

Network Working Group  
Internet-Draft  
Updates: [6126](#) (if approved)  
Intended status: Experimental  
Expires: January 4, 2018

G. Chouasne  
J. Chroboczek  
PPS, University of Paris-Diderot  
July 3, 2017

**TOS-Specific Routing in Babel**  
**draft-chouasne-babel-tos-specific-00**

Abstract

This document describes an extension to the Babel routing protocol to support TOS-specific routing. This version is using mandatory sub-TLVs.

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 January 4, 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

<a href="#">1.</a>	Introduction and background . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Protocol Operation . . . . .	<a href="#">3</a>
<a href="#">2.1.</a>	Conceptual Description . . . . .	<a href="#">3</a>
<a href="#">2.2.</a>	Data Structures . . . . .	<a href="#">3</a>
<a href="#">2.3.</a>	ToS-specific sub-TLV . . . . .	<a href="#">4</a>
<a href="#">2.4.</a>	ToS-specific Messages . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Interaction . . . . .	<a href="#">6</a>
<a href="#">3.1.</a>	ToS interpretation . . . . .	<a href="#">6</a>
<a href="#">3.2.</a>	IP version . . . . .	<a href="#">6</a>
<a href="#">3.3.</a>	Interaction with non-Tos-specific Babel . . . . .	<a href="#">6</a>
3.4.	Interaction with the Source-specific routing extension . . . . .	7
<a href="#">3.5.</a>	Forwarding Behavior . . . . .	<a href="#">7</a>
<a href="#">4.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">5.</a>	Security considerations . . . . .	<a href="#">8</a>
<a href="#">6.</a>	References . . . . .	<a href="#">8</a>
<a href="#">6.1.</a>	Normative References . . . . .	<a href="#">8</a>
<a href="#">6.2.</a>	Informative References . . . . .	<a href="#">9</a>
	Authors' Addresses . . . . .	<a href="#">9</a>

## [1.](#) Introduction and background

The Type of Service (ToS) or Differentiated Services Code Point (DSCP) is a field of the IPv4 and IPv6 headers that can be used to request different per-hop behaviour when forwarding IP packets with identical source and destination. The most common application of the ToS field is to request different queueing policies (priority, drop probability, etc.) from an AQM.

A slightly less common application of the ToS field is to take it into account in addition to the destination address when performing a routing decision. For example, a router that has a low-latency default route with high monetary cost might announce it with a "low-latency" ToS, and thus avoid carrying ordinary best-effort traffic over the expensive route. A router that performs ToS-specific routing maintains a routing table which instead of being merely indexed by destination prefixes is indexed by pairs of a prefix and a ToS value. In order to be routed according to a given entry in the routing table, a packet must match not only the destination prefix but also the ToS value. ToS-specific forwarding is described in more detail in [Section 3.5](#).

Just like ordinary routes, ToS-specific routes can be installed manually but are more commonly installed by a dynamic routing protocol. This document specifies an extension to the Babel routing protocol [[BABEL](#)] that allows it to carry ToS-specific routes. This extension considers the ToS field as an opaque value that is only



compared for equality, but ignores the two bits that have been reserved for Explicit Congestion Notification (ECN) [[RFC3168](#)], and is therefore compatible with the definition of the ToS field used by DSCP [DSCP].

The extended protocol remains interoperable with the unextended Babel protocol in a sense made precise in [Section 3.3](#).

## **[2. Protocol Operation](#)**

### **[2.1. Conceptual Description](#)**

ToS-specific routes are routes that are defined by the same informations as classic routes, plus a ToS. This extension adds ToS-specific routes to the non-ToS-specific routes handled by the original Babel protocol.

This extension doesn't change the processing of non-ToS-specific routes. A node implementing this extension behaves exactly like a node implementing the original protocol as far as non-ToS-specific routes are concerned.

ToS-specific routes are treated analogously to non-ToS-specific routes, except for an additional ToS field in a number of data structures ([Section 2.2](#)) and in the encoding of Update and Request TLVs, which are augmented with a sub-TLV carrying the ToS. This sub-TLV has the mandatory bit set ([\[BABEL\] Section 4.4](#)), and hence, Updates and Requests for ToS-specific routes will be ignored by nodes implementing only the original protocol. Therefore, ToS-specific routes can only consist of a sequence of nodes implementing this extension.

### **[2.2. Data Structures](#)**

An implementation of Babel that implements this extension needs to add a ToS field to a number of data structures it maintains.

#### **[2.2.1. The Source Table](#)**

The source table maintained by any Babel node, as described in [\[BABEL\]](#), Section 3.2.4, is extended with the following field:

- o the ToS specifying the ToS of packets to which this entry applies.



**2.2.2. The Table of Pending Requests**

The table of pending requests, maintained by any Babel node, as described in [BABEL], Section 3.2.4, is extended with the following field:

- o the ToS being requested.


The table is now indexed by (prefix, ToS) pairs.

**2.3. ToS-specific sub-TLV**

This extension defines a new sub-TLV that adds a ToS field to a Babel packet. It turns regular Updates, Route Requests and Seqno Requests into ToS-specific Updates, Route Requests and Seqno Requests. Other TLVs MUST NOT be sent with this sub-TLV attached. If any are received, they will be silently ignored, as described in Section 4.4 of [BABEL].

A node MUST send ToS-specific Update, Route Request and Seqno Request if the route they advertise is ToS-specific. Otherwise, a node MUST send non-ToS-specific Update, Route Request and Seqno Request.

The ToS sub-TLV is defined as follow:



Fields :

Type Set to TBD to indicate a ToS-specific mandatory sub-TLV.

Length The length of the body, exclusive of the Type and Length fields. It MUST be set to 1.

ToS The ToS value of the ToS-specific route. This MUST be a multiple of 4 (i.e. the two least-significant bits MUST be 0).

The value TBD has the mandatory bit set to 1 and this sub-TLV must therefore be handled in accordance with Section 4.4 of [BABEL].



## **2.4. ToS-specific Messages**

This section describes the handling of ToS-specific messages.

### **2.4.1. Updates**

Whenever a ToS-specific Update is sent by a node implementing this extension, the source table MUST be updated as follows : if an entry indexed by (prefix, plen, router-id, ToS), exists, it MUST be modified as described in section 3.7.3 of [[BABEL](#)]. Otherwise, a new entry is created with value (prefix, plen, router-id, ToS, seqno, metric).

### **2.4.2. Requests**

Whenever a node implementing this extension sends a ToS-specific Request, it MUST add its content as an entry to the Pending Requests Table, as described in section 3.2.7 of [[BABEL](#)], with the suitable ToS.

#### **2.4.2.1. Route Requests**

When a node implementing this extension receives a ToS-specific Route Request for a prefix (prefix, plen) and a ToS t, it checks its route table for a selected route with this prefix, plen, and with ToS t in the corresponding source. If such a route doesn't exist, it MUST send a ToS-specific retraction for that prefix.

When a node implementing this extension receives a wildcard Route Request, it SHOULD send a full routing table dump, as described in Section 3.8.1.1 of [[BABEL](#)]. Therefore, it SHOULD also dump its ToS-specific routes.

#### **2.4.2.2. Seqno Requests**

When a node receives a Seqno Request for a prefix (prefix, plen) with a ToS-specific sub-TLV with ToS t, it MUST send a ToS-specific update with ToS t for a selected route specified by the Plen, Prefix and Source ToS field, with either a router-id different from what is specified by the Router-Id field, or a Seqno no less than what is specified by the Seqno field. If there is no such route in the Route Table, it MUST forward the seqno request according to the rules defined in Section 3.8.1.2 of [[BABEL](#)].





### **3. Interaction**

This section describes the interaction of this extension with other protocols and other versions of Babel.

#### **3.1. ToS interpretation**

Several interpretations have been defined for the ToS field.

This extension ignores the last two bits of the field, while the first six bits of the ToS field are opaque to this extension. Hence, it is compatible with all extensions that don't use the last two bits of the field. This is the case of all interpretations known to us. In particular, it is compatible with the Differentiated Service Field interpretation (DSCP), defined in [RFC 2474](#), the original interpretation described in [RFC 791](#), and the ECN protocol, described in [RFC 3168](#).

In the case of the DSCP interpretation, one must note that a packet with a DSCP value will follow the route with the ToS being four times this value.

Details on how packets are forwarded are provided in [Section 3.5](#).

#### **3.2. IP version**

This protocol extension is compatible with IPV4 and IPV6. AE fields MUST be filled accordingly, as described in Section 4.1.3 of [[BABEL](#)].

#### **3.3. Interaction with non-ToS-specific Babel**

The encoding of ToS-specific Updates and Requests is using a reserved sub-TLV. This means that, as defined in section 4.4 of [[BABEL](#)], they will be ignored by nodes that don't implement this extension, which means that non-ToS-specific nodes will not treat ToS-specific routes.

In a topology of routers implementing ToS-specific Babel and non-ToS-specific Babel, all packets will reach their destination if it is reachable. Non-ToS-specific packets will follow the same routes as if ToS-specific routers were non-ToS-specific routers. ToS-specific packets may switch to a non-ToS-specific route if and only if there is no route for the required ToS.

This extension uses a mandatory bit on each sub-TLV. It is necessary that this bit is set to 1 for all ToS-specific sub-TLV to avoid loops, as shown in the following example.



Consider two neighboring routers A and B, A implementing the ToS-specific Babel extension, and B implementing just the base Babel protocol. Suppose that A has a ToS-specific route for a prefix P and ToS t.

B will receive an update from A with this ToS-specific route. Suppose that B reads the update TLV but drops the ToS-specific sub-TLV. It will now forwards packets for P through A.

B will then send an update for the route to (P) to A. A will take B as next-hop for the matching packets.

When a packet with no ToS arrives to A or B, it will travel between A and B indefinitely, as shown in the figure below.



**3.4. Interaction with the Source-specific routing extension**

The ToS-specific routing extension is compatible with the source-specific extension. A node implementing ToS-specific routing and source-specific routing handles ToS-specific routes and source-specific routes. To achieve this, data structures are extended with a ToS field and a source field.

In principle, a node could handle routes that are both ToS- and source-specific. In that case, Updates and Requests advertising source-and-ToS-specific routes would need to be sent with both the source prefix sub-TLV and the ToS sub-TLV, and a route preference ordering for source-and-ToS-specific packets would need to be defined. Until such an ordering is defined, updates with both the source prefix and ToS sub-TLVs MUST NOT be sent, and, if received, MUST be silently dropped.

**3.5. Forwarding Behavior**

When a packet for an address A arrives to a node, it may match several routes. The node MUST choose the route with the destination-first ordering.

This ordering only considers routes that have either no ToS or the Tos of the packet (ignoring the last two bits). It forwards a packet through the route with the most specific prefix. If there are two routes with the same prefix, then one of them has no ToS and the



other has the ToS of the packet (ignoring the last two bits). The packet is following the latter. If there is no route with a matching prefix, the packet is dropped.

When a ToS-specific route is retracted while the router has another non-ToS-specific route, packets should fall back to this non-ToS-specific route as fast as possible, to avoid more delay. Hence, it is RECOMMENDED to use the algorithm described in Section 3.5.5 of [BABEL]

A Babel implementation of this extension must ensure that these semantics are observed. If they aren't, the implementation MUST disambiguate the routing tables by using a suitable algorithm (for example the algorithm of Boutier [SS-ROUTING]).

It may be the case that the forwarding plane cannot handle some ToS values. In that case, routes with these values as ToS MUST NOT be selected and therefore MUST NOT be announced.

**4. IANA Considerations**

IANA is instructed to add the following entry to the "Babel sub-TLV type" registry:

```

+-----+-----+-----+
| Type | Name          | Reference      |
+-----+-----+-----+
| TBD  | ToS-specific TLV | (this document) |
+-----+-----+-----+

```

**5. Security considerations**

The extension defined in this protocol defines three new sub-TLVs for defined TLVs. This extension doesn't alterate any of the security properties of the base protocol.

**6. References**

**6.1. Normative References**

[BABEL] Chroboczek, J., "The Babel Routing Protocol", [draft-chroboczek-babel-rfc6126bis-03](#) (work in progress), July 2016.



## **6.2. Informative References**

[SS-ROUTING]

Boutier, M. and J. Chroboczek, "Source-Specific Routing",  
August 2014.

In Proc. IFIP Networking 2015. A slightly earlier  
version is available online from [http://arxiv.org/  
pdf/1403.0445](http://arxiv.org/pdf/1403.0445).

### Authors' Addresses

Gwendoline Chouasne  
PPS, University of Paris-Diderot  
Case 7014  
75205 Paris Cedex 13  
France

Email: [gwendoline.chouasne-guillon@ens.fr](mailto:gwendoline.chouasne-guillon@ens.fr)

Juliusz Chroboczek  
PPS, University of Paris-Diderot  
Case 7014  
75205 Paris Cedex 13  
France

Email: [jch@irif.fr](mailto:jch@irif.fr)



