

Maintaining more than 255 circuits in IS-IS
<[draft-ietf-isis-wg-255adj-01.txt](#)>

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 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/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

Abstract

This draft describes an optional implementation technique within IS-IS [[IS090](#), [Cal90a](#), [Cal90b](#)] used today by several ISPs for routing within their clouds. IS-IS is an interior gateway routing protocol developed originally by OSI and used with IP extensions as IGP. This draft describes how to effectively bypass the problem of 255 circuits that IS-IS allows to maintain with minor violations of the specification that however does not prevent interoperability with existing implementations.

[1](#). Introduction

IS-IS reserves within its packets only one byte for a circuit ID and the specification requests it to be unique. This ID is used in broadcast networks to identify a PNode. They don't seem to serve any

particular purpose on p2p networks except for some checking in hellos and tie-breaking in removal of excess adjacencies.

[2](#). More than 255 circuits for p2p links

For all p2p links within an Intermediate System, the same circuit ID can be chosen safely to interact with other Intermediate Systems. Even when two such systems are brought up on two ends of the link and both pick the same value, IS-IS specification does not preclude such a configuration. In case of multiple parallel links between the systems the only obvious problem is the impact on tie-breaking in case of excessive adjacencies. However, such a configuration cannot generate forwarding loops.

3. More than 255 circuits for broadcast links

Trickier a case is the one in which an intermediate system has to form adjacencies on a broadcast medium. The decisive insight is the fact that unique circuit ID on a broadcast medium is only needed in the case where the given intermediate system is assuming the role of the DIS for the LAN. As long as the intermediate system has not elected itself DIS, it can use a non-unique circuit ID (1). Therefore, it is enough to change circuit ID in all the packets transmitted to a unique one as soon as DIS election ran and the system must become DIS. Such behavior is basically indistinguishable from a node crashing and coming immediately back with a different circuit ID than the one used before the crash. Implementation experience shows that it is unwise to tear down the adjacencies by sending empty hellos before changing circuit IDs since this can lead to ``ripple''-down behavior in DIS in the case of all routers having the same preference. However, to ensure that all involved parties agree on LAN ID of the media, implementations must interpret in [subsection 8.4.5](#) within 10589 the specified condition

f) the set of Intermediate system adjacency states as including any change in LAN ID to be an indication triggering the recomputation of the according DIS. Additionally, when recycling a previously used circuit ID and reusing it on another LAN special

1. it is important to realize that circuit ID is not a part of tie-breaking rules on DIS election

care has to be taken that the newly generated pseudonode LSPs have sequence numbers high enough as to prevent unnecessary flooding.

3.1. Modification of the Behavior on Broadcast Media

The exact modification of the behavior for broadcast links is given here. It is a modification of chapter 8.4.5 of the original specification that states:

When the broadcast circuit is enabled on an Intermediate system the IS shall perform the following actions.

- a) Commence sending IIH PDUs with the LAN ID field set to the concatenation of its own SystemID and its locally assigned one octet Local Circuit ID.
- b) ... <omitted for clarity>
- c) Start listening for ISO 9542 ISH PDUs and ESH PDUs and acquire adjacencies as appropriate. Do not run the Designated Intermediate System election process.
- d) After waiting ISISHelloTimer * 2 seconds, run the Level 1 and/or the Level 2 Designated Intermediate System election process, depending on the Intermediate system type. This shall be run subsequently whenever an IIH PDU is received or transmitted as described in 8.4.3. (For these purposes, the transmission of the system's own IIH PDU is equivalent to receiving it). If there has been no change to the information on which the election is performed since the last time it was run, the previous result can be assumed. The relevant information is
 - 1) the set of Intermediate system adjacency states,
 - 2) the set of Intermediate System priorities (including this system's) and
 - 3) the existence (or otherwise) of at least one "Up" End system (not including Manual Adjacencies) or Intermediate system adjacency on the circuit.

An Intermediate system shall not declare itself to be a LAN Designated Intermediate system of any type until it has at least one

"Up" End system (not including Manual Adjacencies) or Intermediate system adjacency on the circuit. (This prevents an Intermediate System which has a defective receiver and is incapable of receiving any PDUs from erroneously electing itself LAN Designated Intermediate System.) The LAN ID field in the LAN IIH PDUs transmitted by this system shall be set to the value of the LAN ID reported in the LAN IIH PDU (for the appropriate level) received from the system which this system considers to be the Designated Intermediate System. This value shall also be passed to the Update Process, as the pseudonode ID, to enable Link State PDUs to be issued for this system claiming connectivity to the pseudonode. If this system, as a result of the

Designated Intermediate System election process, considers itself to be designated Intermediate System, the LAN ID field shall be set to the concatenation of this system's own ID and the locally assigned one octet Local Circuit ID.

One additional point has to be introduced:

- a) Commence sending IIH PDUs with the LAN ID field set to the concatenation of its own SystemID and a local non-zero, not necessarily unique circuit ID.
- b) ... <omitted for clarity>
- c) Start listening for ISO 9542 ISH PDUs and ESH PDUs and acquire adjacencies as appropriate. Do not run the Designated Intermediate System election process.
- d) After waiting ISISHelloTimer * 2 seconds, run the Level 1 and or the Level 2 Designated Intermediate System election process depending on the Intermediate system type. If in any point in time the decision process determines the node to be DIS for this LAN:
 - 1) Commence sending IIH PDUs with the LAN ID field set to the concatenation of its own SystemID and a local unique circuit ID.
 - 2) go to step b.

otherwise exhibit standard behavior.

[3.2](#). Multiple levels

Since election of DIS is performed independently at each level, the extensions can be applied to generate local circuit IDs only for the level at which the system has been elected.

[4](#). Acknowledgments

Some smart people probably thought about most of these things before and the p2p case may be documented in the best current practices for IS-IS in the Internet. Mike Shand, Tony Li, Arun Satyanarayana, Rajesh Varadarajan and Christian Hopps provided additional comments.

[5](#). Security Consideration

ISIS security applies to the work presented. No specific security issues with the proposed solutions are known.

6. Intellectual Property Considerations

Lucent may seek patent or other intellectual property protection for some or all of the technologies disclosed in this document. If any standards arising from this document are or become protected by one or more patents assigned to Lucent, Lucent intends to disclose those patents and license them under openly specified and non-discriminatory terms.

References

- [Cal90a] R. Callon. OSI ISIS Intradomain Routing Protocol. INTERNET-RFC, Internet Engineering Task Force, February 1990.
- [Cal90b] R. Callon. Use of OSI ISIS for Routing in TCP/IP and Dual Environments. INTERNET-RFC, Internet Engineering Task Force, December 1990.
- [IS090] ISO. Information Technology - Telecommunications and Information Exchange between Systems - Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service. ISO, 1990.

Przygienda et al.	Expires 1 Dec 1999	[Page 5]
-------------------	--------------------	----------

Internet Draft	255+ Adjacencies in IS-IS	1 Jun 1999
----------------	---------------------------	------------

Authors' Addresses

Tony Przygienda
Bell Labs, Lucent Technologies
[101](#) Crawfords Corner Road
Holmdel, NJ 07733-3030
prz@dnrc.bell-labs.com

Przygienda et al.	Expires 1 Dec 1999	[Page 6]
-------------------	--------------------	----------