



IBM Multi-site Workload Lifeline V2.5 User's Guide

Version 2 Release 5

Note:

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 281.

Fourth Edition (July 2015)

This edition applies to Version 2 Release 5 of IBM Multi-site Workload Lifeline (5655-UM4), and to subsequent releases and modifications until otherwise indicated in new editions.

IBM welcomes your comments. You can send us comments electronically by using one of the following methods:

Internet email:

comsvrcf@us.ibm.com

World Wide Web:

<http://www.ibm.com/systems/z/os/zos/webqs.html>

If you would like a reply, be sure to include your name, address, and telephone number. Make sure to include the following information in your comment or note:

- Title and order number of this document
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corporation 2011, 2015.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables ix

About this document xi

Who should read this document xi

How this document is organized xi

How to use this document xii

 Determining whether a publication is current xii

 How to contact IBM service xii

Conventions and terminology that are used in this document xiii

Qualifying elements xiii

How to read a syntax diagram. xiii

How to send your comments xvi

Summary of changes xvii

Chapter 1. Overview of Multi-site Workload Lifeline 1

Lifeline Advisors 8

Lifeline Agents 9

External load balancers 10

Internal load balancers 12

Intermediary nodes. 13

| IBM MQ Clusters 13

Linux on System z Workloads 13

 Overview of Linux on System z Workloads. 13

 Monitoring Linux on System z application groups 15

 The Linux on System z Management Guest. 15

 The Linux on System z Management Agent 15

 Preparing to use the Linux on System z Management Agent 16

 Installing the Linux on System z Management Agent 17

 Configuring the Linux on System z Management Agent 18

 Starting the Linux on System z Management Agent 18

 Operational characteristics of Linux on System z workloads 19

 Diagnosing Linux on System z Management Agent problems. 20

 Diagnosing Linux on System z workload problems 21

SNA workloads 22

 Overview of Lifeline in a SNA workload environment 22

 Monitoring SNA application groups 24

 Preparing z/OS Agents for monitoring SNA applications 24

 Operational characteristics of SNA workloads 24

 Diagnosing problems when a SNA application is unavailable. 26

 Diagnosing problems when SNA intermediary node servers are unavailable. 27

| IBM MQ Cluster Workloads 28

 Overview of MQ Cluster Workloads 28

 Monitoring MQ cluster queue groups. 29

 Preparing z/OS Agents for monitoring MQ queue managers 30

 Operational characteristics of MQ cluster workloads. 30

 Diagnosing problems when an MQ queue manager is unavailable 31

Chapter 2. Preparing to use Multi-site Workload Lifeline 33

(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent 33

Authorizing the Lifeline Advisors to use BCPii 34

Authorizing the Lifeline Advisor to accept secure peer Advisor connections. 35

	(z/OS systems) Authorizing the Lifeline Advisor to accept secure Agent connections	36
	Authorizing the Lifeline Advisor to accept secure Load Balancer connections	36
	(z/OS systems) Authorizing the Lifeline Advisors and Agents to issue event-driven messages.	37
	(z/OS systems) Authorizing the Lifeline Agents to use WLM services	38
	(z/OS systems) Authorizing the Lifeline Agents to reset active connections	38
	(z/OS systems) Authorizing the Lifeline Agents to monitor SNA applications	39
	(z/OS systems) Enabling the Lifeline Agents to communicate with MQ queue managers	40
	Granting a network management application authority to use the Network Management Interface	40
	Chapter 3. Configuring Multi-site Workload Lifeline	43
	Lifeline Advisor configuration file	43
	Lifeline Agent configuration file	43
	Configuring one Lifeline Advisor as the primary Advisor, and optionally one as the secondary Advisor	44
	Configuring a Lifeline Agent	46
	(z/OS systems) WLM recommendations.	47
	Configuring Lifeline Advisors to automatically restart	48
	(z/OS systems) Recovering from a planned outage of the TCP/IP stack	49
	Recovering from an unplanned outage of the TCP/IP stack	49
	(z/OS systems) Configuring Lifeline Agents to automatically restart	50
	Migrating IBM Multi-site Workload Lifeline	51
	Activating a workload for the first time	51
	Activating a workload when the workload is not deactivated.	51
	Activating or quiescing a workload when not all first-tier load balancers are registered	52
	Activating a workload when not all Lifeline Agents are connected	52
	Deactivating a workload when not all Lifeline Agents are connected	53
	Migrating the Lifeline Advisor	53
	Migrating the Lifeline Agent.	54
	Using the Base Control Program Internal Interface	54
	Configuring and starting syslogd	55
	(z/OS systems) Customizing the TCP/IP profiles of the TCP/IP stacks	55
	CINET considerations	56
	Customizing the VTAM started procedure when configuring SNA workloads	57
	Customizing the NetView for z/OS Message Revision Table when configuring SNA workloads	57
	Configuring external load balancers	58
	Configuring first-tier external load balancers	59
	Configuring second-tier external load balancers	60
	Configuring second-tier internal load balancers	61
	Configuring support for intermediary nodes	61
	Configuring IBM MQ clusters	63
	Chapter 4. Starting the Lifeline Advisors and Agents	65
	Chapter 5. Diagnosing problems in Multi-site Workload Lifeline	67
	Debug levels	67
	Diagnostic data	68
	Interpreting the workload and LB displays	70
	Verifying that the Lifeline Advisor and Agent systems are functioning correctly	72
	Diagnosing problems when a workload is not distributed	74
	Diagnosing why a workload is not being distributed to a server application.	74
	Diagnosing why a workload is not being distributed to an MQ cluster queue	76
	Diagnosing workload distribution problems	77
	Diagnosing workload distribution problems for server applications.	77
	Diagnosing workload distribution problems for IBM MQ clusters	78
	Diagnosing problems when the Lifeline Advisor or Agent appears hung	79
	Diagnosing problems when the group names are not readable	80
	Diagnosing problems when load balancer connection terminates	80
	Diagnosing problems when the Lifeline Agent connection terminates	81
	Diagnosing problems when the Lifeline Advisor to Advisor connection terminates	82
	(z/OS systems) Diagnosing failures of registering with ARM.	82
	Diagnosing problems in the Linux on System z Management Agent	83

Chapter 6. Multi-site Workload Lifeline configuration statements	85
Lifeline Advisor configuration statements	85
<i>advisor_connection_port</i>	87
<i>advisor_id_list</i>	87
<i>advisor_takeover_policy</i>	88
<i>advisor_timeout</i>	89
<i>agent_connection_port</i>	89
<i>agent_id_list</i>	90
<i>cross_sysplex_list</i>	91
<i>debug_level</i>	93
<i>failure_detection_interval</i>	93
<i>intermediary_node_list</i>	94
<i>lb_connection_v4</i>	96
<i>lb_connection_v6</i>	97
<i>lb_id_list</i>	98
<i>management_guest_list</i>	99
<i>mq_manager_list</i>	101
<i>peer_advisor_id</i>	103
<i>query_workload_list</i>	104
<i>sna_application_list</i>	106
<i>update_interval</i>	107
<i>workload_switch_policy</i>	109
Lifeline Agent configuration statements	109
<i>advisor_id</i>	110
<i>advisor_id_list</i>	111
<i>debug_level</i>	111
<i>host_connection</i>	112
<i>site_name</i>	113
Chapter 7. Lifeline Advisor commands	115
MODIFY <i>advisor_proc</i> ,DISPlay command: Display Lifeline Advisor information	115
MODIFY <i>advisor_proc</i> ,ACTivate command: Activate workloads	142
MODIFY <i>advisor_proc</i> ,DEACTivate command: Deactivate workloads	143
MODIFY <i>advisor_proc</i> ,DEBUg command: Change the debug level	145
MODIFY <i>advisor_proc</i> ,DUMPLog command: Dump the internal debug records.	146
MODIFY <i>advisor_proc</i> ,QUIESCE command: Quiesce workloads	146
MODIFY <i>advisor_proc</i> ,REFRESH command: Update configuration	148
MODIFY <i>advisor_proc</i> ,TAKEOVER command: Take over primary Advisor responsibilities	148
Chapter 8. Lifeline z/OS Agent commands	151
MODIFY <i>agent_proc</i> ,DISPlay command: Display Lifeline Agent information	151
MODIFY <i>agent_proc</i> ,ENABLe command: Enable server applications	158
MODIFY <i>agent_proc</i> ,QUIESCE command: Quiesce server applications	160
MODIFY <i>agent_proc</i> ,DEBUg command: Change the debug level.	163
MODIFY <i>agent_proc</i> ,DUMPLog command: Dump the internal trace records.	163
Chapter 9. Lifeline Linux on System z Management Agent commands	165
Display Linux on System z Management Agent information.	165
Change the debug level	169
Dump internal trace records	169
Terminate the Linux on System z Management Agent	170
Chapter 10. Network Management Interface for Multi-site Workload Lifeline	171
NMI message format	173

Chapter 11. Multi-site Workload Lifeline messages	179
Appendix A. A sample started procedure for Lifeline Advisors	253
Appendix B. A sample started procedure for Lifeline z/OS Agents.	255
Appendix C. Sample configuration statements	257
Appendix D. Sample AT-TLS configuration policies	275
Appendix E. Accessibility	279
Notices	281
Policy for unsupported hardware.	283
Trademarks	283
Index	285
Communicating your comments to IBM	289

Figures

1.	Relationships between external load balancers, Lifeline Advisors, and Lifeline Agents	4
2.	Relationships between external load balancers, internal intermediary nodes, Lifeline Advisors, and Lifeline Agents	6
3.	Relationships between external load balancers, external intermediary nodes, Lifeline Advisors, and Lifeline Agents	7
4.	Relationships between external load balancers, Linux on System z intermediary nodes, Lifeline Advisors, and Lifeline Agents for Linux on System z workloads	14
5.	Relationships between external load balancers, internal intermediary nodes, Lifeline Advisors, and Lifeline Agents for SNA workloads	23
6.	Relationships between MQ queue managers, Lifeline Advisors, and Lifeline Agents for MQ cluster workloads	29
7.	Flow between the Lifeline NMI and the network management application	171
8.	Layout of the data response message	172
9.	Sample started procedure AQSADV	253
10.	Sample started procedure AQSAGE	255
11.	Multi-site Workload Lifeline configuration Example 1	257
12.	Sample configuration statements for the Lifeline Advisor	259
13.	Sample configuration statements for the Lifeline Agent	260
14.	Multi-site Workload Lifeline configuration Example 2	261
15.	Sample configuration statements for the Lifeline Advisor, part 1 of 2	262
16.	Sample configuration statements for the Lifeline Advisor, part 2 of 2	263
17.	Multi-site Workload Lifeline configuration Example 3	263
18.	Sample configuration statements for the Lifeline Advisor, part 1 of 2	265
19.	Sample configuration statements for the Lifeline Advisor, part 2 of 2	266
20.	Multi-site Workload Lifeline configuration Example 4	267
21.	Sample configuration statements for the Lifeline Advisor, part 1 of 2	268
22.	Sample configuration statements for the Lifeline Advisor, part 2 of 2	269
23.	Multi-site Workload Lifeline configuration Example 5	270
24.	Sample configuration statements for the Lifeline Advisor, part 1 of 2	271
25.	Sample configuration statements for the Lifeline Advisor, part 2 of 2	272
26.	Sample configuration statements for the Linux on System z Management Agents	273
27.	Multi-site Workload Lifeline configuration Example 6	273
28.	Sample configuration statements for the Lifeline Advisor	274

Tables

1. Lifeline Advisor Configuration Statements	85
2. Lifeline Agent Configuration Statements	109

About this document

This document describes how to use IBM® Multi-site Workload Lifeline to load balance TCP/IP workload requests across two sites at unlimited distances when both sites run the same server applications and have the same data sources.

Who should read this document

This document is written for system administrators who need to understand how to load balance TCP/IP workload requests.

How this document is organized

This document contains the following information:

- Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1 introduces the concept of Lifeline Advisors, Lifeline Agents, external and internal load balancers, as well as an overview of Lifeline support for Linux on System z® workloads, SNA workloads, and IBM MQ Cluster workloads.
- Chapter 2, “Preparing to use Multi-site Workload Lifeline,” on page 33 contains the information about the tasks that you need to complete before you can configure Multi-site Workload Lifeline.
- Chapter 3, “Configuring Multi-site Workload Lifeline,” on page 43 contains the information about the tasks to configure Multi-site Workload Lifeline.
- Chapter 4, “Starting the Lifeline Advisors and Agents,” on page 65 contains the information about how to start the Lifeline Advisors and Agents that make up Multi-site Workload Lifeline.
- Chapter 5, “Diagnosing problems in Multi-site Workload Lifeline,” on page 67 contains the information that you need to diagnose problems.
- Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 describes the configuration statements of the Lifeline Advisors and Agents.
- Chapter 7, “Lifeline Advisor commands,” on page 115 describes the commands that you can use to display Lifeline Advisor information and modify Lifeline Advisor state information.
- Chapter 8, “Lifeline z/OS Agent commands,” on page 151 describes the commands that you can use to display Lifeline z/OS® Agent information and modify this Agent's state information.
- Chapter 9, “Lifeline Linux on System z Management Agent commands,” on page 165 describes the commands that you can use to display Linux on System z Management Agent information and modify this Agent's state information.
- Chapter 10, “Network Management Interface for Multi-site Workload Lifeline,” on page 171 describes the Network Management Interface (NMI) that you can use to monitor Lifeline metrics and statistics.
- Chapter 11, “Multi-site Workload Lifeline messages,” on page 179 contains the Multi-site Workload Lifeline messages.
- Appendix A, “A sample started procedure for Lifeline Advisors,” on page 253 contains a sample procedure that can be used to start the Lifeline Advisor.
- Appendix B, “A sample started procedure for Lifeline z/OS Agents,” on page 255 contains a sample procedure that can be used to start the Lifeline z/OS Agent.

- Appendix C, “Sample configuration statements,” on page 257 provides several examples of workloads that can be managed by Multi-site Workload Lifeline and the Lifeline Advisor and Agent configuration statements for these workloads.
- Appendix D, “Sample AT-TLS configuration policies,” on page 275 provides samples of Policy Agent configuration statements to be used to enable AT-TLS policies between Lifeline Advisors, Lifeline Agents and external load balancers.
- Appendix E, “Accessibility,” on page 279 describes accessibility features to help users with physical disabilities.
- “Notices” on page 281 contains notices and trademarks used in this document.

How to use this document

To use this document, you should be familiar with the network access to your z/OS server applications, either TCP/IP or SNA. If your workload server applications reside on Linux on System z guests, you should also be familiar with Linux on System z and z/VM[®] management of these guests.

Determining whether a publication is current

As needed, IBM updates its publications with new and changed information. For a given publication, updates to the hardcopy and associated BookManager[®] softcopy are usually available at the same time. Sometimes, however, the updates to hardcopy and softcopy are available at different times. The following information describes how to determine if you are looking at the most current copy of a publication:

- At the end of a publication's order number there is a dash followed by two digits, often referred to as the dash level. A publication with a higher dash level is more current than one with a lower dash level. For example, in the publication order number GC28-1747-07, the dash level 07 means that the publication is more current than previous levels, such as 05 or 04.
- If a hardcopy publication and a softcopy publication have the same dash level, it is possible that the softcopy publication is more current than the hardcopy publication. Check the dates shown in the Summary of Changes. The softcopy publication might have a more recently dated Summary of Changes than the hardcopy publication.
- To compare softcopy publications, you can check the last two characters of the publication's file name (also called the book name). The higher the number, the more recent the publication. Also, next to the publication titles in the CD-ROM booklet and the readme files, there is an asterisk (*) that indicates whether a publication is new or changed.

How to contact IBM service

For immediate assistance, visit this website: http://www.ibm.com/support/entry/portal/Overview/Software/Other_Software/z~OS_Communications_Server/

Most problems can be resolved at this website, where you can submit questions and problem reports electronically, and access a variety of diagnosis information.

For telephone assistance in problem diagnosis and resolution (in the United States or Puerto Rico), call the IBM Software Support Center anytime (1-800-IBM-SERV). You will receive a return call within 8 business hours (Monday – Friday, 8:00 a.m. – 5:00 p.m., local customer time).

Outside the United States or Puerto Rico, contact your local IBM representative or your authorized IBM supplier.

If you would like to provide feedback on this publication, see “Communicating your comments to IBM” on page 289.

Conventions and terminology that are used in this document

Samples used in this book might not be updated for each release. Evaluate a sample carefully before applying it to your system.

For definitions of the terms and abbreviations that are used in this document, you can view the latest IBM terminology at the IBM Terminology Web site.

Qualifying elements

Titles

Some topics apply only to Lifeline Advisors and the Lifeline Agents that reside on z/OS systems. Such topics have the feature information presented in parentheses as a qualifying element in the topic title. For example, the topic "(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent" applies only to Lifeline Advisors and those Lifeline Agents that reside on z/OS systems.

Notes

Information traditionally qualified as **Notes** is further qualified as follows:

Note Supplemental detail

Tip Offers shortcuts or alternative ways of performing an action; a hint

Guideline

Customary way to perform a procedure

Rule Something you must do; limitations on your actions

Restriction

Indicates certain conditions are not supported; limitations on a product or facility

Requirement

Dependencies, prerequisites

Result Indicates the outcome

How to read a syntax diagram

This syntax information applies to all commands and statements that do not have their own syntax described elsewhere.

The syntax diagram shows you how to specify a command so that the operating system can correctly interpret what you type. Read the syntax diagram from left to right and from top to bottom, following the horizontal line (the main path).

Symbols and punctuation

The following symbols are used in syntax diagrams:

Required operands

Required operands and values appear on the main path line. You must code required operands and values.




Optional values

Optional operands and values appear below the main path line. You do not have to code optional operands and values.



Selecting more than one operand

An arrow returning to the left above a group of operands or values means more than one can be selected, or a single one can be repeated.



Nonalphanumeric characters

If a diagram shows a character that is not alphanumeric (such as parentheses, periods, commas, and equal signs), you must code the character as part of the syntax. In this example, you must code OPERAND=(001,0.001).



Blank spaces in syntax diagrams

If a diagram shows a blank space, you must code the blank space as part of the syntax. In this example, you must code OPERAND=(001 FIXED).



Default operands

Default operands and values appear above the main path line. Multi-site Workload Lifeline uses the default if you omit the operand entirely.



Variables

A word in all lowercase italics is a *variable*. Where you see a variable in the syntax, you must replace it with one of its allowable names or values, as defined in the text.



Syntax fragments

Some diagrams contain syntax fragments, which serve to break up diagrams that are too long, too complex, or too repetitious. Syntax fragment names are in mixed case and are shown in the diagram and in the heading of the fragment. The fragment is placed below the main diagram.



Syntax fragment:



How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this document or any other z/OS Communications Server documentation, do one of the following:

- Go to the z/OS contact page at <http://www.ibm.com/systems/z/os/zos/webqs.html>. You can enter and submit your comments in the form provided at this Web site.
- Send your comments by e-mail to comsvrcf@us.ibm.com. Be sure to include the name of the document, the part number of the document, the version of z/OS Communications Server, and, if applicable, the specific location of the text that you are commenting on (for example, a section number, a page number or a table number).

Summary of changes

Changes made in IBM Multi-site Workload Lifeline Version 2 Release 5, as updated July 2015

This document contains information previously presented in SC27-4653-02, which supports IBM Multi-site Workload Lifeline Version 2 Release 5.

New Information

- “IBM MQ Clusters” on page 13 and “IBM MQ Cluster Workloads” on page 28 have been added.
- “(z/OS systems) Authorizing the Lifeline Advisors and Agents to issue event-driven messages” on page 37 and “(z/OS systems) Enabling the Lifeline Agents to communicate with MQ queue managers” on page 40 have been added.
- “Customizing the NetView for z/OS Message Revision Table when configuring SNA workloads” on page 57 and “Configuring IBM MQ clusters” on page 63 have been added.
- “Diagnosing why a workload is not being distributed to an MQ cluster queue” on page 76 and “Diagnosing workload distribution problems for IBM MQ clusters” on page 78 have been added.
- “mq_manager_list” on page 101 has been added.
- Message “AQS0010I” on page 182 has been added.
- Message “AQS0191I” on page 229 has been added.

Changed Information

- Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1 has been updated.
- Chapter 2, “Preparing to use Multi-site Workload Lifeline,” on page 33 has been updated.
- “Lifeline Advisor configuration file” on page 43, “Lifeline Agent configuration file” on page 43, “Configuring a Lifeline Agent” on page 46, “Migrating IBM Multi-site Workload Lifeline” on page 51, and “Migrating the Lifeline Advisor” on page 53 have been updated.
- Chapter 5, “Diagnosing problems in Multi-site Workload Lifeline,” on page 67 has been updated.
- Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 has been updated.
- “MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information” on page 115, “MODIFY *advisor_proc*,ACTivate command: Activate workloads” on page 142, “MODIFY *advisor_proc*,DEACTivate command: Deactivate workloads” on page 143, and “MODIFY *advisor_proc*,QUIESCE command: Quiesce workloads” on page 146 have been updated.
- “MODIFY *agent_proc*,DISplay command: Display Lifeline Agent information” on page 151 and “MODIFY *agent_proc*,ENABLE command: Enable server applications” on page 158 have been updated.
- “Change the debug level” on page 169, “Dump internal trace records” on page 169, and “Terminate the Linux on System z Management Agent” on page 170 have been updated.

- Chapter 10, "Network Management Interface for Multi-site Workload Lifeline," on page 171 has been updated.
- Messages "AQS0153I" on page 206 through "AQS0156I" on page 207 have been updated.
- Messages "AQS0158I" on page 208 and "AQS0159I" on page 209 have been updated.
- Message "AQS0162I" on page 210 has been updated.
- Messages "AQS0167I" on page 213 and "AQS0168I" on page 215 have been updated.
- Message "AQS0177I" on page 219 has been updated.
- Messages "AQS0179I" on page 221 and "AQS0180W" on page 223 have been updated.
- Message "AQS0185I" on page 225 has been updated.
- Messages "AQS0201I" on page 229 and "AQS0202I" on page 230 have been updated.
- Message "AQS0207I" on page 233 has been updated.
- Messages "AQS0308W" on page 244 through "AQS0310I" on page 246 have been updated.
- Messages "AQS0313I" on page 247 through "AQS0318W" on page 250 have been updated.
- Appendix C, "Sample configuration statements," on page 257 has been updated.

Changes made in IBM Multi-site Workload Lifeline Version 2 Release 5

This document contains information previously presented in SC27-4653-01, which supports IBM Multi-site Workload Lifeline Version 2 Release 0.

Some topics apply only to Lifeline Advisors and the Lifeline Agents that reside on z/OS systems. Such topics have the feature information presented in parentheses as a qualifying element in the topic title. For example, the topic "(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent" applies only to Lifeline Advisors and those Lifeline Agents that reside on z/OS systems. See "Qualifying elements" on page xiii for more information.

New Information

- "Linux on System z Workloads" on page 13 and "SNA workloads" on page 22 have been added.
- "(z/OS systems) Authorizing the Lifeline Agents to monitor SNA applications" on page 39 has been added.
- "Activating a workload for the first time" on page 51, "Activating a workload when the workload is not deactivated" on page 51, "Activating or quiescing a workload when not all first-tier load balancers are registered" on page 52, "Activating a workload when not all Lifeline Agents are connected" on page 52, "Deactivating a workload when not all Lifeline Agents are connected" on page 53, and "Customizing the VTAM started procedure when configuring SNA workloads" on page 57 have been added.
- "Diagnosing problems in the Linux on System z Management Agent" on page 83 has been added.
- Lifeline Advisor configuration statements "management_guest_list" on page 99 and "sna_application_list" on page 106, and Lifeline Agent configuration statement "site_name" on page 113 have been added.

- Chapter 9, “Lifeline Linux on System z Management Agent commands,” on page 165 has been added. This chapter applies to the Linux on System z Management Agent. For information about issuing commands to the z/OS Agent, see Chapter 8, “Lifeline z/OS Agent commands,” on page 151.
- Message “AQS0009I” on page 181 has been added.
- Messages “AQS0207I” on page 233 through “AQS0214I” on page 238 have been added.
- Sample configuration statements “Multi-site Workload Lifeline configuration Example 4” on page 266 and “Multi-site Workload Lifeline configuration Example 5” on page 269 have been added.

Changed Information

- Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1 has been updated.
- Chapter 2, “Preparing to use Multi-site Workload Lifeline,” on page 33 has been updated.
- “Lifeline Agent configuration file” on page 43, “Configuring a Lifeline Agent” on page 46, “Migrating IBM Multi-site Workload Lifeline” on page 51, “Configuring and starting syslogd” on page 55, “(z/OS systems) Customizing the TCP/IP profiles of the TCP/IP stacks” on page 55, and “Configuring external load balancers” on page 58 have been updated.
- Chapter 4, “Starting the Lifeline Advisors and Agents,” on page 65 has been updated.
- Chapter 5, “Diagnosing problems in Multi-site Workload Lifeline,” on page 67 has been updated.
- Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 has been updated.
- “MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information” on page 115, “MODIFY *advisor_proc*,ACTivate command: Activate workloads” on page 142, “MODIFY *advisor_proc*,DEACTivate command: Deactivate workloads” on page 143, “MODIFY *advisor_proc*,DUMPLog command: Dump the internal debug records” on page 146, and “MODIFY *advisor_proc*,QUIESCE command: Quiesce workloads” on page 146 have been updated.
- “MODIFY *agent_proc*,DISplay command: Display Lifeline Agent information” on page 151 and “MODIFY *agent_proc*,DUMPLog command: Dump the internal trace records” on page 163 have been updated.
- Chapter 10, “Network Management Interface for Multi-site Workload Lifeline,” on page 171 has been updated.
- Messages “AQS0004I” on page 179 through “AQS0008I” on page 181 have been updated.
- Messages “AQS0105I” on page 184 through “AQS0111I” on page 187 have been updated.
- Messages “AQS0114I” on page 189 through “AQS0118I” on page 192 have been updated.
- Message “AQS0120I” on page 192 has been updated.
- Messages “AQS0122I” on page 193 and “AQS0123I” on page 194 have been updated.
- Message “AQS0127W” on page 195 has been updated.
- Message “AQS0130I” on page 197 has been updated.
- Message “AQS0132I” on page 198 has been updated.

- Messages “AQS0135W” on page 200 through “AQS0140I” on page 202 have been updated.
- Message “AQS0144I” on page 204 has been updated.
- Message “AQS0147I” on page 205 has been updated.
- Messages “AQS0156I” on page 207 through “AQS0159I” on page 209 have been updated.
- Message “AQS0162I” on page 210 has been updated.
- Messages “AQS0167I” on page 213 and “AQS0168I” on page 215 have been updated.
- Messages “AQS0174I” on page 218 through “AQS0176I” on page 219 have been updated.
- Messages “AQS0179I” on page 221 and “AQS0180W” on page 223 have been updated.
- Message “AQS0182I” on page 225 has been updated.
- Messages “AQS0185I” on page 225 through “AQS0188W” on page 227 have been updated.
- Messages “AQS0201I” on page 229 and “AQS0202I” on page 230 have been updated.
- Messages “AQS0204I” on page 231 through “AQS0206I” on page 232 have been updated.
- Messages “AQS0300I” on page 239 through “AQS0303I” on page 242 have been updated.
- Messages “AQS0308W” on page 244 through “AQS0318W” on page 250 have been updated.
- Appendix A, “A sample started procedure for Lifeline Advisors,” on page 253 has been updated.
- Appendix B, “A sample started procedure for Lifeline z/OS Agents,” on page 255 has been updated.
- Appendix C, “Sample configuration statements,” on page 257 has been updated.

Removed Information

- Message AQS0191I has been removed.
- Message AQS0200I has been removed.

Changes made in IBM Multi-site Workload Lifeline Version 2 Release 0, as updated June 2014

This document contains information previously presented in SC27-4653-00, which supports IBM Multi-site Workload Lifeline Version 2 Release 0. This document contains minor maintenance updates.

Changed information

- “Diagnosing problems when the Lifeline Agent connection terminates” on page 81 has been updated.
- “MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information” on page 115 has been updated.
- “MODIFY *advisor_proc*,ACTivate command: Activate workloads” on page 142 has been updated.
- NMI message format has been updated.

Changes made in IBM Multi-site Workload Lifeline Version 2 Release 0

This document contains information previously presented in SC23-8600-03, which supports IBM Multi-site Workload Lifeline Version 1 Release 1. This document contains minor maintenance updates and updates pertaining to Active/Query workload support.

New Information

- “ Authorizing the Lifeline Advisor to accept secure peer Advisor connections” on page 35, “(z/OS systems) Authorizing the Lifeline Advisor to accept secure Agent connections” on page 36, and “Authorizing the Lifeline Advisor to accept secure Load Balancer connections” on page 36 have been added.
- “Migrating IBM Multi-site Workload Lifeline” on page 51, “Migrating the Lifeline Advisor” on page 53, and “Migrating the Lifeline Agent” on page 54 have been added.
- Messages “AQS0007I” on page 180, “AQS0008I” on page 181, “AQS0106I” on page 185 have been added.
- Messages “AQS0201I” on page 229 through “AQS0206I” on page 232 have been added.
- Messages “AQS0300I” on page 239 through “AQS0318W” on page 250 have been added.

Changed information

- Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1 has been updated.
- “(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent” on page 33 has been updated.
- Chapter 3, “Configuring Multi-site Workload Lifeline,” on page 43 has been updated.
- “Diagnosing problems when a workload is not distributed” on page 74 and “Diagnosing workload distribution problems” on page 77 have been updated.
- Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 has been updated.
- Chapter 7, “Lifeline Advisor commands,” on page 115 has been updated.
- Chapter 10, “Network Management Interface for Multi-site Workload Lifeline,” on page 171 has been updated.
- The following messages have been updated:
 - Message “AQS0105I” on page 184 has been updated.
 - Message “AQS0124I” on page 194 has been updated.
 - Message “AQS0127W” on page 195 has been updated.
 - Messages “AQS0135W” on page 200 through “AQS0137I” on page 201 have been updated.
 - Message “AQS0153I” on page 206 has been updated.
 - Message “AQS0156I” on page 207 has been updated.
 - Messages “AQS0158I” on page 208 and “AQS0159I” on page 209 have been updated.
 - Message “AQS0162I” on page 210 has been updated.
 - Message “AQS0164W” on page 212 has been updated.

- Messages “AQS0166I” on page 213 through “AQS0168I” on page 215 have been updated.
- Messages “AQS0173I” on page 217 through “AQS0178I” on page 220 have been updated.
- Message “AQS0180W” on page 223 has been updated.
- Message “AQS0182I” on page 225 has been updated.
- Message “AQS0185I” on page 225 has been updated.
- Message “AQS0188W” on page 227 has been updated.
- Message “AQS0189W” on page 228 has been updated.
- Message AQS0191I has been updated.
- Appendix C, “Sample configuration statements,” on page 257 has been updated.

Removed information

- Message AQS0152I has been removed.
- Message AQS0154I has been removed.

Changes made in IBM Multi-site Workload Lifeline Version 1 Release 1

This document contains information previously presented in SC23-8600-01, which supports IBM Multi-site Workload Lifeline Version 1 Release 1. This document contains minor maintenance updates.

New Information

- “Intermediary nodes” on page 13 has been added.
- Lifeline Advisor configuration statement “intermediary_node_list” on page 94 has been added.
- Message “AQS0190I” on page 228 has been added.
- Sample configuration statement “Multi-site Workload Lifeline configuration Example 2” on page 261 has been added.

Changed information

- Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1 has been updated.
- Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 has been updated.
- “Lifeline Advisor configuration statements” on page 85 has been updated.
- “agent_id_list” on page 90 has been updated.
- Message “AQS0158I” on page 208 has been updated.
- Message “AQS0159I” on page 209 has been updated.
- Message “AQS0166I” on page 213 has been updated.
- Appendix C, “Sample configuration statements,” on page 257 has been updated.

Changes made in IBM Multi-site Workload Lifeline Version 1 Release 1

This document contains information previously presented in SC23-8600-00, which supports IBM Multi-site Workload Lifeline Version 1 Release 1. This document contains minor maintenance updates.

New Information

- Lifeline Advisor configuration statement “peer_advisor_id” on page 103 has been added.
- Lifeline Advisor configuration statement “advisor_id_list” on page 111 has been added.
- Message “AQS0190I” on page 228 has been added.
- Message AQS0191I has been added.

Changed information

- “External load balancers” on page 10 has been updated.
- “(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent” on page 33 has been updated.
- “(z/OS systems) Configuring Lifeline Agents to automatically restart” on page 50 has been updated.
- “cross_sysplex_list” on page 91 has been updated.
- “advisor_id” on page 110 has been updated.
- Message “AQS0006I” on page 180 has been updated.
- Message “AQS0116I” on page 190 has been updated.
- Message “AQS0125I” on page 195 has been updated.
- Message “AQS0147I” on page 205 has been updated.
- Message “AQS0168I” on page 215 has been updated.

Chapter 1. Overview of Multi-site Workload Lifeline

Multi-site Workload Lifeline provides the ability to load balance TCP/IP workload requests across two sites at unlimited distances, when both sites are running the same server applications and have the same data sources. To load balance workload requests, Multi-site Workload Lifeline either defines two tiers of load balancing or uses IBM MQ clusters.

A workload consists of one or more groups of applications that access the same data sources to provide related services for a client. All group members for a workload are configured as follows:

- In one or more load balancer application groups within external load balancers as server applications
- In one or more MQ cluster queue groups within Multi-site Workload Lifeline as MQ cluster queues
- A combination of load balancer application groups and MQ cluster queue groups

The server applications or MQ cluster queues in these groups provide the same service for that client.

Rule: All subsequent use of the term “workload groups” refers to both load balancer application groups configured in external load balancers and MQ cluster queue groups configured in Multi-site Workload Lifeline.

Each server application or MQ cluster queue within a workload group can provide the same services as the other applications or MQ cluster queues in the group. You can group server applications or MQ cluster queues that access different data sources into their own workload to provide more granular control of each group of applications or MQ cluster queues.

For a workload consisting of server applications, a first-tier load balancer registers application groups that contain members that are in both sites. These members are considered second-tier load balancers. These members can be external load balancers, sysplex distributor distributed DVIPAs, or intermediary nodes. When the first-tier load balancer distributes connections for an Active/Standby workload, it distributes work only to the second-tier load balancers in one active site. When the first-tier load balancer distributes connections for an Active/Query workload, it distributes work to the second-tier load balancers in the active or both active sites.

An external load balancer, operating as a second-tier load balancer, registers groups of server applications that are on the same site. As this external load balancer receives a workload connection request from a first-tier load balancer, the second-tier external load balancer distributes the request to one of the server applications in the matching application group.

A workload that uses MQ cluster queues is called an MQ cluster workload. An IBM MQ cluster is used to distribute messages for the workload to the MQ queue managers within the MQ cluster. When the MQ cluster distributes connections for an Active/Standby workload, it distributes messages only to the MQ queue managers on the active site.

Rule: MQ cluster workloads can be configured only as Active/Standby workloads.

Multi-site Workload Lifeline consists of Lifeline Advisors and Lifeline Agents. Lifeline Advisors run as processes on z/OS, and Lifeline Agents run as processes on z/OS and Linux on System z Management guests. The primary Lifeline Advisor communicates with external load balancers, if applicable, one or more Lifeline Agents, and possibly a secondary Lifeline Advisor. Lifeline Advisor uses the Server/Application State Protocol (SASP) when communicating with external load balancers.

Each workload that is configured to Multi-site Workload Lifeline is classified as an Active/Standby or Active/Query workload. By definition, an Active/Standby workload can be distributed to only one site at a time, because the workload might contain transactions that update databases. By contrast, an Active/Query workload can be distributed to both sites at the same time, because the workload contains only transactions that query databases. Multi-site Workload Lifeline enforces the single or dual site distribution. If more than one workload is configured, some of the workloads can be distributed to one site, some workloads can be distributed to the other site, and the remaining workloads can be distributed to both sites.

An Active/Query workload is associated with an Active/Standby workload. The database sources that are updated by an Active/Standby workload's transactions are the same database sources that are queried by an Active/Query workload's transactions. Database replication ensures that database updates that occur on the active site are replicated to the standby site. Therefore, the Active/Query workload can query the database on either site. A metric called average replication latency state is used to determine whether the database on the standby site is treated as stale. Average database replication latency is a measure of how long it takes for database updates on the active site to be replicated to the standby site. A configured threshold for average replication latency is specified on the data replication product and is used to determine the average database replication state. When this threshold is exceeded, Lifeline receives an acute latency state, and the database on the standby site is considered stale.

A workload is unavailable on a site if its workload group is unavailable.

When a single workload defines multiple workload groups for an Active/Standby workload, the services that each group provides depend on each other. As a result, if one of the groups is unavailable on that site, the workload is not available on that site.

Rule: For an Active/Query workload, only a single load balancer application group can be defined.

For an Active/Standby workload, a workload failure occurs if at least one of the workload groups on the site where the workload is active has been continuously unavailable for the failure detection interval that is configured for the workload.

An Active/Standby workload is unavailable if either of the following conditions is true:

- All of the workload's server applications within a load balancer application group are unavailable.
- All instances of the workload's MQ cluster queue within an MQ cluster queue group are unavailable.

A server application is considered unavailable if any of the following conditions is true:

- The server application is not active.
- The server application is active but not fully operational.
- The system on which the server application runs is resource constrained.

An MQ cluster queue instance is considered unavailable if any of the following conditions is true:

- The MQ queue manager that hosts the MQ cluster queue is not active.
- The MQ queue manager is active but has either not defined the MQ cluster queue or the MQ cluster queue is not enabled to receive messages.
- The system on which the MQ queue manager runs is resource constrained.

For an Active/Query workload, a workload failure occurs if its load balancer application group on the sites where the workload is active has been continuously unavailable for the failure detection interval that is configured for the workload.

An Active/Query workload is unavailable if all of its server applications within the load balancer application group are unavailable or the database on the standby site is stale. A server application is considered unavailable if any of the following conditions is true:

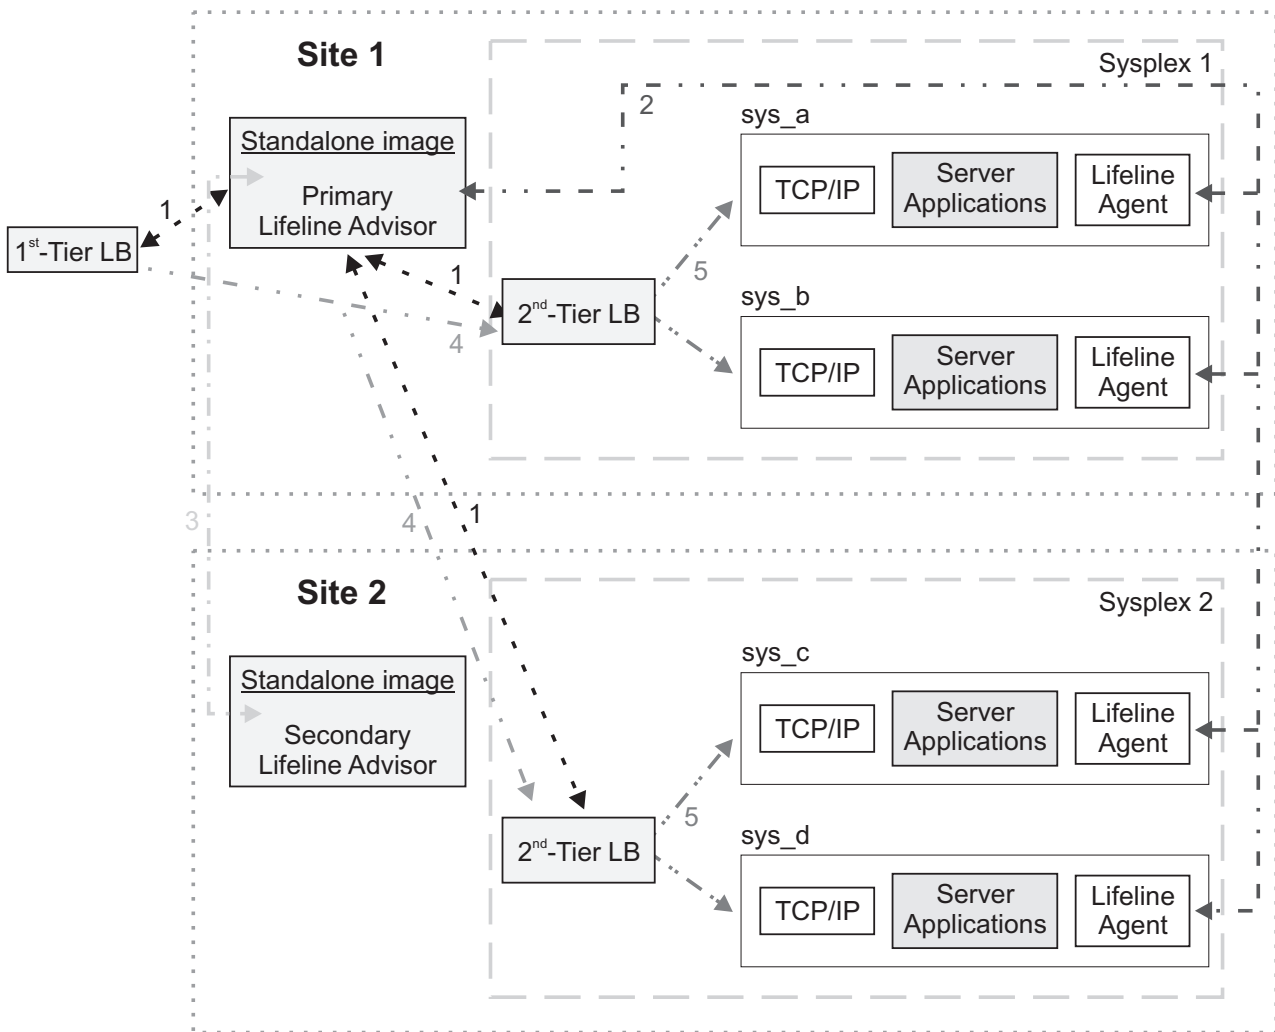
- The server application is not active.
- The server application is active but not fully operational.
- The system on which the server application runs is resource constrained.

Rules:

- When this information describes z/OS target systems, *site* and *sysplex* are used interchangeably.
- All z/OS target systems in a particular site must belong to the same sysplex.
- Linux on System z target systems are not configured in a sysplex. However, they will belong to the same site as the co-located z/OS target systems.

In a typical configuration, the first-tier load balancer determines which site to route a new workload request to and forwards the request to a second-tier load balancer. The second-tier load balancer determines which server application within the site to route the workload request to and forwards it to the z/OS system where the server application is running. Multi-site Workload Lifeline provides the different tiers with routing recommendations based on the availability and health of the server applications and the z/OS systems within the site.

Multi-Site Workload Lifeline



- | | |
|--|--|
| <p>1. Lifeline Advisor 1st and 2nd tier load balancers SASP flows</p> <p>2. Lifeline Advisor to Lifeline Agents communication</p> <p>3. Primary to Secondary Lifeline Advisor state exchange</p> | <p>4. 1st to 2nd tier connection load balancing</p> <p>5. 2nd-tier to server instance connection load balancing</p> |
|--|--|

Figure 1. Relationships between external load balancers, Lifeline Advisors, and Lifeline Agents

In Figure 1, the Lifeline Agents running on each z/OS system in the two sites, Sysplex1 and Sysplex2, monitor the health of the server applications and the system where the Agent is active. The Lifeline Agents periodically send this information to the primary Lifeline Advisor, indicated by the network flows labeled **2**. The primary Lifeline Advisor uses this information to calculate routing recommendations for the workloads that use these server applications. The Lifeline Advisor periodically sends recommendations to first-tier load balancers about which site to route workload requests to and recommendations to second-tier load balancers about which server application within a site to route the workload requests to, indicated by the network flows labeled **1**.

The primary Lifeline Advisor also periodically notifies the secondary Lifeline Advisor about any changes in the workload states, indicated by the network flows labeled **3**, so that the secondary Advisor can take over primary Advisor responsibilities if the current primary Advisor becomes not-operational.

When workload requests are handled by the first-tier load balancers, based on routing recommendations from the Lifeline Advisor, the load balancer selects a site and forwards the request to the second-tier load balancer in that site, indicated by the network flows labeled **4**. When workload requests are handled by the second-tier load balancers, based on routing recommendations from the Lifeline Advisor, the load balancer selects a server application in the site and forwards it to the TCP/IP stack running on the z/OS system, indicated by the network flows labeled **5**. When these second-tier load balancers are external, they receive routing recommendations directly from the Lifeline Advisor. When these second-tier load balancers are internal, also known as z/OS sysplex distributors, they reside on a z/OS system in the site, and receive routing recommendations directly by using z/OS Communications Server functions. For more information about the z/OS sysplex distributor, see *z/OS Communications Server: IP Configuration Guide*.

In another configuration, the second-tier load balancer is not responsible for routing workload requests to a server application within a site. In this case, load balancing within a site will not be accomplished by using a second-tier load balancer such as the z/OS sysplex distributor. Instead, it will be performed by some other node, which is called an intermediary node.

In Figure 2 on page 6, static VIPA addresses might be the preferable method to reach target server applications rather than distributable dynamic VIPAs. These static VIPAs are considered the intermediary node. The first-tier load balancer will distribute workload requests directly to the intermediary node (that is, to server applications using static VIPAs) rather than a second-tier load balancer.

In Figure 3 on page 7, a customized load balancing solution that does not communicate with the Lifeline Advisor might be the preferable method to reach target server applications rather than distributable dynamic VIPAs. These load balancers are considered the intermediary node. The first-tier load balancer will distribute workload requests directly to the intermediary node (that is, to the customized load balancing solution) rather than a second-tier load balancer. The intermediary node will then distribute workload requests to the server applications within the site.

In both cases, the z/OS sysplex distributor is configured as a second-tier load balancer, though it will not distribute workload requests. The z/OS sysplex distributor is used to dynamically determine the set of server applications for each workload, thus these server application instances can be monitored. Multi-site Workload Lifeline provides routing recommendations only to the first-tier load balancer.




Figure 2. Relationships between external load balancers, internal intermediary nodes, Lifeline Advisors, and Lifeline Agents

In Figure 2, the Lifeline Agents, Lifeline Advisors, and first-tier load balancers communicate as is shown in Figure 1 on page 4. However, when a first-tier load balancer handles workload requests based on routing recommendations from the Lifeline Advisor, the load balancer selects a site and forwards the request directly to the intermediary node as indicated by the network flows labeled **4**. In this case, the intermediary node is the TCP/IP stack running on the z/OS system that owns the static VIPA. No second-tier load balancers distribute workload requests within the site. However, these second-tier load balancers that are configured as z/OS sysplex distributors are used to monitor the server applications for the workload.




Figure 3. Relationships between external load balancers, external intermediary nodes, Lifeline Advisors, and Lifeline Agents

In Figure 3, the Lifeline Agents, Lifeline Advisors, and first-tier load balancers communicate as is shown in Figure 1 on page 4. However, when a first-tier load balancer handles workload requests based on routing recommendations from the Lifeline Advisor, the load balancer selects a site and forwards the requests directly to the intermediary node as indicated by the network flows labeled **4**. In this case, the intermediary node is the customized load balancing solution within that site. When the intermediary node handles workload requests, the intermediary node selects a server application in the site and forwards it to the TCP/IP stack that is running on the z/OS system, as indicated by the network flows labeled **5**. No second-tier load balancers distribute workload requests within the site. However, these second-tier load balancers that are configured as z/OS sysplex distributors are used to monitor the server applications for the workload.

See Figure 4 on page 14 in "Overview of Linux on System z Workloads" on page 13 for an example of how a typical workload that contains an application tier on Linux on System z target systems is configured with Multi-site Workload Lifeline.

See Figure 5 on page 23 in "Overview of Lifeline in a SNA workload environment" on page 22 for an example of how a typical workload that contains SNA applications is configured with Multi-site Workload Lifeline.

See Figure 6 on page 29 in “Overview of MQ Cluster Workloads” on page 28 for an example of how a typical workload that uses an IBM MQ cluster is configured with Multi-site Workload Lifeline.

Lifeline Advisors

For a workload consisting of server applications, the primary Lifeline Advisor provides recommendations that are used by a first-tier external load balancer to determine how to route new requests to a second-tier load balancer to select a target site. The second-tier load balancer then routes each request to a target z/OS system in the site to select a target application.

Rule: In this information, whenever the term Lifeline Advisor or Advisor is mentioned, it is referring to the primary Lifeline Advisor. Any references to the secondary Lifeline Advisor are indicated by the terminology secondary Lifeline Advisor or secondary Advisor.

The following are the supported types of second-tier load balancers:

- The second-tier load balancer can be an external load balancer. It uses the Lifeline Advisor recommendations to determine how to route these new requests to the target applications in the z/OS site.
- The second-tier load balancer can be an intermediary node. It uses its own recommendations to determine how to route these new requests to the target applications in the z/OS site.
- The second-tier load balancer can be a z/OS sysplex distributor. It creates its own recommendations that are similar to the ones that the Lifeline Advisor creates to determine how to route these new requests to the target applications in the z/OS site.

The Lifeline Advisor maintains at least two, and up to four, listening sockets or ports; one is for the Lifeline Agents to connect to, up to two are for the external load balancers to connect to, and optionally one is for the secondary Lifeline Advisor to connect to. There are separate IPv4 and IPv6 listening sockets for external load balancers. If your TCP/IP stack is not IPv6 enabled, you will not be able to use the IPv6 listening socket.

The Lifeline Advisor maintains an access control list that allows external load balancers, the secondary Lifeline Advisor, and the Lifeline Agents to connect to it. The access control list specifies the remote IP address of the connecting external load balancers, the remote IP address and port of the secondary Lifeline Advisor, and the Lifeline Agents that are allowed to connect to the Lifeline Advisor.

For a workload consisting of MQ cluster queues, the primary Lifeline Advisor provides recommendations that are used by Lifeline Agents to control the site and target system where MQ cluster queues are enabled for messages. If the site is not active, the MQ cluster queues on that site are disabled for receiving messages. If the site is active, the MQ cluster queues on that site are enabled for receiving messages.

Recommendations provided by the Lifeline Advisor

The recommendations provided by the Lifeline Advisor are dynamic, and can change as the conditions of the target sites, systems, and server applications or MQ cluster queues change. The recommendations include the following information:

- An indication of whether the target server application or MQ cluster queue, the target system, or the target site is currently active. The Advisor uses this information to exclude systems or sites that are not active or do not have the required server applications running or MQ cluster queue available
- For target server applications or MQ cluster queues that reside on z/OS, z/OS Workload Management (WLM) server-specific recommendations. WLM recommendations provide a relative measure of the ability of a target system to handle new workload, as compared to other target systems across both sites. They also provide a measure of the performance of individual server applications or MQ queue managers that host the MQ cluster queue, as compared to the WLM policy goals specified for that workload. The WLM recommendations are derived from the following metrics that are based on the importance of the workload, as defined by the WLM policy:
 - Displaceable general CPU capacity of each system
 - System z Application Assist Processor (zAAP) capacity
 - System z Integrated Information Processor (zIIP) capacity
 - The proportion of each type of CPU that the server application is using
- For target TCP server applications that reside on z/OS, an indication of server application health from a TCP/IP perspective. TCP/IP statistics for target applications are monitored to determine whether specific server applications are encountering problems that prevent the applications from keeping up with the current workload. For example, is a target TCP server application keeping up with TCP connection requests? Are requests being rejected because the backlog queue is full? In such scenarios, the recommendations passed back to the external load balancers are adjusted appropriately, so that the load balancer can direct fewer connections to any server application that is experiencing these problems. These recommendations are provided for TCP server applications and are referred to as Communications Server weights in this information.

Lifeline Agents

The Lifeline Agent receives from the Lifeline Advisor a list of server applications that are configured to receive workload requests from registered external load balancers. The Lifeline Agent gathers data about these server applications and returns the information to the Advisor.

The Lifeline Agent also receives from the Lifeline Advisor, a list of MQ queue managers and MQ cluster queues that are configured to receive workload messages. The Lifeline Agent gathers data about these MQ queue managers and MQ cluster queues and returns the information to the Advisor.

Lifeline Agents can reside on z/OS systems or Linux on System z Management Guests.

- For TCP/IP workloads, the Lifeline Agents reside on z/OS. The Agent monitors the health and availability of both the z/OS system and the workload's TCP/IP server applications on that system.
- For MQ cluster workloads, the Lifeline Agents reside on z/OS. The Agent monitors the health and availability of both the z/OS system and the workload's MQ queue managers and MQ cluster queues on that system.
- For SNA workloads, the Lifeline Agents reside on z/OS. The Agent monitors the health and availability of the z/OS system and the availability of the workload's SNA applications on that system.

- For Linux on System z workloads, the Lifeline Agents reside on both z/OS and Linux on System z. A Linux on System z workload consists of server applications residing on Linux on System z guests, and the database subsystem or SNA applications being accessed by these server applications, residing on z/OS systems.
 - The Agent on Linux on System z runs on a Linux on System z Management Guest. The Agent monitors the availability of the Linux on System z guests configured for the workload that reside on the same z/VM system as the Management Guest. The Agent also monitors the availability of the server applications that run on each of these guests.
 - The Agent on z/OS monitors the health and availability of both the z/OS system and the workload's database subsystem or SNA applications.
 - The Agent on Linux on System z does not monitor any MQ queue managers and MQ cluster queues on Linux on System z guests for MQ cluster workloads.

See “Linux on System z Workloads” on page 13 for details about Lifeline Agents that reside on Linux on System z.

Rule: In this information, the term *z/OS Agent* refers to the Lifeline Agent that resides on z/OS. The term *Linux on System z Management Agent* refers to the Lifeline Agent that resides on the Linux on System z Management Guest. The term *Agent* by itself refers to both the z/OS Agent and the Linux on System z Management Agent.

The Lifeline Agent maintains a single connection to the Lifeline Advisor to communicate status information. This can be either an IPv4 or IPv6 connection. If the TCP/IP stack where the Agent resides is not IPv6-enabled, you will not be able to create the IPv6 connection.

See the Recommendations provided by the Lifeline Advisor section in “Lifeline Advisors” on page 8 for the set of metrics gather by the Agents about the server applications, MQ queue managers and MQ cluster queues, and z/OS system being monitored.

External load balancers

For workloads consisting of server applications, external load balancers are required to distribute workload requests. External load balancers use Server/Application State Protocol (SASP) to obtain recommendations and topology information related to server applications and systems in a clustered environment.

Rule: First-tier load balancers must be external load balancers that are SASP-enabled.

A second-tier load balancer can be an external load balancer that is SASP-enabled, an intermediary node, or a z/OS sysplex distributor on a z/OS system in the site.

External load balancers can obtain detailed information regarding the state of target z/OS applications and systems by communicating with the Lifeline Advisor using the SASP protocol. External load balancers use SASP to do the following things:

- Obtain detailed recommendations about how to distribute workload requests between the two sites, based on the current configuration and workload conditions.

- Obtain detailed recommendations about how to distribute workload requests in the site, based on the current configuration and workload conditions.
- Perform workload load balancing that is optimized for your two sites.

In some configurations, multiple server applications may be listening on the same port. When the TCP connection requests reach a target TCP/IP stack for these server applications sharing the same port, the connections are load balanced by TCP/IP across the multiple server application instances. How this load balancing is performed depends on whether the SHAREPORT or SHAREPORTWLM parameter is specified on the PORT statement. For more details about the PORT statement, see *z/OS Communications Server: IP Configuration Reference*. When you use Multi-site Workload Lifeline, all instances of this server application that share the same port on a target system are reported to external load balancers through a single member entry.

You might be able to customize the features of the way that the external load balancer communicates with the Lifeline Advisor. The SASP protocol defines two features of the external load balancer implementation that you might be able to configure:

- Whether the external load balancer polls the Lifeline Advisor for updated data or whether the updated data is pushed to the external load balancer.
- Whether only members that have updated data should be sent to the external load balancer or whether all members should be sent to the load balancer, regardless of whether their data has changed.

If you configure the external load balancer to request that the Lifeline Advisor push the updated information to the load balancer, the Lifeline Advisor updates the load balancer at least once every update interval. If you configure the external load balancer to poll the Lifeline Advisor for updated information, the Lifeline Advisor passes a recommended polling interval to the load balancer. However, you can configure the external load balancer to disregard this guideline.

Rule: Always configure first-tier load balancers to have updated information pushed to the load balancer.

Guideline: You can configure external second-tier load balancers to either poll for updated information or have updated information pushed to the load balancer.

External load balancers typically use a cluster IP address to represent the server applications that are being load balanced. Client applications use this cluster IP address as the destination IP address for their requests.

In addition, some external load balancers might be able to use either directed or dispatch mode when they forward packets to their destinations.

Directed mode

When an external load balancer uses directed mode to forward packets to their destinations, it can use NAT to convert the destination IP address (the cluster IP address) to an IP address owned by the target z/OS system. When IP packets for these connections are sent back to clients, the external load balancer converts the source IP address (the IP address of the target z/OS system) back to the cluster IP address that the application used on its request. The load balancer either modifies the destination IP address (server NAT) in the packet itself or modifies both the destination and source IP addresses (server NAT and client NAT) in the packet.

The packet must return through the same external load balancer, which recognizes the changes and does the reverse mapping. Then a packet can flow from the original destination to the original source.

Rule: Always configure first-tier load balancers to use directed mode. These load balancers route requests only to second-tier load balancers.

Guideline: You can configure external second-tier load balancers to use either directed or dispatch mode, although directed mode is recommended to avoid having to make server application changes.

Dispatch mode

When an external load balancer uses dispatch mode to forward packets to their destinations, the destination IP addresses for incoming IP packets are not changed. Instead, the external load balancer forwards the packet to a target z/OS system by using the MAC address of a network adapter on that system. The receiving z/OS system inspects the destination IP address of the packet, and accepts the packet if the IP address matches one of the IP addresses in its HOME list. As a result, with dispatch mode, all target z/OS systems must have the external load balancer's cluster IP address defined in their HOME list. However, these addresses must not be advertised externally through dynamic routing protocols. To prevent the addresses from being advertised externally through dynamic routing protocols, define these IP addresses as loopback addresses on z/OS.

Dispatch mode eliminates the need to perform network address translation (NAT). Dispatch mode has the following restrictions, which ensure that the second-tier external load balancer functions correctly:

Restrictions:

- An OSA device can be shared among logical partitions (LPARs) only if virtual MAC (VMAC) addressing is configured for each TCP/IP target stack that shares the OSA device.
- All target applications must bind to the IP address specified by INADDR_ANY or IN6ADDR_ANY, and the cluster IP address of the 2nd-tier load balancer must be defined to the stack. To ensure that this cluster IP address is not advertised, define it as a loopback address.
- First-tier load balancers cannot be configured in dispatch mode.

If you do not meet the previous restrictions, load balancing will not be optimal because work will not be routed to some servers.

Internal load balancers

An internal load balancer, also known as z/OS sysplex distributor, resides on a z/OS system in the site. It uses z/OS Communications Server functions to obtain recommendations and topology information related to server applications in a sysplex environment.

For more information about z/OS sysplex distributor, see *z/OS Communications Server: IP Configuration Guide*.

Tip: You can use z/OS sysplex distributor to perform the second-tier load balancing function.

Intermediary nodes

An intermediary node does not communicate directly with the Lifeline Advisor, but it replaces the function of a second-tier load balancer. An intermediary node is used in any of the following cases:

- Where customized load balancing has been built into a workload
- When the server applications reside on Linux on System z
- When the workload contains SNA applications

When customized balancing is being used, the z/OS sysplex distributor must still be defined, but it is used only for obtaining recommendations and topology information that the first-tier load balancer uses to select the site for a workload connection request.

For Linux on System z applications defined as IP intermediary nodes, the z/OS sysplex distributor must still be defined because it is used to monitor and obtain recommendations on the database subsystem that is referenced by the Linux on System z applications.

For SNA applications, the z/OS sysplex distributor is not needed because recommendations are gathered by using a different method.

The intermediary node makes the load balancing decision to select a server application instance within the site.

IBM MQ Clusters

For workloads that route messages to MQ cluster queues, external load balancers are not used to distribute workload messages. An IBM MQ cluster is instead used to distribute messages destined for an MQ cluster queue to an MQ queue manager instance that is configured in the MQ cluster.

For more information about workloads that uses IBM MQ clusters, see “IBM MQ Cluster Workloads” on page 28.

Linux on System z Workloads

Overview of Linux on System z Workloads

Multi-site Workload Lifeline provides the ability to load balance requests for a workload across two sites at unlimited distances that target TCP/IP applications running on Linux on System z guests. Similar to other workloads defined to Lifeline, a first-tier load balancer registers application groups that contain members that are in both sites.

For Linux on System z workloads, these members are defined as intermediary nodes. The first-tier load balancer distributes TCP/IP connections and the transactions that are sent over these connections to the intermediary node servers. In turn, the intermediary node servers create connections with the server applications on z/OS and send transactions to the z/OS application. These z/OS applications can be TCP/IP server applications (for example, a database subsystem) or SNA applications. Lifeline monitors the availability of the server applications on Linux on System z as well as the health and availability of the server applications on z/OS.

Because a workload consists of one or more groups of applications, it is possible for a workload to consist of a group of TCP/IP z/OS applications, a group of TCP/IP Linux on System z applications, a group of MQ cluster queues, and a group of SNA applications that each access the same data sources on z/OS. Each group of applications must be defined in its own load balancer application group or MQ cluster queue group. Different types of applications cannot be defined in the same workload group.

Linux on System z workloads can only be configured to Lifeline when the Linux on System z server applications reside on Linux on System z guests running on z/VM. Lifeline does not support workloads with server applications that run in a native Linux on System z partition. Linux on System z workloads can be defined as Active/Standby workloads or Active/Query workloads.

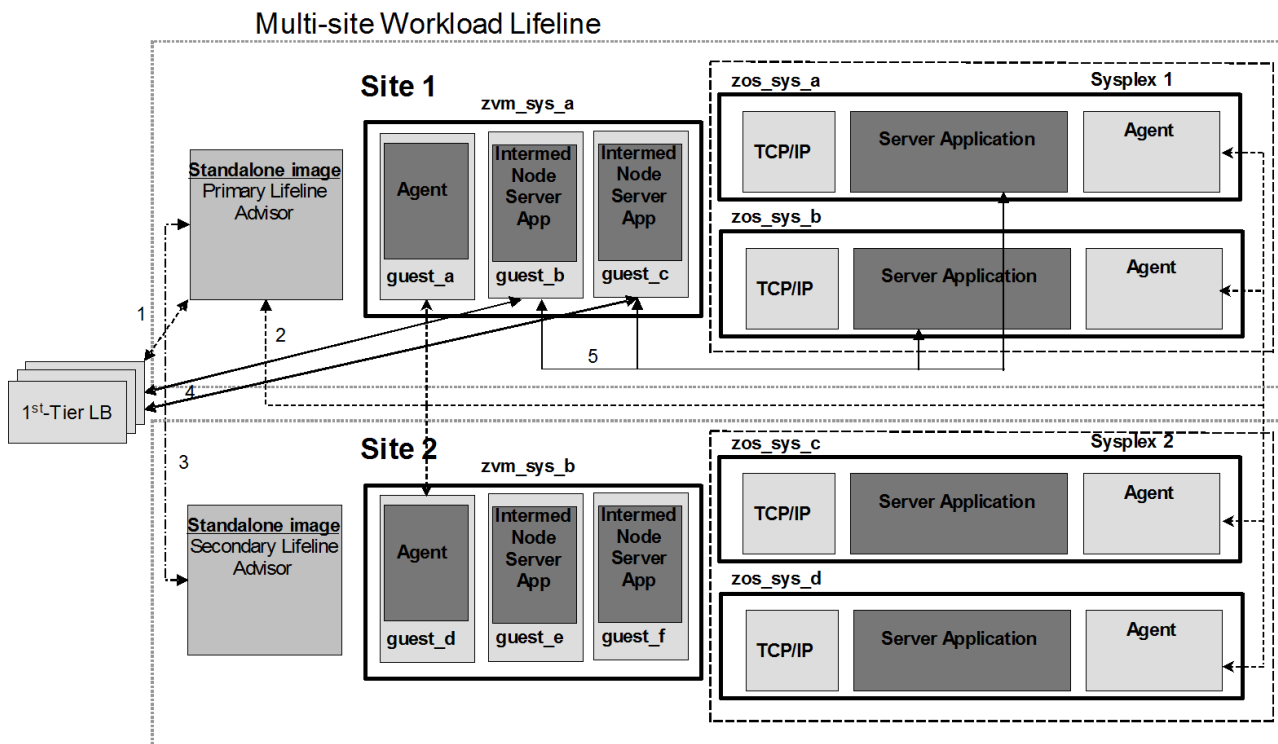


Figure 4. Relationships between external load balancers, Linux on System z intermediary nodes, Lifeline Advisors, and Lifeline Agents for Linux on System z workloads

Figure 4 shows an environment where workloads are processed by Linux on System z server applications. The Lifeline Agents, Lifeline Advisors, and first-tier load balancers communicate as is shown in Figure 1 on page 4. However, the Advisor also communicates with Linux on System z Management Agents in addition to z/OS Agents. When a first-tier load balancer handles workload requests based on routing recommendations from the Advisor, the load balancer selects the active site, which is Site 1 in this example, and forwards the requests directly to the Linux on System z intermediary nodes as indicated by the network flows labeled 4. In this case, the intermediary node is an application-tier Linux on System z server application within that site. The Linux on System z server applications then establish TCP/IP connections to data-tier server applications on z/OS, as indicated by the network flows labeled 5. The Linux on System z server applications typically connect to a z/OS sysplex distributor instance to achieve load balancing across the data-tier server applications within the site.

Monitoring Linux on System z application groups

The first-tier load balancer registers a group of intermediary nodes to the Lifeline Advisor. The Advisor is configured with the z/OS server applications, the list of intermediary nodes, and a corresponding list of Linux on System z guests that use the same workload name. The Advisor sends the intermediary node server group to the Lifeline Linux on System z Management Agents and sends the z/OS applications to the Lifeline z/OS Agents. The Linux on System z Management Agents gather information about the availability of the application-tier intermediary node server applications and return the information to the Advisor. The z/OS Agents gather information about the health and availability of the data-tier server applications and return the information to the Advisor. The Advisor uses the information from both types of Agents to determine the availability of the workload on that site.

See Figure 23 on page 270 for an example of how to configure the Advisor `intermediary_node_list` and `management_guest_list` configuration statements for a Linux on System z workload.

The Linux on System z Management Guest

The Linux on System z Management Guest is a privileged z/VM guest IPLed with Linux on System z. One Linux on System z Management Guest is required per z/VM system that participates in Linux on System z workloads. The Linux on System z Management Agent runs as a Linux on System z daemon on this Management Guest. For more information, see “Management Guest system requirements” on page 16.

The Linux on System z Management Agent

The Linux on System z Management Agent is a Linux on System z daemon that runs on the Linux on System z Management Guest. It monitors the availability of server applications that run on Linux on System z guests on the same z/VM system as the Linux on System z Management Agent. It reports the status of these server applications to the Lifeline Advisor. Only one Linux on System z Management Agent can be active at a time on any given z/VM system. The Linux on System z Management Agent is also responsible for logging off and logging back on monitored Linux on System z guests when the workloads that the guests participate in are deactivated. For more information about the `DEACTIVATE` command, see “`MODIFY advisor_proc,DEACTivate` command: Deactivate workloads” on page 143.

Requirement: The Linux on System z Management Agent has a dependency on the `nmap` package. When the Lifeline Linux on System z RPM is installed, this dependency will be checked. The Linux on System z Management Agent requires a minimum `nmap` version of 4.0 to monitor IPv4 server applications and a minimum `nmap` version of 6.0 to monitor IPv6 server applications.

Restrictions:

- The Linux on System z Management Agent does not support monitoring of UDP applications.
- The Linux on System z Management Agent does not support Transport Layer Security (TLS) connections to the Advisor.

Preparing to use the Linux on System z Management Agent Management Guest system requirements

The Linux on System z Management Guest is supported in a z/VM single system image (SSI) or outside of a single system image.

The Linux on System z Management Guest has several system requirements to allow the Linux on System z Management Agent to function properly.

Requirements:

- The Linux on System z Management Guest must have a minimum of 1 GB of memory.
- The Linux on System z Management Guest must be running a supported Linux on System z distribution. The supported distributions are SUSE Linux Enterprise Server (SLES) 10 SP4 or higher level and Red Hat Enterprise Linux (RHEL) 5 version 5.8 or higher level.
- Monitored server applications must not reside on the Linux on System z Management Guest. The Linux on System z guests for the monitored server applications are logged off when workloads are deactivated.
- The Linux on System z Management Guest must be on the same network as the monitored server applications.
- The /opt/ibm/lifeline/log file system must have 256 MB of available space to prevent truncation of the dumped internal log. If a file system exists for /opt and it does not have 256 MB of available space, an alternate file system can be used as long as a symbolic link is established for /opt/ibm/lifeline/log to point to that file system.
- The Linux on System z Management Guest must have z/VM Class A privileges for the FORCE and XAUTOLOG commands. The Linux on System z Management Agent issues these z/VM commands to log off, log on and re-IPL monitored server applications on Linux on System z guests when workloads are deactivated.

Guideline: For Linux on System z Management Agents, internal trace is used by default. However, syslogd should also be configured to capture ERROR, WARNING, and EVENT log messages from the Linux on System z Management Agent. For more information, see “Configuring and starting syslogd” on page 55.

Rule: Monitored Linux on System z guests can be relocated by using z/VM Live Guest Relocation (LGR). However, the Linux on System z Management Guest must not be relocated. See the publications for the z/VM product for methods to configure z/VM to prevent relocation of the Linux on System z Management Guest.

Monitored server application requirements

To successfully monitor server applications that run on Linux on System z guests, ensure that the following requirement and guidelines are met.

Requirement: Network and host firewalls must not block repeated SYN port scans. The Linux on System z Management Agent relies on SYN port scans to determine the availability of monitored Linux on System z server applications. The Advisor considers the monitored Linux on System z server applications unavailable when firewalls block the SYN port scans.

Guidelines:

- Enable the Linux on System z guests for shutdown signals and define a default system shutdown timeout. This ensures the Linux on System z guests of monitored server applications can shut down the Linux on System z operating system in an orderly manner before the Linux on System z guests are logged off by Lifeline during workload deactivate processing. See the publications for the z/VM product for details.
- Configure the Linux on System z guests with an IPL statement in the z/VM system directories of the guests to load the Linux on System z operating system when the Linux on System z guests are logged on. Linux on System z guests are logged off and restarted by Lifeline during workload deactivate processing. For more information about the DEACTIVATE command, see “MODIFY *advisor_proc*, DEACTivate command: Deactivate workloads” on page 143.
- Configure the Linux on System z operating systems on the monitored Linux on System z guests to automatically start the server applications when the Linux on System z kernel is IPLed.

Installing the Linux on System z Management Agent

The Linux on System z Management Agent is installed on the Linux on System z Management Guest. The recommended method of installation is to use the RPM from the SMP/E installation of the Lifeline product on z/OS. See the program directory that is included with the Multi-site Workload Lifeline installation for information about the location of the RPM and how to transfer the RPM to the Linux on System z Management Guest. The RPM will check prerequisites, create the necessary directories with the proper permissions, and install files and programs in those directories.

Requirement: The Linux on System z Management Agent must be installed from a user ID with root privileges.

Rule: Only one Linux on System z Management Agent is allowed per z/VM system.

The Linux on System z Management Agent is installed as a program named `aqszvm`. The RPM also installs a program named `aqslcmd`, which is used to send commands to the Linux on System z Management Agent. These programs expect to find specific directories and files in locations that are installed by the RPM. If any of these directories or files are moved, symbolic links that represent the original locations must be established to reference the new locations. For more information about the `aqslcmd` program, see Chapter 9, “Lifeline Linux on System z Management Agent commands,” on page 165.

The system administrator needs to be aware of the following files and directories:

`/opt/ibm/lifeline/bin`

The location where the Linux on System z Management Agent program, `aqszvm`, and the `aqslcmd` program are installed. Include this path in the `PATH` environment variable for system administrators.

`/opt/ibm/lifeline/config/agent.conf`

The default configuration file for the Agent. For more information, see “Configuring a Lifeline Agent” on page 46.

/opt/ibm/lifeline/log/agent.trace

The location where internal trace is dumped. For more information, see “Diagnosing problems in the Linux on System z Management Agent” on page 83.

/opt/ibm/lifeline/lib/msg/C/aqsllmsg.cat

The location where the message catalog for the Linux on System z Management Agent can be found.

/opt/ibm/lifeline/man

The location where man pages can be found for aqsagzvm, aqsllcmd, and agent.conf. To display one of these man pages, enter `man -M /opt/ibm/lifeline/man man_page`, where `man_page` is one of these man pages.

Guideline: After you install the Linux on System z Management Agent, it is recommended to configure the Linux on System z Management Guest operating system to automatically start the Linux on System z Management Agent, aqsagzvm, when the Linux kernel is IPLed.

Configuring the Linux on System z Management Agent

The Linux on System z Management Agent requires a configuration file. The default configuration file location is `/opt/ibm/lifeline/conf/agent.conf`. An alternate configuration file can be used by specifying the file and path with the `aqsagzvm -c start` option.

For more information on the Agent configuration file, see “Lifeline Agent configuration statements” on page 109.

Restriction: The Linux on System z Management Agent does not support TLS. If the Advisor is configured to use Application Transparent TLS (AT-TLS) for its Agent connections, the Advisor's AT-TLS policy must ensure that no AT-TLS policy is configured for Linux on System z Management Agent connections.

Starting the Linux on System z Management Agent

The `aqsagzvm` program is used to start the Linux on System z Management Agent. The program must be started by a user with root privileges.

Format

```
▶▶ aqsagzvm [ -c /opt/ibm/lifeline/conf/agent.conf ] [ -c alternate_configuration_file ] [ -h ] [ -? ] ▶▶
```

Parameters

-c *alternate_configuration_file*

Uses the *alternate_configuration_file* instead of the default `/opt/ibm/lifeline/conf/agent.conf` configuration file.

-h, -?

Display a help message and exits.

The Linux on System z Management Agent runs as a daemon and writes all system messages to syslogd.

A man page is available for `aqsagzvm`. To display the man page, enter the following command:

```
man -M /opt/ibm/lifeline/man aqsagzvm
```

Operational characteristics of Linux on System z workloads

Workload failure detection

The first-tier load balancer registers a group of intermediary nodes that the Advisor maps to a set of Linux on System z guest names, a set of Linux on System z server applications, a set of z/OS server applications, and a workload. The workload group's z/OS server applications are monitored by the z/OS agents on the systems where the server applications reside. The Linux on System z intermediary node server applications are monitored by the Linux on System z Management Agents residing on the z/VM systems where the server applications reside.

A load balancer application group for a Linux on System z workload is unavailable on a site when either of the following conditions is true:

- All of the Linux on System z server applications that make up the workload are unavailable on that site.
- All of the z/OS server applications that make up the workload are unavailable on that site.

A Linux on System z workload failure occurs for a site when one of the workload's load balancer application groups is continuously unavailable for the failure detection interval that is configured for that workload.

A Linux on System z workload recovers for a site when all of the workload's load balancer application groups within the workload on that site recover. A load balancer application group for a Linux on System z workload is recovered on a site when at least one of the Linux on System z server applications in the group on that site becomes available and at least one of the z/OS server applications for the group on that site becomes available.

Activating a Linux on System z workload

Linux on System z Management Agents can monitor the availability of Linux on System z server applications, but the Agents do not monitor the active connections to those server applications or terminate the connections. When a Linux on System z workload is activated for the first time, a check for active connections is made to the z/OS Agents on the alternate site for the z/OS server applications, but this check is skipped for the Linux on System z Management Agents. For more information, see “Activating a workload for the first time” on page 51.

Guideline: Before you activate a Linux on System z workload for the first time, verify that no active connections to the server applications that make up the workload on the Linux on System z guests on the alternate site exist. To ensure that the Linux on System z server applications do not have active connections on the alternate site, take either of the following actions:

- Manually determine whether any connections for the workload exist on the monitored Linux on System z guests, and manually terminate them. For example, `netstat` can be issued on the Linux on System z guests to determine

whether connections for the workload exist. If such connections exist, the server applications could be recycled to terminate connections or could be terminated by other means.

- Deactivate the workload before activating it if the workload can tolerate a temporary outage.

Deactivating Linux on System z workloads

Linux on System z Management Agents do not selectively terminate connections to the monitored Linux on System z server applications when a workload is deactivated. Instead, the Linux on System z Management Agents log off the monitored Linux on System z guests that participate in the workload and log them back on. For more information about how to re-IPL the Linux on System z guests and how to restart the server applications, see “Monitored server application requirements” on page 16.

Active/Query workload persistent connection timeout

Restriction: The Linux on System z Management Agent does not monitor how long a persistent Active/Query connection has been established, and therefore, will not terminate these connections after the timeout period has expired. Connections for Active/Query Linux on System z workloads will always behave as if a timeout value of 0 was configured for the workload on the `query_workload_list` Advisor configuration statement. For more information about configuration statements and parameters, see Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85.

Active/Query workload failures

When the acute state for database replication latency is detected for an Active/Query workload, the workload is considered unavailable. Active connections will be reset on that site to prevent queries from accessing stale data. Because the Linux on System z Management Agent does not monitor connections directly, the Linux on System z guests that host these server applications are logged off and logged back on. This process accomplishes the same goal of preventing queries from returning stale data for those workloads. The Active/Query workload will be available on this site when both of the following conditions are true:

- The database is no longer considered stale.
- The Linux on System z guests and server applications are back up and active.

For more information about the database replication latency state, see Chapter 1, “Overview of Multi-site Workload Lifeline,” on page 1.

Diagnosing Linux on System z Management Agent problems

When a Linux on System z server application is not being reported as available, you can use the information in *Diagnosing problems in Multi-site Workload Lifeline* as well as this topic to verify the Linux on System z Management Agent is functioning correctly.

About this task

To diagnose possible problems with the Linux on System z Management Agent, complete the following steps:

Procedure

1. Determine whether the Linux on System z Management Agent has started correctly by using the `ps` command to check whether the `aqszvm` process is running. See the documentation for your Linux distribution for information about the `ps` command.
2. Determine whether the Linux on System z Management Agent is communicating with the Advisor by using the following command on the Linux on System z Management Guest from a user ID with root privileges:
`aqsl1cmd display,members`

Each registered load balancer application group member that is associated with a Linux on System z guest on the z/VM system of the Linux on System z Management Agent will be displayed. Linux on System z guests are associated with registered load balancer application group members in the `management_guest_list` statement of the Advisor's configuration file. For more information about configuration statements and parameters, see Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85.

3. Dump the Linux on System z Management Agent's internal trace and look for `ERROR` or `WARNING` messages in the trace file. The internal trace can be dumped to the `/opt/ibm/lifeline/log/agent.trace` file by issuing the following command on the Linux on System z Management Guest from a user ID with root privileges:
`aqsl1cmd dumplog`

Diagnosing Linux on System z workload problems

When a Linux on System z workload is unavailable, you can use the information in *Diagnosing problems when a workload is not distributed* as well as this topic to verify the Linux on System z workload configuration is correct.

About this task

To diagnose problems with the Linux on System z workload, complete the following steps:

Procedure

1. Determine whether the application-tier Linux on System z server applications for the workload are available on the monitored Linux on System z guests. Use the `netstat` command to display the listening TCP ports on the Linux on System z guests where the server applications are running. Verify that a listening TCP socket on a port and IP address that is compatible with one of the registered members from the first-tier load balancer's load balancer application group exists. See the documentation for your Linux distribution for information about using the `netstat` command.
2. Verify that the application-tier Linux on System z server applications can communicate with the data-tier z/OS server applications. For more information about diagnosing connectivity problems, see *z/OS Communications Server: IP Diagnosis Guide*.
3. Verify that the Lifeline Advisor configuration is correct by checking the following items:
 - Verify that the `intermediary_node_list` entry containing the server application IP address and port is configured with the correct workload name and site.

- Verify that the `management_guest_list` entry containing the Linux on System z guest name where the server application resides is configured with the correct IP address and workload name.

For more information about configuration statements and parameters, see Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85.

4. Verify that one or more data-tier z/OS server applications in the workload is available and one or more application-tier Linux on System z server applications per registered load balancer application group in the workload is available. Display the server applications that make up the workload by issuing the following command on the Advisor:

```
MODIFY procname,DISPLAY,WORKLOAD=workloadname,DETAIL
```

5. If registered load balancer application group IPv6 members are not available, verify that the version of nmap installed on the Linux on System z Management Guest is version 6.0 or higher by issuing the following command on the Linux on System z Management Guest:

```
nmap -V
```

6. Verify that network or host firewalls are not blocking SYN port scans of the application-tier Linux on System z server applications by the Linux on System z Management Agent. Issue the following command on the Linux on System z Management Guest to verify connectivity to a Linux on System z guest's IPv4 address:

```
nmap -sS -send-ip -p portnum ipv4_address
```

Issue the following command on the Linux on System z Management Guest to verify connectivity to a Linux on System z guest's IPv6 address:

```
nmap -sS -6 -send-ip -p portnum ipv6_address
```

where *portnum* is the TCP port the application-tier Linux on System z server application is listening on and *ipv4_address* and *ipv6_address* is the IPv4 or IPv6 address the application-tier Linux on System z server applications is bound to. The ports in the nmap display output should show a status of *open*. If the status for a port is not *open*, verify whether one of the following conditions has occurred:

- If the status of a port is *filtered*, a firewall might be interfering with the SYN scans that the Linux on System z Management Agent performs.
- If the status of a port is *closed*, either a firewall might be interfering with the SYN scans that the Linux on System z Management Agent performs or the application might not be listening on the specified IP address and port.

SNA workloads

Overview of Lifeline in a SNA workload environment

Multi-site Workload Lifeline provides the ability to load balance requests for a workload that targets legacy SNA applications across two sites at unlimited distances. Similar to other workloads defined to Lifeline, a first-tier load balancer registers application groups that contain members that are in both sites.

For SNA workloads, these members are defined as intermediary nodes. The first-tier load balancer distributes TCP/IP connections and the transactions that are sent over these connections to the intermediary node servers. In turn, the intermediary node servers create SNA sessions with the SNA applications in the

same site and send transactions to the SNA applications. These SNA sessions can be established over native SNA networks or Enterprise Extender connections.

Because a workload consists of one or more groups of applications, it is possible for a workload to consist of a group of TCP/IP z/OS applications, a group of TCP/IP Linux on System z applications, a group of MQ cluster queues, and a group of SNA applications that each access the same data sources on z/OS. Each group of applications must be defined in its own load balancer application group or MQ cluster queue group. Different types of applications cannot be defined in the same workload group.

SNA workloads can be configured only in an Active/Standby configuration. When the first-tier load balancer distributes connections on behalf of a SNA workload, it routes the connections only to the intermediary node servers in the active site.

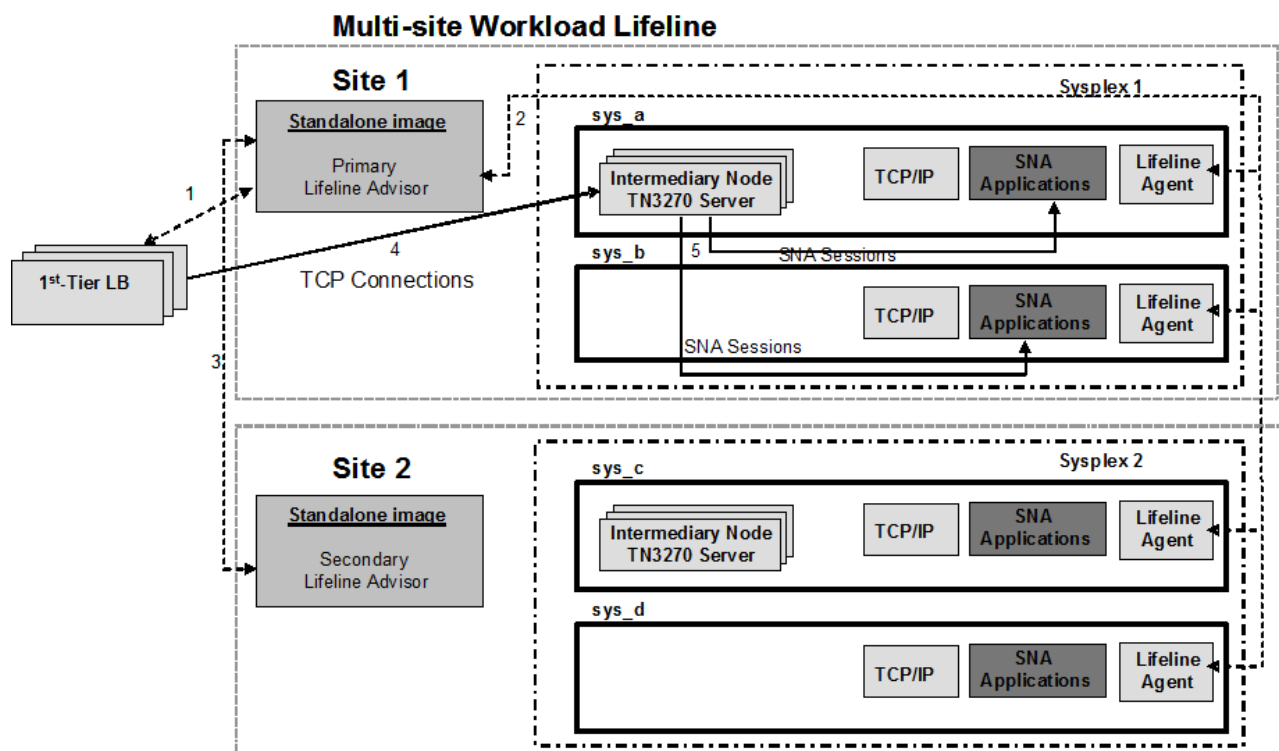


Figure 5. Relationships between external load balancers, internal intermediary nodes, Lifeline Advisors, and Lifeline Agents for SNA workloads

Figure 5 shows an environment in which a workload is processed by SNA applications. The Lifeline Agents, Lifeline Advisors, and first-tier load balancers communicate as is shown in Figure 1 on page 4. When a first-tier load balancer handles TCP/IP workload requests based on routing recommendations from the Advisor, the load balancer selects the active site, which is Site 1 in this example, and forwards the requests directly to the internal intermediary node on that site, indicated by the network flows labeled 4. In this example, the intermediary node is a z/OS Communications Server TN3270 server. When the TN3270 server handles workload requests, it creates a SNA session to the requested SNA application, indicated by the network flows labeled 5. Any subsequent transactions over the TCP/IP connection that are distributed to the TN3270 server are sent over the

existing SNA session to the SNA application. Load balancing across SNA applications within a site can be accomplished by using existing techniques, such as Generic Resource.

Monitoring SNA application groups

The first-tier load balancer registers a group of intermediary node servers to the Lifeline Advisor. The Advisor is configured with the intermediary node servers and a corresponding list of SNA applications that use the same workload name. The Advisor sends the intermediary node server group and its SNA applications to the Lifeline Agents. The Agents gather and return information about the availability of the SNA applications to the Advisor. If the intermediary nodes are internal, the Agents also return the availability information about these intermediary node servers, because these servers are on the same systems as the Agents.

See Figure 20 on page 267 for an example of how to configure the Advisor intermediary_node_list and sna_application_list configuration statements for a SNA workload.

Multi-site Workload Lifeline also supports configurations where the intermediary node servers are external. If these external intermediary node servers reside on Linux on System z guests, they can be monitored by the Linux on System z Management Agents that gather and return the availability information about these intermediary node servers. See “Linux on System z Workloads” on page 13 for details about how Multi-site Workload Lifeline can support these workload types.

Preparing z/OS Agents for monitoring SNA applications

Each Lifeline Agent uses an extended MCS console to monitor SNA applications. See *MVS Programming Authorized Assembler Services Guide* for a description of extended MCS consoles. By using this console, the Agent issues VTAM[®] commands to monitor the status of the SNA applications that make up the workload, and alter the state of these SNA applications. See “(z/OS systems) Authorizing the Lifeline Agents to monitor SNA applications” on page 39 for steps about protecting access to the Agent's extended MCS console.

Guideline: If message flood automation is enabled, the DEFAULT action should be specified as AUTO. This specification allows messages to continue to be sent to the extended MCS console of the Agent. Otherwise, if messages that are destined for the extended MCS console of the Agent are suppressed, the SNA applications might be reported as unavailable, because the Agent cannot monitor them. See *z/OS MVS Initialization and Tuning Reference* for details about the DEFAULT statement in the MSGFLDxx member.

Operational characteristics of SNA workloads

Workload failure detection

The first-tier load balancer registers a group of intermediary node servers which the Lifeline Advisor maps to a set of SNA applications and a workload name. The SNA applications of the load balancer application group are monitored by the Lifeline Agents on the systems where the SNA applications reside. The intermediary node servers for an application group are one of the following types:

Monitored intermediary nodes

- The intermediary node servers are on the systems that are monitored by Lifeline z/OS Agents.

- The intermediary node servers are on the Linux on System z guests that are monitored by Lifeline Linux on System z Management Agents.

Unmonitored intermediary nodes

The intermediary node servers are not on the systems that are monitored by Lifeline Agents.

A load balancer application group for a SNA workload that contains monitored intermediary nodes is unavailable on a site when either of the following conditions occurs:

- All of the SNA applications of the load balancer application group are unavailable on that site.
- All of the intermediary node servers of the load balancer application group are unavailable on that site.

A load balancer application group for a SNA workload that contains unmonitored intermediary nodes is unavailable on a site when all of the group's SNA applications are unavailable on that site.

A SNA workload failure occurs when one of the load balancer application groups of the SNA workload is continuously unavailable for the failure detection interval that is configured for that workload.

The SNA workload recovers for a site when all of the load balancer application groups for the SNA workload on that site recover.

- A load balancer application group for a SNA workload that contains monitored intermediary nodes recovers when both of the following conditions occur:
 - At least one SNA application becomes available on that site.
 - At least one intermediary node server becomes available on that site.
- A load balancer application group for a SNA workload that contains unmonitored intermediary nodes recovers when at least one SNA application becomes available on that site.

Activating a SNA workload

When a SNA workload is activated and the intermediary nodes are internal, monitored servers, the Lifeline Advisor queries the Lifeline Agents on the alternate site to determine whether active TCP/IP connections exist in either of the following situations:

- The SNA workload is activated for the first time.
- The SNA workload is not being activated for the first time and the workload is not in a deactivated state.

If active connections are found, the ACTIVATE command is not successful.

When a SNA workload is activated and the intermediary nodes are external, unmonitored servers:

- If the SNA workload is activated for the first time, the Lifeline Advisor queries the Lifeline Agents on the alternate site to determine whether active SNA sessions exist. If active sessions are found, the ACTIVATE command is not successful.
- If the SNA workload is not being activated for the first time and the workload is not in a deactivated state, the ACTIVATE command is not successful. The SNA workload must be deactivated first.

For more information, see “Activating a workload for the first time” on page 51 and “Activating a workload when the workload is not deactivated” on page 51.

When the intermediary node servers are monitored by Linux on System z Management Agents, see Linux on System z Workloads for more information.

Deactivating a SNA workload

When a SNA workload is deactivated, and the intermediary nodes are internal, monitored servers, the Lifeline Advisor directs the Lifeline Agents on the target site to terminate active TCP/IP connections to the intermediary node servers.

When a SNA workload is deactivated, and the intermediary nodes are external, unmonitored servers, the Lifeline Advisor directs the Lifeline Agents on the target site to terminate active SNA sessions. To do this, the Agents inactivate the SNA application resources that make up the workload.

When the intermediary node servers are monitored by Linux on System z Management Agents, see “Linux on System z Workloads” on page 13 for more information.

Reactivating a SNA workload

When a SNA workload that was deactivated is activated to the same site, and the intermediary nodes are external, unmonitored servers, the Lifeline Advisor directs the Lifeline Agents on this site to activate the workload’s SNA application resources that were previously inactivated. The Agent issues message AQS0214I for each SNA application resource that the Agent successfully activates. After the SNA application resource is activated, the SNA application that uses the resource might need to be restarted as some SNA applications shut down when the associated SNA application resource is inactivated.

Diagnosing problems when a SNA application is unavailable

When a SNA application is not reported as available, you can use the following information to diagnose the problem.

About this task

To diagnose SNA applications that are not available, complete the following steps:

Procedure

1. To determine whether the SNA application resource is active, issue the following command on the system where the SNA application should be active:
`D NET, ID=resource name`

See *z/OS Communications Server: SNA Operation* for more information about this command.

2. Determine whether you need to restart the SNA application. If the SNA application resource was inactivated by the Agent and later activated, the SNA application that uses the resource might have shut down when the associated SNA application resource was inactivated.
3. To determine whether the Lifeline Agent is able to use an extended MCS console to monitor SNA applications, verify that message AQS0009I SNA APPLICATIONS CANNOT BE MONITORED was not issued when the Agent was started. If the message was issued and syslogd was started, check the

syslogd file of the Agent for ERROR or WARNING messages and take the appropriate corrective actions. The syslogd identifier is lagent for the Agent. Otherwise, if the Agent internal trace is being used, dump and check the internal trace for ERROR or WARNING messages and take the appropriate corrective actions.

4. To determine whether messages were suppressed from the Agent's extended MCS console, verify that no message flooding warnings exist on the system console. If warnings exist, see "Preparing z/OS Agents for monitoring SNA applications" on page 24 for information about considerations to take when message flooding automation is enabled.
5. Determine whether the system where the SNA application is active has sufficient system resources. Issue the following command to determine whether the Agent views the SNA application as active:

```
MODIFY procname,DISPLAY,MEMBERS,DETAIL
```

 - If the SNA application is active, issue the following command to determine whether the Advisor views the SNA application as available:

```
MODIFY procname,DISPLAY,WORKLOAD,DETAIL
```
 - If the SNA application is not available, determine whether it is because the reported WLM weight for that system is zero.

Diagnosing problems when SNA intermediary node servers are unavailable

When the monitored intermediary node servers of a SNA workload are not reported as available, you can use the following information to diagnose the problem.

About this task

To diagnose SNA intermediary node servers that are not available, complete the following step:

Procedure

Determine whether the Lifeline Agent views the intermediary node servers as active. Issue the following command on the system where the intermediary node server should be active when the intermediary node resides on z/OS:

```
MODIFY procname,DISPLAY,MEMBERS,DETAIL
```

Issue the following command on the Linux on System z Management Agent when the intermediary node server resides on a Linux on System z guest:

```
aqsl1cmd DISPLAY,MEMBERS,DETAIL
```

If the intermediary node server is not available, use the netstat display on the Agent's system to verify that the intermediary node server is listening on the load balancer application group member's registered port and IP address (or the IPv6 unspecified address or IPv4 inaddr_any address).

- If a z/OS Agent monitors the intermediary node server, see *z/OS Communications Server: IP System Administrator's Commands* for information about using the netstat command.
- If a Linux on System z Management Agent monitors the intermediary node server, see the documentation of your Linux distribution for information about using the netstat command.

IBM MQ Cluster Workloads

Overview of MQ Cluster Workloads

Multi-site Workload Lifeline provides the ability to load balance requests for a workload that targets MQ cluster queues hosted by MQ queue managers on z/OS within an IBM MQ cluster across two sites at unlimited distances. Unlike other workloads defined to Lifeline, a first-tier load balancer is not required to perform the load balancing of workload requests. Lifeline leverages the capabilities of IBM MQ clusters to select which MQ queue managers are eligible to receive messages on behalf of a workload, based on which site the workload is active, and the health and availability of the MQ queue managers to handle additional workload requests.

For MQ cluster workloads, the MQ cluster queues that are defined for the workload are the group members that make up the workload. The IBM MQ cluster distributes messages destined for these MQ cluster queues based on recommendations provided by Lifeline.

Because a workload consists of one or more groups of applications, it is possible for a workload to consist of a group of TCP/IP z/OS applications, a group of TCP/IP Linux on System z applications, a group of MQ cluster queues, and a group of SNA applications that each access the same data sources on z/OS. Each group of applications must be defined in its own load balancer application group or MQ cluster queue group. Different types of applications cannot be defined in the same application/queue group.

MQ cluster workloads can be configured only in an Active/Standby configuration. When messages are targeted to a specific MQ cluster queue, the IBM MQ cluster distributes the messages only to the MQ queue managers in the active site.

When configuring the MQ cluster receiver channels for MQ queue managers that are participating in the IBM MQ cluster, the transport type (TRPTYPE) must be specified as TCP. Configure the connection name (CONNNAME) with either an IPv4 or IPv6 address. Although hostnames are supported by IBM MQ when defining the MQ cluster receiver channel, IBM Multi-site Workload Lifeline only supports IP addresses. For high availability, use a static VIPA as the configured IP address. The CONNNAME port is optional and will default to 1414 if not specified. If using multiple MQ queue managers on the same z/OS system, each MQ queue manager must be configured with a different port.

When configuring the MQ queue managers that are participating in a workload's IBM MQ cluster, the cluster receiver channel must not be automatically started when the MQ queue manager is started, and the rank of the MQ cluster queues for the workload must be set to '0'. This is to prevent MQ queue managers that are being started from erroneously receiving messages based on the current state of the workload. When a Lifeline Agent detects the newly started MQ queue manager, the Agent will ensure that the MQ cluster receiver channel and MQ cluster queues are set to the appropriate state and value, respectively.

By default, when an MQ application connects to a local MQ queue manager and sends messages that target an MQ cluster queue, these messages are routed only to the local MQ queue manager, if it hosts the MQ cluster queue. To ensure that these messages get routed to the correct MQ queue managers, based on the rank of the MQ cluster queue and state of each of the MQ queue manager's cluster receiver

channel, the use-queue attribute of the MQ cluster queue (CLWLUSEQ) must be set to ANY.

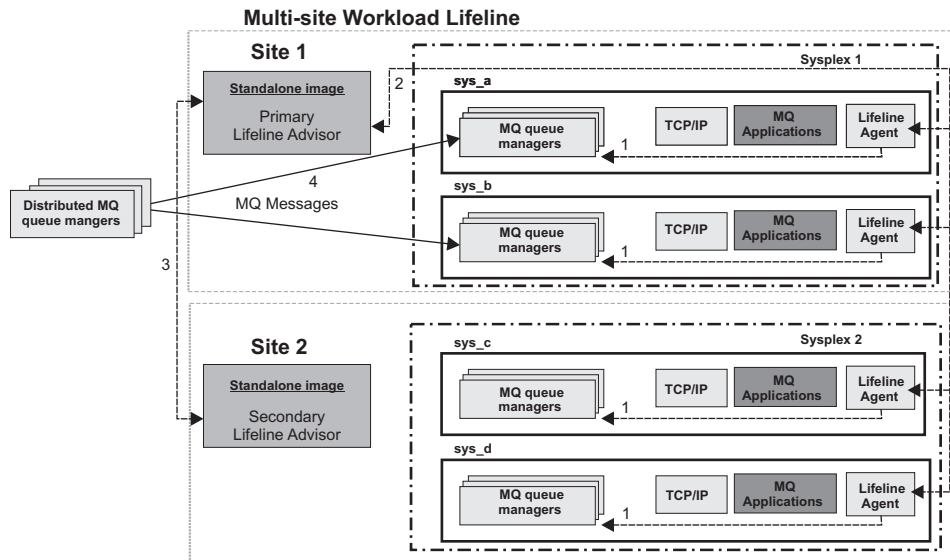


Figure 6. Relationships between MQ queue managers, Lifeline Advisors, and Lifeline Agents for MQ cluster workloads

Figure 6 shows an environment in which a workload is utilizing IBM MQ clusters. The Lifeline Advisors communicate as shown in Figure 1. All MQ queue managers that are depicted in the figure belong to the same MQ cluster. The Lifeline Agents running on each z/OS system in the two sites, Sysplex1 and Sysplex2, monitor the health of the MQ queue managers and the system where the Agent is active. The Lifeline Agents periodically send this information to the primary Lifeline Advisor, indicated by the network flows labeled 2. The primary Lifeline Advisor uses this information to calculate routing recommendations for the workload's messages. The Lifeline Advisor periodically sends recommendations to the Lifeline Agents about which MQ cluster queues are eligible to receive messages. The Lifeline Agent communicates with the MQ queue managers that are on the same system, to make MQ cluster queues eligible or ineligible for receiving messages targeted to that queue, indicated by the flows labeled 1.

A distributed MQ queue manager distributes messages to eligible MQ queue managers based on recommendations from the Advisor. In this example, Site 1 is the active site and messages are routed as shown in the network flows labeled 4. If multiple MQ queue managers are eligible to receive these messages, then the requests are distributed in a round-robin manner between the eligible MQ queue managers.

Monitoring MQ cluster queue groups

When an MQ cluster workload is configured to the Lifeline Advisor, an MQ cluster queue group is dynamically created by the Advisor for each MQ cluster queue defined for the workload. The Advisor sends the MQ cluster queue groups, each consisting of the list of MQ queue managers and the MQ cluster queue they are hosting, to the Lifeline Agents. The Agents gather and return information about the availability and health of the MQ queue managers and the MQ cluster queues to the Advisor.

See Figure 26 for an example of how to configure the Advisor `mq_manager_list` configuration statement for an MQ cluster workload.

Preparing z/OS Agents for monitoring MQ queue managers

Each Lifeline Agent communicates directly with each MQ queue manager on the same system as the Agent. By using this MQ connection, the Agent issues IBM MQ commands to monitor the status and alter the eligibility of the MQ cluster channels hosted on the MQ queue manager. See “(z/OS systems) Enabling the Lifeline Agents to communicate with MQ queue managers” on page 40 for steps on adding the correct IBM MQ libraries in the linklist concatenation.

Operational characteristics of MQ cluster workloads

Workload failure detection

When the Lifeline Advisor parses the entries in the `mq_manager_list`, the Advisor builds MQ cluster queue groups for each unique MQ cluster queue name it finds. The MQ queue managers and the MQ cluster queue of the MQ cluster queue group are monitored by the Lifeline Agents on the systems where the MQ queue managers reside.

An MQ cluster queue group for an MQ cluster workload is unavailable on a site when either of the following conditions occurs:

- All of the MQ queue managers of the MQ cluster queue group are unavailable on that site.
- The MQ cluster queue of the MQ cluster queue group is unavailable on all MQ queue managers on that site.

An MQ cluster workload failure occurs when one of the MQ cluster queue groups of the MQ cluster workload is continuously unavailable for the failure detection interval that is configured for that workload.

The MQ cluster workload recovers for a site when all of the MQ cluster queue groups for the MQ cluster workload on that site recover.

An MQ cluster queue group for an MQ cluster workload recovers when both of the following conditions occur:

- At least one MQ queue manager that hosts the cluster queue becomes available on that site.
- At least one MQ cluster queue instance becomes available on an MQ queue manager on that site.

Activating an MQ cluster workload

When an MQ cluster workload is activated, the Lifeline Advisor queries the Lifeline Agents on the alternate site to determine whether queued messages exist on MQ cluster queues and whether MQ cluster receiver channels are stopped in either of the following situations:

- The MQ cluster workload is activated for the first time.
- The MQ cluster workload is not being activated for the first time and the workload is not in a deactivated state.

If queued messages are found on the MQ cluster queues or not all MQ cluster receiver channels are stopped, the ACTIVATE command is not successful.

When an MQ cluster workload is activated, the Lifeline Advisor directs the Lifeline Agents on the target site to change the queue rank of each MQ cluster queue in the workload to a '1'. The Agents also start each MQ cluster's receiver channel on each MQ queue manager that hosts the workload's MQ cluster.

For more information, see “Activating a workload for the first time” on page 51 and “Activating a workload when the workload is not deactivated” on page 51.

Quiescing an MQ cluster workload

When an MQ cluster workload is quiesced, the Lifeline Advisor directs the Lifeline Agents on the target site to change the queue rank of each MQ cluster queue in the workload to a '0'. The Agents also stop each MQ cluster's receiver channel on each MQ queue manager that hosts the workload's MQ cluster.

Deactivating an MQ cluster workload

When an MQ cluster workload is deactivated, the Lifeline Advisor queries the Lifeline Agents on the previously active site to determine whether MQ cluster receiver channels are stopped and whether queued messages exist on MQ cluster queues hosted by MQ queue managers on the site.

Whether or not the MQ cluster receiver channels are stopped or queued messages are found on the MQ cluster queues, the MQ cluster workload is transitioned into a deactivated state. However, if either condition is true, the DEACTIVATE command is not successful.

Refer to the System Programmer Response for message AQS0207I to determine the possible actions to take if the DEACTIVATE command is not successful.

Reactivating an MQ cluster workload

When an MQ cluster workload that was deactivated is activated to the same site, the Lifeline Advisor directs the Lifeline Agents on this site to change the queue rank of each MQ cluster queue in the workload to a '1'. The Agents also start each MQ cluster receiver channel on each MQ queue manager that hosts the workload's MQ cluster.

Diagnosing problems when an MQ queue manager is unavailable

When an MQ queue manager is not reported as available, you can use the following information to diagnose the problem.

About this task

To diagnose MQ queue managers that are not available, complete the following steps:

Procedure

1. To determine whether the Lifeline Agent views the MQ queue manager as active, issue the following command on the system where the MQ queue manager should be active:

```

|
| MODIFY procname,DISPLAY,MEMBERS,DETAIL
|
| 2. To determine whether the Lifeline Agent can communicate with MQ queue
| managers, verify that message AQS0010I MQ QUEUE MANAGERS CANNOT
| BE MONITORED was not issued when the Agent was started. If the message
| was issued and syslogd was started, check the syslogd file of the Agent for
| ERROR or WARNING messages and take the appropriate corrective actions.
| The syslogd identifier is llagent for the Agent. Otherwise, if the Agent internal
| trace is being used, dump and check the internal trace for ERROR or
| WARNING messages and take the appropriate corrective actions.
|
| 3. To determine whether all MQ cluster queues that were configured for the MQ
| queue manager are defined and enabled to receive messages, issue the
| following Agent command:
|
| MODIFY procname,DISPLAY,MEMBERS,DETAIL
|
| 4. To determine whether the system where the MQ queue manager is active has
| sufficient system resources, issue the following command to determine whether
| the Agent views the MQ queue manager as active:
|
| MODIFY procname,DISPLAY,MEMBERS,DETAIL
|
| • If the MQ queue manager is active, issue the following command to
| determine whether the Advisor views the MQ queue manager as available:
|
| MODIFY procname,DISPLAY,WORKLOAD,DETAIL
|
| • If the MQ queue manager is not available, determine whether it is because
| the reported WLM weight for that system is zero.
|
|

```

Chapter 2. Preparing to use Multi-site Workload Lifeline

You need to evaluate TCP/IP workloads and grant users authority before you can configure Multi-site Workload Lifeline to balance workloads between the two z/OS sites.

Before you begin

You must meet the following requirements:

- All z/OS systems must be at the V1R12 level or higher levels.
- All Linux on System z guest systems must be at the SUSE Linux Enterprise Server (SLES) 10 SP4 or Red Hat Enterprise Linux (RHEL) 5 version 5.8 level or higher levels.
- For workloads consisting of server applications, ensure both of the following statements are true:
 - You must have at least one first-tier external load balancer that supports the Server/Application State Protocol (SASP). The load balancer must have IP connectivity to each second-tier load balancer. The second-tier load balancer is another external load balancer that supports SASP, an intermediary node, and/or a z/OS sysplex distributor that runs on a target system in the site.
 - The applications that are identified to be load balanced must be able to run on target systems that are in different sites, and must be able to access data sources that are synchronized across the sites.
- For workloads utilizing IBM MQ clusters, ensure both of the following statements are true:
 - All MQ queue managers must be at WebSphere® MQ 7.0.1 or higher levels.
 - The MQ queue managers that host the MQ cluster queues for the workload must be configured in an IBM MQ cluster that spans the sites, and must have access to data sources that are synchronized across the sites.

(z/OS systems) Granting users explicit authority to control the Lifeline Advisor and Agent

Rule: This task is only applicable to Lifeline Advisors and the Lifeline Agents residing on z/OS systems.

Grant explicit authority to all users that can start, stop, or modify the Lifeline Advisor and Agent, to prevent unauthorized users from controlling it. If you do not grant explicit authority, any user that is able to issue the START, STOP, or MODIFY command can control the Lifeline Advisor and Agent.

About this task

Restriction: The only commands that you can issue against the Lifeline Advisor and Agent are MODIFY commands, with the exception of the START command that you use to start the Lifeline Advisor or Agent or the STOP command that you use to stop the Lifeline Advisor or Agent. Limit which users are able to issue the START, STOP, and MODIFY commands for these two servers.

To grant users explicit authority to start, stop, or modify the Lifeline Advisor and Agent, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the OPERCMDS class is active and that RACLIST processing is enabled:
SETROPTS CLASSACT(OPERCMDS)
SETROPTS RACLIST (OPERCMDS)
2. Define the following OPERCMDS class profile using a security product like RACF® for the Lifeline Advisor:
RDEFINE OPERCMDS (MVS.SERVGR.LLADVSR) UACC(NONE)
3. Define the following OPERCMDS class profile using a security product like RACF for the Lifeline Agent:
RDEFINE OPERCMDS (MVS.SERVGR.LLAGENT) UACC(NONE)
4. Permit the Lifeline Advisor to the OPERCMDS class profile:
PERMIT MVS.SERVGR.LLADVSR CLASS(OPERCMDS) ACCESS(CONTROL) ID(userid)
5. Permit the Lifeline Agent to the OPERCMDS class profile:
PERMIT MVS.SERVGR.LLAGENT CLASS(OPERCMDS) ACCESS(CONTROL) ID(userid)
6. Issue the following command to refresh the OPERCMDS class profile:
SETROPTS RACLIST(OPERCMDS) REFRESH

Authorizing the Lifeline Advisors to use BCPii

You need to grant Lifeline Advisors general authority to use the Base Control Program Internal Interface (BCPii), including authority to specific resources such as the interconnected Central Processor Complexes (CPCs) and images on those CPCs.

About this task

To allow a Lifeline Advisor to query CPC and image information, complete the following steps. See “Using the Base Control Program Internal Interface” on page 54 for more information.

Procedure

1. Issue the following two commands to ensure that the FACILITY class is active and that RACLIST processing is enabled:
SETROPTS CLASSACT(FACILITY)
SETROPTS RACLIST (FACILITY)
 2. Define the following FACILITY class profile for authority to use BCPii:
RDEFINE FACILITY HWI.APPLNAME.HWISERV UACC(NONE)
 3. Permit the Lifeline Advisor to the FACILITY class:
PERMIT HWI.APPLNAME.HWISERV CLASS(FACILITY) ACCESS(READ) ID(userid)
 4. Define the following FACILITY class profile (one for each CPC) for authority to query the CPC.
RDEFINE FACILITY HWI.TARGET.*netid.nau* UACC(NONE) APPLDATA(*community-name*)
- Where:
- *netid.nau* represents the 3 to 17 character SNA name of the particular CPC.
 - *community-name* is the SNMP community name that was defined on the Support Element for that CPC.
5. Permit the Lifeline Advisor to each of the FACILITY class profiles:
PERMIT HWI.TARGET.*netid.nau* CLASS(FACILITY) ACCESS(READ) ID(userid)

6. Define the following FACILITY class profile (one for each CPC) for authority to query each image of the CPC.

```
RDEFINE FACILITY HWI.TARGET.netid.nau.* UACC(NONE)
```

Where:

- *netid.nau* represents the 3 to 17 character SNA name of the particular CPC.
7. Permit the Lifeline Advisor to each of these FACILITY class profiles:

```
PERMIT HWI.TARGET.netid.nau.* CLASS(FACILITY) ACCESS(READ) ID(userid)
```
 8. Issue the following command to refresh the FACILITY class profile:

```
SETROPTS RACLIST(FACILITY) REFRESH
```

Authorizing the Lifeline Advisor to accept secure peer Advisor connections

If you want to authenticate a peer Lifeline Advisor using Application Transparent TLS (AT-TLS), you need to grant the user ID associated with the peer Advisor access to the AQS.LLADVSR.ADVISORACCESS resource.

About this task

To authenticate a peer Lifeline Advisor by using AT-TLS, grant the peer Advisor access to the AQS.LLADVSR.ADVISORACCESS resource. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.

To allow a Lifeline Advisor to authenticate a peer Advisor by using AT-TLS, rather than the `advisor_id_list` Advisor configuration statement, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the SERVAUTH class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(SERVAUTH)  
SETROPTS RACLIST (SERVAUTH)
```

2. Define the following SERVAUTH class profile:

```
RDEFINE SERVAUTH(AQS.LLADVSR.ADVISORACCESS.sysname) UACC(NONE)
```

Where:

- *sysname* is the MVS system name where the local Advisor resides.
3. Grant the user ID associated with the peer Lifeline Advisor READ access to the SERVAUTH class profile:

```
PERMIT AQS.LLADVSR.ADVISORACCESS.sysname CLASS(SERVAUTH) ACCESS(READ) ID(userid)
```
 4. Issue the following command to refresh the SERVAUTH class profile:

```
SETROPTS RACLIST(SERVAUTH) REFRESH
```

See Appendix D, “Sample AT-TLS configuration policies,” on page 275 for examples of Lifeline Advisor and peer Lifeline Advisor AT-TLS configurations for a peer Lifeline Advisor connection.

(z/OS systems) Authorizing the Lifeline Advisor to accept secure Agent connections

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. Linux on System z Management Agents do not support authenticating to Advisors using AT-TLS.

If you want to authenticate a Lifeline Agent using Application Transparent TLS (AT-TLS), you need to grant the user ID associated with the Agent access to the AQS.LLADVSR.AGENTACCESS resource.

About this task

To authenticate a Lifeline Agent by using AT-TLS, grant the Agent access to the AQS.LLADVSR.AGENTACCESS resource. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.

Note: You must configure each Agent IP addresses in the `agent_id_list` Advisor configuration statement, regardless of whether secure Agent connections are used.

To allow a Lifeline Advisor to authenticate an Agent by using AT-TLS, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the SERVAUTH class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(SERVAUTH)
SETROPTS RACLIST (SERVAUTH)
```

2. Define the following SERVAUTH class profile:

```
RDEFINE SERVAUTH (AQS.LLADVSR.AGENTACCESS.sysname) UACC(NONE)
```

Where:

- *sysname* is the MVS system name where the Advisor resides.

3. Grant the user ID associated with the Lifeline Agent READ access to the SERVAUTH class profile:

```
PERMIT AQS.LLADVSR.AGENTACCESS.sysname CLASS(SERVAUTH) ACCESS(READ) ID(userid)
```

4. Issue the following command to refresh the SERVAUTH class profile:

```
SETROPTS RACLIST(SERVAUTH) REFRESH
```

See Appendix D, “Sample AT-TLS configuration policies,” on page 275 for examples of Lifeline Advisor and peer Lifeline Advisor AT-TLS configurations for a peer Lifeline Advisor connection.

Authorizing the Lifeline Advisor to accept secure Load Balancer connections

If you want to authenticate a load balancer using Application Transparent TLS (AT-TLS), you need to grant the user ID associated with the load balancer access to the AQS.LLADVSR.LBACCESS resource.

About this task

To authenticate a load balancer by using AT-TLS, grant the load balancer access to the AQS.LLADVSR.LBACCESS resource. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.

Note: You must configure each load balancer IP addresses in the `lb_id_list` Advisor configuration statement, regardless of whether secure load balancer connections are used.

To allow a Lifeline Advisor to authenticate a load balancer by using AT-TLS, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the SERVAUTH class is active and that RACLIST processing is enabled:

```
SETOPTS CLASSACT(SERVAUTH)
SETOPTS RACLIST (SERVAUTH)
```

2. Define the following SERVAUTH class profile:

```
RDEFINE SERVAUTH (AQS.LLADVSR.LBACCESS.sysname) UACC(NONE)
```

Where:

- *sysname* is the MVS system name where the Advisor resides.

3. Grant the user ID associated with the load balancer READ access to the SERVAUTH class profile:

```
PERMIT AQS.LLADVSR.LBACCESS.sysname CLASS(SERVAUTH) ACCESS(READ) ID(userid)
```

4. Issue the following command to refresh the SERVAUTH class profile:

```
SETOPTS RACLIST(SERVAUTH) REFRESH
```

See Appendix D, "Sample AT-TLS configuration policies," on page 275 for examples of Lifeline Advisor and peer Lifeline Advisor AT-TLS configurations for a peer Lifeline Advisor connection.

(z/OS systems) Authorizing the Lifeline Advisors and Agents to issue event-driven messages

Rule: This task is only applicable to Lifeline Advisors and the Lifeline Agents residing on z/OS systems.

You need to grant Lifeline Advisors and Agents access to the BPX.CONSOLE resource only if the *userid* that is associated with the OMVS segment created for the Lifeline Advisor and Agent is defined with a non-zero UID. Otherwise, any event-driven message issued by the Lifeline Advisor or Agent will be prefixed with BPXM023I. See "Configuring one Lifeline Advisor as the primary Advisor, and optionally one as the secondary Advisor" on page 44 and "Configuring a Lifeline Agent" on page 46 for more information about OMVS segments.

About this task

If you are using RACF and have already defined the resource profile and enabled the FACILITY class, grant Lifeline Advisors and Agents access to that resource.

If you are using a security product other than RACF that by default denies access to the resource, grant the Lifeline Advisors and Agents access to the resource.

If you are using RACF, complete the following steps to define the resource profile if you have not defined it, and grant Lifeline Advisors and Agents access to the resource:

Procedure

1. Issue the following two commands to ensure that the FACILITY class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(FACILITY)  
SETROPTS RACLIST (FACILITY)
```

2. Define the following FACILITY class profile:

```
RDEFINE FACILITY (BPX.CONSOLE) UACC(NONE)
```

3. Grant the Lifeline Advisor READ access to the FACILITY class profile:

```
PERMIT BPX.CONSOLE CLASS(FACILITY) ACCESS(READ) ID(userid)
```

4. Grant the Lifeline Agent READ access to the FACILITY class profile:

```
PERMIT BPX.CONSOLE CLASS(FACILITY) ACCESS(READ) ID(userid)
```

5. Issue the following command to refresh the FACILITY class profile:

```
SETROPTS RACLIST(FACILITY) REFRESH
```

(z/OS systems) Authorizing the Lifeline Agents to use WLM services

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. You need to grant Lifeline Agents access to the BPX.WLMSEVER resource only if you have already defined the resource profile.

About this task

If you are using RACF and have already defined the resource profile and enabled the FACILITY class, grant Lifeline Agents access to that resource.

If you are using a security product other than RACF that by default denies access to the resource, grant the Lifeline Agents access to the resource.

If you are using RACF, complete the following steps to define the resource profile if you have not defined it, and grant Lifeline Agents access to the resource:

Procedure

1. Issue the following two commands to ensure that the FACILITY class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(FACILITY)  
SETROPTS RACLIST (FACILITY)
```

2. Define the following FACILITY class profile:

```
RDEFINE FACILITY (BPX.WLMSEVER) UACC(NONE)
```

3. Grant the Lifeline Agent READ access to the FACILITY class profile:

```
PERMIT BPX.WLMSEVER CLASS(FACILITY) ACCESS(READ) ID(userid)
```

4. Issue the following command to refresh the FACILITY class profile:

```
SETROPTS RACLIST(FACILITY) REFRESH
```

(z/OS systems) Authorizing the Lifeline Agents to reset active connections

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. You need to grant Lifeline Agents access to the MVS.VARY.TCPIP.DROP resource.

About this task

To process a Lifeline Advisor DEACTIVATE command, the Lifeline Agents must be able to reset active connections that have been identified by the Advisor. See “MODIFY *advisor_proc*,DEACTivate command: Deactivate workloads” on page 143 for more information.

To grant Agents explicit authority to reset active connections, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the OPERCMDS class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(OPERCMDS)
SETROPTS RACLIST (OPERCMDS)
```

2. Define the following OPERCMDS class profile:

```
RDEFINE OPERCMDS (MVS.VARY.TCPIP.DROP) UACC(NONE)
```

3. Permit the Lifeline Agent to the OPERCMDS class profile:

```
PERMIT MVS.VARY.TCPIP.DROP CLASS(OPERCMDS) ACCESS(CONTROL) ID(userid)
```

4. Issue the following command to refresh the OPERCMDS class profile:

```
SETROPTS RACLIST(OPERCMDS) REFRESH
```

(z/OS systems) Authorizing the Lifeline Agents to monitor SNA applications

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. To prevent other applications from accessing the Lifeline Agent extended MCS console, define the MVS.MCSOPER.LLAGNT* generic resource and grant Agents access to this resource.

About this task

To monitor SNA applications, a Lifeline Agent uses an extended MCS console to issue VTAM commands. The console name for this extended MCS console is the text string 'LLAGNT' that is appended with the MVS sysclone value of the system where the Agent resides.

To restrict access to the console name that the Lifeline Agent uses, complete the following steps:

Procedure

1. Issue the following two commands to ensure that the OPERCMDS class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(OPERCMDS)
SETROPTS RACLIST (OPERCMDS)
```

2. Define the following OPERCMDS class profile:

```
RDEFINE OPERCMDS (MVS.MCSOPER.LLAGNT*) UACC(NONE)
```

3. Permit the Lifeline Agent to the OPERCMDS class profile:

```
PERMIT MVS.MCSOPER.LLAGNT* CLASS(OPERCMDS) ACCESS(READ) ID(userid)
```

4. Issue the following command to refresh the OPERCMDS class profile:

```
SETROPTS RACLIST(OPERCMDS) REFRESH
```

(z/OS systems) Enabling the Lifeline Agents to communicate with MQ queue managers

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. To allow a Lifeline Agent to communicate with MQ queue managers, either the IBM MQ SCSQLOAD library must be added to the LNKLST concatenation or added as a STEPLIB DD statement in the Agent's started procedure.

About this task

To monitor MQ queue managers and the MQ cluster queues that they host, the Lifeline Agent creates an MQ connection with each of the MQ queue managers on the system where the Agent resides.

To allow a Lifeline Agent to communicate with MQ queue managers, complete one of the following steps:

Procedure

1. Add the IBM MQ SCSQLOAD dataset name to the LNKLSTxx SYS1.PARMLIB member. See *z/OS Initialization and Tuning Reference* for information on the LNKLSTxx member.
2. Add the IBM MQ SCSQLOAD dataset name as a STEPLIB DD statement in the Agent's started procedure. See *z/OS MVS JCL Reference* for information on the STEPLIB statement.

Granting a network management application authority to use the Network Management Interface

Multi-site Workload Lifeline provides a Network Management Interface (NMI) to allow an application to retrieve workload states, server application metrics, and connection status from the Lifeline Advisor.

About this task

To use this NMI, you must allow the *userid* associated with the application to either have superuser authority or grant access to the AQS.NETMGMT.DISPLAY resource.

To allow the application to use the NMI without having superuser authority, complete the following steps to grant access to the AQS.NETMGMT.DISPLAY resource:

Procedure

1. Issue the following two commands to ensure that the SERVAUTH class is active and that RACLIST processing is enabled:

```
SETROPTS CLASSACT(SERVAUTH)
SETROPTS RACLIST (SERVAUTH)
```
2. Define the following SERVAUTH class profile:

```
RDEFINE SERVAUTH AQS.NETMGMT.DISPLAY UACC(NONE)
```
3. Permit the user ID that is associated with the application READ access to the SERVAUTH class profile:

```
PERMIT AQS.NETMGMT.DISPLAY ACCESS(READ) CLASS(SERVAUTH) ID(userid)
```
4. Issue the following command to refresh the SERVAUTH class profile:

SETROPTS RACLIST(SERVAUTH) REFRESH

Chapter 3. Configuring Multi-site Workload Lifeline

You need to configure Multi-site Workload Lifeline before you can use it to balance TCP/IP workloads among different z/OS sites.

Before you begin

Complete the tasks in Chapter 2, “Preparing to use Multi-site Workload Lifeline,” on page 33.

Lifeline Advisor configuration file

You specify the Lifeline Advisor configuration file on the CONFIG DD statement in the Lifeline Advisor started procedure to configure both the primary and secondary Lifeline Advisors.

The Lifeline Advisor configuration file serves the following basic purposes:

- Defines the listening sockets for the external load balancers, the secondary Lifeline Advisor, and Lifeline Agents
- Provides an access control list for specifying which external load balancers, Lifeline Advisor, and Lifeline Agents can connect to the primary Lifeline Advisor
- Associates each second-tier load balancer to a site for all the defined workloads
- Customizes optional parameters

The Lifeline Advisor reads configuration data from one file, which can be a z/OS UNIX file, a PDS or PDSE member, or a sequential data set.

The configuration file used by the Lifeline Advisor, when the configuration file is a z/OS UNIX file, must meet these requirements:

- The write permission of the configuration file can only be set for the user and group classes.
- The others class can only be granted read permission.

See Appendix A, “A sample started procedure for Lifeline Advisors,” on page 253 for a sample Advisor start procedure.

Lifeline Agent configuration file

For z/OS Agents, specify the Lifeline Agent configuration file on the CONFIG DD statement in the Lifeline Agent start procedure. For Linux on System z Management Agents, specify the Lifeline Agent configuration file by using the **-c configfile** parameter, or use the default location `/opt/ibm/lifeline/conf/agent.conf`. See “Linux on System z Workloads” on page 13 for more information about the Linux on System z Management Agent configuration file.

The Lifeline Agent configuration file serves three basic purposes:

- Defines the IP address and port that the Lifeline Agent binds to for communication with the Lifeline Advisor
- Identifies the IP address and port of the Lifeline Advisor
- Customizes optional parameters

For z/OS Agents, the Lifeline Agent configuration file can be a z/OS UNIX file, a PDS or PDSE member, or a sequential data set.

The configuration file used by the Linux on System z Management Agent or the z/OS Agent, when the configuration file is a z/OS UNIX file, must meet these requirements:

- The write permission of the configuration file can only be set for the user and group classes.
- The others class can only be granted read permission.

See Appendix B, “A sample started procedure for Lifeline z/OS Agents,” on page 255 for a sample Agent start procedure.

Configuring one Lifeline Advisor as the primary Advisor, and optionally one as the secondary Advisor

Use the Lifeline Advisor configuration file to configure one Lifeline Advisor as the primary Advisor. You can also use this file to configure an optional secondary Lifeline Advisor.

About this task

Only one Lifeline Advisor can serve as a primary Advisor at any given time. To start a secondary Lifeline Advisor to take over primary Advisor responsibilities in the case of a failure of the primary Advisor, specify identical statements in the Lifeline Advisor configuration file for both the primary and secondary Lifeline Advisors.

Complete the following steps to configure Lifeline Advisors:

Procedure

1. Define listening sockets, ports, or both.
 - a. Specify the local IPv4 address and port that the Lifeline Advisor listens on for IPv4 load balancer connections on the `lb_connection_v4` configuration statement. The default port for communications with external load balancers is 3860. You can also specify the local IPv6 address and port that the Lifeline Advisor listens on for IPv6 load balancer connections on the `lb_connection_v6` configuration statement.

Guideline: To enable the movement of the primary Lifeline Advisor to another system when the Lifeline Advisor or its underlying system fails, specify a dynamic VIPA (DVIPA) on the `lb_connection_v4` and `lb_connection_v6` statements. Furthermore, make this DVIPA a unique application-instance DVIPA (defined through `VIPARANGE`) rather than a multiple application-instance DVIPA (defined through `VIPADefine`). For more information about configuring DVIPAs, see *z/OS Communications Server: IP Configuration Reference*.

Rules:

- Run the Lifeline Advisor on an INET system. If you run the Lifeline Advisor on a CINET system, the address or addresses that you specify on the `lb_connection_v4` or `lb_connection_v6` statement tie the Lifeline Advisor to the stack owning those addresses. Consequently, the termination of that stack results in the termination of the Lifeline Advisor.

- If the primary Lifeline Advisor is using IPv6 for the external load balancer connections, or if any Lifeline Agents are using IPv6 to connect to the Lifeline Advisor, you must configure the system in which the secondary Lifeline Advisor resides to be enabled for IPv6.
- b. Specify the local port that the Lifeline Advisor listens on for Lifeline Agent connections on the `agent_connection_port` statement. If the TCP/IP stack of the Lifeline Advisor is enabled for IPv6, the Lifeline Advisor opens a listening socket for Lifeline Agents on the IPv6 unspecified address (::) on the port specified by this statement. This enables Lifeline Agents to connect to the Lifeline Advisor using either the IPv4 address specified on the `lb_connection_v4` statement or the IPv6 address specified on the `lb_connection_v6` statement. If the TCP/IP stack is not enabled for IPv6, the Lifeline Advisor opens a listening socket on the IPv4 unspecified address, 0.0.0.0. This enables Lifeline Agents to connect to the Lifeline Advisor using only the IPv4 address specified on the `lb_connection_v4` statement.
 - c. Specify the local port that the Lifeline Advisor listens on for a connection from the secondary Lifeline Advisor on the `advisor_connection_port` statement. The same information specified in the previous step about the use of the `lb_connection_v4` and `lb_connection_v6` statements applies in this step.
2. Define the access control list.

Rule: Specify only complete IP addresses in access control lists. Subnetworks, IP prefixes, or other types of wildcards are not allowed.

- a. Specify the list of external load balancers that are allowed to connect to the Lifeline Advisor on the `lb_id_list` statement. A maximum limit of 100 external load balancers can be connected to a Lifeline Advisor at any given time.

You can also use Application Transparent TLS (AT-TLS) to authenticate external load balancers. If you use AT-TLS for a load balancer, you still need to include the IP address of that load balancer in the `lb_id_list` statement. See “Authorizing the Lifeline Advisor to accept secure Load Balancer connections” on page 36 for more information.
 - b. Specify the list of secondary Lifeline Advisors that are allowed to connect to the primary Lifeline Advisor on the `advisor_id_list` statement. Only a single secondary Lifeline Advisor can be connected to the primary Lifeline Advisor at any given time.

You can also use Application Transparent TLS (AT-TLS) to authenticate secondary Lifeline Advisors. If you use AT-TLS for a secondary Advisor, you do not need to include the IP address of that Advisor in the `advisor_id_list` statement. See Authorizing the Lifeline Advisor to accept secure peer Advisor connections for more information.
 - c. Specify the list of Lifeline Agents that are allowed to connect to the primary Lifeline Advisor on the `agent_id_list` statement.

You can also use Application Transparent TLS (AT-TLS) to authenticate Lifeline z/OS Agents. If you use AT-TLS for a z/OS Agent, you still need to include the IP address of that z/OS Agent in the `agent_id_list` statement. See Authorizing the Lifeline Advisor to accept secure Agent connections for more information.
3. Update the TCP/IP profiles on the systems where the primary and secondary Advisors reside to add the `VIPARANGE` definition that configures the IP address used in the `lb_connection_v4` and `lb_connection_v6` Advisor configuration statements.

4. Define an OMVS segment for the Lifeline Advisor. The Lifeline Advisor accesses files in the UNIX System Services shell, so the Lifeline Advisor needs to have an OMVS segment defined for it. You can use the following JCL snippet to define the OMVS segment:

```
//EZARACF JOB MSGLEVEL=(1,1),MSGCLASS=A,CLASS=A
//DAEMONS EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
ADDUSER userid DFLTGRP(OMVSGRP) OMVS(UID(0) HOME('/'))
RDEFINE STARTED userid.* STDATA(USER(userid))
SETROPTS RACLIST(STARTED) REFRESH
//*
```

5. Optional: Customize optional statements, including the `update_interval`, `failure_detection_interval`, and `debug_level` statements.

Guideline: In most cases, do not customize the `debug_level` statement, unless you are directed to do so by an IBM service representative. Adding additional types of trace data can cause the amount of captured data to become voluminous. Reducing the amount of trace data from the default might make diagnosing a problem more difficult.

6. Optional: Change the location of the Multi-site Workload Lifeline Advisor message catalog. The Lifeline Advisor uses a standard message catalog. The message catalog must be in the z/OS UNIX file system. The default location of the message catalog is `/usr/lib/nls/msg/C/aqsl1msg.cat`. If you want to store the message catalog in another location, you need to change the `NLSPATH` or the `LANG` environment variables. To change the two environment variables, you can update the started procedure of the Lifeline Advisor to pass the `ENVAR` option to set `NLSPATH`. If `aqsl1msg.cat` does not exist, the Lifeline Advisor by default uses the messages hardcoded in the software. These messages duplicate the English message catalog that is shipped with the Multi-site Workload Lifeline product.

Configuring a Lifeline Agent

Use the Lifeline Agent configuration file to configure one Lifeline Agent per z/OS target system across both sites. If workloads contain Linux on System z applications, you also need to use the Lifeline Agent configuration file to configure one Linux on System z Management Agent per target z/VM system across both sites.

About this task

Only one Lifeline Agent can be active per z/OS system at any time. Only one Linux on System z Management Agent per site can be active per z/VM system at any time. If two Linux on System z Management Agents are configured on the same z/VM system, each Agent monitors a different set of Linux on System z guests for the site they are configured. Complete the following steps to configure Lifeline Agents:

Procedure

1. Specify the local IP address and port that the Lifeline Agent binds to for communicating with the Lifeline Advisor on the `host_connection` statement. This step is used as part of the Lifeline Advisor's access control enforcement.

Guideline: For high availability of z/OS Agents, use a static VIPA as the local IP address to tolerate individual link outages on the z/OS system. For more information about configuring static VIPAs, see *z/OS Communications Server: IP Configuration Reference*.

2. Specify the IP address and port of the Lifeline Advisor on the `advisor_id` statement.
3. For Lifeline z/OS Agents, update the TCP/IP profiles on the systems where the Agents reside to add the static VIPA definition that configures the IP address used in the `host_connection` Agent configuration statement.

For Lifeline z/OS Agents, you can also use Application Transparent TLS (AT-TLS) to authenticate the Agents. See *Authorizing the Lifeline Advisor to accept secure Agent connections* for more information.

4. For Lifeline z/OS Agents, define an OMVS segment for the Lifeline Agent. The z/OS Agent accesses files in the UNIX System Services shell, so you need to define an OMVS segment for the Lifeline Agent. You can use the following JCL snippet to define the OMVS segment:

```
//EZARACF JOB MSGLEVEL=(1,1),MSGCLASS=A,CLASS=A
//DAEMONS EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
ADDUSER userid DFLTGRP(OMVSGRP) OMVS(UID(0) HOME('/'))
RDEFINE STARTED userid.* STDATA(USER(userid))
SETROPTS RACLIST(STARTED) REFRESH
//*
```

5. Optional: Use the `debug_level` statement to specify how much trace data is captured in the log file of the Lifeline Agent.

Rule: In most cases, you should not customize this statement, unless you are directed to do so by an IBM service representative. Adding additional types of trace data can cause the amount of data captured to become voluminous. Reducing the amount of trace data from the default might make diagnosing a problem more difficult.

6. Optional: For Lifeline z/OS Agents, change the location of the Multi-site Workload Lifeline Agent message catalog. The z/OS Agent uses a standard message catalog. The message catalog must be in the z/OS UNIX file system. The default location of the message catalog is `/usr/lib/nls/msg/C/aqsl1msg.cat`. If you want to store the message catalog in another location, you need to change the `NLSPATH` or the `LANG` environment variables. To change the two environment variables, you can update the started procedure of the z/OS Agent to pass the `ENVAR` option to set `NLSPATH`. If `aqsl1msg.cat` does not exist, the Lifeline z/OS Agent by default uses the messages hardcoded in the software. These messages duplicate the English message catalog that is shipped with the Multi-site Workload Lifeline product.

For Linux on System z Management Agents, the default location of the Multi-site Workload Lifeline Agent message catalog is `/opt/ibm/lifeline/lib/msg/C/aqsl1msg.cat`. To store the message catalog in another location, you must set the `NLSPATH` environment variable. If `aqsl1msg.cat` does not exist or cannot be found, the Linux on System z Agent by default uses the messages that are hardcoded in the software. These messages duplicate the English message catalog that is shipped with the Multi-site Workload Lifeline product.

(z/OS systems) WLM recommendations

Rule: This section is only applicable to the Lifeline Agents residing on z/OS systems.

The Lifeline z/OS Agents use a WLM recommendation that is referred to as

server-specific WLM. Server-specific WLM provides a relative measure of the ability of a target system to handle additional workload requests, as compared to other systems across both sites. It also indicates how well individual server applications or MQ queue managers are doing compared to the WLM policy goals that have been specified for that workload.

Server-specific WLM recommendations are composed of two key elements:

- The amount of displaceable capacity (general, zAAP, and zIIP) available on the target system, based on the importance level of the server application or MQ queue manager, and the proportion of general, zAAP, and zIIP CPU that is currently being consumed by the workload of the application. For example, if the application is using only general and zAAP CPU, the displaceable zIIP capacity is not considered.
- The performance of the application compared to the WLM goals for that application workload.

WLM provides an interface that enables server applications to report the following additional information:

- The abnormal transaction completion rate, or the rate of abnormal completions per 1000 total transactions
- Application health, which is a value in the range 0 - 100% (100% is optimal), representing the overall health of the server application

Configuring Lifeline Advisors to automatically restart

You can use automation software or define an automatic restart manager (ARM) policy to automatically restart the Lifeline Advisor in case the Advisor or system fails. Although this task is optional, you should restart the Lifeline Advisor as soon as possible after it fails.

About this task

The restart capability covers scenarios where the Lifeline Advisor itself fails; in scenarios where the system the Lifeline Advisor is running on fails, you can configure a secondary Lifeline Advisor to automatically take over primary Advisor responsibilities.

If the Lifeline Advisor or its underlying system fails, and no secondary Lifeline Advisor is configured to automatically take over primary Advisor responsibilities, an external load balancer might do one of the following things:

- Continue to distribute workload requests according to the last set of information received from the Lifeline Advisor
- Resort to pre-configured weights
- Stop distributing new workload requests to second-tier load balancers or server applications

The actual behavior depends on the load balancer implementation. Consult the load balancer documentation for your specific external load balancer for details.

You need to restart the Lifeline Advisor as soon as possible when the Lifeline Advisor itself fails, or configure a secondary Lifeline Advisor to automatically take over primary Advisor responsibilities when the system on which the Lifeline Advisor is running fails.

For more information about defining ARM policies, see *z/OS MVS Setting Up a Sysplex*.

The Lifeline Advisor registers with ARM using the following values:

- ELEMTYPE=SYSLVL1
- ELEMNAME=AQSSLADVSR
- TERMTYPE=ELEMTERM

These values indicate that if the Lifeline Advisor fails on this system, the Lifeline Advisor needs to be restarted only on this system.

To configure Lifeline Advisors to automatically restart using ARM, complete the following step:

Procedure

Update the IXCARM resource. The Lifeline Advisors do not run using a system key. If you are using ARM registration, the started task IDs must be permitted with UPDATE authority to the IXCARM.SYSLVL1.AQSSLADVSR resource profile in the FACILITY class using the SAF product on your system. To enable the Advisor to register with ARM, use the following RACF commands to define the profiles and grant update access:

```
RDEFINE FACILITY IXCARM.SYSLVL1.AQSSLADVSR UACC(NONE)
PERMIT IXCARM.SYSLVL1.AQSSLADVSR CLASS(FACILITY) ID(userid) ACCESS(UPDATE)
SETROPTS RACLIST(FACILITY) REFRESH
```

(z/OS systems) Recovering from a planned outage of the TCP/IP stack

Rule: This task is only applicable to Lifeline Advisors and the Lifeline Agents residing on z/OS systems.

If you use automatic restart manager (ARM) and the TCP/IP stack address space terminates, you need to be sure that the system can recover from a planned outage of the TCP/IP stack.

About this task

When the TCP/IP stack becomes unavailable, the Lifeline Advisor also terminates, because the Advisor can no longer establish any TCP/IP communications. An ARM restart of the Lifeline Advisor will likely fail if the TCP/IP stack is not available when the restart occurs. Complete the following steps to recover from a planned outage of the TCP/IP stack:

Procedure

1. Ensure that a secondary Lifeline Advisor has been started on another system.
2. Configure the secondary Lifeline Advisor to automatically take over primary Advisor responsibilities as soon as the primary Lifeline Advisor terminates on the system where TCP/IP is stopped.

Recovering from an unplanned outage of the TCP/IP stack

If you are using automatic restart manager (ARM) and the TCP/IP stack address space terminates, you need to be sure that the system can recover from an unplanned outage of the TCP/IP stack.

About this task

When the TCP/IP stack becomes unavailable, the Lifeline Advisor also terminates, because the Advisor can no longer establish any TCP/IP communications. An ARM restart of the Lifeline Advisor will likely fail if the TCP/IP stack is not available when the restart occurs. Complete the following steps to recover from an unplanned outage of the TCP/IP stack:

Procedure

1. Ensure that an ARM policy or other automation is in place to quickly restart the TCP/IP stack on the same system.
2. Quickly restart the Lifeline Advisor on the same system. You can restart the Lifeline Advisor by using an automation software package, or by using the TCP/IP profile AUTOLOG statement. For more information about the AUTOLOG statement, see *z/OS Communications Server: Configuration Reference*.

(z/OS systems) Configuring Lifeline Agents to automatically restart

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. You can use automation software or define an automatic restart manager (ARM) policy to automatically restart the Lifeline z/OS Agent in case the Agent or system fails. Although this task is optional, you can complete it to provide high availability to your target applications.

About this task

When a Lifeline Agent fails, the Lifeline Advisor indicates that it has no information for any server applications running on that system. As a result, if the second-tier load balancer responsible for distributing workload requests for this site is an external load balancer, target applications on the failing system will no longer receive new workload requests until the Lifeline Agent is restarted. Automatically restarting the Lifeline Agent on the same system minimizes this perceived outage. For more information about defining ARM policies, see *z/OS MVS Setting Up a Sysplex*.

Tip: If the second-tier load balancer is z/OS sysplex distributor, target applications on the failing system might continue to receive new workload requests, as z/OS sysplex distributor does not rely on the Lifeline Agent to retrieve information about these target applications.

The Lifeline Agent registers with ARM by using the following values:

- ELEMTYPE=SYSLVL1
- ELEMNAME=AQSLLAGENT xx
- TERMTYPE=ELEMTERM

xx in the ELEMNAME value is the MVS sysclone value of the system where the Agent resides.

These values indicate that if the Lifeline Agent fails on this system, the Lifeline Agent needs to be restarted only on this system.

To configure Lifeline Agents to automatically restart using ARM, complete the following step:

Procedure

Update the IXCARM resource. The Lifeline Agents do not run using a system key. If you are using ARM registration, the started task IDs must be permitted with UPDATE authority to the IXCARM.SYSLVL1.AQSLLAGENT* generic resource profile in the FACILITY class using the SAF product on your system. To enable the Agent to register with ARM, use the following RACF commands to define the profiles and grant update access:

```
RDEFINE FACILITY IXCARM.SYSLVL1.AQSLLAGENT* UACC(NONE)
PERMIT IXCARM.SYSLVL1.AQSLLAGENT* CLASS(FACILITY) ID(userid) ACCESS(UPDATE)
SETROPTS RACLIST(FACILITY) REFRESH
```

Migrating IBM Multi-site Workload Lifeline

About this task

When you migrate the Lifeline Advisors and Lifeline Agents to the current level, you do not need to stop all Advisors and Agents and migrate them at the same time.

Tip: Migrate both the primary and secondary Advisors before you migrate the individual Agents.

When migrating to IBM Multi-site Workload Lifeline V2.5, Lifeline performs extra verification checks when processing the ACTIVATE, QUIESCE, and DEACTIVATE commands. For more information about these commands, see Chapter 7, “Lifeline Advisor commands,” on page 115.

Activating a workload for the first time

When an Active/Standby or Active/Query workload is activated to a single site for the first time, the Lifeline Advisor queries the Lifeline z/OS Agents on the alternate site to determine whether at least one of the following conditions is true:

- Active TCP/IP connections to the server applications that make up the workload are present on the alternate site.
- Active SNA sessions, for SNA workloads, to the SNA applications that make up the workload are present on the alternate site.
- For MQ cluster workloads, not all MQ cluster receiver channels are stopped or queued messages exist on MQ cluster queues that are hosted by MQ queue managers on the alternate site.

If one of the above conditions is true, the ACTIVATE command fails for the workload. Message AQS0179I is issued to indicate that connections, sessions, or messages were found. For workloads not specified in the mq_manager_list Advisor configuration statement, the workload must first be deactivated to ensure that active connections or sessions are terminated on the alternate site. For workloads specified in the mq_manager_list Advisor configuration statement, the messages must be processed from the MQ cluster queues and the MQ cluster receiver channels must be stopped.

Activating a workload when the workload is not deactivated

When an Active/Standby workload is activated, and the workload is not in a deactivated state, that is, no previous DEACTIVATE command was issued for the

workload, the Lifeline Advisor queries the Lifeline Agents on the alternate site to determine whether at least one of the following conditions is true:

- Active TCP/IP connections to the server applications that make up the workload are present on the alternate site.
- Active SNA sessions, for SNA workloads, to the SNA applications that make up the workload are present on the alternate site.
- For MQ cluster workloads, not all MQ cluster receiver channels are stopped or queued messages exist on MQ cluster queues that are hosted by MQ queue managers on the alternate site.

If one of the above conditions is true, the ACTIVATE command fails for the workload. Message AQS0179I is issued to indicate that connections, sessions, or messages were found. For workloads not specified in the mq_manager_list Advisor configuration statement, the workload must be deactivated to ensure that active connections or sessions are terminated on the alternate site. For workloads specified in the mq_manager_list Advisor configuration statement, the workload must be deactivated; if queued messages remain on the MQ cluster queues or the MQ cluster receiver channels are not stopped, see the System Programmer Response for message AQS0207I for possible actions to take.

For SNA workloads, see “Activating a SNA workload” on page 25.

For MQ workloads, see “Activating an MQ cluster workload” on page 30.

Activating or quiescing a workload when not all first-tier load balancers are registered

When an Active/Standby or Active/Query workload is activated or quiesced, the Lifeline Advisor determines whether all first-tier load balancers, that are configured in the lb_id_list Advisor configuration statement, have registered. If some of the first-tier load balancers are not registered, although the ACTIVATE or QUIESCE command is successful, warning message AQS0208W is issued to indicate that you should take an action to ensure all configured first-tier load balancers are registered.

Activating a workload when not all Lifeline Agents are connected

When an Active/Standby or Active/Query workload is activated to a single site, the Lifeline Advisor determines whether all Lifeline Agents on the alternate site, that are configured in the agent_id_list Advisor configuration statement, are connected. If some of the Agents are not connected from the alternate site, the ACTIVATE command fails for the workload. Message AQS0179I is issued to indicate that not all Agents were connected on the alternate site. The Agents must be connected or removed from the agent_id_list Advisor configuration statement before the ACTIVATE command can be retried.

When an Active/Standby or Active/Query workload is activated, the Lifeline Advisor determines whether any Lifeline Agents on the target site, that are configured in the agent_id_list Advisor configuration statement, are connected. If none of the Agents are connected from the target site, the ACTIVATE command fails for the workload. Message AQS0179I is issued to indicate that no Agents were connected on the target site. One or more Agents must be connected from the target site before the ACTIVATE command can be retried.

When a SNA workload is activated, the Lifeline Advisor determines whether all Lifeline Agents on the target site, that are configured in the `agent_id_list` Advisor configuration statement, are connected. If some of the Agents are not connected from the target site, although the `ACTIVATE` command is successful, warning message `AQS0213W` is issued to indicate that you should take an action to ensure all SNA applications on the target site are activated.

Deactivating a workload when not all Lifeline Agents are connected

When an Active/Standby or Active/Query workload is deactivated, the Lifeline Advisor determines whether all Lifeline Agents on the single or both target sites, that are configured in the `agent_id_list` Advisor configuration statement, are connected. If some of the Agents are not connected from the target site, although the `DEACTIVATE` command is successful, warning message `AQS0209W` is issued to indicate that you should an action to ensure all of the following conditions are true:

- All active TCP/IP connections to the server applications that make up the workload are terminated on the target sites.
- All active SNA sessions, for SNA workloads, to the SNA applications that make up the workload are terminated on the target site.
- For MQ cluster workloads, see the System Programmer Response for message `AQS0207I` for possible actions to take for MQ cluster receiver channels that are not stopped or queued messages that exist on MQ cluster queues that are hosted by MQ queue managers on the target site.

Migrating the Lifeline Advisor

About this task

To migrate the Lifeline Advisors to the current release, migrate the secondary Advisor first.

Rule: If the primary Advisor is at a later release than the secondary Advisor, the secondary Advisor cannot be started because the connection to the primary Advisor fails. Message `AQS0167I` is issued on the primary Advisor.

Procedure

1. Stop the secondary Advisor.
2. Migrate the secondary Advisor to the current release.
3. Optional: Update the configuration file of the secondary Advisor to enable functions introduced in the release.

Rules:

- If the `query_workload_list` Advisor configuration statement is coded and the `timeout` parameter specified, the `timeout` parameter is ignored and message `AQS0311W` is issued until all Agents are updated to V2R0 or a later release. See Multi-site Workload Lifeline configuration statements for more information about configuration statements and parameters.
- If the `mq_manager_list` Advisor configuration statement is coded, the statement is ignored on all Agents that have not been updated to the same or

| later service level that supports this statement. See Multi-site Workload
| Lifeline configuration statements for more information about configuration
| statements and parameters.

4. Restart the secondary Advisor.
5. Perform an Advisor takeover by using the `MODIFY procname,TAKEOVER` command on the secondary Advisor.

Note: The peer Advisor will not remain active when it tries to become the secondary Advisor.

6. Update the peer Advisor to the current release.
7. Update the configuration file of the peer Advisor to match the primary Advisor's configuration.
8. Restart the peer Advisor as the secondary Advisor.
9. Optional: Perform an Advisor takeover by using the `MODIFY procname,TAKEOVER` command on the peer Advisor, so that it is made the primary Advisor again.

Migrating the Lifeline Agent

About this task

Each Lifeline Agent can be migrated to the current release, independent of the release level of the other Agents.

Procedure

1. Stop the Agent.
2. Update the Agent to the current release.
3. Optional: Update the configuration file of the Agent to enable functions that are introduced in the release.
4. Restart the Agent.

Using the Base Control Program Internal Interface

The Lifeline Advisor uses Base Control Program Internal Interface (BCPii) services to periodically query all interconnected Central Processor Complexes (CPCs) and the images (LPARs) on these CPCs.

About this task

The Hardware Management Console (HMC) and Support Elements on each CPC typically communicate over a network that is separate from the network used for TCP/IP communications. The Advisor queries the images in a site over this network to determine whether the images are available or whether a site failure has occurred.

For these queries to succeed, complete the following steps:

Procedure

1. Configure the Support Elements to support BCPii. You need to define an SNMP community name, and be sure that the BCPii address space is active. See the information about BCPii in *z/OS MVS Programming: Callable Services for High-Level Languages* for more information.
2. Authorize the Lifeline Advisor to use the BCPii services. See "Authorizing the Lifeline Advisors to use BCPii" on page 34 for more information.

Configuring and starting syslogd

For the Lifeline Advisors and Agents to be able to write their log messages and trace data to the syslog daemon (syslogd), you must configure and start syslogd before you start the Advisor and Agent.

Before you begin

Note: This task is optional.

- If the AQSTRACE DD card is specified in the Advisor and z/OS Agent started procedure, an internal trace is used to log messages and trace data.
- For Linux on System z Management Agents, internal trace is used by default.

If syslogd is also started, only log messages at the ERROR, WARNING, and EVENT level are written to syslogd.

About this task

The Lifeline Advisor and z/OS Agent write most log messages and trace data to the syslog daemon (syslogd). A limited number of messages are written to the MVS console, but these are unaffected by syslogd configuration.

As the Lifeline Advisor and z/OS Agent log more data, performance can be adversely affected. The amount of data that is logged by the Lifeline Advisor and Agent is determined by the debug_level statement. If you back up the syslogd output file with a zFS file system instead of a z/OS UNIX file system, you can minimize performance impacts caused by logging.

Tips:

- For Linux on System z, it is possible for the Linux on System z Management Agent to selectively send log messages to the Linux on System z Management Guest console. This is controlled by the syslog daemon's configuration settings on the Linux on System z Management Guest.
- To reduce the performance impact caused by logging to syslogd, add the AQSTRACE DD card in the Advisor and z/OS Agent started procedures.
- Logging to syslogd occurs at a default level until the configuration file has been read in.

For more information about configuring syslogd on z/OS, see *z/OS Communications Server: IP Configuration Guide*. See the documentation for your Linux distribution for information about how to configure syslogd on Linux on System z.

(z/OS systems) Customizing the TCP/IP profiles of the TCP/IP stacks

Rule: This task is only applicable to Lifeline Advisors and the Lifeline Agents residing on z/OS systems.

To accommodate the Lifeline Advisors and Agents, you need to customize the TCP/IP profiles of the TCP/IP stacks that the Lifeline Advisors and Agents will run on.

About this task

Complete the following steps to customize the TCP/IP profiles of the TCP/IP stacks:

Procedure

1. Reserve the ports that the Lifeline Advisor and Agents will use in their respective TCP/IP profiles. All ports for the Lifeline Advisor and Agent use the TCP protocol; reserve all the ports for TCP. The Lifeline Advisor has at least two ports, and potentially four ports, to reserve, including the ports specified on the following statements:

- lb_connection_v4
- lb_connection_v6
- agent_connection_port
- advisor_connection_port

The Lifeline Agent has one port to reserve, which you specify on the following statement:

- host_connection

2. Configure VIPARANGE with the IP addresses specified on the Lifeline Advisor lb_connection_v4 and lb_connection_v6 statements on the TCP/IP stacks of the systems that the Lifeline Advisor can run on, and configure static VIPAs with the IP address specified on the Lifeline Agent host_connection statement on the TCP/IP stack of the system that the corresponding Lifeline Agent can run on.
3. Optional: If you use the SHAREPORT or SHAREPORTWLM parameters on the TCP/IP profile PORT statement to enable multiple TCP server applications to share the same port, some additional considerations might apply to your configuration. For example, if the TCP server applications sharing the same port are also members of groups that are reported to external second-tier load balancers with SASP, ensure that consistent criteria are used by the various load balancing components. When using the Multi-site Workload Lifeline, all instances of a TCP server application that share the same port on a target system are reported to external second-tier load balancers using a single member entry, and therefore, a single recommendation. This recommendation reflects the average net weight calculated for all the server applications sharing the same port on a target system. When the TCP connection requests reach a target TCP/IP stack and multiple server applications are sharing the same port, the connections are then load balanced by TCP/IP across the multiple server application instances.

How this load balancing process is performed depends on whether you specify the SHAREPORT or SHAREPORTWLM parameter on the PORT statement. For more details about the PORT statement, see *z/OS Communications Server: IP Configuration Reference*.

4. Ensure that the maximum TCP send and receive buffer sizes are configured with a value of at least 256 KB. The Lifeline Advisor and Lifeline Agent attempt to internally set their send and receive buffer sizes to 256 KB to optimally handle any large messages that are sent between the Advisor and the Agent. Although the Advisor or Agent can operate with lower values for their send and receive buffers, the Advisor or Agent will shut down if many workloads and first-tier load balancers are configured and buffer sizes are not sufficient. For more information about how to use the TCPCONFIG TCPMAXSENDBUFRSIZE and TCPCONFIG TCPMAXRCVBUFRSIZE statements to configure the maximum send and receive buffer sizes, see *z/OS Communications Server: IP Configuration Reference*.

CINET considerations

Although it is not recommended, the Lifeline Advisor can use multiple TCP/IP stacks in a Common INET (CINET) environment.

To simplify your configuration and to make Lifeline Advisor outages that are the result of a TCP/IP stack failure or termination more predictable and recoverable, configure all incoming connections to the Lifeline Advisor to use a single TCP/IP stack. The addresses that you specify in the `lb_connection_v4` and `lb_connection_v6` statements must belong to the same TCP/IP stack.

Use a `VIPARANGE` statement to define these dynamic VIPAs in the TCP/IP profiles of all the stacks in the CINET environment.

If the Lifeline Advisor is restarted as a result of failure in a given TCP/IP stack, the dynamic VIPAs are then activated on another TCP/IP stack in that system. If you decide to use the IPv4 or IPv6 unspecified addresses for the `lb_connection_v4` and `lb_connection_v6` statements, use the `BIND` parameter on the `PORT` reservation statement to bind these sockets to the dynamic VIPAs on the one TCP/IP stack that you have decided to use.

Customizing the VTAM started procedure when configuring SNA workloads

When SNA workloads are defined to Lifeline, customize your VTAM started procedure to spin its job log off periodically.

About this task

To monitor and influence the availability of the SNA applications that make up each configured SNA workload, the Lifeline z/OS Agents issue `DISPLAY` and `VARY` commands against these SNA applications. The output from these `DISPLAY` and `VARY` commands are written to the VTAM's job log. Over time, the job log can consume a considerable amount of JES2 spool space.

Complete the following step to limit the amount of JES2 spool space that is used by the VTAM job log:

Procedure

Add a `JOB DD` card to the VTAM started procedure, and specify the `JESLOG` parameter. For example, to close the current VTAM job log every 12 hours and create a new job log, use the following `DD` card.

```
//<vtamproc> JOB JESLOG=(SPIN,'+12:00')
```

For more details about the `JESLOG` parameter, see *z/OS MVS JCL Reference*.

Customizing the NetView for z/OS Message Revision Table when configuring SNA workloads

When SNA workloads are defined to Lifeline, customize your Message Revision Table (MRT) to control the logging and automation of VTAM messages.

About this task

To monitor and influence the availability of the SNA applications that make up each configured SNA workload, the Lifeline z/OS Agents issue `DISPLAY` and `VARY` commands against these SNA applications. To issue these commands, a Lifeline Agent uses an extended MCS console. The console name for this extended MCS console is the text string 'LLAGNT' that is appended with the MVS sysclone

value of the system where the Agent resides. The output from these DISPLAY and VARY commands are written to the NetView® for z/OS network log and Canzlog. These messages are also presented to the NetView program for automation. This might result in unnecessary overhead by the NetView program for z/OS to process these messages.

To avoid such overhead, complete the following step to suppress command responses that are sent to the Agent's extended MCS console from being written to the NetView for z/OS network log and Canzlog.

Procedure

Add the following entry to the Message Revision Table to control the logging and automation for the command responses that are sent to the Agent's extended MCS console.

```
UPON(OTHERMSG)
  SELECT
    WHEN (CONSNAME = 'LLAGNT&SYSCLONE.')
      REVISE('NO' AUTOMATE 'NO' SYSLOG 'YES' DISPLAY)
    OTHERWISE
  END
```

Notes:

- This entry does not affect an existing Message Revision Table.
- The Message Revision Table supports only one UPON(OTHERMSG) statement. If the UPON(OTHERMSG) statement exists in the Message Revision Table, add only the SELECT statement. Any number of SELECT statements can be included under that UPON statement.

For more details about the Message Revision Table, see *NetView for z/OS Automation Guide*

Configuring external load balancers

Configure external load balancers with the IP address and port of the Lifeline Advisor.

About this task

Guideline: If the Lifeline Advisor uses IPv6 for the external load balancer connections or if any Lifeline Agents use IPv6 to connect to the Lifeline Advisor, start the secondary Lifeline Advisor only on a system that supports IPv6.

For maximum availability, define the IP address of the Lifeline Advisor as a VIPARANGE DVIPA.

There is a maximum limit of 100 external load balancers that can be connected to a Lifeline Advisor at any given time.

Complete the following tasks to configure external load balancers:

Procedure

1. Configure the first-tier external load balancers.
2. Configure the second-tier load balancers. You can have external load balancers, a z/OS sysplex distributor on one z/OS target server in each site, or an intermediary node in each site.

Configuring first-tier external load balancers

Configure the first-tier external load balancers with the IP address and port of the Lifeline Advisor.

About this task

Complete the following steps to configure the first-tier external load balancers:

Procedure

1. Customize features of the communication of the external load balancer with the Lifeline Advisor. For more information about the features, see “External load balancers” on page 10.
2. Optional: Configure redundant external load balancers for availability reasons. The unique load balancer identifier (LB UID), sometimes referred to as the UID or UUID, of the external load balancer identifies a load balancer. Duplicate LB UIDs are not allowed and connection attempts to the Lifeline Advisor from an external load balancer using the same LB UID as an existing connection force the existing connection to be broken and replaced by the new connection.

Requirements:

- If you want the redundant load balancers to serve as hot standbys that are connected simultaneously with the load balancer that they are backing up, you must specify unique LB UIDs.
 - If you configure the redundant load balancers with the same LB UID, the redundant load balancers must remain unconnected from the Lifeline Advisor until the original load balancer fails.
3. Configure external first-tier load balancers to use directed mode. For more information about directed mode, see “External load balancers” on page 10.

Rule: Always configure first-tier load balancers to use directed mode. These load balancers route workload requests only to second-tier load balancers.

4. Configure each first-tier external load balancer with the members that represent the second-tier load balancers.
 - a. Define members that can share the same type of workload under the same load balancer application group. For example, define members for a TN3270E workload under one group and members for an HTTP workload under another. If server applications that access the same data sources are of different types, configure a single workload for these server applications, but define a different load balancer application group for each server application type. This ensures that the loss of all server applications in one group will result in the entire workload being treated as non-operational.
 - b. Specify valid and reachable IP addresses, a nonzero port, and a nonzero protocol. The IP addresses of the members must represent the IP address of an external second-tier load balancer or the DVIPA of a z/OS sysplex distributor. The port number represents the port number used by that server application for the workload. Typically, all members in a group use the same port number. The protocol that is defined can be specified as either tcp or udp.

Rule: All members must be application members; a nonzero port number must be specified.

Configuring second-tier external load balancers

Configure the second-tier external load balancers with the IP address and port of the Lifeline Advisor.

About this task

Complete the following steps to configure the second-tier external load balancers:

Procedure

1. Customize features of the communication of the external load balancer with the Lifeline Advisor. For more information about the features, see “External load balancers” on page 10.
2. Optional: Configure redundant external load balancers for availability reasons. The unique load balancer identifier (LB UID), sometimes referred to as the UID or UUID, of the external load balancer identifies a load balancer. Duplicate LB UIDs are not allowed and connection attempts to the Lifeline Advisor from an external load balancer using the same LB UID as an existing connection force the existing connection to be broken and replaced by the new connection.

Requirements:

- If you want the redundant load balancers to serve as hot standbys that are connected simultaneously with the load balancer that they are backing up, you must specify unique LB UIDs.
 - If you configure the redundant load balancers with the same LB UID, the redundant load balancers must remain unconnected from the Lifeline Advisor until the original load balancer fails.
3. Configure external second-tier load balancers to use directed mode. For more information about directed mode, see “External load balancers” on page 10.

Guideline: You can configure external second-tier load balancers to use either directed or dispatch mode.

4. Configure each external load balancer with the members that represent the individual target application instances.
 - a. Define members that can share the same type of workload under the same load balancer application group. For example, define TN3270E Telnet servers under one group and HTTP servers under another.
 - b. Specify an IP address, a nonzero port, and a nonzero protocol to define application members. The IP addresses of the members must represent valid and reachable addresses in each site where the server applications reside. The port number represents the port number used by that server application and must match the port number specified in the first-tier load balancer for this workload group. Typically, all members in a group will use the same port number. The protocol that is defined can be specified as either tcp or udp and must match the protocol specified in the first-tier load balancer for this workload.

Rules:

- All IP addresses configured in members belonging to the same group must exist on all second-tier external load balancers.
- All members must be application members; a nonzero port number must be specified.
- Certain classes of IP addresses must not be coded for members in the external load balancer. This includes the following classes of addresses:

- Deprecated IPv6 addresses. These are flagged as such in a NETSTAT HOME display. It is probably safest not to code any auto-configured IPv6 addresses within members.
- Addresses that are not reachable from the load balancer, including loopback addresses.
- Unavailable IPv6 addresses. These addresses might be marked as unavailable if duplicate address detection is in progress, has failed, or the interface ID is unknown. These addresses are displayed in a NETSTAT HOME display, along with the reason that they are marked as unavailable.

Configuring second-tier internal load balancers

Configure the second-tier internal load balancers, that is, the z/OS sysplex distributor, with the IP address and port number configured in the first-tier load balancer application group that represents the DVIPA being distributed.

About this task

For information about configuring z/OS sysplex distributors, see *z/OS Communications Server: IP Configuration Reference*.

Configuring support for intermediary nodes

Intermediary nodes do not communicate directly with the Lifeline Advisor, but they function as a second-tier load balancer. A first-tier load balancer registers a group of intermediary nodes. The members of the group are mapped to entries on the `intermediary_node_list` statement.

Before you begin

- If the intermediary node resides on a Linux on System z guest, the intermediary node entries map to entries configured on the Advisor's `management_guest_list` statement. If the application type for the intermediary node entry is IP, the Advisor's `cross_sysplex_list` statement contains entries that map to the database subsystems referenced by the server applications residing on the Linux on System z guests. If the application type for the intermediary node entry is SNA, the Advisor's `sna_application_list` statement contains entries that map to the SNA applications referenced by the server applications residing on the Linux on System z guests.
- If the intermediary node resides on z/OS and services SNA applications, the intermediary node entries map to entries configured on the Advisor's `sna_application_list` statement. The Advisor's `cross_sysplex_list` statement is not used for SNA workloads.
- Otherwise, the `intermediary_node_list` entries for these registered members are mapped to entries configured on the `cross_sysplex_list` statement.

When used, the `cross_sysplex_list` entries are configured with sysplex distributed DVIPAs and ports. The DVIPAs and ports from the `cross_sysplex_list` entries are monitored to determine the weights and availability of the workload's server applications or database subsystems.

About this task

The members of the registered group are classified as either "Internal" or "External" intermediary nodes.

- A group is an external intermediary group if no member IP addresses configured in the `intermediary_node_list` statement are defined on any TCP/IP stacks monitored by the Lifeline Agents; the intermediary nodes are external to the systems where the workload's server applications reside. The ports of the intermediary nodes might not match the ports of the workload's TCP/IP server applications. The `intermediary_node_list` entries are mapped to `cross_sysplex_list` or `sna_application_list` entries by using the site and workload names.
- A group is an internal intermediary group if all member IP addresses configured in the `intermediary_node_list` statement are defined on the TCP/IP stacks monitored by the Lifeline Agents; the intermediary nodes are on the systems where the workload's server applications reside. The ports of the intermediary nodes match the ports of the workload's server applications. The `intermediary_node_list` entries are mapped to `cross_sysplex_list` entries by using the port, site, and workload names.
 - If a workload consists of TCP/IP server applications that reside on z/OS, the ports of the intermediary nodes match the ports of the workload's server applications. The `intermediary_node_list` entries are mapped to `cross_sysplex_list` entries by using the port, site, and workload names.
 - If a workload consists of SNA applications, the `intermediary_node_list` entries are mapped to `sna_application_list` entries by using the site and workload names.
 - If a workload consists of TCP/IP server applications that reside on Linux on System z, the ports of the intermediary nodes might not match the ports of the workload's database subsystems. The `intermediary_node_list` entries are mapped to `management_guest_list` and the `cross_sysplex_list` entries that use the site and workload names.

Procedure

1. Configure the first-tier external load balancers with the groups of the intermediary nodes.
2. Configure the appropriate Lifeline Advisor configuration statements, based on the type of workload:
 - For workloads consisting of TCP/IP server applications residing on z/OS, configure the `cross_sysplex_list` entries with the sysplex distributed DVIPAs and ports of the monitored server applications.
 - For SNA workloads, configure the `sna_application` list entries with the network qualified names of the SNA applications.
 - For Linux on System z workloads, configure the `management_guest_list` with the server IP addresses and Linux on System z guest name for each of the intermediary nodes. Either configure the `cross_sysplex_list` entries with the sysplex distributed DVIPAs and ports of the monitored database subsystems or configure the `sna_application_list` entries with the SNA application names to be monitored.
3. Configure the `intermediary_node_list` entries for the registered group.
 - The `intermediary_node_list` entry is configured as IP or SNA depending on the type of applications that make up the workload. IP is the default configuration.

Note: `intermediary_node_list` entries for a registered load balancer application group cannot be a mix of IP and SNA entries. Only one type is allowed to be configured for a group.

- If this is an external intermediary group, the site and workload of an `intermediary_node_list` entry must match those of one or more `cross_sysplex_list` entries or `sna_application_list` entries.
- If this is an internal intermediary group for a workload that consists of TCP/IP server applications on z/OS, the port, site, and workload of an `intermediary_node_list` entry must match those of one or more `cross_sysplex_list` entries.
- If this is an internal intermediary group for a SNA workload, the entries must be marked as type SNA, and the site and workload of an `intermediary_node_list` entry must match those of a `sna_application_list` entry.
- If this is an external intermediary group for a Linux on System z workload, the site and workload of an `intermediary_node_list` entry must match those of a `management_guest_list` entry and the intermediary node entries must be configured as either IP or SNA. For intermediary nodes configured as IP, the matching of workload and site is performed against `cross_sysplex_list` entries. For intermediary nodes configured as SNA, the matching of workload and site is performed against `sna_application_list` entries.

For more information about how to configure these statements, see “`cross_sysplex_list`” on page 91, “`intermediary_node_list`” on page 94, “`management_guest_list`” on page 99, and “`sna_application_list`” on page 106 in Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85.

Configuring IBM MQ clusters

MQ queue managers configured in an IBM MQ cluster do not communicate directly with the Lifeline Advisor, but the MQ cluster functions as a load balancer to distribute workload messages destined for an MQ cluster queue to an MQ queue manager instance.

See “Overview of MQ Cluster Workloads” on page 28 for more information on configuring IBM MQ clusters and Lifeline to support MQ workloads. See *WebSphere MQ: Installing IBM MQ* for general information about configuring IBM MQ clusters.

Chapter 4. Starting the Lifeline Advisors and Agents

To distribute TCP/IP workloads between the two sites, you need to start the Lifeline Advisors and Agents that make up Multi-site Workload Lifeline.

Before you begin

You need to first complete the tasks in Chapter 2, “Preparing to use Multi-site Workload Lifeline,” on page 33 and Chapter 3, “Configuring Multi-site Workload Lifeline,” on page 43.

About this task

You can start the Lifeline Agents before the Lifeline Advisors, or start the Lifeline Advisors before the Lifeline Agents. If you start the Lifeline Advisor after you start the Lifeline Agents, the Lifeline Agents periodically attempt to connect to the Lifeline Advisor.

Procedure

1. Start the TCP/IP stacks that the Lifeline Advisors and the Lifeline z/OS Agents use.

Requirement: You must start the TCP/IP stacks that the Lifeline Advisors use before you start the Lifeline Advisors. You can start a Lifeline z/OS Agent before the start of the TCP/IP stack that the Lifeline z/OS Agent uses. If the TCP/IP stack that the Lifeline Agent uses stops, the Lifeline Agent remains active and reestablishes communication with the TCP/IP stack when the TCP/IP stack is restarted.

Note: If you use automation software to start the primary Lifeline Advisor, you need to consider other factors to determine when to start the Advisor.

If you configure the Lifeline Advisor to use dynamic VIPAs, your automation software must perform the following steps:

- Ensure that TCP/IP is started.
- Wait for message EZD1214I to be issued. Message EZD1214I indicates that TCP/IP completed dynamic VIPA processing.

For more information about using dynamic VIPAs, see “Configuring one Lifeline Advisor as the primary Advisor, and optionally one as the secondary Advisor” on page 44.

2. Start the server applications that will be the target of load balancing. You do not need to modify these server applications, their configurations, or started procedures, unless the second-tier external load balancer is using dispatch mode for packet forwarding. For more information about dispatch mode, see “External load balancers” on page 10.
3. Optional: Customize WLM policies for the Lifeline Advisors and Lifeline z/OS Agents, and assign the Lifeline Advisors and Agents to the WLM SYSSTC service class to receive the proper dispatching priority. For more information about categorizing work into service classes, see *z/OS MVS Planning: Workload Management*.

Rule: Ensure that the Lifeline Advisors and Lifeline z/OS Agents receive an adequate amount of system resources to properly balance workloads.

4. Start the primary Lifeline Advisor from a started procedure as a started program (EXEC PGM=). See Appendix A, "A sample started procedure for Lifeline Advisors," on page 253 for a sample Advisor start procedure. Verify that each Lifeline Agent that you expect to connect to the Lifeline Advisor has connected. You can also use the NETSTAT CONN command on the Lifeline Advisor's TCP/IP stack to see which Lifeline Agents are currently connected. As Lifeline Agents connect to the Lifeline Advisor, messages appear on the MVS consoles where the Advisor and Agents are active.

Restriction: You cannot start the Lifeline Advisor under BPXBATCH.

5. Start a Lifeline z/OS Agent and, if required, a Linux on System z Management Agent.
 - Start one Lifeline z/OS Agent from a started procedure as a started program (EXEC PGM=) on each z/OS target system that you want to participate in this method of workload balancing. See Appendix B, "A sample started procedure for Lifeline z/OS Agents," on page 255 for a sample Agent start procedure.
 - If you configure Linux on System z workloads, start one Linux on System z Management Agent on each target z/VM system where the Linux on System z guests reside for the server applications that make up the workload. Start the Linux on System z Management Agent by running the aqsagzvm program. See "Linux on System z Workloads" on page 13 for information about this program.
- a. Optional: Update the SCHEDxx member to add an entry to the program properties table to make the Lifeline Advisors and Lifeline z/OS Agents run non-swappable.

Restrictions:

- You can start only one Lifeline z/OS Agent per z/OS system.
 - You must not start the Lifeline z/OS Agent under BPXBATCH.
 - You can start only one Lifeline Linux on System z Management Agent per site per z/VM system.
6. Optional: Start the secondary Lifeline Advisor. When the secondary Advisor connects to the primary Advisor, messages are issued on the MVS consoles where each Lifeline Advisor is running.
 7. Optional: Start the first-tier and second-tier load balancers. When an external load balancer has connected, a message is issued on the MVS console where the Lifeline Advisor is running.

|
|
|

Chapter 5. Diagnosing problems in Multi-site Workload Lifeline

You can verify the performance and availability of the system and diagnose problems when you use Multi-site Workload Lifeline.

Debug levels

You can update the debug level to change the level of debug messages to be logged.

The following list describes available debug levels and the associated syslogd priority levels:

Debug level

syslogd priority level

- | | |
|----|---|
| 0 | NONE. No messages of any kind are sent to the syslogd file after initialization is complete. |
| 1 | ERROR. Error messages indicate something that requires attention. Messages at this level could be fatal (terminating) or could indicate that an integral part of the workload advising system is not working properly. This information is logged at the syslogd ERROR priority level. |
| 2 | WARNING. Warning messages indicate that an error has occurred, but it is not severe enough to warrant an ERROR. Corrective actions might be necessary because the Advisor or Agent might not be behaving as intended. This information is logged at the syslogd WARNING priority level. |
| 4 | EVENT. Event messages are logged for events that occur periodically, like operator commands, UNIX signals, timer pops, and receipt of a network message. This information is logged at the syslogd NOTICE priority level. |
| 8 | INFO. Informational messages are sent to the syslogd file. These messages do not require corrective actions. This information is logged at the syslogd INFO priority level. |
| 16 | MESSAGE. Message messages log the contents of message packets that are sent between the Advisor and the Load Balancer, the Advisor and peer Advisor, or between the Advisor and Agent. You can use these messages to assist debugging Advisor-to-load balancer, Advisor-to-Advisor, and Advisor-to-Agent communications. This information is logged at the syslogd DEBUG priority level. This level is intended only for IBM service use. |
| 32 | COLLECTION. Collection messages log the details of collecting and manipulating the data that forms the basis of weight calculations. This information is logged at the syslogd DEBUG priority level. COLLECTION is only used by the Agent. This level is intended only for IBM service use. |
| 64 | DEBUG. The intention of this level of message is to provide information that is useful in debugging code, logic, or timing errors. This information is logged at the syslogd DEBUG priority level. This level is intended for IBM service use only. |

- 128 TRACE. Trace messages are intended for development or service to track code processing (footprints). This information is logged at the syslogd DEBUG priority level. This level is intended for IBM service use only.

Diagnostic data

You can collect multiple pieces of diagnostic data to diagnose problems.

For a workload that consists of TCP/IP server applications on z/OS or SNA workloads, the following diagnostic data is helpful:

- Console messages for the Lifeline Advisors and Agents
- Output from the MODIFY command for the Lifeline Advisors and Lifeline z/OS Agents
- Either syslogd log messages or the dumped internal trace file for Lifeline Advisors and Agents
- Lifeline Advisor and Lifeline z/OS Agent address space dumps and snap output
- TCP/IP packet traces of communication with the primary Lifeline Advisor

For Linux on System z workloads, the following diagnostic data is helpful:

- Console messages for the Lifeline Advisors and Agents
- Output from the MODIFY command for the Lifeline Advisors and Lifeline z/OS Agents
- Output from the aqslcmd command for the Lifeline Linux on System z Management Agent
- Either syslogd log messages or the dumped internal trace file for Lifeline Advisors and Agents
- Lifeline Advisor and Lifeline z/OS Agent address space dumps and snap output or core dumps for the Lifeline Linux on System z Management Agent
- TCP/IP packet traces of communication with the primary Lifeline Advisor

Guidelines:

- If the AQSTRACE DD card is not specified in the Lifeline Advisor or z/OS Agent started procedure, syslogd is the only logging facility that either the Advisors or Agents can use. Useful diagnostic information might be lost if syslogd is not running before the Advisor or Agents are started.
- If the AQSTRACE DD card is specified in the Lifeline Advisor and z/OS Agent started procedure, you do not have to run syslogd. The Advisors and z/OS Agents use an internal trace as their logging facility. This internal trace is dumped to the file specified by the AQSTRACE DD card when one of the following conditions is true:
 - The Advisor or z/OS Agent is stopped.
 - The DUMPLOG command is issued against the Advisor or z/OS Agent.

The UNIX System Services file system where the file resides should have 256 MB of available space to prevent truncation of the dumped internal log. A fixed number of debug records are stored in the internal trace. When this number of records has been reached, the logging wraps. See Chapter 7, “Lifeline Advisor commands,” on page 115 for information about the DUMPLOG command for the Advisor and Chapter 8, “Lifeline z/OS Agent commands,” on page 151 for information about the DUMPLOG command for the z/OS Agent.

Rule: The UNIX System Services file used for dumping the Advisor or z/OS Agent internal trace must meet these requirements:

- The write permission of the file can be set only for the user and group classes.
- The others class can be granted only read permission.
- Linux on System z Management Agents use internal trace by default. This internal trace is dumped to the `/opt/ibm/lifeline/log/agent.trace` file when one of the following conditions is true:
 - The Linux on System z Management Agent is stopped.
 - The DUMPLOG command is issued.

The `/opt/ibm/lifeline/log` file system should have 256 MB of available space to prevent truncation of the dumped internal log. A fixed number of trace records are stored in internal memory. When this number of records has been reached, the logging wraps. For more information about the DUMPLOG command, see “Linux on System z Workloads” on page 13.

Rule: The Linux on System z file used for dumping the Linux on System z Management Agent internal trace must meet these requirements:

- The write permission of the file can be set only for the user and group classes.
- The others class can be granted only read permission.
- If both `syslogd` runs and the internal trace is used, log messages at ERROR, WARNING, and EVENT level are written to both logs. Log messages at more detailed levels, INFO, MESSAGE, COLLECT, DEBUG, and TRACE, are written only to internal trace.

The Lifeline Advisor and Lifeline z/OS Agent trigger address space dumps when they encounter certain unexpected error conditions. Both a CEEDUMP and address space snap output are produced and written to the data sets or files that are specified by the started procedure CEEDUMP and CEESNAP DD statements, respectively. The Lifeline Linux on System z Management Agents produce core dumps when they encounter unexpected error conditions.

If the Lifeline Advisor or Lifeline z/OS Agent abnormally stops (for example, an 0C4 abend occurs), an unformatted SYSMDUMP is produced and written to the data set that is specified by the started procedure SYSMDUMP DD statement. If you override the Language Environment® run-time option TERMTHDACT during the installation or started procedure, the SYSMDUMP might not be produced, or a CEEDUMP might be produced instead. Therefore, do not override the TERMTHDACT run-time option. See *z/OS Language Environment Programming Guide* for more information about run-time options. When a Linux on System z Management Agents abnormally stops, it produces a core dump.

In other situations, the z/OS operator needs to dump the address space manually.

You can collect the TCP/IP packet trace data of communication with the primary Lifeline Advisor. For example, you might need to capture the following communication flows:

- Server/Application State Protocol (SASP) messages sent between the primary Advisor and external load balancers
- Messages sent between the primary Advisor and Agents
- Messages sent between the primary and secondary Advisors

For details about how to use the TCP/IP packet trace facility, see *z/OS Communications Server: IP Diagnosis Guide*.

Interpreting the workload and LB displays

The output from the `MODIFY advisor_proc,DISPLAY,WORKLOAD,DETAIL` command shows all of the registered members that are mapped to a workload with each member's site and availability as reported by the Agent from the system where the server application resides.

The output from the `MODIFY advisor_proc,DISPLAY,LB,DETAIL` command shows the load balancer application groups registered by a load balancer with their matching workload. Each member's status is displayed, including member availability and weight information. This status information shows what the Advisor is sending to the load balancer and might not reflect the availability of a member's server application. Member availability depends on several factors as detailed below.

Load Balancer display for second-tier load balancers

If a load balancer application group is registered by a second-tier load balancer, the member status in the `MODIFY advisor_proc,DISPLAY,LB,DETAIL` command always matches the actual status of the members in the corresponding `MODIFY advisor_proc,DISPLAY,WORKLOAD,DETAIL` command.

Load Balancer display for first-tier load balancers

If a load balancer application group is registered by a first-tier load balancer, the member status in the `MODIFY advisor_proc,DISPLAY,LB,DETAIL` command does not always match the actual status shown in the `MODIFY advisor_proc,DISPLAY,WORKLOAD,DETAIL` command:

- If the members in the group are sysplex distributed DVIPAs or external second-tier load balancers:
 - If the workload is not active to a site, the LB display shows the members of the inactive site with a status of unavailable because the first-tier load balancer cannot route workload connections to this site.
 - If a workload is active only to one site and at least one member on the active site is available, the LB display of each member on the site matches the actual status that is shown in the Workload display.
 - If a workload is active only to one site and no members on the active site are available, the LB display shows all members with a status of available. It ensures that the first-tier load balancer routes new workload connections only to the active site.
 - If the members in the group are sysplex distributed DVIPAs, each DVIPA represents a different set of server applications within the site. Each weight, as shown in the LB display, is relative to the other members of this group.
 - If the members in the group are external second-tier load balancers, each second-tier load balancer should be distributing to the same set of server applications within a site. Therefore the weight of each second-tier load balancer within a site is the same value.
 - If an Active/Query workload is active to both sites and static routing is configured in the `query_workload_list` statement, the weights of the members are adjusted so that a site's weight which is the sum of all member weights within the site matches the configured percentages.
- If the members in the group are intermediary nodes:

- If the workload is not active to a site, the LB display shows the members of the inactive site with a status of unavailable because the first-tier load balancer cannot route workload connections to this site.
- If a workload is active only to one site, the LB display shows all available members, that is, the internal intermediary nodes, of the site with a status of available and an equal weight regardless of the actual weight reported by the Agent. This ensures that the first-tier load balancer routes new workload connections to the active site, but does not prefer one internal intermediary node over another within the site. The intermediary node makes the load balancing decision within a site.
- SNA workloads are always Active/Standby workloads. For the active site, the LB display shows the actual availability of each SNA application. Each SNA application is shown with the same weight because the intermediary node makes the load balancing decision within the site. The weights of the available intermediary node servers are equal. This ensures that the first-tier load balancer routes new workload connections to the active site, but does not prefer one member over another within the site. The availability of each intermediary server is shown when either of the following conditions is true:
 - The LB group is an internal intermediary group.
 - The LB group is an external intermediary group that consists of intermediary nodes on Linux on System z guests being monitored by Linux on System z Management Agents.
- If an Active/Query workload is active to both sites, the LB display shows all members within a site with the same weight. However, the weights of the two sites, which are the sum of all member weights, might be different. If dynamic routing is configured on the `query_workload_list` statement, the site weights match the site's relative capacity for processing new connections for this workload. If static routing is configured on the `query_workload_list` statement, the site weights match the configured percentages. It allows the first-tier load balancer to route connections for the workload between the sites that are based on the configured routing type. The intermediary node makes the load balancing decision within a site.
- For an External Intermediary group, the output from this command shows the DVIPAs and ports for the `cross_sysplex_list` entries that are mapped by the intermediary node list. A member is available if a server application is using the `cross_sysplex_list` entry's port on the TCP/IP stack where the entry's DVIPA is active.
- For an Internal Intermediary group, the output from this command shows the IP address and ports for the `intermediary_node_list` entries. A member is available if all of the following conditions are true:
 - The DVIPA on the mapped `cross_sysplex_list` entry is active on the same TCP/IP stack where the IP address on the `intermediary_node_list` entry is active.
 - A server application that is bound to `INADDR_ANY` or `IN6ADDR_ANY` for IPv6 addresses by using the mapped `cross_sysplex_list` entry's port is on the same TCP/IP stack where the IP address on the `intermediary_node_list` entry is active.
 - No server application is bound to the mapped `cross_sysplex_list` entry's DVIPA and port on the same TCP/IP stack where the IP address on the `intermediary_node_list` entry is active.

If the output from the command indicates that an internal intermediary member is not available, check the syslogd file of the Advisor for a warning message to see why the member is not available.

Verifying that the Lifeline Advisor and Agent systems are functioning correctly

View the MVS console and logs of the Lifeline Advisor and Lifeline z/OS Agents to verify that the Advisor and Agent systems are functioning correctly. For Linux on System z workloads, you also need to view the logs of the Lifeline Linux on System z Management Agents to verify these Agent systems.

About this task

View the MVS console of the Lifeline Advisor and Lifeline z/OS Agent systems after they are started to verify that the applications are started and running. For Linux on System z workloads, you also need to check the current processes on the Linux on System z Management Guest for the aqsagzvm process. If internal trace is used, dump and view the Lifeline Advisor and Agent internal log messages to determine whether any error or warning messages were issued. Otherwise, view the syslogd files of the Lifeline Advisor and Agent systems to check for these messages.

If there are any failure messages, see Chapter 11, “Multi-site Workload Lifeline messages,” on page 179 for the corrective actions that you need to take.

Complete the following steps to verify that the Lifeline Advisors and Lifeline Agents are functioning correctly:

Procedure

1. Determine whether the Lifeline Advisor is started and, if applicable, connected to the expected external load balancers by issuing the following command on the Lifeline Advisor:

```
MODIFY procname,DISPLAY,LB
```

Verify that each external load balancer is displayed.

2. Determine whether each external load balancer configured and registered the proper load balancer application groups and members with the Lifeline Advisor.

Tip: The Lifeline Advisor does not check for improperly configured members.

- a. After the entire Multi-site Workload Lifeline system is started and running, display all members registered by each external load balancer by issuing the following command:

```
MODIFY procname,DISPLAY,LB,DETAIL
```

This command displays all groups and members that are defined to each of the load balancers. Verify that each member that you expect to be available is flagged as available.

- b. Display all members for each workload by issuing the following command:

```
MODIFY procname,DISPLAY,WORKLOAD,DETAIL
```

Verify that each member that you expect to be available for a workload is flagged as available.

- c. Screen unavailable IP members for configuration errors, such as incorrect IP addresses, ports, or protocols.
 - d. Screen unavailable SNA members for configuration errors, such as incorrect SNA application names.
3. Determine whether each Lifeline Agent has started correctly and is communicating with the Lifeline Advisor by issuing one of the following commands on each Lifeline Agent:

MODIFY *procname*,DISPLAY,MEMBERS

The command for z/OS Agents. Each registered member, either configured from a load balancer application group or an MQ cluster queue group, is displayed in the output if the member has an IP address owned by the TCP/IP stack where the Lifeline z/OS Agent is active.

aqsl1cmd display,members

The command for Linux on System z Management Agents. Each registered member is displayed in the output if the member configured from a load balancer application group, maps to a Linux on System z guest on the z/VM system where the Linux on System z Management Agent is active, by using the `management_guest_list` statement.

4. Determine whether the target server applications and MQ queue managers for the workload are available for load balancing. On the Lifeline Advisor, issue the following command:

MODIFY *procname*,DISPLAY,WORKLOAD,DETAIL

- a. In the display output, check the state for each member defined for this workload. The state is either `AVAIL`, meaning that the member is available for load balancing, or `UNAVAIL`, meaning that it is not available for load balancing.
- b. For first-tier external load balancers, to make the member available for load balancing, ensure that either of the following condition is true:
 - The external second-tier load balancers must be started and connected to the Lifeline Advisor.
 - The z/OS sysplex distributor that is acting as a second-tier load balancer must be configured on a z/OS image within the site, and the Lifeline Agent on that z/OS image must be active and communicating with the Lifeline Advisor.

If both of the conditions are false, correct the situation and repeat the display command until you are satisfied that all members that you intend to make available for load balancing are displayed as being available.

- c. For second-tier external load balancers, to make the member available for load balancing, ensure that all of the following conditions are true:
 - The Lifeline z/OS Agent must be active on the TCP/IP stack owning the IP address of the member and be communicating with the Lifeline Advisor.
 - The server application must be active.
 - The member must not be quiesced by the Lifeline z/OS Agent operator or the external load balancer. The Lifeline Agent is able to quiesce any member that is owned by the TCP/IP stack where the Agent is active. Depending on the external load balancer implementation, the load balancer administrator might quiesce individual members.

If one of the conditions is false, correct the situation and repeat the display command until you are satisfied that all members that you intend to make available for load balancing are displayed as being available.

- d. For workloads utilizing IBM MQ clusters, to make the member available for load balancing, ensure that all of the following conditions are true:
- The MQ queue manager must be started.
 - The MQ cluster queue is defined and enabled on the MQ queue manager.
 - The MQ cluster receiver channel is started.
 - The TCP listener used by the MQ cluster receiver channel is available.

If any of the conditions is false, correct the situation and repeat the display command until you are satisfied that all members that you intend to make available for load balancing are displayed as being available.

Diagnosing problems when a workload is not distributed

Diagnosing why a workload is not being distributed to a server application

When a workload is not being distributed to a server application, you can use the information in this topic to diagnose problems.

About this task

To determine why a workload is not being distributed to a server application, complete the following steps:

Procedure

1. Determine whether the Advisor is running, a z/OS Agent is running on the MVS system that contains the server application, or a Linux on System z Management Agent is running in a Management Guest on the z/VM system that contains the server application. If they are not running, start the Advisor or Agent.
2. To determine whether any external load balancers have registered the server application, issue the following command on the Lifeline Advisor:
`MODIFY procname,DISPLAY,LB`

Verify that at least one first-tier load balancer is connected to the Advisor.

3. Determine whether the `lb_id_list` statement of the Advisor includes the IP address of the load balancer in question.
If you use Application Transparent TLS (AT-TLS) to authenticate the load balancer, verify that the TLS configuration is correct on the load balancer and Advisor. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.
4. For external first-tier load balancers, determine whether the IP address of the member defined on the load balancer matches the IP address of a second-tier load balancer responsible for load balancing to the server application. If not, correct the definition at the first-tier load balancer.
5. For external second-tier load balancers, determine whether the IP address, port number, and protocol of the member defined on the load balancer match the IP address, port number, and protocol of the server application. If not, correct the definition at the second-tier load balancer.

6. To check the member information on the Lifeline Advisor, issue the following command:

```
MODIFY procname,DISPLAY,LB,DETAIL
```

If the DISTDVIPA flag is displayed for the server application, verify that z/OS sysplex distributor is distributing workload requests to the server application. See *z/OS Communications Server: IP Diagnosis* for more information about diagnosing dynamic VIPA and sysplex problems.

7. Determine whether the `agent_id_list` statement of the Advisor contains the IP address and port that the Agent is bound to on the system where the server application exists. If not, correct the `agent_id_list` statement in the Advisor configuration or the `host_connection` statement in the Agent configuration. If you use Application Transparent TLS (AT-TLS) to authenticate the z/OS Agent, verify that the TLS configuration is correct on the Agent and Advisor. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.
8. To check the member information on the Agent, issue the following command:

```
MODIFY procname,DISPLAY,MEMBERS
```

The command for z/OS Agents

```
aqsl1cmd display,members
```

The command for Linux on System z Management Agents

If the protocol is UDP and the IP address of the server application is a distributed DVIPA address, no workload requests will be distributed for this server application on target TCP/IP stacks. Only the z/OS sysplex distributor TCP/IP stack advertises the DVIPA address, so only server applications running on this TCP/IP stack that binds to the DVIPA for UDP have workload requests distributed to it. Server applications that bind to the distributed DVIPA for protocol UDP on other TCP/IP stacks do not have workload requests distributed to them. On Linux on System z Management Agents, if the member is not displayed, ensure that the Linux on System z guest that hosts the server application is logged on.

9. Determine whether network connectivity exists between the Advisor and the Agent in question. Unexpected loss of network connectivity between the Advisor and the z/OS Agent results in MVS console messages for the Advisor and Agent and related messages in the Advisor and Agent logs. Unexpected loss of network connectivity between the Advisor and the Linux on System z Management Agent results in MVS console messages for the Advisor and related messages in the Agent and Advisor logs. To determine which Agents have connected to the Advisor, and by omission, which have not, issue the following Advisor command:

```
MODIFY procname,DISPLAY,ADVISOR
```
10. To determine whether the server application is available and enabled but not quiesced for the workload, issue the following Advisor command:

```
MODIFY procname,DISPLAY,WORKLOAD,DETAIL
```

 - a. If the member state is UNAVAIL, correct the issue with the server application.
 - b. If the server application is quiesced, enable the server application by using either of the following commands:
 - If the second-tier load balancer is an external load balancer, issue the `MODIFY procname,ENABLE` command on the system where the Agent is active.

- If the second-tier load balancer is z/OS sysplex distributor, issue the VARY TCPIP,,SYSPLEX,RESUME command on the system where the server application is active.
11. If syslogd was started, check the syslogd file for ERROR or WARNING messages and take the appropriate corrective actions. The syslogd identifier is lladvsr for the Advisor and llagent for the Agent. Otherwise, if the Advisor or Agent internal trace is being used, dump and check the internal trace for ERROR or WARNING messages. In both cases, take the appropriate corrective actions. If ERROR and WARNING level log messages are not enabled, enable them and recheck the syslogd file or dump and recheck the internal trace messages later.
 12. Determine whether the first-tier load balancer has connectivity to the second-tier load balancer, and the second-tier load balancer has connectivity to the TCP/IP stack owning the IP address of the server application in question.
 13. Issue the following Advisor command to determine whether the workload has been activated to one of the sites:

```
MODIFY procname,DISPLAY,WORKLOAD
```

If not, activate the workload by issuing the MODIFY *procname*,ACTIVATE command.

Diagnosing why a workload is not being distributed to an MQ cluster queue

When a workload is not being distributed to an MQ cluster queue, you can use the information in this topic to diagnose problems.

About this task

To determine why a workload is not being distributed to an MQ cluster queue, complete the following steps:

Procedure

1. Determine whether the Advisor is running or a z/OS Agent is running on the MVS system that contains the MQ queue manager. If they are not running, start the Advisor or Agent.
2. Determine whether the agent_id_list statement of the Advisor contains the IP address and port that the z/OS Agent is bound to on the system where the MQ queue manager exists. If not, correct the agent_id_list statement in the Advisor configuration or the host_connection statement in the z/OS Agent configuration. If you use Application Transparent TLS (AT-TLS) to authenticate the z/OS Agent, verify that the TLS configuration is correct on the Agent and Advisor. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.
3. To check the member information on the Agent, issue the following command:

```
MODIFY procname,DISPLAY,MEMBERS,DETAIL
```

If the MQ queue manager is not active, restart the MQ queue manager. If not all MQ cluster queues for the workload are defined to the MQ queue manager, ensure these MQ cluster queues are defined. If not all MQ cluster queues for the workload are enabled to receive messages, ensure these MQ cluster queues are enabled.

4. Determine whether network connectivity exists between the Advisor and the z/OS Agent in question. Unexpected loss of network connectivity between the

Advisor and the z/OS Agent results in MVS console messages for the Advisor and Agent and related messages in the Advisor and Agent logs. To determine which z/OS Agents have connected to the Advisor, and by omission, which have not, issue the following Advisor command:

```
MODIFY procname,DISPLAY,ADVISOR
```

5. If syslogd was started, check the syslogd file for ERROR or WARNING messages and take the appropriate corrective actions. The syslogd identifier is lladvr for the Advisor and llagent for the z/OS Agent. Otherwise, if the Advisor or z/OS Agent internal trace is being used, dump and check the internal trace for ERROR or WARNING messages. In both cases, take the appropriate corrective actions. If ERROR and WARNING level log messages are not enabled, enable them and recheck the syslogd file or dump and recheck the internal trace messages later.
6. Issue the following Advisor command to determine whether the workload has been activated to one of the sites:

```
MODIFY procname,DISPLAY,WORKLOAD
```

If not, activate the workload by issuing the `MODIFY procname,ACTIVATE` command.

Diagnosing workload distribution problems

Diagnosing workload distribution problems for server applications

When workload distribution problems occur, you can use the information in this topic to diagnose problems.

About this task

To diagnose workload distribution problems for server applications, complete the following steps:

Procedure

1. To determine whether a first-tier load balancer has connected to the Lifeline Advisor, issue the following command:

```
MODIFY procname,DISPLAY,LB,DETAIL
```

Ensure that there are groups in the display with the `CROSS_SYSPLEX` group flag set, and that the IP addresses listed in those groups are the IP addresses of the second-tier load balancers.

2. Determine whether the update interval value of the Advisor is not abnormally large. The Advisor must wait at least two update intervals before beginning to receive enough data to properly calculate weights when a server application becomes available or when an Agent is started.
 - a. Allow at least three update intervals to expire after a server application is started before re-examining the distribution of workload.
 - b. If the workload is occasionally being sent to overloaded server applications, adjust the `update_interval` downward so workload distribution can react more quickly to the pace of new workload requests.
3. To check routing recommendations, periodically issue the following Advisor command to check the weights of members within the group in question.

```
MODIFY procname,DISPLAY,LB,DETAIL
```

Determine whether the weights are consistent with the expected behavior. If not, see “MODIFY *advisor_proc*,DISPLAY command: Display Lifeline Advisor information” on page 115, and “Interpreting the workload and LB displays” on page 70 for more information about how to analyze the server application weights.

4. Determine whether the `agent_id_list` value of the Advisor contains the IP addresses and ports of the MVS systems that each z/OS Agent is bound to and the Linux on System z Management Guests that each Linux on System z Management Agent is bound to. If not, correct the `agent_id_list` statement in the Advisor configuration or the `host_connection` statement in the Agent configuration.

If you use Application Transparent TLS (AT-TLS) to authenticate the z/OS Agent, verify that the TLS configuration is correct on the Agent and Advisor. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS. The Advisor must be configured not to require Linux on System z Management Agents to authenticate by using AT-TLS.

5. Determine whether members of the group in question using this workload are not unexpectedly quiesced or unexpectedly unavailable by issuing the following Advisor command:

```
MODIFY procname,DISPLAY,WORKLOAD,DETAIL
```

If the member state is UNAVAIL, correct the issue with the server application.

6. If `syslogd` was started, check the `syslogd` file for ERROR or WARNING messages and take the appropriate corrective actions. The `syslogd` identifier is `lladvr` for the Advisor and `llagent` for the Agent. Otherwise, if the Advisor or Agent internal trace is being used, dump and check the internal trace for ERROR or WARNING messages. In both cases, take the appropriate corrective actions. If ERROR and WARNING level log messages are not enabled, enable them and recheck the `syslogd` file or dump and recheck the internal log messages later.

Diagnosing workload distribution problems for IBM MQ clusters

When workload distribution problems occur, you can use the information in this topic to diagnose problems.

About this task

To diagnose workload distribution problems for IBM MQ clusters, complete the following steps:

Procedure

1. Determine whether the update interval value of the Advisor is not abnormally large. The Advisor must wait at least two update intervals before beginning to receive enough data to properly calculate weights when an MQ queue manager or MQ cluster queue becomes available or when a z/OS Agent is started.
 - a. Allow at least three update intervals to expire after the MQ queue manager is started before re-examining the distribution of workload.
 - b. If the workload messages are occasionally being sent to overloaded MQ queue managers, adjust the `update_interval` downward so workload distribution can react more quickly to the pace of new workload requests.

2. To check routing recommendations, periodically issue the following z/OS Agent command to check the availability of members within the MQ cluster queue group in question.

```
MODIFY procname,DISPLAY,MEMBERS,DETAIL
```

3. Determine whether the `agent_id_list` value of the Advisor contains the IP addresses and ports of the MVS systems that each z/OS Agent is bound to. If not, correct the `agent_id_list` statement in the Advisor configuration or the `host_connection` statement in the z/OS Agent configuration. If you use Application Transparent TLS (AT-TLS) to authenticate the z/OS Agent, verify that the TLS configuration is correct on the Agent and Advisor. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring AT-TLS.

4. Determine whether MQ queue managers or MQ cluster queues for the workload in question are not unexpectedly disabled or unexpectedly unavailable by issuing the following Advisor command:

```
MODIFY procname,DISPLAY,WORKLOAD,DETAIL
```

If the member state is UNAVAIL, correct the issue with the MQ queue manager or MQ cluster queue.

5. If `syslogd` was started, check the `syslogd` file for ERROR or WARNING messages and take the appropriate corrective actions. The `syslogd` identifier is `lladvsr` for the Advisor and `llagent` for the z/OS Agent. Otherwise, if the Advisor or Agent internal trace is being used, dump and check the internal trace for ERROR or WARNING messages. In both cases, take the appropriate corrective actions. If ERROR and WARNING level log messages are not enabled, enable them and recheck the `syslogd` file or dump and recheck the internal log messages later.

Diagnosing problems when the Lifeline Advisor or Agent appears hung

When the Lifeline Advisor or Agent appears hung, you can use the information in this topic to diagnose problems.

About this task

To diagnose problems when the Lifeline Advisor or Agent appears hung, complete the following steps:

Procedure

1. To determine whether the Lifeline Advisor or Agent is hung, issue the following command on the system where the Advisor or Agent is running:

```
MODIFY procname,DISPLAY,DEBUG
```

The command for Advisors

```
MODIFY procname,DISPLAY,DEBUG
```

The command for z/OS Agents

```
aqs11cmd display,debug
```

The command for Linux on System z Management Agents

If you do not receive a response, attempt to stop the Advisor or Agent instead of canceling the Advisor or Agent. If the Advisor or Agent does not terminate, it might be hung.

2. If the hang occurred when the DEBUG-level Advisor or Agent trace was in effect, collect the problem documentation by taking the following sub-steps and call IBM Service.
 - a. Take an SVC dump of the z/OS Agent or Advisor address space (depending on which is hung) and of the OMVS address space including its data spaces. For the Linux on System z Management Agent, trigger a core dump, by using the `kill -SIGQUIT <process_id>` command.
 - b. Capture the MVS console messages if applicable.
 - c. If `syslogd` was started, capture the Advisor or Agent messages that are written to `syslogd`. The `syslogd` identifier is `lladvsr` for the Advisor and `llagent` for the Agent.
 - d. If the Advisor or Agent internal trace is used, dump and save the internal trace messages.
3. If DEBUG-level trace was not in effect at the time, turn on DEBUG-level Advisor or Agent trace, reproduce the problem, collect the problem documentation, and call IBM Service.

Diagnosing problems when the group names are not readable

When the load balancer application group names are not readable, you can use the information in this topic to diagnose problems.

About this task

To diagnose problems when the load balancer application group names are not readable, complete the following steps:

Procedure

1. Determine whether the group names defined in the load balancers use characters limited to the ASCII character set. When the load balancers define group names, the names are coded in UTF-8 format. This character set is a superset of the EBCDIC character set; not all characters are translatable to EBCDIC.
2. To make the group names readable, rename the group names in the load balancers to use characters limited to the ASCII character set.

Diagnosing problems when load balancer connection terminates

When the external load balancer connection to the Advisor terminates, you can use the information in this topic to diagnose problems.

About this task

Guideline: The Lifeline Advisor has an internal maximum message size of 128K bytes. If this limit is exceeded, the connection with the external load balancer is closed and an error message is logged stating that the message is too large and was not received.

To diagnose problems when the load balancer connection to the Lifeline Advisor terminates unexpectedly, complete the following steps:

Procedure

1. Determine whether the network administrator has shut down the external load balancer.

2. Determine whether TCP/IP connectivity still exists between the load balancer and the Advisor. For example, from the Advisor system, ping the IP address used by the load balancer.
3. If syslogd was started, check the syslogd file of the Advisor for ERROR or WARNING messages and take the appropriate corrective actions. The syslogd identifier is lladvr for the Advisor. Otherwise, if the Advisor internal trace is being used, dump and check the internal trace for ERROR or WARNING messages and take the appropriate corrective actions.
 - a. If an ERROR message which indicates a send() operation failed with errno = EDC8102I Operation would block is issued, you might have registered too many groups or members from the load balancer. Increase the TCPCONFIG TCPSENDBFRSIZE statement in the TCP/IP profile where the Advisor is active, or register fewer groups and members from the load balancer, and then try restarting the external load balancer. For details about the TCPCONFIG statement, see *z/OS Communications Server: IP Configuration Reference*.
 - b. If ERROR and WARNING level log messages are not enabled, enable them, repeat the operation, and recheck the syslogd file or dump and recheck the internal trace messages again.
4. Check the load balancer for errors.

Diagnosing problems when the Lifeline Agent connection terminates

When the Agent connection to the Advisor terminates, you can use the information in this topic to diagnose problems.

About this task

Guideline: The Advisor and Agent have an internal maximum message size of 128KB. If this limit is exceeded, the connection is closed and an error message is logged stating that the message is too large and was not received.

If the connection between a Lifeline Agent and Advisor terminates unexpectedly, complete the following steps to diagnose problems:

Procedure

1. Determine whether the MVS operator or Linux on System z administrator has shut down this Agent.
2. Determine whether TCP/IP connectivity still exists between the Agent and the Advisor. For example, from the Advisor system, ping the IP address used by the Agent.
3. If syslogd was started, check the syslogd file of the Advisor for ERROR or WARNING messages. The syslogd identifier is lladvr for the Advisor. Otherwise, if internal trace is used, dump and check the internal log for ERROR or WARNING messages. In both cases, take the appropriate corrective actions.
 - a. If an ERROR message which indicates a send() operation failed with errno = EDC8102I Operation would block is issued, you might have registered too many groups or members that belong to the same Agent. Ensure that the TCPCONFIG TCPMAXSENDBFRSIZE and TCPMAXRCVBUFRSIZE statements in the TCP/IP profile are set to 256 KB where the Advisor and z/OS Agent are active, or register fewer groups and members belonging to the Agent, and then try restarting the Agent. For details about the TCPCONFIG statement, see *z/OS Communications Server: IP Configuration*

Reference. For Linux on System z Management Agents, see the documentation for your Linux distribution to determine how to display and set the maximum TCP send and receive buffer sizes.

- b. If ERROR and WARNING level log messages are not enabled, enable them, repeat the operation, and recheck the syslogd file or dump and recheck the internal trace messages again.

Diagnosing problems when the Lifeline Advisor to Advisor connection terminates

When the connection between the peer Lifeline Advisor and this Advisor terminates unexpectedly, you can use the information in this topic to diagnose problems.

About this task

To diagnose problems when the Advisor to Advisor connection terminates, complete the following steps:

Procedure

1. Determine whether the MVS operator has shut down the peer Advisor.
2. Determine whether TCP/IP connectivity still exists between the peer Advisor and the Advisor. For example, from the Advisor system, ping the IP address used by the peer Advisor.
3. If syslogd was started, check the syslogd file of the Advisor for ERROR or WARNING messages and take the appropriate corrective actions. The syslogd identifier is lladvsr for the Advisor. Otherwise, if the Advisor internal trace is being used, dump and check the internal trace for ERROR or WARNING messages and take the appropriate corrective actions.
4. If syslogd was started, check the syslogd file of the peer Advisor for ERROR or WARNING messages and take the appropriate corrective actions. Otherwise, if the Advisor internal trace is being used, dump and check the internal trace for ERROR or WARNING messages and take the appropriate corrective actions.

(z/OS systems) Diagnosing failures of registering with ARM

Rule: This task is only applicable to the Lifeline Agents residing on z/OS systems. When the Lifeline Advisor or Agent fails to register with the Automatic Restart manager (ARM), you can use the information in this topic to diagnose problems.

About this task

To diagnose problems when the Lifeline Advisor or Agent fails to register with ARM, complete the following step:

Procedure

Check the warning-level message written to the syslogd file. This log message is issued when the IXCARM call fails with the return code and the reason code indicated in the log message. See *z/OS MVS Programming: Sysplex Services Reference* for information about interpreting the IXCARM return code and reason code. One of the common causes of failure is the lack of a security profile. See Chapter 3, "Configuring Multi-site Workload Lifeline," on page 43 for information about how to add an ARM security profile for the Advisor or Agent.

Diagnosing problems in the Linux on System z Management Agent

About this task

See “Linux on System z Workloads” on page 13 for information about diagnosing problems with the Linux on System z Management Agent.

Chapter 6. Multi-site Workload Lifeline configuration statements

Multi-site Workload Lifeline configuration statements define the information needed to establish communication with an Advisor and external load balancers, an Advisor and Agents, and between peer Advisors.

Two types of Multi-site Workload Lifeline configuration statements are: Lifeline Advisor configuration statements and Lifeline Agent configuration statements.

General configuration rules for the Multi-site Workload Lifeline are as follows:

- You must specify at least one value for each statement and separate the statement from its values by using one or more blanks.
- You can specify only one statement and one or more values per line.
- Text beyond the specified statement and its values is treated as a comment. If the text does not start with a number sign (#), a warning message is written to the log.
- For statements that contain braces ({ and }), you must specify each brace on a separate line, for example:

```
advisor_id_list
{
  10.10.10.1
  10.10.10.2
}
```

- Any text beyond an opening or closing brace is treated as a comment. If the text does not start with a number sign (#), a warning message is written to the log.
- Text beginning with a number sign (#) is a comment and is ignored. The remainder of the line following the number sign (#) is considered part of the comment.
- For statements with identical labels, a warning message is written to the log, and the last instance of the statement is used.
- When a statement is processed, all of the parameters are examined. Any parameter that is specified incorrectly causes an error. Any inconsistencies between parameters also cause an error.
- Static system symbols can be used for values in a statement. For more information about defining symbols, see *z/OS MVS Initialization and Tuning Reference*.

Lifeline Advisor configuration statements

Use the following configuration statements to configure Lifeline Advisors.

Table 1. Lifeline Advisor Configuration Statements

Configuration statement	Allowed values	Default value	Required or optional	Purpose
advisor_connection_port	1 - 65535	None	Required	Specifies the port on which the primary Advisor listens for connections from a secondary Advisor.
advisor_id_list		None	Optional	Specifies which secondary Advisors are allowed to connect to the primary Advisor.

Table 1. Lifeline Advisor Configuration Statements (continued)

Configuration statement	Allowed values	Default value	Required or optional	Purpose
advisor_takeover_policy	automatic manual	manual	Optional	Specifies the action that a secondary Advisor takes when it determines that a primary Advisor has stopped.
advisor_timeout	10 - 600 (in seconds)	10	Optional	Specifies how often the Advisor checks the availability of the peer Advisor.
agent_connection_port	1 - 65535	None	Required	Specifies the port on which the primary Advisor listens for connections from Agents.
agent_id_list		None	Required	Specifies which Agents are allowed to connect to the primary Advisor.
cross_sysplex_list		None	Required if TCP/IP applications make up at least one workload	Specifies the mapping of second-tier load balancers and their respective site name for a specific workload, and the failure detection interval for the workload.
debug_level	0 - 255	7	Optional	Specifies the level of debug information that is logged.
failure_detection_interval	10 - 600 (in seconds)	60	Optional	Specifies the interval that the primary Advisor waits before determining a site failure.
intermediary_node_list		None	Optional	Specifies the mapping of intermediary nodes and their respective site name for a specific workload.
lb_connection_v4		None	Required if lb_connection_v6 is not specified	Specifies the IPv4 address and port on which the primary Advisor listens for connections from external load balancers. Also specifies the IPv4 address used by the primary Advisor to accept connections from Agents or the secondary Advisor.
lb_connection_v6		None	Required if lb_connection_v4 is not specified	Specifies the IPv6 address and port on which the primary Advisor listens for connections from external load balancers. Also specifies the IPv6 address used by the primary Advisor to accept connections from Agents or the secondary Advisor.
lb_id_list		None	Required if the cross_sysplex_list or sna_application_list is specified	Specifies which external load balancers are allowed to connect to the primary Advisor.
management_guest_list		None	Required if the Advisor communicates with one or more Linux on System z Management Agents	Specifies the Linux on System z guest names and IP addresses that the Linux on System z Management Agents monitor.
mq_manager_list		None	Required if MQ workloads are monitored.	Specifies the MQ queue managers, IBM MQ clusters, and MQ cluster queues that the z/OS Agents monitor.
peer_advisor_id		The IPv4 address on the lb_connection_v4 statement or, if not specified, the IPv6 address on the lb_connection_v6 statement	Optional	Specifies the IP address on the primary Advisor that the secondary Advisor uses to connect to.

Table 1. Lifeline Advisor Configuration Statements (continued)

Configuration statement	Allowed values	Default value	Required or optional	Purpose
query_workload_list		None	Optional	Specifies the Active/Query workload, its associated Active/Standby workload, and the distribution method for the Active/Query workload.
sna_application_list		None	Required if SNA applications make up at least one workload	Specifies the SNA application names and their respective site name for the specific workload.
update_interval	10 - 600 (in seconds)	10	Optional	Specifies how frequently the Agents update the primary Advisor with new information.
workload_switch_policy	automatic manual	manual	Optional	Specifies the action that the primary Advisor takes when it determines either the workload is unavailable on the active site or the active site for a workload is unavailable.

advisor_connection_port

Use the `advisor_connection_port` statement to specify the port that the primary Advisor should listen on for connections from a secondary Advisor.

Syntax

```
►►—advisor_connection_port—host_port—◄◄
```

Parameters

host_port

Specifies which port the primary Advisor listens on for connections from a secondary Advisor. The valid range of port values is 1 - 65535.

Requirement: This port number must match the port number specified in the `advisor_connection_port` configuration statement for the secondary Advisor.

Usage

Restriction: Do not specify the same port that is on the `agent_connection_port`, `lb_connection_v4`, or `lb_connection_v6` configuration statements; otherwise, the secondary Advisor will connect to the primary Advisor's socket that is used for processing Agent connections or load balancer connections.

advisor_id_list

Use the `advisor_id_list` statement to specify which secondary Advisors are allowed to connect to the primary Advisor. This list is also used by the secondary Advisor to select a source IP address when it is connecting to the primary Advisor.

Syntax

```
►►—advisor_id_list—{—advisor_ipaddr—}—◄◄
```

Parameters

advisor_ipaddr

Specifies a list of secondary Advisors that are allowed to connect to the primary Advisor. The list consists of one or more IP address entries, each specified on a separate line and enclosed in braces. Each of these entries represents the IP address of a given secondary Advisor.

Requirement: Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed. Each *advisor_ipaddr* entry must be specified on one line. The *advisor_ipaddr* entry cannot be continued to a subsequent line.

Usage

- You can dynamically update this statement to add or delete IP address entries by using the `MODIFY procname, REFRESH` command. If the secondary Advisor is connected, the entry for that Advisor cannot be deleted.
- For high availability, use a static VIPA as the local IP address to tolerate individual link outages on the z/OS system.

Requirements:

- Do not use VIPARANGE dynamic VIPA.
- You must configure a local IP address for the primary Advisor in this list because it is used during initialization by the primary Advisor to determine whether another primary Advisor is currently active.

Configuring the IP address for the peer Advisor in this list is optional. However, if you do not configure an Application Transparent TLS (AT-TLS) policy to permit Advisor to Advisor communication, the statement is required; otherwise, the peer Advisor cannot connect to this Advisor.

You can configure peer Advisor IP addresses in this statement as well as have an AT-TLS policy enabled, such that a peer Advisor connection is authenticated by using this statement, while a different peer Advisor connection is authenticated by using AT-TLS.

advisor_takeover_policy

Use the *advisor_takeover_policy* statement to specify the action that the secondary Advisor should take when it determines that the primary Advisor has stopped.

Syntax



Parameters

manual

Configures the secondary Advisor to wait for a `MODIFY procname, TAKEOVER` command before taking over as the primary Advisor. The default value is `manual`.

automatic

Configures the secondary Advisor to perform a takeover to become the primary Advisor.

Usage


If automation is already in place to detect a primary Advisor failure, specifying manual will allow automation procedures to handle the takeover of the primary Advisor.

Rule: Keep the configuration files of the primary and secondary Advisor in sync; otherwise, the automatic takeover will not occur even if the automatic parameter is specified.

advisor_timeout

Use the `advisor_timeout` statement to specify how often the Advisor checks the availability of the peer Advisor.

Syntax



Parameters

interval

Specifies how frequently, in seconds, the Advisor checks the availability of the peer Advisor.

The valid range is 10 – 600 (in seconds). The default value is 10 seconds.

The configuration statement is used by the primary Advisor for the following reasons:

- During initialization to determine how long it waits before concluding that another primary Advisor is not active before becoming the primary Advisor.
- To determine how long to wait for the secondary Advisor to signal that it is available before declaring that the secondary Advisor is no longer active.

The configuration statement is used by the secondary Advisor for the following reason:

- To determine how long to wait for the primary Advisor to signal that it is available before declaring that the primary Advisor is no longer active.

agent_connection_port

Use the `agent_connection_port` statement to specify the port on which the primary Advisor should listen for connections from Agents.

Syntax



Parameters

host_port

Specifies the port on which the primary Advisor listens for connections from Agents. The valid range of port values is 1 - 65535.

Requirement: This port number must match the port number specified in the `advisor_id` configuration statement for the Agents.

Usage

Restriction: Do not specify the same port that is on the `advisor_connection_port`, `lb_connection_v4`, or `lb_connection_v6` configuration statement; otherwise, the Agent will connect to the primary Advisor's socket that is used for processing secondary Advisor connections or load balancer connections.

agent_id_list

Use the `agent_id_list` statement to specify which Agents are allowed to connect to the primary Advisor.

Syntax

```
agent_id_list { agent_ipaddr..agent_port }
```

Parameters

agent_ipaddr..agent_port

Specifies a list of Agents that are allowed to connect to the primary Advisor. The list consists of one or more IP address and port pairs, each specified on a separate line and enclosed in braces. Each of these pairs represents the IP address and port of a given Agent. The valid range of port values is 1 - 65535. Do not put any spaces between the IP address, the two periods (..), and the port.

Requirements:

- Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed. Each `agent_ipaddr..agent_port` pair must be specified on one line; it cannot be continued to a subsequent line.
- Each IP address and port pair must match the IP address and port in the `host_connection` configuration statement for an Agent.
- Each IP address in the list can be specified only once.

Usage

If an Agent attempts to connect to the Advisor and the Agent's source IP address and port are not represented in this list, the connection is refused.

Guideline: You can configure this statement as well as have an AT-TLS policy enabled, such that some Agent connections are authenticated only by using this statement, while other Agent connections are also authenticated by using AT-TLS. Linux on System z Management Agents cannot be authenticated by using AT-TLS.

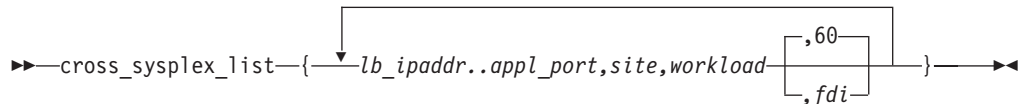
You can dynamically update this statement to add or delete IP address and port pair entries by using the `MODIFY procname,REFRESH` command. If an Agent is connected, the entry for that Agent cannot be deleted.

cross_sysplex_list

Use the `cross_sysplex_list` statement to specify the IP address and site name of each second-tier load balancer, the port of the application group used by the workload, the workload name, and the failure detection interval for the IP workload.

If a first-tier load balancer attempts to connect to the primary Advisor and register group members that are not represented in this list or in the `intermediary_node_list` configuration statement, the connection is refused.

Syntax



Parameters

lb_ipaddr.appl_port,site,workload,fdi

Specifies a list of second-tier load balancers, the site the load balancer resides on, and the workload the load balancer is responsible for. The list consists of one or more IP address, port, site name, workload name, and workload failure detection interval 5-tuples, each specified on a separate line and enclosed in braces. Each of these 5-tuples represents the second-tier load balancer and the mapping of the site name and workload name for that workload. The valid range of port values is 1 - 65535.

Requirements:

- Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed.
- Each *lb_ipaddr.appl_port,site,workload,fdi* 5-tuple must be specified on one line; it cannot be continued to a subsequent line.
- Do not put any spaces between the IP address, the two periods (.), and the port. Do not put spaces between any parameter and the preceding or following comma (,).

Each entry can represent either an Active/Standby workload or an Active/Query workload.

lb_ipaddr.appl_port

Specifies the IP address and port of second-tier load balancers.

Rule: The same IP address is not allowed to be configured in both the `cross_sysplex_list` and the `intermediary_node_list` statement.

site

Specifies the site where the second-tier load balancer resides. The site name corresponds to the MVS sysplex name for that site.

workload

Specifies the workload the load balancer is responsible for. The workload name can consist of alphanumeric characters as well as the @, #, \$, -, and _ characters. The first character must be an alphabetic character. The length of the workload name cannot exceed 64 characters.

Requirement: Each workload in the list must specify the same two unique site names. All workloads must use the same two names.

fdi

Specifies the workload failure detection interval. The failure detection interval specifies how long the Advisor waits to determine that a workload is no longer available on a site. Although an Active/Standby workload might consist of multiple load balancer application groups, the failure detection interval starts when no member of one group is available for that site. For an Active/Query workload or an Active/Standby workload consisting of a single load balancer application group, the failure detection interval starts when no member of the group is available for that site. The failure detection interval resets if at least one of the members of the failing load balancer application group becomes available for that site. The valid range is 10 - 600 in seconds. The default interval is 60 seconds. The failure detection interval value cannot be less than the value that is specified on the `update_interval` statement.

Rule: All entries for a workload, whether in the `cross_sysplex_list`, the `sna_application_list` or the `mq_manager_list`, must specify the same `fdi` value; otherwise a terminating error occurs.

Usage

- The second-tier load balancer IP address can be either the IP address of the external load balancer residing in the site, the IP address of the intermediary node, or the dynamic virtual IP address (VIPA) that is distributed by the z/OS sysplex distributor residing on a system image in the site.
- If an Active/Standby workload consists of multiple load balancer application groups, each load balancer application group must be specified with its own 5-tuple entry, and must be unique with other 5-tuple entries that are defined for the workload by specifying a different second-tier load balancer IP address or a different port.
- The `cross_sysplex_list` statement is required when either of the following conditions occurs:
 - The `intermediary_node_list` and the `mq_manager_list` statements are not configured.
 - The `intermediary_node_list` is configured and any of its entries specify an application type of IP or use an application type of IP by default.
- You can dynamically update this statement by using the `MODIFY procname, REFRESH` command.

Rules:

- A new 5-tuple entry can be added to an existing workload or for a new workload.
- An existing 5-tuple entry for a workload cannot be deleted.
- The workload name and failure detection interval are the only values that can be modified on an existing 5-tuple entry.
- The workload name can be modified only if the workload is quiesced.
- If the workload is also defined in an `intermediary_node_list` entry, `sna_application_list` entry, `mq_manager_list` entry, and/or `management_guest_list` entry, all entries in each list that are using this workload name must be modified.

debug_level

Use the `debug_level` statement to specify the level of debug information that is logged.

Syntax



Parameters

level

Specifies the debug level. If internal trace is used, all trace messages are written to internal trace, but only the trace messages at the ERROR, WARNING, and EVENT level are written to syslogd. If internal trace is not used, all trace messages are written to syslogd. The level value represents a particular debug level or combination of debug levels. See “Debug levels” on page 67 for the possible values. The default level is 7 (1 + 2 + 4 = 7), which means that all error, warning, and event messages are captured. The valid range is 0 - 255.

Usage

- You can dynamically update this statement by using the `MODIFY procname,DEBUG` command.
- To log a combination of debug levels, add the debug level numbers.

Restriction: Do not customize this statement unless you are directed to do so by an IBM Service representative. Adding additional types of trace data can cause the amount of data captured to become voluminous. Reducing the amount of trace data from the default might make diagnosing a problem more difficult.

failure_detection_interval

Use the `failure_detection_interval` statement to specify how long the primary Advisor waits for determining that a site is no longer available.


For site failures, the failure detection interval starts when both of the following conditions are true:

- All Agents on that site are not connected to the primary Advisor
- All LPARs on that site are reported to the primary Advisor, via BCPii services, as not available

The failure detection interval resets when one of the following conditions is true:

- At least one Agent on that site connects to the primary Advisor
- At least one LPAR on that site is reported as available, via BCPii services, to the primary Advisor

Syntax



Parameters

interval

Specifies how long the primary Advisor waits for an indication that a site is available. The valid range is 10 - 600 in seconds. The default interval is 60 seconds.

Requirement: The *interval* value cannot be less than the value specified on the *update_interval* statement.

Usage

You can dynamically update this statement by using the MODIFY *procname*,REFRESH command.

intermediary_node_list

Use the *intermediary_node_list* statement to specify the IP address and site name of each intermediary node, the ports of the load balancer application group used by the workload, and the workload name.

If a first-tier load balancer connects to the primary Advisor and attempts to register group members that are not represented in this list or in the *cross_sysplex_list* configuration statement, the registration is refused and the connection is closed.

Syntax

```
intermediary_node_list { lb_ipaddr..appl_port,site,workload [,IP [,application_type] }
```

Parameters

lb_ipaddr..appl_port,site,workload,application_type

Specifies a list of intermediary nodes, the site that the intermediary node resides on, the workload that the intermediary node is responsible for, and the type of application. The list consists of one or more IP address, port, site name, workload name, and application type 5-tuples, each specified on a separate line and enclosed in braces. Each of these 5-tuples represents the intermediary node and the mapping of the site name and workload name for that workload. The valid range of port values is 1 - 65535.

Requirements:

- Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed.
- Each *lb_ipaddr..appl_port,site,workload,application_type* 5-tuple must be specified on one line; it cannot be continued to a subsequent line.
- Do not put any spaces between the IP address, the two periods (..), and the port. Do not put spaces between any parameter and the preceding or following comma (,).

lb_ipaddr..appl_port

Specifies the IP address and port of the intermediary node.

Rule: The same IP address and port is not allowed to be configured in both the `cross_sysplex_list` and the `intermediary_node_list` statement.

site

Specifies the site where the intermediary node resides. The site name corresponds to the MVS sysplex name for that site or the `site_name` that is specified in the Linux on System z Management Agent's configuration file.

workload

Specifies the workload that the intermediary node is responsible for. The workload name can consist of alphanumeric characters as well as the `@`, `#`, `$`, `-`, and `_` characters. The first character must be an alphabetic character. The length of the workload name must be within 64 characters.

Requirements:

- Each workload in the list must specify the same two unique site names. All workloads must use the same two names.
- If an Active/Standby workload consists of multiple load balancer application groups, each load balancer application group must be specified with its own 5-tuple entry. Each group's 5-tuple entry must also be unique with other 5-tuple entries that are defined for the workload by specifying a different intermediary node IP address or a different port.

application type

Specifies the type of application that processes the workload distributed by the intermediary node.

IP TCP/IP server applications process the workload. This is the default value.

SNA SNA applications process the workload.

Usage

- The intermediary node IP address resides in the site. It uses its own recommendations rather than receive distribution recommendations from the Multi-site Workload Lifeline.
- The members of a load balancer application group must map to either the IP intermediary node entries or the SNA intermediary node entries. A group cannot be used for both IP and SNA applications.
- You can dynamically update this statement by using the `MODIFY procname,REFRESH` command.

Rules:

- A new 5-tuple entry can be added to an existing workload or for a new workload.
- An existing 5-tuple entry for a workload cannot be deleted.
- The workload name is the only value that can be modified on an existing 5-tuple entry.
- The workload name can be modified only if the workload is quiesced.
- If the workload is also defined in a `cross_sysplex_list` entry, `sna_application_list` entry, `mq_management_list` entry, and/or `management_guest_list` entry, all entries in each list that are using this workload name must be modified.

Requirement: A load balancer application group that maps to entries in this list must have corresponding entries in the `cross_sysplex_list` configuration statement if IP is configured or corresponding entries in the `sna_application_list` if SNA is configured.

- When IP is configured:
 - If a workload has at least one intermediary IP address that matches an IP address configured in one of the z/OS Agent's systems, that is, a monitored internal intermediary node configuration, then each `intermediary_node_list` entry for that workload must have a corresponding `cross_sysplex_list` entry with a matching workload name, site, and port.
 - If a workload has at least one intermediary IP address that matches an IP address configured in one of the Linux on System z guests, that is, a monitored external intermediary node configuration, then each `intermediary_node_list` entry for that workload must have a corresponding `cross_sysplex_list` entry with a matching workload name and site.
 - If a workload has no intermediary IP address that matches an IP address configured in the site, that is, an unmonitored external intermediary node configuration, then each `intermediary_node_list` entry for that workload must have a corresponding `cross_sysplex_list` entry with a matching workload name and site.

The corresponding `cross_sysplex_list` entry specifies the distributable DVIPA to be used to monitor the server application for the workload.


- When SNA is configured:
 - Each `intermediary_node_list` entry for that workload must have one or more corresponding `sna_application_list` entries with the same workload and site. The application names of these `sna_application_list` entries represent the SNA applications that process the workload on that site.

lb_connection_v4

Use the `lb_connection_v4` statement to specify the IPv4 address and port that the primary Advisor should listen on for connections from external load balancers. The IP address in this statement is also used by the primary Advisor to accept connections from Agents or the secondary Advisor.

This statement is optional. However, if neither a `lb_connection_v4` nor `lb_connection_v6` statement is present in the configuration file, a terminating error occurs.

Syntax



Parameters

host_ipaddr..host_port

Specifies which IPv4 address, and optionally, the port, on which the primary Advisor listens for IPv4 connections from an external load balancer. This IP address and port must be coordinated on any load balancers that will connect to this Advisor. The port is optional. The default value is 3860. The valid range of port values is 1 - 65535.

Rule: If you specify the port, do not use spaces between the IP address, the two periods (.), and the port.

Requirement: Only a complete IPv4 address is accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) is not allowed.

Usage

- If you specify this statement and also the `cross_sysplex_list` or `sna_application_list` statement, you must specify at least one IPv4 address in the `lb_id_list` statement; otherwise, a terminating error occurs.
- Define the IPv4 address as a VIPARANGE dynamic VIPA. This allows external load balancers and Agents to reconnect to this IP address when the secondary Advisor takes over as the primary Advisor and moves the dynamic VIPA to the host on which the secondary Advisor is active.
- You can dynamically update this statement to add, delete, or modify the IP address and port pair entry on the secondary Advisor by using the `MODIFY procname,REFRESH` command.

Rules:

- If the secondary Advisor is using this entry to connect to the primary Advisor, you cannot delete or modify the entry.
- You cannot dynamically update this statement on the primary Advisor.

Restriction: Do not specify the same port that is on the `advisor_connection_port` or `agent_connection_port` configuration statements; otherwise, the Agents and secondary Advisor will connect to the primary Advisor's socket that is used for processing load balancer connections.

Note: If no `lb_id_list` statement is configured, the `lb_connection_v4` (or `lb_connection_v6`) statement must still be specified as this statement determines the IP address used by Agents and the secondary Advisor to connect to the primary Advisor.

lb_connection_v6

Use the `lb_connection_v6` statement to specify the IPv6 address and port that the primary Advisor should listen on for connections from external load balancers. The IP address in this statement is also used by the primary Advisor to accept connections from Agents or the secondary Advisor.

This statement is optional. However, if neither a `lb_connection_v4` nor `lb_connection_v6` statement is present in the configuration file, a terminating error occurs.

Syntax

►► `lb_connection_v6` `host_ipaddr` ..3860
..host_port ►►

Parameters

`host_ipaddr..host_port`

Specifies which IPv6 address, and optionally the port, that the primary Advisor

listens on for IPv6 connections from an external load balancer. This IP address and port must be coordinated on any load balancers that will connect to this Advisor. The port is optional. The default value is 3860. The valid range of port values is 1 - 65535.

Rule: If you specify the port, do not use spaces between the IP address, the two periods (.), and the port.

Requirement: Only a complete IPv6 address is accepted. Host names, prefixes, or subnets are not allowed. The IPv6 unspecified address (::) is not allowed.

Usage

- If you specify this statement and also the `cross_sysplex_list` or `sna_application_list` statement, you must specify at least one IPv6 address in the `lb_id_list` statement; otherwise, a terminating error occurs.
- Define the IPv6 address as a VIPARANGE dynamic VIPA. This allows external load balancers and Agents to reconnect to this IP address when the secondary Advisor takes over as the primary Advisor and moves the dynamic VIPA to the host on which the secondary Advisor is active.
- You can dynamically update this statement to add, delete, or modify the IP address and port pair entry on the secondary Advisor by using the `MODIFY procname,REFRESH` command.

Rules:

- If the secondary Advisor is using this entry to connect to the primary Advisor, you cannot delete or modify the entry.
- You cannot dynamically update this statement on the primary Advisor.

Restriction: Do not specify the same port that is on the `advisor_connection_port` or `agent_connection_port` configuration statements; otherwise, the Agents and secondary Advisor will connect to the primary Advisor's socket that is used for processing load balancer connections.

Note: If no `lb_id_list` statement is configured, the `lb_connection_v6` (or `lb_connection_v4`) statement must still be specified as this statement determines the IP address used by Agents and the secondary Advisor to connect to the primary Advisor.

lb_id_list

Use the `lb_id_list` statement to specify which external load balancers are allowed to connect to the primary Advisor.

Syntax

```
lb_id_list { lb_ipaddr }
```

Parameters

lb_ipaddr

Specifies a list of external load balancers that are allowed to connect to the primary Advisor. The list consists of one or more IP address entries, each

specified on a separate line and enclosed in braces. Each of these entries represents the IP address of a given load balancer.

Requirement: Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed. You must specify each `lb_ipaddr` entry on one line; it cannot be continued to a subsequent line.

Usage

- If an external load balancer attempts to connect to the primary Advisor and the load balancer's source IP address is not represented in this list, the connection is refused.

Guideline: You can configure this statement as well as have an AT-TLS policy enabled, such that some external load balancer connections are authenticated only by using this statement, while other load balancer connections are also authenticated by using AT-TLS.

- You can dynamically update this statement to add or delete IP address entries by using the `MODIFY procname, REFRESH` command. If a load balancer is connected, you cannot delete the entry for that load balancer.
- You must configure all external load balancers, both first-tier and second-tier, in this list. If a z/OS sysplex distributor is acting as a second-tier load balancer, do not configure it in this list.

Rules:

- If you specify an IPv4 address, you must specify the `lb_connection_v4` statement or a terminating error occurs. Similarly, if you specify an IPv6 address, you must specify the `lb_connection_v6` statement or a terminating error occurs.
- This statement is required if the `cross_sysplex_list` or `sna_application_list` statement is specified.

management_guest_list

Use the `management_guest_list` to specify the IP address, the names of the Linux on System z guests that are used by the workload, and the workload name.

Syntax

►► `management_guest_list` { `guest_name, ipaddr, workload` } ◄◄

Parameters

guest_name, ipaddr, workload

Specifies a list of Linux on System z guest names, IP addresses of the guest, and the workload that the guest belongs to. The list consists of one or more guest name, IP address, and workload name 3-tuples, each specified on a separate line and enclosed in braces.

Requirements:

- Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed.

- Each *guest_name,ipaddr,workload* 3-tuple entry must be specified on one line; it cannot be continued to a subsequent line.
- Do not put spaces between any parameter and the preceding or following comma (,).

guest_name

Specifies the name of the Linux on System z guest that owns the IP address. The guest name can consist of alphanumeric characters as well as the @, #, \$, +, -, and : characters. Contact your z/VM Administrator to understand which Linux on System z guests participate in Lifeline workloads and the IP addresses that are used to reach the server applications on those guests.

ipaddr

Specifies the IP address of a Linux on System z guest to be monitored by the Linux on System z Management Agent.

workload

Specifies the workload that the Linux on System z guest belongs to. The workload name can consist of alphanumeric characters as well as the @, #, \$, -, and _ characters. The first character must be an alphabetic character. The length of the workload name cannot exceed 64 characters.

Usage

- The entries in the *management_guest_list* statement are used to monitor the server applications that run on Linux on System z guests.

Rules:

- Each *ipaddr,workload* pair must also be configured in the *intermediary_node_list*.
 - *management_guest_list* entries that map to *intermediary_node_list* entries must exist on both sites for a workload.
- The workload must be configured in the *cross_sysplex_list*, or the *sna_application_list*, or both lists.
 - Each *management_guest_list* entry must map to one or more *intermediary_node_list* entries with the same workload and site.
 - If the *management_guest_list* entries distribute connections to TCP/IP server applications, mappable entries for the workload in the *cross_sysplex_list* must exist. The *cross_sysplex_list* entries are required to monitor these applications.
 - If the *management_guest_list* entries distribute connections to SNA applications, mappable entries for the workload in the *sna_application_list* must exist. The *sna_application_list* entries are required to monitor these applications.
- A *guest_name* must be associated with only one workload; otherwise, a terminating error occurs.
- Each IP address in the list can be specified only once.
- You can dynamically update this statement by using the *MODIFY procname,REFRESH* command.

Rules:

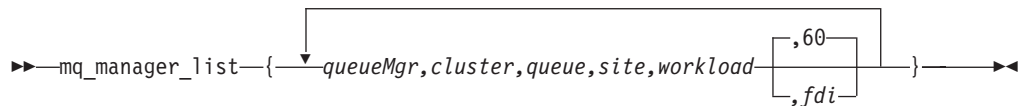
- A new 3-tuple entry can be added to an existing workload or for a new workload.
- An existing 3-tuple entry for a workload cannot be deleted.

- The workload name and guest name are the only values that can be modified on an existing 3-tuple entry.
- The workload name can be modified only if the workload is quiesced.
- If the workload is also defined in a `cross_sysplex_list` entry, `intermediary_node_list` entry, and/or `sna_application_list` entry, all entries in each list that are using this workload name must be modified.

mq_manager_list

Use the `mq_manager_list` statement to specify the MQ queue manager name, the IBM MQ clusters and MQ cluster queues it hosts, the workload associated with each IBM MQ cluster, the site where the MQ queue manager resides, and the failure detection interval for the workload.

Syntax



Parameters

queueMgr, cluster, queue, site, workload, fdi

Specifies a list of the MQ queue managers, the IBM MQ clusters and MQ cluster queues configured on the MQ queue manager, the site where the MQ queue manager resides, and the workload associated with the IBM MQ cluster. The list consists of one or more blank-delimited MQ queue manager name, cluster name, queue name, site name, workload name, and workload failure detection interval 6-tuples, each specified on a separate line and enclosed in braces. Each of these 6-tuples represents the MQ queue manager and the mapping of the cluster name, queue names, site name, and workload name for that workload.

Requirements:

- Each `queueMgr, cluster, queue, site, workload, fdi` 6-tuple must be specified on one line; it cannot be continued to a subsequent line.
- Do not put spaces between any parameter and the preceding or following comma (,).

queueMgr

Specifies the name of the MQ queue manager. The queue manager name can consist of alphanumeric characters only. The length of the queue manager name cannot exceed 4 characters. Only MQ queue managers that reside on z/OS systems should be configured.

cluster

Specifies the name of an IBM MQ cluster that the MQ queue manager belongs to. The cluster name can consist of alphanumeric characters as well as the `_`, `.`, `/`, and `%` characters. The length of the cluster name cannot exceed 48 characters. See *WebSphere MQ: Installing IBM MQ* for information about rules for cluster names.

queue

Specifies the name of an MQ cluster queue that is defined for the IBM MQ cluster. The queue name can consist of alphanumeric characters as well as the

_, ., /, and % characters. The length of the queue name cannot exceed 48 characters. See *WebSphere MQ: Installing IBM MQ* for information about rules for queue names.

site

Specifies the site where the MQ queue manager resides. The site name corresponds to the MVS sysplex name for that site.

workload

Specifies the workload that this cluster belongs to.

The workload name can consist of alphanumeric characters as well as the @, #, \$, -, and _ characters. The first character must be an alphabetic character. The length of the workload name cannot exceed 64 characters.

Rule: A workload cannot span multiple clusters and a cluster must define the queues for only one workload. Therefore, each workload name should have a one-to-one correspondence with a cluster name:

- A cluster name must always be configured with the same workload name.
- A workload name must always be configured with the same cluster name.

Requirement: Each workload in the list must specify the same two unique site names. All workloads must use the same two names.

fdi

Specifies the workload failure detection interval. The failure detection interval specifies how long the Advisor waits to determine that a workload is no longer available on a site. Although a workload's IBM MQ cluster might consist of multiple MQ cluster queues, the failure detection interval starts when one of the MQ cluster queues is not available on any MQ queue manager for that site. The failure detection interval resets when all MQ cluster queues become available for that site. The valid range is 10 - 600 in seconds. The default interval is 60 seconds. The failure detection interval value cannot be less than the value that is specified on the `update_interval` statement.

Rule: All entries for a workload, whether in the `cross_sysplex_list`, the `sna_application_list` or the `mq_manager_list`, must specify the same `fdi` value; otherwise a terminating error occurs.

Usage

- The entries in this list are used to monitor MQ queue managers and the status of their locally defined MQ cluster queues. Each IBM MQ cluster maps to a workload. Only MQ queue managers residing on z/OS systems are monitored.

Rules:

- A given MQ queue manager name can be defined to only one site. Therefore, every entry for that MQ queue manager must specify the same site.
- Only one entry with the same queue manager name, cluster name, and queue name can be configured.
- An MQ cluster's queue must be defined in both sites. Therefore, at least two entries with the same cluster name, queue name and different site name must be configured.
- The cluster name can be defined to only one workload.
- Multiple instances of a queue name must be configured in the same IBM MQ cluster.
- The workload must be an Active/Standby workload.

- You can dynamically update this statement by using the `MODIFY procname,REFRESH` command.

Rules:

- A new 6-tuple entry can be added to an existing workload or a new workload.
- An existing 6-tuple entry cannot be deleted.
- The workload name and failure detection interval are the only values that can be modified on an existing 6-tuple entry.
- The workload name can be modified only if the workload is quiesced.
- If the workload is also defined in a `cross_sysplex_list` entry, `intermediary_node_list` entry, `sna_application_list` entry, and/or `management_guest_list` entry, all entries in each list that are using this workload name must be modified.

peer_advisor_id

Use the `peer_advisor_id` statement to specify the IP address of the primary Advisor that this secondary Advisor uses to communicate.

This statement is optional. If you do not specify the statement, it defaults to either the IPv4 address specified on the `lb_connection_v4` statement, if configured, or the IPv6 address specified on the `lb_connection_v6` statement.

Syntax

►►—`peer_advisor_id—advisor_ipaddr`—————►►

Parameters

advisor_ipaddr

Specifies the IPv4 or IPv6 address which the Advisor, when acting as the secondary Advisor, uses as the destination IP address to connect to the primary Advisor.

Requirements:

- Only a complete IPv4 or IPv6 address is accepted. Hostnames, prefixes, or subnets are not allowed. The IPv4 `INADDR_ANY` address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed.
- The IP address must match either the IPv4 address specified on the `lb_connection_v4` configuration statement or the IPv6 address specified on the `lb_connection_v6` configuration statement for the primary Advisor.

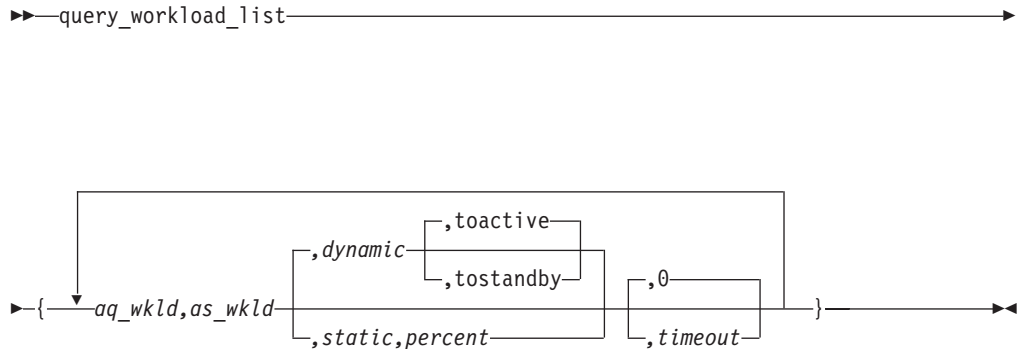
Usage

You can dynamically update this statement only on the primary Advisor by using the `MODIFY procname,REFRESH` command. You cannot dynamically update this statement on the secondary Advisor.

query_workload_list

Use the `query_workload_list` statement to specify the associated Active/Standby workload for an Active/Query workload, the distribution type and options for this Active/Query workload, and the lifetime for persistent connections for this workload.

Syntax



Parameters

aq_wkld, as_wkld, disttype, distoption, timeout

Specifies the Active/Query workload, its associated Active/Standby workload, the distribution type, the distribution-specific option, and the timeout value before a persistent connection for the Active/Query workload is reset. The list consists of one or more Active/Query workload name, associated Active/Standby workload name, distribution type, distribution option, and timeout 5-tuples, each specified on a separate line and enclosed in braces.

Requirements:

- Each *aq_wkld, as_wkld, disttype, distoption, timeout* 5-tuple must be specified on one line; it cannot be continued to a subsequent line.
- Do not put spaces between any parameter and the preceding or following comma (,).

aq_wkld, as_wkld

Specifies the Active/Query workload and its associated Active/Standby workload. The workload names can consist of alphanumeric characters as well as the @, #, \$, -, and _ characters. The first character must be an alphabetic character. The length of each workload name cannot exceed 64 characters.

Requirements:

- The Active/Standby workload must also be defined in at least one of the following statements:
 - `cross_sysplex_list`
 - `mq_manager_list`
 - `sna_application_list`
- The Active/Query workload must also be defined in the `cross_sysplex_list` statement.

Rules:

- An Active/Query workload can be associated with only a single Active/Standby workload.
- An Active/Standby workload can have, at most, two Active/Query workloads associated with it.
- A SNA workload cannot be configured as an Active/Query workload.
- An MQ cluster workload cannot be configured as an Active/Query workload.

disttype/distoption

Specifies the distribution type and the distribution-specific option. The valid Active/Query workload distribution types are either dynamic or static, with the default being dynamic.

dynamic

For the dynamic distribution type, the possible distribution options are **toactive** and **tostandby**. The option **toactive** indicates to gradually favor the same site where the associated Active/Standby workload is active as database replication latency increases. The option **tostandby** indicates to gradually favor the same site where the associated Active/Standby workload is standby as database replication latency increases.

static

For the static distribution type, the distribution option is the percentage, in the range 1 - 99, of new workload connections that are distributed to the active site with the remaining new workload connections that are distributed to the standby site.

timeout

Specifies the timeout value. The timeout value is the elapsed time that an Active/Query workload connection remains established before the connection is reset. The valid range of timeout values, in seconds, is 0 - 3600. The default value is 0, which indicates that persistent Active/Query workload connections are never reset.

Restriction: This value is ignored for Active/Query workload connections to Linux on System z server applications. A timeout value of 0 is used for these workload connections, which indicates that the connections will never be reset.

Usage

You can dynamically update this statement to add, delete, or modify 5-tuple entries for an existing or a new Active/Query workload by using the `MODIFY procname,REFRESH` command.


Rules:

- A new 5-tuple entry can be added or an existing 5-tuple entry can be deleted only if both the Active/Standby and Active/Query workloads have never been active, or if the workloads are deactivated.
- The distribution type for an existing 5-tuple entry can be modified only if the Active/Query workload is quiesced.
- The distribution option and timeout value for an existing 5-tuple entry can be modified without quiescing the Active/Standby and Active/Query workloads.

sna_application_list

Use the `sna_application_list` statement to specify the SNA application name, the site name, the workload name, and the failure detection interval for this workload.

Syntax



Parameters

applname,site,workload,fdi

Specifies a list of SNA applications, the site where each application resides, and the workload that the SNA application belongs to. The list consists of one or more SNA application name, site name, workload name and workload failure detection interval 4-tuples, each specified on a separate line and enclosed by braces.

Requirements:

- Each *applname,site,workload,fdi* 4-tuple entry must be specified on one line; it cannot be continued to a subsequent line.
- Do not put spaces between any parameter and the preceding or following comma (,).

applname

Specifies the SNA application name that is configured on a VTAM APPL definition statement. You must use a network-qualified name by using the form *netid.applname*.

- The *netid* value can consist of alphanumeric characters as well as the @, #, and \$ characters. The first character must be an alphabetic character or the @, #, or \$ character.
- The *applname* value can consist of alphanumeric characters as well as the @, #, \$, *, and ? characters. The first character must be an alphabetic character or the @, #, \$, or ? character.

When the ? or * characters are used, *applname* is a model application program name. Dynamic applications that are created by using the model application name definition will be automatically monitored. For more information about dynamic applications, see “SNA workloads” on page 22. The length of either the *netid* or *applname* value cannot exceed 8 characters.

site

Specifies the site where the SNA application resides. The site name corresponds to the MVS sysplex name for that site.

workload

Specifies the workload that the SNA application belongs to. The workload name can consist of alphanumeric characters as well as the @, #, \$, -, and _ characters. The first character must be an alphabetic character. The length of the workload name cannot exceed 64 characters.

fdi

Specifies the workload failure detection interval. The failure detection interval specifies how long the Advisor waits to determine that a workload is no longer available on a site. Although an Active/Standby workload might consist of

multiple load balancer application groups, the failure detection interval starts when no member of one group is available for that site. The failure detection interval resets if at least one of the members of the failing load balancer application group becomes available for that site. The valid range is 10 - 600 in seconds. The default interval is 60 seconds. The failure detection interval value cannot be less than the value that is specified on the `update_interval` statement.

Rule: All entries for a workload, whether in the `cross_sysplex_list`, the `sna_application_list` or the `mq_manager_list`, must specify the same `fdi` value; otherwise a terminating error occurs.

Usage

- The entries in the `sna_application_list` statement are used to monitor SNA applications that receive workload connections from intermediary nodes.

Rules:

- Each `sna_application_list` entry must map to one or more intermediary node list entries with the same workload, site, and an application type of SNA.
 - Each application name in the list must be network-qualified. The network-qualified application name must be unique within the site.
 - Each application must be associated with only one workload; otherwise, a terminating error occurs.
 - For a given workload, `sna_application_list` entries must exist for both sites. For this type of workload, the members of a first-tier load balancer application group match the IP addresses and ports in the intermediary node list.
 - If the intermediary nodes for the `sna_application_list` entries reside on Linux on System z guests, the `management_guest_list` contains the list of intermediary nodes that the Linux on System z Management Agent monitors. See the “`management_guest_list`” on page 99 statement for more information.
 - The workload must be an Active/Standby workload.
- You can dynamically update this statement by using the `MODIFY procname, REFRESH` command.

Rules:

- A new `sna_application_list` 4-tuple entry can be added to an existing workload or a new workload.
- An existing 4-tuple entry for a workload cannot be deleted.
- The workload name and failure detection interval are the only values that can be modified on an existing 4-tuple entry.
- If the workload is also defined in a `cross_sysplex_list` entry, `intermediary_node_list` entry, `mq_manager_list` entry, and/or `management_guest_list` entry, all entries in each list that are using this workload name must be modified.

update_interval

Use the `update_interval` statement to specify how often Agents update the primary Advisor with new information. You can also use this statement to determine how often the Advisor queries the status of each of the systems in both sites using BCPii communication.

The lower the update interval is, the more up-to-date the recommendations provided to the external load balancer will be with respect to the target's availability and capability to handle additional workload requests. However, the lower the update interval is, the higher the network traffic and CPU overhead are.

Depending on how the external load balancers are configured and implemented, the `update_interval` statement might also determine how often the external load balancer is updated with data from the primary Advisor. If the external load balancer supports the SASP push flag, and if the flag is set in the load balancer, the primary Advisor sends recommendations to the load balancer at least every update interval. Regardless of what value is set for the update interval, if the push flag is supported and configured in the load balancer, the Advisor might update the load balancer with recommendations before the update interval expires. Events that can trigger an earlier update include starting or stopping a target server application, or the addition or deletion of a member's IP address on the Agent host.

The update interval is a key factor in determining the latency period between when changes occur on the target system and when the load balancer is informed of them. Each Agent updates the primary Advisor with new information every update interval. The Advisor, in turn, updates the load balancer with changes in recommendations every update interval, if the load balancer supports the push flag. If the push flag is supported and configured in the load balancer, the Advisor updates the load balancer with any changes in the availability status of the target system and server application when it discovers such a change from the Agent, instead of waiting for the update interval to expire.

When the external load balancer supports and is configured for the push flag, the maximum amount of latency expected between a change in the member's recommendation and when the load balancer is informed of it, is twice the value of the update interval: one update interval for the Agent to report it to the primary Advisor, and one update interval for the Advisor to report it to the load balancer. However, on average, it takes about one update interval for a change in the target server application's recommendation to reach the load balancer.

Syntax



Parameters

interval

Specifies how frequently Agents update the primary Advisor with new information. At each interval, all Agents update the primary Advisor with the following status information for each registered member, which is the application server that the Agent is responsible for:

- Whether the target server application is active or inactive
- Whether the operator quiesced the target server application
- Various metrics that measure the target system
- The ability of the target server application to handle additional workload requests


The valid range is 10 - 600 in seconds. The default interval is 10 seconds.

Requirement: The *interval* value cannot be greater than the value specified on the *failure_detection_interval* statement or the *fdi* value on an entry on the *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* statement.

workload_switch_policy

Use the *workload_switch_policy* statement to specify the action that the primary Advisor should take when it determines either that a workload is no longer available on the active site or the active site for a workload is unavailable.

Syntax



Parameters

manual

Configures the primary Advisor to wait for a *MODIFY procname,ACTIVATE* command before redistributing the workload to the alternate site. The default policy is manual.

automatic

Configures the primary Advisor to activate the redistribution of the workload to the alternate site.

This setting is ignored for the following conditions:

- One or more Agents that are not connected; there is no way to communicate with disconnected Agents on the active site to deactivate the existing connections.
- One or more first-tier load balancers are not registered; they cannot be notified that the workload was redistributed to the alternate site.
- No server applications are available for the workload on the alternate site; redistributing the workload to the alternate site will lead to another workload failure.

Usage

If automation is already in place to detect a workload or site failure, specifying manual will allow automation procedures to handle the redistribution of the workload to the alternate site.

Lifeline Agent configuration statements

Use the following configuration statements to configure Lifeline Agents.

Table 2. Lifeline Agent Configuration Statements

Configuration Statement	Allowed values	Default Value	Required/Optional	Purpose
advisor_id		None	Required, if advisor_id_list is not specified	Specifies the IP address and port of the primary Advisor that the Agent connects to.
advisor_id_list		None	Required, if advisor_id is not specified	Specifies the list of IP address and port pairs of the primary Advisors that the Agent connects to.

Table 2. Lifeline Agent Configuration Statements (continued)

Configuration Statement	Allowed values	Default Value	Required/Optional	Purpose
debug_level	0 - 255	7	Optional	Specifies the level of debug information that is logged.
host_connection		None	Required	Specifies the local IP address and port that the Agent binds to before connecting to the primary Advisor.
site_name		<ul style="list-style-type: none"> For z/OS Agents, the local sysplex name. For Linux on System z Management Agents, no default. 	<ul style="list-style-type: none"> Required for Linux on System z Management Agents. Optional for z/OS Agents 	Specifies the site name that the Agent must be associated with when communicating with the primary Advisor.

advisor_id

Use the `advisor_id` statement to specify the IP address and port of the primary Advisor that this Agent communicates with.

This statement is optional. However, if neither an `advisor_id_list` nor `advisor_id` statement is present in the configuration file, a terminating error occurs. If this statement is specified, the `advisor_id_list` statement cannot be specified, or a terminating error occurs.

Syntax

►►—`advisor_id—advisor_ipaddr..advisor_port`—◀◀

Parameters

advisor_ipaddr..advisor_port

Specifies the IP address and port of the primary Advisor that this Agent communicates with. The IP address must match either the `lb_connection_v4` or `lb_connection_v6` IP address of the Advisor. The valid range of port values is 1 - 65535. Do not use spaces between the IP address, the two periods (.), and the port.

Requirements:

- This port number must match the port number specified in the `agent_connection_port` configuration statement for the Advisor.
- Only a complete IPv4 or IPv6 address is accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed. The IP address must match either the `lb_connection_v4` or `lb_connection_v6` IP address.

Usage

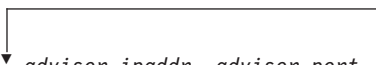
Rule: If you specify an IPv4 address on this statement, you must also specify an IPv4 address on the `host_connection` statement. Similarly, if you specify an IPv6 address on this statement, you must also specify an IPv6 address on the `host_connection` statement.

advisor_id_list

Use the `advisor_id_list` statement to specify the IP address and port pairs of the primary Advisors that this Agent communicates with.

This statement is optional. However, if neither an `advisor_id` nor `advisor_id_list` statement is present in the configuration file, a terminating error occurs. If this statement is specified, the `advisor_id` statement cannot be specified, or a terminating error occurs.

Syntax

```
➤—advisor_id_list—{——advisor_ipaddr..advisor_port—}—➤
```

Parameters

advisor_ipaddr..advisor_port

Specifies the IPv4 or IPv6 addresses and port of the primary Advisors that this Agent communicates with. The Agent loops through all IP addresses until it successfully connects to the primary Advisor using one of these IP addresses. Each IP address must match either the `lb_connection_v4` or `lb_connection_v6` IP address of an Advisor. The valid range of port values is 1 - 65535. Do not put any spaces between the IP address, the two periods (.), and the port.

Requirements:

- Only complete IPv4 or IPv6 addresses are accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed.
- Each `advisor_ipaddr..advisor_port` 2-tuple must be specified on one line; it cannot be continued to a subsequent line.

Usage

Rules:

- The IP addresses specified must be either all IPv4 addresses or all IPv6 addresses. Do not mix IPv4 and IPv6 addresses in the list.
- If you specify IPv4 addresses on this statement, you must also specify an IPv4 address on the `host_connection` statement. Similarly, if you specify IPv6 addresses on this statement, you must also specify an IPv6 address on the `host_connection` statement.
- Do not specify more than two IP address and port pairs.

debug_level

Use the `debug_level` statement to specify the level of debug information that is logged.

Syntax



Parameters

level

Specifies the debug level. If internal trace is being used, all trace messages are written to the internal trace, but only trace messages at the ERROR, WARNING, and EVENT level are written to syslogd. If internal trace is not being used, all trace messages are written to syslogd. The level value represents a particular debug level or combination of debug levels. See “Debug levels” on page 67 for the possible values. The default level is 7 (1 + 2 + 4 = 7), which means that all error, warning, and event messages are captured. The valid range is 0 - 255.

Usage

- You can dynamically update this statement by using the z/OS MODIFY `procname,DEBUG,LEVEL= command` or `aqsllcmd debug,level= command`.
- To log a combination of debug levels, add the specific debug level numbers together.

Restriction: Do not customize this statement unless you are directed to do so by an IBM Service representative. Adding additional types of trace data can cause the amount of data captured to become voluminous. Reducing the amount of trace data from the default might make diagnosing a problem more difficult.

host_connection

Use the `host_connection` statement to specify the local IP address and port that the Agent binds to for communicating with the primary Advisor.

Syntax



Parameters

host_ipaddr..host_port

Use `host_ipaddr..host_port` to specify the IP address and port that this Agent binds to. The valid range of port values is 1 - 65535.

Requirement: Only a complete IPv4 or IPv6 address is accepted. Host names, prefixes, or subnets are not allowed. The IPv4 INADDR_ANY address (0.0.0.0) and the IPv6 unspecified address (::) are not allowed. The IP address and port must match one of the IP address and port pairs specified in the `agent_id_list` configuration statement for the Advisor.

Rule: If you specify an IPv4 address on the `host_connection` statement, you must specify an IPv4 address on the `advisor_id` statement. Similarly, if you specify an IPv6 address on the `host_connection` statement, you must specify an IPv6 address on the `advisor_id` statement.

Usage

Guidelines:

- To ensure high availability, specify the IP addresses in the `agent_id_list` statement for the Advisor and the `host_connection` statement for each Agent as static VIPAs, so that individual link outages on the hosts are tolerated.
- To ensure simplicity and consistency, specify the same port for this statement for each Agent, and reserve this port on each TCP/IP stack where an Agent will be active.
- You can specify an IPv6 address on this statement if the host where the Agent resides is an IPv6-enabled TCP/IP stack, and the Advisor resides on an IPv6-enabled TCP/IP stack.
- For z/OS Agents, you can have an AT-TLS policy enabled in addition to configuring this statement.

site_name

Use the `site_name` statement to specify the site that the Agent is running in.

Syntax

►► `site_name` *site_name* ◀◀

Parameters

site_name

Specifies the 1- through 8-character site name that this Agent is to be associated with when the Agent communicates with the Advisor. The name must match one of the site names in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement for the Agent to successfully connect to the Advisor. The site name can consist of alphanumeric characters as well as the @, #, and \$ characters.

Requirements:

This statement must be specified for the Linux on System z Management Agent. Linux on System z Management Agents do not run in a z/OS sysplex, so the site name must be configured.

This statement is optional for z/OS Agents. By default, the site name for z/OS Agents is the local z/OS sysplex name that the Agent is running in.


Chapter 7. Lifeline Advisor commands

Use the Life Advisor commands to display Lifeline Advisor information and modify Lifeline Advisor state information.

MODIFY *advisor_proc*,DISPlay command: Display Lifeline Advisor information

The MODIFY *advisor_proc*,DISPlay command displays the information related to the Lifeline Advisor.

Format



Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

DISPlay,ADVisor

Displays information about the external load balancers, Agents, and peer Advisor that are connected to this Advisor. When you issue this command on the secondary Advisor with this parameter, information about the load balancer and Agent is not displayed.

DISPlay,ADVisor,DETail

Displays additional information, including the system name and site name where each Agent is and the system name where the peer Advisor is. When you issue this command on the secondary Advisor with this parameter, information about the load balancer and Agent is not displayed.

DISPlay,CONFig

Displays the active configuration for this Advisor.

DISPlay,DEBUg

Displays the debug level in effect for this Advisor.

The debug level displayed is the sum of one or more levels of debug messages being logged. See “Debug levels” on page 67 for more information about the debug levels.

DISPlay, LB

Displays information about the load balancers that are connected to the primary Advisor. You can specify this parameter only on the primary Advisor.

DISPlay, LB, DETai l

Displays additional information, including the register groups and member data for each load balancer. You can specify this parameter only on the primary Advisor.

DISPlay, WORKload

Displays information about the workloads and their current state. You can specify this parameter only from the primary Advisor.

DISPlay, WORKload, DETai l

Displays additional information, including the server applications that make up each workload. You can specify this parameter only from the primary Advisor.

DISPlay, WORKload, STATus

Displays information about the workloads and their current state and status on both sites. You can specify this parameter only from the primary Advisor.

DISPlay, WORKload=workloadname

Displays information about the workload specified by the *workloadname* value and its current state. You can specify this parameter only from the primary Advisor.

Guidelines:

- If **ALL_AS** is specified for the *workloadname* value, information for all Active/Standby workloads specified in the *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* will be displayed.
- If **ALL_AQ** is specified for the *workloadname* value, information for all Active/Query workloads specified in the *cross_sysplex_list* will be displayed.

DISPlay, WORKload=workloadname, DETai l

Displays additional information about the workload specified by the *workloadname* value, including the server applications that make up the workload. You can specify this parameter only from the primary Advisor.

Guidelines:

- If **ALL_AS** is specified for the *workloadname* value, additional information for all Active/Standby workloads specified in the *cross_sysplex*, *mq_manager_list*, or *sna_application_list* will be displayed.
- If **ALL_AQ** is specified for the *workloadname* value, additional information for all Active/Query workloads specified in the *cross_sysplex* will be displayed.

DISPlay, WORKload=workloadname, STATus

Displays information about the workload specified by the *workloadname* value and its current state and status on both sites. You can specify this parameter only from the primary Advisor.

Guidelines:

- If **ALL_AS** is specified for the *workloadname* value, information for all Active/Standby workloads specified in the *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* will be displayed.

- If *ALL_AQ* is specified for the workloadname value, information for all Active/Query workloads specified in the cross_sysplex_list will be displayed.

Examples

The following example shows a sample display when you issue F AQSADV,DISPLAY,ADVISOR on the primary Advisor.

```
F AQSADV,DISPLAY,ADVISOR
AQS0141I ADVISOR SUMMARY
ADVISOR ROLE   : PRIMARY
IPADDR        : 192.10.1.1
LOAD BALANCERS:
IPADDR  : 192.10.1.32          CONNECTED: YES
TIER    : 1
IPADDR  : 192.10.1.64          CONNECTED: YES
TIER    : 2
IPADDR  : 192.10.1.33          CONNECTED: NO
TIER    : 1
IPADDR  : 192.10.1.65          CONNECTED: NO
TIER    : UNKNOWN
AGENTS   :
IPADDR  : 192.10.110.1         CONNECTED: YES
IPADDR  : 192.10.110.2         CONNECTED: YES
IPADDR  : 192.10.110.3         CONNECTED: YES
IPADDR  : 192.20.110.1         CONNECTED: YES
IPADDR  : 192.20.110.2         CONNECTED: NO
IPADDR  : 192.20.110.3         CONNECTED: DOWN
PEER ADVISOR :
IPADDR     : 192.20.130.1
```

The following example shows a sample display when you issue F AQSADV,DISPLAY,ADVISOR on the secondary Advisor.

```
F AQSADV,DISPLAY,ADVISOR
AQS0141I ADVISOR SUMMARY
ADVISOR ROLE   : SECONDARY
IPADDR        : 192.20.130.1
PEER ADVISOR   :
IPADDR         : 192.10.1.1
```

The following information describes the fields in the previous sample display.

Field name

Descriptions

ADVISOR ROLE

The current role and IP address for this Advisor. The IP address can be an IPv4 or an IPv6 address. The Advisor role has the following values:

PRIMARY

Indicates that this Lifeline Advisor is responsible for the following tasks:

- Communicating with the Agents to retrieve information about the systems and server applications that are being monitored
- Communicating with external load balancers to provide routing recommendations for incoming requests for the configured workloads
- Communicating with the secondary Advisor to provide state information for the configured workloads

SECONDARY

Indicates that this Lifeline Advisor is responsible for communicating with the primary Advisor to monitor whether this Advisor needs to take over the primary Advisor responsibilities.

LOAD BALANCERS

The list of external load balancers in this Advisor's configuration, the IP address of each load balancer, whether the specified load balancer is currently connected to the Advisor, and the tier level of each load balancer. The IP address can be an IPv4 or an IPv6 address. The LOAD BALANCERS field is displayed only when you issue the command on the primary Advisor. The tier can be one of the following values:

- 1 Indicates that this load balancer is a first-tier load balancer. The load balancer is responsible for selecting the site that an inbound workload request is routed to.
- 2 Indicates that this load balancer is a second-tier load balancer. The load balancer is responsible for selecting the server application instance in a site that an inbound workload request is routed to.

UNKNOWN

Indicates that the specified load balancer has never connected to the Advisor so its tier type could not be determined.

AGENTS

The list of Lifeline Agents in the Advisor's configuration, the IP address of the Agent, and whether the specified Agent is currently connected to the Advisor. The IP address can be an IPv4 or an IPv6 address. The AGENTS field is displayed only when you issue the command on the primary Advisor.

CONNECTED

Indicates whether the external load balancer or Lifeline Agent is connected to the Advisor. Possible values are:

- YES** Indicates that the external load balancer or Agent is available and connected to the Advisor.
- NO** Indicates that the external load balancer or Agent is not connected to the Advisor.

DOWN

Indicates that the Agent is not connected to the Advisor because the system where the Agent resides is not operational.

PEER ADVISOR

The IP address of the peer Advisor that is connected to this Advisor. The IP address can be an IPv4 or an IPv6 address. If no peer Advisor is connected, the IP address is displayed as NONE.

The following example shows a sample display when you issue F AQSADV,DISPLAY,ADVISOR,DETAIL on the primary Advisor.

```
F AQSADV,DISPLAY,ADVISOR,DETAIL
AQS0142I ADVISOR DETAILS
ADVISOR ROLE   : PRIMARY
IPADDR        : 192.10.1.1
LOAD BALANCERS:
IPADDR   : 192.10.1.32           CONNECTED: YES
TIER     : 1
IPADDR   : 192.10.1.64           CONNECTED: YES
TIER     : 2
```

```

IPADDR : 192.10.1.33          CONNECTED: NO
TIER   : 1
IPADDR : 192.10.1.65          CONNECTED: NO
TIER   : UNKNOWN
AGENTS :
IPADDR : 192.10.110.1        CONNECTED: YES
SYSTEM NAME : SYS1A          SITE : PLEX1
IPADDR : 192.10.110.2        CONNECTED: YES
SYSTEM NAME : SYS1B          SITE : PLEX1
IPADDR : 192.10.110.3        CONNECTED: YES
SYSTEM NAME : LINUXMG1       SITE : PLEX1
IPADDR : 192.20.110.1        CONNECTED: YES
SYSTEM NAME : SYS2A          SITE : PLEX2
IPADDR : 192.20.110.2        CONNECTED: NO
SYSTEM NAME : N/A            SITE : N/A
IPADDR : 192.20.110.3        CONNECTED: DOWN
SYSTEM NAME : LINUXMG2       SITE : PLEX2
PEER ADVISOR :
IPADDR : 192.20.130.1
SYSTEM NAME : CNTL2

```

The following example shows a sample display when you issue F AQSADV,DISPLAY,ADVISOR,DETAIL on the secondary Advisor.

```

F AQSADV,DISPLAY,ADVISOR,DETAIL
AQS0142I ADVISOR DETAILS
ADVISOR ROLE : SECONDARY
IPADDR       : 192.20.130.1
PEER ADVISOR :
IPADDR       : 192.10.1.1
SYSTEM NAME  : CNTL1

```

The following information describes the fields in the previous sample display.

Field name
Descriptions

ADVISOR ROLE

The current role and IP address for this Advisor. The IP address can be an IPv4 or an IPv6 address. The role has the following values:

PRIMARY

- Indicates that this Advisor is responsible for the following tasks:
- Communicating with the Agents to retrieve information about the systems and server applications that are being monitored
 - Communicating with external load balancers to provide routing recommendations for incoming requests for the configured workloads
 - Communicating with the secondary Advisor to provide state information for the configured workloads

SECONDARY

Indicates that this Advisor is responsible for communicating with the primary Advisor to monitor whether this Advisor needs to take over the primary Advisor responsibilities.

LOAD BALANCERS

The list of external load balancers in this Advisor's configuration, the IP address of the load balancer, whether the specified load balancer is currently connected to the Advisor, and the tier level of the load balancer. The IP address can be an IPv4 or an IPv6 address. The LOAD BALANCERS field is only displayed when you issue the command on the primary Advisor. The tier has the following values:

- 1 Indicates that this load balancer is a first-tier load balancer. The load balancer is responsible for selecting the site that an inbound workload request will be routed to.
- 2 Indicates that this load balancer is a second-tier load balancer. The load balancer is responsible for selecting the server application instance within a site that an inbound workload request will be routed to.

UNKNOWN

Indicates that the specified load balancer has never connected to the Advisor so its tier type could not be determined.

AGENTS

The list of Lifeline Agents in this Advisor's configuration to this Advisor, the IP address of the Agent, and whether the specified Agent is currently connected to the Advisor. The IP address can be an IPv4 or an IPv6 address. The AGENTS field is only displayed when you issue the command on the primary Advisor.

SYSTEM NAME

The MVS system name where the z/OS Agent or peer Advisor is active. For Linux on System z Management Agents, the Linux on System z guest name on the Linux on System z Management Agent. If no peer Advisor is connected, this field is not displayed in the PEER ADVISOR section.

SITE

The MVS sysplex name of the system where the connected z/OS Agent exists, or where an unconnected z/OS Agent was previously connected. For Linux on System z Management Agents, the value that is specified on the site_name Agent configuration statement.

CONNECTED

Indicates whether the external load balancer or Lifeline Agent is connected to the Advisor. Possible values are:

YES Indicates that the external load balancer or Agent is available and connected to the Advisor.

NO Indicates that the external load balancer or Agent is not connected to the Advisor.

DOWN

Indicates that the Agent is not connected to the Advisor because the system where the Agent resides is not operational.

PEER ADVISOR

The IP address of the peer Advisor that is connected to this Advisor. The IP address can be an IPv4 or an IPv6 address. If no peer Advisor is connected, the IP address is displayed as NONE.

The following example shows a sample display when you issue F AQSADV,DISPLAY,CONFIG.

```
F AQSADV,DISPLAY,CONFIG
AQS0143I ADVISOR CONFIGURATION
VERSION          : LIFELINE V2.5
CONFIG FILENAME  : USER.PROFILES(LBADVCNF)
DEBUG LEVEL     : 7
ADVISOR CONN PORT : 8200
ADVISOR ID LIST  :
  IPADDR : 192.10.130.1
  IPADDR : 192.20.130.1
ADVISOR TIMEOUT : 10
```



```

ADVISOR TAKEOVER POLICY : MANUAL
AGENT CONN PORT : 8181
AGENT ID LIST :
  IPADDR..PORT : 192.10.110.1..4000
  IPADDR..PORT : 192.10.110.2..4000
  IPADDR..PORT : 192.10.110.3..4000
  IPADDR..PORT : 192.20.110.1..4000
  IPADDR..PORT : 192.20.110.2..4000
  IPADDR..PORT : 192.20.110.3..4000
UPDATE INTERVAL : 10
FAILURE DETECTION INTERVAL : 60
PEER ADVISOR ID : 192.10.1.1
LB CONNECTION V4 : 192.10.1.1..3860
LB CONNECTION V6 : 192::10:1:1..3860
LB ID LIST :
  IPADDR : 192.10.1.32
  IPADDR : 192.10.1.33
  IPADDR : 192.10.1.64
  IPADDR : 192.10.1.65
WORKLOAD SWITCH POLICY : MANUAL
CROSS SYSPLEX LIST :
  IPADDR..PORT : 10.10.1.1..5001
  SITE : PLEX1
  WORKLOAD : WORKLOAD1
  FAILURE DETECTION INTERVAL : 30
  IPADDR..PORT : 10.20.1.1..5001
  SITE : PLEX2
  WORKLOAD : WORKLOAD1
  FAILURE DETECTION INTERVAL : 30
  IPADDR..PORT : 10.10.1.1..6001
  SITE : PLEX1
  WORKLOAD : WORKLOAD1
  FAILURE DETECTION INTERVAL : 30
  IPADDR..PORT : 10.20.1.1..6001
  SITE : PLEX2
  WORKLOAD : WORKLOAD1
  FAILURE DETECTION INTERVAL : 30
  IPADDR..PORT : 10.10.1.1..8020
  SITE : PLEX1
  WORKLOAD : WORKLOAD2
  FAILURE DETECTION INTERVAL : 60
  IPADDR..PORT : 10.20.1.1..8020
  SITE : PLEX2
  WORKLOAD : WORKLOAD2
  FAILURE DETECTION INTERVAL : 60
  IPADDR..PORT : 10.10.1.1..8021
  SITE : PLEX1
  WORKLOAD : WORKLOAD2
  FAILURE DETECTION INTERVAL : 60
  IPADDR..PORT : 10.20.1.1..8021
  SITE : PLEX2
  WORKLOAD : WORKLOAD2
  FAILURE DETECTION INTERVAL : 60
  IPADDR..PORT : 10.10.1.2..7001
  SITE : PLEX1
  WORKLOAD : WORKLOAD3
  FAILURE DETECTION INTERVAL : 60
  IPADDR..PORT : 10.20.1.2..7001
  SITE : PLEX2
  WORKLOAD : WORKLOAD3
  FAILURE DETECTION INTERVAL : 60
SNA APPLICATION LIST:
APPLICATION: NETA.SNAAPPL1
SITE: PLEX1
WORKLOAD: WORKLOAD4
FAILURE DETECTION INTERVAL: 60
APPLICATION: NETB.SNAAPPL1

```

```

SITE:          PLEX2
WORKLOAD:     WORKLOAD4
FAILURE DETECTION INTERVAL: 60
INTERMEDIARY NODE LIST :
IPADDR..PORT : 10.1.1.1..9001
SITE         : PLEX1
WORKLOAD    : WORKLOAD3
APPLICATION TYPE: IP
IPADDR..PORT : 10.2.1.1..9001
SITE        : PLEX2
WORKLOAD    : WORKLOAD3
APPLICATION TYPE: IP
IPADDR..PORT : 10.3.1.1..10000
SITE        : PLEX1
WORKLOAD    : WORKLOAD4
APPLICATION TYPE: SNA
IPADDR..PORT : 10.4.1.1..10000
SITE        : PLEX2
WORKLOAD    : WORKLOAD4
APPLICATION TYPE: SNA
MQ MANAGER LIST :
MANAGER : QMG1
CLUSTER : CLUSTER1
QUEUE   : QUEUE1
SITE    : PLEX1
WORKLOAD : WORKLOAD5
FAILURE DETECTION INTERVAL : 60
MANAGER : QMG2
CLUSTER : CLUSTER1
QUEUE   : QUEUE1
SITE    : PLEX2
WORKLOAD : WORKLOAD5
FAILURE DETECTION INTERVAL : 60
QUERY WORKLOAD LIST :
WORKLOAD : WORKLOAD2
ASSOCIATED WORKLOAD : WORKLOAD1
DISTRIBUTION : DYNAMIC      ACTION : TOACTIVE
PERSISTENT CONNECTION TIMEOUT : 0
MANAGEMENT GUEST LIST :
GUEST    : GUEST1A
IPADDR   : 10.1.1.1
WORKLOAD : WORKLOAD3
GUEST    : GUEST2A
IPADDR   : 10.2.1.1
WORKLOAD : WORKLOAD3

```

The following information describes the fields in the previous sample display.

Field name

Descriptions

VERSION

The current version and release of the Advisor.

CONFIG FILENAME

The configuration file used by this Advisor when the Advisor was initially started.

DEBUG LEVEL

The debug level currently in effect for this Advisor.

ADVISOR CONN PORT

The port number that the primary Advisor listens on for connections from a peer Advisor.

ADVISOR ID LIST

The list of the IP addresses of secondary Advisors that are permitted to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

ADVISOR TIMEOUT

The interval (in seconds) that the primary Advisor uses to determine whether another primary Advisor is active. The interval used by the secondary Advisor to determine whether the primary Advisor is down.

ADVISOR TAKEOVER POLICY

The action that a secondary Advisor takes when it determines that the primary Advisor is down. The takeover policy allows the following values:

AUTOMATIC

Indicates that the secondary Advisor automatically takes over the primary Advisor responsibilities.

MANUAL

Indicates that the secondary Advisor waits until a `MODIFY AQSADV, TAKEOVER` command is issued before the secondary Advisor takes over the primary Advisor responsibilities.

AGENT CONN PORT

The port number that the primary Advisor listens on for connections from Agents.

AGENT ID LIST

The list of IP addresses of Agents that are permitted to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

UPDATE INTERVAL

The interval (in seconds) that an Agent uses to update the primary Advisor with information about the system and server applications being monitored. Depending on how the external load balancer is configured, the `update_interval` statement might also determine how often the load balancer is updated with data from the Lifeline Advisor.

FAILURE DETECTION INTERVAL

The interval (in seconds) that a primary Advisor uses to detect a site failure.

PEER ADVISOR ID

The IPv4 or IPv6 address that the Advisor uses as the destination IP address for connecting to the peer Advisor.

LB CONNECTION V4

The IPv4 address and port number that the primary Advisor listens on for connections from external load balancers.

LB CONNECTION V6

The IPv6 address and port number the primary Advisor listens on for connections from external load balancers.

LB ID LIST

The list of the IP addresses of external load balancers that are permitted to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

WORKLOAD SWITCH POLICY

The action that a primary Advisor takes when it determines that a workload is no longer active. The switch policy allows the following values:

AUTOMATIC

Indicates that the primary Advisor automatically reroutes a failing workload to the alternate site.

MANUAL

Indicates that the primary Advisor waits until a MODIFY AQSADV,ACTIVATE command is issued before the primary Advisor reroutes a failing workload to the alternate site.

CROSS SYSPLEX LIST

The list of both the external and internal second-tier load balancer IP addresses, the site name (i.e. sysplex name) that identifies each second-tier load balancer, the port number of the server application used for a specific workload, the workload name, and the failure detection interval for the workload. The IP address can be an IPv4 or an IPv6 address.

SNA APPLICATION LIST

The list of SNA applications, the site where the application resides, the workload name, and the failure detection interval for the workload.

INTERMEDIARY NODE LIST

The list of IP addresses and ports that identify each intermediary node, the site name where the intermediary node resides, the workload that the intermediary node belongs to, and the type of application that processes the workload. The site name is the same as the sysplex name. The IP address can be an IPv4 or IPv6 address.

MQ_MANAGER_LIST

The list of MQ queue managers, the IBM MQ clusters hosted by the MQ queue manager, the MQ cluster queues defined on the MQ queue manager, the site where the MQ queue manager resides, the workload associated with each IBM MQ cluster, and the failure detection interval for the workload.

QUERY WORKLOAD LIST

The list of the Active/Query and associated Active/Standby workloads, the distribution type, distribution option, and persistent connection timeout for the Active/Query workload. The distribution type can be DYNAMIC or STATIC.

MANAGEMENT_GUEST LIST

The list of Linux on System z guests that the Linux on System z Management Agents monitor, each guest's IP address, and the workload that the guest belongs to.

The following example shows a sample display when you issue F AQSADV,DISPLAY,DEBUG.

```
F AQSADV,DISPLAY,DEBUG
AQS0111I LLADVSR DEBUG LEVEL 7
```

The following example shows a sample display when you issue F AQSADV,DISPLAY,LB.

```
F AQSADV,DISPLAY,LB
AQS0112I LOAD BALANCER SUMMARY
LB INDEX      : 00      UUID      : F956AE65
```

```

IPADDR..PORT : 192.10.1.64..55764
HEALTH       : 7E           FLAGS    : NOCHANGE PUSH
LB INDEX     : 01           UUID     : 9A78BE9E
IPADDR..PORT : 192.10.1.32..40366
HEALTH       : 7E           FLAGS    : NOCHANGE PUSH
2 OF 2 RECORDS DISPLAYED

```

The following information describes the fields in the previous sample display.

Field name

Descriptions

LB INDEX

The reference number used to identify the load balancer. The same reference number is used for a load balancer as long as it remains connected.

UUID A hexadecimal value of the universally unique identifier assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

IPADDR..PORT

The IP address and port number used by the load balancer to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

HEALTH

A hexadecimal value supplied by the load balancer that indicates the general health of the load balancer. Valid values are in the range from 0 to X'7F'.

FLAGS

Indicators set by the load balancer. The flag values are shown as follows:

NOCHANGE

Indicates that the primary Advisor sends only weights that have changed since the last update interval to the load balancer.

PUSH Indicates that the primary Advisor sends weights to the external load balancer when the update_interval has expired. If this flag is not present, the load balancer periodically retrieves weights from the primary Advisor.

TRUST

Indicates that the load balancer trusts member applications to register themselves. This flag is ignored by the primary Advisor.

The following example shows a sample display when you issue F AQSADV,DISPLAY,LB,DETAIL.

```

F AQSADV,DISPLAY,LB,DETAIL
AQS0113I LOAD BALANCER DETAILS
LB INDEX     : 00           UUID     : F956AE65
IPADDR..PORT : 192.10.1.64..55764
HEALTH       : 7E           FLAGS    : NOCHANGE PUSH
GROUP NAME   : WKLD2_GROUP1
WORKLOAD     : WORKLOAD2
GROUP FLAGS  : CROSS_SYSPLEX
IPADDR..PORT : 10.10.1.1..8020
SYSPLEX     : PLEX1
SYSTEM NAME : N/A           PROTOCOL : TCP   AVAIL    : YES
WLM WEIGHT  : N/A           CS WEIGHT : N/A   NET WEIGHT: 00001
FLAGS       : DISTDVIPA
IPADDR..PORT : 10.20.1.1..8020
SYSPLEX     : PLEX2

```

```

SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000
FLAGS       : CMQ DISTDVIPA
GROUP NAME  : WKLD2_GROUP2
WORKLOAD    : WORKLOAD2
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 10.10.1.1..8021
SYSPLEX     : PLEX1
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00001
FLAGS       : DISTDVIPA
IPADDR..PORT: 10.20.1.1..8021
SYSPLEX     : PLEX2
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000
FLAGS       : CMQ DISTDVIPA
GROUP NAME  : WKLD1_GROUP1
WORKLOAD    : WORKLOAD1
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 10.10.1.1..5001
SYSPLEX     : PLEX1
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : NO
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000
FLAGS       : CMQ
IPADDR..PORT: 10.20.1.1..5001
SYSPLEX     : PLEX2
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00001
FLAGS       : DISTDVIPA
GROUP NAME  : WKLD1_GROUP2
WORKLOAD    : WORKLOAD1
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 10.10.1.1..6001
SYSPLEX     : PLEX1
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : NO
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000
FLAGS       : CMQ
IPADDR..PORT: 10.20.1.1..6001
SYSPLEX     : PLEX2
SYSTEM NAME: N/A      PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00001
FLAGS       : DISTDVIPA
GROUP NAME  : WKLD3_GROUP1
WORKLOAD    : WORKLOAD3
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 10.1.1.1..9001
SYSPLEX     : PLEX1
SYSTEM NAME: LINUXMG1 PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00001
FLAGS       : CMQ
IPADDR..PORT: 10.2.1.1..9001
SYSPLEX     : PLEX2
SYSTEM NAME: LINUXMG2 PROTOCOL : TCP  AVAIL      : NO
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000
FLAGS       : CMQ
GROUP NAME  : WKLD4_GROUP1
WORKLOAD    : WORKLOAD4
GROUP FLAGS : CROSS_SYSPLEX
IPADDR..PORT: 10.3.1.1..10000
SYSPLEX     : PLEX1
SYSTEM NAME: SYS1A    PROTOCOL : TCP  AVAIL      : YES
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00001
FLAGS       : CMQ
IPADDR..PORT: 10.4.1.1..10000
SYSPLEX     : PLEX2
SYSTEM NAME: SYS2A    PROTOCOL : TCP  AVAIL      : NO
WLM WEIGHT  : N/A      CS WEIGHT  : N/A  NET WEIGHT: 00000

```

```

      FLAGS      : CMQ
LB INDEX      : 01          UUID      : 9A78BE9E
IPADDR..PORT : 192.10.1.32..40366
HEALTH       : 7E          FLAGS     : NOCHANGE PUSH
GROUP NAME    : TIER2_GROUP1
WORKLOAD     : WORKLOAD1
GROUP FLAGS   :
IPADDR..PORT : 192.10.110.1..5001
SYSPLEX      : PLEX1
SYSTEM NAME: SYS1A        PROTOCOL   : TCP AVAIL      : YES
WLM WEIGHT   : 00058      CS WEIGHT  : 100 NET WEIGHT: 00001
RAW          CP: 58 ZAAP: 00 ZIIP: 00
PROPORTIONAL CP: 58 ZAAP: 00 ZIIP: 00
ABNORM      : 00000      HEALTH    : 100
FLAGS       :
IPADDR..PORT : 192.10.110.2..5001
SYSPLEX      : PLEX1
SYSTEM NAME: SYS1B        PROTOCOL   : TCP AVAIL      : YES
WLM WEIGHT   : 00059      CS WEIGHT  : 100 NET WEIGHT: 00001
RAW          CP: 59 ZAAP: 00 ZIIP: 00
PROPORTIONAL CP: 58 ZAAP: 00 ZIIP: 00
ABNORM      : 00000      HEALTH    : 100
FLAGS       :
GROUP NAME   : TIER2_GROUP2
WORKLOAD    : WORKLOAD2
GROUP FLAGS  :
IPADDR..PORT : 192.10.110.1..6001
SYSPLEX      : PLEX1
SYSTEM NAME: SYS1A        PROTOCOL   : TCP AVAIL      : YES
WLM WEIGHT   : 00058      CS WEIGHT  : 100 NET WEIGHT: 00001
RAW          CP: 58 ZAAP: 00 ZIIP: 00
PROPORTIONAL CP: 58 ZAAP: 00 ZIIP: 00
ABNORM      : 00000      HEALTH    : 100
FLAGS       :
IPADDR..PORT : 192.10.110.2..6001
SYSPLEX      : PLEX1
SYSTEM NAME: SYS1B        PROTOCOL   : TCP AVAIL      : YES
WLM WEIGHT   : 00059      CS WEIGHT  : 100 NET WEIGHT: 00001
RAW          CP: 59 ZAAP: 00 ZIIP: 00
PROPORTIONAL CP: 58 ZAAP: 00 ZIIP: 00
ABNORM      : 00000      HEALTH    : 100
FLAGS       :
2 OF 2 RECORDS DISPLAYED

```

The following information describes the fields in the previous sample display.

Field name

Descriptions

LB INDEX

A reference number used to identify the load balancer. The same reference number is used for a load balancer as long as it remains connected.

UUID A hexadecimal value of the universally unique identifier assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

IPADDR..PORT

- The instance listed under the LB INDEX field represents the IP address and port number used by the load balancer to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.
- The instances listed within the GROUP NAME block for first-tier load balancers represent the IP address and port number of a second-tier load balancer that the first-tier load balancer balances workload requests.

- The instances listed within the GROUP NAME block for second-tier load balancers represent the IP address and port number of a target server application that the second-tier load balancer balances workload requests.

HEALTH

A hexadecimal value supplied by the load balancer that indicates the general health of the load balancer. Valid values are in the range from 0 to X'7F'.

FLAGS

Indicators set by the load balancer. The flag values are shown as follows:

NOCHANGE

Indicates that the primary Advisor sends only weights that have changed since the last update interval to the load balancer.

PUSH Indicates that the primary Advisor sends weights to the external load balancer when the update_interval has expired. If this flag is not present, the load balancer periodically retrieves weights from the primary Advisor.

TRUST

Indicates that the load balancer trusts member applications to register themselves. This flag is ignored by the primary Advisor.

GROUP

The name of a registered group of related target server applications. The group name is a UTF-8 string displayed in EBCDIC on the MVS console. Any non-displayable character is displayed as a question mark (?).

GROUP FLAGS

The flags currently applicable to the group. Group flags are only valid for first-tier load balancers; for second-tier load balancers, this field is left blank. Here is the flag value:

CROSS_SYSPLEX

The group is a cross-sysplex group. This group is used to balance the workload across multiple sysplexes. The IP addresses specified in each cross-sysplex group identify the second-tier load balancers that the first-tier load balancer will balance across.

SYSPLEX

For a first-tier load balancer application group, the value of SYSPLEX indicates the MVS sysplex name where the specified second-tier load balancer will load balance to.

For entries that represent intermediary nodes on a Linux on System z guest, the SYSPLEX name corresponds to the value of the site_name Agent configuration statement on the Linux on System z Management Agent.

For a second-tier load balancer application group, the value of SYSPLEX indicates the MVS sysplex name of the system where the server application being load balanced exists.

SYSTEM NAME

For a first-tier load balancer application group, the system name is N/A, because the registered members are second-tier load balancers.

For entries that represent intermediary nodes on a Linux on System z guest, the system name is the guest name of the intermediary node.

For a second-tier load balancer application group, the system name is the name of the MVS system where the server application exists.

PROTOCOL

The protocol used by the second-tier load balancer for first-tier load balancer application groups or server application for second-tier load balancer application groups. Protocol value is either TCP or UDP.

AVAIL

Indicates whether the second-tier load balancer for a first-tier load balancer application group or the server application for a second-tier load balancer application group is available for workload distribution. The value YES indicates that the second-tier load balancer or the server application is available for load balancing. The value NO indicates that the second-tier load balancer or the server application is not recommended for load balancing.

WLM WEIGHT

The server-specific Workload Manager (WLM) weight value of the server application. For a first-tier load balancer application group, the WLM weight is N/A, because the registered members are second-tier load balancers. For a second-tier load balancer group, the value is a composite weight, in the range 0 - 64, which indicates the sum of the proportional CP, zAAP, and zIIP weights for this server application.

CP Indicates that the RAW value is the WLM server-specific general CP recommendation. This is the amount of displaceable general CPU capacity based on the importance, which is defined by the WLM policy, of the server application as compared to the other target systems within the same MVS sysplex. The PROPORTIONAL value is the RAW value modified by the proportion of general CP capacity that is currently being consumed by the workload of the server application as compared with the other processors (zAAP and zIIP).

zAAP Indicates that the RAW value is the WLM server-specific zAAP recommendation. This is the amount of displaceable zAAP capacity based on the importance, which is defined by the WLM policy, of the server application as compared to the other target systems within the same MVS sysplex. The PROPORTIONAL value is the RAW value modified by the proportion of zAAP capacity that is currently being consumed by the workload of the server application as compared with the other processors (general CPU and zIIP).

zIIP Indicates that the RAW value is the WLM server-specific zIIP recommendation. This is the amount of displaceable zIIP capacity based on the importance, which is defined by the WLM policy, of the server application as compared with the other target systems within the same MVS sysplex. The PROPORTIONAL value is the RAW value modified by the proportion of zIIP capacity that is currently being consumed by the workload of the server application as compared with the other processors (general CPU and zAAP).

ABNORM

Indicates the rate of abnormal transaction completions per 1000 total transaction completions. The value is nonzero if the server application is experiencing conditions in which transactions are not

completed. It is applicable to server applications that report transaction status using Workload Management Services. Under normal conditions or if the server application is not providing this information to WLM, the value of ABNORM is 0. WLM uses abnormal transaction completion rate information to reduce the WLM weight reported for the server application.

HEALTH

The general health of the server application. Values less than 100 indicate that the server application is experiencing conditions that are not enabling it to process new workload requests successfully. It is applicable to server applications that report their health using Workload Management Services. Under normal circumstances or if the server application is not providing this information to WLM, the value of HEALTH is 100. WLM uses health information to reduce the WLM weight reported for the server application.

CS WEIGHT

The Communications Server (CS) weight value recommended by the Lifeline Agent. For a first-tier load balancer application group, the CS weight is N/A, because the registered members are second-tier load balancers. For a second-tier load balancer application group, the CS weight value range is 0 - 100. A higher weight indicates that the server application is able to handle more workload requests than a server application in the same group with a lower weight.

NET WEIGHT

For a first-tier load balancer application group, NET WEIGHT indicates the relative weight of this second-tier load balancer compared to other second-tier load balancers for this workload. For members of an Active/Standby workload, the net weight is either 0 or 1, with a value of 0 indicating that no workload requests are being distributed to this second-tier load balancer, while a value of 1 indicating that workload requests can be distributed to this second-tier load balancer.

For members of an Active/Query workload, the net weight value range is 0 - 64. A higher weight indicates that the site where the second-tier load balancer routes to can handle more requests for the workload than the alternate site. This weight is based on the WLM weight, the CS weight, database replication latency state, and other factors. Weights within a group are normalized to yield the net weight. Normalization involves reducing the weight values while preserving the ratios between the weights.

Additional factors might affect the NET WEIGHT of a first-tier load balancer application group. If the site that a second-tier load balancer is routing to is not active, the net weight for this second-tier load balancer is 0 so that new work requests are not sent to this site. The net weight is also affected by the type of second-tier load balancers in the group; they might be sysplex distributed DVIPAs, external load balancers, or intermediary nodes. For more information, see "Interpreting the workload and LB displays" on page 70.

For a second-tier load balancer application group, NET WEIGHT indicates the relative weight of this server application compared with other server applications within the same site for this workload. The net weight value range is 0 - 64. A higher weight indicates that a server application can handle more workload requests than a lower weight server application in the same group. This weight is based upon the WLM weight, the CS

weight, and other factors. Compare net weights for server applications only with other server applications within the same group. Weights within a group are normalized to yield the net weight. Normalization involves reducing the weight values while preserving the ratios between the weights.

FLAGS

Indicators for the target server applications. The flag values are shown as follows:

CMQ Indicates that the member's site and workload were quiesced with a `MODIFY advisor_proc,QUIESCE` command. No new workload connections are routed to server applications for the specified site and workload.

DISTDVIPA

Indicates that the IP address of the server application is a distributable DVIPA. The IP address can be an IPv4 or an IPv6 address.

LBQ Indicates that an operator at the Load Balancer quiesced the member. No new workload connections are routed to the server application.

NODATA

Indicates that an Agent reported the server application, but at least two update intervals are needed to calculate a valid CS weight.

NOTARGETAPP

Indicates that an Agent reported the member's IP address on a TCP/IP stack but did not find a server application that uses the same port and protocol.

NOTARGETIP

Indicates that an Agent reported the member's IP address on a TCP/IP stack but the IP address is not usable.

NOTARGETSYS

Indicates that no Agent reported the member's IP address on any TCP/IP stacks in either site.

OPQ Indicates that an operator at the owning Agent quiesced the server application or system. No new workload connections are routed to this server application or system.

The following example shows a sample display when you issue `F AQSADV,DISPLAY,WORKLOAD`.

```
F AQSADV,DISPLAY,WORKLOAD
AQS0145I WORKLOAD SUMMARY
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORKLOAD1
STATE : ACTIVE
SITE : PLEX2
WORKLOAD NAME : WORKLOAD3
STATE : ACTIVE
SITE : PLEX1
WORKLOAD NAME : WORKLOAD4
STATE : ACTIVE
SITE : PLEX1
WORKLOAD NAME : WORKLOAD5
STATE : ACTIVE
SITE : PLEX1
TYPE : ACTIVE/QUERY
```

```
WORKLOAD NAME : WORKLOAD2
STATE : ACTIVE
SITE : PLEX1
SITE : PLEX2
ASSOCIATED WORKLOAD : WORKLOAD1
REPLICATION STATE : NORMAL
5 OF 5 RECORDS DISPLAYED
```

The following information describes the fields in the previous sample display.

Field name

Descriptions

TYPE The type of workload. Type values are shown as follows:

ACTIVE/STANDBY

Indicates that the workload has been configured so that it can be activated to only one site at any one time.

ACTIVE/QUERY

Indicates that the workload has been configured so that it can be activated to both sites at the same time.

WORKLOAD NAME

The name of the workload as defined in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement.

STATE

The state of the workload. The state values are shown as follows:

ACTIVE

Indicates that the workload has been activated to the specified site. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies first-tier load balancers to route all inbound requests for this workload to the specified site. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '1' and start the MQ cluster receiver channel.

QUIESCED

Indicates that the workload has been quiesced to all sites. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies first-tier load balancers to not route any inbound requests for this workload. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '0' and stop the MQ cluster receiver channel.

DEACTIVATED

Indicates that the workload has been deactivated. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor resets any existing connections for this workload on systems where active Agents existed when the workload was deactivated. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the previously active site to determine if any MQ cluster receiver channels are not stopped and if any messages exist on any MQ cluster queue for the workload.

SITE The MVS sysplex name where the workload is active. If the state of the workload is QUIESCED or DEACTIVATED, the SITE value is set to N/A. For Active/Query workloads, if the workload is active on both sites, a separate line for each site is displayed.

ASSOCIATED WORKLOAD

For workloads of type Active/Query, the associated Active/Standby workload.

REPLICATION STATE

For workloads of type Active/Query, the current state of data replication for the associated Active/Standby workload. See the specific database replication product documentation for information about how to configure replication latency thresholds. The state values are shown as follows:

NORMAL

Indicates that the replication latency is within the acceptable threshold that was configured for database replication for this workload.

CONSTRAINED

Indicates that the replication latency is within 80% of the maximum threshold that was configured for database replication for this workload.

CRITICAL

Indicates that the replication latency is within 90% of the maximum threshold that was configured for database replication for this workload.

ACUTE

Indicates that the replication latency has exceeded the maximum threshold that was configured for database replication for this workload.

UNKNOWN

Indicates that the replication latency value has either not been reported by the database replication product, or the database replication product cannot determine the current replication latency state.

The following example shows a sample display when you issue F AQSADV,DISPLAY,WORKLOAD,DETAIL.

```
F AQSADV,DISPLAY,WORKLOAD,DETAIL
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORKLOAD1
STATE : ACTIVE
SITE : PLEX2
SERVERS:
IPADDR..PORT : 192.10.110.1..5001
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..5001
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
IPADDR..PORT : 192.10.110.2..5001
SYSTEM NAME : SYS1B   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..5001
SYSTEM NAME : SYS2B   SITE : PLEX2   STATE : UNAVAIL
IPADDR..PORT : 192.10.110.1..6001
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..6001
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
```

```

IPADDR..PORT : 192.10.110.2..6001
SYSTEM NAME : SYS1B   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..6001
SYSTEM NAME : SYS2B   SITE : PLEX2   STATE : AVAIL
WORKLOAD NAME : WORKLOAD3
STATE : ACTIVE
SITE : PLEX1
SERVERS:
IPADDR..PORT : 10.10.1.2..7001
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.2..7001
SYSTEM NAME : N/A     SITE : PLEX2   STATE : UNAVAIL
ZLINUX SERVERS:
HOST : GUEST1A IPADDR..PORT : 10.1.1.1..9001
SYSTEM ID : ZVM1     SITE : PLEX1   STATE : AVAIL
HOST : GUEST2A IPADDR..PORT : 10.2.1.1..9001
SYSTEM ID : ZVM2     SITE : PLEX2   STATE : UNAVAIL
WORKLOAD NAME : WORKLOAD4
STATE : ACTIVE
SITE : PLEX1
SERVERS:
IPADDR..PORT : 10.3.1.1..10000
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.4.1.1..10000
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
SNA APPLICATIONS:
APPL: NETA.SNAPPL1
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
APPL: NETB.SNAPPL2
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
WORKLOAD NAME : WORKLOAD5
STATE : ACTIVE
SITE : PLEX1
MQ QUEUE MANAGERS:
QMGR NAME : QMG1 QUEUE : QUEUE1
IPADDR..PORT : 10.10.1.1..8001
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
QMGR NAME : QMG2 QUEUE : QUEUE1
IPADDR..PORT : N/A
SYSTEM NAME : N/A     SITE : PLEX2   STATE : UNAVAIL
TYPE : ACTIVE/QUERY
WORKLOAD NAME : WORKLOAD2
STATE : ACTIVE
SITE : PLEX1
SITE : PLEX2
ASSOCIATED WORKLOAD : WORKLOAD1
REPLICATION STATE : NORMAL
SERVERS:
IPADDR..PORT : 10.10.1.1..8020
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.10.1.1..8020
SYSTEM NAME : SYS1B   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8020
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8020
SYSTEM NAME : SYS2B   SITE : PLEX2   STATE : AVAIL
IPADDR..PORT : 10.10.1.1..8021
SYSTEM NAME : SYS1A   SITE : PLEX1   STATE : UNAVAIL
IPADDR..PORT : 10.10.1.1..8021
SYSTEM NAME : SYS1B   SITE : PLEX1   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8021
SYSTEM NAME : SYS2A   SITE : PLEX2   STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8021
SYSTEM NAME : SYS2B   SITE : PLEX2   STATE : AVAIL

```

5 OF 5 RECORDS DISPLAYED

The following information describes the fields in the previous sample display.2 OF
2 RECORDS DISPLAYED

Field name

Descriptions

TYPE The type of workload. Type values are shown as follows:

ACTIVE/STANDBY

Indicates that the workload has been configured so that it can be activated to only one site at any one time.

ACTIVE/QUERY

Indicates that the workload has been configured so that it can be activated to both sites at the same time.

WORKLOAD NAME

The name of the workload as defined in the cross_sysplex_list, mq_manager_list, or sna_application_list Advisor configuration statement.

STATE

The state of the workload. The state values are shown as follows:

ACTIVE

Indicates that the workload has been activated to the specified site. For workloads defined in the cross_sysplex_list or sna_application_list Advisor configuration statement, the Advisor notifies first-tier load balancers to route all inbound requests for this workload to the specified site. For workloads defined in the mq_manager_list Advisor configuration statement, the Advisor notifies Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '1' and start the MQ cluster receiver channel.

QUIESCED

Indicates that the workload has been quiesced to all sites. For workloads defined in the cross_sysplex_list or sna_application_list Advisor configuration statement, the Advisor notifies first-tier load balancers to not route any inbound requests for this workload. For workloads defined in the mq_manager_list Advisor configuration statement, the Advisor notifies Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '0' and stop the MQ cluster receiver channel.

DEACTIVATED

Indicates that the workload has been deactivated. For workloads defined in the cross_sysplex_list or sna_application_list Advisor configuration statement, the Advisor resets any existing connections for this workload on systems where active Agents existed when the workload was deactivated. For workloads defined in the mq_manager_list Advisor configuration statement, the Advisor notifies Lifeline Agents on the previously active site to determine if any MQ cluster receiver channels are not stopped and if any messages exist on any MQ cluster queue for the workload.

SITE The MVS sysplex name where the workload is active. If the state of the workload is QUIESCED or DEACTIVATED, the SITE value is set to N/A. For Active/Query workloads, if the workload is active on both sites, a separate line for each site is displayed.

ASSOCIATED WORKLOAD

For workloads of type Active/Query, the associated Active/Standby workload.

REPLICATION STATE

For workloads of type Active/Query, the current state of data replication for the associated Active/Standby workload. See the specific database replication product documentation for information about how to configure replication latency thresholds. The state values are shown as follows:

NORMAL

Indicates that the replication latency is within the acceptable threshold that was configured for database replication for this workload.

CONSTRAINED

Indicates that the replication latency is within 80% of the maximum threshold that was configured for database replication for this workload.

CRITICAL

Indicates that the replication latency is within 90% of the maximum threshold that was configured for database replication for this workload.

ACUTE

Indicates that the replication latency has exceeded the maximum threshold that was configured for database replication for this workload.

UNKNOWN

Indicates that the replication latency value has either not been reported by the database replication product, or the database replication product cannot determine the current replication latency state.

IPADDR..PORT

The IP address and port number that the server application is listening on. The IP address can be an IPv4 or an IPv6 address.

SYSTEM NAME

The MVS system name of the system where the server application exists. This value is N/A if the system is not known.

SITE The MVS sysplex name of the site where the server application exists.

STATE

The status of the server application. The state values are shown as follows:

AVAIL

Indicates that the server application is available on the specified system and is able to process workload requests.

UNAVAIL

Indicates that the server application is not available on the specified system and is not able to process workload requests. Possible reasons for the server application being unavailable include not being active on the specified system or quiesced by the MVS operator.

ZLINUX SERVERS

Information about the Linux on System z server applications that process a workload.

HOST The name of the Linux on System z guest where the server application exists.

IPADDR..PORT

The IP address and port number that the server application is listening on. The IP address can be an IPv4 or an IPv6 address.

SYSTEM ID

The z/VM system identifier of the system where the Linux on System z guest resides.

SITE The name of the site where the Linux on System z server application exists. This value is blank when the Linux on System z guest is not available.

STATE

The status of the Linux on System z server application. The state values are shown as follows:

AVAIL

Indicates that the Linux on System z server application is available and active on the specified host and is able to process workload requests.

UNAVAIL

Indicates that the Linux on System z server application is not available on the specified host and is not able to process workload requests. If the system identifier is not available (N/A), the server application was not reported by any Linux on System z Management Agent in the specified site.

SNA APPLICATIONS

Information about the SNA applications that process a workload.

APPL The SNA application name.

SYSTEM NAME

The MVS system name of the system where the SNA application exists. This value is N/A if the system is not known.

SITE The MVS sysplex name of the site where the SNA application exists.

STATE

The status of the SNA application. The state values are shown as follows:

AVAIL

Indicates that the SNA application is available and active on the specified system and is able to process workload requests.

UNAVAIL

Indicates that the SNA application is not available on the specified host and is not able to process workload requests. If the system name is not available (N/A), the SNA application was not reported by any z/OS Agent in the specified site.

MQ QUEUE MANAGERS

Information about the MQ queue managers that process a workload.

QMGR NAME

The name of the MQ queue manager.

|
| **QUEUE**

| The name of a MQ cluster queue that is defined on this MQ queue
| manager.

| **IPADDR..PORT**

| The IP address and port number that the MQ queue manager is listening
| on for this workload's MQ cluster receiver channel. The IP address can be
| an IPv4 or an IPv6 address.

| **SYSTEM NAME**

| The MVS system name of the system where the MQ queue manager exists.
| This value is N/A if the system is not known.

| **SITE** The MVS sysplex name of the site where the MQ queue manager exists.

| **STATE**

| The status of the MQ queue manager. The state values are shown as
| follows:

| **AVAIL**

| Indicates that the MQ queue manager is available on the specified
| system and is able to process messages for the workload on the
| specified MQ cluster queue.

| **UNAVAIL**

| Indicates that the MQ queue manager is not able to process
| messages for the workload on the specified system and MQ cluster
| queue. If the system name is not available (N/A), the MQ queue
| manager was not reported by any z/OS Agent in the specified site.

| The following example shows a sample display when you issue F
| AQSADV,DISPLAY,WORKLOAD,STATUS.

| F AQSADV,DISPLAY,WORKLOAD,STATUS
| AQS0149I WORKLOAD STATUS
| TYPE : ACTIVE/STANDBY
| WORKLOAD NAME : WORKLOAD1
| SITE : PLEX1 STATE : DEACTIVATED STATUS : AVAIL
| SITE : PLEX2 STATE : ACTIVE STATUS : AVAIL
| WORKLOAD NAME : WORKLOAD3
| SITE : PLEX1 STATE : ACTIVE STATUS : AVAIL
| SITE : PLEX2 STATE : DEACTIVATED STATUS : AVAIL
| WORKLOAD NAME : WORKLOAD4
| SITE : PLEX1 STATE : ACTIVE STATUS : AVAIL
| SITE : PLEX2 STATE : DEACTIVATED STATUS : UNAVAIL
| WORKLOAD NAME : WORKLOAD5
| SITE : PLEX1 STATE : ACTIVE STATUS : AVAIL
| SITE : PLEX2 STATE : DEACTIVATED STATUS : AVAIL
| TYPE : ACTIVE/QUERY
| WORKLOAD NAME : WORKLOAD2
| SITE : PLEX1 STATE : ACTIVE STATUS : AVAIL
| SITE : PLEX2 STATE : ACTIVE STATUS : AVAIL
| 5 OF 5 RECORDS DISPLAYED

| The following information describes the fields in the previous sample display.

| **Field name**

| **Descriptions**

| **TYPE** The type of workload. Type values are shown as follows:

| **ACTIVE/STANDBY**

| Indicates that the workload has been configured so that it can be
| activated to only one site at any one time.

|
| **ACTIVE/QUERY**
|

Indicates that the workload has been configured so that it can be activated to both sites at the same time.

|
| **WORKLOAD NAME**
|

The name of the workload as defined in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement.

| **SITE** The MVS sysplex name where the workload is configured. A separate line for each site is displayed.
|

|
| **STATE**
|

The state of the workload. The state values are shown as follows:

| **ACTIVE**
|

Indicates that the workload has been activated to the specified site. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies first-tier load balancers to route all inbound requests for this workload to the specified site. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '1' and start the MQ cluster receiver channel.

| **QUIESCED**
|

Indicates that the workload has been quiesced on the specified site. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies first-tier load balancers to not route any inbound requests for this workload to the specified site. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the specified site to change the queue rank of each MQ cluster queue in the workload to a '0' and stop the MQ cluster receiver channel.

| **DEACTIVATED**
|

Indicates that the workload has been deactivated on the specified site. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor resets any existing connections for this workload on systems where active Agents existed when the workload was deactivated. For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies Lifeline Agents on the specified site to determine if any messages exist on any MQ cluster queue for the workload.

| **STATUS**
|

The status of the workload. The status values are shown as follows:

| **AVAIL**
|

Indicates one of the following:

- For workloads defined in the `cross_sysplex_list` Advisor configuration statement, both of the following conditions are true:
 - At least one server application is available for the workload and is able to process workload requests on the specified site.
 - If the workload is specified in the `query_workload_list` Advisor configuration statement, the replication latency state

for the associated Active/Standby workload is NORMAL, CONSTRAINED, or CRITICAL.

- For workloads defined in the `sna_application_list` Advisor configuration statement, both of the following conditions are true:
 - At least one SNA application is available for the workload and is able to process workload requests on the specified site.
 - If the workload is configured with monitored intermediary nodes, at least one intermediary node is available for the workload and is able to process workload requests on the specified site.
- For workloads defined in the `management_guest_list` Advisor configuration statement, both of the following conditions are true:
 - At least one server application on z/OS is available for the workload and is able to process workload requests on the specified site.
 - At least one server application on Linux on System z is available for the workload and is able to process workload requests on the specified site.
- For workloads defined in the `mq_manager_list` Advisor configuration statement, at least one MQ cluster queue for each configured MQ cluster queue is available for the workload and is able to receive workload messages on the specified site.

UNAVAIL

Indicates one of the following:

- For workloads defined in the `cross_sysplex_list` Advisor configuration statement, either of the following conditions is true:
 - No server application is available for the workload to process workload requests on the specified site.
 - If the workload is specified in the `query_workload_list` Advisor configuration statement, the replication latency state for the associated Active/Standby workload is ACUTE or UNKNOWN.
- For workloads defined in the `sna_application_list` Advisor configuration statement, either of the following conditions is true:
 - No SNA application is available for the workload to process workload requests on the specified site.
 - If the workload is configured with monitored intermediary nodes, no intermediary node is available for the workload to process workload requests on the specified site.
- For workloads defined in the `management_guest_list` Advisor configuration statement, either of the following conditions is true:
 - No server application on z/OS is available for the workload to process workload requests on the specified site.
 - No server application on Linux on System z is available for the workload to process workload requests on the specified site.

- For workloads defined in the mq_manager_list Advisor configuration statement, no MQ cluster queue for at least one configured MQ cluster queue is available for the workload to receive workload messages on the specified site.

The following example shows a sample display when you issue `F AQSADV,DISPLAY,WORKLOAD=WORKLOAD1`.

```
F AQSADV,DISPLAY,WORKLOAD=WORKLOAD1
AQS0145I WORKLOAD SUMMARY
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORKLOAD1
STATE : ACTIVE
SITE : PLEX2
1 OF 1 RECORDS DISPLAYED
```

See the example when a `F AQSADV,DISPLAY,WORKLOAD` command is issued for information that describes the fields in this sample display.

The following example shows a sample display when you issue `F AQSADV,DISPLAY,WORKLOAD=*ALL_AQ*,DETAIL`.

```
F AQSADV,DISPLAY,WORKLOAD=*ALL_AQ*,DETAIL
AQS0146I WORKLOAD DETAILS
TYPE : ACTIVE/QUERY
WORKLOAD NAME : WORKLOAD2
STATE : ACTIVE
SITE : PLEX1
SITE : PLEX2
ASSOCIATED WORKLOAD : WORKLOAD1
REPLICATION STATE : NORMAL
SERVERS:
IPADDR..PORT : 10.10.1.1..8020
SYSTEM NAME : SYS1A SITE : PLEX1 STATE : AVAIL
IPADDR..PORT : 10.10.1.1..8020
SYSTEM NAME : SYS1B SITE : PLEX1 STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8020
SYSTEM NAME : SYS2A SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8020
SYSTEM NAME : SYS2B SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 10.10.1.1..8021
SYSTEM NAME : SYS1A SITE : PLEX1 STATE : UNAVAIL
IPADDR..PORT : 10.10.1.1..8021
SYSTEM NAME : SYS1B SITE : PLEX1 STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8021
SYSTEM NAME : SYS2A SITE : PLEX2 STATE : AVAIL
IPADDR..PORT : 10.20.1.1..8021
SYSTEM NAME : SYS2B SITE : PLEX2 STATE : AVAIL
1 OF 1 RECORDS DISPLAYED
```

See the example when a `F AQSADV,DISPLAY,WORKLOAD,DETAIL` command is issued for information that describes the fields in this sample display.

The following example shows a sample display when you issue `F AQSADV,DISPLAY,WORKLOAD=WORKLOAD1,STATUS`.

```
F AQSADV,DISPLAY,WORKLOAD=WORKLOAD1,STATUS
AQS0149I WORKLOAD STATUS
TYPE : ACTIVE/STANDBY
WORKLOAD NAME : WORKLOAD1
SITE : PLEX1 STATE : DEACTIVATED STATUS : AVAIL
SITE : PLEX2 STATE : ACTIVE STATUS : AVAIL
1 OF 1 RECORDS DISPLAYED
```


See the example when a the example when a **F AQSADV,DISPLAY,WORKLOAD,STATUS** command is issued for information that describes the fields in this sample display.

MODIFY *advisor_proc*,ACTivate command: Activate workloads

Use the **MODIFY *advisor_proc*,ACTivate** command to activate a workload. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, force the Lifeline Advisor to signal first-tier load balancers that new connections for the specified workload or for all the workloads can be routed to the requested site. For workloads defined in the `mq_manager_list` Advisor configuration statement, force the Lifeline Advisor to signal Lifeline Agents on the active site to change the queue rank of each MQ cluster queue in the workload to a '1' and start the workload's MQ cluster receiver channel.

Rule: You can issue this command only from the primary Advisor.

Format



Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

ACTivate

Signals the first-tier load balancers that a workload can be distributed to a site.

WORKload=

workloadname

Specifies the name of the workload to be activated. The name must match one of the workload names specified in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement. The *workloadname* value must not exceed 64 characters in length.

Rules:

- You must quiesce the specified workload before you issue the command to activate it.
- You must have at least one active Agent on the target site.
- For Active/Standby workloads, if active connections remain after you quiesce the specified workload on the alternate site, you must deactivate the workload before you issue the command to activate it on the target site.

ALL_AS

Indicates that all the Active/Standby workloads specified in the `cross_sysplex_list`, `mq_manager_list`, and `sna_application_list` Advisor configuration statements are to be activated on the requested site.

Rules:

- You must quiesce all the Active/Standby workloads before you issue the command to activate them.
- You must have at least one active Agent on the target site.
- If active connections remain for one or more of the Active/Standby workloads after you quiesce the workloads on the alternate site, you must deactivate the applicable workloads before you issue the command to activate them on the target site.

ALL_AQ

Indicates that all the Active/Query workloads specified in the `cross_sysplex_list` Advisor configuration statement are to be activated on the requested site if specified, or on both sites.

Rule: You must have at least one active Agent on the requested target sites.

Guideline: `*ALL*` is deprecated and, if specified, is processed like `*ALL_AS*`.

Site=*sitename*

Specifies the name of the site where the workload is to be distributed. The name must match one of the site names specified in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement. The *sitename* value must not exceed 8 characters in length.

For an Active/Query workload, the Site keyword is optional. If not specified, the workload is activated on both sites.

Examples

The following example shows a sample display when you activate an Active/Standby workload by issuing F AQSADV,ACTIVATE,WORKLOAD=WORKLOAD1,SITE=PLEX1 on the primary Advisor.

```
F AQSADV,ACTIVATE,WORKLOAD=WORKLOAD1,SITE=PLEX1
AQS0153I WORKLOAD WORKLOAD1 ACTIVATED ON SITE PLEX1
```

The following example shows a sample display when you activate an Active/Query workload by issuing F AQSADV,ACTIVATE,WORKLOAD=WORKLOAD2 on the primary Advisor.

```
F AQSADV,ACTIVATE,WORKLOAD=WORKLOAD2
AQS0155I WORKLOAD WORKLOAD2 ACTIVATED
```

The following example shows a sample display when you activate all Active/Standby workloads by issuing F AQSADV,ACTIVATE,WORKLOAD=*ALL_AS*,SITE=PLEX1 on the primary advisor.

```
F AQSADV,ACTIVATE,WORKLOAD=*ALL_AS*,SITE=PLEX1
AQS0201I ALL MATCHING WORKLOADS ACTIVATED ON SITE PLEX1
```


MODIFY *advisor_proc*,DEACTivate command: Deactivate workloads

Use the `MODIFY advisor_proc,DEACTivate` command to deactivate a workload. For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, force the Lifeline Advisor to reset all the existing connections for the specified workload or for all workloads. For workloads defined in the `mq_manager_list` Advisor configuration statement, force the Lifeline Advisor

to signal Lifeline Agents on the previously active site to verify each MQ cluster receiver channel is stopped and to check for existing messages on each MQ cluster queue in the workload.

Rule: You can issue this command only from the primary Advisor.

Format



Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

DEACTivate

Signals the Advisor that a workload needs to be deactivated.

WORKload=

workloadname

Specifies the name of the workload to be deactivated. The name must match one of the workload names specified in the *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* Advisor configuration statement. The *workloadname* value must not exceed 64 characters in length.

Rule: You must quiesce the specified workload on the target site before you issue the command to deactivate it.

ALL_AS

Indicates that all the Active/Standby workloads specified in the *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* Advisor configuration statement are to be deactivated.

Rule: You must quiesce all the Active/Standby workloads on the target site before you issue the command to deactivate them.

ALL_AQ

Indicates that all the Active/Query workloads specified in the *cross_sysplex_list* Advisor configuration statement are to be deactivated on the requested site if specified, or on both sites.

Rule: You must quiesce all the Active/Query workloads on the target sites before you issue the command to deactivate them.

Guideline: **ALL** is deprecated and, if specified, is processed like **ALL_AS**.

Site=*sitename*

Specifies the name of the site where the Active/Query workload is to be deactivated. The name must match one of the site names specified in the

cross_sysplex_list Advisor configuration statement. The *sitename* value must not exceed 8 characters in length. The Site keyword is optional. If not specified, the workload is deactivated on both sites.

For an Active/Standby workload, the Site keyword is not allowed. The workload is deactivated on the site where the workload was last quiesced.

Examples

The following example shows a sample display when you deactivate an Active/Standby workload by issuing F AQSADV,DEACTIVATE,WORKLOAD=WORKLOAD1 on the primary Advisor.

```
F AQSADV,DEACTIVATE,WORKLOAD=WORKLOAD1
AQS0155I WORKLOAD WORKLOAD1 DEACTIVATED
```

The following example shows a sample display when you deactivate an Active/Query workload by issuing F AQSADV,DEACTIVATE,WORKLOAD=WORKLOAD2 on the primary Advisor.

```
F AQSADV,DEACTIVATE,WORKLOAD=WORKLOAD2
AQS0155I WORKLOAD WORKLOAD2 DEACTIVATED
```

The following example shows a sample display when you issue F AQSADV,DEACTIVATE,WORKLOAD=*ALL_AS* on the primary Advisor.

```
F AQSADV,DEACTIVATE,WORKLOAD=*ALL_AS*
AQS0202I ALL MATCHING WORKLOADS DEACTIVATED
```

MODIFY *advisor_proc*,DEBUG command: Change the debug level

Use the MODIFY *advisor_proc*,DEBUG command to change the debug level that is in effect for the Lifeline Advisor.

Format

```
►►—[MODIFY]—advisor_proc—,DEBUG—,Level=debuglevel—►►
    |
    |—F—
```

Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

DEBUG,Level=*debuglevel*

Changes the debug level in effect for this Advisor. The debug level that you can set is the sum of one or more levels of the debug messages to be logged. See “Debug levels” on page 67 for information about debug level values.

Example

The following example shows a sample display when you issue F AQSADV,DEBUG,LEVEL=127 on the Advisor.

```
F AQSADV,DEBUG,LEVEL=127
AQS0138I LLADVSR DEBUG LEVEL UPDATED
```

MODIFY *advisor_proc*,DUMPLog command: Dump the internal debug records

Use the MODIFY *advisor_proc*,DUMPLog command to dump the internal trace to the UNIX System Services file specified by the AQSTRACE DD card in the Advisor started procedure.

Format

►►—MODIFY—*advisor_proc*—,DUMPLog—◄◄
 └─┬─┘
 F

Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

DUMPLog

Signals the Advisor that the internal trace records need to be dumped. Any existing data in the target UNIX System Services file will be overwritten.

Example

The following example shows a sample display when you issue F AQSADV,DUMPLOG on the Advisor.

```
F AQSADV,DUMPLOG
AQS0204I LLADVSR DUMP INTERNAL LOG COMMAND SUCCESSFUL
```

MODIFY *advisor_proc*,QUIESCE command: Quiesce workloads


Use the MODIFY *advisor_proc*,QUIESCE command to quiesce a workload. For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, force the Lifeline Advisor to signal the first-tier load balancers that no new connections are to be routed for the specified workload or for all workloads. For workloads defined in the *mq_manager_list* Advisor configuration statement, force the Lifeline Advisor to signal Lifeline Agents to change the queue rank of each MQ cluster queue in the workload to a '0' and stop the workload's MQ cluster receiver channel.

Notes:

- For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, existing connections for this workload will continue to be routed.
- For workloads defined in the *mq_manager_list* Advisor configuration statement, no new messages will be sent to the MQ queue managers that host the MQ cluster queues for the workload. MQ server applications will continue to be able to process messages from the MQ cluster queues for the workload.

Rule: You can issue this command only from the primary Advisor.

Format



Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

QUIESCE

Signals the first-tier load balancers that a workload can no longer be distributed to a site.

WORKload=

workloadname

Specifies the name of the workload to be quiesced. The name must match one of the workload names specified in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement. The *workloadname* value must be less than or equal to 64 characters in length.

ALL_AS

Indicates that all the Active/Standby workloads specified in the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` Advisor configuration statement are to be quiesced.

ALL_AQ

Indicates that all the Active/Query workloads specified in the `cross_sysplex_list` Advisor configuration statement are to be quiesced on the requested site if specified, or on both sites.

Guideline: *ALL* is deprecated and, if specified, is processed like *ALL_AS*.

Site=*sitename*

Specifies the name of the site where the Active/Query workload is to be quiesced. The name must match one of the site names specified in the `cross_sysplex_list` Advisor configuration statement. The *sitename* value must not exceed 8 characters in length. The Site keyword is optional. If not specified, the workload is quiesced on both sites.

For an Active/Standby workload, the Site keyword is not allowed. The workload is quiesced on the site where it was previously activated.

Examples

The following example shows a sample display when you quiesce an Active/Standby workload by issuing `F AQSADV,QUIESCE,WORKLOAD=WORKLOAD1` on the primary Advisor.

```

  F AQSADV,QUIESCE,WORKLOAD=WORKLOAD1
  AQS0155I WORKLOAD WORKLOAD1 QUIESCED
  
```

The following example shows a sample display when you quiesce an Active/Query workload by issuing F AQSADV,QUIESCE,WORKLOAD=WORKLOAD2 on the primary Advisor.

```
F AQSADV,QUIESCE,WORKLOAD=WORKLOAD2
AQS0155I WORKLOAD WORKLOAD2 QUIESCED
```

The following example shows a sample display when you quiesce all Active/Standby workloads by issuing F AQSADV,QUIESCE,WORKLOAD=*ALL_AS* on the primary Advisor.

```
F AQSADV,QUIESCE,WORKLOAD=*ALL_AS*
AQS0202I ALL MATCHING WORKLOADS QUIESCED
```

MODIFY *advisor_proc*,REFRESH command: Update configuration

Use the MODIFY *advisor_proc*,REFRESH command to force the Lifeline Advisor to read the configuration file again and apply any changes to its active configuration.

See “Lifeline Advisor configuration statements” on page 85 for the configuration statements that are refreshable.

Format

```
►►—MODIFY—advisor_proc—,REFRESH—►►
   └─┬──┘
     F
```

Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

REFRESH

Forces this Advisor to read the configuration file again and apply any changes to its active configuration.

Example

The following example shows a sample display when you issue F AQSADV,REFRESH.

```
F AQSADV,REFRESH
AQS0181I LLADVSR CONFIGURATION UPDATED
```

MODIFY *advisor_proc*,TAKEOVER command: Take over primary Advisor responsibilities

Use the MODIFY *advisor_proc*,TAKEOVER command to force the Lifeline Advisor to take over primary Advisor responsibilities.

Rule: You can issue this command only from the secondary Advisor.

Format

```
►►—MODIFY—advisor_proc—,TAKEOVER—►►
   └─┬──┘
     F
```

Parameters

advisor_proc

Indicates the member name of the cataloged procedure used to start the Advisor.

TAKEOVER

Forces this Advisor to takeover primary Advisor responsibilities.

Example

The following example shows a sample display when you issue F AQSADV,TAKEOVER on the secondary Advisor.

```
F AQSADV,TAKEOVER
AQS0169I ADVISOR TAKEOVER SCHEDULED
AQS0166I ADVISOR IS NOW IN PRIMARY ROLE
```


Chapter 8. Lifeline z/OS Agent commands

Use the Lifeline z/OS Agent commands to display Lifeline z/OS Agent information and modify Lifeline z/OS Agent state information.

MODIFY *agent_proc*,DISplay command: Display Lifeline Agent information

Use the MODIFY *agent_proc*,DISplay command to display Lifeline Agent information.

Format



Parameters

agent_proc

Indicates the member name of the cataloged procedure used to start the Agent.

DISplay,CONFig

Displays the active configuration for this Agent.

DISplay,DEBug

Displays the debug level in effect for this Agent.

The debug level displayed is the sum of one or more levels of debug messages being logged. See “Debug levels” on page 67 for more information about the debug levels.

DISplay,MEMbers

Displays information about registered local server applications.

DISplay,MEMbers,DETail

Displays detailed information about registered local server applications, including the job name and ASID of the server application.

DISplay,MEMbers,DETail,PORT=portnum

Displays detailed information about registered local server applications that are bound to the specified port number. The valid range of the *portnum* is 0 - 65535.

DISplay,MEMbers,DETail,TCPname=tcpname

Displays detailed information about registered local server applications that are associated with the specified TCP/IP stack. The *tcpname* value must be less than or equal to 8 characters in length.

Examples

The following example shows a sample display when you issue F AQSAGE,DISPLAY,CONFIG.

```

F AQSAGE,DISPLAY,CONFIG
AQS0144I AGENT CONFIGURATION
VERSION          : LIFELINE V2.5
CONFIG FILENAME  : USER.PROFILES(LBAGECNF)
DEBUG LEVEL      : 7
ADVISOR ID       : 192.10.1.1..8181
HOST CONNECTION  : 192.10.110.1..4000
SITE NAME        : PLEX1

```

The following information describes the fields in the previous sample display.

Field name

Descriptions

VERSION

The current version and release of the Agent.

CONFIG FILENAME

The configuration file used by this Agent when the Agent was initially started.

DEBUG LEVEL

The debug level currently in effect for this Agent.

ADVISOR ID

The IP address and port number that the primary Advisor listens on for connections from this Agent. The IP address can be an IPv4 or an IPv6 address.

HOST CONNECTION

The source IP address and port number that this Agent uses to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

SITE NAME

The site where the Agent resides. If the site_name statement is not configured on a z/OS Agent, this information will not be displayed.

The following example shows a sample display when you issue F AQSAGE,DISPLAY,DEBUG.

```

F AQSAGE,DISPLAY,DEBUG
AQS0111I LLAGENT DEBUG LEVEL 7

```

The following example shows a sample display when you issue F AQSAGE,DISPLAY,MEMBERS.

```

F AQSAGE,DISPLAY,MEMBERS
AQS0114I MEMBER SUMMARY
LB INDEX      : 00          UUID       : A67B6699
GROUP NAME    : WKLD2_GROUP1
IPADDR..PORT : 10.10.1.1..8020
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        : ANY DISTDVIPA
GROUP NAME    : WKLD2_GROUP2
IPADDR..PORT : 10.10.1.1..8021
MATCHES      : 000          PROTOCOL   : TCP
FLAGS        : DISTDVIPA
LB INDEX      : 01          UUID       : 9A78BE9E
GROUP NAME    : TIER2_GROUP1
IPADDR..PORT : 192.10.110.1..5001
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        :
GROUP NAME    : TIER2_GROUP2
IPADDR..PORT : 192.10.110.1..6001
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        :

```



```

GROUP NAME      : SNAGROUP1
SNA APPL       : NETA.SNAAPPL1  AVAIL : YES
MQ CLUSTER     : CLUSTER1
GROUP NAME     : CLUSTER1-QUEUE1
IPADDR..PORT   : 10.10.1.1..1414
MATCHES       : 001             PROTOCOL : TCP
QMGR NAME     : QMG1
6 OF 6 RECORDS DISPLAYED

```

The following information describes the fields in the previous sample display.

Field name
Descriptions

LB INDEX

The reference number used to identify the load balancer. The same reference number is used for a load balancer as long as the load balancer remains connected.

UUID A hexadecimal value of the universally unique identifier assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

MQ CLUSTER

The name of an IBM MQ Cluster that is hosted by an MQ queue manager on this z/OS system.

GROUP NAME

For load balancer application groups, the name of a group of related target server applications registered on the external load balancer.

For MQ cluster queue groups, the name of the group that the Advisor creates for each MQ cluster queue defined for the workload. The group name is the configured MQ cluster name, followed by a hyphen, "-", followed by the configured MQ cluster queue name.

IPADDR..PORT

For load balancer application groups, the IP address and port number of a member defined in the registered group.

For MQ cluster queue groups, the IP address and port number that the MQ queue manager is listening on.

The IP address can be an IPv4 or an IPv6 address.

MATCHES

For load balancer application groups, the number of server applications that match the IP address and port number. For the server applications sharing a port, this value can be more than 1. If the value of matches is 0, the Agent finds the IP address of the member reported on an active TCP/IP stack, but does not find a server application listening on that IP address or port number.

For MQ cluster queue groups, matches is 0 if the group's IBM MQ cluster or MQ cluster queue is not defined on the MQ queue manager. The IBM MQ clusters and MQ cluster queues defined on each MQ queue manager are configured on the Advisor in the mq_manager_list configuration statement.

PROTOCOL

The protocol that the member is using. The protocol value is TCP or UDP.

|
| **QMGR NAME**
|

The name of an MQ queue manager hosting the IBM MQ cluster and MQ cluster queue mapped by the group name.

FLAGS

Indicators for the member. The flag values are shown as follows:

ANY Indicates that the server application is bound to INADDR_ANY or the unspecified IPv6 address.

DISTDVIPA

Indicates that the IP address of the member is a distributable DVIPA. The IP address can be an IPv4 or an IPv6 address.

NODATA

Indicates that the Agent is reporting a temporary value for the Communications Server weight. Two update intervals are needed before the Agent can calculate a proper value for the Communications Server weight.

SDQ

Indicates that an operator at the owning Agent quiesced the sysplex distributed DVIPA server application's port, jobname, or TCP/IP stack with a VARY TCPIP,,SYSPLEX,QUIESCE command. No new workload connections are routed to this server application. Use the VARY TCPIP,,SYSPLEX,RESUME command to allow new workload connections to be routed to this server application.

SYSQ, TCPQ, or APPQ

Indicates that the operator has requested that no more work be routed to the quiesced server application. The different flags reflect the highest level of the quiesce command that applies, and indicates the type of the enable command that must be used to enable the server application.

| **SYSQ** Indicates that the server application was quiesced with the
| MODIFY *agent_proc*, QUIESCE, SYSTEM command, and that
| you must use the MODIFY *agent_proc*, ENABLE, SYSTEM
| command to enable the server application.

| **TCPQ** Indicates that the server application was quiesced with the
| MODIFY *agent_proc*, QUIESCE, TCPNAME=*tcpname* command,
| and that you must use the MODIFY
| *agent_proc*, ENABLE, TCPNAME=*tcpname* command to enable
| the server application.

| **APPQ** Indicates that the server application was quiesced with the
| MODIFY *agent_proc*, QUIESCE, PORT=*portnum* command, and
| that you must use the MODIFY
| *agent_proc*, ENABLE, PORT=*portnum* command to enable the
| server application.

V6

Indicates that the server application has specified the IPv6_V6ONLY socket option. Only workload requests from IPv6-enabled clients are able to communicate with the server application.

SNA APPL

The SNA application name.

AVAIL

Indicates whether the SNA application is available to process

workload requests. The value YES indicates that the application is available. The value NO indicates that the application is not available.

The following example shows a sample display when you issue F AQSAGE,DISPLAY,MEMBERS,DETAIL.

```
F AQSAGE,DISPLAY,MEMBERS,DETAIL
AQS0115I MEMBER DETAILS
LB INDEX      : 00      UUID      : A67B6699
GROUP NAME    : WKLD2_GROUP1
IPADDR..PORT : 10.10.1.1..8020
MATCHES      : 001      PROTOCOL   : TCP
FLAGS        : ANY DISTDVIPA
TCPNAME      : TCPIP
JOBNAME      : JOB1      ASID       : 0034 RESOURCE  : 0000096B
GROUP NAME    : WKLD2_GROUP2
IPADDR..PORT : 10.10.1.1..8021
MATCHES      : 000      PROTOCOL   : TCP
FLAGS        : DISTDVIPA
TCPNAME      : TCPIP
JOBNAME      : N/A      ASID       : N/A RESOURCE  : N/A
LB INDEX      : 01      UUID      : 9A78BE9E
GROUP NAME    : TIER2_GROUP1
IPADDR..PORT : 192.10.110.1..5001
MATCHES      : 001      PROTOCOL   : TCP
FLAGS        :
TCPNAME      : TCPIP
JOBNAME      : JOB3      ASID       : 0036 RESOURCE  : 0000096D
GROUP NAME    : TIER2_GROUP2
IPADDR..PORT : 192.10.110.1..6001
MATCHES      : 001      PROTOCOL   : TCP
FLAGS        :
TCPNAME      : TCPIP
JOBNAME      : JOB4      ASID       : 0037 RESOURCE  : 0000096E
GROUP NAME    : SNAGROUP1
SNA APPL     : NETA.SNAAPPL1          AVAIL : YES  JOBNAME : JOB5
SESSION COUNT : 0000000001
MQ CLUSTER   : CLUSTER1
GROUP NAME    : CLUSTER1-QUEUE1
IPADDR..PORT : 10.10.1.1..1414
MATCHES      : 001      PROTOCOL   : TCP
QMGR NAME    : QMG1
FLAGS        : ANY
TCPNAME      : TCPIP
JOBNAME      : AQSICHIN ASID       : 003D RESOURCE  : 0000040E
6 OF 6 RECORDS DISPLAYED
```

The following information describes the fields in the previous sample display.

Field name
Descriptions

LB INDEX

Reference number used to identify the load balancer. The same reference number is used for a load balancer as long as the load balancer remains connected.

UUID

A hexadecimal value of the universally unique identifier assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

MQ CLUSTER

The name of an IBM MQ Cluster hosted by an MQ queue manager on this z/OS system.

GROUP NAME

For load balancer application groups, the name of a group of related target server applications registered on the external load balancer.

For MQ cluster queue groups, the name of the group the Advisor creates for each MQ cluster defined for the workload. The group name is the configured MQ cluster name, followed by a hyphen, "-", followed by the configured MQ cluster queue name.

IPADDR..PORT

For load balancer application groups, the IP address and port number of a member defined in the registered group.

For MQ cluster queue groups, the IP address and port number that the MQ queue manager is listening on.

The IP address can be an IPv4 or an IPv6 address.

MATCHES

For load balancer application groups, the number of server applications that match the IP address and port number. For server applications sharing a port, this value can be more than 1. If the value of matches is 0, the Agent found the IP address of the member reported on an active TCP/IP stack, but did not find a server application listening on that IP address or port number.

For MQ cluster queue groups, matches is 0 if the group's IBM MQ cluster or MQ cluster queue is not defined on the MQ queue manager. The IBM MQ clusters and MQ cluster queues defined on each MQ queue manager are configured on the Advisor in the mq_manager_list configuration statement.

PROTOCOL

The protocol that the member is using. The protocol value will be TCP or UDP.

QMGR NAME

The name of an MQ queue manager hosting the IBM MQ cluster and MQ cluster queue mapped by the group name.

FLAGS

Indicators for the member. The flag values are shown as follows:

ANY Indicates that the server application is bound to INADDR_ANY or the unspecified IPv6 address.

DISTDVIPA

Indicates that the IP address of the member is a distributable DVIPA. The IP address can be an IPv4 or an IPv6 address.

NODATA

Indicates that the Agent is reporting a temporary value for the Communications Server weight. Two update intervals are needed before the Agent can calculate a proper value for the Communications Server weight.

SYSQ, TCPQ, or APPQ

Indicates that the operator has requested that no more work be routed to the quiesced server application. The different flags reflect

the highest level of the quiesce command that applies, and indicates the type of the enable command that must be used to enable the server application.

SYSQ Indicates that the server application was quiesced with the MODIFY *agent_proc*, QUIESCE, SYSTEM command, and that you must use the MODIFY *agent_proc*, ENABLE, SYSTEM command to enable the server application.

TCPQ Indicates that the server application was quiesced with the MODIFY *agent_proc*, QUIESCE, TCPNAME=*tcpname* command, and that you must use the MODIFY *agent_proc*, ENABLE, TCPNAME=*tcpname* command to enable the server application.

APPQ Indicates that the server application was quiesced with the MODIFY *agent_proc*, QUIESCE, PORT=*portnum* command, and that you must use the MODIFY *agent_proc*, ENABLE, PORT=*portnum* command to enable the server application.

V6 Indicates that the server application has specified the IPv6_V6ONLY socket option. Only workload requests from IPv6-enabled clients are able to communicate with the server application.

TCPNAME

The name of the TCP/IP Server stack that owns the IP address for the member.

JOBNAME

The MVS job name of the server application. If there are no server applications matching this member, N/A is displayed.

ASID The MVS address space identifier of the server application. If there are no server applications matching this member, N/A is displayed.

RESOURCE

The identifier used by the TCP/IP stack that uniquely identifies one instance of the server application. The resource identifier is also displayed by the NETSTAT CONN command. If there are no server applications matching this member, N/A is displayed.

SNA APPL

The SNA application name.

AVAIL

Indicates whether the SNA application is available to process workload requests. The value YES indicates that the SNA application is available. The value NO indicates that the SNA application is not available.

JOBNAME

The MVS job name of the SNA application. If no SNA applications match this member, N/A is displayed.

SESSION COUNT

The number of active sessions for the SNA application.

The following example shows a sample display when you issue F AQSAGE, DISPLAY, MEMBERS, DETAIL, PORT=5001.

```

F AQSAGE,DISPLAY,MEMBERS,DETAIL,PORT=5001
AQS0115I MEMBER DETAILS
LB INDEX      : 01          UUID       : 9A78BE9E
GROUP NAME    : TIER2_GROUP1
IPADDR..PORT: 192.10.110.1..5001
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        :
TCPNAME      : TCPIP
JOBNAME      : JOB3          ASID       : 0036 RESOURCE : 0000096D
1 OF 1 RECORDS DISPLAYED

```

The following example shows a sample display when you issue F AQSAGE,DISPLAY,MEMBERS,DETAIL,TCPNAME=TCPIP.

```

F AQSAGE,DISPLAY,MEMBERS,DETAIL,TCPNAME=TCPIP
AQS0115I MEMBER DETAILS
LB INDEX      : 00          UUID       : A67B6699
GROUP NAME    : WKLD2_GROUP1
IPADDR..PORT: 10.10.1.1..8020
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        : ANY DISTDVIPA
TCPNAME      : TCPIP
JOBNAME      : JOB1          ASID       : 0034 RESOURCE : 0000096B
GROUP NAME    : WKLD2_GROUP2
IPADDR..PORT: 10.10.1.1..8021
MATCHES      : 000          PROTOCOL   : TCP
FLAGS        : DISTDVIPA
TCPNAME      : TCPIP
JOBNAME      : N/A          ASID       : N/A RESOURCE : N/A
LB INDEX      : 01          UUID       : 9A78BE9E
GROUP NAME    : TIER2_GROUP1
IPADDR..PORT: 192.10.110.1..5001
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        :
TCPNAME      : TCPIP
JOBNAME      : JOB3          ASID       : 0036 RESOURCE : 0000096D
GROUP NAME    : TIER2_GROUP2
IPADDR..PORT: 192.10.110.1..6001
MATCHES      : 001          PROTOCOL   : TCP
FLAGS        :
TCPNAME      : TCPIP
JOBNAME      : JOB4          ASID       : 0037 RESOURCE : 0000096E
GROUP NAME    : SNAGROUP1
SNA APPL     : NETA.SNAAPPL1          AVAIL : YES  JOBNAME : JOB5
SESSION COUNT : 0000000001
MQ CLUSTER   : CLUSTER1
GROUP NAME    : CLUSTER1-QUEUE1
IPADDR..PORT: 10.10.1.1..1414
MATCHES      : 001          PROTOCOL   : TCP
QMGR NAME    : QMG1
FLAGS        : ANY
TCPNAME      : TCPIP
JOBNAME      : AQS1CHIN ASID       : 003D RESOURCE : 0000040E
6 OF 6 RECORDS DISPLAYED


```

MODIFY *agent_proc*,ENABLE command: Enable server applications

Use the MODIFY *agent_proc*,ENABLE command to mark the server applications as enabled. The Lifeline Agent will notify the Advisor to recommend this server application for new workload requests.

Tip: If more than one server application is sharing a port, all of the server applications are enabled.

Format



Parameters

agent_proc

Indicates the member name of the cataloged procedure used to start the Agent.

ENABLE,PORT=*portnum*

Marks the server applications bound to the specified port number as enabled. The range of the *portnum* value is 0 - 65535.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can enable the server applications only by the VARY TCPIP,,SYSPLEX,RESUME,PORT=*portnum* command.

ENABLE,PORT=*portnum*,PROTOcol=*protocol*

Marks the server applications that are bound to the specified port number and use the specified protocol as enabled. The range of the *portnum* value is 0 - 65535. The *protocol* value is TCP or UDP.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can enable the server applications only by the VARY TCPIP,,SYSPLEX,RESUME,PORT=*portnum* command.

ENABLE,PORT=*portnum*,PROTOcol=*protocol*,IPaddr=*ipaddr*

Marks the server applications that are bound to the specified port number and use the specified protocol and the specified IP address as enabled. The range of the *portnum* value is 0 - 65535. The *protocol* value is TCP or UDP. The IP address can be an IPv4 or an IPv6 address.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can enable the server applications only by the VARY TCPIP,,SYSPLEX,RESUME,JOBNAME=*jobname* command.

ENABLE,TCPname=*tcpname*

Marks the server applications associated with the specified TCP/IP stack as enabled. The *tcpname* value must be less than or equal to 8 characters in length.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can enable the server applications only by the VARY TCPIP,,SYSPLEX,RESUME,TARGET command.

ENABLE,SYStem

Marks the server applications on this system as enabled.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can enable the server applications only by the VARY TCPIP,,SYSPLEX,RESUME,TARGET command.

Examples

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=5001.

```
F AQSAGE,ENABLE,PORT=5001
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE ENABLED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=8021, but the server application was registered for a distributable DVIPA.

```
F AQSAGE,ENABLE,PORT=8021
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE ENABLED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP.

```
F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE ENABLED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=8021,PROTOCOL=TCP, but the server application was registered for a distributable DVIPA.

```
F AQSAGE,ENABLE,PORT=8021,PROTOCOL=TCP
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE ENABLED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=5001,PROTOCOL=UDP, but no matching server applications are listening on this port or protocol.

```
F AQSAGE,ENABLE,PORT=5001,PROTOCOL=UDP
AQS0122I NO MEMBERS MATCH SELECTION CRITERIA
AQS0123I ENABLE APPLICATION COMMAND REJECTED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP,IPADDR=192.10.110.1.

```
F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP,IPADDR=192.10.110.1
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE ENABLED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=8021,PROTOCOL=TCP,IPADDR=10.10.1.1, but the server application was registered for a distributable DVIPA..

```
F AQSAGE,ENABLE,PORT=8021,PROTOCOL=TCP,IPADDR=10.10.1.1
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE ENABLED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP,IPADDR=10.10.1.2, but no matching server applications are listening on this port or protocol.

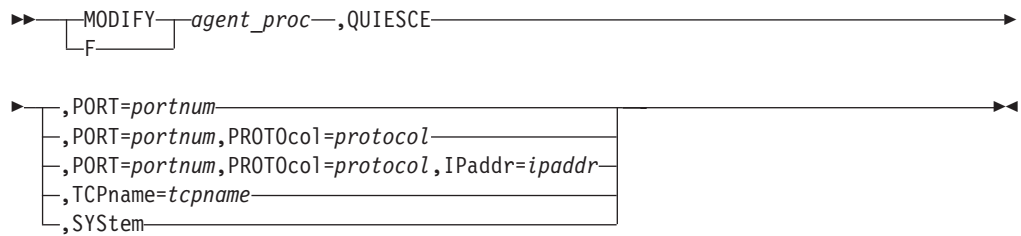
```
F AQSAGE,ENABLE,PORT=5001,PROTOCOL=TCP,IPADDR=10.10.1.2
AQS0122I NO MEMBERS MATCH SELECTION CRITERIA
AQS0123I ENABLE APPLICATION COMMAND REJECTED
```

MODIFY *agent_proc*, QUIESCE command: Quiesce server applications

Use the MODIFY *agent_proc*, QUIESCE command to mark the server applications as quiesced. The Lifeline Agent will notify the Advisor to not recommend this server application for new workload requests.

Tip: If more than one server application is sharing a port, all of the server applications are quiesced.

Format



Parameters

agent_proc

Indicates the member name of the cataloged procedure used to start the Agent.

QUIESCE, PORT=portnum

Marks the server applications bound to the specified port number as quiesced. The range of the *portnum* value is 0 - 65535.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can quiesce the server applications only by the VARY TCPIP, ,SYSPLEX, QUIESCE, PORT=portnum command.

QUIESCE, PORT=portnum, PROTOcol=protocol

Marks the server applications that are bound to the specified port number and use the specified protocol as quiesced. The range of the *portnum* value is 0 - 65535. The *protocol* value is TCP or UDP.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can quiesce the server applications only by the VARY TCPIP, ,SYSPLEX, QUIESCE, PORT=portnum command.

QUIESCE, PORT=portnum, PROTOcol=protocol, IPAddr=ipaddr

Marks the server applications that are bound to the specified port number and use the specified protocol and the specified IP address as quiesced. The range of the *portnum* value is 0 - 65535. The *protocol* value is TCP or UDP. The IP address can be an IPv4 or an IPv6 address.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can quiesce the server applications only by the VARY TCPIP, ,SYSPLEX, QUIESCE, JOBNAME=jobname command.

QUIESCE, TCPname=tcpname

Marks the server applications associated with the specified TCP/IP stack as quiesced. The *tcpname* value must be less than or equal to 8 characters in length.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can quiesce the server applications only by the VARY TCPIP, ,SYSPLEX, QUIESCE, TARGET command.

QUIESCE, SYSTEM

Marks the server applications on this system as quiesced.

Rule: If the IP address of the member is a distributable DVIPA when you specify this parameter, you can quiesce the server applications only by the VARY TCP/IP, ,SYSPLEX,QUIESCE,TARGET command.

Examples

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=5001.

```
F AQSAGE,QUIESCE,PORT=5001
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE QUIESCED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=8021, but the server application was registered for a distributable DVIPA.

```
F AQSAGE,QUIESCE,PORT=8021
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE QUIESCED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP.

```
F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE QUIESCED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=8021,PROTOCOL=TCP, but the server application was registered for a distributable DVIPA.

```
F AQSAGE,QUIESCE,PORT=8021,PROTOCOL=TCP
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE QUIESCED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=UDP, but no matching server applications are listening on this port or protocol.

```
F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=UDP
AQS0122I NO MEMBERS MATCH SELECTION CRITERIA
AQS0123I QUIESCE APPLICATION COMMAND REJECTED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP,IPADDR=192.10.110.1.

```
F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP,IPADDR=192.10.110.1
AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE QUIESCED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=8021,PROTOCOL=TCP,IPADDR=10.10.1.1, but the server application was registered for a distributable DVIPA.

```
F AQSAGE,QUIESCE,PORT=8021,PROTOCOL=TCP,IPADDR=10.10.1.1
AQS0140I 0 OF 1 MATCHING REGISTERED APPLICATIONS ARE QUIESCED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP,IPADDR=10.10.1.2, but no matching server applications are listening on this port or protocol.

```
F AQSAGE,QUIESCE,PORT=5001,PROTOCOL=TCP,IPADDR=10.10.1.2
AQS0122I NO MEMBERS MATCH SELECTION CRITERIA
AQS0123I QUIESCE APPLICATION COMMAND REJECTED
```

The following example shows a sample display when you issue F AQSAGE, QUIESCE, TCPNAME=TCPIP.

```
F AQSAGE, QUIESCE, TCPNAME=TCPIP
AQS0140I 2 OF 4 MATCHING REGISTERED APPLICATIONS ARE QUIESCED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

The following example shows a sample display when you issue F AQSAGE, QUIESCE, SYSTEM, but two of the server applications were registered for a distributable DVIPA.

```
F AQSAGE, QUIESCE, SYSTEM
AQS0140I 2 OF 4 MATCHING REGISTERED APPLICATIONS ARE QUIESCED -
DISTRIBUTED DVIPA MEMBERS ARE IGNORED
```

MODIFY *agent_proc*, DEBbug command: Change the debug level

Use the MODIFY *agent_proc*, DEBbug command to change the debug level that is in effect for the Lifeline Agent.

Format

```
➤—[MODIFY]—agent_proc—, DEBbug—[ ,Level=debuglevel]—➤
```

Parameters

agent_proc

Indicates the member name of the cataloged procedure used to start the Agent.

DEBbug, Level=*debuglevel*

Changes the debug level in effect for this Agent. The debug level that you can set is the sum of one or more levels of the debug messages to be logged. See “Debug levels” on page 67 for debug level values.

Example

The following example shows a sample display when you issue F AQSAGE, DEBUG, LEVEL=127.

```
F AQSAGE, DEBUG, LEVEL=127
AQS0138I LLAGENT DEBUG LEVEL UPDATED
```

MODIFY *agent_proc*, DUMPLog command: Dump the internal trace records

Use the MODIFY *agent_proc*, DUMPLog command to dump internal trace to the UNIX System Services file specified by the AQSTRACE DD card in the Agent started procedure.

Format

```
➤—[MODIFY]—agent_proc—, DUMPLog—➤
```

Parameters

agent_proc

Indicates the member name of the cataloged procedure used to start the Agent.

DUMPLog

Signals the Agent that the internal trace records need to be dumped. Any existing data in the target UNIX System Services file will be overwritten.

Example

The following example shows a sample display when you issue F AQSAGE,DUMPLOG.

```
F AQSAGE,DUMPLOG
AQS0204I LLAGENT DUMP INTERNAL LOG COMMAND SUCCESSFUL
```

Chapter 9. Lifeline Linux on System z Management Agent commands

Use the Linux on System z Management Agent commands to display Lifeline Linux on System z Management Agent information and modify Linux on System z Management Agent state information.

Use the `aqslcmd` program to issue commands to the Linux on System z Management Agent. You must run the program on the Management Guest where the Linux on System z Management Agent runs and only by a user with root privileges.

Format

`aqslcmd command`

command is one of the following options:

```
▶▶—DISplay, —CONFig—————▶▶
      | —DEBUg—————▶▶
      | —MEMbers—————▶▶
      | —MEMbers,DETail————▶▶
      | —MEMbers,DETail,PORT=portnum————▶▶
      | —MEMbers,DETail,TCpname=guestname——▶▶
      |
▶▶—DEBUg, —,Level=debuglevel——▶▶
      |
▶▶—DUMPLog—————▶▶
      |
▶▶—STOP—————▶▶
```

Guideline: The value for command is not case sensitive.


A man page is available for the `aqslcmd` command. To display the man page, enter the following command:

```
man -M /opt/ibm/lifeline/man aqslcmd
```

Display Linux on System z Management Agent information

Use the `aqslcmd DISplay` command to display Linux on System z Management Agent information. Use the `MODIFY agent_proc,DISplay` command to display z/OS Agent information.

Format



Parameters

DISplay,CONFIg

Displays the active configuration for this Agent.

DISplay,DEBUg

Displays the debug level in effect for this Agent.

The displayed debug level is the sum of one or more levels of debug messages being logged. See “Debug levels” on page 67 for more information about the debug levels.

DISplay,MEMbers

Displays information about registered server applications on this z/VM system.

DISplay,MEMbers,DETAil

Displays detailed information about registered server applications on this z/VM system, including the host name of the server application.

DISplay,MEMbers,DETAil,PORT=portnum

Displays detailed information about registered server applications on this z/VM system that are bound to the specified port number. The valid range of the *portnum* is 0 - 65535.

DISplay,MEMbers,DETAil,TCPname=guestname

Displays detailed information about registered server applications on this z/VM system that are associated with the Linux on System z guest, *guestname*. The *guestname* value must be less than or equal to 8 characters in length and is not case sensitive.

Examples

The following example shows a sample display when you issue `aqs11cmd display,config`.

```
aqs11cmd display,config
AQS0144I AGENT CONFIGURATION
VERSION          : Lifeline V2.5
CONFIG FILENAME  : /opt/ibm/lifeline/conf/agent.conf
DEBUG LEVEL     : 7
ADVISOR ID      : 192.10.1.1..8181
HOST CONNECTION  : 192.10.110.3 ..4000
SITE NAME       : PLEX1
Command complete
```

The following information describes the fields in the sample display.

Field name

Descriptions

VERSION

The current version and release of the Agent.

CONFIG FILENAME

The configuration file that this Agent used when the Agent was initially started.

DEBUG LEVEL

The debug level currently in effect for this Agent.

ADVISOR ID

The IP address and port number that the primary Advisor listens on for connections from this Agent. The IP address can be an IPv4 or an IPv6 address.

HOST CONNECTION

The source IP address and port number that this Agent uses to connect to the primary Advisor. The IP address can be an IPv4 or an IPv6 address.

SITE NAME

The site where the Agent resides.

The following example shows a sample display when you issue `aqsllcmd display,debug`.

```
aqsllcmd display,debug
AQS0111I LLAGENT DEBUG LEVEL 7
Command complete
```

The following example shows a sample display when you issue `aqsllcmd display,members`.

```
aqsllcmd display,members
AQS0114I MEMBER SUMMARY
LB INDEX      : 00          UUID      : 8F659238
GROUP NAME    : WKLD3_GROUP1
IPADDR..PORT : 10.1.1.1..9001
MATCHES      : 001        PROTOCOL   : TCP
1 of 1 RECORDS DISPLAYED
Command complete
```

The following information describes the fields in the sample display.

Field name

Descriptions

LB INDEX

Reference number that is used to identify the load balancer. The same reference number is used for a load balancer as long as the load balancer remains connected.

UUID A hexadecimal value of the universally unique identifier that is assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

GROUP NAME

The name of a group of related target server applications that are registered on the external load balancer.

IPADDR..PORT

The IP address and port number of a member defined in the registered group. The IP address can be an IPv4 or an IPv6 address.

MATCHES

The number of available server applications that match the IP address and port number. If the value of matches is 0, the Agent found the Linux on

System z guest that hosts this member logged on in this z/VM system, but did not find a server application that listens on that IP address and port number. This number will never be greater than 1.

PROTOCOL

The protocol that the member is using. The protocol value will always be TCP.

The following example shows a sample display when you issue `aqsllcmd display,members,detail`.

```
aqsllcmd display,members,detail
AQS0115I MEMBER DETAILS
LB INDEX      : 00          UUID       : 8F659238
GROUP NAME    : WKLD3_GROUP1
IPADDR..PORT : 10.1.1.1..9001
MATCHES      : 001        PROTOCOL    : TCP
TCPNAME      : GUEST1A
1 of 1 RECORDS DISPLAYED
Command complete
```

The following information describes the fields in the sample display.

Field name

Descriptions

LB INDEX

Reference number that is used to identify the load balancer. The same reference number is used for a load balancer as long as the load balancer remains connected.

UUID A hexadecimal value of the universally unique identifier that is assigned by the load balancer. This byte array can be up to 64 bytes in length. Only one load balancer with the same UUID can be connected to the primary Advisor at the same time.

GROUP NAME

The name of a group of related target server applications that are registered on the external load balancer.

IPADDR..PORT

The IP address and port number of a member defined in the registered group. The IP address can be an IPv4 or an IPv6 address.

MATCHES

The number of available server applications that match the IP address and port number. If the value of matches is 0, the Agent found the Linux on System z guest that hosts this member logged on in this z/VM system, but did not find a server application that listens on that IP address and port number. This number will never be greater than 1.

PROTOCOL

The protocol that the member is using. The protocol value will always be TCP.

TCPNAME

The guest name of the Linux on System z guest that hosts this member.

The following example shows a sample display when you issue `aqsllcmd display,members,detail,port=9001`.

```
aqsllcmd display,members,detail,port=9001
AQS0115I MEMBER DETAILS
LB INDEX      : 00          UUID       : 8F659238
```



```
GROUP NAME   : WKLD3_GROUP1
IPADDR..PORT : 10.1.1.1..9001
MATCHES      : 001          PROTOCOL   : TCP
TCPNAME      : GUEST1A
1 of 1 RECORDS DISPLAYED
Command complete
```

The following example shows a sample display when you issue `aqsllcmd display,members,detail,tcpname=guest1a`.

```
aqsllcmd display,members,detail,tcpname=guest1a
AQS0115I MEMBER DETAILS
LB INDEX     : 00          UUID       : 8F659238
GROUP NAME   : WKLD3_GROUP1
IPADDR..PORT : 10.1.1.1..9001
MATCHES      : 001          PROTOCOL   : TCP
TCPNAME      : GUEST1A
1 of 1 RECORDS DISPLAYED
Command complete
```

Change the debug level

Use the `aqsllcmd DEBug` command to change the debug level that is in effect for the Linux on System z Lifeline Agent.

Format

```
►► aqsllcmd DEBug, [ ,Level=debuglevel ] ◀◀
```

Parameters

DEBug,Level=debuglevel

Changes the debug level in effect for this Agent. The debug level that you can set is the sum of one or more levels of the debug messages to be logged. See “Debug levels” on page 67 for debug level values.

Example

The following example shows a sample display when you issue `aqsllcmd debug,level=127`.

```
aqsllcmd debug,level=127
AQS0138I LLAGENT DEBUG LEVEL UPDATED
Command complete
```

Dump internal trace records

Use the `aqsllcmd DUMPLog` command to dump internal trace records to the `/opt/ibm/lifeline/log/agent.trace` file.

Format

```
►► aqsllcmd DUMPLog ◀◀
```

Parameters

DUMPLog

Signals the Agent that internal trace records need to be dumped. Any existing data in the `/opt/ibm/lifeline/log/agent.trace` file will be overwritten.

Example

The following example shows a sample display when you issue `aqsl1cmd dumplog`.

```
aqsl1cmd dumplog
AQS0204I LLAGENT DUMP INTERNAL LOG COMMAND SUCCESSFUL
Command complete
```

Terminate the Linux on System z Management Agent

Use the `aqsl1cmd STOP` command to terminate the Linux on System z Management Agent.

Format

▶▶—`aqsl1cmd STOP`—▶▶

Parameters

STOP

Terminates the Linux on System z Management Agent.

Example

The following example shows a sample display when you issue `aqsl1cmd stop`.

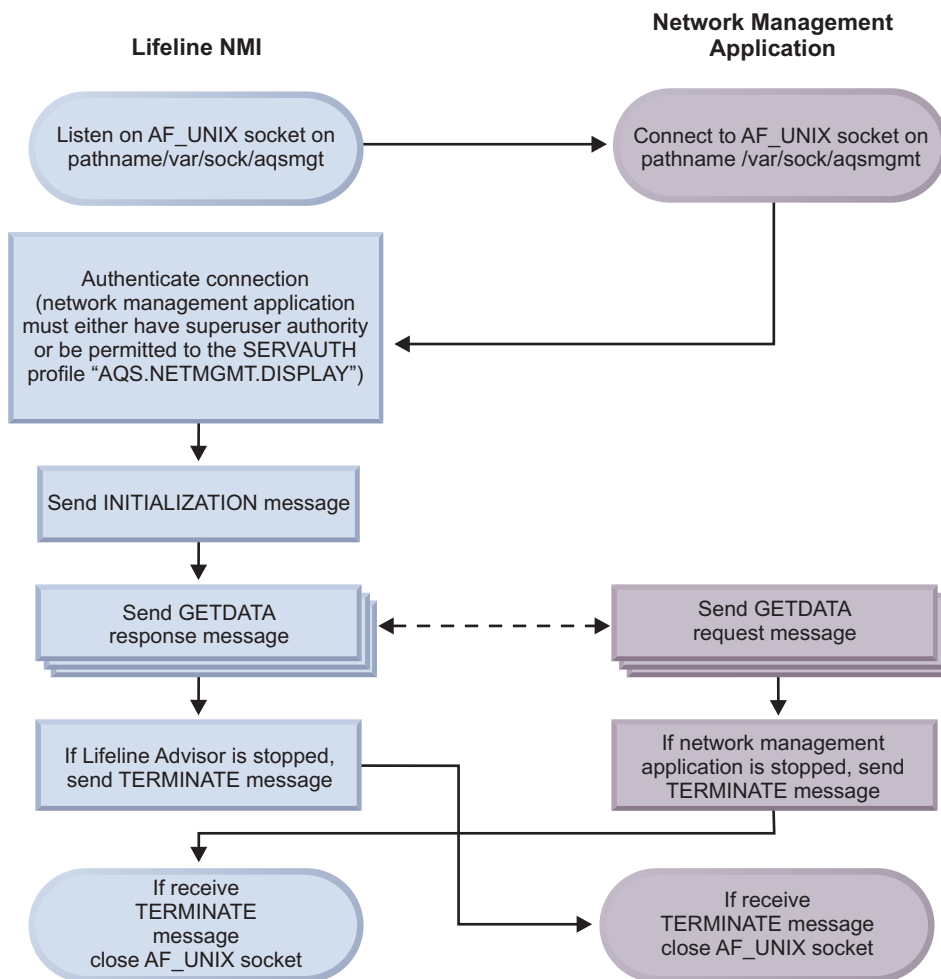
```
aqsl1cmd stop
Command complete
```

Chapter 10. Network Management Interface for Multi-site Workload Lifeline

Multi-site Workload Lifeline provides a Network Management Interface (NMI) that you can use to monitor Lifeline metrics and statistics. Network management applications can communicate with the Lifeline NMI through an AF_UNIX socket by using a series of messages.

The flow between the Lifeline NMI and the network management application is described in Figure 7.

Figure 7. Flow between the Lifeline NMI and the network management application



There are three possible message types: initialization message, data message, and termination message. The initialization and termination message contain only a message header. The data message is either a request or a response. A data request message contains only a message header.

The data response message contains the following headers and section descriptors:

1. A message header.

2. A record header for the overall response.
3. A section descriptor for Lifeline Advisors, which maps one section for the primary Advisor, and optionally, one section for the secondary Advisor.
4. A section descriptor for workloads, which maps sections for each of the defined workloads.
5. A section descriptor for sites, which maps sections for each of the defined sites.
6. A section descriptor for Lifeline Agents, which maps each of the Agents that connect to the primary Advisor.
7. A section descriptor for load balancers, which maps each of the load balancers (either external or z/OS sysplex distributor) that connect to the primary Advisor.
8. A section descriptor for load balancer application groups, which maps the groups that are defined to external load balancers.
9. A section descriptor for load balancer application group IP addresses, which maps the IP addresses that are configured for a specific load balancer application group.
10. A section descriptor for z/OS TCP/IP server applications, which maps all the z/OS TCP/IP server applications that make up each of the defined workloads.
11. A section descriptor for servers that are monitored by the Linux on System z Management Agent, which maps all the Linux on System z server applications that make up each of the defined workloads.
12. A section descriptor for SNA servers, which maps all the SNA applications that make up each of the defined workloads.
13. A section descriptor for MQ queue managers, which maps all of the MQ queue managers that host the IBM MQ clusters and MQ cluster queues for each defined workload.

Figure 8 describes the layout of the data response message.

Figure 8. Layout of the data response message

Message Header
Record Header Response
Section Descriptor for Lifeline Advisors
Section Descriptor for Workloads
Section Descriptor for Sites
Section Descriptor for Agents
Section Descriptor for Load Balancers
Section Descriptor for Load Balancers Groups
Section Descriptor for Load Balancers Group IP Addresses
Section Descriptor for z/OS TCP/IP server applications
Section Descriptor for servers monitored by Linux on System z Management Agents
Section Descriptor for SNA servers
Section Descriptor for MQ queue managers
Lifeline Advisor Sections (one for primary Advisor, and one for secondary Advisor – if active)
Zero or more Workload Sections
Zero or more Site Sections
Zero or more Lifeline Agent Sections
Zero or more Load Balancer Sections (both external and Sysplex Distributor load balancers)
Zero or more Load Balancer Group Sections
Zero or more Load Balancer Group IP Address Sections
Zero or more z/OS TCP/IP application Sections
Zero or more servers monitored by Linux on System z Management Agent Sections
Zero or more SNA server Sections
Zero or more MQ queue managers Sections

Use the record header descriptor in the message header to locate the response record header. You can locate the offset of the first section descriptor through the header. All eight section descriptors are contiguous. You can locate the data for each of the sections through the respective section descriptor. See “NMI message format” for more details about the layout of the NMI messages.

NMI message format

The following information shows the format of an NMI message.

```

/*
** Message Header
*/
typedef struct _AQSMMessageHdr {
    char          MsgHdrID[4];          /* Identifier for Message Header      */
#define AQSMMessageHdrID "AQSM"
    unsigned int  MsgHdrLength;        /* Length of message header          */
    unsigned short MsgHdrVersion;      /* Interface version                  */

#define AQSVersion1 1 /* AQS NMI version 1 */
#define AQSVersion2 2 /* AQS NMI version 2 */
#define AQSVersion3 3 /* AQS NMI version 3 */
#define AQSVersion4 4 /* AQS NMI version 4 */
#define AQSVersion 4 /* AQS NMI version for V2R5 */

    unsigned short MsgHdrType;         /* Message type                      */
#define AQSSInitType 1 /* Initialization message */
#define AQSTermType 2 /* Termination message */
#define AQSGetDataType 3 /* Get Advisor Data message */
    char          MsgHdrRsvd1[4];      /* Reserved - set to 0 */
    char          MsgHdrCorrelator[16]; /* User-defined correlator */
}

```

```

    unsigned int    MsgHdrRc;                /* Return code */
#define AQRcOK      0                       /* Successful NMI call */
#define AQRcError  8                       /* Unsuccessful NMI call */
    unsigned int    MsgHdrRsn;              /* Reason code */
#define AQSBadID   1                       /* Incorrect message identifier */
#define AQSBadVersion 2                   /* Incorrect message version */
#define AQSBadType 3                       /* Incorrect message type */
#define AQSBadSize 4                       /* Incorrect message size */
#define AQSBadFields 5                   /* Reserved fields are not zero */
#define AQSConnsExceeded 6               /* Too many active connections */
#define AQSInternalError 7               /* Internal error */
#define AQSNotAuth 8                      /* Authorization failure */
    unsigned int    MsgHdrMsgLength;        /* Message length (including header) */
    unsigned int    MsgHdrTime;             /* Timestamp when message was created */
    char            MsgHdrRsvd2[16];        /* Reserved - set to 0 */
    AQRRecordDesc   MsgHdrRecDesc;         /* Record Descriptor - only
                                          meaningful for GETDATA responses */

} AQSMMessageHdr;
/*
** Section Correlator
*/
typedef struct _AQSCorrelator {
    char            SecID[4];               /* Identifier for the section */
    unsigned int    SecInstance;           /* Instance count of this section */
} AQSCorrelator;

#define AQSAdvisorID      "AQSA"          /* Identifier for Advisor section */
#define AQSWorkloadID     "AQSW"          /* Identifier for Workload section */
#define AQSsiteID         "AQSI"          /* Identifier for Site section */
#define AQSAGentID        "AQSC"          /* Identifier for Agent section */
#define AQSLoadBalancerID "AQSL"          /* Identifier for Load Balancer section */
#define AQSGroupID        "AQSG"          /* Identifier for Group section */
#define AQSGrpIPID        "AQSP"          /* Identifier for Group IP Address section */
#define AQSServerID       "AQSS"          /* Identifier for Servers section */
#define AQSMgmtGuestServerID "AQSZ"      /* Identifier for Linux on System z
                                          Management Guest Servers section */
#define AQSNAServerID     "AQSN"          /* Identifier for SNA Servers section */
#define AQSMMQmgrID       "AQSQ"          /* Identifier for MQ queue managers section*/
/*
** Message Section for Lifeline Advisors (1st section)
*/
typedef struct _AQSAAdvisorSec {
    AQSCorrelator   AdvCorr;               /* Advisor correlator */
    struct in_addr   AdvIpAddr4;           /* Advisor's IPv4 Address */
    struct in6_addr  AdvIpAddr6;           /* Advisor's IPv6 Address */
    char            AdvSysName[8];         /* System Name where Advisor resides */
    unsigned char    AdvRole;              /* Advisor role */
#define AQSAdvisorPrimary 1               /* Advisor in primary role */
#define AQSAdvisorSecondary 2             /* Advisor in secondary role */
    char            AdvRsvd[3];            /* Reserved - set to 0 */
} AQSAAdvisorSec;
/*
** Message Section for Workloads (2nd section)
*/
typedef struct _AQSWWorkloadSec {
    AQSCorrelator   WkldCorr;              /* Workload correlator */
    char            WkldWorkloadName[64]; /* Name of workload */
    unsigned char    WkldStatus;           /* Routing status of workload */
#define AQSWorkloadOk     1               /* Site and servers for workload
                                          are available */
#define AQSWorkloadNotOk 2               /* Site or servers for workload
                                          are not available */
    char            WkldRsvd[3];           /* Reserved - set to 0 */
} AQSWWorkloadSec;
/*
typedef struct _AQSWWorkloadSecAQ {
    unsigned char    WkldType;              /* Workload type */
#define AQSWorkloadTypeAS 1               /* Active/Standby workload */
#define AQSWorkloadTypeAQ 2               /* Active/Query workload */
    unsigned char    WkldDistType;         /* Workload distribution type for Ac
                                          tive/Query workloads */
#define AQSWorkloadTypeNotUsed 0          /* Active/Standby workload */
#define AQSWorkloadTypeStatic 1           /* Static distribution */
#define AQSWorkloadTypeDynamic 2         /* Dynamic distribution */
    char            WkldRsvd2[2];          /* Reserved - set to 0 */
    AQSCorr_t       WkldAssocWorkload;     /* For Active/Query workload,
                                          correlator of the associated
                                          Active/Standby workload */
    unsigned short   WkldAssocNum;         /* For Active/Standby workload, the

```

```

        number of Active/Query workloads
        associated with this workload
        char WkldRsvd3[2];
        unsigned int WkldLatencyFlags;
        #define AQSWkldLatencyAcute 0x00000001
        #define AQSWkldLatencyNormal 0x00000002
        #define AQSWkldLatencyConstrained 0x00000004
        #define AQSWkldLatencyCritical 0x00000008
        #define AQSWkldLatencyUnknown 0x00000010
        unsigned int WkldStatusRsn;
        #define AQSWkldStatusOk 1
        #define AQSWkldStatusQuiesced 2
        #define AQSWkldStatusWorkFailure 3
        #define AQSWkldStatusSiteFailure 4
        #define AQSWkldStatusRouting 5
    } AQSWorkloadSecAQ;

typedef struct _AQSWorkloadSecV2 {
    AQSWorkloadSec AQSWorkloadSection1;
    AQSWorkloadSecAQ AQSWorkloadSection2;
} AQSWorkloadSecV2;

/*
** Message Section for Sites (3rd section)
*/
typedef struct _AQSSiteSec {
    AQSCorrelator SiteIdCorr;
    AQSCorrelator SiteWkldCorr;

    char SiteRsvd1[16];

    char SiteName[8];
    unsigned char SiteWorkloadState;
    #define AQSSiteWorkloadActive 1
    #define AQSSiteWorkloadQuiesce 2
    #define AQSSiteWorkloadDeactivated 3
    char SiteRsvd2[3];
    unsigned short SiteWeight;

    /* If AQS NMI version is less than version 4, SiteFlags is reserved */
    unsigned short SiteFlags;
    #define AQSSiteClusterNmValid 2

    /* SiteClusterName field is included when this is AQS NMI version 4 */
    char SiteClusterName[48];

} AQSSiteSec;

typedef struct _AQSSiteSecAQ {
    unsigned int SitePercentWeight;

    unsigned int SitePreferredFlag;

    #define AQSSiteNotPreferred 1
    #define AQSSitePreferred 2
    char SiteReserved2[3];
} AQSSiteSectionAQ;

typedef struct _AQSSiteSectionV2 {
    AQSSiteSection AQSSiteSection1;
    AQSSiteSectionAQ AQSSiteSection2;
} AQSSiteSectionV2;

/*
** Message Section for Lifeline Agents (4th section)
*/
typedef struct _AQSAgentSec {
    AQSCorrelator AgntCorr;

    union {
        struct in_addr in;
        struct in6_addr in6;
    } AgntIpAddr;
    char AgntSysName[8];
    char AgntSiteName[8];

```

```

        unsigned char    AgntState;           where Agent resides
#define AQSAgentActive  1                    /* State of Agent
#define AQSAgentInactive 2                  /* Agent is active
        char            AgntType[1];        /* Agent is inactive
#define AQSAgentZOS     1                    /* Agent type
#define AQSAgentZLinuxMgmt 2                /* z/OS Agent
        unsigned short  AgntFlags;         /* Linux on System z Management Agent
#define AQSAgentIPv4   1                    /* Agent flags
        } AQSAgentSec;                     /* AgntIpAddr is IPv4 address
        /*
** Message Section for Load Balancers (5th section)
*/
typedef struct _AQSLBSec {
    AQSCorrelator    LBCorr;                /* Correlator for Load Balancer
    union {
        struct in_addr  in;                 /* Load Balancer's IPv4 address
        struct in6_addr in6;                /* Load Balancer's IPv6 address
    } LBIPAddr;
    unsigned char    LBTier;                /* Load Balancer Tier
#define AQSLoadBalancerTier1  1            /* First-tier Load Balancer
#define AQSLoadBalancerTier2  2            /* Second-tier Load Balancer
#define AQSLoadBalancerTierUnk 3           /* Unknown tier
#define AQSLoadBalancerIntermed 4          /* Intermediary tier Load Balancer
    unsigned char    LBState;               /* State of load balancer
#define AQSLoadBalancerActive  1           /* Load balancer is active
#define AQSLoadBalancerInactive 2         /* Load balancer is not active
    unsigned short   LBFlags;               /* Load balancer flags
#define AQSLoadBalancerInt     1           /* Internal load balancer (i.e.
                                           Sysplex Distributor)
#define AQSLoadBalancerIPv4    2           /* LBIPAddr is IPv4 address
    char             LBRsvd[8];            /* Reserved - set to zero
} AQSLBSec;
/*
** Message Section for Groups (6th section)
*/
typedef struct _AQSGroupSec {
    AQSCorrelator    GrpCorr;                /* Correlator for Group
    AQSCorrelator    GrpLBCorr;              /* Correlator for Load Balancer
                                           this group is defined to
    AQSCorrelator    GrpWkldCorr;           /* Correlator for Workload this
                                           group belongs to
    unsigned short   GrpPortNum;           /* Port number for the servers
                                           configured for this group
    char             GrpName[255];         /* Group name
    unsigned short   GrpFlags;             /* Group flags
#define AQSGroupInt 1                    /* Simulated group for internal load
                                           balancers (i.e. Sysplex Distributor) -
                                           there is not a corresponding GroupIP
                                           section for this group
    char GrpRsvd[1];                       /* Reserved - set to zero
} AQSGroupSec;
/*
** Message Section for Group IP Addresses (7th section)
*/
typedef struct _AQSGrpIpSec {
    AQSCorrelator    GrpIpCorr;             /* Correlator for GrpIp
    AQSCorrelator    GrpIpGrpCorr;         /* Correlator for Load Balancer
                                           Group this IP Address is
                                           defined to
    union {
        struct in_addr  in;                 /* IPv4 address for a server
                                           configured for this group
        struct in6_addr in6;                /* IPv6 address for a server
                                           configured for this group
    } GrpIpAddr;
    unsigned short   GrpIpFlags;           /* Group IP flags
#define AQSGrpIpIPv4  1                    /* GrpIpAddr is IPv4 address
    char             GrpIpRsvd[2];        /* Reserved - set to zero
} AQSGrpIpSec;
/*
** Message Section for Servers (8th section)
*/
typedef struct _AQSServerSec {
    AQSCorrelator    SrvrCorr;              /* Correlator for Server
    AQSCorrelator    SrvrLBCorr;           /* Correlator for Load Balancer this
                                           server is defined to
    AQSCorrelator    SrvrWkldCorr;        /* Correlator for Workload this
                                           server belongs to
    union {

```



```

        struct in_addr  in;          /* Server's IPv4 address      */
        struct in6_addr in6;       /* Server's IPv6 address      */
    } SrvrIpAddr;
    char      SrvrJobName[8];      /* The MVS jobname of server  */
    char      SrvrSysName[8];      /* System name where server resides */
    char      SrvrSiteName[8];    /* Name of site (i.e. sysplex)
                                   where server resides        */
    unsigned int  SrvrNormWeight; /* Net weight for this server
                                   relative to other servers on this
                                   site - calculated by applying CS
                                   health as percentage of WLM
                                   weight for this server          */
    unsigned int  SrvrWlmWeight; /* ServerWLM weight for this
                                   server - measure of how well
                                   server is meeting its WLM
                                   policies and represents the
                                   displaceable processor capacity
                                   based on server's importance
                                   relative to other servers on this
                                   site                            */
    unsigned int  SrvrAbTerms;    /* Server abnormal terminations -
                                   number of abnormal transaction
                                   completions per 1000
                                   transactions                    */
    unsigned int  SrvrHealth;     /* Health indicator for server -
                                   used to adjust the reported
                                   ServerWLM weight              */
    unsigned int  SrvrCPUWeight; /* CPU factor in WLM weight -
                                   amount of displaceable general
                                   CPU capacity based on server's
                                   importance relative to other
                                   servers on this site and
                                   modified by proportion of
                                   general CPU capacity being
                                   consumed compared to zAAP
                                   and zIIP processors            */
    unsigned int  SrvrZAAPWeight; /* zAAP factor in WLM weight -
                                   amount of displaceable zAAP
                                   capacity based on server's
                                   importance relative to other
                                   servers on this site and modified
                                   by proportion of zAAP capacity
                                   being consumed compared to
                                   general CPU and zIIP
                                   processors                    */
    unsigned int  SrvrZIIPWeight; /* zIIP factor in WLM weight -
                                   amount of displaceable zIIP
                                   capacity based on server's
                                   importance relative to other
                                   servers on this site and modified
                                   by proportion of zIIP capacity
                                   being consumed compared to
                                   general CPU and zAAP
                                   processors                    */
    unsigned int  SrvrCSHealth; /* Health of server from a
                                   Communications Server
                                   perspective - calculated based
                                   on how well server is keeping up
                                   with connections in backlog, how
                                   well it is establishing new
                                   connections, and whether the
                                   server is dropping connections */
    unsigned short SrvrPortNum; /* Port number where server is
                                   listening                    */
    unsigned short SrvrFlags; /* Server flags          */
#define AQSServerIPv4 1 /* SrvrIPAddr is Ipv4 address */
    unsigned char SrvrState; /* State of server      */
#define AQSServerUp 1 /* Server is up        */
#define AQSServerDown 2 /* Server is down      */
    char SrvrRsvd[3]; /* Reserved - set to zero */
} AQSServerSec;
/*
** Message Section for Linux on System z Guest Servers (9th section, optional)
**
typedef struct _AQSMgmtGuestServerSection
{
    AQSCorrelator MgmtGuestServer; /* Correlator for Management Guest Server */
    AQSCorrelator MgmtGuestServerLB; /* Correlator for Load Balancer this
                                       server is defined to */
    AQSCorrelator MgmtGuestServerWork; /* Correlator for Workload this
                                       server is defined to */
}

```

```

server belongs to */
union {
    struct in_addr in; /* Server's IPv4 address */
    struct in6_addr in6; /* Server's IPv6 address */
} MgmtGuestServerIPAddr;
char MgmtGuestSystemID[8]; /* z/VM System ID where the server resides */
char MgmtGuestSiteName[8]; /* Name of site where server resides */
char MgmtGuestName[8]; /* Linux on System z Guest name where the
server resides */
unsigned int MgmtGuestServerWt; /* Net weight for this server relative to
other Linux on System z servers on this
site */
unsigned short MgmtGuestServerPort; /* Port number where server is listening */
unsigned short MgmtGuestServerFlags; /* Server flags */
/* #define AQSServerIPv4 1 MgmtGuestServerIPAddr is IPv4 address */
unsigned char MgmtGuestServerState; /* State of server */
/* #define AQSServerUp 1 Server is up */
/* #define AQSServerDown 2 Server is down */
char MgmtGuestServerRsvd[3]; /* Reserved - set to zero */
} AQSMgmtGuestServerSection;
/*
** Message Section for SNA Application Servers (10th section, optional)
*/
typedef struct _AQSSNAServerSection
{
    AQSCorrelator SNAServer; /* Correlator for SNA Server */
    AQSCorrelator SNAServerLB; /* Correlator for Load Balancer this
server is defined to */
    AQSCorrelator SNAServerWork; /* Correlator for Workload this
server belongs to */
    char SNAJobName[8]; /* The MVS jobname of server */
    char SNASystemName[8]; /* System name where server resides */
    char SNASiteName[8]; /* Name of site (i.e sysplex) where server
resides */
    char SNAServerAppName[8]; /* The VTAM name of the SNA application */
    char SNAServerNetid[8]; /* The VTAM network ID of the SNA
application */
    unsigned int SNAServerNormWt; /* Net weight for this server relative to
other SNA servers on this site */
    unsigned char SNAServerState; /* State of server */
#define AQSSNAServerUp 1 /* Server is up */
#define AQSSNAServerDown 2 /* Server is down */
    unsigned char SNAFlags; /* Server flags */
#define AQSSNADuplicateSvr 1 /* Server SNA application already specified
under a different Load Balancer
correlator. */
    char SNAServerRsvd[2]; /* Reserved - set to zero */
} AQSSNAServerSection;
/*
** Message Section for MQ queue managers (11th section, optional)
*/
typedef struct _AQSMQMGrSection
{
    AQSCorrelator AQSMQMGr; /* Correlator for MQ queue manger */
    AQSCorrelator AQSMQMGrWork; /* Correlator for a workload this MQ queue
manager processes messages for */
    char AQSMQMGrName[4]; /* MQ queue manager name */
    char AQSMQMGrIPAddr[16]; /* MQ queue manager IP Addr */
    char AQSMQMGrJobName[8]; /* MQ queue manager job name */
    char AQSMQMGrSystemName[8]; /* MQ queue manager System name */
    char AQSMQMGrSysplexName[8]; /* MQ queue manager Sysplex name */
    char AQSMQMGrClusterName[48]; /* MQ cluster name */
    unsigned int AQSMQMGrNormWt; /* Normalized MQ queue manager weight */
    unsigned int AQSMQMGrWLMWt; /* WLM MQ queue manager server specific
weight */
    unsigned short AQSMQMGrPort; /* MQ queue manager port */
    unsigned short AQSMQMGrFlags; /* MQ queue manager flags */
#define AQSMQMGrIPv4 1 /* AQSMQMGrIPAddr is IPv4 */
    unsigned char AQSMQMGrState; /* MQ queue manager state */
#define AQSMQMGrUp 1 /* MQ queue manager is up */
#define AQSMQMGrDown 2 /* MQ queue manager is down */
    char AQSMQMGrRsvd[3]; /* Reserved - set to zero */
} AQSMQMGrSection;

```

Chapter 11. Multi-site Workload Lifeline messages

AQS0001I MULTI-SITE WORKLOAD LIFELINE REQUIRES THE TCP/IP BASE FEATURE OF z/OS

Explanation: The Multi-site Workload Lifeline function is disabled because the TCP/IP base feature of z/OS is not enabled.

System action: Processing ends.

Operator response: Contact the system programmer.

System programmer response: The IFAPRDxx parmlib member specifies whether or not the TCP/IP base feature is enabled. Correct the IFAPRDxx parmlib member if it is in error. IBM has defined the appropriate product enablement settings in the IFAPRD00 member of SYS1.IBM.PARMLIB.

If TCP/IP_BASE is not part of your system, no action is required; the Multi-site Workload Lifeline function cannot be used.

Module: AQSUTUTI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0004I *applname* IS NOT SUPPORTED IN THIS ENVIRONMENT

Explanation: The Multi-site Workload Lifeline application was started in an unsupported environment. It is supported only in the z/OS UNIX shell or batch environments.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline z/OS Agent

System action: Processing ends.

Operator response: Start the Multi-site Workload Lifeline Advisor or z/OS Agent in a supported environment.

System programmer response: None.

Module: AQSUTUTI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0004I LLADVSR IS NOT SUPPORTED IN THIS ENVIRONMENT

AQS0005I *applname* MUST BE APF-AUTHORIZED

Explanation: The Multi-site Workload Lifeline application requires APF authorization to function correctly, but the application was not running as APF authorized.

In the message text:

AQS0006I • AQS0007I

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline z/OS Agent

System action: Processing ends.

Operator response: Contact the system programmer to ensure that the Advisor or Agent was installed correctly.

System programmer response: Ensure that the Advisor or z/OS Agent was installed correctly. If the Advisor or z/OS Agent executable application resides in the z/OS UNIX file system, ensure that the APF-authorized attribute is on. If the Advisor or z/OS Agent resides in a load module, ensure that the library that contains it is APF authorized.

Module: AQSUTUTI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0005I LLADVSR MUST BE APF-AUTHORIZED
```

AQS0006I MESSAGE CATALOG *filename* COULD NOT BE OPENED - DEFAULT MESSAGES WILL BE USED

Explanation: The Multi-site Workload Lifeline message catalog file could not be opened. The command will use the default message strings instead of the message strings in the message catalog file.

In the message text:

filename

The name of the message catalog that could not be opened.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Ensure that the message catalog was installed properly. If the message catalog is installed in a location other than the default, verify that the NLSPATH and LANG environment variables are correctly defined in your Lifeline Advisor and Lifeline Agent started procedures. For Linux on System z Management Agents, verify that the catalog is located in the /opt/ibm/lifeline/lib/msg/C directory.

Module: AQSUTUTI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0006I MESSAGE CATALOG aqs11msg.cat COULD NOT BE OPENED - DEFAULT MESSAGES WILL BE USED
```

AQS0007I MESSAGE CATALOG *filename* AT INCORRECT SERVICE LEVEL - DEFAULT MESSAGES WILL BE USED

Explanation: The Multi-site Workload Lifeline message catalog is not at the same service level as the Lifeline Advisor or Lifeline Agent. The Lifeline application will use the default message strings instead of the message strings in the message catalog.

In the message text:

filename

The name of the message catalog that could not be used.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Ensure that the correct service level of the message catalog was installed. If the updated message catalog was installed in a location other than the default location, verify that the NLSPATH and LANG environment variables are correctly defined in your Lifeline Advisor and Lifeline Agent started procedures. For Linux on System z Management Agents, verify that the updated catalog is installed in the /opt/ibm/lifeline/lib/msg/C directory.

Module: AQSUTUTI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0007I MESSAGE CATALOG aqs11msg.cat AT INCORRECT SERVICE LEVEL - DEFAULT MESSAGES WILL BE USED
```

AQS0008I AGENT UNABLE TO PROCESS DEACTIVATE COMMANDS

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message when it does not have authority to drop connections as a result of a DEACTIVATE command issued from the Lifeline Advisor.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Using the SAF product on your system, ensure the following:

- The OPERCMDS class is active.
- The MVS.VARY.TCPIP.DROP profile is defined.
- The Lifeline z/OS Agent has CONTROL access to this profile.

When the errors are corrected, stop and restart the Lifeline z/OS Agent. See “(z/OS systems) Authorizing the Lifeline Agents to reset active connections” on page 38 for more information about configuring the correct authorization.

Module: LAMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0009I SNA APPLICATIONS CANNOT BE MONITORED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message when it does not have authority to open an extended MCS console. The z/OS Agent uses this console to manage SNA applications.

System action: Processing continues; however, SNA applications on this z/OS Agent's system will not be monitored. Multi-site Workload Lifeline Advisor displays will show these SNA applications as not active.

Operator response: Contact the system programmer.

System programmer response: Using the SAF product on your system, ensure the following conditions are true:

- The OPERCMDS class is active.
- The MVS.MCSOPER.LLAGNT* profile is defined.
- The Lifeline z/OS Agent has READ access to this profile.

When the errors are corrected, stop and restart the Lifeline z/OS Agent. See “(z/OS systems) Authorizing the Lifeline Agents to monitor SNA applications” on page 39 for more information about configuring the correct authorization.

Module: LASNA

Routing code: 8, 10

Descriptor code: 12

AQS0010I • AQS0101I

Automation: Not applicable for automation.

| AQS0010I MQ QUEUE MANAGERS CANNOT BE MONITORED

| **Explanation:** The Multi-site Workload Lifeline z/OS Agent issues this message when one of the following conditions is true:

- | • The MQ APIs used to connect to a local MQ queue manager cannot be loaded.
- | • The MQ queue manager is not at a supported level.

| The z/OS Agent uses this connection to monitor MQ queue managers.

| **System action:** Processing continues; however, MQ queue managers on this z/OS Agent's system will not be monitored. Multi-site Workload Lifeline Advisor displays will show these MQ queue managers as unavailable.

| **Operator response:** Contact the system programmer.

| **System programmer response:** To determine why the MQ queue manager cannot be monitored:

- | • If internal trace is being used, dump and examine the Agent's internal trace messages.
- | • If internal trace is not being used, examine the Agent syslogd file. The syslogd identifier is llagent for the Agent.

| If the MQ APIs could not be loaded, ensure one of the following conditions is true:

- | • The IBM MQ SCSQLOAD dataset name is in the LNKSTxx SYS1.PARMLIB member.
- | • The IBM MQ SCSQLOAD dataset name is added as a STEPLIB DD statement in the z/OS Agent's started procedure.

| If an unsupported version of IBM MQ is being used, ensure that WebSphere MQ version 7.0.1 or higher is installed.

| When the errors are corrected, stop and restart the Lifeline z/OS Agent. See "(z/OS systems) Enabling the Lifeline Agents to communicate with MQ queue managers" on page 40 for more information about adding the IBM MQ SCSQLOAD dataset name to the z/OS Agent's linklist concatenation.

| **Module:** LAMQ

| **Routing code:** 8, 10

| **Descriptor code:** 12

| **Automation:** Not applicable for automation.

AQS0101I *applname* STARTING

Explanation: The Multi-site Workload Lifeline application is starting.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LAMAIN, LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0101I LLADVSR STARTING

AQS0102I *applname* **INITIALIZATION COMPLETE****Explanation:** The Multi-site Workload Lifeline application completed initialization.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing continues.**Operator response:** None.**System programmer response:** None.**Module:** LAMAIN, LMMAIN**Routing code:** 8, 10**Descriptor code:** 12**Automation:** Not applicable for automation.**Example:**

AQS0102I LLADVSR INITIALIZATION COMPLETE

AQS0103I *applname* **SHUTDOWN IN PROGRESS****Explanation:** The Multi-site Workload Lifeline application is shutting down in response to a STOP command.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing continues.**Operator response:** None.**System programmer response:** None.**Module:** LAMAIN, LMMAIN**Routing code:** 8, 10**Descriptor code:** 12**Automation:** Not applicable for automation.**Example:**

AQS0103I LLADVSR SHUTDOWN IN PROGRESS

AQS0104I • AQS0105I

AQS0104I *applname* SHUTDOWN COMPLETE

Explanation: The Multi-site Workload Lifeline application ended in response to a STOP command.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing ends.

Operator response: None.

System programmer response: None.

Module: LAMAIN, LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0104I LLADVSR SHUTDOWN COMPLETE

AQS0105I *applname* CONFIGURATION ERRORS DETECTED

Explanation: The configuration file for the Multi-site Workload Lifeline application cannot be opened, or contains one or more errors.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing ends.

Operator response: Contact the system programmer.

System programmer response: If syslogd was started, examine the syslogd file for configuration error messages. The syslogd identifier is lladvsr for the Advisor and llagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages for configuration error messages. Correct the configuration file errors and restart the application. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

Module: LACONFIG, LMCONFIG

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0105I LLADVSR CONFIGURATION ERRORS DETECTED

AQS0106I ADVISOR REFRESH ERRORS DETECTED

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*, REFRESH command when the Advisor configuration file contains changes to statements that cannot be refreshed.

System action: The MODIFY command is ignored.

Operator response: Contact the system programmer.

System programmer response: If syslogd was started, examine the syslogd file for configuration error messages. The syslogd identifier is lladvr for the Advisor. Otherwise, if internal trace is being used, dump and examine the internal trace messages for configuration error messages. Correct the configuration file errors and reissue the MODIFY *procname*, REFRESH command. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0107I *applname* MODIFY COMMAND SYNTAX ERROR AT *location*

Explanation: The Multi-site Workload Lifeline application detected a syntax error on a z/OS MODIFY command or the Linux on System z aqslcmd command.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

location

The portion of the command that is not syntactically valid.

System action: The command is rejected.

Operator response: If syslogd was started, examine the syslogd file to determine why the MODIFY command failed. The syslogd identifier is lladvr for the Advisor and lagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages to determine why the command failed. Re-enter the command using the valid syntax. See Chapter 7, "Lifeline Advisor commands," on page 115, Chapter 8, "Lifeline z/OS Agent commands," on page 151, and Chapter 9, "Lifeline Linux on System z Management Agent commands," on page 165 for more information.

System programmer response: None.

Module: LACMD, LCCMD, LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0107I LLADVSR MODIFY COMMAND SYNTAX ERROR AT 'BAD'

AQS0108I *applname* **MODIFY COMMAND PARAMETER** *parm* **INCORRECT VALUE** *value*

Explanation:

The Multi-site Workload Lifeline application detected an incorrect value for a parameter on a z/OS MODIFY command or the Linux on System z aqslcmd command.

In the message text:

applname

The application name. Possible values:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

parm

The name of the command parameter that is in error.

value

The incorrect value entered on the command.

System action: The command is rejected.

Operator response: Re-enter the MODIFY command with a valid value for the indicated parameter. See Chapter 7, "Lifeline Advisor commands," on page 115, Chapter 8, "Lifeline z/OS Agent commands," on page 151, and Chapter 9, "Lifeline Linux on System z Management Agent commands," on page 165 for more information.

System programmer response: None.

Module: LCCMD, LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0108I LLADVSR MODIFY COMMAND PARAMETER LEVEL INCORRECT VALUE 2000

AQS0109I *syscall* **CALL FAILED** *errno/errnojr* **FOR PORT** *port*

Explanation: The Multi-site Workload Lifeline Advisor or Agent application was unable to initialize a listening or client socket.

In the message text:

syscall

The type of call that failed. For example, the system call might be a SOCKET, SETSOCKOPT, or BIND call.

errno

The UNIX System Services return code. See *return codes (errno)* in *z/OS UNIX System Services Messages and Codes* for more information about these return codes.

errnojr

The hexadecimal UNIX System Services reason code. See *return codes (errno)* in *z/OS UNIX System Services Messages and Codes* for more information about the format of the 4-byte reason code.

port

The port number on which the call failed.

System action: Processing ends.

Operator response: Determine whether TCP/IP is available and restart if necessary. Restart the application. If TCP/IP is available when the application issues this message, take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If the internal trace is not being used, save the syslogd file.

- Contact the system programmer.

System programmer response: If TCP/IP is available when the application issues this message and if syslogd was started, examine the syslogd file and correct the error. The syslogd identifier is lladvrs for the Advisor and lagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages, and correct the error.

Module: LMLISTEN, LASOCKET

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0109I BIND CALL FAILED 6F/744C7246 FOR PORT 3860

AQS0110I UNABLE TO ESTABLISH *endpoint* LISTENING SOCKET

Explanation: The Multi-site Workload Lifeline Advisor was unable to initialize because it could not create a listening socket.

In the message text:

endpoint

The listening socket that could not be created. Possible values are LOAD BALANCER, ADVISOR, or AGENT.

System action: Processing ends.

Operator response: Determine whether TCP/IP is available and restart if necessary. Restart the application. If TCP/IP is available when the application issues this message, take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, dump and save the syslogd file.
- Contact the system programmer.

System programmer response: If TCP/IP is available when the Advisor issues this message, see message AQS0109I that might precede this message, or if syslogd was started, examine the syslogd file and correct the error. The syslogd identifier is lladvrs for the Advisor. Otherwise, if internal trace is being used, dump and examine the internal trace messages, and correct the error.

Module: LMLISTEN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0110I UNABLE TO ESTABLISH LOAD BALANCER LISTENING SOCKET

AQS0111I *applname* DEBUG LEVEL *level*

Explanation: This message is issued in response to a z/OS MODIFY *procname*, DISPLAY, DEBUG command or a Linux on System z aqs11cmd display, debug command.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

AQS0112I

level

The current debug level in effect. Valid debug level values are:

- 1 Errors are logged.
- 2 Warnings are logged.
- 4 Significant events are logged.
- 8 Informational messages are logged.
- 16 Debug data messages are logged. This level is for IBM service use only. If the *applname* value is LLADVSR, messages related to TCP/IP messages sent between the Advisor and load balancers, between the primary and secondary Advisors, and between the Advisor and Agent are logged.
- 32 Debug data messages are logged. This level is for IBM service use only. If the *applname* value is LLAGENT, messages related to data collection and manipulation that support weight calculations are logged.
- 64 Internal debug data is logged. This level is for IBM service use only.
- 128 Function entry and exit tracing are logged. This level is for IBM service use only.

Individual values can be added together. For example, if the debug level displayed is 7, all ERROR, WARNING, and EVENT messages are logged. If the AQSTRACE DD card was specified in the Advisor or z/OS Agent started procedure, these messages are logged in an internal trace, and only messages at the ERROR, WARNING, and EVENT levels are logged to the syslogd file. For Linux on System z Agents, internal trace is being used by default. If internal trace is not being used, all messages are logged to the syslogd file.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD, LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0111I LLADVSR DEBUG LEVEL 7
```

AQS0112I LOAD BALANCER SUMMARY

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,LB command. This message is followed by summary information about connected load balancers. See "MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0113I LOAD BALANCER DETAILS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,DISPLAY,LB,DETAIL` command. This message is followed by detailed information about connected load balancers. See “*MODIFY advisor_proc,DISplay* command: Display Lifeline Advisor information” on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0114I MEMBER SUMMARY

Explanation: The Multi-site Workload Lifeline Agent issues this message in response to a z/OS `MODIFY procname,DISPLAY,MEMBERS` command or a Linux on System z `aqsl1cmd display,members` command. This message is followed by summary information about members (applications being load balanced) that one or more load balancers registered with the Multi-site Workload Lifeline Advisor. See “*MODIFY agent_proc,DISplay* command: Display Lifeline Agent information” on page 151 or “*Display Linux on System z Management Agent* information” on page 165 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0115I MEMBER DETAILS

Explanation: The Multi-site Workload Lifeline Agent issues this message in response to a `MODIFY procname,DISPLAY,MEMBERS,DETAIL` command or a Linux on System z `aqsl1cmd display,members,detail` command. This message is followed by detailed information about members (applications being load balanced) that one or more load balancers registered with the Multi-site Workload Lifeline Advisor. See “*MODIFY agent_proc,DISplay* command: Display Lifeline Agent information” on page 151 or “*Display Linux on System z Management Agent* information” on page 165 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0116I

AQS0116I *applname* **INITIALIZATION ERROR - REASON CODE** *reason*

Explanation: The Multi-site Workload Lifeline application could not initialize because it detected an error.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

reason

The code for the error encountered. Possible values for the reason code are:

- 1 Another copy of the application is already active. Only one copy of the Advisor can be active within the MVS sysplex. Only one copy of the z/OS Agent can be active on an MVS system. Only one copy of the Linux on System z Management Agent can be active on a z/VM system for a given site.
- 2 The Advisor or Agent application cannot open the configuration file, or found an error in the configuration file. Message AQS0105I might precede this message.
- 3 An internal error was encountered in the application.
- 4 The Advisor or Agent application found an error in a start option parameter in the started procedure.
- 5 The user ID associated with the Advisor or z/OS Agent started task is not authorized, or the Advisor or z/OS Agent was not started as a started procedure.
- 6 A required resource is not available. For example, the TCP/IP stack is not started, or a configured IP address is not defined or available on the TCP/IP stack.
- 7 One of the following conditions has occurred:
 - The Advisor or z/OS Agent application was started on an unsupported z/OS release.
 - The Linux on System z Management Agent was started and not all of the required components or packages were found.
 - The Linux on System z Management Agent was started and not all of the required components or packages are at the required level.

System action: Processing ends.

Operator response:

- If the *reason* value is 1, verify that the instance of the active application is the one that you want. If a previous instance of the application is stopped but has not ended, issue a CANCEL command to cancel the previous instance of the Advisor or z/OS Agent before you start the new instance, or kill the previous instance of the Linux on System z Management Agent.
- If the *reason* value is 6, start the TCP/IP stack if it is not already started.
- If the *reason* value is 7, verify that Advisor or z/OS Agent is being started on a supported z/OS release. Multi-site Workload Lifeline is supported on z/OS V1R13 and higher releases. For the Linux on System z Management Agent, take the following actions:
 - If internal trace is being used, dump and save the internal trace messages.
 - If internal trace is not being used, save the syslogd file.
 - Contact the system programmer.
- If the stack is already started, and for all other reason values, take the following actions:
 - If internal trace is being used, dump and save the internal trace messages.
 - If internal trace is not being used, dump and save the syslogd file.
 - Contact the system programmer.

System programmer response: Ensure that there is sufficient information to debug this problem. Configure the application to use a debug level of 127. See “Debug levels” on page 67 for information about how to configure the debug level for the specified application. Take the appropriate action, based on the *reason* value:

- If the *reason* value is 2, examine the application's dumped internal trace file or syslogd file for errors. Correct the configuration file as needed. See Chapter 3, "Configuring Multi-site Workload Lifeline," on page 43 for information about configuring the Advisor and Agent.
- If the *reason* value is 3, contact the IBM software support center. The application's dumped internal trace file or syslogd file is the minimum diagnostic data that must be provided. See "Diagnostic data" on page 68 for more information about collecting diagnostic data.
- If the *reason* value is 4, examine the application's dumped internal trace file or syslogd file for errors. Correct the started procedure as needed.
- If the *reason* value is 5, examine the security product profiles that have been established for the Advisor and z/OS Agent. Ensure that the Advisor or z/OS Agent is started from a started procedure and that the user ID in the started procedure is permitted to the appropriate profiles if they are defined. See "(z/OS systems) Authorizing the Lifeline Agents to use WLM services" on page 38 for more information about setting up the authorization profiles for the security product. Ensure that the Linux on System z Management Agent is started with a user ID with root authority.
- If the *reason* value is 6 and the TCP/IP stack is already started, and if syslogd was started, examine the syslogd file for errors. The syslogd identifier is lladvrs for the Advisor and llagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages for errors. Verify that the correct IP addresses are configured in the configuration file. Correct the configuration file as needed. See Chapter 3, "Configuring Multi-site Workload Lifeline," on page 43 for more information about configuring the Advisor and Agent.
- If the *reason* value is 7, for the Linux on System z Management Agent, see the installation documentation and verify that the following conditions are true:
 - All of the required packages and system software are installed and included in the PATH environment variable.
 - The required packages meet the minimum version and release specifications.
 - The Linux on System z Management Agent user ID has the required z/VM command privileges.

Module: LAMAIN, LMADV, LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0116I LLADVSR INITIALIZATION ERROR - REASON CODE 6

AQS0117I *applname* CONFIGURATION WARNINGS DETECTED

Explanation: The Multi-site Workload Lifeline application configuration file contains one or more statements that resulted in warning messages being logged. These warnings do not cause the application to end, but might indicate conditions that affect the normal operation of the application.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing continues.

Operator response: Take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, dump and save the syslogd file.
- Contact the system programmer.

System programmer response: If syslogd was started, examine the syslogd file for the application configuration warning messages. The syslogd identifier is lladvrs for the Advisor and llagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages for the application configuration warning

AQS0118I • AQS0120I

messages. Correct the configuration file definitions that generated the warning messages and either restart the application or, if the *applname* value is LLADVSR and the corrections are for modifiable configuration definitions, issue the `MODIFY procname,REFRESH` command. See Chapter 3, “Configuring Multi-site Workload Lifeline,” on page 43 for more information about configuring the Advisor and Agent.

Module: LACONFIG, LMADV, LMCMD, LMCONFIG

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0117I LLADVSR CONFIGURATION WARNINGS DETECTED
```

AQS0118I ALL APPLICATIONS ARE QUIESCED BY OPERATOR

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a request to quiesce or enable a subset of the target applications. A `MODIFY procname,QUIESCE` or `MODIFY procname,ENABLE` command was issued. A previous `MODIFY procname,QUIESCE,SYSTEM` command was issued to quiesce all applications for this z/OS Agent. The z/OS Agent enforces a hierarchy so that a previous SYSTEM action cannot be overridden by a TCPNAME or PORT request.

System action: The MODIFY command is rejected because all applications have been quiesced by the `MODIFY procname,QUIESCE,SYSTEM` command.

Operator response: Issue the `MODIFY procname,DISPLAY,MEMBERS` command to list the registered members and the status flags for each member. The operator quiesce flag (SYSQ) indicates that an MVS operator quiesced all members. If you do not want some of the members to be quiesced, issue the `MODIFY procname,ENABLE,SYSTEM` command to re-enable all members, and the `MODIFY procname,QUIESCE,TCPNAME=stackname` command or the `MODIFY procname,QUIESCE,PORT=portnum` command to quiesce a subset of the members. See Chapter 8, “Lifeline z/OS Agent commands,” on page 151 for more information.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0120I ALL APPLICATIONS FOR *stackname* ARE QUIESCED BY OPERATOR

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a request to quiesce or enable a specific target application. A `MODIFY procname,QUIESCE` or `MODIFY procname,ENABLE` command was issued. A previous `MODIFY procname,QUIESCE,TCPNAME=stackname` command was issued to quiesce all applications for the TCP/IP stack. The z/OS Agent enforces a hierarchy so that a previous TCPNAME action cannot be overridden by a PORT request.

In the message text:

stackname

The name of the TCP/IP stack.

System action: The MODIFY command is rejected because all applications for this TCP/IP stack have been quiesced by the `MODIFY procname,QUIESCE,TCPNAME=stackname` command.

Operator response: Issue the `MODIFY procname,DISPLAY,MEMBERS` command to list the registered members and the status flags for each member. The operator quiesce flag (TCPQ) indicates that an MVS operator quiesced all members for the TCP/IP stack. If you do not want some of the members to be quiesced, issue the `MODIFY procname,ENABLE,TCPNAME=stackname` command to re-enable all members for this TCP/IP stack, and the `MODIFY procname,QUIESCE,PORT=portnum` command to quiesce a subset of the members. See Chapter 8, “Lifeline z/OS Agent commands,” on page 151 for more information.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0120I ALL APPLICATIONS FOR TCPIP ARE QUIESCED BY OPERATOR

AQS0121I NETWORK ACCESS TO PEER ADVISOR ON *sysname* HAS RECOVERED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it has reestablished network communication with its peer Advisor.

In the message text:

sysname

The MVS system name where the peer Advisor is active.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0121I NETWORK ACCESS TO PEER ADVISOR ON CNTL2 HAS RECOVERED

AQS0122I NO MEMBERS MATCH SELECTION CRITERIA

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a MODIFY *procname*, QUIESCE or a MODIFY *procname*, ENABLE command. The MODIFY command is rejected because no active registered target applications match the criteria specified in the MODIFY command.

This message is also issued in response to a MODIFY *procname*, QUIESCE, TCP=*stackname* or MODIFY *procname*, ENABLE, TCP=*stackname* command when the *stackname* value does not match a TCP/IP stack that is active on the z/OS Agent system.

System action: The MODIFY command is ignored. Processing continues.

Operator response: Issue the MODIFY *procname*, DISPLAY, MEMBERS command to list the registered members on the local MVS system and the status flags for each member. The operator quiesce flag (SYSQ, TCPQ, or APPQ) indicates that an MVS operator quiesced the member. The absence of this flag indicates that the member is enabled from an MVS operator perspective. If the flag value is not the value that you want, issue a MODIFY *procname*, QUIESCE or a MODIFY *procname*, ENABLE command to change the flag value. See Chapter 8, "Lifeline z/OS Agent commands," on page 151 for more information.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0123I *cmd* COMMAND REJECTED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a MODIFY *procname*, QUIESCE or a MODIFY *procname*, ENABLE command. The MODIFY command is rejected. A more specific error message precedes this message and explains the reason for the command rejection.

In the message text:

cmd

The command that was rejected. Possible values are:

- QUIESCE TCPNAME
- ENABLE TCPNAME
- QUIESCE APPLICATION
- ENABLE APPLICATION

System action: The MODIFY command is ignored. Processing continues.

Operator response: Look for a previous message that provides more specific information about why the MODIFY command was rejected. Issue the MODIFY *procname*, DISPLAY, MEMBERS command to list the registered members on the local MVS system and the status flags for each member. The operator quiesce flag (SYSQ, TCPQ, or APPQ) indicates that an MVS operator quiesced the member. The absence of this flag indicates that the member is enabled from an MVS operator perspective. If the flag value is not the value that you want, issue a MODIFY *procname*, QUIESCE or a MODIFY *procname*, ENABLE command to change the flag value. See Chapter 8, "Lifeline z/OS Agent commands," on page 151 for more information.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0123I QUIESCE APPLICATION COMMAND REJECTED

AQS0124I NETWORK ACCESS TO PEER ADVISOR ON *sysname* NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it can no longer communicate with its peer Advisor over the network. If the Lifeline Advisor that issues this message is the secondary Advisor, AQS0164W was issued before this message.

In the message text:

sysname

The MVS system name where the peer Advisor was last active.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Verify the network connectivity between this Advisor and the peer Advisor. Ensure that the Advisor configuration files are identical between this Advisor and the peer Advisor. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing connectivity problems.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically issue the MODIFY *procname*, TAKEOVER command to have this Advisor assume the role of primary Advisor.

Example:

AQS0124I NETWORK ACCESS TO PEER ADVISOR ON CNTL2 NOT AVAILABLE

AQS0125I BCPII ACCESS TO PEER ADVISOR ON *sysname* NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it can no longer communicate with the image (LPAR) where the peer Advisor was active by using the z/OS Base Control Program Internal Interface (BCPii) service.

In the message text:

sysname

The MVS system name where the peer Advisor was last active.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Verify that z/OS is operational on the peer Advisor's image. Verify the connectivity used by BCPii to communicate with the peer Advisor's image. See *z/OS MVS Programming: Callable Services for High-Level Languages* for information about diagnosing BCPii issues.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0125I BCPII ACCESS TO PEER ADVISOR ON CNTL2 NOT AVAILABLE

AQS0126I BCPII ACCESS TO PEER ADVISOR ON *sysname* HAS RECOVERED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it has reestablished communication with the image (LPAR) where the peer Advisor was active by using the z/OS Base Control Program Internal Interface (BCPii) service.

In the message text:

sysname

The MVS system name where the peer Advisor is active.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0126I BCPII ACCESS TO PEER ADVISOR ON CNTL2 HAS RECOVERED

AQS0127W *applname* CONNECTION TO ADVISOR AT *ipaddress* CLOSED DUE TO PROTOCOL ERROR

Explanation: The Multi-site Workload Lifeline Advisor that was started as the secondary Advisor or Agent issues this message when it closes its connection to the primary Advisor because of an internal protocol error. This message will be deleted and message AQS0129I will be issued when the secondary Advisor or Agent successfully reconnects to the primary Advisor.

In the message text:

applname

The application name. Possible values are:

AQS0128I

LLADVSR

For the Multi-site Workload Lifeline secondary Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

ipaddress

The IP address of the primary Advisor.

System action: The secondary Advisor or Agent writes messages to internal trace or the syslogd file and attempts to reconnect to the primary Advisor.

Operator response:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, dump and save the syslogd file.
- If the packet trace or network sniffer trace is active, dump and save the trace.
- Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace message. Otherwise, examine the syslogd files for the Advisor and Agent. The syslogd identifier is lladvsr for the Advisor and llagent for the Agent. If the debug level includes messages at the message level, the file will show the data that each application sent or received. If the data that was sent by one application was the same data received by the corresponding application, contact IBM software support services. If the data that was sent was not the same data that was received, this might indicate that the network integrity has been compromised. If the packet trace is active, examine the trace for the data that was sent and received. If the packet trace or network sniffer trace is not active, try to recreate the problem with active packet trace or network sniffer trace. See *z/OS Communications Server: IP Diagnosis Guide* for information about the packet trace.

Module: LMAARP, LAMAIN

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically create snapshots of the Advisor and Agent syslogd or internal trace files for later problem determination.

Example:

```
AQS0127W LLADVSR CONNECTION TO ADVISOR AT 10.10.10.1 CLOSED DUE TO PROTOCOL ERROR
```

AQS0128I *peername* CONNECTION FROM *ipaddress* CLOSED DUE TO INACTIVITY

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the Advisor closes its connection to the Multi-site Workload Lifeline Agent or peer Advisor application because this Advisor did not receive a message from the application within the expected time.

In the message text:

peername

The application name. Possible values are:

ADVISOR

For the Multi-site Workload Lifeline Advisor

AGENT

For the Multi-site Workload Lifeline Agent

ipaddress

The IP address of the peer Advisor or Agent

System action: The Advisor closes the connection with the application. The peer Advisor or Agent attempts to reconnect to the Advisor until it is successful or is stopped. If the situation is temporary, the peer Advisor or Agent will successfully reconnect to the Advisor.

Operator response: If the peer Advisor or Agent is not active, start it. If the error persists, contact the system programmer.

System programmer response: Network connectivity problems, routing problems, slow system performance on the

peer Advisor or Agent system, and low MVS dispatching priority for the application might cause this problem. If the error persists, increase the value of the `advisor_timeout` or `updated_interval` statement in the Advisor configuration file and restart the Advisor. See Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 for more information about configuration statements. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems.

Module: LMADV, LMAGNT

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0128I ADVISOR CONNECTION FROM 10.10.10.1 CLOSED DUE TO INACTIVITY

AQS0129I *applname* CONNECTED TO ADVISOR AT *ipaddress*

Explanation: The Multi-site Workload Lifeline secondary Advisor or Agent issues this message when it is connected to the Multi-site Workload Lifeline primary Advisor.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline secondary Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

ipaddress

The IP address of the primary Advisor

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LAADV, LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0129I LLADVSR CONNECTED TO ADVISOR AT 10.10.10.1

AQS0130I *applname* CONNECTION TO ADVISOR AT *ipaddress* IS NO LONGER ACTIVE

Explanation: The Multi-site Workload Lifeline Advisor or Agent issues this message when it has lost its connection to the peer Multi-site Workload Lifeline Advisor.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline primary or secondary Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

ipaddress

The IP address of the primary or secondary Advisor

AQS0131I • AQS0132I

System action: The Advisor or Agent attempts to reconnect to the other Advisor. If the situation is temporary, the Advisor or Agent will successfully reconnect to the other Advisor.

Operator response:

- If the other Advisor is not active, start it.
- If the Advisor is active, take the following actions:
 - If the internal trace is being used, dump and save the internal trace messages.
 - If the internal trace is not being used, save the syslogd file.
 - Contact the system programmer.

System programmer response: If syslogd was started, examine the syslogd file for Advisor or Agent error or warning messages. The syslogd identifier is lladvsr for the Advisor and llagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages for error or warning messages. Correct any errors and restart the application that is in error.

Module: LAMAIN, LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0130I LLADVSR CONNECTION TO ADVISOR AT 10.10.10.1 IS NO LONGER ACTIVE

AQS0131I AGENT CONNECTED FROM *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it is connected to a Multi-site Workload Lifeline Agent.

In the message text:

ipaddress

The IP address of the Agent.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMAGNT

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0131I AGENT CONNECTED FROM 10.10.10.1

AQS0132I AGENT CONNECTION FROM *ipaddress* IS NO LONGER ACTIVE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it lost its connection to the Multi-site Workload Lifeline Agent. This message usually indicates that the Agent has been stopped.

In the message text:

ipaddress

The IP address of the previously connected Agent.

System action: Processing continues. If the Agent is not stopped, the Agent attempts to reconnect to the Advisor. If the situation is temporary, the Agent will reconnect to the Advisor.

Operator response: Determine whether the Agent has ended. Restart the Agent if necessary. If the Agent cannot reconnect to the Advisor, take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, save the syslogd file.
- Contact the system programmer.

System programmer response: No action is required if the Agent was intentionally stopped; otherwise, if syslogd was started, examine the syslogd file for Advisor or Agent errors, or warning messages. The syslogd identifier is lladvsr for the Advisor and llagent for the Agent. Otherwise, if internal trace is being used, dump and examine the internal trace messages for Advisor or Agent errors, or warning messages. Correct any errors and restart the application that is in error.

Module: LMAGNT

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0132I AGENT CONNECTION FROM 10.10.10.1 IS NO LONGER ACTIVE

AQS0133I LOAD BALANCER CONNECTED FROM *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it is connected to a load balancer.

In the message text:

ipaddress

The IP address of the load balancer.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0133I LOAD BALANCER CONNECTED FROM 10.10.10.1

AQS0134I LOAD BALANCER CONNECTION FROM *ipaddress* IS NO LONGER ACTIVE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it lost its connection with a load balancer. This message usually indicates that the load balancer has been stopped or the load balancer was incorrectly configured.

In the message text:

ipaddress

The IP address of the previously connected load balancer.

System action: Processing continues. If the load balancer is still active, it might attempt to reconnect to the Advisor.

Operator response: Determine whether the load balancer is active. Restart the load balancer, if necessary. If the load balancer is active, contact the system programmer.

System programmer response: Check the load balancer and correct any configuration or connectivity problems. Restart the load balancer, if necessary. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems.

Module: LMLB

AQS0135W • AQS0136W

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0134I LOAD BALANCER CONNECTION FROM 10.10.10.1 IS NO LONGER ACTIVE

AQS0135W ADVISOR CONNECTION FROM *ipaddress* CLOSED DUE TO PROTOCOL ERROR

Explanation: The Multi-site Workload Lifeline Advisor that was started as the primary Advisor issues this message when it closes its connection to the secondary Advisor because of an internal protocol error. This message will be deleted when the secondary Advisor successfully reconnects to the Advisor.

In the message text:

ipaddress

The IP address of the secondary Advisor.

System action: The primary Advisor writes messages to internal trace or to the syslogd file and waits for the secondary Advisor to reconnect.

Operator response:

- If internal trace is being used, dump and save the internal trace message.
- If internal trace is not being used, save the syslogd file.
- If packet trace is active, save the trace.
- Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages. Otherwise, examine the syslogd files for the Advisor. The syslogd identifier is lladvr for the Advisor. If the debug level includes messages at the message level, the file will show the data that each application sent or received. If the data that was sent by one application was the same data received by the corresponding application, contact IBM software support services. If the data that was sent was not the same data that was received, this might indicate that the network integrity has been compromised. If the packet trace is active, examine the trace for the data that was sent and received. If the packet trace is not active, try to recreate the problem with active packet trace. See *z/OS Communications Server: IP Diagnosis Guide* for information about the packet trace.

Module: LMAARP

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically create snapshots of the Advisor and peer Advisor syslogd or internal trace files for later problem determination.

Example:

AQS0135W ADVISOR CONNECTION FROM 10.10.10.1 CLOSED DUE TO PROTOCOL ERROR

AQS0136W AGENT CONNECTION FROM *ipaddress* CLOSED DUE TO PROTOCOL ERROR

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it closes its connection to the Multi-site Workload Lifeline Agent because of an internal protocol error. This message will be deleted when the Agent successfully reconnects to the Advisor.

In the message text:

ipaddress

The IP address of the previously connected Agent.

System action: The Advisor writes messages to internal trace or to the syslogd file. If the Agent is still active, it attempts to reconnect to the Advisor until the Agent connects with the Advisor or the Agent is stopped. If the situation is temporary, the Agent will reconnect to the Advisor.

Operator response: Take the following actions:

- If internal trace is being used, dump and save the internal trace messages.

- If internal trace is not being used, save the syslogd file.
- If packet trace is active, save the trace.
- Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages. Otherwise, examine the syslogd files for the Advisor and Agent. The syslogd identifier is lladvr for the Advisor and llagent for the Agent. If the debug level includes messages at the message level, the file will show the data that each application sent or received. If the data that was sent by one application was the same data received by the corresponding application, contact IBM software support services. If the data that was sent was not the same data that was received, this might indicate that the network integrity has been compromised. If the packet trace is active, examine the trace for the data that was sent and received. If the packet trace is not active, try to recreate the problem with active packet trace. See *z/OS Communications Server: IP Diagnosis Guide* for information about the packet trace.

Module: LCZAP, LMAGNT

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically create snapshots of the Advisor and Agent syslogd or internal trace files for later problem determination.

Example:

```
AQS0136W AGENT CONNECTION FROM 10.10.10.1 CLOSED DUE TO PROTOCOL ERROR
```

AQS0137I *applname* ENDED ABNORMALLY

Explanation: The Multi-site Workload Lifeline application ended in response to an unexpected error.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: Processing ends.

Operator response: Take the following actions:

- Save the CEEDUMP and snap output or core dump.
- If internal trace is not being used, save the syslogd file.
- Contact the system programmer.

System programmer response: Contact IBM software support services.

Module: LCASSERT, LCERROR

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically create a snapshot of the Advisor or Agent syslogd or internal trace file for later problem determination.

Example:

```
AQS0137I LLADVSR ENDED ABNORMALLY
```

AQS0138I *applname* DEBUG LEVEL UPDATED

Explanation: The Multi-site Workload Lifeline application issues this message in response to a z/OS MODIFY *procname*,DEBUG,LEVEL= command or a Linux on System z `aq$llcmd debug,level=` command.

In the message text:

AQS0139I • AQS0140I

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: The command is accepted. Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD, LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0138I LLADVSR DEBUG LEVEL UPDATED

AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE *state*

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a MODIFY *procname*, ENABLE|QUIESCE, SYSTEM command.

In the message text:

state

The state of the registered applications on the system where the command was issued. Possible values are:

- ENABLED
- QUIESCED

System action: The MODIFY command is accepted. Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0139I ALL MATCHING REGISTERED APPLICATIONS ARE QUIESCED

AQS0140I *subtotal* OF *total* MATCHING REGISTERED APPLICATIONS ARE *state* - DISTRIBUTED DVIPA MEMBERS ARE IGNORED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message in response to a MODIFY *procname*, ENABLE|QUIESCE, PORT|TCPNAME command.

In the message text:

subtotal

The count of matching applications that can be quiesced or enabled by the z/OS Agent.

total

The number of matching applications. This count includes the applications that can be quiesced or enabled by the z/OS Agent, and the applications that can only be quiesced or enabled by the VARY TCP/IP, SYSPLEX, QUIESCE|RESUME command.

state

The state of the registered applications on the system where the command was issued. Possible values are:

- ENABLED
- QUIESCED

System action: The MODIFY command is accepted. Processing continues.

Operator response: If the *subtotal* value does not equal the *total* value, the remaining matching applications can be quiesced or enabled by TCPIP using only the VARY TCPIP,,SYSPLEX,QUIESCE|RESUME command.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0140I 4 OF 5 MATCHING REGISTERED APPLICATIONS ARE QUIESCED - DISTRIBUTED DVIPA MEMBERS ARE IGNORED

AQS0141I ADVISOR SUMMARY

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,ADVISOR command. This message is followed by summary information about this Advisor. See "MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0142I ADVISOR DETAILS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,ADVISOR,DETAIL command. This message is followed by detailed information about this Advisor. See "MODIFY *advisor_proc*,DISplay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0143I ADVISOR CONFIGURATION

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,CONFIG command. This message is followed by the active configuration for this Advisor. See "MODIFY *advisor_proc*,DISPlay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0144I AGENT CONFIGURATION

Explanation: The Multi-site Workload Lifeline Agent issues this message in response to a z/OS MODIFY *procname*,DISPLAY,CONFIG command or a Linux on System z *qsllcmd display,config* command. This message is followed by the active configuration for this Agent. See "MODIFY *agent_proc*,DISPlay command: Display Lifeline Agent information" on page 151, or "Display Linux on System z Management Agent information" on page 165 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LACMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0145I WORKLOAD SUMMARY

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,WORKLOAD command. This message is followed by summary information about the workloads defined to this Advisor. See "MODIFY *advisor_proc*,DISPlay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0146I WORKLOAD DETAILS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DISPLAY,WORKLOAD,DETAIL command. This message is followed by detailed information about the workloads defined to this Advisor. See "MODIFY *advisor_proc*,DISPlay command: Display Lifeline Advisor information" on page 115 for the explanation of the display output.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 5, 8, 9

Automation: Not applicable for automation.

AQS0147I BCPII ACCESS TO SITE *sitename* NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it can no longer communicate with any of the images (LPARs) in the site specified by the *sitename* value by using the z/OS Base Control Program Internal Interface (BCPii) service.

In the message text:

sitename

The sysplex name of the site to which BCPii access is not available.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response:

- Verify that z/OS is operational on at least one of the site's images.
- Verify the connectivity used by BCPii to communicate with each of the site's images.

See *z/OS MVS Programming: Callable Services for High-Level Languages* for information about diagnosing BCPii issues.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0147I BCPII ACCESS TO SITE SITE02 NOT AVAILABLE

AQS0148I BCPII ACCESS TO SITE *sitename* HAS RECOVERED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it detects that it has reestablished communication with at least one image (LPAR) in the site specified by the *sitename* value by using the z/OS Base Control Program Internal Interface (BCPii) service.

In the message text:

sitename

The sysplex name of the site to which BCPii access has recovered.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0153I • AQS0155I

AQS0148I BCPII ACCESS TO SITE SITE02 HAS RECOVERED

AQS0153I WORKLOAD *workloadname* *state* ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE|DEACTIVATE|QUIESCE,WORKLOAD=*workloadname*,SITE=*sitename* command.

In the message text:

workloadname

The name of the workload specified on the MODIFY command.

state

The state of the workload. Possible values are:

- ACTIVATED
- DEACTIVATED
- QUIESCED

sitename

The sysplex name of the site.

System action:

- When the *state* value is QUIESCED,
 - For workloads defined in the cross_sysplex_list Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the workload specified by the *workloadname* value must not be routed to the site specified by the *sitename* value. Connection timeouts or resets might occur on the client applications that initiated the requests while the workload is in a quiesced state.
- When the *state* value is DEACTIVATED,
 - For workloads defined in the cross_sysplex_list Advisor configuration statement, all active connections to server applications on the site specified by the *sitename* value that are defined in the workload are reset.
- When the *state* value is ACTIVATED,
 - For workloads defined in the cross_sysplex_list or sna_application_list Advisor configuration statement, the Advisor notifies load balancers that new connection request to server applications for the workload specified by the *workloadname* value must be routed to the site specified by the *sitename* value.
 - For workloads defined in the mq_manager_list Advisor configuration statement, the Advisor notifies z/OS Lifeline Agents on the site, specified by the *sitename* value, to allow new messages for the workload, specified by the *workloadname* value, to be routed to this site.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0153I WORKLOAD WORK01 ACTIVATED ON SITE SITE01

AQS0155I WORKLOAD *workloadname* *state*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE|DEACTIVATE|QUIESCE,WORKLOAD=*workloadname* command.

In the message text:

workloadname

The name of the workload specified on the MODIFY command.

state

The state of the specified workload. Possible values are:

- ACTIVATED
- QUIESCED
- DEACTIVATED

System action:

- When the *state* value is QUIESCED,
 - For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the workload specified by the *workloadname* value must not be routed to any site. Connection timeouts or resets might occur on the client applications that initiated the requests while the workload is in a quiesced state.
 - For workloads defined in the *mq_manager_list* Advisor configuration statement, the Advisor notifies z/OS Lifeline Agents on the currently active site to prevent new messages for the workload, specified by the *workloadname* value, to be routed to any site.
- When the *state* value is DEACTIVATED,
 - For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, all active connections to server applications that are defined in the workload are reset.
 - For workloads defined in the *mq_manager_list* Advisor configuration statement, the Advisor notifies z/OS Agents on the previously active site to check for existing messages on each MQ cluster queue in the workload, specified by the *workloadname* value.
- When the *state* value is ACTIVATED,
 - For workloads defined in the *cross_sysplex_list* Advisor configuration statement, the Advisor notifies load balancers that new connection request to server applications for the workload specified by the *workloadname* value must be routed to both sites.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0155I WORKLOAD WORK01 QUIESCED
```

AQS0156I WORKLOAD *workloadname* IS NOT DEFINED TO SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE command. The Advisor configuration file contains a *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* definition for the workload specified by the *workloadname* value, but there is no definition of this workload for the site specified by the *sitename* value.

In the message text:

workloadname

The name of the workload specified on the MODIFY command.

sitename

The name of the site specified on the MODIFY command.

System action: The MODIFY command is ignored. Processing continues.

Operator response: Take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, save the syslogd file.
- Contact the system programmer.

System programmer response:

- If internal trace is being used, dump and examine the internal trace messages.

AQS0157I • AQS0158I

- If internal trace is not being used, examine the syslogd file. The syslogd identifier is lladvsr for the Advisor.
- Examine the internal trace messages or syslogd file for the Advisor configuration warning messages that were generated while the `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` definitions were being processed. Correct the configuration file warnings and either restart the Advisor or issue the `MODIFY procname,REFRESH` command. See Chapter 3, “Configuring Multi-site Workload Lifeline,” on page 43 for more information about configuring the Advisor.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0156I WORKLOAD WORK01 IS NOT DEFINED TO SITE SITE01
```

AQS0157I NETWORK ACCESS TO SITE *sitename* NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it can no longer communicate over the network with any Multi-site Workload Lifeline Agents in the site specified by the *sitename* value.

In the message text:

sitename

The sysplex name of the site to which network access is not available.

System action: Processing continues. If the site specified by the *sitename* value represents the active site for some workloads, the Advisor continues to direct load balancers to route new connection requests to server applications for these workloads to this site.

Operator response: Contact the system programmer.

System programmer response: Verify network connectivity between the Advisor and the site specified by the *sitename* value. Verify that Agents are running on each LPAR and Linux on System z Management Guest in the site where a server application for the workloads resides. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0157I NETWORK ACCESS TO SITE SITE01 NOT AVAILABLE
```

AQS0158I STANDBY SITE *sitename* HAS RECOVERED FOR WORKLOAD *workloadname*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Standby workload in the Advisor configuration file.
- Message AQS0159I was previously issued for the workload.
- If the workload is specified in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, at least one server application is available in the standby site specified by the *sitename* value for each load balancer application group defined for the workload specified by the *workloadname* value.
- If the workload is specified in the `mq_manager_list` Advisor configuration statement, at least one MQ cluster queue is available in the standby site specified by the *sitename* value for each MQ cluster queue group defined for the workload specified by the *workloadname* value.

In the message text:

sitename

The sysplex name of the standby site.

workloadname

The name of the workload that is defined in the Advisor configuration file.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0158I STANDBY SITE SITE02 HAS RECOVERED FOR WORKLOAD WORK01

AQS0159I STANDBY SITE *sitename* HAS UNAVAILABLE SERVERS FOR WORKLOAD *workloadname*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Standby workload in the Advisor configuration file.
- If the workload is specified in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, no server applications are available in the standby site specified by the *sitename* value for at least one load balancer application group defined for the workload specified by the *workloadname* value.
- If the workload is specified in the *mq_manager_list* Advisor configuration statement, no MQ cluster queue is available in the standby site specified by the *sitename* value for at least one MQ cluster queue group defined for the workload specified by the *workloadname* value.

In the message text:

sitename

The sysplex name of the standby site.

workloadname

The name of the workload that is defined in the Advisor configuration file.

System action: The site is no longer a usable target for the workload specified by the *workloadname* value.

Operator response: Notify the system programmer.

System programmer response: Ensure that server applications for the workload are active on the standby site specified by the *sitename* value.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0159I STANDBY SITE SITE02 HAS UNAVAILABLE SERVERS FOR WORKLOAD WORK01

AQS0160I NETWORK ACCESS TO SITE *sitename* HAS RECOVERED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it reestablishes a connection with at least one Multi-site Workload Lifeline Agent in the site specified by the *sitename* value.

In the message text:

AQS0161I • AQS0162I

sitename

The sysplex name of the site to which network access has recovered.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LMAGNT, LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0160I NETWORK ACCESS TO SITE SITE01 HAS RECOVERED

AQS0161I UNABLE TO CONNECT TO PEER ADVISOR AFTER *time* SECONDS

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it is started as the primary Advisor and it does not detect another Advisor acting as the primary. When a configured primary Advisor is started, it always verifies that there is no other primary Advisor active before becoming the primary Advisor.

In the message text:

time

The number of seconds the Advisor waited trying to connect to the peer Advisor.

System action: The Advisor assumes the role of primary Advisor.

Operator response: None.

System programmer response: None.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0161I UNABLE TO CONNECT TO PEER ADVISOR AFTER 10 SECONDS

AQS0162I MODIFY *cmd* COMMAND QUEUED - *reason*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE|DEACTIVATE|QUIESCE command. Processing of the command is delayed until the specified reason has been resolved.

In the message text:

cmd

The MODIFY command that was issued. Possible values are:

- ACTIVATE
- DEACTIVATE
- QUIESCE

reason

The reason that the command was not immediately processed. Possible values are:

NO FIRST-TIER LOAD BALANCER IS REGISTERED

A MODIFY *procname*,ACTIVATE or QUEISCE command was issued, but no first-tier load balancer is currently connected and registered with the Advisor.

DEACTIVATE COMMAND IN PROGRESS

A MODIFY *procname*,ACTIVATE or MODIFY *procname*,QUIESCE command was issued, but a prior MODIFY DEACTIVATE command has not yet completed.

UNKNOWN LATENCY STATE FOR WORKLOAD *workloadname*

A MODIFY *procname*,ACTIVATE command was issued for the workload, specified by the *workloadname* value, but the database replication latency state for the workload is unknown.

QUIESCE FOR WORKLOAD *workloadname* QUEUED

A MODIFY *procname*,DEACTIVATE command was issued for the workload, specified by the *workloadname* value, but a prior MODIFY QUIESCE command has not been processed.

QUIESCE FOR ALL MATCHING WORKLOADS QUEUED

A MODIFY *procname*,DEACTIVATE command was issued for all Active/Standby or all Active/Query workloads, but a prior MODIFY QUIESCE command has not been processed.

ACTIVATE OR DEACTIVATE COMMAND IN PROGRESS

A MODIFY *procname*,ACTIVATE command was issued, but a prior ACTIVATE or DEACTIVATE command has not yet completed.

System action: Processing continues.

A MODIFY QUIESCE or MODIFY ACTIVATE command is not processed until the following conditions are true:

- A first-tier load balancer registers.
- A prior MODIFY DEACTIVATE command completes.

A MODIFY ACTIVATE command for an Active/Query workload is delayed if the database replication latency state is unknown. The workload will be activated when either of the following conditions occurs:

- The Advisor receives the first event over its Network Management Interface for the Active/Query workload. The event indicates that the database replication latency state for the workload specified by the *workloadname* value is no longer unknown.
- The failure detection interval for the workload specified by the *workloadname* value expires. Because the database replication latency state is still unknown, the state will be treated as acute.

A MODIFY DEACTIVATE command is not processed until a queued QUIESCE command for the workload specified by the *workloadname* value completes.

Operator response: Take the appropriate action, based on the *reason* value:

NO FIRST-TIER LOAD BALANCER IS REGISTERED

Contact the system programmer.

DEACTIVATE COMMAND IN PROGRESS

The MODIFY command will be processed when the DEACTIVATE command completes.

UNKNOWN LATENCY STATE FOR WORKLOAD *workloadname*

Contact the system programmer.

QUIESCE FOR WORKLOAD *workloadname* QUEUED

The MODIFY command will be processed when the QUIESCE command completes.

QUIESCE FOR ALL MATCHING WORKLOADS QUEUED

The MODIFY command will be processed when the QUIESCE command completes.

ACTIVATE OR DEACTIVATE COMMAND IN PROGRESS

The MODIFY ACTIVATE command will be processed when the ACTIVATE or DEACTIVATE command completes.

System programmer response: Take the appropriate action, based on the *reason* value:

NO FIRST-TIER LOAD BALANCER IS REGISTERED

Verify the network connectivity between the Advisor and the first-tier load balancer. Check the load balancer and correct any configuration or connectivity problems. If necessary, restart the load balancer. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing network connectivity problems.

AQS0163W • AQS0164W

UNKNOWN LATENCY STATE FOR WORKLOAD *workloadname*

Determine why the database replication latency state is unknown, and correct, if possible.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0162I MODIFY QUIESCE COMMAND QUEUED - DEACTIVATE COMMAND IN PROGRESS
```

AQS0163W WAITING FOR FIRST-TIER LOAD BALANCER REGISTRATION

Explanation: The Multi-site Workload Lifeline Advisor issues this message when no first-tier load balancers are connected and registered to it. This message will be deleted after the first first-tier load balancer connects and registers with the Advisor.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Verify network connectivity between the Advisor and first-tier load balancer. Check the load balancer and correct any configuration or connectivity problems. Restart the load balancer if necessary. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing network connectivity problems.

Module: LMMAIN

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

AQS0164W UNABLE TO CONNECT TO PEER ADVISOR AFTER *time* SECONDS

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it is started as the secondary Advisor and is unable to connect to another Advisor acting as the primary Advisor. The message is deleted when either the peer Advisor is contacted or this Advisor assumes the role of primary Advisor.

In the message text:

time

The number of seconds the Advisor waited to try to connect to the peer Advisor.

System action: If the Advisor `advisor_takeover_policy` configuration statement is configured as manual, the Advisor remains in the secondary role. If `advisor_takeover_policy` is configured as automatic, the Advisor assumes the role of primary Advisor.

Operator response: If `advisor_takeover_policy` is configured as manual, issue the `MODIFY procname,TAKEOVER` command to have this Advisor assume the role of primary Advisor; otherwise, contact the system programmer.

System programmer response: Ensure that the Advisor started procedure is configured correctly (started in primary or secondary role). Verify the network connectivity between this Advisor and the peer Advisor. Ensure that the Advisor configuration files are identical between this Advisor and the peer Advisor. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing network connectivity problems.

Module: LMADV

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically issue the `MODIFY procname,TAKEOVER` command to have this Advisor assume the role of primary Advisor.

Example:

AQS0164W UNABLE TO CONNECT TO PEER ADVISOR AFTER 10 SECONDS

AQS0165I PEER ADVISOR CONNECTED FROM *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when a peer Advisor has established a connection with it.

In the message text:

ipaddress

The IP address of the peer Advisor.

System action: Processing continues if the peer Advisor is configured as the secondary Advisor. If the peer Advisor is configured as the primary Advisor, this Advisor will assume the role of secondary Advisor.

Operator response: None.

System programmer response: None.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0165I PEER ADVISOR CONNECTED FROM 10.10.10.1

AQS0166I ADVISOR IS NOW IN PRIMARY ROLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it assumes the role of primary Advisor. The Advisor assumes the role in one of the following situations:

- The Advisor procedure was configured to start the Advisor as the primary Advisor.
- Communication was lost with the peer Advisor that is acting as the primary Advisor and this Advisor was configured to automatically take over the role of primary Advisor.
- A MODIFY *procname*, TAKEOVER command was issued.

System action: The Advisor assumes the role of primary Advisor and is able to accept connections from load balancers and Multi-site Workload Agents.

Operator response: None.

System programmer response: None.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically determine when to use MODIFY commands that are valid only for the primary Advisor.

AQS0167I CONNECTION FROM *ipaddress* REJECTED - *reason*

Explanation: The Multi-site Workload Lifeline primary Advisor issues this message when it receives a connection request from a peer Advisor or Agent and the request is not accepted.

In the message text:

ipaddress

The IP address of the peer Advisor or Agent.

reason

The reason the connection was rejected. Possible values are:

AQS0167I

UNAUTHORIZED ADVISOR

One of the following conditions is true:

- The IP address of the peer Advisor is not included in the Advisor `advisor_id_list` configuration statement.
- The IP address used to connect to this Advisor does not match the IP address in the Advisor `lb_connection_v4` or `lb_connection_v6` configuration statement.
- The Advisor requires a secure connection with the peer Advisor by using an Application Transparent TLS policy.

SECONDARY ADVISOR ALREADY REGISTERED

This Advisor is already communicating with another peer Advisor; only one peer connection is allowed.

SECONDARY ADVISOR AT UNSUPPORTED LEVEL

The secondary Advisor is at an earlier version or service level than the primary Advisor. The secondary Advisor must be at the same or later version or service level.

AGENT AT UNSUPPORTED LEVEL

The Agent is at an earlier version than the primary Advisor no longer supports.

MULTIPLE AGENTS ACTIVE ON SYSTEM ID *system_id*

A Linux on System z Management Agent attempts to connect to the Advisor while another Linux on System z Management Agent on the same system ID, specified by *system_id*, and configured with the same site name is already connected.

System action: The connection from the primary Advisor is reset.

Operator response: Contact the system programmer.

System programmer response: Take the appropriate action, based on the *reason* value:

UNAUTHORIZED ADVISOR

For a non-secure connection, ensure that the IP address specified by the *ipaddress* value is included in the Advisor `advisor_id_list` configuration statement. Ensure that the IP address being used by the peer Advisor to connect to this Advisor matches the IP address in the Advisor `lb_connect_v4` or `lb_connection_v6` configuration statement. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

For a secure connection, verify that the Application Transparent TLS policy on both Advisors is correctly defined. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring Application Transparent TLS.

SECONDARY ADVISOR ALREADY REGISTERED

Ensure that the peer Advisor currently connected to this Advisor is the correct one. If not, stop this peer Advisor and start the correct peer Advisor.

SECONDARY ADVISOR AT UNSUPPORTED LEVEL

Upgrade the secondary Advisor to a version or service level that is the same with or later than the primary Advisor. See "Migrating the Lifeline Advisor" on page 53 for more information.

AGENT AT UNSUPPORTED LEVEL

Upgrade the Agent to a version that is supported by the primary Advisor. See "Migrating the Lifeline Agent" for more information.

MULTIPLE AGENTS ACTIVE ON SYSTEM ID *system_id*

Ensure that the Linux on System z Management Agent that is connected from the system ID, specified by *system_id*, is the correct one. If not, stop the existing Linux on System z Management Agent and start the correct Linux on System z Management Agent.

Module: LMADV, LMAGNT, LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0167I CONNECTION FROM 10.10.10.1 REJECTED - UNAUTHORIZED ADVISOR

AQS0168I CONNECTION FROM *ipaddress* REJECTED - *reason*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it receives a connection request from a load balancer or Agent and it is not accepted.

In the message text:

ipaddress

The IP address of the load balancer or Agent.

reason

The reason the connection was rejected. Possible values are:

UNAUTHORIZED AGENT

The IP address of the Agent is not included in the Advisor *agent_id_list* configuration statement, the IP address used to connect to this Advisor does not match the IP address in the Advisor *lb_connection_v4* or *lb_connection_v6* configuration statement, or for z/OS Agents, the Advisor requires a secure connection with the Agent by using an Application Transparent TLS policy.

UNAUTHORIZED LB

The IP address of the load balancer is not included in the Advisor *lb_id_list* configuration statement or the Advisor requires a secure connection with the load balancer by using an Application Transparent TLS policy.

UNAUTHORIZED SITE

The sysplex name of the site where the Agent resides is not defined in the Advisor *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* configuration statement.

INCORRECT LB REGISTRATION

One of the following events occurred:

- A first-tier load balancer attempted to register a group with an IP address or port that is not included in the Advisor *cross_sysplex_list* or *intermediary_node_list* configuration statement.
- A first-tier load balancer attempted to register a group with members for only one of the sites that are defined in the Advisor *cross_sysplex_list* or *intermediary_node_list* configuration statement.
- A first-tier load balancer attempted to register more than one group for an Active/Query workload.
- A second-tier load balancer attempted to register before a first-tier load balancer has registered.
- A second-tier load balancer attempted to register a group name that does not match any group name registered by a first-tier load balancer.
- A load balancer attempted to register a group with no associated port.

System action: The connection from the load balancer or Agent is reset.

Operator response: Contact the system programmer.

System programmer response: Take the appropriate action, based on the *reason* values:

UNAUTHORIZED AGENT

Ensure that the value that *ipaddress* specifies is included in the Advisor *agent_id_list* configuration statement. Ensure that the IP address being used by the Agent to connect to this Advisor matches the IP address in the Advisor *lb_connection_v4* or *lb_connection_v6* configuration statement.

For a secure connection with a z/OS Agent, verify that the Application Transparent TLS policy on both the Advisor and Agent is correctly defined. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring Application Transparent TLS.

UNAUTHORIZED LB

Ensure that the value that *ipaddress* specifies is included in the Advisor *lb_id_list* configuration statement.

For a secure connection, verify that the Application Transparent TLS policy on the Advisor and security configuration on the load balancer is correctly defined. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring Application Transparent TLS.

UNAUTHORIZED SITE

Ensure that the sysplex name where the Agent resides is included in the Advisor *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* configuration statement. Ensure that a Linux on System z Management Agent is not connecting to an Advisor that does not support Linux on System z workloads.

INCORRECT LB REGISTRATION

- Dump and examine the internal trace messages, if being used, or examine the syslogd file to determine which group is being registered incorrectly, and modify the first-tier load balancer to remove the group or ensure that the group is included in the Advisor `cross_sysplex_list` or `intermediary_node_list` configuration statement and contains members for each of the sites.
- Dump and examine the internal trace messages, if being used, or examine the syslogd file to determine which groups are being registered incorrectly for an Active/Query workload, and modify the first-tier load balancer to ensure all members for the workload are configured in a single group.
- Ensure that a first-tier load balancer is configured to connect to the Advisor.
- Dump and examine the internal trace messages, if being used, or examine the syslogd file to determine which first-tier load balancer application groups are registered, and modify the second-tier load balancer configuration to ensure that its groups match the registered first-tier load balancer application groups.
- Dump and examine the internal trace messages, if being used, or examine the syslogd file to determine which group is being registered without a port and modify the load balancer configuration to specify a port for that group.

See Chapter 6, “Multi-site Workload Lifeline configuration statements,” on page 85 for more information about configuration statements and parameters.

Module: LMAGNT, LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0168I CONNECTION FROM 10.10.10.1 REJECTED - UNAUTHORIZED AGENT
```

AQS0169I PRIMARY ADVISOR TAKEOVER SCHEDULED

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,TAKEOVER` command.

System action: The Advisor attempts to communicate with the peer Advisor to ensure that this Advisor will be the only one to assume the role of primary Advisor. Message AQS0166I will be issued when the takeover has completed.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0170I MODIFY TAKEOVER COMMAND IGNORED -reason

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,TAKEOVER` command.

In the message text:

reason

The reason the `MODIFY TAKEOVER` was ignored. Possible values are:

ALREADY PRIMARY

The Advisor is already the primary Advisor.

ALREADY SCHEDULED

A previous `MODIFY procname,TAKEOVER` command was issued and has not yet completed.

System action: The `MODIFY` command is ignored.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0170I MODIFY TAKEOVER COMMAND IGNORED - ALREADY PRIMARY

AQS0172I TAKEOVER PROCESSING INCOMPLETE - UNRESPONSIVE PEER AT *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*, TAKEOVER command. The Advisor attempted to notify the peer Advisor that this Advisor will become the primary Advisor. The connection was lost before the notification completed.

In the message text:

ipaddress

The last known IP address of the peer Advisor.

System action: The Advisor assumes the role of primary Advisor, but the role of the peer Advisor is unknown. Message AQS0166I will be issued when the takeover has completed.

Operator response: Contact the system programmer.

System programmer response: Verify the network connectivity between this Advisor and the peer Advisor. If the peer Advisor is not active, restart it. If the peer Advisor is active, verify that it is running as the secondary Advisor. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing network connectivity problems.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0172I TAKEOVER PROCESSING INCOMPLETE - UNRESPONSIVE PEER AT 10.10.10.1

AQS0173I ADVISOR IS NOW IN SECONDARY ROLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it assumes the role of secondary Advisor. The Advisor assumes the role of secondary Advisor when one of the following has occurred; the Advisor procedure was configured to start the Advisor as the secondary Advisor or a MODIFY *procname*, TAKEOVER command was issued on the peer Advisor.

System action: The Advisor assumes the role of secondary Advisor and attempts to connect to the peer Advisor.

Operator response: None.

System programmer response: None.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically determine when to use MODIFY commands that are valid only for the secondary Advisor.

AQS0174I TAKEOVER FAILED - ERROR ON PRIMARY ADVISOR AT *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*, TAKEOVER command. While attempting to communicate with the peer Advisor to ensure that this Advisor will be the only one to assume the role of primary Advisor, the peer Advisor did not relinquish its role as primary Advisor.

In the message text:

ipaddress

The IP address of the peer Advisor.

System action: This Advisor remains in the role of secondary Advisor and the peer Advisor is in an unknown role.

Operator response: Contact the system programmer.

System programmer response: Check for console messages. If syslogd was started, examine the syslogd file on the system of the peer Advisor. The syslogd identifier is lladvr for the Advisor. If internal trace is being used, dump and examine the internal trace messages on the system of the peer Advisor. If the peer Advisor is not active, restart it in the role of primary Advisor. If the peer Advisor is active, verify that it is running as the primary Advisor.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically stop the peer Advisor and restart this Advisor as the primary Advisor.

Example:

AQS0174I TAKEOVER FAILED - ERROR ON PRIMARY ADVISOR AT 10.10.10.1

AQS0175I TAKEOVER FAILED - ERROR ESTABLISHING LISTENING SOCKETS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*, TAKEOVER command. While attempting to become the primary Advisor, this Advisor was unable to establish listening sockets to allow for load balancers, Multi-site Workload Lifeline Agents, or a peer Advisor to connect to it.

System action: The Advisor remains in the role of secondary Advisor.

Operator response: Determine whether TCP/IP is available and restart if necessary. Reissue the MODIFY *procname*, TAKEOVER command. If TCP/IP is available when this message is issued, save the syslogd file. Otherwise, if internal trace is being used, dump and save the internal trace messages. Contact the system programmer.

System programmer response: To determine which listening sockets cannot be established, take the appropriate action:

- If TCP/IP is available when the Advisor issues this message and internal trace is being used, dump and examine the internal trace messages.
- If TCP/IP is available when the Advisor issues this message and internal trace is not being used, examine the syslogd file.
- The syslogd identifier is lladvr for the Advisor.

Correct the error.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically start the peer Advisor as the primary Advisor.

AQS0176I *type* STATE NOT REPLICATED TO SECONDARY ADVISOR

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it cannot send updated state information to the secondary Advisor.

In the message text:

type

The type of state. Possible values are:

- WORKLOAD
- BCPII

System action: Processing continues.

- If the *type* value is WORKLOAD, the primary and secondary Advisors are no longer in sync. If the secondary Advisor were to take over primary Advisor responsibilities, the state of the workloads is unknown.
- If the *type* value is BCPII and the secondary Advisor were to take over primary Advisor responsibilities, the system images that reside within the sites are unknown, but will be rediscovered as Agents from these system images reconnect to the Advisor.

Operator response: Contact the system programmer.

System programmer response: Verify the network connectivity between this Advisor and the secondary Advisor. If internal trace is being used, dump and examine the internal trace messages to determine the cause of the communication failure. Otherwise, examine the syslogd file to determine the cause of the communication failure. The syslogd identifier is lladvr for the Advisor. Restart the secondary Advisor to enable the two Advisors to resync. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0176I WORKLOAD STATE NOT REPLICATED TO SECONDARY ADVISOR

AQS0177I ADVISOR STOPPING – *reason*

Explanation: The Multi-site Workload Lifeline secondary Advisor issues this message during initialization when the peer Advisor resets the connection with this Advisor.

In the message text:

reason

The reason the secondary Advisor is stopping. Possible values are:

NOT AUTHORIZED TO PRIMARY ADVISOR

One of the following conditions is true:

- The IP address of this Advisor is not included in the peer Advisor `advisor_id_list` configuration statement.
- The IP address specified in the Advisor `lb_connection_v4` or `lb_connection_v6` configuration statement does not match what is defined in the peer Advisor's configuration.
- The peer Advisor requires a secure connection with this Advisor by using an Application Transparent TLS policy.

PRIMARY ADVISOR ALREADY HAS A SECONDARY REGISTERED

The peer Advisor is already communicating with another peer Advisor; only one peer connection is allowed.

UNSUPPORTED LEVEL

The secondary Advisor is at an earlier version or service level than the peer Advisor. The secondary Advisor must be at the same or a later version or service level.

System action: Processing ends. The connection with the peer Advisor is reset.

AQS0178I

Operator response: Contact the system programmer.

System programmer response: Take the appropriate action, based on the *reason* value:

NOT AUTHORIZED TO PRIMARY ADVISOR

For a non-secure connection, ensure that the Advisor's IP address is included in the peer Advisor's *advisor_id_list* configuration statement. Ensure that the Advisor's *lb_connection_v4* and *connection_v6* configuration statements match the peer Advisor's configuration statements. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

For a secure connection, verify that the Application Transparent TLS policy on both Advisors is correctly defined. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring Application Transparent TLS.

PRIMARY ADVISOR ALREADY HAS A SECONDARY REGISTERED

Ensure that the secondary Advisor connected to the peer Advisor is the correct one. If not, stop the current secondary Advisor and restart this Advisor as the secondary Advisor.

UNSUPPORTED LEVEL

Upgrade the secondary Advisor to a version or service level that is the same with or later than the peer Advisor. See "Migrating the Lifeline Advisor" on page 53 for more information.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0177I ADVISOR STOPPING – NOT AUTHORIZED TO PRIMARY ADVISOR

AQS0178I WORKLOAD *workloadname* MUST BE QUIESCED BEFORE ISSUING A *cmd* COMMAND

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,DEACTIVATE|ACTIVATE` command. The workload must be quiesced before you can deactivate it or reroute it to an alternate standby site.

In the message text:

cmd

The MODIFY command that was issued. Possible values are:

- ACTIVATE
- DEACTIVATE

workloadname

The name of the workload specified on the MODIFY command.

System action: The MODIFY command is ignored.

Operator response: Issue a `MODIFY procname,QUIESCE,WORKLOAD=workloadname` to quiesce the workload, and reissue the `MODIFY procname,DEACTIVATE|ACTIVATE` command.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0178I WORKLOAD WORK01 MUST BE QUIESCED BEFORE ISSUING A DEACTIVATE COMMAND

AQS0179I MODIFY *cmd* COMMAND IGNORED - *reason*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when a MODIFY *procname* command is ignored.

In the message text:

cmd

The MODIFY command that was issued. Possible values are:

- ACTIVATE
- DEACTIVATE

reason

The reason that the MODIFY command was ignored. Possible values are:

NO GROUP IS REGISTERED FOR WORKLOAD

The MODIFY *procname*,DEACTIVATE command was issued, but the load balancer that registered the load balancer application group that is used by that workload is not currently registered.

ADVISOR MUST BE PRIMARY

The MODIFY command specified by the *cmd* value was issued on the Advisor that was acting as the secondary Advisor.

NO AGENT IS CONNECTED

The MODIFY *procname*,DEACTIVATE command was issued, but no Multi-site Workload Lifeline Agents from either site are available to reset active connections on their systems.

NOT ALL REQUIRED AGENTS CONNECTED ON SITE *sitename*

The MODIFY *procname*,ACTIVATE command was issued to a specific site, specified by the *sitename* value, or the MODIFY *procname*,ACTIVATE command was issued without specifying a site, but not all required Multi-site Workload Lifeline Agents are available in the site, specified by the *sitename* value.

SNA WORKLOAD NOT DEACTIVATED ON SITE *sitename*

The MODIFY *procname*,ACTIVATE command was issued for a SNA workload, but the workload is not deactivated on the alternate site specified by the *sitename* value.

WORKLOAD *workloadname* HAS CONNECTIONS ON SITE *sitename*

The MODIFY *procname*,ACTIVATE command was issued for the workload specified by the *workloadname* value, but active TCP connections or SNA sessions for the workload exist on the alternate site, specified by the *sitename* value.

WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename*

The MODIFY *procname*,ACTIVATE command was issued for the workload specified by the *workloadname* value, but messages exist on at least one of the workload's MQ cluster queues hosted by an MQ queue manager on the alternate site specified by the *sitename* value.

WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename*

The MODIFY *procname*,ACTIVATE command was issued for the workload specified by the *workloadname* value, but the workload's MQ cluster receiver channel is not stopped on at least one of the MQ queue managers processing the workload on the alternate site specified by the *sitename* value.

System action: The MODIFY command is ignored.

Operator response: Take the following actions based on the *reason* value:

NO GROUP IS REGISTERED FOR WORKLOAD

Contact the system programmer.

ADVISOR MUST BE PRIMARY

Reissue the MODIFY command specified by the *cmd* value on the system where the primary Advisor is running.

NO AGENT IS CONNECTED

Contact the system programmer.

NOT ALL REQUIRED AGENTS CONNECTED ON SITE *sitename*

Issue a MODIFY *procname*,DISPLAY,ADVISOR command to determine which configured Agents are currently not connected. Contact the system programmer with the list of Agents that are not connected.

SNA WORKLOAD NOT DEACTIVATED ON SITE *sitename*

Issue a `MODIFY procname,DEACTIVATE,WORKLOAD=workloadname` command to inactivate SNA applications on the alternate site specified by the *sitename* value, and reissue the `MODIFY procname,ACTIVATE` command.

WORKLOAD *workloadname* HAS CONNECTIONS ON SITE *sitename*

Issue the `MODIFY procname,DEACTIVATE,WORKLOAD=workloadname` command to reset all active connections on the target site specified by the *sitename* value, and reissue the `MODIFY procname,ACTIVATE` command.

| **WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename***

| Contact the system programmer.

| **WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename***

| Contact the system programmer.

System programmer response: Take the following actions based on the *reason* value:

NO GROUP IS REGISTERED FOR WORKLOAD

Verify the network connectivity between the Advisor and the first-tier load balancer that contains the group for the workload. Check the load balancer and correct any configuration or connectivity problems. If necessary, restart the load balancer. See *z/OS Communications Server: IP Diagnosis Guide* for more information about diagnosing network connectivity problems. When the load balancer is registered, reissue the `MODIFY` command specified by the *cmd* value.

NO AGENT IS CONNECTED

Verify that at least one z/OS Agent is connected on both sites, and for Linux on System z workloads, at least one Linux on System z Management Agent is connected on both sites. Verify that the Agents on each site are active by issuing the `MODIFY procname,DISPLAY,ADVISOR` command and, if necessary, restart the Agents. Verify network connectivity between the Advisor and Agents. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing z/OS networking connectivity problems. See the documentation for your Linux distribution for more information about diagnosing network connectivity problems on Linux on System z.

NOT ALL REQUIRED AGENTS CONNECTED ON SITE *sitename*

Issue the `MODIFY procname,DISPLAY,ADVISOR` command to determine which required Agents are currently not connected. Verify that at least one z/OS Agent is connected on both sites, and for Linux on System z workloads, at least one Linux on System z Management Agent is connected on both sites. Verify that the Agents on each site are active and, if necessary, restart the Agents. Verify network connectivity between the Advisor and Agents. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing z/OS network connectivity problems. See the documentation for your Linux distribution for more information about diagnosing network connectivity problems on Linux on System z.

| **WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename***

| For each z/OS Agent on the site specified by the *sitename* value:

- | • If internal trace is being used, dump and examine the internal trace messages.
- | • Otherwise, examine the `syslogd` file. The `syslogd` identifier is `llagent` for the Agent.

| Check for error messages indicating that messages were found on an MQ cluster queue for the workload specified by the *workloadname* value. If MQ server applications are active and processing messages from these MQ cluster queues, wait until all queued messages are processed and reissue the `MODIFY procname,ACTIVATE` command. If the messages remain on the MQ cluster queues, perform one of the following steps:

- | • Start the MQ server applications needed to process these messages.
- | • Issue the `MODIFY procname,DEACTIVATE,WORKLOAD=workloadname` command to transition the workload into the DEACTIVATED state and reissue the `MODIFY procname,ACTIVATE` command. MQ message processing will start on the activated site even though messages remain on the quiesced site.

| **WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename***

| For each z/OS Agent on the site specified by the *sitename* value:

- | • If internal trace is being used, dump and examine the internal trace messages.
- | • Otherwise, examine the `syslogd` file. The `syslogd` identifier is `llagent` for the Agent.

| Check for error messages indicating that an MQ cluster's receiver channel has not stopped. Wait until the MQ cluster receiver channel stops and reissue the `MODIFY procname,ACTIVATE` command. If the MQ cluster receiver channel does not stop, verify that the MQ cluster receiver channel is correctly configured. If

possible, manually stop the MQ cluster receiver channel using the STOP CHANNEL command. See *WebSphere MQ: Administering IBM MQ* for information about IBM MQ commands.

Optionally, issue the MODIFY *procname*,DEACTIVATE,WORKLOAD=*workloadname* command to transition the workload into the DEACTIVATED state and reissue the MODIFY *procname*,ACTIVATE command. Note that although MQ message processing will start on the activated site, additional messages might continue to be received on the quiesced site.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0179I MODIFY QUIESCE COMMAND IGNORED - ADVISOR MUST BE PRIMARY

AQS0180W WORKLOAD *workloadname* **HAS UNAVAILABLE SERVERS ON SITE** *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Standby workload in the Advisor configuration file.
- For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, no server application is available in the site specified by the *sitename* value for at least one load balancer application group defined for the Active/Standby workload specified by the *workloadname* value.

A server application is considered unavailable in any of the following conditions is true:

- The server application is not active.
- The server application is active but not fully operational.
- The system on which the server application runs is resource constrained.
- For SNA and Linux on System z workloads, when either condition is true:
 - No members in the load balancer application group are available on the site specified by the *sitename* value.
 - No applications that the members in the load balancer application group distribute connections to are available on the site specified by the *sitename* value.
- For workloads defined in the *mq_manager_list* Advisor configuration statement, at least one MQ cluster queue is not available on any MQ queue manager in the site specified by the *sitename* value for the Active/Standby workload specified by the *workloadname* value. The *mq_manager_list* Advisor configuration statement specifies each MQ queue manager on the site that should have the MQ cluster queue defined.

A MQ cluster queue is considered unavailable if either of the following conditions is true:

- No MQ queue managers that host the MQ cluster queue are active.
- No active MQ queue managers have the MQ cluster queue defined.

For workloads defined in the *cross_sysplex_list* Advisor configuration statement that are not Linux on System z workloads, this message will be deleted when at least one server application in each load balancer application group becomes available on the site specified by the *sitename* value, or the workload is rerouted to the alternate standby site.

For workloads defined in the *mq_manager_list* Advisor configuration statement, the message will be deleted when all MQ cluster queues become available on the site specified by the *sitename* value, or the workload is rerouted to MQ queue managers on the alternate site.

For SNA and Linux on System z workloads, this message will be deleted when both conditions are true:

- At least one member in the load balancer application group is available on the site specified by the *sitename* value.
- At least one server application that the members in the load balancer application group distribute connections to becomes available on the site specified by the *sitename* value.

In the message text:

AQS0181I

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the active site where no servers are available for the workload.

System action: The site is no longer a usable target for the workload specified by the *workloadname* value.

- | If the Advisor *workload_switch_policy* configuration statement is configured as automatic:
- | • For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, the Advisor will automatically notify load balancers that new connection requests to server applications for the specified workload that are currently being routed to the site specified by the *sitename* value must be routed to the alternate standby site.
 - | • For workloads defined in the *mq_manager_list* Advisor configuration statement, the Advisor will automatically notify z/OS Agents that MQ cluster queue ranks should be altered so that messages for the specified workload that are currently being routed to MQ queue managers in the site specified by the *sitename* value must be routed to the MQ queue managers in the alternate standby site.
- | If the *workload_switch_policy* configuration statement is configured as manual:
- | • For workloads defined in the *cross_sysplex_list* or *sna_application_list* Advisor configuration statement, the Advisor continues notifying load balancers to route new connection requests to server applications for the workload specified by the *workloadname* value to the site specified by the *sitename* value.
 - | • For workloads defined in the *mq_manager_list* Advisor configuration statement, the Advisor continues notifying z/OS Agents that MQ cluster queue ranks should be unchanged so that messages continue to be routed for the workload specified by the *workloadname* value to the site specified by the *sitename* value.

Operator response: If the *workload_switch_policy* configuration statement is configured as manual, issue the *MODIFY procname,ACTIVATE* command for the workload specified by the *workloadname* value currently being routed to the site specified by the *sitename* value to notify the load balancers that new connection requests must be routed to the specified site. See Chapter 7, "Lifeline Advisor commands," on page 115 for more information. Contact the system programmer.

System programmer response: Ensure that the server applications for the workload are available on the active site specified by the *sitename* value. Verify that each server application is active and is fully operational, and that the system where it resides has sufficient resources. See *z/OS Communications Server: IP Diagnosis Guide* and *z/OS Communications Server SNA Diagnosis Vol 1, Techniques and Procedures* for more information about verifying z/OS server application availability. See the documentation for your Linux distribution for more information about verifying Linux on System z server application availability.

Module: LMLB

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically issue the *MODIFY procname,ACTIVATE* command to have the workload routed to the alternate site.

Example:

```
AQS0180W WORKLOAD WORK01 HAS UNAVAILABLE SERVERS ON SITE SITE01
```

AQS0181I ADVISOR CONFIGURATION UPDATED

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a *MODIFY procname,REFRESH* command.

System action: The Advisor configuration is updated.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0182I ADVISOR CONFIGURATION NOT UPDATED

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*, REFRESH command. The Advisor configuration file cannot be opened, or contains one or more errors.

System action: The MODIFY command is ignored.

Operator response: Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages for configuration error messages. Otherwise, examine the syslogd file for configuration error messages. The syslogd identifier is lladvr for the Advisor. Correct the configuration file errors and reissue the MODIFY *procname*, REFRESH command. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0184W PRIMARY AND SECONDARY ADVISOR CONFIGURATION NOT IN SYNC

Explanation: The Multi-site Workload Lifeline Advisor issues this message in one of the following cases:

- The secondary Advisor is started using a different configuration file than the primary Advisor.
- A MODIFY *procname*, REFRESH command issued from the peer Advisor results in the active configuration for the two Advisors to no longer be the same.

This message will be deleted when the active configuration of both Advisors is back in synchronization.

System action: Processing continues.

Operator response: Ensure that the Advisor configuration file is updated to match the peer Advisor and issue the MODIFY *procname*, REFRESH command from this Advisor.

System programmer response: None.

Module: LMADV

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

AQS0185I AGENT CONNECTION TO ADVISOR AT *ipaddress* REJECTED - *reason*

Explanation: The Multi-site Workload Lifeline Agent issues this message during initialization when the Multi-site Workload Lifeline Advisor resets the connection with this Agent.

In the message text:

ipaddress

The IP address of the Advisor.

reason

The reason the Agent connection was rejected. Possible values are:

SITE NOT DEFINED TO ADVISOR

The site where this Agent resides is not included in the Advisor *cross_sysplex_list*, *mq_manager_list*, or *sna_application_list* configuration statement.

NOT AUTHORIZED TO ADVISOR

One of the following conditions occurred:

AQS0186I

- The IP address defined in the Agent `host_connection` configuration statement is not included in the Advisor `agent_id_list` configuration statement.
- The IP address defined in the Agent `advisor_id` or `advisor_id_list` configuration statement does not match the Advisor `lb_connection_v4` or `lb_connection_v6` configuration statement.
- For z/OS Agents, the Advisor requires a secure connection with this Agent by using an Application Transparent TLS policy.

UNSUPPORTED LEVEL

The version of the Agent is unknown or no longer supported by the Advisor.

MULTIPLE AGENTS ACTIVE ON THIS SYSTEM ID

Another Linux on System z Management Agent from this system and site is already connected to the Advisor.

System action: Processing ends. The connection with the primary Advisor is reset.

Operator response: Contact the system programmer.

System programmer response: Take the appropriate action, based on the *reason* value:

SITE NOT DEFINED TO ADVISOR

Ensure that the Agent's site is included in the Advisor's `cross_sysplex_list`, `mq_manager_list`, or `sna_application_list` configuration statement. Ensure that a Linux on System z Management Agent is not connecting to an Advisor that does not support Linux on System z workloads. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

NOT AUTHORIZED TO ADVISOR

Ensure that the Agent's IP address defined in the `host_connection` configuration statement is included in the Advisor's `agent_id_list` configuration statement. Ensure that the Agent's IP address defined in the `advisor_id` or `advisor_id_list` configuration statement matches the Advisor's `lb_connection_v4` or `lb_connection_v6` configuration statement. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters.

For z/OS Agents that require a secure connection, verify that the Application Transparent TLS policy on both the Advisor and Agent is correctly defined. See *z/OS Communications Server: IP Configuration Guide* for more information about configuring Application Transparent TLS.

UNSUPPORTED LEVEL

Upgrade the Agent to a version supported by the Advisor. See "Migrating the Lifeline Agent" on page 54 for more information.

MULTIPLE AGENTS ACTIVE ON THIS SYSTEM ID

Ensure that the Linux on System z Management Agent currently connected from this system is the correct one. If not, stop the existing Linux on System z Management Agent and start the correct Linux on System z Management Agent.

Module: LAADV, LAMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0185I AGENT CONNECTION TO ADVISOR AT 10.10.10.1 REJECTED - NOT AUTHORIZED TO ADVISOR
```

AQS0186I ADVISOR NMI INTERFACE IS NO LONGER AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the Network Management Interface (NMI) is stopped.

System action: Processing continues. The NMI connections with network management agents are reset.

Operator response: Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages for

NMI error messages and correct the error. Otherwise, examine the syslogd file for NMI error messages and correct the error. The syslogd identifier is lladvr for the Advisor. Recycle the Advisor to restart the NMI.

Module: LMNMI

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0187I BCPII MONITORING IS NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when it is unable to use the Base Control Program Internal Interface (BCPii) service.

System action: Processing continues.

Operator response: Take the following actions:

- If internal trace is being used, dump and save the internal trace messages.
- If internal trace is not being used, save the syslogd file.
- Contact the system programmer.

System programmer response: Verify that the BCPii address space was successfully started. Ensure that the Lifeline Advisor is authorized to use BCPii services. If internal trace is being used, dump and examine the internal trace messages to determine the cause of the BCPii failure. Otherwise, examine the syslogd file to determine the cause of the BCPii failure. The syslogd identifier is lladvr for the Advisor. Correct the problem, and then stop and restart the Lifeline Advisor. See *z/OS MVS Programming: Callable Services for High-Level Languages* for information about diagnosing BCPii issues.

Module: LMBCPII

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0188W SITE *sitename* IS NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when both of the following conditions occur:

- The Advisor can no longer communicate over the network with any Multi-site Workload Lifeline Agents in the specified site.
- The Advisor can no longer communicate with any of the images (LPARs) in the specified site by using the Base Control Program Internal Interface (BCPii) service.

This message will be deleted when at least one Agent from this site reconnects to the Advisor or communication is reestablished with at least one image from this site by using BCPii services.

In the message text:

sitename

The sysplex name of the unresponsive site.

System action: The site is no longer considered responsive and is no longer a usable target for any defined workloads. If the Advisor `workload_switch_policy` configuration statement is configured with a value of `automatic`, the Advisor will automatically direct load balancers that new connection requests to server applications for workloads currently being routed to the site specified by the *sitename* value should be routed to the alternate, standby site. If `workload_switch_policy` is configured with a value of `manual`, the Advisor continues to direct load balancers to route new connection requests to server applications for these workloads to the site specified by the *sitename* value.

Operator response: If `workload_switch_policy` is configured with a value of `manual`, issue the `MODIFY procname,ACTIVATE` command for each workload currently being routed to the site specified by the *sitename* value. This will direct the load balancers that new connection requests should be routed to the alternate, standby site. See Chapter 7, "Lifeline Advisor commands," on page 115 for more information.

AQS0189W • AQS0190I

System programmer response: Ensure that the site specified by the *sitename* value is available. If the site is available, verify network connectivity between the Advisor and the site. Verify that Agents are running on each LPAR and Linux on System z Management Guest in the site where a server application for the workloads resides. Verify the connectivity used by BCPii to communicate with each of the site's images. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing z/OS network connectivity problems. See the documentation for your Linux distribution for more information about diagnosing network connectivity problems on Linux on System z. See *z/OS MVS Programming: Callable Services for High-Level Languages* for High-Level Languages for information about diagnosing BCPii issues.

Module: LMBCPII, LMLB

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically issue the `MODIFY procname,ACTIVATE` command to have workloads routed to the alternate site.

Example:

```
AQS0188W SITE SITE01 IS NOT AVAILABLE
```

AQS0189W PEER ADVISOR IS NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when both of the following conditions occur:

- The Advisor can no longer communicate over the network with the peer Advisor.
- The Advisor can no longer communicate with the image (LPAR) where the peer Advisor was active by using the Base Control Program Internal Interface (BCPii) service.

This message will be deleted when the peer Advisor reconnects to this Advisor or communication is reestablished with the image where the peer Advisor was active by using BCPii services.

System action: Processing continues. The peer Advisor site is no longer considered responsive.

Operator response: If this Advisor is the secondary Advisor and `advisor_takeover_policy` is configured with a value of `manual`, issue the `MODIFY procname,TAKEOVER` command to make this Advisor assume the role of primary Advisor; otherwise, contact the system programmer.

System programmer response: Ensure that the peer Advisor is available. Check for prior AQS0124I and AQS0125I messages to determine the system name where the peer Advisor was active. If the peer Advisor is available, verify network connectivity between the Advisor and its peer Advisor. Verify the connectivity used by BCPii to communicate with each of the site's images. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems. See *z/OS MVS Programming: Callable Services for High-Level Languages* for information about diagnosing BCPii issues.

Module: LMBCPII

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically issue the `MODIFY procname,TAKEOVER` command to have this Advisor assume the role of primary Advisor.

AQS0190I MULTIPLE CONNECTION ATTEMPTS FROM *ipaddress* REJECTED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when a load balancer repeatedly tries to connect. This usually indicates that the load balancer was configured incorrectly and that the Advisor does not permit any of the connection attempts.

In the message text:

ipaddress

The IP address of the load balancer.

System action: Processing continues. If the load balancer is still active, it might continue to attempt to reconnect to the Advisor.

Operator response: Locate the prior AQS0168I message that precedes this message and contact the system programmer. Message AQS0168I describes the reason that the load balancer connection was rejected.

System programmer response: See message AQS0168I for information about how to correct the problem.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0190I MULTIPLE CONNECTION ATTEMPTS FROM 10.10.10.1 REJECTED

| **AQS0191I ALL LB CONNECTION ATTEMPTS REJECTED - NO LB_ID_LIST CONFIGURED**

| **Explanation:** The Multi-site Workload Lifeline Advisor issues this message when the `lb_id_list` statement is not configured but load balancers try to connect.

| **System action:** Processing continues. If the load balancers are still active, they might continue to attempt to reconnect to the Advisor.

| **Operator response:** Contact the system programmer.

| **System programmer response:** Stop the load balancers from trying to connect to the Advisor.

| **Module:** LMLB

| **Routing code:** 8, 10

| **Descriptor code:** 12

| **Automation:** Not applicable for automation.

| **Example:**

| AQS0191I ALL LB CONNECTION ATTEMPTS REJECTED - NO LB_ID_LIST CONFIGURED

AQS0201I ALL MATCHING WORKLOADS *state* ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,ACTIVATE,WORKLOAD=*ALL_AS*,SITE=sitename` or a `MODIFY procname,ACTIVATE|DEACTIVATE|QUIESCE,WORKLOAD=*ALL_AQ*,SITE=sitename` command.

In the message text:

state

The state of the workloads. Possible values are:

- ACTIVATED
- DEACTIVATED
- QUIESCED

sitename

The sysplex name of the site.

| **System action:** Processing continues based on the *state* value:

- When the *state* value is ACTIVATED,
 - | – For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the matching workloads are to be routed to the site specified by the *sitename* value.
 - | – For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies z/OS Lifeline Agents on the specified site to allow new messages for the matching workloads to be routed to the site specified by the *sitename* value.
- When the *state* value is DEACTIVATED,

AQS0202I

- | – For workloads defined in the `cross_sysplex_list` Advisor configuration statement, all active connections to server applications that are defined for the matching workloads on the site specified by the `sitename` value are terminated.
- When the `state` value is `QUIESCED`,
 - | – For workloads defined in the `cross_sysplex_list` Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the matching workloads must not be routed to the site specified by the `sitename` value. Connection timeouts or resets might occur on the client applications that initiated the requests while these workloads are in a quiesced state.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0201I ALL MATCHING WORKLOADS QUIESCED ON SITE SITE01
```

AQS0202I ALL MATCHING WORKLOADS *state*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,DEACTIVATE|QUIESCE,WORKLOAD=*ALL_AS*` or a `MODIFY procname,ACTIVATE|DEACTIVATE|QUIESCE,WORKLOAD=*ALL_AQ*` command.

In the message text:

state

The state of the workloads. Possible values are:

- `ACTIVATED`
- `DEACTIVATED`
- `QUIESCED`

System action:

- When the `state` value is `ACTIVATED`,
 - | – For workloads defined in the `cross_sysplex_list` Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the matching workloads are to be routed to both sites
- When the `state` value is `DEACTIVATED`,
 - | – For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, all active connections to server applications that are defined for the matching workloads are terminated.
 - | – For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies z/OS Agents on the previously active site to verify that all MQ cluster receiver channels are stopped and to check for existing messages on each MQ cluster queue for each matching workload.
- When the `state` value is `QUIESCED`,
 - | – For workloads defined in the `cross_sysplex_list` or `sna_application_list` Advisor configuration statement, the Advisor notifies load balancers that new connection requests to server applications for the matching workloads must not be routed to any site. Connection timeouts or resets might occur on the client applications that initiated the requests while these workloads are in a quiesced state.
 - | – For workloads defined in the `mq_manager_list` Advisor configuration statement, the Advisor notifies z/OS Lifeline Agents on the currently active site to prevent new messages for all matching workloads to be routed to any site.

Operator response: None.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0202I ALL MATCHING WORKLOADS QUIESCED

AQS0203W LOAD BALANCER CONNECTION FROM *ipaddress* IS UNRESPONSIVE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when a load balancer is configured to request Server Application State Protocol (SASP) weights but the Advisor has not received any SASP message from this load balancer for more than twice the interval that is configured on the *failure_detection_interval* Advisor configuration statement. The message will be deleted when the Advisor receives a new SASP message from the load balancer.

In the message text:

ipaddress

The IP address of the load balancer.

System action: Processing continues. If the load balancer is still active, it might attempt to send an SASP message.

Operator response: Determine whether the load balancer is active. Restart the load balancer, if necessary. Otherwise, contact the system programmer.

System programmer response: Check the load balancer and correct any configuration or connectivity problems. Restart the load balancer, if necessary. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems.

Module: LMLB

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

Example:

AQS0203W LOAD BALANCER CONNECTION FROM 10.10.10.1 IS UNRESPONSIVE

AQS0204I *applname* DUMP INTERNAL LOG COMMAND SUCCESSFUL

Explanation: The Multi-site Workload Lifeline Advisor or Agent issues this message in response to a z/OS MODIFY *procname*, DUMPL0G command or Linux on System z `aqsl1cmd dumplog` command when the command completes successfully.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: The internal trace messages are dumped to the file specified by the AQSTRACE DD card in the Advisor or z/OS Agent started procedure, or in the `/opt/ibm/lifeline/log/agent.trace` file for the Linux on System z Management Agent.

Operator response: None.

System programmer response: None.

Module: LACMD, LMCMD

Routing code: 8, 10

AQS0205I • AQS0206I

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0204I LLADVSR DUMP INTERNAL LOG COMMAND SUCCESSFUL
```

AQS0205I *applname* DUMP INTERNAL LOG COMMAND NOT SUCCESSFUL

Explanation: The Multi-site Workload Lifeline Advisor or Agent issues this message in response to a z/OS MODIFY *procname*, DUMPLOG command or Linux on System z *aqsllcmd dumplog* command when the command does not complete successfully.

In the message text:

applname

The application name. Possible values are:

LLADVSR

For the Multi-site Workload Lifeline Advisor

LLAGENT

For the Multi-site Workload Lifeline Agent

System action: The internal trace messages are not successfully dumped.

Operator response: Contact the system programmer.

System programmer response: Examine the syslogd file for the reason the internal trace messages could not be successfully dumped. The syslogd identifier is *lladvrs* for the Advisor and *llagent* for the Agent. Verify that the file system where the internal trace messages are to be dumped has sufficient space. If necessary, on z/OS, you can take a console dump of the application to capture the internal trace messages. On Linux on System z, you can issue a kill -SIGQUIT command against the process ID of the application to capture the internal trace messages.

Module: LACMD, LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0205I LLADVSR DUMP INTERNAL LOG COMMAND NOT SUCCESSFUL
```

AQS0206I *peername* AT *ipaddress* IS AT A level LEVEL - SOME FUNCTIONS MIGHT NOT BE AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the Multi-site Workload Lifeline Agent or peer Advisor that connects to this Advisor is not at the same communication protocol version. This occurs when one of the following conditions is true:

- The peer Advisor or Agent is not at the same version level as this Advisor.
- The peer Advisor or Agent is at the same version level, but not at the same service level.

In the message text:

peername

The application name. Possible values are:

ADVISOR

For the Multi-site Workload Lifeline Advisor

AGENT

For the Multi-site Workload Lifeline Agent

ipaddress

The IP address of the peer Advisor or Agent

level

The indication of how the communication level is different. Possible values are:

LOWER

The peer Advisor or Agent is at a lower version or service level

HIGHER

The peer Advisor or Agent is at a higher version or service level

System action: Processing continues. The communication protocol version used will be the lower of the protocol versions between the Advisor and the other application. Any functions that are configurable at the higher communication protocol version are not available because the negotiated protocol version is at a lower level.

Operator response: Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages. Otherwise examine the syslogd files for the Advisor. The syslogd identifier is lladvr for the Advisor. A WARNING message is logged which indicates the communication protocol version supported by the Advisor and the negotiated communication protocol version. The negotiated version will be the lower of the protocol versions supported by the Advisor and application. Ensure that all Advisors and Agents are at the same version and service level. See "Migrating IBM Multi-site Workload Lifeline" on page 51 for information on migrating to a newer version.

Module: LMAGNT, LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0206I AGENT AT 10.10.10.10 IS AT A LOWER LEVEL - SOME FUNCTIONS MIGHT NOT BE AVAILABLE

AQS0207I MODIFY *cmd* COMMAND MIGHT NOT BE SUCCESSFUL - *reason*

| **Explanation:** The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY
| *procname*,ACTIVATE|DEACTIVATE|QUIESCE command.

In the message text:

cmd

The MODIFY command issued. Possible values are:

- ACTIVATE
- DEACTIVATE
- QUIESCE

reason

The reason that the command might not be successful. Possible values are:

NOT ALL SNA APPLICATIONS WERE ACTIVATED

A MODIFY *procname*,ACTIVATE command was issued for a SNA workload and not all SNA applications could be activated.

NOT ALL SNA APPLICATIONS WERE INACTIVATED

A MODIFY *procname*,DEACTIVATE command was issued for a SNA workload and not all SNA applications could be inactivated.

NOT ALL LINUX ON SYSTEM z GUESTS WERE LOGGED OFF

A MODIFY *procname*,DEACTIVATE command was issued for a Linux on System z workload and not all Linux on System z guests could be logged off.

NOT ALL AGENTS WERE CONNECTED

One of the following occurred:

- A MODIFY *procname*,DEACTIVATE command was issued and not all required Multi-site Workload Lifeline z/OS Agents and Multi-site Workload Lifeline Linux on System z Management Agents, if the workload contains Linux on System z server applications, are connected.

AQS0207I

- A MODIFY *procname*, QUIESCE command was issued and not all required z/OS Agents, if the workload is defined in the mq_manager_list Advisor configuration statement, are connected.

WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename*

A MODIFY *procname*, DEACTIVATE command was issued for a workload defined in the mq_manager_list Advisor configuration statement, but messages exist on at least one of the workload's MQ cluster queues hosted by an MQ queue manager running on the site specified by the *sitename* value.

WORKLOAD *workloadname* MQ MESSAGING NOT ENABLED ON ALL QUEUE MANAGERS

A MODIFY *procname*, ACTIVATE command was issued for a workload defined in the mq_manager_list Advisor configuration statement, but the MQ cluster receiver channel could not be started on at least one of the workload's MQ cluster queues hosted by an MQ queue manager running on the site specified on the ACTIVATE command.

WORKLOAD *workloadname* MQ MESSAGING NOT DISABLED ON ALL QUEUE MANAGERS

A MODIFY *procname*, QUIESCE command was issued for a workload defined in the mq_manager_list Advisor configuration statement, but the MQ cluster receiver channel could not be stopped on at least one of the workload's MQ cluster queues hosted by an MQ queue manager running on the site where the workload was previously active.

WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename*

The MODIFY *procname*, DEACTIVATE command was issued for the workload specified by the *workloadname* value, but the workload's MQ cluster receiver channel is not stopped on at least one of the MQ queue managers processing the workload on the alternate site specified by the *sitename* value.

System action: The MODIFY command is processed by the Advisor, but any processing on the Agent systems is skipped for Agents currently not connected to the Advisor.

Operator response: Take the appropriate action, based on the *reason* value:

NOT ALL SNA APPLICATIONS WERE ACTIVATED

Contact the system programmer.

NOT ALL SNA APPLICATIONS WERE INACTIVATED

Contact the system programmer.

NOT ALL LINUX ON SYSTEM z GUESTS WERE LOGGED OFF

Contact the system programmer.

NOT ALL AGENTS WERE CONNECTED

Issue a MODIFY *procname*, DISPLAY, ADVISOR command to determine which configured Agents are currently not connected. Contact the system programmer with the list of Agents that are not connected.

WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename*

Contact the system programmer.

WORKLOAD *workloadname* MQ MESSAGING NOT ENABLED ON ALL QUEUE MANAGERS

Contact the system programmer.

WORKLOAD *workloadname* MQ MESSAGING NOT DISABLED ON ALL QUEUE MANAGERS

Contact the system programmer.

WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename*

Contact the system programmer.

System programmer response: Take the appropriate action, based on the *reason* value:

NOT ALL SNA APPLICATIONS WERE ACTIVATED

See message AQS0212I for information about how to correct the problem.

NOT ALL SNA APPLICATIONS WERE INACTIVATED

See message AQS0211I for information about how to correct the problem.

NOT ALL LINUX ON SYSTEM z GUESTS WERE LOGGED OFF

See message AQS0210I for information about how to correct the problem.

NOT ALL AGENTS WERE CONNECTED

See message AQS0209W for information about how to correct the problem.

WORKLOAD *workloadname* HAS QUEUED MQ MESSAGES ON SITE *sitename*

For each z/OS Agent on the site specified by the *sitename* value:

- If internal trace is being used, dump and examine the internal trace messages.
- Otherwise, examine the syslogd file. The syslogd identifier is *lagent* for the Agent.

Check for error messages indicating that messages were found on an MQ cluster queue for the workload specified by the *workloadname* value. If MQ server applications are active and processing messages from these MQ cluster queues, wait until all queued messages are processed and reissue the MODIFY *procname*,DEACTIVATE command. Message AQS0155I will be issued in response to the command when messages are no longer found on any MQ cluster queue for the workload specified by the *workloadname* value.

WORKLOAD *workloadname* MQ MESSAGING NOT ENABLED ON ALL QUEUE MANAGERS

For each z/OS Agent on the site specified by the *sitename* value:

- If internal trace is being used, dump and examine the internal trace messages.
- Otherwise, examine the syslogd file. The syslogd identifier is *lagent* for the Advisor.

Check for error messages that indicate why starting the MQ cluster receiver channel failed, or error messages that indicate the MQ queue manager configuration does not match the Advisor's configuration for this MQ cluster. Correct the problem and reissue the Activate command. Message AQS0153I will be issued in response to the command when all problems have been corrected.

WORKLOAD *workloadname* MQ MESSAGING NOT DISABLED ON ALL QUEUE MANAGERS

For each z/OS Agent on the site specified by the *sitename* value:

- If internal trace is being used, dump and examine the internal trace messages.
- Otherwise, examine the syslogd file. The syslogd identifier is *lagent* for the Agent.

Check for error messages that indicate why stopping the MQ cluster receiver channel failed, or error messages that indicate the MQ queue manager configuration does not match the Advisor's configuration for this MQ cluster. Correct the problem and reissue the Quiesce command. Message AQS0155I will be issued in response to the command when all problems have been corrected.

WORKLOAD *workloadname* CHANNEL NOT STOPPED ON SITE *sitename*

For each z/OS Agent on the site specified by the *sitename* value:

- If internal trace is being used, dump and examine the internal trace messages.
- Otherwise, examine the Agent's syslogd file. The syslogd identifier is *lagent* for the Agent.

Check for error messages indicating that an MQ cluster's receiver channel has not stopped. Wait until the MQ cluster receiver channel stops and reissue the MODIFY *procname*,DEACTIVATE command. Message AQS0155I will be issued in response to the command when all MQ cluster receiver channels are stopped for the workload. If the MQ cluster receiver channel does not stop, verify that the MQ cluster receiver channel is correctly configured. If possible, manually stop the MQ cluster receiver channel using the STOP CHANNEL command. See *WebSphere MQ: Administering IBM MQ* for information about IBM MQ commands.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0207I MODIFY DEACTIVATE COMMAND MIGHT NOT BE SUCCESSFUL - NOT ALL AGENTS WERE CONNECTED

AQS0208W ACTIVATE AND QUIESCE COMMANDS QUEUED FOR UNREGISTERED LOAD BALANCERS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE|QUIESCE command. Processing of the command is delayed for unregistered first-tier load balancers. The message will be deleted when all first-tier load balancers register or all unregistered load balancers are removed from the Advisor's *lb_id_list* configuration statement and a MODIFY *procname*,REFRESH command is issued.

System action: Processing continues. The command will be processed for registered first-tier load balancers. All

AQS0209W • AQS0210I

queued commands will be processed for an unregistered first-tier load balancer when it registers. Automatic workload switch is disabled while there are unregistered load balancers.

Operator response: Issue the `MODIFY procname,DISPLAY,ADVISOR` command to determine which configured load balancers are currently not registered. Contact the system programmer with the load balancers that are not registered.

System programmer response: Verify the network connectivity between the Advisor and each unregistered load balancer. Check the load balancer and correct any configuration or connectivity problems. Restart the load balancer if necessary. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems. If these disconnected load balancers are no longer required, remove them from the Advisor's `lb_id_list` configuration statement and issue the `MODIFY procname,REFRESH` command.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0209W DEACTIVATE COMMANDS IGNORED FOR DISCONNECTED AGENTS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a `MODIFY procname,DEACTIVATE` command. Processing of the command is ignored for disconnected Multi-site Workload Lifeline Agents. The message will be deleted when all Agents connect to the Advisor or all disconnected Agents are removed from the Advisor's `agent_id_list` configuration statement and a `MODIFY procname,REFRESH` command is issued.

System action: Processing continues. On systems where the Agent is disconnected, no active workload connections will be reset. Automatic workload switch is disabled while there are disconnected Agents.

Operator response: Issue the `MODIFY procname,DISPLAY,ADVISOR` command to determine which Agents are not connected. Contact the system programmer with the list of Agents that are not connected.

System programmer response: Verify network connectivity between the Advisor and disconnected Agents. Verify that the Agents are active and, if necessary, restart the Agents. When this message is deleted, reissue the `MODIFY procname,DEACTIVATE` command. The workload can still be activated on the alternate site, but you should verify that all active connections for this workload on the target site are reset before you attempt to issue a `MODIFY procname,ACTIVATE` command for this workload to the alternate site. Verify network connectivity between the Advisor and Agents. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems. See the documentation for your Linux distribution for more information about diagnosing network connectivity problems on Linux on System z. If these disconnected Agents are no longer required, remove them from the Advisor's `agent_id_list` configuration statement and issue the `MODIFY procname,REFRESH` command.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0210I NOT ALL LINUX ON SYSTEM z GUESTS FOR WORKLOAD *workloadname* WERE LOGGED OFF

Explanation: The Multi-site Workload Lifeline Linux on System z Management Agent issues this message while a `MODIFY procname,DEACTIVATE` command is processed by the Multi-site Workload Lifeline Advisor and some of the Linux on System z guests were not logged off.

In the message text:

workloadname

The name of the workload that is specified on the `MODIFY` command.

System action: Processing continues. Not all of the Linux on System z guests were logged off for the workload specified by the *workloadname* value on the z/VM system where the Linux on System z Management Agent is running.

Operator response: Contact the system programmer.

System programmer response: Dump and examine the Agent internal trace messages for error messages. The trace

will provide more information about the specific guests that could not be logged off. Log off these guests manually by using the z/VM FORCE command from the Linux on System z Management Guest.

Module: LAADV

Routing code: N/A

Descriptor code: N/A

Automation: Not applicable for automation.

Example:

AQS0210I NOT ALL LINUX ON SYSTEM z GUESTS FOR WORKLOAD WORK03 WERE LOGGED OFF

AQS0211I SOME SNA APPLICATIONS FOR WORKLOAD *workloadname* MIGHT NOT BE INACTIVATED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message while a MODIFY *procname*,DEACTIVATE command is processed by the Multi-site Workload Lifeline Advisor and some of the SNA applications might not be inactivated.

In the message text:

workloadname

The name of the workload that is specified on the MODIFY command.

System action: Processing continues. Some of the SNA applications might not have been inactivated for the workload specified by the *workloadname* value.

Operator response: Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages for error messages. Otherwise, examine the syslogd file for error messages. The syslogd identifier is llagent for the Agent. The trace will provide more information about the specific SNA applications that could not be inactivated.

If the message was issued because the SNA application was manually inactivated prior to the MODIFY *procname*,DEACTIVATE command, no further action is needed for this SNA application. The SNA application will be activated by the z/OS Agent if the workload, specified by the *workloadname* value, is later activated on this site.

Verify that each SNA application for this workload is not active by issuing the DISPLAY NET, ID= command for each SNA application on the systems where the SNA applications reside. Manually inactivate any SNA applications that are not inactive by using the VARY NET, INACT, ID= command.

Module: LASNA

Routing code: 8,10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0211I SOME SNA APPLICATIONS FOR WORKLOAD WORK04 MIGHT NOT BE INACTIVATED

AQS0212I NOT ALL SNA APPLICATIONS FOR WORKLOAD *workloadname* WERE ACTIVATED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message while a MODIFY *procname*,ACTIVATE command is processed by the Multi-site Workload Lifeline Advisor and some of the SNA applications were not activated.

In the message text:

workloadname

The name of the workload that is specified on the MODIFY command.

System action: Processing continues. Not all of the SNA applications were activated for the workload specified by the *workloadname* value.

Operator response: Contact the system programmer.

System programmer response: If internal trace is being used, dump and examine the internal trace messages for

AQS0213W • AQS0214I

error messages. Otherwise, examine the syslogd file for error messages. The syslogd identifier is llagent for the z/OS Agent. The trace will provide more information about the specific SNA applications that could not be activated. Manually activate these SNA applications by using the VARY NET,ACT,ID= command.

Module: LASNA

Routing code: 8,10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0212I NOT ALL SNA APPLICATIONS FOR WORKLOAD WORK04 WERE ACTIVATED
```

AQS0213W ACTIVATE COMMANDS IGNORED FOR DISCONNECTED AGENTS

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE command for a SNA workload. Processing of the command is ignored for disconnected Multi-site Workload Lifeline z/OS Agents. The message will be deleted when all z/OS Agents connect to the Advisor or all disconnected z/OS Agents are removed from the Advisor's agent_id_list configuration statement and a MODIFY *procname*,REFRESH command is issued.

System action: Processing continues. Not all of the SNA applications were activated for the workload specified by the *workloadname* value. None of the SNA applications were activated on a system where a z/OS Agent is disconnected. Automatic workload switch is disabled while there are disconnected Agents.

Operator response: Issue the MODIFY *procname*,DISPLAY,ADVISOR command to determine which z/OS Agents are not connected. Contact the system programmer with the list of z/OS Agents that are not connected.

System programmer response: Verify that the z/OS Agents are active and, if necessary, restart the z/OS Agents. Manually activate the SNA applications that were not activated by using the VARY NET,ACT,ID= command. Verify network connectivity between the Advisor and z/OS Agents. See *z/OS Communications Server: IP Diagnosis Guide* for information about diagnosing network connectivity problems. If these disconnected z/OS Agents are no longer required, remove them from the Advisor's agent_id_list configuration statement and issue the MODIFY *procname*,REFRESH command.

Module: LMCMD

Routing code: 8,10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0214I APPLICATION *applname* HAS BEEN ACTIVATED

Explanation: The Multi-site Workload Lifeline z/OS Agent issues this message while a MODIFY *procname*,ACTIVATE command is processed by the Multi-site Workload Lifeline Advisor and the SNA application is activated.

In the message text:

applname

The name of the SNA application that was activated.

System action: Processing continues.

Operator response: None.

System programmer response: None.

Module: LASNA

Routing code: 8,10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically restart any SNA applications that were previously terminated when the SNA application was inactivated.

Example:

AQS0214I APPLICATION SNAAPPL1 HAS BEEN ACTIVATED

AQS0300I SECURE CONNECTION REQUEST RECEIVED FROM USER *userid* AT *ipaddress*

Explanation: The Multi-site Workload Lifeline Advisor received a Transport Layer Security (TLS) secure connection request originating at the IP address specified by the *ipaddress* value. The request was sent by a load balancer, Multi-site Workload Lifeline peer Advisor, or Multi-site Workload Lifeline Agent.

In the message text:

userid

The user ID that is associated with the client certificate of the load balancer, peer Advisor, or Agent that requests access to the Advisor.

ipaddress

The IP address of the load balancer, peer Advisor, or Agent that sent the request.

System action: Processing continues. SAF authorization for access to the Advisor will be checked for the user ID specified by the *userid* value.

Operator response: None.

System programmer response: None.

Module: LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0300I SECURE CONNECTION REQUEST RECEIVED FROM USER IBMUSER AT 10.10.10.1

AQS0301I *client* CONNECTION ATTEMPT FROM USER *userid* AT *ipaddress* FAILED - REASON CODE *reason*

Explanation: The Multi-site Workload Lifeline Advisor received a connection request originating at the IP address specified by the *ipaddress* value. The request was sent by a load balancer, Multi-site Workload Lifeline peer Advisor, or Multi-site Workload Lifeline Agent.

Application Transparent Transport Layer Security (AT-TLS) was enabled in the TCP/IP stack and no matching *lb_id_list* (for load balancer connections), *advisor_id_list* (for peer Advisor connections), or *agent_id_list* (for Agent connections) entry was configured in the Advisor profile. Authorization for the connection to the Advisor failed for this client.

In the message text:

client

The type of client that attempted to connect to the Advisor. Possible values are:

ADVISOR

The client is a peer Advisor.

AGENT

The client is an Agent.

LB The client is a load balancer.

userid

The user ID of the load balancer, peer Advisor, or Agent that is requesting access to the Advisor. If the user ID is not obtained from AT-TLS, the value is UNKNOWN.

ipaddress

The IP address of the load balancer, peer Advisor, or Agent.

reason

The code that explains the failure. Possible values are:

AQS0301I

- 1 A usable AT-TLS policy does not exist for this connection. For example, the policy agent is not active, or the AT-TLS policy for this connection specifies the wrong port.
- 2 The AT-TLS policy that is defined for this connection does not enable AT-TLS. In the policy, the TTLSEnvironmentAction statement is not configured with TTLSEnabled set to On.
- 3 The AT-TLS policy that is defined for this connection does not define the Advisor as a controlling application or require client authentication with SAF checking. In the policy for the Advisor, the TTLSEnvironmentAction statement is not configured correctly. The HandshakeRole must be set to ServerWithClientAuth. The TTLSEnvironmentAdvancedParms parameter must be configured with ApplicationControlled set to On and ClientAuthType set to SAFCheck.
- 4 The AT-TLS handshake failed for this connection.
- 5 System authorization facility (SAF) authorization failed for this connection. The SERVAUTH class profile AQS.LLADVSR.LBACCESS.*sysname* (for a load balancer connection), AQS.LLADVSR.ADVISORACCESS.*sysname* (for a peer Advisor connection), or AQS.LLADVSR.AGENTACCESS.*sysname* (for an Agent connection) exists but the userid that is associated with the client certificate is not permitted with READ access to this profile.
- 6 The Advisor was unable to obtain storage for processing an AT-TLS connection request.
- 7 The Advisor call to the SIOCTLCTL IOCTL failed unexpectedly.
- 8 System authorization facility (SAF) authorization failed for this connection. The SERVAUTH class profile AQS.LLADVSR.LBACCESS.*sysname* (for a load balancer connection), AQS.LLADVSR.ADVISORACCESS.*sysname* (for a peer Advisor connection), or AQS.LLADVSR.AGENTACCESS.*sysname* (for an Agent connection) is not defined.

System action: Processing continues. The Advisor does not allow connections from this client. The client that attempted to connect to the Advisor might continue to attempt to connect.

Operator response: If the client is a z/OS Agent or load balancer, ensure that IP address of the client is defined in the *lb_id_list*, or *agent_id_list* statement in the Advisor configuration file. If the client is a peer Advisor and you are not using AT-TLS for this connection, ensure that IP address of the client is defined in the *advisor_id_list*.

If you are using AT-TLS for this connection, take the appropriate action, based on the *reason* value:

- 1 Start the Policy Agent if it is not started. If the AT-TLS policy for the Advisor connections has changed, refresh the Policy Agent. If the problem is not corrected, take the following actions:
 - If internal trace is being used, dump and save the Advisor internal trace messages.
 - If internal trace is not being used, save the Advisor syslogd file.
 - Save the AT-TLS syslogd file.
 - Save the policy agent syslogd file.
 - Contact the system programmer.
- 6 If the storage problem cannot be corrected, dump and save the Advisor internal trace messages, if internal trace is being used, or save the Advisor syslogd file. If a dump was not created, take a dump of the Advisor address space, and then contact the system programmer.

For all other reasons, take the following actions:

- Save the system console.
- If internal trace is being used, dump and save the Advisor internal trace messages.
- If internal trace is not being used, save the Advisor syslogd file.
- Save the AT-TLS syslogd file.
- Save the policy agent syslogd file.
- Contact the system programmer.

System programmer response: Examine the Advisor internal trace file or syslogd file for errors. The syslogd identifier is *lladvsr* for the Advisor. Correct the *lb_id_list*, *advisor_id_list*, or *agent_id_list* statements in the configuration file as needed. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for information about configuring the Advisor and Agent.

If you are using AT-TLS for this connection, take the appropriate action, based on the *reason* value:

- 1 If the Policy Agent is active and has been refreshed since the last change to the AT-TLS policy, examine the system console or internal trace file, the Advisor syslogd file, the AT-TLS syslogd file, and the policy agent syslogd file for errors. Correct the AT-TLS policy for this connection. See *z/OS Communications Server IP Configuration Reference* for more information about the Policy Agent and policy statements for AT-TLS. Refresh the Policy Agent after changing the policy.
- 2 Change the AT-TLS policy for this connection in the TTLSEnvironmentAction statement to TTLS-enabled On. See *z/OS Communications Server IP Configuration Reference* for more information about the Policy Agent and policy statements for AT-TLS. Refresh the Policy Agent after changing the policy.
- 3 If the client is a Linux on System z Management Agent, ensure that no AT-TLS policy is configured for the IP address and port of this client. For other clients, change the AT-TLS policy for this connection in the TTLSEnvironmentAction statement to ensure the following:
 - The HandshakeRole is set to ServerWithClientAuth.
 - The TTLSEnvironmentAdvancedParms parameter is configured with ApplicationControlled set to On and ClientAuthType set to SAFCheck.

See *z/OS Communications Server IP Configuration Reference* for more information about the Policy Agent and policy statements for AT-TLS. Refresh the Policy Agent after changing the policy.

- 4 Correct the TLS handshake parameters in the AT-TLS policy for this connection. For example, ensure the following:
 - The HandshakeTimeout value for the Advisor policy is sufficient.
 - The HandshakeRole value for the Advisor is ServerWithClientAuth or Server.
 - The HandshakeRole value for the Agent and load balancers is Client.

See *z/OS Communications Server IP Configuration Reference* for more information about the Policy Agent and policy statements for AT-TLS. Refresh the Policy Agent after changing the policy.

- 5 Ensure that the user ID specified by the *userid* value has at least read access to the correct SERVAUTH class profile:
 - AQS.LLADVSR.LBACCESS.*sysname* for a load balancer connection
 - AQS.LLADVSR.ADVISORACCESS.*sysname* for a peer Advisor connection
 - AQS.LLADVSR.AGENTACCESS.*sysname* for an Agent connection

For more information, see *z/OS Security Server RACF Command Language Reference*.

- 6 If the storage problem cannot be corrected, contact IBM software support services with all supporting documentation. The application syslogd file or internal trace file is the minimum diagnostic data that should be provided.
- 7 Examine the system console, the Advisor syslogd file, the AT-TLS syslogd or internal trace file, and the policy agent syslogd file for errors. Ensure that the certificate is correct. For more information, see *z/OS Security Server RACF Command Language Reference*. If the problem is not corrected, contact IBM software support services with all supporting documentation.
- 8 Define necessary SAF profiles on each system where the Advisor can run. Ensure that the user ID specified by the *userid* value has at least read access to the correct SERVAUTH class profile:
 - AQS.LLADVSR.LBACCESS.*sysname* for a load balancer connection
 - AQS.LLADVSR.ADVISORACCESS.*sysname* for a peer Advisor connection
 - AQS.LLADVSR.AGENTACCESS.*sysname* for an Agent connection

For more information, see *z/OS Security Server RACF Command Language Reference*.

Module: LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0302I • AQS0303I

AQS0301I AGENT CONNECTION ATTEMPT FROM USER UNKNOWN AT 10.10.10.1 FAILED - REASON CODE 1

AQS0302I WORKLOAD *workloadname1* MUST BE ACTIVE BEFORE ACTIVATING WORKLOAD *workloadname2*

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE,WORKLOAD=*workloadname2* command when *workloadname1* is not active. The Advisor configuration file contains a query_workload_list definition that associates the workload specified by the *workloadname2* value with the workload specified by the *workloadname1* value.

In the message text:

workloadname1

The name of the Active/Standby workload associated with the workload specified on the MODIFY command.

workloadname2

The name of the Active/Query workload specified on the MODIFY command.

System action: The MODIFY command is ignored.

Operator response: Issue a MODIFY *procname*,ACTIVATE,WORKLOAD=*workloadname1*,SITE= command to activate the associated workload. Reissue the previous MODIFY command to activate the workload name specified by the *workloadname2* value. See Chapter 7, "Lifeline Advisor commands," on page 115 for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0302I WORKLOAD WORK01 MUST BE ACTIVE BEFORE ACTIVATING WORKLOAD WORK02

AQS0303I WORKLOAD *workloadname* MUST BE QUIESCED ON SITE *sitename* BEFORE ISSUING A DEACTIVATE COMMAND

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DEACTIVATE,WORKLOAD=*workloadname* command when the workload specified by the *workloadname* value has not been quiesced on the site specified by the *sitename* value.

In the message text:

workloadname

The name of the workload specified on the MODIFY command.

sitename

The sysplex name of the site where the workload is active.

System action: The MODIFY command is ignored.

Operator response: Issue a MODIFY *procname*,QUIESCE command to quiesce the workload specified by the *workloadname* value on the target site. Reissue the previous MODIFY command to deactivate the workload. See Chapter 7, "Lifeline Advisor commands," on page 115 for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0303I WORKLOAD WORK02 MUST BE QUIESCED ON SITE SITE01 BEFORE ISSUING A DEACTIVATE COMMAND

AQS0304I *cmd* **COMMAND FOR WORKLOAD** *workloadname* **MUST NOT SPECIFY A SITE**

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DEACTIVATE|QUIESCE,WORKLOAD=*workloadname*,SITE= *command* for an Active/Standby workload. Workloads that are defined in the Advisor configuration file as Active/Standby workloads cannot be deactivated or quiesced specifying the SITE keyword. The SITE keyword is valid on this MODIFY command only for workloads that are defined in the Advisor Configuration file as Active/Query workloads.

In the message text:

cmd

The MODIFY command that was issued. Possible values are:

- DEACTIVATE
- QUIESCE

workloadname

The name of the workload specified on the MODIFY command.

System action: The MODIFY command is ignored.

Operator response: Reissue the MODIFY command without specifying the SITE keyword. See the MODIFY command section for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0304I QUIESCE COMMAND FOR WORKLOAD WORK01 MUST NOT SPECIFY A SITE

AQS0305I **ACTIVATE COMMAND FOR WORKLOAD** *workloadname* **MUST SPECIFY A SITE**

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE,WORKLOAD=*workloadname* command for an Active/Standby workload. The workload specified by the *workloadname* value is defined in the Advisor configuration file as an Active/Standby workload and cannot be activated without specifying the SITE keyword. Omitting the SITE keyword is valid on this MODIFY command only for workloads that are defined in the Advisor Configuration file as Active/Query workloads.

In the message text:

workloadname

The name of the workload specified on the MODIFY command.

System action: The command is ignored.

Operator response: Reissue the MODIFY command specifying the SITE keyword. See the MODIFY command section for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0305I ACTIVATE COMMAND FOR WORKLOAD WORK01 MUST SPECIFY A SITE

AQS0306I ACTIVATE COMMAND FOR MATCHING WORKLOADS MUST SPECIFY A SITE

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,ACTIVATE,WORKLOAD=*ALL_AS* command. All workloads that are defined in the Advisor configuration file as Active/Standby workloads cannot be activated without specifying the SITE keyword. Omitting the SITE keyword is valid on this MODIFY command only for workloads that are defined in the Advisor Configuration file as Active/Query workloads.

System action: The command is ignored.

Operator response: Reissue the MODIFY command specifying the SITE keyword. See the MODIFY command section for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

AQS0307I *cmd* COMMAND FOR MATCHING WORKLOADS MUST NOT SPECIFY A SITE

Explanation: The Multi-site Workload Lifeline Advisor issues this message in response to a MODIFY *procname*,DEACTIVATE|QUIESCE,WORKLOAD=*ALL_AS*,SITE= command. All workloads that are defined in the Advisor configuration file as Active/Standby workloads cannot be deactivated or quiesced specifying the SITE keyword. The SITE keyword is valid on this MODIFY command only for workloads that are defined in the Advisor Configuration file as Active/Query workloads.

In the message text:

cmd

The MODIFY command that was issued. Possible values are:

- DEACTIVATE
- QUIESCE

System action: The MODIFY command is ignored.

Operator response: Reissue the MODIFY command without specifying the SITE keyword. See the MODIFY command section for more information.

System programmer response: None.

Module: LMCMD

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0307I QUIESCE COMMAND FOR MATCHING WORKLOADS MUST NOT SPECIFY A SITE

AQS0308W WORKLOAD *workloadname* ACCESSING STALE DATA ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- Static routing is configured for the workload, and the workload_switch_policy Advisor configuration statement is defined as manual.
- The Advisor receives an event over its Network Management Interface which indicates that the database replication latency state for the workload is acute or unknown.

The message will be deleted when the database replication for the workload is no longer in the acute or unknown state or the workload is quiesced on the site specified by the *sitename* value.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where the database is not current.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Determine why database replication is in the acute or unknown state, and correct, if possible. If the workload specified by the *workloadname* value is active on the alternate site, quiesce the workload on the site specified by the *sitename* value, using the `MODIFY procname, QUIESCE, WORKLOAD=workloadname, SITE=` command, to ensure that the workload is routed only to the site where the database is current.

Module: LMADV

Routing code: 1, 8

Descriptor code: 2, 7

Automation: This message is a possible candidate for automation if you want to automatically issue the `MODIFY procname, QUIESCE` command to quiesce routing to the site where the database is not current.

Example:

```
AQS0308W WORKLOAD WORK02 ACCESSING STALE DATA ON SITE SITE01
```

AQS0309I WORKLOAD *workloadname* DATA ACCESS ON SITE *sitename* NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- Dynamic routing is configured for the workload or static routing is configured for the workload and the `workload_switch_policy` Advisor configuration statement is defined as automatic.
- The Advisor receives an event over its Network Management Interface which indicates that the database replication latency state for the workload is acute or unknown.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where the database is not current.

System action: For workloads defined in the `cross_sysplex_list` Advisor configuration statement, the Advisor notifies load balancers to route connection requests to the alternate site and deactivates the workload on the site specified by the *sitename* value to prevent accessing stale data.

Operator response: Contact the system programmer.

System programmer response: Determine why database replication is in the acute or unknown state, and correct, if possible.

Module: LMADV

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0309I WORKLOAD WORK02 DATA ACCESS ON SITE SITE01 NOT AVAILABLE
```

AQS0310I WORKLOAD *workloadname* DATA ACCESS ON SITE *sitename* HAS RECOVERED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- Dynamic routing is configured for the workload or static routing is configured for the workload and the *workload_switch_policy* Advisor configuration statement is defined as automatic.
- Message AQS0309I was previously issued for the workload.
- The Advisor receives an event over its Network Management Interface which indicates that the database replication latency state for the workload is no longer acute or unknown.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where the database is now current.

System action: The Advisor re-activates the workload on the site specified by the *sitename* value and notifies load balancers to route new connection requests to both sites.

Operator response: None.

System programmer response: None.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0310I WORKLOAD WORK02 DATA ACCESS ON SITE SITE01 HAS RECOVERED
```

AQS0311W PERSISTENT CONNECTION FUNCTION IS DISABLED - BACK LEVEL AGENT FOUND

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the timeout parameter was enabled on the *query_workload_list* Advisor configuration statement, and a Lifeline Agent from an earlier release of Multi-site Workload Lifeline connects to the Advisor. See Chapter 6, "Multi-site Workload Lifeline configuration statements," on page 85 for more information about configuration statements and parameters. The message will be deleted when the Agent, that is not at the current release, is stopped and the connection is closed with the Advisor.

System action: The persistent connection function is disabled.

Operator response: Contact the system programmer.

System programmer response: To determine which Agent is not at the current release:

- If internal trace is being used, dump and examine the internal trace messages.
- If internal trace is not being used, examine the Advisor syslogd file. The syslogd identifier is lladvsr for the Advisor.

Upgrade the Agent to the current release and restart it.

Module: LMAGNT, LMCMD

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

AQS0312W ROUTING CONFIGURATION FOR WORKLOAD *workloadname* IGNORED - NOT ACTIVE ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is configured as an Active/Query workload in the Advisor configuration file.
- Static routing is configured for the workload, and the *workload_switch_policy* Advisor configuration statement is defined as manual.
- The workload is activated to only a single site.

The message will be deleted when the workload is activated to the alternate site or the workload is quiesced.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where the workload is not activated.

| **System action:** For workloads defined in the *cross_sysplex_list* Advisor configuration statement, the Advisor notifies load balancers to route new connection requests for the workload only to the alternate site.

Operator response: Issue the `MODIFY procname,ACTIVATE,WORKLOAD=workloadname,SITE=sitename` command to ensure the workload is activated to both sites.

System programmer response: None.

Module: LMADV

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

Example:

AQS0312W ROUTING CONFIGURATION FOR WORKLOAD WORK02 IGNORED - NOT ACTIVE ON SITE SITE01

AQS0313I ROUTING CONFIGURATION FOR WORKLOAD *workloadname* IGNORED - NOT ACTIVE ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is configured as an Active/Query workload in the Advisor configuration file.
- Static routing is configured for the workload, and the *workload_switch_policy* Advisor configuration statement is defined as automatic.
- The workload is activated to only a single site.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where the workload is not activated.

| **System action:** For workloads defined in the *cross_sysplex_list* Advisor configuration statement, the Advisor notifies load balancers to route new connections for the workload only to the alternate site.

Operator response: Issue the `MODIFY procname,ACTIVATE,WORKLOAD=workloadname,SITE=sitename` command to ensure the workload is activated to both sites.

System programmer response: None.

Module: LMADV

AQS0314I • AQS0315W

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0313I ROUTING CONFIGURATION FOR WORKLOAD WORK02 IGNORED - NOT ACTIVE ON SITE SITE01

AQS0314I ROUTING CONFIGURATION FOR WORKLOAD *workloadname* RESTORED

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is configured as an Active/Query workload in the Advisor configuration file.
- Static routing is configured for the workload, and the *workload_switch_policy* Advisor configuration statement is defined as automatic.
- Message AQS0313I was previously issued for the workload.
- The workload is now activated to both sites.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

| **System action:** For workloads defined in the *cross_sysplex_list* Advisor configuration statement, the Advisor notifies
| load balancers to route new connection requests for the workload to both sites.

Operator response: None.

System programmer response: None.

Module: LMMAIN

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

AQS0314I ROUTING CONFIGURATION FOR WORKLOAD WORK02 RESTORED

AQS0315W WORKLOAD *workloadname* ATTEMPTING TO ACCESS UNAVAILABLE SERVERS ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- Static routing is configured for the workload, and the *workload_switch_policy* Advisor configuration statement is defined as manual.
- The workload is active on the site specified by the *sitename* value.
- | • No server applications are available on the site specified by the *sitename* value for the single load balancer application group defined for the workload. A server application is considered unavailable if it is not active, it is active but not fully operational, or the system where it resides is resource constrained.

The message will be deleted when at least one server application in the load balancer application group is active on the site or the workload is quiesced on the site specified by the *sitename* value.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where no server applications are available for the workload.

System action: The Advisor continues notifying load balancers to route new connection requests for the workload to both sites.

Operator response: Contact the system programmer.

System programmer response: Ensure that the server applications for the workload are available on the site specified by the *sitename* value. Verify that each server application is active and is fully operational, and that the system where it resides has sufficient resources. If the workload specified by the *workloadname* value is active on the alternate site, quiesce the workload on the site specified by the *sitename* value, by using the `MODIFY procname, QUIESCE, WORKLOAD=workloadname, SITE= command`. This ensures the workload is routed only to the site where the server applications are active. See Chapter 7, "Lifecycle Advisor commands," on page 115 for more information.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: This message is a possible candidate for automation if you want to automatically issue the `MODIFY procname, QUIESCE` command to quiesce routing to the site where no servers for the workload are available.

Example:

```
AQS0315W WORKLOAD WORK02 ATTEMPTING TO ACCESS UNAVAILABLE SERVERS ON SITE SITE01
```

AQS0316I WORKLOAD *workloadname* HAS UNAVAILABLE SERVERS ON SITE *sitename*

Explanation: The Multi-site Workload Lifecycle Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- The workload is not active on the site specified by the *sitename* value, or dynamic routing is configured for the workload, or static routing is configured for the workload and the `workload_switch_policy` Advisor configuration statement is defined as automatic.
- No servers applications are available on the site specified by the *sitename* value for the single load balancer application group defined for the workload. A server application is considered unavailable if it is not active, it is active but not fully operational, or the system where it resides is resource constrained.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where no server applications are available for the workload.

System action: For workloads defined in the `cross_sysplex_list` Advisor configuration statement, the Advisor notifies load balancers to route new connection requests for the workload only to the alternate site.

Operator response: Contact the system programmer.

System programmer response: Ensure that the server applications for the workload are available on the site specified by the *sitename* value. Verify that each server application is active and is fully operational, and that the system where it resides has sufficient resources.

Module: LMLB

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

Example:

```
AQS0316I WORKLOAD WORK02 HAS UNAVAILABLE SERVERS ON SITE SITE01
```

AQS0317I WORKLOAD *workloadname* HAS RECOVERED ON SITE *sitename*

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- Dynamic routing is configured for the workload or static routing is configured for the workload and the *workload_switch_policy* Advisor configuration statement is defined as automatic.
- Message AQS0316I was previously issued for the workload.
- At least one server application is available on the site specified by the *sitename* value for the single load balancer application group defined for the workload.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

sitename

The sysplex name of the site where a server application is now available.

System action: For workloads defined in the *cross_sysplex_list* Advisor configuration statement, the Advisor notifies load balancers to route new connection requests for the workload to both sites.

Operator response: None.

System programmer response: None.

Module: LMLB

Routing code: 1, 8

Descriptor code: 2, 7

Automation: Not applicable for automation.

Example:

```
AQS0317I WORKLOAD WORK02 HAS RECOVERED ON SITE SITE01
```

AQS0318W WORKLOAD *workloadname* IS NOT AVAILABLE

Explanation: The Multi-site Workload Lifeline Advisor issues this message when the following conditions are true:

- The workload specified by the *workloadname* value is defined as an Active/Query workload in the Advisor configuration file.
- The workload is unavailable on both sites. A workload is unavailable on a site when one of the following conditions is true:
 - No server applications are available on that site for the single load balancer application group defined for the workload. A Linux on System z workload is unavailable if either no Linux on System z server applications are available for the load balancer application group defined for the workload or no z/OS server applications that these Linux on System z applications distribute to are available on that site. A server application is considered unavailable when one of the following conditions is true:
 - The server application is not active.
 - The server application is active but not fully operational.
 - The system where the server application resides is resource constrained.
 - The workload is not activated on the site.
 - The Advisor receives an event over its Network Management Interface which indicates that the database replication latency state for the workload is acute or unknown on that site.

The message will be deleted when at least one server application in each load balancer application group is active on one of the sites. A Linux on System z workload becomes available when at least one Linux on System z server application and one of its z/OS server applications is available on that site.

In the message text:

workloadname

The name of the workload that is defined in the Advisor configuration file.

System action: Processing continues.

Operator response: Contact the system programmer.

System programmer response: Issue `MODIFY procname,DISPLAY,WORKLOAD,DETAIL` to determine whether the workload failure on a site is caused by server application unavailability, the workload not being activated on the site, or replication latency state.

- If the server application is unavailable, verify that each server application is active and is fully operational, and that the system where it resides has sufficient resources.
- If the workload is not activated on the site, issue the `MODIFY procname,ACTIVATE,WORKLOAD=` command to activate the workload on the site.
- If the replication latency state is causing the failure, determine why database replication is in the acute or unknown state, and, if possible, resolve the cause of the replication latency state.

Module: LMLB

Routing code: 8, 10

Descriptor code: 12

Automation: Not applicable for automation.

Example:

```
AQS0318W WORKLOAD WORK02 IS NOT AVAILABLE
```

Appendix A. A sample started procedure for Lifeline Advisors

You can use the sample started procedure to understand how to set up the environment for the Lifeline Advisor.

```
//AQSADV PROC ROLE=S
//*
//*****
//* Function: Sample procedure for running the Multi-site
//* Workload Lifeline Advisor
//*
//* Parameters: [ROLE={S | P}]
//* S ==> Secondary Advisor (Default Role)
//* P ==> Primary Advisor
//*****
//AQSADV EXEC PGM=AQSLBADV,REGION=0K,TIME=NOLIMIT,
// PARM='/ -R &ROLE'
//*****
//* The STEPLIB card below points to the PDS/E where the Lifeline
//* load modules are kept.
//*****
//STEPLIB DD DSN=USER.LOADLIB,DISP=SHR
//*****
//* The CONFIG card below points to the member where the Lifeline
//* Advisor configuration file is located
//*****
//CONFIG DD DSN=USER.PROFILES(ADVCONF),DISP=SHR
//STDENV DD DUMMY
//SYSPRINT DD SYSOUT=*,DCB=(RECFM=F,LRECL=80,BLKSIZE=80)
//SYSIN DD DUMMY
//SYSERR DD SYSOUT=*
//SYSOUT DD SYSOUT=*,DCB=(RECFM=F,LRECL=80,BLKSIZE=80)
//CEEDUMP DD SYSOUT=*,DCB=(RECFM=FB,LRECL=132,BLKSIZE=132)
//SYSDUMP DD SYSOUT=*,DCB=(RECFM=FB,LRECL=132,BLKSIZE=132)
//*****
//* The SYSTCPD card below points to the member where the TCPIP.DATA
//* file is located
//*****
//SYSTCPD DD DSN=USER.PROFILES(TCPDATA),DISP=SHR
//*****
//* The AQSTRACE card below points to the UNIX System Services
//* file where the Advisor dumps internal trace messages.
//*****
//AQSTRACE DD PATH='/tmp/advisor.trace',PATHOPTS=(OWRONLY,OCREAT),
// PATHMODE=(SIRUSR,SIWUSR,SIRGRP,SIROTH)
```

Figure 9. Sample started procedure AQSADV

Appendix B. A sample started procedure for Lifeline z/OS Agents

You can use the sample started procedure to understand how to set up the environment for the Lifeline Agent.

```
//AQSAGE      PROC
//*
//*****
//*
//* Function: Sample procedure for running the Multi-site
//* Workload Lifeline Agent
//*
//*****
//AQSAGE EXEC PGM=AQSLBAGE,REGION=0K,TIME=NOLIMIT,
//  PARM='POSIX(ON) ALL31(ON)'/
//*****
//* The STEPLIB card below points to the PDS/E where the Lifeline
//* load modules are kept.
//*****
//STEPLIB DD DSN=USER.LOADLIB,DISP=SHR
//*****
//* The CONFIG card below points to the member where the Lifeline
//* Agent configuration file is located
//*****
//CONFIG DD DSN=USER.PROFILES(AGECONF),DISP=SHR
//STDENV DD DUMMY
//SYSPRINT DD SYSOUT=*,DCB=(RECFM=F,LRECL=80,BLKSIZE=80)
//SYSIN DD DUMMY
//SYSERR DD SYSOUT=A
//SYSOUT DD SYSOUT=A
//CEEDUMP DD SYSOUT=A
//SYSDUMP DD SYSOUT=A
//*****
//* The SYSTCPD card below points to the member where the TCP/IP.DATA
//* file is located
//*****
//SYSTCPD DD DSN=USER.PROFILES(TCPDATA),DISP=SHR
//*****
//* The AQSTRACE card below points to the UNIX System Services
//* file where the Agent dumps internal trace messages.
//*****
//AQSTRACE DD PATH='/tmp/agent.trace',PATHOPTS=(OWRONLY,OCREAT),
//          PATHMODE=(SIRUSR,SIWUSR,SIRGRP,SIROTH)
```

Figure 10. Sample started procedure AQSAGE

Appendix C. Sample configuration statements

You can use the following examples to understand how to configure the Lifeline Advisor and the Lifeline Agent for different workload environments.

Multi-site Workload Lifeline configuration Example 1

The following figure shows an example of the Multi-site Workload Lifeline configuration where an Active/Standby workload uses the z/OS sysplex distributor to route requests within a site. It also shows how to configure the Lifeline Advisors and Lifeline Agents to work in such an environment.




Figure 11. Multi-site Workload Lifeline configuration Example 1

To configure the previous sample environment shown in Figure 11, complete the following tasks:

1. Configure and start the Lifeline Advisor on the primary controller in the primary role. The primary Advisor listens on dynamic VIPA (DVIPA) address 10.10.10.1, which is defined in the VIPARANGE statement.
2. Configure and start the Lifeline Advisor on the secondary controller, in the secondary role. The secondary Advisor uses static VIPA address 192.10.20.1 to communicate with the primary Advisor.
3. Define distributable DVIPA 10.10.1.1 on system `sys_a` in sysplex Site1, and allow both `sys_a` and `sys_b` to be targets of this DVIPA.
4. Define distributable DVIPA 10.10.2.1 on system `sys_c` in sysplex Site2, and allow both `sys_c` and `sys_d` to be targets of this DVIPA.

5. Configure and start Lifeline Agents on each of the images in Site 1 and Site 2, by using the static VIPA that is defined in that image (192.10.1.1 on sys_a, 192.10.1.2 on sys_b, 192.10.2.1 on sys_c, and 192_10.2.2 on sys_d).

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 1. Use this configuration file for the Advisors on both the primary and secondary controllers.

```

#####
#
# This file contains sample configuration statements for the Lifeline
# Advisor - using default settings in most cases.
#
#####

# Following specifies the port which the primary Advisor will listen
# on for connections from Agents.
agent_connection_port      8181

# Following specifies which Agents are allowed to connect to the
# primary Advisor.
agent_id_list
{
    192.10.1.1..4000      # IP address and port for sys_a
    192.10.1.2..4000      # IP address and port for sys_b
    192.10.2.1..4000      # IP address and port for sys_c
    192.10.2.2..4000      # IP address and port for sys_d
}

# Following specifies the port which the primary Advisor will listen
# on for connections from a secondary Advisor.
advisor_connection_port    8000

# Following specifies which secondary Advisors are allowed to connect
# to the primary Advisor (should contain at least the IP address of this
# host).
advisor_id_list
{
    192.10.10.1          # IP address for Advisor on primary controller
    192.10.20.1          # IP address for Advisor on secondary controller
}

# Following specifies the IPv4 address and port the primary Advisor
# listens on for connections from load balancers. The IP address in this
# statement is also used by the primary Advisor to accept connections
# from Agents or the secondary Advisor.
lb_connection_v4          10.10.10.1..3860

# Following specifies which load balancers are allowed to connect to
# the primary Advisor.
lb_id_list
{
    192.100.1.1          # IP address of LB in Site1
    192.100.2.1          # IP address of LB in Site2
}

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each z/OS sysplex distributor
# used to route the workload, the port number of the server
# application used for the workload, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
cross_sysplex_list
{
    10.10.1.1..80,SITE1,WORKLOAD1      # HTTP Server on Site1
    10.10.2.1..80,SITE2,WORKLOAD1      # HTTP Server on Site2
}

```

Figure 12. Sample configuration statements for the Lifeline Advisor

The Lifeline Advisor accesses files within the UNIX System Services shell, so the Lifeline Advisor needs to have an OMVS segment defined for it. Assuming that the Lifeline Advisor is started using a procedure called AQSADV, use the following

JCL snippet to define the OMVS segment:

```
//EZARACF JOB MSGLEVEL=(1,1),MSGCLASS=A,CLASS=A
//DAEMONS EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
    ADDUSER AQSADV DFLTGRP(OMVSGRP) OMVS(UID(0) HOME('/'))
    RDEFINE STARTED AQSADV.* STDATA(USER(AQSADV))
    SETROPTS RACLIST(STARTED) REFRESH
//*
```

Based on this example, you need to update the TCPIP profile on the two stand-alone images to contain the following information:

```
VIPADYNAMIC
VIPARANGE DEFINE 255.255.255.255 10.10.10.1
ENDVIPADYNAMIC
```

Sample configuration statements for the Lifeline Agent

You can use the following sample configuration file to configure the Lifeline Agent for Example 1.

```
#####
#
# This file contains sample configuration statements for the Lifeline
# Agent - most defaults are selected.
#
#####

# Following specifies the IP address and port of the primary Advisor
# that this Agent connects to.
advisor_id          10.10.10.1..8181

# Specifies the local IP address and port this Agent binds to before
# connecting to the primary Advisor. Uncomment the statement that
# applies to the image that the Agent is running on.
#host_connection    192.10.1.1..4000 # sys_a
#host_connection    192.10.1.2..4000 # sys_b
#host_connection    192.10.2.1..4000 # sys_c
#host_connection    192.10.2.2..4000 # sys_d
```

Figure 13. Sample configuration statements for the Lifeline Agent

The Lifeline Agent accesses files within the UNIX System Services shell, so you need to define an OMVS segment for the Lifeline Agent. Assuming that the Lifeline Agent is started using a procedure called AQSAGE, you can use the following JCL snippet to define the OMVS segment:

```
//EZARACF JOB MSGLEVEL=(1,1),MSGCLASS=A,CLASS=A
//DAEMONS EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
    ADDUSER AQSAGE DFLTGRP(OMVSGRP) OMVS(UID(0) HOME('/'))
    RDEFINE STARTED AQSAGE.* STDATA(USER(AQSAGE))
    SETROPTS RACLIST(STARTED) REFRESH
//*
```

According to previous samples, you need to update the TCPIP profile on the image owning the distributable dynamic VIPA in the following way:

```

VIPADYNAMIC
; Uncomment statement that applies to the image where the z/OS sysplex distributor
; is being started
;
; Site1 definitions
; VIPADEFINE 255.255.255.255 10.10.1.1
; VIPADIST DISTMETHOD SERVERWLM 10.10.1.1 PORT 80 DESTIP ALL
;
; Site2 definitions
; VIPADEFINE 255.255.255.255 10.10.2.1
; VIPADIST DISTMETHOD SERVERWLM 10.10.2.1 PORT 80 DESTIP ALL
ENDVIPADYNAMIC

```

Multi-site Workload Lifeline configuration Example 2

The following figure shows an example of the Multi-site Workload Lifeline configuration where an Active/Standby workload uses a customized load balancer to route requests within a site. It also shows how to configure the Lifeline Advisors and Lifeline Agents to work in such an environment.




Figure 14. Multi-site Workload Lifeline configuration Example 2

To configure the previous sample environment shown in Figure 14, first repeat the initial five tasks from Example 1.

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 2. Use this configuration file for the Advisors on both the primary and secondary controllers.

```

=====
#
# This file contains sample configuration statements for the Lifeline
# Advisor - using default settings in most cases.
#
=====

# Following specifies the port which the primary Advisor will listen
# on for connections from Agents.
agent_connection_port      8181

# Following specifies which Agents are allowed to connect to the
# primary Advisor.
agent_id_list
{
    192.10.1.1..4000      # IP address and port for sys_a
    192.10.1.2..4000      # IP address and port for sys_b
    192.10.2.1..4000      # IP address and port for sys_c
    192.10.2.2..4000      # IP address and port for sys_d
}

# Following specifies the port which the primary Advisor will listen
# on for connections from a secondary Advisor.
advisor_connection_port    8000

# Following specifies which secondary Advisors are allowed to connect
# to the primary Advisor (should contain at least the IP address of this
# host).
advisor_id_list
{
    192.10.10.1          # IP address for Advisor on primary controller
    192.10.20.1          # IP address for Advisor on secondary controller
}

# Following specifies the IPv4 address and port the primary Advisor
# listens on for connections from load balancers. The IP address in this
# statement is also used by the primary Advisor to accept connection
# from Agents or the secondary Advisor.
lb_connection_v4          10.10.10.1..3860

# Following specifies which load balancers are allowed to connect to
# the primary Advisor.
lb_id_list
{
    192.100.1.1          # IP address of LB in Site1
    192.100.2.1          # IP address of LB in Site2
}

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each z/OS sysplex distributor
# used to route the workload, the port number of the server
# application used for the workload, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
cross_sysplex_list
{
    10.10.1.1..80,SITE1,WORKLOAD1      # HTTP Server on Site1
    10.10.2.1..80,SITE2,WORKLOAD1      # HTTP Server on Site2
}

```

Figure 15. Sample configuration statements for the Lifeline Advisor, part 1 of 2

```

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each intermediary node used
# to route the workload, the port number to reach the
# intermediary node, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
intermediary_node_list
{
  192.10.100.1..8000,SITE1,WORKLOAD1 # Intermediary node on Site1
  192.10.200.1..8000,SITE2,WORKLOAD1 # Intermediary node on Site2
}

```

Figure 16. Sample configuration statements for the Lifeline Advisor, part 2 of 2

Follow the same steps in Example 1 to define an OMVS segment for the Advisor and the VIPARANGE DVIPA on both stand-alone LPARs.

Sample configuration statements for the Lifeline Agent

Use the same Lifeline Agent configuration file from Example 1, and follow the same steps in Example 1 to define an OMVS segment for the Agent and the z/OS sysplex distributor in both sites.

Multi-site Workload Lifeline configuration Example 3

The following figure shows an example of a configuration of Multi-site Workload Lifeline where two workloads are defined, one as an Active/Standby workload and the other as an Active/Query workload. Both workloads use the z/OS sysplex distributor to route requests within a site. The following figure describes how to configure the Lifeline Advisors and Lifeline Agents to work in such an environment.




Figure 17. Multi-site Workload Lifeline configuration Example 3

To configure the sample environment shown in Figure 17 on page 263, complete the following tasks:

1. Configure and start the Lifeline Advisor on the primary controller in the primary role. The primary Advisor listens on dynamic VIPA (DVIPA) address 10.10.10.1, which is defined in the VIPARANGE statement.
2. Configure and start the Lifeline Advisor on the secondary controller, in the secondary role. The secondary Advisor uses static VIPA address 192.10.20.1 to communicate with the primary Advisor.
3. Define two distributable DVIPAs, 10.10.1.1 and 10.10.1.2, on system `sys_a` in sysplex Site1, and allow both `sys_a` and `sys_b` to be targets of these DVIPAs. 10.10.1.1 will be used for the Active/Standby workload and 10.10.1.2 will be used for the Active/Query workload.
4. Define two distributable DVIPAs, 10.10.2.1 and 10.10.2.2, on system `sys_c` in sysplex Site2, and allow both `sys_c` and `sys_d` to be targets of these DVIPAs. 10.10.2.1 will be used for the Active/Standby workload and 10.10.2.2 will be used for the Active/Query workload.
5. Configure and start the Lifeline Agents on each of the images in Site 1 and Site 2, by using the static VIPA that is defined in that image (192.10.1.1 on `sys_a`, 192.10.1.2 on `sys_b`, 192.10.2.1 on `sys_c`, and 192.10.2.2 on `sys_d`).

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 3. Use this configuration file for the Advisors on both the primary and secondary controllers.


```

#=====
#
# This file contains sample configuration statements for the
# Lifeline Advisor - using default settings in most cases.
#
#=====

# Following specifies the port which the primary Advisor will
# listen on for connections from Agents.
agent_connection_port      8181

# Following specifies which Agents are allowed to connect to
# the primary Advisor.
agent_id_list
{
    192.10.1.1..4000      # IP address and port for sys_a
    192.10.1.2..4000      # IP address and port for sys_b
    192.10.2.1..4000      # IP address and port for sys_c
    192.10.2.2..4000      # IP address and port for sys_d
}

# Following specifies the port which the primary Advisor will
# listen on for connections from a secondary Advisor.
advisor_connection_port    8000

# Following specifies which secondary Advisors are allowed to
# connect to the primary Advisor (should contain at least the IP
# address of this host).
advisor_id_list
{
    192.10.10.1 # IP address for Advisor on primary controller
    192.10.20.1 # IP address for Advisor on secondary controller
}

# Following specifies the IPv4 address and port the primary Advisor
# listens on for connections from load balancers. The IP address
# in this statement is also used by the primary Advisor to accept
# connections from Agents or the secondary Advisor.
lb_connection_v4          10.10.10.1..3860

# Following specifies which load balancers are allowed to
# connect to the primary Advisor.
lb_id_list
{
    192.100.1.1           # IP address of LB in Site1
    192.100.2.1           # IP address of LB in Site2
}

```

Figure 18. Sample configuration statements for the Lifeline Advisor, part 1 of 2

```

# Following specifies the IP address and site name (i.e.
# sysplex name) that identifies each z/OS sysplex distributor
# used to route the workload, the port number of the server
# application used for the workload, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
cross_sysplex_list
{
  10.10.1.1..80,SITE1,WORKLOAD1      # HTTP Server on Site1 -
                                     # Active/Standby
  10.10.2.1..80,SITE2,WORKLOAD1      # HTTP Server on Site2 -
                                     # Active/Standby
  10.10.1.2..80,SITE1,WORKLOAD2      # HTTP Server on Site1 -
                                     # Active/Query
  10.10.2.2..80,SITE2,WORKLOAD2      # HTTP Server on Site2 -
                                     # Active/Query
}
# Following specifies the Active/Query workload, its
# associated Active/Standby workload, the distribution type
# and distribution option, and timeout for persistent
# connections. These entries represent the workloads defined
# in the cross_sysplex_list that are classified as
# Active/Query workloads.
query_workload_list
{
  WORKLOAD2,WORKLOAD1,DYNAMIC,TOACTIVE,0
}

```

Figure 19. Sample configuration statements for the Lifeline Advisor, part 2 of 2

Follow the same steps in Example 1 to define an OMVS segment for the Advisor and the VIPARANGE DVIPA on both stand-alone LPARs.

Sample configuration statements for the Lifeline Agent

Use the same Lifeline Agent configuration file from Example 1 to define an OMVS segment for the Agent.

Based on this example, you need to update the TCPIP profile on the image owning the distributable dynamic VIPAs in the following way:

```

VIPADYNAMIC
; Uncomment the Site1 statements or Site2 statements that apply to the image where the
; z/OS sysplex distributor is being started and the site of the z/OS sysplex distributor
;
; Site1 definitions
; VIPADEFINE 255.255.255.255 10.10.1.1
; VIPADIST DISTMETHOD SERVERWLM 10.10.1.1 PORT 80 DESTIP ALL
; VIPADEFINE 255.255.255.255 10.10.1.2
; VIPADIST DISTMETHOD SERVERWLM 10.10.1.2 PORT 80 DESTIP ALL
;
; Site2 definitions
; VIPADEFINE 255.255.255.255 10.10.2.1
; VIPADIST DISTMETHOD SERVERWLM 10.10.2.1 PORT 80 DESTIP ALL
; VIPADEFINE 255.255.255.255 10.10.2.2
; VIPADIST DISTMETHOD SERVERWLM 10.10.2.2 PORT 80 DESTIP ALL
ENDVIPADYNAMIC

```

Multi-site Workload Lifeline configuration Example 4

The following figure shows an example of the Multi-site Workload Lifeline configuration where an Active/Standby SNA workload is defined. It also shows how to configure the Lifeline Advisors and Lifeline Agents to work in such an

environment.

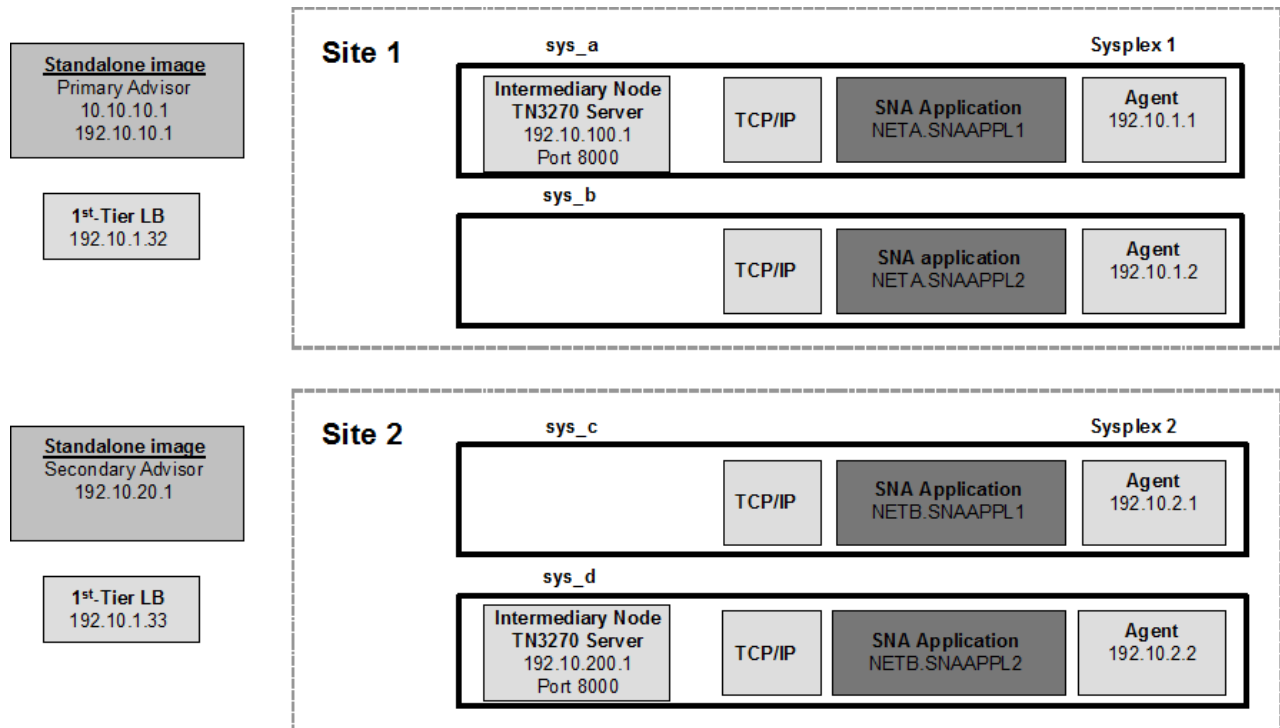


Figure 20. Multi-site Workload Lifeline configuration Example 4

To configure the previous sample environment shown in Figure 20, complete the following tasks:

1. Configure and start the Lifeline Advisor on the primary controller in the primary role. The primary Advisor listens on dynamic VIPA (DVIPA) address 10.10.10.1, which is defined in the VIPARANGE statement.
2. Configure and start the Lifeline Advisor on the secondary controller, in the secondary role. The secondary Advisor uses static VIPA address 192.10.20.1 to communicate with the primary Advisor.
3. Configure and start Lifeline Agents on each of the images in Site 1 and Site 2, using the static VIPA that is defined in that image (192.10.1.1 on sys_a, 192.10.1.2 on sys_b, 192.10.2.1 on sys_c, and 192.10.2.2 on sys_d).

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 4. Use this configuration file for the Advisors on both the primary and secondary controllers.

```

#####
#
# This file contains sample configuration statements for the Lifeline
# Advisor - using default settings in most cases.
#
#####

# Following specifies the port which the primary Advisor will listen
# on for connections from Agents.
agent_connection_port      8181

# Following specifies which Agents are allowed to connect to the
# primary Advisor.
agent_id_list
{
    192.10.1.1..4000      # IP address and port for sys_a
    192.10.1.2..4000      # IP address and port for sys_b
    192.10.2.1..4000      # IP address and port for sys_c
    192.10.2.2..4000      # IP address and port for sys_d
}

# Following specifies the port which the primary Advisor will listen
# on for connections from a secondary Advisor.
advisor_connection_port    8000

# Following specifies which secondary Advisors are allowed to connect
# to the primary Advisor (should contain at least the IP address of this
# host).
advisor_id_list
{
    192.10.10.1          # IP address for Advisor on primary controller
    192.10.20.1          # IP address for Advisor on secondary controller
}

# Following specifies the IPv4 address and port the primary Advisor
# listens on for connections from load balancers. The IP address in
# this statement is also used by the primary Advisor to accept
# connections from Agents or the secondary Advisor.
lb_connection_v4          10.10.10.1..3860

# Following specifies which load balancers are allowed to connect to
# the primary Advisor.
lb_id_list
{
    192.100.1.1          # IP address of LB in Site1
    192.100.2.1          # IP address of LB in Site2
}

```

Figure 21. Sample configuration statements for the Lifeline Advisor, part 1 of 2

```

# Following specifies the network qualified application
# name, site of each SNA application, and the workload that
# the SNA application belongs in.
# These entries are mapped by intermediary_node_list
# entries of type SNA.
sna_application_list
{
  NETA.SNAAPPL1,SITE1,WORKLOAD4
  NETA.SNAAPPL2,SITE1,WORKLOAD4
  NETB.SNAAPPL1,SITE2,WORKLOAD4
  NETB.SNAAPPL2,SITE2,WORKLOAD4
}

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each intermediary node used
# to route the workload, the port number to reach the
# intermediary node, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
# The application type is SNA because SNA applications process
# this workload.
intermediary_node_list
{
  192.10.100.1..8000,SITE1,WORKLOAD4,SNA # Intermediary node on Site1
  192.10.200.1..8000,SITE2,WORKLOAD4,SNA # Intermediary node on Site2
}

```

Figure 22. Sample configuration statements for the Lifeline Advisor, part 2 of 2

Follow the same steps in Example 1 to define an OMVS segment for the Advisor and the VIPARANGE DVIPA on both stand-alone LPARs.

Sample configuration statements for the Lifeline Agent

Use the same Lifeline Agent configuration file from Example 1 to define an OMVS segment for the Agent.

Multi-site Workload Lifeline configuration Example 5

The following figure shows an example of the Multi-site Workload Lifeline configuration where an Active/Standby Linux on System z workload is defined. It also shows how to configure the Lifeline Advisors and Lifeline Agents to work in such an environment.

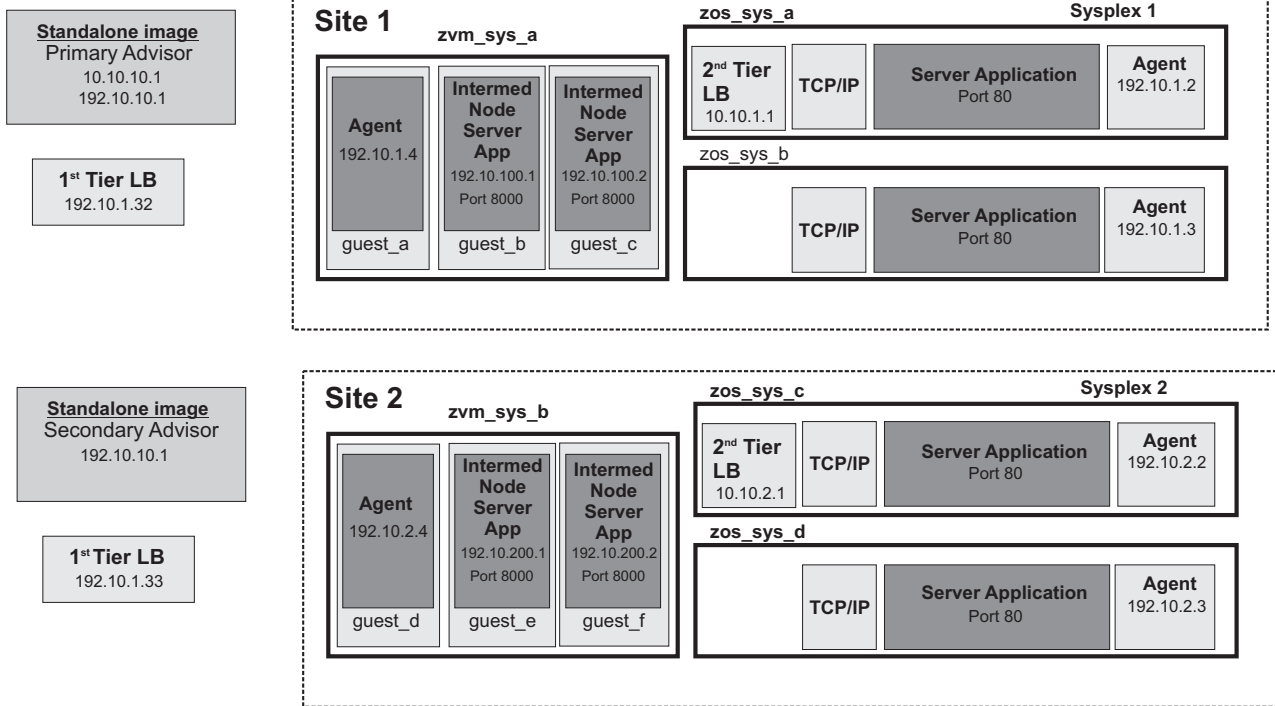


Figure 23. Multi-site Workload Lifeline configuration Example 5

To configure the previous sample environment shown in Figure 23, first repeat the initial five tasks in Figure 11 on page 257.

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 5. Use this configuration file for the Advisors on both the primary and secondary controllers.

```

#####
#
# This file contains sample configuration statements for the Lifeline
# Advisor by using default settings in most cases.
#
#####

# Following specifies the port which the primary Advisor will listen
# on for connections from Agents.
agent_connection_port      8181

# Following specifies which Agents are allowed to connect to the
# primary Advisor.
agent_id_list
{
    192.10.1.2..4000      # IP address and port for sys_a
    192.10.1.3..4000      # IP address and port for sys_b
    192.10.2.2..4000      # IP address and port for sys_c
    192.10.2.3..4000      # IP address and port for sys_d
    192.10.1.4..4000      # IP address and port for guest_a
    192.10.2.4..4000      # IP address and port for guest_d
}

# Following specifies the port which the primary Advisor will listen
# on for connections from a secondary Advisor.
advisor_connection_port    8000

# Following specifies which secondary Advisors are allowed to connect
# to the primary Advisor (should contain at least the IP address of this
# host).
advisor_id_list
{
    192.10.10.1          # IP address for Advisor on primary controller
    192.10.20.1          # IP address for Advisor on secondary controller
}

# Following specifies the IPv4 address and port the primary Advisor
# listens on for connections from load balancers. The IP address in
# this statement is also used by the primary Advisor to accept
# connections from Agents or the secondary Advisor.
lb_connection_v4          10.10.10.1..3860

# Following specifies which load balancers are allowed to connect to
# the primary Advisor.
lb_id_list
{
    192.100.1.1          # IP address of LB in Site1
    192.100.2.1          # IP address of LB in Site2
}

```

Figure 24. Sample configuration statements for the Lifeline Advisor, part 1 of 2

```

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each z/OS sysplex distributor
# used to route the workload, the port number of the server
# application used for the workload, and the workload name.
# These entries represent the load balancer server groups that are
# registered by the first-tier load balancers.
cross_sysplex_list
{
  10.10.1.1..80,SITE1,WORKLOAD5    # HTTP Server on Site1
  10.10.2.1..80,SITE2,WORKLOAD5    # HTTP Server on Site2
}

# Following specifies the IP address and site name (i.e. sysplex name)
# that identifies each intermediary node used
# to route the workload, the port number to reach the
# intermediary node, and the workload name.
# These entries represent the load balancer server groups
# registered by the first-tier load balancers.
# The application type is IP because TCP/IP applications process
# this workload.
intermediary_node_list
{
  192.10.100.1..8000,SITE1,WORKLOAD5,IP # Linux on System z server on Site1
  192.10.100.2..8000,SITE1,WORKLOAD5,IP # Linux on System z server on Site1
  192.10.200.1..8000,SITE2,WORKLOAD5,IP # Linux on System z server on Site2
  192.10.200.2..8000,SITE2,WORKLOAD5,IP # Linux on System z server on Site2
}

# These entries represent the Linux on System z z/VM guests hosting the
# application-tier intermediary node server applications.
management_guest_list
{
  guest_b,192.10.100.1,WORKLOAD5
  guest_c,192.10.100.2,WORKLOAD5
  guest_e,192.10.200.1,WORKLOAD5
  guest_f,192.10.200.2,WORKLOAD5
}

```

Figure 25. Sample configuration statements for the Lifeline Advisor, part 2 of 2

Follow the same steps in Example 1 to define an OMVS segment for the Advisor and the VIPARANGE DVIPA on both stand-alone LPARs.

Sample configuration statements for the Lifeline z/OS Agent

Use the same Lifeline Agent configuration file from Example 1 for the z/OS Agents, and follow the same steps in Example 1 to define an OMVS segment for the z/OS Agents and the z/OS sysplex distributor in both sites.

Sample configuration statements for the Lifeline Linux on System z Management Agent

You can use the following sample configuration file to configure the Lifeline Linux on System z Management Agents for Example 5.


```

=====
#
# This file contains sample configuration statements for the
# Lifeline Linux on System z Management Agent - most
# defaults are selected.
#
=====
# Following specifies the IP address and port of the primary Advisor
# that this Agent connects to.
advisor_id 10.10.10.1..8181

# Specifies the local IP address and port this agent binds to before
# connecting to the primary advisor. Uncomment the statement that
# applies to the image that the Agent is running on.
#host_connection 192.10.1.4..4000 # guest_a
#host_connection 192.10.2.4..4000 # guest_d

# Specifies the site name the Agent should be associated with when
# communicating with the primary Advisor. Uncomment the statement that
# corresponds to the site that the Agent is running in.
#site_name SITE1
#site_name SITE2

```

Figure 26. Sample configuration statements for the Linux on System z Management Agents

Multi-site Workload Lifeline configuration Example 6

The following figure shows an example of the Multi-site Workload Lifeline configuration where an Active/Standby MQ cluster workload is defined. It also shows how to configure the Lifeline Advisors and Lifeline Agents to work in such an environment.

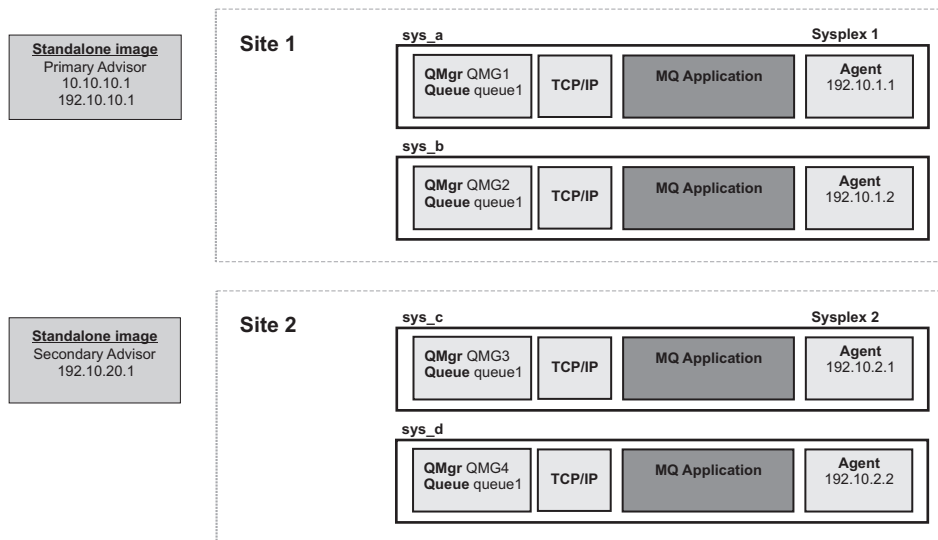


Figure 27. Multi-site Workload Lifeline configuration Example 6

To configure the previous sample environment shown in Figure 27, first repeat the initial five tasks in Figure 11 on page 257.

Sample configuration statements for the Lifeline Advisor

You can use the following sample configuration file to configure the Lifeline Advisor for Example 6. Use this configuration file for the Advisors on both the

primary and secondary controllers.

```
#####  
#  
# This file contains sample configuration statements for the Lifeline  
# Advisor - using default settings in most cases.  
#  
#####  
  
# Following specifies the port which the primary Advisor will listen  
# on for connections from Agents.  
agent_connection_port      8181  
  
# Following specifies which Agents are allowed to connect to the  
# primary Advisor.  
agent_id_list  
{  
    192.10.1.1..4000      # IP address and port for sys_a  
    192.10.1.2..4000      # IP address and port for sys_b  
    192.10.2.1..4000      # IP address and port for sys_c  
    192.10.2.2..4000      # IP address and port for sys_d  
}  
  
# Following specifies the port which the primary Advisor will listen  
# on for connections from a secondary Advisor.  
advisor_connection_port    8000  
  
# Following specifies which secondary Advisors are allowed to connect  
# to the primary Advisor (should contain at least the IP address of this  
# host).  
advisor_id_list  
{  
    192.10.10.1          # IP address for Advisor on primary controller  
    192.10.20.1          # IP address for Advisor on secondary controller  
}  
  
# Following specifies the IPv4 address the primary Advisor uses  
# to accept connections from Agents or the secondary Advisor.  
lb_connection_v4          10.10.10.1..3860  
  
# Following specifies the MQ Queue Manager name, MQ cluster name,  
# MQ cluster queue, and the site name (i.e. sysplex name) that  
# identifies each MQ Queue Manager used to route the workload.  
mq_manager_list  
{  
    QMG1,CLUSTER1,QUEUE1,SITE1,WORKLOAD1 # MQ Queue Manager on Site1  
    QMG2,CLUSTER1,QUEUE1,SITE1,WORKLOAD1 # MQ Queue Manager on Site1  
    QMG3,CLUSTER1,QUEUE1,SITE2,WORKLOAD1 # MQ Queue Manager on Site2  
    QMG4,CLUSTER1,QUEUE1,SITE2,WORKLOAD1 # MQ Queue Manager on Site2  
}
```

Figure 28. Sample configuration statements for the Lifeline Advisor

Follow the same steps in Example 1 to define an OMVS segment for the Advisor and the VIPARANGE DVIPA on both stand-alone LPARs.

Sample configuration statements for the Lifeline z/OS Agent

Use the same Lifeline Agent configuration file from Example 1 for the z/OS Agents, and follow the same steps in Example 1 to define an OMVS segment for the z/OS Agents and the z/OS sysplex distributor in both sites.

Appendix D. Sample AT-TLS configuration policies

Use AT-TLS to enable TLS before starting the TCP/IP stacks and the Lifeline Advisor and Lifeline Agent applications.

Example AT-TLS Policy for a Lifeline Advisor to peer Lifeline Advisor connection

```
TTLRule AdvisorToPeerRule
{
  LocalPortRange 8000
  Direction Inbound
  TTLGroupActionRef AdvisorGroup
  TTLEnvironmentActionRef AdvisorToPeerEnvironment
}
TTLGroupAction AdvisorGroup
{
  TTLEnabled On
}
TTLEnvironmentAction AdvisorToPeerEnvironment
{
  TTLSKeyRingParms
  {
    Keyring advisor_key_ring
  }
  TTLEnvironmentAdvancedParms
  {
    # TTLS will verify a user ID is associated with certificate
    ClientAuthType SAFCheck
    ApplicationControlled On
  }
  HandshakeRole ServerWithClientAuth
  Trace 7
}

TTLRule PeerToAdvisorRule
{
  RemotePortRange 8000
  Direction Outbound
  TTLGroupActionRef AdvisorGroup
  TTLEnvironmentActionRef PeerToAdvisorEnvironment
}
TTLGroupAction AdvisorGroup
{
  TTLEnabled On
}
TTLEnvironmentAction PeerToAdvisorEnvironment
{
  TTLSKeyRingParms
  {
    Keyring advisor_key_ring
  }
  HandshakeRole Client
  Trace 7
}
```

Example AT-TLS Policy for a Lifeline Advisor to Lifeline Agent connection

```
TTLRule AdvisorToAgentRule
{
  LocalPortRange 8181
  Direction Inbound
```

```

TTLSTLSGroupActionRef AdvisorGroup
TTLSEnvironmentActionRef AdvisorToAgentEnvironment
}
TTLSTLSGroupAction AdvisorGroup
{
TTLSEnvironmentEnabled On
}
TTLSEnvironmentAction AdvisorToAgentEnvironment
{
TTLSTLSKeyRingParms
{
Keyring advisor_key_ring
}
}
TTLSEnvironmentAdvancedParms
{
# TTLS will verify a user ID is associated with certificate
ClientAuthType SAFCheck
ApplicationControlled On
}
HandshakeRole ServerWithClientAuth
Trace 7
}

```

```

TTLSTLSRule AgentToAdvisorRule
{
RemotePortRange 8181
Direction Outbound
TTLSTLSGroupActionRef AgentGroup
TTLSEnvironmentActionRef AgentToAdvisorEnvironment
}
TTLSTLSGroupAction AgentGroup
{
TTLSEnvironmentEnabled On
}
TTLSEnvironmentAction AgentToAdvisorEnvironment
{
TTLSTLSKeyRingParms
{
Keyring agent_key_ring
}
HandshakeRole Client
Trace 7
}

```

Example AT-TLS Policy for a Lifeline Advisor to load balancer connection

```

TTLSTLSRule AdvisorToLBRule
{
LocalPortRange 3860
LocalAddr 10.10.10.1
Direction Inbound
TTLSTLSGroupActionRef AdvisorGroup
TTLSEnvironmentActionRef AdvisorToLBEnvironment
}
TTLSTLSGroupAction AdvisorGroup
{
TTLSEnvironmentEnabled On
}
TTLSEnvironmentAction AdvisorToLBEnvironment
{
TTLSTLSKeyRingParms
{
Keyring advisor_key_ring
}
}
TTLSEnvironmentAdvancedParms
{

```

```
# TTLS will verify a user ID is associated with certificate
ClientAuthType SAFCheck
ApplicationControlled On
}
HandshakeRole ServerWithClientAuth
Trace 7
}
```

Appendix E. Accessibility

Publications for this product are offered in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience problems, send an e-mail to mhvrcfs@us.ibm.com or write to:

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Mail Station P181
2455 South Road
Poughkeepsie, NY 12601-5400
U.S.A.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using such products to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. See *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Vol I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Notices

This information was developed for products and services offered in the USA.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 and cannot confirm the accuracy of performance compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe and PostScript are registered trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft 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.

Other product and service names might be trademarks of IBM or other companies.

Index

Numerics

1st-tier load balancers 10

A

a SNA application is unavailable 26
accessibility 279
activate 51, 52
activating workloads 142
advisor_connection_port 87
advisor_id 110
advisor_id_list 87, 111
advisor_takeover_policy 88
advisor_timeout 89
agent_connection_port 89
agent_id_list 90
AT-TLS
 Sample configuration policies 275
Authorizing
 Lifeline Agents 39
 monitor SNA applications 39

B

BCPii 34

C

changing the debug level 145, 163
CINET Considerations 57
configuration statements 85
 Lifeline Advisor 85
 Lifeline Agent 109
configure 18
configuring
 external load balancers 58
 first-tier external load balancers 59
 Lifeline Advisor 44, 48, 50
 Lifeline Agent 46
 Lifeline Agents
 WLM recommendations 48
 primary Advisor 44
 restart 48, 50
 second-tier external load balancers 60
 second-tier internal load balancers 61
 secondary Advisor 44
 support for Intermediary nodes 61
 syslogd 55
Configuring 43
 Base Control Program Internal Interface 54
 BCPii 54
cross_sysplex_list 91
customizing
 Message Revision Table 57
 NetView for z/OS 57
 TCP/IP profiles 55, 57
 TCP/IP stacks 55, 57

D

deactivate 53
deactivating workloads 144
debug levels 67
debug_level 93, 112
diagnosing
 debug levels 67
 diagnostic data 68
 failure of registering with ARM 82
 Lifeline Advisor connection termination 82
 Lifeline Advisor hung 79
 Lifeline Agent hung 79
 Lifeline Agent termination 81
 Linux on System z Management Agent 20
 Linux on System z workload 21
 load balancer connection termination 80
 Multi-site Workload Lifeline 67
 problems in the Linux 83
 SNA application 26
 SNA intermediary node servers are unavailable 27
 unreadable group names 80
 workload distribution problems 77, 78
 workload not distributed 74, 76
diagnostic data 68
directed mode 10
disability 279
dispatch mode 10
display command
 Lifeline Advisor
 configuration information 115, 169, 170
 debug level 115, 169, 170
 external load balancers 115, 169, 170
 Lifeline Agents 115, 169, 170
 peer Advisor 115, 169, 170
 work load 115, 169, 170
 Lifeline Agent
 configuration 151
 debug levels 151
 server applications 151
display information 165
dumping the internal debug records 146, 163

E

enabling server applications 158
external load balancers
 configuring 58
External load balancers
 directed mode 10
 dispatch mode 10

F

failure of registering with ARM 82
failure_detection_interval 93
first-tier external load balancers
 configuring 59

I

- IBM Software Support Center, contacting xii
- install 17
- Internal load balancers 12
- Interpreting
 - Workload and LB displays 70

L

- lb_connection_v4 96
- lb_connection_v6 97
- lb_id_list 98
- license, patent, and copyright information 281
- Lifeline 165
- Lifeline Advisor
 - activating workloads 142
 - changing the debug level 145
 - deactivating workloads 144
 - display command 115, 169, 170
 - dumping the internal debug records 146
 - quiescing workloads 146
 - sample configuration statements 257
 - taking over primary Advisor responsibilities 148
 - updating configuration 148
- Lifeline Advisor commands 115
- Lifeline Advisor configuration file 43
- Lifeline Advisor configuration statements 85
 - advisor_connection_port 87
 - advisor_id_list 87
 - advisor_takeover_policy 88
 - advisor_timeout 89
 - agent_connection_port 89
 - agent_id_list 90
 - cross_sysplex_list 91
 - debug_level 93
 - failure_detection_interval 93
 - lb_connection_v4 96
 - lb_connection_v6 97
 - lb_id_list 98
 - peer_advisor_id 103
 - update_interval 108
 - workload_switch_policy 109
- Lifeline Advisor connection termination 82
- Lifeline Advisor hung 79
- Lifeline Advisors 8
 - sample started procedure 253
 - starting 65
 - verifying 72
- Lifeline Agent
 - changing the debug level 163
 - display command 151
 - dumping the internal debug records 163
 - enabling server applications 158
 - quiescing server applications 161
 - sample configuration statements 257
- Lifeline Agent commands 151
- Lifeline Agent configuration file 43
- Lifeline Agent configuration statements 109
 - advisor_id 110
 - advisor_id_list 111
 - debug_level 112
- Lifeline Agent hung 79
- Lifeline Agents 9
 - sample started procedure 255
 - starting 65
 - verifying 72

- Linux on System z 13, 15, 18
- Linux on System z application groups 15
- Linux on System z Management Agent 16, 17, 18, 165
- Linux on System z Management Agent commands 165
- Linux on System z workload 13
- load balancer connection termination 80

M

- Management Agent 15, 18
- Management Guest 15
- management_guest_list 99
- messages 179
- migrate 51
- modify command
 - Lifeline Advisor
 - activating workloads 142
 - changing the debug level 145
 - deactivating workloads 144
 - Dumping the internal debug records 146
 - quiescing workloads 146
 - taking over primary Advisor responsibilities 148
 - updating configuration 148
 - Lifeline Agent
 - changing the debug level 163
 - Dumping the internal debug records 163
 - enabling server applications 158
 - quiescing server applications 161
- MQ cluster queue 76
- mq_manager_list 101
- Multi-site Workload Lifeline 171
 - configuration example 257
 - configuration statements 85
 - configuring 43
 - diagnosing 67
 - messages 179
 - overview 1
 - preparing 33

N

- Network Management Interface 171
 - message format 173

O

- operational characteristics 19, 24
- overview
 - Multi-site Workload Lifeline 1

P

- peer_advisor_id 103
- preparing
 - active connections 39
 - authorizing the Lifeline Advisors 34
 - authorizing the Lifeline Advisors and Agents 37
 - authorizing the Lifeline Agents 38, 39
 - BCPii 34
 - event-driven messages 37
 - granting users explicit authority 33
 - prerequisites 33
 - WLM services 38
- prerequisites 43

Q

- quiesce 52
- quiescing server applications 161
- quiescing workloads 146

R

- RACF 33, 38
- recommendations 8
- recovering
 - planned outage 49
 - TCP/IP stack 49, 50
 - unplanned outage 50

S

- Sample configuration policies 275
- sample configuration statements 257
- sample started procedure
 - Lifeline Advisors 253
 - Lifeline Agents 255
- second-tier external load balancers
 - configuring 60
- second-tier internal load balancers
 - configuring 61
- second-tier load balancers 10
- server application 74
- SNA 22
- SNA application 24
- SNA application groups 24
- SNA workload 22
- SNA workloads 24
- sna_application_list 106
- starting
 - Lifeline Advisors 65
 - Lifeline Agents 65
 - syslogd 55
- support for Intermediary nodes
 - configuring 61
- syntax diagram, how to read xiii
- syslogd 55
- syslogd priority level 67
- System z Management Agent 83

T

- taking over primary Advisor responsibilities 148
- the is not deactivated
 - workload 51
- trademark information 283

U

- unreadable group names 80
- update_interval 108
- updating configuration 148
- users explicit authority 33

V

- verifying
 - Lifeline Advisors 72
 - Lifeline Agents 72

W

- when not all are connected
 - Lifeline Agents 52, 53
- when not all first-tier load balancers are registered 52
- WLM recommendations 48
- WLM services 38
- workload 13, 22, 51, 52, 53
- workload distribution problems 77, 78
- workload not distributed 74, 76
- workload_switch_policy 109

Z

- z/OS Agents 24
- z/OS sysplex distributor 12

Communicating your comments to IBM

If you especially like or dislike anything about this document, you can send us comments electronically by using one of the following methods:

Internet email:

comsvrcf@us.ibm.com

World Wide Web:

<http://www.ibm.com/systems/z/os/zos/webqs.html>

If you would like a reply, be sure to include your name, address, and telephone number. Make sure to include the following information in your comment or note:

- Title and order number of this document
- Page number or topic related to your comment

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this document. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.



Product Number: 5655-UM4

Printed in USA

SC27-4653-03

