Title: Successful Business Continuity – Part 3 of 5

This is the third in a series of articles discussing how to implement AIX in an environment dedicated to business continuity. The topic of this article is the assignment of volume groups, physical volumes, logical volumes, JFS log logical volumes, file system mount points, resource groups, and major numbers. The name assignments of each of these entities should be unique enterprise wide and be conducive to high availability and disaster recovery. Much of the content of this article assumes that a SAN environment is available to the AIX systems. In a SAN environment, enterprise wide unique names allow the AIX system administrator to move a resource group between AIX systems, between clusters, or even between data centers without conflict. The point of ensuring that all resources are addressed by enterprise wide unique values is to eliminate the conflicts that must be resolved during a high availability fail-over or a disaster recovery implementation.

Definition: Enterprise wide unique - refers to a parameter that has one distinct value across any or all platforms throughout the entire enterprise.

This series of articles defines many IT areas that require enterprise wide policies, guidelines, standards, and procedures be defined, and offers recommended solutions for those defined areas. The areas discussed includes:

- Part 1:
  - User Names and UID Numbers
  - Group Names and GID Numbers

- Part 2:
  - Machine names
  - Hostnames
  - Boot adapter and service names
  - Resource group names
  - Aliases

- Part 3:
  - Volume Groups
  - Major Numbers
  - Logical Volumes
  - JFS Log Logical volume names
  - Mount points

- Part 4:
    - MQ Series Queue names and aliases
    - Resource Group start/stop scripts
    - Error logging
    - Error Notification

- Part 5:
    - Automated Documentation
    - Console Access
    - Job Scheduling
    - Project Planning

The purpose of this series of articles is to provide a foundation for business continuity. In support of that purpose, each topic discussed in this article is divided into the following:

- Policies
- Guidelines
- Standards
- Procedures

Each organization should define their own set of policies, guidelines, standards and procedures that define their enterprise wide rules of design and implementation. These rules ensure the ability of an organization to operate on a day-to-day basis as well as in a disaster recovery effort.

On an AIX system the Logical Volume Manager (LVM) provides the system administrator with the ability to allocate storage space for various purposes. The default volume group created when the AIX operating system is installed is called "rootvg". In order to facilitate high availability fail-overs and disaster recovery, the "rootvg" volume group is used by many organizations only for the storage of operating system specific files, programs, etc. All other volume groups are created on a SAN environment for the storage of applications and data.

Many organizations configure their AIX systems with two internal hard drives, both drives allocated for storage of the operating system in the "rootvg" volume group, the second drive being a mirror of the first.

An extension of this technique for creation of the boot device(s) is to install AIX on one internal hard drive, mirror the boot disk to a second internal hard drive, and

create a second mirror on a SAN disk. Then change the boot list to boot from the SAN disk first, the internal disks second, and third.

This (recommended) technique requires the system administrator to install the operating system onto a single internal disk, mirror the boot disk onto a second internal disk and onto a SAN disk. Example commands are shown where "hdisk0" is the boot disk on which AIX is initially installed, "hdisk1" is the second internal disk, and "hdisk2" is the SAN disk.

```
# mirrorvg rootvg hdisk1
# mirrorvg rootvg hdisk2
# bosboot –ad /dev/hdisk1
# bosboot –ad /dev/hdisk2
# bootlist –m –o hdisk2 hdisk0 hdisk1
```

The result of this technique is a primary boot disk on the SAN with two mirrors of the boot disk stored internally. One of the benefits of this technique is the "rootvg" can be booted from any AIX system that can access the SAN environment. In the event of a machine failure, it's "rootvg" volume group can simply be booted from another AIX system.

When upgrading AIX to a newer OS level, one of the "rootvg" mirrored disks can be used to accept the upgrade. This technique, called upgrading an "alternate clone", is outside the scope of this discussion but is well documented on the IBM web site. Utilizing this technique allows the system administrator to perform a system upgrade to an alternate "boot" disk while the system is up and running. Once the upgrade is complete the system is rebooted from the alternate "boot" disk, thus reducing the downtime to the time it takes to perform a single reboot. An added benefit of this technique is if the upgrade results in an unstable system, it can be rebooted from the original, unchanged "boot" disk.

In support of a business continuity environment, all application programs, files, and data should be located in non-rootvg volume groups, on disks in the SAN environment . Maintenance, backups, fail-overs, system recovery, and disaster recovery procedures will be made easier, predictable, repeatable by adhering to this rule.

In a previous article of this series, the concept of "**Resource Group**" was used to define any logical collection of resources, which may include disk, I/O, users, applications, etc. **A resource group should be viewed as independent from any machine, server, or data center.** In this context, the resource group name is used as

the basis of all other naming structures for all entities whether or not they are controlled by HACMP. The resource group name shall be an enterprise wide unique value in order to eliminate conflicts during manual, automated, or disaster recovery fail-overs.

Resources, such as disks, volume groups, NFS mounts, etc., needed to support applications should be allocated and provided with a "**resource group**" name. These names will be used as the basis for volume group names. The previous article in this series discussed the creation of resource group names and described a simple technique for defining the resource group names. A more complex technique will be described here for the purpose of integrating the "resource group" name into the naming structure of volume groups, logical volumes, directory mount points, and other entities to be discussed later.

The resource group naming structure described here is an actual implementation in a service provider's data center environment, where AIX systems for multiple companies may be simultaneously housed on a single piece of hardware separated by LPAR's.

The supporting principles of business continuity provide the basis for defining the standardized naming structures mentioned thus far. Those principles are:

- Policies – *those things that shall be done*
- Guidelines – *those things that should be done or can be expected*
- Standards – *the structure, based on policies and guidelines*
- Procedures – *how to implement the structure*

The following policies, guidelines, standards, and procedures describe the rules imposed in an actual data center for building, supporting and maintaining a large multi-company, multi-data center, AIX environment.

Policies:

All system resources shall be divided into logical resource groups, each resource group shall have an enterprise wide unique name. The resource group name shall be the basis of numerous naming structures and policies.

Resource groups shall be defined for standalone and high availability systems. All AIX systems will be designed as though they participate in a high availability environment, regardless of whether or not they actually do. This does not mean that all AIX systems will have HACMP or other automated high availability software installed.

A centralized repository (database) shall contain a list of all configured resource groups enterprise wide. New resource groups shall be entered into this repository.

The boot disk for each AIX system should be located on the SAN environment. Redundant mirrors may be located internally in each machine.

All data and applications will be stored separately from the "rootvg" volume group. Additional volume groups will be configured for storage of data and applications as necessary.

Each volume group shall have an enterprise wide unique name for the purpose of eliminating naming conflicts during manual, HACMP, or Disaster Recovery fail-overs.

Each volume group volume shall have a cluster wide unique major number for the purpose of eliminating numbering conflicts during manual, HACMP, or Disaster Recovery fail-overs. If possible, make the major number unique enterprise wide.

Each logical volume shall have an enterprise wide unique name for the purpose of eliminating naming conflicts during manual, HACMP, or Disaster Recovery fail-overs.

Each volume group that contains file systems will require at least one JFS log. Each JFS Log logical volume shall have an enterprise wide unique name for the purpose of eliminating naming conflicts during manual, HACMP, or Disaster Recovery fail-overs.

Each filefile system shall have an enterprise wide unique mount point directory for the purpose of eliminating directory conflicts during manual, HACMP, or Disaster Recovery fail-overs.

Guidelines:

Each AIX system will host one or more resource groups.

There should be a unique DNS entry based on each resource group name. These names shall be aliases pointing to IP names that are assigned to IP addresses.

Application volume group, logical volume, JFS log logical volumes, and file system mount point names should be enterprise wide unique values based on the resource group name.

Where possible create volume groups with an enterprise wide unique major number. At a bare minimum, the major number should be cluster wide unique.

More than one JFS log may be configured for a volume group if necessary, dependent upon file system load.

In order to ensure successful installation of an application into unique mount points, the application team must work closely with the system engineer.

Standards: Resource Group Name

A single standard shall be used for standalone, High Availability, and Disaster Recovery environments. This will eliminate naming conflicts in the event of a manual or automated fail-over, or if multiple instances of an application are running on a single AIX system.

The concept of **Resource Group** (RG) is used here in a larger scope than it is used in HACMP. In an enterprise environment, a resource group is any logical collection of resources, this may include disk, I/O, users, applications, etc. **A resource group should be viewed as being independent from any AIX system or data center.** The resource group name is used as the basis of all other naming structures for all entities whether or not they are controlled by HACMP. The resource group name shall be an enterprise wide unique value in order to eliminate conflicts during manual, automated, or disaster recovery fail-overs.

When designing any new system, the first step is to determine the resource group name(s). The names of volume groups, logical volumes, mount points, major numbers, WLM classes, etc, are all derived from the resource group name(s).

The resource group name shall consist of exactly 8 characters with the following structure:

ApplicationCode + Environment + Function + Custom + Sequence ID
3 char + 1 char + 1 char + 2 char + 1 char

The detailed information for each component of the resource group name is described below:

RG Name Component	Number of Characters	Values
Application Code	3	atl = Atlas db2 = DB2 nim = NIM ora = Oracle peo = PeopleSoft sap = SAP tps = Maximo vio = Virtual I/O
Environment	1	a = acceptance g = pre-production/Gold d = test/development p = production t = test x = disaster recovery
Function	1	a = application c = combination/multi-purpose d = database m = management u = utility
Company or other identifier	2	ac = Acme mx = Mt Xia ib = IBM
Sequence ID	1	0-9,A-Z,a-z

A single AIX system may contain multiple resource groups, and typically there will be one volume group defined per resource group. However, a resource group may contain several volume groups depending upon the requirements of the application.

Standards: Volume Group Name

To assign enterprise wide unique volume group (VG) names, the system administrator must first define the resource group names. Once the resource group names have been defined, then a VG name may be defined based on the resource group name.

In order to facilitate normal maintenance, disaster recovery and business continuity, it is recommended that each enterprise wide unique volume group name be assigned enterprise wide unique major number.

A single AIX system may contain multiple resource groups, and typically there will be one VG defined per resource group. However, a resource group may contain several VG's, depending upon the requirements of the application.

To define a VG name, obtain the 8 character resource group name. then add a 2 digit volume group sequence number that will uniquely identify the VG, followed by the characters "vg". The VG name will always end with the characters "vg".

The VG name shall consist of exactly 12 characters with the following structure:

```
ApplicationCode + Environment + Function + Company + Sequence ID + VG Sequence ID + "vg"
    3 char  + 1 char   + 1 char + 2 char + 1 char  +   2 char    + 2 char
```

As an example, a resource group named "db2apmx0", may have multiple associated VG's:

VG Name Component	VG Sequence Identifier	LV Identifier	VG Name
db2apmx0	00	vg	db2apmx000vg
db2apmx0	01	vg	db2apmx001vg
db2apmx0	02	vg	db2apmx002vg

Each VG also requires a server or cluster wide unique Major Number. A unique major number can be generated using the following algorithm:

```
VGNAME=”db2apmx000vg”
MajorNbr=$( print "${VGNAME}" | sum -o | awk '{ print $1 }' )
print ${MajorNbr}
```

To reiterate, before creating a VG, first establish an enterprise wide unique resource group name, a VG name, and a major number. Then create the VG.

Standards: Logical Volume Name

To assign enterprise wide unique logical volume (LV) names, the system administrator must first define the resource group names. Once the resource group names have been defined, then a VG must be defined based on the RG name. After the VG has been created, LV's can be assigned. A VG will typically contain several LV's, and each LV will be named based on the resource group to which it is associated.

To define a LV name, obtain the 8 character resource group name, then add a 4 digit logical volume sequence identifier that will uniquely identify the LV, followed by the characters "lv". The 4 digit LV sequence identifier will consist of alpha-numeric characters and must always be exactly 4 characters in length. The LV name will always end with the characters "lv".

The LV name shall consist of exactly 14 characters with the following structure:

```
ApplicationCode + Environment + Function + Company + Sequence ID + LV Sequence ID + "lv"
   3 char    + 1 char   + 1 char+ 2 char+ 1 char    +   4 char    + 2 char
```

As an example, a resource group named "db2apmx0", may have a volume group named "db2apmx00vg". This volume group may have multiple LV's associated with it:

VG Name Component	LV Sequence Identifier	LV Identifier	LV Name
db2apmx0	db20	lv	db2apmx0db20lv
db2apmx0	db21	lv	db2apmx0db21lv
db2apmx0	db22	lv	db2apmx0db22lv

JFS file systems will require a logical volume for the JFS log. This must also have an enterprise wide unique name.

Standards: JFS Log Logical Volume Name

To assign enterprise wide unique JFS Log LV names, the system administrator must first define the resource group names. Once the resource group names have been defined, then a VG must be defined based on the RG name. After the VG has been created, JFS Log LV's can be assigned. A VG will typically contain one JFS Log LV's, however multiple JFS Log LV's can exist.

To define a JFS Log LV name, obtain the 8 character resource group name, then add the 4 digit logical volume sequence identifier that will uniquely identify the JFS Log LV, followed by the characters "lv". The 4 digit JFS Log LV sequence identifier will consist of the characters "jfs" followed by a single digit to uniquely identify the JFS Log LV. The JFS Log LV name will always end with the characters "lv".

The JFS Log LV name shall consist of exactly 14 characters with the following structure:

```
ApplicationCode + Environment + Function + Company + Sequence ID +    "jfs"   + JFS Log Sequence ID + "lv"
   3 char    +   1 char + 1 char+ 2 char+ 1 char  + 3 char+    1 char       + 2 char
```

As an example, a resource group named "db2apmx0", may have a volume group named "db2apmx00vg". This volume group may have multiple JFS Log LV's associated with it:

JFS Name Component	JFS Log LV Sequence ID	JFS Log LV ID	JFS Log LV Name
db2apmx0	jfs0	lv	db2apmx0jfs0lv
db2apmx0	jfs1	lv	db2apmx0jfs1lv
db2apmx0	jfs2	lv	db2apmx0jfs2lv

JFS file systems will require a logical volume for the JFS log. This must also have an enterprise wide unique name.

Standards: file system Directory Mount Points

To assign enterprise wide unique LV names, the system administrator must first define the Resource Groups, Volume Groups and Logical Volumes. Once these have been defined, the file system mount point directory names can be assigned. Typically a file system mount point is required for each logical volume, therefore the mount point can usually be based on the logical volume name, or at a minimum the resource group name.

To define a file system mount point directory name, obtain the 8 character resource group name, then depending upon the applications file system requirements, use the RG name as the mount point, or add sub-directories to make it enterprise wide unique.

The file system mount point directory name shall consist of at least 8 characters, but may be of a variable length:

/ + ApplicationCode + Environment + Function + Company + Sequence ID + (LV Sequence ID or Directory Structure)

```
   3 char     + 1 char  + 1 char+ 2 char+ 1 char  + 4 or more characters
```

As an example, a resource group named "db2apmx0", may have multiple file systems associated with it:

RG Name Component	Optional Logical Volume Sequence ID	Optional Sub-Directories	file system Mount Point
db2apmx0		db2_08_01/bin	/db2apmx0/db2_08_01/bin
db2apmx0		db2_08_01/data	/db2apmx0/db2_08_01/data
db2apmx1	mq01		/db2apmx1mq01
db2apmx1	mq02		/db2apmx1mq02
db2apmx1	mq03		/db2apmx1mq03

Procedures:

The procedures associated with implementing the standards described here require strict attention to detail when creating volume groups, logical volumes and file systems. The system administrator must plan the implementation and define a "resource group" to associate with the naming structures. For example, the steps in defining a new file system are as follows:

1. Define an enterprise wide unique resource group name.
2. Define an enterprise wide unique volume group name.
3. Define a cluster wide unique (enterprise wide unique if possible) major number for the new volume group.
4. Define an enterprise wide unique JFS log logical volume name.
5. Define an enterprise wide unique logical volume name.
6. Define an enterprise wide unique mount point directory name.
7. Create the volume group using the enterprise wide unique name and major number.
8. Create the JFS log logical volume using the enterprise wide unique name.
9. Create the logical volume using the enterprise wide unique name.
10. Create the file system on the logical volume using the enterprise wide unique mount point directory and JFS log.

System design and implementation in a business continuity environment should be regarded as a “non-standard” design. Normally a “standard” implementation does not account for the ability to perform fail-overs from any AIX system to any other AIX system, as is described here. When installing applications in this “business continuity” environment recognize that the application vendor will almost always provide “standard” installation instructions, not “business continuity” installation instructions. For the system and application administrators, there is a learning curve associated with moving from a “standard” environment to a “business continuity” environment because the “standard” installation instructions for applications must usually be modified to fit this enhanced environment.

The procedures for each application to be installed in the “business continuity” environment will be specific to each environment and will likely be different from the documentation provided by the application vendor. It will be the job of the system and application administrators to create a procedure to fit the environment.

The next article in this series will discuss enterprise level techniques for defining resource group start/stop scripts, MQ Series queue names and aliases, error logging, and error notification.