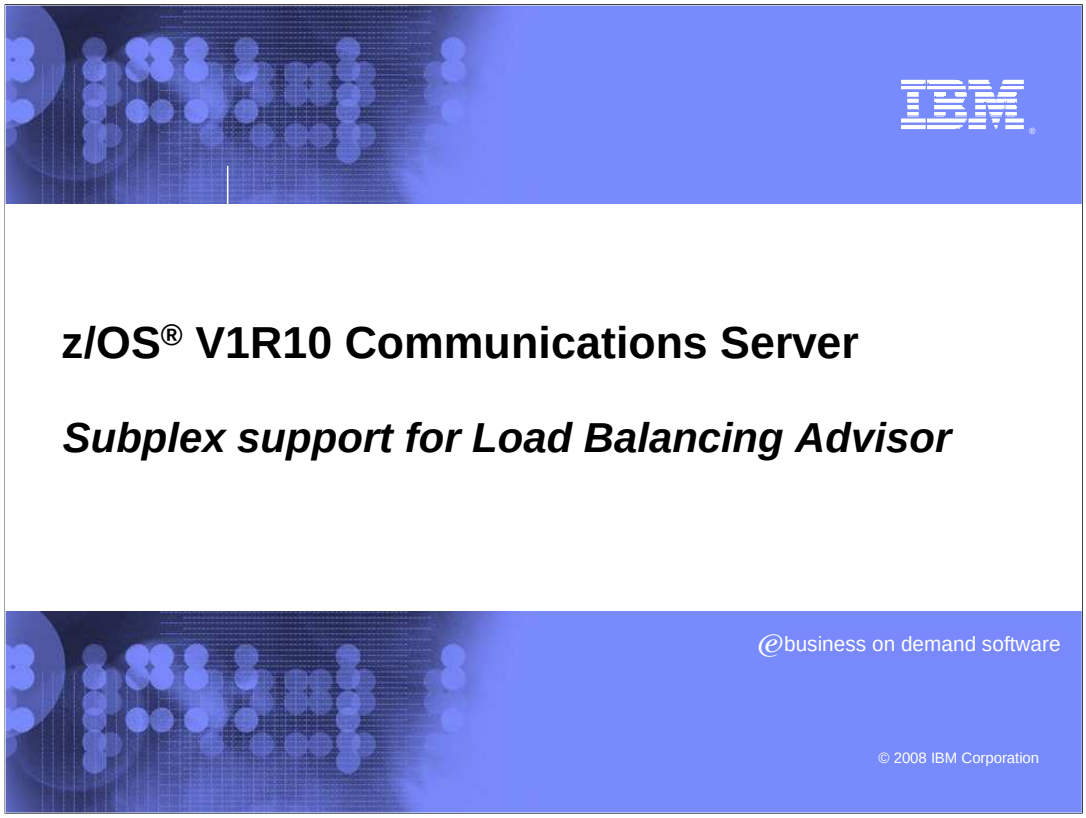

This presentation covers subplex support for load balancing advisor

Background

- A subplex is a subset of a sysplex
  - Members communicate through cross-system coupling facility (XCF) groups
  - Members communicate with each other, and not with members outside the subset.
- Useful when multiple networks with different attributes are attached to specific systems in the sysplex
  - Example: internal-only and external networks
  - Still in one sysplex but isolated

A subplex is a subset of a sysplex that consists of all the members in the sysplex that communicate through cross-system coupling facility (XCF) groups with each other, and not with members outside the subset. TCP members of a sysplex communicate through the coupling facility using Dynamic XCF interfaces - either IUTSAMEHOST, hipersockets, or XCF signaling. Subplexing was introduced in V1R8. It allows a sysplex to be partitioned into sysplex subsets. TCP members within a subplex communicate only with each other using Dynamic XCF.

Typically subplexing is used to isolate members with the same security characteristics. For example, a subplex can be created for a group of systems that communicate on an internal network with each other. Another subplex can contain systems that communicate with external networks and therefore need a higher level of security.

Defining multiple subplex scopes within a sysplex can be useful in scenarios where multiple networks with different security or functional attributes are attached to specific systems in the sysplex. For example, consider the scenario where some systems in the sysplex are connected to an internal enterprise network and other systems in the sysplex are connected to external networks. In this example, you can partition the sysplex into two subplexes from a sysplex network function perspective. One subplex can include the systems connected to the internal network. The other subplex can include the systems that are connected to the external networks. By defining these separate subplexes, users can still exploit some of the sysplex network functions while preserving the isolation of the networks. They will not automatically enable dynamic XCF connectivity across the entire sysplex.

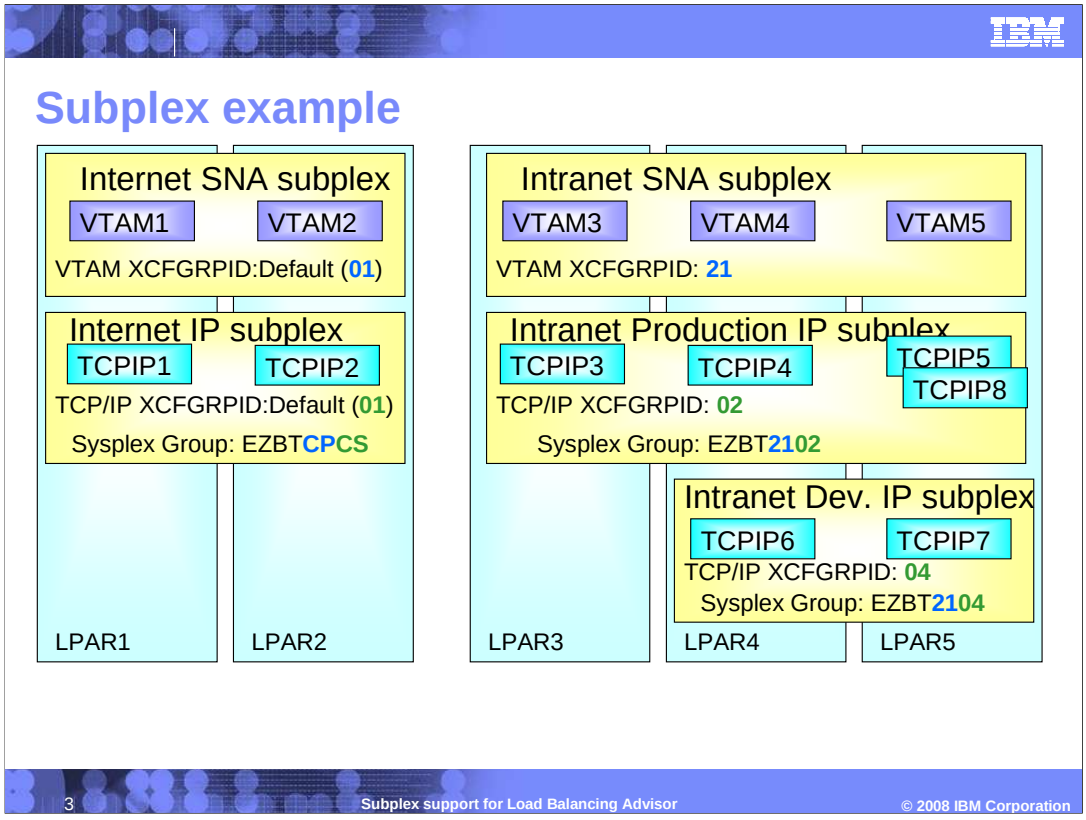

This page shows a sysplex environment divided into three TCP/IP subplexes.

The subplex on the left spans LPAR1 and LPAR2. There are two stacks in this subplex: stack TCPIP1 on LPAR1 and stack TCPIP2 on LPAR2. This subplex uses both the VTAM and TCP/IP default XCF group IDs of 01. The TCP/IP subplex ID is 0101 is formed by concatenating the VTAM XCF group ID 01 (default) and the TCP/IP XCF group ID 01 (default). The TCP/IP sysplex group name is EZBTCPCS, which is the default name when you use the VTAM and TCP/IP XCF group ID defaults. In a subplex, the TCP/IP sysplex group name is EZBTvvtt, where vv is the VTAM XCFGRPID and tt is the TCP/IP XCFGRPID. When VTAM XCFGRPID is not specified, the name is EZBTCPtt. When the TCP/IP XCFGRPID is not specified, the name is EZBvvCS; therefore the default TCP/IP sysplex group name when neither XCFGRPID is specified is EZBTCPCS.

There are two subplexes in the three LPARs on the right side of the slide. The production IP subplex has TCP/IP subplex ID 2102 because the VTAM XCF group ID is 21 and the TCP/IP XCF group ID is 02. Subplex 2102 spans LPAR3, LPAR4, and LPAR5. The TCP/IP sysplex group name is EZBT2102. There are four stacks in this subplex: stack TCPIP3 on LPAR3, stack TCPIP4 on LPAR4, and stacks TCPIP5 and TCPIP8 on LPAR5.

The Development IP subplex spans only LPAR4 and LPAR5. This subplex has a TCP/IP subplex ID of 2104 which is VTAM XCF group ID 21 and TCP/IP XCF group ID 04. The TCP/IP sysplex group name EZBT2104. There are two stacks in this subplex: stack TCPIP6 on LPAR4 and stack TCPIP7 on LPAR5.

A TCP/IP subplex can not span multiple VTAM subplexes because all TCP/IP stacks on an LPAR use the same VTAM for their dynamic XCF communication.

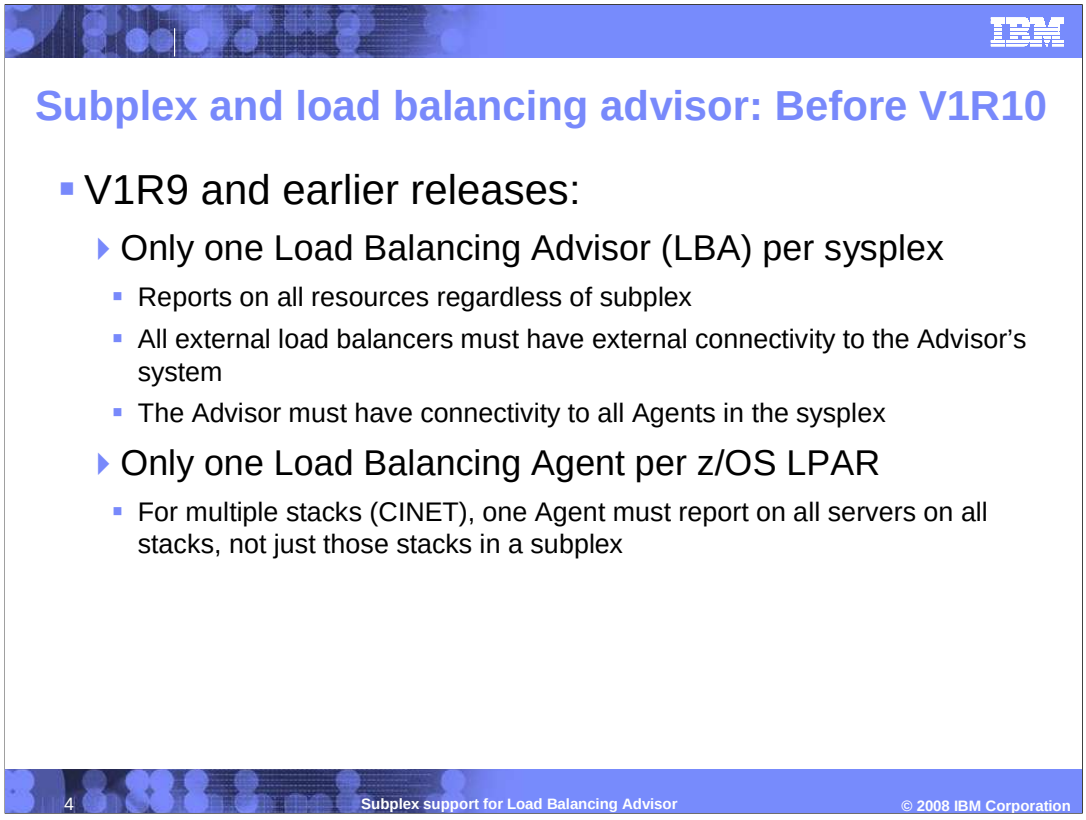

Before z/OS V1R10, the z/OS Load Balancing Advisor and Agent did not formally support the TCP/IP and SNA subplex function that was introduced in z/OS V1R8. While the Load Balancing Advisor and Agents do not explicitly make use of XCF services, they do not fit as neatly into a TCP/IP subplex environment in their pre-V1R10 implementation.

In V1R9 and earlier releases, only a single Advisor can be started within the sysplex. This is enforced by using a sysplex-wide enqueue. The Advisor has access to information for all systems/applications in the sysplex. Unaware of subplex scope, it reports on all resources in the sysplex. Subplexing can be accomplished only by configuring the external load balancer (that is, only specifying applications/systems that belong to a specific subplex). The external load balancers must all have external IP connectivity to the system that the Advisor is running on. The single Advisor must have IP connectivity to all Agents running in the sysplex, regardless of which TCP/IP subplex the stacks that they are executing on belong to.

Also in V1R9 and earlier, only a single Load Balancing Agent can be started within a z/OS LPAR. This is enforced by using a system-wide enqueue. If the z/OS LPAR has multiple TCP/IP stacks active, a single Agent interacts with all of them. In scenarios where TCP/IP subplexes are deployed in a CINET environment, it is desirable to have the ability to start an Agent instance per subplex.

Subplex and automated domain name registration: V1R9 and earlier

- V1R9 and earlier releases:
  - Only one Automated Domain Name Registration (ADNR) server per sysplex can be automatically recovered

In V1R9, you can start multiple ADNR instances in a sysplex, but only one is recoverable using the automated restart manager (ARM). This inhibits recoverability, particularly in a subplex environment.

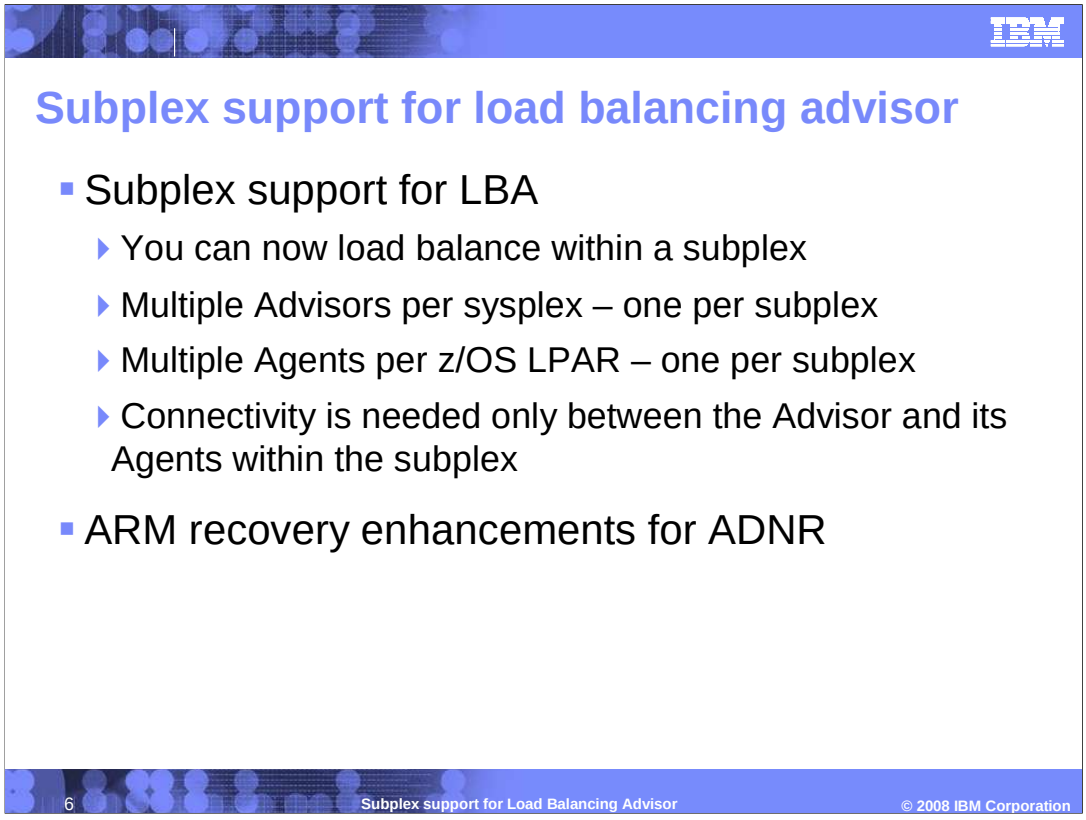

In z/OS V1R10, you can now load balance within a subplex. The Advisor allows multiple instances of itself to be present in the sysplex environment. Each Advisor is configured with the TCP/IP sysplex group name of the subplex that Advisor is handling.

Also, Load Balancing Agent allows multiple instances of itself to be present in the same MVS image or LPAR. Each Agent is configured with the TCP/IP sysplex group name of the subplex that Agent handles.

Additionally, for ADNR you can specify an ARM element name so that each instance of ADNR in a sysplex can be restarted.

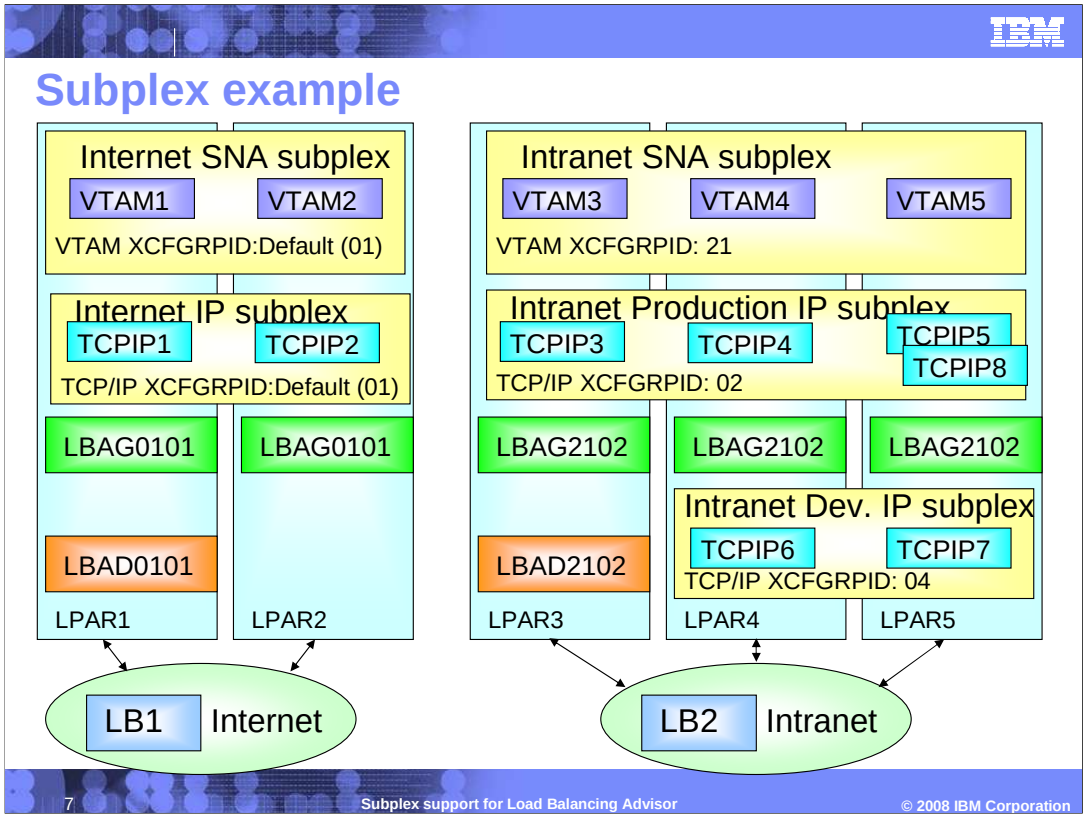

The previous figure is now shown with LBA applications configured to allow external Load Balancers to connect to the Internet IP subplex and the Intranet Production IP subplex.

LB1 is balancing connections to applications running on TCP/IP stacks in the Internet IP subplex on LPAR1 and LPAR2. The TCP/IP sysplex group name is EZBTCPCS (VTAM XCFGRPID 01 and TCP/IP XCFGRPID 01); this is the default TCP/IP sysplex group name when the TCP/IP subplex ID is 0101 (the default VTAM and TCP/IP XCFGRPID). LB1 connects to the Load Balancing Advisor in this subplex. The Advisor job, LBAD0101, is configured to use stacks that are members of TCP/IP subplex ID of 0101. A single instance of this Advisor can run in LPAR1 or LPAR2; it is currently running in LPAR1. Two Agents are configured to use the stacks that are members of TCP/IP subplex ID of 0101. The Agent job names are LBAG0101 on LPAR1 and LBAG0101 on LPAR2.

LB2 is balancing connections to applications running TCP/IP stacks in the Intranet Production IP subplex on LPAR3, 4, and 5. TCP/IP sysplex group name is EZBT2102 (VTAM XCFGRPID 21 and TCP/IP XCFGRPID 02). The TCP/IP subplex ID is 2102. LB2 connects to a Load Balancing Advisor in this subplex. The Advisor, LBAD2102, is configured to use stacks that are members of TCP/IP subplex ID of 2102. A single instance of this Advisor can run in LPAR3, LPAR4, or LPAR5; it is currently running in LPAR3. Three Agents are configured to use stacks that are members of TCP/IP subplex ID of 2102. The three Agent job names are LBAG02102 on LPAR3, LBAG2102 on LPAR4, and LBAG2102 on LPAR5. Note that, even though there are two TCP/IP stacks in LPAR5 in subplex 2102, there is only one Load Balancing Agent for that subplex on that LPAR. The one Agent will report on all servers in that LPAR in that subplex.

There is no Load Balancer for applications running in the Intranet Development IP subplex. Therefore no Advisor and no Agents need to run in this subplex. ***If*** you want to load balance in the Intranet Development IP subplex, configure an Advisor instance to run on either LPAR4 or LPAR5. Also, configure an Agent instance to run on both LPAR4 and LPAR5. Configure the Advisor and Agent applications to use stacks that are members of TCP/IP subplex ID 2104 (TCPIP6 and TCPIP7).

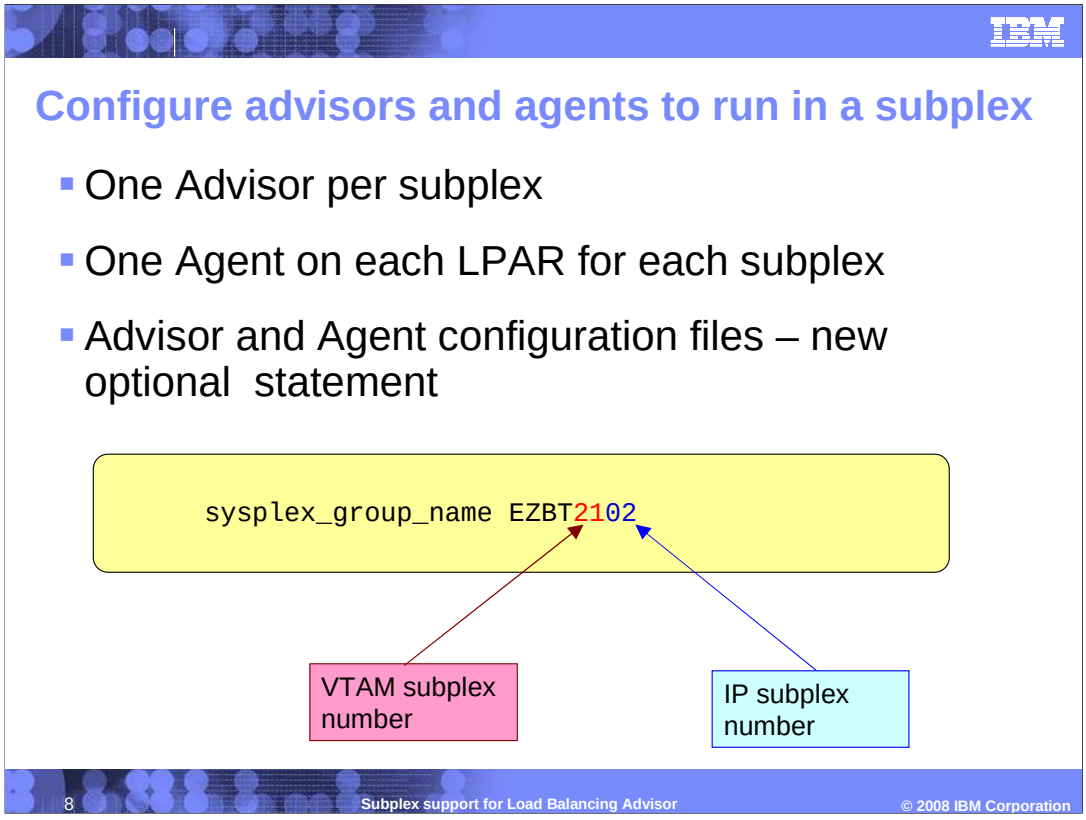

You configure one Advisor per subplex, and you configure one Agent for each subplex on each LPAR within that subplex. As of z/OS V1R10, more than one Agent can run in an LPAR.

Specify which subplex the Advisor and Agent is associated with using the sysplex_group_name statement in their configuration files.

The TCP/IP sysplex group name is in the format EZBT*vvtt*. *vv* is the VTAM subplex group ID, as specified on the VTAM XCFGRPID start option. If VTAM was started with no subplex group ID specified, *vv* is 'CP'. *tt* is the TCP/IP subplex group ID, as specified on the XCFGRPID parameter of the GLOBALCONFIG statement in the TCP profile. If the TCP/IP stack was started with no subplex ID specified, *tt* is 'CS'. Therefore, the default TCP/IP sysplex group name is EZBTCPCS when the TCP/IP subplex ID is 0101 (the default VTAM and TCP/IP XCFGRPIDs).

Each Load Balancing Advisor and Agent operating in a sysplex subplexing environment should specify a sysplex_group_name statement in its configuration file. The statement is optional. If it is omitted, it is assumed that the Advisor or Agent is not running in a subplexing environment.

Tip: Use the DISPLAY TCPIP,,SYSPLEX,GROUP command to display the current TCP sysplex group name.

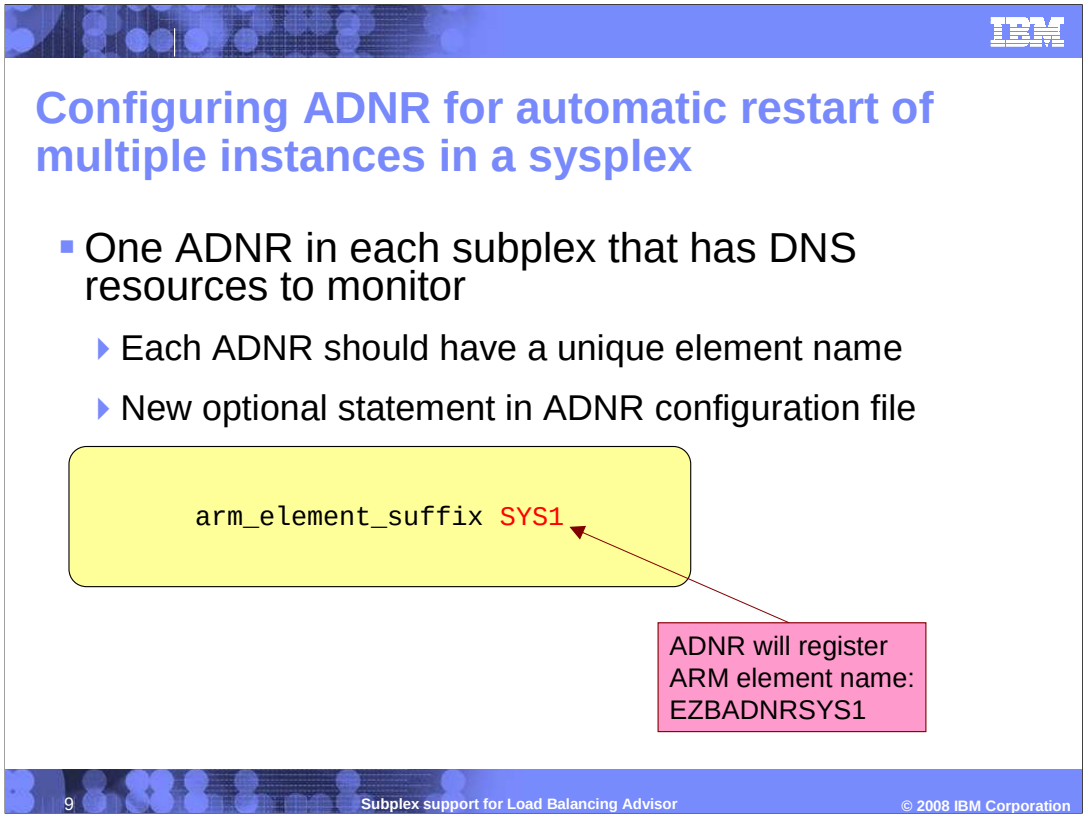

ADNR supports only one connection to an Advisor. If you are using ADNR in a subplexing environment, you will need a separate ADNR instance in each subplex that contains name server resources being monitored by ADNR.

You should configure each ADNR instance to have a unique Automatic Restart Manager (ARM) element name. If you do not specify a unique element name for each instance of ADNR, only the first instance of ADNR that registers with ARM will recover automatically.

ADNR registers the ARM element name EZBADNR concatenated with the arm_element_suffix value. If there is no arm_element_suffix statement, the element name is EZBADNR. For example, if you configure this statement arm_element_suffix SYS1, ADNR registers ARM element name EZBADNRSYS1.

Restriction: There must be a one-to-one correspondence between a subplex and a DNS zone.

Recommendation: Each ADNR instance should have a unique ARM element name within a sysplex.

For more information, refer to *z/OS Communications Server: IP Configuration Reference,* chapter ‘Automated domain name registration’, section ‘arm_element_suffix’.

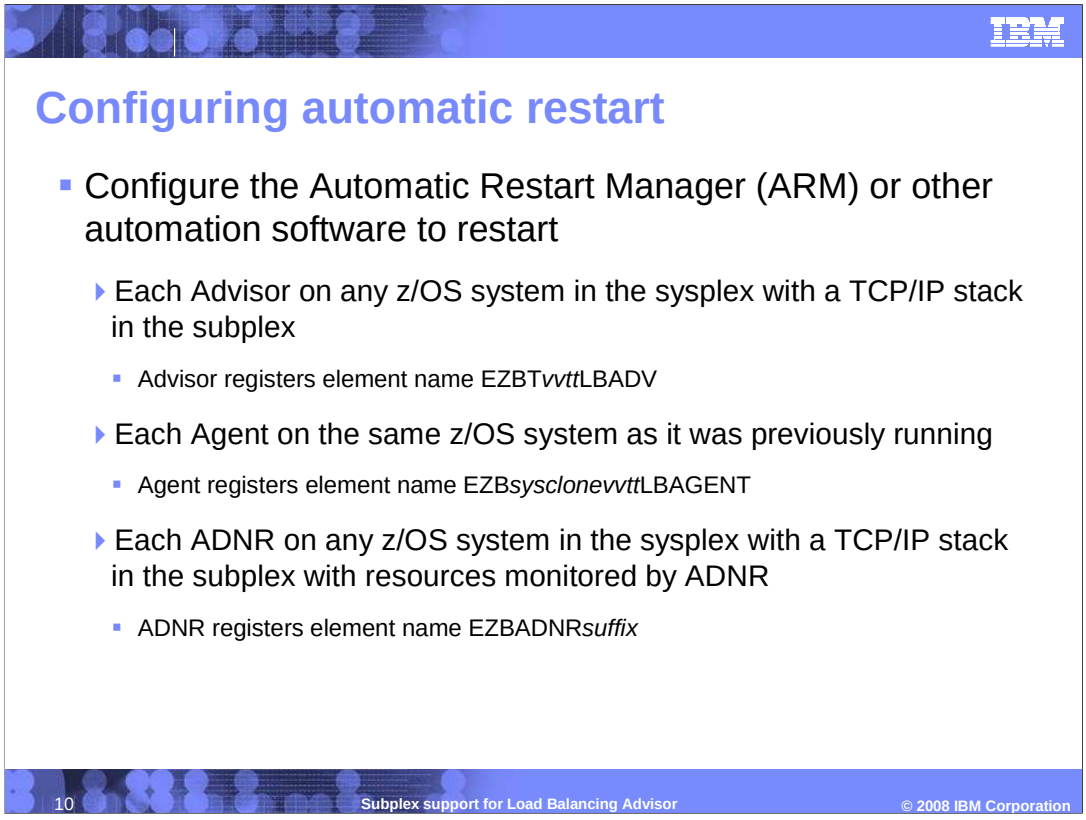

Configure the Automatic Restart Manager (ARM) or other automation software to restart each Advisor on any z/OS system in the sysplex that supports the subplex that was configured for that Advisor

Configure the ARM to restart each Agent on the same z/OS system as it was previously running, and to restart each ADNR on any z/OS system in the sysplex that supports the subplex that the ADNR is associated with.

The subplex is determined by both the VTAM and TCP/IP subplex group IDs, which are denoted by *vv* and *tt*. These subplex group IDs are reflected in the TCP/IP sysplex group name, which is of the format EZBT*vvtt*. Each Load Balancing Advisor registers with ARM with element name EZBT*vvtt*LBADV

The Agent registers with ARM with element name EZB*sysclonevvtt*LBAGENT. *sysclone* is a 1- or 2-character shorthand notation for the name of the MVS system. vvtt is the last four characters of the sysplex_group_name parameter in the Agent configuration file. If this parameter is not specified, vvtt is omitted.

For example, if the SYSCLONE value is 03 and sysplex_group_name is not specified in the Agent configuration file, the resulting element value is EZB03LBAGENT. If the SYSCLONE value is 03 and sysplex_group_name is EZBT2102, the resulting element value is EZB032102LBAGENT.

ADNR registers with ARM with element name EZBADNR*suffix* where *suffix* is a 1 to 8 character string specified on the arm_element_suffix configuration statement.

Configuring automatic restart

- ARM policy TARGET_SYSTEM keyword
  - Ensure application restarts in correct subplex
- SAF profile
  - In the FACILITY class, give UPDATE authority for user ID of each started task to IXCARM.SYSTCPIP.*elemname*

You should define an ARM policy with the TARGET_SYSTEM keyword to indicate which systems the element can be restarted on in order to ensure that the application is restarted only on a system that is in the same subplex. The Advisor and Agent should be restarted on a system with a VTAM that has been started with the same XCFGRPID (*vv*) and has an available TCP/IP stack with XCFGRPID *tt*.

When you use ARM registration, you must permit the started task IDs for each Agent, each Advisor, and each ADNR instance to have UPDATE authority to the correct profiles. These profiles are the IXCARM.SYSTCPIP.elemname profiles in the FACILITY class in the SAF-compliant security product on your system.

For example, if the VTAM XCFGRPID is 21 and the TCP/IP XCFGRPID is 02, the facility name is IXCARM.SYSTCPIP.EZBT2102LBADV.

For a complete description of the SYSCLONE static system symbol, refer to *z/OS MVS Initialization and Tuning Guide*.

For information about automatic restart, refer to *z/OS MVS Setting Up a Sysplex.*

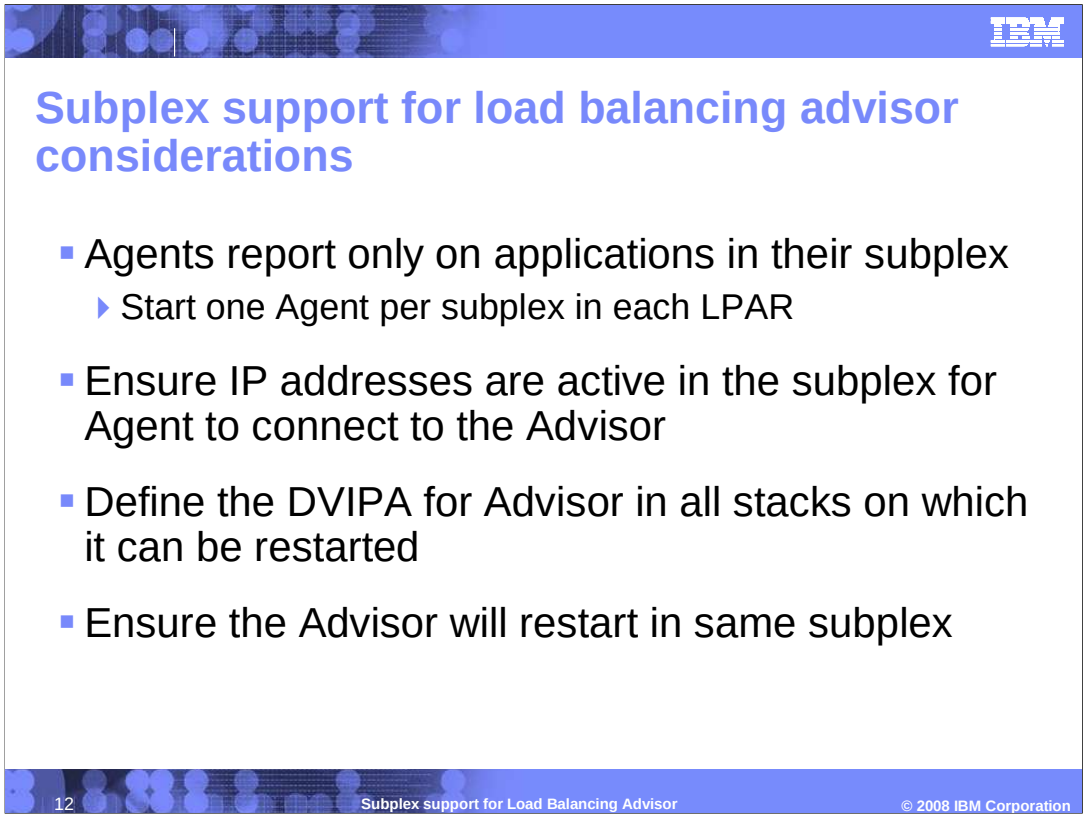

Ensure that the Advisor and Agent are configured for a subplexed environment.

There should be one Agent on each LPAR in the subplex. The Agents report to the Advisor only about applications within their subplex. If the Agent is configured with *sysplex_group_name* EZBT*vvtt*, the Agent reports only applications that are on the VTAM subplex *vv* and TCP/IP stacks with subplex *tt*. When configured for subplexing, the Agents do not report on other applications in the same LPAR. As of z/OS V1R10, there can be more than one Agent in an LPAR if they are in different subplexes.

IP addresses used as the source IP address for outbound Agent connection to the Advisor should be configured/owned by the proper stacks.

The DVIPA for the Advisor needs to be defined in all the stacks associated with the subplex (and where a restart of the Advisor can occur).

For recovery and restart operations, the Advisor should be restarted on a TCP/IP stack that belongs to the same subplex.

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 z/OS

A current list of other IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

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 (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

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

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

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

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