HACMP for AIX: Version 4.1.1 Update

By Daniel P. Cox

The newly announced HACMP for AIX, Version 4.1.1, has expanded support for the entire RISC System/6000 family—from the C20 to Scalable POWERparallel systems. It also supports AIX Version 4.1.4. The new features including the Visual System Manager, Cluster Node Snapshot utility, and Quick Configuration utility make installation, administration, and management easier than ever.

IBM's High Availability Cluster Multiprocessing for AIX (HACMP for AIX) now supports up to eight RISC System/6000 (RS/6000) uniprocessors, SMP processors, and Scalable Parallel (SP2™) thin or wide nodes in a highly available cluster. This function provides horizontal scalable growth and high availability in an open systems environment.

High availability, a computing configuration that recovers automatically from a single or multiple points of failure, provides better assurance against application and system downtime than standard hardware and software alone. Together, the HACMP software and a cluster of loosely coupled processors called *nodes* provide application availability by transferring control from a failed processor to a backup that has redundant capabilities. When a failure occurs or when components are restored to operation, HACMP executes test and recovery scripts tailored by the system administrator to specify the actions to be taken. HACMP software detects and recovers from failures of disks, disk adapters, networks, network adapters, and processors.

HACMP relies on the application, however, to make any failure recovery or fallover transparent to external users and client machines. If a node fails, nominal recovery time is approximately 30 to 300 seconds. Actual recovery time is a function of the system configuration, the application configuration, the size of the user's databases, and the user's recovery script (if any).

HACMP can be an alternative to upgrading a processor in the RS/6000 product line, such as an RS/6000 POWERserver™ 570 to 591. It provides horizontal cluster performance scalability with applications distributed across RS/6000 processors, sharing the disk and/or CPU resource of the cluster as independent or concurrent operating machines.

Figure 1 shows the HACMP system architecture.

HACMP for AIX Version 4

The latest version of HACMP for AIX provides complete high availability for the RISC System/6000 from the entry-level Model 7009 C10 to the 7015 R21 models, as well as the Scalable POWERparallel (SP) thin and wide node models. It supports the RS/6000 SMP Models G30, J30, and R30 from 2-way to 8-way processors. Today, IBM is the only vendor offering this range of system support in the UNIX industry.

This new version, HACMP 4.1.1, provides enhanced concurrent access support for serial-link and Serial Storage Architecture (SSA) disk subsystems, ease-of-configuration enhancements for HACMP installations, a new High Availability for Network File System for AIX (HANFS for AIX) feature, and support for new IBM processors, disk subsystems, and devices.

Key features of HACMP include the following:

- RS/6000 hardware processor models can be mixed and matched.
- RS/6000 SP thin or wide node models can be used together.
- Disk subsystems from SCSI differential fast-wide, high-speed 9333 Serial Disk and Serial Storage Architecture 7133 can be used in the same cluster.

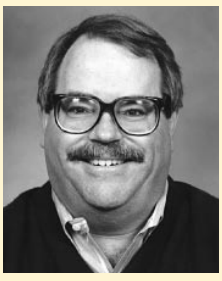

Daniel P. Cox

HACMP System Architecture

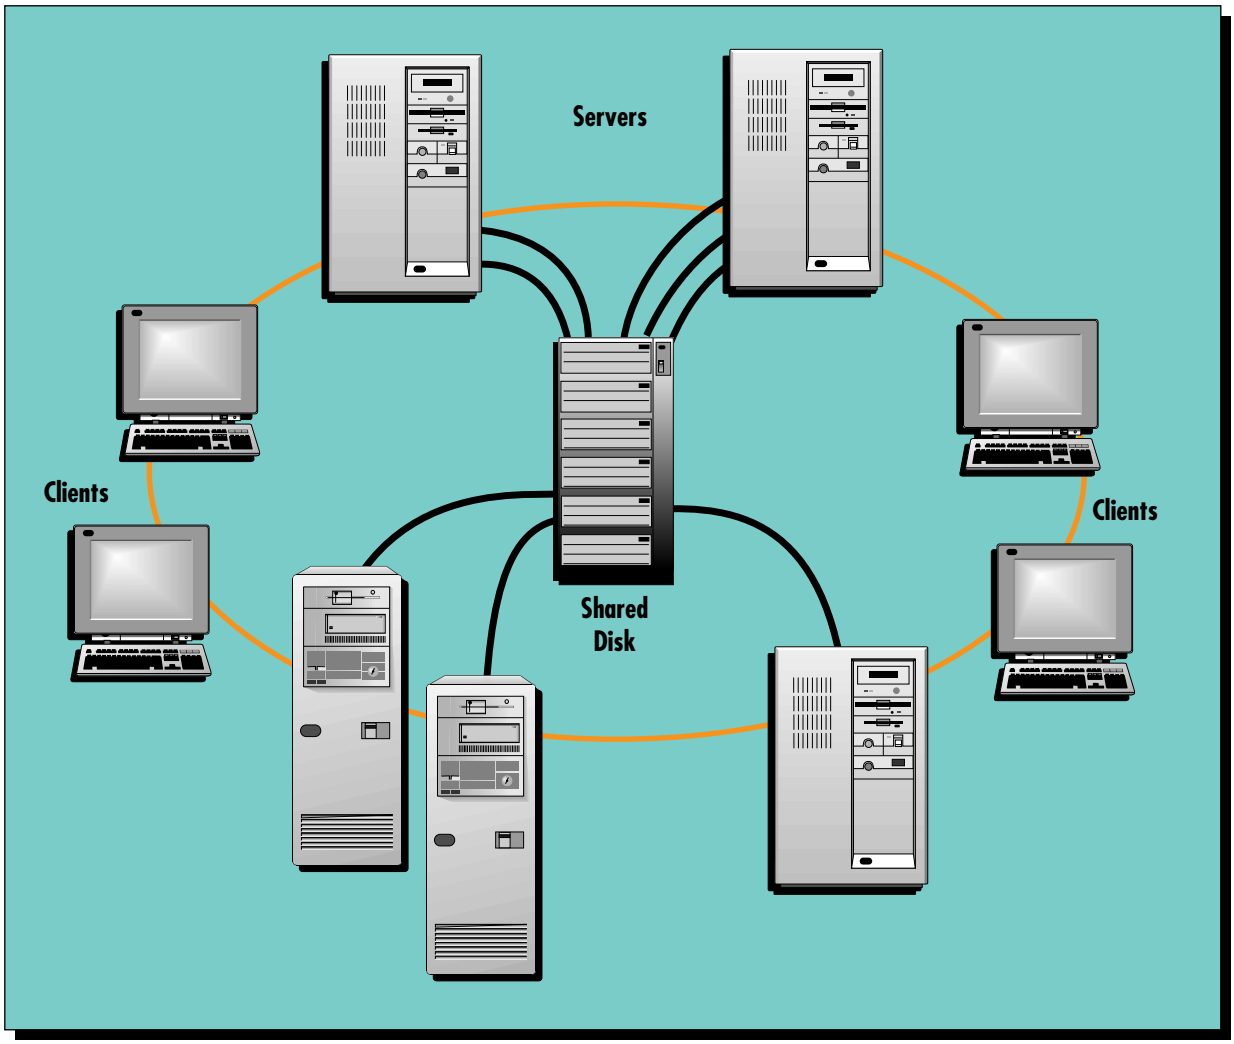

Figure 1. HACMP system architecture

- The hardware system configuration is flexible. Various LAN networks with either Ethernet, Token-Ring, Fiber-optic Data Distribution Interface (FDDI), or Fibre Channel Standard adapters with the same level of IP Address Take-Over (IP@TO) support can be used together.
- Users can operate a cluster of mixed AIX system levels, which optimizes configuration and application flexibility in a mixed cluster. This provides the flexibility to seamlessly migrate from previous HACMP Version 3 clusters on AIX Version 3 to HACMP for AIX Version 4 without taking the HACMP cluster out of service to upgrade the system and application software.
- System administrators can install, customize, and manage the HACMP cluster using the Visual System Manager (VSM), a drag-and-drop interface based on an AIX industry-standard from Common Open Software Environment (COSE).
- The Cluster Node Snapshot utility extends the HACMP cluster environment.
- Concurrent resource management for the 9333 High Performance Serial and the 7133 SSA disk subsystems now provides the horizontal scalability to produce up to an estimated 18,000 TPMs on an 8-node cluster of 8-way 7015-R30 SMP machines using the PowerPC 601™. This scalable performance has multisecond application recovery—unmatched in the industry for performance and availability.

HACMP for AIX Version 4 Highlights

Several highlights of this new version are described in the following sections.

High Availability Subsystem

HACMP 4.1.1 marks a significant advance in ease-of-configuration and system management enhancements for high-availability clusters. This subsystem provides the base services for cluster membership, system management, configuration integrity and control, and base services for fallover and recovery. It also provides programmers and system administrators with cluster status and monitoring facilities. The High Availability Subsystem contains all functions described in the remainder of this section except the Concurrent Resource Manager (CRM) and HANFS for AIX, which are separate features.

Visual System Manager

HACMP 4.1.1 makes it easy to install and configure an HACMP cluster. With AIX Common Desktop Environment (CDE) and a graphics display, the new Visual System Manager with its intuitive drag-and-drop interface allows system administrators to visualize the relationships between hardware and software resources that define the high-availability cluster, and drag screen objects from one location to another within a cluster.

Drag-and-Drop Graphical Interface

In addition to the standard SMIT interface for configuring HACMP clusters, HACMP 4.1.1 now provides a Graphical User Interface (GUI) for drag-and-drop configuration. The AIX VSM icons make it easier to visualize relationships between resources within the cluster while configuring new systems. The VSM template objects can be used to create multiple copies of the same configuration for multiple clusters. Users can load a saved configuration and make minor cluster-specific modifications to the image. VSM also makes migration and reconfiguration tasks easier.

The icon-based interface in HACMP 4.1.1 enables the user to filter subsets of data about the cluster configuration. By using a filter, users can focus on specific resources within a cluster, deliberately limiting the display to specific highly available resources, networks, or subsystems within a cluster.

Quick Configuration

The Quick Configuration utility, another new function in HACMP 4.1.1, enables installation and configuration of HACMP clusters using a few keystrokes. Selecting from a menu of hardware configuration options, users can quickly and easily select an appropriate hardware configuration, letting HACMP execute the logic necessary for configuring the new cluster, including set up of the initial AIX environment.

Common hardware configurations included in HACMP 4.1.1 are those using Redundant Array of Independent Disks (RAID), Small Computer Systems Interface (SCSI), and serial devices. With the Cluster Node Snapshot utility, users can create and save additional customized "quick configurations."

Cluster Node Snapshot Utility

The Cluster Node Snapshot utility can "capture" the configuration of an existing cluster in an ASCII representation or restore a captured configuration to a cluster, providing an invaluable tool for quick analysis of a cluster during migration or cloning of additional clusters.

Support organizations can use this facility to track changes to a particular cluster; system administrators might use it to maintain multiple cluster configurations, swapping configurations as needed.

The Cluster Node Snapshot utility can also be used with the HACMP GUI: any snapshot can be viewed as a simple text file or through its graphical representation in VSM, enabling quick analysis and diagnosis.

HACMP 4.1.1 now provides a Graphical User Interface (GUI) for drag-and-drop configuration.

Concurrent Resource Manager

HACMP 4.1.1 supports concurrent access to RAID, serial-link and SSA disk subsystems. This allows multiple RS/6000 systems (up to eight are supported for a serial-link-attached disk and up to two are supported for an SSA-attached disk) to access a shared disk simultaneously. HACMP has distributed locking facilities that support this function. For example, software such as the HACMP Distributed Lock Manager, can coordinate the use of the shared disk, which allows the workload to be spread across the cluster.

HANFS for AIX

High Availability for Network File System for AIX (HANFS for AIX) provides high availability for data accessed via the Network File System (NFS). It replaces HANFS Version 3 for those who are using AIX Version 4.1.4 or later. HANFS for AIX supports the configurations, the takeover, and the functionality of HANFS Version 3. The HANFS for AIX interface for administrators closely resembles the interface for the corresponding components of HACMP.

Two AIX systems running HANFS for AIX form a single highly available NFS server. This server can survive hardware and software outages—plus certain types of planned outages—that would make a single system NFS server unavailable. A

properly configured HANFS for AIX system that has no single points of failure can continue to serve data to NFS clients despite the loss of a single processor, disk, adapter, or network.

Additionally, HANFS for AIX allows one processor to take over for another with no disruption of service to the NFS clients—the NFS clients see only a temporary delay in the response by the server. HANFS for AIX supports all RS/6000 server systems in two node clusters along with a wide range of disk and network choices. One system can be the active server and the other an active standby. The workload can also be split between the two systems, in which each RS/6000 functions as the backup for the other. HANFS for AIX supports all NFS clients that are supported by AIX 4.1.4.

Cluster Manager monitors the state of the nodes, interfaces, and networks that comprise a cluster.

System Management

HACMP Version 4.1.1 provides enhancements for system administrators such as the following:

- Installation verification services allow the system administrator to confirm the installation procedures.
- Conversion tools ease the migration from a previous version of HACMP directly to HACMP for AIX Version 4.1.1.
- Comprehensive, data-driven scripts in HACMP minimize the need to modify the HACMP fallover scripts, simplifying cluster management and configuration.
- Cluster status and monitoring tools allow the system administrator to pass cluster status information to a single node in the cluster.
- Upgrade functions include version compatibility, a new function in this release. An existing cluster running HACMP for AIX Version 3.1 can be upgraded to HACMP Version 4.1 without taking the entire cluster off-line. During the upgrade process, individual nodes in the cluster can be removed from the cluster, upgraded one at a time, then reintegrated into the cluster. Nodes running both HACMP for AIX Version 3.1 and Version 4.1 can coexist while the rest of the nodes are upgraded.

For greater scalability, different instances of an application such as order entry can be run on multiple HACMP cluster nodes. CRM provides access controls and locking mechanisms for a database manager and an application in different nodes to share common disk storage. In addition, CRM provides several operator commands to control a shared disk environment.[1]

How Does HACMP Work?

HACMP 4.1.1 software automatically reconfigures the available replicated resources when hardware failures or outages occur, while allowing for hardware and software resources to be fully utilized when there are no hardware malfunctions in place. This fallover and recovery/restart is end-user generated through the supplied configuration and administrative scripts that can be modified to meet application configuration requirements.

In addition to providing high availability, HACMP 4.1.1 can also be configured to provide loosely coupled multiprocessing services. These configurations, whether concurrent access or parallel distributed access, allow a workload to be spread across multiple RISC System/6000 processors, sharing the disk and/or CPU resource of clustered processors.

HACMP for AIX focuses on a cluster as a set of RS/6000 and AIX resources. The cluster can include one or more of the following: disks, volume groups, filesystems, network addresses, or applications. This clustered approach, together with the capability of applications fallover and recovery/restart of HACMP 4.1.1, provides additional levels of high-availability processing for mission-critical applications, such as retail point-of-sale systems, banking credit verification, travel reservation, order-entry/inquiry, or health and hospital care systems.

Cluster Manager

The Cluster Manager (CM) component of HACMP monitors the state of the nodes, interfaces, and networks that comprise a cluster. It also provides highly available access to these resources and to critical disk data and software resources running on the cluster. To perform these functions, CM consists of three functions:

- **Heartbeat**. The CM monitors the nodes and network interfaces associated with a cluster using a heartbeat protocol. Heartbeat information received from other nodes in the cluster is interpreted and used to identify the failure of nodes, networks, and communications interfaces.

[1] CRM configurations commonly require an application that is designed to take advantage of concurrent access configurations. Not all disk subsystems sold with the RISC System/6000 servers operate in CRM environments. Consult your IBM Marketing Representative or Business Partner for current supported disk subsystems.

- **Membership**. The nodes in a cluster vote to admit or expel member nodes when changes in cluster state are detected. A node's membership status in the cluster determines its availability to manage cluster resources and to participate in failure recovery scenarios.
- **Synchronization**. When nodes are admitted or expelled, resources associated with those nodes can be moved from or added to other nodes, and these changes are made in a controlled, synchronized way. This ensures that the cluster resources are unavailable for the smallest possible time.

The CM and related services run on each of the nodes within the cluster. When the cluster is online, the CM monitors its health by sending heartbeat data over the interfaces that connect the nodes. If the CM detects a failure in part of the cluster, it reconfigures the cluster to enable it to direct data traffic through a path that avoids the failed cluster component.

The CM also triggers events based on changes in the state of the cluster. These events cause AIX shell scripts to run on the cluster, providing both system- and user-level responses to the event.

Cluster Configurations

An HACMP cluster can operate with up to eight nodes in numerous configurations, including hot standby, rotating standby, mutual takeover, and concurrent access. The following examples describe some configuration possibilities using the simplest 2-node cluster, but larger numbers of nodes can be readily clustered by defining the resources accordingly.

Hot standby: One node owns all resources; if it fails, the standby node takes over the resources. When the failed node rejoins the cluster, the resources are returned to the original node that owned them.

Rotating standby: This is identical to a hot standby, except that when the failed node rejoins the cluster, the resources are not returned to the node until the standby node fails.

Mutual takeover: Resources are divided among the nodes; some are owned by each node. If either node fails, the other node takes over all the resources. When the failed node rejoins the cluster, the resources are returned to the original owning node.

Concurrent access: Two or more nodes are active simultaneously, sharing the same physical disk resources. Both nodes own the disk resources. Other resources are divided between the two nodes, each owning some of them.

HACMP for AIX, Version 4.1.1

Some features of the latest version of HACMP include the following:

- Full concurrent access support for all IBM types of disk subsystems, including RAID subsystems, Serial-link disks, and Serial Storage Architecture disks
- Common support across the entire RISC System/6000 family, from C20 to Scalable POWERparallel systems
- Additional system management functions that make it easy to use for cluster creation, administration, and replication
- Visual System Manager with a graphical interface for defining and customizing clusters
- Cluster Node Snapshot utility that captures a specific cluster configuration and makes replication of clusters easier
- Quick Configuration utility that allows installation and configuration of HACMP clusters using just a few keystrokes

Together these features provide more flexibility and extendability as well as enhanced installation, administration, and management features.

Resources not owned by each node are designated as takeover. If either node fails, the other node takes over all resources. When the failed node rejoins the cluster, the resources are returned to the original owning node.

Clusters of up to eight nodes offer great configuration flexibility:

- One node can back up seven nodes.
- Although eight nodes can operate during peak operations, that number can be less during non-peak hours with minimal interruption of service.
- All resources could go to one node upon failure or be split among several nodes.
- Using cascading takeover, resources could be passed to either the primary standby or to its secondary standby node.

It is important to remember that an HACMP configuration is simply a definition of how the resources are defined on each node within the

cluster. The number of resources, combined with how they can be assigned, gives the cluster designer nearly unlimited flexibility in laying out an HACMP configuration.

Alternate Data Access Configurations

HACMP supports alternate data access configurations, addressing the need to run applications that share the same data on multiple nodes within a cluster.

In a concurrent access configuration, each multiple system has its own path to the disks holding the data. Any system in the cluster can physically access the data, but the systems must cooperate to ensure coordinated accesses to the data and to preserve data integrity. For these configurations, the HACMP Concurrent Resource Manager provides the necessary access controls and locking mechanisms for a database manager and an application on different nodes to share access to the common storage. Such configurations provide a high degree of scalability, limited only by the number of systems that can simultaneously attach to the shared disks.

HACMP provides the event management and application fallover in 30–300 seconds for simple applications.

Additionally, HACMP provides the mechanisms to allow a node to take over the function of another that fails. Since the shared disks can be accessed from any system, HACMP can restart the application that was running on the failed system on one of the surviving nodes.

Another solution to the problem of multiple systems accessing the same data is to have the data logically partitioned by the system within the cluster with a Database Manager (DBM) providing access to both partitions. Each system within the cluster has sole access to a dedicated partition of the total data set. The system nodes themselves are interconnected through a fast communications link. When a request is made to a particular system, the system decides whether or not it can locally access the data. If not, either the request is forwarded by the DBM to the owning system, or the data is retrieved by the DBM from the owning system.

This configuration is called *parallel distributed access* because the individual systems can process requests in parallel and access the distributed data across the cluster. The cluster scales to the degree that the communications path between the system does not become a bottleneck.

HACMP provides two advantages to a cluster built around parallel distributed access. First, if the disks can be physically connected to two or more systems, then HACMP can restart an application from a failed system on another system that can access the same disks. Second, HACMP's Distributed Lock Manager provides locking mechanisms for a database manager running on multiple nodes, which allows coordinated access to the shared data.

In both types of data access configurations, HACMP uses the mirroring capabilities of the base AIX operating system to make the data highly available and accessible. The configuration flexibility of HACMP allows users to choose the cluster topology and database manager that best suits the requirements of their computing environment. In fact, HACMP can support both types of data access within a common cluster.

Why Use HACMP?

HACMP is designed to maintain a highly available application and data server for those applications that cannot fail for any length of time. On average, a system outage is approximately 4.5 hours and will cost approximately $200,000 in lost revenue.

HACMP provides the event management and application fallover in 30–300 seconds for simple applications. It may take longer if complex database processor transactions or large database files must be re-created because of a system catastrophe caused by a natural disaster such as a flood, fire, earthquake, or a man-made disaster such as a power outage caused by terrorist acts.

The HACMP control system provides flexibility for users to design their own system for event control and recovery/restart that meets their processing and management needs across a wide spectrum of hardware and software offerings both by IBM and other OEM hardware and software providers. Some key points below address the HACMP benefits.

Improves Application Availability. HACMP for AIX improves systems reliability of RS/6000 systems in high-availability application environments. HACMP for AIX can be added to RS/6000 and AIX application configurations. The incremental costs to implement HACMP for AIX compared to reinvesting in new and additional hardware provides an inexpensive solution to increased system availability that may be required in certain business-critical application areas.

Provides Scalability. HACMP for AIX can extend the current system configuration. By using the clustering facility, the number of users on the system can be expanded. HACMP provides scalability without replacing current hardware and software solutions if processing capability exceeds the current system capacity.

Scalability is an important use of HACMP for AIX in a mutual takeover configuration, which extends

system capacity by splitting the workload on up to two to eight systems. The control system provides increased system availability as the number of users and applications increase, and system outages become more critical to the business process.

The Concurrent Resource Manager subsystem also provides availability and true horizontal scalability. This is apparent when a single image of the data is required for the application to execute and the data cannot be partitioned as in the mutual takeover configurations.

Provides Backup/Recovery. HACMP for AIX provides the basic service when business-critical applications require availability and recovery/restart. This service can be defined using any HACMP for AIX environments with multiple disk and network attachments available. In single-system environments, no shared backup is available when a processor failure occurs. Based on the user's definition of the configuration, an AIX HACMP for AIX cluster can operate in many different configurations, such as an Idle Standby or Simple Fallover to an idle machine.

HACMP for AIX provides additional capability in its inter-LAN network fallover. When a primary network fails, HACMP for AIX will fall over to a backup LAN on a separate LAN topology. For example, the primary LAN is Token Ring and the fallover backup LAN is Ethernet. This applies to disk subsystems as well, where the primary disk device is 7135-110 RAIDiant Array and the backup mirrored device is the 9333 High Performance Serial Disk subsystem.

Provides Investment Protection. HACMP for AIX can provide users with a no single-point-of-failure in an HACMP for AIX server configuration, without the extensive investment required for fault-tolerant hardware. HACMP for AIX and RS/6000 have many hardware redundant features. HACMP for AIX requires only replicated devices to eliminate a single-point-of-failure.

Enables Horizontal Growth. HACMP for AIX allows the total application performance to expand without removing or replacing the current processor base. It provides scalable growth by increasing the number of RS/6000 processors as an alternative to vertical upgrades from one model to another. When an RS/6000 model is reaching capacity with the current workload, adding a second or even third processor can be almost transparent in an HACMP for AIX Mutual Takeover configuration.

Applications and data can be transferred from one RS/6000 to another, without the need to migrate data or port applications. A variety of RS/6000 models can be attached in an HACMP for AIX cluster with the disk subsystem as the common link.

Provides An Alternative to SMP. HACMP for AIX provides an alternative to stand-alone SMP configurations.

When an application or a business processing environment grows and the number of users increases, the requirement for availability of the processing system can outweigh the increased speed of the application offered in today's symmetric machines. HACMP for AIX provides a solution for both requirements: application scalability and higher availability.

Application scalability can be handled through Mutual Takeover or Partitioned Workload, where one measure of scalable performance can be achieved. It can also be handled in a Concurrent I/O Access environment that uses a single-system data image. These HACMP for AIX configurations provide availability and recovery/restart for any number of system outage conditions, such as disk, disk adapter, LAN network, LAN adapter, or processor node failure.

HACMP for AIX configurations require minimal customization when failures occur in either the operating or application system.

HACMP/6000 configurations require minimal customization when failures occur in either the operating or application system.

I**ncreases User Productivity**. High availability allows business-critical processes to operate when components or subsystems fail or during maintenance downtime.

HACMP for AIX allows users to shutdown a system for scheduled maintenance and restart the system with automatic resynchronization of applications and data from the backup or fallover system. When a system or subsystem fails, HACMP for AIX provides automatic fallover with application and data restart. In this case, the failing system can be repaired and put back into service, and will relink into the HACMP for AIX cluster without loss of data.

HACMP for AIX also can detect multiple events. Depending on the event (such as a network failure), a unique fallover scenario can be executed, such as reinitiating service on a replicated network adapter that does not abort the cluster processor, but shifts service to the replicated or backup network adapter.

Allows Minimal Impact on End Users. The high-availability features of HACMP for AIX allow system administrators to make changes, apply upgrades, and provide scheduled hardware and/or software maintenance and backup services with minimal impact to the end users.

A backup system allows end users who are normally assigned to the machine that is out of service to continue their work. HACMP for AIX automatically integrates the processor back into the cluster and resychronizes the applications and data onto the repaired machine.

Enables Horizontal Growth. HACMP for AIX enables non-disruptive horizontal growth in the server complex through a loosely coupled cluster implementation. It provides scalable growth through its cluster configuration environments.

For scalable growth, Mutual Takeover configurations can be created that result in almost 100% CPU efficiency in a cluster of two or more processors. These configurations deal with applications that can split data independently between machines—sharing only data that is common for all applications.

HACMP/6000 uses UNIX interfaces that make it easy to port to future UNIX-based systems and subsystems.

In data sharing, LAN and filesystem overhead can reduce CPU efficiency. If data must be shared, the Concurrent I/O Access environment can be used, but it is less efficient than a single system because distributed locks must be maintained between cluster processors.

Provides Easier Portability. HACMP for AIX uses UNIX interfaces that make it easy to port to future UNIX-based systems and subsystems. HACMP for AIX uses industry-standard TCP/IP communication protocols as the transport mechanism among machines in the cluster as well as clients that may be attached to one or more cluster processors. HACMP for AIX can use multiple TCP/IP interfaces on the current RISC System/6000 adapters, such as Ethernet, Token-Ring, Fiber-optic Data Distribution Interface (FDDI), and Fibre Channel adapters and switches.

Provides Incremental Systems Facilities. HACMP for AIX provides incremental systems facilities—for adding new disk subsystems and adding new processors without stopping the cluster—to the existing application base for 8-way scalability through clustering as well as high-availability fallover.

HACMP for AIX allows different RS/6000 models to be clustered together to meet users' requirements for processing capacity and availability. Users can upgrade a current machine vertically (for example, 7013-591 to 7013-591) or purchase a second machine (for example, an additional 7013-580) and cluster the application and data workload onto both machines. This option increases CPU efficiency, improves system performance, and provides a higher level of availability than can be achieved in a stand-alone configuration.

Provides Common Interface. HACMP for AIX uses System Management Interface Tool (SMIT) for the basic installation, configuration, and administration in an HACMP for AIX cluster environment. It also uses SMIT to configure and update a cluster from a single node, providing a common management focal point from previous HACMP for AIX configurations.

Provides Cluster Monitoring. If HACMP for AIX detects a server failure, it provides cluster monitoring and automatic fallover to backup processors. It alerts either a local or a remote services console if a cluster fails. HACMP for AIX can use the NetView/6000 control system, sending a Simple Network Management Protocol (SNMP) alert record to the NetView/6000 host reporting on cluster status if an out-of-service condition occurs. If NetView/6000 is not used, this can also be monitored in the HACMP for AIX cluster console service.

Allows Script Customization. System administrators and application programmers can extend and customize HACMP for AIX scripts for use in configuration and applications to meet specific availability and/or scalability requirements. Although scripts are provided with the system, it is necessary to customize them to meet any specific requirements. HACMP for AIX now includes comprehensive scripts to ease the customization process. Users can define pre- and post-scripts that will minimize modifications to the HACMP for AIX-provided scripts.

Summary

HACMP for AIX provides flexibility and scalability across all AIX systems. It runs on all AIX versions and supports all the RISC System/6000 servers including the 9076 SP models. It supports the largest number of active processor nodes, most disk subsystem variety, all LAN topologies, and provides the richest set of Relational Database Management System (RDBMS) support by the RDBMS vendor community.

Daniel P. Cox, IBM Corporation 11400 Burnet Road, Austin, TX 78758. Mr. Cox is the brand manager for clustered systems in the RISC System/6000 Division. After joining IBM in 1968 as a programmer in the IBM San Jose Laboratory, he has held positions in development programming, systems engineering, product planning, and management. He has been involved with HACMP for AIX since its beginning—in both planning and implementation. Mr. Cox has a BS and an MS from San Jose State University.