This presentation covers core group failure detection on the z/OS platform.

Agenda

- Current core group detection protocol
- XCF as the core group detection protocol

The current core group detection protocol used on the z/OS platform is covered first, after which you will see how to replace this with XCF in the version 7 product.

Core group discovery/failure detection, V6

- Heartbeats are used to detect core group member startups and failures
- Heartbeat interval set to 30 seconds by default
- Core group custom properties available to fine-tune this interval
  - IBM_CS_FD_PERIOD_SECS, which specifies the time interval, in seconds, between consecutive heartbeats. The default value for this property is 30 seconds
  - IBM_CS_FD_CONSECUTIVE_MISSED, which specifies the consecutive number of heartbeats that must be missed before the core group member is considered failed. The default value for this property is 6.

Starting with V6, a high availability manager is available to ensure that components within the WebSphere Application Server are always available. A high availability manager instance runs on each server in a cell, including application servers, deployment managers and node agents. A cell can be divided into multiple high availability domains, known as core groups, and every server that belongs to that core group is monitored for its startup and failure. In V6, this is done solely using a heartbeat mechanism. At regularly scheduled intervals, each core group member sends a ping packet on every open core group connection. If the packet is acknowledged, all is assumed to be well. If no response is received from a given member for a certain number of consecutive pings, the member is marked as failed. The heartbeat interval is set to 30 seconds by default, but this can be tuned with custom properties on the core group. The IBM_CS_FD_PERIOD_SECS custom property allows you to change the time interval between heartbeats and the IBM_CS_FD_CONSECUTIVE_MISSED custom property allows you to change the number of missed heartbeats before considering the member as failed

Core group discovery/failure detection, V7.0

- WebSphere for z/OS V7.0 now allows XCF to be used as the protocol for discovery and failure
  - Reduces processor usage compared to heartbeat mechanism
  - Allows more realistic ‘interval’ for failure determination
- Heartbeats are still available as mechanism to detect core group member status

WebSphere for z/OS Version 7.0 adds the ability to specify that XCF should be used to determine core group member startup and failure. Use of the XCF system services to provide the core group member status significantly reduces processor usage compared to the heartbeat mechanism. This is particularly noticeable during processor idle times. A 30 second heartbeat interval was necessary to determine member failure so as not to consume a prohibitive amount of processor time. Using XCF as the notification mechanism allows more timely notification in addition to significantly reduced processor usage.

The heartbeat mechanism is still available and is currently the default.

Specifying XCF…requirements

- All servers that are members of the core group must be Version 7.0 or higher
- All core group members must be running the z/OS operating system
- z/OS VTAM® component must be configured to start XCFINIT = YES

In order to take advantage of XCF for the discovery and failure protocol, all servers in the target core group must be at Version 7.0 or higher. If you have a mixed cluster with different WebSphere levels, you must continue to use the default discovery and failure detection protocol until you have migrated all servers in the core group. You must also have all members of the core group configured on the z/OS operating system. The XCF protocol can only be used in a homogenous environment. The XCFINIT VTAM parameter must be specified in order to enable TCP/IP to use the services that the z/OS Cross-system Coupling Facility provides.

Specifying XCF for discovery and failure protocol

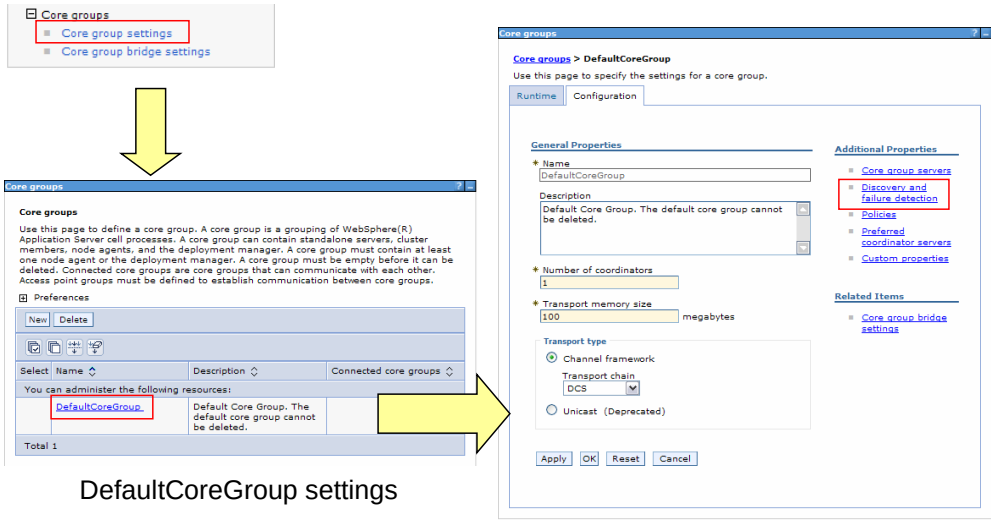

DefaultCoreGroup settings

Under core group settings, you will find the core groups defined to your cell. This slide shows only the DefaultCoreGroup which is created for you automatically on cell creation. Looking at the core group you want to change, you will find a Discovery and failure detection option, which is new for Version 7.0. This is where you specify that you want to use XCF for the discovery and failure protocol.

Specifying XCF for discovery and failure protocol

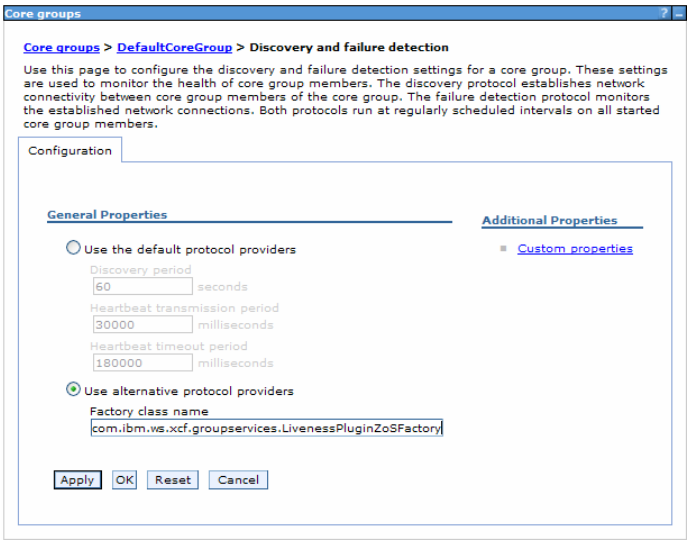

- **com.ibm.ws.xcf.groupservices.LivenessPluginZoSFactory.**

To use XCF instead of the default protocol provider for discovery and failure detection, check the 'Use alternative protocol providers' radio button. The Factory class name is blank, so you must provide the factory class name that is used to create the alternative protocol provider, which is **com.ibm.ws.xcf.groupservices.LivenessPluginZoSFactory.**

Default protocol settings…

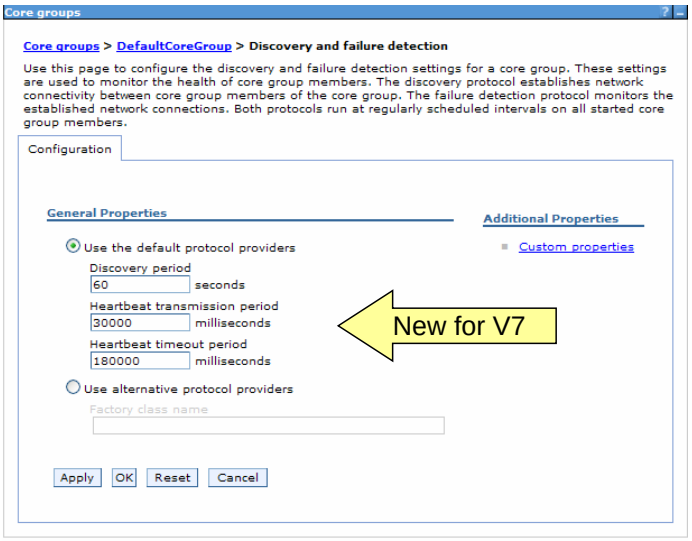

- Custom properties still needed if any V6 core group members

If you are using the default protocol provider for core group member discovery and failure, you can now specify values for the heartbeat time and the timeout interval in the administrative console. These were previously specified as custom properties on the core groups, IBM_CS_FD_PERIOD_SECS and IBM_CS_FD_CONSECUTIVE_MISSED. You must still specify the custom properties if any V6 members exist.

Summary

- XCF can be used as the protocol for core group detection failure on z/OS

In summary, the new XCF protocol can now be used to detect a core group failure. This presentation showed you how and when you can specify the new protocol.

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 e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WASv7zOS_DCS_XCFPlug-in.ppt

This module is also available in PDF format at: ../WASv7zOS_DCS_XCFPlug-in.pdf

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

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 VTAM WebSphere 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.

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.