IBM

AIX 5L Differences Guide Version 5.1 Edition

- AIX 5L - The industrial strength UNIX operating system
- Intel Itanium-based and IBM POWER-based platform support
- Version 5.0 and Version 5.1 enhancements explained

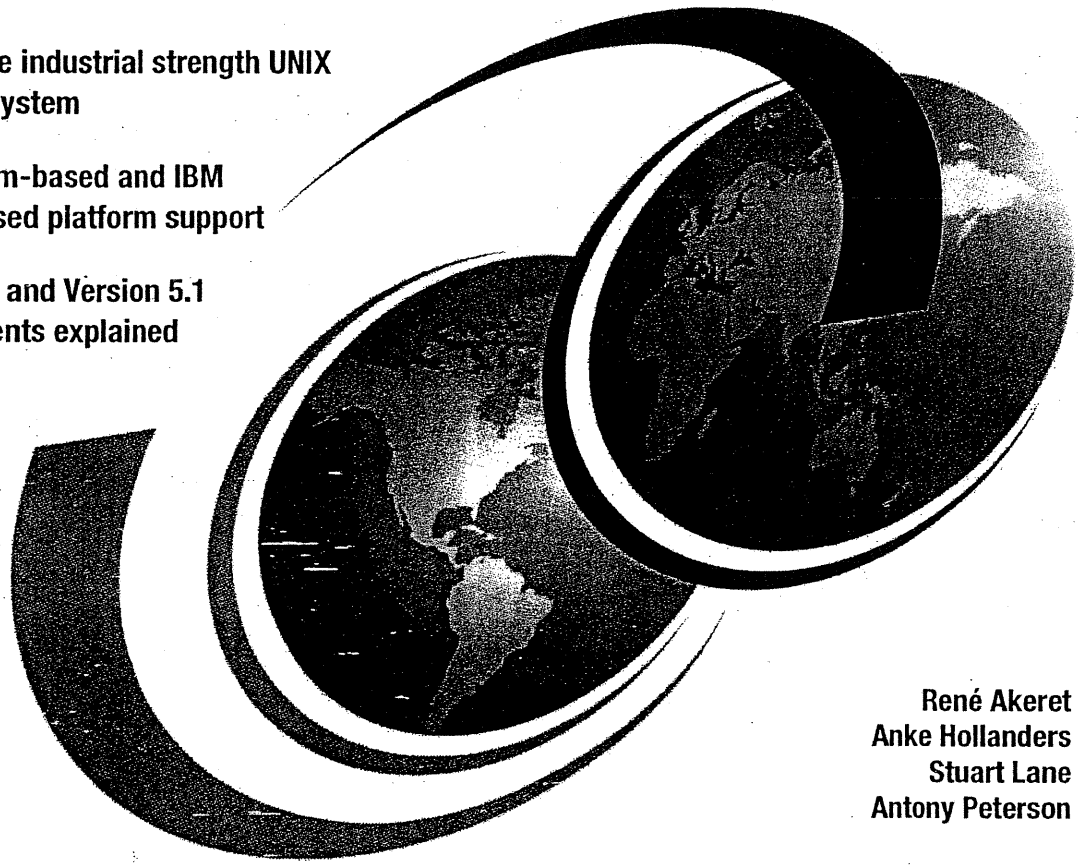

René Akeret
Anke Hollanders
Stuart Lane
Antony Peterson

Redbooks

ibm.com/redbooks

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, "Special notices" on page 473.

Second Edition (June 2001)

This edition applies to AIX 5L for POWER Version 5.1, program number 5765-E61 and for Itanium-based systems Version 5.1, program number 5799-EAR available as an PRPQ.

This document was updated on September 28, 2001.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JN9B Building 003 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Chapter 1. AIX 5L introduction and overview

AIX is IBM's strategic UNIX operating system for mission-critical, core business applications. The industrial-strength features and functions of AIX have been well proven over the years in a wide variety of server environments, from relatively small, single-processor systems through IBM's massively parallel Scalable POWERParallel servers. These features include:

- State-of-the-art 32-bit and 64-bit kernels
- 32-bit and 64-bit application programming interfaces (APIs) support
- Linux affinity that allows customers to realize a smooth technology transition between two of the industry's most open, standards-based operating environments, AIX and Linux.
- Workload Manager to balance the most complex workloads.
- Simplified system management - System Management Interface Tool (SMIT) and Web-based System Management

AIX 5L represents the next generation of AIX. Fortified with open technologies from some of the world's top providers, AIX 5L builds on a solid heritage of supplying integrated, enterprise-class support for RS/6000 and IBM @server pSeries systems.

With AIX 5L Version 5.1, IBM provides an industrial strength UNIX operating system with increased levels of integration, flexibility, and performance for meeting the high demands of today's mission-critical workloads.

The following list is a quick description of the enhancements and differences available in this release. For further information, consult the references provided.

- AIX 5L kernel and application development differences

  A summary of these differences can be found in Section 1.1, "AIX 5L kernel and application development differences summary" on page 6.

- Development environment and tools enhancements
  - An improved print function for DBX that provides more legible output is explained in Section 2.2, "DBX enhancements" on page 15.
  - Pthread enhancements, including application-level access to the pthread debug library, a new method to unregister atfork handlers, and a read/write locking enhancement are explained in Section 2.3, "Pthread differences and enhancements" on page 17.

- Core file enhancements that allow an application to core dump without termination are discussed in Section 2.11, "Lightweight core file support" on page 34.
- Enhancements to the KDB kernel debugger including a new way to load it and additional subcommands are discussed in Section 2.5, "KDB kernel debugger and kdb command enhancements" on page 22.
- Enhancements that allow application level control over the scheduler during critical sections to prevent loss of context are explained in Section 2.9, "Context switch avoidance" on page 32.
- 32-bit application scaling enhancements are discussed in Section 2.10, "Very large program support" on page 33.
- A new Korn shell, ksh93, is discussed in Section 2.20, "KornShell enhancements" on page 45.
- Enhancements in malloc provide faster access to free memory for applications is discussed in Section 2.15, "Malloc enhancements" on page 38.
- An improved `restore` command helps you recover sparse database files, as explained in Section 2.17, "Non-sparseness support for the restore command" on page 41.
- The `pax` command includes support for large files, such as dumps greater than 2 GB, as discussed in Section 2.18, "The pax command enhancements" on page 42.
- AIX 5L introduces the IBM AIX Developer Kit, JAVA 2 Technology Edition, Version 1.3.0, as discussed in Section 2.24, "Java currency" on page 49.

• LVM and file system enhancements
- New LVM hot-spare disk support, new `redefinevg`, `migratelp`, and `recreatevg` commands, new logical track group sizes, and hot spot management are discussed in Section 3.1, "Summary of the enhancements" on page 51.
- The /proc file system is discussed in Section 3.3, "The /proc file system" on page 77.
- The JFS2 is introduced in Section 3.4, "The enhanced Journaled File System" on page 81. It provides the capability to store much larger files than JFS, in a more efficient manner.
- NFS statd, AutoFS, and CacheFS enhancements are discussed in Section 3.5, "NFS statd multithreading" on page 96, Section 3.6,

"Multithreaded AutoFS" on page 97, and Section 3.7, "Cache file system enhancements" on page 97.

- A new passive mirror write consistency check can improve disk mirroring performance as discussed in Section 3.8, "Passive mirror write consistency check" on page 99.
- Updates to LVM libraries for multithreaded applications are discussed in Section 3.9, "Thread-safe liblvm.a" on page 100.

- System management and utility enhancements
  - An expanded set of devices that allow for simultaneous multiple device configuration during system startup is discussed in Section 4.9, "Fast device configuration enhancement" on page 140.
  - New ways for you to dynamically manage your paging areas, such as deactivating a paging space with the `swapoff` command or decreasing its size, is discussed in Section 4.11, "Paging space enhancements" on page 146.
  - Updates to the error log provide a more concise view of system errors, such as a link between the error log and diagnostics, or the elimination of duplicate errors, are described in Section 4.13, "Error log enhancements" on page 151.
  - AIX 5L provides a set of resources to be monitored and actions to be taken at defined events providing automatic monitoring and recovery of select critical system resources. For more information, see Section 4.14, "Resource Monitoring and Control (RMC)" on page 154.
  - Shutdown logging is available, as described in Section 4.15, "Shutdown enhancements" on page 171.
  - New methods to diagnose system errors through dump improvements are described in Section 4.18, "System dump enhancements" on page 179.
  - The ability to recover from certain system hangs is covered in Section 4.19, "System hang detection" on page 182.
  - Enhancements to performance tools, including the `truss`, `iostat`, and `vmstat` commands, are discussed in Section 4.21, "Performance Analysis Tools" on page 186.
  - Workload Manager continues to receive improvements, as discussed in Chapter 7, "Workload Manager" on page 391.
  - The new System V Release 4 print subsystem is discussed in Section 4.25, "System V Release 4 print subsystem" on page 204.

- Web-Based System Manager receives major usability improvements with a much improved architecture and usability enhancements, such as accelerator keys. A discussion of all the enhancements can be found in Section 4.27, "Web-based System Manager for AIX 5L" on page 236.
- Security and User authentication and LDAP enhancements are discussed in Sections 4.28, "User and group integration" on page 260; 4.31, "IBM SecureWay Directory Version 3.2" on page 271; and 4.32, "LDAP name resolution enhancement" on page 275.
- A new documentation search engine to handle single- and double-byte searches together is discussed in Section 4.36, "Documentation search-engine enhancement" on page 286.
- AIX is Tivoli ready, as discussed in Section 4.41, "Tivoli readiness" on page 291.
- An updated Welcome Center available with CATIA will teach you what is available for AIX in the CATIA market. For more information, see Section 4.42, "CATIA Welcome Center" on page 291.

• Networking Enhancements

- The demand for QoS arises from applications such as digital audio/video or real-time applications and the need to manage bandwidth resources for arbitrary administratively-defined traffic classes. For more information, see Section 5.1, "Quality of service support" on page 295.
- Together, multipath routing and dead gateway detection provide automatic selection of alternate network pathways that provide significant improvements in network availability. For more information, see Section 5.2, "TCP/IP routing subsystem enhancements" on page 299.
- With Virtual IP Address, the application is bound to a virtual IP address, not a real network interface that can fail. When a network or network interface failure is detected (using routing protocols or other schemes), a different network interface can be used by modifying the routing table without affecting application operation. For more information, see Section 5.5, "Virtual IP address support" on page 331.
- Dynamic Feedback Protocol (DFP) is a way to provide load statistics to a load manager so that load can be balanced by sending future connections to available servers. For more information, see Section 5.13, "Dynamic Feedback Protocol (5.1.0)" on page 361.

- Sendmail Version 8.11 improves performance by having multiple queues, memory-buffered pseudo-files, and more control over resolver time-outs. For more information, see Section 4.20, "Sendmail upgrade enhancements (5.1.0)" on page 184.
- TCP/IP performance over congested networks is improved through increased initial windows, explicit congestion notification, and limited transmit mechanism functions, which are configurable by a system administrator. For more information, see Section 5.2, "TCP/IP routing subsystem enhancements" on page 299.
- TCP splicing helps push the data-relaying function of a proxy application (from server-side socket to the client-side socket or vice versa) into the kernel. For more information, see Section 5.3.2, "TCP splicing" on page 322.
- Network Interface Takeover is a new option allowing the configuration of multiple adapters, including IBM 10/100 Mbps Ethernet PCI adapter, Gigabit Ethernet-SX PCI adapter, and 10/100/1000 Base-T Ethernet PCI adapter, allowing one or more to be designated as a backup. For more information, see Section 5.16, "Etherchannel enhancements (5.1.0)" on page 369.
- Virtual LAN (VLAN) provides the ability to create virtual LANs across multiple physical LANs or segment and/or divide physical LAN segments into virtual LANs. For more information, see Section 5.17, "Virtual local area network (VLAN) (5.1.0)" on page 374.
- Enhancements to the Network Buffer Cache and HTTP GET kernel extension provide class leading Web server performance. For more information, see Sections 5.6, "Network Buffer Cache dynamic data support" on page 335, and 5.7, "HTTP GET kernel extension enhancements" on page 338.
- Applications can be modified to capture network data packets through a new interface, as explained in Section 5.8, "Packet capture library" on page 342.
- To allow more flexible development of firewall software, AIX provides additional hooks, as described in Section 5.9, "Firewall hooks enhancements" on page 343.
- PC Interoperability using Fast Connect File and Print Services provides support for Windows 2000, improved user and name mapping, share options, WTS support, better performance, and more, as discussed in Section 5.10, "Fast Connect enhancements" on page 345.

- Enhancements to increase affinity with Linux
  - A set of Linux-compatible routines has been added to AIX 5.1 so that Linux applications using these routines do not have to supply their own libraries. For more information, see Section 6.2, "AIX source affinity for Linux applications (5.1.0)" on page 388.
  - AIX Toolbox for Linux Applications, delivered on a supplemental CD, which contains a collection of open source and GNU software built for AIX and packaged in RPM format. For more information, see Section 6.1, "AIX Toolbox for Linux Applications" on page 379.
- A list of packages and filesets that are not part of the AIX 5L Itanium-based offering is provided in Appendix A, "AIX 5L POWER and Itanium-based differences" on page 471.

1.1 AIX 5L kernel and application development differences summary

The AIX development team made every effort possible to make AIX 5L for the POWER and Itanium-based platforms appear and function identically; there are, however, a few unavoidable differences due to the underlying hardware.

The following list provides a summary of the major differences, from a kernel and application development point of view, between POWER and Itanium-based systems:

- The most influential difference is the use of the IA64 instruction set architecture (ISA). Itanium-based platforms operate in little endian mode. The Itanium-based ISA arranges instructions into bundles and groups. It also contains instruction *predication* to enable explicit parallelism in instruction execution.
- Itanium-based AIX has a 64-bit kernel. There is no 32-bit kernel for Itanium-based systems.
- Common header files contain #ifdef _ia64 to denote differences between Itanium-based and POWER structures.
- Itanium-based systems have a different machine register context. The MST, signal context, and jump buffers all contain different context. The user space debugger for Itanium-based systems also displays the IA64 registers.
- Itanium-based systems have a different application binary interface (ABI) than POWER. Linkage and parameter passing conventions are different from POWER due to the machine register context differences between the platforms.