Internet Routing Protocol Testing

QA Robot

Introduction

The exponential growth in the use of Internet-related technologies is by now such a well recognized phenomenon that we can read about this in publications ranging from research papers to supermarket tabloids. Many readers recognize that at the core of this information superhighway lies a complex network of ever-more-powerful switches and routers. The primary impetus for this increasing power has been frequent and significant technological leaps in hardware design and the usable bandwidth of communications media. Evolving at the same time as these have been the protocols and methods for managing these switches and routers and their operation together as a cohesive network.

This paper focuses on what is arguably the most complex component of managing these networks - the exchange and maintenance of routing information. These protocols are implemented in complex software. The relative immaturity of some of these protocols, and especially the unforeseen scale of the applications, have created an unprecedented need for the capability to test IP routing protocols.

Agilent Technologies
Innovating the HP Way

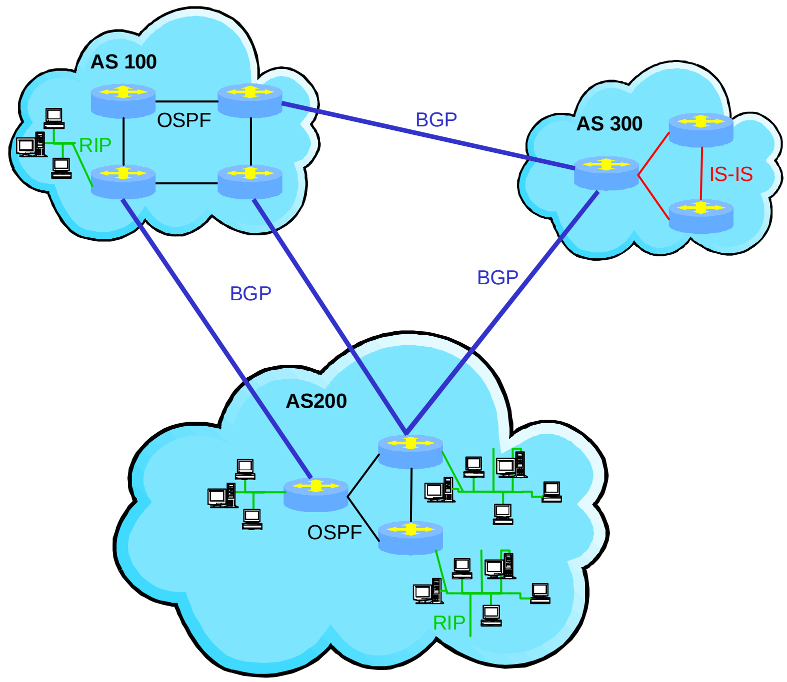

Figure 1: Autonomous systems (AS) can be viewed as super-nodes in a higher-level topology, while the interior of an AS shows the more traditional topological graph where routers are nodes and networks are edges.

Popular Routing Protocols

Autonomous Systems, EGPs and IGPs

The public Internet is operated by many different Internet Service Providers (ISPs) in countries around the world. In the United States alone, there are numerous ISPs of diverse sizes. It is important to acknowledge the conundrum that the ISPs comprising the Internet are at once in collaboration and in competition with one another. While two ISPs in one sense compete for customers' traffic, they certainly recognize that providing unfettered worldwide connectivity bridging all ISPs is perhaps the biggest attraction of the Internet.

It is this competitive/collaborative dichotomy that gives birth to the concept of Autonomous Systems. An ISP maintains control over the topology of the network that it administers. The terminology for such a grouping is an *autonomous system* (AS). If the number of routers being administered by an ISP is very large, it is possible that the ISP may divide its network into multiple autonomous systems. While detailed knowledge about network topology is both available and required for optimal routing within an AS, such knowledge is simply not available between different autonomous systems. There are at least two clear reasons for this:

- Different AS's are frequently administered by competing organizations and this information would be considered proprietary.
- Calculating routing tables needs to be done at multiple levels of abstraction in order to be computable in a reasonable period of time

The complexity of a shortest paths computation (e.g., Dijkstra's algorithm) for the graph of all the routers and networks in the public Internet is almost beyond comprehension. Thehe complexity of a shortest paths computation (e.g., Dijkstra's algorithm) for the graph of all the routers and networks in the public Internet is almost beyond comprehension. The practical solution to this is to abstract out this level of detail and to view the Internet as a graph of autonomous systems at the world-wide level and only compute router-to-router paths within a given autonomous system.

Since much more detailed topological information is used for computing routes within an AS than between ASs, routing protocols can be divided into two classes:

- those used for routing within an AS (interior gateway protocols), and
- those used for routing between ASs (exterior gateway protocols).

Before providing further detail about these two classes of routing protocols, we should first document another important routing classification: unicast routing vs. multicast routing. An extremely simplistic description is that a unicast network address denotes a single destination for a packet whereas a multicast network address denotes a potentially large set of destination addresses. This is such a fundamental difference that these address spaces are treated distinctly by routers and indeed have their own unique sets of routing protocols. It is important, though, to understand that the fundamental notions of the autonomous system, IGPs and EGPs apply to multicast routing protocols as well.

Unicast Routing Protocols

The vast majority of the Internet network addresses exchanged between and within autonomous systems are unicast addresses. In the Internet, and even within a single autonomous system, more than one protocol participates in this exchange. For example, a very reasonable configuration would be similar to that depicted for AS 200 in Figure 1 where we see a combination of RIP and OSPF serving the role of IGP. This hybrid use of two different IGPs is justified since a simple routing protocol like RIP is adequate for accumulating and propagating network routes at the topological periphery of an AS, whereas the capabilities of a link-state IGP such as OSPF or ISIS are required as path computation becomes more complex away from the periphery.

Well-known unicast protocols

Exterior Gateway Protocols

- BGP
- EGP (obsolete)

Interior Gateway Protocols

- OSPF
- ISIS
- RIP
- IGRP (Cisco proprietary)
- EIGRP (Cisco proprietary)

Multicast Routing Protocols

Multicast addresses occupy a separate address space from unicast addresses. This special characteristic is denoted by the setting of special bits in the address. For IP version 4, this can be recognized by IP addresses that look like 224.*.*.*, etc. Because this is a separate address space with distinct forwarding characteristics, these routes are collected, computed, and re-advertised using distinct routing protocols. A partial list of these routing protocols is provided below:

- PIM (sparse and dense mode)
- DVMRP
- MOSPF

These protocols have yet to see wide use in the public Internet. As such, the balance of this paper will focus on unicast issues. Because of anticipated applications requiring broadcast transmission, these may become very important in the future and do require careful watching.

Advanced Router Architecture

User-plane vs. Control-plane

The most touted attribute of giga- and tera-bit routers is their data forwarding capacity, which seems to outstrip former top-of-the-line equipment at an ever-increasing rate. This user-plane aspect of the engineering problem is the one most focussed on in the new router architectures and is the one most likely to perform in keeping with predictions. What represents a much greater unknown is the handling of control-plane traffic. The user-plane in a modern router is generally a hardware-based switching matrix of one of a variety of designs. All IP datagrams not addressed to the router itself are forwarded exclusively by this hardware.

The control-plane traffic is comprised of packets addressed to the router itself, which includes, among others:

- routing protocol packets
- SNMP packets for network management
- HTTP packets for network management

Some of the reasons that control plane traffic needs special attention are:

- Control-plane packets must be "stolen" from data path and delivered to a software-based packet processor, usually more susceptible to implementation errors than the hardware-based user-plane forwarding mechanism.

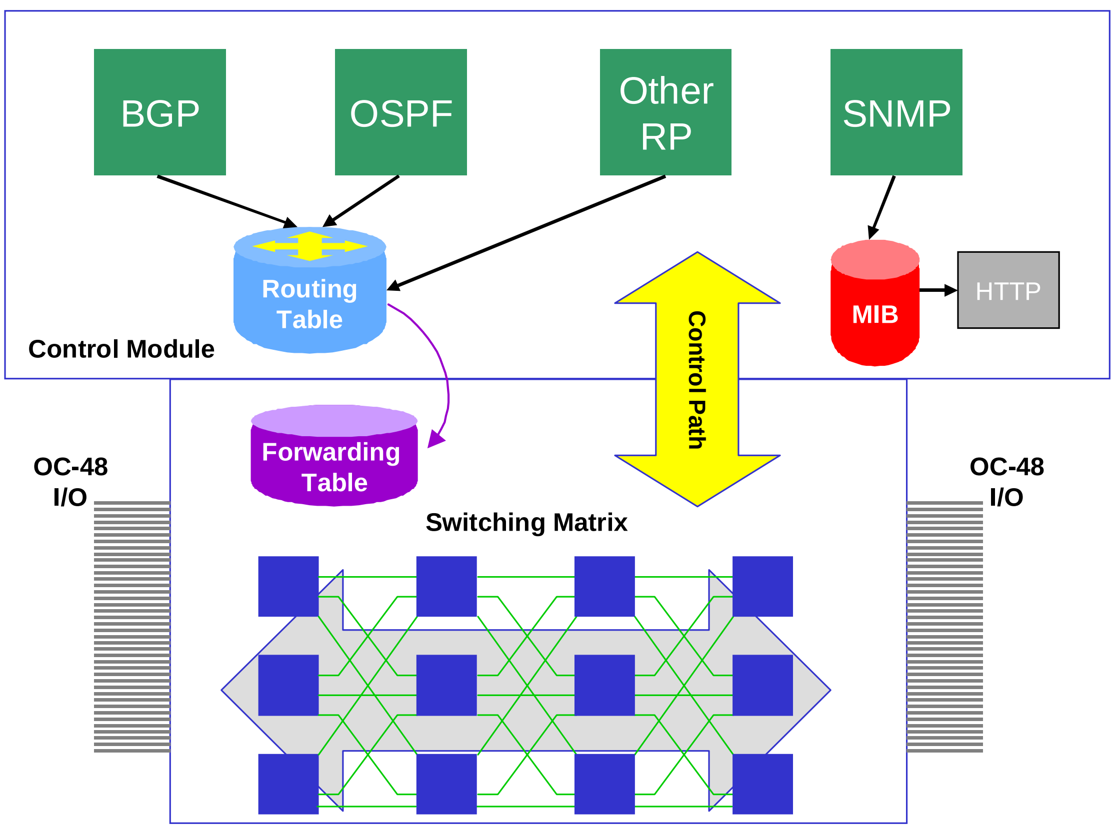

Figure 2: An idealized view of a modern high-performance router

- Multiple routing protocols all need to contribute to and be driven by a unified routing table
- These routing protocols and the associated route computations are complex, and implementing these is error-prone
- Since a single router's routing table is actually part of a distributed database of routing tables in all the routers in the network, they must be kept consistent and therefore the timing and synchronization of this information, not only within a single router but between neighboring routers is, critical
- Since it is very likely that the router will simultaneously be handling routing protocols, SNMP and possibly HTTP traffic, and these are likely processed by very different software components, it is important to ascertain that these components interact in a manner giving correct and reasonable responsiveness to each.

In addition to these formal control-plane packets, there is additional control interaction between the software-based packet processor and the hardware-based switching engine. While this interaction takes many forms, the one most relevant to this paper is the programming of the forwarding table accessed by the hardware as part of its regular forwarding operations. We noted above that multiple protocols interact to maintain a unified routing table consistent with the routing tables in the router's neighbors. It is at least as important that the forwarding table in the router be consistent with its routing table. Since these are separate databases, it is imperative to ascertain how well the forwarding table tracks the changing routing table in periods of routing change.

Classes of Testing

Network Equipment Manufacturers (NEMs) and ISPs have potentially different goals in their testing. A manufacturer may be very concerned about protocol compliance while an ISP may not. An ISP or even large enterprise network managers may often be involved in performing comparative analysis of many different manufacturers platforms. Both NEMs and ISPs are interested in being able to confirm that a particular box or network of boxes performs according to advertised specifications. Since modern router architecture portends very high performance levels, this requires test equipment that can operate at levels above the maximum expected performance of the network itself.

In the following paragraphs we describe some of the basic different classes of testing related to routing protocols.

RFC Compliance

Compliance testing is somewhat different for IETF protocols because:

- RFCs are often not expressed with such rigid and formal specifications as, for example, ATM forum documents, and
- the rapidly changing nature of the Internet itself forces these protocols to be in a constant state of flux, requiring ever-evolving compliance test tools.

For example, the original BGP-4 standard RFC1771 has been augmented by RFCs 1965,1966, 1997, 2283, 2385, 2545 and 2547. The original OSPF specification 1583 was obsoleted by 2328 and has been extended in 2370 and 2676. ISIS, originally published as an ISO standard for propagating routing information for OSI networks, was replicated in RFC 1142 and extended in RFC 1195 to allow it to be used for IP networks. There are a number of newer Internet-draft RFCs that

continue to extend ISIS so that it can keep pace with the state of the art of IP routing.

Compliance (or conformance) testing ensures that an implementation claiming to support one or more of these specifications has done so completely and correctly.

Data Stress

This class of testing is used to ascertain that the router is able to receive packets at some rate K packets per second (of some statistically described packet size) on N inputs, and forward them on M outputs without loss and perhaps within some latency and jitter constraints. Pure data stress testing normally requires a type of test equipment sometimes called *wire-rate* testers. Of course, a wire-rate tester for OC-192 speeds will require hardware-based packet generation/processing, whereas a 100 Mb/s wire-rate tester may reasonably be software-based. Wire-rate data stress testing may range from filling the available bandwidth with packets that are of a uniform size and traffic characteristics to a realistic traffic mix including the different packet sizes, protocols, and other aspects of the diverse traffic mix that actually exists in a real Internet link.

Routing Stress

As powerful as the modern router designs are, since they are machines limited by physical realities, there are limits to how fast they can perform and how much data they can store. Since it is a virtual given that these limits will continually be pushed, it is important to have test tools that can determine what these operational limits are for a given router or network of routers. Some of the operating parameters that should be verified while operating under periods of intense routing traffic are:

- maximum routing table size
- maximum rate of learning new routes
- maximum rate of changes in existing routes (route flaps)
- that routes are consistent between neighbors
- that all advertised routes are present.

Interaction Between Routing Protocols

Normally more than one routing protocol will contribute to the routing information used and propagated by the router. This is illustrated in Figure 2 where we see BGP, OSPF, as well as potentially other routing protocols (*RP* in Figure 2) contributing to and taking input from the single routing table. The most obvious case of this is in routers that are autonomous system border routers. Using an EGP-class protocol, these routers teach external autonomous systems about networks within their own autonomous system, and, conversely, must propagate information about external networks throughout their own autonomous systems. This is accomplished using an IGP-class protocol. The mechanisms via which this interaction takes place are not specified as part of the RFCs and thus are an especially important area to test.

Interaction Between Routing Table and Forwarding Table

Figure 2 indicates that the routing table and forwarding table are likely to be implemented in separate hardware modules. The operations by which changes propagate from the routing table to the forwarding table are potentially complex and certainly very time-critical. These two databases must remain closely synchronized or routing loops and packet loss are likely.

Tests that verify the timing and accuracy of data forwarding with respect to routing changes will test this interaction between routing and forwarding tables.

In the following sections we present six sample test scenarios that provide examples of some of the basic classes of routing protocol testing that we have presented here.

Test Scenario #1: BGP Stress

Figure 3 depicts a router maintaining BGP sessions with numerous peers in different autonomous sessions. The BGP protocol should be stressed along two different dimensions:

- the number of peer sessions in use, and
- the number of routes/route changes processed per unit of time.

Clearly, to set up hundreds of routers in a laboratory is unwieldy at best. Moreover, it is very difficult to 'trick' these routers into advertising millions of network routes unless one of the routers has actually advertised them in the first place. In addition, it is virtually impossible to verify that the DUT actually propagates the routes correctly and in a timely manner.

In Figure 4 we show how such a scenario can be simulated and controlled in the laboratory. All 300 peer sessions are created from a single test box which injects one million network routes on one of the peer sessions. The test equipment then automatically verifies that the DUT propagates the learned routes to the other 299 peers within a defined time window.

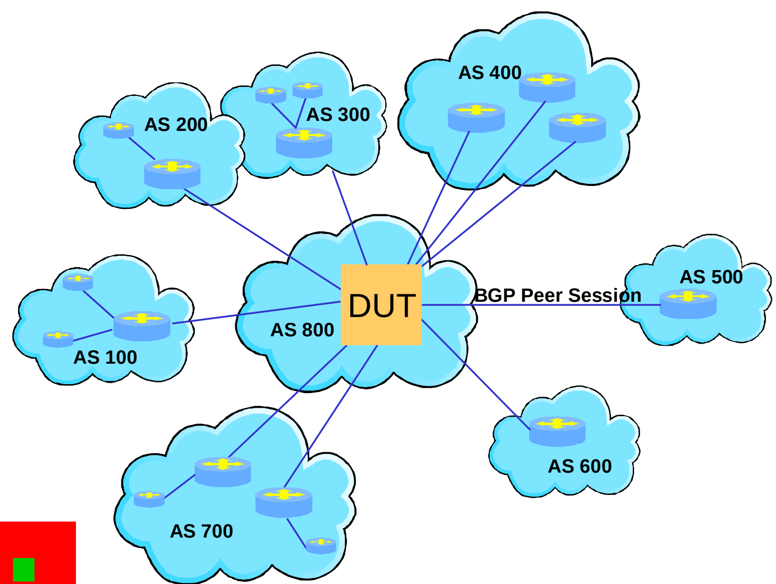

Fig. 3: 300 peers, millions of routes

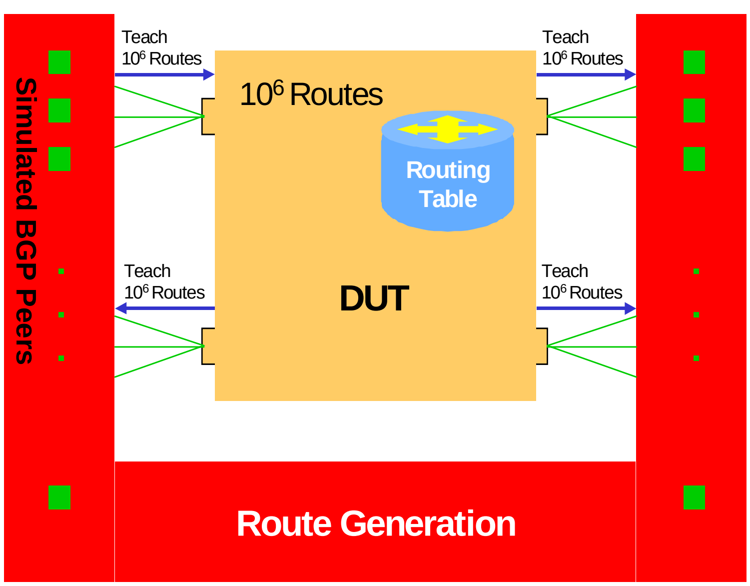

Fig. 4: Test scenario #1 in the lab

Test Scenario #2: OSPF Stress

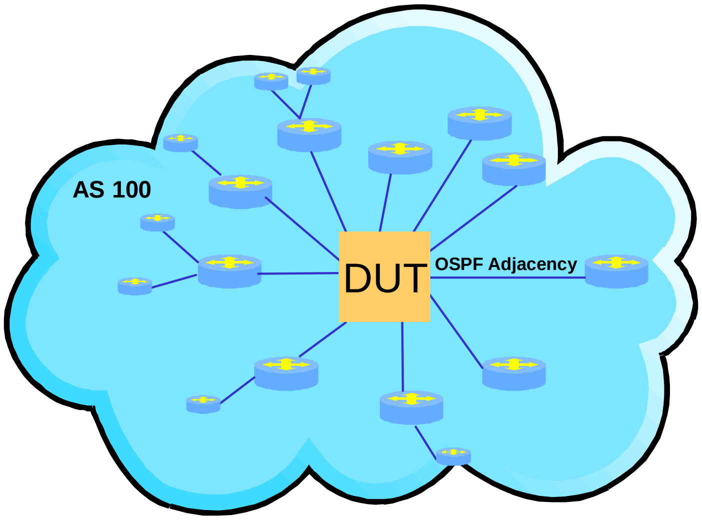

Fig. 5: 256 adjacencies, thousands of network routes

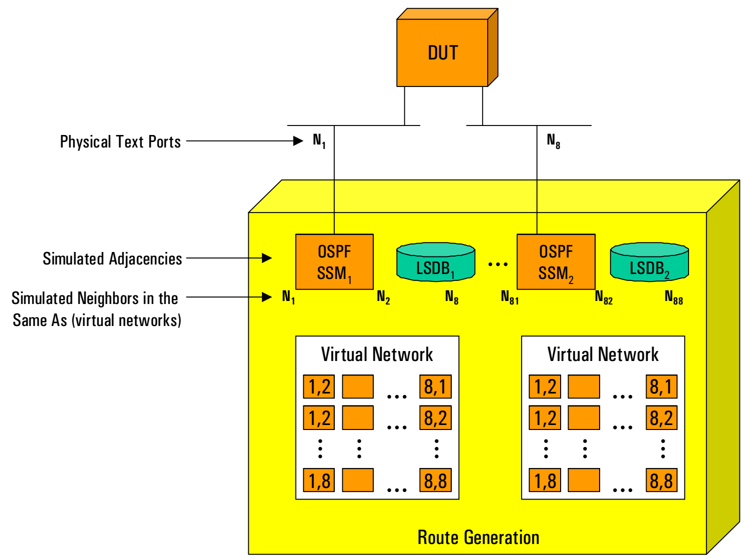

Fig. 6: Test scenario #2 in the lab

Figure 5 depicts a router maintaining OSPF adjacencies with a large number of neighbors in the same autonomous system. Like BGP, an OSPF (or ISIS for that matter) implementation should be stressed along at least two different dimensions:

- the number of adjacencies in use, and
- the number of LSAs/LSA (i.e., route) changes processed per unit of time.

Ideally, it should be possible to connect the DUT to a single piece of test equipment that will appear to the DUT as if it were connected to 256 OSPF-enabled routers. In Figure 6 we show how this can be achieved in the laboratory. All 256 adjacencies are maintained in a single test box which injects thousands of LSAs. The test equipment then automatically verifies that the LSAs advertised by each of the 256 simulated OSPF neighbors are correctly propagated to the other 255 neighbors. As one can imagine, this exchange is extremely computationally intense for the DUT and provides an opportunity to determine the efficiency and robustness of an OSPF implementation.

Test Scenario #3: BGP, OSPF and Data Interaction

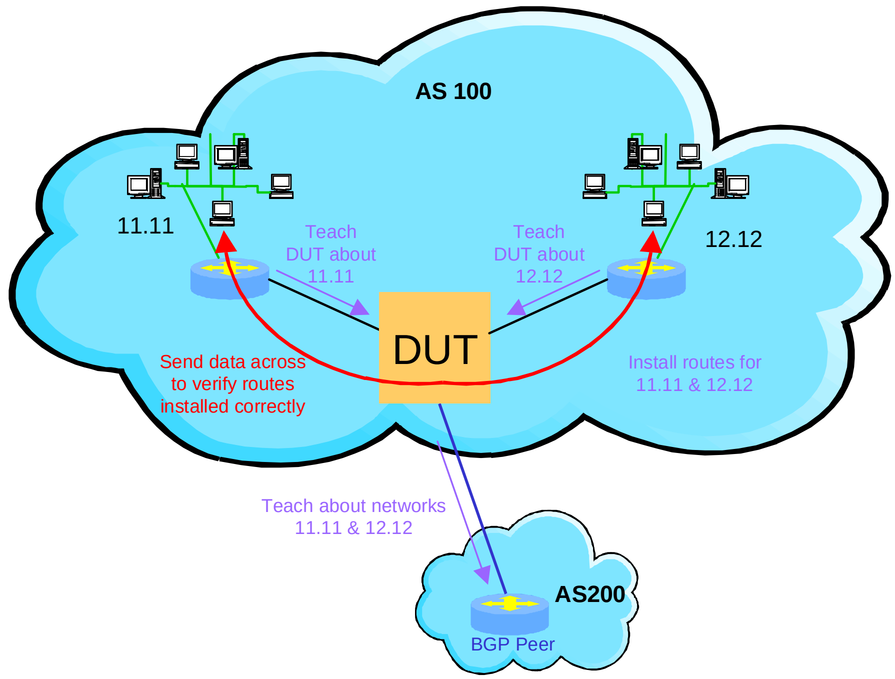

Fig. 7

In Figure 7 we depict a test scenario that examines the interaction between different routing protocols as well as basic relation between the routing table and the data forwarding table. The basic test steps are:

- Establish two OSPF adjacencies and one BGP peering session with the DUT.
- The first OSPF adjacency advertises a route to network 11.11 and the second OSPF adjacency advertises a route to network 12.12.
- The test confirms that routes to networks 11.11 and 12.12 are advertised to the BGP peer in AS200.
- IP datagrams are forwarded from a simulated client in network 11.11 to a simulated host in network 12.12, confirming that the data forwarding table was correctly updated as a result of the routing table change.

Test Scenario #4: Basic Route Flap Test

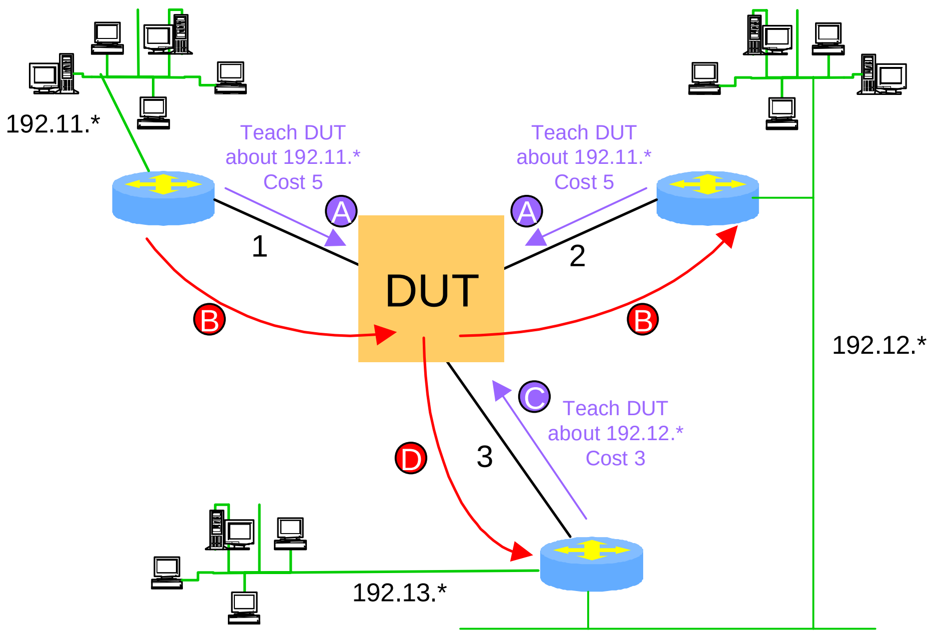

Fig. 8

In Figure 8 we depict a test scenario that stresses the interaction between the routing table and the data forwarding table. The basic test steps are:

- Establish three ISIS adjacencies with the DUT.
- The first ISIS adjacency advertises routes to 254 networks ranging from 192.11.1 to 192.11.254 with cost 5. (step A)
- The second ISIS adjacency advertises routes to 254 networks ranging from 192.12.1 to 192.12.254 with cost 5.(step A)
- The third ISIS adjacency advertises routes to 254 networks ranging from 192.13.1 to 192.13.254 with cost 5. (step A)
- The test confirms that the DUT has propagated the routes to all neighbors.
- IP datagrams are forwarded from simulated clients in networks 192.11.* to simulated hosts in networks 192.12.*, confirming that the data forwarding table was correctly updated as a result of the routing table change. (step B)
- Now the third ISIS adjacency advertises routes to 254 networks ranging from 192.12.1 to 192.12.254 with cost 3, creating a better route than that over adjacency 2. (step C)
- The test confirms that the IP datagrams being sent from networks 192.11.* are now forwarded via simulated router number 3.(step D)

This basic test can be extended in many useful ways. For example, the test can confirm that the route change occurs within acceptable time limits and that no data is lost. Also, the test may repeatedly “flap” (i.e., change) the route such that the preferred route oscillates between router number two and router number three.

QA Robot

Test Scenario #5: Wire Rate Route Flap Test

Figure 9 illustrates a more complex test scenario that subjects the DUT(s) to routing stress and data forwarding stress over different interfaces using separate pieces of test equipment. This scenario utilizes the traffic generation capability of basic wire-rate traffic generator/capture equipment and injects route advertisements and route changes via separate route generators. This scenario is not completely realistic in that in a live network control plane traffic invariably is delivered and received over the same links as the data plane traffic. Nevertheless, such a test scenario can be used effectively to stress data forwarding performance at high data rates when faced with rapid reprogramming of the data forwarding table. The fundamental differences between this test scenario and that depicted in Figure 8 are:

- Data is being forwarded over many different physical paths.
- Data is being forwarded at high (i.e., > 100 Mbps) bit rates.
- The test est is somewhat unrealistic in that routing traffic is being sent and received "out of band".
- The device under test is actually a network rather than a single router.
- The fine degree of control over the timing of the routing change and the resulting data forwarding change that was shown in Figure 8 is lost here due to the use of uncoupled test equipment.

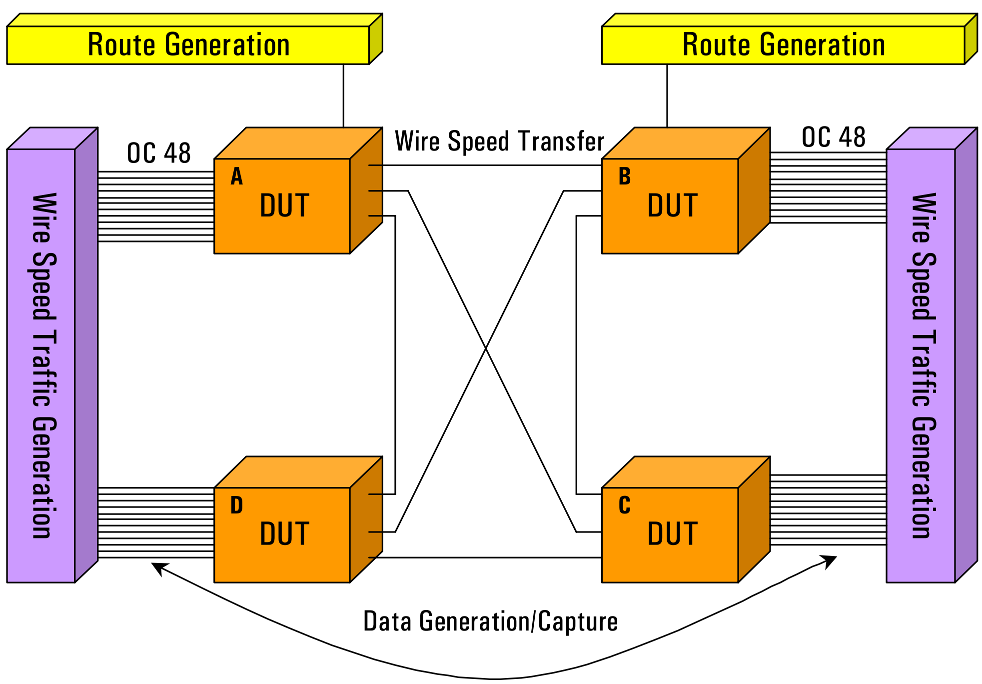

Fig. 9

Test Scenario #6: Integrated Wire Rate Route Flap Test

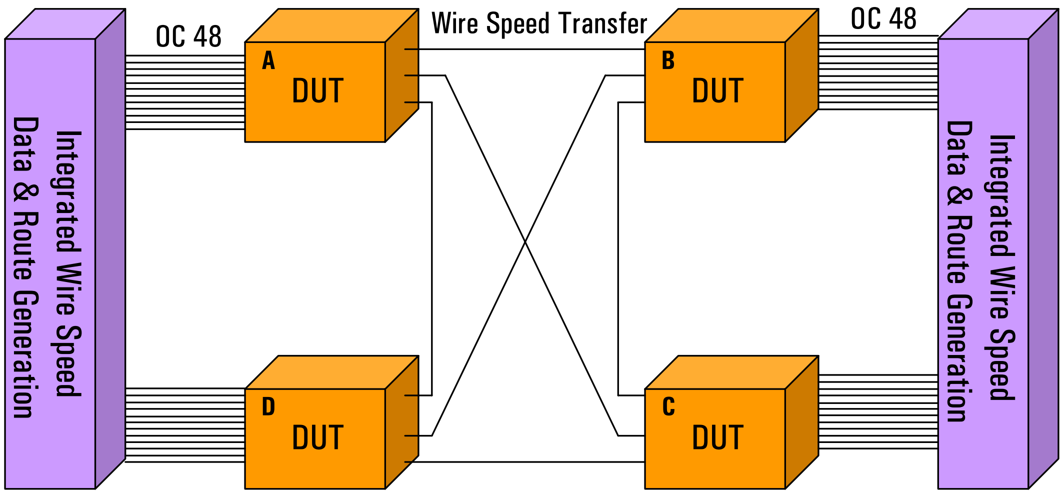

Fig. 10

Figure 10 illustrates a similar but much more realistic test scenario than Figure 9. Here, using a single, more sophisticated piece of test equipment, the DUT(s) are subjected to routing and high rate data forwarding stress over the *same* interfaces. This scenario requires a hybrid test device that has both wire-rate traffic generator/capture capability as well as routing stress capability. The fundamental differences between this test scenario and that depicted in Figure 9 are:

- Data is being forwarded at even higher bit rates (e.g., OC48 or beyond)
- The test is *realistic* in that routing traffic is being sent and received "in band".
- The test requires a hybrid test device combining features of both types of test equipment in Figure 9.
- The test restores the fine degree of control over the timing of the routing change and the corresponding data forwarding change that was shown in Figure 8 but lost in Figure 9. This is due to a single test device controlling the whole test.

Summary

In summary, the testing of routing protocols testing represents one of the more complex aspects of quality assurance and performance testing for Internet products. We have inventoried the most widely used routing protocols in the Internet today and explained how the interaction between them is critical to the correct operation of the Internet. The correctness of operation of these protocols internally, the interaction between them, and the rate at which this processing can occur is part of the routing protocol test paradigm. In addition, through the use of example test scenarios, we illustrated how user-plane data forwarding testing can be interwoven with control-plane routing protocol testing to achieve true system testing of modern router implementations.

This page intentionally left blank.

Agilent QA Robot

The Agilent Technologies QA Robot is the industry's only comprehensive test system which combines router protocol conformance, validation and stress testing with superior intelligent traffic generation capabilities in a single system. The QA Robot simulates real-time network configurations and tests high volumes of routing traffic.

www.Agilent.com/comms/QARobot

United States:

Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:

Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-9999

Japan:

Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 267-4245
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 2599-7889
Fax: (852) 2506-9233

Australia/New Zealand:

Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Specifications subject to change.

5968-9391E 03/00 Rev A

Agilent Technologies
Innovating the HP Way