

Benchmarking Applicability for Basic OSPF Convergence
draft-ietf-bmwg-ospfconv-applicability-00.txt

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

2. Abstract

This draft describes the applicability of [2] and similar work which may be done in the future. Refer to [3] for terminology used in this draft and [2]. The draft defines the advantages as well as limitations of using the method defined in [2], besides describing the pitfalls to avoid during measurement.

3. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [1].

4. Motivation

There is a growing interest in testing SR-Convergence for routing protocols, with many people looking at testing methodologies which can provide information on how long it takes for a network to converge after various network events occur. It is important to consider the framework within which any given convergence test is executed when attempting to apply the results of the testing, since the framework can have a major impact on the results. For instance, determining when a network is converged, what parts of the router's operation are considered within the testing, and other such things will have a major impact on what apparent performance routing protocols provide.

This document describes in detail the various benefits and pitfalls of tests described in [2]. It also explains how such measurements can be useful for providers and the research community.

5. Advantages of Such Measurement

- o To be able to compare the iterations of a protocol implementation. It is often useful to be able to compare the performance of two iterations of a given implementation of a protocol to determine where improvements have been made and where further improvements can be made.
- o To understand, given a set parameters (network conditions), how a particular implementation on a particular device is going to perform. For instance, if you were trying to decide the processing power (size of device) required in a certain location within a network, you can emulate the conditions which are going to exist at that point in the network and use the test described to measure the performance of several different routers. The results of these tests can provide one possible data point for an intelligent decision.

If the device being tested is to be deployed in a running network, using routes taken from the network where the equipment is to be deployed rather than some generated topology in these

tests will give results which are closer to the real performance of the device. Care should be taken to emulate or take routes from the actual location in the network where the device will be (or would be) deployed. For instance, one set of routes may be taken from an abr, one set from an area 0 only router, various sets from stub area, another set from various normal areas, etc.

- o To measure the performance of an OSPF implementation in a wide variety of scenarios.
- o To be used as parameters in OSPF simulations by researchers. It may some times be required for certain kinds of research to measure the individual delays of each parameter within an OSPF implementation. These delays can be measured using the methods defined in [2].
- o To help optimize certain configurable parameters. It may some times be helpful for operators to know the delay required for individual tasks so as to optimize the resource usage in the network i.e. if it is found that the processing time is x seconds on a router, it would be helpful to determine the rate at which to flood LSAs to that router so as to not overload the network.

6. Assumptions Made and Limitations of such measurements

- o The interactions of SR-Convergence and forwarding; testing is restricted to events occurring within the control plane. Forwarding performance is the primary focus in [4] and it is expected to be dealt with in work that ensues from [5].
- o Duplicate LSAs are Acknowledged Immediately. A few tests rely on the property that duplicate LSA Acknowledgements are not delayed but are done immediately. However if some implementation does not acknowledge duplicate LSAs immediately on receipt, the testing methods presented in [2] could give inaccurate measurements.
- o It is assumed that SPF is non-preemptive. If SPF is implemented so that it can (and will be) preempted, the SPF measurements taken in [2] would include the times that the SPF process is not running ([2] measures the total time taken for SPF to run, not the amount of time that SPF actually spends on the device's processor), thus giving inaccurate measurements.
- o Some implementations may be multithreaded or use a

multiprocess/multirouter model of OSPF. If because of this any of the assumptions taken in measurement are violated in such a model, it could lead to inaccurate measurements.

- o The measurements resulting from the tests in [2] may not provide the information required to deploy a device in a large scale network. The tests described focus on individual components of an OSPF implementation's performance, and it may be difficult to combine the measurements in a way which accurately depicts a device's performance in a large scale network. Further research is required in this area.

7. Observations on the Tests Described in [2]

Some observations taken while implementing the tests described in [2] are noted in this section.

7.1. Measuring the SPF Processing Time Externally

The most difficult test to perform is the external measurement of the time required to perform an SPF calculation, since the amount of time between the first LSA which indicates a topology change and the duplicate LSA is critical. If the duplicate LSA is sent too quickly, it may be received before the device under test actually begins running SPF on the network change information. If the delay between the two LSAs is too long, the device under test may finish SPF processing before receiving the duplicate LSA. It is important to closely investigate any delays between the receipt of an LSA and the beginning of an SPF calculation in the device under test; multiple tests with various delays might be required to determine what delay needs to be used to accurately measure the SPF calculation time.

7.2. Noise in the Measurement Device

The device on which measurements are taken (not the device under test) also adds noise to the test results, primarily in the form of delay in packet processing and producing output from which measurements are taken. The largest source of noise is generally the delay between the receipt of packets by the measuring device and the information about the packet reaching the device's output, where the event can be measured. The following steps may be taken to reduce this sampling noise:

- o Take lot of samples. The more samples which are taken, the less

that noise in the measurements will impact the overall measurement, as noise will tend to average out over a large number of samples.

- o Try to take time-stamp for a packet as early as possible. Depending on the operating system being used on the box, one can instrument the kernel to take the time-stamp when the interrupt is processed. This does not eliminate the noise completely, but at least reduces it.
- o Keep the measurement box as lightly loaded as possible, unless the loading is part of the test itself.
- o Having an estimate of noise can also be useful.

The DUT also adds noise to the measurement. The first and third points also apply to the DUT.

7.3. Gaining an Understanding of the Implementation Improves Measurements

While the tester will (generally) not have access to internal information about the OSPF implementation being tested using [2], the more thorough the tester's knowledge of the implementation is, the more accurate the results of the tests will be. For instance, in some implementations, the installation of routes in local routing tables may occur while the SPF is being calculated, dramatically impacting the time required to calculate the SPF.

7.4. Gaining an Understanding of the Tests Improves Measurements

One method which can be used to become familiar with the tests described in [2] is to perform the tests on an OSPF implementation for which all the internal details are available, such as GateD. While there is no assurance that any two implementations will be similar, this will provide a better understanding of the tests themselves.

8. Acknowledgements

Thanks to Howard Berkowitz, (hcb@clark.net) and the rest of the BGP benchmarking team for their support and to Kevin Dubray(kdubray@juniper.net) who realized the need of this draft.

9. References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC2119](#), March 1997.
- [2] Manral, V., "Benchmarking Methodology for Basic OSPF Convergence", [draft-bmwg-ospfconv-intraarea-00.txt](#), May 2002
- [3] Manral, V., "OSPF Convergence Testing Terminology and Concepts", [draft-bmwg-ospfconv-term-00.txt](#), My 2002
- [4] Bradner, S., McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", [RFC2544](#), March 1999.
- [5] Trotter, G., "Terminology for Forwarding Information Base (FIB) based Router Performance", [RFC3222](#), October 2001.

10. Authors' Addresses

Vishwas Manral
Netplane Systems
189 Prashasan Nagar
Road number 72
Jubilee Hills
Hyderabad, India

vmanral@netplane.com

Russ White
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709

riw@cisco.com

Aman Shaikh

University of California
School of Engineering
1156 High Street
Santa Cruz, CA 95064

aman@soe.ucsc.edu