



Standardize IBM DB2 HADR Configuration Using HP DMA

HP Database and Middleware Automation version 10.30

Table of Contents

Purpose	2
Prerequisites	6
Prerequisites for the DB2 - Configure HADR Database workflow	6
Prerequisites for the DB2 - Configure Tivoli SAMP On HADR Database workflow	7
Process Overview	9
Workflow 1: DB2 - Configure HADR Database	10
Workflow 2: DB2 - Configure Tivoli SAMP On HADR Database	12

Purpose

This paper describes how to use HP Database and Middleware Automation (HP DMA) to create, automate and standardize configuring an IBM DB2 high availability disaster recovery (HADR) environment on an existing DB2 software, instance, and database using default or custom "golden user templates" and configuring IBM DB2 HADR with Tivoli System Automation for Multiplatforms (TSAMP).

Benefits of HP DMA

HP DMA automates many of the daily administrative tasks required to manage the lifecycle of relational databases and J2EE application servers. These tasks are complex, often manual, typically time-consuming, and frequently error-prone. HP DMA improves the efficiency of these administrative tasks, enabling administrators to deliver change faster with higher quality, better consistency, and improved reliability.

HP DMA equips you to do the following:

- Define and enforce standards for software installation
- Define an installation process once and reuse it repeatedly
- Leverage knowledge enterprise-wide
- Avoid human error

DB2 HADR

The DB2 High Availability Disaster Recovery (HADR) is a database replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby or the secondary. The standby database is kept in synch with the primary database through log data that is generated on the primary and shipped to the standby. The standby constantly rolls forward through the logs.

HADR and HP DMA

An HADR primary database does not automatically switch over to its standby database in the event of a failure. You must perform a manual takeover when the primary database has failed. HP DMA automates this task by using the Automatic Client Reroute to handle failover. The Automatic Client Reroute feature transfers client application requests from a failed database server to a standby database server through floating virtual IP address. User can also configure primary (public) and secondary (private interconnect) network for setting up the failover. Currently DMA does not support multiple standby configuration, however, the workflow can be customized to setup the failover.

HP DMA validates that the DB2 software is installed, the instance already exists, and is up and running on both primary and standby. It also validates and ensures that the HADR database only exists on the primary instance node before configuring HADR.

HP DMA shuts down the primary DB2 instance and the licensing daemons that are running for the primary instance on which the DB2 - Configure HADR Database workflow is deployed. It kills all application user connections that have DB2 instances that are to be upgraded. It validates the eligibility for the instance(s) and database(s) to be configured for HADR.

IBM Documentation

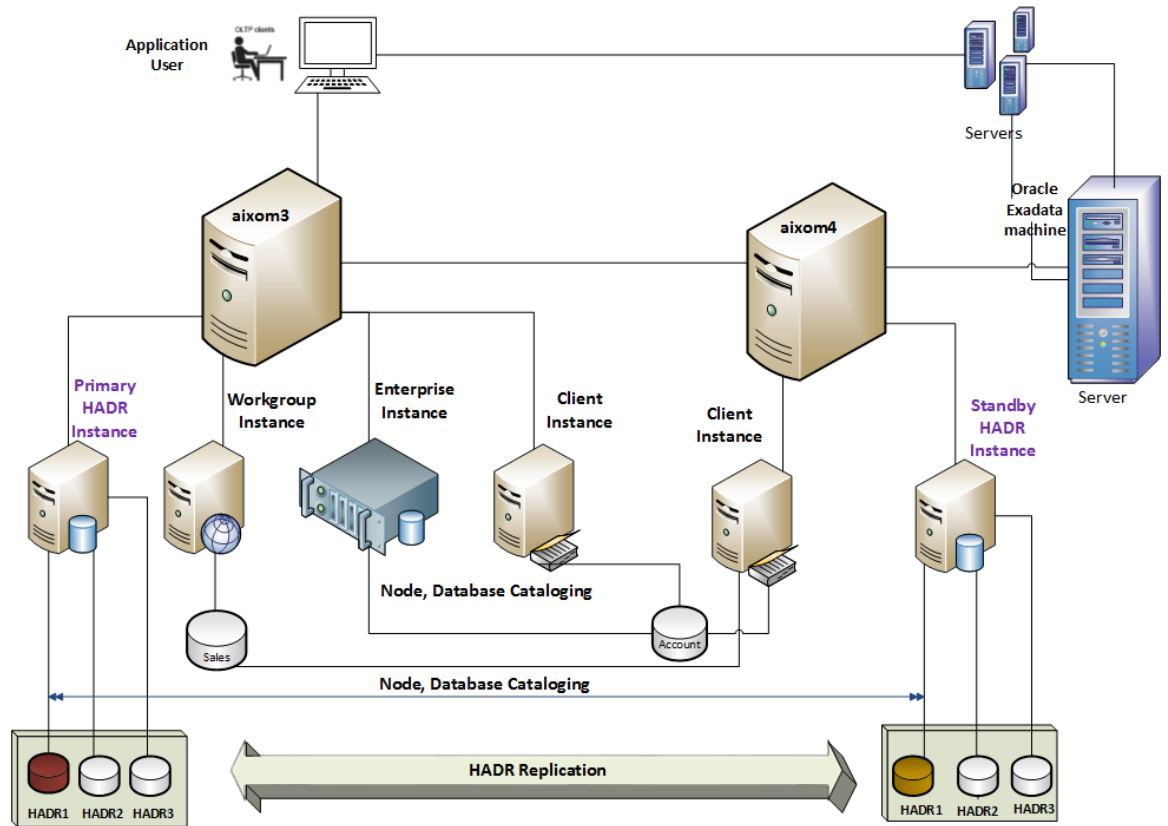
For more information about IBM DB2, refer to the following IBM documentation:

DB2 version	IBM documentation
10.1	IBM DB2 10.1
10.5	IBM DB2 10.5

Goal

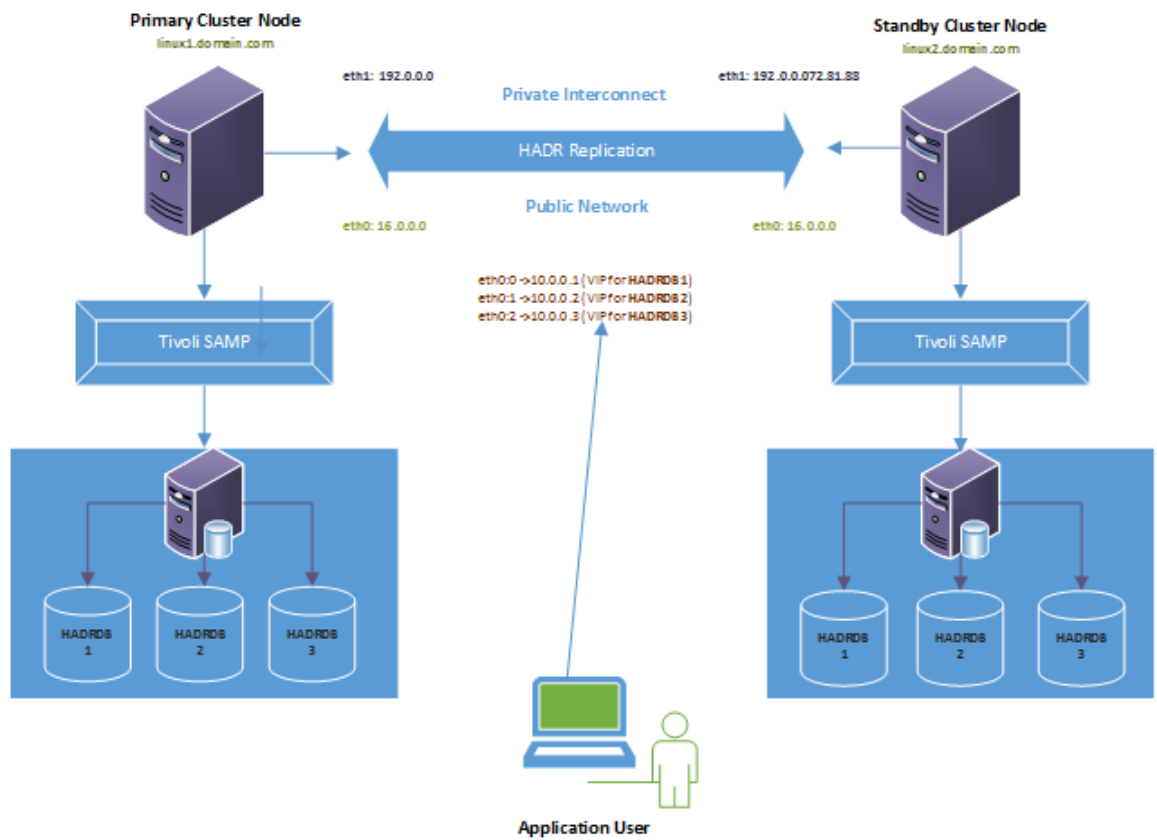
The DB2 - Configure HADR Database configures DB2 database(s) for IBM DB2 for Linux and AIX operating systems on the target source and destination servers. This workflow currently supports DB2 versions 10.1 and 10.5 on RedHat Linux and AIX servers. This is an instance level workflow. This workflow configures the database with the HADR feature.

High Availability Data Replication Using HADR Port



The DB2 - Configure Tivoli SAMP On HADR Database workflow configures DB2 database(s) for IBM DB2 for Linux and AIX operating systems on the target source and destination servers with Tivoli System Automation for Multiplatforms (samp) where this workflow is deployed.

High Availability with Tivoli System Automation for Multiplatforms



Prerequisites

Before performing the procedures in this paper, your environment must meet the following minimum requirements:

Prerequisites for the DB2 - Configure HADR Database workflow

- SSH service must be turned on for both primary and standby host computers.
- The source and destination host computer is configured with SSH password-less login across the nodes (primary to standby and vice versa).
- A TCP/IP interface must be available between the HADR host computers, and a high-speed, high-capacity network is recommended.
- Use identical host computers for the HADR primary and standby databases. That is, they should be from the same vendor and have the same architecture.
- Both the primary and standby host computers must run one of the following operating systems (that is supported by IBM DB2 10.1 or 10.5 and HP DMA):
 - Linux
 - AIX

See the *HP DMA Support Matrix* for specific operating system versions, available at: softwaresupport.hp.com

- The operating system on the primary and standby host computer must be the same version, including patches.
- DB2 software must be provisioned on both the primary and standby host computer.

Tip: You can use [Workflow 1: DB2 - Provision Software](#) workflow to accomplish this.

- The DB2 instance must be provisioned on both primary and standby host computer.

Tip: You can use [Workflow 2: DB2 - Provision Instance](#) workflow to accomplish this.

- The DB2 database must be created on the instance at the primary host computer on which the workflow will be deployed.

Tip: You can use [Workflow 3: DB2 - Provision Database](#) workflow to accomplish this.

- DB2 instance on primary host computer must be up and running on both the primary and standby host computer.
- Installation media:
 - The DB2 server installation software binary file from IBM.
 - Installation software binary file must be available locally or available for download from the software repository.¹
- Storage:
 - 4-6 GB to provision the DB2 software.

¹For additional information, see [Alternative methods for specifying input files](#).

- 1 GB for each DB2 instance.
- 1 GB for each DB2 database (more may be required for your configuration).
- At least 1 GB for Catalog tablespace.
- If automatic storage is on, 1 GB on the default directory where the default tablespace will be created.
- Unchallenged ability to become the DB2 database user.
- The operating system kernel parameters and shared memory is properly configured.
- License for HP DMA.
- License for DB2 Database version 10.1 or 10.5.

Note: You have 90 days before you are required to purchase a DB2 license.

- The following workflow requirements:

Workflow	Requirements
DB2 - Configure HADR Database	<ul style="list-style-type: none"> - The <code>sudo</code> package is installed on the target servers. - The target servers have the <code>gunzip</code> and <code>tar</code> utilities in the environment path.

Prerequisites for the DB2 - Configure Tivoli SAMP On HADR Database workflow

- SSH service must be turned on for both primary and standby host computers.
- The source and destination host computer is configured with SSH password-less login across the nodes (primary to standby and vice versa).
- DB2 software must be provisioned on both the primary and standby host computer.

Tip: You can use [Workflow 1: DB2 - Provision Software](#) workflow to accomplish this.

- The DB2 instance must be provisioned on both primary and standby host computer.

Tip: You can use [Workflow 2: DB2 - Provision Instance](#) workflow to accomplish this.

- The DB2 database must be available on primary and standby instance cluster node with primary and standby HADR configuration state.
- DB2 instance on primary host computer must be up and running on both the primary and standby host computer.
- Tivoli System Automation for Multiplatforms (TSAMP) must be installed and on both primary and standby cluster nodes.
- Valid license to activate the TSAMP on primary and standby cluster nodes.
- Guidelines from IBM to provide the correct input parameter values for the steps Gather Parameters For Configure Tivoli SAMP on HADR Database and Gather Advanced Parameters For Configure Tivoli SAMP on HADR Database in order to prepare correct XML file for TSAMP.

- The following workflow requirements:

Workflow	Requirements
DB2 - Configure Tivoli SAMP On HADR Database	<ul style="list-style-type: none">- The <code>sudo</code> package is installed on the target servers.- The target servers have the <code>gunzip</code> and <code>tar</code> utilities in the environment path.

Refer to the [IBM Documentation](#) for the following:

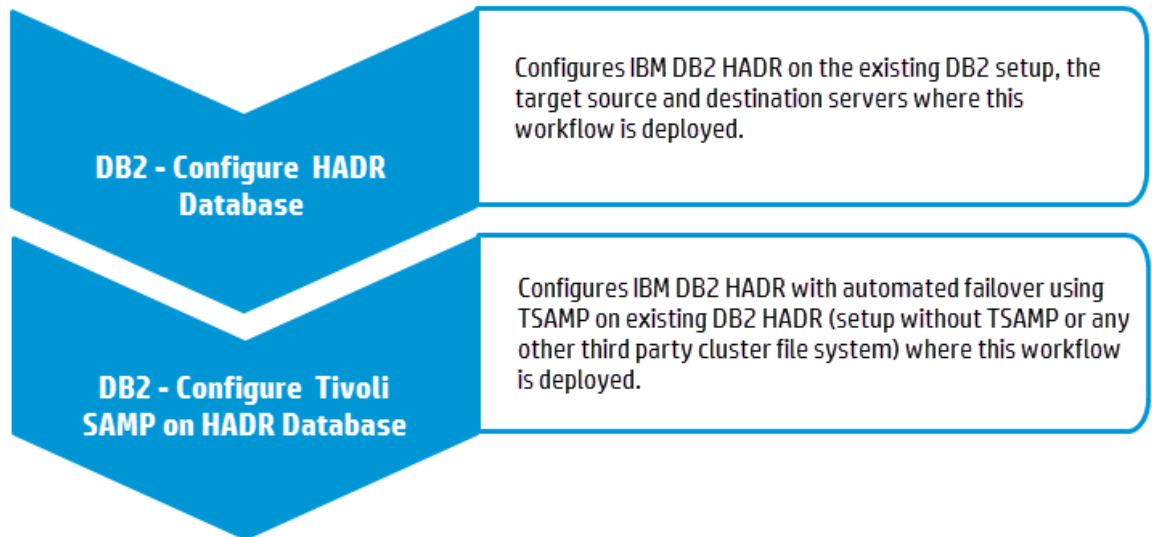
- Complete installation and infrastructure requirements for IBM DB2.
- Acceptable types and range of values when using HP DMA advanced parameters to configure IBM DB2 HADR settings.

Tip: For examples of how to meet these general requirements for DB2 10.5 on either a Red Hat Enterprise Linux server or an AIX server, see [How to meet prerequisites](#). You can adapt these instructions for other DB2 versions.

Process Overview

Configuration of DB2 HADR is done in two steps. In the first step, DB2 - Configure HADR Database workflow configures the HADR with manual failover. In the second step, the DB2 - Configure Tivoli SAMP On HADR Database workflow configures HADR using Tivoli SAMP which allows to have automatic failover and failback. These workflows validate prerequisites, configures, and verifies HADR configuration.

Use the following HP DMA workflows to standardize the process of configuring DB2 HADR:



Note: These workflows support DB2 version 10.1 or 10.5 on a Red Hat Linux or AIX server.

The examples given are appropriate for configuring DB2 10.5 HADR on an AIX server with an 'Enterprise' edition instance.

Before running the DB2 - Configure HADR Database workflow, the DB2 license must be activated for the instances that you create. For more information, see [Where can I learn more about IBM licenses?](#)

The following sections provides detailed information required to run the workflow.

Note: For additional information, see [How to run an HP DMA workflow.](#)

Workflow 1: DB2 - Configure HADR Database

This section provides detailed information required to run the DB2 - Configure HADR Database workflow.

Solution pack

This workflow requires the HP DMA Database Provisioning Solution Pack.

Parameters to expose

None

Input parameters

When you deploy the DB2 - Configure HADR Database workflow, specify input parameter values for the following steps.

The parameters in the "Gather Advanced Parameters for Configure DB2 HADR" step provide options to customize the workflow deployment.

Note: Bold text in the following tables indicates that you must specify a value for the parameter.

Step: Gather DB2 Source and Destination Instances

Parameter	Description	Example Value
DB2 Destination HADR Instance	Required: The standby node instance name of the remote server. Administration tools, such as the DB2 Control Center, use this parameter to contact the remote server. In the bridge execution workflow, this instance will be used at the run time to run the specific steps configured to run on the standby instance node. The value will be set at the run time.	hadr105 [aixom02.mycompany.com]
DB2 Source HADR Instance	Required: The primary node instance name of the local server. Administration tools, such as the DB2 Control Center, use this parameter to contact the local server. In the bridge execution workflow, this instance will be used at the run time to run the specific steps configured to run on the primary instance node. The value will be set at the run time.	hadr105 [aixom01.mycompany.com]

Step: Gather Parameters for Configure DB2 HADR

Parameter	Description	Example Value
DB2 HADR Database Name	Required: The database name for which the High Availability Disaster Recovery will be configured. The database name must be available on the primary instance node.	DB2_HADR
DB2 HADR Local Service Name	Required: This parameter specifies the TCP service name for which the local high availability disaster recovery (HADR) process accepts connections. Example: DB2_HADR_SVC.	DB2_HADR_SERVICE_P1
DB2 HADR Port Number	Required: This parameter specifies the TCP service port number for which the high availability disaster recovery (HADR) process accepts connections. The same port will be used in the primary and standby node for HADR communication.	58234
DB2 HADR Remote Instance Name	Required: This parameter specifies the instance name of the remote server. Administration tools, such as the DB2 Control Center, use this parameter to contact the remote server. High availability disaster recovery (HADR) also checks whether a remote database requesting a connection belongs to the declared remote instance.	hadr105
DB2 HADR Remote Service Name	Required: This parameter specifies the TCP service name for which the remote high availability disaster recovery (HADR) process accepts connections. Example: DB2_HADR_SVC.	DB2_HADR_SERVICE_P1
DB2 HADR Standby Host Name	Required: This parameter specifies the TCP/IP host name or IP address of the remote high availability disaster recovery (HADR) database server.	aixom2.mycompany.com

Workflow 2: DB2 - Configure Tivoli SAMP On HADR Database



This section provides detailed information required to run the DB2 - Configure Tivoli SAMP On HADR Database workflow.

Solution pack

This workflow requires the HP DMA Database Provisioning Solution Pack.

Parameters to expose

None

Input parameters

When you deploy the DB2 - Configure Tivoli SAMP On HADR Database workflow, specify input parameter values for the following steps.

The parameters in the Gather Advanced Parameters For Configure Tivoli SAMP on HADR Database step provide options to customize the workflow deployment.

Note: Bold text in the following tables indicates that you must specify a value for the parameter.

Step: Gather DB2 Source and Destination Instances

Parameter	Description	Example Value
DB2 Destination HADR Instance	Required: The standby node instance name of the remote server. Administration tools, such as the DB2 Control Center, use this parameter to contact the remote server. In the bridge execution workflow, this instance will be used at the run time to run the specific steps configured to run on the standby instance node. The value will be set at the run time.	hadr105 [aixom02.mycompany.com]

Step: Gather DB2 Source and Destination Instances, continued

Parameter	Description	Example Value
DB2 Source HADR Instance	The primary node instance name of the local server. Administration tools, such as the DB2 Control Center, use this parameter to contact the local server. In the bridge execution workflow, this instance will be used at the run time to run the specific steps configured to run on the primary instance node. The value will be set at the run time.	hadr105 [aixom01.mycompany.com]

Step: Gather Parameters For Configure Tivoli SAMP on HADR Database

Parameter	Description	Example Value
Database Name	Required: The database name for which the High Availability Disaster Recovery will be configured. The database must be available on the primary and standby instance cluster nodes.	DB2HADR
IP Of Primary Cluster Node	Required: Internet Protocol Address (IP address) for primary cluster node machine where the HADR Instance and database is configured.	16..0.0.1
IP Of Standby Cluster Node	Required: Internet Protocol Address (IP address) for primary cluster node machine where the HADR Instance and database is configured.	16.0.0.2
Local Instance Name	Required: This parameter specifies the instance name of the local cluster node. Administration tools, such as the DB2 Control Center, use this parameter to contact the local server. High availability disaster recovery (HADR) also checks whether a local database requesting a connection belongs to the declared local instance. Default, it is configured to use the instance name on which this workflow is deployed.)	DB2_105HADR_SVC1
Local Instance Port Number	Required: DB2 connection port number for the local instance on primary cluster node where HADR database is mounted.	51000

Step: Gather Parameters For Configure Tivoli SAMP on HADR Database, continued

Parameter	Description	Example Value
Primary Cluster Node Name	Required: This parameter specifies the local host (primary cluster node name) name for high availability disaster recovery (HADR) TCP communication.	aixom01.mycompany.com
Quorum Device Name	Required: A network quorum is an IP address that can be pinged from both the primary and the standby nodes. In the event of a site failure, the quorum decides which node serves as the active node and which node goes offline. When you are choosing the network quorum, ensure that the IP remains active all the time. The DNS server IP is always a good choice for the network quorum.	16.0.1.1
Remote Instance Name	Required: This parameter specifies the instance name of the remote cluster node (server). Administration tools, such as the DB2 Control Center, use this parameter to contact the remote server. High availability disaster recovery (HADR) also checks whether a remote database requesting a connection belongs to the declared remote instance.	hadr105
Remote Instance Port Number	Required: DB2 connection port number for the remote instance on standby cluster node where HADR database is mounted.	51000
Standby Cluster Node Name	Required: This parameter specifies the remote host(standby cluster node name) name for high availability disaster recovery (HADR) TCP communication.	aixom02.mycompany.com
Subnetmask Of Primary Cluster	Required: Subnet Mask Address(IP address) for primary cluster node machine where the HADR Instance and database is configured.	255.255.240.0
Subnetmask Of Standby Cluster	Required: Subnet Mask Address(IP address) for standby cluster node machine where the HADR Instance and database is configured.	255.255.240.0

To learn more about HP Database and Middleware Automation visit

hp.com/go/dma

© Copyright 2014-2015 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

May 2015



