

**BROCADE**



**To CUP, or Not To CUP? That is the  
(FICON) Question!**

Brazil Computer Measurement Group  
August 2008 Conference

Dr. Steve Guendert  
Global Solutions Architect  
Brocade Communications

# Abstract

- 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 presentation 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.



# Legal Disclaimer

All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.

NOTHING IN THIS PRESENTATION SHALL BE DEEMED TO CREATE A WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD PARTY RIGHTS WITH RESPECT TO ANY PRODUCTS AND SERVICES REFERENCED HEREIN.

Brocade, Fabric OS, File Lifecycle Manager, MyView, and StorageX are registered trademarks and the Brocade B-wing symbol, DCX, and SAN Health are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.



# Agenda

- What is CUP?
- History of CUP
- ESCON use
- FICON use
- FICON Director Activity Report (RMF 74-7)




# Control Unit Port – CUP

- CUP on FICON is a legacy of CUP on ESCON
- The 9032-5 ESCON Directors have an in-band management capability that utilized an embedded port in the control processing cards to provide a communications path to an MVS console
  - It was used to report hardware errors up to MVS (Helpdesk)
  - It was used to block and unblock ports (PDCMs)
  - It was used to monitor performance
- When switched-FICON was being engineered, IBM wanted to be sure that its users would have a consistent look and feel between ESCON and FICON including CUP
- FICON directors have an internal node port (N\_Port) that supports the CUP function.



# CUP Under ESCON



- ESCON manager
- Embedded port FE is also called the 'Control Unit Port' or CUP
- MVS Consoles or ESCON workstation can control the ESCON Director
- Best practice was to have 2 or more CHPIDs with access to the CUP
- Safe switching

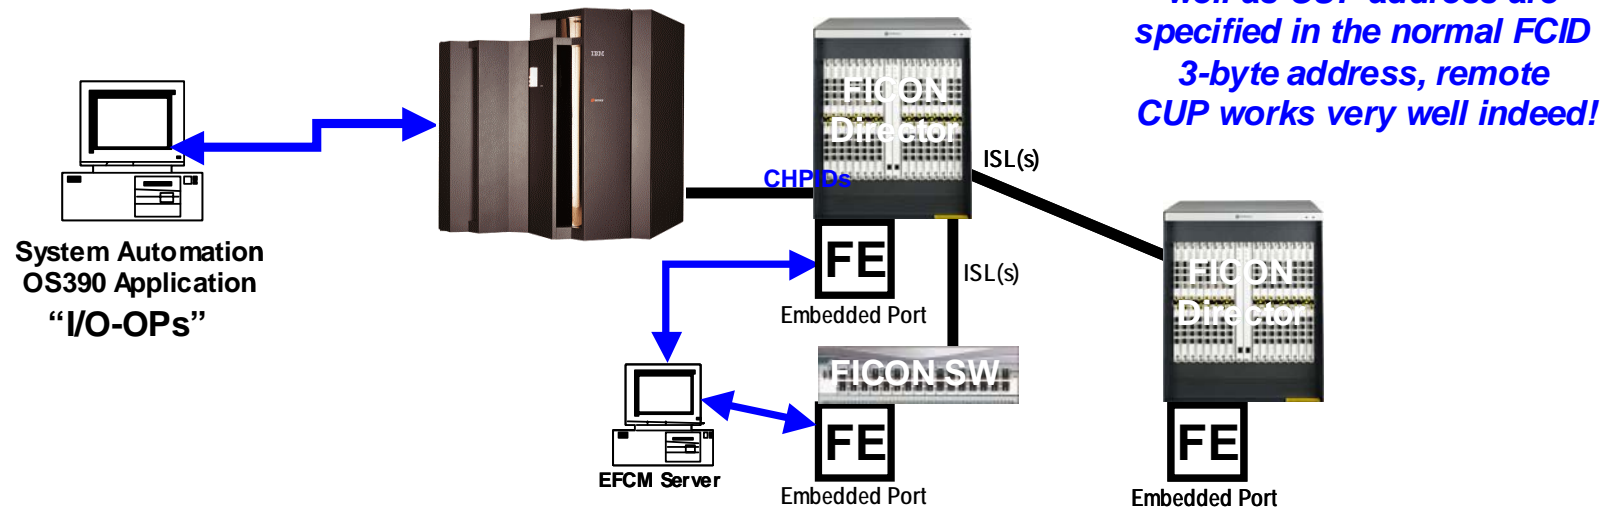
# Safe Switching

- Describes the SA OS/390 process of path removal.
- Process involves contacting all systems in a sysplex and verifying a path can be removed before it is removed
- Each system must logically remove the path and respond positively to the requesting system before the path can be taken out of service.
- Systems usually will only respond negatively if the path in question is the last path for an application.



# CUP Under FICON

## Control Unit Port



- Switches/Directors have an embedded port “FE” in the control processor
- On 256 port Directors, this logical “FE” overlaps the physical port “FE”
- Using CUP, on 256 port boxes, physical ports “FE” and “FF” cannot be used
  - In that case, use “FE” and “FF” for port swapping and for intermix ports
- Only one RMF should attempt to access “CUP” at any one time
- Too much activity to “CUP” can cause missing interrupts – boxed device
- Best practice is still to have 2 or more CHPIDs with access to the CUP

# CUP definition

- Optional for defining in IOCP/HCD
- Without the definition, the following functions do not work:
  - System Automation for z/OS (S/A) I/O-Ops which is used to manage and display the status of the FICON Director.
  - Some RMF reports, such as the RMF 74-7 record (FICON Director Activity Report).




# Recommendations for defining a FICON director

- Define the FICON director as a device
  - Error reporting: switch related HW 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 well as displaying device routing information.
    - All switches must be online as I/O devices on all the systems where S/A I/O-ops manager is running.
- Define at least two paths to the FICON director



# Generating a CUP Device Address



*Because Domain ID as well as CUP address are specified in the normal FCID 3-byte address, remote CUP works very well indeed!*

- 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 switch, **FE** is the embedded port address and **xx** we do not care about currently

- CHPID PATH=(07,17),SHARED,TYPE=FC,SWITCH=7A
- CNTLUNIT CUNUMBR=F008,PATH=(07,17),UNIT=2032,
- UNITADD=((00,1)),LINK=(7AFE,7AFE)
- IODEVICE ADDRESS=(F008,1),UNITADD=00,CUNUMBER=(F008),UNIT=2032

*← ALWAYS the unit type!*

- With multiple 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) or LINK=(74FE) etc.

# Using CUP



## Genning CUP

```
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
```

- From one, and only one LPAR path, RMF communicates with each FICON switch (the CTP card) 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
- A FICON Switching Device report is generated for each FICON device with CUP
- If necessary, timeout values for the Missing Interrupt Handler (MIH) need to be changed from the default 30 seconds to 3 minutes

# FICON Director CUP resources

<b>Resource</b>	<b>Number</b>
Maximum number of logical paths	256
Maximum number of devices	1
Unit address (device addresses)	x'00'
Link address for the CUP	x'FE'
Number of logical control units	1
Logical control unit address	00



# FICON CUP Support

- CUP = Control Unit Port (an embedded not physical port)
  - Allows in-band management of Directors from the management applications on a Mainframe
  - Used for configuration, monitoring, and error handling
- CUP support is provided by all vendors
  - Usually it is an optional licensed feature
  - Supported in single or cascaded FICON environments
- CUP is typically used by customers with:
  - SA/zOS (System Automation software) configuration tool
  - RMF (Resource Measurement Facility) monitoring tool
    - There are a lot of things going on in a FICON storage network
    - Some very important information is gathered by RMF through CUP



# Details of CUP

- There are 39 CUP commands, or Channel Command Words (CCWs), for monitor and control of switch functions
- Due to the historic nature and use of ESCON CUP, FICON CUP commands are oriented toward management of a single switch
- Although the use of CUP in a cascaded environment is supported, CUP is primarily limited to controlling a single switch at a time



# Initial Program Load (IPL)

- The IPL is a special configuration file identified by the unique file name “IPL”
- The parameters contained in the IPL file are the same as other configuration files, but there are special uses for this file
- The configuration defined by the IPL will be applied upon reboot of the FICON Director/switch
  - This could be a scheduled downtime reboot
  - Or this could be an unscheduled power failure for example (U.S. August 2003)
- Active=Saved Setting (a Mode Register bit - discussed later):
  - When the Active=Saved Mode is enabled, changes made to the active configuration (HCD settings) are also stored in the IPL
  - When Active=Saved Mode is disabled, at reboot the HCD settings will be restored to the FICON switching device which might not accurately reflect the hardware state of the box
    - Especially if port swaps have not been updated in the HCD

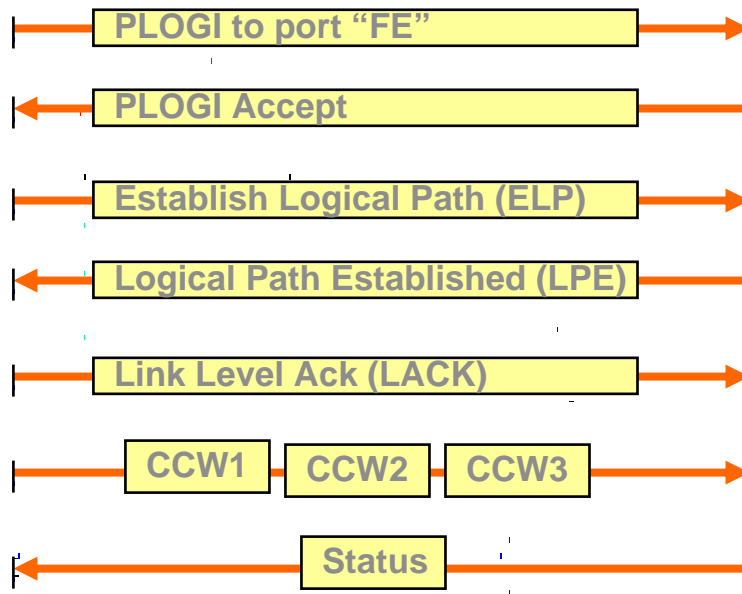



# CUP Logical Path Establishment and CCW Flow

Mainframe (Channel)

FICON Director

Control Unit



# Using CUP

## In-band Monitoring of the FICON Director

- Port Statistics
  - Number of words transmitted, received, etc
  - These counters are already available in the switch
- Switch Node Identifier
  - Serial Number, Manufacturer, etc
  - The information is the same Switch Node ID provided in the RNID ELS
- Configuration File Information
  - List of Configuration files residing on the switch
  - Actual file content, including port address name and port connectivity
- History Summary (Directory History Buffer)
  - Each change in status or configuration of the ports is logged in a history buffer
  - The history buffer can be retrieved through CUP
- Switch Configuration Data
  - Timeout values, number ports per card, etc



# PDCMs

## Prohibit Dynamic Connectivity Mask and Zoning

- The Prohibit Dynamic Connectivity Mask (PDCM) is a mechanism to define port connectivity (also referred to as prohibit/allow and blocking and unblocking)
- Prohibiting ports using PDCM across zones is redundant
  - Zoning rules already prohibit the devices from communicating
  - In the event of a conflict between zoning and PDCM definitions, the most restrictive rules are automatically applied
- E-Ports and PDCM
  - Prohibiting certain device types from using an ISL
    - Example: not allowing your DASD and tape traffic to traverse the same ISL(s)



# FICON Configure Allow/Prohibit Matrix

## PDCMs

- Notice the tool tip is showing that the mouse is over the cell that intersects 0C and 08.
- Using the PDCM matrix, we can physically block the connection of F-port(s) to E-port(s)
- Physically cannot send a frame to that E-port

ISL Port

Addr	Port Name	Blocked	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F	10	11	12	13	14
04	port04	<input checked="" type="checkbox"/>	⊘	⊘	⊘				⊘	⊘	⊘								
05	port05	<input checked="" type="checkbox"/>	⊘						⊘	⊘	⊘								
06	port06	<input checked="" type="checkbox"/>	⊘			⊘			⊘	⊘	⊘								
07		<input type="checkbox"/>			⊘						⊘								
08		<input type="checkbox"/>									⊘								
09		<input type="checkbox"/>									⊘								
0A	port0A	<input type="checkbox"/>	⊘	⊘	⊘						⊘								
0B	port0B	<input type="checkbox"/>									⊘								
0C	neverUsePort	<input type="checkbox"/>	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘	⊘
0D		<input type="checkbox"/>									⊘								
0E		<input type="checkbox"/>									⊘								
0F		<input type="checkbox"/>									⊘								
10		<input type="checkbox"/>									⊘								
11		<input type="checkbox"/>									⊘								
12		<input type="checkbox"/>									⊘								
13		<input type="checkbox"/>									⊘								
14		<input type="checkbox"/>									⊘								

These F-ports cannot be allocated to this E-port

Active=Saved (The activated configuration will also get copied to the IPL file.)


Activate Save As... Close Help



# Additional Benefit - Port Swapping

## Move FICON Cables without IOCP Update

- Swaps port assignments for a pair of switch ports
- Enables FICON storage devices to be moved onto different ports without the need to update the mainframe I/O configuration file (via HCD usually)
- Supports Mainframe zOS use of static I/O definitions
- Users can swap ports across all ports within a switch or Director
- NOTE: If the optical transceiver has failed, and not the port, just swap out the optics and re-cable back into the original port
  - A port swap would not be necessary!




Defective port is address-swapped to a good port

# EFC Manager for FICON Management

## “Port Swap”



# FICON Management - Port Swap



**Port address swapping**

# Reasons That CUP Is Very Useful

- If the user wants to be able to do in-band management from their Hardware management consoles, they need CUP
  - If they want to be able to use the management tools they are familiar with from ESCON, be it ESCON manager, System Automation OS/390 I/O-Ops, etc.
- If it is important to them to get the Service Information Messages (SIM) to the MVS console for FICON device hardware failures
- If the user plans on using the Dynamic Channel Path Management function of Workload Manager and IRD (DCM) when it becomes available for FICON
- If they want to do RMF Monitor I reports for reporting on FICON Director activity statistics in the SMF Record Type 74 Subtype 7, they need CUP
  - RMF Monitor I does long term data collection for resource utilization



# Differences in CUP from ESCON to FICON

- **ESCON did not use buffer credits**
  - ESCON was half-duplex so an acknowledgement had to be received after each and every frame was sent
- **FICON does use buffer credits**
  - FICON is full-duplex, buffer credits provide flow control
- **RMF needs to report on these fancy, new “Buffer Credits”**
  - IBM chose to do this by creating RMF 74 subtype 7 records
  - RMF 74 subtype 7 records are new with FICON
  - RMF 74 subtype 7 records are only provided via FICON CUP
  - RMF 74 subtype 7 records are needed to understand FICON performance
  - RMF 74 subtype 7 records create a “FICON Director Activity Report” during each reporting interval for each CUP found



# CUP for Use With RMF

- Much has changed in FICON

- Multiple ESCON links can be aggregated onto a single FICON link
- FICON allows full-duplex operations rather than the unidirectional technology of ESCON
- FICON allows Fan In and Fan Out to really be effective (ESCON did not)
- ESCON did not have a function like FICON Cascading
- And lets not forget that all of these could be out of tune at the same time!



# CUP for Use With RMF

- In **MY** opinion, RMF is the **KEY** motivator for using CUP:
  - RMF is a mainframe-centric view of the channel and device activity
  - Without CUP, RMF has no idea how fabric connectivity is affecting the I/O activity that it sees on any given CHPID
  - With CUP, an RMF FICON switch shares the gathered statistics about each of its ports enabling RMF to accurately report I/O channel and device activity as well as network timings
- CUP is A Fee-based software key that must be configured on each and every FICON switch for which CUP will be enabled
- What has changed that makes CUP more important in FICON?



# ESCON Device Information Blocks (DIBs)

## -a brief historical perspective

- ESCON DIBs contain control, status, or user data.
- Storage director's (control unit) adapters define maximum DIB size employed by configuration.
  - 8,16,32,64,128,256,512, or 1024 bytes long
- A microprocessor on the ESCON adapter manages the adapter's data buffer.
  - 3990-3/6 ESCON adapter held (8) 256 byte DIBs.



# ESCON Device Information Blocks (DIBs)

## -a brief historical perspective (2)

- During initialization, a subsystem's ESCON adapters notify the channel subsystem of DIB and data buffer sizes to be employed.
- The channel subsystem employs these values for all future communications with the subsystem.
- Assuming ESCON adapter is receiving data, it must successfully pass each DIB it receives to buffer areas in the storage director before it can request another DIB be sent by the channel to avoid potential of data overrun.
- While the channel may initially schedule a full data buffer of DIBs be sent to a device, further DIBs may only be transmitted after a request has been received from the ESCON adapter signalling that a prior DIB has been successfully passed along to the storage director.




# Buffer Credit Concepts

- Define the maximum amount of data that can be sent prior to an acknowledgement
- Buffer credits are physical ASIC port or card memory resources and are finite in number as a function of cost
- Within a fabric, each port may have a different number of buffer credits
- The number of available buffer credits is communicated at fabric logon (FLOGI)
- One buffer credit allows a device to send one 2112 byte frame of data (2K usable for z/OS data)
- Assuming that each credit is completely full, you need one credit for every 1 KM of link length over a 2 Gbit fibre
- Unfortunately, z/OS disk workload rarely produce full credits. For a 4K transfer, the average frame size for a 4K transfer is 819 bytes
  - Hence, five credits would be required per KM over a 2 Gbit fibre



# CUP for Use With RMF

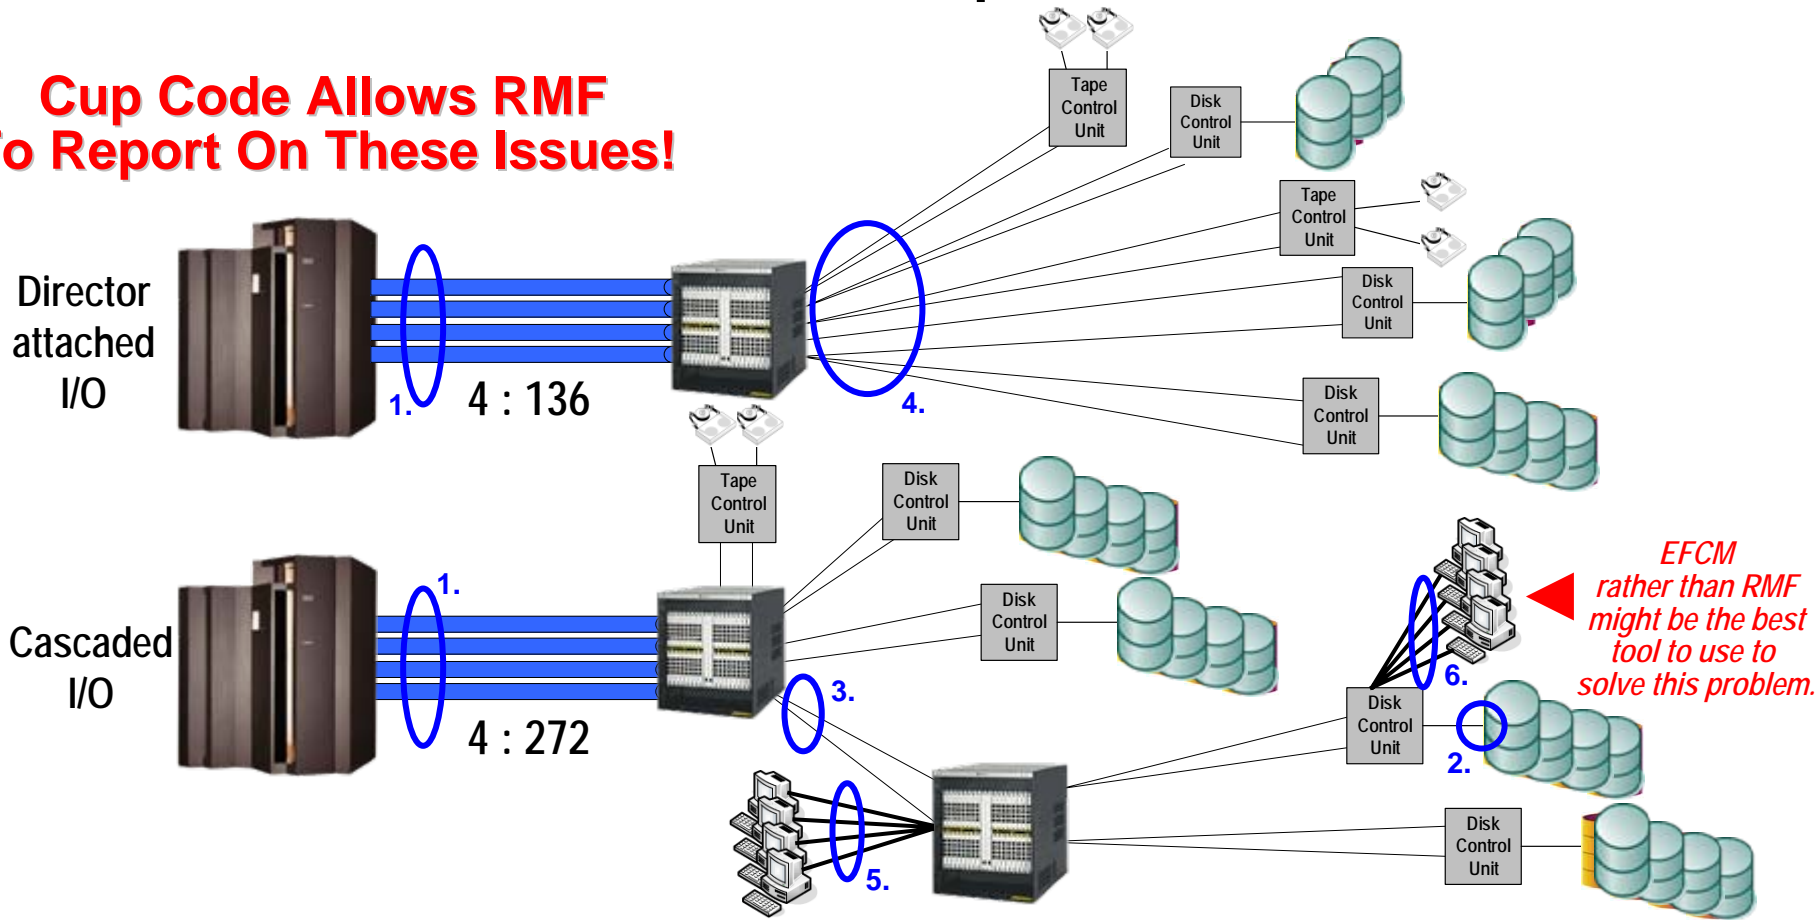


**Does EXACTLY what RMF Should Do!**

- For RMF reporting, constructs such as Unit Control Block (UCB) and Logical Control Unit (LCU) are used to store gathered statistics (RMF Monitor I)
- A UCB is a device that is accepting or creating frame traffic
- A LCU is a control unit definition that encompasses one or more UCBs at a higher level
- RMF assumes he owns and controls everything in the universe (what SAN?)
- Can only measure and report on what is told to him about devices and paths
- Does not really understand anything about channel aggregation, fan in-fan out, and ISLs
- Most z/OS customers live and die by their RMF performance tuning!

# Networked FICON Has Special Issues


**Cup Code Allows RMF  
To Report On These Issues!**



1. Could have FICON CHPID oversubscription due to incorrect aggregation
2. Could have Storage Port oversubscription due to incorrect Fan-Out aggregation
3. Could have ISL congestion due to changing workloads or improper FanIn-FanOut aggregation
4. Could have a permanent or intermittent frame pacing delay situation affecting performance
5. Maybe in a FICON and FC intermix structure, the SAN is causing CU busy problems out in storage
6. Maybe FICON and FC are using the same DASD and Open Systems is causing CU busy problems

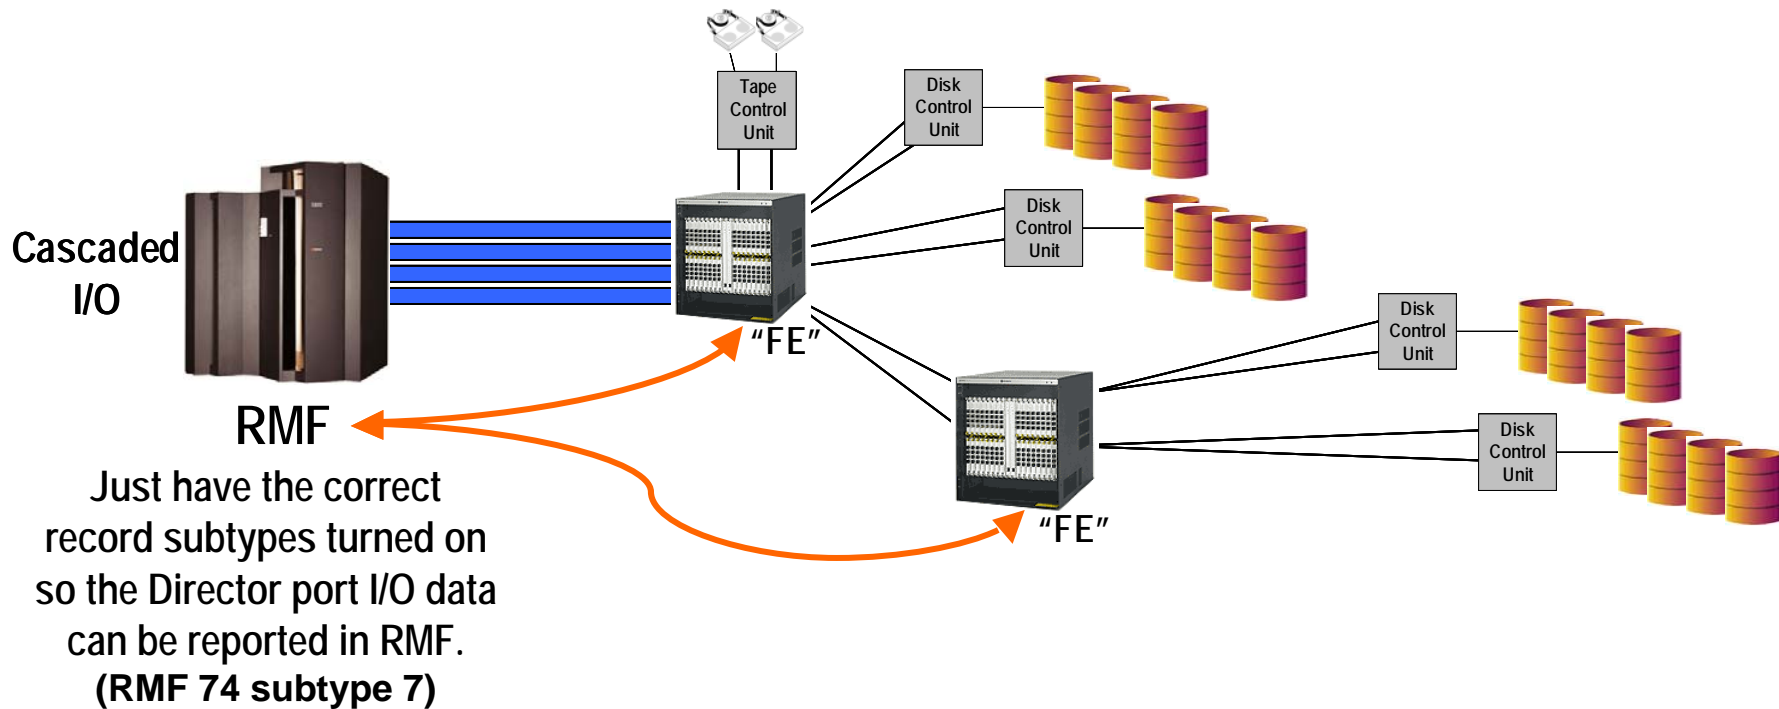


# CUP for Use With RMF



Just have the correct record subtypes turned on so the Director port I/O data can be reported in RMF.  
**(RMF 74 subtype 7)**


# CUP for Use With RMF



- From one, and only one LPAR's path, RMF communicates with each FICON switch in turn, requesting that its "FE" port dump all FICON port statistics down to the mainframe.
- Then RMF distributes that information to all the other LPARs on the mainframe
- A FICON Switching Device report is generated for each FICON device with CUP

Also, timeout values need to be changed from 30 seconds to 3 minutes on the Missing Interrupt Handler (MIH) on older FICON directors


**BROCADE**



# Frame Pacing Delay



# RMF 74 Subtype 7 Records



- Four data classes of data are reported by the 74 subtype 7
- Port data includes average read/write frame sizes, average bandwidth, error count, and pacing delays for each port. *Frame pacing occurs when a director port exhausts its available credits.* Frame pacing delays are measured in 2.5 micro-second units
- Data is collected for each RMF interval if FCD is specified in your *ERBRMFnn* parmlib member


# Frame Pacing Delay

- AVG FRAME PACING
  - Defined by RMF as the average number of time intervals of 2.5 microseconds that a frame has to wait before it could be transmitted due to no buffer credits being available on a given director port.
- You always want to see a zero value in this field!
  - Reporting on this value was one of the primary reason that the RMF 74-7 record was developed – it was not needed for ESCON
  - A non-zero value in the AVG FRAME PACING field indicates that you have an issue with insufficient BB Credits
  - It is critical to use CUP in any FICON environment in which distance extension is being utilized
  - 4Gbps may create more Frame Pacing Delay issues than 2Gbps
- z/OS disk workloads rarely use a "full" 2148 byte credit
  - For example, with a 4k block transfer, the average frame size for each 4k transfer is typically about 819 bytes



# Where does frame pacing occur?

- Incorrect number of BB credits (not enough) assigned on a port
- Poorly architected environment:



# FICON Director Activity Report

F I C O N   D I R E C T O R   A C T I V I T Y										
z/OS V1R4			SYSTEM ID KS01		START	10/03/2005-13.55.00		INTERVAL 000.05.00		
IODF = 0C			RPT VERSION V1R2		END	10/03/2005-14.00.00		CYCLE 1.000 SECONDS		
SWITCH DEVICE: 0001			CR-DATE: 07/07/2005		CR-TIME: 18.47.41		ACT: POR		MODEL: 001	
PORT			SWITCH ID: **		TYPE: 006140		MAN: MCD		PLANT: 01	
-CONNECTION-			AVG FRAME		AVG FRAME SIZE		PORT BANDWIDTH (MB/SEC)		ERROR	
ADDR	UNIT	ID	PACING	READ	WRITE	-- READ --	-- WRITE --	COUNT		
04	CHP-H	47	0	433	593	1.01	1.52	0		
05	CU	----	0	830	1195	2.30	5.81	0		
06	CU	----	0	140	69	0.00	0.00	0		
07	CU	C052	0	591	85	0.00	0.00	0		
	CU	C050								
08	CHP-H	45	0	400	577	0.98	1.55	0		
09	CU	----	0	1355	374	8.42	1.00	0		
0A	-----	----	0	0	0	0.00	0.00	0		
0B	-----	----	0	0	0	0.00	0.00	0		
0C	CHP-H	48	0	431	601	1.02	1.56	0		
0D	CU	----	0	776	374	3.57	1.41	0		
0E	CU	----	0	0	0	0.00	0.00	0		
0F	CU	C053	0	1773	78	0.08	0.00	0		
	CU	C051								
10	CHP-H	46	0	366	716	1.06	2.32	0		
11	CU	----	0	1099	393	4.62	0.97	0		
12	-----	----	0	0	0	0.00	0.00	0		
13	-----	----	0	0	0	0.00	0.00	0		
14	CHP-H	50	0	533	832	0.32	0.58	0		
15	CU	----	0	868	1223	2.25	5.55	0		
16	CU	----	0	158	72	0.00	0.00	0		
17	CU	C053	0	1761	77	0.09	0.00	0		
	CU	C051								
18	CHP-H	49	0	378	745	1.04	2.33	0		
19	CU	----	0	1118	399	4.83	0.99	0		
1A	-----	----	0	0	0	0.00	0.00	0		
1B	-----	----	0	0	0	0.00	0.00	0		
1C	CHP-H	51	0	737	535	0.34	0.17	0		
1D	CU	----	0	877	1230	3.22	7.91	0		
1E	CU	----	0	0	0	0.00	0.00	0		
1F	CU	C052	0	590	82	0.00	0.00	0		
	CU	C050								
20	CHP-H	4A	0	374	756	1.04	2.40	0		
21	CU	----	0	1472	413	3.51	0.36	0		
22	-----	----	0	0	0	0.00	0.00	0		
23	-----	----	0	0	0	0.00	0.00	0		



# Frame Pacing Delay Being Reported

FICON DIRECTOR ACTI

z/OS V1R7

SYSTEM ID PDM1

DATE 11/28/2006

RPT VERSION V1R7 RMF

TIME 21.44.00

IDDF = 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 ADDR	-CONNECTION-	ID	AVG FRAME PACING	AVG FRAME SIZE	READ	WRITE	PORT BANDWIDTH (MB/SEC)	
	UNIT						-- 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



# Local Frame Pacing Delay

- How can you run out of buffer credits inside a datacenter?
  - Frame pacing delays occur when multiple, heavily used paths merge into a single FICON link
  - Frame pacing delays can contribute to PEND, DISC, and CONN time measurements



**Frame Pacing Delay  
is caused  
by running out  
of buffer credits!**

*UCBs serviced by these storage ports  
are probably experiencing additional delays  
usually reported as PEND Time and CONN Time  
and sometimes as DISC time*

**Frame Pacing Delay came about with  
FC and FICON so it is not a factor in  
ESCON performance!**



# Suggestion

- Use the RMF 74-7 record as a way to help narrow down/troubleshoot performance problems in your environment.
- The RMF 74-7 record in and of itself will not allow you to precisely diagnose what is causing a performance problem, other than to indicate BB Credit issues.



# Access CUP via Low Utilized CHPIDs


## EXAMPLE



Questions??



**BROCADE**



**THANK YOU**

