XipLink™

# TechnologyInformationBrief

## Scaling XipLink XA Deployments

*Load distribution through multiple virtually inline appliances*

## Introduction

XipLink XA appliances are among the most scalable optimization solutions available. They are generally designed such that an XA appliance can singularly handle the bandwidth load of most satellite networks. However there are high bandwidth networks, both in point to point links and at the hub of a star network that can exceed the capabilities of a single XA 30K unit. This document will briefly introduce means to distribute the load among multiple appliances to achieve scalability.

## XipLink Deployment Modes

Satellite and wireless networks are not done with a universal design. XipLink addresses different network topologies, designs and priorities by being exceptionally flexible with how the XA appliance can be added to the network and intercept traffic for interception.

XA appliances achieve this flexibility in part being based on SCPS and it's advantages such as it's dynamic capability negotiation, support for meshing, and IP and TCP level transparency. To install on a network to perform optimization, XA appliances use mechanisms such as:

- bridging
- routing
- tunneling
- encryption
- NAT
- dynamic routing
- selective optimization
- firewalling
- DSCP remarking
- DHCP client or server
- the ability to function with just one or several network interfaces
- fail-to-wire and more.

Further, most of those mechanisms can be used simultaneously.

Central to most high bandwidth networks is the ability to deliver very high availability. XipLink supports high availability through the use of hot redundancy. XipLink supports a hot redundant solution, allowing one or more units to be used as backup units which can automatically, and very quickly, take over for primary units in the event of hardware, software, cabling or network failures.

In the simplest setup of an XA device, it is installed in the network path, either as a router or a bridge. An XA-30K is capable of 155Mbps and 30,000 connections. Importantly, TCP connections in excess of 30,000 are still allowed through the device -- they are simply not intercepted and then compressed, accelerated or otherwise enhanced.

However, on larger networks which exceed the abilities of an XA-30K appliance, either through session limitations, PPS or bandwidth limitations, a single appliance is no longer viable as an option. In this situation, load must be distributed among multiple devices, or a limited portion of the network selected for optimization.

## Distributing Load with Policy Based Routing

The most straightforward mechanism to achieve this is with Policy Based Routing ("PBR"). This is not so much a function of the XA appliance, but rather the capability of virtually all modern IP router appliances to intelligently route IP traffic based on the source IP addresses to the optimization device, and not allow the IP packets to route in a loop on the network. This configuration can be used both with single-sided (hub-only) deployment of XHO and also with in two-way solutions.

- In a PBR configuration with no scaling and no redundancy, the router would set the IP address of the XA appliance as the next hop route to the remote network with a standard route, and set a policy based route such that traffic arriving from the remote network is first passed to the XA appliance.

- In a PBR configuration with redundancy, the same approach is used, however instead of specifying the IP address of the XA appliance directly, the router uses a route towards the virtual IP address that is shared among the redundant units.
- In a PBR configuration with multiple units, the network is segmented into pieces, and traffic is distributed among the multiple XA units. The traffic would be directed to the virtual IPs of the individual redundant pairs of appliances. In this way, the traffic can be distributed among multiple appliances to achieve the necessary scalability.

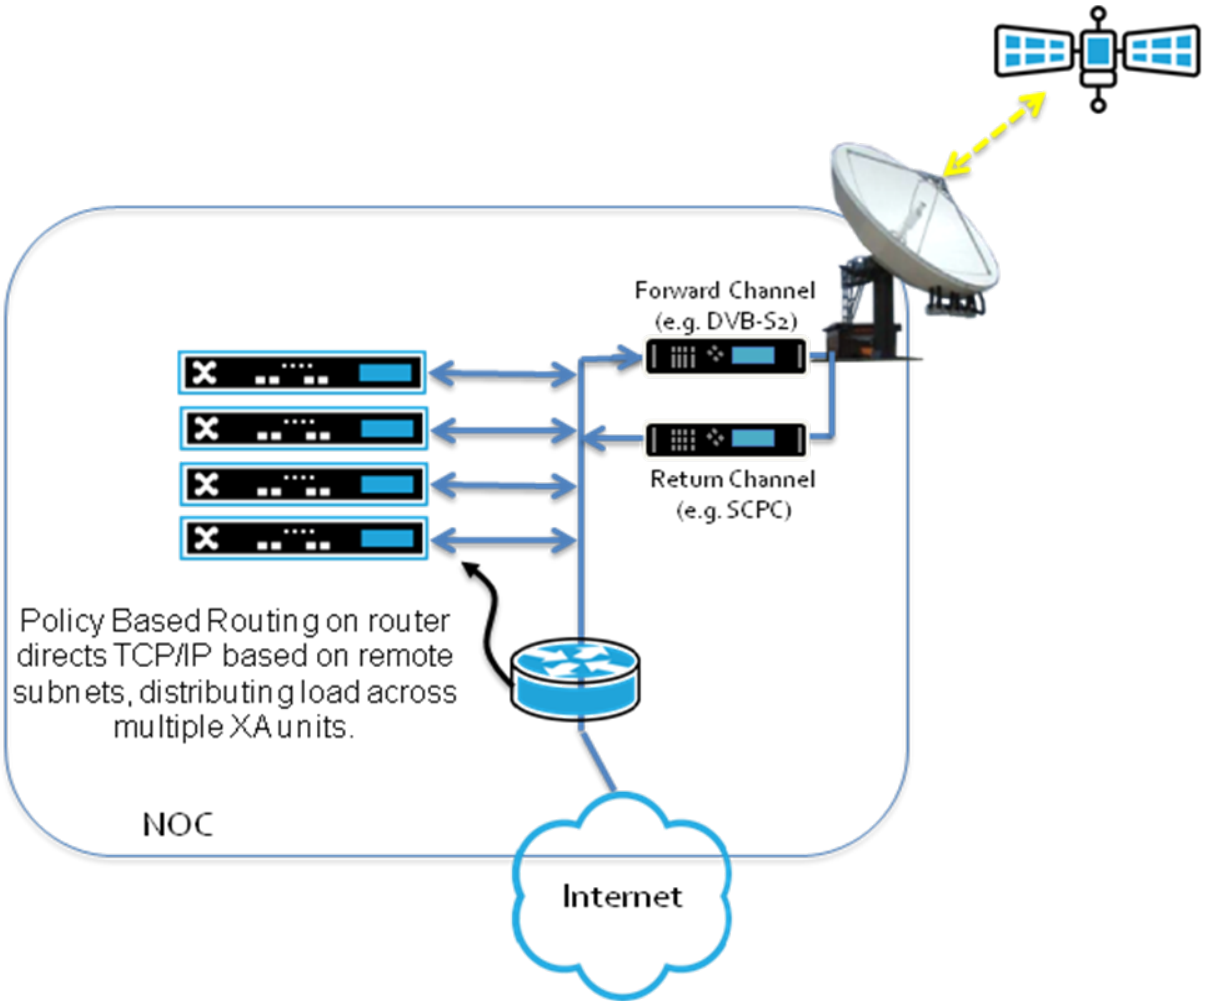


This PBR solution also allows incremental deployment through steadily increasing the portion of the network to the XA appliance(s) once the optimizations are confirmed to be functional. For instance, just a single IP address or network link could initially be optimized. From the perspective of the XA appliance, the routing tables will remain the same through the scaling process.

The PBR can also act with differentiators such as port numbers. For instance, when deploying the XipLink XHO solution for one-way compression, the installer could choose to only forward port 80 traffic. However, to achieve TCP acceleration and QoS of all streams with the XHO solution, all traffic should be directed to the XA appliance(s).

PBR is a proven way to divert IP traffic for optimization without having to re-cable the critical network path. However, it does require the installer to be familiar with Policy Based Routing.

## Distributing Load with WCCP

If you are using an advanced Cisco router at the hub, another alternative to PBR is to use Web Cache Communication Protocol (WCCP). WCCP follows a similar concept in that the XA Appliance is installed not in the direct network path but rather the router directs IP traffic using a GRE tunnel or Layer 2 (L2) technique to the XA, and the XA appliance sends it back to the router after optimization.

Importantly, WCCP provides automatic load sharing among multiple XA appliances.

WCCP can require less configuration than PBR, and while there are ways with PBR to automatically route around the XipLink device in case of failure, processes like heartbeats and automatic disconnection in the event of failure are integrated automatically with WCCP.

It also supports operation with multiple routers concurrently.

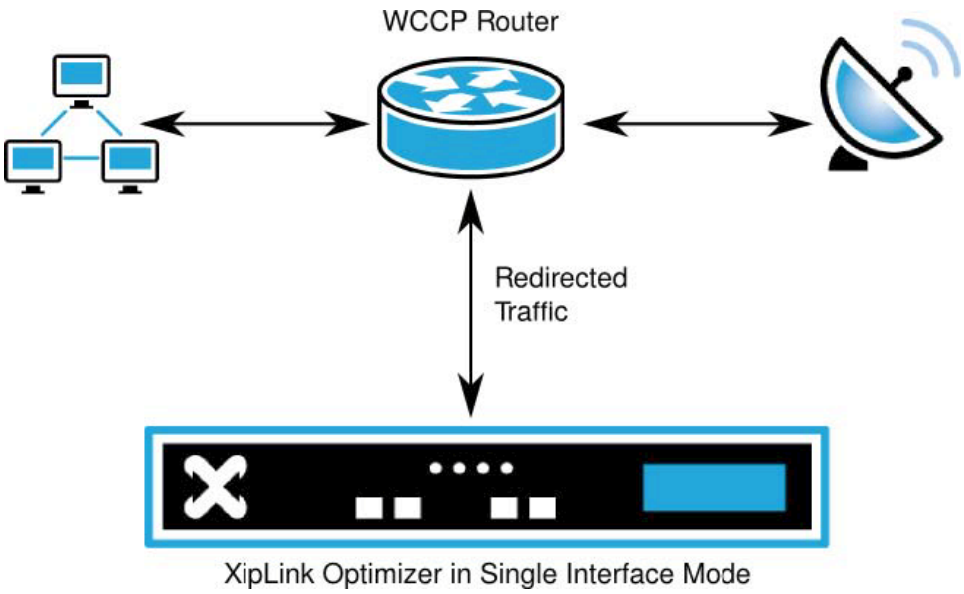


For further information on WCCP and XA Appliances, please review the Technical Information Bulletin written specifically on the topic, *"Deploying XipLink XA Units with WCCP"*.