TRANSPORTING MAINFRAME DATA OVER DISTANCE

The purpose of this paper is to discuss various factors relative to transporting mainframe data over distance and to alert the reader that CTP has services available to help customers evaluate their networks when dealing with ESCON traffic and projected performance analysis.

In a local computing environment the efficient flow of data within the mainframe and into and out of the mainframe is generally taken for granted. Only when someone complains do the red flags go up. The process of preventing or correcting such complaints has many names. Capacity planning, configuration management, performance evaluations are a few that come to mind. The names are not nearly as important as the timely resolution of complaints. Unfortunately, determining the cause or causes of complaints, performance problems, can be a complex undertaking. Is it the application? Is it the operating system? Is it the mainframe? Is it the DWDM transport equipment? Is it the channel or channels? Are there too few peripheral devices? There are many questions to be addressed. In truth, the inefficient flow of data can characterize the root cause of performance problems.

When peripheral devices (printers, terminals, tape drives, disks, another mainframe etc) are placed at distance from the local mainframe an additional set of considerations are added to the complexity of a computing environment. This paper will ultimately concentrate on this area. Before we get too specific let us spend some time discussing the various components involved.

Let us start with an application. Typically an application can be viewed as a program that accomplishes a business need. The term application is used broadly and includes business areas like, payroll, customer inquiry, engineering design, print services, web browser, backup and recovery, data mirroring, etc, etc. Application design is an important contributor to overall computing system performance. Application dependent is a term that refers to the application design and how sensitive the application may be to external influences such as distance. The host operating system can be considered an application. The host computer can be described in many ways. Large MVS mainframe, open system mainframe, NT machine, AS400, UNIX processor, SP2 and distributed system represent a sample of terms used to describe the various types of host computers. Additionally, certain peripheral control units can be viewed as host computers in some data mirroring applications. Data mirroring refers to a group of applications that continually update data at a remote location to keep that data in sync with a corresponding set of local production data. All host computers share a common thread. They are all collectively a tool, a tool that relies on the efficient flow of data both inbound and outbound to do work. A channel is an entity that connects a peripheral device to a host computer. It is that part of the host computer system that enables and manages the efficient flow of data both inbound and outbound. The number of channels associated with a given host computer can vary depending on the host computer. Like host computers there are many types of channels. Some of the more common type of channels are referred to as ESCON, Bus Tag, SCSI and Fiber Channel. Each channel with it's associated protocol is unique. Not only do channels provide for the flow of data into and out of the host computer, they are responsible for data integrity. CRC, LRC and Parity are some of the terms used when discussing channel data integrity. Channels are also encumbered with distance limitations. The distance limitation defines how far from the host computer a device or devices connected via the channel can be located. The distance limitation for a given channel type is dependent on many factors. Host computers and their associated channels are designed and optimized for local attachment of peripheral devices. Peripheral devices connect to channels. Peripheral devices are products that directly support applications. Printers, tape drives, disks, terminals, teleprocessing control units, disk imagers, represent some but certainly not all of the peripheral devices available.

The remote placement of peripheral devices is not only desirable but in some instances a requirement in today's business environments. The fundamental technology to support remote peripheral installation is certainly available. Let us define remote peripheral installation. We will use the term Extended Distance. Extended distance is any distance greater than the local attachment specification for the host computer, type of channel and peripheral device to be installed at an extended distance. Extended distance can range to thousands of miles. The term channel extension is typically used to describe an existing part of the information systems industry. In this discussion we will use the term channel extension in a generic sense to describe the capability to install a peripheral device at an extended distance from the host computer. The channel extension device connects to the host channel on one side and the network on the other side. The term network will also be used in a generic sense. For our discussion network refers to the path from the peripheral device to the host channel. To enable this extended distance the channel extension device may convert the channel protocol to some communication protocol and assume responsibility for data integrity and efficiency (performance) over the network link, or the channel extension device may not do any communication protocol conversion as exemplified by many types of dense wave division multiplexer (DWDM) equipment often labeled as protocol independent. Extended distance also introduces the term Latency in a new context. In a local environment latency may be measured in nano-seconds or microseconds. In an extended distance environment, latency is typically measured in milliseconds. One millisecond for every one hundred twenty fives of electrical distance is a generally used metric for latency over extended distance. Latency can be a significant factor in the performance of extended distance peripherals. Fundamentally there are two types of generic channel extension, transparent mode and emulation mode. DWDM is a form of transparent channel extension. In DWDM implementations the channel data is sent across the network without any manipulation, in this mode all error conditions are managed by the host computer. Also, in transparent mode the local channel data, (command, status and data), may be converted to a higher layer communication protocol and transmitted across the network. Obviously the process is reversed at the remote end of the network. In this case the network might be frame relay, ATM, IP etc. Transparent mode performance is extremely sensitive to latency. Emulation mode is quite different from transparent mode. Emulation mode mitigates the negative effect of latency relative to performance and throughput. In emulation mode the local channel extension device becomes (emulates) the device we are installing at an extended distance. Accordingly, the local channel extension device interacts with the host computer in exactly the same way as the remote device. The important point here is the channel protocol, chatter, that would normally traverse the link is significantly reduced. At the remote end of the network the process is once again reversed. The remote channel extension device assumes the identity of the host computer and attempts to drive the peripheral device at rated speed. There is an important point to understand in Emulation Mode. The channel extension subsystem designer must know as much, if not more, about the host computer, channel, and device as the original equipment (OEM) designer. In either mode the channel extension subsystem designer must understand the application and which mode is appropriate.

So, what is the big deal about distance or latency? The big deal is time. Perhaps an example is the best way to discuss this. We will start with the application. The application is Synchronous Data Mirroring. Remember Synchronous Data Mirroring is implemented with interlock between the local and remote peripheral devices. Synchronous Data Mirroring can adversely affect I/O response time, especially over distance. The channel protocol is ESCON. The application was first installed in a local environment for proof of concept. The application performed well. The average update I/O response time was 3.7 milliseconds. Then the mirror peripherals were installed at a bunker site approximately eighty geographic miles away. The average update I/O response time increased to more than 30 milliseconds. The application had to be stopped because the performance was terrible. Why did the update I/O response time increase to an unacceptable level? Latency is the answer. Latency has two sources: equipment and distance. The distance latency in this case was eighty miles which equals 640 microseconds (refer to the earlier metric of one millisecond for each one hundred twenty five miles). This assumes a straight line connection between the local and remote location. The one way distance latency of 640 microseconds must be multiplied by 2 which yield a roundtrip distance latency of 1.28 milliseconds. The equipment latency for the roundtrip is an additional 2 milliseconds. The total roundtrip latency is 3.28 milliseconds. Unfortunately this example gets worse. Since the metric here is update I/O response time, we must take a closer look at the anatomy of an ESCON update I/O. Based on the amount of data being sent from the local to the remote site for each update I/O there are ten ESCON turn around protocol sequences required. Each of the ten ESCON sequences must traverse the entire link. 10 ESCON sequences times 3.28 milliseconds each yields a total latency of 32.8 milliseconds for additional update I/O response time. This is an increase of over 800% in update I/O response time and is directly related to the ESCON protocol sequencing and distance.. This additional transfer time caused the information to be stored locally until the network could pass it along to the far end. You could look as this as leaky bucket. The leak being the data sent to the remote site and the bucket is being filled by real time transactions. The bucket will eventually fill and bring the system to a halt. Wouldn't it have been better to have known this before starting the data mirroring application?

At this point you may start to see the complexity of extended distance placement of peripheral devices and the potential for performance problems. Where do we go from here? How do we evaluate the efficiency of transporting ESCON mainframe data over distance? Can you picture a potential installation where a peripheral device is connected to a host over a metropolitan area network with a local connection at each end of the network covering a distance of approximately seventy miles? The best answer may be, it depends. It certainly depends on the reader’s job and area of responsibility. As stated earlier there are many components in a geographically dispersed computing complex. There are also many individuals associated with this type of installation. Data center manager, system programmer, disaster recovery planner, computer repairman, communication specialist, network engineer, network manager, IS salesperson, storage specialist are some, but certainly not all inclusive, that come to mind. Simply stated, we start at the beginning with the application. Once we understand the application we can start to drill down, utilizing various diagnostic tools and test equipment as appropriate, until enough data has been gathered to support a logical conclusion and potential corrective action. Can it be done and who does this? In an ideal world this process would be a collaborative effort. The world is not always ideal. Sometimes a person, although expert in one ~~a~~ specific area, is charged with an overall system evaluation. Is there a place that this person can turn to?

CTP can help, because in addition to providing specific test equipment, CTP provides consulting services ranging from specific education on CTP test equipment to application education and insight, CTP can provide the type of information required for success in this area. Depending on your job area and responsibility you may benefit from a discussion with CTP. Contact us at , www.c2p.com or mktg@c2p.com .

About the Author:

Paul Carrick with 38 years in IS industry, 33 with IBM and 5 with Inrange Technologies Corp., has conducted assignments in the IT industry for project planning, justification and implementation. Paul has developed and implemented many large information system solutions with a variety of hardware and software platforms. He has trained and developed individuals to understand the intricate details associated with Disaster Recovery products and their solutions involving Disks and Tape mediums in wide area networks.

.