

Individual  
Internet-Draft  
Expires: April 4, 2005

R. White  
R. Adams  
Cisco Systems  
V. Manral  
SiNett Corp.  
October 4, 2004

Considerations in Benchmarking Routing Protocol Network Convergence  
draft-white-network-benchmark-01

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 4, 2005.

Copyright Notice

## Abstract

This document attempts to discuss some of the definitions required to undertake the specifications of benchmarks measuring routing protocols convergence within a network, and also to discuss some of the possible ways to benchmark a routing protocol performance within a network, and some of the implications of those benchmarks. The

White, et al.

Expires April 4, 2005

[Page 1]

---

Internet-Draft

Considerations in Benchmarking RPs

October 2004

definition of convergence is discussed first, then polling network devices. Several tests which are commonly used to measure network convergence are examined.

This draft does not attempt to define what techniques should be used to benchmark network convergence, but only to provide considerations testers should consider when attempting to measure network convergence using various methods.

## 1. Motivation

As the ability to benchmark components within a network appears to be coming under greater scrutiny, and specifications are being written to standardize ways to measure the performance of individual components within given frameworks, the next level of benchmarking has not been approached, that of measuring the performance of networks. But what is meant when we say the performance of a network, from the perspective of routing protocols? Various tests have been used in the past to measure the convergence of a network, some of which actually measure completely different things than others.

It's important to attempt to examine the measurement of network convergence in a way that exposes these differences, and helps

vendors, end users, and those in the research community have some common ground when discussing network convergence.

## [2.](#) A Problem of Definitions

As we examine the issues and concepts surrounding the measurement of network performance in terms of convergence, we find that most of the basic problems we face surround defining the terms in use. For instance, what is convergence, exactly? What is a network? In the following sections, we discuss each of these concepts, and attempt to address each one.

### [2.1](#) Networks

In its most nominal form, a network is composed of a group of devices interconnected in some way, which send data over these interconnections for various purposes. But, when we discuss the concept of routing protocol convergence within a network, the definition needs to be more precise. For instance, since hosts do not, generally, participate in routing, should they be considered a part of the network when benchmarking the performance of a routing protocol? The obvious answer appears to be a resounding no, but, in some possible tests types, hosts, acting as testing devices, play a large part in the test itself.

When considering tests in which hosts participate as traffic or route generators, or other testing devices, we must consider the impact these hosts have on the test results, although we may not consider them a part of the network.

### [2.2](#) Convergence

Convergence is probably one of the hardest words in networking to define. Just about everyone who has worked on networks for a period

of time knows what it means, but no-one can explain it sufficiently to someone who doesn't understand how a network works for it to be understood. In fact, this is because there are several different meanings attributed to convergence, and which meaning is intended depends on the context in which the word is set. Convergence can mean:


- o The time at which all the routing protocol processes running on devices which participate in routing in the network agree on the best path to each reachable destination in the network.
- o The time at which the best path to each reachable destination in the network has been loaded into some local table which may then be used to forward packets (the routing information base, or RIB).
- o The time at which each router in the network has built the tables necessary to actually forward packets through the network, so that a packet transmitted from one part of the network would actually reach any given reachable destination within the network.

For instance, on a Cisco router, `show ip ospf stats` would allow the tester to see the time of the last completed SPF, `show ip route` would allow the tester to see what routes are installed in the RIB, and `show ip cef` would allow the tester to see the forwarding information which has been built from the RIB. Each test designed to measure the performance of routing protocols within a network must determine which type of convergence is being measured, if that measurement is acceptable to the information being gathered, and which test will actually measure the desired type of convergence.

### [2.3](#) Polling Devices in a Network

One common way to measure network convergence is to poll the devices in the network, using some command supplied within the routing software, to determine when particular events have occurred, or particular pieces of information have reached all the routers in the network. Polling eliminates the need for the clock of each device within the network to be synchronized for the test to have meaningful results. However, there are some issues with the rate of polling devices within the network which need to be addressed in any test which polls devices for this information; the first is the rate at which polling takes place.

If, in a test, you are attempting to measure some parameter to within one second of its occurrence, then you would need to poll at a rate much higher than once per second.




For instance, in this time line, suppose a polling event is set up to take place every second. An event is started just after some polling event takes place, but the polling process doesn't recognize the test as starting until the 1 second poll. An event occurs just before the 2 second poll, and the polling process detects this at the 2 second poll. The polling process would indicate that from the time the event started until the time the event has finished, one second has elapsed. In reality, closer to two seconds has elapsed.

The interval of the polling process can be reduced until the measurement is felt to be accurate, but it should be at least half of the desired accuracy. Common practice actually shows that it should be about one tenth of the desired accuracy.

A second consideration when polling for network events is the performance of the device running the polling process. If the process cannot poll each device at the scheduled interval, or the polling is "jittered," the time between each actual poll varies by some amount, the accuracy of the tests will be called into question. The amount of jitter introduced by the polling device, and the rate at which the device can effectively poll, should be measured in some way, and this measurement should be taken into account when designing tests which rely on polling.

Finally, when polling devices to determine when a network event occurs, issues with serialization must be considered. Most devices used for polling will not be able to poll several devices within the network at once, and will thus serialize the polling of devices.



Suppose, for instance, a single device is polling ten devices in the network. If it can poll five devices per second, it will take a full two seconds for it to detect any event on all ten devices, giving an effective accuracy of about four seconds. The amount of time required for a polling device to serialize through all the devices it is polling needs to be considered when polling a very large number of devices.

#### [2.4](#) Passing Traffic Through the Network to Determine Convergence

One of the most widely used tests for determining network convergence is starting some traffic stream at one end of a network, disrupting or completing the network, and determining how long the traffic stream is either not delivered, or takes to be delivered. For instance:

Source----R1----R2----Sink

A traffic stream is generated on Source, and the link between R1 and R2 is connected in some way. The time between the connection of this


link and the arrival of the traffic at the Sink is measured as network convergence. This type of test is extremely useful in testing real the response of a network to changing conditions. There are some considerations which should be examined when using this sort of test, or examining the results of this sort of test

#### [2.4.1](#) The Various Elements of Performance Cannot Be Separated

Using this sort of testing, there is no way to separate the performance of a routing protocol from the performance the interaction between the routing protocol and the forwarding engine, nor from the performance of the forwarding engine itself. In many tests, this is acceptable, since these are all elements of the network in total, but if specific elements of routing protocol performance are being measured, such tests can be problematic when attempting to analyze the results.

#### [2.4.2](#) The Total Convergence of the Network May Not Be Measured

If you have the following topology:



Suppose a traffic stream is sourced from Source, and then all the devices in the network are brought up (R1 through R6). The time from the device startup to the traffic stream reaching the sink is measured as network convergence.

As soon as the path Source, R1, R2, Sink converges, the Sink will begin receiving traffic, and the network will be considered to be converged by the test. However, without polling the remaining routers, R3 through R6, there is no way to know if those routers have also converged on the best path to the Sink and the Source. While this example may be considered extreme, there are many complex topologies where:

- o The path chosen by the traffic stream may not be the path expected.
- o The path chosen by the traffic stream may switch during network convergence, with the stream taking some secondary path at first, and the successively better paths converging over the life of the test.
- o The path chosen by the traffic stream switches so quickly that no traffic is lost, while the routing protocols still take some time to converge.

Tests which rely on traffic passing through the network to determine network convergence times should thoroughly examine the way in which the test topology converges, and examine the consistency of that convergence, with enough test runs to get a good feel for the range of possible results. Examining the same test sequence with slight changes in the network topology may help to provide an understanding of how the network under test converges, and also may help to provide more insight into the factors impacting convergence in the test network.

It's also possible that if the test network does not converge completely for some time after the test traffic successfully passes through the topology, the continuing convergence could impact the results of a second test run, if the test runs are placed too closely together. If a first test is run, and a second test is started immediately on traffic making it through the test topology, the

results of the second test may be skewed by convergence which is still taking place from the first test run.



These are important considerations which should be noted when examining or performing tests which rely on the presence of a data stream within a routing system to measure convergence.

### [3](#) Informative References

[OSPF-BENCH]

Manral, V., White, R. and A. Shaikh, "Benchmarking Basic OSPF Single Router Control Plane Convergence", [draft-ietf-bmwg-ospfconv-intraarea-10](#) (work in progress), July 2004.

### Authors' Addresses

Russ White  
Cisco Systems  
riw@cisco.com

Robert Adams  
Cisco Systems  
robeadam@cisco.com

Vishwas Manral  
SiNett Corp.  
vishwas@sinett.com

## Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED

WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

## Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.