Oracle® SD-WAN Edge

TCP Termination

Original Publication Date: Nov 1, 2019

Table of Contents

About This Document

This document discusses the concept and operation of TCP Termination within the Talari Adaptive Private Network (APN), and provides configuration commands and design recommendations to assist you with implementation.

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.

References

The following documents are available: *Talari Glossary*

Introduction to TCP Termination

Without TCP Termination enabled, a single TCP connection is established between Host A and Host B, where the two hosts reside on separate network segments across the WAN, as illustrated in Figure 1. TCP Termination provides the ability to split a single TCP connection across the WAN into three separate TCP Connections, all managed and maintained by the Talari APN, as shown in Figure 2. TCP Termination is only used for Conduit traffic.

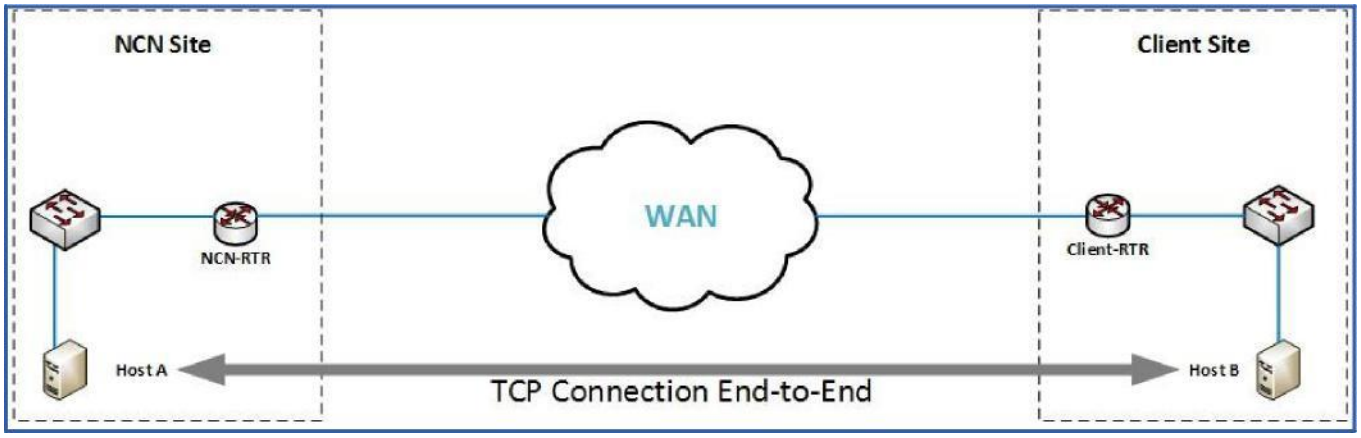

Figure 1

Figures 1 and 2 depict a typical two-site Talari APN with traffic flowing from the NCN site to a Client site. However, TCP Termination can also operate between two Client sites with an established conduit.

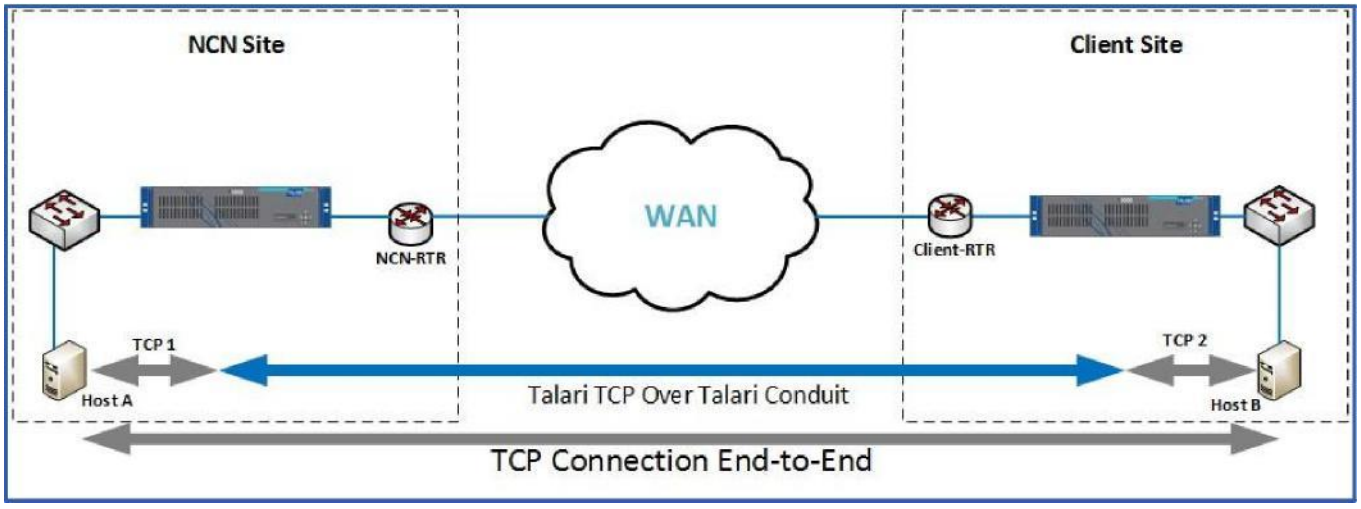

Figure 2

Figure 2 assumes that a Talari conduit between the two Talari Appliances exists, and depicts three separate connections used for TCP Termination. These connections can be defined as:

- Host A to NCN
- NCN to Client – Conduit services
- Client to Host B

Benefits of TCP Termination include increased throughput across the Conduit/WAN, significant improvement in performance when there is loss on the link and a high round trip time (RTT) across the WAN, and the ability to achieve maximum throughput through the Talari conduit while locally terminating the TCP session.

Functionality in Detail

The functionality, as described above, creates three TCP sessions. The initial TCP handshake is from **Host A** to **Host B**. As Host A communicates with Host B, the Talari Appliances monitor the TCP flow and support a modified end-to-end three way handshake, creating three separate TCP connections. Once the separate TCP sessions are established, data transfer can begin. The Talari Appliances will maintain individual TCP sessions between local Hosts, and also maintain a third TCP session across the conduit between Talari Appliances. These sessions will be established for any TCP flow that is identified as a TCP terminated flow.

For conduit traffic, a separate Talari-TCP will be used. This Talari-TCP will identify each unique flow and allow the Talari Appliances to maintain multiple sessions across the conduit. This Talari-TCP is encapsulated in the conduit and not seen by the user. The Talari Appliances also have built-in support for failure scenarios. In the event of a host or Talari failure, the TCP connections will be reset gracefully. If the conduit is down, the Talari Appliances will terminate the TCP-terminated connection and the Hosts will have to re-establish their TCP session.

The system has a dynamic capability which can be used to disable TCP Termination if system resources are getting low. In this scenario, the Talari Appliance notifies all other Talari Appliances to which it has conduits that is low on resources, and directs them to disable TCP Termination for the immediate time frame. Once system resources become

available on the Talari Appliance, the TCP Termination functionality will be re-enabled and communicated to all other Talari Appliances within the network. This function is an internal component of the TCP Termination capability and protects the system from potential catastrophic events.

Note: When transferring files using FTP or SCP with TCP Termination enabled, the reported rate of transfer is the rate between local host machine and local Talari Appliance. Since TCP Termination buffers numerous TCP packets and acknowledges incoming packets locally, the transfer rate can be much higher than the user's WAN link bandwidth.

The transfer is reported complete only when all the packets are sent to the destination and acknowledgement is received. Therefore, there may be some delay between seeing a message that the files are 100% sent and the transfer completing.

Configuring TCP Termination

TCP Termination can be enabled on the Talari in two ways: as a Conduit Default Set Rule that can be applied to any Conduit Service in the APN, or as a site-specific local Rule applied only on matching flows at the site where it is configured.

Conduit Default Set Rule

To enable TCP Termination as a **Conduit Default Set Rule**, navigate to **Global > Default Sets > Conduit Default Sets > [conduit default set]**.

Navigate to **Rules** and click **Add** (“+” icon) to open a new rule configuration menu.

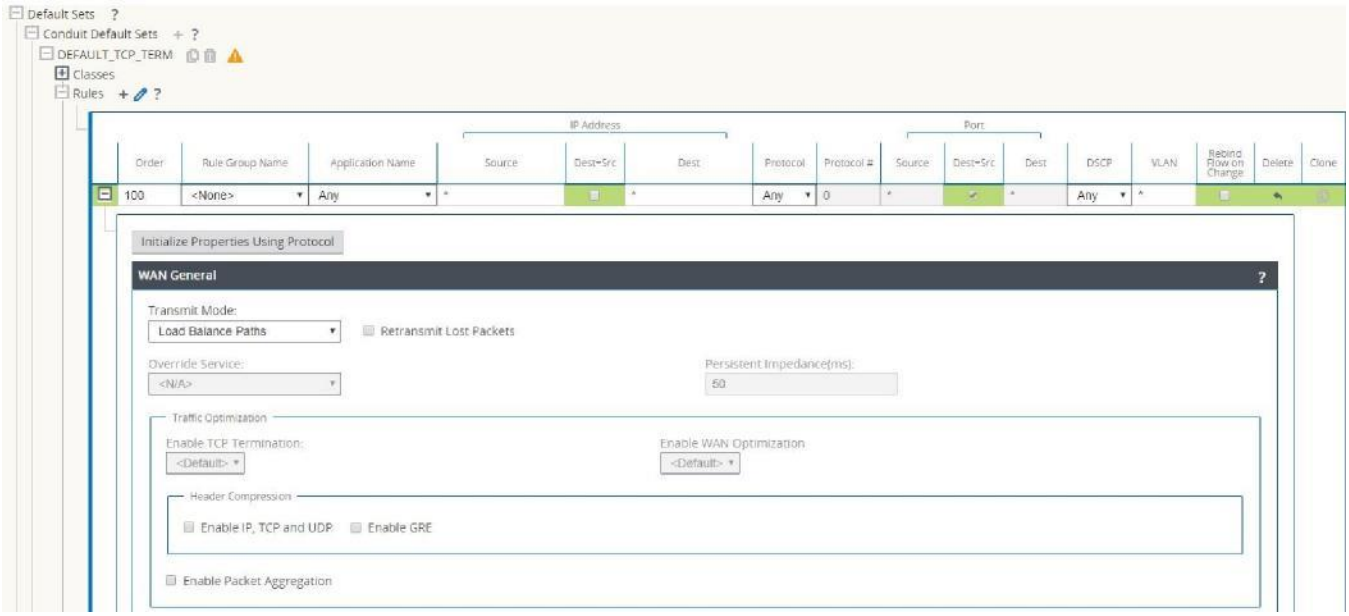

Configure a rule for the type of traffic you wish to enable TCP Termination on. Expand the rule to view the WAN General Properties. TCP Termination will be configurable only if a TCP-based protocol has been selected as the Rule Protocol.

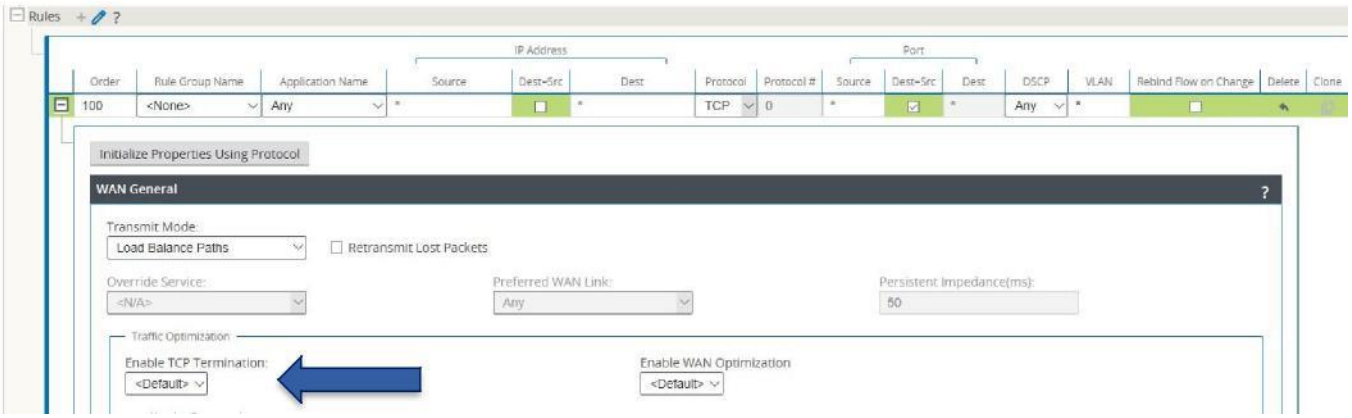

Select “On” from the “Enable TCP Termination” dropdown menu and **Click** “Apply” to confirm the new TCP Termination rule.

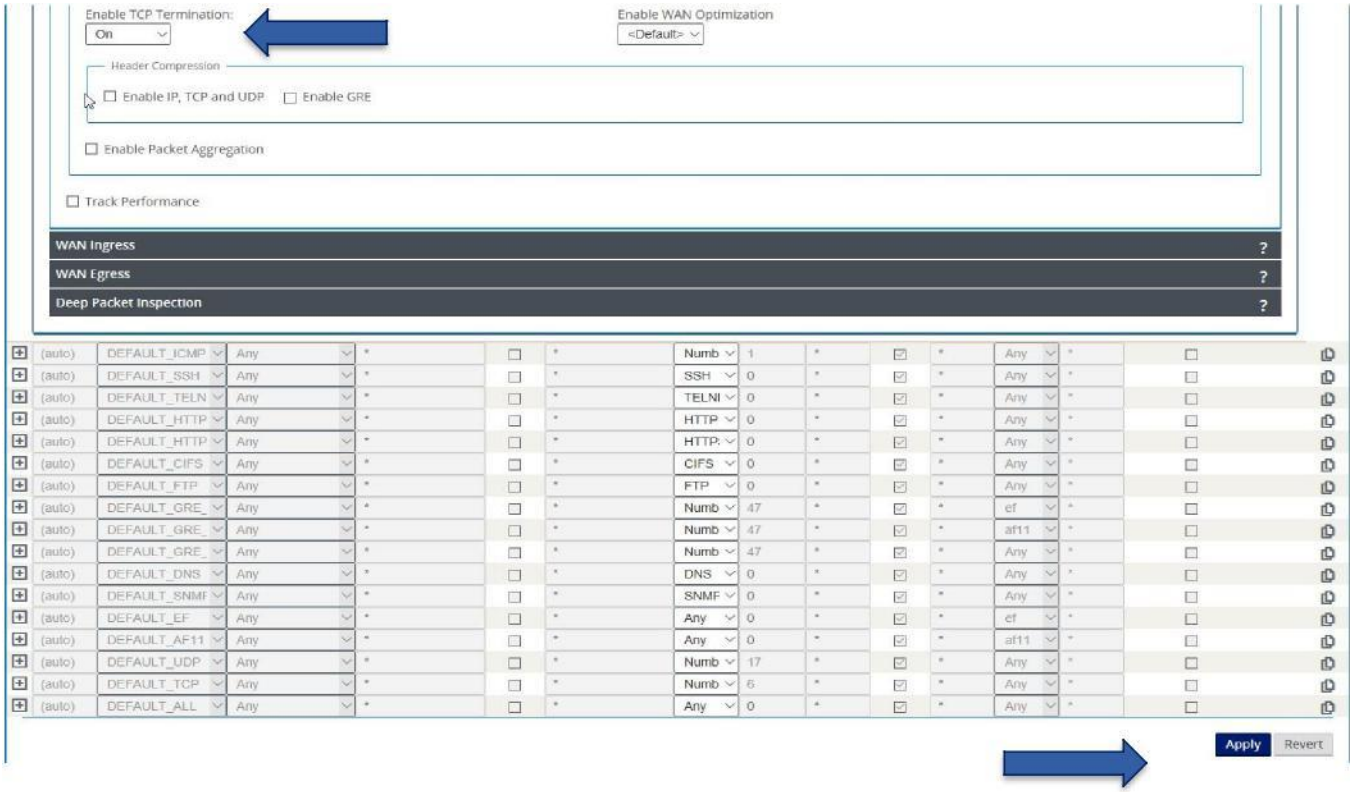

All TCP traffic that matches this Rule on this Conduit will now be TCP terminated locally to improve throughput, reducing the round-trip times for acknowledgment packets.

Site-Specific Rule

To enable TCP Termination as a site-specific local **Rule**, rather than applying TCP Termination to all Conduit flows, navigate to **Connections > [site name] > Conduits > [conduit name] > Local site > Rules** and click **Add** (“+” icon).

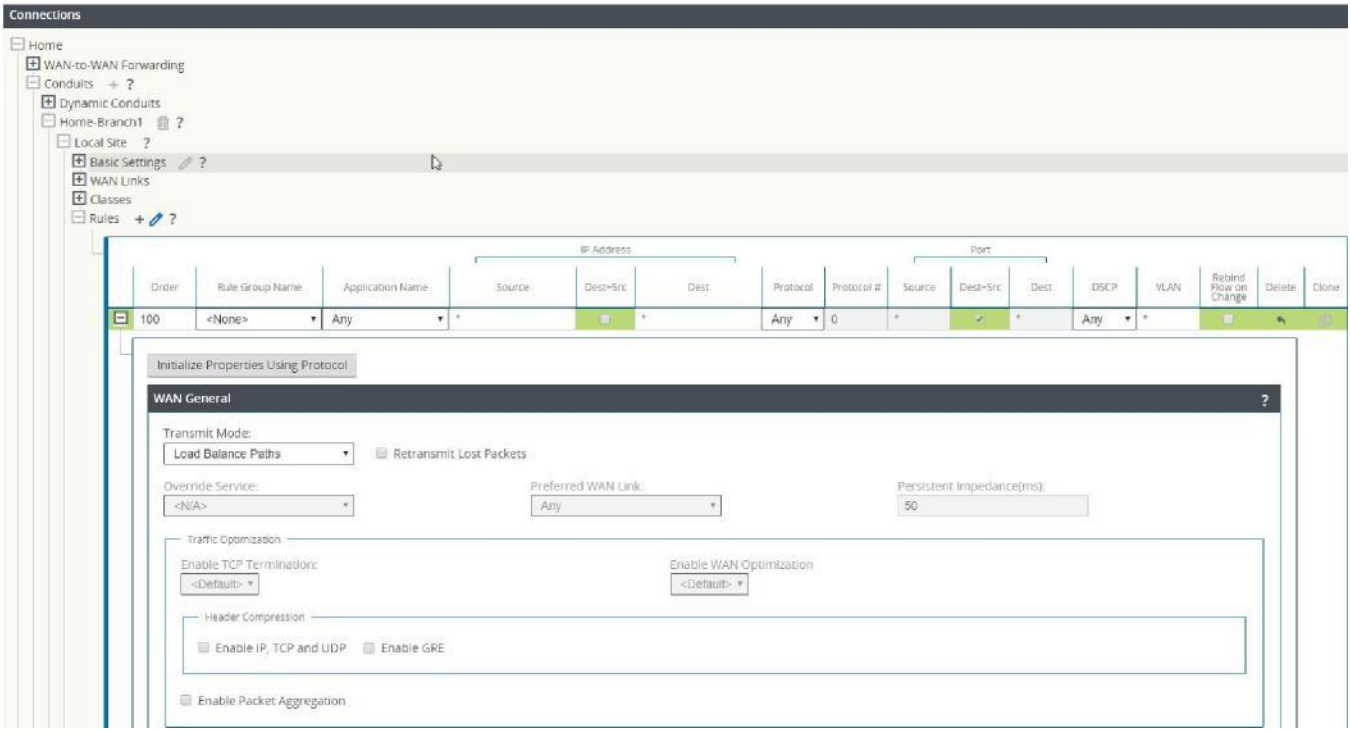

Configuration as shown above in the Conduit Default Set to enable local TCP Termination for flows matching the rule parameters.

While no other configuration is required to enable TCP Termination, there are several design considerations to be aware of when implementing TCP Termination. These are described in the next section.

Design Considerations

There are many design considerations to be aware of when implementing TCP Termination, including:

- When enabled, TCP Termination will use the maximum allowable resources per platform (defined below).
- Use Rules to guarantee TCP Termination is not starved out of the conduit. It should be used with TCP_ACK Class (see note below).
- Total TCP flow numbers are based on Inbound and Outbound flows per platform.

- Supports High Availability (HA) configuration (sessions not maintained across HA failure)

Note: If only the default bulk class is defined, TCP-terminated traffic could use all available bandwidth and starve out any potential new TCP-terminated flows. The recommendation is to define a specific rule and class for TCP-terminated traffic.

Additionally, TCP Termination is currently not supported with Riverbed implementations.

Talari TCP Termination has a limit to the number of TCP-terminated flows that may be supported based on the APN appliance used. These limits are listed below and are based on the hardware capabilities of the individual platform. Once the supply of flows has been exceeded, any new flows will not be TCP-terminated until pre-existing flows end and TCP Termination resources are freed. The platform numbers for TCP Termination are as follows:

- T750 APNA supports a maximum of 8000 TCP flows.
- E100 APNA supports a maximum of 8000 TCP flows.
- T860 APNA supports a maximum of 8000 TCP flows.
- T3000 APNA supports a maximum of 16000 TCP flows.
- T3010 APNA supports a maximum of 16000 TCP flows.
- T5000 APNA supports a maximum of 16000 TCP flows.
- T5200 APNA supports a maximum of 16000 TCP flows.

For example: The APN T750 appliance supports a total of 8000 TCP terminated flows. If a rule is defined for ssh using TCP Termination, with a set minimum resource usage of ten percent, at least 800 ssh sessions can be used for TCP Termination. The remaining TCP Termination sessions (7200) would be used for any other TCP session sourced or destined for the same conduit.

When the Talari APNA has multiple conduits defined, TCP Termination must allocate resources for each conduit. The method TCP Termination uses to allocate resources to a specific conduit is as follows:

- Determine if a conduit has TCP Termination enabled (default rule).
- Determine if there is any rule defined for TCP Termination.
- Allocate resources based on minimum allocation defined per rule for a conduit (if rule applies to two sites the min_resrc_pct is divided by two). See Note below.

- Any remaining resources are allocated on a first come first serve basis until platform resources are depleted.

Troubleshooting TCP Termination

The Talari Appliance can provide useful information regarding TCP Terminated flows. This information can be viewed from the appliance Web Console. Navigate to **Monitor > Flows,** click the “TCP Termination Table” check box, then click Refresh to display the TCP Termination Flows tables below the WAN Ingress and WAN Egress Flows. All Conduit Flows that are using TCP Termination will be displayed on this page. The page will include relevant information on a per-flow basis.

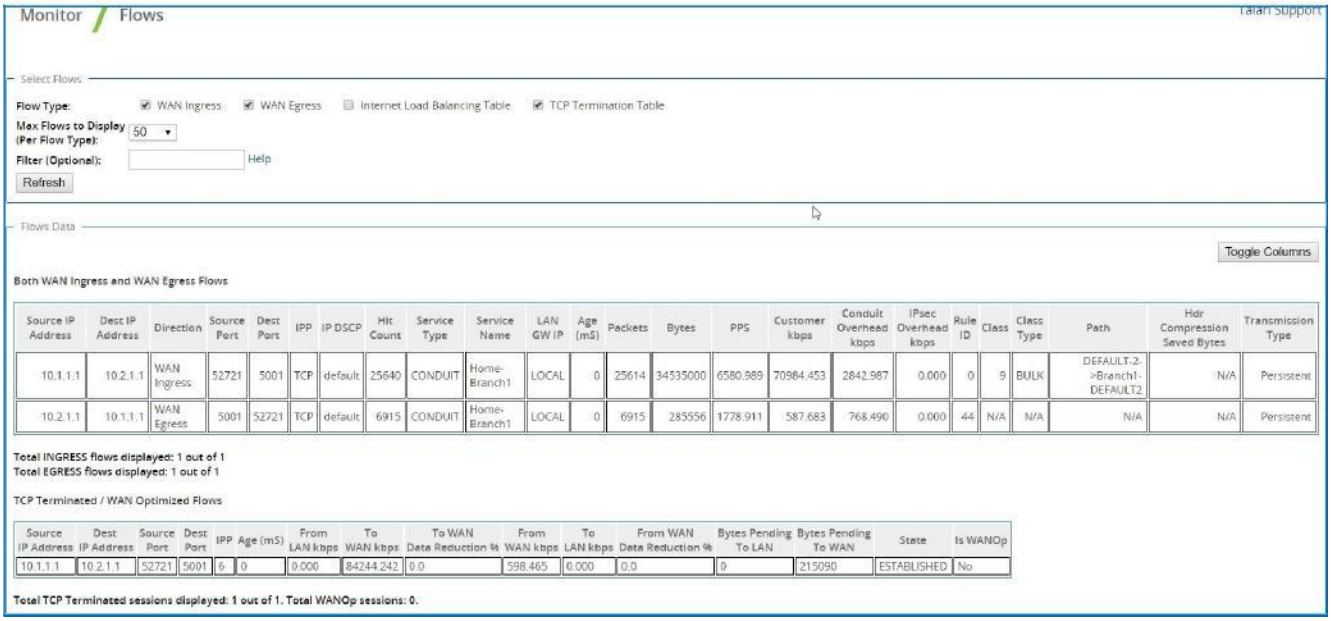

Additionally, TCP Termination will add entries into the APN_common.log file. The example shown below indicates that there has been a reset sent from the client Host A to the remote Host B notifying Host B that the connect has been reset. The local APNA instructs the remote APNA to tear down or reset the TCP connection.

tcp_do_segment@forward/tcp_input.c:1813 tp:0x29aeabfc 10.30.10.21:445 --> 10.10.10.21:49330 got reset, close the connection

These are logged to assist the user in monitoring the state of the terminated flows. The TCP-terminated flows are conduit flows only and consist solely for traffic between APN sites. This can simplify the troubleshooting process. Any other issues related to troubleshooting TCP Termination would require Talari Support personnel to assist. Prior to contacting Talari Support, it is recommended to collect a diagnostic log file from the APNAs in question, using the APNA Web console **Diagnose** pull-down menu. A diagnostic data capture tool will collect log files as well as low level debug information from the APNA and save it to a file, which can be forwarded to your Talari support representative for review.

Summary

The addition of TCP Termination to the Talari product line provides additional throughput based on the WAN link RTT and any potential circuit loss. This capability can increase throughput multiple times, depending on circuit characteristics. TCP Termination is easily configured and maintained within the Talari APNA. For additional questions, please contact your local Talari representative.