be routed to that application automatically

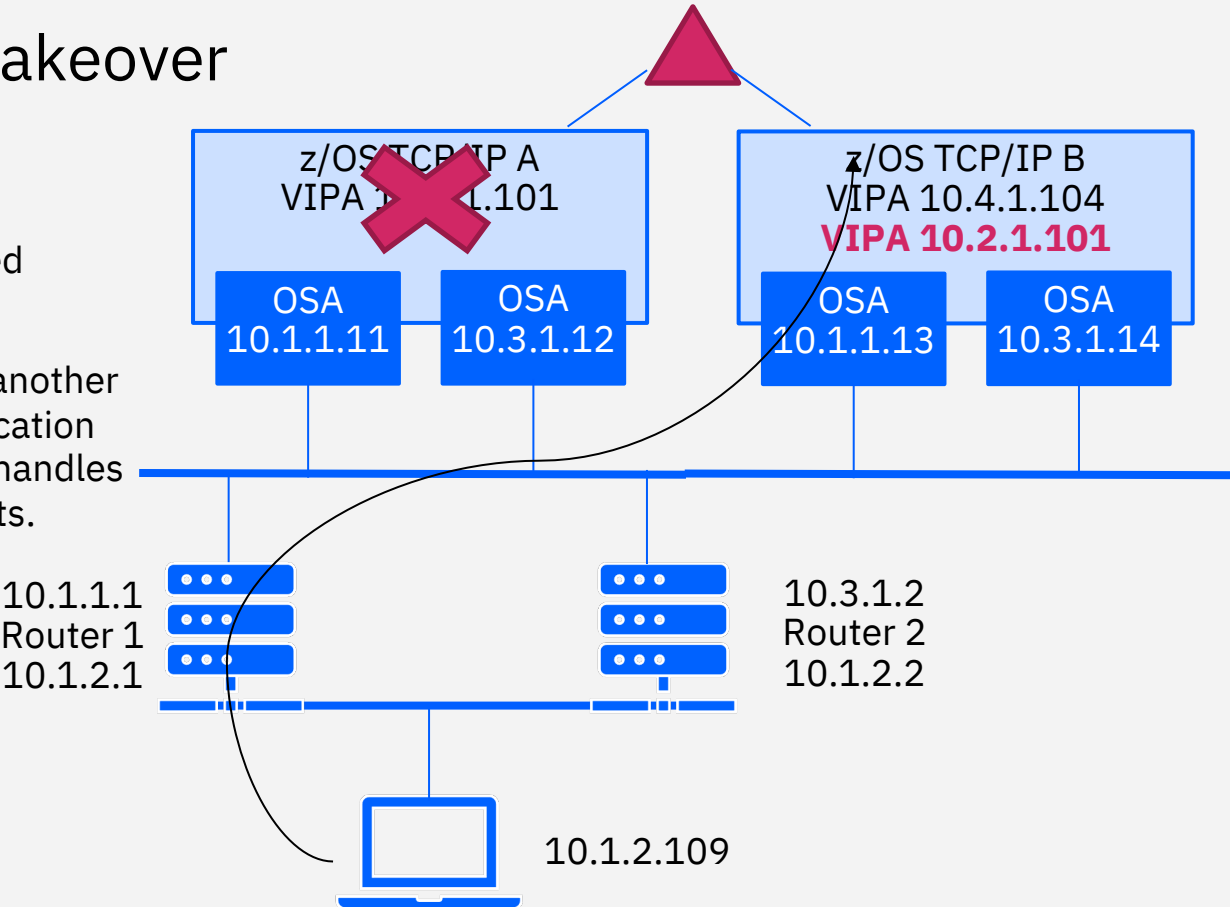
Benefit: TCP connections will be broken but can automatically reconnect through remaining stacks

Dynamic VIPA Takeover



Dynamic VIPA Takeover

- 1 TCP/IP stack A has failed
- 2 DVIPA 10.2.1.101 is automatically started on another node with the same application environment, which now handles all the failing node's clients.



What is Sysplex Distributor?

Base Sysplex	00
Parallel Sysplex	00
Benefits of Parallel Sysplex	00
Sysplex Distributor Overview	00


Journey to Sysplex

1964




Single System

1990



Base Sysplex

1994



Parallel Sysplex

Base Sysplex



A cluster of IBM mainframes acting together as a single system image with z/OS

- Sysplex – short for “Systems Complex” was introduced in 1990
- Data sharing and parallel computing to share workload for high performance and availability
- Communicates using channel-to-channel (CTC) connections between LPARs
- **Why was it created?** To solve the complexity of managing multiple z/OS systems

Parallel Sysplex



A set of systems within a sysplex that all have access to the same one or more Coupling Facilities

- In 1994, Coupling Facility (CF) was introduced to improve sysplex performance
- Allows you to increase the number of CPCs that can directly share work
- **Why was it developed?** Base sysplex was insufficient to provide speed and integrity necessary for data sharing

Benefits of Parallel Sysplex

Compatible change and nondisruptive growth

New IBM applications, software and hardware can be introduced non-disruptively and can coexist with previous release

Capacity and scaling

Can scale linearly from 2 to 32 systems and capacity meets every processing requirement known today

Dynamic workload balancing

Work can be dynamically distributed across any node in sysplex with available capacity

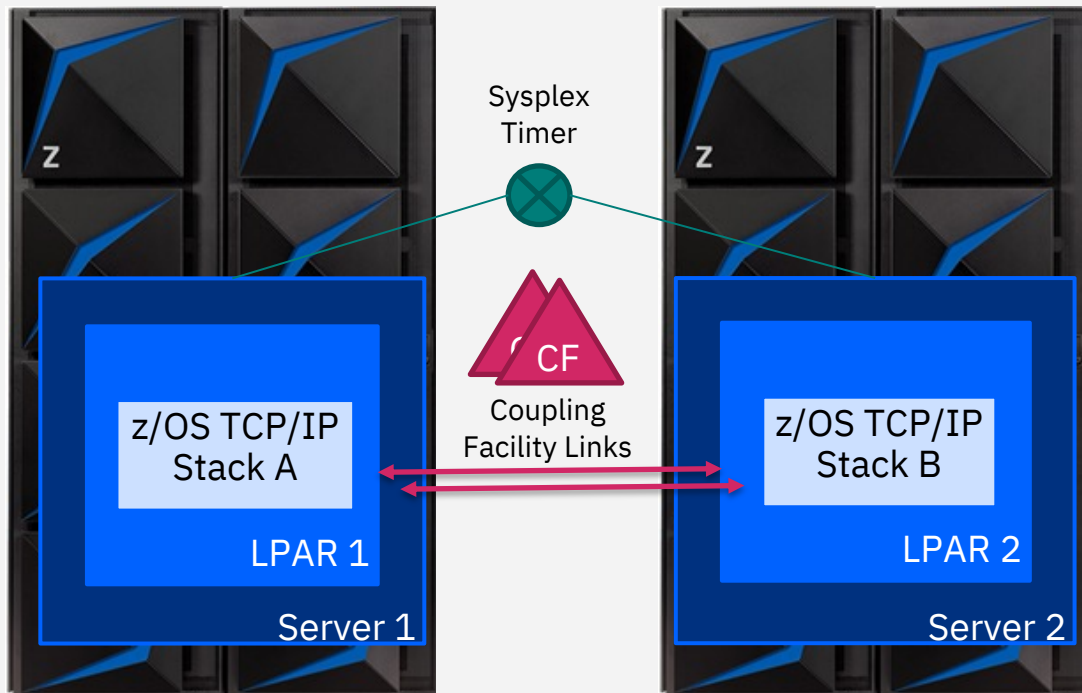
Single System Image

The collection of systems in the Parallel Sysplex appear as a single entity to the end user

What You'll Need

What you'll need to take advantage of Parallel Sysplex:

- z/OS
- IBM Z Server
- Two Coupling Facilities
- Coupling links
- Sysplex Timer



Sysplex Distributor Overview

- Builds on dynamic VIPA and automatic VIPA takeover for load distribution among servers in a sysplex
- Provides the benefits of Workload Management (WLM) without requiring WLM support in the application
- Adds the ability to nondisruptively move an application back to its original location after Dynamic VIPA takeover has been used to move the app after a failure
- Also known as a Distributed DVIPA and is just a special of DVIPA

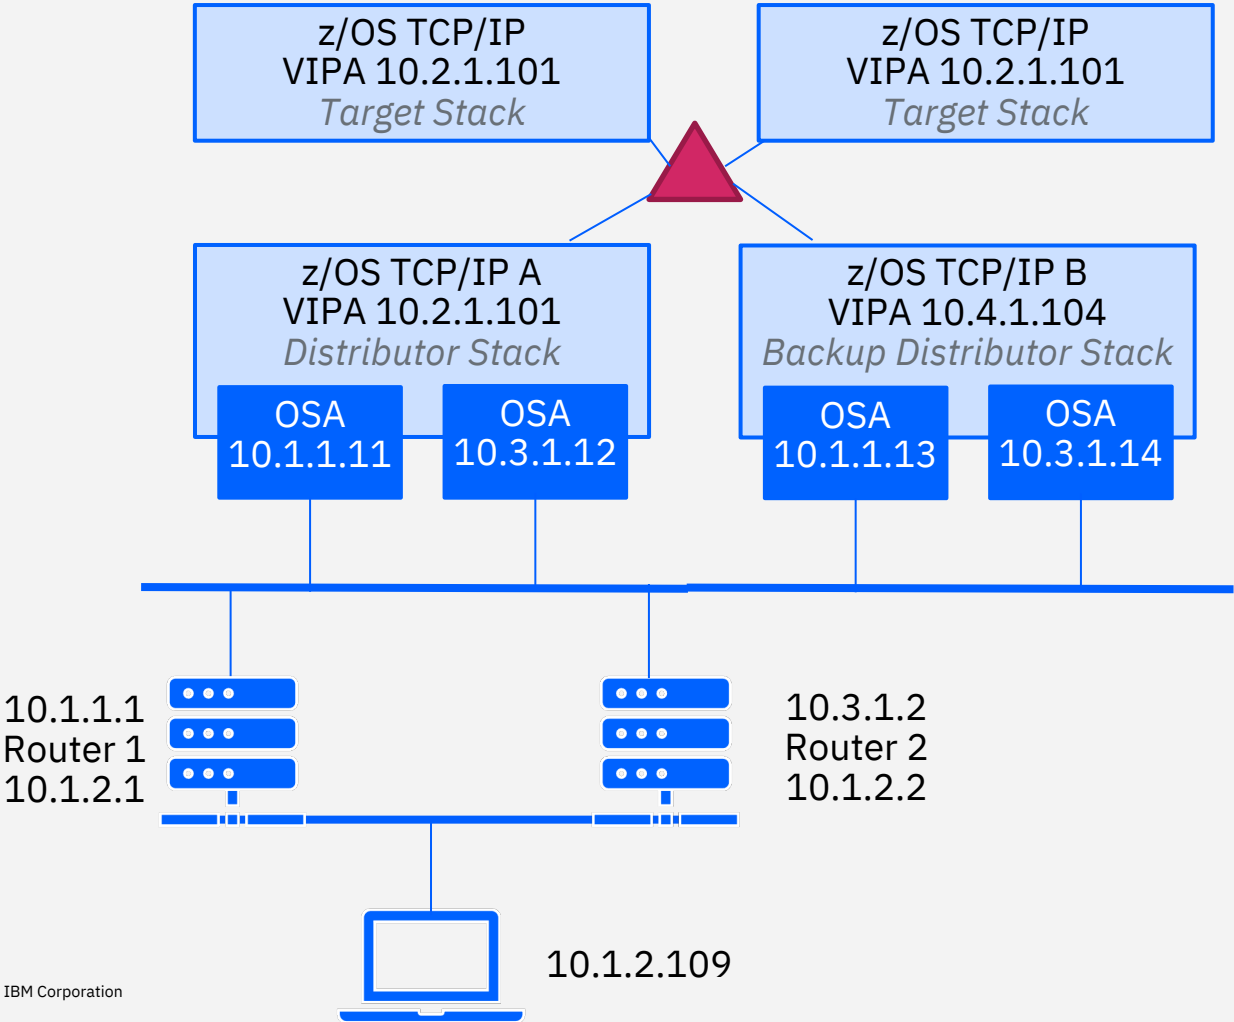
How It Works

Sysplex Distributor


A **distributed DVIPA** exists on several stacks but is advertised outside of the sysplex by one stack, the distributing stack

Single system image of sysplex distributor is represented by a special IP address, distributed DVIPA

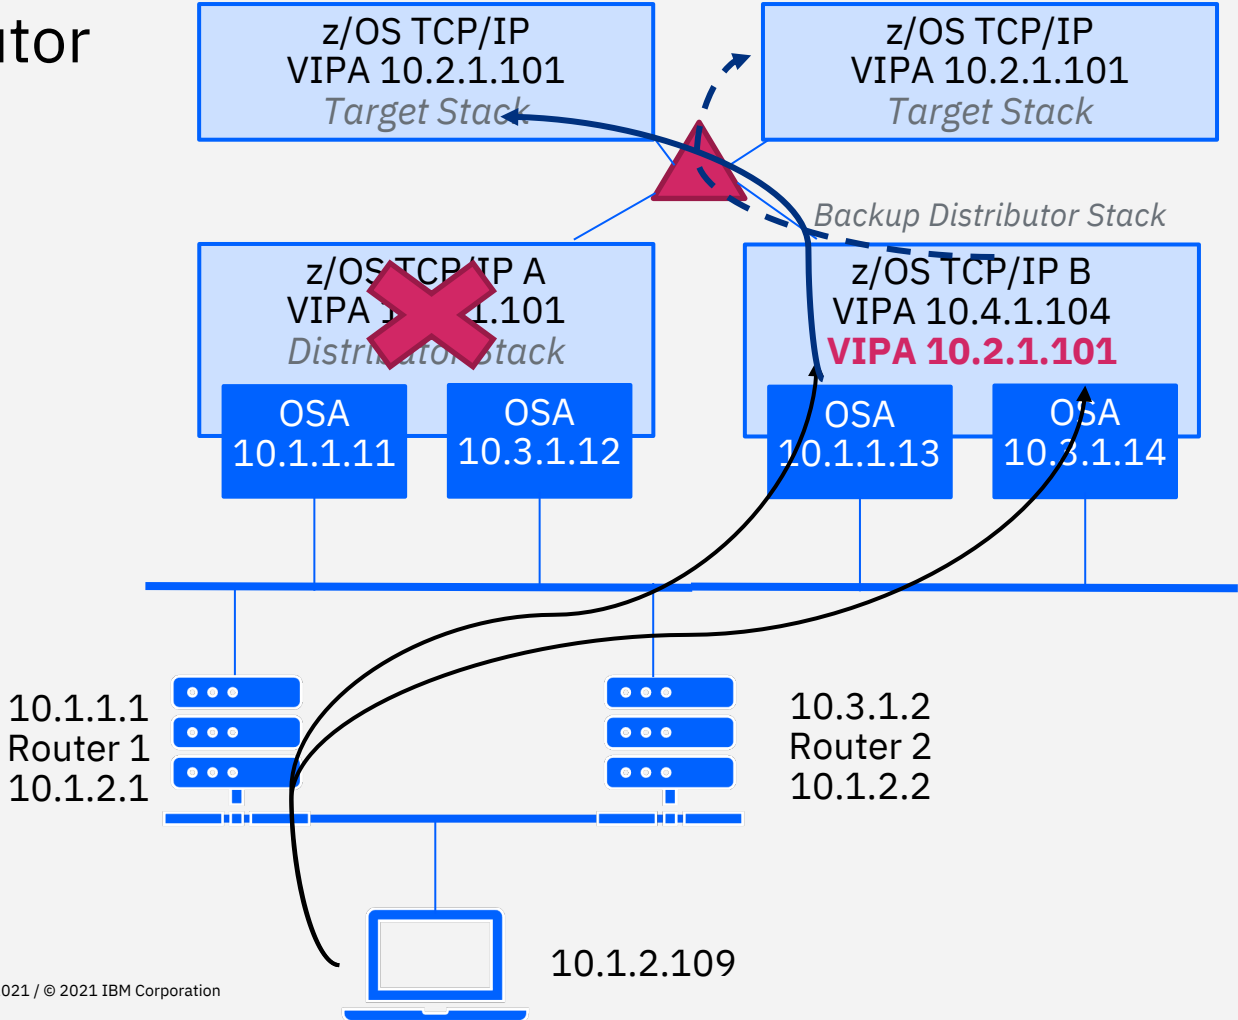
WLM provides the distributing stack with a WLM recommendation for each target system



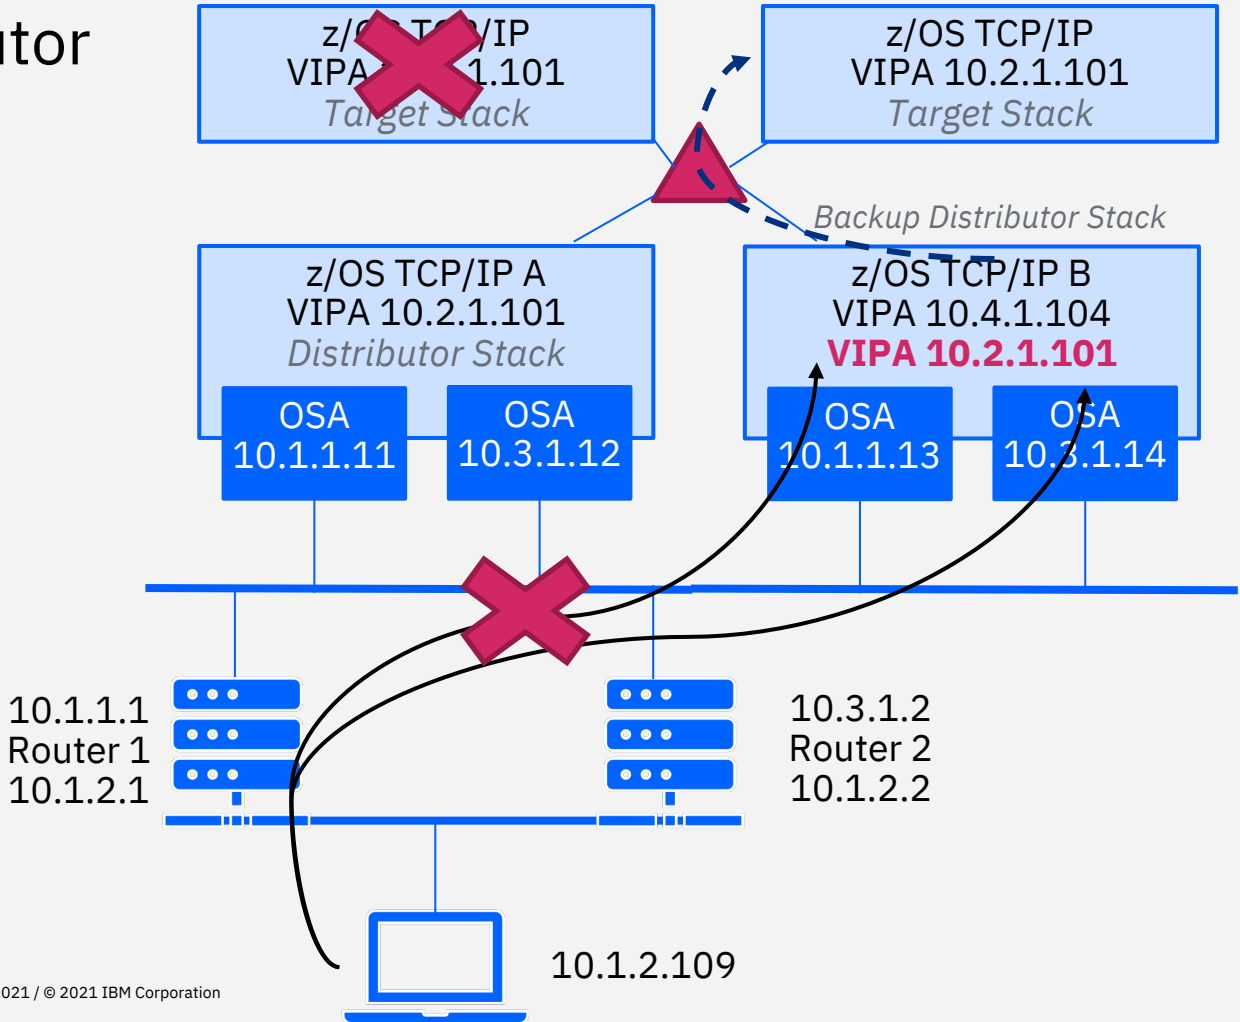
Sysplex Distributor



Sysplex Distributor



Sysplex Distributor



Workload Distribution Methods

Distribution Methods

	Distribution using	Ideal scenarios
BASEWLM	WLM system weights	Server acts as an access point to applications that runs in other address spaces (i.e. different server class) Ex. z/OS Connect
SERVERWLM	WLM server-specific weights	For most applications
WEIGHTEDACTIVE	A relative preference	When you have dissimilar targets and connections
ROUNDROBIN	Work is sent to each target in order	When you have like targets and like connections OR When even distribution of connections is needed to all target servers
HOTSTANDBY	Single target with one or more backup (hot standby) target servers	When you only want an application to be used on one system versus another AND want instant failover to a backup

Distribution Methods

Based on:

- Target Server Responsiveness (TSR)
 - Fraction, compound health-metric per Target Server (from 0 (bad) to 100 (good))
 - TCSR (Target Connectivity Success Rate)– indicates connectivity between distributing stack and target stack
 - SEF (Server Accept Efficiency Fraction)– indicates the target accept efficiency
 - QoS (Quality of Service)– takes retransmits and packet loss into consideration
- CER (Connection Establishment Rate) – indicates network connectivity between Server and Client
- Local Target Preference – prefer local server
- Timed Affinity – connections from the same client needs to be routed to the same server instance

Implementation Steps

How to Configure

1. Define the VIPA with a VIPADEFINE statement and include on a VIPADISTRIBUTE statement
2. Another TCP/IP stack can act as a backup for the distributed DVIPA by properly coding a VIPABACKUP statement
3. The options specified on a VIPADISTRIBUTE statement are inherited by a backup stack

How to Configure

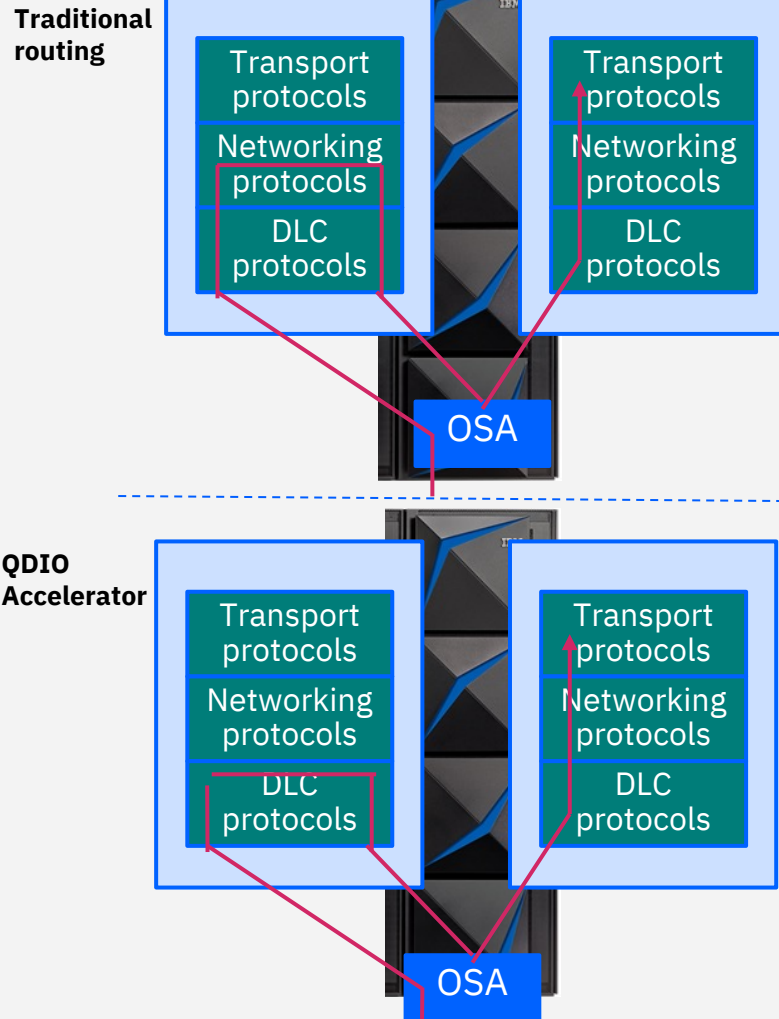
```
IPCONFIG SYSPLEXROUTING DYNAMICXCF 193.9.200.4 255.255.255.240 1
IPCONFIG6 DYNAMICXCF 2000::93:9:200:4
VIPADYNAMIC
  VIPADEFINE 255.255.255.192 9.67.240.2
  VIPADISTRIBUTE DEFINE 9.67.240.2 PORT 20 21 8000 9000 DESTIP
    193.9.200.2
    193.9.200.4
    193.9.200.6
  VIPADEFINE V6DVIPA1 2000::9:67:240:2
  VIPADISTRIBUTE DEFINE V6DVIPA1 PORT 20 21 8000 9000 DESTIP
    2000::93:9:200:2
    2000::93:9:200:4
    2000::93:9:200:6
ENDVIPADYNAMIC
```

The example shown to the left is a properly configured distributed DVIPA.

Performance Improvements

QDIOACCELERATOR

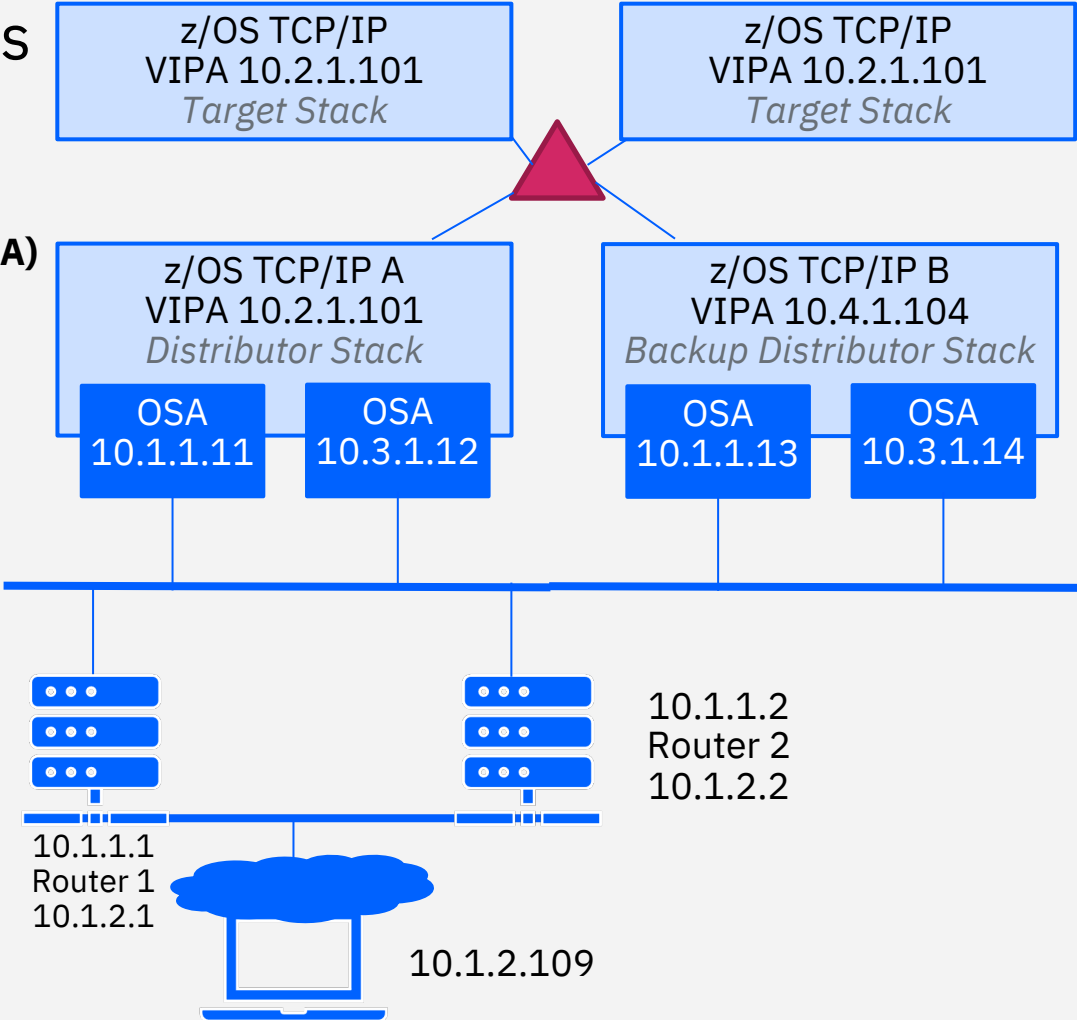
- A single z/OS stack may be used as a “router” between remote devices and other LPARs
- Bypasses the IP layer in z/OS, reducing path length and improving performance
- Supports Sysplex Distributor (SD) when traffic to target stack is sent over HiperSockets Dynamic XCF or QDIO as a result of VIPAROUTE definition



Performance Improvements

IP Syplex-Wide Security Associations (SWSA)

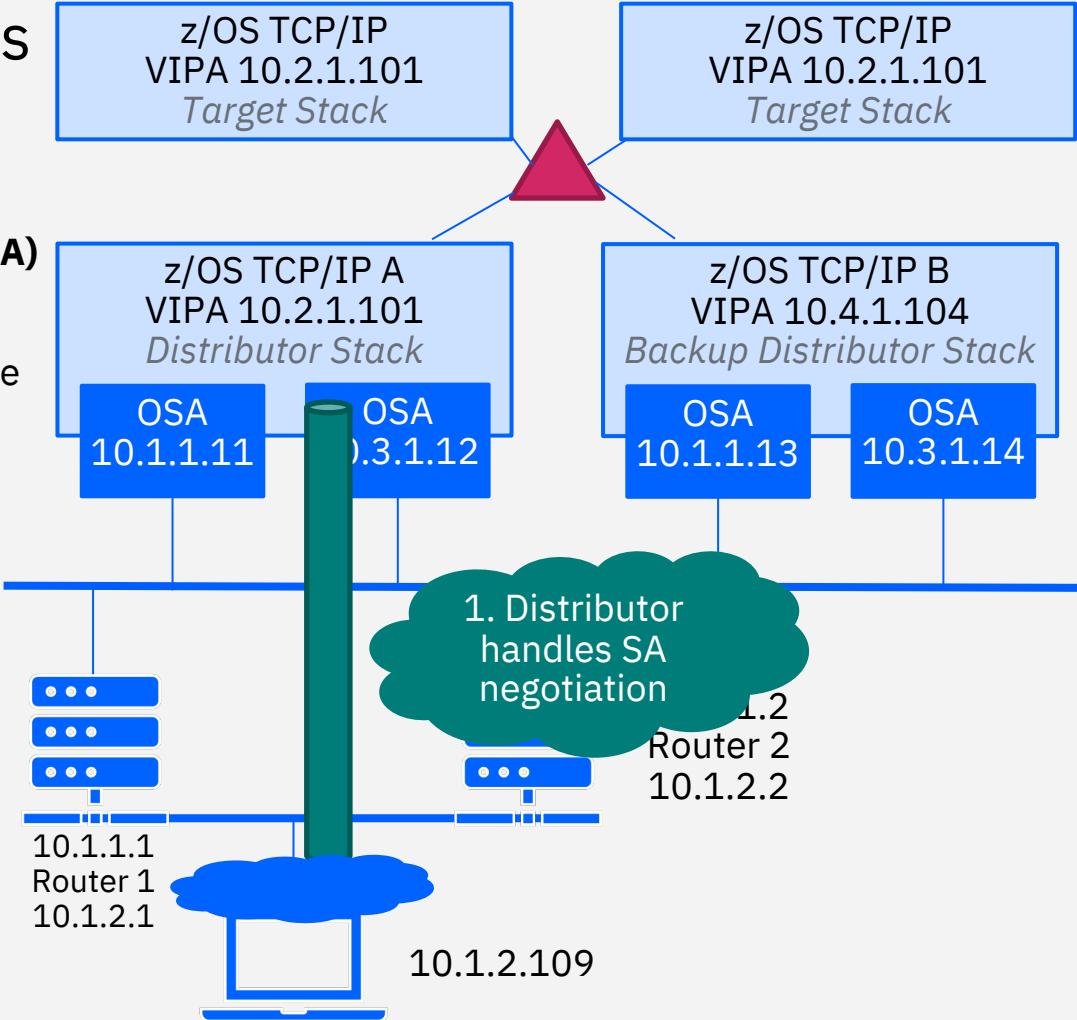
- If you're using Sysplex Distributor and IPsec and you lost connection, you won't have to repeat entire authentication
- When DVIPA is moved during DVIPA takeover, SWSA re-establishes new IPsec SAs with the same security characteristics as the SA's that existed on the host that previously owned the DVIPA
- Distributor handles SA negotiation and refresh



Performance Improvements

IP Syplex-Wide Security Associations (SWSA)

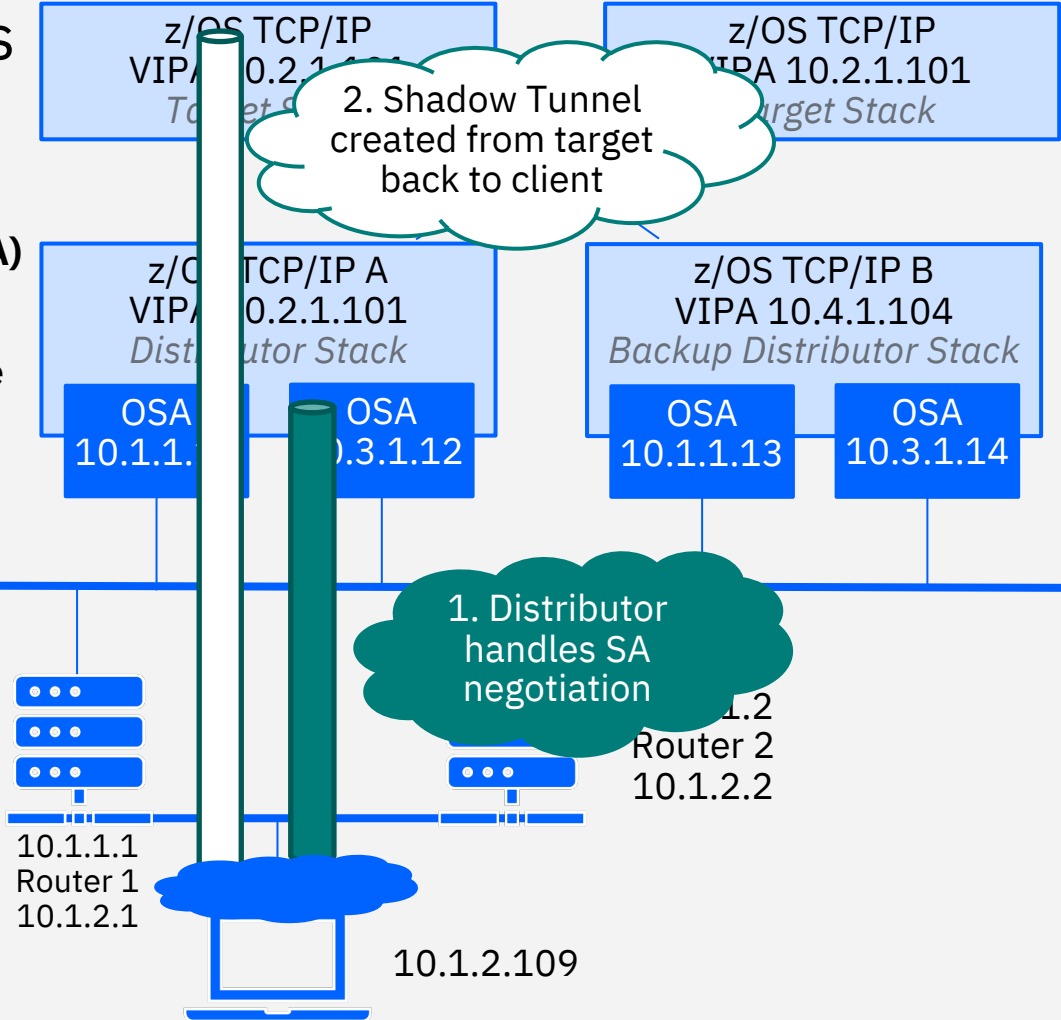
- If you're using SD and IPsec and you lost connection, you won't have to repeat entire authentication
- When DVIPA is moved during DVIPA takeover, SWSA re-establishes new IPsec SAs with the same security characteristics as the SA's that existed on the host that previously owned the DVIPA
- Distributor handles SA negotiation and refresh



Performance Improvements

IP Syplex-Wide Security Associations (SWSA)

- If you're using SD and IPsec and you lost connection, you won't have to repeat entire authentication
- When DVIPA is moved during DVIPA takeover, SWSA re-establishes new IPsec SAs with the same security characteristics as the SA's that existed on the host that previously owned the DVIPA
- Distributor handles SA negotiation and refresh



Performance Improvements

SYSPLEXMONITOR

- Detects problems that affect TCP/IP Sysplex Health
- Issues actionable message, leaves sysplex group
- Monitors for recoverable problems

Troubleshooting

Troubleshooting

Diagnosing Sysplex
Distributor problems
presents some
unique challenges

Problem

- A DVIPA is associated with multiple stacks in a parallel sysplex
- Locating where a problem is located can be more difficult

Solution

- Netstat command
- Display sysplex commands



Steps for Diagnosing Sysplex Problems

1. Determine that all stacks you expect to be communicating are in the same subplex (if subplex is being used).
2. Confirm the DVIPA definitions on a stack.
3. For Sysplex workload monitoring, run `netstat VDPT/-O` command and `netstat VCRT/-V`
4. If the output is not what you expected, run the sysplex command `VIPADyn`
5. If the output is not what you expected, run the display `netstat CONFIG/-f`
6. Run `Netstat ALLCONN(/-a)` with an IP address filter display command to verify that a server application has been activated and bound to the correct port.
7. Do a `CTRACE` with options `XCF`, `TCP`, and `SYSTCPD` on the stacks.

Issue the z/OS UNIX **netstat** command from the z/OS UNIX shell

Issue the `NETSTAT` command from TSO

Issue the `DISPLAY TCPIP,,NETSTAT` command from the system console

Troubleshooting

VARY TCPIP,,SYSPLEX

- Stop a particular server application on a target stack from receiving new DVIPA sysplex distributor workload but does not affect existing connections
- Use the command along with `PORT=portnum`, `JOBNAME=jobname`, or `TARGET` parameter

Thank you!

Chelsea Jean-Mary
Client Technical Specialist

—

chelsea.t.jean-mary@ibm.com
ibm.com

© Copyright IBM Corporation 2021. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., 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 at [Copyright and trademark information](#).



References

OS/390 Parallel Sysplex Configuration Volume 1: Overview

<http://www.redbooks.ibm.com/redbooks/pdfs/sg245637.pdf>

Chapter 2. Network Connectivity for Parallel Sysplex

<http://www.redbooks.ibm.com/redbooks/pdfs/sg245639.pdf>

SNA and TCP/IP Integration

[http://ps-2.kev009.com/p390/\\$IBMmanuals/TCPIP&SNA_Integration.pdf](http://ps-2.kev009.com/p390/$IBMmanuals/TCPIP&SNA_Integration.pdf)

z/OS Communications Server: IP Configuration Guide

<https://www.ibm.com/docs/en/zos/2.4.0?topic=addressing-configuring-distributed-dvipas-sysplex-distributor>

Networking on z/OS

<https://www.ibm.com/docs/en/zos-basic-skills?topic=distributor-sysplex-roles>

References

Configuring distributed DVIPAs – Sysplex Distributor

<https://www.ibm.com/docs/en/zos/2.1.0?topic=addressing-configuring-distributed-dvipas-sysplex-distributor>

Parallel Sysplex on IBM Z

https://www.ibm.com/it-infrastructure/z/technologies/parallel-sysplex?lnk=STW_US_STESCH&psrc=NONE&pexp=DEF&lnk2=learn_ITInfrastructuregen

Overview of diagnosing sysplex distributor problems

<https://www.ibm.com/docs/en/zos/2.1.0?topic=problems-overview-diagnosing-sysplex-distributor>

DYNAMICXCF Interfaces

Dynamic connectivity is automatically established between all sysplex members.

Using XCF messaging, each stack learns of other sysplex members

DYNAMICXCF interfaces are created using three transport technologies:

- Inside same LPAR: **IUTSAMEH** links or interfaces
- Inside same CPC: **HiperSockets**
- Between CPCs – **XCF Signaling**

