

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 24, 2014

E. Osborne
Cisco Systems, Inc.
JD. Ryoo
ETRI
H. van Helvoort
Huawei Technologies
A. D'Alessandro
Telecom Italia
T. Cheung
ETRI
October 21, 2013

Capabilities and Modes in PSC draft-or-psc-cap-mode-01

Abstract

This document introduces capabilities and modes to PSC. A capability is an individual behavior, and a node's set of capabilities are signalled using the method given in this draft. A mode is a particular combination of capabilities.

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](#)].

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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.	Capabilities	3
2.1.	Sending the Capabilities TLV	3
2.1.1.	What to send	3
2.1.2.	When to send it	4
2.2.	Receiving the Capabilities TLV	4
2.2.1.	Comparing Capabilities TLVs	4
2.2.2.	Handling a missing TLV	4
2.3.	Handling errors	4
2.3.1.	TLV Timeout	4
2.3.2.	TLV Mismatch	5
2.3.3.	Handling an error	5
3.	Modes	5
3.1.	PSC Mode	5
3.2.	APS Mode	6
3.2.1.	SF-P and FS priority swap	6
3.2.2.	Modified non-revertive behavior and MS-W	6
3.2.3.	Signal degrade	6
3.2.4.	Exercise	6
4.	Backward compatibility	6
5.	IANA Considerations	7
6.	Security Considerations	7
7.	Acknowledgements	7
8.	Normative References	7
	Authors' Addresses	7

[1.](#) Introduction

This document brings two things to PSC - Capabilities, and Modes. A Capability is an individual behavior whose use is signalled in a Capabilities TLV inside PSC, and a Mode is a predefined set of Capabilities.

2.1.2. When to send it

PSC sends messages for two reasons - messages in response to external inputs, and periodic retransmission of current status. It is not necessary to send the Capabilities TLV with every PSC packet. Indeed, it may be expensive to send and to parse an Capabilities TLV attached to a packet intended to trigger a protection switch or other real-time behavior. However, it is also true that if a node does not periodically send its Capabilities TLV, the receiving node cannot tell the difference between a deliberate omission of the Capabilities TLV for performance reasons versus an accidental omission due to an implementation issue. To guard against this, a node **MUST** send its Capabilities TLV in every periodic retransmission of current status, and **MAY** send its Capabilities TLV more frequently than that.

2.2. Receiving the Capabilities TLV

A node **MUST** establish a receive timer for the Capabilities TLV. By default this **MUST** be three times the periodic retransmission timer of five seconds - so, fifteen seconds. Both the periodic retransmission time and the timeout **SHOULD** be configurable by the operator. When a node receives a Capabilities TLV it resets the timer to fifteen seconds. If the timer expires, the node behaves as in [Section 2.3](#).

2.2.1. Comparing Capabilities TLVs

When a node receives a Capabilities TLV it **MUST** compare it to its most recent transmitted Capabilities TLV. If the two are equal, the protection domain is said to be running in the mode indicated by that set of capabilities (see [Section 3](#)).

2.2.2. Handling a missing TLV

As mentioned in [section 2.1.2](#), a node may not send the Capabilities TLV with every PSC message. However, a node **MUST** send the Capabilities TLV as part of its periodic state retransmission.

2.3. Handling errors

This section covers the two possible errors - a TLV timeout and a TLV mismatch - and the error handling procedures in both cases.

2.3.1. TLV Timeout

If the Capabilities TLV receive timer expires, a node is said to have timed out. When this happens, the node **MUST** alert the operator and **MUST** behave as in [Section 2.3.4](#), "Handling an error".

2.3.2. TLV Mismatch

If the sent and received Capabilities TLVs are not equal, this indicates a capabilities mismatch. For timer and timeout purposes, a capabilities mismatch is treated as a missed TLV. That is, the received TLV is not taken as an input to the receive timer for Capabilities TLV refresh purposes. A node MAY retain the received TLV for logging, alert or debug purposes.

2.3.3. Handling an error

If a node decides that it is in an error state in respect to the Capabilities TLV, it MUST do two things:

- 1) Indicate the detected mismatch to the operator by the usual alert mechanisms (e.g. syslog).
- 2) Not make any state transitions based on the contents of any PSC messages

To expand on point 2 - assume node A is receiving NR(0,0) from its PSC peer Z and is also receiving a mismatched set of capabilities (e.g. received 0x4, transmitted 0x5). If Z detects a local SF-W and wants to initiate a protection switch (that is, by sending SF(1,1)), Node A MUST NOT react to this input by changing its state. A node MAY increase the severity or urgency of its alarms to the operator, but until the operator resolves the mismatch in the Capabilities TLV the protection domain will likely operate in an inconsistent state.

3. Modes

A Mode is a given set of Capabilities. Modes are shorthand; referring to a set of capabilities by their individual values or by the name of their mode does not change the protocol behavior. This document defines two modes - PSC and APS.

<< are those the right names? TBD. >>

<< how to define future modes. TBD. >>

3.1. PSC Mode

PSC Mode is defined as the lack of any Capabilities - that is, a Capabilities set of 0x0. It is the behavior specified in [RFC6378](#). There are two ways to declare PSC Mode. A node can send a Capabilities TLV of 0x0, or it can send no Capabilities TLV at all. This is further explored in [section 4](#).

3.2. APS Mode

APS Mode is defined as the use of five specific capabilities. These capabilities are defined in other documents, but are repeated here in order to define the mode. They are listed below. APS Mode is indicated with a Capabilities TLV of 0x1F.

3.2.1. SF-P and FS priority swap

This feature is defined in [draft-rhd-mpls-tp-psc-priority](#) and is assigned bit 0x1.

3.2.2. Modified non-revertive behavior and MS-W

These are both defined in [draft-cdh-mpls-tp-psc-non-revertive](#) but are negotiated separately. Modified NR behavior is assigned bit 0x2 and MS-W is assigned bit 0x4.

3.2.3. Signal degrade

This feature is defined in [draft-rhd-mpls-tp-psc-sd](#). It is assigned bit 0x8.

3.2.4. Exercise

This feature is defined in [draft-dj-mpls-tp-exer-psc](#). It is assigned bit 0x10.

4. Backward compatibility

As defined in [Section 3](#), PSC Mode is indicated either with a Capabilities TLV of 0x0 or the lack of any Capabilities TLV. This is to allow backward compatibility between two nodes - one which can send the Capabilities TLV, and one which cannot.

[RFC6378](#) does not define how to handle an unrecognized TLV. There may be some implementations that silently discard an unrecognized TLV, and some that take more drastic steps like refusing to allow PSC to operate. Thus, a node which has the ability to send and receive the PSC Mode Capabilities TLV MUST be able to both send the PSC Mode Capabilities TLV and send no Capabilities TLV at all. An implementation MUST be configurable between these two choices.

One question that arises from this dual definition of PSC Mode is, what happens if a node which was sending a non-null Capabilities TLV (e.g. APS Mode) sends PSC packets without any Capabilities TLV? This case is handled as follows:

If a node has never, during the life of a PSC session, received a Capabilites TLV from a neighbor, the lack of a Capabilites TLV is treated as receipt of a PSC Capabilites TLV. This allows for interop between nodes which support the PSC Mode TLV and nodes which do not, and are thus implicitly operating in PSC Mode.

If a node has received a non-null Capabilites TLV (e.g. APS Mode) during the life of a PSC session and then receives a PSC packet with no Capabilites TLV, the receiving node MUST treat the lack of Capabilites TLV as simply a lack of refresh. That is, the receipt of a PSC packet with no Capabilites TLV simply does not reset the receive timer defined in [Section 2.2](#)

5. IANA Considerations

- A value for the Cap TLV
- A registry for the capabilities inside the Cap TLV

6. Security Considerations

None.

7. Acknowledgements

8. Normative References

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

Authors' Addresses

Eric Osborne
Cisco Systems, Inc.

Email: eosborne@cisco.com

Jeong-Dong Ryoo
ETRI

Huub van Helvoort
Huawei Technologies

Email: Huub.van.Helvoort@huawei.com

Alessandro D'Alessandro
Telecom Italia

Email: alessandro.dalessandro@telecomitalia.it

Taesik Cheung
ETRI

Email: cts@etri.re.kr