



Installation Guide | PUBLIC

Document Version: 1.1 – 2023-02-14

Running an SAP System on IBM Db2 10.1 with the Db2 pureScale Feature (Out of Mainstream Maintenance)


Content

1	Introduction.	4
1.1	Document History.	5
1.2	Architectural Overview of DB2 pureScale.	5
1.3	SAP Specifics.	7
2	Planning.	10
2.1	Installation Restrictions.	10
2.2	Installation Requirements.	11
3	Conversion of an Existing SAP System to a System Running on DB2 pureScale.	14
3.1	Introduction.	14
3.2	Overview of a DB2 File System Before and After the Conversion to DB2 pureScale.	15
3.3	Conversion to DB2 pureScale – Step by Step.	17
	Introduction.	17
	Upgrading the Database to DB2 10.1.	17
	Applying the Latest Kernel Patch.	18
	Performing an Offline Backup.	18
	Setting Required Configuration Parameters.	18
	Adding the DB2 pureScale Feature to the Existing DB2 10.1 Software Installation.	19
	Preparing the Cluster for the Conversion.	20
	Manually Setting Up Passwordless Access for User Root.	23
	Running the db2checkSD Utility.	24
	Preparing the GPFS Cluster.	25
	Creating the GPFS File Systems.	27
	Rebalancing Tablespaces to the New File System.	28
	Moving Remaining Data to the GPFS.	31
	Mounting GPFSs Under New Mount Points.	32
	Converting the DB2 10.1 ESE Instance to a DB2 10.1 pureScale Instance.	35
	Adapting the JDBC URL (AS Java Only).	39
	Testing the DB2 pureScale Installation.	40
	Adding Members and CFs.	41
	Running the db6_update_db Script.	42
	Checking Configuration Settings.	43
	Installing the DB2 pureScale License.	43
4	Installation of a New SAP System on DB2 pureScale.	44
4.1	Introduction.	44
4.2	Overview of the DB2 File Systems Required for a new Installation.	45

4.3	Installing a New SAP System – Step by Step.	46
	Introduction.	46
	Installing the DB2 10.1 Software.	47
	Preparing the Cluster for the Installation.	47
	Manually Setting Up Passwordless Access for User Root.	50
	Preparing the GPFS Cluster.	51
	Creating the GPFS File Systems.	53
	Installing the SAP System on a DB2 10.1 ESE Instance.	55
	Applying the Latest Kernel Patch.	60
	Setting Required Configuration Parameters.	61
	Converting the DB2 10.1 ESE Instance to a DB2 pureScale Instance.	61
	Adapting the JDBC URL (AS Java Only).	63
	Testing the New Installation.	64
	Adding Members and CFs.	65
	Installing the DB2 pureScale License.	66
	Running the db6_update_db Script.	66
	Checking Configuration Settings.	67
5	Implementation of SAP-Specific Enhancements.	68
5.1	Overview.	68
5.2	Adapting the DB2 Client Connectivity Setup.	68
	Introduction.	68
	Using the Round-Robin Connectivity Setup.	69
	Using a User-Defined Connectivity Setup.	71
	Adapting the DB2 Client Configuration After Installation or Deinstallation of SAP Application Servers.	73
5.3	Partitioning the Update Tables (AS ABAP Only).	74
5.4	Using DB2 pureScale Monitoring Enhancements in the DBA Cockpit (AS ABAP Only).	78
	Introduction.	78
	Monitoring the Cluster Status.	79
	Monitoring the Cluster Performance.	80
	Monitoring Buffer Pools in the DB2 pureScale Cluster.	81
	Monitoring the Package Cache in the DB2 pureScale Cluster.	82
	Monitoring the Cluster Caching Facility (CFs).	84
	Checking for Page Contention.	85
	Checking the Client Affinity.	87
6	Removing a DB2 pureScale Installation.	88
7	Disclaimer.91

1 Introduction

i Note


SAP systems running on IBM Db2 10.1 are out of mainstream maintenance. For more information, see SAP Note [1168456](#) .

This document explains how you create an SAP system running on IBM DB2 for Linux, UNIX, and Windows with the IBM DB2 pureScale Feature. It describes how you can do one of the following:

- Convert the DB2 10.1 database of an existing SAP system to an IBM DB2 10.1 database with the DB2 pureScale Feature
- Install a new SAP system on IBM DB2 10.1 with the DB2 pureScale Feature

This document contains information and recommendations specifically for SAP system environments.

i Note

Before you start, make sure that you read **SAP Note** [1718559](#) . This SAP Note contains the most recent information about the pureScale installation, as well as corrections to this document. Make sure that you always have the most recent version of this SAP Note.

Terminology and Variables Used in This Document

IBM Terminology

- *IBM DB2 Version 10.1 for Linux, UNIX, and Windows* is referred to as *DB2 10.1*.
- *IBM DB2 pureScale Feature for Enterprise Server Edition* is referred to as *DB2 pureScale*.
- *IBM cluster caching facility* is referred to as *CF*.
- *IBM Tivoli System Automation for Multiplatforms* is referred to as *SA MP*.
- *IBM General Parallel File System* is referred to as *GPFS*.

SAP Terminology

- *SAP NetWeaver application server ABAP* is referred to as *AS ABAP*.
- *SAP NetWeaver application server Java* is referred to as *AS Java*.

Variables

The following variables are used in this document:

Name of Variable	Description
<SAPSID>	SAP system ID in upper case

Name of Variable	Description
<sapsid>	SAP system ID in lower case
<DBSID>	Database name in upper case
<dbsid>	Database name in lower case

i Note

The database name is not necessarily the same name as the SAP system ID. For example, the database name is not necessarily the same as the SAP system ID in an environment with multiple components in one database (MCOD).

1.1 Document History

i Note

Before you start the implementation, make sure you have the latest version of this document. You can find it at https://help.sap.com/viewer/db6_purescale_10_1 on SAP Help Portal.

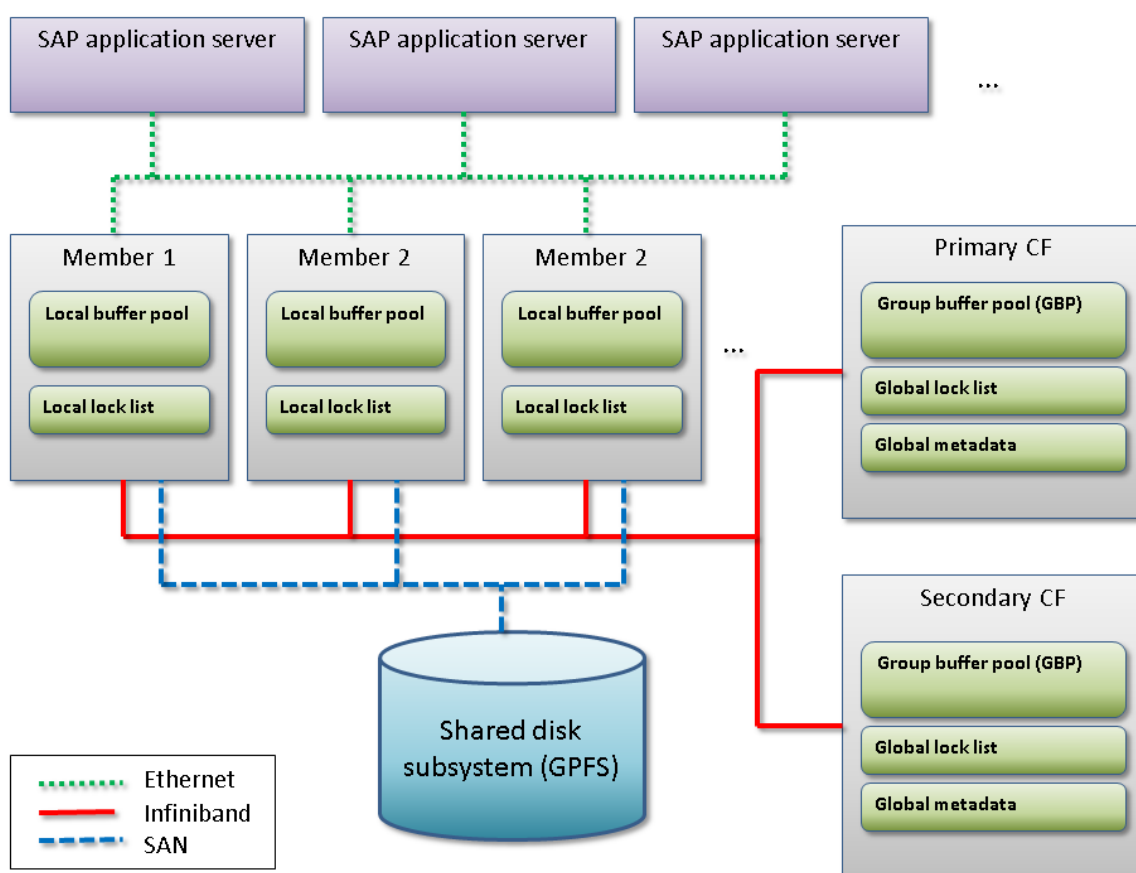
The following table provides an overview of the most important document changes:

Version	Date	Description
1.1	2023-02-14	SAP systems running on IBM Db2 10.1 are out of mainstream maintenance.
1.02	2021-05-31	Link correction
1.01	2018-05-11	Minor correction in section Installing the SAP System on a DB2 10.1 ESE Instance [page 55]
1.0	2012-08-16	First version

1.2 Architectural Overview of DB2 pureScale

The DB2 pureScale feature is an extension to the existing DB2 for Linux, UNIX, and Windows product. You can use DB2 pureScale to create a database cluster using the shared disk approach. The main focus of DB2 pureScale is on scalability and high availability.

The following figure provides an overview of the architecture of DB2 pureScale in an SAP environment:



Overview of DB2 pureScale Architecture

A node in a DB2 pureScale cluster is called a **member**. Each member runs its own DB2 instance using the DB2 server software, and has access to the complete database.

To assist in sharing data between the members and to particularly address the problems of concurrency control and cache coherency, DB2 pureScale introduces the **cluster caching facility** (also known as CF). The CF typically runs on a dedicated physical or logical host and is connected to all members of the cluster using a high speed interconnect, for example, Infiniband, which supports **remote direct memory access (RDMA)** and the **user direct access programming library API (uDAPL)**. On **Linux**, you can also use a 10 Gigabit Ethernet (10GE) network as high speed interconnect. This setup ensures that the memory of each member and the CF can be accessed without involvement of the respective CPU.

Since one CF would constitute a single point of failure, two CFs can be installed in a DB2 pureScale cluster, a **primary CF** and a **secondary CF**. The primary CF performs all CF workload for the cluster and the secondary CF can take over if the primary CF fails. Through a process called **duplexing**, the secondary CF is kept up-to-date so that a fast takeover is ensured and no vital information from the primary CF is lost.

To ensure concurrency control in the cluster, the **global lock manager (GLM)** running on the CF was introduced with DB2 pureScale. In addition, a new type of database locks is introduced. The locks that are used in a normal non shared database environment to ensure transaction serialization are now called **logical locks**. To prevent conflicting accesses to the same data from different pureScale members, **physical locks** are introduced.

Physical locks are held by the DB2 members, not by transactions. The **local lock manager (LLM)**, which is the lock manager of every member, requests a physical lock from the global lock manager before it grants the logical lock to the transaction. The GLM uses a global lock list to track the lock request made by the local lock managers.

Since physical locks work on pages, they can also be used to address the problem of cache coherency. The **group buffer pool (GBP)** in the CF is a single buffer pool that is used by all members. The members cache pages in their own **local buffer pools (LBPs)** and use the GBP to maintain page consistency between members. A set of protocols is used to determine:

- Which copy of a particular page that exists in multiple local buffer pools and the global buffer pool is the latest
- How to make changes to this page
- How to propagate these changes to other members

Finally, the CF holds a memory area that contains cluster-wide metadata that needs to be accessed and updated by all DB2 pureScale members.

The disk subsystem is accessed by all members of the pureScale cluster, and must therefore allow for concurrent read and write access while maintaining the integrity of the files. To achieve this, a **clustered file system** (also known as shared disk file system) is used. DB2 pureScale uses IBM's **general parallel file system (GPFS)**. GPFS is a high-performance, scalable file management solution that provides fast and reliable access to a common set of files. Management and installation of GPFS is embedded in DB2 pureScale. To prevent data loss if a single disk array fails and to avoid another single point of failure, an appropriate RAID level like RAID10 is recommended. The disk subsystem is attached to the DB2 pureScale cluster using a **storage area network (SAN)**, which again allows for fast communication with high-speed protocols, such as Fibre Channel.

Besides storing database data, the shared disk is also used to store the log files of the members. The log files of a member are called **log stream**, and every member writes its log stream into its own directory.

Another component that is required for DB2 pureScale (and which is not shown in the figure above) is the **cluster manager**. The cluster manager is responsible for checking the availability of all cluster components. If it detects the failure of a component, it automatically performs all necessary actions so that the impact on the overall system is as minimal as possible. For example, if the primary CF fails, the cluster manager automatically assigns the primary role to the secondary CF and redirects all requests to it. DB2 pureScale uses IBM Tivoli System Automation for Multi Platforms (SA MP) as cluster manager. SA MP was integrated in DB2 starting with in DB2 Version 9.5.

The SAP application servers (AS ABAP or AS Java) act as clients of the DB2 pureScale cluster. Normally, they are connected to the cluster using a TCP/IP network connection. The database shared library (DBSL) of the AS ABAP kernel uses the IBM Data Server Driver for ODBC and CLI to connect to DB2 pureScale. The AS Java uses the type 4 JDBC driver for the same purpose. The clustering of the database is transparent to its client, so that DB2 pureScale appears like a single database to the SAP system.

1.3 SAP Specifics

Various enhancements were made to SAP NetWeaver 7.0 SR3 and higher to deal with the following SAP-specific aspects in a DB2 pureScale cluster:

- Client affinity concept

- Partitioning of update tables
- Monitoring enhancements in the DBA Cockpit

Client Affinity Concept

Experience from running SAP applications on DB2 z/OS Parallel Sysplex has shown that the overall performance of the DB2 pureScale cluster increases if a similar workload is processed on every member. To realize a similar workload on each member, you can do the following:

- On SAP application level, you can use logon groups to direct SAP users with similar tasks to a specified set of SAP application servers.
- On database level, you can by default assign every member to one or more SAP application servers.

It is beneficial if you can define an affinity of each SAP application server to a specific DB2 pureScale member. Since every member can fail, you can enhance the concept of **client affinity** (keeping in mind that the SAP application servers act as clients of the DB2 pureScale cluster) by defining a list of members instead of a single member. If the first defined member in this list fails, the SAP application server connects to the second member, and so on. This concept does not imply that there is an equal number of SAP application servers and DB2 members available. Several SAP applications servers can, of course, connect to the same DB2 member.


An SAP system installation on DB2 pureScale uses the `db2dsdriver.cfg` configuration file of the DB2 client to define such an affinity. This file contains a list of all application servers as well as a client affinity list that controls to which DB2 member the respective SAP application server connects. The SAP installation tool creates an initial configuration in the `db2dsdriver.cfg` that you can later adapt to define certain affinities manually. For more information about configuration details regarding client affinity, see [Adapting the DB2 Client Connectivity Setup \[page 68\]](#).

Partitioning of Update Tables

The programming model in the AS ABAP defers updates to the database to the end of a business transaction. At that time, all updates to the database are performed asynchronously in update tasks. During a business transaction, all update requests are stored in the tables VBHDR, VBMOD, and VBDATA. These tables are critical for the overall performance of the system.

In a DB2 pureScale cluster, all members must access these tables frequently so that contention due to locking can occur. Contention can result in reduced system performance. To avoid contention and the resulting reduced system performance, you can partition the update tables VBHDR, VBMOD, and VBDATA so that the SAP application server operates on its own data partition for these tables.

Monitoring Enhancements in the DBA Cockpit

The DBA Cockpit (SAP transaction DBACOCKPIT) was enhanced with monitoring functions that are specific to DB2 pureScale. To use these new monitoring functions, you have to apply a special transport together with correction instructions as described in [SAP Note 1409540](#) .

For more information, see [Using DB2 pureScale Monitoring Enhancements in the DBA Cockpit \[page 78\]](#).

2 Planning

2.1 Installation Restrictions

The first versions of DB2 pureScale are shipped with several restrictions. The following section describes only the most important ones, such as the following:

- Operating system-specific restrictions
- Database-specific restrictions
- SAP system-specific restrictions

Operating System-Specific Restrictions

- The installation of DB2 pureScale is **only** supported on AIX and Linux.
- DB2 pureScale **cannot** run across AIX workload partitions (WPARs).

Database-Specific Restrictions

The following restrictions apply to the DB2 10.1 pureScale release:

- DB2 pureScale supports only single-node databases. With DB2 pureScale, you cannot use the DB2 database partitioning feature (DPF).
- Multidimensional clustering (MDC) and insert time clustering tables (ITC tables) cannot be used.
- Inplace table reorganization and online index reorganization are not supported. Instead, you can use the stored procedure `ADMIN_MOVE_TABLE`.
- Incremental backups and tablespace backups are not supported. Instead, you can perform full database backups.
- You can apply a backup only to a system with the same cluster topology. For example, a backup that was taken on a DB2 pureScale cluster with two CFs and five members can only be applied to another DB2 pureScale cluster with two CFs and five members.
- Only one DB2 pureScale instance can run on each host.
- DB2 pureScale does not support high availability disaster recovery (HADR). Continuous availability is already provided by DB2 pureScale. To set up a disaster recovery site, you can use storage replication.

SAP System-Specific Restrictions

- DB2 pureScale is **only** supported for SAP systems based on at least **SAP NetWeaver 7.0 SR3**.
- The host names of all SAP application servers must be different within the first 8 characters. This is a requirement for the partitioning of the update tables, which is described later in this guide.
- The installation of multiple components in one database (MCOD) is **not** supported with DB2 pureScale.
- The AS ABAP of the SAP system must use the new DB2 client connectivity.
In this scenario, the [IBM Data Server Driver for CLI and ODBC](#) (CLI Driver) is used as DB2 client software. The old client connectivity (that is, the DB2 Runtime Client) is **not** supported.

2.2 Installation Requirements

Make sure that the following requirements are met when you plan the installation or the upgrade to DB2 pureScale:

- Hardware and operating system requirements
- SAP system-specific requirements
- Space requirements
- Other requirements

Hardware and Operating System Requirements

For more information about the required hardware and software requirements, see the following information:

- In the IBM DB2 10.1 Information Center, see:

Section	URL
<i>Installation prerequisites for DB2 pureScale Feature (AIX)</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/r0054850.html ➡
<i>Pre-installation checklist for DB2 pureScale Feature (AIX)</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/r0056077.html ➡
<i>Installation prerequisites for DB2 pureScale Feature (Linux)</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/r0057441.html ➡
<i>Pre-installation checklist for DB2 pureScale Feature (Linux)</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/r0057204.html ➡

- [Preparing the Cluster for the Conversion \[page 20\]](#)
- [Preparing the Cluster for the Installation \[page 47\]](#)

SAP System Requirements

- Make sure that you have applied the correct SAP kernel patch.
For SAP systems based on SAP NetWeaver 7.0 SR3 and higher, a specific DBSL (`dbdb6slib`) version is required. The following table shows the minimal DBSL version that is required for your SAP system release:

SAP Kernel Release	DBSL Patch Level
7.00	206
7.01	42
7.10	152
7.11	38

For SAP releases higher than the ones listed in this table, you do not need a specific SAP kernel patch.

→ Recommendation

We strongly recommend that you apply the latest available kernel patch to your system.

To check the current patch level of your DBSL, proceed as follows:

- Log on to an application server as user `<sapsid>adm`.
- Enter the following command:

`disp+work -v`

You can find the DBSL patch information at the end of the output. For more information about how to download and apply the latest SAP kernel patch, see [SAP Note 19466](#).

- To be able to use the new functionality of DB2 10.1 with the SAP NetWeaver Application Server ABAP, you need a minimal SAP_BASIS support package level. These support packages contain adaptations in the area of monitoring (SAP transaction `DBACOCKPIT`) and in the ABAP data dictionary. The following table lists the minimal required SAP_BASIS support packages for the different SAP releases:

SAP Basis Release	SAP_BASIS Support Package
7.00	SP16
7.01	SP5
7.02	SP9
7.10	SP9
7.11	SP4
7.20	n/a
7.30	SP5


SAP Basis Release	SAP_BASIS Support Package
7.31	SP2

- In addition to the support packages mentioned above, you must implement the following SAP Notes:

SAP Note	Description
1354186	DB6: LONG/LOB type mapping and database object check
1701181	DB6: ABAP DDIC: Enhancements for DB2 10.1
1409540	DB6: CCMS Adoptions for DB2 pureScale Feature
1677675	DB6: Compatibility Patches for new DB2 Releases
1721095	DB6: DB2 10.1 pureScale Adoptions

- The following SAP tools for DB2 require specific patch levels to be able to work with DB2 pureScale:

Tool	Patch Level
db6pmudf	19
brdb6brt	27

For more information about how to download and apply the latest patch for these tools, see **SAP Note 19466** . The dmdb6bcp tool is not relevant for DB2 pureScale.

Space Requirements

- For the installation of the DB2 10.1 software (including GPFS and SA MP), you need 4 GB on the local file system of each participating host.
- For the shared disk size, see the following sections for details about the required disk space:
 - [Preparing the Cluster for the Conversion \[page 20\]](#)
 - [Preparing the Cluster for the Installation \[page 47\]](#)
- For the remote installation directory of the SAP installation tool, you need approx. 500 MB of free disk space.

Other Requirements

For the latest additions and corrections to this document, see **SAP Note 1718559** .

3 Conversion of an Existing SAP System to a System Running on DB2 pureScale


3.1 Introduction

This section describes how you can convert an existing SAP system that does not use the DB2 pureScale Feature to a system running on DB2 pureScale.

Converting a non-pureScale instance to a DB2 pureScale instance includes the following steps:

- Adding the DB2 pureScale Feature to the existing DB2 software installation
- Creating GPFS file systems
- Moving existing data to the new file systems
- Converting the DB2 10.1 ESE instance to DB2 pureScale
- Adding more members and CFs to the DB2 pureScale cluster

i Note

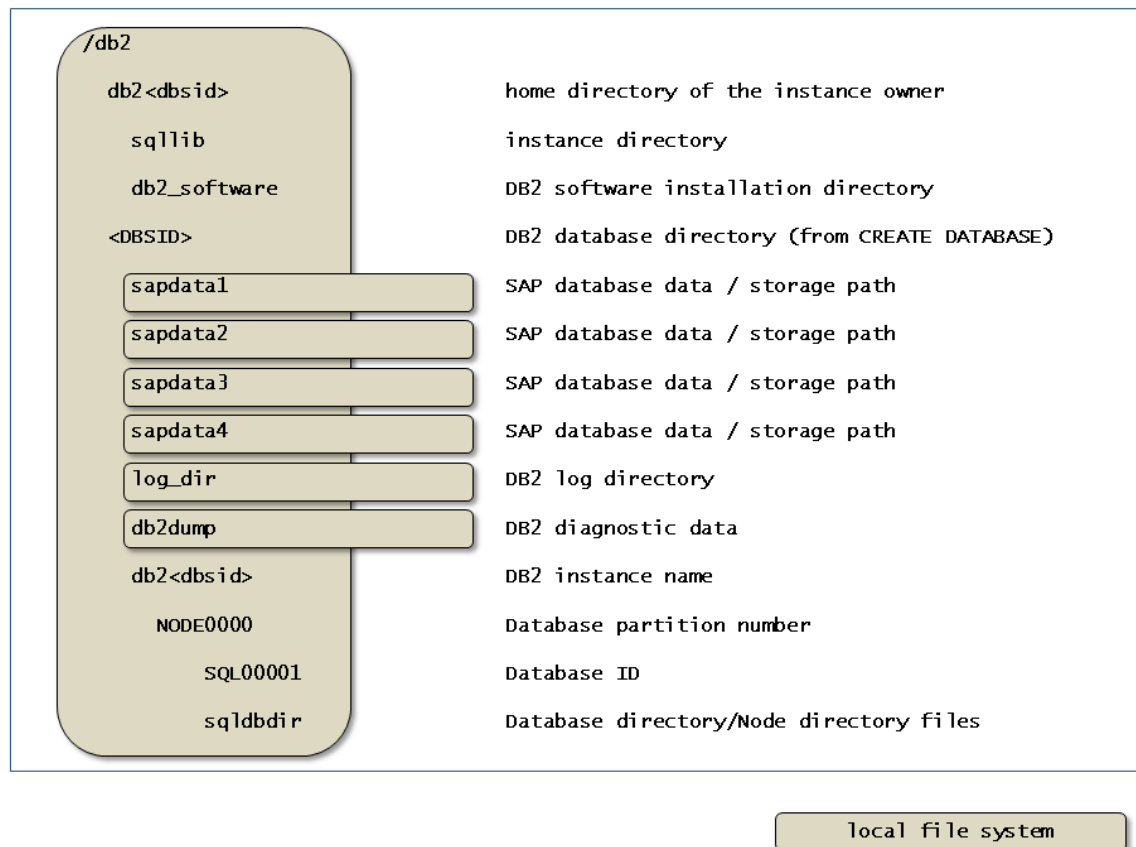
During the procedure described in the following sections, an existing autostorage or non-autostorage database is converted into an **autostorage database with only one storage path**. By default, DB2 pureScale uses parallel I/O by setting the DB2 environment variable `DB2_PARALLEL_IO` to `"*"` (see `DB2_PARALLEL_IO` in *System environment variables* in the IBM DB2 10.1 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.regvars.doc/doc/r0005658.html> ). In this guide, we therefore use only one storage path per automatic storage database. At the time of writing this document, there was no sound evidence that providing more than one storage path or file system for DB2 data or log files results in better overall system performance.

There are various options to move the existing data to the new GPFS file systems. In this document, we describe how you can move data by rebalancing tablespaces. The advantage of tablespace rebalancing is that it requires the least system downtime. This means that during the tablespace rebalancing process, the system can stay online, but the tablespace rebalancing operations require additional CPU and disk capacity. Furthermore, this procedure also allows you to convert a non-automatic storage database to a database that uses automatic storage.

Alternatively, you can move data offline using operating system means (for example, by copying data using the `cp` command) or a redirected restore. The redirected restore requires longer system downtime but is usually easier to perform than tablespace rebalancing.

3.2 Overview of a DB2 File System Before and After the Conversion to DB2 pureScale

The following figure shows a typical DB2 directory hierarchy after the installation of SAP NetWeaver 7.0 SR3 (or higher):



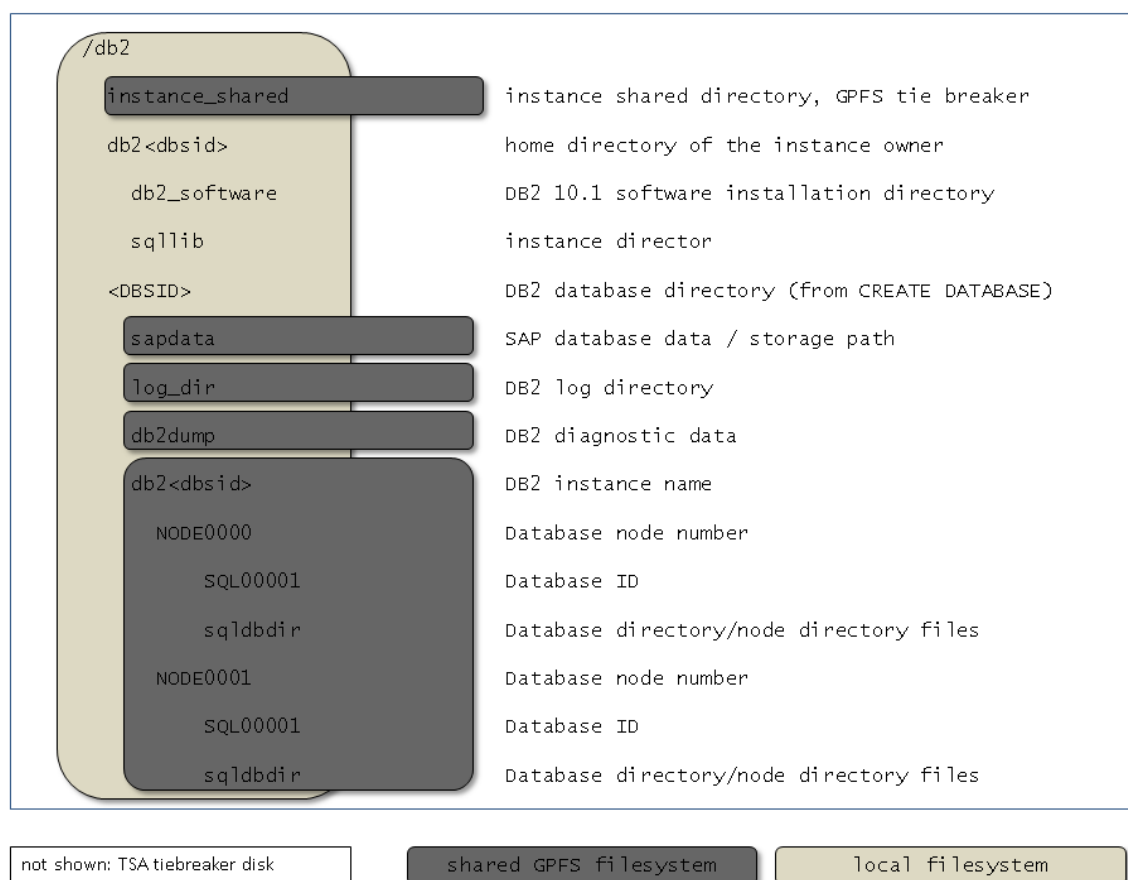
DB2 Directory Hierarchy after SAP System Installation

If you installed your SAP system as described in the SAP NetWeaver installation guide, you created separate file systems for the DB2 database, the **db2dump** directory (that contains DB2 diagnostic data) and the logging directory. All file systems reside on local disks. In the following, we assume that this is the case in your actual system configuration.

i Note

If you did **not** install your SAP system using separate file systems, you can still use this document. However, you have to adapt some of the steps outlined later on. For example, if you are asked to unmount the **sapdata** file systems, you have to rename the corresponding directories instead.

During the conversion procedure, the DB2 file system is converted as shown in the following figure:



DB2 Directory Hierarchy after the Upgrade to DB2 10.1 with pureScale

Only the home directory of the instance owner and the DB2 10.1 software installation must stay on the local file system. Everything else is moved to several shared file systems. A new extra GPFS file system is used for the instance-shared directory. This directory holds files that are relevant to all installed instances. It also hosts files that must be available to all members and CFs at all times – for example, the `db2nodes.cfg` file – and it serves as a tie breaker for the GPFS cluster.

Another GPFS file system holds all DB2 data except the logs. There is only one `sapdata` directory and, therefore, only one storage path due to the reasons outlined earlier in this section. We recommend that you use an extra GPFS file system only for DB2 logging. This file system contains the log streams of all cluster members. Two more file systems are needed for the diagnostic data (`db2dump` directory) and the database directory. The database directory contains configuration files, the history file, and so on.

i Note

In theory, you can limit the number of GPFS file systems by using nested mount points. For example, you can create a large GPFS and mount it under `/db2/<DBSID>`. Inside this file system, you can mount another GPFS for the log streams under `/db2/<DBSID>/log_dir`. This way, you do not need additional file systems for the diagnostic data and the database directory.

However, GPFS mounts its file systems in random order. If a mount point does not exist, it is automatically created. For the example shown in the figure above, this means that on system start or after a failover, it is possible that GPFS first mounts `/db2/<DBSID>/log_dir` and afterwards `/db2/<DBSID>`. In this

situation, the last mount effectively hides the file system that was mounted first, which leads to an unusable cluster.

3.3 Conversion to DB2 pureScale – Step by Step

3.3.1 Introduction

The database conversion to DB2 pureScale consists of the following steps that you **must** perform in the **specified** order:

1. [Upgrade your Database to DB2 10.1 \[page 17\]](#).
2. [Apply the Latest Kernel Patch \[page 18\]](#).
3. [Perform an offline backup \[page 18\]](#).
4. [Set the required configuration parameters \[page 18\]](#).
5. [Add the DB2 pureScale Feature to the existing DB2 10.1 software \[page 19\]](#).
6. [Prepare the cluster for the upgrade \[page 20\]](#).
7. [Manually set up passwordless access for user root \[page 23\]](#).
8. [Run the db2checkSD utility \[page 24\]](#).
9. [Prepare the GPFS cluster \[page 25\]](#).
10. [Create the necessary GPFS file systems \[page 27\]](#).
11. [Rebalance tablespaces to the new file system \[page 28\]](#).
12. [Move remaining data to the GFPSSs \[page 31\]](#).
13. [Mount GFPSSs under new mount points \[page 32\]](#).
14. [Convert the DB2 10.1 ESE instance to a DB2 pureScale instance \[page 35\]](#).
15. [Adapt the JDBC URL \(AS Java only\) \[page 39\]](#).
16. [Test the DB2 pureScale installation \[page 40\]](#).
17. [Add members and CFs \[page 41\]](#)
18. [Run the db6_update_db script \[page 42\]](#).
19. [Check configuration settings \[page 43\]](#).
20. [Install the DB2 pureScale license \[page 43\]](#).

3.3.2 Upgrading the Database to DB2 10.1

Use

The procedures in the following sections assume that your existing SAP system is running on DB2 10.1. If this is not the case, first upgrade your system to DB2 10.1.

Procedure

For more information, see the upgrade guide *Upgrading to Version 10.1 of IBM DB2 for Linux, UNIX, and Windows* at https://help.sap.com/viewer/db6_upgrade_10_1.

3.3.3 Applying the Latest Kernel Patch

Use

Make sure that you apply the latest available kernel patch for your SAP system before you continue.

Procedure

For more information about how to download and install a kernel patch, see **SAP Note 19466** .

3.3.4 Performing an Offline Backup

Before you start any conversion tasks on your existing systems, we strongly recommend that you perform a full database offline backup.

3.3.5 Setting Required Configuration Parameters

Use

DB2 with pureScale does not support the health monitor. Therefore, you have to set the database manager configuration parameter `HEALTH_MON` to `OFF`. Furthermore, DB2 pureScale does not support incremental backups. Therefore, you must set the database configuration parameter `TRACKMOD` to `OFF`.

Procedure

1. Log on to the database server as user `db2<dbsid>`.
2. Update the database manager configuration using the following commands:
`db2 "update dbm cfg using HEALTH_MON OFF"`
3. Update the database configuration using the following command:
`db2 "update db cfg using TRACKMOD OFF"`

4. To activate the changes, stop and restart the database manager.
5. To check that the parameters are set correctly, enter the following commands:

```
db2 "get dbm cfg" | grep HEALTH_MON
db2 "get db cfg" | grep TRACKMOD
```

3.3.6 Adding the DB2 pureScale Feature to the Existing DB2 10.1 Software Installation

Use

By default, your SAP system uses a DB2 software installation that is suitable for DB2 ESE instances. To create a DB2 pureScale instance, the DB2 pureScale Feature must be added to the existing DB2 software installation. During the conversion to DB2 pureScale, the `db2iupdt` command automatically installs the DB2 software on the other hosts in the cluster.

Procedure

1. Stop the SAP system and the DB2 database manager (`db2stop`).

i Note

Make sure that the database manager is really not running. If you perform the following procedure while the database manager is still running, the update of the database instance cannot complete successfully. As a consequence, you need to perform manual steps to bring the database instance back to a normal state.

2. Log on to the database server as user `root`.
3. Insert and mount the database DVD to `<DVD_mount>`.
4. Change to the following directory:

```
cd <DVD_mount>/platform/ESE/disk1
```
5. Run the installation prerequisite check using the following command:

```
./db2prereqcheck
```

The output of the command should contain a message such as the following:

```
DBT3533I The db2prereqcheck utility has confirmed that all installation prerequisites were met for DB2 database server with DB2 pureCluster feature. Version: "10.1.0.0".
```

If the check returns errors, you need to perform the necessary actions before you can continue with the installation of the DB2 10.1 software (for example, upgrading the maintenance level of AIX or applying an AIX service pack).
6. Ensure that your X Windows display is set up correctly and call the graphical installation tool `db2setup`.
7. On the *Welcome* screen, choose *Install a Product*.
8. On the next screen, choose *Work with Existing* under *DB2 Enterprise Server Edition 10.1 with the IBM DB2 pureScale Feature*.

⚠ Caution

Make sure that you choose *DB2 Enterprise Server Edition 10.1 with the IBM DB2 pureScale Feature*, **not** *DB2 Advanced Enterprise Server Edition 10.1 with the IBM DB2 pureScale Feature*

9. In the dialog box *Select the DB2 copy to work with*, select the DB2 10.1 database software installation that is used by your SAP system and choose *Launch DB2 Setup wizard*.
10. On the *Introduction* screen, choose *Next*.
11. On the *Installation action* screen, choose *Install DB2 Enterprise Server Edition with the IBM DB2 pureScale Feature* and choose *Next*.
12. Confirm the installation directory with *Next*.
13. On the *Languages* screen, choose *Next*.
14. On the *Instance setup* screen, select *Do not create a DB2 instance*.
15. On the *Host list* screen, confirm the host list with *Next*.

⚠ Caution

Do not add any members or cluster caching facilities (CFs). The host list must contain one entry. Note that, with this step, your DB2 instance is **not** automatically converted to DB2 pureScale.

16. On the *Summary* screen, choose *Next* and wait until the installation finishes.
17. Confirm the *Setup complete* dialog box with *Finish*.
18. Restart the DB2 instance and the SAP system.

3.3.7 Preparing the Cluster for the Conversion

3.3.7.1 Introduction

To prepare the cluster, you perform the following steps:

1. [Prepare the general parallel file system \[page 20\]](#).
2. [Configure IO completion ports \[page 22\]](#).
3. [Identify the cluster interconnect name \[page 23\]](#).
4. [Set up identical users on all nodes \[page 23\]](#).

For more information, see *Preparing to install the DB2 pureScale Feature for DB2 Enterprise Server Edition* in the IBM DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/c0060061.html> ➦.

3.3.7.2 Preparing the General Parallel File System (GPFS)

For a DB2 pureScale cluster, you need at least six physical or virtual shared disks. This corresponds to the five required GPFS file systems that are mentioned in [Overview of a DB2 File System Before and After the Conversion to DB2 pureScale \[page 15\]](#) plus an additional disk for the TSA tiebreaker.

It is important that you check that these six disks are actually shared on all cluster nodes. To do so, log on to each cluster node as user **root** and run the **lspv** command. The following shows an example of the output from the **lspv** command:

```
root@db2dsf1 / > lspv
hdisk0      00c5cc1479a245dd  rootvg      active
hdisk1      00c5ccf479278ca7    None
hdisk2      00c5ccf42dae2dbe  None
hdisk3      00c5ccf478d1e31e  None
hdisk4      00c5ccf478d1e49a  None
hdisk5      00c5ccf4d5b68397  None
hdisk6      00c5ccf4d5b68539  None
root@db2dsf2 / > lspv
hdisk0      00c5cc1479a245dd  rootvg      active
hdisk1      00c5ccf479278ca7    None
hdisk2      00c5ccf42dae2dbe  None
hdisk3      00c5ccf478d1e31e  None
hdisk4      00c5ccf478d1e49a  None
hdisk5      00c5ccf4d5b68397  None
hdisk6      00c5ccf4d5b68539  None
```

The shared disks must appear with the same device number on every node in the cluster. The device number is reported in the second column of the **lspv** output. In the example above, the devices with the IDs **00c5ccf479278ca7**, **00c5ccf42dae2dbe**, **00c5ccf478d1e31e**, **00c5ccf478d1e49a**, **00c5ccf4d5b68397**, and **00c5ccf4d5b68539** are shared on both nodes. Depending on the system setup, the device names (hdisk1, hdisk2,...,hdisk6 as shown in the example) for the shared disks can differ on the nodes.

i Note

If you want to determine the size of a disk in MB on AIX, you can use the following command:

```
getconf DISK_SIZE /dev/<disk_name>
```

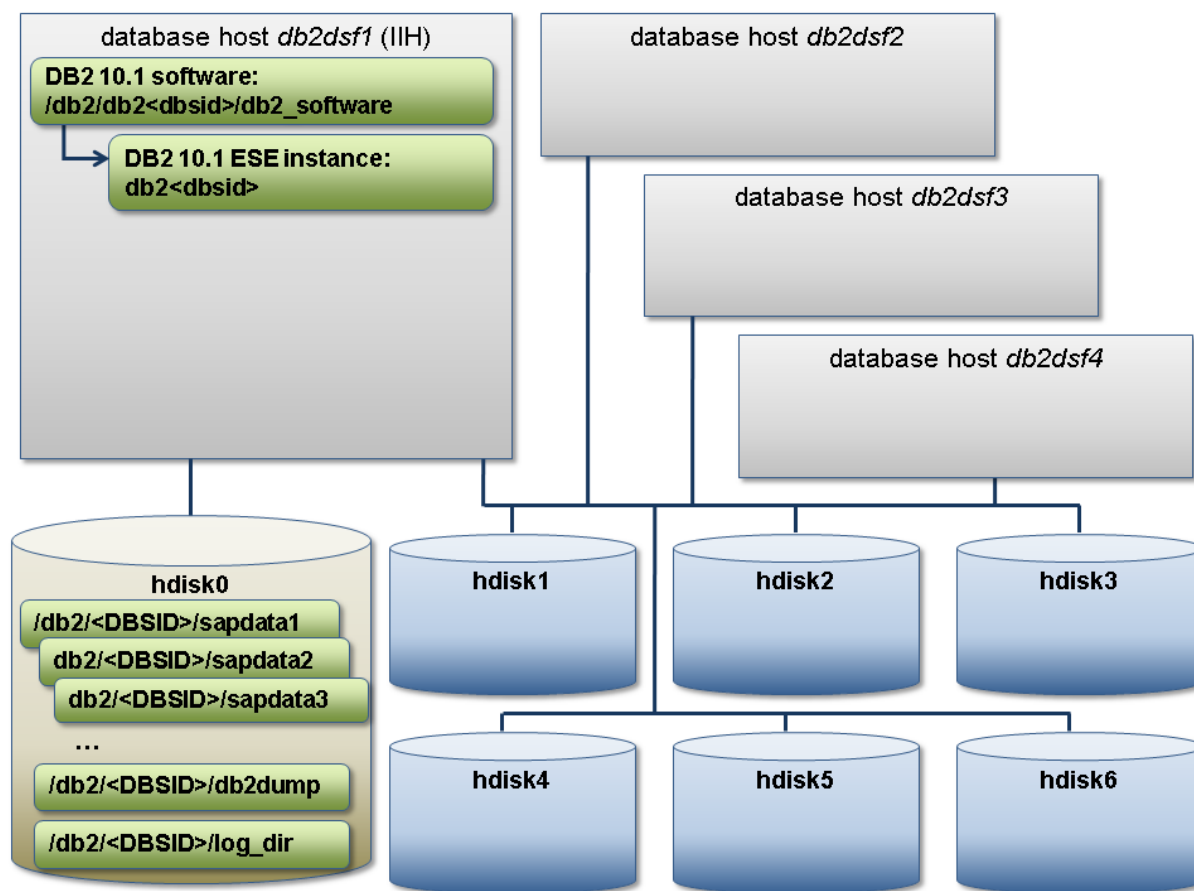
For example:

```
getconf DISK_SIZE /dev/hdisk2
```

For more information, see *Configuring PVIDs for a DB2 pureScale instance (AIX)* in the IBM DB2 Information Center at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/t0056777.html> 📄

The following figure shows a typical configuration with an installed SAP system and the shared disks in place:



DB2 Configuration in an SAP Environment

The DB2 10.1 software is already installed and three additional hosts are available for the DB2 pureScale cluster. The host where the DB2 installation starts (in this case **db2dsf1**) is called the **install-initiating host (IIH)**.

3.3.7.3 Configuring I/O Completion Ports (IOCP)

Use

On AIX, you must ensure that I/O completion ports (IOCPs) are used.

To do so, configure IOCPs as described in the IBM DB2 Information Center *Configuring IOCP (AIX)* at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.perf.doc/doc/t0054518.html> .

3.3.7.4 Identifying the Cluster Interconnect Netname

Use

For some upgrade tasks, you have to know the correct cluster interconnect netname of the cluster hosts. To find out this netname, run the following command as user `root` on each host:

`<inst_dir>/instance/native/install/db2getnetwork`

❖ Example

The following shows an example of the output from this command:

```
db2dsf1,22.22.22.22,db2dsf1-ib0,ib0
```

```
db2dsf1,33.33.33.33,db2dsf1.wdf.sap.corp,en0
```

In this example, `db2dsf1-ib0` is the cluster interconnect netname of `db2dsf1`.

3.3.7.5 Setting Up Identical Users on all Nodes

Use

The following users must be available on all cluster nodes:

- Database administration user `db2<dbsid>`
- DB2 connect user `sap<sapsid>` (AS ABAP) and `sap<sapsid>db` (AS Java)
- SAP system administration user `<sapsid>adm`

These users must belong to the same groups, the user and group IDs must match, and they must use the same shell and the same path to the home directory.

You can use a central user management for user authorization and authentication, for example, by using the Lightweight Directory Access Protocol (LDAP) or the Network Information Service (NIS). Alternatively, you can also use local user authentication and authorization. In this case, the SAP installer makes sure that all users are created with the correct properties on all members and CFs.

3.3.8 Manually Setting Up Passwordless Access for User Root

Use

In a DB2 pureScale cluster, user `root` must be able to log on to all cluster nodes without a password using secure shell (open `ssh`). The following procedure describes the steps required to configure passwordless access for user `root` between two cluster nodes.

Procedure

1. Log on to the first node of the cluster as user `root`.
2. Generate an RSA key for user `root` using the following command:
`ssh-keygen -t dsa`
3. Specify the file name for the key and a passphrase by entering the default file name for the key without any passphrase:
`~/.ssh/id_dsa`
4. Enable key authentication for `ssh` on the same node using the following command:
`cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys`
5. Transfer the public key to the second node using the following command:
`scp ~/.ssh/id_dsa.pub root@<hostname>:~/id_dsa.pub`
6. Log on to the second node of the cluster as user `root`.
7. Enable key authentication for `ssh` on the second node using the following command:
`cat ~/id_dsa.pub >> ~/.ssh/authorized_keys`
8. Remove the public key file from the home of user `root` on the second node using the following command:
`rm ~/id_dsa.pub`
9. Add both cluster nodes to the list of known hosts using the following command:
**`ssh-keyscan -t dsa <host1_short_name>, <host1_full_name>,
ssh-keyscan -t dsa <host1_short_name>, <host1_full_name>,
<host1_ip> >> ~/.ssh/known_hosts
ssh-keyscan -t dsa <host2_short_name>, <host2_full_name> ,`**
10. Set up passwordless access for user `root` on the second node:
To do so, repeat this procedure, but starting from the second node and making sure that you reverse the nodes.

More Information

Installing and setting up OpenSSH in the IBM DB2 Information Center at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/t0055342.html> ➡

3.3.9 Running the db2checkSD Utility

Use

You use the `db2checkSD` utility to check if the prerequisites for the conversion to DB2 pureScale are met.

Procedure

1. Log on to the database server as user **db2<dbsid>**.
2. Since the GPFS file system has not yet been set up, pass the **-nogpfs** option to the **db2checkSD** utility by running it as follows:

```
<inst_dir>/bin/db2checkSD <DBSID> -l <name of logfile> -nogpfs
```

The output of the **db2checkSD** utility should look as follows:

```
DBT5000I The db2checkSD utility completed successfully. The specified database  
can be upgraded to a data-sharing environment. The output log file is named  
"/tmp/checksd.log".
```

If your database is not yet enabled for automatic storage management, the **db2checkSD** utility reports an error. You can ignore this error message because during the upgrade to DB2 pureScale, a non-autostorage database is converted to an automatic storage database.

If the **db2checkSD** utility reports other errors, you need to review the log file and take appropriate actions. Re-run the **db2checkSD** utility until all problems have been solved.

3.3.10 Preparing the GPFS Cluster

Use

Before you can convert the DB2 10.1 ESE instance to DB2 pureScale instance, you have to move the database to a GPFS file system. This file system can be managed either manually using GPFS-specific commands (user-managed file system) or automatically by DB2 (DB2-managed file system).

→ Recommendation

We strongly recommend that you use the DB2-managed file system approach. This way, DB2 performs all necessary GPFS configuration tasks, for example, while new cluster members are added.

To make sure that DB2 pureScale automatically manages the GPFS file system, you need to run the **db2cluster_prepare** command. Among other things, the GPFS file system is then created for the instance-shared directory and certain entries in the DB2 global registry are made.

Procedure

1. Log on to the database server as user **root**.
2. Run the **db2cluster_prepare** command as follows:

```
<inst_dir>/instance/db2cluster_prepare -l <name of logfile> -instance_shared_dev  
<shared disk> -instance_shared_mount /db2/instance_shared
```

❖ Example

```
/db2/db2<dbsid>/db2_software/instance/db2cluster_prepare -l /tmp/  
clusterprep.log -instance_shared_dev /dev/hdisk3 -instance_shared_mount /db2/  
instance_shared
```

If the **db2cluster_prepare** command reports errors, you need to review the log file and take appropriate actions.

3. To check that the instance-shared file system was created, you can use the **db2cluster** command as in the following example:

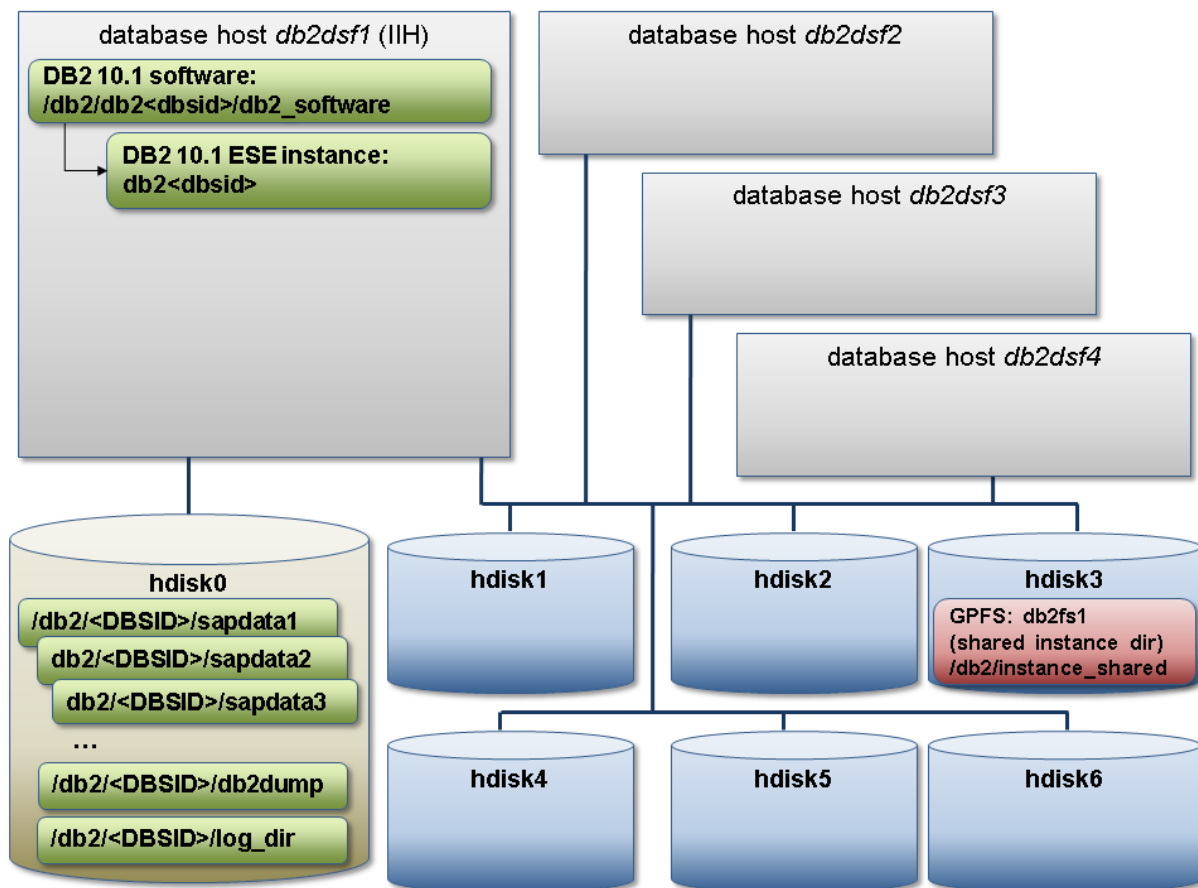
```
<inst_dir>/bin/db2cluster -cfs -list -filesystem
```

An example of the output of this command looks as follows:

FILE SYSTEM NAME	MOUNT_POINT
db2fs1	/db2/instance_shared

Result

At this stage of the overall upgrade process, your system looks as shown in the following figure:



Status of Cluster After First Preparation Steps

3.3.11 Creating the GPFS File Systems

Use

In contrast to a standard SAP system installation, only one file system for the SAP database is used in this cluster scenario.

i Note

You can extend this file system by adding more disks to it.

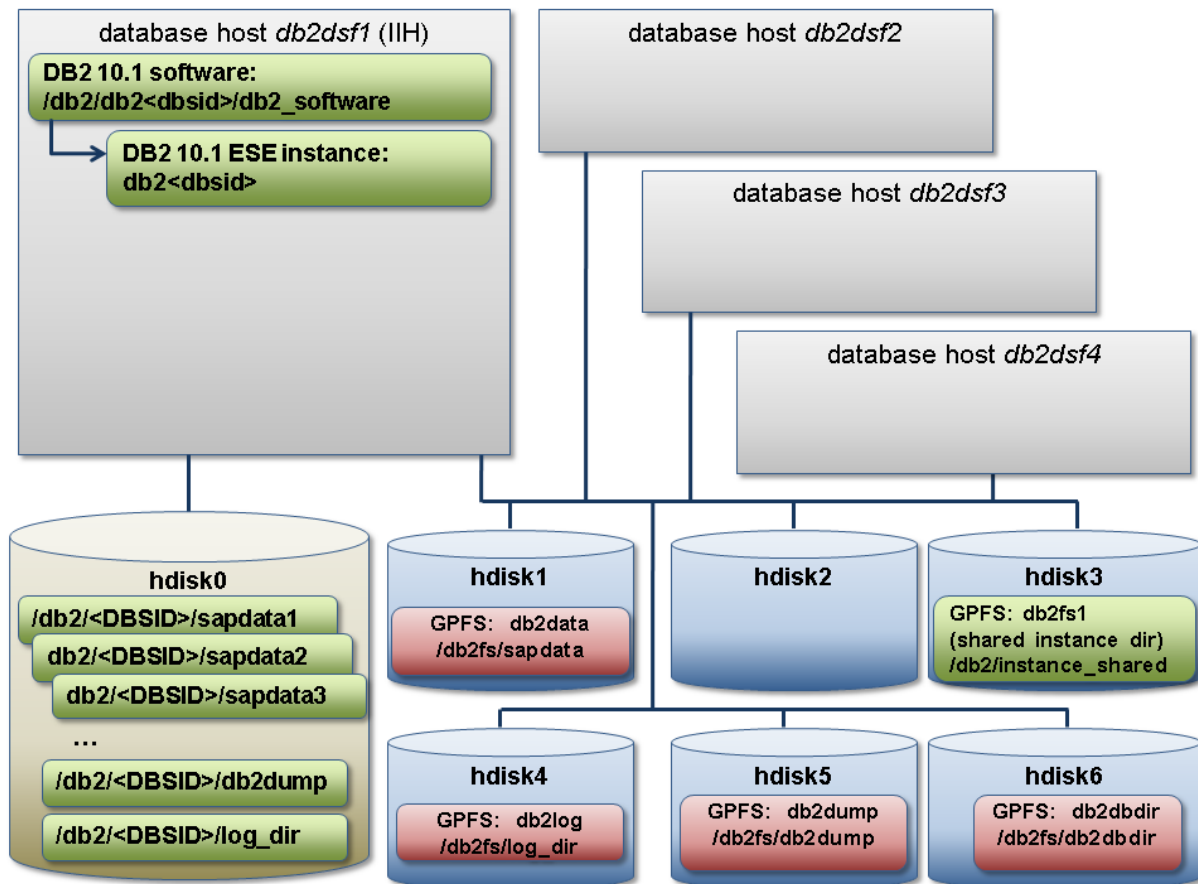
Procedure

1. Log on to the database server as user **root**.
2. Create the GPFS file system for the SAP database using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2data -disk <sapdata_disk>
where **<sapdata_disk>** is the disk for the file system that is going to be used by the SAP database, for example, **/dev/hdisk1**.
3. Create the GPFS file system for the DB2 log files using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2log -disk <log_disk>
where **<log_disk>** is the disk for the file system that is used by the DB2 log files, for example, **/dev/hdisk4**.
4. Create the GPFS file system for the DB2 diagnostic data using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2dump -disk <dump_disk>
where **<dump_disk>** is the disk for the file system that is used by the DB2 diagnostic data, for example, **/dev/hdisk5**.
5. Create the GPFS for the DB2 database directory using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2dbdir -disk <dbdir_disk>
where **<dbdir_disk>** is the disk for the file system that is going to be used by the DB2 database directory, for example **/dev/hdisk6**.
6. Since the newly created file systems must be accessible by the instance owner, change their permissions as follows:
chown db2<dbsid>:db<dbsid>adm /db2fs/db2data
chown db2<dbsid>:db<dbsid>adm /db2fs/db2log
chown db2<dbsid>:db<dbsid>adm /db2fs/db2dump
chown db2<dbsid>:db<dbsid>adm /db2fs/db2dbdir
7. Check that the required file systems are created by running the following command:
<inst_dir>/bin/db2cluster -cfs -list -filesystem
The output should look similar to the following example:

FILE SYSTEM NAME	MOUNT_POINT
db2data	/db2fs/db2data
db2dbdir	/db2fs/db2dbdir
db2dump	/db2fs/db2dump
db2fs1	/db2/instance_shared
db2log	/db2fs/db2log

Result

After the creation of the GPFS file systems, your system looks as shown in the following figure:



System Status after Creation of GPFSs

3.3.12 Rebalancing Tablespaces to the New File System

Use

You use the following online procedure to move the DB2 database to the GPFS file systems. In addition, you can use this procedure to enable a non-autostorage database for automatic storage management.

i Note

Tablespace rebalancing affects the performance of the system. Make sure that you perform the following steps **only** during low system activity.

Procedure

1. Log on to the database server as user `db2<dbsid>`.
2. List the existing storage paths with the following command:

db2pd -d <DBSID> -storagepaths

The following shows an example of the output from this command:

```
Database Partition 0 -- Database DSJ -- Active --
Up 4 days 03:08:31
Database Storage Paths:
Number of Storage Paths          4
Address          PathID PathState PathName
0x0700000020754F20 0      InUse   /db2/DSJ/sapdata1
0x0700000020755200 1      InUse   /db2/DSJ/sapdata2
0x07000000207554E0 2      InUse   /db2/DSJ/sapdata3
0x07000000207557C0 3      InUse   /db2/DSJ/sapdata4
```

If there are no storage paths listed (for example, the number of storage paths is 0), your database is not enabled for automatic storage.

3. If your database is enabled for **automatic storage management**, perform the following steps:
 1. Add the storage path of the new GPFS db2data file system and drop the existing storage paths using the following command:

db2 "ALTER DATABASE DROP STORAGE ON '<old_path_1>', '<old_path_2>', ... ADD STORAGE ON '<new_path>'"

❖ Example

db2 "ALTER DATABASE DROP STORAGE ON '/db2/DSJ/sapdata1', '/db2/DSJ/sapdata2', '/db2/DSJ/sapdata3', '/db2/DSJ/sapdata4' ADD STORAGE ON '/db2fs/db2data'"

DB2 issues the following warning that you can ignore:

SQL2095W Storage path "/db2/DSJ/sapdata1" is in the drop pending state because one or more automatic storage table spaces reside on the path.
SQLSTATE=01691

2. Run the following command again:

db2pd -d <DBSID> -storagepaths

The output now looks as follows:

```
Database Partition 0 -- Database DSJ -- Active -- Up 4 days 03:24:07
Database Storage Paths:
Number of Storage Paths          5
Address          PathID PathState PathName
0x0700000020754F20 0      DropPending /db2/DSJ/sapdata1
0x0700000020755200 1      DropPending /db2/DSJ/sapdata2
0x07000000207554E0 2      DropPending /db2/DSJ/sapdata3
0x07000000207557C0 3      DropPending /db2/DSJ/sapdata4
0x0700000037F3EE60 4      NotInUse   /db2fs/db2data
```

4. If the database is **not** enabled for automatic storage management, perform the following steps:
 1. Enable your database for automatic storage by adding an automatic storage path with the following command

db2 "ALTER DATABASE <DBSID> ADD STORAGE ON '/db2fs/db2data'"

2. Convert all DMS tablespaces of the database to automatic storage as follows:

- To convert a single tablespace, enter the following command:
db2 "ALTER TABLESPACE <tablespace_name> MANAGED BY AUTOMATIC STORAGE"
- To generate a script that contains the statements required to rebalance **all** tablespaces, use the following SQL statement
db2 -x "select 'alter tablespace ' || CHR(34) || TBSP_NAME || CHR(34) || ' managed by automatic storage;' from SYSIBMADM.SNAPTbsp where TBSP_USING_AUTO_STORAGE != 1 and TBSP_TYPE = 'DMS' and TBSP_CONTENT_TYPE in ('ANY', 'LARGE')" > /tmp/convert_tbs.sql
- To execute the generated script, use the following command
db2 -tvf /tmp/convert_tbs.sql

5. Rebalance all tablespaces (**except** the temporary tablespaces) by doing one of the following:

- Rebalance a **single** tablespace by entering the following command:
ALTER TABLESPACE <tablespace_name> REBALANCE
- Generate a script that contains the statements required to rebalance **all** tablespaces using the following SQL statement:

```
db2 -x "select 'ALTER TABLESPACE ' || CHR(34) || tbspace || CHR(34) || ' REBALANCE;' from syscat.tablespaces where TBSPACETYPE = 'D'" >/tmp/tbsprebalance.sql
```

To execute the generated script, use the following command:

```
db2 -tvf /tmp/tbsprebalance.sql
```

Tablespace rebalancing is an asynchronous action. You can monitor the rebalancing process using the following command:

```
db2 "LIST UTILITIES SHOW DETAIL"
```

6. Move the temporary tablespaces to the new storage path using the following commands:

```
db2 "rename tablespace PSAPTEMP16 to oldTEMP16"
db2 "rename tablespace SYSTOOLSTMPSPACE to oldTOOLSTMPSPACE"
db2 "create temporary tablespace PSAPTEMP16 in nodegroup IBMTEMPGROUP pagesize 16k extentsize 2 prefetchsize automatic no file system caching dropped table recovery off"
db2 "create user temporary tablespace SYSTOOLSTMPSPACE in nodegroup IBMCATGROUP pagesize 16k extentsize 2 prefetchsize automatic no file system caching dropped table recovery off"
db2 "drop tablespace oldTEMP16"
db2 "drop tablespace oldTOOLSTMPSPACE"
```

7. Check that the database was moved to the new storage paths using the following command:

```
db2pd -d <DBSID> -storagepath
```

The output now looks as follows:

```
Database Partition 0 -- Database DSJ -- Active -- Up 0 days 03:40:20
Storage Group Paths:
Address          SGID PathID PathState PathName
0x0700000060464400 0 4 InUse /db2fs/db2data
```

3.3.13 Moving Remaining Data to the GPFS

Use

You can move remaining data from the old file systems to the new shared GPFS using the **cp** command.

i Note

You **must** perform the following steps **offline**. Do **not** start the database again until you have [mounted GPFS files systems under new mount points \[page 32\]](#).

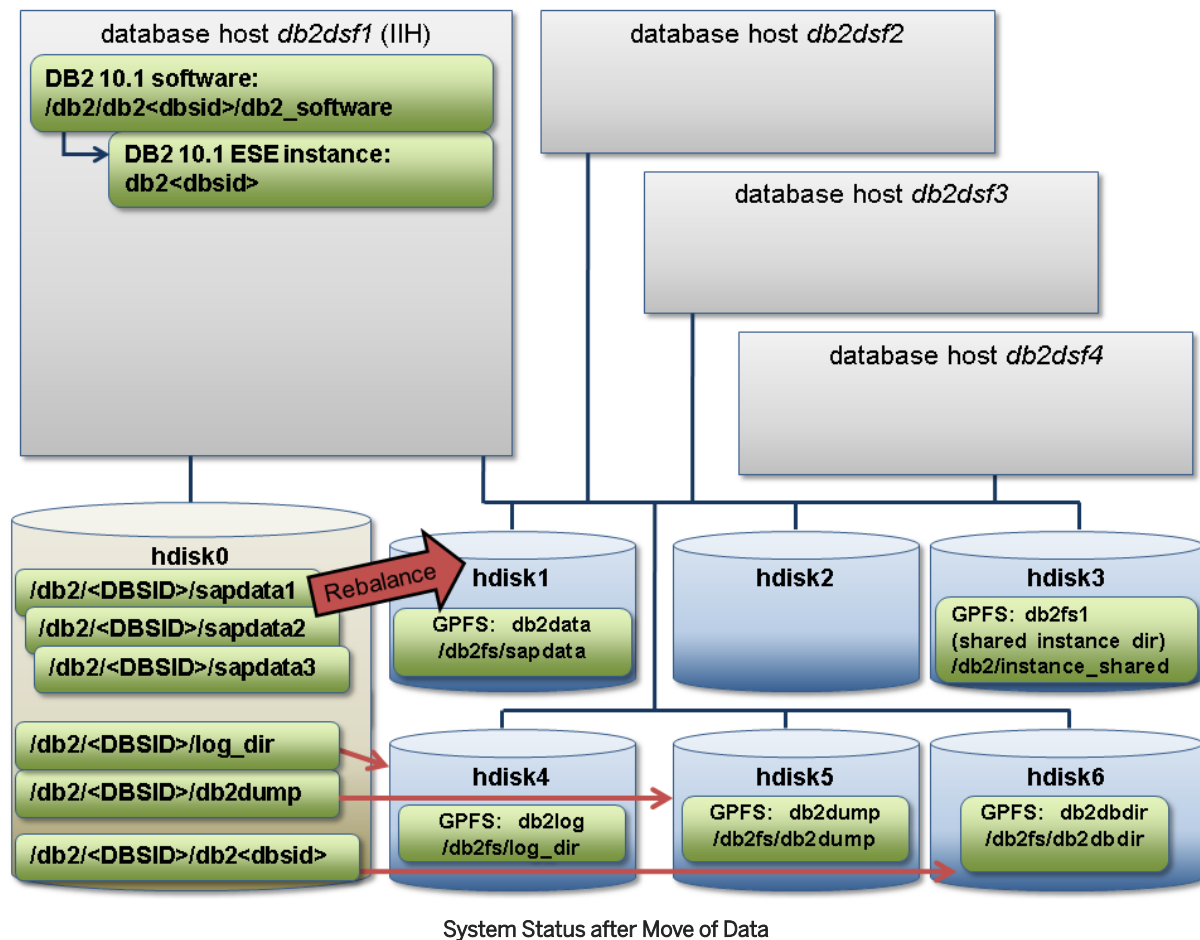
Procedure

1. Stop the SAP system and the database.
2. Log on to the database server as user **root**.
3. Copy all data inside the directories – except the directory that contains database data – to the new GPFS using the following commands:

```
cd /db2/<DBSID>/log_dir
cp -pR * /db2fs/db2log
cd /db2/<DBSID>/db2dump
cp -pR * /db2fs/db2dump
cd /db2/<DBSID>/db2<dbsid>
cp -pR * /db2fs/db2dbdir
```

Result

The move processes that are described in this section are depicted in the following figure:



3.3.14 Mounting GPFSs Under New Mount Points

Use

You use the following procedure to integrate the GPFS into the standard SAP/DB2 directory hierarchy.

Procedure

1. Stop the SAP system and the DB2 database.
2. Log on to the database server as user `root`.
3. Unmount the original SAP file systems using the following commands:

```
unmount /db2/<DBSID>/sapdata1
unmount /db2/<DBSID>/sapdata2
```

```

unmount /db2/<DBSID>/sapdata3
unmount /db2/<DBSID>/sapdata4
unmount /db2/<DBSID>/db2dump
unmount /db2/<DBSID>/log_dir

```

i Note

Make sure that these file systems are not remounted during a system restart. If you did not use separate file systems for `sapdata(n)`, `log_dir` and `db2dump`, rename these directories.

4. Create a new `sapdata` directory and change the permissions using the following commands:

```

mkdir /db2/<DBSID>/sapdata
chown db2<dbsid>:db<dbsid>adm /db2/<DBSID>/sapdata

```

5. Assign the correct permissions to the instance-shared directory using the following commands:

```

chown db2<dbsid>:db<dbsid>adm /db2/instance_shared

```

To change the mount point of a GPFS, you need to use GPFS commands. Before you change the mount point, run the following command to list all GPFS and their mount points:

```

/usr/lpp/mmfs/bin/mmlsfs all -T

```

The following shows an example of the output from this command:

```

File system attributes for /dev/db2data:
flag value      description
-----
-T /db2fs/db2data Default mount point
File system attributes for /dev/db2dbdir:
flag value      description
-----
-T /db2fs/db2dbdir Default mount point
File system attributes for /dev/db2dump:
flag value      description
-----
-T /db2fs/db2dump Default mount point
File system attributes for /dev/db2fs1:
flag value      description
-----
-T /db2/instance_shared Default mount point
File system attributes for /dev/db2log:
flag value      description
-----
-T /db2fs/db2log Default mount point

```

6. Unmount the GPFS for the DB2 database using the following command:

```

/usr/lpp/mmfs/bin/mmumount /db2fs/db2data

```

7. Change the mount point of the GPFS file system for the DB2 data using the following command:

```

/usr/lpp/mmfs/bin/mmchfs /dev/db2data -T /db2/<DBSID>/sapdata

```

8. Remount the GPFS file system for the DB2 database using the following command:

```

/usr/lpp/mmfs/bin/mmmount /dev/db2data

```

9. Change the mount points of the other file systems in the same way by using the following commands:

```

/usr/lpp/mmfs/bin/mmumount /db2fs/db2dbdir
/usr/lpp/mmfs/bin/mmchfs /dev/db2dbdir -T /db2/<DBSID>/db2<dbsid>
/usr/lpp/mmfs/bin/mmmount /dev/db2dbdir
/usr/lpp/mmfs/bin/mmumount /db2fs/db2dump
/usr/lpp/mmfs/bin/mmchfs /dev/db2dump -T /db2/<DBSID>/db2dump
/usr/lpp/mmfs/bin/mmmount /dev/db2dump
/usr/lpp/mmfs/bin/mmumount /db2fs/db2log
/usr/lpp/mmfs/bin/mmchfs /dev/db2log -T /db2/<DBSID>/log_dir
/usr/lpp/mmfs/bin/mmmount /dev/db2log

```

10. Check if the GPFS are available at the new location using the following command:

```
/usr/lpp/mmfs/bin/mmlsfs all -T
```

The following shows an example of the output from this command:

```
File system attributes for /dev/db2data:
flag value      description
-----
-T /db2/DSJ/sapdata Default mount point
File system attributes for /dev/db2dbdir:
flag value      description
-----
-T /db2/DSJ/db2dsj Default mount point
File system attributes for /dev/db2dump:
flag value      description
-----
-T /db2/DSJ/db2dump Default mount point
File system attributes for /dev/db2fs1:
flag value      description
-----
-T /db2/instance_shared Default mount point
File system attributes for /dev/db2log:
flag value      description
-----
-T /db2/DSJ/log_dir Default mount point
```

11. Adapt DB2s autostorage paths using the **db2relocatedb** command as follows:

1. Log on as user **db2<dbsid>** and create a **db2relocatedb** configuration file **relocASpath.cfg** that contains the following lines:

```
DB_NAME=<DBSID>
```

```
DB_PATH=/db2/<DBSID>
```

```
INSTANCE=db2<dbsid>
```

```
STORAGE_PATH=/db2fs/db2data,/db2/<DBSID>/sapdata
```

2. Run the **db2relocatdb** command using the following command:

```
db2relocatedb -f relocASpath.cfg
```

The following example shows the message that should appear after you issued this command:

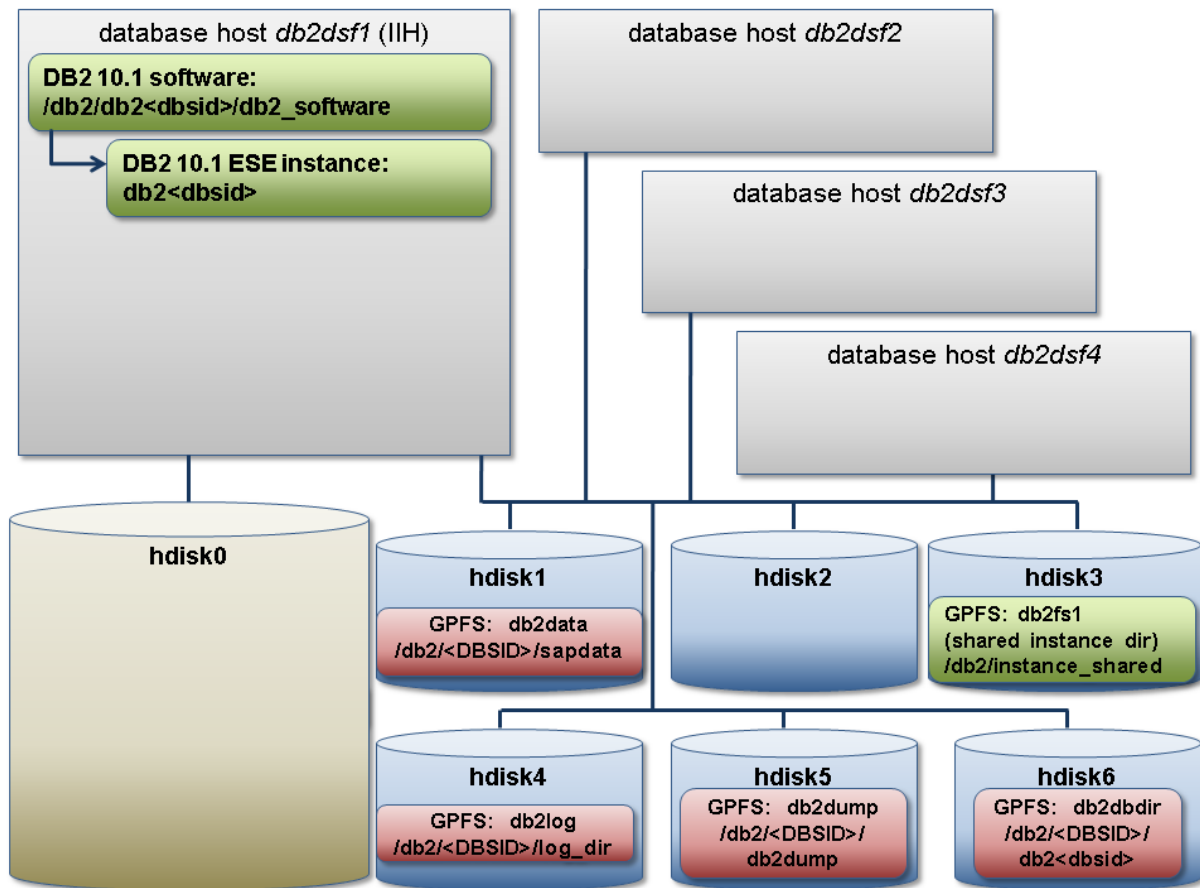
```
Files and control structures were changed successfully.
```

```
DBT1000I The tool completed successfully.
```

12. You can now start the DB2 database and the SAP system.

Result

The following figure shows the configuration of your system:



System Configuration Under New Mount Points

The database of your SAP system is now located on GPFS.

3.3.15 Converting the DB2 10.1 ESE Instance to a DB2 10.1 pureScale Instance

Use

The following section describes how you convert your instance from DB2 10.1 ESE to DB2 pureScale using the SAP installer.

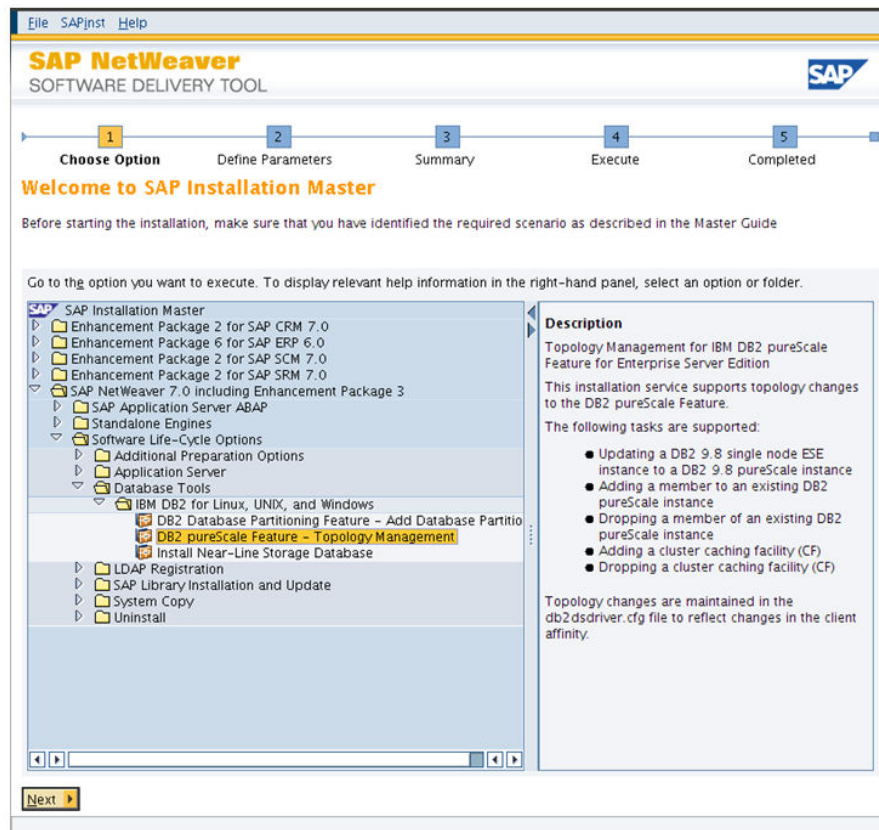
Prerequisites

- Before you start the installer, make sure that the `/sapmnt` directory is available on all participating hosts. The installer modifies the DB2 client configuration that is located in the SAP global directory `/sapmnt/<SAPSID>/global/db6`.
- Make sure that you have 500 MB of free space available for the remote installation directory that the installer prompts you to specify during the dialog phase.

Procedure

1. Stop the SAP system and the database server.
2. Log on to the database server as user `root`.
3. For the following steps of this conversion task, use the SAP installer of an SAP NetWeaver 7.3 release with enhancement package 1. You can use this installer for this task even if you use a different SAP release. (This installer is called software provisioning manager 1.0.)
4. Start the installer as described in the [appropriate installation guide for your SAP system release](#).
5. On the *Welcome* screen, choose ► *<SAP NetWeaver Release>* ► *Generic Installation Options* ► *Database Tools* ► *DB2 pureScale Feature – Topology Management* ▾.

If you use the software provisioning manager 1.0 with support package 5 for the installation, the description of this option only mentions DB2 9.8, but you can also use it for DB2 10.1.



Installer Welcome Screen

6. Choose *Next* and follow the instructions provided on the installer screens.
7. On the *Updating the Instance* dialog, add one member and one CF on different hosts and specify the instance shared directory and the SA MP tiebreaker device.

File SAPInst Help

SAP NetWeaver
SOFTWARE DELIVERY TOOL

1 Choose Option 2 **Define Parameters** 3 Summary 4 Execute 5 Completed

IBM DB2 pureScale – Updating the Instance

Update the DB2 9.8 ESE instance to a DB2 9.8 pureScale instance.

Instance-Shared Directory
Specify the instance-shared directory.
Instance Shared Directory: /db2/instance_shared

Member
Specify the cluster interconnect netname for the displayed database host.

ID	Type	Host Name	Netname
0	MEMBER	db6c31p04	db6c31p04-1b0

Cluster Caching Facility
Specify the cluster caching facility (CF) parameters and the DB2 cluster services tiebreaker.

Type	Host Name	Netname	Tiebreaker Device
CF	db6c11p04	db6c11p04-1b0	/dev/hd1sk2

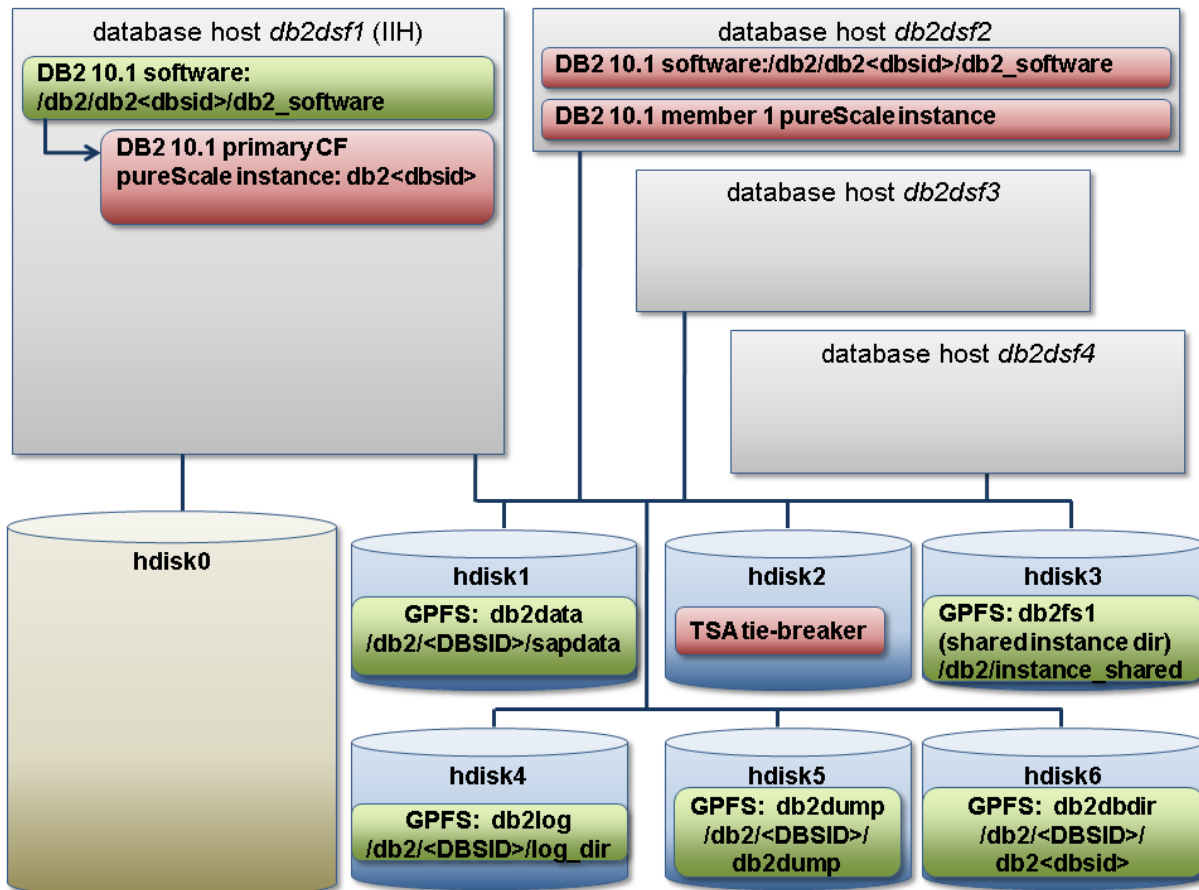
Back Next

Installation Dialog - Updating the Instance

8. Leave the *Changing the Topology* screen as it is.
9. On the *Parameter Summary* screen, you can check the specified parameters and – if required – change them by selecting the parameters and choosing the *Revise* button.
10. To start the installation, choose *Next* on the *Parameter Summary* screen.
11. After the installation has finished successfully, log on to the DB2 pureScale member as user `db2<dbsid>` and start the DB2 pureScale cluster using the following command:
db2start

More Information

After the conversion of the DB2 10.1 instance to DB2 pureScale, you have a functional DB2 pureScale cluster looking as follows:



System Status after Conversion to DB2 pureScale

3.3.16 Adapting the JDBC URL (AS Java Only)

Use

If your SAP system is a Java only system or contains an AS Java (for example, an SAP dual stack system), you have to adapt the JDBC URL that is used to establish a connection to the database.

i Note

You need to perform the following procedure **only** on one AS Java instance.

Procedure

1. Log on to the AS Java as user `<sapsid>adm`.

Note

Make sure that your X Window settings allow this user to open a new X Window.

2. Start the AS Java configuration tool using the following command:
**`cd /usr/sap/<SAPSID>/J{C}<instance-no>/j2ee/configtool
./configtool.sh`**
3. Confirm the dialog box *Do you want to use the default DB settings?* with *Yes*.
4. In the navigation frame of the configuration tool, choose *secure store*.
5. In the *Secure Store Data* group box, select the *jdbc/pool/<DBSID>/Url*.
6. In the *Value* field, change the value for this key as follows:

`jdbc:db2:///`
`<DBSID>:dsdriverConfigFile=<path_to_db2dsdriver.cfg>;<existing_jcc_properties>;`







Example

The existing JDBC URL looks as follows:

`jdbc:db2://saphost1:5912/DSJ:deferPrepares=0;`

You change this JDBC URL to the following:

**`jdbc:db2://saphost1:5912/DSJ:dsdriverConfigFile=/sapmnt/DSJ/global/db6/
db2dsdriver.cfg;deferPrepares=0;`**

7. From the menu, choose  *File*  *Apply* .
8. Exit the configuration tool by choosing  *File*  *Exit*  from the menu.
The changed JDBC URL is used upon the next restart of the AS Java.

3.3.17 Testing the DB2 pureScale Installation

Use

The following section describes how you test the installation of your DB2 pureScale system.

Procedure

1. If the DB2 pureScale cluster has not started yet, log on to the database server as user `db2<dbsid>`.
2. Start the DB2 pureScale cluster using the following command:
`db2start`
3. Log on to an SAP application server as user `<sapsid>adm`.

4. Test the database connection using the following command:

R3trans -d

This should result in the following output:

...

R3trans finished (0000).

If R3trans finishes with a return code other than 0000, review the trans.log file that was written by R3trans and correct the problem.

5. Start the SAP system.
6. Log on to the SAP system and perform some basic tasks.
7. Check the DB2 diagnostic log db2diag.log at /db2/<DBSID>/db2dump/db2diag.log for errors.

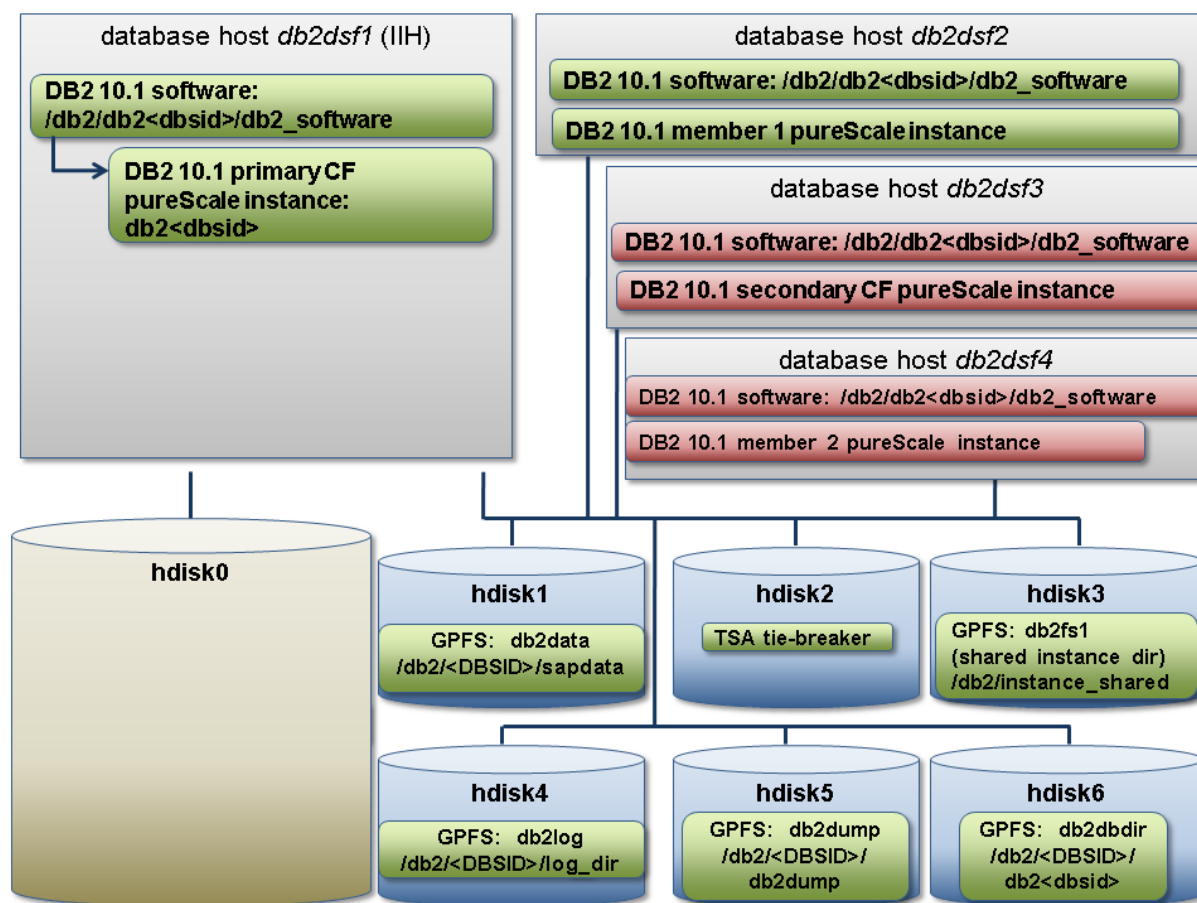
3.3.18 Adding Members and CFs

To add members and a secondary CF to your DB2 pureScale cluster, you follow the procedure described in [Converting the DB2 10.1 ESE Instance to a DB2 10.1 pureScale Instance \[page 35\]](#) and make appropriate changes on the *Changing the Topology* screen of the SAP installer.

i Note

Changes to the DB2 pureScale topology like adding and dropping a member or a CF are currently offline actions. Before you perform such an operation, stop the SAP system and the database.

The following figure shows what your system looks like, for example, after you have created a secondary CF on host db2dsf3 and after you have added an additional member on host db2dsf4:



System Status After Adding Members and CFs

3.3.19 Running the db6_update_db Script

Use

The db6_update_db script can be used to perform certain configuration tasks that are necessary in a SAP environment. This includes granting authorizations to the DB2 connect user that are needed so that applications like the pureScale monitoring extension for the DBA Cockpit work properly.


Procedure

Download the latest db6_update_db script from SAP Note [1365982](#) and run it as described there.

i Note

Make sure that you always use the latest version of this script from this SAP Note. You might face problems later if you use an older version of this script, such as the version that is provided on the DB2 RDBMS DVD.

3.3.20 Checking Configuration Settings

Compare the settings for the database and database manager configuration parameters in your database with the values suggested in **SAP Note 1692571**  that always contains the most up-to-date proposals for these parameters.

Furthermore, review and implement the changes suggested in *Adjusting database configuration parameters to meet DB2 pureScale environment requirements* in the IBM DB2 Information Center at:

[!\[\]\(feabb98897b440bc8695a03336a6e2df_img.jpg\)](http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/t0056787.html).

3.3.21 Installing the DB2 pureScale License


Use

You must have the correct license for your DB2 pureScale installation.

i Note

The following procedure applies **only** if you are an SAP customer and you are entitled to apply for a DB2 OEM license from SAP. In all other cases, you must contact your IBM sales representative to receive the DB2 pureScale licenses.

Procedure

Follow the instructions of SAP Note **816773** , sections *Version 10.1 of DB2 for UNIX and Windows* and *IBM DB2 10 pureScale Feature*.

4 Installation of a New SAP System on DB2 pureScale

4.1 Introduction

Before you install a new SAP system on DB2 pureScale, you must consider the following:

- The database of an SAP system **must** be located on a GPFS file system.

i Note

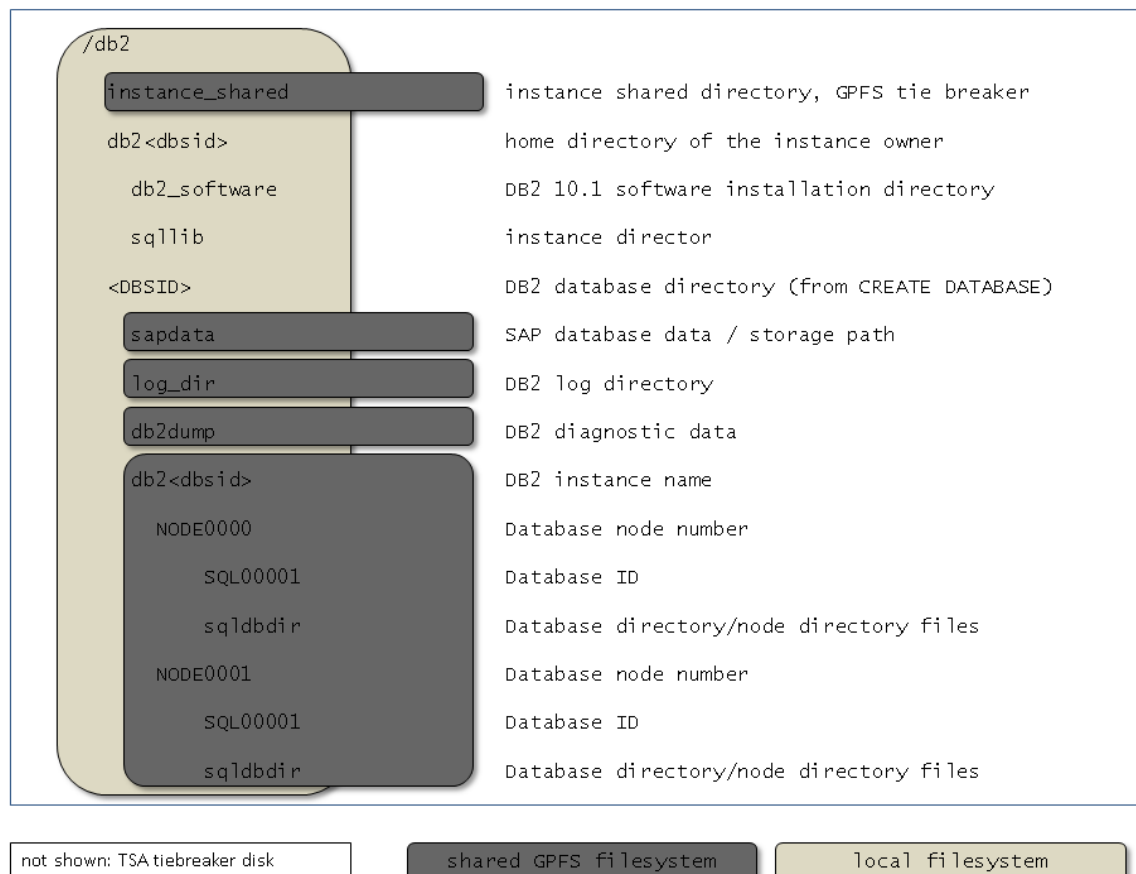
To avoid the overhead of copying or rebalancing the database and other files to the GPFS, the procedure described in this section creates the necessary GPFSs **before** you install the SAP system.

- You must first install the SAP system with DB2 10.1 and then convert the DB2 10.1 ESE instance to a DB2 pureScale instance.

The following sections describe the steps that are required for the installation of a new SAP system on DB2 pureScale in more detail.

4.2 Overview of the DB2 File Systems Required for a new Installation

The following figure shows the recommended directory and file system structure for a new SAP system installation on DB2 10.1 with pureScale:



Overview of Recommended Directory and File System Structure

The following GPFSs are required:

GPFS	Description
/db2/instance_shared	Instance-shared directory containing files that are relevant to all installed instances. It also serves as a tie breaker for the GPFS cluster.
/db2/<DBSID>/sapdata	Database file system containing the SAP database data

i Note

There is only **one** sapdata directory and, therefore, only **one** storage path.

GPFS	Description
db2/<DBSID>/log	DB2 log directory containing the log streams of the DB2 members
db2/<DBSID>/db2dump	DB2 dump directory containing the diagnostic data
db2/<DBSID>/db2<dbsid>	Database directory containing configuration files, the history file, and so on.

Note

In theory, you can limit the number of GPFS file systems by using nested mount points. You create, for example, a large GPFS and mount it under /db2/<DBSID>. In this file system, you can mount another GPFS for the logs under /db2/<DBSID>/log_dir. This way, you do not need additional file systems for the diagnostic data and the database directory.

However, GPFS mounts its file systems in random order. If a mount point does not exist, it is automatically created. For the example shown in the above figure, this means that upon system start or after a failover, it is possible that GPFS mounts first /db2/<DBSID>/log_dir and afterwards /db2/<DBSID>. In this situation, the last mount effectively hides the file system that was mounted first, which leads to an unusable cluster.

4.3 Installing a New SAP System – Step by Step

4.3.1 Introduction

The SAP system installation on DB2 pureScale consists of the following steps that you **must** perform in the **specified** order:

1. [Install the DB2 10.1 software \[page 47\].](#)
2. [Prepare the cluster for the installation \[page 47\].](#)
3. [Manually set up passwordless access for user root \[page 50\].](#)
4. [Prepare the GPFS cluster \[page 51\].](#)
5. [Create the GPFSs \[page 53\].](#)
6. [Install the SAP System on DB2 10.1 ESE \[page 55\].](#)
7. [Apply the latest kernel patch \[page 60\].](#)
8. [Set required configuration parameters \[page 61\].](#)
9. [Convert the DB2 10.1 ESE instance to a DB2 pureScale instance \[page 61\].](#)
10. [Adapt the JDBC URL \(AS Java only\) \[page 63\].](#)
11. [Test your new installation \[page 64\].](#)
12. [Add members and CFs \[page 65\].](#)
13. [Install the DB2 pureScale license \[page 66\].](#)
14. [Run the db6_update_db script \[page 66\].](#)
15. [Check configuration settings \[page 67\].](#)

4.3.2 Installing the DB2 10.1 Software

Use

The DB2 10.1 database software must be available locally on all nodes of the cluster. You install the DB2 software only on one host, which becomes the **install-initiating host (IIH)**. During the creation of additional members and CFs, the DB2 software is automatically installed on the other hosts in the cluster by the **db2iupdt** command.

Procedure

1. Log on to the database server as user **root**.
2. Insert and mount the database DVD to **<DVD_mount>**.
3. Change to directory **<DVD_mount>/<platform>/ESE/disk1**.
4. Run the installation prerequisite check using the following command:

```
./db2prereqcheck
```

The output of the command should contain a line like the following:

```
DBT3533I The db2prereqcheck utility has confirmed that all installation prerequisites were met for DB2 database server with DB2 pureCluster feature. Version: "10.1.0.0".
```

If the check returns any errors, you need to perform the necessary actions before you can continue with the installation of the DB2 10.1 software (for example, upgrading the maintenance level of AIX or applying an AIX service pack).

5. Install the software for DB2 10.1 including the DB2 pureScale Feature using the following command:

```
./db2_install -p ese -f PURESACLE -b <inst_dir>
```

⚠ Caution

Make sure that the DB2 software installation path that you are using is **not** later located on a shared file system. We recommend that you install the DB2 software in the home directory of the instance owner (as in a standard SAP installation), for example, under **/db2/db2<dbsid>/db2_software**.

4.3.3 Preparing the Cluster for the Installation

4.3.3.1 Introduction

Preparing the cluster consists of the following steps:

1. You [prepare the general parallel file system \[page 48\]](#).
2. [Configure IO completion ports \[page 49\]](#).
3. You [identify the cluster interconnect name \[page 50\]](#).
4. You [set up identical users on all nodes \[page 50\]](#).

4.3.3.2 Preparing the GPFS for the SAP System Installation

For a DB2 pureScale cluster, you need at least six physical or virtual shared disks. This corresponds to the five required GPFS file systems as described in [Overview of the Required DB2 File System \[page 45\]](#) plus an additional disk for the TSA tiebreaker.

It is important that you check that these six disks are actually shared on all cluster nodes. To do so, log on to each cluster node as user **root** and run the **lspv** command:

The following shows an example of the output from the **lspv** command:

```
root@db2dsf1 / > lspv
hdisk0      00c5cc1479a245dd  rootvg      active
hdisk1      00c5ccf479278ca7  None
hdisk2      00c5ccf42dae2dbe  None
hdisk3      00c5ccf478d1e31e  None
hdisk4      00c5ccf478d1e49a  None
hdisk5      00c5ccf4d5b68397  None
hdisk6      00c5ccf4d5b68539  None
root@db2dsf2 / > lspv
hdisk0      00c5cc1479a245dd  rootvg      active
hdisk1      00c5ccf479278ca7  None
hdisk2      00c5ccf42dae2dbe  None
hdisk3      00c5ccf478d1e31e  None
hdisk4      00c5ccf478d1e49a  None
hdisk5      00c5ccf4d5b68397  None
hdisk6      00c5ccf4d5b68539  None
```

The shared disks must appear with the same device number on every node in the cluster. The device number is reported in the second column of the **lspv** output. In the example above, the devices with the IDs **00c5ccf479278ca7**, **00c5ccf42dae2dbe**, **00c5ccf478d1e31e**, **00c5ccf478d1e49a**, **00c5ccf4d5b68397**, and **00c5ccf4d5b68539** are shared on both nodes. Depending on the system setup, the device names for the shared disks (such as **hdisk1**, **hdisk2**, . . . , **hdisk6** as shown in the example) can differ on the nodes.

For more information, see *Configuring PVIDs for a DB2 pureScale instance (AIX)* in the IBM DB2 10.1 Information Center.

i Note

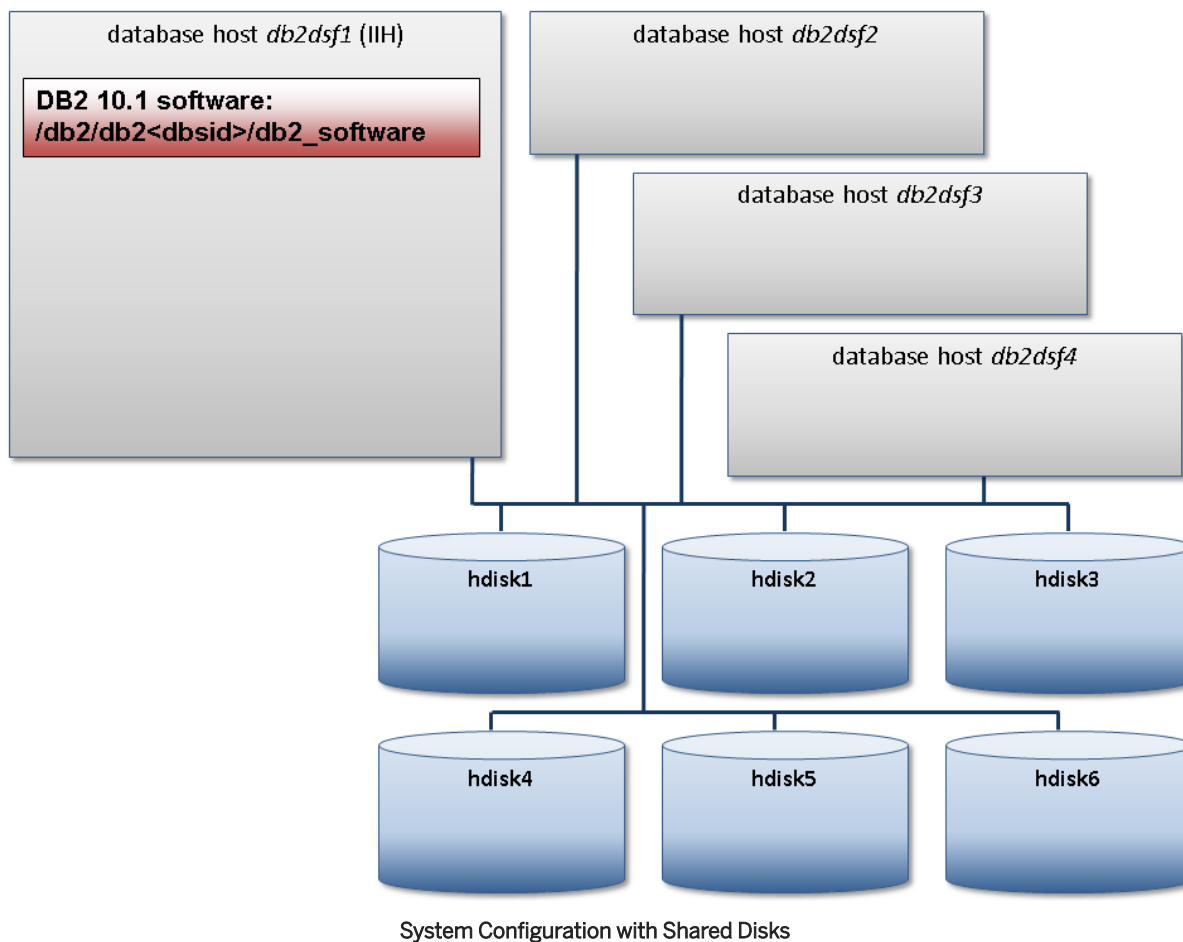
If you want to determine the size of a disk in MB on AIX, you can use the following command:

```
getconf DISK_SIZE /dev/<disk_name>
```

For example:

```
getconf DISK_SIZE /dev/hdisk2
```

The following figure shows a typical system configuration after the installation of the DB2 software and with the shared disks in place:



The DB2 10.1 software is already installed and three additional hosts are available for the DB2 pureScale cluster.

4.3.3.3 Configuring I/O Completion Ports (IOCP)

Use

On AIX, you must ensure that I/O completion ports (IOCPs) are used.

To do so, configure IOCPs as described in the IBM DB2 Information Center *Configuring IOCP (AIX)* at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.perf.doc/doc/t0054518.html> ➡ .

4.3.3.4 Identifying the Cluster Interconnect Netname

Use

For some upgrade tasks, you have to know the correct cluster interconnect netname of the cluster hosts. To find out this netname, run the following command as user `root` on each host:

`<inst_dir>/instance/native/install/db2getnetwork`

❖ Example

The following shows an example of the output from this command:

```
db2dsf1,22.22.22.22,db2dsf1-ib0,ib0
```

```
db2dsf1,33.33.33.33,db2dsf1.wdf.sap.corp,en0
```

In this example, `db2dsf1-ib0` is the cluster interconnect netname of `db2dsf1`.

4.3.3.5 Setting Up Identical Users on all Nodes

Use

The following users must be available on all cluster nodes:

- Database administration user `db2<dbsid>`
- DB2 connect user `sap<sapsid>` (AS ABAP) and `sap<sapsid>db` (AS Java)
- SAP system administration user `<sapsid>adm`

These users must belong to the same groups, the user and group IDs must match, and they must use the same shell and the same path to the home directory.

You can use a central user management for user authorization and authentication, for example, by using the Lightweight Directory Access Protocol (LDAP) or the Network Information Service (NIS). Alternatively, you can also use local user authentication and authorization. In this case, the SAP installer makes sure that all users are created with the correct properties on all members and CFs.

4.3.4 Manually Setting Up Passwordless Access for User Root

Use

In a DB2 pureScale cluster, user `root` must be able to log on to all cluster nodes without a password using secure shell (open ssh). The following procedure describes the steps required to manually configure passwordless access for user `root` between two cluster nodes.

Procedure

1. Log on to the first node of the cluster as user `root`.
2. Generate an RSA key for user `root` using the following command:
`ssh-keygen -t dsa`
3. Specify the file name for the key and a passphrase by entering the default file name for the key without any passphrase:
`~/.ssh/id_dsa`
4. Enable key authentication for `ssh` on the same node using the following command:
`cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys`
5. Transfer the public key to the second node using the following command:
`scp ~/.ssh/id_dsa.pub root@<hostname>:~/id_dsa.pub`
6. Log on to the second node of the cluster as user `root`.
7. Enable key authentication for `ssh` on the second node using the following command:
`cat ~/id_dsa.pub >> ~/.ssh/authorized_keys`
8. Remove the public key file from the home of user `root` on the second node using the following command:
`rm ~/id_dsa.pub`
9. Add both cluster nodes to the list of known hosts using the following command:
**`ssh-keyscan -t dsa <host1_short_name>,<host1_full_name>,
ssh-keyscan -t dsa <host1_short_name>,<host1_full_name>,
<host1_ip> >> ~/.ssh/known_hosts
ssh-keyscan -t dsa <host2_short_name>,<host2_full_name>,`**
10. Set up passwordless access for user `root` on the second node:
To do so, repeat this procedure, but starting from the second node and making sure that you reverse the nodes.

More Information

Installing and setting up OpenSSH in the IBM DB2 Information Center at:

<http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.qb.server.doc/doc/t0055342.html> ➤

4.3.5 Preparing the GPFS Cluster

Use

The database of a DB2 pureScale cluster must be located on a GPFS file system. This file system can be managed either manually using GPFS-specific commands (user-managed file system) or automatically by DB2 (DB2-managed file system).

→ Recommendation

We strongly recommend that you use the **DB2-managed file system** approach. In this way, DB2 performs all necessary GPFS configuration tasks, for example, while new cluster members are added.

To make sure that DB2 pureScale automatically manages the GPFS file system, you need to run the **db2cluster_prepare** command. Among other things, the GPFS file system is then created for the instance-shared directory and certain entries in the DB2 global registry are made.

Procedure

1. Log on to the database server as user **root**.
2. Run the **db2cluster_prepare** command as follows:
<inst_dir>/instance/db2cluster_prepare -l <name of logfile> -instance_shared_dev <shared disk> -instance_shared_mount /db2/instance_shared

❖ Example

```
/db2/db2<dbsid>/db2_software/instance/db2cluster_prepare -l /tmp/  
clusterprep.log -instance_shared_dev /dev/hdisk3 -instance_shared_mount /db2/  
instance_shared
```

If the **db2cluster_prepare** command reports errors, you need to review the log file and take appropriate actions.

3. To check that the instance-shared file system was created, you can use the **db2cluster** command as in the following example:

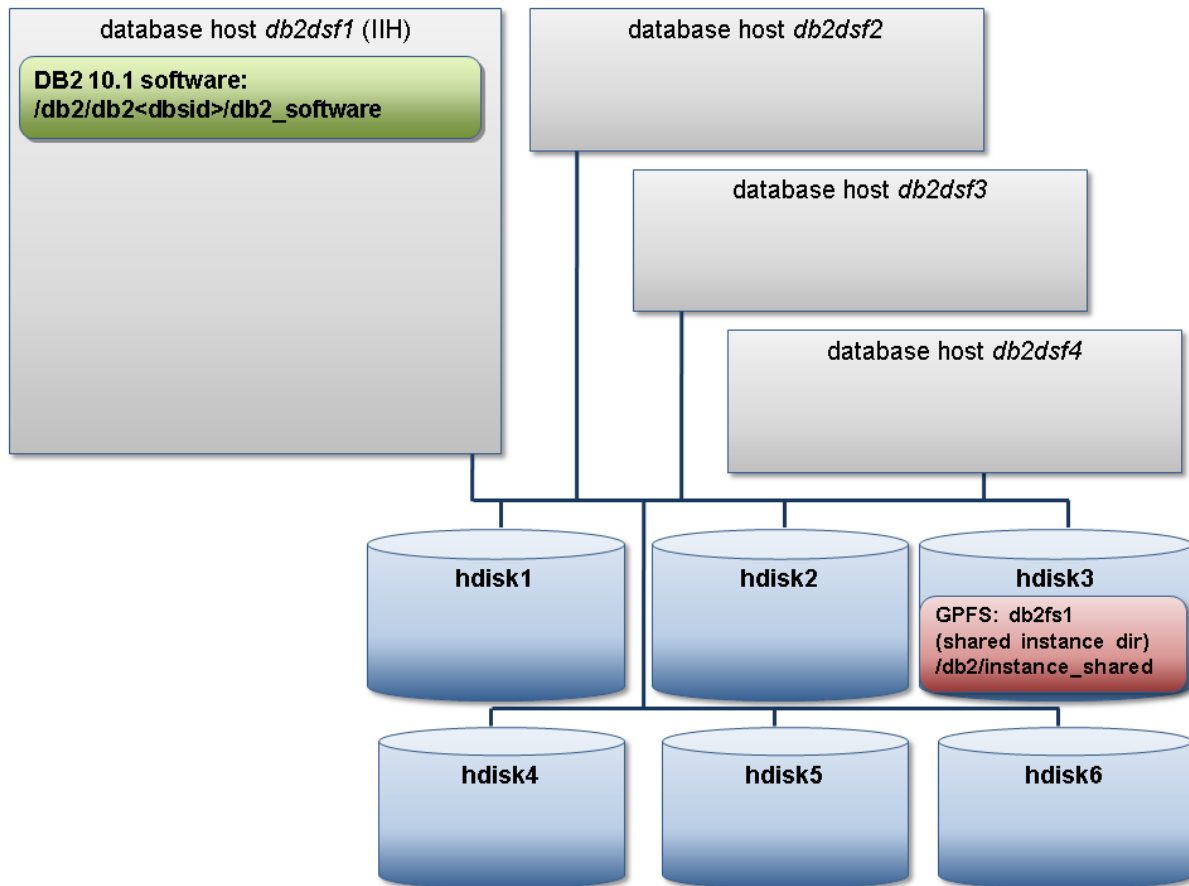
```
<inst_dir>/bin/db2cluster -cfs -list -filesystem
```

An example of the output of this command looks as follows:

FILE SYSTEM NAME	MOUNT_POINT
db2fs1	/db2/instance_shared

Result

At this stage of the overall installation process, your system looks as shown in the following figure:



System Status after Cluster Preparation

4.3.6 Creating the GPFS File Systems

Use

With the following procedure, you create the GPFS file systems.

Note that, in contrast to a standard SAP system installation, we use only one file system for database data in this guide. The reason for this is that, by default, DB2 pureScale uses parallel I/O by setting the DB2 environment variable **DB2_PARALLEL_IO** to **"*"** (see **DB2_PARALLEL_IO** in *System environment variables* in the IBM DB2 10.1 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.regvars.doc/doc/r0005658.html>). At the time of writing this document, there was no sound evidence that providing more than one file system for DB2 data and log files will result in a better overall system performance. You can extend the file system later by adding more disks to it.

i Note

In the following procedure, the **-mount** option of the **db2cluster** command is used. Make sure that the path after the **-mount** option does not end with a slash.

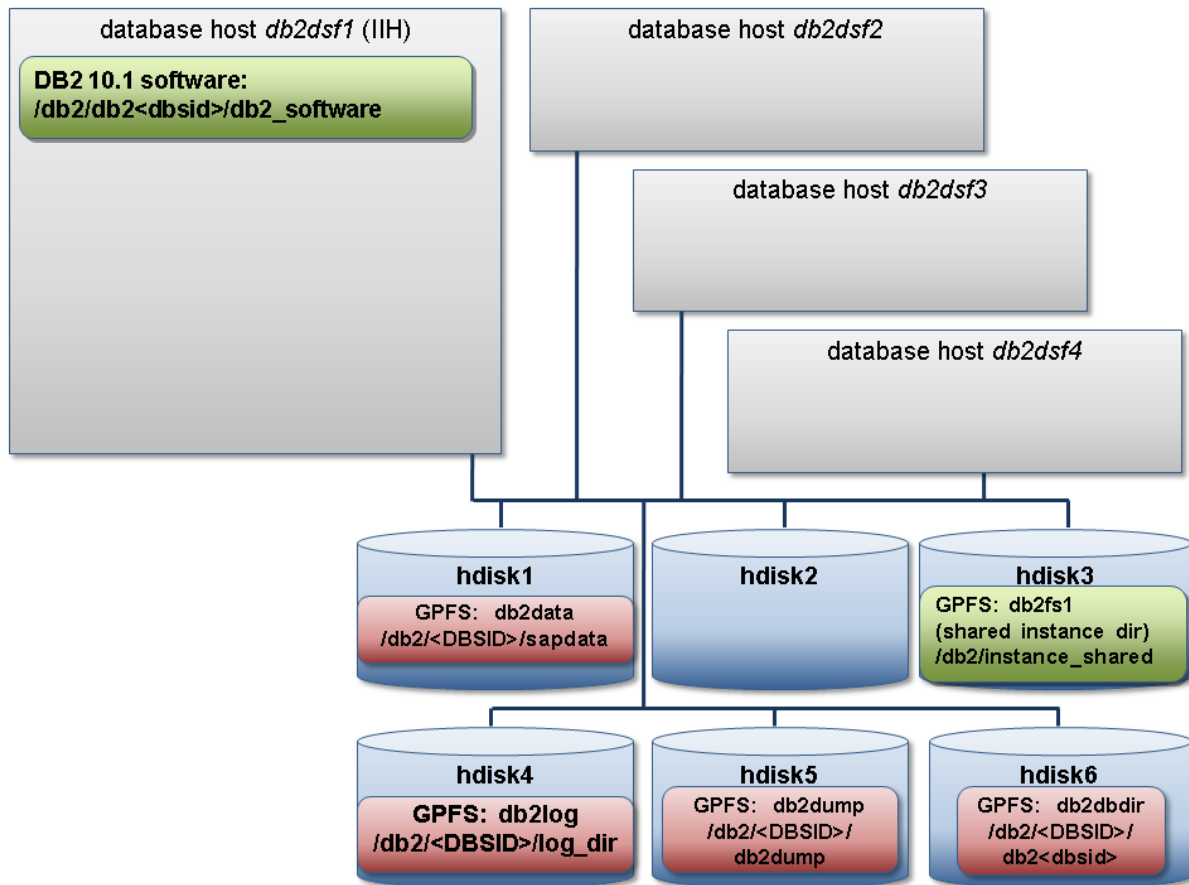
Procedure

1. Log on to the database server as user **root**.
2. Create the GPFS for the SAP database using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2data -disk <sapdata_disk> -mount /db2/<DBSID>/sapdata
where **<sapdata_disk>** is the disk for the file system that is used by the SAP database, for example, **/dev/hdisk1**.
3. Create the GPFS file system for the DB2 log files using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2log -disk <log_disk> -mount /db2/<DBSID>/log_dir
where **<log_disk>** is the disk for the file system that is used by the DB2 log files, for example, **/dev/hdisk4**.
4. Create the GPFS file system for the DB2 diagnostic data using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2dump -disk <dump_disk> -mount /db2/<DBSID>/db2dump
where **<dump_disk>** is the disk for the file system that is used by the DB2 diagnostic data, for example, **/dev/hdisk5**.
5. Create the GPFS for the DB2 database directory using the following command:
<inst_dir>/bin/db2cluster -cfs -create -filesystem db2dbdir -disk <dbdir_disk> -mount /db2/<DBSID>/db2<dbsid>
where **<dbdir_disk>** is the disk for the file system that is used by the DB2 database directory, for example **/dev/hdisk6**.
6. Check that the required file systems are created by running the following command:
<inst_dir>/bin/db2cluster -cfs -list -filesystem
The output should look similar to the following example:

FILE SYSTEM NAME	MOUNT_POINT
db2data	/db2/DSJ/sapdata
db2dbdir	/db2/DSJ/db2dsj
db2dump	/db2/DSJ/db2dump
db2fs1	/db2/instance_shared
db2log	/db2/DSJ/log_dir

Result

At this level of the overall installation process, your system looks as shown in the following figure:



System Status After Creating the GPFS File Systems

4.3.7 Installing the SAP System on a DB2 10.1 ESE Instance

Use

You perform a standard SAP system installation on a DB2 10.1 ESE instance using the SAP installer. During the installation you make sure that the existing DB2 software is used for the creation of the database instance.

Prerequisites

Make sure that you use a DB2 10.1 client DVD that matches the DB2 10.1 RDBMS DVD used during the DB2 software installation (see [Overview of the DB2 File Systems Required for a New Installation \[page 45\]](#)).

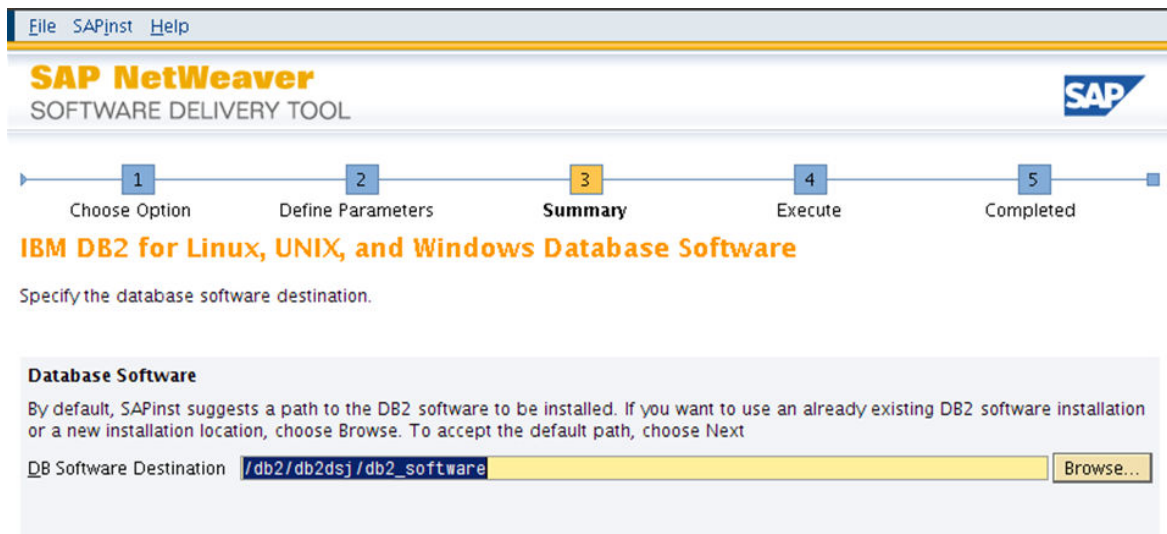
Procedure

When you run the SAP installer to create the SAP database instance, pay attention to the following:

- **Use of the existing DB2 10.1 software installation**

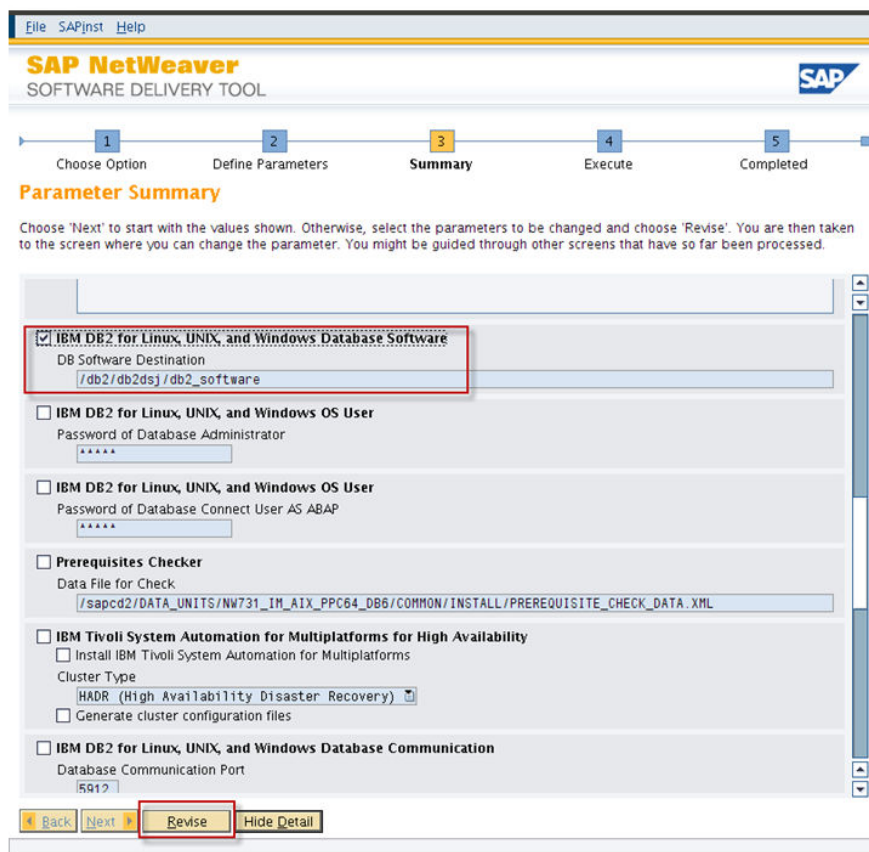
If the installer asks for the DB2 9 RDBMS and DB2 9 client DVDs, you must specify the DB2 10.1 DVDs that you used during the installation of the database software (see [Installing the DB2 10.1 Software \[page 47\]](#)). By default, the installer has an integrated feature to install the DB2 database software during the installation of an SAP system. For this particular installation, however, you need to use the existing database software. To do this, proceed as follows:

Specify the path to the existing database software in the following dialog:



DB Software Destination Dialog

You can find this dialog either by choosing *Custom* (instead of *Typical*) as parameter mode in the installer or by revising the setting of the database software installation directory using the *Revise* button on the *Parameters Summary* screen at the end of the installation dialog phase:



Parameter Summary Screen

- **Specification of only one sapdata directory**

During the database instance installation, specify **only** one sapdata directory named sapdata as follows:

- In custom installation mode
If you run the installer in parameter mode *Custom*, on the screen *IBM DB2 for Linux, UNIX, and Windows Sapdata Directories*, remove the directories suggested by the installer (for example, sapdata1, ..., sapdata4). Instead, add a single sapdata directory. In addition, select the *Use Automatic Storage for Tablespaces* checkbox.
- In typical installation mode
If you run the installer in parameter mode *Typical*, choose the *IBM DB2 for Linux, UNIX, and Windows Sapdata Directories* parameter and then *Revise* on the *Parameter Summary* screen. Then you can change the parameters for IBM DB2 for Linux, UNIX, and Windows sapdata directories. After that, proceed as in custom mode.

The following figures provide examples of the required selections during the installation in custom mode and in typical mode:

IBM DB2 for Linux, UNIX, and Windows Sapdata Directories

Rename, add, or remove sapdata directories for storage paths or tablespace containers.

Attention
We strongly recommend that each sapdata directory resides on a separate physical device.

Storage Management and Tablespaces

Create Tablespaces During the Installation Procedure ☒

Use Automatic Storage Management for Tablespaces (AutoStorage) ☒

Sapdata Directories

Directory
sapdata

Add Remove

Selections for Sapdata Directory (Custom and Typical Installation Mode)

File Help

1 Choose Service 2 Define Parameters 3 Check Parameters 4 Execute Service 5 Completed

Parameter Summary

Choose 'Next' to start with the values shown, or select the parameters you want to change and choose 'Revise'.

☐ IBM DB2 for Linux, UNIX, and Windows Database Memory
Instance Memory [MB]
11618

☐ IBM DB2 for Linux, UNIX, and Windows Minimize Database Size
☐ Use DB2's Row Compression.
☐ Use Deferred Table Creation.

☒ IBM DB2 for Linux, UNIX, and Windows Sapdata Directories
☒ Create Tablespaces During the Installation Procedure
☒ Use Automatic Storage Management for Tablespaces (AutoStorage)
Directory
sapdata

☐ SAP System > Database Import
Number of Parallel Jobs
3

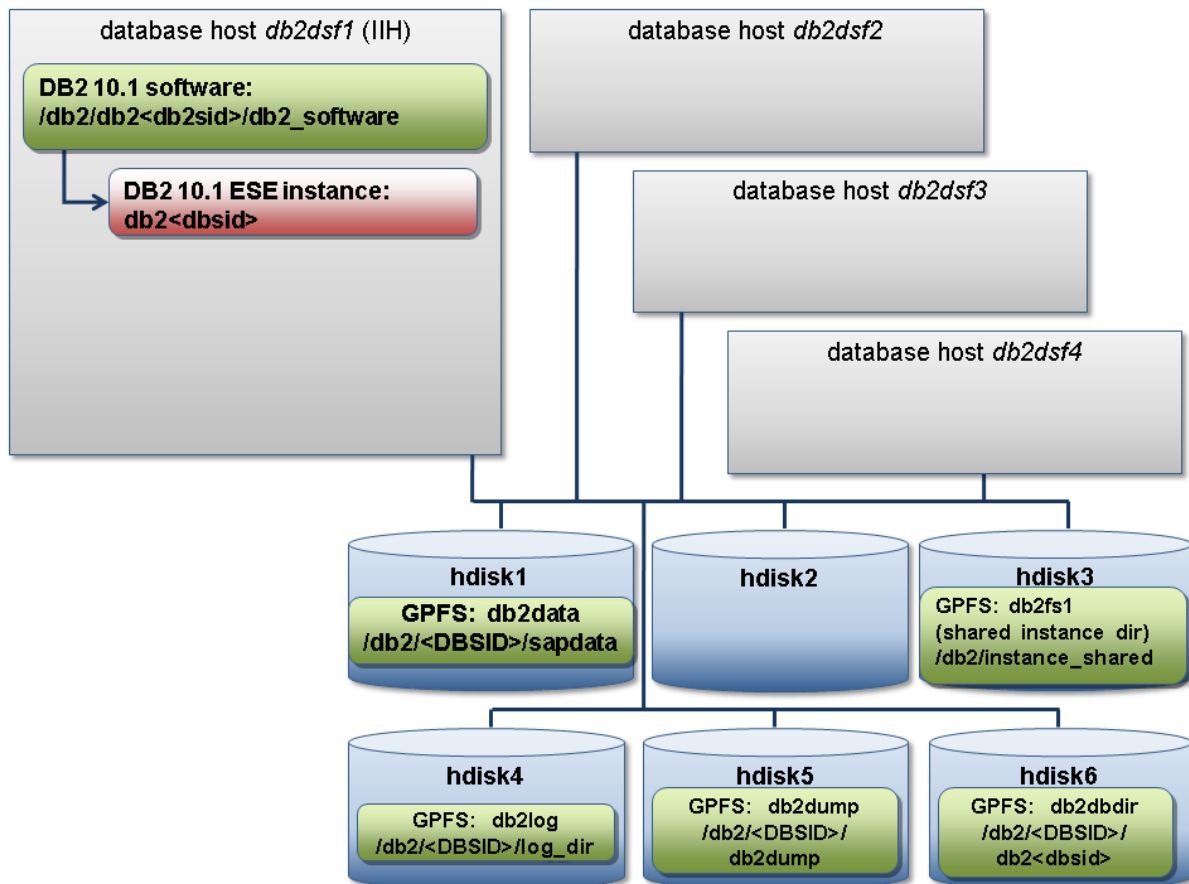
☐ SAP System > Secure Store Settings
Key Phrase (8-30 characters)

Back Next Revise

Selections on the "Parameter Summary" Screen for Typical Installation Mode

Result

The installer has performed a standard SAP system installation on a DB2 10.1 ESE instance, as shown in the following figure:



System Status After Installation of SAP System on DB2 Instance

4.3.8 Applying the Latest Kernel Patch

Use

Make sure that you apply the latest available kernel patch for your SAP system before you continue.

Procedure

For more information about how to download and install a kernel patch, see [SAP Note 19466](#).

4.3.9 Setting Required Configuration Parameters

Use

DB2 with pureScale does not support the health monitor. Therefore, you have to set the database manager configuration parameter `HEALTH_MON` to `OFF`. Furthermore, DB2 pureScale does not support incremental backups. Therefore, you must set the database configuration parameter `TRACKMOD` to `OFF`.

Procedure

1. Log on to the database server as user `db2<dbsid>`.
2. Update the database manager configuration using the following commands:
`db2 "update dbm cfg using HEALTH_MON OFF"`
3. Update the database configuration using the following command:
`db2 "update db cfg using TRACKMOD OFF"`
4. To activate the changes, stop and restart the database manager.
5. To check that the parameters are set correctly, enter the following commands:
`db2 "get dbm cfg" | grep HEALTH_MON`
`db2 "get db cfg" | grep TRACKMOD`

4.3.10 Converting the DB2 10.1 ESE Instance to a DB2 pureScale Instance

Use

The following section describes how you convert your instance from DB2 10.1 ESE to DB2 pureScale using the SAP installer.

Prerequisites

- Before you start the installer, make sure that the `/sapmnt` directory is available on all participating hosts. The installer modifies the DB2 client configuration that is located in the SAP global directory `/sapmnt/<SAPSID>/global/db6`.
- Make sure that you have 500 MB of free space available for the remote installation directory that the installer prompts you to specify during the dialog phase.

Procedure

1. Stop the SAP system and the database server.
2. Log on to the database server as user `root`.
3. For the following steps, use the SAP installer of a SAP NetWeaver 7.3 release with SAP enhancement package 1. (This installer is called software provisioning manager 1.0.) You can use this installer for this task even if you use a different SAP release.
Start the installer as described in the [appropriate installation guide for your SAP system release](#).
4. On the *Welcome* screen, choose **<SAP NetWeaver Release> > Generic Installation Options > Database Tools > DB2 pureScale Feature – Topology Management**.
If you use the software provisioning manager 1.0 with support package 5 for the installation, the description of this option only mentions DB2 9.8, but you can also use it for DB2 10.1.
5. Choose *Next* and follow the instructions provided on the installer screens.
6. On the *Updating the Instance* dialog, add one member and one cluster caching facility (CF) on different hosts and specify the instance shared directory and the SA MP tiebreaker device.

SAP NetWeaver
SOFTWARE DELIVERY TOOL

1 Choose Option 2 **Define Parameters** 3 Summary 4 Execute 5 Completed

IBM DB2 pureScale – Updating the Instance
Update the DB2 9.8 ESE instance to a DB2 9.8 pureScale instance.

Instance-Shared Directory
Specify the instance-shared directory
Instance Shared Directory:

Member
Specify the cluster interconnect netname for the displayed database host.

ID	Type	Host Name	Netname
0	MEMBER	db6c31p04	db6c31p04-1b0

Cluster Caching Facility
Specify the cluster caching facility (CF) parameters and the DB2 cluster services tiebreaker.

Type	Host Name	Netname	Tiebreaker Device
CF	db6c11p04	db6c11p04-1b0	/dev/hd1sk2

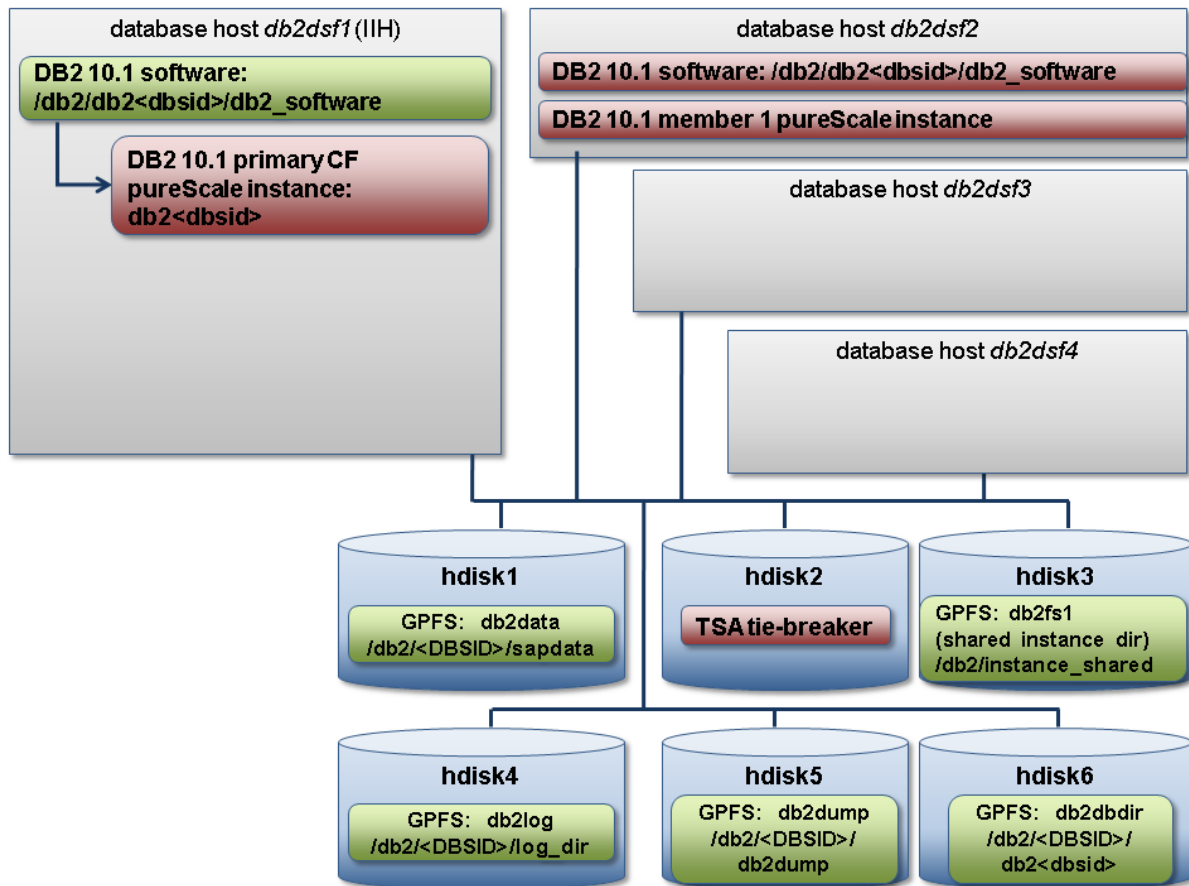
Back Next

Installation Dialog - Updating the Instance

7. Leave the *Changing the Topology* screen as it is.
8. On the *Parameter Summary* screen, you can check the specified parameters and – if required – change them by selecting the parameters and choosing the *Revise* button.
9. To start the installation, choose *Next* on the *Parameter Summary* screen.
10. After the installation has finished successfully, log on to the DB2 pureScale member as user `db2<dbsid>` and start the DB2 pureScale cluster using the following command:
db2start

Result

After the conversion of the DB2 10.1 instance to DB2 pureScale, you have a functional DB2 pureScale cluster that looks as follows:



System Status after Conversion to DB2pureScale

4.3.11 Adapting the JDBC URL (AS Java Only)

Use

If your SAP system is a Java only system or contains an AS Java (for example, an SAP dual stack system), you have to adapt the JDBC URL that is used to establish a connection to the database.

i Note

You need to perform the following procedure only on one AS Java instance.

Procedure

1. Log on to the AS Java as user `<sapsid>adm`.

Note

Make sure that your X Window settings allow this user to open a new X Window.

2. Start the AS Java configuration tool using the following command:
`cd /usr/sap/<SAPSID>/J{C}<instance-no>/j2ee/configtool
./configtool.sh`
3. Confirm the dialog box *Do you want to use the default DB settings?* with *Yes*.
4. In the navigation frame of the configuration tool, choose *secure store*.
5. In the *Secure Store Data* group box, select the *jdbc/pool/<DBSID>/Url*.
6. In the *Value* field, change the value for this key as follows:

```
jdbc:db2:///
<DBSID>:dsdriverConfigFile=<path_to_db2dsdriver.cfg>;<existing_jcc_properties>;
```







Example

The existing JDBC URL looks as follows:

```
jdbc:db2://saphost1:5912/DSJ:deferPrepares=0;
```

You change this JDBC URL to the following:

```
jdbc:db2://saphost1:5912/DSJ:dsdriverConfigFile=/sapmnt/DSJ/global/db6/
db2dsdriver.cfg;deferPrepares=0;
```

7. From the menu, choose  *File*  *Apply* .
8. Exit the configuration tool by choosing  *File*  *Exit*  from the menu.
The changed JDBC URL is used upon the next restart of the AS Java.

4.3.12 Testing the New Installation

Use

To test the installation of your DB2 pureScale system, you perform the following steps.

Procedure

1. If the DB2 pureScale cluster has not started yet, log on to the database server as user `db2<dbsid>`.
2. Start the DB2 pureScale cluster using the following command:
`db2start`
3. Log on to an SAP application server as user `<sapsid>adm`.

4. Test the database connection using the following command:

R3trans -d

This should result in the following output:

...

R3trans finished (0000).

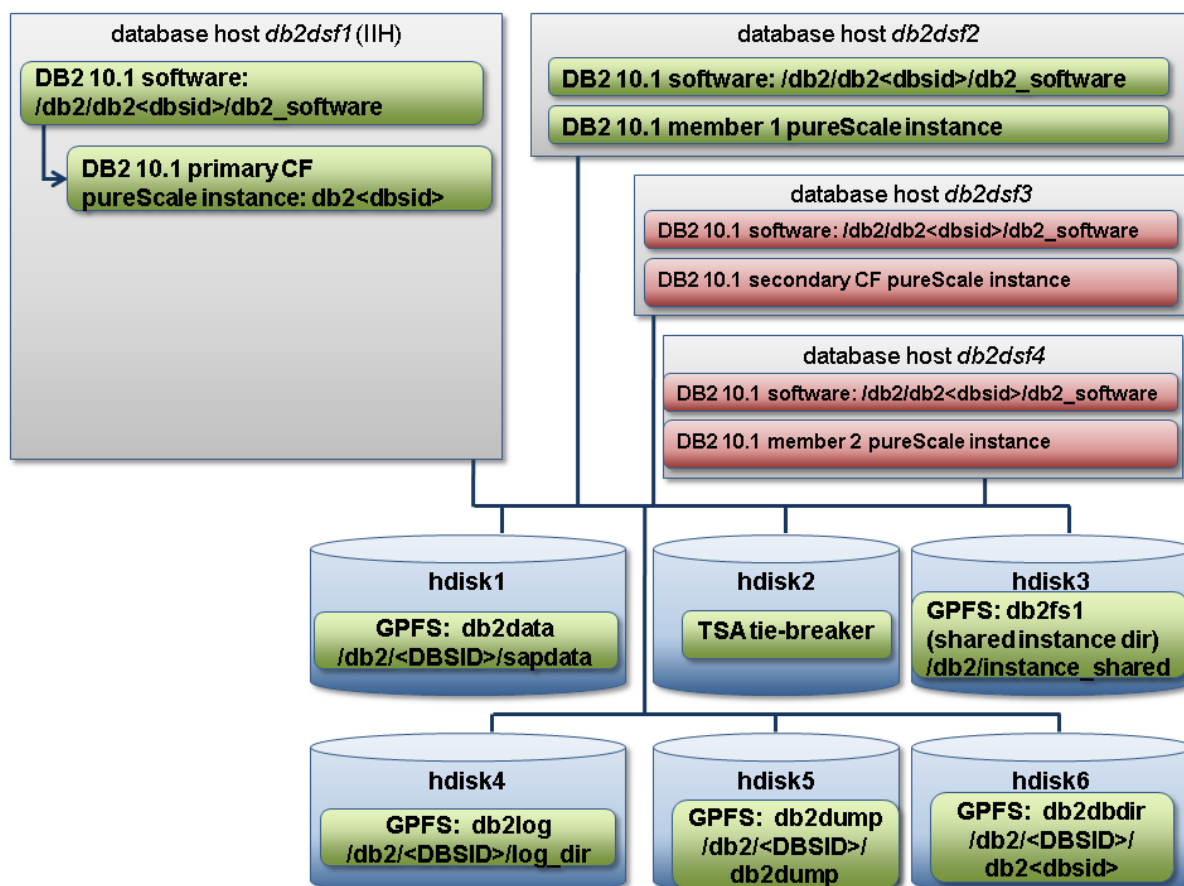
If R3trans finishes with a return code other than 0000, review the trans.log file that was written by R3trans and correct the problem.

5. Start the SAP system.
6. Log on to the SAP system and perform some basic tasks.
7. Check the DB2 diagnostic log db2diag.log at /db2/<DBSID>/db2dump/db2diag.log for errors.

4.3.13 Adding Members and CFs

To add members and a secondary CF to your DB2 pureScale cluster, you use the same procedure as described in [Converting the DB2 10.1 ESE Instance to a DB2 10.1 pureScale Instance \[page 61\]](#) and make appropriate changes on the [Changing the Topology](#) screen of the SAP installer.

After creating a secondary CF on host db2dsf3 and adding an additional member on host db2dsf4, your system looks as shown in the following figure:



System Status After Adding Members and CFs

4.3.14 Installing the DB2 pureScale License

Use

You must have the correct license for your DB2 pureScale installation.

i Note

The following procedure applies **only** if you are an SAP customer and you are entitled to apply for a DB2 OEM license from SAP. In all other cases, you must contact your IBM sales representative to receive the DB2 pureScale licenses

Procedure

Follow the instructions of SAP Note [816773](#), sections *Version 10.1 of DB2 for UNIX and Windows* and *IBM DB2 10 pureScale Feature*.

4.3.15 Running the db6_update_db Script

Use

The `db6_update_db` script can be used to perform certain configuration tasks that are necessary in a SAP environment. This includes granting authorizations to the DB2 connect user that are needed so that applications like the pureScale monitoring extension for the DBA Cockpit work properly.


Procedure

Download the latest `db6_update_db` script from SAP Note [1365982](#) and run it as described there.

i Note

Make sure that you always use the latest version of this script from this SAP Note. You might face problems later if you use an older version of this script, such as the version that is provided on the DB2 RDBMS DVD.

4.3.16 Checking Configuration Settings

As a final step of the overall installation process, you compare the settings for the database and database manager configuration parameters in your upgraded database with the values suggested in **SAP Note 1692571** . This SAP Note always contains the most up-to-date proposals for these parameters.

5 Implementation of SAP-Specific Enhancements

5.1 Overview

The following sections provide information about SAP-specific enhancements that you have to implement after an upgrade to DB2 pureScale or a new installation of an SAP system on DB2 pureScale:

- [Adapting the DB2 Client Connectivity Setup \[page 68\]](#)
- [Partitioning of the Update Tables \(AS ABAP Only\) \[page 74\]](#)
- [Using the DB2 pureScale-Specific Monitoring Enhancements in the DBA Cockpit \[page 78\]](#)

5.2 Adapting the DB2 Client Connectivity Setup

5.2.1 Introduction

In a non pureScale environment, all SAP application servers connect to one database server. In a DB2 pureScale environment with multiple DB2 members on different hosts, additional questions arise, such as:

- To which DB2 member should an SAP application server connect initially (first connect)?
- How should an SAP application server distribute its workload to the different DB2 members?
- What should happen if a DB2 member fails?

As of SAP NetWeaver 7.0 SR3, the AS ABAP uses the IBM Data Server Driver for ODBC and CLI (also known as CLI driver) whereas the AS Java uses the IBM Data Server Driver for JDBC and SQLJ (also known as JDBC driver) to connect to the database. Both clients are installed once in the `/sapmnt/<SAPSID>/global/db6` directory. This `/sapmnt/<SAPSID>/global` directory is shared between all SAP application servers.

In a regular SAP system installation **without DB2 pureScale**, the `db2cli.ini` file is used for the configuration of the client connectivity. The installer creates the `db2cli.ini` file during the SAP system installation. It contains the necessary connection information as shown in the following example:

```
; Comment lines start with a semi-colon.
[DSJ]
Database=DSJ
Protocol=tcPIP
Hostname=db2dsf1
Servicename=5912
[COMMON]
Diagpath=/usr/sap/DSJ/SYS/global/db6/db2dump
```

This client configuration is sufficient in an environment outside DB2 pureScale and in a DB2 pureScale environment with **only** one member.

In an SAP system **with DB2 pureScale**, you configure the CLI driver using the file `db2dsdriver.cfg`. If you use the installer (as described in *Updating the DB2 10.1 ESE Instance to a DB2 pureScale Instance*, it creates a default `db2dsdriver.cfg` file.

The following sections describe the resulting default client connectivity behavior in a DB2 pureScale environment, that is, the round-robin connectivity setup and how you can implement a user-defined client connectivity.

5.2.2 Using the Round-Robin Connectivity Setup

During the upgrade to DB2 pureScale or during the installation of additional DB2 members, the SAP installation tool modifies the client configuration as follows:

- If necessary, comments out the connection information in the `db2cli.ini` configuration file
- If it does not exist, creates the `db2dsdriver.cfg` file
- Creates or updates a list with all DB2 members together with the respective connection information (that is, the host name and the port number) in the `db2dsdriver.cfg` file
- Creates or updates a list with all SAP application servers in the `db2dsdriver.cfg` file

As a result, the SAP applications servers connect in a round-robin way to the available DB2 members.

For example, an SAP system consists of six SAP application servers (`sapdsf1`,..., `sapdsf6`) and three DB2 members (`db2dsf1`,...,`db2dsf3`). During startup of the SAP application servers, all connections from `sapdsf1` are made to `db2dsf1`, all connections from `sapdsf2` to `db2dsf2`, all connections from `sapdsf3` to `db2dsf3`, and all connections from `sapdsf4` to `db2dsf1` (beginning again with the first DB2 member in the list), and so on. All further connects (for example, secondary connections or during an SAP work process restart) are made to the same DB2 member.

If a DB2 member fails, new connections are made again in a round-robin like fashion. If `db2dsf1` fails, `sapdsf1` and `sapdsf4` connect to `db2dsf2` and check regularly if `db2dsf1` is online again. As soon as `db2dsf1` is back online, they redirect their connections at the next transaction boundary back to `db2dsf1`. If `db2dsf1` is not online and `db2dsf2` also fails, all workload is directed to the third member `db2dsf3`. The SAP application servers `sapdsf1` and `sapdsf4` now check regularly whether either `db2dsf1` or `db2dsf2` has come back online. If possible, the SAP application servers `sapdsf1` and `sapdsf4` connect directly to their home host `db2dsf1`. If `db2dsf1` is not available but `db2dsf2` comes back online, they connect to `db2dsf2` and check regularly if `db2dsf1` comes back online. Once the first member `db2dsf1` is available, the SAP application servers `sapdsf1` and `sapdsf4` finally connect to this member.

The following table illustrates this connection scenario. The first fallback indicates the preferred DB2 member for the fallback, and second fallback indicates the DB2 member to which an SAP application server falls back if the first fallback is not available.

SAP Application Server	Primary Member	1st Failover	2nd Failover	1st Fallback	2nd Fallback
sap_as_1	member_1	member_2	member_3	member_1	member_2
sap_as_2	member_2	member_3	member_1	member_2	member_3

SAP Application Server	Primary Member	1st Failover	2nd Failover	1st Fallback	2nd Fallback
sap_as_3	member_3	member_1	member_2	member_3	member_1
sap_as_4	member_1	member_2	member_3	member_1	member_2
sap_as_5	member_2	member_3	member_1	member_2	member_3
sap_as_6	member_3	member_1	member_2	member_3	member_1




The following is an example of a `db2dsdriver.cfg` file that represents this scenario described in the above table:

```
<configuration>
  <dnscollection>
    <dsn alias="DSJ" name="DSJ" host="db2dsf1.wdf.sap.corp" port="5912" />
  </dnscollection>
  <databases>
    <database name="DSJ" host="db2dsf1.wdf.sap.corp" port="5912">
      <acr>
        <parameter name="enableAcr" value="true" />
        <parameter name="enableSeamlessAcr" value="true" />
        <parameter name="affinityFailbackInterval" value="60" />
        <alternateserverlist>
          <server name="server_1" hostname="db2dsf1.wdf.sap.corp" port="5912" />
          <server name="server_2" hostname="db2dsf2.wdf.sap.corp" port="5912" />
          <server name="server_3" hostname="db2dsf3.wdf.sap.corp" port="5912" />
        </alternateserverlist>
        <clientaffinityroundrobin>
          <client name="sap_as_1" hostname="sapdsf1.wdf.sap.corp"/>
          <client name="sap_as_2" hostname="sapdsf2.wdf.sap.corp"/>
          <client name="sap_as_3" hostname="sapdsf3.wdf.sap.corp"/>
          <client name="sap_as_4" hostname="sapdsf4.wdf.sap.corp"/>
          <client name="sap_as_5" hostname="sapdsf5.wdf.sap.corp"/>
          <client name="sap_as_6" hostname="sapdsf6.wdf.sap.corp"/>
        </clientaffinityroundrobin>
      </acr>
    </database>
  </databases>
  <parameters>
    <parameter name="CommProtocol" value="TCPIP" />
  </parameters>
</configuration>
```

The relevant configuration is made in the automatic client reroute (ACR) section of the `db2dsdriver.cfg` file. ACR is enabled and the fallback interval is set, for example, to 60. The value 60 means that 60 seconds after the first transaction boundary was reached, the client checks if it can fall back to the primary DB2 member.

The `db2dsdriver.cfg` file also contains a list of all DB2 members in `<alternateserverlist>` and a list of all SAP application servers in `<clientaffinityroundrobin>`.

For more information about the elements in the automatic client reroute (ACR) section of the `db2dsdriver.cfg` configuration file, see the following sections in the IBM DB2 10.1 Information Center:

Topic	URL
<i>Client affinities for clients that connect to DB2 Database for Linux, UNIX, and Windows</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.apdv.cli.doc/doc/c0056508.html 
<i>Configuration of client affinities for non-Java clients for DB2 Database for Linux, UNIX, and Windows connections</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.apdv.cli.doc/doc/c0056258.html 
<i>Example of enabling client affinities for non-Java clients for DB2 Database for Linux, UNIX, and Windows connections</i>	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.apdv.cli.doc/doc/c0056272.html 

5.2.3 Using a User-Defined Connectivity Setup


The round-robin connectivity setup provides a good distribution of the database workload to the DB2 members. For the scalability of the DB2 pureScale cluster, it is important that the DB2 members work as much as possible on their own objects. If multiple DB2 members access a database object at the same time, the access to these database objects must be serialized with the help of locks. To avoid contention, it is important that the workload is distributed to the DB2 members in such a way that locking is minimized. In an SAP DB2 pureScale environment, you can avoid locking as much as possible by doing the following:

1. Distribute the SAP workload to the different SAP application servers by using logon groups.

Example

An SAP system with six applications servers is mostly used for HR- and CRM-related workload. To evenly distribute the workload, you create logon groups so that all HR-related workload is directed to SAP application servers `sapdsf1` to `sapdsf3` and all CRM-related workload is directed to SAP application servers `sapdsf4` to `sapdsf6`.

The definition and creation of SAP logon groups is **not** described in this document. For more information, see *Creating a New Logon Group*, for example in the SAP NetWeaver 7.0 library at:

http://help.sap.com/saphelp_nw70/helpdata/en/f3/795a421b5ec153e10000000a1550b0/frameset.htm 

2. For every SAP application server, you define manually preferred DB2 members (which we call a user-defined connectivity setup).

Example of a User-Defined Connectivity Setup

A DB2 pureScale cluster consists of two members. If you continue the example from above, you can define that the SAP application servers `sapdsf1` to `sapdsf3` primarily connect to the first DB2 member and the SAP application servers `sapdsf4` to `sapdsf6` primarily connect to the second DB2 member. It is now likely that

both DB2 members work on different database objects most of the time, which avoids locking and improves scalability.

The following is an example of the specification of such a setup in the `db2dsdriver.cfg` file:

```
<configuration>
<dsnccollection>
  <dsn alias="DSJ" name="DSJ" host="db2dsf1.wdf.sap.corp" port="5912" />
</dsnccollection>
<databases>
  <database name="DSJ" host="db2dsf1.wdf.sap.corp" port="5912">
    <acr>
      <parameter name="enableAcr" value="true" />
      <parameter name="enableSeamlessAcr" value="true" />
      <parameter name="affinityFailbackInterval" value="5" />
      <alternateserverlist>
        <server name="member_1" hostname="db2dsf1.wdf.sap.corp" port="5912" />
        <server name="member_2" hostname="db2dsf2.wdf.sap.corp" port="5912" />
      </alternateserverlist>
      <affinitylist>
        <list name="as_group_1" serverorder="member_1,member_2" />
        <list name="as_group_2" serverorder="member_2,member_1" />
      </affinitylist>
      <clientaffinitydefined>
        <client name="sap_as_1" hostname="sapdsf1.wdf.sap.corp"
          listname="as_group_1" />
        <client name="sap_as_2" hostname="sapdsf2.wdf.sap.corp"
          listname="as_group_1" />
        <client name="sap_as_3" hostname="sapdsf3.wdf.sap.corp"
          listname="as_group_1" />
        <client name="sap_as_4" hostname="sapdsf4.wdf.sap.corp"
          listname="as_group_2" />
        <client name="sap_as_5" hostname="sapdsf5.wdf.sap.corp"
          listname="as_group_2" />
        <client name="sap_as_6" hostname="sapdsf6.wdf.sap.corp"
          listname="as_group_2" />
      </clientaffinitydefined>
    </acr>
  </database>
</databases>
<parameters>
  <parameter name="CommProtocol" value="TCPIP" />
</parameters>
</configuration>
```

The `<alternateserverlist>` contains all DB2 members again with their host name and communication port. Several ordered lists of these DB2 members are defined in the `<affinitylist>`. The order of the members in these lists defines the failover strategy. One of these lists is afterwards assigned to every SAP application server in the `<clientaffinitydefined>` section.

For example, if member `db2dsf1` fails, all affected SAP application servers follow the order as specified in affinity list `as_group_1` and fail over to member `db2dsf2`. If member `db2dsf2` fails, all affected SAP application servers follow the order as specified in affinity list `as_group_2` and fail over to member `db2dsf1`.

The following table shows this connectivity setup:

SAP Application Server	List	Primary Member	1st Failover
sap_as_1	as_group_1	member_1	member_2

SAP Application Server	List	Primary Member	1st Failover
sap_as_2	as_group_1	member_1	member_2
sap_as_3	as_group_1	member_1	member_2
sap_as_4	as_group_2	member_2	member_1
sap_as_5	as_group_2	member_2	member_1
sap_as_6	as_group_2	member_2	member_1

This user-defined connectivity setup allows for a detailed assignment of SAP application servers to the members in a DB2 pureScale cluster. Together with SAP logon groups of the AS ABAP, you can effectively minimize locking situations in the DB2 pureScale cluster.

5.2.4 Adapting the DB2 Client Configuration After Installation or Deinstallation of SAP Application Servers

As outlined in *Using the Round-Robin Connectivity Setup*, the SAP installation tool automatically adapts the `db2dsdriver.cfg` file if you install additional or drop existing members of the DB2 pureScale cluster.

However, if you install additional or drop existing SAP application servers, this file must also be adapted.

For SAP NetWeaver 7.0 with Enhancement Package 3 and SAP NetWeaver 7.3 with Enhancement Package 1 and Higher

The SAP installer automatically adapts the `db2dsdriver.cfg` file during the installation of a new dialog instance. After the installation, check the `db2dsdriver.cfg` file and make further adaptations as necessary, for example, in a user-defined connectivity setup.

For All Other SAP NetWeaver Releases

After installing or dropping an SAP application server, you must adapt the `db2dsdriver.cfg` file manually as follows:

- If you use the round-robin client connectivity setup, you must insert or remove an entry in the `<clientaffinityroundrobin>` section of the `db2dsdriver.cfg` file for the respective SAP application server.
- If you use the user-defined client connectivity setup, you must create an additional entry in the `<clientaffinitydefined>` section of the `db2dsdriver.cfg` file for a newly installed SAP application

server. With the `listname` attribute, you specify to which DB2 pureScale members the application server should connect. If you need to drop an SAP application server, remove the corresponding entry in the `<clientaffinitydefined>` section.

After every SAP application server restart, the modified version of the `db2dsdriver.cfg` file is automatically read and the applied changes are also automatically considered. You can manually reload an updated configuration in the `db2dsdriver.cfg` file by using the [Reload Client Affinity File on all Application Servers](#) pushbutton on the *DB2 pureScale Feature: Client Affinity* screen in the DB2 pureScale monitoring extension of the DBA Cockpit.

For more information, see [Checking the Client Affinity \[page 87\]](#).

5.3 Partitioning the Update Tables (AS ABAP Only)

Use

A central component of AS ABAP is the update management. To perform updates to the database, the SAP application servers use special work processes of type UPD and UP2. Database updates are deferred until the end of an SAP transaction. During an SAP transaction, update requests are stored in the tables VBHDR, VBMOD, and VBDATA.

These tables are also called update tables. In highly parallel environments where many applications servers read from and write to these tables, the overall performance of the system can decrease due to locking issues. To overcome these problems, we recommend that you partition the update tables based on the SAP application server name if there is a considerable amount of update task activity in your system. If the update tables are partitioned, database locks are performed on table partition level and contention does not occur.

Prerequisites

1. Log on to the SAP system.
2. Make sure that your SAP system includes all the Support Packages or corrections mentioned in SAP Note [1379362](#) and SAP Note [1594952](#).
3. Make sure that you have installed the latest version of the DB6CONV report (see SAP Note [1513862](#)) on your system.
4. Call SAP transaction SM13 and ensure that there are no open updates in the SAP system.

Procedure

i Note

The following procedures are based on the assumption that the system consists of SAP application servers `sapdsf1` and `sapdsf2`. Both servers are configured for update processing.

Configuring the SAP System

1. For performance and high-availability reasons, we strongly recommend that you set up [local update processing](#), that is, every SAP application server can process updates. To do so, configure UPD and UPD2 work processes on each application server.
2. If local update processing is configured, load balancing of update tasks between the application servers takes place by default. To avoid lock situations, switch off this update task dispatching by setting the following SAP profile parameters on all SAP application servers:

```
rdisp/vb_dispatching = 0
```

```
rdisp/vbname = <name of the local application server with update services>
```

3. The [VBKEY](#) field serves as partitioning key for all three update tables. To switch the [VBKEY](#) to a format that allows table partitioning based on the host name, specify the following SAP profile parameters for all application servers:

```
dynp/luw_id_format = 2
```

```
rdisp/vb_key_comp = HOST/SYNR/WPNR/DATE/TIME/STMP
```

4. To represent the host name in the [VBKEY](#) field, you can choose between the first 8 characters of the textual representation of the host name or the IPv4 address of the host. We recommend that you use the host name here. In this way, you can, for example, avoid problems with IPv6 addresses. To do so, specify the following SAP profile parameter for all SAP application servers:

```
rdisp/vb_key_use_hostname = 1
```

5. For reasons of SAP system high availability, it can be useful to specify a virtual host name and a virtual IP address for each SAP application server. These virtual addresses decouple the SAP application server (the running ABAP or Java kernel) from the underlying host. In this case, an application server can be moved from one host to another host if necessary, for example, due to a system outage. The following profile parameters specify the virtual host name:

```
SAPLOCALHOST = <virtual host name>
```

```
SAPLOCALHOSTFULL = <full qualified virtual host name>
```

The SAP installation already sets these profile parameters if called with the installation option

[SAPINST_USE_HOSTNAME](#). For more information, see SAP Notes [962955](#) and [1564275](#).

In such an environment, you must use the virtual host names for table partitioning of the update tables.

In this way, after a failover of one SAP application server to another host, this server can still access its own update records in the update tables (in its own partition). To use the virtual IP address or virtual host name as specified in [SAPLOCALHOST](#) in the update key, set the following SAP profile parameter on all SAP application servers:

```
rdisp/vb_key_use_saplocalhost = 1
```

Note

You can set all profile parameters except for `rdisp/vbname` once in the `DEFAULT . PFL` profile (if it exists in your SAP release) so that they are valid for each application server.

6. Restart your SAP system so that all changes take effect.

Partitioning the Update Tables

The update tables are created unpartitioned during the SAP system installation. To perform the table partitioning, use the `DB6CONV` report together with the online table move UDF. Before starting the `DB6CONV` conversion, enter the storage parameters for the table partitioning as described in the following procedure:

1. Start transaction SE14, enter [VBHDR](#) as table name, and choose the [Edit](#) button.
2. Choose the [Storage Parameters](#) button and then the [For New Creation](#) button.
3. In the dialog window, choose the [Current Database Parameters](#) radio button.

4. To insert new *OPTIONS* lines for the partition clause, place the cursor on an *OPTIONS* line and use **[SHIFT F5]** (or the relevant button).
5. Enter the partitioning clause as shown in the following example:

Storage parameters: (display and maintain)

Table: VBHDR
Maintain marked parameters for next new creation

Storage	Table	Parameter	Value
		TABLESPACE	DSJ#PROTD
		INDEXSPACE	DSJ#PROTI
		LONGSPACE	DSJ#PROTD
		LOCKSIZE	ROW
		OPTIONS	PARTITION BY RANGE
		OPTIONS	("VBKEY")
		OPTIONS	(
		OPTIONS	PARTITION "PART_sapdsf1"
		OPTIONS	STARTING('sapdsf100000000000000000000000000000000') INCLUSIVE
		OPTIONS	ENDING ('sapdsf12ZZZZZZZZZZZZZZZZZZZZZZZZZZZZ') INCLUSIVE,
		OPTIONS	PARTITION "PART_sapdsf2"
		OPTIONS	STARTING('sapdsf200000000000000000000000000000000') INCLUSIVE
		OPTIONS	ENDING ('sapdsf22ZZZZZZZZZZZZZZZZZZZZZZZZZZZZ') INCLUSIVE
		OPTIONS)

Storage Parameters for Table Partitioning

In the partitioning clause, specify the following:

- Field *VBKEY* as partitioning key
 - One partition for every SAP application server. We recommend that you choose *PART_<hostname>* as partition name.
 - The low value and high value (as defined after *STARTING* and *ENDING*) for each partition must be a string literal with 32 characters. The first eight characters are determined by the respective application server host name. The low value is then filled with "0"s (the number zero), the high value is filled with "Z". If the host name of the application server is less than 8 characters, it needs to be right-padded with "0" or "Z".
6. Write down the regular tablespaces and index tablespaces in which the tables are located, for example, <SID>#PROTD, <SID>#PROTI (you must remember them in step 9).
 7. Save the storage parameters.
 8. In transaction SE38, start the DB6CONV report.
 9. Enter the tablespaces from step 6 and perform a table move.

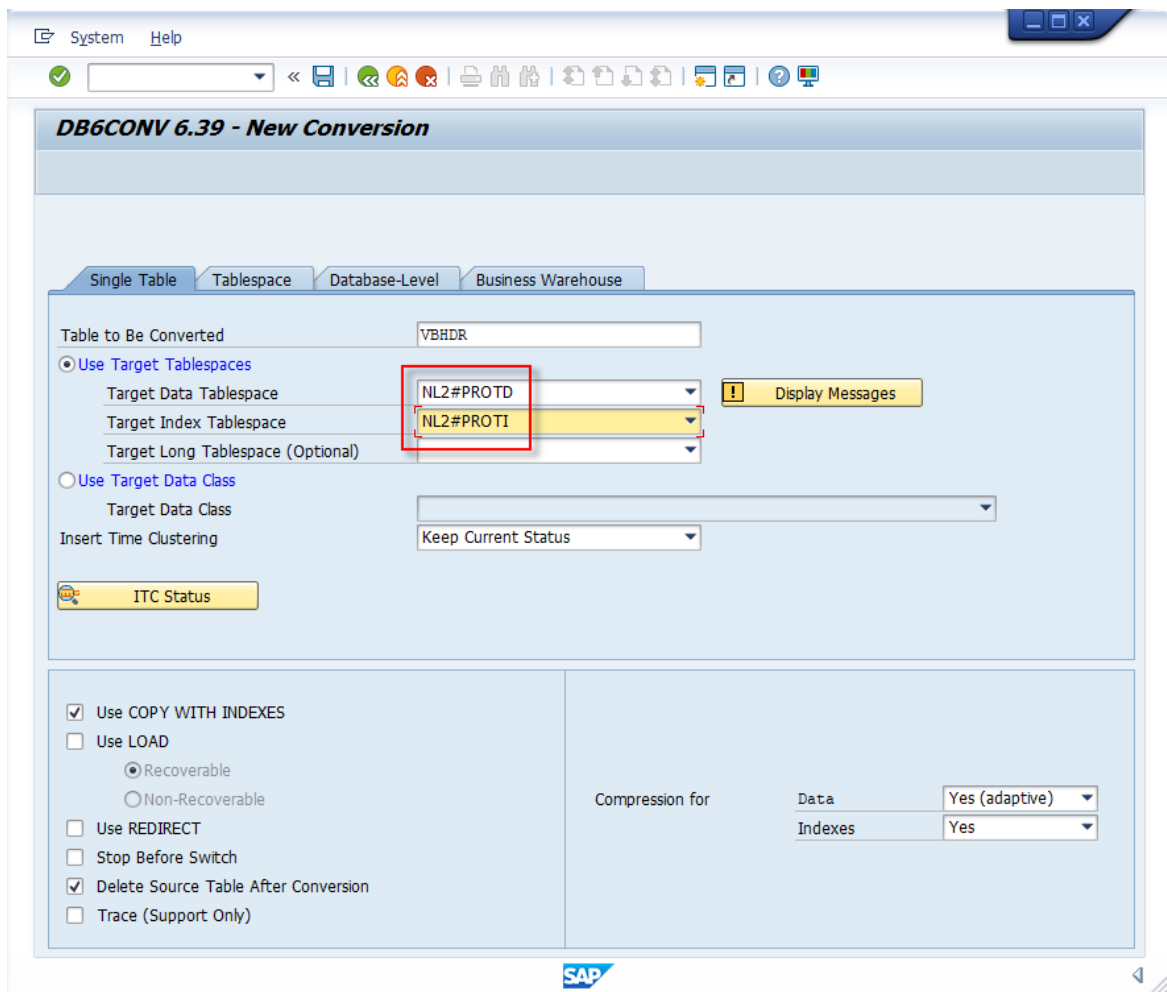


Table Move Using DB6CONV

10. Perform the same steps for tables VBMOD and VBDATA.
11. To check the update functionality of a specific SAP application server, proceed as follows:
 1. Log on to the application server and call transaction SM12.
 2. From the menu, choose ► *Extras* ► *Diagnosis in Update* ►. The log file that is displayed should contain no errors.

→ Recommendation

We strongly recommend that you perform this test on all SAP application servers.

More Information

- SAP Note [1379362](#): DB6: Support for partitioned tables in the ABAP DDIC
- SAP Note [1513862](#): DB6: Table conversion using DB6CONV version 6 or higher
- SAP Note [191191](#): Partitioning update tables
- SAP Note [962955](#): Use of virtual TCP/IP host names

- SAP Note [1564275](#): *Install SAP Systems Using Virtual Host Names on Windows*
- *Update Management* on SAP Help Portal for your SAP NetWeaver version, for example, at <https://help.sap.com/viewer/979cf1522d164bf7a781796efd8850ee/7.4.19/en-US/078cb02dc14d497f9779f7a309c1a7bc.html>

5.4 Using DB2 pureScale Monitoring Enhancements in the DBA Cockpit (AS ABAP Only)

5.4.1 Introduction

The DBA Cockpit was enhanced with specific monitoring functions that let you monitor DB2 pureScale members and CFs. That is, these DB2 pureScale-specific enhancements provide the following information:

- An overview of the cluster topology and cluster alerts
- An overview of the DB2 member performance in the DB2 pureScale cluster
- Information about the group buffer pool and local buffer pool hit ratio
- The content of the (global) package cache
- Information about the memory configuration and consumption of the cluster caching facilities (CFs)
- Information about page contention in the DB2 pureScale cluster
- Information about connected clients and the client configuration (that is, the content of the `db2dsdriver.cfg` file)

Note

The DB2 pureScale-specific monitoring functions are only available in the SAP GUI-based user interface of the DBA Cockpit.

To use the enhanced version of the DBA Cockpit, you have to check whether your SAP system has the appropriate Support Package level as described in **SAP Note 1409540**. If this is not the case, you have to apply the correction as described in this SAP Note and then import the attached transport.

Most of the DB2 pureScale functions in the DBA Cockpit are available via remote monitoring. That is, you can monitor multiple DB2 pureScale systems by using a DBA Cockpit that contains the DB2 pureScale-specific enhancements and that uses remote database connections. For more information, see **SAP Note 1409540**.

After you have successfully imported the transport attached to **SAP Note 1409540**, you can access the additional monitoring functions by calling the DBA Cockpit (SAP transaction DBACOCKPIT) and choosing **DB2 pureScale Feature > <monitoring function>** in the navigation frame of the SAP GUI-based user interface of the DBA Cockpit.

More Information

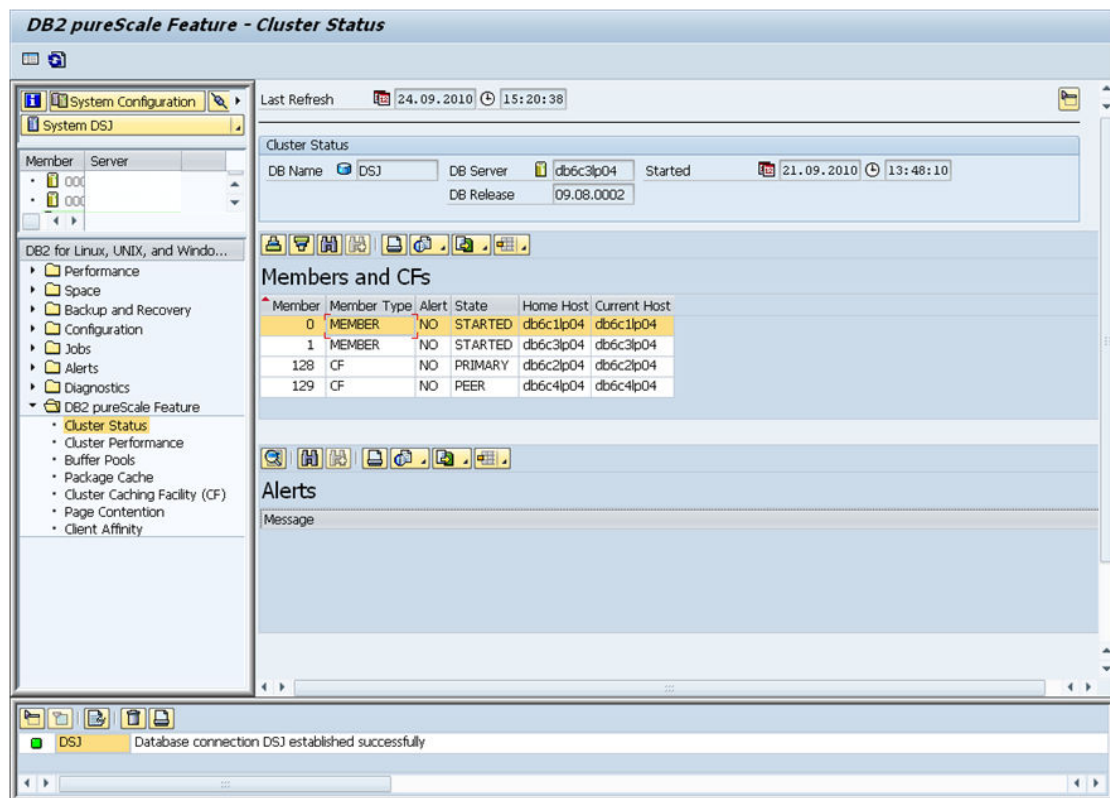
- [Monitoring the Cluster Status \[page 79\]](#)

- [Monitoring the Cluster Performance \[page 80\]](#)
- [Monitoring the Buffer Pools in the DB2 pureScale Cluster \[page 81\]](#)
- [Monitoring the Package Cache in the DB2 pureScale Cluster \[page 82\]](#)
- [Monitoring the Cluster Caching Facility \(CF\) \[page 84\]](#)
- [Checking for Page Contention \[page 85\]](#)
- [Checking the Client Affinity \[page 87\]](#)

5.4.2 Monitoring the Cluster Status

The *Cluster Status* screen provides an overview of the members or CFs that are part of the current DB2 pureScale cluster. In addition, information about current or home host names and alerts per member is displayed as well as the current state of a member.

You can display detailed alert messages by double-clicking the member to be analyzed in the *Members and CFs* overview table.

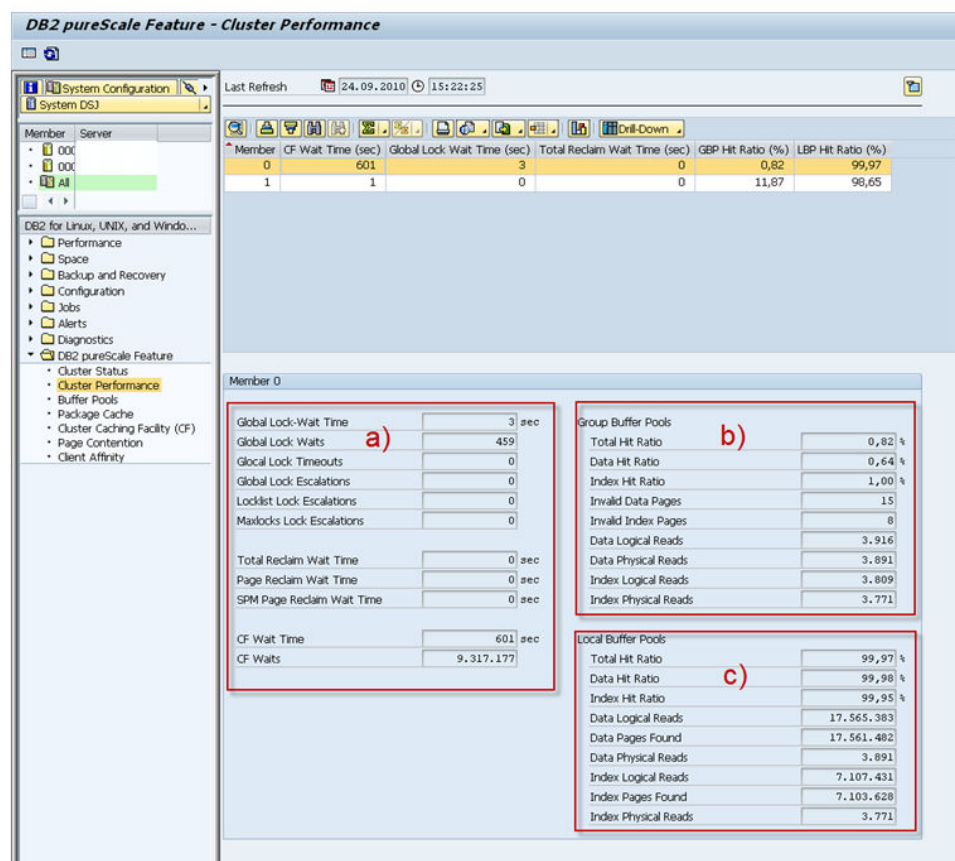


DBA Cockpit - Cluster Status Screen

5.4.3 Monitoring the Cluster Performance

The *Cluster Performance* screen provides information about current wait times and buffer pool hit ratios for each member in the DB2 pureScale cluster. In addition, global metrics that are drilled down to a specific member are displayed.

You can display additional member-related performance metrics by double-clicking a line in the overview table. The information is displayed in the detail area below as shown in the following figure:



DBA Cockpit - Cluster Performance Screen

The following information is displayed in the left group area (a):

- Global lock-wait times (physical, page-related wait times)
- Number of global lock waits
- Number of global lock timeouts and lock escalations
- Number of local locklist escalations and maxlock lock escalations
- Local page and space mapping page (SPM) reclaim wait times
- CF wait time (total) and number of CF waits

In the *Group Buffer Pools* area (b), information about the group buffer pool hit ratio, invalid pages as well as page or read statistics are displayed.

The hit ratio metrics are calculated as follows:

Hit Ratio Type	Formula
Total global hit ratio (global BP hit ratio)	$((\text{global logical data reads} + \text{global logical index reads}) - (\text{global physical data reads} + \text{global physical index reads})) / (\text{global logical data reads} + \text{global logical index reads}) * 100$
Global data hit ratio	$(\text{global logical data reads} - \text{global physical data reads}) / \text{global logical data reads} * 100$
Global index hit ratio	$(\text{global logical index reads} - \text{global physical index reads}) / \text{global logical index reads} * 100$

In the [Local Buffer Pools](#) area (c), information about the local buffer pool hit ratio is displayed as well as page or read statistics per member.

The hit ratio metrics are calculated as follows:

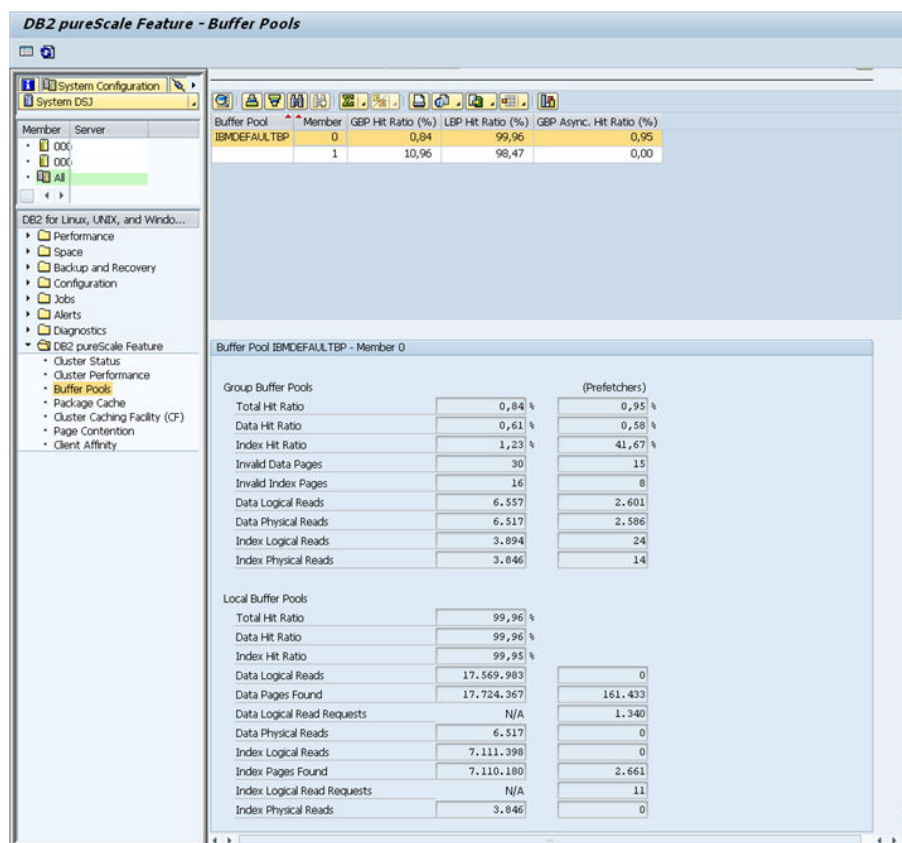
Hit Ratio Type	Formula
Total local hit ratio	$((\text{local logical data reads} + \text{local logical index reads}) - (\text{local physical data reads} + \text{local physical index reads})) / (\text{local logical data reads} + \text{local logical index reads}) * 100$
Local data hit ratio	$(\text{local logical data reads} - \text{local physical data reads}) / \text{local logical data reads} * 100$
Local index hit ratio	$(\text{local logical index reads} - \text{local physical index reads}) / \text{local logical local reads} * 100$

5.4.4 Monitoring Buffer Pools in the DB2 pureScale Cluster

The [Buffer Pool](#) screen provides information about computed hit ratio metrics of the buffer pools in the DB2 pureScale cluster. These metrics are grouped by the buffer pool name. Global and local hit ratios are drilled down to each member.

You can display detailed information about the metrics by double-clicking a line in the overview table.

The information is displayed in the *Buffer Pool <buffer pool name> – Member* detail area as shown in the following figure:



DBA Cockpit - Buffer Pools Screen

Note

The values displayed in the *Prefetchers* area are drilled down to the local and global prefetcher processes.

The metrics used are calculated in a similar way to the metrics described in [Monitoring the Cluster Performance \[page 80\]](#).

5.4.5 Monitoring the Package Cache in the DB2 pureScale Cluster

The *Package Cache* screen displays statement texts in the package cache that are related to a specific DB2 pureScale member as well the related package cache statistics.

When you access the *Package Cache* screen, a *Selection Criteria* dialog box appears where you can limit the result set by specifying the number of executions, the total activity time and the statement text.

The following information is displayed in the overview table:

Column	Description
<i>Member</i>	Member ID
<i>Statement Text</i>	Statement text
<i>Executions</i>	Number of executions per statement/member combination
<i>Total Activity Time</i>	Overall time of executions and prepares (in s)
<i>Avg. Activity Time</i>	Total activity time divided by the number of executions (in ms)
<i>CF Wait Time</i>	Overall CF wait time in seconds to process statement (including communication overhead with CF and lock waits)
<i>Global Lock Wait Time</i>	Global wait time in seconds that occurred on logical lock situations
<i>Total Reclaim Wait Time</i>	Sum of wait times for physical page reclaim in seconds
<i>GBP Hit Ratio (%)</i>	Percentage of the group buffer pool hit ratio that is related to statement

To display the complete statement text and other detailed metrics, such as additional global or local buffer pool statistics, double-click a line in the overview table.

The complete statement text is displayed as well as detailed package cache statistics as shown in the following figure:

DB2 pureScale Feature - Package Cache

Member	Statement Text	Executions	Total Activit	Avg. Activity Time (ms)	CF Wait Time (sec)	Glob	Total Recla	GBP Hit Ratio (%)
0	UPDATE SYSTOOLS.HMON_	629,300	137	0	0	0	0	100,00
0	UPDATE SYSTOOLS.HMON_	36	67	1.861	48	0	0	100,00
0	SELECT CREATE_TIME FROM	647,928	45	0	14	0	0	100,00
0	DELETE FROM SYSTOOLS.H	36	39	1.083	24	0	0	0,00
0	UPDATE SYSTOOLS.HMON_	35	33	943	23	0	0	100,00
0	SELECT STATS_TIME, INDE	629,300	22	0	0	0	0	0,00
0	SELECT COUNT(*) FROM S	35	12	343	8	0	0	100,00
0	SELECT COUNT(*) FROM S	35	12	343	8	0	0	100,00
0	SELECT TABNAME, BLOCINI	3,082	12	4	0	0	0	0,00

Statement Details

SELECT STATS_TIME, INDEXTYPE FROM SYSIBM.SYSINDEXES WHERE TBCREATOR = ? AND TENAME = ? WITH
UR

Coordinator Execution Time	22 sec	Group Buffer Pools	Total Hit Ratio	0,00 %
Number of Executions	629,300		Data Hit Ratio	0,00 %
Global Lock-Wait Time	0 sec		Index Hit Ratio	0,00 %
Global Lock Waits	0		Invalid Data Pages	0
Global Lock Timeouts	0		Invalid Index Pages	0
Global Lock Escalations	0		Data Logical Reads	532
Locklist Lock Escalations	0		Data Physical Reads	532
Maxlocks Lock Escalations	0		Index Logical Reads	49
			Index Physical Reads	49
Total Reclaim Wait Time	0 sec	Local Buffer Pools	Total Hit Ratio	0,00 %
Page Reclaim Wait Time	0 sec		Data Hit Ratio	0,00 %
SPM Page Reclaim Wait Time	0 sec		Index Hit Ratio	0,00 %
CF Wait Time	0 sec		Data Logical Reads	812,397
CF Waits	1,674		Data Pages Found	811,865

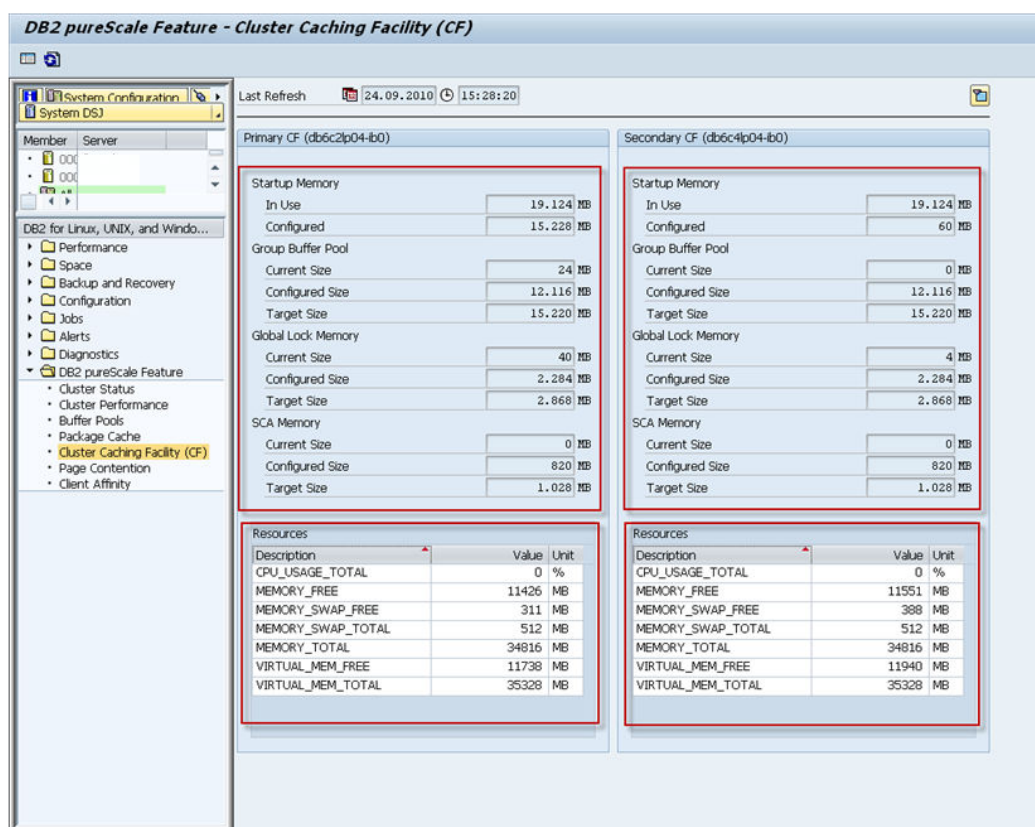
DBA Cockpit - Package Cache

The metrics used are calculated in a similar way to the metrics described in [Monitoring the Cluster Performance](#) [page 80].

5.4.6 Monitoring the Cluster Caching Facility (CFs)

The *Cluster Caching Facility (CF)* screen provides information about the current status of the cluster caching facilities (CFs) in the DB2 pureScale cluster as well as resource metrics that belong to a configured parameter setting per CF.

The information is displayed as shown in the following figure:



DBA Cockpit - Cluster Caching Facility (CF) Screen

In the [Resources](#) detail area, the overall host-related resources are displayed.

5.4.7 Checking for Page Contention

The [Page Contention](#) screen provides a detailed overview of events and statistics if competitive situations occur, for example, if different members try to access the same physical page at the same time.

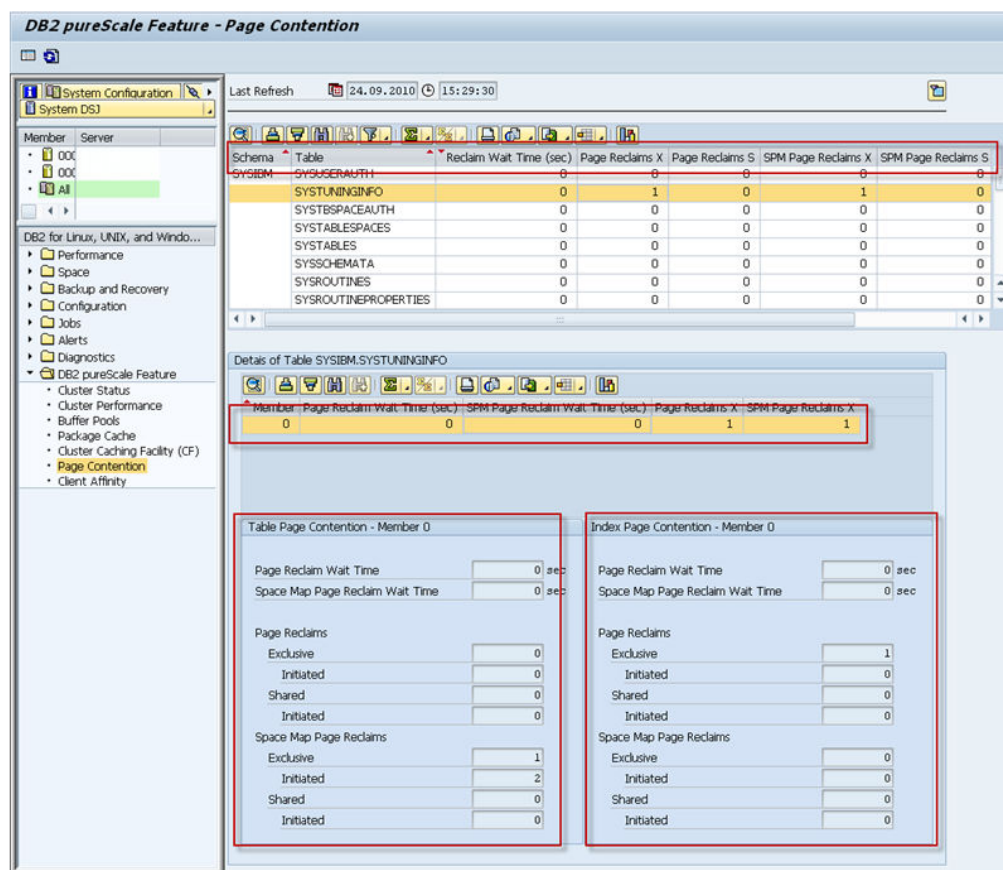
The overview table on the [Page Contention](#) screen displays the following key figures about page contention:

Column	Description
Schema	Name of database schema
Table	Name of table belonging to the schema
Reclaim Wait Time (sec)	Overall page reclaim wait times that are aggregated for each table

Column	Description
Page Reclaims X	Total number of page reclaims per table (exclusive)
Page Reclaims S	Total number of page reclaims per table (shared)
SPM Page Reclaims X	Total number of space map page reclaims per table (exclusive)
SPM Page Reclaims S	Total number of space map page (SMP) reclaims per table (shared)

To access detailed statistics per DB2 pureScale members, double-click a line in the overview table. To drill down the metrics for a specific member, double-click the member to be analyzed in the [Details of Tables](#) area.

The detailed information is displayed as in the following figure:



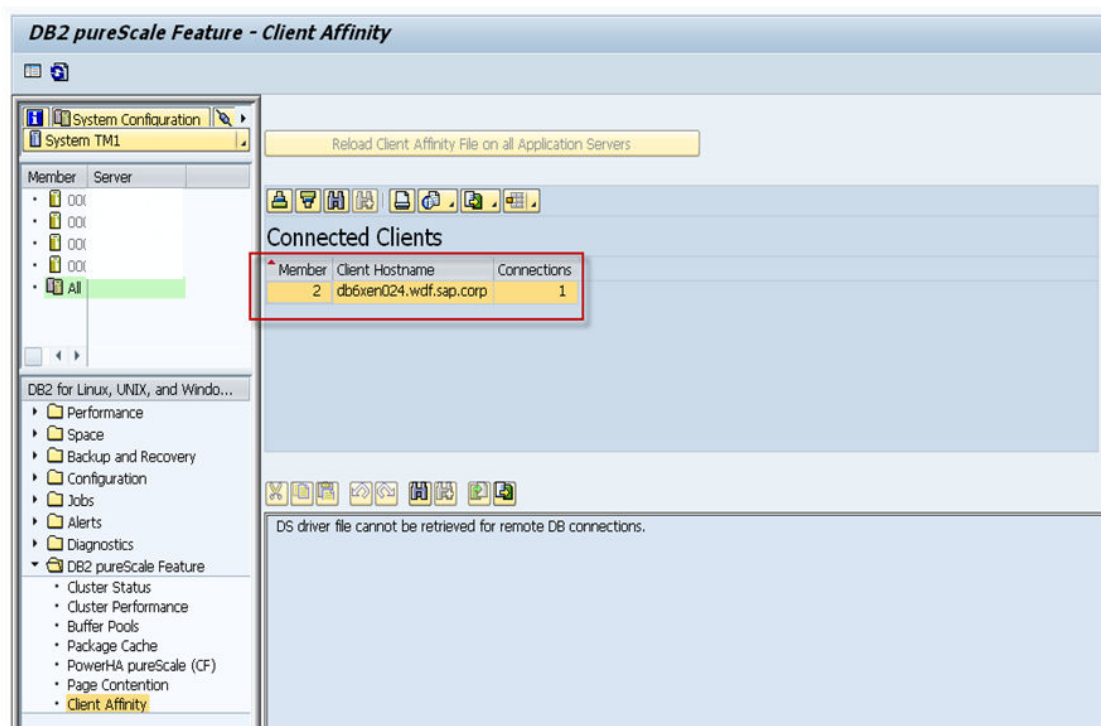
DBA Cockpit - Page Contention Screen

In the [Table Page Contention](#) and [Index Page Contention](#) area, the number of shared and exclusive reclaims and currently initiated reclaims is displayed. These values relate to the selected member and table.

5.4.8 Checking the Client Affinity

The *Client Affinity* screen analyzes the client affinity file and displays which database connection is currently open on a specific member.

To check the current configuration parameters in the `db2dsdriver.cfg` file, double-click a line in the overview table. This function is only supported if a local database connection exists (local monitoring).



DBA Cockpit - Client Affinity Screen

6 Removing a DB2 pureScale Installation

Use

To uninstall a DB2 pureScale cluster, you perform the following steps.

Note

Before you install DB2 pureScale again, it is important that you perform all the steps of this procedure. If you do not start in a clean environment, you might encounter problems.

Procedure

1. Back up your database.
2. Stop the SAP system.
3. If you want to uninstall the SAP instances, you can do so using the SAP installation tool.

Caution

You must **not** use the SAP installation tool to uninstall the database instance and the database. The DB2 pureScale instance and database must be uninstalled **manually**. You also have to make sure that the SAP installation tool does **not** delete any users during the uninstallation of the SAP system.

4. Log on to a DB2 pureScale member as user `db2<dbssid>` and stop the cluster using the following command:
db2stop
5. Check if all members and CFs are in status STOPPED using the following command:
db2instance -list
If this is not the case, you can repeat the **db2stop** command with the force option. If the state of a member or CF is ERROR, use the following commands to display and to clear the alert:
db2cluster -list -alert
db2cluster -clear -alert
For more information, see `db2cluster` command in the IBM DB2 Information Center.
6. Clean all DB2-related IPC resources on all hosts using the following command:
ipclean -a
7. Switch to user `root`.
8. Check on all hosts that no processes and no IPC resources are left over using the following commands:
ps -ef | grep db2<dbssid>
ipcs -a | grep db2
9. Terminate leftover processes manually and remove remaining IP resource with using the **ipcrm** command.
10. Drop the DB2 pureScale instance on one of the hosts of the DB2 pureScale cluster using the following command:
<inst_dir>/instance/db2idrop -g db2<dbssid>

i Note

You must **only** perform the following steps 11 to 17 on the hosts where you just dropped the instance.

11. Gather information about the GPFS cluster tie breaker, the GPFS domain, and the GPFS file systems that you require later in this procedure using the following commands:

1. To retrieve information about the GPFS cluster tie breaker, enter the following command:

```
<inst_dir>/bin/db2cluster -cfs -list -tiebreaker
```

The output of this command might look as follows:

The current quorum device is of type Disk with the following specifics: /dev/hdisk3.

2. To list the GPFS domain name, enter the following command:

```
<inst_dir>/bin/db2cluster -cfs -list -domain
```

The output of this command might look as follows:

Domain Name: db2cluster_20100217154241.wdf.sap.corp

3. To list the existing GPFS file systems, enter the following command:

```
<inst_dir>/bin/db2cluster -cfs -list -filesystem
```

The output of this command might look as follows:

FILE SYSTEM NAME	MOUNT_POINT
db2data	/db2/DSJ/sapdata
db2dbdir	/db2/DSJ/db2dsj
db2dump	/db2/DSJ/db2dump
db2fs1	/db2/instance_shared
db2log	/db2/DSJ/log_dir

12. The disk that contains the instance-shared directory (/db2/instance_shared) serves also as a tie breaker for the GPFS cluster. To remove this file system, you must first change the tiebreaker using the following commands:

```
<inst_dir>/bin/db2cluster -cfs -stop
```

```
<inst_dir>/bin/db2cluster -cfs -set -tiebreaker -majority
```

13. Start the GPFS cluster using the following command:

```
<inst_dir>/bin/db2cluster -cfs -start
```

14. To remove a file system with the **db2cluster** command, it must be empty. Remove file system with the instance-shared directory using the following commands:

```
rm -rf /db2/instance_shared
```

```
<inst_dir>/bin/db2cluster -cfs -delete -filesystem db2fs1
```

15. Remove the remaining file systems using the following commands:

```
rm -rf /db2/<DBSID>/sapdata
```

```
<inst_dir>/bin/db2cluster -cfs -delete -filesystem db2data
```

```
rm -rf /db2/<DBSID>/db2dsj
```

```
<inst_dir>/bin/db2cluster -cfs -delete -filesystem db2dbdir
```

```
rm -rf /db2/<DBSID>/db2dump
```

```
<inst_dir>/bin/db2cluster -cfs -delete -filesystem db2dump
```

```
rm -rf /db2/<DBSID>/log_dir
```

```
<inst_dir>/bin/db2cluster -cfs -delete -filesystem db2log
```

16. Remove the GPFS cluster using the following commands:

```
<inst_dir>/bin/db2cluster -cfs -stop
```

```
<inst_dir>/bin/db2cluster -cfs -delete -domain <domain_name>
```

i Note

As `<domain_name>`, use the name gathered in substep 2 of step 11.

17. Check if the system automation domain (`<SAMP_domain>`) still exists using the following command:
lsrpdomain
If it exists, the **lsrpdomain** command reports the domain name. Remove the domain using the following command:
rmrpdomain -f <SAMP_domain>
18. Remove the DB2 10.1 software by running the following command on each host:
<inst_dir>/install/db2_deinstall -a
19. Check that GPFS and TSA are removed correctly on all hosts:
 - For AIX, enter the following commands:
lslpp -l | grep sam
lslpp -l | grep gpfs
The first command should not return package names `sam.core`, `sam.rte`, or `sam.msg` (it might return packages with samples in the name, for example, `X11.samples.common`). The second command should not return any package names.
 - For Linux, enter the following commands:
rpm -qa sam
rpm -qa sam.*
rpm -qa gpfs*
All three commands should not return any package names.
20. Make sure that directory `/var/db2` is cleaned from all DB2 pureScale entries on all hosts. If no other DB2 installations reside on your hosts, remove directory `/var/db2` on all hosts.
21. Remove all log files that are related to DB2 pureScale from the `/tmp` directory on each host. This step is required so that if you install DB2 pureScale again, you see only newly generated log files. You have to delete the following files:
 - `db2*`
 - `ibm.db2.cluster*`
 - Any `installGPFS` or `uninstallGPFS` files
 - Any `installSAM` or `uninstallSAM` files.
22. Check the file `/etc/services` on each host and make sure that all port entries that are associated with DB2 (for example, `DB2CF_<instname>*`, `sapdb2<dbsid>`) are removed.

7 Disclaimer

By following links to IBM Documentation you are leaving the SAP product documentation and entering a site that is not hosted by SAP. By using the link, YOU AGREE that unless expressly stated otherwise in your agreements with SAP you are about to access an external webpage which is not part of SAP's offering:

(i) the content of the linked-to site and any further external site is not product documentation and you may not infer any product documentation claims against SAP based on this information;

(ii) the fact that SAP provides links to external sites does not imply that SAP agrees or disagrees with the contents and information provided on such sites. SAP does not guarantee the correctness of the information provided.

(III) SAP DOES NOT GIVE ANY REPRESENTATION REGARDING THE QUALITY, SAFETY, SUITABILITY, ACCURACY OR RELIABILITY OF ANY EXTERNAL WEBPAGE OR ANY OF INFORMATION, CONTENT AND MATERIALS PROVIDED THEREON;



(IV) YOU VISIT THOSE EXTERNAL WEBPAGES ENTIRELY AT YOUR OWN RISK. SAP SHALL NOT BE DIRECTLY OR INDIRECTLY RESPONSIBLE OR LIABLE FOR ANY DAMAGE OR LOSS CAUSED OR ALLEGED TO BE CAUSED BY OR IN CONNECTION WITH YOUR USE OF OR RELIANCE ON ANY CONTENT, GOODS OR SERVICES AVAILABLE ON OR THROUGH ANY SUCH LINKED WEBPAGE.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

© 2023 SAP SE or an SAP affiliate company. 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 SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company 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.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.