



IBM @server

IBM DLPAR Tool Set for pSeries v2.0.2.0



Joefon Jann, Niteesh Dubey, R. Sarma Burugula - IBM Research
Other contributors: Bob Minns - IBM Austin

Last Updated: September 2004

© 2004 IBM Corporation

Special Notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Many of the pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special Notices (continued)

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L(logo), alphaWorks, AS/400, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DFDSM, DirectTalk, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM, IBM(logo), ibm.com, IBM TotalStorage Proven, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, POWERparallel, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RISC System/6000, RS/6000, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SP1, SP2, System/390, The Engines of e-business, THINK, ThinkPad, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: AIX/L(logo), AIX 5L, AIX PVM, AS/400e, BladeCenter, Chipkill, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, Domino, e-business(logo), e-business on demand, eServer, GigaProcessor, HACMP, HACMP/6000, IBMLink, IMS, Intelligent Miner, iSeries, NUMACenter, POWER, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, S/390 Parallel Enterprise Server, ThinkVision, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Visualization Data Explorer, X-Architecture, z/Architecture.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Intel, Itanium and Pentium are registered trademarks and Intel Xeon and MMX are trademarks of Intel Corporation in the United States and/or other countries.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Other company, product and service names may be trademarks or service marks of others.

Contents

- Overview
- Demonstration Configuration
- Setup Required
- ***lparLsCfgs.pl, lparLsCfgsAll.pl, lparLsCfgsIO.pl***
- ***lparLsLoads.pl***
- ***lparSetCfgs.pl***
- ***lparLoadRM.pl***
- ***moveSlot.pl***
- Additional Tools
- Feedback

Overview

The **DLPAR Tool Set for pSeries** contains a set of scripts, mostly written in Perl :

1. for **displaying (active) partition configurations (CPUs, LMBs, I/O Slots)**
2. for **monitoring the CPU & memory loads on multiple partitions on a single screen**
3. for **time-based automation of DLPAR (Dynamic LPAR) operations** across a group of partitions
4. for **load-based resource sharing** via DLPAR between a group of partitions running AIX V5.2 or later releases
5. for moving a **physical I/O slot** (e.g. **IDE or SCSI controller with CD/DVD/Tape, ethernet adapters, empty slots**) between partitions
6. for testing and exercising these features

Download from IBM AlphaWorks:

- ▶ <http://www.alphaworks.ibm.com/tech/dlpar>

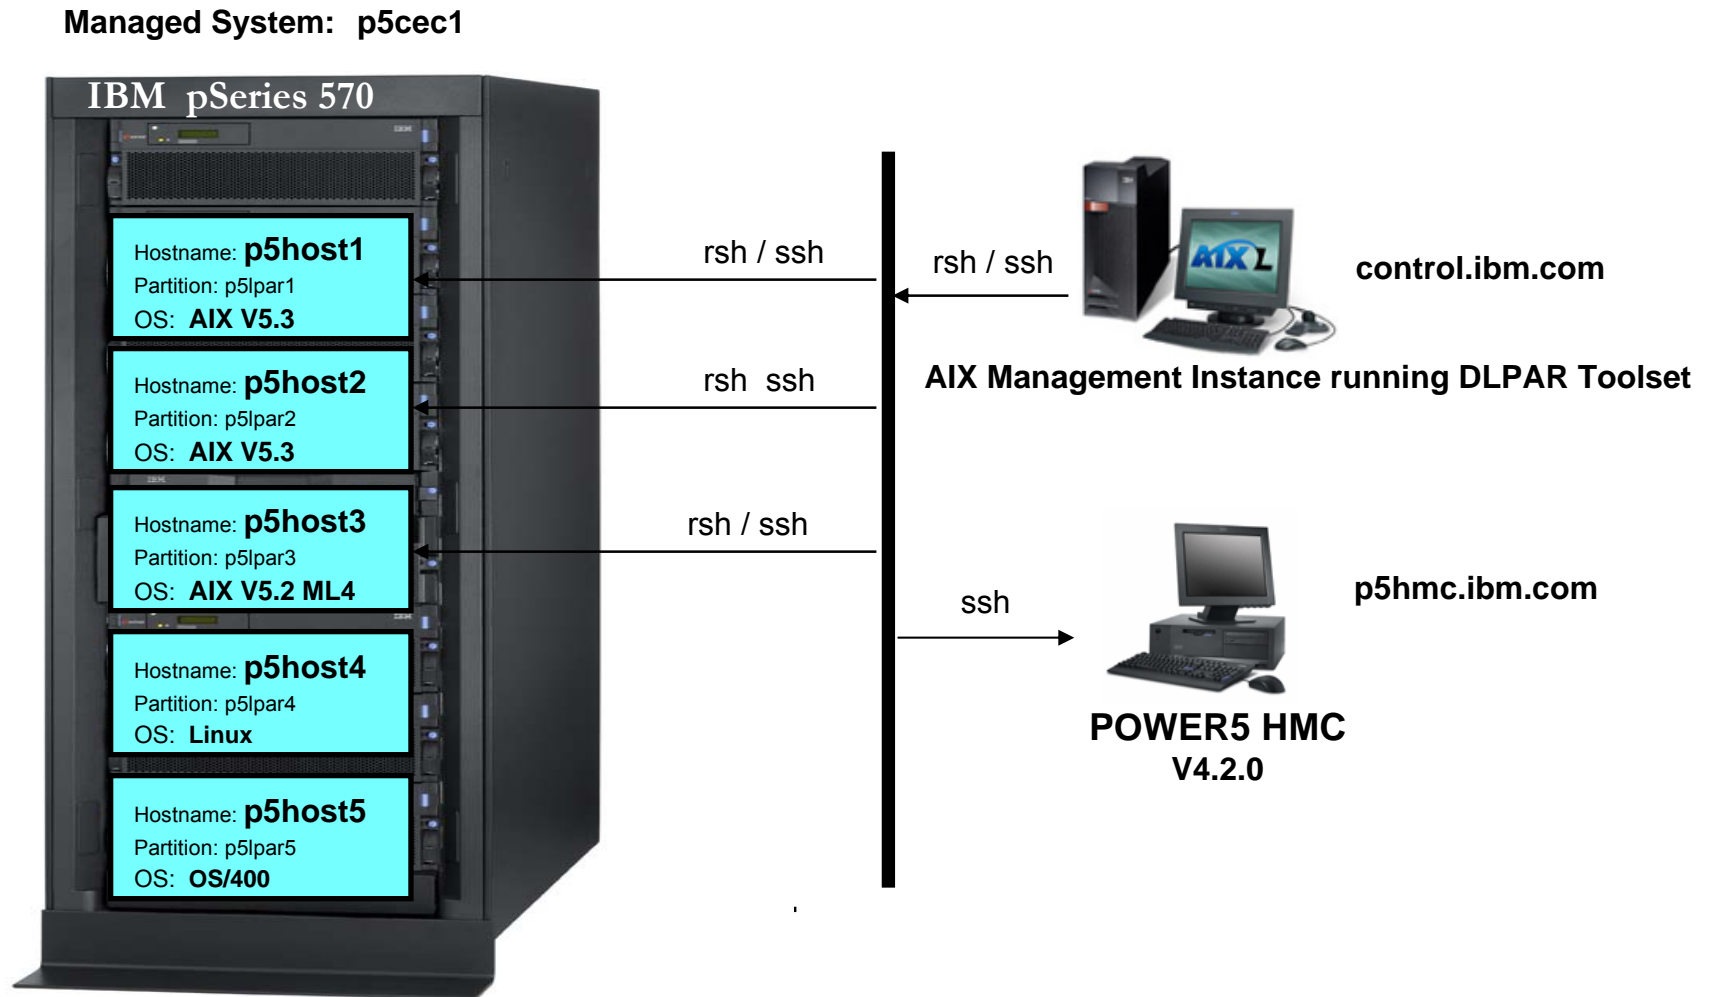
Latest features (V2.0.2.0):

- ▶ Support for **POWER5 pSeries p5 servers** running AIX V5.2 ML4 or AIX V5.3
- ▶ Support for **ssh** or rsh for executing remote commands on AIX hosts
- ▶ Additional scripts - **lparLsCfgsIO.pl, lmbSz.pl, copy_ssh_key.ksh**

Limitations

- ▶ Single managed POWER4 or POWER5 system (CEC)
- ▶ Single HMC (running R3.2.6/V3.3.0 or higher, V4.2.0 or higher)
- ▶ Limited support for micro-partitioning (shared processor partitions)

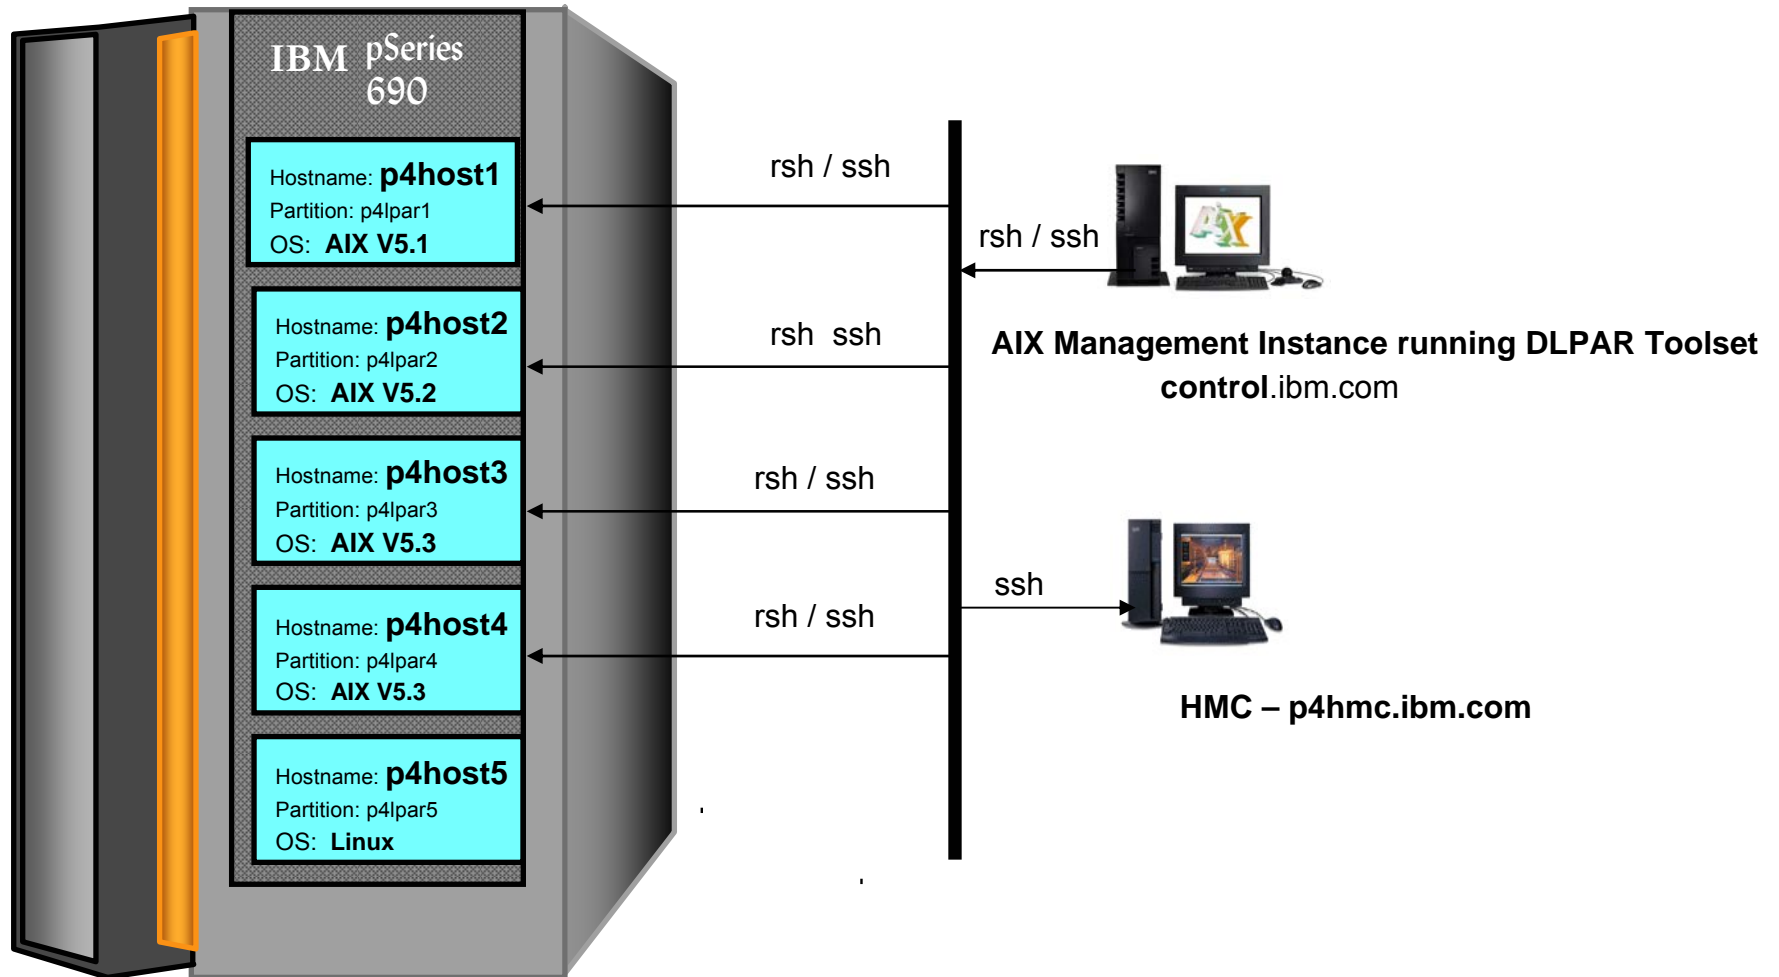
Demo Configuration with LPARs on a p5 Server



The DLPAR toolset is running on an AIX instance to manage the resources (CPU, Memory, IO slots) for LPARs that are running AIX, on a p5 server managed by an HMC.

Demo Configuration of a P4 SMP with LPARs

Managed System: p4cec1



The DLPAR toolset is running on an AIX instance to manage the resources (CPU, Memory, IO slots) for LPARs that are running AIX, on a P4 server managed by an HMC.

Setup Required

1. **An AIX Management Instance** which can run either in a separate PPC-workstation or in one of the LPARs.
2. **ssh** connection from a sysadmin's loginid in the AIX Management Instance to a sysadmin loginid in the HMC.
3. **A hostList file** to contain a subset of AIX LPAR hostnames that you choose to monitor and/or resource sharing.
4. **rsh** or **ssh** privileges from the Management Instance to all the AIX LPARs in your subset.
5. Set up **DR_ environment variables** (in the file **setEnv** or **setEnvc**) for:
 - ▶ The fullpathname of the Directory containing the Tool Set (usually /opt/dlparToolset/)
 - ▶ HMC *hostname*
 - ▶ HMC sysadmin's SSH *userid*
 - ▶ Name of the Managed system (CEC = Central Electronics Complex)
 - ▶ *hostList* file
 - ▶ A snapshot interval value, a history interval value
 - ▶ Trigger values
 - ▶ ...

lparLsCfgs.pl

Lists the configuration of all active partitions running AIX in the **<hostList>** file

- Supports AIX 5.1, 5.2 and 5.3 partitions
- Output includes:
 - ▶ Hostname (OS view)
 - ▶ Partition Name (HMC view)
 - ▶ #CPUs : *minimum#, current#, maximum#*
 - ▶ #Memory Blocks (LMBs) : *minimum#, current#, maximum#*
 - ▶ DLPAR[-capable] : YES/NO
 - ▶ #CPUs in the free pool of the CEC
 - ▶ #LMBs in the free pool of the CEC
- If a partition is running in “shared processor mode” on a p5 server, #CPUs represents the number of virtual processors on that partition
- Output is automatically saved to file (e.g. ~/lparLsCfgs.outFile) for future reference

Sample lparLsCfgs.pl

This script is being run by user "root" on host "control".

INPUT:

```
Host list           = /opt/dlparToolset/bin/hostList
hmcHost            = p5hmc.ibm.com
hmcUser            = hscroot
Managed System    = p5cecl
RemoteShell for AIX hosts = rsh
output File        = /lparLsCfgs.outFile
```

Checking ping and ssh to p5hmc.ibm.com.....

HMC Commands are running under Restricted Bash shell on POWER5 HMC (V4R2.0).....

Reading the /opt/dlparToolset/bin/hostList file

Checking ping and rsh to p5host1.....

Checking ping and rsh to p5host2.....

Checking ping and rsh to p5host3.....

Getting the Current Configurations

NOTE: It takes roughly 30 to 45 seconds to get the current configuration of each LPAR from the HMC.

Hostname	LPARname	CPUmin	CPUcur	CPUmax	LMBmin	LMBcur	LMBmax	DLPAR
p5host1	p5lpar1	1	1	6	8	64	80	YES
p5host2	p5lpar2	1	2	6	16	64	128	YES
p5host3	p5lpar3	1	2	8	8	12	32	YES

Available #CPUs in the Free Pool = 3

Available #LMBs in the Free Pool = 13

NOTE: 1 LMB of Memory = 128 MB

lparLsCfgsAll.pl

Lists the configurations (CPUs, Memory & I/O slots) of all the active partitions, irrespective of the OS type of the LPAR. Resources in the free pool are also listed

- Output includes:
 - ▶ Partition Name (HMC view)
 - ▶ #CPUs : *minimum#, current#, maximum#*
 - ▶ #Memory Blocks (LMBs) : *minimum#, current#, maximum#*
 - ▶ DLPAR[-capable] : YES/NO
 - ▶ #CPUs in the free pool of the CEC
 - ▶ #LMBs in the free pool of the CEC
 - ▶ All the physical I/O slots : *lparName/FreePool, [DRCIndex], physicalLocations, Descriptions*
 - ▶ It displays the DRCIndex of the slot only if the managed system is a p5 server
- If a partition is running in “shared processor mode” on a p5 server, #CPUs represents the number of virtual processors in that LPAR.
- Output is automatically saved to file (e.g. ~/lparLsCfgsAll.outFile) for future reference

Sample lparLsCfgsAll.pl Output:

This script is being run by user "root" on host "control".

INPUT:

```
hmcHost          = p5hmc.ibm.com
hmcUser          = hscroot
Managed System  = p5cec1
output File      = /lparLsCfgsAll.outFile
```

Checking ping and ssh to p5hmc.watson.ibm.com.....

HMC Commands are running under **Restricted Bash shell on POWER5 HMC (V4R2.0)**.....

Retrieving the current configuration from HMC

NOTE: It takes roughly 30 to 45 seconds to get the current configuration of each LPAR from the HMC.

```
===== CPUs and Memory =====
LPARname          CPUmin  CPUcur  CPUmax  LMBmin  LMBcur  LMBmax  DLPAR
=====
p5lpar1           1        1        6        8       64      80      YES
p5lpar2           1        2        6       16      64     128      YES
p5lpar3           1        2        8        8       12     32      YES
```

Available #CPUs in the Free Pool = 3

Available #LMBs in the Free Pool = 13

NOTE: 1 LMB of Memory = 128 MB

Sample IparLsCfgsAll.pl Output (cont.) :

```

===== IO Slots =====
  DRCIndex      PhysicalLocations      Descriptions
=====
p5lpar1:
  21010001      U7879.001.DQD05WZ-P1-T14      Storage controller
  21030001      U7879.001.DQD05WZ-P1-T6      PCI 10/100/1000Mbps Ethernet UTP 2-port
-----
p5lpar2:
  21030019      U7879.001.DQD05WX-P1-T6      PCI 10/100/1000Mbps Ethernet UTP 2-port
  2101001B      U7879.001.DQD05WX-P1-T12      Storage controller
-----
p5lpar3:
  21010003      U7879.001.DQD05WZ-P1-T12      Storage controller
  21030003      U7879.001.DQD05WZ-P1-C1      PCI 1Gbps Ethernet UTP
-----
Free Pool:
  21010019      U7879.001.DQD05WX-P1-T14      Storage controller
  21020019      U7879.001.DQD05WX-P1-T4      Universal Serial Bus UHC Spec
  2101001A      U7879.001.DQD05WX-P1-C3      Fibre Channel Serial Bus
  2104001A      U7879.001.DQD05WX-P1-C6      Empty slot
  2102001B      U7879.001.DQD05WX-P1-T15      Other Mass Storage Controller
  21020001      U7879.001.DQD05WZ-P1-T4      Universal Serial Bus UHC Spec
  21010002      U7879.001.DQD05WZ-P1-C3      PCI 10/100Mbps Ethernet w/ IPsec
  21020002      U7879.001.DQD05WZ-P1-C4      PCI 1Gbps Ethernet UTP
  21040003      U7879.001.DQD05WZ-P1-C2      PCI 1Gbps Ethernet UTP
=====

```

lparLsCfgsIO.pl

Sorts and displays all the physical I/O slots of the managed system

- ▶ POWER4 -- Sort by PhysicalLocations
 - ▶ POWER5 – Sort by the last 4 digits of DRC Index
The last 4 digits of a DRC index is the “*bus id*”.
- Useful for physically locating the I/O slots on the managed system
 - Output includes:
 - ▶ All the physical I/O slots : *[DRCIndex]* , *physicalLocations*, *Descriptions*
 - ▶ It displays the *DRCIndex* of the slot only if the managed system is a p5 server
 - ▶ Partition Name (HMC view)
 - Output is automatically saved to file (e.g. ~/lparLsCfgsIO.outFile) for future reference

Sample IparLsCfgsIO.pl Output:

DRCIndex	PhysicalLocations	Descriptions	LPARName
21010001	U7879.001.DQD05WZ-P1-T14	Storage controller	p5lpar1
21020001	U7879.001.DQD05WZ-P1-T4	Universal Serial Bus UHC Sp	null
21030001	U7879.001.DQD05WZ-P1-T6	PCI 10/100/1000Mbps Etherne	p5lpar1
21010002	U7879.001.DQD05WZ-P1-C3	PCI 10/100Mbps Ethernet w/	null
21020002	U7879.001.DQD05WZ-P1-C4	PCI 1Gbps Ethernet UTP	null
2103	U7879.001.DQD05WZ-P1-C5	PCI 1Gbps Ethernet UTP	null
21040002	U7879.001.DQD05WZ-P1-C6	Fibre Channel Serial Bus	null
21010003	U7879.001.DQD05WZ-P1-T12	Storage controller	p5lpar3
21020003	U7879.001.DQD05WZ-P1-T15	Other Mass Storage Controll	null
21030003	U7879.001.DQD05WZ-P1-C1	PCI 1Gbps Ethernet UTP	p5lpar3
21040003	U7879.001.DQD05WZ-P1-C2	PCI 1Gbps Ethernet UTP	null
21010019	U7879.001.DQD05WX-P1-T14	Storage controller	null
21020019	U7879.001.DQD05WX-P1-T4	Universal Serial Bus UHC Sp	null
21030019	U7879.001.DQD05WX-P1-T6	PCI 10/100/1000Mbps Etherne	p5lpar2
2101001A	U7879.001.DQD05WX-P1-C3	Fibre Channel Serial Bus	null
2102001A	U7879.001.DQD05WX-P1-C4	Fibre Channel Serial Bus	null
2103001A	U7879.001.DQD05WX-P1-C5	Empty slot	null
2104001A	U7879.001.DQD05WX-P1-C6	Empty slot	null
2101001B	U7879.001.DQD05WX-P1-T12	Storage controller	p5lpar2
2102001B	U7879.001.DQD05WX-P1-T15	Other Mass Storage Controll	null
2103001B	U7879.001.DQD05WX-P1-C1	Fibre Channel Serial Bus	null
2104001B	U7879.001.DQD05WX-P1-C2	Empty slot	null

lparLsLoads.pl

Lists the current & historical CPU & Memory utilizations of all LPARs in 1 window

- Talks to HMC to get the LMB size on the managed system
- Monitors AIX V5.1, V5.2 and V5.3 partitions.
- The user on the AIX management instance executes the AIX commands on all the AIX hosts in the hostList via rsh or ssh
- Can specify your own **snapshot interval value** and **history interval value**
- **Single Display includes** both Snapshot utilization & Historical utilization of CPU & memory:
 - ▶ #P -- number of online CPUs
 - ▶ freP -- number of free CPUs available in LPAR
 - ▶ CPU Utilization %

 - ▶ #LMB -- number of LMBs
 - ▶ freL -- number of free LMBs available in LPAR
 - ▶ Memory Utilization %
- If a LPAR is running in “shared processor mode” on a p5 server, #P represents the number of online virtual processors on that host
- Historical load is automatically saved to ~/lparLsLoads.histOut file

Sample IparLsLoads.pl Output:

CURRENT LOADs (CPU, MEMORY) -- snapshot taken every 2 seconds.

Hostname	#P	#RTh/P	% CPU Load	#LMB	#MB	#FR/s/L	% Memory Load
p5host1	1	4.000	100	64	8192	0.000	8
p5host2	2	1.500	99	64	8192	0.000	9
p5host3	3	0.000	1	13	1664	0.000	100

NOTE: 1 LMB of Memory = 128 MB

HISTORICAL LOADs -- averaged over the last 5 snapshots.

Hostname	avg #RTh/P	% cpuLoad Hist	freP	freL	Avg#FR/s/L	% memLoad Hist
p5host1	3.600	100	0	57	0.000	8
p5host2	1.500	100	0	56	0.000	9
p5host3	0.000	1	2	0	0.000	83

For definitions of the headings, see the **DEFS file** in the DOC/ subdirectory.
Press Control-C to exit!

lparSetCfgs.pl – *Time-based Resource Management*

Reconfigures multiple partitions using a single target configuration file

NOTE: This script is not intended for moving resources on partitions running in “shared processor mode” on p5 servers

- Can use **absolute values** or **delta values** in the target configuration file
 - ▶ Can also use the string "nochange" to indicate that no change is required

- In a lot of installations, one can use 2 configuration files to reconfigure the LPARs in a Managed System based on time of day
 - ▶ Can use this script in a **cron entry** to automate the daily (or weekly) cycle, moving resources between LPARs based on known load requirements

- Extensive checking:
 - ▶ Validates hostnames & LPAR names
 - ▶ Validates that resources are available within the managed system
 - ▶ Validates minimum/maximum limits for DLPAR operations
 - ▶ Removes resources before adding resources to other partitions

Sample IparSetCfgs.pl Configuration Files:

shift1 file:

# HostName	CPUs	LMBs
# =====	=====	=====
p5host1	-1	nochange
p5host2	1	8
p5host3	4	12

shift2 file:

# HostName	CPUs	LMBs
# =====	=====	=====
p5host1	2	8
p5host2	4	12
p5host3	1	8

Sample IparSetCfgs.pl Output:

CURRENT CONFIGURATIONS:

Hostname	LPARname	#CPUs	#LMBs	DLPAR
p5host1	p5lpar1	2	64	YES
p5host2	p5lpar2	3	64	YES
p5host3	p5lpar3	3	21	YES

Available (=unreserved) #CPUs in the Free Pool = 0

Available (=unreserved) #LMBs in the Free Pool = 4

NOTE: 1 LMB of Memory = 128 MB

TARGET CONFIGURATIONS:

Hostname	LPARname	#CPUs	#LMBs	DLPAR
p5host1	p5lpar1	3	68	YES
p5host2	p5lpar2	2	68	YES
p5host3	p5lpar3	2	12	YES

NOTE: 1 LMB of Memory = 128 MB

Checking if the target configurations are possible

Changing to the target configurations please wait.

```
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -d 0 -r proc -o m -p p5lpar3 -t p5lpar1 --procs 1
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -d 0 -r proc -o r -p p5lpar2 --procs 1
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -d 0 -r mem -o m -p p5lpar3 -t p5lpar2 -q 512
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -d 0 -r mem -o m -p p5lpar3 -t p5lpar1 -q 512
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -d 0 -r mem -o r -p p5lpar3 -q 128
```

IparLoadRM.pl – *Load-based Resource Management*

- Administrator selects a subset of LPARs to participate in CPU and/or Memory resource sharing in the **<hostWts> file**

NOTE: This script is not intended to be used for balancing loads on partitions running in “shared processor mode” on p5 servers

- Resources are pulled from the free pool of the CEC first, then from other partitions
- Can limit or disable use of the free pool for resource sharing by using env. vars:
 - ▶ DR_RSVD_FREP_CPUS (# reserved-CPU in free pool) not used for resource sharing
 - ▶ DR_RSVD_FREP_LMBS (# reserved-LMB in free pool)
- Resources are moved only when required
Design Point: Free resources are not given back to the free pool, so as to minimize thrashing.
- Default Trigger values** for CPU and Memory resource-movement-evaluation can be overridden by settings for individual LPARs in the (<hostWts>) configuration file
- Use **WEIGHT** settings to give relative priority to LPARs
Weights are used only when there is resource contention.
 - ▶ A weight can be any positive integer (Weights do not have to add up to 100)

lparLoadRM.pl – Usage Example:

Customer Scenario:

1. CPU-only resource sharing
2. Only share the resources already allocated to the partitions,
(no stealing of additional resources from the free pool)
3. p5host2 has a higher priority for resources than p5host1

Solution:

1) To Disable Memory Resource Sharing:

```
export DR_TRGR_MEM_LOAD=100 (before running the lparLoadRM.pl script)
```

2) To Disable Free Pool:

```
export DR_RSVD_FREP_CPUS=999 (before running the lparLoadRM.pl script)
```

3) hostsWts - Configuration File:

```
# The following line is the header line:
```

#hostName	WEIGHT	CPU_LOAD	RQ_DEPTH	MEM_LOAD	PGSTEALS
p5host1	20	95	1	85	0
p5host2	80	95	1	85	0

Sample lparLoadRM.pl Output: (1 of 3)

```
[root@control:/opt/dlparToolset/bin] ./lparLoadRM.pl -f hostsWts
```

This script is being run by user "root" on host "control".

INPUT:

```
hostsWts file           = /opt/dlparToolset/bin/hostsWts
Refresh interval       = 4 seconds
Show history data for the past = 5 refreshes
HMC Hostname           = p5hmc.ibm.com
HMC Username (for ssh) = hscroot
Managed System Name   = p5cec1
RemoteShell for AIX hosts = rsh
Reserved #CPUs in the Free Pool = 0
Reserved #LMBs in the Free Pool = 0
```

Checking ping and ssh to p5hmc.ibm.com.....

HMC Commands are running under Restricted Bash shell on POWER5 HMC (V4R2.0).....

Reading the file /opt/dlparToolset/bin/hostsWts

Checking ping and rsh to p5host1.....

Checking ping and rsh to p5host2.....

EFFECTIVE LOAD BALANCE POLICY:

#hostName	WEIGHT	CPU_LOAD	RQ_DEPTH	MEM_LOAD	PGSTEALS
p5host1	20	95	1	85	0
p5host2	80	95	1	85	0

Retrieving the current configuration from HMC

Reading the data

Sample IparLoadRM.pl Output: (2 of 3)

Note: Both partitions are now CPU-loaded to 100%, then p5host2 gets additional PLANNED resources based on the relative WEIGHT settings (see PLANNED section).

=====

CURRENT CONFIGURATIONS:

Hostname	LPARname	CPUmin	CPUcur	CPUmax	LMBmin	LMBcur	LMBmax
p5host1	p5lpar1	1	4	32	1	12	32
p5host2	p5lpar2	1	1	4	1	8	32

Available (=unreserved) #CPUs in the Free Pool = 0
 Available (=unreserved) #LMBs in the Free Pool = 12
 NOTE: 1 LMB of Memory = 128 MB

DESIRED RESOURCE CHANGES based on (min,max) values and historical Loads of CPU & memory (Negative numbers indicate unused resources):

Hostname	Change in #CPUs	Change in #LMBs	Weight
p5host1	+1	-10	20
p5host2	+3	-6	80

PLANNED RESOURCE CHANGES based on FREE POOL, WEIGHTs and historical

Loads for CPU and Memory (Negative numbers indicate unused resources):

Hostname	Change in #CPUs	Change in #LMBs	Weight
p5host1	-3	0	20
p5host2	+3	0	80

Moving the resources.....

```
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cecl -d 0 -r proc -o m -p p5lpar1 -t p5lpar2 --procs 3
```

Sample IparLoadRM.pl Output: (3 of 3)

Results after Planned Resources have been Moved to p5host2:

CURRENT CONFIGURATIONS:

Hostname	LPARname	CPUmin	CPUcur	CPUmax	LMBmin	LMBcur	LMBmax
p5host1	p5lpar1	1	1	32	1	12	32
p5host2	p5lpar2	1	4	4	1	8	32

Available (=unreserved) #CPUs in the Free Pool = 0
 Available (=unreserved) #LMBs in the Free Pool = 12

DESIRED RESOURCE CHANGES based on (min,max) values and historical Loads of CPU & memory (Negative numbers indicate unused resources):

Hostname	Change in #CPUs	Change in #LMBs	Weight
p5host1	+4	-10	20
p5host2	0	-6	80

PLANNED RESOURCE CHANGES based on FREE POOL, WEIGHTS and historical Loads for CPU and Memory (Negative numbers indicate unused resources):

Hostname	Change in #CPUs	Change in #LMBs	Weight
p5host1	0	0	20
p5host2	0	0	80

moveSlot.pl

A generic script to **move a physical IO slot (e.g. the media slot with CD/Tape, SCSI controller, Ethernet Adapter, Empty Slot)** between **DLPAR-capable AIX partitions or the free pool**

- Uses the AIX command `“/usr/sbin/rmdev”` to remove the parent and child devices attached to the slot being removed
- Requires 10/2003 HMC software (R3V2.4 or higher)
- Uses the DR_ environment variables, in particular, DR_IO_SLOT:
 - ▶ The `DR_IO_SLOT` env. variable is used to specify the **DRC Index (in POWER5) or the PhysicalLocation (in POWER4)** .
 - ▶ The script `lparLsCfgrIO.pl` can be used to locate the slot with DRC Index (in POWER5), or the PhysicalLocation (in POWER4)

moveSlot.pl performs the following steps:

1. Checks in the HMC to see if the requested slot is already assigned to another LPAR or in the free pool (unassigned)
2. If assigned, checks if the slot contains an adapter
3. If adapter, accesses the source partition to unconfigure the slot and its child devices
4. Moves the slot to the target LPAR
5. Accesses the target LPAR to configure the slot and the new devices
6. If a target LPAR is not specified, the slot is returned to the Free Pool of the CEC

Sample moveSlot.pl Output:

```
[root@control:/opt/dlparToolset/bin] ./moveSlot.pl -host p5host3
```

This script is being run by user "root" on host "control".

INPUT:

```
Host list           = /opt/dlparToolset/bin/hostList
IO Slot            = 21020003
Target Host        = p5host3
hmcHost            = p5hmc.ibm.com
hmcUser            = hscroot
Managed System    = p5cec1
RemoteShell for AIX hosts = ssh
```

Checking ping and ssh to p5hmc.ibm.com.....

HMC Commands are running under Restricted Bash shell on POWER5 HMC (V4R2.0).....

Checking ping and ssh to p5host1.....

Checking ping and ssh to p5host2.....

Checking ping and ssh to p5host3.....

Checking ping to p5host3.....

Checking ssh as a root user to p5host3.....

Finding out which LPAR contains the slot with drcIndex 21020003

The slot with DRC Index 21020003 is in free pool!

Adding the slot with DRC index 21020003 from free pool to p5host3.....

```
/usr/bin/ssh hscroot@p5hmc.ibm.com chhwres -m p5cec1 -r io --rsubtype slot -o a -p p5lpar3 -l 21020003
```

Configuring the slot with DRC index 21020003 and its devices onto the host p5host3.....

```
/usr/bin/ssh p5host3 /usr/sbin/cfgmgr
```

Additional Tools

- **nCPUs** - displays the number of CPUs in the LPAR
- **nLMBs** - displays the number of Logical Memory Blocks in the LPAR

- **onHosts** - executes an rsh/ssh command on each host in the script
- **onHostsBg** - executes an rsh/ssh command in the background, on each host in the script
- **onHostList** - executes an rsh/ssh command on a list of hosts in a file
- **onHostListBg** - executes an rsh/ssh command on a list of hosts in a file in the background

- **topp** - gives the top <n> CPU intensive processes
- **topm** - gives the top <n> memory intensive processes
- **cpuStress** - easy-to-use program for stressing CPUs
- **memStress** - easy-to-use program for stressing memory
- **kill.procs** - kills processes whose names contains the input string. Use it to kill the stress programs

- **copy_ssh_key.ksh** - can be used for copying the SSH security key from the ~/.ssh directory on the current host to the file ~/.ssh/authorized_keys2 on a remote host for a given loginID
- **lmbSz.pl** - gets the LMB size on the Managed System (SMP)

rbash (Restricted bash shell) on the HMC

- rbash is mandatory in the 10/2003 HMC software release (R3V2.4 or higher).
- DLPAR Tool Set v2.0.2.0 scripts support BOTH
 - the new HMC software releases with mandatory rbash, as well as
 - older HMC software releases that do not enforce rbash.

Feedback

Please post questions or constructive comments to the Discussion Forum at:

<http://www.alphaworks.ibm.com/tech/dlpar>

Thank you !