



FCoE MAC Addressing

Uri Elzur

Pat Thaler

V 0.1

December 4, 2007

Considerations

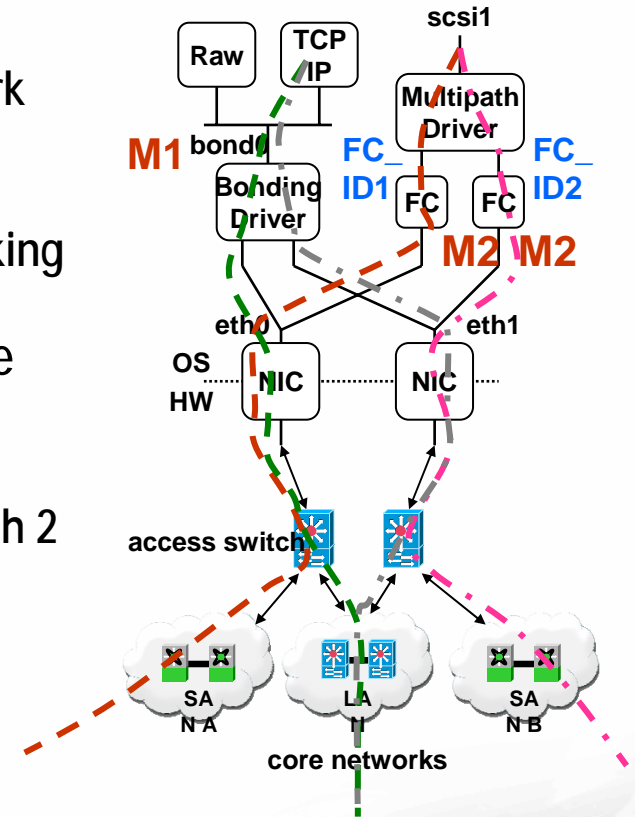
- Complexity NIC and Switch
- Scalability
 - Virtualization
- Security
- Provisioning

NPMA – NIC side

- New mechanism for Ethernet NIC Hardware
 - Frame filtering uses high speed circuitry, not easy to touch in operation
- Proposed behavior
 - Discovery – use the Ethernet MAC
 - May get multiple responses
 - After receiving the fabric FLOGI-Response
 - Switch to using the new FCoE MAC address on the fly (while HW is actively Tx/Rx)
 - Keep both Ethernet MAC and new FCoE MAC (FCoE MAP || FC_ID) operational
 - May receive more discovery responses after new MAC is used
 - Should not drop them – alternate switch for other VM, redundancy? LB? etc.
 - **NIC Must filter based on EtherType AND MAC !**
 - » **Adds more fields to wire speed hardware!**
- Scalability – one MAC per OS/GOS per stack, 2 per OS/GOS per stack for Failover / Load Balancing

SPMA – NIC side

- One MAC address shared between Ethernet and FCoE – possible, but not practical
 - Ties Failover and Load Balancing together for Network and Storage stacks
 - Administrator used to separate configuration
 - Have to allow for OS driven MAC override for networking => 2 MAC addresses
- Two MAC addresses (one per stack) have none of the issues
 - Addresses issues mentioned before (e.g. 07-546v1)
 - Note also, FCoE allows one MAC to be associated with 2 FCFs with 2 distinct FC-ID
 - OK if SAN and LAN fabrics are separated
 - Higher risk of miss-configuration
- Scalability – one MAC per OS/GOS per stack, 2 per OS/GOS per stack for Failover / Load Balancing
- T11 FC-BB-5 should not eliminate shared MAC option, but should also allow for separate server provided MAC address for LAN and for FCoE



Switch side - NPMA

- Edge Ethernet switch(/s) ACL Learning
 - Any MAC can be used for Discovery
 - Capture and parse the FC FLOGI Response to get the FC_ID
 - Translate FC_ID into FC_MAP || FC_ID as a MAC address
 - How does the switch learn the MAP?
 - or: Pushed by an FC or FCF switch to the edge (more?) Ethernet switch
 - Out of scope?
- Requires Dynamic update of ACL (
 - As described in 07-656v0 slide 7, with new SA = FC_OUI || XX.XX.XX.XX entry
 - ACL are not standardized, but **static ACL are broadly supported in switch ASICs**
- Potential race with multiple VN_Node receiving FLOGI response simultaneously e.g. blade server with virtualization
 - VN_Node may try to start an FC Exchange, while switch ACL has not yet been updated
- Tear down of ACL
 - E.g. server reboot may leave the old FC_MAP || FC_ID enabled by the switch ACL
- Scalability
 - Requires ONE entry per FC_ID for access control and zoning
 - Each entry is 48 bit, as the Ethernet switch supports a mix of Eth (must be 48bit) and FCoE (may be 24 only + one MAP for the whole switch) ACL
 - Not clear that one MAP per switch is sufficient for Virtualization

Switch side - SPMA

- Separate MAC for FCoE simplifies the ACL on the switch
 - ACL can be based on MAC address rather than MAC address plus type.
 - Static configuration of ACLs
- Edge switch FCoE Link setup and data transfer – regular MAC learning and forwarding
- Switch ACL Learning
 - Any MAC can be used for Discovery (similar to NPMA)

Switch side - SPMA

- Background info: In 802.1Q, two kinds of address learning are defined:
 - SVL - for a set of VLANs, when an address is learned in one VLAN in the set, it is used for all VLANs in the set
 - IVL - for a set of VLANs, when an address is learned in one VLAN, that learning is not used in any other VLANs in the set
 - IVL Bridges are common in the data center
- Edge Switch Scalability Per Server
 - Common Switch ASICs support large MAC table entries

SPMA – Switch Hardening

- Goal: allow the Ethernet switches to easily identify FCoE traffic to extend the notion of the “Channel” to the FCoE HBA over an Ethernet network
 - One way
 - Use of a well known FCoE MAC addresses with a fixed (within an FC fabric or virtual fabric) upper 24 bits - the FC-MAP.
 - Then, allow switches to have an ACL entry rejecting all packets that use those upper 24 bits from other sources
 - Another way
 - Use existing 802.1Q defined behavior to protect against this attack
 - FCoE traffic uses FCoE dedicated VLAN (or VLANs)
 - All switches supporting IVL can now block non-FCoE traffic except for the designated VLAN/s
 - In operation can support
 - Static enabling of addresses for VLAN access or
 - Learning addresses from login snooping:
 - » Use the port MAC address and Login VLAN for Discovery and FLOGI
 - » For an FCoE VLAN, usage is enabled for the MAC address used for FLOGI response
 - Now a rogue cannot attack with a non-FCoE packet

Virtualization

- Virtualization requires trust in HW i.e. HBA.
- For SPMA, that HBA can (and should) prevent "rogue" host from using a legal FC-ID on the single FCoE MAC for non-authorized traffic
- NPIV
 - Arguably, the switch ACL can fulfill its role when inspecting one FCoE MAC for that physical HBA
 - Possible with SPMA
 - NPMA requires one FCoE MAC per NPIV!
- Migration
 - WWNN and WWPN are migrated. FC_ID is not carried over
 - New FLOGI is required, new FC_ID provided. No need to migrate MAC
 - Similar switch mechanisms for learning and ACL can be used as normally used
- Scalability - Storage in HV/VMK
 - SPMA may use one MAC and one switch ACL entry per HBA
 - NPMA requires multiple MAC and multiple ACL entries, one per FC_ID
- Scalability - Storage in VM
 - SPMA may use one MAC and one switch ACL entry per HBA
 - NIC may filter based on FC_ID
 - NPMA requires multiple MAC and multiple ACL entries, one per FC_ID

Summary

- SPMA
 - Eliminates NIC complexity in dealing with changing address on the fly
 - Simpler Switch learning and ACL
 - Allows better Scalability
 - Virtualization
 - But requires, a translation table for the FCF
 - Addresses Security concerns
 - VLANs with IVL switches can protect FCoE traffic paths
 - Can be hassle free Provisioning
 - Static ACL configuration or Dynamic learning from FLOGI
- T11 FC-BB-5 should not eliminate shared MAC option, but should also allow for separate server provided MAC address for LAN and for FCoE
- T11 FC-BB-5 should mandate use of SPMA