

*High Availability Cluster
Multi-Processing for AIX*

HACMP Master Glossary

Version 5.4.1

Ninth Edition (October 2007)

Before using the information in this book, read the general information in [Notices for Master Glossary](#).

This edition applies to HACMP for AIX, version 5.4.1 and to all subsequent releases of this product until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 1998, 2007. All rights reserved.

Note to U.S. Government Users Restricted Rights—Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Master Glossary

This Glossary contains definitions of terms used throughout the HACMP for AIX (HACMP™), Version 5.4.1 documentation set. General familiarity with AIX and networking terms and commands is assumed. HACMP usage of some specific AIX terms and commands is explained here.

A

ACD

See *active configuration directory*.

active configuration directory (ACD)

The directory in which a copy of the HACMP-specific ODM object classes are stored. All the HACMP daemons, scripts, and utilities reference the ODM data stored in the ACD at runtime. See also *default configuration directory*, *staging configuration directory*, and *HACMP configuration database*.

activation code

A password used to activate inactive processors or memory in CUoD. Each activation code is uniquely created for a system and requires the system Vital Product Data (VPD) to ensure correctness.

In HACMP SMIT and in HACMP guides, the activation code is also referred to as the *license key*.

active varyon

A state in which an enhanced concurrent volume group can be activated on a cluster node. Is equivalent to “ordinary” varyon.

See also *passive varyon* and *fast disk takeover*.

adapter

See *network interface*.

Address Resolution Protocol (ARP)

The Internet communication protocol used to dynamically map Internet addresses to physical (hardware) addresses on local area networks.

The `/usr/sbin/cluster/etc/clinfo.rc` script, which is called by the `clinfo` utility whenever a network or node event occurs, updates the system’s ARP cache. This ensures that the IP addresses of all cluster nodes are updated after an IP address takeover. The script can be further customized to handle site-specific needs.

See also *gratuitous ARP*.

AIX Fast Connect

An application that allows you to share services between an AIX station and PC-based clients. This application is integrated with HACMP so that it can be configured as a highly available resource, without an application server or additional scripts.

aliasing

A capability of TCP/IP that allows configuring of multiple IP addresses on a physical network interface. The second (and subsequent) addresses on an interface are the “alias” addresses.

See also *IP address takeover* and *IP address takeover through IP aliases*, and *heartbeat over IP aliases*.

alternate hardware address (MAC address)

A logical datalink layer address that can be assigned to the network interface. It is not preserved during IPL. Another term used for Alternate Hardware Address is Locally Administered Address (LAA, as opposed to Universal Administered Address). See also *hardware address swapping*.

anti-collocation

A resource distribution policy for HACMP service labels on IPAT via IP aliasing networks. Much like the node distribution policy available for resource groups, the options allow you to specify the placement of individual resources with relation to other resources.

The *anti-collocation* resource level location policy ensures all resources of this type are allocated on the first physical resource that is

not already serving, or serving the least number of, a resource of the same type. See also *collocation*, *anti-collocation with persistent labels* and *collocation with persistent*.

anti-collocation with persistent labels

A resource distribution policy for HACMP service labels on IPAT via IP aliasing networks. Much like the node distribution policy available for resource groups, the options allow you to specify the placement of individual resources with relation to other resources.

The *anti-collocation with persistent labels* policy ensures that all resources of this type are not placed on the same physical interface as the persistent IP label configured on the node. See also *collocation*, *anti-collocation*, and *collocation with persistent labels*.

APAR

IBM Authorized Program Analysis Report. A report of a problem caused by a suspected defect in a current unaltered release of an IBM program.

application

A software that runs on top of HACMP and uses the services HACMP provides. An application can incorporate calls from the HACMP application programming interfaces (APIs).

Applications can use the *Clinfo* API to obtain cluster status, or *Simple Network Management Protocol* (SNMP) operations to access the *HACMP Management Information Base* (MIB).

See also *multi-tiered application*, *application monitoring* and *application server*.

application monitoring

A function that lets you specify applications to be monitored by multiple monitors in HACMP, and have the applications restarted when a problem is detected. You can configure two types of application monitors:

- *Process application monitors* use the *RSCT* subsystem *RMC* to detect the unexpected termination of one or more processes of an application.
- *User-defined application monitors* use custom scripts and polling intervals to check the health of an application at a user-defined polling interval.

application server

A collection of the application's start and stop scripts that is defined to HACMP. The application here is the software that runs on top of HACMP and uses HACMP services. Application servers are configured as cluster *resources* that are made highly available.

application provisioning

A task of configuring the minimum and the desired amounts of CPU and memory that you would like to provision for your application server with the use of *DLPAR* function and *Capacity Upgrade on Demand* (CUoD).

See also *Capacity Upgrade on Demand* (CUoD), and *DLPAR*.

automatic discovery and import of volume groups

An option for configuring shared resources in SMIT. When adding a volume group to a resource group, HACMP can gather information about all shared volume groups available for import on all nodes defined to the resource group, and import the volume groups to all the destination nodes, if needed.

automatic discovery of cluster configuration

An HACMP process of gathering information about cluster components. For instance, when you add or change initial communication paths (to new nodes), or add networks, communication interfaces/devices and volume groups, HACMP can gather information about all defined host names, subnets, network interfaces and volume groups available in the existing cluster configuration. You can then configure cluster components from picklists, with some SMIT fields automatically filled in for you. We recommend to run the HACMP discovery process each time you add a volume group. This way, HACMP picks up the volume group's name for display in picklists. Discovery is especially useful if you plan to use C-SPOC after adding a volume group.

Automatic Error Notification

An HACMP utility that configures AIX Error Notification methods automatically for certain device errors. You can enable automatic error notification through SMIT.

B
boot address

No longer used as a term in HACMP.

bring resource groups offline: method of stopping cluster services

In HACMP 5.4.1, one method of immediately stopping HACMP cluster services on a cluster node.

When a resource group is intentionally taken out of the cluster and placed in the OFFLINE state, the Cluster Manager on that node releases the resources. Other nodes do not take over the resources.

Stopping the cluster services with the option to bring resource groups offline is equivalent to the cluster stop mode “graceful” used in releases prior to HACMP 5.4.

See also these terms:

- *stopping HACMP cluster services*
- *move resource groups: method of stopping cluster services*
- *unmanage a resource group: method of stopping cluster services.*

broadcast

A network packet delivery option where a given packet can be sent to all hosts attached to the network.

C
Capacity Upgrade on Demand (CUoD)

A facility in some System p™ processors that lets you acquire—but not pay for—a fully configured system. The additional CPUs and memory, while physically present, are not used until you decide that the additional capacity you need is worth the cost. This provides you with a fast and easy upgrade in capacity to meet peak or unexpected loads.

cascading resource group

Term no longer used in HACMP. See *resource group*, *resource group policies* and *node priority policy*.

Cascading without Fallback (CWOFF)

Term no longer used in HACMP. Instead, for a non-concurrent resource group, you can specify a Never Fallback policy. See also *fallback*, *resource group* and *resource group policies*.

child resource group

A child resource group depends on a parent resource group. It is acquired on any node in the cluster only after the parent resource group has been activated. The child resource group depends on some application services that the parent resource group provides.

Upon resource group release (during failover or stopping cluster services, for example)

HACMP brings offline a child resource group before a parent resource group is taken offline.

Upon failover of a parent resource group, a child resource group is temporarily stopped and restarted on the same or another node in the cluster.

If a parent resource group fails to get started in the cluster, a child resource group goes into the ERROR state (its dependency is not met).

See also *dependent resource groups*, and *parent resource group*.

clcmd

See *Cluster Communication Daemon*.

clconvert

A utility that converts the configuration of an earlier version of HACMP to the current version. The utility creates new data structures and objects, redefines field names and values within data structures, and ensures data integrity between versions of the HACMP software. The **clconvert** utility is called automatically when HACMP is installed on a system that has a previous version of HACMP installed. Additionally, a *clconvert_snapshot* utility is available to convert snapshot files to the current version of HACMP.

clconvert_snapshot

A utility that converts cluster snapshot files from an earlier version to the current version of HACMP when upgrading.

See also *clconvert*, *cluster snapshot*.

cldare utility

The utility used for *DARE* (dynamic automatic reconfiguration) tasks: topology and resource configuration changes that can be made without stopping cluster services.

client

A machine connected to the cluster network so it can access data or services maintained in the cluster.

clstat

See *cluster status utility*.

cluster

A loosely-coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other. HACMP defines relationships among cooperating systems where peer *cluster nodes* provide the services offered by a cluster node should that node be unable to do so.

Cluster Communication Daemon (clcomd)

This utility provides a common communication path for all HACMP utilities. This communication path is secured using either Standard Security, or Enhanced Security on the systems that run *PSSP*, using Kerberos.

For Standard Security, the principle of least privilege is used for remote command execution for commands in */usr/es/sbin/cluster*. This ensures that no arbitrary command can run on a remote node with root privilege. A select set of HACMP commands are considered trusted and allowed to run as root; all other commands run as user nobody.

See also *Kerberos*.

cluster configuration

A user definition of all cluster components. Component information is stored in the ODM. Components include cluster name and ID and information about member nodes, *network interface*, and *network module*. See also *dynamic reconfiguration*.

cluster configuration database

See *HACMP Configuration Database*.

cluster diagnostic utility

An HACMP utility available from SMIT that provides a common interface to several HACMP and AIX diagnostic tools for troubleshooting an HACMP cluster.

cluster environment

See *cluster configuration*.

cluster event

Represents a change in a cluster's composition that the Cluster Manager recognizes and can respond to. The following are some examples of the major cluster events:

- **node_up**—a node joins or rejoins the cluster
- **node_down**—a node detaches from the cluster
- **network_down**—a network failure
- **network_up**—a network is added or restored
- **swap_adapter**—a service network interface card fails.
- **site_up**—a site joins or rejoins the cluster
- **site_down**—a site detaches from the cluster.

cluster hardware

Hardware included in the cluster: disks and disk devices, processors, network interfaces and networks.

The Cluster Manager monitors all but disks and disk devices. For disk monitoring, configure disk heartbeating networks. The RSCT software components monitor networks. Use user-defined events for monitoring of other hardware-related parameters, such as CPU.

Cluster Information Program (Clinfo)

SNMP monitor program. **clinfo**, running on a client machine or on a cluster node, obtains updated cluster information from the Cluster Manager. It then makes information about the state of an HACMP cluster, nodes, and networks available to clients and applications running in a clustered environment. **clinfo** and its associated APIs enable developers to write applications that recognize and respond to changes within a cluster.

Cluster Manager (clstrmgr)

An HACMP utility responsible for monitoring local hardware and software subsystems, tracking the state of the cluster peers, and triggering cluster events when there is a change in the status of the cluster.

The Cluster Manager starts and runs independently of the RSCT stack and therefore starts immediately after installation. The HACMP MIB is accessible as soon as the system comes up, even without a configured cluster. When started, the Cluster Manager responds to SNMP (Simple Network Monitoring Control) requests; however, until a full configuration is read, only a subset of MIB variables is available.

In HACMP clusters, the RSCT software components—Group Services, Resource Monitoring and Control (RMC), and Topology Services—are responsible for most of the cluster monitoring tasks.

cluster monitoring with Tivoli

A function that allows monitoring of HACMP clusters through the Tivoli management interface.

cluster name

A user-defined ASCII text string that uniquely identifies a cluster in an HACMP system.

cluster node

A system unit that participates in an HACMP cluster as a server. It can be an IBM System p™ system unit, or a *logical partition (LPAR)* within a partitioned System p™ server system unit; a Cluster 1600 unit, or an RS/6000 unit that is designed for server applications and that meets the minimum requirements for internal memory, internal disk, and I/O slots.

Cluster Resource Manager

The Cluster Resource Manager centralizes the storage of and publishes updated information about HACMP-defined resource groups. The Resource Manager on each node coordinates information gathered from the HACMP global ODM, the Cluster Manager, and other Resource Managers in the cluster to maintain updated information about the content, location, and status of all HACMP resource groups.

cluster services

The cluster high availability services, such as the Cluster Manager and other services that run on the nodes, when the HACMP software is running on the node. When the HACMP cluster services are running, HACMP is monitoring the cluster resources.

Also, the term cluster services is used to describe the resources and data maintained on the cluster for access by clients (Clinfo) and their applications.

Cluster-Single Point of Control (C-SPOC)

An HACMP administrative utility that lets you perform command maintenance tasks on all nodes in a cluster by running operations from a single node in SMIT (the **System Management** panel) or from the command line on any node. C-SPOC provides commands to manage cluster services, shared LVM components, and user and group accounts.

Cluster SMUX Peer (clsmuxpd)

Obsolete. Starting with HACMP v. 5.3, **clinfo** accesses the MIB directly. An HACMP daemon that provides Simple Network Management Protocol (SNMP) support to client applications. It maintains cluster status information in a special HACMP MIB. Developers can use one of the Clinfo APIs or standard SNMP routines to access the cluster information in this MIB.

cluster snapshot

An ASCII file containing a record of all the data that defines a particular cluster configuration. The cluster snapshot utility lets you save and restore a particular cluster configuration, by applying a cluster snapshot. Optionally, you can save a snapshot without logs. The *clconvert_snapshot* utility is available to convert snapshot files to the current version of HACMP during an upgrade. The cluster snapshot utility is a shell script that you can launch via SMIT or from the command line.

cluster status utility (clstat)

An HACMP utility that reports on cluster, node, network interface, and resource group status. Runs on ASCII and X terminals and as a web-based interface.

Cluster Test Tool

A utility that evaluates the recovery capabilities of an HACMP cluster. The Cluster Test Tool is available for testing a new cluster before it becomes part of a production environment, and for testing configuration changes to an existing cluster, when the cluster is not in service.

cluster topology

See *topology*.

cluster verification utility

An HACMP utility available from SMIT that verifies the HACMP-specific modifications to AIX software and checks for correct HACMP configuration of many cluster components.

collocation

A resource distribution policy for HACMP service labels on IPAT via IP aliasing networks. Much like the node distribution policy available for resource groups, the options allow you to specify the placement of individual resources with relation to other resources.

The *collocation* resource level location policy ensures that all resources of this type are placed on the same physical resource. See also *anti-collocation*, *anti-collocation with persistent labels* and *collocation with persistent labels*.

collocation with persistent labels

A resource distribution policy for HACMP service labels on IPAT via IP aliasing networks. Much like the node distribution policy available for resource groups, the options allow you to specify the placement of individual resources with relation to other resources.

The *collocation with persistent labels* resource level location policy ensures that all service labels will be placed on the same interface as the persistent IP label. See also *collocation*, *anti-collocation*, and *anti-collocation with persistent labels*.

Command Execution Language (CEL)

The language used to create commands that work across cluster nodes. HACMP supports 32-node clusters. *C-SPOC* commands use CEL.

communication device

The device entry that identifies a network interface, for example tty0, en0.

communication interface

The term used to define the relationship between a network interface and the associated IP label/address, function and other parameters necessary for HACMP. It is important to distinguish the communication interface, a grouping of IP-related items, and a *network interface* (for example, en0), which is a single, logical unit associated with a particular network device.

communication path

The *IP label/address* or a Fully-Qualified Domain Name that can be used to communicate with the remote node. The communication path lets HACMP automatically discover the cluster configuration from a single node.

concurrent access

Simultaneous access to a shared volume group or a raw disk by two or more nodes. In this configuration, all the nodes defined for concurrent access to a shared volume group are owners of the shared resources associated with the volume group or raw disk.

If a node in a concurrent access environment fails, it releases the shared volume group or disk, along with all its resources. Access to the shared volume group or disk is, however, continuously available, as long as another node is active. Applications can switch to another node immediately.

connection authentication

A process that verifies information about the origination of a message. HACMP verifies the source IP Address and port number.

concurrent capable volume group

A volume group that can be varied on in either non-concurrent or concurrent access mode in an HACMP cluster environment.

concurrent resource group

A concurrent resource group attempts to become activated on all the nodes in the node list; there is no priority among the owner nodes. A concurrent resource group has the startup policy Online On All Available Nodes.

If one node fails, the other nodes continue to offer the service; the group does not move. You must make use of concurrent access software to configure this type of group.

See also *resource group* and *non-concurrent resource group*.

Configuration Assistant

A utility for quickly configuring HACMP. A *Two-Node Configuration Assistant* is responsible for configuring HACMP once you give it the name of the application, the service IP label, the pathnames to the start and stop scripts and the communication path to the takeover node for the application.

control packet

A message used for communication between Cluster Managers. A means of communication for coordinating the process of nodes joining or leaving the cluster, and for coordinating event handling.

corrective action

A user-authorized action taken during the cluster verification when certain configuration errors can be corrected.

cross-mounting

See *NFS cross-mounting*.

cross-site LVM mirroring

A mechanism to replicate data between the disk subsystems located and configured at two remote sites, to ensure disaster recovery. Disks use remote LVM mirroring (via a SAN, for example). HACMP automatically synchronizes mirrors after a disk or node failure and subsequent reintegration.

C-SPOC

See *Cluster-Single Point of Control*.

custom resource group

A term no longer used in HACMP. See *resource group* and *resource group policies*.

D

DARE (Dynamic Automatic Reconfiguration)

Changes in the cluster configuration that you can make dynamically, without stopping the cluster services. The changes take effect upon the next cluster synchronization.

data link control (DLC)

A set of rules used by nodes in an SNA network on a data link (such as Token Ring or Ethernet) to accomplish an orderly exchange of information.

data path device (VPATH)

A virtual device that represents a single *ESS* volume (hdisk device) that can be accessed via one or more disk adapters. The VPATH device manages the multiple connection paths to automatically provide load balancing and path fallover protection to its associated volume device.

DCD

See *default configuration directory*.

deadman switch (DMS)

The deadman switch assures that if the system pauses long enough it shuts itself down. It is a kernel routine that solves the occasional problem caused by a system that experiences CPU saturation or long I/O cycles and then continues normal processing.

The DMS runs on a timer with:

```
reset interval = 1 second
trip interval = 60 seconds.
```

If the timer runs out, the node shuts down.

RSTC topology services resets the timer every 1 second. By default, the deadman switch is enabled when you start the cluster services. The following commands provide operational information:

```
/usr/sbin/rsct/bin/hatsdmsinfo
lssrc -lstopsvcs
```

default configuration directory (DCD)

The directory in which the system default ODM object classes are stored, including HACMP for AIX specific object classes. For more information, see the *Concepts and Facilities Guide*.

delayed fallback timer

An attribute of a resource group. A timer that you can configure that lets the resource group fall back to a higher priority node at a specified date, or using the recurring daily, weekly, monthly or yearly delayed fallback times.

D

The delayed fallback timer is used when the resource group has the fallover policy Fallback to the Higher Priority Node in the List and the timer is first defined and then assigned to a resource group.

See also *resource group*, *resource group policies*, and *settling time*.

dependent resource groups

A relationship between two or more resource groups that you can specify to control the timing of events or the location of the resource groups. For timing of events, the dependency is parent/child, where services (or an application) in one resource group must be started *after* services in another resource group are running successfully.

If resource group A depends on resource group B, resource group A is brought online on a node only after resource group B is online on any node in the cluster. In other words, a child resource group A depends on a parent resource group B. Clusters with dependent resource groups allow including multi-tiered applications in resource groups.

For location dependencies, you can specify that certain resource groups must always be online on the same node or online on different nodes.

See also *location dependency*, *multi-tiered application*, *resource group*, *parent resource group*, and *child resource group*.

Diagnostic Group Shutdown Partition (DGSP)

See *Merge* and *Domain Merge*.

discovery

See *Automatic Discovery of Cluster Configuration*.

direct serial connection

A point-to-point serial connection between cluster nodes. Prevents node isolation (cluster partitioning) and the spurious disk takeovers that can result from node isolation. The direct serial network can be the SCSI-3 bus using Target Mode SCSI, Target Mode SSA, or an RS232 serial network.

disk events

Disk and disk adapter failures. HACMP does not directly handle disk events. Rather, these failures are handled by AIX through LVM mirroring on SCSI disk devices, SSA disks, and

by internal data redundancy on RAID device disk arrays and the IBM *ESS* storage device family.

The *error notification* facility can be used if additional action is desired in the case of a disk event.

disk heartbeat (diskhb) network

A non-IP point-to-point network over a disk in an *enhanced concurrent mode* volume group. Two nodes, connected to the disk, periodically write heartbeat messages and read heartbeat messages (written by the other node) on a small, non-data portion of the disk. Used for error detection. See also *heartbeat over disk*, and *enhanced concurrent mode*.

disk management

The AIX LVM manages disks at the logical level. HACMP uses AIX LVM facilities to provide high availability—in particular, volume groups and disk mirroring.

disk mirroring

A method of minimizing the effect of a disk failure by duplicating the contents of the disk. If a disk fails, the node can access the mirrored disk and continue to work.

disk takeover

Occurs when the node that currently owns the disk fails and a surviving node assumes control of the shared disk so that the disk remains available (for non-concurrent access configurations). See also *fast disk takeover*.

distribution policy

A load-balancing policy that you can specify for the resource group, and that HACMP uses to place resource groups on a node when a node joins a cluster.

In HACMP, you can select a node-based distribution policy. HACMP ensures that only one resource group is acquired on a node during node startup, if that node belongs to the nodelist.

Note: Network-based distribution is no longer supported.

See also *resource group*, *resource group policies* and *startup*.

distribution preference

A network-wide attribute that you set in HACMP to control the placement of the service IP label aliases on the physical network interface cards on the nodes in the cluster. Configuring a distribution preference lets you customize the load balancing for service IP labels in the cluster, taking into account the persistent IP labels previously assigned on the nodes. It also lets you configure the type of distribution preference suitable for the VPN firewall external connectivity requirements. The four types of distribution preference are: collocation, anti-collocation, anti-collocation with persistent, and collocation with persistent.

See also *collocation*, *anti-collocation*, *anti-collocation with persistent* and *collocation with persistent*.

DMS

See *deadman switch*.

domain merge

When communications paths between the nodes are restored in a partitioned cluster, the domain merge process halts all nodes on one side of the partition to avoid data corruption.

Dynamic Logical Partitioning (DLPAR)

A facility in some System p™ processors that provides the ability to logically attach and detach a *managed system's* resources to and from a logical partition's operating system without rebooting. Some of the features of DLPAR include:

- *Capacity Upgrade on Demand (CUoD)*, a feature of the System p™, which allows you to activate preinstalled but inactive processors as resource requirements change.
- The Dynamic Processor Deallocation feature of the System p™ servers, and on some SMP models. It lets a processor be taken offline dynamically when an internal threshold of recoverable errors is exceeded. DLPAR substitutes the inactive processor, if one exists, for the processor that is suspected of being defective. This online switch does not impact applications and kernel extensions. This function is not supported by HACMP.

- Cross-partition workload management, which is particularly important for server consolidation, and can be used to manage system resources across partitions. This function is not supported by HACMP.

See also *Capacity Upgrade on Demand (CUoD)* and *managed system*.

dynamic node priority policy

A user-defined method for the HACMP Cluster Manager to select takeover nodes to receive a *resource group*. Policies are based on the value of an RSCT *Resource Monitoring and Control resource attribute*, such as most available disk space.

See also *resource group policies*.

dynamic automatic reconfiguration (DARE)

The process where changes made to the cluster configuration on one node are synchronized across all cluster nodes and the changed configuration becomes the currently active configuration. Dynamic reconfiguration is triggered by synchronization when cluster services are running on the local node, and occurs while the cluster is running.

E**enhanced concurrent mode**

Added in AIX. Instances of CLVM (Concurrent LVM) coordinate changes between nodes through the Group Services component of the Reliable Scalable Cluster Technology (RSCT) facility in AIX. Group Services protocols flow over the communications links between the cluster nodes. This implies a number of differences between the support for enhanced concurrent mode and prior concurrent mode capabilities.

Any disk supported by HACMP for attachment to multiple nodes can be used in enhanced concurrent mode; the special facilities of Serial Storage Architecture (SSA) disks are not required.

Enterprise Storage Server (ESS)

An IBM disk storage system that provides enhanced scalability and high availability, and integrates with Storage Area Networks (SAN).

See also *Peer-to-Peer Remote Copy (PPRC)*, *site*.

F

Eprimary

A command provided with PSSP to designate a primary switch node on systems that run PSSP.

Error Emulation

The function that allows you to test notify methods defined in the AIX Error Notification facility by simulating an error. Once an error has been emulated, your customized response method is run, thereby allowing you to check that the notify method will work as planned in the event of an actual error.

Error Notification

An AIX facility that lets you detect an event not specifically monitored by the HACMP software—such as a disk adapter failure—and program a response to the event.

ESS

See *Enterprise Storage Server*.

Ethernet

A type of network with 10 Megabits/second, 100 Megabits/second, or 1000 Megabits/second. Nodes and clients that communicate over an Ethernet network require the Ethernet High-Performance LAN Adapter.

Event Emulator utility

A utility that emulates cluster events and dynamic reconfiguration synchronization events. The Event Emulator runs event scripts that produce output, but do not affect the cluster configuration or status.

Event Management Services (EMS)

An *RSCT* distributed subsystem providing a set of high availability services. Replaced by *Resource Monitoring and Control (RMC)*.

Note: Oracle registers still depend on network-related notification provided by EMS.

event notification

A command that sends mail indicating that an event is about to happen or that an event has just occurred, along with the success or failure of the event.

event recovery

A command that attempts to recover from an event failure.

event retry

The number of times an event command may be rerun if it fails. You can also specify the number of times the system will attempt to execute the recovery command.

event summary

An easy-to-read summary of important event details that appears after each event in the **hacmp.out** file. You can view in one place a compilation of event summaries from current and past **hacmp.out** files. Or, you can save event summaries to a file. Use the **Display Event Summaries** panel in SMIT.

events

See *cluster event*.

F

failure detection rate

The amount of time the cluster takes to detect a failure. You can affect this time by setting a shorter or longer heartbeat rate for a Network Module. Failure detection rate has default values for each network type, or it can be customized.

fallback

The process when a *joining or reintegrating* node acquires resources previously owned by another node. You can specify a fallback policy for a resource group. See also *startup, fallover* and *resource group policies*.

fallback timer

See *delayed fallback timer*.

fallover

The process of an *active* node acquiring resources previously owned by another node, in order to maintain availability of those resources. You can specify a fallover policy for a resource group. See also *startup, fallback* and *resource group policies*.

fast disk takeover

A mechanism of taking over *enhanced concurrent volume groups*. During fast disk takeover, HACMP skips the extra processing needed to break the disk reserves, or update and synchronize the LVM information by running *lazy update*. HACMP automatically detects enhanced concurrent volume groups and

ensures that the faster option for volume group takeover is launched in the event of a node failure.

Fast disk takeover is used for enhanced concurrent volume groups that are included in *non-concurrent* resource groups.

See also *enhanced concurrent mode* and *disk takeover*.

Fast Node Failure Detection

A process where HACMP can quickly detect or a node failure in the cluster. When a node fails, HACMP 5.4.1 uses disk heartbeating to place a *departing* message on the shared disk, so that neighboring nodes are aware of the node failure within one heartbeat period (*hbrate*) and can distribute this information throughout the cluster. Remote nodes that share the disks receive this message and broadcast that the node has failed. Directly broadcasting the node failure event greatly reduces the time it takes for the entire cluster to become aware of the failure compared to waiting for the missed heartbeats, and therefore HACMP can take over critical resources faster.

See also *disk heartbeat (diskhb) network*.

fast recovery

A method speeding up failover in large clusters. Configure your system for fast recovery through the SMIT HACMP Extended Resources Configuration screen.

fault tolerance

A model for maintaining computer data availability. It relies on specialized hardware to detect a hardware fault and to instantly switch to a redundant hardware component—whether the failed component is a processor, memory board, power supply, I/O subsystem, or storage subsystem.

FDDI (Fiber Distributed Data Interface)

An industry-standard fiber optic network technology that provides high-speed (100 Megabits per second) communication among cluster members. Nodes and clients that communicate over a FDDI network require either the FDDI Dual Ring Adapter, or the FDDI Single Ring Adapter and (in most cases) the FDDI Concentrator.

FDDI hardware address takeover

An HACMP function that allows takeover of *FDDI* hardware addresses, preserving *Address Resolution Protocol (ARP) caches* in the event of a failover.

See also *ARP (Address Resolution Protocol)*, and *FDDI (Fiber Distributed Data Interface)*.

file collection

See *HACMP file collection*.

forced shutdown

No longer used in HACMP. In previous releases, forced shutdown was a method of immediately shutting down HACMP on a cluster node. The node retains control of its *resources*; other nodes view it as a *graceful shutdown*.

See *move resource groups: method of stopping cluster services*.

forced stop

See *move resource groups: method of stopping cluster services*.

forced varyon

A method of forcing an activation of a volume group that you can specify in SMIT. You can specify that HACMP attempts a forced varyon operation for all volume groups in the resource group in the event that a normal varyon operation for the volume group fails due to a lack of quorum, and if HACMP finds at least one complete copy of every logical volume for this volume group.

See also *enhanced concurrent mode*.

free pool

The number of CPUs and the amount of memory that can be dynamically allocated by HACMP through HMC to the LPAR, should the LPAR require additional resources. The free pool is the difference of the total amount of resources on the frame minus the resources that are currently used by the LPARs.

The free pool includes resources on a particular frame *only*. For instance, if a cluster is configured with LPARs that reside on frames A and B, HACMP does not request resources from a pool on frame B for an LPAR that resides on frame A.

See also *DLPAR, Hardware Management Console (HMC) and Capacity Upgrade on Demand (CUoD)*.

Fully Qualified Domain Name (FQDN)

The *IP label* and the domain name of a host, for example, “server.ibm.com”.

G

General Parallel File System (GPFS)

IBM shared disk file system. Allows users shared access to files that may span multiple disk drives on multiple nodes.

Geographic Logical Volume Manager (GLVM)

Increases data availability for IBM volumes that use it to copy data to a remote site for disaster recovery purposes. HACMP/XD for GLVM takes advantage of the following components to reduce downtime and recovery time during disaster recovery:

- AIX GLVM data mirroring and synchronization
- TCP/IP-based unlimited distance network support
- HACMP cluster management.

global network

A logical collection of multiple HACMP networks. This reduces the probability of network partitions that can cause the cluster nodes on one side of a partition to go down. All networks grouped in a global network must be of the same type (all Ethernet, for example).

global ODM

Software that extends the ODM. It enables one system to update the ODM entries of other systems. The global ODM allows an entire HACMP cluster to be configured from a single node.

GSM

Global Systems for Mobile Communications. A leading digital cellular system. GSM was first introduced in 1991, and by the end of 1997 was available in more than 100 countries; GSM is the standard in Europe and Asia.

GPFS

See *General Parallel File System*.

GPFS cluster data

GPFS configuration data stored on the primary and secondary GPFS cluster data servers.

grace period

The time period during a cluster event when interface reconfiguration may be in progress and networks are not monitored. Network Grace Period is network specific and has default values for each network type, or can be customized.

graceful shutdown

No longer used in HACMP. In previous releases, when a node leaves the cluster intentionally, the Cluster Manager on that node runs the `node_down_local` script to release its resources. Other nodes do *not* take over the resources.

See *bring resource groups offline: method of stopping cluster services*

graceful stop

No longer used in HACMP. See *graceful shutdown*.

graceful shutdown with takeover

No longer used in HACMP. Instead, the following option is used in SMIT: “stop cluster services: move resource groups”.

See *move resource groups: method of stopping cluster services*.

gratuitous ARP (Address Resolution Protocol)

Typically, entries in an ARP cache expire after about twenty minutes. A gratuitous ARP is an ARP reply when there was no ARP request. A gratuitous ARP reply recommends that all hosts on the network receive the ARP reply and refresh their ARP cache. For IP address takeover through IP aliasing to be successful, systems and devices connected to the network must be configured to support gratuitous ARP in HACMP.

See also *IP address Takeover through IP aliasing*, and *Address Resolution Protocol (ARP)*.

Group Services

The Group Services RSCT subsystem provides client subsystems with a distributed coordination and synchronization service. Each client subsystem forms one or more **groups** by having its processes connect to the Group Services subsystem and uses the various Group

Services interfaces. A process of a client subsystem is called a **GS client**. For example, *Resource Monitoring and Control (RMC)* is a Group Services client subsystem. The RMC daemon on each node is a GS client.

H

HACMP (High Availability Cluster Multi-Processing) for AIX

Licensed Program (LP) that provides custom software that recognizes changes within a cluster and coordinates the use of AIX features to create a highly available environment for critical data and applications.

HACMP Configuration Database

The data repository for HACMP configuration information. This data repository is the AIX Object Data Model (ODM).

HACMP/ES

Licensed Program (LP) that provides Enhanced Scalability to the *HACMP for AIX* LP. An HACMP Enhanced Scalability cluster can include up to 32 nodes.

Starting with version 5.1, HACMP/ES is referred to as HACMP. The first three of the four features that existed up through 4.5 (HAS, CRM, ES, ESCRM) are discontinued. ESCRM (Enhanced Scalability Concurrent Resource Manager) includes all features with enhanced scalability plus concurrent resource management.

HACMP/XD

The HACMP/XD feature provides software solutions for disaster recovery. Added to the base HACMP software, they enable a cluster to operate over extended distances at two sites. The solutions include:

- HACMP/XD for GLVM
- HACMP/XD for Metro Mirror
- HACMP/XD for HAGEO Technology

See also *Metro Mirror*, *GLVM*, *HAGEO*.

HACMP file collection

A list of files that can be automatically propagated and synchronized across the cluster. This ensures each node has the same version of key files, and that HACMP and applications run correctly. Such files include event scripts,

application scripts, and some AIX and HACMP configuration files. Optionally, you can specify that HACMP synchronizes a specific file collection on all nodes to ensure that it is the same, before running the rest of the cluster verification process.

HAGEO (High Availability Geographic Cluster) for AIX

Feature that extends an HACMP cluster to encompass two physically separate data centers. Data entered at one site is sent across a point-to-point TCP/IP network and mirrored at a second, geographically distant location. This software provides highly available disaster recovery.

hardware address

A unique identifier (for example, 00:06:29:ac:fb:11) that is programmed into every network card for an Ethernet interface. See also *alternate hardware address*.

hardware address swapping

Works in conjunction with *IP address takeover via IP Replacement*. When an IP address is taken over, its hardware address is also taken over by the network interface. It can then continue to provide service to clients without requiring an ARP cache update.

Hardware Management Console (HMC)

An interface that lets you manage all DLPAR operations on several or all LPARs created on the frame, collect CUoD system profile information and enter *activation codes* for CUoD.

For integration with HACMP, HMC should have a TCP/IP connection to the LPAR and a configured IP label through which a connection will be established. The **lshmc** command displays the HMC configuration.

See also *Capacity Upgrade on Demand (CUoD)*.

HAView utility

Cluster monitoring utility that lets system administrators view HACMP clusters through the TME 10 NetView for AIX graphical network management interface. HAView extends NetView services to let users monitor HACMP clusters and cluster components across a network from a single node. HAView creates symbols that reflect the state of all

nodes, networks, and network interface objects associated in a cluster. In HACMP, the HAView utility also monitors resource groups.

heartbeat

State-of-health message exchanged between network modules. The Cluster Manager uses these to track the membership and status of the cluster nodes. A means of detecting failure in the cluster.

heartbeat over disk

Device-based point-to-point communication path for sending *heartbeat* messages between two nodes using disk I/O. Uses external disk bus and shared disk. Volume groups must be in *enhanced concurrent mode*. Configure to use **diskhb** as the network type.

See also *enhanced concurrent mode*.

heartbeat over IP aliases

Process used by RSCT to send *heartbeat* messages between cluster nodes over *IP alias addresses*. You configure a private address range and HACMP generates an address for each interface on the network to use for heartbeat messages.

high availability

A model for maintaining computer service availability. It views service availability not as a series of replicated physical components, but rather as a set of system-wide, shared resources that cooperate to guarantee essential services.

home node

The highest priority node for a resource group; it is the first node that is listed in the default nodelist for any non-concurrent resource group.

hostname

A system in a non-HACMP environment is typically identified by a hostname, which is itself tied to a network interface. Since a system usually has only one network interface, the hostname can uniquely identify a system.

A cluster node has more than one network interface and therefore more than one hostname. To avoid ambiguity, each node in an HACMP cluster is identified by a user-defined *node name* which is unique to that cluster.

IBM TotalStorage Disk Systems (DS Series)

The IBM TotalStorage DS6000 and DS8000 series are high-performance, high-capacity enterprise disk storage systems. Both support HACMP/XD for Metro Mirror.

See also *Metro Mirror*, *PPRC*

inactive takeover

No longer used. Resource group/node attribute set by the user for some non-concurrent resource groups in releases prior to HACMP 5.2.

Now you can select the resource group startup, fallover and fallback policies that replicate this behavior. See the chapter on upgrading to HACMP v.5.4.1 in the *Installation Guide*.

intentional fallover

An administrative procedure for allowing the cluster services to continue running while servicing a cluster node. To cause intentional fallover, stop the HACMP cluster services with the option of moving resource groups.

See *move resource groups: method of stopping cluster services*.

interface address

See *network mask*.

interface ID

See *IP label*.

interface name

See *IP label*.

IP Address

A unique identifier, for example 10.10.10.10, for an end-point for the Internet Protocol.

IP address takeover (IPAT)

A mechanism of the acquisition of a service IP label belonging to a cluster node being taken over by another node. This assures the cluster will continue providing network service to clients. See also *network mask*, *Address Resolution Protocol (ARP)*.

IP address takeover (IPAT) via IP Aliases

A mechanism of placing an IP service label of the failed node onto a network interface of another node *as an IP alias*, while *keeping* the interface's original IP label and hardware address.

See also *Gratuitous ARP*.

IP address takeover (IPAT) via IP Replacement

A *service IP label* of the failed network interface card replaces the *IP label* that was previously configured on the network interface card on another node.

IP aliasing

The method by which one or more IP Addresses can be configured to concurrently reside on a single *network interface*, using the **ifconfig** command.

IP label

For TCP/IP networks, the name specified in the */etc/hosts* file or by the Domain Name Service for a specific IP address.

IP Replacement

The method by which one IP address is replaced by another IP address on the same network interface.

IW mode

The mode HACMP uses for the *RMC Independent Workstation* (stand alone) mode.

J**Journaled File System (JFS)**

AIX facility that uses database journaling techniques to protect the integrity of file system meta-data. This facility cannot be used in HACMP concurrent access environments. Use JFS2 in this environment.

K**keepalive**

See *heartbeat*.

Kerberos

An authentication mechanism used on the systems that run *PSSP* that centralizes command authority via one authentication server.

L**label**

See *IP label*.

lazy update

During a lazy update, HACMP does not update on the cluster nodes the definition of those LVM components that do not have an LVM component activated. This information is not updated until failover occurs.

Alternatively, you can de-activate the volume group on the local node and export and import the volume group on all the other cluster nodes.

Lazy update is only an option for those LVM components that are under the control of HACMP.

license key

See *activation code*, and *Capacity Upgrade on Demand (CUoD)*.

Licensed Program (LP)

A software application or operating system program licensed from IBM.

local node

Node on which you are logged in. See also *cluster node*.

location dependency

Resource group location dependency offers you an explicit way to specify that a set of certain resource groups will always be online on the same node, or online on the same site, or that a set of resource groups will always be online on different nodes. You can combine these location policies with parent/child dependencies, to have all children run on the same node while the parent runs on a different node, or, to have all children run on different nodes for better performance.

See also *dependent resource groups*, *multi-tiered application*, *resource group*.

M

logical link

Cluster network, or a connection between nodes that allows messages and other information to pass between cluster nodes.

logical partition (LPAR)

Logical partition within a System p™ server system unit. You can install HACMP on the logical partitions of a single frame and configure each partition as a *cluster node*.

The advanced partitioning features of AIX version 5.2 and up provide the ability to dynamically allocate system CPU, memory, and I/O slot resources (DLPAR). Combining HACMP with DLPAR lets you dynamically redistribute CPU and memory resources to logical partitions that are currently executing application workload, to further support application transition within a single frame.

Do not have all your cluster nodes configured as LPARs within *the same* physical server. This configuration could potentially be a significant single point of failure.

See also *Dynamic Logical Partitioning (DLPAR)*.

Logical Volume Manager (LVM)

AIX facility that manages disks at the logical level. HACMP uses AIX LVM facilities to provide high availability—in particular, volume groups, disk mirroring and fast disk takeover.

LP

See *Licensed Program*.

M

Management Information Base (MIB)

See *Simple Network Management Protocol*.

managed system

A System p™ frame that is LPAR-capable and that is managed by Hardware Management Console (HMC).

See also *Hardware Management Console (HMC)*.

map

A set of related submaps that provides a graphical and hierarchical presentation of a network and its systems.

membership protocol

Mechanism whereby all cluster managers running in a cluster determine which nodes are members of the cluster and handle membership transitions.

merge

See also *domain merge*.

message authentication

A process that verifies the identity of the sender and the integrity of the data. In HACMP, the message is signed and the signature encrypted by a shared (symmetrical) key when sent, and the signature decrypted and verified when received.

HACMP message authentication uses Message Digest version 5 (MD5) to create the digital signatures for the message digest. Message authentication uses the following types of keys to encrypt and decrypt signatures and messages: Data Encryption Standard (DES), Triple DES, or Advanced Encryption Standard (AES).

message encryption

A process that renders data in a packet sent over a network as unreadable. In HACMP, the encryption algorithm uses a shared (symmetrical) key to make data unreadable. The message is encrypted when sent and decrypted when received.

Message encryption uses the same encryption key as message authentication (as selected): Data Encryption Standard (DES), Triple DES, or Advanced Encryption Standard (AES).

Metro Mirror (Synchronous PPRC)

Provides real-time mirroring of logical volumes between two ESS 800 or DS6000s or DS8000s that can be located at a distance of up to 300km away from each other.

It guarantees that the secondary copy is up-to-date by ensuring that the primary copy will be written only if the primary system receives acknowledgment that the secondary copy has been written.

If the acknowledgment of the remote write is *not* received within a fixed period of time, the write is considered to have failed, and is rendered ineligible for storage to disk. At this

point, the application receives an I/O error, and in due course, the failed write I/O is “aged-out” of each NVS.

mirroring

AIX facility for maintaining more than one copy of stored data, to prevent loss of data.

move resource groups: method of stopping cluster services

In HACMP 5.4.1, one method of immediately stopping HACMP cluster services on a cluster node without stopping the applications.

Other nodes take over the resources. Moving resource groups is similar to the cluster stop mode “graceful with takeover” in previous releases.

multicast

A process of logically sending a single message to one or more peer nodes. Not necessarily implemented as a multicast at the network level. The Cluster Manager multicasts control messages at times to speed the distribution of information needed to coordinate action when a node failure occurs.

multithreaded application

An application written using threads. See also *thread-safe*.

multi-tiered application

An application that depends on the startup of another application. Such applications can be included as resources in dependent resource groups. The dependent resource groups may or may not reside on the same node.

See also *dependent resource groups*.

N

network

In an HACMP environment, connects multiple nodes and allows clients to access these nodes. HACMP defines IP and non-IP (device-based) networks. See also *serial network* and *disk heartbeat network*.

network attribute

In HACMP, a network is defined by its name and its attribute: IP-based or non IP-based.

network failure

An event that occurs when a network fails for all the nodes in a cluster. That is, no node is able to communicate across that network.

Network File System (NFS)

A distributed file system for computer networks. Independent of operating system, type of system, and network architecture.

network interface

Connects a node to a network. In AIX: The device entry that identifies a communication interface, for example, *en0*. In HACMP: You configure and identify “network interfaces” that can be used, and potentially can be made highly available by HACMP. You identify such interfaces by configuring *IP addresses/labels* for HACMP, and HACMP finds the underlying interface based on the IP address configured in the AIX ODM.

A node must be configured with at least two network interfaces for each network to which it connects: a service interface that handles cluster traffic, and one or more non-service interfaces (previously referred to as standby interfaces).

network interface attribute

The attributes that define a network interface include an *IP label/IP address* and a function (service, non-service, or persistent). See also *network interface*.

network interface card

The physical network adapter card through which a node is connected to the physical network.

network interface event

Network interface card failure and recovery. HACMP processes these events.

network mask

Determines which bits in an IP address refer to the network and which refer to the host.

An IP address consists of 32 bits divided into four octets. Some of the bits form the network address, the remainder form the host address.

The HACMP software uses the subnet feature of TCP/IP to divide a single physical network into separate logical subnets. In order to use subnets, you must define a network mask for your system.

N

See also *AIX Version 6.1 Networks and communication management* for more information about classes of addresses. Also, ask your network administrator about the class and subnets used at your site.

network module

Cluster Manager interface module. Monitors a network and its interfaces. Each network module monitors one kind of hardware interface using one kind of communication protocol (for example, FDDI, Ethernet). Provides data transport for messages between nodes.

Configuring includes defining a network module name, description, address (IP or /dev), pathname, and heartbeat rate.

network monitoring

See *cluster status utility*.

network name

A symbolic value that identifies a network participating in an HACMP system. Cluster processes need this information to determine which interfaces are connected to the same physical network. The network name is arbitrary, but must be used consistently.

NFS

See *Network File System*.

NFS cross-mounting

An HACMP-specific NFS configuration that describes the situation in which one node is both the NFS server for one file system, *and* the NFS client for another file system, while a second node is the NFS client for the first file system, *and* the NFS server for the second file system. Essentially, each node is part of a “mutual takeover” or “active-active” cluster configuration, both providing and mounting an NFS file system.

node

See *cluster node*.

node attachment

Node event that occurs when a node joins or rejoins the cluster. See also *cluster event*.

node-by-node migration

The process of migrating a cluster from HACMP (base) to the HACMP/ES subsystem on a node-by-node basis without stopping the

cluster services. This migration is not necessary if you are upgrading from a version higher than HACMP 5.1. Starting with HACMP 5.1, the only mode of HACMP available is HACMP/ES.

node detachment

Node event that occurs when a node leaves the cluster voluntarily or fails. See also *cluster event*.

node environment

Includes the definition and distribution of *resources* among participating nodes in a cluster and the setting of *runtime parameters* per node.

node isolation

Complete loss of communication between two or more parts of the cluster. Occurs when all TCP/IP-based networks connecting cluster nodes fail, if a global network is not defined. Each group (one or more) of nodes is completely isolated from the other groups. A cluster experiencing node isolation is referred to as a partitioned cluster.

nodelist

Nodes assigned to participate in the takeover of a given *resource group*. See also *resource group*.

node name

A symbolic value (string) that identifies a node participating in an HACMP system. (This is not the *hostname*.)

node priority policy

User-selected method that the HACMP Cluster Manager uses to determine the node order for takeover of a non-concurrent resource group. The method can be the order of the participating nodelist, or a *dynamic node priority policy* that uses an RSCT *Resource Monitoring and Control (RMC) resource attribute*.

See also *resource group policies*.

node state

Current situation of a cluster node. States include: Stable (a node is known to be up or down), reconfiguring (event scripts are handling an event that calls for reorganization of one or more cluster resources), unstable (a node is joining or leaving the cluster). Monitored by *cluster status utility*.

non-concurrent access

Term used when referring to shared resources, especially disks. In non-concurrent environments, only one connection is active at any given time, and the node with the active connection owns the disk.

non-concurrent resource group

A resource group that does not have the startup policy Online on All Available Nodes.

See *resource group* and *concurrent resource group*.

nondisruptive startup of HACMP cluster services

A function of HACMP that lets you start and restart HACMP cluster services, and manage applications that are already running without stopping the applications first.

nondisruptive upgrade of HACMP software

A function of HACMP that lets you upgrade HACMP software, while keeping cluster services running and the application continuously available throughout the upgrade

non-persistent priority override location

No longer used in HACMP. In previous releases, a priority override location (POL) could remain in effect only until the cluster services are stopped on all nodes.

non-service network interface

Backs up the service interface. If the service interface fails, the Cluster Manager directs the traffic for that node to the non-service interface. Using another interface on the non-aliased networks ensures that a network interface is not a single point of failure.

non-sticky migration

Term no longer used in HACMP.

O
object

In HAView, a generic term for any entity that HAView discovers and displays on the topology map, or any entity that you add to the topology map. In the AIX object data manager, an instance or member of an object class, conceptually similar to a structure that is a member or array of structures.

In HAView, objects convey to the symbol various semantic attributes that represent an entity.

See also *symbol*.

Object Data Manager (ODM)

AIX facility that stores objects describing HACMP entities such as nodes and resources. HACMP includes a global ODM so that changes made to the system on one node can be synchronized on the whole cluster.

OFFLINE state of a resource group

A resource group is in the OFFLINE state when the resources of the resource group are not active on the node or when HACMP sees them as inactive.

See also *bring resource groups offline: method of stopping cluster services*.

ONLINE state of a resource group

A resource group is in the ONLINE state when HACMP is actively managing the cluster resources on a node.

offline migration

A term used when HACMP is brought offline on all nodes prior to performing the upgrade. During this time, resources are not available.

P
partitioned cluster

See *node isolation*.

parent resource group

A parent resource group must be in the online state before the resource group that depends on it (the child) can be brought online.

A child resource group depends on a parent resource group. It is acquired on any node in the cluster only after the parent resource group has been activated. Typically, the child resource group depends on some application services that the parent resource group provides.

Upon resource group release (during fallover or stopping cluster services, for example)

HACMP brings offline a child resource group before a parent resource group is taken offline.

See also *child resource group* and *dependent resource groups*.

passive varyon

A state in which an enhanced concurrent volume group can be activated on a cluster node. When an enhanced concurrent volume group is varied on in passive mode, the LVM allows only a limited number of read-only operations on the volume group. It does not allow mounting file systems, opening logical volumes, or synchronizing volume groups.

See also *active varyon*, and *fast disk takeover*.

Peer-to-Peer Remote Copy (PPRC)

Synchronous protocol that allows real-time mirroring of data from one Logical Unit (LUN) in an *ESS* system to another LUN in a second *ESS* system. The second *ESS* can be located at a remote site.

Also, see *Enterprise Storage Server*

persistent migration

No longer used in HACMP.

persistent node IP label

An IP alias that can be assigned to a service interface on a specified node. A persistent node IP label is an address that always stays on the same node, co-exists on a network interface that already has a service label defined, does not require installing an additional physical interface card on that node, and *is not* part of any resource group. Assigning a persistent node IP label to a network on a node allows you to have a node-bound address on a cluster network that you can use for administrative purposes to access a specific node in the cluster.

persistent priority override location

No longer used in HACMP. In previous releases, a priority override location (POL) could “persist” after a cluster reboot.

point-to-point

Direct communications link between two nodes. See also *serial network* and *disk heartbeat network*.

post-events

One or more commands specified by the system administrator that execute after the Cluster Manager runs a particular event script. Usually these are site-specific scripts integrated into the system when it is installed.

pre-events

One or more commands specified by the system administrator that execute before the Cluster Manager runs a particular event script. Usually these are site-specific scripts integrated into the system after it is installed.

primary node

User-designated node in a cluster. Typically, it is a node that owns the resource group and that is the first node in the nodelist for that resource group. It also is referred to as a home node for a resource group.

primary instance

The online instance of a *replicated resource group*. The node that has a resource group in the online state will activate all the resources in the resource group. The mirror copy is the *secondary instance*, hosted on a node at the *secondary site*.

primary site

The highest priority site for a *resource group*. A resource group may or may not have a backup site (*secondary site*).

priority

See *takeover priority*, *node priority policy* and *priority override location*.

priority override location (POL)

No longer used in HACMP. Instead, you can specify in SMIT to which destination node to move the resource group, and the group stays on this node until it has to change its host, either due to a failover, or due to your explicit request.

There is one exception to this rule: For a resource group that has a home node specified in its nodelist, if you move it manually to another node, or take it offline on another node, it remains in that state on the destination node either until the next cluster reboot (after which it returns to the home node), or until you manually change its location again with **clRGmove**.

process takeover

After the *takeover* process, the node configured to take over any application server resources runs the script to restart those applications (for example, a database back end process).

PSSP

Parallel System Support Program that runs on IBM System p™ Cluster 1600.

PTF

Program Temporary Fix. A correction for a software problem reported in an *APAR*.

Q**quorum**

An LVM facility that must be considered when configuring logical volume components in an HACMP environment. Quorum determines whether a volume group can be placed online, using the `varyonvg` command or whether it can remain online after a failure of one or more of the physical volumes in the volume group. Quorum checking is enabled by default.

If a processor fails before all copies of the Volume Group Descriptor Area (VGDA) and Volume Group Status Area (VGSA) have been updated or if a disk drive fails while writing, all copies of the information within the volume group may not be the same. The Logical Volume Manager uses a voting scheme called quorum where “majority rules” to determine the correct mapping and status information. With a volume group, quorum is achieved when more than half of the VGDA and VGSA are accessible and identical in content. Having exactly half of the VGDA and VGSA accessible and current does not achieve quorum.

AIX provides an option to disable quorum checking. With quorum disabled, a volume group remains varied on until the last physical volume in the volume group becomes unavailable. Quorum should never be disabled in concurrent access configurations as disk subsystem failures could result in nodes accessing diverged database copies.

R**reconfiguration**

Process whereby the cluster adjusts to changes in membership. See also *dynamic reconfiguration*.

reentrant

See *thread-safe*.

reintegration

Refers to the actions that occur within the cluster when a node that had previously detached from the cluster returns to the cluster. These actions are controlled by the event scripts and, when necessary, by manual intervention.

remote node

Node in the cluster other than the *local node*. See also *cluster node*.

replicated resource

A volume group resource that is mirrored to another disk subsystem in real-time. HACMP supports the configuration of sites that can be configured with geographically mirrored volume groups (GLVM replicated resources) and *HAGEO* GeoMirror devices as replicated resources.

Also, any resource is contained in a *replicated resource group*.

replicated resource group

A resource group that is configured to have both primary and secondary instances, distributed on nodes located at two different sites. See also *primary instance*, *primary site*, *secondary instance*, *secondary site*.

resource

Cluster entity, such as a disk, file system, network interface, or application server, that is made highly available in the cluster.

resource attribute

Term defined and used by the *Resource Monitoring and Control* subsystem of *RSCT*. A characteristic of a system resource. There are two types of resource attributes: persistent attributes and dynamic attributes. Persistent attributes describe enduring characteristics of the resource. Dynamic attributes represent changing characteristics of the resource, as tracked by the *Resource Monitoring and Control* subsystem of *RSCT*.

Dynamic resource attributes defined for *node priority policies* include:

- highest percentage of free memory
- most available processor time
- the disk that is least busy.

See also *resource group policies*.

R

resource class

Term defined and used by the *Resource Monitoring and Control* subsystem of *RSCT*. A set of system resources of the same type. A resource class defines the common characteristics that instances of the resource class can have (in terms of services provided, configuration parameters, and so on).

resource group

Set of cluster resources handled as one unit by HACMP. Configured by the user. You can configure startup, fallover and fallback policies customizing the resource group behavior to your needs. You can also configure the settling time and the distribution policy that affect the resource group startup, and the delayed fallback timer that affects the resource group fallback. Resource groups can be concurrent and non-concurrent.

See also *concurrent resource group*, *non-concurrent resource group*, *distribution policy*, *delayed fallback timer*, *resource group policies* and *settling time*.

resource group management

The process by which you can move resource groups from one node to another, or take them offline or online while the cluster is running, using the *clRGmove* utility from the SMIT interface or from the command line. See also *resource group migration*.

resource group migration

Movement of a resource group from one node to another without an actual fallover.

resource group policies

A set of predefined policies that define how a resource group behaves during startup, fallover and fallback. You can specify different combinations of startup, fallover and fallback policies for a resource group and HACMP verifies that the selected combination is valid.

Startup policies include: Online on Home Node Only, Online on First Available Node, Online Using Distribution Policy, Online on All Available Nodes.

Fallover policies include: Fallover to Next Priority Node in the List, Bring Offline (on Error Node Only), Fallover Using Dynamic Node Priority.

Fallback policies include: Fallback to Higher Priority Node in the List, Never Fallback.

See also *resource group*, *concurrent resource group*, *dependent resource group*, *distribution policy*, *dynamic node priority*, *delayed fallback timer*, *non-concurrent resource group*, and *settling time*.

Resource Group Policy Administrator (RGPA)

A component within the *cluster manager* that calculates what actions should be taken on each *resource group* during a specific *event*.

Resource Monitoring and Control (RMC)

The *RSCT* RMC subsystem replaces *Event Management* resource variables with *resource classes* and *resource attributes* associated with the resource class. The RMC subsystem runs on each HACMP cluster node and provides global access to subsystems and resources throughout the cluster, thus providing a single monitoring and management infrastructure.

See also *Event Management*.

rolling migration

When upgrading HACMP, cluster services are stopped on one node at a time, that node is upgraded and reintegrated into the cluster before the next node is upgraded.

Starting with HACMP 5.4, when using rolling migration to upgrade HACMP on an individual node, you may choose to keep the applications and the resources running continuously on that node, although they will not be highly available during the upgrade.

rotating resource group

A term no longer used in HACMP.

See also *distribution policy*, *resource group* and *resource group policies*.

RS232 serial line

See *serial network*.

RSCT (Reliable Scalable Cluster Technology)

RSCT services provide scalability, notify distributed subsystems of software failure, and coordinate recovery and synchronization among all subsystems in the software stack. This package of services is a prerequisite for HACMP and is packaged with AIX. The package includes *Topology Services*, *Group*

Services, Event Management Services (obsolete since HACMP v. 5.2), and *Resource Monitoring and Control (RMC)*.

runtime parameters

Environmental conditions set per node. For example, activating an NIS server, setting Cluster Manager debug level.

S

SCD

See *staging configuration directory*.

secondary instance

Refers to the backup copy of the *primary instance* of a *replicated resource group*. The secondary instance of the resource group is intended for data backup purposes. HACMP tries to keep the secondary instance in ONLINE_SECONDARY state.

secondary site

The lower priority site for a replicated resource group. Also called the backup site for a resource group.

selective Fallover

An automatically launched function of HACMP that attempts to selectively move only a resource group that has been affected by an individual resource failure to another node in the cluster, rather than moving all resource groups.

serial connections

See *serial network*.

serial network

Non-IP-based network. A SCSI bus using *Target Mode SCSI*, *Target Mode SSA*, or an RS232 serial line that may be used to connect pairs of nodes in a cluster. The purpose of the serial network is to prevent node isolation.

Serial Optical Channel Converter (SOCC)

A proprietary IBM serial optical link that provides high-speed communication between two nodes. Not supported in HACMP starting with version 5.1.

service address

The connection (that has an IP label) between clients and applications. The service address is the address published by the Cluster

Information Program (Clinfo) to application programs that want to use cluster services. HACMP handles the service IP Label/Address as a resource. This *IP Label/Address* may move from one node to another along with the other configured resources in a *resource group*.

settling time

A specified time that HACMP waits before activating an offline resource group on a higher priority node joining the cluster. As a result, after the settling time expires, the resource group gets started on the node that has the highest priority among the nodes that joined the cluster during the settling time interval. If there are no such nodes, the resource group remains offline.

If the highest priority node for this resource group joins, the resource group gets started without delay.

See also *resource group*.

shared disks

Disks configured to serve more than one node. In the HACMP system, shared disks are physically connected to multiple nodes.

shared file systems

File systems configured to serve more than one node. See also *shared disks*.

shared IP address

IP address configured so that it may rotate among cluster nodes. See also *non-concurrent resource group*.

shared logical volume

Logical volume configured to allow access by more than one node.

shared volume group

Volume group configured to allow access by more than one node. Resides entirely on the external disks shared by cluster nodes.

Simple Network Management Protocol (SNMP)

In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached TCP/IP-based networks. SNMP is an application layer protocol. It includes a protocol, a database specification, and a set of data objects. A set of data objects forms a Management Information Base (MIB). Information on devices managed is defined and

stored in the application's MIB. HACMP supplies its own MIB, accessed by the Clinfo program.

single-node cluster

In AIX, you can configure dependencies between applications running on a single node, and application monitors for the applications.

When you install HACMP and add nodes to this configuration, the configuration created on a single node is preserved. You can extend your single-node configuration to a multiple-node HACMP cluster.

single point of failure

A single point of failure exists when a critical cluster function is provided by a single component. If that component fails, the cluster has no alternative way to provide that function and essential services become unavailable.

site

A data center that is running software that will be mirrored over some distance to another node or *ESS* that is defined as belonging to a second site. When nodes are defined in two sites the site becomes a topology component that is known to the HACMP software. Requires additional software such as GLVM, HAGEO or PPRC, or customization to use sites.

site policy

Inter-site management policy for startup, failover, and fallback for *replicated resource groups*. Options are:

Ignore (default): Use unless sites are defined.

Prefer Primary Site: The *primary instance* of the resource group is brought online on a node at this site and will fallback to this site. The *secondary instance* is maintained on the other site.

Online on Either Site: The *primary instance* of the resource group is brought online on the first node in its nodelist, no matter which site the node is on. It will not fallback to the other site. The *secondary instance* is maintained on the other site.

Online on Both Sites: The *primary instance* of the resource group is brought online on all nodes at both sites. The resource group node policy must be Online on All Available Nodes. This configuration has no failover or fallback.

Smart Assist

A utility for assisting the user in making an application highly available using HACMP. It incorporates knowledge of the application within itself. Depending on the type of the Smart Assist, it can either discover information about a running application to configure HACMP, or assist you in the configuration of that application.

Smart Assist for DB2

HACMP Smart Assist for DB2 extends an existing HACMP configuration to include monitoring and recovery support for DB2 Universal Database (UDB) Enterprise Server Edition.

This software lets you automatically configure HACMP in an environment where DB2 is already configured. This makes non-partitioned DB2 instances highly available.

Smart Assist for Oracle

Provides assistance for defining Oracle Application Server components to HACMP in AIX. HACMP provides a SMIT interface and includes support for configuring Oracle Application Server for Cold Failover Cluster, Oracle Application Server for Active Failover Cluster (AFC) and Oracle Database Instances (CFC only). In addition, HACMP Smart Assist for Oracle configures HACMP to monitor the Oracle processes and ensure that the applications remain available.

HACMP Smart Assist for Oracle helps you to configure Oracle components in one of multiple high availability configurations. You can choose the best high availability configuration for your environment based upon Oracle's recommendations, and then use the SMIT screens to implement the configuration in HACMP under the AIX environment.

This Smart Assist for Oracle also monitors Oracle AS and Oracle database instances and processes. In order to bring the Application Server and database under the control of HACMP, you must have pre-installed the respective Oracle components.

Smart Assist for WebSphere

An application that extends an existing HACMP configuration to include monitoring and recovery support for any of the following components in an IBM WebSphere Application Server Network Deployment (ND)

environment: WebSphere Application Server, WebSphere Deployment Manager, DB2 Universal Database Enterprise Server Edition, and Tivoli Directory Server. HACMP Smart Assist for WebSphere automates the HACMP configuration process for supported applications to minimize the time and effort required to integrate an application into an existing HACMP cluster.

See also *WebSphere*.

Smart Assist Framework

A set of tools, APIs and documentation that provides functions to *Smart Assist* developers and to Smart Assists themselves. See the *HACMP Smart Assist Developer's Guide*.

SMIT

System Management Interface Tool. An IBM utility. HACMP menus are accessible through SMIT.

SMS

The Short Message Service (SMS) provides the ability to send and receive text messages to and from mobile telephones. This text may contain words or numbers, or an alphanumeric combination.

Snapshot

See *cluster snapshot*.

SNMP

See *Simple Network Management Protocol*.

SOCC

See *Serial Optical Channel Converter*.

SP Switch (SP_switch)

A network used on the Cluster 1600 systems that run *PSSP*. See also *aliasing*.

SRC

See *System Resource Controller*.

staging configuration directory (SCD)

The directory in which the HACMP for AIX-specific ODM object classes are stored temporarily during a dynamic reconfiguration operation. See also *active configuration directory* and *default configuration directory*.

standby

An idle resource available to replace another equal resource currently in use. For example, a processor or a network interface.

startup

For a node, a process of a *joining or reintegrating* node acquiring a resource group that previously remained offline.

Also, for a resource group, you can specify a *startup policy*. The settling time and the distribution policy (if specified) affect the resource group startup.

See also *distribution policy*, *failover*, *fallback*, *resource group*, *resource group policies* and *settling time*.

sticky location

The term is no longer used.

sticky migration

The term is no longer used.

sticky resource group

No longer used. Instead, the term *priority override location* is used to describe a node to which a resource group has been moved.

See *priority override location*.

stopping HACMP cluster services

HACMP 5.4.1 introduces new terminology for the methods of stopping cluster services on a cluster node: bring resource groups offline, move resource groups and unmanage resource groups.

See:

- *bring resource groups offline: method of stopping cluster services*
- *move resource groups: method of stopping cluster services*
- *unmanage resource groups: method of stopping cluster services*.

submap

A particular view of some aspect of a network that displays symbols that represent objects. Some symbols may explode into other submaps, usually having a more detailed view than their parent submap. The application that creates a submap determines what part of the network the submap displays.

swapping hardware address

See *hardware address swapping*.

T

symbol

In HAView, a picture or icon that represents an object. Each symbol has an outside and inside component.

The outside component differentiates the object classes; the inside component differentiates the objects within the class.

synchronize cluster

A command that propagates the local node's definition of the cluster topology and resources to the rest of the cluster. A synchronization process sends the information contained on the current node to all defined cluster nodes. If a node attempts to join a cluster with a node environment that is out-of-sync with the active cluster nodes, it is denied. You must synchronize the node environment to the joining member by running the cluster verification and synchronization process in SMIT.

System Resource Controller (SRC)

Standard AIX facility used to start and stop the HACMP daemons.

T

takeover

See *fallover*.

takeover priority

A value assigned to each configured cluster resource on a per-node basis. In the event of a takeover, the active node with the highest priority will acquire the resource. A node with a higher priority is listed before the node with lower priority. See also *resource group*, *node priority policy* and *resource group policies*.

target mode SCSI

A *serial*, or non-IP network that may be used to connect pairs of nodes in a cluster over a SCSI bus.

target mode SSA

A *serial* network that may be used to connect nodes in an HACMP cluster using a Serial Storage Architecture (SSA) loop.

TCP/IP

Transmission Control Protocol/Internet Protocol. A communications subsystem that allows you to set up local area and wide area

networks. HACMP for AIX uses TCP/IP facilities to maintain communication among cluster members.

thread-safe

AIX includes support for threads in user programs. A thread is a path of execution: a program counter, registers, and a stack. You can have one or more threads per process, and they all share the same address space. On a multiprocessor machine these threads can be running simultaneously. HACMP for AIX supplies versions of the HACMP libraries for multithreaded applications to use.

A program is thread-safe when multiple threads in a process can be running that program successfully without data corruption. A library is thread-safe or *reentrant* when multiple threads can be running a routine in that library without data corruption.

tm SCSI

See *target mode SCSI*.

tm SSA

See *target mode SSA*.

Token-Ring

A type of network. Can be configured to run at either 4 megabits per second (Mbps), 16 Mbps, or 100 Mbps. Nodes and clients that communicate over a Token-Ring network require the Token-Ring High-Performance Network interface.

topology

Cluster layout and connections. Cluster, node, network, and network interface information.

topology services

An RSCT subsystem that generates heartbeats over multiple networks and provides information about network interface membership, node membership, and routing. Network interface and node membership provide indication of physical network interface and node failures respectively. Reliable Messaging uses the routing information to route messages between nodes around network interface card failures.

tracing

Diagnostic aid. HACMP tracing applies to the following daemons: **clstrmgr** and **clinfo**. Their execution is logged and reported by the system trace utility, **trcrpt**.

Two-Node Cluster Configuration Assistant

A wizard-like application that lets you configure a two-node cluster definition by providing the minimum information required by HACMP. The application is designed for users with little knowledge of HACMP who want to quickly set up a basic HACMP configuration.

U**unmanage a resource group: method of stopping cluster services**

In HACMP 5.4.1, one method of immediately stopping HACMP cluster services on a cluster node without stopping the applications or any other monitored resources.

If you use this method to stop the cluster services, the Cluster Manager will not react to resource failures nor run any event scripts for the resources associated with the application.

When you restart the cluster services for an “unmanaged” resource group with the option to automatically manage resources, HACMP will acquire the resources and start monitoring them. If you restart the cluster services with the option of manually managing resources, the group and its resources will remain unmanaged and not monitored by HACMP.

Note that you can always manually bring the resource group and its resources online on any node using **clRGmove**.

Placing a resource group in an “unmanaged” state is equivalent to a “forced shutdown” term for stopping HACMP cluster services in previous releases.

V**verifying**

Checking software modifications and cluster configuration, using the *cluster verification* utility.

Verifying software assures you that the HACMP-specific modifications to AIX system files exist and are correct.

Verifying the cluster topology and resource configuration assures you that all resources used by HACMP are validly configured and that ownership and takeover of those resources are defined and agreed upon by all nodes.

See also *cluster verification utility*.

version compatibility

Online update of software. HACMP 4.5 and up lets you install future versions of hardware or software without total loss of cluster availability during the replacement period.

Virtual SCSI (VSCSI)

Virtual SCSI is based on a client and server relationship. The Virtual I/O Server owns the physical resources and acts as server, or target, device. Physical adapters with attached disks on the Virtual I/O Server partition, may be shared by one or more partitions. These partitions contain a virtual SCSI client adapter that sees these virtual devices as standard SCSI compliant devices and LUNs.

See *VIO Server*

VIO (virtual I/O) Server

An appliance that provides virtual storage and shared Ethernet adapter capability to client logical partitions on POWER5 systems. It allows a physical adapter with attached disks on the VIO Server partition to be shared by one or more partitions, enabling clients to consolidate and potentially minimize the number of physical adapters required.

VLAN

VLAN is a technology used for establishing virtual network segments on top of physical switch devices.

volume group

An AIX LVM facility logical storage entity. Allows file systems to be spread across multiple disks and dynamically expanded during run time. A volume group can represent one or more physical disks. Each group contains a table of information about the disks it contains (for example, the file system, physical location, contents of the disk). If a

W

node detaches from the cluster, the surviving nodes can provide exactly the same view of disk services.

VPATH

See data path device.

W

WebSMIT

A web-enabled user interface that provides consolidated access to the HACMP SMIT configuration and management functions, cluster navigation, interactive cluster status information. It shows graphical representation of sites, nodes and networks, resource groups, and resource group dependencies. It also has links to online HACMP user documentation.

WebSphere

The IBM's infrastructure software platform for e-business on demand[™]. It lets you develop, deploy and integrate e-business applications, including Web publishing and enterprise-scale transaction processing. You can use it to create Web content, extend applications to incorporate mobile devices, and build electronic e-marketplaces.

Workload Manager (WLM)

A system administration tool provided with AIX. Workload Manager classes and their associated rules can be configured as resources in HACMP resource groups, so that the starting, stopping, and reconfiguration of WLM is under HACMP control.

Notices for Master Glossary

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept. LRAS / Bldg. 003
11400 Burnet Road
Austin, TX 78758-3493
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.