



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**A PERFORMANCE ANALYSIS OF BGP/MPLS VPN
FAILOVER FUNCTIONALITY**

by

Guan Chye Tan

December 2006

Thesis Advisor:
Second Reader:

Geoffrey Xie
John Gibson

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 2006	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE: A Performance Analysis of BGP MPLS VPN Failover Functionality			5. FUNDING NUMBERS	
6. AUTHOR(S) Guan Chye Tan				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) Future military systems, many of which have unique timing requirements, will rely on the Global Information Grid (GIG) as the core data communication infrastructure. The GIG currently uses the BGP/MPLS VPN technology to provide secure and robust IP-level connectivity. This technology supports the provisioning of IP connectivity by a service provider to multiple customers over a common physical IP backbone while allowing complete logical separation of customer traffic and routing information. This research focuses on evaluating and validating the performance characteristic of BGP/MPLS VPN to determine if the use of this technology can provide the necessary performance guarantees required by military applications. A set of experiments have been performed to identify the key factors that affect the time delay of a network failure and recovery. The results show that reducing the ISIS SPF interval and Hello interval could shorten the failover latency while decreasing the ISIS SPF interval and TDP Hello interval could reduce the restoration delay, hence improving the BGP/MPLS VPN failover functionality.				
14. SUBJECT TERMS Border Gateway Protocol (BGP), Multi-Protocol Label Switching (MPLS), Virtual Private Network (VPN), Failover Functionality			15. NUMBER OF PAGES 111	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**A PERFORMANCE ANALYSIS OF
BGP/MPLS VPN FAILOVER FUNCTIONALITY**

Guan Chye Tan
Civilian, Ministry of Defense, Singapore
B.S., National University of Singapore, 1999

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
December 2006**

Author: Guan Chye Tan

Approved by: Geoffrey Xie
Thesis Advisor

John Gibson
Second Reader

Peter J. Denning
Chairman, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Future military systems, many of which have unique timing requirements, will rely on the Global Information Grid (GIG) as the core data communication infrastructure. The GIG currently uses the BGP/MPLS VPN technology to provide secure and robust IP-level connectivity. This technology supports the provisioning of IP connectivity by a service provider to multiple customers over a common physical IP backbone while allowing complete logical separation of customer traffic and routing information.

This research focuses on evaluating and validating the performance characteristic of BGP/MPLS VPN to determine if the use of this technology can provide the necessary performance guarantees required by military applications. A set of experiments have been performed to identify the key factors that affect the time delay of a network failure and recovery. The results show that reducing the ISIS SPF interval and Hello interval could shorten the failover latency while decreasing the ISIS SPF interval and TDP Hello interval could reduce the restoration delay, hence improving the BGP/MPLS VPN failover functionality.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	MOTIVATION.....	1
B.	THESIS OBJECTIVE	2
C.	ORGANIZATION OF THESIS.....	3
II.	BACKGROUND.....	5
A.	CHAPTER OVERVIEW.....	5
B.	OVERVIEW OF MPLS	5
1.	What is MPLS?	5
2.	MPLS vs IP	6
3.	Elements of MPLS	6
4.	MPLS Forwarding	7
5.	Advantages of MPLS	8
C.	MPLS LABEL DISTRIBUTION PROTOCOL.....	9
1.	What is Label Distribution?	9
2.	Label Distribution Methods.....	9
3.	Label Retention.....	11
4.	Label Distribution Control.....	11
D.	BGP/MPLS VIRTUAL PRIVATE NETWORK	12
1.	BGP/MPLS VPN Overview.....	12
2.	BGP/MPLS VPN Features.....	13
3.	How BGP/MPLS VPN Operates?	15
E.	MPLS TRAFFIC ENGINEERING	16
F.	MPLS PROTECTION MECHANISM	17
1.	Types of Protection Mechanisms.....	17
2.	MPLS Protection Mechanisms Options	19
G.	PRIOR RELATED WORK ON MPLS PROTECTION	21
1.	Efficiency of Routing and Resilience Mechanisms	22
2.	A Fast and Scalable Inter-Domain MPLS Protection Mechanism	23
H.	SUMMARY	25
III.	LAB SETUP AND PROCEDURES.....	27
A.	CHAPTER OVERVIEW.....	27
B.	PERFORMANCE METRICS	27
1.	Failover Time.....	27
2.	Restoration Time	28
C.	SCOPE OF EXPERIMENT.....	29
D.	LABORATORY SET-UP	30
E.	BASIC ROUTER CONFIGURATIONS.....	32
1.	Installing PE Router.....	32
2.	Installing P Router	35
3.	Installing CE Router	37

F.	EXPERIMENTAL TESTING AND CONFIGURATIONS.....	37
1.	Varying ISIS Metric Value.....	38
2.	Varying ISIS SPF Intervals	41
3.	Varying ISIS Hello Intervals	43
4.	Varying TDP Discovery Hello Intervals.....	46
5.	Varying MPLS TE Tunnel Configuration Options	49
6.	Varying Static and Non-static Routing Configuration	52
G.	TOOLS AND PROCEDURES FOR DATA COLLECTION PROCESS.....	54
1.	Traffic Generation.....	54
2.	Failure and Recovery Simulation	55
3.	Packet Capturing	57
3.	Traffic Monitoring	60
H.	SUMMARY	62
IV.	EXPERIMENTAL RESULTS AND ANALYSIS.....	63
A.	CHAPTER OVERVIEW.....	63
B.	RESULTS AND ANALYSIS.....	63
1.	Varying ISIS Metric Value.....	63
2.	Varying ISIS SPF Intervals	67
3.	Varying ISIS Hello Intervals	70
4.	Varying TDP Discovery Hello Intervals.....	72
5.	Varying MPLS TE Tunnel Configuration Options	75
6.	Varying Static and Non-Static Routing Configuration.....	79
C.	SUMMARY	82
V.	CONCLUSION AND FURTHER RESEARCH AREAS.....	85
A.	CHAPTER OVERVIEW.....	85
B.	CONCLUSION	85
C.	FURTHER RESEARCH AREAS.....	88
1.	Examining additional parameters of BGP/MPLS VPN	89
2.	Expanding the size of the laboratory network set up.....	89
3.	Examining the MPLS Fast Reroute	89
4.	Examining the prioritization of multiple VPNs	90
	LIST OF REFERENCES.....	91
	INITIAL DISTRIBUTION LIST	95

LIST OF FIGURES

Figure 1.	MPLS Label (From Ref. [7].).....	7
Figure 2.	Routing via Label Switched Path (After Ref. [8].)	8
Figure 3.	Downstream Unsolicited Label Distribution (From Ref. [9].)	10
Figure 4.	Downstream-on-Demand Label Distribution (From Ref. [9].).....	10
Figure 5.	VPNs with a Service Provider Backbone (From Ref. [11].).....	12
Figure 6.	Provider Edge/Customer Edge Router Relationship (From Ref. [12].)	14
Figure 7.	Basic Topology For Dynamic Headend Reroute	20
Figure 8.	Basic Topology For Fast Reroute	21
Figure 9.	Laboratory Set-Up of BGP/MPLS VPN Network	31
Figure 10.	Screenshot of the GUI of the Bricks Program.....	55
Figure 11.	Modification to Network Topology for Test Case 14 - 18.	57
Figure 12.	Screenshot of the GUI of the Wireshark Program.	58
Figure 13.	Modification to Network Topology to Allow Data Capturing	59
Figure 14.	Screenshot of the GUI of the Bandwidth Gauges Feature in the SolarWinds Network Management Software.	60
Figure 15.	Screenshot of the GUI of the Kiwi Syslog Daemon	61
Figure 16.	Mean Failover Time for Test Cases 1 - 8	65
Figure 17.	Mean Restoration Time for Test Cases 1–3 & 5-7	67
Figure 18.	Mean Failover and Restoration Time with respect to SPF interval	69
Figure 19.	Mean Failover and Restoration Time with respect to Hello interval	72
Figure 20.	Mean Failover and Restoration Time with respect to TDP Discovery Hello Interval	74
Figure 21.	Sample Details in MPLS Forwarding Table	74
Figure 22.	Comparison of Failover and Restoration Time among different TE tunnel configuration options.....	76
Figure 23.	MPLS-TE system block diagram with tunnels introduced into IGP (After Ref. [7].).....	78
Figure 24.	Comparison of Failover and Restoration Time based on tunnels being established and different spf-initial-wait intervals.....	79
Figure 25.	Comparison of Mean Failover Time between Static and Non-Static Routing Configuration.....	81
Figure 26.	Comparison of Mean Restoration Time between Static and Non-Static Routing Configuration.....	82

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Basic Configuration of PE Router – PE1	34
Table 2.	Configuration of P Router – P1.....	36
Table 3.	Configuration of CE Router – CE1A	37
Table 4.	Router Command to Assign a Different Metric to an Interface	39
Table 5.	Sample Router Configuration for ISIS metric.....	40
Table 6.	Experimental Parameters for Test 1 - 8.....	40
Table 7.	Sample Router Command to Configure SPF Intervals	42
Table 8.	Sample Router Configuration for ISIS SPF Interval.....	42
Table 9.	Experimental Parameters for Test 9 – 13	43
Table 10.	Sample Router Command to Configure Hello Interval.....	44
Table 11.	Sample Router Configuration for ISIS Hello Interval.....	45
Table 12.	Experimental Parameters for Test 14 – 18	46
Table 13.	Sample Router Command to Configure TDP Discovery Intervals	47
Table 14.	Sample Router Configuration for TDP Hello Interval	48
Table 15.	Experimental Parameters for Test 19 - 24.....	49
Table 16.	Sample Router Command to Configure a MPLS TE Tunnel	50
Table 17.	Sample Router Configuration for MPLS Traffic Engineering	51
Table 18.	Experimental Parameters for Test 25 - 28.....	52
Table 19.	Sample Router Command to Configure a Static Route	52
Table 20.	Sample Router Configuration for Static Routing	53
Table 21.	Experimental Parameters for Test 29 - 32.....	54
Table 22.	Router Commands to Simulate a Link/Node Failure and Recovery ...	56
Table 23.	Statistical Results on Failover Time for Test Cases 1 – 8	63
Table 24.	Statistical Results on Restoration Time for Test Cases 1 - 8.....	64
Table 25.	Statistical Results on Failover Time for Test Cases 9 – 13	67
Table 26.	Statistical Results on Restoration Time for Test Cases 9 - 13.....	68
Table 27.	Statistical Results on Failover Time for Test Cases 14 – 18	70
Table 28.	Statistical Results on Restoration Time for Test Cases 14 - 18.....	71
Table 29.	Statistical Results on Failover Time for Test Cases 19 – 24	73
Table 30.	Statistical Results on Restoration Time for Test Case 19 – 24.....	73
Table 31.	Statistical Results on Failover Time for Test Cases 25 – 28	75
Table 32.	Statistical Results on Restoration Time for Test Cases 25 - 28.....	76
Table 33.	Statistical Results on Failover Time for Test Cases 29 – 32	80
Table 34.	Statistical Results on Restoration Time for Test Cases 29 - 32.....	80

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I would like to express my heartfelt thanks to the following people for their kind support and assistance that leads to the successful completion of this thesis.

Many thanks to my thesis advisor, Prof. Geoffrey Xie for his helpful advice, support and encouragement that kept me on track throughout my entire thesis work. I would also like to extend my gratitude to my second reader, Prof. John Gibson, for his time and effort to assist me.

Finally, I would like to thank my family and friends for their unfailing patience and support.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. MOTIVATION

The GIG [1] is a globally interconnected end-to-end set of information capabilities, associated processes, and personnel for collecting, processing, storing, disseminating, and managing information on demand to war fighters, policy makers, and support personnel. Future military systems will rely on GIG as the core data communication infrastructure. Many of these systems have unique timing requirements in order to accomplish their missions. It is important, therefore, that the infrastructure is able to support these systems in achieving their goals.

DoD uses the IP-based, Internet architecture as the GIG architecture to allow inter-connection of hundreds of individually administrated DoD networks[2]. Like the Internet backbone, the GIG network backbone - the Global Information Grid-Bandwidth Expansion (GIG-BE), employs commercial-off-the-shelf (COTS) switches and routers to facilitate information exchanges. It also adopts common Internet standards and protocols to allow seamless inter-operability among the GIG's systems. One of them is the BGP/MPLS Virtual Private Network (VPN) technology.

The GIG currently uses the BGP/MPLS VPN technology to provide secure and robust IP-level connectivity. BGP/MPLS VPN uses a variety of commercially-available switching and routing technologies. It supports the provision of IP connectivity by a service provider to multiple customers over a common physical IP backbone while allowing complete logical separation of customer traffic and routing information. BGP/MPLS VPN allows network traffic from different organizations and of different classifications to flow across the same backbone by logically separating them into different VPNs.

Military applications have unique requirements on fault tolerance and fast recovery. Real-time communication services provided by the network infrastructure are therefore essential to many of such applications. Ideally, their

performance should be minimally impacted by link or node failures as long as the network is still physically connected with the remaining components. Currently, BGP/MPLS VPN provides a failover service to route VPN traffic around failed links or routers using different types of protection mechanisms. It is the interest of this thesis research, therefore, to find out, in detail, the performance of this failover service, in terms of its impact on the message delivery latency. Also of interest are the contributing factors that affect the performance. In addition, this thesis research aims to establish whether the BGP/MPLS VPN can be configured to implement policy-based rerouting whereby certain traffic flows can be given priority, either statically or dynamically on the fly, to be rerouted ahead of other flows.

B. THESIS OBJECTIVE

The objective of this thesis is to study and analyze the performance of the BGP MPLS VPN that is used by GIG and most civilian Internet Service Providers' network backbone infrastructure. A BGP MPLS VPN network backbone using CISCO hardware and software will be set up under laboratory conditions. The configuration of the network backbone will be based on the recommendations gathered from the literature research. Network traffic will be generated across the network backbone using a network traffic generation software tool to simulate traffic flow generated by the military system applications. Performance statistics will be collected and a statistical analysis will then be performed.

In essence, the thesis aims to answer the following questions:

- (i) What is the fastest failover time achievable for a BGP MPLS VPN backbone in the event of a link or router failure?
- (ii) What are the determining factors for this delay?
- (iii) What are all the possible router configuration options that can be used to reduce the failover time?

C. ORGANIZATION OF THESIS

The next chapter provides an overview of BGP MPLS VPN technology and discusses prior related work on MPLS protection mechanisms. Chapter III describes the laboratory setup of the BGP MPLS VPN network backbone using CISCO equipment. Chapter IV presents the results and findings from the statistical analysis of experimentation performed on the lab setup of Chapter III. Finally, Chapter V concludes the research work and provides suggestions on further research areas.

THIS PAGE INTENTIONALLY LEFT BLANK

II. BACKGROUND

A. CHAPTER OVERVIEW

This chapter begins with an overview of the MPLS. It then continues to provide readers with insight about the various MPLS key components that are relevant to the study of the BGP/MPLS VPN failover functionality, namely the label distribution protocol, MPLS virtual private network, MPLS traffic engineering and MPLS protection mechanisms. The chapter ends with a discussion on some related work on the MPLS protection mechanism by reviewing two papers; the first paper provides some findings on the efficiency of the protection mechanisms in an intra-domain environment and the second paper presents a MPLS protection mechanism inter-domain solution.

B. OVERVIEW OF MPLS

1. What is MPLS?

In a Multiprotocol Label Switching (MPLS) network [3], incoming packets are assigned a "label" by a "label edge router (LER)." Packets are forwarded along a "label switch path (LSP)" in which each "label switch router (LSR)" makes forwarding decisions based solely on the contents of the label. At each hop, the LSR strips off the existing label and applies a new label, which tells the next hop how to forward the packet. Label Switch Paths (LSPs) are established by network operators for a variety of purposes, such as to guarantee a certain level of performance, to route around network congestion, or to create IP tunnels for network-based virtual private networks. In many ways, LSPs are no different than circuit-switched paths in ATM or Frame Relay networks, except that they are not dependent on a particular Layer 2 technology.

MPLS is based on the notion of label switching. The initial intention of using label switching is to increase the forwarding speed at Layer 3 to a level close to Layer 2. Label switching methods allow routers to forward packets based on the contents of a simple label, rather than by performing a complex route loop

up, based on IP destination address. Although the main benefit from the initial intention is no longer valid today due to better Layer 3 hardware, there are still other benefits that one can reap from using MPLS.

2. MPLS vs IP

MPLS has some circuit-switch properties over a packet-switch network. In conventional IP packet forwarding [4], an independent forwarding decision is made at each router based on the IP destination address in the packet's header. Routing protocols such as Intermediate System-To-Intermediate System (ISIS), Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) are used to assist in the hop-by-hop decisions. In MPLS, the forwarding decision is made once, that is, when the packet enters the network [5]. The packet is assigned to a label based on its IP destination at the entry of the network. When the packet is forwarded to the next hop, the label is sent along with it. There is no further analysis required to make the forwarding decision at subsequent hops until the packet leaves the MPLS network.

3. Elements of MPLS

A forwarding equivalence class (FEC) [5] is a group of IP packets which are forwarded in the same manner (e.g., over the same path, with the same forwarding treatment) by a router. It can be based on the network prefix, quality of service, and so on.

A MPLS label is defined as a short, fixed length, locally significant identifier which is used to identify a FEC. The label which is put on a particular packet represents the "Forwarding Equivalence Class" to which that packet is assigned. The format of the MPLS label differs depending on the mode that MPLS operates. In the frame-mode [6], the label is carried as a "shim" layer between the Layer 2 and Layer 3 headers. MPLS labels are 4 octets long and consist of a 20-bit label, a 3-bit Experimental (EXP) field, a bottom of label stack (S) bit, and an 8-bit Time-to-Live (TTL) field. This is illustrated in Figure 1.




Figure 1. MPLS Label (From Ref. [7].)

A label switch router (LSR) is a device in a MPLS network that performs MPLS control and forwarding components. It forwards a packet based on the value of a label encapsulated in the packet. A LSR, however, can also forward native Layer 3 packets. A label edge router (LER) is an ingress or egress LSR that resides at the exit points of the MPLS network. A label switched path (LSP) is a specific traffic path through an MPLS network. It is the path through one or more LSRs at one level of the hierarchy, followed by a packet in a particular FEC.

4. MPLS Forwarding

An ingress LSR will classify an IP packet into a FEC when it arrives at the entry point of the MPLS network [6]. The ingress LSR will then tag a label based on the FEC. Each label is unique to its router. The label, which serves as an identifier, will enable a LSR to forward the packet without having to do a lookup in the IP routing table. The label will be swapped at each hop along the LSP until it reaches the penultimate LSR. The penultimate LSR will either pop or remove the label before forwarding the IP packet to the egress LSR. If the label is removed at the penultimate LSR, then the egress LSR will simply do a lookup at the IP routing table and forward the packet accordingly, skipping the step of label

lookup. Figure 2 shows the routing of a packet via the LSP and the label switching table of LSR2.




Figure 2. Routing via Label Switched Path (After Ref. [8].)

5. Advantages of MPLS

There are many advantages to using MPLS [4]. It enables a single converged network to support both new and legacy services, allowing efficient migration to an IP-based infrastructure. MPLS operates over legacy infrastructures such as DS3 and SONET and new infrastructures (10/100/1000/10G Ethernet) and networks (IP, ATM, Frame Relay, Ethernet and TDM). The word “Multiprotocol” indicates that MPLS has the ability to carry multiple network protocols.

Another advantage of MPLS is that it does not require high degrees of router processing from the label-switch routers for the forwarding since the most intensive part of the process, which is the assignment decisions, has been made at the label edge routers. Less high-end routers and switches can be used to perform the forwarding instead.

In addition, MPLS also allows traffic engineering. It can force a packet to follow a certain route based on some decisions other than the IP destination. In conventional IP forwarding, this would require the use of some encoding at the packet header to indicate the route the packet wants to travel (source routing). This increases the packet size, thus resulting in additional network load. On the other hand, a MPLS label, which is relatively smaller, can be used to represent the route. MPLS traffic engineering will be discussed in detail in the third section of this chapter.

Since MPLS can isolate traffic within its network by means of LSPs, it can also make IP as secure as frame relay in the wide area network with the appropriate level of security, without the need for encryption over public IP networks. As such, many service providers use MPLS for provision of Virtual Private Network services. MPLS VPN will be discussed in details in the fourth section of this chapter.

C. MPLS LABEL DISTRIBUTION PROTOCOL

1. What is Label Distribution?

Label distribution ensures that adjacent routers have a standardized view of the FEC [10]. The LSRs will have a common understanding regarding to which FEC the label is referring. Label distribution can either ride on an existing routing protocol or use a dedicated label distribution protocol. A label distribution protocol [5] is a set of procedures by which one LSR informs another of the label/FEC bindings it has made. Some of these label distribution protocols are Label Distribution Protocol (LDP), Tag Distribution Protocol (TDP), RSVP, Multiprotocol Extensions for BGP-4(MP-BGP), and Protocol Independent Multicast (PIM), which can be employed in a MPLS network.

2. Label Distribution Methods

There are two methods for label distribution. They are downstream unsolicited label distribution and downstream-on-demand label distribution. For

the former method, LSRs do not have to wait for label bindings to be requested before advertising them to their upstream neighbors. In Figure 3, LSR2 is the downstream LSR of LSR1. They have a LDP adjacency. When LSR2 discovers a “next-hop” for a particular FEC, it generates a label for that FEC and communicates the bindings to LSR1. Upon receiving the label binding, LSR1 inserts it into its forwarding tables. If LSR2 is the next hop for the FEC, LSR1 can use that label, knowing that its meaning is understood.



Figure 3. Downstream Unsolicited Label Distribution (From Ref. [9].)

On the other hand, a downstream-on-demand label distribution will allow a LSR to request a label for a prefix from its downstream peer. In Figure 4, LSR1 recognizes LSR2 as its next-hop for a FEC. It then requests to LSR2 for a binding between the FEC and a label. If LSR2 recognizes the FEC and has a next-hop for it, it creates a binding and replies to LSR1. In this way, both LSRs will have a common understanding of which FEC the label generated is referring to.




Figure 4. Downstream-on-Demand Label Distribution (From Ref. [9].)

3. Label Retention

When a LSR receives all the bindings from its LSR peers, it decides whether to retain all the bindings or to discard some of them. In such cases, only those that correspond to best routes will be retained while the rest will be removed. There are two modes of label retention that a LSR can operate in: liberal label retention and conservative label retention. In the liberal label retention mode, a LSR will retain all the bindings received from its peers. In contrast, a LSR in a conservative retention mode will only retain those label bindings that correspond to the best route for a FEC.

More memory is needed when the liberal label retention is used. However, it takes shorter time to failover to an alternate path if the original LSP failed, as compared to the conservative retention mode.

4. Label Distribution Control

The communication between LSRs can take place at two planes in a MPLS network. The control plane is where the exchange of routing information and label bindings occurs. The receiving and sending of labeled packets are carried out at the data plane.

There are two types of LSP controls: Independent LSP Control and Ordered LSP Control. The independent LSP control allows LSRs to assign labels to prefixes independently. Labels are assigned regardless of whether other LSRs have assigned labels. However, in an ordered LSP control, a label-FEC binding is communicated if the LSR is the egress LSR to a particular FEC. The formation of the LSP flows from the egress to the ingress.

There are advantages to adopting the independent LSP control. Labels are exchanged in shorter time. There will be more delay in the packet forwarding in the ordered LSP controls. In addition, the independent LSP control does not depend on the availability of the egress LSR. However, the ordered LSP control ensures consistent granularity and freedom from loops. The ordered LSP control is mostly used in explicit routing and multicast.

D. BGP/MPLS VIRTUAL PRIVATE NETWORK

1. BGP/MPLS VPN Overview

BGP/MPLS VPN is fast becoming a popular choice for many service providers to provide IP-based VPN services to customers. It supports the provision of IP connectivity by a service provider to multiple customers over a common physical IP backbone, while allowing complete logical separation of customer traffic and routing information [6]. Interconnection of different sites belonging to the same customer is provided over the MPLS backbone. Figure 5 shows an example of a VPN with a service provider (P) backbone network, service provider edge routers (PE), and customer edge routers (CE). In this instance, a customer device attaching to the CE router at VPN 1 Site 1 is able to communicate with a customer device attaching to the CE router at VPN 1 Site 2.




Figure 5. VPNs with a Service Provider Backbone (From Ref. [11].)

The key network components of the BGP/MPLS VPN are the provider edge (PE) routers, the provider (P) routers and the customer edge (CE) routers.

The PE routers are routers within the service provider backbone that connect to customer sites. In a MPLS network, a PE router also performs as an edge LSR. The P routers are routers within the service provider backbone that do not connect directly to customer sites. They are the LSRs in a MPLS network. The CE routers are routers at the customer sites that are directly connected to the service provider network. They connect directly to the PE routers.

2. BGP/MPLS VPN Features

In order to establish confidentiality over a shared network infrastructure, the most common method that many current VPN solutions adopt is the use of encrypted tunnels [11]. These connection-oriented and point-to-point tunnels are established in a packet-based, connectionless network backbone to provide the VPN services. However, this characteristic limits the ability to leverage the benefits that packet-based, connectionless network architecture like the Internet offers, such as ease of connectivity and multiple services. MPLS VPN, conversely, is connection-less. It does not require tunnels and encryption to provide confidentiality. It can provide the same level of security that connection-oriented VPNs offer and yet reduce the complexity required to implement the VPN service as a result of using tunnels and encryption.

MPLS Layer 3 VPN, in particular, adopts the peer model in which routing information is exchanged between the customer and the service provider. However, each customer's routing information is maintained in separate forwarding tables known as the virtual routing and forwarding tables (VRF). Figure 6 shows the VRF for each VPN. Packets are then uniquely identified with the associated labels and LSPs for each VRF. As such, the traffic for each VPN is kept separated from the rest. Devices in one VPN are unable to access any device from another VPN, unless due to misconfiguration or deliberate configuration for inter-connection between them.




Figure 6. Provider Edge/Customer Edge Router Relationship (From Ref. [12].)

One of the main issues with connection-oriented VPN is scalability. VPNs must scale to support hundreds of thousands of sites. With connection-oriented, point-to-point implementation, this is not optimal, especially when one customer site has to connect to all other sites. With the peer model, MPLS VPN only requires the customer site to connect to one provider edge router in order to establish connectivity to the rest of the customer sites within the same VPN. In addition, the VRFs are maintained by the PE routers. The P routers only maintain the routes to the PE routers. Hence, the scalability of the provider's core increases and the support for increasing number of VPNs is not constrained by any device within the provider's network.

3. How BGP/MPLS VPN Operates?

Each PE router has a default forwarding table and many VRFs, depending on the number of customers' sites connecting to it [13]. The default forwarding table contains "public" routes and each VRF contains its very own "private" routes. Hence, non-communication between VPN sites and non-VPN sites are ensured by leaving the "private" routes out of the default forwarding table. The VRF differs among routers from different network vendors. For a Cisco router [12], each VRF includes an IP routing table, a derived Cisco forwarding (CEF) table, a set of interfaces that use the forwarding table, and a set of rules and routing parameters that control the information that is included in the routing table. Each customer site can only be associated with one VRF even if the site is a member of multiple VPNs.

A PE router can learn an IP prefix from a CE router either through static routing configuration, BGP or some other IGP protocols like OSPF and RIP. The IP prefix is based on IPv4 address family. The PE router will convert the IP prefix it has learned from the CE router into a VPN-IPv4 value by combining it with a 8-byte route distinguisher (RD). The new prefix belongs to the VPN-IPv4 address family. It uniquely identifies the customer address, even if the customer site is using globally non-unique IP addresses like the ones in the private address spaces.

The new VPN routing information is then injected into BGP. The external gateway protocol is used to propagate the new VPN routing information to the rest of the PE routers in the MPLS network backbone that have a need to know. The distribution is performed using the BGP multiprotocol extensions that support address spaces other than IPv4. The process of converting the newly learned IP prefix into VPN-IPv4 values and using BGP multiprotocol extensions allows routes for a particular VPN to be learned only by other members of the VPN, ensuring that only members of the same VPB can communicate with one another.

E. MPLS TRAFFIC ENGINEERING

Traffic engineering (TE) [14] is the technique or process of steering traffic across to the backbone to facilitate efficient use of available bandwidth between pairs of routers. It aims to balance the traffic load on the various links and routers in the network. TE can be best utilized in a network where multiple parallel or alternate paths are available.

TE is essential to service providers as it enables the service providers to maximize the utilization of network resources, as well as enhance the quality-of-service they offer [15]. In a large network, the available network bandwidth may not be efficiently utilized due to the routes computed by the interior gateway protocols (IGP), such as OSPF and IS-IS. On the one hand, the “optimal paths” computed based on least cost metrics may not have the sufficient resources to carry all the traffic through the backbone. Traffic congestion at some chokepoints may occur as a result. On the other hand, the suboptimal paths are under-utilized. Hence, the use of traffic engineering can help to steer some of the traffic destined to follow the optimal path to a sub-optimal path, in order to enable better bandwidth management and utilization.

MPLS TE, like other traffic engineering techniques, allows traffic to flow through a path that is different from the IGP destination-based hop-by-hop routing. MPLS TE uses extended Resource Reservation Protocol (RSVP) to signal and establish TE tunnels. The path used by the TE tunnel can be explicitly configured or can be based on the path defined by the IGP in the core. In the latter case, the tunnel is not tied to any specific path through the backbone. The TE tunnel can reroute packets via any available path through the network in the event of a link or router failure. The tunnel path is pre-established at tunnel setup time. Based on the bandwidth requirements, class of service for the data traffic, or administrative policies, reservations for the TE tunnels are performed using RSVP. The ingress LSR calculates and establishes the tunnel, depending on the requirements and the available resources. Traffic using the tunnel is then forwarded along the defined path through the network using MPLS. Unlike typical RSVP for Quality-of-Service (QoS), the admission control for the reservation is

done once, during the setting up of the tunnel, and not at the packet forwarding time. The control of bandwidth reservation is also carried out during this time.

Other than path controlling, MPLS TE also provides a means of resilience for the network backbone. A primary and backup path can be configured for a TE tunnel. In the event the primary path fails, the backup path can be used.

F. MPLS PROTECTION MECHANISM

One of the primary concerns of a network service provider is the service availability to the customers. In order to ensure that the disruption of service is minimized, resilience mechanisms are deployed in the service provider's network. These mechanisms will ensure that network service to the customers will continue in the event of a network failure.

1. Types of Protection Mechanisms

In general, the protection mechanisms can be classified into three categories, namely the restoration mechanisms, end-to-end protection switching mechanisms, and the Fast Reroute [16].

A restoration mechanism has no pre-established backup paths. A typical example of such a mechanism is the IP rerouting. A new route is established only after a network failure occurs. In IP routing, packets are forwarded on a per-hop basis. Routing tables are constructed usually using some interior gateway routing protocol, such as OSPF or IS-IS. Administrative link cost is assigned to each link in the network and the path to the destination is calculated based on least cost. The routing can be classified into single path routing and multi-path routing. The commonly known routing for each type is the shortest path routing (SSP) and the equal-cost multi-path (ECMP) routing, respectively. SSP uses the same least cost link throughout to forward the traffic to the destination. ECMP splits the traffic equally among the next hop for each equal cost path.

The key strength of IP rerouting is its robustness against network failures. It is capable of surviving multiple network failures so long as the network is

physically connected. However, the key weakness of this protection mechanism is its slow failure recovery. The failure recovery latency is typically higher than the other types of protection mechanisms. It has been widely reported that the failure recovery latencies range from seconds to minutes for Layer 3 routing protocols [17]. Real time applications, such as military applications, often require additional QoS criteria such as delay, delay variations, packet drop rate, and, most importantly, fault tolerance and fast recovery. As such, this protection mechanism is deemed as too slow to protect traffic of real time services.

The other type of protection mechanism is the end-to-end protection switching mechanism. This mechanism is based on the notion of primary and backup paths. Primary paths and disjoint backup paths are pre-established for the connection set up. The traffic will always travel along the primary path except in the event of a network failure. In this case, the head-end router of the connection will switch the traffic from the primary path to the backup path. "Hello" messages are sent by each node along the connection at regular intervals to assess the status of the connection. When the head-end router does not receive a certain number of hello messages within the stipulated time period, it will deem that the primary path is down and then switch the traffic to the next identified backup path.

The end-to-end protection mechanism supports real time applications better than the restoration mechanism. Due to the pre-establishment of backup paths, this mechanism is able to switch the traffic faster than the latter. Thus, the failure recovery latency is shorter. However, in the case where the primary and the backup path both fail, the network connectivity will be lost. This can be a significant detriment in terms of service availability. The workaround for such a situation is to construct more backup paths. However, this requires additional network resources and additional link management.

The third type of protection mechanism, Fast Reroute, is a special type of protection switching mechanism. The end-to-end protection mechanism makes use of hello messages to detect path failure. As the number of nodes increase

along the path and the point of failure occurs further away from the head-end router, the failure detection time increases. The fast reroute mechanism is able to detect failure at its location and redirect the traffic from there. This greatly reduces the failure detection time and reaction time.

2. MPLS Protection Mechanisms Options

The above three types of protection mechanisms can be deployed in the MPLS-enabled network[19]. In a MPLS-based network where traffic engineering is not utilized, IP rerouting is the default protection mechanism. The IP reroute relies on the underlying Layer 3 routing protocol to establish a new route after a network failure occurs. Depending on the label retention mode for which the MPLS network is configured, a new LSP will be established (if it is in conservative mode) after the network failure occurs or the head-end router (if it is in liberal mode) will use the next available LSP in the MPLS forwarding table to forward the traffic.

The head-end reroute is a member of the IP reroute protection mechanism family that uses MPLS TE. It is also the default protection mechanism for MPLS TE. Head-end reroute establishes a backup LSP that is dynamically signaled after a network failure occurs. One advantage of this option is that the backup LSP will not consume any network resources until it is utilized as a result of the failure. However, it incurs a long failover time. The packet loss during failure can be higher than that of an IGP convergence. Figure 7 shows the basic topology for the dynamic head-end reroute.




Figure 7. Basic Topology For Dynamic Headend Reroute

The standby LSP is the option that uses end-to-end protection mechanisms. It uses RSVP-TE signaling to signal a backup path, in advance, from the ingress to the egress nodes. The failure recovery latency is shorter than the head-end reroute. Nevertheless, it requires high resource utilization at the ingress to maintain unused backup LSPs. The topology is the same as the one shown in Figure 7. The only difference, in this case, is that a backup path is pre-established during the tunnel setup.

The MPLS Fast Reroute is a type of a fast reroute protection mechanism which provides very fast failure recovery capability. The MPLS Fast Reroute uses a “local repair principle” [16] that allows traffic to be rerouted at any of the “points-of-local repair” (PLR) along the path, instead of only rerouting at the head-end router. This is done by pre-establishing bypass tunnels for any LSR that is a potential point of failure along the LSP. It has been reported that the MPLS Fast Reroute can switch traffic on a failed link to a recovery path within 20ms. However, that response is limited to the global label assignment case [17]. The failover time is certainly much faster than what the other two protection mechanisms can offer, though. In addition, it helps to reduce the processing load

of the head-end router. It should be noted that this protection mechanism requires greater network configuration and increases signaling complexity.

Like the standby LSP, the MPLS Fast Reroute uses RSVP-TE to signal the backup LSPs. There are two backup options: namely, the one-to-one backup, which is also known as the detour mode, and the facility backup, which is the bypass mode. The former provides a separate backup path for each PLR of every path, while the latter provides protection switching for every network element instead. Figure 8 shows the basic topology for a Fast Reroute.




Figure 8. Basic Topology For Fast Reroute

G. PRIOR RELATED WORK ON MPLS PROTECTION

The wide acceptance of MPLS among the network service providers has provided motivation for research and academic communities to conduct further research and development with respect to this technology. One of the research areas is the MPLS protection mechanism. Detailed studies of this protection mechanism were conducted to examine the key strengths and weaknesses of this resilience mechanism and to explore solutions to optimize its potential. In this section, we will examine two papers that were published related to MPLS protection mechanisms. The first paper, entitled “Efficiency of Routing and

Resilience Mechanisms” [16], was published by a group of academic researchers from University of Wurzburg. The paper provides a comprehensive study regarding the efficiency of standard and improved routing and resilience mechanisms. The second paper, entitled “A Fast and Scalable Inter-Domain MPLS Protection Mechanism” [17], proposed a solution for the inter-domain MPLS recovery problem, which is a result of multiple independent domain administrations. The first paper looks at the different resilience mechanisms designed to address failure recovery problems in the intra-domain context, while the second paper addresses the issue from the inter-domain perspective.

1. Efficiency of Routing and Resilience Mechanisms

In the first paper, the primary focus is the link utilization in the network backbone as a result of employing various types of protection mechanisms to provide service continuity to the customers [16]. Often these protection mechanisms merely preserve connectivity by switching traffic to backup paths when the primary path fails. As a result, overload may occur at certain parts of the network due to the traffic redirection. The authors investigated how well the different protection mechanisms are able to redirect traffic in a way that the link utilization is maintained at the minimum in both failure-free and protected failure scenarios.

The different types of protection mechanisms discussed in the earlier section were examined. The authors highlighted the different ways to optimize these protection mechanisms to carry more protected traffic. They used some heuristic algorithms to optimize the various protection mechanisms in order to increase the spreading of backup traffic and decrease the required backup capacity. Then they carried out an experiment to compare the efficiency of the resilience mechanisms in different network topologies and with different resilience requirements. The key performance indicator for each resilience mechanism was the minimization of the maximum link utilization under the different stated conditions. In the experiment, different sizes of networks, ranging from 10 to 50 nodes and different average node degrees, from 3 – 6, were

simulated to determine if the size and the mesh level (highly or sparsely meshed) of the network play a part in the efficiency of the resilience mechanisms. They also compared the resilience mechanisms in different protection scenarios to determine if the protection variation has any impact on the efficiency of each resilience mechanism.

The experimental results showed that the size of the network and the average network node degrees do play a part in the efficiency of the resilience mechanisms. In addition, the difference in performance of each resilience mechanism is quite significant under different protected scenarios. Nevertheless, the focus of this paper, on the efficiency of the resilience mechanism, is very much on the link utilization in the network backbone, which is looking from the perspective of a network service provider. Alternatively, the use of packet loss as a performance metric to compare the efficiency of the resilience mechanism may provide another dimension for examination, since packet loss is an externally observable event and has direct impact on a customer's application performance.

2. A Fast and Scalable Inter-Domain MPLS Protection Mechanism

The second paper focuses on the failure recovery issue in the non-homogenous, independent inter-domain context [17]. Several MPLS protection mechanisms have been proposed over the years to address the issue of fault tolerance and fast recovery as a result of the slow layer three protection and recovery mechanisms. However, these mechanisms are designed for intra-domain recoveries. They do not address failure recovery in the inter-domain environment, especially when the domains are under different administrations.

The authors first discussed the different types of MPLS protection mechanisms, as well as their key strengths and weaknesses in the intra-domain environment. Some of the key strengths include fast recovery times and scalability, while the key weaknesses include inefficient use of bandwidth and long failover time. They then explained the ineffectiveness of the MPLS protection mechanisms when the LSP spanned across multiple domains that are not under a single administration. The difficulty arises from the unwillingness to

share information among the service providers. When a network failure occurs in an independent domain, the service provider of that domain will likely hide the failure information in fear of negative image and exploitation of such information by competitors. The service provider will contain this information and try to recover the failure by itself. However, the MPLS end-to-end protection mechanisms would require some kind of failure signaling to all the upstream domains. As such, the signaling containment of the originating domain will render the resilience mechanism ineffective.

The proposed solution uses concatenated primary and backup LSP, protection signaling, and a domain boundary protection scheme to provide protection across multiple, independent domains. The domain boundary protection scheme includes the introduction of some new protection elements to pre-establish inter-domain local bypass tunnels. The proposed solution relies on some basic amount of information from neighboring domains and makes no assumption regarding protection mechanisms of other domains or levels of cooperation. A simulation experiment using OPNet was conducted. Three models were constructed and compared. The baseline model was based on the traditional layer three inter-domain routing protocol, BGP. The second model implemented MPLS recovery using an end-to-end path protection mechanism, and the third model used the proposed solution. The simulation results revealed the potential for this proposed solution for MPLS inter-domain protection.

The primary focus of this thesis research is to determine the factors involved in the MPLS failure recovery, particularly in the context of an intra-domain failure recovery. Nevertheless, these two papers have shown the ongoing effort by academic communities and industries to further research and develop the potential of MPLS, including the failure recovery capability.

H. SUMMARY

An overview of MPLS and its various key components such as label distribution protocol, MPLS virtual private network, MPLS traffic engineering and MPLS protection mechanisms were presented. Some related work on the MPLS protection mechanisms was also discussed.

In the next chapter, the laboratory set-up of the BGP/MPLS VPN network backbone, including the equipment configuration, is described. The tools and procedures for the data collection process are also detailed.

THIS PAGE INTENTIONALLY LEFT BLANK

III. LAB SETUP AND PROCEDURES

A. CHAPTER OVERVIEW

This chapter describes the laboratory set-up of the BGP/MPLS VPN network used for this thesis. The first section discusses the performance metrics to be evaluated. The next section covers the scope of the experimentation for this study. The chapter then presents the overall network architecture of the BGP/MPLS VPN laboratory setup and the basic configuration for the PE, P and CE routers. This is followed by a detailed description of the various parameters of interest, the test cases and the required router configuration for the testing. The last Section describes the tools used to collect the required raw data and explains the procedures to perform network failover and link restoration experiments.

B. PERFORMANCE METRICS

The objective of this thesis research is to study the performance of the BGP/MPLS VPN failover functionality. There are two quantitative metrics that can be used to evaluate this performance. They are the failover time and restoration time.

1. Failover Time

In an event of a network failover, the network traffic will be redirected to the next best available LSP provided by the routing tables. Alternatively, it will be rerouted from the primary LSP to the backup LSP if there are any pre-established alternate paths. One of the main concerns the customers of a service provider has is the packet delay experienced during the failover. Packet delay has direct impact on customers' applications. In some instances, the impact can be so severe that it causes degradation in the application's performance. In the case of real time applications such as military and financial systems, it is critical that the impact is negligible and will not affect these systems in meeting their goals.

Since packet delay is an externally observable event, the failover time can be measured in terms of the application's perspective. The network backbone can be treated as a black box. In this study, the failover time was based on the time difference between when the application stops receiving traffic due to the occurrence of the failure event and when it starts receiving traffic again.

2. Restoration Time

In the event of a network restoration, the traffic flow will be redirected back to the primary LSP from the backup LSP upon failure recovery. In this study, the restoration time is based on the time from the recovery of the link or the node to the time where the traffic starts traversing the original path, i.e. the time the first packet travels on the original path after the switch over.

This event should be transparent to the customer's applications. Traffic disruption is expected to be very negligible since both primary and backup path are working and care using techniques such as "make before break" method or synchronous switch-over [18], are usually taken to ensure the traffic disruption is well under control. Hence, the traffic should traverses smoothly from the backup path to the primary path and the customers' applications should not experience any packet delay in such circumstances.

Nonetheless, it is still important to keep the restoration time to the minimum. One of the reasons is because the backup path often does not support the same amount of traffic as the primary path. Furthermore, it might also not provide as good quality of service as the preferred working path. In addition, the backup path itself is seldom protected as well. In protection scenarios where the the number of backup paths are less than the primary paths and one backup path is allocated to support more than one primary LSPs, it is vital that the traffic is reverted back to the primary path at the earliest possible time to avoid creating choke points in the network backbone.

There are two mode of restoration, namely the manual mode and the automatic mode. The former mode would require the network administrator to

manually configure the switch over. As such, the network administrator can decide when would be the best time to perform the operation. In the situation where the switch over might result in traffic disruption, the network administrator can choose to perform the switch over at a later time. In the latter mode, the router would automatically switch back the traffic to the primary path upon detecting the primary path is working. In this case since there is no human intervention, it is important to ensure that the restoration is performed at the shortest possible time to ensure an optimized network backbone. As such, it is the interest of this study to investigate how much time is taken for a failure recovery to take place under different conditions. This would allow service provider to make the necessary configurations to ensure that the traffic is redirected back to the desired path at the earliest possible time.

C. SCOPE OF EXPERIMENT

In order to identify the contributing factors to the time delay of a network failure and recovery, and to determine the possible router configuration options that could be used to reduce the failover time and restoration time, it is the interest of this study to examine some of the key components in BGP/MPLS VPN, as highlighted in Chapter 2. This means manipulating some of their parameters or attributes to see if they have direct impact on the time delay. Some of the components to be investigated include the interior gateway protocol, the label distribution protocol and the MPLS traffic engineering. Due to time and resource constraints, it is not the aim of this study to examine each and every parameter/attribute of the various components of the BGP/MPLS VPN. However, through the initial literature research of this thesis work, the scope of the experimentation was limited to a few identified parameters for each component. These parameters are the ISIS metric assignment, ISIS Shortest Path First (SPF) intervals, TDP Discovery Hello intervals and the various MPLS TE tunnel configuration options. The details of these parameters will be elaborated upon in a subsequent section.

D. LABORATORY SET-UP

Figure 9 shows the implementation of the BGP/MPLS VPN network set-up in the laboratory. The BGP/MPLS VPN network is formed by 4 routers. Two P routers, whose main functionality was to perform label switching, were set up to form the core of the network backbone. Two PE routers, which were the main workhorse for this network backbone, were installed at the edge of the network. They formed the entry and exit points to the network. Each PE router was connected to two CE routers. Each CE router served as a link from the customer's network to the service provider's network. The four customer networks, as shown in the figure, are grouped into two virtual private networks, VPN A and VPN B.




Figure 9. Laboratory Set-Up of BGP/MPLS VPN Network

E. BASIC ROUTER CONFIGURATIONS

This Section shows the basic configurations required to set up the PE, P and CE routers. Other additional router commands were included subsequently into the router configuration to facilitate the experimentation. Nonetheless, these basic configurations were good enough to create the laboratory BGP/MPLS VPN network backbone.

1. Installing PE Router

The two PE routers were installed using Cisco 3620 routers running on Cisco Internetworking Operating System (IOS) version 12.2(3). Each PE router was connected to two P routers and two CE routers. All connections were established using Ethernet interfaces. Table 1 shows the basic configuration for one of the PE routers – PE1.

Basic Configuration of PE Router – PE1
Current configuration : ! version 12.2 ! hostname PE1 ! ! ! <i>define VRF instances</i> ip vrf Customer_A rd 100:110 route-target export 100:1000 route-target import 100:1000 ! ip vrf Customer_B rd 100:120 route-target export 100:2000 route-target import 100:2000 ! ! ! <i>enable CEF</i> ip cef ! ! ! <i>configure the loopback interface to be used as the</i>


```

! BGP update source and TDP router ID
interface Loopback0
ip address 10.10.10.4 255.255.255.255
ip router isis
!
!
! configure Ethernet interfaces for MPLS & IS-IS
interface Ethernet0/0
description Link to P2
ip address 10.1.1.5 255.255.255.252
ip router isis
tag-switching ip
!
interface Ethernet0/1
description Link to P1
ip address 10.1.1.14 255.255.255.252
ip router isis
tag-switching ip
!
!
! Configure VRF interfaces
interface Ethernet0/2
description Link to VPN A
ip vrf forwarding Customer_A
ip address 200.0.4.1 255.255.255.0
!
interface Ethernet0/3
description Link to VPN B
ip vrf forwarding Customer_B
ip address 200.0.4.1 255.255.255.0
!
!
! configure IS-IS as the MPLS VPN backbone
router isis
net 49.0001.0000.0000.0004.00
is-type level-1
metric-style wide
!
!
! configure global BGP parameters
router bgp 100
bgp log-neighbor-changes
neighbor 10.10.10.6 remote-as 100
neighbor 10.10.10.6 update-source Loopback0
!
! configure for PE-CE routing session

```

```

address-family ipv4 vrf Customer_B
redistribute connected
redistribute static
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf Customer_A
redistribute connected
redistribute static
no auto-summary
no synchronization
exit-address-family
!
! activate multiprotocol BGP route exchange
address-family vpnv4
neighbor 10.10.10.6 activate
neighbor 10.10.10.6 send-community both
exit-address-family
!
!
! configure static routes for the PE-CE connectivity
ip classless
ip route vrf Customer_A 172.120.0.0 255.255.255.0 200.0.4.2
ip route vrf Customer_B 172.120.0.0 255.255.255.0 200.0.4.2
!
!
end

```

Table 1. Basic Configuration of PE Router – PE1

The “tag-switching ip” command on the router configuration was used to enable MPLS for a network interface. It is an alternative to the command “mpls ip” command that is used in other Cisco IOS versions. In the sample configuration below, the “tag-switching ip” command was issued on all the network interfaces that were connected to the P routers.

The label distribution protocol that was used in the laboratory set-up was Cisco’s proprietary Tag Distribution Label (TDP) protocol. TDP is the default label distribution protocol on Cisco routers. It is very similar to the IETF standard protocol for label distribution – LDP. However, there are some minor differences

between the two protocols. Unlike TDP, LDP provides MD5 authentication. LDP uses multicast for neighbor discovery while TDP uses local broadcast. The protocols use different ports for neighbor discovery and session establishment [6].

IS-IS and OSPF are the two most commonly adopted IGPs for BGP/MPLS VPN network as they are the only two IGPs that support MPLS traffic engineering. There was no strong preference to use one IGP over the other for the laboratory set-up. IS-IS was chosen as the IGP for the experimentation. Like the “tag-switching ip” command, the “ip router isis” command was issued on all the network interfaces that were connected to the P routers to enable the IGP protocol in the various interfaces.

Different routing protocols such as RIP2, OSPF, eBGP or even static routing can be used for the connectivity between the PE and CE routers. Since the selection of any of these protocols had no impact on the performance analysis of the BGP/MPLS VPN, static routing was chosen for this case. Static routes were configured for each VRF in the PE router. The PE router then advertised the routes across the backbone using the multiprotocol BGP to the other PE router.

The configuration of the other PE router – PE2 is similar to the one in Table 1 except for the values of some parameters such as the IP addresses for the loopback address, network interfaces etc.

2. Installing P Router

In this laboratory set-up, each P router was connected to the PE routers only. The installation of the P routers was more straightforward than the PE routers mainly because the VRF and BGP configurations were not required for the P router. Like the PE routers, Cisco 3620 routers with Cisco IOS 12.2(3) were used for the P routers. Table 2 shows the basic configuration for one of the P router – P1.

Basic Configuration of P Router – P1

```
Current configuration :
!
version 12.2
!
hostname P1
!
!
! enable CEF
ip cef
!
!
! configure the loopback interface
interface Loopback0
ip address 10.10.10.1 255.255.255.255
ip router isis
!
!
! configure Ethernet interfaces for MPLS & IS-IS
interface Ethernet0/1
description Link to PE1
ip address 10.1.1.13 255.255.255.252
ip router isis
tag-switching ip
!
interface Ethernet0/2
description Link to PE2
ip address 10.1.1.21 255.255.255.252
ip router isis
tag-switching ip
!
!
! configure IS-IS as the MPLS VPN backbone
router isis
net 49.0001.0000.0000.0001.00
is-type level-1
metric-style wide
!
!
! ip classless
!
!
end
```

Table 2. Configuration of P Router – P1

3. Installing CE Router

The hardware specification requirements for the four CE routers are less stringent than the PE and P routers. As such, Cisco 2610 routers running on Cisco IOS 12.3(5) were used for the CE routers. Different routing protocols can be used to exchange routing information between the CE router and the PE router. In this laboratory set-up, the routing information was based on static routes. Routes to the other networks in the same VPN were manually configured. Table 3 shows the basic configuration for one of the CE routers – CE1A.

Basic Configuration of CE Router – CE1A
Current configuration: ! version 12.0 ! hostname CE1A ! interface Ethernet0/0 description Link to PE1 (Customer A) ip address 200.0.4.2 255.255.255.0 ! interface Ethernet1/1 ip address 172.120.0.1 255.255.255.0 ! ip classless ip route 192.168.0.0 255.255.255.0 Ethernet0/0 ip route 200.0.6.0 255.255.255.0 Ethernet0/0 ! end

Table 3. Configuration of CE Router – CE1A

F. EXPERIMENTAL TESTING AND CONFIGURATIONS

A total of 23 test cases were constructed to examine the parameters of the various key components of the BGP/MPLS VPN. Each test case differed based on either the location of the failure, the type of failure event, or the value for a particular parameter of a component. In this experiment, the location of the

failure was limited either at the headend router, PE1, or at the midpoint routers, P1 or P2. The type of failure event was either a link failure or a node failure.

Each test case consists of a set of 15 readings for both failover timing and restoration timing. Basic statistical measures such as the mean and standard deviation were generated from each set of readings to provide better understanding on the data collected. The test cases were grouped according to the parameters of interest.

Below describes the identified parameters of the various key components for the experiment, their associated test cases and the router configurations to make the necessary changes to the values of the parameters.

1. Varying ISIS Metric Value

The ISIS metric is the metric assigned to the links and it is used to calculate the cost from each other router via the links in the network to other destinations. The metric can be configured for Level 1 or Level 2 routing. The range is from 0 to 63 and the default value is 10 [19].

The symmetrical network topology in the laboratory setup had provided two equal-link-cost LSPs from PE1 to PE2 by default. There was no preference over any of the two available LSPs when forwarding traffic from PE1 to PE2. However, manipulating the metric parameter of any link by setting a different value would result in the headend router preferring a particular LSP at all time until the metric parameter was reconfigured again. It is the interest of this study to investigate how the values of this parameter affect the performance of the BGP/MPLS VPN failover functionality, in particularly the failover time and the restoration time.

In the experimental testing of this parameter, two sets of test cases were performed to determine the difference in the failover time and restoration time when the links were configured with different metric values. The first set of test cases (Test Cases 1-4), were based on the ISIS default metric value which is 10. The second set of test cases (Test Cases 5-8) were configured such that one of

the links, which is the link connecting PE1 to P2, had a higher metric value, say 40. This would result in PE1 preferring the path to PE2 via P1. Table 4 shows the router command to assign a different link metric to a particular interface.

```
PE1 (config) # interface ethernet0/1

PE1 (config-if) # isis metric 40 level-1

PE1 (config-if) # exit
```

Table 4. Router Command to Assign a Different Metric to an Interface

The Table below shows a sample router configuration with the changes made to the ISIS metric for one of the Ethernet interfaces.

Sample Router Configuration for ISIS metric
<pre>Current configuration : ! version 12.2 ! hostname P1 ! ! enable CEF ip cef ! interface Loopback0 ip address 10.10.10.1 255.255.255.255 ip router isis ! interface Ethernet0/1 ip address 10.1.1.13 255.255.255.252 ip router isis isis metric 40 level-1 tag-switching ip ! interface Ethernet0/2 ip address 10.1.1.21 255.255.255.252 ip router isis tag-switching ip ! router isis net 49.0001.0000.0000.0001.00 is-type level-1 metric-style wide</pre>

```

!
ip classless
!
end

```

Table 5. Sample Router Configuration for ISIS metric

Each set of test cases measured the failover time and restoration time at different locations of failure as well as different types of failure events except for Test Cases 4 and 8 where only the failover time was measured. The restoration time for a node failure was briefly measured and it took a few minutes for the restoration process to complete. It was mainly due to the startup of the operating system which was dependent on the IOS and hardware used. Hence, the scope of experimentation did not cover the restoration time for node failure. MPLS TE was not used for this testing. There was no pre-configured standby LSP. A backup LSP was established dynamically only after the failure had occurred. In addition, the TDP discovery intervals and the ISIS SPF intervals were configured at their default values. Table 6 below describes the experimental parameters for each test case.

Experimental Parameters			
	Metric Assignment	Type of Failure	Location of Failure
Test Case 1	Equal	Link Failure	Interface at PE1, linking to P1
Test Case 2			Interface at P1, linking to PE1
Test Case 3			Interface at P1, linking to PE2
Test Case 4		Node Failure	P1
Test Case 5	Unequal	Link Failure	Interface at PE1, linking to P1
Test Case 6			Interface at P1, linking to PE1
Test Case 7			Interface at P1, linking to PE2
Test Case 8		Node Failure	P1

Table 6. Experimental Parameters for Test 1 - 8

2. Varying ISIS SPF Intervals

Unlike distance vector protocol, one of the unique characteristic of a link state protocol such as ISIS is to hold routing information for some period of time after receiving them before processing it [20]. This is mainly due to the shortest path first algorithm. A router running ISIS would hold on the SPF computation after receiving a link state packet. The longer the wait period is, the number of times the algorithm has to be executed would be reduced since more update packets are allowed to arrive before performing the calculation. As such, the wait period for the SPF computation would help to reduce the overall load on the processor and memory.

There are three shortest path first interval timers in the Cisco IOS software [20] namely the `spf-max-wait`, the `spf-initial-wait` and the `spf-second-wait`. The first interval is the minimum time that should elapse between consecutive shortest path first computations. The range is from 1 to 120 seconds and the default value is 10 seconds. The second timer is the number of milliseconds between the receipt of new link state information and running SPF. It indicates the initial SPF calculation delay after receiving an update of a topological change. The range is 1 to 120000 milliseconds. The default is 5500 milliseconds. The third timer is the minimum wait time between the first and subsequent SPF calculations. The range and default values are the same as the second timer. It is the interest of this study to investigate how each of these SPF intervals impacts the performance of the BGP/MPLS VPN failover functionality, in particularly the failover time and the restoration time.

In the experimental testing of the SPF intervals, five test cases (Test 9 - 13) were performed to determine the failover time and restoration time when the SPF intervals were configured with different values. Table 7 shows the router command to configure the SPF intervals. In this instance, the `spf-max-wait` timer was set at 1 second; the `spf-initial wait` and the `spf-second-wait` timer were set at 500 milliseconds respectively.

```

PE1 (config) # router isis

PE1 (config-router) # spf-interval 1 500 500

PE1 (config-router) # exit

```

Table 7. Sample Router Command to Configure SPF Intervals

The Table below shows a sample router configuration with the changes made to the SPF interval.

Sample Router Configuration for ISIS SPF Interval
<pre> Current configuration : ! version 12.2 ! hostname P1 ! ip cef ! interface Loopback0 ip address 10.10.10.1 255.255.255.255 ip router isis ! interface Ethernet0/1 ip address 10.1.1.13 255.255.255.252 ip router isis tag-switching ip ! interface Ethernet0/2 ip address 10.1.1.21 255.255.255.252 ip router isis tag-switching ip ! router isis net 49.0001.0000.0000.0001.00 is-type level-1 spf-interval 1 500 500 metric-style wide ! end </pre>

Table 8. Sample Router Configuration for ISIS SPF Interval

Each test case measured the failover time and restoration time based on the link failure located at the interface of P1 connecting to PE2. MPLS TE was not used for this testing. There was no pre-configured standby LSP and a backup LSP was established dynamically only after the failure had occurred. In addition, the TDP discovery intervals, the ISIS metric and hello interval were configured at their default values. Table 9 below describes the difference in the SPF interval settings for Test 9 -13.

Experimental Parameters for Test 9 - 13		
Test 9	spf-max-wait	1 seconds
	spf-initial-wait	500 milliseconds
	spf-second-wait	500 milliseconds
Test 10	spf-max-wait	1 seconds
	spf-initial-wait	1000 milliseconds
	spf-second-wait	1000 milliseconds
Test 11	spf-max-wait	3 seconds
	spf-initial-wait	1000 milliseconds
	spf-second-wait	1000 milliseconds
Test 12	spf-max-wait	3 seconds
	spf-initial-wait	3000 milliseconds
	spf-second-wait	3000 milliseconds
Test 13	spf-max-wait	10 seconds
	spf-initial-wait	10000 milliseconds
	spf-second-wait	10000 milliseconds

Table 9. Experimental Parameters for Test 9 – 13

3. Varying ISIS Hello Intervals

The ISIS Hello packets are used to initialize and maintain adjacencies between neighboring routers [21]. There are three types of IS-IS Hello packets namely the Level 1 LAN IS to IS Hello Protocol Data Units (PDUs), the Level 2 LAN IS to IS Hello PDUs and the Point-to-Point IS to IS Hello PDUs. The Level 1 Hello packets are used by Level 1 routers on broadcast LANs; Level 2 LAN Hello

packets are used by Level 2 routers on broadcast LANs; Point-to-Point IS to IS Hello packets are used on non-broadcast media, such as point-to-point links. In this study, the focus is on the Level 1 Hello packets.

There are two parameters for configuring the ISIS Hello interval. The ISIS hello interval parameter specifies the length of time between hello packets that the router sends. The ISIS hello multiplier parameter specifies the number of hello packets a neighbor must miss before the router should declare the adjacency as down. In essence, the hello interval is multiplied by the hello multiplier to determine the hold interval. The default value for the ISIS hello interval is 10 seconds. The default multiplier value is 3 which is also the minimum value that can be set.

It is the interest of this study to investigate how the hello interval impacts the performance of the BGP/MPLS VPN failover functionality, in particularly the failover time and the restoration time. In the scenario where the link fault detection is superseded by the line protocol error message, fast hello intervals will most unlikely improve the detection of a network topological change since the router will be immediately notify of the loss of line protocol on point-to-point links. However, in the absence of the line protocol assistance, we would like to study the response of the failover and restoration time with respect to the hello interval.

In the experimental testing of the hello intervals, five test cases (Test 14 - 18) were performed to determine the failover time and restoration time when the ISIS hello intervals were configured with different values. Table 10 shows the router command to configure the hello interval on a particular interface. In this instance, the hello interval was set at 1 second and the hello multiplier at 4.

```
PE1 (config) # interface ethernet0/1
PE1 (config-if) # isis hello-interval 1
PE1 (config-if) # isis hello-multiplier 4
PE1 (config-if) # exit
```

Table 10. Sample Router Command to Configure Hello Interval

The Table below shows a sample router configuration with the changes made to the Hello interval.

Sample Router Configuration for ISIS Hello Interval
<pre> Current configuration : ! version 12.2 ! hostname P1 ! ! enable CEF ip cef ! interface Loopback0 ip address 10.10.10.1 255.255.255.255 ip router isis ! interface Ethernet0/1 ip address 10.1.1.13 255.255.255.252 ip router isis isis hello-interval 1 isis hello-multiplier 4 tag-switching ip ! interface Ethernet0/2 ip address 10.1.1.21 255.255.255.252 ip router isis tag-switching ip ! router isis net 49.0001.0000.0000.0001.00 is-type level-1 metric-style wide ! ip classless ! end </pre>

Table 11. Sample Router Configuration for ISIS Hello Interval

Each test case measured the failover time and restoration time based on the link failure located between P1 and PE2. The next Section would elaborate the set up to simulate such a link failure where no line protocol error message

would be detected by the routers. MPLS TE was not used for this testing. There was no pre-configured standby LSP and a backup LSP was established dynamically only after the failure had occurred. In addition, the TDP discovery intervals, the ISIS metric and SPF Intervals were configured at their default values. Table 12 below describes the difference in the Hello interval settings for Test 14 -18.

Experimental Parameters for Test 14 - 18	
Test Case	Hello Interval
Test 14	1 second
Test 15	2 second
Test 16	3 second
Test 17	5 second
Test 18	10 second

Table 12. Experimental Parameters for Test 14 – 18

4. Varying TDP Discovery Hello Intervals

LSR uses the LDP/TDP discovery mechanism to discover potential LDP/TDP peers. In the laboratory setup, the PE and P routers use this mechanism to discover one another by periodically sending TDP Hello messages, in the form of UDP packets, to a specific port number, port number 711. Upon receipt of the TDP Hello messages from its neighbors, the router would establish the corresponding TDP adjacencies.

There are two parameters in the TDP discovery mechanism: the hello interval and the hello holdtime interval. The hello interval is the period of time between the sending of consecutive Hello messages. The default interval is 5 seconds. The hello holdtime interval is the period of time a discovered TDP neighbor is remembered without the receipt of a TDP Hello message from the neighbor. The default interval is 15 seconds. This study investigated how each of

these TDP intervals impacts the performance of the BGP/MPLS VPN failover functionality.

In the experimental testing of the TDP discovery intervals, six test cases (Test 19 -24) were performed to determine the failover time and restoration time when the TDP discovery intervals were configured with different values. Table 13 shows the router command to configure the TDP discovery intervals for a particular interface. In this instance, the hello interval was set at 1 second; the hello holdtime was set at 3 seconds.

```
PE1 (config) # interface ethernet0/0
PE1 (config-if) # tag-switching tdp discovery hello interval 1
PE1 (config-if) # tag-switching tdp discovery hello holdtime 3
PE1 (config-if) # exit
```

Table 13. Sample Router Command to Configure TDP Discovery Intervals

The Table below shows a sample router configuration with the changes made to the TDP Hello interval.

Sample Router Configuration for TDP Hello Interval
<pre>Current configuration : ! version 12.2 ! hostname P1 ! ! enable CEF ip cef ! tag-switching tdp discovery hello interval 1 tag-switching tdp discovery hello holdtime 3 ! interface Loopback0 ip address 10.10.10.1 255.255.255.255 ip router isis ! interface Ethernet0/1 ip address 10.1.1.13 255.255.255.252</pre>

```

ip router isis
tag-switching ip
!
interface Ethernet0/2
 ip address 10.1.1.21 255.255.255.252
 ip router isis
 tag-switching ip
 !
router isis
 net 49.0001.0000.0000.0001.00
 is-type level-1
 metric-style wide
 !
ip classless
!
end

```

Table 14. Sample Router Configuration for TDP Hello Interval

Each test case measured the failover time and restoration time based on the link failure located at the interface of P1 connecting to PE2. MPLS TE was not used for this testing. There was no pre-configured standby LSP and a backup LSP was established dynamically only after the failure had occurred. In addition, the SPF intervals, Hello interval and the ISIS metric were configured at their default values. Table 15 below describes the difference in the TDP discovery interval settings for Test 19 -24.

Experimental Parameters for Test 14 - 19		
Test 19	hello interval	1 seconds
	holdtime interval	3 seconds
Test 20	hello interval	2 seconds
	holdtime interval	6 seconds
Test 21	hello interval	3 seconds
	holdtime interval	9 seconds
Test 22	hello interval	10 seconds
	holdtime interval	30 seconds

Test 23	hello interval	15 seconds
	holdtime interval	45 seconds
Test 24	hello interval	20 seconds
	holdtime interval	60 seconds

Table 15. Experimental Parameters for Test 19 - 24

5. Varying MPLS TE Tunnel Configuration Options

The following test cases were constructed to determine the performance of the BGP/MPLS VPN failover functionality when MPLS traffic engineering was deployed. The MPLS VPN traffic can be carried over MPLS TE tunnels. In this laboratory set-up, TE tunnels were configured to carry traffic between the two PE routers. From a Layer 2 perspective, a MPLS tunnel interface, which is configured with a set of requirements such as bandwidth and media requirements, denotes the head of an LSP. Whereas from a Layer 3 perspective, a LSP tunnel interface represents the head-end of a unidirectional virtual link to the tunnel destination [20].

As highlighted in Chapter 2, the LSP tunnels are computed at the LSP head based on a fit between required and available resources. There are two options to determine the path for a tunnel. The first option uses the MPLS TE calculation module to determine the path. The calculation module uses the TE link state database which contains flooded topology and resource information to determine the best path [22]. This database is updated by the IGP flooding whenever a change such as the establishment of new LSP or change of bandwidth, occurs. Alternatively, the path for a LSP tunnel can be determined by explicit routing, by which the users are allowed to dictate the path. A tunnel can contain multiple explicit path options but at most only one path determined by the TE path calculation module is allowed. After the path calculation, RSVP is used to signal and maintain the LSP tunnel at each hop along the LSP.

In this testing of the MPLS traffic engineering, four test cases (Test 25 -28) were performed to determine the failover time and restoration time when the MPLS tunnels were used. As the deployment of MPLS TE involves very detailed

router configuration for every router along the LSP tunnel, this section will not provide the complete list of router commands for enabling MPLS TE in the network. For details, please refer to Cisco website: http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120lmit/120s/120s5/mpls_te.htm. Table 16 shows the router command to configure a MPLS TE tunnel.

```

PE1 (config) # interface tunnel0

PE1 (config-if) # ip unnumbered loopback0

PE1 (config-if) # tunnel destination 10.10.10.6

PE1 (config-if) # tunnel mode mpls traffic-eng

PE1 (config-if) # tunnel mpls traffic-eng priority 1 1

PE1 (config-if) # tunnel mpls traffic-eng bandwidth 512

PE1 (config-if) # tunnel mpls traffic-eng path-option 1 dynamic

PE1 (config-if) # tunnel mpls traffic-eng path-option 2 explicit name BK

PE1 (config-if) # tunnel mpls traffic-eng autoroute announce

PE1 (config-if) # exit

```

Table 16. Sample Router Command to Configure a MPLS TE Tunnel

The Table below shows a sample router configuration with the MPLS Traffic Engineering enabled

Sample Router Configuration for MPLS TE
<pre> Current configuration : ! version 12.2 ! hostname P1 ! ! enable CEF ip cef mpls traffic eng tunnels ! </pre>

```

!
interface Loopback0
 ip address 10.10.10.1 255.255.255.255
 ip router isis
!
interface Tunnel0
ip unnumbered Loopback0
tunnel destination 10.10.10.6
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 1 1
tunnel mpls traffic-eng path-option 1 dynamic
tunnel mpls traffic-eng path-option 2 explicit name BK
tunnel MPLS traffic-eng bandwidth 512
!
interface Ethernet0/1
 ip address 10.1.1.13 255.255.255.252
 ip router isis
 tag-switching ip
 mpls traffic eng tunnels
 ip rsvp bandwidth 1024
!
interface Ethernet0/2
 ip address 10.1.1.21 255.255.255.252
 ip router isis
 tag-switching ip
 mpls traffic eng tunnels
 ip rsvp bandwidth 1024
!
router isis
 net 49.0001.0000.0000.0001.00
 is-type level-1
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-1
!
ip explicit-path name BK enable
 next-address 10.1.1.6
 next-address 10.1.1.10
 next-address 10.10.10.6

ip classless
!
end

```

Table 17. Sample Router Configuration for MPLS Traffic Engineering

Each test case measured the failover time and restoration time based on the link failure located at the interface of P1 connecting to PE2. The number of tunnels and the path options were configured differently for each test case. Table 18 below describes the configuration of each test case.

Experimental Parameters for Test 20 - 23			
	No of Tunnels	No & Type of Path Options	SPF Intervals
Test 25	2	1 explicit path for each tunnel	Default
Test 26	1	2 explicit paths	Default
Test 27	1	1 dynamic path	Default
Test 28	2	1 explicit path for each tunnel	spf-initial-wait – 0.5s

Table 18. Experimental Parameters for Test 25 - 28

6. Varying Static and Non-static Routing Configuration

The following test cases were constructed to determine the performance of the BGP/MPLS VPN failover functionality in the case where static routing is preferred instead of IGP. Although static routing requires manual reconfiguration in the event of a network change and it is not as robust as IGP since there is no automatic routing around a network outage [23], it does have its own advantages. Firstly, it is easy to configure and secondly, there is no routing protocol overhead.

In the experimental testing of the static routing, four test cases (Test 29 - 32) were performed to determine the failover time and restoration time when static routes were used instead of using an IGP. Table 19 shows the router command to configure a static route.

```
PE1 (config) # ip route 10.10.10.3 255.255.255.255 10.1.1.6
```

Table 19. Sample Router Command to Configure a Static Route

The Table below shows a sample router configuration with the static routes used

Sample Router Configuration for Static Routing
<pre> Current configuration : ! version 12.2 ! hostname P1 ! ! enable CEF ip cef ! interface Loopback0 ip address 10.10.10.1 255.255.255.255 ! interface Ethernet0/1 ip address 10.1.1.13 255.255.255.252 tag-switching ip ! interface Ethernet0/2 ip address 10.1.1.21 255.255.255.252 tag-switching ip ! ip classless ip route 10.10.10.6 255.255.255.255 10.1.1.6 ip route 10.10.10.6 255.255.255.255 10.1.1.13 ! end </pre>

Table 20. Sample Router Configuration for Static Routing

Like Test Cases 1-4, Test Cases 29 – 32 measured the failover time and restoration time at different location of failures as well as different types of failure events. MPLS TE and IGP were not used for this testing. There was two manually configured static LSPs from the headend router PE1 to the tailend router PE2. Hence, in the event where one of the LSP fails, the headend router would switch the traffic over to the other available LSP. In this test scenario, the TDP discovery interval was configured at its default value. Table 21 below describes the experimental parameters for each test case.

	Type of Failure	Location of Failure
Test Case 29	Link Failure	Interface at PE1, linking to P1
Test Case 30		Interface at P1, linking to PE1
Test Case 31		Interface at P1, linking to PE2
Test Case 32	Node Failure	P1

Table 21. Experimental Parameters for Test 29 - 32

G. TOOLS AND PROCEDURES FOR DATA COLLECTION PROCESS

A variety of software programs were used to collect the raw data. These tools provided the ability to perform traffic generation, failure and recovery simulation, packet capturing and traffic monitoring. They allowed data to be collected in a precise manner so that the statistical analysis based on these data would be correct.

1. Traffic Generation

In order to simulate real traffic across the network backbone between two applications, a traffic generation program named Bricks was used. The program was installed in two workstations, each located at a different side of the network backbone but on the same VPN. The program provides two modes of operations, one as the data transmitter and the other as the data receiver. During the data collection, one workstation would assume the role of the transmitter while the other workstation would assume the role of the receiver and the traffic transmitted would flow across the network backbone.

The program supports the transmission of a variety of layer four protocols including TCP and UDP packets. In this study, UDP was selected as the type of packets for transmission. In addition, the program also allows the specification of the transmission rate and packet size. Figure 10 shows the screenshot of the traffic generation program for the transmitter mode.




Figure 10. Screenshot of the GUI of the Bricks Program.

2. Failure and Recovery Simulation

There are several possible causes that would lead to network failover, such as link or node failure, administration change, setting of IGP overload bit and path optimization [19]. In this study, the network failover events were limited to link and node failures.

There were two options to simulate a link failure in the laboratory set-up. The first option is by pulling the network cable off the interface. The second option is by issuing a “shutdown” command on the specific network interface through the router command line interface. Both options were tested and the results were indifferent. As such, the option of issuing the shutdown command to the network interface was preferred. Similarly, a node failure can be triggered by the turning off the power button of the router or by issuing a “reload” command through the command line interface. Likewise, the latter option was chosen. The

Table below shows the procedures to simulate a link and node failure as well as recovery.

<u>To simulate a link failure on interface ethernet0/0</u> PE1 (config) # interface ethernet0/0 PE1 (config-if) # shutdown PE1 (config-if) # exit
<u>To simulate a link recovery on interface ethernet0/0</u> PE1 (config) # interface ethernet0/0 PE1 (config-if) # no shutdown PE1 (config-if) # exit
<u>To simulate a node failure on a router</u> PE1 # reload

Table 22. Router Commands to Simulate a Link/Node Failure and Recovery

In order to simulate a link failure where the fault detection is not based on the line protocol signaling, a slight modification to the network topology was made specifically for Test Case 14 – 18. Figure 11 shows the modification. For Test Case 14 – 18, the link failure would not be initiated based on the steps as described above. Instead, the link failure would be initiated by pulling the cable at the hub that is closer to the router PE2. In this case, the midpoint router, P1 would rely on the ISIS hello packets to detect the topological changes.




Figure 11. Modification to Network Topology for Test Case 14 - 18.

3. Packet Capturing

For the data capturing process, an open source packet sniffer program, Wireshark was utilized. It was used to measure the failover time and recovery time, up to the precision of milliseconds. Figure 12 shows the screenshot of the GUI of the Wireshark program where the details of the packets that had been captured were displayed. This included the time of capture, the source and destination IP addresses and ports, the protocols used and the contents of the packets.




Figure 12. Screenshot of the GUI of the Wireshark Program.

In order to measure the failover and restoration time, Wireshark was installed at two workstations. One of them was the same workstation that had Bricks installed with the data receiver mode. Prior to the trigger of the link or node failure, the packet sniffer program was set to monitor all incoming traffic to the workstation including the UDP packets send from the Bricks program that was running on the other workstation. Upon the occurrence of the link or node failure, the time gap between the halt and the re-admission of the incoming UDP packets were captured and displayed by the program. This time gap was recorded as the failover time for the particular measurement.

The second workstation was located inside the BGP/MPLS VPN network. This was to allow the measurement of the restoration time upon the node or link failure recovery. Some minor topological changes were made to the network backbone to allow the sniffing of packets along the original path. Figure 13

shows the minor changes to the original network topology so as to allow the data capturing process.




Figure 13. Modification to Network Topology to Allow Data Capturing

An Ethernet switch was deployed between PE1 router and P1 router. It had been identified that the primary path would traverse from PE1 to PE2 via P1. A workstation installed with Wireshark was connected to the Ethernet switch. One of the ports in the Ethernet switch was configured as a SPAN port to allow the packet sniffer program to capture all packets traversing between PE1 and P1. Using this set-up, the measurement of the time where the link or node failure recovery was activated to the time where the first packet of the traffic flow generated from the Bricks program began to traverse the original path was made.

3. Traffic Monitoring

In addition to the traffic generation program and the packet sniffer program, two other tools were used to assist the monitoring of the traffic along the paths in the network backbone and the activities taking place in the routers. They were SolarWinds Network Management Software (standard edition) and Kiwi Syslog Daemon. The SolarWinds Network Management Software provides a suite of network performance monitoring tools. One of them, the “Bandwidth Gauges” tool was mostly used to monitor the amount of data being received and transmitted by the routers. It also served the purpose of verifying the traffic flow along paths at any instant.




Figure 14. Screenshot of the GUI of the Bandwidth Gauges Feature in the SolarWinds Network Management Software.

The Kiwi Syslog Daemon captures all the Simple Network Management Protocol (SNMP) traps sent by the routers. The PE and P routers were

configured with the SNMP traps and message logging enabled. All the relevant SNMP traps and debug messages were sent to the workstation installed with the Kiwi Syslog Daemon. The information captured provided very useful insights on the activities happening inside the PE and P routers. Observing the activities such as the signaling of fault detection due to link failure and other control plane signaling facilitates the analysis of the behavior of the routers in the event of network failover and restoration. Figure 15 shows the screenshot of the GUI of the Kiwi Syslog Daemon that was used to capture all the SNMP and debugging messages.




Figure 15. Screenshot of the GUI of the Kiwi Syslog Daemon

H. SUMMARY

The performance metrics identified for the experimental testing and the scope of the experimentation were discussed. The configuration of the PE, P and CE routers for the BGP/MPLS VPN test-bed network were also presented. In addition, a detailed description of the various parameters of interest, the test cases and the required router configurations for the testing were also provided. Finally, the tools and procedures for the data collection process were also described.

The next chapter shows the results obtained from the data collection process and presents a detailed statistical analysis on the results.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

A. CHAPTER OVERVIEW

This chapter presents the statistical results generated from each test case and provides a detailed analysis of the statistical results that was performed to determine the contributing factors for the time delay in a network failure or recovery.

B. RESULTS AND ANALYSIS

Based on the 15 readings collected each for both failover and restoration time in every test case, the corresponding statistical results – the mean and the standard deviation were generated using Microsoft Excel program. The results are presented according to the associated parameter of interest. Each set of results is followed by a detailed analysis.

1. Varying ISIS Metric Value

Table 20 and 21 show the statistical results on the failover time and restoration time for Test Cases 1-8.

	Mean	Std Dev	Range
Test Case 1	0.030s	0	0
Test Case 2	1.064s	0.062s	1.007 - 1.204s
Test Case 3	5.598s	0.086s	5.590 - 5.614s
Test Case 4	1.117s	0.011s	1.103 - 1.140s
Test Case 5	6.138s	0.259s	5.654 - 6.530s
Test Case 6	6.050s	0.255s	5.431 - 6.258s
Test Case 7	5.619s	0.005s	5.609 - 5.626s
Test Case 8	7.621s	0.011s	7.061 – 8.279s

Table 23. Statistical Results on Failover Time for Test Cases 1 – 8

	Mean	Std Dev	Range
Test Case 1	7.814s	0.986s	6.165 - 9.721s
Test Case 2	6.443s	1.630s	4.744 - 8.877s
Test Case 3	8.081s	0.727s	7.106 - 9.547s
Test Case 4	-	-	-
Test Case 5	7.769s	1.072s	6.276 - 9.724s
Test Case 6	6.883s	0.666s	4.805 - 7.592s
Test Case 7	7.755s	0.864s	6.302 – 9.483s
Test Case 8	-	-	-

Table 24. Statistical Results on Restoration Time for Test Cases 1 - 8

The protection mechanism in this experimental case was based on headend reroute where the backup path would only be established after a link or node failure had occurred. As mentioned earlier, due to the symmetrical topology of BGP/MPLS VPN network set-up, the primary and backup LSP from PE1 to PE2 had the same link cost by default. In order to make one LSP to have a higher link cost than the other, the ISIS metric for one of the links was configured with a higher value. In that case, the primary LSP had a lower link cost than the backup path. This would result in the network topology being “non-symmetrical”. In this particular experimental case, the author compared the failover time and restoration time between the symmetrical and “non-symmetrical” network at different locations of failure, as well as the different failure type. Figure 8 shows the mean failover time for Test Cases 1 – 8. Two test cases having the same failure location or same failure type were grouped together, in which one had a symmetrical topology while the other had a “non-symmetrical” topology.




Figure 16. Mean Failover Time for Test Cases 1 - 8

Except for Test Cases 3 and 7, where the link failure occurred at P1 interface linking to PE2, which was non-adjacent to the headend router, the results from the rest of the test cases showed that the failover time was much faster for a symmetrical network than a “non-symmetrical” network. That’s to say that when both the primary and backup paths had the same link cost, the failover time from the primary to the backup path was shorter. In the case where the link failure occurred at PE1 and the network was symmetrical, the average failover time was 30 milliseconds. The average failover time, however, was about 6 seconds for a “non-symmetrical” network, with the failure happening at the same location.

Upon investigation, it was discovered that the routing table contained two routing entries to the same destination, PE2, for the symmetrical network. One of them indicated P1 as the next hop destination while the other pointed to P2. Both entries were derived from the ISIS protocol. As for the “non-symmetrical” network, there was only one routing entry at any one time for destination PE2, which was the route that had the lower metric. As such, when the headend router

in a symmetrical network detected that the existing route could no longer be used due to the link or node failure, it was able to use the next available LSP from the routing table immediately. Hence, it could switch the traffic over in a shorter period of time. On the contrary, the headend router in a “non-symmetrical” network would be required to compute a new route to the destination since there was no available routes to be used in the routing table.

There was no significant difference in the failover time between Test Cases 3 and 7. Unlike the rest of the test cases, the failure in Test Cases 3 and 7 occurred away from the headend router. There were no line protocol error messages received by PE1 when the link failure occurred. As such, the fault detection at the PE1 was based on the signaling from the IGP. This could possibly contribute to the time delay for the failover time in Test Case 3.

Figure 17 shows the mean restoration time for Test Case 1 – 8. Similar to the figure above, two test cases with the same failure location or same failure type were grouped together, in which one had a symmetrical topology while the other had a “non-symmetrical” topology. The results showed that the difference in mean restoration time between a symmetrical and “non-symmetrical” network was trivial. In both cases, the PE1 required some lead time to compute the new route after the change in network topology as a result of the link or node recovery.




Figure 17. Mean Restoration Time for Test Cases 1–3 & 5-7

2. Varying ISIS SPF Intervals

Table 22 and 23 show the statistical results on the failover time and restoration time for Test Cases 9 - 13.

	Mean	Std Dev	Range
Test Case 9	0.599s	0.005s	0.586 – 0.606s
Test Case 10	1.104s	0.006s	1.098 – 1.117s
Test Case 11	1.096s	0.022s	1.018 – 1.117s
Test Case 12	3.102s	0.007s	3.090 – 3.114s
Test Case 13	10.100s	0.010s	10.087 – 10.123s

Table 25. Statistical Results on Failover Time for Test Cases 9 – 13

	Mean	Std Dev	Range
Test Case 9	5.033s	0.294s	4.503 – 5.542s
Test Case 10	5.062s	0.463s	4.445 – 5.799s
Test Case 11	5.354s	0.636s	4.442 – 6.708s
Test Case 12	6.121s	1.142s	5.180 – 9.793s

Test Case 13	17.764s	2.369s	12-117 – 19.872s
---------------------	---------	--------	------------------

Table 26. Statistical Results on Restoration Time for Test Cases 9 - 13

Figure 18 below shows the mean failover time and restoration time with respect to the spf-initial-wait interval. The graph clearly showed the direct correlation between the mean failover time and the spf-initial-wait interval. The mean failover time decreased as the spf-initial-wait interval decreased. The default value of the spf-initial-wait interval is 5.5 seconds. The mean failover time measured in Test Case 3 with the spf-initial-wait interval configured at the default value was 5.598 seconds. In the extreme case where the spf-initial-interval interval was configured at 500 milliseconds, it took the 599 milliseconds for the headend router, PE1, to switch the traffic over to the backup LSP.

The spf-initial-wait interval is the time the router will wait after receiving new link state information, before performing the shortest path first algorithm to determine the new best route. It can be set as low as one millisecond although the recommended lowest spf-initial-wait is 40 milliseconds [20].




Figure 18. Mean Failover and Restoration Time with respect to SPF interval

Most routers run on single processor. If the router runs the SPF computation immediately after receiving a new link state packet and if the second link state packet arrives at the time when the first SPF computation is still running, then the second link state packet will be put on hold. It will wait until the first SPF computation is finished and the new information is flooded to the router's neighbors before the running the SPF computation again. After the second SPF computation, the router will flood the next new information again. However, if the wait period is sufficiently long, this will allow the link state packets to be processed in batch and reduce the number of flooding. As such, the router would have a chance to gather information from several updates before running the SPF algorithm so that it does not have to run the algorithm more than absolutely necessary. This will not only assist to reduce the router processing and memory load, but also avoid causing slow network convergence if the network is showing instability. Therefore, a balance of setting the appropriate spf-

initial-wait interval is required to ensure no long failover delay or slow network convergence.

The graph also indicated a correlation between the spf-initial-wait interval and the restoration time. The time required to switch the traffic back to the primary LSP increased as the spf-interval-wait interval increased. When the spf-initial interval was set at 3 seconds, the mean restoration time was measured at about 6 seconds. It increased to around 8 seconds when the spf-interval-wait was 5.5 seconds. It then took about 18 seconds to switch the traffic over when the spf-interval-wait interval was configured at 10 seconds. However, it was also observed that the restoration time decreased at a slower rate as the spf-initial-wait interval was reduced. For instance, the differences in restoration times between Test Cases 9, 10 and 11 were negligible.

A comparison was made between Test Cases 10 and 11 to determine the impact spf-max-wait interval had on the failover and restoration times. The insignificant results from the two test cases, however suggested no strong correlation among them.

3. Varying ISIS Hello Intervals

Table 24 and 25 show the statistical results on the failover time and restoration time for Test Cases 14 - 18.

	Mean	Std Dev	Range
Test Case 14	8.191s	0.391s	6.969 – 8.495s
Test Case 15	9.066s	0.491s	7.890 – 9.477s
Test Case 16	9.951s	0.432s	8.953 – 10.424s
Test Case 17	11.185s	1.160s	8.734 – 12.422s
Test Case 18	12.335s	1.746s	10.592 – 15.808s

Table 27. Statistical Results on Failover Time for Test Cases 14 – 18

	Mean	Std Dev	Range
Test Case 14	4.931s	0.271s	4.637 – 5.677s
Test Case 15	9.646s	0.902s	8.448 – 11.401s
Test Case 16	15.702s	2.403s	11.161 – 18.338s
Test Case 17	15.877s	0.530s	15.120 – 17.009s
Test Case 18	18.565s	0.624s	17.775 – 19.687s

Table 28. Statistical Results on Restoration Time for Test Cases 14 - 18

Figure 19 below shows the mean failover time and restoration time with respect to the hello interval. The graph clearly showed the linear relationship between the mean failover time and the hello interval as well as the mean restoration time and the hello interval. Both mean failover time and restoration time decreased as the hello interval decreased. In this set of test cases, the multiplier value was left at its default value – 3 which is also the lowest possible value that can be configured. At the default hello interval of 10 seconds, the mean failover time and restoration time were about 12 and 18 seconds respectively. At the extreme case where the hello interval was configured at 1 second, the mean failover time was around 8 seconds and the mean restoration time was about 5 seconds.




Figure 19. Mean Failover and Restoration Time with respect to Hello interval

In the absence of the line protocol signaling, the routers would rely on the ISIS hello packets to detect its neighbors. As such, fast hellos with a low hold time would allow a quick network convergence, resulting in shorter failover and restoration time. However, setting the hello interval too low might produce counter-effects since it might generate unnecessary amount of control packets, thus overloading the network traffic.

4. Varying TDP Discovery Hello Intervals

Table 26 and 27 show the statistical results on the failover time and restoration time for Test Cases 19 - 24.

	Mean	Std Dev	Range
Test Case 19	5.597s	0.004s	5.590 – 5.606s
Test Case 20	5.595s	0.014s	5.561 – 5.608s
Test Case 21	5.600s	0.005s	5.589 – 5.609s

Test Case 22	5.599s	0.007s	5.584 – 5.613s
Test Case 23	5.599s	0.006s	5.590 – 5.609s
Test Case 24	5.600s	0.008s	5.587 – 5.613s

Table 29. Statistical Results on Failover Time for Test Cases 19 – 24

	Mean	Std Dev	Range
Test Case 14	5.224s	2.080s	2.823 – 8.165s
Test Case 15	6.150s	2.963s	2.963 – 8.676s
Test Case 16	7.354s	1.446s	3.268 – 8.663s
Test Case 17	9.978s	1.922s	8.539 – 15.207s
Test Case 18	13.796s	0.911s	12.463 – 15.373s
Test Case 19	18.753s	1.595s	15.863 – 22.868s

Table 30. Statistical Results on Restoration Time for Test Case 19 – 24

Figure 20 below shows the mean failover time and restoration time with respect to the TDP discovery hello interval. From the graph, it can be seen that the TDP discovery hello interval did not affect the failover time. Regardless of the TDP discovery hello interval values, which ranged from 1 to 20 seconds, the failover time remained around 6 seconds. On the contrary, different TDP discovery hello intervals resulted in different restoration times. The time required to switch the traffic back to the primary LSP from the backup LSP increased as the TDP discovery hello interval increased. At the extreme case when the TDP discovery hello interval was set at 20 seconds with the holdtime interval configured, as recommended, as three times the value of the hello interval, the restoration time was more than 18 seconds. However, the headend router took only 5 seconds to switch over the traffic when the interval was set at 1 second.




Figure 20. Mean Failover and Restoration Time with respect to TDP Discovery Hello Interval

For label retention, Cisco routers adopt the liberal label retention mode, where the LSR will retain all the bindings received from its TDP peers. When PE1 received the label bindings generated for PE2 from P1 and P2, it would retain both bindings in its MPLS forwarding table. The figure below shows a sample of the details in a MPLS forwarding table of PE1

```
PE1#show mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	10.1.1.8/30	0	Et0/0	10.1.1.6
17	Pop tag	10.10.10.3/32	619	Et0/0	10.1.1.6
18	Aggregate	200.0.4.0/24[V]	10090		
19	Aggregate	200.0.4.0/24[V]	1776		
20	Pop tag	10.1.1.20/30	0	Et0/1	10.1.1.13
21	17	10.10.10.6/32	597	Et0/0	10.1.1.6
	19	10.10.10.6/32	0	Et0/1	10.1.1.13
23	Pop tag	10.10.10.1/32	615	Et0/1	10.1.1.13

Figure 21. Sample Details in MPLS Forwarding Table

The local tag number 21 in the MPLS forwarding table for PE1 was for destination PE2, which had the IP address of 10.10.10.6. There were two outgoing tags, tag 17 and 19, bound to this local tag. These two tags were provided by the two P routers connected to it. In the event a link or node failure occurred resulting in one of the outgoing tags no longer being valid, there would be another tag immediately available for use. Hence, this explained why the different settings of the discovery hello interval had no impact on the failover time in such a case. However, the failover time would have increased if the conservative label retention mode was used.

Although the TDP discovery hello intervals had no impact on the failover time, it did affect the restoration time. This is because when a link or a node recovers from its failure, there will be time incurred to re-establish the TDP adjacency between the LSRs. This is determined by the discovery hello interval. Hence, as the hello interval gets larger, the time to re-establish the TDP adjacency between PE1 and the other LSRs will get longer, which in turn will cause the time traffic to switch over to the new established LSP to be longer.

5. Varying MPLS TE Tunnel Configuration Options

Table 28 and 29 show the statistical results on the failover time and restoration time for Test Cases 25 - 28.

	Mean	Std Dev	Range
Test Case 25	5.608s	0.017s	5.592 – 5.668s
Test Case 26	5.902s	1.774s	2.765 – 9.700s
Test Case 27	4.956s	1.809s	2.351 – 11.461s
Test Case 28	0.610s	0.010s	0.593 – 0.625s

Table 31. Statistical Results on Failover Time for Test Cases 25 – 28

	Mean	Std Dev	Range
Test Case 25	7.424s	0.974s	5.235 – 8.886s
Test Case 26	-	-	-
Test Case 27	7.354s	1.446s	3.268 – 8.663s

Test Case 28	5.978s	0.306s	5.431 – 6.457s
---------------------	--------	--------	----------------

Table 32. Statistical Results on Restoration Time for Test Cases 25 - 28

A comparison between the different MPLS TE tunnel configuration options was made to determine their impact on the failover functionality. Figure 22 below shows the comparison of failover and restoration time among different MPLS TE tunnel configuration options. The graph did not showed conclusive results in regard to the impact the different TE tunnel configurations had on the failover and restoration time.




Figure 22. Comparison of Failover and Restoration Time among different TE tunnel configuration options

In Test Case 25, two tunnels, denoting the primary and the backup LSPs, were configured. The two paths were pre-established such that resources, like the bandwidth, were reserved for both paths before any traffic flow took place. The failover and restoration times were very close to the ones that were configured without tunnels. In this case, there were additional resources utilized as a result of the additional tunnel although there were no clear advantages in the improvement of the failover time and restoration times.

In Test Case 26 and 27, only one tunnel was configured. For the former test case, there were two explicitly configured path options: the first path option via P1 and the other via P2. In the event of a link or node failure, the traffic would be redirected to the second path option. However, one observation made was the traffic was not rerouted back to the first path option in the event of a link or node recovery. This explained why there was no result for restoration time in Test Case 26. In Test Case 27, there was only one path option which was determined based on the TE calculation module. The module would establish a backup path after a link or node failure had occurred.

In both of these test cases, the mean failover time was about 5 and 6 seconds respectively, similar to the result for Test Case 3 and 25. However, it was observed that from the set of readings taken for the two test cases, the variation among the readings were quite high. There were instances for both test cases where the failover time could reach as low as around 2 seconds, unlike the Test Case 3 and 25 where the failover timings were consistently around 5.6 seconds.

The Cisco MPLS Autoroute Announce feature was used for this experiments. The router command to enable this feature was illustrated in Chapter 3. This feature specifies the IGP, in this case ISIS, to use the tunnel (provided that the tunnel is up) in its enhanced shortest path first calculation. Currently, the only way to forward traffic onto a tunnel is by enabling this feature or by explicitly configuring forwarding, such as using an interface static route [24]. The figure below shows the MPLS-TE system block diagram of a headend router where the tunnels were introduced into the IGP shortest path calculation.




Figure 23. MPLS-TE system block diagram with tunnels introduced into IGP (After Ref. [7].)

As such, the TE tunnels were treated exactly like any normal IGP routes in the SPF calculation. The tunnels can also be configured with different metric values to give more or less preference over the other IGP routes in the shortest path selection. Hence, the behavior of the LSP tunnels in regard to the BGP/MPLS VPN failover functionality would be expected to be the same as the other (non tunneled) LSPs.

A comparison was made among Test Case 3, 9, 25 and 28 to determine the impact of the spf-initial-wait intervals on the failover and restoration time for both tunnel and non-tunnel enabled scenarios. For the case of the TE-enabled scenarios - Test Case 25 and 28, two tunnels were pre-established but the spf-initial-wait interval was different from each. For the non-TE enabled scenarios - Test 3 and 9, there were no pre-established tunnels. The spf-initial-wait-interval between the two test cases was different as well. Figure 24 shows the mean failover time and restoration time among the selected test cases.




Figure 24. Comparison of Failover and Restoration Time based on tunnels being established and different spf-initial-wait intervals.

The diagram above indicates the indifference of the failover and recovery performance behaviors between TE enabled and non-TE enabled scenarios. In the case where the spf-initial-wait interval was configured at 500 milliseconds, Test Cases 9 and 28, the failover time for either was about 600 milliseconds. The period of time required to switch over the traffic from the primary tunnel to the backup tunnel was about the same as the time required for a “normal” primary LSP to reroute to a backup LSP.

6. Varying Static and Non-Static Routing Configuration

Table 33 and 34 show the statistical results on the failover time and restoration time for Test Case 29 - 32.

	Mean	Std Dev	Range
Test Case 29	0.689s	0.482s	0.030 – 1.031s

Test Case 30	1.311s	0.276s	1.016 - 1.813s
Test Case 31	NA	NA	NA
Test Case 32	1.876s	0.413s	1.105 – 2.545s

Table 33. Statistical Results on Failover Time for Test Cases 29 – 32

	Mean	Std Dev	Range
Test Case 29	6.931s	1.168s	5.719 – 8.804s
Test Case 30	5.182s	0.411s	4.437 – 5.707s
Test Case 31	NA	NA	NA
Test Case 32	-	-	-

Table 34. Statistical Results on Restoration Time for Test Cases 29 - 32

A comparison was made between the results derived from Test Cases 29 – 32 and the results derived from Test Case 1 - 4 to determine the impact that each routing configuration had on the failover functionality. Given that we have already know how the failover time and restoration time fare in the case of using a IGP where the LSPs are of equal cost, we would also want to investigate how much the performance would differ if static routing is used instead. Figure 22 below shows the comparison of mean failover time between the static and non-static routing configuration.

From the graph, there is no significant difference in the mean failover time between the static routing and the ISIS at the various failure locations except for the link failure that is located at the interface of P1 connecting to PE2. When the link failure occurred at the headend router, both type of routing configurations had the failover operation completed below one second. However, in the exceptional case, Test Case 31, there is no failover performed by the headend router. The midpoint router did not provide any signaling to the headend router to inform regarding the link failure. As such, the headend router continued to send the traffic to the primary path which had been broken where the packets were dropped at the midpoint router.




Figure 25. Comparison of Mean Failover Time between Static and Non-Static Routing Configuration.

Except for Test Case 31 where there was no restoration time collected as a result of no failover capability, the mean restoration time between the static routing configuration and the non-static routing configuration were very close, similar to the mean failover time. The mean restoration time for the link failure that occurred at the headend router was about seven to eight seconds for both routing configuration types. For link failure that occurred at the midpoint router in which the link interface was adjacent to the headend router, the mean restoration time for both types were about five to six seconds.

The results for both mean failover and restoration time showed that static routing would also provide the same BGP/MPLS VPN failover performance as compared to one based on IGP routing when the link failure occurred at or next to the headend router. However, if the link failure occurs away from the headend router where there is no signaling of failure to the headend router, the traffic would be disrupted and service availability to the customer would be affected.




Figure 26. Comparison of Mean Restoration Time between Static and Non-Static Routing Configuration.

C. SUMMARY

In this chapter, the statistical results derived from the data collected were presented. A detailed analysis of the results was conducted and discussed. The results had identified some of the attributes in each of the identified key components of the BGP/MPLS VPN that had affected the time delay of a network failure or recovery. The appropriate setting of a lower ISIS SPF interval and hello interval would reap a shorter time delay in the network failover in the event of a link or node failure. The configuration of multiple LSP/LSP tunnels to the same destination with equal link cost would also allow a lower failover time. As for the restoration time, a low TDP Hello interval, a low ISIS SPF interval or a low ISIS Hello interval would allow the traffic to redirect back to the original path in a shorter span of time.

The following chapter summarizes and concludes the thesis research. It discusses the research areas that have not been explored due to the lack of time and resources and provides recommendation for further research in related areas.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION AND FURTHER RESEARCH AREAS

A. CHAPTER OVERVIEW

This Chapter presents the main conclusions drawn from the performance analysis described in Chapter IV. Possible extensions to the thesis study are also proposed.

B. CONCLUSION

This thesis experimentally examined some of the key parameters in the BGP/MPLS VPN protocol to identify the contributing factors to the latency of a network failover or recovery event and further determine the possible router configuration options that could be used to reduce the failover time and restoration delays. These parameters include the ISIS distance metric, the ISIS Hello interval, the ISIS SPF interval, the TDP Hello interval and the various MPLS TE tunnel configuration options.

A BGP/MPLS VPN network consisting of four Cisco label switched routers were set up to facilitate the experimentation. Additional routers and workstations were also deployed to represent customers' networks. To facilitate raw data collection, software tools were used to generate traffic, capture packets and monitor traffic. 32 test cases were conducted to examine some of the parameters of the identified key components of the BGP/MPLS VPN. Basic statistical measures were generated based on the readings collected from each test case and a detailed statistical analysis was then performed on the results.

The conclusions drawn from the statistical analysis are as follows:

- o A symmetrical and a “non-symmetrical” topology were constructed by means of manipulating the ISIS metric and were compared in terms of the failover and restoration time. The failover time was faster for a symmetrical network than a “non-symmetrical” network when the link failure occurred near the headend router. This is because there is more than one routing entry in the routing table in the headend router

for the same prefix as a result of configuring multiple LSPs with equal link cost. As such, the headend router could switch the traffic over in a shorter period of time. Hence, configuring multiple LSPs with equal link cost in the headend router, not only allows load balancing of the network traffic, it also allows faster network failover in the event of a network link or node failure.

- o The ISIS SPF Interval is the time period the router running ISIS would hold on the SPF computation after detecting a network topological change. The results had shown that a shorter SPF interval would trigger the router to react faster to a network topological change. As such, we can manipulate the value of this IGP attribute to allow a quick failover in the event of a link or node failure and a quick restoration to the primary path from the backup path in the event of a link or node recovery. Nonetheless, oversetting of this attribute might bring counter effects, destabilizing the network.
- o The ISIS Hello packets are used to initialize and maintain adjacencies between neighboring routers. The ISIS hello interval specifies the length of time between hello packets that the router sends. Like the ISIS SPF interval, a shorter hello interval would result in a quicker failover and recovery as a result of detecting the topological change faster. However, this is only true in the case where the fault detection relies on the IGP signaling. A fast hello would not result in a faster convergence time if the fault detection was provided by the line protocol since the line protocol alerts are typically faster than the ISIS Hello interval.
- o The routers used the TDP discovery mechanism to detect its TDP neighbors. The TDP hello interval is the period of time between the sending of consecutive Hello messages. The TDP Hello interval does not impact the failover time. This is because Cisco routers adopt the liberal label retention mode and as such, there might be more than one

label available for a particular prefix. In the event of a link or a node failure, the router does not have to re-establish the label distribution session to obtain a new label. Instead, the additional label in the MPLS forwarding table would allow the traffic to be redirected immediately to the next available LSP. However, a shorter TDP Hello interval would result in a shorter restoration time. As such, we can adjust a shorter TDP Hello interval to allow the traffic to be redirected back to the primary path in a shorter frame of time.

- o A comparison between the different MPLS TE tunnel configuration options was made to determine their impact on the failover functionality. There were no conclusive results in regard to the impact the different TE tunnel configurations had on the failover and restoration time. However, the behavior of the tunnels in terms of the failover functionality was similar to that of normal LSPs where the attributes like the SPF interval affected both LSPs and the LSP tunnels.
- o A comparison between the static and non-static routing configuration was made to determine their impact on the failover functionality. Static routing provides similar failover performance as compared to one based on IGP routing when the link failure occurred at or next to the headend router. However, it was observed that if the link failure occurs away from the headend router where there is no signaling of failure to the headend router, service availability to the customer would be resulted.

In essence, this thesis study had identified some of the parameters in the BGP/MPLS VPN protocol that would affect the latency of a network failover or recovery event and determine the router configuration options that could be used to reduce the failover and restoration delays. The experimentation had shown that the failover time can be achieved below one second by configuring the ISIS SPF interval to be below one second. Assuming that the routers have the high

processing capacity to process the additional routing load, the network bandwidth has the spare resources to accommodate the additional traffic generated and the network is very much stable and does not experience very frequent topological changes, the SPF interval should be configured as low as below one second to achieve a rapid failover. The ISIS hello interval can also be used as a backup signaling capability to assist the detection of a network failure to help to improve the failover time in the event where the line protocol signaling is absent. The shortest possible hello interval of 1 second and holdtime of 3 seconds was experimented and an improvement of 4 seconds was achieved as compared to one configured with the default hello interval value of 10 seconds. In addition, by configuring multiple LSPs with the same link cost metric, using dynamic or static routing can also help to reduce the failover time. Sub-second failover was achieved in the experiment for a link failure that occurred at the headend router.

To achieve a short restoration delay, the ISIS SPF interval and the TDP Hello interval can be configured to improve the latency. A 5 second restoration delay was achieved in the experimentation when the SPF interval was configured below one second or the TDP Hello interval was configured at its minimum value of 1 second as compared to 8 seconds from the default settings.

In the past where hardware posed some limitations, the configuration options mentioned above had to be done cautiously to avoid destabilizing the network. However, given that the routers nowadays come with high processing and memory capabilities, the findings from this thesis research should provide insights on what can be configured to reduce the failover and restoration latency and how best can it be achieved.

C. FURTHER RESEARCH AREAS

This work is far from determining all the contributing factors to the latency of a network failure or recovery event. A number of possible future research areas have been identified as follows.

1. Examining additional parameters of BGP/MPLS VPN

Due to the lack of resources and time, this study only covered a set of parameters deemed most important. Other parameters might contribute to the time delay of a network failure and recovery in an unanticipated way. Some of these parameters include the RSVP signaling in the case where LSP tunnels are deployed, other ISIS intervals such as the partial route calculation (PRC) interval etc or OSPF parameters, in the case if OSPF is chosen as the IGP instead. One experiment worth performing is to set all IGP timing parameters to their minimum value in the nonsymmetrical topology used for test cases 5-8 and then compare the results against those of the static routes (Test Cases 29-32). The performance gap will represent a lower bound on the latency incurred by dynamic routing.

2. Expanding the size of the laboratory network set up

The laboratory network backbone set up for this study was based on a simple deployment of four label switched routers. The number of available LSPs from the ingress to the egress router was very limited. In addition, the locations of the link or node failure were also constrained by the size of the network backbone. As such, with the increase in the number of nodes in the MPLS network backbone, we can examine how the existing identified factors or even other contributing factors affect the failover and recovery performance as the complexity of the network topology increases.

3. Examining the MPLS Fast Reroute

The scope of the experimentation only covers the IP reroute and end-to-end protection mechanisms. Till date, it is documented that the MPLS Fast Reroute protection mechanism offers the fastest failover and recovery capability as compared to the other protection mechanisms. It would be beneficial to determine how the contributing factors identified in this thesis research on the time delay of a network failure and recovery play a part in the MPLS Fast Reroute protection mechanism.

4. Examining the prioritization of multiple VPNs

Since MPLS allows the complete logical separation of network traffic and routing information by means of multiple VPNs, it is also useful to found out if the MPLS network backbone supports the prioritization of network traffic based on different VPNs such that in the event of a network component failure, the network traffic of a higher prioritized VPN would be preferred. And if the prioritization is supported, we would also like to examine if the preference can be set dynamically, .i.e. without requiring the network administrator to manually configure to preempt one network from another in the event of resource competition. And also, to determine the smallest granularity of traffic that can be given preference.

LIST OF REFERENCES

1. U.S. Department of Defense. 2002. DoD Directive 8100.1. Global Information Grid (GIG) Overarching Policy. Washington, D.C.
2. Juniper Networks. "Juniper Networks Wins IP Routing for GIG-BE, Global Information Grid-Bandwidth Expansion."
<http://www.juniper.net/company/presscenter/pr/2003/pr-031230.html>, last accessed on 20 Oct 2006
3. MPLS Resource Center. "MPLS FAQ",
<http://www.mplsrc.com/mplsfaq.shtml>, last accessed on 21 Oct 2006
4. IXIA. "Multi-Protocol Label Switching (MPLS) Conformance and Performance Testing."
http://www.ixiacom.com/library/white_papers/display?skey=mpls, last accessed on 21 Oct 2006
5. Rosen, E. Viswanathan, A. Callon, R. "Multiprotocol Label Switching Architecture", RFC 3031, The Internet Society, January 2001.
<http://www.ietf.org/rfc/rfc3031.txt>, last accessed on 21 Oct 2006
6. Lewis, M. *Troubleshooting Virtual Private Networks (VPN)*, Edition 1, Cisco Press, 2001.
<http://www.ciscopress.com/articles/article.asp?p=391649&seqNum=1&rl=1>, last accessed on 21 Oct 2006
7. London's Global University. "MPLS - Label Switching".
http://www.hep.ucl.ac.uk/~ytl/qos/mpls_01.html, last accessed on 21 Oct 2006
8. Xie, G. Notes for CS4550 (Computer Networks II), Naval Postgraduate School, 2006 (unpublished)
9. MPLS Resource Centre. "Advanced MPLS".
<http://www.mplsrc.com/WhitePapers/advancedmpls3.pdf>, last accessed on 22 Oct 2006
10. Nortel Networks. "MPLS Tutorial and Operational Experiences".
<http://www.nanog.org/mtg-9905/ppt/mpls.ppt>. last accessed on 22 Oct 2006
11. Cisco. "MPLS Virtual Private Networks".
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/vpn.pdf>, last accessed on 22 Oct 2006

12. Wu, T. "MPLS VPNs - Layer 2 or Layer 3, Understanding the Choice". Riverstone Networks.
http://www.riverstonenet.com/pdf/mpls_vpns_layer2_or_layer3.pdf, last accessed on 23 Oct 2006
13. Rosen, T. Rekhter, T. "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, The Internet Society, February 2006.
<http://www.ietf.org/rfc/rfc4364.txt>, last accessed on 23 Oct 2006
14. Lakshman, T. Lobo, L. *MPLS Configuration on CISCO IOS Software*, Edition 1, Cisco Press, 2005.
<http://www.ciscopress.com/articles/article.asp?p=426640&seqNum=1&rl=1>, last accessed on 23 Oct 2006
15. Li, Y. Panwar, S. Liu, C.J. "Performance Analysis of MPLS TE Queues for QoS Routing", http://www.cie-gnyc.org/convention_2003/poster/li.pdf, last accessed on 24 Oct 2006
16. Menth, M. Martin, R. Hartmann, M. Spoerlein, U. "Efficiency of Routing and Resilience Mechanisms". University of Wuerzburg. Germany
17. Huang, C. Messier, D. "A Fast and Scalable Inter-Domain MPLS Protection Mechanism". <http://www.sce.carleton.ca/faculty/huang/R-3-114.pdf>, last accessed on 24 Oct 2006
18. Sharma, V. Hellstrand, F. "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, The Internet Society, February 2003. <http://www.faqs.org/ftp/rfc/pdf/rfc3469.txt.pdf>, last accessed on 24 Oct 2006
19. Poretsky, S. Papneja, R. Rao, S. Roux, J-L. L. "Benchmarking Methodology for MPLS Protection Mechanisms", Internet-Draft, <http://tools.ietf.org/wg/mpls/draft-papneja-mpls-protection-meth-merge-00.txt>, last accessed on 24 Oct 2006
20. Paquet, C. Teare, D. *Building Scalable Cisco Networks*, Edition 1, Cisco Press, 2000.
<http://www.ciscopress.com/articles/article.asp?p=31319&seqNum=4&rl=1>, last accessed on 9 Nov 2006
21. Avici Systems. "Integrated IS-IS".
http://www.avici.com/documentation/HTMLDocs/03675-02_revBA/ISIS.html, last accessed on 10 Nov 2006
22. White, R. Retana, A. *IS-IS: Deployment in IP Networks*, Edition 1, Addison Wesley Professional, 2003.
<http://www.informit.com/articles/article.asp?p=30305&seqNum=2&rl=1>, last accessed on 10 Nov 2006

23. TECHFAQ. "What is a static route?". <http://www.tech-faq.com/static-route.shtml>, last accessed on 10 Nov 2006
24. Cisco. "MPLS Traffic Engineering and Enhancements". http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801039b6.html#wp1020161, last accessed on 14 Nov 2006

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Professor Geoffrey Xie
Naval Postgraduate School
Monterey, California
4. Professor John Gibson
Naval Postgraduate School
Monterey, California
5. Yeo Tat Soon
Director, Temasek Defence Systems Institute (TDSI)
National University of Singapore
Singapore
6. Tan Guan Chye
Defence Science & Technology Agency
Singapore