

High Availability in an SAP Business Warehouse Multi-Partition DB2 Environment Using SA MP



Applies to:

SAP products based on SAP NetWeaver 7.0 SR3 or higher running on IBM DB2 for Linux, UNIX, and Windows databases. For more information, visit the [Landscape Design and Architecture homepage](#).

Summary

This document describes how to set up a shared disk cluster for an SAP Business Warehouse system on a multi-partition IBM DB2 for Linux, UNIX, and Windows (IBM DB2 for LUW) database using IBM Tivoli System Automation for Multiplatforms (SA MP). The document describes how you must configure your file systems for database partitions according to a high available cluster node model and its corresponding failover policy. In addition, the document describes the SAP installation tool *SAPinst* and the SAP cluster setup tool that helps you to set up and to configure your database partition cluster with just a few steps in an automated manner.

Using this document, you learn how to setup a highly available SAP NetWeaver BW system on IBM DB2 for LUW using the Database Partition Feature (DPF) together with Tivoli SA MP. The use of DPF together with SA MP requires disks for each partition that can be moved between nodes. This is called a shared disk cluster that stays not in conflict with the shared-nothing architecture of DPF, as you will see in the following. You will learn how to configure and set up your database partitions and how your cluster manager can handle the partitions in the case of a failure.

This document contains examples of how to use the SAP and IBM utilities to set up your specific cluster environment. The example is related to AIX systems, but can easily be adapted to other platforms, such as Linux.

Author: Hinnerk Gildhoff (IBM) , Brigitte Bläser (IBM)

Company: IBM Deutschland Research & Development GmbH

Created on: 15 November 2011

Author Bio



Hinnerk Gildhoff works as a software engineer for IBM Deutschland R&D GmbH in the SAP DB2 development team. He focuses on SAP and DB2 high availability concepts, automatization, and distributed technologies. His experience covers various platforms like Linux, UNIX, and Windows as well as technologies such as IBM Tivoli System Automation for Multiplatforms, Microsoft Cluster Services, Virtualization and more.



Brigitte Bläser is a senior software engineer at IBM Deutschland Research & Development GmbH in Böblingen. She is a member of the joint IBM/SAP development team that enables SAP applications for DB2 for Linux, UNIX, and Windows and is responsible for integrating DB2 for LUW solutions into SAP NetWeaver BW.

Table of Contents

Introduction	3
DB2 Database Partition Feature (DPF)	3
Volume Groups, Logical Volumes, and File Systems	3
SAP Data Directories and DB2 Tablespaces with Automatic Storage	4
Local Database Directory and Database Log Directory	6
Shared Home Directory of the DB2 Administrator	6
Shared or Local Directories: DB2 Dump Directory and DB2 Software Directory	7
Installing SAP NetWeaver BW on DB2 with DPF	8
The SAP Installation Tool SAPInst.....	8
Creating a Multi Database Partition Cluster	13
Using the SAP IBM Cluster Setup Tool	13
Customizing Failover Policy	15
Validating Cluster Functionality	16
Database Partition Failover.....	16
Database Host Failover.....	16
List of Abbreviations	16
Table of Figures.....	17
Related Content.....	18
Copyright.....	19

Introduction

This document describes how to set up a shared disk cluster for an SAP Business Warehouse system on a multi-partition IBM DB2 for Linux, UNIX, and Windows (IBM DB2 for LUW) database using IBM Tivoli System Automation for Multiplatforms (SA MP).

In this section, you will learn how to prepare an operating system before using the SAP installation tool *SAPinst* to install a SAP Business Warehouse system on a multi-partition database environment with IBM DB2 for LUW. This section gives an overview, but it is not a complete guide about storage layout, database partitions, and other operating system-relevant topics. It requires a certain knowledge and experience with these topics. This section highlights only important design decisions that should be kept in mind to avoid typical pitfalls.

DB2 Database Partition Feature (DPF)

The database partitioning feature allows the segmentation of an IBM DB2 for LUW database into several database partitions that can be located on a single large Symmetric Multiprocessing (SMP) database server (scale-up) or on several smaller discrete servers (scale-out).

DPF follows the shared-nothing architecture. That is, each database partition is a part of the database and consists of its own data, indexes, configuration files, and transaction logs. Tables can be distributed automatically over several database partitions using a hash key that is defined when the table is created.

Since data is distributed across database partitions, the throughput and bandwidth of many processors is aggregated and thus reduces the runtime of SQL access and database maintenance operations. Data retrieval and data manipulation requests are automatically decomposed into sub-requests, and executed in parallel on the database partitions where data resides. The fact that databases are split across database partitions is transparent to users issuing SQL statements.

The database partitioning feature is supported for SAP NetWeaver BW systems and SAP systems based on SAP NetWeaver BW, like SAP Supply Chain Management (SAP SCM) and SAP Strategic Enterprise Management (SAP SEM).

Volume Groups, Logical Volumes, and File Systems

Before installing an SAP NetWeaver BW system on DB2 using the Database Partition Feature, think carefully about your storage paths with respect to your preferred high availability failover policy.

In most cases, customers do not use cluster file systems like GPFS. They use normal file systems, such as JFS2 or EXT3. For such non-cluster file systems, a defined volume group can only be active on one node at the same time. This is important to understand because it implies the need to group logical volumes (LV) in logical volume groups (VG) that belong to only one database partition. Thus, we can move all volume groups belonging to one database partition from one node to a spare node and activate them without interrupting other database partition processing.

Each database partition uses the following directories: one local database directory, one directory for transaction logs, and `<n> sapdata` directories for the data. Each of these directories can be created as a separate file system on a logical volume that resides in a separate volume group.

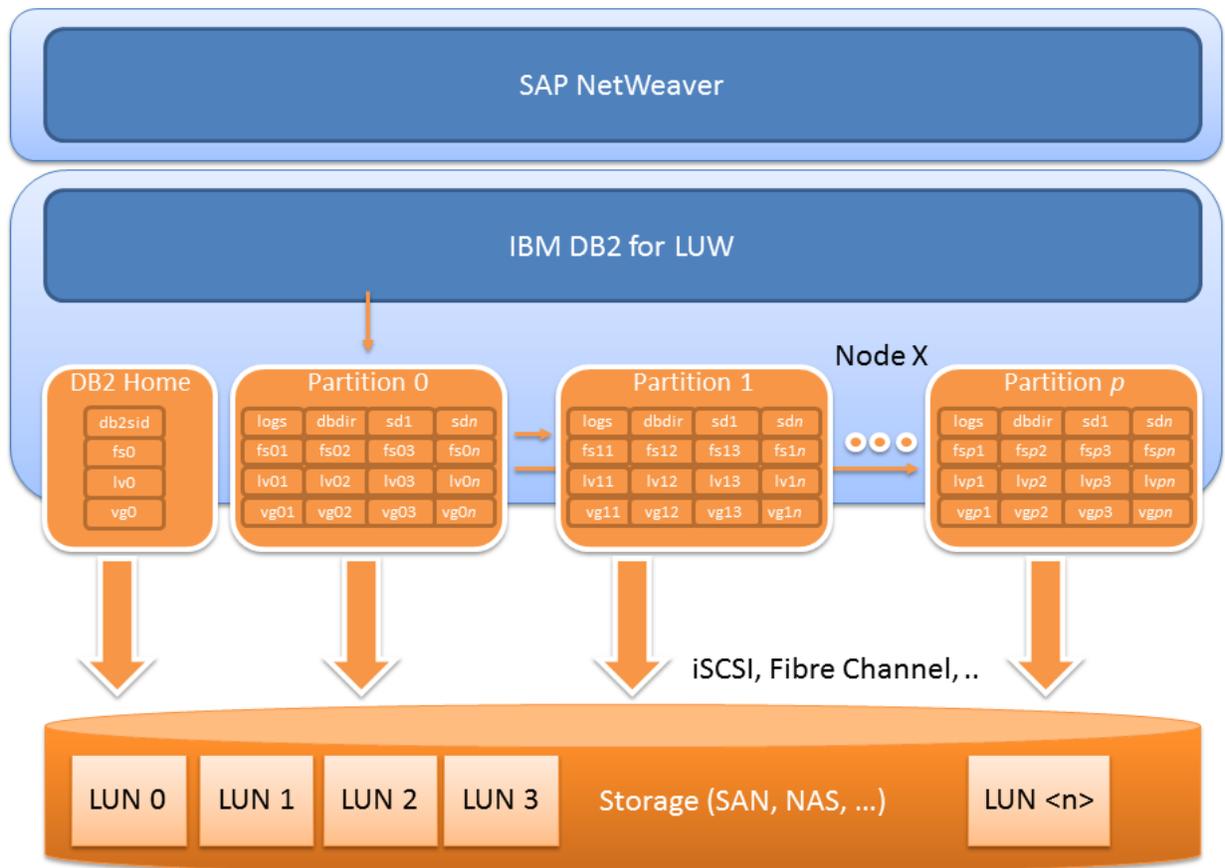
Note: Make sure that all mount points are not automatically mounted at system restart because this would result in a conflict with the cluster manager.

Thus, if you have, for instance, four `sapdata` directories, we recommend that you create the following volume groups, each with one logical volume and one file system:

- One volume group for the instance directory

- Six volume groups for each database partition:
 - One for the local database directory
 - One for the transaction logs
 - One for each of the four *sapdata* directories

This ensures optimal performance and provides you with the benefit that every database partition can be moved with a planned failover to a spare node (see Figure 1 - Storage Layout). Be aware that in the case of an unplanned failover (due to a node failure), all node partitions have to move to a different spare node.



fs<n> := file system <n> lv<n> := logical volume <n> vg<n> := volume group <n> sd<n> := sapdata <n>

Figure 1 - Storage Layout

SAP Data Directories and DB2 Tablespaces with Automatic Storage

Today, SAP and IBM recommend automatic storage for DB2 tablespaces. This is the default value in *SAPinst*, which means that when creating a database with automatic storage, you specify a set of storage paths that are called *sapdata* directories in the SAP world. *SAPinst* uses these *sapdata* paths to create the DB2 tablespaces without specifying the location of the containers explicitly. The DB2 engine automatically creates one DMS file container in each storage path for each tablespace.

Example of first *sapdata* path specified in *SAPinst*:

```
/db2/<SID>/sapdata1
```

Resulting container path created by DB2 automatically:

```
/db2/<SID>/sapdata1/db2<sid>/NODE0000/<SID>/T0000000/C0000000.LRG
```

Our recommended approach uses the default values of *SAPinst*, which means automatic storage in combination with unspecific storage paths. The default layout has four *sapdata* directories for every database partition. The following describes the resulting path structure for unique mount paths:

```
/db2/<DBSID>/sapdata1/db2<sid>/NODE0000
/db2/<DBSID>/sapdata1/db2<sid>/NODE0001
/db2/<DBSID>/sapdata1/db2<sid>/NODE0002
```

```
/db2/<DBSID>/sapdata2/db2<sid>/NODE0000
/db2/<DBSID>/sapdata2/db2<sid>/NODE0001
/db2/<DBSID>/sapdata2/db2<sid>/NODE0002
```

```
/db2/<DBSID>/sapdata3/db2<sid>/NODE0000
/db2/<DBSID>/sapdata3/db2<sid>/NODE0001
/db2/<DBSID>/sapdata3/db2<sid>/NODE0002
```

```
/db2/<DBSID>/sapdata4/db2<sid>/NODE0000
/db2/<DBSID>/sapdata4/db2<sid>/NODE0001
/db2/<DBSID>/sapdata4/db2<sid>/NODE0002
```

It is important to understand the resulting storage path structure because in our planned cluster setup, we need mount points to unique file systems (see Figure 2 – File Systems) for each database partition, which can failover to a different cluster node. So we have to manually create file systems on every node path and mount them before we start *SAPinst*.

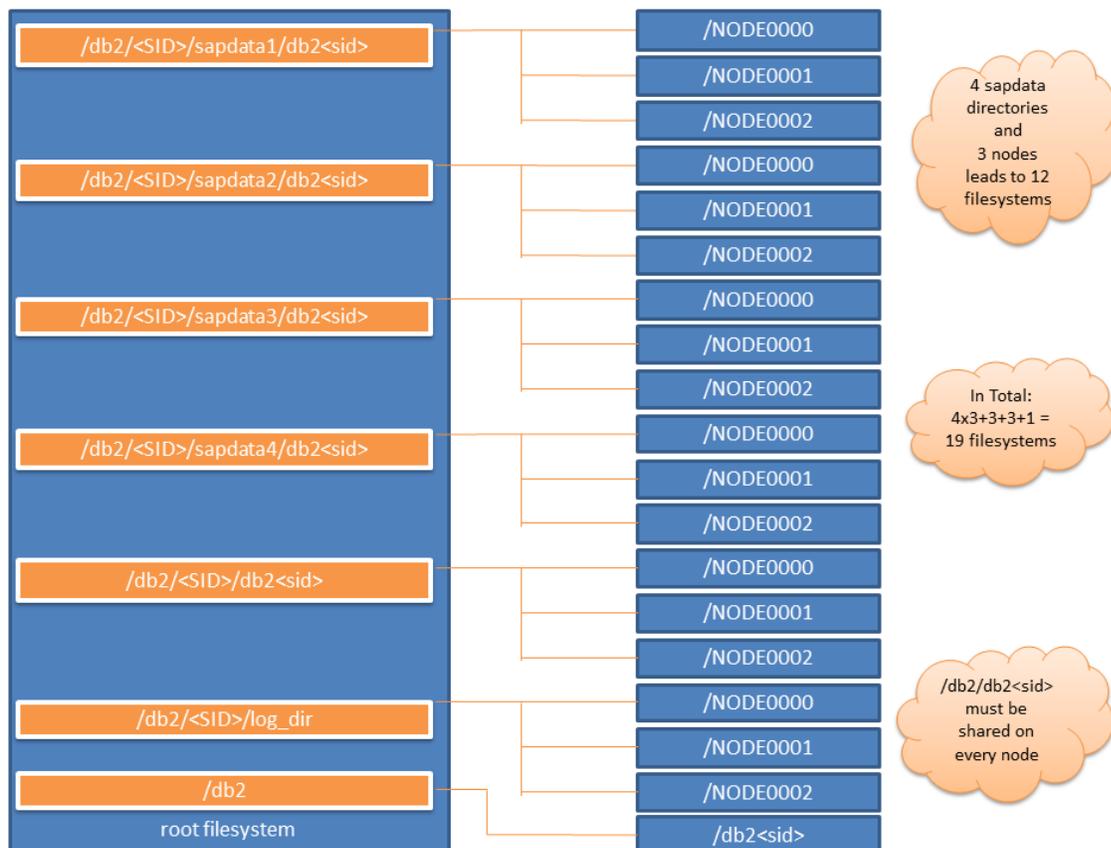


Figure 2 – File Systems

Thus, in a cluster environment with four *sapdata* directories and three nodes, we will have $4 \times 3 = 12$ mount points (file systems) just for the partition data because the *sapdata* directory appears before the *node* directory in the path structure! In addition, we will have three more file systems for each database directory, three for each log directory and one shared filesystem for `/db2/db2<sid>`, as you can see in Figure 2 – File Systems.

Note: Before running `SAPinst`, ensure that the storage paths that you manually created have no typos and are mounted on the correct file system. Otherwise, DB2 will store the data on a different file system that cannot be used for high availability.

Note: Do not use nested link structures or other types of links in the storage path. This will not work with the DB2 High Availability Instance Configuration Utility (`db2haicu`).

Local Database Directory and Database Log Directory

The local database directory and the local log directory exist for every database partition. They need a mount point to a unique file system as well as the database partition data described above, so that we can mount these file systems on different cluster nodes in the case of a failure. Ensure that you always create and mount the file systems before you run `SAPinst`.

The default path structure for these directories is adequate for high availability so that we can use the following mount points to unique file systems:

```
/db2/<DBSID>/db2<sid>/NODE0000
/db2/<DBSID>/db2<sid>/NODE0001
...
/db2/<DBSID>/db2<sid>/NODE000<n>
```

Database partitions have their own log directories, as shown in the following:

```
/db2/<DBSID>/log_dir/NODE0000
/db2/<DBSID>/log_dir/NODE0001
...
/db2/<DBSID>/log_dir/NODE000<n>
```

Shared Home Directory of the DB2 Administrator

The home directory of the database instance owner `db2<dbsid>` must be shared between all cluster nodes because the following, for example, must be shared:

- The database instance owner environment (e.g. DB2 registry)
- The `db2nodes.cfg` file that are automatically handled by the cluster manager.

The SAP default path for the database instance owner home directory is as follows:

```
/db2/db2<dbsid>/
```

Shared or Local Directories: DB2 Dump Directory and DB2 Software Directory

In addition, there are two other directories for our cluster setup which could be shared between the nodes.

The first directory is the db2dump directory for diagnose log files and database notifications. You can share this directory for all nodes, or use a local directory for every node in your cluster environment. We suggest a local dump directory on its own file system, but this is your decision. The SAP default path for the *db2dump* directory is the following:

```
/db2/<DBSID>/db2dump
```

The second directory is the database software directory. If you use the SAP default installation structure, your database software directory is installed in the home directory of the database instance owner and is already shared. In this case, there is nothing to do. Alternatively, you can use a different directory for your database software (e.g. /opt/IBM/db2/V9.7/FP4/), which enables you to decide on whether you want to share the directory or whether you want to use local database software directories. The SAP default path for the software directory is as follows:

```
/db2/db2<dbsid>/db2_software
```

Installing SAP NetWeaver BW on DB2 with DPF

After configuring your operating system - especially the underlying storage layout - you have to enable *ssh* or *rsh* login for user *root* without password authentication. This enables user *root* to connect to all nodes of the cluster, which is required to create and manage the cluster environment. Afterwards, you can install an SAP NetWeaver Business Warehouse system, add database partitions, and create the cluster resources.

The SAP Installation Tool SAPinst

You can use *SAPinst* to install a fresh SAP NetWeaver system, or perform a SAP system copy to export one of your existing SAP systems into a new SAP system. It is not relevant for now which method you choose; we only recommend that you use the high-availability installation option on the SAPinst welcome screen (see Figure 3 - SAPinst High-Availability System). This enables you to separate your SAP system into different SAP instances, each running on its own virtual host name (see usage of SAPinst parameter `SAPINST_USE_HOSTNAME`), which is very recommendable for cluster systems.

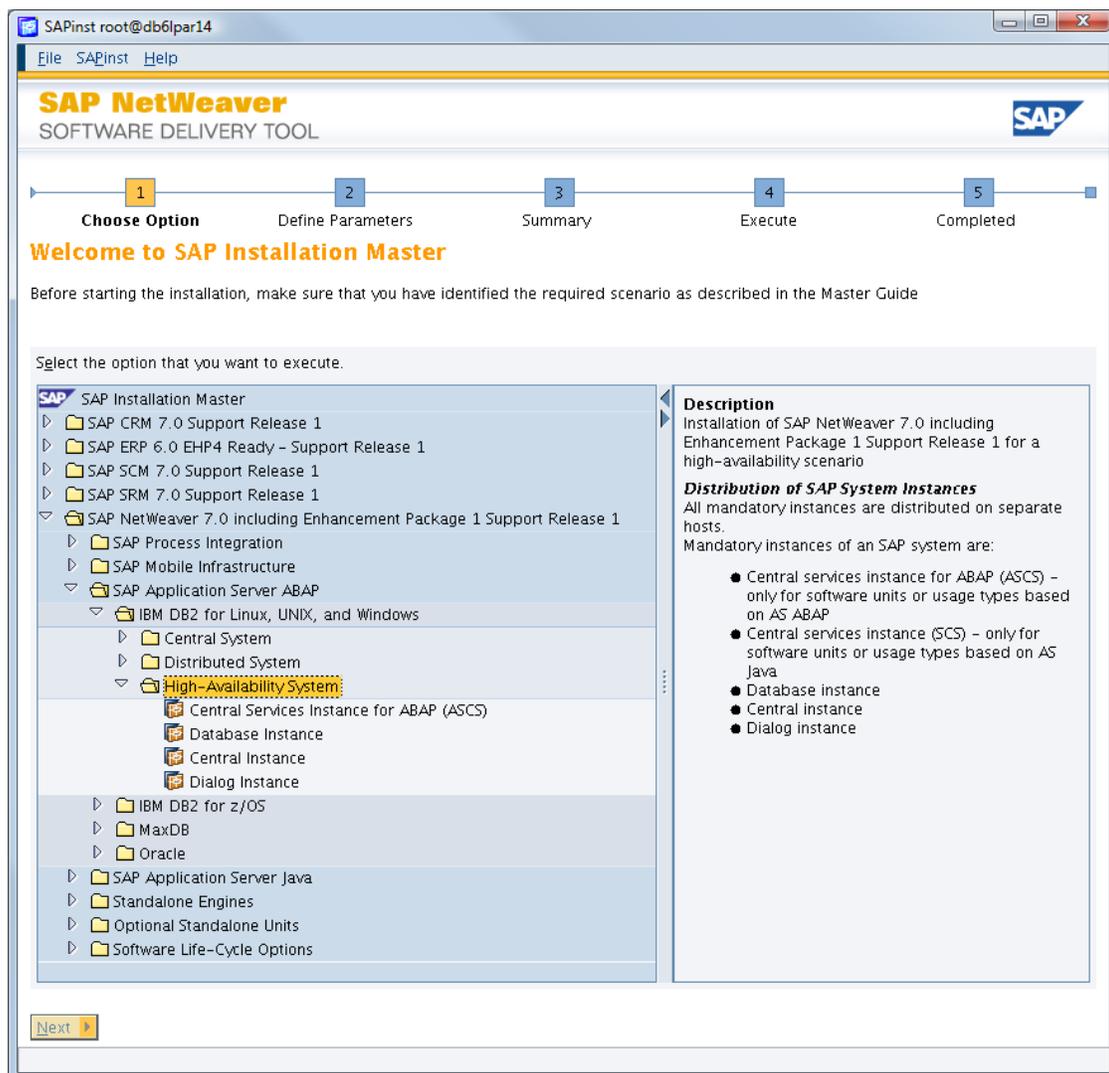


Figure 3 - SAPinst High-Availability System

One of the relevant dialogs regarding a high-availability DPF system appears in the dialog phase of the database installation process (Figure 4 - SAPinst Sapdata Directories). Here, you can enable and disable automatic storage and specify the *sapdata* directories. Ensure that you have selected automatic storage and

that the number of sapdata directories corresponds to the already created directories. The *sapdata* directory path structure must exist (end with the node directory) and must be mounted on a unique file system.

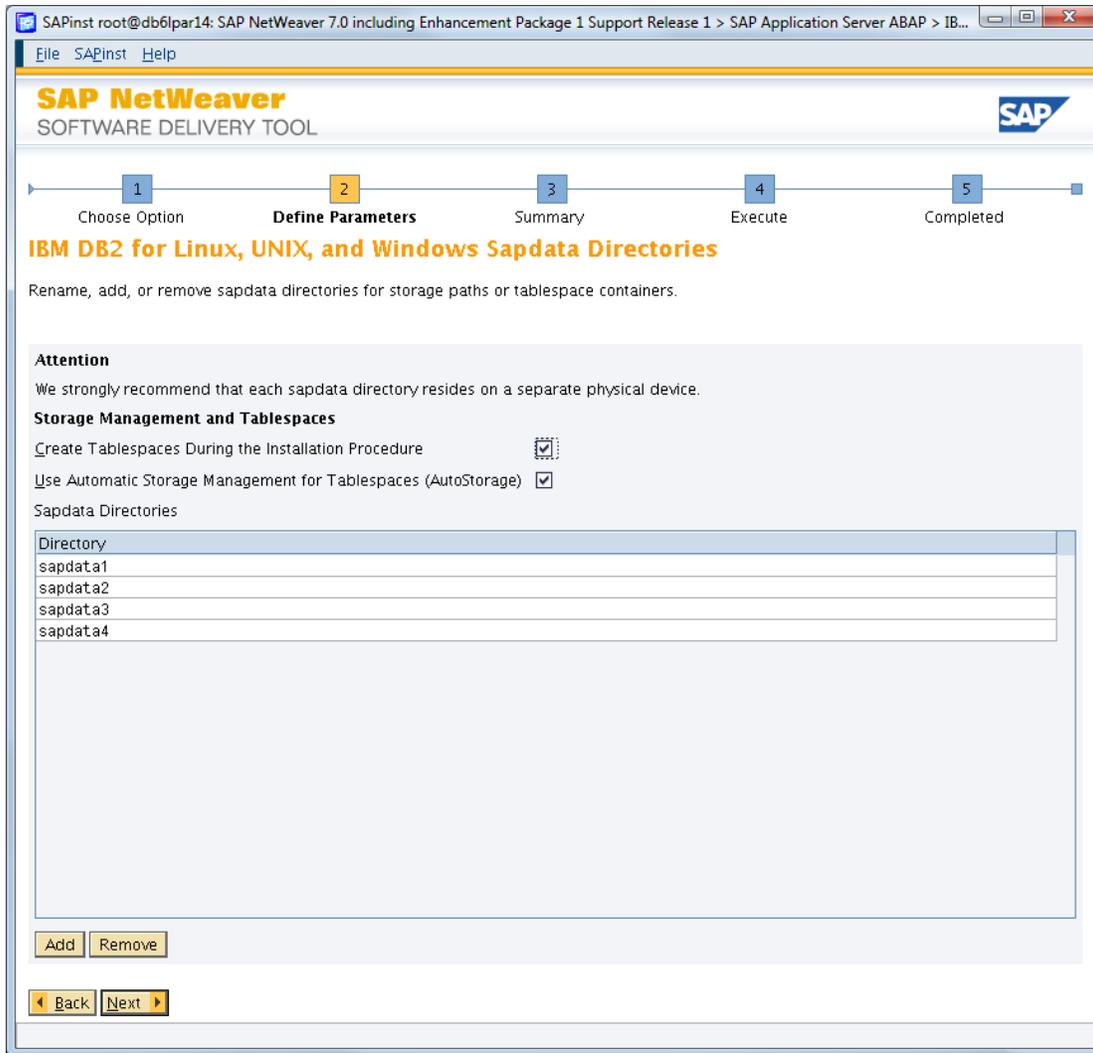


Figure 4 - SAPinst Sapdata Directories

In addition, the database installation phase offers you another special dialog that is important for our cluster system. This dialog requires input for your high availability setup. Choose the option to install IBM Tivoli System Automation for Multiplatforms (see Figure 5 - SAPinst SA MP) and choose shared disk instead of HADR as your cluster type. Shared disk as a cluster type just means that your data is stored on a disk that can be mounted on different nodes, in contrast to an HADR cluster, where the data is synchronized between two separate disks (disaster recovery concept). The required Tivoli *SA MP* software and the corresponding Tivoli *SA MP* policies are installed on your current host only where *SAPinst* currently runs.

The software and the policies are required on every node of your cluster. You have to install the *SA MP* software on all other cluster nodes manually (use *installSAMP.sh* from the RDBMS-DVD). The *SA MP* policies are distributed automatically from our Cluster Setup Tool as described in the chapter *Using the SAP IBM Cluster Setup Tool*.

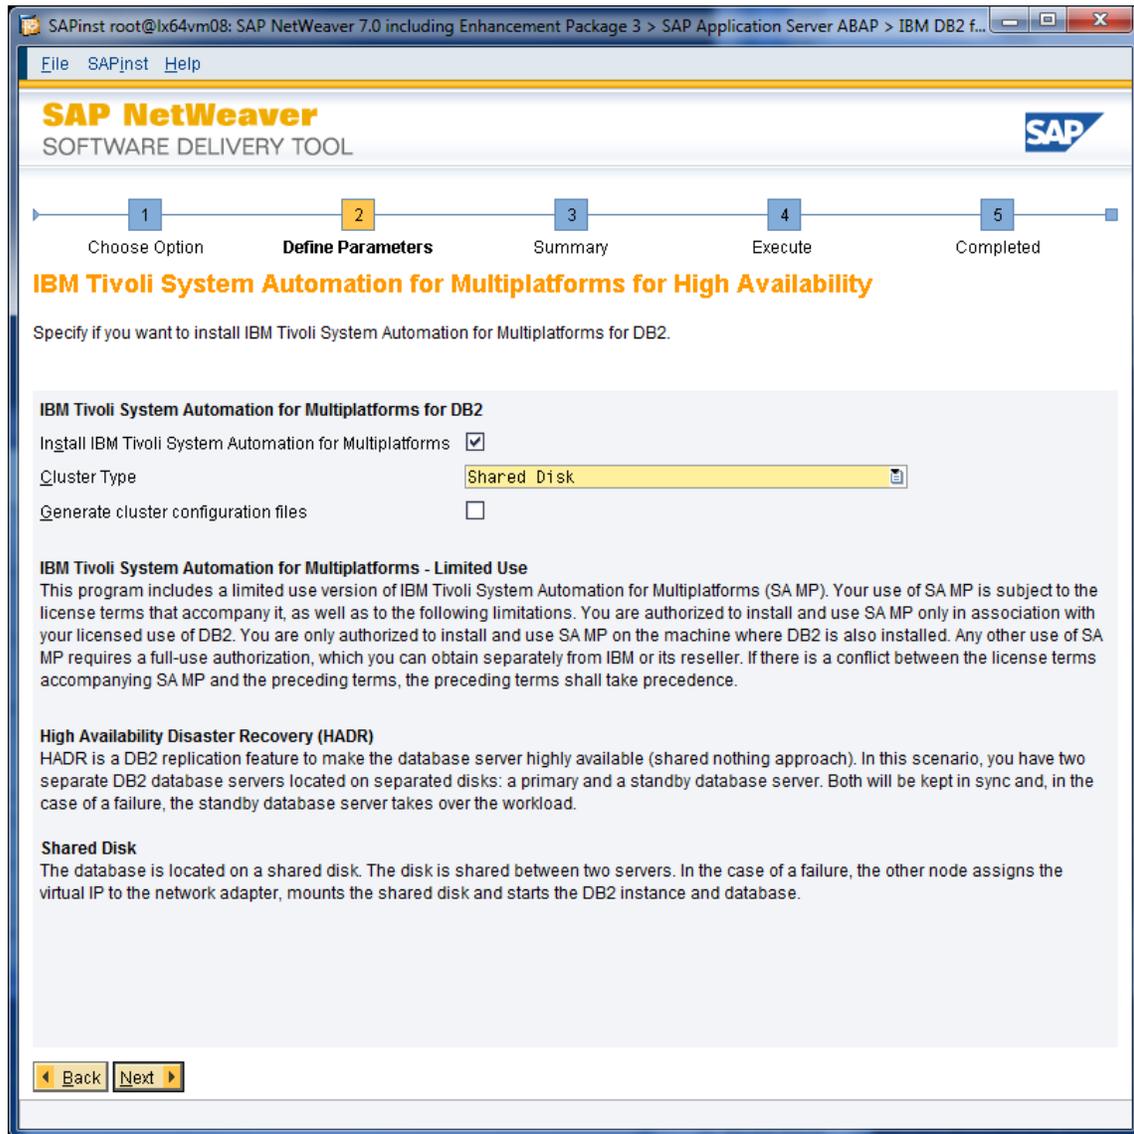


Figure 5 - SAPinst SA MP

So, after having installed the central services of SAP, the database instance, and the central instance of SAP, you can use *SAPinst* again to add physical or logical database partitions to your SAP system (see Figure 6 – SAPinst: Additional Database Partition).

Note: If you migrate your SAP system to IBM DB2 for LUW, you have to use an exit step in SAPinst to create additional database partitions. This is a slightly different process from installing a new single SAP system and adding database partitions afterwards.

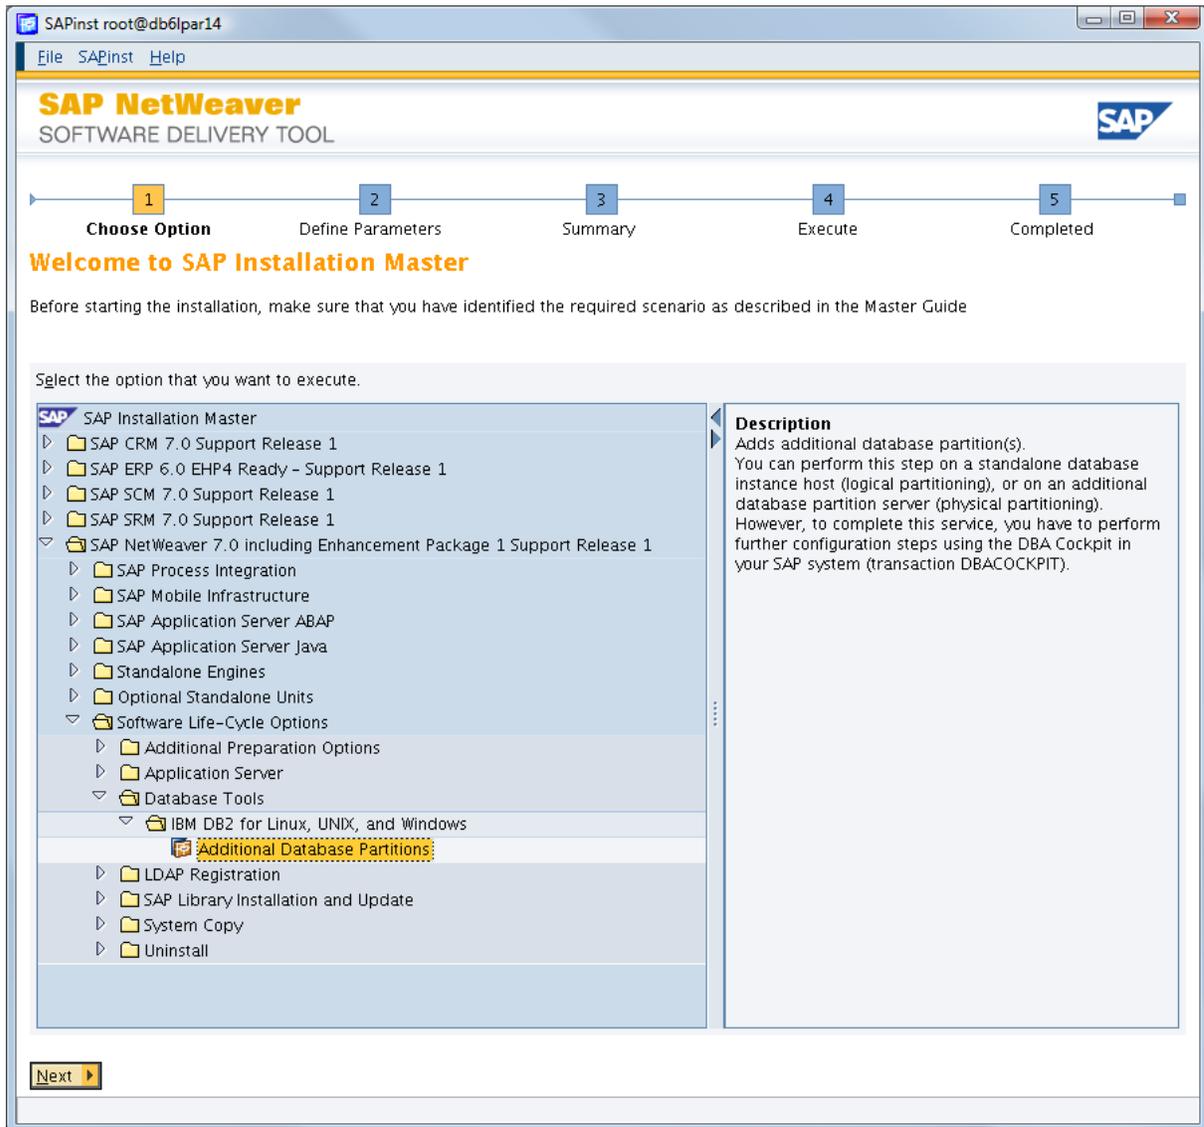


Figure 6 – SAPinst: Additional Database Partition

The following step only needs to be executed if you install a new SAP NetWeaver BW system but not if you create a system copy:

After your database partitions have been installed, log on to your SAP system and call the DBA Cockpit (SAP transaction DBACOCKPIT). In the navigation frame, choose *BW Administration* and *BI Data Distribution* to distribute your data to the newly created database partitions.

SAP and IBM recommend to distribute all ODS and FACT partition groups to all database partitions except partition zero, which has a special role in a multi partition environment (see Figure 7 - BW Data Distribution).

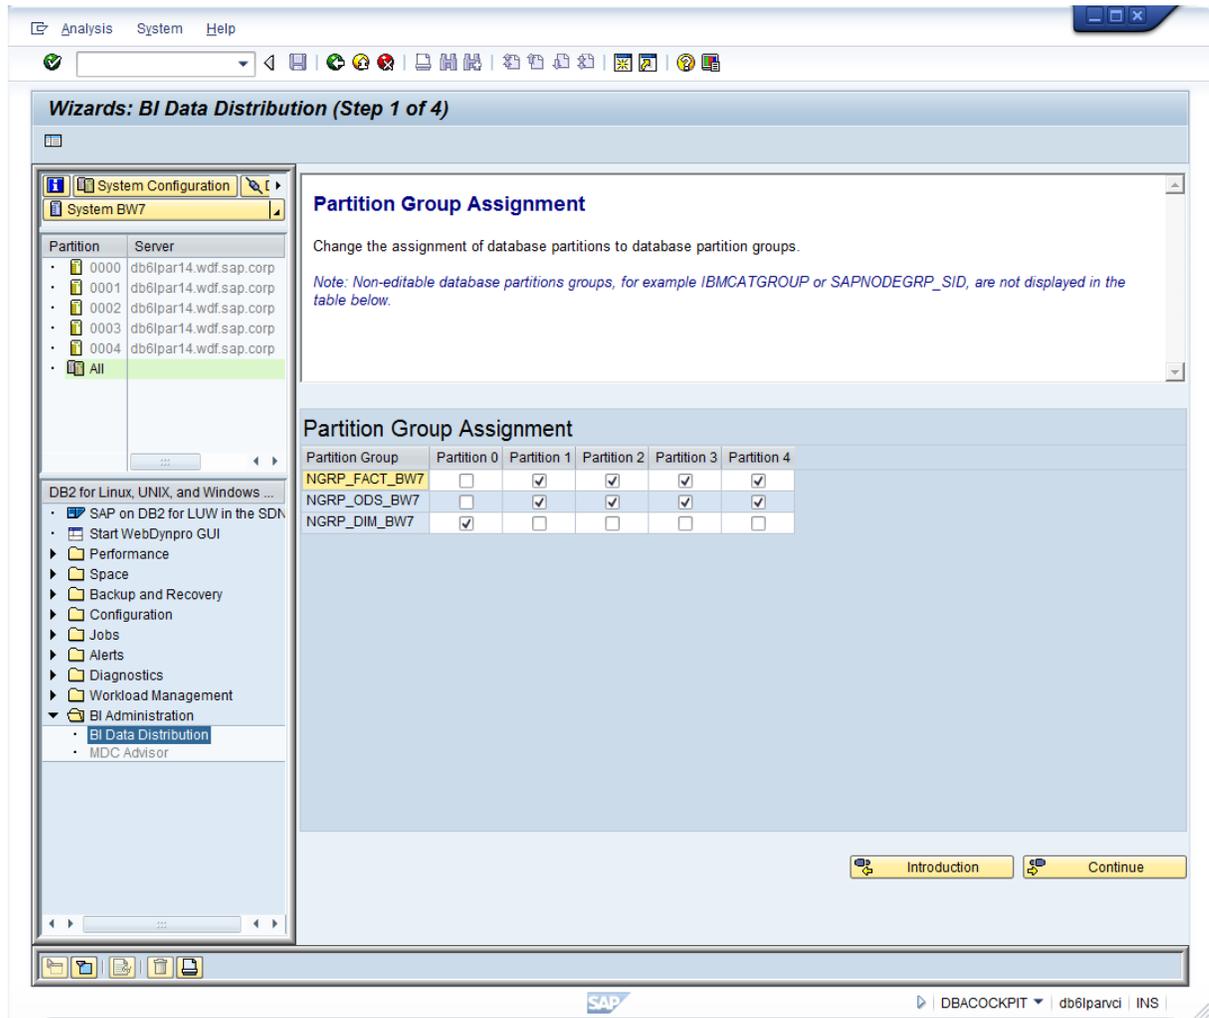


Figure 7 - BW Data Distribution

After the data distribution, your SAP system is a normal multi-partition business warehouse system. In the next chapter, you learn how to activate the cluster manager and how to apply the cluster policies to your environment to achieve a fully automated shared disk cluster scenario.

Creating a Multi Database Partition Cluster

The following section explains how to set up and configure Tivoli System Automation for your database cluster.

Using the SAP IBM Cluster Setup Tool

You can download the Cluster Setup Tool from SAP Note 960843. The tool automates the cluster setup process for HADR-, single database partition-, and multi database partition cluster.

Figure 8 - *Create, Show or Edit Database Configuration* - shows a generated cluster configuration for a multi partition cluster with two nodes (see DB2_HOSTNAME_LIST) and two database partitions (see DB2_DIR_LIST), each with only one *sapdata* directory.

```

General Cluster Configuration

[1]  SAP_SID                =  BW7
[2]  SAP_CI_HOSTNAME       =  db61parvci
[3]  TSA_DOMAIN_NAME      =  sap_bw7db2
[4]  TSA_TIEBREAKER_IP_ADDRESS =  10.17.202.1
[5]  TSA_REMOTE_CMD       =  ssh

Database Cluster Configuration

[6]  DB2_HOSTNAME_LIST    =  db61par14
                               db61par15
[7]  DB2_CLUSTER_TYPE     =  SD_DPF
[8]  DB2_INST_DIR        =  /db2/db2bw7/db2_software
[9]  DB2_DB2INSTANCE     =  db2bw7
[10] DB2_HA_HOSTNAME     =  db61parvdb
[11] DB2_HA_IP_ADDRESS   =  10.17.202.99
[12] DB2_NETWORK_INTERFACE_LIST =  en0:10.17.202.217:255.255.254.0:db61par14
                               en0:10.17.202.218:255.255.254.0:db61par15
[13] DB2_HA_IP_MASK      =  255.255.254.0
[14] TSA_LICENSE_FILE    =  /data/samp_dev/sam31.lic
[15] TSA_USER_LIST       =  bw7adm
                               db2bw7
[16] DB2_DIR_LIST        =  /db2/BW7/sapdata1/db2bw7/NODE0001
                               /db2/BW7/db2bw7/NODE0001
                               /db2/BW7/log_dir/NODE0001
                               /db2/BW7/sapdata1/db2bw7/NODE0000
                               /db2/BW7/db2bw7/NODE0000
                               /db2/BW7/log_dir/NODE0000
[17] TSA_USER_GROUP_NAME =  s agrp
[18] TSA_USER_GROUP_ID   =  230
[19] DB2_COMM_PORTS     =  sapdb2BW7:6612/tcp
                               DB2_db2bw7:6614/tcp
                               DB2_db2bw7_1:6615/tcp
                               DB2_db2bw7_2:6616/tcp
                               DB2_db2bw7_3:6617/tcp
                               DB2_db2bw7_4:6618/tcp
                               DB2_db2bw7_END:6619/tcp
[20] DB2_GROUP_LIST     =  dbbw7adm:207:true
                               dbbw7ctl:210:true
                               dbbw7mnt:211:true
                               sapsys:217:true
                               sapinst:209:true
[21] DB2_USER_LIST      =  sapadm:225:/home/sapadm:true:/bin/csh:true:sapsys:sapinst
                               bw7adm:212:/home/bw7adm:true:/bin/csh:true:sapsys:sapinst:dbbw7ctl
                               db2bw7:215:/db2/db2bw7:false:/bin/csh:true:dbbw7adm:sapinst:
                               sapbw7:216:/home/sapbw7:true:/bin/csh:true:dbbw7mnt:staff

```

Figure 8 - Create, Show or Edit Database Configuration

In the case of a shared disk cluster, the tool executes many of replication tasks on all your cluster nodes, such as replicating ports, user groups, users, and SA MP policies.

Figure 9 - *Create Database Cluster* - shows how the Cluster Setup Tool uses this information to create a running cluster environment. The tool iterates over all nodes of your cluster and prepares them for the shared disk cluster. The last step is calling the db2 executable *db2haicu* that creates the DB2 SA MP resources.

```

Check general configuration
  Check database cluster configuration : OK

Create cluster domain and nodes
  Create SA MP domain
    Check for SA MP software on cluster nodes
      Check for SA MP on node db61par14 : OK
      Check for SA MP on node db61par15 : OK
    Prepare cluster nodes : OK
    Create cluster domain sap_bw7db2 : OK [skipped]

Create database cluster resources
  Create database cluster (shared disk)
    Check DB2 cluster : OK
    Prepare DB2 cluster
      Distributing DB2 policy : OK [skipped]
      Mount shared filesystems to node db61par14 : OK [skipped]
      Replicate user groups
        Create user group dbbw7adm on node db61par15 : OK [skipped]
        Create user group dbbw7ctl on node db61par15 : OK [skipped]
        Create user group dbbw7mnt on node db61par15 : OK [skipped]
        Create user group sapsys on node db61par15 : OK [skipped]
        Create user group sapinst on node db61par15 : OK [skipped]
      Replicate users
        Create user sapadm on node db61par15 : OK [skipped]
        Create user bw7adm on node db61par15 : OK [skipped]
        Create user db2bw7 on node db61par15 : OK [skipped]
        Create user sapbw7 on node db61par15 : OK [skipped]
      Grant SA MP control to bw7adm,db2bw7 on all nodes
        Create user group sagrp on node db61par14 : OK [skipped]
        Grant SA MP control on node db61par14 : OK [skipped]
        Create user group sagrp on node db61par15 : OK [skipped]
        Grant SA MP control on node db61par15 : OK [skipped]
      Replicate DB2 communication ports : OK [skipped]
      Disable DB2 Fault Monitor : OK [skipped]
    Configuring SAP for HA DB2 cluster
      Modifying User Environment : OK [skipped]
      Update SAP profiles with virtual database hostname : OK [skipped]
      Update thin client configuration with virtual database hostname : OK [skipped]
    Start database server on node db61par14 : OK [skipped]
  Create DB2 Cluster using db2haicu
    Generate db2haicu configuration file (/tmp/cluster_config.xml) : OK
    Removing virtual IP address from cluster nodes : OK [skipped]
    Execute db2haicu at node db61par14 : OK

```

Figure 9 - Create Database Cluster

In this example, the script is executed on host *db61par14* and all replication tasks are done with the spare node *db61par15*.

Figure 10 - *Show Database Cluster State* - shows you the SA MP perspective of the generated cluster spanning both nodes (*db61par14* and *db61par15*). Both resource groups represent one database partition with more than one resource. There is a resource for the partition logs, the local database directory, the partition data, and the database partition itself. In addition, database partition zero has a special role and contains a virtual IP resource in the resource group so that all database requests contact partition zero first.

By default, the Cluster Setup Tool defines a mutual cluster failover policy and lets the database partitions fail over from the first to the second node (which is defined in DB2_HOSTNAME_LIST). If you prefer a different setup, you have to call db2haicu again with a customized XML file (see section *Customizing Failover Policy*).

```

-----
| IBM Tivoli System Automation for Multiplatforms 05/05/11 15:39:08 |
-----
Online IBM.ResourceGroup:db2_db2bw7_0-rg Nominal=Online
|- Online IBM.Application:db2_db2bw7_0-rs
|   |- Online IBM.Application:db2_db2bw7_0-rs:db61par14
|   '--- Offline IBM.Application:db2_db2bw7_0-rs:db61par15
|- Online IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0000-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0000-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0000-rs:db61par15
|- Online IBM.Application:db2mnt-db2_BW7_log_dir_NODE0000-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_log_dir_NODE0000-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_log_dir_NODE0000-rs:db61par15
|- Online IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0000-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0000-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0000-rs:db61par15
'--- Online IBM.ServiceIP:db2ip_10_17_202_99-rs
|   |- Online IBM.ServiceIP:db2ip_10_17_202_99-rs:db61par14
|   '--- Offline IBM.ServiceIP:db2ip_10_17_202_99-rs:db61par15
Online IBM.ResourceGroup:db2_db2bw7_1-rg Nominal=Online
|- Online IBM.Application:db2_db2bw7_1-rs
|   |- Online IBM.Application:db2_db2bw7_1-rs:db61par14
|   '--- Offline IBM.Application:db2_db2bw7_1-rs:db61par15
|- Online IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0001-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0001-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_db2bw7_NODE0001-rs:db61par15
|- Online IBM.Application:db2mnt-db2_BW7_log_dir_NODE0001-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_log_dir_NODE0001-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_log_dir_NODE0001-rs:db61par15
'--- Online IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0001-rs
|   |- Online IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0001-rs:db61par14
|   '--- Offline IBM.Application:db2mnt-db2_BW7_sapdata1_db2bw7_NODE0001-rs:db61par15
Online IBM.Equivalency:db2_db2bw7_0-rg_group-equ
|- Online IBM.PeerNode:db61par14:db61par14
'--- Online IBM.PeerNode:db61par15:db61par15
Online IBM.Equivalency:db2_db2bw7_1-rg_group-equ
|- Online IBM.PeerNode:db61par14:db61par14
'--- Online IBM.PeerNode:db61par15:db61par15
Online IBM.Equivalency:db2network
|- Online IBM.NetworkInterface:en0:db61par14
'--- Online IBM.NetworkInterface:en0:db61par15

```

Figure 10 - Show Database Cluster State

Customizing Failover Policy

Before customizing the failover policy, you still need to run the Cluster Setup Tool at least once to prepare all cluster nodes for the shared disk scenario. Afterwards, you can repeat the last step of the Cluster Setup Tool with a manually customized failover policy. The last step of the Cluster Setup Tool is the creation of the SA MP resources with *db2haicu*.

After running the Cluster Setup Tool, call *db2haicu* manually again to delete the recently created resources. Your nodes stay prepared for the cluster.

```
#su - db2<dbsid>
#db2haicu -delete
```

Next, customize the XML configuration file for *db2haicu* according to your failover policy. For details about different possible configurations, check the documentation about sample input files for the DB2 High Availability Instance Configuration Utility (*db2haicu*) at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.ha.doc/doc/r0052514.html>

You can find the XML file already generated in the temporary directory on the node where you executed the Cluster Setup Tool.

After you have customized the configuration file, call `db2haicu` again to create the *SA MP* cluster resources with your failover policy as follows:

```
#db2haicu -f /tmp/cluster_config.xml
```

Validating Cluster Functionality

The following section explains how to validate your cluster setup. Before you leave your system under the control of SA MP, you have to test some basic cluster behavior to ensure that your cluster setup is valid and works as expected.

Database Partition Failover

Perform a manual failover of a specific cluster resource group and monitor with *Issam* if the resource successfully moves to a different node. Afterwards, log on to your SAP system and check if everything works as expected.

```
# rgreq -o move <rg>
# Issam
```

Database Host Failover

Reboot the node where database partitions are online, monitor the cluster resources on a different node, and check on this node with *Issam* if the resources successfully move to the defined spare node. Afterwards, log on to your SAP system and validate if everything works as expected.

```
node1# reboot
node2# Issam
```

List of Abbreviations

SAPinst	SAP Installation Tool
SAP NW BW	SAP NetWeaver Business Warehouse
DPF	Database Partition Feature
SMP	Symmetric Multiprocessing
SQL	Structured Query Language
GPFS	General Parallel File System
JFS2	Enhanced Journaled File System
EXT3	Third Extended File System
FS	File System
VG	Volume Group
LV	Logical Volume
LUN	Logical Unit Number
P0	Partition 0
iSCSI	Internet Small Computer System Interface
SMS	System-Managed Storage Tablespace
DMS	Database-Managed Storage Tablespace
SA MP	Tivoli System Automation for Multiplatforms
db2haicu	DB2 High Availability Instance Configuration Utility

Table of Figures

Figure 1 - Storage Layout	4
Figure 2 – File Systems	5
Figure 3 - SAPinst High-Availability System	8
Figure 4 - SAPinst Sapdata Directories	9
Figure 5 - SAPinst SA MP	10
Figure 6 – SAPinst: Additional Database Partition	11
Figure 7 - BW Data Distribution	12
Figure 8 - Create, Show or Edit Database Configuration	13
Figure 9 - Create Database Cluster	14
Figure 10 - Show Database Cluster State	15

Related Content

For more information, visit the [Landscape Design and Architecture homepage](#)

Copyright

© Copyright 2011 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Oracle Corporation.

JavaScript is a registered trademark of Oracle Corporation, used under license for technology invented and implemented by Netscape.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects S.A. in the United States and in other countries. Business Objects is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.