CS z/OS V1R7 Enhancements to IP Workload in a z/OS Sysplex

CS z/OS V1R7 Enhancements to IP Workload in a z/OS Sysplex

- **Sysplex autonomics phase II**
  - **Rejoin the Sysplex group**
- **Deactivate/reactivate stack-managed Dynamic VIPAs (VIPADEFINE/VIPABACKUP)**
- **Quiesce/resume target applications for Sysplex Distributor**

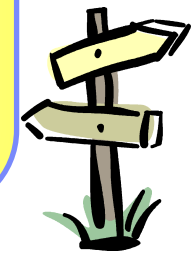

Sysplex autonomics phase II

Rejoin the Sysplex group

Background information: TCP/IP Sysplex recovery functions before z/OS V1R6

- **TCP/IP Sysplex recovery functions to protect against major hardware and software failures are triggered when a TCP/IP stack leaves the TCP/IP XCF group (terminates)**
  - If the leaving stack was a DVIPA owner, a backup stack will take over the DVIPA along with any associated Sysplex Distributor responsibilities - and new workload will continue to be processed by the Sysplex
  - If the leaving stack was a target stack for distributed workload, the distributing stack will remove it from its list of candidate target stacks - stop sending more connections to it

- **If a TCP/IP stack doesn't terminate, but enters an "unresponsive" condition, recovery functions are not triggered**
  - If the unresponsive stack is a Sysplex Distributor stack, no new connections to the distributed application will be processed and routing of inbound data through the distributing stack to target stacks for existing connections will cease
  - If the unresponsive stack is a target stack, the distributing stack will continue to send new connections to it and since WLM may see the target stack as lightly loaded, that stack may even be seen as a preferred stack for new workload - sending even more workload down the drain

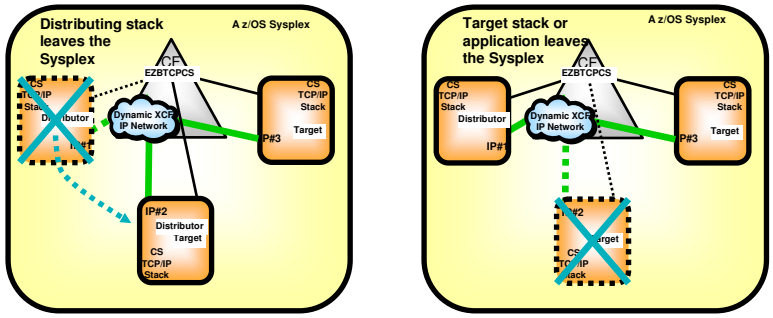

Background information: What can cause an unresponsive condition of TCP/IP?

- **There are a few known error conditions that can cause TCP/IP to become unresponsive or appear to be hanging - without actually terminating:**
  - The downstream network lost visibility of the distributing stack due to an OMPROUTE outage or malfunction and the network routers do not know how to reach the destination DVIPA addresses
  - VTAM is malfunctioning, data link control services are not working properly, and IP packets cannot be received or sent
  - TCP/IP is in a critical storage constraint situation
  - XCF IP network connectivity (Dynamic XCF) between the distributing stack and the target stacks is not functioning
  - Abends/errors in the TCP/IP Sysplex code components

Background information:
TCP/IP Sysplex autonomics phase I - at a z/OS V1R6 level

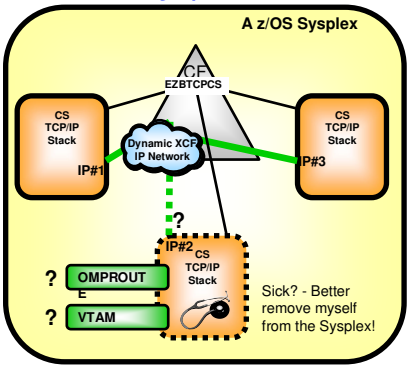

The assumption is that if a TCP/IP stack determines it can no longer perform its Sysplex functions correctly, it is better for it to leave the TCP/IP XCF group and by doing so, signal the other TCP/IP stacks in the Sysplex that they are to initiate whatever recovery actions have been defined, such as moving dynamic VIPA addresses or removing application instances from distributed application groups.

- Autonomic functions to reduce single point of failure for distributed applications in a Sysplex
  - Monitor CS health indicators
    - Storage usage - CSM, TCPIP Private & ECSA
  - Monitor dependent networking functions
    - OMPROUTE availability
    - VTAM availability
    - XCF links available
  - Monitor Communications Server component-specific functions
- Monitors determine if this TCPIP stack will remove itself from the Sysplex and allow a healthy backup to take ownership of the Sysplex duties (own DVIPAs, distribute workload)
- Monitoring is always done, but configuration controls in the TCPIP Profile determine if the TCPIP stack will remove itself from the Sysplex.

```
GLOBALCONFIG SYSPLEXMONITOR TIMERSECS
seconds RECOVERY|NORECOVERY
DELAYJOIN|NODELAYJOIN
```

- *Timersecs* - used to determine duration of the troubling condition before issuing messages or leaving the Sysplex (if Recovery)
- *RECOVERY* - TCPIP removes itself from the Sysplex. Recommended and is the default value.
- *NORECOVERY* - TCPIP does not remove itself from the Sysplex.
- *DELAYJOIN* - Delay joining Sysplex until OMPROUTE is up
- *NODELAYJOIN* - Join Sysplex immediately

Messages are always issued to the console when these conditions are detected regardless of SYSPLEXMONITOR Recovery specification
Messages are eventual action (deleted when the action is taken or problem is resolved)

New operator command is provided to allow TCPIP to leave the sysplex (ie. EZBTCPCS xcf group)
Vary TCPIP,,SYSPLEX,LEAVEGROUP

To have TCPIP rejoin the sysplex group, a Vary Obey of the TCPIP profile with sysplex configuration statements is needed.
Severe problems may require a TCPIP stack restart

Background information: TCP/IP Sysplex autonomics phase I: DELAYJOIN

- **Delay joining the Sysplex:**
  - The case that is being addressed by this function is where a primary stack is restarted and attempts to take back the dynamic VIPA addresses for which it is the primary owner.
  - If it attempts to do so before its OMPROUTE is up and active, a window may occur where the dynamic VIPA address have been taken back from the backup stack, but hasn't yet been advertised by the restarted stack - resulting in a time window where that address isn't available
    - That window would normally be very short (a few seconds), but could be longer depending on local operations procedures
  - The base idea is that it is better to leave the address with the backup stack until the restarting stack is fully ready to take over its responsibilities

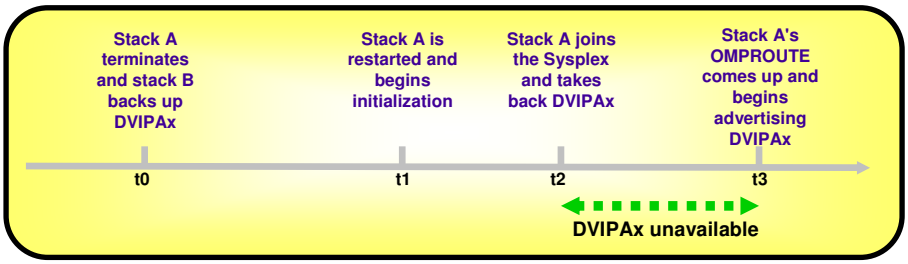

TCP/IP Sysplex autonomics phase II
New functions added in z/OS V1R7

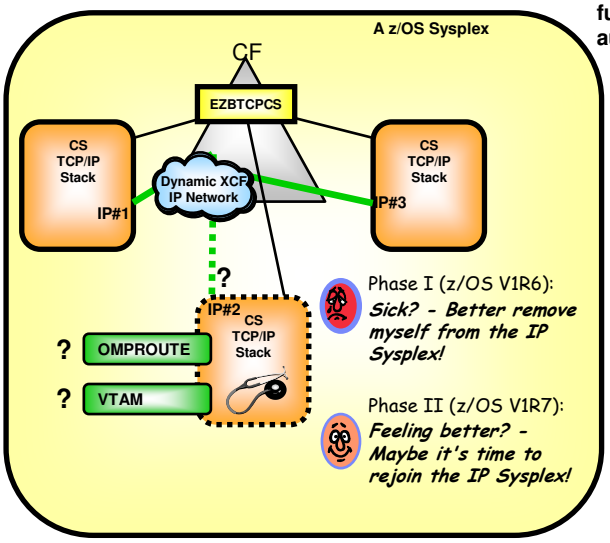

- **z/OS V1R7 adds the following functions to the TCP/IP Sysplex autonomics:**
  - Retain the current Sysplex configuration data in an inactive state when a stack leaves the Sysplex
  - Reactivate the currently inactive Sysplex configuration when a stack rejoins the Sysplex
  - New options for rejoining the Sysplex:
    - Via an operator command
    - Automatically when the error condition that caused the stack to leave the Sysplex has been cleared

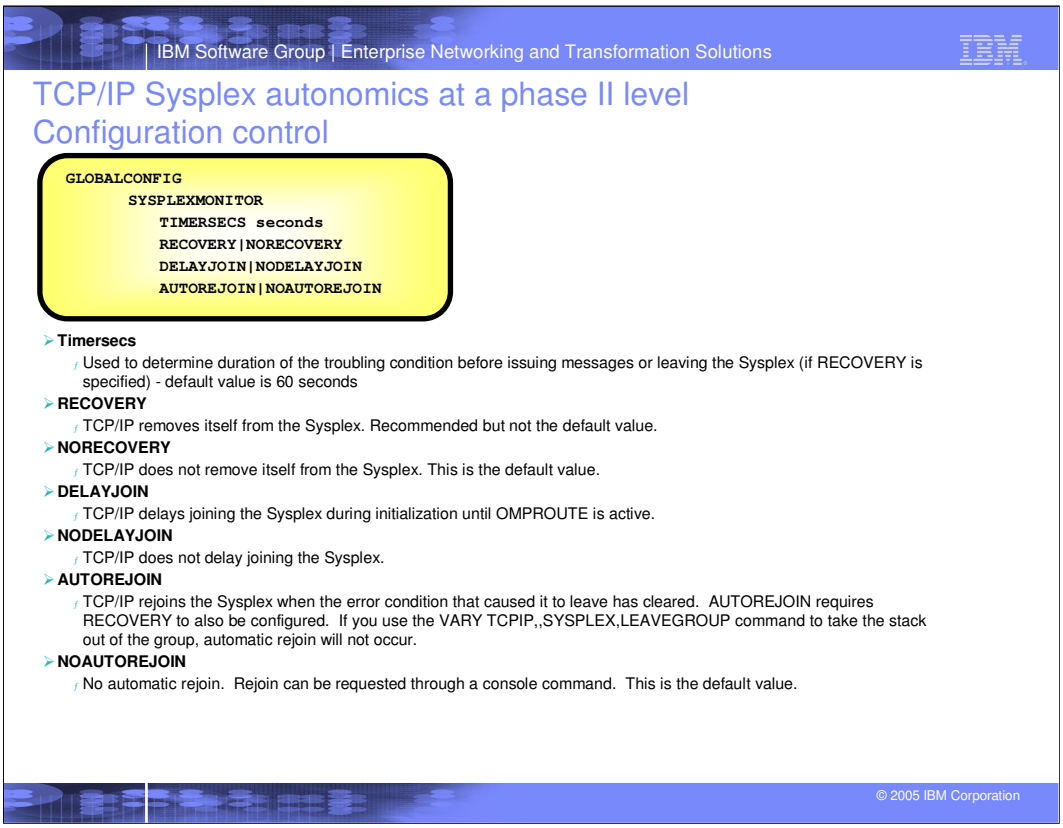

Messages are always issued to the console when these conditions are detected regardless of SYSPLEXMONITOR Recovery specification
Messages are eventual action (deleted when the action is taken or problem is resolved)

New operator command is provided to allow TCP/IP to leave the sysplex (i.e.. EZBTCPCS XCF group)
Vary TCPIP,,SYSPLEX,LEAVEGROUP

To have TCP/IP rejoin the sysplex group, a Vary Obey of the TCP/IP profile with sysplex configuration statements is needed.
Severe problems may require a TCP/IP stack restart

Recoverable problem conditions

- **EZZ9671E - tcpstackname DETERMINED THAT VTAM WAS INACTIVE FOR AT LEAST timevalue SECONDS**
  - The problem is cleared when VTAM is started
- **EZZ9672E or EZZ9678E - tcpstackname DETERMINED THAT OMPROUTE WAS NOT RESPONSIVE FOR AT LEAST timevalue SECONDS**
  - The problem is cleared when OMPROUTE is restarted
- **EZZ9673E - tcpstackname DETERMINED THAT DYNAMIC XCF CONNECTIVITY TO ALL PARTNERS WAS NOT AVAILABLE FOR AT LEAST timevalue SECONDS**
  - The problem is cleared when any XCF route is successfully activated
- **EZZ9679E - tcpstackname DETERMINED THAT CSM WAS CRITICAL FOR AT LEAST timevalue SECONDS**
  - The problem is cleared when CSM storage is no longer critical
- **EZD1172E - tcpstackname DETERMINED THAT ALL PARTNERS WERE UNREACHABLE FOR AT LEAST timevalue SECONDS**
  - The problem is cleared when any configured route (VIPAROUTE or XCF) to a partner is activated
- **EZD1187E - tcpstackname WAS NOT ABLE TO GET TCP/IP storagetype STORAGE**
  - Due to storage limits set by GLOBALCONFIG, the requested storage (ECSA or PRIVATE) was not available
  - The problem is cleared when the requested storage is no longer critical

Non-recoverable problem conditions

NOTES

- EZD1170E - tcpstackname WAS NOT ABLE TO GET TCP/IP storagetype STORAGE
  - The requested storage was not available and GLOBALCONFIG storage limits were not set for the requested storage
- EZZ9670E - tcpstackname SYSPLEX PROCESSING ENCOUNTERED A NONRECOVERABLE ERROR - abendcode - abendreasoncode
- EZZ9674E - tcpstackname SYSPLEX PROCESSING WAS NOT RESPONSIVE FOR AT LEAST timevalue SECONDS

OBEYFILE processing of the REJOIN options

NOTES

- AUTOREJOIN/NOAUTOREJOIN can be changed via a VARY TCPIP,,OBEYFILE command whether or not the stack is currently in the Sysplex group.
  - Changing from AUTOREJOIN to **NO**AUTOREJOIN will prevent the stack from automatically rejoining the Sysplex group when a problem detected by the Sysplex Autonomics function is relieved.
  - Changing from **NO**AUTOREJOIN to AUTOREJOIN will allow the stack to automatically rejoin the Sysplex group when all problems detected by the Sysplex Autonomics function are relieved.
    - If you change from **NO**AUTOREJOIN to AUTOREJOIN *after* the stack has left the Sysplex and *before the problem which caused it to leave has been relieved*, the stack will automatically rejoin the Sysplex group when the problem is relieved.
    - However, if you change from **NO**AUTOREJOIN to AUTOREJOIN *after the problem which caused the stack to leave the group has been relieved*, a VARY TCPIP,,SYSPLEX,JOINGROUP command will be needed to cause the stack to rejoin the Sysplex.

TCP/IP stacks leaving and rejoining the Sysplex

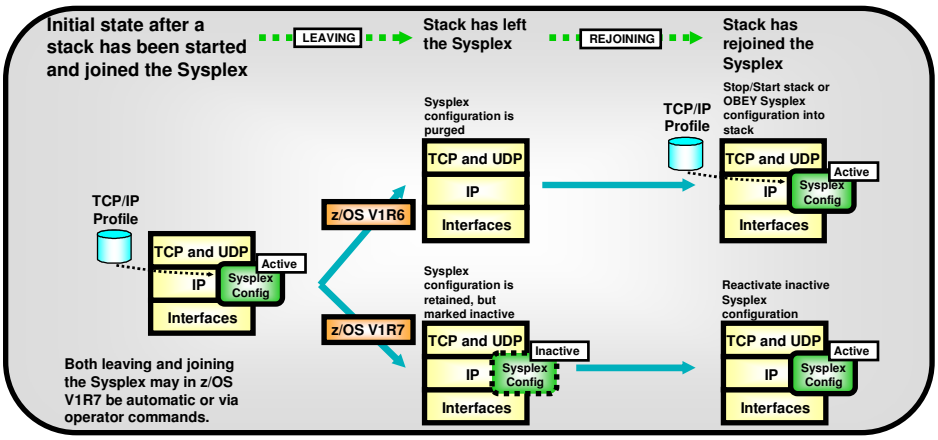

- **Leaving the Sysplex in z/OS V1R6, purges Sysplex configuration data from the stack's internal configuration blocks.**
  - To rejoin the Sysplex, the Sysplex configuration data must be reapplied to the stack's active configuration through a restart or an OBEY command
- **In z/OS V1R7, a stack's Sysplex configuration data will be retained in an inactive status when a stack leaves the Sysplex**
  - The inactive Sysplex configuration data will be shown on the NETSTAT VIPADCFG report as inactive
  - Rejoining the Sysplex will then reactivate the currently inactive Sysplex configuration data

Details on what Sysplex configuration data is retained

- **The VIPADYNAMIC configuration information that will be saved when the stack leaves the group, and reprocessed when the stack rejoins the group, includes the following VIPADYNAMIC definitions:**
  - VIPADEFINE
  - VIPABACKUP
  - VIPADISTRIBUTE
  - VIPARANGE
  - VIPASMPARMS
  - VIPAROUTE

- **The following Dynamic VIPA definitions are not saved when the stack leaves the Sysplex group:**
  - Target DVIPAs
    - These will be automatically re-created when the stack rejoins the group if this stack is still a target for that DVIPA from another (distributing) stack
  - BIND- or IOCTL-created DVIPAs
    - These must be re-created by the applications or by the MODDVIPA utility after the stack has rejoined the group

- **When a stack has left the Sysplex group, saved VIPADYNAMIC configuration (if any) can be displayed by the Netstat VIPADCFG/-F command.**

- **VIPADEFINE and VIPABACKUP definitions can be deactivated while the stack is out of the group. If a VIPADEFINE or VIPABACKUP configuration definition is deactivated at the time the stack rejoins the group, it will remain deactivated when the VIPADYNAMIC profile is reprocessed.**

How to rejoin

- **Rejoin can be**
  - Automatic
    - GLOBALCONFIG SYSPLEXMONITOR AUTOREJOIN
      - The stack will rejoin the Sysplex, when the problem that caused it to automatically leave the Sysplex has been relieved
      - Is only supported in combination with the SYSPLEXMONITOR RECOVERY option (leave the Sysplex automatically if a problem is detected)
      - Automatic rejoin is triggered by the events that clear the error condition (XCF links back up, OMPROUTE restarted, etc.)
      - Bounce prevention logic built into the storage condition logic if storage limits are set on GLOBALCONFIG
  - Operator command initiated
    - VARY TCPIP,[stackname],SYSPLEX,JOINGROUP
      - Matching the vary command to leave the Sysplex that was introduced in z/OS V1R6
      - Allowing full operator control over when to leave and when to rejoin the Sysplex
  - OBEYing a new Sysplex configuration into a stack that currently has left the Sysplex
    - VARY TCPIP,[stackname],OBEY,DSN=my.sysplex.conf
      - Will override (replace) currently inactive Sysplex configuration and rejoin the Sysplex
      - Provides compatibility for operations procedures that were established prior to z/OS V1R7
- **Rejoin will work under the same conditions as the initial join**
  - If DELAYJOIN is configured, the stack will ensure OMPROUTE is up and fully functional before the rejoin will take place

Operator commands to request a TCP/IP stack to leave or rejoin the Sysplex

- **Operator command to leave and rejoin the Sysplex:**

```
>>--Vary TCPIP,-|----------|-,SYSplex,->
                |          |
                +-PROCNAME-+

->|-LEAVEgroup-------------|---->
  |-JOINgroup--------------|
```

- Operator command must be issued on the system where the stack that is to leave or rejoin the Sysplex is running
- Allows an operator-initiated recovery of an error condition:
  - Sysplex autonomics monitoring functions issue error messages, but NORECOVERY was configured:
    - Command to leave the Sysplex
  - Sysplex autonomics monitoring functions issue messages about problem conditions being cleared:
    - Command to rejoin the Sysplex

```
VARY TCPIP,[stackname],SYSPLEX,LEAVEGROUP
.....
EZZ0053I COMMAND SYSPLEX,LEAVEGROUP COMPLETED SUCCESSFULLY

VARY TCPIP,[stackname],SYSPLEX,JOINGROUP
.....
EZD1176I TCPCS HAS SUCCESSFULLY JOINED THE TCP/IP SYSPLEX GROUP
EZD1192I THE VIPADYNAMIC CONFIGURATION WAS SUCCESSFULLY RESTORED FOR TCPCS
```

Notes on rejoining the Sysplex

NOTES

- **If the rejoin command is issued after the stack has left the Sysplex group, it will also reprocess the stack's saved VIPADYNAMIC configuration.**
  - Tip: Before issuing the VARY TCPIP,,SYSPLEX,JOINGROUP command, use the VARY TCPIP,,SYSPLEX,DEACTIVATE command to deactivate any DVIPA that you do not want restored when the stack rejoins the Sysplex
- **When this command is issued, the following message is displayed:**
  - EZD1178I THE VARY TCPIP,,SYSPLEX,JOINGROUP COMMAND WAS ACCEPTED
- **If VTAM is not running, or if the DELAYJOIN parameter is configured for GLOBALCONFIG SYSPLEXMONITOR and OMPROUTE is not initialized, the join will not take place until after VTAM (and OMPROUTE, if DELAYJOIN is configured) are initialized.**
- **When the join has completed, the following message is displayed:**
  - EZD1176I TCPCS HAS SUCCESSFULLY JOINED THE TCP/IP SYSPLEX GROUP
- **If the stack had previously left the group and VIPADYNAMIC configuration had been saved, you will see either:**
  - EZD1192I THE VIPADYNAMIC CONFIGURATION WAS SUCCESSFULLY RESTORED FOR TCPCS
    - or
  - EZD1193I ALL OF THE VIPADYNAMIC CONFIGURATION DEFINITIONS FOR TCPCS COULD NOT BE RESTORED
    - and specific error messages for the configuration conflict(s).
- **If the VARY TCPIP,,SYSPLEX,LEAVEGROUP command is used to take the stack out of the Sysplex group, a VARY TCPIP,,SYSPLEX,JOINGROUP command is required to bring the stack back into the group.**
  - The Sysplex autonomics function will not automatically bring the stack back into the group after a VARY TCPIP,,SYSPLEX,LEAVEGROUP command has been issued.
- **The VARY TCPIP,,SYSPLEX,JOINGROUP command will not be accepted if the Sysplex problem detection cleanup function was unsuccessful and message EZZ9675E was issued.**
  - If this has occurred, you have to restart the stack before it will be able to rejoin the Sysplex group.
- **If any configuration conflict is detected while reprocessing the saved VIPADYNAMIC configuration, specific informational messages are issued and the saved VIPADYNAMIC definitions that fail are discarded.**

Netstat command output when a stack has left the Sysplex group

- **When a stack has left the Sysplex group**
  - Netstat VIPADYN/-v report will show no entries ... unchanged from V1R6
  - Netstat VIPADCFG/-F report will show the saved configuration ... new
    - This is the configuration information that will be reprocessed when (if) the stack rejoins the Sysplex group
    - The Netstat report will be preceded by new informational messages to indicate that the stack has left the group and the VIPADYNAMIC configuration is not active

```
# netstat -F
EZZ2502I TCPCS is not a member of the TCP/IP sysplex group
EZZ2503I ALL VIPADYNAMIC configuration for TCPCS is currently inactive
MVS TCP/IP NETSTAT CS V1R7       TCPIP Name: TCPCS           19:07:50
Dynamic VIPA Information:

  VIPA Define:
    IpAddr/PrefixLen: 203.1.1.140/24
      Moveable: Immediate  SrvMgr: No
```

Deactivate/reactivate stack-managed Dynamic VIPAs (VIPADEFINE/VIPABACKUP)

Background information: DVIPA categories

- **Types of dynamic VIPA addresses:**
  - Stack-managed
    - Defined through VIPADEFINE and VIPABACKUP
    - All stack-managed DVIPAs are activated/deactivated when a stack joins/leaves the Sysplex
  - Event-managed
    - Defined through VIPARANGE
    - Individual DVIPAs are activated/deactivated when an application binds to one, or a MODDVIPA command is issued against one, or an application issues an IOCTL call for one
  - Distributed DVIPAs
    - Defined on distributing (owning) stack through VIPADEFINE/VIPABACKUP
    - Activated/deactivated on owning stack as other stack-managed DVIPAs
    - Defined on target stacks through VIPADISTRIBUTE statement on distributing stack
- **Event-managed DVIPAs can be moved around the Sysplex individually based on one of the events listed above**
- **Stack-managed DVIPAs (all of them) can only be moved when a stack leaves or joins the Sysplex**
  - There is no mechanism to request movement of an individual stack-managed DVIPA except through dynamic configuration changes - OBEYFILE processing
- **z/OS V1R7 implements a new operator command to request movement of individual stack-managed DVIPAs**

Improved operations: operator-initiated movement of individual stack-managed dynamic VIPA addresses

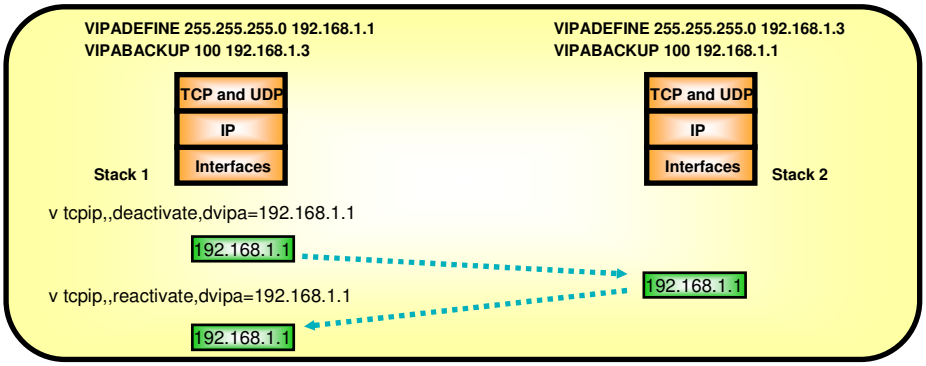

- **Deactivate**
  - DVIPA is deactivated and a configured backup stack will takeover the DVIPA
  - Backup DVIPA can be deactivated also, removing eligibility as a backup
- **Reactivate**
  - Original owner can regain ownership
  - Can also reactivate a backup DVIPA that's been deactivated
  - Prior to these commands, Vary OBEY files were needed to cause a DVIPA takeover
  - These commands can't be used on a DVIPA that was created from VIPARANGE with bind, ioctl(), or the MODDVIPA utility

Deactivate a stack-managed DVIPA

- **Removes the DVIPA resources from that stack, but saves the DVIPA's configuration information**
- **Behaves as if the DVIPA has been deleted on that stack**
- **If DVIPA is active:**
  - Stack stops advertising it
  - DVIPA is removed from the HOME list
  - Backup stack, if any, will activate the DVIPA (takeover)
- **Netstat VIPADYN/-v and DISPLAY TCPIP,,SYSPLEX,VIPAD will not show this DVIPA**
- **Netstat VIPADCFG/-F report will include this DVIPA's information under a new heading**
- **Deactivated DVIPA will appear in the Home list if:**
  - DVIPA has existing connections
    - Netstat VIPADYN/-v and DISPLAY TCPIP,,SYSPLEX,VIPAD will show DVIPA in QUIESCING status. DVIPA will disappear when last connection ends.
  - This stack is a target from the taking-over stack
    - Netstat VIPADYN/-v and DISPLAY TCPIP,,SYSPLEX,VIPAD will show DVIPA as Active and Dest

Reactivate a stack-managed DVIPA

- **Reprocesses the deactivated configuration definition**
- **Reactivated DVIPA will be active if:**
  - DVIPA's origin was VIPADEFINE
  - DVIPA's origin was VIPABACKUP, MOVEABLE parameter was configured, and the DVIPA is not active elsewhere in the Sysplex
- **If reactivated DVIPA is active:**
  - Stack advertises the DVIPA
  - DVIPA is added to the HOME list (if not already there)
  - Current owner, if any, will give up the DVIPA (takeback)
- **Reactivation of a DVIPA reprocesses the deactivated configuration**
  - Reactivation will fail if there is conflicting configuration on this stack or in the Sysplex. If this occurs:
    - The DVIPA remains deactivated
    - You can either
      - Delete the deactivated DVIPA (VIPADELETE will delete the DVIPA definition and any VIPADISTRIBUTE definitions for that DVIPA), or
      - Remove the conflicting configuration and reissue the VARY TCPIP,,SYSPLEX,REACT command

Deactivate/reactivate stack-managed DVIPAs

- **Deactivating/reactivating can be done when stack is not in the Sysplex group.**
  - DEACTivate marks the DVIPA as deactivated
  - REACTivate will unmark the DVIPA
  - Deactivated DVIPAs remain deactivated when stack rejoins the group and the VIPADYNAMIC configuration is reprocessed
- **Deactivation/reactivation only applies to VIPADEFINE or VIPABACKUP DVIPAs.**
  - You cannot deactivate a target DVIPA (unless it is also configured as VIPADEFINE or VIPABACKUP), or a VIPARANGE DVIPA created by BIND, SIOCSVIPA or SIOCSVIPA6 ioctl, or the MODDVIPA utility.

```
>>--Vary TCPIP,-|----------|-,SYSplex,->
                |          |
                +-PROCNAME-+

->|-DEACTivate,DVIPA=dvipa-|
->|-REACTivate,DVIPA=dvipa-|
```

- ***dvipa*** **is the IPv4 address, IPv6 address, or IPv6 interface name**

Netstat output when stack-managed DVIPAs are deactivated

Example of a Netstat VIPADCFG/-F report after a DVIPA is deactivated via VARY TCPIP,,SYSPLEX,DEACT,DVIPA=INTF1

```
# netstat -F
MVS TCP/IP NETSTAT CS V1R7       TCPIP Name: TCPCS           18:07:00
Dynamic VIPA Information:

  VIPA Define:
    IpAddr/PrefixLen: 203.1.1.140/24
      Moveable: Immediate  SrvMgr: No

Deactivated Dynamic VIPA Information:

  VIPA Define:
    IntfName: INTF1
      IpAddr: 206::1401:1
        Moveable: Immediate  SrvMgr: n/a

  VIPA Distribute:
    DestIntf:    INTF1
      Dest:      206::1401:1..6000
        DestXCF: ALL

          SysPt: No   TimAff: No    Flg: BaseWLM

#
```

Quiesce/resume target applications

Background information: How to quiesce a distributed application or a single instance of a distributed application

- **Target stacks configured to be eligible for Sysplex distribution (identified as targets on the DESTIP parameter of the VIPADISTRIBUTE statement) and that have an application listening on the distributed port will receive work.**

- **A distributed application can be quiesced completely:**
  - Issue a Vary TCPIP,,OBEY command on the distributing stack containing a VIPADIST DELETE statement for the distributed application you want to quiesce.
    - Stops new connections from being distributed while preserving existing connections.
    - Requires the operator to configure a profile with the VIPADIST DELETE statement.
    - Requires one VIPADIST DELETE statement for each DVIPA being distributed to the target application.
    - Not really useful to completely remove a target stack from the Sysplex distributor environment.

- **Individual target application instances can be removed from Sysplex distribution with a VARY TCPIP,,OBEYFILE command on the distributing stack:**
  - Replacing the list of destination XCF addresses that connections can be distributed to - removing the one target stack on which the application instance, you want to remove, executes.
  - No new connections for this application instance will be sent to the target stack. Existing connections will continue until they are closed.

- **No easy way to remove all applications instances on a target stack from Sysplex distribution.**
  - Only way is again to use OBEYFILE processing to replace the destination XCF addresses on all VIPADIST statements that distribute work to the target stack in question.

Improved operations: operator-initiated quiesce and resume of individual server applications or full target systems

- **Ability to quiesce a target system or an application instance prior to shutdown**
  - Planned maintenance scenarios of system or application
    - Allows existing systems or applications to drain work queue prior to shutdown
  - Relieve temporary constraints of resources on target system
  - Temporary - Does not affect Sysplex Distributor's permanent configuration
  - Issued on target system being affected
  - Can also be used to control individual server applications in a SHAREPORT group
  - Only way to achieve similar capability earlier was via temporary configuration changes based on OBEYFILE commands
- **VARY TCPIP,,SYSPLEX,QUIESCE,options**
  - TARGET - Quiesces all applications on target stack.
  - PORT=xxx - Quiesce all applications bound to the specified port on this stack
    - JOBNAME=jobname - Allows quiesce of a single application in SHAREPORT group
    - ASID=asid - Further qualify job being quiesced (such as when dealing with duplicate jobnames)
  - No new TCP connections sent to the quiesced target (stack or application)
    - For all Distributed DVIPAs that the entity is a target for
  - Existing TCP connections are maintained (or in other words, the process is non-disruptive)
- **VARY TCPIP,,SYSPLEX,RESUME,options**
  - TARGET|PORT|JOBNAME|ASID
  - Allows identified target stacks and/or applications to once again be targets for distribution

VARY TCPIP,,SYSPLEX,QUIESCE or RESUME

- **These commands are always issued on the target stack or system and impacts the ability for application instances on that target stack or system to receive workload from Sysplex Distributor.**

```
>>--Vary TCPIP,-|----------|-,SYSplex,->
                |          |
                +-PROCNAME-+

  |-QUIesce,POrt=portnum-|--------------------|
                         |                    |
                           +-,JOBNAME=jobname--+-|----------------|
                                                                  +,ASID=asid------+

  |-QUIesce,TARGET---------|

  |-RESUME,POrt=portnum-|----------------------|
                        |                      |
                                                                +-,JOBNAME=jobname---+-|-
----------------|
                                                                    +,ASID=asid-------+

  |-RESUME,TARGET----------|
```

Guidelines for use of QUIESCE and RESUME commands

- **These commands must be issued on the system and TCP/IP stack where the application instance is running.**
- **The commands apply to a single TCP/IP stack's application instance.**
  - If the server needs to be quiesced or resumed over multiple stacks in a CINET environment, the command would need to be issued on each stack.
- **Any Sysplex distributor timed affinities with the application instance being quiesced will be terminated.**
  - Existing connections are not affected.
- **The quiesce state is associated with the application's active listening socket.**
  - If the application is recycled or if the application closes and opens a new listening socket on the specified port, it will no longer be in a quiesced state.
- **If the application is bound to the unspecified address, it may continue to receive connection requests that are not using a distributed DVIPA as the destination IP address.**
- **The QUIESCE state for a TARGET persists for all application instances (existing and new) running on this TCP/IP stack, until the TCP/IP stack is recycled or a V TCPIP,,RESUME,TARGET command is issued.**
- **When an entire TCP/IP stack is quiesced via the TARGET option, you cannot resume individual applications for workload distribution.**
  - You can, however, resume distribution for the entire TCP/IP stack using the V TCPIP,,RESUME,TARGET command.
- **RESUME with the TARGET option is the only valid command following a QUIESCE with the TARGET option command.**
- **When a TCP/IP stack is quiesced, the "ready count" (Rdy) field that appears on the Netstat VDPT display (issued on the Sysplex Distributor routing stack) will be zero for all entries associated with this target TCP/IP stack.**

Netstat

- **Use the Netstat ALL/-A report to display applications status:**

```
MVS TCP/IP NETSTAT CS V2R7       TCPIP NAME: TCPCS            17:40:36
Client Name: CICS1                    Client Id: 0000004A
Local Socket: 0.0.0.0..27             Foreign Socket: 0.0.0.0..0
  Last Touched:       17:09:22          State:              Listen


   CurrentBacklog:     0000000000        MaximumBacklog:     0000000010
   CurrentConnections:    0000000300     SEF:                098
   SharePort: WLM
      RawWeight:         02              NormalizedWeight:   01
   Quiesced:                        Dest
```

The Quiesced state indicates if this server application has been quiesced for DVIPA Sysplex Distributor workload balancing. If the value is Dest then this server will receive no new DVIPA Sysplex Distributor workload connections until the server application has been resumed. When the server application is resumed, the Quiesced value will change to No.

Trademarks, Copyrights and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linus Torvalds.

Other company, product and 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 and/or changes in the product(s) and/or program(s) 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.

Information is provided "AS IS" without warranty of any kind. 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 (e.g., 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.