

Oracle® SD-WAN Edge

Best Path Determination



Original Publication Date: October 29, 2019



Copyright © 2019, 2007 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. Windows® 7 and Windows® XP are trademarks or registered trademarks of Microsoft Corporation.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Table of Contents

How Talari Determines Path State & Best Path Available 4
Standby WAN Links in Best Path Selection 6

My Oracle Support

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support (CAS) can assist you with My Oracle Support registration.

Call the CAS main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

1. Select 2 for New Service Request.
2. Select 3 for Hardware, Networking, and Solaris Operating System Support.
3. Select one of the following options:
 - For technical issues such as creating a new Service Request (SR), select 1.
 - For non-technical issues such as registration or assistance with My Oracle Support, select 2.

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system

- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center (OHC) site, <http://docs.oracle.com>. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at <http://www.adobe.com>.

1. Access the Oracle Help Center site at <http://docs.oracle.com>.
2. Click Industries.
3. Click the Oracle Communications link.

Under the SD-WAN header, select a product.

4. Select the Release Number.

A list of the entire documentation set for the selected product and release appears.

5. To download a file to your location, right-click the PDF link, select Save target as (or similar command based on your browser), and save to a local folder.

How Talari Determines Path State & Best Path Available

- Talari APN Appliances (APNA) have the ability to determine which path to send traffic over on a per-packet basis.
- This is not necessarily the most expensive link or the one with the most bandwidth.
- There are a variety of factors that a Talari APNA considers when determining which path to put a packet on as it adapts to the conditions of the WAN.

When a Talari APNA receives a packet from the Conduit (a WAN Egress packet), it calculates the quality of the path from which it was received, based on one-way trip latency, loss, bandwidth, and jitter. When the Talari APNA has a packet to send on the Conduit (a WAN Ingress packet), the Talari APNA checks the quality of all paths and sends the packet over the best available path.

All paths within the Talari APNA are classified with a use state of GOOD, BAD, or DEAD. A path is declared DEAD if there are no Conduit packets received on the path for 1.5 seconds and the Talari APNA will not attempt to use a DEAD path.

To declare a path BAD is a bit more complicated and the Talari APNA uses a sophisticated algorithm to calculate the loss, latency, and jitter before making this determination. For example, 3 lost packets out of 4, 4 lost packets out of 10, are examples of what will cause a path to be marked BAD. The Talari APNA will send packets over a BAD path only if that is the only path available, or if the configuration requires the Talari APNA to do so, such as disabling Bad Loss Sensitivity (this will force the Talari APNA to keep a path marked GOOD even when loss is detected). If all paths are BAD, the Talari APNA will use the BAD path containing the most bandwidth. In addition, when a path transitions from BAD back to GOOD, there is a 10 second hold time on the path prior to customer frames traversing this path.

For all GOOD paths, the Talari APNA uses this same method of calculation to determine which of the GOOD paths is declared the best path. This is done on a per-packet basis, not measured by amount of time. One packet may be sent using an Internet path and the next may be sent via MPLS, unless Persistent Path is configured in the Talari APNA (this configuration option will be explained below). MTU and number of hops are not taken into consideration for making this decision. The remote Talari

APNA, or WAN Egress processor, is the brains reporting back to the local Talari APNA, or WAN Ingress processor, about the condition of the network. Every 2-3 packets the WAN Egress processor receives is the sample size used to report conditions back to the local device. This process is analogous to predicting the weather. You are more likely to be correct in attempting to predict what the weather will be in the upcoming 20 minutes than to try to predict what it will be a few hours from now.

Once the Talari APNA determines the best path, traffic may be influenced further depending on specific options enabled within the Talari APNA config. The four options for mode of transmission are, Duplicate Path, Load Balancing, Persistent Path, and Override Service.

After making a rule for traffic within the Talari APNA configuration you will be prompted to select the Transmit Mode. If selecting the Duplicate Path option, the Talari APNA will make a copy of a packet and send both packets down two separate paths at the same time. The packet to make it to the destination first will be used, the other discarded. The Talari APNA will select the best path using the algorithm discussed above, and select the second path based on the least amount of similarities to the first one, to guarantee variety between them. For example, two AT&T Links would not be selected. This option is typically chosen for VoIP traffic due to the small nature of the UDP packets and should only be considered if you have the additional bandwidth required to send duplicate packets.

With the Load Balancing option, traffic is put on the best path until that path runs out of available bandwidth. The additional traffic will then spill over to the next best path. For example with two WAN Links, one containing 1.5M of bandwidth 10ms of latency, and the other with 5M of bandwidth but 20ms latency detected. Path 1 will be considered the best path by the Talari APNA and used initially when trying to push 4M of traffic. Talari APNA will fill this pipe with 1.5M, but after 5 ms Talari APNA will spill the remaining traffic over onto path 2. Once path 1 has the bandwidth necessary, if it is still considered the better path due to the lower latency value, traffic will spill back from path 2 to path 1 seamlessly. Load Balancing is the default option when creating rules in the Talari APNA config.

With Persistent Path, the latency is allowed to build for 50ms on the best path before traffic jumps onto the next best path. Traffic will begin and continue on the best path until the latency has built up to the point where the second path is now considered

better than the first. Traffic will continue on path 2 until the latency on path 1 scales back down to a value below that of path 2. Once this level has been reached, the Talari APNA will again consider path 1 to be more reliable than path 2. Persistent path should be configured for real-time and interactive traffic only, not for bulk data unless the exact amount of bandwidth per application is known and available.

Override Service is an option for traffic you do not want put into the Conduit. Once you select Override Service, you must then specify Intranet, Internet, or Passthrough service for this traffic to use as an alternative to Conduit service. For example, if you have specific traffic you do not want Talari APNA to manipulate in any way, you can create a rule with specific IP addresses or port numbers, then select Override Service as your transmit mode and Passthrough as the service you want this traffic overridden to.

Standby WAN Links in Best Path Selection

WAN Links configured to use the criteria explained above for packet path selection purposes are referred to as Regular Active WAN Links. However, WAN Links are provided via many mediums, such as Wireless, Microwave, LTE, and Satellite, as well as cost that can be fixed or on a per bit basis. Customers may require these types of WAN Links only be used when no other bandwidth is available or if the site has multiple links down at one time. This is especially true if the link is billed on a cost per bit basis; for example when running a backup LTE link at a site that has a data cap.

The Standby WAN Link option is an available configuration setting for such scenarios, but this creates a set of paths that differ in the path selection algorithm from those of Regular Active WAN Links.

Standby WAN Links can function as a Last-Resort Standby to be turned on when all paths are dead, or as an On-Demand Standby to be turned on when some paths are dead. Standby WAN Links can be configured with heartbeat traffic disabled completely, or heartbeat traffic enabled and sent at a specific interval. If Standby WAN Links are configured with heartbeats disabled, the path is assumed to be in a GOOD state, however no data traffic traverses the link until the configured criteria is met. If On-Demand Standby Links are configured, an additional small latency penalty is added so traffic prefers Regular Active WAN Links.

For more information concerning Standby WAN Link configuration and capability, please refer to the APN Release 5.2 New Features Guide.

