


IBM Tivoli Netcool/OMNIbus V7.3.1

Architectural redundancy options for probe data flow resilience



© 2013 IBM Corporation

In this training module, you learn about the architecture redundancy options that provide resiliency to the flow of data from IBM Tivoli® Netcool/OMNIbus V7.3.1 probes.

Objectives

- After completing this training module, you can accomplish these tasks:
- Explain probe resilience
- Configure a peer-to-peer probe pair
- Decide when to use a peer-to-peer probe pair

The purpose of this IBM Tivoli Netcool/OMNIbus V7.3.1 training module is to provide an understanding of probe resilience within a multitier system architecture.

Resilient probe considerations

- Peer-to-Peer
 - P2P with duplicate event sources provides probe resilience
- Cold Standby, when duplicate event sources are not available
 - You can use custom triggers to manage cold stand-by probe failover
- Object Server failover
 - Using the ServerBackup property provides object server dual resilience
- Store and forward (SAF)
 - When both object servers are unavailable for short periods, using SAF provides a solution.
 - During longer object server outage periods, SAF files roll over

Peer-to-peer requires two identical event sources. This is because the slave probe discards all events when the master probe is reachable from the slave probe's physical host.

The peer-to-peer feature is not available for all probes. To check if the feature is available, run the probe with the `-dumpprops` or the `-help` option and check for the peer-to-peer settings.

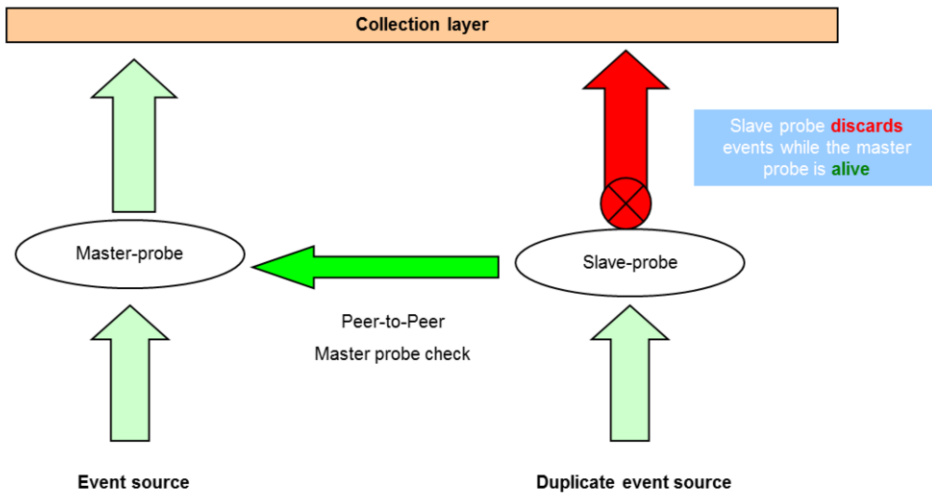
When only one event source is available, keep the backup probe in cold standby. This behavior requires custom configuration. All probes support `Server` and `ServerBackup` properties that allow probes to failover and failback between a primary and backup object server. All probes support store and forward. SAF stores event data in files on the probe server when the object server is unavailable. The SAF files have a finite size and event data is lost when the SAF files overflow.

Peer-to-Peer properties

- Master probe properties
 - Mode : 'master'
 - PeerHost : 'slave_host'
 - Peerport : 9001
 - BeatInterval : 30
 - BeatThreshold : 1
- Slave probe properties
 - Mode : 'slave'
 - PeerHost : 'master_host'
 - Peerport : 9001
 - BeatInterval : 30
 - BeatThreshold : 1

The peer-to-peer properties are Mode, PeerHost, PeerPort, BeatInterval, and BeatThreshold. Mode and PeerHost interchange between the probe's positions in the peer-to-peer pair, with PeerHost being the IP Address or FQDN of the other probe server in the peer-to-peer pair. Peerport is a port that is available for use on both servers. After you configure Mode, PeerHost, and PeerPort, the peer-to-peer feature functions under normal conditions. The properties BeatInterval and BeatThreshold are provided to allow you to configure the peer-to-peer pair to tolerate network and system conditions. Setting BeatInterval to larger values, reduces the load on the probes and the network. Setting the BeatThreshold to a larger value allows for longer delays in the probes processing of peer-to-peer messages.

Peer-to-Peer: When the master probe is alive



5

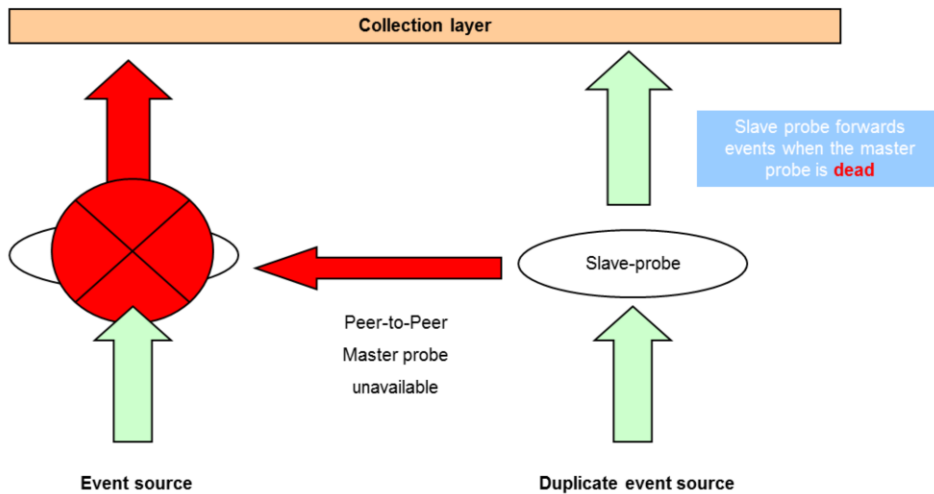
Architectural redundancy options for probe data flow resilience

© 2013 IBM Corporation

Under normal circumstances, the slave probe checks that the master probe is available by using the Peerport. When the master probe is available, the slave probe discards all of the events it receives.

The master probe forwards events regardless of the status of the slave probe.

Peer-to-Peer: When the master probe is dead



6

Architectural redundancy options for probe data flow resilience

© 2013 IBM Corporation

When the master probe is shutdown, the slave probe is unable to check its status with the Peerport. The slave probe then forwards the events it has stored, since the last successful check of the master probe, to the object server it is connected to.

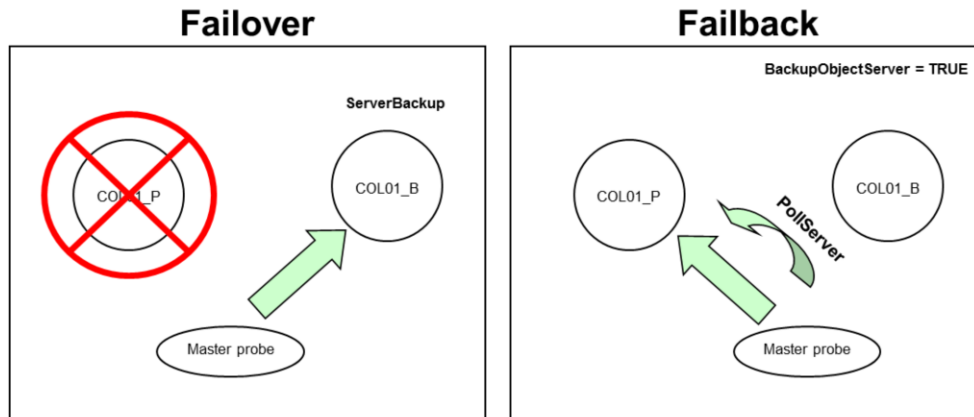
ObjectServer probe failover and failback

- For failover
 - Set **ServerBackup** to the backup object server
- For failback
 - **PollServer** must be larger than **NetworkTimeout**
 - Backup Object servers must have the property **BackupObjectServer** set to TRUE

```
Server           : 'COL_P'  
ServerBackup     : 'COL_B'  
NetworkTimeout   : 10  
PollServer       : 20
```

All probes support object server failover and failback. In a multitier environment **Server** and **ServerBackup** are set to collect object server names, or when you are not using a collection layer, to aggregate object server names. The use of a virtual object server, such as **AGG_V**, is not a desirable configuration, as probe failback is unmanaged, and requires a restart of the backup aggregation object server. In the example property settings, **NetworkTimeout** is set to ten seconds, and **PollServer** to twenty seconds. Typically, **PollServer** is set to twice the value of **NetworkTimeout**, where **NetworkTimeout** is a value that allows detecting that the connection to the object server is unavailable.

Object Server failover and failback



When the **Server** is down the probe forwards events to the **ServerBackup**

When the Server returns to operation the probe detects it within **PollServer** seconds

8

Architectural redundancy options for probe data flow resilience

© 2013 IBM Corporation

Under normal conditions the master probe is connected to the primary collection object server. When the object server is detected as being unavailable, the master probe switches to the ServerBackup object server. This switchover happens when an event is ready to send to the object server, and within the period NetworkTimeout defines. If the primary object server, that Server defines, becomes available, the master probe fails back within the period the property PollServer specifies. If NetworkTimeout is greater than PollServer, the master probe never fails back automatically.

Store and Forward

- Probes
 - AutoSAF
 - Stores data to files if the object server is unavailable at start-up
 - StoreAndForward
 - 1 – Normal SAF
 - Stores data to files if the object server is unavailable up to file size limits
 - 2 – Circular SAF
 - Always stores data in a set of rotated files

In addition to ServerBackup, all probes support store and forward. There are two types of store and forward, AutoSAF and SAF. AutoSAF works automatically, and does not require the object servers to be available to be initiated. SAF requires that the object servers are running before it is initiated. SAF has two modes, normal and circular. The administrator decides if store and forward is required, under what circumstances SAF needs to activate, and how much event data it needs to preserve.

Summary

Now that you have completed this training module, you can accomplish these tasks:

- Decide which type of probe resilience is suitable for your installation
- Configure a peer-to-peer master and slave probe in your environment
- Configure a probe for ObjectServer failover and failback

Now that you have completed this training module, you can accomplish these tasks:

Decide which type of probe resilience is suitable for your installation

Configure a peer-to-peer master and slave probe in your environment

Configure a probe for ObjectServer failover and failback

Useful reference information

- To check which features are available for a probe you can use the command line switches
 - -dumpprops
 - -help
- The online probes and gateway manual covers how the command probe properties affect the probe's behavior

To check which features are available for a probe you can use the command line switches

-dumpprops

-help

The online probes and gateway manual covers how the command probe properties affect the probe's behavior



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and Tivoli 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>

Other company, product, or service names may be trademarks or service marks of others.

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 2013. All rights reserved.