

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: September 6, 2012

W. Lu  
A. Tian  
Ericsson  
March 5, 2012

ISIS Transaction TLV  
draft-lu-isis-transaction-tlv-00

## Abstract

ISIS local updates may require multiple LSPs to convey. Receiving routers, whose decision processes are without such knowledge, may generate incorrect routing table updates based on the partial set of LSPs it receives and hence the traffic outage before they are corrected by another run of the decision process. This memo describes a method that makes the decision process more informed so that the interim results can be minimized or avoided.

## 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 September 6, 2012.

## Copyright Notice

Copyright (c) 2012 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

Internet-Draft

ISIS Transaction TLV

March 2012

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

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language . . . . .</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Acronyms . . . . .</a>	<a href="#">3</a>
<a href="#">1.3.</a>	<a href="#">Implicit Purging . . . . .</a>	<a href="#">3</a>
<a href="#">1.4.</a>	<a href="#">PDU Transaction Knowledge . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Transaction TLV . . . . .</a>	<a href="#">5</a>
<a href="#">2.1.</a>	<a href="#">LSP Transaction Set . . . . .</a>	<a href="#">5</a>
<a href="#">2.2.</a>	<a href="#">TLV Format . . . . .</a>	<a href="#">5</a>
<a href="#">2.3.</a>	<a href="#">T-TLV Count . . . . .</a>	<a href="#">6</a>
<a href="#">2.4.</a>	<a href="#">Transaction ID . . . . .</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">Operation . . . . .</a>	<a href="#">7</a>
<a href="#">3.1.</a>	<a href="#">Originator . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.</a>	<a href="#">Receiver . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.1.</a>	<a href="#">Opening the Transaction . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.2.</a>	<a href="#">Invalid Transaction . . . . .</a>	<a href="#">8</a>
<a href="#">3.2.3.</a>	<a href="#">Processing T-TLV . . . . .</a>	<a href="#">8</a>
<a href="#">3.2.4.</a>	<a href="#">Closing the Transaction . . . . .</a>	<a href="#">8</a>
<a href="#">3.2.5.</a>	<a href="#">Aborting the Transaction . . . . .</a>	<a href="#">9</a>
<a href="#">3.2.6.</a>	<a href="#">Exit Transaction . . . . .</a>	<a href="#">9</a>
<a href="#">3.2.7.</a>	<a href="#">Timer Expiry . . . . .</a>	<a href="#">9</a>
<a href="#">3.2.8.</a>	<a href="#">TID Wrap . . . . .</a>	<a href="#">9</a>
<a href="#">4.</a>	<a href="#">Multiple Transaction Sets . . . . .</a>	<a href="#">9</a>
<a href="#">5.</a>	<a href="#">Use Cases . . . . .</a>	<a href="#">10</a>
<a href="#">5.1.</a>	<a href="#">Avoid Unwanted Purging . . . . .</a>	<a href="#">10</a>
<a href="#">5.2.</a>	<a href="#">Allow reordering of TLVs in GR case . . . . .</a>	<a href="#">10</a>
<a href="#">5.3.</a>	<a href="#">Help Precise SPF Scheduling . . . . .</a>	<a href="#">10</a>
<a href="#">5.4.</a>	<a href="#">Other than SPFs . . . . .</a>	<a href="#">11</a>
<a href="#">6.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">8.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">11</a>
<a href="#">9.</a>	<a href="#">References . . . . .</a>	<a href="#">12</a>
<a href="#">9.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">12</a>
<a href="#">9.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">12</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">12</a>

## [1.](#) Introduction

Link state protocols run on the knowledge of the entire topology. Incomplete topology information, even temporary, can result in traffic outage or routing loop. While transitional routing changes are inevitable and common to both OSPF [[RFC2328](#)] and ISIS [[RFC1195](#)][ISO.10589.1992], impacts to unchanged network connectivity are unnecessary and should be minimized if not totally avoidable.

### [1.1.](#) 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 [RFC 2119](#) [[RFC2119](#)].

### [1.2.](#) Acronyms

IS-IS - Intermediate System to Intermediate System

OSPF - Open Shortest Path First

TLV - Type Length Value

PDU - Protocol Data Unit

LSP - Link State PDU

SPF - Shortest Path First

### [1.3.](#) Implicit Purging

Compared to OSPF, these impacts are unique to ISIS. There are two reasons. One is that ISIS LSPs use implicit TLV purging. Although LSPs do have age field which can be used for purging purpose, ISIS does not have the age granularity down to TLV level which is the atomic unit of ISIS link state information. If for some reason a TLV

needs to be relocated to a different LSP fragment (e.g. TLV-B in Figure 1 and Figure 2), this TLV can be perceived as being purged from the original LSP fragment. And if the receiving ISIS starts its decision process before it sees the second LSP fragment, the reachability via this TLV, if any, will be lost.

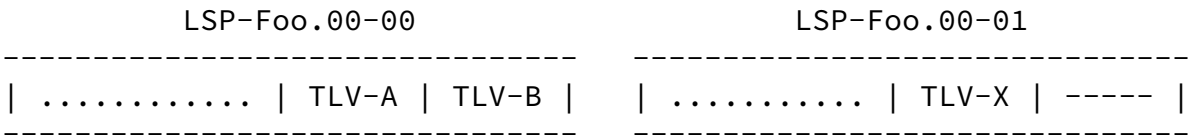


Figure 1: Frag 00 Almost Full

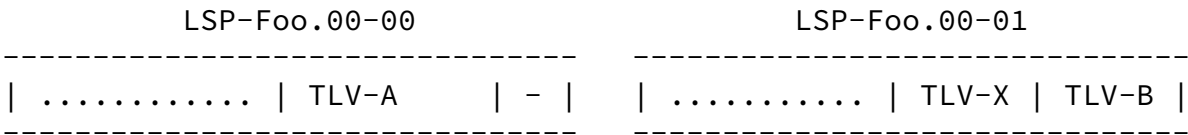


Figure 2: TLV-A grow out TLV-B

The second reason is that operating directly on data link layer, ISIS cannot extend the LSP size beyond the MTU limit as opposed to OSPF which can leverage the IP fragmentation capability to extend its LSU size. The consequence is that when an LSP fragment is full or nearly full, if some of its TLVs need to expand, they will have to be relocated to other LSP fragment. Alternatively some other TLVs can be moved out of this LSP fragment to make room for the needy. Either way an implicit purge condition is created.

1.4. PDU Transaction Knowledge

The above issues can be mitigated if the receiving routers are provided with LSP transaction information. In other words, if the receivers know how many LSPs they should expect from a particular

originating Intermediate System, so that they acquire complete topology updates from that System, the receivers should be able to avoid running their decision process based on the incomplete transitional link state information.

There can be many ways to accomplish the purpose. However to be practical the solution should meet following requirements:

1. It must be backward compatible. Adding a new TLV can easily fulfill this;
2. The TLV should be simple and short, that it does not take significant LSP space;
3. The solution should be fallback-able. That is, in case of errors or mistakes, it can fallback to the operation state without such solution;

4. It can be implemented easily without adding much burden on the originator and its update process. In particular it should not delay or change the timing of LSP flooding;
5. On the receiver side, the new logic should be simple and can be easily integrated to the existing logic, such as SPF scheduler. Performance wise, the new addition should be negligible.

This document describes a transaction knowledge based TLV that can be used by the receiving routers to make informed decision.

## [2.](#) Transaction TLV

A new TLV is introduced to indicate that the carrying LSP needs additional LSPs to complement. For example, in Figure 2 LSPs "00" and "01" both have to be included in the SPF to reflect the correct change. If the receiver kicks off its SPF right after receiving LSP "00" and before seeing LSP "01", the reachability pertaining to TLV-B will be incorrectly removed, cause temporary traffic loss. The TLV is called Transaction TLV as it provides the transaction knowledge of the changes in which a set of LSPs are involved.

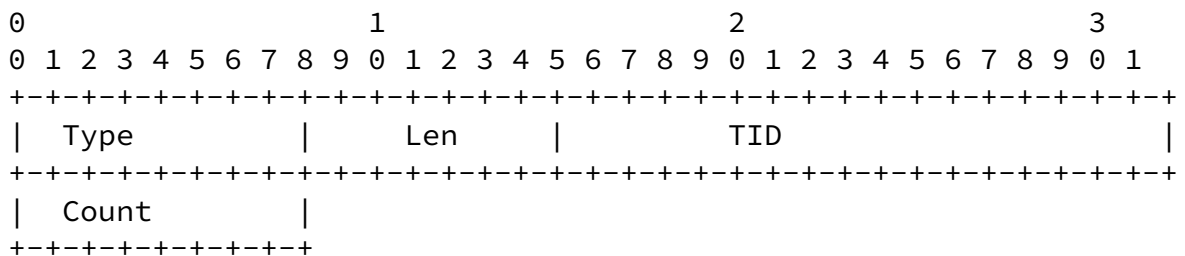
## 2.1. LSP Transaction Set

LSPs which are coherent in contents are called LSP Transaction Set. These LSPs must be processed atomically by the receiver's decision process to avoid incorrect result. In Figure 2, LSPs "00" and "01" form the LSP transaction set.

The LSP transaction set is a subset of all LSPs originated by an IS. In other words, the transaction set belongs to a single IS, and shares the same LSP ID which is the System ID plus the pseudo node ID.

## 2.2. TLV Format

Figure 3 describes the Transaction TLV data fields:



### Figure 3: Transaction TLV Format

Type

1 byte, value TBD (T-TLV);

Len

1 byte, 4 or 5;

TID

4 bytes, Transaction ID;

Count

1 byte, optional.

The TLV can appear at most once in an LSP (fragment). Each LSP in an LSP transaction set is encoded with the same TLV except the last LSP which MUST include the "count" in this TLV to tell the total number LSPs in this transaction set.

### [2.3.](#) T-TLV Count

The count field is of 1 byte size, same as the size for LSP number field. The first "count-1" LSPs use the shorter T-TLV which has length 4. The last LSP use the longer T-TLV which contains the count field which counts in itself. This design is to make the originator easy to encode the TLV without having to know the count beforehand.

### [2.4.](#) Transaction ID

Along the time, an IS may generate and flood a number of LSP transaction sets. To differentiate one set from another, a monotonically increased Transaction ID (TID) is used. Each IS maintains and manages its TID. If an IS is also a DIS in one or more of its interfaces, each pseudo node has its own TID which is independent of the TID of the non-pseudo node, or other pseudo node.

In case of TID conflict (due to race condition in flooding, or errors), the higher TID invalidates the lower TID. When a TID reaches the maximum, the TID wrap mechanism is used, which is detailed in [Section 3.2.8](#).

## [3.](#) Operation

### [3.1.](#) Originator

When an IS decides that its update process will have to use multiple LSPs to convey some information atomically, it labels these LSPs with the T-TLV and follows the procedures below:

1. The T-TLV should be used for SPF-sensitive changes only;
2. It starts a new TID number which is per LSP ID based. How to choose the initial TID number is a local decision though the natural choice would be 1. The number MUST be incremented from the existing one, and monotonically increased ever since;
3. T-TLV is recommended to be added to the end of an LSP. The TID is stored in the LSP set space (containing all LSP fragments), as opposed to the individual LSP space;
4. Repeat step 3 until there is no more to pack, at which time a T-TLV with count field is inserted to this last LSP in the transaction set.

### 3.2. Receiver

Per ISIS protocol nature, if the receiver does not understand and support the T-TLV, the TLV is silently ignored. This ensures the backward compatibility.

Otherwise the receiving IS enters into the transaction procedure.

An IS will not engage its decision process into such procedure for T-TLVs whose carrying LSPs are already installed in the database. In other words, the procedure is activated only upon the receiving of the T-TLVs whose carrying LSPs are new.

#### 3.2.1. Opening the Transaction

When a T-TLV is received, the receiving IS enters the LSP transaction procedure.

##### Type

The TID is recorded to indicate that the current TID is active;

##### Len

A protection timer is started to prevent the error case where the transaction cannot close in time;

LSPs in a transaction set may not arrive in the order they are sent. Whichever arrives the first opens the transaction;

Count

The transaction record (TID/count/status) is maintained under the LSP set space (SystemID + PseudoNodeID).

### 3.2.2. Invalid Transaction

The transaction is invalid, and MUST be aborted (exit) if:

Type

TID is outdated. This occurs if a higher TID is found, or the same TID is closed in the past transaction;

Len

more than one T-TLV is found in an LSP;

TID

more than one T-TLV with count field is found in an LSP transaction set;

### 3.2.3. Processing T-TLV

When a T-TLV is received, following rules apply:

Type

Open the transaction if not yet, also increment the corresponding local count. If the received TLV contains the count, note down the announced count.

Len

If the local count equals the announced count, close the transaction.

### 3.2.4. Closing the Transaction

If the received T-TLV causes the local count to match the announced count:

Type

Change the current transaction TID from active to closed;

Len

Cancel the protection timer;

Len

Exit the transaction.

Note that any T-TLV can close the transaction as long as it causes the match of counters. Implementation should not assume that the T-TLV with count field comes the last.

#### [3.2.5.](#) Aborting the Transaction

Any error condition can abort the current transaction. The handling procedure is the same as the one in [Section 3.2.4](#).

#### [3.2.6.](#) Exit Transaction

A transaction can be terminated normally (closing) or abnormally due to error conditions.

Closing and aborting the transaction are technically the same operation. The difference is that closing the transaction fulfills the purpose of T-TLVs for avoiding unnecessary packet loss.

Either way after the transaction is terminated, the decision process MUST no longer block its SPF and should start the computation immediately or follow whatever SPF scheduling mandates.

#### [3.2.7.](#) Timer Expiry

The expiry of the protection timer indicates that some transaction error has occurred. The receiving IS MUST abort the transaction.

The length of the timer is a local decision.

#### [3.2.8.](#) TID Wrap

When a TID reaches the maximum (0xFFFFFFFF), the originating IS will have to refrain from using T-TLV for LSP maximum age (21 minutes usually). The logic is similar to that of LSP sequence number wrapping.

### [4.](#) Multiple Transaction Sets

The count field is of 1 byte size, same as the size for LSP number field. The first "count-1" LSPs use the shorter T-TLV which has length 4. The last LSP use the longer T-TLV which contains the count field which counts in itself. This design is to make the originator

easy to encode the TLV without having to know the count beforehand.

## [5.](#) Use Cases

The count field is of 1 byte size, same as the size for LSP number field. The first "count-1" LSPs use the shorter T-TLV which has length 4. The last LSP use the longer T-TLV which contains the count field which counts in itself. This design is to make the originator easy to encode the TLV without having to know the count beforehand.

### [5.1.](#) Avoid Unwanted Purging

The unwanted purging described in [Section 1.3](#) can be avoided using T-TLVs. The originator can add T-TLV to LSP-Foo.00-00 and T-TLV (count=2) to LSP-Foo.00-01. The receivers will withhold the SPF till both LSPs are received. The missing TLV-B in the first LSP as shown in Figure 2 will not be treated as an implicit purging, as it will be found in the second LSP.

### [5.2.](#) Allow reordering of TLVs in GR case

If an IS advertises lots of redistributed routes in its LSPs, it is not trivial to maintain its TLV (like TLV 135) orders.

This is especially true when an IS has just gone through the graceful restart process. Because the RIB does not necessarily supply the redistributed routes the same order in the Pre-Restart time, reconstruct LSPs will result in LSPs with TLVs reordered.

And if the number of redistributed routes is high, they spread over multiple LSPs. When the set of LSPs reaches other ISes, the same issue of 5.1 can arise, even if there is no change to redistributed routes at all.

It is not impossible for the originating IS to use sophisticated means to keep those TLVs in their original order. Nevertheless this issue can easily be addressed with the T-TLVs.

The restarting IS can add T-TLVs to all LSPs that are subject to TLV reordering, and transmit them upon exit of its graceful restart process. Thus the receiving ISes will not mistakenly purge some IS

external reachability prefixes.

### [5.3.](#) Help Precise SPF Scheduling

As a link state protocol, ISIS has to two conflict goals. One is to be fast responsive to the network changes. The other is the network stability.

If the SPF is scheduled too swiftly, the system (and even the

network) can melt down for some storm activities. On the other hand if there is lot of delays, the network becomes too slow adapting changes.

For example, if an IS receives a burst of redistributed routes from BGP, it may send out dozens of LSPs for advertising all those routes. The receiving ISes, upon received first several LSPs, usually start the decision process to compute the new routing table. The routing table is incomplete and will be soon overwritten by another SPF run.

If the number of routes (and hence LSPs) is high, most such SPF runs will be useless and wasteful. Only the last SPF will contribute to the final and correct routing table.

The T-TLV if used will provide guidance to the receiving ISes to run SPF only when all LSPs are in place. This SPF is equivalent to the final SPF mentioned above. Therefore it saves a lot of SPF runs and network churns. What is more is that the T-TLV driven SPF can be kick started immediately, compared to the final SPF which usually has some amount of delay.

### [5.4.](#) Other than SPFs

The T-TLV may also be used for some non-SPF related operation. For example, the receiving ISes may choose to defer its TE database uploading process until all LSPs that carry the TE information are received.

## [6.](#) Security Considerations

This proposal does not introduce additional issues on security

condition.

## [7.](#) IANA Considerations

A new ISIS T-TLV is introduced. The type is TBD by IANA.

## [8.](#) Acknowledgements

TBD

## [9.](#) References

Lu & Tian

Expires September 6, 2012

[Page 11]

---

Internet-Draft

ISIS Transaction TLV

March 2012

### [9.1.](#) Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### [9.2.](#) Informative References

[ISO.10589.1992]

International Organization for Standardization,  
"Intermediate system to intermediate system intra-domain-  
routing routine information exchange protocol for use in  
conjunction with the protocol for providing the  
connectionless-mode Network Service (ISO 8473)",  
ISO Standard 10589, 1992.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and  
dual environments", [RFC 1195](#), December 1990.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.

## Authors' Addresses

Wenhu Lu  
Ericsson

300 Holger Way  
San Jose, California 95134  
USA

Email: [Wenhu.Lu@ericsson.com](mailto:Wenhu.Lu@ericsson.com)

Albert Tian  
Ericsson  
300 Holger Way  
San Jose, California 95134  
USA

Email: [Albert.Tian@ericsson.com](mailto:Albert.Tian@ericsson.com)