



# I/O Priority Queueing Channel Subsystem Channels I/O Control Units



© Copyright IBM Corp. 2000, 2001

## CSS I/O Priority Queueing

New function on 2064

Allows OS to specify priority value to channel subsystem for I/O requests

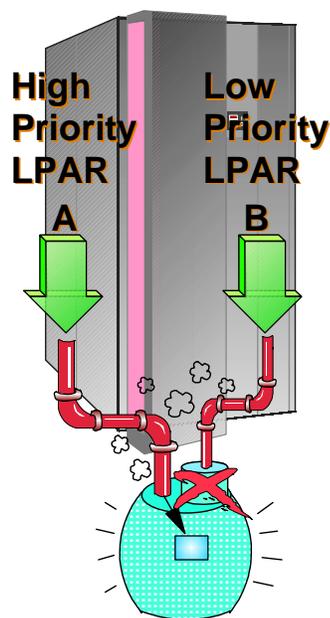
z/OS in goal mode sets priority value based on goals defined in WLM

Complementary to current goal mode priority management:

- IOS UCB queues
- CSS dequeuing priority
- Channel queuing and dequeuing priority
- Shark control unit

FICON adapter priority support

Function can be enabled/disabled on CEC via operator control



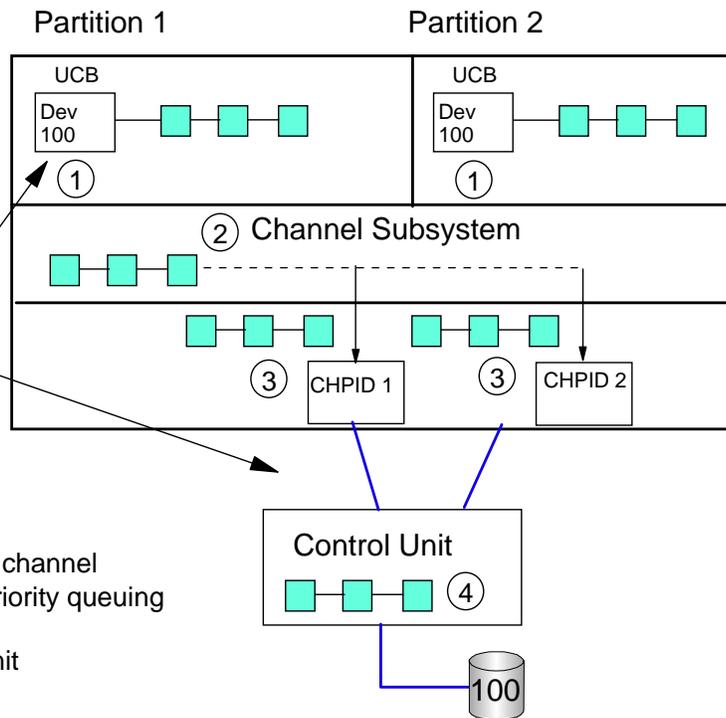
# I/O Queueing Points

End-to-end I/O Priority Management on z/OS in WLM Goal Mode, 2064 and ESS

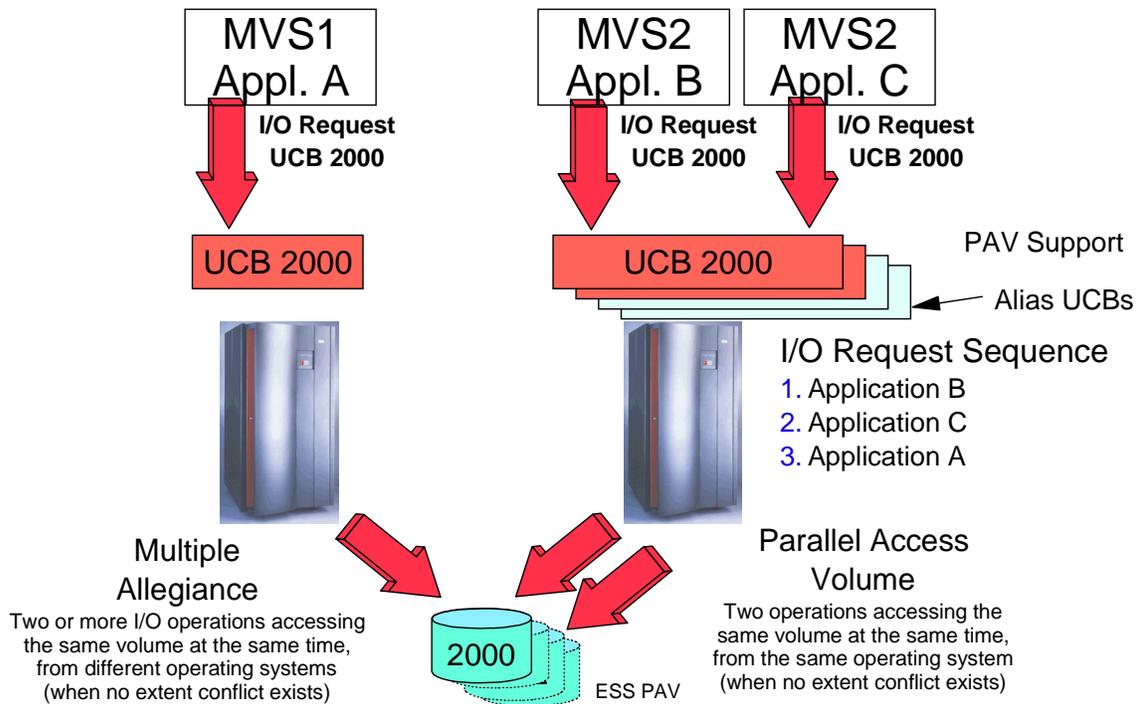
Same value used for UCB and control unit queueing

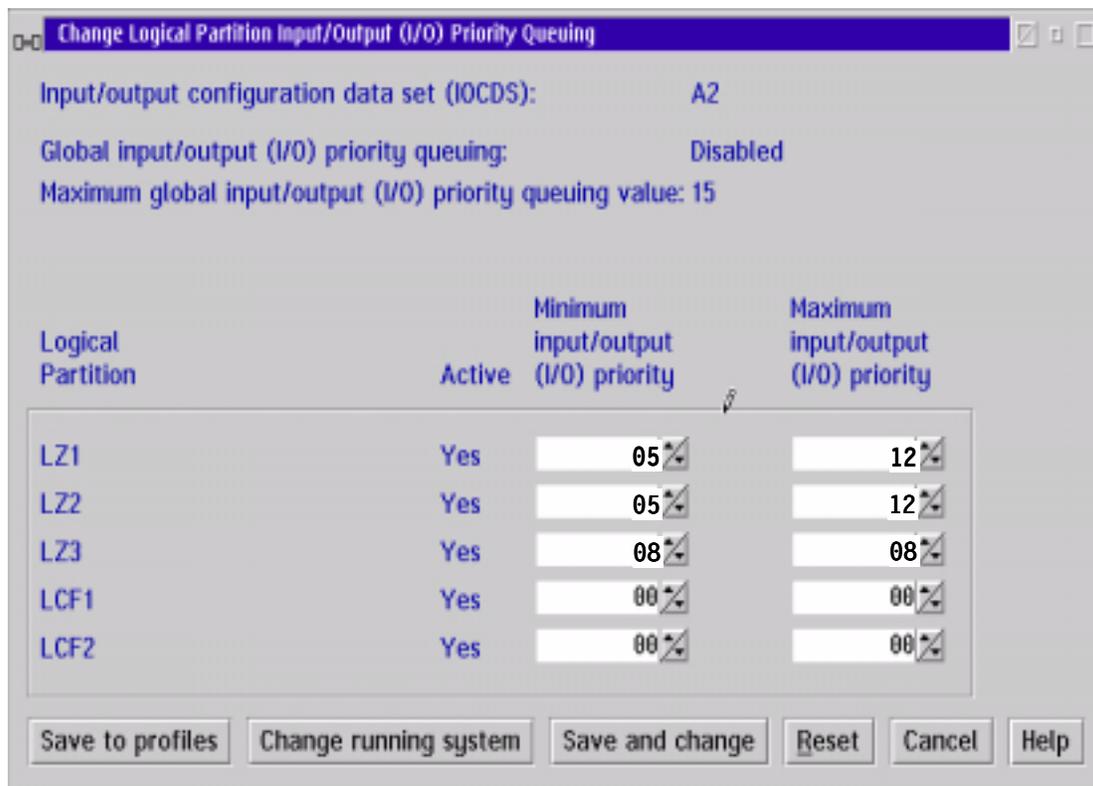
## Queueing Points

1. IOS UCB Queue
2. In CSS waiting for channel
3. FICON Channel priority queuing and dequeuing
4. In Shark control unit



# Queueing at the UCB





## CSS I/O Priority Queueing

CSS I/O Priority Queueing is switched on or off at the CPC level.

*Each* LPAR has a range of priorities, assigned on the HMC

Non-z/OS operating systems use the fixed priority assigned on the HMC. If a range is specified, it will select the lower number

All LPARs within the same LPAR Cluster should have the same range of priorities assigned

## CSS I/O Priority Queueing

**WLM assigns UCB and CU I/O priority so that:**

- System related SCPs are assigned highest priority
- An SCP missing its goal because of I/O delay gets helped
- An SCP competing for the same devices with a more important SCP that is missing its goal ( $PI > 1$ ) is the donor

**WLM assigns CSS I/O priorities so that:**

- System related SCPs are assigned highest priority
- A High Importance SCP missing its goal has next highest
- SCPs meeting goals are managed so that those using fewer I/O resources are assigned higher I/O priorities than heavy I/O users
- Discretionary work has the lowest priority
- WLM knows the range of priorities specified in HMC and maps to this range

**I/O weight is ratio of Connect Time to Elapsed Time**

## I/O Priorities Compared

CSS Priority	Type of Work
FF	System Work
FE	Importance 1 and 2 missing goals
FD	Importance 3 and 4 missing goals
F9-FC	Meeting goals - Adjusted by ratio of connect time to elapsed time
F8	Discretionary

UCB Priority	Type of Work
FF	SYSTEM service class
FE	SYSSTC service class
F9-FD	Managed by WLM Policy Adjustment
F8	Discretionary

## I/O Priorities Compared

	CSS Priority	Type of Work
	FF	System Work
2	FE	Importance 1 and 2 missing goals
3	FD	Importance 3 and 4 missing goals
1	F9-FC	Meeting goals - Adjusted by ratio of connect time to elapsed time
4	F8	Discretionary

Actual ORB value is x'0x'

UCB Priority	Type of Work
FF	SYSTEM service class
FE	SYSSTC service class
F9-FD	Managed by WLM Policy Adjustment
F8	Discretionary



## Implementing CSS I/O Priority Queueing



© Copyright IBM Corp. 2000, 2001

## Implementation Steps

IBM 2064 running z/OS in 64-bit mode

z/OS must be running in WLM goal mode

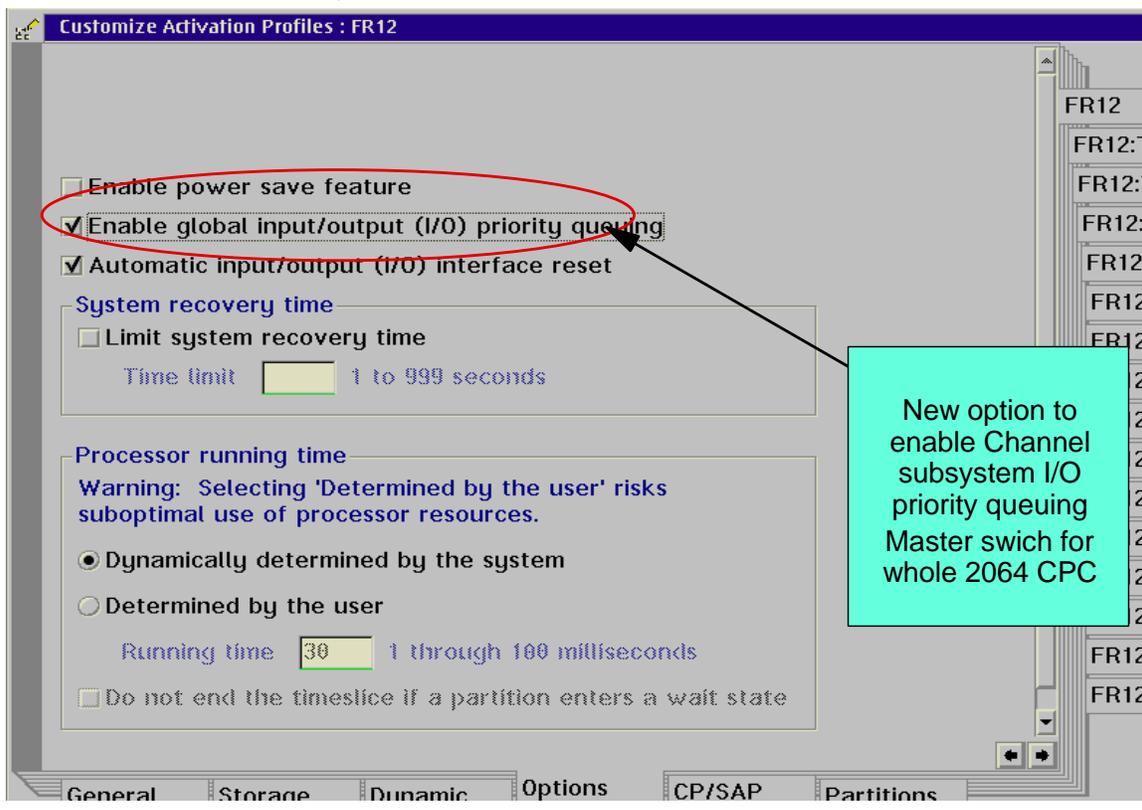
'I/O Management' set to YES in WLM policy

CSS I/O Priority Management enabled in HMC

Valid range of priorities specified in HMC (for each LPAR)

Specify default I/O priority for non-z/OS images

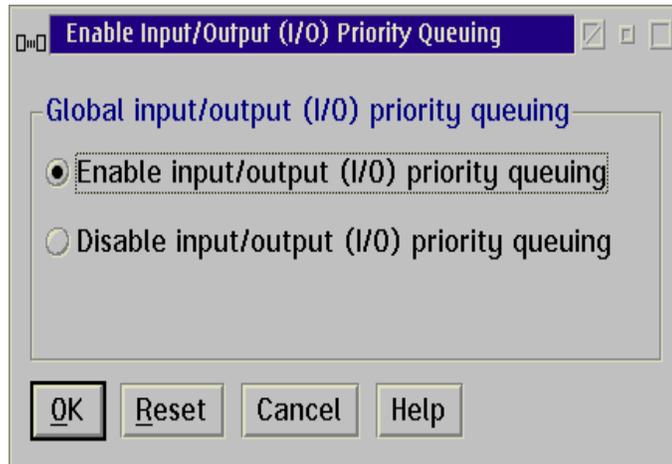
## HMC CSS IOPQ Enablement



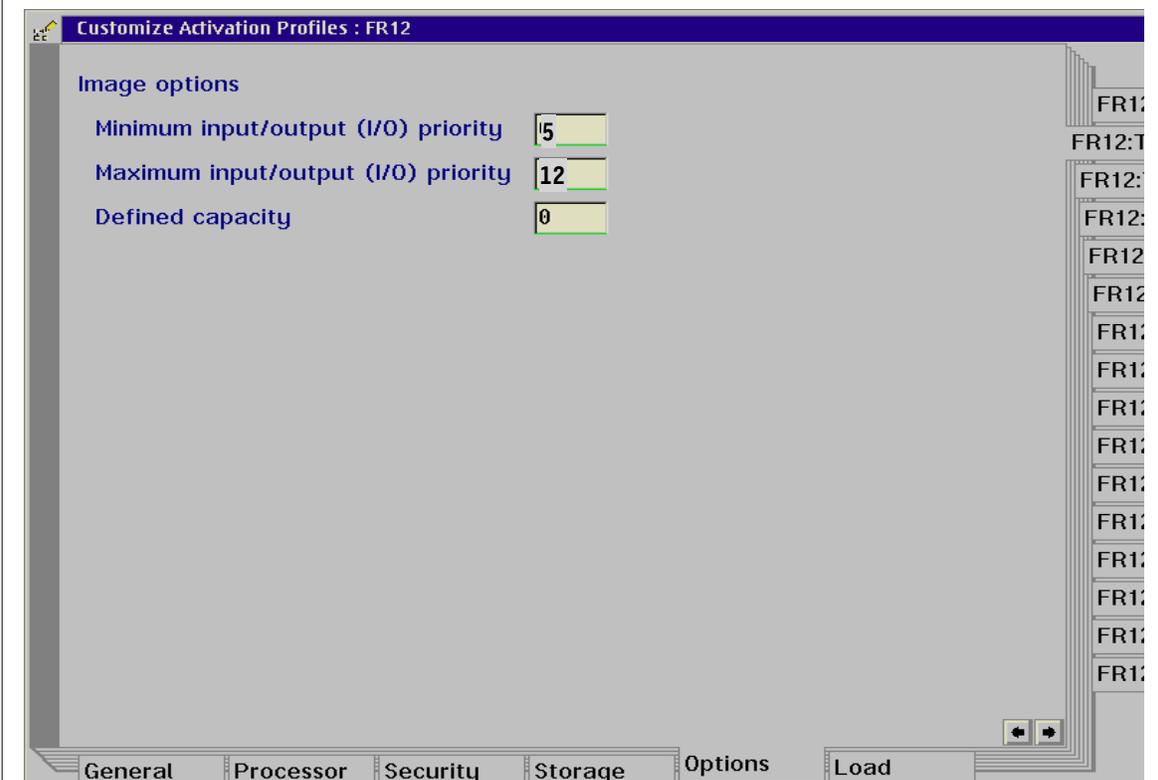
The screenshot shows the 'Customize Activation Profiles : FR12' window in the HMC. The 'Enable global input/output (I/O) priority queuing' checkbox is checked and circled in red. A callout box points to this checkbox with the text: 'New option to enable Channel subsystem I/O priority queuing Master switch for whole 2064 CPC'. Other options include 'Enable power save feature' (unchecked), 'Automatic input/output (I/O) interface reset' (checked), 'System recovery time' (with a 'Limit system recovery time' section containing a 'Time limit' field set to 1 to 999 seconds), and 'Processor running time' (with a 'Warning: Selecting 'Determined by the user' risks suboptimal use of processor resources.' and radio buttons for 'Dynamically determined by the system' and 'Determined by the user', plus a 'Running time' field set to 30 milliseconds and a 'Do not end the timeslice if a partition enters a wait state' checkbox).

## Enable/Disable dynamically

As well as defining IOPQ in the activation profile, you can also dynamically enable it or disable it at the CPC level....



## IOPQ Range for LP



Change Logical Partition Input/Output (I/O) Priority Queuing

Input/output configuration data set (IOCDs): A2

Global input/output (I/O) priority queuing: Disabled

Maximum global input/output (I/O) priority queuing value: 15

Logical Partition	Active	Minimum input/output (I/O) priority	Maximum input/output (I/O) priority
TC4S0	Yes	05	12
TC4S1	Yes	05	12
TC4S2	Yes	05	12
TC4S3	Yes	00	00
TC4S4	Yes	00	00
TC4S5	Yes	00	00
TC4S6	Yes	00	00
TC4S7	Yes	00	00
TC4S8	Yes	00	00
TC4S9	Yes	00	00
TC4SA	Yes	00	00
TC4SB	Yes	00	00
TC4SC	Yes	00	00
CF2	Yes	00	00
...	..	..	..

Buttons: Save to profiles, Change running system, Save and change, Reset, Cancel, Help



## I/O Priority Queueing Priority Flow



© Copyright IBM Corp. 2000, 2001

# I/O Queueing Controls and Flow

End-to-end I/O Priority Management on z/OS in WLM Goal Mode, 2064 and ESS

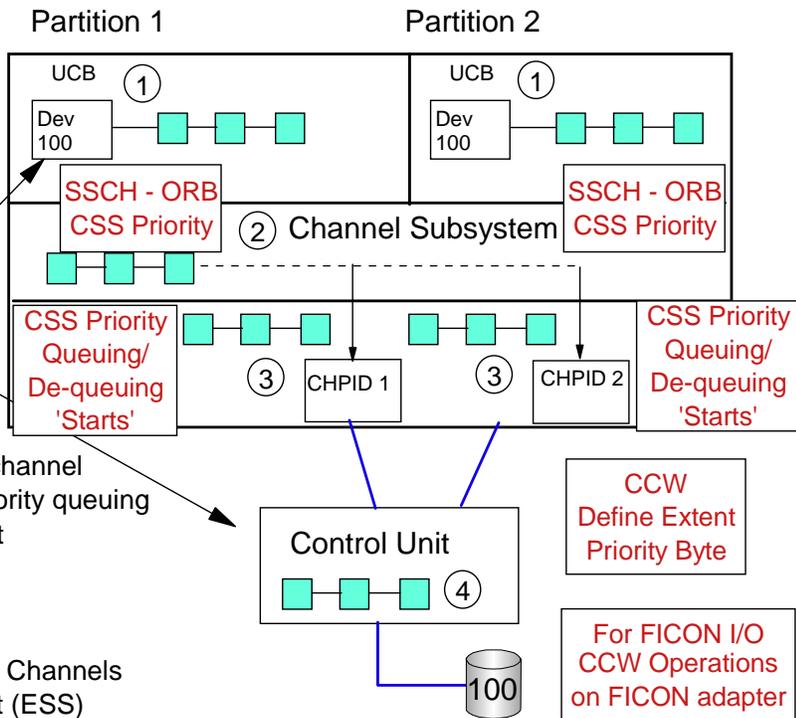
Same value used for UCB and control unit queueing

Queueing Points

1. IOS UCB Queue
2. In CSS waiting for channel
3. FICON Channel priority queuing
4. In Shark control unit

Queueing Support

1. z/OS
2. z900 CSS and z900 Channels
3. Storage Control Unit (ESS)



## Implementing CSS IOPQ

And that's it!  
Any (more) questions

Thank you Frank Kyne for the initial presentation