


BROCADE



FIP Keep Alive

Some Issues and Solutions

T11/08-434v0

Anoop Ghanwani
anoop@brocade.com

August 6, 2008

Overview

- The need for FKA
- FKA mechanism
- Scalability concerns
- Do we always need FKA?
- Optimized FKA
- Summary

The need for FKA

- FKA was introduced in FCoE because the VE_Port and the VF_Port need not be connected by a single physical link
- The primary function of FKA is a mechanism for the VF_Port to know that the VN_Port is not reachable because the physical link is no longer a reliable indicator
 - A unidirectional message flow; no response from VF_Port is needed
 - VN_Port finds out about unreachable VF_Ports by the absence of periodic advertisements multicast from the FCF
- Additionally, FKA has been found to be useful for
 - Refreshing the forwarding tables in learning bridges in the absence of actual exchange traffic
 - As a mechanism of timing out automatically generated ACLs



FKA mechanism

- An ENode needs to send an FKA for every MAC address that it transmits with
- FKAs are periodic
 - T11/08-263v1 mentions an interval of 30 sec
 - The current specified default interval is 8 sec
 - There has been some discussion about reducing it even further to 4 sec
 - A VN_Port is detected as being non-reachable after missing 3 consecutive messages

Scalability analysis

- An ENode performs a FLOGI followed by ~200 FDISCs
 - With FPMA, this translates to ~200 FKAs every 4 sec
 - Probably not a big deal for the server
- An FCF allows ~4000 FLOGI/FDISCs
 - Results in ~4000 FKA messages every 4 sec
 - Could be a big deal for some implementations
- A high number of FKAs is also burdensome for intermediate bridges that need to snoop these

Do we always need FKA?

- FKAs are not useful when the ENode is directly connected to the FCF
 - Directly attached devices can be detected by LLDP (802.1AB)
- FKA can be turned off by the FCF by setting FKA_ADV_PERIOD to zero in the FIP Advertisement
 - However, we need a capability for an end station to request a FKA interval of 0

Optimized FKA

- A single FKA may be used between an ENode and an FCF
- Bridges at the edge (i.e. directly attached to the ENode) in which ACLs are implemented can snoop on FLOGI/FDISC responses to associate the list of MAC addresses with a given FKA
 - When a FKA timeout is detected, the ACLs for all MAC addresses associated with that FKA are removed
 - This doesn't work for bridges in the core, but neither does the installation of ACLs
 - See T11/ 08-283v1
- This won't keep the forwarding state in intermediate bridges refreshed
 - We could use a Keep Alive that the FCF doesn't care about
 - Can be sent with a much larger interval, e.g. once every 60 sec
 - The frequency for this need not be negotiated




Summary

- It should be possible for either end to disable FKA
 - FKA not useful when the ENode and FCF are attached to the same physical link
- Optimize FKA
 - Adopt a mechanism where a single FKA is used between an ENode and FCF
 - For keeping the bridge forwarding tables refreshed, we can use lower frequency keep alives

Proposed Changes

- Add FKA timer TLV in the FLOGI/FLOGI response messages
 - If set to 0 by either end, it means FKA is disabled
 - ENode gets to pick the value first
 - The FCF can respond with
 - The same timer value that it receives, or
 - A larger timer value, or
 - Zero
- Use of per-MAC-address Keep Alive
 - Local decision, so no change required to Advertisement/FLOGI/FDISC
 - FCF simply ignores these if it receives them

BROCADE



THANK YOU

