XipLink™

TechnologyInformationBrief

Deploying XipLink XA Units with WCCP

Virtual inline deployment alongside Cisco Routers

Introduction

The Web Cache Communication Protocol (WCCP) allows a Cisco router to automatically redirect network traffic to an XA optimizer, permitting a non-inline deployment that incorporates heartbeat capabilities and possible load distribution among multiple units for increased scale. It can also minimize network disruption when installing it to a few seconds. However, users must be cautious of the network topology and also aware of the router limitations for handling the load of performing WCCP.

XipLink optimizers implement WCCP version 2. WCCPv2-compliant routers are only available from Cisco (WCCPv2 was first introduced in IOS 12.0(5)T).

Capabilities

XipLink optimizers currently implement the following WCCP features:

- Support for HTTP and non-HTTP traffic, including UDP
- GRE and Layer 2 packet forwarding
- Multiple routers
- Multiple optimizers
- Automatic failover
- Limited load distribution (see below)

XipLink optimizers currently do not support the following WCCP features:

- GRE or Layer 2 packet return (the optimizer forwards its packets normally using ARP lookups)
- MD5 shared security
- Multicast groups

WCCP Deployment

WCCP allows you to deploy XipLink optimizers without interrupting your network. Unlike a typical inline installation, with WCCP you simply connect the XipLink optimizer directly to your network's router, then configure both the optimizer and your router to redirect traffic through the XipLink device.

WCCP can also be used to progressively deploy XipOS technology by starting with a small remote network to optmize (for example, based on a limited IP range) and then incorporating more and more traffic redirection as confidence with the installation progresses.

The optimizer can also connect to the WCCP router in Router mode, as shown in figure 2. This requires two network ports on the router, one to send redirected traffic to the optimizer and another to receive traffic from the optimizer. The router redirects traffic to the optimizer's Routed interface, and the optimizer returns traffic from its Wireless interface, which the router routes normally.

Virtually anything can be forwarded, not just TCP. The deployments described here can also be achieved without WCCP using policy-based routing.

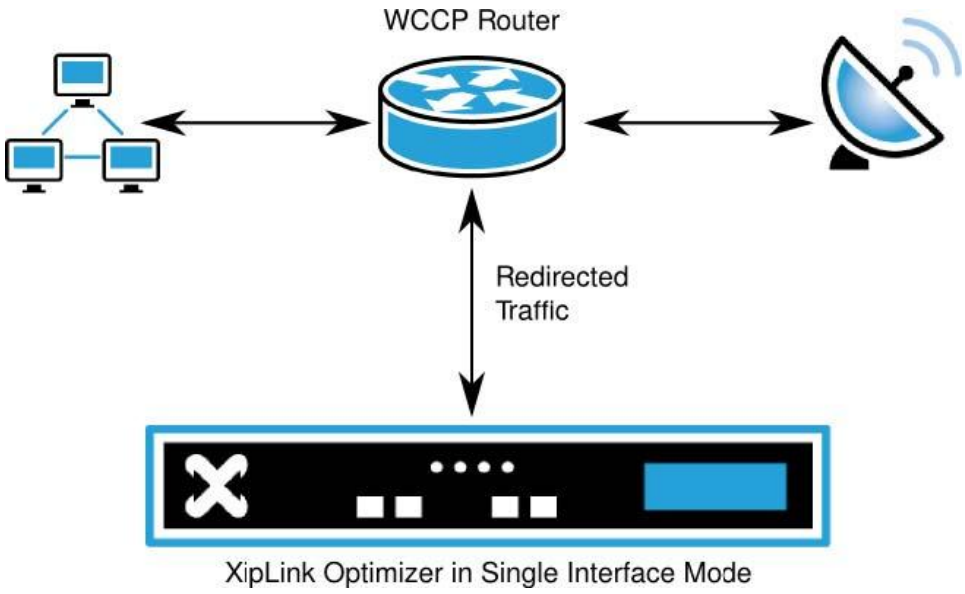

Figure 1: Deployment in Single Interface Mode

Figure 1 depicts the simplest WCCP deployment. The optimizer's Wireless interface is connected to a WCCPv2-capable router that lies between the local network and the wireless uplink. In this example the optimizer is configured in Single Interface mode so all redirected traffic is sent to the optimizer's Wireless interface.

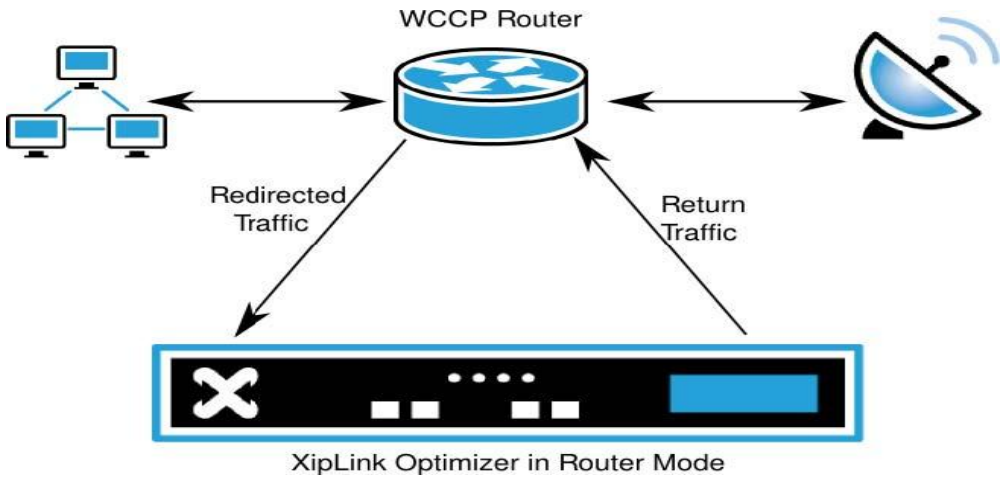

Figure 2: Deployment with Two Network Interfaces

The rest of this section discusses WCCP deployments with the optimizer in Single Interface mode, but all the deployments are also applicable when the optimizer in is Router mode.

Hub Deployments with Multiple Routers

A single optimizer can support more than one router with WCCP. This is most useful for hub deployments, where the optimizer can accelerate traffic with several remote sites each connected to the hub network via its own router.

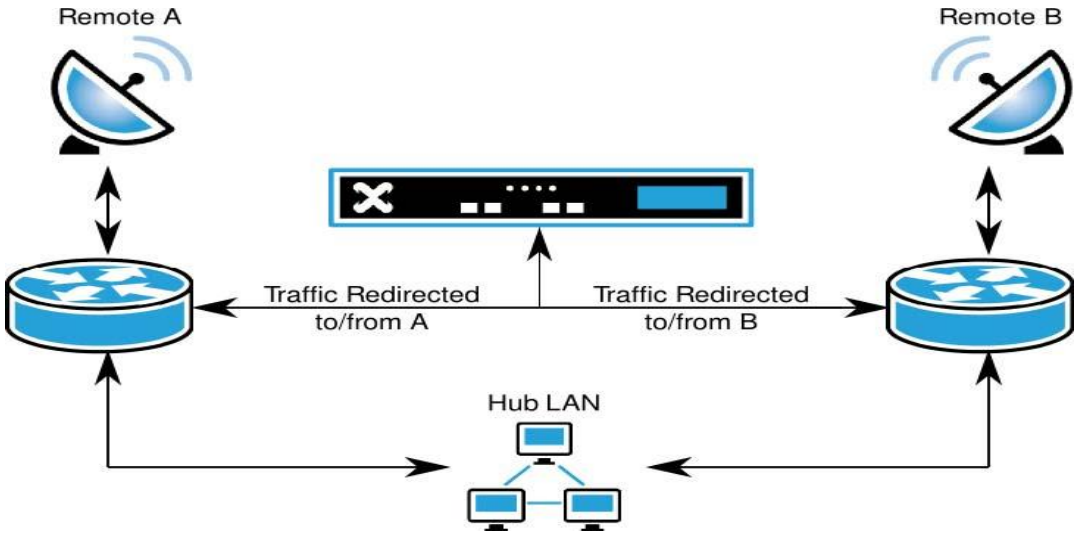

Figure 3: Multiple Router Deployment

Load Distribution with WCCP

WCCP deployments can also make use of more than one XipLink optimizer. This can be used to share load for increased scalability or to distribute applications. It is also a alternate redundancy strategy. In this scenario, multiple XA Appliances register with a router to handle IP traffic, and a algorithmic strategy is used to distribute.

Registering multiple units can achieve increased scale beyond the capabilities of a single unit, while operating on the same traffic definition. This requires the router(s) to use GRE packet forwarding. WCCP's load distribution algorithm is based on source (or destination) IP addresses and ports, and so it will not share the load equally among the optimizers.

Another deployment strategy is to allocate different types of traffic to different optimizers. For example, HTTP traffic can be handled by one optimizer (with an XHO license), all other TCP traffic by another, and all UDP traffic by yet a third optimizer.

There is no way for multiple XipLink optimizers to share their quality-of-service (QoS) states, and so any multiple-optimizer deployment may not correctly enforce QoS settings. Therefore the recommendation is to only use WCCP load distribution when deployed alongside a solution with an integrated PEP / TCP acceleration solution. Load balancing is supported in GRE redirection mode.

WCCP Configuration

The configuration of WCCP within the XA Appliance is generally straightforward and the user interface makes it easy. Service IDs are used to identify the kinds of traffic to be forwarded.

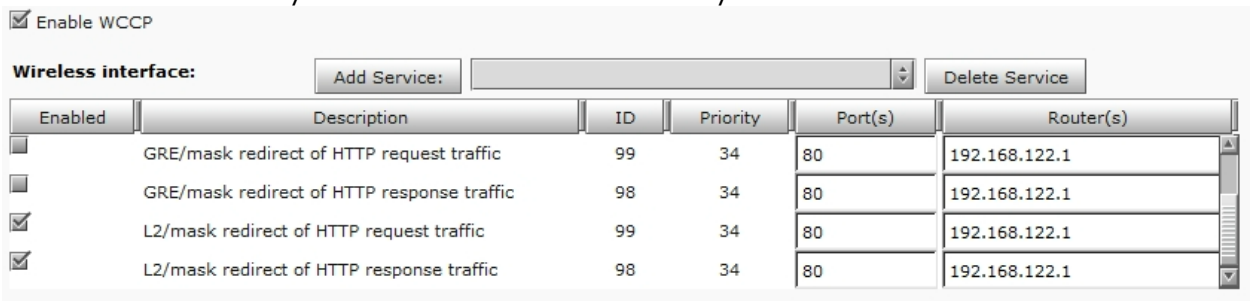

Figure 4: Snapshot of WCCP configuration within XipOS User Inteface

More complicated is the configuration within the Cisco router to properly enable WCCP. In deployment, it is important on most routers to ensure that the WCCP forwarding is occurring in hardware, as software-based operation is generally far less scalable. The use of hardware-based forwarding depends entirely on the router selected and it is best to consult with the XipLink technical team before planning a high bandwidth deployment.

Consult the XipLink user manual for detailed configuration information.