
z/OS Communications Server

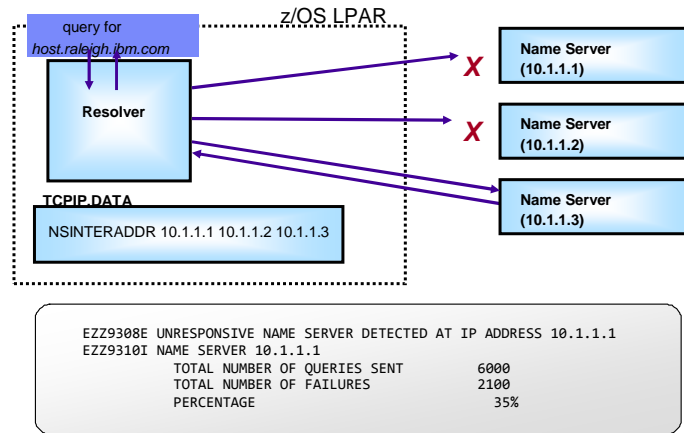
System resolver autonomics



© 2011 IBM Corporation

This presentation describes the updates to the system resolver for autonomics in z/OS® V1R13 Communications Server.

Background




In V1R12, z/OS Communications Server introduced the ability to have the resolver monitor for unresponsive name servers. The resolver determines the responsiveness of a name server based on statistics accumulated over 5-minute windows of activity.

In this example, you see three name servers have been defined in the TCPIP.DATA file. The first two name servers are not responding to resolver queries, delaying the completion of the application query. The resolver maintains statistics of how many queries are sent to each name server, and how many of those queries did not receive a response. The resolver uses these statistics to determine if a name server is unresponsive.

When the resolver determined that a name server was unresponsive, the resolver notified the network operator. You see an example of the messages generated to alert the network operator of the situation.

Stop using unresponsive servers

- Optionally stop forwarding new queries to an unresponsive name server
- Network operator notification function still supported



3

System resolver autonomies

© 2011 IBM Corporation

The monitoring support introduced in V1R12 issued messages but took no other action when a name server was determined to be unresponsive.

While the monitoring function remains available, starting in z/OS V1R13 Communications Server, you can optionally configure the resolver to stop forwarding queries to unresponsive name servers. In this example, the same three name servers are defined, and the first two name servers are again unresponsive. You see that with the autonomic quiescing function, the resolver does not forward the query to these unresponsive name servers. Instead, the resolver sends the request to the third name server in the list immediately, since that name server is the only responsive name server currently. Meanwhile, the resolver periodically sends a poll DNS query to the unresponsive name servers. When the name servers respond to the resolver polls at an acceptable rate, the resolver will resume sending application queries to the name servers. Network operator messages are issued to indicate when the resolver stops and resumes forwarding queries to a name server.

This function is optional. By default, the resolver performs the network operator notification level of monitoring introduced by z/OS V1R12.

Comparison of monitoring functions

	Network Operator Notification (V1R12)	Autonomic Quiescing (V1R13)
Measurement interval	60 seconds	30 seconds
Intervals considered	5	1 or 2
Minimum sample size	1 to become unresponsive, 0 to become responsive	10
Polling?	No	Yes
Default setting	Active with 25% threshold	Function not active
Application queries sent to unresponsive name server?	Yes	No, (with one exception)
Vehicle for reporting statistics	New EZZ9310I message issued every five minutes that a name server remains unresponsive	MODIFY RESOLVER command displays most recent failure rate for each name server
GLOBALTCPIPDATA	Not required	Required

4

System resolver autonomies

© 2011 IBM Corporation

You can use this table to compare the two levels of name server monitoring performed by the resolver.

The new autonomic quiescing function is more responsive to name server conditions, using only 30 or 60 seconds worth of statistics, compared to the previous 5-minute window. The autonomic quiescing function requires a larger sample size before declaring that a name server is unresponsive. The main reason for the larger sample size requirement is that the ramifications of declaring a name server to be unresponsive are more severe with the autonomic quiescing function. Note that you can switch from one method of monitoring to the other at any time. This will cause any statistics to be reset.

Because of the need for a larger sample size, the autonomic quiescing function can perform polling to supplement information gathered from application queries. Polling continues to occur even while the server is deemed unresponsive, so that the resolver can determine when the server becomes responsive again. Polling consists of UDP queries sent every six seconds. The resolver only polls name servers that are defined in the global TCPIP.DATA file.

The network operator notification function remains the default level of monitoring, using a threshold value of 25% failure rate. The two functions provide statistics related to the name servers in different ways. The network operator notification function reports statistics only for unresponsive name servers, while you can display the overall failure rate for every name server being monitored when the autonomic quiescing function is active.

Perhaps the biggest difference between the functions is that the autonomic quiescing function requires you to define a global TCPIP.DATA file, while the network operator notification function does not. The use of a global TCPIP.DATA file allows the resolver to accurately determine name server responsiveness without concern for different timeout values being used by different applications.

Enablement

- UNRESPONSIVETHRESHOLD option on the resolver setup statement
 - Default is 25% without AUTOQUIESCE
 - Specify a percentage
 - AUTOQUIESCE is an optional operand
 - Required to activate the autonomic quiescing function

You must code the AUTOQUIESCE operand on the UNRESPONSIVETHRESHOLD resolver setup statement to activate the autonomic quiescing function. The default is that autonomic quiescing is not active. In addition to coding AUTOQUIESCE, you must specify a percentage value on the UNRESPONSIVETHRESHOLD statement, even if you want to use 25% as the threshold value.

A percentage of zero disables all monitoring functions, regardless of whether you code AUTOQUIESCE or not.

AUTOQUIESCE is ignored if GLOBALTCPIPDATA is not specified. In this case, the resolver will perform only the network operator notification level of monitoring.

Operator notifications

```
EZZ9311E STOPPED USING NAME SERVER AT IP ADDRESS 10.1.1.1
EZZ9313I NAME SERVER 10.1.1.1
TOTAL NUMBER OF QUERIES SENT      510
TOTAL NUMBER OF FAILURES          306
TOTAL NUMBER OF RESOLVER POLLS SENT 10
TOTAL NUMBER OF POLL FAILURES      6
PERCENTAGE                        60%
```

```
EZZ9312I RESUMED USING NAME SERVER AT IP ADDRESS 10.1.1.1
```

In addition to no longer forwarding queries to the name server, the resolver alerts the network operator of the unresponsive name server using two messages. The first message, EZZ9311E, is an eventual action message that remains on the screen until you clear it or the resolver determines the name server is responsive again. The second message, EZZ9313I, provides the statistics used by the resolver to determine that the name server was unresponsive.

This pair of messages is generated for each name server it believes is unresponsive. Generally the statistics represent a 30-second interval. In the case where there are very few application queries, the statistics might cover an interval of 60 seconds in order to have a sufficient sample size.

When the resolver determines using polling that a name server is responsive again, the resolver issues message EZZ9312I to the operator console and application queries once again can be forwarded to the server. The resolver uses the threshold percentage specified on the UNRESPONSIVETHRESHOLD statement to determine if the name server is responding at a satisfactory rate to polling. Note that the resolver might continue to poll the name server even after it is deemed responsive, in order to have a sufficient sample size to detect a return to an unresponsive state.

If you disable autonomic quiescing or remove the name server from the global TCPIP.DATA file, the resolver will clear message EZZ9311E from the console but will not issue message EZZ9312I.

Displaying name-server status

```

F RESOLVER,DISPLAY
EZZ9298I DEFAULTTCPIPDATA - None
EZZ9298I GLOBALTCPIPDATA - SYS1.TCPPARMS(TCPDATA)
EZZ9298I DEFAULTIPNODES - USER55.ETC.IPNODES
EZZ9298I GLOBALIPNODES - None
EZZ9304I CACHE
EZZ9298I CACHESIZE - 200M
EZZ9298I MAXTTL - 2147483647
EZZ9298I UNRESPONSIVETHRESHOLD - 25
EZZ9304I AUTOQUIESCE
EZD2305I NAME SERVER 10.1.1.1
                STATUS: ACTIVE           FAILURE RATE: 0%
EZD2305I NAME SERVER 10.1.1.2
                STATUS: QUIESCED        FAILURE RATE: 100%
EZD2305I NAME SERVER 10.1.1.3
                STATUS: ACTIVE           FAILURE RATE: *NA*
EZZ9293I DISPLAY COMMAND PROCESSED

```

When the autonomic quiesce function is enabled, you can use the `MODIFY RESOLVER,DISPLAY` or the `MODIFY RESOLVER,REFRESH` command to display the current statistics pertaining to each name server in the global `TCPIP.DATA` file.


A name server can be either `ACTIVE`, in which case application queries are forwarded to it, or `QUIESCED`, in which case application queries are not forwarded to it but the resolver will continue polling. However, if all name servers are `QUIESCED` then application queries are forwarded.

The calculated failure rate might appear as “NA” in cases where there are not yet a sample size of 10 queries to measure the failure rate. This can occur for up to eight minutes after a configuration change or after the resolver or TCP/IP is started.

Note that this information is not available either when using the network operator notification function, or when all monitoring is disabled.

Health Checker support

- CSRES_AUTOQ_GLOBALTCPIPDATA
 - Checks for GLOBALTCPIPDATA if AUTOQUIESCE is enabled
- CSRES_AUTOQ_TIMEOUT
 - Checks for effective RESOLVERTIMEOUT if AUTOQUIESCE is enabled
- CSRES_AUTOQ_RESOLVEVIA
 - Checks for RESOLVEVIA UDP if AUTOQUIESCE is enabled
- Checks are performed
 - At resolver start
 - For MODIFY RESOLVER,REFRESH



The autonomic quiescing function has interactions with other resolver configuration settings. To help identify situations where these interactions are incorrect or inefficient, three resolver autonomic quiescing Health Checker tests are defined. The checks are performed by default whenever Health Checker is active and the resolver is started. The same checks are also performed when a `MODIFY RESOLVER,REFRESH` command is issued.

The first check verifies that the `GLOBALTCPIPDATA` statement is specified when the autonomic quiescing function is requested.

The second check verifies that the `RESOLVERTIMEOUT` value is less than or equal to a defined “maximum”. By the default, the timeout “maximum” value is five seconds, which is the recommended upper limit for this setting, especially for use with `AUTOQUIESCE`. You can change the value to be checked to the value you use in your installation by altering the check to have a `PARM` of `TIMEOUT(n)`, where *n* is the number of seconds to check for.

The last check verifies that the `RESOLVEVIA` value is coded as `UDP`. Using `UDP` is recommended, especially for use with `AUTOQUIESCE`, since the resolver uses `UDP` protocols for poll queries and this ensures that application requests and polling have consistent results.

Each of these three checks can be successful, unsuccessful, or not applicable (in case `AUTOQUIESCE` is not enabled). Appropriate messages are issued for each case.

Diagnosis

- If you suspect a problem related to resolver polling of unresponsive name servers, collect:
 - Resolver CTRACE records
 - Dump of resolver address space information
- If you suspect a problem related to resolver determination of name server responsiveness, collect:
 - Trace resolver of application resolver queries
 - Resolver CTRACE records
 - Dump of resolver address space information



In general, if you suspect there is a problem with resolver processing, Resolver CTRACE and trace resolver information should be collected for IBM Service to use. If you suspect a problem with the resolver polling processing, however, there is no trace resolver information to collect. The resolver creates only CTRACE records in the new polling processing.

A dump of the system, including the resolver address space data, is also very useful to IBM Service when debugging resolver problems.

Trace resolver example

```
res_init Resolver values:  
Global Tcp/Ip Dataset = SYS1.TCPPARMS(RESGLOBL)  
. . .  
UnresponsiveThreshold = 25,AUTOQUIESCE  
. . .  
res_init Succeeded
```

```
res_send Name Server Capabilities  
Monitoring intervals used = 1  
Name server 10.1.1.1  
EDNS0 Support = unknown  
Queries sent = 0  
Failures = 0  
Percentage = 0%  
Skip this name server, auto-quieted  
Name server 10.1.1.2  
EDNS0 Support = unknown  
Queries sent = 200  
Failures = 20  
Percentage = 10%  
Use this name server for queries  
Name server 10.1.1.3  
EDNS0 Support = unknown  
Queries sent = 20  
Failures = 0  
Percentage = 0%  
Use this name server for queries  
res_send DNS is unresponsive, skipping 10.1.1.1  
res_send Sending query to Name Server 10.1.1.2
```

Trace resolver output is updated to reflect autonomic quiescing information.

As shown on this slide, the `res_init` trace resolver output displays whether `AUTOQUIESCE` is enabled and the configured unresponsive threshold percentage.

Additionally, the `res_send` trace resolver output is updated to display autonomic quiescing information for each name server, and to indicate when name servers are skipped because they are deemed to be unresponsive by the resolver. You can tell whether `AUTOQUIESCE` is in effect by examining the number of "Monitoring intervals used." If this value is one, then `AUTOQUIESCE` is enabled. If this value is five, then only the network operator notification function is enabled. No statistics are displayed if all monitoring is disabled. The statistics displayed here do not include polls that are sent to the name server.

Things to think about

- Consider using network operator notification to estimate proper threshold
- SMTP and BIND DNS utilities (dig, nslookup) are not affected
- CINET configurations require special consideration
- MAXSOCKETS might need to be adjusted
- The resolver is now an OMVS process:
it now stops when OMVS stops



To estimate a proper unresponsive threshold for your environment, use the network operator notification function to measure typical failure rates in your environment. Start with a 1% threshold and monitor EZZ9310I messages to determine the highest error rate experienced during normal operation. When you enable AUTOQUIESCE, set the failure threshold to a value greater than this normal failure rate.

Applications with their own resolver code are not covered by the autonomic quiescing function. This includes SMTP and the BIND DNS utility applications. DNS queries generated by these applications are sent to the name server even if the system resolver believes the name server is unresponsive.

If you use common INET (CINET), autonomic quiescing might not be usable in your configuration. If a name server can be reached from only selected stacks, the resolver might erroneously conclude that the name server is unresponsive if requests are sent from stacks that have no connectivity to the name server. This can result in the name server being skipped even for queries sent from stacks that do have connectivity to the name server. If your installation does not have connectivity to all name servers from all stacks, use the network operator notification level of monitoring.

Because resolver opens a new socket for each name server it polls, you might need to adjust your MAXSOCKETS setting. If you closely control the number of sockets on your system, ensure that the resolver can open a polling socket for each name server. If the resolver cannot obtain a socket, the resolver will not poll the name server and the name server might be deemed unresponsive indefinitely.

In z/OS V1R13, the resolver is now an OMVS process. If you shutdown OMVS using a MODIFY command, the resolver now stops automatically as part of the shutdown processing. If you then restart OMVS using a MODIFY command, the resolver restarts automatically. This is similar to how OMVS automatically starts resolver during system initialization. The same BPXPRMxx definition used to start resolver during initialization is used during MODIFY OMVS,RESTART processing.



Feedback

Your feedback is valuable

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

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

Click to send email feedback:

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

This module is also available in PDF format at: [../resolver.pdf](..../resolver.pdf)

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



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and z/OS 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 other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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

© Copyright International Business Machines Corporation 2011. All rights reserved.