


# IBM i Suspend and Resume

Jeff Blight  
Power CTSS



# Agenda

*What we will talk about today:*


**1 Suspend / Resume**

**2 What implications are there on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Resuming an IBM I partition**

**5 Migrating Aix LPAR aware**



## Suspend / Resume

- With IBM v7.1 TR2 IBMi started to support the capability to suspend an active partition and resume that partition to state it was in at the point of suspension
  
- Once partition is suspended
  - System resources to be freed up and reused by other partitions on the system  
Virtual server adapter configuration entries with associated VIOS partitions are removed and saved with the image in a storage device
  - Negates or reduces the requirement to shutdown and restart partitions should a system need to under go maintenance
  
- Suspending partitions can be:
  - Resumed – returns the partition to the state it was in when suspended.
  - Shutdown – Invalidates suspend state and moves partition to a state of powered off. If the storage device that contains the partition state is available then all saved virtual server adapter configuration entries are restored.
  - Migrated (Requires Power VM Enterprise Edition) \*Future stated direction for IBM i

## Suspend / Resume - Partition Hibernation Applicability / Benefits

- Resource balancing
  - e.g. suspend lower priority and/or long running workloads to free resources. Useful for performance management and energy management.
  
- Planned CEC outages for maintenance/upgrades
  - Hibernation may be used in place of or in conjunction with partition mobility.
  - Hibernate/resume may require less time and effort than manual database shutdown and restart, for example.
  
- Potential future usage: Cloud based image capture and deployment

# Suspend / Resume for IBM i - Minimum Requirements

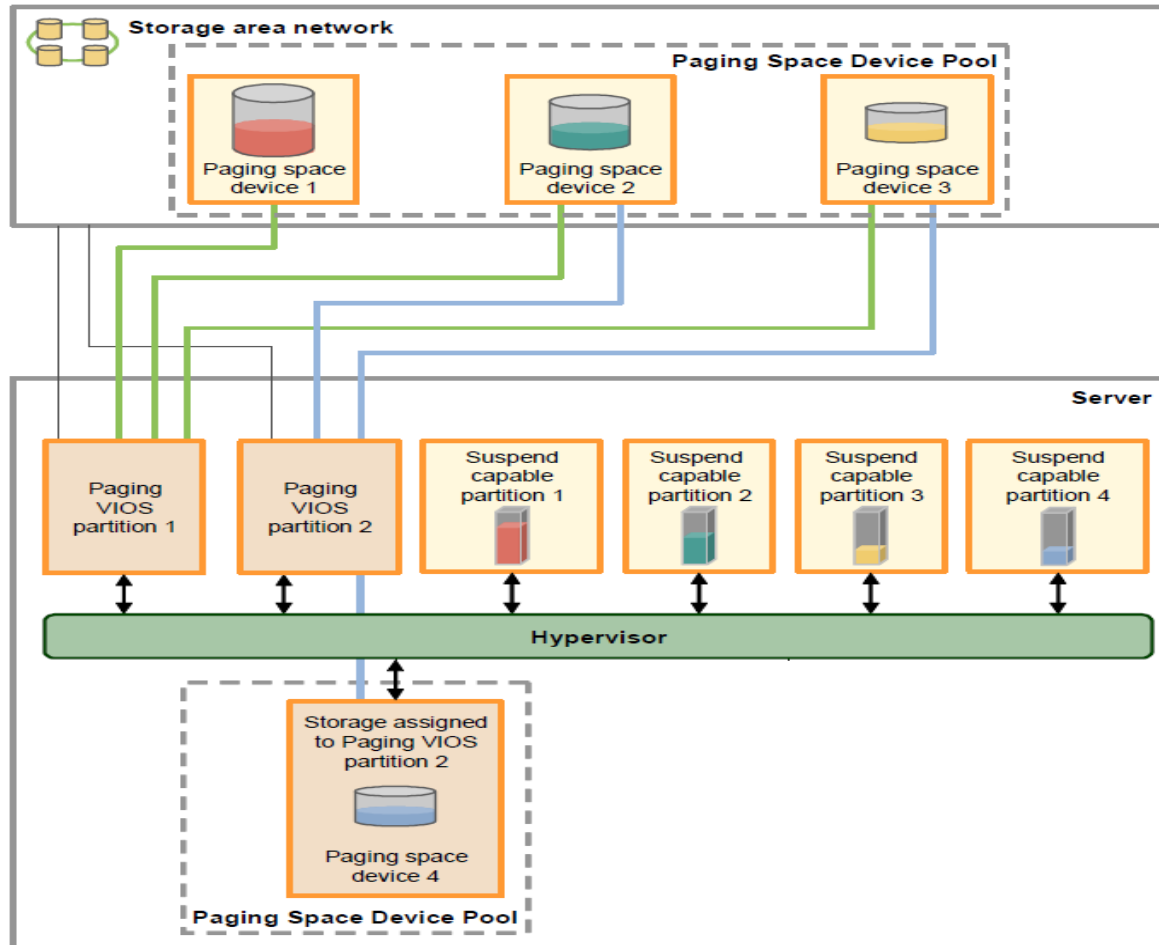
- IBM v7.1 TR2 or above
- Power VM Standard Edition or Enterprise Edition
- POWER 7 Firmware 7.2.0 SP1 or above
- HMC v7 r7.2.0 or above
- VIOS 2.2 0.11-FP24 SP01
- Paging space storage devices
- Suspend capable partitions

Note: At the software levels above, the maximum supported concurrent operations for Suspend/Resume and Partition Mobility is limited to 4. However, there is no limitation for the maximum number of partitions that can be in a suspended state.

## Suspend / Resume – Configuration Requirements

- The reserved storage device must be kept persistently associated with the logical partition.
- The HMC ensures that the Reserved Storage Device Pool is configured with at least one active VIOS partition available in the pool.
- The logical partition must not have physical I/O adapters assigned.
- The logical partition must not be a full system partition, a VIOS partition or a service partition.
- The logical partition must not be an alternative error logging partition.
- The logical partition must not have a Barrier Synchronization Register (BSR).
- The logical partition must not have huge memory pages.
- Monitoring systems should be manually stopped/resumed while suspending and resuming logical partitions.
- Two WWPNs are required for NPIV and must be zoned in the switch.
- Virtual Media Library must be manually removed from the Virtual I/O Server configuration before suspending a partition.

# Suspend / Resume - pictorial representation

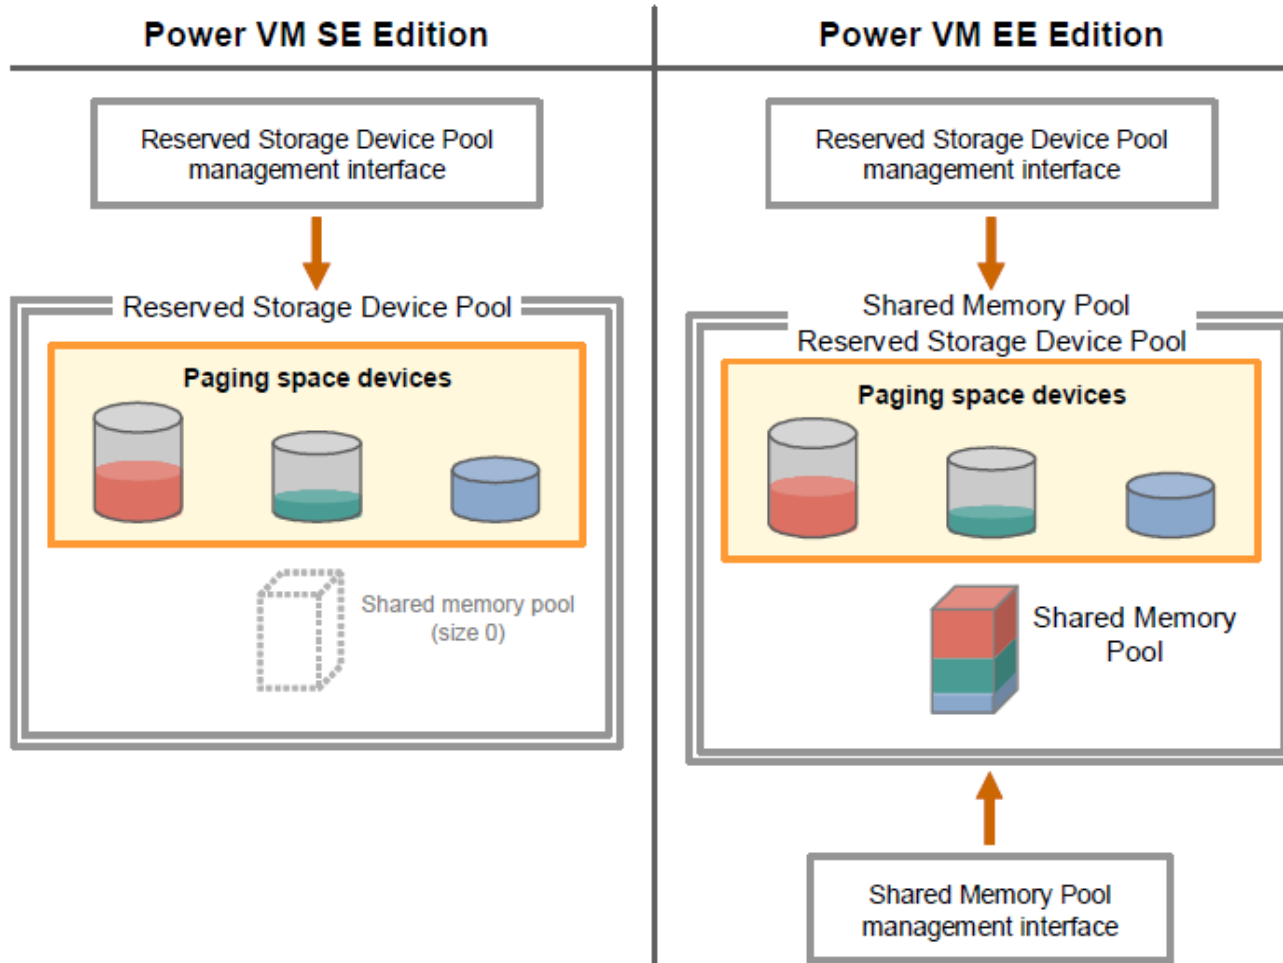


## Suspend / Resume – Reserved Storage Device Pools and Devices

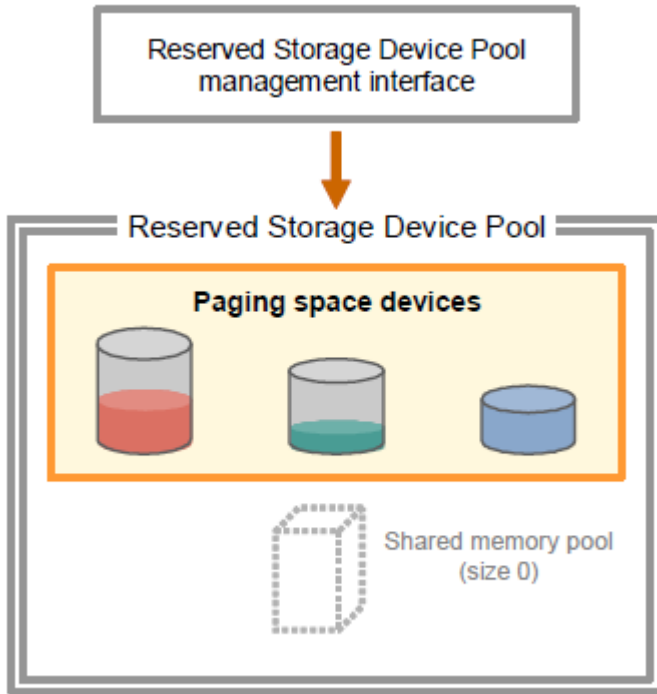
- Exclusive storage provider for Active Memory Sharing (AMS), Suspend/Resume, remote restart and future capabilities
- A Reserved Storage Device Pool has the reserved storage devices (storage devices) and it is basically like a Shared Memory Pool of memory size 0
- When a partition is suspended its state is stored on a reserved area of disk termed a storage device, the storage space required to hold a suspend partition is approximately 110% of the partition's configured maximum memory size
- Reserved Storage Device Pools can be assigned to at most two VIOS for redundant path to the storage devices
- The Reserved Storage Device Pool is managed via the HMC interface at the HMC.



# Suspend / Resume - Reserved Storage Pool Management



# Suspend / Resume - PowerVM SE: RS Pool Management




- Create/Delete the RS device pool
- Add/Remove VIOS to/from the pool
- Add/Remove reserved storage devices to/from the pool

HMC auto picks an unused and suitable\* device from this pool to store partition suspend data

\* Size suggested by PHYP

## Suspend / Resume - PowerVM EE: Two management interfaces. AMS pool and RS device pool interactions



- When AMS pool is created, RS pool automatically gets created
- When AMS pool is deleted, RS pool is **not** automatically deleted
- When RS pool is created, AMS pool is not automatically visible to user
- When RS pool already exists, user to specify the memory attributes only to create AMS pool
- When AMS pool exists, RS pool can't be deleted as it serves as storage provider to AMS pool
- A pool can't be deleted when the devices in pool are in use by a partition

# Suspend / Resume - Reserved Storage Device Pool Management

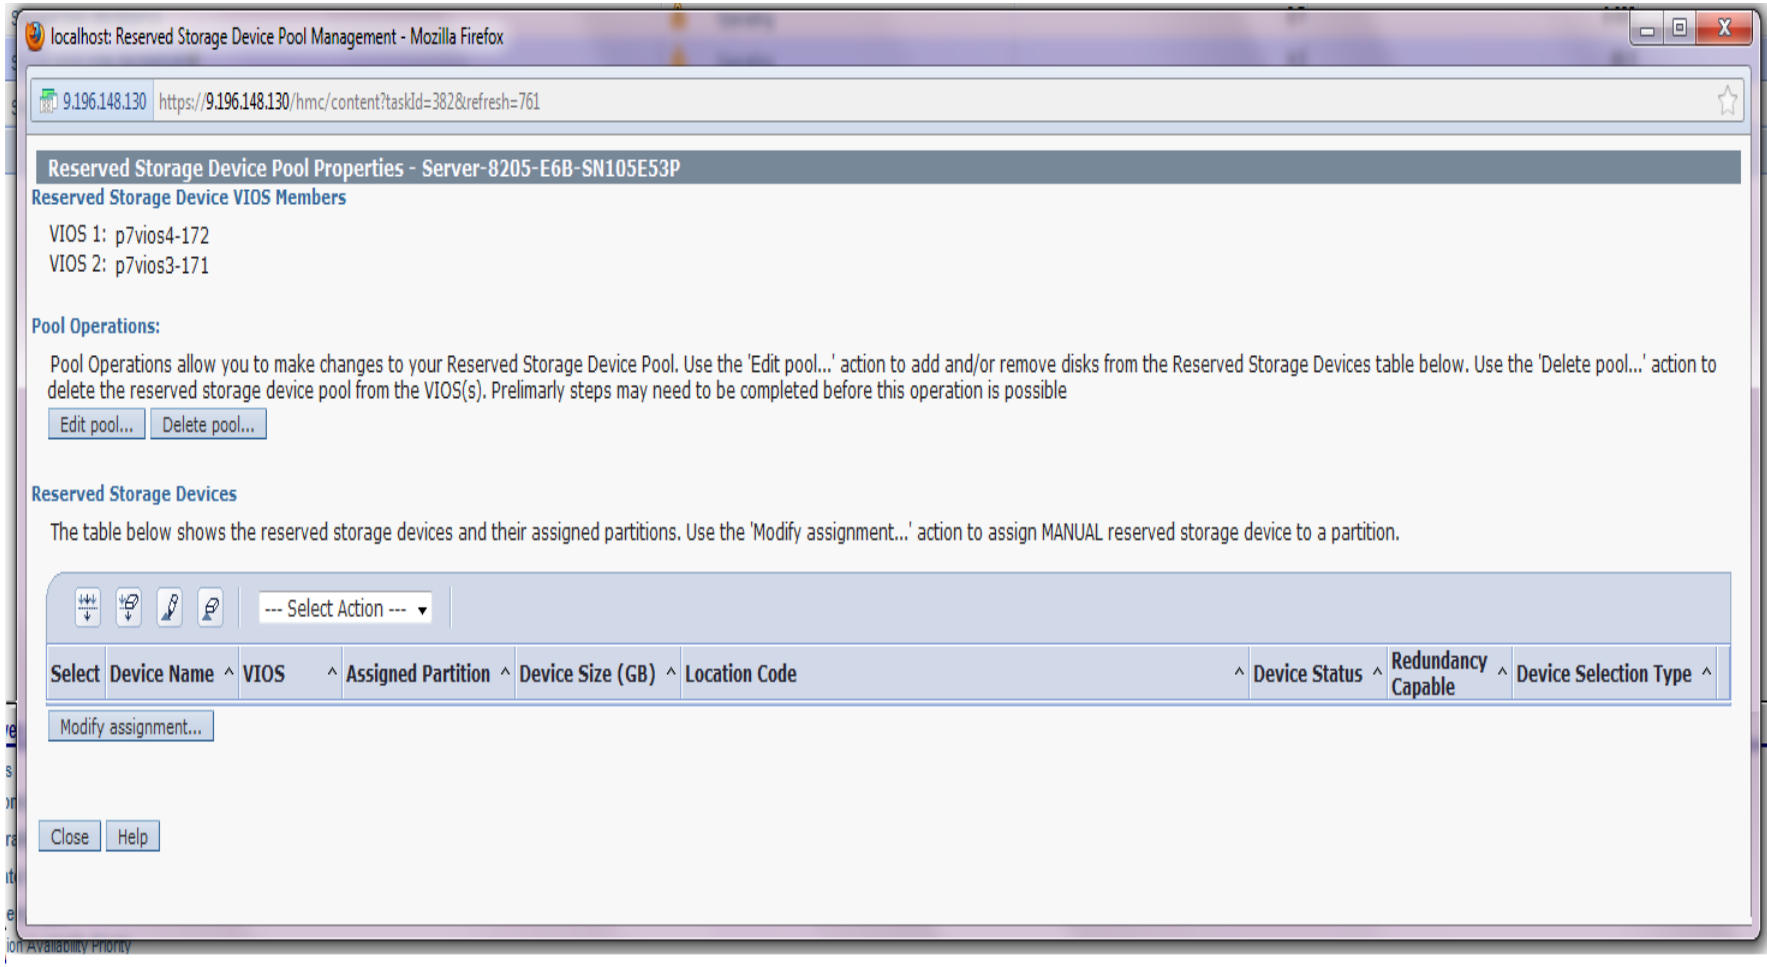
The screenshot displays the Hardware Management Console (HMC) interface. The main window shows a list of servers under the 'Servers' tab. The selected server is 'Server-8205-E6B-SN105E53P'. Below the server list, the 'Tasks' menu is expanded, showing various management options for the selected server.

Select	Name	Status	Available Processing Units	Available Memory (GB)	Reference Code
<input type="checkbox"/>	Server1-MMA-SN100ED60	Operating		2.6	4.25
<input type="checkbox"/>	Server2-MMA-SN100ED90	Operating		2	3.625
<input type="checkbox"/>	Server3-E4A-SN100DF11	Operating		1.9	2.8125
<input type="checkbox"/>	Server4-E4A-SN100DF21	Operating		2.7	0.625
<input checked="" type="checkbox"/>	Server-8205-E6B-SN105E53P	Operating		4.2	49.5
<input type="checkbox"/>	Server-8231-E2B-SN0676FP	Operating		2.3	17.25

Tasks: Server-8205-E6B-SN105E53P

- Properties
- Operations
- Configuration
  - Create Logical Partition
  - System Plans
    - Partition Availability Priority
    - View Workload Management Groups
    - Manage Custom Groups
  - Manage Partition Data
    - Manage System Profiles
  - Virtual Resources
    - Shared Processor Pool Management
    - Shared Memory Pool Management
    - Virtual Storage Management
    - Virtual Network Management
    - Reserved Storage Device Pool Management
- Connections
- Serviceability
- Hardware Information
- Capacity On Demand (CoD)
- Updates

# Suspend / Resume - RS Device Pool Creation



localhost: Reserved Storage Device Pool Management - Mozilla Firefox

9.196.148.130 https://9.196.148.130/hmc/content?taskId=382&refresh=761

### Reserved Storage Device Pool Properties - Server-8205-E6B-SN105E53P

#### Reserved Storage Device VIOS Members

VIOS 1: p7vios4-172  
VIOS 2: p7vios3-171





#### Pool Operations:

Pool Operations allow you to make changes to your Reserved Storage Device Pool. Use the 'Edit pool...' action to add and/or remove disks from the Reserved Storage Devices table below. Use the 'Delete pool...' action to delete the reserved storage device pool from the VIOS(s). Preliminary steps may need to be completed before this operation is possible

[Edit pool...](#) [Delete pool...](#)

#### Reserved Storage Devices

The table below shows the reserved storage devices and their assigned partitions. Use the 'Modify assignment...' action to assign MANUAL reserved storage device to a partition.





 --- Select Action ---

Select	Device Name ^	VIOS ^	Assigned Partition ^	Device Size (GB) ^	Location Code	Device Status ^	Redundancy Capable	Device Selection Type ^
<a href="#">Modify assignment...</a>								

[Close](#) [Help](#)

# Suspend / Resume - RS Device Pool Device Selection

samvednaHMC: Reserved Storage Device Pool Management - Mo...

https://9.126.134.58/hmc/content?taskId=33&refresh=59

### Reserved Storage Device Selection - HV4

To display the available reserved storage devices in the device lists, you must first select filter parameters and then click Refresh. You can list all available reserved storage devices by selecting All as the device type, or you can narrow your search by selecting a device type, maximum size, or minimum size.

Device Type:

Maximum Size (in GBs):

Minimum Size (in GBs):

Choose from the following list of devices. You can choose more than one reserved storage device to be added to the pool. After you have made your selections, select the OK button to assign the selected devices to the pool.

**VIOS 1 device list:**

Select	VIOS ^	Device Name ^	Device Size (GBs) ^	Redundancy Capable ^
<input type="checkbox"/>	vio1	hdisk2154	1.2	True
<input type="checkbox"/>	vio1	hdisk2204	1.2	False
<input type="checkbox"/>	vio1	hdisk2215	1.2	False
<input type="checkbox"/>	vio1	hdisk295	1.2	True
<input type="checkbox"/>	vio1	hdisk256	1.2	True
<input type="checkbox"/>	vio1	hdisk238	1.2	True
<input type="checkbox"/>	vio1	hdisk2160	1.2	False
<input type="checkbox"/>	vio1	hdisk2171	1.2	False
<input type="checkbox"/>	vio1	hdisk2182	1.2	False
<input type="checkbox"/>	vio1	hdisk2193	1.2	False

**VIOS 2 device list:**

Select	VIOS ^	Device Name ^	Device Size (GBs) ^	Redundancy Capable ^
<input type="checkbox"/>	vio2	hdisk2220	1.2	False

Done | 9.126.134.58 | 00:00

# Suspend / Resume - RS Device Pool Properties

localhost: Reserved Storage Device Pool Management - Mozilla Firefox

9.196.148.130 https://9.196.148.130/hmc/content?taskId=382&refresh=761

### Reserved Storage Device Pool Properties - Server-8205-E6B-SN105E53P

**Reserved Storage Device VIOS Members**

VIOS 1: p7vios4-172  
 VIOS 2: p7vios3-171

**Pool Operations:**

Pool Operations allow you to make changes to your Reserved Storage Device Pool. Use the 'Edit pool...' action to add and/or remove disks from the Reserved Storage Devices table below. Use the 'Delete pool...' action to delete the reserved storage device pool from the VIOS(s). Preliminary steps may need to be completed before this operation is possible

**Reserved Storage Devices**

The table below shows the reserved storage devices and their assigned partitions. Use the 'Modify assignment...' action to assign MANUAL reserved storage device to a partition.

Select	Device Name ^	VIOS ^	Assigned Partition ^	Device Size (GB) ^	Location Code ^	Device Status ^	Redundancy Capable ^	Device Selection Type ^
<input type="radio"/>	hdisk15 hdisk15	p7vios4-172 p7vios3-171	TCR-cts200(9)	10.0	U78AA.001.WZSG6SJ-P1-C6-T1-W201800A0B86E391A-LE000000000000 U78AA.001.WZSG6SJ-P1-C4-T1-W201800A0B86E391A-LE000000000000	Active Active	true	AUTO
<input type="radio"/>	hdisk16 hdisk16	p7vios4-172 p7vios3-171	AIXdemo-202(8)	10.0	U78AA.001.WZSG6SJ-P1-C6-T1-W201800A0B86E391A-LF000000000000 U78AA.001.WZSG6SJ-P1-C4-T1-W201800A0B86E391A-LF000000000000	Active Active	true	AUTO
<input type="radio"/>	hdisk17 hdisk17	p7vios4-172 p7vios3-171	AIXdemo-201(7)	10.0	U78AA.001.WZSG6SJ-P1-C6-T1-W201800A0B86E391A-L1000000000000 U78AA.001.WZSG6SJ-P1-C4-T1-W201800A0B86E391A-L1000000000000	Active Active	true	AUTO
<input type="radio"/>	hdisk14 hdisk14	p7vios4-172 p7vios3-171	AIXdemo-210(14)	10.0	U78AA.001.WZSG6SJ-P1-C6-T1-W201800A0B86E391A-LD000000000000 U78AA.001.WZSG6SJ-P1-C4-T1-W201800A0B86E391A-LD000000000000	Active Active	true	AUTO

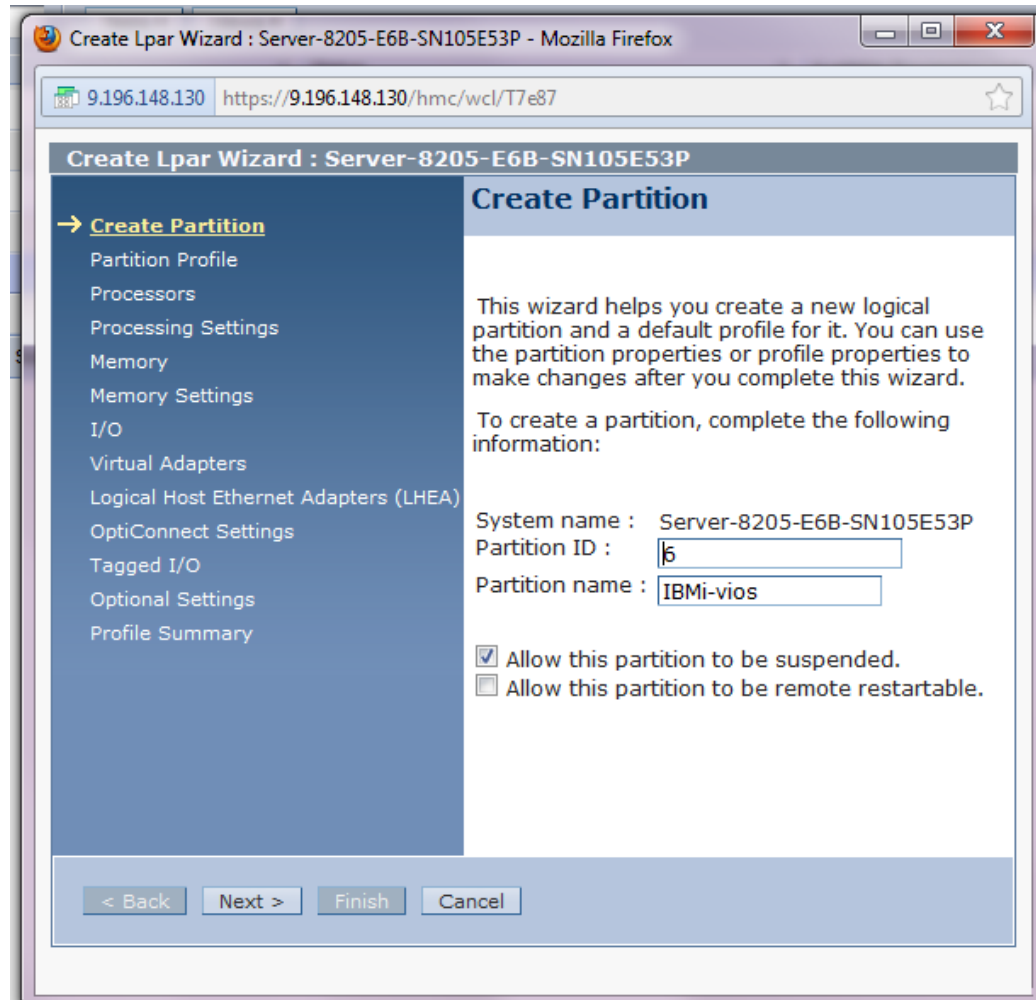
## Suspend / Resume – Setting a new partition to be suspend capable

### **User enabled the setting at create partition**

- At create partition,
  - Checkbox in GUI and optional in parameter in CLI.
  - Defaults to NOT suspend capable for no user input.
  - No restrictions/changes in create/add profile and modify profile.
  
- HMC action - At partition activation partition,
  - HMC validates for incompatible resource and settings if suspend capable is enabled.
  - If validation fails → activation fails.
  - User action: Fix the incompatibilities or disable suspend capable setting.



# Suspend / Resume - Creating a Suspend Capable Partition – GUI example

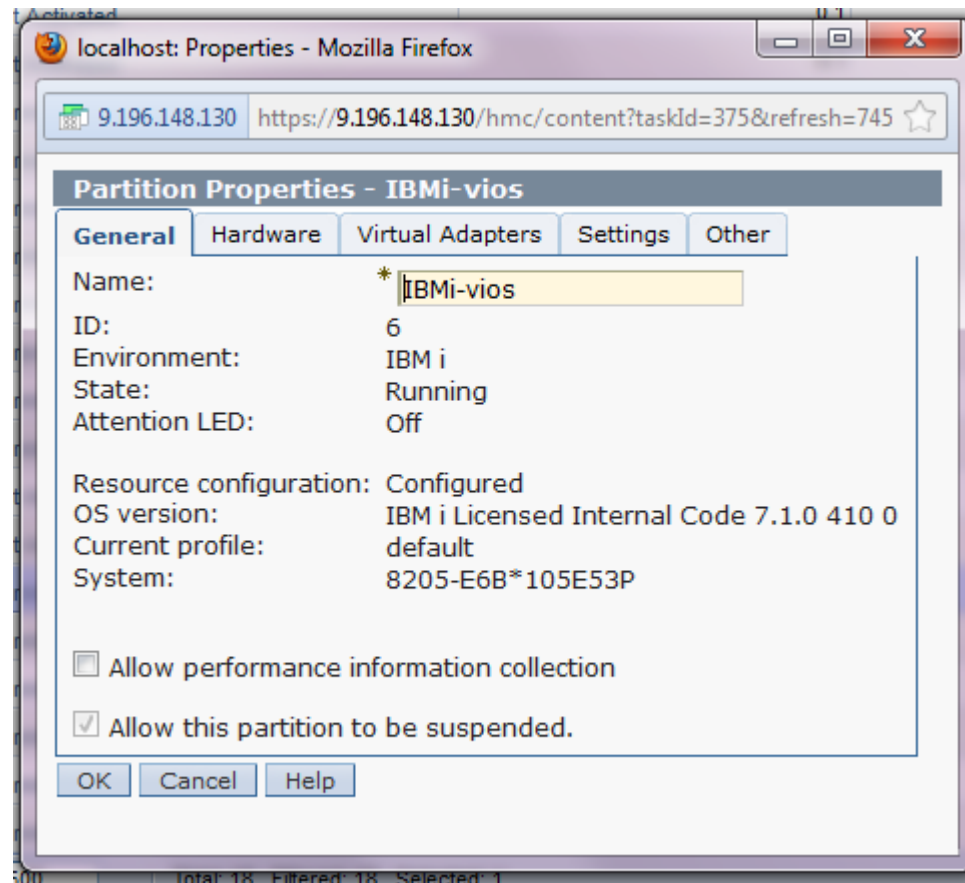


## Suspend / Resume – Changing an existing partition to be suspend capable

### **User enabled the setting on existing partition**

- On active partition
  - HMC validates for resource and settings restrictions.
  - HMC enables the setting only if above validations succeed
  
- On in-active partition
  - HMC enables the suspend capable setting. NO validation
  - HMC marks partition as non-bootable. This forces validation at HMC at the next activation (and prevents auto boot)

## Suspend / Resume - Setting Suspend Capability on an existing Partition



# Suspend / Resume – Partition suspension validation (HMC)

## HMC high level steps

*Explicit HMC user interfaces are available for validation only.*

- *Validation is also part of actual suspend operation.*
  
- Checks if CEC is suspend capable and partition is active has suspend capable enabled.
- Checks for max number concurrent suspend operations in progress.
- Checks for presence of reserved storage device pool, and at least one active VIOS with Resource Monitoring and Control (RMC) connection.
- Checks for presence of restricted resources, restricted settings at the partition.
- If partition already has device, checks the size requirement. If partition doesn't have device, checks for availability of suitable device in pool.
- Checks with OS (using RMC) if it is capable of suspend and if it is ready for suspend.

# Suspend / Resume – Partition suspension (HMC)

## HMC high level steps

- Performs validation (as discussed in previous slide).
- Associates the storage device to the partition (if not already associated).
- Initiates the suspend process.
- Keeps note of progress (at both PHYP and in HMC) based on PHYP async messages.
- Displays the progress information to the user.
  - GUI – progress bar with %complete
  - CLI – Total and remaining MB with lssyscfg command
- User has option to stop the suspend operation.
- If the suspend operation fails, HMC will auto recover from the operation.
- If HMC auto recover fails, user can initiate recover explicitly.
- After partition is suspended, HMC does cleanup operation.
- After HMC cleanup, the **partition power state** is changed to **Suspended**.

# Suspend / Resume – Partition suspension (HMC)

## Additional details

- Progress States (visible in UI)
  - Starting
  - Validating
  - Saving HMC data
  - Saving partition data
  - Completing
  
- All HMC data transfer happens via PHYP through VIOS to the storage device using Virtual Asynchronous Services Interface (VASI) channel.
- HMC cleanup activity: This involves removing all the 1-1 virtual server adapters from VIOSeS and updating their last activated profiles.
- User initiated cancel of suspend operation is accepted until PHYP completes its work.

# Suspend / Resume – Suspend menu option

localhost: Hardware Management Console Workplace (V7R7.7.0.0) - Mozilla Firefox  
 https://9.196.148.130/hmc/connects/mainuiFrameset.jsp

## Hardware Management Console

hscroot | Help | Logoff

Systems Management > Servers > Server-8205-E6B-SN105E53P

Select	Name	ID	Status	Processing Units	Memory (GB)	Active Profile	Environment	Reference Code
<input type="checkbox"/>	AXG1-cts180	16	Not Activated		0.1	1.25 default	ADX or Linux	00000000
<input type="checkbox"/>	AXG1-cts181	17	Not Activated		0.1	1.25 default	ADX or Linux	00000000
<input type="checkbox"/>	AXdemo-201	7	Running		0.2	2 AMS	ADX or Linux	
<input type="checkbox"/>	AXdemo-202	8	Running		0.2	2 AMS	ADX or Linux	
<input type="checkbox"/>	AXdemo-210	14	Running		0.4	2 AMS	ADX or Linux	
<input type="checkbox"/>	AX-ha61a-206	10	Running		0.2	2 default	ADX or Linux	
<input type="checkbox"/>	AX-ha61b-207	11	Running		0.2	2 default	ADX or Linux	
<input type="checkbox"/>	AX-ha71a-208	12	Running		0.2	3 default	ADX or Linux	
<input type="checkbox"/>	AX-ha71b-209	13	Running		0.2	3 default	ADX or Linux	
<input type="checkbox"/>	Dir3-152	15	Running		0.4	8 default	ADX or Linux	
<input type="checkbox"/>	IBM-6	5	Not Activated		0.1	4 default	IBM i	00000000
<input type="checkbox"/>	IBM-7	4	Not Activated		0.1	4 default	IBM i	00000000
<input checked="" type="checkbox"/>	IBM-vios-6	6	Running		0.1	4 default	IBM i	00000000
<input type="checkbox"/>	IBM-vios-L1	18	Running		0.1	4 LPM	IBM i	00000000
<input type="checkbox"/>	p7vios3-17				0.4	4 default	Virtual I/O Server	
<input type="checkbox"/>	p7vios4-17				0.4	4 default	Virtual I/O Server	
<input type="checkbox"/>	p7vios5-17				0.2	4 default	Virtual I/O Server	
<input type="checkbox"/>	TCR-cts200				0.2	6 AMS	ADX or Linux	

Max Page Size: 500 | Page 18 | Filtered: 18 | Selected: 1

# Suspend / Resume – Suspend validation and OS confirmation

The screenshot shows the Hardware Management Console interface. A modal dialog titled "Partition Suspend/Resume - IBMi-vios" is open, displaying the following information:

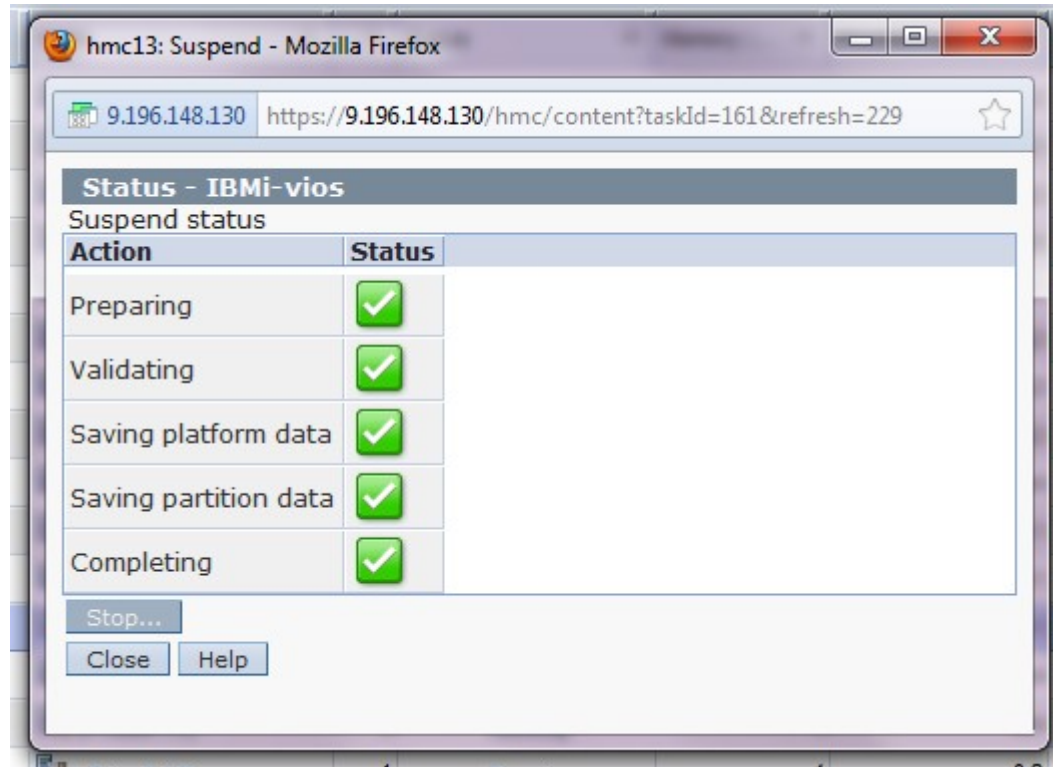
- System name:** Server-8205-E6B-SN105E53P
- Partition name:** IBMi-vios
- Message:** "If supported, a second VIOS can be used to provide a redundant path and higher availability to the device"
- VIOS1:** p7vios4-172
- VIOS2:** p7vios3-171
- Wait Time (minutes):** 5
- Detail Level:** 1

Buttons at the bottom of the dialog include "Validate", "Suspend", "Cancel", and "Help".

Select	Name	ID	Status	Processing Units
<input type="checkbox"/>	ADX61-cts180		16 Not Activated	0.1
<input type="checkbox"/>	ADX71-cts180			0.1
<input type="checkbox"/>	ADXdemo-20			0.2
<input type="checkbox"/>	ADXdemo-20			0.2
<input type="checkbox"/>	ADXdemo-21			0.4
<input type="checkbox"/>	ADX-ha61a-2			0.2
<input type="checkbox"/>	ADX-ha61b-2			0.2
<input type="checkbox"/>	ADX-ha71a-2			0.2
<input type="checkbox"/>	ADX-ha71b-2			0.2
<input type="checkbox"/>	Dir63-152			0.4
<input type="checkbox"/>	IBMI-6			0.1
<input type="checkbox"/>	IBMI-7			0.1
<input checked="" type="checkbox"/>	IBMI-vios			0.1
<input type="checkbox"/>	IBMI-vios-LP			0.1
<input type="checkbox"/>	p7vios3-171			0.4
<input type="checkbox"/>	p7vios4-172	3	Running	0.4
<input type="checkbox"/>	p7vios5-173	1	Running	0.2
<input type="checkbox"/>	TCR-cts200	9	Running	0.2



# Suspend / Resume – Suspend status



# Suspend / Resume – Partition resumption validation (HMC)

## HMC high level steps

*Explicit HMC user interfaces are available for validation only.*

- *Validation is also part of actual resume operation.*
  
- User to specify the options if partition's old (i.e. at suspend time) mpio and vlan bridge settings have to be reconfigured exactly or reconfigured as much as possible.
- Checks for presence of reserved storage pool, and at least one active VIOS with RMC connection.
  
- Reads the partition configuration data from the storage device and checks
  - Partition compatibility (ex: LMB size, proc compatibility etc)
  - If all the virtual io adapters can be re-configured (VFC, VETH, VSCSI)
  - If (proc, mem) type are supported and the quantity of resources for the partitions can be re-allocated
- If validation fails, the error information is displayed to the user. User can decide on appropriate corrective action.

# Suspend / Resume – Partition resumption (HMC)

## HMC high level steps

- Performs validation (as shown in previous slide).
- Initiates the resume process.
- Keeps note of progress (at both PHYP and in HMC) based on PHYP async messages.
- Displays the progress information to the user.
  - GUI – progress bar with %complete
  - CLI – Total and remaining MB with `lssyscfg` command
- User has option to cancel the resume operation.
- If the resume operation fails, HMC will auto recover from the operation.
- If HMC auto recover fails, user can initiate recover explicitly.
- After partition is resumed, the ***partition power state*** is changed to ***Running***.

# Suspend / Resume – Partition resumption validation (HMC)

## Additional details

- Progress States (visible in UI)
  - Starting
  - Reading HMC data
  - Validating
  - Restoring partition configuration
  - Reading partition data
  - Completing
  
- HMC resource allocation activity: This involves allocating the processor, memory resources and re-configuring the partition's virtual adapters.
- The VIOS last activated profiles are updated after virtual server adapters are created.
- User initiated cancel of suspend operation is accepted until PHYP completes its work.
- The storage device is released (only if NOT AMS partition) after resume is complete.

# Suspend / Resume – Resume menu option

localhost: Hardware Management Console Workplace (V7R7.7.0.0) - Mozilla Firefox

9.196.148.130 https://9.196.148.130/hmc/connects/mainuiFrameset.jsp

## Hardware Management Console

hscroot | Help | Logoff

Systems Management > Servers > Server-8205-E6B-SN105E53P

Filter [ ] Tasks Views

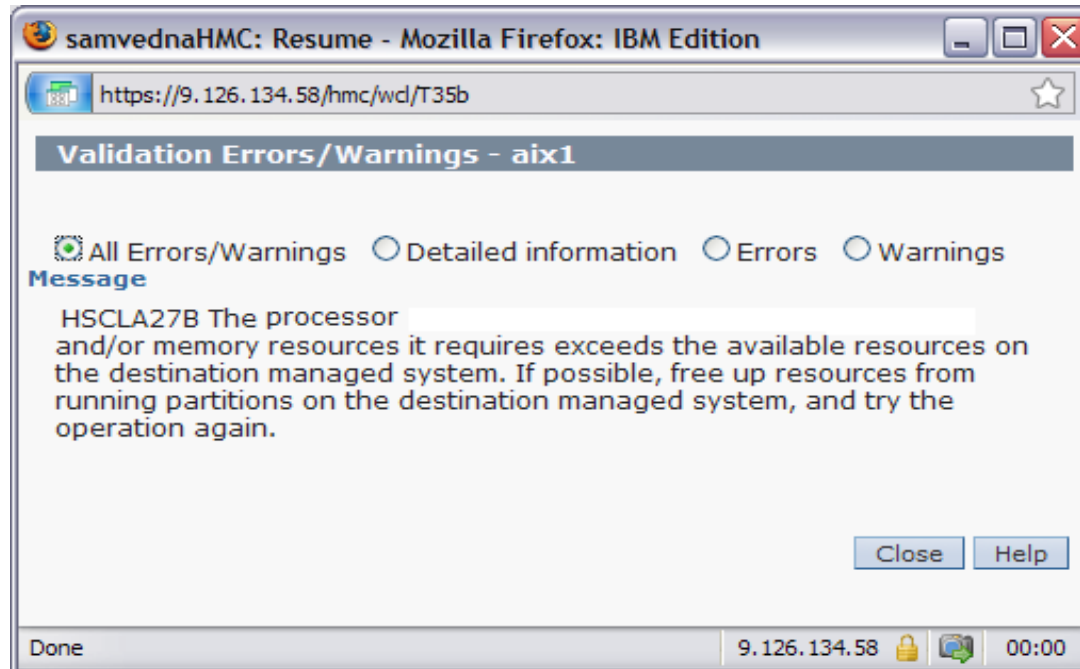
Select	Name	ID	Status	Processing Units	Memory (GB)	Active Profile	Environment	Reference Code
<input type="checkbox"/>	ADX61-cts180	16	Not Activated		0.1	1.25 default	AIX or Linux	00000000
<input type="checkbox"/>	ADX71-cts181	17	Not Activated		0.1	1.25 default	AIX or Linux	00000000
<input type="checkbox"/>	ADXdemo-201	7	Running		0.2	2 AMS	AIX or Linux	
<input type="checkbox"/>	ADXdemo-202	8	Running		0.2	2 AMS	AIX or Linux	
<input type="checkbox"/>	ADXdemo-210	14	Running		0.4	2 AMS	AIX or Linux	
<input type="checkbox"/>	ADX-ha61a-206	10	Running		0.2	2 default	AIX or Linux	
<input type="checkbox"/>	ADX-ha61b-207	11	Running		0.2	2 default	AIX or Linux	
<input type="checkbox"/>	ADX-ha71a-208	12	Running		0.2	3 default	AIX or Linux	
<input type="checkbox"/>	ADX-ha71b-209	13	Running		0.2	3 default	AIX or Linux	
<input type="checkbox"/>	Dir63-152	15	Running		0.4	8 default	AIX or Linux	
<input type="checkbox"/>	IBMI-6	5	Not Activated		0.1	4 default	IBM i	00000000
<input type="checkbox"/>	IBMI-7	4	Not Activated		0.1	4 default	IBM i	00000000
<input checked="" type="checkbox"/>	IBMI-vios-20	6	Running		0.1	4 default	IBM i	00000000
<input type="checkbox"/>	IBMI-vios-L	18	Running		0.1	4 LPM	IBM i	00000000
<input type="checkbox"/>	p7vios3-17				0.4	4 default	Virtual I/O Server	
<input type="checkbox"/>	p7vios4-17				0.4	4 default	Virtual I/O Server	
<input type="checkbox"/>	p7vios5-17				0.2	4 default	Virtual I/O Server	
<input type="checkbox"/>	TCR-cts200				0.2	6 AMS	AIX or Linux	

Max Page Size: 500 | 18 Filtered: 18 Selected: 1

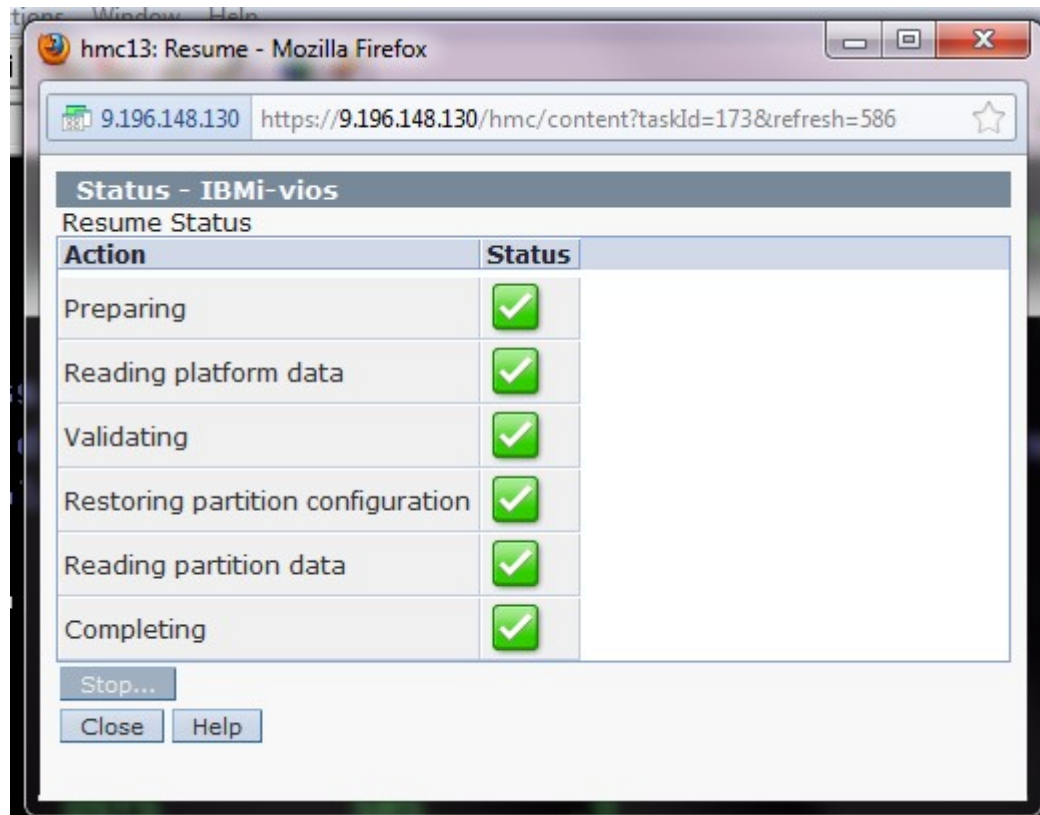
Context menu for IBMI-vios-20:

- Properties
- Change Default Profile
- Operations
  - Restart
  - Shut Down
  - Deactivate Attention LED
  - Schedule Operations
  - Suspend Operations
  - Recover
- Configuration
- Hardware Information
- Dynamic Logical Partitioning
- Serviceability

# Suspend / Resume – Resume validation error example



# Suspend / Resume – Resume status



The screenshot shows a Mozilla Firefox browser window titled 'hmc13: Resume - Mozilla Firefox'. The address bar displays '9.196.148.130 https://9.196.148.130/hmc/content?taskId=173&refresh=586'. The main content area is titled 'Status - IBMi-vios' and shows the 'Resume Status' of a partition. A table lists the actions and their status, all of which are completed with green checkmarks.

Action	Status
Preparing	✓
Reading platform data	✓
Validating	✓
Restoring partition configuration	✓
Reading partition data	✓
Completing	✓

Below the table, there are three buttons: 'Stop...', 'Close', and 'Help'.

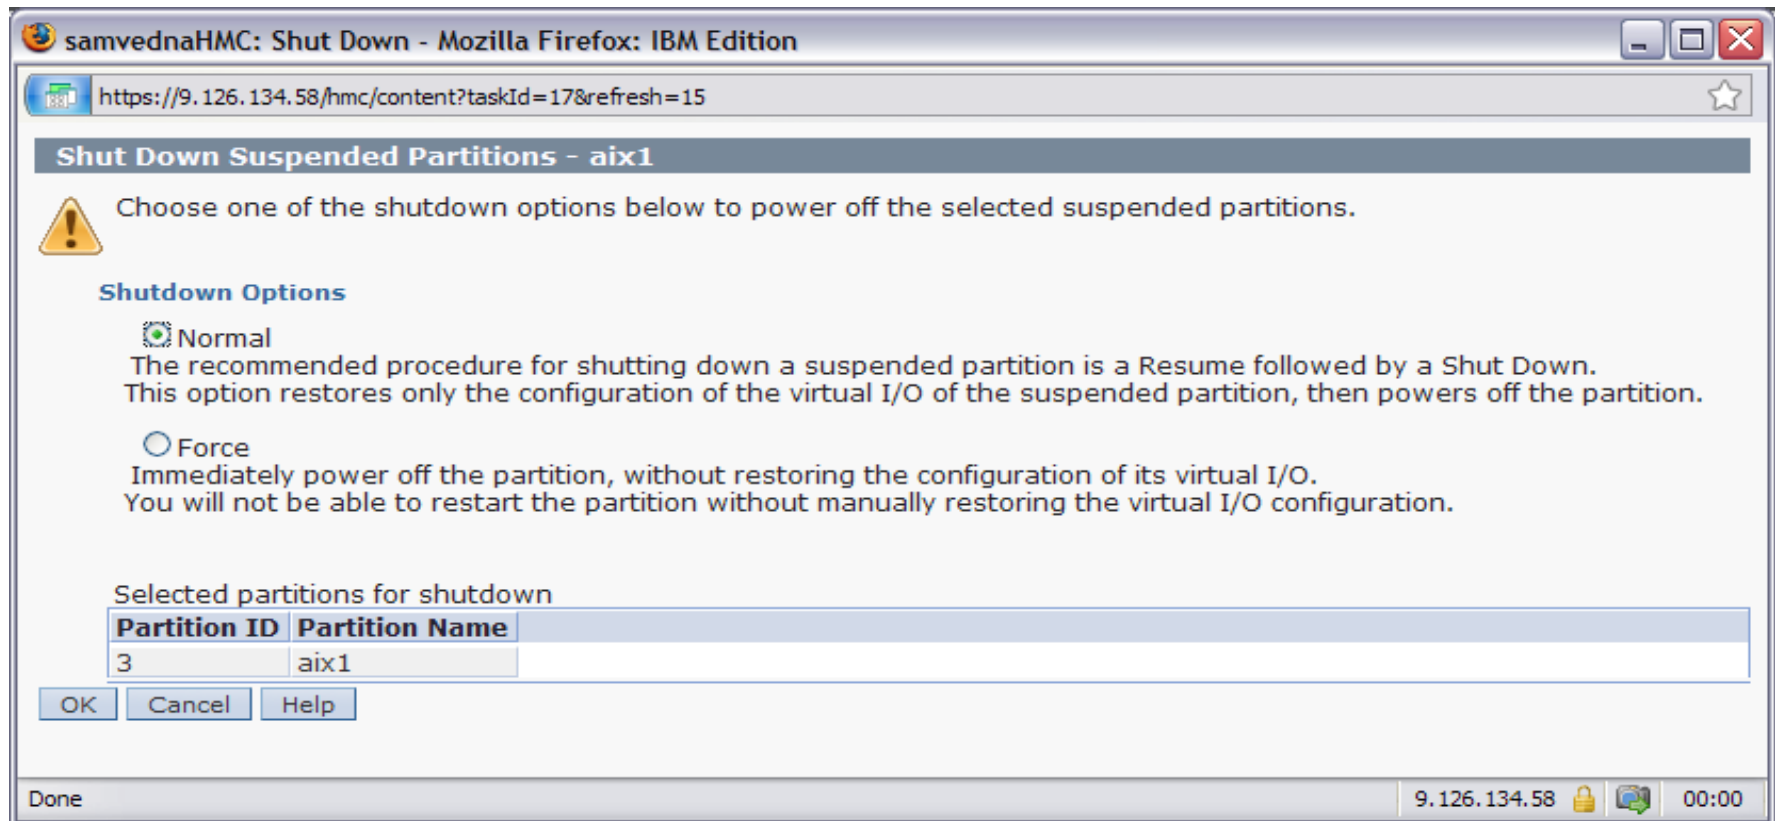
## Suspend / Resume – Shutdown of a suspended partition (HMC)

### HMC high level steps

- *Shutdown of a suspended partition is NOT generally recommended.*
- *HMC reconfigures all virtual server adapters at shutdown of suspended partition (hence follows resume flow partially).*
- *This ensures subsequent activation of the partition with last activated profile succeeds.*
- A force shutdown option is available if virtual server adapter reconfiguration faces unrecoverable error.
- If the partition is NOT AMS, the storage device is released from the partition.



# Suspend / Resume – Shutdown of a suspended partition (HMC)



The screenshot shows a web browser window titled "samvednaHMC: Shut Down - Mozilla Firefox: IBM Edition". The address bar contains the URL "https://9.126.134.58/hmc/content?taskId=17&refresh=15". The main content area is titled "Shut Down Suspended Partitions - aix1". A warning icon is present next to the instruction: "Choose one of the shutdown options below to power off the selected suspended partitions." Under the heading "Shutdown Options", there are two radio buttons: "Normal" (selected) and "Force". The "Normal" option is described as the recommended procedure, restoring the virtual I/O configuration before powering off. The "Force" option is described as immediately powering off the partition without restoring the configuration. Below the options, a table lists the selected partitions for shutdown:

Partition ID	Partition Name
3	aix1

At the bottom of the dialog are buttons for "OK", "Cancel", and "Help". The status bar at the bottom of the browser window shows "Done", the IP address "9.126.134.58", a lock icon, a chat icon, and the time "00:00".

## Suspend / Resume – Recovery of a failed suspend or resume

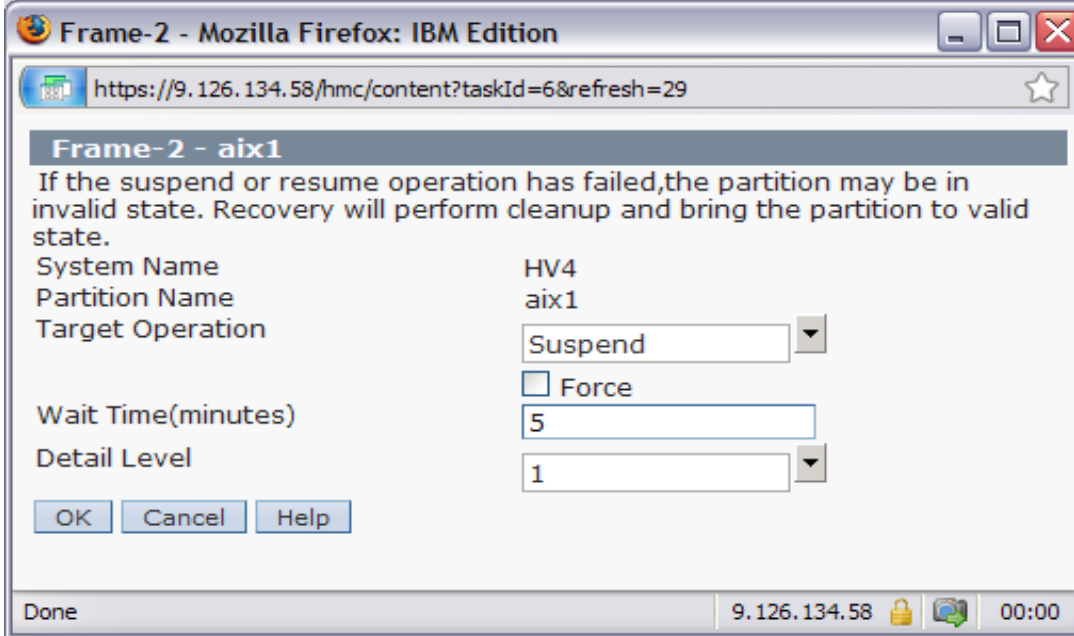
When can User issue recover:

- Suspend/Resume is taking long time and user ends the operation abruptly.
- User is not able to abort suspend/resume successfully.
- Initiating a Suspend/Resume operation has resulted in an extended error indicating partition's hibernation state is not valid.
- User can recover operation using the same HMC or different HMC.

What does HMC recover:

- Determines the last successful step in the operation from progress data. (remember that both suspend and resume store the operation progress in both HMC and PHYP)
- Based on the last successful step, HMC will try to
  - either proceed further to complete the operation
  - or rollback the operation.
- If no progress data is available, user to recover with force option. HMC will recover as much as possible.

# Suspend / Resume – Recovery of a failed suspend or resume



The screenshot shows a Mozilla Firefox browser window titled "Frame-2 - Mozilla Firefox: IBM Edition". The address bar contains the URL "https://9.126.134.58/hmc/content?taskId=6&refresh=29". The page content is titled "Frame-2 - aix1" and contains the following text: "If the suspend or resume operation has failed, the partition may be in invalid state. Recovery will perform cleanup and bring the partition to valid state." Below this text are several configuration fields:

System Name	HV4
Partition Name	aix1
Target Operation	Suspend
	<input type="checkbox"/> Force
Wait Time(minutes)	5
Detail Level	1

At the bottom of the form are three buttons: "OK", "Cancel", and "Help". The status bar at the bottom of the browser window shows "Done", the IP address "9.126.134.58", a lock icon, a globe icon, and the time "00:00".

## Suspend / Resume – Configuration Gotcha's

- Configuration, management and usage for suspend resume is common across both IBM i and AIX, however there is some information missing in the current online documentation
- All disk units must have the reserve policy set to no\_reserve, you will receive the following error when you attempt to suspend the partition:

```
HSCLA27C The operation to get the physical device location for adapter U9117.MMB.102709P-V20-C6 on the virtual I/O server partition VIOS1-Dilling has failed. The partition command is: migmgr -f get_adapter -t vscsi -s U9117.MMB.102709P-V20-C6 -d 1 The partition standard error is: child process returned error
HSCLA9C1 The partition suspend or resume operation has stopped unexpectedly. If necessary, perform a suspend or resume recovery operation for the partition.
```

- When suspending a partition the HMC does not pre-validate storage pool capabilities to ensure that there is enough space to hold the suspended state information, so the suspend will fail part way through.
- By default the HMC will attempt to provide a second VIOS server to provide a redundant path and higher availability, you will need to ensure that there is the secondary VIOS server is configured correctly to support the suspension of your partition if you wish to utilize this capability. Otherwise you will receive a message similar to :

There is no redundant device available in the reserved storage device pool that can be used by this partition. This partition requires a device with a size of at least 6563 MB. Add a device of at least that size to the reserved storage device pool, then try the operation again.

# Agenda

*What we will talk about today:*

**1 Suspend / Resume**

**2 What implications are there on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Resuming an IBM I partition**

**5 Migrating Aix LPAR aware**



## Suspend / Resume and Power-HA

- Currently Power-HA and Suspend/Resume are independent capabilities and are unaware of each other, this means:
  - On Suspension
    - When you suspend a high available partition it will automatically fail over
    - If you do not wish for a failover to occur then you must manually disable High Availability before suspending
  - On Resuming
    - When you resume a highly available suspend partition then you need to fail it back over to become the active primary partition
    - If you manually disabled the fail over, you must manually re-enable the partition to ensure that fail over to a standby environment occurs going forward

## Suspend / Resume and TCP/IP (or APPC) communications

- When you suspend a partition that is configured for communication then you need to be aware of the following:
  - If there is active communication ongoing, at the time of suspension, then this will be interrupted part way through. Depending on what is ongoing there may be a business impact
  - At resumption, it is likely that the connection with the server will need to be re-established before inter-communication can begin again. Depending on the application, it is likely that there may need to be some clean-up undertaken when communication is reestablished
  
- If your system is acting as a server in a communication process:
  - If there is active communication ongoing, at the time of suspension, then this will be interrupted part way through. Depending on what is ongoing there may be a business impact
  - At resumption, it is likely that the connection with the communication initiator will need to be re-established before inter-communication can begin again. It is likely that there will be some cleanup activity required

## Suspend / Resume and TCP/IP (or APPC) communications

- Software likely to be impacted:
  - MQ: Impact will be dependent on the MQ configuration
  - WebSphere Application Server: Application dependent
  - Sockets programs: Application dependent
  - APPC programs: Application dependent
  - FTP: Will require manual cleanup, may also impact automated processes




## Suspend / Resume and Personal Communications 5250 sessions

- Suspending an IBM i partition persists the active 5250 sessions
- What happens when a partition is resumed is dependant on a number of factors including:
  - Whether the 5250 session is terminated whilst the partition is suspended
  - Whether the 5250 session have named terminals specified in the profile
- Unsurprisingly the majority of the behavior can be found when you consider what happens when you lose connectivity due to a network issue

Note: Our investigation was undertaken with IBM Personal Communications as the terminal emulator, the findings may differ when utilising other manufactires emulators

# Suspend / Resume and Personal Communications 5250 sessions

- If you leave a 5250 session active for the duration of the suspension then the 5250 session should automatically reconnect to the job
- During testing we found that the user typically would need to press a key (any key) for their 5250 session to redisplay their screen




## Suspend / Resume and Personal Communications 5250 sessions

- If you shutdown a 5250 session during a suspension of a partition, when you resume the partition what happens on resume is dependant on:
  - How you configured your 5250 session in personal communications
  - Your system values for system control
    - QSEVRCYACN – Device I/O error action
    - QDSCJOBITV – Time interval before disconnected job ends
    - .....

# Suspend / Resume and Personal Communications 5250 sessions

- If your 5250 session definition does not have a Workstation ID defined, when you resume the partition a new job is started for your session and the original session will follow the behavior specified in your system values

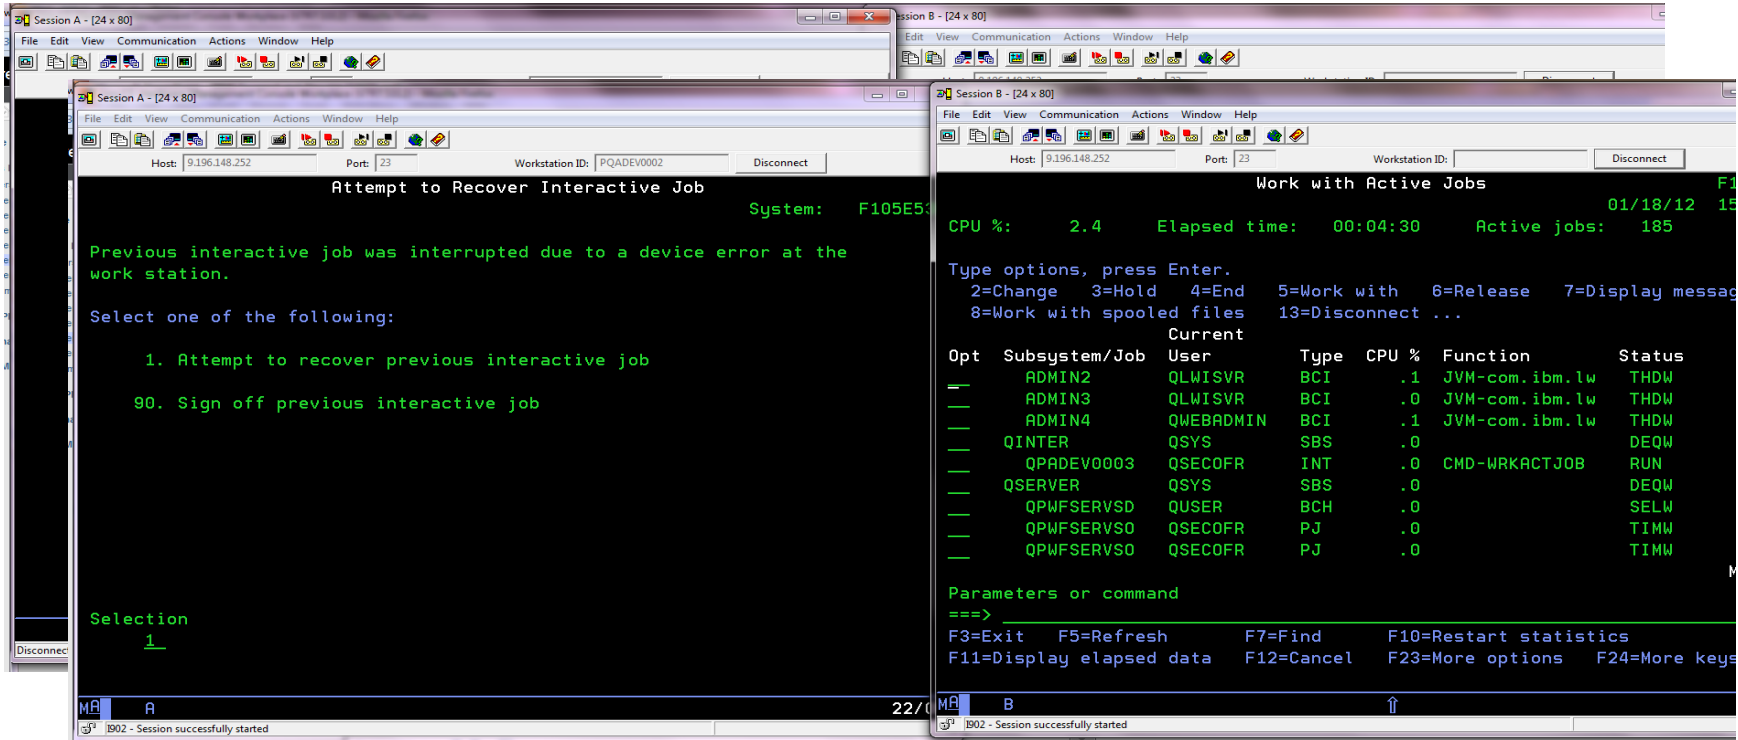


## Suspend / Resume and Personal Communications 5250 sessions

- If your 5250 session definition does not have a Workstation ID defined, when you resume the partition a new job is started for your session and the original session will follow the behavior specified in your system values
  
- For discussion purposes let's assume the following:
  - QDEVRCVACT is set to \*DSCJOB
  
  - QDSCJOBTV is set to 30 (minutes)

# Suspend / Resume and Personal Communications 5250 sessions

- If the 5250 session definition does have a Workstation ID defined and the 5250 session is reconnected after the time limit, then the original job would have been ended and the a new job within the time limits specified by QDSCJOBITV then your user would have an option to reconnect to the previous job



## Suspend / Resume and Personal Communications 5250 sessions

- If your 5250 session definition does have a Workstation ID defined, and the 5250 session is connected after the QDSCJOBTV time limit then
  - The original job would be ended by the Operating System
  - The a new job would be started for the 5250 session
  
- What happens if you have other values specified in the QDEVRCVACT system value?
  - The same thing that would happened when you have a communications issues with a 5250 session
  
  - Refer to the [V7R1 Infocenter](#) for additional information

## Suspend / Resume and Personal Communications 5250 sessions

- If your 5250 session definition does have a Workstation ID defined, then what happens is dependant on your system values and how quickly the 5250 session is restarted in relation to the partition resumption.
- Assuming that you disconnect job with I/O errors and the 5250 session is restarted in the time limit, then the 5250 session would reconnect otherwise a new job will be started.



# Suspend / Resume and 5250 applications

- So that's great, but what about if my users are running applications on their screens?
- If you reconnect to a disconnected job then you typically get an escape message, the behavior is application dependant

Session A - [24 x 80]

```

File Edit View Communication Actions Window Help
Host: 9.196.148.252 Port: 23 Workstation ID: PQADEV0002 Disconnect
Display Program Messages
Job 019973/QSECOFR/PQADEV0002 started on 01/18/12 at 13:06:02 in subsystem Q
Session or device error occurred in file WDWDDEM02 (C G D F).

Type reply, press Enter.
Reply . . . g

F3=Exit F12=Cancel
  
```

Session B - [24 x 80]

```

Edit View Communication Actions Window Help
Host: 9.196.148.252 Port: 23 Workstation ID: Disconnect
Work with Active Jobs
CPU %: 2.4 Elapsed time: 00:06:07 Active jobs: 185
Type options, press Enter.
1=Change 2=Change 3=Hold 4=End 5=Work with 6=Release 7=Display message
8=Work with spooled files 13=Disconnect ...

Current
Host Subsystem/Job User Type CPU % Function Status
-----
ADMIN2 QLWISVR BCI .1 JVM-com.ibm.lw THDW
ADMIN3 QLWISVR BCI .1 JVM-com.ibm.lw THDW
ADMIN4 QWEBADMIN BCI .1 JVM-com.ibm.lw THDW
QINTER QSYS SBS .0 DEQW
PQADEV0002 QSECOFR INT .0 PGM-WDWDDEM02 DSPW
PQADEV0003 QSECOFR INT .0 CMD-WRKACTJOB RUN
QSERVER QSYS SBS .0 DEQW
QPWFSEVRSD QUSER BCH .0 SELW
QPWFSEVRVSO QSECOFR PJ .0 TIMW

Parameters or command
=>
F3=Exit F5=Refresh F7=Find F10=Restart statistics
F11=Display elapsed data F12=Cancel F23=More options F24=More keys
  
```

# Suspend / Resume and IBM i interactive PASE applications

- We have experienced a mixed set of states, on a few occasions we have found that our PASE program ended abnormally, however for the majority of tests we have reconnected directly into the QP2TERM environment and our application is still running

The screenshot displays two overlapping windows from an IBM i management console. The left window, titled 'hmc13: Resume - Mozilla Firefox', shows a 'Resume Status' window for the 'IBM-i-vios' profile. It contains a table with the following data:

Action	Status
Preparing	✓
Reading platform data	✓
Validating	✓
Restoring partition configuration	✓
Reading partition data	✓
Completing	✓

The right window, titled 'Session B - [27 x 132]', shows a terminal session with the following output:

```

/00OpenSys/usr/bin/-sh
#
> cd /home/BLIGHTJ
#
> ./t_sigwaitinfo
./t_sigwaitinfo: PID is 765
./t_sigwaitinfo: signals blocked
  
```

At the bottom of the terminal window, a status bar indicates 'ib902 - Session successfully started'.

## Suspend / Resume and Open Access

- Open Access is a capability that allows the programmatic emulation of the RPG File I/O. As there are many potential ways in which Open Access can be used it is difficult to say what effect that Suspending and resuming a partition would have on behavior.
  
- Two scenarios spring to mind:
  - Using Open Access with RPG code being called from a J2E, web facing PHP or communications program:  
Very environment and application specific. If database or printer I/O is being manipulated then there most like be a set of clean up required here
  - Using Open Access with RPG code called from a 5250 session  
This is most likely database or printer I/O manipulation. This should be covered in standard procedures used to handle disconnection 5250 sessions
  
- Ideally your Open Access handlers would be Suspend/Resume aware

# Agenda

*What we will talk about today:*

**1 Suspend / Resume**

**2 What implications are there on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Resuming an IBM I partition**

**5 Migrating Aix LPAR aware**



## Allowing or preventing a Suspend on IBM i

- During a partition suspension, or a standalone partition suspension validation, from a management console (HMC or SDMC) there is a communication between the hypervisor and the IBM i partition to ensure that partition is in a state that it can be suspended
- What does this mean? In a state that can be suspended? What are you on about...
  - All systems at some point will be undertaking some business critical function that should not, or can not, interrupted at that point in time. If we were to suspend the partition at this time then there would be a dire set of business consequences that we wish to avoid at all times.
- The partition suspension validation utilizes a capability on the IBM i partition that can be programmatically exploited to ensure that the business critical functionality will be completed in a time window, has been is quiesced or can block the suspension of the the IBM i partition

## Allowing or preventing a Suspend on IBM i

- This capability is surfaced on the IBM i by a Work Management exit added with PTF SI42815 called QIBM\_QWC\_SUSPEND

```

Session D - [24 x 80]
File Edit View Communication Actions Window Help
Host: ibmi-7vios.aixncc.uk.ibm.c Port: 23 Workstation ID: Disconnect
Work with Registration Information
Type options, press Enter.
5=Display exit point      8=Work with exit programs

Exit
Point
Format
Registered
Text

Opt  Exit Point      Format      Registered  Text
--  -
  _  QIBM_QWC_JOBITPPGM  JITP0100   *YES        Job interrupt program exit po
  _  QIBM_QWC_PRERESTRICT  PRER0100  *YES        ENDSYS ENDSBS *ALL pre restri
  _  QIBM_QWC_PWRDWSYS    PWRD0100  *YES        Prepower down system exit poi
  _  QIBM_QWC_PWRDWSYS    PWRD0200  *YES        Prepower down system exit poi
  _  QIBM_QWC_QSTGLOWACN  STGL0100  *YES        Auxiliary storage lower limit
  _  QIBM_QWC_RESUME     RSMS0100  *YES        Resume system
  _  QIBM_QWC_SUSPEND    SSPS0100  *YES        Suspend system
  _  QIBM_QWT_JOBNOTIFY  NTFY0100  *YES        Job Notification
  _  QIBM_QWT_PREATTPGMS  ATTN0100  *YES        Preattention program exit poi
  _  QIBM_QWT_SYSREQPGMS  SREQ0100  *YES        Presystem request pgm exit po
  _  QIBM_QYIV_INVGTRSRV  GTRS0100  *YES        Inventory gathering services
                                     More...

Command
===>
F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel

MA  D  MW  09/003
IBM i - Session successfully started
Smarter Commerce

```

## Allowing or preventing a Suspend on IBM i

- At this moment there is no publically available document on how you would code to this exit point, however we will discuss what we need to think about when utilising this exit point

## Allowing or preventing a Suspend on IBM i

- Whilst at first glance this seems like a pretty straight forward exit point, there are several things to be aware of:
  - The exit point calls a standalone program
  - This program runs in a standalone job
- What are the implications of this?
  - The prime implication is that it is responsibility of the exit point implementer to ensure that the system (or application's that are suspend/resume aware) are in a state to allow the partition to be suspended, however as the program runs in its own job there is a burden on the exit point developer to have some type of inter process communication, or interaction, with other jobs and processes on the system and to respond back to the operating system in a timely manner
- This could potentially mean retro-fitting the required Suspend/Resume awareness into existing applications and environments



## Allowing or preventing a Suspend on IBM i

- What mechanism are available to an Exit Point developer to do this, here are some with a discussion to follow:
  - Interrupt Job Exit - Invoking a Program in Another Job
  - Inter process communications via signals and semaphores
  - Inter process communications via queues
  - Inter process communications via locks
  - Inter process communications via data area/system area
- Ideally you would develop capabilities suitable to your organization or applications

## Allowing or preventing a Suspend on IBM i

- Interrupt Job Exit
  - This allows you to invoke a program in another job, allowing you therefore to validate some job specific environment setting such as a QTEMP object or environment variable
  - To use this capability you need to know information about each job on the system you need to interrogate, this would need to be designed into your environment
  - You would also potentially be stopping a critical piece of functionality whilst this check occurs, which may be problematic for time sensitive communications activities

## Allowing or preventing a Suspend on IBM i

- Inter-process communications (POIX style semaphores and signals)
  - There are many options here

When some critical activity is occurring a semaphore could be create, the exit program checks for the semaphore and if found blocks the suspend operation. If no such semaphore exists then the exit program creates one which indicates that the system is about to be suspended and no critical activity should occur. On resume another exit prgram could clear the semaphore

When the exit program starts execution it sends out a signal and waits for a response (semaphore or other IPC) to indicate whether there is critical system activity at that time

- Be aware of the implications of user activities
- 
- Ideal you would implement this via a set of library functions in service programs
- Downside is awareness of POSIX capabilities on IBM i

## Allowing or preventing a Suspend on IBM i

- Inter-process communications (Queues)
  - Same concepts as for the previous example, however using queues and non destructive reads for the Inter-process communication between the environment and the exit program
- Inter-process communications (Locks)
  - Utilizing locking on objects to indicate that critical system activity is occurring or that a partition is being suspended
- Inter process communications via data area/system area
  - Utilize a global data area(s) or system area(s) in a common library to indicate critical system activity or being suspended
- These options are better understood by the broader IBM i community
- Ideally implemented via a set of functions with a service library

# Agenda

*What we will talk about today:*


**1 Suspend / Resume**

**2 What implications are there on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Resuming a suspended IBM i partition**

**5 Migrating Aix LPAR aware**



## Resuming a suspended IBM i partition

- At any given point a suspended partition can be resumed and it will continue as if had not been suspended.
- Depending on what was happening on the system at the time of suspension, there may be a set of cleanup that needs to occur!
- It is possible to utilize a capability on the IBM I to programmatically undertake some cleanup or undertake other activities that may be desirable

## Resuming a suspended IBM i partition

- This capability is surfaced on the IBM i by a Work Management exit added with PTF SI42815 called QIBM\_QWC\_RESUME

```

Session D - [24 x 80]
File Edit View Communication Actions Window Help
Host: ibmi-7vios.aixncc.uk.ibm.c Port: 23 Workstation ID: Disconnect
Work with Registration Information
Type options, press Enter.
5=Display exit point 8=Work with exit programs

Exit
Point
Format Registered Text
Opt Exit Point
- QIBM_QWC_JOBITPPGM JITP0100 *YES Job interrupt program exit poi
- QIBM_QWC_PRERESTRICT PRER0100 *YES ENDSYS ENDSBS *ALL pre restri
- QIBM_QWC_PWRDWSYS PWRD0100 *YES Prepower down system exit poi
- QIBM_QWC_PWRDWSYS PWRD0200 *YES Prepower down system exit poi
- QIBM_QWC_QSTGLOWACN STGL0100 *YES Auxiliary storage lower limit
- QIBM_QWC_RESUME RSMS0100 *YES Resume system
- QIBM_QWC_SUSPEND SSPS0100 *YES Suspend system
- QIBM_QWT_JOBNOTIFY NTFY0100 *YES Job Notification
- QIBM_QWT_PREATTPGMS ATTN0100 *YES Preattention program exit poi
- QIBM_QWT_SYSREQPGMS SREQ0100 *YES Presystem request pgm exit poi
- QIBM_QYIV_INVGTRSRV GTRS0100 *YES Inventory gathering services
More...

Command
===>
F3=Exit F4=Prompt F9=Retrieve F12=Cancel

MA D MW 09/003
1902 - Session successfully started
Smarter Commerce

```

## Allowing or preventing a Suspend on IBM i

- At this moment there is no publically available document on how you would code to this exit point, however we will discuss what we need to think about when utilising this exit point



## Resuming a suspended IBM i partition

- Whilst at first glance this seems like a pretty straight forward exit point, there are several things to be aware of:
  - The exit point calls a standalone program
  - This program runs in a standalone job
- What are the implications of this?
- The prime implication is that it is the responsibility of the exit point implementer to ensure that the system, or application's that are Suspend/Resume aware, that the system is being resumed. However as the program runs in its own job there is a burden on the exit point developer to have some type of inter process communication, or interaction, with other jobs and processes on the system and to respond to back in a timely manner
- This could potentially mean retro-fitting the required Suspend/Resume awareness into existing applications and environments

## Resuming a suspended IBM i partition

- What mechanisms are available to an Exit Point developer to do this, here are some with a discussion to follow:
  - Interrupt Job Exit - Invoking a Program in Another Job
  - Inter process communications via signals and semaphores
  - Inter process communications via queues
  - Inter process communications via locks
  - Inter process communications via data area/system area
- Ideally you would develop capabilities suitable to your organisation or applications

## Resuming a suspended IBM i partition

- Interrupt Job Exit
  - This allows you to invoke a program in another job, allowing you therefore to validate some job specific environment setting such as a QTEMP object or environment variable
  - To use this capability you need to know information about each job on the system you need to interrogate, this would need to be designed into your environment
  - You would also potentially be stopping a critical piece of functionality whilst this check occurs, which may be problematic for time sensitive communications activities

# Agenda

*What we will talk about today:*

**1 Suspend / Resume**

**2 What really happens on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Migrating Aix LPAR aware**



## Resuming a suspended IBM i partition

- Whilst at first glance this seems like a pretty straight forward exit point, there are several things to be aware of:
  - The exit point calls a standalone program
  - This program runs in a standalone job
- What are the implications of this?
- The prime implication is that it is the responsibility of the exit point implementer to ensure that the system, or application's that are Suspend/Resume aware, that the system is being resumed. However as the program runs in its own job there is a burden on the exit point developer to have some type of inter process communication, or interaction, with other jobs and processes on the system and to respond to back in a timely manner
- This could potentially mean retro-fitting the required Suspend/Resume awareness into existing applications and environments

## Resuming a suspended IBM i partition

- What mechanisms are available to an Exit Point developer to do this, here are some with a discussion to follow:
  - Interrupt Job Exit - Invoking a Program in Another Job
  - Inter process communications via signals and semaphores
  - Inter process communications via queues
  - Inter process communications via locks
  - Inter process communications via data area/system area
- Ideally you would develop capabilities suitable to your organisation or applications

## Resuming a suspended IBM i partition

- Interrupt Job Exit
  - This allows you to invoke a program in another job, allowing you therefore to validate some job specific environment setting such as a QTEMP object or environment variable
  - To use this capability you need to know information about each job on the system you need to interrogate, this would need to be designed into your environment
  - You would also potentially be stopping a critical piece of functionality whilst this check occurs, which may be problematic for time sensitive communications activities

## Resuming a suspended IBM i partition

- Inter-process communications (POIX style semaphores and signals)
  - There are many options here

When some critical activity is occurring a semaphore could be create, the exit program checks for the semaphore and if found blocks the suspend operation. If no such semaphore exists then the exit program creates one which indicates that the system is about to be suspended and no critical activity should occur. On resume another exit program could clear the semaphore

When the exit program starts execution it sends out a signal and waits for a response (semaphore or other IPC) to indicate whether there is critical system activity at that time

- Be aware of the implications of user activities
- 
- Ideal you would implement this via a set of library functions in service programs
- Downside is awareness of POSIX capabilities on IBM i



## Resuming a suspended IBM i partition

- Inter-process communications (Queues)
  - Same concepts as for the previous example, however using queues and non destructive reads for the Inter-process communication between the environment and the exit program
- Inter-process communications (Locks)
  - Utilizing locking on objects to indicate that critical system activity is occurring or that a partition is being suspended
- Inter process communications via data area/system area
  - Utilize a global data area(s) or system area(s) in a common library to indicate critical system activity or being suspended
- These options are better understood by the broader IBM i community
- Ideally implemented via a set of functions with a service library

# Agenda

*What we will talk about today:*

**1 Suspend / Resume**

**2 What implications are there on IBM i**

**3 Allowing or preventing a Suspend on IBM i**

**4 Resuming an IBM I partition**

**5 Migrating Aix LPAR aware**



# Migrating Aix LPAR aware applications

- AIX uses a different mechanism to inform applications that there is a Suspend/Resume or Live Partition Mobility activity occurring for a partition
- The AIX kernel sends a SIGRECONFIG signal out to all processes
- The process can then use the dr\_reconfig() api to gain insight into what what is occurring within the partition and can use the same api to veto the activity

The screenshot displays the IBM Power Systems management console interface. A window titled 'hmc13: Suspend - Mozilla Firefox' shows the status of a partition named 'AIXdemo-142'. The status is 'Suspend status' and the action 'Preparing' is marked with a green checkmark. Below this, a table shows the progress of various steps: 'Validating' (In progress), 'Saving platform data', 'Saving partition data' (MB total = 4910 MB remaining = 4676 Percent complete =), and 'Completing'. A 'Stop...' button is visible.

In the foreground, a terminal window titled 'Session C - [24 x 80]' shows the execution of a program named 't\_sigwaitinfo'. The terminal output includes the following text:

```

150 Opening data connection for t_sigwaitinfo (66466 bytes).
226 Transfer complete.
66466 bytes received in 0.01698 seconds (3824 Kbytes/s)
local: t_sigwaitinfo remote: t_sigwaitinfo
ftp> quit
221 Goodbye.
$ ls -la
total 152
drwxr-xr-x  2 blightj  system    256 Dec  2 13:39 .
drwxr-xr-x  3 blightj  staff     256 Dec  1 16:26 ..
-rw-r--r--  1 blightj  system   1453 Dec  1 16:29 dspsigset.c
-rw-r--r--  1 blightj  system   66466 Dec  2 13:39 t_sigwaitinfo
$ chmod 777 t_sigwaitinfo
$ ./t_sigwaitinfo
./t_sigwaitinfo: PID is 7471302
./t_sigwaitinfo: signals blocked
before for...before sigwaitinfobefore sig check..before term check...got si
58 (SIG-58)
si_signo=58, si_code=8 (other), si_value=0
si_pid=0, si_uid=0
before sigwaitinfobefore sig check..before term check...got signal: 58 (SIG
si_signo=58, si_code=8 (other), si_value=0
si_pid=0, si_uid=0

```

# Migrating Aix LPAR aware applications

- The IBMi PASE environment has the signal and the api defined within the open system include files

The image shows two terminal windows side-by-side, both titled 'Session D - [27 x 132]'. The left window shows the contents of the file `/QOpenSys/usr/include/sys/dr.h`, with the cursor on line 169. The right window shows the contents of the file `/QOpenSys/usr/include/sys/dr.h`, with the cursor on line 66. The right window also shows the command `migrate` being entered at the prompt.

```

Browse : /QOpenSys/usr/include/sys/dr.h
Record : 66 of 807 by 18
Control : migrate

.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3.
/*
 * Signals 50 to 57 are as
 * See SIGRTMIN and SIGRTM
 */
#define SIGRECONFIG 58 /*
#define SIGCPUFAIL 59 /*
/* RESERVED - DON'T USE
#define SIGKAP 60 /*
#define SIGGRANT SIGKAP /*
#define SIGRETRACT 61 /*
#define SIGSOUND 62 /*
#define SIGSAK 63 /* se
/*
 * additional signal names
 */
#define SIGIOINT SIGURG /*
#define SIGAIO SIGIO /* be
#define SIGPTY SIGIO /* p
#define SIGIOT SIGABRT /*

unsigned int cap_constrained : 1; /* Set to 1 if requested capacity
 * update is constrained by PHYP to
 * be within partition capacity bounds.
 */
unsigned int migrate : 1;
unsigned int hibernate : 1;
unsigned int partition : 1;

unsigned int workload_partition : 1; /* workload partition event */
unsigned int checkpoint : 1; /* checkpoint operation */
unsigned int restart : 1; /* restart operation */

/* The following fields are filled out for cpu based requests */
int lcpu; /* logical id of target cpu */
int bcpu; /* bind-id of target cpu */
uint phase_gen; /* Reserved for kernel use.

F3=Exit F10=Display Hex F12=Cancel F15=Services F16=Repeat find F19=Left F20=Right
14/012
1902 - Session successfully started
  
```

# Migrating AIX LPAR aware applications

- Running the same AIX LPAR aware application inside the IBMi PASE QP2TERM environment, we see that the IBMi does not send a SIGRECONFIG signal to listening AIX DLPAR aware applications.

The screenshot displays the HMC console interface. On the left, a table shows the status of various processes. A modal window titled 'hmc13: Suspend - Mozilla Firefox' is open, showing the suspend status for 'IBMI-vios'. The status is 'In progress' with a progress bar at 4% completion. The terminal window on the right shows the following output:

```

/00openSys/usr/bin/-sh
#
> cd /home/BLIGHTJ
#
> ls -la
total 224
drwxrws--- 2 qsecofr  0      8192 Dec 02 14:56 .
drwxrwsrwx  5 qsys    0      8192 Dec 02 14:51 ..
-rwxrwx---  1 qsecofr  0     66466 Dec 02 14:56 t_sigwaitinfo
#
> ./t_sigwaitinfo
./t_sigwaitinfo: PID is 620
./t_sigwaitinfo: signals blocked

```

## Migrating Aix LPAR aware applications

- What does this mean
- Any application vendor who has an existing IBMi PASE application or is in the process of porting an application which is DLPAR aware needs to recode their application(s)
- They need to utilize the existing IBMi OS/400 exit points as described in the previous section
- Multiple choices on how to achieve this:
  - Redesign the application to utilize the IBMi exit point
  - or
  - Custom IPC mechanism such as
    - Have an exit point program that sends out the SIGRECONFIG signal to the relevant applications (there may be additional work to have this list of interested applications available to the exit program)
    - Provide a dr\_reconfig api wrapper that handles the investigatory and veto process via the external exit program