

VDP Transport Options

Presentation to 802.1Qbg Team

*Rakesh Sharma (IBM), Chuck Hudson (HP), Renato Recio (IBM), Manoj
Wadekar (Qlogic), Anoop Ghanwani (Brocade), Srikanth Kilaru (Juniper),
Vivek Kashyap (IBM), Vijoy Pandey (BNT)*

01/12/2010

VSI Discovery Protocol needs a Trivial TLV Transport Protocol

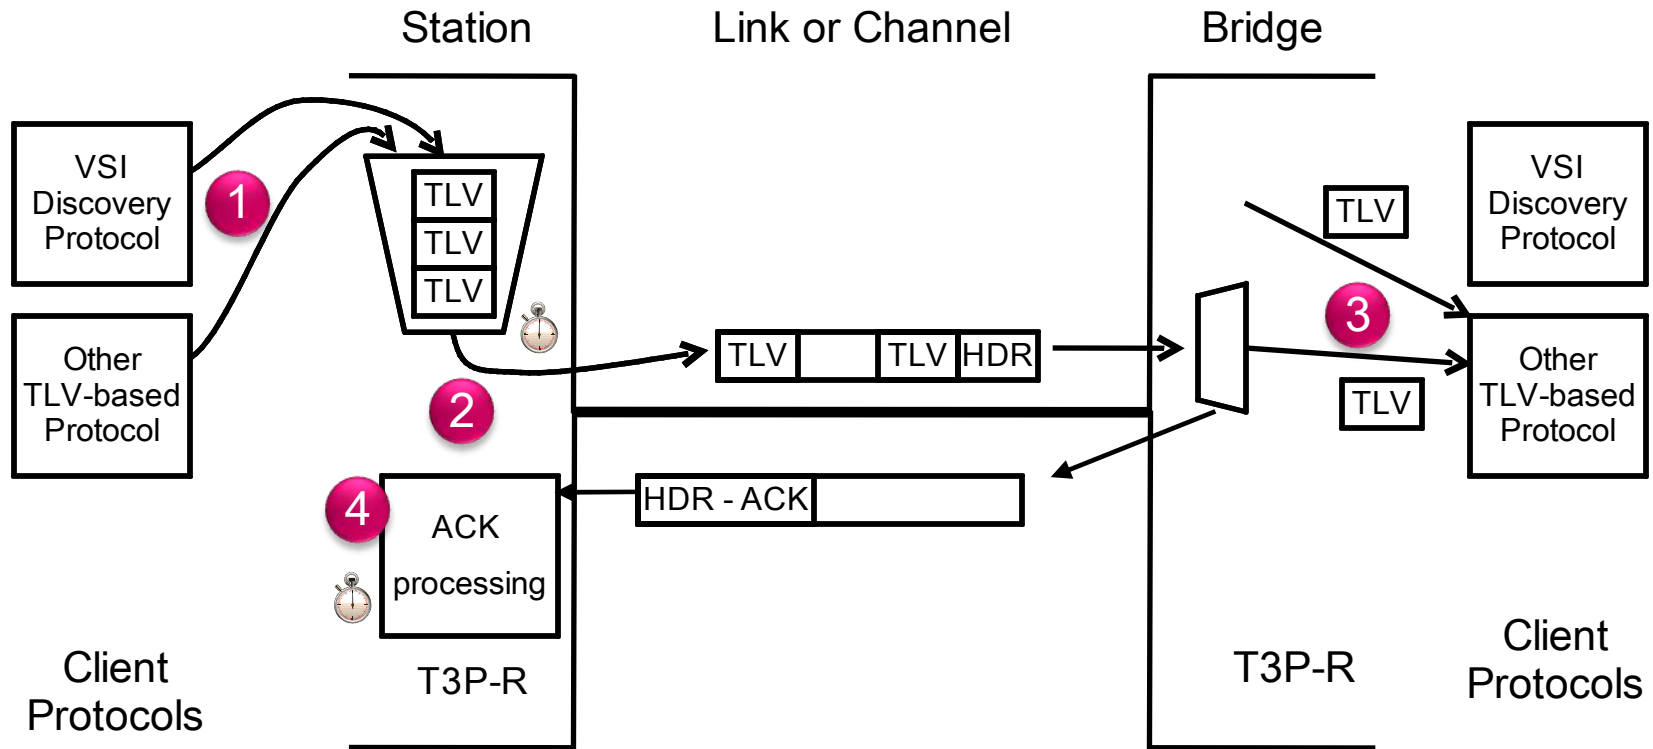
Classic LLDP	T3P-R
Requires ALL TLVs in every LLDP exchange. Limited to 1500 bytes, no fragmentation. Limits number of TLVs i.e. VSIs. With current definition of TLV, we get about 32 VSIs per channel.	Assumes subsets of TLVs can be passed in a frame. Therefore, does not put limit on number of VSIs
Assumes that information is exchanged (no state maintained).	Assumes that TLV exchanges can be stateful and that state can be maintained for each TLV type.
<ol style="list-style-type: none"> 1. Delivery is confirmed by regularly retransmitting (and fast retransmit) the full set of information. If a frame gets lost, next retransmit will send the lost information. Retransmission mechanism is slow. 2. Frame transmissions are limited by available credits (1 new credit per second and max of 5) 	Successful delivery is confirmed by other approaches: <ul style="list-style-type: none"> • ACK support in T3P (T3P-R). T3P-R can enable simpler and more scalable protocol. • TLV-specific ACKs/NACKs may not be required • Transmission of a signature of the last TLV sent or of current state may not be required.
Requires repetition of information that is not changing.	Only transmits TLVs that are new or changed.
Current IEEE standard (802.1AB) and widely implemented	New protocol but appear to have wide support in 802.1bg/h standardization team.

Conclusion: Classic LLDP can be used for VSI Discovery and Configuration with limitations. T3P-R does not have the limitations.

Backup

Trivial TLV Transport - R (single direction shown)

Notes:
 Reliability using ACK
 Single writer, multiple readers
 Sent/received in the same order as sent.



1 Client protocols pass outgoing TLVs to T3P (indicating whether the TLV is 'queued' or 'immediate') Reliability option can be set. First in queue sets timer.

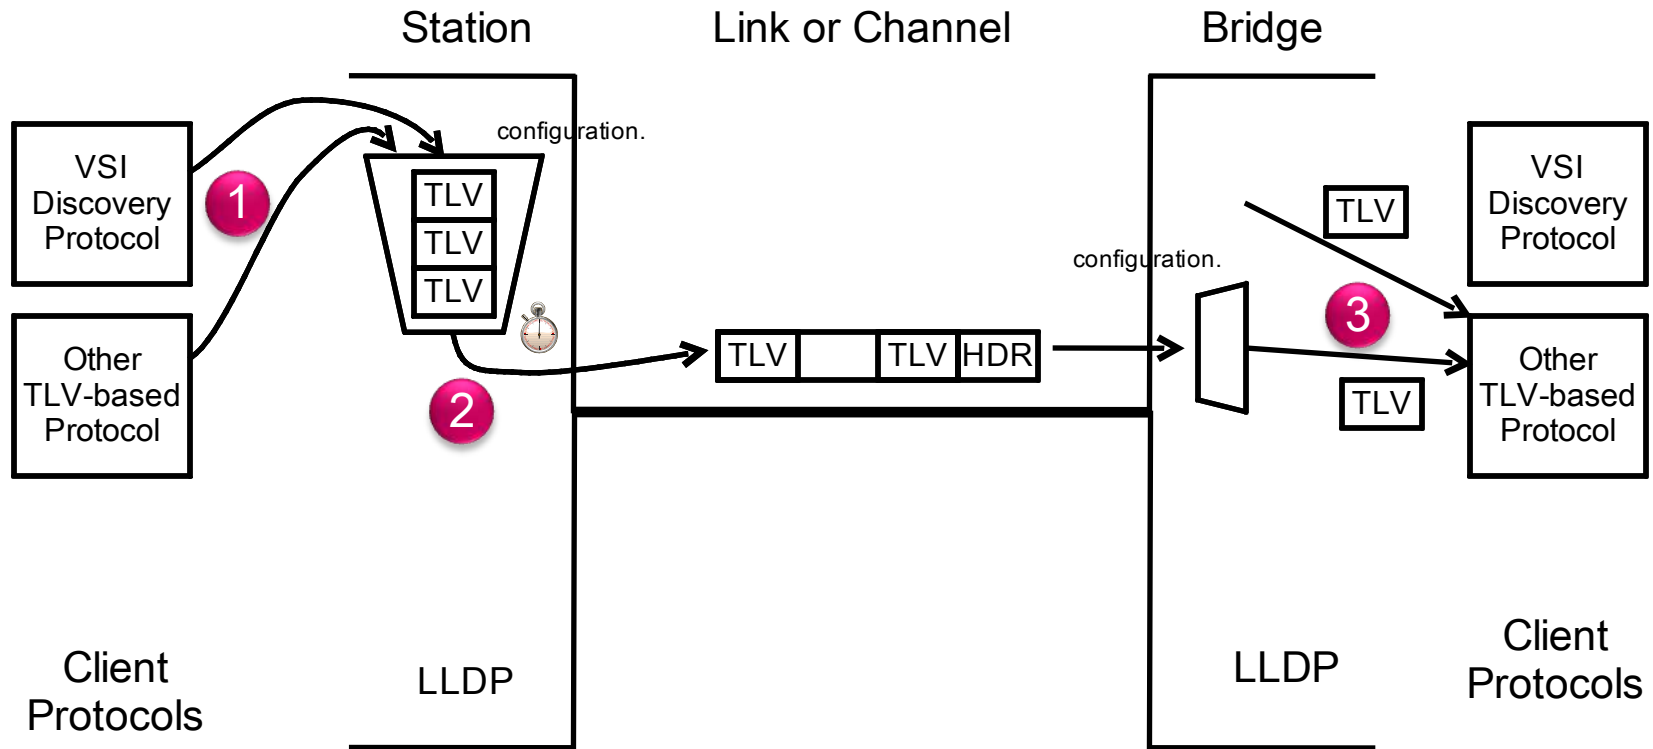
2 If ((queue > 1500) | (immediate TLV) | (timer expired)) Then Send T3P frames until queue emptied

3 Client protocols subscribe to specific TLVs, so the appropriate TLVs are forwarded to the correct client protocol

4 ACK Processing to Provide reliability. If ACK Timer expires, re-xmit. Only 1 TLV block tx at time.

Classic LLDP (single direction shown)

Notes:
Reliability using LLDP retransmits
Single writer, multiple readers
Sent/received in the same order as sent.



1

Client protocols pass outgoing TLVs to LLDP.

2

Client protocols TLVs are stored as local configuration

3

Client protocols access remote TLVs from LLDP configuration..