


# IBM Workload Deployer

## Database-as-a-Service (DBaaS) patterns



This presentation covers the Database-as-a-Service (DBaaS) support found in the IBM Workload Deployer product.

## Table of contents

- Overview
- Enabling database patterns
- Creating database patterns
- Summary

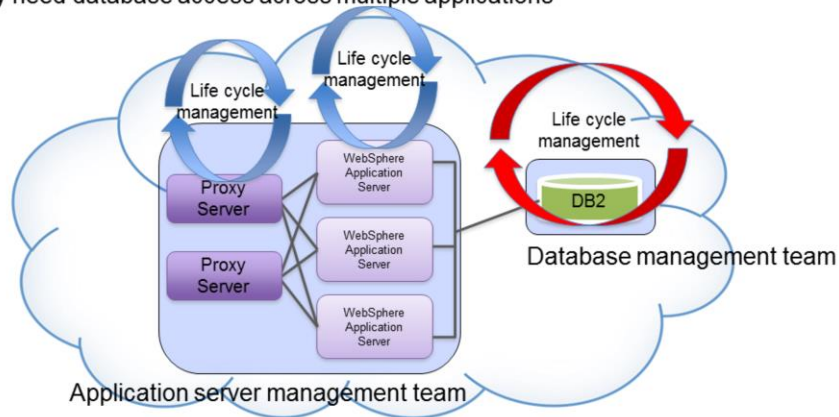
Looking at the agenda, you will see an overview of the DBaaS support and some information on its actual use. This includes information about enabling and creating database patterns.

## ***Overview***

In this section you will see an overview of the DBaaS function available in the IBM Workload Deployer appliance.

## Overview

- Separate need for Databases in many usage scenarios:
  - Separate management team
  - Life cycle independent of application life cycle
  - May need database access across multiple applications



4

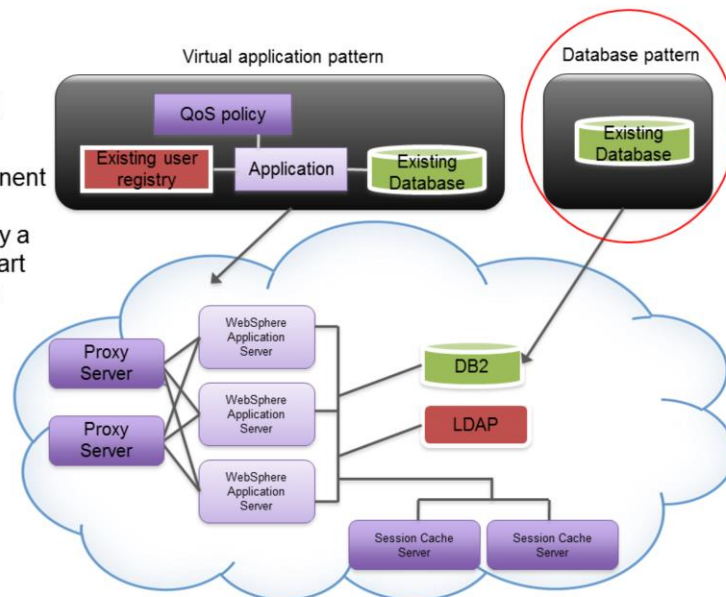
DBaaS patterns

© 2012 IBM Corporation

In many usage scenarios the database is a distinct entity with its own administrative team and its own life cycle. IBM Workload Deployer models this behavior with database patterns and database instances.

## Database patterns

- Created and managed independent of any virtual application pattern
- Existing Database component (shown)
  - Also possible to deploy a database pattern as part of a virtual application pattern



5

DBaaS patterns

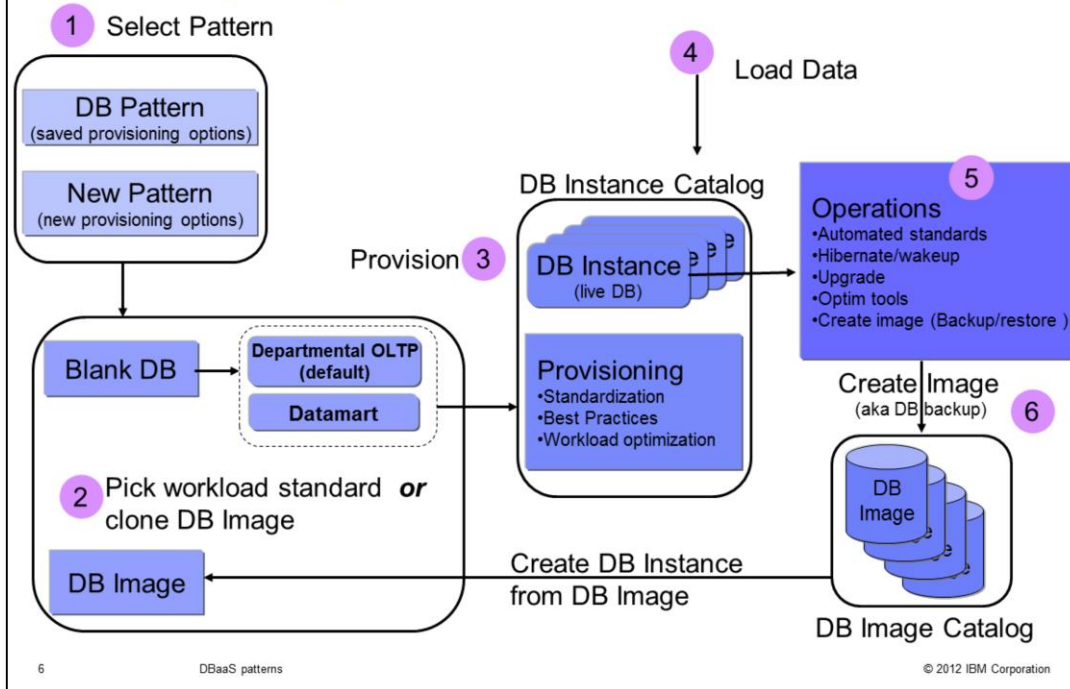
© 2012 IBM Corporation

In most “real world” usage scenarios databases are managed independent of any one application. To more closely model this paradigm, “Database patterns” were introduced. You can create, delete, update, backup and restore databases created by a database pattern. These management activities are independent of your virtual application. Deleting your virtual application has no impact on the deployed database or database pattern.

This slide shows that you can include an existing database as part of your virtual application pattern. This is considered an “existing database component” and is not actually deployed as part of the virtual application. The database already exists and is being used by the application deployed in the virtual application pattern. The “existing database component” can be created as a database instance in Workload Deployer, as shown in the red circle labeled as “Database Pattern.” In this case, the database is managed within Workload Deployer. It can also be an already existing remote database that was created and is being managed outside of the Workload Deployer environment. In both of these cases, where an “existing database component” is being used, deletion of the database does not affect the state of the virtual application instance, and vice versa.

While this is not shown on the slide, note that it is also possible to deploy a database pattern as part of a virtual application. Instead of including an “existing database” in the virtual application pattern, you can instead include a “database component” that becomes a pattern-deployed database service. In the case of a virtual application pattern-deployed database service, the database is deployed as part of the virtual application so when the virtual application is deleted, the database is deleted as well. You might find this useful in some instances but for most scenarios, you will most likely want to use an “existing database component” that is managed independently.

## Provisioning and operations flow



This slide shows the overall flow for provisioning a database from IBM Workload Deployer. You see that the first thing you need to do is select a pattern. You can use a pattern that was saved OR create a new one here. When creating a pattern you have two provisioning options as seen in number two. You can select the type of workload standard to apply to the DB instance to be created OR you can select to clone an existing DB image that has been backed up to your DB image catalog repository. Tivoli Storage Manager is used by Workload Deployer to store these DB images. If selecting a workload standard from which to create your pattern, you will have two options: Departmental OLTP and Dynamic Data Mart. Departmental OLTP is the default. Once you have your pattern set up as you want, you will deploy the database pattern which causes a DB instance to be provisioned as seen in number three on the slide. The provisioning process will make many decisions for you to tailor your DB instance. The provisioned DB instance will include standardization, best practices and workload optimization that correspond to the workload standard selected. Once provisioned, you have a fully functional robust database for your IT environment. Number five on the slide shows some operations that are available to you from the IBM Workload Deployer Database Service Console. This includes the backup of your image to the DB Image Catalog as seen in number six. Note that this operation requires Tivoli Storage Manager.

## Workload standard provisioning

- **Workload standard:**
  - Specified as part of the database pattern
  - Set of predefined database configurations used to provision a database
    - A set of scripts runs to tune the operating system and instance configuration, create the database and accompanying objects and load the initial data
- **Two workload standards available:**
  - Transactional standard (OLTP)
    - Primarily used for online transaction processing and is optimized for transactional applications
  - Data Mart standard
    - Primarily used for data warehousing and is optimized for reporting applications

The previous slide mentioned selecting a workload standard to apply to the DB instance to be created as an option when creating a pattern. Workload standards are a set of predefined database configurations used as a provisioning approach for creating database patterns. When a workload standard is selected, a set of scripts runs to tune the operating system and instance configuration, create the database and accompanying objects and load the initial data.

There are two workload standards available with IBM Workload Deployer. The Transactional standard is primarily used for online transaction processing and is optimized for transactional applications. The Data Mart standard is primarily used for data warehousing and is optimized for reporting applications.

## Clone provisioning

- Define a database pattern using an existing database image
  - Uses existing image as a model for creating the database pattern
    - Existing image has been previously created
    - Existing image can be manually created or automatically created by DBaaS
      - Recommend using one that has been manually created
- Virtual machine is created with the same resource settings as the original
  - Uses same licenses and configuration

Cloning is the other provisioning approach available to you with IBM Workload Deployer. Cloning uses an existing database image as a model for creating database patterns. The existing database image must have been previously backed up to Tivoli Storage Manager. When an image is selected, the metadata stored during backup is retrieved. A new virtual machine is created with the same resource settings. The DB2 Restore command creates a new database with the same license and configurations. This cloned database then sits on the newly created virtual machine.

You should use manually created images in preference to automatically created backups for this task. You can manually create a database image with the Create a Database Image function in the Database Service Console. Images created this way are labeled as “Manual” under Image Type in the list of images.

IBM Tivoli Storage Manager must be configured to use this function.



## ***Enabling database patterns***

This section covers the enablement of the database patterns.

## Enabling database patterns

- Licenses for virtual images and pattern types **MUST** be enabled and system plug-ins configured before you can start creating database patterns
  - Accept licenses for base OS virtual images
  - Accept licenses for pattern types and enable patterns
  - Configure plug-ins for database workload patterns
- Configure Tivoli Storage Manager plug-in (for database backups) - optional

Before using database patterns, there is some administrative work that must be done. The virtual image used for database patterns is the base OS virtual image. You must accept the license for that in order to deploy database patterns. There are also some pattern types that you must accept licenses for and enable. Finally, there are some plug-ins that must be configured. You will see these on the next few slides. There is also an optional administrative task that might be needed if you plan to create database backups. If you plan to create database backups, you need to configure the Tivoli Storage Manager plug-in to point to an existing Tivoli Storage Manager installation.

## Enable virtual image licenses for base OS (Catalog > Virtual Images)

**Needs enabling**

Virtual Images		IBM OS Image for AIX Systems - Large	
AIX Systems	↑↓	Description:	None provided
IBM OS Image for AIX Systems	🌱	Created on:	Nov 19, 2011 2:28:01 AM
<b>IBM OS Image for AIX Systems - Large</b>	<b>🌱</b>	Current status:	🌱 License not accepted
IBM OS Image for AIX Systems - Medium	🌱	Updated on:	Nov 19, 2011 5:31:47 AM
IBM OS Image for AIX Systems - Small	🌱	License agreement:	🌱 Not accepted [accept...]
IBM OS Image for AIX Systems - Tiny	🌱	Hypervisor type:	PowerVM
IBM OS Image for AIX Systems - Xlarge	🌱	Operating system:	AIX, version 6100-05 (IBM AIX 6100-05)

**Accept license**

Virtual Images		IBM Workload Deployer Image for x86 Systems	
x86	↑↓	Description:	IBM Workload Deployer Image for x86 Systems
IBM Workload Deployer Image for x86 Systems	<b>🌱</b>	Created on:	Nov 19, 2011 2:22:17 AM
		Current status:	🌱 Read-only
		Updated on:	Jan 7, 2012 11:11:10 PM
		License agreement:	🌱 Accepted [view...]
		Hypervisor type:	ESX
		Operating system:	RedHat Enterprise Linux 64-Bit, version 5.7 (RHEL 5.7 X64)

**License accepted**

11 DBaaS patterns © 2012 IBM Corporation

The top picture on this slide is showing the base OS images for AIX Systems that are used for database pattern deployments on PowerVM hypervisors. All but one of them has had their license accepted. In order to use that one for deployments, you need to accept the license from the hyperlink highlighted on the right. If you are using ESX for deployments, there is only one image whose license needs to be accepted for database deployments. That is shown on the bottom of the slide and you see that its license has already been accepted. Note that z/VM does not currently support database patterns so there are no base OS virtual image licenses that you need to accept on that platform.

**Enabling pattern types (Cloud > Pattern Types)**

**Pattern Types**

- Foundation Pattern Type 2.0 (Already enabled)
- Foundation Pattern Type 2.0.0.0
- IBM Data Mart Pattern 1.1 (Needs enabling)**
- IBM Data Mart Pattern 1.1.0.0
- IBM Database Patterns 1.1 (Needs enabling)**
- IBM Database Patterns 1.1.0.0
- IBM Transactional Database Pattern 1.1 (Needs enabling)**
- IBM Transactional Database Pattern 1.1.0.0
- Web Application Pattern Type 2.0
- Web Application Pattern Type 2.0.0.0
- WebApp Pattern Type 1.0
- WebApp Pattern Type 1.0.0.3

**IBM Transactional Database Pattern**

Name: IBM Transactional Database Pattern  
 Description: IBM Transactional Database Pattern  
 License Agreement: Not Accepted [Accept...]  
 Status: Not Accepted [Enable All...]  
 Required: IBM Database Patterns (1.1.0.0 or above)  
 Plug-ins: Show me all Plug-ins in this pattern type  
 Dependency: This pattern type has one or more pre-requirements that need to be met. If you want to change the status of the current one, the listed pattern type must be enabled. Please press "Enable All" to enable the pattern types. Please

Pattern Type name	Status
Foundation Pattern Type 2.0.0.0	Available
Database Patterns 1.1.0.0	Unavailable

License Agreement: Accepted [View...]  
 Status: Available [Disable...]

**Success**

12 DBaaS patterns © 2012 IBM Corporation

This slide shows the pattern types that you need to enable. The three pattern types needed for database patterns are highlighted in the red box on the left and include the Data Mart Pattern, the Database Pattern and the Transactional Database Pattern. On the right you see the steps outlined that need to be completed for each pattern type to be enabled. In step one, you need to make sure that any pattern dependencies have been taken care of. In this case the Foundation Pattern and the Database Pattern are required. The Database Pattern has had its license accepted but has not been enabled. The Foundation Pattern is already enabled and is marked "Available" as shown. In step two, you need to click the "Accept" hyperlink and you are given the opportunity to read and accept this particular license. Finally in step 3, you need to click the "Enable all" hyperlink to enable this pattern type for use. Notice that the hyperlink says "Enable All". This allows you to enable the Database Pattern when you are enabling the Transactional Database Pattern. This is only possible if the license for the Database Pattern has already been accepted.

## Configure system plug-ins that make up pattern types (Cloud > System Plug-ins )

**System Plug-ins**

- Plug-in Name
- IBM Database Patterns 1.1
- activator (2.0.0.0) ✓
- agent (2.0.0.0) ✓
- autoscaling (2.0.0.0) ✓
- com.ibm.maestro.plugin.debug (2.0.0.0) ✓
- datamart (1.1.0.0)** ⚠
- db2 (1.1.0.0) ✓
- deployinlet (2.0.0.0) ✓
- dswc (1.1.0.0) ✓
- firewall (2.0.0.0) ✓
- foundation-license (2.0.0.1) ✓
- iaas.WCA (2.0.0.0) ✓
- logbackup (2.0.0.0) ⚠
- monitoring (2.0.0.0) ⚠
- oltp (1.1.0.0)** ⚠
- ssh (2.0.0.0) ✓

**datamart (1.1.0.0)** Configure Delete

Plug-in Name: datamart (1.1.0.0)  
 Created by: cbadmin  
 Last Modified on: Dec 13, 2011 10:39:59 AM  
 This plug-in has no components

**oltp (1.1.0.0)** Configure Delete

Plug-in Name: oltp (1.1.0.0)  
 Created by: cbadmin  
 Last Modified on: Dec 13, 2011 10:39:16 AM  
 This plug-in has no components

Configuration

Environment:

- None(disabled)
- Only IBM Transactional Database Pattern
- Only IBM Transactional Database Pattern for Non-Production Environment
- Both

**Select one**

13 DBaaS patterns © 2012 IBM Corporation

The IBM Data Mart Pattern and IBM Transactional Database Pattern each has a system plug-in that needs additional configuration before use. This slide shows the configuration that needs to be done. The needed plug-ins are found under the Cloud > System Plug-ins menu. Select the “IBM Database Patterns 1.1” and configure the ‘datamart’ and the “oltp” plug-ins as highlighted on the slide. You see that you need to select how you want to enable it. Your choices are to enable them for a production environment, enable them for a non-production environment or both. The default is that both are disabled. Until you have configured at least one plug-in, you are not able to deploy any database patterns.

## Tivoli Storage Manager configuration - optional

- Tivoli Storage Manager is required for backup and restore of DB2 image
- Tivoli Storage Manager integration requires system plug-in to be configured from:
  - Cloud > System Plug-ins, and then select **IBM Database Patterns 1.1.0.0** and then **tsm (1.1.0.0)** System plug-in

The screenshot displays the 'System Plug-ins' management interface. On the left, a list of plug-ins includes 'IBM Database Patterns 1.1.0.0', 'db2 (1.1.0.0)', 'dswc (1.1.0.0)', and 'tsm (1.1.0.0)'. The 'tsm (1.1.0.0)' plug-in is selected and highlighted. The main panel shows details for 'tsm (1.1.0.0)', including its name, creator (cbadmin), and last modified date (Oct 26, 2011 2:09:19 PM). A 'Configure' button is located in the top right corner of this panel. Below the main panel, a configuration dialog box is open, titled 'Configuration'. It contains the following fields: 'TSM server address' (empty), 'TSM server TCP/IP port' (1500), 'TSM server administrator user' (admin), 'TSM server administrator password' (empty), and 'Domain for DB2' (DB2DOMAIN). 'OK' and 'Cancel' buttons are at the bottom of the dialog.

To use Tivoli Storage Manager's backup functionality you must first configure the corresponding "tsm" system plug-in. To configure the system plug-in, navigate to Cloud and then System Plug-ins. Select the "IBM Database Patterns 1.1" and click "tsm (1.1.0.0)." Next, click the "Configure" button located in the upper right corner to bring up the configurable attributes panel. In this panel you enter the address, port and user credential information for the Tivoli Storage Manager. While the "default." "Domain for DB2" is DB2DOMAIN, this value should be obtained from your Tivoli Storage Manager administrator. Once these steps are complete you can use the database backup capabilities from Workload Deployer. The restore is done from Tivoli Storage Manager.

## ***Database patterns***

This section covers database pattern creation.

## Database patterns



**testSample** Deploy Open Export Delete Clone

**Application ID:** a-9ad8f936-4e2b-4837-a8b9-a64cf87952ac

**Description:** Blank application

**Created by:** cbadmin

**Last Modified by:** cbadmin

**Created on:** Sep 23, 2011 12:48:32 AM

**Last Modified on:** Sep 23, 2011 12:54:36 AM

**Database Pattern**

**Specify options for your database pattern.**

**Database pattern name:** testSample

**Database pattern description:**

**Purpose:** Production

**Source:** Apply a database workload standard

Name	Workload Type
<input checked="" type="radio"/> Departmental Transactional	Departmental Transactional
<input type="radio"/> Data Mart	Data Mart

**Maximum User Data Space (GB):** 10

**Database Compatibility Mode:** DB2 (Default)

**Schema File:** artifacts/setup\_db.sc Browse Delete

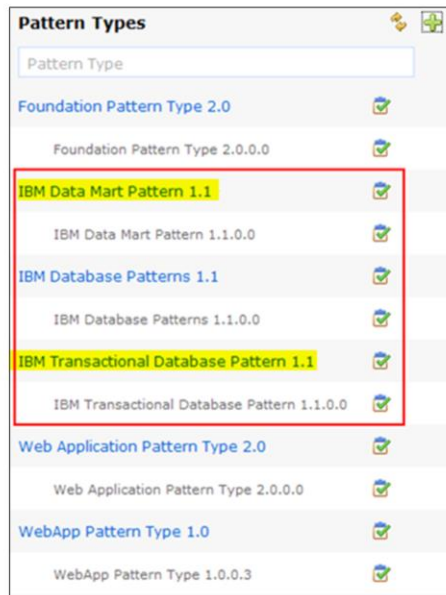
- Can be created from
  - Workload standard
  - Clone

Here you see the two different pattern types using the IBM Database Patterns that can be created for deployment to a ESX or PowerVM hypervisor. The first type of pattern is a “Virtual Application pattern” that includes a database as seen in the “testSample” virtual application shown in the virtual application builder. The second type of pattern that uses the IBM Database Patterns is the “Database Pattern” itself. This allows you to deploy just a database that can then be managed independent from any applications that might use it. In either case, a database can be created from a “Workload standard” or a “Clone.” You will see this on later slides.



## Associated pattern types

- IBM database workload patterns
  - Transactional Database Pattern
  - Data Mart Pattern



To support the workload standards, as you saw earlier, two database workload patterns are provided. They are the “Transactional Database Pattern” and the ‘Data Mart Pattern’. The Transactional Database Pattern provides a set of capabilities that are essential to the provisioning and management of transactional database infrastructure for datacentric applications. The Data Mart Pattern is similar to the Transactional Database Pattern, but deploys data marts including configuring the unique memory and disk requirements for this type of workload. The Data Mart Pattern includes data compression for improved performance in a shared disk environment typical to cloud data mart deployment.

## Create new database from workload standards in Database Patterns

My Database Patterns

Search...

- When 'Source' is set to 'Apply a database workload standard':
  - Departmental Transactional
  - Data Mart

**Database Pattern**

**Specify options for your database pattern.**

Database pattern name:

Database pattern description:

Purpose:  **Production or Non-production**

Source:  **From workload standard**

Name	Description	Workload Type
<input checked="" type="radio"/> Departmental Transactional	For databases primarily used for online transaction processing (OLTP). The database will be optimized for transactional applications.	Departmental Transactional
<input type="radio"/> Data Mart	For databases primarily used for data warehousing. The database will be optimized for reporting applications.	Data Mart

**Select workload type**

Maximum User Data Space (GB):  **Data space dictates VM size**

Database Compatibility Mode:  **DB2 or Oracle**

Schema File:    **Schema file used for DB2**

18 DBaaS patterns © 2012 IBM Corporation

To create a database from the workload standards in “Database Patterns,” you will see a “source” field where you can select to “Apply a database workload standard” to the database you will provision. When that ‘source’ is selected, you have the option of choosing from the two workload standards. The first choice, Departmental Transactional, is the transactional standard for online transaction processing. The second choice here, Data Mart, can be used for a database that is used primarily for data warehousing. When your source is “Apply a database workload standard,” you can specify a schema file to format the database that is created; this is not a required parameter however. Another piece of information that is asked for when building your database patterns is “Purpose.” The choices there are “production” and “non-production.” This is used to determine which license should be applied to the DB2 instance to be provisioned. The ‘maximum user data space’ value dictates the size of the virtual machine that is created to host your database.

## Create new database from workload standards in virtual application builder

- Also available when deploying a database as part of a virtual application

The screenshot displays the 'Add policy for application' window in the IBM virtual application builder. On the left, an 'Enterprise Application' icon is connected to a 'Database' icon. The right pane shows the configuration for the 'Database' component. The following fields are highlighted with red boxes:

- Purpose:** Non-Production
- Source:** Apply a database workload standard
- Workload Type:** Transactional

Other visible configuration options include:

- Name:** Database
- Database Name:** mydb
- Maximum User Data Space (GB):** 10
- Database Compatibility Mode:** DB2 (Default)
- Schema File:** artifacts/setup\_db.sql

19

DBaaS patterns

© 2012 IBM Corporation

You can also deploy databases as part of a virtual application. In this case, the database is configured in the virtual application builder as seen on the slide. The workload standards are shown over on the right, highlighted in red.

## Create new database from clone

- When 'Clone' is chosen as source, a list of available images is displayed

Database Pattern

Specify options for your database pattern.

Database pattern name:

Database pattern description:

Purpose:

Source:

**From clone (DB image)**

	Image Name	Database Name	OS
<input checked="" type="radio"/>	System-created image 'oltpdb'	oltpdb	Linux
<input type="radio"/>	System-created image 'dmdb'	dmdb	Linux
<input type="radio"/>	System-created image 'martdb3'	martdb3	Linux
<input type="radio"/>	myimage1	martdb3	Linux

20 DBaaS patterns © 2012 IBM Corporation

The other choice for “Source” when creating a database, is “Clone from database image.” If that is selected, you are presented with a list of images to choose from. This is shown by the red box on the slide. In this case, there are four images available to choose from.

Note that inputs for Maximum User Data Space size, Database Compatibility and Schema are not required when using cloning. For a “clone” source, DB2 Restore is instead used to create a new database with the same license and configuration as the selected image. You will find the same “clone” option available when creating a database pattern in the virtual application builder.

## Database Patterns - operations

- From Patterns → Database Patterns, select the appropriate database pattern
- Can Deploy, Open (edit) or Delete database pattern
- Grant others access to the selected database pattern

Test DB		Operations
Database Pattern ID:	a-f1ccd8f9-49be-4452-9799-6a206b0f00cc	Deploy Open Delete
Created by:	cbadmin	
Last Modified by:	cbadmin	
Access granted to:	Administrator [owner] Add more...	Grant access to others
Database Pattern Description:	Test database	
Purpose:	production	
Source:	Apply a database workload standard	

21

DBaaS patterns

© 2012 IBM Corporation

The operations available for a Database Pattern are shown on the slide. When you select a particular database pattern that you have created, you have the option to deploy an instance of it into your cloud, open it for editing or delete the database pattern. You can also grant others access to the selected database pattern.

## Database Patterns – open (edit)

Test DB

**Database Pattern**

**Specify options for your database pattern.**

Database pattern name:

Database pattern description:

Purpose:

Source:

Name	Description	Workload Type
<input checked="" type="radio"/> Departmental Transactional	For databases primarily used for online transaction processing (OLTP). The database will be optimized for transactional applications.	Departmental Transactional
<input type="radio"/> Data Mart	For databases primarily used for data warehousing. The database will be optimized for reporting applications.	Data Mart

Maximum User Data Space (GB):

Database Compatibility Mode:

Schema File:

Same panel as when creating pattern

22
DBaaS patterns
© 2012 IBM Corporation

This slide shows the edit option for database patterns. You see you are shown the same panel as when you created the pattern.

## Database Patterns - deploy

Deployment of a database pattern creates a database instance

Test DB Deploy Open Delete

**Deploy Database from Database Pattern**

Name:

Description:

Select target environment profile:

Filter by IP type:  IPv4  IPv6

Filter by profile type:

Profile:

Cloud group:

IP group:

Select target cloud group:

IP Version:  IPv4  IPv6

Cloud group:

23

DBaaS patterns

© 2012 IBM Corporation

The deploy panel is similar to the panel you see when deploying virtual applications. Note that there is no option here to attach an SSH public key, however, as there is with virtual application pattern deployments. An SSH key can be added later from the Database Service Console if needed. Here, you need only to specify the target environment where you want to have the database deployed. This can be specified with an environment profile or a target cloud group. Any database pattern that is part of a virtual application pattern is deployed as part of the virtual application deployment. You can see an example of that in the virtual application administration education module.

## Reference

- IBM Database Patterns information center
  - [http://publib.boulder.ibm.com/infocenter/worlodep/v3r1m0/topic/com.ibm.dbapp.doc/db/dbc\\_iwd-db.html](http://publib.boulder.ibm.com/infocenter/worlodep/v3r1m0/topic/com.ibm.dbapp.doc/db/dbc_iwd-db.html)
- IBM Transactional Database Pattern and IBM Data Mart Pattern announcement
  - <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS211-474>

For additional information on the database patterns in Workload Deployer, you can reference the information center whose URL is shown on the slide. Also shown is a link to the announcements for the workload standards introduced.



---

Section

## ***Summary***

This section covers the summary.

## Summary

- Real life database usage scenarios modeled in Workload Deployer with Database Patterns
  - Database managed independent of virtual application
- Database also available as part of virtual applications
  - Life cycle is tied to the virtual application

To align with real world database usage scenarios, databases can be managed independent of any virtual application. This is modeled in IBM Workload Deployer with the Database Pattern. This is only one available database management scenario. You can also create a database belonging to your virtual application. In this case the database's life cycle will follow your virtual application's life cycle. In this presentation you saw how to enable the database patterns and how to create them. To see information on the management of deployed database patterns, see DBaaS Administration presentation.

## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, AIX, DB2, DB, PowerVM, System p, and Tivoli are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.