

An Oracle White Paper

March 2014

Dead Connection Detection

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Executive Overview	2
Introduction	2
Pre-12c Mechanism.....	2
12c Mechanism	3
Parameters to enable Dead Connection Detection	4
Conclusion	4

Executive Overview

Dead Connection Detection (DCD) is a feature of Oracle Net which helps in recovering resources allocated for a connection that is no longer usable. DCD has been enhanced in Oracle Database 12c to reduce detection time drastically from about 15 minutes to as low as 2-3 minutes. This paper provides an overview of 12c and pre-12c DCD mechanisms, configuration details and benefits of the 12c mechanism.

Introduction

DCD is intended primarily for environments in which clients power down their systems or client machines crash unexpectedly without disconnecting from Oracle Database sessions. An early detection of such scenarios enables speedy recovery of database resources. A less common usage scenario for DCD is to keep database connections alive when an external firewall timeout is configured to terminate idle connections.

This feature is configured on Oracle Database server side and is controlled/enabled using parameter `sqlnet.expire_time` in `sqlnet.ora`. In 12c as well as pre-12c releases, probe packets are sent to the client by server to check if the connection is still usable; the difference being that in pre-12c, these probe packets are initiated at NS (Network Session) layer of Oracle Net, whereas in 12c, the probes are initiated at TCP level.

Pre-12c Mechanism

Pre-12c Server process sends SQL*Net probe packets to verify connectivity if the connection has been idle for more than the interval specified in `sqlnet.expire_time`. If sending of the probe fails, an error is returned, causing the server process to exit. The time taken for send failures to unresponsive node in TCP depends on system-wide TCP parameters. The relevant parameters as per Linux `tcp(7)` man page are:

`tcp_retries1`

The number of times TCP will attempt to retransmit a packet on an established connection normally, without the extra effort of getting the network layers

involved. Once we exceed this number of retransmits, we first have the network layer update the route if possible before each new retransmit. The default is the RFC specified minimum of 3.

`tcp_retries2`

The maximum number of times a TCP packet is retransmitted in established state before giving up. The default value is 15, which corresponds to a duration of approximately between 13 to 30 minutes, depending on the retransmission timeout. The RFC 1122 specified minimum limit of 100 seconds is typically deemed too short.

12c Mechanism

Dead Connection Detection has been enhanced in Oracle Database Release 12c to reduce the amount of time taken to detect terminated client connections. If the system supports TCP keep-alive tuning per socket, Oracle Net Services automatically uses the enhanced detection model. TCP keep-alive parameters are tuned at a per-connection level to detect connection death. With this approach, TCP keep-alive probes are sent after the connection has been idle for `sqlnet.expire_time` and within maximum 1 minute the health of connection is ascertained.

The following three parameters associated with TCP keep-alive probes are tuned:

1. `TCP_KEEPIDLE`: specifies the timeout, with no activity until the first keep-alive packet is sent. Default 2hrs
This parameter takes its value from `SQLNET.EXPIRE_TIME`.
2. `TCP_KEEPCNT`: The number of keep-alive probes sent. (Linux default 9)
`TCP_KEEPCNT` is always set to 10.
3. `TCP_KEEPINTVL`: specifies the interval, between when successive keep-alive packets are sent if no acknowledgement is received. (Linux default 75)
`TCP_KEEPINTVL` is always set to 6.

Another advantage of this mechanism is that the probes are implemented by TCP stack and do not need to be consumed by application at the receiver side. So even in cases

where the receiver is busy and not in a position to read incoming data from the network, there would not be any buffering of probe packets.

The new 12c mechanism is currently supported on all platforms except Solaris.

Parameters to enable Dead Connection Detection

`SQLNET.EXPIRE_TIME` parameter can be set in `sqlnet.ora` to specify a time interval, in minutes. This value denotes the idle timeout after which the probes for DCD are sent. If the system supports TCP keep-alive tuning, Oracle Net automatically uses the enhanced detection model, and tunes the TCP keep-alive parameters.

Default Value: 0

Recommended Value: 10

12c users can revert to the pre-12c mechanism of sending Oracle Net NS probe packets by setting `USE_NS_PROBES_FOR_DCD=true` in `sqlnet.ora`.

Note that, even though it is very small, probe packets generate additional traffic on the network.

Conclusion

12c Dead Connection Detection mechanism is faster as it does not rely on TCP send failures. It also reduces the overhead of probe packet processing on client side as well as overhead of sending Oracle Net NS probe packets on the server side.



Dead Connection Detection
March 2014
Authors: Bhaskar Mathur, Feroz Khan,
Kant Patel

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2014, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.