


BROCADE



FIP Snooping

T11/ 08-283v1

Anoop Ghanwani
(anoop@brocade.com)

June 4, 2008

Overview

- What is snooping?
- Examples of snooping in existing bridged networks
- FIP snooping – what works and what doesn't work?
- Summary

What is Snooping?

- According to Merriam-Webster
snoop: to look or pry especially in a sneaking or meddlesome manner
- Normally bridges forward frames except those addressed to its management entity or those addressed to the well-known bridge address
- Snooping in the bridged world refers to a bridge control plane picking up certain frames not addressed to it, processing the frame, and only then forwarding it
- Snooping is typically done for the purpose of optimizing forwarding or for security reasons

Examples of Snooping

- IGMP snooping
 - Bridge snoops on IGMP messages to determine where to forward frames destined to a multicast group
 - Optimizes forwarding to only those parts of the spanning tree where there are interested receivers
 - When IGMP snooping is enabled
 - In the absence of any IGMP messages for a given group, traffic for that group is flooded or dropped, depending on configuration
 - Once an IGMP message is received then the forwarding is restricted to only those ports from where “joins” were seen
 - The “snooped state” times out after typically 210 seconds
 - Typically enabled on all switches – edge and core – in the network

Examples of Snooping (2)

- DHCP rogue server prevention
 - Used to prevent a rogue server from responding to a DHCP request
 - A bridge has “trusted” and “untrusted” ports
 - Ports with legitimate DHCP servers are configured to be trusted
 - Only trusted ports are allowed to respond to DHCP requests
 - State may be maintained for the following
 - We want to limit the number of DHCP responses from the server
 - We want to maintain a table of MAC/IP and lease times on untrusted ports so that access control can be enforced
 - State needs to be preserved across switch reboots
 - This function would typically be enabled only at edge switches
 - All ports connected to the core would be configured as trusted
 - The function would not be enabled in core switches

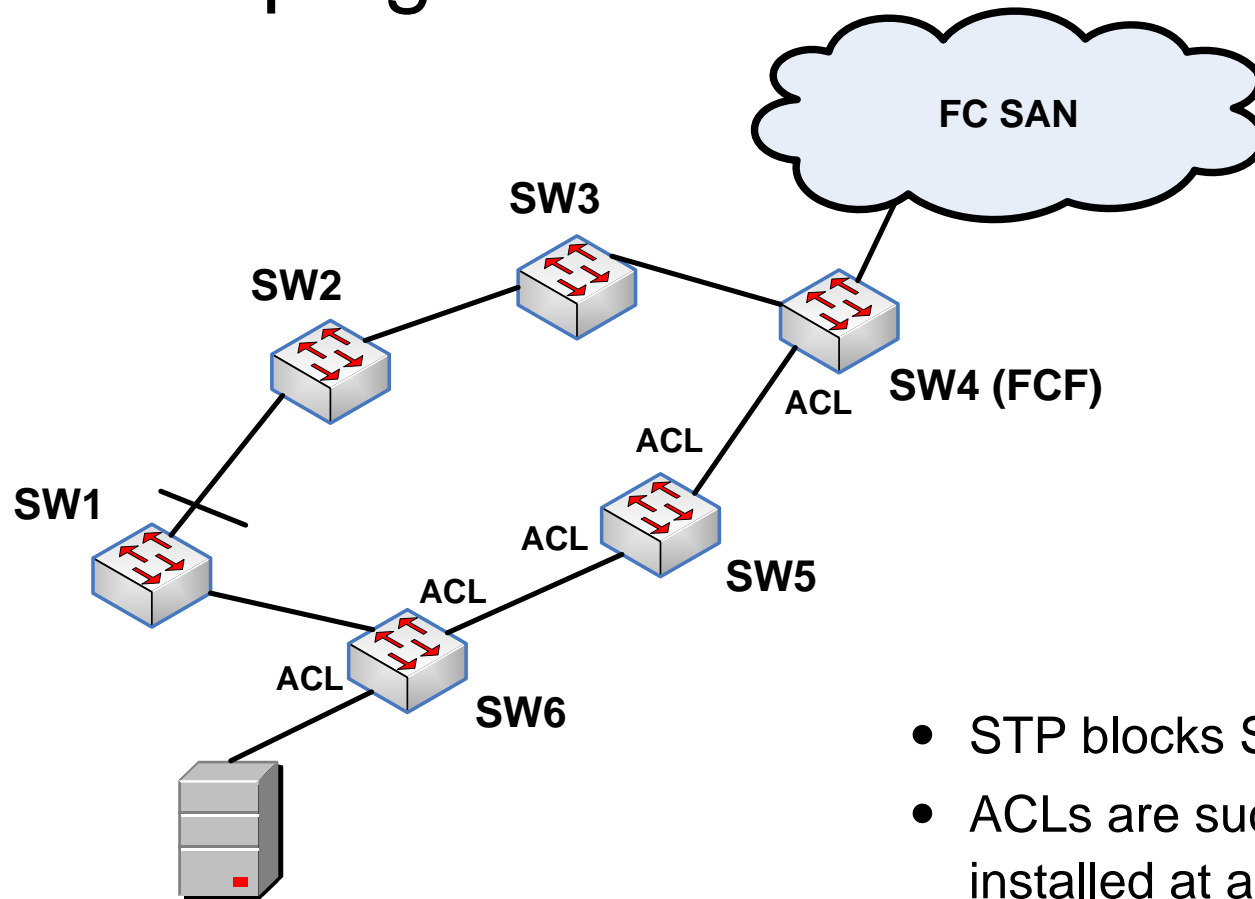
Some Observations About the Snooping Examples

- IGMP snooping
 - Any state installed will age out in the absence of messages that caused the state to get installed
 - The aging out of state results in default configured behavior – either flood or filter
- DHCP snooping
 - State is held for as long as the DHCP lease
 - Enabled only on edge switches
 - A switch reboot would result in traffic being affected unless the state is preserved
- For both IGMP snooping and DHCP snooping
 - The data must use the same path as the control traffic
 - Forward and reverse paths must be symmetric

Snooping and FCoE


- FIP has its own protocol type for easy snooping
- The expectation is that FIP messages are snooped at all bridges along the path between the ENode and FCF
 - Install an ACL that permits FCoE traffic for which a successful FLOGI was performed
- What happens in the following scenarios?
 - The ENode fails and the FLOGI session is terminated by the FCF
 - A link in the middle of the network goes down and STP reconfigures
 - A multipath-capable technology such as TRILL is in use

FIP Snooping



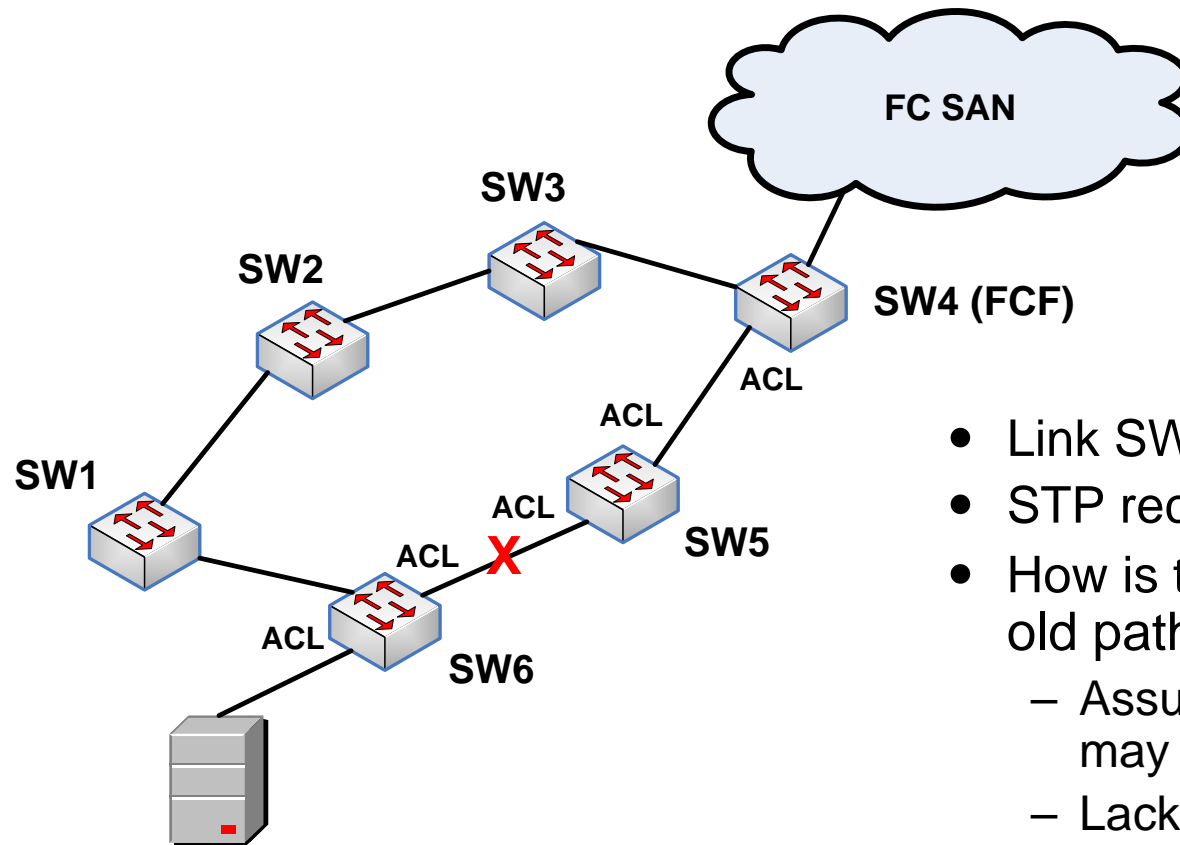
- STP blocks SW1-SW2
- ACLs are successfully installed at all switches along the path SW6-SW5-SW4

FIP Snooping (2)



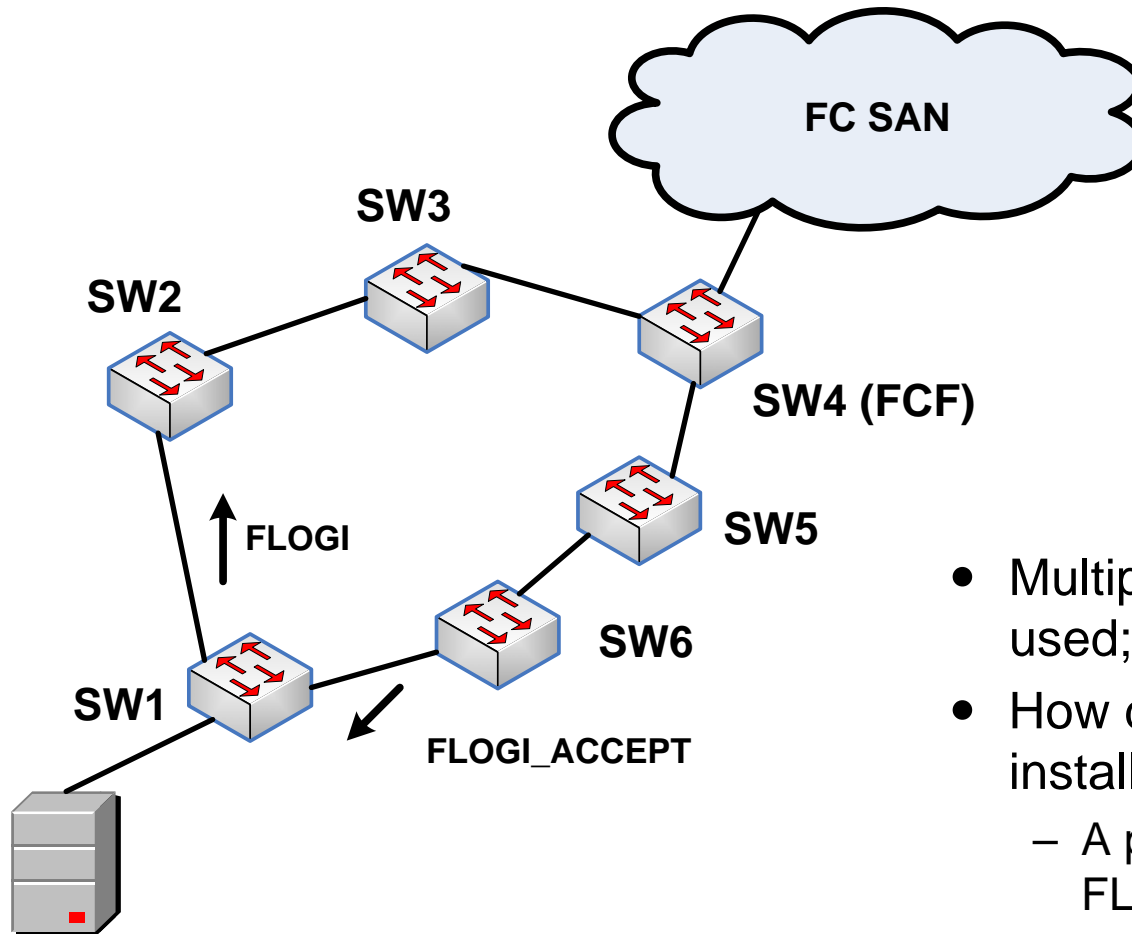
- Link to the server fails
- It's easy for SW6 to recognize this and remove the ACL
- FIP reset/LKA timeout could remove the state from SW6-SW5-SW4

FIP Snooping (3)



- Link SW6-SW5 fails
- STP reconfigures
- How is the ACL removed from the old path?
 - Assuming FIP reset is unicast it may not get to SW5
 - Lack of LKA could be used
- How is the ACL installed on the new path?
 - A problem because there is no FLOGI!

FIP Snooping (4)



- Multipath-capable technology is used; e.g. TRILL
- How does the ACL get installed/removed?
 - A problem because FLOGI, FLOGI_ACCEPT, and FCoE data may all go on different paths (asymmetric routes)

What Works?

- Snooping at the edge
 - Connected to either an ENode or FCF
 - Edge port can easily be determined by LLDP
 - ACL gets installed during FLOGI
 - ACL gets removed when the physical link goes down

Problem Scenarios for Snooping

- Installation of state at intermediate bridges along a path is a problem
 - STP reconvergence after a link failure
- Removal of state is less of a problem
 - Can be addressed by lack of LKA
 - However, it requires switches to snoop **a lot** of messages – all LKAs from all sessions going through it
- Asymmetric paths
 - Not a problem in networks with xSTP since learning would break as well
 - Will be a problem in networks deploying TRILL/SPBB with ECMP
 - Forward and return paths may be different
 - Even in the same direction FIP and FCoE data may use different paths



Summary

- There are open issues for snooping in FCoE switches
 - It is possible to solve them
 - Requires end station awareness of changes to network topology, multiple paths, etc.
- Recommendations for this group
 - FC-BB-5 should recommend snooping to be performed only on edge ports similar to what is done with access control problems in existing bridged LANs
 - Defer solutions for dealing with snooping and ACLs at intermediate switches to FC-BB-6