

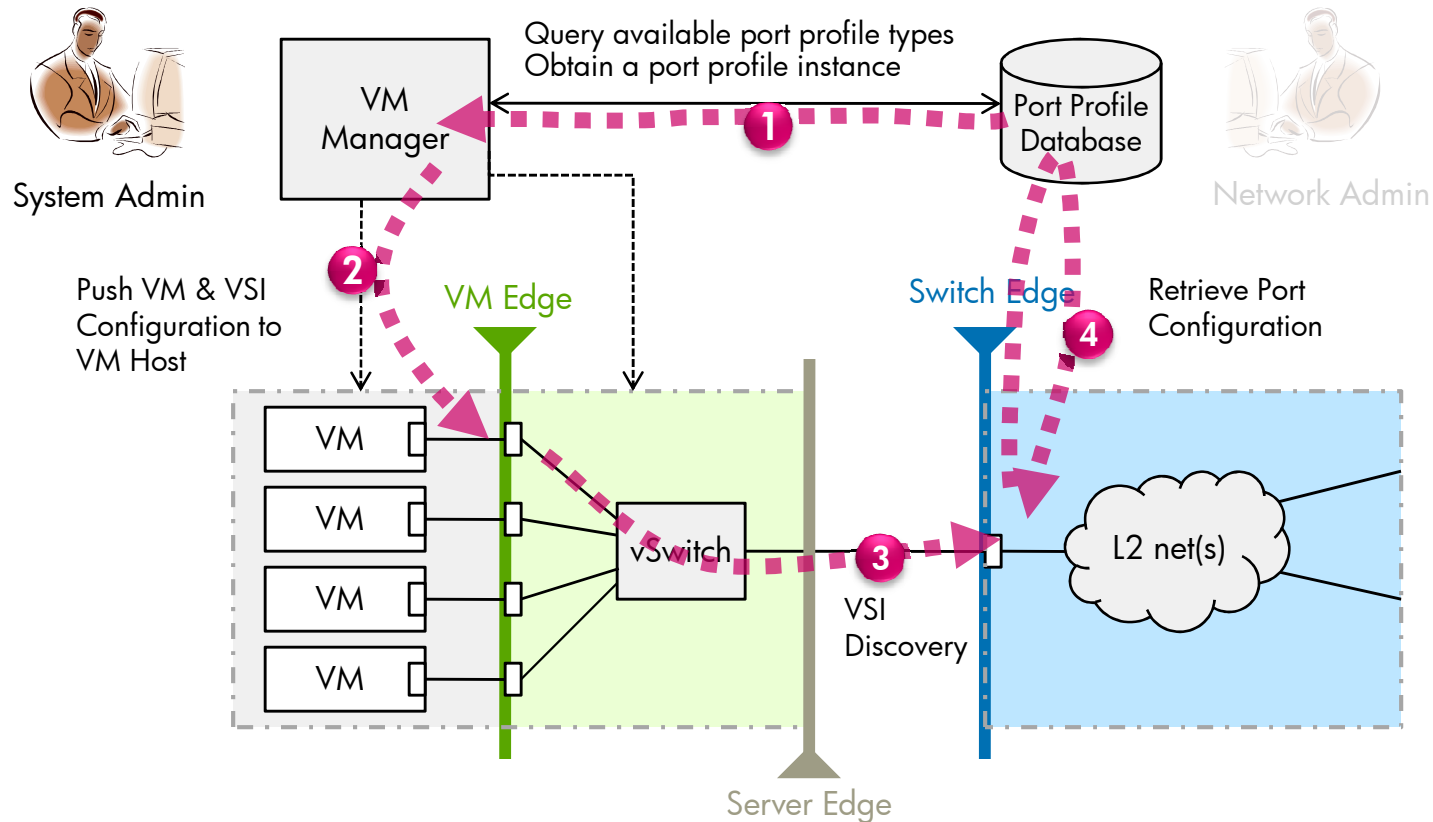
# VSI Discovery Requirements & Trivial TLV Transport Protocol (T<sup>3</sup>P)

Draft v0.1

November 23, 2009

Chuck Hudson with input from Anoop Ghanwani, Manoj Wadekar, Mike Krause, Rakesh Sharma, Srikanth Kilaru

# Steps for Configuring Edge Connections (vPorts)



# VSI Discovery Protocol needs an alternative to LLDP

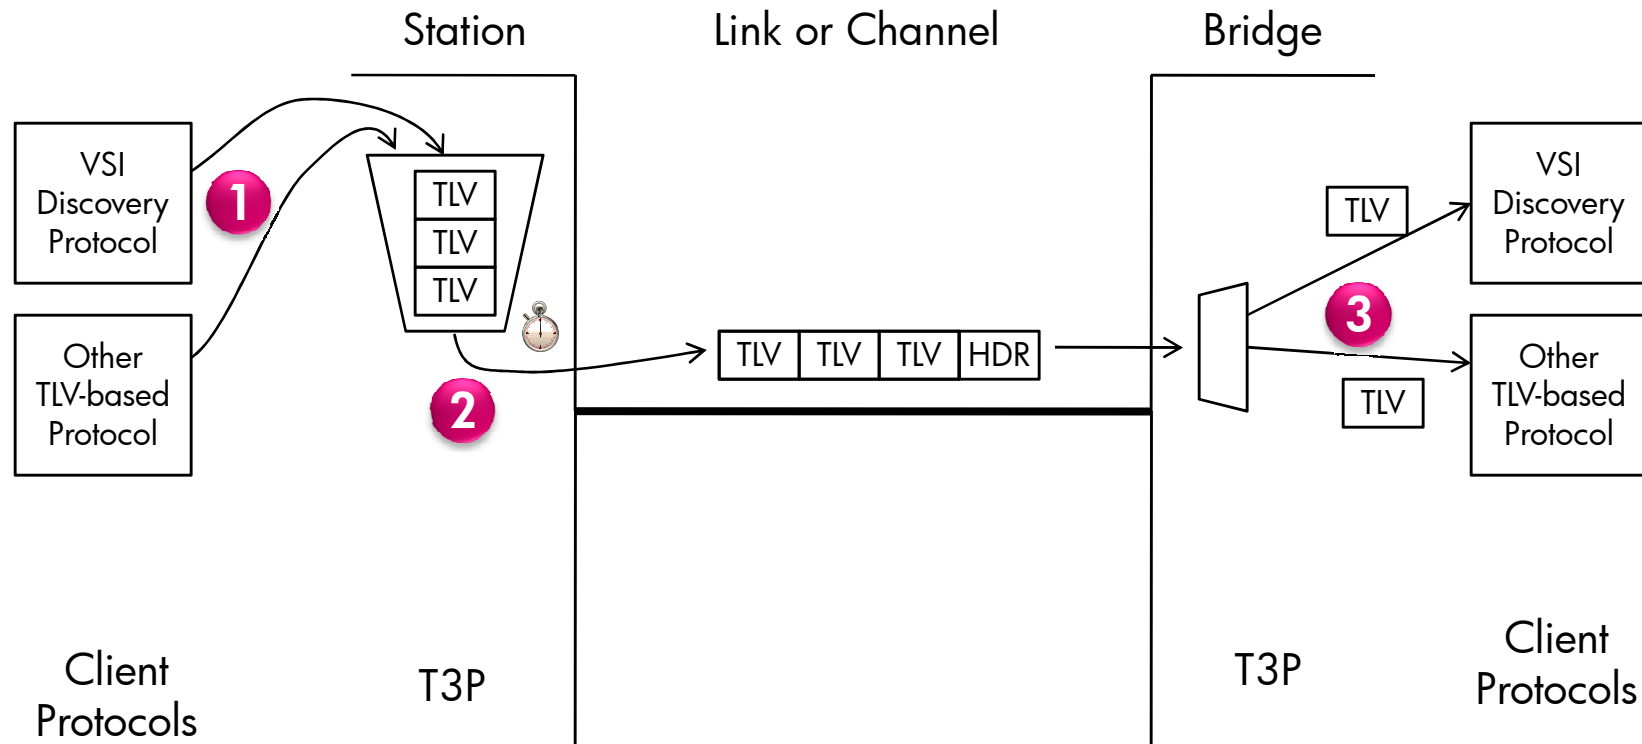
LLDP	VDP/T3P
Requires ALL TLVs in every LLDP exchange. Limited to 1500 bytes, no fragmentation.	Assumes subsets of TLVs can be passed in a frame. Total size of all TLVs can significantly exceed 1500 bytes.
Assumes that information is exchanged (no state maintained).	Assumes that TLV exchanges can be stateful and that state can be maintained for each TLV type.
Delivery is confirmed by regularly retransmitting (and fast retransmit) the full set of information.	Successful delivery is confirmed by other approaches: <ul style="list-style-type: none"> <li>• TLV-specific ACKs/NACKs</li> <li>• Transmission of a digest of the last TLVs sent or of current state.</li> <li>• Etc.</li> </ul>
Requires repetition of information that is not changing.	Only transmits TLVs that are new or changed.

# What should stay the same

- TLV format
  - But allow multiple of the same TLV type
- Addressing
- Strong push to minimize traffic

# Trivial TLV Transport (single direction shown)

Notes:  
 Reliability done in client protocol  
 Single writer, multiple readers  
 Sent/received in the same order as sent.



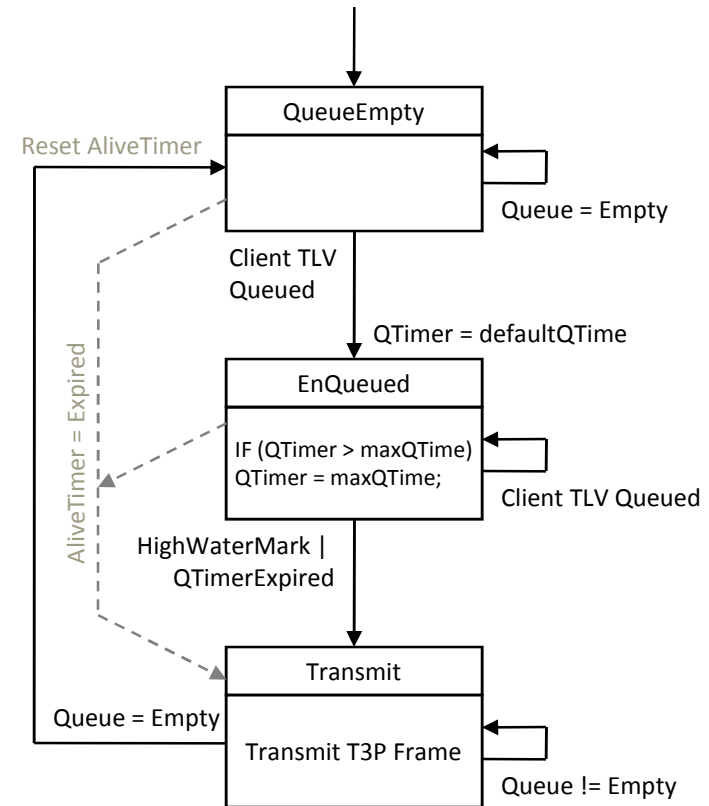
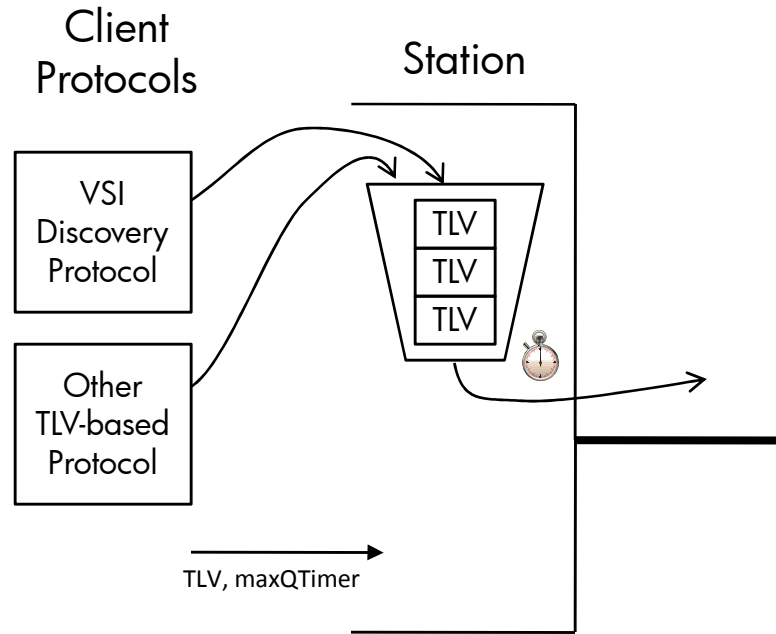
**1** Client protocols pass outgoing TLVs to T3P (indicating whether the TLV is 'queued' or 'immediate'). First TLV in queue sets queue timer.

**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. (\*single sender, but multiple listeners allowed)

# Trivial TLV Transport

## Transmit State Machine (rough)



- May want to add Keep Alive Timer for T3P link
  - Reset when entering QueueEmpty state
  - Force transmission of a T3C TLV
- May want to allow client protocols to specify maxQueueTime
  - When TLV queued, verify Queue timer is  $\leq$  maxQTime for new TLV
  - Setting maxQTime to zero is asking for immediate transmission

# Trivial TLV Transport Protocol

- Very simple
- Meets needs of VSI Discovery Protocol
- Minimizes traffic for new protocols
- Easy to learn (re-uses LLDP knowledge)
- Could carry other client protocols
  - Example: LLDP TLVs
    - LLDP 'client protocol' holds inventory of LLDP TLVs
    - Use TLV transmit timers to match LLDP intervals
    - Add timers for each TLV received (to clear if no longer sent)
    - Allows LLDP for