

**Calculating Fibre Channel switch port Buffer Credit  
requirements for Distance SAN's (v6)**

Noelle Milton Vega  
Rensselaer Technology Group, LTD.

## Introduction

When purchasing a fibre channel switch (and licenses) for the purpose of implementing a multi-site SAN fabric, one switch feature you will need to consider is its buffer credit capability. In particular, how much of it will you need to ensure the optimal streaming of data over the distance covered by the two sites? This is important to know, since having too few buffers can cause application I/O to block as it waits for acknowledgment from the far-end; while having too many buffers will cause excess I/O to queue up in the switch (more on that later), not to mention unnecessarily increase the cost of the switch. The aim of this white paper is to demonstrate how to right size a FC switch with respect to buffer credits.

The key to understanding Fibre Channel buffers, and how many are required to accommodate adjacent FC ISL connected ports, is to first understand the concept of "line frame-length". **Line frame-length** is the number of frames (of whatever type) that can exist in transit, at any one time, along the fibre optic path that connects two adjacent nodes/sites. As we'll see, although the linear distance between ends of a fibre optic path is fixed, it's line frame-length actually varies as follows:

- (1) It increases linearly as the transmission speed of the equipment in use **increases** (e.g. 1.0 Gbps<sub>fc</sub> -> 2.0 Gbps<sub>fc</sub> -> 10.0 Gbps<sub>fc</sub>), and vice versa,
- (2) It increases as the size of the frame decreases, and vice versa. For FC frames, the frame size is determined by its payload and frame header sizes.


If a fibre optic path happens to represent an FC ISL link, then, in knowing its **line bit-length**, we can calculate the maximum number Fibre Channel frames that can simultaneously fit (source to destination and back) along that link. And that is the number of buffers your ISL switch ports will require, and for the following reasons:

- (1) To safely store-and-forward frames at the source end, while they are in transit (on a piece of glass) to the destination end, and ultimately acknowledged as having been received,
- (2) To, in the meantime, provide an immediate "successfully sent" I/O acknowledgment to the sending application so it does not have to block waiting for the frame to reach the opposite end, and then for the acknowledgment coming back.

Consider the following practical example. A company decides to locate its disaster recovery site 100 miles away from its primary operating site. It purchases Fibre Channel Switch and DWDM equipment with 2Gbps<sub>fc</sub> interfaces, and with an appropriate number of buffers to accommodate continuous streaming of data between end ISL ports. Several months after the implementation, monitoring of the ISL links reveals congestion between the two sites. As a remedy, the company considers upgrading its end-to-end Switch and DWDM equipment to support 10Gbps<sub>fc</sub> transmission rates. However, doing so, as indicated previously, will increase the "effective distance" or **line bit-length** between the two sites by 5 fold. Stated differently, it's as if the DR site were moved from being 100 miles away to being 500 miles away.

## Calculating Fibre Channel switch port Buffer Credit requirements for Distance SAN's

Why? Because, as illustrated in the diagram below, relative to the  $2\text{Gbps}_{\text{FC}}$  speed, at  $10\text{Gbps}_{\text{FC}}$ , the physical distance between each inserted bit is 5 times smaller. This *relative compression* is due to the faster rate at which optical variations (bits) are now being introduced into the front end of the DWDM pipe by the FC switch ISL port (5 times as fast in our example).



*Note that the diagram above is not drawn to scale.*

Above, the gray boxes for the  $2\text{Gbps}_{\text{FC}}$  and  $10\text{Gbps}_{\text{FC}}$  cases contain the same number of optical variations (bits) in transit. However, the gray box for the  $10\text{Gbps}_{\text{FC}}$  case consumes  $1/5^{\text{th}}$  the physical fibre distance relative to its  $2\text{Gbps}_{\text{FC}}$  counterpart. Stated differently, relative to the  $2\text{Gbps}_{\text{FC}}$  case, at  $10\text{Gbps}_{\text{FC}}$ , 5 times the amount of bits (data frames) can fit along the same fibre optic path.

This means that for the  $10\text{Gbps}_{\text{FC}}$  case, the same frame, while traveling the same physical distance, will now have 5 times as many frames behind it (in transit on the glass) when it reaches the opposite end (again, because of the compression relative to the  $2\text{Gbps}_{\text{FC}}$  case). If, for example, the data in the gray boxes happened to represent one Fibre Channel frame then, relative to the  $2\text{Gbps}_{\text{FC}}$  case, at  $10\text{Gbps}_{\text{FC}}$  the application could issue the equivalent of 5 times as many FC frames before ever receiving an I/O acknowledgement for the first frame. That's because, at the input end, each frame is being inserted and acknowledged locally in  $1/5^{\text{th}}$  the time. The resulting additional frames must be buffered at the source end to prevent the application from blocking while it waits for I/O acknowledgment.

Note that for a variety of reasons, the remote (destination) site storage array and/or host equipment can cause the source site application I/O to block. For example, if the destination site employed slower storage and/or host equipment relative to the source site, then during data bursts the destination equipment may not have enough horsepower to process the additional incoming I/O's fast enough. So, while allocating an appropriately sufficient number of switch buffer credits cannot guarantee continuous streaming of data between two sites, it does guarantee that the end-to-end ISL based fabric is not the cause of such a shortcoming.

## Calculating Fibre Channel switch port Buffer Credit requirements for Distance SAN's

Before we begin an example of how to calculate buffer requirements, it is important to know the numerical definition of a Fibre Channel Gigabit, as well as to understand the structure of a Fibre Channel Frame.

In the Fibre Channel world, one gigabit is defined to be 1,062,500,000 bits (which is not  $(1024)^3$ ). Other Fibre Channel gigabit values are then derived from this reference definition. For example, two (Fibre Channel) gigabits =  $2 \times 1,062,500,000$  bits = 2,125,000,000 bits. To avoid confusion with the traditional (non Fibre Channel) definition of a gigabit, throughout this document I will use the symbol  $Gb_{fc}$  to mean "1,062,500,000 bits", or 1 Fibre Channel Gigabit.

In summary:

- 1  $Gb_{fc}$  = 1,062,500,000
- 2  $Gb_{fc}$  = 2,125,000,000
- 10  $Gb_{fc}$  = 10,625,000,000

Next, we show the anatomy of a Fibre Channel Frame with notes.

<b>Start of Frame</b>	4 bytes	32 bits
<b>Standard Frame Header</b>	24 bytes	192 bits
<b>Data (payload)</b>	[0 - 2,112] bytes	[0 - 16,896] bits
<b>CRC</b>	4 bytes	32 bits
<b>End of Frame</b>	4 bytes	32 bits
<b>TOTAL (Nbr bits/frame):</b>	[36 - 2,148] bytes	288 - 17,184 bits

### Notes:

The term byte used here, and in the Table 3 above means 8 bits (not the 10 bits that result from 8/10 bit encoding).

The maximum Fibre Channel frame size is 2,148 bytes.

The final frame size must be a multiple of 4 bytes. Thus the Data (payload) segment will, as necessary, be padded with 1 to 3 "fill-bytes" to achieve an overall 4 byte frame alignment.

The standard Frame Header size is 24 bytes. However, up to 64 additional bytes (for a total of an 88 byte header) can be included for applications that need extensive control information. Since the total frame size cannot exceed the maximum of 2,148 bytes, these additional Header bytes will subtract from the Data segment size by as much as 64 bytes (per frame). This is why the maximum Data (payload) size is 2,112 (because  $[2,112 - 64] = 2,048$ , which is exactly 2K-bytes of data).

The final frame, once constructed, is passed through the 8byte to 10byte conversion process.

In the FC world, 1 Word =  $4 \times 8/10$  bit encoded bytes (40 bits).

## Calculating Fibre Channel switch port Buffer Credit requirements for Distance SAN's

### Our Example

Let say, for the purposes of discussion, that we have redundant (i.e. two) fibre optic paths between a primary and secondary site.<sup>1</sup> The linear distance (as opposed to displacement) is 91.732608 Km (or about 57 miles) for one path, and 43.452288 Km (or about 27 miles) for the other path.

We start first with the speed at which all electrical variations (i.e. baud)<sup>2</sup> propagate through transmission mediums. This is the speed of light:

299,792.458 km / s == .00000333564095 s / km
--

Speed of Light

Next we take the linear distance between the two sites and determine:

- (1) How many seconds it takes for 1-bit to travel the one-way distance linear between the two sites; that is, express the "distance" between the two sites in seconds. (Column 2 below). This is determined by the speed of light.
- (2) Having determined the one-way distance in seconds between the two sites (a fixed number), we can now determine the maximum number of bits that can exist, in transit, between the two sites at any one time. In other words, we can calculate the "distance" between the two sites in bits (as opposed to miles). This is determined by the speed of light, as well as by the rate at which the transmitting equipment (e.g. a fibre channel port) can create electrical variations onto the medium (i.e. fibre optic line). In other words, how fast can it push bits onto one end of the fibre optic line (usually 1 Gb<sub>fc</sub>, 2 Gb<sub>fc</sub>, or 10 Gb<sub>fc</sub>). (Columns 3 & 4).

Table 1

Distance (Km)	Distance (secs)	# of Bits (1Way)	8/10 FC "bytes"
91.732608 (57mi)	0.000305987044	325,111.234	32,511.1234
43.452288 (27mi)	0.000144941231	154,000.058	15,400.0058

Values for bits, bytes, & frames were calculated assuming a **1** Gb<sub>fc</sub> bit insertion rate (i.e. the rate at which bits are introduced into the frond end of the DWDM fibre optic line).

Table 2

Distance (Km)	Distance (secs)	# of Bits (1Way)	8/10 FC "bytes"
91.732608 (57mi)	0.000305987044	650,222.468	65,022.2468
43.452288 (27mi)	0.000144941231	308,000.116	30,800.0116

Values for bits, bytes, & frames were calculated assuming a **2** Gb<sub>fc</sub> bit insertion rate (i.e. the rate at which bits are introduced into the frond end of the DWDM fibre optic line).

<sup>1</sup> In DWDM terminology, redundant fibre optic paths is sometimes referred to as path protection.

<sup>2</sup> Note: 1-Baud (electrical variation) will represent a different number of bits depending on the compression codecs used. In our example, 1 baud represents 1 bit.

## Calculating Fibre Channel switch port Buffer Credit requirements for Distance SAN's

### Calculations notes for the tables above:

Column 2 indicates the amount of time (in seconds) it takes one **OPTICAL VARIATION** (however many number of bits that optical variation happens to represent) to travel the one-way distance specified in column one. It comes from multiplying the number of seconds it takes **light** to travel 1 km (i.e. .00000333564095 s / km), by the number of km traveled which, again, is specified in Column 1.

Column 3 is the product of column 2 (i.e. the line distance in seconds) and the FC bit insertion rate (either 1,062,500,000 for Table 1, or 2,125,000,000 for Table 2). This column effectively represents the amount of additional I/O that could have been processed at the host, had the response to that I/O been instantaneous; (actually twice that amount, since acknowledgment time for those I/O's, coming back the other way, have to be accounted for as well).

Column 4 is Column 3 divided by 10. This is done to group single bits in transit into equivalent 8/10 bit "byte" quantities. The Fibre Channel protocol converts every 8 bit byte into a 10 bit equivalent (via the 8/10 bit encoding algorithm) before transmitting it. So the value in this Column 4 will determine the number of buffers needed for an ISL switch port.

Comparing the two tables above, we can see that the faster the transmitting Fibre Channel port is, the more total number of bits that particular port can "insert" into the line before the very first one it inserted reaches the opposite end. That's because the time in between bit insertions is reduced for faster ports (i.e. an increased rate at which bits are being introduced into the front end of the DWDM pipe). Therefore, relatively speaking, a fixed length fibre optic path *effectively* gets longer (bit length wise) as the speed of the transmit/receive equipment increases. The line (bit) length for 2 Gbps<sub>FC</sub> xmit/recv equipment is twice as long as it is for 1 Gbps<sub>FC</sub> xmit/recv equipment. Similarly, the line (bit) length for 10 Gbps<sub>FC</sub> xmit/recv equipment is five times as long as it is for 2 Gbps<sub>FC</sub> xmit/recv equipment. The net/practical effect of this is that the faster the FC port, the larger the number of buffers that will be required for your switch ISL ports.

Given the **linear distance** between the two datacenters (in Kilometers) and the **speed of the fibre channel** equipment (1 Gbps<sub>FC</sub>, 2 Gbps<sub>FC</sub>, 10 Gbps<sub>FC</sub>), we have thus far been able calculate "**line bit length**" of the link between the two sites (i.e. the number of bits that can fit on the line, end to end, at any one time). Knowing the line bit length, we can divide this value by the number of bits in a Fibre Channel frame to determine the equivalent "**line Fibre Channel frame length**" (in other words, the number of FC frames that can fit on the line, end to end, at any one time). As we saw from table 1, the number of bits contained in a FC frame varies with: (1) the Data/Payload size, (2) padding, which ensures a FC frame whose final size is a multiple of 32 bits (4 bytes), and (3) the size of the Frame header, which can range from 24 to 88 bytes.

Calculating Fibre Channel switch port Buffer Credit requirements for Distance SAN's

$$\text{Frame Length8 (in km)} = (299,792.458 \text{ km/s}) \times (\text{Seconds-Between-Inserted-Bits } \textit{secs/bit}) \times (\text{Number-Of-Bits-Per-Frame } \textit{bits})$$

$$\text{Frame Length10 (in km)} = (299,792.458 \text{ km/s}) \times (\text{Seconds-Between-Inserted-Bits } \textit{secs/bit}) \times (\text{Number-Of-Bits-Per-Frame } \textit{bits}) \times 10/8$$

Where:

$$\text{Seconds-Between-Inserted-Bits} = 1/1,062,500,000 \text{ (for 1Gbps FC)} \\ = 1/2,125,000,000 \text{ (for 2Gbps FC)}$$

Number-Of-Bits-Per-Frame = Variable depending on Data payload and/or Header size. See table and notes above. In most cases the Header will be the Standard size of 24 bytes.

Table 4

Number of bits / frame (after 8/10 bit encoding)	Frame Length (km) @ 1 Gbps FC (.000282157607529411 km/bit)	Number of in-transit frames (1-way) for 91.732608 km DWDM leg.
17,184 / 21,480	6.0607 km (2112 PL-bytes)	15.1355 frames (buffers)
16,672 / 20,840	5.8801 km (2048 PL-bytes)	15.6003 frames (buffers)
8,480 / 10,600	2.9908 km (1024 PL-bytes)	30.6708 frames (buffers)
4,384 / 5,480	1.5462 km (0512 PL-bytes)	59.3268 frames (buffers)

**1Gbps (1,062,500,000 bps) data input rate.**

Table 5

Number of bits / frame (after 8/10 bit encoding)	Frame Length (km) @ 2 Gbps FC (.000141078803764705 km/bit)	Number of in-transit frames (1-way) for 91.732608 km DWDM leg.
17,184 / 21,480	3.0303 km (2112 PL-bytes)	30.2170 frames (buffers)
16,672 / 20,840	2.9400 km (2048 PL-bytes)	31.2006 frames (buffers)
8,480 / 10,600	1.4954 km (1024 PL-bytes)	61.3417 frames (buffers)
4,384 / 5,480	0.7731 km (0512 PL-bytes)	118.6537 frames (buffers)

**2Gbps (2,125,000,000 bps) data input rate.**

Notes:

- Column 1 represents the total number of 8/10 bits in a Fibre Channel frame (with a standard header size of 24 bytes) at varying Data payloads (PL). From top to bottom, the payloads used to calculate each row of column 1 are, respectively: 2112 bytes, 2048 bytes, 1024 bytes, and 512 bytes.
- Column 2 represents the product of the 10 bit value in Column 1, and (1/1,062,500,000) for the 1Gbps (Table 4), and (1/2,125,000,000) for the 2Gbps (Table 5). Thus, this Column (Column 2) essentially represents the linear distance that 1 (one) single frame consumes for the specified payload (PL).
- Column 3 represents the quotient derived by dividing the longest (worst case) DWDM line distance (in kilometers), by the number of kilometers per frame (calculated in Column 2). Thus, this column essentially indicates how many additional ONE-WAY frames worth of data could have been processed by the host/application, had the response to the first frame been instantaneous. In other words, this is how many ONE-WAY (not round trip) switch buffers you would need to allow non-stop transmission. Double (i.e. round trip) the values in this column 3 to yield the number of buffers required of your ISL switch ports.