FORMULATION AND BENEFIT ANALYSIS OF OPTIMIZATION MODELS FOR NETWORK RECOVERY DESIGN

Approved by:

Dr. Richard Barr

Dr. Eli Olinick

Dr. Marion Sobol

Dr. Jerrell Stracener

Dr. Stephen A. Szygenda

FORMULATION AND BENEFIT ANALYSIS OF OPTIMIZATION MODELS FOR NETWORK RECOVERY DESIGN

A Praxis Presented to the Graduate Faculty of the
School of Engineering
Southern Methodist University
in
Partial Fulfillment of the Requirements
for the degree of
Doctor of Engineering
with a
Major in Engineering Management
By

M. Scott Kingsley

(B.S., Engineering Technology, 1989)
(M.S., Southern Methodist University, 1993)

Date

COPYRIGHT

ACKNOWLEDGEMENTS

Kingsley, M. Scott

B.S., Edison State College, 1982
M.S., Southern Methodist University, 1993

TITLE

Advisor: Professor Richard S. Barr

Doctor of Engineering Management conferred

Praxis completed

Overview:

Table of Contents

1 OVERVIEW OF PROBLEM 2

LIST OF FIGURES

LIST OF TABLES

DEDICATION

1 Overview of Problem

Telecommunications service providers must provide reliable networks. Traditionally their goal has been to provide “Five 9’s”, or 99.999%, network up time on a link-by-link basis. This translates to less than five minutes of down time per year on any link. *Traffic* is any combination of voice, data or video transmitted over such a network. The total interruption of traffic is a *failure* and can result from either inoperable equipment or communications links between equipment locations. Failure *recovery* is a generic term describing various methods to reroute traffic from failed *working* (primary) equipment or link capacity onto operating *spare* (backup) equipment or link capacity.

It is commonplace that spare equipment and capacity cost exceed working network cost. The challenge for telecommunications service providers is to minimize the cost of recovering failed working traffic onto spare capacity (equipment, links, or both) given a large number of possible network failure possibilities. Unfortunately, little conclusive research exists to guide service providers given different network scenarios and constraints. The problem is amplified by the relatively recent introduction of non-uniform Quality of Service (QoS) requirements whereby real time services, such as Voice over IP (VoIP) and streaming video, must be recovered immediately while recovering less critical services, for example web searches, can be delayed or eliminated completely.

Generally, though, service providers are far more concerned with link than equipment failures because link failures occur more frequently and take longer to repair than equipment failures. In addition, only single failure scenarios, such as a single failed link, are usually considered in analysis. These relaxations significantly reduce the complexity problem.

1.1 Network Models

To understand the functionality of telecommunications networks the Open Systems Interconnection (OSI) model and Internet framework were developed. Although there are subtle differences between the two, at a higher level they are accomplished identically.

The OSI model and Internet framework both subdivide a conceptual communications system into smaller parts called layers. Each layer has defined responsibilities and provides connection and services to the layers above and below it.

1.1.1 OSI Model

The International Standards Organization (ISO) Open Systems Interconnection model was developed to standardize and separate networking functions into seven different layers (Medhi & Ramasamy, 2007). There were actual protocols developed and used as part of the OSI model. However, today the OSI model is predominantly that—a model. The original protocols developed and used as part of the OSI model are virtually obsolete except in federal government applications. Protocols used today in most networks have been replaced with those developed by the Internet Engineering Task Force (IETF) which also modified the OSI model with the Internet framework (Huitema, 1995). However, the use of the OSI model layers to describe network functions is a de facto standard in the industry even though it may describe actual protocols that are part of the TCP/IP framework (Table 1).

OSI Model Layer Number	OSI PDU Name	OSI Model Layer Name	Protocols	Internet Framework Layer Name	Internet Framework Layer Number
7		Application	DNS, DHCP, FTP, Telnet, SMTP, HTTP	Application	4
6		Presentation	EBCDIC, ASCII, JPEG, SSL		
5		Session	RPC, SSH, SQL		
4	Segment	Transport	TCP, UDP, RTP	Transport	3
3	Packet	Network	IP	Internetwork	2
2	Frame	Data Link	Frame Relay, ATM, Ethernet	Network Access	1
1	Bit	Physical	SONET, T1, T3		

Table 1: Comparison of OSI model and Internet framework

It is important to observe that the top four layers of the OSI model (Application, Presentation, Session and Transport) are only used in end devices, such as personal computers, while the bottom three layers are the only layers used in network routers, switches, and multiplexers.

1.1.2 Internet framework

The Internet framework combines the Application, Presentation and Session layers into a common Application layer while keeping the Transport layer separate. Below the Transport layer is the Internetwork Layer.

The OSI Data Link and Physical Layer are combined into a single Internet framework Network Access Layer. In the Internet framework only the Internetwork and Network Access layers are used in network equipment while the upper Transport and Applications layers reside in end devices.

1.1.3 Encapsulation

Each layer of the OSI model or Internet framework performs has specific responsibilities and provides services to layers above and below it. This is accomplished using "encapsulation" and is described using the OSI model (Figure X).

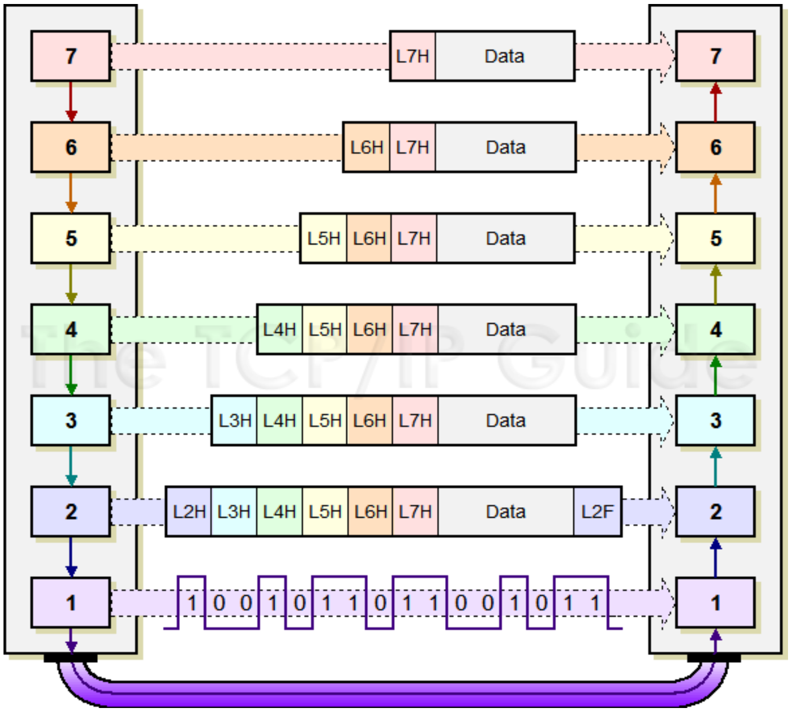

http://www.tcpipguide.com/free/t_DataEncapsulationProtocolDataUnitsPDUsandServiceDa.htm

Figure X

Data from the end user starts at the top, or Layer 7, of the OSI model. At each layer an applicable protocol is applied and instruction included on what to do with the data by placing them in *overhead* in front of the data. The result is called the "Protocol Data Unit" (PDU) for Layer 7. Once this is accomplished, the combination is passed down to the next lower layer which sees only the received layer PDU only as data. The next lower layer cannot access the information it receives from the layer above. The next lower layer then again applies a layer specific protocol and adds instructions on what to accomplish in the layer overhead and passes it to the next lower layer.

The PDU's of the bottom four layers have specific names that correspond to the layer they are accomplished at. As the data is encapsulated and handed down to the OSI model Layer 4, the Transport layer, and encapsulated. It is then referred to as a Layer 4 PDU, or *segment.* The segment is handed down to the OSI model Layer 3, the Network Layer, where it is also encapsulated and referred to as a Layer 3 PDU, or *packet*. The packet is handed to the OSI model Layer 2 where the encapsulation process is again repeated and the resulting structure is called a Layer 2 PDU or *frame*. Finally at Layer 1 the encapsulation sequence is converted to *bits* for transmission.

1.1.4 Network Recovery and OSI Model

Recovery occurs only at the bottom Layer 1 and at Layer 4 of each. Optical technologies—including Synchronous Optical Network (SONET), Wavelength Division Multiplexing (WDM) and the Optical Transport Network (OTN), an improvement on conventional WDM—all operate at Layer 1. Recovery is typically accomplished at this layer due to the simplicity and speed by which it can be implemented. However, recovery at Layer 1 affects large volumes of traffic when, for example, traffic is switched from working to spare fiber capacity when a working fiber fails. It is also expensive, requiring at least as much spare capacity as working capacity to be accomplished.

Alternatively, recovery can be accomplished at Layer 4 at the granular level of individual traffic streams. Recovery at Layer 4 is accomplished by requesting retransmission of failed traffic and usually implemented using the Transmission Control Protocol (TCP). The advantage is it works at a granular level and is inexpensive. Layer 4 recovery is limited in scale and is time-consuming, making it impractical for large outages and useless for failed real-time traffic.

However, there is a recent technology development that combines the advantages of many of the layers including Layer 1 and Layer 4. Multiprotocol Label Switching (MPLS) has successfully integrated rapid recovery at a granular, individual traffic stream level. MPLS proponents also claim large cost reductions over traditional Layer 1 service-provider recovery. Unfortunately, as with other recovery methods, there is little conclusive MPLS network-design information available to guide service providers.

This praxis analyzes the cost of conventional Layer 1 recovery methods to MPLS-based methods. Details of this summary are presented below.

1.2 Introduction

Telecommunications service providers are mandated by the Federal Communications Commission (FCC) to provide reliable public networks(Network Reliability and Interoperability Council V1, 2003). Network failures are defined as either complete loss of link connectivity, such as a cut fiber, or as individual streams of data within links. They are impossible to prevent but can and must be recovered from. However, failure recovery requires expensive additional capacity and equipment be added to network infrastructures. Failure recovery is typically accomplished in two ways: protection or restoration.

"Protection" means allocating network capacity to recover from failures before they occur. For several decades network recovery has been accomplished primarily by switching all traffic on to reserved spare capacity on separate links, or "spans," so this method is often called *span protection*.[1] Span protection is accomplished locally by switching demands, or traffic, onto spare unused spans at the nodes directly on either side of the failure.[2] The overwhelming advantage of span protection is it can restore traffic quickly, generally within 50 milliseconds after a failure occurs, and has become the de facto industry standard for recovery speed (Figure 1). A disadvantage of span protection is that it is expensive, generally requiring at least twice the working traffic capacity to be accomplished. While real-time services such as video and VoIP require 50-millisecond- recovery capability, many less-critical services do not (Alvarez, 2001). Therefore, span protection may be unnecessary in every part of service-provider networks.

Another recovery method, "restoration", determines and enables alternate traffic routing after failures occur. This recovery method is granular in that it can restore either individual IP data streams or all traffic on a link and is accomplished remotely from the failure at traffic origination and destination endpoint nodes. It is most commonly implemented as "path

---

[1] The terms "link," "line," and "span" are used interchangeably in the industry and in this praxis.
[2] The terms "traffic" and "demands" are used interchangeably in this praxis.

restoration," which deploys an alternative path or paths for the traffic affected by the failure. While path restoration may be more efficient than span protection, it may not be fast enough because of the time-consuming signaling and connection setup required after a failure occurs.

	Detection	Recovery
OTN/WDM	1-10 milliseconds	10-30 milliseconds
SONET	.1 milliseconds	50 milliseconds
TCP-based	40 seconds	1-10 seconds

Figure 1: Recovery time by technology (Dirceu, June 2000)

While service providers face intense pressure to provide reliable networks, consumers continually demand more advanced services at lower prices. Implementing new advanced services is more easily accomplished using Multiprotocol Label Switching (Morrow & Sayeed, 2007) . MPLS can be applied to many networking technologies but is presently implemented only in IP networks.

MPLS provides three major advantages to service providers. First, it enables the convergence of voice, data and video over a single IP network. Second, it also enables differing Quality of Service levels for different services. In other words, time-sensitive services, like voice and video, are given priority over less-critical services, such as email and web searches. Finally, MPLS also allows for traffic engineering, which controls traffic flow to achieve performance and resource optimization.

MPLS also provides enhanced recovery capabilities. MPLS Fast Reroute (FRR) is primarily based on conventional span restoration methodologies. As in conventional span protection, FRR restoration can generally be accomplished in less than 50 milliseconds. However, like conventional path restoration, MPLS span protection is accomplished at a granular level by recovering only individually affected IP traffic streams (if needed) rather than all traffic on the span. It also shares capacity between different possible failures scenarios, which reduces the required restoration capacity required by conventional span protection. MPLS can also implement path restoration capabilities and, as in MPLS FRR, shares restoration capacity.

There is long-held—but not conclusively proven opinion by service providers—that path restoration is more efficient that span protection, but slower. Attempting to achieve efficiencies of path-restoration efficiencies with span-protection speed, many service providers implement a hybrid recovery strategy. One form of *hybrid recovery* first applies span protection to quickly restore traffic while more time-consuming path-restoration paths are calculated and set up. When path restoration preparation is concluded, the traffic originally restored onto the span protected capacity is moved to path-restoration paths. The goal of hybrid recovery is to accomplish path restoration capacity efficiency with the speed of span protection. Hybrid recovery by applying span protection and path restoration sequentially has yet to be proven more efficient than path restoration alone.

1.3 Motivation

A critical area of concern and investigation by service providers is restoration cost. Even though restoration methods have been analyzed for many years, there is no conclusive research that shows, under varying conditions and constraints, which method is most effective. There is only the telecommunications-industry folklore that path restoration is more cost-efficient that span protection. The advantage of one method over the other is further blurred when considering shared or joint-capacity designs, MPLS recovery modifications, or the use of hybrid recovery methods that use span protection and path restoration. As a result, the most common service-provider practice is to over-provision, or intentionally adds extra capacity to networks in the hope it is sufficient to safely accomplish recovery. For example, one service provider does not use their spans beyond 40% of their capacity(Award Solutions, 2008).

The motivation for this praxis is to provide service providers with well-researched guidance in designing network recovery capacity so as to reduce escalating restoration costs, which currently tally in the billions of dollars. An analysis of the effectiveness and efficiency of span, path, MPLS FRR, and hybrid restoration methods are presented to accomplish that goal.

1.4 Network Recovery Overview

As with other networking function, recovery responsibility is delegated to the bottom four layers of the OSI model. A description of these responsibilities, their implementation in the OSI layers and the methods used to implement them are discussed below.

1.4.1 OSI Layer 1: Synchronous Optical Networks (SONET)

Virtually all telecommunications networks today use Synchronous Optical Network (SONET) technology. SONET operates at Layer 1 of the OSI model. As the name implies, SONET networks are built on fiber optics and transfer bits using light. Data from the layers above is given to SONET which creates the SONET frame and converts the bit stream to optical. SONET is a vast improvement over earlier technologies. Since it is synchronous, embedded multiplexed signals can be added and removed easily (Black & Waters, 2002). It is also very fast and, when coupled with Wave Division Multiplexing (WDM), is capable of transferring billions of bits per second on a single fiber optic cable.

One of the greatest advantages of SONET is the rapid protection of large amounts of network traffic. It is accomplished in SONET using span protection or path restoration and implemented in point-to-point, ring, or mesh topologies by duplicating the capacity of the carried traffic. However, protection and restoration are different and often confused.

1.4.1.1 Protection versus Restoration

Protection is pre-planned. It is accomplished first by having a pool of spare capacity available when a failure occurs. All necessary cross-connections or other mechanisms necessary to switch working traffic to spare capacity are implemented *before* the failure occurs. Restoration is also designed to provide for a pool of capacity to recover from failures but all recovery paths and necessary cross-connections are made *after* the failure occurs.

Protection is accomplished quickly while restoration takes considerably more time. However, from a network-capacity design perspective, protection and restoration result in identical capacity requirements. The differences are in how the recovery is accessed and accomplished.

1.4.1.2 Layer 1: Span Protection

The most prevalent form of Layer 1 network recovery used in telecommunications networks today is span protection, also called link or line protection. Simple in concept, span protection provides enough capacity to overcome any single failure in the network. It is implemented on a span-by-span basis in point-to-point, ring, and mesh topologies. The physical capacity is simply duplicated between all connected nodes in a network by adding new spans with sufficient “spare” capacity. The spans between the nodes may be physically diverse so a single event, such as a fiber cut, does not disable both spans.

Span protection is accomplished locally. In other words, when a span failure occurs, the span’s endpoint nodes divert all traffic to spare spans connecting the node pair. A common misconception is that traffic restored using span protection can only be accomplished with a single backup span. In reality span restoration can be demand “granular”— demands can be restored with any available spans linking the failure’s endpoint nodes. However, it is most commonly implemented using a single backup path for all demands, switching all traffic onto spare spans en masse, because it is simple to implement.

A major disadvantage of span protection is it requires doubling the network capacity, since spare capacity must equal or exceed the amount of working traffic carried on the network. However, the primary advantage of span protection is it can be accomplished quickly, generally in less than the 50 milliseconds required for critical and real-time services.

1.4.1.3 Layer 1: Path Restoration

Another method of network recovery is path restoration, which can be implemented point-to-point, ring, and mesh topologies. Functionally, path restoration involves the coordination of many network elements, some local to, and some remote from, the failure. The process requires that failure notification signaling must be transmitted to the origin, destination, and intermediate nodes of each affected traffic path after a failure occurs. Then, alternate paths must be calculated by each demand’s end nodes (or an offline system) and necessary intermediate network connections made(Placeholder1). Unfortunately, these steps can rarely be accomplished quickly, especially when networks are large.

Path restoration is believed to require less capacity than span protection by many service provider network engineering groups because individual demands are rerouted in a distributed manner across the entire network. Theoretically each restored demand could take a different path across the network requiring less capacity than span protection. However, this assertion has not been conclusively proven under varying circumstances and constraints.

An addition to path-restoration design is called "stub release," whereby capacity on paths unaffected by failures can be reused in the calculation of new paths(Placeholder2).

1.4.1.4 Layer 1: Hybrid Recovery and Design Enhancements

Service providers sometimes attempt to combine the speed of span protection with the perceived efficiency of path restoration by using hybrid recovery methods. Hybrid recovery can be accomplished in two ways. First, service providers sometimes implement span protection and path restoration separately in different parts of their networks. Span protection is used where restoration speed is paramount and path restoration is implemented for non-critical services where the delay necessary to calculate new paths and signal new connections is acceptable. Second, hybrid recovery can also be implemented by sequentially applying span protection and path restoration. Span protection is accomplished first for speed; later, after paths are calculated and established, the traffic is switched again based on path protection. The goal is to complete restoration within 50 milliseconds using span protection and then achieve system efficiency by switching to path restoration. However, the efficiency of hybrid recovery over span protection or path restoration alone has not been documented.

In addition to hybrid schemes, engineers have developed two design techniques to enhance both span protection and path restoration. The first is shared capacity which analyzes multiple failures and seeks to avoid duplicated protection capacity for failed common spans of different failure scenarios. The second is joint capacity design whereby the working and protection/restoration capacity are optimized simultaneously rather than individually. Both methods have been shown to be more efficient than without their use (Grover, Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking, 2003).

1.4.2 Layer 2: Ethernet, Frame Relay and ATM

Protocols that operate at Layer 2, the Data Link layer, of the OSI model include Ethernet, Frame Relay, and Asynchronous Transfer Mode (ATM). All are still popular technologies in service provider networks.

Layer 2 is “connection-oriented” in that a confirmed connection is established between the end clients before data is exchanged. Layer 2 detects failures by using a Frame Check Sequence (FCS) algorithm which calculates a checksum or validation number based on the data frame contents, and appends the number to the frame prior to transmission. When the frame is received, the arriving data is once again processed using the same algorithm. If the results match the number in the received FCS, then the frame is considered error-free. If errors are detected, Layer 2 simply drops the frame and no recovery of data is attempted. This responsibility is delegated, as will be shown, to Layer 4.

1.4.3 Layer 3: Internet Protocol (IP)

Since its inception, the public Internet has used the Internet Protocol (IP) for transmitting data. However, with the introduction of Voice over IP (VoIP) technology, IP is now also used extensively in service-provider networks and has become the standard for all modern networking. IP and other related protocols are embedded in network equipment such as routers.

Layer 3 is responsible for routing data across networks and is accomplished using several methods. However, most commonly used is the Open Shortest Path First (OSPF) protocol (Moy, 1998). By using OSPF, each router creates its own view of the network topology by exchanging information with all other routers in the network. Each router then creates its own routing table of the shortest route to the other routers. IP is "connectionless" because no end-to-end connection is set up before data is transferred. Data is encapsulated in packets and forwarded to the destination IP address on a hop-by-hop basis until the packet is delivered. In other words, at each router each packet is disassembled, the destination address examined and compared to the routing table to determine the next hop route, and then forwarded on.

IP is a "best effort" protocol. Packets may arrive successfully or they may not. If a transmission is unsuccessful, IP provides no response indicating a transmission failure. However, IP does have a mechanism to reroute subsequent packets around failures in the network. If a failure occurs, network routers on both sides of the failure detect the failure and trigger OSPF to update all router routing tables and a new path is selected for packets to follow. This process is effective but slow and cumbersome. Data may be lost while the router topology and routing tables "converge", or complete the update process. However, Cisco Networks and other network equipment vendors have developed proprietary methods to speed up the convergence and re-routing processes. Meanwhile, packets lost before the OSPF update process is completed are recovered via retransmission using Layer 4 mechanisms (Graziani & Johnson, 2009).

Native IP, or IP service not enhanced by proprietary vendor additions, has many disadvantages. The first is it does not use network resources efficiently. All packets are forwarded using the shortest path. If there is more traffic than can be carried over the shortest path it will be dropped even though other routes have abundant capacity. IP does not consider traffic load or congestion in deciding whether to forward packets but just blindly does so. Also, IP has very little ability to provide for different Classes of Service (CoS) with related Quality of Service (QoS) levels. Without such capability voice, video, and other real-time applications cannot be accomplished consistently. Even with vendor enhancements, IP requires extensive maintenance to force packets across underutilized spans in the network.

1.4.4 Layer 4: Transmission Control Protocol (TCP)

As discussed above, Layer 2 and Layer 3 depend on other Layer 4 to recover lost data by using the Transmission Control Protocol (TCP). TCP is used to number packets at the sender. Frames or packets may get lost or be corrupted. TCP is used at the receiver to detect the failure and request the data be retransmitted from the sender. TCP is effective for recovering through retransmission occasional lost frames or packets, but not for bulk recovery of data as would occur if, for example, if a fiber were cut. In that case, Layer 1 SONET recovery must be initiated.

1.4.5 Legacy Recovery Limitations

As discussed, Layer 4 has functionality to recover small amounts of lost data through retransmission. SONET technology has traditionally been the only method to restore large amounts of data en masse as the result of major event such as a failed span. The need existed for recovery methods that could:

1. Be accomplished quickly at Layer 2 at a granular demand level with various QoS levels,
2. Provide connection-oriented capability to connectionless Layer 3 IP streams, and
3. Be accomplished at the speed of Layer 1 SONET.

Multiprotocol Label Switching (MPLS) filled these needs. Therefore, it is often referred to as a “Layer 2 ½” technology.

1.4.6 MPLS Overview

MPLS is an evolutionary technology that is rapidly being deployed in service- provider networks worldwide (Minei & Lucek, 2008). As the name implies, it is applicable to many protocols, but is primarily being applied to IP, where it enhances IP without replacing it. In fact, IP and IP/MPLS can coexist simultaneously in the same network. MPLS enhances IP by appending to each IP packet a “shim” header, which includes an MPLS label and QoS designator.

1.4.6.1 MPLS Implementation

A key element of MPLS is the Forwarding Equivalency Class (FEC). A FEC is an aggregate of traffic between two routers that requires similar treatment. Packets with the same destination, QoS, and performance requirements are part of the same FEC. This FEC concept is significant because instead of switching individual packets across the network, aggregate streams of packets with similar characteristics are switched together, which reduces the workload on routers and allows dynamic routing of packets based on QoS, span congestion, and other requirements.

The packets in FECs are implemented as a Label Switched Paths (LSPs). As IP packets enter the IP/MPLS network they do so at a Label Edge router (LER), where the packets are assigned to a FEC, then labeled, and forwarded across an LSP. MPLS/IP routers in the core of the network are called Label Switch Routers (LSRs) and forward the aggregate LSP streams across the network by calculating the end-to-end paths using a Constrained Open Shortest Path First (COSPF) algorithm. COSPF is a significant advantage because overcomes the limitation of conventional OSPF by considering QoS parameters and other metrics.

MPLS provides significant advantages to IP besides more-effective routing for packets. In general, it adds Layer 2 functionality to IP by providing connection-oriented capability. As a result, connections are verified before data is transferred and lost packets are recovered via retransmission over LSPs. It also provides for QoS capabilities to IP streams so that VoIP, video, and other real-time applications are given delivery with priority over less-time-sensitive applications. It is also capable of encapsulating and transporting any Layer 2 protocol with the advantages discussed. Finally, and most significant for this praxis, MPLS provides additional recovery capabilities.

1.4.6.2 MPLS Recovery Mechanisms

MPLS can recover from either span or node failures. However, since node failures are far less common than span failures (Grover, Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking, 2003), this praxis addresses only span failures.

MPLS recovery can be accomplished in two ways. The first is with span protection via MPLS Fast Reroute (FRR), which itself can be implemented in two ways: Facility and 1:1 (Vassuer, Pickavet, & Demeester, 2004). In MPLS Facility FRR, protection acts like the basic span protection most commonly implemented by service providers whereby all traffic is detoured across a single backup path. However, it can be accomplished with per-LSP granularity; all demands riding the span do not have to be rerouted. In MPLS 1:1 FRR, individual IP streams bypass the failure across <u>any</u> path that is available between the two nodes on either side of the failure. Both MPLS Facility and 1:1 FRR require that protection paths be pre-determined and pre-provisioned.

One significant advantage of MPLS FRR is that traffic can be restored at Layer 2 (at a demand granular level if desired) in less than 50 milliseconds. By providing Layer 1 SONET restoration speeds, this has the potential of saving service providers billions of dollars in recovery capacity costs.

Recovery can be accomplished using traditional path restoration. However, failure detection, restoration path calculation and circuit set-up requires far too much time to be effective. MPLS has also implemented path ***protection*** where recovery paths are predetermined and set up in advance in anticipation of failures; hence it uses path protection rather than path restoration. Path protection is much faster than path restoration, but still may be too slow on long paths.

Hybrid MPLS recovery can be implemented using FRR and path restoration or protection sequentially. First, FRR is accomplished for rapid recovery. When this is complete, MPLS path restoration moves traffic from the FRR paths to the path-restoration LSPs. As in SONET, the conventional opinion is that path restoration is more efficient that FRR. However, this assertion has not been conclusively prove correct.

1.5 Drivers for Efficient Network Recovery

Telecommunication is an extremely competitive and capital-intensive business. Billions of dollars are spent on equipment and facilities as well as the personnel to manage these massive networks. Service providers must find way to be innovative and efficient to survive in this volatile environment. The need to understand how to effectively and efficiently apply recovery methods is being influenced by both external and internal drivers.

1.5.1 External Drivers

External drivers include consumers and government regulators. Consumers dominate service-provider product strategies and constantly demand more advanced and higher-bandwidth services at lower prices. Today, different services require Quality of Service requirements. Voice over IP, video, and other time-sensitive services must be recovered quickly, whereas web searches and other less critical services have less stringent recovery needs.

Presently, service providers predominantly use SONET span protection to accomplish rapid, all-or-none recovery which for some services is excessive or unnecessary. At the same time, the Federal Communications Commission mandates recovery requirements for different services(Placeholder3). Failure to comply can have costly, business affecting consequences.

1.5.2 Internal Drivers

Service-provider network-design engineers constantly struggle to create networks that meet recovery requirements while attempting to minimize equipment and facilities costs. However, conclusive research that demonstrates how different recovery methods operate under various conditions and constraints is scarce. As a result, network designers overbuild, adding excessive capacity to their networks to ensure effective recovery. Another strategy is to combine different hybrid recovery methods. For instance, using span protection first followed by path protection, to reduce overbuild costs. However, this strategy may have the opposite effect of increasing recovery costs. In addition, the overwhelming benefits of MPLS and more advanced recovery methods have rapidly accelerated global deployment. However, this technology has complicated the task of network engineering and increased the uncertainty of their design results.

Once network services are provisioned, or installed, on network capacity they are turned over to the service-provider network-management function. Historically, managing networks requires enormous computing and human resources. Today, with many different services, varied

QoS requirements, and multiple recovery methods in use, the process is further complicated. In addition, MPLS provides for granular, per-IP-stream recovery but this often results in additional layers of recovery to manage, since previous methods may remain. For instance, MPLS FRR may be implemented but, because of service provider uncertainty of its effectiveness and efficiency, SONET may be left in place as well.

Service-provider executive managers realize that network recovery has far- reaching effects on network costs. With billions of dollars spent on equipment, facilities, and personnel, they are under constant pressure to reduce costs, increase revenue, and improve their competitive position. They are without doubt highly motivated to benefit from the enormous savings possible from efficient recovery strategies.

1.6 Example of Problem

The network in Figure 1 illustrates how span protection, path restoration, Fast Reroute, and hybrid recovery are accomplished and the differences between each method. The network consists of eight nodes. Nodes 1 through 4 are access, or edge, nodes that originate and terminate traffic. Nodes 5 through 8 are core, or backbone, nodes that only switch traffic. All spans are bidirectional and uncapacitated. Although not shown, each span duplicates the working capacity with restoration capacity.

For this example, bidirectional traffic is being exchanged between nodes 1 and 3. Data is routed along the path 1-7-8-3 when a complete failure occurs between nodes 7 and 8.

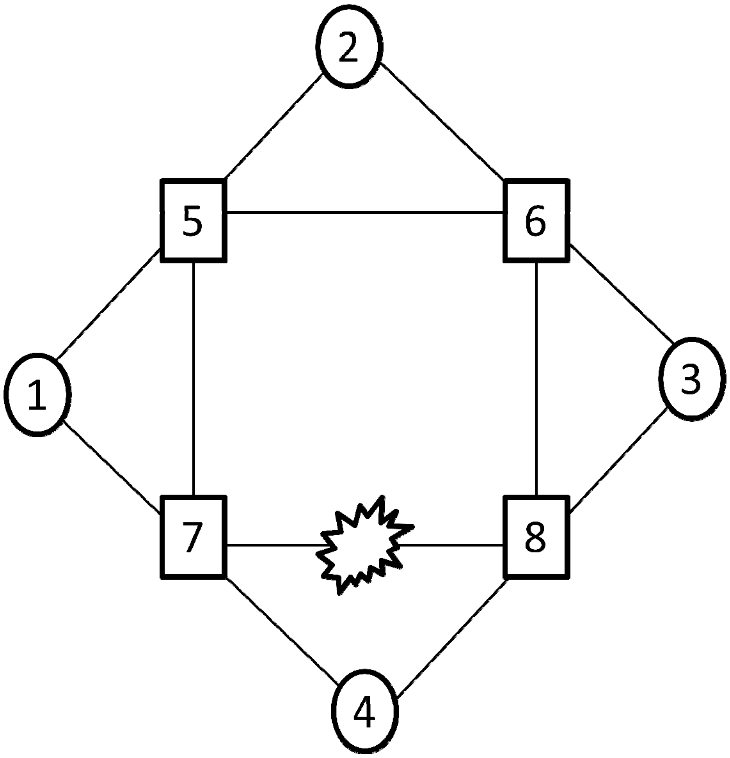

Figure 2: Example failure

1.6.1 Span Protection

With span protection, calculation of recovery routes is accomplished in advance. The required capacity is set aside until the failure occurs at which point all traffic carried is moved to the predetermined and preconfigured paths. Therefore, restoration is fast, usually in less than 50 milliseconds.

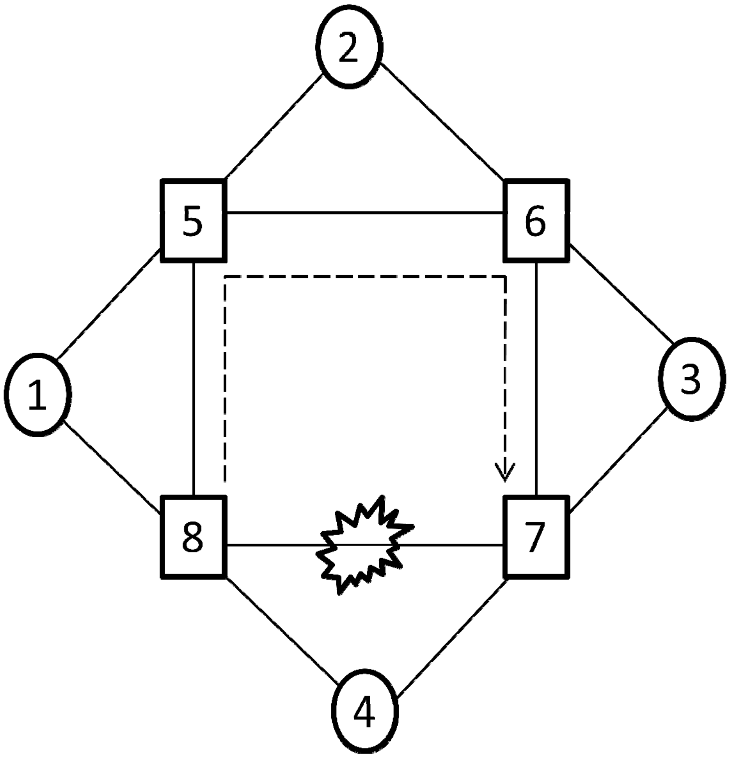

Figure 3: Span failure

Figure 3 illustrates how span protection operates. Span protection occurs locally. Therefore, the core nodes on both sides of the failure sense and initiate recovery and must be included in the restoration route. For example, for traffic between node 1 and node 3, when span 7-8 fails, nodes 7 and 8 both quickly sense the loss of data and switch all working capacity onto the spare capacity via paths 8-5-6-7 or 8-4-7. All traffic between nodes 1 and 3 now follows the either path 1-8-5-6-7-3 or path 1-8-4-7-3 or could be split between the two paths.

Which route or routes will be selected depends on the route selection algorithm used. The shortest path, least hops, bandwidth, delay, cost, and reliability are all metrics used by different public and privately developed routing protocols. In addition, network design policy may also determine which recovery route is selected. In this example path 8-4-7 would typically not be allowed to participate in the recovery of data from failed span 7-8 because the span capacities connecting the edge to core nodes generally have much less capacity than core spans and therefore may be insufficient to recover the entire capacity of span 7-8 via node 4.

1.6.2 Path Restoration and Protection

Path restoration operates much differently than span protection. Using the same example, bidirectional traffic is connected between access nodes 1 and 3 when the span between core nodes 7 and 8 fails. The recovery does not originate locally at nodes 7 and 8 but at the edge access nodes where the traffic enters and leaves the network. In this case, core nodes 7 and 8 signal access nodes 1 and 3 that the failure has occurred. Access nodes 1 and 3 then must calculate an alternate path (or paths) to restore the traffic on. If only the core nodes are allowed to be used for restoration, the restoration path established will be 1-5-6-3. However, path restoration is granular, meaning individual traffic demands can be routed separately across different paths. Therefore, access nodes 2 and 4 might participate in the restoration of part of the traffic even though their spans have less capacity than core-router spans.

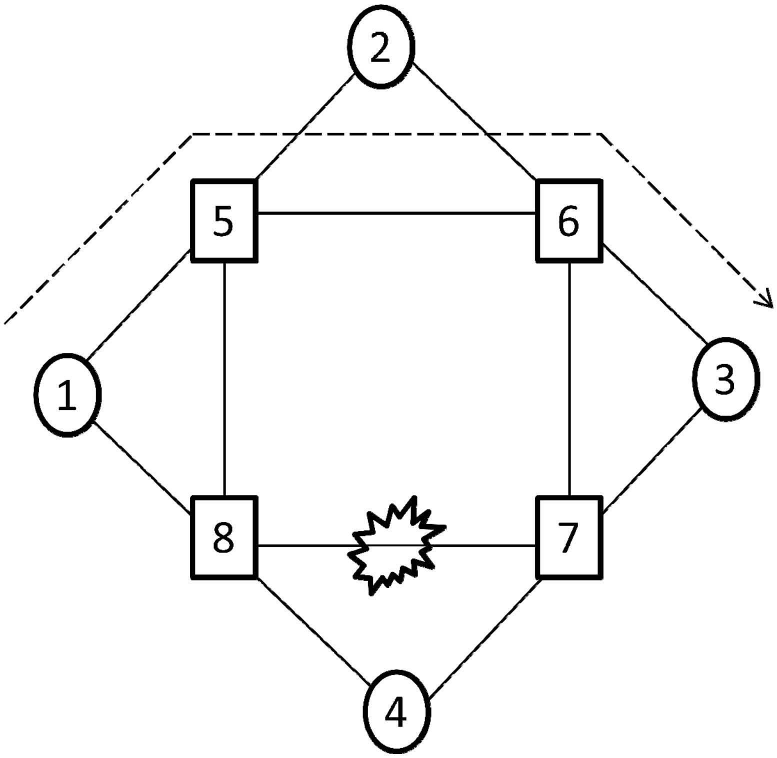

Figure 4: Path recovery

As a result many other paths are available including the following: 1-5-2-6-7-3, 1-8-5-6-3, or any and all combinations of routes between access nodes 1 and 3. However, all traffic could be forced on a single backup path or divided among all available paths. As with span protection, the path or paths taken depends on network management policy as well as one or more routing protocol metrics such as shortest path, most available bandwidth and others. However, path restoration is usually accomplished on a single backup path due to the implementation complexity of doing so over multiple paths.

The primary disadvantage of path restoration is that the failures must be detected and signaled back to the origination and destination nodes of the traffic and restoration paths determined after the failure occurs. Therefore, path restoration is time-consuming and can seldom be completed in less than 50 milliseconds in networks of any significant size, particularly if the number of traffic demands is high. MPLS has the ability to provide path *protection* where recovery path are predetermined and set up in advance. This significantly speeds up the path recovery process. However, it still may too slow when longer paths are involved.

Figure 5: Recovery sequence

The ability of path recovery methods to distribute restored traffic across many nodes, spans, and paths has, perhaps erroneously, led service providers to believe that path restoration is more efficient than span protection. However, this assertion has not been conclusively proven for different scenarios and constraints.

1.6.3 Hybrid Recovery

Service providers may implement both span protection and path recovery in conventional or MPLS enables networks. Hybrid recovery methods exist because service providers want the speed of span protection and the perceived efficiency of path restoration or protection.

As Figure 6 demonstrates, in order to provide enough capacity for restoration, hybrid recovery requires that each span have restoration capacity of the maximum capacity required by either method. However, the capacity required using hybrid recovery has not been conclusively demonstrated or proven. Also, using hybrid recovery methods may create a computationally intensive problem that is difficult to solve in a timely manner.

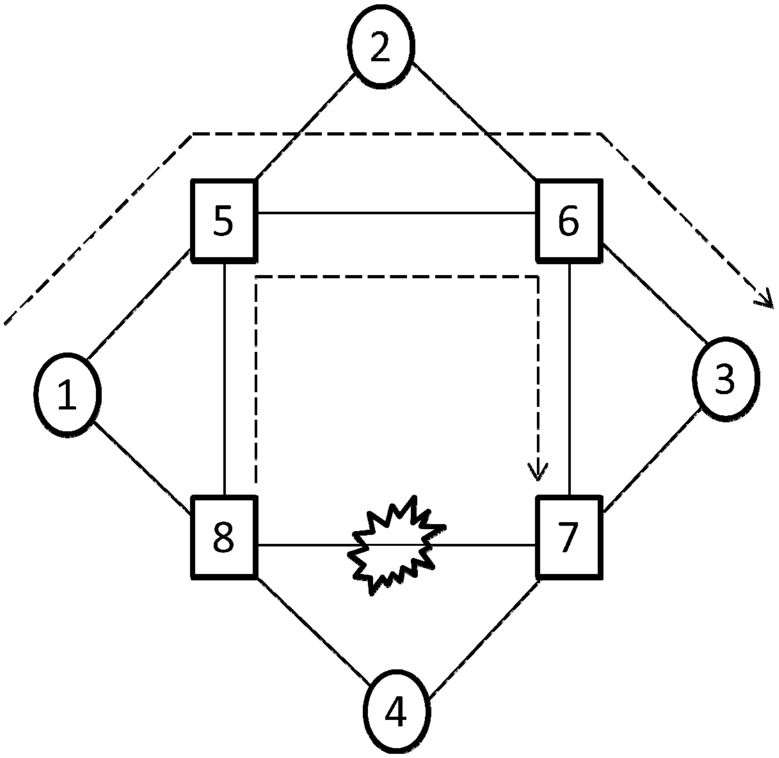

Figure 6: Hybrid recovery

1.7 Praxis Overview

This praxis addresses the modeling, analysis, and optimization of span, path, Fast Reroute, and hybrid recovery methods used in modern telecommunications networks. Specifically, it determines how much capacity is required to carry working and restoration traffic in the event of the failure of any span in the network.

Libraries of realistic networks are analyzed. The physical network topologies are defined in terms of nodes, spans and span costs. All spans are uncapacitated but the results are also modularized to reflect typically used span capacity sizes. Also given are traffic demand matrixes that are varied by number and size of demands.

Chapter 2 formulates recovery solutions using four methods: span protection, path restoration, Fast Reroute, and hybrid recovery. The models minimize the cost of working and spare capacity. The basic models are later enhanced by applying span modularity; instead of leaving spans uncapacitated, actual link capacities used by service providers are applied. Chapter 3 describes the hypotheses tested and the computational experiments conducted. Presented is: the overall approach and test networks, the factors and levels explored, the performance criteria, and the results of the computational experiments. Chapter 4 shows the results of integrated analysis over different factor levels. Chapter 5 is an analytical summary of the experimental results. Chapter 6 shows the cost effects of link density, demand size, and restoration migrations. Chapter 7 presents recommendations to service providers and network managers and Chapter 8 identifies possibilities for future research.

1.7.1 Approach and Methodology

Network flow models and algorithms provide the theoretical foundation for this praxis. The primary area of importance is the calculation and comparison of working and spare capacity and related costs for several network recovery methods.

Several models using Linear Programming (LP) and Mixed Integer Programming (MIP) are demonstrated. The models are implemented using the Generalized Algebraic Modeling System (GAMS) mathematical modeling language and solved using the CPLEX optimization software package. The models are applied to a suite of networks that represents realistic service provider scenarios. The computational experiment results are evaluated using statistical Analysis of Variance (ANOVA).

1.7.2 Significance of Contributions

Comprehensive and conclusive research has not been accomplished that compares the efficiency of conventional span, path or hybrid recovery combinations. Service provider uncertainty is further increased with the widespread acceptance of Multiprotocol Label Switching (MPLS) and related Fast Reroute (FRR) restoration.

Therefore, the praxis results will have a far-reaching impact on the telecommunications industry at many levels. Network design engineers will be able to understand recovery methods and how they operate under different conditions and therefore more efficiently design their network. Network management will be simplified by eliminating unnecessary restoration technology duplications. It also provides a roadmap to increased service provider profitability by dramatically reducing network costs by decreasing required network transmission facilities and equipment. Equipment manufacturers will also benefit by being able to design or modify their equipment to implement the most efficient and effective recovery methods under varying circumstances.

2 Formulation and Benefit Evaluation of Optimization Models for Network Recovery Design

2.1 Problem Statement

This praxis addresses the problem of engineering and analyzing telecommunications networks that can recover from a single span failure. The objective is to minimize system cost through optimal allocation and use of spare recovery capacity.

In this study, the physical topology and link attributes of networks are given. The links are bidirectional, uncapacitated, and have costs that are directly related to their length. Also given are traffic matrixes, whose elements represent individual working demands between network node pairs. Several potential connection paths are generated for each such demand.

A series of optimization models are presented for different design assumptions. In each case, the objective is to assign spare capacity to a network's spans at minimum cost while ensuring that the system can recover from a failure of any individual span by rerouting all affected demands. In some instances, spare capacity may be shared by multiple rerouted demands. Link modularity sizes of 10Gbps and 40Gbps are later applied and analyzed.

2.1.1 Survey of Literature

While the research in this area does not compare the different technical approaches (nor does what is done herein) this section summarizes the published work for these types of recovery schemes:

- Span protection only
- Path restoration only

- Span protection and Path restoration
- MPLS Fast Reroute, and
- Internet Engineering Task Force (IETF) Requests for Comments (RFC’s).

2.1.1.1 Span Protection

Dunn et al. assert that using disjoint successive k-shortest paths (KSPs) for span protection is faster and easier than other methods, such as maximum flow algorithms (Dunn, Grover and MacGregor 1994). The primary finding is that solutions using KSP restoration are nearly equal to maximum flow in typical network models.

Grover et al. finds that span restoration based on KSP can be accomplished with between 50% and 70% redundancy (Grover, Bilodeaux and Venables, Near optimal spare capacity planning in a mesh restorable network 1991). Their approach uses successive k-shortest paths, considers span modularity, and limits path lengths.

Venables et al. consider two heuristic strategies for the placement of spare capacity in span-protected networks (Venables, Grover and MacGregor 1993). The first is the Spare Link Placement Algorithm (SLPA) and is based on the principle of iterative span addition to produce the greatest incremental change in network restorability. The other, the Iterated Cutsets Heuristic (ICH) formulates spare capacity placement as a linear programming problem subject to constraints based on a subset of cutsets of the network. Iteration and heuristic rules are used to develop the constraint set required by ICH. The results indicate that the ICH method is approximately five percent more efficient than the SLPA method.

2.1.1.2 Path Restoration

Kennington et al. analyze spare capacity for mesh networks using path restoration while considering modular span capacities and single span failures (Kennington, Nair and Spiride, Optimal spare capacity assignment for path restorable mesh networks: cuts, decomposition, and an empirical analysis 1998). The authors develop a branch-and-cut-approach based on an arc-flow formulation for the failure routing. They report span restoration requires on average 12% more spare capacity than path restoration without stub release.

Menth et al. propose and model using linear programming an end-to-end (path) restoration mechanism that uses multipath routing and load balancing (Menth, Martin, et al. 2009). They found that the structure of the traffic matrix, selection of the primary paths as well as network topology have significant effects on network restoration capacity. However, the size of the network did not. Also, the benefit of multiple paths for working and spare capacity was maximized at three paths (one primary and two spare). Their approach showed as little as 17% extra capacity could restore all working capacity.

Li et al. propose an algorithm for shared path protection that is applicable to a variety of network technologies, including MPLS (Li, Kalmanek and Doverspike 2002). The results are compared to non-shared shortest-path approach and another shared shortest-path restoration method. The measure of effectiveness is the restorations overbuild, or additional capacity, to protect working traffic. Their results showed that that their shared-path restoration algorithm is up to 22% more efficient than the other two methods.

2.1.1.3 Span Protection and Path Restoration

Doucette et al. present what they claim to be the first comprehensive study of mesh network restoration capacity across a range of node degrees (Doucette and Grover 2001). Six types of restorable mesh network designs are compared: 1+1 non-shared backup protection, span restoration with and without joint optimization of working and spare paths, chain-optimized ("meta-mesh") span restoration, Shared Backup Path Restoration (SBPP), and true dynamic path restoration. Model formulations for each of the six designs reference previous works of the authors. Their results show that true path restoration outperforms all other methods. Span protection is at least 10% more expensive than path restoration with stub release, independent of network density. Meta-mesh and SBPP are almost equally efficient and both are almost as efficient as true path restoration. Joint capacity allocation is much more efficient than separate single capacity allocations. In general, total working and spare capacity costs decrease as node degree increases.

Murakami et al. address the optimal span capacity design of Asynchronous Transfer Mode (ATM) networks based on span protection and end-to-end (path) restoration methods (Murakami and Kim 1998). They formulate the problem as a large-scale linear program. A row generation and deletion mechanism is used to handle the explosive number of path possibilities with path restoration. Their study focuses on varying network topologies. The authors acknowledge the popular conception that path restoration is more capacity efficient than span protection and agree in many cases. However, their results show that there is little advantage to path restoration over span restoration when networks are well-connected physically or have an unbalanced demand matrix.

Orlowski et al. present a mixed-integer-programming (MIP) model for solving network restoration problems that integrates topology, hardware (routers, Digital Cross-Connect systems, various interface cards, etc.), span capacities, and restoration method simultaneously (Orlowski and Wessaly 2003). Based on a branch-and-cut algorithm with a column generation procedure (to cope with the large number of paths generated with path restoration), they compare the optimal network costs using span protection and path restoration with and without stub release. Single span and node failure are also compared. Their results showed that network cost is mostly related to other considered factors and independent of restoration method.

Xiong et al. compare path restoration without stub release and span protection under a single span failure scenario (Xiong and Mason 1997). Instead of using column generation to define paths, a set of hop-limited path variables are pre-calculated. With joint working and spare capacity optimization, they obtain almost the same network cost for span protection and path restoration.

Caenegem et al. suggest a simulated annealing algorithm to compute low-cost solutions to span protection and path restoration without stub release in single span failures for a given set of working paths (Caenegem, Wauters and Demeester 1997). In this limited study, span protection is found to be 20-25% more expensive than path restoration.

Doverspike et al. compare the efficiency of "patch" (span) protection to end-to-end (path) restoration in restoring failed working capacity on capacitated networks (Doverspike and Wilson 1994). To simplify the problem, they formulate a lower bound heuristic of six hops and an LP formulation for the upper bound and compare the two with weighted averages. They find that the two methods have a negligible performance gap and that path restoration is generally more efficient than span protection. They also observe that the efficiency of path restoration is greater for span failures than node failures. However, the efficiency difference between is negligible for path restoration over span protection at low levels of network traffic congestion.

Ambs et al. describe the improvement of the AT&T span protection strategy (Ambs, et al. 2000). Previously, span protection was implemented at AT&T using the FASTAR system. However, allowing span protection across multiple paths is considered by AT&T to be more efficient. AT&T researched and developed their Restnet system to incorporate multiple path span protection with a shared restoration capacity. The system is based on an arc-path linear-programming formulation. Paths are pre-selected based on a column-generating and path-pruning methodology. AT&T reported savings of 35% or more over the older FASTAR system.

2.1.1.4 MPLS Fast Reroute

Martin et al. calculate the capacity required to restore working paths based only on shortest path routing for Facility and 1:1 MPLS Fast Reroute restoration (FRR) and compare the results to end-to-end (path) restoration using Shortest Path Rerouting (SPR) (Martin and Menth 2006). Their results conclude that Facility and 1:1 FRR requires 47-86% and 21-26% more capacity respectively than SPR.

Menth et al. compare Single Shortest Path (SSP) and Shortest Multipath (SMP) [3] to Facility and 1:1 MPLS FRR facility restoration (Menth, Martin, et al. 2009). They report that the path-restoration efficiency under normal conditions was similar between SSP and FRR, but under failures that shortest path was much more efficient. MPLS 1:1 protection is more efficient than Facility backup. They also find that SMP path restoration was much more efficient than

---

[3] SMP is an end-to-end (path) restoration mechanism that allows traffic to be split among many different paths

either FRR Facility backup for span failures but not for node failures. SMP efficiency increases as network size increases whereas the efficiency for each MPLS mechanism does not. However, the authors do not develop any algorithms or other methods in this paper to explain how their results are achieved.

2.1.1.5 Internet Engineering Task Force

The Internet Engineering Task Force (IETF) is the authority for Internet protocol standards and has developed a series of Requests for Comment (RFC) related to recovery in MPLS and GMPLS networks. A summary of pertinent RFCs follows.

RFC 3031“MPLS Architecture” provides an overview of the MPLS network architecture. It defines the functions of MPLS-enabled Label Switched Routers (LSR’s) and describes the encoding, operation and distribution of labels (Rosen, Viswanathan and Callon 2001). It covers both control and traffic forwarding functions. Control involves segmenting traffic in Forwarding Equivalency Classes (FECs), assigning labels to FECs and distributing them among LSRs, and establishing Label Switched Paths (LSPs) on which FEC associated packets are forwarded. It also describes the integration of IP-based hop-by-hop routing as well as explicit and multipath routing.

RFC 2702 “Requirements for Traffic Engineering over MPLS” presents requirements for traffic engineering over MPLS networks and details the capabilities necessary to design efficient and reliable MPLS networks (Awduche, et al. 1999). Improved bandwidth utilization and operational performance are the intended results of MPLS-TE. The document describes the traffic trunk as an abstraction of aggregate traffic flows mapped to an LSP. MPLS-TE assigns constraints to traffic trunks, such as traffic rate and priority, and to resources including span bandwidth and resource class. MPLS-TE uses a Constraint-based Shortest Path First (CSPF) to select paths considering both trunk and resource constraints.

RFC 4090 “Fast Reroute Extensions to RSVP-TE for LSP Tunnels” defines the RSVP-TE extensions for local repair of LSP tunnels including bypass tunnels using facility and detour tunnels using 1:1 restoration methods (Pan, Swallow and Atlas 2005). The behavior of Point of Local Repair (PLR) and merge node behavior is described as well as procedures for backup path computation.

RFC 4427 “Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)” defines common terminology for GMPLS-based protection and restoration (Mannie and Papadimitriou 2006). RFC 4426 “Generalized Multi-Protocol Label Switching (GMPLS) Recovery Functional Specification” outlines the protocol extensions required to implement GMPLS recovery (protection and restoration) mechanisms (Lang, Rajagopalan and Papadimitriou 2006). Span protection, end-to-end (path) protection and restoration and shared mesh restoration methods are detailed.

RFC 4428 “Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-Based Recovery Mechanisms (including protection and restoration)” provides an analysis grid to evaluate and compare GMPLS transport plane recovery techniques (Papadimitriou and Mannie 2006). Different recovery phases are investigated. It focuses on transport and not the control plane survivability and recovery. Recovery resources sharing using Shared Risk Link Groups (SRLG’s) are discussed.

2.2 Mathematical Formulations

This section presents the mathematical formulations for span protection, path restoration and Fast Reroute recovery methods. All models are modified formulations from Pioro and Medhi (Placeholder1). Hybrid recovery, a calculated process, is also described.

2.2.1 Notation and Conventions

A *node* represents an instance of network equipment, such as a router. A *span* is a physical transmission line between nodes is represented by a single link but traffic flow is bidirectional. The term *span, link* and *line* are used interchangeably. All spans are assumed to be uncapacitated, however, module sizes commonly used in the industry are applied and compared to the uncapacitated results. The cost of a span is related to the distance between its two endpoint nodes. A demand is an instance of required capacity, or bidirectional data flow, between a pair of network nodes. The terms *traffic* and *demand(s)* are used interchangeably.

The following notations and conventions are used. The collection of demands forms the set D, and each individual demand $d \in D$ is routed across one or more spans or edges from the set of network spans *E*. Following a failure of span $e \in E$, the remaining operable spans form L_e , where $L_e = E \setminus \{e\}$ .Each demand *d* can be routed via the candidate paths in set P_d. Associated with each span $e \in E$ is a set of candidate restoration paths denoted by set Q_e. States result when a single span fails; S is the sum of all states while s is an instance of a single failed span. Modularity set M contains the possible link capacities, which are 10Gbps and 40Gbps in this study.

In the mathematical programming models that follow, the following symbols are used:

Constants

δ_{edp}	= 1 if span e belongs to path p realizing demand d ; 0, otherwise
h_d	volume on demand $d \in D$
$h_{d\,s}$	volume on demand d in state s
ξ_e	unit cost of span e
β_{eql}	= 1 if span l belongs to path q restoring span e; 0, otherwise
χ_{ds}	demand coefficient of demand d in state s, $h_{ds} = \chi_{ds} h_d$
α_{es}	fractional availability coefficient of span e in state s $(0 < \alpha_{es} < 1)$

Variables

x_{dp0}	normal flow allocated to path p of demand d
x_{dps}	normal flow allocated to path p of demand d in state s
y_e	normal capacity of span e
$y_e^{'}$	protection capacity of span e
$y_l^{'}$	protection capacity of span *not* e
z_{eq}	flow restoring normal capacity of span e on restoration path q

u_{eq} binary flow variable associated with z_{eq}

2.2.2 Span Protection

Span protection is accomplished when a span fails and all the capacity on the original (working) span is switched to another dedicated standby (protection) span. All traffic is restored on a single backup path. Typical span protection does not consider shared capacity. Span protection designs allocate enough capacity to recover if all spans failed simultaneously.

The span-protection problem can be modeled as the following mixed-integer programming problem.

2.2.3 Span Protection Formulation

$$\text{Minimize } F=\sum_{e\in E}\xi_e\,(y_e+y_e^{'}) \tag{2.2.2a}$$

subject to:

$$\sum_{p\in P_d} x_{dp0}=h_d,\forall d\in D \tag{2.2.2b}$$

$$\sum_{d\in D}\sum_{p\in P_d}\delta_{edp}x_{dp0}\le y_e,\forall e\in E \tag{2.2.2c}$$

$$\sum_{d\in D}\sum_{p\in P_d}\delta_{edp}x_{dp0}\le my_e,\forall e\in E, m\in MGbps \tag{2.2.2d}$$

$$\sum_{q\in Q} z_{eq}=y_e,\forall e\in E \tag{2.2.2e}$$

$$\sum_{q \in Q_e} u_{eq} = 1, \forall e \in E \qquad 2.2.2f$$

$$z_{eq} \leq u_{eq}, \forall e \in E, q \in Q_e \qquad 2.2.2g$$

$$\sum_{q \in Q_e} \beta_{eql} z_{eq} \leq y_l', \forall e \in E, l \in L_e \qquad 2.2.2h$$

2.2.3.1 Span Protection Formulation Explanation

The objective function (2.2.2a) minimizes the cost of network spans considering both working and restoration capacity needs. Constraints 2.2.2b and 2.2.2c assure that working span demands are carried only on normal span capacities. Constraint 2.2.2d is the same as 2.2.2c except it applies span modularity of 10Gbps and 40Gbps to be applied later. Constraints 2.2.2e and 2.2.2h assure that the working capacity of link e is recovered using only the protection capacity of the remaining spans *l*. Constraint 2.2.h does not assume shared capacity. Constraints 2.2.2f and 2.2.2g are equivalent to $z_{eq} = u_{eq} y_e$ which states that the flow restoring normal capacity of span e on restoration path q must equal the normal capacity of span e given that restoration span q is available as per u_{eq}, the binary availability variable of z_{eq}. However, the constraint is written in two steps because multiplying two variables is not permitted on the right-hand side of the MIP formulation.

2.2.4 Fast Reroute Facility Restoration

Fast Reroute Facility restoration is a form of span protection which restores all traffic from a failed span onto a single backup span. However, unlike conventional span protection, FRR Facility restoration does consider shared capacity. Fast Reroute Facility restoration is also modeled as a Mixed Integer Program (MIP) problem.

2.2.4.1 Fast Reroute Facility Restoration Formulation

$$\text{Minimize F}=\sum_{e\in E} e\xi_e\,(y_e + y_e^{'}) \tag*{2.2.3a}$$

subject to:

$$\sum_{p=P_d} x_{dp0} = h_d, \forall d \in D \tag*{2.2.3b}$$

$$\sum_{d\in D}\sum_{p\in P_d} \delta_{edp} x_{dp0} \le y_e, \forall e = E \tag*{2.2.3c}$$

$$\sum_{d\in D}\sum_{p\in P_d} \delta_{edp} x_{dp0} \le m y_e, \forall e \in E, m \in MGbps \tag*{2.2.3d}$$

$$\sum_{q\in Q} z_{eq} = y_e, \forall e \in E \tag*{2.2.3e}$$

$$\sum_{q\in Q_e} u_{eq} = 1, \forall e \in E \tag*{2.2.3f}$$

$$z_{eq} \le u_{eq}, \forall e \in E, q \in Q_e \tag*{2.2.3g}$$

$$\sum_q \beta_{leq} z_{eq} \le y_l^{'}, \forall l \ne E \qquad 2.2.3h$$

2.2.4.2 *Fast Reroute Facility Restoration Explanation*

The objective function 2.2.3a minimizes the cost of network spans considering both working and restoration capacity needs. Constraints 2.2.3b and 2.2.3c assure than working span demands are carried only on normal span capacities. Constraint 2.2.2d is the same as 2.2.3c except it applies span modularity of 10Gbps and 40Gbps to be applied later. Constraints 2.2.3e and 2.2.3h assure that the working capacity of e is recovered using only the protection capacity of the remaining spans *l*. Constraint 2.2.h does assume shared capacity. Constraints 2.2.3f and 2.2.3g are equivalent to $z_{eq} = u_{eq} y_e$ which states that the flow restoring normal capacity of span e on restoration path q must equal the normal capacity of span e given that restoration span q is available as per u_{eq}, the binary availability variable of z_{eq}. However, the constraint is written in two steps because multiplying two variables is not permitted on the right-hand side of the MIP formulation.

2.2.5 Path Restoration

Path restoration restores individual traffic flows, or demands, rather than entire span capacity en masse. Path restoration not only considers shared spare capacity, but can also use a procedure called commonly called “stub release” whereby the capacity released by the failed flows on those spans of the failed paths that survived the failure is reused. Since the objective of path restoration is often to re-optimize network capacity, this formulation of path restoration is unrestricted which allows connected flows to be moved when a failure occurs if it will result in a more optimum solution. Path restoration is an extension of the multi-commodity flow problem and is modeled as a Linear Program (LP).

2.2.5.1 Path Restoration Formulation

$$\text{Minimize F}=\sum_{e\in E} e\xi_e\,(y_e + y_e^{'}) \qquad (2.2.4a)$$

subject to:

$$\sum_{p=P_{dps}} x_{dps} = h_{ds}, \forall d \in D, s \in S \qquad (2.2.4b)$$

$$\sum_{d\in D}\sum_{p\in P_d} \delta_{dpe} x_{dps} \le \alpha_{es} y_e, \forall e \in E \qquad (2.2.4c)$$

$$\sum_{d\in D}\sum_{p\in P_d} \delta_{edp} x_{dps} \le \alpha_{es} m y_e, \forall e \in E, m \in MGbps \qquad (2.2.4d)$$

2.2.5.2 Path Restoration Formulation Explanation

The objective function 2.2.4a is straight-forward in that it attempts to minimize the cost of the network spans considering both working and restoration capacity needs. The first constraint 2.2.4b says that for all paths that the demands on the paths using available, non-failed spans must equal the volume of demand on the working spans. In other words, normal demand flows are carried using only normal span capacities.

One of the key considerations in the path restoration model is identifying the state of network spans and is identified by α_{es}. Each network span is assigned a span number and is paired with an equivalent state number. A failed span is designated by making the state number associated with a span number a “0”. For instance, S1 indicates that all network spans are available except E1. However, state S0 indicates that the network is completely operational, that all network spans are available. Therefore, the second constraint 2.2.4c says for all demands and paths that if a demand uses a path that is carried on a working span then the capacity must be less than or equal to the capacity of the span if that span is available. Equation 2.2.4d is included to show how span modularity can be included in the path restoration formulation. By including m in the right-hand side of the equation, 9.3.1c, modularity M can be applied with capacities of 10Gbps and 40Gbps.

2.2.6 Hybrid Recovery

Hybrid recovery can be accomplished by selective or sequential application of span protection and path restoration. The capacity required for hybrid recovery is calculated by using the maximum capacity required by either span protection or path restoration on each span.

3 Experimental Design

3.1 The Experiment

In previous sections, the different network failure recovery methods have been explained and related models for minimizing the cost of working and spare network capacities were presented. This section continues with details of a series of statistical experiments designed to help service providers determine which of the recovery methods should be used. Service providers continue to design networks based on rules-of-thumb developed in the absence of conclusive research that shows network costs under various conditions. This study addresses this shortcoming through a rigorous statistical comparison of optimally engineered networks to give practitioners insights into the best recovery method or methods under a variety of situations and assumptions commonly found in practice.

The problem addressed is: given a mesh network with a specified number of links, demand level, and traffic modularity, which recovery technique should be deployed: span protection, path restoration, fast-reroute restoration, or hybrid recovery? The dependent variable is the cost of an optimally designed network, since this is the primary decision-making metric in practice.

Many factors can affect the cost of a network, including the number of nodes, number of links, network topology, number and size of the demands to be carried, the total amount of traffic on the network, costs of component equipment, cost of capital equipment, data modularity, and capacities of each span. Based on the recommendations of industry professionals, many of these variables are assumed to be fixed. The study assumes an existing network with 20 nodes (roughly the size of a significant national network currently implemented by a service provider) carrying a total of 3.8 terabits per second of traffic, unlimited bandwidth in each span, and an equipment cost that is a linear function of its associated bandwidth. All demands in a given instance are assumed to be the same size, but each connection is between a different pair of nodes.

Factor	Factor levels
Number of network links	31 (1 instance)
	57 (3 intances)
	95 (3 instances)
	133 (3 instances)
Numberof demand x demand unit	380 demands x 10Gbps (6 replications)
	95 demands x 40Gbps (6 replications)
	38 demands x 100Gbps (6 replications)
Link modularity	None
	Modularity at 10Gbps and 40Gbps
Recovery method	Span protection
	Path restoration
	Fast Reroute restoration
	Hybrid recovery

Experiment factors and leves

In this experiment, the factors to be varied are: the size and number of demands, number of network spans, and link modularity. The levels for these factors are shown in Table 3-1, and are described in detail in the sections to follow. There are ten network topologies, nine of which are randomly generated. For each of the three demand levels, the number of individual demands multiplied by the demand size (in gigabits per second) equals a total traffic load of 3,800 Gbps, so that the total network loads were comparable across levels. Span capacities are evaluated both with and without modularity. With modularity, capacity is assigned to spans in multiples of 10 and 40 Gbps; without modularity, span capacity can take on any integer value.

3.1.1 Response Variable: Optimal Network Cost

The major costs in service-provider MPLS core networks are fiber-optic links and router optical interfaces. However, as a result of the well-publicized fiber-optic surplus and the ability to enable high-capacity wavelengths on existing fiber-optic links when additional capacity is required, managers consider fiber-optics expenses to be sunk costs and generally not a part of their network cost calculations. Therefore, the cost applied by service providers to their network designs is based almost exclusively on routers' optical interface costs.

The optimized network models previously presented also use router optical interface cost as their cost basis which reflects this industry practice. These results are also compared to the "raw" cost of uncapacitated links using an average interface cost based on 1Gbps of traffic. This analysis is included because it is the method typically used by service providers; they traditionally do not integrate modular link sizes in their network design processes but do so by simply dividing anticipated raw bandwidth by an interface speed. The motivation for this approach is that service providers are anxious to discover if applying link modularity has a significant effect on network cost or not.

Router optical interface speeds of 10Gbps and 40Gbps are used in the models. The 10Gbps interfaces have traditionally been used in service provider networks but 40Gbps interfaces are now being implemented. Service providers prefer to use 40Gbps router optical interfaces because they are easier to manage but are unsure whether they are more cost effective than 10Gbps interfaces.

The 10Gbps router optical interfaces cost approximately $10,000 each and 40Gbps interfaces cost roughly four times that. It can therefore be reasonably assumed that the average cost per 1Gbps is $1000 for either the 10Gbps or 40Gbps interface speeds. Since there is little apparent cost efficiency for choosing one interface speed over the other, other factors may have unexpected influences on the cumulative network cost of each interface type.

The network models are run first using uncapacitated links. The unit cost per 1Gbps demand unit in the objective function of each model, ξ_e, is therefore $1000. Later, the models are run using modularity (m). These models will select the number of 10Gbps and 40Gbps interfaces for each link required to carry the demands based on the minimized cost results of the model calculations. However, since 40Gbps interfaces are preferred, the cost of the 40Gbps interface is decreased by one dollar to slightly favor choosing 40Gbps over 10Gbps interfaces.

Without doubt, the cost of the router optical interfaces will change dramatically in the future. However, the cost of the interfaces is irrelevant to the relative results provided by the models as long as the ratio between the 10Gbps and 40Gbps interfaces is accurate. (This ratio is presently 1:4.) The same logic applies to the development of higher-speed interfaces. For example, 100Gbps optical interfaces are becoming available but are not expected to be widely implemented for several years. In any case, changes like these can be easily and quickly accomplished in the models and the results recalculated.

3.1.2 Factor: Number of Links

Tests are performed on a library of ten networks, labeled N0-N9. One network, N0 (Figure 7), is an actual network while networks N1-N9 are randomly generated. All networks have 20 nodes. Except for N0 with 31 uncapacitated bidirectional spans, the numbers of spans are chosen and assigned to groups of randomly generated networks to provide a varied representation of span density and related percentage of mesh and node degree (Table 2). In addition to Network N0, three groups of three networks (N1-N3, N4-N6, and N7-N9) are assigned 57, 95 and 133 uncapacitated bidirectional spans respectively.

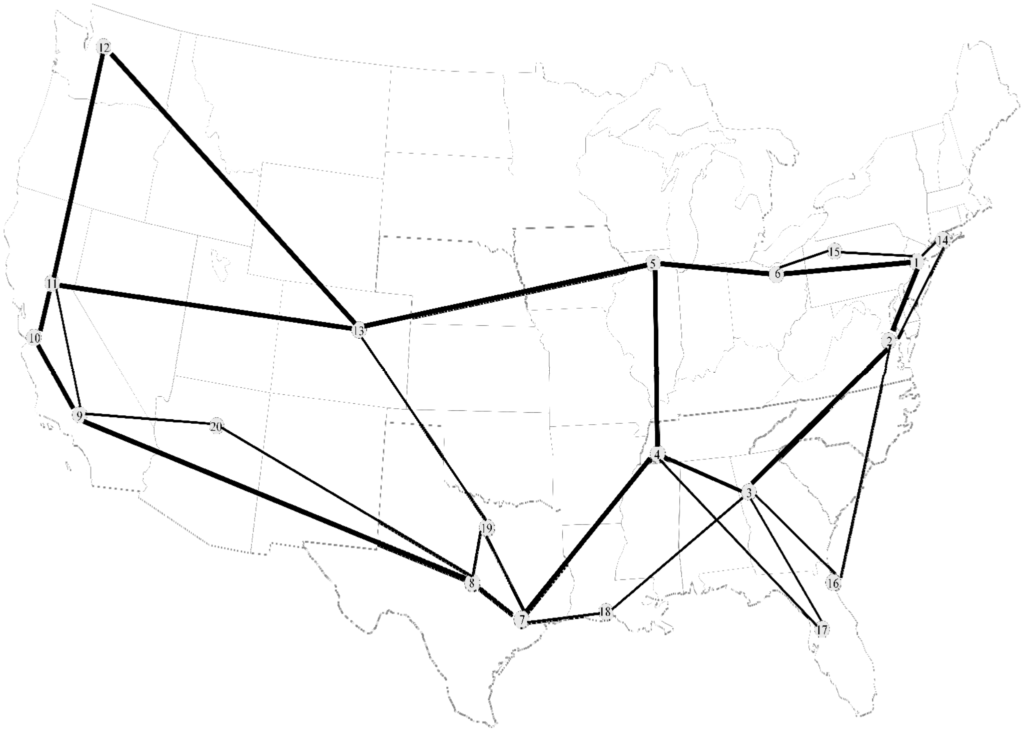

Figure 7: N0 Network Map

Network (N0-N9)	Number of links	Percent of Mesh	Average node degree
N0	31	16.32%	1.55
N1-N3	57	30.00%	2.85
N4-N6	95	50.00%	4.75
N7-N9	133	70.00%	6.65

Table 3-1: Number of links, percent of mesh and node degree for networks N0-N9

The networks N0-N9 all differ by the span connectivity between the twenty nodes. Number of links refers to the number of spans connecting network nodes. Four link densities (31, 57, 95 and 133 spans) are used in the studies. Network N0, the real network, has 31 spans. Networks N1-N3 use 57 spans; Networks N4-N6 have 95 spans while Networks N7-N9 are each connected by 133 spans. With the exception of 31 links, the link densities were chosen based on their mesh percentage and resulting node degree which are common metrics used in service provider networks. Mesh percentage is the ratio of links that actually connect nodes to the number of links that would be needed to fully mesh all nodes with links. The number of links required to fully mesh all nodes can be found by the formula n(n-1)/2. Node degree is the ratio of the number of nodes to the number of links. 57, 95 and 133 links have mesh percentages of 30%, 50% and 70% and node degrees of 2.85, 4.75 and 6.65 respectively.

3.1.3 Factor: Demand Size

Bidirectional O-D pair demands are randomly generated between the twenty nodes used for the test networks N0-N9. The numbers of demands generated are 38, 95 or 380 demands. Each group of generated demands is a demand set. Eighteen unique demand sets are generated: six replications for each of the 38, 95, and 380 demand sizes.

The traffic load on all networks is constant at 3800Gbps (3.8Tbps). This is typical of the network demand on several national service provider networks (Guichard, Le Faucheur, & Vasseur, 2005). The traffic load on the test networks is characterized by a demand set and a related demand size.

Demand sets contain 38, 95, or 380 demands. Demand sizes of 10Gbps, 40Gbps and 100Gbps are used because they are common capacities used in service provider networks. The number of demands are paired with only one demand size such that that the product of each pair is 3800Gbps, the constant traffic load on all test networks (Figure XX).

Number of demands x Demand size	Network traffic load
380 demands x 10Gbps	3800Gbps
95 demands x 40Gbps	3800Gbps
38 demands x 100Gbps	3800Gbps

Figure XX: Traffic load

3.1.4 Factor: Modularity

In real networks the connecting spans are only available in certain capacities, or modules. Tests are performed with and without modularity applied. When modularity is considered, two levels are used: 10Gbps and 40Gbps. These are used because they are common span modules sizes used in the industry. Results may use a combination of 10Gbps and 40Gbps spans. However, models are slightly biased to favor 40Gbps spans over 10Gbps spans because this is the preference in the industry.

In this testing, span modularity is applied and compared to raw capacity needs. Span modularity typically used in service provider networks is 10Gbps and 40Gbps. The 40Gbps interface costs today are roughly four times the 10Gbps interface costs. The 40Gbps spans are preferred by service providers since they are easier to manage than a larger number of 10Gbps spans. Therefore, the cost function in the models of 40Gbps span modularity is slightly favored to force selection of 40Gbps spans over 10Gbps spans.

3.1.5 Factor: Restoration Method

Four restoration methods are compared including span protection, path restoration, Fast Reroute Facility protection and Hybrid recovery.

3.1.6 Hypotheses Investigated

The study's goal of comparing recovery techniques under combinations of experimental factors is achieved by gathering evidence and statistically analyzing the results to determine the truth or falsity of the following hypotheses.

1. *H_0 #1:* The network costs are equal for all link densities regardless of restoration method.
2. *H_0 #2:* The network costs are equal for all demand sizes regardless of restoration method.
3. *H_0 #3:* The network costs are equal for all restoration methods regardless of demand size and number of links.
4. *H_0 #4:* The network costs are equal for all demand sizes and link densities regardless of restoration method.
5. *H_0 #5:* The network costs are equal for all link densities and restoration method regardless of demand size.
6. *H_0 #6:* The network costs are equal for all demand sizes, restoration method regardless of number of links.
7. *H_0 #7:* The network costs are equal for all link densities, demand sizes and restoration methods regardless of modularity.
8. *H_0 #8:* The network costs are equal for all restoration methods, link densities and demand size regardless of modularity.

By rigorously addressing these postulations, key factors in selecting the appropriate restoration technology should emerge. The next section describes the evidence that is used and how it was collected for analysis.

3.2 The Design

3.2.1 What evidence is used

To create optimized networks for each combination of factors, a mathematical program must be solved. To this end, a computer model has been developed for each of the restoration methods described in Chapter 2. The models are solved using state-of-the-art optimization software and the minimum network cost determined for each instance is recorded for analysis.

3.2.2 How evidence is gathered

The evidence gathering process involves the following steps:

1. Generating networks;
2. Generating demand sets;
3. Finding paths through links on networks;
4. Generating input files for mathematical language program and optimizer;
5. Optimizing routing of demand on paths.

3.2.2.1 Network Generation

Random networks N1-N9 are generated using the RGEN network generator developed by McLoud (McLoud). Parameters accepted by RGEN relevant to these studies include the number of nodes and number of links. In this case, all networks have 20 links but three networks each are assigned 57, 95, and 133 links. All networks are generated with a mesh topology and include links distances.

3.2.2.2 Demand Generation

Demands are randomly generated between the twenty nodes such that duplications are minimized. Three sets of six demand sets are generated containing 380, 95, and 38 demands. Therefore, eighteen unique demand sets are generated. Each network is tested with all eighteen demand sets.

3.2.2.3 Path Generation

Once the networks and demands are generated, candidate paths for demands are found through the networks using a path finding program developed by Olinick (Olinick) for use in the path-restoration models. The path-finding program is written with AMPL mathematical programming language. Inputs to the path finding program include the number of nodes, links and link distances which are the output of RGEN and the demand data from the demand generation. These inputs are prepared using Excel to develop the necessary text file for the path-finding program. The path-finding program finds the least costly candidate paths through the test networks based on link distance. Paths are span and node diverse.

3.2.2.4 File Generation

Once the test networks, demand sets and candidate paths are created, several files are generated. δ_{edp} is a “1” if span e belongs to path p realizing demand d ; “0”, otherwise for all *e,d,p*. For span protection and Fast Reroute restoration, x_{dp0} is generated which is the normal flow associated with path p of demand d for all d,p. For path restoration, x_{dps} is generated instead of x_{dp0} so the state of links can be considered for the normal flow associated with path p of demand d for all d,p,s. Finally, β_{eql} is generated showing a "1" if span l belongs to path q restoring span e; "0", otherwise. All files are generated using Excel and are inputs to the solver.

3.2.3 Software and Computing Environment

GAMS is used as the mathematical programming language to represent the models developed in Chapter 2. The files generated are referenced and inputs in GAMS. GAMS then uses CPLEX as the solver to find the cost optimized solutions to the tests.

All tests are performed on a Hewlett Packard Compaq 6715b equipped with an AMD Turion 64 dual core processor running at 2.3 GHz with 4GB of RAM. The models are implemented using the General Algebraic Modeling System (GAMS) model description language (Brooke; Kendrick; Meeraus; Raman;). Solutions are generated using CPLEX. The SAS (Norusis, 2010) and SPSS (George & Mallery, 2010) software packages are used to perform the statistical analysis.

3.2.4 Number of observations

Eighteen unique demand sets are tested on Networks N0-N9. Demand sets contain 380, 95, or 38 demands. Six demand sets using the 380, 95, and 38 demands was generated for a total of eighteen demand sets. Although the combination of sources and destinations for each of the eighteen demand sets is unique, all demand sets consume the same amount of bandwidth across the network to which they are applied. Demand units used for each demand are 100Gbps, 40Gbps and 10Gbps because they are common modules used within the industry. To maintain a constant bandwidth usage across the networks, the number of demands is varied: 38 demands at 100Gbps for each of demand sets 1-6; 95 demands at 40Gbps each for each of demand sets 7-12 and 380 demands at 10Gbps each for demand sets 13-18 are used. Each demand module times the number of demands in every instance consumes a total network bandwidth of 3800Gbps (3.8 Tbps).

Each of the ten networks are tested against each of the eighteen unique demand sets for a total of 180 separate examinations across Networks N0-N9 for each restoration method for 720 tests. In addition, modularity is tested on all networks and demands sets with the exception of hybrid recovery for 540 more tests. (Hybrid recovery is found by a manual process and is not optimized using GAMS.) Therefore, there is a total 1260 test accomplished.

3.2.5 Randomization

To avoid unknown biases in the results, traditional experimentation protocol requires randomization in the observation collection order. In this situation, however, the results from the computer runs are the same no matter what order the problems are solved.

As a result, randomization was not required for this design. While the software was applied to the various problem sets in no particular order, a strict randomization process was not followed.

3.3 The Analysis

3.3.1 Data Analysis Method

Hypotheses are tested on one or more of the network sets N0, N1-N3, N4-N6 and N7-N9. Eight hypotheses are tested.

3.3.2 Test Statistics Used

The tests are performed using the ANOVA statistical analysis. One of four restoration methods (span protection, path restoration, Fast Reroute Facility protection and hybrid recovery) are one factor and one or more of the other investigated factors previously listed are the others. One-factor, two-factor, three-factor tests are used in the analysis.

SAS statistical analysis software was for testing. Fill in something here to describe procesdurre. Enterprise edition

3.3.3 Significance Levels

For all tests the level of significance is set to $\alpha = 0.05$ for Type I error. In addition, Tukey's Honest Significance Test (HST), also called the Tukey range test, is conducted when the results are significant (National Institute of Science and Technology (NIST)). Tukey's HST is a multiple comparison statistical test which compares every possible pair of means often used with ANOVA to determine which means are significantly different from each other greater than the standard error would allow. The output from Tukey's is a grouping of means that are the same.

Network	Number of Links	Number of demands x Demand unit	RECOVERY METHOD				Modularity (not for Hybrid)	
			1:1	Path	FRR	Hybrid		
N0	31	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N1	57	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N2	57	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N3	57	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N4	95	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N5	95	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N6	95	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N7	133	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N8	133	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
N9	133	380 demands x 10Gbps (6)	6	6	6	6	18	
		95 demands x 40Gbps (6)	6	6	6	6	18	
		38 demands x 100Gbps (6)	6	6	6	6	18	
			180	180	180	180	540	**1260**

Table XX:

4 Experiment Test Results

Experiment tests are presented by one-, two- , and three-factor ANOVA statistical analysis results. Each is accomplished over several hypotheses developed earlier.

4.1 4.1 One-Factor Analyses

One-factor ANOVA analyses are presented first. Hypotheses *H_0 #1, H_0 #2, and H_0 #3* are tested to find the cost effect of recovery methods, number of links, and demand units when considered individually.

4.1.1 Cost Effect of Recovery Method

***H_0 #1:* Total cost is the same across all recovery methods.**

This statistical analysis seeks to determine if the average cost is influenced by the four recovery methods (1:1, Path, FRR and Hybrid) used across all numbers of links, demand units and modularity (Table 1). A one-factor analysis of variance shows that recovery method is a statistically significant factor ($p < .0001$) based on a 5% significance level (□his statistical analy*H_0 #1* is rejected; the average cost is *not* the same across all recovery methods.

In addition, Tukey's HSD test is conducted, and the results are summarized in Table 1. Based on the results of this test, there is a statistically significant difference in average cost between the four recovery methods.

	Recovery Method			
	1:1	**Path**	**FRR**	**Hybrid**
Tukey Grouping	D	A	B	C
Average Cost	$14,891,522	$11,223,871	$11,949,451	$13,586,896
Data Points	180	180	180	180

Table 2: Cost effect of recovery method

The analysis also shows, in general, the average cost of 1:1 > Hybrid > FRR > Path (Table 1 and Figure 1). Therefore, 1:1 protection is the by far the most costly recovery method. Path restoration is, on average, the most efficient when not considering demand size, link density, or modularity.

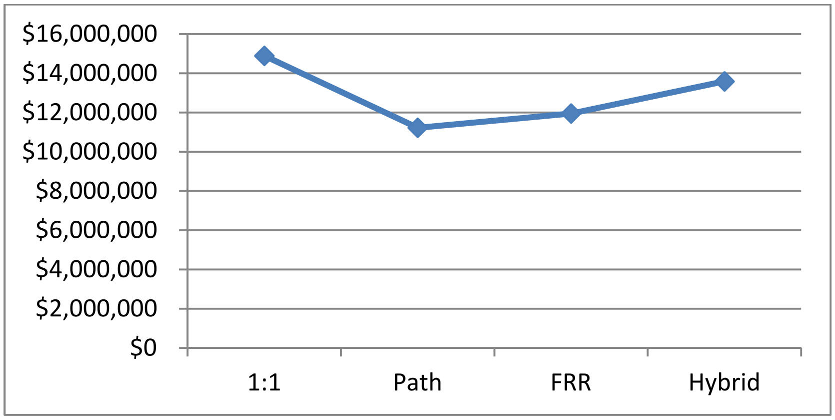

Figure 8: Cost effect of recovery method

4.1.2 Cost Effect of Number of Links

H_0 #2: **The total cost is the same across all numbers of links.**

This analysis is based on the average cost of the number of links (31, 57, 95 and 133) tested across all recovery methods, demand units, and modularity (Table 2). A one-factor analysis of variance shows the number of links are statistically significant factors (p < .0001) based on a 5% significance level (T□□□̣□□□□□□□̣□□̣□□□□□ H_0 *#2* is rejected; the average cost is *not* the same across all number of links.

In addition, Tukey’s HSD test is conducted and the results are summarized in Table 2. Based on the results of this test, there is a statistically significantly difference in average cost of the number of links used.

	Number of Links			
	31	**57**	**95**	**133**
Tukey Grouping	D	C	B	A
Total Cost	$21,698,326	$14,591,898	$11,317,316	$9,901,128
Data Points	72	216	216	216

Table 3: Cost effect of the number of demand units

The analysis also shows, in general, network cost increases as number of links decreases; 31 links > 57 links > 95 links > 133 links (Table 2 and Figure 2). The cost is least for 133 links and most for 31 links when not considering recovery method, demand units or modularity.

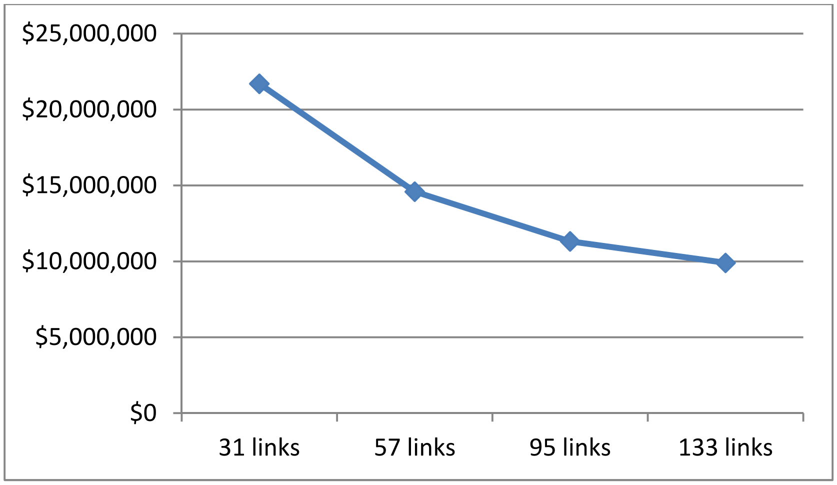

Figure 9: Cost effect of the number of number of links

4.1.3 Cost Effect of Number of Demand Units

***H_0 #3:* The total cost is the same across all demand units.**

This analysis is based on the average cost of the number of demand units (100Gbps, 40Gbps, and 10Gbps) tested across all recovery methods, number of links and modularity (Table 3). A one-factor analysis of variance shows demand units are statistically significant factors ($p < .0001$) based on a 5% significance level (□his analysis is basedH_0 *#3* is rejected; the average cost is *not* the same across all demand units.

In addition, Tukey's HSD test is conducted, and the results summarized in Table 3. Based on the results of this test, there is a statistically significant difference in average cost between demand units.

	Demand Units		
	100Gbps	**40Gbps**	**10Gbps**
Tukey Grouping	C	B	A
Total Cost	$13,826,258	$12,637,511	$12,249,311
Data Points	240	240	240

Table 4: Cost effect of the number of demand units

The analysis also shows, in general, 100Gbps > 40Gbps>10Gbps demand units. 100Gbps demand sizes are most costly (Table 3 and Figure 3.) Decreasing demand size to 40Gbps or 10Gbps results in cost savings. Also, the cost difference between 100Gbps and 40Gbps is greater than that of 40Gbps and 10Gbps.

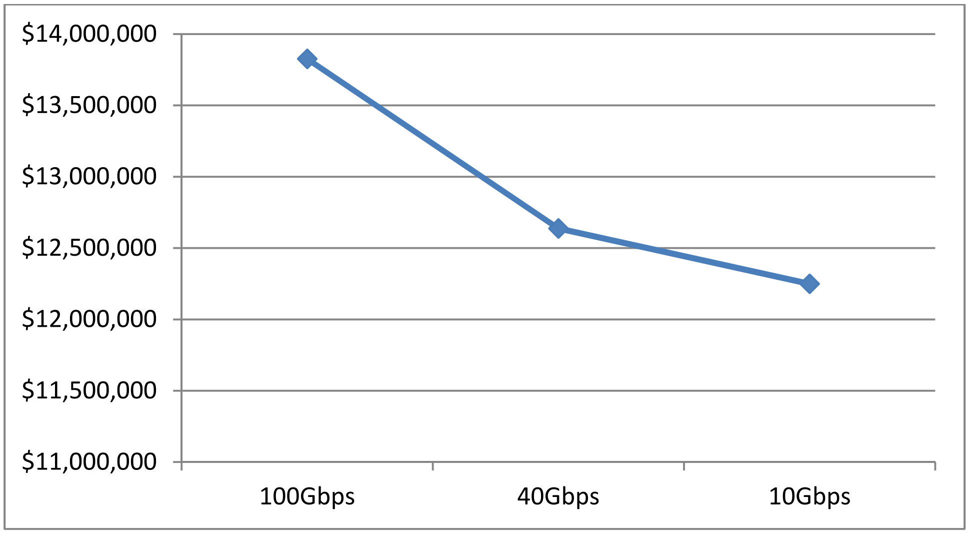

Figure 10: Cost effect of the number of demand units

4.2 Two-Factor Analyses

Two-factor ANOVA analyses are presented next. Hypotheses *H_0 #4, H_0 #5, and H_0 #6* are tested to find the cost effect of the number of links and demand units, recovery method and number of links, and recovery method and demand units when considered together.

4.2.1 Cost Effect of Number of Links and Demand Units

***H_0 #4:* The total cost is the same across all numbers of links and demand units.**

This analysis is based on the average cost of the number of links (31, 57, 95, and 133) and demand units (100Gbps, 40Gbps, and 10Gbps) tested across all recovery methods and modularity (Table 4). A two-factor analysis of variance shows the number of links and demand units are statistically significant factors ($p < .6505$) based on a 5% significance level (□his analysis is basll*H_0 #4* is not rejected.

		Demand Units		
Number of links	**Parameter**	**100Gbps**	**40Gbps**	**10Gbps**
31	**Average Cost**	\$22,412,556	\$21,359,903	\$21,064,137
	Data Points	24	24	24
57	**Average Cost**	\$15,451,625	\$14,365,612	\$13,836,202
	Data Points	72	72	72
95	**Average Cost**	\$12,232,198	\$11,027,129	\$10,781,516
	Data Points	72	72	72
133	**Average Cost**	\$10,932,850	\$9,612,327	\$9,191,941
	Data Points	72	72	72

Table 5: Cost effect of the number of links and demand units

4.2.2 Cost Effect of Recovery Method and Number of Links

***H_0 #5:* The total cost is the same for all recovery methods with the same number of links.**

This analysis is based on the average cost of the recovery method (1:1, Path, FRR, and Hybrid) and the number of links (31, 57, 95, and 133) tested across all demand units and modularity (Table 5). A two-factor analysis of variance shows recovery method and number of links are statistically significant factors ($p < .0001$) based on a 5% significance level (□his analysis is basedH_0 *#5* is rejected; the average cost is *not* the same for all recovery methods and number of links.

		Recovery Method			
Number of links	**Parameter**	**1:1**	**Path**	**FRR**	**Hybrid**
31	**Average Cost**	$23,928,889	$20,772,583	$19,428,200	$22,663,632
	Data Points	18	18	18	18
57	**Average Cost**	$17,131,296	$12,231,481	$13,845,931	$15,158,882
	Data Points	54	54	54	54
95	**Average Cost**	$13,310,444	$9,550,169	$10,441,506	$11,967,146
	Data Points	54	54	54	54
133	**Average Cost**	$11,220,370	$8,707,060	$9,068,000	$10,609,080
	Data Points	54	54	54	54

Table 6: Cost effect of recovery method and number of links

Links	Recovery Method	Tukey's Grouping (least to most cost)												
133	**Path**	A												
133	**FRR**	A	B											
95	**Path**	A	B	C										
95	**FRR**			C	D	E								
133	**Hybrid**				D	E								
133	**1:1**					E	F							
95	**Hybrid**						F							
57	**Path**						F							
95	**1:1**							G						
57	**FRR**							G						
57	**Hybrid**								H					
57	**1:1**									I				
31	**FRR**										J			
31	**Path**											K		
31	**Hybrid**												L	
31	**1:1**													M

Table 7: Tukey's grouping for cost effect of recovery method and number of links

In addition, Tukey's HSD test is conducted, and the results summarized in Table 6. Based on the results of this test, there is a statistically significant difference in average cost between recovery methods and number of links combinations.

The analysis in H_0 *#2* showed, in general, the average costs decreases as the number of links increases for each recovery method. In other words, the average cost of 31 > 57 > 95 > 133 links for all recovery methods. In addition, H_0 *#5* results show the cost difference between the recovery methods is greatest at 31 and 57 links; at 95 and 133 links the cost averages are closer together (Table 5 and Figure 4). Also, the cost for 95 and 133 links is almost the same for path restoration.

The analysis in H_0 *#1* shows 1:1 > Hybrid > FRR > Path. However, H_0 *#5* results show the cost relationship between Path and FRR recovery changes at 31 links. The average cost is less using FRR for 31 links (Table 5, highlighted in yellow). This is reflected in Table 6 as well as the equivalent average costs between recovery method and number of links pairs.

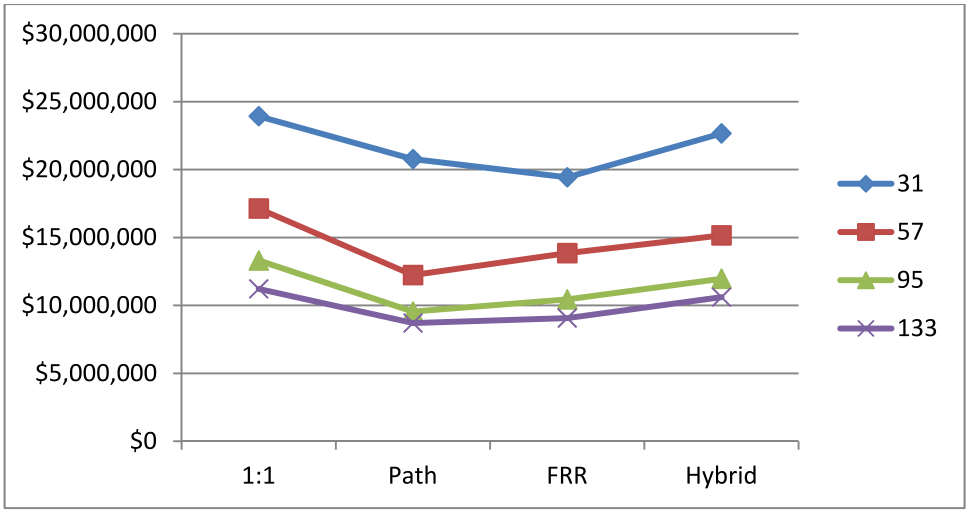

Figure 11: Cost effect of recovery method and number of links

4.2.3 Cost Effect of Recovery Method and Demand Units

***H_0 #6:* The total cost is the same for all recovery methods on networks with the same demand units.**

This analysis is based on the average cost of the four recovery methods (1:1, Path, FRR, and Hybrid) and the demand units (100Gbps, 40Gbps, and 10Gbps) tested across all number of links and modularity (Table 7). A two-factor analysis of variance shows recovery method and demand units are statistically significant factors ($p < .0001$) based on a 5% significance level t□□□̣□□□□□□□̣□□̣□□□□□̣*H_0 #6* is rejected; the average cost is *not* the same for all recovery methods and demand units.

Number of links	**Parameter**	**Recovery Method** **1:1**	**Path**	**FRR**	**Hybrid**
100Gbps	**Average Cost**	$16,578,611	$14,010,926	$13,880,239	$16,969,032
	Data Points	60	60	60	60
40Gbps	**Average Cost**	$16,280,472	$12,567,483	$12,982,963	$14,688,941
	Data Points	60	60	60	60
10Gbps	**Average Cost**	$16,334,167	$11,867,561	$12,724,525	$13,642,083
	Data Points	60	60	60	60

Table 8: Cost effect of recovery method and demand units

In addition, Tukey's HSD test is conducted, and the results summarized in Table 8. Based on the results of this test, there is a statistically significant difference in average cost between recovery methods and demand unit combinations.

Demand Unit	**Recovery Method**	**Tukey's Grouping**									
10Gbps	**Path**	A									
40Gbps	**Path**	A	B	C							
10Gbps	**FRR**			C	D						
40Gbps	**FRR**			C	D	E					
10Gbps	**Hybrid**				D	E	F	G			
100Gbps	**FRR**					E	F	G	H		
100Gbps	**Path**						F	G	H	I	
40Gbps	**Hybrid**									I	

40Gbps	**1:1**										J
10Gbps	**1:1**										J
100Gbps	**1:1**										J
100Gbps	**Hybrid**										J

Table 9: Grouping for cost effect of recovery method and demand units

The analysis in H_0 *#3* showed that, in general, the larger the demand units the more costly the network (100Gbs > 40Gbps > 10Gbps). H_0 *#6* results show an exception is 1:1 restoration where the 40Gbps demand units are on average the least costly (100Gbs > 10Gbps > 40Gbps) (Table 7, highlighted in yellow, and Figure 5). However, the cost differences between all demand units with 1:1 protection are much smaller than for other recovery methods. Also, 1:1 and Hybrid are more clearly separated and defined by cost grouping than Path and FRR, which tend to overlap more. The 10Gbps and 40Gbps are close together for Path and FRR.

The analysis in H_0 *#1* shows 1:1 > Hybrid > FRR > Path. However, H_0 *#6* results show the cost relationship between 1:1 and Hybrid, FRR and Path recovery changes at 100Gbps. Hybrid is on average more costly than 1:1, and Path is more costly than FRR (Table 7, highlighted in blue). In other words, Hybrid > 1:1 > Path > FRR. This is also reflected in Table 8 as well as the equivalent average costs between recovery method and demand units' pairs.

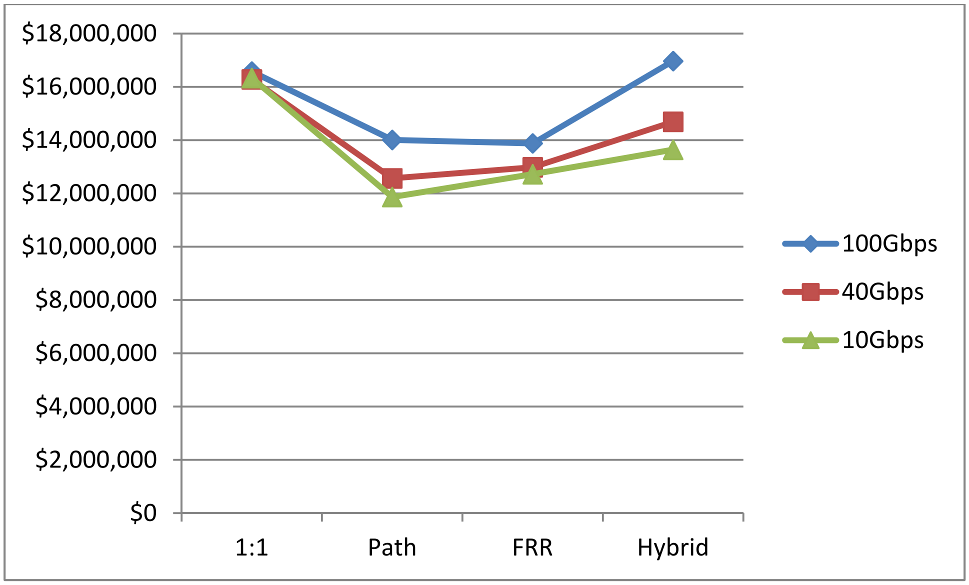

Figure 12: Cost effect of recovery method and demand units

4.3 Three-Factor Analyses

Three-factor ANOVA analyses are presented next. Hypotheses H_0 *#7* is tested to find the cost effect of recovery methods, number of links, and demand units when considered together.

4.3.1 Cost Effect of Recovery Method, Number of Links, and Demand Units

***H_0 #7:* The total costs are the same for all recovery methods on networks with the same number of links and demand units.**

This analysis is based on the average cost of the four recovery methods (1:1, Path, and FRR), number of links (31, 57, 95, and 133) and demand units (100Gbps, 40Gbps, and 10Gbps) tested across modularity (Table 9). A three-factor analysis of variance shows recovery method, demand units, and numbers of links are not statistically significant factors (p = .9992) based a 5% significance level (□emand units, and numbH_0 *#7* is *not* rejected.

LINKS	DEMAND SIZE	1:1	Path	FRR	Hybrid
31	100Gbps	$24,133,333	$22,008,333	$20,086,891	$24,964,279
	Data points	6	6	6	6
	40Gbps	$23,760,000	$20,356,667	$19,144,630	$21,917,864
	Data points	6	6	6	6
	10Gbps	$23,893,333	$19,952,750	$19,053,079	$21,108,755
	Data points	6	6	6	6
57	100Gbps	$17,322,222	$13,544,167	$14,512,226	$16,871,650
	Data points	18	18	18	18
	40Gbps	$17,151,667	$11,960,090	$13,631,917	$15,128,888
	Data points	18	18	18	18
	10Gbps	$16,920,000	$11,190,187	$13,393,650	$13,476,109
	Data points	18	18	18	18
95	100Gbps	$13,494,444	$10,532,407	$11,081,098	$13,659,438
	Data points	18	18	18	18
	40Gbps	$13,112,444	$9,445,821	$10,256,269	$11,541,735
	Data points	18	18	18	18
	10Gbps	$13,324,444	$8,672,277	$9,987,149	$10,700,265
	Data points	18	18	18	18
133	100Gbps	$11,364,444	$9,958,796	$9,840,741	$12,376,759
	Data points	18	18	18	18
	40Gbps	$11,097,778	$8,507,353	$8,899,037	$10,167,277
	Data points	18	18	18	18
	10Gbps	$11,198,889	$7,655,032	$8,464,222	$9,283,204
	Data points	18	18	18	18

Table 10: Effect of recovery method, number of links, and demand units

4.3.2 Cost Effect of Modularity

H_0 #8: **The total cost is the same for all restoration methods on networks with the same number of links, demand units, and modularity.**

Due to the complexity of a four factor analysis, data from 4.3.1 was run with modularity of 10Gbps and 40Gbps instead of doing a four factor analysis and the cost difference between the two results was compared.[4] As Table 10, 11, and 12 illustrate, the cost difference between the non-modular and modular runs is minimal. Modularity does not affect the average network cost.

Restoration Method/Modularity (Y/N)						
Link density/Demand size	**1:1/Y**	**1:1/N**	**Path/Y**	**Path/N**	**FRR/Y**	**FRR/N**
31 links-100Gbps	$24,339,405	$24,133,333	$21,354,474	$22,008,333	$20,001,180	$20,086,891
31 links-40Gbps	$23,766,073	$23,760,000	$21,322,889	$20,356,667	$19,251,197	$19,144,630
31 links-10Gbps	$24,316,070	$23,893,333	$19,980,422	$19,952,750	$19,144,547	$19,053,079
57 links-100Gbps	$17,562,354	$17,322,222	$13,579,667	$13,544,167	$14,769,092	$14,512,226
57 links-40Gbps	$16,926,244	$17,151,667	$11,911,371	$11,960,090	$13,849,107	$13,631,917
57 links-10Gbps	$17,090,142	$16,920,000	$11,194,181	$11,190,187	$13,589,144	$13,393,650
95 links-100Gbps	$14,055,773	$13,494,444	$11,121,393	$10,532,407	$11,680,830	$11,081,098
95 links-40Gbps	$13,113,006	$13,112,444	$9,366,989	$9,445,821	$10,353,636	$10,256,269
95 links-10Gbps	$13,535,816	$13,324,444	$8,788,139	$8,672,277	$10,462,521	$9,987,149
133 links-100Gbps	$11,894,718	$11,364,444	$10,589,742	$9,958,796	$10,504,748	$9,840,741
133 links-40Gbps	$11,097,500	$11,097,778	$8,509,790	$8,507,353	$9,007,557	$8,899,037
133 links-10Gbps	$11,359,779	$11,198,889	$7,716,510	$7,655,032	$8,665,951	$8,464,222

Table 11: Modular and non-modular costs

---

[4] Hybrid recovery was not included in the analysis because to determine hybrid modularity costs requires extensive manual calculation and was not accomplished.

Links/Demand Unit	1:1 Modular Ratio	Path Modular Ratio	FRR Modular Ratio
31 links-100Gbps	0.85%	-3.06%	-0.43%
31 links-40Gbps	0.03%	4.53%	0.55%
31 links-10Gbps	1.74%	0.14%	0.48%
57 links-100Gbps	1.37%	0.26%	1.74%
57 links-40Gbps	-1.33%	-0.41%	1.57%
57 links-10Gbps	1.00%	0.04%	1.44%
95 links-100Gbps	3.99%	5.30%	5.13%
95 links-40Gbps	0.00%	-0.84%	0.94%
95 links-10Gbps	1.56%	1.32%	4.54%
133 links-100Gbps	4.46%	5.96%	6.32%
133 links-40Gbps	0.00%	0.03%	1.20%
133 links-10Gbps	1.42%	0.80%	2.33%

Table 12: Ratio of modular to non-modular costs

Modularity (Y/N)	Average Cost ($)
Y	$14,326,999
N	$14,326,999

Table 13: Average cost comparison between modular and non-modular networks

4.4 Analysis Summary

This section presents a summarization of the analyses previously accomplished. Summaries for numbers of links demand unit, recovery method, and modularity are detailed.

4.4.1 Number of Links Summary

The results of H_0 *#2* showed the average cost of 31 > 57 > 95 > 133 links. These results were consistent throughout all other the tests which were statistically significant.

4.4.2 Demand Units Summary

The results of H_0 *#3* showed the average cost of 100Gbps > 40Gbps > 10Gbps demand units. However, H_0 *#5* shows that relationship holds for all recovery methods except 1:1 protection where 100Gbps demand units are still the most costly, but 40Gbps demand units are less costly than 10Gbps demand units. These results are summarized in Table 13. Cells that are not highlighted show results from H_0 *#3* while highlighted cells reflect differences found in H_0 *#6*.

Links	Recovery Method	Least cost	Middle cost	Most cost
31	**1:1**	40Gbps	10Gbps	100Gbps
	Path	10Gbps	40Gbps	100Gbps
	FRR	10Gbps	40Gbps	100Gbps
	Hybrid	10Gbps	40Gbps	100Gbps
57	**1:1**	40Gbps	10Gbps	100Gbps
	Path	10Gbps	40Gbps	100Gbps
	FRR	10Gbps	40Gbps	100Gbps
	Hybrid	10Gbps	40Gbps	100Gbps
95	**1:1**	40Gbps	10Gbps	100Gbps
	Path	10Gbps	40Gbps	100Gbps
	FRR	10Gbps	40Gbps	100Gbps
	Hybrid	10Gbps	40Gbps	100Gbps
133	**1:1**	40Gbps	10Gbps	100Gbps
	Path	10Gbps	40Gbps	100Gbps
	FRR	10Gbps	40Gbps	100Gbps
	Hybrid	10Gbps	40Gbps	100Gbps

Table 14: Demand unit cost summary

4.4.3 Restoration Summary

The results of H_0 *#1* show the average cost of 1:1 > H > FRR > Path recovery. However, H_0 *#5* found FRR is less costly than 1:1 protection (Path > FRR) for all demand units of 31 links as well as for 100Gbps demand units for 57, 95, and 133 links when the number of links was considered with recovery method. In addition, H_0 *#6* showed hybrid recovery is more costly than 1:1 protection (Hybrid > 1:1) for all number of links at 100Gbps when demand units were considered with recovery method. Notice that if path restoration is not an option due to speed of recovery that FRR is *always* the least costly recovery method. These results are summarized in Table 14.

Links	Demand units	Least costly	Most costly
31	100Gbps	FRR	HYBRID
	40Gbps	FRR	1:1
	10Gbps	FRR	1:1
57	100Gbps	FRR	HYBRID
	40Gbps	PATH	1:1
	10Gbps	PATH	1:1
95	100Gbps	FRR	HYBRID
	40Gbps	PATH	1:1
	10Gbps	PATH	1:1
133	100Gbps	FRR	HYBRID
	40Gbps	PATH	1:1
	10Gbps	PATH	1:1

Table 15: Restoration Summary

4.4.4 Modularity Summary

The results of H_0 *#8* show that modularity has no effect on the average network cost. These results were consistent throughout all other the tests which were statistically significant.

4.5 Service Provider Recommendations

Given the results obtained in this praxis, service provider managers can make informed financial decisions regarding the networks whether they are already operating or are in the planning stages. Recommendations regarding recovery methods, number of network links, demand unit sizes, and modularity are presented.

In some cases, recommendations are made regarding migrating to different recovery methods, number of links and demand units. Specific cost analysis of migrations is the topic of the next chapter.

4.5.1 Service Provider Recommendations: Number of Links

The average network cost is 31 > 57 > 95 > 133 links in every circumstance even when considering recovery method and demand size. As previously discussed, fiber-optics are considered sunk costs and not included in the models. Service providers that use less than 133 links should consider migrating to 133 links.

4.5.2 Service Provider Recommendations: Recovery Method

When number of links and demand size are not considered, 1:1 protection is the most and path restoration the least costly such that 1:1 > Hybrid > FRR > Path. This agrees with long-held service provider opinion and is accurate until the number of links and demand units are considered. Exceptions occur for all demand units at 31 links as well as at 100Gbps demand units for 57, 95 and 133 links: path restoration is more costly than FRR. Therefore, in these situations service providers should only consider FRR, *not* path restoration. Also at 100Gbps demand units for *all* numbers of links hybrid recovery is on average *less* costly that 1:1 protection.

Path restoration, and therefore hybrid recovery, being a combination of 1:1 and Path recovery, are inherently too slow for larger networks and may not be an option for some service providers. In this case the clear choice for the least average cost for all networks regardless of the number of links and demand units is *always* FRR.

4.5.3 Service Provider Recommendations: Demand Units

On average the demand units of 100Gbps > 40Gbps > 10Gbps when not considering recovery method or number of links. Service providers that use 10Gbps demand units will be most cost efficient *except* for those using 1:1 protection. In this case, the service provider should use 40Gbps demand sizes since 100Gbps > 10Gbps > 40Gbps. Migration from 100Gbps or 10Gbps demand units is recommended for 1:1 protected networks. In addition, hybrid restoration is on average more costly than 1:1 protection when using 100Gbps for *any* number of links. Service providers should not consider migrating from 1:1 protection to hybrid restoration for any network that uses 100Gbps demand units.

4.5.4 Service Provider Recommendations: Modularity

Service providers are uncertain whether applying modularity to their network designs has any cost impact. This results previously presented show that modularity has no effect on the resulting costs.

5 Costs of Link Density, Demand Size and Recovery Method Migrations

Knowing the cost benefit or penalty for migrating from one network configuration to another is an important asset for service providers and network managers. However, reliable research and data is scarce or non-existent. This chapter investigates this important topic and provides conclusive insight to solutions.

Information from the test previously performed in compiled and used as the data for this chapter. First, a cost ranking of networks using combinations of number of links, demand units, and recovery method was created and are found in Table 20 in Appendix A. Also, the costs of migrating from different number of links, demand units, and recovery methods is presented in Figures 13 and 14 and can be found in Appendix B. Finally, a comprehensive listing of the costs of migrating from any number of links, demand unit, and recovery methods to any other is presented in Appendix C.

As found in H_0*#7*, the combination of the number of links, demand units and recovery methods is not statistically significant. Therefore, the results of that test cannot be trust to be accurate. Regardless, these analyses are completed anyway. The results that do not agree with previous statistically significant tests are noted.

5.1 Link Density Migration

Table 17 shows in detail the cost difference percentages between different number of links (31-57, 57-95, 95-133) for each demand unit and restoration method. A negative percentage indicates a cost savings.

A migration from a lower to a higher number of links always results in average network cost savings. There is an average cost savings to migrate the number of links from 31-57 > 57-95 > 95-133 regardless of demand size or restoration method. In addition, the cost savings *magnitude* is greatest for 31-57 and decreases incrementally for 57-95 and 95-133 links. The most savings to migrate the number of links is at 43.92% for 31-57 links at 10Gbps using path restoration and least at 5.45% for 95-133 links at 100Gbps using path restoration.

Link density/Demand size	**1:1**	**Path**	**FRR**	**Hybrid**
31-57 links/100Gbps	-28.22%	-38.46%	-27.75%	-32.42%
57-95 links/100Gbps	-22.10%	-22.24%	-23.64%	-19.04%
95-133 links/100Gbps	-15.78%	-5.45%	-11.19%	-9.39%
31-57 links/40Gbps	-27.81%	-41.25%	-28.80%	-30.97%
57-95 links/40Gbps	-23.55%	-21.02%	-24.76%	-23.71%
95-133 links/40Gbps	-15.36%	-9.94%	-13.23%	-11.91%
31-57 links/10Gbps	-29.19%	-43.92%	-29.70%	-36.16%
57-95 links/10Gbps	-21.25%	-22.50%	-25.43%	-20.60%
95-133 links/10Gbps	-15.95%	-11.73%	-15.25%	-13.24%

Table 16: Cost effect of number of links migrations considering demand unit and recovery method.

Previously, the results of H_0*#2* found showed the average cost to be 31 > 57 > 95 >133 link. The result of migrating the number of links from 31-57, 57-95, and 95-133 links show a cost savings in all cases and therefore agrees with H_0 *#2.*

5.2 Demand Unit Migration

Table 18 shows how the numbers of links and recovery combinations affected the migration of demand units cost difference percentages. A negative percentage represents a cost savings while a positive percentage indicates a cost increase.

All 100Gbps-40Gbps migrations result in cost savings. The results of H_0 *#6* show that for 40Gbps-10Gbps migrations savings result for all numbers of links using 1:1 protection where 10Gbps demand sizes are more costly than 40Gbps demand sizes (represented by a positive percentage with a highlighted red numbers in the 40Gbps-10Gbps column). The one exception is 57 links which shows a 1.35% cost savings of 10Gbps over 40Gbps demand sizes. However, this is not a substantial variance from H_0 *#6*. Also, the magnitude of saving is greater for 100Gbps-40Gbps migrations than for 40Gbps-10Gbps migrations.

In general, the “Cost savings ranking” of Table 18 shows the cost savings is Hybrid > Path > FRR > 1:1. The exceptions (highlighted in yellow) are 57 links where for 100Gbps-40Gbps migration, and 95 and 133 links for 40Gbps-10Gbps migrations where hybrid recovery and path restoration reverse such that Path > Hybrid > FRR > 1:1.

The cost savings is always greater for 100Gbps-40Gbps with one exception. For 57 links using 1:1 protection the cost savings is less for 100Bbps-40Gbps than for 40Gbps-10Gbps migrations (highlighted in blue cells).

The most saving for 100Gbps-40Gbps is for 133 links using hybrid recovery; the lease is for 57 links using path protection. For 40Gbps-10Gbps, the most savings occurs at 57 links using hybrid recovery.

Links/Demand unit/Recovery method	100Gbps-40Gbps migration	40Gbps-10Gbps migration	Cost savings ranking ("1" is the most savings)
31 links/1:1	-1.55%	0.56%	4
31 links/Path	-7.50%	-1.98%	2
31 links/FRR	-4.69%	-0.48%	3
31 links/Hybrid	-12.20%	-3.69%	1
57 links/1:1	-0.98%	-1.35%	4
57 links/Path	-11.70%	-6.44%	2
57 links/FRR	-6.07%	-1.75%	3
57 links/Hybrid	-10.33%	-10.92%	1
95 links/1:1	-2.83%	1.62%	4
95 links/Path	-10.32%	-8.19%	2
95 links/FRR	-7.44%	-2.62%	3
95 links/Hybrid	-15.50%	-7.29%	1
133 links/1:1	-2.35%	0.91%	4
133 links/Path	-14.57%	-10.02%	2
133 links/FRR	-9.57%	-4.89%	3
133 links/Hybrid	-17.85%	-8.70%	1

Table 17: Cost effect of migrating demand units

5.3 Recovery Migration

Table 19 illustrates the cost difference parentages of migrating from one restoration method to another for any given number of links and demand unit combinations. A negative percentage indicates a cost saving of the current restoration method to the transition restoration method. A positive percentage represents a cost increase over the current restoration method. Migration duplications are not repeated because they are simply the opposite of each other, result in cost increases, and are not recommended. For instance, 1:1/Path migration is negative and therefore a cost savings whereas Path/1:1 migration is positive and results in a cost increase of equal magnitude.

5.3.1 1:1 to Path Migration (1:1/Path)

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. This is true for 1:1 to path migration. The most cost savings is accomplished at 95 links at 10Gbps with 31.64%.; the least cost savings results at 31 links and 100Gbps at 8.8% savings.

5.3.2 1:1 to FRR Migration (1:1/FRR)

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. This is true for 1:1 to hybrid migration.The most cost savings occurs with 95 links at 10Gbps with 25.05% savings. The least cost savings is with 133 links at 100Gbps with a 13.41% cost savings.

5.3.3 1:1 to Hybrid Migration (1:1/Hybrid)

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. However, H_0 *#6* showed hybrid recovery is more costly than 1:1 protection (Hybrid > 1:1) for all number of links at 100Gbps when demand units were considered with recovery method. These the results found in Table xx for 1:1/Hybrid, shown with highlighted yellow background, agree with one exception. Using 57 links at 100Gbps with hybrid recovery is 2.6% less cost (highlighted with a blue number) than 1:1 protection. However, this difference is not substantial.

The most cost savings accomplished by 1:1/Hybrid migration occurs with 57 links at 10Gbps with 20.35%. The largest cost increase occurs with 133 links at 100Gbps with a cost increase of 8.91% cost increase migrating to hybrid recovery.

5.3.4 Path to FRR (Path/FRR) Migration

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. However, H_0 *#5* found FRR is less costly than 1:1 protection (Path > FRR) for all demand units of 31 links (blue highlighted background).

In addition, H_0 *#6* showed FRR recovery is less costly than 1:1 protection (Hybrid > 1:1) for all number of links at 100Gbps when demand units were considered with recovery method (highlighted with red numbers). The results found in Table xx for Path/FRR only conclusively agree with H_0 *#6* for 31 links at 100Gbps since 133 links at 100Gbps only shows a slight cost savings of 1.19%. Both 57 and 95 links at 100Gbps both show a significant cost increase of Path/FRR which is usually expected.

5.3.5 Hybrid to Path Migration (Hybrid/Path)

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. This is true for hybrid to path migration in Table XX. The most cost savings occurs with 95 links at 100Gbps with a 22.89% savings. The least cost savings is with 31 links at 10Gbps with a 5.48% savings.

5.3.6 Hybrid to FRR Migration (Hybrid/FRR)

The results of H_0*#2* showed the cost of 1:1 > Hybrid > FFR > Path. This is true for hybrid to FRR migration.The most cost savings occurs with 133 links at 100Gbps with 20.49% savings. The least savings is with 57 links at10Gbps with only a 0.61% savings.

	Recovery Migration From/To					
Links/Demand unit	**1:1/Path**	**1:1/FRR**	**1:1/Hybrid**	**Path/FRR**	**Hybrid/Path**	**Hybrid/FRR**
31 links-100Gbps	-8.81%	-16.77%	3.44%	-8.73%	-11.84%	-19.54%
31 links-40Gbps	-14.32%	-19.42%	-7.75%	-5.95%	-7.12%	-12.65%
31 links-10Gbps	-16.49%	-20.26%	-11.65%	-4.51%	-5.48%	-9.74%
57 links-100Gbps	-21.81%	-16.22%	**-2.60%**	7.15%	-19.72%	-13.98%
57 links-40Gbps	-30.27%	-20.52%	-11.79%	13.98%	-20.95%	-9.89%
57 links-10Gbps	-33.86%	-20.84%	-20.35%	19.69%	-16.96%	-0.61%
95 links-100Gbps	-21.95%	-17.88%	1.22%	5.21%	-22.89%	-18.88%
95 links-40Gbps	-27.96%	-21.78%	-11.98%	8.58%	-18.16%	-11.14%
95 links-10Gbps	-34.91%	-25.05%	-19.69%	15.16%	-18.95%	-6.66%
133 links-100Gbps	-12.37%	-13.41%	8.91%	-1.19%	-19.54%	-20.49%
133 links-40Gbps	-23.34%	-19.81%	-8.38%	4.60%	-16.33%	-12.47%
133 links-10Gbps	-31.64%	-24.42%	-17.11%	10.57%	-17.54%	-8.82%

Table 18: Cost effect comparison of link density and demand size pairs to restoration method migrations

5.3.7 Detailed Migration Cost Results

Not all possible migration combinations could be presented in this praxis. However, a detailed compilation of migration costs from number of links, demand units, and recovery method to a different number of links, demand units, and recovery method is presented in Appendix XX.

6 FUTURE RESEARCH

100Gbps

MPLS TP (Cisco Systems. Inc.)

100Gbps

SRLG’sulti-layer recovery1:1 versus facility(Graziani & Johnson, 2009)

BIBLIOGRAPHY

Alvarez, S. (2001). *High Availability Network Fundamentals: A Practical Guide to Predicting Network Availability.* Indianapolis: Cisco Press.

Ambs, K. S., Cwilich, S., Deng, M., Houck, D. J., Lynch, D. F., & Yan, D. (2000). Optimizing restoraton capacity in the ST&T network. *AT&T Interfaces , 30*, 26-44.

Award Solutions. (2008). *Exploring MPLS.* Richardson, TX.

Award Solutions. (2008). *Verizon Wireless OTNGN and MPLS.* Richardson, TX.

Awduche, D. J., Malcolm, J., O'Dell, M., & McManus, J. (1999, September). *RFC2702: Requirements for traffic engineering over MPLS.* Retrieved March 1, 2009, from Internet Engineering Task Force: http:\\ietf.org/rfc/rfc2702.txt

Black, U., & Waters, S. (2002). *SONET and T-1: Architecture for Digital Transport Networks* (2nd ed. ed.). Upper Saddle River, NJ: Prentice-Hall.

Brooke; Kendrick; Meeraus; Raman;. (n.d.). *GAMS-A User's Guide.* Retrieved January 5, 2009, from GAMS Online Documentation: http://www.gams.com/docs/document.htm

Caenegem, B. V., Wauters, N., & Demeester, P. (1997). Spare capacity assignment for different restoration strategies in mesh survivable networks. *IEEE International Conference on Communications 1997*, (pp. 288-292).

Cisco Systems. Inc. (n.d.). *Understanding MPLS-TP and Its Benefits.* Retrieved Februrary 21, 2010, from Cisco Systems: http://www.cisco.com/en/US/technologies/tk436/tk428/white_paper_c11-562013.html

Davie, B. S., & Farrell, A. (2008). *MPLS: Next Steps.* San Francisco: Morgan Kaufmann.

Dirceu, C. (June 2000). Evolution of Optical Transport Technologies from SONET/SDH to WDM. *IEEE Communications* .

Doucette, J., & Grover, W. D. (2001). Comparison of mesh protection and restoration schemes and dependency on graph connectivity. *3rd International Workshop on Design of Reliable Communications Networks*, (pp. 121-128).

Doverspike, R. D., & Wilson, B. (1994). Comparison of capacity efficiency of DCS network restoration routing techniques. *Journal of Network Sytems Management , 2* (2), 95-123.

Dunn, D. A., Grover, W. D., & MacGregor, M. H. (1994). A comparison of k-shortest path selection for shared restoration connections. *IEEE Journal on Selected Areas in Communications , 12* (1), 88-99.

Farrel, A., & Bryskin, I. (2006). *GMPLS: Architecture and Applications.* San Francisco: Morgan Kaufmann.

Fourer, R., Gay, D. M., & Kernighan, B. W. (2003). *AMPL: A Modeling Language for Mathematical Programming* (2nd ed.). Pacific Grove, CA: Brooks/Cole-Thomson Learning.

George, D., & Mallery, P. (2010). *SPSS for Windows* (10th ed.). Boston: Allyn & Bacon.

Graziani, R., & Johnson, A. (2009). *Routing Protocols and Concepts: CCNA Exploration Companion Guide.* Indianapolis: Cisco Press.

Grover, W. D. (2003). *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking.* Engelwood Cliffs, NJ: Prentice-Hall.

Grover, W. D., Bilodeaux, T. D., & Venables, B. D. (1991). Near optimal spare capacity planning in a mesh restorable network. *Proceedings of IEEE GLOBECOM '91* , 2007-2012.

Gubbins, E. (2009, September 14). *Debating 40G and 100G Economics.* Retrieved September 14, 2009, from Connected Planet: http://connectedplanetonline.com/business_services/news/40g-100g-debate-0914/

Guichard, J., Le Faucheur, F., & Vasseur, J. P. (2005). *Definitive MPLS Network Designs.* Indianapolis: Cisco Press.

Huitema, C. (1995). *Routing in the Internet.* Upper Saddle River, NJ: Prentice-Hall.

Iraschko, R. M., MacGregor, W. D., & Grover, W. D. (1998). Optimal capacity placement for path restoration in STM or ATM mesh survivable networks. *IEEE/ACM Transactions on Networking , 6* (3), 325-336.

Kennington, J. L., Nair, V. S., & Spiride, G. (1998). *Optimal spare capacity assignment for path restorable mesh networks: cuts, decomposition, and an empirical analysis.* Dallas: Southern Methodist University Department of Computer Science and Engineering.

Kennington, J. L., Olinick, E. V., & Spiride, G. (2006). Basic mathematical programming models for capacity allocation in mesh-based survivable networks. *International Journal of Management Science , 35*, 1-16.

Lang, J. P., Rajagopalan, B., & Papadimitriou, D. (2006, March). *RFC4426: Generalized Multiprotocol Label Switching (GMPLS) recovery functional specification.* Retrieved April 15, 2009, from Internet Engineering Task Force: http://ietfreport.isoc.org/idref/rfc4427/index.html

Li, G., Kalmanek, C., & Doverspike, R. (2002). Efficient distributed path selection for shared restoration connections. *Proceedings of IEEE INFOCOM*, (pp. 140-149).

Mannie, E., & Papadimitriou, D. (2006, March). *RFC4427: Recovery terminology for Generalized Multiprotocol Label Switching (GMPLS)* . Retrieved May 7, 2009, from Internet Engineering Task Force.

Martin, R., & Menth, M. (2006, April). Backup capacity requirements for MPLS fast reroute. *7. ITG Fachtagung Photonische Netze (VDE)* .

McLeod, T. W. *Constrained Integer Network Flows.* Dallas: Southern Methodist University.

Medhi, D., & Ramasamy, K. (2007). *Network Routing: Algorithms, Protocols, and Architectures.* San Francisco: Morgan Kaufmann Publishers.

Menth, M., Martin, R., Hartmann, M., & Sporlein, U. (2009). Effiency of routing and resilience mechanisms in packet-switched communication networks. *European Transactions on telecommunications , 21* (2), 108-120.

Menth, M., Reifert, A., & Milbrandt, J. (2004). *Optimization of End-to-End Protection Switching Mechanisms for MPLS Networks.* Wurzburg: University of Wurzburg Institute of Computer Science.

Minei, I., & Lucek, J. (2008). *MPLS-Enabled Application: Emerging Developments and New Technologies* (2nd ed. ed.). Chichester, West Sussex, England: John Wiley & Sons.

Morrow, M. J., & Sayeed, A. (2007). *MPLS and Next Generation Networks.* Indianapolis: Cisco Press.

Moy, J. T. (1998). *Anatomy of an Internet Routing Protocol.* Boston: Addison-Wesley-Longman.

Murakami, K., & Kim, H. (1998). Optimal capacity and flow assignment for self-healing ATM networks based on line and end-to-end restoration. *IEEE/ACM Transaction on Networking , 6* (2), 207-221.

National Institute of Science and Technology (NIST). (n.d.). *Engnineering Statisitics Handbook.* Retrieved March 8, 20009, from http://www.itl.nist.gov/div898/handbook/prc/section4/prc471.htm

Network Reliability and Interoperability Council V1, F. G. (2003). *Network Reliability Final Report.* Washington, D.C.: Federal Communications Commission (FCC).

Norusis, M. J. (2010). *PASW Statistics 18 Guide to Data Analysis.* Upper Saddle River, NJ: Prentice Hall.

Olinick, E. *Path selection algorithm.*

one, N. t. (n.d.).

Orlowski, S., & Wessaly, R. (2003). Comparing restoration concepts using optimal network configurations with integrated hardware and routing decisions. *Proceedings of the Fourth International Workshop on the Design of Reliable Communication Networks.*

Osborne, E., & Simha, A. (2003). *Traffic Engineering with MPLS.* Indianapolis: Cisco Press.

Pan, P., Swallow, G., & Atlas, A. (2005, May). *RFC4090: Fast reroute extensions to RSVP-TE for LSP tunnels.* Retrieved April 5, 2009, from Internet Engineering Task Force: http:\\tools.ietf.org/html/rfc4090

Papadimitriou, D., & Mannie, E. (2006, March). *Analysis of Generalized Multiprotocol Label Switching (GMPLS)-based recovery mechanisms (including protection and restoration).* Retrieved April 15, 2009, from Internet Engineering Task Force: http://ietfreport.isoc.org/idref/rfc4428/

Perrin, S. (2010, April 7). *100G vs. 40G and 100G Economics.* Retrieved April 7, 2010, from Lightreading: http://www.lightreading.com/webinar.asp?webinar_id=29379

Pioro, M., & Medhi, D. (2004). *Routing, Flow and Capacity Design in Communication and Computer Networks.* San Francisco: Morgan Kaufmann.

Rosen, E. C., Viswanathan, A., & Callon, R. (2001, January). *RFC3031: MPLS Architecture.* Retrieved March 1, 2009, from Internet Engineering Task Force: http://www.ietf.org/rfc/rfc3031.txt

Vassuer, J.-P., Pickavet, M., & Demeester, P. (2004). *Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS.* San Francisco: Morgan Kaufmann.

Venables, B. D., Grover, W. D., & MacGregor, M. H. (1993). Two stategies for spare capacity placement in mesh restorable networks. *Proceedings of IEEE ICC '93*, (pp. 267-271).

Xiong, Y., & Mason, L. (1997). Restoration strategies and spare capacity requirements in self-healing ATM networks. *IEEE INFOCOM '97*, (pp. 352-360).

Yamanaka, N., Shiomoto, K., & Oki, O. (2006). *GMPLS Technologies: Broadband Backbone Networks and Systems.* Boca Raton: CRC Press.

Zhuo, F. X. (2010). *Designing and Implementing IP/MPLS-Based Ethernet Layer 2 VPN Services: An Advanced Guide for VPLS and VLL.* Indianapolis: Wiley .

7 Appendix

7.1 Appendix A: Cost ranking of networks

Links/Demand unit/Recovery method	**Cost**
133 links/10Gbps/Path	$7,655,032
133 links/10Gbps/FRR	$8,464,222
133 links/40Gbps/Path	$8,507,353
95 links/10Gbps/Path	$8,672,277
133 links/40Gbps/FRR	$8,899,037
133 links/10Gbps/Hybrid	$9,283,204
95 links/40Gbps/Path	$9,445,821
133 links/100Gbps/FRR	$9,840,741
133 links/100Gbps/Path	$9,958,796
95 links/10Gbps/FRR	$9,987,149
133 links/40Gbps/Hybrid	$10,167,277
95 links/40Gbps/FRR	$10,256,269
95 links/100Gbps/Path	$10,532,407
95 links/10Gbps/Hybrid	$10,700,265
95 links/100Gbps/FRR	$11,081,098
95 links/40Gbps/1:1	$11,097,778
57 links/10Gbps/Path	$11,190,187
133 links/10Gbps/1:1	$11,198,889
57 links/100Gbps/1:1	$11,364,444
95 links/40Gbps/Hybrid	$11,541,735
57 links/40Gbps/Path	$11,960,090
133 links/100Gbps/Hybrid	$12,376,759
95 links/100Gbps/1:1	$13,112,444
133 links/40Gbps/1:1	$13,324,444
57 links/10Gbps/FRR	$13,393,650
57 links/10Gbps/Hybrid	$13,476,109
31 links/10Gbps/1:1	$13,494,444
57 links/100Gbps/Path	$13,544,167
57 links/40Gbps/FRR	$13,631,917
95 links/100Gbps/Hybrid	$13,659,438
57 links/100Gbps/FRR	$14,512,226
57 links/40Gbps/Hybrid	$15,128,888
57 links/100Gbps/Hybrid	$16,871,650

133 links/100Gbps/1:1	$16,920,000
57 links/10Gbps/1:1	$17,151,667
31 links/40Gbps/1:1	$17,322,222
31 links/10Gbps/FRR	$19,053,079
31 links/40Gbps/FRR	$19,144,630
31 links/10Gbps/Path	$19,952,750
31 links/100Gbps/FRR	$20,086,891
31 links/40Gbps/Path	$20,356,667
31 links/10Gbps/Hybrid	$21,108,755
31 links/40Gbps/Hybrid	$21,917,864
31 links/100Gbps/Path	$22,008,333
57 links/40Gbps/1:1	$23,760,000
95 links/10Gbps/1:1	$23,893,333
31 links/100Gbps/1:1	$24,133,333
31 links/100Gbps/Hybrid	$24,964,279

Table 19: Cost ranking of networks when considering number of links, demand unit, and recovery method

7.2 Appendix B: Network migration cost sorted by demand unit and number of links

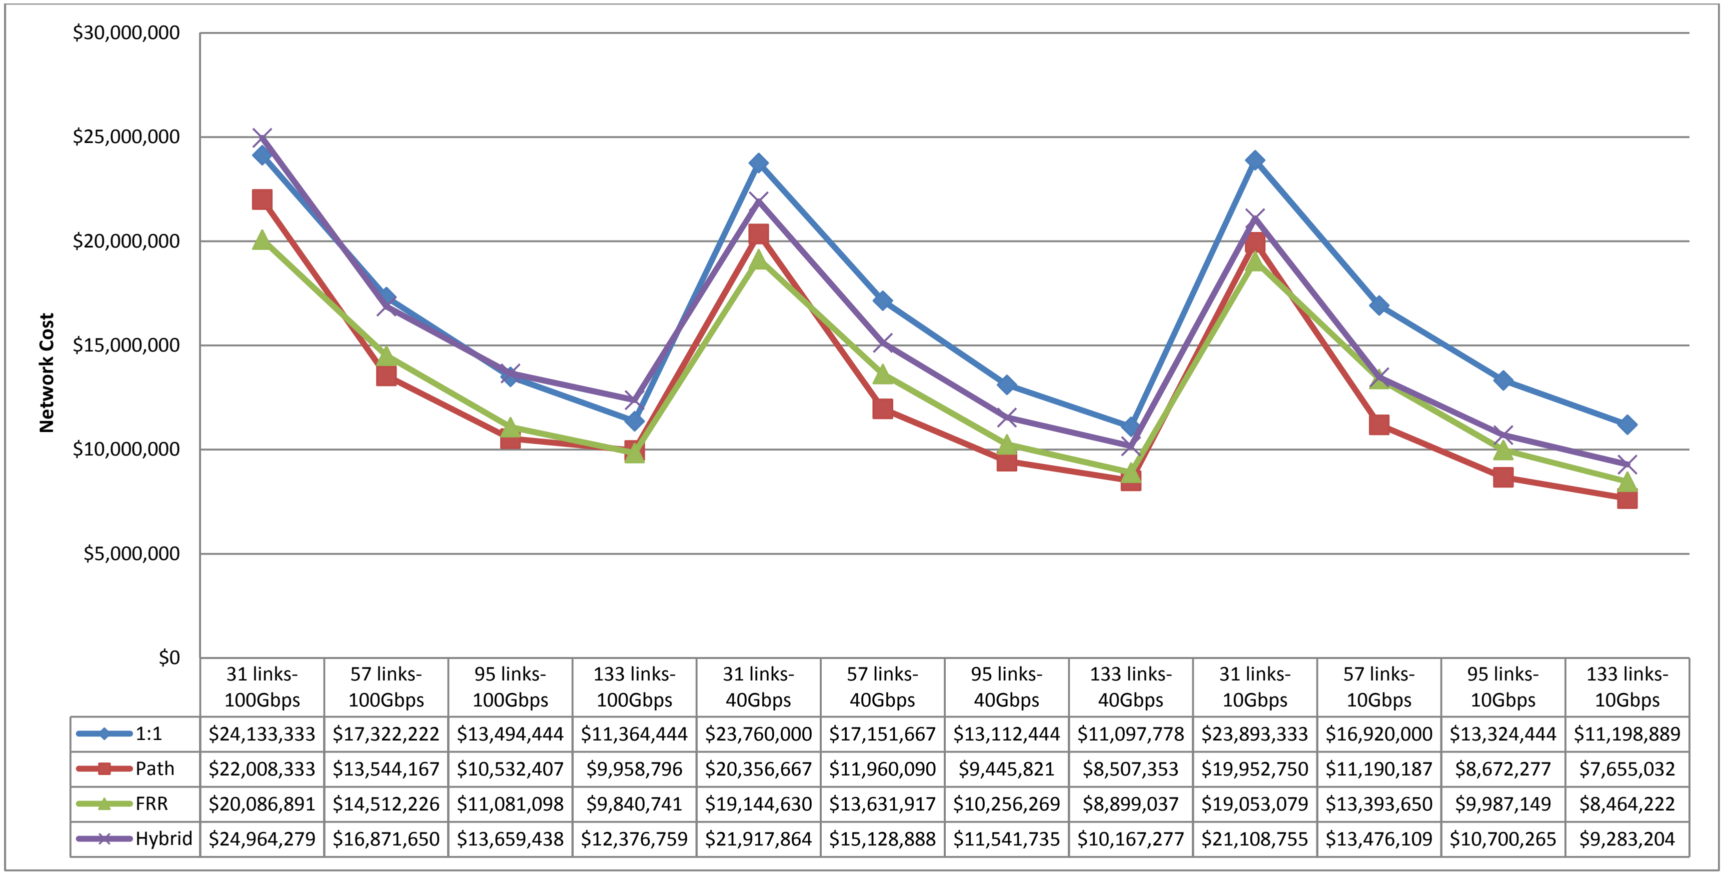

	31 links-100Gbps	57 links-100Gbps	95 links-100Gbps	133 links-100Gbps	31 links-40Gbps	57 links-40Gbps	95 links-40Gbps	133 links-40Gbps	31 links-10Gbps	57 links-10Gbps	95 links-10Gbps	133 links-10Gbps
◆ 1:1	$24,133,333	$17,322,222	$13,494,444	$11,364,444	$23,760,000	$17,151,667	$13,112,444	$11,097,778	$23,893,333	$16,920,000	$13,324,444	$11,198,889
■ Path	$22,008,333	$13,544,167	$10,532,407	$9,958,796	$20,356,667	$11,960,090	$9,445,821	$8,507,353	$19,952,750	$11,190,187	$8,672,277	$7,655,032
▲ FRR	$20,086,891	$14,512,226	$11,081,098	$9,840,741	$19,144,630	$13,631,917	$10,256,269	$8,899,037	$19,053,079	$13,393,650	$9,987,149	$8,464,222
× Hybrid	$24,964,279	$16,871,650	$13,659,438	$12,376,759	$21,917,864	$15,128,888	$11,541,735	$10,167,277	$21,108,755	$13,476,109	$10,700,265	$9,283,204

Figure 13: Cost effect comparison of different link densities and demand units to restoration methods (grouped by demand units)

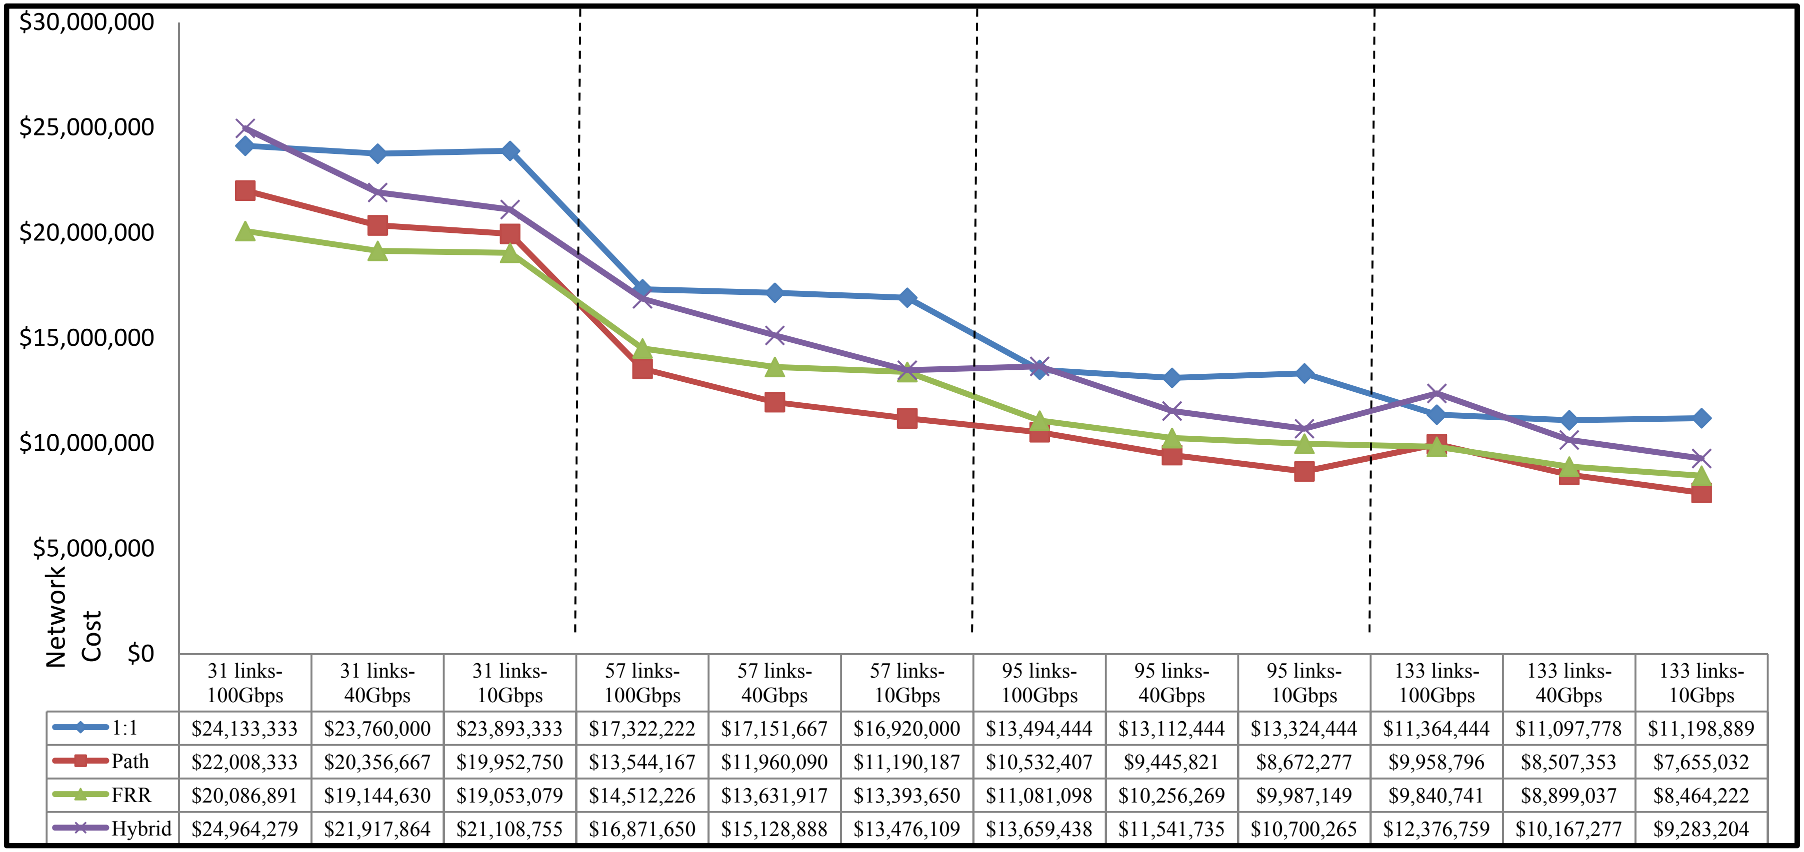

	31 links-100Gbps	31 links-40Gbps	31 links-10Gbps	57 links-100Gbps	57 links-40Gbps	57 links-10Gbps	95 links-100Gbps	95 links-40Gbps	95 links-10Gbps	133 links-100Gbps	133 links-40Gbps	133 links-10Gbps
1:1	$24,133,333	$23,760,000	$23,893,333	$17,322,222	$17,151,667	$16,920,000	$13,494,444	$13,112,444	$13,324,444	$11,364,444	$11,097,778	$11,198,889
Path	$22,008,333	$20,356,667	$19,952,750	$13,544,167	$11,960,090	$11,190,187	$10,532,407	$9,445,821	$8,672,277	$9,958,796	$8,507,353	$7,655,032
FRR	$20,086,891	$19,144,630	$19,053,079	$14,512,226	$13,631,917	$13,393,650	$11,081,098	$10,256,269	$9,987,149	$9,840,741	$8,899,037	$8,464,222
Hybrid	$24,964,279	$21,917,864	$21,108,755	$16,871,650	$15,128,888	$13,476,109	$13,659,438	$11,541,735	$10,700,265	$12,376,759	$10,167,277	$9,283,204

Figure 14: Cost effect comparison of number of links at different demand units to restoration methods (grouped by number of links)

7.3 Appendix C: Detailed migration comparisons

	31 links	31 links	31 links	57 links	57 links	57 links	95 links	95 links	95 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	1:1	1:1	1:1	1:1	1:1	1:1	1:1	1:1	1:1
31 links/100Gbps/1:1		-1.55%	-0.99%	-28.22%	-28.93%	-29.89%	-44.08%	-45.67%	-44.79%
31 links/40Gbps/1:1	1.57%		0.56%	-27.10%	-27.81%	-28.79%	-43.21%	-44.81%	-43.92%
31 links/10Gbps/1:1	1.00%	-0.56%		-27.50%	-28.22%	-29.19%	-43.52%	-45.12%	-44.23%
57 links/100Gbps/1:1	39.32%	37.16%	37.93%		-0.98%	-2.32%	-22.10%	-24.30%	-23.08%
57 links/40Gbps/1:1	40.71%	38.53%	39.31%	0.99%		-1.35%	-21.32%	-23.55%	-22.31%
57 links/10Gbps/1:1	42.63%	40.43%	41.21%	2.38%	1.37%		-20.25%	-22.50%	-21.25%
95 links/100Gbps/1:1	78.84%	76.07%	77.06%	28.37%	27.10%	25.38%		-2.83%	-1.26%
95 links/40Gbps/1:1	84.05%	81.20%	82.22%	32.11%	30.80%	29.04%	2.91%		1.62%
95 links/10Gbps/1:1	81.12%	78.32%	79.32%	30.00%	28.72%	26.98%	1.28%	-1.59%	
133 links/100Gbps/1:1	112.36%	109.07%	110.25%	52.42%	50.92%	48.89%	18.74%	15.38%	17.25%
133 links/40Gbps/1:1	117.46%	114.10%	115.30%	56.09%	54.55%	52.46%	21.60%	18.15%	20.06%
133 links/10Gbps/1:1	115.50%	112.16%	113.35%	54.68%	53.16%	51.09%	20.50%	17.09%	18.98%
31 links/100Gbps/Path	9.66%	7.96%	8.56%	-21.29%	-22.07%	-23.12%	-38.68%	-40.42%	-39.46%
31 links/40Gbps/Path	18.55%	16.72%	17.37%	-14.91%	-15.74%	-16.88%	-33.71%	-35.59%	-34.55%
31 links/10Gbps/Path	20.95%	19.08%	19.75%	-13.18%	-14.04%	-15.20%	-32.37%	-34.28%	-33.22%
57 links/100Gbps/Path	78.18%	75.43%	76.41%	27.89%	26.64%	24.92%	-0.37%	-3.19%	-1.62%
57 links/40Gbps/Path	101.78%	98.66%	99.78%	44.83%	43.41%	41.47%	12.83%	9.63%	11.41%
57 links/10Gbps/Path	115.67%	112.33%	113.52%	54.80%	53.27%	51.20%	20.59%	17.18%	19.07%
95 links/100Gbps/Path	129.13%	125.59%	126.86%	64.47%	62.85%	60.65%	28.12%	24.50%	26.51%
95 links/40Gbps/Path	155.49%	151.54%	152.95%	83.39%	81.58%	79.13%	42.86%	38.82%	41.06%
95 links/10Gbps/Path	178.28%	173.98%	175.51%	99.74%	97.78%	95.10%	55.60%	51.20%	53.64%
133 links/100Gbps/Path	142.33%	138.58%	139.92%	73.94%	72.23%	69.90%	35.50%	31.67%	33.80%
133 links/40Gbps/Path	183.68%	179.29%	180.86%	103.61%	101.61%	98.89%	58.62%	54.13%	56.62%
133 links/10Gbps/Path	215.26%	210.38%	212.13%	126.29%	124.06%	121.03%	76.28%	71.29%	74.06%

	31 links	31 links	31 links	57 links	57 links	57 links	95 links	95 links	95 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	1:1	1:1	1:1	1:1	1:1	1:1	1:1	1:1	1:1
31 links/100Gbps/FRR	20.14%	18.29%	18.95%	-13.76%	-14.61%	-15.77%	-32.82%	-34.72%	-33.67%
31 links/40Gbps/FRR	26.06%	24.11%	24.80%	-9.52%	-10.41%	-11.62%	-29.51%	-31.51%	-30.40%
31 links/10Gbps/FRR	26.66%	24.70%	25.40%	-9.08%	-9.98%	-11.20%	-29.17%	-31.18%	-30.07%
57 links/100Gbps/FRR	66.30%	63.72%	64.64%	19.36%	18.19%	16.59%	-7.01%	-9.65%	-8.18%
57 links/40Gbps/FRR	77.04%	74.30%	75.27%	27.07%	25.82%	24.12%	-1.01%	-3.81%	-2.26%
57 links/10Gbps/FRR	80.18%	77.40%	78.39%	29.33%	28.06%	26.33%	0.75%	-2.10%	-0.52%
95 links/100Gbps/FRR	117.79%	114.42%	115.62%	56.32%	54.78%	52.69%	21.78%	18.33%	20.24%
95 links/40Gbps/FRR	135.30%	131.66%	132.96%	68.89%	67.23%	64.97%	31.57%	27.85%	29.92%
95 links/10Gbps/FRR	141.64%	137.91%	139.24%	73.45%	71.74%	69.42%	35.12%	31.29%	33.42%
133 links/100Gbps/FRR	145.24%	141.45%	142.80%	76.03%	74.29%	71.94%	37.13%	33.25%	35.40%
133 links/40Gbps/FRR	171.19%	167.00%	168.49%	94.65%	92.74%	90.13%	51.64%	47.35%	49.73%
133 links/10Gbps/FRR	185.12%	180.71%	182.29%	104.65%	102.64%	99.90%	59.43%	54.92%	57.42%
31 links/100Gbps/Hybrid	-3.33%	-4.82%	-4.29%	-30.61%	-31.30%	-32.22%	-45.94%	-47.48%	-46.63%
31 links/40Gbps/Hybrid	10.11%	8.40%	9.01%	-20.97%	-21.75%	-22.80%	-38.43%	-40.17%	-39.21%
31 links/10Gbps/Hybrid	14.33%	12.56%	13.19%	-17.94%	-18.75%	-19.84%	-36.07%	-37.88%	-36.88%
57 links/100Gbps/Hybrid	43.04%	40.83%	41.62%	2.67%	1.66%	0.29%	-20.02%	-22.28%	-21.02%
57 links/40Gbps/Hybrid	59.52%	57.05%	57.93%	14.50%	13.37%	11.84%	-10.80%	-13.33%	-11.93%
57 links/10Gbps/Hybrid	79.08%	76.31%	77.30%	28.54%	27.27%	25.56%	0.14%	-2.70%	-1.13%
95 links/100Gbps/Hybrid	76.68%	73.95%	74.92%	26.82%	25.57%	23.87%	-1.21%	-4.00%	-2.45%
95 links/40Gbps/Hybrid	109.10%	105.86%	107.02%	50.08%	48.61%	46.60%	16.92%	13.61%	15.45%
95 links/10Gbps/Hybrid	125.54%	122.05%	123.30%	61.89%	60.29%	58.13%	26.11%	22.54%	24.52%
133 links/100Gbps/Hybrid	94.99%	91.97%	93.05%	39.96%	38.58%	36.71%	9.03%	5.94%	7.66%
133 links/40Gbps/Hybrid	137.36%	133.69%	135.00%	70.37%	68.69%	66.42%	32.72%	28.97%	31.05%
133 links/10Gbps/Hybrid	159.97%	155.95%	157.38%	86.60%	84.76%	82.26%	45.36%	41.25%	43.53%

	133 links	133 links	133 links	31 links	31 links	31 links	57 links	57 links	57 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	1:1	1:1	1:1	Path	Path	Path	Path	Path	Path
31 links/100Gbps/1:1	-52.91%	-54.01%	-53.60%	-8.81%	-15.65%	-17.32%	-43.88%	-50.44%	-53.63%
31 links/40Gbps/1:1	-52.17%	-53.29%	-52.87%	-7.37%	-14.32%	-16.02%	-43.00%	-49.66%	-52.90%
31 links/10Gbps/1:1	-52.44%	-53.55%	-53.13%	-7.89%	-14.80%	-16.49%	-43.31%	-49.94%	-53.17%
57 links/100Gbps/1:1	-34.39%	-35.93%	-35.35%	27.05%	17.52%	15.19%	-21.81%	-30.96%	-35.40%
57 links/40Gbps/1:1	-33.74%	-35.30%	-34.71%	28.32%	18.69%	16.33%	-21.03%	-30.27%	-34.76%
57 links/10Gbps/1:1	-32.83%	-34.41%	-33.81%	30.07%	20.31%	17.92%	-19.95%	-29.31%	-33.86%
95 links/100Gbps/1:1	-15.78%	-17.76%	-17.01%	63.09%	50.85%	47.86%	0.37%	-11.37%	-17.08%
95 links/40Gbps/1:1	-13.33%	-15.36%	-14.59%	67.84%	55.25%	52.17%	3.29%	-8.79%	-14.66%
95 links/10Gbps/1:1	-14.71%	-16.71%	-15.95%	65.17%	52.78%	49.75%	1.65%	-10.24%	-16.02%
133 links/100Gbps/1:1		-2.35%	-1.46%	93.66%	79.13%	75.57%	19.18%	5.24%	-1.53%
133 links/40Gbps/1:1	2.40%		0.91%	98.31%	83.43%	79.79%	22.04%	7.77%	0.83%
133 links/10Gbps/1:1	1.48%	-0.90%		96.52%	81.77%	78.17%	20.94%	6.80%	-0.08%
31 links/100Gbps/Path	-48.36%	-49.57%	-49.12%		-7.50%	-9.34%	-38.46%	-45.66%	-49.15%
31 links/40Gbps/Path	-44.17%	-45.48%	-44.99%	8.11%		-1.98%	-33.47%	-41.25%	-45.03%
31 links/10Gbps/Path	-43.04%	-44.38%	-43.87%	10.30%	2.02%		-32.12%	-40.06%	-43.92%
57 links/100Gbps/Path	-16.09%	-18.06%	-17.32%	62.49%	50.30%	47.32%		-11.70%	-17.38%
57 links/40Gbps/Path	-4.98%	-7.21%	-6.36%	84.01%	70.20%	66.83%	13.24%		-6.44%
57 links/10Gbps/Path	1.56%	-0.83%	0.08%	96.68%	81.92%	78.31%	21.04%	6.88%	
95 links/100Gbps/Path	7.90%	5.37%	6.33%	108.96%	93.28%	89.44%	28.60%	13.56%	6.25%
95 links/40Gbps/Path	20.31%	17.49%	18.56%	133.00%	115.51%	111.23%	43.39%	26.62%	18.47%
95 links/10Gbps/Path	31.04%	27.97%	29.13%	153.78%	134.73%	130.08%	56.18%	37.91%	29.03%
133 links/100Gbps/Path	14.11%	11.44%	12.45%	120.99%	104.41%	100.35%	36.00%	20.10%	12.36%
133 links/40Gbps/Path	33.58%	30.45%	31.64%	158.70%	139.28%	134.54%	59.21%	40.59%	31.54%
133 links/10Gbps/Path	48.46%	44.97%	46.29%	187.50%	165.93%	160.65%	76.93%	56.24%	46.18%

	133 links	133 links	133 links	31 links	31 links	31 links	57 links	57 links	57 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	1:1	1:1	1:1	Path	Path	Path	Path	Path	Path
31 links/100Gbps/FRR	-43.42%	-44.75%	-44.25%	9.57%	1.34%	-0.67%	-32.57%	-40.46%	-44.29%
31 links/40Gbps/FRR	-40.64%	-42.03%	-41.50%	14.96%	6.33%	4.22%	-29.25%	-37.53%	-41.55%
31 links/10Gbps/FRR	-40.35%	-41.75%	-41.22%	15.51%	6.84%	4.72%	-28.91%	-37.23%	-41.27%
57 links/100Gbps/FRR	-21.69%	-23.53%	-22.83%	51.65%	40.27%	37.49%	-6.67%	-17.59%	-22.89%
57 links/40Gbps/FRR	-16.63%	-18.59%	-17.85%	61.45%	49.33%	46.37%	-0.64%	-12.26%	-17.91%
57 links/10Gbps/FRR	-15.15%	-17.14%	-16.39%	64.32%	51.99%	48.97%	1.12%	-10.70%	-16.45%
95 links/100Gbps/FRR	2.56%	0.15%	1.06%	98.61%	83.71%	80.06%	22.23%	7.93%	0.98%
95 links/40Gbps/FRR	10.80%	8.20%	9.19%	114.58%	98.48%	94.54%	32.06%	16.61%	9.11%
95 links/10Gbps/FRR	13.79%	11.12%	12.13%	120.37%	103.83%	99.78%	35.62%	19.75%	12.05%
133 links/100Gbps/FRR	15.48%	12.77%	13.80%	123.65%	106.86%	102.76%	37.63%	21.54%	13.71%
133 links/40Gbps/FRR	27.70%	24.71%	25.84%	147.31%	128.75%	124.21%	52.20%	34.40%	25.75%
133 links/10Gbps/FRR	34.26%	31.11%	32.31%	160.02%	140.50%	135.73%	60.02%	41.30%	32.21%
31 links/100Gbps/Hybrid	-54.48%	-55.55%	-55.14%	-11.84%	-18.46%	-20.07%	-45.75%	-52.09%	-55.18%
31 links/40Gbps/Hybrid	-48.15%	-49.37%	-48.91%	0.41%	-7.12%	-8.97%	-38.20%	-45.43%	-48.94%
31 links/10Gbps/Hybrid	-46.16%	-47.43%	-46.95%	4.26%	-3.56%	-5.48%	-35.84%	-43.34%	-46.99%
57 links/100Gbps/Hybrid	-32.64%	-34.22%	-33.62%	30.45%	20.66%	18.26%	-19.72%	-29.11%	-33.67%
57 links/40Gbps/Hybrid	-24.88%	-26.65%	-25.98%	45.47%	34.55%	31.89%	-10.47%	-20.95%	-26.03%
57 links/10Gbps/Hybrid	-15.67%	-17.65%	-16.90%	63.31%	51.06%	48.06%	0.51%	-11.25%	-16.96%
95 links/100Gbps/Hybrid	-16.80%	-18.75%	-18.01%	61.12%	49.03%	46.07%	-0.84%	-12.44%	-18.08%
95 links/40Gbps/Hybrid	-1.54%	-3.85%	-2.97%	90.68%	76.37%	72.87%	17.35%	3.62%	-3.05%
95 links/10Gbps/Hybrid	6.21%	3.71%	4.66%	105.68%	90.24%	86.47%	26.58%	11.77%	4.58%
133 links/100Gbps/Hybrid	-8.18%	-10.33%	-9.52%	77.82%	64.47%	61.21%	9.43%	-3.37%	-9.59%
133 links/40Gbps/Hybrid	11.77%	9.15%	10.15%	116.46%	100.22%	96.24%	33.21%	17.63%	10.06%
133 links/10Gbps/Hybrid	22.42%	19.55%	20.64%	137.08%	119.28%	114.93%	45.90%	28.84%	20.54%

	95 links	95 links	95 links	133 links	133 links	133 links	31 links	31 links	31 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	Path	Path	Path	Path	Path	Path	FRR	FRR	FRR
31 links/100Gbps/1:1	-56.36%	-60.86%	-64.07%	-58.73%	-64.75%	-68.28%	-16.77%	-20.67%	-21.05%
31 links/40Gbps/1:1	-55.67%	-60.24%	-63.50%	-58.09%	-64.19%	-67.78%	-15.46%	-19.42%	-19.81%
31 links/10Gbps/1:1	-55.92%	-60.47%	-63.70%	-58.32%	-64.39%	-67.96%	-15.93%	-19.87%	-20.26%
57 links/100Gbps/1:1	-39.20%	-45.47%	-49.94%	-42.51%	-50.89%	-55.81%	15.96%	10.52%	9.99%
57 links/40Gbps/1:1	-38.59%	-44.93%	-49.44%	-41.94%	-50.40%	-55.37%	17.11%	11.62%	11.09%
57 links/10Gbps/1:1	-37.75%	-44.17%	-48.75%	-41.14%	-49.72%	-54.76%	18.72%	13.15%	12.61%
95 links/100Gbps/1:1	-21.95%	-30.00%	-35.73%	-26.20%	-36.96%	-43.27%	48.85%	41.87%	41.19%
95 links/40Gbps/1:1	-19.68%	-27.96%	-33.86%	-24.05%	-35.12%	-41.62%	53.19%	46.00%	45.31%
95 links/10Gbps/1:1	-20.95%	-29.11%	-34.91%	-25.26%	-36.15%	-42.55%	50.75%	43.68%	42.99%
133 links/100Gbps/1:1	-7.32%	-16.88%	-23.69%	-12.37%	-25.14%	-32.64%	76.75%	68.46%	67.66%
133 links/40Gbps/1:1	-5.09%	-14.89%	-21.86%	-10.26%	-23.34%	-31.02%	81.00%	72.51%	71.68%
133 links/10Gbps/1:1	-5.95%	-15.65%	-22.56%	-11.07%	-24.03%	-31.64%	79.37%	70.95%	70.13%
31 links/100Gbps/Path	-52.14%	-57.08%	-60.60%	-54.75%	-61.34%	-65.22%	-8.73%	-13.01%	-13.43%
31 links/40Gbps/Path	-48.26%	-53.60%	-57.40%	-51.08%	-58.21%	-62.40%	-1.33%	-5.95%	-6.40%
31 links/10Gbps/Path	-47.21%	-52.66%	-56.54%	-50.09%	-57.36%	-61.63%	0.67%	-4.05%	-4.51%
57 links/100Gbps/Path	-22.24%	-30.26%	-35.97%	-26.47%	-37.19%	-43.48%	48.31%	41.35%	40.67%
57 links/40Gbps/Path	-11.94%	-21.02%	-27.49%	-16.73%	-28.87%	-36.00%	67.95%	60.07%	59.31%
57 links/10Gbps/Path	-5.88%	-15.59%	-22.50%	-11.00%	-23.97%	-31.59%	79.50%	71.08%	70.27%
95 links/100Gbps/Path		-10.32%	-17.66%	-5.45%	-19.23%	-27.32%	90.72%	81.77%	80.90%
95 links/40Gbps/Path	11.50%		-8.19%	5.43%	-9.94%	-18.96%	112.65%	102.68%	101.71%
95 links/10Gbps/Path	21.45%	8.92%		14.83%	-1.90%	-11.73%	131.62%	120.76%	119.70%
133 links/100Gbps/Path	5.76%	-5.15%	-12.92%		-14.57%	-23.13%	101.70%	92.24%	91.32%
133 links/40Gbps/Path	23.80%	11.03%	1.94%	17.06%		-10.02%	136.11%	125.04%	123.96%
133 links/10Gbps/Path	37.59%	23.39%	13.29%	30.09%	11.13%		162.40%	150.09%	148.90%

	95 links	95 links	95 links	133 links	133 links	133 links	31 links	31 links	31 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	Path	Path	Path	Path	Path	Path	FRR	FRR	FRR
31 links/100Gbps/FRR	-47.57%	-52.98%	-56.83%	-50.42%	-57.65%	-61.89%		-4.69%	-5.15%
31 links/40Gbps/FRR	-44.99%	-50.66%	-54.70%	-47.98%	-55.56%	-60.01%	4.92%		-0.48%
31 links/10Gbps/FRR	-44.72%	-50.42%	-54.48%	-47.73%	-55.35%	-59.82%	5.43%	0.48%	
57 links/100Gbps/FRR	-27.42%	-34.91%	-40.24%	-31.38%	-41.38%	-47.25%	38.41%	31.92%	31.29%
57 links/40Gbps/FRR	-22.74%	-30.71%	-36.38%	-26.95%	-37.59%	-43.84%	47.35%	40.44%	39.77%
57 links/10Gbps/FRR	-21.36%	-29.48%	-35.25%	-25.65%	-36.48%	-42.85%	49.97%	42.94%	42.25%
95 links/100Gbps/FRR	-4.95%	-14.76%	-21.74%	-10.13%	-23.23%	-30.92%	81.27%	72.77%	71.94%
95 links/40Gbps/FRR	2.69%	-7.90%	-15.44%	-2.90%	-17.05%	-25.36%	95.85%	86.66%	85.77%
95 links/10Gbps/FRR	5.46%	-5.42%	-13.17%	-0.28%	-14.82%	-23.35%	101.13%	91.69%	90.78%
133 links/100Gbps/FRR	7.03%	-4.01%	-11.87%	1.20%	-13.55%	-22.21%	104.12%	94.54%	93.61%
133 links/40Gbps/FRR	18.35%	6.14%	-2.55%	11.91%	-4.40%	-13.98%	125.72%	115.13%	114.10%
133 links/10Gbps/FRR	24.43%	11.60%	2.46%	17.66%	0.51%	-9.56%	137.32%	126.18%	125.10%
31 links/100Gbps/Hybrid	-57.81%	-62.16%	-65.26%	-60.11%	-65.92%	-69.34%	-19.54%	-23.31%	-23.68%
31 links/40Gbps/Hybrid	-51.95%	-56.90%	-60.43%	-54.56%	-61.19%	-65.07%	-8.35%	-12.65%	-13.07%
31 links/10Gbps/Hybrid	-50.10%	-55.25%	-58.92%	-52.82%	-59.70%	-63.74%	-4.84%	-9.30%	-9.74%
57 links/100Gbps/Hybrid	-37.57%	-44.01%	-48.60%	-40.97%	-49.58%	-54.63%	19.06%	13.47%	12.93%
57 links/40Gbps/Hybrid	-30.38%	-37.56%	-42.68%	-34.17%	-43.77%	-49.40%	32.77%	26.54%	25.94%
57 links/10Gbps/Hybrid	-21.84%	-29.91%	-35.65%	-26.10%	-36.87%	-43.20%	49.06%	42.06%	41.38%
95 links/100Gbps/Hybrid	-22.89%	-30.85%	-36.51%	-27.09%	-37.72%	-43.96%	47.06%	40.16%	39.49%
95 links/40Gbps/Hybrid	-8.75%	-18.16%	-24.86%	-13.71%	-26.29%	-33.68%	74.04%	65.87%	65.08%
95 links/10Gbps/Hybrid	-1.57%	-11.72%	-18.95%	-6.93%	-20.49%	-28.46%	87.72%	78.92%	78.06%
133 links/100Gbps/Hybrid	-14.90%	-23.68%	-29.93%	-19.54%	-31.26%	-38.15%	62.30%	54.68%	53.94%
133 links/40Gbps/Hybrid	3.59%	-7.10%	-14.70%	-2.05%	-16.33%	-24.71%	97.56%	88.30%	87.40%
133 links/10Gbps/Hybrid	13.46%	1.75%	-6.58%	7.28%	-8.36%	-17.54%	116.38%	106.23%	105.24%

	57 links	57 links	57 links	95 links	95 links	95 links	133 links	133 links	133 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	FRR	FRR	FRR	FRR	FRR	FRR	FRR	FRR	FRR
31 links/100Gbps/1:1	-39.87%	-43.51%	-44.50%	-54.08%	-57.50%	-58.62%	-59.22%	-63.13%	-64.93%
31 links/40Gbps/1:1	-38.92%	-42.63%	-43.63%	-53.36%	-56.83%	-57.97%	-58.58%	-62.55%	-64.38%
31 links/10Gbps/1:1	-39.26%	-42.95%	-43.94%	-53.62%	-57.07%	-58.20%	-58.81%	-62.76%	-64.57%
57 links/100Gbps/1:1	-16.22%	-21.30%	-22.68%	-36.03%	-40.79%	-42.34%	-43.19%	-48.63%	-51.14%
57 links/40Gbps/1:1	-15.39%	-20.52%	-21.91%	-35.39%	-40.20%	-41.77%	-42.63%	-48.12%	-50.65%
57 links/10Gbps/1:1	-14.23%	-19.43%	-20.84%	-34.51%	-39.38%	-40.97%	-41.84%	-47.41%	-49.98%
95 links/100Gbps/1:1	7.54%	1.02%	-0.75%	-17.88%	-24.00%	-25.99%	-27.08%	-34.05%	-37.28%
95 links/40Gbps/1:1	10.68%	3.96%	2.14%	-15.49%	-21.78%	-23.83%	-24.95%	-32.13%	-35.45%
95 links/10Gbps/1:1	8.91%	2.31%	0.52%	-16.84%	-23.03%	-25.05%	-26.15%	-33.21%	-36.48%
133 links/100Gbps/1:1	27.70%	19.95%	17.86%	-2.49%	-9.75%	-12.12%	-13.41%	-21.69%	-25.52%
133 links/40Gbps/1:1	30.77%	22.83%	20.69%	-0.15%	-7.58%	-10.01%	-11.33%	-19.81%	-23.73%
133 links/10Gbps/1:1	29.59%	21.73%	19.60%	-1.05%	-8.42%	-10.82%	-12.13%	-20.54%	-24.42%
31 links/100Gbps/Path	-34.06%	-38.06%	-39.14%	-49.65%	-53.40%	-54.62%	-55.29%	-59.57%	-61.54%
31 links/40Gbps/Path	-28.71%	-33.03%	-34.21%	-45.57%	-49.62%	-50.94%	-51.66%	-56.28%	-58.42%
31 links/10Gbps/Path	-27.27%	-31.68%	-32.87%	-44.46%	-48.60%	-49.95%	-50.68%	-55.40%	-57.58%
57 links/100Gbps/Path	7.15%	0.65%	-1.11%	-18.19%	-24.28%	-26.26%	-27.34%	-34.30%	-37.51%
57 links/40Gbps/Path	21.34%	13.98%	11.99%	-7.35%	-14.25%	-16.50%	-17.72%	-25.59%	-29.23%
57 links/10Gbps/Path	29.69%	21.82%	19.69%	-0.97%	-8.35%	-10.75%	-12.06%	-20.47%	-24.36%
95 links/100Gbps/Path	37.79%	29.43%	27.17%	5.21%	-2.62%	-5.18%	-6.57%	-15.51%	-19.64%
95 links/40Gbps/Path	53.64%	44.32%	41.79%	17.31%	8.58%	5.73%	4.18%	-5.79%	-10.39%
95 links/10Gbps/Path	67.34%	57.19%	54.44%	27.78%	18.27%	15.16%	13.47%	2.61%	-2.40%
133 links/100Gbps/Path	45.72%	36.88%	34.49%	11.27%	2.99%	0.28%	-1.19%	-10.64%	-15.01%
133 links/40Gbps/Path	70.58%	60.24%	57.44%	30.25%	20.56%	17.39%	15.67%	4.60%	-0.51%
133 links/10Gbps/Path	89.58%	78.08%	74.97%	44.76%	33.98%	30.47%	28.55%	16.25%	10.57%

	57 links	57 links	57 links	95 links	95 links	95 links	133 links	133 links	133 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	FRR	FRR	FRR	FRR	FRR	FRR	FRR	FRR	FRR
31 links/100Gbps/FRR	-27.75%	-32.14%	-33.32%	-44.83%	-48.94%	-50.28%	-51.01%	-55.70%	-57.86%
31 links/40Gbps/FRR	-24.20%	-28.80%	-30.04%	-42.12%	-46.43%	-47.83%	-48.60%	-53.52%	-55.79%
31 links/10Gbps/FRR	-23.83%	-28.45%	-29.70%	-41.84%	-46.17%	-47.58%	-48.35%	-53.29%	-55.58%
57 links/100Gbps/FRR		-6.07%	-7.71%	-23.64%	-29.33%	-31.18%	-32.19%	-38.68%	-41.68%
57 links/40Gbps/FRR	6.46%		-1.75%	-18.71%	-24.76%	-26.74%	-27.81%	-34.72%	-37.91%
57 links/10Gbps/FRR	8.35%	1.78%		-17.27%	-23.42%	-25.43%	-26.53%	-33.56%	-36.80%
95 links/100Gbps/FRR	30.96%	23.02%	20.87%		-7.44%	-9.87%	-11.19%	-19.69%	-23.62%
95 links/40Gbps/FRR	41.50%	32.91%	30.59%	8.04%		-2.62%	-4.05%	-13.23%	-17.47%
95 links/10Gbps/FRR	45.31%	36.49%	34.11%	10.95%	2.69%		-1.47%	-10.90%	-15.25%
133 links/100Gbps/FRR	47.47%	38.53%	36.10%	12.60%	4.22%	1.49%		-9.57%	-13.99%
133 links/40Gbps/FRR	63.08%	53.18%	50.51%	24.52%	15.25%	12.23%	10.58%		-4.89%
133 links/10Gbps/FRR	71.45%	61.05%	58.24%	30.92%	21.17%	17.99%	16.26%	5.14%	
31 links/100Gbps/Hybrid	-41.87%	-45.39%	-46.35%	-55.61%	-58.92%	-59.99%	-60.58%	-64.35%	-66.09%
31 links/40Gbps/Hybrid	-33.79%	-37.80%	-38.89%	-49.44%	-53.21%	-54.43%	-55.10%	-59.40%	-61.38%
31 links/10Gbps/Hybrid	-31.25%	-35.42%	-36.55%	-47.50%	-51.41%	-52.69%	-53.38%	-57.84%	-59.90%
57 links/100Gbps/Hybrid	-13.98%	-19.20%	-20.61%	-34.32%	-39.21%	-40.81%	-41.67%	-47.25%	-49.83%
57 links/40Gbps/Hybrid	-4.08%	-9.89%	-11.47%	-26.76%	-32.21%	-33.99%	-34.95%	-41.18%	-44.05%
57 links/10Gbps/Hybrid	7.69%	1.16%	-0.61%	-17.77%	-23.89%	-25.89%	-26.98%	-33.96%	-37.19%
95 links/100Gbps/Hybrid	6.24%	-0.20%	-1.95%	-18.88%	-24.91%	-26.88%	-27.96%	-34.85%	-38.03%
95 links/40Gbps/Hybrid	25.74%	18.11%	16.05%	-3.99%	-11.14%	-13.47%	-14.74%	-22.90%	-26.66%
95 links/10Gbps/Hybrid	35.62%	27.40%	25.17%	3.56%	-4.15%	-6.66%	-8.03%	-16.83%	-20.90%
133 links/100Gbps/Hybrid	17.25%	10.14%	8.22%	-10.47%	-17.13%	-19.31%	-20.49%	-28.10%	-31.61%
133 links/40Gbps/Hybrid	42.73%	34.08%	31.73%	8.99%	0.88%	-1.77%	-3.21%	-12.47%	-16.75%
133 links/10Gbps/Hybrid	56.33%	46.84%	44.28%	19.37%	10.48%	7.58%	6.01%	-4.14%	-8.82%

	31 links	31 links	31 links	57 links	57 links	57 links	95 links	95 links	95 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid
31 links/100Gbps/1:1	3.44%	-9.18%	-12.53%	-30.09%	-37.31%	-44.16%	-43.40%	-52.18%	-55.66%
31 links/40Gbps/1:1	5.07%	-7.75%	-11.16%	-28.99%	-36.33%	-43.28%	-42.51%	-51.42%	-54.97%
31 links/10Gbps/1:1	4.48%	-8.27%	-11.65%	-29.39%	-36.68%	-43.60%	-42.83%	-51.69%	-55.22%
57 links/100Gbps/1:1	44.12%	26.53%	21.86%	-2.60%	-12.66%	-22.20%	-21.15%	-33.37%	-38.23%
57 links/40Gbps/1:1	45.55%	27.79%	23.07%	-1.63%	-11.79%	-21.43%	-20.36%	-32.71%	-37.61%
57 links/10Gbps/1:1	47.54%	29.54%	24.76%	-0.29%	-10.59%	-20.35%	-19.27%	-31.79%	-36.76%
95 links/100Gbps/1:1	85.00%	62.42%	56.43%	25.03%	12.11%	-0.14%	1.22%	-14.47%	-20.71%
95 links/40Gbps/1:1	90.39%	67.15%	60.98%	28.67%	15.38%	2.77%	4.17%	-11.98%	-18.40%
95 links/10Gbps/1:1	87.36%	64.49%	58.42%	26.62%	13.54%	1.14%	2.51%	-13.38%	-19.69%
133 links/100Gbps/1:1	119.67%	92.86%	85.74%	48.46%	33.12%	18.58%	20.19%	1.56%	-5.84%
133 links/40Gbps/1:1	124.95%	97.50%	90.21%	52.03%	36.32%	21.43%	23.08%	4.00%	-3.58%
133 links/10Gbps/1:1	122.92%	95.71%	88.49%	50.65%	35.09%	20.33%	21.97%	3.06%	-4.45%
31 links/100Gbps/Path	13.43%	-0.41%	-4.09%	-23.34%	-31.26%	-38.77%	-37.94%	-47.56%	-51.38%
31 links/40Gbps/Path	22.63%	7.67%	3.69%	-17.12%	-25.68%	-33.80%	-32.90%	-43.30%	-47.44%
31 links/10Gbps/Path	25.12%	9.85%	5.79%	-15.44%	-24.18%	-32.46%	-31.54%	-42.15%	-46.37%
57 links/100Gbps/Path	84.32%	61.83%	55.85%	24.57%	11.70%	-0.50%	0.85%	-14.78%	-21.00%
57 links/40Gbps/Path	108.73%	83.26%	76.49%	41.07%	26.49%	12.68%	14.21%	-3.50%	-10.53%
57 links/10Gbps/Path	123.09%	95.87%	88.64%	50.77%	35.20%	20.43%	22.07%	3.14%	-4.38%
95 links/100Gbps/Path	137.02%	108.10%	100.42%	60.19%	43.64%	27.95%	29.69%	9.58%	1.59%
95 links/40Gbps/Path	164.29%	132.04%	123.47%	78.61%	60.16%	42.67%	44.61%	22.19%	13.28%
95 links/10Gbps/Path	187.86%	152.73%	143.40%	94.55%	74.45%	55.39%	57.51%	33.09%	23.38%
133 links/100Gbps/Path	150.68%	120.09%	111.96%	69.41%	51.91%	35.32%	37.16%	15.89%	7.45%
133 links/40Gbps/Path	193.44%	157.63%	148.12%	98.32%	77.83%	58.41%	60.56%	35.67%	25.78%
133 links/10Gbps/Path	226.12%	186.32%	175.75%	120.40%	97.63%	76.04%	78.44%	50.77%	39.78%

	31 links	31 links	31 links	57 links	57 links	57 links	95 links	95 links	95 links
	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps	100Gbps	40Gbps	10Gbps
	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid	Hybrid
31 links/100Gbps/FRR	24.28%	9.12%	5.09%	-16.01%	-24.68%	-32.91%	-32.00%	-42.54%	-46.73%
31 links/40Gbps/FRR	30.40%	14.49%	10.26%	-11.87%	-20.98%	-29.61%	-28.65%	-39.71%	-44.11%
31 links/10Gbps/FRR	31.02%	15.04%	10.79%	-11.45%	-20.60%	-29.27%	-28.31%	-39.42%	-43.84%
57 links/100Gbps/FRR	72.02%	51.03%	45.45%	16.26%	4.25%	-7.14%	-5.88%	-20.47%	-26.27%
57 links/40Gbps/FRR	83.13%	60.78%	54.85%	23.77%	10.98%	-1.14%	0.20%	-15.33%	-21.51%
57 links/10Gbps/FRR	86.39%	63.64%	57.60%	25.97%	12.96%	0.62%	1.98%	-13.83%	-20.11%
95 links/100Gbps/FRR	125.29%	97.80%	90.49%	52.26%	36.53%	21.61%	23.27%	4.16%	-3.44%
95 links/40Gbps/FRR	143.41%	113.70%	105.81%	64.50%	47.51%	31.39%	33.18%	12.53%	4.33%
95 links/10Gbps/FRR	149.96%	119.46%	111.36%	68.93%	51.48%	34.93%	36.77%	15.57%	7.14%
133 links/100Gbps/FRR	153.68%	122.73%	114.50%	71.45%	53.74%	36.94%	38.80%	17.29%	8.73%
133 links/40Gbps/FRR	180.53%	146.29%	137.20%	89.59%	70.01%	51.43%	53.49%	29.70%	20.24%
133 links/10Gbps/FRR	194.94%	158.95%	149.39%	99.33%	78.74%	59.21%	61.38%	36.36%	26.42%
31 links/100Gbps/Hybrid		-12.20%	-15.44%	-32.42%	-39.40%	-46.02%	-45.28%	-53.77%	-57.14%
31 links/40Gbps/Hybrid	13.90%		-3.69%	-23.02%	-30.97%	-38.52%	-37.68%	-47.34%	-51.18%
31 links/10Gbps/Hybrid	18.27%	3.83%		-20.07%	-28.33%	-36.16%	-35.29%	-45.32%	-49.31%
57 links/100Gbps/Hybrid	47.97%	29.91%	25.11%		-10.33%	-20.13%	-19.04%	-31.59%	-36.58%
57 links/40Gbps/Hybrid	65.01%	44.87%	39.53%	11.52%		-10.92%	-9.71%	-23.71%	-29.27%
57 links/10Gbps/Hybrid	85.25%	62.64%	56.64%	25.20%	12.26%		1.36%	-14.35%	-20.60%
95 links/100Gbps/Hybrid	82.76%	60.46%	54.54%	23.52%	10.76%	-1.34%		-15.50%	-21.66%
95 links/40Gbps/Hybrid	116.30%	89.90%	82.89%	46.18%	31.08%	16.76%	18.35%		-7.29%
95 links/10Gbps/Hybrid	133.31%	104.83%	97.27%	57.68%	41.39%	25.94%	27.66%	7.86%	
133 links/100Gbps/Hybrid	101.70%	77.09%	70.55%	36.32%	22.24%	8.88%	10.36%	-6.75%	-13.55%
133 links/40Gbps/Hybrid	145.54%	115.57%	107.61%	65.94%	48.80%	32.54%	34.35%	13.52%	5.24%
133 links/10Gbps/Hybrid	168.92%	136.10%	127.39%	81.74%	62.97%	45.17%	47.14%	24.33%	15.26%

	133 links 100Gbps Hybrid	133 links 40Gbps Hybrid	133 links 10Gbps Hybrid
31 links/100Gbps/1:1	-48.72%	-57.87%	-61.53%
31 links/40Gbps/1:1	-47.91%	-57.21%	-60.93%
31 links/10Gbps/1:1	-48.20%	-57.45%	-61.15%
57 links/100Gbps/1:1	-28.55%	-41.31%	-46.41%
57 links/40Gbps/1:1	-27.84%	-40.72%	-45.88%
57 links/10Gbps/1:1	-26.85%	-39.91%	-45.13%
95 links/100Gbps/1:1	-8.28%	-24.66%	-31.21%
95 links/40Gbps/1:1	-5.61%	-22.46%	-29.20%
95 links/10Gbps/1:1	-7.11%	-23.69%	-30.33%
133 links/100Gbps/1:1	8.91%	-10.53%	-18.31%
133 links/40Gbps/1:1	11.52%	-8.38%	-16.35%
133 links/10Gbps/1:1	10.52%	-9.21%	-17.11%
31 links/100Gbps/Path	-43.76%	-53.80%	-57.82%
31 links/40Gbps/Path	-39.20%	-50.05%	-54.40%
31 links/10Gbps/Path	-37.97%	-49.04%	-53.47%
57 links/100Gbps/Path	-8.62%	-24.93%	-31.46%
57 links/40Gbps/Path	3.48%	-14.99%	-22.38%
57 links/10Gbps/Path	10.60%	-9.14%	-17.04%
95 links/100Gbps/Path	17.51%	-3.47%	-11.86%
95 links/40Gbps/Path	31.03%	7.64%	-1.72%
95 links/10Gbps/Path	42.72%	17.24%	7.04%
133 links/100Gbps/Path	24.28%	2.09%	-6.78%
133 links/40Gbps/Path	45.48%	19.51%	9.12%
133 links/10Gbps/Path	61.68%	32.82%	21.27%

	133 links 100Gbps Hybrid	133 links 40Gbps Hybrid	133 links 10Gbps Hybrid
31 links/100Gbps/FRR	-38.38%	-49.38%	-53.78%
31 links/40Gbps/FRR	-35.35%	-46.89%	-51.51%
31 links/10Gbps/FRR	-35.04%	-46.64%	-51.28%
57 links/100Gbps/FRR	-14.71%	-29.94%	-36.03%
57 links/40Gbps/FRR	-9.21%	-25.42%	-31.90%
57 links/10Gbps/FRR	-7.59%	-24.09%	-30.69%
95 links/100Gbps/FRR	11.69%	-8.25%	-16.22%
95 links/40Gbps/FRR	20.68%	-0.87%	-9.49%
95 links/10Gbps/FRR	23.93%	1.80%	-7.05%
133 links/100Gbps/FRR	25.77%	3.32%	-5.67%
133 links/40Gbps/FRR	39.08%	14.25%	4.32%
133 links/10Gbps/FRR	46.22%	20.12%	9.68%
31 links/100Gbps/Hybrid	-50.42%	-59.27%	-62.81%
31 links/40Gbps/Hybrid	-43.53%	-53.61%	-57.65%
31 links/10Gbps/Hybrid	-41.37%	-51.83%	-56.02%
57 links/100Gbps/Hybrid	-26.64%	-39.74%	-44.98%
57 links/40Gbps/Hybrid	-18.19%	-32.80%	-38.64%
57 links/10Gbps/Hybrid	-8.16%	-24.55%	-31.11%
95 links/100Gbps/Hybrid	-9.39%	-25.57%	-32.04%
95 links/40Gbps/Hybrid	7.23%	-11.91%	-19.57%
95 links/10Gbps/Hybrid	15.67%	-4.98%	-13.24%
133 links/100Gbps/Hybrid		-17.85%	-24.99%
133 links/40Gbps/Hybrid	21.73%		-8.70%
133 links/10Gbps/Hybrid	33.32%	9.52%	

Appendix A

Links/Demand unit/Recovery method	Cost
133 links/10Gbps/Path	$7,655,032
133 links/10Gbps/FRR	$8,464,222
133 links/40Gbps/Path	$8,507,353
95 links/10Gbps/Path	$8,672,277
133 links/40Gbps/FRR	$8,899,037
133 links/10Gbps/Hybrid	$9,283,204
95 links/40Gbps/Path	$9,445,821
133 links/100Gbps/FRR	$9,840,741
133 links/100Gbps/Path	$9,958,796
95 links/10Gbps/FRR	$9,987,149
133 links/40Gbps/Hybrid	$10,167,277
95 links/40Gbps/FRR	$10,256,269
95 links/100Gbps/Path	$10,532,407
95 links/10Gbps/Hybrid	$10,700,265
95 links/100Gbps/FRR	$11,081,098
95 links/40Gbps/1:1	$11,097,778
57 links/10Gbps/Path	$11,190,187
133 links/10Gbps/1:1	$11,198,889
57 links/100Gbps/1:1	$11,364,444
95 links/40Gbps/Hybrid	$11,541,735
57 links/40Gbps/Path	$11,960,090
133 links/100Gbps/Hybrid	$12,376,759
95 links/100Gbps/1:1	$13,112,444
133 links/40Gbps/1:1	$13,324,444

57 links/10Gbps/FRR	$13,393,650
57 links/10Gbps/Hybrid	$13,476,109
31 links/10Gbps/1:1	$13,494,444
57 links/100Gbps/Path	$13,544,167
57 links/40Gbps/FRR	$13,631,917
95 links/100Gbps/Hybrid	$13,659,438
57 links/100Gbps/FRR	$14,512,226
57 links/40Gbps/Hybrid	$15,128,888
57 links/100Gbps/Hybrid	$16,871,650
133 links/100Gbps/1:1	$16,920,000
57 links/10Gbps/1:1	$17,151,667
31 links/40Gbps/1:1	$17,322,222
31 links/10Gbps/FRR	$19,053,079
31 links/40Gbps/FRR	$19,144,630
31 links/10Gbps/Path	$19,952,750
31 links/100Gbps/FRR	$20,086,891
31 links/40Gbps/Path	$20,356,667
31 links/10Gbps/Hybrid	$21,108,755
31 links/40Gbps/Hybrid	$21,917,864
31 links/100Gbps/Path	$22,008,333
57 links/40Gbps/1:1	$23,760,000
95 links/10Gbps/1:1	$23,893,333
31 links/100Gbps/1:1	$24,133,333
31 links/100Gbps/Hybrid	$24,964,279