

**Moonv6 Test Suite**  
*Routing Transition Mechanism*  
*Interoperability Test Suite*

**DRAFT**

**Technical Document**

Revision 0.1



---

*IPv6 Consortium*  
*InterOperability Laboratory*  
*Research Computing Center*  
*University of New Hampshire*

*121 Technology Drive, Suite 2*  
*Durham, NH 03824-3525*  
*Phone: (603) 862-2804*  
*Fax: (603) 862-4181*  
*<http://www.iol.unh.edu>*

## TABLE OF CONTENTS

ACKNOWLEDGEMENTS .....	4
INTRODUCTION .....	5
TEST ORGANIZATION.....	6
REFERENCES .....	7
GROUP 1: OSPF Functionality with Tunnels.....	8
Test RTM.1.1: OSPFv3 Convergence with Static IPv6 Tunnels .....	9
Test RTM.1.2: OSPFv3 Convergence with 6over4 IPv6 Tunnels .....	11
Test RTM.1.3: Multi-Area OSPFv3 Convergence with Static IPv6 Tunnels .....	13
Test RTM.1.4: Multi-Area OSPFv3 Convergence with 6over4 IPv6 Tunnels .....	15
GROUP 2: BGP Functionality with Tunnels .....	17
Test RTM.2.1: Static IPv6 Tunnel BGP-4+ Convergence .....	18
Test RTM.2.2: 6over4 Tunnel BGP-4+ Convergence .....	20
Test RTM.2.3: BGP-4 and BGP-4+ with IPv4 and IPv6 Transport.....	22

**MODIFICATION RECORD**  
Draft Version Complete

September 29, 2004

## ACKNOWLEDGEMENTS

**The University of New Hampshire would like to acknowledge the efforts of the following individuals in the development of this test suite. This test suite is written for the Moonv6 Project and is an incremental development of the Routing Convergence test suite. The people who significantly contributed to the Routing Convergence test suite are listed below in addition to the authors of this plan.**

Eric Barrett	University of New Hampshire
Don Heidrich	Cisco Systems
Yasuyuki Matsuoka	NTT Network Systems Service Laboratories
Kari Revier	University of New Hampshire
Kazunori Sakumoto	Fujitsu Ltd.
Benjamin Schultz	University of New Hampshire

# INTRODUCTION

## Overview

The University of New Hampshire's InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards-based products by providing an environment where a product can be tested against other implementations of a standard. This suite of tests has been developed to help implementers evaluate the functioning of their Internet Protocol, version 6 capable products. The tests do not determine if a product conforms to the IPv6 routing specifications, nor are they purely interoperability tests. Rather, they provide one method to isolate problems within a device. Successful completion of all tests contained in this suite does not guarantee that the tested device will interoperate with other IPv6 devices. However, combined with satisfactory operation in the IOL's semi-production environment, these tests provide a reasonable level of confidence that the Router Under Test will function well in many multi-vendor IPv6 environments.

The definition of "convergence" for this test suite means that the traffic flow from ingress to egress is restored after the primary link that the traffic flow was upon is altered due to a change in topology.

The tests in this suite are written in a generic manner that does not define the specific speed of different link types. When traffic is passed through the network to measure convergence times, it is assumed that the network will converge from the time the first egress data packet is dropped until the next egress data packet is received. It should be remembered that the faster traffic is passed through the network, the higher resolution of convergence time will be produced. These tests will be optimal if performed on the same link type, but it is not a requirement.

## Acronyms

<b>ABR:</b> Area Border Router	<b>AS:</b> Autonomous System	<b>MED:</b> Multi-Exit Discriminator
<b>RD:</b> Route Distinguisher	<b>RT:</b> Route Target	
<b>RUT:</b> Router Under Test	<b>TR:</b> Testing Router	<b>G:</b> Traffic Generator

When several entities of the same type are present in a test configuration, a number is appended to the acronym to yield a label for each entity. For example, if there were three testing routers in the test configuration, they would be labeled TR1, TR2 and TR3. The traffic Generators can also perform routing emulation and inject routes into the TR routing tables.

## TEST ORGANIZATION

This document organizes tests by group based on related test methodology or goals. Each group begins with a brief set of comments pertaining to all tests within that group. This is followed by a series of description blocks; each block describes a single test. The format of the description block is as follows:

Test Label:	The test label and title comprise the first line of the test block. The test label is composed by concatenating the short test suite name, the group number, and the test number within the group, separated by periods. The Test Number is the group and test number, also separated by a period. So, test label RTM.1.2 refers to the second test of the first test group in the routing transition mechanism test suite. The test number is 1.2.
Purpose:	The Purpose is a short statement describing what the test attempts to achieve. It is usually phrased as a simple assertion of the feature or capability to be tested.
References:	The References section lists cross-references to the specifications and documentation that might be helpful in understanding and evaluating the test and results.
Resource Requirements:	The Resource Requirements section specifies the software, hardware, and test equipment that will be needed to perform the test.
Discussion:	The Discussion is a general discussion of the test and relevant section of the specification, including any assumptions made in the design or implementation of the test as well as known limitations.
Test Setup:	The Test Setup section describes the configuration of all devices prior to the start of the test. Different parts of the procedure may involve configuration steps that deviate from what is given in the test setup. If a value is not provided for a protocol parameter, then the protocol's default is used for that parameter.
Procedure:	This section of the test description contains the step-by-step instructions for carrying out the test. These steps include such things as enabling interfaces, unplugging devices from the network, or sending packet from a test station. The test procedure also cues the tester to make observations, which are interpreted in accordance with the observable results given for that test part.
Observable Results:	This section lists observable results that can be examined by the tester to verify that the RUT is operating properly. When multiple observable results are possible, this section provides a short discussion on how to interpret them. The determination of a pass or fail for each test is usually based on how the RUT's behavior compares to the results described in this section.
Possible Problems:	This section contains a description of known issues with the test procedure, which may affect test results in certain situations.

## **REFERENCES**

The following documents are referenced in this text:

## **GROUP 1: OSPF Functionality with Tunnels**

### **Scope:**

These following tests are designed to verify basic interoperability of OSPF in an IPv6 Transition environment.

### **Overview:**

OSPFv2 and OSPFv3 have been tested in a multi-vendor environment at a very basic level with both convergence and interoperability being demonstrated. It is still an unknown to see how the OSPF routing algorithm will run over various IPv6 Tunnels that span an IPv4-only network. While it is an ideal that native IPv6 will be everywhere, it is a necessity to transport IPv6 over IPv4. It is a network requirement to demonstrate that routing functionality is stable in this environment.

## Test RTM.1.1: OSPFv3 Convergence with Static IPv6 Tunnels

**Purpose:** To verify that OSPF convergence can occur when there are static IPv6 tunnels in the network.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 27, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. This should remain consistent when there are static IPv6 tunnels in the network.

**Test Setup:** Connect the network as described in the figure below. Enable OSPFv2 and OSPFv3, set all link costs to 4 (including those for the IPv6 tunnels) and all routers to be in area 0. The links that G1 and G2 are on, in addition to Link 2 and Link 4 are IPv6 enabled. Links 1, 3, 5, 6, 7 and 8 are IPv4 only. Create a full mesh of static IPv6 tunnels over all links except Links 2 and 4. Also create static IPv6 tunnels from TR1-TR3 and TR4 to TR2.



### Procedure:

#### Part A: Link Cost Change

1. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
2. Pass continuous IPv6 traffic from G3 to G4 and from G4 to G3.
3. On TR1 and TR4, change the cost from 4 to 40 for Link 1.
4. Allow the network to converge and observe the egress data traffic.
5. On TR1 and TR4, change the cost from 40 to 4 for Link 1.
6. Allow the network to converge and observe the egress data traffic.
7. On TR2 and TR3, change the cost from 4 to 40 for Link 1.
8. Allow the network to converge and observe the egress data traffic.
9. On TR2 and TR3, change the cost from 40 to 4 for Link 1.
10. Allow the network to converge and observe the egress data traffic.

#### Part B: Link Down

11. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
12. Pass continuous IPv6 traffic from G3 to G4 and from G4 to G3.
13. Disconnect Link 1.
14. Allow the network to converge and observe the egress data traffic.
15. Reconnect Link 1.
16. Allow the network to converge and observe the egress data traffic.
17. Disconnect Link 3.

18. Allow the network to converge and observe the egress data traffic.
19. Reconnect Link 3.
20. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, when the metric on the tunnel over Link 1 is set to 40, IPv6 traffic flow changes from (TR1-TR4) to (TR1-TR5-TR4). When the cost is restored, traffic flow is restored to its original path (TR1-TR4). When the metric on the tunnel over Link 3 is set to 40, IPv6 traffic flow changes from (TR2-TR3) to (TR2-TR5-TR3). When the cost is restored, traffic flow is restored to its original path (TR2-TR3).
- In Part B, when Link 1 is disconnected, IPv6 traffic flow changes from (TR1-TR4) to (TR1-TR5-TR4). When link 1 is reconnected, traffic flow is restored to its original path (TR1-TR4). When Link 3 is disconnected, IPv6 traffic flow changes from (TR2-TR3) to (TR2-TR5-TR3). When link 3 is reconnected, traffic flow is restored to its original path (TR2-TR3).

**Possible Problems:** None.

## Test RTM.1.2: OSPFv3 Convergence with 6over4 IPv6 Tunnels

**Purpose:** To verify that OSPF convergence can occur when there are 6over4 IPv6 tunnels in the network.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 27, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. This should remain consistent when there are static IPv6 tunnels in the network.

**Test Setup:** Connect the network as described in the figure below. Enable OSPFv2 and OSPFv3, set all link costs to 4 (including those for the IPv6 tunnels) and all routers to be in area 0. The links that G1 and G2 are on, in addition to Link 2 and Link 4 are IPv6 enabled. Links 1, 3, 5, 6, 7 and 8 are IPv4 only. Create a full mesh of 6over4 IPv6 tunnels over all links except Links 2 and 4. Also create static IPv6 tunnels from TR1-TR3 and TR4 to TR2.



### Procedure:

#### Part A: Link Cost Change

1. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
2. Pass continuous IPv6 traffic from G3 to G4 and from G4 to G3.
3. On TR1 and TR4, change the cost from 4 to 40 for Link 1.
4. Allow the network to converge and observe the egress data traffic.
5. On TR1 and TR4, change the cost from 40 to 4 for Link 1.
6. Allow the network to converge and observe the egress data traffic.
7. On TR2 and TR3, change the cost from 4 to 40 for Link 1.
8. Allow the network to converge and observe the egress data traffic.
9. On TR2 and TR3, change the cost from 40 to 4 for Link 1.
10. Allow the network to converge and observe the egress data traffic.

#### Part B: Link Down

11. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
12. Pass continuous IPv6 traffic from G3 to G4 and from G4 to G3.
13. Disconnect Link 1.
14. Allow the network to converge and observe the egress data traffic.
15. Reconnect Link 1.
16. Allow the network to converge and observe the egress data traffic.
17. Disconnect Link 3.

18. Allow the network to converge and observe the egress data traffic.
19. Reconnect Link 3.
20. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, when the metric on the tunnel over Link 1 is set to 40, IPv6 traffic flow changes from (TR1-TR4) to (TR1-TR5-TR4). When the cost is restored, traffic flow is restored to its original path (TR1-TR4). When the metric on the tunnel over Link 3 is set to 40, IPv6 traffic flow changes from (TR2-TR3) to (TR2-TR5-TR3). When the cost is restored, traffic flow is restored to its original path (TR2-TR3).
- In Part B, when Link 1 is disconnected, IPv6 traffic flow changes from (TR1-TR4) to (TR1-TR5-TR4). When link 1 is reconnected, traffic flow is restored to its original path (TR1-TR4). When Link 3 is disconnected, IPv6 traffic flow changes from (TR2-TR3) to (TR2-TR5-TR3). When link 3 is reconnected, traffic flow is restored to its original path (TR2-TR3).

**Possible Problems:** None.

### Test RTM.1.3: Multi-Area OSPFv3 Convergence with Static IPv6 Tunnels

**Purpose:** To verify that OSPF convergence can occur with static IPv6 Tunnels.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 29, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. OSPFv2 and OSPFv3 are two separate protocols that route IPv4 and IPv6 traffic respectively. These protocols use separate software processes and should operate like “ships in the night”, in that they do not interfere with each other.

**Test Setup:** Connect the network as described in the figure below. Enable OSPF, set all link costs to 4 for OSPFv2 and OSPFv3. Configure TR1 to be in area 0. Configure TR3 and TR4 to be in area 1. Configure TR2 and TR5 to be ABRs between area 0 and area 1. Disable IPv6 on Links 2, 3, 5 and 6 so only IPv4 is supported. Create static IPv6 tunnels TR1-TR4 and TR1-TR3. It is recommended that this test is run with a larger number of routers and tunnels to demonstrate the functionality of static IPv6 tunnels. This second run should also use G1 and G2 to inject a realistic number of routes into this network. The minimum topology is covered here to allow laboratories with limited resources to reproduce this test.



**Procedure:**

*Part A: Link Down, same topologies*

1. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.
2. Disconnect Link 5.
3. Allow the network to converge and observe the egress data traffic.
4. Reconnect Link 5.
5. Allow the network to converge and observe the egress data traffic.
6. Disconnect Link 6.
7. Allow the network to converge and observe the egress data traffic.
8. Reconnect Link 6.
9. Allow the network to converge and observe the egress data traffic.
10. Disconnect Link 2.
11. Allow the network to converge and observe the egress data traffic.
12. Reconnect Link 2.

13. Allow the network to converge and observe the egress data traffic.

*Part B: Link Cost Change, same topologies*

14. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.

15. On TR5's interfaces on Link 5 and Link 6, change the cost from 4 to 40 for both OSPFv2 and OSPFv3.

16. Allow the network to converge and observe the egress data traffic.

17. On TR5's interfaces on link 5 and Link 6, change the cost from 40 to 4 for both OSPFv2 and OSPFv3.

18. Allow the network to converge and observe the egress data traffic.

19. On TR2's interfaces on Link 2 and Link 3, change the cost from 4 to 40 for both OSPFv2 and OSPFv3.

20. Allow the network to converge and observe the egress data traffic.

21. On TR2's interfaces on Link 2 and Link 3, change the cost from 40 to 4 for both OSPFv2 and OSPFv3.

22. Allow the network to converge and observe the egress data traffic.

*Part C: Link Cost Change, different topologies*

23. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.

24. On TR5's interfaces on Link 5 and Link 6, change the cost from 4 to 40 for OSPFv2.

25. Allow the network to converge and observe the egress data traffic.

26. On TR5's interfaces on link 5 and Link 6, change the cost from 40 to 4 for OSPFv2.

27. Allow the network to converge and observe the egress data traffic.

28. On TR2's interfaces on Link 2 and Link 3, change the cost from 4 to 40 for OSPFv3.

29. Allow the network to converge and observe the egress data traffic.

30. On TR2's interfaces on Link 2 and Link 3, change the cost from 40 to 4 for OSPFv3.

31. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, when Link 5 or Link 6 is removed, IPv4 and IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4). When the link is restored, traffic flow is restored to its original path (TR1-TR5-TR4). When Link 2 is removed, there should be no effect on the traffic flow.
- In Part B, when the cost on Link 5 and Link 6 is adjusted, IPv4 and IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4). When the cost is restored, traffic flow is restored to its original path (TR1-TR5-TR4). When the cost on Link 2 and Link 3 is adjusted, there should be no effect on the traffic flow.
- In Part C, when the cost on Link 5 and Link 6 is adjusted for OSPFv2, IPv4 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4) and IPv6 traffic flow remains the same. When the cost is restored, IPv4 traffic flow is restored to its original path (TR1-TR5-TR4). When the cost on Link 5 and Link 6 is adjusted for OSPFv3, IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4) and IPv4 traffic flow remains the same. When the cost is restored, IPv6 traffic flow is restored to its original path (TR1-TR5-TR4).

**Possible Problems:** None.

## Test RTM.1.4: Multi-Area OSPFv3 Convergence with 6over4 IPv6 Tunnels

**Purpose:** To verify that OSPF convergence can occur with 6over4 IPv6 Tunnels.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 29, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. OSPFv2 and OSPFv3 are two separate protocols that route IPv4 and IPv6 traffic respectively. These protocols use separate software processes and should operate like “ships in the night”, in that they do not interfere with each other.

**Test Setup:** Connect the network as described in the figure below. Enable OSPF, set all link costs to 4 for OSPFv2 and OSPFv3. Configure TR1 to be in area 0. Configure TR3 and TR4 to be in area 1. Configure TR2 and TR5 to be ABRs between area 0 and area 1. Disable IPv6 on Links 2, 3, 5 and 6 so only IPv4 is supported. Create 6over4 IPv6 tunnels TR1-TR4 and TR1-TR3. It is recommended that once this test is run with success, it should be rerun with a larger number of routers and tunnels to demonstrate the functionality of 6over4 IPv6 tunnels. This second run should also use G1 and G2 to inject a realistic number of routes into this network. The minimum topology is covered here to allow laboratories with limited resources to reproduce this test.



**Procedure:**

*Part A: Link Down, same topologies*

1. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.
2. Disconnect Link 5.
3. Allow the network to converge and observe the egress data traffic.
4. Reconnect Link 5.
5. Allow the network to converge and observe the egress data traffic.
6. Disconnect Link 6.
7. Allow the network to converge and observe the egress data traffic.
8. Reconnect Link 6.
9. Allow the network to converge and observe the egress data traffic.
10. Disconnect Link 2.
11. Allow the network to converge and observe the egress data traffic.
12. Reconnect Link 2.

13. Allow the network to converge and observe the egress data traffic.

*Part B: Link Cost Change, same topologies*

14. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.

15. On TR5's interfaces on Link 5 and Link 6, change the cost from 4 to 40 for both OSPFv2 and OSPFv3.

16. Allow the network to converge and observe the egress data traffic.

17. On TR5's interfaces on link 5 and Link 6, change the cost from 40 to 4 for both OSPFv2 and OSPFv3.

18. Allow the network to converge and observe the egress data traffic.

19. On TR2's interfaces on Link 2 and Link 3, change the cost from 4 to 40 for both OSPFv2 and OSPFv3.

20. Allow the network to converge and observe the egress data traffic.

21. On TR2's interfaces on Link 2 and Link 3, change the cost from 40 to 4 for both OSPFv2 and OSPFv3.

22. Allow the network to converge and observe the egress data traffic.

*Part C: Link Cost Change, different topologies*

23. Pass continuous IPv4 and IPv6 traffic from G1 to G2 and from G2 to G1.

24. On TR5's interfaces on Link 5 and Link 6, change the cost from 4 to 40 for OSPFv2.

25. Allow the network to converge and observe the egress data traffic.

26. On TR5's interfaces on link 5 and Link 6, change the cost from 40 to 4 for OSPFv2.

27. Allow the network to converge and observe the egress data traffic.

28. On TR2's interfaces on Link 2 and Link 3, change the cost from 4 to 40 for OSPFv3.

29. Allow the network to converge and observe the egress data traffic.

30. On TR2's interfaces on Link 2 and Link 3, change the cost from 40 to 4 for OSPFv3.

31. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, when Link 5 or Link 6 is removed, IPv4 and IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4). When the link is restored, traffic flow is restored to its original path (TR1-TR5-TR4). When Link 2 is removed, there should be no effect on the traffic flow.
- In Part B, when the cost on Link 5 and Link 6 is adjusted, IPv4 and IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4). When the cost is restored, traffic flow is restored to its original path (TR1-TR5-TR4). When the cost on Link 2 and Link 3 is adjusted, there should be no effect on the traffic flow.
- In Part C, when the cost on Link 5 and Link 6 is adjusted for OSPFv2, IPv4 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4) and IPv6 traffic flow remains the same. When the cost is restored, IPv4 traffic flow is restored to its original path (TR1-TR5-TR4). When the cost on Link 5 and Link 6 is adjusted for OSPFv3, IPv6 traffic flow changes from (TR1-TR5-TR4) to (TR1-TR2-TR3-TR4) and IPv4 traffic flow remains the same. When the cost is restored, IPv6 traffic flow is restored to its original path (TR1-TR5-TR4).

**Possible Problems:** None.

## **GROUP 2: BGP Functionality with Tunnels**

### **Scope:**

These following tests are designed to verify basic interoperability of BGP in an IPv6 Transition environment.

### **Overview:**

BGP-4 and BGP-4+ have been tested in a multi-vendor environment at a very basic level with both convergence and interoperability being demonstrated. It is still an unknown to see how the BGP routing algorithm will run over various IPv6 Tunnels that span an IPv4-only network. While it is an ideal that native IPv6 will be everywhere, it is a necessity to transport IPv6 over IPv4. It is a network requirement to demonstrate that routing functionality is stable in this environment. In the tests in this section, it is assumed that BGP is configured to peer only with directly connected neighbors.

## Test RTM.2.1: Static IPv6 Tunnel BGP-4+ Convergence

**Purpose:** To verify that BGP-4+ Convergence can occur with Static IPv6 Tunnels.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 29, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. This test is a demonstration of network operation when Static IPv6 Tunnels are used in a network.

**Test Setup:** Connect the network as described in the figure below. Enable BGP and configure each router to be in the proper AS. Only enable IPv4 on Links 1, 2, 3 and 4. Create static IPv6 tunnels on each of the links. It is recommended that once this test is run with success, it should be rerun with a larger number of routers and tunnels to demonstrate the functionality of static IPv6 tunnels. This second run should also use G1 and G2 to inject a realistic number of routes into this network. The minimum topology is covered here to allow laboratories with limited resources to reproduce this test.



### Procedure:

#### Part A: AS Path Prepend

1. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
2. Configure a policy to prepend an outbound AS Path attribute on TR1 between TR3 and TR1.
3. Configure a policy to prepend an outbound AS Path attribute on TR4 between TR3 and TR4.
4. Allow the network to converge and observe the egress data traffic.
5. Configure a policy to prepend two outbound AS Path attributes on TR1 between TR2 and TR1.
6. Configure a policy to prepend two outbound AS Path attributes on TR4 between TR2 and TR4.
7. Allow the network to converge and observe the egress data traffic.

#### Part B: Link Down

8. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
9. Disconnect Link 2.

10. Allow the network to converge and observe the egress data traffic.
11. Reconnect Link 2.
12. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, at Step 4, traffic is observed flowing from (TR1-TR2-TR4). When the AS Path information is changed in Step 7, IPv6 traffic flow changes from (TR1-TR2-TR4) to (TR1-TR3-TR4).
- In Part B, when the link is removed, IPv6 traffic flow changes from (TR1-TR3-TR4) to (TR1-TR2-TR4). When the link or cost is restored, traffic flow is also restored to its original path.

**Possible Problems:** None.

## Test RTM.2.2: 6over4 Tunnel BGP-4+ Convergence

**Purpose:** To verify that BGP-4+ Convergence can occur with 6over4 Tunnels.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 29, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. When the topology changes, routing protocols update their tables to reflect the change and optimize traffic flow. This test is a demonstration of network operation when Static IPv6 Tunnels are used in a network.

**Test Setup:** Connect the network as described in the figure below. Enable BGP and configure each router to be in the proper AS. Only enable IPv4 on Links 1, 2, 3 and 4. Create 6over4 IPv6 tunnels on each of the links. It is recommended that once this test is run with success, it should be rerun with a larger number of routers and tunnels to demonstrate the functionality of 6over4 IPv6 tunnels. This second run should also use G1 and G2 to inject a realistic number of routes into this network. The minimum topology is covered here to allow laboratories with limited resources to reproduce this test.



### Procedure:

#### Part A: AS Path Prepend

1. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
2. Configure a policy to prepend an outbound AS Path attribute on TR1 between TR3 and TR1.
3. Configure a policy to prepend an outbound AS Path attribute on TR4 between TR3 and TR4.
4. Allow the network to converge and observe the egress data traffic.
5. Configure a policy to prepend two outbound AS Path attributes on TR1 between TR2 and TR1.
6. Configure a policy to prepend two outbound AS Path attributes on TR4 between TR2 and TR4.
7. Allow the network to converge and observe the egress data traffic.

#### Part B: Link Down

8. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
9. Disconnect Link 2.

10. Allow the network to converge and observe the egress data traffic.
11. Reconnect Link 2.
12. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, at Step 4, traffic is observed flowing from (TR1-TR2-TR4). When the AS Path information is changed in Step 7, IPv6 traffic flow changes from (TR1-TR2-TR4) to (TR1-TR3-TR4).
- In Part B, when the link is removed, IPv6 traffic flow changes from (TR1-TR3-TR4) to (TR1-TR2-TR4). When the link or cost is restored, traffic flow is also restored to its original path.

**Possible Problems:** None.

## Test RTM.2.3: BGP-4 and BGP-4+ with IPv4 and IPv6 Transport

**Purpose:** To verify that BGP-4+ can operate with both IPv4 and IPv6 to transport the routing information.

**References:**


**Resource Requirements:**

- Monitor to capture packets

**Last Modification:** September 29, 2004

**Discussion:** Routing protocols share link information to determine the best path through the network. In current practice, IPv4 BGP routing information is exchanged between nodes with IPv4 transport and IPv6 BGP routing information is exchanged between nodes with IPv6 transport. In theory, nodes should support exchange of both IPv4 and IPv6 routing information over IPv4 transport or both IPv4 and IPv6 routing information over IPv6 transport.

**Test Setup:** Connect the network as described in the figure below. Enable BGP and configure each router to be in the proper AS. It is recommended that once this test is run with success, it should be rerun with a larger number of routers and tunnels to demonstrate the functionality of IPv4 only transport or IPv6 only transport. This second run should also use G1 and G2 to inject a realistic number of routes into this network. The minimum topology is covered here to allow laboratories with limited resources to reproduce this test.



### Procedure:

#### Part A: IPv4 only Transport

1. Only enable IPv4 on Links 1, 2, 3 and 4. Configure BGP to transport all routing information using IPv4 transport.
2. Configure a policy to prepend an outbound AS Path attribute on TR1 between TR3 and TR1.
3. Configure a policy to prepend an outbound AS Path attribute on TR4 between TR3 and TR4.
4. Allow the network to converge and observe the egress data traffic.
5. Inject IPv4 and IPv6 route information from G1 and G2.
6. Pass continuous IPv4 traffic from G1 to G2 and from G2 to G1.
7. Allow the network to converge and observe the egress data traffic.

#### Part B: IPv6 only Transport

8. Only enable IPv6 on Links 1, 2, 3 and 4. Configure BGP to transport all routing information using IPv6 transport

9. Configure a policy to prepend an outbound AS Path attribute on TR1 between TR3 and TR1.
10. Configure a policy to prepend an outbound AS Path attribute on TR4 between TR3 and TR4.
11. Allow the network to converge and observe the egress data traffic.
12. Inject IPv4 and IPv6 route information from G1 and G2.
13. Pass continuous IPv6 traffic from G1 to G2 and from G2 to G1.
14. Allow the network to converge and observe the egress data traffic.

**Observable Results:**

- In Part A, the routing information should be propagated between all devices. IPv4 traffic should be properly distributed end to end.
- In Part B, the routing information should be propagated between all devices. IPv6 traffic should be properly distributed end to end.

**Possible Problems:** None.