



PORTrockIT Topology Overview Guide Eli-v5.03.148

Bridgeworks

Unit 1, Aero Centre, Ampress Lane,
Ampress Park, Lymington,
Hampshire SO41 8QF
Tel: +44 (0) 1590 615 444
Email: support@4bridgeworks.com

Table of Contents

1	Introduction	3
1.1	Definition Of Terms	3
1.2	Supported Protocols Per Mode Of Operation	3
2	Topology: Bridged In-Path	4
2.1	When	4
2.2	Benefits	4
2.3	Limitations	4
3	Topology: Policy Routed Logical-In-Path	5
3.1	When	5
3.2	Benefits	5
3.3	Limitations	5
3.4	Routing Policies	6
3.4.1	Routing At The Host	6
3.4.2	Example	6
3.5	Routing At The Gateway	7
3.5.1	Example	7
3.5.2	Policy	8
4	Topology: WCCPv2	10
4.1	When	10
4.2	Benefits	10
4.3	Limitations	11
4.4	Example configuration with Redirect Lists	11
5	Topology: Out-Of-Path	13
5.1	Out-Of-Path Server Side Asymmetric	13
5.2	When	13
5.3	Benefits	13

5.4	Limitations	13
5.4.1	Example	14

Introduction

The intention of this topology document is to guide the user on which of the current modes that PORTrockIT can be deployed in would best suit their environment, their needs and WAN configurations.

It is not intended to be a definitive description of all the possible configuration but as a starting point upon which to build the final configuration.

It will in each circumstance provide a description as well as the benefits and advantages of each of the modes of operation. It is important to note that each mode of operation can be intermixed with any other combination of modes.

Definition Of Terms

An endpoint is defined as a server/host/client that selected traffic will be accelerated to or from within a PORTrockIT deployed topology.

A DMZ (demilitarized zone) is a physical or logical subnetwork that contains and exposes external-facing services to a usually larger and untrusted network.


A PORTrockIT Node is either a physical or virtual image of Bridgeworks PORTrockIT technology.

A mode of operation refers to the different modes that PORTrockIT can operate in.

Supported Protocols Per Mode Of Operation

Due to the rate of protocol adoption the supported protocol list is always increasing, please contact a Bridgeworks sales representative by email to sales@4bridgeworks.com for an up to date list.

Topology: Bridged In-Path



The PORTrockIT Nodes are deployed physically in-line between the client and server clusters and the WAN-facing router. In this configuration two network Interfaces become bridged together, a WAN port and a LAN port. All traffic flows through those two Network Interfaces, and traffic to be accelerated is diverted based on TCP port numbers to the acceleration engine. As the PORTrockIT Node in essence acts as a switch each port in the bridge has to be on its own isolated network, to prevent a 'switching loop'.

For non-accelerated traffic the PORTrockIT is transparent at the Ethernet frame level (layer 2), for all accelerated traffic the PORTrockIT will rewrite the layer 4 information (such as TCP sequence numbers) but the destination and source addresses of the traffic will remain those of the endpoints. QoS is preserved for non-accelerated traffic while traffic using tagged VLANs is passed through the Ethernet bridge and can be intercepted to accelerate traffic.

When

- If both sites already have a connection to each other without the need for routing rules. For example, if both sites are on the same subnet (L2 Adjacent).
- For a small installation with a single endpoint.


Benefits

- This does not require any routing configuration to the environment.
- No ongoing maintenance of the routing policies is required.

Limitations

- Physical access to cables and switches is required to the existing network infrastructure.
- For a Virtual Instance in ESXi, vSwitches must be set to have "Promiscuous Mode", "MAC Address Changes" and "Forged Transmits" enabled.
- Without a hardware bypass card this can become a single source of failure.
- This mode can only be deployed on ESXi or hardware due to the limitations of the environment.

Topology: Policy Routed Logical-In-Path



In this mode the PORTrockIT Node is no longer physically in the path, the Node acts as an additional router and local network policy is responsible for forwarding the traffic to the PORTrockIT.

PORTrockIT is transparent at the IP layer (layer 3) for non-accelerated traffic. For all accelerated traffic, we rewrite the layer 4 information (TCP sequence numbers/etc.) but the destination and source addresses remain those of the endpoints. QoS is preserved for non-accelerated traffic.

When

- If the PORTrockIT Nodes being deployed are on two different networks (i.e. a different CIDR block range). This is important to route traffic.
- If PORTrockIT is to become a VPN between sites.

Benefits

- For a Virtual Instance in ESXi, vSwitches can have “Promiscuous Mode” disabled.
- The only physical change to the network required is to plug the network ports from the PORTrockIT Node into a switch, ensuring that the LAN port is on the same network as the endpoint.
- Routing Policy can become specific enough to only route traffic to be accelerated, this will decrease the risk of failure for other traffic.
- WAN path failover and dead gateway detection is available allowing the transfer between sites to continue even if one WAN link connected to the PORTrockIT fails.
- This topology will work on all platforms.

Limitations

- Routing Policies must be created and maintained.
- Endpoints must be on different network subnets.

Routing Policies

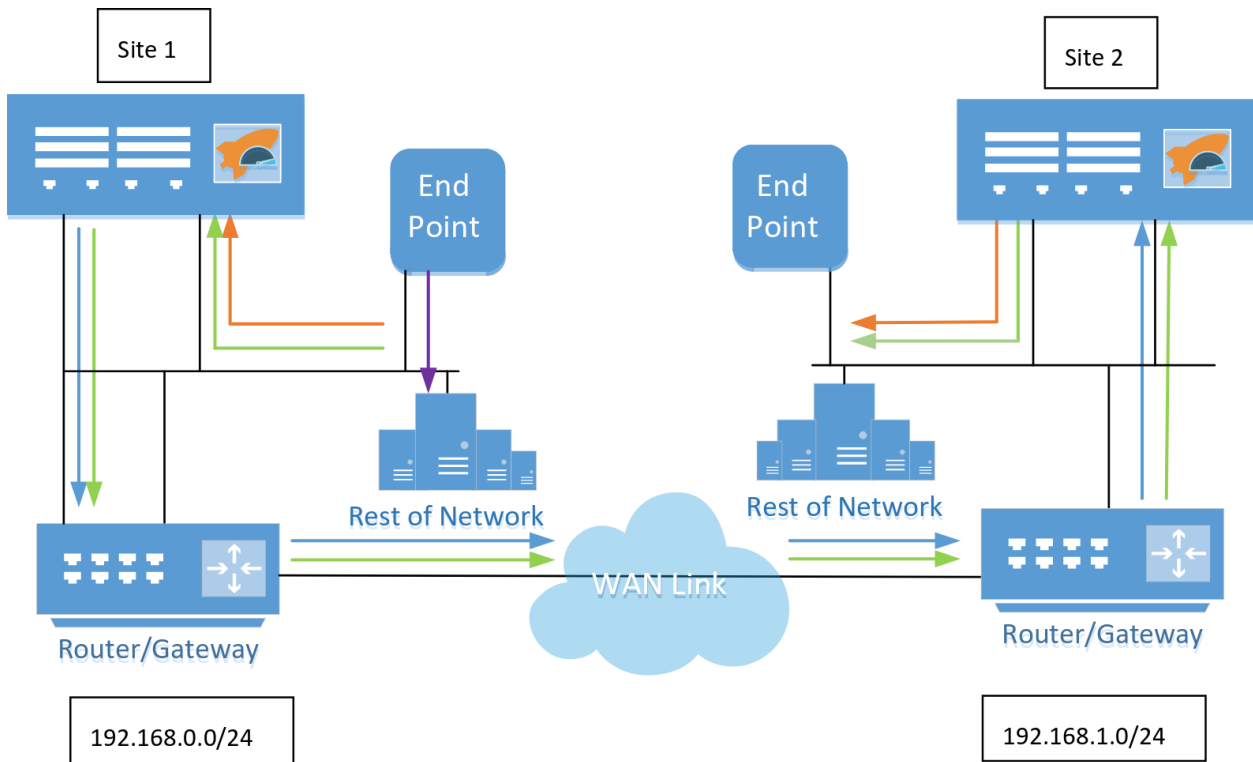
There are two primary locations for configuring network routing policies.





Routing At The Host

In this configuration, specific routing rules are added to the host using tools such as “ip route” or changing the default gateway to be the LAN address of the PORTrockIT network port. It is recommended to be as specific as possible when routing traffic, only sending traffic to be accelerated to the PORTrockIT.

Example

The following topology diagram contains arrows that represent the flow of different forms of traffic from “Site 1” to “Site 2”. This traffic flow occurs in both directions but for clarity the diagram has been simplified.



Traffic from the server destined for site 1 (Not to be accelerated)	
Traffic from the server to be accelerated destined for site 2	
Traffic from the server destined for site 2 (Not to be accelerated)	
PORTrockIT Accelerated Traffic Now on PORT 16665	

If routing policies are not specific on the server, all traffic destined for “Site 1” can be routed to the PORTrockIT Node which will forward the traffic to “Site 1”’s primary router (Green arrows). For example: The IP address of the LAN port of the PORTrockIT Node at “Site 1” is set to “192.168.0.10”. The following policy routing rule is added to the server at “Site 1”.

```
ip route add default via 192.168.0.10
```

This could be made more specific for example:

```
ip route add 192.168.1.0/24 via 192.168.0.10
```

In this example a VPN is set up between the sites, allowing all traffic destined for “Site 2” to be routed through the PORTrockIT Node, including all non-accelerated traffic such as, ICMP, UDP and any TCP ports that are not subject to acceleration on the PORTrockIT Node (Green arrows). The VPN connection can be used by both the “Server” and the “Rest of Network”.


All traffic to be accelerated is diverted to the Acceleration engine and optimised (Orange Arrows). Whilst the data is in flight between the sites (Blue Arrows), the packets are not in the same format as those received from the server until they are reassembled at the second PORTrockIT Node. No data is changed during this process.

Routing At The Gateway

Routing Policies can be set on a firewall or a gateway that will redirect traffic that is subject for acceleration to the PORTrockIT.

Example

The following topology diagram contains arrows that represent the flow of different forms of traffic from “Site 1” to “Site 2”. This traffic flow occurs in both directions but for clarity the diagram has been simplified.



Traffic from the server destined for site 1 (Not to be accelerated)	
Traffic from the server to be accelerated destined for site 2	
Traffic from the server destined for site 2 (Not to be accelerated)	
PORTrockIT Accelerated Traffic Now on PORT 16665	

In this example routing policy on the “Router/Gateway” is responsible for forwarding the traffic destined for “Site 2” to the LAN port of the PORTrockIT Node. The PORTrockIT Node then re-directs that traffic back to the router to get to “Site2”.

All traffic to be accelerated is diverted to the Acceleration engine and optimised (Orange Arrows). Whilst the data is in flight between the sites (Blue Arrows), the packets are not in the same format as those received from the server until they are reassembled at the second PORTrockIT Node. No data is changed during this process.

A VPN is set up between the sites allowing all traffic destined for “Site 2” to be routed through the PORTrockIT Node, including all non-accelerated traffic such as, ICMP, UDP and any TCP ports that are not subject to acceleration on the PORTrockIT Node (Green arrows).

Policy

To illustrate the policy on the router/gateway the following example shows some basic config options within a Cisco routers config. In this example the LAN port of the PORTrockIT Node I “Site 1” is 192.168.0.10 and the endpoint to be accelerated is a NetApp appliance with its replication running on TCP port 11104.

First match the source of 192.168.0.0/24 destined for 192.168.1.0/24 for traffic running on TCP port 11104.


```
access-list 101 permit tcp 192.168.0.0 0.0.0.255 192.168.1.0 0.0.0.255 eq 11104
```

Once complete, create a route map called ‘clients’ with next hop for matching access-list as

192.168.0.10 (this redirects traffic matching the access-list rule to 192.168.0.10, the IP address of the PORTrockIT Node).

```
route-map clients permit 10  
match ip address 101  
set ip next-hop 192.168.0.10
```

Topology: WCCPv2



This mode can be considered an extension of the “Policy Routed Logical In-Path” mode, where the PORTrockIT communicates with the L3 switch or router that is implementing the routing policies. This communication allows the router to ensure the PORTrockIT is operational prior to sending traffic to be accelerated, and allows automatic load-distribution for PORTrockIT units with multiple LAN ports.

In this mode, similar to the “Policy Routed Logical In-Path” mode the PORTrockIT Node is not physically in the path, but unlike “Policy Routed Logical In-Path” mode, packets that cannot be accelerated are not routed by the PORTrockIT but instead returned to the L3 switch or router to perform the next routing step.

It is recommended that PORTrockIT Nodes in this topology are matched with other PORTrockIT Nodes in the same topology. But it may be inter-mixed with nodes in “In-Line bridged” and “Out-of-Path” mode.

PORTrockIT is transparent at the IP layer (layer 3) for non-accelerated traffic. For all accelerated traffic, we rewrite the layer 4 information (e.g. TCP sequence numbers) but remain transparent at layer 5 and above. There is no support for VLAN tagging but QoS is preserved for non-accelerated traffic.

When

- Deploying into an existing network which uses Cisco® switches or routers.
- Physically out of path failover is required.
- If the PORTrockIT Nodes being deployed are on two different networks (i.e. a different CIDR block range). This is important to route traffic.
- This topology will run within AWS using the Cisco® Cloud Services Router.

Benefits


- For a Virtual Instance in ESXi, vSwitches can have “Promiscuous Mode” disabled.
- The only physical change to the network required is to plug in the three network ports from the PORTrockIT Node into a switch, the LAN port must be on the same network as the endpoint.
- Routing Policy can become specific enough to only route traffic to be accelerated, this will decrease the risk of failure for other traffic.

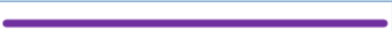



- WAN path failover is available allowing the transfer between sites to continue even if one WAN link connected to the PORTrockIT fails.
- When multiple interfaces are configured on the PORTrockIT, WCCPv2 will ensure the switch or router load balances over these interfaces.
- Allows for failover between PORTrockIT Nodes, although the endpoint must re-establish its connection as connections aren't statefully failed over to a partner PORTrockIT.

Limitations

- Only works with Cisco® routers or switches.
- The PORTrockIT must be located on the same L2 network as the switch or router that is configured with WCCPv2.
- 2 IP addresses must be allocated at each site, one for the LAN port and one for the WAN port.
- No support for VLAN tagging.
- This mode is not compatible with the PORTrockIT integrated VPN.
- Endpoints must be on different network subnets.

Example configuration with Redirect Lists



Traffic from the server destined for site 1 (Not to be accelerated)	
Traffic from the server to be accelerated destined for site 2	
Traffic from the server destined for site 2 (Not to be accelerated)	
PORTrockIT Accelerated Traffic Now on PORT 16665	

To illustrate the WCCP configuration on the switch/router the following example shows a basic configuration for a Cisco® router. In this example the LAN port of the PORTrockIT Node at “Site 1”

is 172.16.0.10 and the endpoint to be accelerated is a NetApp appliance with its replication running on TCP port 11104-11105.

```
ip access-list extended WCCP_NETAPP
    remark NetApp Appliance
    permit tcp 192.168.0.0 0.0.0.255 192.168.1.0 0.0.0.255 range 11104 11105

ip wccp check services all
ip wccp check acl outbound
ip wccp 42 redirect-list WCCP_NETAPP

interface FastEthernet0/0
    ip address 192.168.0.1 255.255.255.0
    ip wccp 42 redirect in
    duplex auto
    speed auto

interface FastEthernet0/1
    ip address 172.16.0.1 255.255.0.0
    duplex auto
    speed auto
```


Configuration on the PORTrockIT side is as simple as providing the switch/router's WCCP IP address and the service ID to use.

When using multiple PORTrockIT units, each connected remote PORTrockIT requires its own unique WCCPv2 service group to provide per-site failover in the event of network disruption.

Cisco is a registered trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

Topology: Out-Of-Path

Out-Of-Path Server Side Asymmetric



In this mode, the client side PORTrockIT at “Site 1” can be configured to be in “In-Line bridged” or “Policy Routed Logical-in-Path”. It is important to note that only the server side endpoint can be in the “Out-Of-Path” mode. Unlike the previous modes the source IP address of the traffic has changed from the client side endpoint to the IP address of the PORTrockIT Node. This requires no configuration change to the server environment. This configuration is suitable for traditional client-server relationships where connections are established from clients to the server.

When

- If the chosen protocol to accelerate is a client server relationship protocol. For example, sending and retrieving data from an Object Store.
- If you can't logically put your PORTrockIT on the same network as the endpoints to be accelerated.

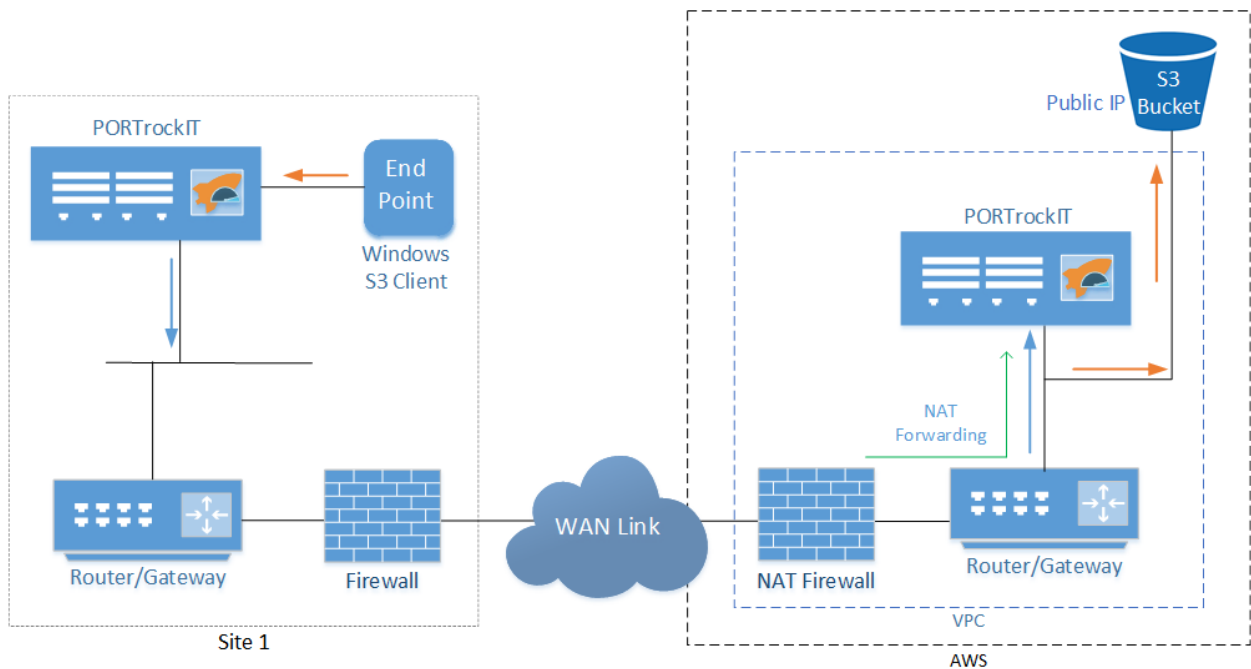
Benefits

- For a Virtual Instance in ESXi, vSwitches can have “Promiscuous Mode” disabled.
- This topology will work on all platforms.
- The only physical change to the network required is to plug in the three network ports from the PORTrockIT Node into a switch which is in a network which has a route defined, to the endpoint.
- No Routing Policies must be set on the server site “Site 2”, nor does it have to be physically in path.
- WAN path failover is available allowing the transfer between sites to continue even in one WAN link connected to the PORTrockIT goes down.

Limitations

- Only compatible with client-server relationships.

Example



In this example, “Site 1” is configured to be in “Bridged-in-Path”, in “AWS” the PORTrockIT Node is configured to be in “Out-of-Path” mode and is deployed as an EC2 image. The object store and the PORTrockIT are not in the same network domain and the object store is an externally hosted service preventing its routing from being modified.

All traffic that is to be accelerated is diverted to the Acceleration engine and optimised (Orange Arrows). Whilst the data is in flight between the sites (Blue Arrows), the packets are not in the same format as those received from the server, until they are reassembled at the second PORTrockIT Node. No data is changed during this process.