
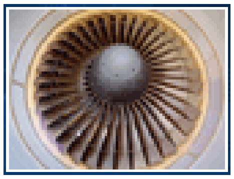
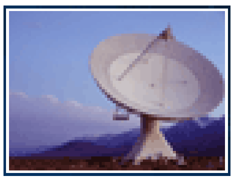


High Speed Digital Radar Recording

White Paper

prepared for
MIT – Lincoln Laboratory

by
DSPCon, Incorporated
Bridgewater, New Jersey

October 15, 2004

Preface

Purpose

Industry has been providing "high speed" digital recorders for many years for use in a wide variety of different applications. As processors, networks and storage become faster and less expensive, industry is able to offer faster and faster digital recording systems. While this technology is certainly moving very quickly, the market is demanding the capability to record even faster, and less expensively. Communication signals and Radar signals are constantly growing in terms of bandwidth. This technology shift along with today's political climate is resulting in a huge market demand for digital systems that can record much faster than yesterday's systems.

The purpose of this paper is to provide a board overview of the state of the art in the area of high-speed digital recording and the various key, commercially available technologies being used in industry. While the techniques and technologies discussed with have a board range of potential applications, DSPCon's primary application focus, during the time of this writing was Radar recording.

Audience

This document is intended for Engineers, Scientist or various staff level personnel who are interested in current, commercially available techniques for high-speed, digital data collection.

Contents

Introduction

Regardless of how you look at the digital recording challenge, the basic building blocks of such a system are the same today as they have been for years. This is particularly true if there is a desire to use Commercial Off The Shelf (COTS) building blocks to construct a recording system. A high-level block diagram of these basic building blocks is shown in Figure 1 below.

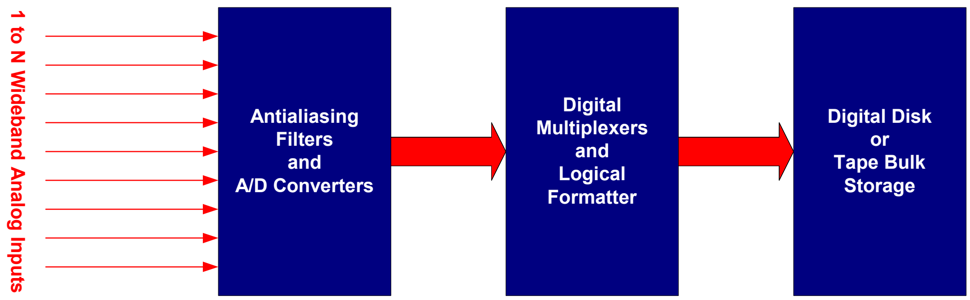

Figure 1–Basic System Blocks

These basic building blocks consist of the following major functions.

1. Filtering and Analog to Digital Conversion
2. A Digital Multiplexer and Logical Data Formatter
3. A Disk Array, Solid State Memory and/or some kind of Removable Media Device (DVD, Tapes, etc).

Please note that in general, the number of input channels can be traded off against the input bandwidth (sampling rate) as long as the aggregate system data rate (number of channels multiplied by the digitized byte rate of each channel) does not exceed the bandwidth of any bus in the system or the input bandwidth of the bulk storage system. For example, DSP*Con* has delivered 4-channel systems with 70 MHz of input bandwidth for each channel. DSP*Con* has also delivered 128-channel systems, with a very similar architecture, in which the input bandwidth was 105 kHz.

This following section discusses each of these blocks and describes the state of the art technology of each.

Filtering and A/D Conversion

Arguably, this is the most important part of the recorder. It is the only part of the recorder that handles the analog signals that will be captured. The first part of the block is responsible for removing spectral content on the signal, which is beyond our frequency of interest. This step is necessary to prevent an effect known as “aliasing”, in which these higher (unwanted) frequencies may corrupt the signal of interest as the data is being digitized. While a discussion of digitizing techniques is beyond the scope of this paper, it is a very important part of the front end of a digital recorder. A system that lacks proper anti-aliasing filters results in the recording of corrupted signals. It is important to note that this filtering must be done with analog components. There is no technique for removing the effects of aliasing once the data has been digitized.

Once the analog signals have been sent through the proper anti-aliasing filters, the data is ready to be converted into digital form. The concept here is actually quite simple. A/D converter chips are fed with a clock. Each clock edge causes the converter to quantize the analog level that is present at the converter at the time of the clock edge. The faster the provided clock, the more often the A/D quantizes. In order to capture all of the desired spectral content, one must sample (clock the A/D) at a theoretical minimum rate of twice the maximum spectral frequency bandwidth of interest–known as the “Nyquist rate”. For example, if we are interested in recording all analog content from DC, up to 2 MHz, we must sample at the rate of at least 4 MHz. As actual silicon does not behave at exactly the Nyquist rate, most engineers sample at a minimum sample rate of 2.5 times the maximum frequency of interest. This accounts for imperfect filter roll-off and other real-world effects. While many engineers believe that sampling at rates much higher than the Nyquist rate is required in order to capture all of the spectrum of interest, a DSP*Con* White Paper entitled “*Interpolation of Sparse Time History Data*” explains why this is not the case. It can be found in its entirety on the DSP*Con* web site (www.dspcon.com). To access the paper, choose “Download” from the menu choices, then choose “Tech Papers” and scroll to the document.

By quantizing (sampling) often enough, we can assure that we capture all of the spectral content of interest. Since we are using a finite number of bits in each sample to characterize an infinite number of points on each analog signal, how accurately can we capture a complex analog wave? The answer is driven by the number of bits in each sample. This is measured in a ratio of the actual signal to noise (or error) in the sample. This ratio is called the SNR (Signal-to-Noise Ratio) of the system. Years ago, if you wanted to sample at 20 MHz, you would be “stuck” with 8-bit samples (or less) that resulted in an SNR of about 42 dB. Today, it is very common to sample 20 MHz signals with 16-bit A/D converters, which yields a SNR of almost 80 dB. Compare this to the less than 30 dB achieved from any reel-to-reel analog recorder. This performance is far better than any analog recorder that has ever been made and usually far better than is required (given the noise associated with sensors, cabling, and other parts of the system).

One final note about the front end of the recorder. Analog recorders maintain a poor time correlation between each channel that is being recorded. Static and dynamic inter-channel time displacement errors create many microseconds of time discrepancy between tracks. Typically**,** DSP*Con* recording systems maintain a channel-to-channel time correlation better than .1 degrees at 1 kHz**.**

Digital Multiplexer and Formatter

Assuming that we are interested in 14 channels and 2 MHz analog bandwidth, we must sample at 5 MHz. In addition, assume that each sample contains 16 bits (2 bytes). This results in 140 Mbytes/second. As little as two years ago this would have overwhelmed even the most powerful processors. Today, several processor vendors, for example, Intel, Texas Instruments, and Motorola, are shipping processors capable of billions of operations per second. These powerful processors are able to manipulate this amount of data all in real time. At a minimum, this block must input all of the data from the A/D front end, and supply this data as a single time-packed stream to a disk or tape device. Better systems can actually analyze this data and provide useful statistics to the operator, while also streaming the data to the bulk storage. Arguably, the most important task that this block should perform is logical formatting of the data that will be written to the storage. This logical format allows end users to read the data and process it with algorithms that they may have purchased or developed. Typically, such formats include a specification of the number of channels being recorded, the rate at which they are being recorded, and time tagging of the data plus significant events.

One of the most important aspects of the digital multiplexer is its ability to process the data in addition to formatting the data. This offers a powerful solution to a large problem. The problem is how to manage the enormous quantity of data that is collected and needs to be archived by wide band recorders. Most often, these types of recorders are used to record long periods of data in order to be sure that specific events of interest are recorded. The powerful real-time data processors within a digital recorder offer two powerful ways to address the data reduction problem.

- Users may annotate blocks of interest with data marks that include long, text-based comments that can be searched for off line. As these marks can be searched for rapidly on either disk- or tape-based systems, data reduction can be performed quickly and inexpensively.
- The real-time data formatter can be programmed to search automatically for characteristics of signals of interest in the time and/or frequency domains. This applies to both the start and end of these events. Using this technique allows the recorder to record for much longer periods of time because only events of interests are recorded. Moreover, when a recording mission is concluded, there is only a small fraction of the data that needs to be archived.

Potential users of these all-digital, wideband recording systems must understand that planning is necessary to allow for recovery of specific data in a time-efficient manner. As the actual recording bandwidth per inch of media is increased by analog to digital conversion, users will not have the luxury of playing back data at increased speeds. There must be a tight firmware link between the real-time indexing of important data as it is recorded.

Bulk Storage

There are two ways that bulk storage systems can be characterized. The first is the interface standard that is used to get data in and out of the unit. The second is the type of physical media within the unit that is used to record the data.

Media Types

There are basically three different types of digital storage system media. Each has its benefits and its drawbacks. All, however, are digital and therefore offer all of the benefits detailed in the first part of this paper. Cost factors vary widely with each type of system, depending on environmental factors. As the military or intelligence markets no longer drive recording technology, systems tailored to those markets can be inordinately expensive. DSP*Con* strives to use COTS products wherever practical, but can use any digital medium which best fits the application.

- ***Digital Tapes***—These devices offer the obvious benefit of having removable media. They also can be flown at relatively high altitudes and still work well. Lastly, most modern tape devices hold large quantities of data up to 660 Gbytes per tape. One large disadvantage of tape devices is that they tend to be slow. Even the faster tape units are capable of handling only 32 Mbytes/second.

- ***Solid-state Devices***—Basically, a solid-state bulk storage device is a large amount of RAM configured to look like a disk or tape. These devices tend to be very mechanically robust and extremely fast. They do, however, also tend to be very expensive and are generally not removable.

- ***Disk-based Systems***—Disk-based storage systems are made from piecing together two or more individual disk drives to produce a higher-capacity device, a faster device, or both. Disk-based storage systems tend to be fast and inexpensive. In addition, depending on how the disks are physically mounted with the storage device, they can also be very mechanically robust and even removable (like tapes). There are basically two types of disk-based storage systems.

  JBOD—A JBOD (Just a Bunch Of Disks) is two or more individually addressable disks within a box. This is roughly equivalent to hooking up multiple disks to one interface. These devices are extremely inexpensive and offer large capacities as a result of adding together the individual capacities of all of the disks contained within the JBOD. Since many of today's disks are capable of speeds close to 60 Mbytes/second and 144 Gbytes/disk, a 60 Mbytes/second, 720 Gbyte (5-disk) JBOD can be purchased very inexpensively. Moreover, if software based striping techniques are used, sustained rates of well over 160 Mbytes per second are achievable.

  RAID—The problem with a JBOD is that while it allows you to add the capacities of each disk together, the slowest disk device within the JBOD determines the speed of the entire device. This is due to the fact that each disk is individually addressed for each write of data. A RAID is a variation on this theme. A RAID is also made up of a bunch of disks. However, unlike a JBOD, which writes its data to each disk in series, a RAID contains a piece of hardware that spreads each data byte across multiple disks. This allows us to add not only the capacities of each disk device, but also the individual speeds. If each disk is capable of 60 Mbytes/second, when five such devices are put inside of a RAID, the resulting speed is 300 Mbytes/second (5 times 60 Mbytes/second). Typically, RAIDs also contain an additional disk used to store the parity of each transfer. The benefit of this is that each data transfer to the RAID is read back and the parity is checked against the expected parity. If an individual device fails, that device will be immediately replaced with the Parity disk. As a result, RAIDs tend to be one of the most reliable types of bulk storage devices. Today, it is common to see RAID devices that have a capacity of 720 Gbytes, and are capable of write rates in excess of 300 Mbytes/second.

- ***Emerging Technologies*****—**As the past several years have shown, there is a good deal of innovation still taking place in data storage technology. These advancements include new media types, cost reductions (for both the recorder and the media), increases in recording rates, and significant increases in recording densities. In fact, over the past several years, recording densities have grown at the rate of approximately 30% roughly every 6 months. An important consideration to end users is how to stay current with technology trends without assuming the overwhelming financial burden associated with replacing all of the fielded recorders and dubbing all of the archived data from the outgoing media to the new media.

  Regardless of how one chooses to handle this problem, there is always a cost associated with staying current with technology–-digital recording systems tend to make this outlay less traumatic. This is because digital recording systems see bulk storage simply as a peripheral that hangs off of an industry standard interface. This is true regardless of the type of storage. Therefore, the storage device itself can get bigger, faster or even change its nature and not impact the recording system. This very important fact is also true for replaying archived data. As long as the data is digital, and the storage device has a common digital interface supported by recorder/replayer, a variety of media types can be archived effectively. DSP*Con* has had a great deal of success in swapping digital tape recorders with RAID, JBODS and vice versa. Moreover, DSP*Con* fully intends to continue to support new devices as long as these devices are digital and obey open interface standards.

The next section of this paper discusses some popular I/O to storage interface standards.

Interface Types

Today, there are two different (albeit similar) methods of moving digital data to and from digital bulk storage devices. The fact that these are standards is very important because, in theory, this makes replacing one unit from one vendor with another unit from perhaps a different vendor easy. Moreover, in theory, upgrading a unit to the next generation should be relatively inexpensive. Each of these methods is described below.

- ***SCSI***—A SCSI is basically a parallel bus with address lines that can attach to up to fifteen devices. A processor (called a host) may initiate a transfer of a block of data to or from any device on the bus. A SCSI Ultra Wide bus is 64 bits and has a maximum sustained transfer rate of 320 Mbytes/second. SCSI buses come in two electrical types—*Single-ended* and *Differential*. The maximum length of a differential SCSI bus is approximately 20 meters–-far less for single-ended. Please note that for each type of device that can be on a SCSI bus (fixed disk, removable media, etc.), there exists a very well defined and rigid protocol for communication with these devices.

- ***Fibre Channel***—Unlike SCSI buses, Fibre Channel is a serial interface that can be run up to 10 kilometers. Each "loop" will support up to 2 Gbits/second. At the time of this writing, several vendors are already starting to ship 4 Gbit/second devices with talk in the near term of increasing this to 10 Gbit/Second. There is no practical limit to how many devices can be present on a Fibre Channel loop. Most importantly, Fibre Channel supports two well-known protocols; SCSI and TCP/IP. This turns out to be an incredibly important fact. It means that Fibre Channel can be used for networking computers together and/or for interfacing computers to bulk storage. There are two major reasons why this is significant. The first is that since Fibre Channel is a networking standard, this technology is well supported by hundreds of vendors, and that peripherals, including switches and I/O interfaces, are inexpensive. In addition, there is a great deal of commercial industry focus for enhancing this technology. The second reason is that people are now building standard networks that consist of processors and bulk storage devices. As long as the bulk storage devices have a Fibre Channel interface, they appear to the network as simply another client device. This is important because it means that the wideband recorder may readily be connected to many back office networks for data extraction.

Current Alternative for High-Speed Digital Recording

Industry can now, cost effectively provide high-speed multi-channel data acquisition systems that have an aggregate throughput, which is almost unlimited. In order to have an integrated system in which a single computer can obtain high speed access to the recorded data it would be necessary to use either Fibre Channel or SCSI with Gigabit Ethernet. Gigabit Ethernet in concert with SCSI is a more complicated and a lower performance solution, but it is also cheaper. For this reason, Fibre Channel will be the only transfer medium considered throughout the rest of this paper.

Fibre Channel has a limit on a single link of close to 200 Mbytes per second. The practical maximum number of links per channel is 2. Being conservative, this would limit the single channel/data stream rate to approximately 320Mbytes per second, or 160 Million, 16 bit samples per second. This transfer rate can only be accomplished with the proper hardware. Issues of RAM transfer rate, and having separate isolated transfer buses have to be taken into consideration. Using a Motorola 5500 Single processor Power PC board would allow for a transfer rate of over 800 Mbytes per second throughput, which would allow the board to transfer 400Mbytes per second of input data and 400Mbytes per second of output data. A second consideration also has to be made of the disk configuration. Each disk drive used in a recording system can transfer approximately 50 megabytes per second. This would indicate a need for up to 8 drives on 2 separate Fibre Channel loops in order to sustain 320 Mbytes per second on a JBOD. It should also be noted, that if RAIDS are required for either redundancy or recording size, it would take 4 controllers to sustain this transfer rate. For sizing a RAID's performance is approximately 100Mbytes per second,

Each of the configurations presented below can be expanded to as many channels as is required. IRIG/GPS time stamping is also a possibility for all systems.

System 1 - High Speed, Continuous IF Recorder

A system can be configured to continuously record IF range signals. In this recording example the maximum sample rate is 105 MHz per channel, 14-bit samples with a SINAD response of better than 72dB. The dual A/D converter has a bandwidth of greater than 250Mhz, meaning that if the analog bandwidth of the input signal is limited to 52.5 MHz, sub-sampling techniques can be used to retrieve signals at a center frequency approaching 250MHz. A block diagram of such a system is shown below in Figure 2.

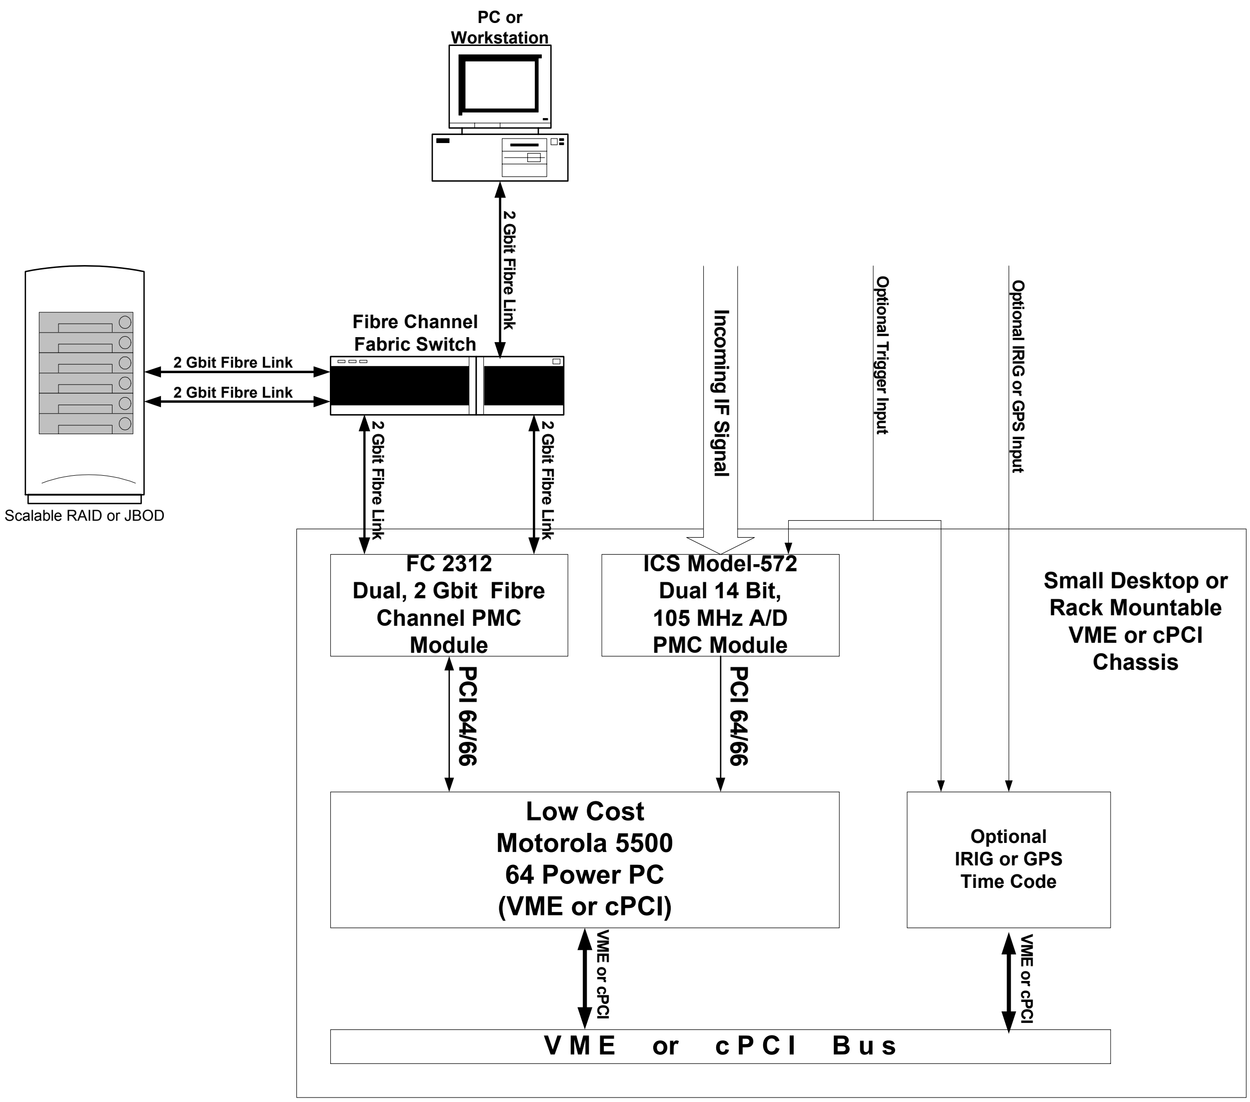

Figure 2–Continuous IF Recorder Block Diagram

System 2 - High Speed, Continuous IF Recorder

A system can be configured to continuously record IF range signals. In this recording example the maximum sample rate is 210 MHz per channel, 12-bit samples with a SFDR of better than 75dB. The dual A/D converter has a bandwidth of greater than 700Mhz, meaning that if the analog bandwidth of the input signal is limited to 105 MHz, sub-sampling techniques can be used to retrieve signals at a center frequency approaching 700MHz. A block diagram of such a system is shown below in Figure 3.

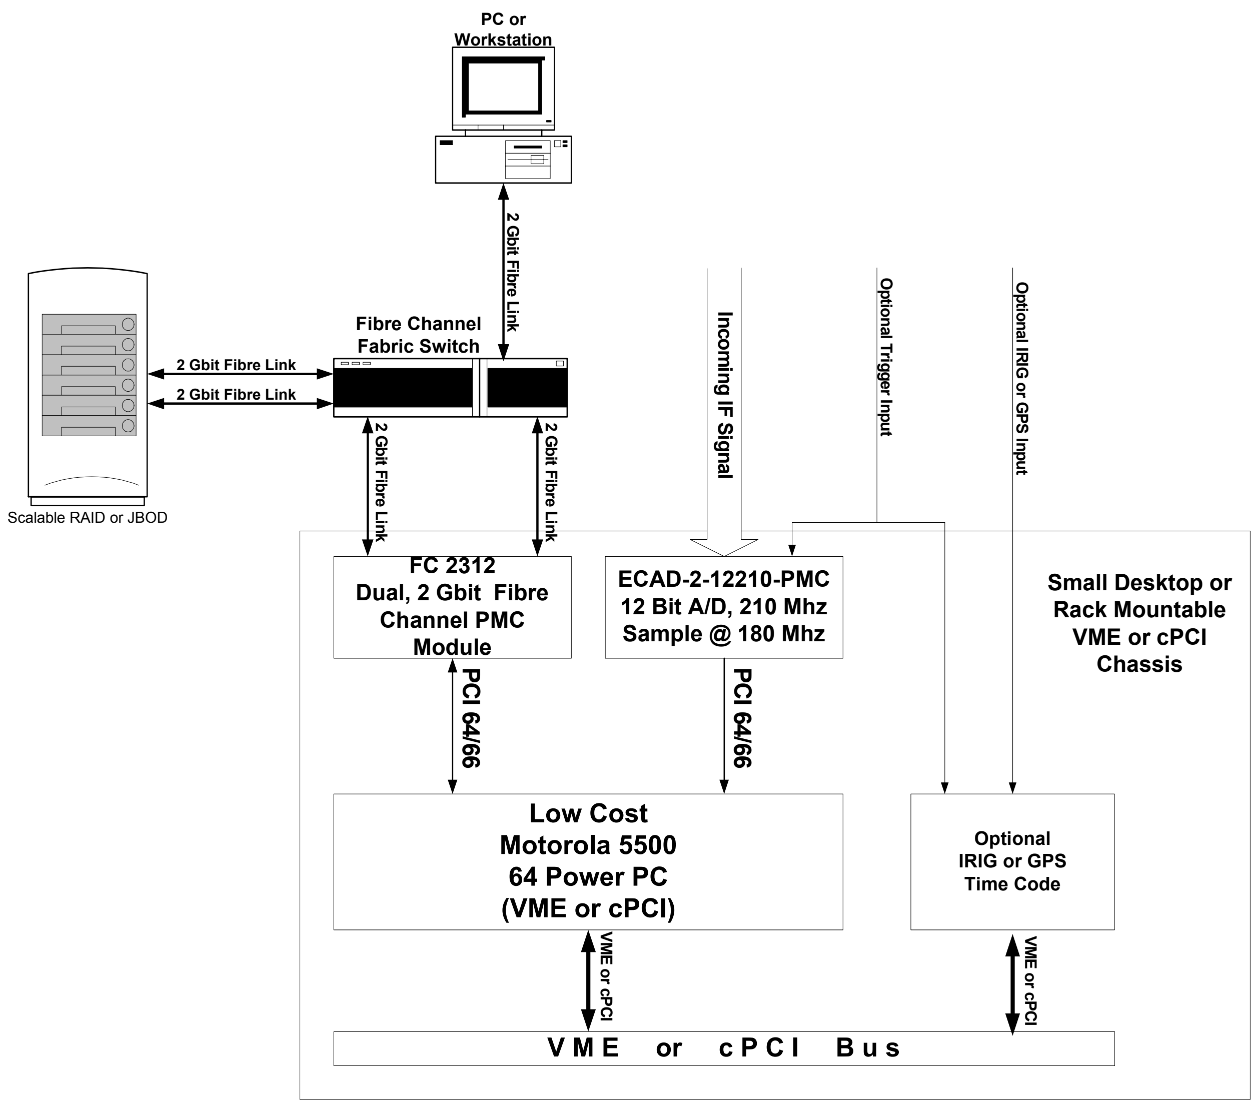

Figure 3–Continuous IF Recorder Block Diagram

System 3 - High Speed Triggered Recording

In many case such as pulsed Radar systems, there is no requirement to continuously record data, because there may be only a small window in which the data is interesting. This next system can sample based on a trigger to sample data in a gated range. The maximum sample rate is 2GHz per channel, 10 bit samples. The A/D has a bandwidth of 3GHz, so signals that are band limited to less than 1GHz could be sub-sampled. The length of each capture can be up to 96 million samples. Delayed triggers are also possible to allow for the system to capture traces in any time window. The trigger rate must be limited so that the transfer rate for a single channel is less than 320Mbytes per second. The transfer rate is calculated from the trigger rate times the capture length in samples. It should be noted that most RADAR systems have an aggregate transfer rate, which is nowhere near the maximum transfer rate of 320 Mbytes per second. It should be noted that a single I/O module can have up to 2 channels, but when configured with 2 channels each channel has a maximum sample rate of 1 GHz. A block diagram of such a system is shown below in Figure 4.

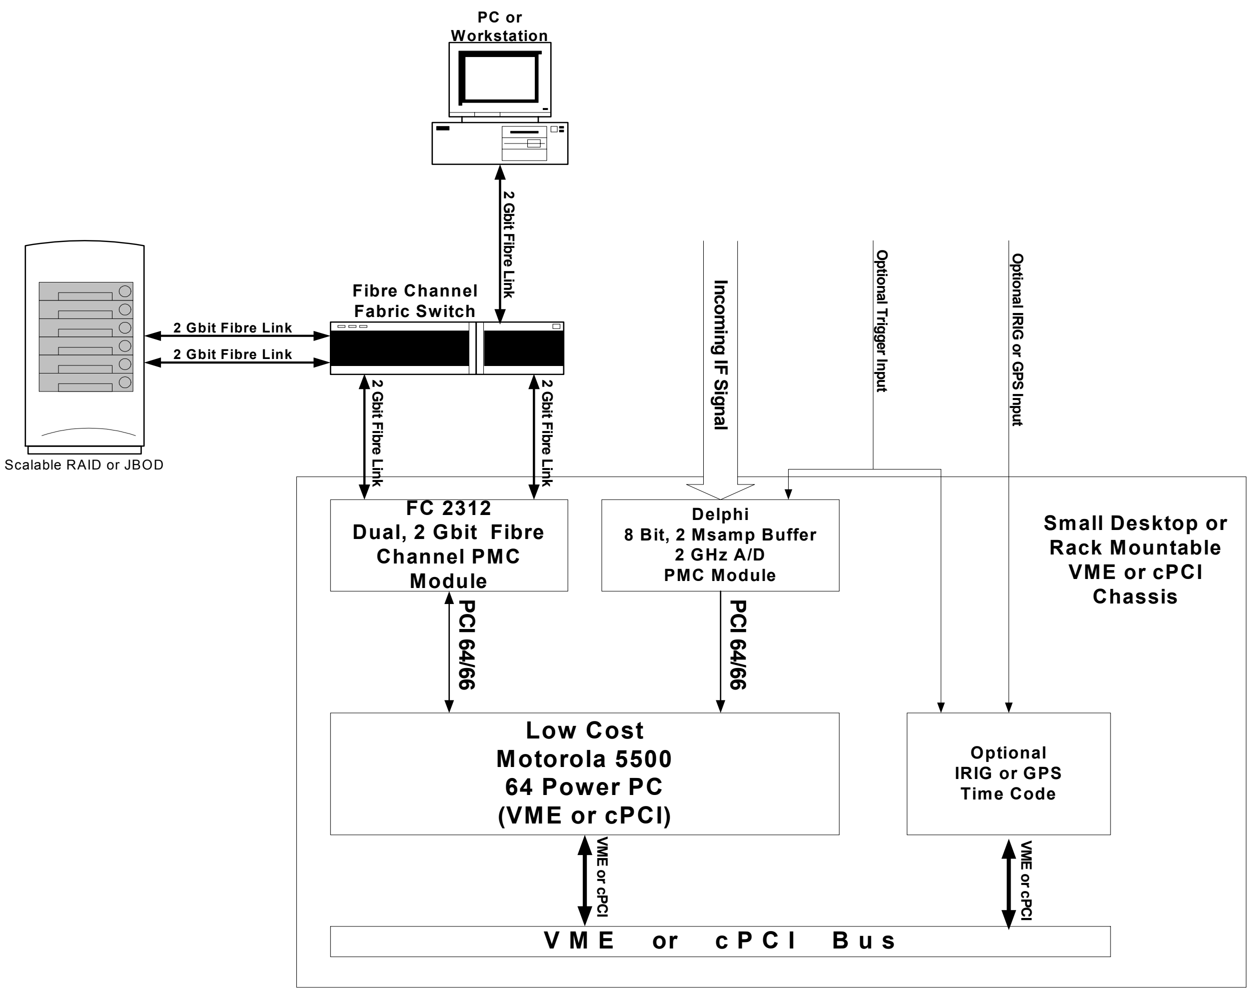

Figure 4–High Speed Triggered Recorder Block Diagram

System 4 - High Speed IF Recorder

In many cases a user is only interested in a slice of data in an IF bandwidth. This happens many times when a signal only has a small bandwidth, and the user does not want to do direct sampling of the IF signal just to have to filter and decimate once the data has been recorded. This system is similar to System 1, but instead of recording direct samples from an A/D, the samples from the A/D are sent through a Digital Drop Receiver. The Digital Drop Receiver can sample data at 100 MHz with 14 bit precision, and then Bandshift, Filter, and Decimate the data. The system here has up to 2 input channels per board. The Bandshift capability can be precision controlled so that the result can be centered directly above the center of an IF signal. This configuration is many times preferred if the IF signal is band limited. Decimation factors of 2, 4, 8, 16, and 32 are available for each channel. Advantages of the Digital Drop receiver are that the system can record Complex data instead of just real data, which in many cases is easier to process, and the DDR actually increases the SNR of the signal, through processing gain. It should also be noted that when using a decimation by a factor of 2, the data rate is the same as the raw data stream, but now the data can be in complex format. A block diagram of such a system is shown below in Figure 5.

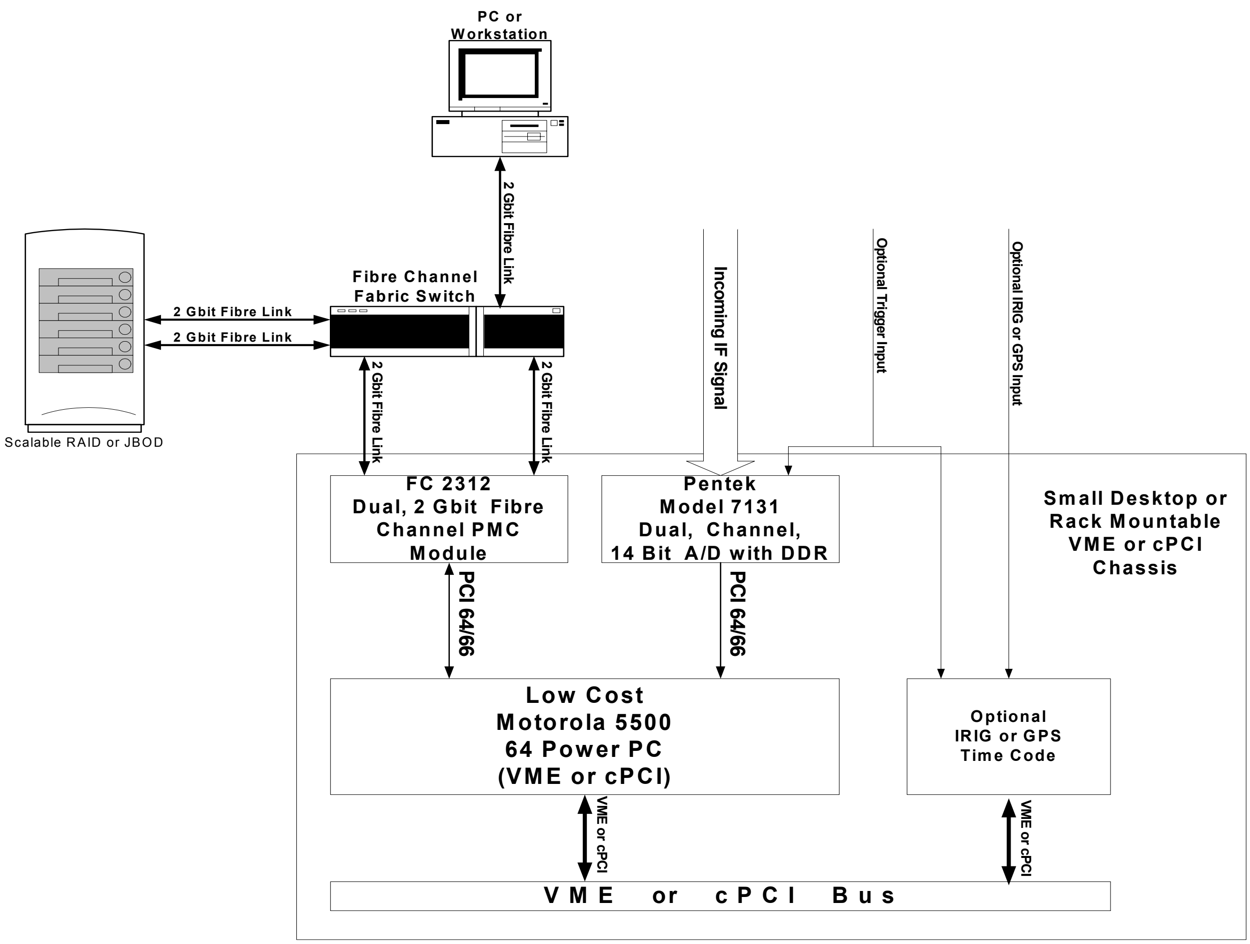

Figure 5–High Speed IF Recorder Block Diagram

System 5 - Digital Input Recorder

Many times users have digital events, which need to be sampled for long periods of time as in a logic analyzer. This system can record up to a 32 bit digital input stream at 40 Million samples per second. FPDP will allow for a clock and data digital input stream. All channels need to be sampled on the same clock edge. This system only creates a maximum data rate of 160Mbytes per second, which indicates a need for a single Fibre Channel Connection, and 4 disk drives, or 2 RAIDs. A block diagram of such a system is shown below in Figure 6.

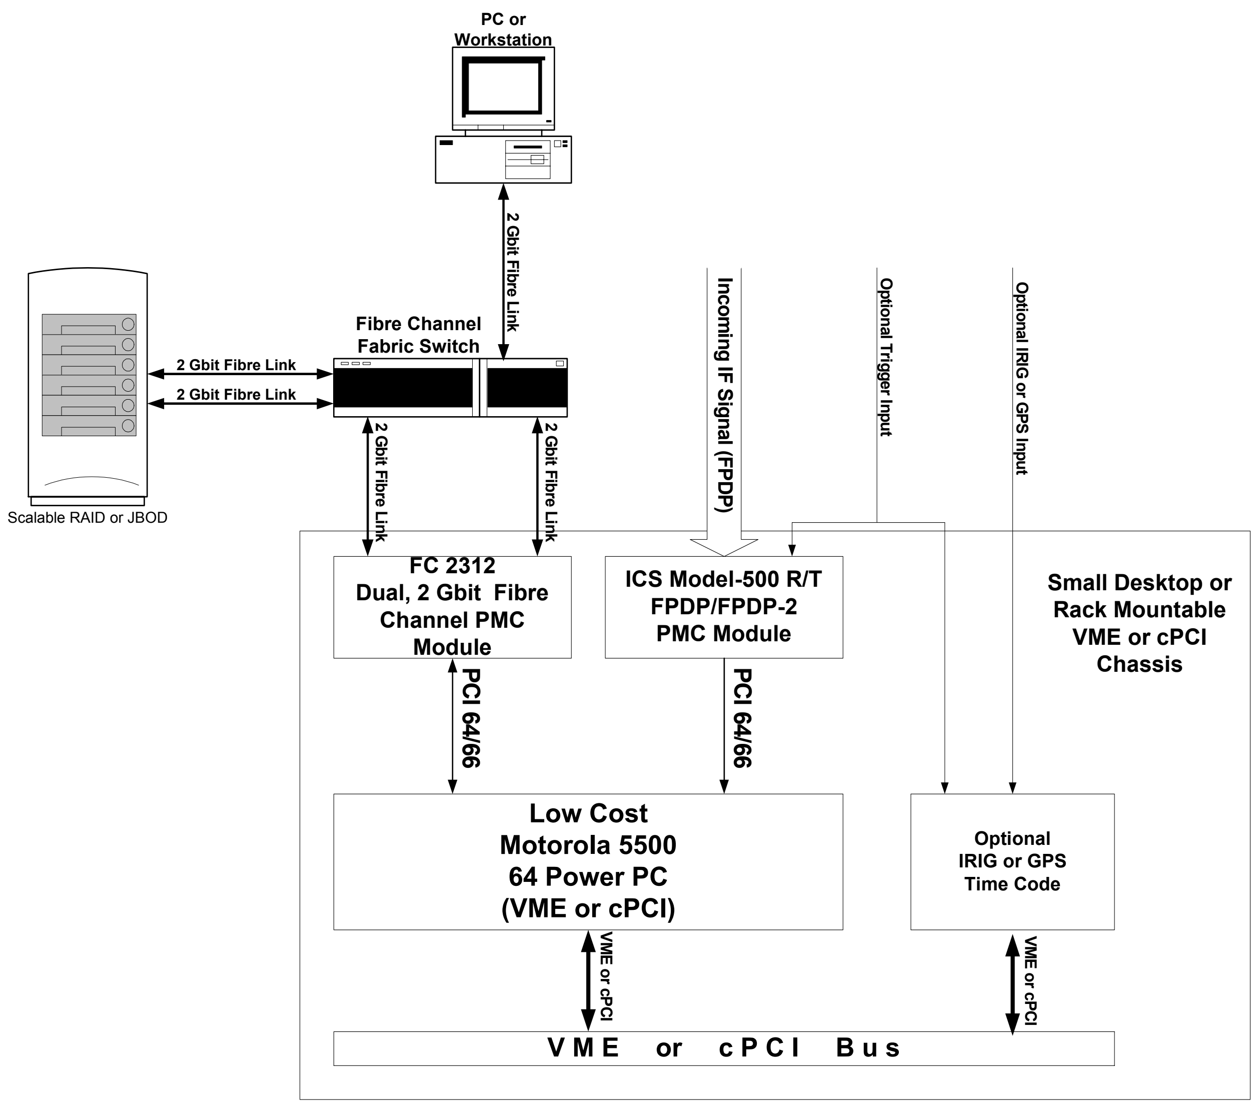

Figure 6–Digital Input Recorder Block Diagram

System 6 - Serial FPDP Recorder

Serial FPDP is now becoming a widely used protocol for transporting digital data from analog front-end systems to digital back-end systems. This system has a serial FPDP Interface, which is capable of up to 200Mbytes per second. The standard, Serial FPDP, is a growing communication standard, which has a qualified communication protocol. Many serial standards in the past failed to be accepted because although the hardware level was defined, the protocol level of the interface was not. There is a growing segment of companies using this protocol. Companies have created hardware, which can convert directly from FPDP to this serial standard, and vice versa. A block diagram of such a system is shown below in Figure 7.

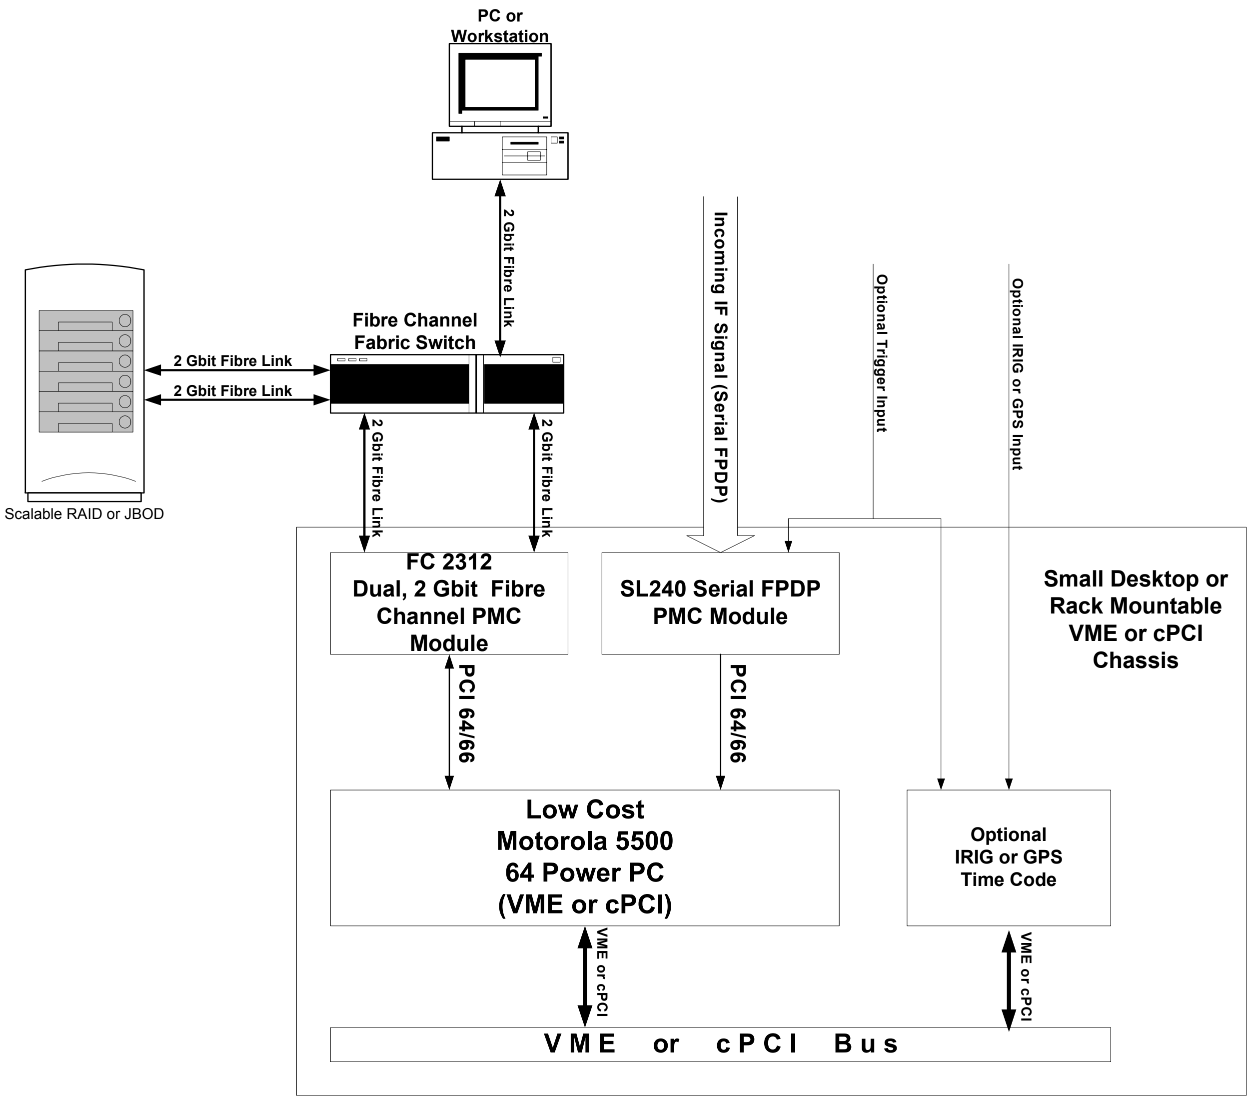

Figure 7–Serial FPDP Recorder Block Diagram

Summary

The following table provides a basic summary of some of the major characteristics of the systems and technologies discussed in this paper.

Config	Channels	Bits	Bandwidth	Storage	Max Acquisition
1	2/slot	14	Digital 52.5MHz Analog 250MHz	Expandable Fibre Channel	Continuous
2	2/slot	12	Digital 105MHz Analog 700MHz	Expandable Fibre Channel	Continuous
3	1/slot	10	Digital 1.0GHz Analog 3.0GHz	Expandable Fibre Channel	Up to 96 million samples per pulse
4	2/slot	14	Digital 42MHz Analog 250MHz	Expandable Fibre Channel	Continuous
5	32/slot	1	Each bit can be sampled at up to 40MHz	Expandable Fibre Channel	Continuous
6	1/slot	1	2 Gbits	Expandable Fibre Channel	Continuous