# To CUP, or Not To CUP? That is the FICON Question!

Dr. Stephen R. Guendert
Brocade Communications

*CUP, or Control Unit Port is a holdover from ESCON directors. In a FICON environment, CUP allows for in-band management, and opens the door to FICON director performance metrics via the RMF 74-7 record, more commonly known as the FICON Director Activity Report. In an effort to reduce acquisition costs and be more competitive on price, many vendors will try and make the case why you do not need CUP on FICON directors. This paper will present the reasons why, from a performance management perspective, "that not to CUP" is the wrong answer to the question posed by the title.*

## Introduction-CUP concepts

CUP on FICON is a legacy of CUP on ESCON. The 9032-5 ESCON directors had an in-band management capability that utilized an embedded port in the control processing cards on the ESCON director to provide a communications path to an MVS console. The CUP functionality was used for reporting hardware errors up to MVS (Helpdesk), blocking and unblocking ports via the PDCM (Prohibit Dynamic Connectivity Mask) and for monitoring performance. When switched FICON was being engineered, IBM wanted to make certain that its mainframe customers would have a consistent look and feel between ESCON and FICON, including the CUP functionality.

In a FICON environment, CUP is primarily used for in-band management of the FICON directors. Host based management programs manage the FICON directors/switches using CUP protocol by sending commands to the emulated control device implemented in HCD. A FICON director or switch that supports CUP can be controlled by one or more host based management programs or director consoles. Control of the FICON directors can be shared between these options. There are 39 CUP commands, or channel command words (CCWs), for monitor and control of the FICON director functions. The CUP function and command set are defined in an IBM proprietary document "FICON Director Programming Interface", also commonly referred to as the CUP specification. CUP commands are oriented towards management of a single switch, even though the use of CUP in a cascaded FICON environment is fully supported.

**CUP under ESCON**

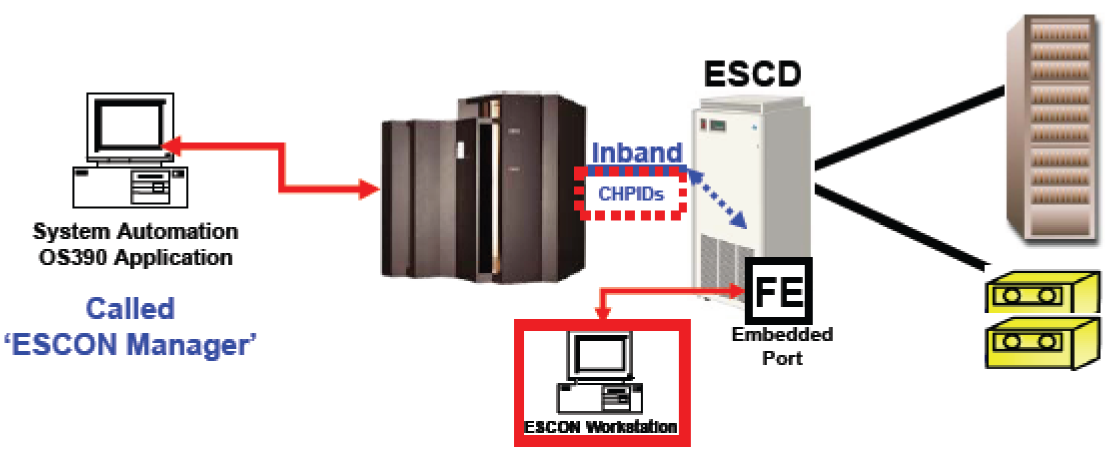


**Figure 1-ESCON CUP**

Figure 1 above is a simplified representation of an ESCON environment and how CUP fits into that environment. You can see the management stations, and the embedded "FE" port on the ESCON director. This embedded port is the CUP. Best practices were to have two or more CHPIDS with access to the CUP, and to also implement safe switching. Safe switching describes the System Automation for OS/390 process of path removal. Simply stated, safe switching is a process that involves contacting all systems in a sysplex and verifying a path can be removed before it is actually removed. Each system involved must logically remove the path and response positively to the requesting system before the path can be taken out of service. Systems will usually only respond negatively if the path in question is the last path for an application.

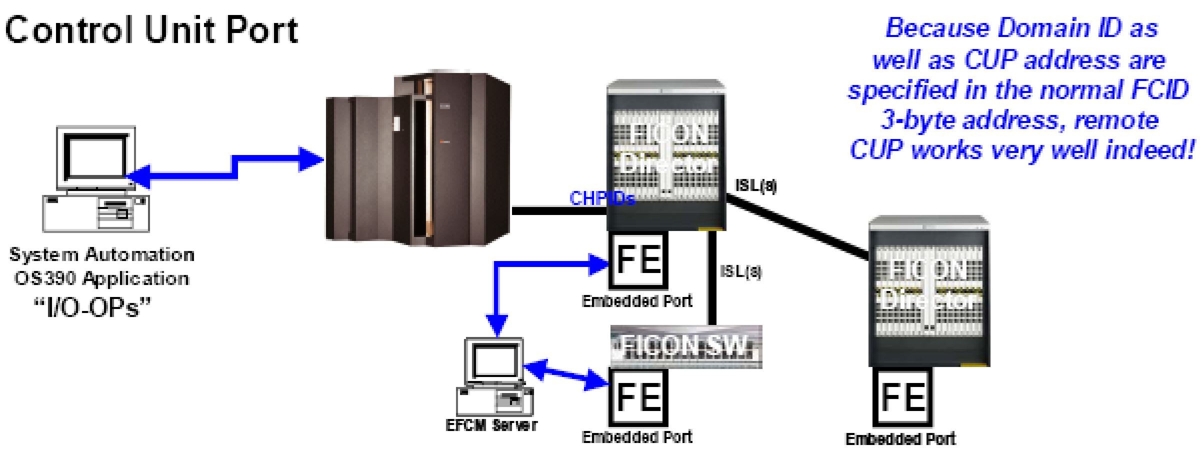


**Figure 2-FICON CUP**

Figure 2 is a simplified representation of a FICON environment and how CUP fits into that environment. Similar to ESCON, FICON directors have the embedded port "FE" for the CUP port. On FICON directors that have 256 or more ports, this has caused a slight dilemma. This logical "FE" overlaps the physical "FE" port, so the physical "FE" and "FF" ports cannot be used on these directors for FICON connectivity. They may still be used for port swaps or for FICON/FCP intermix, however, they cannot be "genned" in HCD for FICON connectivity. In a FICON environment, only one RMF LPAR should attempt to access the CUP port at any one time, and it is still best practice to have two or more CHPIDs with access to the CUP. However, too much activity to the FICON CUP can cause missing interrupts, leading to the potential for having a boxed device.

**Defining CUP in HCD**

Defining CUP in HCD/IOCP is optional; however, if CUP is not defined, the following functions will not work:

1) System Automation (SA) for z/OS I/O Ops which is used to manage and display the status of the FICON director.
2) RMF reports, such as the RMF 74-7 (FICON director activity report)

The FICON director will be defined in HCD as a device. This has to be done for error reporting purposes. Switch related hardware errors are reported to z/OS against a device number. If the switch is not defined as an I/O device, and that I/O device is not online to z/OS, then switch related errors cannot be surfaced and actioned. System automation for z/OS access provides operational tools for "safe switching" as described earlier, as well as displaying device routing information. For this to function, all switches must be online as I/O devices on all the systems where S/A I/O ops manager is running.

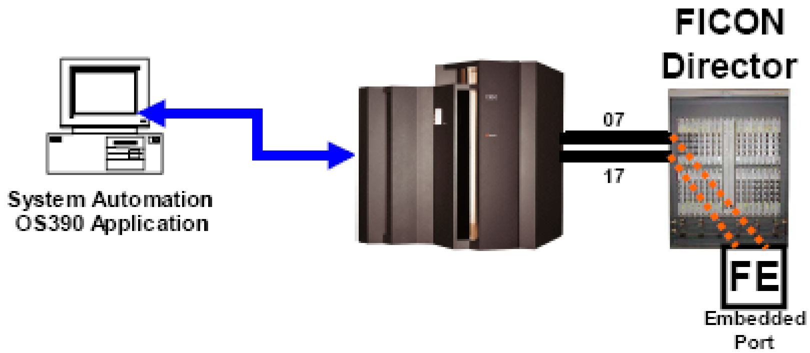


```
CHPID PATH=(07,17),SHARED,TYPE=FC,SWITCH=7A
CNTLUNIT CUNUMBR=xxxx,PATH=(07,17),UNIT=2032,
            UNITADD=((00,1)),LINK=(7AFE,7AFE)
IODEVICE ADDRESS=(xxxx,1),UNITADD=00,CUNUMBER=(xxxx),UNIT=2032
```

**Figure 3-Defining CUP**

Figure 3 is an example of how CUP would be defined for our earlier sample FICON environment. The assumption made is the FICON director has a Domain ID (DID) of 7A. The CUP address is always “FE” and a specific CUP address notation in a fabric is dd-FE-xx, where dd is the domain ID of the director/switch, FE is the embedded CUP port address, and xx is irrelevant to the discussion. With multiple directors/switches in a fabric with CUP, define them all in the same way, but the LINK statement will have the different domain IDs of the different switches, i.e. LINK=(61FE) or LINK=(6EFE), etc. From one, and only one LPAR path, RMF communicates with each FICON director/switch in turn, requesting that its “FE” port dump all FICON port statistics down to the mainframe. The “master” RMF then distributes that information to all the other LPARs on the mainframe and a FICON switching device report is generated for each FICON director/switch with CUP.

**Using CUP**

CUP as an in-band management tool provides a great deal of useful information to the end-user. Chief amongst this useful data are the following:

- Port statistics-the number of words/frames transmitted and received
- Switch node identifier-serial number, manufacturer, etc
- Configuration file information-the list of configuration files residing on the director including actual file content with port address names and port connectivity
- History summary (director history buffer)-each change in status or configuration of the ports is logged in a history buffer which is retrievable via CUP.
- Switch configuration data-timeout values, ports, etc.

CUP also provides the prohibit dynamic connectivity mask (PDCM) as a management tool. The PDCM is a mechanism to define port connectivity such as prohibits, allows, and blocking/unblocking of ports. Use of the PDCM can be a complex undertaking. However, it provides a highly useful mechanism for segregating traffic, in effect, providing a very rudimentary quality of service (QoS) functionality.

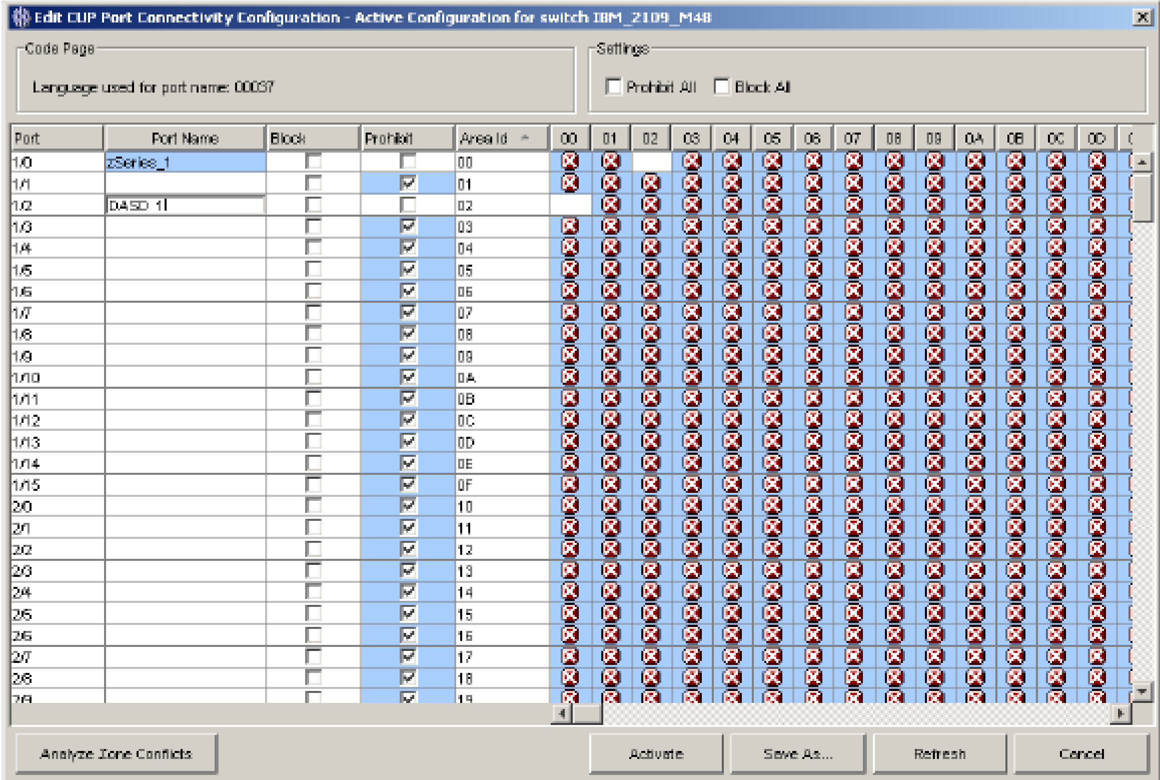


**Figure 4-PDCM**

CUP is also essential for in getting Service Information Messages (SIMs) to the MVS consoles for FICON device failures. The introduction of the FICON I/O protocol to the mainframe I/O subsystem ushered in a new era in our ability to process data rapidly and efficiently. The FICON protocol is significantly different than its predecessor ESCON protocol. As a result of two main changes that FICON made to the mainframe channel I/O infrastructure, the requirements for a new RMF record came into being. The first significant change was that unlike FICON, ESCON did not use buffer credits to account for packet delivery. The second significant change was the introduction of "FICON cascading" which was not possible with ESCON. If the end user wants to do RMF Monitor I reports for reporting on FICON director activity statistics via the RMF 74-7 record, they need to use CUP. This is typically the most important reason the end user needs to have CUP on each FICON director in their configuration. The RMF 74-7 record provides the end user with the only mechanism for detecting buffer to buffer credit shortages. ESCON did not use buffer to buffer credits, it used DIBs.

**ESCON DIBs-A brief historical perspective**

The device information block (DIB) contains control, status or user data. The maximum DIB size to be employed by a configuration is defined by the storage director's (control unit's) ESCON adapters. DIBs may be 8, 16, 32, 64,128, 256, 512 or 1024 bytes long. During initialization (defined as a subsystem or path being varied online or processor IML time) a subsystem's ESCON adapters notify the channel subsystem of the DIB and data buffer size that is to be employed. The channel subsystem then employs these values for all future communications with the subsystem.

An ESCON adapter is a specialized hardware interface which manages communication between the channel subsystem and the control unit. While the ESCON adapter is a part of the control unit, we need to consider its functions distinct from those performed by the channel and the control unit during the transmission of data. The ESCON socket connects to the 8b/10b converter, which in turn is connected to a microprocessor which manages the connection and schedules DIBs of data to and from the underlying storage director. This microprocessor also manages the data buffer which is incorporated in the ESCON adapter. The 3990-3/6 ESCON adapter's data buffer held eight 256 byte DIBs, i.e. 2048 bytes of storage. The DIB and data buffer sizes may be selected by the storage subsystem manufacturer within the constraints imposed by the ESCON architecture. Assuming that the ESCON adapter is receiving data, it must successfully pass each DIB it receives to buffer areas in the underlying storage director before it can request another DIB of data to be sent by the channel to avoid the potential of data overrun. While the channel may initially schedule a full data buffer of DIBs to be sent to the device, further DIBs may only be transmitted after a request has been received from the ESCON adapter signaling that a prior DIB has successfully been passed to the storage director.

**Packet Flow and Credits**

The fundamental objective of flow control is to prevent a transmitter from over-running a receiver by allowing the receiver to pace the transmitter, managing each I/O as a unique instance. At extended distances, pacing signal delays can result in degraded performance. Buffer-to-buffer credit flow control is employed to transmit frames from the transmitter to the receiver and pacing signals back from the receiver to the transmitter. The basic information carrier in the fibre channel protocol is the frame. Other than ordered sets, which are used for communication of low-level link conditions, all information is contained within the frames. To prevent a target device (either host or storage) from being sent more frames than it has buffer memory to store (overrun), the fibre channel architecture provides a flow control mechanism based on a system of credits. Each credit represents the ability of the receiver to accept a frame. Simply stated, a transmitter cannot send more frames to a receiver than the receiver can store in its buffer memory. Once the transmitter exhausts the frame count of the receiver, it must wait for the receiver to credit-back frames to the transmitter. A good analogy would be a pre-paid calling card: there are a certain amount of minutes to use, and one can talk until there is no more time (minutes) on the card.

**Buffer-to-Buffer Flow Control**

Buffer-to-buffer flow control is flow control between two optically adjacent ports in the I/O path (i.e. transmission control over individual network links). Each fibre channel port has dedicated sets of hardware buffers for send and receive operations. These buffers are more commonly known as buffer-to-buffer credits (BB_credits).

The number of available BB_credits defines the maximum amount of data that can be transmitted prior to an acknowledgment from the receiver. BB_credits are physical memory resources incorporated in the Application Specific Integrated Circuit (ASIC) that manages the port. It is important to note that these memory resources are limited. Moreover, the cost of the ASICs increases as a function of the size of the memory resource. One important aspect of fibre channel is that adjacent nodes do not have to have the same number of credits. Rather, adjacent ports communicate with each other during fabric login (FLOGI) and port login (PLOGI) to determine the number of credits available for the send and receive ports on each node.

A BB_credit can transport a 2,112 byte frame of data. The FICON FC-SB-2 and FC-SB-3 ULPs employ 64 bytes of this frame for addressing and control, leaving 2K available for z/OS data. In the event that a 2 Gbit transmitter is sending full 2,112 byte frames, one credit is required for every 1 KM of fibre between the sender and receiver. Unfortunately, z/OS disk workloads rarely produce full credits. For a 4K transfer, the average frame size is 819 bytes, therefore five credits would be required per KM of distance as a result of the decreased average frame size. It is important to note that increasing the fibre speed increases the number of credits required to support a given distance. In other words, every time the distance doubles, the number of required BB_credits doubles to avoid transmission delays for a specified distance.

**Exhaustion of BB_credits and Frame Pacing Delay**

Among the many functions of the CUP feature is an ability to provide host control functions such as blocking and unblocking ports, safe switching and in-band host communication functions such as port monitoring and error reporting. Enabling CUP on FICON directors while also enabling RMF 74 subtype 7 (RMF 74-7) records for your z/OS system yields a new RMF report called the “FICON Director Activity Report”. Data is collected for each RMF interval if FCD is specified in your *ERBRMFnn* parmlib member. RMF will format one of these reports per interval per each FICON director that has CUP enabled and the parmlib specified. This RMF report is often overlooked but contains very meaningful data concerning FICON I/O performance—in particular, frame pacing delay. **It is extremely important to note that indications of frame pacing delay are the only indication available to indicate a BB_credit starvation issue on a given port.**

Frame pacing delay has been around since fibre channel SAN was first implemented in the late 1990s by our open systems friends. But until the use of cascaded FICON increased, its relevance in the mainframe space has been completely overlooked. If frame pacing delay is occurring then the buffer credits have reached zero on a port for an interval of 2.5 microseconds and no more data can be transmitted until a credit has been added back to the buffer credit pool for that port. Frame pacing delay causes unpredictable performance delays, and those delays generally result in elongated FICON connect time and/or elongated PEND times that show up on the volumes attached to those links. **It is important to note that only when using switched FICON and only when CUP is enabled on the FICON switching device(s) can RMF provide the report that provides frame pacing delay information. Only the RFM 74-7 FICON Director Activity Report provides FICON frame pacing delay information. You cannot get this information from any other source today.**

```
                                          F I C O N   D I R E C T O R   A C T I V I T Y
                                                                                                                     PAGE    1
        z/OS V1R6                    SYSTEM ID SC64            DATE 10/06/2004            INTERVAL 10.00.001
                                     RPT VERSION V1R5 RMF      TIME 09.10.00              CYCLE 1.000 SECONDS
IODF = 58   CR-DATE: 09/23/2004   CR-TIME: 15.35.18      ACT: ACTIVATE
SWITCH DEVICE: 0061   SWITCH ID: 61     TYPE: 006064    MODEL: 001   MAN: MCD    PLANT: 01   SERIAL: 0000000119D3
PORT     -CONNECTION-   AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC)     ERROR
ADDR     UNIT     ID      PACING       READ   WRITE      -- READ --  -- WRITE --     COUNT
 04      SWITCH  ----          0        579     889          0.04          0.03          0
 05      CHP      5A           0         71     238          0.07          0.21          0
 06      CHP      80           0         68     175          0.07          0.16          0
 07      CU      ----          0          0       0          0.00          0.00          0
 08      CU      ----          0        886      73          0.03          0.00          0
 09      CHP      5C           0        171     129          0.17          0.15          0
 0A      CHP      81           0        165      85          0.13          0.08          0
 0B      ------  ----     P O R T    O F F L I N E
 0C      CU      ----          0        829      86          0.05          0.00          0
 0D      CHP      5E           0         73     888          0.00          0.03          0
 0E      CHP      82           0        112     720          0.00          0.02          0
 0F      ------  ----     P O R T    O F F L I N E
 10      CU      ----          0        826      89          0.05          0.00          0
 11      CHP      60           0          0       0          0.00          0.00          0
```

**Figure 5-Sample FICON Director Activity Report (RMF 74-7)**

The fourth column from the left in figure 5 is the column where frame pacing delay is reported. Any number other than zero in this column is an indication of frame pacing delay occurring. If there is a non-zero number it reflects the number of times that I/O was delayed for 2.5 microseconds or longer due to buffer credits falling to zero. Figure 5 shows what you would always hope to see, zeros down the entire column indicating that enough buffer credits are always available to transfer FICON frames.

```
                                     F I C O N   D I R E C T O R   A C T I
      z/OS V1R7               SYSTEM ID PDM1           DATE 11/28/2006
                              RPT VERSION V1R7 RMF     TIME 21.44.00

IODF = 70   CR-DATE: 09/20/2006   CR-TIME: 10.49.34     ACT: POR

SWITCH DEVICE: 006E   SWITCH ID: **    TYPE: 006140   MODEL: 001   MAN: MCD

PORT    -CONNECTION-   AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC
ADDR    UNIT     ID      PACING       READ   WRITE      -- READ --  -- WRITE -
 04     SWITCH  ----          3         71    1715           0.32        41.7
 05     CHP      5E           0          0       0           0.00         0.0
 06     CHP      C0           0        259     839           0.01         0.0
 07     CHP      C0           0        678     631           0.05         0.0
 08     SWITCH  ----          0         71    1689           0.38        39.0
```

**Figure 6-Frame Pacing Delay Indications in RMF 74-7 record**

But in figure 6 above, you can see that on the FICON Director Activity Report for switch ID 6E, there were at least three instances when port 4, a cascaded link, suffered frame pacing delays during this RMF reporting interval. This would have resulted in unpredictable performance across this cascaded link during this period of time.

This type of reporting on buffer to buffer credit starvation is not available on open systems. However, in a FICON/FCP Intermix environment, a director with CUP installed can also report on the open systems attached ports.

**How does frame pacing delay occur?**

There are two typical reasons that frame pacing delay occurs. The first is that a long distance link runs out of buffer credits, where possibly not enough buffer credits were specified to keep the link fully utilized. This can also happen even when maximum buffer credits are specified if the distance is too long for the speed of the link and the size of the frames being sent. In a local environment frame pacing delay can occur, for example, when a FICON CHPID has issued open exchanges to a number of control units that are all responding back to the host simultaneously as shown in figure 7.

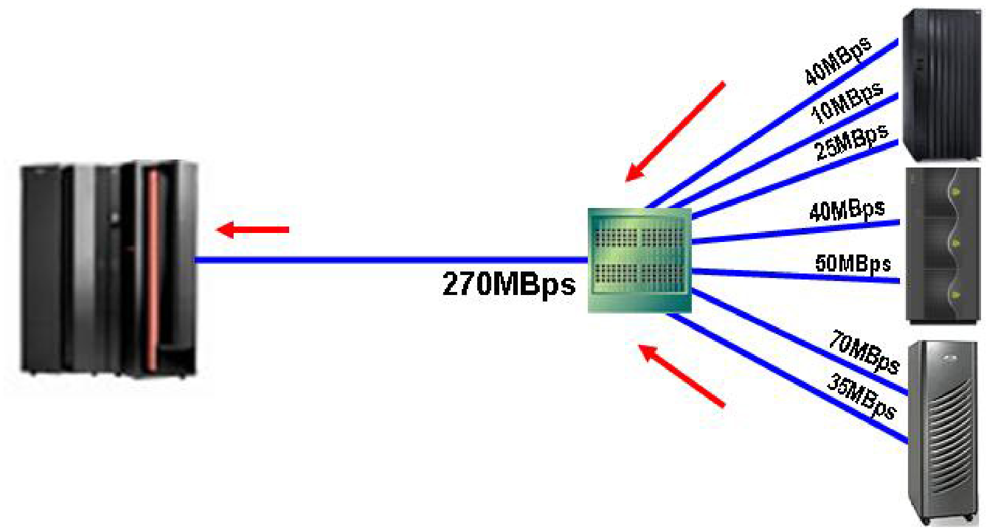


**Figure 7 – Frame Pacing Delay from storage to host**

In figure 7, the links between storage and the FICON Director are all running well within the limits of 2Gbps or 4Gbps links. But all of these links are attempting to converge into a single 2Gbps link that runs back to the host CHPID from the FICON Director.

From the FICON Director to storage there is no performance problem, but the 2Gbps link between the FICON Director and the mainframe has become overwhelmed. At its very best this link should operate at no more than about 150MBps and is simply unable to handle a 270MBps unidirectional frame flow on a 2Gbps link.

The port in the FICON Director for this CHPID connection quickly runs out of buffer credits. This signals all of the storage links converging to this FICON Director port to suspend I/O activity until buffer credits have been refreshed on that port.

**How can things be improved?**

It would appear that even with the new FICON directors and the ability to assign BB_credits to each port from a pool of available credits on each port card, that the end user is still stuck. The end user can best hope they allocate correctly, and then monitor their RMF 74-7 report for indications of frame pacing delay to indicate they are out of BB_credits. They can then go ahead and make the necessary adjustments to their BB_credit allocations for crucial ports such as the ISL ports on either end of a cascaded link. However, any adjustments made will merely be a better guestimate since the exact number being used is not indicated. Imagine driving a car without a fuel gauge, and having to rely on EPA miles per gallon estimates so you could calculate how many miles you could drive on a full tank of gas. Of course, this estimate would not reflect driving characteristics. And in the end, the only accurate indication you get that you are out of gas is a coughing engine that stops running.

Similar to open exchanges, the data needed to calculate BB_credit usage is currently available in RMF. All that would be needed is some mathematical calculations be performed. **The author respectfully submits that the RMF 74-7 record (FICON Director Activity Report) should be updated with these 2 additional fields and the appropriate interfaces be added between the FICON directors and CUP code. Director management software could also be enhanced to include these two valuable metrics.**

**Observations and Conclusion**

FICON CUP is a licensed feature on a FICON director. Since the CUP code is IBM proprietary, IBM charges the FICON director vendors a licensing fee for every director it is sold on. This cost is passed along to the end user and to the storage OEMs. In an effort to make their FICON director bids more price competitive, often times the OEMs will leave CUP off the FICON directors. Even worse, they will try and convince the end use they do not need CUP. In so doing, they perform a grave disservice to their customers. The RMF 74-7 record is currently the only way to have any understanding on buffer to buffer credit usage. This information is very important to have, and in a cascaded FICON environment it is absolutely critical. While the RMF 74-7 record is not the definitive source for solving performance problems in a FICON environment, its use can help expedite the troubleshooting process and it can definitely help narrow down the things impacting performance. **Therefore, the answer to the question posed by the title of the paper is To CUP!**

**Key References**


1) FICON Native and Implementation Reference Guide. IBM Corporation
2) FICON Management Server in-Brief. McDATA Corporation.
3) Understanding the Performance Implications of Buffer to Buffer Credit Starvation: Frame Pacing Delay. Steve Guendert. CMG 2007.