Out-of-Path Deployment

Platform: https://racks.uninets.com
Lab Name: Riverbed

Topology

Out-of-Path Deployment

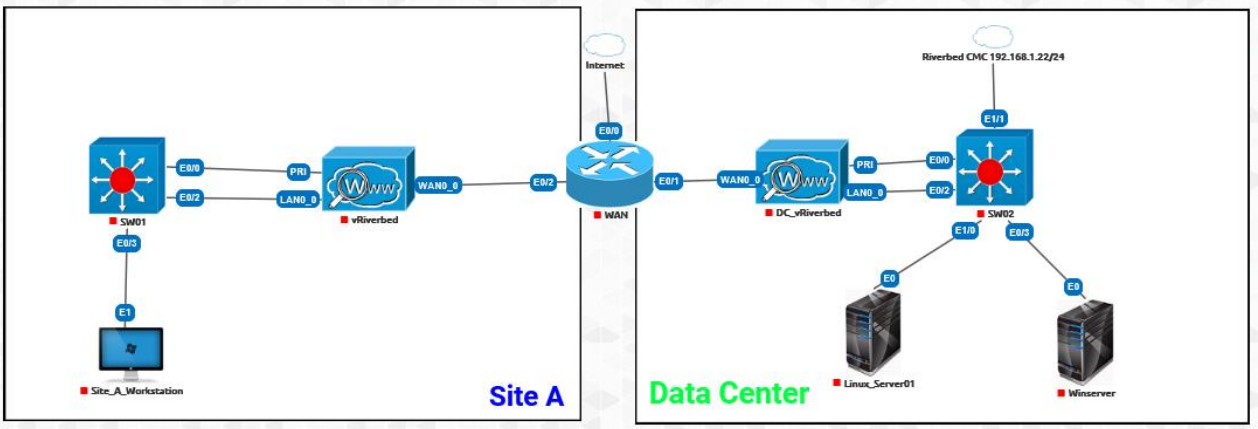

Tasks

- Connect to Datacenter Steelhead CLI via SSH
- Steelhead Initial Configuration Jumpstart
- Verifying IP Address with CLI Configuration
- Verify Time and Date Settings using CLI
- Clear Secure Vault
- Connect to the Steelhead Management Console (UI)
- Enable Datacenter Steelhead Out-of-Path Support
- Configure Branch Steelhead Fixed Target Rule
- Map a Network Drive for Optimization Tests
- Using the Peer and Current Connections Report

Explanation

In this lab scenario, you will configure the branch and datacenter Steelhead Appliances to peer using the Server-Side Out-of-Path (SSOOP) configuration.

Before proceeding, take a few moments to review the out-of-path topology diagram. A Server-side out-of-path Steelhead appliance is deployed physically out-of-path in the Datacenter with only the Primary interface being connected to the network, as shown in the out-of-path topology diagram. In the lab environment, you will be using the N21-PC in your San Francisco branch to Connect to the N29-FIL file server in the Sydney office. The Sydney office was part ofa recent acquisition and they do not have any Riverbed products yet, however you don't have the luxury of deploying the Steelhead in an in-path deployment, so you will be using an out-of-path deployment with fixed target rules. Keep in mind that this will not allow any users at the Sydney office to have optimized connections, however it.

will allow optimized access to the Sydney file server while users are slowly relocating to the Hong Kong and San Francisco office. For the purposes of this lab scenario, your N21-SH Steelhead appliance has been initially configured using the configuration jump-start wizard. Once the configuration is completed, you will initiate multiple CIFS file transfers between your Windows 7 client PC and the file server in the Sydney office while monitoring and recording the transfer times.

The SSOOP deployment is an excellent choice for proof-of-concept and lab-based Steelhead implementation since no network downtime or additional networking equipment is required for optimization to take place. Remember that Steelhead appliances are still required to form a symmetric relationship, even if the server-side datacenter Steelhead is out-of-path. In this deployment scenario, the branch Steelhead appliance is configured as an in-path appliance and uses a fixed-target.

In path rule to peer with the server-side out-of-path Steelhead appliance. A fixed-target rule will ensure peering (inner channel setup) between both Steelhead appliances. Since the server-side Steelhead appliance is out-of-path, any TCP connections from the client-side Steelhead must match the fixed-target in-path rule for optimizations to take place.

Task 1

Configuration

Connect to Datacenter Steelhead CLI via SSH You can connect to a new Steelhead appliance using the serial console (RS-232 or Virtual machine console) or using SSH. It is best practice to use the serial console for the initial configuration jumpstart and connect with SSH for subsequent administration

1. Open Putty Connection Manager and select N29 Sydney > N29-SH Steelhead (OOP) in the sidebar. In the lab environment, you will be accessing the Steelhead appliance using Preconfigured bookmarks in Putty Connection Manager. To simplify Connectivity in the lab environment, the AUX network interface has been preconfigured with an IP address of 10.99.29.20 /16. In your environment, The Steelhead will acquire an address on the Primary interface using DHCP.

2. Click inside of the black console window to activate the SSH session. You will Need to select YES to accept the risk of using a self-signed certificate with SSH.

3. Login using the username admin and password password.

Task 2

Steelhead Initial Configuration Jumpstart The next step is to use the configuration jumpstart wizard to establish basic network connectivity for the primary interface of your out-of-path Steelhead appliance. In newer Steelhead versions, you will need to disable Simplified Routing by using the

(conf t) # in-path simplified routing none command before using the

Configuration jumpstart for an out-of-path deployment. To show how to configure the out-of-path configuration manually (for example, after you've been testing a Steelhead using the in-path deployment method), we will not Activate the out-of-path configuration using the wizard.

1. Use Putty Connection Manager from the previous task to enter the following

Commands.

amnesiac > en

amnesiac # {You are now in enable mode} amnesiac

conf t

amnesiac (config) # {You are now in configuration terminal mode}

amnesiac (config) # in-path simplified routing none

Warning: By turning off Simplified Routing, VLAN connection-based

mapping (currently enabled) will not work properly. amnesiac

(config) # conf j

Riverbed Steelhead configuration wizard.

Step 1: Hostname? [amnesiac] n29-sh

Step 2: Use DHCP on primary interface? [yes] no

Step 3: Primary IP address? 10.1.29.20

Step 4: Netmask? [0.0.0.0] 255.255.255.0

Step 5: Default gateway? 10.1.29.1

Step 6: Primary DNS server? <leave blank> Step

7: Domain name? <leave blank>

Step 8: Admin password? <leave blank>

Step 9: SMTP Server? <leave blank>

Step 10: Notification Email Address? <leave blank>

Step 11: Set the primary interface speed? <leave blank>

Step 12: Set the primary interface duplex? <leave blank>

Step 13: Would you like to activate the in-path configuration? no

Step 14: Would you like to activate the out-of-path configuration? no

To change an answer, enter the step number to return to.

Otherwise hit <enter> to save changes and exit.

Choice: <press ENTER>

Configuration changes saved.

The initial configuration of your N29-SH datacenter out-of-path SteelHead is

now complete.

2. Perform a ping test to make sure that we can reach the datacenter router. To verify that your network interfaces have been configured properly, you

should use a simple ping (ICMP) test to test connectivity to the router.

n29-sh > en

n29-sh # ping -I primary 10.1.29.1

PING 10.1.29.1 (10.1.29.1) from 10.1.29.20 : 56(84) bytes of data.

64 bytes from 10.1.29.1: icmp_seq=1 ttl=64 time=0.467 ms

64 bytes from 10.1.29.1: icmp_seq=1 ttl=64 time=0.271 ms

64 bytes from 10.1.29.1: icmp_seq=1 ttl=64 time=0.381 ms

64 bytes from 10.1.29.1: icmp_seq=1 ttl=64 time=0.296 ms

{Press Ctrl+C to escape the ping test}

If your ping test is successful, you're good to go. When performing a ping test from the Steelhead CLI, you can use the

Following syntax to specify the source network interface. ping -I

<interface_name:primary,aux,inpath0_0> <destination address>

If not, you can run the configuration-jumpstart wizard again to check for any misc configured addresses. You will simply need to press Enter to accept the previously entered values for each question. You will also want to check your VLANs or network cabling and isolate any network problems. Troubleshooting Tip: When using a Steelhead virtual appliance in an SSOOP deployment, the LAN and WAN interfaces (3rd and 4th Network Adapter) should not be on the same VLAN or a network loop will be created that will cause Instability in your network. When using VMware ESXi, even network loops on

Isolated vSwitches without associated physical network adapters can cause CPU overload and instability when using the vSphere Client or vCenter. If this occurs, it is best practice to uncheck the Connected checkbox. To avoid accidental loops in the future, you may also want to create two separate port groups with different VLAN numbers for the LAN (NIC3) and WAN (NIC4) respective interfaces. In the lab environment, we use v999 as our "Discard" VLAN. Since we uncheck the connected box on the LAN and WAN network adapters, we can use v999 for both interfaces to avoid the loop.

Task 3

Verify Network Settings using CLI

It's a good idea to verify that the network settings for the Primary and AUX interface have been set properly. Keep in mind that the AUX IP address has been pre-configured in the lab environment.

1. Verify that the Primary and AUX interface configuration match the example below.Verify that the IP address that we configured in the configuration jump-start wizard matches the example below. All of the IP addresses for the VMs in the lab can be found on the topology diagram.

n29-sh # show int primary brief

Interface primary state

Up: yes

IP address: 10.1.29.20

Netmask: 255.255.255.0

MTU: 1500

HW address: 02:50:08:09:DD:C0

n33-sh # show int aux brief Interface

aux state

Up: yes

IP address: 10.99.29.20

Netmask: 255.255.0.0

MTU: 1500

HW address: 02:50:08:09:DD:C1

You can use the following syntax in the CLI to show interface status. The commands below show the shorthand syntax for your convenience.

2. If your primary interface's IP address and net mask do not match the example above, let's go ahead and set the proper IP address now. Note: You do not need to perform this step if the current primary address is 10.1.29.20.

Show interface <interface_name:primary,aux,inpath0_0> brief

If you need to change an interface IP address from the CLI, you can use the following command.

interface <interface_name> ip address <ip_address> <subnet_mask>

n29-sh # conf t

n29-sh (config) # no int pri dhcp n29-sh (config) #

int pri ip address 10.1.29.20 /24 n29-sh (config) #

wr mem