

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: August 17, 2008

A. Atlas  
BT  
S. Bryant  
M. Shand  
Cisco Systems  
February 14, 2008

**Synchronisation of Loop Free Timer Values**  
**draft-atlas-bryant-shand-lf-timers-04**

Status of this Memo

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 becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/1id-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 August 17, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This draft describes a mechanism that enables routers to agree on a common convergence delay time for use in loop-free convergence.

## Requirements Language

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 [RFC2119](#) [1].

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Required Properties . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Mechanism . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Protocol Details . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	ISIS . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	OSPF . . . . .	<a href="#">5</a>
<a href="#">5.</a>	IANA considerations . . . . .	<a href="#">5</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">7.</a>	References . . . . .	<a href="#">6</a>
<a href="#">7.1.</a>	Normative References . . . . .	<a href="#">6</a>
<a href="#">7.2.</a>	Informative References . . . . .	<a href="#">6</a>
	Authors' Addresses . . . . .	<a href="#">6</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">8</a>



## **1. Introduction**

Most of the loop-free convergence mechanisms [3] require one or more convergence delay timers that MUST have a duration that is consistent throughout the routing domain. This time is the worst case time that any router will take to calculate the new topology, and to make the necessary changes to the FIB. The timer is used by the routers to know when it is safe to transition between the loop-free convergence states.

The time taken by a router to complete each phase of the loop-free transition will be dependent on the size of the network and the design and implementation of the router. It can therefore be expected that the optimum delay will need to be tuned from time to time as the network evolves.

Manual configuration of the timer is fraught for two reasons, firstly it is always difficult to ensure that the correct value is installed in all of the routers, and secondly, if any change is introduced into the network that results in a need to change the timer, for example due to a change in hardware or software version, then all of the routers need to be reconfigured to use the new timer value.

It is therefore desirable that a means be provided by which the convergence delay timer can be automatically synchronized throughout the network.

## **2. Required Properties**

The timer synchronization mechanism MUST have the following properties:

- o The convergence delay time must be consistent amongst all routers that are converging on the new topology.
- o The convergence delay time must be the highest delay required by any router in the new topology.
- o The mechanism must increase the delay when a new router is introduced to the network that requires a higher delay than is currently in use.
- o When the router that had the longest delay requirements is removed from the topology, the convergence delay timer value must, within some reasonable time, be reduced to the longest delay required by the remaining routers.



- o It must be possible for a router to change the convergence delay timer value that it requires.
- o A router which is in multiple routing areas, or is running multiple routing protocols may signal a different loop-free convergence delay for each area, and for each protocol.

How a router determines the time that it needs to execute each convergence phase is an implementation issue, and outside the scope of this specification. However a router that dynamically determines its proposed timer value must do so in such a way that it does not cause the synchronized value to continually fluctuate.

### **3. Mechanism**

The following mechanism is proposed.

A new information element is introduced into the routing protocol that specifies the maximum time (in milliseconds) that the router will take to calculate the new topology and to update its FIB as a result of any topology change.

When a topology change occurs, the largest convergence delay time required by any router in the new topology is used by the loop-free convergence mechanism.

If a routing protocol message is issued that changes the convergence delay timer value, but does not change the topology, the new timer value **MUST** be taken into consideration during the next loop-free transition, but **MUST NOT** instigate a loop-free transition.

If a routing protocol message is issued that changes the convergence timer value and changes the topology, a loop-free transition is instigated and the new timer value is taken into consideration.

The loop-free convergence mechanism should specify the action to be taken if a timer change (only) message and a topology change message are independently generated during the hold-off time. A suitable action would be to take the same action that would be taken if two uncorrelated topology changes occurred in the network.

All routers that support loop-free convergence **MUST** advertise a loop-free convergence delay time. The loop-free convergence mechanism **MUST** specify the action to be taken if a router does not advertise a convergence delay time.



## **4. Protocol Details**

This section describes the protocol changes needed to implement the timer synchronization function.

### **4.1. ISIS**

The controlled convergence timer value will be carried in a new Sub-TLV of the capability TLV as defined in [2] .

This draft defines one such SUB-TLV where the type is for the worst-case FIB compute/install time, the value is 16 bits and is specified in milliseconds; this gives a max value of about 65s.

The format of the Sub-TLV is as shown below.

Sub-TLV FIB-Convergence Timer

TYPE: <TBD>

Length: 2 octets

Value: <16-bit timer value expressed in milliseconds>

This MUST be carried in a capability TLV with the S-bit set to zero (indicating that it MUST NOT be leaked between levels).

### **4.2. OSPF**

A new type-10 opaque LSA (the controlled convergence LSA) will be defined as part of OSPF changes needed to define the loop-free convergence mechanism. This will consist of one or more TLVs. This draft defines one such TLV where the type is for the worst-case FIB compute/install time, the value is 16 bits and is specified in milliseconds; this gives a max value of about 65s.

## **5. IANA considerations**

There will be IANA considerations that arise as a result of this draft, but they are not yet determined.

## **6. Security Considerations**

If an abnormally large timer value is proposed by a router, there is a danger that the loop-free convergence process will take an excessive time. If during that time the routing protocol signals the





need for another transition, the loop-free transition will be abandoned and the default best case (traditional) convergence mechanism used.

It is still undesirable that the routers select a convergence delay time that has an excessive value. The maximum value that can be specified in the LSP/LSA is limited through the use of a 16 bit field to about 65 seconds. When sufficient implementation experience is gained, an architectural constant will be specified which sets the upper limit of the convergence delay timer.

## **7. References**

### **7.1. Normative References**

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Vasseur, J., "IS-IS Extensions for Advertising Router Information", [draft-ietf-isis-caps-07](#) (work in progress), February 2007.

### **7.2. Informative References**

- [3] Bryant, S. and M. Shand, "A Framework for Loop-free Convergence", [draft-bryant-shand-lf-conv-frmwk-03](#) (work in progress), October 2006.

## Authors' Addresses

Alia K, Atlas  
BT

Email: [akatlas@bt.com](mailto:akatlas@bt.com)

Stewart Bryant  
Cisco Systems  
250, Longwater Ave,  
Reading, RG2 6GB,  
UK

Email: [stbryant@cisco.com](mailto:stbryant@cisco.com)



Mike Shand  
Cisco Systems  
250, Longwater Ave,  
Reading, RG2 6GB,  
UK

## Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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, THE IETF TRUST 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.

## Intellectual Property

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).

## Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

