

FCOE MULTI-HOP; DO YOU CARE?

There is a lot of discussion in the industry around FCoE's current capabilities, and specifically around the ability to perform multi-hop transmission of FCoE frames and the standards required to do so. A recent discussion between Brad Hedlund at Cisco and Ken Henault at HP (<http://bit.ly/9Kj7zP>) prompted me to write this post. Ken proposes that FCoE is not quite ready and Brad argues that it is.

When looking at this discussion remember that Cisco has had FCoE products shipping for about 2 years, and has a robust product line of devices with FCoE support including: UCS, Nexus 5000, Nexus 4000 and Nexus 2000, with more products on the road map for launch this year. No other switching vendor has this level of current commitment to FCoE. For any vendor with a less robust FCoE portfolio it makes no sense to drive FCoE sales and marketing at this point and so you will typically find articles and blogs like the one mentioned above. The one quote from that blog that sticks out in my mind is:

“Solutions like HP's upcoming FlexFabric can take advantage of FCoE to reduce complexity at the network edge, without requiring a major network upgrades or changes to the LAN and SAN before the standards are finalized.”

If you read between the lines here it would be easy to take this as 'FCoE isn't ready until we are.' This is not unusual and if you take a minute to search through

articles about FCoE over the last 2-3 years you'll find that Cisco has been a big endorser of the protocol throughout (because they actually had a product to sell) and other vendors become less and less anti-FCoE as they announce FCoE products.

It's also important to note that Cisco isn't the only vendor out there embracing FCoE: NetApp has been shipping native FCoE storage controllers for some time, EMC has them road mapped for the very near future, Qlogic is shipping a 2nd generation of Converged Network adapter, and Emulex has fully embraced 10Gig Ethernet as the way forward with their OneConnect adapter (10GE, iSCSI, FCoE all in one card.) Additionally support for FCoE switching of native Fibre Channel storage is widely supported by the storage community.

Fibre Channel over Ethernet (FCoE) is defined in IEEE FC-BB5 and requires the switches it traverses to support the IEEE Data Center Bridging (DCB) standards for proper traffic treatment on the network. For more information on FCoE or DCB see my previous posts on the subjects

(FCoE: <http://www.definethecloud.net/?p=80>,

DCB: <http://www.definethecloud.net/?p=31>.)

DCB Has four major components, and the one in question in the above article is Quantized Congestion Notification (QCN) which the article states is required for multi-hop FCoE. QCN is basically a regurgitation of FECN and BECN from

frame relay. It allows a switch to monitor its buffers and push congestion to the edge rather than clog the core. In the comments Brad correctly states that QCN is not required for FCoE, the reason for this is that Fibre Channel operates today without any native version of QCN, therefore when placing it on Ethernet you will not need to add functionality that wasn't there to begin with, remember Ethernet is just a new layer 1-2 for native FC layers 2-4, the FC secret sauce remains unmodified. Remember that not every standard defined by a standards body has to be adhered to by every device, some are required, some are optional. Logical SANs are a great example of an optional standard.

Rather than discuss what is or isn't required for multi-hop FCoE I'd like to ask a more important question that we as engineers tend to forget: Do I care? This question is key because it avoids having us argue the technical merits of something we may never actually need, or may not have a need for today.

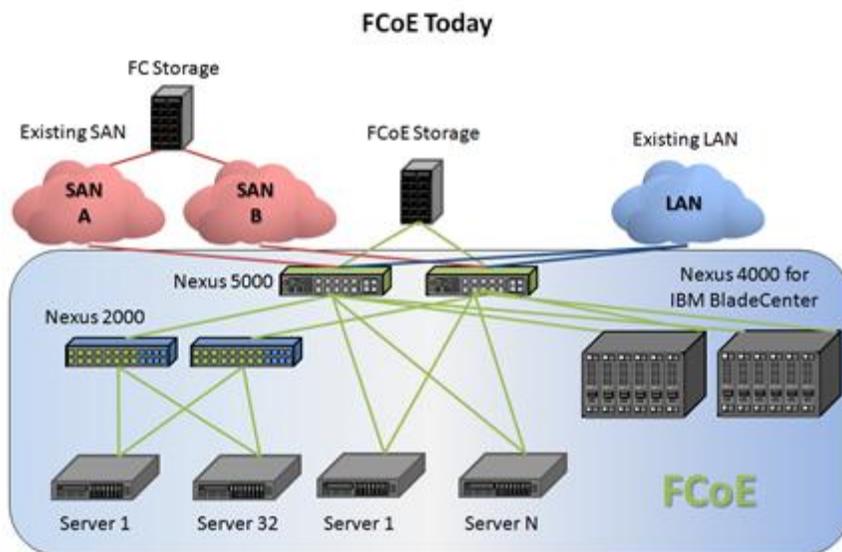
Do we care?

First let's look at why we do multi-hop anything: to expand the port-count of our network. Take TCP/IP networks and the internet for example, we require the ability to move packets across the globe through multiple routers (hops.) This is in order to attach devices on all corners of the globe.

Now let's look at what we do with FC today: typically one or two hop networks (sometimes three) used to connect several hundred devices (occasionally but rarely

more.) It's actually quite common to find FC implementations with less than 100 attached ports. This means that if you can hit the right port count without multiple hops you can remove complexity and decrease latency, in Storage Area Networks (SAN) we call this the collapsed core design.

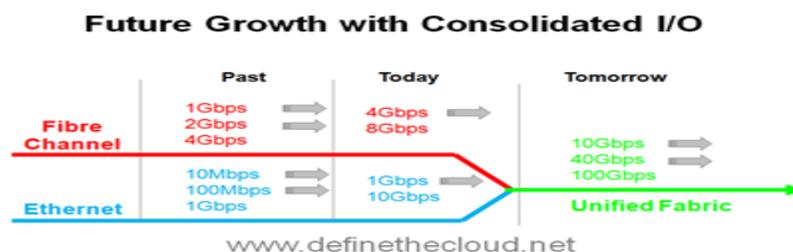
The second thing to consider is a hypothetical question: If FCoE were permanently destined for single hop access/edge only deployments (it isn't) should that actually stop you from using it? The answer here is an emphatic no, I would still highly recommend FCoE as an access/edge architecture even if it were destined to connect back to an FC SAN and Ethernet LAN for all eternity. Let's jump to some diagrams to explain. In the following diagrams I'm going to focus on Cisco architecture because as stated above they are currently the only vendor with a full FCoE product portfolio.



In the above diagram you can see a fairly dynamic set of FCoE connectivity options. Nexus 5000 can be directly connected to servers, or to Nexus 4000 in IBM BladeCenter to pass FCoE. It can also be connected to 10GE Nexus 2000s to increase its port density.

To use the nexus 5000 + 2000 as an example it's possible to create a single-hop (2000 isn't an L2 hop it is an extension of the 5000) FCoE architecture of up to 384 ports with one point of switching management per fabric. If you take server virtualization into the picture and assume 384 servers with a very modest V2P ratio of 10 virtual machines to 1 physical machine that brings you to 3840 servers connected to a single hop SAN. That is major scalability with minimal management all without the need for multi-hop. The diagram above doesn't include the Cisco UCS product portfolio which architecturally supports up to 320 FCoE connected servers/blades.

The next thing I've asked you to think about is whether or not you should implement FCoE in a hypothetical world where FCoE stays an access/edge architecture forever. The answer would be yes. In the following diagrams I outline the benefits of FCoE as an edge only architecture.



The first benefit is reducing the networks that are purchased, managed, power, and cooled from 3 to 1 (2 FC and 1 Eth to 1 FCoE.) Even just at the access layer this is a large reduction in overhead and reduces the refresh points as I/O demands increase.



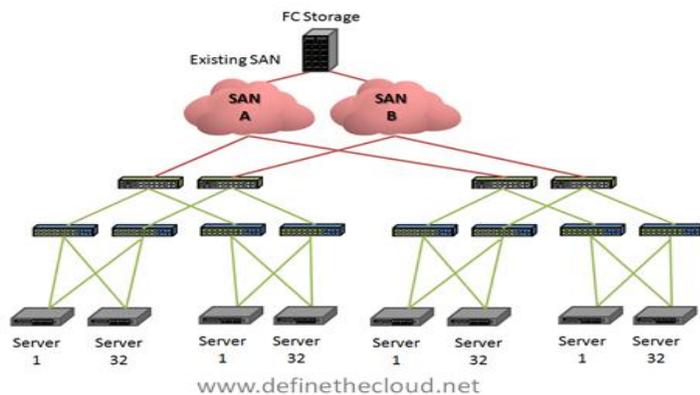
The second benefit is the overall infrastructure reduction at the access layer.

Taking a typical VMware server as an example we reduce 6x 1GE ports, 2x 4GFC ports and the 8 cables required for them to 2x 10GE ports carrying FCoE. This increases total bandwidth available while greatly reducing infrastructure. Don't forget the 4 top-of-rack switches (2x FC, 2x GE) reduced to 2 FCoE switches.

Since FCoE is fully compatible with both FC and pre-DCB Ethernet this requires 0 rip-and-replace of current infrastructure. FCoE is instead used to build out new application environments or expand existing environments while minimizing infrastructure and complexity.

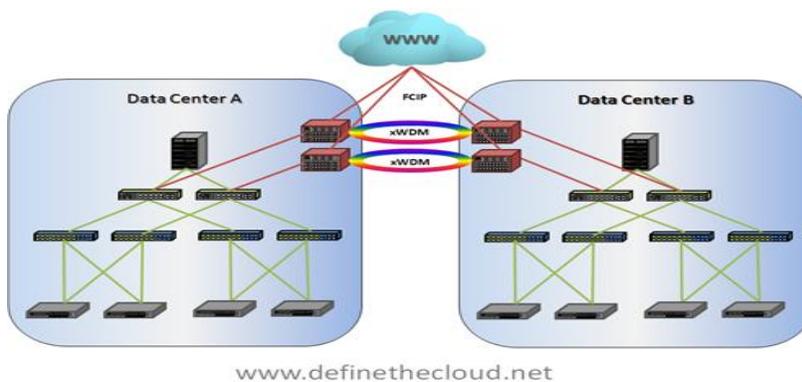
What if I need a larger FCoE environment?

If you require a larger environment than is currently supported extending your SAN is quite possible without multi-hop FCoE. FCoE can be extended using existing FC infrastructure. Remember customers that require an FCoE infrastructure this large already have an FC infrastructure to work with.



What if I need to extend my SAN between data centers?

FCoE SAN extension is handled in the exact same way as FC SAN extension, CWDM, DWDM, Dark Fiber, or FCIP. Remember we're still moving Fibre Channel frames.



Source: <http://www.definethecloud.net/fcoe-multi-hop-do-you-care/>