

Oracle® SD-WAN Edge

Features Guide



Release 9.0

F29168-04

June 2021

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2019, 2021, Oracle and/or its affiliates.

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, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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 Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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.

Contents

About This Guide

My Oracle Support x

Revision History

1 Release 2.3 Features

Geographic Redundancy	1-1
GRE Support	1-3
GRE Header Inspection Support	1-3
GRE Header Compression	1-4
Oracle Route Support	1-4
Multiple Intranet Services Defined	1-4
Route Learning via SNMP	1-5
SNMPv2 Route Polling Configuration	1-6
Include/Exclude rules	1-6
Included Routes	1-6
Excluded Routes	1-6
Intranet & Internet Enhancements	1-7
Intranet Name Support	1-7
Rule and Class Improvements	1-7
Default Classes:	1-8
Default Rules per Conduit:	1-8
WAN Link and Path Enhancements	1-13
Reserve Minimum Bandwidth for Conduit WAN Links	1-13
Configurable Congestion Control per WAN Link	1-14
Path Eligible Setting for Traffic Types	1-15
Reporting Enhancements	1-16
Availability Report	1-16
QoS Reports	1-17
Usage Report	1-17
Periodic Status Reports	1-18

Double Event Triggers	1-19
Observed Protocols	1-19
Enhanced Network Change Management	1-20
The Change Management Workflow	1-21

2 Release 2.4 Features

Network Functionality and Deployability	2-1
route_eligibility_based_on_path =Boolean	2-2
route_eligibility_to_wan_link_name =Text	2-2
L2 MAC Learning for Multiport Bridging	2-2
Design Considerations	2-3
udp_port_num =Number (2156)	2-4
udp_port_num_alt =Number (2156)	2-4
udp_port_switch_interval_minutes =Number (1440)	2-4
Network Topology	2-4
Support for Serial High Availability Appliances	2-5
Oracle Serial HA	2-6
Design Considerations:	2-7
Multiple VLAN Segment on Common WAN link	2-7
Design Considerations:	2-9
Path MTU Discovery	2-9
Design Considerations:	2-11
Usability	2-11
The Site Cloning Process	2-11
Next Steps in the Site Cloning Process	2-12
Improved Site WAN Link Provisioning	2-13
WAN Link Rates	2-14
Provisioning Groups	2-14
The Concept of Using Shares	2-15
Services	2-16
Shares of Group	2-16
Glossary	2-16

3 Release 2.5 Features

Oracle Hardware Support	3-1
Introducing the T5000 Appliance	3-1

4 Release 3.0 Features

Dynamic Conduits	4-1
------------------	-----

Design Considerations	4-1
Dynamic Conduit Configuration	4-2
Dynamic Conduit Configuration Creation	4-2
WAN To WAN Forwarding Enhancements	4-6
Routing Enhancements	4-7
Intranet or Internet Fallback Routes	4-7
Additional Enhancements	4-7

5 Release 3.1 Features

Default Configuration Parameter Change	5-1
Oracle SD-WAN Edge Configuration Editor	5-1
Oracle SD-WAN Edge Configuration Editor and Oracle SD-WAN Edge Aware	5-2

6 Release 4.0 Features

256-Site Adaptive Private Networks	6-1
Changes to Data Storage on Oracle SD-WAN Edge Appliances	6-2
Changes to Local Route Scale	6-2

7 Release 4.1 Features

Oracle Virtual Appliance CT800	7-1
MOS Estimation	7-1
How to Configure	7-2
How to View	7-3
Security Enhancements	7-3
Summary	7-3
How to Configure	7-4
SNMP Polling for ARP Table	7-4
Appliance Settings from Aware	7-4

8 Release 4.2 Features

Non-Resetting Configuration Updates	8-1
Configuration Updates	8-1
Impact of Common Configuration Updates	8-2
Software Updates	8-3

9 Release 4.3 Features

Configure Private MPLS WAN Links	9-1
----------------------------------	-----

Add Private MPLS WAN LINK	9-2
Define WAN Link Basic Properties (Private MPLS)	9-3
Assign Autopath Group to Conduit-WAN Link	9-4
Verify Autopath Creation	9-4
View Permitted Rate and Congestion for WAN Links	9-5
View Permitted Rate	9-5
View Congestion	9-5
Configuration Versioning	9-6
Support for Installing User-Generated Certificates on Appliances	9-7

10 Release 4.4 Features

LAN GRE Tunnels	10-1
Monitor LAN GRE Tunnels	10-2
IPsec Encryption in Conduit	10-2
Monitoring IPsec	10-3
Path State Configurability and Monitoring	10-4
Monitor Statistics	10-5
Availability Reports	10-7
Additional Enhancements	10-9
Appliance T5200 Support	10-10
Oracle Virtual Appliance VT500 Support	10-10

11 Release 5.0 Features

Enhanced Match Criteria for Rules	11-1
Virtual Routing and Forwarding (VRF)	11-1
Monitoring	11-6
Dynamic Routing	11-8
Virtual IP Address Identity	11-8
Open Shortest Path First (OSPF) Routing Protocol	11-8
Interior Border Gateway Protocol (IBGP)	11-10
Filters	11-12
Network Objects	11-13
Monitoring	11-14
WAN Link IP Address Learning (DHCP Client)	11-14
Monitoring	11-15
IPsec VPN Termination	11-16
Monitoring	11-20
Standby WAN Links	11-21

12 Release 5.1 Features

Virtual Appliance VT800	12-1
Alarm System	12-1
Diagnose Alarms	12-2
Route Export Filters	12-3
Operating System Patching	12-4
Customizable Web Console	12-4
DHCP Relay and DHCP Server	12-6

13 Release 5.2 Features

Support for 550 Sites	13-1
Stateful Firewall	13-1
DHCP Relay & DHCP Server	13-1
Standby WAN Link (VSAT)	13-5
Adaptive Bandwidth Detection	13-8
Active Bandwidth Testing	13-9
SNMPv3 Polling and Trap Capability	13-11
Eligibility for IPsec Non-Conduit Routes	13-11
Additional Enhancements	13-12
Routing Enhancements	13-12

14 Release 6.0 Features

Application Packet Filtering	14-1
Applications	14-1
Apply the Application to Firewall Policies	14-1
Apply the Application to QoS Rules	14-2
Tracking Based on Firewall Policy	14-3
Tracking Based on QoS Rule	14-3
VRF Firewall Enhancement	14-4
Easy First Install Simplified Appliance Installation	14-6
Configuration using Templates	14-7
WAN Link Templates	14-7
Basic Configuration Mode	14-8
Service Chaining	14-13

15 Release 6.1 P2 Features

Site Templates	15-1
Additional Features in Edge 6.1 GA P2	15-4

16 Release 7.0 Features

WAN Optimization	16-1
Zscaler Integration	16-3
Customer Edge (CE) Router Replacement Within the APN	16-7
E100 as an NCN	16-14
Capacity Report for the E100 NCN	16-16
NetFlow (Support for Version 9 and IPFIX)	16-16
Additional Features in 7.0 GA	16-17

17 Release 7.1 Features

E1000 Hardware Options	17-1
Interactive Dashboard	17-3
WAN Optimization on Virtual Appliances	17-6
WAN Optimization Reporting Enhancements	17-7
Additional Features in 7.1 GA	17-8

18 Release 7.2 Features

User Interface Enhancements	18-1
WAN Optimization Dashboard and Reporting Enhancements	18-5
Enhanced DHCP Relay	18-10
Client Private Subnet Reuse for Untrusted Segment	18-11
Palo Alto GlobalProtect Cloud Integration	18-11
Private Cloud Path Enhancement	18-12
Additional Features in 7.2 GA	18-13

19 Release 7.2 P3 Features

Configuration Versioning and Comparison	19-1
---	------

20 Release 7.3 Features

Enhanced Application Identification	20-1
E500 Appliance (7.3 GA P3)	20-3
Private Registration Server (7.3 GA P3)	20-3

Threshold Alerting (7.3 GA P4)	20-3
Additional Features in 7.3	20-5

21 Release 8.0 Features

22 Release 8.1 Features

23 Release 8.2 and later

24 Release 9.0 Features

Selective Software Update	24-1
Create and Update Plan	24-1
Monitor Plan Progress	24-2
Microsoft Azure Virtual WAN	24-3
Prerequisites	24-3
Configuring Azure Virtual WAN	24-3
Create a Virtual WAN	24-3
Create a Hub	24-4
Connect a Virtual Network to the Hub	24-4
Create an Application ID and Secret Key	24-5
Create an Azure Configuration	24-5
Add Service to the Site	24-6
SD-WAN Edge Configuration	24-6
Audits	24-8
Configuration Modes	24-8
All Sites	24-8
Global	24-43

25 Release 9.0M1 Features

Wake on LAN	25-1
Configuring Directed Broadcast for Wake on LAN	25-1

About This Guide

The purpose of this document is to describe features for all incremental releases of Oracle SD-WAN Edge.

Documentation Set

This table lists related documentation.

Document Name	Document Description
Oracle SD-WAN Edge Release Notes	Contains information about added features, resolved issues, requirements for use, and known issues in the latest Oracle SD-WAN Edge release.
Oracle SD-WAN OS Release Notes and Upgrade Guide	Contains information about inserting an OS Partition Image or OS Patch on an appliance in order to migrate to a new OS version or apply fixes to an existing version.
Oracle SD-WAN Security Guide	Contains information about security methods within the Oracle SD-WAN solution.
Oracle SD-WAN Edge Features Guide	Contains feature descriptions and procedures for all incremental releases of Oracle SD-WAN Edge. This guide is organized by release version.
Oracle SD-WAN Edge High Availability Guide	Contains information about implementing High Availability, as well as deployments and configuration.
Oracle SD-WAN Edge Virtual Appliance Installation Guide	Contains information about how to install a Virtual Appliance on a supported hypervisor.
Oracle SD-WAN Edge Configuration File Reference	Contains information about the structure, language, and defaults of the Oracle SD-WAN Configuration File in detail.

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.

Revision History

This section provides a revision history for this document

Date	Description
May 2020	• Initial release
June 2020	• Post-GA work
October 2020	• Updates for the 9.0M1 release, including addition of "Wake on LAN" section.
June 2021	• Adding additional steps to "Create an Application ID and Secret Key."

1

Release 2.3 Features

This chapter includes features and enhancements released in 2.3.

Geographic Redundancy

Currently, an Oracle Adaptive Private Network supports the concept of a single Network Control Node (NCN), which can be deployed in a High Available (HA) configuration. This allows for local redundancy, meaning that both appliances are deployed locally. With Edge g2.3 and the new Geographic Redundancy feature, the NCN and secondary NCN will not reside at the same location; they will reside in two separate data centers or locations. Typically, the second site would be some form of a disaster recovery facility. In the event of a primary data center failure, the backup data center should be operational and the secondary Oracle appliance would act as the NCN for Oracle SD-WAN Edge. There are a number of considerations to be aware of with this design:

- Oracle SD-WAN Edge supports a primary and secondary NCN
- HA is supported at primary and secondary NCN sites
- A secondary appliance will function as a client appliance when configured for Geographic Redundancy
- The active NCN is the clock source for the Oracle SD-WAN Edge
- The active NCN is the administration point for the Oracle SD-WAN Edge
- The active NCN will synchronize its database with the secondary NCN
- All client sites **MUST** have a conduit to the active and secondary NCNs
- Extra precautions must be taken when configuring routes
- The secondary NCN site should have static IP's for public Internet links
- If WAN-to-WAN forwarding is enabled on the Geographic Redundancy Oracle, the route cost will be the same for both NCN and Geographic Redundant NCN appliances. This can impact Oracle SD-WAN Edge routes and should be reviewed.

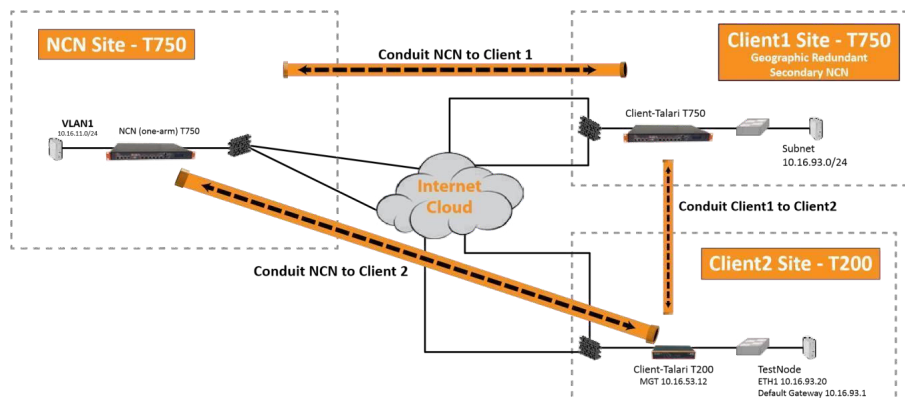


Figure 1

In Figure 1, we have the NCN site and two client sites, Client1 and Client2. Client1 is the Geographic Redundancy NCN site. As depicted in Figure 1, there must be a conduit between all client sites and the NCN site, as well as the Secondary NCN site. With this design, there are a number of recommendations:

- Design the redundant NCN site first and then design the routes
- Plan the Oracle SD-WAN Edge configuration file before deployment using the Oracle SD-WAN Edge configuration editor
- There must be enough WAN capacity for the required conduits
- Geographic NCN appliance requirements support hardware T3000 or T750 platforms only
- Be aware of the number of Oracle clients required, this will dictate NCN hardware requirements
- Be aware of failover times if local HA is deployed
- Since multiple conduits are now built from client sites, UDP hole punching may not work properly on all firewalls or NAT devices (two conduits)

The commands required to enable the Geographic Redundancy capability are provided below.

With the Geographic, Redundant NCNs, a new command was required to differentiate Oracle appliances. The new command is “appliance_mode,” and is described in detail below. The NCN primary would be configured with the following options:

```
add appli                               model=t750
{
  ance name=NCN
  set appliance_properties
  secure_key=0xacbf1332
  enable_wan_to_wan_forwarding=yes
  appliance_mode=primary_ncn;
```

The Oracle Client configuration would be configured as:

```
add applia                               model=t750
{
  nce name=Client1
  set appliance_properties
  secure_key=0xacbf1332
  enable_wan_to_wan_forwarding=yes
  appliance_mode=secondary_ncn;
```

appliance_mode = Text

Specifies the appliance's role in the Oracle SD-WAN Edge. It can be set as “primary_ncn”, “secondary_ncn,” or “client.” The Primary NCN would be set to primary, the secondary NCN would be set to secondary and a traditional client would be set to client.

GRE Support

In Oracle SD-WAN Edge release 2.3, Oracle is adding support for Generic Routing Encapsulation (GRE) header inspection and GRE header compression. GRE header inspection allows the appliance to look inside the IP GRE header and determine the protocol that resides inside the frame (inner protocol). Based on the inner protocol, Oracle APN will apply any defined rule and classify traffic accordingly. This capability simplifies the Oracle deployment within an infrastructure where IP GRE tunneling is currently in use.

GRE Header Inspection Support

Oracle SD-WAN Edge release 2.3 supports the ability to inspect IP GRE frames, and apply the corresponding rule and class based on the inner IP datagram. This support is based on RFC 2784 (Generic Routing Encapsulation). This capability eliminates the complexities of configuring external DSCP or ToS reflection for IP datagrams at the endpoint of the

IP GRE tunnel. Eliminating these complexities simplifies the Oracle's implementation process and requires user to only define the appropriate rules and classes in the configuration file. Certain WAN optimization devices also use IP GRE to encapsulate their traffic. This allows the APNA to identify certain WAN optimization flows that utilize the GRE encapsulation and classifying that traffic type. There are a few design considerations when implementing this feature:

- GRE uses IP protocol 47
- Does not support (Oracle) TCP termination
- Available on all Oracle appliances
- Supports only the inner IP datagram
- Supports GRE Header checksums

The GRE header inspection is enabled by default. The user must define the rules to map the inner protocol to a specific class. To monitor this capability the user should be aware of any IP GRE that exists in their infrastructure. The appliance will automatically identify these flows and display them in the flow table. Based on the Inner protocol, the appliance may apply any configured rule or the flow may default to an existing default rule. Figure 2 illustrates the flow page of an appliance where a GRE flow has been identified.

Figure 2 illustrates a flow that is encapsulated in a GRE tunnel. From the flow displayed, the IP Protocol (IPP) field is defined as GRE/TCP. This indicates that this is a GRE encapsulated flow and an inner IP field (protocol 1, ICMP). If the inner IP protocol was telnet, for example, the system would display GRE/TCP, and 23 would be under the destination port column. Now that the application is known, an existing rule for TCP telnet can be applied. This defined rule would classify the flow as an interactive flow, which would map to the Interactive class (Class 11).

In previous releases, the user would only see IPP 47. Any rules used to classify the traffic would require the traffic to be marked by either ToS or DSCP. The GRE header inspection now simplifies the deployment of the Oracle appliance when GRE tunneling is used within an

infrastructure. This eliminates the need for marking the GRE frames based on the inner IP protocol.

GRE Header Compression

Oracle SD-WAN Edge release 2.3 provides GRE header compression support, allowing for less additional overhead per packet when the customer uses the GRE protocol through the Oracle SD-WAN Edge. When compression is enabled, the GRE header compression will reduce the GRE overhead of packets from 24 bytes per packet to 7 bytes per packet. GRE header compression is performed by default, and is supported with and without GRE checksums.

Oracle Route Support

Oracle has provided enhancements to the existing route support within the Oracle SD-WAN Edge. This has been expanded to include the following enhancements:

- Route learning via SNMPv2
- Multiple Intranet defined services and gateways
- Intranet/Internet Route Failover based on path state

Each of these new features will be described in detail below with a focus on Route learning via SNMPv2.

Multiple Intranet Services Defined

In prior releases (2.2 and prior) Oracle only supported a single defined Intranet service. This is a WAN service the APNA could be configured to use if traffic was not conduit based. The Intranet services were defined on a WAN link to be either the primary WAN link or the secondary WAN link. The new enhancement now allows a user the flexibility to configure up to 32 Intranet services.

Each service may have a primary and a secondary defined WAN link. When adding routes to the Oracle SD-WAN Edge configuration file, the user can assign a route to a defined Intranet service if multiple Intranet services exist. A single WAN link can have multiple Intranet services defined. The commands required for this capability are provided below:

Service definition would use the following commands:

```
add intranet_service name=Intranet-1
{
}
add intranet_service name=Intranet-0
{
}
}
```

Adding the service to a specific wan link would use the following commands:

```
add net_usage      intranet_service_name=Intranet-0
                  service_type=intranet
                  wan_egress_rate_pct=10
                  wan_ingress_rate_pct=10;
```

If a route was added for the service, it would appear like the following:

```
add route         net=192.168.80.0/24
                  intranet_service_name=Intranet-0
                  cost=6 service=INTRANET;
```

These commands could be repeated for the individual routes and services as required.

Route Learning via SNMP

In Oracle SD-WAN Edge release 2.3, the APNA allows a user to define routers so that they can be polled for routes using SNMPv2. Once the routes have been learned by the APNA, the user can define rules which will include or exclude the routes from the

APNA route table. These routes will then be advertised or propagated to other APNAs within the Oracle SD-WAN Edge. Additional capabilities include the ability to continue polling the router for routes and, if a route is removed from the routing table of the router, propagate the topology change across the Oracle SD-WAN Edge. The polling intervals supported are “poll now,” “every 30 seconds,” “every minute,” or “every 5 minutes.” If the router that is probed is not reachable, the APNA can also be configured to purge the learned routes or maintain them. When using this capability, care must be taken when adding the routes. The user must define the routes properly to avoid any routing loop or problems. Routes included must be assigned to the correct Oracle SD-WAN Edge service; if Intranet(Internet) service is selected it must match the service defined in the configuration. This capability has the following design considerations:

- Currently only support for SNMPv2 is provided
- Probes the interface using the MIB: RFC 2096 IP Forwarding Table MIB and RFC1213-MIB
- Uses specific intervals to poll the router
- Router must support (configured) the MIB defined above (some routers do not)
- Can purge routes if router is not reachable
- Can include or exclude routes as required, excluded by default
- When defining the include/exclude rule, a user must assign a route to the correct service – local, intranet etc.
- Routes are local to each appliance, so the process is performed per appliance
- If no community string is defined, “public” is used
- Static routes learned from a polled router are displayed with the unknown interface
- Community string only supports alphanumeric characters
- Only Local routes are propagated via the Oracle SD-WAN Edge
- The polling takes place through the APNA management interface

- All straddle segments are added to the Oracle's route table by default

Log into the web console of the appliance and proceed to **Manage Network -> SNMP Route Learning**. Shown in Figure 3 are four sections to the web page: **Configuration, Include/Exclude Rules, Included Routes, Excluded Routes**.

SNMPv2 Route Polling Configuration

The first step is to complete the configuration section of the web page. First, define the router and community string under the configurations section of the page. The configuration section allows the user to perform the following:

- Add a router to poll
- Include the Community string
- Purge routes if the router is unreachable
- Polling time frame
- If Propagate Routes is set to YES, add routes to Oracle SD-WAN Edge route table
- Add multiple routes if required

The router that is typically added would be the LAN-side router. Subnets learned from this router would typically be local routes from a Oracle appliance perspective. The user must know which subnets are local subnets and which subnets are Intranet or Internet routes.

Include/Exclude rules

The user must then create a rule set which defines rules that would either include routes or exclude routes from the APNA route table. Once the rule is defined the user would hit the apply button to filter the corresponding routes and add them to the include route table. These routes could then be propagated to the Oracle SD-WAN Edge routing table, if desired. To propagate the routes, the user would have to have the option "Propagated routes" option set to "yes". Defining include routes requires knowledge of the network infrastructure. Any route included must be assigned to the correct service for proper Oracle SD-WAN Edge routing.

When adding routes to include in the Oracle SD-WAN Edge routing table, there is significance to the order of the defined rule. The rules are processed in a top-down method. With that in mind, the user must be aware of the defined rules for including/excluding routes. The preferred procedure would be to have more general rules defined first with more specific rules defined later in the rules list.

Included Routes

The Included Routes section displays the included routes. These are routes that will be propagated across the Oracle SD-WAN Edge from the local appliance. The local appliance will propagate these routes to all other appliances with which it maintains a conduit. These routes are then reachable from the conduit.

Excluded Routes

Excluded Routes are routes not included in the Oracle SD-WAN Edge route table, and are not propagated to other APNAs. By default, all routes learned from SNMP are

Excluded Routes. The user must define Include Rules to add any Learned Route to the Oracle SD-WAN Edge route table.

Intranet & Internet Enhancements

Intranet/Internet Route Failover Based on Path State

In previous releases when there was a WAN link failure of the Intranet service, the appliance would forward traffic based on the WAN link defined for the Intranet service. There was no searching through the route table to determine if there was an alternate route to reach the destination. In release 2.3, an added enhancement now allows the appliance to continue searching through the Oracle SD-WAN Edge route table for a second route in the event of a WAN link failure.

For example, if there are two WAN links, WAN Link A and WAN link B, WAN link A is the MPLS circuit/WAN link while WAN link B is the Internet circuit. If the Intranet service is defined for WAN link A and that router becomes unreachable, the appliance will now continue searching the Oracle SD-WAN Edge route table for an alternative route when there is a new flow. If the user defined a second route to be a conduit route, the flow could traverse the conduit in the event of an Intranet WAN link failure. This provides a backup route in the event of a WAN link failure.

Intranet Name Support

In Oracle SD-WAN Edge release 2.2 and prior there is the concept of a single Intranet service. This single service was defined and all Intranet routes used this defined WAN link and gateway. There was also the ability to have a primary and secondary WAN link defined for the Intranet service. In Edge 2.3 the number of Intranet services has been increased to 32, allowing a user to define up to 32 separate Intranet services. The user would define separate services in the case where they had multiple MPLS routers each supporting specific subnets. The user could then define the separate Intranet service and point a specific route/subnet to the corresponding service. This allows a much more flexible solution when supporting Intranet traffic.

Rule and Class Improvements

In Oracle SD-WAN Edge release 2.2 and prior releases, the typical Oracle configuration file supported the concept of a single default class. This default class was designed to support all different traffic flows without a specific advantage to a certain traffic type – Real-time verses Interactive Verses Bulk, for example. Any traffic type that did not have a user defined rule would match the default class 9. The default rules have now been enhanced to support different rules and classes for the most common traffic flows seen on networks today. Before a user implements the Oracle SD-WAN Edge, they must review the current rule set to understand what is enabled by default. The rules and classes are not all encompassing, but do cover a wide range of application and traffic flows. If the user has a custom application that is critical to the success of the deployment, they should define the application characteristics to a Oracle representative. The representative can then discuss the options with the user to define the correct rule set (class) for the customer application.

Shown below are the new default rules and classes as of Edge release 2.3:

Default Classes:

Class 0-9:

- User settable class
- Default: Bulk class
- Default: 1% share

Class 10 (*udp_ef_realtime_class*):

- Default class for user-defined UDP rules
- Realtime class
- 50% share

Class 11 (*control_tcp_ack_af11_int_class*):

- Default class for TCP Standalone ACK traffic.
- Interactive class
- 50% share

Class 12 (*ssh_telnet_interactive_class*):

- Interactive class
- 30% share

Class 13 (*gre_tcp_other_interactive_class*):

- Interactive class
- 20% share

Class 14 (*http_https_interactive_class*):

- Interactive class
- 10% share

Class 15 (*cifs_bulk_class*):

- Bulk class
- 45% share

Class 16 (*ftp_bulk_class*):

- Bulk class 45% share

Default Rules per Conduit:

- ICMP

(Assigned to class 11)

1. protocol_str=ICMP
2. class_name=control_tcp_ack_af11_int_class
3. transmit_mode=PERSISTENT_PATH
4. resequence_packets=YES

5. resequence_holdtime_ms=set rule_default
6. nontcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350
8. class_tail_drop_small_packet_bytes=30000
 - SSH

(Assigned to class 12)

1. protocol_str=SSH
2. class_name=ssh_telnet_interactive_class
3. transmit_mode=LOAD_BALANCE_PATHS
4. retransmit_lost_packets=YES
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350
8. class_tail_drop_small_packet_bytes=65000
9. reassign_flow_if_packet_exceeds_size_bytes=512 // for SCP
10. reassign_flow_if_packet_exceeds_size_class_name=ftp_bulk_class // for SCP
11. reassign_class_tail_drop_small_packet_bytes=(*-1/2 second based on WAN ingress bandwidth for the conduit*)
 - Telnet

(Assigned to class 12)

1. protocol_str=TELNET
2. class_name=ssh_telnet_interactive_class
3. transmit_mode=LOAD_BALANCE_PATHS
4. retransmit_lost_packets=YES
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350
8. class_tail_drop_small_packet_bytes=65000
 - HTTP

(Assigned to class 14)

1. protocol_str=HTTP
2. class_name=http_https_interactive_class
3. transmit_mode=LOAD_BALANCE_PATHS
4. retransmit_lost_packets=YES
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350

8. class_tail_drop_small_packet_bytes=100000
 - HTTPS

(Assigned to class 14)

1. protocol_str=HTTPS
2. class_name=http_https_interactive_class
3. transmit_mode=LOAD_BALANCE_PATHS
4. retransmit_lost_packets=YES
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350
8. class_tail_drop_small_packet_bytes=100000
 - CIFS

(Assigned to class 15)

1. protocol_str=CIFS
2. class_name=cifs_bulk_class
3. tcp_standalone_ack_class_name=control_tcp_ack_af11_int_class
4. tcp_standalone_ack_class_tail_drop_small_packet_ms=350
5. tcp_standalone_ack_class_tail_drop_small_packet_bytes=30000
6. transmit_mode=LOAD_BALANCE_PATHS
7. retransmit_lost_packets=YES
8. resequence_packets=YES
9. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
10. class_tail_drop_small_packet_bytes=(~2 seconds based on WAN ingress bandwidth for the conduit)
 - FTP

(Assigned to class 16)

1. protocol_str=FTP
2. class_name=ftp_bulk_class
3. tcp_standalone_ack_class_name=control_tcp_ack_af11_int_class
4. tcp_standalone_ack_class_tail_drop_small_packet_ms=350
5. tcp_standalone_ack_class_tail_drop_small_packet_bytes=30000
6. transmit_mode=LOAD_BALANCE_PATHS
7. retransmit_lost_packets=YES
8. resequence_packets=YES
9. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
10. class_tail_drop_small_packet_bytes=(~2 seconds based on WAN ingress bandwidth for the conduit)

- GRE_EF

(Assigned to class 10)

1. protocol_str=GRE
2. dscp_tag=ef
3. class_name=udp_ef_realtime_class
4. gre_header_compression_enabled=YES
5. transmit_mode=LOAD_BALANCE_PATHS
6. retransmit_lost_packets=YES
7. resequence_packets=YES
8. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
9. class_tail_drop_small_packet_ms=100
10. class_tail_drop_small_packet_bytes=15000

- GRE_AF11

(Assigned to class 11)

1. protocol_str=GRE
2. dscp_tag=af11
3. class_name=control_tcp_ack_af11_int_class
4. gre_header_compression_enabled=YES
5. transmit_mode=LOAD_BALANCE_PATHS
6. retransmit_lost_packets=YES
7. resequence_packets=YES
8. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
9. class_tail_drop_small_packet_ms=350
10. class_tail_drop_small_packet_bytes=65000

- GRE

(Assigned to class 13)

1. protocol_str=GRE
2. class_name=gre_tcp_other_interactive_class
3. gre_header_compression_enabled=YES
4. transmit_mode=LOAD_BALANCE_PATHS
5. retransmit_lost_packets=YES
6. resequence_packets=YES
7. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
8. class_tail_drop_small_packet_ms=350
9. class_tail_drop_small_packet_bytes=200000

- EF

(Assigned to class 10)

1. protocol_str=*
2. dscp_tag=ef
3. class_name=udp_ef_realtime_class
4. transmit_mode=DUPLICATE_PATHS
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=100
8. class_tail_drop_small_packet_bytes=15000
 - AF11

(Assigned to class 11)

1. protocol_str=*
2. dscp_tag=af11
3. class_name=control_tcp_ack_af11_int_class
4. transmit_mode=PERSISTENT_PATH
5. resequence_packets=YES
6. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
7. class_tail_drop_small_packet_ms=350
8. class_tail_drop_small_packet_bytes=30000
 - UDP

(Assigned to class 10)

1. protocol_str=UDP
2. class_name=udp_ef_realtime_class
3. transmit_mode=PERSISTENT_PATH
4. resequence_packets=YES
5. resequence_holdtime_ms=set rule_default nontcp_resequence_holdtime_ms
6. class_tail_drop_small_packet_ms=100
7. class_tail_drop_small_packet_bytes=15000
 - TCP

(Assigned to class 13)

1. protocol_str=TCP
2. class_name=gre_tcp_other_interactive_class
3. tcp_standalone_ack_class_name=control_tcp_ack_af11_int_class
4. tcp_standalone_ack_class_tail_drop_small_packet_ms=350
5. tcp_standalone_ack_class_tail_drop_small_packet_bytes=30000
6. transmit_mode=LOAD_BALANCE_PATHS

7. retransmit_lost_packets=YES
8. resequence_packets=YES
9. resequence_holdtime_ms=set rule_default tcp_resequence_holdtime_ms
10. class_tail_drop_small_packet_ms=350
11. class_tail_drop_small_packet_bytes=300000

- Other

(Assigned to class 13)

1. protocol_str=*
2. class_name=gre_tcp_other_interactive_class
3. transmit_mode=PERSISTENT_PATH
4. resequence_packets=NO
5. class_tail_drop_small_packet_ms=350
6. class_tail_drop_small_packet_bytes=200000

When the upgrade is performed, the Oracle SD-WAN Edge editor and compiler will automatically assign a percentage of the available bandwidth to each class. This will eliminate any potential issues when upgrading if there are many rules defined in an existing configuration file.

WAN Link and Path Enhancements

In Oracle SD-WAN Edge release 2.3, WAN link enhancements include a number of configurable options that were classified as “debug options” in previous releases, and were used at a number of customer installations. These options are now easily accessible by all customers, and include reserving minimum bandwidth for a conduit link, and congestion control per WAN link. These options may be configured in the Oracle SD-WAN Edge configuration editor, or as specific configuration options. These will be defined in more detail below.

Reserve Minimum Bandwidth for Conduit WAN Links

Certain network conditions cause congestion on a defined WAN link or path. When this congestion occurs, the appliance would reduce the amount of conduit data being forwarded on that WAN link. The appliance would reduce the usage rate down to its lowest defined rate (in certain cases) which was 80 kbps. When this occurred, there could be a performance impact to user data that was inside the conduit. For instance, 80 kbps would typically not even support a typical VoIP phone call. This unforeseen state could be caused by some external WAN-facing router that is oversubscribed with no configuration options of guarantying the appliance its defined usable rate. To overcome this issue, Oracle SD-WAN Edge release 2.3 allows the user to define a WAN link minimum rate - which is the lowest usage rate a WAN link will use during times of congestion. This will guarantee the appliance will continue to send no less than this defined value on the WAN link.

There are a number of design considerations to be aware of for this capability:

- Can be applied to a specific conduit on a WAN link only
- Must consider defined usage rates

- If congestion is consistent, the user should resolve the congestion issue on the router or firewall

An example of defining this minimum usage rate:

```
add virtual_wan_link name=Client-1-Link2
{
set properties
-----
virtual_interface_name=CL1VL11
virtual_ip_addr=10.1.20.12
gw_ip_addr=10.1.20.5
wan_ingress_physical_rate_kbps=3000
wan_egress_physical_rate_kbps=3000
wan_ingress_permitted_rate_kbps=3000
wan_egress_permitted_rate_kbps=3000
enable_public_ip_learning=true
tracking_ip_addr=10.1.20.15; add conduit_usage
remote_site_name=NCN-Site
wan_egress_rate_pct=100.0
wan_ingress_rate_pct=100.0
} minimum_reserved_bandwidth_kbps=400;
-----
```

minimum_reserved_bandwidth_kbps = Number (80)

The minimum amount of bandwidth that this usage will be reduced to during on demand scheduling. The default value for conduit links is shown above.

As we can see from the above configuration example, this option is applied under the WAN link and conduit usage. Once the specific WAN link is defined, the user can proceed to the conduit usage section where the minimum bandwidth option can be defined.

Configurable Congestion Control per WAN Link

The congestion threshold defines the period of time the appliance determines the WAN link is congested – if there is congestion for the defined period, it will then proceed into a congestion avoidance state. This is accomplished by reducing the amount of data sent on a WAN link. Sending less conduit data on a WAN link should result in the probability of the congested state clearing itself. The default value for this configuration option is 20 ms. (20000- value – default).

If the user chooses to change this value, it is recommended they contact a Oracle representative for specific values to use. If they decide to change the value to avoid the appliance from acting on a congestion state, the user would increase the value for example, if the value is increased to 2000000 – this would instruct the appliance that congestion must occur for a period of 2 seconds before the congestion avoidance algorithm is invoked. In traditional networks, 2 seconds of congestion would not happen. Since congestion is not detected the congestion avoidance algorithm is not enabled and the appliance continues sending data up to the defined usable rate.

```
add virtual_wan_link name=Client-1-Link2
```

```
{
  set properties
  virtual_interface_name=CL1VL11
  virtual_ip_addr=10.1.20.12
  gw_ip_addr=10.1.20.5
  wan_ingress_physical_rate_kbps=3000
  wan_egress_physical_rate_kbps=3000
  wan_ingress_permitted_rate_kbps=3000
  wan_egress_permitted_rate_kbps=3000
  enable_public_ip_learning=true
  congestion_threshold_us_per_s_us=200000 tracking_ip_addr=10.1.20.15; .
  .
}
```

With the congestion avoidance algorithm disabled, the Oracle will send data up to the configured usage rate. There could be a negative impact to other data flows within the network infrastructure depending on bandwidth allocation. The Oracle assumes that the defined usage rate is a guaranteed rate and attempts to send data at that defined rate as required. Users must consider this when using this new parameter.

Path Eligible Setting for Traffic Types

The Path Eligible/Ineligible options allow users to specify a certain traffic class to be eligible or ineligible for a specific WAN link. By default, all WAN links are eligible for all traffic classes. When a WAN link is ineligible for a traffic class internally, the appliance will add 150 ms latency to the defined WAN link. Depending on the WAN links in the configuration, this may only reduce the use of that WAN link under normal conditions. When traffic is queued up, the WAN link could be used by the appliance for the ineligible class if circumstances dictate.

The configuration parameters are shown below:

```
set properties
virtual_interface_name=CL4VL44
virtual_ip_addr=10.4.50.12
gw_ip_addr=10.4.50.5
wan_ingress_physical_rate_kbps=10000
wan_egress_physical_rate_kbps=10000
wan_ingress_permitted_rate_kbps=10000
wan_egress_permitted_rate_kbps=10000
wan_ingress_realtime_eligible=true
wan_ingress_interactive_eligible=true
```

```
wan_egress_realtime_eligible=true
wan_egress_interactive_eligible=true
wan_ingress_bulk_eligible=true

wan_egress_bulk_eligible=true;
```

This configuration example displays the default option. When reviewing a standard configuration with all paths eligible, these options will not appear in the configuration file.

They are displayed above so the user can understand what commands to use for these options. To verify the settings, the user can use the Oracle SD-WAN Edge QoS reports to view and verify that a specific WAN link is limiting a traffic class from that WAN link. Users should consult a Oracle representative if they believe they need this feature enabled for a specific WAN link.

Reporting Enhancements

Oracle SD-WAN Edge release 2.3 now has additional reports that assist the user in understanding the availability of the Oracle SD-WAN Edge and quantifying the value of the Oracle SD-WAN Edge. These new reports provide data on Oracle SD-WAN Edge availability, QoS usage, periodic network status, and a method to reduce outdated information using double triggers for event notification. The following table correlates the old reports with the new reports in 2.3.

Edge Release Reports	2.2	Edge Release 2.3 Reports
Reports		Performance Reports
Appliance Graphs		Appliance Reports
n/a		QoS Reports
n/a		Usage Reports
n/a		Availability Reports
n/a		Periodic Status Reports

Table 2

Availability Report

The Network availability report provides comprehensive data regarding uptime, goodtime, badtime, downtime and incidents per WAN link, conduit, and path. This data is available on a per site basis at the NCN and at any of the client sites. The client site report is based on that client's conduits and WAN links. Here is an example of a client site report:

The definition of availability report terms follows:

Term	Definition
Oracle SD-WAN Edge Object	A path, conduit or WAN Link.
Incidents	A counter for the number of periods of downtime.
Goodtime	The total amount of time that an object has been in a good state.
Badtime	The total amount of time that an object has been in a bad state.
Uptime	The total amount of time that an object has been in a state that is greater than dead.
Downtime	The total amount of time that an object has been in a dead state.

Table 3

The report displays up to 24 hours of data by default. Other display options are available including, 1 hour, 24 hours, 7 days, or "All Available Data." To access this report, login into the web console and select **Monitor** -> **Availability Reports** from the pull-down menu.

QoS Reports

In prior release of Edge, the user did not have insight into the application classes that traverse a specific conduit, WAN link or path. There were only generic counters that could be used which displayed traffic on a per class basis. This data was provided as counters for total traffic only. In the Edge release 2.3, the appliance now provides class statistics based one of the following options: WAN link, Conduit, Path, and Site. The Report will look like the

following:

From the above screen capture we can see the different classes of traffic on the path

"PPCBL - Colo-L3. From the screen capture we can see the applications classified as Realtime (blue), Interactive (green) and Bulk (red). Traditionally, Realtime is VoIP or time sensitive traffic, Interactive is http/https/telnet and Bulk is any type of file transfer application (ftp). From this report the user can determine which traffic class was using what portion of the bandwidth on a conduit or WAN link. To view this report log in into the web console of the appliance and proceed to **Monitor** -> **QoS Reports**. Insight into this type of data allows the network administrator to plan accordingly for WAN link and bandwidth usage, and potential WAN link expansion.

Usage Report

The Oracle SD-WAN Edge Usage report allows the user to display usage for services that include Conduit, Intranet, and Internet. By default, this report displays all conduit and services information in a stacked fashion. The user can then disable conduit statistics from being displayed and only view Intranet and Internet traffic, if desired. This data can be displayed for a site or a WAN link, ingress direction or egress direction. To view this report login into the web console and proceed to **Monitor** -> **Usage Reports**. The default graph for the Oracle SD-WAN Edge Usage report is shown below.



This screen capture is displaying usage for a specific WAN link (PP-CBL) in Oracle ingress direction (LAN to WAN). The different colors represent the conduits, as well as Internet and Intranet services provisioned on the PP-CBL WAN link. The user also has multiple options on time frame for data that is to be displayed, as well as the ability to display any available archived database. To view this report login into the web console of the appliance and proceed to **Monitor -> Usage Reports**

Periodic Status Reports

The Periodic Status reports are provided to the user automatically, once configured. These reports provide details regarding the status of the underlying network. The underlying network is considered to be the wide area network – MPLS cloud or ISP (internet) cloud, as seen from a Oracle appliance perspective. By default, the status report is not sent. To configure these reports, the user would login to the web console and proceed to **Integrate -> Periodic Status Reports** page. Configurable options include: details can be turned off or on for each individual Site, WAN link, Conduit, Path, Internet service, and intranet service.

To simplify the configuration process, the “Select All” option may be used, or just select individual items as appropriate. Once defined, the user can preview the report before the actual update is emailed. The email can be sent based on the user defined time criteria

(every day, specific days) as well as defining a specific time of day. There is also a “Send Now” option that can be used to test the capability. For this to work properly, the user must have the “Email Alerts” defined properly in the web console **Integrate**, and then **Configure Events and Alerts** section of the web console.

Once emailed, the report would look like Figure 7 below (depending on properties selected).

Double Event Triggers

In previous releases, the appliance would send a notification immediately after an event occurred under any circumstances. To enhance this capability and reduce the number of event notifications that users receive, a double event trigger capability has now been added to the appliance. This allows certain event emails to only be sent when they persist over a user-defined period of time. For example, a conduit is dead - and is dead for a certain period of time. Once the pre-defined time limit is reached and the conduit is still dead, the event notification would be sent to the defined email address. This will reduce the number of emails sent when a conduit goes down and comes back up immediately. Typically, path state change events generate more state changes which result in more email notifications. The double event trigger capability reduces the number of notifications dramatically while still allowing the user the ability to monitor any path state changes. This capability is disabled by default. The user has the following options when configuring this capability:

Supports Event Types: Timeframe in seconds:

SERVICE	2
CONDUIT	5
WANLINK	10
PATH	30
	60

To configure this capability, log in to the web console of the appliance and proceed to

Integrate -> Configure Events and Alerts. Scroll down to the "General Event Configuration" section and look for the section labeled "Alert if State Persists." By default, the user will see the behavior set to "Alert Immediately", this pull-down supports the timeframe options shown above. Once all options are configured select the "Apply Settings" to complete any changes.

Observed Protocols

In Oracle SD-WAN Edge release 2.3, Oracle now provides users with a list of protocols that are traversing the Oracle appliance. The protocols are displayed to assist the user in verifying the correct rule set is applied, as well as learn what protocols may not be IANA-based protocols that reside within their infrastructure. The following screen shot illustrates this data within the Oracle appliance.

Talari Statistics

Show: **Observed Protocols** Enable Auto Refresh 5 seconds Refresh Clear Table

Observed Protocol Statistics

Filter: in Any column Apply

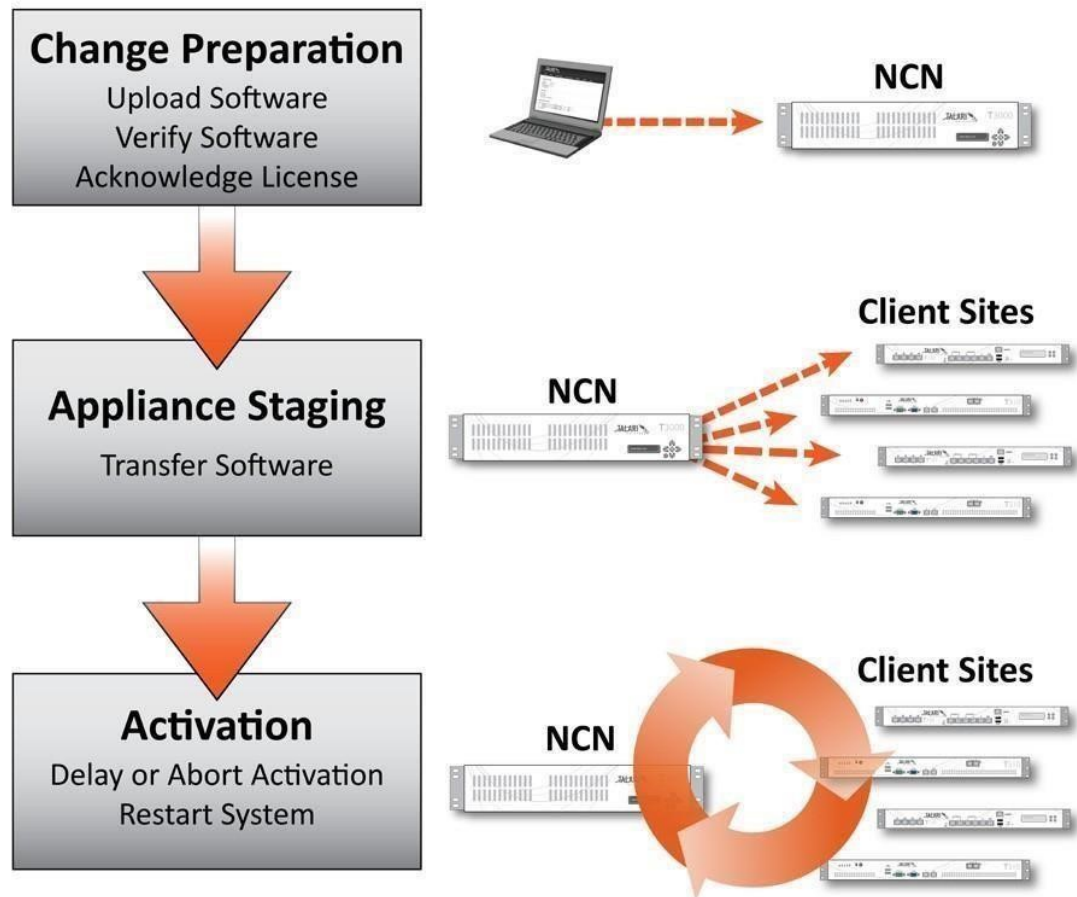
Application	Rule	Protocol	Port	Service Type	Service Instance	WAN Ingress Packets	WAN Ingress Bytes	WAN Ingress kbps	WAN Egress Packets	WAN Egress Bytes	WAN Egress kbps
ftp	(7)	TCP	21	CONDUIT	NCN-Colo-Client-PPark	2	120	0.00	0	0	0.00
ssh	(98)	TCP	22	CONDUIT	NCN Colo Client-PPark	0	0	0.00	2422799	121557652	0.00
ssh	(103)	TCP	22	CONDUIT	NCN-Colo-Client-PPark	2583476	3347498442	0.00	0	0	0.00
ssh	(110)	TCP	22	CONDUIT	NCN-Colo-Client-PPark	0	0	0.00	11183	407726	0.00
ssh	(111)	TCP	22	CONDUIT	NCN-Colo-Client-PPark	0	0	0.00	59332	3861518	0.00
ssh	(114)	TCP	22	CONDUIT	NCN-Colo-Client-PPark	557	7162680	0.00	0	0	0.00
ssh	(115)	TCP	22	CONDUIT	NCN-Colo-Client-PPark	347	790026	0.00	0	0	0.00
telnet	(123)	TCP	23	CONDUIT	NCN-Colo-Client-PPark	4	240	0.00	0	0	0.00
smtp	(122)	TCP	25	CONDUIT	NCN-Colo-Client-PPark	0	0	0.00	1824	284780	0.00
smtp	(123)	TCP	25	CONDUIT	NCN-Colo-Client-PPark	508	100028	0.00	0	0	0.00
tacacs	(122)	TCP	49	CONDUIT	NCN-Colo-Client-PPark	0	0	0.00	1164	66374	0.00
tacacs	(123)	TCP	49	CONDUIT	NCN-Colo-Client-PPark	85	29038	0.00	0	0	0.00
http/www/www/http	(122)	TCP	80	CONDUIT	NCN-Colo-Client-PPark	0	0	0.00	14560	9213620	0.00

To view the observed protocols within the web console of the appliance, log in and proceed to **Monitor-> Statistics**. From the pull-down list “**Show**” select the “Observed Protocols” option. Depending on the data flow through the Oracle, the table may take a few second to populate. The table will display known and unknown protocols. This data displayed will provide which rule the flow is matching on, as well as other data regarding the flow.

Enhanced Network Change Management

The enhanced network change management process allows a user to upload a new file package to the Oracle SD-WAN Edge. This new package can be a software update, or a configuration update, or both. The three-step workflow is a set of checks and processes that ensure that configuration changes and software updates are applied to the network in a reliable, fail-safe way. Go to **Manage Network -> Change Management** to access this utility.

The Change Management Workflow



In addition to using change management for network software distribution and activation, a twostep local change management utility has been introduced as well, allowing for an easier and more intuitive process for updating software and configuration files on individual appliances. Local change management is located at the **Manage Appliance -> Local Change Management** page.

2

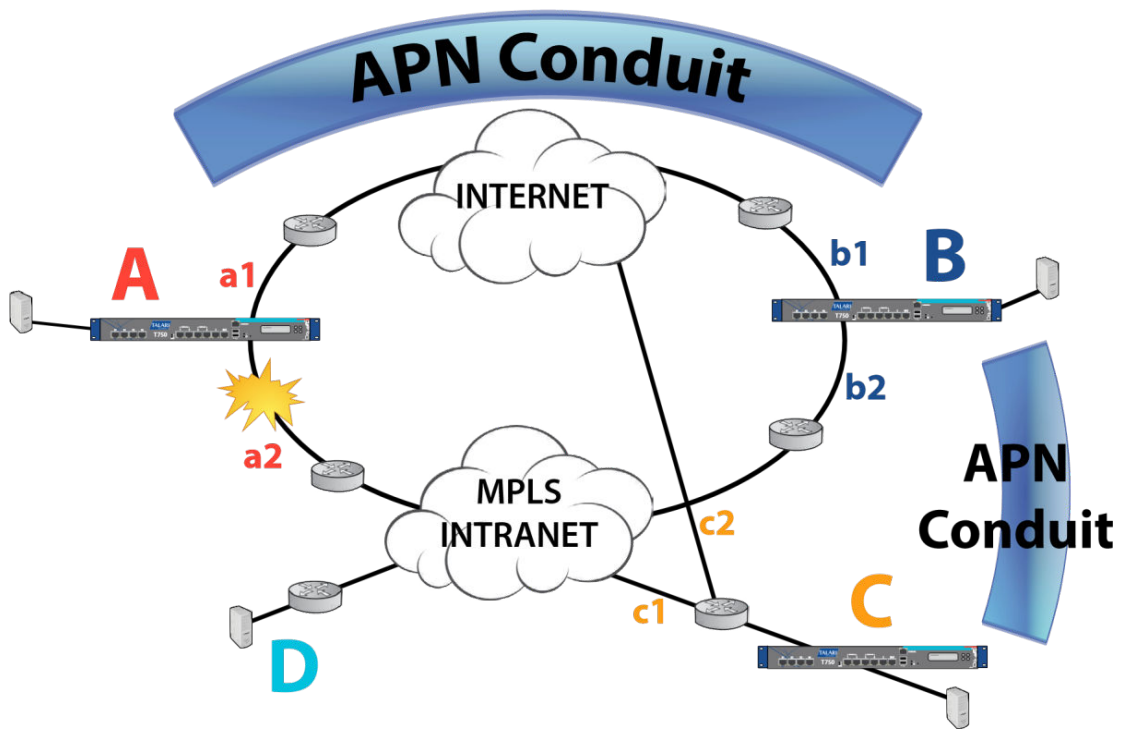
Release 2.4 Features

This chapter includes features and enhancements released in 2.4.

Network Functionality and Deployability

Intranet Route Eligibility Determined by Path State

A new feature has been added in Oracle SD-WAN Edge Release 2.4 to determine Intranet route eligibility by path state as opposed to WAN link state. In the past, a route was ineligible if the WAN link was down, but there were cases where the WAN link was only down at one site. This caused the Intranet data to not flow between sites. To resolve this issue the Oracle appliance now can make a path eligible/ineligible decision based on path state and not just WAN link state. The **“Route Eligible Path”** is a path whose status is used to determine whether a route is eligible to be used when making routing decisions.



As shown in Figure 1 above, there are four sites in a sample prior release network. There is a conduit from A to B and a conduit from B to C. WAN links a2 and b2 are also configured with intranet service. When a2 fails, all paths using a2 are dead, so a2 is in WAN link DEAD state. Intranet traffic going from site A to site B will then skip the intranet route and use the available conduit route. At site B, since b2 is used for a path for the B-C conduit, it remains in a GOOD state. Intranet traffic from site B to site A will continue to use the intranet route that enters the network on the b2 link and exits the network on the a2 link. This traffic will fail to reach the destination since a2 is down.

To solve this problem, this new feature allows a user to add an eligible path to the intranet route. When doing a route lookup, the intranet route is now skipped if the eligible path is DEAD.

One Route Eligible Path can be configured for each dynamic and/or user-configured static intranet route but there is no limit on the number of intranet routes that can be configured to use a Route Eligible Path. The configuration editor allows a path to be configured as a Route Eligible Path for user's static intranet routes. Adding, deleting, or changing the Route Eligible Path on an intranet route does not require a reset of the

Oracle SD-WAN Edge (please see Figure 2). For more information on the Oracle SD-WAN Edge Configuration Editor, please see the Oracle SD-WAN Edge Configuration Editor User's Guide.

Route Eligible Path can also be set by editing the configuration file itself. The command line options would include the following:

```
add route
net=10.0.0.0/8 intranet_service_name=Intranet-1
route_eligibility_based_on_path=true
route_eligibility_from_wan_link=a2
route_eligibility_to_wan_link=b2 service=INTRANET;
```

`route_eligibility_based_on_path = Boolean`

This feature allows a user to add an eligible path to the Intranet route. For Intranet Services, enable the Intranet route failover feature. Route eligibility will be based on the state of an associated path.

`route_eligibility_from_wan_link_name = Text` The “from” WAN link name for the path that determines whether to mark this route as ineligible, based on the state of the specified path.


`route_eligibility_to_wan_link_name = Text`

The “to” WAN link name for the path that determines whether to mark this route as ineligible, based on the state of the specified path.

As defined above, the configuration requires the user to specify the WAN links and to enable the feature. For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference*.

L2 MAC Learning for Multiport Bridging

Mac address learning allows a user to define three ports in a bridge group. This simplifies many deployments by allowing the user to connect an MPLS router and a Firewall directly to the appliance without requiring a switch. Figure 3 provides two examples. The new bridging capability in Oracle SD-WAN Edge R2.4 reduces the need for infrastructure changes when deploying Oracle SD-WAN Edge appliances.



Multiple LAN Routers in Same LAN Subnet Multiple LAN Routers in Different LAN Subnet

MAC address learning for multiport bridging stores the source MAC address of each received packet so that future packets destined for that address can be forwarded only to the port on which that address is located. Packets destined for unrecognized addresses are forwarded out of every port. There are no options to configure this command. When three or more ports are configured in an Interface group this feature is enabled.

Design Considerations

- Oracle does not support user spanning tree with this feature. The infrastructure must therefore be designed accordingly
- This feature is enabled by default

Port Switching

Some WAN service providers do not allow long duration UDP sessions and block them in the Cloud. To avoid such issues, the Oracle SD-WAN Edge R2.4 introduces a new feature allowing the user to specify an alternate UDP port for the Oracle SD-WAN Edge conduit packets. UDP Port Switching is a preventative measure to change the source UDP port at specified intervals. The Alternate UDP port number and the port switch interval are user settable. Port Switching can be set using the Oracle SD-WAN Edge Configuration Editor (Figure 4, below). For more information on the Oracle SD-WAN Edge Configuration Editor, please see the *APN Configuration Editor User's Guide*.

Port Switching can also be set by editing the configuration file itself. The command line options would include the following:

```
add conduit_usage

remote_site_name=NCN-Site wan_egress_rate_fair_share=800000
wan_ingress_rate_fair_share=800000 service_group_name=Default
udp_port_num=2156

udp_port_num_alt=2157

udp_port_switch_interval_minutes=1500;
```

udp_port_num =*Number* (2156)

This will be used as the source UDP port for all WAN ingress packets sent from this link. The Oracle SD-WAN Edge will also only accept WAN Egress packets at this link with dst_port set to this port number.

udp_port_num_alt =*Number* (2156)

This will be used as the alternate source UDP port for all WAN ingress packets sent from this link. The Oracle SD-WAN Edge will also only accept WAN Egress packets at this link with dst_port set to this port number, or the udp_port_num_alt value.

udp_port_switch_interval_minutes =*Number* (1440)

if udp_port_num and udp_port_num_alt are both set and are not equal)

Interval in minutes to be used when switching between the two values of udp_port_num and udp_port_num_alt. Allowed values are from 1 minute to 8640 minutes (6 days).


For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference* available from the Oracle support site.

Network Topology

Support for Multiple Routers on LAN-side Subnets

With Oracle SD-WAN Edge Release 2.4, the local routes selection process allows the user to configure the local route to be eligible only when the gateway is reachable. If the gateway is unreachable, the route will be skipped and the next available route will be selected to forward the traffic.

With previous releases, the selection process for local routes only allows for the selection of the route with the lowest cost. Even if this local route's gateway is unreachable, and there are other routes available to get to the local network, this local route will still be selected, causing the traffic to the local network to be dropped. The new feature supports multiple LAN routers in the same or different subnets as shown in Figures 5 and 6 below.



Multiple LAN Routers in Same LAN Subnet Multiple LAN Routers in Different LAN Subnet

Figure 5 Figure 6

This feature is disabled by default but can optionally be enabled once the routes are added, using the Oracle SD-WAN Edge Configuration Editor as shown in Figure 7 on the following page.

Multiple routers on LAN-side subnets can also be set by editing the configuration file itself. The command for this feature is:

```
add route
net=10.3.50.0/24
gw_ip_addr=10.3.10.65
cost=6
route_eligible_on_gw=true
service=LOCAL;
```

route_eligible_on_gw = Boolean (NO)

Enabling this option will cause a route to only be valid if the gateway specified in this route is reachable. This parameter is for use in local routes only.

For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference*.

Support for Serial High Availability Appliances

In Release 2.4, Oracle has enhanced its HA support by expanding it to support the Serial Inline HA Topology. This allows a simple Fail-to-Wire based serial HA option that simplifies deployments for end users. The Serial Inline HA feature also includes supporting additional

Fail-to-Block groups or nonHA traffic on the control segment. Figure 8 (below) shows an example of how Serial Inline HA could be configured.

Oracle Serial HA

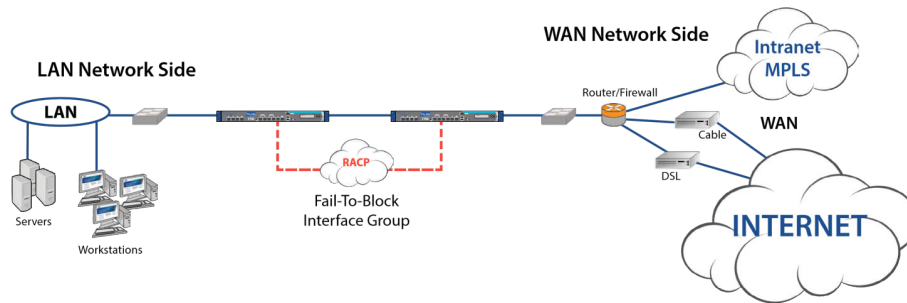


Figure 8

Serial Inline HA can be enabled in the Oracle SD-WAN Edge configuration file using the Oracle SD-WAN Edge

Configuration Editor GUI (shown in Figure 9 below). For more information on the Oracle SD-WAN Edge Configuration Editor, please see the *APN Configuration Editor User's Guide*.

The Serial Inline HA feature may also be set in the Oracle SD-WAN Edge configuration file manually, using the `use_serial_ha` parameter. The user would configure the HA in the standard method but select the option for “`use_serial_ha`.” This feature can also be added to the configuration manually. A sample of the configuration setting is shown below:

```
Define site name=NCN1
{
  add appliance name=ncn1
  {
    ...
  } add ha_appliance name=ncn1
  - HA; { add ha_service set properties
  primary_appliance_name=ncn1
  HA secondary_appliance_name=ncn1 primary_reclaim=false
  use_serial_ha=true; add interface_group
  {
    set interface_properties viprimary_ip_addr=10.40.10.13
    rtual_interface_name=VLAN1
    secondary_ip_addr=10.40.10.14; }
  }
}
```

For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference*.

Design Considerations:

- Keepalives must be configured on a Fail-To-Block (FTB) interface group
- FTB interface can be directly connected between appliances
- Supports untrusted interface groups for conduit traffic
- Link health not used for HA priority
- Must have a Fail to Wire bypass group defined for conduit traffic
- Standby appliance allows packet to pass through
- Spanning tree must be considered when deploying this topology
- Sends HA protocol across the FTB HA interface group only
- Configuration updates traverse the FTB HA interface group from active to backup appliance

Multiple VLAN Segment on Common WAN link

Oracle SD-WAN Edge Release 2.4 allows multiple VLANs and gateways to be configured on the same WAN link by supporting multiple Access Interfaces. Each “Access_Interface” will support a name, a VIP address, a gateway and the option to enable Proxy ARP. An example of these topologies is shown in Figure 10, below, where with the addition of this new feature Access Interfaces have been configured for VLANs g and b and Edge can perform Proxy ARP for routers g and b.

In previous Edge releases, the WAN link and WAN link gateway were on a single VLAN, causing problems for customers that segmented traffic (e.g. Intranet vs. Internet) on different VLANs for security reasons. Additionally, if Proxy ARP was enabled on a WAN link, it would only work properly for hosts that resided on the corresponding VLAN segment. If the WAN link was assigned to VLAN A and the gateway went down, the users on VLAN A would have connectivity but those on other VLAN would not.




Figure 10

This feature may be set in the Configuration Editor, as shown in Figure 11 below. Note that the first configured WAN link access interface is set to be the Primary by default and additional access interfaces are defined to be excluded. The user would then be expected to verify the primary access interface link, as well as any interfaces to be configured as secondary. When the primary access interface gateway is not reachable, the appliance would use the access interface configured as secondary.

For more information on the Oracle SD-WAN Edge Configuration Editor, please see the *APN Configuration Editor User's Guide*.

These interfaces are used to define which gateway the appliance would use to forward conduit frame for the corresponding WAN link. These options are defined under the WAN link, Set Properties field of a WAN link definition.

To edit the configuration manually, the options associated with this feature are included below:

```
add access_interface name=CL3_WL0_access_interface_1
```

	virtual_interface_name=CL3VL0 gw_ip_addr=10.3.10.2	virtual_ip_addr=10.3.10.12
set properties	enable_proxy_arp=true;	
	primary_conduit_access_interface=Cogent-NS_175Federal-AutoAI-0	
	wan_ingress_physical_rate_kbps=2000	
	wan_egress_physical_rate_kbps=2000	
	wan_ingress_permitted_rate_kbps=2000	
	wan_egress_permitted_rate_kbps=2000	

The Proxy ARP capability is configured under the Access Interface and is shown above. “True” would indicate Proxy ARP is enabled and “false” would indicate that the feature is disabled for the Access Interface. Proxy ARP needs to be enabled for local side subnets.

primary_conduit_access_interface = Text

The name of the access interface to be used as the primary access interface for this WAN Link. Mandatory.

secondary_conduit_access_interface = Text

The name of the access interface to be used as the secondary access interface for this WAN Link.

For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference*.

Design Considerations:

- Enable Proxy ARP for local side subnets

Path MTU Discovery

Path MTU discovery allows the sender of IP packets to discover the Maximum Transmission Unit (MTU) of packets that it is sending to a given destination. The MTU is the largest packet that can be sent through the network along a path without requiring fragmentation. Previous releases supported ICMP to adjust the conduit MTU.

To overcome these limitations, Oracle SD-WAN Edge Release 2.4 introduces a new path MTU discovery method that will actively probe each sending path of each conduit to find out the current MTU, and adjust the conduit MTU accordingly. This feature must be enabled by the user and will then probe each path within a conduit every 10 minutes.

This feature can be configured in the Oracle SD-WAN Edge Configuration Editor (See Figure 12, below). Additionally, the Oracle SD-WAN Edge Web Console and CLI will display the current MTU that is being used by each WAN Ingress path on the conduit (See Figure 13 for an example) For more information on the Oracle SD-WAN Edge Configuration Editor, please see the *APN Configuration Editor User's Guide*. For more information on using the Oracle SD-WAN Edge Web Console, please refer to the *APN Appliance Operation Guide*.

Figure 12. Active MTU Discovery in the Oracle SD-WAN Edge Configuration Editor.

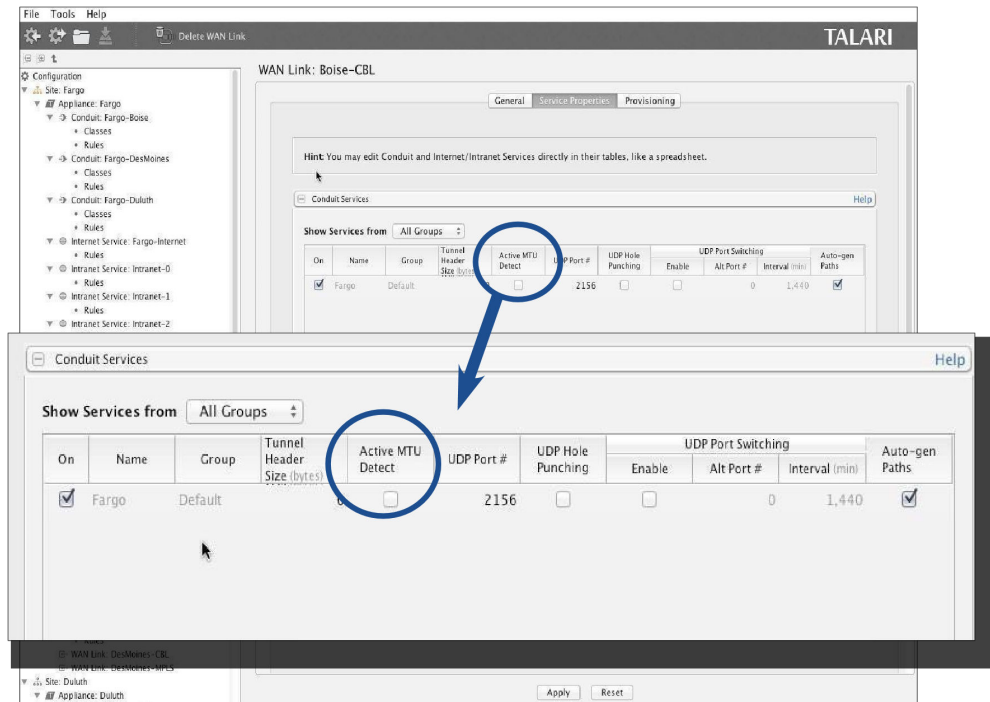


Figure 13. Active MTU Discovery Shown in the APN Web Console

From Link	To Link	Path State	Conduit State	Source Port	Destination Port	Discovered MTU	Latency BOWT	Statistical Jitter (mS)	Packets Received	Packets Out of Order	Packets Lost %	kbps
Colo-L3	wl-km	GOOD	GOOD	2156	2156	1488	12	6	165	0	0	7.3
Colo-TW	wl-km	GOOD	GOOD	2156	2156	1488	11	6	175	0	0	7.9
Colo-L3	CL-AD	GOOD	GOOD	2156	5917	1488	31	6	161	0	0	7.8
Colo-TW	CL-AD	GOOD	GOOD	2156	5917	1488	9	6	183	0	0	9.3
Colo-L3	CL-AG	GOOD	GOOD	2156	2156	1488	15	6	164	0	0	7.5
Colo-TW	CL-AG	GOOD	GOOD	2156	2156	1488	8	8	192	0	0	9.7

If editing the configuration file manually, the command options for this capability in the configuration file are enabled under the WAN link, conduit services:

```
remote_site_name=Client-test1
wan_egress_rate_fair_share=100000
add conduit_usage wan_ingress_rate_fair_share=100000
service_group_name=Default udp_port_num=2156
active_path_mtu_discovery_enable=true;
```

For more information on the Oracle SD-WAN Edge configuration file, please see the *APN Configuration Reference*.

Design Considerations:

- Currently if a lower MTU is detected, MTUs of all paths of the corresponding WAN link are reduced

Usability

Clone Site Configuration Wizard

To make it easier to add a new client site to a configuration, Oracle SD-WAN Edge Configuration Editor now allows the user to clone a pre-existing site's configuration as a new site.

To clone an existing site using the Configuration Editor, select the site that you would like to clone, and click the **Clone Site** button. This will bring up a dialog screen containing the cloned site information that must be changed before being allowed to save the new site information. See Figure 14 below. For more information on the Oracle SD-WAN Edge Configuration Editor, please see the *APN Configuration Editor User's Guide*.

The Site Cloning Process

A new cloned site must contain its own valid site configuration. In other words, the existing information from the site being cloned (presented on the screen in red) must be changed before the new site can be created. As you make the needed changes, the text will change from red to black. Any particular portion of the site being cloned that is unneeded in the new cloned site may be excluded by clicking on that line item and then clicking the **Exclude** button, and graying out the unneeded information. Please see Figure 15 below.

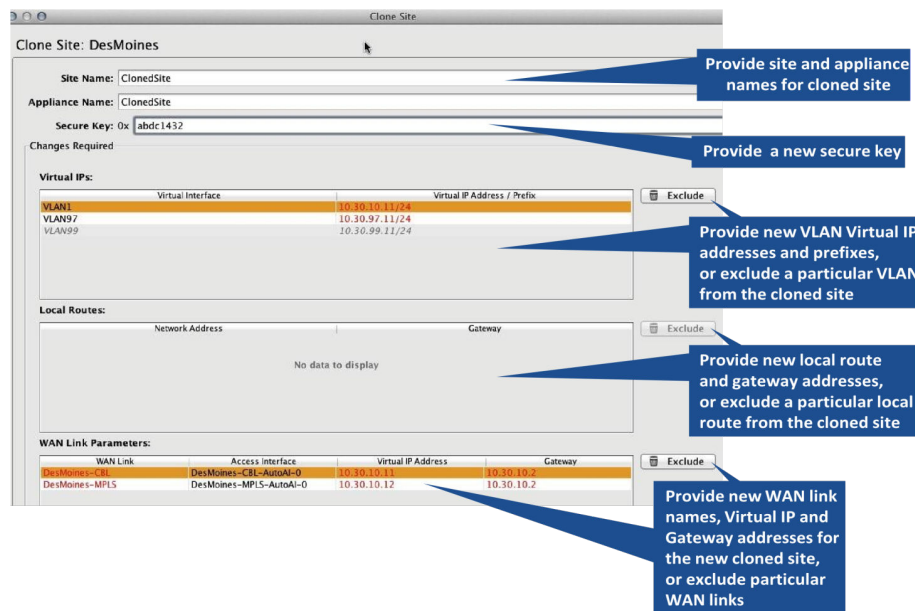


Figure 15

Once you have edited the screen making sure that all fields are valid and unique from their original values, the OK button will be enabled (Figure 16). Clicking the OK button will allow an audit of the new site to make sure the information will be valid for your configuration. If any errors are found, you will receive an error message pointing out that needs to be corrected.

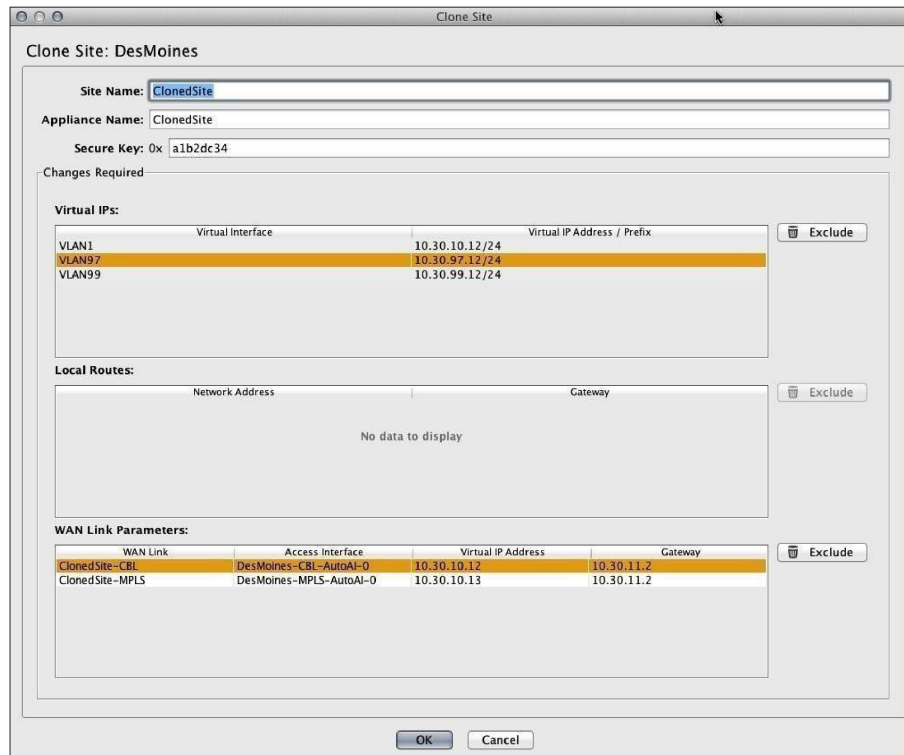


Figure 16. New site ready for audit.

If the audit succeeds, you will be presented with a screen detailing the next steps that need to be performed in making your new site functional in your network (see Figure 17).

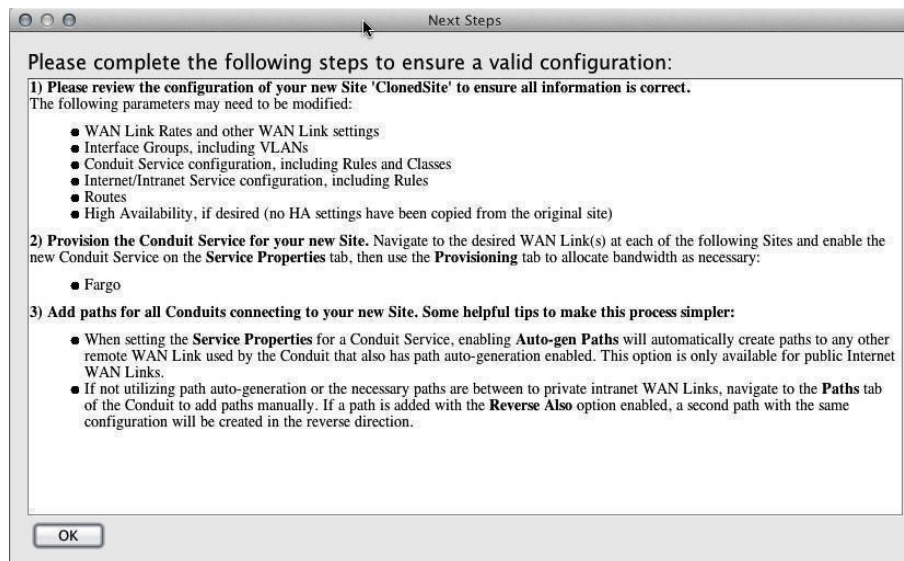


Figure 17: Next steps in the site cloning process message

Next Steps in the Site Cloning Process

Sample Next Steps for the newly created client site called “ClonedSite:”

1. **Please review the configuration of your new Site 'ClonedSite' to ensure all information is correct.**

The following parameters may need to be modified:

- WAN Link Rates and other WAN Link settings
 - Interface Groups, including VLANs
 - Conduit Service configuration, including Rules and Classes
 - Internet/Intranet Service configuration, including Rules
 - Routes
 - High Availability, if desired (no HA settings have been copied from the original site)
- Provision the Conduit Service for your new Site.** Navigate to the desired WAN Link(s) at each of the following Sites and enable the new Conduit Service on the **Service Properties** tab, then use the **Provisioning** tab to allocate bandwidth as necessary:
- Fargo
1. **Add paths for all Conduits connecting to your new Site.** Some helpful tips to make this process simpler:
 - When setting the **Service Properties** for a Conduit Service, enabling **Auto-gen Paths** will automatically create paths to any other remote WAN Link used by the Conduit that also has path auto-generation enabled. This option is only available for public Internet WAN Links.
 - If not utilizing path auto-generation or the necessary paths are between to private intranet WAN Links, navigate to the Paths tab of the Conduit to add paths manually. If a path is added with the **Reverse Also** option enabled, a second path with the same configuration will be created in the reverse direction.

More detailed information on the elements of these steps is available in the *APN Configuration Editor User's Guide*.

Improved Site WAN Link Provisioning

An Overview of Provisioning

Provisioning allows for the automatic bidirectional (Ingress/Egress) distribution of bandwidth for a

WAN link among the various services associated with that WAN link. Using the Oracle SD-WAN Edge

Configuration Editor, the user can enter the values in the text fields and directly into the cells of the table like a spreadsheet. The Provisioning page is shown below in Figure 18. There are three steps to Provisioning that provide for this bandwidth distribution in a simple and effective way:

1. WAN Link Rates (Setting the WAN link physical and permitted rates)
2. Provisioning Groups (Create and edit groups of shares of bandwidth)
3. Services (View and edit services for groups or individual site WAN links)

With Oracle SD-WAN Edge Release 2.4, we introduce the concept of Fair Shares to the provisioning process. Shares are used to distribute the permitted bandwidth between the provisioning groups. The bandwidth calculated is based on the shares allocated for a

particular group, divided by the total shares for all groups. A separate pool of shares is used for both Ingress and Egress traffic.

This area allows the user to set both the WAN link

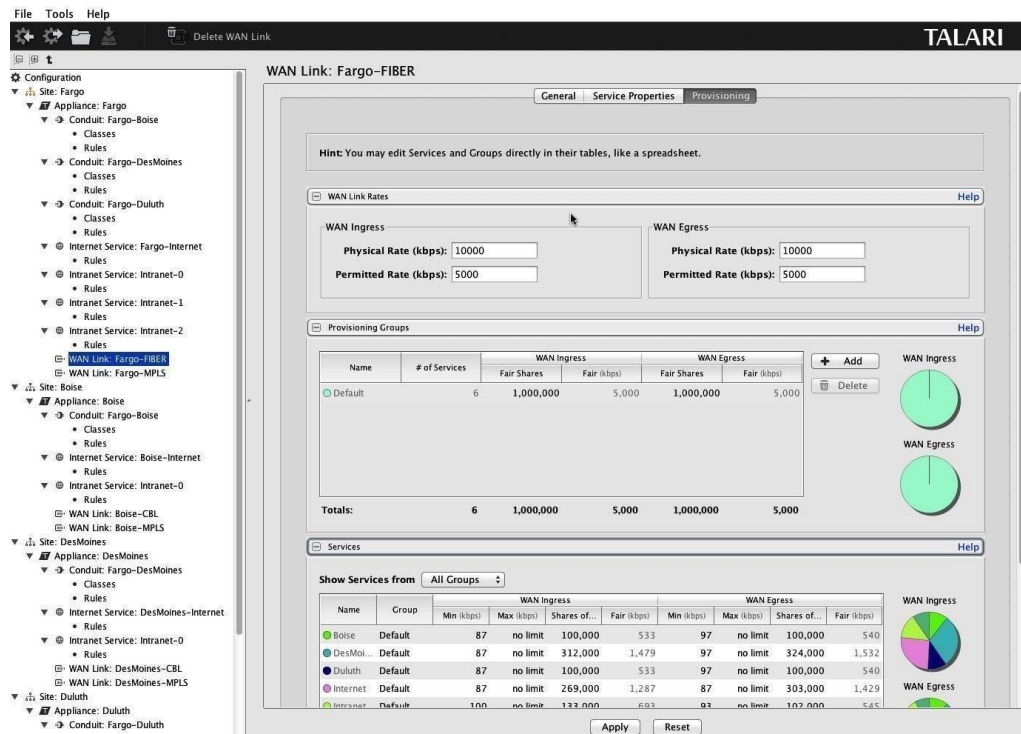


Figure 18

WAN Link Rates

Physical Rate (the raw bit rate for the incoming/ outgoing traffic), and the WAN link Permitted Rate (the available rate for incoming/outgoing traffic). Please see Figure 19 on the next page.

Provisioning Groups

A Provisioning Group contains a collection of WAN Link bandwidth usages for any given WAN Link. This allows the user to allocate and distribute the shares of bandwidth among a smaller set of services at a high level before drilling down to the individual services for finetuning. They also provide a boundary for the automatic redistribution of bandwidth within the child Services of the Provisioning Group.

In the **Provisioning Groups** table, shares are used to distribute the WAN Ingress/ Egress eligible bandwidth, which is the **Permitted Rate** minus the total **Min** reserved bandwidth of all Services on the WAN Link. All Services are initially assigned to a "Default" Group that is allocated all the eligible bandwidth. The user can create additional Groups and allocate bandwidth to its members by giving that Group a number of **Fair Shares**. The resulting total bandwidth for all Services in the Group is then shown in the **Fair (kbps)** column. Please see Figure 20 for a view of the Provisioning Groups section.

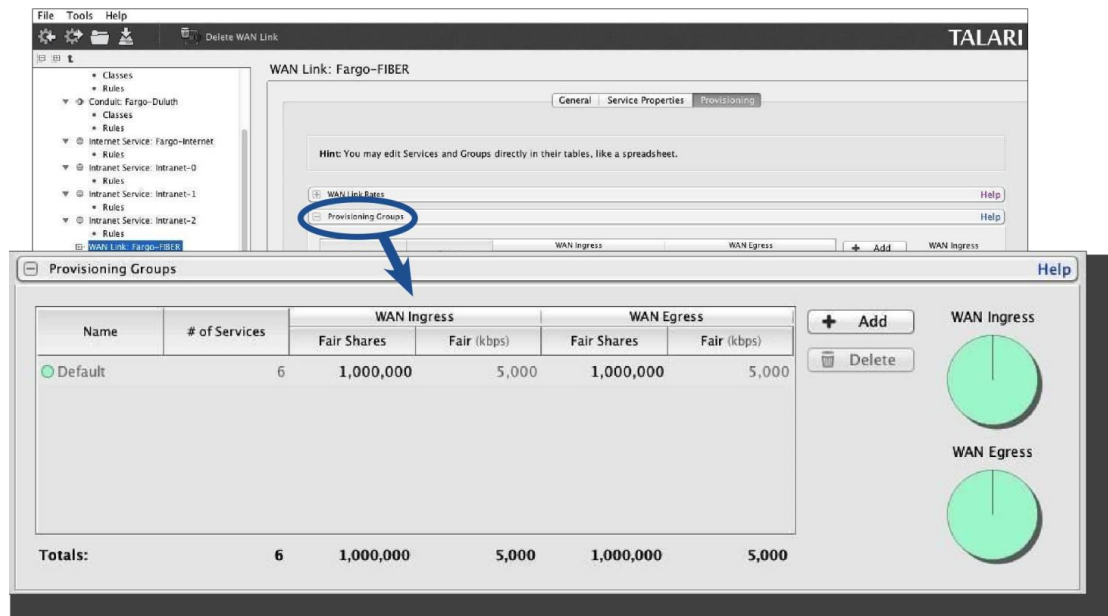


Figure 20

To create a Provisioning Group:

- Click the **Add** button
- Provide a Group name
- Provide the number of WAN Ingress and WAN Egress **Fair Shares** required for the new group
- Reassign Services to the new Group using the **Group** column in the **Services** table

Provisioning Groups are available to simplify the provisioning process and are not required if they are not needed.

The Concept of Using Shares

When provisioning bandwidth for Oracle SD-WAN Edge with a large number of sites, using percentages does not allow for enough granularity as the site count increases.

Oracle has instituted the use of shares for each of the Services or Groups of Services within the WAN Link. The total number of shares is up to the user, allowing any amount of granularity or precision when allocating bandwidth among the different Services. There are two distinct pools of these shares: WAN Ingress and WAN Egress.

Note:

All Services receive their Min Reserved Bandwidth before Fair distribution, which could result in Groups with equal Fair Shares having disparate Fair Rates. Fair Rates can also be affected by Service Maximums, if defined.

Services

The services definition for a WAN link are determined in this section (see Figure 21 below). For conduit services, the user would define the fair shares allocated to a client site. By default, all sites are placed in the “Default” group with the fair shares divided evenly. Services for individual site WAN links are shown and may be edited here.

- Display desired Provisioning Group or all groups by using the pull-down menu Add individual services to a Provisioning Group by selecting the service name and choosing the desired Group
- Set the desired WAN Ingress and Egress minimum and maximum rates, and Group shares for the service by double-clicking the cell to change the rate or number of shares
- To set an unlimited maximum rate, enter “0” or “no limit” into the cell
- Click the **Apply** button to save the settings

Shares of Group

On this table, the shares are used in the same way as above, but in this case, it is a new pool of shares within each group used. These shares are used to divide up the bandwidth among the members of a group based on the ratio of the current service divided by the total number of shares for the group in which it is a member. The Minimum rate acts as a base bandwidth allocation for each service, and the amount of bandwidth available for fair allocation is based on the total permitted for the group minus the sum of the minimums for each service in the group.

Glossary

- Adaptive Private Networking (Oracle SD-WAN Edge)

As used in this guide, the name for the whole network that includes the Adaptive Private Networking Appliances, the Wide Area Network, the conduits between peer APNAs, as well as other network application services. Oracle SD-WAN Edge is configured from a single APNA, which is the Network Control Node (NCN).

- Adaptive Private Networking Appliance (APNA)

The general name for a specific Oracle network appliance, also occasionally referred to as a Appliance.

- Client Node (Client)

A Oracle Client Node is an Oracle SD-WAN Edge appliance that is located across the Oracle network from the NCN. Although an NCN may potentially have multiple clients, each client has only one NCN.

- Conduit Service

The Conduit service is a logical combination of one or more paths, and is the typical mode for enterprise site-to-site intranet traffic, utilizing the full value of the Oracle's

Adaptive Private Networking. In this mode, depending on configuration, the traffic is actively managed across multiple WAN links to create an end to end conduit.

- Ethernet Interface

A physical or configurable interface of the APNA. For example, the T730 has nine userdefined Ethernet Gigabit interfaces, plus a predefined Management interface.

- Flow

A flow is a stateful instance (memory) used to track and treat application traffic from its source to its destination across Oracle SD-WAN Edge. The properties of a particular flow are derived from the routes, rules, and service that the traffic flow matches.

- Internet Network Service

The Internet Service is for traffic between an enterprise user and sites on the public Internet. Traffic of this type is not encapsulated. During times of congestion, Oracle SD-WAN Edge does actively manage bandwidth by rate-limiting Internet traffic relative to the conduit and intranet traffic as per the configuration established by the administrator.

- Intranet Network Service

The Intranet Service is for any portion of enterprise Intranet traffic that has not been defined for transmission across an Oracle SD-WAN Edge conduit. As with Internet traffic, it remains unencapsulated, and Oracle SD-WAN Edge manages bandwidth by rate-limiting this traffic relative to other service types during times of congestion. Note that under certain conditions, and if configured for Intranet Fallback on the Conduit, traffic between a pair of APNAs that ordinarily travels via a conduit may instead be treated as Intranet to maintain network reliability.

- Network Control Node (NCN)

The NCN is the central APNA that acts as the master controller of Oracle SD-WAN Edge, as well as the central point of administration for the client nodes. The NCN's primary purpose is to establish and utilize a conduit with one or more Oracle Client Nodes across the network for enterprise site-to-site communications. A particular NCN can administer and have conduits to multiple Client Nodes.

- Network Service

A logical set of operations performed on the traffic as it uses Oracle SD-WAN Edge. The set of services supported are Bypass, Passthrough, Internet, Intranet, and Conduit.

- Passthrough Network Service

Traffic directed to the Passthrough service includes broadcasts, ARPs and other nonIPv4 traffic, as well as traffic on the APNA's local subnet, specifically configured subnets, or rules applied by the network administrator. The APNA does not delay, shape or modify

this traffic. Because the Oracle service does not hinder this traffic, the network administrator must be sure that Passthrough traffic does not consume substantial resources on the WAN links which the APNA is configured to use for other services. Example: Passthrough may be used if a host is located on the WAN side of the APNA, but access to the host does not impact the APNA's specific WAN links. Think of the special management IP of the WAN link router as a typical example of a proper explicit use of Passthrough.

- Redundant Oracle SD-WAN Edge Control Protocol (RACP)

The protocol developed by Oracle to provide functionality for two high availability (HA) APNA's to communicate availability information.

- Rule

A Oracle Networking Service equivalent of a typical router access control list or filter mask. A rule defines match criteria and properties for IP flows. Flows that match those criteria use the service with which the rule is associated.

- Oracle Path Oracle Conduit Class

A Oracle Path is a logical link between two Oracle Virtual IP addresses (VIP). A Class is a queued service point into a Oracle conduit. The Class to which traffic is assigned determines its share of the conduit bandwidth, permitted queue depth, and its priority, relative to other traffic, for Oracle Network resources.

- **TCP Termination**

TCP termination provides the ability to split a single TCP connection into three separate TCP connections all managed and maintained by the Oracle SD-WAN Edge. TCP termination is only used for conduit traffic.

- Traffic Service Types

Traffic Service Types apply while the system is in the Active state noted above.

- Trust Relay Points (TRP)

A Cisco Systems software function implemented in voice over IP networks that provides multiple voice capabilities, such as transversing trusted firewalls.

- Trusted WAN Port

Appliance port processing network traffic that is protected by a firewall, performing as if it were a traditional WAN port.

- Untrusted WAN Interface

Appliance interface processing network segment traffic that is not being protected by a firewall. Non-conduit traffic from the WAN is unable to communicate to any network interface inside of the appliance. The segment is entirely isolated from the rest of the network with the exception of the APNs own 128 bit AES encrypted paths.

- WAN Link

The general term for an enterprise's connection to a WAN. These WAN links are typically connected to router ports. Some examples of WAN Links are T1, DSL, or Frame Relay.

3

Release 2.5 Features

This chapter includes features and enhancements released in 2.5.

Oracle Hardware Support

Oracle SD-WAN Release 2.5 incorporates the T5000 seamlessly into the entire family of appliances offered by Oracle. Basic Oracle SD-WAN Edge appliances platform capabilities are listed below.


Appliance Model	Conduits	WAN Ingress Paths	WAN Egress Paths	Flows	Flows with TCP Termination
T5000	128	576	576	256,000	16,000
T3000	128	576	576	256,000	16,000
T750	32	216	216	64,000	8,000
T730	16	72	72	64,000	4,000
T510	8	36	36	32,000	500

Table 1

Table 1 provides a detailed view of the supported hardware maximum capabilities. The number of conduits can be used to derive the number of sites if used as an NCN or as a client in a meshed configuration. If any of the appliances are used as a client device, the hardware can still support the number of conduits defined, but will be dependent on the Oracle SD-WAN Edge architecture deployed. Table 1 is provided so users understand the conduits supported by all platforms in Edge release 2.5.

Introducing the T5000 Appliance

Designed to bring WAN reliability and higher bandwidth to large data centers and call centers, the 2U rack-mountable Mercury T5000 appliance affordably delivers up to 3.0 Gbps uplink/3.0 Gbps downlink (6 Gbps total) across up to 8 WAN connections. T5000 appliance can easily communicate with other Oracle appliances such as the Mercury T510, T730, T750, and T3000 models. T5000 runs the same software as other Oracle appliances while taking performance and scalability to the next level, supporting gigabits of WAN bandwidth across the union of private WAN links and public Internet connections and providing support for up to 128 branch connections.



For information on installing the T5000, please reference the Oracle SD-WAN Edge T5000 Getting Started Guide and the Oracle SD-WAN Edge T5000 Hardware Guide.

4

Release 3.0 Features

This chapter includes features and enhancements released in 3.0.

Dynamic Conduits

The dynamic conduit is a conduit between two Oracle SD-WAN Edge client sites that is not predefined in the Oracle SD-WAN Edge configuration file, but is created on-demand based on network traffic. From a user perspective, the advantage of a conduit between client sites is that traffic can flow directly from one client Oracle SD-WAN Edge site to a second client site without having to traverse the NCN or two conduits. In addition, the conduit is built and removed dynamically based on user defined traffic thresholds. These thresholds are defined in either packets per second (pps) or bandwidth (kbps). From a configuration perspective, the Dynamic Conduit requires some up front configuration time. Another benefit of the Dynamic conduit is the ability for the any client to dynamically build a conduit to any other client.

This allows a dynamic full Mesh configuration for customer traffic flows. Once a threshold for the Dynamic Conduit is reached and the dynamic conduit is created, the appliances test the dynamic conduit before making full use of it in the following manner:

- Send Bulk data if any exists and verify no loss, then
- Send Interactive data and verify no loss, then
- Send Real Time data after the Bulk and Interactive data are considered stable (no loss or acceptable levels)
- If there is no Bulk or interactive data send Real Time Data after the conduit has been stable for a period of time

If the user data falls below the configured thresholds for a user defined period of time, the dynamic conduit is torn down.

Design Considerations

Based on the above traffic flows as well as the nature of dynamic conduits the user should be aware of certain Design considerations. These considerations are as follows:

- For voice traffic across the Dynamic Conduit be aware of WAN link limitations/quality
 - Loss
 - Latency of the WAN link
- In certain cases WAN link may not be recommended as a path for a Dynamic Conduit if there is a high loss or latency that would impact certain traffic types. In this case, do not configure the WAN link as part of the Dynamic Conduit
- How often a WAN link transitions from good to bad
- Ideally there is traffic between the Oracle SD-WAN Edge sites that is non-voice traffic Bulk or Interactive before Voice

- WAN link thresholds are based on all traffic on a WAN link, including conduit, Intranet, and Internet
- In this release Dynamic Conduits support a single “Dynamic Conduit Default set” for rules and classes.
- When using Dynamic conduits, the user should have consisted rules for the following options: Header compression and TCP Termination.

Adding a site to the Dynamic Conduit is a service reset for all site participating in the Dynamic Conduit.

Dynamic Conduit Configuration

Dynamic Conduits have the concept of an Intermediate site; this site could be an NCN site. If the NCN site has two client sites connected, Client A and Client B, with WAN-To-WAN forwarding enabled Client A would communicate with Client B through the NCN site (Intermediate). Any site configured as the Intermediate site monitors traffic flowing through sites that are configured to support Dynamic Conduits. In this example the NCN site is monitoring traffic levels between Client A and Client B. Once the configured threshold is reached through the Intermediate site the Dynamic Conduit is built between Client A and Client B.

The other high level design consideration is related to WAN-To-WAN Forwarding

Groups. By default, all Oracle SD-WAN Edge sites reside in the default forwarding group. When WAN-To-WAN forwarding is enabled all routes from all sites are known throughout the Oracle SD-WAN Edge. This may not be desired. Because of this, the concept of WAN-To-WAN Forwarding Groups was added in the 3.0 release. The user now has the ability to create multiple WAN-To-WAN Forwarding Groups that do not share routes. In this release the Intermediate Oracle will forward between WAN-To-WAN Forwarding Groups. In future software, a user will have an option to forward between WAN-To-WAN Forwarding Group or not.

A high-Level description of the configuration process follows.

The process for configuring a Dynamic Conduits is as follows:

- Identify intermediate Site
 - Enable intermediate site (use default WAN-To-WAN Forwarding Group)
 - Enable WAN to WAN forwarding at intermediate site
- Identify Client sites for Dynamic Conduits
 - Enable dynamic conduits at the clients site
- Enable the Dynamic conduit service (WAN Link – Service properties)
 - Provision WAN Link resources for the Dynamic Conduit (shares)
 - Identify threshold used for Dynamic Conduit Creation
 - Define using Dynamic Conduit Default Set
 - Define threshold at Intermediate Site WAN Link

Dynamic Conduit Configuration Creation

For the Dynamic Conduits to be created the user would define a site, typically the NCN (site but not required to be the NCN site) site, to act as the intermediate node for the client sites. These options are configured at the appliance level. At this site the user

would enable “WAN-To-WAN forwarding”, as well as select the “Intermediate site” option, see figure 1.

Appliance: t-ncn-colo

Name: t-ncn-colo

Secure Key: 0x ef18171625

Model: T5000 Mode: primary_ncn

WAN To WAN Forwarding Group: [Dropdown]

Default Direct Route Cost: 5

Gateway ARP Timer (ms): 1000

WAN to WAN Forwarding

Enable WAN to WAN Forwarding Route Cost: 10

Dynamic Conduit Settings

Enable Dynamic Conduits Set as Intermediate Site

Source MAC Learning

Enable Source MAC Learning

Interface Groups Virtual IP Addresses Conduits WAN Links Routes Network Services

VLAN1

+ Add

Edit

Delete

Apply Reset

Figure 1

In this example the NCN with the WAN-To-WAN option enabled will forward all routes to all client appliances within a WAN-To-WAN Forwarding Group. By default, all sites reside in the default WAN To-WAN Forwarding Group. The user can also define additional WAN-To-WAN Forwarding Groups as required. The groups are defined at the Global configuration level. The intermediate site will monitor the traffic flow between Client Sites to determine if the traffic level reaches the user defined threshold. If the traffic flow reaches the defined threshold, the Intermediate node will instruct the client nodes to establish a Dynamic Conduit. The sample thresholds will be described later in this document.

At the client node, the user would enable the dynamic conduits option at the appliance level.

Appliance: t-client-Bangalore

Name: t-client-Bangalore

Secure Key: 0x f98fcb99

Model: T730 Mode: client

WAN To WAN Forwarding Group:

Default Direct Route Cost: 5

Gateway ARP Timer (ms): 1000

WAN to WAN Forwarding

Enable WAN to WAN Forwarding Route Cost: 10

Dynamic Conduit Settings

Enable Dynamic Conduits Set as Intermediate Site

Source MAC Learning

Enable Source MAC Learning

Interface Groups Virtual IP Addresses Conduits WAN Links Routes Network Services

Banga-bypass

+ Add

Edit

Delete


Apply Reset

Figure 2

There are two methods for configuring thresholds for a dynamic conduit. The Dynamic Conduit will be created if any of the configured values (thresholds) are reached. The options can be configured at a global level or based on a WAN Link configured at the Intermediate node. If the user does not want to match on the WAN link they would only have to configure the thresholds at the global level. Currently if configured at the WAN link level all traffic accounted for is counted as the threshold value, so conduit traffic, intranet traffic and Internet traffic are all count towards the WAN link threshold. Examples of these options are defined below:

Option 1:

The advantage of this option is to offload bandwidth on one of the intermediate node WAN links. As clients communicate to each other through the intermediate node there may be a requirement to remove this traffic from one (or multiple) of its local WAN links. This can be accomplished by defining a threshold on the local WAN link. If one of the thresholds is reached the Dynamic Conduit will be established between client sites. The key design point when using the WAN link threshold option is that this is total traffic on the WAN link. This includes conduit, internet, and intranet traffic, not just client to client traffic. The option is defined under the appliance – WAN Link – General-Property's tab. Figure 3 displays the options to configure the available threshold options.

**Figure 3**


Option 2:

Once the Dynamic Conduit is enabled at a client site there is a Dynamic Conduit

Default Set defined. Within this default set is a properties tab which includes

"creation limits". The values for the conduit create are Sample time in seconds (default value 10 seconds), Throughput in kbps (default value 250 kbps), and Throughput in pps (default value 10 pps).

From the global level once a client site has "Dynamic Conduits" enabled look for "Dynamic Conduit Default Set: Default". Figure 4 displays these settings.

**Figure 4**

Once the intermediate site is defined, and WAN To WAN forwarding is enabled and Client sites have dynamic conduits enabled if any Creation limit is reached the dynamic Conduit is created. Again, this can be pps or kbps.

For each client site the user would also have to provision the Bandwidth shares for the Dynamic Conduit. These provisioned Fair shares per WAN link are used by all dynamic conduits on that WAN link. The allocated minimum reserve shares are per Dynamic Conduit for the WAN Link. Figure 5 shows an example of enabling the service on a WAN Link:

General Service Properties Provisioning

Hint: You may edit Conduit and Internet/Intranet Services directly in their tables, like a spreadsheet.

Conduit Services

Show Services from All Groups

On	Name	Group	Tunnel Header Size (bytes)	Active MTU Detect	UDP Port #	UDP Hole Punching	UDP Port Switching			Autopath Group
							Enable	Alt Port #	Interval (min)	
<input checked="" type="checkbox"/>	<DYNAMIC>	Default	0	<input type="checkbox"/>	2156	<input type="checkbox"/>	<input type="checkbox"/>	0	1,440	<Default>
<input checked="" type="checkbox"/>	Client-GEU	Default	0	<input type="checkbox"/>	2156	<input type="checkbox"/>	<input type="checkbox"/>	0	1,440	<Default>
<input checked="" type="checkbox"/>	Client-WH	Default	0	<input type="checkbox"/>	2156	<input type="checkbox"/>	<input type="checkbox"/>	0	1,440	<Default>
<input checked="" type="checkbox"/>	NCN-Colo	Default	0	<input type="checkbox"/>	2156	<input type="checkbox"/>	<input type="checkbox"/>	0	1,440	<Default>

Figure 5

While Figure 6 shows how to define the shares for a WAN Link once the service is enabled.

Services

Show Services from All Groups

Name	Group	WAN Ingress				WAN Egress			
		Min (kbps)	Max (kbps)	Shares of Group	Fair (kbps)	Min (kbps)	Max (kbps)	Shares of Group	Fair (kbps)
<DYNAMIC>	Default	80	no limit	100	354	80	no limit	100	354
Client-GEU	Default	80	800	100	114	80	800	100	114
Client-WH	Default	80	no limit	100	114	80	no limit	100	114
NCN-Colo	Default	80	900	400	216	80	900	400	216
Internet	Default	100	no limit	300	202	100	no limit	300	202
Totals:		660	1,700	1,000	1,000	660	1,700	1,000	1,000

Dynamic Conduit Summary

Max Dynamic Conduits Possible	WAN Ingress					WAN Egress				
	Min Per Conduit (kbps)	Min Total (kbps)	Max Per Conduit (kbps)	Max Total (kbps)	Fair Per Conduit (kbps)	Min Per Conduit (kbps)	Min Total (kbps)	Max Per Conduit (kbps)	Max Total (kbps)	Fair Per Conduit (kbps)
4	80	320	no limit	no limit	88	80	320	no limit	no limit	88

Figure 6

There is also dynamic conduit remove settings that are user definable. In addition to the above, the web console allows the user to delete a dynamic conduit or Freeze a dynamic conduit. The freeze option allows the user to keep the conduit up and ignore the remove conduit. This feature would be used for testing a dynamic conduit as well as for troubleshooting purposes. These options reside under the Manage Network Dynamic Conduits tab in the web console.

WAN To WAN Forwarding Enhancements

To provide the flexibility required for Dynamic Conduits to operate, WAN-to-WAN forwarding was enhanced to allow for multiple groups. All sites that are part of a WAN-to-WAN Forwarding Group with WAN-To-WAN forwarding enabled have a common

routing table. The routing table consists of routes to all other sites in the group. Many customers in the past did not enable WAN-To-WAN forwarding because of this fact. By default, all APNA's are applied to the default group. If the requirement is for only certain sites to support dynamic Conduits the user would define a new WAN-To-WAN

Forwarding Group, then at the appliance level assign the appliance to the correct WAN-To-WAN Forwarding Group.

In addition, APNA's in one group will not have direct routes of an APNA that resides in another WAN-To-WAN Forwarding Group. The user also has the flexibility in allowing or excluding Internet routes and Intranet routes in the routing table. The

Internet/Intranet routes are considered local routes from a WAN-To-WAN forwarding perspective and included in the routing table unless otherwise configured. When configuring or planning to deploy Dynamic Conduits contact your Oracle representative for any additional information.

Routing Enhancements

Intranet or Internet Fallback Routes

There were certain configurations when route eligibility was used that a Conduit Fallback route was not hit because the conduit was down, but the gateway was still reachable. When this occurred, the traffic would hit a pass-through route which in certain designs was then dropped by the Oracle. To eliminate Oracle from dropping frames the user can now select Intranet/Internet fallback routes such that if the conduit fails the Oracle will forward traffic to the defined Intranet router/gateway. See figure 7 for details.



From figure 5 we can see that this option is configured under "Appliance – Internet Service – Ignore WAN link Status". The same option exists under Intranet service and is enabled via the check box.

Additional Enhancements

- **Auto-Path Groups for WAN Links**
In previous releases the user would have to create a manual path and then edit the advanced attributes for the path for private WAN links. Because there was no auto-path for private (intranet) WAN links the user was forced to define the paths and attributes for the path. This new feature allows a user to define an "Auto-Path-group" at the global level and assign attributes to the group.

An "Auto-Path-group" defines a set of WAN links in the Oracle SD-WAN Edge that are reachable to each other. Figure 8 displays the options for the group.



Figure 8

The options other than the group name are standard Oracle options:

- DSCP setting for the path
- Encryption enabled or disabled
- Bad loss sensitivity
- Instability Sensitive

Once the group is defined it can be applied to as many WAN Links as needed reducing the configuration time.

Figure 9 displays where to apply this option.

The group is applied at the WAN link – Service properties – Conduit Level, under the Autopath group pull down menu.

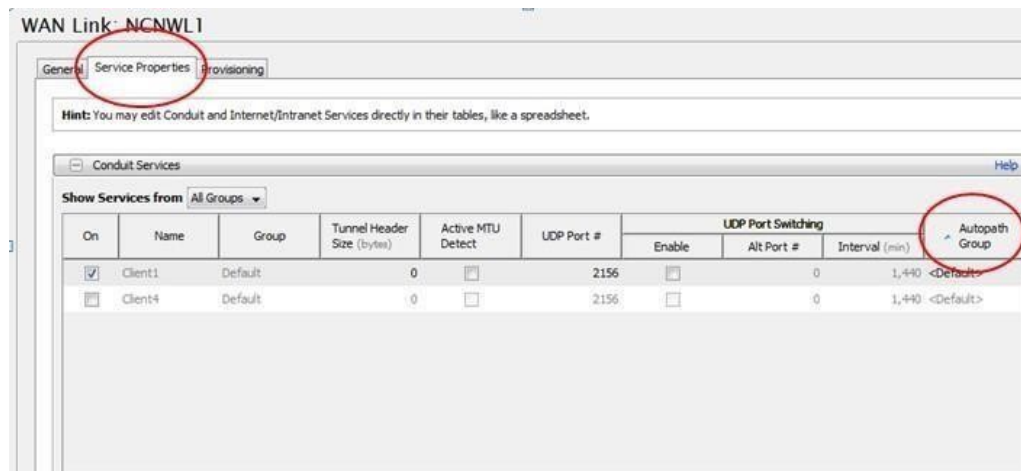


Figure 9

- SNMP
 - Within the Oracle SNMP MIB there were enhancements and bug fixes added to the Oracle MIB. These changes include the following:
 - Add a Last updated date for the MIB
 - OracleNumEvents is now a gauge with events being clean up after 30 days
 - When in standby mode for HA counters are now accurate
 - All rule statistics types have changed from counters (64) to gauge
 - Dynamic Conduit statistics table is a separate table as there is only data when the conduit is in use
 - The appliance serial number is viewable from SNMP
 - The MIB now allows customers to define a “Contact”, “Location”, and or Description for an appliance. This is configured under Integrate >>Configure Events and Alerts in the web console of the appliance
 - tnStatsRuleTable now reports data correctly and will match the web console displayed data for rule statistics
 - tnStatsConduitEntry –now records BOWT (Best One Way Time), Jitter, Packets Lost uni- directionally
 - The ability to query Oracle version information was added to the MIB, including APNware version, OS version etc.
 - tnstatsConduitClassType values now match the web console classes as expected
 - Cisco NetFlow MIB is no longer supported
 - * ARP Timeout Setting

In the past this option was set-able through a debug level. This has now been moved to a user configurable parameter. Any previously configured ARP timeout value set using the debug level will be lost and must now be configured using the Gateway ARP Timer setting in the Configuration Editor. It defines the amount of time between ARPs for a Oracle WAN link gateway.

This option is changed when devices do not handle the 1 second ARP value used by default in the Oracle appliance. The user can now define the number of seconds for ARP on an appliance at the global level. Figure 8 displays where to define this option.

Appliance: t-ncn-colo

Name: t-ncn-colo

Secure Key: 0x ef18171625

Model: T5000 Mode: primary_ncn

WAN To WAN Forwarding Group: [Dropdown]

Default Direct Route Cost: 5

Gateway ARP Timer (ms): 1000

WAN to WAN Forwarding
 Enable WAN to WAN Forwarding Route Cost: 10

Dynamic Conduit Settings
 Enable Dynamic Conduits Set as Intermediate Site

Source MAC Learning
 Enable Source MAC Learning

Interface Groups Virtual IP Addresses Conduits WAN Links Routes Network Services

VLAN1

+ Add
 Edit
 Delete

Apply Reset

Figure 10

- WAN Link Disable Paths Option

In previous releases, a user wishing to prevent traffic from using a WAN link would have to individually disable each path that used the WAN link. Based on user requests, this has been enhanced to allow a user to select a WAN link to disable, which will automatically disable all paths on that WAN link.

- Ping Enhancements

Ping has now been enhanced to include the following new options:

- Ping count
- Packet size

In the past the system would only ping 5 times and the user could not define the size of a ping frame. This has been enhanced for troubleshooting purposes.

- Security Enhancements

Certain security enhancements are now included in APNware 3.0. These are listed below.

- Within the web console or the CLI the admin level user can now change the root level password if they know the current root password.

- Users accessing the appliance from SSH now can be authenticated by RADIUS or TACACS (Release 3.0P1)
- Additional RADIUS and TACACS server are now configurable, if the first server fails then the next server in the list will be used to authenticate the user. (release 3.0P1)

5

Release 3.1 Features

This chapter includes features and enhancements released in 3.1.

Default Configuration Parameter Change

In the Oracle SD-WAN Edge Software Release 3.1 there has been a change to a default configuration parameter. In previous software releases the Oracle SD-WAN Edge configuration editor/compiler would add in a bridge pair when configured for Fail-to-Wire (FTW) or provide a warning message that no bridge pair had been configured. Typically, the user would then add in the bridge pair definition if required. This would lead to a configuration file having the following entry for a bridge pair (as an example):

```
add bridge_pair device_one=1 device_two=2
```

This allows the Oracle appliance to bridge traffic between ports 1 and port 2, which also creates a FTW pair. The change is that now the system will NOT automatically add into the configuration file the above bridge commands. Instead the following warning will be provided to the user during validation of creating a configuration file:

```
* ---> WARNING: EC329: in define site 'Client---Test' ---> add
appliance ---> add interface_group: Interfaces in this interface
group are not part of a bridge pair. Devices connected to these
interfaces will not be reachable without the aid of an external
device.
```

To resolve the configuration warning the user should identify the site with the configuration issue and perform one of the following:

- Add in the bridge pair
- Enable Source MAC Learning

If the issue is not resolved, traffic that is considered pass--- through will not flow between port 1 and port 2. This may be required behavior, which is the reason for the default configuration change.

When upgrading from a previous release the user should not receive these warnings, but there are always exceptions. If the user does receive the warning message and has questions please contact support for clarifications on the message, configuration assistance and help in resolving the warning message.

Oracle SD-WAN Edge Configuration Editor

The new Oracle SD-WAN Edge Configuration Editor is a web--- based tool incorporated with Oracle SD-WAN Edge

Software that delivers Oracle SD-WAN Edge configuration editing capabilities, configuration compiling, and the ability to create and edit network maps. The Configuration Editor operates on Configuration Packages, which consist of an Oracle SD-WAN Edge configuration and one or more optional network maps. These Configuration Packages can be saved, re--- opened,

exported, or imported from the Configuration Editor. Unlike the legacy Configuration Editor, this new tool works without Java, making it a more secure choice for network development.

You will find the Configuration Editor on the web interface by navigating to **Manage Network**, and then **Configuration Editor**.

Once inside the Configuration Editor, you can select “New” to start building your Configuration Package or you can select “View Tutorial” to be guided through an introduction of the new tool.

The Configuration Tree (located in the left--- side pane of the Configuration Editor) is where you can add sites, define connections between those sites, and perform provisioning activities. Figure 3 illustrates an example of the Configuration Tree expanded to edit the Basic Settings for a site. All items denoted with “+” can be expanded to show more detail.

After building the network configuration, a Network Map can be customized by clicking & dragging individual sites onto the map, or by auto populating the map with all sites. Every new

Configuration Package has a default map called “Network Map” associated with it. This map can be re-named and additional maps can be created using the “+” tab.

To further customize your map, you can also add backgrounds to it. Click the gear symbol on the map and select “Set Background”. As shown in the example below, you can then place your Sites on the background where you wish. (See figure 5)

Oracle SD-WAN Edge Configuration Editor and Oracle SD-WAN Edge Aware

Oracle SD-WAN Aware is a new network management system to be release in 2014. In order to deploy Oracle SD-WAN Aware in an Adaptive Private Network, each Oracle *Mercury* Oracle SD-WAN Edge Appliance in the network must have its software updated to R3.1 or higher. The new Oracle SD-WAN Edge Configuration Editor is the

foundation of the configuration model that will allow an Oracle SD-WAN Aware Node to perform Management and Monitoring functions for the entire Oracle SD-WAN Edge. For more information on Oracle SD-WAN Aware and its implementation please contact your local Oracle representative.

6

Release 4.0 Features

This chapter includes features and enhancements released in 4.0.

256-Site Adaptive Private Networks



Note:

This feature is only supported for APNs in which Oracle SD-WAN Aware R1.0 GA P1 (or later) has been deployed. See *APN Software R4.0 GA Release Notes* for more details.

In Oracle SD-WAN Edge Software R4.0, a feature has been added to enable a single T5000, configured as an NCN, to create up to 256 static conduits. This doubles the previously supported scale. This will, in the general case, enable a T5000 NCN to govern an Oracle SD-WAN Edge with up to 256 Client Sites.

The T5000 can now also support a greater number of WAN Paths and Flows. See the below table for details on supported capacity for each NCN-capable appliance model:

Appliance Model	T750	T3000	T3010	T5000 (w/o Aware)	T5000 (w/ Aware)
Max Static Conduits	32	128	128	128	256
Max Dynamic Conduits	16	32	32	32	32
Max WAN Ingress Paths	216	576	576	576	1152
Max WAN Egress Paths	216	576	576	576	1152
Max Flows (TCP Term off)	64,000	256,000	256,000	256,000	512,000
Max Flows (TCP term on)	8,000	16,000	16,000	16,000	16,000
Max Public WAN Links	8	8	8	8	8
Max Private WAN Links	32	32	32	32	32

Configuration Editor

In Oracle SD-WAN Edge Software R4.0, the user will continue to be able to create and edit Oracle SD-WAN Edge configuration files from the NCN web interface via **Manage Network > Configuration Editor**. However, the Oracle SD-WAN Edge Configuration Editor on the NCN will only support creating and editing Oracle SD-WAN Edge configuration files for APNs with up to 128 Client Sites. To scale past 128 Clients, the Oracle SD-WAN Edge configuration file must be managed from Oracle SD-WAN Aware.

Changes to Data Storage on Oracle SD-WAN Edge Appliances



Note:

The following changes apply to all APNs, regardless of Oracle SD-WAN Aware deployment.

In Oracle SD-WAN Edge Software R4.0, it is guaranteed that each APNA participating in a worst-case typical 256-Site Oracle SD-WAN Edge will be able to store at least 7 days of statistical and event data.

Previously, 30 days of statistical storage was guaranteed on each APNA in an Oracle SD-WAN Edge. With the introduction of Oracle SD-WAN Aware, storage of statistical data on the APNAs is not the preferred method for network monitoring. An Oracle SD-WAN Edge equipped with Aware can offload statistics from the APNAs to the Aware Node and accumulate up to a year's worth of network-wide data.

Changes to Local Route Scale



Note:

The following changes apply to all APNs, regardless of Oracle SD-WAN Aware deployment.

In Oracle SD-WAN Edge Software R4.0, support for unique local routes was increased for all Oracle SD-WAN Edge appliances. Previously, unique local route scale was capped at 128 for all appliances models. See the below table for details on increased capacity for each appliance model:

Appliance Model	T510	T750	T750	T3000	T3010	T5000
Max Unique Local Routes	512	512	512	512	2048	2048

7

Release 4.1 Features

This chapter includes features and enhancements released in 4.1.

Oracle Virtual Appliance CT800

Oracle SD-WAN Edge Software R4.1 supports the deployment of a Oracle Virtual Appliance CT800 within Amazon Web Services (AWS). This feature allows the user to provide Conduit connectivity to the AWS cloud. Using the Oracle CT800 solution, users will be able to leverage the typical Oracle features providing a more reliable and secure connection to the AWS cloud. The solution also allows users to take advantage of the quality of service offering provided today by Oracle Conduit. As cloud demands grow, users can leverage Oracle Conduits to allocate WAN bandwidth to more critical applications when demand dictates.

Oracle CT800 features within AWS:

- Supports up to 100 Mbps of Conduit throughput
- Supports Static and Dynamic Conduits
- Supports Internet and AWS Direct Connect WAN services
- Supports typical FTB configurations
- Supports configuration of WAN/LAN/Management ports as needed

Oracle CT800 requirements within AWS:

- AWS EC2 Instance Type: c3.2xlarge
- # of vCPUs: 8
- RAM: 15 GB
- Storage: 160 GB
- # Network Interfaces: 2¹

For installation details, see the *Oracle Virtual Appliance CT800 Getting Started Guide*.

As your cloud services grow, deploy the Oracle CT800 for a more robust and secure connection from your enterprise network to services in the AWS cloud.

MOS Estimation

In conjunction with Oracle SD-WAN Aware 1.1, Oracle SD-WAN Edge Software R4.1 supports the ability to provide Mean Opinion Score (MOS) estimates for defined applications traversing the WAN via a Oracle Conduit. This will assist the user in problem isolation for voice or other critical traffic flowing across the Conduits.

It is important to note that Oracle MOS estimation is based on WAN Egress flow processing. As traffic matching a defined rule is received at the WAN Egress side of the Conduit, the

¹ At minimum, the Oracle CT800 requires 2 Network Interfaces (1 for MGT and 1 for LAN/WAN). However, the CT800 can support up to 4 Network Interfaces.

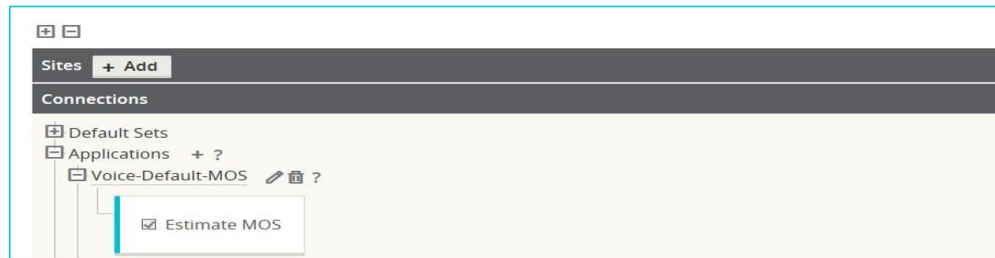
values for calculating the MOS are saved in the database file for reporting in Oracle SD-WAN Aware. The MOS calculation is based on the E-Model or ITU-T G.107. The following describes how a user would enable the MOS capability in the Oracle SD-WAN Edge Configuration Editor, then view calculated MOS values in Oracle SD-WAN Aware.

How to Configure

1. Create an Application with "Estimate MOS" Enabled

In the NCN web console, under **Manage Network**, and then **Configuration Editor**:

- Open a configuration to edit (or create a new configuration)
- Navigate to **Connections -> Applications** and select the "+" to add a new application
- Give the new application a name and make sure that "Estimate MOS" is enabled



1. Add the Application to a Rule for a Static Conduit

In the NCN web console, under **Manage Network -> Configuration Editor**:

- Open a configuration to edit (or create a new configuration)
- Navigate to **Connections -> [Site Name] -> Conduits -> [Conduit Name] -> Local Site**

Under the **Rules** section, add a new rule with the previously-defined Application Name.

Note:

The rule can be added to a Default Set or to a specific Conduit.



3. Verify the Traffic Matches the Defined Rule

In the web console, under **Monitor -> Flows**:

- Review relevant traffic flows to ensure they are matching the defined rule

How to View

Once an application has been configured with "Estimate MOS" enabled, calculated MOS values for the application can be viewed in Oracle SD-WAN Aware. MOS values are calculated over one-minute intervals. The user can view both Average Oracle SD-WAN Edge MOS and Lowest Oracle SD-WAN Edge MOS for each interval.

- **Average APN MOS:** Calculated using the average latency for packets observed on WAN Egress that match the application over one minute, and an average loss percentage (sampled every second) for all flows matching the application over multiple minutes.
- **Lowest APN MOS:** Calculated using the average latency for packets observed on WAN Egress that match the application over one minute, and an average loss percentage (sampled every second) for all flows matching the application over multiple minutes.

To generate a graph of MOS values, in the Oracle SD-WAN Aware web console navigate to **Monitor > Graphs -> [Site Name] -> Conduits -> Applications -> Application Name -> [Average APN MOS | Lowest APN MOS]**.

To generate a report of MOS values, in the Oracle SD-WAN Aware web console navigate to **Monitor > Reports -> Applications**.

Security Enhancements

Oracle SD-WAN Edge Software R4.1 supports several security enhancements. Although not required for deployments, these features are available for environments that require an additional level of encryption or security in general.

Summary

A summary of these features is provided below. See *APN Security Technical Paper* for more details. Each feature is configurable at a global level, for the entire Adaptive Private Network (Oracle SD-WAN Edge).

256-bit AES Encryption

256-bit AES Encryption is now supported, in addition to the previously supported 128-bit AES Encryption. 256-bit AES Encryption is *not* enabled by default.

Enhanced Encryption Key Generation/Rotation

Per-session encryption keys are generated and automatically rotated (when Encryption Key Rotation is enabled) using an Elliptic Curve Diffie-Hellman algorithm. Encryption Key Rotation is enabled by default.

Extended Packet Authentication Trailer

To provide users with the ability to have strong message authentication, an optional trailer inside the encrypted payload can now be enabled. By default, this optional trailer is composed of a 4-byte checksum of the unencrypted packet data, which acts like a standard Hashed Message Authentication Code (HMAC). While a standard HMAC would impact performance significantly, this checksum trailer provides a similar benefit while minimizing processing overhead. If use of a standard HMAC is required, the optional trailer can be configured to use a 16-byte SHA-256 HMAC in place of the 4-byte packet checksum.

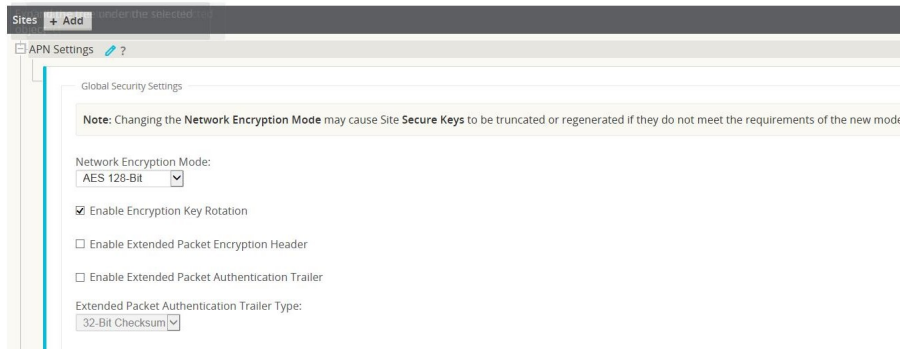
Extended Packet Encryption Header

To provide users with the ability to have the highest level of packet uniqueness and protection against Frequency Analysis, an optional 16-byte counter can now be prefixed inside the encrypted payload to act as a rotating, cryptographically random Initialization Vector.

How to Configure

In the NCN web console, under **Manage Network -> Configuration Editor**:

- Open a configuration to edit (or create a new configuration)
- Navigate **Sites -> Settings**



Note:

Enabling some of the enhanced security options could impact performance. Contact your Oracle Sales representative or Oracle Support for help in understanding the potential performance impact.

SNMP Polling for ARP Table

Oracle SD-WAN Edge Software R4.1 supports retrieval of ARP entries and associated statistics via SNMP polling. Polling will retrieve the data for existing ARP table entries, along with associated interfaces and statistics for each entry. The data for this SNMP entry is updated once a minute on the Oracle SD-WAN Edge Appliance.

For more data on this specific MIB enhancement please review the Oracle MIB and from the Oracle SD-WAN Edge Appliance web console under **Integrate -> Download/View Oracle MIB**.

Appliance Settings from Aware

In conjunction with Oracle SD-WAN Aware R1.1, Oracle SD-WAN Edge Software R4.1 supports the ability to configure appliance settings from Oracle SD-WAN Aware. Appliance settings include options associated with DNS, NTP, Time Zone, User Authentication, FTP server configuration, Notification settings,

Netflow, etc. These options are not governed by the Oracle SD-WAN Edge Configuration Editor. With Oracle SD-WAN, Aware R1.1, these settings can now be pushed from Oracle SD-WAN Aware to any user-selected group of Oracle SD-WAN Edge Appliances in the network. Additionally, templates of appliance settings can be

created, edited, and saved on Oracle SD-WAN Aware to streamline future Oracle SD-WAN Edge Appliance installs.

From a deployment perspective, the following should be considered:

- Settings defined locally will override settings that were previously pushed from Oracle SD-WAN Aware
- Settings pushed from Oracle SD-WAN Aware will override settings that were previously defined locally
- Settings pushed from Oracle SD-WAN Aware will take effect once received on the target Oracle SD-WAN Edge Appliance(s)
- Settings pushed from Oracle SD-WAN Aware are received on the target Oracle SD-WAN Edge Appliance(s) via the management interface
- Settings can only be pushed from Oracle SD-WAN Aware to those Oracle SD-WAN Edge Appliances that are reachable (i.e. displayed on the **Manage -> Discovery** screen)
- Blank settings in a template will not be pushed to the Oracle SD-WAN Edge Appliance(s)

In the Oracle SD-WAN Aware web console, under **Manage -> Appliance Settings**:

- Select "New" to create a new Appliance Settings template
- For each section, select "Include in File" to edit the options in that section
- Once all desired options have been set, select "Save" to save the template
- Select "Export" to export the template to the desired Oracle SD-WAN Edge Appliance(s)

Manage / APNA Appliance Settings

New Open... Save Save As... Import... Export...

General ■ Include in File ?

Inactivity Timeout:

DNS ■ Include in File ?

Primary Nameserver: Secondary Nameserver:

NTP ■ Include in File ?

Use NTP Server: Host:

Timezone ■ Include in File ?

Time Zone: UTC

8

Release 4.2 Features

This chapter includes features and enhancements released in 4.2.

Non-Resetting Configuration Updates

Oracle SD-WAN Edge Software R4.2 now supports all network configuration updates as non-resetting updates. This process change is inherent in all configuration updates and does not require any specific command to be enabled. When a new Oracle SD-WAN Edge Configuration is applied to the network via Change Management, the Oracle Service will not be forced to restart for any combination of configuration parameter changes. In addition, the web interface will provide detailed information on the time required for a staged configuration update to be performed on the network, allowing the user to understand the potential impact of an update to their Appliances and to their network as a whole.

Configuration Updates

Under **Manage Network -> Change Management** on the web interface of the Network Control Node (NCN), there are two new columns, *Traffic Interruption - Expected* and *Traffic Interruption - Actual*. These columns display values that indicate the following for each Site in the Oracle SD-WAN Edge:

Traffic Interruption - Expected: The expected worst traffic interruption time for the Site, based on the difference between the active and staged configurations. For more detail on the impact to each Conduit, refer to APN_traffic_impact.log.

- 0 sec - No service interruption for this Site
- <1 sec - Traffic interruption time will be less than 1 second
- <1 min - Traffic interruption time will be less than 1 minute
- <3 min - Traffic interruption time will be less than 3 minutes (only relevant for software updates, which still require a restart of the Oracle Service)
- Loc Chg Mgt - Package must be applied via that Site's local change management

Traffic Interruption - Actual: The time it took to perform a configuration or software update at the Site. This does not include the time it took for full network convergence during an update. Updates to Sites may occur in parallel.

- N ms/s - Actual traffic interruption time at the Site when the update was performed
- Err - An error occurred when performing the update

Manage Network / Change Management Talari Support

Activate ?

You may now activate the changes that have been distributed across your network. Each appliance will apply the changes and restart the Talari service.

Note: A reboot or loss of power during this operation may result in an incomplete installation which could require manual installation to resolve.

Click **Activate Staged** to begin.

Activate Staged In: 10 seconds

Activation Complete.
The network change process has finished. Click **Done** to exit this screen.
To undo your changes, click the **Revert** button.

Revert Abort Done

Currently Prepared: Configuration - aws-ct800-testbedNewCfg.zip Software - Current Running

Configuration Filenames: Active - aws-ct800-testbedNewCfg.zip Staged - aws-ct800-testbed.zip

Site-Appliance	Model	State	Currently Active		Currently Staged		Traffic Interruption		Download Package
			Software	Config	Software	Config	Expected	Actual	
NCN-ct800-ncn	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
AU-ct800-AU	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
BRA-ct800-BRA	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
Client-ct800-client	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
GEO-ct800-GEO	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
IRL-ct800-IRL	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged
SNG-ct800-SNG	CT800		R4_2_QA_D1_11132014	21:52 on 11/21/14	R4_2_QA_D1_11132014	19:31 on 11/21/14	0 sec	0 ms	active / staged

From the above, the expected traffic interruption time for the configuration update was 0 seconds (i.e. no expected interruption) for each Site, and that the actual interruption time met this expectation. Certain configuration updates will impact traffic for a period of time. When this is the case, the expected traffic interruption time will show the maximum interruption time that can be expected in the form of <1 sec, <1 min, or <3 min. Prior to activating a staged configuration update, the user can use the expected traffic interruption time for the staged configuration to determine if a maintenance window is required for the update. Please note that an appliance with traffic load may require more time to complete the update process.

Impact of Common Configuration Updates

The expected traffic interruption times for various common configuration changes are outlined below. These are estimates; the actual interruption times may be less and will be displayed in the web interface and logged for help with future planning.

No Interruption:

- NCN mode (primary, secondary) is changed
- Gateway ARP Timer is changed or Proxy ARP is enabled/disabled
- WAN-to-WAN Forwarding is enabled/disabled
- Intermediate Node attribute is enabled/disabled
- The attributes of a Route, Rule, Class or Autopath are changed
- Maximum Dynamic Conduits is changed or Dynamic Conduit Thresholds are changed
- WAN Link Conduit, Intranet, or Internet usage is changed (but not added/removed)
- Remote Site that has no Conduit to the local Site is added/removed
- Encryption key is changed or key rotation is enabled/disabled

Interruption of Less Than 1 Second:

- Site is added/removed or Site name is changed
- Conduit is added/removed/changed
- WAN Link or WAN Path is added/removed/changed
- Appliance Name is changed
- HA attributes are changed
- Interface Group is added/removed/changed
- WAN-to-WAN Forwarding Group is changed
- Dynamic Conduits are enabled/disabled
- WAN Link Conduit, Intranet, or Internet usage is added/removed
- Source MAC Learning is enabled/disabled
- Extended Packet Encryption Header is enabled/disabled
- Extended Packet Authentication Trailer is enabled/disabled or Trailer Type is changed

Interruption of Less Than 1 Minute:

- Encryption mode for the Oracle SD-WAN Edge is changed
- Any routing table update change when Local Route Eligibility is enabled at a site.
WAN Path Encryption is enabled/disabled

Software Updates

As in previous releases, software updates will require a restart of the Oracle Service. The new display will instruct the user of the expected traffic interruption time. The following screen shows the expected and actual traffic interruption times for a software update:

Manage Network / Change Management Talari Support

Change Process Overview ?

The Change Management process allows a user to upload changes to the network, whether it be a configuration, software, or both. This three-step workflow is a set of checks and processes that ensure that configuration changes and software updates are applied in a reliable, fail-safe way.

Step 1
Change Preparation
Upload Files to NCN

■■■■■■■■■■> NCN

Step 2
Appliance Staging
Transfer Files to Clients

■■■■> NCN Clients

Step 3
Activation
Activate Change

■■■■> NCN Clients

Clicking the Activate Staged button will skip to the Appliance Staging step, where you may switch to a previously-staged appliance package (if present).

Activate Staged Begin →

Configuration Filenames: Active - aws-ct800-testbed.zip Staged - aws-ct800-testbed.zip

Site-Appliance	Model	State	Currently Active		Currently Staged		Traffic Interruption		Download Package
			Software	Config	Software	Config	Expected	Actual	
NCN-ct800-ncn	CT800		R4_2_QA_D1_11132014	19:31 on 11/21/14	R4_1_DeveloperSpecific_P1_10012014	22:07 on 10/2/14	<3 min	76 s	active / staged
AU-ct800-AU	CT800		R4_2_QA_D1_11132014	19:31 on 11/21/14	R4_1_DeveloperSpecific_P1_10012014	22:07 on 10/2/14	<3 min	60 s	active / staged
BRA-ct800-BRA	CT800		R4_2_QA_D1_11132014	19:31 on 11/21/14	R4_1_DeveloperSpecific_P1_10012014	22:07 on 10/2/14	<3 min	62 s	active / staged
Client-ct800-client	CT800		R4_2_QA_D1_11132014	19:31 on 11/21/14	R4_1_DeveloperSpecific_P1_10012014	22:07 on 10/2/14	<3 min	67 s	active / staged

From the above, the expected traffic interruption time for the software update was less than 3 minutes for each Site, but that the actual interruption time was less than 100 ms. The user traffic would have been impacted for a bit longer than this as Conduits between the Sites must be operational or in the “GOOD” state before user traffic can flow across them.

9

Release 4.3 Features

This chapter includes features and enhancements released in 4.3.

Configure Private MPLS WAN Links

This feature simplifies creating Oracle SD-WAN Edge configurations when adding a Multiprotocol Layer Switching (MPLS) WAN Link. Previously, users created a WAN Link for each MPLS queue. Each WAN Link required a unique Virtual IP Address (VIP) to create the WAN Link and a unique Differentiated Services Code Point (DSCP) tag corresponding to the provider's queuing scheme. Once users defined a WAN Link for each MPLS queue, they defined the Intranet Service to map to a specific queue.

In Oracle SD-WAN Edge 4.3, a new MPLS specific WAN Link definition (i.e., Access Type) is available. Once the user selects the new Access Type, **Private MPLS**, they can define MPLS queues associated with the WAN Link. This allows users to define a single VIP with multiple DSCP tags that correspond to the provider's queuing implementation for the MPLS WAN Link. This also allows users to map the Intranet Service to multiple MPLS Queues on a single MPLS WAN Link.



Note:

If you have existing MPLS configurations and would like to implement the **Private MPLS** Access Type, please contact Oracle Support for assistance.

The high-level steps to configure this enhancement from within the Oracle SD-WAN Edge Configuration Editor are:

1. Define the WAN Link Access Type as **Private MPLS**.
2. Define the MPLS Queues corresponding to the Service Provider MPLS queues.
3. Enable the WAN Link for Conduit Service (enabled by default for Private MPLS WAN Links).
4. From **Conduit** ▾ **WAN Link**, assign an Autopath group.



Note:

If the Autopath Group is assigned from the WAN Link level, Oracle SD-WAN Edge will build paths automatically between the NCN and Client MPLS Queues based on matching DSCP tags. If the Autopath Group is assigned from the MPLS Queue level, Oracle SD-WAN Edge will build paths automatically regardless of whether the DSCP tags match.

1. Ensure that the same Autopath Group is configured at the NCN and Client.

2. Verify that the Paths for the WAN Link are built automatically.
3. Assign Intranet Service to a specific queue if needed.

 **Note:**

The Oracle configuration may not have a one-to-one mapping for provider-based queues. This is based on specific deployment scenarios.

 **Note:**

You cannot create Autopath Groups between different Private Access Types. For instance, you cannot create Autopath Groups between a Private Internet Access Type and a Private MPLS Access Type.

Add Private MPLS WAN LINK

In the Oracle SD-WAN Edge Configuration Editor, once users click **+ (Add)** under **Sites** ▾ **[Site Name]** ▾ **WAN Links**, the **Add WAN Link** pop-up appears.

Figure 1 illustrates how to configure the new WAN Link Access Type of **Private MPLS**.

Selecting a WAN Link Access Type of **Private MPLS** defines the Basic Settings for the WAN Link. The configurable settings include the physical (permitted) rate for WAN Ingress and WAN Egress. The MPLS Queues cannot exceed the physical (permitted) rate values. There are also no Audit errors if the sum of all MPLS Queues is below the WAN Link physical rates. See Figure 2 for a screen shot of the defined MPLS Queues.

Figure 2: Defined MPLS Queues

Figure 2 illustrates the Basic Settings for a Private MPLS WAN Link. Under the Basic Settings, there is now a new MPLS Queues Tab. Users click **+ Add** to add specific MPLS Queues. These should correspond with the queues defined by the Service Provider.

Users must define the following attributes for the MPLS Queues option:

- **MPLS Queue Name**
- **DSCP Tag:** This setting should correspond to the Service Provider's DSCP tag setting for the queue.
- **Unmatched:** When enabled, any frames arriving that do not match defined tags within the configuration file are mapped to this queue and the bandwidth defined for this queue.

- **WAN Ingress Permitted Rate:** The amount of bandwidth that Oracle devices are permitted to use for upload, which cannot exceed the defined physical upload rate of the WAN Link.
- **WAN Egress Permitted Rate:** The amount of bandwidth that Oracle devices are permitted to use for download, which cannot exceed the defined physical download rate of the WAN Link.

When users expand the MPLS Queue definition (by clicking the +), additional options appear.

These options include:

- **Tracking IP Address:** WAN Link tracking address
- **Congestion Threshold:** The defined amount of time for congestion (in microseconds) after which the MPLS Queue will throttle packet transmission to avoid additional congestion. When congestion exceeds the set Threshold, Oracle will back off the sending rate.
- **Eligibility:** The MPLS Queue's eligibility to process specific classes of traffic. When eligibility is disabled for a specific class of traffic, that class of traffic is unlikely to route through the MPLS Queue unless network conditions require it.

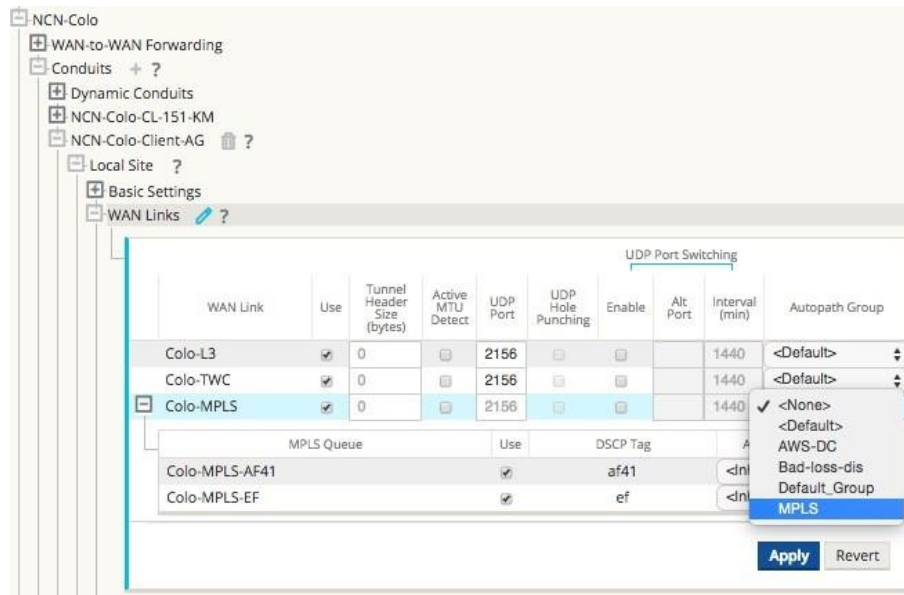
Users should configure the MPLS Queues that correspond to the existing Service Provider WAN Link queue definitions.

Since this is a new WAN Link Access Type, any existing MPLS WAN Links that are configured prior to Oracle SD-WAN Edge 4.3 are not impacted. Users should discuss migrating to the 4.3 features if desired, but migration is not required. Please contact your implementation team or support for additional recommendations when migrating to the new WAN Link Access Type.

Define WAN Link Basic Properties (Private MPLS)

Once the Private MPLS WAN Link with its MPLS Queues is defined, users should assign an Autopath Group to the WAN Link under a specific Conduit definition. Go to **Connections**

▢ [Site Name] ▢ **WAN Links** ▢ [MPLS WAN Link Name] ▢ **Conduits** ▢ [Conduit Name] ▢ [Local Site] ▢ **WAN Links** and click Edit (✎). Click the **Autopath Group** drop-down menu and choose from the available groups. By default, MPLS Queues inherit the Autopath Group assigned to the MPLS WAN Link. You may choose to set the individual MPLS Queues to **Inherit** the chosen **Autopath Group** or choose an alternate from the **Autopath Group** dropdown menu for each MPLS Queue. Figure 3 illustrates this process.



Figure

3: Conduit WAN Link Autopath Group Drop-Down Menu

Note:

If there is not a one-to-one mapping, based on DSCP tag, between queues at the local site and the remote site, users must map MPLS Queues to specific Autopath Groups. Inheriting an Autopath Group from the MPLS WAN Link will only automatically generate paths between queues with matching DSCP tags.

Assign Autopath Group to Conduit-WAN Link

Figure 4 and Figure 5 illustrate that the Autopath Group defined is the same for the NCN and Client appliance. This allows the system to build the Paths automatically. Figure 5 displays the NCN and Client settings. At the NCN site users can also expand the WAN Link associated with the conduit.

Verify Autopath Creation

Once users build the Path, the configuration is complete and they can follow the Change Management procedure to activate the configuration. Figure 6 illustrates the automatically generated paths.

View Permitted Rate and Congestion for WAN Links

In Oracle SD-WAN Edge 4.3, the Web console now allows users to view the permitted rate for WAN Links and WAN Link Usages and whether a WAN Link, Path, or Conduit may be in a congested state. In past Oracle SD-WAN Edge releases, this information was typically only available in Oracle SD-WAN Edge log files and via CLI-based commands. These options are now available in the Web console to assist users in problem isolation and troubleshooting.

View Permitted Rate

Permitted Rate is the amount of bandwidth that a particular WAN Link, Conduit Service, Intranet Service, or Internet Service is permitted to use at a given point in time. The permitted rate for a WAN Link is static, and is defined explicitly in the Oracle SD-WAN Edge configuration. The permitted rate for a Conduit Service, Intranet Service, or Internet Service will fluctuate over time, in response to congestion, user demand, and Fair Shares, but will always be greater than or equal to the Minimum Reserved Bandwidth for the Service.

Go to **Monitor** ▾ **Statistics** and select **WAN Link Usage** from the **Show** drop-down menu to display the page in Figure 7 including the **Permitted Kbps** information.

Under **Local WAN Links** users can see the configured permitted rates for each WAN Link defined at the local site. For example, RJS-NCN-WL2 has a configured permitted rate of 100Mbps. The **Usages and Permitted Rates** table displays the actual, real-time permitted rates for each Service that the WAN Link is used for (individual Conduit Services and combined Internet-Intranet Services). This information can assist users in troubleshooting a specific WAN Link problem or throughput associated with a WAN Link.

View Congestion

Oracle SD-WAN Edge 4.3 enhances the Web console to display a WAN Egress congestion state if it occurs within the site. There are three states associated with congestion:

- **UNKNOWN:** The WAN Link, Path, or Conduit is down so there is no congestion state
- **NO:** The WAN Link, Path, or Conduit is not congested
- **YES:** The WAN Link, Path, or Conduit is congested

In addition to the Web console displaying a congested state, there are also event notifications that users can enable and that the appliance generates when the congested state occurs. The event options include:

- **WAN_LINK_CONGESTION:** A WAN Link has become congested or un-congested
- **USAGE_CONGESTION:** A Conduit has become congested or un-congested

 **Note:**

Congestion is detected if packets in the WAN are delayed more than 100ms from the expected time of arrival.

To view congestion on a WAN Link, go to **Monitor** ▾ **Statistics** and choose **WAN Link Usage** from the **Show** drop-down menu. See Figure 8 for a screen capture of a congested WAN Link. The congested WAN Link is highlighted in red with the Congestion state as **YES**.

To view congestion on a Path, go to **Monitor** ▾ **Statistics** and choose **Paths** from the **Show** drop-down menu. See Figure 9 for a congested Path.

Monitor / Statistics Talari Support

Talari Statistics

Show: **Paths** Enable Auto Refresh **5** seconds **Refresh** Show latest data.

Path Statistics

Filter: in **Any column** **Apply**

Show **100** entries Showing 1 to 4 of 4 entries

Num	From Link	To Link	Path State	Conduit State	Source Port	Destination Port	Discovered MTU	Latency BOWT	Statistical jitter (mS)	Packets Received	Packets Out of Order	Packets Lost %	kbps	Conduit Type	Congestion
1	CL1-3010-WL1	RJS-NCN-WL1	GOOD	GOOD	2156	2156	1488	2	108	42411	0	0.00	1688.41	Static	YES
2	CL1-3010-WL2	RJS-NCN-WL2	GOOD	GOOD	2156	2156	1488	2	2	231725	0	0.00	10017.02	Static	NO
3	RJS-NCN-WL2	CL1-3010-WL2	GOOD	GOOD	2156	2156	N/A	2	2	69515	0	0.00	3727.39	Static	NO
4	RJS-NCN-WL1	CL1-3010-WL1	GOOD	GOOD	2156	2156	N/A	2	2	335889	0	0.00	21576.51	Static	NO

Showing 1 to 4 of 4 entries

Bandwidth calculated over the last 137 seconds

Figure 9: Congested Path

As illustrated in Figure 9, congestion can occur in one direction of a Path. In the above Oracle SD-WAN Edge, congestion is occurring from CL1-3010-WL1 ▾ RJS-NCNWL1. So, congestion (WAN Egress) is occurring from the Client site to the NCN site. Figure 9 was taken at the Client site. The Oracle NCN uses Oracle encapsulation to notify the Client appliance that it has congestion on the defined Path above. With this information, the Client appliance can also display that congested state.

To view congestion on a Conduit, go to **Monitor** ▾ **Statistics** and choose **Conduit** from the **Show** dropdown menu.



Note:

If a WAN Link, Path, or Conduit is congested, the Web console will display this state while the congestion persists and for an additional 15 seconds after the congestion clears. This is to provide users the ability to troubleshoot the event when it occurs.

Configuration Versioning

To reduce the chance of an Oracle SD-WAN Edge misconfiguration from multiple users exporting changes from the Configuration Editor in Oracle SD-WAN Edge, Oracle introduced Configuration Versioning in release 4.3.

The Oracle SD-WAN Edge Configuration Editor applies versioning metadata to configuration packages when the user saves a configuration or exports a configuration to Change Management. Upon export to Change Management, Oracle SD-WAN Edge

detects whether the configuration the user is attempting to export is derived from the running configuration. If the configuration is not derived from the current, running configuration, Oracle SD-WAN Edge presents the warning illustrated in Figure 10.



Figure 10: Export

Configuration Dialog with Warning Message

The user can choose to proceed and overwrite the running configuration or cancel the Export.

Support for Installing User-Generated Certificates on Appliances

Currently, all major browsers present a warning screen to users when they attempt to access the Web console of an appliance for the first time stating that the SSL certificate is invalid. Some browsers allow the user to add an exception to avoid the warning in the future, but the exception is specific to the appliance, the workstation, and the browser. The certificate is invalid for two reasons:

- The identity on the certificate does not match the URL of the appliance (typically an IP address).
- An authority trusted by the user's system did not sign the certificate.

Oracle SD-WAN Edge 4.3 allows users to upload generated certificates to the Web console of the appliance.

The user should generate the certificate for the appliance's IP address and the appropriate Certificate Authority should sign the certificate prior to installation. If the certificate is generated properly, it will be trusted by the systems on the user's network.

 **Note:**

For User-Generated certificates, there is also a root certificate that is loaded into the user's Web browser.

To upload the certificate, log into the appliance and proceed to **Manage Appliance** ▢ **HTTPS Certificate**. Users can upload a certificate and key file as required. There is no procedure to delete a certificate that was uploaded, but the user can regenerate a Oracle certificate by selecting **Regenerate HTTPS Certificate** as illustrated in Figure 11.

Figure 11: HTTPS Certificate

10

Release 4.4 Features

This chapter includes features and enhancements released in 4.4.

LAN GRE Tunnels

Oracle SD-WAN Edge 4.4 introduces LAN GRE Tunnels and allows you to configure Appliances to terminate GRE

Tunnels on the LAN. For example, in certain environments it may be advantageous to create a GRE Tunnel between a Appliance and a LAN side Linux host or router. This allows the Appliance to pass Conduit traffic into a GRE Tunnel terminated on the host or router for forwarding or processing. LAN GRE Tunnels can be used in the AWS environment where no Layer 2 support is available to simplify the deployment process.

To configure a LAN GRE Tunnel:

1. Log into your Appliance's web console.
2. Click on **Manage Network**, and then **Configuration**.
3. Open **Sites** → **[Site Name]**, and then **LAN GRE Tunnels** and click **+** to add a new tunnel.
4. Enter a **Name** and select a **Source IP** from the list of configured Virtual IPs.
5. Enter the tunnel's **Destination IP** and prefix (e.g., 10.4.0.20).
6. Click the **Checksum** checkbox if a checksum in the header is required.
7. Enter the **Keepalive Period** in seconds.

 **Note:**

If the Keepalive Period is set to **0**, no keepalive packets will be sent, but the tunnel will stay up even if the other end of the tunnel is unreachable.

1. Enter the number of **Keepalive Retries**.

 **Note:**

This is the number of times that the Appliance sends keepalive packets without a response before it brings the tunnel down.

1. Click **Apply**.



Note:

If the packet size including the GRE header exceeds the MTU and if "don't fragment" is set in the IP header, the packet will be dropped, but if "don't fragment" is **not** set in the IP header, the packet will be fragmented.



Figure 1: Create a LAN GRE Tunnel

Monitor LAN GRE Tunnels

To monitor configured LAN GRE Tunnels, go to **Monitor**, and then **Statistics** and choose **LAN GRE Tunnel** from the **Show** drop-down menu.



Figure 2: LAN GRE Statistics

IPsec Encryption in Conduit

Oracle SD-WAN Edge 4.4 introduces the ability to secure Conduit user data with IPsec encapsulated by the Oracle Reliable Protocol (TRP). This is executed using a 140-2 Level 1 FIPS-certified IPsec cryptographic library using Suite B algorithms and 256-bit ECP.




Figure 3: IPsec Encryption in a Conduit To implement IPsec Encryption on a Conduit:

1. Log into your Appliance's web console.
2. Click on **Manage Network**, and then **Configuration**.
3. Open **Connections** → **Default Sets** → **Conduit Default Sets** → [Site Name], and then **IPsec Settings**.
4. Click the edit icon (✎) and click the checkbox next to **Secure Conduit User Data with IPsec** to enable IPsec on the conduit.
5. Choose **ESP**, **ESP+Auth**, or **AH** from the **Tunnel Mode** drop-down menu.
6. If the Tunnel Mode is ESP or ESP+Auth, choose **AES 128-bit** or **AES 256-bit** from the **Encryption Mode** drop-down menu.
7. If the Tunnel Mode is AH, choose **SHA1** or **SHA-256** from the **Hash Algorithm** drop-down menu.
8. Click **Apply**.



Figure 4: Configure IPsec on a Conduit

Monitoring IPsec

To monitor Conduits secured with IPsec go to **Monitor**, and then **Statistics** and choose **Conduit** from the **Show** drop-down menu. The **IPsec Tunnel State** column indicates whether a Conduit's IPsec tunnel state is **GOOD**, **DEAD**, or **NEG** (tunnel is being negotiated).



Figure 5: IPsec Tunnel State

Path State Configurability and Monitoring

Oracle SD-WAN Edge 4.4 allows users to control when a Path is marked bad and how long a bad Path is kept on probation. It also gives users greater visibility into why a bad or dead Path is in its current state.

1. Oracle SD-WAN Edge 4.4 gives you greater control of the Bad Loss Sensitivity feature with the ability to manually configure the threshold for loss before a Path is marked BAD.
2. Now you can manually configure Silence Period, or the time that must elapse before a Path is marked BAD after packets are determined to be overdue.
3. Now you can also manually configure the Path Probation Period, or the time that must elapse before a Path is marked GOOD, after the symptom (e.g., loss or silence) clears.

The following new attributes are configurable:

- **Bad Loss Sensitivity:** The Path state transitions from **GOOD** to **BAD** when a specified amount of loss is observed. When set to **On**, loss is evaluated based on an internal formula. When set to **Custom**, loss is evaluated based on user-defined threshold. When set to **Off**, loss does not affect Path state.
- **Percent Loss Over Time:** Designate the Percent Loss (via drop-down menu) that is tolerable Over Time (via drop-down menu). Together these attributes establish what percentage of loss is tolerable over a specified period of time. Once exceeded, the Path State transitions from GOOD to BAD.
- **Silence Period (ms):** The Path state transitions from **GOOD** to **BAD** when no packets are received within the specified amount of time.
- **Path Probation Period (ms):** The probation period before changing the Path state from **BAD** to **GOOD**, after the symptom (e.g., loss or silence) clears.

To configure Bad Loss Sensitivity:

1. Log into your Appliance's web console.
2. Click on **Manage Network**, and then **Configuration**.
3. Open **Connections** → **Autopath Groups** → **[Autopath Group Name]**.

- Click the edit icon and choose **Custom** from the **Bad Loss Sensitivity** drop-down menu

Figure 6: Configuring New Bad Loss Sensitivity Attributes

- Choose a **Percent Loss** and a corresponding **Over Time** amount from their respective dropdown menus to establish a threshold for loss over time before the Path state will transition from GOOD to BAD.
- Define a **Silence Period** (in milliseconds) from the drop-down menu. If no packets are received within the Silence Period you designated, the Path state will transition from GOOD to BAD.
- Define a **Path Probation Period** (in milliseconds) from the drop-down menu. After the symptom (e.g., loss or silence) clears, the Path state will transition from BAD to GOOD once the Path Probation Period you designate here elapses.
- Click the checkbox next to **Instability Sensitive** if you want latency penalties and spikes considered in the Path scoring algorithm.
- Click **Apply**.

Monitor Statistics

The **Monitor**, and then **Statistics** default **Paths** screen was updated and renamed to increase usability. The **Source Port**, **Destination Port**, **Discovered MTU**, **Packets Received**, and **Packets Out of Order**, columns were moved to the new **Paths (Advanced)** screen (see Figure 8) to save space. Thus, the new, default **Paths (Summary)** screen is streamlined and easier to read.



Figure 7: Paths (Summary) Screen

On the new **Monitor**, and then **Statistics**, and then **Paths (Advanced)** screen, some column headings have been condensed to make the screen more readable.

The following header names have changed:

- Congestion is now **Cong**
- Source Port is now **Src Port**
- Destination port is now **Dst Port**
- Discovered MTU is now **MTU**
- Statistical Jitter (mS) is now **Jitter (mS)**
- Packets Out of Order is now **OOO**



Figure 8: New Paths (Advanced) Screen

The following new columns were also added to the **Paths (Advanced)** screen:

- **Reason:** This column indicates why a Path is marked BAD or DEAD.

- **Path State Duration:** This column indicates how long a Path has been in its current state.



Figure 9: Reason and Duration

Availability Reports

The Availability Reports screen (**Monitor**, and then **Availability Reports**) was reorganized and updated for both usability and readability.



Figure 10: Availability Reports

The **Incidents** column was enhanced to include a cluster of sub-columns that convey the following information:

- **Total:** The number of times a Path or Conduit has transitioned to a BAD or DEAD state.
- **Loss:** The number of times a Path has been marked BAD due to packet loss.
- **Silence:** The number of times a Path has been marked BAD or DEAD due to packet silence (i.e., no packets are received).
- **Peer:** The number of times a Path has been marked BAD or DEAD because the remote site indicated it was BAD or DEAD.

The **Badtime** column was enhanced to include a cluster of sub-columns that convey the following information:

- **Total:** The total time a Path or Conduit has been in a BAD state.
- **Loss:** The total time a Path has been marked BAD due to packet loss.
- **Silence:** The total time a Path has been marked BAD due to packet silence (i.e., no packets are received).
- **Peer:** The total time a Path has been marked BAD because the remote site indicated it was BAD.

The **Downtime** column was enhanced to include a cluster of sub-columns that convey the following information:

- **Total:** The total time a Path or Conduit has been marked DEAD.
- **Silence:** The total time a Path has been marked DEAD (i.e., no packets are received).

- **Peer:** The total time a Path has been marked DEAD because the remote site indicated it was DEAD.

Additional changes to Availability Reports:

- Path and Conduit information in Availability Reports is also included in the Periodic Status Reports when they are configured to include the Path or the Conduit in the report.

Additional Enhancements

The following enhancements were also rolled into Oracle SD-WAN Edge 4.4:

- Perform Diagnostic Dumps on Remote Client Appliances from the NCN

If you can still reach a Client via the NCN but cannot access that Client via its management port, the new `background_diagnostics` command allows administrators to perform a diagnostic dump on the remote Client via the NCN command line. When administrators execute the `tcon debug` command then execute the `remote_cmd [Remote`

`Site] ``tcon background_diagnostics`` (e.g., remote_cmd Omaha-CL1 ``tcon background_diagnostics``) command from a debug shell, a diagnostic dump will be started as a background process at the remote site specified.`

Administrators can track the execution of the current remote diagnostics operation and avoid spawning separate diagnostic dumps using `remote_cmd [Remote Site] ``tcon background_diagnostics_status``` (e.g., `tcon remote_cmd Omaha_CL1 ``tcon background_diagnostics_status```). This way administrators will know when the remote

diagnostic dump is complete.



Note:

The `background_diagnostics` and `background_diagnostics_status` commands can only be executed from within the debug shell.

- Export Authentication Logs to Syslog

When syslog is available for an appliance, administrators can now navigate to **Integrate** →

Configure Events and Alerts and click the **Authentications to Syslog** checkbox to forward user login events to a remote syslog server.



Figure 11: Syslog Settings

- Route Serviceability Enhancements

The **Routes** screen under **Monitor**, and then **Statistics** now displays the **Maximum allowed routes** as well as the routes in use. It also houses a new **Purge dynamic routes** button that you can use to clear dynamic routes and refresh the route table if you suspect it is corrupted.



Figure 12: Changes to Routes Screen

Appliance T5200 Support

Oracle SD-WAN Edge 4.4 introduces support for the new Appliance T5200. The T5200 is the first Appliance with 10G fiber connectivity. The 2U rack-mountable T5200 delivers up to 3Gbps across eight public and 32 private WAN Links and connects up to 256 sites, all of which you can manage through the Oracle SD-WAN Edge 4.4 interface. Refer to the *Appliance T5200 Hardware Guide* for additional details.

Oracle Virtual Appliance VT500 Support

Oracle SD-WAN Edge 4.4 introduces support for the new Oracle Virtual Appliance VT500. The VT500 is the first Oracle Virtual appliance built to run on VMware's vSphere virtual server environment. It delivers up to 40Mbps across three public and 32 private WAN Links and connects up to eight sites that you can manage through the Oracle SD-WAN Edge 4.4 interface. Refer to the *Oracle Virtual Appliance VT500 Getting Started Guide* for additional details.

Release 5.0 Features

This chapter includes features and enhancements release in 5.0.

Enhanced Match Criteria for Rules

In support of the new Virtual Routing and Forwarding (VRF) feature set, Edge 5.0 has significantly enhanced the criteria for Rule matching. The VLAN ID and the newly introduced Routing Domain may be used as match criteria for Rules in addition to the previously supported match criteria.

- **Routing Domain:** You can now choose one of the available, configured Routing Domains from the drop-down menu when creating a new Rule.



Figure 1: Routing Domain Drop-Down Menu

- **VLAN ID:** You can now enter one of the configured VLAN IDs in the VLAN ID field when creating a new Rule.



Figure 2: VLAN ID Field

Virtual Routing and Forwarding (VRF)

Edge 5.0 introduces Virtual Routing and Forwarding (VRF) to empower network administrators by giving them tools to segment their network for additional security and manageability. You can now separate guest network traffic from employee traffic, create distinct routing domains to segment large corporate networks, support multiple tenants at a Client Site, and segment traffic to support multiple customer networks.

Here are the touch points in Edge 5.0 to add, configure, and use Routing Domains to control and segment network traffic. Routing Domains can be used with both the Open Shortest Path First (OSPF) and Interior Border Gateway Protocol (IBGP) protocols.

In the Configuration Editor under **Global**, and then **Routing Domains** click Add (+) and enter a Name for your new Routing Domain. If you want to default to this Routing Domain, click the Default checkbox. Click Apply to save the changes. If you plan to implement a single Routing Domain, no explicit configuration is required. All new configurations are automatically populated with a default Routing Domain.




Figure 3: Add a Routing Domain

Under **Sites**, and then **[Client Site Name]**, and then **Routing Domains** click the Enable checkbox to enable a configured Routing Domain for the Site. Click the Default checkbox to make that Routing Domain the default for the Site. Click Apply to save the changes.

Note: Unchecking Enable for a Routing Domain will make it unavailable for use at the Site.

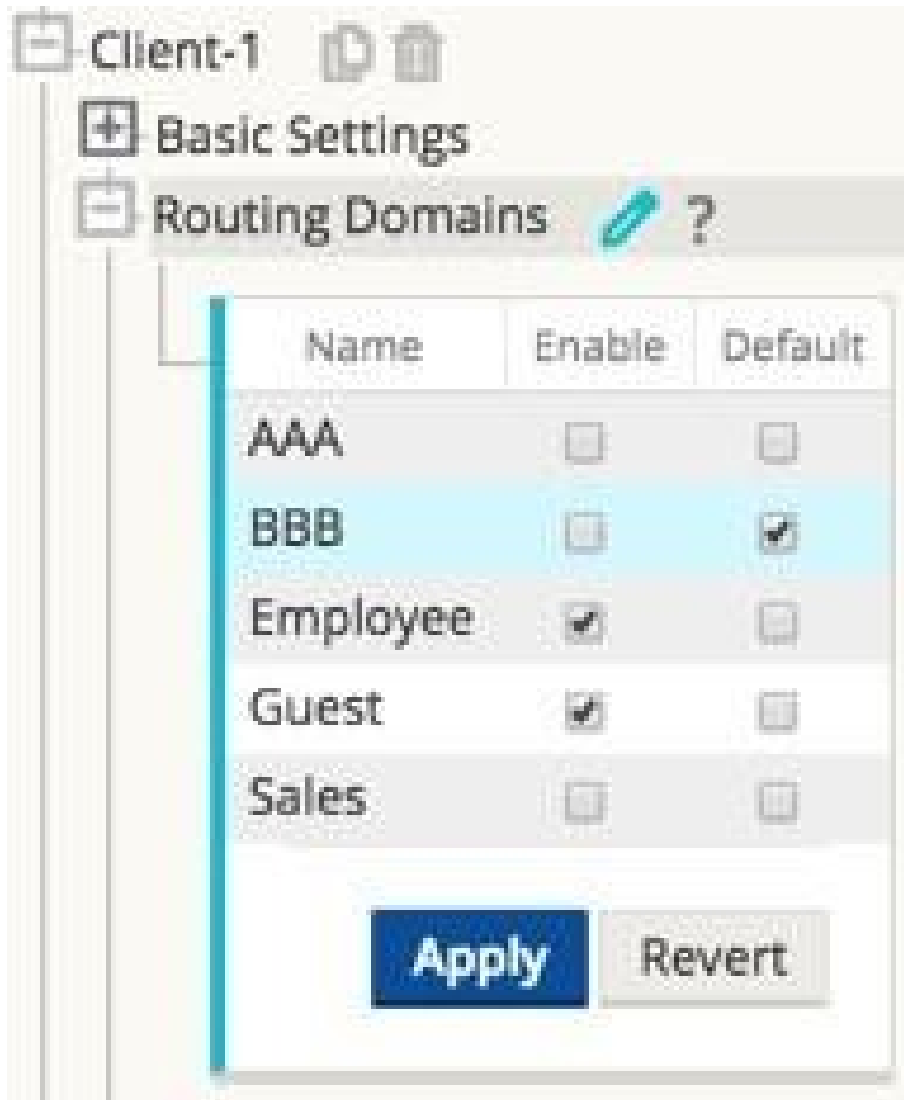


Figure 4: Site Level Routing Domain Management

Under **Sites**, and then **[Client Site Name]**, and then **Interface Groups** choose a Routing Domain from the drop-down menu when configuring Virtual Interfaces.

Note: Once Virtual Interfaces are associated with a specific Routing Domain, only those interfaces will be available when using that Routing Domain.

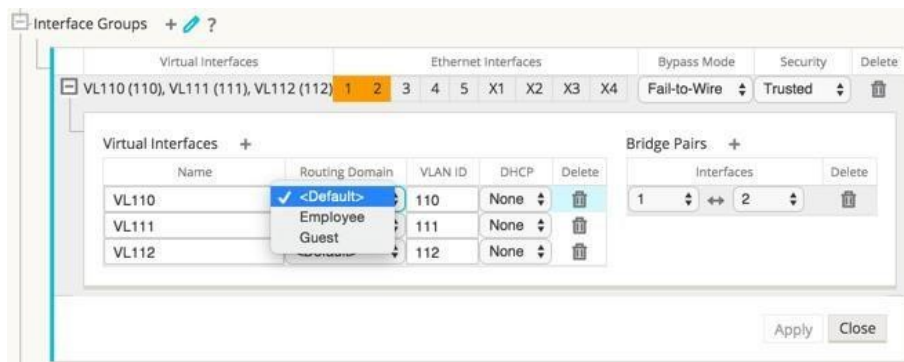


Figure 5: Configuring Interface Groups

From **Sites**, and then **[Client Site Name]**, and then **Virtual IP Addresses** choose a Routing Domain from the drop-down menu when configuring Virtual IP Addresses. The Routing Domain you choose determines which Virtual Interfaces are available from the drop-down menu.

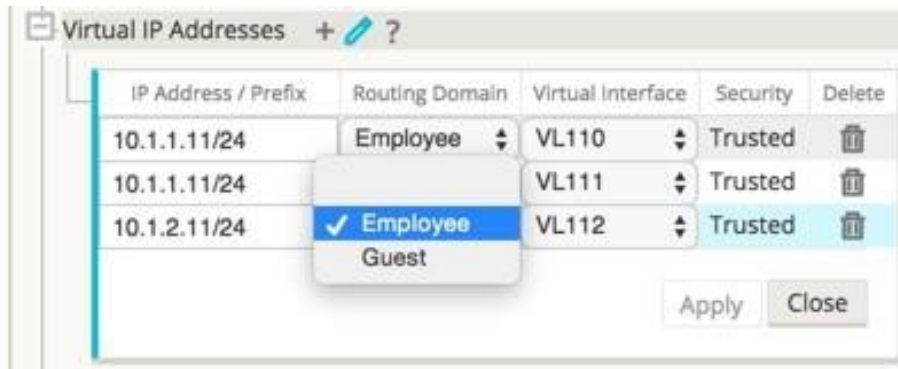


Figure 6: Configure Virtual IP Addresses

Under **Sites**, and then **[Client Site Name]**, and then **LAN GRE Tunnels** choose a Routing Domain from the drop-down menu when configuring a LAN GRE Tunnel. The Routing Domain you choose determines which Source IP Addresses are available from the drop-down menu.



Figure 7: Configure a LAN GRE Tunnel

From **Sites**, and then **[Client Site Name]**, and then **WAN Links**, and then **[WAN Link Name]**, and then **Access Interfaces** choose a Routing Domain from the drop-down menu when configuring an Access Interface. The Routing Domain you choose determines which Virtual Interfaces are available from the drop-down menu.

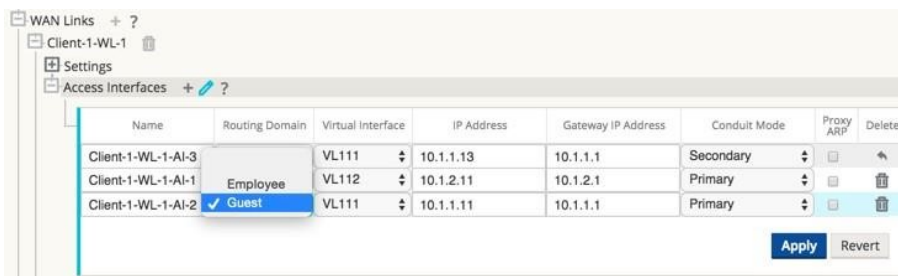


Figure 8: Configure an Access Interface

Under **Connections**, and then **[Site Name]**, and then **Intranet Services**, and then **[Intranet Service Name]**, and then **Basic Settings** click the Edit (✎) icon. Choose a Routing Domain from the drop-down menu.




Figure 9: Select Routing Domain for Intranet Service

From **Connections**, and then **[Site Name]**, and then **Intranet Services**, and then **[Intranet Service Name]**, and then **Rules** choose a Routing Domain from the drop-down menu.




Figure 10: Configure a Rule

From **Connections**, and then **[Site Name]**, and then **IPsec Tunnels** when you choose LAN as the Service Type, choose a Routing Domain from the drop-down menu. The Routing Domain will determine which Local IP Addresses are available.

Note: If the Service Type is Intranet, the Routing Domain is pre-determined by the chosen Intranet Service.




Figure 11: Configure IPsec Tunnel

From **Connections**, and then **[Site Name]**, and then **Routes** choose a Routing Domain from the drop-down menu. New Routes are automatically associated with the default Routing Domain.

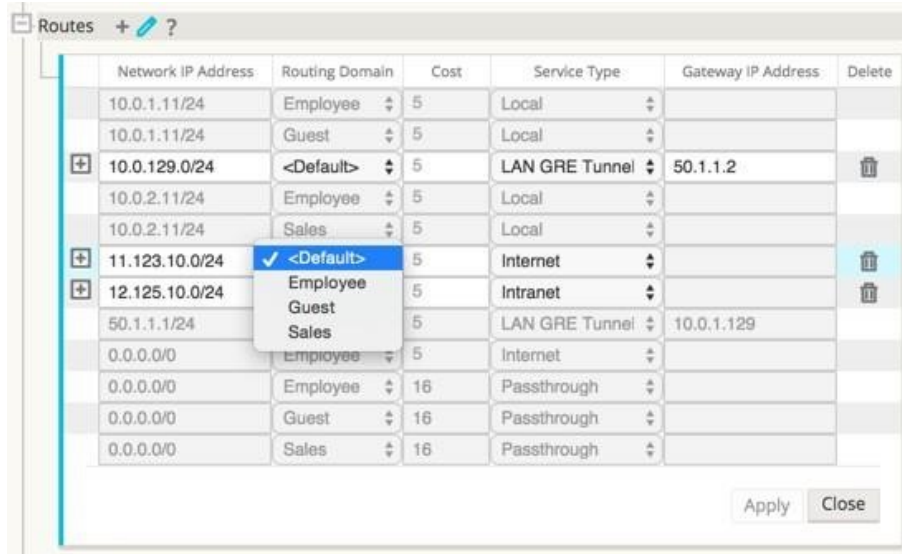


Figure 12: Configure a Route

Monitoring

From the Oracle Talari Appliance's home page, Routing Domain names are displayed in the System Status area of the screen.



Figure 13: Appliance Home Page

Under **Monitor**, and then **Statistics**, Routing Domain information is displayed, and results can be filtered by Routing Domain for the following criteria:

- Access Interfaces
- WAN Links
- MPLS Queues
- Intranet Services
- ARP
- LAN GRE
- Tunnels
- IPsec
- Routes
- Flows

WAN Link Statistics

Filter: Employee in Routing Domain Apply

Show 100 entries Showing 1 to 4 of 4 entries (filtered from 15 total entries)

WAN Link	Access Interface	Routing Domain	IP Address	Proxy Address	Proxy ARP State	MAC	Last ARP Reply Age (ms)
Client-1-WL-2 (standby)	N/A	Employee	10.1.1.11	11.11.11.11	N/A	N/A	N/A
Client-2-WL-2	N/A	Employee	10.2.1.11	20.20.20.20	N/A	N/A	N/A
GEO-Site-WL-2	N/A	Employee	10.100.2.11	7.7.7.7	N/A	N/A	N/A
NCN-Primary-WL-2 (standby)	NCN-Primary-WL-2-AI-1	Employee	10.0.2.11	3.3.3.3	DISABLED	N/A	N/A

Showing 1 to 4 of 4 entries (filtered from 15 total entries)

Figure 14: WAN Link Statistics Filtered by Employee Routing Domain

Under **Manage Network**, and then **View Configuration**, wherever configuration information for the following attributes is displayed, the Routing Domain is also displayed:

- Sites
- WAN Links
- Intranet Services
- LAN GRE
- Tunnels
- IPsec
- Routes
- Flows

Route Configuration

Routes for routing domain 'Employee' :

Num	Network Addr	Gateway IP Address or Next_Hop	Service	Site	Cost	Type	Protocol	Neighbor Direct	Route Eligibility Type	Route Eligible Based on
0	10.0.2.11/32	*		NCN-Primary	5	Static	-	-	-	-
1	10.0.2.0/24	*	IPHost	NCN-Primary	5	Static	-	-	-	-
2	10.100.1.0/24	*	NCN-Primary-GEO-Site	GEO-Site	5	Static	-	-	-	-
3	10.100.2.0/24	*	NCN-Primary-GEO-Site	GEO-Site	5	Static	-	-	-	-
4	10.2.1.0/24	*	NCN-Primary-Client-2	Client-2	5	Static	-	-	-	-
5	10.2.2.0/24	*	NCN-Primary-Client-2	Client-2	5	Static	-	-	-	-
6	10.1.1.0/24	*	NCN-Primary-Client-1	Client-1	5	Static	-	-	-	-
7	10.1.2.0/24	*	NCN-Primary-Client-1	Client-1	5	Static	-	-	-	-
8	11.123.10.0/24	*	Internet	*	5	Static	-	-	-	-
9	0.0.0.0/0	*	Internet	*	5	Static	-	-	-	-
10	0.0.0.0/0	*	Passthrough	*	16	Static	-	-	-	-
11	0.0.0.0/0	*	Discard	NCN-Primary	16	Static	-	-	-	-

Dynamic Routing

Edge 5.0 introduces support for Dynamic Routing protocols. This feature enables your Oracle Talari Appliance to discover LAN subnets, advertise Conduit routes, work more seamlessly within networks using the Interior Border Gateway Protocol (IBGP) and Open Shortest Path First (OSPF) protocols, potentially eliminate redundant equipment (branch routers), and support graceful router failover.

Note: Edge 5.0 uses Interior BGP (IBGP) and OSPF as an Interior Gateway Protocol (IGP).

Virtual IP Address Identity

To use a Virtual IP Address for Dynamic Routing, go to **Sites**, and then **[Site Name]**, and then **Virtual IP Addresses**. Click the Identity checkbox for a Virtual IP Address to use it for IP services.

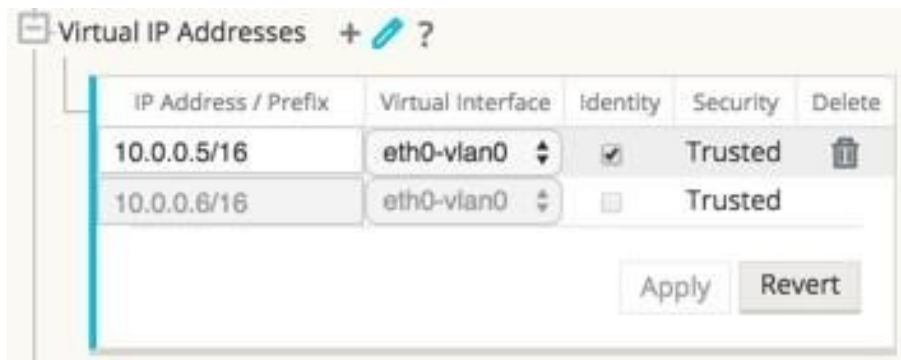


Figure 15: Virtual IP Address Identity

Open Shortest Path First (OSPF) Routing Protocol

OSPF is an interior routing protocol that uses a link state algorithm to exchange routing information between routers within a single routing domain (i.e., autonomous system). You can now configure Oracle Talari Appliances to learn routes and advertise routes using OSPF.

To configure OSPF:

1. Under **Connections**, and then **[Site Name]**, and then **Route Learning**, and then **OSPF**, and then **Basic Settings** click the Edit (✎) icon.
2. Click the Enable checkbox, enter an optional Router ID, click the Advertise APN Routes checkbox if you wish to advertise Routes, and click Apply to enable OSPF.

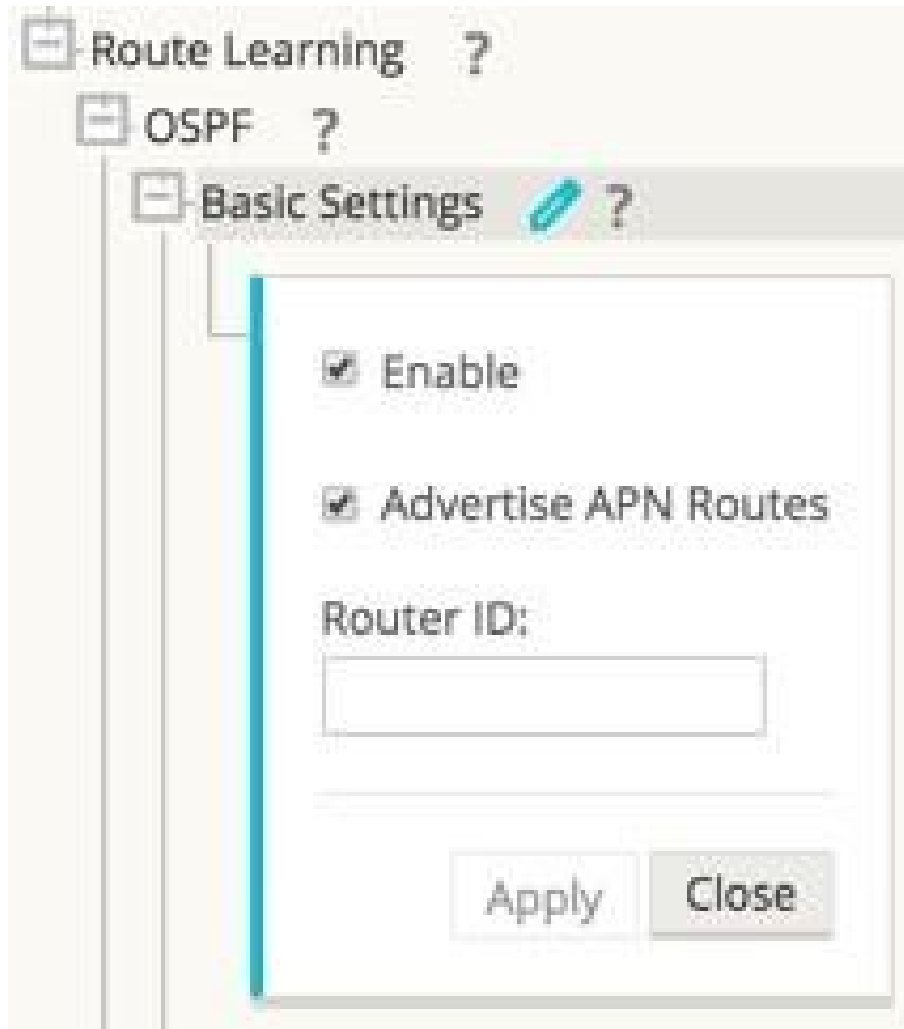


Figure 16: Enable OSPF

- Expand **OSPF > Areas** and click the Edit icon.



Figure 17: Configuring OSPF

- Enter an area ID to learn from and advertise to.
- For sites with multiple Routing Domains, under Virtual Interfaces choose a Routing Domain from the drop-down menu as illustrated in Figure 17. The Routing Domain determines which Virtual Interfaces are available.

Note: If there is only one Routing Domain configured, the Routing Domain column will not appear. If Identity is not checked for a specific Virtual IP Address (see the Virtual IP Address Identity section for more details), the associated Virtual Interface will not be available for IP services.

1. Choose one of the available Virtual Interfaces from the Name drop-down menu. The Virtual Interface will determine the Source IP Address.
2. Enter the Interface Cost (10 is the default).
3. Choose an Authentication Type from the drop-down menu.
4. If you chose Password or MD5 in step 8, enter the Password associated text field.
5. In the Hello Interval field, enter the amount of time to wait between sending Hello protocol packets to directly connected neighbors (10 seconds is the default).
6. In the Dead Interval field, enter the amount of time to wait to receive a Hello protocol packet before marking a router as dead (40 seconds is the default).
7. Click Apply to save your changes.

Interior Border Gateway Protocol (IBGP)

Interior BGP (IBGP) is an exterior routing protocol designed to exchange routing information between routing domains (i.e., autonomous systems). However, BGP may be used for routing within a domain. In this application, it is referred to as Interior BGP. You can now configure Oracle Talari Appliances to learn routes and advertise routes using Interior BGP.

To configure Interior BGP (IBGP):

1. Under **Connections**, and then **[Site Name]**, and then **Route Learning**, and then **IBGP**, and then **Basic Settings** click the Edit (✎) icon.
2. Click the Enable checkbox, click the Advertise APN Routes checkbox if you wish to advertise Routes, enter an optional Router ID, and enter the number of the Local Autonomous System to learn routes from and advertise routes to in the Local Autonomous System field. Click Apply to enable IBGP.

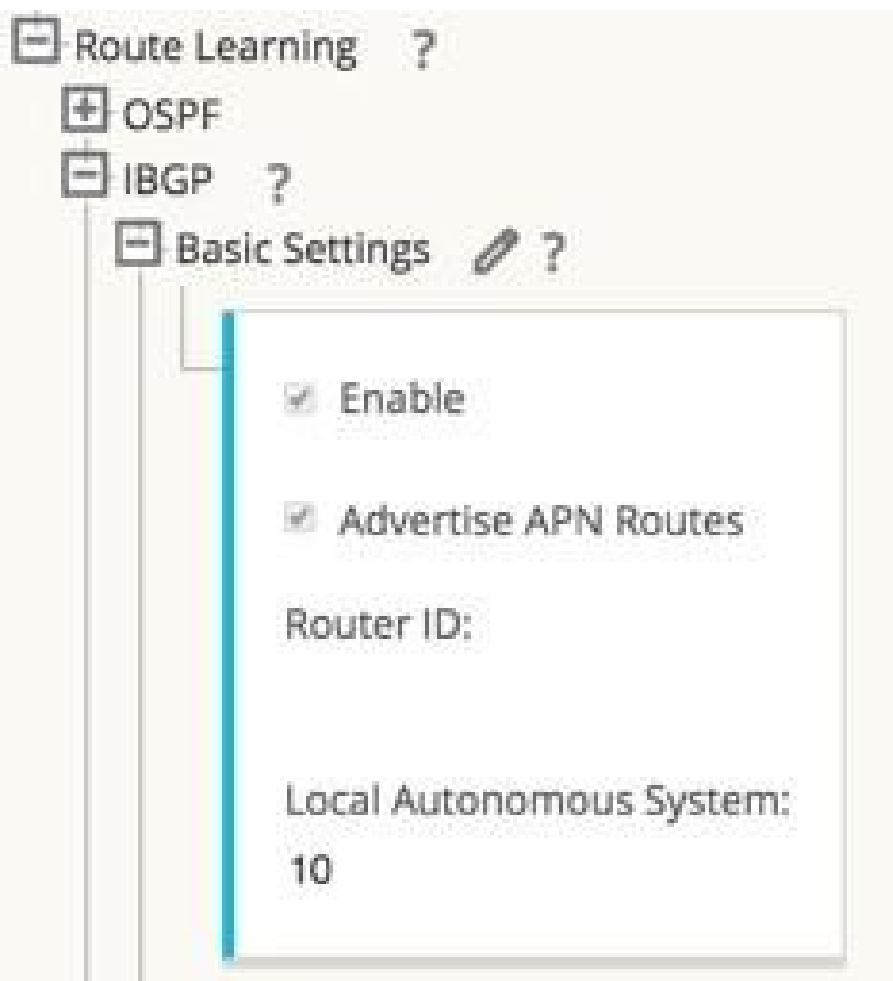


Figure 18: Enable IBGP

1. Expand **IBGP**, and then **Basic Settings**, and then **Neighbors** and click the Add (+) icon.



Figure 19: Add a Neighbor to a Site with a Single Routing Domain

Note: If there is only one Routing Domain configured, the Routing Domain column will not appear. If Identity is not checked for a specific Virtual IP Address (see the Virtual IP Address Identity section for more details), the associated Virtual Interface will not be available for IP services

1. For Sites with multiple Routing Domains, choose a Routing Domain from the drop-down. The Routing Domain determines which Virtual Interfaces are available.
2. Choose a Virtual Interface from the drop-down menu. The Virtual Interface will determine the Source IP Address.
3. Enter the IP Address of the IBGP Neighbor router in the Neighbor IP field.
4. In the Hold Time (s) field, enter the Hold Time, in seconds, to wait before declaring a neighbor down (the default is 180).

- In the Local Preference (s) field, enter the Local Preference value, in seconds, which is used for selection from multiple IBGP routes (the default is 100).
- Click the IGP Metric checkbox to enable the comparison of internal distances to calculate the best route.
- In the Password field, enter a password for MD5 authentication of IBGP sessions (authentication is not required).

Filters

Filters are used to import or exclude routes learned via OSPF and IBGP based on specific match criteria.

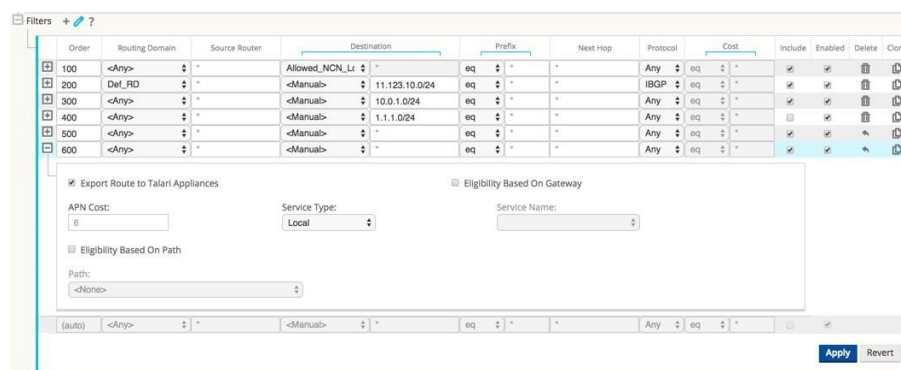


Figure 20: Add a Filter on a Site with Multiple Routing Domains

- Expand **Route Learning**, and then **Filters** and click the Add (+) icon.

Note: If there is only one Routing Domain configured, the Routing Domain column will not appear.

- Click the + next to your new Filter to expand the settings.
- You can use the following criteria to construct each Filter that you create.
 - Order: The Order in which filters are prioritized. The first filter that a route matches to will be applied to that route.
 - Routing Domain: To match routes from a specific routing domain, choose one of the configured Routing Domains from the drop-down menu.
 - Source Router: To match routes from a specific source router, enter the IP address of the Source Router.
 - Destination: To match routes by destination, choose Manual from the drop-down menu and enter an IP Address and Netmask in the adjacent field or choose from the list of available Network Objects.
 - Prefix: To match routes by prefix, choose a match predicate from the drop-down menu and enter a Route prefix in the adjacent field.
 - The predicates are:
 - Eq: Equal to
 - lt: Less than
 - le: Less than or equal to
 - gt: Greater than

- ge: Greater than or equal to
 - Next Hop: To match routes by next hop, enter the IP address of the Next Hop.
 - Protocol: To match routes by protocol, choose the protocol from the drop-down menu (Any, OSPF, or IBGP) to learn routes from.
 - Cost: If the protocol for your filter is OSPF, to match routes by cost, choose a match predicate from the drop-down menu and enter a route cost in the adjacent field.
 - The predicates are:
 - eq: Equal to
 - lt: Less than
 - le: Less than or equal to
 - gt: Greater than
 - ge: Greater than or equal to
 - Include: Click the checkbox to Include routes that match this filter. Otherwise matching routes are ignored.
 - Enabled: Click the checkbox to Enable this filter. Otherwise the filter is ignored.
 - Clone: Click the Clone icon to make a copy of an existing Filter.
 - Export Route to Oracle Talari Appliances: Click the checkbox to export matching routes to Oracle Talari Appliances at other Sites when WAN-to-WAN Forwarding is enabled. This functionality is enabled by default and only applies for the following Service Types: Local, LAN GRE Tunnel, and LAN IPsec Tunnel.
 - Eligibility Based on Gateway: Click the checkbox to ensure that a matching route is not used if its Gateway is unreachable.
 - Cost: Enter the cost that the Oracle Talari Appliance applies to matching routes (the default is 6).
 - Service Type: Select the Service Type (e.g., Local, Internet, Intranet, LAN GRE Tunnel, LAN IPsec Tunnel, or Passthrough) that will be assigned to matching routes.
 - Service Name: For Intranet, LAN GRE Tunnel, and LAN IPsec Tunnel, specify the name of the configured Service Type to use.
 - Eligibility Based on Path: Click the checkbox to ensure that a matching route is not used if a chosen Path is dead. Choose a Path from the list of available Paths on the drop-down menu below.
4. Once you have configured your filter, click Apply.

Network Objects

Edge 5.0 introduces Network Objects, a new option under the Global section in the Configuration Editor. Now you can group multiple subnets together, and reference a single Network Object when defining a Route Filter rather than creating a filter for each subnet.

1. If you plan to use Network Objects, navigate to **Global**, and then **Network Objects** click Add (+).
2. Click Add (+) under Networks.
3. Enter the IP Address and Subnet of the new Network Object.
4. Click Apply to save the settings.

- To edit the Network Object's name, double-click on the name of the Network Object and enter a new name.

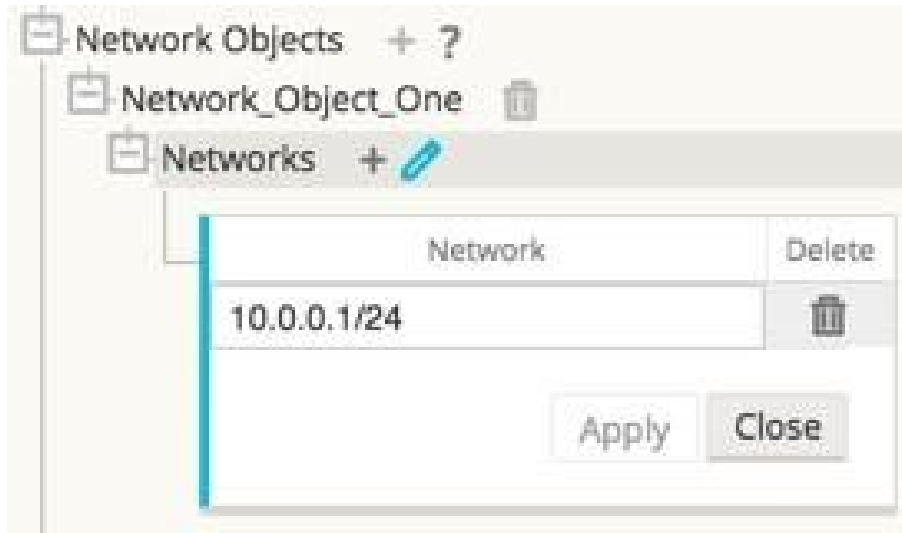


Figure 21: New Network Object

Monitoring

Under **Monitor**, and then **Statistics** all functions for Routes supported in Edge 4.4 are supported in Edge 5.0 regardless of whether a Route is Dynamic or Static.

Route Statistics
Maximum allowed routes: 16000

Routes for routing domain : CUSTA

Filter: in

Show entries Showing 1 to 5 of 5 entries

Num [#]	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
0	100.100.100.0/24	*	JM-NCN-JM-GEO	YES	*	JM-GEO	Static	-	-	5	0	YES	N/A	N/A
1	10.0.1.0/24	*	Local	YES	*	JM-NCN	Static	-	-	5	0	YES	N/A	N/A
2	10.0.1.0/24	*	Local	NO	*	JM-NCN	Dynamic	OSPF	-	6	0	NO	GW	*
3	0.0.0.0/0	*	Passthrough	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
4	0.0.0.0/0	*	Discard	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

Showing 1 to 5 of 5 entries

Figure 22: Route Statistics for Dynamic and Static Routes

WAN Link IP Address Learning (DHCP Client)

Edge 5.0 introduces WAN Link IP Address Learning via DHCP Clients. This functionality reduces the amount of manual configuration to deploy Oracle Talari Appliances and reduces customers' ISP costs by eliminating the need to purchase static IP Addresses. Now Oracle Talari Appliances can obtain dynamic IP Addresses for WAN Links on untrusted interfaces eliminating the need for an intermediary WAN router to perform this function or a static IP.

Note: DHCP Client can only be configured for Oracle Talari Appliances configured as Client Nodes.

To Configure DHCP for an Untrusted Virtual Interface, choose Client from the DHCP dropdown menu under **Sites**, and then **[Client Name]**, and then **Interface Groups**, and then **Virtual Interfaces**.

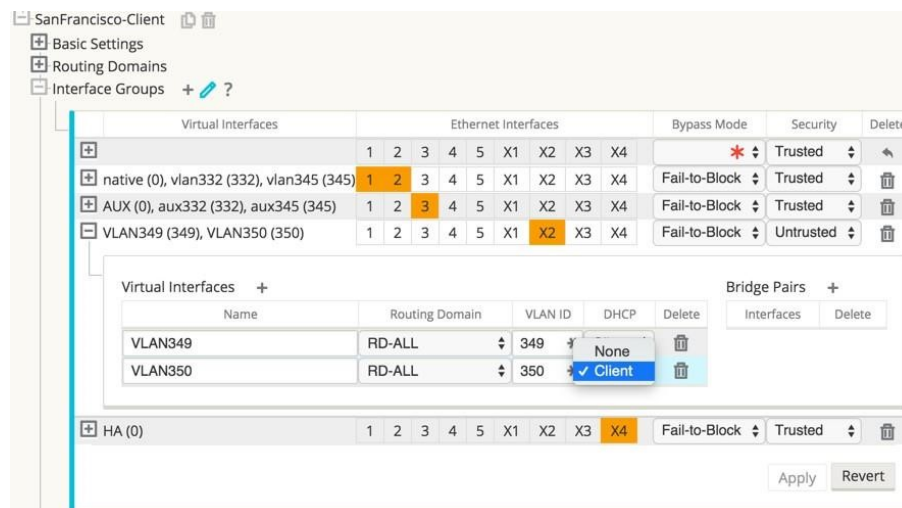


Figure 23: Enable DHCP for a Virtual Interface

Under **WAN Links**, and then **[WAN Link Name]**, and then **Settings**, and then **Basic Settings** click the **Autodetect Public IP** checkbox to enable the Network Control Node (NCN) to detect the Public IP Address to be used by the Public Internet WAN Link.

Note: This is required when DHCP Client mode is configured for the WAN Link.

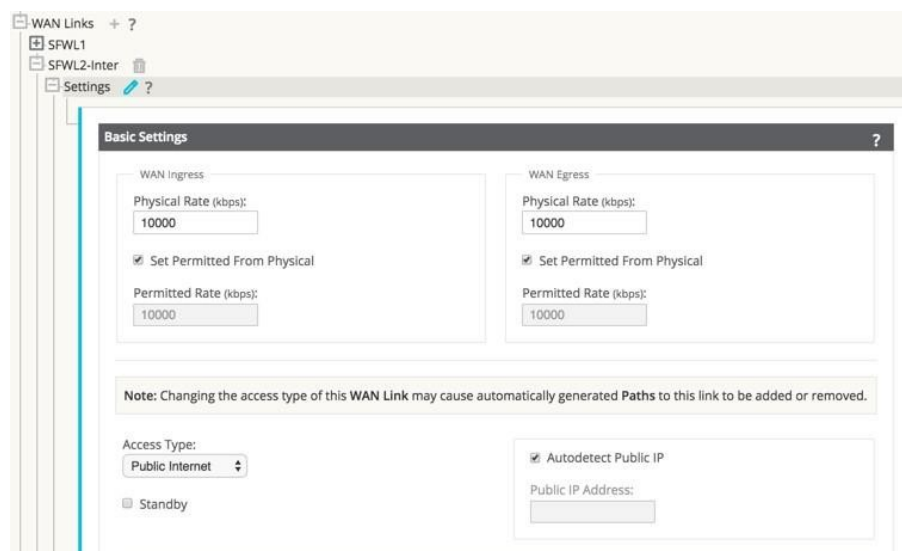


Figure 24: Autodetect Public IP

Monitoring

The runtime Virtual IP Address, Subnet Mask, and Gateway settings are logged as in 4.4. Events are generated when Dynamic Virtual IPs are learned, released, or expired; when there is a communication issue with the learned Gateway or DHCP server; or when duplicate IPs are detected. If duplicate IPs are detected at a Site, Dynamic Virtual IPs are released and renewed until all Virtual Interfaces at the site have unique Virtual IP Addresses.

Under **Manage Network > Enable/Disable/Purge Flows** the DHCP Client WAN Links table provides the status of learned IPs. From here you can request to **Renew** the IP, which will refresh the lease time. You can also choose to **Release & Renew**, which will get a new IP address with a new lease.

DHCP Client WAN Links									
Ethernet Interface	Virtual Interface	WAN Link	IP Address / Prefix	Gateway IP Address	Lease Duration Seconds	Remaining Seconds	Expiration Date	Action	
X2	VLAN349	SFWL3-Inter	10.30.30.55/24	10.30.30.2	1800	1640	9:13 on 1/8/2016	Renew	Submit
X2	VLAN350	SFWL4-Inter	10.20.20.53/24	10.20.20.2	86400	71035	4:29 on 1/9/2016	Renew	Submit

Figure 25: DHCP Client WAN Links

IPsec VPN Termination

Expanding on the IPsec in Conduit feature introduced in 4.4, Edge 5.0 release allows third party devices to terminate IPsec VPN Tunnels on the LAN or WAN side of Oracle Talari Appliances. Now you can secure site-to-site IPsec Tunnels terminating on an Oracle Talari Appliance using a 140-2 Level 1 FIPS certified IPsec cryptographic binary.

Note: Bandwidth provisioning is not available for LAN side IPsec VPN termination.

If you plan to implement Certificates for IKE negotiation, navigate to **Sites**, and then **Certificates** and add any necessary certificates.

Certificates ?			
Identity + ?			
Name	Fingerprint	Delete	
talari-id	0F:31:F3:4E:B2:D4:31:75:AD:70:AD:7D:D8:35:64:47:8A:D0:68:95		
Trusted + ?			
Name	Fingerprint	Delete	
talari-ca	36:D8:B7:5F:2A:BE:02:EB:5F:DE:45:B7:88:21:7F:60:59:6A:50:32		
talari-root	81:C3:1D:51:EA:59:DB:B7:BA:78:D1:D0:FF:2F:9D:35:46:D3:58:88		

Figure 26: Certificates for IKE Negotiation

Navigate to **Connections**, and then **[Site Name]**, and then **IPsec Tunnels** to create an IPsec Tunnel.

You can configure the following criteria, and click Apply to save your settings:

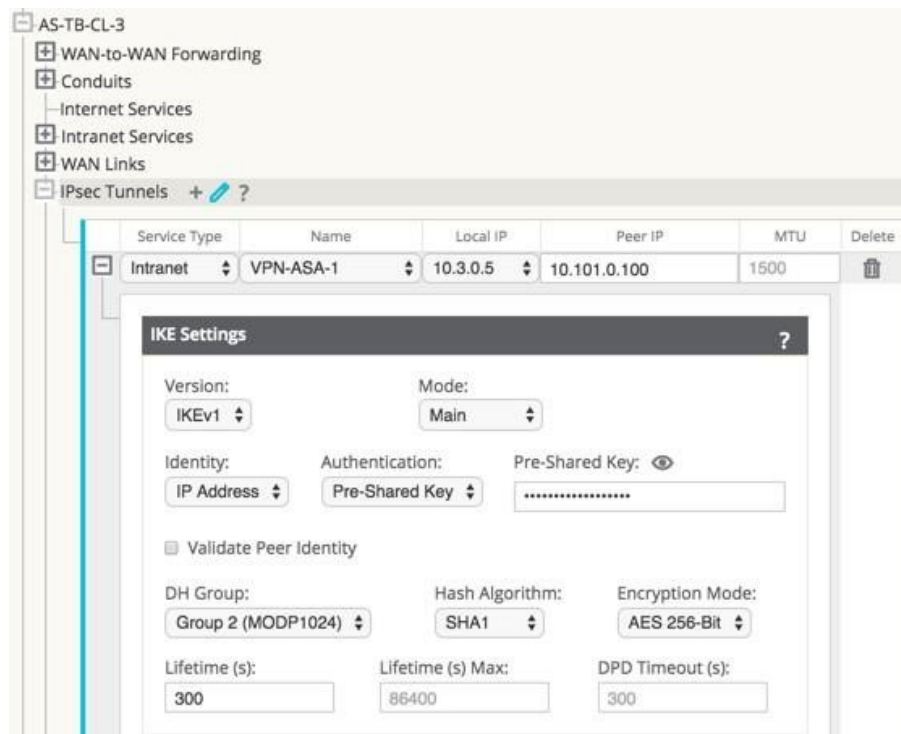


Figure 27: IPsec IKEv1 Settings

- Service Type: Choose either Intranet or VPN from the drop-down menu.
- Name: If the Service Type is Intranet, choose from the list of configured Intranet Services in the drop-down menu. If the service type is LAN, enter a unique Name.
- Local IP: Choose the Local IP address of the IPsec Tunnel from the drop-down menu of available Virtual IP Addresses configured at this Site.
- Peer IP: Enter the Peer IP address of the IPsec Tunnel.
- MTU: The default is 1500, but you can enter a different value.
- IKE Settings
 - Version: Choose either IKEv1 or IKEv2 from the drop-down menu.
 - Mode: For IKEv1, choose either Main or Aggressive from the Mode drop-down menu.
 - Identity: Choose either Auto or IP Address from the Identity drop-down.
 - Authentication: Choose either Pre-Shared Key or Certificate from the Authentication drop-down menu.
 - * Pre-Shared Key: If you are using a Pre-Shared Key, copy and paste it into this field. Click on the Eyeball (👁) icon to view the Pre-Shared Key.
 - * Certificate: If you are using an Identity Certificate, choose it from the drop-down menu.
 - Validate Peer Identity: Click the Validate Peer Identity checkbox to validate the IKE's Peer Identity. If the peer's ID type is not supported, do not enable this feature.
- DH Group: Choose the Diffie–Hellman group (Group 1, Group 2, or Group 5) to use for IKE key generation from the drop-down menu.

- Hash Algorithm: Choose MD5, SHA1, or SHA-256 from the drop-down menu to authenticate IKE messages.
- Encryption Mode: Choose AES 128-bit, AES 192-bit, or AES 256-bit as the Encryption Mode for IKE messages from the drop-down menu.
- Lifetime (s): Enter the preferred duration, in seconds, for an IKE security association to exist. The default is 3600 seconds.
- Lifetime Max (s): Enter the maximum preferred duration, in seconds, to allow an IKE security association to exist. The default is 86400 seconds.
- DPD Timeout (s): Enter the Dead Peer Detection timeout, in seconds, for VPN connections. The default is 300 seconds.
- IKEv2 Settings

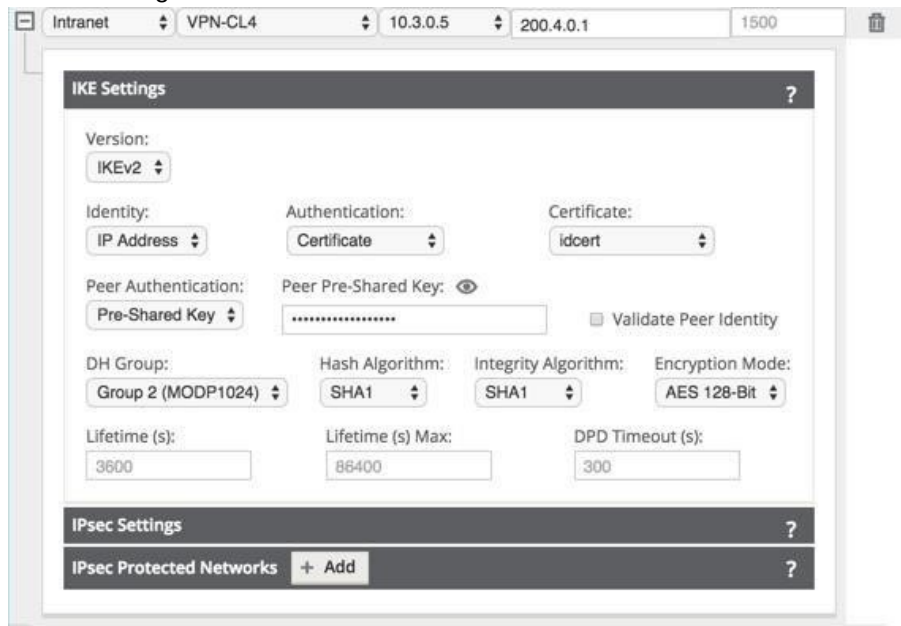


Figure 28: IPsec IKEv2 Settings

- Peer Authentication: Choose Mirrored, Pre-Shared Key, or Certificate Peer Authentication from the drop-down menu.
- Peer Pre-Shared Key: Paste the IKEv2 Peer Pre-Shared Key into this field for authentication. Click on the Eyeball (👁) icon to view the Pre-Shared Key.
- Integrity Algorithm: Choose MD5, SHA, or SHA-256 as the hashing algorithm to use for HMAC verification from the drop-down menu.
- IPsec Settings

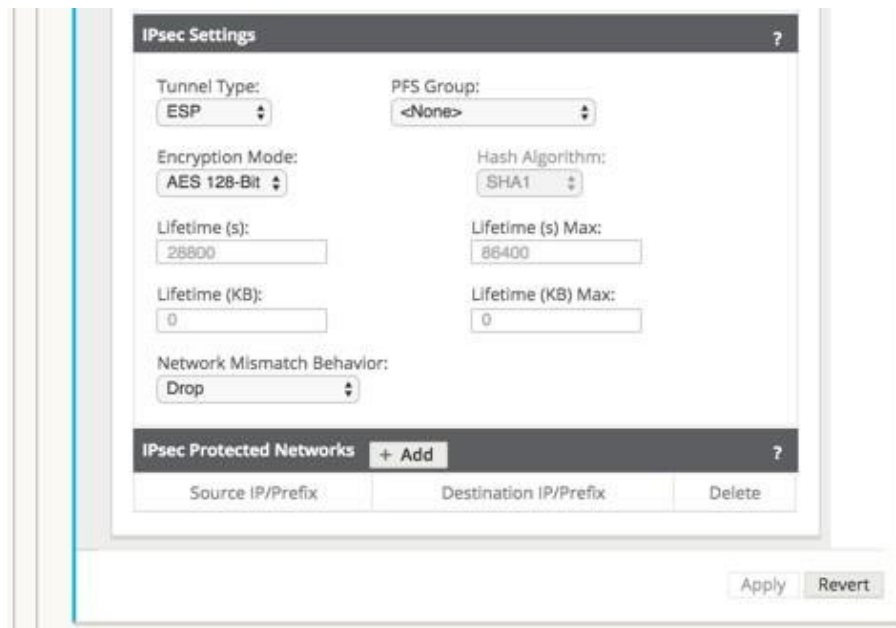


Figure 29: IPsec Settings and IPsec Protected Networks

- Tunnel Type: Choose ESP, ESP+Auth, or AH as the Tunnel Type from the dropdown menu.
- PFS Group: Choose the Diffie–Hellman group (Group 1, Group 2, or Group 5) to use for perfect forward secrecy key generation from the drop-down menu.
- Encryption Mode: If you chose ESP or ESP+ Auth, choose AES 128-bit, AES 192bit, or AES 256-bit as the Encryption Mode for IPsec messages from the drop-down menu.
- Hash Algorithm: If you chose ESP+Auth or AH as the Tunnel Type, choose MD5, SHA1, or SHA-256 from the Hash Algorithm drop-down menu to use for HMAC verification.
- Lifetime (s): Enter the amount of time, in seconds, for an IPsec security association to exist. The default is 28800 seconds.
- Lifetime Max (s): Enter the maximum amount of time, in seconds, to allow an IPsec security association to exist. The default is 86400 seconds.
- Lifetime (KB): Enter the amount of data, in kilobytes, for an IPsec security association to exist.
- Lifetime Max (KB): Enter the maximum amount of data, in kilobytes, to allow an IPsec security association to exist.
- Network Mismatch Behavior: Choose Drop, Send Unencrypted, or Use Non-IPsec Route as the desired action for your Talari WAN to take if a packet does not match the IPsec Tunnel's Protected Networks from the drop-down menu.
- IPsec Protected Networks
 - Source IP/Prefix: After clicking the Add (+ Add) button, enter the Source IP and Prefix of the network traffic the IPsec Tunnel will protect.
 - Destination IP/Prefix: Enter the Destination IP and Prefix of the network traffic the IPsec Tunnel will protect.

Monitoring

Under **Monitor**, and then **Statistics** when you choose IPsec Tunnel from the Show drop-down menu, you can see the following criteria:

- Tunnel Name
- State
- Service Type
- Packets Received
- Packets Sent
- Kbps Received
- Kbps Sent
- Packets Dropped
- Bytes Dropped
- MTU

Statistics

Show: **IPsec Tunnel** Enable Auto Refresh **5** seconds (refresh) Show latest data.

IPsec Tunnel Statistics

Filter: in **Any column**

Show **100** entries Showing 1 to 8 of 8 entries **1**

Name	State	Service Type	Packets Received	Kbps Received	Packets Sent	Kbps Sent	Packets Dropped	Bytes Dropped	MTU
AS-TB-NCN-AS-TB-CL-1	DEAD	Conduit	0	0	0	0	0	0	1359
AS-TB-NCN-AS-TB-CL-2	GOOD	Conduit	0	0	0	0	0	0	1359
AS-TB-NCN-AS-TB-CL-3	GOOD	Conduit	0	0	0	0	0	0	1359
AS-TB-NCN-AS-TB-CL-4	GOOD	Conduit	0	0	0	0	0	0	1359
VPN-ASA-1	DEAD	Intranet	0	0	0	0	0	0	1427
VPN-ASA-2	DEAD	LAN	0	0	0	0	0	0	1377
VPN-PaloAlto	DEAD	Intranet	0	0	0	0	0	0	1439
VPN-SonicWall	DEAD	Intranet	0	0	0	0	0	0	1456

Showing 1 to 8 of 8 entries **1**

Figure 30: IPsec Tunnel Monitoring

Under **Manage Network**, and then **View Configuration** when you choose IPsec Tunnel from the Show dropdown menu, you can view the IPsec Tunnel configuration:

```
IPsec Tunnel Configuration
=====
Name: VPN-ASA-1
=====
ipsec_service_type=intranet
ike_local_ip_addr=10.0.0.6
ike_remote_ip_addr=10.101.0.100
network_mtu=1500
ike_version=2
ike_auth=psk
ike_identity=auto
ike_peer_auth=cert
ike_validate_peer_identity=1
ike_hash_algorithm=sha256
ike_integ_algorithm=sha256
ike_encryption_mode=aes256
ike_dhgroup=group2
ike_lifetime_s=300
ike_lifetime_s_max=86400
ike_dpd_s=300
ipsec_tunnel_mode=tunnel
ipsec_tunnel_type=esp_auth
ipsec_encryption_mode=aes128
ipsec_hash_algorithm=sha
ipsec_pfsgroup=none
ipsec_lifetime_s=28800
ipsec_lifetime_s_max=86400
ipsec_lifetime_kb=0
ipsec_lifetime_kb_max=0
ipsec_mismatch_behavior=drop
Protected Networks:
  [1] 10.0.0.0/16 -> 10.101.0.0/16
  [2] 10.4.0.0/16 -> 10.101.0.0/16
  [3] 10.3.0.0/16 -> 10.101.0.0/16
  [4] 10.2.0.0/16 -> 10.101.0.0/16
  [5] 10.1.0.0/16 -> 10.101.0.0/16
=====
```

Figure 31: IPsec Tunnel Configuration

Standby WAN Links

Edge 5.0 introduces Standby WAN Links, which you can configure so user traffic will only be transmitted on that WAN Link when all other available WAN Links are dead or disabled. This feature can *only* be configured for Private Intranet and Public Internet Access Types. Simply click the Standby checkbox when you configure a Private Intranet or Public Internet WAN Link in the Configuration Editor.

The screenshot shows the 'Basic Settings' tab for a WAN Link configuration. The 'WAN Ingress' and 'WAN Egress' sections both have 'Physical Rate (kbps)' set to 6000 and 'Permitted Rate (kbps)' set to 6000, with the 'Set Permitted From Physical' checkbox checked. The 'Access Type' is set to 'Public Internet'. The 'Standby' checkbox is checked. The 'Public IP Address' is set to 3.3.3.3. A note states: 'Note: Changing the access type of this WAN Link may cause automatically generated Paths to this link to be added or removed.'

Figure 32: Standby WAN Link on Public Internet Connection

Monitoring

A Path that has at least one Standby WAN Link as an endpoint is considered a backup Path. Under **Monitor**, and then **Statistics** all functions for Paths are supported regardless of whether a Path is configured as a backup Path in Edge 5.0.

Statistics

Show: **Paths (Summary)** Enable Auto Refresh **5** seconds **Stop** Show latest data.

Path Statistics Summary

Show: **100** entries Showing 1 to 16 of 16 entries Processing...

Filter: in (Any column)

Num	From Link	To Link	Path State	Conduit State	Conduit Type	BOWT	Jitter (ms)	Loss %	kbps	Congestion
1	NCN_PRIV	CL1-PRV	GOOD	GOOD	Static	5	2	0.00	16.15	NO
2	NCN_PUB	CL1-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.38	NO
3	CL1-PRV	NCN_PRIV	GOOD	GOOD	Static	4	4	0.00	19.53	NO
4	CL1-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	5	2	0.00	0.38	NO
5	NCN_PRIV	CL2-PRV	GOOD	GOOD	Static	5	2	0.00	16.55	NO
6	NCN_PUB	CL2-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.34	NO
7	CL2-PRV	NCN_PRIV	GOOD	GOOD	Static	4	4	0.00	19.00	NO
8	CL2-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	4	4	0.00	0.34	NO
9	NCN_PRIV	CL3-PRV	GOOD	GOOD	Static	5	2	0.00	15.31	NO
10	NCN_PUB	CL3-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.34	NO
11	CL3-PRV	NCN_PRIV	GOOD	GOOD	Static	4	4	0.00	18.75	NO
12	CL3-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	5	2	0.00	0.34	NO
13	NCN_PRIV	CL4-PRV	GOOD	GOOD	Static	5	2	0.00	12.57	NO
14	NCN_PUB	CL4-PUB	GOOD	GOOD	Static	5	2	0.00	15.56	NO
15	CL4-PRV	NCN_PRIV	GOOD	GOOD	Static	4	4	0.00	20.02	NO
16	CL4-PUB	NCN_PUB	GOOD	GOOD	Static	4	4	0.00	11.00	NO

Showing 1 to 16 of 16 entries

Bandwidth calculated over the last 5 seconds

Figure 33: Paths (Summary) Including Standby WAN Links

Statistics

Show: **Paths (Summary)** Enable Auto Refresh **5** seconds **Stop** Show latest data.

Path Statistics Summary

Show: **100** entries Showing 1 to 6 of 6 entries (filtered from 16 total entries)

Filter: **standby** in (Any column)

Num	From Link	To Link	Path State	Conduit State	Conduit Type	BOWT	Jitter (ms)	Loss %	kbps	Congestion
2	NCN_PUB	CL1-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.38	NO
4	CL1-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	5	2	0.00	0.38	NO
6	NCN_PUB	CL2-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.34	NO
8	CL2-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	5	2	0.00	0.28	NO
10	NCN_PUB	CL3-PUB (standby)	GOOD	GOOD	Static	5	2	0.00	0.34	NO
12	CL3-PUB (standby)	NCN_PUB	GOOD	GOOD	Static	5	2	0.00	0.34	NO

Showing 1 to 6 of 6 entries (filtered from 16 total entries)

Bandwidth calculated over the last 5 seconds

Figure 34: Paths (Summary) Including Standby WAN Links Filtered

12

Release 5.1 Features

This chapter includes features and enhancements released in 5.1.

Virtual Appliance VT800

Edge 5.1 introduces support for the new Virtual Appliance VT800. This new virtual appliance supports different performance levels depending on how it is licensed. The VT800 supports up to 200 Mbps of full-duplex performance, 8 Public WAN Links, 32 Private WAN Links, and scales higher than the VT500 to support more Conduits, Paths, and tunnels.

Alarm System

Edge 5.1 introduces a new Alarm System that streamlines the configuration and number of severity based alerts for network administrators. Now you can configure contextually-based alarms with specific criteria for triggering and clearing alarm states. To configure an Alarm:

1. Under **Integrate** > **Configure Alarms**, click the **Add Alarm** button.
2. Select an **Event Type** from the drop-down menu.

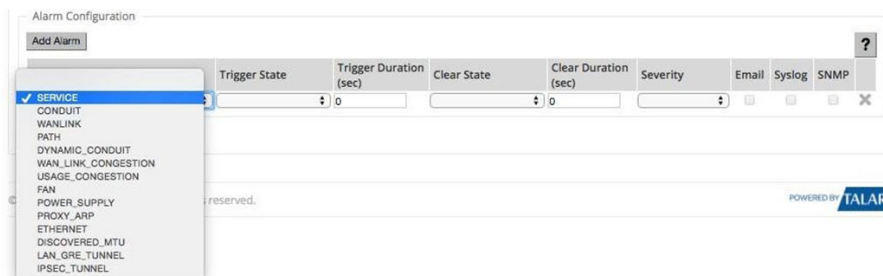


Figure 1: Select Event Type

1. Choose a **Trigger State** from the drop-down menu. When the Event Type enters this state an Alarm is triggered. The options available on the Trigger State drop-down menu are determined by the Event Type.




Figure 2: Select Trigger State

1. Enter the amount of time (in seconds) in the **Trigger Duration** field that the Event Type must remain in the Trigger state to trigger the Alarm. The default is 0 seconds, which would trigger the alarm immediately.

Note:

The Trigger Duration field is not available for some Event Types.

1. Choose a **Clear State** from the drop-down menu. When the Event Type enters this state the existing Alarm is cleared. The options available on the Clear State dropdown menu are determined by the Trigger State.

The screenshot shows the 'Alarm Configuration' window. The 'Event Type' is set to 'SERVICE' and the 'Trigger State' is 'STANDBY_HA'. The 'Trigger Duration (sec)' is 0. The 'Clear State' dropdown menu is open, showing 'ACTIVE_HA' and 'STANDBY_HA' as options. The 'Clear Duration (sec)' is 0. The 'Severity' dropdown is also visible. There are checkboxes for 'Email', 'Syslog', and 'SNMP'. An 'Apply Settings' button is at the bottom left. The footer includes '© 2007-2016 Talari Networks. All rights reserved.' and 'POWERED BY TALARI'.

Figure 3: Select a Clear State

1. Enter the amount of time (in seconds) in the **Clear Duration** field that the Event Type must remain in the Clear State to clear the Alarm. The default is 0 seconds, which would clear the alarm immediately.

Note:

The Clear Duration field is not available for some Event Types.

1. Choose a **Severity** from the drop-down menu based on the urgency of the alarm. The Severity is displayed in the alert that is sent out when the Alarm is triggered and cleared and is also displayed with the Alarm under **Diagnose > View/Clear Alarms**.

The screenshot shows the 'Alarm Configuration' window. The 'Event Type' is 'SERVICE' and the 'Trigger State' is 'STANDBY_HA'. The 'Trigger Duration (sec)' is 0. The 'Clear State' dropdown is set to 'ACTIVE_HA'. The 'Clear Duration (sec)' is 0. The 'Severity' dropdown menu is open, showing options: 'DEBUG', 'INFO', 'NOTICE', 'WARNING', 'ERROR', 'CRITICAL', 'ALERT', and 'EMERGENCY'. There are checkboxes for 'Email', 'Syslog', and 'SNMP'. An 'Apply Settings' button is at the bottom left. The footer includes '© 2007-2016 Talari Networks. All rights reserved.' and 'POWERED BY TALARI'.

Figure 4: Choose a Severity

1. Select the alert delivery method by clicking the **Email**, **Syslog**, and **SNMP** checkboxes. You can select multiple delivery methods, however, even if you do not choose a delivery method, an alarm is produced that you can view on the **View/Clear Alarms** page.
2. Click **Apply** to save the alarm.
3. Repeat steps 1 through 8 to add additional Alarms.

Diagnose Alarms

To diagnose network issues based on current Alarms, you can use the Diagnose Alarm page to see a list of all current Alarms. Click **Diagnose**, and then **View/Clear**

Alarms to sort and filter the list of Alarms, or clear them by clicking the **Clear Action** checkbox at the end of an Alarm row then clicking the **Clear Checked Alarms** button near the top of the page. Click the **Clear All Alarms** button to clear all current alarms.

Diagnose / View/Clear Alarms

Alarms

Enable Auto Refresh Time Interval 5 seconds Refresh

Clear Checked Alarms Clear All Alarms ?

Triggered Alarms Summary

Filter: in Any column Apply

Show 100 entries Showing 1 to 12 of 12 entries

Severity	Object Type	Object Name	Trigger State	Trigger Duration (sec)	Clear State	Clear Duration (sec)	Clear Action
CRITICAL	CONDUIT	Pittsburgh:Raleigh	DEAD	0	GOOD	0	<input type="checkbox"/>
CRITICAL	CONDUIT	Pittsburgh:Portland	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PittsburghLink->RaleighLink	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PittsburghLink2->RaleighLink2	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	RaleighLink->PittsburghLink	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PortlandLink->PittsburghLink	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PittsburghLink->PortlandLink	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PittsburghLink2->PortlandLink2	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	RaleighLink2->PittsburghLink2	DEAD	0	GOOD	0	<input type="checkbox"/>
ERROR	PATH	PortlandLink2->PittsburghLink2	DEAD	0	GOOD	0	<input type="checkbox"/>
WARNING	WANLINK	PittsburghLink	DEAD	0	GOOD	0	<input type="checkbox"/>
WARNING	WANLINK	PittsburghLink2	DEAD	0	GOOD	0	<input type="checkbox"/>

Showing 1 to 12 of 12 entries

Figure 5: View or Clear Alarms

Route Export Filters

For networks in which Route Learning has been enabled, Edge 5.1 provides more fine grained control over which Edge routes are advertised to routing neighbors rather than advertising all or no routes. Export Filters are used to include or exclude routes for advertisement via OSPF and IBGP based on specific match criteria.

Under **Connections**, and then (**Site Name**), and then **Route Learning** in the Configuration Editor, **Import Filters** are separate and distinct from **Export Filters**. You may configure up to 32 Export Filters.

Route Learning ?

OSPF
IBGP

Import Filters + ?

Order	Source Router	Destination	Prefix	Next Hop	Protocol	Cost	Include	Enabled	Delete	Clone
100	*	Allowed_CL1_Local_Networks	eq *	*	Any	eq *	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
200	<Manual>	11.123.0.0/16	eq *	*	Any	eq *	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(auto)	<Manual>	*	eq *	*	Any	eq *	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Export Filters + ?

Order	Network Address	Prefix	APN Cost	Service Type	Site/Service Name	Gateway IP Address	Include	Enabled	Delete	Clone
100	<Manual>	*	eq *	Conduit	JM-CL2	*	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
200	<Manual>	0.0.0.0/32	eq 0	Any	<Any>	*	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
(auto)	<Manual>	*	eq *	Any	<Any>	*	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

You can use the following criteria to construct each Export Filter that you create:


- **Order:** The Order in which filters are prioritized. The first filter that a route matches to will be applied to that route.
- **Service Type:** To match filters from a specific routing domain, choose one of the configured Routing Domains from the drop-down menu.
- **Network Address:** Enter the IP Address and Netmask or configured Network Object that describes the route's network.
- **Prefix:** To match routes by prefix, choose a match predicate from the drop-down menu and enter a Route prefix in the adjacent field.

- **APN Cost:** The method (predicate) and the APN Route Cost that are used to narrow the Selection of routes exported.
- **Service Type:** To match routes by Service Type, select the Service Type (e.g., Local, Internet, Intranet, LAN GRE Tunnel, LAN IPsec Tunnel, or Passthrough) from a list of existing, supported Services.
- **Site/Service Name:** If you select a Service Type, you may also need to select a specific Site or Service Name
- **Gateway IP Address:** If you choose LAN GRE Tunnel as the Service Type, enter the Gateway IP for the tunnel.
- **Include:** Click the checkbox to **Include** routes that match this filter. Otherwise matching routes are ignored.
- **Enable:** Click the checkbox to **Enable** this filter. Otherwise the filter is ignored.
- **Clone:** Click the **Clone** icon to make a copy of an existing Filter.

Operating System Patching

To facilitate the expeditious distribution of Debian patches to customers, those patches will now be bundled in new OS patches that are separate from Edge Software updates and full OS updates. OS patches can be independently uploaded and installed to the active OS partition on Oracle Talari Appliances running OS 4.1 or later and each patch builds on the previously uploaded patch.

To see the Currently Installed Patch Level, navigate to **Manage Appliance** ▢ **OS Partitions**. To upload new OS patches, scroll down to the **OS Patch Network Upload** area of the page.



OS Patch Network Upload

This section allows a Talari OS Patch to be uploaded from the web browser.

Currently Installed Patch Level: 4.5P0

Filename: No file chosen

Figure 7: OS Patch Network Upload

Once you download a new patch to your local machine, click the **Choose File** button to select the file you downloaded and click **Upload** to install the patch to the active OS partition on your Oracle Talari Appliance. For a more in depth explanation of OS patching, please refer to the *OS Partition Update Guide*.

Customizable Web Console

Now you can customize the look and feel of your Oracle Talari Appliance's Web Console. Edge 5.1 allows network administrators to add a Custom Login Message, a Custom Support Link, and Upload a Custom Logo to brand their Oracle Talari Appliances' web interfaces.

From the **Manage Appliance** ▢ **HTTPS Settings** in the **Custom Login Message** area, enter a message to appear on the login page for appliance users. Click the **Allow HTML** box to format and style your message with HTML. When you are done, click the **Save Login Message** button to save the message.



Figure 8: Custom Login Message

In the **Custom Support Link** area of the **HTTPS Settings** screen enter a **Support Link Name** and your organization's **Support Link URL** to create a link on the appliance login page.



Figure 9: Custom Support URL

From the **Upload Custom Logo** section on the **HTTPS Settings** screen, you can upload a logo to replace the Talari logo on your appliance. Click the **Choose File** button, choose the logo image you want to upload, and click **Upload Custom Logo**. If you need to remove the logo you updated, click **Remove Custom**.




Figure 10: Custom Logo

Here is an example of the login screen of an Oracle Talari Appliance with a Custom Logo and Custom Login Message.



Figure 11: Customized Web Console

DHCP Relay and DHCP Server


Edge 5.1 introduces the ability to use your Oracle Talari Appliances as either DHCP Servers or DHCP Relay Agents to simplify your network's configuration. Now you can use your Oracle Talari Appliances to issue IP Addresses via DHCP or forward DHCP packets between clients and servers where necessary.

 **Note:**

DHCP Relay and DHCP Server require appliances to be running OS 4.5 or later.

Management Interface DHCP Server

From the **Manage Appliance** > **Local Network Settings** screen you can now configure the **Management Interface DHCP Server**. Click the **Enable DHCP Server** checkbox to start the server, then enter the **Lease Time** (in minutes), the **Domain Name**, and define the IP Address range by entering a **Start IP Address** and an **End IP Address**.



Management Interface DHCP Server

The Management Interface DHCP Server will use the current Management Interface IP settings (gateway, subnet mask, and DNS servers) for DHCP offers. The DHCP Server IP range, defined by Start and End IP Address, must be valid in the Management Interface subnet.

DHCP Server Status: running

Enable DHCP Server:

Lease Time (minutes):

Domain Name:

Start IP Address:


End IP Address:

Figure 12: Configure DHCP Server

Click the **Change Settings** button to finish configuring the DHCP Server. Click the **Show Clients** button to view the current DHCP clients, and click the **Clear Clients** button to release the current DHCP Client Leases.

 **Note:**

If you plan to use DHCP Server on an Oracle Talari Appliance configured for High Availability (HA), do not configure the service on both the Active and Standby appliance. Doing so will lead to duplicate IP Addresses on the defined management network.



Client IP Address	Lease Start Time	Lease End Time	Client MAC Address	Client Hostname	State
10.3.1.2	Tue Feb 9 00:00:27 2016	Tue Feb 9 00:00:27 2016	36:5e:b5:f3:43:99	AS-TB-CL3-TN	active

Figure 13: DHCP Server Client Database

DHCP Relay

Network administrators can use the DHCP Relay service on the management port of Oracle Talari Appliances to relay requests and replies between local DHCP Clients and a remote DHCP server. This allows local hosts to acquire dynamic IP Addresses from the remote DHCP Server. For a more in depth explanation of DHCP Relay, please refer to *Using Oracle Talari Appliances as DHCP Relay Agents*.

From the **Manage Appliance** > **Local Network Settings** Screen you can configure the **Management Interface DHCP Relay**. Click the **Enable DHCP Relay** checkbox to enable the service. Enter the **DHCP Server IP Address** and click the **Change settings** button to begin using your appliance as a DHCP Relay Agent.

 **Note:**

If you plan to use DHCP Relay on an Oracle Talari Appliance configured for High Availability

(HA), do not configure the service on both the Active and Standby appliance. Doing so will lead to duplicate IP addresses on the defined management network.



Management Interface DHCP Relay

Enable DHCP Relay:

DHCP Server IP Address:

Figure 14: Enable DHCP Relay

13

Release 5.2 Features

This chapter includes features and enhancements released in 5.2.

Support for 550 Sites

Edge 5.2 supports the ability for a T5200, functioning as an NCN, to create up to 550 static Conduits without loss of performance and with full support for all previously existing features, including up to 16,000 Routes. The user also gains the ability to have up to 23,000 WAN Paths, 11,000 WAN Links, 512,000 Flows, and 200,000 Rules.

Stateful Firewall

Edge 5.2 provides a firewall built into the Oracle Talari Application. The firewall allows Policies between Services and Zones, and supports Static NAT, Dynamic NAT (PAT), and Dynamic NAT with Port Forwarding. Additional firewall capabilities include:

- Filtering traffic flows between Zones
- Filtering traffic between services within a Zone
- Filtering traffic between services that reside in different Zones
- Filtering traffic between services at a site
- Defining Filter Policies to Allow, Deny, or Reject flows
- Tracking flow state for selected flows
- Applying Global Policy Templates
- Support for Port Address Translation for traffic to the Internet on an untrusted port, as well as port forwarding inbound and outbound

To simplify the configuration process, firewall Policies are created at the Global level. This Global configuration consists of Pre-Appliance and Post-Appliance site Policy Templates that can be applied to all sites within Edge. For a more in-depth explanation of the Stateful Firewall feature in Oracle SD-WAN Edge 5.2 GA, please refer to *SD-WAN Firewall Configuration Guide*.

DHCP Relay & DHCP Server

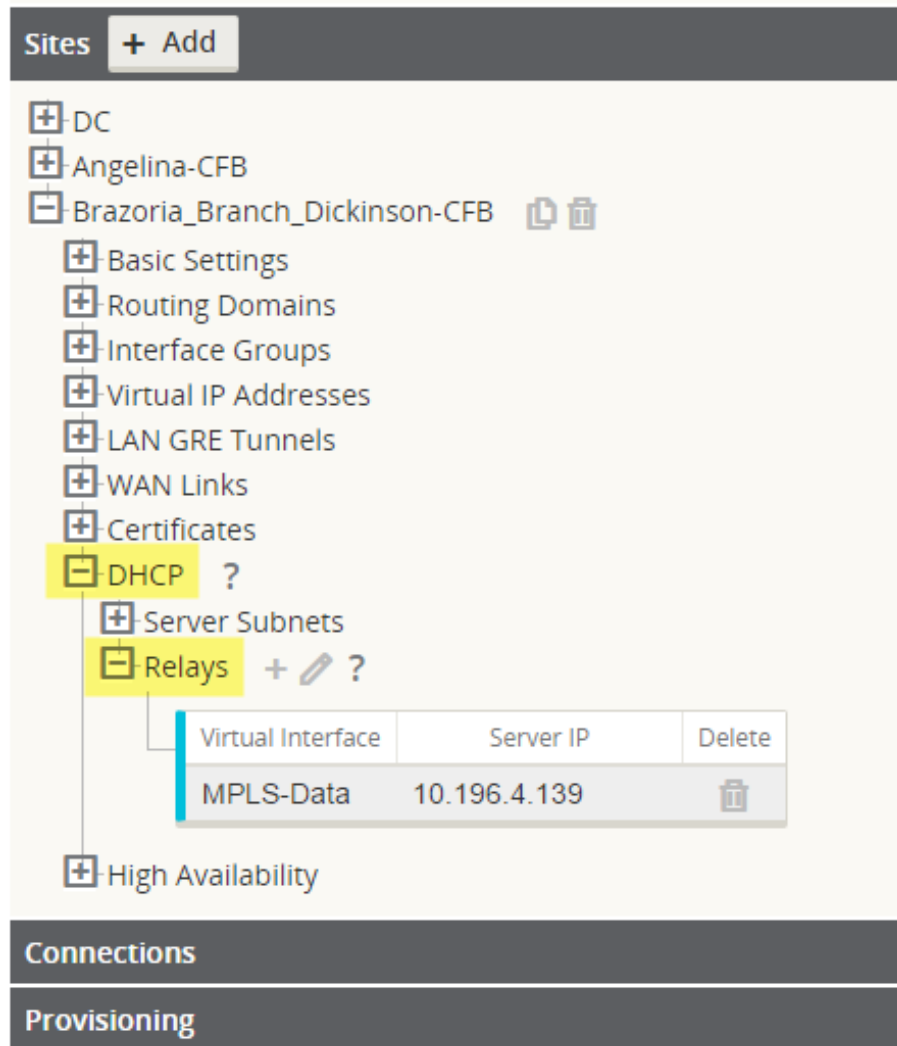
Devices on the same network as the Oracle Talari Appliance's LAN/VLAN interface may now use the DHCP Relay & DHCP Server features to provide those devices with their IP configuration. These features help to simplify the client site network by reducing the amount of equipment necessary.

- DHCP Relay

Network administrators can now use the DHCP Relay service on data ports of Oracle Talari Appliances to relay requests and replies between local DHCP Clients and a remote DHCP

Server. This allows local hosts to acquire dynamic IP addresses from the remote DHCP Server.

To configure DHCP Relay, navigate to **Manage Network > Configuration Editor > Sites > [Site Name] > DHCP**. Expand **Relays** then specify the data ports to be used and the Server IP address.



- DHCP Server


Network administrators can now also use the DHCP Server feature on data ports of Oracle Talari Appliances to allow local hosts to acquire dynamic or static IP addressing directly from the Oracle Talari Appliance.

To configure DHCP Server:

1. Navigate to **Manage Network > Configuration Editor > Sites > [Site Name] > DHCP** and expand **Server Subnets**.
2. Select the Virtual Interface to be used and specify the range of IP addresses allowed to be dynamically assigned to local hosts.

Users may also choose to enter additional information in this section that hosts will then be configured with as well, such as gateway IP, DNS, and an Option Set (described below).

The **Hosts** option of this drop down allows users to manually tie specific IP addresses to specific hosts via host MAC address if desired.



The screenshot shows the DHCP configuration interface. The left sidebar lists various configuration sections: Global, Sites (+ Add), DC, Angelina-CFB (with expand and delete icons), Basic Settings, Routing Domains, Interface Groups, Virtual IP Addresses, LAN GRE Tunnels, WAN Links, Certificates, DHCP (?), Server Subnets (+ ?), Relays, High Availability, and Brazoria_Branch_Dickinson-CFB. The main area displays a table of Virtual Interfaces:

Virtual Interface	IP Subnet	Domain Name	Primary DNS	Secondary DNS	Enable	Delete
MPLS-Data	10.200.247.41/24		192.168.51.5		<input checked="" type="checkbox"/>	

Below the table, the 'MPLS-Data' interface is expanded to show its configuration:

Ranges +

Range Start IP	Range End IP	Gateway IP	Option Set	Delete
10.200.247.200	10.200.247.205	10.200.247.1	New DHCP Option Set	

Hosts +


Fixed IP Address	MAC Address	Option Set	Delete
10.200.247.206	1a:0a:45:f4:e1:52	<None>	

At the bottom of the interface, there are sections for Connections, Provisioning, and dits: 0.

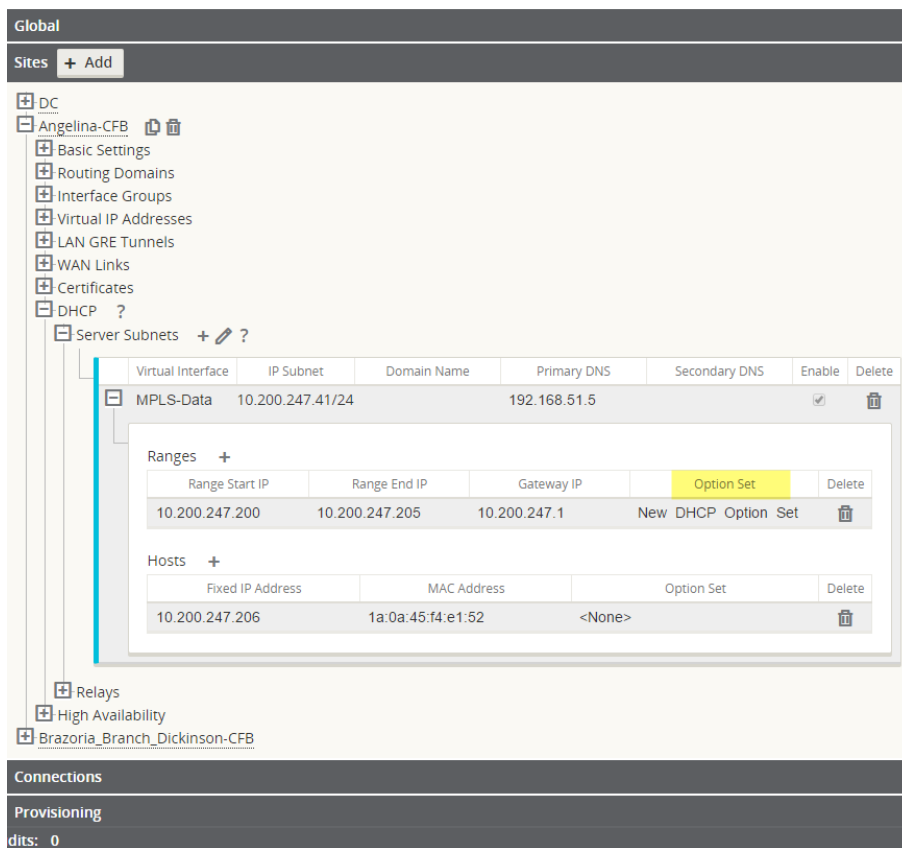
 **Note:**

The following feature is optional, not required.

DHCP Option Sets are a group of DHCP settings or parameters that can be applied to individual IP address ranges. To create DHCP Option Sets, navigate to the **Global** section of the configuration and expand **Options**. Enter the required settings you would like to include in the set, then click **Apply**.



Your DHCP Option Set must then be tied to a DHCP range and is done so in the **Sites** section where the IP address range was defined.



To view a list of Clients from the DHCP Server Database, navigate to **Monitor > DHCP** from the web UI.

Routing Domain	Client IP Address	Lease Start Time	Lease End Time	Client MAC Address	Client Hostname	State
Default_RoutingDomain	10.200.247.200	Mon Jul 11 15:23:23 2016	Mon Jul 11 15:29:23 2016	3a:1a:dc:67:ca:b4	TexasF_Angelina2_TN	active

Standby WAN Link (VSAT)

Introduced in Edge 5.2, this feature gives users the ability to have as many as three Standby WAN Links with customizable priorities per location, providing users the flexibility to use the more expensive links only when needed. The Standby WAN Links may be activated to supplement Conduit bandwidth when specified thresholds are met (On-Demand Standby) or when all primary WAN Links are DEAD or Disabled (Last-Resort Standby).

Below are steps to enable this feature. This example chooses the On-Demand Standby option:

1. Set the WAN Link mode using the Configuration Editor under **Sites > [Site Name] > WAN Links > [WAN Link Name] > Settings > Advanced Settings > WAN Link Mode**.

The screenshot shows the configuration editor for a WAN Link. The left sidebar shows a tree view with 'Settings' selected. The main content area shows the 'Advanced Settings' tab. The 'WAN Link Mode' is set to 'On-Demand Standby'. Other settings include:

- Provider ID: (empty)
- Frame Cost (bytes): 0
- Congestion Threshold (μs): 20000
- MTU Size (bytes): 1500
- WAN Link Mode: On-Demand Standby
- Priority: 1
- Heartbeat Interval: 1 second

The **Priority** option is a value to indicate which Standby WAN Link will be activated in which order and the **Heartbeat Interval** can either be set or disabled.


 **Note:**

A more detailed definition of the three modes available can be found by clicking the ? icon to display the help text.

 **Note:**

A WAN Link configured in Standby mode can not have Internet or Intranet Services enabled on it, this will result in a Configuration Audit Error.


2. Create a Default Set in the **Global** section that will be used for Conduits using the Standby WAN Link.



The screenshot shows the configuration interface for the **Global** section. The left sidebar lists various configuration categories: APN Settings, Routing Domains, Firewall, Applications, Network Objects, Default Sets (with a help icon), Conduit Default Sets (with a plus and help icon), and New_Conduit_Default_Set (with copy and delete icons). Under New_Conduit_Default_Set, there are sub-entries for Classes, Rules, IPsec Settings, and Advanced Settings (with edit and help icons). The Advanced Settings sub-entry is selected, and a dialog box is open. The dialog box has a title 'Standby WAN Link(s) Activation' and contains the text: 'Bandwidth threshold in a percentage of the total egress permitted rates of regular WAN links in the conduit (%):'. Below this text is a text input field containing the value '20'. At the bottom right of the dialog box are 'Apply' and 'Close' buttons.

Under **Advanced Settings**, the user is able to specify a bandwidth threshold in terms of a percentage of the total WAN Egress Permitted Rates of regular WAN Links. If the available bandwidth provided by the regular WAN Links in the conduit falls below this bandwidth threshold, On-Demand Standby WAN Links in the Conduit will be activated to supplement bandwidth.

Apply the Default Set to specific Conduits under **Connections > [Site Name] > Conduits > [Conduit Name] > Local Site > Basic Settings > Default Set.**



Note:

Step 2 is only required when choosing the On-Demand Standby option and is not applicable for Last-Resort Standby WAN Links.

Output from the **Monitor > Statistics** page of the web UI will let you know which WAN Links are in Standby mode. The user will observe minimal amounts of traffic traversing such links, depending on how the Heartbeat Interval and Activation thresholds have been configured.

Show 100 entries Showing 1 to 12 of 12 entries Processing...

Num	From Link	To Link	Path State	Conduit State	Conduit Type	BOWT	Jitter (mS)	Loss %	kbps	Congestion
1	BestEffort	NYC-MPLS	GOOD	GOOD	Static	5	3	0.00	24.80	NO
2	EF	NYC-MPLS	GOOD	GOOD	Static	5	3	0.00	28.28	NO
3	SYR-Charter (standby)	NYC-DSL	GOOD	GOOD	Static	5	4	0.00	0.39	NO
4	NYC-DSL	SYR-Charter (standby)	GOOD	GOOD	Static	2	4	0.00	0.39	NO
5	NYC-MPLS	BestEffort	GOOD	GOOD	Static	6	3	0.00	16.04	NO
6	NYC-MPLS	EF	GOOD	GOOD	Static	6	3	0.00	18.38	NO
7	BestEffort	BUFF-MPLS	GOOD	GOOD	Static	5	5	0.00	19.73	NO
8	EF	BUFF-MPLS	GOOD	GOOD	Static	5	5	0.00	22.95	NO
9	SYR-Charter (standby)	BUFF-CLink	GOOD	GOOD	Static	5	4	0.00	0.39	NO
10	BUFF-CLink	SYR-Charter (standby)	GOOD	GOOD	Static	2	4	0.00	0.55	NO
11	BUFF-MPLS	BestEffort	GOOD	GOOD	Static	2	6	0.00	28.93	YES
12	BUFF-MPLS	EF	GOOD	GOOD	Static	6	5	0.00	93.45	YES


Showing 1 to 12 of 12 entries

Adaptive Bandwidth Detection

This feature is introduced in Edge 5.2 for users with VSAT, LOS, Microwave, 3G/4G/LTE WAN Links, whose available bandwidth varies based upon weather and atmosphere conditions, location, line of site obstructions, etc. It allows the Oracle Talari Appliance to adjust the bandwidth rate on the WAN Link dynamically based on a defined bandwidth range, to use the maximum amount available without marking the paths BAD.

To enable this feature:

1. In the Configuration Editor, navigate to **Sites > [Site Name] > WAN Links > [WAN Link Name] > Settings > Advanced Settings**.
2. Check the **Adaptive Bandwidth Detection** box and enter in the **Minimum Acceptable Bandwidth**.



The screenshot shows the 'Settings' window for a WAN Link. The 'Advanced Settings' section is expanded, showing the following configuration:

Basic Settings	
Advanced Settings	
Provider ID:	Frame Cost (bytes): 0
Congestion Threshold (µs): 20000	MTU Size (bytes): 1500
WAN Link Mode: Regular Active	
<input checked="" type="checkbox"/> Adaptive Bandwidth Detection	Minimum Acceptable Bandwidth (Percent Conduit Egress %): 30
Eligibility	
Cell Networking	

Note:

There is no specific logging or event alerts for this feature, but users may refer to **Monitor > Performance Reports** for a historical trend in bandwidth rates.

To schedule recurring bandwidth tests after enabling Adaptive Bandwidth Detection:

1. Navigate to **Diagnose > Path Bandwidth**.
2. Under **Schedule Path Bandwidth Testing**, click the Add button. Select the Path Name to test on, Frequency, Day of Week (if applicable), Hour (if applicable), and Minute, then click Apply Settings.

Diagnose / Path Bandwidth Talari Support

Instant Path Bandwidth Testing

Path: NCN-site-WL-1->CL1-site-WL-1 ?

Schedule Path Bandwidth Testing

Add

Path Name	Frequency	Day of Week	Hour	Minute	
CL1-site-WL-1->NCN-site-WL-1	every 4 hours	Sunday	0	0	X

History Path Bandwidth Testing Result

Show 50 entries Showing 0 to 0 of 0 entries Search

Num	From Link	To Link	Test Time	Min Bandwidth (kbps)	Max Bandwidth (kbps)	Avg Bandwidth (kbps)
No data available in table						

Showing 0 to 0 of 0 entries



Note:

If Adaptive Bandwidth Detection is configured but recurring bandwidth testing is not scheduled, the bandwidth test will run once and the Oracle Talari appliance will use that one-time result. Recurring bandwidth testing is required for Adaptive Bandwidth Detection to function as intended.

Users may monitor the bandwidth detected on these links from the web UI under **Monitor > Statistics > WAN Link Usage > Local WAN Egress On Demand WAN Link Usages**.

Local WAN Egress On Demand WAN Link Usages

Filter: in Any column

Show 100 entries Showing 1 to 2 of 2 entries

WAN Link	WAN Link Mode	Standby Priority	Adaptive Bandwidth Detection				Conduit Name	Conduit Activate On Demand Threshold Kbps	Conduit Usable Bandwidth Kbps	In Use
			Configured	Minimum Acceptable BW Kbps	Maximum Allowed BW Kbps	Current Allowed BW Kbps				
Client-Inet (standby)	On-Demand	1	Local	2940	9800	3940	NCN-site-Client-test	2940	3940	Yes
Client-Satcom	Regular-Active	N/A	No	N/A	N/A	N/A	NCN-site-Client-test	2940	3940	No

Showing 1 to 2 of 2 entries

Active Bandwidth Testing

Edge 5.2 provides users the ability to issue an instant path bandwidth test, or to schedule such testing to be completed at specific times on a recurring basis. This feature will be useful for demonstrating how much bandwidth the user has between two locations during new and existing installations, also for testing paths to determine the outcome of setting and confirmation changes, such as adjusting DSCP tag settings or bandwidth Permitted Rates.

To use this tool:

1. Navigate to **Diagnose > Path Bandwidth**.
2. Select the desired Path and click **Test**.

Diagnose / Path Bandwidth Talari Support

Instant Path Bandwidth Testing

Path: NYC-DSL->BUFF-CLink ?

Results

Minimum Bandwidth:19850 kbps
Maximum Bandwidth:20663 kbps
Average Bandwidth:20266 kbps

Schedule Path Bandwidth Testing

Path Name	Frequency	Day of Week	Hour	Minute	
NYC-MPLS->BUFF-MPLS	every day	Sunday	10	30	✕
BUFF-CLink->NYC-DSL	every 12 hours	Sunday	0	15	✕

History Path Bandwidth Testing Result

Filter: Any column

Show 10 entries Showing 11 to 20 of 34 entries

1 2 3 4

Num	From Link	To Link	Test Time (hh:mm on mm/dd/yyyy)	Min Bandwidth (kbps)	Max Bandwidth (kbps)	Avg Bandwidth (kbps)
11	NYC-DSL	BUFF-CLink	10:18 on 7/7/2016	1962	2026	1999
12	NYC-DSL	BUFF-CLink	10:34 on 7/7/2016	1986	2018	2004

The output will display the minimum, maximum, and average bandwidth results of the test. Along with the ability to test the bandwidth, the user can now change the configuration file to use the learned bandwidth. This is accomplished via the Auto Learn option is under **Site > [Site Name] > WAN Links > [WAN Link Name] > Settings** and if enabled, the system will use the learned bandwidth.

Users may also schedule reoccurring tests of path bandwidth in weekly, daily, or hourly intervals.

Diagnose / Path Bandwidth Talari Support

Instant Path Bandwidth Testing

Path: NYC-DSL->BUFF-CLink ?

Results

Minimum Bandwidth:19850 kbps
Maximum Bandwidth:20663 kbps
Average Bandwidth:20266 kbps

Schedule Path Bandwidth Testing

Path Name	Frequency	Day of Week	Hour	Minute	
NYC-MPLS->BUFF-MPLS	every day	Sunday	10	30	✕
BUFF-CLink->NYC-DSL	every 12 hours	Sunday	0	15	✕

History Path Bandwidth Testing Result

Filter: Any column

Show 10 entries Showing 11 to 20 of 34 entries

1 2 3 4

Num	From Link	To Link	Test Time (hh:mm on mm/dd/yyyy)	Min Bandwidth (kbps)	Max Bandwidth (kbps)	Avg Bandwidth (kbps)
11	NYC-DSL	BUFF-CLink	10:18 on 7/7/2016	1962	2026	1999
12	NYC-DSL	BUFF-CLink	10:34 on 7/7/2016	1986	2018	2004

 **Note:**

A history of the path bandwidth testing results will be displayed at the bottom of this page and results will archive every 7 days.

SNMPv3 Polling and Trap Capability



Note:

The platform only supports a single user account for each SNMPv3 capability.

To configure SNMPv3 Polling and Traps, navigate to the SNMPv3 section of the **Integrate > Configure Events and Alerts** page and fill in the fields as required.

SNMP v3

Enable v3 Agent

User Name:

Password:

Verify Password:

Authentication:

Encryption:

Enable v3 Traps

Destination IP Address(es):

Port:

User Name:

Password:

Verify Password:

Authentication:

Encryption:

Eligibility for IPsec Non-Conduit Routes

Prior to Edge R5.2, IPsec tunnel routes would remain in the route table even if the tunnel became unavailable.

Statistics

Show: Routes Enable Auto Refresh 5 seconds Refresh Clear Counters on Refresh Purge dynamic routes

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: DC1-Intrane

Show 100 entries Showing 1 to 2 of 2 entries (filtered from 27 total entries)

Num	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible
2	10.3.10.30/32	*	DC1-Intranet	YES	*	DC1	Static			5	22	YES
23	10.4.10.0/24	*	DC1-Intranet	YES	*	DC1	Static			5	11	YES

Showing 1 to 2 of 2 entries (filtered from 27 total entries)

Routing Table Example A

Using the **Keepalive** option under **Connections > [Site Name] > IPsec Tunnels** enhances such behavior so that the IPsec Non-Conduit Routes will now be considered ineligible when the IPsec tunnel is no longer available.

Connections

- DC1
 - WAN-to-WAN Forwarding
 - Conduits
 - Internet Services
 - Intranet Services
 - WAN Links
 - IPsec Tunnels** + ?

Service Type	Name	Routing Domain	Firewall Zone	Local IP	Peer IP	MTU	Keepalive	Delete
Intranet	DC1-Intranet	<Default>	<Default>	10.0.10.11	10.3.10.30	1500	<input checked="" type="checkbox"/>	

IKE Settings ?

IPsec Settings ?

IPsec Protected Networks + Add ?

Source IP/Prefix	Destination IP/Prefix	Delete
10.0.10.0/24	10.4.10.0/24	

Statistics

Show: Routes Enable Auto Refresh 5 seconds Refresh Clear Counters on Refresh Purge dynamic routes

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: DC1-Intrane

Show 100 entries Showing 1 to 2 of 2 entries (filtered from 27 total entries)

Num	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible
2	10.3.10.30/32	*	DC1-Intranet	YES	*	DC1	Static			5	11	YES
23	10.4.10.0/24	*	DC1-Intranet	NO	*	DC1	Static			5	0	NO

Showing 1 to 2 of 2 entries (filtered from 27 total entries)

Routing Table Example B

Additional Enhancements

Routing Enhancements

- OSPF Type 5 to Type 1

Users now have the ability to decide whether learned OSPF routes are exported as external Type 5 or intra-area Type 1.

- Hairpin from non-WAN-to-WAN Forwarding Site

Users may now configure a 0.0.0.0/0 route to hairpin traffic between two locations without impacting any additional locations. If used for Intranet traffic, specific Intranet

routes will be added to the Client site to forward Intranet traffic through the Conduit to the hairpin site. Enabling WAN-to-WAN Forwarding to accomplish this is no longer necessary.

14

Release 6.0 Features

This chapter includes features and enhancements released in 6.0.

Application Packet Filtering



Note:

Prior to Edge 6.0 GA, the objects that perform MOS scoring were originally called “Applications” but have been renamed “Rule Groups” in this, and future, releases.

In Edge 6.0 GA, Applications are a set of one or more rule match criteria, such as IP address, Protocol, DSCP, or Port Number. An Application is a way to put an identifier on a packet when it enters the system to track it. Once a flow has been matched to an Application type, the Application identifier can be used either on the rule or firewall filter as possible match criteria to handle this type of traffic as needed.

Applications

In the Configuration Editor under **Advanced > Global > Applications**, click **Add (+)** to create a new Application that will allow for multiple different criteria.




Figure 1: Add a new Application

Apply the Application to Firewall Policies

Once an Application is created you can then make a firewall policy that will treat all specified match criteria the same way. This can be done from a Global level via **Global > Firewall > Firewall Policy Templates**. This will apply to all firewalls within the network.

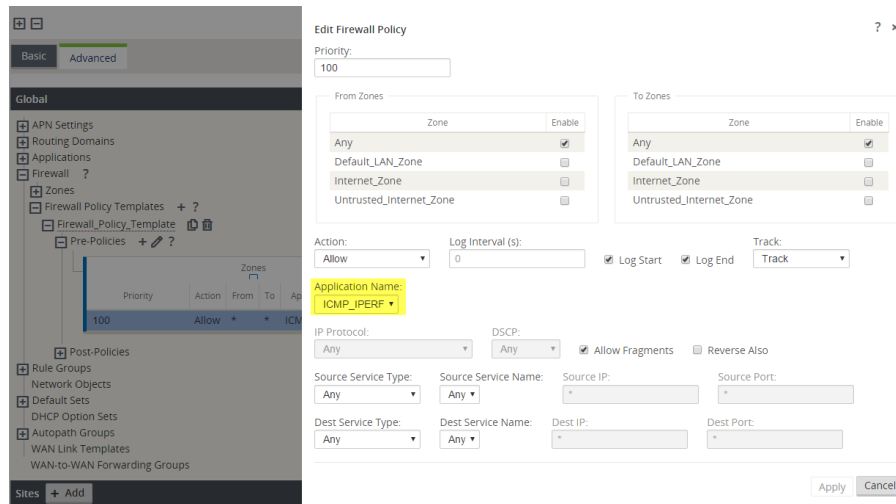


Figure 2: Associating a Global Firewall Policy with an Application

Firewall policies can also be configured from a Site level via **Connections > [Site Name] > Firewall > Policies**. These will only affect traffic at that site.




Figure 3: Associating a Firewall Policy with an Application at the Site Level

Apply the Application to QoS Rules

Once an Application is created you can then make a single QoS rule that will treat all specified match criteria the same. This can be done from a Global level under **Global > Default Sets > Conduit Default Sets > Rules**.

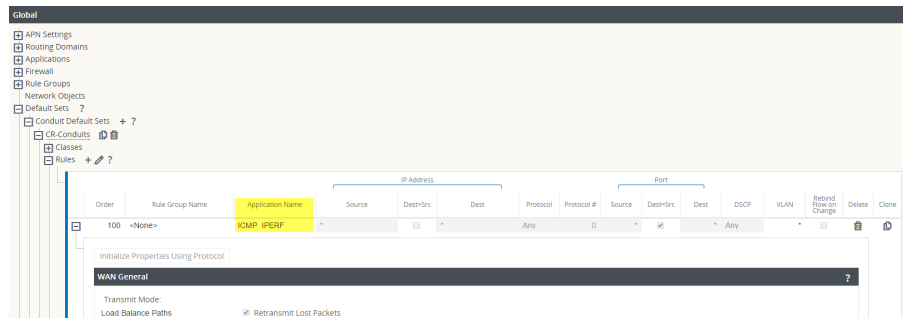


Figure 4: Adding Application to a Global QoS Rule

QoS rules can also be configured, and Applications can be added to them, at a site level under **Connections > [Site Name] > Conduits > [Path Name] > Local Site > Rules**. These will only affect the traffic at that site.

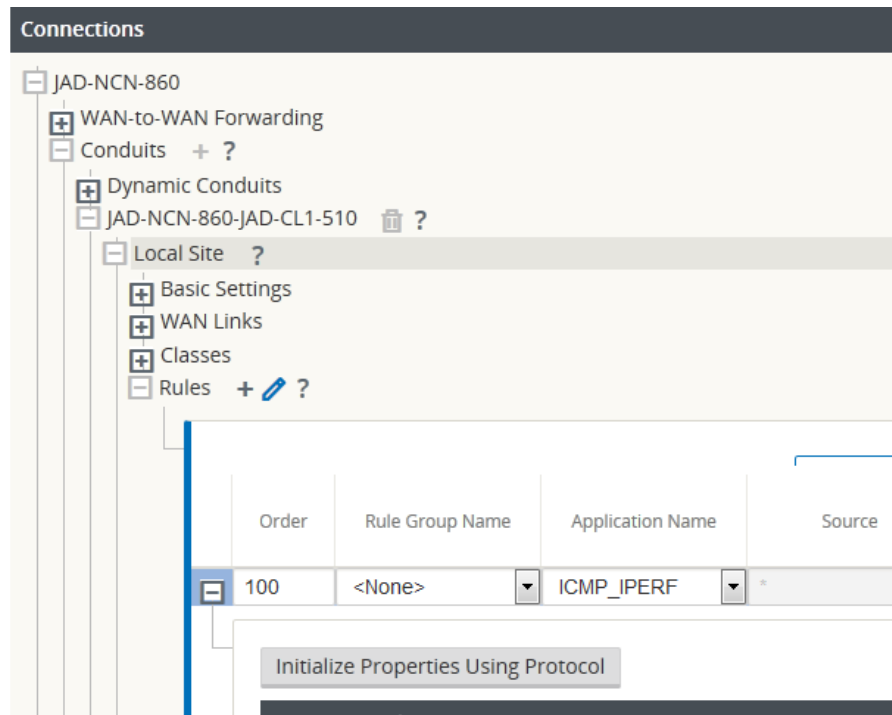


Figure 5: Adding Application to a Local QoS Rule

Tracking Based on Firewall Policy

Users can check to see the statistics for Applications for the Firewall Policy under **Monitor > Firewall** in the web UI and select Applications from the dropdown. This allows users to easily see all connections that match to the selected Application, where they are coming from, where they are going to, and how much traffic they are generating. With this, the user can easily see how their Firewall policies are acting on the traffic for each Application.




Figure 6: Filtering Firewall statistics by Application

Tracking Based on QoS Rule

Users can check to see the status of the current Application for the Rule created under **Monitor > Statistics** in the web UI and select Applications from the dropdown. This allows the user to be able to see at a glance the amount of traffic being generated by a specific Application, and how many sessions are generating it. This can be useful to track bandwidth utilization for specific application types.

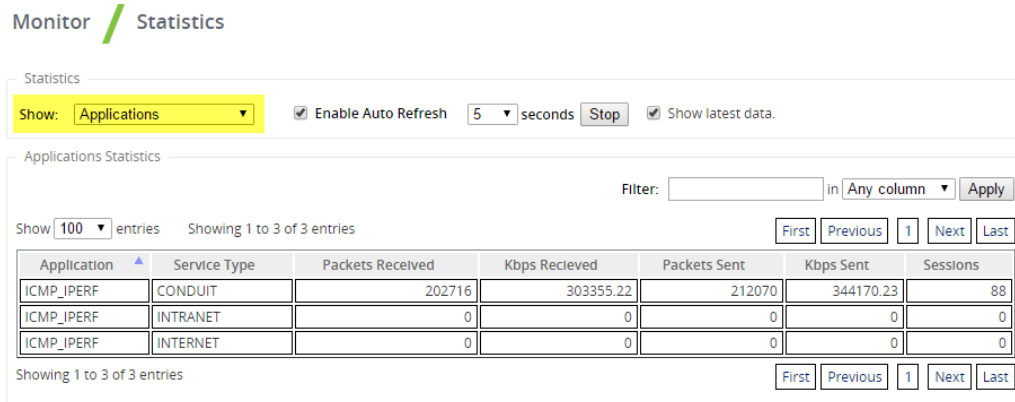


Figure 7: Tracking QoS by Application

VRF Firewall Enhancement

Edge 6.0 GA introduces VRF Firewall enhancements to allow for multiple VRFs, each having access to the Internet. Each VRF is configured to be associated with a different user group, for example, employee or guests, while keeping the traffic from each isolated. This feature allows each Routing Domain (user group) access to the Internet through a common Access Interface. This provides the following capability:

- Local guest-user Internet access
- Employee-user Internet access for defined applications
- Employee-users may continue hairpin all other traffic to the NCN
- Allow the user to add specific routes per Routing Domain, if required
- When enabled, this feature applies to all Routing Domains

Users may also create multiple access interfaces to accommodate separate public facing IP addresses. Either option provides the required security necessary per user group.



Note:

Detailed instructions for how to configure VRFs can be found in the Edge 5.0 *New Features Guide*.

Below are the steps to configure this option:

1. Create Internet Service for a Site under **Connections > [Site Name] > Internet Services** and enable the **Use** checkbox under **WAN Links**.
2. Enable the checkbox labeled **Internet Access for All Routing Domains** under **Sites > [Site Name] > WAN Links > [WAN Link Name] > Access Interfaces**.



Figure 8: Enabling Internet access for All Routing Domains

Selecting this checkbox allows the Edge to use this Access Interface for Internet Service on all configured Routing Domains.

Users may choose to configure either a shared Access Interface or one Access Interface for each group (separate public facing IP addresses).



Note:

After completing the following steps you should see 0.0.0.0/0 routes added, one per Routing Domain, under **Connections > [Site Name] > Routes**.

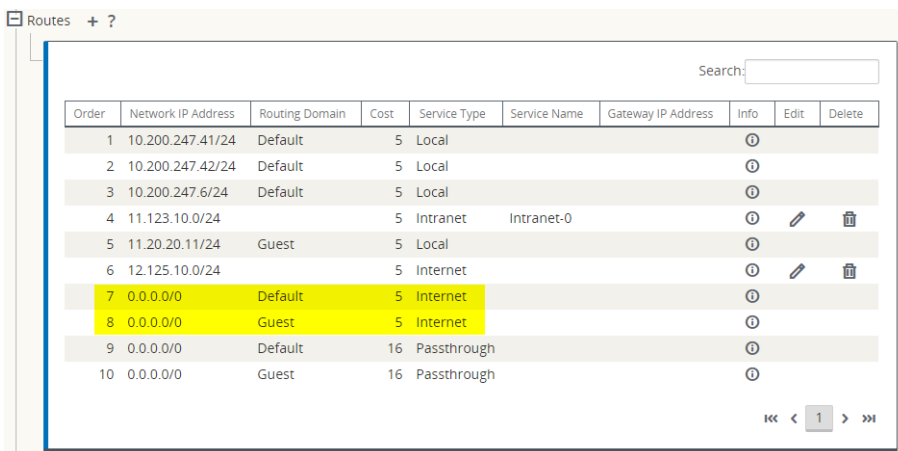


Figure 9: Verifying Routes Added for Each Routing Domain



Note:

It is no longer required to have all Routing Domains enabled at the NCN. Disabling Routing Domains at the NCN that are in use at a Branch site will produce a popup message:




Figure 10: Removing a Routing Domain

Users may confirm that each Routing Domain is using the Internet Service by checking the Routing Domain column in the Flows table of the web UI under **Monitor > Flows**.

Flows List

Both WAN Ingress and WAN Egress Flows

Routing Domain	Source IP Address	Dest IP Address	Direction	Source Port	Dest Port	IP DSCP	HT Count	Service Type	Service Name	LAN GW IP	Age (mS)	Packets	Bytes	PPS	Customer kbps	Conduct Overhead kbps	IPsec Overhead kbps	Rule ID	Class Type	Path	Hdr Compression Saved Bytes	Transmission Type
Guest	11.20.20.20	12.125.10.20	WAN Ingress	8	3355	icmp	default	62	INTERNET	-	LOCAL	74	62	5208	1.013	0.681	0.000	202	N/A	N/A	N/A	N/A
Default	10.200.247.200	12.125.10.20	WAN Ingress	8	16185	icmp	default	66	INTERNET	-	LOCAL	311	66	5544	1.009	0.678	0.000	202	N/A	N/A	N/A	N/A
Guest	12.125.10.20	11.20.20.20	WAN Egress	9	18458	icmp	default	62	INTERNET	-	LOCAL	94	62	5208	1.013	0.681	0.000	202	N/A	N/A	N/A	N/A
Default	12.125.10.20	10.200.247.200	WAN Egress	9	3968	icmp	default	66	INTERNET	-	LOCAL	328	66	5544	1.008	0.678	0.000	202	N/A	N/A	N/A	N/A

Total INGRESS flows displayed: 2 out of 2
Total EGRESS flows displayed: 2 out of 2

Figure 11: Flows in Routing Domains

Users may also check the routing table for each Routing Domain under **Monitor > Statistics > Routes**.

Routes for routing domain: Guest

Filter: [] in [Any column] [Apply]

Show: 100 [v] entries Showing 1 to 5 of 5 entries

Num	Network Addr	Gateway IP Address	Service	Firewall Zone	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	HT Count	Eligible	Eligibility Type	Eligibility Value
0	11.20.20.0/24	*	Local	Default_LAN_Zone	YES	*	Angelina-CFB	Static	-	-	1	0	YES	N/A	N/A
1	11.10.10.0/24	*	DC-Angelina-CFB	Default_LAN_Zone	YES	*	DC	Static	-	-	5	0	YES	N/A	N/A
2	0.0.0.0/0	*	Internet	Untrusted_Internet_Zon	YES	*	*	Static	-	-	5	159	YES	N/A	N/A
3	0.0.0.0/0	*	Passthrough	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
4	0.0.0.0/0	*	Discard	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

Showing 1 to 5 of 5 entries

Figure 12: Flows for a Routing Domain

Easy First Install Simplified Appliance Installation

Oracle Talari Appliances going into new sites and RMA replacement appliances going into existing site can now be implemented with ease. A new Edge can now be connected and powered on by a non-technical person. A Network Administrator can add a new site, configure it at the NCN, and upload the software and configuration package to a registration server. This registration server will reside in the cloud. The Client Appliance is then installed and will acquire an IP address from DHCP. Once the Client Appliance has a valid IP address it will contact the registration server, and based on its serial number, download the corresponding package. Once the package is downloaded, it is activated automatically and the site will become operational.

Configuration using Templates


In Edge 6.0 GA, new customers with five or more basic sites to setup for the first time will enjoy time savings while setting up new sites and WAN Links. With the use of templates, users may configure certain settings one-time and then duplicate the settings across more than one site as needed. This functionality is presented to the user in two key ways. First, the ability to create and administer WAN Link templates. Second a tab, which simplifies the setup of basic sites. Each of these is accessed via the Basic tab. Under the Basic tab, you have the Network option used for WAN Link templates and the Site option, which simplifies the configuration process for a site.

WAN Link Templates

The WAN Link Templates functionality provides users with a way to setup basic configuration for WAN Links and reuse these across the network to save time. The WAN Link Templates feature exists within both the Basic configuration mode and the Advanced configuration mode, with minimal differences between the two modes in Edge 6.0 GA.

Below are the steps to use this feature through the Basic configuration mode:


Manage Network > Configuration Editor > New > Basic.




Click Network to change from the (default) Sites view to Network view.



Click **+ WAN Link Template** to view the Add WAN Link Template screen shown below.



Once a WAN Link Template is added, it will be displayed as one of the WAN Link Templates on the Network view within Basic mode.



Basic Configuration Mode

Edge 6.0 GA introduces the Basic configuration mode as our first step in a larger ease of use evolution. Network administrators with basic sites will be able to reduce repetitive tasks and configure new sites with minimal clicks. Combined with WAN Link Templates (see above) the Basic configuration is a very powerful tool to be up and running with minimal manual configuration.

The concept of the **Basic > Sites** view is to simplify the configuration process to allow the user to create a configuration file, which will generate a Conduit between the defined sites. The required configuration properties for a Conduit between sites include:

- **Appliance**
- **Interface**
- **WAN Links**
- **Static Routes**

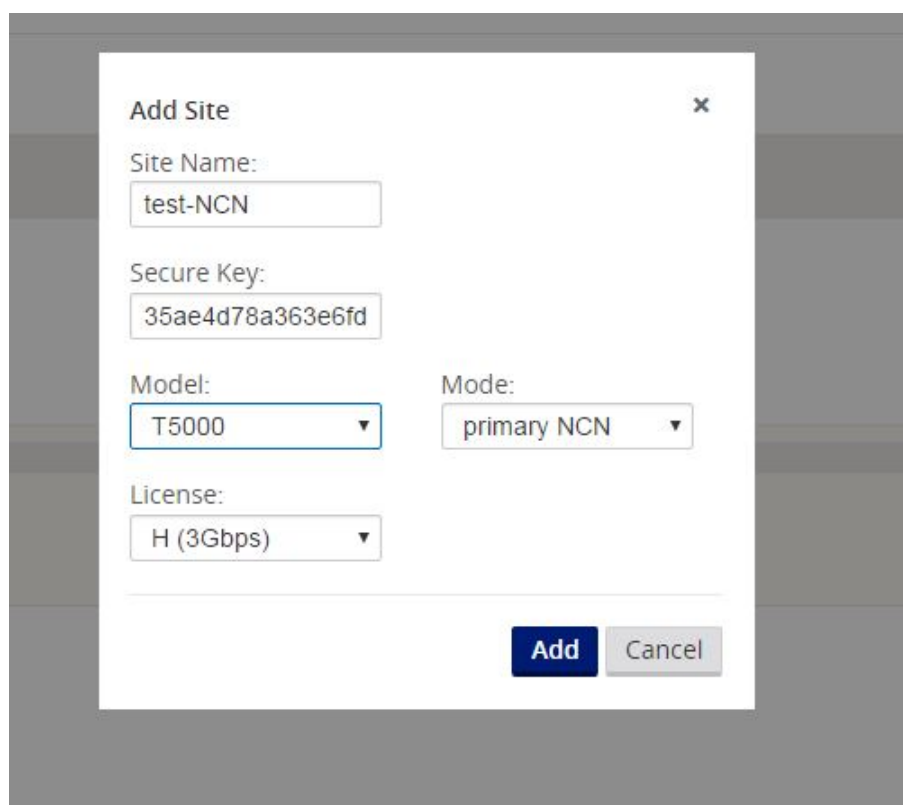
Existing users will observe that one configuration change on the Basic mode view may in fact modify or change more than one setting in Advance mode. Basic mode does allow the Import of existing configurations, and allows the user to move between Basic and Advanced modes.

Below are the steps to use the Basic configuration mode.

Manage Network > Configuration Editor > New > Basic.



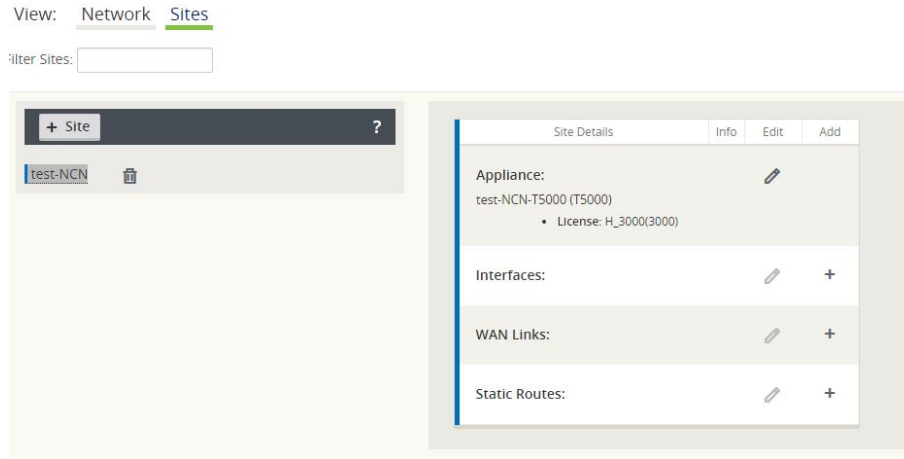
Click **+ Site** and enter basic site details.



Add from the Add Site Dialog will present the basic site details in the site list to the left and display a Site Summary to the right. The Site Summary provides the ability to add, view, and edit site details for interfaces, WAN Links, and Static Routes.

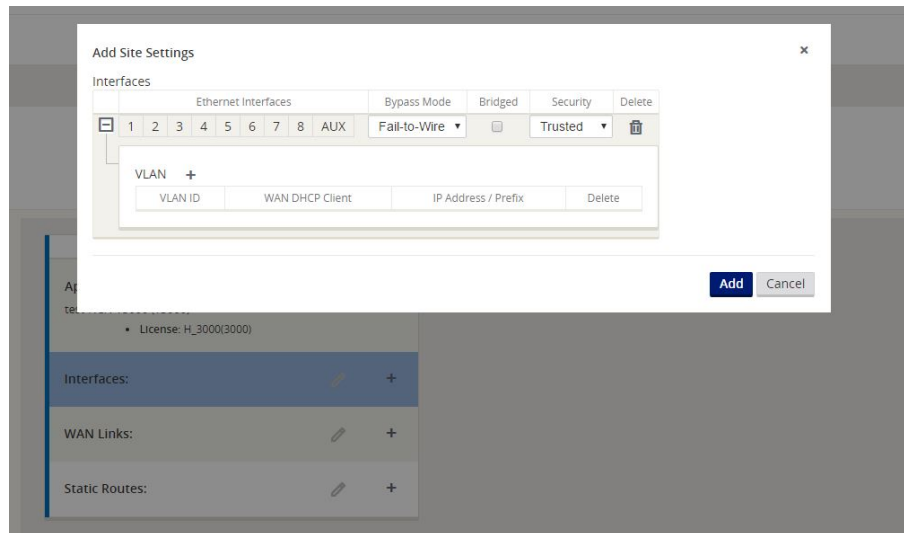
Appliance

From this point forward if the user desires to edit Appliance information just entered in the previous step, they can click the Edit icon to the right of the Appliance settings in the summary view.



Interfaces

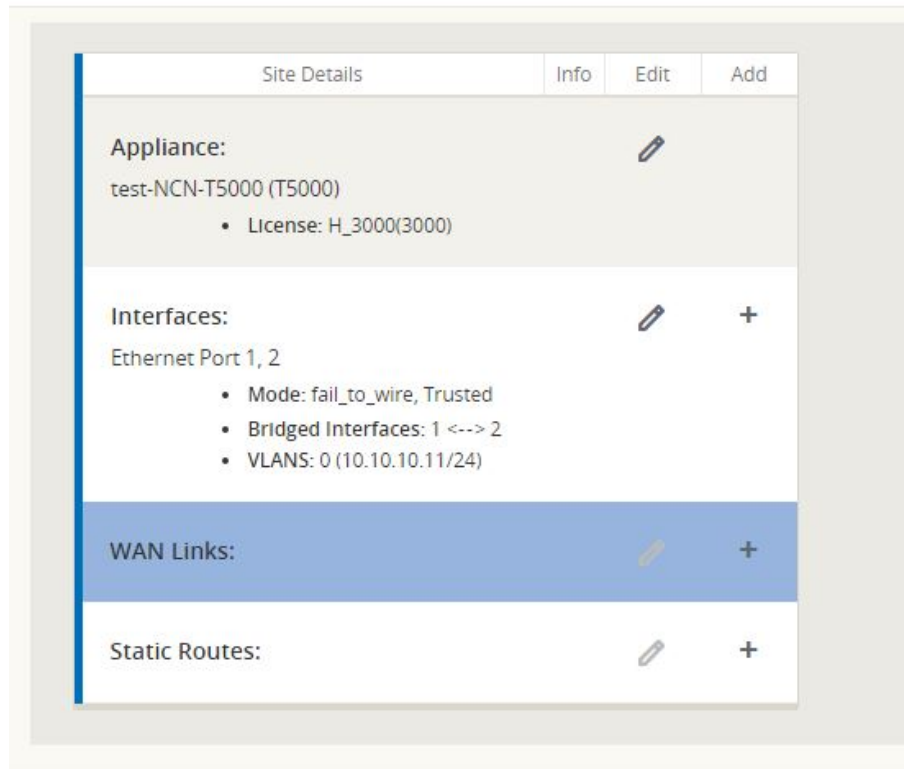
Clicking the Add / Edit Icon to the right of the Interfaces summary view shown for the site will provide the ability to add, edit, and delete Interfaces.



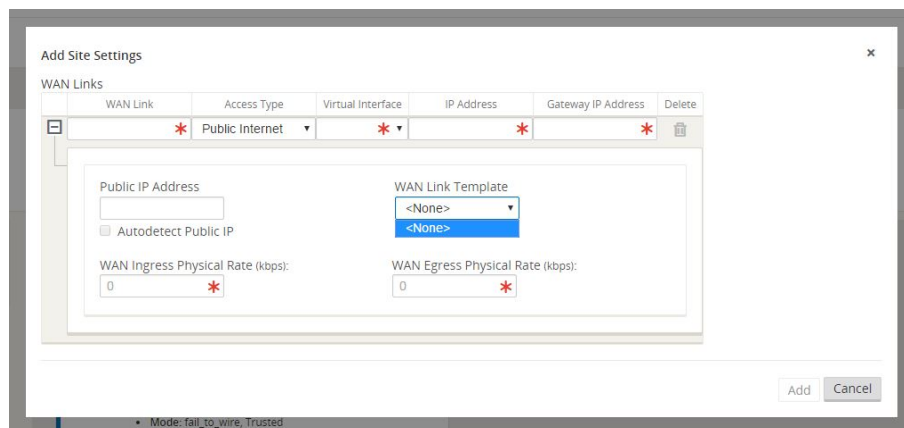
The Interface option allows the user to define the physical topology of the site, such as the ports, logical VLANs and security level for the physical ports. At this level, the user can also define if the WAN interface will use DHCP for an IP address, or they may statically assign an IP address. This allows the user to configure multiple options under the same panel.

WAN Links

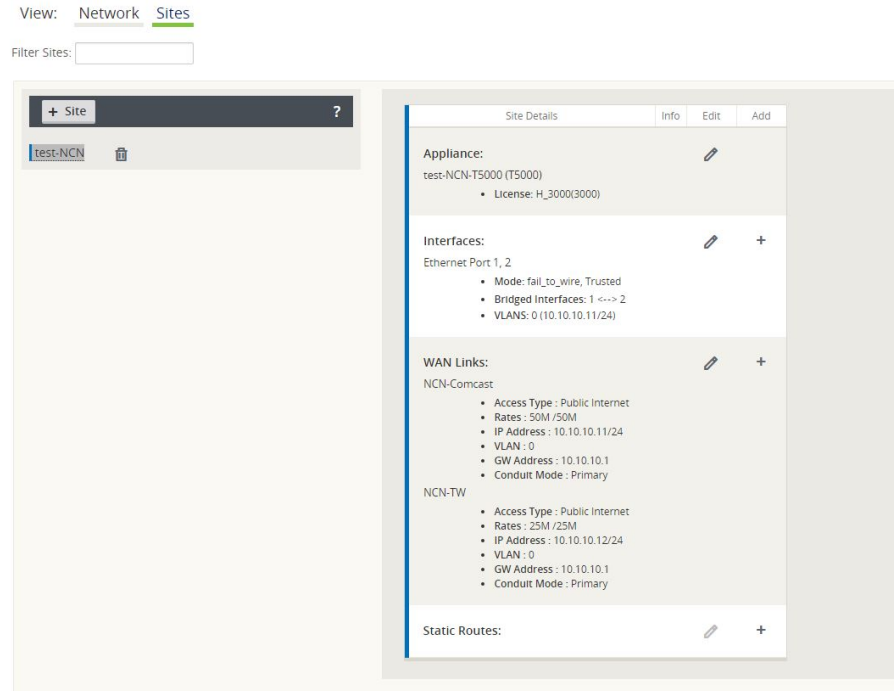
Clicking the Edit Icon to the right of the WAN Links summary view shown for the site will provide the ability to add, edit, and delete WAN Links.



While Adding / Editing a WAN Link, the option to use a WAN Link Template is provided. After selecting a WAN Link Template, the WAN Link will be configured using the WAN Link Template values. The user has the option to overwrite the Template values if desired. Additionally once the Virtual Interface is selected, the IP address is automatically provided from the interface configuration.

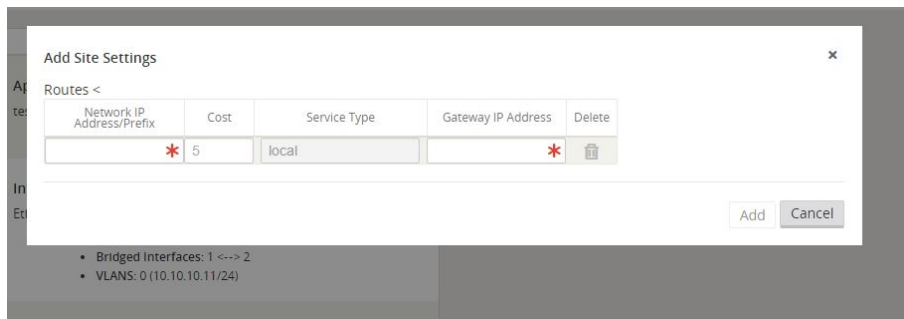


A summary view of WAN Links is then displayed in Basic mode after the initial configuration is complete.



Static Routes

Clicking the Add / Edit icon to the right of the Static Routes area will take the user to the Add / Edit Static Routes dialog. Currently the user can only add local routes within the Basic configuration view.



After configuration, the summary view will display the site information configured and provide the ability to edit all items, as well as add more Interfaces, WAN Links, or Static Routes as needed.

Site Details	Info	Edit	Add
Appliance: test-NCN-T5000 (T5000) <ul style="list-style-type: none">License: H_3000(3000)			
Interfaces: Ethernet Port 1, 2 <ul style="list-style-type: none">Mode: fail_to_wire, TrustedBridged Interfaces: 1 <--> 2VLANs: 0 (10.10.10.11/24)			
WAN Links: NCN-Comcast <ul style="list-style-type: none">Access Type : Public InternetRates : 50M /50MIP Address : 10.10.10.11/24VLAN : 0GW Address : 10.10.10.1Conduit Mode : Primary NCN-TW <ul style="list-style-type: none">Access Type : Public InternetRates : 25M /25MIP Address : 10.10.10.12/24VLAN : 0GW Address : 10.10.10.1Conduit Mode : Primary			
Static Routes: <ul style="list-style-type: none">10.100.10.0/24 via 10.10.10.50 (Local) cost 5			

The Basic view is intended to simplify the configuration process and provide the user the ability to create a configuration file quickly and easily. For more complicated configurations, the user may create a Basic configuration using this mode, then proceed to the Advanced mode to complete the configuration.

Service Chaining

Edge 6.0 GA now provides support for service chaining on the T860 Appliance with OS 5.0. This capability allows the T860 Appliance to run the application natively and support a Guest VM via KVM. This capability is intended for sophisticated partners. For more information on this capability please contact your representative.

15

Release 6.1 P2 Features

This chapter includes features and enhancements released in 6.1 P2.

Site Templates

Users now have the ability to configure Bridge Pairs, VLANs, and Ethernet Interfaces using Site Templates. This reduces configuration complexity when adding branch locations with similar topologies and saves the user time.

To create a Site Template, begin on the Basic tab, select the Network view, and click the + **Site Template** button.




Figure 1

Here, you will select which Ethernet Interfaces should be used, set the Bypass Mode, choose whether to bridge the Interfaces, pick a Security setting, add any required VLANs, and set the WAN DHCP Client option. Click the **Add** button and observe that the New_Site_Template now appears on the left-hand side of the page. To change the template name or edit any of its settings, simply click on it.

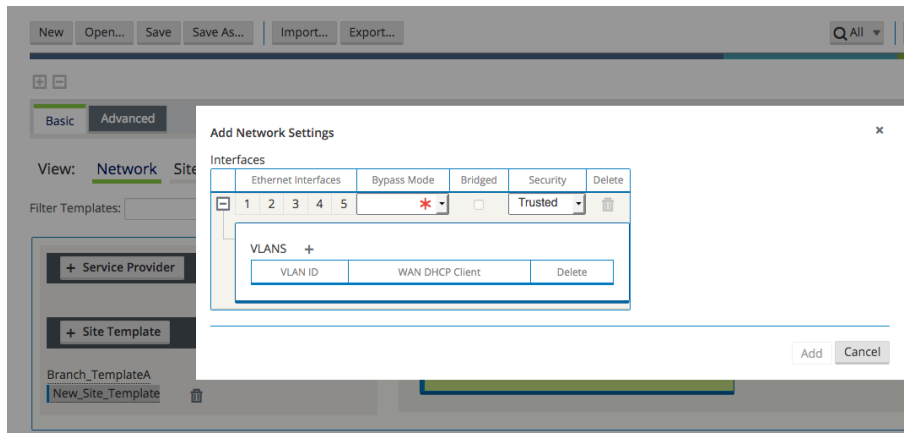


Figure 2

The user will then select which sites should use the template.

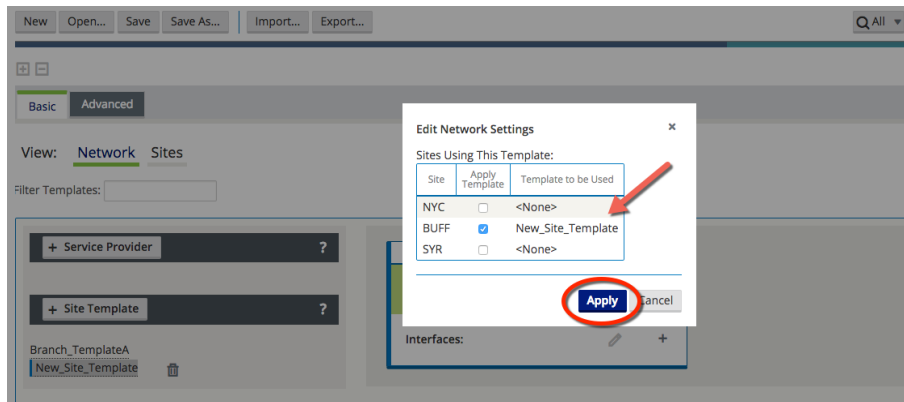


Figure 3

Additionally, users may create a Site Template based on an existing site within the configuration. To do so, change the view to Sites, select the branch site name, and click **Generate Site Template**.

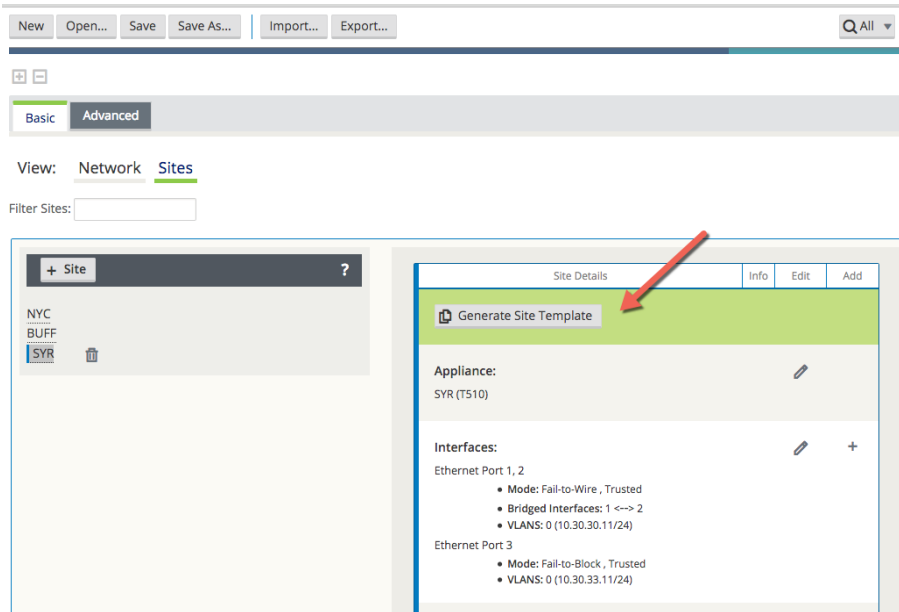


Figure 4

In this example a new Site Template will be generated from the existing site, "SYR". You can confirm the settings in the pop up window and change the template name. This new template has been named Branch_TemplateA.

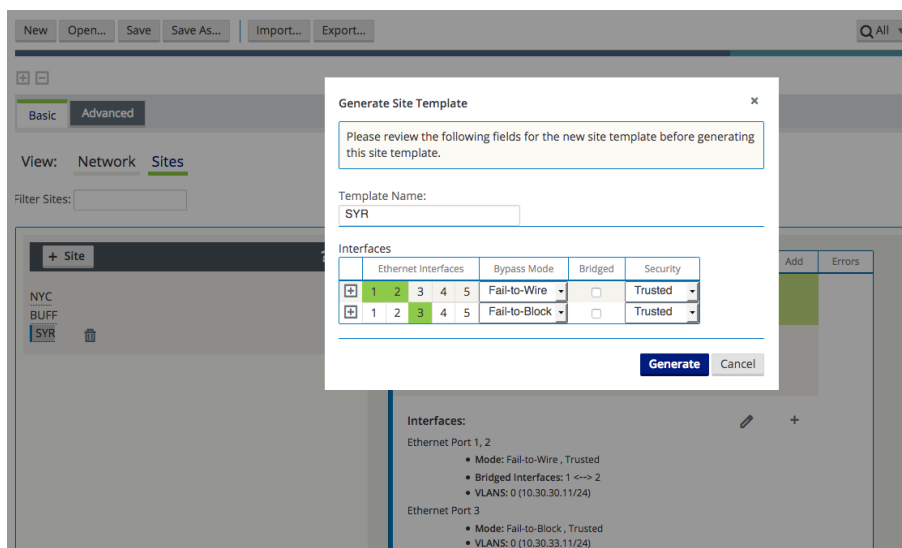


Figure 5

Observe that Branch_TemplateA now appears on the Network view page.

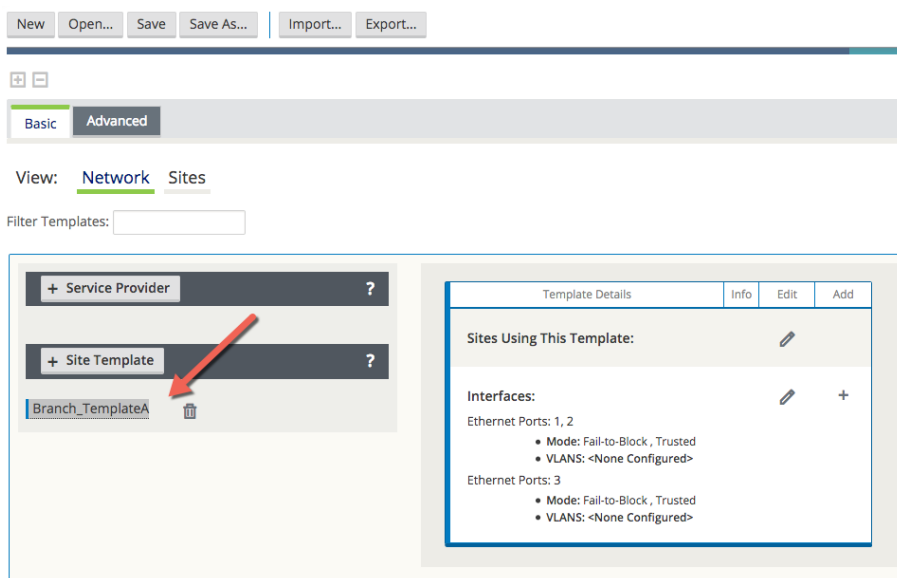


Figure 6

Assign the new template to a site. In the example below, Branch_TemplateA has been assigned to branch site “BUFF”.

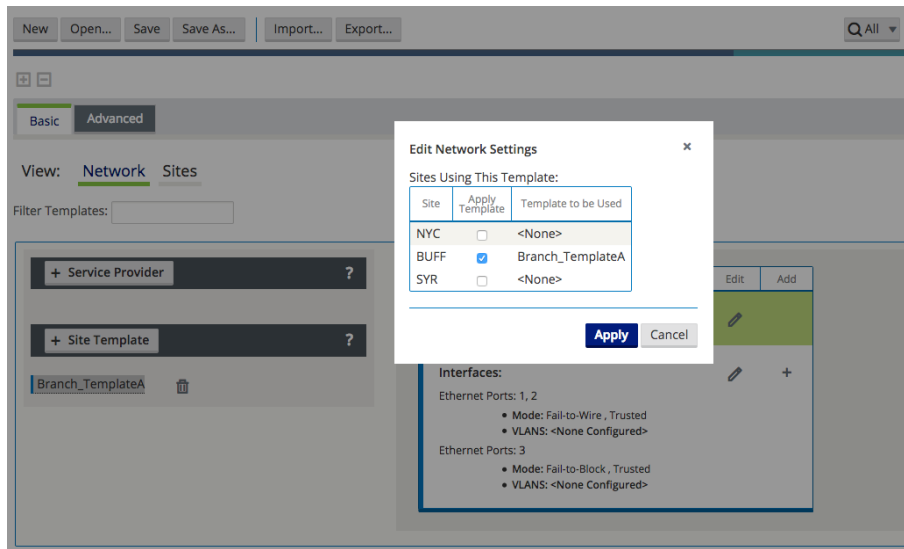


Figure 7

Back on the Sites view, site “BUFF” shows it has been assigned to Site Template Branch_TemplateA.

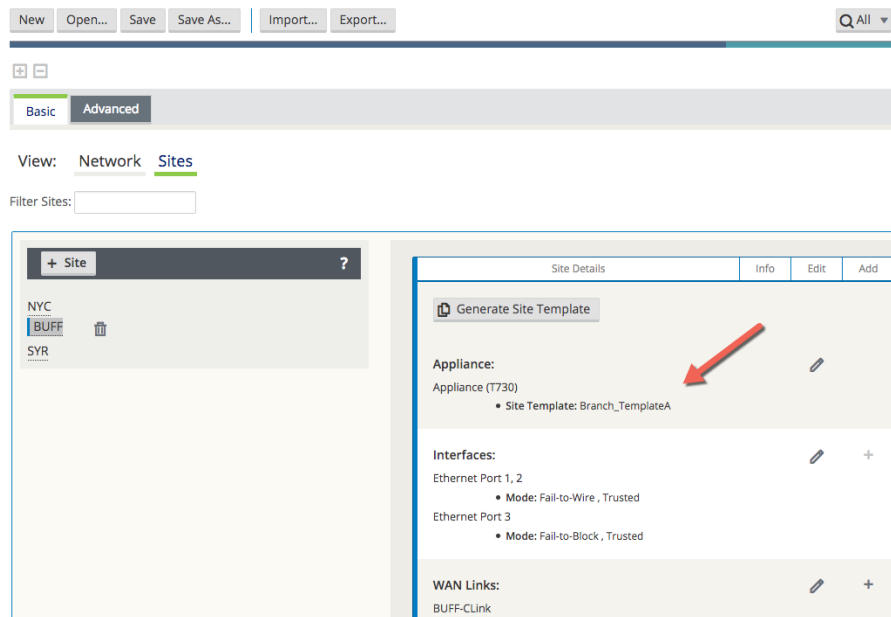


Figure 8

Additional Features in Edge 6.1 GA P2

Additional features included in Edge 6.1 GA P2 include Service Provider – Aware (SP-Aware), OpenDaylight API for service provider configuration, and Restful APIs for service provider network Change Management.

16


Release 7.0 Features

This chapter includes features and enhancements released in 7.0.

WAN Optimization

7.0 GA introduces the capability to perform WAN Optimization on TCP flows, allowing users to simplify branch network infrastructure by consolidating SD-WAN and WAN Optimization services on a single device. WAN Optimization (WANOp) increases efficiency across the WAN for bulk file-transfer traffic, specifically for data requested by more than one user at the same location.


When WAN Optimization is enabled for a flow, TCP Termination is automatically enabled as well. This feature splits a single TCP connection into 3 separate TCP connections, all managed and maintained by the , in to offer maximum throughput and reliable transfer across the WAN via the conduit. The diagram below shows an example WANOp flow between two sites.



WAN Optimization is supported on the E100, T3010v2, T5000v2, and T5200 Oracle Talari Appliance models.

Note:

A one-time reboot is required on all WANOp-capable Oracle Talari Appliances after upgrading to 7.0. The following banner will be displayed if the reboot is necessary:



Session Capacity for Supported Models

Appliance Model	Number of Sessions
E100	8000
T3010v2	8000
T5000v2	16000
T5200	16000

The WANOp solution is configured on a per-rule basis and performs deduplication and compression on TCP Conduit traffic.

Configuring WAN Optimization

Pull up the web UI for the NCN appliance, navigate to **Manage Network > APN Configuration Editor** and **Import** the current configuration file. On the **Advanced** tab, under **Global > Default Sets > Conduit Default Sets > [Conduit Default Set] > Rules**, click the **(+)** icon to create a rule for the type of traffic to be optimized.

Order	Rule Group Name	Application Name	IP Address			Protocol	Protocol #	Port			DSCP	VLAN	Rebind flow on Change	Delete	Clone
			Source	Dest+Src	Dest			Source	Dest+Src	Dest					
100	<None>	Any				NFS	0								

Expand the rule properties. WAN Optimization is enabled via a dropdown menu under the **WAN General** section. When WANOp is enabled, TCP Termination is also enabled by default.

Initialize Properties Using Protocol

WAN General

Transmit Mode:
Load Balance Paths Retransmit Lost Packets

Override Service: <N/A> Preferred WAN Link: Any Persistent Impedance(ms): 50

Traffic Optimization

Enable TCP Termination: On

Enable WAN Optimization: On

Note:

When WANOp is enabled, TCP Termination is enabled for WAN Optimization to function as designed. If desired, the user can also enable TCP Termination independently from the WAN Optimization capability.

A reciprocal rule enabling WANOp will be generated automatically at the remote site of the selected Conduit.

Once your configuration is complete, **Export** it to **Change Management** and follow the prompts through the Change Management process until the new configuration has been Activated.

Verification

To verify that traffic flow is being optimized, navigate to the **Monitor > Flows** page on the NCN. Uncheck the WAN Ingress and WAN Egress Flow Types, and check TCP Termination, then click the refresh button to display only TCP Terminated flows.

The flows table will show detailed information about all TCP Terminated flows, including their WANOp state, as shown below:

Select Flows

Flow Type: WAN Ingress WAN Egress Internet Load Balancing Table TCP Termination Table

Max Flows to Display (Per Flow Type): 50

Filter (Optional): Help

Refresh

Flows Data

TCP Terminated Flows

Source IP Address	Dest IP Address	Source Port	Dest Port	IPP	Age (mS)	From Wan kbps	To Wan kbps	From Lan kbps	To Lan kbps	Bytes Pending To LAN	Bytes Pending To WAN	State	Is WANOp
10.1.1.1	10.1.9.1	80	53258	6	15	1307.620	26042.475	0.000	0.000	0	528	ESTABLISHED	Yes

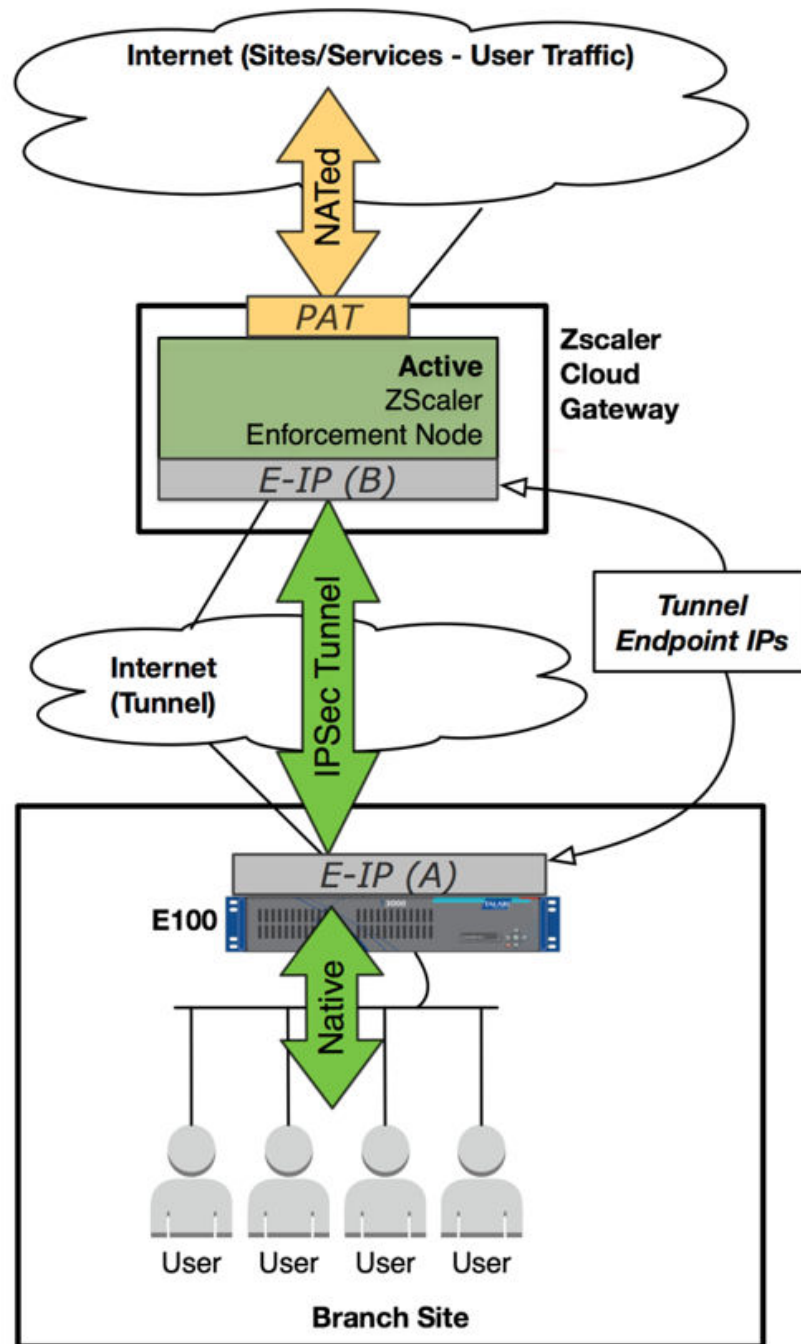
Total TCP Terminated sessions displayed: 1 out of 1. Total WANOp sessions: 1.

For more information on the WAN Optimization solution, including more detailed capabilities, performance, and monitoring options, please see the WAN Optimization Guide.

Zscaler Integration

Zscaler is a Cloud Security Provider (CSP) that delivers key Next Generation Firewall features including Intrusion Detection System (IDS), Intrusion Prevention System (IPS), Data Loss Prevention (DLP), and Sandboxing.

Introduced in APN R7.0 GA, users can now integrate a branch office Oracle Talari Appliance with the Zscaler Cloud Security Gateway via IPSec tunneling, for the purposes of tunneling Internet-destined traffic to Zscaler for cloud-hosted filtering and security services.



Configuration

To configure a Zscaler IPsec tunnel, navigate to **Manage Network > Configuration Editor** on the NCN and **Import** the current configuration file. Click on the **Advanced** tab, expand **Connections > [Site Name] > IPsec Tunnels** and click the **(+)** icon.

Select Zscaler as the **Service Type**, select the **Local IP** address, fill in the **Peer IP** address of the Zscaler Enforcement Node (ZEN), enter the IKE **Pre-Shared Key**, and click **Apply**.

The screenshot shows the configuration page for an IPsec Tunnel. At the top, a table lists the tunnel details:

Service Type	Name	Routing Domain	Firewall Zone	Local IP	Peer IP	MTU	Keepalive	Delete
Zscaler	ShapelHill-Client_Zscaler1	RD-BMC	Internet_Zone	10.60.20.12	104.129.194.39	400		

Below the table are the configuration sections:

- IKE Settings:**
 - Version: IKEv1
 - Mode: Main
 - Identity: Auto
 - Authentication: Pre-Shared Key
 - Pre-Shared Key: [Redacted]
 - Validate Peer Identity:
 - DH Group: Group 2 (MODP1024)
 - Hash Algorithm: SHA1
 - Encryption Mode: AES 128-Bit
 - Lifetime (s): 86400
 - Lifetime (s) Max: 86400
 - DPD Timeout (s): 20
- IPsec Settings:**
 - Tunnel Type: ESP-NULL
 - PFS Group: <None>
 - Hash Algorithm: SHA1
 - Lifetime (s): 28800
 - Lifetime (s) Max: 86400
 - Lifetime (KB): 0
 - Lifetime (KB) Max: 0
 - Network Mismatch Behavior: Drop
- IPsec Protected Networks:** (Empty table with columns for Source IP/Prefix, Destination IP/Prefix, and Delete)

The 'Apply' button at the bottom right is highlighted with a red circle.

Note:

When you add an IPsec tunnel with a Service Type of **Zscaler**, the following default configurations will be applied that are not applied when selecting **LAN** or **Intranet Service Types**.

- Firewall – Adds a Deny policy from Default_LAN_Zone to Untrusted_Internet_Zone.
- NAT – Deletes the default outbound PAT policy, if one exists.
- Routing – Adds a 0.0.0.0/0 route over the Zscaler tunnel and a /32 host-route of the tunnel Peer IP to the gateway.

Save the configuration and **Export** it to **Change Management**. Follow the Change Management process to **Stage** and **Activate** the new configuration.

Verification

Once the new configuration is running, follow the steps below to verify functionality.

1. Generate Internet traffic from a host on the LAN to a URL that has been blocked by Zscaler.
2. Verify the Zscaler IPsec Tunnel status in the web UI of the Oracle Talari Appliance under **Monitor > Statistics > IPsec Tunnel**.

Statistics

Show: IPsec Tunnel Enable Auto Refresh 5 seconds Refresh Show latest data.

IPsec Tunnel Statistics

Show 100 entries Showing 1 to 1 of 1 entries Filter: in Any column Apply

Name	Routing Domain	State	Service Type	Packets Received	Kbps Received	Packets Sent	Kbps Sent	Packets Dropped	Bytes Dropped	MTU
CH-Zscaler	RD-BMC	GOOD	Internet	565	1151.35	636	235.46	0	0	1348

Showing 1 to 1 of 1 entries

3. Verify the flow of the generated traffic through the Oracle Talari Appliance via **Monitor > Flows**. Once you have identified the flow, confirm the Service Type as INTERNET.

Monitor / Flows Talari Support

Select Flows

Flow Type: WAN Ingress WAN Egress Internet Load Balancing Table TCP Termination Table

Max Flows to Display (Per Flow Type): 50

Filter (Optional): 443 Help

Refresh

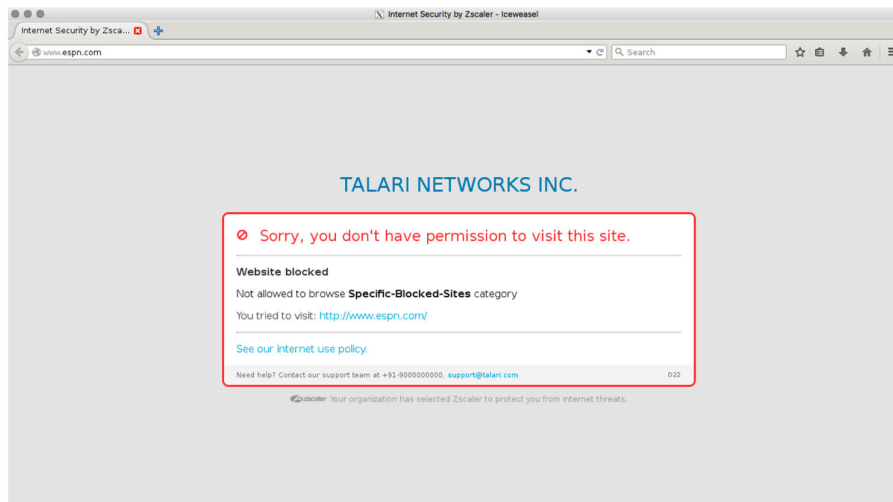
Flows Data Toggle Columns

Both WAN Ingress and WAN Egress Flows

Routing Domain	Source IP Address	Dest IP Address	Direction	Source Port	Dest Port	IP	DSCP	HL Count	Service Type	Service Name	LAN GW IP	Age (mS)	Packets	Bytes	PPS	Customer kbps	Conduit Overhead kbps	IPsec Overhead kbps	Rule ID	Class Type	Class	Path	Hdr Compression Saved Bytes	Transmission Type
RD-BMC	10.60.130.20	172.217.8.174	WAN Ingress	53020	443	TCP	default	14	INTERNET	CH-Zscaler	LOCAL	16640	0	0	0.000	0.000	0.000	260	N/A	N/A	N/A	N/A	N/A	
RD-BMC	10.60.130.20	34.223.209.188	WAN Ingress	43903	443	TCP	default	16	INTERNET	CH-Zscaler	LOCAL	3951	0	0	0.000	0.000	0.000	0.000	260	N/A	N/A	N/A	N/A	N/A
RD-BMC	10.60.130.20	35.161.11.107	WAN Ingress	46251	443	TCP	default	44	INTERNET	CH-Zscaler	LOCAL	2063	0	0	0.000	0.000	0.000	0.000	260	N/A	N/A	N/A	N/A	N/A
RD-BMC	172.217.8.174	10.60.130.20	WAN Egress	443	53020	TCP	default	9	INTERNET	CH-Zscaler	LOCAL	16654	9	1590	0.039	0.012	0.000	0.016	260	N/A	N/A	N/A	N/A	N/A
RD-BMC	34.223.209.188	10.60.130.20	WAN Egress	443	43903	TCP	default	13	INTERNET	CH-Zscaler	LOCAL	3942	13	4617	0.116	0.048	0.000	0.048	260	N/A	N/A	N/A	N/A	N/A
RD-BMC	35.161.11.107	10.60.130.20	WAN Egress	443	46251	TCP	default	4	INTERNET	CH-Zscaler	LOCAL	2053	4	6739	0.176	0.046	0.000	0.046	260	N/A	N/A	N/A	N/A	N/A

Total INGRESS flows displayed: 3 out of 8
Total EGRESS flows displayed: 3 out of 7

4. Verify Zscaler is blocking the traffic.



In the event the IPsec Tunnel between the Oracle Talari Appliance and the Zscaler ZEN goes down, the 0.0.0.0/0 route through the tunnel will become unreachable and pulled from the routing table. Traffic will hit the next available, reachable 0.0.0.0/0 route out to the Internet. Route reachability can be verified in the Oracle Talari Appliance's web UI under **Monitor > Statistics > Routes**.

Note:

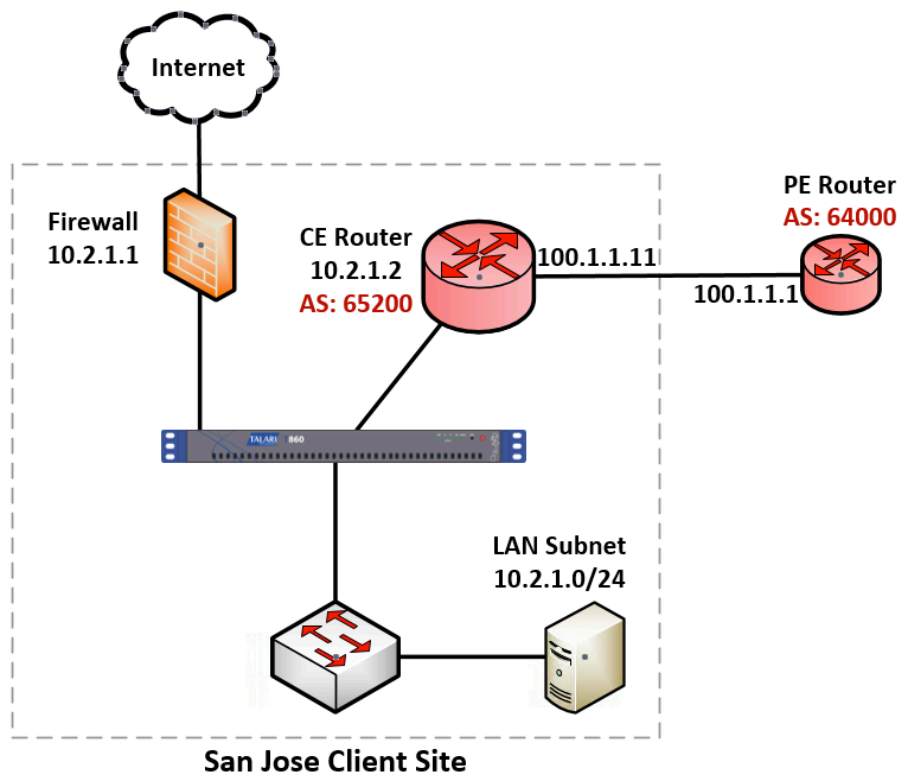
R7.0 GA only supports a single VRF/routing domain for Zscaler.

Customer Edge (CE) Router Replacement Within the APN


Oracle SD-WAN Edge 7.0 introduces the ability to replace a Customer Edge Router with a Adaptive Private Network Appliance. This is accomplished by leveraging the APNA's ability to masquerade its Local Autonomous System (AS) number (on a per-neighbor basis) so that it can peer with a Provider Edge (PE) Router in a manner consistent with a traditional Customer Edge (CE) Router. The APNA can peer with other BGP neighbors as well, using either its true Local AS number or a masqueraded AS number.

Installation Summary

Sample APN site before replacing the CE Router with the APNA.



Sample Edge site after replacing the CE Router with the APNA.

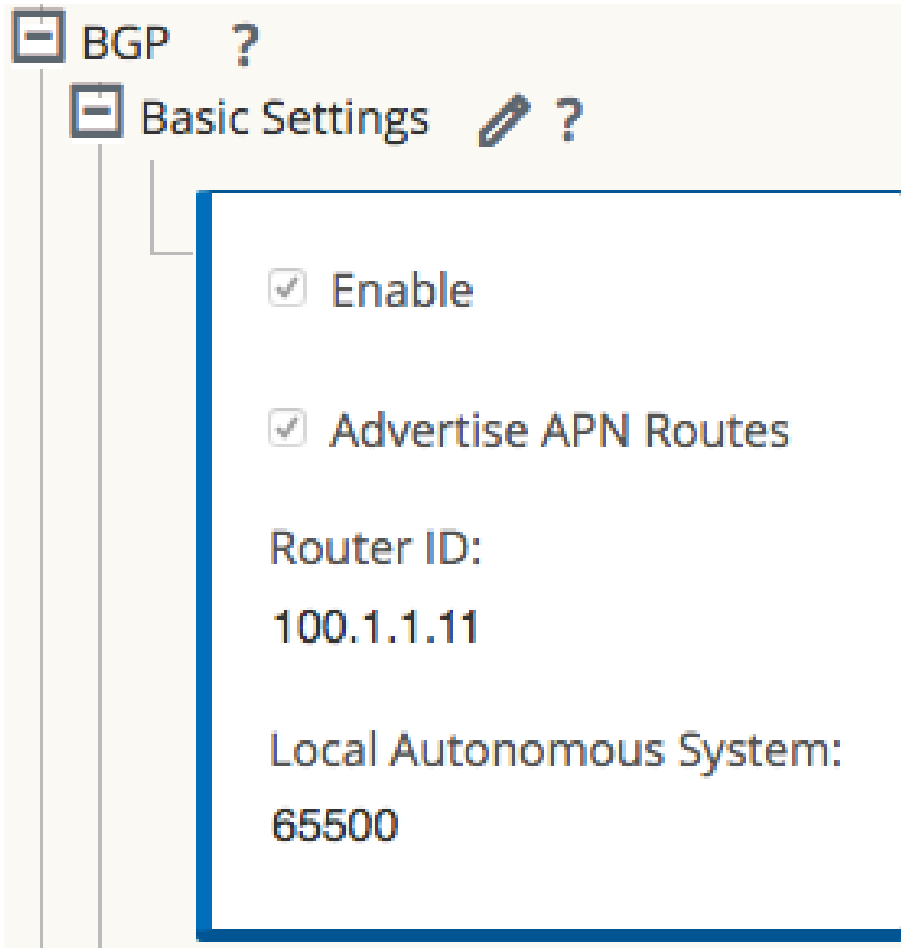


The CE router is removed and the APNA peers directly with the PE router via eBGP by masquerading its AS number as the replaced CE router's AS number (AS 65200). The APNA's actual Local AS number is 65500 and it can peer via iBGP with local routers in this AS.

If desired, APNAs can also peer with each other via iBGP over a Conduit. This allows Edge to act as an Autonomous System. The primary use-case intended for Edge as an Autonomous System consists of the primary NCN, and secondary NCN if required, are configured as Route-Reflectors, and Clients using an iBGP peering session to the NCN(s) for BGP reachability information.

BGP Configuration

Using the Configuration Editor, navigate to **Connections > [Site Name] > Route Learning > BGP > Basic Settings** and click the pencil icon to edit.



Check the **Enable** box to enable BGP on the APNA. If it is desirable to advertise Edge routes to BGP peers, check the **Advertise Routes** box. Enter an optional **Router ID** and enter the **Local Autonomous System** number.

Neighbors

Use the (+) icon to the right of the **Neighbors** section to add BGP neighbor entries.

Virtual Interface	Source IP	Local AS (AS Masquerade)	Neighbor IP	Remote AS	Hold Time(s)	Local Preference	IGP Metric	Route Reflector Client	Disable Local AS Loop Protection	Password	Delete
VI_port_3	100.1.1.11	65200	100.1.1.1	64000	180	100	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		<input type="button" value="Apply"/> <input type="button" value="Revert"/>

Choose the appropriate **Virtual Interface**, enter the **Local AS** number or **AS Masquerade number**, and enter the **Neighbor IP** address. In this example, we are using the **AS Masquerade** number 65200, to match the AS Number of the former CE Router.

Note: If the Local AS field in the **Neighbors** section is left blank, the default behavior is to use the Local AS defined in the previous step under **Basic Settings**. If no Local AS is defined in either of these sections, no AS number will be used.

The following options may also be set:

- **Hold Time(s)** - Time in seconds to wait before declaring a neighbor as DOWN.
- **Local Preference** - Sets the BGP attribute Local Preference for routes learned from the neighbor specified.

- **Route Reflector Client** - The APNA will act as a Route Reflector and the neighbor will be treated as a Route Reflection Client.
- **Disable Local AS Loop Protection** - By default, BGP routes learned that contain the APNA's Local AS number in the AS path will be rejected to guard against routing loops. This can be disabled for situations in which learned routes are prepended with the APNA's Local AS number to influence path selection in BGP.
- **Password** - Used if the BGP session requires MD5 authentication.

Import and Export Filters

Now that BGP is enabled and neighbors have been configured, the Import Filters can be configured under **Connections > [Site Name] > Route Learning > Import Filters**.

By default, no routes will be imported until Import Filters have been added, as the default filter rejects all route advertisements. Expand the Import Filters section and use the **(+)** icon to add a filter.

Order	Source Router	Destination	Prefix	Next Hop	Protocol	Route Tag	Cost	Include	Enabled	Delete	Clone
100	*	<Manual>	eq *	*	Any	*	eq *	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		
(auto)	*	<Manual>	eq *	*	Any	*	eq *	<input type="checkbox"/>	<input type="checkbox"/>		

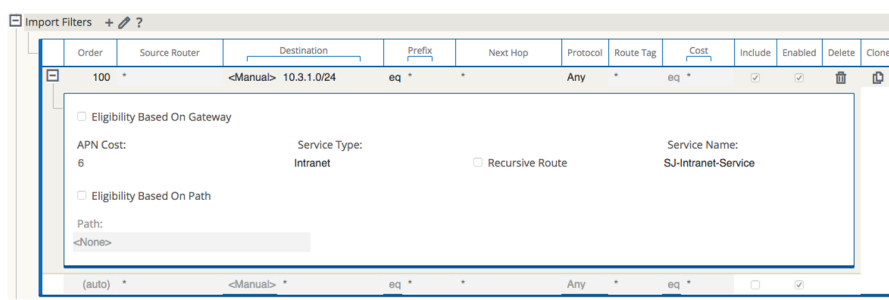
Note:

For each added filter, use any combination of the **Destination**, **Prefix**, and **Next Hop** fields to match desired BGP routes to learn. If these fields are left with their default value of **(*)**, all advertised BGP routes will be imported. Additionally, it is important to understand the impact of the **Include** and **Enabled** checkboxes. If **Include** is checked, routes that match the filter will be imported. On the same filter, if **Include** is not checked, then routes that match the filter will not be imported. The **Enabled** checkbox simply enables or disables the filter entirely.

Use the **(+)** icon to the left of the **Order** column to reveal Edge specific options. Click the **Service Type** dropdown box to expose the available options. Depending on the Service Type chosen, various additional options will be available and are listed below.

- **Export Route to Oracle Talari Appliances:** If the Export Route to Oracle Talari Appliances checkbox is enabled, the Oracle Talari Appliance will communicate route data to Oracle Talari Appliances at other sites if WAN-to-WAN forwarding is enabled. This functionality is enabled by default but only applies to the following Service Types: Local and LAN GRE Tunnel.
- **Eligibility Based on Gateway:** If the gateway becomes unreachable, this feature will ensure that traffic is not sent to matching routes.
- **APN Cost:** The cost will be applied to the matched routes when importing into the Oracle Talari Appliance's route table. The default APN Cost is 6.
- **Service Type:** Choose a Service Type from all the existing, supported Services.
- **Recursive Route:** When the Service Type is Conduit, check this option to find the Conduit name from an imported route's source router automatically.
- **Service Name:** The name of the service that matching routes will use.

- **Eligibility Based on Path:** If enabled, Path state becomes criteria for filters.



Once configuration of the APN is complete, the configuration should be saved and **Change Management** should be used to push the configuration changes to the APNAs.

Static Routes File

Oracle Talari Appliances provide a **Static Routes** file that can be edited to define routes that should persist through software and configuration changes made to the APN. This is used for inserting static routes into the dynamic routing table, not the APN routing table. It ensures that any necessary static routes are advertised to the PE router after the CE router replacement, regardless of changes to the APN configuration. By default, static routes defined in this file will be advertised to all neighbors within the specified routing domain.

BGP Verification and Troubleshooting

After the replacement, login to the web UI of the APNA and navigate to **Monitor > Statistics** to verify that the change is successful.



This will bring up the **Paths (Summary)** statistics page. Verify that **Path State** and **Conduit State** report GOOD for each WAN Link as shown in the image below.

Num	From Link	To Link	Path State	Conduit State	Conduit Type	BOWT	Jitter (ms)	Loss %	kbps	Congestion
1	RAL-NCN-Intet-WL	Sanjose-CLI-Intet-WL	GOOD	GOOD	Static	2	2	0.00	13.49	NO
2	RAL-NCN-Intet-WL	Sanjose-CLI-Intet-WL	GOOD	GOOD	Static	2	2	0.00	14.40	NO
3	Sanjose-CLI-Intet-WL	RAL-NCN-Intet-WL	GOOD	GOOD	Static	2	2	0.00	18.31	NO
4	Sanjose-CLI-Intet-WL	RAL-NCN-Intet-WL	GOOD	GOOD	Static	2	2	0.00	14.10	NO

Next, use the dropdown menu to select **Routes** to verify that the expected routes are properly being learned via BGP. In the example below, notice the 10.3.1.0/24 route shows **Type** as Dynamic and **Protocol** as BGP.

Num	Network Addr	Gateway IP Address	Service	Firewall Zone	Reachable	Site IP Address	Site	Type	Protocol
0	100.1.1.0/24	*	Local	Default_LAN_Zone	YES	*	Sanjose-CLI	Static	-
1	10.2.1.0/24	*	Local	Default_LAN_Zone	YES	*	Sanjose-CLI	Static	-
2	10.1.1.0/24	*	RAL-NCN-Sanjose-CLI	Default_LAN_Zone	YES	*	RAL-NCN	Static	-
3	10.3.1.0/24	100.1.1.12	SJ-Intranet-Service	Default_LAN_Zone	YES	*	*	Dynamic	BGP

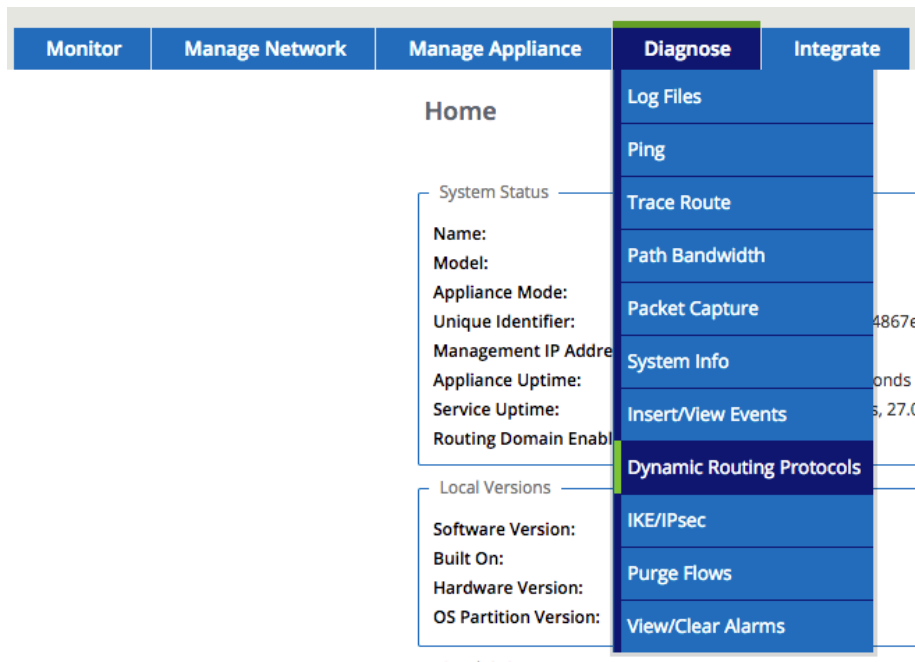


Note:

The route must also be considered reachable for it to be used.

BGP Troubleshooting Enhancements

The Oracle Talari Appliance's web UI provides tools to gather information about the Dynamic Routing Protocols you have enabled. These tools can be found under **Diagnose > Dynamic Routing Protocols**.



Below are descriptions of each option. When a view allows filtering, enter the Network Address and Mask in the format shown below.

Dynamic Routing Protocol

View: BGP Show Route Table Protocol for NWAddress/Mask Routing Domain: Default_RoutingDomain

Network Address: * 10.1.1.0 Mask: * 24 Submit Refresh

- **BGP State** - Shows an overview of the current state of each Dynamic Routing Protocol instance.
- **BGP Show Route Table Protocol** - Shows prefixes associated with each BGP instance/neighbor.
- **BGP Show Route NWAddress/Mask Table** – Shows prefixes associated with each BGP instance/neighbor and allows filtering for specific prefixes. Will provide APN and BGP routes.
- **BGP Show Route Table Protocol NWAddress/Mask** - Shows prefixes associated with each BGP instance/neighbor and allows filtering for specific prefixes. Provides BGP routes only.
- **BGP Show Route Export** - Shows routes being advertised from the Oracle Talari Appliance.
- **BGP Show Route Export (detailed)** - Shows routes being advertised from the Oracle Talari Appliance, as well as routing protocol attributes.
- **BGP Show Route Preexport** - Shows all applicable routes for advertisement.

- **BGP Show Route Preexport (detailed)** - Shows all applicable routes for advertisement, as well as routing protocol attributes.
- **Show Route Table**- Provides an overview of each route prefix.
- **Show Route Table (detailed)** - Provides an overview of each route prefix and protocol-specific attributes such as Next Hop, Local Preference, AS Path, etc.
- **Show Route Count in Table** - Gives a count of all entries in the routing table (BGP and APN).
- **Show Protocol** - Outputs a list of routing protocols that are currently running and their states.
- **Talari Protocol Table** - Shows only the Edge routing table.
- **Appliance ifconfig** - Shows the output of the “ifconfig” command to provide the user detailed information about each active interface port.
- **BGP Configure** - Reloads the advanced routing configuration.
- **BGP Restart** - Restarts all routing protocols.

For additional information on this topic (including how to edit the Static Routes file) please refer to the CE Router Replacement Guide on the Support Portal section of our website under Documentation.

E100 as an NCN

R7.0 GA now supports deployment of the E100 Oracle Talari Appliance as a primary and secondary NCN for up to 8 Client sites (9 total sites per-network). This is done in the Configuration Editor from the **Advanced** tab under **Sites > [Site Name] > Basic Settings** where the **Model** should be the E100 and the **Mode** can now be either primary NCN or secondary NCN.

Global

Sites + Add

NCN1

Basic Settings

Appliance Name: NCN Secure Key: 2311243cff Regenerate

Model: E100 Mode: primary NCN

Site Template: <None>

Default Direct Route Cost: 5

Gateway ARP Timer (ms): 1000

Enable Source MAC Learning

After completing the configuration, the user will **Export** it to **Change Management** and follow the prompts to create a package for the E100 appliance. Once you have uploaded the package to the E100, the Home Page will reflect that the E100 is functioning as the NCN Appliance.

TALARI Networks Monitor Manage Network Manage Appliance Diagnose Integrate

Home

System Status

Name: E100-1000

Model: E100

Appliance Mode: **NCN**

Unique Identifier: 615270160079


Management IP Address: 192.168.8.90

Appliance Uptime: 1 hours, 45 minutes, 56.9 seconds

Service Uptime: 41 minutes, 53.0 seconds

Routing Domain Enabled: Default_RoutingDomain

Note: If you try to push a configuration through Change Management where the primary or secondary NCN is an E100 and you have defined more than 8 Client sites (resulting in more than 9 total sites per-network), the configuration will not pass the Validation Check.



Capacity Report for the E100 NCN

Appliance Model	T510	T730	T750	T860	E100	T3010	T5000	T5200
Supported as NCN	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Max Client Sites as NCN	N/A	N/A	32	32	8	128	256	550
Max Static Conduits	8	16	32	32	32	128	256	550
Max Dynamic Conduits	4	8	16	16	16	32	32	32
Max WAN Ingress Paths	36	72	216	216	216	576	1152	5500
Max WAN Egress Paths	36	72	216	216	216	576	1152	5500
Max Flows (TCP Term off)	64,000	64,000	64,000	64,000	64,000	256,000	512,000	512,000
Max Flows (TCP term on)	500	4,000	8,000	8,000	8,000	16,000	16,000	16,000
Max Public WAN Links	3	8	8	8	8	8	8	8
Max Private WAN Links	32	32	32	32	32	32	32	32
Max Routes (Static & Dynamic)	16,000	16,000	16,000	16,000	16,000	16,000	16,000	16,000
Max Recommended Routing Domains ¹	16	16	16	16	16	16	16	16

1

NetFlow (Support for Version 9 and IPFIX)

While NetFlow v5 is the default setting, users now have the ability to export flow information using NetFlow v9 and IPFIX. To enable NetFlow and select the version, navigate to **Integrate > NetFlow Host Settings** from the web UI of any Talari Appliance, click the **Enable NetFlow** button, then select the preferred version from the **NetFlow Version** drop down.

TALARI Networks | Monitor | Manage Network | Manage Appliance

Integrate / NetFlow Host Settings

NetFlow Host Settings

Enable NetFlow

NetFlow Version: 5, 9 (selected), IPFIX

NetFlow Host 1: IP Address [] Port []

NetFlow Host 2: (Optional - can be left blank.) IP Address [] Port []

NetFlow Host 3: (Optional - can be left blank.) IP Address [] Port []

Apply Settings | Clear Settings

© 2017 Talari Networks

Note:

To complete the configuration, you must enter in a **NetFlow Host IP Address** and **Port** number, then click the **Apply Settings** button.

Additional Features in 7.0 GA

Additional features included in 7.0 GA include support for the VT800 at the data rate of 500Mbps through a Conduit on the Azure cloud and the VT800 at the data rate of 2Gbps for ESXi on the Intel Xeon E7-8870v4.

Appliance Model	Throughput 1400B at MOS score of 4.3 or Better
T510	2 x 40 Mbps
T730	2 x 80 Mbps
T750	2 x 120 Mbps
E100	2 x 200 Mbps
T860	2 x 800 Mbps
T3010	2 x1 Gbps
T5000	2 x 3 Gbps
T5200	2 x 5 Gbps
VT800 (ESXi)	2 x 2 Gbps
VT800 (Azure)	2 x 500 Mbps
VT800 (Hyper-V)	2 x 200 Mbps
CT800	2 x 100 Mbps

17

Release 7.1 Features

This chapter includes features and enhancements released in 2.3.

E1000 Hardware Options

7.1 introduces three hardware variations for the E1000 in the form of optional expansion cards. Customers may order either four additional fail-to-wire Gigabit Ethernet ports or two 10 Gigabit Ethernet fiber ports. To determine which expansion card (if any) is installed on the appliance, check the number of Ethernet interfaces under **Manage Appliance > Local Network Settings**:


Hardware Option	Ports
E1000 without expansion card	AUX, MGT, interfaces 1-8
E1000 with 10G expansion card	AUX, MGT, interfaces 1-10
E1000 with FTW expansion card	AUX, MGT, interfaces 1-12

Port Labelling for Expansion Cards

Port 9

Port 10

10G Fiber (2 Port) Expansion Card:



The E1000 with 10G fiber expansion card does not ship with SFPs. The following modules are supported in conjunction with this card:

Description	Intel Part
Intel (Short Range) Dual Rate 10GBASE-SR/1000BASE-SX (Supplier Part FTLX8571D3BCVIT1 or AFBR-709DMZ-IN2)	E10GSFPSR

Intel (Long Range) Dual Rate 10GBASE-LR/1000BASE-LX (Supplier Part FTLX1471D3BCVI31)	E10GSFPLR
Intel Ethernet SFP+ 10GbE direct attach passive copper Twinaxial Cable (Available in 1 Meter, 3 Meter, and 5 Meter lengths)	1 Meter: XDACBL1M 3 Meter: XDACBL3M 5 Meter: XDACBL5M

Port 9

Port 12

Port 11

Port 10

Fail to Wire Copper (4 Port) Expansion Card:



Note:

The configuration editor will not detect which expansion card (if any) is installed on an E1000, and will offer port 1 – 12 for all E1000s. Before beginning configuration for an E1000, please verify the physical ports on the appliance. For an **E1000 with no expansion card**, only configure ports 1 – 8. For an **E1000 with the 10G fiber** expansion card, only configure ports 1 – 10; ports 9 and 10 should not be configured for FTW. For an **E1000 with the FTW** expansion card, all 12 ports may be configured.

If a configuration that does not match the available hardware is applied to an E1000, the Talari service will be disabled. Once a mismatched configuration has been applied to an E1000, a corrected package must be applied through Local Change Management before the service will start. Alternately, the appliance may be factory defaulted and a corrected configuration applied using the Easy 1st Install process. **The Talari service will be disabled until a corrected package is applied.**

For more information about the E1000, available hardware options, and special configuration considerations, please see the *E1000 Installation Guide* and the *E1000 Hardware Guide*.

Interactive Dashboard



7.1 enhances the Oracle Talari Appliance Home page, providing users with an interactive dashboard which provides at-a-glance insight into the APN and quick access to areas of interest:

Name: PPark-Primary

Talari Support

System Status			
OS Version:	5.1	Hardware Version:	E1000v1
Appliance Mode:	Active NCN	Software Version:	R7_1_GA_11142017
Management IP Address:	192.168.20.10	Appliance Uptime:	2 weeks, 4 days, 1 hours, 39 minutes, 12.9 seconds
Unique Identifier:	E20109170022	Service Uptime:	20 minutes, 24.0 seconds
Local Appliance:	Active	Peer Appliance:	Standby
		Last Update Received:	0 seconds ago

- Paths Summary
- Classes
- WAN Link Usage
- Diagnostic Logs



Recent Event Errors
No messages to display.

Dead Conduits

- Conduit 'PPark-AZ-TNET-East' is currently dead.
- Conduit 'PPark-Azure-WestUS' is currently dead.
- Conduit 'PPark-kgault' is currently dead.

Live Conduits

Conduit PPark-ahodge:	Uptime: 18 minutes, 41.0 seconds.
Conduit PPark-AWS-VA-USA:	Uptime: 19 minutes, 53.0 seconds.
Conduit PPark-Azure-EastUS:	Uptime: 15 minutes, 7.0 seconds.
Conduit PPark-Bangalore:	Uptime: 19 minutes, 48.0 seconds.
Conduit PPark-cparsons:	Uptime: 19 minutes, 49.0 seconds.
Conduit PPark-dpereira:	Uptime: 19 minutes, 51.0 seconds.
Conduit PPark-dwilhelm:	Uptime: 19 minutes, 51.0 seconds.
Conduit PPark-jdickey:	Uptime: 19 minutes, 51.0 seconds.
Conduit PPark-jhill:	Uptime: 19 minutes, 51.0 seconds.
Conduit PPark-josswald:	Uptime: 19 minutes, 50.0 seconds.



System Status and Quick Links

Name: PPark-Primary

Talari Support

System Status			
OS Version:	5.1	Hardware Version:	E1000v1
Appliance Mode:	Active NCN	Software Version:	R7_1_GA_11142017
Management IP Address:	192.168.20.10	Appliance Uptime:	2 weeks, 4 days, 1 hours, 39 minutes, 12.9 seconds
Unique Identifier:	E20109170022	Service Uptime:	20 minutes, 24.0 seconds
Local Appliance:	Active	Peer Appliance:	Standby
		Last Update Received:	0 seconds ago

[Paths Summary](#)
[Classes](#)
[WAN Link Usage](#)
[Diagnostic Logs](#)



Recent Event Errors
No messages to display.

Dead Conduits

- Conduit 'PPark-AZ-TNET-East' is currently dead.
- Conduit 'PPark-Azure-WestUS' is currently dead.
- Conduit 'PPark-kgault' is currently dead.

Live Conduits

- Conduit PPark-ahodge: Uptime: 18 minutes, 41.0 seconds.
- Conduit PPark-AWS-VA-USA: Uptime: 19 minutes, 53.0 seconds.
- Conduit PPark-Azure-EastUS: Uptime: 15 minutes, 7.0 seconds.
- Conduit PPark-Bangalore: Uptime: 19 minutes, 48.0 seconds.
- Conduit PPark-cparsons: Uptime: 19 minutes, 49.0 seconds.
- Conduit PPark-dpereira: Uptime: 19 minutes, 51.0 seconds.
- Conduit PPark-dwilhelm: Uptime: 19 minutes, 51.0 seconds.
- Conduit PPark-jdickey: Uptime: 19 minutes, 51.0 seconds.
- Conduit PPark-jhill: Uptime: 19 minutes, 51.0 seconds.
- Conduit PPark-josswald: Uptime: 19 minutes, 50.0 seconds.


The new

landing page moves the appliance name to the top of the page. Local system information has been consolidated under the System Status heading. Additional system information (including model, service uptime, and the build date for software) has been moved to the **Diagnose > System Info** tab.

Just below the System Status information, quick links are provided to commonly used screens:


- The Paths Summary button takes users to the **Monitor > Statistics** Path Summary report.
- The Classes button takes users to the **Monitor > Statistics** Classes report, with the Conduit Filter pre-set to the first Conduit in the list.
- The WAN Link Usage button takes users to the **Monitor > Statistics** WAN Link Usage report.
- The Diagnose button takes users to the **Diagnose > Log Files** page.

Site Usage



The Site Usage graphs show an overview of the past 24 hours of WAN Ingress and WAN Egress bandwidth usage for the site, broken down by Conduit, Intranet, and Internet. Hover over a point on the graph to view a tooltip showing the date, time, and exact values for Conduit, Intranet, and Internet bandwidth at that time. These graphs are linked to the **Monitor > Usage Reports** page. Clicking on a graph will take users to the Usage Reports page to view the requested report in more detail.

Path Status



This bar chart provides a visual display of Conduit and Path states. Red (Dead) segments are clickable. A Dead Conduits segment links to **Monitor > Statistics > Conduits**, pre-filtered to show only dead Conduits. A Dead Paths segment links to **Monitor > Statistics > Paths (Summary)**, pre-filtered to show only dead Paths.

Recent Event Errors

Recent Event Errors	
Event Time	Description
2017-11-02 03:53:47	Watchdog update too late, interval 41248313 uSecs, limit 2000000 uSecs

Displays up to 10 recent errors, if any Events with an Event Type of "Error" are available to display.

Conduits

Dead Conduits	
Conduit 'PPark-e1000-test' is currently dead.	
Conduit 'PPark-josswald' is currently dead.	

Live Conduits	
Conduit PPark-AWS-VA-USA:	Uptime: 16 hours, 15 minutes, 12.0 seconds.
Conduit PPark-Azure-EastUS:	Uptime: 16 hours, 15 minutes, 19.0 seconds.
Conduit PPark-Azure-WestUS:	Uptime: 16 hours, 15 minutes, 9.0 seconds.
Conduit PPark-Bangalore:	Uptime: 16 hours, 15 minutes, 9.0 seconds.

Any dead Conduits are displayed above the list of live Conduits. Clicking on a dead Conduit will take the user to **Monitor > Statistics > WAN Link Usage**, where the Usage and

Permitted Rates table will be automatically filtered based on the service name for the Conduit.

Live Conduits provide hyperlinks to the remote site (on the NCN only).

WAN Optimization on Virtual Appliances

7.1 expands support for WAN Optimization to the VT800 and CT800 platforms. WAN Optimization is supported on these platforms at the following levels, with the following resources:



Platform	License Level	WANOp Capacity	VCPUs	RAM	Max WANOp Sessions	Disk Size	Cloud Instance Type
VT800 for ESXi	20 Mbps	8 Mbps	2	8GB	1,500	160GB	
VT800 for ESXi	2 Gbps	200 Mbps	14 (2.10GHz)	16GB	10,000	160GB	
VT800 for Azure	20 Mbps	8 Mbps	4	28GB	10,000	160GB	DS12_v2
VT800 for Azure	500 Mbps	100 Mbps	8 (2.4GHz)	56GB	16,000	160GB	DS13_v2
VT800 for Hyper-V	20 Mbps	8 Mbps	2	8GB	1,500	160GB	
VT800 for Hyper-V	200 Mbps	100 Mbps	10 (2.10GHz)	10GB	5,000	160GB	
CT800 for AWS	20 Mbps	8 Mbps	8	15GB	5,000	160GB	c3.2xlarge
CT800 for AWS	200 Mbps	50 Mbps	8	15GB	5,000	160GB	c3.2xlarge

Note:

The maximum number of WANOp sessions is scaled based on available memory. If a virtual appliance has insufficient dedicated RAM, the maximum number of WANOp sessions will be lower. Provisioning a virtual appliance below recommended system specifications will not disable WANOp, but will impact WANOp performance. Provisioning a virtual appliance below the defined minimum specifications is not supported.

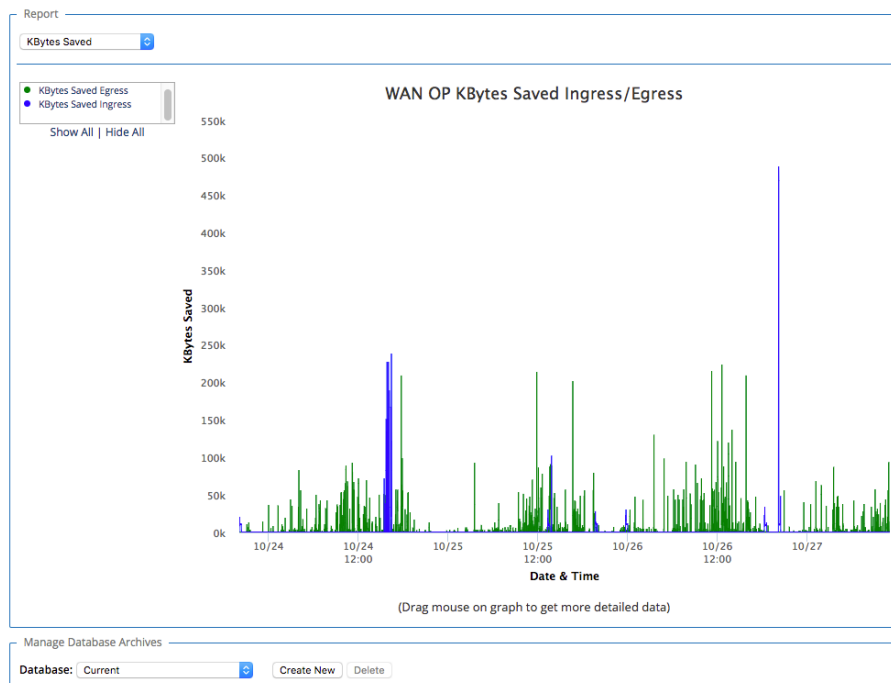
A warning banner will be displayed in the Web Console if WANOp is enabled on a VT800 or CT800 that does not meet the minimum recommended system specifications. An example is shown below, on a VT800 with insufficient RAM and VCPUs:

Warning:

-  WAN Optimization will likely have degraded performance unless at least 8 GB of RAM are allocated to the appliance. The system currently only has 4.06 GB.
-  WAN Optimization will likely have degraded performance unless at least 2 cores are allocated to the appliance. The system currently only has 1.

WAN Optimization Reporting Enhancements

7.1 enhances the existing WAN Optimization monitoring facilities with graphical reports to display key WANOp data over time. To view the reports, navigate to **Monitor > WAN Optimization**.




Different reports may be selected from the dropdown in the upper left-hand corner of the Report pane, and filtered using the criteria to the left of the graph. Available reports include:

- Kilobytes Saved (Egress/Ingress)
- Compression Ratio (Egress/Ingress)
- Deduplication Ratio (Egress/Ingress)
- Data Reduction Percentage (Egress/Ingress)
- Deduplication Cache Percentage (Hit/New)
- Deduplication Cache Count (Hit/New)
- Deduplication Cache Kilobytes (Hit/New)

For in-depth descriptions of the information provided in these reports, please see the *WANOp Setup and Configuration Guide*.

Users can zoom in to view a period of time in greater detail by dragging on the graph to select the timeframe of interest:



To zoom out to the original graph data, click the Zoom Reset button in the upper right corner of the pane.

Users may also view reports for archived databases by selecting the report they wish to view, then scrolling down to the Manage Database Archives pane and selecting an archived database from the dropdown. Changing reports after selecting an archived database will reset the report to the current database.

Additional Features in 7.1 GA

Increased Maximum Bandwidth on AUX Port

7.1 increases the maximum WAN link bandwidth for interface groups including the AUX port to 500Mbps on the T3010, T5000, and T5200.

Monitor > Statistics Enhancements

7.1 introduces some slight changes to the Monitor > Statistics page in the Web Console:

- In the Paths (Summary) view, the Congestion column has been removed.
- In the Classes view, the Conduit Filter field has been replaced with a dropdown menu. Additionally, any entry in the Dropped Packets column with a value greater than 0 will be highlighted.
- In the ARP view, entries are automatically sorted by reply state.
- In the WAN Link Usage view, the Usage % column has been added to the “Local WAN Links” and “Usage and Permitted Rates” tables. The Usage % is Kbps/ Permitted Kbps.

WANOp Intelligent Cache

7.1 optimizes the performance of WANOp caching for items accessed multiple times, ensuring faster speeds for frequently accessed files.

18

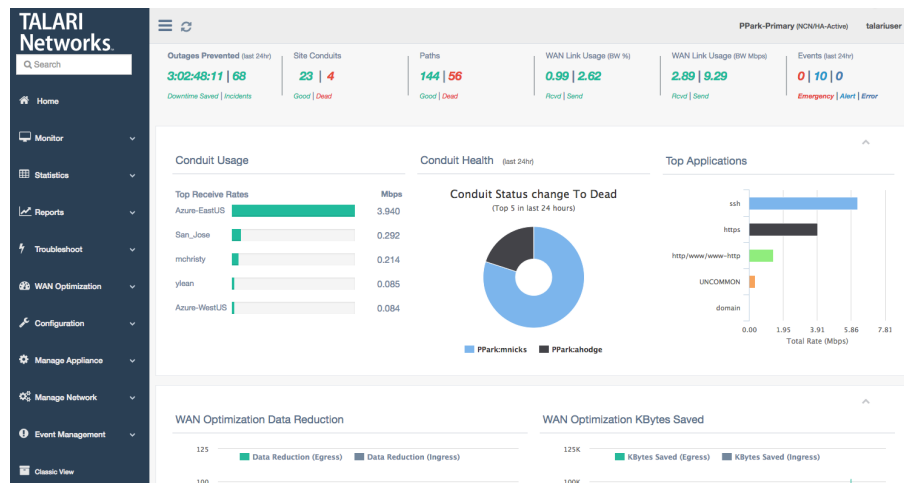
Release 7.2 Features

This chapter includes features and enhancements released in 7.2.

User Interface Enhancements

7.2 GA introduces a new and improved user interface, including a new landing dashboard and updated navigation. The new dashboard and all statistics screens are responsive for easier viewing on varying screen sizes.


Navigation menus have been moved to the sidebar, and reorganized into logical groupings to make the navigation experience more intuitive. Additionally, users may quickly locate any navigation link using the Search bar at the top of the navigation sidebar.



New Landing Dashboard

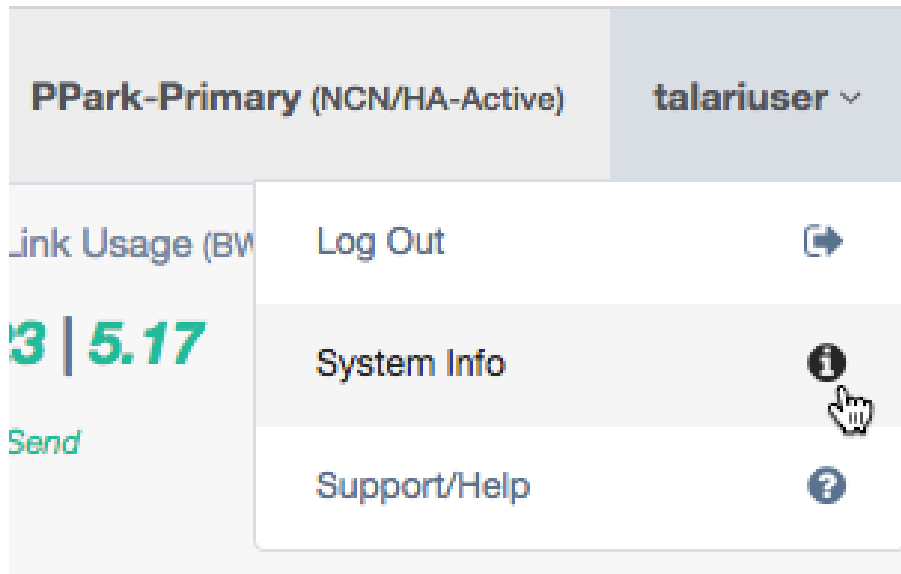
Landing Dashboard Components

The top bar will be visible on every page of the new User Interface:




Toggle visibility of the navigation menu and refresh the current page using the buttons on the left-hand side of the top bar. On the right-hand side, the site name and logged-in username are displayed.


Click on the username to display a dropdown menu with links for logout, system information, and support:



Click System Info in the dropdown to display the system status information, including OS Version, Hardware Version, Software Version, and Management IP:



Appliance alerts will also be shown in the top bar, if any exist:



Click the alert icon to display a popup listing all appliance alerts.

The first section of the new dashboard presents summary information about the health of the site:



Outages Prevented (Last 24 hours): The amount of downtime saved (calculated as time when Paths went dead but the Conduits stayed up) and number of incidents (times Paths went dead) detected in the last 24 hours. Links to **Reports > Availability**.

Site Conduits: The number of Good and Dead Conduits for the site. Each number links to **Statistics > WAN > Conduits**, filtered by the appropriate state.

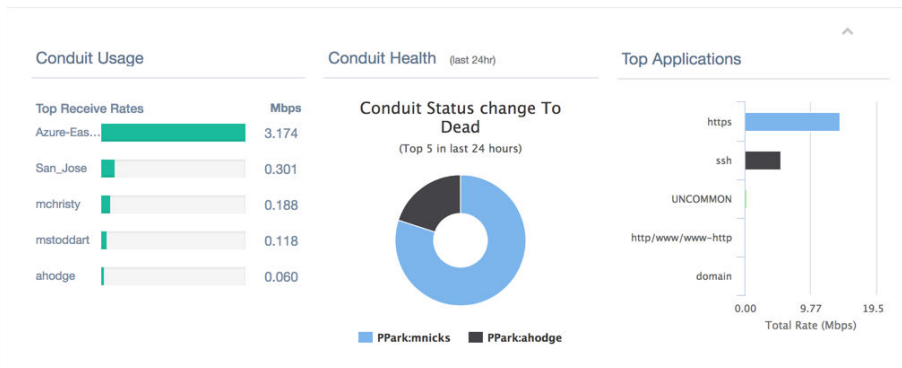
Paths: The number of Good and Dead paths for the site. Each number links to **Statistics > WAN > Paths (Summary)**, filtered by the appropriate state. Any Bad paths are included in the count of Good paths, as this is a transitory path state.

WAN Link Usage (BW %): The local WAN Link Usage by percentage of total permitted rate, for receive (download) and send (upload). Links to **Statistics > WAN > WAN Link Usage**.

WAN Link Usage (Mbps): The local WAN Link Usage by Mbps, for receive (download) and send (upload). Links to **Statistics > WAN > WAN Link Usage**.

Events (Last 24 hours): The number of emergency/alert/critical-level severity events in the last 24 hours. Links to **Event Management > Insert/View Events**.

The second section of the dashboard provides information about Conduit health:



Note:


This section and all of the following sections can be collapsed and expanded using the ^ button in the upper right corner.

Conduit Usage: The top five receive rates in Mbps for the local site, sorted in descending order. The report title links to **Statistics > WAN > Conduits**.

Conduit Health: The five Conduits which have gone dead the most in the last 24 hours. If fewer than five Conduits have gone dead in the last 24 hours, only those which have gone dead will be shown. Users may hover over each section of the graph to see total state changes for that site within the last 24 hours. The report title links to **Statistics > WAN > Conduits**.

Top Applications: The top five observed protocols and the total rate for each. The report title links to **Statistics > QOS > Observed Protocols**.

The third section of the dashboard provides at-a-glance information about WAN Optimization, and will only be displayed on appliances that support WAN Optimization:



WAN Optimization Data Reduction: Report displaying data reduction percentages for WAN Egress and WAN Ingress WANOp traffic at the site on a per-minute basis. Users may hover over a point on the chart to display detailed data. Select the legend headers to turn data points off or on. The report title links to **WAN Optimization > Statistics**, with the data reduction report displayed.


WAN Optimization Kbytes Saved: Report displaying kilobytes saved using WAN Optimization for WAN Egress and WAN Ingress. Users may hover over a point on the chart to display detailed data. Select the legend headers to turn data points off or on. The report title links to **WAN Optimization > Statistics**, with the Kbytes Saved report displayed.



Note:

Data is only displayed if WAN Optimization is enabled. Otherwise, the section will report “No Data to Display”.

The fourth section of the dashboard presents site bandwidth reports:



Total WAN Ingress (Upload): Report displaying total WAN Ingress bandwidth usage for Conduit, Intranet, and Internet services for the last 24 hours. Users may hover over a point on the chart to display detailed data. Select the legend headers to turn data points off or on. The report title links to **Reports > Usage**, with WAN Ingress pre-selected in the direction dropdown.

Total WAN Egress (Download): Report displaying total WAN Egress bandwidth usage for Conduit, Intranet, and Internet services for the last 24 hours. Users may hover over a point on the chart to display detailed data. Select the legend headers to turn data points off or on. The report title links to **Reports > Usage**, with WAN Egress pre-selected in the direction dropdown.

The fifth and final section displays Conduit status information:

Conduit Status	
Live Conduit(s)	Dead Conduit(s)
Conduit PPark-ahodge: Uptime: 22 hours, 14 minutes, 30.0 seconds.	Conduit PPark-AZ-TNET-East is currently dead.
Conduit PPark-AWS-VA-USA: Uptime: 1 days, 16 hours, 22 minutes, 30.0 seconds.	Conduit PPark-jhill is currently dead.
Conduit PPark-Azure-EastUS: Uptime: 1 days, 16 hours, 22 minutes, 30.0 seconds.	Conduit PPark-josswald is currently dead.
Conduit PPark-Azure-WestUS: Uptime: 1 days, 16 hours, 22 minutes, 28.0 seconds.	Conduit PPark-MP is currently dead.
Conduit PPark-Bangalore: Uptime: 1 days, 16 hours, 22 minutes, 26.0 seconds.	
Conduit PPark-cparsons: Uptime: 1 days, 16 hours, 22 minutes, 23.0 seconds.	

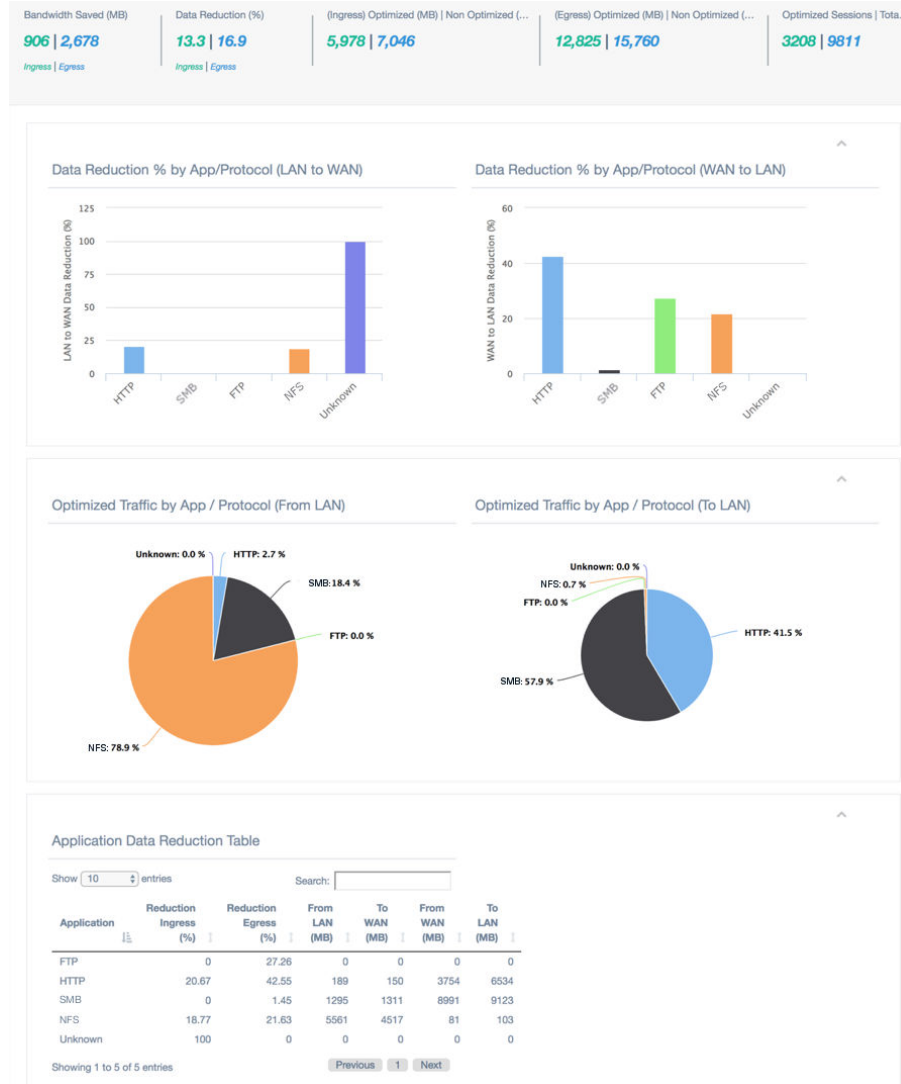
Live Conduits: A list of all live Conduits at the site. On the NCN, live conduits link to the remote site.

Dead Conduits: A list of all dead Conduits at the site. Dead Conduits link to **Statistics > WAN > Conduits**, filtered for dead conduits.

WAN Optimization Dashboard and Reporting Enhancements

7.2 provides a new at-a-glance dashboard for WAN Optimization with more detailed reports and more data about the protocols being optimized. The dashboard refreshes automatically every minute to provide up-to-date information. Additionally, all WANOp pages have been consolidated into the new WAN Optimization menu in the sidebar for ease of location. The WAN Optimization menus and dashboard will only be displayed on appliances that support WAN Optimization.

To view the WAN Optimization dashboard, navigate to **WAN Optimization > Dashboard**:



Bandwidth Saved (MB) 4 4,925 <small>Ingress Egress</small>	Data Reduction (%) 0 14 <small>Ingress Egress</small>	(Ingress) Optimized (MB) Non Optimized (MB) 9,516 10,234	(Egress) Optimized (MB) Non Optimized (MB) 29,072 35,268	Optimized Sessions Total Flows 2651 9232
--	---	---	---	---

Bandwidth Saved (MB): Total bandwidth saved using WAN Optimization in MB, for WAN Ingress and WAN Egress. Links to **WAN Optimization > Monitor WANOp** with the Kbytes Saved report pre-selected.

Data Reduction (%): Data reduction percentage using WAN Optimization, for WAN Ingress and WAN Egress. Links to **WAN Optimization > Monitor WANOp** with the Data Reduction % report pre-selected.

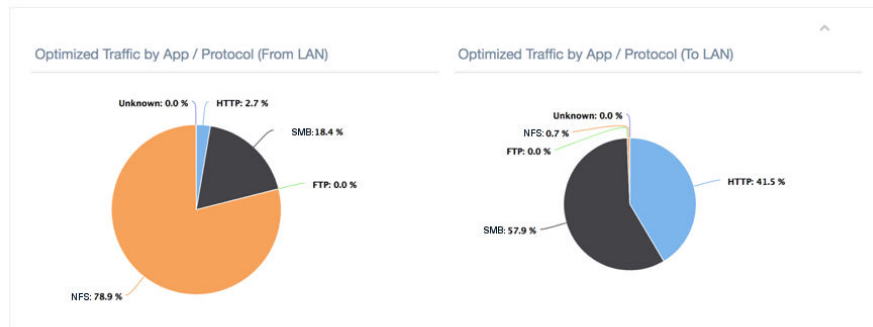
(Ingress) Optimized (MB) | Non Optimized (MB): Total WAN Ingress bandwidth optimized vs non-optimized, in MB. Links to **WAN Optimization > Statistics**.

(Egress) Optimized (MB) | Non Optimized (MB): Total WAN Egress bandwidth optimized vs non-optimized, in MB. Links to **WAN Optimization > Statistics**.

Optimized Sessions | Total Flows: Total number of active WAN Optimized sessions vs total active sessions. Links to **WAN Optimization > Flows**, with WAN Ingress, WAN Egress, and TCP Termination Table pre-selected.

The second section of the WAN Optimization dashboard provides data reduction reports by application/protocol:

Bandwidth Saved (MB) 906 2,678 <small>Ingress Egress</small>	Data Reduction (%) 13.3 16.9 <small>Ingress Egress</small>	(Ingress) Optimized (MB) Non Optimized (...) 5,978 7,046	(Egress) Optimized (MB) Non Optimized (...) 12,825 15,760	Optimized Sessions Total... 3208 9811
---	---	--	---	---



Application Data Reduction Table

Show entries Search:

Application	Reduction Ingress (%)	Reduction Egress (%)	From LAN (MB)	To WAN (MB)	From WAN (MB)	To LAN (MB)
FTP	0	27.26	0	0	0	0
HTTP	20.67	42.55	189	150	3754	6534
SMB	0	1.45	1295	1311	8991	9123
NFS	18.77	21.63	5561	4517	81	103
Unknown	100	0	0	0	0	0


Showing 1 to 5 of 5 entries Previous 1 Next

Data Reduction % by App/Protocol (LAN to WAN): Data reduction percentage for each protocol or application (for applications defined in the configuration), for WAN Ingress (upload). Users may hover over each bar to see the exact percentage.

Data Reduction % by App/Protocol (WAN to LAN): Data reduction percentage for each protocol or application (for applications defined in the configuration), for WAN Egress (download). Users may hover over each bar to see the exact percentage.

The third section of the WAN Optimization dashboard provides information about optimized traffic by application/protocol:

Bandwidth Saved (MB) 906 2,678 <small>Ingress Egress</small>	Data Reduction (%) 13.3 16.9 <small>Ingress Egress</small>	(Ingress) Optimized (MB) Non Optimized (...) 5,978 7,046	(Egress) Optimized (MB) Non Optimized (...) 12,825 15,760	Optimized Sessions Total... 3208 9811
---	---	--	---	---



Application Data Reduction Table

Show entries Search:

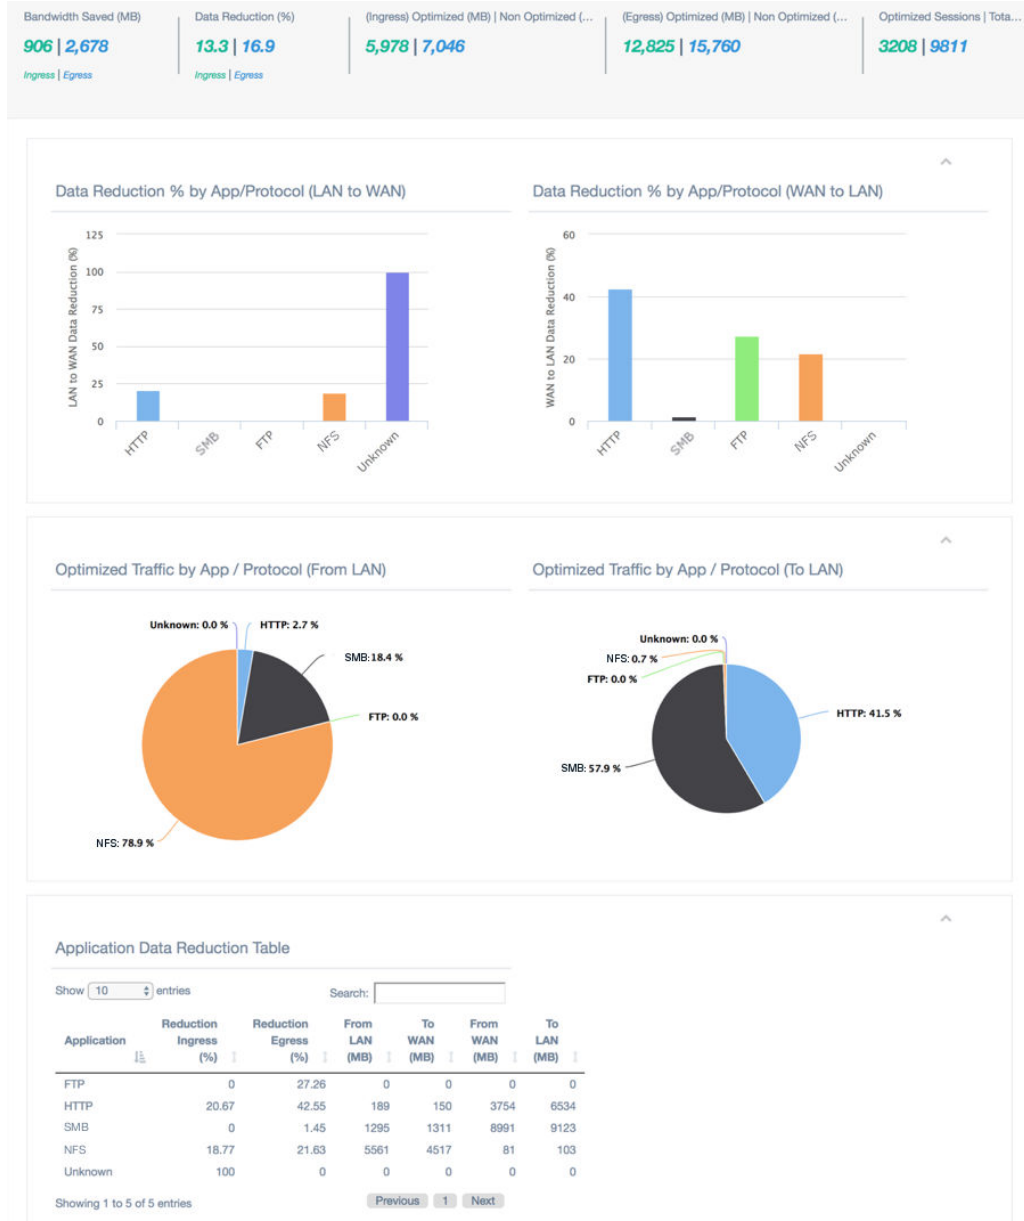
Application	Reduction Ingress (%)	Reduction Egress (%)	From LAN (MB)	To WAN (MB)	From WAN (MB)	To LAN (MB)
FTP	0	27.26	0	0	0	0
HTTP	20.67	42.55	189	150	3754	6534
SMB	0	1.45	1295	1311	8991	9123
NFS	18.77	21.63	5561	4517	81	103
Unknown	100	0	0	0	0	0

Showing 1 to 5 of 5 entries [Previous](#) [1](#) [Next](#)

Optimized Traffic by App/Protocol (From LAN): Chart displaying the percentage of total optimized upload/WAN Ingress traffic for each protocol or application (for applications defined in the configuration). Users may hover over each bar to see the exact percentage.

Optimized Traffic by App/Protocol (To LAN): Chart displaying the percentage of total optimized download/WAN Egress traffic for each protocol or application (for applications defined in the configuration). Users may hover over each bar to see the exact percentage.

The final section of the WAN Optimization dashboard is the **Application Data Reduction Table**:




The Application Data Reduction table shows flow statistics for WAN Optimized sessions that match user-defined Application objects in the configuration, as well as WAN Optimized HTTP, HTTPS, FTP, SSH, and Telnet sessions that do not match a defined Application object.

For more information about WAN Optimization, please see the WAN Optimization Guide.

Enhanced DHCP Relay

7.2 introduces the ability for users to configure up to four DHCP server relay addresses per virtual interface, allowing users with multiple DHCP servers at their NCN site to take advantage of increased redundancy.

DHCP relays may be configured in the Advanced view of the Configuration Editor, under **Sites > [site name] > DHCP > Relays**, as shown below:




Virtual Interface	Server IP 1	Server IP 2	Server IP 3	Server IP 4	Delete
E4-GUEST	192.168.49.13	192.168.49.14			

When configuring DHCP Relays, the Virtual Interface and Server IP 1 are required. Server IPs 2 through 4 are optional.

Monitoring information for DHCP Relay is available under **Monitor > DHCP**.

Client Private Subnet Reuse for Untrusted Segment


7.2 introduces the ability to set duplicate Virtual IPs at multiple sites when the Virtual IP Address is Private and the associated Interface Group is defined as Untrusted. This feature is intended for use in situations where multiple sites are being deployed with the same WAN link provider, with provider equipment pre-configured for the same IP address/subnet at every site.



Virtual Interfaces	Ethernet Interfaces	Bypass Mode	Security	Delete
Pt12 (0)	1 2 3 4 5	Fail-to-Wire	Trusted	
Pt5 (0)	1 2 3 4 5	Fail-to-Block	Untrusted	

IP Address / Prefix	Virtual Interface	Firewall Zone	Identity	Private	Security	Delete
192.168.0.1/24	Pt12	Default_LAN_Zone	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Trusted	
10.0.1.1/24	Pt5	Untrusted_Internet_Zone	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Untrusted	

If one or more of the duplicate Virtual IPs is not private, an Audit Error will be displayed.




	EC319	10.0.1.1/24	At Site CL2 in 'virtual_ip_addrn->ip_addrn': '10.0.1.1' is not unique throughout 'Default_RoutingDomain'. The duplicate IP Address was found at 'site CL1: virtual_ip_addrn->ip_addrn'
--	-------	-------------	--

Palo Alto GlobalProtect Cloud Integration

7.2 adds support for integration of branch office Oracle Talari Appliances with the Palo Alto GlobalProtect cloud service via IPsec tunneling, enabling users to tunnel Internet-directed traffic to GPCS for cloud-hosted filtering and security services.

To configure a Palo Alto GlobalProtect cloud IPsec tunnel, navigate to **Configuration > Configuration Editor** on the NCN and **Import** the current configuration file. Click on the **Advanced** tab, expand **Connections > [Site Name] > IPsec Tunnels**, and click the **(+)** icon.

Select Palo Alto as the **Service Type**, select the **Local IP** address from the dropdown, fill in the **Peer IP** address of the GlobalProtect cloud service IKE Gateway, enter the IKE **Pre-Shared Key**, add the local Protected Networks for the IPsec tunnel, and click **Apply**.



If no options are available in the Local IP dropdown, ensure Internet Service is enabled on at least one WAN link at the site under **Connections > [Site Name] > Internet Services**.

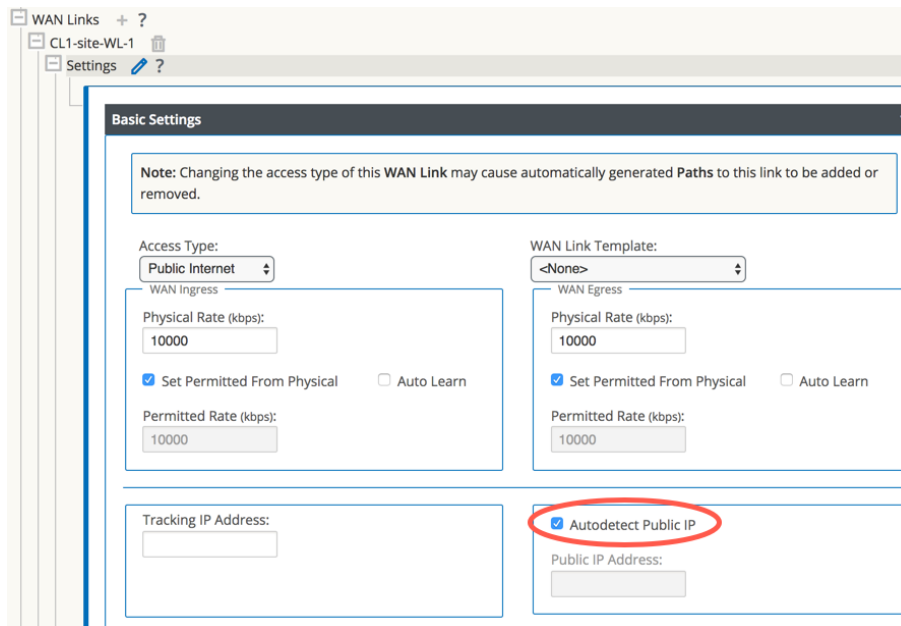
Private Cloud Path Enhancement

In certain cases, service providers have a private cloud which is separate from the public Internet. Within their environment they use PAT (Port Address Translation) to forward user traffic from their private cloud to the Internet. In these cases, the service provider will have a limited number of public IP address for NATing. When deployed for an enterprise customer, if they select one of these providers for multiple client sites, there is the possibility that multiple Client WAN links could be PATed/NATed to the same public IP address. Prior to 7.2, Talari would validate/learn a path based on the source IP address of the received frame (at the NCN for example). The end result is that the first site brought online would function as expected, with a Path in the GOOD state. However, at the second Client site using the same public IP address, the Path would be in the DEAD state. To resolve this issue, 7.2 has been enhanced to use the source IP address and source port for path learning validation. With this enhancement Talari has expanded its ability to interoperate with multiple additional Service Provider WAN environments.

Note:

Conduits between Client sites with the same shared public IP are not supported at this time.

All WAN links which may reside behind the same public IP must have Autodetect Public IP enabled in the configuration under **Sites > [Site Name] > WAN Links > [WAN Link] > Settings > Basic Settings**, as shown below:



Remote sites other than the NCN will not be able to bring up paths to a client using a shared public IP unless UDP Hole Punching is enabled in the configuration under **Connections > [Site Name] > Conduits > [Conduit] > Local Site > WAN Links** at the client sites which share the public IP, as shown below:

The screenshot shows the 'WAN Links' configuration table. The 'UDP Hole Punching' checkbox for 'CL1-site-WL-2' is highlighted with a red circle.

WAN Link	Use	Tunnel Header Size (bytes)	Active MTU Detect	UDP Port	UDP Port Switching			Autopath Group
					UDP Hole Punching	Enable	Alt Port	
CL1-site-WL-1	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	2156	<input type="checkbox"/>	<input type="checkbox"/>	1440	<Default>
CL1-site-WL-2	<input checked="" type="checkbox"/>	0	<input checked="" type="checkbox"/>	2156	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1440	<Default>

Additional Features in 7.2 GA

7.2 introduces the following additional features:

- Configuration Editor:

A note has been added in the Configuration Editor at all locations where a Rule may be configured to clarify that Drop Limit and Disable Limit values in milliseconds are not valid for Bulk Classes. These values will automatically be set to 0. Drop Depth (bytes) and Disable Depth (bytes) values should be used for Bulk Classes instead.

19

Release 7.2 P3 Features


This chapter includes features and enhancements released in 7.2 P3.

Configuration Versioning and Comparison

Configuration Versioning

Beginning in 7.2 GA P3, configurations will be versioned automatically whenever a change is applied via Change Management. Each new version contains additional metadata which allows network administrators to quickly audit changes to the configuration. The information collected includes: the user who edited the configuration and when the edit was made, the user who activated the configuration and when the configuration was activated, the name of the configuration, and a user-generated comment describing the reason for the change. All of the information except the user-generated comment will be collected automatically without any required action from the user or network administrator.

Whenever a configuration is saved or exported, the user will be prompted to add a comment, as shown below:



 **Note:**

Adding or updating the user comment is not enforced by the configuration editor. Users may save or export a configuration with no added comment, or without updating the comment text.

 **Note:**

Adding a user comment when exporting a saved configuration will overwrite any user comment added when the configuration was previously saved.

Configuration Comparison

To review and compare configurations, navigate to **Configuration > Compare Archived Configurations** and click **Select Configurations to Compare**.

 **Note:**

Only configurations which have already been activated via Change Management are available for comparison.

Click to select or deselect a configuration for comparison. Selected configurations will be highlighted. Click Compare to view the selected configurations.

Activated Time	Activated By	Edited Time	Edited By	Configuration	Comment
2018-05-21 15:05:14	talariuser_192-168-50-79	2018-05-21 14:52:57	talariuser_192-168-50-79	example_configuration3_20180521_1505.cfg	*Increased AVL INET bandwidth.*
2018-05-21 14:50:29	talariuser_192-168-50-79	2018-05-21 13:46:58	talariuser_192-168-50-79	example_configuration2_20180521_1450.cfg	*Added site CLT*
2018-05-21 13:10:49	talariuser_172-16-42-1	2018-05-21 12:04:21	talariuser_172-16-42-1	example_configuration1_20180521_1310.cfg	*Removed Facebook application, made AVL Geo NCN*
2018-05-04 15:00:57	talariuser_172-16-42-1	2018-05-04 14:42:48	talariuser_172-16-42-1	testEdit2_20180504_1500.cfg	*add facebook*
2018-04-05 12:43:05	talariuser	2018-04-05 12:38:45	talariuser	AVL-intra-04052018_20180405_1243.cfg	*add intranet service at AVL*

 **Note:**

More than two configurations may be selected, but only two configurations may be compared at a time. The Compare button will be greyed out until only two configurations are selected.

The newer configuration will be displayed on the left, and the older configuration will be displayed on the right. In addition to the configuration name and user-generated comment, the header for each configuration will show the activating/editing username, IP address, date, and time:

Compare Archived Configurations

Select Configurations to Compare

Newer	Older
Config Name: <code>example_configuration4_20180521_1530.cfg</code>	Config Name: <code>example_configuration3_20180521_1505.cfg</code>
Activated By: <code>snorris_192-168-50-79 @ 2018-05-21 15:30:56</code>	Activated By: <code>talaruser_192-168-50-79 @ 2018-05-21 15:05:14</code>
Edited By: <code>snorris_192-168-50-79 @ 2018-05-21 15:24:40</code>	Edited By: <code>talaruser_192-168-50-79 @ 2018-05-21 14:52:57</code>
Comment: <code>"Removed site CLT."</code>	Comment: <code>"Increased AVL INET bandwidth."</code>

Changes in newer config: `example_configuration4_20180521_1530.cfg` vs older config: `example_configuration3_20180521_1505.cfg`
 Legend: **blue** => modified, **red** => removed, **green** => added

```

set apn_properties
  encryption_mode=aes128
define application_match_collection
{
}
define application
  name=DEFAULT_AF11
{
  set application_properties
}
define application
  encryption_mode=aes128
define application_match_collection
{
}
define application
  name=DEFAULT_AF11
{
  set application_properties
}
define application
  
```

A side-by-side comparison of the configuration text file is displayed below the headers. Objects that have been removed will be highlighted in red and marked with a “-”:

Changes in newer config: `example_configuration4_20180521_1530.cfg` vs older config: `example_configuration3_20180521_1505.cfg`
 Legend: **blue** => modified, **red** => removed, **green** => added

```

wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
autopath_group_name=Default
add conduit_usage
-
-
-
-
-
-
remote_site_name=G50
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
+
+
+
+
+
+
remote_site_name=CLT
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
autopath_group_name=Default
add conduit_usage
remote_site_name=G50
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
  
```

Objects that have been added will be highlighted in green and marked with a “+”:

Changes in newer config: `example_configuration2_20180521_1450.cfg` vs older config: `example_configuration1_20180521_1310.cfg`
 Legend: **blue** => modified, **red** => removed, **green** => added

```

service_group_name=Default
autopath_group_name=Default
add conduit_usage
+
+
+
+
+
+
remote_site_name=CLT
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
autopath_group_name=Default
add conduit_usage
remote_site_name=G50
wan_egress_rate_fair_share=1000
wan_ingress_rate_fair_share=1000
service_group_name=Default
  
```

Objects that have been modified will be highlighted in blue and marked with a “<”:

Changes in newer config: `example_configuration3_20180521_1505.cfg` vs older config: `example_configuration2_20180521_1450.cfg`
 Legend: **blue** => modified, **red** => removed, **green** => added

```

1 set conduit_properties
}
add virtual_wan_link
  name=AVL-INET
{
  set properties
  access_type=public_internet
  wan_ingress_physical_rate_kbps=10000 <<
  wan_egress_physical_rate_kbps=50000 <<
  wan_ingress_permitted_rate_kbps=10000 <<
  wan_egress_permitted_rate_kbps=50000 <<
add access_interface
  name=AVL-INET-AI-1
  virtual_interface_name=AVL-INET-v11
1 set conduit_properties
}
add virtual_wan_link
  name=AVL-INET
{
  set properties
  access_type=public_internet
  wan_ingress_physical_rate_kbps=3000
  wan_egress_physical_rate_kbps=15000
  wan_ingress_permitted_rate_kbps=3000
  wan_egress_permitted_rate_kbps=15000
add access_interface
  name=AVL-INET-AI-1
  virtual_interface_name=AVL-INET-v11
  
```

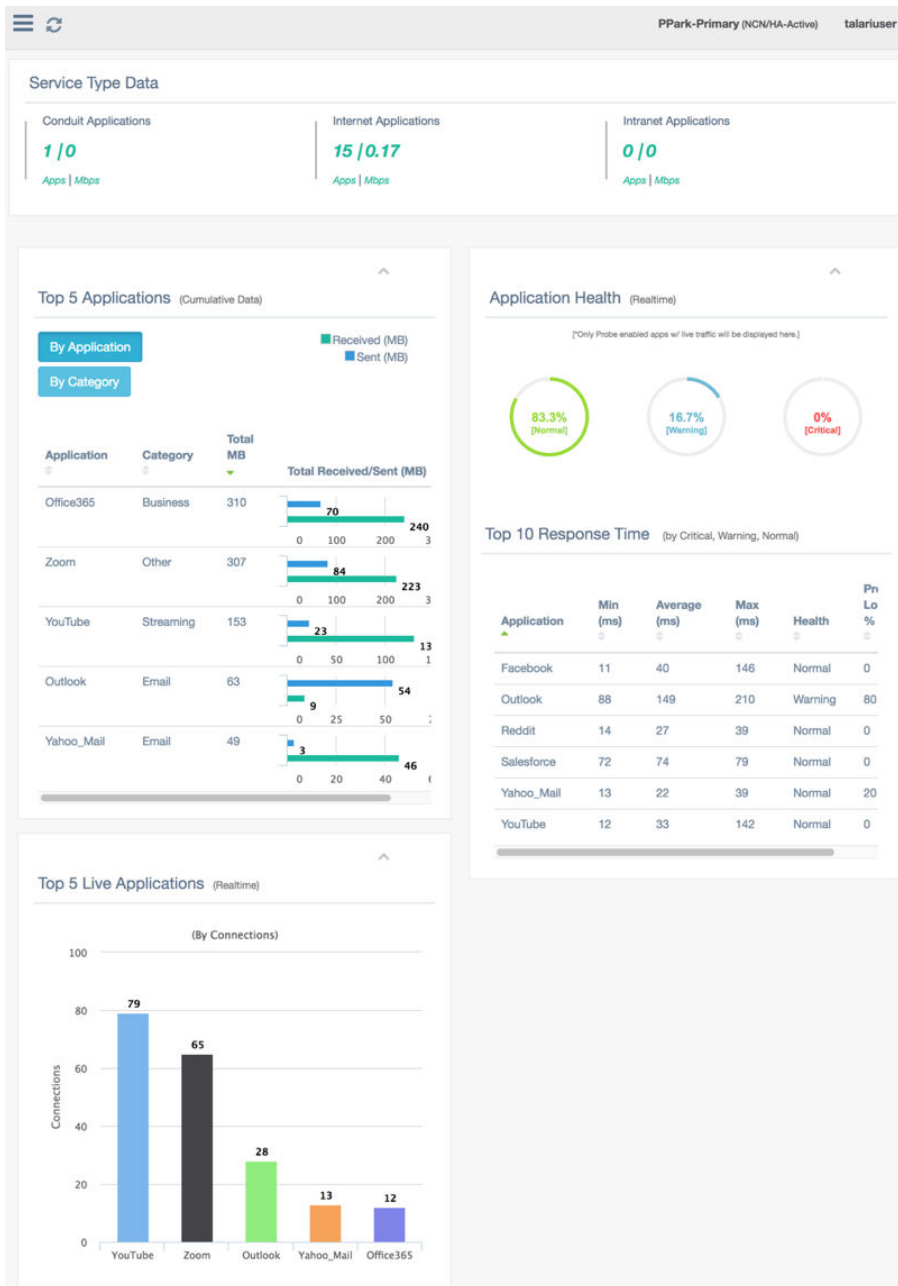
Release 7.3 Features

This chapter includes features and enhancements released in 7.3.

Enhanced Application Identification

7.3 GA introduces Enhanced Application Identification, which offers a significant improvement to how Oracle Talari Appliances identify and forward applications. This release introduces the following new application identification enhancements:

- DNS snooping, a less intrusive application identification technique when compared to our existing DNS proxy or manual six-tuple identification mechanisms.
- Simplified application policy configuration, with a default signature library (the Application Signature Library) with over 100 application entries included. Preset application signatures are modular and can be downloaded and upgraded independently of software packages via the regular Change Management process. Talari will provide updates to the Application Signature Library moving forward based on customer feedback.
- A dedicated application dashboard which allows the user to view top cumulative and live applications, bandwidth usage by service, and application health information. This information helps administrators perform common tasks such as troubleshooting and capacity planning:



- Streamlined configuration elements that make creating an application policy fast and easy. The Enhanced Application Identification is extensible and supports the addition of user-defined categories and applications.
- Applications are assigned to a pre-defined application category, or users may configure additional application categories as required.

By combining all of these capabilities, users can create granular application policies such as steering a single application (e.g., Microsoft Office 365) out the local internet service while forwarding all other SaaS application(s) back to the data center or NCN site. The user can also define the scope of the application policy which could include a single location, all Edge sites or a subset of sites depending on user needs. Traditional QOS services are applied for conduit services where the user can map an application to a pre-defined classification or select their own classification from a pre-defined list.

For information on configuring and monitoring Enhanced Application Identification, please see the *Enhanced Application Identification & Application Signatures Guide*.

E500 Appliance (7.3 GA P3)

7.3 GA P3 adds support for the E500 appliance. The E500 is an extension of the E-series of Oracle Talari Appliances. The E500 is intended for use in mid-sized branch or regional offices that require higher performance and port density than the E100 provides. The E500 supports WAN Optimization and Easy 1st Install. For more information on this platform, please see the *E500 Installation Guide* and the *E500 Hardware Guide*.

Private Registration Server (7.3 GA P3)

Beginning in 7.3 GA P3, customers who do not wish to depend on the public Registration Server may host a Private Registration Server for use during the Easy 1st Install process. The private registration server may be deployed for access via an incumbent private intranet or for access via the public Internet, and may use either a static IP host or a DNS-resolvable Fully Qualified Domain Name (FQDN).

Once the Private Registration Server (PRS) is installed and operational, the high-level data flow for the Easy 1st Install process to complete properly is as follows:

- User provides the serial number for the site being deployed to the NCN
- The NCN uploads package to the PRS
 - Connectivity must exist to the PRS from the NCN management port IP
 - The NCN pushes the client package to the PRS via HTTPS
- Once the client appliance is powered on and has an IP address/gateway/DNS for the management port, the following occurs:
 - The client will attempt to establish an HTTPS session via its management port to the PRS and provide its serial number
 - * Once the serial number is validated via the HTTPS session, the PRS will provide a URL for the client to download the appliance package
 - The client appliance will establish a second HTTPS session to retrieve the appliance package based on the validated serial number

For detailed information on deploying and using a Private Registration Server, please see the *Private Registration Server Installation and Deployment Guide*.

Threshold Alerting (7.3 GA P4)

7.3 GA P4 introduces the ability to monitor WAN link usage and trigger an alert if a user-defined usage threshold is exceeded. Threshold Alerting can provide insight into situations wherein the failure of one WAN link in a Conduit would result in the remaining WAN link(s) being oversubscribed, allowing customers to resolve potential issues before they arise.

Threshold Alerting is configured using the Advanced view of the Configuration Editor, and is disabled by default. To enable Threshold Alerting at a site, go to **Sites > [Site] > Basic Settings** and enter a non-zero value for at least one threshold:

Basic Settings
✎ ?

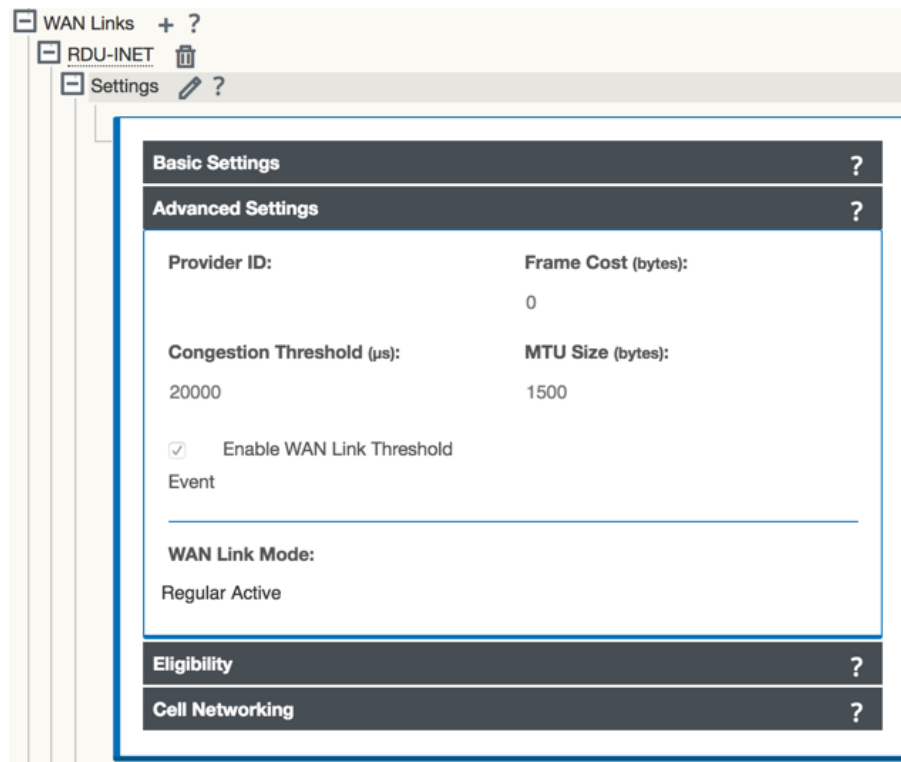
Appliance Name: E100	Secure Key: 132bb462cbc68809 Regenerate
Model: E100	Mode: primary NCN
Site Template: <None>	
<hr/>	
Default Direct Route Cost: 5	
<hr/>	
Gateway ARP Timer (ms): 1000	
<hr/>	
<input type="checkbox"/> Enable Source MAC Learning	
<hr/>	
Application Normal RTT adjust time (ms): 0	Application Warning RTT adjust time (ms): 0
<hr/>	
WAN Ingress Lower Threshold (kbps): 8000	WAN Egress Lower Threshold (kbps): 18000
WAN Ingress Higher Threshold (kbps): 11000	WAN Egress Higher Threshold (kbps): 20000

When a threshold is configured and the combined WAN link usage for the site exceeds the configured value, an event will be generated.

Note:

When the combined WAN link usage exceeds the Lower Threshold value, an event with a severity of “Notice” will be generated. When the combined WAN link usage exceeds the Higher Threshold value, an event with a severity of “Warning” will be generated.

By default, all WAN links at a site are used for threshold calculations. To exclude a WAN link from threshold calculations, go to **Sites > [Site] > WAN Links > [WAN Link] > Settings > Advanced Settings** in the Advanced view of the Configuration Editor and uncheck the “Enable WAN Link Threshold” box:



Additional Features in 7.3

- **CT800-128 and VT800-128 Appliances (7.3 GA P4)**
7.3 GA P4 introduces support for the CT800-128 and the VT800-128. These new virtual appliances build on the CT800 and VT800 appliances to support up to 128 conduits in AWS, ESXi, Azure, and HyperV.
- **Increased Throughput in AWS (7.3 GA P4)**
The CT800-128 supports a new maximum performance level of 500Mbps full-duplex for AWS.

21

Release 8.0 Features

This chapter includes features and enhancements released in 8.0.

ID	Description	Release
16569	Old log files are now compressed to permit longer retention.	8.0 GA
16551	8.0 GA introduces Cloud Connect, which allows customers to connect to participating Cloud Connect providers from the enterprise platforms via Cloud Conduits.	8.0 GA
19119	The licensing requirement for virtual appliances (VT800, VT800 - 128, CT800, and CT800 - 128) has been removed.	8.0 GA P1

22

Release 8.1 Features

This chapter includes features and enhancements released in 8.1.

ID	Description	Release
29989632 (19500)	User Names can now contain several special characters that were previously disallowed: @, /, and \.	APN 8.1 P1
29989624 (19486)	T5200 CPU profile optimizations as well as general packet scheduler enhancements have been made to significantly improve performance and stability during heavy load across large networks. Advanced logging and alerting tools are included.	APN 8.1 P1
29989439 (19230)	When configuring email alerts with SMTP Authentication enabled, event emails may not be sent.	APN 8.1 P1
29989448 (19245)	8.1 introduces support for the D6000. The D6000 is the replacement appliance for the T5X00 series, with improved performance.	APN 8.1 GA
29989220 (18936)	8.1 introduces support for the D2000. The D2000 is the replacement appliance for the E1000.	APN 8.1 GA

23

Release 8.2 and later

Feature descriptions for release 8.2 and later are available in the documentation set that corresponds to each release. See the Release Notes for the release to see a description of the feature.

Release 9.0 Features

This chapter includes features released in 9.0.

Selective Software Update

You can deploy two major releases of Oracle SD-WAN Edge within a single network using Selective Software Update. This avoids a global outage when updating the entire SD-WAN; some appliances update, while other devices continue to operate. Selective Software Update lets you update the Edge devices in your network on a rolling basis, by selecting sites and identifying update time. Each software release has a compatibility matrix that details backwards compatibility for prior releases of Oracle SD-WAN Edge. This compatibility matrix can be found in the Release Notes.

You may upgrade and downgrade between two sequential major releases, and can also schedule upgrades/downgrades on a per-site or a site-group level. When upgrading an HA pair by Selective Software upgrades, both peers will be upgraded to the same target version.

When two versions of the Edge software are deployed within your APN, the Configuration Editor will only allow you configure settings that are common to both versions. Later, when all sites are upgraded to the same Edge version, all features may be configured.

Create and Update Plan

Update Plan allows you to create and upload Oracle software packages to the SD-WAN Controller and create/modify a software update schedule.

Upload Software Packages

Click on the **Choose Files** button to upload a new software package. Once uploaded, you can also view the configuration.

Software Update Settings

You can set the amount of time the system will ignore incomplete wait times after no transfer update. Choose from 10 seconds, 30 seconds, 1 minute, or 2 minutes. You can also disable this option completely.

You can also configure the installation countdown time. Choose from 10 seconds, 30 seconds, 1 minute, or 2 minutes.

Software Update Schedule

The Software Update Schedule table contains the following fields:

- Site/Region: Site name or Region name
- Model: Device model
- Current Active Software: Currently running software
- Target Update Time: Planned date and time for the software update

To cancel any changes and continue with the previously created plan, click **Cancel**. This action will return to the Status page where any previously created plan can be restarted.

To clear out all uploaded software and scheduled updates from the plan, click **Clear Plan**.

To verify the updated plan, click **Verify Plan**. Plan verification can be done throughout the scheduling process, but must be done before activating the plan. If there are any errors in the verification process, you will be prompted with corrective steps to take.

To activate the verified plan, click **Apply Plan**. This action will activate the updated plan and return to the Status page to monitor the plan's progress. Once the software is updated on a site, the Oracle service for that site will be restarted.

 **Note:**

Selective software update uses the time and date set on the active SD-WAN Controller appliance for updates. To ensure that updates are properly scheduled, the time, timezone, and date must be set properly on the active SD-WAN Controller, as well as all GEO and HA SD-WAN Controller appliances using the 'Date and Time Settings' screen in the 'Manage SD-WAN Edge' menu.

Monitor Plan Progress

The Monitor Plan Progress page allows you to monitor the status of the selective software update plan.

Options

- **Update Plan:** Allows you to upload Oracle software (.tar.gz) for the SD-WAN Controller and all SD-WAN Edge models in the network. You can also create or modify the schedule for a selective software update plan. Selecting this will stop any running plan.
- **Stop Plan:** Stops the currently running plan.
- **Resume Plan:** Resumes the stopped plan.

Selective Software Status Table

Site-Appliance: 'Site name'-'Appliance name'

Model: device model

State selective update plan status:

- **Stopped:** software update is stopped
- **Pending:** software update is pending
- **In Progress:** software update is currently in progress
- **Finished:** software update has successfully completed
- **Cancelled:** software update was cancelled
- **Failed:** software update has failed

Current Active Software: currently running software

Current Active Config: date of running configuration

Target Software: software targeted for the update

Target Update Time: planned date and time for the software update

Microsoft Azure Virtual WAN

Microsoft Azure Virtual WAN is a cloud networking service that allows you to create a connection between site networks, internet, and Azure Virtual Networks (VNets). It provides optimized and automated branch connectivity through and to Azure. Azure Virtual WAN gateways are hosted and managed by Azure, and the connections are managed by the user.

Microsoft has created a partner program that allows for Oracle SD-WAN to be integrated. This connection is created using IPsec tunnels.

For more information, see the [Azure website](#).

Prerequisites

In order to integrate Oracle SD-WAN with Microsoft Azure, you must have the following:

- An Azure account, which includes an associated Subscription ID and Tenant ID, as well as access to the Azure portal.
- An Oracle SD-WAN Edge system.
- Direct internet access through the tn-mgt0 interface.

Configuring Azure Virtual WAN

There are two parts to the configuration process:

1. In the Azure Portal, create a new Virtual WAN, hub, and connect the Virtual Network to the hub. You must also obtain an application ID and a secret key.
2. In Oracle SD-WAN Edge, use your Azure credentials, application ID, and secret key to create a new Azure service, and connect it back to Azure using IPsec tunnels.

Create a Virtual WAN

Follow these steps to create a new Virtual WAN from the Azure portal.

1. Navigate to the [Azure portal](#) and sign in with your Azure account.
2. Click on **+Create a resource**.
3. Type **Virtual WAN** into the search box and then click on **Enter**.
4. Select **Virtual WAN**. On the **Virtual WAN** page, click on **Create** to bring up the **Create WAN** page.
5. Click on the **Basics** tab and fill in the following fields:
 - **Subscription:** Select the subscription you want to use.
 - **Resource Group:** Create a new resource group or choose an existing one.

 **Note:**

It is recommended that you create a new resource group.

- **Resource Group Location:** Select a resource location of your WAN.
 - **Name:** The name of your WAN.
 - **Type:** Choose Basic or Standard. If you create a basic WAN, you can only create a basic hub.
6. Click on **Review +Create**.
 7. Once validation passes, click on **Create** to create your virtual WAN.

Create a Hub

Follow these steps to create a hub in the Azure portal:

1. On the page for the Virtual WAN you created, click on **Hubs** in the **Connectivity** section.
2. Click on **+ New Hub**.
3. Enter the following on the **Basics** tab:
 - Region
 - Name
 - Hub private address space. The minimum address space is /24.
4. Click on the **Next: Site to site** button.
5. On the Site to site tab, fill out the following fields:
 - **Creating a Site to Site VPN:** Select **Yes**.
 - **Gateway scale units:** Select the number of units from the dropdown.
6. Click on **Review + Create**.
7. Once validated, click on **Create**. It will take 30 minutes to process.

Connect a Virtual Network to the Hub

Follow these steps to create a connection between a Virtual Network and your hub in the Azure portal.

1. On your Virtual WAN page, click on **Virtual Network Connections**.
2. Click on **+ Add connection**.
3. Fill in the following fields:
 - **Connection name:** The name of your connection.
 - **Hubs:** Select the hub to connect to.
 - **Subscription:** Verify the subscription.
 - **Virtual network:** Select the virtual network to connect to the hub. This cannot have a pre-existing virtual network gateway.
4. Click on **OK** to create.

Create an Application ID and Secret Key

Follow these steps to register your app in the Azure portal and create an application ID and secret key.

1. Log in to the Azure portal.
2. Navigate to **Azure Active Directory, App registrations, New registration**.
3. Enter a name for the app.
4. Select the **supported account types**.
5. Choose **Web** as the app type under **Redirect URI**.
6. Click on **Register**.
7. Make note of the application ID (also referred to as the client ID).
8. To create a secret key, click on the **Certificates & secrets** page.
9. Click on **New client secret**.
10. Enter a description and an expiration date.
11. Click on **Add**.
12. Make note of the secret key, as you cannot retrieve it later.

Perform the following steps to assign the appropriate roles for authentication purpose:

1. In the Azure portal, navigate to the **Resource Group** where the Virtual WAN was created.
2. Navigate to **Access control (IAM)**.
3. Click **+ Add** and select **Add role assignment**.
Provide values for the following fields:
 - Role: Select Owner from the drop-down list. This role allows management of everything including access to resources.
 - Assign access to select Azure AD user, group, or service principal.
 - Select Provide the name of the registered application created earlier and select the corresponding entry when it appears.
4. Click **Save**.

Create an Azure Configuration

Once you have a subscription ID, tenant ID, application ID, and secret key, you can enter these into the Oracle SD-WAN Edge Configuration Editor to begin the automation process. Follow these steps:

1. From your Edge system, navigate to **Configuration, Configuration Editor**.
2. Click on the **Global** option.
3. On the menu on the lefthand side, click on **Services**.
4. Click on the arrow next to **Microsoft Azure Virtual WAN Services**.
5. Enter your subscription ID, tenant ID, application ID, and secret key.
6. Click on **Sync with Azure** to validate these credentials.

7. Once validated, choose a available WAN hub to connect to from the dropdown.
8. Click on **Save**.

Add Service to the Site

Once you've configured the Microsoft Azure service, you can add the service to the site. In order to connect to Azure Virtual WAN, there must be 2 active/active IPsec tunnels.

If you use one WAN link, the local WAN link serves as the end point of both tunnels connected to the branch site.

1. From your site, navigate to the **Services** page.
2. Click on the arrow next to **Azure Virtual WAN service**.
3. Click on **Add**.
4. Enter 1 WAN link to use.
5. Enter which hub to use on the list of available hubs.
6. Save once completed.

SD-WAN Edge Configuration

The SD-WAN Edge has a set of tools that can be used to reliably deliver application services over a combination of networks. In order for these tools to work optimally and deliver the expected results, it is important to understand the applications that will be used, to have a well-designed network, and to apply the tools in the right situation.

Configuration Package

The Configuration Package is an archive that contains the SD-WAN Edge Configuration file and meta data about the SD-WAN Edge Configuration. This includes the network maps that are constructed using the Configuration.

The Configuration Package is used to save and export the configuration file. To apply the changes, Export the Configuration Package to the active SD-WAN Controller - Change Management Inbox, compile the new Configuration and then push the changes to the network.

The actions available from the Configuration Management screen are:

- Add - Create a new configuration
- Import - Select a file to import
- Lock/Unlock
- Download - Download the configuration to the users local system
- Publish to SD-WAN Controller
- Modify - Add a comment or change the name of the configuration
- Clone - Create a copy of the configuration
- Delete

New Configuration Package

Allows the user to create a new Configuration Package.

Modify

Allows the user to modify the name of an existing configuration.

Import SD-WAN Edge Configuration

You can use Change Management or Local to import a configuration. Import configuration always imports with the name of the configuration being imported. If another configuration file already exists with the same name, the user will be prompted select override option or to change the name of the configuration. If no Configuration Package is open, you must create a new one using the Add button.

Selecting Import will display a dialog that allows for an SD-WAN Edge Configuration or Configuration Package to be imported for editing from external sources:

- **From Change Management** displays a list of current and previous Configurations on this SD-WAN Edge from the Change Management archive.
- **From File** allows a Configuration to be uploaded from the user's system

Name Conflict

When importing a package or configuration file, if another package with the same name exists, then the imported file will need to be renamed.

- **Package Name:** The new name for the imported file.
- **Allow Override:** Enable replacing the contents of existing package with the contents of imported file. The contents of the existing package will be lost.

Export Configuration Package

Selecting Export will display a dialog that allows the user to export an open Configuration Package to Change Management.

A dialog will show and it allows for the current Configuration Package to be exported to a **Destination:**

- **Download File:** The Configuration Package will be compressed into an archive and a download will be initiated.

Import Application Preset

This will import a new Oracle defined application signature file into the current open config.

User can import from file or import from signature files saved on the SD-WAN Edge. If importing from file, the file will be saved on the SD-WAN Edge and be used for later import.

When importing the new signature file into current config, there are two modes:

- **Merge:** user changes to the Oracle defined applications will stay and new additions from the new signatures will be added to the current config file.
- **Overwrite:** user changes to the Oracle defined applications will be lost and the new application signatures will be added to the current config file.

Import Cloud Config

This will import cloud services configuration into the global Cloud Services section of the current configuration. The import file is usually provided by the cloud services provider in JSON format with the .json file extension.

Click **Clear** at any time to reset the filter and display the entire configuration.

Note:

While any named object may be filtered, the tree will only expand to the object types shown in the Type field. For example, you may filter for the Virtual Interface "VLAN1", but only the Site(s) containing matching Virtual Interfaces will be automatically expanded.

Global Actions

A list of options available for changing the settings over the entire configuration.

- **Regenerate All Secure Keys:** In the event that all of the secure keys in the configuration need to be changed, this tool will randomly generate new secure keys for all sites in the configuration, instead of requiring a manual change of all keys for each site individually.

Audits

You can configure the Configuration Editor performs an audit on the Configuration by clicking on the Audit button. The Audit Panel, located at the bottom of the page, displays any errors or warnings in the Configuration.

All the errors must be resolved for the Configuration file to compile properly in Change Management.

Configuration Modes

All Sites

Sites

The concept of the **Sites** view is to simplify the configuration process to allow the user to create a configuration file, which will generate a Conduit between the defined sites.

The required configuration properties for a Conduit between sites include:

- Basic Settings
- Define Interfaces
- Define WAN Links
- Routing

Basic Settings

When adding a site the user would configure the type of SD-WAN Edge for the site, as well as the option to select from a site template which would typically be used for branch office (clinet) sites. The Site Summary provides the ability to add, view, and edit site details for interfaces, WAN Links, and Static Routes.

Define Interfaces

The Interface option allows the user to define the physical topology of the site, such as the ports, logical VLANs and security level for the physical ports. At this level, the user can also define if the WAN interface will use DHCP for an IP address, or they may statically assign an IP address. This allows the user to configure multiple options under the same panel.

Define WAN Links

While Adding / Editing a WAN Link, the option to use a WAN Link Template is provided. After selecting a WAN Link Template, the WAN Link will be configured using the WAN Link Template values. The user has the option to overwrite the Template values if desired. Additionally once the Virtual Interface is selected, the IP address is automatically provided from the interface configuration.

WAN Link Templates

The WAN Link Templates functionality provides users with a way to setup basic configuration for WAN Links and reuse these across the network to save time.

Below are the steps to use this feature through the Basic configuration mode:

- The WAN Link Template feature has been enhanced to include the ability to configure Service Provider-specific WAN Link Templates for Broadband, MPLS, and Private Intranet connections.
- This allows for a quicker site configuration by applying a WAN Link Template for newly created sites, as well as an easier way to clone branch locations with similar Service Provider attributes.
- To create a WAN Link Template based on Service Provider attributes, navigate to **SD-WAN Edge Configuration > Import** the current configuration to get started.
- Under the Basic tab, select the Network view and click (**Service Provider**).

Routing

Click the icon to the right of Static Routes: to add Static Routes. Click the icon to the right of the Static Routes to edit the Static Routes. Clicking will take the user to the Add / Edit Static Routes dialog. Currently the user can only add local routes within the Basic configuration view.

Site Configuration

The Sites Configuration section allows the user to define each SD-WAN Edge in the network and to configure it.

Configuration options for each site include: Basic Settings for the SD-WAN Edge, Interface Groups that will be utilized by the SD-WAN Edge, Virtual IP Addresses on the Interfaces that enable the SD-WAN Edge to communicate with other networks, WAN Link properties that enable the SD-WAN Edge to connect to other locations, Routes that enable the SD-WAN Edge to direct traffic to other networks, and High Availability (HA)

Basic Settings

Basic Settings allows the Configuration of the SD-WAN Edge Name, Secure Key, Model, Mode of the physical SD-WAN Edge at a specific site.

SD-WAN Edge Name:

Name for this SD-WAN Edge. The two SD-WAN Edge deployed in High Availability configuration can be named different.

Model:

The Model is the specific hardware located at the site and the Mode determines how the site will participate in the Adaptive Private Network. Note, not all models can operate in all modes.

Site Template:

A Site Template can be applied to the site if they exist.

Regions:

Regions are used as metadata for sites. You can associate one region per site. Each region can be associated with more than one site. Once a region is created, it cannot be removed if it is mapped to a site. You can also add a region directly from this screen if you do not see the one you want.

User Tags

User tags are used as metadata for sites. Each site can have one user tag associated with it. User tags can be associated with more than one site. You can also add a user tag from this screen if you do not see the one you want.

Secure Key:

Secure Key is used to encrypt and decrypt data exchanged between sites using conduit service.

Regenerate

Regenerate can be used to auto generate a new Secure Key.

Network Roles:

- **Primary SD-WAN Controller:** The primary SD-WAN Controller for the SD-WAN Edge network. The SD-WAN Controller is responsible for managing SD-WAN Edge configurations and software versions for all clients and serves as a mediator between clients.
- **Secondary SD-WAN Controller:** Typically a Geo-located client in the SD-WAN Edge network that has the ability to take over management functions of the SD-WAN Edge network in the event the Primary SD-WAN Controller becomes unavailable. Note, the Secondary SD-WAN Controller does not provide High Availability or Failover capabilities for an individual site.
- **Client:** A client will receive its SD-WAN Edge Configuration from an SD-WAN Controller and participates in the SD-WAN Edge network as described by the Configuration.

Default Direct Route Cost:

The Default Direct Route Cost (1 to 15) can be set that will be used for all routes added to this SD-WAN Edge.

Gateway ARP Timer (ms):

Gateway ARP Timer (ms) is the time can be set to adjust the time (100 to 20,000 milliseconds), between ARP requests for the configured Gateway IP Addresses.

Enable Source MAC Learning:

Enable Source MAC Learning when checked the SD-WAN Edge will store the Source MAC Address of received packets so that outgoing packets to the same destination can be sent to the same port.

Application Normal RTT adjust time (ms):

Application Normal RTT adjust time (ms), is the time (0 to 500 milliseconds), to adjust network-wide application normal round trip time.

Application Warning RTT adjust time (ms):

Application Warning RTT adjust time (ms), is the time (0 to 500 milliseconds), to adjust network-wide application normal round trip time.

WAN Threshold Overview:

At a site, all WAN Links will have the threshold event enabled by default. (Advanced->Sites->SiteName)->WAN Links->WANLinkName)->Settings->Advanced Settings->Enable WAN Link Threshold Event). Enabling it means that WAN link will be included in the computation of a threshold event. When the average combined WAN Egress or Ingress usages in a minute crosses this set threshold, an event will be generated. The maximum allowed configured threshold value is 100,000,000 Kbps and a value of 0 means it is disabled. There are 4 configurable thresholds:

- **WAN Ingress Lower Threshold (Kbps):**

For example, this may be used in a WAN Ingress Usage Threshold warning alarm.

- **WAN Ingress Higher Threshold (Kbps):**

For example, this may be used in a WAN Ingress Usage Threshold critical alarm.

- **WAN Egress Lower Threshold (Kbps):**

For example, this may be used in a WAN Egress Usage Threshold warning alarm.

- **WAN Egress Higher Threshold (Kbps):**

For example, this may be used in a WAN Egress Usage Threshold critical alarm.

Routing Domains

Routing Domains can be enabled, disabled, or set as the default on a Site by Site basis.

Interface Groups

An Interface Group allows one or more **Ethernet Interfaces** that share a common subnet to be configured together. If the subnet is behind a firewall or other secure device, the **Security Zone** should be set to **Trusted**. Untrusted interfaces will permit only Conduit, ICMP and ARP traffic.

Pass-through traffic may be enabled between two Ethernet Interfaces by creating a **Bridge Pair**. Setting the **Bypass Mode** to **Fail-to-Wire** will enable a physical connection between the

bridge pairs, allowing traffic to flow in the event of SD-WAN Edge restart or failure. Setting the Bypass Mode to **Fail-to-Block** will disable the physical connection between the bridge pairs, preventing traffic from flowing in the event of SD-WAN Edge restart or failure. Only interfaces forming a hardware bypass pair are eligible for **Fail-to-Wire**.

VLAN traffic may be routed by creating a **Virtual Interface**. Traffic matching the given VLAN ID will be routed by the Oracle SD-WAN Edge based on user configuration while undefined VLAN traffic will pass through. Each Virtual Interface must have an associated Virtual IP Address. Each Virtual Interface is automatically associated with the default **Routing Domain**, but you can choose a different one from the drop-down list of configured Routing Domains for each site.

Enabling **Port State Reflection** on a Bridge Pair forces the SD-WAN Edge to administratively take the WAN-side port of a bridge pair down when its corresponding LAN-side port goes down and vice versa. This completely stops the flow of traffic through the bridge pair. This value can only be set when the automatic bridging has not been enabled on the Interface Group via the Bridged attribute.

Virtual IP Addresses are IP addresses for a Site on a particular Virtual Interface. The Virtual IP address is used for communications between Sites across the Oracle SD-WAN and can be used as next-hop routes for traffic transmitted across the Oracle WAN. Each Virtual IP Address must be associated with a **Routing Domain**. The Routing Domain determines which Virtual Interfaces can be used.

- **IP Address / Prefix:** The full host and netmask of the Virtual IP Address.
- **Routing Domain:** A drop-down menu of available Routing Domains.
- **Virtual Interface:** A drop-down menu of available Virtual Interfaces determined by the Routing Domain.
- **Identity:** If you click the Identity checkbox, the Virtual IP Address will be used as the peering IP for use with IP services (e.g., Used as the Source IP Address when peering with devices participating in dynamic routing).
- **Private:** If you click the Private checkbox, the Virtual IP Address remains local to the SD-WAN Edge and is not shared with remote SD-WAN Edge.
- **Security:** The security of the Virtual Interface Group's segment of the network. Security is either Trusted (i.e., protected by a firewall) or Untrusted.

Define WAN Links

A WAN Link provides for the description of individual Internet or Intranet connections of a Site to the WAN or direct connections to other Sites. Individual uses of a WAN Link would be used to describe connections such as: individual Cable, DSL, fiber or other Internet Service Providers; MPLS, IPSec or other site-to-site VPN connections; backup links such as Cellular or Advanced Wireless providers.

The WAN Link Settings configures the properties and behavior of a WAN Link

WAN Link - Basic Settings

The WAN Link Basic Settings allows for the description of the type of the link and any Public IP Address if available.

WAN Ingress

- **Physical Rate** is the bit rate limit of the WAN Link for the traffic traveling from the LAN into the WAN. Configuration should match the physical capacity of the WAN Link purchased from the service provider.

- **Permitted Rate** is the bit rate that the SD-WAN Edge is allowed to use from the physical rate. Configuration should not be more than the physical rate.
- **Auto Learn** indicates whether the permitted rate of the WAN Link will be automatically adjusted based on bandwidth test results. Before a valid test is completed, the physical rate will be used. No matter what the bandwidth test result is, the applied permitted rate will not exceed the physical rate.

WAN Egress

- **Physical Rate** is the bit rate limit of the WAN Link for the traffic traveling from the WAN into the LAN. The Configuration should match the physical capacity of the WAN Link purchased from the service provider.
- **Permitted Rate** is the bit rate that the SD-WAN Edge is allowed to use from the physical rate. The Configuration should not be more than the physical rate.
- **Auto Learn** indicates whether the permitted rate of the WAN Link will be automatically adjusted based on bandwidth test results. Before a valid test is completed, the physical rate will be used. No matter what the bandwidth test result is, the applied permitted rate will not exceed the physical rate.

Access Types:

- **Public Internet:** A Public Internet WAN Link is one that is connected to the Internet via an ISP.
- **Private Intranet:** A Private Intranet WAN Link is one that only connects to one or more sites within the SD-WAN Edge network and can not connect to locations outside of it.
- **Private MPLS:** A Private MPLS WAN Link is a one that uses one or more DSCP tags to control the quality of service between two or more points on an Intranet and cannot connect to locations outside of the SD-WAN Edge network.

Autodetect Public IP: If enabled, the SD-WAN Edge will automatically detect the Public IP Address. Autodetection may not be used on SD-WAN Edge configured as an SD-WAN Controller.

Public IP Address: The IP Address of the NAT or proxy server. Public IP Address is not available when Autodetect IP Address is selected.

Tracking IP Address: A virtual IP address on the WAN Link that can be pinged to determine the state of the WAN Link.

WAN Link - MPLS Queues

The WAN Link MPLS Queues allow for the definition of service queues using standard DSCP tags. At least one Class must be defined for each MPLS link.

MPLS Queue Name: The name of the MPLS Queue.

DSCP Tag: The DSCP tag assigned to the Oracle Conduit Path packets and handled by the carrier for that Queue policy enforcement.

- When creating Conduit Paths between two Private MPLS Wanlinks, paths between each MPLS queues will be automatically generated.
- User traffic may be selected for any tagged Conduit Path during transmission regardless of the original packet DSCP.
- Autopath Groups may be created to associate multiple paths for consolidated configuration.

Unmatched: If enabled, Ingress user traffic with DSCP tags not defined as MPLS queues will use this queue for provisioning and will be re-tagged to with this queue's DSCP. Egress Intranet packets are not re-tagged.

WAN Ingress Permitted Rate (Kbps): The available or allowed rate, in Kbps, for WAN Ingress traffic. The sum of WAN Ingress Permitted Kbps for all queues in a Private MPLS WAN Link may not exceed the WAN Ingress Permitted Kbps for the Private MPLS WAN Link.

WAN Egress Permitted Rate (Kbps): The available or allowed rate, in Kbps, for WAN Egress traffic. The sum of WAN Egress Permitted Kbps for all queues in a Private MPLS WAN Link may not exceed the WAN Egress Permitted Kbps for the Private MPLS WAN Link.

Clicking the Expand Icon () on any row will show the following advanced options:

Tracking IP Address: A virtual IP address on the WAN Link that can be pinged to determine the state of the MPLS Queue.

Congestion Threshold: The amount of congestion (in microseconds) after which the MPLS Queue will throttle packet transmission to avoid further congestion.

Eligibility: The Eligibility settings for an MPLS Queue allow the administrator to influence the usage of the Wanlink, per direction for each Oracle Class (Realtime, Interactive, & Bulk). If Eligibility is disabled, a latency calculation penalty (150ms by default) is applied during path selection to this Wanlink for traffic in that class for the selected direction.

WAN Link - Advanced Settings

The WAN Link Advanced Settings allows the configuration of the ISP specific attributes.

Provider ID: An optional unique ID number, 1-100, to designate WAN Links connected to the same service provider. The Provider ID is used to differentiate Paths when sending duplicate packets.

Frame Cost: Additional header/trailer bytes added to every packet, such as for Ethernet IPG or AAL5 trailers.

Congestion Threshold: The amount of congestion (in microseconds) after which the WAN Link will throttle packet transmission to avoid further congestion.

MTU Size: The largest raw packet size in bytes, not including the **Frame Cost**.

Enable WAN Link Threshold Event Enabling this will include this WAN Link in generating threshold event.

WAN Link Mode:

- **Regular Active:** A regular active WAN link is a primary WAN link on which user traffic is transmitted.
- **On-Demand Standby:** An on-demand standby WAN link is a standby WAN link that may be activated to supplement conduit bandwidth when the bandwidth provided by the primary WAN links in the conduit falls below the bandwidth threshold configured in conduit QoS Policy. In addition, if all primary WAN links are dead or disabled, an on-demand standby WAN link becomes active and user traffic is transmitted on it.

- **Last-Resort Standby:** A last-resort standby WAN link is a standby WAN link on which user traffic is transmitted only when all regular active WAN links and all on-demand standby WAN links are dead or disabled.

NOTE: Only **Public** and **Private Intranet** WAN Links can be configured in **last-resort** or **on-demand** standby mode.

Priority: The configured priority value indicates the order in which a standby WAN link is activated. A priority 1 standby WAN link is activated before a priority 2 standby WAN link. A priority 3 standby WAN link is the last to be activated.

If there are both on-demand standby WAN links and last-resort standby WAN links configured for a conduit, on-demand standby WAN links are always activated before last-resort standby WAN links. Thus, a priority 3 on-demand standby WAN link is activated before a priority 1 last-resort standby WAN link.

Heartbeat Interval: While a standby WAN link is inactive, a heartbeat message is transmitted at this configured interval. If the heartbeat interval is set to "disabled", heartbeat messages are not sent at all while the standby WAN link is inactive. Without heartbeat messages, the actual state of the link and the paths using this link is unknown. This setting is meant for links that are known to be always GOOD. As such, the state of the link is assumed and shown as GOOD in statistics monitoring reports when the link is not active. When such standby WAN link becomes active, heartbeat messages are sent at 25ms or 50ms interval when there is no other traffic.

WAN Link - Eligibility

The Eligibility settings for a WAN Link allows the user to add an extra penalty for using the WAN Link for certain Classes of traffic. When a Class of traffic is marked as not-eligible for the WAN Link, a penalty is added that makes the WAN Link unlikely to be used unless network conditions require it.

WAN Link - Cell Network

The WAN Link Cell Network settings allows for configuration of settings necessary for cell instead of packet based networks, like ATM.

WAN Link - Access Interfaces

An Access Interface defines the **IP Address** and **Gateway IP Address** for a WAN Link. At least one Access Interface is required for each WAN Link. To add an Access Interface, you must choose a Routing Domain which determines which Virtual Interfaces are available for use.

A site must have an Internet Service defined before the **Default Internet Access** checkbox can be enabled.

When using the **Default Internet Access** be certain that configured Routing Domains have non-overlapping network spaces. Any non-directly connected subnets must also have non-overlapping network spaces.

Certificates

Identity Certificates are used to sign or encrypt data to validate the contents of a message and the identity of the sender. Trusted certificates are used to verify message signatures. Oracle SD-WAN Edge accept both Identity Certificates and Trusted Certificates. Administrators can manage certificates in the Configuration Editor.

Identity Certificates

Identity Certificates require that the certificate's private key be available to the signer. Identity Certificates or their certificate chains must be trusted by a peer to validate the contents and identity of the sender. The configured Identity Certificates and their respective Fingerprints are displayed in the Configuration Editor.

Trusted Certificates

Trusted Certificates are self-signed, intermediate certificate authority (CA) or root CA certificates used to validate the identity of a peer. No private key is required for a Trusted Certificate. The configured Trusted Certificates and their respective Fingerprints are listed here.

DHCP

Configure DHCP Server and DHCP relay on the virtual interfaces. The DHCP Server will assign dynamic IP addresses to the connected clients, while the Relay will relay the DHCP requests to the configured DHCP servers.

Server Subnets

NOTE: It is recommended to specify ranges that do not overlap with statically configured IP Addresses or Access Interface IP Addresses for the site.

NOTE: Do not specify a subnet address or broadcast address as part of your DHCP address range.

Configure different virtual interfaces that will be used by the DHCP Server.

- **Routing Domain:** Select a configured Routing Domain when multiple domains are present.
- **Virtual Interface:** Select a Virtual Interface that is configured in server mode.
- **IP Subnet:** The IP Subnet used by the DHCP server to provides addresses for.
- **Domain Name:** Enter the Domain Name that will be sent by the DHCP Server to the clients.
- **Primary DNS:** Enter the Primary DNS that will be sent by the DHCP Server to the clients.
- **Secondary DNS:** Enter the Secondary DNS that will be sent by the DHCP Server to the clients.
- **Enable:** Enable the subnet for use.

Ranges

Configure dynamic IP address pools that will be used to allocate IP addresses to clients.

- **Range Start IP:** The first IP Address in the pool that will be allocated.
- **Range End IP:** The last IP Address in the pool that will be allocated.
- **Gateway IP:** Optional Gateway IP Address that will be sent to the client.
- **Option Set:** Select an option set that will be used to assign various parameters to the server or the connected clients.

Hosts

Configure individual hosts that require a fixed IP address based on the mac address.

- **Fixed IP Address:** Select a fixed IP Address to allocate to the Host.
- **MAC Address:** Enter the MAC address to identify the host.

- **Option Set:** Select a option set that will be used to assign various parameters to the connected host.

Relays

Configure DHCP Relay for individual interfaces that will communicate to remote DHCP Server(s)

- **Routing Domain:** Select a configured Routing Domain when multiple domains are present.
- **Virtual Interface:** Select a Virtual Interface that is configured in relay mode.
- **Server IP 1:** Enter the first DHCP Server IP Address that the Relay will use to forward the request and response from the clients.
- **Server IP 2:** Enter the optional second DHCP Server IP Address that the Relay will use to forward the request and response from the clients.
- **Server IP 3:** Enter the optional third DHCP Server IP Address that the Relay will use to forward the request and response from the clients.
- **Server IP 4:** Enter the optional fourth DHCP Server IP Address that the Relay will use to forward the request and response from the clients.

DNS Proxy

Configure DNS Proxy for each routing domain

- **Routing Domain:** Select a configured Routing Domain when multiple domains are present.
- **IP Address:** Specify the IP address of the primary/secondary DNS server.
- **Use DHCP Client DNS:** Use DHCP client learned DNS server as primary/secondary DNS server.

Override Proxies

Configure DNS servers for DNS request matches certain domain name.

- **Domain Match:** DNS request matching the configured domain will be sent to the override DNS servers.
- **Primary DNS Server IP:** Primary override DNS server.
- **Secondary DNS Server IP:** Secondary override DNS server.

High Availability

When using High Availability (HA), 2 SD-WAN Edge are deployed as a pair with one designated as the Primary and the other as the Secondary. Data traffic is routed through the Primary SD-WAN Edge while the Secondary remains passive. The Secondary SD-WAN Edge monitors the health of the Primary SD-WAN Edge and **if any failures occur** takes over full support of the network services after the designated Failover Time.

HA SD-WAN Edge may be deployed one-arm, or fully-inline in a parallel or serial configuration. When deployed serially, Use Serial Configuration should be checked.

By default, the SD-WAN Edge specified in the Site's Basic Settings will be considered the Primary SD-WAN Edge and the HA SD-WAN Edge, the Secondary. If Swap Primary/Secondary is checked, this designation is flipped and any configuration parameters that apply to the Primary SD-WAN Edge will instead apply to the Secondary.

NOTE: When HA is enabled, all of the Primary SD-WAN Edge's IP addresses are virtualized so that they may be shared with the Secondary SD-WAN Edge.

HA IP Interfaces designate a Virtual Interface and a pair of Virtual IP Addresses over which the two SD-WAN Edge in the HA pair will communicate with one another. Each interface may optionally reference one or more L2 External Trackers—such as a router—that can be used to determine the health of the Primary SD-WAN Edge. Each External Tracking IP Address must reside on the subnet associated with the given Virtual Interface.

Conduits

The Dynamic or Static Conduit Services can be defined in the Service section.

Dynamic Conduits

Dynamic Conduits allows for the enabling and disabling of Dynamic Conduits on the Site. Dynamic Conduits are Conduits that are established directly between sites based on a configured threshold. They are only operational when the defined threshold is reached. The threshold is typically based on traffic. They are not required for normal operation.

Conduit name (static Conduits)

The Conduit Service between two sites that can be configured in this section. The system will add a static Conduit between a client site and the SD-WAN Controller as this Conduit is required. Any additional static Conduit will have to be added manually by the user.

Click the option to add a static Conduit. Click the option next to the Conduit name to delete the Conduit.

Local Site

Local Site allows the user to view and configure the Conduit settings from local Site's perspective. This includes any additional Class or Rules changes the user required for this specific Conduit. The user can also add paths if required.

Remote Site

Remote Site allows the user to view and configure the Conduit settings from remote Site's perspective.

Basic Settings

Disable Reverse Also: Click this button to disable the mirroring of Rules and Classes to both ends of the Conduit. This action can not be undone and the Conduit must be recreated to re-enable Reverse Also.

Tracking IP Address: A virtual IP address on the Path that can be pinged to determine the state of the Path:

- **GOOD:** Reply immediately
- **BAD:** Reply in >100ms (milliseconds)
- **DEAD:** No reply

QoS Policy: Name of the Conduit QoS Policy that will be used to populate Rules and Classes for the Conduit on the Site. This setting will be mirrored when Reverse Also is enabled. When a Conduit QoS Policy is applied to Conduit Service the Classes can only be edited at the Conduit QoS Policy level.

Unlink Classes from QoS Policy: When Conduit QoS Policy is set for Conduit Service and this button is checked and Apply button is clicked, then the Conduit QoS Policy Classes will be copied to the Conduit Service scope where they can be locally edited without affecting the Conduit QoS Policy. This button will not have any effect on the Rules. This button will only be enabled when a Conduit QoS Policy name is set for the Conduit Service. After this if the button is unchecke and Apply button is clicked,

then the Classes at the local scope will be removed and only the Classes defined at the Conduit QoS Policy will be applied to this service. Also the Classes can only be edited at the Conduit QoS Policy. This setting will be mirrored when Reverse Also is enabled.

WAN Links

WAN Link: Name of the WAN Link

Use: Allow the Conduit Service to use the WAN Link. When Use is not enabled, all other options will be unavailable.

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Active MTU Detect: If enabled, all WAN Ingress Paths for Dynamic Conduits will be actively probed for MTU.

UDP Port: The specified port will be used for WAN Ingress packets and required for WAN Egress packets.

UDP Hole Punching: If enabled, the SD-WAN Controller will assist UDP connectivity between compatible NAT-protected client sites.

UDP Port Switching:

- **Enable:** If enabled, the WAN Link will alternate its UDP port at the specified interval. When UDP Port Switching is not enabled, Alt Port and Interval will be unavailable.
- **Alt Port:** The alternate UDP Port to be used when UDP Port Switching is enabled and active.
- **Interval (min):** The interval, in minutes, that the WAN Link will alternate its UDP Port.

Auto-Path Group: The group used to determine what Paths may be automatically generated between the WAN Link and remote WAN Links and what default Path settings to use.

- **<None>** indicates that no group is desired and will prevent Paths from being automatically generated to or from the WAN Link.
- **<Default>** uses the group currently marked as default and is automatically updated when the default group changes.

If the WAN Link is a Private MPLS, then enabling the WAN Link for a service will also allow the row to expand and show options for the individual MPLS Queues. When enabled, clicking the Expand Icon () will show the following options:

- **Use:** Allow the Conduit Service to use the MPLS Queue. When Use is not enabled, all other options will be unavailable. An MPLS Queue may not be used for a service unless the service is first enabled for the Private MPLS WAN Link.
- **DSCP Tag:** The DSCP Tag applied to the Oracle Conduit Path.
- **Auto-Path Group:** The group used to determine what Paths may be automatically generated between the MPLS Queue and remote MPLS Queues and what default Path settings to use. For MPLS Queues, an additional option of <Inherit> is present and will use the following rules to generate paths:
 - **<None>**: no Paths will be created.
 - **<Inherit>**: the Private MPLS' Auto-Path Group setting will be used to create Paths. This MPLS Queue will generate Paths to remote MPLS Queues if the remote Auto-Path Group setting, even if inherited, matches the local setting. If a remote MPLS Queue' Auto-Path Group setting is also <Inherit>, a Path will only be generated if the local and remote DSCP tags are the same.

- **<Default> or a specific group:** This MPLS Queue will generate Paths to remote MPLS Queues if the remote Auto-Path Group setting, even if inherited, matches the local setting, regardless of DSCP tag.

Classes

Any Class specific changes to this Conduit can be entered here and only impact this Conduit. Class options are consistent with a previous description.

Rules

Any rule specific changes to this Conduit can be entered here and only impact this Conduit. Rule options are consistent with a previous description.

Remote Site: typically unused but available if any change is required. This option is consistent with the option available at the local site.

Local Site: allows the user to view and configure the Conduit settings from remote Site's perspective. This includes any additional Class or Rules changes the user requires for this specific Conduit.

Paths

A Path can be created by clicking the Add button next to the Paths category. By default the system will generate paths between WAN Links defined as access type **Public Internet**. The user would be required to use the auto-path group function or enable paths manually for WAN Links with an access type of **Private Intranet**.

Convert to Static Path: Convert Path, and all other Paths associated by WAN Link, generated by an Autopath Group, to a Static Path. This action cannot be undone.

From Site: Source Site for the Path (Read Only).

From WAN Link: Originating WAN Link for the Path (Read Only).

From DSCP Tag: If the From WAN Link is an MPLS Queue, the DSCP Tag associated with the Class (Read Only).

To Site: Destination Site for the Path (Read Only).

To WAN Link: Terminating WAN Link for the path (Read Only).

To DSCP Tag: If the To WAN Link is an MPLS Queue, the DSCP Tag associated with the Class (Read Only).

Reverse Also: If enabled, a Path with the same WAN Links will be built in the opposite direction.

IP DSCP Tagging: DSCP Tag to set in the IP header for Path traffic.

Enable Encryption: If enabled, packets sent in this Path will be encrypted.

Bad Loss Sensitive: If enabled, packet loss will cause the Path to transition to the BAD state and will incur a latency penalty in Path scoring. Disabling this option may be useful when the loss of bandwidth is intolerable.

Percent Loss (%) (default:DEFAULT): This can only be set when Bad Loss Sensitive is set to Custom. If packet loss exceeds the set percentage over the configured time, path state will transit from "GOOD" to "BAD". Default is to use Oracle hard coded algorithm.

Over Time (ms) (default:1000): This can only be set when Bad Loss Sensitive is set to Custom and Percent Loss is set to value other than DEFAULT. If packet loss exceeds

the set percentage over this configured time, path state will transit from "GOOD" to "BAD".

Silence Period (ms) (default:DEFAULT): Path state transitions from "GOOD" to "BAD" when no packets have been received for the specified amount of time. When not specified, the silence period will be automatically calculated according to ongoing network measurements and will transition to BAD after the receiving appliance sends 3 unanswered keepalive requests.

Path Probation Period (ms) (default:10000): Probation period to wait before moving path state from "BAD" to "GOOD". Default is 10000 ms.

Instability Sensitive: If enabled, Latency penalties due to "BAD" state and other spikes in latency are considered in the Path scoring algorithm. Disabling this option may be useful when the loss of bandwidth (if Bad Loss Sensitive enabled) or latency spikes are intolerable.

There are 4 combinations for the bad loss sensitive and instability sensitive settings:

- **Option 1:** When **Bad Loss Sensitive** is set to **Enable** or **Custom** and **Instability Sensitive** is enabled, a Path may be marked as **BAD** and incur a latency penalty so it is only used as a last resort. In the event multiple Paths are marked **BAD**, there is still competition among them based on regular Path scoring.
- **Option 2:** When **Bad Loss Sensitive** is set to **Enable** or **Custom** and **Instability Sensitive** is **disabled**, a Path may be marked as **BAD** and only used as a last resort, however the latency spikes are not considered. In the event multiple paths are marked **BAD**, the ones with **Instability Sensitive disabled** will likely be used first.
- **Option 3:** When **Bad Loss Sensitive** is set to **Disable** and **Instability Sensitive** is **enabled**, a Path remains **GOOD** in spite of loss, however latency spikes are considered, so that Path is only likely to be used after Paths without latency spikes are exhausted.
- **Option 4:** When **Bad Loss Sensitive** is set to **Disable** and **Instability Sensitive** is **disabled**, a Path remains **GOOD** and latency spikes are not considered, therefore the Path will likely remain in constant use.

Tracking IP Address: A virtual IP address on the Path that can be pinged to determine the state of the Path:

- **GOOD:** Reply immediately.
- **BAD:** Reply in >100ms (milliseconds)
- **DEAD:** No reply.

Reverse Tracking IP Address: If Reverse Also in enabled, a virtual IP address on the reverse Path that can be pinged to determine the state of the reverse Path.

- **GOOD:** Reply immediately.
- **BAD:** Reply in >100ms (milliseconds)
- **DEAD:** No reply.

Cloud Services

Cloud Service name
The Cloud Service between the SD-WAN Edge network and the Cloud Gateway that can be configured in this section.

Click the () option to add a Cloud Service. Click the () option next to the Cloud Service name to delete the Cloud Service.

- **Cloud Service:** Select from the list of the Cloud Services that were configured in the **Global** section.

- **Cloud Service QoS Policy:** One and only one Cloud Service QoS Policy is created automatically by the system.

Internet Services

An Internet Service can be created by clicking the Add button next to the Internet Services category. Note, only one Internet Service may exist on a Site.

Internet

An Internet Service can be deleted by clicking the Delete button next to the Internet Service name.

Basic Settings

NOTE: Firewall Zone is not configurable for an Internet Service. If the Service is trusted, it will be assigned to the Internet_Zone. If the Service is untrusted, it will be assigned to the Untrusted)Internet_Zone.

Enable Primary Reclaim: If enabled, the (use = primary) Internet Usage associated with this service on a WAN Link will forcefully reclaim status as the active service on that WAN Link.

Cost of Default Route: If needed the user can change the cost of the default Internet Route 0.0.0.0/0 to a valid value other than the default cost.

QoS Policy: Name of the Internet QoS Policy that will be used to populate Rules for the Internet Service on the Site.

Ignore WAN Link Status: If enabled, packets destined for this service will still choose this service even if all WAN Links for this service are unavailable.

Export Default Route: If enabled, the default route for the Internet Service, 0.0.0.0/0, will be exported to other Sites if WAN-to-WAN Forwarding has been enabled.

WAN Links

Use: Allow the Service to use this WAN Link. When Use is not enabled, all other options will be unavailable.

NOTE: If a last-resort standby WAN link is configured for Internet Service and it is configured with disabled heartbeats (in the Site/WAN Link/Settings/Advanced Settings section), its configured priority value must be the highest among all last-resort standby WAN links that are used by Internet Service. On-demand standby WAN links cannot be configured for Internet Service.

Mode: The Service's mode for traffic redundancy or load balancing

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Access Interface Failover: If enabled, Internet/Intranet packets with mismatched VLAN can still use the service.

WAN Ingress:

- **Tagging:** The DSCP tag to apply to WAN Ingress packets on the Service.
- **Max Delay (ms):** The maximum time, in milliseconds, to buffer packets when the WAN Links bandwidth is exceeded.

WAN Egress:

- **Tagging:** The DSCP tag to apply to WAN Egress packets on the Service.
- **Matching:** Internet WAN Egress packets matching this tag will be assigned to the Service.

- **Grooming:** If enabled, packets will be randomly discarded to prevent WAN Egress traffic from exceeded the Service's provisioned bandwidth.

Rules

The ability to identify traffic based on a rule is the same as previously described. The rule definition will be used to match a specific traffic flow. Once matched the user must define the action to take for the traffic flow. The available actions are described below.

WAN Link: assign the WAN Link that has Internet Service enabled

Override Service:

- **Intranet Service:** override to a defined Intranet service;
- **Discard:** drop the traffic.

Pass-through: map the flow to pass-through and allow the traffic to flow through SD-WAN Edge unchanged.

Intranet Services

An Intranet Service can be created by clicking the () button next to the Intranet Services category.

Intranet

An Intranet Service can be deleted by clicking the () button next to the Intranet Service name.

Basic Settings

Routing Domain: The Routing Domain chosen for the Intranet Service.

Firewall Zone: The Firewall Zone chosen for the Intranet Service. By default, the Service is placed into the Default_LAN_Zone.

Enable Primary Reclaim: If enabled, the (use = primary) Internet Usage associated with this service on a WAN Link will forcefully reclaim status as the active service on that WAN Link.

QoS Policy: Name of the Internet QoS Policy that will be used to populate Rules for the Internet Service on the Site.

Ignore WAN Link Status: If enabled, packets destined for this service will still choose this service even if all WAN Links for this service are unavailable.

WAN Links

Use: Allow the Service to use the WAN Link. When Use is not enabled, all other options will be unavailable.

NOTE: On-demand standby WAN links cannot be configured for Intranet Service.

Mode: The Service's mode for traffic redundancy or load balancing

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Access Interface Failover: If enabled, Internet/Intranet packets with mismatched VLAN can still use the service.

WAN Ingress:

- **Tagging:** The DSCP tag to apply to WAN Ingress packets on the Service.
- **Max Delay (ms):** The maximum time, in milliseconds, to buffer packets when the WAN Links bandwidth is exceeded.

WAN Egress:

- **Tagging:** The DSCP tag to apply to WAN Egress packets on the Service.

- **Matching:** Internet WAN Egress packets matching this tag will be assigned to the Service.
- **Grooming:** If enabled, packets will be randomly discarded to prevent WAN Egress traffic from exceeding the Service's provisioned bandwidth.

If the WAN Link is a Private MPLS, then enabling the WAN Link for a service will also allow the row to expand and show options for the individual MPLS Queues. When enabled, clicking the Expand Icon () will show the following options:

Use: Allow the Service to use this MPLS Queue. When Use is not enabled, all other options will be unavailable. An MPLS Queue may not be used for a service unless the service is first enabled for the Private MPLS WAN Link. Classes marked for unmatched tags must be enabled for Intranet Services.

Unmatched: If enabled, DSCP tags not matched by other MPLS Queues will use this Class. This field is for information purposes only and must be edited in WAN Link -> Settings.

WAN Ingress:

- **Tagging:** The DSCP tag to apply to WAN Ingress packets on the Service. This field is not editable for MPLS Queues.
- **Max Delay (ms):** The maximum time, in milliseconds, to buffer packets when the WAN Links bandwidth is exceeded.

WAN Egress:

- **Tagging:** The DSCP tag to apply to WAN Egress packets on the Service.
- **Matching:** Internet WAN Egress packets matching this tag will be assigned to the Service. This field is not editable for MPLS Queues.
- **Grooming:** If enabled, packets will be randomly discarded to prevent WAN Egress traffic from exceeded the Service's provisioned bandwidth.

Rules

The ability to identify traffic based on a rule is the same as previously describe. The rule definition will be used to match a specific traffic flow. Once matched the user must define the action to take for the traffic flow. The available actions are described below.

WAN Link: assign the WAN Link that has Internet service enabled

Override Service:

- **Intranet Service:** override to a defined Intranet service;
- **Discard:** drop the traffic.

Pass-through: map the flow to pass-through and allow the traffic to flow through SD-WAN Edge unchanged.

LAN GRE Tunnel

The LAN GRE Tunnel feature allows you to configure Oracle SD-WAN Edge to terminate GRE tunnels on the LAN.

- **Routing Domain:** The Routing Domain chosen for the LAN GRE Tunnel.
- **Firewall Zone:** The Firewall Zone chosen for the Tunnel. By default, the Tunnel is placed into the Default_LAN_Zone.

- **Source IP:** The source IP address of the tunnel. This is one of the Virtual Interfaces configured at this site. The available Source IP addresses are determined by the Routing Domain selected.
- **Destination IP:** The destination IP address of the tunnel.
- **Tunnel IP/Prefix:** The tunnel IP address and prefix.
- **Checksum:** Enable or disable Checksum for the tunnel's GRE header.
- **Keepalive Period (s):** The period of time between sending keepalive messages. If configured to 0, no keepalive packets will be sent, but the tunnel will stay up.
- **Keepalive Retries:** The number of times that the Oracle SD-WAN Edge sends keepalive packets without a response before it brings the tunnel down.

IPsec Tunnels

An IPsec Tunnel secures both user data and header information. Oracle SD-WAN Edge can negotiate fixed IPsec Tunnels on the LAN or WAN side with non-Oracle peers. For IPsec Tunnels over LAN, a Routing Domain must be selected. If the IPsec Tunnel uses an Intranet Service, the Routing Domain is pre-determined by the chosen Intranet Service.

- **Service Type:** Choose either **Intranet**, **LAN**, **Palo Alto** or **Zscaler**.
- **Firewall Zone:** The Firewall Zone chosen for the Tunnel. By default, for Service Type **Palo Alto** or **Zscaler** the Tunnel is placed into **Internet Zone** otherwise the **Default_LAN_Zone**.
- **Name:** When **Service Type** is Intranet, choose an Intranet Service the tunnel will protect. Otherwise, enter a name for the service.
- **Local IP:** Choose a local **Virtual IP Address** to use as the local tunnel end point.
- **Peer IP:** Enter the remote peer's IP address.
- **MTU:** Enter the maximum IKE or IPsec packet size between 576 and 1500.
- **Keepalive:** Click the checkbox to keep the tunnel active and enable route eligibility for routes to the **Intranet Service** or **LAN** IPsec tunnel.

IKE Settings

Internet Key Exchange (IKE) is an IPsec protocol used to create a security association (SA). Oracle SD-WAN Edge supports both the IKEv1 and IKEv2 protocols. The Configuration Editor allows you to modify the following IKE settings:

- **Version:** Choose either the **IKEv1** or **IKEv2** protocol.
- **Mode:** Choose either **Main Mode** or **Aggressive Mode**.
- **Identity:** Choose **Auto** to automatically identify the peer, or choose IP Address to specify the peer's IP.
- **Authentication:** Choose either **Pre-Shared Key** or **Certificate** as the method of authentication.
- **Pre-Shared Key:** If you choose **Pre-Shared Key** authentication, enter the key into this required field. To reveal the text you entered into the Pre-Shared Key field, click the Eye icon (). To hide the text, click the Eye icon () again.
- **Certificate:** If you choose **Certificate** authentication, you can choose from the existing, configured certificates. The default is **None**.
- **Validate Peer Identity:** Enable or disable validation of the IKE's Peer Identity if the peer's ID type is supported, otherwise do not enable this feature.
- **DH Group:** The following Diffie-Hellman groups are available for IKE key generation:

- **Group 1:** 768-bit group
- **Group 2:** 1024-bit group
- **Group 5:** 1536-bit group
- **Hash Algorithm:** The **MD5**, **SHA1**, and **SHA-256** Hash Algorithms are available for IKE messages.
- **Encryption Mode:** **AES-128**, **AES-192**, **AES-256**, and **GCMAES 256 Bit** Encryption Modes are available for IKE messages.
- **Lifetime (s):** Your preferred duration, in seconds, for an IKE security association to exist. Enter 0 for unlimited.
- **Lifetime Max (s):** Your maximum preferred duration, in seconds, to allow an IKE security association to exist. Enter 0 for unlimited.
- **DPD Timeout (s):** The time, in seconds, after receiving no packets or DPD replies to an IKE peer is considered DEAD. Enter 0 to disabled Dead Peer Detection.
- **IKEv2 Settings**
- **Peer Authentication:** **Mirrored**, **Pre-Shared Key**, and **Certificate** modes are available for Peer Authentication.
- **Peer Pre-Shared Key:** If you choose **Pre-Shared Key** authentication, enter the key into this required field.
- **Integrity Algorithm:** The **MD5**, **SHA1**, and **SHA-256** hashing algorithms are available for HMAC verification.

IPsec Settings

The Configuration Editor allows you to modify the following IPsec settings:

- **Tunnel Type:** Choose **ESP**, **ESP+Auth**, **AH** or **ESP-NULL** as the Tunnel Encapsulation Type.
- **ESP:** Encrypts the user data only
- **ESP+Auth:** Encrypts the user data and includes an HMAC
- **AH:** Only includes an HMAC
- **ESP-NULL:** This is the default setting for Zscaler Internet Service connecting to Cloud Security Provider
- **ESP+Auth:** This is the default setting for Palo Alto Internet Service connecting to Cloud Security Provider
- **PFS Group:** The following Diffie-Hellman groups are available for perfect forward secrecy key generation:
 - **None**
 - **Group 1:** 768-bit group
 - **Group 2:** 1024-bit group
 - **Group 5:** 1536-bit group
- **Encryption Mode:** **AES-128**, **AES-192**, **AES-256**, and **GCMAES 256 Bit** Encryption Modes are available for IPsec messages. Not applicable for Internet service with Tunnel Type **ESP-NULL**.
- **Hash Algorithm:** The **MD5**, **SHA1**, and **SHA-256** hashing algorithms are available for HMAC verification. Applicable for Tunnel Type **ESP+Auth**.

- **Lifetime (s):** Your preferred duration, in seconds, for an IPsec security association to exist. Enter 0 for unlimited.
- **Lifetime Max (s):** Your maximum preferred duration, in seconds, to allow an IPsec security association to exist. Enter 0 for unlimited.
- **Lifetime (KB):** The amount of data, in kilobytes, for an IPsec security association to exist. Enter 0 for unlimited.
- **Lifetime Max (KB):** The maximum amount of data, in kilobytes, to allow an IPsec security association to exist. Enter 0 for unlimited.
- **Network Mismatch Behavior:** The action for the Oracle WAN to take if a packet does not match the IPsec Tunnel's Protected Networks from the drop-down menu.
- LAN Tunnels can **Drop** the packets or **Use Non-IPsec Routes** to transmit them.
- Intranet Tunnels can **Drop** the packets or **Send Unencrypted** packets.
- Palo Alto or Zscaler Internet Tunnels can **Drop** the packets or **Use Non-IPsec Routes** to transmit them.

IPsec Protected Networks

- **Source IP/Prefix:** The Source IP and Prefix of the network traffic the IPsec Tunnel will protect.
- **Destination IP/Prefix:** The Destination IP and Prefix of the network traffic the IPsec Tunnel will protect.

Firewall

Firewall allows for the filtering and translation of traffic in the the SD-WAN Edge network.

Settings

Settings allows for the configuration of Policy Templates for the Site and other settings that apply to only an individual SD-WAN Edge.

Policy Templates

Policy Templates allows users to deploy Firewall Policy Templates to a Site.

Click the () option to add a Template.

The Template consists of the following options:

Priority: The order/precedence in which templates will be applied.

Name: The name of the Policy Template to use at the Site.

Advanced

Advanced allows users to modify certain behaviors for the site. The following options can be changed:

Default Firewall Action:

- **Use Global Setting:** Use the Global setting configured in SD-WAN Edge Network Settings
- **Allow:** Packets not matching any filter policy is permitted.
- **Drop:** Packets not matching any filter policy is dropped.

Default Connection State Tracking:

- **Use Global Setting:** Use the Global setting configured in SD-WAN Edge Network Settings

- **No Tracking:** Bidirectional connection state tracking will not be performed on packets not matching any filter policy.
- **Track:** Bidirectional connection state tracking will be performed on TCP, UDP and ICMP packets not matching any filter policy or NAT rule. This feature will block flows which appear illegitimate, due to asymmetric routing or failure of checksum, protocol specific validation -- proceed with caution if the Oracle SD-WAN Edge is not fully inline. For conduit to conduit TCP flows, sequence window check is ignored. The recommendation is to enable this at both end sites.

Untracked and Denied Timeout (s): The time, in seconds, to wait for new packets before closing Untracked or Denied Connections.

TCP Initial Timeout (s): The time, in seconds, to wait for new packets before closing a TCP session that has not completed a handshake.

TCP Idle Timeout (s): The time, in seconds, to wait for new packets before closing an active TCP session.

TCP Closing Timeout: The time, in seconds, to wait for new packets before closing a TCP session after a request to terminate.

TCP Time Wait Timeout (s): The time, in seconds, to wait for new packets before closing a terminated TCP session.

UDP Initial Timeout (s): The time, in seconds, to wait for new packets before closing a UDP session that has not seen traffic in both directions.

UDP Idle Timeout (s): The time, in seconds, to wait for new packets before closing an active UDP session.

ICMP Initial Timeout (s): The time, in seconds, to wait for new packets before closing an ICMP session that has not seen traffic in both directions.

ICMP Idle Timeout (s): The time, in seconds, to wait for new packets before closing an active ICMP session.

Generic Initial Timeout (s): The time, in seconds, to wait for new packets before closing a generic session that has not seen traffic in both directions.

Generic Idle Timeout (s): The time, in seconds, to wait for new packets before closing an active generic session.

Policies

Policies allows for the configuration of Filtering policies on the local SD-WAN Edge as well as the display of policies as they apply from the Global or Site Policy Templates. Policies for an SD-WAN Edge will be applied in the following order:

- Pre-Policies from Templates configured in Firewall Settings according to the order of the templates.
- Pre-Policies from the Global Template configured in SD-WAN Edge Network Settings.
- Policies configured locally.
- Policies automatically created to support NAT or Port Forwarding policies.
- Post-Policies from Templates configured in Firewall Settings according to the order of the templates.
- Post-Policies from the Global Template configured in SD-WAN Edge Network Settings.

Pre Template Policies

The Pre Template Policies are policies from the Templates configured in Firewall Settings or the Global Policy Template configured in SD-WAN Edge Network Settings. These policies will apply before policies statically configured for the Site in the order dictated by the Firewall Settings followed by the policies from the Global Policy Template.

Filter Policies

Filter Policies allows for the configuration of packet filtering for the Site.

NOTE: When filtering using Zones, traffic using a Conduit Route manually configured in the **Routes** section does not know the **To Zone** until the traffic arrives at the remote Site. Filter Policies for this traffic must be configured at the remote Site.

NOTE: When filtering using Zones, traffic using a Conduit Route generated by a Discard Route from a remote Site does not know the **To Zone** until the traffic arrives at the remote Site. Filter Policies for this traffic must be configured at the Site where the Discard Route is configured.

NOTE: When filtering using Zones, traffic from a private VIP may only be filtered at the local Site using the Zone for the **Private** VIP. Similarly, if the Source IP address for a packet is translated using NAT, the original **Inside Zone** can only be filtered locally. All remote SD-WAN Edge must use the **Outside Zone**.

Select the () option to add a Policy. The Policy consists of the following options:

Priority: The order/precedence in which Filters are applied (automatically redistributed on Apply).

Routing Domain: If selected, the Routing Domain this Filter will apply to.

From Zones: The Zone(s) a packet originates from that the Filter will match.

To Zones: The Zone(s) a packet is destined to that the Filter will match.

Action:

- **Allow:** Traffic matching this rule is permitted.
- **Drop:** Traffic matching this rule is dropped.
- **Reject:** Traffic matching this rule is rejected. For TCP, a reset message is sent; For UDP, an ICMP port unreachable message is sent; For other traffic, an ICMP protocol unreachable message is sent.
- **Count and Continue:** Traffic matching this rule is counted on the rule but no action is taken.

Log Interval (s): The time, in seconds, between logging the number of packets matching the filter (0 = disabled, valid settings are 60-600).

Log Start: Click the checkbox to generate a log when a new Connection is created by a packet matching this Filter.

Log End: Click the checkbox to generate a log when a Connection matching this Filter is deleted.

Connection State Tracking:

- **Use Site Setting:** Use the Site setting configured in Firewall -> Settings
- **No Tracking:** Bidirectional connection state tracking will not be performed on packets matching the filter policy.

- **Track:** Bidirectional connection state tracking will be performed on TCP, UDP and ICMP packets matching the filter policy. This feature will block flows which appear illegitimate, due to asymmetric routing or failure of checksum, protocol specific validation -- proceed with caution if the Oracle SD-WAN Edge is not fully inline.

IP Protocol: The IP Protocol that the Filter will match.

DSCP: The DSCP tag that the Filter will match.

Allow Fragments: Click the checkbox to allow fragmented packets matching the Filter.

Reverse Also: Click the checkbox to automatically add a copy of this Filter with the Source (including From Zones) and Destination (including To Zones) settings reversed. The new policy will be created immediately after the original policy in the SD-WAN Edge's filter table.

Source Service Type: The Service Type that the Filter will match.

Source Service Name: The Service that the Filter will match.

Source IP: The Source IP Address and Subnet Mask that the Filter will match.

Source Port: The Source Port or Port Range that the Filter will match.

Dest Service Type: The Destination Type that the Filter will match.

Dest Service Name: The Destination that the Filter will match.

Dest IP: The Destination IP Address and Subnet Mask that the Filter will match.

Dest Port: The Destination Port or Port Range that the Filter will match.

Post Template Policies

The Post Template Policies are policies from the Templates configured in Firewall Settings or the Global Policy Template configured in SD-WAN Edge Network Settings. These policies will apply after policies statically configured for the Site in the order dictated by the Firewall Settings followed by the policies from the Global Policy Template.

Static NAT Policies

Static NAT Policies allows for the configuration of Network Address Translation policies between individual hosts or subnets.

NOTE: NAT translations are not permitted if the Inside and Outside Zones are the same.

NOTE: While both Inbound and Outbound translations can be configured simultaneously for a Service, only the first to match will be used. Multiple translations may occur if a rule exists on the Service a packet is received on and the Service a packet is sent on.

Select the () option to add a Policy. The Policy consists of the following options:

Priority: The order/precedence in which translations are applied (automatically redistributed on Apply).

Routing Domain: If selected, the Routing Domain this translation will apply to.

Allow Return Flow: If enabled, this policy would allow return flow traffic also.

Direction:

- **Inbound:** The source address for a packet will be translated for packets received on the Service. The destination address will be translated for packets transmitted on the Service.
- **Outbound:** The destination address for a packet will be translated for packets received on the Service. The source address will be translated for packets transmitted on the Service.

Service Type: The Service Type that the translation applies to.

Service Name: The Service Name that the translation applies to.

Inside Zone: The Zone a packet must be from to allow translation. The Inside Zone is inferred from the configured Service for inbound rules.

Inside IP Address: The Inside IP Address and Prefix to translate (Source IP Address in the direction selected).

Outside Zone: The Zone a packet must be destined for to allow translation. The Outside Zone is inferred from the configured Service for outbound rules.

Outside IP Address: The Outside IP Address and Subnet Mask packets will be translated to (Source IP Address in the direction selected).

Dynamic NAT Policies

Dynamic NAT Policies allows for the configuration of Network Address Port Translation policies between an inside network and an outside IP address.

NOTE: Dynamic NAT translations allow all reciprocal traffic for session initiated from the Inside Network. To filter these connections, add filter Policies for the outbound traffic.

NOTE: NAT translations are not permitted if the Inside and Outside Zones are the same.

NOTE: While both Inbound and Outbound translations can be configured simultaneously for a Service, only the first to match will be used. Multiple translations may occur if a rule exists on the Service a packet is received on and the Service a packet is sent on.

Select the () option to add a Policy. The Policy consists of the following options:

Priority: The order/precedence in which translations are applied (automatically redistributed on Apply).

Direction:

- **Inbound:** The source address for a packet will be translated for packets received on the Service. The destination address will be translated for packets transmitted on the Service.
- **Outbound:** The destination address for a packet will be translated for packets received on the Service. The source address will be translated for packets transmitted on the Service.

Type:

- **Port Restricted:** Port Restricted NAT uses the same outside port for all translations related to an Inside IP Address and Port pair. This mode is typically used to allow Internet P2P applications (hole punching).
- **Symmetric:** Symmetric NAT uses the same outside port for all translations related to an Inside IP Address, Inside Port, Outside IP Address and Outside Port tuple. This mode is typically used to enhance security or expand the maximum number of NAT sessions.

Service Type: The Service Type that the translation applies to.

Service Name: The Service Name that the translation applies to.

Inside Zone: The Zone a packet must be from to allow translation. The Inside Zone is inferred from the configured Service for inbound rules.

Inside IP Address: The Inside IP Address and Prefix to translate (Source IP Address in the direction selected).

Outside Zone: The Zone a packet must be destined for to allow translation. The Outside Zone is inferred from the configured Service for outbound rules or the Outside IP Address for inbound rules.

Outside IP Address: The Outside IP Address packets will be translated to (Source IP Address in the direction selected).

Allow Related: Click the checkbox allow packets related to a Connection (ICMP error packets).

IPsec Passthrough: Click the checkbox to allow an IPsec (AH/ESP) session to be translated. Only a single session from the inside network will be permitted.

GRE/PPTP Passthrough: Click the checkbox to allow a GRE/PPTP session to be translated. Only a single session from the inside network will be permitted.

Note: When the Internet Service is added with an untrusted WAN Link (WANLink usage that have Untrusted Ports in the Interface Group), the system by default adds a Dynamic NAT Policy for Outbound Direction and Service Type Internet (unless the user has already created one). This Dynamic NAT policy is editable by the user. Also if the user deletes this policy the system will add one back. The Dynamic NAT policy that is created by the system will be removed when the Internet Service is trusted or removed.

Port Forwarding Rules

Port Forwarding Rules allow traffic from an Outside network to access specific hosts and ports on the Inside network without the session being initiated from the inside.

Select the option to add a Rule. The Rule consists of the following options:

Routing Domain: The Routing Domain this Rule will match. For Port Forwarding Rules on Local and Intranet Services, the Routing Domain is inferred from the Service.

Protocol:

- **Both:** Both TCP and UDP ports will be forwarded.
- **TCP:** Only TCP ports will be forwarded.
- **UDP:** Only UDP ports will be forwarded.

Outside Port: The Outside port or port range to forward.

Inside IP Address: The Inside IP Address to forward to.

Inside Port: The Inside port or port range to forward to. If a range is configured, it must define the name number of ports as the **Outside Port**.

Fragments: Click the checkbox to enable forwarding of packet fragments.

Log Interval (s): The time, in seconds, between logging the number of packets matching the rule (0 = disabled, valid settings are 60-600).

Log Start: Click the checkbox to generate a log when a new Connection is created by a packet matching this Rule.

Log End: Click the checkbox to generate a log when a Connection matching this Rule is deleted.

Connection State Tracking:

- **Use Site Setting:** Use the Site setting configured in Firewall -> Settings
- **No Tracking:** Bidirectional connection state tracking will not be performed on packets matching the Rule.
- **Track:** Bidirectional connection state tracking will be performed on TCP, UDP and ICMP packets matching the Rule. This feature will block flows which appear illegitimate, due to asymmetric routing or failure of checksum, protocol specific validation -- proceed with caution if the Oracle SD-WAN Edge is not fully inline.

WAN Link

Based on the site the configured WAN Links will be displayed. The Configuration options include:

Dynamic Conduit Thresholds

The Dynamic Conduit Thresholds describe conditions for which an Intermediate Site will trigger the creation or destruction of a Dynamic Conduit between two adjacent sites. Creating a Dynamic Conduit may be triggered by exceeding either packets per second or throughput. Deleting a Dynamic Conduit is triggered when both thresholds are no longer met.

NOTE: Additional thresholds are defined in the Dynamic Conduit QoS Policy. A Dynamic Conduit will be created when the thresholds are met for either the WAN Link thresholds defined here or the thresholds defined in the Dynamic Conduit QoS Policy.

Ability to configure thresholds for Dynamic Conduits. The options allow the user to configure a threshold based on packets per second or bytes per second. Once the threshold is reached the Dynamic Conduit will be created between the appropriate sites.

- **WAN Ingress**
- **Enable Kbps Threshold:** If enabled, allows setting of Throughput (Kbps) threshold trigger
- **Throughput (Kbps):** The threshold, in Kbps, on the intermediate site at which Dynamic Conduits will be triggered on WAN Ingress.
- **Enable pps threshold:** If enabled, allows setting of Throughput (pps) threshold trigger
- **Throughput (pps):** The threshold, in packets per second, on the intermediate site at which Dynamic Conduits will be triggered on WAN Ingress.
- **WAN Egress**
- **Enable Kbps Threshold:** If enabled, allows setting of Throughput (Kbps) threshold trigger
- **Throughput (Kbps):** The threshold, in Kbps, on the intermediate site at which Dynamic Conduits will be triggered on WAN Ingress.
- **Enable pps threshold:** If enabled, allows setting of Throughput (pps) threshold trigger
- **Throughput (pps):** The threshold, in packets per second, on the intermediate site at which Dynamic Conduits will be triggered on WAN Egress.

Conduit Services

Conduit Service: Name of the Conduit Service

Use: Allow the Conduit Service to use this WAN Link. When Use is not enabled, all other options will be unavailable.

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Active MTU Detect: If enabled, all WAN Ingress Paths for Dynamic Conduits will be actively probed for MTU.

UDP Port: The specified port will be used for WAN Ingress packets and required for WAN Egress packets.

UDP Hole Punching: If enabled, the SD-WAN Controller will assist UDP connectivity between compatible NAT-protected client sites.

UDP Port Switching:

- **Enable:** If enabled, the WAN Link will alternate its UDP port at the specified interval. When UDP Port Switching is not enabled, Alt Port and Interval will be unavailable.
- **Alt Port:** The alternate UDP Port to be used when UDP Port Switching is enabled and active.
- **Interval (min):** The interval, in minutes, that the WAN Link will alternate its UDP Port.

Auto-Path Group: The group used to determine what Paths may be automatically generated between the WAN Link and remote WAN Links and what default Path settings to use.

- **<None>** indicates that no group is desired and will prevent Paths from being automatically generated to or from the WAN Link.
- **<Default>** uses the group currently marked as default and is automatically updated when the default group changes.

If the current WAN Link is a Private MPLS, then enabling the WAN Link for a service will also allow the row to expand and show options for the individual MPLS Queues. When enabled, clicking the Expand Icon () will show the following options:

- **Use:** Allow the Conduit Service to use the MPLS Queue. When Use is not enabled, all other options will be unavailable. An MPLS Queue may not be used for a service unless the service is first enabled for the Private MPLS WAN Link.
- **DSCP Tag:** The DSCP Tag applied to the Oracle Conduit Path..
- **Auto-Path Group:** The group used to determine what Paths may be automatically generated between the MPLS Queue and remote MPLS Queues and what default Path settings to use. For MPLS Queues, an additional option of <Inherit> is present and will use the following rules to generate paths:
 - **<None>:** no Paths will be created.
 - **<Inherit>:** the Private MPLS' Auto-Path Group setting will be used to create Paths. This MPLS Queue will generate Paths to remote MPLS Queues if the remote Auto-Path Group setting, even if inherited, matches the local setting. If a remote MPLS Queue' Auto-Path Group setting is also <Inherit>, a Path will only be generated if the local and remote DSCP tags are the same.
 - **<Default> or a specific group:** This MPLS Queue will generate Paths to remote MPLS Queues if the remote Auto-Path Group setting, even if inherited, matches the local setting, regardless of DSCP tag.

Cloud Services

Cloud Service: Name of the Cloud Service

Use: Allow the Cloud Service to use this WAN Link. When Use is not enabled, all other options will be unavailable.

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Active MTU Detect: If enabled, all WAN Ingress Paths for Cloud Conduits will be actively probed for MTU.

UDP Port: The specified port will be used for WAN Ingress packets and required for WAN Egress packets.

UDP Hole Punching: If enabled, the SD-WAN Controller will assist UDP connectivity between compatible NAT-protected client sites.

UDP Port Switching:

- **Enable:** If enabled, the WAN Link will alternate its UDP port at the specified interval. When UDP Port Switching is not enabled, Alt Port and Interval will be unavailable.
- **Alt Port:** The alternate UDP Port to be used when UDP Port Switching is enabled and active.
- **Interval (min):** The interval, in minutes, that the WAN Link will alternate its UDP Port.

Auto-Path Group: The group used to determine what Paths may be automatically generated between the WAN Link and remote WAN Links and what default Path settings to use.

- **<None>** indicates that no group is desired and will prevent Paths from being automatically generated to or from the WAN Link.
- **<Default>** uses the group currently marked as default and is automatically updated when the default group changes.

Internet/Intranet Service Usage

Use: Allow the Service to use this WAN Link. When Use is not enabled, all other options will be unavailable.

Mode: The Service's mode for traffic redundancy or load balancing

Tunnel Header Size (bytes): The size of the tunnel header, in bytes, if applicable.

Access Interface Failover: If enabled, the Service will fail over to the secondary Access Interface when the primary is unavailable.

WAN Ingress:

- **Tagging:** The DSCP tag to apply to WAN Ingress packets on the Service.
- **Max Delay (ms):** The maximum time, in milliseconds, to buffer packets when the WAN Links bandwidth is exceeded.

WAN Egress:

- **Tagging:** The DSCP tag to apply to WAN Egress packets on the Service.
- **Matching:** Internet WAN Egress packets matching this tag will be assigned to the Service.
- **Grooming:** If enabled, packets will be randomly discarded to prevent WAN Egress traffic from exceeded the Service's provisioned bandwidth.

If the current WAN Link is a Private MPLS, then enabling the WAN Link for a service will also allow the row to expand and show options for the individual MPLS Queues. When enabled, clicking the Expand Icon () will show the following options:

Use: Allow the Service to use this MPLS Queue. When Use is not enabled, all other options will be unavailable. An MPLS Queue may not be used for a service unless the service is first enabled for the Private MPLS WAN Link. Classes marked for unmatched tags must be enabled for Intranet Services.

Unmatched: If enabled, DSCP tags not matched by other MPLS Queues will use this Class. This field is for information purposes only and must be edited in WAN Link -> Settings.

WAN Ingress:

- **Tagging:** The DSCP tag to apply to WAN Ingress packets on the Service. This field is not editable for MPLS Queues.
- **Max Delay (ms):** The maximum time, in milliseconds, to buffer packets when the WAN Links bandwidth is exceeded.

WAN Egress:

- **Tagging:** The DSCP tag to apply to WAN Egress packets on the Service.
- **Matching:** Internet WAN Egress packets matching this tag will be assigned to the Service. This field is not editable for MPLS Queues.
- **Grooming:** If enabled, packets will be randomly discarded to prevent WAN Egress traffic from exceeded the Service's provisioned bandwidth.

Routes

Ability to add or remove Oracle routes from a site. Once expanded the user can view the configured routes for the site.

Select the () option to add a route. The route consists of the following option:

Network IP Address: Route to be added. Requires the network address and mask.

Routing Domain: The Routing Domain chosen for the Route. New Routes are automatically associated with the default Routing Domain.

Cost: Oracle cost for the route

Service types:

- **Conduit:** Identifies IP traffic as Conduit traffic and matches a Conduit based on Conduit Rules.
- **Internet:** Identifies IP traffic as Internet traffic and matches the Internet Service.
- **Intranet:** Identifies IP traffic as Intranet traffic and matches an Intranet Service based on the Intranet Rules.
- **Cloud:** Identifies IP traffic to a Cloud Service.
- **Pass-through:** Identifies IP traffic as Pass-through and matches the Pass-through Service.
- **Local:** Identifies IP traffic as local to the site and matches no service. Traffic sourced and destined to a local route will be ignored.
- **LAN GRE Tunnel:** Identifies IP traffic as local to the site and matches LAN-side GRE tunnel service.
- **LAN IPsec Tunnel:** Identifies IP traffic as local to the site and matches an LAN IPsec Tunnel.

- **Discard:** Identifies IP traffic as local to the site and should be discarded. Discard routes are exported to remote Sites as Conduit Routes and can be used to facilitate routing of NAT traffic to specific Sites.

Gateway IP address: define the gateway/router to reach this route. Certain routes require a gateway.

Delete: ability to delete a route.

Route Learning

Oracle SD-WAN Edge have the ability to learn routes, gather link state information, construct a map of the network, and implement changes dynamically.

Open Shortest Path First (OSPF)

A routing protocol, supported by Oracle SD-WAN Edge, that uses a link state routing algorithm to detect changes in the network topology and re-route packets by computing the shortest path free for each route.

Basic Settings

- **Enable:** Enable or disable OSPF with this checkbox.
- **Advertise SD-WAN Edge Routes:** Enable advertisement of SD-WAN Edge routes via OSPF.
- **Router ID:** A Router ID, in IPv4 Format, used for OSPF advertisements.
- **Export OSPF Route Type:** Advertise the SD-WAN Edge route to OSPF neighbors as type 1 Intra-area route or type 5 External route.
- **Export OSPF Route Weight:** The cost advertised to OSPF neighbors is the original SD-WAN Edge cost plus the weight configured here.

OSPF Areas

- **ID:** The IP Address or Area ID of the network that OSFP will learn routes from and advertise routes to.
- **Stub Area:** Enabling the Stub Area feature ensures that this Area will not receive route advertisements from outside of the designated Autonomous System.
- **Virtual Interfaces**
- **Routing Domain:** A configured Routing Domain.
- **Name:** A configured Virtual Interface.
- **Source IP Address:** The IP address used to send OSPF messages for this interface.
- **Interface Cost:** The base cost for routes learned on the interface.
- **Authentication Type:** **None**, **Plain Text**, and **MD5** are supported.
- **Hello Interval:** The amount of time to wait between sending Hello protocol packets to directly connected neighbors (10 seconds is the default).
- **Dead Interval:** The amount of time to wait to receive a Hello protocol packet before marking a router as dead (40 seconds is the default).

Internal Border Gateway Protocol (BGP)

A routing protocol, supported by Oracle SD-WAN Edge, that is capable of making routing decisions based on Paths determined by ISPs.

Basic Settings

- **Enable:** Enable or disable BGP with this checkbox.

- **Advertise SD-WAN Edge Routes:** Enable advertisement of SD-WAN Edge routes via BGP.
- **Router ID:** A Router ID, in IPv4 Format, used for BGP advertisements.
- **Local Autonomous System:** The Local Autonomous System number.

BGP Neighbors

All of the configured BGP peer routers that are scrutinized to find the shortest paths for data. All of the neighbors must be part of the same Autonomous System.

- **Routing Domain:** A configured Routing Domain.
- **Virtual Interface:** A configured Virtual Interface.
- **Source IP:** The Source IP address for the BGP session.
- **Neighbor IP:** The IP address of the BGP Neighbor router.
- **Hold Time(s):** The Hold Time, in seconds, that elapses before a route is declared dead (the default is 180).
- **Local Preference:** The Local Preference value used for selecting from multiple BGP routes (the default is 100).
- **IGP Metric:** The IGP Metric checkbox enables the comparison of internal distances to calculate the best route.
- **Route Reflector:** The Route Reflector checkbox enables local site to be a route reflector and treat the neighbor as a route reflection client.
- **Next Hop Self:** The Next Hop Self checkbox allows local site to advertise own address as route's next hop.
- **Disable Local AS Loop Protection:** BGP prevents routing loops by rejecting received routes with the local AS number in the AS path. The checkbox disables the check.
- **Password:** The password for MD5 authentication of BGP sessions which is not required.

Import Filters

Network administrators can configure Filters to fine tune how route learning determines the shortest Path for data.

- **Order:** The Order in which filters are applied. Once a filter is applied, the Order is automatically sorted.
- **Source Router:** The IP address of the Source Router.
- **Destination:** The IP Address and Netmask or Network Objects that describe the route's destination.
- **Prefix:** The method (predicate) and prefix length. The predicates are:
 - **eq:** Equal to
 - **lt:** Less than
 - **le:** Less than or equal to
 - **gt:** Greater than
 - **ge:** Greater than or equal to
- **Next Hop:** The IP address of the Next Hop.
- **Protocol:** The protocol to learn routes from (**Any, OSPF, or BGP**).

- **Route Tag:** The 32-bit value attached to routes for redistribution.
- **Cost:** The method (predicate) and route cost. The predicates are:
 - **eq:** Equal to
 - **lt:** Less than
 - **le:** Less than or equal to
 - **gt:** Greater than
 - **ge:** Greater than or equal to
- **Include:** If you do not Include routes that match a filter, those routes are ignored.
- **Enabled:** A filter that is not Enabled has no effect.
- **Delete:** Delete a filter from the configuration.
- **Clone:** Administrators can Clone existing filters to work more efficiently.
- **Export Route to Oracle SD-WAN Edge:** If the **Export Route to Oracle SD-WAN Edge** function is not enabled, the Oracle SD-WAN Edge will not communicate route data to Oracle SD-WAN Edge at other Sites. This functionality is enabled by default and only applies for the following Service Types: Local, and LAN GRE Tunnel.
- **Eligibility Based On Gateway:** If a Gateway is unreachable, this feature will ensure that traffic is not sent to matching routes.
- **SD-WAN Edge Cost:** The cost will be applied to the matched routes when importing into Oracle SD-WAN Edge route table(the default is 6).
- **Service Type:** Choose a Service Type from all of the existing, supported Oracle SD-WAN Edge Services.
- **Recursive Route:** When service type is conduit, check this to allow SD-WAN Edge to find conduit name from imported route's source router automatically.
- **Use Next Hop:** When recursive route is checked, check this to allow SD-WAN Edge to find conduit name from imported routes's next hop instead of source router.
- **Service Name:** The name of the service that matching routes will use.
- **Eligibility Based on Path:** If enabled, Paths become criteria for filters.

Export Filters

Network administrators can configure up to 32 Route Export Filters to narrow the selection of routes to export for advertisement to neighboring routers.

- **Order:** The Order in which filters are exported. Once a filter is applied, the Order is automatically sorted.
- **Routing Domain:** If multiple Routing Domains are configured, choose a Routing Domain from the drop-down menu to narrow the available results or choose Any to include results from all Routing Domains.
- **Network Address:** The IP Address and Netmask or configured Network Object that describes the route's network.
- **Prefix:** The method (predicate) and prefix length. The predicates are:
 - **eq:** Equal to
 - **lt:** Less than
 - **le:** Less than or equal to

- **gt:** Greater than
- **ge:** Greater than or equal to
- **SD-WAN Edge Cost:** The method (predicate) and SD-WAN Edge Route Cost that are used to narrow the selection of routes exported. The predicates are:
- **eq:** Equal to
- **lt:** Less than
- **le:** Less than or equal to
- **gt:** Greater than
- **ge:** Greater than or equal to
- **Service Type:** The Service Type to export from a list of the existing, supported Oracle Services. Choose Any and all matching Service Types are exported.
- **Site/Service Name:** The Site or Service name to export that is determined by the Service Type. Choose Any and all available instances of the chosen Service Type are exported.
- **Gateway IP Address:** If you chose LAN GRE Tunnel or Local as the Service Type, enter the Gateway IP for the tunnel.
- **Include:** Include routes that match the filter for advertisement. If this box is not checked, matching routes are ignored.
- **Enable:** Enable or disable the filter.
- **Clone:** Administrators can Clone existing filters to work more efficiently.
- **Export OSPF Route Type:** Advertise the SD-WAN Edge route to OSPF neighbors as type 1 Intra-area route or type 5 External route.
- **Export OSPF Route Weight:** The cost advertised to OSPF neighbors is the original SD-WAN Edge cost plus the weight configured here.

Site Template

In the **Basic > Sites** view there is also the option to "**Generate Site Template**" based on the selected Site. Clicking this button generates a Site Template based on the current site. You can select which Ethernet Interfaces should be used, set the Bypass Mode, choose whether to bridge the Interfaces, pick a Security setting, add any required VLANs, and set the WAN DHCP Client option. `New_Site_Template` will appear under the Network tab of Basic view on the left-hand side of the page. To change the template name or edit any of its settings, simply click on it.

Provisioning

Provisioning allows for the bidirectional (WAN Ingress/WAN Egress) distribution of bandwidth for a WAN Link among the various services associated with that WAN Link.

There are two steps to Provisioning that provide for this bandwidth distribution in a simple and effective way:

- **Provisioning Groups:** Create and edit groups of bandwidth. (Optional)
- **Services:** View and edit bandwidth settings for services within a bandwidth group

The Concept of Using Shares

When provisioning bandwidth for networks with a large number of sites, using percentages does not allow for enough granularity as the site count increases. With the SD-WAN Edge provisioning process, we introduce the concept of Fair Shares. Shares are used to distribute the permitted bandwidth over groups, and services within groups.

With shares, the total number of shares is up to the user, allowing any amount of granularity or precision when allocating bandwidth among the different Groups and Services.

Shares are discussed in more detail in the '**Services**' sections.

Provisioning Groups

A Provisioning Group is a container for an arbitrary collection of Services on any given WAN Link. They allow the user to allocate bandwidth at a high-level before drilling down to the individual Services within the Group for fine-tuning. They also provide a boundary for the automatic redistribution of bandwidth within the child Services of the Provisioning Group.

NOTE: Provisioning Groups are available to simplify the provisioning process and are not required if they are not needed.

Fair Shares

In the Provisioning Groups table, shares are used to distribute the WAN Ingress/Egress eligible bandwidth, which is the Permitted Rate minus the total Min reserved bandwidth of all Services on the WAN Link. All Services are initially assigned to a **Default Group** that is allocated all of the eligible bandwidth. The user can create additional Groups and allocate bandwidth to its members by giving that Group some number of Fair Shares. The resulting total bandwidth for all Services in the Group is then shown in the Fair (Kbps) column.

NOTE: All Services receive their Min Reserved Bandwidth before Fair distribution, which could result in Groups with equal Fair Shares having disparate Fair Rates. Fair Rates can also be affected by Service Maximums, if defined.

Services

The **Services** section allows the user to further fine-tune bandwidth allocation. Services that are assigned to the same group contend for the bandwidth allocated to that group. The services shown in the **Services** section of the selected WAN Link have been enabled on that WAN Link by the current Configuration.

By default, all services are assigned to the **Default** group with a default number of fair shares divided evenly among them. The default number of shares serves as a starting point and is not restricted by (nor related to) the number of shares set in the **Provisioning Groups** section.

To compute a service's actual fair share of the permitted bandwidth, the following formula is used:

$$s_fs / sg_fs * g_fs$$

- s_fs is the service's fair share as shown in the "Shares of Group" column.
- sg_fs is the sum of the "Shares of Group" for all the services in the same group.
- g_fs is the "Fair Share" value in the Groups table for the service's group.

The result is rounded.

 **Note:**

This is the value that is shown in the View Configuration->WAN Links page as "Computed wan_egress_rate_fair_share" and "Computed wan_ingress_rate_fair_share".

Default minimum rates by service type:

- **Conduit:** 80 Kbps (including Dynamic Conduit type)
- **Internet/Intranet:** 100 Kbps

 **Note:**

To set an unlimited **Max (Kbps)** rate, enter '0' (zero) into the cell.

For the **Dynamic Conduits** service entry (if configured on the selected WAN Link), the **Min (Kbps)** and **Max (Kbps)** fields are variable. Therefore, the range of values that can be expected is shown.

Shares of Group

In the **Services** section, shares are again used to distribute the eligible WAN Ingress/Egress bandwidth. The Group that a service is assigned to determines the eligible bandwidth (listed in the Fair (Kbps) column in the **Provisioning Groups** section) for all services assigned to the same Group. The Shares of Group are used to divide up the eligible bandwidth among the members of a group based on the ratio of the current service divided by the total number of shares for the group in which it is a member.

The Minimum rate acts as a base bandwidth allocation for each service, and the amount of bandwidth available for fair allocation is based on the total permitted for the group minus the sum of the minimums for each service in the group.

In the case of the **Dynamic Conduits** service, the **Shares of Group** is divided among all Dynamic Conduits. Please refer to the context help for **Dynamic Conduit Provisioning** for more information.

New services enabled on a WAN Link will be placed in the **Default** Group with a **Shares of Group** value of 0 (zero). The **Shares of Group** must be configured to a non-zero value in order to be valid.

When moving a service between Groups, the Service will keep its configured amount of shares. The shares will be removed from the old group and taken to the new Group.

When deleting or disabling a service on a WAN Link, that service's shares will be removed as well. The shares will not be distributed over the remaining services.

Dynamic Conduit Provisioning

The **Dynamic Conduit Provisioning** worksheet is for configuring the parameters of an individual Dynamic Conduit. If a **Dynamic Conduit** service is not enabled on the selected WAN Link, the worksheet will be hidden.

The settings in this worksheet should be treated in the same way as a static Conduit service. Set the **Min (Kbps)** and **Max (Kbps)** for an individual Dynamic Conduit here. Each Dynamic Conduit will use the settings provisioned here. Once the **Min (Kbps)** and **Max (Kbps)** rates have been configured, the **Fair (Kbps)** per Dynamic Conduit will be recalculated to reflect the new settings.

Using the **Possible Dynamic Conduits** and the individual Dynamic Conduit settings in this worksheet, worst-case usages of Min, Max and Fair bandwidth will be calculated. The worst-case Min and Max bandwidth will be shown in the **Min Total (Kbps)** and **Max Total (Kbps)** columns (respectively) in the **Dynamic Conduit Provisioning** worksheet. The worst-case fair bandwidth will be shown in the **Fair (Kbps)** column of the main **Services** table for the **Dynamic Conduits** service.

NOTE: To set an unlimited **Max (Kbps)** rate, enter '0' (zero) into the cell.

About the Possible Dynamic Conduits column:

The **Possible Dynamic Conduits** value represents the total number of Dynamic Conduits that could exist simultaneously (based on the current Configuration described momentarily). It is either the total number of sites reachable via Dynamic Conduit OR the maximum number of Dynamic Conduits supported by the platform being configured, whichever is fewer. Several Configuration parameters cooperate to determine if Dynamic Conduits can be created between sites. Sites that have: Dynamic Conduits enabled AND share a common intermediate site with W-T-W Forwarding enabled AND are in the same W-T-W Forwarding group; can establish Dynamic Conduits between each other and are factored into the **Possible Dynamic Conduits** value.

Fair (Kbps)

The fair bandwidth is based on a worst-case scenario in which all accounted Dynamic Conduits are up simultaneously. The number of shares for an individual Dynamic Conduit are used in the calculation. The number of shares for an individual Dynamic Conduit receives in the worst-case is equal to the number of **Shares of Group** defined for the **Dynamic Conduits** service divided by the **Possible Dynamic Conduits**:

For example, if the **Dynamic Conduits** service has 100,000 shares defined for it in the **Shares of Group** column in the Services table, and if the current Configuration accounted for a **Possible Dynamic Conduits** of 4, then 25,000 shares ($100,000 / 4$) will be used as the number of fair shares for an individual Dynamic Conduit.

Once the worst-case number of shares has been calculated, the Fair (Kbps) rate is calculated in the same manner as the other service types.

(Required) Enter introductory text here, including the definition and purpose of the concept.

Global

SD-WAN Edge Network Settings

SD-WAN Edge Network Settings allows for the configuration of global parameters that impact the operation of the entire SD-WAN Edge network. These settings may impact the way individual Site settings are interpreted or applied.

Global Security Settings

Network Encryption Mode: Defines the algorithm used for all encrypted Paths in the SD-WAN Edge network. This setting does not apply to non-encrypted Paths. When changing this setting, the Secure Key for Sites in the SD-WAN Edge network may be modified.

- When changing to AES 128-Bit, keys longer than 16 characters will be truncated to 16 characters.
- When changing to AES 256-Bit, keys shorter than 16 characters will be regenerated as 32 characters.

Enable Encryption Key Rotation: If enabled, Encryption Keys will be regenerated for every Conduit with encryption enabled using an Elliptic Curve Diffie-Hellman key exchange at intervals of 10 - 15 minutes.

Enable Extended Packet Encryption Header: If enabled, a 16 byte, randomly seeded counter will be prepended to the beginning of every encrypted message. When encrypted, this counter will serve as a random Initialization Vector, deterministic only with the encryption key. This will randomize the output of encryption, providing strong message indistinguishability. Note, when enabled, this option will increase packet overhead by 16 bytes.

Enable Extended Packet Authentication Trailer: If enabled, an authentication code will be appended to the end of every encrypted message. This trailer allows for the verification that packets are not modified in transit. Note, when enabled, this option will increase packet overhead

Extended Packet Authentication Trailer Type: The type trailer to include in encrypted messages

WARNING: Using SHA-256 may significantly impact network performance.

- 32-Bit Checksum: A 4 byte value calculated by computing the ones-complement checksum of the encrypted packet's contents.
- SHA-256: A 16 byte value calculated using SHA-256 over the encrypted packets contents.

Global Firewall Settings

Global Firewall Settings allows for the configuration of global parameters that impact the operation of the Firewall on individual SD-WAN Edge.

Global Policy Template: A Firewall Policy template to be applied to all SD-WAN Edge in the network.

Default Firewall Action:

- Allow: Packets not matching any filter policy is permitted.
- Drop: Packets not matching any filter policy is dropped.

Default Connection State Tracking: Click the checkbox to enable bidirectional connection state tracking for TCP, UDP and ICMP flows that do not match a filter policy or NAT rule. Asymmetric flows will be blocked when this is enabled even when there are no Firewall policies defined. The settings may be defined at the site level which will override the global setting. If there is the possibility of asymmetric flows at a site, the recommendation is to enable this at a site or policy level and not globally. For

conduit to conduit TCP flows, sequence window check is ignored. The recommendation is to enable this at both end sites.

Global Path Bandwidth Testing Settings

Global Path Bandwidth Testing Settings allows for the configuration of global parameters that impacts path bandwidth testing operations on individual SD-WAN Edge.

Path Bandwidth Test Time (ms): The maximum length of time that each packet bandwidth burst test will last in ms. The minimum value is 20ms and maximum value is 10000ms. For each bandwidth test, if the configured time is 5 seconds or less, the test will be run 10 times to get the final test result. If the configured time is greater than 5 seconds, the test will be run 5 times to get the final test result.

WARNING: Setting test time to be more than 200ms may cause path to go bad and impact normal user traffic.

Cloud Security Settings

Cloud Security Settings allows for the configuration of global parameters that impact the operation of the Cloud service of the entire SD-WAN Edge network. These settings may impact the way individual Site settings are interpreted or applied.

Network Encryption Mode: Defines the algorithm used for all encrypted Paths in the SD-WAN Edge network. This setting does not apply to non-encrypted Paths. When changing this setting, the **Secure Key** for Sites in the SD-WAN Edge network may be modified.

- When changing to AES 128-Bit, keys longer than 16 characters will be truncated to 16 characters.
- When changing to AES 256-Bit, keys shorter than 16 characters will be regenerated as 32 characters.

Enable Encryption Key Rotation: If enabled, Encryption Keys will be regenerated for every Conduit with encryption enabled using an Elliptic Curve Diffie-Hellman key exchange at intervals of 10 - 15 minutes.

Enable Extended Packet Encryption Header: If enabled, a 16 byte, randomly seeded counter will be prepended to the beginning of every encrypted message. When encrypted, this counter will serve as a random Initialization Vector, deterministic only with the encryption key. This will randomize the output of encryption, providing strong message indistinguishability. Note, when enabled, this option will increase packet overhead by 16 bytes.

Enable Extended Packet Authentication Trailer: If enabled, an authentication code will be appended to the end of every encrypted message. This trailer allows for the verification that packets are not modified in transit. Note, when enabled, this option will increase packet overhead

Extended Packet Authentication Trailer Type: The type trailer to include in encrypted messages

WARNING: Using SHA-256 may significantly impact network performance.

- 32-Bit Checksum: A 4 byte value calculated by computing the ones-complement checksum of the encrypted packet's contents.
- SHA-256: A 16 byte value calculated using SHA-256 over the encrypted packets contents.

Cloud Services

This allows for the settings of a service provider's cloud server. It is normally provided as a json file and imported by using the "Import Cloud Config..." button.

- Service Name - This is the name of the cloud service.
- Subscriber - This is the subscriber to the service provider's cloud service.
- SD-WAN Edge Network ID - This is the subscriber's SD-WAN Edge network ID.
- IP/Domain - This is the public IP or domain of the cloud service.
- Port 1 / 2 - These are the UDP public port numbers of the cloud service for data traffic.
- Mgt IP/Domain - This is the management IP or domain of the cloud service. It is also used for REST API communication. If it is not specified, the cloud service IP/Domain is used.
- Mgt Port - This is the management port of the cloud service.
- Network BW Limit - This is the bandwidth(Kbps) limit of the network.
- Per Site BW Limit - This is the bandwidth(Kbps) limit of each site's Cloud Service and applies to each of its egress and ingress separately.
- Max Number of Sites - The limit of sites that can connect to this Cloud Service.
- Cloud Service BW Limit - The total bandwidth(Kbps) limit for the cloud service in the network.
- Service Provider Provided Shared Key - This is the pre-shared key provided by the service provider for Rest API authentication.
- Subscriber Generated Secure Key - This is the subscriber secure hexadecimal key used for encryption and membership verification in the SD-WAN Edge network.

Routing Domains

Routing domains are networked systems that include a set of routers that are used to segment network traffic. New Sites are automatically associated with the default Routing Domain.

Application Categories

Application Categories provide a list of Oracle predefined categories which the user can add, delete, or edit. The application categories are used on the application dashboard to view top categories from a usage perspective. The user can add new Categories to be used in User-Defined Applications. The column In Oracle Preset shows if the Category is Pre-Defined (checked) or User-Defined. An Application Category can be used as a match criteria in Application Policy, which effectively includes all Applications that have this Category attribute as the match.

User Defined Applications

Please see Oracle Defined Applications

Oracle Defined Applications

An Application has a set of one or more match criteria which are described below. As new flows arrive, they can be tagged as belonging to a specific application. Only the

first match will apply to any given flow, so applications should be ordered in a way that the most specific, and/or most desired match is higher in the list (lower index) so that it matches before wider catch-all matches. The Oracle pre-defined list is structured in this manner.

Once a packet has been classified, the application identifier can be used either on a rule or firewall filter as a possible match criteria to handle this type of traffic. User Defined Applications provide the ability for the user to define custom applications. The user defined applications take precedence over the Oracle Pre-defined application list.

Applications are a convenient way to manage large complex combinations of match criteria, or managing large numbers of rules, or policies where match criteria may change. These changes can then be made to the application without having to go find each rule or policy where they may have been used.

The options include:

- **Priority:** The order/precedence in which policy will be applied, lower numbered policies are applied first.
- **Name:** The customer defined name for the application.
- **Category:** The Application Category this application belongs to. If no Category is selected, the application is assigned to the Other category.
- **Enable:** Check to enable the application as a match criteria for Application Policies, Rules, and Firewall Policies.
- **Classification:** When an application policy steers traffic to a conduit the classification option provides the pre-defined QoS rules associated with the Application. These will map to the standard Oracle SD-WAN Classes and Rules - Real-Time, Interactive or Bulk as shown in the following table. The priority level are straightforward with P1 being the highest priority in a Class and serviced first from a traffic perspective.

Application Classification	Default Rule Used	Class of Service
realtime_p1	Default_EF	Real-Time 10
realtime_p2	Default_UDP	Real-Time 10
interactive_p1	Default_ICMP	Interactive 11
interactive_p2	Default_ssh	Interactive 12
interactive_p3	Default_HTTP	Interactive 14
interactive_p4	Default_Telnet	Interactive 12
bulk_p1	Default_CIFS	Bulk 15
bulk_p2	Default_FTP	Bulk 16

- **Response Time Normal:** The application Dashboard provides the user with the Health of an application. The health is calculated when the Probing Interval is defined (Basic timeframe seconds) and is compared to default values for Normal and Warning. When defining an application, the user can use the default values or configure these values. These are applied at the global level. when the application is used in an application policy and applied to multiple sites this response time is applied for the application to all source sites defined. Additionally, there is the option to change these values at the site level if required. In the configuration go to the site and select Basic Settings. If configured at a site level the values are added together for the site.
- **Response Time Warning:** The application Dashboard provides the user with the Health of an application. The health is calculated when the Probing Interval is defined (Basic timeframe seconds) and is compared to default values for Normal and Warning. When defining an application, the user can use the default values or configure these values.

These are applied at the global level. when the application is used in an application policy and applied to multiple sites this response time is applied for the application to all source sites defined. Additionally, there is the option to change these values at the site level if required. In the configuration go to the site and select Basic Settings. If configured at a site level the values are added together for the site.

- **Probing Interval:** To probe an application for health checking purposes define a timeframe for the probe to be sent. The system will establish a TCP session with the domain and calculate the response time and then close the TCP session. The Health of an application is displayed on the application Dashboard.
- **Application Match Criteria:** Allows the user to define a method to match an application, either a 5 tuple or a Domain Name. When Domain Name is selected as Match, the user is required to enter a domain name which will then be matched by either DNS Proxy or DNS snooping.
- **Port –** When only 1 port number is specified, this port must match either the source and destination port in the packet
- **Port –** When one port range is specified, i.e 20-21, either the source or destination port must fall into this range.
- **Ports –** When 2 port numbers are specified, each port number must match at least the source port and the other match the destination port of in the packet.
- **Ports –** When 2 port ranges are specified, source port in the packet must fall in one range and the destination port in the packet must fall in the other range.
- **Network IP Address 1 –** When only 1 IP address and mask is specified, this address needs match either the source and destination IP address in the packet.
- **Network IP address 2 –** When both IP addresses and masks are specified, both addresses must match source and destination IP addresses in the packet.
- **DSCP match a specific DSCP value**

Note: If a single port or network IP address is defined the system will check the source and destination for a match:

The user also has the ability to edit, delete, or clone application if required.

Application Policies

An application policy provides the user the ability to select a defined (Oracle defined or user defined) application (or category) and steer the application to a Oracle service. This provides the ability for users to steer certain applications to the local Internet service, while hair-pinning others as needed. The configuration of a policy requires certain properties which will be defined below in detail. Once the policy is defined and the configuration activated the SD-WAN Edge will use DNS snooping (disabled by default) to match the Domain Name defined for the application matching. Once a match is found the application will be steered to the defined service as long as the service is available. If the service is down the routing table is used to forward the packets to their destination.

Application policies - allows the user to view defined policies with configured attributes. This view just displays based on the defined application name, not the underlying application. It also provides the destination service and classification for the application name.

Configuration Properties for Adding an Application Policy:

- **Priority:** The order/precedence in which policies will be applied, low numbered policies applied first.
- **Name(application):** The name of the application policy which will be used to display the policy.
- **Enable:** When selected the system will look to match on the parameters define for the application - 5 tuple, or dns matching domain name. If a domain name is defined and this is the first policy configured DNS snooping will be enabled for all SD-WAN Edge the policy applies to.
- **Dest Site:** This is the destination site for the policy, the destination site could be a site the application is sent to. The dest site could be a site name and combined with a service to achieve hair-pinning the INTERNET service at the SD-WAN Controller as an example. The destination site could also be the local option, where the user defines the local INTERNET service which must be available at the site to steer traffic to.
- **Dest Service Type:** The service type is an available Oracle service. The service must be available at the (dest) site otherwise the user will receive a warning during the configuration process, this can be an issue when selecting a site group for the dest site. Available service types will include: INTERNET , INTRANET.
- **Dest Service Name:** If a service is enabled at a site and multiple services can be supported then each service must have a unique name. Multiple service names can be listed for the user to choose from. The user must select the correct service name for application steering to work properly.
- **Classification:** When an application policy steers traffic to a conduit the classification option provides the pre-defined QOS rules associated with the application. The default Application Classification can be used or the user can select from a pre-defined list of Classifications to override the value defined for the Application. These will map to the standard Oracle classes and rules - Real-time, Interactive or Bulk.
- **Application Category Match:** Each Pre-Defined and User-Defined application map to a category. If preferred the user can select a category (group of applications) and steer the category to a service.
- **Application Match:** select a specific application and steer it to a service.
- **Source Network Match:** The user can define source group address prefixes and the use them as a source match for the application policy. They must first be configured under the "Source Group objects" tab.
- **Site Group Match:**This is the destination site the policy will be applied to. The options will include all client sites, SD-WAN Controller sites or used defined sites. This allows the user as much flexibility as possible when assigning policies to sites. This can be accomplished at the group level or site level. The dest could also be a site name and combined with a service to achieve hair-pinning the INTERNET service at the SD-WAN Controller. The application policy is then applied to all sites in the defined site group.
- **Site Match:** This option is used when the user require a single source site match for a policy. If the user has a single site with unique services this option can be selected for the specific site. The application policy is then only applied to that site.

Site Group Objects

Site group Objects: The Site group options allows the user to group sites together for use in the application policy. By default there are default groups for - all client(branch) sites, SD-WAN Controller based sites and all sites. If the user needs any other groupings they have the ability to create them as needed.

NOTE: Global Firewall options can be configured in the Global SD-WAN Edge Network Settings section.

Firewall

Firewall allows users to configure global firewall objects, including defining Zones and Firewall Policy Templates.

NOTE: Global Firewall options can be configured in the Global **SD-WAN Edge Network Settings** section.

Zones

Zones define a logical security grouping of networks connected to the SD-WAN Edge network. Zones can be applied to Virtual Interfaces, Intranet Services, LAN GRE Tunnels and LAN IPsec Tunnels. Intranet Services automatically determine a Zone based on the configuration. Three Zones are automatically defined and always present in the SD-WAN Edge network:

- **Default_LAN_Zone:** This Zone is applied to Virtual Interfaces, Intranet Services, LAN GRE Tunnels and LAN IPsec Tunnels if no Zone is specifically configured.
- **Internet_Zone:** This Zone is applied to Internet Services that do not have a usage for a WAN Link on an untrusted interface.
- **Untrusted_Internet_Zone:** This Zone is applied to Internet Services that have at least one usage for a WAN Link on an untrusted interface.

Firewall Policy Templates

Firewall Policy Templates can be used to simplify the Firewall configuration for similar sites within the SD-WAN Edge network or for all Sites simultaneously. Each Site can have zero or more Templates applied allowing the Site to share Firewall roles in the network. Under the Global SD-WAN Edge Network Settings, a single Template can be applied for all Sites simultaneously.

Policies from the Templates will be applied at each Site in the following order:

- Pre-Policies from Templates configured in Firewall Settings according to the order of the templates.
- Pre-Policies from the Global Template configured in SD-WAN Edge Network Settings.
- Policies configured at the Site.
- Policies automatically created to support NAT or Port Forwarding policies for the Site.
- Post-Policies from Templates configured in Firewall Settings according to the order of the templates.
- Post-Policies from the Global Template configured in SD-WAN Edge Network Settings.

Network Objects

Named groups of network elements that allow network administrators to more efficiently manage network configurations.

DHCP Option Sets

DHCP Option sets are a group of DHCP Options or Parameters that can be applied to individual IP Address ranges or a single host.

Options

These are different options that can be configured and sent to DHCP clients.

- Option Name : Select the DHCP option that needs to be configured.
- Option Number : Enter the option number(224 - 254) for Custom option. This field is pre-configured for well known options.
- Data Type : Select the data type for the value field for Custom option. This field is pre-configured for well known options
- Value: Enter the value for the selected option.

QoS Policies

A QoS Policy defines a global set of Conduit, Cloud Service, Internet and Intranet QoS Policies that can be applied in the SD-WAN Edge network. In addition to these QoS Policies, a Site may override or add to definitions contained in the QoS Policy.

Conduit QoS Policies

A Conduit QoS Policy defines a global set of Classes and Rules that can be applied to any Conduit Service in the SD-WAN Edge network. The Conduit QoS Policy allows the user to define a Conduit's **Rules** and **Classes** and then apply them in the Conduit Service. This allows the Classes and Rules to be declared and audited in one central location. In addition to the Rules defined in the Conduit QoS Policy, a Conduit Service may add to the definitions contained in the Conduit QoS Policy.

Classes (Realtime, Interactive, Bulk)

These are tools that the user can employ to classify a specific type of traffic on the Conduit and then apply Rules as to how that traffic is handled. Conduit traffic is scheduled according to its Class type and parameters. Traffic is assigned to a specific class using the Class Identifier parameter within the Rules. Each Conduit Service can have up to sixteen Classes of Service. Class Identifier does not necessarily imply scheduling priority: although an intuitive convention may be to have Class priority decrease with increasing number, e.g., Class 0 has highest priority and Class 9 has lowest priority, though this is not necessary. When a Conduit QoS Policy is applied to Conduit Service, the Classes can only be edited from the scope of the Conduit QoS Policy. If there is a need to edit Classes at the Conduit Service when a Conduit QoS Policy is applied, then the Unlink Classes from QoS Policy button can be checked in the Basic Settings of Conduit Service at Connections tab. Please see more details for this button at that section

The sixteen Classes of Services supported the SD-WAN Edge network are condensed to four Classes at the path level (in descending order of priority).

The scheduler allocates the highest priority to the Reserved Class and the lowest priority to the Bulk Class. The Reserved Class is not visible to the user as it is reserved for SD-WAN Edge use only. The remaining three Classes can be mapped to any of the sixteen Classes of Service, which are configured by the user.

Initial and **Sustained** columns are used to describe the parameters for Realtime and Interactive traffic.

The **Initial Period** is the duration in milliseconds that the Class will apply the **Initial Rate** for the Realtime Class flows or the **Initial Share** for the Interactive Class flows. The **Sustained Rate** and **Sustained Share** are used to limit the flow after the Initial Period has ended.

The SD-WAN Edge Class supports three types:

- **Realtime Class**

Best used for low latency, low bandwidth, time-sensitive traffic. Applications that are time sensitive but don't really need high bandwidth, such as voice over IP networks, can be categorized as Realtime. These applications are very sensitive to latency and jitter, but may tolerate some loss. Sometimes it is better to lose a few packets but not so many that it causes distortion. Realtime Classes provide a per-packet drop policy if the Conduit Class queue depth exceeds an estimated queue time. Small Packet Max Delay and Large Packet Max Delay parameters are used to configure the queue depth.

The Guaranteed Rate is based on the parameters set for the Realtime Class. The scheduler guarantees the **Initial Rate** and the Sustained Rate configured by the user. The **Initial Rate** determines how fast the packets can get out of the queue in a given period of time, called **Initial Period**. After the Initial Period is over, the **Sustained Rate** determines the rate at which the packets leave the Conduit. Typically initial rate and sustained rates are set to 50% of the Conduit bandwidth.

When in contention, Realtime Class will receive guaranteed rate plus a small percentage of the available bandwidth, which is shared with the remaining two Classes; Interactive and Bulk.

- **Interactive Class**

Best used for interactive traffic with low to medium latency requirements and low to medium bandwidth requirements. These applications typically have a server-client relationship; they involve human input in the form of mouse clicks or cursor moves from the client side and display graphics sent from the server to the client. Although client to server communication may not need high bandwidth, it is sensitive to loss and latency. Similarly, communication in the direction of server to client may not be sensitive to loss but does need high bandwidth to transfer graphical information. Examples include: Interactive Video, Remote Desktop, SSH, HTTPS, CICS, SQL, and VNC. Interactive Classes provide a per-packet drop policy if the queue depth exceeds a user configured byte count threshold, and the estimated time a packet will be pending to the Conduit in its Class queue exceeds a user configured time duration (in milliseconds). Small Packet Max Delay and Large Packet Max Delay parameters are used to configure the queue depth.

The Sustained Share (%) (m2) bandwidth remaining after the Realtime traffic has been serviced is available for Interactive Class to be used on a fair share basis. In order to service Interactive Class packets that are starved due to Realtime Class, the parameter **Initial Share (%)** (m1) determines the rate at which these packets will be serviced quickly during a given time, called the **Initial period** (x1). Typically, Initial Period (x1) is set at 20ms. The Sustained Share (m2), determines the rate at which these packets are serviced after the initial period is complete.

- **Bulk Class**

Best used for high bandwidth but high-latency tolerant traffic. Applications that handle file transfer and need high bandwidth are categorized as Bulk Class. Such applications are not very sensitive to loss or latency. Typically TCP will retransmit lost packets, but this will also cause too many retransmissions, thereby affecting application

performance. These applications involve very little human interference and are mostly handled by the systems themselves. Examples of Bulk applications include FTP, TFTP, CIFS, and rsync. Bulk Classes provide a per-packet drop policy if the Conduit Class queue depth exceeds an estimated queue byte count. The Delay Min Depth parameter can be used to configure the queue depth.

The Sustaining Share (%) bandwidth remaining after Realtime and Interactive has been serviced is available for Bulk Class to be used on a fair share basis. These packets get serviced last and they do not receive initial share percentage like Interactive Class. However, this Class does share the remaining bandwidth with the Interactive Class on a fair share basis. The parameter Bulk Share (m2, in percentage) determines the remaining Conduit bandwidth the Bulk Class will receive. Typically, Interactive Class gets a higher share than bulk.

NOTE: When not in contention, the Classes will be serviced at Conduit Rate.

Conduit/Conduit QoS Policy/Cloud Service QoS Policy Rules

Rules are shown as a list view on this screen, organized by their relative **Order** and shown with their match criteria in the table. Rules are checked and matched for the current service in the order shown in the table.

Rules that share an **MOS Groups Name** may be monitored collectively in Reports or Graphs, if **Track Performance** option is enabled.

Note: Use **MOS Groups** section to define new **MOS Group Name**

Clicking the **Clone** button will insert a copy of the selected rule at the end of the list. After the rule shows and all the information are added to the rule, when apply the changes, the table will be renumbered back to hundred numbers, in the same order. For example if the order number is 100 and 200 and a new rule is added as order number 150, after the apply the order numbers will be: 100, 200, 300. The new rule with order 150 is now 200 and the order 200 becomes order 300.

Once the rule is defined for matching criteria the user can set rule specific properties by selecting the () option.

The rule options are consistent for any rule defined in the system. Once open the user can set the following:

Initialize Properties Using Protocol

- WAN General
- WAN Ingress
- WAN Egress
- Deep Packet Inspection

These options are describe in the following section and are set for each rule defined. The user only has to select options that pertain to the specific rule.

Initialize Properties Using Protocol button will fill the Rule properties using recommended settings for this protocol after a protocol is selected.

WAN General

The WAN General tab provides the user with the ability to configure general operations on the flows matching this rule. These include the following types of operations:

- **Load balancing:** Traffic for the flow will be balanced across multiple paths for this service. Aggregate all the paths for the flows. With this option configured, packets are

sent across the best path until it is completely used. The remaining packets are then sent across the next best path.

- **Duplicate paths:** Traffic for the flow will be duplicated across multiple paths for this flow to increase reliability.
- **Persistent paths:** Traffic for the flow will remain on the same path. Maintain the same path for the flows if possible. This only changes when the path is not available.
- **Preferred WAN Link:** Traffic for the flow will prefer paths using this WAN link. This only applies when transmit mode is set to persistent path and for rules specific to a conduit.
- **Persistent Impedance(ms):** Traffic for the flow will stay on one path until the wait time is longer than the configured value. This only applies when transmit mode is set to persistent path.
- **Override services:** Traffic for the flow will override to a different service. In the case of a Conduit, it could override to Intranet, Internet, Pass-through, or Discard. For an Intranet Service, it can override to Internet, Pass-through, or Discard. For an Internet Service, it can override to Intranet, Pass-through or Discard. This feature allows you to select the destination service that the flows should go to, if enabled. In other words, this allows you to drop traffic out of the Conduit.
- **Retransmit Lost Packets:** The SD-WAN Edge will re-transmit any frames lost the in the cloud.
- **TCP Termination:** Traffic for this flow will be TCP terminated locally to improve throughput, reducing the round-trip times for acknowledgement packets. This functionality allows you to extend the end station TCP windows for high latency and high bandwidth networks.
- **WAN Optimization:** Traffic for this flow will be cached and De-Duplicated locally to improve throughput, reducing the amount of traffic sent over the WAN.
- **Header Compression:** Headers on this flow will be compressed to improve throughput. Support for IP,TCP and UDP as well as GRE frames can be enabled.
- **Packet Aggregation:** Small packets on this flow will be aggregated together into larger packets to improve throughput and reduce the impact of the headers on the bandwidth usage.
- **Track Performance:** If enabled, performance of a rule over time will be recorded in a session DB. Recorded attributes are loss, latency, jitter and bandwidth used.

Conduit/Advanced Settings - Conduit Class Policing Threshold Settings

In general the default settings should be accepted for these parameters but under some circumstances it may be appropriate to use these tuning settings.

These thresholds indicate the maximum queue depth of packets and data pending for a class within a conduit.

Entry should be set as the maximum latency difference anticipated between WAN paths with additional allocation + additional margin for occasional bursts.

- **Enable Policing Action** When checked, policing action will be taken when traffic passes threshold. Default is disabled.
- **Entry (Higher) Threshold in ms** When packet and latency exceeds this threshold, the classes will have more aggressive policing until the congestion is mitigated. The default value is 200ms. Set to 0ms to disable events raised with policing state transition.

- **Exit (Lower) Threshold percentage** The percentage of the entry threshold below which the policy will no longer be enforced. The default value is 20%
- **Average Packet Size in bytes** An estimate of the anticipated packet sizes to be use for policing. For voice intensive use cases, this should be lower. For bulk traffic use cases it should be set high. The default value is 750 bytes.

Why would these ever need to be set:

If you have eligibility disabled on WAN links that have substantial capacity relative to the conduit.

If you have very high latency WAN links, such as satellite combined with low latency network where the difference is above the default setting. For example: Satellite path is 450ms combined with a 20ms wireline path. In this case the max threshold should be set to $450-20=430$, plus an addition 40 ms for bursts, resulting in 470ms.

WAN Ingress

The WAN Ingress section provides the user with the ability to configure WAN Ingress behavior for this rule on the matching flow data. These properties are related to setting or reassigning the Class of WAN Ingress packets and controlling duplication and dropping packets due to queue depth values.

- **Class:** The Class that is to service traffic flows that match this Rule. The default value is Class 9.
- **Large Packet Size:** Packets destined for this Class which are larger than or equal to this size will follow large packet drop policy. Packets which are smaller than this size will follow small packet drop policy. If this size is set to 0, all packets will be treated as small packets.
- **Drop limit:** The maximum amount of estimated time that packets smaller than the Large Packet Size will have to wait in the Class scheduler. If the estimated time exceeds this threshold, the packet will be discarded and statistics will be counted. Not valid for Bulk Classes.
- **Drop Depth:** If the queue depth exceeds this threshold, the packet will be discarded and statistics will be counted.

NOTE: Either value Drop Limit or Drop Depth will allow a frame to be dropped.

- **Enable RED:** Random Early Detection (RED) will help promote the fair sharing of Class resources by judiciously discarding packets as worsening congestion is encountered. Works best with protocols/applications that back off when they detect loss, like TCP.
- **Reassign Size:** This is used to define a packet length. When exceeded, a flow will be reassigned to a different Class defined by the Reassign Class id.
- **Reassign Class:** This Class is used when the packets in a given flow exceed a defined length. If the default option is selected, packets will not be assigned to an alternate class based on packet size, and will continue to be mapped to the class specified in the "General" section.
- **Duplicate Packets:** Values used to determine when to not duplicate a flow
- **TCP Standalone ACK Class:** Allows the responding TCP Standalone ACK's to be mapped to a higher priority Class when a large file transfer is taking place. Used to improve performance of a file transfer. If the default option is selected, TCP Standalone ACK's will continue to be mapped to the class specified in the "General" section.

WAN Egress

The WAN Egress section provides the user with the ability to set rule properties controlling operations to packets received via the WAN Egress - DSCP tagging and packet resequencing operations.

- **Re-sequence Packets:** Simply means that it puts the packets back in order at the destination.
- **Re-sequence Hold Time:** Amount of time a packet can be held for re-sequencing before being sent to the LAN.
- **Discard Late Re-sequence Packets:** If an out-of-order packet arrives late and the dependent packet has already been sent to the LAN, then discard it.
- **DSCP Tag:** Remark packets in a given flow with a new DSCP Tag.

Deep Packet Inspection

The Deep Packet Inspection section provides the user with the ability to configure rule properties related to operations based on the contents of the matching packets.

Enable Passive FTP Detection: If enabled, this parameter will make processing decisions based upon user data. The rule will learn the port used for FTP data transfer and apply the rule properties to the learned port.

IPsec Properties

The IPsec properties section allows you to enable IPsec protection for data in the Conduit. If enabled, an IPsec tunnel is established across the Conduit before data can flow.

- **Secure Conduit User Data with IPsec:** If enabled, user data transmitted using the Conduit is secured using an IPsec tunnel.
- **Tunnel Mode:** The available IPsec protocols you can choose from.
- **ESP:** Data is encapsulated and encrypted.
- **ESP+Auth:** Data is encapsulated, encrypted, and validated with an HMAC.
- **AH:** Data is validated with an HMAC.
- **Encryption Mode:** The encryption algorithm used when ESP is enabled.
- **Hash Algorithm:** The hash algorithm used to generate an HMAC.
- **Lifetime (s):** Your preferred duration, in seconds, for an IPsec security association to exist. Enter 0 for unlimited.

Advanced Settings

This section allows you to specify a bandwidth threshold in terms of a percentage of the total egress permitted rates of regular WAN links. If the available bandwidth provided by the regular WAN links in the conduit falls below this bandwidth threshold, on-demand standby WAN links in the conduit will be activated to supplement bandwidth.

Dynamic Conduit QoS Policies

A Dynamic Conduit QoS Policy defines a global set of Classes, Rules and Dynamic Conduit properties that are applied to all Dynamic Conduits in the SD-WAN Edge network.

Dynamic Conduit Properties

A Dynamic Conduit QoS Policy is created automatically in the Configuration, when Dynamic Conduits is configured in the system. Currently the values defined are used for all Dynamic Conduits defined in the system

- **Basic Settings:** Properties used to establish/remove Dynamic Conduits. Value descriptions are provided to the user if they hover over the specific command.
- **Classes:** Same values used for static Conduits.
- **Rules:** Same values used for static Conduits.

Internet QoS Policies

An Internet QoS Policy defines a global set of Rules that can be applied to any Internet Service in the Adaptive Private Network Network. In addition to the Rules defined in the Internet QoS Policy, an Internet Service may override or add to definitions contained in the QoS Policy.

Internet Service Rules

Rules are shown as a list view on this screen, organized by their relative **Order** and shown with their match criteria in the table. Rules are checked and matched for the current service in the order shown in the table. Rules are typically used to tie services to a specific WAN Link. Rule matching options are the same as previously defined.

Mode: defined what action will be taken for the defined rule.

Override Service: which service to override the flow to.

Enable Passive FTP Detection: If enabled, this parameter will make processing decisions based upon user data. The rule will learn the port used for FTP data transfer and apply the rule properties to the learned port.

Clicking the **Clone** button will insert a copy of the selected rule directly above it.

Intranet QoS Policies

An Intranet QoS Policy defines a global set of Rules that can be applied to any Intranet Service in the Adaptive Private Network. In addition to the Rules defined in the Intranet QoS Policy, an Intranet Service may override or add to definitions contained in the QoS Policy. Rules are typically used to tie services to a specific WAN Link. Rule matching options are the same as previously defined.

Override Service: which service to override the flow to.

Intranet Service Rules

Rules are shown as a list view on this screen, organized by their relative **Order** and shown with their match criteria in the table. Rules are checked and matched for the current service in the order shown in the table.

Override Service: which service to override the flow to.

Enable Passive FTP Detection: If enabled, this parameter will make processing decisions based upon user data. The rule will learn the port used for FTP data transfer and apply the rule properties to the learned port.

Clicking the **Clone** button will insert a copy of the selected rule directly above it.

Cloud Service QoS Policies

A Cloud Service QoS Policy defines a global set of Classes, Rules and Service properties that are applied to all Cloud Services in the SD-WAN Edge network. It is automatically created if none exists in the configuration.

Cloud Service Properties

Currently the values defined are used for all Cloud Services defined in the system

- **Basic Settings:** Properties used to establish/remove Cloud Services. Value descriptions are provided to the user if they hover over the specific command.

- **Classes:** Same values used for static Conduits.
- **Rules:** Same values used for static Conduits.

MOS Groups

Note: In previous releases, MOS Groups were called Applications. This change in terminology occurred in version 5.2 of the SD-WAN Edge Software.

The MOS Group is a gathering of rules which can define a particular application in the network. To create a new MOS Group in the configuration, simply click the add icon from this section. Once the MOS Group is created from this scope, individual rules can be tagged as belonging to a particular MOS Group by setting the "MOS Group Name" field for that rule to the newly created MOS Group. MOS Groups for the existing Default Rules have already been created for you.

From the scope of the MOS Group section, MOS Groups have the option to "Estimate MOS". Enabling this setting will cause the Aware to calculate a MOS Score passively for existing Traffic that passes through the Conduit. This MOS Score is a quality assessment of the MOS Group traffic judged as if it were a VOIP phone call. This statistical data is only visible from the SD-WAN Aware.

Note: **Track Performance** option needs to be enabled in order to "Estimate MOS" for the "MOS Groups Name" in the Rules.

Autopath Groups

Autopath Groups automatically generate Paths between WAN Links using preset parameters. When a pair of local and remote WAN Links of the same Access Type (Public Internet, Private Intranet, or Private MPLS) reference the same Autopath Group, a Path is created in both directions between the links using the Autopath Group settings. By Default, Paths between Private MPLS WAN Links are only created between MPLS Queues with matching DSCP tags.

A Default Autopath Group must always exist and is denoted as <DEFAULT> in Conduit Usages.

The Configuration options for Autopath groups are:

- **IP DSCP Tagging:** Provides a tag for the external IP header of the Talari Reliable Protocol (TRP) frame.
- **Enable Encryption:** Encrypts the TRP frame.
- **Bad Loss Sensitive:** A Path may be marked as **BAD** due to loss and will incur a latency penalty in Path scoring. Disabling this option may be useful when the loss of bandwidth is intolerable.
- **Percent Loss (%):** When **Bad Loss Sensitive** is set to **Custom**, if packet loss exceeds the set percentage over the configured time, the **GOOD** Path state will change to **BAD**. The default setting uses an internal Oracle algorithm.
- **Over Time (ms):** When **Bad Loss Sensitive** is set to **Custom** and Percent Loss is set to a value other than **DEFAULT**, if packet loss exceeds the set percentage over this configured time, the Path state is marked as **BAD**.
- **Silence Period (ms):** The Path state transitions from **GOOD** to **BAD** when no packets are received within the specified amount of time.
- **Path Probation Period (ms):** The period to wait before changing the Path state from **BAD** to **GOOD**.

- **Instability Sensitive:** Latency penalties due to **BAD** state and other spikes in latency are considered in the Path scoring algorithm when this is enabled. Disabling this option may be useful when the loss of bandwidth (If Bad Loss Sensitive enabled) or latency spikes are intolerable.

There are four combinations for the Bad Loss Sensitive and Instability Sensitive settings:

- **Option 1:** When **Bad Loss Sensitive** is set to **Enable** or **Custom** and **Instability Sensitive** is enabled, a Path may be marked as **BAD** and incur a latency penalty so it is only used as a last resort. In the event multiple Paths are marked **BAD**, there is still competition among them based on regular Path scoring.
- **Option 2:** When **Bad Loss Sensitive** is set to **Enable** or **Custom** and **Instability Sensitive** is **disabled**, a Path may be marked as **BAD** and only used as a last resort, however latency spikes are not considered. In the event multiple paths are marked **BAD**, the ones with **Instability Sensitive disabled** will likely be used first.
- **Option 3:** When **Bad Loss Sensitive** is set to **Disable** and **Instability Sensitive** is **enabled**, a Path remains **GOOD** in spite of loss, however latency spikes are still considered, so that Path is only likely to be used after Paths without latency spikes are exhausted.
- **Option 4:** When **Bad Loss Sensitive** is set to **Disable** and **Instability Sensitive** is **disabled**, a Path remains **GOOD** and latency spikes are not considered, therefore the Path will likely remain in constant use.

Service Providers

Service providers are the container objects for WAN Link Templates. The intended abstraction is, to place a WAN Link Template in the scope of a particular service provider to state that that particular template defines a WAN Link provided by that particular service provider.

For example, adding a broadband WAN Link Template under a service provider named **TimeWarnerCable** implies that the wan link template will be applied to a Time Warner Cable public internet WAN Link.

Where the broadband example is not too powerful, the real value of the Service Provider is in MPLS WAN Links.

For example, adding an MPLS WAN Link Template under a service provider of "Verizon" implies that MPLS WAN Links using this template will be denoted as a Verizon MPLS WAN Link. The underlying functionality will associate all of WAN Links using that template with an autogenerated "Verizon_MPLS" autopath group, thereby removing the need for the user to configure and associate autopath groups.

WAN Link Templates

WAN Link Templates can be used to simplify the WAN Link configuration for similar sites within the SD-WAN Edge network. Each Site who shares the same WAN Link characteristics can select the WAN Link Template to be applied to its WAN Link properties.

WAN Link Template - Basic Settings

The WAN Link Template Basic Settings allow for the description of the type of the link available.

Link Type is the type of link for this template, which may be Broadband, Private Link, or MPLS.

Auto-Path Group is the group used to determine what Paths may be automatically generated between the WAN Link and remote WAN Links and what default Path settings to use.

- **<None>** indicates that no group is desired and will prevent Paths from being automatically generated to or from the WAN Link.

WAN Ingress

- **Physical Rate** is the bit rate limit of the WAN Link for the traffic traveling from the LAN into the WAN. Configuration should match the physical capacity of the WAN Link.
- **Auto Learn** indicates whether the permitted rate of the WAN Link will be automatically adjusted based on bandwidth test results. Before a valid test is completed, the physical rate will be used. No matter what the bandwidth test result is, the applied permitted rate will not exceed the physical rate.

WAN Egress

- **Physical Rate** is the bit rate limit of the WAN Link for the traffic traveling from the WAN into the LAN. The Configuration should match the physical capacity of the WAN Link purchased from the service provider.
- **Auto Learn** indicates whether the permitted rate of the WAN Link will be automatically adjusted based on bandwidth test results. Before a valid test is completed, the physical rate will be used. No matter what the bandwidth test result is, the applied permitted rate will not exceed the physical rate.

WAN Link Template - MPLS Queues

The WAN Link Template MPLS Queues allow for the definition of service queues using standard DSCP tags.

DSCP Tag: The DSCP Tag applied to the Oracle Conduit Path..

WAN Ingress Permitted Rate (Kbps): The available or allowed rate, in Kbps, for WAN Ingress traffic. The sum of WAN Ingress Permitted Kbps for all queues in a Private MPLS WAN Link may not exceed the WAN Ingress Permitted Kbps for the Private MPLS WAN Link.

WAN Egress Permitted Rate (Kbps): The available or allowed rate, in Kbps, for WAN Egress traffic. The sum of WAN Egress Permitted Kbps for all queues in a Private MPLS WAN Link may not exceed the WAN Egress Permitted Kbps for the Private MPLS WAN Link.

WAN to WAN Forwarding Groups

The WAN To WAN (W-T-W) Forwarding Group is used to allow client sites to communicate through an intermediary site with each other. In previous releases only a single W-T-W forwarding groups was allowed. In current releases this number is unrestricted. When enabled the routing tables are shared between the site with W-T-W forwarding enabled and all client site in the specific W-T-W group. Additionally when using Dynamic Conduits W-T-W forwarding must be enabled. By default all sites are in a default W-T-W forwarding group.

Site Name

As a site is added to the Configuration file the site name is automatically added. The user can open the site name and configure details for the site including the following:

- **WAN-To-WAN:** Forwarding - assign the site to a W-T-W Forwarding Group other than the default.

- **Conduit:** Define Conduit information
- **Internet Services:** Defined Internet properties for the site
- **Intranet Services:** Define Intranet properties for the site
- **Routes:** Define routes for the site.
- **WAN Links:** Define WAN Link properties for a site.

Conduit-to-Conduit Forwarding

Conduit-to-Conduit Forwarding allows the Site to act as an intermediate hop between two adjacent Sites for any Site-to-Site traffic. Unlike enable WAN-to-WAN Forwarding, this will not export any routes from one site to other sites. Conduits includes Static Conduits, Dynamic Conduits and Cloud Services.

Conduit-to-Internet/Intranet Forwarding

Conduit-to-Internet/Intranet Forwarding allows the Site to act as an intermediate hop between any Site-to-Internet or Intranet traffic. Unlike enable WAN-to-WAN Forwarding, this will not export internet/intranet routes from to other sites.

WAN-to-WAN Forwarding

WAN-to-WAN Forwarding allows the Site to act as an intermediate hop between two adjacent Sites for any Site-to-Site, Internet or Intranet traffic and to act as a mediator for Dynamic Conduits. Allows the user to add multiple groups, each groups is defined by a group name. When defining client sites the client site would then be associated with a specific group name if desired.

Release 9.0M1 Features

This chapter includes features released in 9.0M1.

Wake on LAN

SD-WAN Edge supports subnet directed broadcasts on virtual interfaces, which are used primarily to enable Wake on LAN-type packets.

Wake on LAN (WOL) is a combination of hardware and software technologies to wake up sleeping computers. It allows the administrator to remotely power up all sleeping machines so that they can receive updates. WOL is based on the principle that when the host shuts down, the NIC still receives power, and keeps listening on the network for the required packet to arrive.

To allow WOL packets to traverse your SD-WAN Network, the SD-WAN devices must be configured to allow directed broadcasts.

Configuring Directed Broadcast for Wake on LAN

To enable Directed Broadcast on a new or existing virtual interface:

1. From your SD-WAN Edge Controller, open the Configuration Editor.
2. Click on the **All Sites** section.
3. Configure a new site, or modify an existing one.
4. On **Define Interfaces**, click on **Add** or edit an existing one.
5. Click on the toggle button under the **Directed Broadcast** column.
6. Click on **Submit**.