z/OS Communications Server

Usability

Simplification and ease-of-use are enhancements that in some way or another make life a little easier for all of you. They are improvements in how you configure or operate components of the Communications Server.

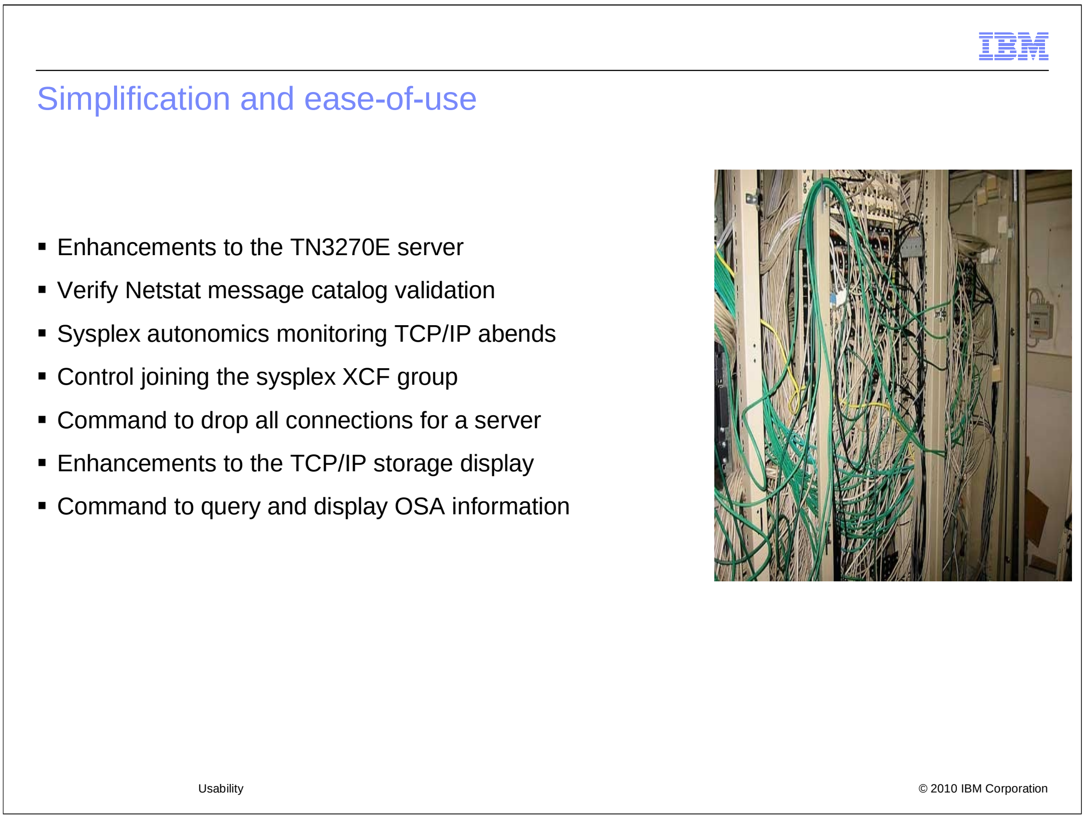

Simplification and ease-of-use

- Enhancements to the TN3270E server
- Verify Netstat message catalog validation
- Sysplex autonomics monitoring TCP/IP abends
- Control joining the sysplex XCF group
- Command to drop all connections for a server
- Enhancements to the TCP/IP storage display
- Command to query and display OSA information

Usability

Several changes were made in the z/OS V1R12 Communications Server to make the product more consumable. This is especially true in the areas of configuration and administration.

Several enhancements were made to the TN3270E server.

A new check has been added to ensure the correct Netstat catalogs are being used.

New sysplex autonomics and a new sysplex command simplify management of TCP/IP stacks in a sysplex.

A new drop-connections command simplifies dropping all connections from a server.

The TCP/IP storage display has been changed to clarify the type of storage used in ECSA.

A new command to query the OSA will simplify getting OSA information that isn’t available using OSA/SF.

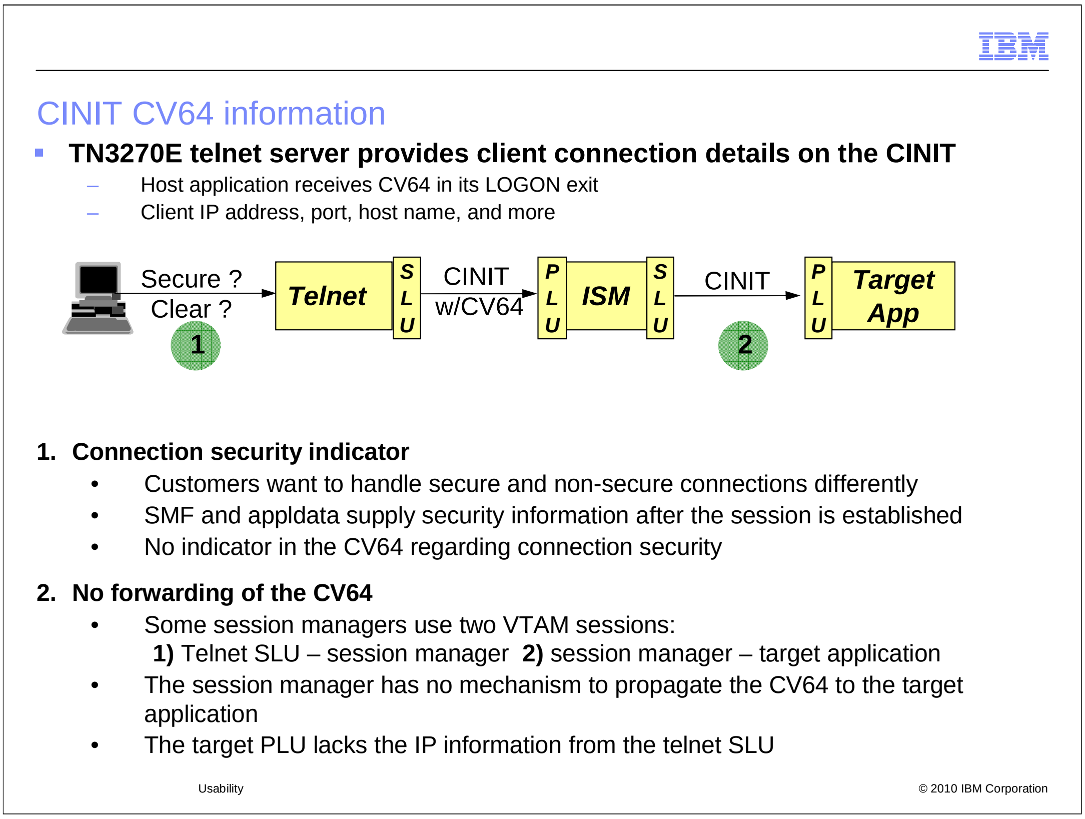

CINIT CV64 information

- **TN3270E telnet server provides client connection details on the CINIT**
  - Host application receives CV64 in its LOGON exit
  - Client IP address, port, host name, and more

Secure ?
Clear ?
1
Telnet
SLU
CINIT
w/CV64
PLU
ISM
SLU
CINIT
2
PLU
Target App

1. **Connection security indicator**
   - Customers want to handle secure and non-secure connections differently
   - SMF and appldata supply security information after the session is established
   - No indicator in the CV64 regarding connection security
2. **No forwarding of the CV64**
   - Some session managers use two VTAM sessions:
     1) Telnet SLU – session manager **2)** session manager – target application
   - The session manager has no mechanism to propagate the CV64 to the target application
   - The target PLU lacks the IP information from the telnet SLU

Usability

© 2010 IBM Corporation

When you establish a session using telnet, Telnet creates a secondary logical unit (SLU) or telnet LU name. Telnet uses the SLU to initiate a session with the target primary logical unit (PLU) by sending a CINIT to the PLU. Telnet builds a control vector, CV64, with IP and other information and appends it to the CINIT that is delivered to the PLU in its LOGON exit. Some PLUs have logic to parse the CV64, if it is present. Once the session is established, you can issue a VTAM display of the telnet application SLU which shows IP information associated with the telnet client.

Currently there is no real-time indicator that can be used to determine the security of the TCPIP connection while the session is being established. Telnet provides the security level of the TCP/IP connection using SMF records and appldata on the TCP connection termination SMF record. However, this information is available too late to be useful as a real-time security check.

A target application PLU trying to parse a CV64 might encounter a problem if there is a session manager between telnet and the PLU. Some session managers actually have two sessions for the session from the telnet SLU to the target application PLU. One of the sessions is from telnet to the session manager and the other session is from the session manager to the target application. The session manager receives the CINIT with the CV64 in its LOGON exit, but it has no way to pass the CV64 on to the target application PLU. The target PLU does not receive the CV64 that carries the IP information for the telnet client and any functionality that the PLU provided based on the IP information is lost. The session manager in this picture is IBM Session Manager (ISM). A session manager that does CLSDST PASS to establish the sessions does not have this problem.

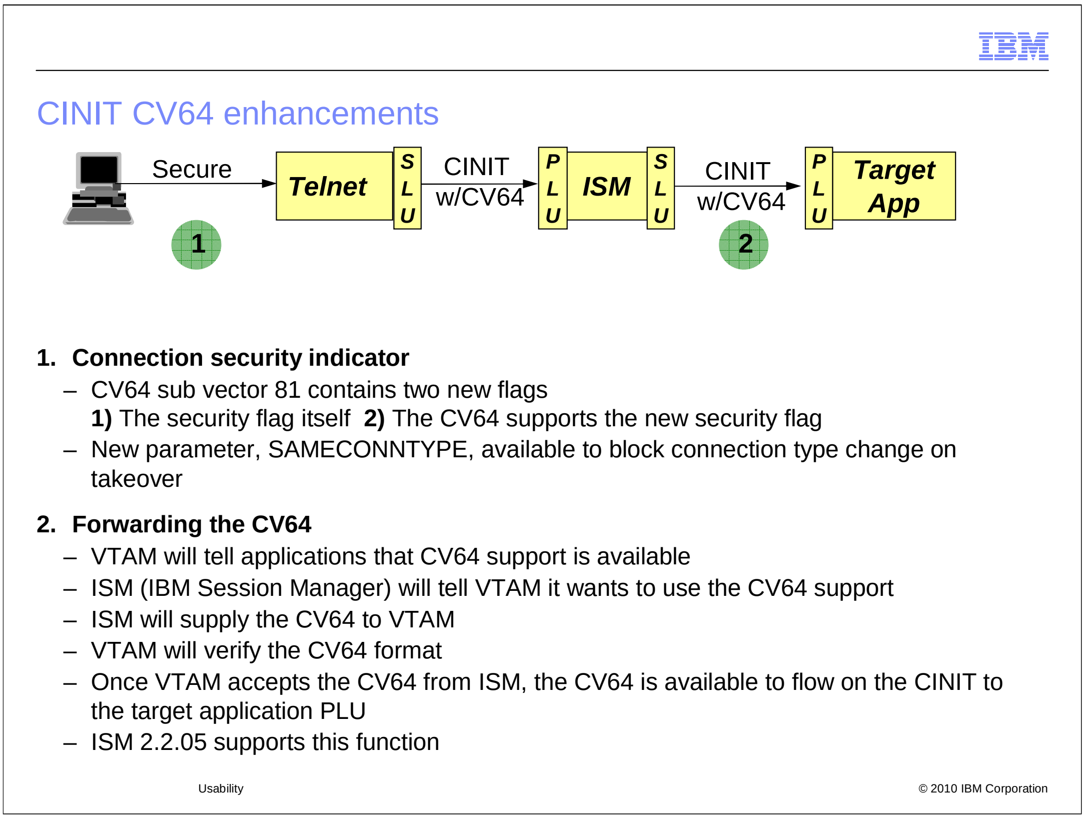
IBM
Secure
1
Telnet
SLU
CINIT w/CV64
PLU
ISM
SLU
CINIT w/CV64
2
PLU
Target App

CINIT CV64 enhancements

1. **Connection security indicator**
   - CV64 sub vector 81 contains two new flags
     1) The security flag itself **2)** The CV64 supports the new security flag
   - New parameter, SAMECONNTYPE, available to block connection type change on takeover
2. **Forwarding the CV64**
   - VTAM will tell applications that CV64 support is available
   - ISM (IBM Session Manager) will tell VTAM it wants to use the CV64 support
   - ISM will supply the CV64 to VTAM
   - VTAM will verify the CV64 format
   - Once VTAM accepts the CV64 from ISM, the CV64 is available to flow on the CINIT to the target application PLU
   - ISM 2.2.05 supports this function

Usability

To provide security information, a flag that indicates whether the TCP/IP connection is secure has been added to the CV64. An additional flag must be checked to ensure the new CV64 security flag is being set properly. You must first check this flag to verify the new function is supported on the CV64 you are processing before you check the flag to determine the security level of the TCP/IP connection. If a telnet connection is taken over with the reconnect option by another connection, the type of connection can be identical or at a higher level of security. Today it is possible for a secure telnet connection to take over a basic connection. If this occurs the CV64 information for the connection is not correct. A new session is not established and no new CINIT is sent to the target application. The application will continue to have security information relating to the original connection and not the new connection. A new parameter, SAMECONNTYPE, can be specified indicating the taker connection must be the same CONNTYPE as the target connection.

To solve the CV64 forwarding problem, VTAM has added API support that allows an application to provide a CV64 to associate with a session. When the application OPENs its ACB, VTAM sets a bit to indicate that the support is available on this level of VTAM. Also, during the OPEN ACB, the application tells VTAM that it is exploiting the new function. Later, during SETLOGON processing, the application provides the CV64 to VTAM. VTAM checks the CV64 to make sure it is built correctly. Then, VTAM stores the CV64 so that later it can append it to a CINIT for a session with the application. The CINIT being delivered to the target application PLU has a CV64 with it. ISM 2.2.05, which was available in March 2010, supports this new function. There are no known plans for other session managers to exploit this function at this time.

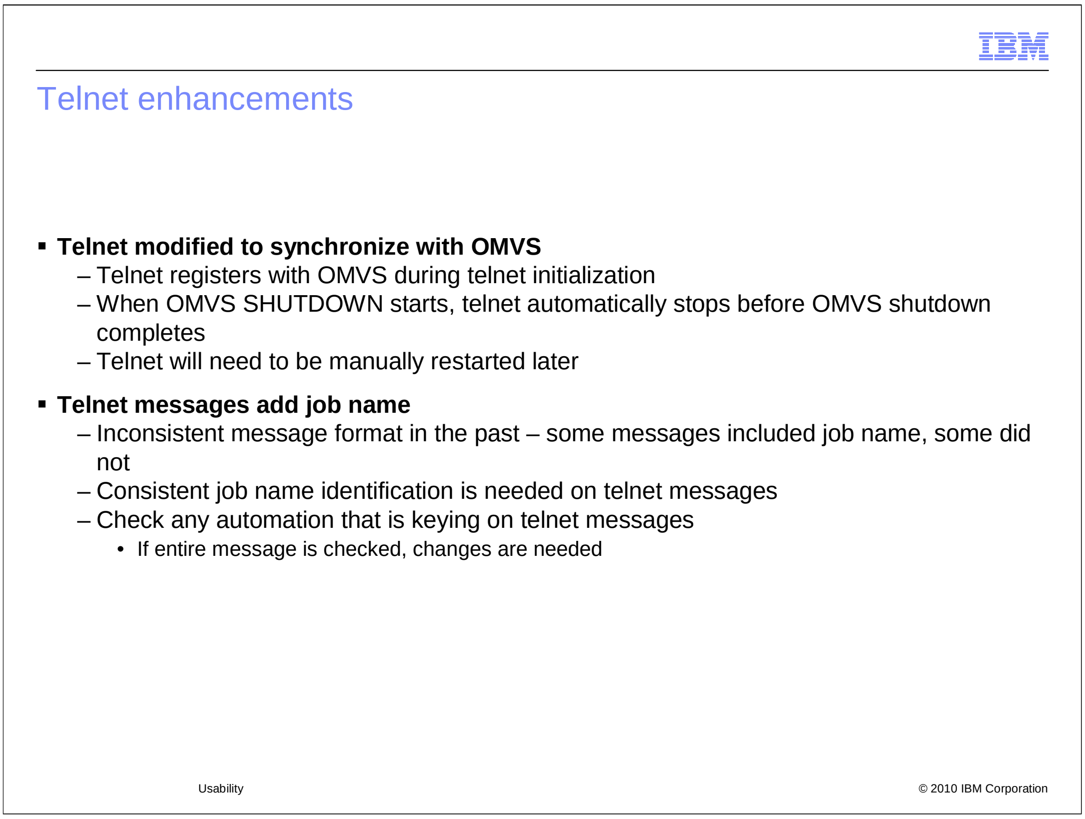

Telnet enhancements

- **Telnet modified to synchronize with OMVS**
  - Telnet registers with OMVS during telnet initialization
  - When OMVS SHUTDOWN starts, telnet automatically stops before OMVS shutdown completes
  - Telnet will need to be manually restarted later
- **Telnet messages add job name**
  - Inconsistent message format in the past – some messages included job name, some did not
  - Consistent job name identification is needed on telnet messages
  - Check any automation that is keying on telnet messages
    - If entire message is checked, changes are needed

Usability

© 2010 IBM Corporation

There are times when OMVS needs to be shut down and restarted to implement a change. Because telnet uses UNIX System Services, when OMVS is shut down telnet should be shut down first. Otherwise telnet might abnormally terminate and not be useable when OMVS is restarted.

Telnet now registers with OMVS. When OMVS shutdown is issued, OMVS calls a telnet routine that will shut down telnet. OMVS waits for telnet to stop. The new telnet exit drives the same process as if an operator issued a stop telnet console command. Once telnet is stopped, OMVS can continue its shutdown. The TCP/IP stack does not automatically shut down. It remains active and blocks OMVS shutdown. Don’t forget to stop the TCP/IP stack. Telnet, and the stack, will need to be restarted manually later.

Some customers run multiple telnet servers for potential recovery scenarios. Since LUNS/LUNR function was introduced, more configurations will run multiple telnet servers on the same system. Close to 50 telnet messages do not include the job name of the issuing telnet server. All messages use “TELNET” as the first word. When the LUNS/LUNR function was introduced, job name identification on telnet messages became more important and those messages used the second word of the message as the job name. Customers using multiple telnet servers need a way to determine, for all messages, which telnet server issued the message.

In z/OS V1R12, the first word is now always the job name. Every telnet message has changed in some way. If you have automation, you need to verify that it will still behave correctly given the changes made to the telnet messages. If your automation checks values by position, most messages are OK because the word TELNET was replaced with the job name. The LUNS/LUNR messages do change, however, because the TELNET word was removed and the already existing job name now becomes the first word causing all other words to shift. If your automation checks the entire message, the automation will have to be changed to understand the new format.

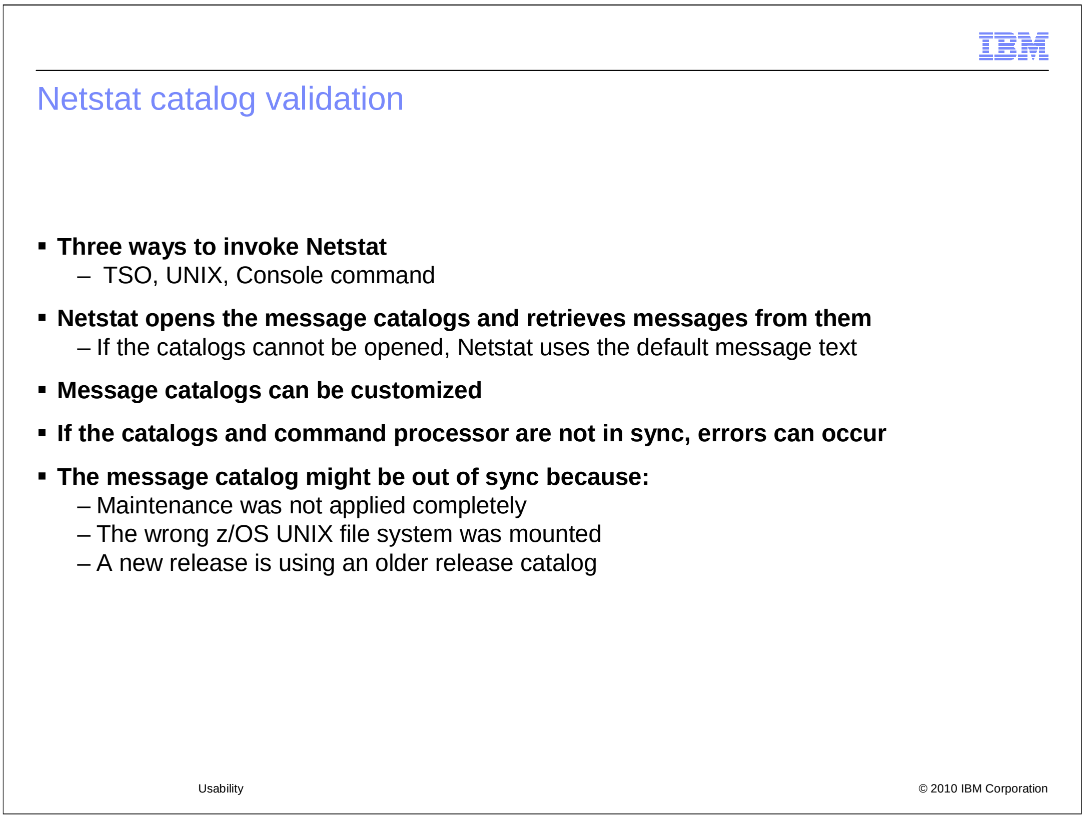

Netstat catalog validation

- **Three ways to invoke Netstat**
  - TSO, UNIX, Console command
- **Netstat opens the message catalogs and retrieves messages from them**
  - If the catalogs cannot be opened, Netstat uses the default message text
- **Message catalogs can be customized**
- **If the catalogs and command processor are not in sync, errors can occur**
- **The message catalog might be out of sync because:**
  - Maintenance was not applied completely
  - The wrong z/OS UNIX file system was mounted
  - A new release is using an older release catalog

Usability

There are three methods for invoking Netstat; from TSO, from UNIX, and from the MVS console. Netstat uses message catalogs, which are installed in the z/OS UNIX file system directory, /usr/lpp/tcpip/lib/nls/msg/C. These message catalogs contain all the messages Netstat will issue. A message catalog can be customized to change the text or order of variables in the messages. Netstat uses two message catalogs, netmsg.cat (for IPv4 messages) and netmsg6.cat (for IPv6 messages). When Netstat is invoked, it will try to open both message catalogs. If a message catalog cannot be opened, default messages are used. The default message text is included in the Netstat command.

Since Netstat retrieves messages from the catalog and then inserts variable text into the messages, various errors can occur if the message catalog is not at the same level as Netstat expects. Netstat might abnormally end while trying to issue the message, possibly causing the configuration component to terminate. Netstat might also issue messages indicating the message is reserved or might use the wrong message. This can leave Netstat unusable.

The message catalog might be out of sync for various reasons. Maintenance might not be applied to both the Netstat program and the message catalog. The wrong z/OS UNIX file system might be mounted. The message catalogs might have been modified in one release and not updated when a new release is installed.

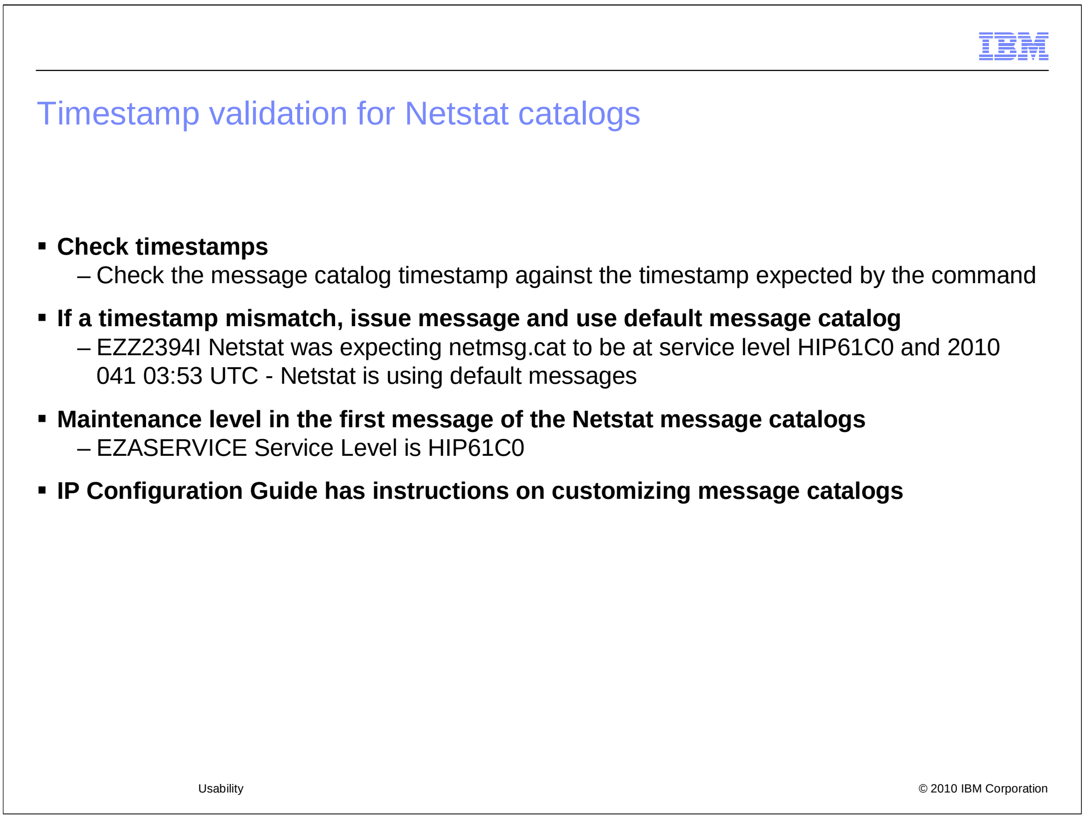

Timestamp validation for Netstat catalogs

- **Check timestamps**
  - Check the message catalog timestamp against the timestamp expected by the command
- **If a timestamp mismatch, issue message and use default message catalog**
  - EZZ2394I Netstat was expecting netmsg.cat to be at service level HIP61C0 and 2010 041 03:53 UTC - Netstat is using default messages
- **Maintenance level in the first message of the Netstat message catalogs**
  - EZASERVICE Service Level is HIP61C0
- **IP Configuration Guide has instructions on customizing message catalogs**

To avoid the Netstat catalog mismatch problem, Netstat will now verify the timestamp of the message catalog against the timestamp that Netstat expects. If the timestamps do not match, new message EZZ2394I is issued which contains the message catalog Netstat tried to open, the service level Netstat expected and the timestamp Netstat expected. Netstat will use the default messages for that invocation of the command. This logic is very similar to logic added to Omproute and the FTP client and server. Netstat message catalogs also have been updated to include the service level in the first message of the message catalog. The service level message starts with EZASERVICE and contains the FMID or PTF level of the message catalog. This message is never displayed by the Netstat command but can be viewed by browsing the message catalog.

The IP Configuration Guide has instructions for including the timestamp and general information on customizing the message catalogs.

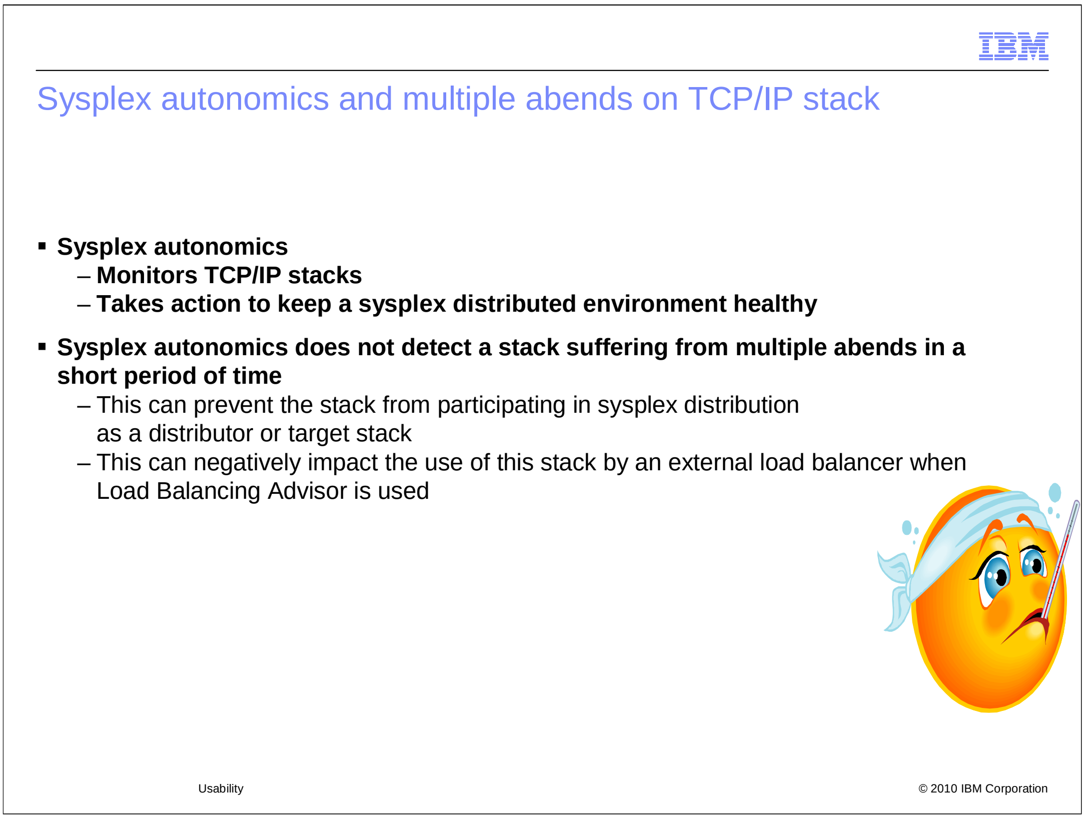

Sysplex autonomics and multiple abends on TCP/IP stack

- **Sysplex autonomics**
  - **Monitors TCP/IP stacks**
  - **Takes action to keep a sysplex distributed environment healthy**
- **Sysplex autonomics does not detect a stack suffering from multiple abends in a short period of time**
  - This can prevent the stack from participating in sysplex distribution as a distributor or target stack
  - This can negatively impact the use of this stack by an external load balancer when Load Balancing Advisor is used

Usability

© 2010 IBM Corporation

Sysplex autonomics was designed to detect issues with a TCP/IP stack that will prevent it from serving as a sysplex distributor or valid target node. Storage constraints, lock contention, VTAM inactive, and no dynamic routing capability (OMPROUTE inactive) are all reasons that the TCP/IP stack cannot participate in sysplex networking operations. When such conditions are encountered, the TCP/IP stack can remove itself from the TCP/IP sysplex group. If the TCP/IP stack is serving as a sysplex distributor, distributor responsibilities are passed to a backup TCP/IP stack. If the TCP/IP stack is serving as a target for distributed connections, that TCP/IP stack is no longer eligible to receive new connection requests.

Sysplex autonomics currently does not detect a stack suffering from multiple abends in a short period of time. The abends consume a large portion of the stacks resources which can prevent the stack from serving as a sysplex distributor or target node.

Additionally, a stack experiencing multiple abends in a short period of time is a poor choice for use by an external load balancer.

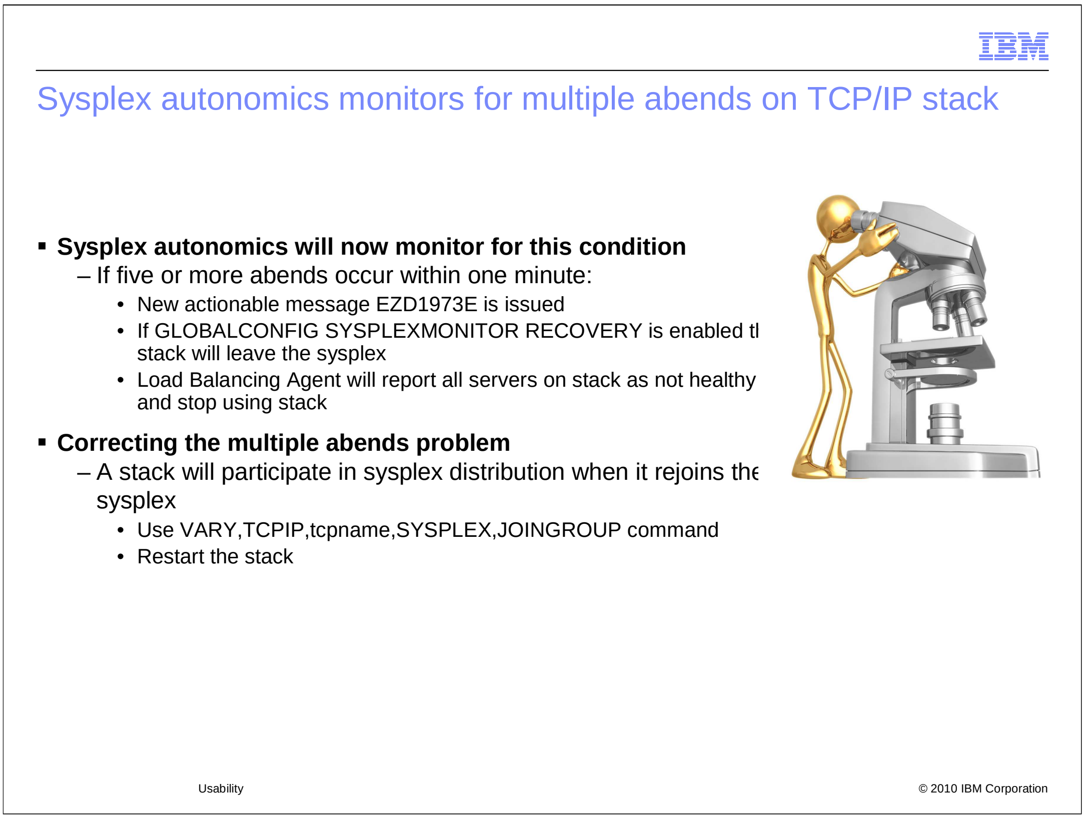

Sysplex autonomics monitors for multiple abends on TCP/IP stack

- **Sysplex autonomics will now monitor for this condition**
  - If five or more abends occur within one minute:
    - New actionable message EZD1973E is issued
    - If GLOBALCONFIG SYSPLEXMONITOR RECOVERY is enabled tl stack will leave the sysplex
    - Load Balancing Agent will report all servers on stack as not healthy and stop using stack
- **Correcting the multiple abends problem**
  - A stack will participate in sysplex distribution when it rejoins the sysplex
    - Use VARY,TCPIP,tcpname,SYSPLEX,JOINGROUP command
    - Restart the stack

Usability

© 2010 IBM Corporation

Sysplex autonomics will now monitor stacks for multiple abends in a short time. There is nothing new that needs to be configured for this new monitoring. If a stack suffers five or more abends within one minute these actions are taken. First, Message EZD1973E is issued indicating multiple nonrecoverable errors are adversely affecting sysplex processing. Second, if SYSPLEXMONITOR RECOVERY is enabled in GLOBALCONFIG the stack is removed from the sysplex. Third, the stack’s load balancing agent will stop reporting this stack to the load balancing advisor preventing its use by an external load balancer.

After a stack experiences the multiple abends problem, there are several ways the stack can rejoin the sysplex distributed or external load balancer environment. If the stack was removed from the sysplex by sysplex autonomics you can manually rejoin the sysplex. Use the VARY,TCPIP,tcpname,SYSPLEX,JOINGROUP command. You can restart the stack. If the stack suffers no further abends in a ten minute period it will rejoin the external load balancer environment. Servers are reported as healthy to an external load balancer in one of two ways. After ten minutes with no further abends, the Load Balancing Agent will report this stack’s servers as healthy. The second way is when a stack is restarted, the agent will immediately report this stack’s servers as healthy.

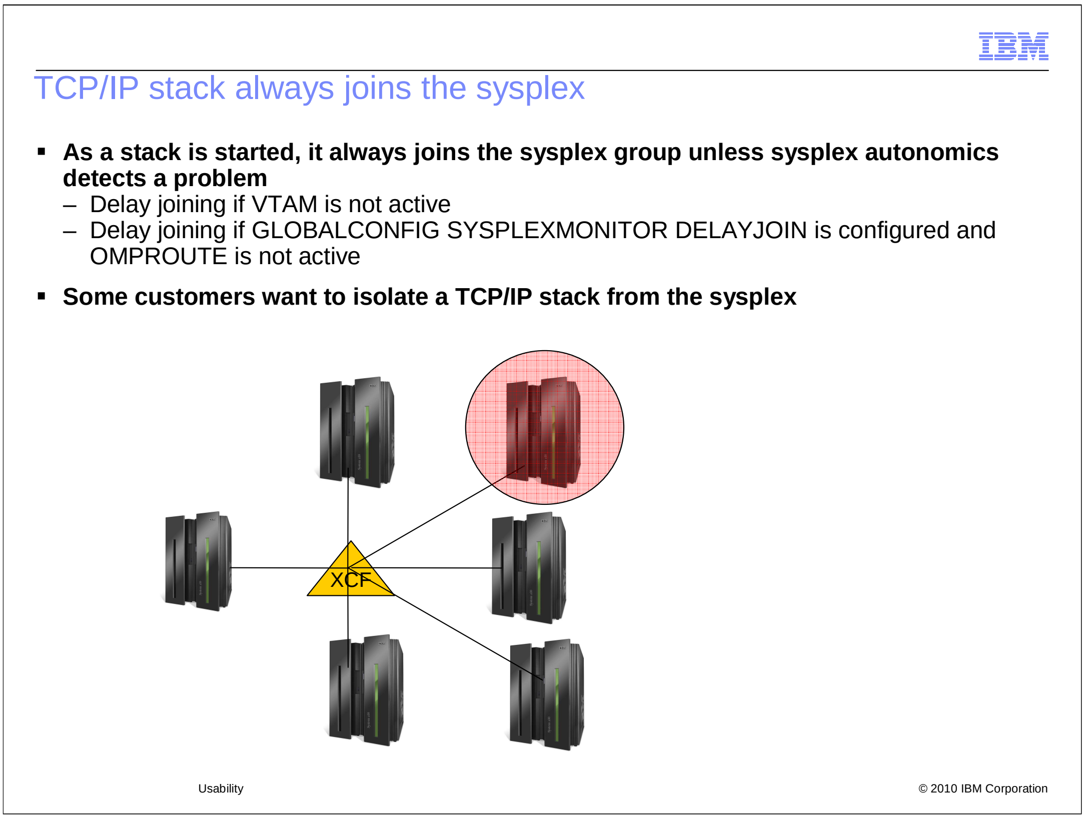

TCP/IP stack always joins the sysplex

- **As a stack is started, it always joins the sysplex group unless sysplex autonomics detects a problem**
  - Delay joining if VTAM is not active
  - Delay joining if GLOBALCONFIG SYSPLEXMONITOR DELAYJOIN is configured and OMPROUTE is not active
- **Some customers want to isolate a TCP/IP stack from the sysplex**

Usability

As a stack is started, it always joins the sysplex group unless sysplex autonomics detects a problem such as VTAM not active or OMPROUTE is not active and SYSPLEXMONITOR DELAYJOIN is configured in GLOBALCONFIG.

Some customers want to isolate a TCP/IP stack from other stacks in a sysplex.

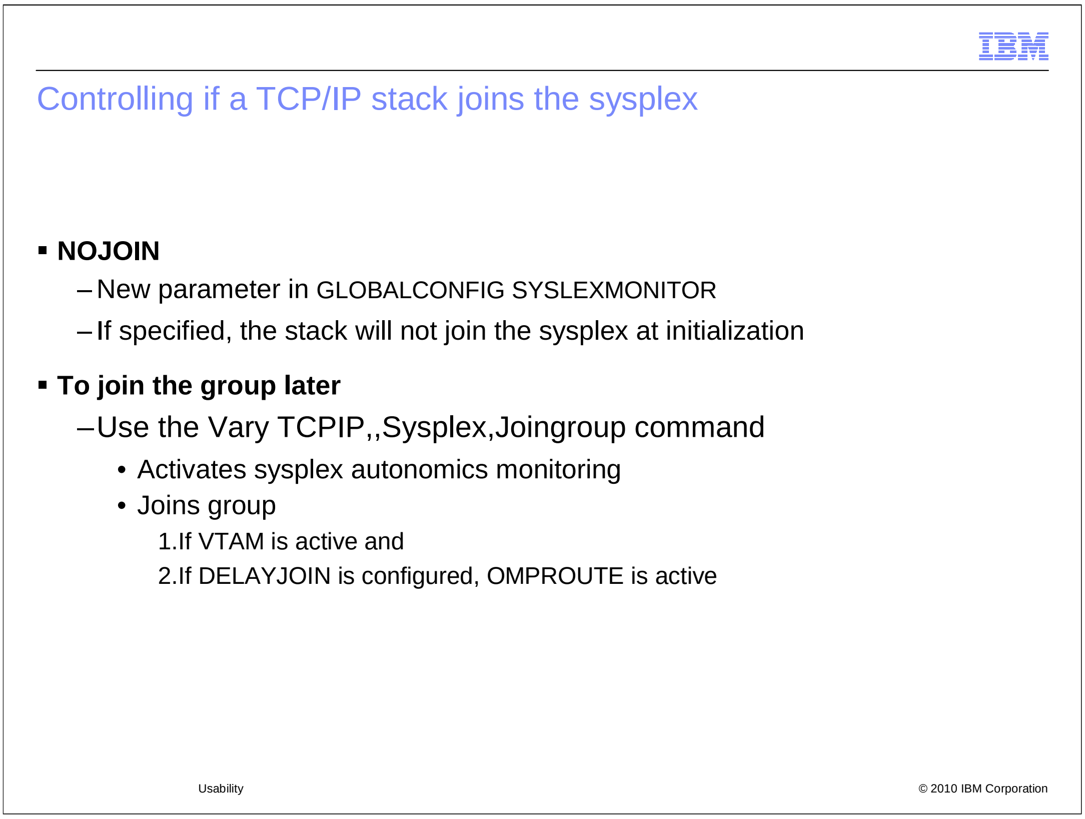

Controlling if a TCP/IP stack joins the sysplex

- **NOJOIN**
  - New parameter in GLOBALCONFIG SYSLEXMONITOR
  - If specified, the stack will not join the sysplex at initialization
- **To join the group later**
  - Use the Vary TCPIP,,Sysplex,Joingroup command
    - Activates sysplex autonomics monitoring
    - Joins group
      1. If VTAM is active and
      2. If DELAYJOIN is configured, OMPROUTE is active

A new configuration parameter in GLOBALCONFIG SYSLEXMONITOR called NOJOIN will control if a stack joins the sysplex at initialization. If it is in the initial profile, the stack will not join the sysplex. If the parameter is not coded, you get existing behavior. The stack joins the sysplex group if sysplex autonomics allows it.

The existing command, Vary TCPIP,,Sysplex,Joingroup can override this parameter. The command will not override existing Sysplex autonomics problem detection functions and parameters. The stack will not join the sysplex group until VTAM is active. Remaining SYSPLEXMONITOR parameters are activated as the command is issued; if DELAYJOIN is configured, the stack will not join the sysplex group until OMPROUTE is active.

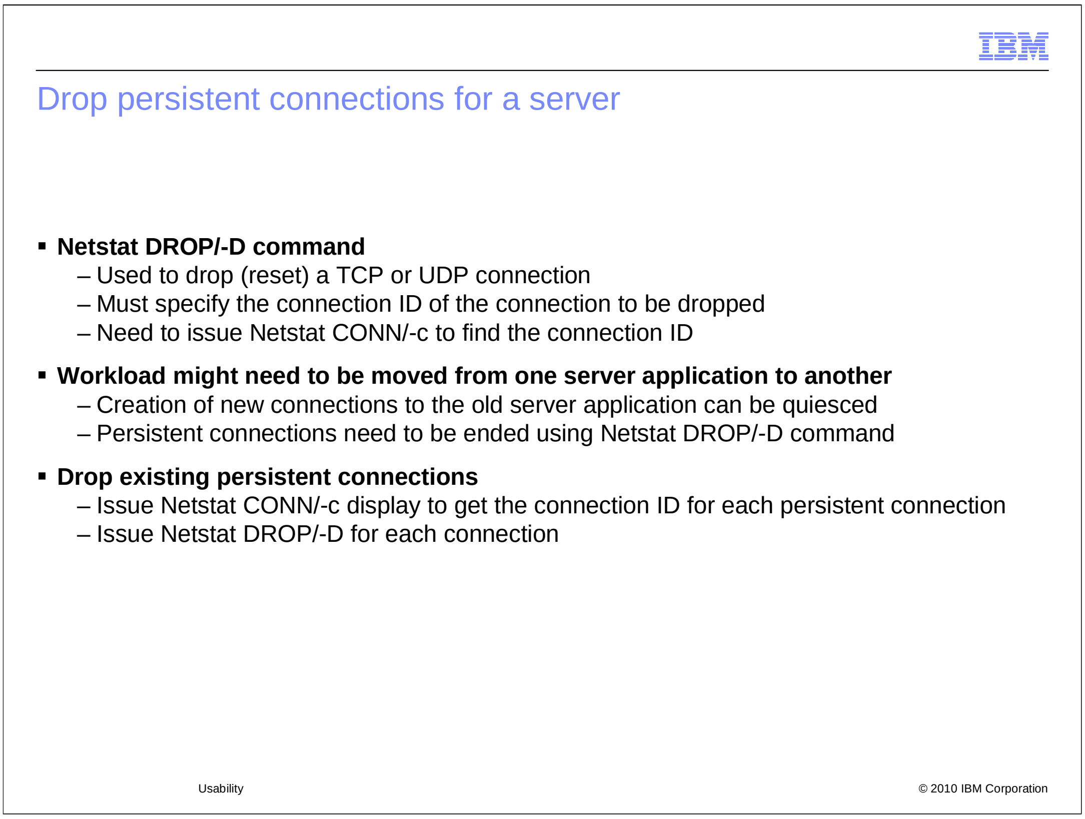

Drop persistent connections for a server

- **Netstat DROP/-D command**
  - Used to drop (reset) a TCP or UDP connection
  - Must specify the connection ID of the connection to be dropped
  - Need to issue Netstat CONN/-c to find the connection ID
- **Workload might need to be moved from one server application to another**
  - Creation of new connections to the old server application can be quiesced
  - Persistent connections need to be ended using Netstat DROP/-D command
- **Drop existing persistent connections**
  - Issue Netstat CONN/-c display to get the connection ID for each persistent connection
  - Issue Netstat DROP/-D for each connection

Usability

The Netstat DROP/-D command can be used to drop (reset) a TCP or UDP connection. This includes the VARY TCPIP,,DROP MVS command, the TSO NETSTAT DROP command and the z/OS UNIX netstat –D command. The connection ID of the connection to be dropped must be specified. Generally, users need to issue a Netstat CONN/-c command to find the appropriate connection ID.

If you want to move workload from one server application to another you can quiesce the creation of new connections to the old server, but persistent connections need to be ended using the Netstat DROP/-D command.

You need to issue a Netstat CONN/-c display command to get the connection ID of each persistent session to be reset, and issue a Netstat DROP/-D command for each connection. If a server has dozens of persistent connections, this can be tedious.

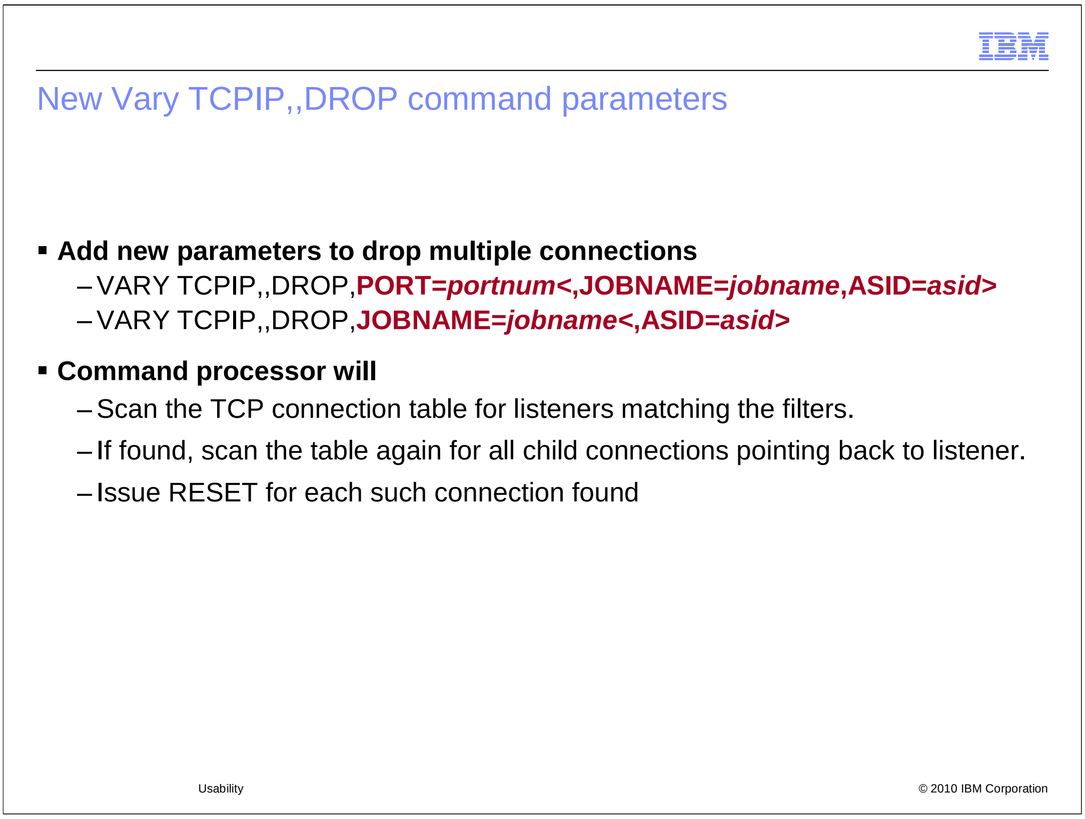

New Vary TCPIP,,DROP command parameters

- **Add new parameters to drop multiple connections**
  - VARY TCPIP,,DROP,**PORT=*portnum*<,JOBNAME=*jobname*,ASID=*asid*>**
  - VARY TCPIP,,DROP,**JOBNAME=*jobname*<,ASID=*asid*>**
- **Command processor will**
  - Scan the TCP connection table for listeners matching the filters.
  - If found, scan the table again for all child connections pointing back to listener.
  - Issue RESET for each such connection found

Usability

To address this problem, the existing VARY TCPIP,,DROP command is extended with new parameters to allow all TCP connections associated with a server matching the specified filter to be reset. The structure of the new parameters are modeled after the parameters of the existing VARY TCPIP,,SYSPLEX,QUIESCE command. You can specify the job name and optionally the address space ID or the port number and optionally job name and address space ID for the servers. The command processor will scan the TCP connection table for listeners matching the supplied filters. If a match is found, it will scan the table again for all child connections associated with that listener. For each one found, it will reset the connection.

If more than one server application is found to match the input filter values, the command is failed. You can re-issue the command specifying additional filter parameters to identify a specific server application.

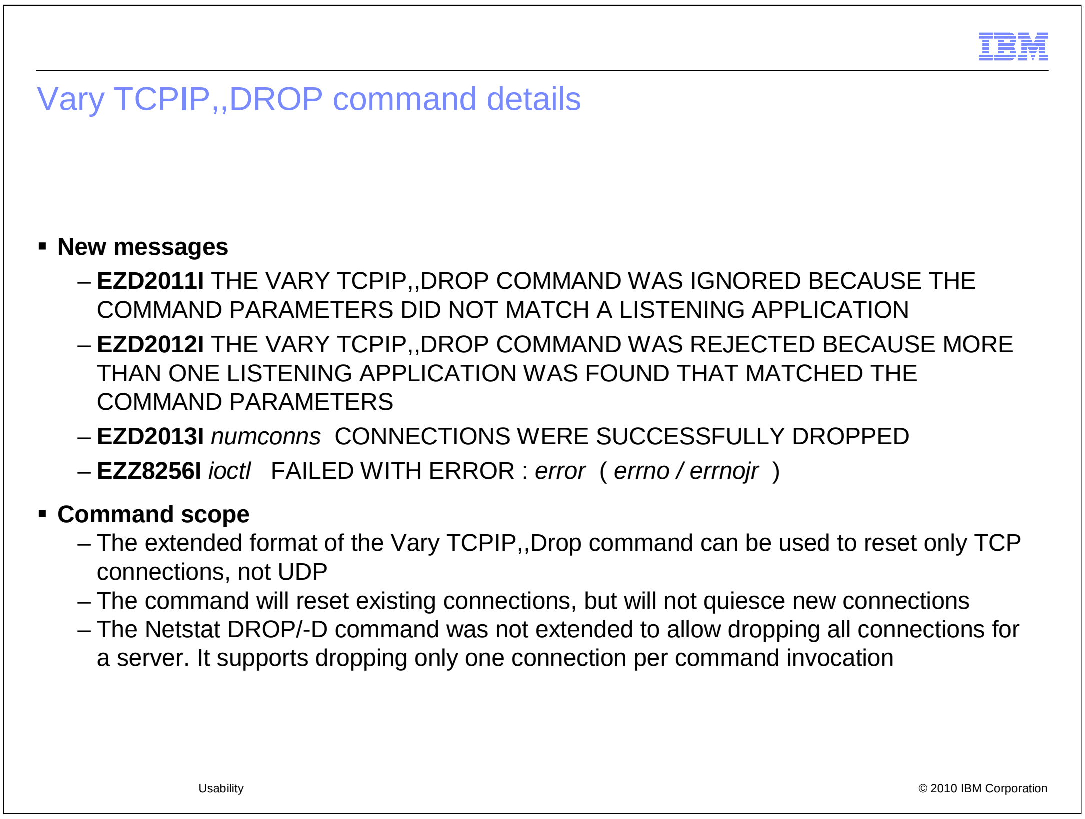

Vary TCPIP,,DROP command details

- **New messages**
  - **EZD2011I** THE VARY TCPIP,,DROP COMMAND WAS IGNORED BECAUSE THE COMMAND PARAMETERS DID NOT MATCH A LISTENING APPLICATION
  - **EZD2012I** THE VARY TCPIP,,DROP COMMAND WAS REJECTED BECAUSE MORE THAN ONE LISTENING APPLICATION WAS FOUND THAT MATCHED THE COMMAND PARAMETERS
  - **EZD2013I** *numconns* CONNECTIONS WERE SUCCESSFULLY DROPPED
  - **EZZ8256I** *ioctl* FAILED WITH ERROR : *error* (*errno / errnojr*)
- **Command scope**
  - The extended format of the Vary TCPIP,,Drop command can be used to reset only TCP connections, not UDP
  - The command will reset existing connections, but will not quiesce new connections
  - The Netstat DROP/-D command was not extended to allow dropping all connections for a server. It supports dropping only one connection per command invocation

Usability

There are new messages that might be issued from the Vary Drop command processor in connection with the new function. EZD2011I is issued if no listening server application matching the input parameters can be found. EZD2012I is issued if more than one listening application is found that matches the input parameters. EZD2013I is issued when all the connections associated with the server application that matched the input parameters have been reset, and includes the number of connections that were reset. If an error occurs while processing the IOCTL for this function, EZZ8256I is issued with an indication of the error and the errorno and errnojr numbers.

Only TCP connections associated with the server application matching the supplied parameters are dropped. UDP connections are not affected.
Existing TCP connections are reset by this command, but new connection requests are not quiesced. New connections might or might not be reset depending on the timing of when the connections are established. You might want to quiesce new connection requests to the server application before issuing this command.

The Netstat DROP/-D command was not extended in this release, and can still reset only one connection per command invocation.

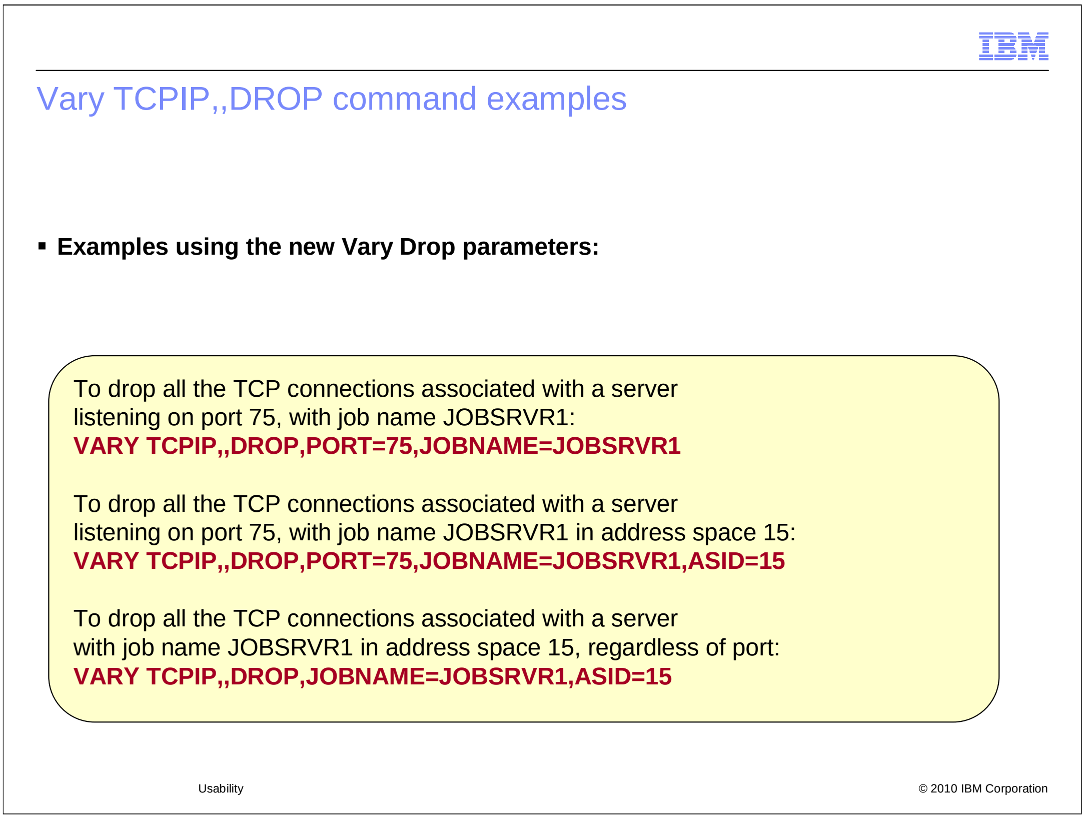

Vary TCPIP,,DROP command examples

- **Examples using the new Vary Drop parameters:**

To drop all the TCP connections associated with a server
listening on port 75, with job name JOBSRVR1:

```
VARY TCPIP,,DROP,PORT=75,JOBNAME=JOBSRVR1
```

To drop all the TCP connections associated with a server
listening on port 75, with job name JOBSRVR1 in address space 15:

```
VARY TCPIP,,DROP,PORT=75,JOBNAME=JOBSRVR1,ASID=15
```

To drop all the TCP connections associated with a server
with job name JOBSRVR1 in address space 15, regardless of port:

```
VARY TCPIP,,DROP,JOBNAME=JOBSRVR1,ASID=15
```

Here are some examples of how to specify the new parameters on the Vary Drop command.

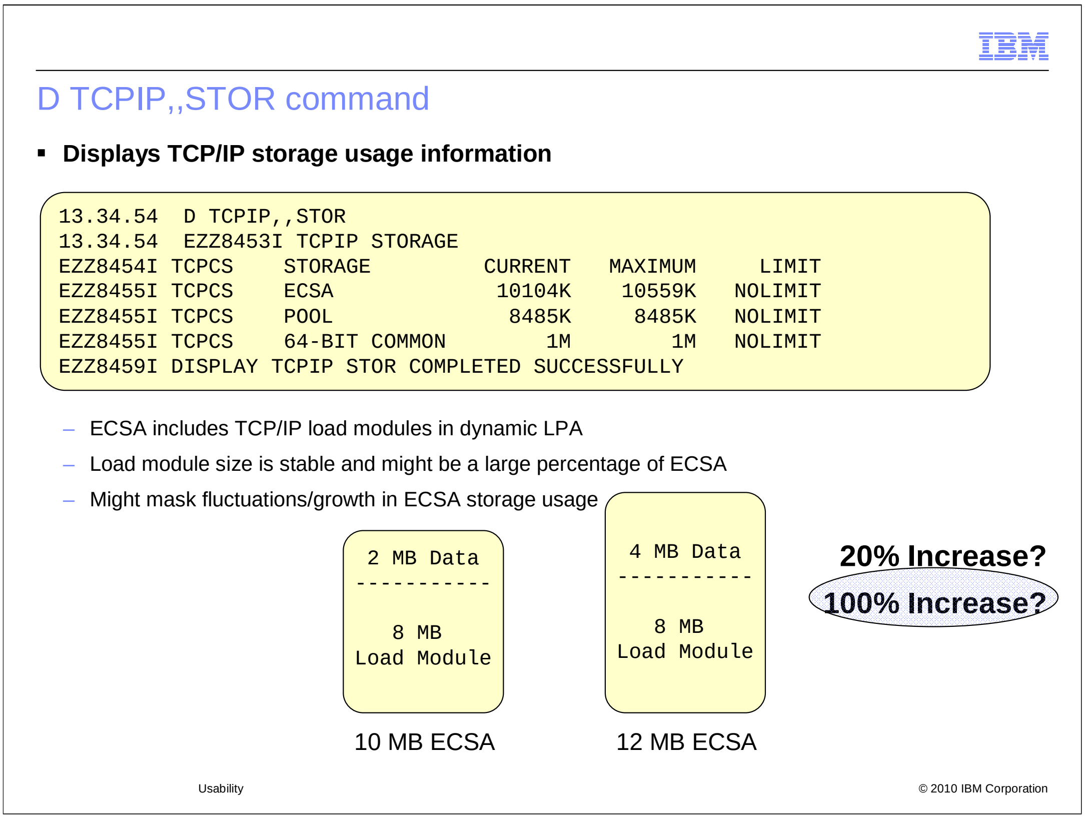

D TCPIP,,STOR command

- **Displays TCP/IP storage usage information**

```
13.34.54  D TCPIP,,STOR
13.34.54  EZZ8453I TCPIP STORAGE
EZZ8454I TCPCS    STORAGE         CURRENT   MAXIMUM     LIMIT
EZZ8455I TCPCS    ECSA             10104K    10559K   NOLIMIT
EZZ8455I TCPCS    POOL              8485K     8485K   NOLIMIT
EZZ8455I TCPCS    64-BIT COMMON        1M        1M   NOLIMIT
EZZ8459I DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY
```

- ECSA includes TCP/IP load modules in dynamic LPA
- Load module size is stable and might be a large percentage of ECSA
- Might mask fluctuations/growth in ECSA storage usage

2 MB Data
8 MB Load Module
10 MB ECSA

4 MB Data
8 MB Load Module
12 MB ECSA

20% Increase?
100% Increase?

Usability

© 2010 IBM Corporation

The “D TCPIP,,STOR” command displays information about the use of storage by z/OS Communications Server. The amounts of extended common storage (ECSA) in use, pooled private storage in use, and 64-bit common storage in use are displayed. The information is also available through the Network Management Interface using the GetStorageStatistics request.

The value displayed for ECSA storage includes the size of the TCP/IP load modules which are loaded using dynamic LPA functions. The size of these load modules is a stable value and might be a large percentage of the value displayed for common storage usage. This makes it difficult to recognize significant storage growth in common storage.

For example, assume the current ECSA usage value is 10 megabytes, of which eight megabytes is load module storage. That leaves two megabytes actually being used for control blocks. Assume the ECSA storage usage increases by two megabytes to 12 megabytes. Using the current display this looks like a 20 percent increase in ECSA storage usage. But it is actually a 100 percent increase in ECSA storage used for control blocks.

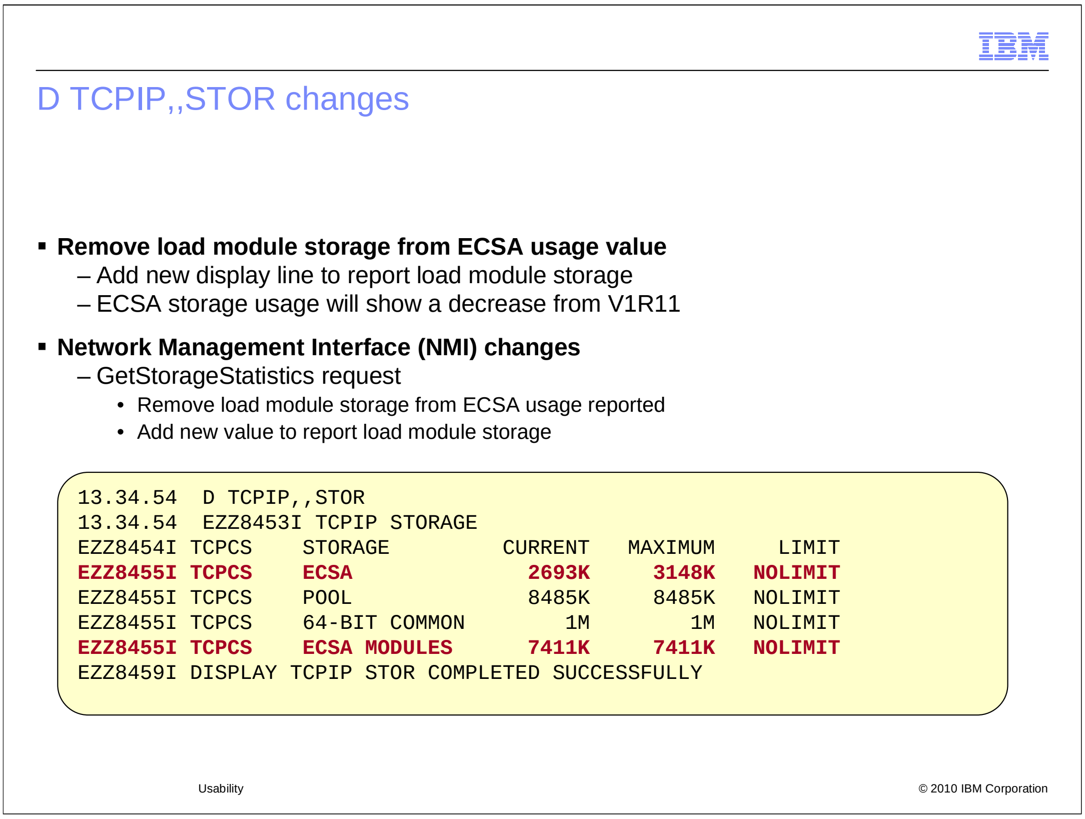

D TCPIP,,STOR changes

- **Remove load module storage from ECSA usage value**
  - – Add new display line to report load module storage
  - – ECSA storage usage will show a decrease from V1R11
- **Network Management Interface (NMI) changes**
  - – GetStorageStatistics request
    - Remove load module storage from ECSA usage reported
    - Add new value to report load module storage

```
13.34.54  D TCPIP,,STOR
13.34.54  EZZ8453I TCPIP STORAGE
EZZ8454I TCPCS    STORAGE          CURRENT   MAXIMUM     LIMIT
EZZ8455I TCPCS    ECSA               2693K     3148K   NOLIMIT
EZZ8455I TCPCS    POOL               8485K     8485K   NOLIMIT
EZZ8455I TCPCS    64-BIT COMMON         1M        1M   NOLIMIT
EZZ8455I TCPCS    ECSA MODULES       7411K     7411K   NOLIMIT
EZZ8459I DISPLAY TCPIP STOR COMPLETED SUCCESSFULLY
```

Usability

© 2010 IBM Corporation

The ECSA storage value is updated to reflect only the amount of storage used for control blocks and no longer includes the size of the load modules loaded into common storage. The display is updated to include a new line which shows the amount of common storage used for load modules. For similar workloads, the value of common storage usage in V1R12 is reduced from V1R11 since the value will no longer include the amount of storage used for load modules.

The NMI is also updated to remove the load module storage from the common storage value and to add a new value for load modules in common storage.

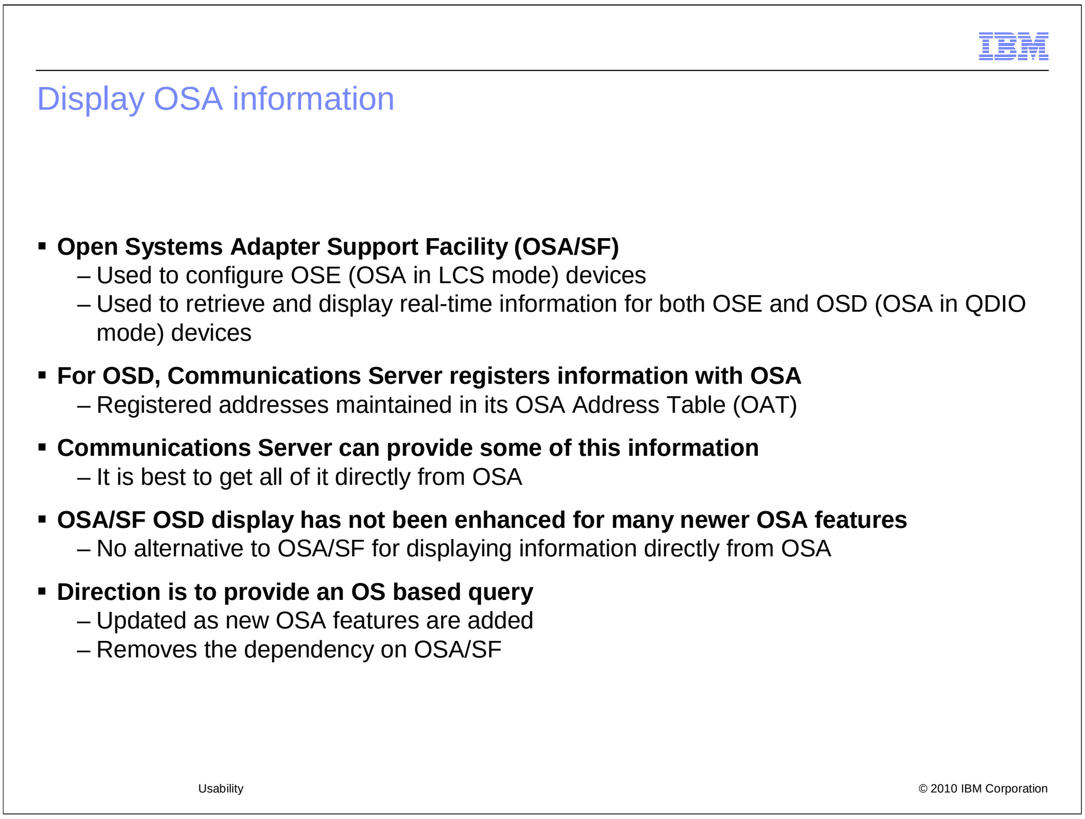

Display OSA information

- **Open Systems Adapter Support Facility (OSA/SF)**
  - Used to configure OSE (OSA in LCS mode) devices
  - Used to retrieve and display real-time information for both OSE and OSD (OSA in QDIO mode) devices
- **For OSD, Communications Server registers information with OSA**
  - Registered addresses maintained in its OSA Address Table (OAT)
- **Communications Server can provide some of this information**
  - It is best to get all of it directly from OSA
- **OSA/SF OSD display has not been enhanced for many newer OSA features**
  - No alternative to OSA/SF for displaying information directly from OSA
- **Direction is to provide an OS based query**
  - Updated as new OSA features are added
  - Removes the dependency on OSA/SF

Usability

OSA/SF has been used for years to configure OSA and display the configuration. OSA/SF has played a more central role for pre-QDIO OSE devices than for today's QDIO OSD devices. For OSD, Communications Server registers information (such as IP addresses) exclusively use Interprocess Architecture signals (IPAs) exchanged with the host to enable and configure features and register IP addresses to OSA.

The OSA/SF display does not support many of these newer OSA features and there is currently no other alternative to OSA/SF for displaying OSA information. The current direction is to provide an operating system based query to display OSA information that is updated when new OSA features are made available.

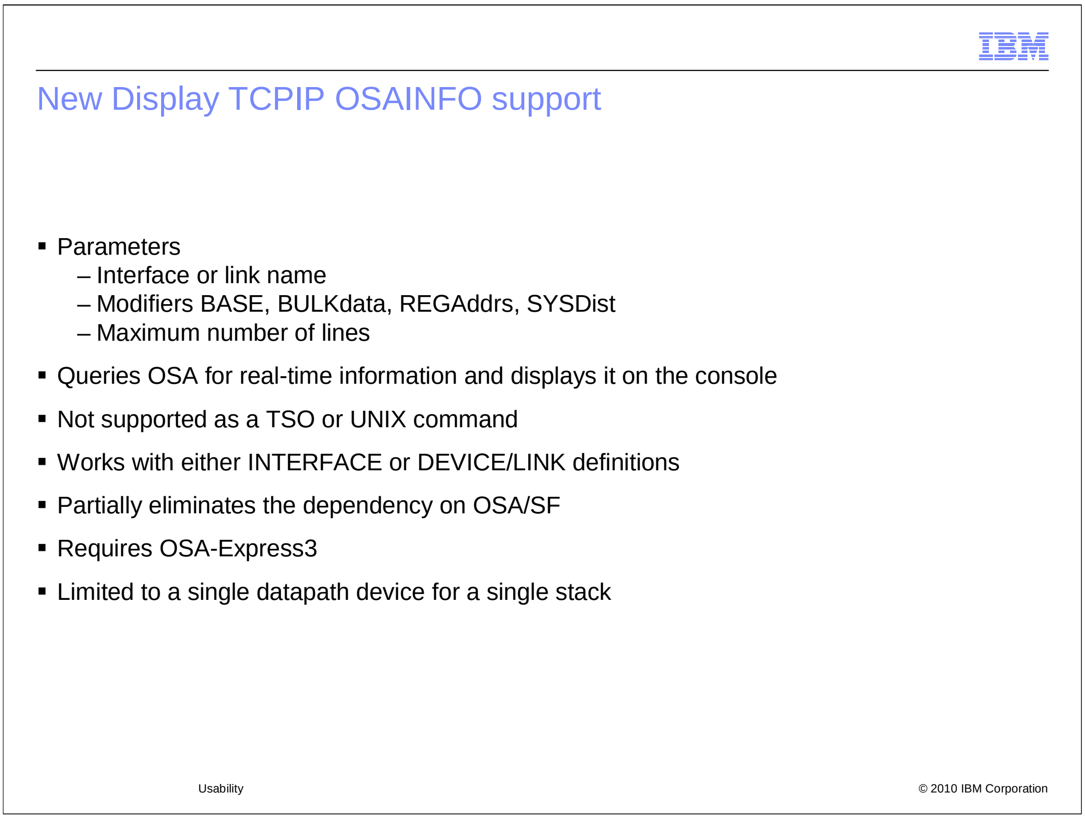

New Display TCPIP OSAINFO support

- Parameters
  - Interface or link name
  - Modifiers BASE, BULKdata, REGAddrs, SYSDist
  - Maximum number of lines
- Queries OSA for real-time information and displays it on the console
- Not supported as a TSO or UNIX command
- Works with either INTERFACE or DEVICE/LINK definitions
- Partially eliminates the dependency on OSA/SF
- Requires OSA-Express3
- Limited to a single datapath device for a single stack

Usability

A new console display, D,TCPIP,OSAinfo, provides the information missing from the OSA/SF command. The parameters are defined here.

The interface name can be specified as either an interface name or link name. The *Display OSAINFO* implementation is currently limited to a single datapath device (hence the required interface name parameter).

BASE, BULKdata, REGAddrs, and SYSDist are modifiers supplied to customize and limit output.

BULKdata and SYSDist are the queue types of QDIO inbound workload queuing and when specified, they request the registered routing variables (RVs) for those queue types. QDIO inbound workload queuing is covered in another presentation.

The maximum number of lines can optionally be specified to limit the amount of output to the console. By default 200 lines are displayed. All lines can be displayed.

The command is a console command that retrieves information from the OSA and then displays the results on the console. There are no TSO or UNIX equivalent commands. Because either INTERFACE or DEVICE/LINK statements define OSAs, the display command will show results for either definition choice. At this time the command is a partial replacement for OSA/SF because it will work only for OSA-Express3 and can display information for only a single datapath device for a single stack.

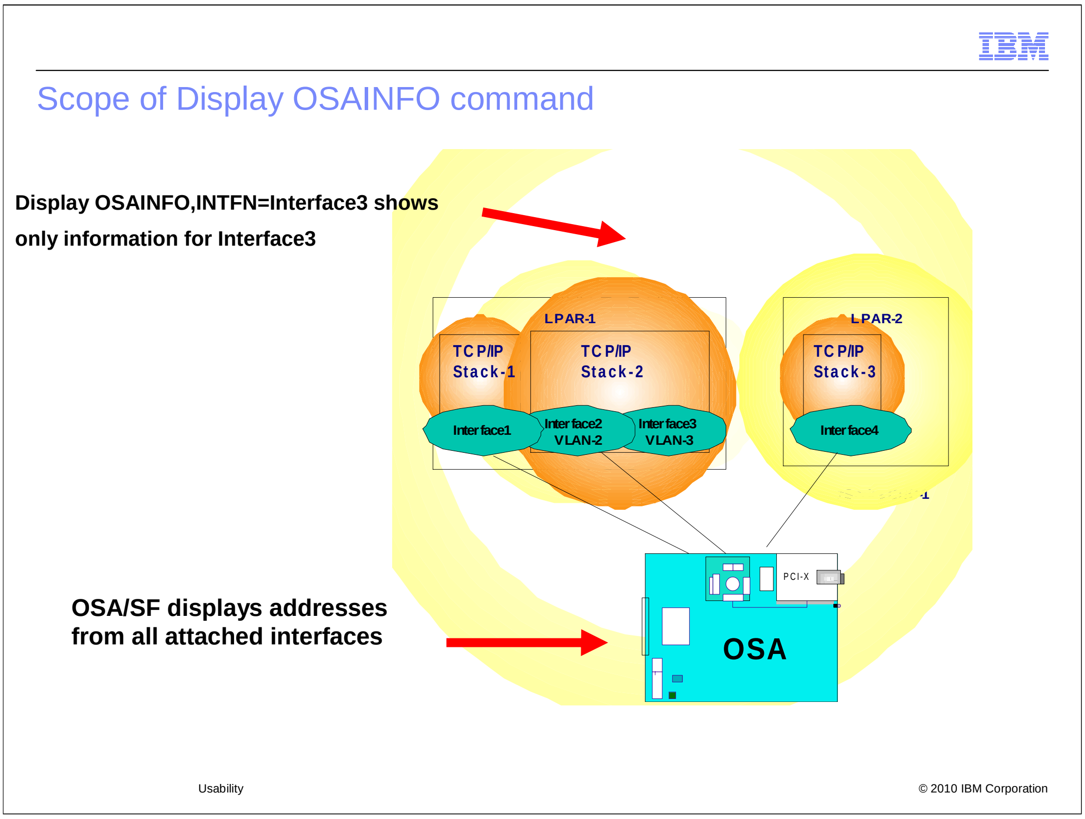

Scope of Display OSAINFO command

Display OSAINFO,INTFN=Interface3 shows only information for Interface3

OSA/SF displays addresses from all attached interfaces

Usability

© 2010 IBM Corporation

This diagram shows a multiple VLAN definition in Stack-2 and each VLAN is a separate interface.

The scope of a single *Display OSAINFO* command is limited to the specified interface on the stack on which the command was issued. Information for interfaces in other stacks cannot be queried unless the command is issued there.

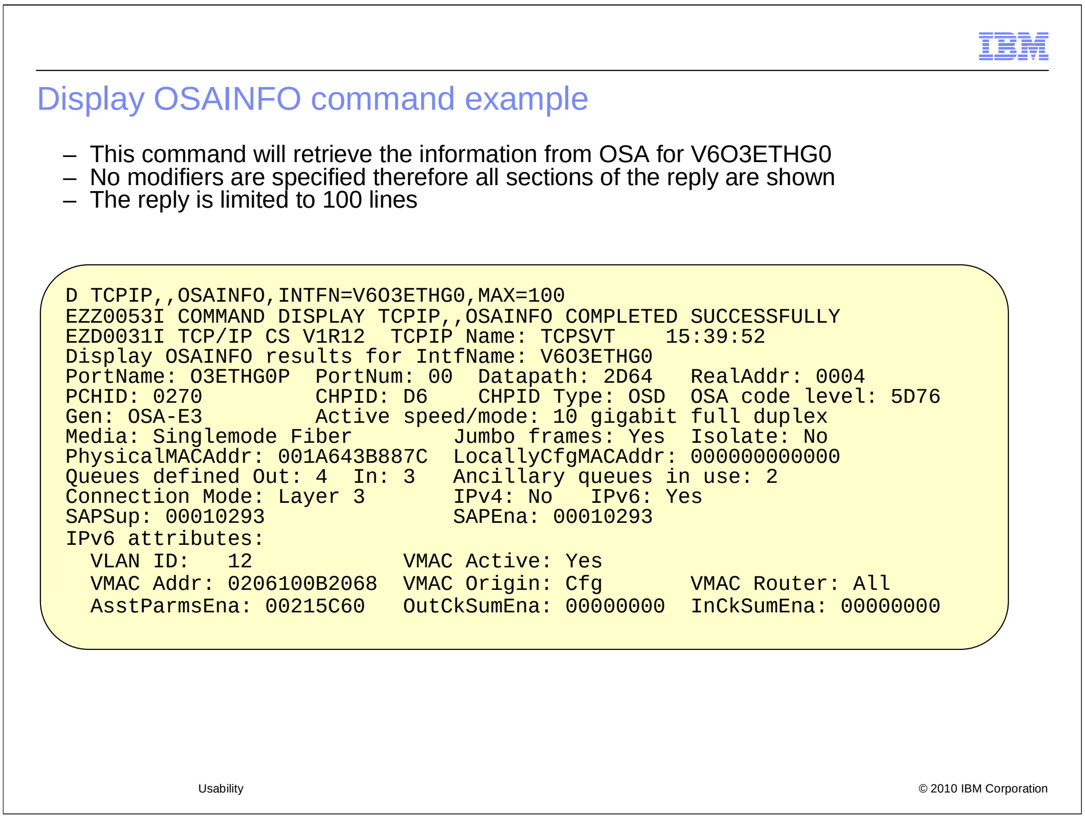

Display OSAINFO command example

– This command will retrieve the information from OSA for V6O3ETHG0
– No modifiers are specified therefore all sections of the reply are shown
– The reply is limited to 100 lines

```
D TCPIP,,OSAINFO,INTFN=V6O3ETHG0,MAX=100
EZZ0053I COMMAND DISPLAY TCPIP,,OSAINFO COMPLETED SUCCESSFULLY
EZD0031I TCP/IP CS V1R12  TCPIP Name: TCPSVT    15:39:52
Display OSAINFO results for IntfName: V6O3ETHG0
PortName: O3ETHG0P  PortNum: 00  Datapath: 2D64   RealAddr: 0004
PCHID: 0270         CHPID: D6    CHPID Type: OSD  OSA code level: 5D76
Gen: OSA-E3         Active speed/mode: 10 gigabit full duplex
Media: Singlemode Fiber        Jumbo frames: Yes  Isolate: No
PhysicalMACAddr: 001A643B887C  LocallyCfgMACAddr: 000000000000
Queues defined Out: 4  In: 3   Ancillary queues in use: 2
Connection Mode: Layer 3       IPv4: No   IPv6: Yes
SAPSup: 00010293               SAPEna: 00010293
IPv6 attributes:
  VLAN ID:   12            VMAC Active: Yes
  VMAC Addr: 0206100B2068  VMAC Origin: Cfg       VMAC Router: All
  AsstParmsEna: 00215C60   OutCkSumEna: 00000000  InCkSumEna: 00000000
```

Usability © 2010 IBM Corporation

This is an example of the display command. In this example V6O3ETHG0 can be either an INTERFACE or LINK name and the number of lines that are displayed is limited to 100.

This part of the sample reply is the start of the BASE section. The BASE section shows general information about the OSA such as the CHPID (in this sample the CHPID is D6).

All of the fields displayed in the reply are documented in z/OS Communications Server IP System Administrator's Commands Version 1 Release 12.

Message EZZ0053I is not part of the report but instead it is issued when the display command is *accepted*.

Message EZD0031I is the 1st message in the reply and is issued when all information has been received from OSA

This part of the sample reply is the end of the BASE section. This sample shows information about the IPv6 Layer 3 attributes such as the Global VLAN ID and VMAC information.

If the data device has IPv4 enabled (which this sample does not), the IPv4 Layer 3 attributes are displayed.

If the data device has IPv6 enabled (which this sample does), the IPv6 Layer 3 attributes are displayed.

If the data device has IPv4 and IPv6 enabled , the IPv4 Layer 3 attributes are displayed first, followed by the IPv6 Layer 3 attributes.

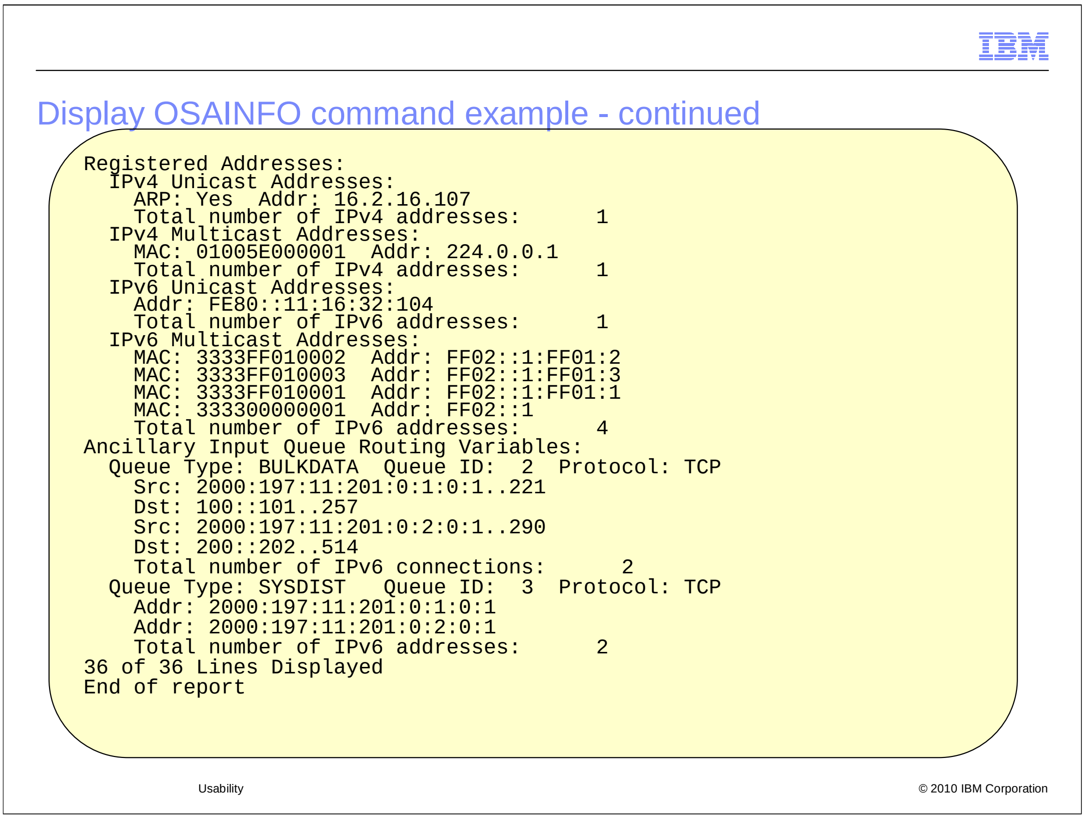

Display OSAINFO command example - continued

```
Registered Addresses:
  IPv4 Unicast Addresses:
    ARP: Yes  Addr: 16.2.16.107
    Total number of IPv4 addresses:      1
  IPv4 Multicast Addresses:
    MAC: 01005E000001  Addr: 224.0.0.1
    Total number of IPv4 addresses:      1
  IPv6 Unicast Addresses:
    Addr: FE80::11:16:32:104
    Total number of IPv6 addresses:      1
  IPv6 Multicast Addresses:
    MAC: 3333FF010002  Addr: FF02::1:FF01:2
    MAC: 3333FF010003  Addr: FF02::1:FF01:3
    MAC: 3333FF010001  Addr: FF02::1:FF01:1
    MAC: 333300000001  Addr: FF02::1
    Total number of IPv6 addresses:      4
Ancillary Input Queue Routing Variables:
  Queue Type: BULKDATA  Queue ID:  2  Protocol: TCP
    Src: 2000:197:11:201:0:1:0:1..221
    Dst: 100::101..257
    Src: 2000:197:11:201:0:2:0:1..290
    Dst: 200::202..514
    Total number of IPv6 connections:      2
  Queue Type: SYSDIST   Queue ID:  3  Protocol: TCP
    Addr: 2000:197:11:201:0:1:0:1
    Addr: 2000:197:11:201:0:2:0:1
    Total number of IPv6 addresses:      2
36 of 36 Lines Displayed
End of report
```

Usability

The Registered Addresses section displays all the IPv4 and IPv6 unicast and multicast addresses registered with the OSA. Note that the IPv4 information was inserted here for illustration purposes only. If the interface has IPv4 enabled, the IPv4 registered unicast and multicast addresses are displayed. The ARP field indicates if the OSA is performing ARP for an IPv4 unicast address. If the interface has IPv6 enabled, the IPv6 registered unicast and multicast addresses are displayed.

The Bulk Data section displays the source and destination IP addresses and ports of the TCP connections for which OSA is performing QDIO Inbound Workload Queuing for streaming connections. If the interface has QDIO Inbound Workload Queuing enabled for BULKDATA and there is at least one connection, the BULKDATA section is displayed. Note that you can see IPv4 or IPv6 addresses here but not both as QDIO Inbound Workload Queuing is not allowed when a single datapath device is used for both IPv4 and IPv6.

The Ancillary Input Queue section displays the destination IP address for which OSA is performing QDIO Inbound Workload Queuing for sysplex distributor. If the interface has QDIO Inbound Workload Queuing enabled for sysplex distributor and at least one destination address, the SYSDIST section is displayed. Note that you can see IPv4 or IPv6 addresses here but not both as Inbound Workload Queuing is not allowed when a single datapath device is used for both IPv4 and IPv6.

The total number of lines is displayed along with the total number of lines possible. The MAX parameter can be specified to limit the total number of lines displayed. If MAX=* is specified and more than 65,535 lines are required, the report is truncated. “Max lines limit reached” is displayed instead of the message with the counts.

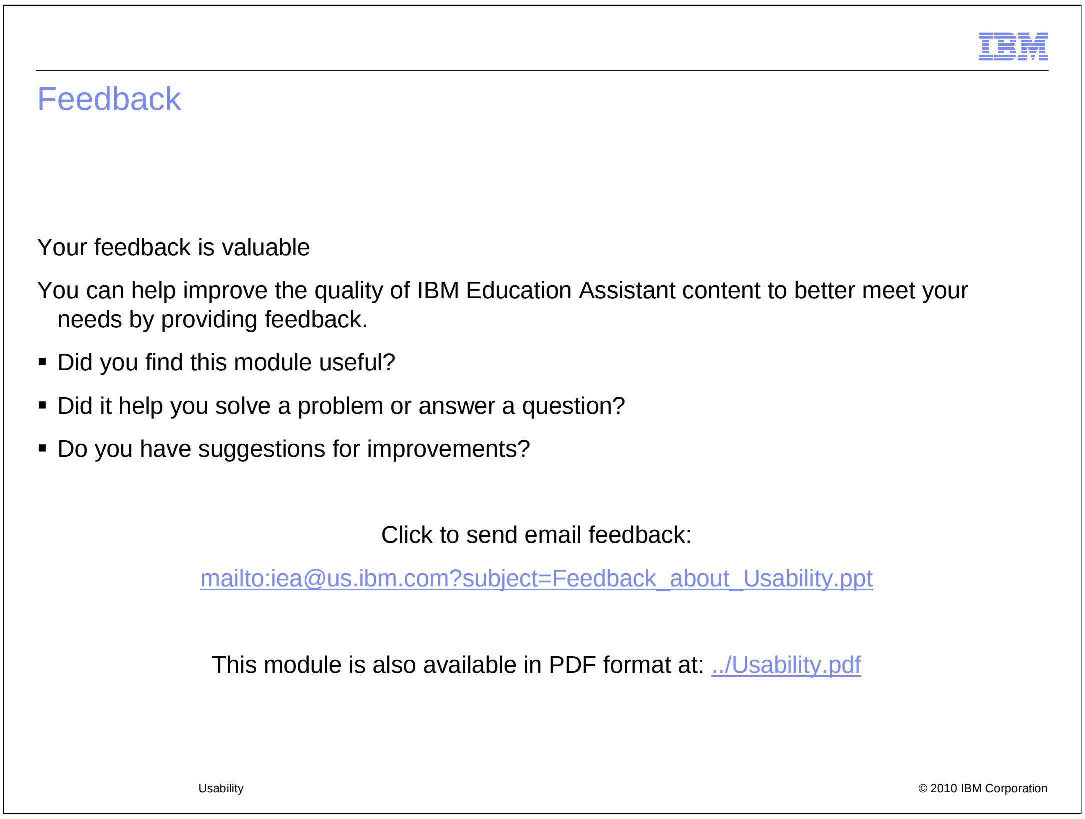

Feedback

Your feedback is valuable

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

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

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_Usability.ppt

This module is also available in PDF format at: ../Usability.pdf

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

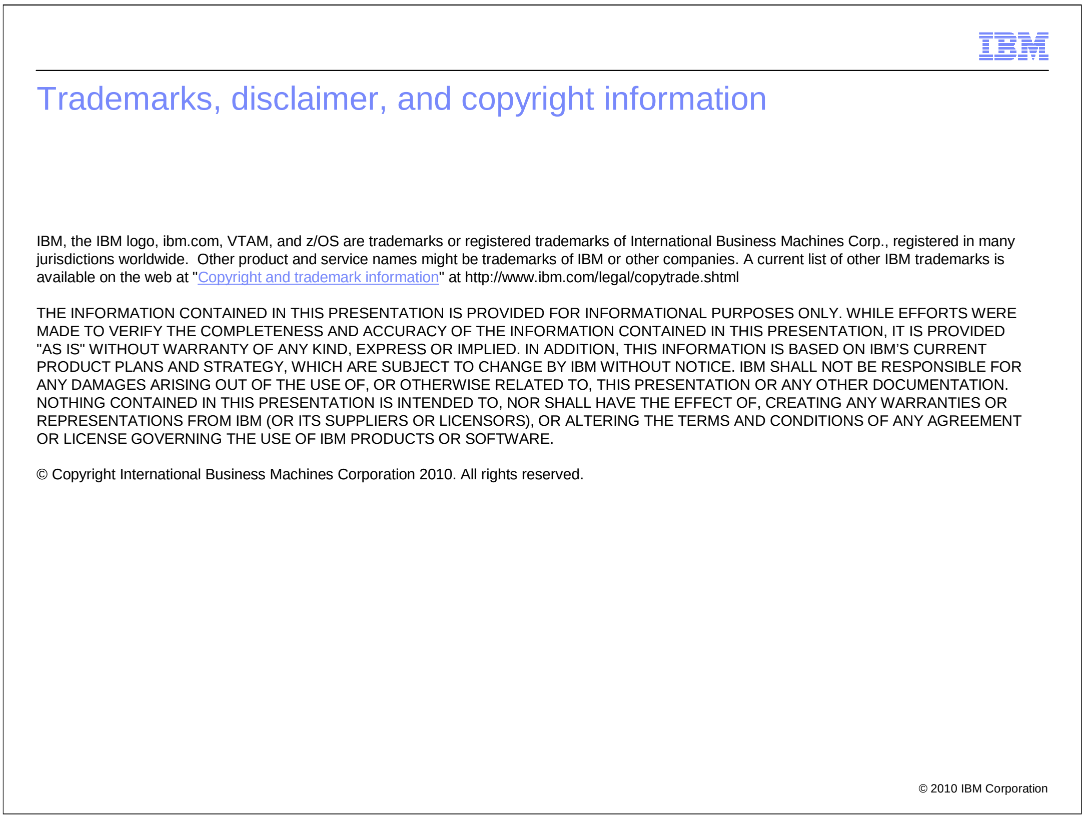

Trademarks, disclaimer, and copyright information

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