Redbooks Paper

David Toussaint

Installation and Configuration of Tivoli System Automation for High Availability of DB2 UDB BCU on AIX

This IBM® Redpaper documents an installation and configuration of IBM Tivoli® System Automation on IBM AIX® to provide high availability for a DB2® Universal Database™ (UDB) Balanced Configuration Unit (BCU). It provides the information that is needed to enable high availability quickly for an AIX DB2 environment. It is of special interest to those who intend to provide high availability for a DB2 database BCU that is running on the AIX platform.

Introduction

This is a quick-start guide to IBM Tivoli System Automation. Using this guide, you can begin an AIX Tivoli System Automation implementation with a minimal amount of research. The objective is to summarize the various commands, queries, and procedures that are useful in the processes of configuring and troubleshooting a high availability AIX BCU environment.

Tivoli System Automation manages the availability of application resources. Tivoli System Automation policies provide scripts to detect system components, create the appropriate Reliable Scalable Cluster Technology (RSCT)/ Resource Monitoring and Control (RMC) resources, and define relationships or dependencies between the components and resources. Tivoli System Automation provides policy-based automation to maintain availability by:

- Reducing the frequency and duration of incidents that impact IT availability
- Establishing a single point of high availability within the IT organization
- Addressing shortcomings of high availability with best practices
- Easing management of complex IT infrastructures
- Reducing costs
- Moving the IT organization from reactive error correction to proactive service enhancement

Preparation

Our DB2 home directory was mounted using the network file system (NFS). To make it highly available, we used Tivoli System Automation to automate the failover of the IP address, NFS server process, exported data, and the AIX control files as follows:

- A specific service IP was created for the NFS service.
- A specific service IP was created for the DB2 administration interface, data partition0.
- The volume groups (disks) that contain the NFS and DB2 data must share the same major number on both potential hosts. A separate logical volume and file system, /nfsctrl, was created within the “shared” NFS data volume group, to store our configuration files.
- Links to /nfsctrl/xtab and /nfsctrl/rmtab were used to replace /etc/xtab and /etc/rmtab.

Failover testing

Before starting work with Tivoli System Automation, we tested hosting both NFS server and the DB2 resources on their failover partners. In other words, we did a manual failover. The manual failover required that we import and test all the required volume groups, the mounting of file systems, and the starting of NFS and DB2 services to expose any problems with the configuration before Tivoli System Automation was introduced. After the manual failover was successful, we began to automate the process with Tivoli System Automation.

Downloading, installing, and updating

Tivoli System Automation builds on RSCT. Tivoli System Automation 2.1.1.1 requires updates to RSCT level 2.4.5.2. You can accomplish this by:

- Downloading the latest base components from Xtreme Leverage Downloads portal:

  `http://w3-103.ibm.com/software/xl/portal`

- Installing the Tivoli System Automation for Mulitplatforms V2.1.1 base component as follows:

  ```
  SAM2110- from C891KML.tar
  ./installSAM
  ```

- Downloading Fixpack 3, IBM Tivoli System Automation for Multiplatforms, 2.1.0-TIV-ITSAMP-FP0003 from:

  http://www-1.ibm.com/support/docview.wss?uid=swg24012715&rs=0&cs=utf-8&context=SSRM2X&dc=D400&loc=en_US&lang=en&cc=US

- Installing Fixpack 3 as follows:

  ```
  SAM2111- from 2.1.0-TIV-SABASE-AIX-FP0003.tar
  ./installSAM
  ```

- Downloading policies. Policies are the prepackaged scripts that you can use to create resources and relationships. They also contain the scripts that start, stop, and monitor the created resources.

  You can download and install sam.policies.aix.1.2.2.1 from:

  ftp://ftp.software.ibm.com/software/tivoli/products/sys-auto-linux/

All of these packages should be installed on each of the nodes in the HA cluster.

Preparing the nodes for Tivoli System Automation domain

After installing the Tivoli System Automation code and policies, you must prepare your environment and create the peer domain and nodes as follows:

1. Tivoli System Automation requires that you export the CT_MANAGEMENT_SCOPE environment variable. Add the environment variable to /.profile of each node in your HA cluster as follows:

   ```
   CT_MANAGEMENT_SCOPE=2
   export ENV DSHPATH PRINTER MAIL CT_MANAGEMENT_SCOPE
   ```

2. Run the following command on each of the domain nodes to initialize the communication between the nodes:

   ```
   r56n03,04,06,08,>  preprpnode r56g03 r56g04 r56g06 r56g08
   ```

3. Create the peer domain by running this command on one node *only*:

   ```
   r56n06>   mkrpdomain db2ha r56g03 r56g04 r56g06 r56g08
   ```

4. Start the peer domain by running this command on one node *only*::

   ```
   r56n06> startrpdomain db2ha
   ```

Useful query commands include:

lsrpnodes	Nodes may take a moment to reach the online state
lsrpdomain	Domain may take a moment to reach the online state
lssamctrl	SAMControl status
samctrl	M T (manual = True) Turn off automation (Tivoli System Automation)
samctrl	M F (manual = False) Turn on automation (Tivoli System Automation)

Tivoli System Automation configuration resources

This section demonstrates some tips and techniques that we used during this installation.

Tie breaker

A domain with an even number of nodes requires a tie breaker resource. There are several types. Complete the procedure in Example 1 to configure the network tie breaker.

Example 1 Configuring a tie breaker

```
CONFIG
r56n06 and n08 :/bm1 # vi /usr/sbin/cluster/netmon.cf
29.40.9.254
29.40.98.254

r56n06: mkrsrc IBM.TieBreaker Type="EXEC" Name="NetTieBkr"
DeviceInfo='PATHNAME=/usr/sbin/rsct/bin/samtb_net Address=129.40.98.254 Log=1
Count=2' PostReserveWaitTime=30;
r56n06:/bm1 # chrsrc -c IBM.PeerNode OpQuorumTieBreaker="NetTieBkr"
```

Volume groups

To avoid logging errors, we prevented the HA volume groups from automatically starting by running:

```
chvg -a'n' -Q'n' vg_p##
```

NFS resources

The NFS scripts are found in /usr/sbin/rsct/sapolicies/nfsserver. You can access them as follows:

1. Find and populate the sa-nfsserver.conf file with the parameters of your environment. Example 2 shows a modified configuration file.

Example 2 Modified config file

```
# --directory for control scripts
script_dir="/usr/sbin/rsct/sapolicies/nfsserver"

# --prefix of all NFS server resources
prefix="SA-nfsserver-"

# --list of nodes in the NFS server cluster
nodes="r56n06 r56n08"

# --IP address and netmask for NFS server,
#   If more instances of <ip_>, add more rows, like: ip_2 ip_3...
ip_1="129.40.9.12,255.255.255.0"
# --List of network interfaces ServiceIP ip_x depends on.
#   Entries are lists of the form <network-interface-name>:<node-name>,...
#   If more instances of <nieq_>, add more rows, like: nieq_2 nieq_3...
#   Default: create 'static equivalencies'
#   to create 'dynamic equivalencies' use keyword nieq_1_dyn ...
nieq_1="en0:r56n06,en0:r56n08"
```

```
# --common local mountpoint for shared data
#   If more instances of <data_>, add more rows, like: data_tmp, data_proj...
#   Note: the keywords need to be unique!
data_varlibnfs="/nfsctrl"
data_work="/shared_db2home"

# --LVM definitions: VG and optional hdisk  (for AIX only)
#   one entry allowed, like: myvg  ...  with hdisk like: myvg hdisk5
lvm="vgdb2nhome hdisk10"
```

2. Run `cfgnfsserver` to preview the commands used to create the resources.
3. If you are satisfied with the preview, you can create the resources and relationships. To do this, run:

   `cfgnfserver -p (perform)`

   The resources and resource group should resemble those in Example 3. To verify that the resource group is online, ensure that the NominalState is `Online`, as shown.

Example 3 NominalState is online

```
#lsrg -g SA-nfsserver-rg

Displaying Resource Group information:
For Resource Group "SA-nfsserver-rg".

Resource Group 1:
        Name                = SA-nfsserver-rg
        MemberLocation      = Collocated
        Priority            = 0
        AllowedNode         = ALL
        NominalState        = Online
        ExcludedList        = {}
        Subscription        = {}
        Owner               =
        Description         =
        InfoLink            =
        ActivePeerDomain    = db2ha
        OpState             = Online
        TopGroup            = SA-nfsserver-rg
        ConfigValidity      =
        TopGroupNominalState = Online
```

To view the individual resources (Example 4), use:

`/usr/sbin/rsct/sapolicies/bin # getstatus`

Example 4 Individual resources

```
/usr/sbin/rsct/sapolicies/bin # getstatus

-- Resource Groups and Resources --

                            Group Name           Resources
                            ----------           ---------
                       SA-nfsserver-rg SA-nfsserver-data-varlibnfs
                       SA-nfsserver-rg SA-nfsserver-data-work
```

```
                        SA-nfsserver-rg    SA-nfsserver-ip-1
                        SA-nfsserver-rg     SA-nfsserver-lvm
                        SA-nfsserver-rg  SA-nfsserver-server
                                      -                    -
-- Resources --

                           Resource Name            Node Name          State
                           -------------            ---------          -----
               SA-nfsserver-data-varlibnfs               r56n05        Offline
               SA-nfsserver-data-varlibnfs               r56n07         Online
                                      -                    -              -
                    SA-nfsserver-data-work               r56n05        Offline
                    SA-nfsserver-data-work               r56n07         Online
                                      -                    -              -
                         SA-nfsserver-ip-1               r56n05        Offline
                         SA-nfsserver-ip-1               r56n07         Online
                                      -                    -              -
                          SA-nfsserver-lvm               r56n05        Offline
                          SA-nfsserver-lvm               r56n07         Online
                                      -                    -              -
                     SA-nfsserver-server                 r56n05        Offline
                     SA-nfsserver-server                 r56n07         Online
```

4. Example 4 on page 5 shows the result of the getstatus query for a healthy cluster. Further test the function and availability of the remote NFS by listing the files at the NFS clients. If a problem exists with the status or function of the NFS resources, you must resolve it before attempting the failover. If no problems exist, you can attempt the failover using:

   ```
   rgreq -o Move -n r56n07 SA-nfsserver-rg
   ```

   This moves the NFS resource group from r56n07 (current host) to its failover partner, r56n05.

To recover from a failed failover attempt, take all Tivoli System Automation resources offline. The getstatus query results must reflect `Offline` for *all* resources before you attempt to bring them back online. Follow these steps:

1. Take all NFS resources offline:

   ```
   chrg -o Offline -s 'Name like "SA-nfsserver-%"'
   ```

2. If the *state* of any resources, per the getstatus command, is `Unknown`, a resource reset is needed. To reset NFS resources, use:

   ```
   resetrsrc -s "Name like 'SA-nfsserver-%'" IBM.Application (if necessary)
   resetrsrc -s "Name like 'SA-nfsserver-%'" IBM.ServiceIP (if necessary)
   ```

3. After all the resources are offline, bring all NFS resources online, using:

   ```
   chrg -o Online -s 'Name like "SA-nfsserver-%"'
   ```

DB2 resources

The DB2 scripts (regdb2salin) must be modified significantly for AIX. These changes should be incorporated into the next release. You can access DB2 resources as follows:

1. Find the tarball directory, including working modified scripts, called db2.sam.scripts.tar.Z. Extract script files from the tarball.

2. Modify and run the regdb2aixbcu command or use the following command syntax, which creates the DB2 “rs” resources:

```
regdb2salin -a InstanceName -n NetMask -i Admin FCM address
PrimaryNode,StandbyNode:PrimaryNode,StandbyNode

/opt/IBM/db2/V9.1/ha/tsa/regdb2salin -a db2inst2 -n 255.255.255.0 -i
129.40.98.10 -l
r56g06,r56g08:r56g03,r56g04:r56g03,r56g04:r56g03,r56g04:r56g03,r56g04:r56g03,r5
6g04:r56g03,r56g04:r56g03,r56g04:r56g03,r56g04:r56g04,r56g03:r56g04,r56g03:r56g
04,r56g03:r56g04,r56g03:r56g04,r56g03:r56g04,r56g03:r56g04,r56g03:r56g04,r56g03
```

3. Run mkrsrc_containers. This creates the logical volume manager (LVM) and mount resources.

If the expected resources do not appear in the resource list, it is possible that you need to create Admin. Run mkrsrc_cont_admin to create Admin. You may also rerun mkrsrc_containers if the expected resources do not appear in the resource list.

Figure 1 shows the seven resources created for Balanced Processing Unit 1 (BPU1). Each data partition has these seven resources. You should create resources like those in Figure 1 for each data partition. To view the resources, use:

```
/usr/sbin/rsct/sapolicies/bin/getstatus
```

```
Group Name              Resources
db2_db2inst2_1-rg       b2_db2inst2_1-rs_mount_4      partition fs mounts
db2_db2inst2_1-rg       db2_db2inst2_1-rs_mount_3     partition fs mounts
db2_db2inst2_1-rg       db2_db2inst2_1-rs_mount_2     partition fs mounts
db2_db2inst2_1-rg       db2_db2inst2_1-rs_mount_1     partition fs mounts
db2_db2inst2_1-rg       db2_db2inst2_1-rs_mount_0     partition fs mounts
db2_db2inst2_1-rg       db2_db2inst2_1-lvm            volume group
db2_db2inst2_1-rg       db2_db2inst2_1-rs             db2 partition
```

Figure 1 The seven resources created for BPU 1

The resources in Figure 2 must be created for each BPU 0, the Admin node.

```
Group Name              Resources
db2_db2inst2_0-rg       db2_db2inst2_0-rs_mount_1
db2_db2inst2_0-rg       db2_db2inst2_0-rs_mount_0
db2_db2inst2_0-rg       db2_db2inst2_0-lvm
db2_db2inst2_0-rg       db2_db2inst2_0-rs_ip            floating ip address
db2_db2inst2_0-rg       db2_db2inst2_0-rs
```

Figure 2 Resources to be created for each BPU 0

Example 5 shows a healthy cluster with both NFS and DB2 BCU resources.

Example 5 A healthy cluster

```
r56n01:/ # /usr/sbin/rsct/sapolicies/bin/getstatus
-- Resource Groups and Resources --
                            Group Name           Resources
                           ----------           ---------
                      SA-nfsserver-rg SA-nfsserver-data-varlibnfs
                      SA-nfsserver-rg SA-nfsserver-data-work
                      SA-nfsserver-rg    SA-nfsserver-ip-1
                      SA-nfsserver-rg     SA-nfsserver-lvm
```

```
                    SA-nfsserver-rg  SA-nfsserver-server
                                  -                    -
                   db2_db2inst5_0-rg    db2_db2inst5_0-rs
                   db2_db2inst5_0-rg db2_db2inst5_0-rs_ip
                                  -                    -
                   db2_db2inst5_1-rg    db2_db2inst5_1-rs
                                  -                    -
                  db2_db2inst5_10-rg   db2_db2inst5_10-rs
                                  -                    -
                  db2_db2inst5_11-rg   db2_db2inst5_11-rs
                                  -                    -
                  db2_db2inst5_12-rg   db2_db2inst5_12-rs
                                  -                    -
                  db2_db2inst5_13-rg   db2_db2inst5_13-rs
                                  -                    -
                  db2_db2inst5_14-rg   db2_db2inst5_14-rs
                                  -                    -
                  db2_db2inst5_15-rg   db2_db2inst5_15-rs
                                  -                    -
                  db2_db2inst5_16-rg   db2_db2inst5_16-rs
                                  -                    -
                   db2_db2inst5_2-rg    db2_db2inst5_2-rs
                                  -                    -
                   db2_db2inst5_3-rg    db2_db2inst5_3-rs
                                  -                    -
                   db2_db2inst5_4-rg    db2_db2inst5_4-rs
                                  -                    -
                   db2_db2inst5_5-rg    db2_db2inst5_5-rs
                                  -                    -
                   db2_db2inst5_6-rg    db2_db2inst5_6-rs
                                  -                    -
                   db2_db2inst5_7-rg    db2_db2inst5_7-rs
                                  -                    -
                   db2_db2inst5_8-rg    db2_db2inst5_8-rs
                                  -                    -
                   db2_db2inst5_9-rg    db2_db2inst5_9-rs
                                  -                    -
-- Resources --
                        Resource Name           Node Name           State
                        -------------           ---------           -----
                SA-nfsserver-data-varlibnfs              r56n05          Online
                SA-nfsserver-data-varlibnfs              r56n07         Offline
                                  -                      -               -
                    SA-nfsserver-data-work               r56n05          Online
                    SA-nfsserver-data-work               r56n07         Offline
                                  -                      -               -
                      SA-nfsserver-ip-1                  r56n05          Online
                      SA-nfsserver-ip-1                  r56n07         Offline
                                  -                      -               -
                       SA-nfsserver-lvm                  r56n05          Online
                       SA-nfsserver-lvm                  r56n07         Offline
                                  -                      -               -
                     SA-nfsserver-server                 r56n05          Online
                     SA-nfsserver-server                 r56n07         Offline
                                  -                      -               -
                       db2_db2inst5_0-rs                 r56n05          Online
                       db2_db2inst5_0-rs                 r56n07         Offline
                                  -                      -               -
                   db2_db2inst5_0-rs_ip                  r56n05          Online
                   db2_db2inst5_0-rs_ip                  r56n07         Offline
```

```
                    -               -               -
 db2_db2inst5_1-rs                  r56n01          Online
 db2_db2inst5_1-rs                  r56n02          Offline
                    -               -               -
db2_db2inst5_10-rs                  r56n01          Offline
db2_db2inst5_10-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_11-rs                  r56n01          Offline
db2_db2inst5_11-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_12-rs                  r56n01          Offline
db2_db2inst5_12-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_13-rs                  r56n01          Offline
db2_db2inst5_13-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_14-rs                  r56n01          Offline
db2_db2inst5_14-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_15-rs                  r56n01          Offline
db2_db2inst5_15-rs                  r56n02          Online
                    -               -               -
db2_db2inst5_16-rs                  r56n01          Offline
db2_db2inst5_16-rs                  r56n02          Online
                    -               -               -
 db2_db2inst5_2-rs                  r56n01          Online
 db2_db2inst5_2-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_3-rs                  r56n01          Online
 db2_db2inst5_3-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_4-rs                  r56n01          Online
 db2_db2inst5_4-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_5-rs                  r56n01          Online
 db2_db2inst5_5-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_6-rs                  r56n01          Online
 db2_db2inst5_6-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_7-rs                  r56n01          Online
 db2_db2inst5_7-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_8-rs                  r56n01          Online
 db2_db2inst5_8-rs                  r56n02          Offline
                    -               -               -
 db2_db2inst5_9-rs                  r56n01          Offline
 db2_db2inst5_9-rs                  r56n02          Online
```

Troubleshooting

Sometimes the regdb2salin scripts do not run cleanly. That can explain why one of the resources for partition 0 is missing. If this happens, Example 6 shows how to create mount point resources, add them to a resource group, and define a dependency relationship of mounts to a volume group (LVM) resource.

Example 6

```
To manually add the missing db2fsp1 resource.
```

```
# mkrsrc IBM.Application \
    Name="db2_db2inst2_0-rs_mount_1" \
    ProtectionMode=1 \
    StartCommand="/opt/IBM/db2/V9.1/ha/tsa/mount_start.ksh /db2fs1p0" \
    StartCommandTimeout=300 \
    StopCommand="/opt/IBM/db2/V9.1/ha/tsa/mount_stop.ksh /db2fs1p0" \
    StopCommandTimeout=600 \
    MonitorCommand="/opt/IBM/db2/V9.1/ha/tsa/mount_monitor.ksh /db2fs1p0" \
    MonitorCommandPeriod=300 \
    MonitorCommandTimeout=290 \
    UserName=root \
    NodeNameList={'r56g06','r56g08'}

#   addrgmbr -g db2_db2inst2_0-rg IBM.Application:db2_db2inst2_0-rs_mount_1
#   mkrel -S IBM.Application:db2_db2inst2_0-rs_mount_1  -G
IBM.Application:db2_db2inst2_0-lvm -p DependsOn db2_db2inst2_0-rs_mount_1_on_lvm
```

When we were performing the Admin node failover testing, at first, the failover was not triggered by disabling the Fast Communication Manager (FCM) network interface. This was resolved by creating a dependency on that resource. The steps to follow for creating a dependency are:

1. Verify or define nic equivalency for ServiceIP:

   ```
   # mkequ db2_db2inst2_0_nieq IBM.NetworkInterface:en4:r56g06,en4:r56g08
   ```

2. Define the relationship for equiv and serviceIP:

   ```
   # mkrel -S IBM.ServiceIP:db2_db2inst2_0-rs_ip -G
   IBM.Equivalency:db2_db2inst2_0_nieq -p DependsOn db2_db2inst2_0_nieq
   ```

3. Define the dependency of volume groups on network equivalency:

   ```
   # mkrel -S IBM.Application:db2_db2inst2_0-lvm -G
   IBM.Equivalency:db2_db2inst2_0_nieq -p DependsOn db2_db2inst2_0_lvm_on_nieq
   ```

Error logging

By default, Tivoli System Automation logs errors using /usr/es/adm/cluster.log and errpt.

Failover testing

When testing the DB2 BCU2 HA configuration with Tivoli System Automation, we had to modify several of the parameters significantly to avoid exceeding the default timeout parameters and to minimize the CPU load. This was because of the relatively large numbers of resources. Before you modify the parameters, you must take all resources offline by running:

```
chrg -o  Offline -s 'Name like "db2_db2inst2_%"'
```

To modify each parameter, follow these steps:

1. For the health query interval of each resource, use:

   ```
   chrsrc -s 'Name like "db2_db2inst2%"' IBM.Application MonitorCommandPeriod=300
   ```

2. For the health query timeout, use:

   ```
   chrsrc -s 'Name like "db2_db2inst2%"' IBM.Application MonitorCommandTimeout=290
   ```

3. For the resource startup script timeout, use

```
chrsrc -s 'Name like "db2_db2inst2%"' IBM.Application StartCommandTimeout=300
```

4. For the Resource Stop script timeout, use:

```
chrsrc -s 'Name like "db2_db2inst2%"' IBM.Application StopCommandTimeout=720
```

When testing goes wrong, you are often left with resources in various states such as online, offline, and unknown. When the state of a resource is unknown, before attempting to restart it, you must issue:

```
resetrsrc
```

When you are restarting DB2, you must verify that all the resources are offline before attempting to bring them online again. You must also correct the db2nodes.cfg file. Use the getstatus command to see the state of each resource on all the nodes and to ensure that all indicate offline (Example 7).

Example 7

```
chrg -o Offline -s 'Name like "db2_db2inst2%"'
resetrsrc -s "Name like 'db2_db2inst2_%'" IBM.Application ( if necessary )
resetrsrc -s "Name like 'db2_db2inst2_%'" IBM.ServiceIP ( if necessary )
verify db2nodes.cfg for accuracy, modify to suit your test requirements.
/db2home/db2inst2/sqllib/db2nodes.cfg
 chrg -o Online -s 'Name like "db2_db2inst2%"'
```

Tip: When we were testing the NFS failover, we were able to move the server over successfully, but the existing NFS client mounts stopped functioning. We solved this problem by unmounting and remounting the NFS volume. We also discovered that the major number of the NFS volume group was not identical on the two host servers. After this was changed, the existing NFS mounts continued to function after the failover without intervention.

Tivoli System Automation and RSCT commands

The following Tivoli System Automation and RSCT commands are commonly used for basic queries, for checking component status, for listing the status of group members, and other helpful tasks:

- The basic query commands are:

```
lsrg
lsrel
lsequ
lsrsrc
```

- Use this command to check the status of components:

```
getstatus ( /usr/sbin/rsct/sapolicies/bin/getstatus  )
```

- Use this command to list group members and resource status:

```
lsrg -m
```

- Use this command to bring all resources online

```
chrg -o Online -s 'Name like "db2_db2inst2_%"'
```

- Use this command to take all resources offline

  ```
  chrg -o Offline -s 'Name like "db2_db2inst2_%"'
  ```

- Use this command for a soft failover of the db2_db2inst2_9-rg resource that is currently hosted by r56g04:

  ```
  rgreq -o Move -n r56g04 db2_db2inst2_9-rg
  ```

- Use these commands to reset resources (which is necessary when the resource state is unknown):

  ```
  resetrsrc -s "Name like 'db2_db2inst2_%'" IBM.Application
  resetrsrc -s "Name like 'db2_db2inst2_%'" IBM.ServiceIP
  ```

- Use this command to view resource details, including the path to start, stop, and monitor scripts, and timeouts:

  ```
  lsrsrc -s "Name like 'SA-nfsserver-data%'" IBM.Application
  ```

Summary

Tivoli System Automation is a powerful tool that you can use to provide high availability to critical resources. Tivoli System Automation provides an alternative to the traditional HACMP™ implementations of the past, allowing a more granular, resource-driven approach to high availability. We have discussed how to set up the environment, download the Tivoli software and fixpacks, configure the resources, issue commands, and test the failover capability.

The author of this redpaper

David Toussaint is a Systems Management Specialist in the BICoC and System p™ Commercial Benchmark Center of the IBM Systems and Technology Group in Poughkeepsie.

Thanks to Ken Delathouwer, IBM Toronto, for sharing his experiences and assisting with this implementation.

Resources

- IBM Web site:

  http://www-306.ibm.com/software/tivoli/products/sys-auto-multi/

- *IBM Tivoli Systems Automation for Multiplatforms.* Application Enablement of NFS File Server, Enrico Joedecke:

  ftp://ftp.software.ibm.com/software/tivoli/products/sys-auto-linux/

- *IBM Tivoli System Automation for MultiplatformsBase Component User's Guide,* SC33-8210
- *Balanced Configuration Unit for AIX with IBM System Storage DS4800: Design and Implementation Version 2.1,* SC10-4381

Notices

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 IBM 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.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS 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.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo)  ™
AIX®
DB2 Universal Database™
DB2®
HACMP™
IBM®
Redbooks™
System p™
System Storage™
Tivoli®

Other company, product, or service names may be trademarks or service marks of others.