PFC and Untagged Frames

bb-ghanwani-pfc-untagged-0509-v1

Anoop Ghanwani (anoop@brocade.com)
Dan Eisenhauer (dge@us.ibm.com)
Rakesh Sharma (rakesh@us.ibm.com)

May 2009

Overview

- Motivation
- Types of untagged frames
- Problems with untagged frames on links with PFC
- Options for solving these problems

Motivation

- PFC operates on a per-priority basis
- The priority for a frame is usually determined by looking at the priority code point (PCP) bits in the VLAN tag
- How should untagged frames be handled?
- P802.1Qbb is silent on this issue

Types of Untagged Frames

- Control
  - BPDU: RSTP/MSTP, LLDP, 802.1X, MVRP
  - Slow protocols: LACP, 802.3ah
- User data

- Control frames are always untagged
- We have some control over whether or not user data are transmitted as untagged

Untagged Frames in 802.1Q

- Ingress operation
  - Untagged frames are assigned the default port priority
  - The VLAN is determined from the protocol in the frame or the PVID of the receiving port
- Egress operation
  - The port may be in the *tagged set* or *untagged set* for a VLAN
  - If in the untagged set, the frame is transmitted untagged and both the VLAN and priority information are lost

- It is assumed that if preservation of the priority is important, then the VLAN tag will be retained
- An end station which doesn't know which VLAN to use can send priority-tagged frames (VLAN ID = 0)
- A bridge always knows what VLAN ID to use (VLAN ID != 0)
  - As a result, a bridge never sends priority-tagged frames

PFC and Untagged Frames

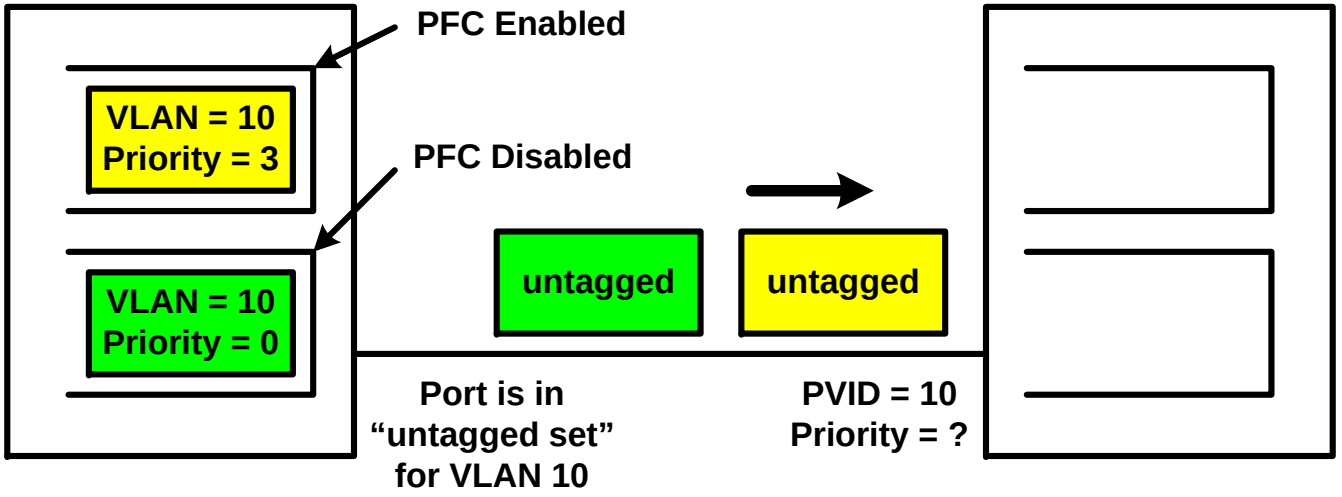

- An untagged frame for a VLAN could be transmitted from any priority
- These would be classified as being on the same priority (the port default priority) at the receiving end

Problem Scenarios

- Untagged frames are sent from a priority with PFC enabled and the default port priority at the remote end has PFC disabled
  - Lossy service
- Untagged frames are sent from a priority with PFC enabled and the default port priority at the remote end has PFC enabled, but is a different value
  - PFC messages hit the wrong queue
- Untagged frames are sent from a priority with PFC disabled and the default port priority at the remote end has PFC enabled
  - PFC messages hit the wrong queue

Option 1

- Control frames are sent untagged
  - An implementation must ensure that the receiver must ensure that they do not consume resources that are for use by PFC frames
- The port default priority must not have PFC enabled
- Always require tagging of data frames when PFC is enabled
  - If VLAN unaware, end stations must send traffic as priority-tagged and ignore the VLAN ID in received frames

- This option has the least impact on the existing bridge specification
- But it is a problem for existing network stacks, e.g. Linux, to deal with VLAN tags they do not know about
  - Such frames are usually discarded

Option 2

- Control frames are sent untagged (Same as Option 1)
  - An implementation must ensure that the receiver must ensure that they do not consume resources that are for use by PFC frames
- The port default priority must not have PFC enabled (Same as Option 1)
- Require priority-tagging for untagged VLANs on links with PFC
  - Bridges and end stations would transmit untagged frames as priority-tagged
  - Alternatively, require this behavior only if transmitting from a priority that has PFC enabled

- Requires modification to 802.1Q for bridges to transmit untagged frames as priority-tagged when PFC is negotiated on a link

Summary

- Untagged frames pose a problem when using PFC
- Two options presented
  - Based on existing behavior of networking stacks, Option 2 is more appealing
- Any other options?

BROCADE

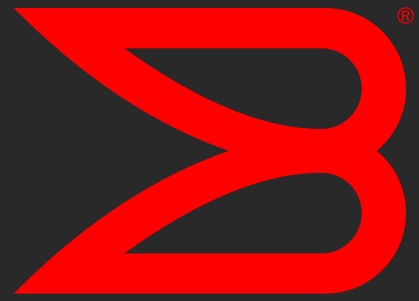

THANK YOU