An FCoE Direct End-to-End Connectivity Scheme

Michael Ko (michael@huaweisymantec.com), Li Ke, Wang Yuchen
Wu Xueping, Yang Qiang, Liu Lifeng, Meng Jian, Yang Qinqin
Huawei Symantec Technologies Co., Ltd.

Fibre Channel over Ethernet (FCoE) is a new standard that maps Fibre Channel frames directly over enhanced Ethernet. Central to the FCoE standard is the FCoE Forwarder which converts Fibre Channel frames to/from FCoE frames. Since the FCoE standard maintains the switch centric paradigm of Fibre Channel, all FCoE end nodes only know the Ethernet address of the FCoE Forwarder. Therefore all control and data communication with other FCoE end nodes must pass through the FCoE Forwarder. This bottleneck situation is unavoidable when there are legacy Fibre Channel devices requiring an FCoE Forwarder to act as a conduit between the Fibre Channel SAN and the Ethernet network. However, this problem will become intolerable when frame conversion is no longer needed as users replace existing Fibre Channel devices with FCoE ones on the Ethernet network.

The solution proposed in this report is to retain the use of the existing FCoE Forwarder for managing the control plane functions. The data plane functions, i.e., sending and receiving data packets, will be conducted directly between the end nodes, bypassing the FCoE Forwarder. We will call such end nodes that can support direct end-to-end connections without the FCoE Forwarder as shortcut capable.

To be connected to the fabric, an FCoE end node will perform fabric login with the FCoE Forwarder to have a Fibre Channel identifier assigned. An end node can then proceed to register itself with the name server and to query other objects in the name server. So far we are just following the current FCoE standard.

With the current FCoE standard, since an end node only knows the Ethernet address of the FCoE Forwarder, a means is needed to discover other FCoE end nodes that have logged in to the fabric and are shortcut capable. Accordingly, an FCoE end node will send a solicitation request to the multicast address of the shortcut capable group address. All shortcut capable end nodes will enable reception of frames sent to this multicast address. The solicitation request identifies the originator as shortcut capable. It also contains its Ethernet address and its Fibre Channel identifier. Each shortcut capable end node that receives such a solicitation request will add the originator to its list of shortcut capable end nodes. It will also reply with an advertisement identifying itself as shortcut capable. Likewise, the advertisement contains its Ethernet address and its Fibre Channel identifier. When the originator end node receives such an advertisement, it will add the other end node to its list of shortcut capable end nodes.

With the Ethernet address information available, a shortcut capable end node can now send a Fibre Channel port login directly to the other end node without requiring the FCoE Forwarder to act as an intermediary. After port login, the two end nodes can then send FCoE frames directly to each other and bypass the FCoE Forwarder. For interoperability with end nodes which are not shortcut capable, connectivity is through the FCoE Forwarder as currently defined in the FCoE standard.

For access control, Fibre Channel uses zoning, a function normally managed by switches. A shortcut capable end node can support zoning by consulting the fabric zone server in the FCoE Forwarder for the necessary information. This allows a target to determine if an initiator belongs to an authorized zone, and what volumes are visible to the initiator. Because the target no longer relies on the switches for zoning support, FCoE switches are not required. Instead, cheaper switches that simply support enhanced Ethernet features can be used for interconnection.

Looking towards the future, when bridging to legacy Fibre Channel devices are no longer necessary as native FCoE devices take over, the functions required of an FCoE Forwarder degenerate to that of just providing the control plane support. This includes the Fibre Channel control plane functions such as handling fabric login, assignment of Fibre Channel identifiers, name server, fabric zone server, etc., and the FCoE control plane functions such as the handling of FCoE Forwarder discovery and the Keep Alive messages, etc. Since the FCoE Forwarder is no longer in the time critical data path, it can be implemented in software running on a server that is connected to the fabric, thereby providing a cost effective alternative.