


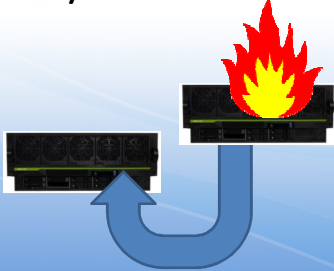



Simplified Remote Restart (SRR)
via HMC and/or PowerVC



 Nigel Griffiths
POWER Advanced Technology Support
IBM Europe

© 2015 IBM Corporation

(complex) Remote Restart

- Originally RR needs Active Memory Sharing (AMS) memory & backing paging devices
- This was used to store the VIOS & LPAR details to allow a remotely rebuild LPAR/VM
- AMS use was not widely implemented
- This (complex) Remote Restart rarely used

Simplified Remote Restart

- Much simpler!!
→ VM details captured to the HMC
- Need to tell HMC to do this
& reboot the VM ☹ sorry about that!
- Worth doing at VM (LPAR) creation time
or set it for next reboot

As simple as 1-2-3

1. Get Live Partition Mobility working
 - If LPM does not work then RR never will!
2. Set Remote Restart - ASAP
3. LPM Validate + SRR Validate
4. Test it works
 - HMC CLI
 - PowerVC GUI
5. Prepare for the worse: pre-decide VM priority & targets
6. BANG!! Don't Panic – do it for real!!! (unlikely)

Live Partition Mobility Reminder

- **Live Partition Mobility** = while running
 - Many years experience now
- **Static Partition Mobility** = while shutdown
 - Quick as no memory to move
- **Dead Partition Mobility** = from the grave!
 - Simple Remote Restart



Live Partition Mobility - Reminder

- Requires PowerVM Enterprise Edition
- Requires “spare” capacity on the target server
- Keep HMC & VIOS’s up to date
- Pure virtual network (SEA) & disks (vSCSI or vFC)
- Source + Target need same subnet & disks/LUN access
- Gotcha!!
 1. Virtual optical media from VIOS → can just delete it
 2. Logical Memory Block size → ASMI + Server reboot
 3. Processor Mode to older box (can’t move P8 VM to a P7 box)
 4. Linux on POWER OK but missing the IBM RPMs



Live Partition Mobility - Best Practice

- 10G to 10G dedicated Ethernet connection for LPM (if possible)
- 1. Keep up to date using Fix Level Recommendation Tool (FLRT)
 - <https://www14.software.ibm.com/webapp/set2/flrt/home>
- 2. LPM setup **checklist** for first time
 - <http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/tips1184.html>
- 3. LPM prep **checklist** if its been a while since LPM for a partition
 - <http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/tips1185.html>
- 4. Follow VIOS performance guidelines for LPM
 - http://www-01.ibm.com/support/knowledgecenter/9119-MHE/p8hc3/p8hc3_viostone_lpmperf.htm?cp=9119-MHE&lang=en



2 Set Simplified Remote Restart flag

2

- Assuming LPM is a “go”




2 Set Simplified Remote Restart flag

2

- Machine Properties Capable

Capability	Value
Inactive Partition Mobility Capable	True
IBM i Partition Mobility Capable	True
Partition Processor Compatibility Mode Capable	True
Partition Availability Priority Capable	True
Electronic Error Reporting Capable	True
Active Partition Processor Sharing Capable	True
Firmware Power Saver Capable	True
Hardware Power Saver Capable	True
Virtual Switch Capable	True
Virtual Fibre Channel Capable	True
Active Memory Expansion Capable	True
Hardware-Accelerated Active Memory Expansion Capable	True
Partition Suspend Capable	True
Partition Remote Restart Capable	True
PowerVM Partition Remote Restart Capable	True
Virtual Trusted Platform Module Capable	True
SR-IOV Capable	True
Dynamic Platform Optimization Capable	True
Virtual Server Network Phase 2 Capable	True
PowerVM Partition Simplified Remote Restart Capable	True



2 Set Simplified Remote Restart flag




2


- Can't set Remote Restart flag at the
 - VM / LPAR OS
 - HMC GUI LPAR profile
 - HMC Enhanced+ LPAR profile
- Current only via the HMC command line



2 Set Simplified Remote Restart flag

2

- Assuming LPM is a “go” 
- Machine Properties Capable is “go” 
- Access to the HMC command line is “go” 



HMC CLI Check (if capable)

2

lssyscfg -r lpar -m **machine** --filter lpar_names="LPARname"

```
hmc> lssyscfg -r lpar -m P8-lime-8284-22A-SN215296V --filter lpar_names="vm61"
name=vm61,lpar_id=7,lpar_env=aixlinux,state=Running,
resource_config=1,os_version=AIX 7.1 7100-03-05-1524,
logical_serial_num=215296V7,default_profile=default_profile,
curr_profile=default_profile,work_group_id=none,
shared_proc_pool_util_auth=1,allow_perf_collection=1,
power_ctrl_lpar_ids=none,boot_mode=norm,lpar_keylock=norm,
auto_start=0,redundant_err_path_reporting=0,rmc_state=active,
rmc_ipaddr=9.137.62.61,time_ref=0,lpar_avail_priority=127,
desired_lpar_proc_compat_mode=POWER8,
curr_lpar_proc_compat_mode=POWER8,suspend_capable=0,
remote_restart_capable=0, ← NOT this option
simplified_remote_restart_capable=1,
remote_restart_status=Remote Restartable,
sync_curr_profile=0,affinity_group_id=none,vtpm_enabled=0
```



HMC CLI Check (if current flags)

2

Check whole machine:

```
hmc> lssyscfg -r lpar -m P8-lime-8284-22A-SN215296V \
-F simplified_remote_restart_capable,name
0,limevios1
0,limevios2
0,vm36_Ubuntu1504
0,vm26-ubuntu1504
0,vm35_SLES12
0,vm20-SLES-11.3
0,vm22-RHEL7-GA
0,vm112-64d8b471
1,vm61
hmc>
```



HMC CLI set

2

- To enable the remote restart feature:

- `chsyscfg -r lpar -m server`
- `-i "name=partition name, simplified_remote_restart_capable=1"`

- Example:

- `chsyscfg -r lpar -m P8-lime \`
- `-i "name=vm61,simplified_remote_restart_capable=1"`
- Takes a couple of seconds

- To disable similar but “=1” → “=0”



Warning:

2

- I am told this adds some CPU overhead on the HMC
- It regularly collects VIOS config details for SRR VM's
- Not seen how a much CPU%
 - It is not large = occasional



HMC Classic GUI (once set on the CLI)

Partition Properties - vm61_AIX735

General Hardware Virtual Adapters Settings Other

Name: *vm61_AIX735

ID: 7

Environment: AIX or Linux

State: Running

Attention LED: Off

Resource configuration: Configured

OS version: AIX 7.1 7100-03-05-1524

Current profile: default_profile

System: 8284-22A*215296V

Allow performance information collection

Allow this partition to be suspended.

Virtual Trusted Platform Module (VTPM)

Warning: VTPM Trusted Key is the default key.

Sync current configuration Capability: Sync turned OFF

Remote Restartable (Simplified): Activated

OK Cancel Help

View only - Not settable here



HMC Enhanced+ (once set on the CLI)

The screenshot shows the HMC interface for partition `vm61_AIX735`. The 'General Properties' section is active, displaying various system details. Under the 'Virtualization Capabilities' section, the 'Remote Restart' checkbox is checked. A red arrow points to this checkbox with the text 'View only - Not settable here'.

Property	Value
Partition Name:	vm61_AIX735
OS Type / Environment:	AIX/Linux
OS Version:	AIX 7.1 7100-03-05-1524
IP Address:	9.137.62.61
Boot Mode:	Normal
Key Lock Position:	Manual (selected), Normal
Serial Number * System Machine Type:	8286-42A*100EC7V
Description:	
Group Tags:	
Virtualization Capabilities:	<input checked="" type="checkbox"/> Suspend / Resume <input checked="" type="checkbox"/> Remote Restart

3 Validate for LPM & SRR

3


- Live Partition Mobility – you should know this!



3 Validate for LPM

3

- Live Partition Mobility – you should know this!



3 Validate for LPM

3

Partition Migration Validation - P8-lime-8284-22A-SN215296V - vm61_AIX735

Fill in the following information to set up a migration of the partition to a different managed system. Click Validate to ensure that all requirements are met for this migration. You cannot migrate until the migration set up has been verified.


Source system : P8-lime-8284-22A-SN215296V
 Migrating partition: vm61_AIX735
 Remote HMC:
 Remote User:
 Destination system: P8-emerald-8286-42A-SN100EC7V
 Destination profile name:
 Destination shared processor pool:
 Source mover service partition:
 Destination mover service partition:
 Wait time (in min):
 Override virtual network errors when possible:
 Override virtual storage errors when possible:
 Override partition UUID:
 Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS
<input type="checkbox"/>			



3 Validate for LPM

3



This is OK & LPM/SRR ready



3 Validate for Simplified RR

3

Back on the HMC CLI

```
hmc> rrstartlpar -o validate
```

```
    -m Source-box
```

```
    -p LPAR-name
```

```
    -t Target-box
```

3 Validate for Simplified RR

3

```
hmc> rrstartlpar -o validate -m P8-lime-8284-22A-SN215296V
-p vm61 -t P8-emerald-8286-42A-SN100EC7V
```

Warnings:

HSCLB504 The migrating partition cannot use hardware-accelerated Active Memory Expansion on the destination managed system because the destination managed system does not support hardware-accelerated Active Memory Expansion.

HSCLA4CC The management console cannot maintain the source Virtual I/O Server (VIOS) slot number 14 for virtual SCSI adapter 5 on the destination VIOS partition 2*8286-42A*100EC7V.

HSCLA4CC The management console cannot maintain the source Virtual I/O Server (VIOS) slot number 3 for virtual SCSI adapter 4 on the destination VIOS partition 3*8286-42A*100EC7V.

```
hmc>
```



Does

3

rrstartlpar validate = LPM validate ?

- No as it also checks Remote Restart flag is set
- If Flag not set, you get:
 - HSCLA9B9 Partition LPARNAME is not remote restart capable.
- It might check other things too.

rrstartlpar Complete Syntax

- **rrstartlpar**
 - o { **restart** | **validate** | **cancel** | **cleanup** | **recover** }
 - m *managed-system* [-t *target-managed-system*]
 - {-p *partition-name* | --id *partition-ID*}
 - [--redundantvios {0 | 1 | 2}] ← old
 - [--mpio {1 | 2}]
 - [--vlanbridge {1 | 2}]
 - [--retaindev]
 - [--usecurrdata]
 - [-w *wait-time*] ← default 3 minutes
 - [-d *detail-level*] ← amount of output
 - [--force] ← cleanup/recover
 - [-v] ← verbose

SRR Official Pre-Requisties Briefly

Machine Level

1. LPM pre-reqs = access to same external storage & sub-net
2. The HMC 820 SP1 or later (with latest PTF) + 820 firmware
3. Machines are simplified remote restart capable
4. Both hosts must be managed by the same HMC
5. HMC ←→FSP connection (needs to definitely know the box is off)
6. The source host must be in **Error**, **Power Off**, or **Error - dump in progress** state on the HMC. (NOTE: Power off from the HMC is OK)

LPAR / VM Level

1. VM must be Simplified Remote Restart capability enabled
2. Remote restart state of the VM must be "Remote restartable"

Note:

- SSP not officially supported ("Mine works fine!!" says Nigel)

Lets pretend the machine crashed

The word "BANG!" is written in a large, bold, red, sans-serif font with a thick black outline. The text is slanted slightly to the right and has a dynamic, explosive feel.

Example from my machines

```
hmc> rrstartlpar -o restart  
-m P8-lime-8284-22A-SN215296V  
-p vm61  
-t P8-emerald-8286-42A-SN100EC7V
```

Example using Shell variables

```
SOURCE=P8-lime-8284-22A-SN215296V
```

```
VM=vm61
```

```
TARGET=P8-emerald-8286-42A-SN100EC7V
```

Alternative


```
hmc>
```


```
> rrstartlpar -o restart -m $SOURCE -p $VM -t $TARGET
```

```
...
```

```
>
```

It just works fine

- Like LPM you can watch it on the HMC go through various phases of creating the VM 

- Restarts the VM automatically at the end for you 

Clean up the debris!



VM / LPAR Definition on the **source** machine
is still there → unlike LPM

Why?

- Source Machine was powered-off so
Not possible to remove the LPAR
- VIOS was shutdown = no virtual I/O deconfig

Clean up later - after the PANIC!



```
SOURCE=P8-lime-8284-22A-SN215296V
```

```
VM=vm61
```

```
hmc> rrstartlpar -o cleanup -m $SOURCE -p $VM
```

```
HSCLA9CE The managed system is not in a valid state to support  
partition remote restart operations.
```

```
[[ Need to power-on the machine ]]
```

```
hmc> rrstartlpar -o cleanup -m $SOURCE -p $VM
```

```
HSCLA928 The Virtual I/O Server (VIOS) partition limevios1 is not  
in the Running state. This operation is only allowed when the  
VIOS partition is running.
```

```
[[ Need to start the VIOS\(s\) ]]
```


Clean up later - after the PANIC!



After machine power up, VIOS(s) started and settled down


```
hmc> rrstartlpar -o cleanup -m $SOURCE -p $VM
```

[[for my small VM about 10 second]]


```
hmc>
```

All Done !

So how to test RR?




Test Prep setup



- Get set up (reminder)
 - Set up the Virtual Machine with SRR flag
 - HMC GUI: Operations → Mobility → Validate
 - HMC CLI: `rrstartlpar -o verify . . .`

Test Prep setup




Get to a machine state were SSR works
= **Power Off**

1. Shutdown Virtual Machine
 2. Shutdown VIO Server(s)
 3. Power Off the machine
- **Power Off** state → ready for SRR

Note: This is slightly cheating!!

- In a real crash the VM + VIOS don't cleanly stop & flush disks

Test Prep setup



Get to a machine state were SSR works
= **Power Off**


- **Ugly Alternatives:**
- **1) Yank the power cords! * ****
- **2) HMC Power-off Server - VIOS & VM running***

Note:

* Now VM + VIOS have to crash recover filesystems


** Need the FSP running afterwards, so power up to PowerOff

Testing via HMC CLI



- HMC CLI: `rrstartlpar -o restart ...`
 - Should return fairly quickly (after validate phase)
- Watch the HMC status to see it working on the target machine
- The VM will start automatically
- If it fails: `rrstartlpar -o recover ...`

Post Testing



Clean up source machine


1. Restart machine
2. Restart VIOS
3. Wait 5 minutes
4. `rrstartlpar -o cleanup ...`

But there is another way ...

PowerVC can do Simplified Remote Restart

Same as pocess but don't use HMC rrestartlpar

Don't worry that you can't find any SRR clues in PowerVC

- SRR button  **Remotely Restart Virtual Machines**
only appears when your click a server in the right state
(like Power-off) and the Host machine is SRR ready
- Warning PowerVC pre-reqs for SSR
 - HMC 830 + Firmware 830 – IMHO recommend latest update
 - VIOS 2.2.3.4+
 - Current PowerVC is 1.2.3.2

PowerVC

IBM PowerVC Users Configuration Messages 2 nag IBM

Hosts

Hosts Host Groups HMC Connections

Refresh Add Host Remove Host Enter Maintenance Mode Change HMC

Manage Existing Virtual Machines Remotely Restart Virtual Machines

Name	Virtual Machines	State	Maintenance Status	Health	HMC Connection
P6-bronze-8203-E4A-SN10E0A21	3	Operating	OK	OK	HMC14
P6-gold-8203-E4A-SN10E0A11	0	Operating	OK	OK	HMC14
P6-orange-8203-E4A-SN10E0A51	0	Operating	OK	OK	HMC14
P6-red-8203-E4A-SN10E0A41	2	Operating	OK	OK	HMC14
P6-silver-8203-E4A-SN10E0A31	1	Operating	OK	OK	HMC14
P7-diamond-8233-E9B-SN100271P	0	Operating	OK	OK	HMC14
P7-indigo-8231-E1C-SN0659FDR	2	Operating	OK	OK	HMC14
P7-purple-9117-MMB-SN100525P	1	Operating	OK	OK	HMC14
P8-emerald-8286-42A-SN100EC7V	11	Unknown	OK	Attention	HMC14
P8-lime-8284-22A-SN215296V	18	Power Off	OK	Warning	HMC14

Click

PowerVC

Remotely Restart Virtual Machines

After selecting a virtual machine to remotely restart, you can specify the destination host.

Select a virtual machine

Remotely restart all virtual machines.

Filter

Name	State	Processors	Memory
vm61_AX735	Error	2	4096

Total: 1 Selected: 0

Then Click

Remote Restart Close

PowerVC Remotely Restart Virtual Machines
 After selecting a virtual machine to remotely restart, you can specify the destination host.

Selected virtual machine
 vm61_AIX735

Select a destination host
 * Host P8-emerald-8286-42A-SN100EC7V

Then Click Remote Restart Close

First Host is determined by PowerVC as the best target or you select a target machine from the list

PowerVC

It starts a background job for that task

✓ The selected virtual machines are being remotely restarted from host P8-lime-8284-22A-SN215296V.
 View All (296)...

On the screen it shows the status

vm61_AIX735	P8-lime-8284-22A-SN215296V	Rebuilding	Critical
-------------	----------------------------	------------	----------

Later is just shows the VM running on the target machine

Rather an anticlimax !!

It will clean up the original VM once the machine + VIOS started
 → It checks every 2 minutes

Further Information



- HMC RR – Knowledge Center → [Manual](#)
- HMC Community Files – DevelopWorks → [Link](#)
- Whitepaper & rrMonitor script
→ [Simplified Remote Restart Info](#)
- PowerVC Overview of Remote Restart → [Link](#)
- PowerVC Deep Dive by Christine Wang → [Link](#)
- Mr chmod666 Articles → [PowerVCHints](#) [UsingSRR](#)

But ...


We have only covered the Trivial Case

- Manually
 - Setup the machine / VM(s) for SRR
 - Later: Noticed a machine has stop unexpectedly
 - Selected a high priority VM to restart
 - Selected a target with enough resources
 - Kick off the recovery
- Back in the real world
 - 10's to 100's of machines some with
100's to 1000's of LPARs



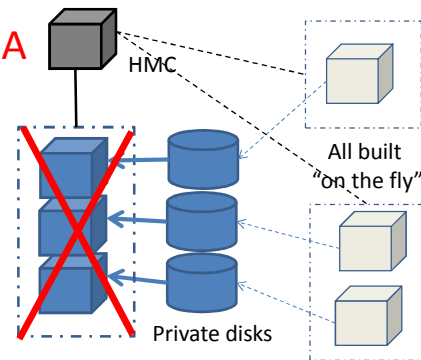
Isn't SRR PowerHA on the cheap?

No it is not like PowerHA



PowerHA

Standby VM constantly monitoring
 One per LPAR
 A second running hot OS
 Already to take-over / fail-over
 Disks already configured
 Works without HMC, FSP etc.
 Completely automatic



Simplified Remote Restart

Only possible when machine = Error/PowerOff
 Can't demine if OS failed / stuck / halted
 HMC monitoring FSP
 No HMC or no FSP = no SRR possible
 Nothing ready to take-over=has to be built/LPM
 Needs common HMC +cold reboot
 Default is manual

SRR scaled up


- Automate

- Setup → HMC CLI = fairly easy then LPM = BAU
- Noticed crash → HMC will notice or other Alert Tool
- Selected priority → Your job!
- Selected target → Your job! (also what can be sacrificed)
- Kick off SRR → a script or tool



Manual SRR ?

Operator Driven




HMC reported HW Events (email)

1. HMC informs "operations team" of machine failure
2. They run procedure to recover production VMs ASAP
3. Check the Server is still down
4. Runs rrstartlpar for selected VM's to selected targets

- IMHO scaling to a few dozen high priority VMs - OK

Automatic SRR ?

Script(s) Driven



Via remote ssh commands to the HMC

1. Script polls for the machine state
2. Check it says down for, say, 5 minutes
 - i.e. not a reboot
3. Run a remote ssh rrstartlpar for selected VM's to selected targets

Automatic SRR ?

Script(s) Driven


Via remote ssh commands to the HMC

1. Script polls for the machine state
2. Check it says down for, say, 5 minutes
 - i.e. not a reboot
3. Run run remote ssh rrestartlpar for selected VM's to selected targets


- Complications:

- How to update this, if using LPM regularly!
- Needs some way to switch off for maintenance
- Multiple Server failures? Might be cause a right mess

- IMHO but not scalable to dozens of Servers with 1000's of VMs



SRR Script !



rrMonitor – not officially supported

HMC Community Files

- https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=0196fd8d-7287-4dff-8526-102b5bcf0df5#fullpageWidgetId=W395818bd593b_487f_a7ec_79c3c27093f8
 - rrMonitor_readme.txt
 - [rrMonitor](#) Perl Script Download ~140 lines
- It does a number of things on a regular basis but not the cleanup function
- Parameters: source-server, target-server,
LPAR-name, time interval (seconds) between steps
- Limitation: 1 script for 1 LPAR
- I am told that this has to be run on HMC, which is normally BANNED!!

56¹⁸³= 1e+320
Molecules=1e+82
In Universe

SRR Script !




rrMonitor fixed

- **POWER guru Benoit** benoit.creau@chmod666.org
- Modified the original script = 217 lines of Perl
- To run as a ssh remote command from another machine
- Find it here
<http://chmod666.org/wp-content/uploads/2015/06/rrMonitor.txt>


Simplified Remote Restart (RR) via HMC and/or PowerVC

Simple to configure – assumes pure virtual & LPM
Requires POWER8 & HMC/FW 820+ [PowerVC 830]
Excellent “get out of jail free” card
No further costs




Effort required for large scale deployment
Not a PowerHA replacement








Power Systems 

Next Time

Nov 18th - VIOS Shared Storage Pool phase 5 and SSP Update



<p>Previous Sessions: Linux on POWER Field Exp POWER8 E850 from exper PowerSC Tools for IBM i HMC 8.20 Tech Preview The "Key" to IBM i Licensing POWER8 E870 from exper Linux on Power: Best Pract Linux for AIX/IBM i guys PowerKVM Deep Dive More Tricks Power Masters Power8 from hands-on Power up your Linux And more.....</p>	<p>Future Sessions →</p> <ul style="list-style-type: none">▪ Dec 2nd - Boost IBM i performance with IBM FlashSystem▪ Suggestions Welcome <div style="text-align: right;"></div> <p>Webinar wiki: http://tinyurl.com/PowerSystemsTechnicalWebinars Youtube Channel: http://tinyurl.com/IBMPowerVUGYoutubeChannel</p> <p>Twitter:  Gareth Coates @power_gaz Nigel Griffiths @mr_nmon Jyoti Dodhia @JyotiDodhia Mandie Quartly @mandieq</p>
--	---