



Installation Guide | PUBLIC

Document Version: 2.4 – 2024-11-27

IBM Db2 High Availability Solution: IBM Tivoli System Automation for Multiplatforms

Content

- 1 Introduction. 4**
- 1.1 SA MP With Shared Disk. 5
- 1.2 SA MP With HADR. 6
- 1.3 High Availability (HA) Feature. 8
- 2 Document History. 10**
- 3 Naming Conventions. 12**
- 4 Planning. 14**
- 4.1 Hardware Requirements. 14
- 4.2 Software Requirements. 15
- 4.3 Virtual Host Name and Virtual IP Address. 16
 - Reusing the Host Name and IP Address of a Single Server. 16
 - Setting Up a New Virtual Host Name and IP Address. 17
- 4.4 Restrictions. 18
 - Restrictions for the HADR Setup. 18
- 5 Preparation. 20**
- 5.1 Setting Up File Systems for Shared Disk Scenario. 20
- 5.2 Setting Up the ssh/rsh Connection. 22
- 5.3 Changing the Password of the SAP Connect User. 23
- 5.4 Configuration of Log Archiving for HADR Setup. 24
- 6 Installation. 25**
- 6.1 Installation Overview. 25
- 6.2 Preparing the Global Host and Installing the SAP Central Services. 25
- 6.3 Downloading the Latest Db2 Software and SA MP License. 26
- 6.4 Installing the Database Server. 26
- 6.5 Setting Up a Standby Database Server Using Homogeneous System Copy. 27
- 6.6 Installing the Primary Application Server. 29
- 6.7 Installing the SA MP Software. 29
- 6.8 Setting Up the High-Availability Database Cluster. 30
 - Cluster Setup Tool. 30
 - Installing the SA MP License for the Database Cluster. 31
 - Copying the Files and Scripts for the Cluster Setup to a Local Directory. 31
 - Setting Up the Database Cluster. 31
 - Changing References to the Virtual Host Name or IP Address. 33

7	Post-Installation Activities	35
7.1	Validating the Database Cluster	35
	Checking the Database Cluster Setup	35
	Checking the Database Cluster Failover	37
7.2	Backing Up the SA MP Core Policy	38
8	Updating the Software	40
8.1	Updating the Database Fix Packs	40
8.2	Upgrading from Db2 V9.1 to V9.5 or V9.7	41
9	Uninstall	43
9.1	Uninstalling the Database Cluster	43
9.2	Uninstalling the SA MP Software	44
9.3	Deleting the Database Software	44
A	Appendix	45
A.1	SA MP Policy for Db2	45
	Cluster Configuration File sapdb2cluster.conf	45
	Cluster Setup Script sapdb2cluster.sh	47
A.2	Syntax of SA MP Setup Scripts	48
	Script prereqSAM	48
	Script installSAM	49
	Script uninstallsAM	51
A.3	Troubleshooting	52
	Symbol Resolution Failed for db2start on AIX	52
	Db2 Immediately Stops After the db2start/startdb/startjeedb Command	52
	Db2 Immediately Starts After the db2stop/stopdb/stopj2eedb Command	53
A.4	Graceful Maintenance Tool (Former Graceful Cluster Switch)	53
A.5	References	54
A.6	Disclaimer	54

1 Introduction

Purpose

This document describes how to set up a high-availability (HA) solution for IBM Db2 for Linux, UNIX, and Windows using IBM Tivoli System Automation for Multiplatforms (SA MP) when your operating system is AIX, Solaris SPARC (only up to and including Db2 10.5), or Linux (only up to and including Db2 11.5).

SA MP is a high-availability cluster solution that provides several monitoring mechanisms to detect system failures and a set of rules to initiate the correct action without any user intervention. The set of rules is called "policy", which describes the relationships between applications or resources. This provides SA MP with extensive up-to-date information about the system landscape so that it can restart the resource on the current node or move the whole application to another cluster node.

The SA MP license that is bundled with Db2 is restricted to the use with Db2. It must not be used to automate other components like the SAP application server. The license scope is described in the IBM Knowledge Center at https://www.ibm.com/support/knowledgecenter/SSEPGG_11.1.0/com.ibm.db2.luw.qb.server.doc/doc/r0051894.html .

Due to the database reconnect feature of SAP NetWeaver, a database server failure and a subsequent failover controlled by SA MP within the database cluster to another node is almost transparent to the clients. The running transactions of the work processes are terminated. The users of transactions within the affected dialog processes receive a message. SAP NetWeaver reconnects to the database server as soon as it is available again.

Setup Types for Database Clusters

There are two setup types to make a Db2 server highly available using SA MP:

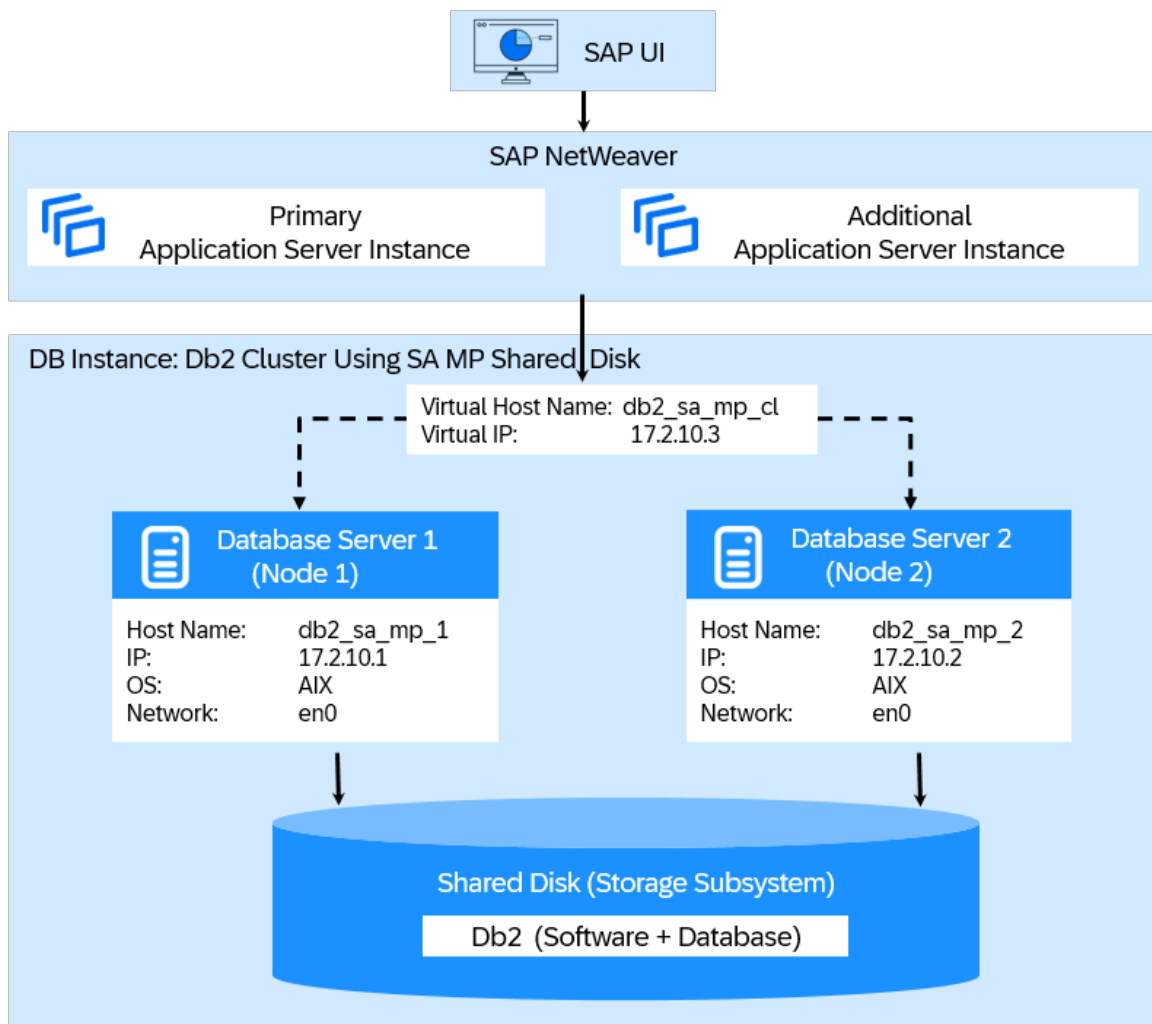
- [SA MP With Shared Disk \[page 5\]](#)
The database is located on a shared disk. The disk is shared between two servers. When the first server fails, the second server assigns the virtual IP to the network adapter, mounts the shared disk, and starts the Db2 instance and database.
- [SA MP With HADR \[page 6\]](#) (High Availability and Disaster Recovery):
HADR is a Db2 replication feature to make the database server highly available (shared nothing approach). In this scenario, you have two separate Db2 database servers: a primary and a standby database server. The servers are kept in sync and, in the event of failure, the standby database server takes over the workload.

1.1 SA MP With Shared Disk

Learn about the shared-disk approach to make a Db2 server highly available using SA MP.

The figure below shows a setup of two database servers that share the database software and the database itself on a shared disk. Both database servers have access to the shared disk, but only one of them is running at a time and is connected to the shared disk. The SAP NetWeaver application server instances are running on separate machines. SAP NetWeaver only knows the virtual host name or the virtual IP address of the database cluster. SA MP assigns the virtual IP address to the node on which the Db2 software is currently running.

In the event of failure of node 1, SA MP detects the failure and assigns the virtual IP address to a network adapter on node 2 (en0), mounts the shared disk and starts the Db2 instance on node 2. During the activation of the database, Db2 triggers a crash recovery to bring the database into a consistent and usable state. All open transactions from node 1 are rolled back and all committed transactions that were still in the memory when the crash occurred are completed. This process can take a while.



Shared-Disk Approach

1.2 SA MP With HADR

Learn about how you can use HADR to make the database server highly available ("shared nothing" approach).

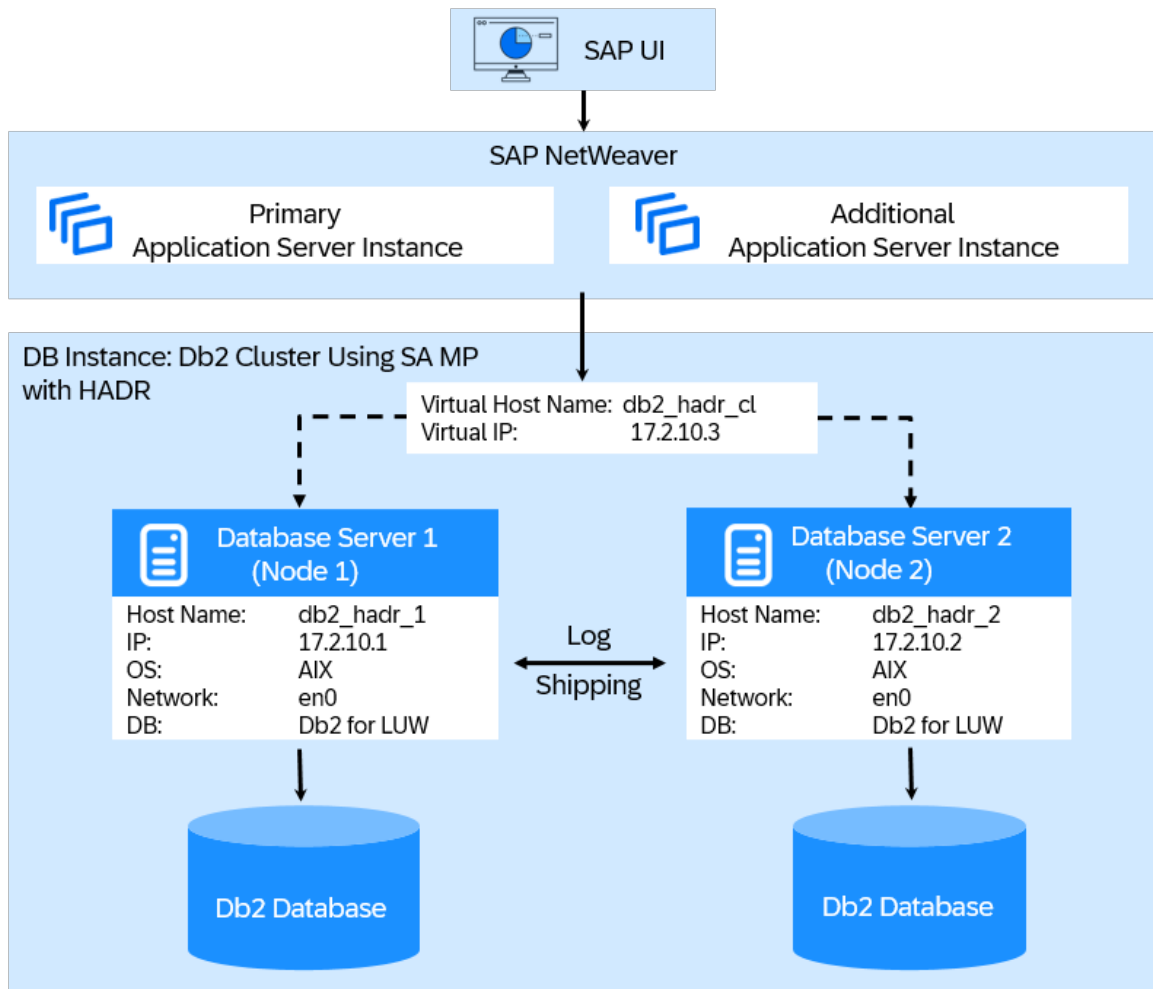
The figure below shows a setup of two database servers. Both database servers have their own storage and are up and running. In HADR, one server has the role of the primary server. This means that all clients are connected to this server. All transactions are written to logs. The log data is transferred via TCP/IP to the second database server, the standby server. This standby server updates the local database by rolling forward the transferred log files. So the standby server is kept in sync with the primary server.

HADR is a replication feature only. This means that it has no failure detection and no automation facilities. SA MP monitors the two database servers. When the primary database server crashes, SA MP initiates the HADR takeover by the standby server and also ensures that the virtual IP address is assigned to the new primary server.

Db2 also provides a client reroute feature, which ensures that all clients know the standby database server. In the event of a failure, the clients reroute their connections to the standby server. We recommend that you use a virtual IP address instead of the Db2 client reroute feature where possible. For more information, see SAP Note [1568539](#).

→ Recommendation

We recommend that you use SA MP for monitoring the database servers and moving the virtual IP address since this is transparent to all application servers and the monitoring infrastructure.



SA MP with HADR

If the primary server crashes, the takeover only requires a few seconds since the second database server is already running and the database is already in a consistent state.

However, the synchronization algorithm takes some additional time for updates. You can choose between the following synchronization modes for HADR:

- **SYNC** (synchronous)
This mode provides the greatest protection against transaction loss, and using it results in the longest transaction response time among the three modes. In this mode, log writes are considered successful only if logs have been written to log files on the primary database and if the primary database has received acknowledgment from the standby database that the logs have also been written to log files on the standby database. The log data is guaranteed to be stored at both sites.
- **NEARSYNC** (near synchronous)
While this mode has a shorter transaction response time than the synchronous mode, it also provides slightly less protection against transaction loss. In this mode, log writes are considered successful only if the log records have been written to the log files on the primary database and if the primary database has received acknowledgment from the standby system that the logs have also been written to the main

memory on the standby system. Data is only lost if both sites fail simultaneously and if the target site has not transferred all of the log data that it has received to nonvolatile storage.

- **ASYNC** (asynchronous)
This mode has a high chance of transaction loss if the primary system fails. It also has a short transaction response time. In this mode, log writes are considered successful only if the log records have been written to the log files on the primary database and have been delivered to the TCP layer of the primary system's host machine. Since the primary system does not wait for acknowledgment from the standby system, transactions might be considered committed when they are still on their way to the standby server.
- **SUPERASYNC** (super asynchronous)
This mode is available for IBM Db2 as of version 9.5 FP8. The super asynchronous mode complements the existing set of synchronization modes by ensuring that transactions can never be blocked or experience elongated response times due to network interruptions or congestion, therefore allowing transactions to be processed more quickly than in any other HADR synchronization mode.

→ Recommendation

We recommend that you use **NEARSYNC** for the HADR synchronization mode.

⚠ Caution

You must ensure that your network can rapidly manage the synchronization workload. If you do not ensure a rapid synchronization, the network becomes a bottleneck for the performance of the database server:

- The primary database has to wait for acknowledgment of the shipped log data in the **SYNC** and **NEARSYNC** modes.
- The primary database has to wait to send new messages to the standby server in the **ASYNC** mode when the TCP/IP buffers are full.

The distance between the two database servers can be much larger than in the shared disk setup since the two database servers only need a network connection to each other. So you can place them at different locations to protect your data, for example, against fire.

1.3 High Availability (HA) Feature

As of Db2 V9.5, Db2 can communicate with the associated cluster manager to change the cluster configuration according to actions and state changes of Db2.

For example, if the high availability (HA) feature is not enabled, you have to start and stop Db2 using cluster manager commands. With the Db2 HA feature, you can use native `db2start` or `db2stop` commands to manage your Db2 database. The commands change the state of the Db2 resource in the cluster configuration and prevent the cluster manager from bringing the resource back online or offline. From the user's point of view, the whole cluster handling is transparent. By using this feature, you reduce the risk of wrong configuration of your cluster when you change the configuration of Db2. For example, if you add a new container to your cluster on a new shared disk, Db2 adds this shared disk to your resource group to make sure that the failover also mounts the new disk at the other node.

In addition, the `db2haicu` tool (Db2 high-availability instance configuration utility) was introduced that you can use to enable HA and to configure the cluster manager in your system environment.

→ Recommendation

The `db2haicu` tool also provides an interactive mode for cluster setup. For SAP systems, you can use the `sapdb2cluster.sh` script to set up the cluster. The `sapdb2cluster.sh` uses `db2haicu` for the cluster setup and ensures that the cluster is set up in a way that `startdb` and `stopdb` scripts work correctly. We recommend that you use the `sapdb2cluster.sh` script.

2 Document History

⚠ Caution

Before you start the implementation, make sure you have the latest version of this document. It's available at https://help.sap.com/viewer/db6_samp on SAP Help Portal.

Here's an overview of the most important document changes:

Document History

Version	Date	Description
2.4	2024-11-27	Update due to new database version Db2 12.1
2.3.2	2024-07-16	Update of some graphics; no content changes
2.3.1	2024-03-25	Update of Setting Up a Standby Database Server Using Homogeneous System Copy [page 27]
2.3	2023-03-28	Info: SAP systems running on IBM Db2 10.1 and lower are out of mainstream maintenance. Update of section References [page 54]
2.2	2019-11-05	IBM Db2 11.5 supported; minor updates
2.01	2018-06-19	Small update of References [page 54]
2.0	2018-02-22	Minor updates only, for example, links to SAP Service Marketplace replaced by links to SAP Help Portal . and information about Db2 9.1 and 9.5 removed because these database versions are no longer supported.
1.9	2017-06-22	Graphics update (no content-related changes)
1.8	2016-08-01	Update due to new DB2 for Linux, UNIX, and Windows Version 11.1
1.71	2014-08-29	Minor correction of graphic in chapter <i>Setting Up a New Virtual Host Name and IP Address</i>
1.7	2014-05-27	Update with minor corrections and changes, for example, regarding access to the cluster setup tool; DB2 10.5 added.
1.6	2013-12-20	<ul style="list-style-type: none">Chapter <i>Graceful Cluster Switch (Micro-Outage)</i> was changed and information about the graceful maintenance tool was added.Chapter <i>Setting Up the High-Availability SAP Cluster</i> was removed.
1.5	2012-10-22	Update: Addition of DB2 for Linux, UNIX, and Windows Version 10.1, introduction of SAP cluster components

Version	Date	Description
1.04	2010-11-30	Update of section <i>Updating the Database Fix Packs</i>
1.03	2010-07-22	Updated version (including graceful cluster switch)
1.02	2008-03-10	Updated version
1.0	2006-08-11	Initial version

3 Naming Conventions

In this document, the following naming conventions apply:

IBM Terminology

Term	Abbreviation
IBM Db2 12.1 for Linux, UNIX, and Windows	Db2 12.1
IBM Db2 11.5 for Linux, UNIX, and Windows	Db2 11.5
IBM Db2 11.1 for Linux, UNIX, and Windows	Db2 11.1
IBM Db2 10.5 for Linux, UNIX, and Windows	Db2 10.5
IBM Db2 10.1 for Linux, UNIX, and Windows	Db2 10.1 (out of mainstream maintenance)
IBM Db2 V9.7 for Linux, UNIX, and Windows	Db2 V9.7 (out of mainstream maintenance)
IBM Db2 High Availability Disaster Recovery	HADR
IBM Tivoli System Automation for Multiplatforms	SA MP
IBM Reliable Scalable Cluster Technology	RSCT

With Db2 11.1, IBM introduced the concept of Modification Packs. A Modification Pack (also referred to as Mod or MP) introduces new functions to the Db2 product. For the IBM Db2 Modification Packs and Fix Packs, we mostly use abbreviations such as Db2 11.1 Mod 2 Fix Pack 2, or even shorter, simply Db2 11.1 MP2 FP2.

Renaming: IBM DB2 for Linux, UNIX, and Windows is now IBM Db2

IBM has changed its database name from IBM DB2 for Linux, UNIX, and Windows to simply IBM Db2 (with a lowercase 'b' in Db2). In older SAP publications, you'll still find the old product name, but in more recent documentation, we use the new term, sometimes extended by 'for Linux, UNIX, and Windows' to avoid confusion with other products of the Db2 family, such as Db2 for z/OS or Db2 for i.

SAP Terminology

As of SAP NetWeaver 7.1, the following terminology changes apply:

Up to and Including SAP NetWeaver 7.0	New Naming as of SAP NetWeaver 7.1
Central instance	Primary application server instance
Dialog instance	Additional application server instance

Variables

Variable	Description
<DBSID>	Name of the database in upper case
<dbsid>	Name of the database in lower case
	System ID of the SAP system in upper case
<sapsid><SAPSID>	System ID of the SAP system in lower case

4 Planning

4.1 Hardware Requirements

To set up the Db2 server in an HA environment using the free two-node license of SA MP, the following hardware requirements must be met:

Requirement Type	Requirement
Server Hardware	<p>Two separate servers for the database.</p> <p>Both servers must meet the Db2 requirements as described in <i>System requirements for IBM Db2 for Linux, UNIX, and Windows</i> at http://www.ibm.com/support/docview.wss?uid=swg27038033 .</p>
Physical Disks	<p>For shared disk setup only:</p> <p>One shared disk, which you can mount from both database servers. No synchronization is required since the disk is mounted by only one server at a time. NFS is not supported.</p> <div><p>→ Recommendation</p><p>We recommend that you use Storage Area Network (SAN).</p></div>
Memory	<p>For shared disk setup, both servers should have at least the memory configured for Db2 since the configuration is moved between the two nodes.</p> <p>For HADR setup, both servers must have the same amount of memory since buffer pool operations are also replayed by the standby database server.</p>
Disk Space	<p>300 MB additional free disk space on each node within the HA cluster for the SA MP Software</p> <div><p>ⓘ Note</p><p>Make sure that the directories <code>/usr/sbin</code>, <code>/opt</code> and <code>/var</code> have at least 100 MB of free disk space available.</p></div>

→ Recommendation

We recommend that you use two identical machines for the cluster setup. Since in a shared disk setup, the database manager configuration and the database registry are transferred during movement of Db2 from one node to the other, the configuration must suit both of them.

In the HADR setup, the primary server has to wait until the standby database server sends the acknowledgment back to the primary server. This means that the standby server also impacts the performance of the primary server.

4.2 Software Requirements

Cluster Setup Tool

To set up a high-availability solution for IBM Db2 for Linux, UNIX, and Windows using IBM Tivoli System Automation for Multiplatforms (SA MP), you require the cluster setup tool that you can download from SAP Note [960843](#). The cluster setup tool comprises the script `sapdb2cluster.sh` (also referred to as *cluster setup script* in this document) and a configuration file.

Note

Before you start the installation, see SAP Note [960843](#). Always use the latest version of the cluster setup tool that is attached to this SAP Note.

Supported Operating Systems

The following operating systems are supported with SA MP:

- AIX 64 bit
- Only up to and including Db2 11.5: Red Hat Enterprise Linux
- Only up to and including Db2 11.5: SuSE Linux Enterprise Server
- Only up to and including Db2 10.5: Solaris on SPARC

For more information about the required operating system versions for your SAP system running on IBM Db2, see the product availability matrix (PAM) for your SAP system release at <https://support.sap.com/en/release-upgrade-maintenance.html>. Make sure that the operating system that you use is both supported for your SAP system and IBM Db2 version and for SA MP.

Additional Requirements

- Perl must be installed on the database servers.
- On Linux, Public Domain Korn Shell (`pdksh`) must also be installed.
- All nodes must run on the same operating system.
- All host machines must have the same operating system patch level.
- You must use the Db2 CLI driver for client connectivity.

4.3 Virtual Host Name and Virtual IP Address

The virtual host name and the virtual IP address are used to access the cluster. The virtual host name is a reference on the DNS server or to the virtual IP address in the `/etc/hosts` files. The cluster binds the virtual IP address to the active cluster node. All clients have to refer either to the virtual host name or directly to the virtual IP address. Therefore, the clients always connect to the node of the cluster where the clustered Db2 or the SAP instance is currently running.

To set up such an environment, you can choose between the following options:

- [Reuse the host name and IP address of a single server as virtual host name and IP address \[page 16\]](#).
- [Set up a new virtual host name and IP address \[page 17\]](#).

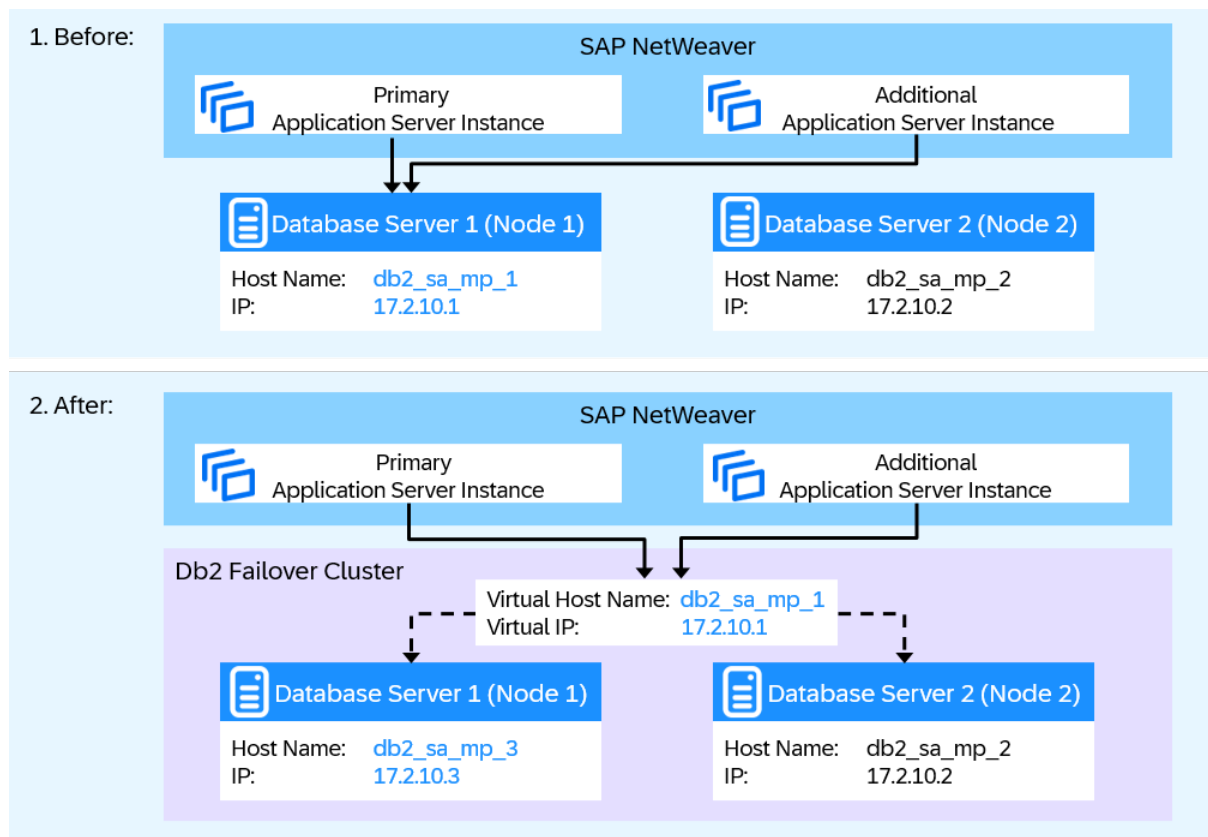
For a new installation or system copy, we always recommend the second approach. The first option is only recommended for an existing SAP system that you want to make highly available.

4.3.1 Reusing the Host Name and IP Address of a Single Server

To reuse the host name and the IP address of the single database server as virtual host name and virtual IP address, you have to replace the physical host name and the physical IP address of the single database server with a new host name and a new IP address because you specify the old host name and the old IP address of the database server as the new virtual host name and virtual IP address in the cluster setup configuration file `sapdb2cluster.conf`.

The advantage of this approach is that you do not have to change the references to the database host or SAP instance host of your SAP system landscape and of additional monitoring tools. However, additional services that are installed on the server are not available anymore on the specified host name or IP address.

The following figure provides an example of such a reuse before and after the database cluster has been set up:



Reuse of Host Name and IP Address of Single Database Server

The host name db2_sa_mp_1 of the single database server is reused as virtual host name; the IP address 17.2.10.1 is reused as virtual IP address. To avoid naming conflicts, the database server needs a new network identity. Therefore, the database host name is changed to db2_sa_mp_3 and the IP address to 17.2.10.3.

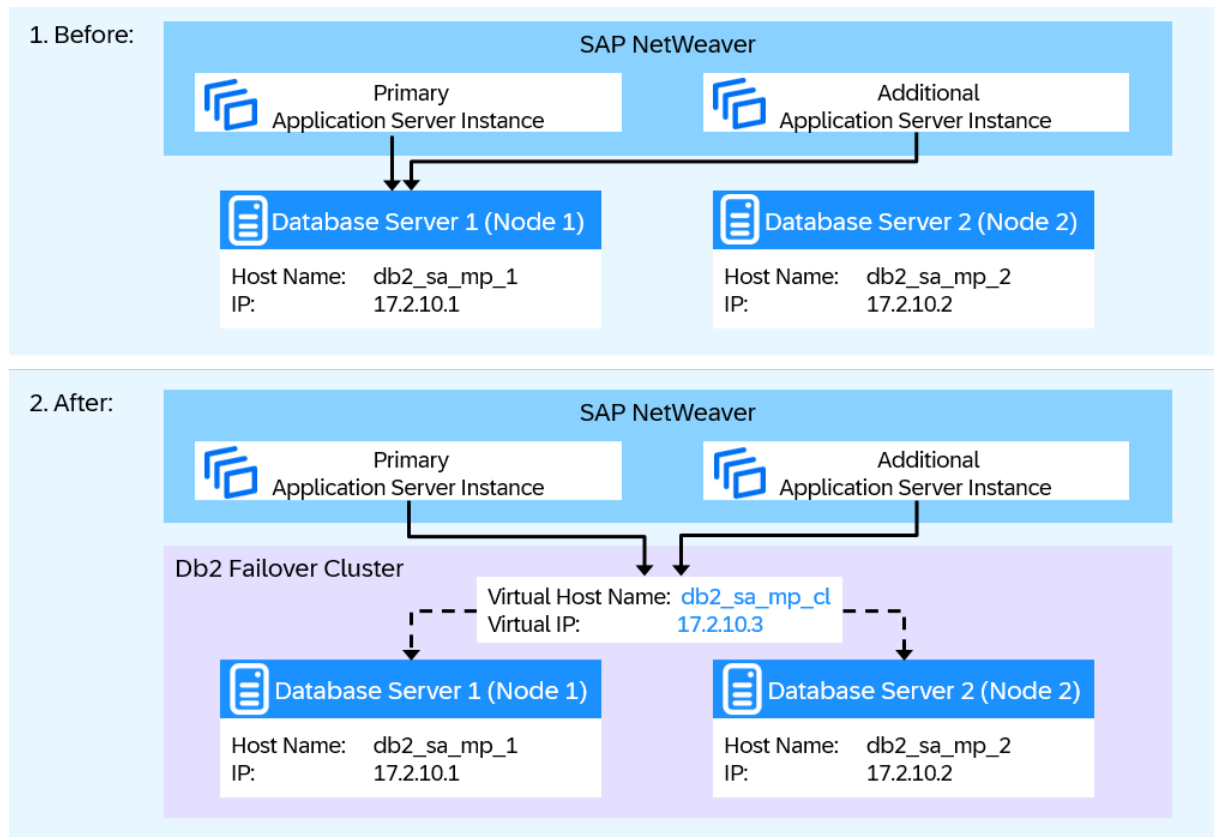
The host name and IP address of the second database server (node 2) remain the same and both database servers are now addressed with the new virtual address and IP address.

You have to change the entries of the host name in db2nodes.cfg to the new physical host name of node 1. This file is located in the Db2 instance directory.

4.3.2 Setting Up a New Virtual Host Name and IP Address

To set up a new virtual host name and a new virtual IP address, you specify the new host name and the new IP address in the cluster setup configuration file sapdb2cluster.conf.

The following figure shows an example of how to set up a new virtual host name or IP address before and after the database cluster has been set up:



Setup of New Virtual Host Name and IP Address

The database server keeps its host name `db2_sa_mp_1` and its IP address `17.2.10.1`. A new virtual host name `db2_sa_mp_cl` and a new virtual IP address `17.2.10.3` are chosen for the failover cluster. If you already have an SAP NetWeaver system running on this database server (node 1), you have to change all references to the database host in the SAP system landscape and in the monitoring tools. Although the `sapdb2post.sh` script automatically replaces some of the references in your SAP system landscape, you have to make some of the changes manually.

For more information, see [Changing References to the Virtual Host Name or IP Address \[page 33\]](#).

4.4 Restrictions

4.4.1 Restrictions for the HADR Setup

For the HADR setup, the following restrictions apply:

- HADR is not supported in a partitioned database environment.
- The primary and standby databases must have the same operating system version and the same version of the Db2 database system, except for a short time during a rolling upgrade.

- Log archiving can only be performed by the current primary database.
- You can use Self-Tuning Memory Manager (STMM) only on the current primary database.
- Backup operations are not supported on the standby database.
- Non-logged operations, such as changes to database configuration parameters and to the recovery history file, are not replicated to the standby database.
- Load operations with the `COPY NO` option specified are not supported.
- HADR does not support the use of raw I/O (direct disk access) for database log files. If HADR is started via the `START HADR` command or the database is activated (restarted) with HADR configured and raw logs are detected, the associated command fails.

5 Preparation

Make sure that you perform the following steps before you start the installation:

1. [Set up file systems for shared disk scenario \[page 20\]](#).
2. [Set up the ssh/rsh connection \[page 22\]](#).
3. [Change the password of the SAP connect user \[page 23\]](#).
4. [Configure the log archiving for HADR setup \[page 24\]](#).

5.1 Setting Up File Systems for Shared Disk Scenario

Before you install SA MP, you have to ensure the required file systems are set up as follows:

- All directories under `/db2` must be located on the shared disk but not on NFS shares. For example, the following file systems are required for a standard SAP system with a database that uses Db2's automatic storage management:

File System	Description
<code>/db2/db2<dbsid></code>	Contains the home directory of <code>db2<sapsid></code>
<code>/db2/<DBSID>/log_dir</code>	Contains at least the online database log files
<code>/db2/<DBSID>/db2dump</code>	Contains Db2 diagnostic log files, Db2 dump files, and further service engineer information
<code>/db2/<DBSID>/db2<dbsid></code>	Contains the local database directory
<code>/db2/<DBSID>/saptemp1</code>	Contains the temporary tablespace(s)
<code>/db2/<DBSID>/sapdata1</code>	SAP data for container type database managed space (DMS) FILE or for use of Db2's automatic storage management

To make sure that either of the mentioned file systems is stored on the shared disk, choose one of the following options (where you replace `<dir>` with any of the directories mentioned here):

- Mount the shared disk on `/db2` and create an appropriate directory, for example, `/<dir>`.
- Create a partition on the shared disk and mount the partition on `/db2/<dir>`.
For more information about the recommended file system sizes, see your relevant installation guide at https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html and database upgrade guide at <https://help.sap.com/viewer/p/DB6>.
- All directories under `/usr/sap/<SAPSID>/<INSTANCE>` must be located on the shared disk but not on NFS shares.

- Make sure that all required mount points to the shared disk are not mounted at system boot time and that the file systems check at boot time is disabled.
 - On Linux, you have to set the `noauto` option. Furthermore, set the sixth column `fs_passno` to `0` in file `/etc/fstab` as follows:


```
/dev/sdc1 /db2 reiserfs noauto,acl,user_xattr 0 0
/dev/sdc2 /usr/sap/<SAPSID>/<INSTANCE> reiserfs noauto,acl,user_xattr 0 0
```
 - On AIX, you have to set the options `check` and `mount` to `false`.

```
/db2:
dev      = /dev/fslv01
vfs      = jfs2
log      = /dev/loglv00
mount    = false
check    = false
options  = rw
account  = false
/usr/sap/<SAPSID>/<INSTANCE>:
dev      = /dev/fslv02
vfs      = jfs2
log      = /dev/loglv02
mount    = false
check    = false
options  = rw
account  = false
```

⚠ Caution

You must make sure that none of the file systems are mounted automatically during system reboot. SA MP is the only authority allowed to mount the file systems.

- If you have to mount the file systems listed above manually, you **must** disable SA MP Automation first, that is, set it to **manual** mode or lock the resource. Otherwise, SA MP unmounts the file systems immediately. To set SA MP to manual mode or to automation mode, use the following commands as user `root`, `db2<dbsid>` or `<sapsid>adm`:

Mode	Command
Manual Mode (disabled):	<code>samctrl -MT</code>
Automation Mode (enabled):	<code>samctrl -MF</code>

To lock the resource or resource group, use the following commands as user `root`, `db2<dbsid>` or `<sapsid>adm`:

- Lock (automation disabled):
 - Resource group:


```
rgreq -o lock <resource group>
```

🔗 Example

```
rgreq -o lock db2_db2ha3_0-rg
```

- Resource:


```
rgmbrreq -o lock <resource>
```

❖ Example

```
rgmbrreq -o lock IBM.Application:db2mnt-db2-rs
```

- Unlock (automation enabled):
 - Resource group:
rgreq -o unlock <resource group>

❖ Example

```
rgreq -o unlock db2_db2ha3_0-rg
```

- Resource:
rgmbrreq -o unlock <resource>

❖ Example

```
rgmbrreq -o unlock IBM.Application:db2mnt-db2-rs
```

- All types of local file systems are supported by SA MP, for example, `reiserfs`, `jfs2`, and so on.

5.2 Setting Up the ssh/rsh Connection

For the installation of SA MP, you have to enable `ssh` or `rsh` login for user `root` without password authentication. This enables user `root` to connect to all nodes of the cluster including itself on any node of the cluster.

To configure the nodes of the cluster, the SA MP cluster setup script uses `ssh` or `rsh`.

📌 Note

By default, the setup script uses `ssh` on Linux and `rsh` on AIX.

Procedure

Setting Up the Connection Using ssh

For example, to use RSA authentication instead of password authentication for `ssh`, proceed as follows:

1. Log on to the first node of the cluster as user `root`.
2. Generate an RSA key for user `root` by entering the following command:

```
ssh-keygen -t dsa
```

You are now asked to specify the file name for the key and a passphrase:

- Enter the default file name for the key:
~/.ssh/id_dsa
- Do not specify any passphrase.

3. To enable key authentication for `ssh` on the same node, enter the following command:
`cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys`
4. To transfer the public key to the second node, enter the following command:
`scp ~/.ssh/id_dsa.pub root@<hostname>:~/id_dsa.pub`
5. Log on to the second node of the cluster as user `root`.
6. To enable key authentication for `ssh` on the second node, enter the following command:
`cat ~/id_dsa.pub >> ~/.ssh/authorized_keys`
7. To remove the public key file from the home of user `root` on the second node, enter the following command:
`rm ~/id_dsa.pub`
8. To add both cluster nodes to the list of known hosts, enter the following command:
`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>, <host2_ip> >> ~/.ssh/known_hosts`
9. Follow the same procedure in the opposite way on the second node.

Setting Up the Connection Using `rsh`

If you want to use `rsh` and avoid password authentication, add the two nodes to the `.rhosts` file in the home directory of the user `root`.

1. Log on to first node of the cluster as user `root`.
2. Add the two nodes to the file `.rhosts` and the user `root` separated by a space.

❄ Example

If you use the example setup shown in the [Introduction \[page 4\]](#), the content of `.rhosts` should look as follows:

```
db2_sa_mp_1.company.net root
db2_sa_mp_1 root
17.2.10.1 root
db2_sa_mp_2.company.net root
db2_sa_mp_2 root
17.2.10.2 root
```

5.3 Changing the Password of the SAP Connect User

If you change the password of the database connect user with the `dscdb6up` tool, the password in the `dscdb6.conf` file and the operating system password of the connect user on the database host are changed.

On the standby database server, the password is not changed. As a result, the SAP systems cannot connect to the standby database server if Db2 is started on this node. If you use Network Information Service (NIS) for central user management, the password is not changed in the NIS directory.

This means that if you change the password of connect user `sap<sapsid>` or `<sapsid>adm`, you have to manually change the operating system password on the standby database server in the failover cluster or in the NIS directory.

5.4 Configuration of Log Archiving for HADR Setup

To configure log archiving for HADR setup, we recommend that you configure both the primary and the standby database to have automatic log retrieval capability from all log archive locations. Both the primary and the standby database must be able to retrieve log files from all the log archive locations to which either of the databases might archive log files.

The log archiving is only performed by the primary database. If you change the HADR roles of the database servers or if a failure occurs, the new primary database is responsible for log archiving. If you have set up different log archive locations, your logs might be archived twice and, in the case of local or remote catch-up, you might have to manually copy the archived logs from the old primary server to the active log location of the new primary server.

6 Installation

6.1 Installation Overview

The following sections describe how to install and set up the database failover cluster.

To set up SA MP in a distributed SAP NetWeaver system, proceed as follows:

1. [Prepare the global host and install the central services \[page 25\]](#).
2. Install the database instance:
 1. [Download the latest Db2 software and SA MP license \[page 26\]](#).
 2. [Install the database server \[page 26\]](#).
 3. [Set up a standby database server using homogeneous system copy \[page 27\]](#).
3. [Install the primary application server \[page 29\]](#).
4. [Install the SA MP software \[page 29\]](#).
5. Set up the high-availability database cluster:
 1. [Install the SA MP Db2 license for the database cluster \[page 31\]](#).
 2. [Set up the database cluster \[page 31\]](#).
 3. [Change references to the virtual host name or IP address \[page 33\]](#).
6. Perform [post-installation activities \[page 35\]](#).

Note

You can also install SA MP after you have completed a distributed SAP NetWeaver installation. In this case, you only need to install the second database server as described in [Installing the Database Server \[page 26\]](#) and continue to set up the high availability cluster as described in [Setting Up the Database Cluster \[page 31\]](#).

If your database server is not addressed using a virtual host name, you have to change all references to it as described in [Changing References to the Virtual Host Name or IP Address \[page 33\]](#).

6.2 Preparing the Global Host and Installing the SAP Central Services

If you want to make the SAP central services instance highly available, you must install the instance on a virtual host name. You can do this by using the parameter `SAPINST_USE_HOSTNAME` when you start the SAP installer. For more information about the installation, see your relevant installation guide at https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html.

6.3 Downloading the Latest Db2 Software and SA MP License

To be able to install the latest version of Db2 and SA MP, download the required software from SAP Support Portal at <https://launchpad.support.sap.com/#/softwarecenter/database> > Database and Database Patches > IBM Db2.

6.4 Installing the Database Server

Prerequisites

If you want to set up an HADR cluster and you do not use NIS/NFS for central user management, you must ensure that the group and user IDs are identical on the cluster nodes. For the shared disk setup, the script `sapdb2cluster.sh` copies the user groups, users, and home directories to the second node.

Note

- For SAP system releases up to and including SAP Basis 4.6D, the connect user is `sapr3`.
- The user `sapadm` is used for the SAP Host Agent.

Procedure

Install the database server as described in the appropriate installation guide at https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html.

→ Recommendation

We recommend that you install Db2 on the shared disk in the home directory `/db2/db2<dbSid>/db2_software` of the `db2<dbSid>` user. This is the default for SAP NetWeaver 7.0 SR3 and higher.

You have to install the OEM license once on one of the nodes.

6.5 Setting Up a Standby Database Server Using Homogeneous System Copy

If you set up the switchover cluster based on the Db2 feature HADR, you have to create a standby database as a copy of the primary database. You can use the SAP homogeneous system copy for setting up the standby database server.

Context

⚠ Caution

The following section **only** applies to the HADR-based cluster setup. For shared-disk setup, you must **omit** this section.

Keep in mind that you do the final setup of the HADR cluster using the script `sapdb2cluster.sh` which is provided by SAP.

Procedure

1. Log on to the first node as user `db2<dbsid>`.
2. Enable archiving for the logs, for example, to a disk folder using the following command:

```
db2 "update db cfg for <dbsid> using LOGARCHMETH1 DISK:/db2/<DBSID>/log_archive/"
```

⚠ Caution

Both the primary and the standby database need to be able to retrieve log files from all the log archive locations to which either of the databases might archive log files.

The log archiving is only performed by the primary database. If you change the HADR roles of the database servers or if a failure occurs, the new primary database is responsible for log archiving. If you have set up different log archive locations, your logs might be archived twice and, in the case of local or remote catch-up, you might have to copy the archived logs manually from the old primary server to the active log location of the new primary server.

To configure log archiving for HADR, we recommend that you configure both the primary and the standby database to have automatic log retrieval capability from all log archive locations.

3. Create the directory `/db2/<DBSID>/backup` using the following command:

```
mkdir /db2/<DBSID>/backup
```

4. Create a full database backup:

```
db2 "backup db <dbsid> to /db2/<DBSID>/backup compress"
```

5. Log on to the second node as user `root`.
6. Create the directory `/db2/<DBSID>/backup`:

```
mkdir /db2/<DBSID>/backup
```

7. Transfer the backup from the first node to the second node to the directory `/db2/<DBSID>/backup`:

```
scp root@<hostname1>:/db2/<DBSID>/backup/* /db2/<DBSID>/backup/
```

8. Start the SAP software provisioning manager (SWPM).

Note

If you are using virtual host names, the virtual IP address must be bound to the current host before starting the installer.

1. Choose one of the following options, depending on which SAP NetWeaver version you are using:
 - For SAP NetWeaver 7.3EHP1-7.52: **▶ <Product> ▶ IBM Db2 for Linux, UNIX, and Windows ▶ System Copy ▶ Target System ▶ Distributed System ▶ <Technical_Stack> ▶ Database Instance ▶**
 - For SAP NW 7.0-7.03: **▶ <Product> ▶ Software Life-Cycle Options ▶ IBM Db2 for Linux, UNIX, and Windows ▶ System Copy ▶ Target System Installation ▶ Distributed ▶ <Technical_Stack> ▶ Database Instance ▶**
2. For the parameter settings mode, choose *custom* mode.
3. For the copy method, choose *homogeneous system copy* so that you can use your backup to restore the backup on the standby server.
4. When you're asked for passwords, make sure to use the same passwords that were used on the primary database server.
5. On the *System Copy for HADR Purposes* screen, check the box that you're using system copy with backup/restore for HADR purposes. This will cause the SWPM to exit at the point where you must restore your database.
6. On the *Database Communication* screen, make sure that you have the same database communication parameter values as on the primary database server.
7. When you reach the exit step to restore the database for the homogeneous system copy, choose *Cancel* to exit the SWPM.
9. Restore the database on your secondary host using the backup that you created on the primary host.

Note

You do not need to restart the SWPM after you have restored your database because the subsequent installation phases have already been executed on the primary database server.

1. Log on to the second node as user `db2<dbsid>`.
2. Create the `sapdata*` and `saptmp*` directories, using the same directory structure as on the primary node. If you're using recovery control files, create this directory, too.
3. Restore the database backup using the following command:


```
db2 "restore db <dbsid> from /db2/<DBSID>/backup replace history file"
```
4. Verify that your database is in rollforward pending state using the following command:


```
db2 "get db cfg for <dbsid>" | grep Rollforward
```

6.6 Installing the Primary Application Server

The installation of SAP application servers is not covered in this document. For more information, see the respective installation guide that you use to install the database server. You can find the installation guides at https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html.

6.7 Installing the SA MP Software

You can find the SA MP software in the following directory on the RDBMS DVD of IBM Db2:

```
<dvd_mount>/<platform>/ESE/disk1/db2/<platform>/tsamp
```

Procedure

1. Log on to the first node as user `root`.
2. Insert and mount the database software DVD to `<dvd_mount>`.
3. To change to the mount directory, enter the following command:
`cd <dvd_mount>/<platform>/ESE/disk1/db2/<platform>/tsamp`
4. To check that the prerequisites are met, enter the following command:
`./prereqSAM`
5. To start the installation, enter the following command:
`./installSAM --noliccheck`

Note

You can ignore the warning `"CT_MANAGEMENT_SCOPE not set"`.
It is automatically set during the cluster setup.

6. Install the SA MP software on the second node.

More Information

[Script prereqSAM \[page 48\]](#)

[Script installSAM \[page 49\]](#)

6.8 Setting Up the High-Availability Database Cluster

6.8.1 Cluster Setup Tool

To set up the high availability cluster, you use the cluster setup tool. It comprises the cluster setup configuration file and the setup scripts that are described in the following:

Name	Description
File <code>sapdb2cluster.conf</code>	Contains all data that is required for the setup and maintenance of the cluster. It's created by the script <code>sapdb2cluster.sh</code> . During the creation process, <code>sapdb2cluster.sh</code> asks for all information required for your cluster setup and finally generates <code>sapdb2cluster.conf</code> .
Script <code>sapdb2cluster.sh</code>	Sets up the cluster using the data contained in file <code>sapdb2cluster.conf</code> . During the setup process, the script <code>sapdb2cluster.sh</code> stops the database server, consecutively starts and stops both nodes, and finally activates the database on the original node. The script is also used to reset or delete the cluster configuration. For more information about the syntax of <code>sapdb2cluster.sh</code> , see Cluster Setup Script <code>sapdb2cluster.sh</code> [page 47] .

The following files are replaced during the setup process and are required later.

File Name	Description
<ul style="list-style-type: none"><code>startdb</code><code>startj2eedb</code><code>stopdb</code><code>stopj2eedb</code><code>sampdbcmd</code>	Replace the former SAP scripts <code>startdb</code> and <code>stopdb</code> because you have to start and stop the database using the SA MP scripts listed here.

Related Information

[Installing the SA MP License for the Database Cluster \[page 31\]](#)

[Copying the Files and Scripts for the Cluster Setup to a Local Directory \[page 31\]](#)

[Setting Up the Database Cluster \[page 31\]](#)

[Changing References to the Virtual Host Name or IP Address \[page 33\]](#)

6.8.2 Installing the SA MP License for the Database Cluster

Prerequisites

- You have the SA MP license file at hand. If you have licensed Db2 from SAP, download the appropriate license file as described in SAP Note [816773](#).
- Check if the Db2 license is installed. If not, install it as described in [SAP Note 816773](#).

Procedure

1. Log on to the first node as user `root`.
2. Move the license file to the home directory of user `root`:
 - Linux: `/root`
 - AIX: `/home/root`

The cluster setup tool transfers the license to all other database cluster nodes and installs the license on each node.

6.8.3 Copying the Files and Scripts for the Cluster Setup to a Local Directory

Download the latest version of the SA MP cluster setup files and scripts from [SAP Note 960843](#).

Before you set up the cluster, you have to extract all files to a local directory, for example, `~/smp_setup`, as follows:

Procedure

1. Log on to the first node as user `root`.
2. Create a directory for the scripts using the following command:
`mkdir ~/smp_setup`
3. Download and extract the setup files archive to this directory.

6.8.4 Setting Up the Database Cluster

You set up the high availability cluster using the cluster setup configuration file and the setup script `sapdb2cluster.sh` (see also [The Cluster Setup Tool \[page 30\]](#)).

Prerequisites

- You have to use at least version 6.31 of the cluster setup scripts that are attached to [SAP Note 960843](#).
- If you want to [reuse the host name and IP address of a single server \[page 16\]](#), you have to change the host name and the IP address of the database server before creating the cluster.

Procedure

Generating the Cluster Configuration

1. Log on to the first database cluster node as user `root`.
2. Mount the shared disk (for HADR setup, you do not have to mount a shared disk).
3. Change to the directory to which you copied the cluster setup tool and the setup files and scripts:

```
cd ~/samp_setup
```
4. Execute the cluster setup script using the following command:

```
./sapdb2cluster.sh
```
5. In the main menu, choose option *1 – DB2 Database Cluster*.
6. In the database menu, choose option *1 – Create, Show or Edit Database Configuration*.
This starts the configuration generator, which creates the configuration file `sapdb2cluster.conf` (or uses the existing file).
7. Provide all required parameters that could not be detected automatically.

⚠ Caution

Make sure that you keep the file `sapdb2cluster.conf` because it is required to reset or uninstall the cluster.

Setting Up the Database Cluster

The final cluster setup is performed by the setup cluster script `sapdb2cluster.sh` based on the information contained in file `sapdb2cluster.conf` as follows:

1. Log on to the first database cluster node as user `root`.
2. Mount the shared disk (for HADR setup, you do not have to mount a shared disk).
3. Change to the directory to which you copied the scripts using the following command:

```
cd ~/samp_setup
```
4. Execute the cluster setup script – assuming that you have created the cluster configuration file in `~/samp_setup` – using the following command:

```
./sapdb2cluster.sh
```
5. In the main menu, choose option *1 – Create, Show or Edit Database Configuration*.
6. In the database menu, choose option *2 – Create Database Cluster*.
This starts the cluster creation process, which is based on the information specified in the configuration file `sapdb2cluster.conf`.

📄 Note

The cluster setup tool creates the log file `sapdb2cluster.log` with detailed information about the creation process.

6.8.5 Changing References to the Virtual Host Name or IP Address

Note

You can omit this procedure if one of the following applies:

- You are reusing the old host name as the virtual host name.
- You have a new installation, that is, you directly set up the failover cluster after the installation of the database instance before installing the application servers.

If you have chosen to [set up a new virtual host name and a new virtual IP address \[page 17\]](#), you must perform the following steps. Although the script `sapdb2post.sh` performs some of the steps, you have to perform other steps manually as described in the following.

Process

1. `sapdb2cluster.sh` changes the default profile:
In the default profile `/sapmnt/<SAPSID>/profile/DEFAULT.PFL`, the database host name has to be replaced for `SAPDBHOST` and `j2ee/dbhost`.
2. If the Db2 CLI driver is used, `sapdb2post.sh` changes the Db2 CLI driver settings in file `db2cli.ini`, that is, the database host name is replaced.
3. `sapdb2cluster.sh` changes the environment of users `<sapsid>adm` and `sap<sapsid>`:
The environment of the users `<sapsid>adm` and `sap<sapsid>` has to be adapted because most of the environment source files contain the host name as part of their file name. For all other nodes of the cluster, `sapdb2cluster.sh` creates the appropriate links as shown in the following example.

Example

```
.dbenv_db2_sa_mp_2.sh -> .dbenv_db2_sa_mp_1.sh
```

4. You change the catalog entries on all SAP NetWeaver instances:
 1. Log on to the host of an SAP NetWeaver instance as user `db2<dbsid>`.
 2. Generate a list of all catalog entries and look for an entry `NODE<SAPSID>` by entering the following command:

```
db2 "list node directory"
```
 3. Uncatalog the reference to the old host name using the following command:

```
db2 "uncatalog node NODE<SAPSID>"
```
 4. Catalog the reference to the new virtual host name using the following command:

```
db2 "catalog tcpip node NODE<SAPSID>  
remote <virtual_hostname> server <service_name>"
```

where `<servicename>` stands for `sapdb2<SAPSID>`.
 5. Repeat this procedure for all SAP NetWeaver instances.
5. You change the JDBC URL:

Note

The following steps **only** apply for ABAP+Java and Java systems.

1. Log on to the host of an SAP NetWeaver instance as user `root`.
2. Change to the directory of the configuration tool using the following command:
`cd /usr/sap/<SAPSID>/<Instance>/j2ee/configtool`
3. Run the configuration tool using the following command:
`./configtool.sh`
4. In the left frame, choose *security store*.
5. In the right frame, choose the key *jdbc/pool/<SAPSID>/url*.
6. Change the host name in the JDBC URL to the virtual host name.

🔗 Example

- URL with the old host name:
`jdbc:db2://db2_sa_mp_1:5912/TSP:deferPrepares=0`
- URL with the new virtual host name:
`jdbc:db2://db2_sa_mp_c1:5912/TSP:deferPrepares=0`

7. Choose *Add*.
 8. To save your changes, click the disk icon in the upper left corner.
 9. Close the configuration tool.
 10. Restart the Java instance.
6. Change the references in remotely monitored systems.

📌 Note

Make sure that you also change the references in your monitoring tools to the new virtual host name.

7 Post-Installation Activities

After you have successfully completed the installation and the cluster setup, perform the following post-installation activities:

1. Validate the data cluster:
 - [Check the database cluster setup \[page 35\]](#).
 - [Check the database cluster failover \[page 37\]](#).
2. [Back up the SA MP core policy \[page 38\]](#).

7.1 Validating the Database Cluster

7.1.1 Checking the Database Cluster Setup

After the installation, you have to check if the failover cluster is working properly.

Procedure

1. Log on to one of the cluster nodes as user `root`, `<sapsid>adm`, or `db2<dbsid>`.
2. To check if both nodes are online, enter the following command:

lsrpnode

The following is an example output that can look as follows (RSCT versions can differ):

Name	OpState	RSCTVersion
db2_sa_mp_1	Online	2.4.8.0
db2_sa_mp_2	Online	2.4.8.0

3. To check if the domain is online, enter the following command:

lsrpdomain

The following is an example output that can look as follows (RSCT versions can differ):

Name	OpState	RSCTActiveVersion	MixedVersions	TSPort	GSPort
sap_tsmdb2...	Online	2.4.8.0	No	12347	12348

4. To check the resource groups, enter the following command:

lsrg

The following is an example output that can look as follows:

```
Resource Group names:  
db2_db2tsm_0-rg
```

The naming convention for the resource group name is as follows:

`db2_db2<dbsid>_0-rg`

5. To check the status of the resource group, enter the following commands:

- Shared disk setup:
`lsrg -g db2_db2<dbsid>_0-rg`
- HADR setup:
`lsrg -g db2hadr_<dbsid>-rg`

The output looks as follows:

```

Displaying Resource Group information:
For Resource Group "db2_db2tsm_0-rg".
Resource Group 1:
  Name                = db2_db2tsm_0-rg
  MemberLocation      = Collocated
  Priority             = 0
  AllowedNode         = ALL
  NominalState       = Online
  ExcludedList        = {}
  Subscription        = {}
  Owner               =
  Description         =
  InfoLink            =
  ActivePeerDomain    = sap_tsmdb2
  OpState           = Online
  TopGroup            = db2_db2tsm_0-rg
  ConfigValidity      =
  TopGroupNominalState = Online
  
```

Since most of the SA MP commands are asynchronous, there are two states for the resource group:

State of Resource Group	Description
<i>NominalState</i>	Indicates the status that the resource group is supposed to have
<i>OpState</i>	Indicates the status that the resource group currently has

If you just started the resource group, the *NominalState* is *Online* but the *OpState* is *Offline* or *Pending Online*. If the resource group is online, *OpState* turns into *Online* as well. Shutting down the resource group causes the same effect but vice versa.

If the resource group is not yet online, start it using the following command:

db2start

To receive more information about resource members of the cluster, use the following SA MP command:

`lsrg -Ab -m -V -g db2_db2<dbsid>_0-rg`

In addition, you can use the **lssam** command to display all registered resources and the status on the node where the script is issued.

The output of the **lssam** command looks as follows:

```

Online IBM.ResourceGroup:db2_db2djd_0-rg Nominal=Online
  |- Online IBM.Application:db2_db2djd_0-rs
    |- Offline IBM.Application:db2_db2djd_0-rs:db6xen003
    '- Online IBM.Application:db2_db2djd_0-rs:db6xen004
  |- Online IBM.ServiceIP:db2ip_10_17_202_162-rs
    |- Offline IBM.ServiceIP:db2ip_10_17_202_162-rs:db6xen003
    '- Online IBM.ServiceIP:db2ip_10_17_202_162-rs:db6xen004
  '- Online IBM.Application:db2mnt-db2-rs
    |- Offline IBM.Application:db2mnt-db2-rs:db6xen003
    '- Online IBM.Application:db2mnt-db2-rs:db6xen004
  
```

6. For HADR, check the state of the Db2 HADR database servers:

1. Log on to one of the cluster nodes as user `db2<dbsid>`.
2. Check if the database server is up and running.
3. Check the HADR state:

```
db2pd -db <dbsid> -hadr
```

In the sample output, the database ID is `SJH`. The output looks as follows:

```
Database Partition 0 -- Database SJH -- Active -- Up 2 days 22:47:43
HADR Information:
Role      State                SyncMode HeartBeatsMissed   LogGapRunAvg
(bytes)
Primary Peer                Nearsync 0                191289728
ConnectStatus ConnectTime          Timeout
Connected   Mon Mar 12 15:48:06 2007 (1173732486) 120
LocalHost    LocalService
is0081       SJH_HADR_1
RemoteHost   RemoteService        RemoteInstance
is0082       SJH_HADR_2           db2sjh
PrimaryFile  PrimaryPg            PrimaryLSN
S0000018.LOG 6030                 0x00000001F253E876
StandByFile  StandByPg            StandByLSN
S0000018.LOG 6030                 0x00000001F253E876
```

7. Check if you can start and stop Db2 in the cluster.

- To start and stop the members of the resource group – taking into account the relationship between the members – use the following command:

```
rgmbrreq -o [start|stop|cancel] <resource_class>:<rgmbr>
```

- To explicitly start and stop the Db2 instance, only use the following command:

```
rgmbrreq -o [start|stop|cancel]
          IBM.Application:db2_db2<dbsid>_0-rs
```

7.1.2 Checking the Database Cluster Failover

You can check the cluster failover by moving the Db2 instance from one node to the other.

Procedure

1. Log on to one cluster node as user `root`, `<dbsid>adm`, or `db2<dbsid>`.
2. To move Db2 to another node, enter the following command:

- Shared disk setup:

```
rgreq -o move db2_db2<dbsid>_0-rg
```

- HADR setup:

You could use the same command as in the shared disk setup, but we recommend that you use the `db2 takeover` command as user `db2<sid>`:

```
rgreq -o move db2hadr_<dbsid>-rg
db2 takeover hadr on db <dbsid>
```

Result

In the shared disk setup, SA MP shuts down Db2 on the specified node, unmounts the shared disk, removes the virtual IP address from the network adapter of this node, and sets up everything on the other node in the cluster.

In the HADR setup, SA MP removes the virtual IP address from the primary host, issues the command `db2 takeover HADR on db <db sid>`, and finally assigns the virtual IP to the new primary server. In this way, the standby server becomes the new primary server and vice versa.

You can monitor the process using the command `getstatus` or `lssam -top`.

However, this method does not check the monitoring and failure detection features of SA MP. To check the failover behavior, you can do the following:

- Shut down the operating system on the node where the Db2 instance is currently running.
- Alternatively, kill the Db2 instance using the `db2_kill` command.

SA MP detects that the resources of the resource group are not online anymore and starts all required resources on the other node.

7.2 Backing Up the SA MP Core Policy

Use

SA MP stores all its information about the cluster in a policy that contains all definitions of resources, resource groups, and their relationships. Therefore, this policy contains the core configuration of the SA MP cluster. After the cluster setup, we recommend that you back up the SA MP core policy. In this way, you can easily restore the initial cluster setup if you accidentally change the SA MP core policy.

Procedure

Backing Up the SA MP Core Policy

1. Log on to a cluster node as user `root`.
2. Enter the following command:
`sampolicy -s /bkp/samp/bkpPolicy.xml`

Restoring the SA MP Core Policy

1. Log on to a cluster node as user `root`.
2. Enter the following command:
`sampolicy -a /bkp/samp/bkpPolicy.xml`

ⓘ Note

The restore process replaces the current policy by the one specified and finally activates this policy.

8 Updating the Software

8.1 Updating the Database Fix Packs

Use


Follow the steps below to update the database Fix Packs in an SA MP environment. After a Db2 database upgrade or Fix Pack update, you have to update the Db2 HA precanned policy files in SA MP that are stored in the subdirectory of the Db2 software directory `ha/τsa`.

These policy files are required by SA MP for monitoring, starting and stopping the database server, and transferring the registry.

Procedure

1. Log on to the first node as user `<sapsid>adm`.
2. To shut down the Db2 instance, enter the following command:
`stopdb`
3. Log on to the first node as user `root`.
4. To check if Db2 has been stopped, enter the following command:
 - Shared disk setup:
`lsrg -g db2_db2<dbsid>_0-rg`
 - HADR setup:
`lsrg -g db2hadr_<dbsid>-rg`
In the output, check if the *OpState* is *Offline*.
5. To disable SA MP, enter the following command:
`samctrl -MT`
6. If the Db2 software is located on the shared disk, mount it.
7. Apply the Db2 Fix Pack to the first node as described in the appropriate SAP Note for the Fix Pack installation:

Db2 Version	SAP Note Number
Db2 11.1	2303756
Db2 10.5	1871003
Db2 10.1	1708037

Db2 Version	SAP Note Number
Db2 V9.7	1363169 

8. If the Db2 software is located on the shared disk, unmount it.
9. Log on to the second node as user `root`.
10. If the Db2 software is located on the shared disk, mount it.
11. If the Db2 software is not located on the shared disk, you must apply the Db2 Fix Pack to this node as described in the appropriate **SAP Note** (see the table in step 7).

Caution

You must not update the Db2 instance (`db2iupdt`) again because it was already updated in step 6.

12. To remove the old HA policy, enter the following command:


```
rm /usr/sbin/rsct/sapolicies/db2/*
```
13. To update to the new HA policy, enter the following command:


```
cp <DB2_software_dir>/ha/tsa/* /usr/sbin/rsct/sapolicies/db2/.
```
14. If the Db2 software is located on the shared disk, unmount it.
15. To enable SA MP, enter the following command:


```
samctrl -MF
```
16. Log on to the first node as user `<sapsid>adm`.
17. To start the Db2 instance, enter the following command:


```
startdb
```
18. To check if Db2 is running again, enter the following command:
 - Shared disk setup:


```
lsrg -g db2_db2<sapsid>_0-rg
```
 - HADR setup:


```
lsrg -g db2hadr_<sapsid>-rg
```

OpState and *Nominal State* of the resource group come online in a short time.

Caution

If you share the Db2 software, you have to install the Fix Pack only once, but you must replace the Db2 HA precanned policy on both nodes.

8.2 Upgrading from Db2 V9.1 to V9.5 or V9.7

Use

You use the following procedure to upgrade your database from Db2 V9.1 to Db2 V9.5 or V9.7.

Note

Note that Db2 V9.1 and V9.5 are no longer supported.

Prerequisites

If you enabled your Db2 V9.1 installation for high availability, you have to delete the cluster before you upgrade. Otherwise, you have to reset the cluster afterwards. This step is necessary because `db2haicu` has to set up the cluster to enable the Db2 HA feature.

Procedure

1. Log on to the **first node** as user `<sapsid>adm`.
2. To shut down the Db2 instance, enter the following command:
`stopdb`
3. Log on to the **first node** as user `root`.
4. Change to the directory of SA MP setup scripts using the following command:
`cd ~/samp_setup`
5. Run the cluster setup tool in interactive mode:
`./sapdb2cluster.sh -f ~/samp_setup/sapdb2cluster.conf`
6. Choose *4 – Delete Database Cluster*.
7. Perform the upgrade to Db2 V9.5 or higher using the appropriate database upgrade guide.

Note

You have to ensure the SA MP software is updated according to the relevant upgrade guide, too.

8. Download and extract the latest setup files archive to the `~/samp_setup/` directory.
9. Run the cluster setup tool in interactive mode:
`./sapdb2cluster.sh -f ~/samp_setup/sapdb2cluster.conf`
10. Choose *2 – Create Database Cluster*.

Note

If your cluster configuration file has to be updated as well, `sapdb2cluster.sh` performs this step and prompts for additional parameters.

9 Uninstall

To uninstall the cluster and the software from the two database hosts, you have to perform the following steps:

1. [Uninstall the database cluster \[page 43\]](#).
2. [Uninstall the SA MP software \[page 44\]](#).
3. [Delete the Db2 software \[page 44\]](#).

9.1 Uninstalling the Database Cluster

To uninstall the cluster and to separate the two nodes, use the `sapdb2cluster.sh` script.

Procedure

1. Log on to the first node where you initially transferred the SA MP scripts to as user `root` (for more information, see [Setting Up the Database Cluster \[page 31\]](#)).
2. If the second node becomes the server where the database is running (that is, the single database server), move the directory `~/smp_setup` to the other node and log on to the other node as user `root`.
3. Change to the directory using the following command:

```
cd ~/smp_setup
```
4. Run the cluster setup tool in interactive mode:

```
./sapdb2cluster.sh -f ~/smp_setup/sapdb2cluster.conf
```
5. Choose *4 – Delete Database Cluster*.
6. Change to the home directory using the following command:

```
cd ~
```
7. Delete the setup scripts using the following command:

```
rm -R ~/smp_setup
```

Result

The cluster has been uninstalled and the node where you performed the above mentioned steps is able to run Db2. Since the virtual host name and virtual IP address are no longer valid, you have to replace all references to it with the host name or IP address of the single database server.

The `sapdb2cluster.sh` script has updated the references in the SAP profile (`db2cli.ini` file), and has restored the SAP `startdb/stopdb` scripts when you deleted the Db2 resource group. In addition, you have to replace all references listed for manual actions as described in [Changing the References to Virtual Host Name or IP Address \[page 33\]](#).

You can also remove the users `db2<dbsid>`, `sap<sapsid>`, and `<sapsid>adm` from the host where the database instance is no longer running.

9.2 Uninstalling the SA MP Software

You uninstall the SA MP software using the `uninstallSAM` script from the SA MP software directory on the DVD.

Procedure

1. Log on to the first node as user `root`.
2. Insert and mount the database software DVD to `<dvd_mount>`.
3. Change to the directory where the SA MP software is located using the following command:
`cd <dvd_mount>/<platform>/ESE/disk1/db2/<platform>/tsamp`
4. To uninstall the SA MP software, enter the following command:
`./uninstallSAM`

9.3 Deleting the Database Software

Depending on the database version that you are using, you have to delete the Db2 software as described in the appropriate installation guide that you can find in SAP Help Portal at https://help.sap.com/viewer/30839dda13b2485889466316ce5b39e9/CURRENT_VERSION/en-US/c8ed609927fa4e45988200b153ac63d1.html.

A Appendix

A.1 SA MP Policy for Db2

A.1.1 Cluster Configuration File `sapdb2cluster.conf`

The `sapdb2cluster.conf` policy file contains all parameters that are required to set up the Db2 cluster and the SAP cluster components properly. The following is a list of all parameters in the cluster configuration file.

Note

The cluster setup tool creates the cluster configuration file and fills the parameters with values. You do not need to edit the configuration file manually.

General Cluster Configuration Parameters

Parameter	Description	Value
<code>TSA_DOMAIN_NAME</code>	Domain name of the cluster	<code>sap_<dbsid>db2</code>
<code>TSA_LICENSE_FILE</code>	Location of the downloaded SA MP license file	<code><home-root>/sam32.lic</code>
<code>TSA_TIEBREAKER_IP_ADDRESS</code>	IP address of the network tie breaker (see <i>Network Tiebreaker</i> section in this chapter)	<code><ip-address></code>
<code>TSA_DISK_HEARTBEAT</code>	Disk heartbeat, which separates network failures from node failures	<code>[<raw-device> <LVID> <MPATH> <PVID>]</code>
<code>TSA_REMOTE_CMD</code>	Remote shell that you use for the setup	<code>[ssh rsh]</code>
<code>TSA_USER_LIST</code>	Comma-separated list of users with the authority to start or stop the cluster and single resources in the cluster.	<code>db2<dbsid>,<sapsid>adm</code>
<code>TSA_USER_GROUP_NAME</code>	The name of the group of TSA operators	<code>sagrp</code>

Db2 Cluster Configuration Parameters

Parameter	Description	Value
<code>DB2_HOSTNAME_LIST</code>	Comma-separated list of the host names, separated by commas	<code><host1>[,<host2>]</code>

Parameter	Description	Value
DB2_CLUSTER_TYPE	Cluster type: shared disk, HADR, or none	[SAMP HADR NONE]
DB2_INST_DIR	Software installation directory for Db2	<db2-software-path>
DB2_COMM_PORTS	Comma-separated list of Db2 communication ports	[<name>:<port_num>[,<name>:<port_num>]
DB2_HADR_PORTS	Comma-separated list of the two Db2 HADR log shipping ports	[<name>:<port_num>[,<name>:<port_num>]
DB2_HADR_SYNC_MODE	Synchronisation mode for HADR	[SYNC NEARSYNC ASYNC SUPERASYNC]
DB2_HA_HOSTNAME	Virtual host name of the database cluster	<virtual-hostname>
DB2_HA_IP_ADDRESS	Virtual IP address of the database cluster	<virtual-ip>
DB2_DB2INSTANCE	Db2 instance owner	<db2sid>
DB2_DIR_LIST	Comma-separated list of Db2 directories on shared disks	[<mnt1>[,<mnt2>]]
DB2_LV_LIST	Comma-separated list of logical volumes of raw shared disks (AIX only)	[<lv1>[,<lv2>]]
DB2_NETWORK_INTERFACE_LIST	Comma-separated list of public HA network interfaces for Db2, at least one per node	<net>:<ip>:<sub>:<host>[,<net>:<ip>:<sub>:<host>]
DB2_GROUP_LIST	List of user groups	<group>:<gid>:<local_group>[,<group>:<gid>:<local_group>]
DB2_USER_LIST	Comma-separated list of users	<user>:<uid>:<home_dir>:<local_home_dir>:<shell>:<local_user> <prim_group>[:<sec_group>] [,<user>:<uid>:<home_dir>:<local_home_dir>:<shell>:<local_user> <prim_group>[:<sec_group>]]

SAP Cluster: Configuration Parameters

Parameter	Description	Value
SAP_SID	SAP system ID	<SAPSID>

Parameter	Description	Value
SAP_CI_HOSTNAME	Host name of the SAP central instance	<sap-ci-hostname>
SAP_VERSION	SAP kernel version	<sap-kernel-version>

Network Tiebreaker

If the cluster is split into two or more subclusters (for example, because of network problems), only one subcluster must have the authority to manage the resources of the cluster. Otherwise, data might be corrupted when it is accessed by several nodes at the same time and the READ or WRITE operations are not synchronized in any way.

To ensure correct behavior in the event of a cluster split, you need a tiebreaker. Using the tiebreaker technology, only one subcluster gets the quorum to manage the resources.

The network tiebreaker uses the following rule:

The subcluster that can be accessed from outside of the cluster gets the quorum. The network tiebreaker uses an external IP address to resolve a tie. It uses ICMP echo requests (`ping`) to check the network connectivity to an IP address that is outside of the cluster.

If the network tiebreaker is able to ping the IP address, this part of the cluster gets the quorum and manages the resources. We recommend that you use the IP address of the gateway for the network tiebreaker. By choosing the gateway as the tiebreaker, you ensure that all nodes that can ping the tiebreaker can communicate with each other via the gateway. These nodes should be in the same subcluster after a cluster split.

Note

Make sure that the ICMP echo request and ICMP echo reply messages are not rejected by any firewall settings for the tiebreaker and the cluster nodes.

A.1.2 Cluster Setup Script `sapdb2cluster.sh`

Use

The script `sapdb2cluster.sh` sets up the cluster according to the cluster configuration file. It configures SA MP according to the data contained in the `sapdb2cluster.conf` file and creates all required resources, resource groups, relationships between the resources, and installs the high-availability policy files on the cluster nodes.

```
./sapdb2cluster.sh [-h] [-tl <n>] [-l <LogFile>] [-f <ConfigFile>] [-[db|wd|sr] -[show|create|state|switch|delete]]
```

Options

Parameter	Description
-h	Prints the usage syntax
-l <logfile>	Specifies the log file name. The default log name is <code>sapdb2cluster.log</code> .
-f <configfile>	Specifies the configuration file name. By default, the script looks for <code>sapdb2cluster.conf</code> in the current directory.
-tl <n>	Specifies the trace level: 1 – Debug (Default), 2 – Warning, 3 – Error
-db -show	Show and edit database cluster configuration
-db -create	Create database cluster
-db -state	Show database cluster state
-db -switch	Graceful database cluster switch (micro-outage)
-db -delete	Delete database cluster

Example

The following are examples of how you can use the script `sapdb2cluster.sh` and its options:

- Start cluster setup menu: `./sapdb2cluster.sh`
- Show database cluster configuration immediately: `./sapdb2cluster.sh -db -show`
- Start graceful cluster switch immediately: `./sapdb2cluster.sh -db -switch`

A.2 Syntax of SA MP Setup Scripts

A.2.1 Script prereqSAM

The `prereqSAM` script determines the prerequisites of SA MP on this node. It checks the level of the operating system, architecture, and distribution as well as the presence of required software packages.

```
prereqSAM    [--nodeleteinstlog] [--silent]  
             [-d <inst-pkg-dir>] [-l <log-file>]
```


Options

Parameter	Description
<code>--nodeleteinstlog</code>	Prevents removal of old logs. Since logs carry a process ID (PID) in their name, it is not likely that old logs are overwritten.
<code>--silent</code>	Suppresses output to the command line
<code>-d <inst-pkg-dir></code>	Allows prerequisite checking based on NLS files in <code><inst-pkg-dir></code> although <code>prereqSAM</code> is not in the package directory. <code>prereqSAM</code> is not dependent on the packages to be installed, but needs the directory for its message files.
<code>-l <logfile></code>	Logs to file <code><log-file></code> instead of <code>/tmp/prereqSAM.<pid>.log</code>

Return Codes

Return Code	Description
0	The version of the operating system is supported and all prerequisite packages were found installed at the correct version. The log file contains names and versions of the installed packages.
20	An installed package does not have the correct version. The log file contains the names and versions of the missing packages.
21	Package has not been found installed
22	Operating system version is not supported
23	Unable to conduct prerequisite checking due to a reason described in more detail in the log. Reason might be that a file is missing.

A.2.2 Script installSAM

The `installSAM` script checks the prerequisites once again and finally installs the SA MP software on this machine.

```
installSAM [--force] [--noliccheck]
           [--nodeleteinstlog] [--silent] [--upgrade]
           [-d <inst-pkg-dir>] [-l <logfile>]
```

Options

Parameter	Description
<code>--force</code>	Allows downgrading package <code>sam</code> if the installed version is higher than the version of the supplied package

Parameter	Description
<code>--noliccheck</code>	Allows first time installation even if no license file is supplied, or allows upgrade if no installed license is found. The fact is recorded in the log.
<code>--nodeleteinstlog</code>	Prevents removal of old logs. Since logs contain a PID in their name, it is unlikely that old logs are overwritten.
<code>--silent></code>	Suppresses output to the command line
<code>--upgrade</code>	Allows to upgrade SA MP if <code>--silent</code> is specified
<code>-d <inst-pkg-dir></code>	Allows installing packages from directory <code><inst-pkg-dir></code> although <code>installSAM</code> is not in the package directory
<code>-l <logfile></code>	Logs to file <code><logfile></code> instead of <code>/tmp/installSAM.<pid>.log</code>

Return Codes

Return Code	Description
0	<code>installSAM</code> including prerequisite checking and installation completed successfully
1	Package installer returned with return code not 0, return code and message can be found in the log. On AIX , the package installer is <code>installp</code> . On Linux , the package installer is <code>rpm</code> .
2	Package <code>sam</code> is already installed at the same version. If <code>--force</code> is specified, the installation overwrites the same version. A warning is logged and the installation status is returned.
3	Package <code>sam</code> is already installed at a higher version. If <code>--force</code> is specified, the installation overwrites the higher version. A warning is logged and the installation status is returned.
4	Package <code>sam</code> is already installed at a lower version than package version and option <code>--silent</code> is specified. If <code>--upgrade</code> is specified, an upgrade from a lower version is performed and the upgrade status is returned.
5	The node on which the installation task is required is online. The task is not performed.
6	SA MP base license file is not found, or no installed license is detected.

Return Code	Description
7	installSAM was unable to continue because directories or files could not be detected or other conditions make it fail.
2x	Prerequisite checking conducted by prereqSAM failed (see Script prereqSAM [page 48])

A.2.3 Script uninstallSAM

The uninstallSAM script removes the SA MP software.

```
uninstallSAM [--nodeleteinstlog] [--silent] [-d <inst-pkg-dir>] [-l <log-file>]
```

Options

Parameter	Description
--nodeleteinstlog	Prevents removal of old logs. Since logs carry a PID in their name, it is not likely that old logs are overwritten.
--silent	Suppresses output to the command line
-d <inst-pkg-dir>	Allows uninstall from directory <inst-pkg-dir> although uninstallSAM is not in the package directory
-l <logfile>	Logs to file <log-file> instead of /tmp/prereqSAM.<pid>.log

Return Codes

Return Code	Description
0	uninstallSAM completed successfully
1	An error occurred during uninstall
2	The domain is still online. Uninstall stopped.
3	Uninstall is not possible

A.3 Troubleshooting

A.3.1 Symbol Resolution Failed for db2start on AIX

If you have installed Db2 on a shared disk on AIX, the following error can occur when you start Db2:

```
> db2start
Could not load program db2start:
Symbol resolution failed for /usr/lib/threads/libc.a[aio_64.o] because:
Symbol kaio_rdwr64 (number 0) is not exported from dependent module /unix.
Symbol listio64 (number 1) is not exported from dependent module /unix.
Symbol acancel64 (number 2) is not exported from dependent module /unix.
Symbol iosuspend64 (number 3) is not exported from dependent module /unix.
Symbol aio_nwait (number 4) is not exported from dependent module /unix.
Symbol aio_nwait64 (number 5) is not exported from dependent module /unix.
Symbol aio_nwait_timeout (number 6) is not exported from dependent module /unix.
Symbol aio_nwait_timeout64 (number 7) is not exported from dependent module /
unix.
System error: Error 0
Examine .loader section symbols with the 'dump -Tv' command.
```

The error occurs if you do not install Db2 on the node where you did not use `db2setup` and you do not have asynchronous I/O (AIO) enabled. Normally, `db2setup` enables AIO for you during setup.

Solution

1. Check if AIO is enabled using the following command:
`lsdev -Cc aio`
If AIO is enabled, the output is as follows:
aio0 Available Asynchronous I/O (Legacy)
If AIO is not enabled, the output is as follows:
aio0 Defined Asynchronous I/O (Legacy)
2. Enable AIO:
 1. Log on to the node as user `root`.
 2. Open the AIO configuration in `smitty` using the following command:
`smitty aio`
 3. Enable AIO as shown in *Configure defined aio* in the `smitty` menu.

A.3.2 Db2 Immediately Stops After the db2start/startdb/ startjeedb Command

Problem

After you have installed SA MP and configured SA MP to manage Db2, the native Db2 command `db2start` becomes obsolete. If you use `db2start` to start a Db2 instance, SA MP does not know that you want to start

Db2 on this node. Once SA MP monitors that Db2 is started on this node, it immediately stops Db2 according to the rule *“this resource is supposed to be offline”*.

On Linux and AIX (if `syslogd` is configured), an entry in the system log shows that `db2_stop.ksh` is invoked to bring Db2 down.

During the setup of the cluster, the SAP scripts `startdb` and `startj2eedb` are replaced by the script `sapdb2post.sh`. The new scripts use SA MP commands to start Db2 via SA MP commands.

Solution

Make sure that the scripts that you use to start Db2 are the same as on the SA MP setup images.

A.3.3 Db2 Immediately Starts After the `db2stop/stopdb/stopj2eedb` Command

Problem

When you have installed SA MP and configured SA MP to manage Db2, the native Db2 command `db2stop` becomes obsolete. If you use `db2stop` to stop a Db2 instance, SA MP does not know that you want to stop Db2 on this node. Once SA MP monitors that Db2 is stopped on this node, it immediately starts Db2 according to the rule *“this resource is supposed to be online”*. On Linux and AIX (if `syslogd` is configured), an entry in the system log shows that `db2_start.ksh` is invoked to start Db2. During the setup of the cluster, the SAP scripts `stopdb` and `stopj2eedb` are replaced by the script `sapdb2post.sh`. The new scripts use SA MP commands to stop Db2 using SA MP commands.

Solution

Make sure that the scripts that you use to stop Db2 are the same as on the SA MP setup images.

A.4 Graceful Maintenance Tool (Former Graceful Cluster Switch)

The graceful maintenance tool (GMT) script replaces what was formerly known as *graceful cluster switch*. The graceful cluster switch has become a feature of the graceful maintenance tool (GMT). The GMT script can be used with the `DB2_GMT_CMD=CLUSTER_FAILOVER` command and performs an automated Db2 HADR cluster failover. This maintenance task can be executed from any Db2 cluster server.

For more information about the GMT, see the following documentation:

Database Administration Guide for SAP on IBM Db2 for Linux, UNIX, and Windows on SAP Help Portal at https://help.sap.com/viewer/db6_admin

SAP Notes [1530812](#) and [1907533](#)

A.5 References

Are you looking for more documentation? Here's an overview of information sources that are available for SAP systems on IBM Db2 for Linux, UNIX, and Windows.

Documentation by SAP

For central access to all our documentation, use our [SAP on IBM Db2 overview page](#) on SAP Help Portal.

Also check out and participate in our [SAP community for IBM Db2](#). Here you'll find blog posts, Q&As, whitepapers, videos, and guides.

To find SAP Notes, use the [SAP Support portal](#).

If you're looking for installation guides for SAP systems on Db2, you'll find them [here](#) on SAP Help Portal.

Documentation by IBM

For product documentation on your Db2 version, go to the [IBM Db2 documentation](#).

A.6 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.

© 2024 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.