

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 3, 2018

S. Bryant
Huawei Technologies
A. Atlas
C. Bowers
Juniper Networks
August 02, 2017

Routing Timer Parameter Synchronization
draft-bryant-rtgwg-param-sync-03

Abstract

This document describes a mechanism for a link state routing protocol to coordinate the value of a routing timer parameter amongst routers in the flooding domain. The document also defines the solution to one specific case: the agreement of a common convergence timer value for use by routers in network convergence.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on February 3, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Requirements Language	3
3.	Overview of Mechanism	3
4.	Protocol Details	4
4.1.	ISIS	4
4.2.	OSPF	5
5.	Convergence Time	6
5.1.	Required Properties	6
5.2.	Definition of the Convergence Timer	7
6.	IANA considerations	7
6.1.	ISIS	7
6.2.	OSPF	8
6.3.	Routing Timer Parameter Synchronization Registry	8
7.	Security Considerations	8
8.	Acknowledgments	9
9.	Contributing Authors	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

There exist use cases where it desirable for a network to use a common value for a routing timer parameter across all nodes. In the past, these types of use case have been addressed by setting the parameter to a constant value in the protocol definition itself, or by requiring that the same value of the parameter be configured at every node.

Setting the routing timer parameter to a constant value in the protocol definition makes it difficult to change the parameter, since a change would require formal modification to the protocol. In practice, such a change is impractical, so the constant value needs to be chosen conservatively. This may impose a fundamental restriction on the eventual use of the protocol.

Manual or "static" configuration of the timer parameter is fraught

for two reasons. First, it is can be difficult to ensure that the correct, identical, value is installed in all of the routers. Second, if any change is introduced into the network that results in a need to change the value (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 parameter value. Such consistency may be ensured by deploying automated means such as enforcing the new value by invoking the management interface of all involved routers. For example, a central management entity may be responsible for communicating the new configuration value by means of vendor-specific command line interface (CLI), NETCONF[RFC6241], etc. This approach may be attractive if all involved nodes expose technology-agnostic and vendor-independent interfaces to modify a given network-wide configuration parameter.

This document describes a protocol extension that propagates a routing timer parameter throughout the flooding domain, which can be used as an alternative to the centralized approach described above. The method of choosing between one or more different advertised values, the flooding scope, and the action to be taken when the parameter changes MUST be provided in the definition of the parameter type.

This document also creates one parameter type: Convergence Timer intended for use in IP Fast-reroute applications [[RFC5714](#)] [[RFC5715](#)].

Note that this protocol is only intended to be used for the propagation of parameters needed to support the operation of the routing system. It MUST NOT be used as a general purpose parameter exchange protocol, and in particular it MUST NOT be used as a parameter negotiation protocol, since such use may degrade the ability of the underlying link-state routing protocol to carry out its essential purpose.

2. 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](#) [[RFC2119](#)].

3. Overview of Mechanism

Routing Timer parameter values that can be disseminated by means of the attribute defined in this specification MUST be defined as a configurable parameter or a default parameter in the corresponding specification.

A new information element is introduced into the routing protocol that specifies the parameter. Each router taking part in the parameter synchronization is expected to advertise a specific value of the parameter, which that router determines based mainly on considerations local to that router. In general, different routers in the flooding domain may advertise different values of the

parameter. How the values advertised by a router are determined is out of scope of this document.

A router receiving the parameter values advertised by all routers in the flooding domain will use a well-defined method to select the operational value of the parameter that it uses in the running of the protocol. All routers MUST use the same method applied to the same set of advertised parameter values. All routers SHALL therefore choose the same operational value for the parameter.

Note the operational value for the parameter selected SHOULD NOT directly affect the value for the parameter advertised by a router, since this introduces a form of negotiation leading to additional routing protocol traffic and possibly to instability in the routing protocol.

The method of selecting from a range of advertised parameter values MUST be provided in the parameter definition.

The definition of the parameter MUST specify the action to be taken when a new parameter value is advertised that would cause a change in the selected value.

The definition of the parameter MUST specify the action to be taken in the legacy/migration case, where not all routers advertise the parameter.

[4. Protocol Details](#)

This section describes the protocol extensions needed to implement this functionality.

4.1. ISIS

A new Routing Timer Parameter Synchronization (RTPS) sub-TLV is introduced into the IS-IS Router CAPABILITY TLV (TLV #242 defined in [RFC7981]). The setting of the S-bit in TLV #242 (indicating whether the parameter should be leaked between levels) MUST be included in the specific routing parameter definition.




Figure 1: Routing Timer Parameter Synchronization ISIS Sub-TLV

The Type (1 octet) of this sub-TLV to be assigned by IANA.

Length is variable (minimum value 5, multiple of 5 octets) and represents the total Length of the field.

Sub-Type consists of a one octet identifier of the timer type.

Duration is a 32 bit value representing is the length of the timer in milliseconds. This is capable of expressing a time in the range 1ms to just under 50 days.

4.2. OSPF

A new Routing Timer Parameter Synchronization TLV is defined for the OSPF Router Information LSA. This new TLV may be carried in a type 10 or type 11 OSPF Opaque LSA depending on the required flooding scope.

The specification of a the specific routing timer parameter MUST define the appropriate flooding scope(s) for that parameter.




Figure 2: Routing Timer Parameter Synchronization OSPF TLV

The Type (2 octets) of this sub-TLV to be assigned by IANA.

Length is variable (minimum value 5, multiple of 5 octets) and represents the total Length of the field.

Sub-Type consists of a one octet identifier of the timer type.

Duration is a 32 bit value representing is the length of the timer in milliseconds. This is capable of expressing a time in the range 1ms to just under 50 days.

5. Convergence Time

Routers running a fast-reroute mechanism such as Maximally Redundant Tree (MRT) [RFC7812] fast re-route require a network wide convergence time value so that know how long they need continue using the repair path before it is safe to use the base path. This time is set to be the worst case time that any router will take to calculate the new topology, and to make the necessary changes to the FIB.

The time taken by a router to complete each phase of the 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.

5.1. Required Properties

The Convergence Time mechanism MUST have the following properties:

- o The operational convergence delay time MUST be consistent among all routers that are converging on the new topology.
- o The operational convergence delay time MUST be the highest delay time advertised 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.

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 delay value must do so in such a way that it does not cause the synchronized value to continually fluctuate.

5.2. Definition of the Convergence Timer

It is RECOMMENDED that the routing convergence timer be limited to a maximum of 60 seconds.

The routing convergence timer value selected is the largest value advertised.

If a routing protocol message is issued that changes the Convergence Timer value, but does not change the topology, the new timer value MUST be taken into consideration during the next network transition, but MUST NOT instigate a new transition.

If a routing protocol message is issued that changes both the Convergence Timer value and the topology, a transition is instigated and the new timer value MUST be taken into consideration.

The convergence mechanism MUST 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 router running ISIS that supports convergence timer synchronization SHOULD advertise the Routing Timer Parameter Synchronization sub-TLV with the value of the Convergence Timer in the Duration field of this sub-TLV. The S-bit is set to zero indicating that the Convergence Timer RPTS sub-TLV MUST NOT be leaked between levels.

A router running OSPF that supports convergence timer synchronization SHOULD advertise the Routing Timer Parameter Synchronization TLV with the value of the Convergence Timer in the Duration field of this TLV. The TLV SHOULD only be carried in the type 10 Opaque LSA which prevents propagation outside the OSPF area.

[6.](#) IANA considerations

[6.1.](#) ISIS

IANA is requested to allocate a new Sub-TLVs for TLV 242 from the IS-IS TLV Codepoints name space.

Value	Description	Reference
TBD	Routing Timer Parameter Synchronization	This Document

[6.2.](#) OSPF

IANA is requested to allocate a new OSPF Router Information (RI) TLV from the Open Shortest Path First (OSPF) Parameters name space

Value	TLV Name	Reference
TBD	Routing Timer Parameter Synchronization	This document

A value in the range 12 to 32767 is requested.

[6.3.](#) Routing Timer Parameter Synchronization Registry

IANA is requested to create a new Routing Timer Parameter Synchronization Registry within its own name space, and to allocate one value from it.

Value	Name	Reference
0	Reserved	This document
1	Convergence Timer	This document
2..255	Reserved	This document

Allocations within this registry require IETF Consensus. This link state protocol extension MUST NOT be used for any purpose other than one associated with the routing system timer parameters.

[7.](#) Security Considerations

The introduction of this parameter advertising mechanism does not introduce a significant vulnerability into the base routing protocol and is secured in exactly the same way as the other TLVs that are carried.

In specifying a new parameter, consideration must be given to the impact of the additional parameter, and in particular the rate of change of that parameter, on the dynamics of the link-state routing protocol in use. In the specific case of the Convergence Timer, the amount of data being carried and the rate of change of the parameter value will have a negligible impact on the link-state routing protocol in use.

A rouge router deliberately introducing an anomalous parameter value is just as capable of introducing many other anomalies into the routing domain.

As far as possible, care should be taken to validate that the parameter is reasonable.

In the specific case of the Convergence Time RTPS, the following considerations apply.

If an abnormally large timer value is proposed by a router, there is a danger that the convergence process will take an excessive time. If during that time the routing protocol signals the need for another transition, the transition will be abandoned and the default best case (traditional) convergence mechanism used.

It is RECOMMENDED that implementations prohibit the configuration of a router convergence timer value in excess of 60 seconds.

8. Acknowledgments

The authors thank Les Ginsberg and the other authors of [[I-D.ietf-isis-segment-routing-msd](#)] and [[I-D.ietf-ospf-segment-routing-msd](#)], Mohamed Boucadair for their review comments and proposed text, and Tom Petch for his review comments.

9. Contributing Authors

Mike Shand
Independent
mike@mshand.org.uk

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", [RFC 7981](#), DOI 10.17487/RFC7981, October 2016, <<http://www.rfc-editor.org/info/rfc7981>>.

10.2. Informative References

Internet-Draft

Routing Timer Parameter Sync

August 2017

[I-D.ietf-isis-segment-routing-msd]

Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling MSD (Maximum SID Depth) using IS-IS", [draft-ietf-isis-segment-routing-msd-04](#) (work in progress), June 2017.

[I-D.ietf-ospf-segment-routing-msd]

Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak, "Signaling MSD (Maximum SID Depth) using OSPF", [draft-ietf-ospf-segment-routing-msd-05](#) (work in progress), June 2017.

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", [RFC 5714](#), DOI 10.17487/RFC5714, January 2010, <<http://www.rfc-editor.org/info/rfc5714>>.

[RFC5715] Shand, M. and S. Bryant, "A Framework for Loop-Free Convergence", [RFC 5715](#), DOI 10.17487/RFC5715, January 2010, <<http://www.rfc-editor.org/info/rfc5715>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", [RFC 6241](#), DOI 10.17487/RFC6241, June 2011, <<http://www.rfc-editor.org/info/rfc6241>>.

[RFC7812] Atlas, A., Bowers, C., and G. Enyedi, "An Architecture for IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)", [RFC 7812](#), DOI 10.17487/RFC7812, June 2016, <<http://www.rfc-editor.org/info/rfc7812>>.

Authors' Addresses

Stewart Bryant
Huawei Technologies

Email: stewart.bryant@gmail.com

Alia Atlas

Juniper Networks

Email: akatlas@gmail.com

Chris Bowers

Juniper Networks

Email: cbowers@juniper.net

Bryant, et al.

Expires February 3, 2018

[Page 10]