

PW usage nits

Access

# LDP status messages

There is a mathematical theorem that states

**4447 < 5885** (Note: *less than*, not just *less than or equal to*)

This is important since the *earlier* RFC 4447 states :

If the PW Status TLV is not present following the FEC TLV in the initial PW Label Mapping message received by a PE, then the PW Status TLV will not be used, and both PEs supporting the pseudowire will revert to the label withdraw procedure for signaling status changes.

(this is actually stated twice, and I understand `will=MUST`)

which forced the *later* RFC 5885 to state :

BFD CV Types used for fault detection and status signaling (i.e., CV Types 0x08 and 0x20) SHOULD NOT be used when a control protocol such as LDP [RFC4447] or L2TPV3 [RFC3931] is available that can signal the AC/PW status to the remote endpoint of the PW.

# LDP status messages (cont.)

Does this make *any* sense ?

VCCV-BFD diagnostic messages  
are in-band and fast

LDP status messages (and label withdrawal messages for that matter)  
are out-of-band and TCP-based (i.e., slow)

Statically configured PWs can use BFD diagnostic codes

When using the PWE3 control protocol  
we are forced to use LDP messaging

*Maybe we can work around the limitation  
by claiming not to support the LDP status message ?*

# LDP status messages (cont.)

## RFC 4447 states :

When a PW is first set up, the PEs MUST attempt to negotiate the usage of the PW status TLV. This is accomplished as follows: A PE that supports the PW Status TLV MUST include it in the initial Label Mapping message following the PW FEC and the interface parameter sub-TLVs. The PW Status TLV will then be used for the lifetime of the pseudowire.

**So, the PE is not allowed to “lie” about its support and even if it did - it would revert to label withdrawal !**

# LDP status messages (cont.)

## Proposals:

Update RFC 4447 to :

- allow PEs to state that they don't support status TLV
- allow PWs that
  - support VCCV-BFD
  - do not support (or claim not to support) status TLV to use BFD diagnostic codes

# CC type usage

There was a lengthy discussion on the email list about the use of Associated Channel CC types :

1. CW nibble
2. router alert label
3. TTL expiry

and now there is a new draft proposing using the GAL

RFC 5085 says that

- only a single CC type may be used per PW
- it is *not recommended* to change CC as that could lead to interop issues

couldn't that lead to interop issues?

It also doesn't say what a receiver should do if it receives an unexpected CC type

# CC type usage (cont.)

Quick review of the arguments that were brought up

1. RFC 4385 says that the ACH *is* identified by the CW nibble
2. RA is an essential MPLS architectural feature and thus never always has to be accepted
1. TTL expiry is necessarily supported as any expired packet **MUST** be sent for special processing

Conclusion:

- There is never really any reason to signal CC type.

Questions:

- Must statically provisioned PWs accept all CC types ?
- If a certain CC type was negotiated should other CC types received be discarded ?
- Is it a general principle (applying equally to static configuration) that both directions use the same CC type ?

# CC type usage (cont.)

## Proposals

Update RFC 4447 to state

- CC negotiation is optional
- if not negotiated, receiver **MUST** accept any CC type (*conservative sending liberal accepting*)
- if CC *is* negotiated
  - it must be the same in both directions
  - if a non-negotiated CC type *is* received then it is silently discarded

Update RFC 5085 to state

- static PWs must support any CC type received



# CW usage (when LDP is used for setup)

RFC 4447 section 6.2 and Appendix A give detailed explanations how to force both PW directions to either *use* or *not use* the CW

If both endpoints prefer the use of the control word, this procedure will cause it to be used. If either endpoint prefers not to use the control word or does not support the control word, this procedure will cause it not to be used. If one endpoint prefers to use the control word but the other does not, the one that prefers not to use it has no extra protocol to execute; it just waits for a Label Mapping message that has  $c=0$ .

# CW usage (when BGP is used for setup)

## RFC 4761 – section 3.2.4 says (read the fine print)

### 3.2.4 Signaling PE Capabilities

The following extended attribute, the "Layer2 Info Extended Community", is used to signal control information about the pseudowires to be setup for a given VPLS. ... This information includes the Encaps Type (type of encapsulation on the pseudowires), Control Flags (control information regarding the pseudowires), and the Maximum Transmission Unit (MTU) to be used on the pseudowires. ...

```

0 1 2 3 4 5 6 7
+-----+-----+
|  MBZ      |C|S|      (MBZ = MUST Be Zero)
+-----+-----+

```

Figure 4: Control Flags Bit Vector

Name	Meaning
C	A Control word [7] MUST or MUST NOT be present when sending VPLS packets to this PE, depending on whether C is 1 or 0, respectively
S	Sequenced delivery of frames MUST or MUST NOT be used when sending VPLS packets to this PE, depending on whether S is 1 or 0, respectively

But nothing explicit is said about  
both directions *using* or *not using* the CW

# CW usage (for general PWs)

Is it a general principle of CWs that both directions either use or don't use the CW ?

Very little is said in RFCs about static PWs

When manually provisioning PWs do we need to ensure CW usage consistency ?

## Proposal

Update RFC 3985 to state that for ALL PWs

- both directions **MUST** either *use* or *not use* the CW
- require enforcement of this at setup (*however accomplished*)
- usage **MUST** not change during the life of the PW