

Part 2 in the Reliability Series

Reliable Backbone Bandwidth using Composite Links™

Abstract

As IP networks evolve from carrying commodity transport to carrying delay-sensitive and mission critical services, carriers require new backbone protection solutions. This paper examines the traditional Active/Standby redundancy model and identifies a new cost-effective alternative to backbone protection based on loadsharing redundancy.

Authors

*Adam Dunstan
Hudson Gilmer*

Reliable Backbone Bandwidth using Composite Links

Introduction

Today's leading Service Providers are motivated by financial necessity to upgrade their IP networks from providing "best efforts" commodity transport to delivering differentiated high-quality, high-margin services.

Part 1 of the Reliability series investigates the causes of IP network downtime and the financial impact to carriers of improving network availability. As detailed in this paper, 32% of downtime in IP networks is due to link failures.

In order for carriers to minimize the financial impact of inevitable fiber cuts and optical component failures, backbone protection solutions must be implemented which can rapidly reroute traffic across redundant links to minimize/eliminate service disruption.


This paper details strategies for carriers to provide rapid, cost-effective and protocol-independent mechanisms to protect against link failures in the optical core.

Traditional "Active/Standby" approaches to Backbone Reliability

For the last 10-15 years, most carriers have utilized Switched Optical Network (SONET/SDH)/Switched Digital Hierarchy (SDH) rings for backbone protection in the optical core. In a typical SONET/SDH ring infrastructure, two physical paths (active and standby) between SONET/SDH terminals are provisioned to provide redundancy. This design ensures that should link failure occur in the active path, traffic is redirected over standby path with a switchover time of 50ms. A 50ms switchover time is fast enough to prevent data and voice service disruptions under normal circumstances.

While SONET/SDH active/standby redundancy has effectively served its purpose in protecting optical links, it is too expensive for widespread deployment. In addition to the cost of deploying SONET/SDH equipment, fully 50% of transmission capacity is consumed by standby links. This doubling of optical backbone cost prevents most carriers from deploying SONET/SDH rings on all but their most critical routes.

Figure 1: Active/Standby Facilities Redundancy



Another drawback is that a SONET/SDH ring implementation portrayed in figure 1 provides redundancy in the optical core, but does not protect the link from failures at the router interfaces, SONET/SDH Add-Drop Multiplexer (ADM), or intra-PoP links connecting them.

Automatic Protection System - Extending Active/Standby protection to the Router


In order to address this vulnerability, some carriers have chosen to implement Automatic Protection System (APS) to extend Active/Standby redundancy beyond the optical fiber to the router interfaces on either end (figure 2).

APS relies on router ports which switch between working / protect ports based on header indicators produced by the SONET/SDH ADM. Depending on the configuration, the two circuits may be terminated in the same router or in different routers.

Reliable Backbone Bandwidth using Composite Links

This approach eliminates the single points of failure in figure 1, but further increases the cost due to duplication of expensive line cards on each router.

Figure 2: Active/Standby Facility + Port Redundancy



Needed: A widely deployable backbone protection architecture

Carriers require a protection solution which is cost-effective enough for deployment throughout their backbone network – not only critical routes. This requires a fundamentally different approach than the traditional 1:1 Active/Standby protection model.

MPLS Fast Reroute/Local Protection offers some advantages – rapid restoration and flexibility to define backup routes per LSP/traffic class. However, because Fast Reroute/Local Protection operates at layer 3 and applies only to MPLS traffic, it cannot provide a total replacement for the layer-1 protection of SONET/SDH rings.

Equal Cost Multi-Path (ECMP) is a failover feature of OSPF which allows similar links to be grouped to protect each other in the event of failure. The drawbacks of ECMP are that it cannot match SONET/SDH protection times and cannot be used in conjunction with MPLS or IS-IS.

Introducing Composite Links

Composite Links were developed by Avici Systems to simplify bandwidth management, improve bandwidth utilization, and deliver cost-effective backbone protection.

A Composite Link is comprised of up to 64 physical links or fiber interfaces between 2 Avici TSRs. These links are viewed by the higher layer routing and switching protocols as a single physical interface consisting of the aggregate capacity of all member links.

Each Composite Link member is a POS (packet over SONET/SDH/SDH) interface and is established using the PPP protocol. This standard encapsulation enables multi-protocol support and is used by Avici to provide transport for IP traffic, OSI control traffic used to deliver the IS-IS routing protocol and MPLS traffic.

The interfaces used in Composite Links can have a 4:1 speed mismatch. This gives carriers flexibility in migrating to higher speed trunks or to provide greater granularity of capacity. For example, a carrier can create an OC-768 composite link consisting of a combination of three OC-192c and four OC-48c members.

If a member of a Composite Link fails, all data that was being transmitted on that link is automatically re-distributed to the other Composite Link members. This re-distribution occurs in less than 45ms and does not impact traffic traversing other members of the Composite Link.

To ensure traffic flows are delivered in the order they were received, a hash value is calculated from the IP source and destination address, this value is used to select the Composite Link member. This ensures that a specific dataflow is always sent down the same Composite Link member therefore maintaining packet ordering.

Engineering Enhanced Reliability with Composite Links

The advent of Composite Links fundamentally changes the economics of backbone protection by enabling protocol independent backbone and port protection


Reliable Backbone Bandwidth using Composite Links

on a loadsharing (1:n) basis rather than on an Active/Standby (1:1) basis.

Loadsharing redundancy means that 1 unit provides redundancy for n units of traffic. In normal operation, n units of traffic are spread across n+1 active units. In the event of an outage on any of those units, each of the n units carries 1/n of the total traffic.


For example, a carrier with 3 x 2.5G of traffic between two PoPs could define a composite link consisting of 4 OC-48c members to carry the traffic. (figure 3)

Figure 3: Composite Links Normal Operation



In the event that any individual member failed, traffic would be distributed across the three surviving links without impacting the upper layer routing (IP) or switching (MPLS) protocol. (figure 4)

Figure 4: Composite Links After Link Failure




In this example, $n=3$ and cost premium to protect each link is equivalent to 33% of the unprotected link cost. Contrast this with Active/Standby redundancy portrayed in Figure 3 where the cost premium to protect each link is 100% of the unprotected link cost. The cost-effectiveness of loadsharing redundancy increases with the value of n , as depicted in the figure 4.

Even when only two physical routes exist, carriers can more efficiently utilize available bandwidth by combining MPLS and Composite Links. If a define composite link loses a member, lower priority “best efforts” traffic is rerouted or dropped while higher priority SLA-backed traffic is unimpacted.

This approach allows carriers to fill the otherwise empty standby route with “best-efforts” traffic and increase total revenues without risking SLA penalties.

Figure 4: Cost of Backbone Protection



The application of Composite Links in DWDM environments

In order to accommodate rapidly increasing traffic requirements, carriers have begun migrating from SONET/SDH to Dense Wave Division Multiplexing (DWDM) in the optical core. DWDM allows multiple channels of 2.5Gbps or 10Gbps to be transmitted on discrete wavelengths of a single fiber.

The combination of DWDM and Composite Trunks offers tremendous data transmission capacity between two TSR with manageable growth granularity and flexible protection options.

For carriers migrating or planning to migrate from SONET/SDH rings to DWDM rings or meshed optical networks, Composite Links provides a consistent, flexible protection mechanism for smooth migration regardless of the optical transmission technology.

Summary

With IP networks carrying increasing quantities of delay-sensitive and SLA-backed services, backbone protection is becoming a necessity throughout the network – not a luxury for selected routes.

Reliable Backbone Bandwidth using Composite Links

The advent of DWDM and mesh optical networks is providing further stimulus for carriers to revisit their bandwidth protection strategies.

In light of these developments, carriers are concluding that traditional models of Active/Standby protection are not cost-effective for widespread deployment, and are beginning to turn to Composite Links as a cost-effective mechanism to ensure protection throughout their backbone.

Table 1: Protection Type Comparison

	SONET/ SDH Ring	SONET/ SDH Ring w. APS	ECMP	Comp. Link
Layer-1 independent	N	N	Y	Y
MPLS/IP interoperable	Y	Y	N	Y
Protects ports?	N	Y	N	Y
Port protection cost	-	1/1	-	1/n
Link protection cost	1/1	1/1	1/n	1/n