

AIX, VIOS and HMC Facts and Features



Ravi G Singh
Consulting IT Specialist – System p
rsingh@us.ibm.com

Release date: Nov 2007, latest updates in Red.

Notice: This document is presented "As-Is" and IBM does not assume any responsibility for the statements expressed herein. Any unannounced feature/product listed here is only a statement of direction by IBM and subject to change or withdrawal without notice. Author is not responsible for errors in this document that may result in any kind of inaccuracies.

Acknowledgement: Thanks to all the IBMers who contributed to this document and readers for their valuable feedback and comments.

Updates: The latest document can be downloaded from the URL below. One or more of the following URL links (depending on the selected audiences) will show this document.

Intranet: <http://w3.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD103130>
Internet: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD103130>
BPs: <http://partners.boulder.ibm.com/src/atmastr.nsf/WebIndex/TD103130>

Trademark and Copyright

IBM, the IBM logo, the e-business logo, System p™, System p5™, pSeries®, eServer™, OpenPower®, RS/6000®, POWER5™, POWER5+™, POWER6™, POWER™, PowerPC®, Power Architecture™, AIX®, AIX 5L™, BladeCenter®, IntelliStation®, EnergyScale™, HACMP, OS/2, PartnerWorld, pSeries, RS/6000, SecureWay, Tivoli and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both.

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

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

Microsoft, Windows and Windows NT are registered trademarks of the Microsoft Corporation.

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.

Copyright

© Copyright IBM Corp. 2007, All Rights Reserved.

Contents

AIX 5.x and 6.x	4
AIX 4.x	12
VIOS	15
HMC and Firmware	16
HMC Ports	17
HBA/FC Adapter Code	17
Notes	18
AIX 6 Preview features	26

AIX 6.x and 5.x Facts and Features

Announcement/Offering	5.1	5.2	5.3	6.1
General availability	2001-05-04	2002-10-18	2004-08-20	2007-11-06
Withdrawn from marketing	2005-04-29	TBA	TBA	TBA
Withdrawn from support	2006-04-01	2009-04-30	2010-09	TBA
Latest Technology Level (TL)	5100-09	5200-10	5300-07	6100
Latest Service Pack (SP)		SP3	SP1	SP1
Price included with H/W	Yes	Yes	No	No
License Transferable– Internal	Y-P4 & P5	Y-P4 & P5	Y-P4 & P5	
External	Y-P4 & P5	Y-P4 & P5	N-P4 & P5	

Hardware Support	5.1	5.2	5.3	6.1
POWER6	No	Yes-p570	Yes	Yes
POWER5	No	Yes	Yes	Yes
POWER4	Yes	Yes	Yes	Yes
POWER3	Yes	Yes	Yes	No
PowerPC970	Yes	Yes	Yes	Yes
RS64, 604	Yes	Yes	Yes	No
Kernel Support	32b, 64b	32b, 64b	32b, 64b	64b
UP & MP – Note 28 and 35	UP & MP	UP & MP	MP	MP
CHRP	Yes	Yes	Yes	
PCI, MCA	Yes	No	No	No
Server/Processor				
Max. Processors	32	32	64-P5	64-P5, P6
HMT	Yes	Yes	No	No
SMT	No	No	Yes-P5	Yes-P5, P6
Dynamic CPU guard		Yes	Yes	
UEGard (Uncorrectable Error)		Yes	Yes	
Dynamic Processor sparing	No	Yes	Yes	Yes
CUoD				
Permanent	Yes	Yes	Yes	Yes
Mobile - SOD	No	No	No	Yes
CoD - Temporary				
Utility CoD	No	No	Yes	Yes
Reserve CoD for shared proc pool	No	No	Yes	No
On/Off CoD post pay	Yes	Yes	Yes	Yes
Trial CoD	Yes	Yes	Yes	Yes
Capacity BackUp (CBU)	Yes	Yes	Yes	Yes
Vector Support (Altivec)			Yes	Yes
Decimal Floating Point	No	No	Yes-P6	Yes-P6
Processor instruction retry				Y-P6
Processor retry & recovery				Y-P6
Cold Node repair - SOD				Y-P6
Hot Node Add - SOD				Y-P6

Advanced POWER Virtualization	5.1	5.2	5.3	6.1
Micro-partitions max.-Note10	No	No	254-P5	254-P5, P6
Virtualization – Processors, Memory, IO, Network and storage	No	No	Y-P6,P5	Y-P6,P5
Dedicated Processor Sharing	No	No	Yes-P6	Yes-P6
Virtual processor pool or Multiple Shared processor pool	No	No	Yes-P6	Yes-P6
Virtual Partition memory or Shared Memory pool - SOD	No	No	No	Yes-P6
Virtual I/O server	No	No	Y	Y
IBM Virtualization Manager				Y
Integrated Virtualization Manager (IVM) – Note 24	No	No	Yes-P5	Y
Partition Load Manager (PLM)	No	Yes-ML F	Yes	No
Virtual Ethernet – Note 13	No	No	Yes-P5	Y
Integrated Virtual Ethernet (IVE) – Note 38	No	No	Yes-P6	Yes-P6
WPARs	No	No	No	Yes->=P4
WPAR Manager	N	N	N	Y
Live Application Mobility	No	No	No	Yes->=P4
Live Partition Mobility	No	No	Yes-P6	Yes-P6
Partition Hibernation - SOD	No	No		
Memory Affinity (relocate pages dynamically for optimum performance) - SOD	No	No		

LPAR	5.1	5.2	5.3	6.1
Dynamic LPAR-Processor, Memory and I/O	No	Yes	Yes	Yes
Multiple Instances of AIX on a Single rootvg – Note23	No	No	Yes	Yes
Maximum LPARs – Note10	32	64	254	254
Max. Virtual processors/LPAR - Note 19	No	No	64	64
Min. memory/LPAR in MB-Note12	256	256	128	
Memory incr./decrement in MB	256	256	16	
Application & LPAR Accounting	No	No	Yes	

I/O	5.1	5.2	5.3	6.1
Infiniband (PCI & GX) – Note 22	No	No	Yes-64b	Yes
PCIe	No	Y-p570	Yes	Yes
PCI-X	Yes	Yes	Yes	Yes
High Performance Switch (HPS)	No	Yes	Yes	
IP over FC, iSCSI	No	Y-ML03	Yes	Yes
MPIO	No	Yes	Yes	Yes
DMP-IO SCSI	No	Yes	Yes	
Virtual I/O and Ethernet	No	No	Yes	Yes
IO Drawer Dynamic Configuration	No	No	Yes	Yes

MEMORY	5.1	5.2	5.3	6.1
Max. Real Memory in GB-Note10	512	1024	2048-64b 96-32b	
4KB, 64KB, 16MB & 16GB Memory page – Note 26		Yes	64b- Dynamic	Dynamic
Dynamic Variable page size (4K to 64K) for data, heap, stack and shared memory	N	N	N	Y-P6
Dynamic Memory Sparing	No	No	Static-P5	
CUoD (Permanent) – P4, P5, P6	No	Yes	Yes	Yes
Reserve CoD	No	No	No	No
On/Off CoD post pay – P5, P6		Yes	Yes	Yes
Trial CoD) – P4, P5, P6		Yes	Yes	Yes
Capacity BackUp (CBU)	Permanent	Permanent	Permanent	
Address Space segments	Unlimited	Unlimited	Unlimited	
Data heap Segment size in GB – Note 33	2	3.25	3.25-32b Pbytes-64b	
Program size in 256MB segments	10	13	16-32b 64M-64b	

System Management & RAS	5.1	5.2	5.3	6.1
Dump analysis	Y	Y	Y	
Live System Dump	N	N	N	Y-P6
DVD Support for system dump	N	N	Y	
SUMA – Note 4	Y	Y	Y	
System hang detection	Y	Y	Y-Enhance	
Probevue	N	N	N	Y
Concurrent AIX Updates – eFix	N	N	N	Y
Concurrent Firmware upgrade	No	Yes-P5	Yes-P5	Yes-P5, P6
EEH	Y	Y	Y-Enhance	
FFDC enhancements - LMT, CT®, and RTEC – Note 37	N	LMT	Y	
Inventory Scout	2.1	2.2-P5	2.2-P5	
WLM- CPU, Memory, I/O - TOD Management	N N	Y Y	Y Y	
Multiple instances of AIX in one rootvg (multibos)	N	N	Y	
Thin Server & Common boot image – Note30	N	N	Y	
NIM	Y	Y	Y	Y+NFS4
GUI	N	N	N	Y
Solution Performance Tuning – Note 36	N	N	N	Y
Systems Director for AIX	N	N	N	Y
X11/CDE	5.1	5.2	5.3	
X Release	X11R6.1	X11R6.1	X11R6.1	
GNOME Desktop	N	N	Y-V2.4	

TCP/IP and networking	5.1	5.2	5.3	6.1
IP	V4/V6	V4/V6	V4/V6	
DHCP			V4/V6	
Sendmail	8.11.0	8.11.0	8.3.14 8.11.6p2	
BIND	8.8.0, 4.9.3	9.02.0, 8.8.0, 4.9.3	9.02.0, 8.8.0	
Multipath routing and DGD, VIPA	Yes	Yes	Yes	
SLP V2 and SCTP – Note 3	No	No	Yes	
SNMP	V1	V3	V3	
NFS	V2 & 3	V2 & 3	V4	
NFSV3 with Kerberos	No	No	Yes	
NFS-statd Threading	Mutli	Mutli		
AutoFs	Mutli	Mutli		
Max for thewall	Note6	Note6	Note6	
Max. no. of interfaces for ENT	1024	1024	No limit	
Network Data Administration Facility (NDAF)	No	No	Yes	
NFS Proxy Serving with Caching	No	No	Yes	

Security & Compliance	5.1	5.2	5.3	6.1
CAPP/EAL4+ - Note31	No	Yes	Yes	
LSLPP/EAL4+ - Note31		Yes	Yes	
UNIX03	No	Y-ML04	Y	
UNIX98 & C99		Yes		
ISO/IEC 9899:1999		Yes	Yes	
IEEE 1003.1-2001		Yes	Yes	
PAM		Yes	Yes	
Kerberos		V5	V5	
SecureWay IBM Dir (LDAP)	3.2.2/4.1	NA	NA	
POSIX Realtime API	No	Y-ML03	Yes	
No. of logged users max			32,767	
User/group name length (chars)	8	8	255	
Max. Group line length (chars)	16000	16000	16000	
Password length (chrs)	8	8	8	Upto 255
Role Based Access Control (RBAC)	N	N	N	Y-P6, P5, P4
Trusted AIX	N	N	N	Y
AIXPert	N	N	Y	Y+LDAP
Trusted Execution	N	N	N	Y
Secure by Default	N	N	N	Y
File Permissions Manager (fpm)	N	Y	Y	Y

LVM	5.1	5.2	5.3	6.1
Max PVs in a VG	128	128	1024	
Max LVs in a VG – Note 5	512	512	4096	
Max PPs/disk – Note 5	Note 5	Note 5	NA	
PP Size in GB	1	1	128	
Max PPs/LV – Note5	32,512	32,512	NA	
Max PPs/VG	130,048	130,048	2048K	
Max Major nos.				
Hot spare disk in a VG	Yes	Yes	Yes	
Logical Track Grp Size (LTG) KB	128-1024	128-1024	Variable	
AIO – Legacy and POSIX	Legacy	Both	Both	
Paging space Increase/Decrease	Incr only	Both	Both	
Boot image Size – Note 27	12MB	12MB	32MB	

FILE SYSTEM	5.1	5.2	5.3	6.1
Filesystem Architectural Max	4PB	J2-4PB J-16TB/64b J-1TB/32b	J2-4PB	
Filesystem Max tested – Note 20	J2-1TB J-1TB	J2-16TB J-1TB	J2- 32TB	
File Size Architectural Max		J2-1PB J-64GB	J2-1PB J-64GB	
File size Max tested	J2-1TB J-64GB	J2-16TB J-64GB	J2 > 16TB	
File Descriptors or open files	3,145,728	3,145,728	3,145,728	
File Descriptors/process	32,768	65,534		
Max. sub directories/dir–J & J2	32, 766	32, 766	32, 766	
Max file system inodes	Fixed	J2- Dynamic	J2- Dynamic	
JFS and JFS2	Both	Both	Both	
JFS2 Quota	No	No	Yes	
JFS2 Shrink	No	No	Yes	
RAM disk or cachefs in GB	2	> 2		
Outline JFS2log Max	1GB	Dynamic 32b-1GB 64b-64GB	Dynamic	
Inline JFS2log Max	32MB	256KB to 16GB		
CDROM FS- cdromd DVDRAM FS	Yes	Automount Yes	Automount Yes	
Freeze and Thaw – Note 29	No	No	Yes	
Encrypted Filesystem	N	N	N	Y-J2

SEMAPHORES	5.1	5.2	5.3	6.1
Max Semaphore IDs	131,072	131,072	128K-32b 1024K-64b	
Max Semaphores per semph. ID	65,535	65,535	65,535	
Semaphore Max value	32,767	32,767	32,767	
MESSAGE QUEUES				
Max Message size	4MB	4MB	4MB	
Max bytes on queue	4MB	4MB	4MB	
Max message queue IDs	131,072	131,072	128K-32b 1024K-64b	
Max messages per queue ID	524,288	524,288	524,288	
SHARED MEMORY				
Max Segment size , 32b process	2GB	2GB	2GB	
Max Segment size , 64b process	64GB	1TB	1TB-32b 32TB-64b	
Max Shared memory IDs	131,072	131,072	128K-32b 1024K-64b	
Max segments/process 64b	268435456	268435456	268435456	
Max segments/process 32b	11	11	11	
Min. segment size	1	1	1	

KERNEL/PROCESS	5.1	5.2	5.3	6.1
Max Threads/process	32,767	32,767	32,767	
Max Threads/System				
Max processes/user				
Max Processes/system		256,000		
Memory/Storage Keys Application	N	N	Y-P6	Y-P6
Memory/Storage Keys Kernel	N	N	N	Y-P6
Resources (/etc/security/limits)				
CPU time		2 [^] 31-1	2 [^] 31-1	
File size		4,194,303	2 [^] 31-1	
Stack Size		523,264	2 [^] 31-1	
Memory Size		2 [^] 31-1	2 [^] 31-1	
Open Files		2 [^] 31-1	2 [^] 31-1	
Data Size		4,194,303	2 [^] 31-1	
Core Size		2 [^] 31-1	2 [^] 31-1	
IO Pacing tuning level	System	System	Filesystem	
Max. no. of devices – Note 32		5K	25K	
ksh	ksh93/ ksh88	ksh93/ ksh88	ksh93/ ksh88	
Core File naming	core.pid. date,	core.pid. date,	core.pid. date,	
Enhancements	snapcore	snapcore	snapcore	

Expansion pack/Web download /Documentation – Note9	5.1	5.2	5.3	6.1
RSCT	NA	2.3.9.0 2.3.7.1	2.4.5.0 2.4.3.1	
OpenSSH	3.8.0.0	4.3.0.52 4.1.0.52, 3.8.0.52, 3.7.0.52	4.3.0.53 4.1.0.53, 3.8.0.53	
Adobe Acrobat Reader	5.09	5.09	N/A	
IBM HTTP Server 128-bit	1.3.19.4	2.0.47.1, 1.3.19.4	2.0.47.1	
Netscape Communicator 128-b	4.8	4.8/7.0	7.0	
Mozilla	1.7.5, 1.4.2	1.7.13, 1.7.5, 1.4.2	1.7.13, 1.4.2	
Mozilla FireFox 64b		1.5.0.6	1.5.0.6	
Cryptographic Library	NA	5.2-256-bit	5.3	
Perl	5.6.0	5.8.0	5.8.2	
Documentation Library	Y	Y	Y	
AIX Information Center	Y	Y	Y	
IBM Director Server	N	5.1	5.2, 5.1	
IBM Director Console	N	5.1	5.1	
IBM Director Agent	N	5.2, 5.1	5.2, 5.1	
IBM System Director Console	N	N	N	Y
Active Energy Manager (AEM)	N	N	3.1-P6	3.1-P6
Service Agent	3.3.0.0	3.3.0.0	3.3.0.0	
JDK				
Supported Version/Release	142, 131	5, 142, 131	5, 142, 131	
Unsupported Version/Release	141,140, 130, 122, 118	141,140, 130, 122, 118	141,140, 130, 122, 118	

Licensed Software Support	5.1	5.2	5.3	6.1
VisualAge C++	6.0.0/5.0.2	6.0.0/5.0.2	6.0	
XL C/C++	8.0, 7.0	9.0, 8.0, 7.0	9.0, 8.0, 7.0	
C for AIX	6.0.0/5.0.2	6.0.0/5.0.2	6.0	
XL C	8.0, 7.0	9.0, 8.0, 7.0	9.0, 8.0, 7.0	
XL Fortran	10.1, 9.1, 8.1, 7.1.1	11.1, 10.1, 9.1, 8.1, 7.1.1	11.1, 10.1, 9.1, 8.1.1	
COBOL	2.0	2.0	2.0	
PL/1	2.0	2.0	2.0	
AIXlink/X.25	2.1	2.1, 2.0.2	2.1, 2.0.2	
AIX Fast Connect	3.2.0, 3.1.2	3.2.0, 3.1.2	3.2.0	
Communication Server	6.1	6.3, 6.1	6.3, 6.1.0.5	
DCE	3.1-32b	3.2, 3.1-32b	3.2	

DFS			3.1	
CSM-Mgmt. Server – Note25 & 25A	N	1.6, 1.5, 1.4.1, 1.3.3	1.6, 1.5, 1.4.1, 1.3.3	
CSM-managed Nodes – Note 25 & 25A	>=1.3	1.6, 1.5, 1.4.1, 1.3.3	1.6, 1.5, 1.4.1, 1.3.3	
GPFS Ver Max Filesystem size	2.1 Note 7	3.1,2.3, 2.2, 21 Note 7	3.1, 2.3 Note7	
PSSP – Note17	3.4 or 3.5	3.5	No	
HACMP	5.2, 5.1, 4.5, 4.4.1,	5.4, 5.3, 5.2, 5.1, 4.5	5.4, 5.3, 5.2, 5.1	5.4, 5.3,
eWLM Server Agent	No	No 2.1	2.1 2.1	
System Provisioning			Yes-P5	
Grid Tool Kit		3	3.3	

AIX 4.x Facts and Features

Announcement/Offering	4.1.4/5	4.2.0	4.2.1	4.3.1	4.3.2	4.3.3
General availability	25-Apr-97		25-Apr-97	24-Apr-98	23-Oct-98	17-Sep-99
Withdrawn from marketing	31-Dec-98		31-Dec-99		Sep-99	30-Jun-03
Support withdrawn	31-Mar-99		31-Mar-00	31-Dec-01	31-Dec-01	31-Dec-03
Latest Maintenance Level (ML)						4330-11
Hardware Support						
CHRP	Y	Y	Y	Y	Y	Y
PCI, MCA	Y	Y	Y	Y	Y	Y
LVM						
Max PVs in a VG	32	32	32	32	128	128
Max LVs in a VG	256	256	256	256	512	512
Max PPs/disk	1016	1016	Note 5	Note 5	Note 5	Note 5
Max PPs/LV						
Max Major nos.	256	256	256	256	256	
Max VMM Mapped devices	512	512	512	512	512	
Hot spare disk in a VG						No
Logical Track Grp Size (LTG) KB						128
AIO	Legacy	Legacy	Legacy	Legacy	Legacy	Legacy
Paging space Increase/Decrease	Incr only	Incr only	Incr only	Incr only	Incr only	Incr only
FILE SYSTEM						
Filesystem Architectural Max	64GB	128GB	1TB	1TB		1TB
Filesystem Max tested						1TB
File Size Architectural Max	2GB	64GB	64GB	64GB	64GB	64GB
File size Max tested						64GB
File Descriptors or open files	200,000	200,000	200,000	1000,000	1,048,576	1,048,576
File Descriptors/process	2,000	2,000	2,000	32,767	32,767	32,768
Max file system inodes			2 ²⁴	2 ²⁴		
Max file system fragments			2 ²⁸	2 ²⁸		
JFS and JFS2	JFS	JFS	JFS	JFS	JFS	JFS
RAM disk or cachefs in GB						2
Outline JFSlog Max	256MB	256MB	256MB	256MB	256MB	1GB
Inline JFSlog Max						32MB
MEMORY						
Real Memory	2GB	4GB	4GB	16GB	32GB	96GB
Size of Executable in Bytes				7*2 ⁵⁶		
Size of Text, Data & BSS in Bytes				7*2 ⁵⁶		
Symbol Values (address)	2 ³²	2 ³²	2 ³²	2 ⁶⁴		
Address Space segments	256MBx16	256MBx16	256MBx16	256MBx16	256MBx16	256MBx16
Data heap Segment size	2GB	2GB	2GB	2GB	2GB	2GB
Program size in 256MB segments	10	10	10	10	10	10
SEMAPHORES						
Max Semaphore IDs	4,096	4,096	4,096	4,096	131,072	131,072
Max Semaphores per ID	65,535	65,535	65,535	65,535	65,535	65,535
Semaphore Max value	32,767	32,767	32,767	32,767	32,767	32,767

Announcement/Offering	4.1.4/5	4.2.0	4.2.1	4.3.1	4.3.2	4.3.3
MESSAGE QUEUES						
Max Message size	65,535	65,535	65,535	65,535	4MB	4MB
Max bytes on queue	65,535	65,535	65,535	65,535	4MB	4MB
Max message queue IDs	4,096	4,096	4,096	4,096	131,072	131,072
Max messages per queue ID	8,192	8,192	8,192	8,192	524,288	524,288
	4.1.4/5	4.2.0	4.2.1	4.3.1	4.3.2	4.3.3
SHARED MEMORY						
Max Segment size , 32-bit	256MB	256MB	256MB	2GB	2GB	2GB
Max Segment size , 64-bit	NA	NA	NA	2GB	2GB	2GB
Max Shared memory IDs	4,096	4,096	4,096	65,538	131,072	131,072
Max segments/process 64b						268435456
Max segments/process	10	10	10	11	11	11
Min. segment size	1	1	1	1	1	1
KERNEL/PROCESS						
Max Kernel Threads/process	512	512	512	32,767	32,767	32,767
Max Threads/System			262,143	262,143		512,000
Max processes/user	65,538	65,538	131,072	131,072	131,072	
Max Processes/system	131,072	131,072	131,072	131,072	131,072	170,000
Max concurrent groups/process	32	32	32	32	32	32
Max for thewall	64MB	64MB	64MB	128MB	Note6	Note6
ksh			ksh88	ksh88	ksh88	ksh88
Core File naming	core	core	core	core	core	core
TCP/IP and networking						
IP	V4	V4	V4	V4/V6	V4/V6	V4/V6
Sendmail		8.7	8.7	8.8.8	8.8.8	8.9.3
BIND			4.9.3	8.1.1	8.8	
Multipath routing and DGD, VIPA	No	No	No	No	No	No
SNMP	V1	V1	V1	V1	V1	V1
NFS-statd Threading						Single
AutoFs						Single
X11						
X Release						X11R6.1
Security						
C2		Yes	Yes	Yes	Yes	Yes
JDK						
Supported Version/Release						131, 130, 122, 118
New Features/Enhancements						
WLM- CPU, Memory I/O						Y N
Bonus and Expansion Pack						
Adobe Acrobat Reader						4.05
IBM HTTP Server 128-bit						1.3.12.4
Netscape Communicator 128-b						4.79
JDK						1.3.1
Cryptographic Library						5.5.3.75

Licensed Software Support						
VisualAge C++						6.0/5.0/3.6
C for AIX						8.1/7.1.1
XL Fortran						3.1.1
AIXlink/X.25						6.1
AIX Fast Connect						
Communication Server						
DCE						

VIOS Facts and Features

Features	V1.1	V1.2	V1.3	V1.4	V1.5
General availability	2005.04	2005.10	2006.07	2007.06	2007.11
Latest Fix pack					1.1
Hardware Support					
P6	N	N	N	Y	Y
P5	Y	Y	Y	Y	
P4, P3	N	N	N	N	N
JS21	N	Y	Y	Y	
AIX	5.3	5.3	5.3	>=5.3	>=5.3
Virtualization					
Dedicated LPAR	Y	Y	Y	Y	Y
Shared Processor LPAR	Y	Y	Y	Y	Y
Multiple Shared Processor pools	Y	Y	Y	Y-P6	Y-P6
Shared Dedicated processor Capacity	Y	Y	Y	Y-P6	Y-P6
Integrated Virtualization Manager – Note 38	Y	Y	Y	Y	Y
IVM – DLPAR (Proc & Memory)	N	N	Y	Y	
Shared Ethernet	Y	Y	Y	Y	Y
SEA Failover	N	Y	Y	Y	Y
VSCSI	Y	Y	Y	Y	Y
Virtual Ethernet	Y	Y	Y	Y	Y
GVRP (GARP VLAN Registration Protocol)	N	N	N	Y	Y
Live Partition Mobility	N	N	N	Y	Y
Maximum VIOS	10	10	10	10	10
LDAP	N	N	N	Y	Y
SNMP	N	N	N	Y	Y
SPT	N	N	N	Y	Y
CIM (Common Inf. Model)-SOD	N	N	N	Y	Y
System Management					
Performance Management Topas & viostat	N	N	Y	Y	Y
IBM Tivoli Integration (TSM, ITUAM, TADDM, ITM)	N	N	N	Y	Y
Browser based HMC GUI	N	N	N	Y	Y
Storage Support nSeries, NetApp, iSCSI and Fibre Channel attach	N	N	N	Y	Y
SAS and SATA	N	N	N	Y	Y

IBM Tivoli Integration:

IBM Tivoli Monitoring
 IBM Tivoli Storage Manager
 IBM Tivoli Usage and Accounting Manager
 IBM Tivoli Application Dependency Discovery Manager

HMC and Firmware Facts & Features

HMC Code	V3R3.x	V4R5.x	V5R1.x	V5R2.x	V6R1.x	V7R3.x
General availability	2005.07.12	2005.06.06	2005.10.14	2006.02.28	2006.08.22	2007.06.10
Latest Maintenance Level (ML)	R3.7 2006.03.31	R5.0 2005.06.06	R1.1 2006.01.12	R2.1 2006.04.19	R1.3 2007.09.26	R3.2.0 2007.11.28
Max. non P5-59x servers	16	48	48	48	48	Inf. not available
Max. managed systems	32	32	32	32	32	available
Max. LPARs – Note18	64 for 7315	254	254	254	254	
Hardware Support						
7310-C06 for P5					Y	Y - Opt
7042-CR4 & C06 –P6, P5					N	Y
7310-CR4 for P5					Y	Y – Opt
7310-C04 for P5		Y	Y	Y	Y	Y – Opt
7310-CR3 for P5		Y	Y	Y	Y	Y – Opt
7310-C03 for P5		Y	Y	Y	Y	Y – Opt
7310-CR2 for P5		Y	Y	Y	Y	Y - Opt
7315-C04 for P4	Y					N
7315-CR3 for P4	Y					N
7315-C03 for P4	Y					N
7315-CR2 for P4	Y					N
7315-C02		Y	Y	Y	Y	N
7315-C01		Y	Y	Y	Y	N
6878 V3						N
POWER6 (P6)						Y
POWER5 (P5)	N	Y	Y	Y	Y	Y
POWER4 (P4)	Y	N	N	N	N	N
CFM Support – Note 21	N	Y	Y	Y	Y	
Recommended - Sys FW BPA		SF230 BP230	SF235 BP235	SF240 BP240	SF240 BP240	SF240 BP240
Browser based GUI	N	N	N	N	N	IE >=6.0 or FF >= 1.506

HMC Support URLs

<http://www14.software.ibm.com/webapp/set2/sas/f/hmc/power5/download/v61.Update.html>

<http://www14.software.ibm.com/webapp/set2/sas/f/power5cm/supportedcode.html>

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Firmware Levels

- GA3-SF222, GA4-SF225, GA5-SF230, GA6-SF235, GA7-SF240 (last release for P5).
- Proxy-HTTP support for call-home is provided in 01SF240-258 (GA7SP3) and HMC V6.1.

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

<http://www14.software.ibm.com/webapp/set2/sas/f/power5cm/home.html>

Power 5 HMC port information

Port	Protocol	Application	Enabled by Default	Modification allowed in Network Config	Security	Notes
9090, 9940, 30000 -to- 30009	tcp tcp	WebSM WebSM	no	yes	SSL	Listen only Pair of ports used
22 1024-65535	tcp tcp	ssh ssh	no	yes	3DES	HMC Admin WS
512 & 514 1024-65535	tcp tcp	rexec rexec				HMC Admin WS
1024-65535 808	tcp tcp	Inventory Scout				HMC LPARs
80	tcp	http	yes	yes		
443	tcp	https	yes	yes	SSL	
657 697 1697	udp/tcp udp/tcp	RMC RMC	yes	yes		DLPAR LPARs ???
9,920	tcp	FCS	yes	yes		Call Home
9,900	udp	FCS	yes	yes		Call Home
4,411	tcp	Web server	yes	yes		InfoCenter (V4)
4,412	tcp	Web server	yes	yes		InfoCenter (V5)
9,443	tcp	Secure Web Server	yes	yes	SSL	Remote ASM(V5)
9,735	tcp	Vtty	yes	yes		
2300, 2301	tcp	I5250 console	yes	yes	2301 - SSL	
6,000	tcp	X11	yes	no	Xhost -	
5,988	tcp	CIM	yes	yes	SSL	CIMOM provides inf. to CIM clients
9197, 9198	tcp	CIM Indication	yes	yes	SSL	
123	udp	NTP	no	yes		
1,701	udp	I2tp	yes	yes		
427	udp	SLP	no	yes		Used in Cluster
2,049	tcp	NFS	no	no		
69	tcp	TFTP	no	no		
n/a	icmp	Ping	yes	yes		
500, 4500	udp	IPSec	no	no		VPN

HBA/FC Adapter Code

FC	Gb/Sec	Ports	Part #	ROS	Z9	ZA	Z8
5773	4	2	PCIe x4				
5774	4	1	PCIe x4				
1905 5758	4 4	1 1	03N5005, 03N5014 03N5014	02C82135	BS2.10X8	BS2.10XB	C9488F4C
1910 5759	4 4	2 2	03N5029	02C82132	BS2.10X8	BS2.10XB	C94ABA08
5716 1957 1977	2 2 2	1 1 1	80P4543,03N6441,03N7069 03N4698,03N6440,03N7068 80P6101, 80P6455, 03N6439, 03N7067	02881914	TS1.91A5	T1D1.91A5	C93AF096
6239	2	1	00P4295, 80P4381, 80P6415	02E01971	HS1.91X4	H1D1.91X4	C9332A79
6228	2	1	03N2452, 09P0102, 00P2995, 09P5079, 00P4494, 09P5080, 80P3388, 80P4383	02C03951	CS3.91X4	C1D3.91X4	C2D3.91X4
6227	1	1	09P4038, 09P1162, 03N4167, 24L0023	02903331	SS3.30X1, SS3.22A1	S1F3.30X1, S1F3.22A1	S2F3.30X1, S2F3.22A1

Notes

1. SOD – Statement of Direction, DGD-Dead Gateway Detection, VIPA-Virtual IP Address, WLM-Work Load Manager.
2. J-JFS, J2-JFS2, HMT-Hardware Multi Threading, SMT–Simultaneous Multi Threading.
CAPP/EAL4+: Controlled Access Protection Profile and Evaluation Assurance Level 4+ (CAPP/EAL4+)
3. SLP – Service Location Protocol, SCTP – Stream Control Transmission Protocol
4. SUMA – System Update Management Assistant for policy based automatic download of updates.
5. Max PPs/disk: It is a combination of PVs in a VG. Alternatives are 32 disks with 1016 PPs, 1 disk with 1016*32 PPs, 16 disks with 1016*2 PPs.

In AIX 53, Scalable VG removes the limit on the no. of PPs in a PV instead the limit is at the VG level.

VG type	Maximum PVs	Maximum LVs	Maximum PPs per VG	Maximum PP size
Normal VG	32	256	32512 (1016 * 32)	1 GB
Big VG	128	512	130048 (1016 * 128)	1 GB
Scalable VG	1024	4096	2097152	128 GB

6. Max value for thewall : AIX5.1 and later – smaller of ½ of RAM or 64GB for 64b kernel, smaller of ½ of RAM or 1GB for 32b kernel, 1GB or half the memory for CHRP, 256MB or ½ the memory for non CHRP.
7. Tested max. GPFS filesystem size: depends on the block size of the filesystem
 - 16 KB block size, one or more filesystems with a total size of 1 TB mounted.
 - 64 KB block size, one or more filesystems with a total size of 10 TB mounted
 - 256 KB or greater block size, one or more filesystems with a total size of not greater than 200 TB where no single filesystem exceeds 100 TB mounted.
 - GPFS 2.3 or later, file system architectural limit 2⁴⁹ bytes, GPFS 2.2 file system architectural limit 2⁵¹ bytes (2 Petabytes) Current tested limit approximately 2 PB.
 - No. of filesystems: GPFS v3.1.0.5 or later 64, GPFS v3.1.0.1 thru v3.1.0.4 32 GPFS v2.3 all service levels 32.
 - No. of files in a filesystem: The architectural limit of the number of files in a file system is determined by the file system format. For file systems created prior to GPFS V2.3, the limit is 268,435,456. For file systems created with GPFS V2.3 or later, the limit is 2,147,483,648. Please note that the effective limit on the number of files in a file system is usually lower than the architectural limit, and could be adjusted using the **-F** option of the mmchfs command.
 - Disk size: The maximum disk size supported by GPFS depends on file system format and the underlying device support. For file systems created prior to GPFS version 2.3, the maximum disk size is 1 TB due to disk format limitations. For file systems created with GPFS 2.3 or later, these limitations have been removed, and the maximum disk size is only limited by the device driver support.
AIX 5L with 64-bit kernel, GPFS supports disks larger than 2 TB (provided the disk device supports it), up to the operating system limit. On other supported platforms, GPFS supports disks up to 2 TB in size.
8. IBM's intent is to offer this feature in future and a statement of direction only.
9. Bonus Pack withdrawn as of May 21, 2004, products that previously existed on the AIX 5L Bonus Pack now resides on the AIX 5L Expansion Pack or the Web Download Pack.
10. Maximum limit depends on the pSeries server Type-Model, no. of processors and memory configured.
11. Designed to be compliant.
12. If system firmware and HMC code is pre 10/2002, then minimum physical memory for each partition is 1 GB.
13. VELAN/VEL - Virtual Ethernet LAN, SEA-Shared Ethernet Adapter in a VIO, No. of SEA/VIO – unlimited, 16 VE trunk adapter/SEA, 20 VLAN/VE trunk adapter, 320 VLAN/physical ENT, 256 Virtual Ethernet connects/LPAR.

17. PSSP 3.5 support is withdrawn effective Apr 30, 2008.
 18. Maximum limit depends on the pSeries server Type-Model, pl. refer to the announcement letters or HMC code information URL for specifics.
 19. Physical, Logical and Virtual Processors:

Processor Terminology	Dedicated Partition	Micro-Partition
Physical Processors (PhPr) in the server or CEC.	Total no. of Processors configured in the server including COD option, LiPr+COD processors in the server. <i>Example: If a server is configured with a total of 64 processors, 32 Licensed and 32 COD. PhPr is 64.</i>	Total no. of Processors configured in the server including COD option, LiPr+COD processors in the server. <i>Example: If a server is configured with a total of 64 processors, 32 Licensed and 32 COD. PhPr is 64.</i>
Licensed Processors (LiPr) in the server or CEC.	No. of Processors Licensed, max 64. <i>Example: A server configured with a total of 64 processors, 32 Licensed and 32 COD. LiPr is 32.</i>	No. of Processors Licensed, max 64. <i>Example: A server configured with a total of 64 processors, 32 Licensed and 32 COD. LiPr is 32.</i>
Shared Processors (ShPr) in a pool or Dedicated Processors (DePr) in a dedicated LPAR	No. of LiPr allocated to a dedicated LPAR, max 64. <i>Example: Out of 32 LiPr, if 8 processors allocated, DePr is 8.</i>	No. of LiPr allocated to a shared processor pool, max 64. <i>Example: Out of 32 LiPr, if 8 processors allocated, ShPr is 8.</i>
Virtual Processors (ViPr) in a LPAR	Ratio of 1 to 1 with the DePr, can range from one to DePr. <i>Example: If DePr is 8, ViPr can be from 1 to 8. Let ViPr be 8.</i>	Ratio of 10 to 1 with the ShPr, can range from 1 to ShPr upto a max. of 64. <i>Example: If ShPr is 8, ViPr can be from 1 to 64. Let ViPr be 8.</i>
Logical Processors SMT on (LoPr-On) in a LPAR	Ratio of 2 to 1 with ViPr, LoPr-On is 2xViPr upto a max of 128. <i>Example: DePr is 8, LoPr-On is 16.</i>	Ratio of 2 to 1 with ViPr, LoPr-On is 2xViPr upto a max of 128. <i>Example: ViPr is 8, LoPr-On is 16.</i>
Logical Processors SMT off (LoPr-Off) in a LPAR	Ratio of 1 to 1 with ViPr, LoPr-Off is 1 x ViPr upto a max of 64. <i>Example: ViPr is 8, LoPr-Off is 8.</i>	Ratio of 1 to 1 with ViPr, LoPr-Off is 1 x ViPr upto a max of 64. <i>Example: ViPr is 8, LoPr-Off is 8.</i>
Entitled Capacity	DePr	ShPr
No. of processors in the LPAR (No. of procs shown in lscfg or lsdev commands)	ViPr	ViPr

20. JFS2 filesystem size depends on the block size.

Filesystem Block Size	Max. Filesystem Size in TB
512	4
1024	8
2048	16
4096	32

21. HMC V4 R5.0 and FW 01SF230_120_120 are the minimum code levels required to support CFM and IO Drawer concurrent maintenance. P5-570 and 59x also require power subsystem code at 02BP230_125 prior to upgrading LIC.

22. IB HCA Speeds & Feeds

1x-2.5 Gb/sec, 4x-10 Gb/sec, 12x-30Gb/sec, DDR-60Gb/sec, QDR-120Gb/sec
 IBM IB HCAs - PCI-4x only, GX-4x & 12x with Virtualization across 64 LPARs. There are two InfiniBand device drivers: one for the GX bus and the one for the PCIX bus. Both of these device drivers support only 64-bit kernel mode. Concurrent mode diagnostic support for the PCIX adapter is not provided.

23. Multiple Instances of AIX on a Single Root Volume Group In AIX 5.3, the root user can create multiple instances of AIX on a single root volume group (rootvg). A new utility, /usr/sbin/multibos, is supplied in AIX 5L with 5300-03 to create and manage a new instance of the operating system within the running rootvg. The multibos utility provides the root user operations to setup, access, maintain, update, and customize this new instance of the Base Operating System (BOS). The result of creating a new instance of the BOS with multibos is a rootvg with two distinct and bootable instances of the operating system within a single rootvg. The running instance, called the active BOS, can be in production while multibos operations are used to modify the non-running instance, called the standby BOS. The multibos command and corresponding man page in the *AIX 5L Version 5.3 Commands Reference* incorrectly lists the supported level for multibos as 5300-02. You must run multibos with maintenance level 5300-03. For more detailed information, refer to the latest /usr/lpp/bos/README.multibos file, the multibos man page, and documentation regarding multibos in the AIX Information Center.

24. IVM – Integrated Virtualization Manager

IVM may be used to complete the following tasks:

- Create and manage logical partitions
- Configure the virtual Ethernet networks
- Manage storage in the Virtual I/O Server
- Create and manage user accounts
- Create and manage serviceable events through Service Focal Point
- Download and install updates to device microcode and to Virtual I/O Server software
- Back up and restore logical partition configuration information
- View application logs and the device inventory

Restrictions and Limitations

Because the Integrated Virtualization Manager provides a subset of the HMC functionality, there are some key restrictions and limitations that are worth highlighting.

- Full dynamic LPAR support for VIOS partition only:** Dynamically adding or removing memory or processing resources from a running client partition is not supported with the IVM. The partition should be powered off first. Keep in mind that the POWER5 Hypervisor allows partitions to use more than their entitled processing capacity via the shared processing pool, lessening the importance of processing dynamic LPAR in most environments.
- All physical I/O is owned by the VIOS partition:** This statement means that the VIOS partition is a single point of failure - if it fails, all client partitions will also fail as their virtual disk, optical, and Ethernet devices will not be accessible.
- Limited service function:** There is no call-home support, and concurrent maintenance support is limited to adapters owned by the VIOS.
- No redundant VIOS partition support:** Because all physical I/O must be owned by a single partition, it is not possible to have more than one VIOS partition for redundancy.

25. Limitations

Scaling limitations: Your CSM cluster must adhere to all of the following applicable limits.

Your cluster size limit will be the smallest number of all applicable limits. Support for clusters above such limits will be considered via Special Bid.

Maximum number of operating system images:

- For supported pSeries AIX 5L or Linux systems, the cluster scaling limit is 512 operating system images.
- For supported xSeries based Linux systems, the cluster scaling limit is 1024 operating system images.
- For combined environments, the limit is 1024 operating system images, with no more that 512 of those images being on pSeries systems.

Cluster limits for all platforms and situations:

Situation	Maximum number
LPARs per Cluster	512
Servers per Cluster	128

Number of LPARs per HMC	128
Number of eServer(R) servers per HMC	32

Cluster limits for specific processor types: with Industry Standard Switch Interconnects Hardware
Max.number

p690 servers per cluster	32
p670 servers per cluster	32
p655 servers per cluster	64
p650 servers per cluster	64
p630 servers per cluster	64
p615 servers per cluster	64
System p5 and eServer p5 servers per cluster except	128
p5-590 servers per cluster	16
p5-595 servers per cluster	16

Other limitations

CSM for AIX 5L:

- The management server needs to be at the highest CSM level of any node in the cluster.
- CSM for AIX 5L does not run on the RS/6000 SP Switch or SP Switch2.
- PSSP and CSM will not coexist in the same cluster.

CSM for Linux on POWER:

- The management server needs to be at the highest CSM level of any node in the cluster.
- SUSE Linux ES 9 and Red Hat 3 or 4 (ES/AS) are the only Linux distributions supported.
- Note: To see if CSM runs on the desired combination of Linux Distributions and hardware, refer to the Frequently Asked Questions section at
- <http://techsupport.services.ibm.com/server/csm/documentation>
- The p655 should not be used as a management server.

CSM for Linux Multiplatform

- The following operating systems are supported on the cluster management server:
- Red Hat EL 3 (AS/ES/WS)
- Red Hat EL 4 (AS/ES/WS)
- SUSE Linux ES 9
- Additional information is available at
- <http://techsupport.services.ibm.com/server/csm/documentation>
- The management server needs to be at the highest CSM level of any node in the cluster.
- There is limited hardware control support for the x440.
- Professional versions of Linux are not supported with CSM V1.5.
- For xSeries, if you do not have a terminal server, the remote console function (rconsole) will not work and automatic MAC address collection will not work for the full installation of Linux nodes.

25A. Software requirements for CSM 1.6 for AIX 5L

For all AIX 5L servers, the following AIX 5L levels are required for CSM V1.6:

- AIX 5L V5.3: Technology Level 5, with RSCT V2.4.6
- AIX 5L V5.2: Technology Level 9, with RSCT V2.3.1

The CSM management server must be running AIX 5L V5.3, AIX 5L V5.2, SLES, or RHEL, with all recommended maintenance packages. The other machines within the cluster are referred to as managed nodes and can be running any supported edition of CSM, AIX 5L, or Linux. Certain CSM for AIX 5L functions require non-IBM software. The following non-IBM software is required: · openCIMOM V0.7 or 0.8 This should be used if you want to perform remote hardware control operations for System p servers attached with an HMC. Obtain the software from the “AIX toolbox for Linux applications” CD (refer to Software Announcement 201-090, dated April 17, 2001), or to obtain the download, visit <ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/>

For your convenience, the following required IBM and non-IBM software is included on the CSM CD-ROM: · conserver 8.1.

CSM for Linux on POWER

- The default base versions of SLES 9 or SLES 10 for POWER systems, as distributed and supported by Novell Corporation
- The default base version of RHEL AS 4 for POWER, as distributed and supported by Red Hat Certain CSM for Linux on POWER functions require non-IBM software. The following non-IBM software is required and can be obtained from the listed sources:
- AutoUpdate V4.3.4, or later levels

This should be used if you want to perform the software maintenance installation and upgrade of non-CSM RPMs on your Linux-managed nodes from the management server. Download the software at <http://freshmeat.net/projects/autoupdate>

For your convenience, the following required IBM and non-IBM software is included on the CSM CD-ROM:

- tftp-HPA 0.34
- fping-2.4b2-5
- conserver 8.1
- IBMJava2-JRE 1.4.2

The following non-IBM software is required if you want to perform remote hardware control operations for POWER servers attached with an HMC:

- openCIMOM V0.7 or 0.8. To download, visit <ftp://ftp.software.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/>

The following non-IBM software is required if you want to install Linux diskless nodes:

- Warewulf V2.6.2, or later, open source clustering solution software, available from <http://www.warewulf-cluster.org>
- Yum automatic updater and package installer/remover for RPM systems, available from <http://www.linux.duke.edu/projects/yum>

CSM Highly Available Management Server (HA MS) feature (all platforms)

CSM HA MS requires the same operating systems and level as the base CSM.

For CSM V1.6 on SLES 10 on System x, an additional SLES service fix is required to use HA MS. Refer to Bugzilla defect number 26256.

The shared disks supported by HA MS do not always support these operating systems. Use of HA MS is restricted to operating systems supported by the chosen shared disk.

CSM HA MS also requires that CSM V1.6 be installed on both management servers. CSM HA MS does not require the installation of any other HA product. All required software is included in CSM and HA MS.

Service

Service for CSM for AIX and Linux can be obtained at

<http://www14.software.ibm.com/webapp/set2/sas/f/csm/home.html>

Limitations

Scaling limitations

Your CSM cluster must adhere to all of the following applicable limits. Your cluster size limit will be the smallest number of all applicable limits. Support for clusters above such limits will be considered via special bid.

Maximum number of operating system images:

- For supported System p AIX 5L or Linux systems, the cluster scaling limit is 512 operating system images.
- For supported System x based Linux systems, the cluster scaling limit is 1,024 operating system images.
- For combined environments, the limit is 1,024 operating system images, with no more than 512 of those images being on System p systems.

Cluster limits for System p Situation Maximum number

- LPARs per cluster 512
- Number of LPARs per HMC 128
- Number of IBM eServer servers per HMC 32

Cluster limits for specific processor types Hardware Maximum number

- p690 servers per cluster 32
- p670 servers per cluster 32
- p655 servers per cluster 64
- p650 servers per cluster 64
- p630 servers per cluster 64
- p615 servers per cluster 64
- System p5 and IBM eServer p5 128 servers per cluster except p5-590 servers per cluster 16
- Other limitations

CSM for AIX 5L

- AIX 5L V5.1 is no longer supported.
- The management server needs to be at the highest CSM level of any node in the cluster.
- CSM for AIX 5L does not run on the RS/6000 SP Switch or SP Switch2.
- PSSP and CSM will not coexist in the same cluster.

CSM for Linux on POWER

- The management server must be at the highest CSM level of any node in the cluster.
- SLES 9 and 10, and RHEL 4 AS are the only Linux distributions supported.

Note: To see if CSM runs on the desired combination of Linux Distributions and hardware, refer to the Frequently Asked Questions section at <http://www14.software.ibm.com/webapp/set2/sas/f/csm/home.html>

- Support for diskless Linux nodes requires that all such nodes have the same version and release of Linux.

CSM for Linux Multiplatform

- The following operating systems are supported on the cluster management server:
 - RHEL 4 (AS/ES/WS) & - SLES 9 or 10
- The management server must be at the highest CSM level of any node in the cluster.
- RHEL 4 WS is supported only on managed nodes.
- There is limited hardware control support for the System x 440.
- Professional versions of Linux are not supported with CSM V1.6.
- For System x, if you do not have a terminal server, the remote console function (rconsole) will not work and automatic MAC address collection will not work for the full installation of Linux nodes. Refer to the **Hardware requirements** section for CSM for Linux Multiplatform.
- Support for diskless Linux nodes requires that all such nodes have the same version and release of Linux.

26. Supported Memory page sizes

AIX 5L Version 5.3 with the 5300-04 Recommended Maintenance Package supports up to four different page sizes, but the actual page sizes supported by a particular system will vary based on processor type. The following table lists the page sizes supported by AIX 5L Version 5.3 with the 5300-04 Recommended Maintenance Package and required System p™ hardware:

Page Size	Required Hardware	Requires User Configuration	Restricted	Kernel
4KB	ALL	No	No	64 & 32
64KB	IBM POWER5+™ or later	No	No	64 only
16MB	POWER4™ or later	Yes	Yes	64 & 32
16GB	POWER5+ or later	Yes	Yes	64 only

- As with all previous versions of AIX® and AIX 5L, 4KB is the default page size for AIX 5L Version 5.3 with the 5300-04 Recommended Maintenance Package. A process will continue to use 4KB pages unless a user specifically requests another page size be used.
- No system configuration changes are necessary to enable a system to use 64KB pages. On systems that support 64KB pages, the AIX 5L kernel will automatically configure 64KB pages for the system. 64KB pages are fully pageable, and the size of the pool of 64KB page frames on a system is dynamic and fully managed by AIX 5L.
- AIX 5L will vary the number of 4KB and 64KB page frames on a system to meet demand on the different page sizes. Both the svmon and vmstat commands can be used to monitor the number of 4KB and 64KB page frames on a system.

The specific page sizes supported on a system depends on a system's processor type. You can use the **pagesize -af** command to display all of the virtual memory page sizes supported by AIX on a system. You can specify the page sizes to use for three regions of a process's address space using an

environment variable or settings in an application's XCOFF binary with the **ldedit** or **ld** commands as shown in the following table:

Region	ld / ldedit option	LDR_CNTRL environment variable	Description
Data	-bdatapsize	DATAPSIZE	Initialized data, bss, heap
Stack	-bstacksize	STACKPSIZE	Initial thread stack
Text	-btextpsize	TEXTPSIZE	Main executable text

For example, the following command causes **mpsize.out** to use 64 KB pages for its data, 4 KB pages for its text, and 64 KB pages for its stack on supported hardware:

```
$ LDR_CNTRL=DATAPSIZE=64K@TEXTPSIZE=4K@STACKPSIZE=64K mpsize.out
```

Unless page sizes are selected using one of the above mechanisms, a process will continue to use 4 KB pages for all three process memory regions by default. Using 64 KB pages rather than 4 KB pages for a multi-threaded process's data can reduce the maximum number of threads a process can create. Applications that encounter this limit can reduce internal pthread library memory usage and allow for more threads to be created by setting the environment variable **AIXTHREAD_GUARDPAGES** to 0. In addition to these three memory regions of a process's address space, you can select the page size for system V shared memory regions by using the **SHM_PAGESIZE** command to the **shmctl()** system call. The 4 KB and 64 KB page sizes are intended to be general-purpose, and no system configuration changes are necessary to enable a system to use these page sizes. The 16 MB large page size and 16 GB huge page size are intended only to be used in very high performance environments, and a system administrator must configure a system to use these page sizes. Furthermore, the support for 16 MB large pages and 16 GB huge pages is limited. 16 MB large pages are only supported for process data and shared memory, and 16 GB huge pages are only supported for shared memory. The **ps -Z** command displays the page sizes being used for the data, stack, and text memory regions of a running process. The **vmstat** command is enhanced to display information about multiple page sizes. The **-p** and **-P** options to the **vmstat** command displays VMM statistics for each supported page size. Finally, the following **vmo** command can be used to disable all kernel support for 64 KB and 16 GB pages: **vmo -r -o vmm_mpsize_support=0**

27. Maximum size of boot image increased: For AIX 5L Version 5.3, the maximum size of the boot image has changed from the previous value of 11,984 KB (12 MB minus 16 KB) to 31,984 KB (32 MB minus 16 KB).

28. The AIX 5L operating system previously contained both a uniprocessor 32-bit kernel and a 32-bit multiprocessor kernel. Effective with AIX 5L Version 5.3, the operating system supports only the multiprocessor kernel. The AIX 5L Version 5.3 32-bit multiprocessor kernel supports the following systems: RS/6000, System p, or OEM hardware based on the Common Hardware Reference Platform (CHRP) architecture, regardless of the number of processors. The maximum real memory supported by a 32-bit kernel system (or partition) is 96 GB. AIX 5L Version 5.2 is the last release of AIX that supports the uniprocessor 32-bit kernel.

The AIX 5L Version 5.3 kernels provide the same functionality, regardless of which kernel is being used. The 32-bit and 64-bit kernel systems have common base libraries, commands, utilities, and header files. Differences between 32-bit and 64-bit kernel systems are limited to the following:

System and I/O Support. The 64-bit kernel limits support to 64-bit POWER-based systems, while the 32-bit kernel supports both 32-bit and 64-bit POWER-based systems. In addition, the 64-bit kernel does not support all I/O that is supported by the 32-bit kernel.

Application Support. The 64-bit kernel supports both 32-bit and 64-bit applications. Application source and binaries are portable between AIX 5L Version 5.3 64-bit and 32-bit kernel systems, in the absence of any application dependencies on internal kernel details or on kernel extensions that are not supported under the 64-bit kernel but are supported under the 32-bit kernel.

- **Binary Compatibility.** Binary compatibility is provided for 32-bit applications running on earlier versions of AIX on POWER-based systems, except for applications linked statically or applications dependent on undocumented or unsupported interfaces. In addition, some system

file formats have changed, and 32-bit applications processing these files might need to be recompiled.

- **Application Scalability.** AIX 5L Version 5.3 provides a more scalable application binary interface (ABI) for 64-bit applications. To take advantage of the scalability improvements to 64-bit programs, all 64-bit applications and libraries must be recompiled on AIX 5L Version 5.3. In addition, existing 32-bit kernel extensions and device drivers used by 64-bit applications might have to be modified in order to support the new 64-bit ABI.v

Kernel Extensions. Kernel extensions for the 64-bit kernel run in 64-bit mode and have the scalability of the larger kernel address space. Some kernel services available in the 32-bit kernel are no longer provided by the 64-bit kernel, so existing 32-bit kernel extensions may have to be ported in order to be used with the 64-bit kernel. Existing 32-bit kernel extensions continue to be supported by the 32-bit kernel, but these kernel extensions are not usable by the 64-bit kernel. Not all of the kernel extensions supported for the 32-bit kernel are supported for the 64-bit kernel, particularly the device drivers for the I/O.

- **Dual-mode Kernel Extensions.** AIX 5L Version 5.3 supports dual-mode kernel extensions, which can be loaded by a common configuration method, regardless of which kernel is being used. A dual-mode kernel extension is an archive file that contains both the 64-bit and 32-bit versions of the kernel extension as members.

Installation and Enablement. The 32-bit and 64-bit kernels are provided as part of the AIX 5L Version 5.3 base media and are installed on all supported hardware systems. The default kernel enabled during installation is dependent on the hardware system being installed. On POWER5 systems, the 64-bit kernel is enabled during base system installation. On all other systems, the 32-bit kernel is enabled. However, you can override this default option at installation time through the system installation panels. You can switch between the 32-bit and 64-bit kernels without reinstalling the operating system. 1. Modify the **/usr/lib/boot/unix** directory and the **/unix** directory to be a symbolic link to the binary for the desired kernel. 2. Run the **bosboot** command to write a new system boot image. 3. Reboot the system. The path name of the 64-bit kernel is **/usr/lib/boot/unix_64**, and the path name of the multiprocessor versions of the 32-bit kernel is **/usr/lib/boot/unix_mp**.

29. JFS2 file system freeze and thaw feature: A new feature for the JFS2 file system is added to AIX 5L Version 5.3 with the 5300-01 Recommended Maintenance package. This feature provides an external interface whereby an application can request that a JFS2 file system freeze, or stay quiescent. After the freeze operation, the file system must remain quiescent until it is thawed or until the specified timeout has past. The request for freeze or thaw can be performed from the command or from the API as follows:

Command: `chfs -a freeze=<timeout or "off"> <file system name> , chfs -a refreeze=<timeout> <file system name>`

API: `fscntl() fscntl(vfs, FSCNTL_FREEZE, (caddr_t)timeout, 0); fscntl(vfs, FSCNTL_REFREEZE, (caddr_t)timeout, 0); fscntl(vfs, FSCNTL_THAW, NULL, 0);`

30. Common Boot Image Management

_ **New commands and new SMIT interfaces**

– `mkts`, `swts`, `dbts`, `lstts`, `rmtts`, `mkcosi`, `cpcosi`, `chcosi`, `lscosi`, and `rmcosi`

_ **New functionality**

- Clone a common image and modify the clone image.
- Allow a thin server to switch to a different common image
- Specify a time to switch to different common OS image.
- Backup user data in `/home`, `/tmp`, and certain files in `/root` and restore this data after switching new common image.
- Allow thin servers to run a debug boot
- Reduced user interface complexity

31. Security Certifications:

> Labeled Security Protection Profile (LSPP) to meet the assurance requirements of Evaluation Assurance Level 4 augmented (EAL4+).

> Controlled Access Protection Profile (CAPP) under the Common Criteria for Information Security Evaluation (CC) at the Evaluation Assurance Level 4+ level (EAL4+).

<http://www-03.ibm.com/servers/aix/products/aixos/certifications/>

32. Ref: AIX 5.3 Operating system and device management, SC23-5204-00, Chapter 17.

Devices include hardware components such as, printers, drives, adapters, buses, and enclosures, as well as pseudo-devices, such as the error special file and null special file. Device drivers are located in the **/usr/lib/drivers** directory. The number of devices that AIX can support can vary from system to system, depending on several important factors. The following factors have an impact on the file systems that support the devices:

- Configuring a large number of devices requires storage of more information in the ODM device-configuration database. It can also require more device special files. As a result, more space and more i-nodes are required of the file system.
- Some devices require more space than others in the ODM device-configuration database. The number of special files or i-nodes used also varies from device to device. As a result, the amount of space and i-nodes required of the file system depends on the types of devices on the system.
- Multipath I/O (MPIO) devices require more space than non-MPIO devices because information is stored in the ODM for the device itself as well as for each path to the device. As a rough guideline, assume that each path takes up the space of one-fifth of a device. For example, an MPIO device with five paths will have the space equivalent to two non-MPIO devices.
- AIX includes both logical devices and physical devices in the ODM device-configuration database. Logical devices include volume groups, logical volumes, network interfaces, and so on. In some cases, the relationship between logical and physical devices can greatly affect the total number of devices supported. For example, if you define a volume group with two logical volumes for each physical disk that is attached to a system, this will result in four AIX devices for each disk. On the other hand, if you define a volume group with six logical volumes for each physical disk, there will be eight AIX devices for each disk. Therefore, only half as many disks could be attached.
- Changing device attributes from their default settings results in a larger ODM device-configuration database and could lead to fewer devices that can be supported.
- More devices require more real memory.

With AIX 5.2 and the minimum RAM file system size, it is likely that up to 5000 AIX devices can be configured. With the maximum RAM file system size, it is likely that up to 25,000 AIX devices could be configured. These numbers include both physical and logical devices. Depending on the various factors mentioned in this section, your system might be able to configure more or fewer devices than this number. With a large number of devices in the system, the longer configuration time contributes to a longer boot time.

33. Memory Segments

- **32 bit address space:** 16x256 MB for 32 bit, max shared memory segments is 11 or 2.75GB
 - Seg 0 : Kernel text & data
 - Seg 1 : User Process text
 - Seg 2 : User data, heap & stack,
 - Seg 3 to C & E : Shared memory segments for user processes
 - Seg D : Shared library text
 - Seg F : Per process shared library data
 - **64 address space**
 - (0 - 4 GB) - First 16 segments are exempt from general use to keep compatibility with 32-bit user process model.
 - (4 GB - 448 PB) -Application text, data, and heap, (448 - 512 PB) - Default shared memory segments,
 - (512 - 576 PB) - Privately loaded objects, (576 - 640 PB) - Shared text and data
 - (640 - 960 PB) - System reserved, (960 PB - 1 EB) - User process stack
- (Ref: Developing and Porting C and C++ Applications on AIX and General Programming Concepts: Writing and Debugging Programs).

34. AIX 6.1 Preview and New Features

Workload Partition: Workload Partition (WPAR) is a software-base virtualization capability of AIX 6 that will provide a new capability to reduce the number of AIX operating system images that need to be maintained when consolidating multiple workloads on a single server. WPARs will provide a way for clients to run multiple applications inside the same instance of an AIX operating system while providing security and administrative isolation between applications. WPARs complement logical partitions and can be used in conjunction with logical partitions if desired.

WPAR can improve administrative efficiency by reducing the number of AIX operating system instances that must be maintained and can increase the overall utilization of systems by consolidating multiple workloads on a single system and is designed to improve cost of ownership.

Application Mobility: Application Mobility is a new capability that allows a client to relocate a running WPAR from one system to another, without requiring the workload running in the WPAR to be restarted. Application Mobility is intended for use within a data center and requires the use of the new Licensed Program Product, the IBM AIX Workload Partitions Manager.

Partition Mobility: Live Partition Mobility will allow clients to move a running partition from one physical System p POWER6 server to another System p POWER6 server without application downtime, helping clients to avoid application interruption for planned system maintenance, provisioning, and workload management.

Security: AIX 6 will also include significant new capabilities to enhance the already strong security available with AIX. Some of these capabilities include:

- **Role Based Access Control:** Role Based Access Control (RBAC) is designed to improve security and manageability by allowing administrators to delegate system administrative duties to nonroot users. RBAC in AIX has been enhanced to provide very fine granular authorizations that by name identify the privileged operation which they control. These authorizations can be used to create the required roles necessary and assign those roles to the users required to manage the system. Such nonroot users will be able to assume the role and perform the allowed privileged operations.

- **Trusted AIX:** Trusted AIX extends the security capabilities of the AIX operating system by supplying integrated multi-level security. Trusted AIX is implemented as an installation option that can provide the highest levels of label-based security to meet critical government and private industry security requirements.

- **Encrypting filesystem:** The IBM Journaled Filesystem Extended (JFS2) provides for even greater data security with the addition of a new capability to encrypt the data in a filesystem. Clients can select from a number of different encryption algorithms. The encrypted data can be backed up in encrypted format, reducing the risk of data being compromised if backup media is lost or stolen. The JFS2 encrypting filesystem can also prevent the compromise of data even to root-level users.

- **Enhancements to the AIX Security Expert:** The AIX Security Expert was introduced with Technology Level 5 update to the AIX V5.3 operating system, and provides clients with the capability to manage more than 300 system security settings from a single interface and the ability to export and import those security settings between systems. AIX 6 includes an enhancement to the Security Expert to store security templates in a Lightweight Directory Protocol (LDAP) directory for use across a client's enterprise.

- **Trusted Execution:** The Trusted Execution (TE) feature provides for an advanced mechanism for checking and maintaining system integrity. A signature (SHA256/RSA) database for the important system files is created automatically as part of regular AIX install.

The TE tool can be used to check the integrity of the system against the database. Also administrators can define policies such that the loads of files listed in the database are monitored and execution/loads not allowed if hashes do not match. Additionally, administrators can lock the signature database or the files in the database from being modified by any one in the system, including root.

- **Secure by Default:** The AIX 6 installation process will offer a new option, Secure by Default, that installs only the minimal number of services to provide the maximum amount of security. The Secure by Default option works particularly well when used in conjunction with the AIX Security Expert to only enable the system services required for the system's intended purpose.

- **Continuous availability:** Improved reliability, availability, and serviceability have become the most important requirements for many clients, particularly clients that have consolidated multiple workloads onto a single system. Over the past several years, IBM has included many continuous availability features in the AIX operating system. AIX 6 includes many mainframe-inspired continuous availability features, including:

- **Concurrent AIX updates:** Concurrent AIX updates provides a new capability to deliver some kernel updates as Interim Fixes that will not require a system reboot to be put into effect. This new capability will provide IBM with a tool to reduce the number of unplanned outages required to maintain a secure, reliable system.

- **Kernel Storage Keys:** Kernel exploitation of the POWER6™ processor storage key feature brings a mainframe-inspired reliability capability to the UNIX market for the first time. Storage keys can reduce the number of intermittent outages associated with undetected memory overlays inside the kernel. Applications can also use the POWER6 storage key feature to increase the reliability of large, complex applications running under the AIX V5.3

or AIX V6.1 operating systems.

– **Dynamic tracing with probevue:** AIX 6 will provide a new dynamic tracing capability that can simplify debugging complex system or application code without requiring code changes and recompilation. This dynamic tracing facility will be introduced via a new tracing command, probevue, that allows a developer or system administrator to dynamically insert trace breakpoints in existing code without having to recompile the code. A developer or system administrator can use probevue to dynamically place probes in existing code, and specify the data to be captured at probe point.

– **Live dump:** AIX 6 continues to build upon the first failure data capture and nondisruptive service aid features introduced in prior AIX releases. A new live dump feature allows selected subsystems to dump their memory state and traces to the filesystem for subsequent service analysis, without requiring a full system dump and outage.

• **Improving manageability:** AIX 6 includes many new capabilities to improve the manageability of the AIX operating system, including NFSv4 support for the Network Installation Manager (NIM), a new, graphical installation tool and a new graphical systems console, the Systems Director Console for AIX. The Systems Director Console for AIX provides a responsive Web access to common systems management tools such as the Systems Management Interface Tool (SMIT) and offers integration into the IBM Systems Director. The Systems Director Console for AIX is included with AIX 6.

Note 35: AIX 6.1 kernel will provide 32 bit binary compatibility for applications running on AIX 5L on POWER4 systems and later. **The 32-bit kernel will no longer be available with AIX 6, 32-bit kernel extensions and device drivers will not run and this does not affect support for user space 32 and 64 bit applications.**

Application (user) Space	<input checked="" type="checkbox"/>	32-bit Applications
	<input checked="" type="checkbox"/>	64-bit Applications
Kernel Space	<input checked="" type="checkbox"/>	64-bit Device Drivers
	<input checked="" type="checkbox"/>	64-bit Kernel Extensions
	<input checked="" type="checkbox"/>	32-bit Device Drivers
	<input checked="" type="checkbox"/>	32-bit Kernel Extensions

Note 36: Solution Performance Tuning (out of the box performance)

- Selected changes to AIX tuning defaults to optimize on common System p workloads
- Primarily expected to benefit workloads sensitive to VMM tuning
- Based on extensive benchmarking and tuning experience from AIX performance
- New categorization of tunables into restricted and non-restricted types

Note 37: **Advanced First Failure Data Capture features**

AIX 5L Version 5.3 with the 5300-05 Technology Level package provides many advanced First Failure Data Capture (FFDC) features. These features include Lightweight Memory Trace (LMT), Component Trace (CT[®]), and Run-Time Error Checking (RTEC). These features are enabled by default, at levels that provide valuable FFDC information with minimal performance impacts. The advanced FFDC features can be individually manipulated. To enable or disable all three advanced FFDC features, enter the following command: `smit ffdc` You can then choose to enable or disable FFDC features. Note that a **bosboot** and reboot are required to fully enable or disable all FFDC features. Any change to LMT will not take effect until the next boot.

System trace The system trace facility has been enhanced to support process and thread-based tracing. You can restrict the tracing to a process and capture the events in relation to the process for better debugging. For more information, see the **trace** command documentation. The **trace** command supports settings of larger trace buffers for regular users. For more information, see the **trcctl** command

documentation. The system trace can be used to trace processor utilization register (PURR) to provide more accurate event timings in a shared processor partition environment.

Lightweight Memory Trace The Lightweight Memory Trace (LMT) provides system trace information for First Failure Data Capture (FFDC). It is a constant kernel trace mechanism that records software events occurring during system life. The system activates LMT at initialization, then tracing runs continuously. Recorded events are saved into per processor memory trace buffers. There are two memory trace buffers for each processor, one to record common events, and one to record rare events. The memory trace buffers can be extracted from system dumps and accessed on a live system by service personnel. The impact on the throughput of a kernel-intensive benchmark is one percent, and is much less for typical user workloads. LMT requires the consumption of a small amount of pinned kernel memory. The default amount of memory required for the trace buffers is calculated based on factors that influence software trace record retention. For the 64-bit kernel, the default calculation is additionally limited such that no more than 1/128th of system memory can be used by LMT, and no more than 256 MB by a single processor. The 32-bit kernel uses the same default buffer memory size calculation, but restricts the total memory allocated for LMT (all processors combined) to 16 MB. The 64-bit kernel resizes the LMT trace buffers in response to dynamic reconfiguration events, the 32-bit kernel does not. The following table shows some examples of default LMT memory consumption:

Machine	Number of CPUs	System Memory	Total LMT Memory: 64-bit Kernel	Total LMT Memory: 32-bit Kernel
POWER3™ (375 MHz CPU)	1	1 GB	8 MB	8 MB
POWER3 (375 MHz CPU)	2	4 GB	16 MB	16 MB
POWER5 (1656 MHz CPU, SPLPAR, 60% ent cap, SMT)	8 logical	16 GB	120 MB	16 MB
POWER5 (1656 MHz CPU)	16	64 GB	512 MB	16 MB

To determine the amount of memory being used by LMT, enter the following shell command: `echo mtrc | kdb | grep mt_total_memory` The **raso** tunable command can be used to disable LMT. It can also be used to increase or decrease the memory trace buffer sizes. For more information, see the **raso** command documentation.

Component Trace The Component Trace (CT) facility provides system trace information for specific system components. This information allows service personnel to access component state information through either in-memory trace buffers or through traditional AIX system trace. CT is enabled by default. The use of in-memory CT buffers can be persistently disabled across reboots by using the **ctctrl -P memtraceoff** command. CT can be persistently enabled by running the **ctctrl -P memtraceon** command. **Note:** A **bosboot** is required to make the command persistent on the next boot. Information on these and other CT commands can be found in the **ctctrl** command documentation.

Run-Time Error Checking The Run-Time Error Checking (RTEC) facility provides service personnel with a method to manipulate debug capabilities that are already built into product binaries. RTEC provides service personnel with powerful first failure data capture and second failure data capture error detection features. All Run-Time Error Checking can be persistently disabled across reboots by running the **errctrl -P errcheckoff** command. RTEC can be re-enabled persistently by running the **errctrl -P errcheckon** command. **Note:** A **bosboot** is required to make the command persistent on the next boot. For more information on the **errctrl** command, see *AIX 5L Version 5.3 Commands Reference, Volume 2*. RTEC features include:

- Xmalloc debug** In AIX 5L Version 5.3 with 5300-05 Technology Level, random sampling of xmalloc allocations is enabled to catch memory leaks, buffer overruns and accesses to freed data. Xmalloc debug is similar to the previous memory overlay detection system (MODS). To specifically disable the xmalloc debug RTEC feature, run the **errctrl errcheckoff -c alloc.xmdbg -r** command. To enable xmalloc debug, run the **errctrl errcheckon -c alloc.xmdbg -r** command. For more information, see the MODS and **errctrl** command documentation.
- Excessive Interrupt Disablement Detection** The Excessive Interrupt Disablement Detection mechanism in AIX can detect whether or not privileged

code remains disabled for interrupts for too long. Because excessive disablement might lead to performance problems, AIX writes an error log record to report this detection: IDENTIFIER TIMESTAMP T C RESOURCE_NAME DESCRIPTION A2205861 0705170705 P S SYSPROC Excessive interrupt disablement time Report these error logs to IBM service. The detailed report contains additional information including a stack traceback and LMT (trace) data that can be used by IBM to identify the source of the problem.

Only one period of interrupt disablement that exceeds .5 seconds is logged per boot (default). Note that each of these error log entries might identify a unique potential problem. These error reports are persistently disabled if RTEC is globally disabled. On a per-boot basis, disablement detection can be disabled by running the following command: `errctrl errcheckoff -c proc.disa` Finally, the following functions can be called from a disabled code section of a detected kernel extension to exempt the section from future excessive disablement reporting: `disablement_checking_suspend` `disablement_checking_resume` For more information about disablement checking, see "disablement_checking_suspend Kernel Service" and "disablement_checking_resume Kernel Service" in the *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1*. Also see the `errctrl` command documentation.

Other RAS enhancements The `chcore` command provides for management of location of core files. For more information, see the `chcore` command documentation. AIX error logging now supports up to 4096 bytes of event data (see the `/usr/include/sys/err_rec.h` file). However, this size error log entry is intended only for restricted system use and general error log entries should continue to contain 2048 bytes or less of event data. While up to 4096 bytes of detail data is allowed, this size entry may be truncated across a reboot in certain circumstances. The largest detail data size guaranteed not to be truncated is 2048 bytes. A large error log entry reduces the non-volatile storage available to the system dump facility in the event of a system crash.

Note 38: Restrictions using HEA under EtherChannel

Host Ethernet Adapter logical ports are only supported under EtherChannel if all adapters within the EtherChannel are HEA logical ports (including the backup adapter, if any). Consequently, having HEA logical port adapters intermixed with physical Ethernet adapters or Virtual I/O Ethernet Adapters in the same EtherChannel is not supported.

When using multiple HEA logical ports as primary adapters in an EtherChannel, the physical ports associated with the HEA logical ports must also be placed in an EtherChannel in the Ethernet switch. Consequently, all partitions that use HEA logical ports going to the same HEA physical ports must also be placed in an EtherChannel.

For example, assume that Partition 1 is configured as follows:

- * A logical HEA port out of physical HEA port 0
- * A logical HEA port out of physical HEA port 1
- * An EtherChannel created using the logical HEA ports listed above

If another partition on the same system that needs to use a logical HEA port out of physical HEA port 0 or out of physical HEA port 1, you must create an EtherChannel for the partition over both of the logical HEA ports, similar to the configuration of Partition 1. Attempting to use either of those logical HEA ports as stand-alone in other partitions might cause connectivity problems, because packets might not be delivered to the correct logical HEA port.

The aforementioned configuration restriction does not exist for using logical HEA ports in a Network Interface Backup configuration (1 primary and 1 backup), since the physical HEA ports do not require specific configuration on the Ethernet switch.