

IBM Business Intelligence Pattern with BLU Acceleration
Version 1.0.0

Installation and Administration Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 105.

Product Information

This document applies to IBM Business Intelligence Pattern with BLU Acceleration Version 1.0.0 and may also apply to subsequent releases.

Licensed Materials - Property of IBM

© **Copyright IBM Corporation 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	vii
Chapter 1. Product overview	1
Prerequisites	4
Release notes	4
Chapter 2. Installation checklists	7
Preinstallation checklist	7
Installation checklist	7
Post-installation checklist	8
Chapter 3. Configuring the pattern components	9
IBM Business Intelligence Pattern with BLU Acceleration component	9
User registry components	12
Other components	18
Chapter 4. Planning the deployment	19
Deploying shared services	19
Reviewing security requirements	20
Enable SSL between Cognos BI components	20
Enabling SSL for use with an LDAP directory server	20
Use IBM Cognos Application Firewall	22
Techniques for Securing Access to IBM Cognos Connection	22
Securing Data Source Signons	25
Setting the HTTPOnly parameter	25
Preparing to add supplementary fonts	26
Chapter 5. Installing and deploying the pattern	27
Downloading the command-line tool	27
Uploading the pattern type	28
Enabling the pattern	29
Creating the virtual application pattern	30
Deploying virtual application patterns	32
Provisioned virtual machines	35
Viewing virtual machine details	37
Verifying the deployed BI application.	38
Chapter 6. Deployed virtual applications	41
Post-deployment tasks.	41
Accessing the deployed BI application	41
Identifying the active Content Manager	41
Installing end user modeling tools.	42
Installing translated product documentation	42
Configuring IBM Cognos BI products to connect to a BI virtual application	43
Importing or exporting BI deployment archives	44
Rendering images in HTML and PDF reports	45
Moving a BI virtual application to a production environment.	45
Maintenance of BI virtual applications	46
Updating IBM DB2 with BLU Acceleration settings	46
Viewing virtual application instance log files	47
Determine a backup strategy	48
Diagnosing query and reporting problems	49
Updating user credentials	50
Updating the list of valid domain names and host names	51

Updating notification mail server information	51
Updating fonts	52
Updating the list of primary BlackBerry push servers	53
Performance monitoring	53
Monitoring virtual application instances	53
Monitoring virtual machines	54
Monitoring databases	54
Monitoring dynamic cubes	55
Chapter 7. Data preparation	57
ETL scenarios and best practices	57
Creating and loading data into columnar tables	57
Loading atomic facts and dimensions	57
Use stored procedures to load aggregate tables	59
Stored procedures	59
Updating IBM DB2 with BLU Acceleration settings	59
Dynamic cube performance	60
Enabling the result set cache	60
Near real-time analysis	61
Create a recent fact table to store recent transactions	61
Updating a model to merge historical and today's transactions	62
Refreshing the data cache in the recent transaction cube	63
Chapter 8. IBM Cognos Business Intelligence samples	65
Importing the IBM Cognos Dynamic Cubes sample deployments	65
Dynamic cube samples	66
Deploying sample dynamic cubes	66
Dynamic query mode sample reports	67
Cognos Mobile samples	68
The Sample Outdoors Company	69
The Sample Outdoors Group of Companies	70
Employees	72
Sales and marketing	72
Sample Outdoors database, models, and packages	73
Sample Outdoors organization and schemas	74
Warehouse Schema	77
Transaction Schemas	79
Chapter 9. Fix packs and upgrades	83
Combined product fix packs	84
Adding combined or Cognos Business Intelligence fix packs to the catalog	84
Adding combined or Cognos Business Intelligence fix packs to the catalog by using the command-line tool	85
Applying combined fix packs	86
Applying Cognos Business Intelligence fix packs	87
Uninstalling a Cognos Business Intelligence fix pack	87
DB2 fix packs	88
Adding DB2 fix packs to the catalog	88
Adding DB2 fix packs to the catalog by using the command-line tool	89
Applying DB2 fix packs	90
Uninstalling a DB2 fix pack	91
Base operating system image fix packs	91
Business Intelligence Pattern with BLU Acceleration upgrade	92
Upgrading Business Intelligence Pattern with BLU Acceleration	92
Appendix A. Accessibility features	93
Appendix B. Troubleshooting and Support	95
Troubleshoot a problem	95
Searching knowledge bases	97
Getting fixes	98

Contacting IBM Support	98
Exchanging information with IBM.	99
Troubleshoot IBM Business Intelligence Pattern with BLU Acceleration	101
Cognos user cannot log in to virtual machines	102
Deployment fails when starting IBM Business Intelligence Pattern with BLU Acceleration.	102
Notices	105
Index	109

Introduction

Use this document with IBM® Business Intelligence Pattern with BLU Acceleration, a high performance pattern that simplifies the installation and configuration of an IBM Cognos® Business Intelligence in IBM PureApplication™ System environments. The deployed virtual application instance is highly optimized to improve query performance and deliver insights for large data volumes at unprecedented speeds.

This document contains instructions for installing and deploying the pattern, maintaining the virtual application instance, and preparing data to load into the analytic repository.

Audience

Installers can use this document to install and configure BI Pattern with BLU Acceleration. Administrators can use this document to monitor their analytic system, improve model and report performance, and to configure near real-time analysis.

To use this guide, you should be familiar with the following technology:

- IBM PureApplication System environments
- database and data warehouse concepts
- reporting concepts

Finding information

To find BI Pattern with BLU Acceleration product documentation on the web, access the IBM Business Intelligence Pattern with BLU Acceleration Information Center (pic.dhe.ibm.com/infocenter/cbiima/v1r0m0/index.jsp). Release Notes are published directly to Information Centers, and include links to the latest technotes and APARs.

You can also read PDFs of the product documentation from the IBM Cognos product documentation web page (www.ibm.com/support/docview.wss?uid=swg27039326).

To find IBM PureApplication System product information on the web, access the IBM PureApplication System information Center (pic.dhe.ibm.com/infocenter/psappsys/v1r1m0/index.jsp).

To find IBM DB2® and IBM DB2 with BLU Acceleration product information on the web, access the IBM DB2 information Center (pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp).

Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

Samples disclaimer

The Sample Outdoors Company, Great Outdoors Company, GO Sales, any variation of the Sample Outdoors or Great Outdoors names, and Planning Sample depict fictitious business operations with sample data used to develop sample applications for IBM and IBM customers. These fictitious records include sample data for sales transactions, product distribution, finance, and human resources. Any resemblance to actual names, addresses, contact numbers, or transaction values is coincidental. Other sample files may contain fictional data manually or machine generated, factual data compiled from academic or public sources, or data used with permission of the copyright holder, for use as sample data to develop sample applications. Product names referenced may be the trademarks of their respective owners. Unauthorized duplication is prohibited.

Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products. BI Pattern with BLU Acceleration has accessibility features. For information about these features, see Appendix A, “Accessibility features,” on page 93.

Chapter 1. Product overview

IBM Business Intelligence Pattern with BLU Acceleration is a pattern that simplifies the installation and specialized in-memory configuration of an IBM Cognos Business Intelligence environment. It uses reusable and customizable components that collectively make an automation package for the solution deployment. This automation package decreases the time and costs of a solution deployment and shortens the time to value by automatically creating middleware assets.

The pattern deploys a virtual application instance that consists of the following products:

Table 1. Summary of products in the deployed environment

Product	Description
IBM Cognos Business Intelligence	Web-based business intelligence solution with integrated reporting, analysis, and event management features.
IBM WebSphere® Application Server	Web application server for hosting Java™ based web applications.
IBM Tivoli® Directory Integrator	Synchronizes data across multiple repositories.
IBM DB2	Content store database that is used to store customer application data, including security, configuration data, models, metrics, report specifications, and report output.
IBM DB2 with BLU Acceleration	Scalable high performance database that is used for the BI Pattern with BLU Acceleration analytic repository.

The resulting IBM Cognos BI environment integrates many business intelligence activities in one web-based solution that is optimized to improve query performance and quickly deliver insights for large data volumes.

Before you use IBM Cognos BI, understand how each of the components can help you to do your job.

Table 2. Summary of IBM Cognos BI components and what each one is used for

Component	Activity
IBM Cognos Connection	Publishing, managing, and viewing content
IBM Cognos Administration	Administering IBM Cognos software
IBM Cognos Dynamic Cubes	Modeling dynamic cubes in IBM Cognos Cube Designer
IBM Cognos Dynamic Query Analyzer	Optimizing dynamic cube performance
IBM Cognos Workspace	Providing interactive workspaces
IBM Cognos Workspace Advanced	Ad hoc querying and data exploration
IBM Cognos Report Studio	Managed reporting
IBM Cognos Event Studio	Event management and alerting

Table 2. Summary of IBM Cognos BI components and what each one is used for (continued)

Component	Activity
IBM Cognos Software Development Kit	A platform-independent automation interface for working with IBM Cognos BI services and components.

IBM Cognos Connection

IBM Cognos Connection is the portal to IBM Cognos software. IBM Cognos Connection provides a single access point to all corporate data available in IBM Cognos software.

You can use IBM Cognos Connection to work with entries such as reports, agents, and packages. You can use IBM Cognos Connection to create shortcuts, URLs, and pages, and to organize entries. You can personalize IBM Cognos Connection for your own use.

You can use IBM Cognos Connection to create and run reports and cubes and distribute reports. You can also use it to create and run agents and schedule entries.

For more information about IBM Cognos Connection, see the *IBM Cognos Connection User Guide*.

IBM Cognos Administration

Use IBM Cognos Administration to do server administration, data management, security and content administration, activities management, and portal services administration. You can access IBM Cognos Administration from the **Launch** menu in IBM Cognos Connection.

For more information about IBM Cognos Administration, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

IBM Cognos Dynamic Cubes

IBM Cognos Dynamic Cubes adds an in-memory relational OLAP component to the dynamic query mode server to provide a multidimensional view of a relational data warehouse with accelerated performance. You can then do OLAP analysis by using the Cognos Dynamic Cubes server.

For more information about IBM Cognos Dynamic Cubes, see the *IBM Cognos Dynamic Cubes User Guide*.

IBM Cognos Dynamic Query Analyzer

IBM Cognos Dynamic Query Analyzer is a query visualization tool. With it, you can visualize and troubleshoot the query logs generated when you run dynamic query reports in IBM Cognos Business Intelligence. The Aggregate Advisor component of IBM Cognos Dynamic Query Analyzer analyzes dynamic cubes to recommend aggregates that improve cube performance. Aggregate Advisor can

apply the in-database aggregates to your database tables and model automatically. You can decide whether you want to apply the in-database aggregates immediately or schedule them for a later time.

For more information about IBM Cognos Dynamic Query Analyzer, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

IBM Cognos Workspace

You can view and open favorite workspaces and reports, manipulate the content in the workspaces, and email your workspaces. You can also use comments, activities, and social software such as IBM Connections for collaborative decision making.

For more information about using IBM Cognos Workspace, see the *IBM Cognos Workspace User Guide*.

IBM Cognos Workspace Advanced

With IBM Cognos Workspace Advanced, you can do more advanced data exploration, such as adding more measures, conditional formatting, and advanced calculations. You can create reports with relational or dimensional data sources, and that show data in lists, crosstabs, and charts. You can also use your own external data source.

For more information about working in Cognos Workspace Advanced, see the *IBM Cognos Workspace Advanced User Guide*.

IBM Cognos Report Studio

Using IBM Cognos Report Studio, report authors create, edit, and distribute a wide range of professional reports. Report Studio provides powerful functionality, such as bursting, prompts, maps, and advanced charting, and provides many ways to customize reports.

Use Report Studio for following types of reports:

- Reports that are intended for a wide audience
- Reports that exist long enough to require maintenance for changing requirements and data
- Reports that require detailed control over the appearance

For more information about working in Report Studio, see the *IBM Cognos Report Studio User Guide*.

IBM Cognos Event Studio

IBM Cognos Event Studio can notify decision-makers in your organization of events as they happen so that they can make timely and effective decisions.

For information about how to use Cognos Event Studio, see the *IBM Cognos Event Studio User Guide*.

When it applies in-database recommendations, Aggregate Advisor can schedule an event to load data later. For more information about the events that are created by

Aggregate Advisor, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

IBM Cognos Software Development Kit

Using the IBM Cognos Software Development Kit, you can manage IBM Cognos processes and implement custom reporting solutions by using the cross-platform web services, libraries, and programming interfaces that are provided with the IBM Cognos Software Development Kit. As a developer, you can use the Software Development Kit to create and modify reports and queries, schedule, and deploy reports and other objects, and administer IBM Cognos Business Intelligence servers and security.

For more information about using the IBM Cognos Software Development Kit, see the *IBM Cognos Software Development Kit Developer Guide*.

Prerequisites

Review the minimum hardware and software requirements to install and deploy IBM Business Intelligence Pattern with BLU Acceleration.

Hardware requirements

You require access to:

- IBM PureApplication System W1700 Version 1.1.0.1
- a computer that uses either a Microsoft Windows, or Red Hat Enterprise Linux operating system where you can download the installation files and connect to IBM PureApplication System.

Software requirements

The official set of software requirements are available on the IBM Business Intelligence Pattern with BLU Acceleration 1.0 software environments page (www.ibm.com/support/docview.wss?uid=swg27039392). If there is a conflict between the information that is provided in the product documentation and the information in the software environments page, the information in the software environments page takes precedence.

- The base operating system image on the system must be AIX® V2.0.0.1 or later.
- For user authentication, LDAP Version 3 Compliant Directory Servers, such as IBM Tivoli Directory Server. You can also use a custom authentication provider, if one exists.
- The following shared services must be deployed in the cloud group
 - Database Performance Monitoring service
 - System Monitoring service
 - ELB proxy service

Release notes

Release notes provide a list of topics and links that are useful when you install IBM Business Intelligence Pattern with BLU Acceleration.

Contents

- “Release Notes”
- “Deploying IBM Business Intelligence Pattern with BLU Acceleration”
- “Software environments”

Release Notes

To review release notes for installation and configuration issues for IBM Business Intelligence Pattern with BLU Acceleration, see the document titled Release Notes for IBM Business Intelligence Pattern with BLU Acceleration v1.0 (www.ibm.com/support/docview.wss?uid=swg21647202).

To view other known issues and their workarounds, see the IBM Support portal (www.ibm.com/support).

For more information about finding information on the Support portal, see the IBM Support portal assistance (www.ibm.com/software/support/portal/sp-help.html).

Deploying IBM Business Intelligence Pattern with BLU Acceleration

For instructions on how to install and deploy IBM Business Intelligence Pattern with BLU Acceleration, see Chapter 5, “Installing and deploying the pattern,” on page 27.

BI Pattern with BLU Acceleration documentation is available on the IBM Business Intelligence Pattern with BLU Acceleration Information Center (pic.dhe.ibm.com/infocenter/cbiima/v1r0m0/index.jsp).

Software environments

To review an up-to-date list of environments that are supported by BI Pattern with BLU Acceleration, including information about operating systems, patches, browsers, web servers, directory servers, database servers, and application servers, see the Software Environments site (www.ibm.com/support/docview.wss?uid=swg27039392).

Chapter 2. Installation checklists

Use the installation checklists to plan for IBM Business Intelligence Pattern with BLU Acceleration installation and deployment.

Preinstallation checklist

There are some tasks that you must do before you install IBM Business Intelligence Pattern with BLU Acceleration.

- ___ • Review the supported software environments information at IBM Support Portal (www.ibm.com/support/docview.wss?uid=swg27039392).
- ___ • Review the product “Release notes” on page 4 for the latest information about issues that might impact the product installation or configuration.
- ___ • Verify that your hardware and software meet the minimum requirements. For more information, see “Prerequisites” on page 4.
- ___ • If you do not have the BI Pattern with BLU Acceleration part numbers, review the download instructions in the IBM Business Intelligence Pattern with BLU Acceleration download document. (www.ibm.com/support/docview.wss?uid=swg24035593).
- ___ • Obtain the required licenses so that you can download the software from IBM Passport Advantage®.
- ___ • Review the Chapter 4, “Planning the deployment,” on page 19 documentation.

Installation checklist

This checklist assists you in preparing to install IBM Business Intelligence Pattern with BLU Acceleration.

- ___ • Complete the following table with information for the IBM PureApplication System:

Table 3. IBM PureApplication System information required during installation

System information	Sample value	Your value
The host name or the IP address of the target system	9.24.207.202	
The user ID and password of a user that can install the BI Pattern with BLU Acceleration type on the target system.	psadmin psadminpwd	

- ___ • Get access to a computer that uses either a Windows or Linux operating system where you can access the target IBM PureApplication System and where you can copy the installation files.

Note: If you are using a computer with a Linux operating system, log in with a non-root user account.

- ___ • Verify that the computer has Java Runtime Environment (JRE) 1.6 or higher and that the PATH environment variable includes the bin directory of the JRE.

- ___ • Access Passport Advantage and download the IBM Business Intelligence Pattern with BLU Acceleration version 1.0.0 package.
- ___ • Uncompress the IBM Business Intelligence Pattern with BLU Acceleration package into an empty directory.
- ___ • Download the command line tool and upload the pattern type to the system.
- ___ • Enable the pattern type.
- ___ • Create the virtual application pattern. For information about the pattern configuration options, see Chapter 3, “Configuring the pattern components,” on page 9.
- ___ • Deploy the virtual application to create the BI Pattern with BLU Acceleration virtual application instance.

Post-installation checklist

After you install IBM Business Intelligence Pattern with BLU Acceleration on your system, you can access the virtual application instance to verify the pattern deployment and to continue setting up your environment.

- ___ • Access the virtual application instance and verify the deployment.
- ___ • Apply fix packs and interim fixes.
- ___ • Secure your virtual application instance. For more information, see “Reviewing security requirements” on page 20.
- ___ • Install and configure client modeling tools.
- ___ • “Determine a backup strategy” on page 48.
- ___ • Depending on your environment, perform other post deployment tasks as required. For more information, see “Post-deployment tasks” on page 41.

Chapter 3. Configuring the pattern components

IBM Business Intelligence Pattern with BLU Acceleration virtual applications contain middleware components that represent the Business Intelligence service and the IBM DB2 services for the BI virtual application instance.

When creating the virtual application pattern, you can choose the following components:

- “IBM Business Intelligence Pattern with BLU Acceleration component,” which represents a service to deploy IBM Cognos Business Intelligence 10.2.1.1 and IBM DB2 10.5 with BLU Acceleration. The pattern must include one IBM Business Intelligence Pattern with BLU Acceleration component.
- “User registry components” on page 12 components, which provide a user registry (namespace) for container-managed security. If you do not configure and connect a user registry to a Business Intelligence component, anonymous login to your application is enabled by default. If you use multiple user registries, each one must be connected to the IBM Business Intelligence Pattern with BLU Acceleration component that exists in your pattern.

These components include the configuration specifications and user registry specifications required for the application deployment. They also contain artifacts such as the drivers needed to connect to a content store and data sources. The configuration specifications that you choose determine how your virtual application instance is deployed, and some of these settings cannot be changed after the pattern is deployed. To ensure that you are prepared with the appropriate information, review these configuration inputs before you start the installation.

IBM Business Intelligence Pattern with BLU Acceleration component

To create an IBM Business Intelligence Pattern with BLU Acceleration virtual application, you configure the component in the Virtual Application Builder. You can also connect the IBM Business Intelligence Pattern with BLU Acceleration component to multiple existing user registry components.

General configuration settings

You can specify the following general configuration settings:

Name A unique name that identifies the BI component in the Virtual Application Builder.

Virtual host name

A unique name that the ELB Proxy Service uses to relay incoming requests to an IBM Business Intelligence Pattern with BLU Acceleration virtual application instance. This value represents the web alias name that is used to construct entry points to the application instance. For example, if the virtual host name is `ibmcognos`, then the entry points to the application would be `http://ibmcognos/p2pd` and `http://ibmcognos/p2pd/servlet/dispatch`.

To change this setting after the pattern is deployed, you must back up any data on the active Content Manager server, terminate your deployment, change the name, and then redeploy the pattern.

Enable SSL

Select this setting to enable the Secure Sockets Layer (SSL) protocol for communication between the ELB Proxy Service and all Cognos BI components.

Valid domains or hosts

The list of domain names or host names that are separated by commas. This list is used by IBM Cognos Application Firewall (CAF) to validate requests with host names and domain names. For example, *.mycompany.com, hostname, hostname:80.

Notification

You can set up your BI virtual application for email notification by selecting the **Notification** check box and specifying the following attributes:

Notification mail server and port number

The host name and port number of the mail server computer. The syntax is host:port. For example, mailservername:25.

Notification user ID

The user ID required to log in to the email account. If the mail server does not require authentication, leave it blank.

Notification password

The password that is required for the email account that is used to log in to the mail server. If the mail server does not require authentication, leave it blank.

Notification default sender

The email address of the sender of outgoing messages. For example, cognosbireports@ibm.com.

Administrative credentials

You must provide user credentials with administrative privileges. BI Pattern with BLU Acceleration virtual applications use these credentials to administer, optimize, and maintain the deployed virtual applications. To specify administrative credentials, use the following attributes:

Administrative namespace ID

The identifier of the namespace that is associated with a user with administrative privileges. The namespace ID must be the same as the namespace ID defined in the Existing User Registry component; or if a user registry is not configured yet, specify NONE.

Administrative user name

The user name with administrative privileges for deploying and administering BI Pattern with BLU Acceleration applications. The user name must already be defined on the LDAP server. If you define multiple user registry components, the user name must be valid on only one of them.

Administrative password

The password for the user name that is used for deploying and administering BI Pattern with BLU Acceleration application instances.

Mobile configuration

The pattern includes IBM Cognos Mobile files that enable your application to run on mobile devices.

To enable BlackBerry users to receive content notification, you must specify the following attribute:

BlackBerry MDS host names and port numbers

Specifies the host names and port numbers for the BlackBerry Mobile Data Service (MDS) Connection Service primary push servers. This information enables content notifications to be delivered to BlackBerry users. For example: server1:8080, 1.1.1.1:5050. You can update this value after deployment.

Fonts

If your environment requires fonts that are not installed with IBM Cognos Business Intelligence, you can install a set of supplementary fonts. You must package these additional fonts in a compressed file and then specify the file name in the following attribute:

Supplementary fonts

The name of the compressed (zip) file containing additional font files that can be used when authoring reports. The file must contain the complete list of additional fonts. The supported font file extensions are ttf or ttc.

For more information about fonts, see “Preparing to add supplementary fonts” on page 26.

External data source

The pattern deployment creates a data source connection to the source database.

External data source name

Specifies the name of the data source that is used to define the physical connection to the source database for the IBM Business Intelligence virtual application. Data sources are stored in the Cognos namespace and must be unique.

Internal data store

The pattern deployment creates a database and data source connection to the target database.

Internal data store name

Specifies the name of the database that is used for reporting and data analysis.

Deployment

You can set limitations for the allocation of server resources by specifying the following attributes:

Maximum number of cores

Specifies the maximum number of virtual processors that are used by your application. Acceptable value range is 12 - 156, in increments of 4.

Maximum amount of memory (GB)

Specifies the maximum amount of virtual memory, in gigabytes, that is used by your application. Acceptable value range is 100 - 740, in increments of 20.

IBM Cognos BI Samples Deployment

You can install the IBM Cognos Business Intelligence Samples as part of the pattern deployment.

Note: The documentation refers to the samples to explain how to use some product features. It is helpful to install the samples when you are setting up a development environment.

IBM Cognos BI Samples deployment

Specifies whether to install the IBM Cognos BI Samples with the pattern deployment.

The pattern installs the samples, configures the datasource connection and the package in Cognos Connection, and sets up the database for reporting. To work with the samples model in IBM Cognos Cube Designer, you must download the model from the DownloadClient endpoint and deploy the cube to the COGP virtual machine. For more information, see “Dynamic cube samples” on page 66.

User registry components

You define the user authentication sources for your BI applications by configuring and then connecting user registry components, also referred to as a namespace, to a IBM Business Intelligence Pattern with BLU Acceleration component. When user registry components are not defined in your application pattern, anonymous login to your virtual application is enabled by default.

External authentication resources must exist in your environment and be configured for use before you can use them in the pattern. For information about LDAP configuration and custom Java authentication provider, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

If your applications require LDAP attributes that are not included in the IBM Tivoli Directory Server or the Microsoft Active Directory user registry component, use the Generic LDAP Directory Server component.

IBM Tivoli Directory Server

An IBM Tivoli Directory Server user registry component represents an existing Lightweight Directory Access Protocol (LDAP) service that you can connect to a BI component. This LDAP service provides a user registry, also referred to as a namespace, for container-managed security.

You can specify the following attributes for each user registry:

Name A unique name that represents the name of the namespace. Specify this name for maintenance operations on a user registry.

Namespace ID

A value that identifies this namespace. This value is used internally at run time by the IBM Business Intelligence component and must be unique across all user registries.

Server and port number

The host name and the port number for the directory server, which is separated by a colon. For example, `ldapservname:389`. Instead of the host name, you can specify the IP address.

Base distinguished name

The base distinguished name of the LDAP server. For example, `o=acme,c=us`. These entries must exist in the LDAP server.

User lookup

The user lookup information that is used for binding to the LDAP directory server. For example, `(cn=${userID})`.

Bind user DN

The user ID used for binding to the LDAP server when you search using the user lookup property, or when you perform operations that use external identity mapping.

Password

The password for the user ID used for binding to the LDAP server.

Use bind credentials for search

Specifies whether to use the bind credentials to perform a search. If this property is checked, the bind credentials provided in the namespace configuration will be used to perform a search in the LDAP directory server. If this property is unchecked, or bind credentials are not presented, the authenticated user credentials will be used for searching.

SSL certificate database

Specifies the name of the `.zip` file containing the certificate database files required by the directory server for establishing SSL connections. You can update this field after deployment.

Microsoft Active Directory

A Microsoft Active Directory user registry component represents an existing LDAP service that can be connected to a BI component. This LDAP service provides a user registry, also referred to as a namespace, for container-managed security.

You can specify the following attributes for each user registry:

Name A unique name that represents the name of the namespace. Specify this name for maintenance operations on a user registry.

Namespace ID

A value that identifies this namespace. This value is used internally at run time by the IBM Business Intelligence component and must be unique across all user registries.

Server and port number

The host name and the port number for the directory server, which is separated by a colon. For example, `ldapservname:389`. Instead of the host name, you can specify the IP address.

Base distinguished name

The base distinguished name of the LDAP server. For example, `o=acme,c=us`. These entries must exist in the LDAP server.

User lookup

The user lookup information that is used for binding to the LDAP directory server. For example, `(cn=${userID})`.

Bind user DN

The user ID used for binding to the LDAP server when you search using the user lookup property, or when you perform operations that use external identity mapping.

Password

The password for the user ID used for binding to the LDAP server.

Use bind credentials for search

Specifies whether to use the bind credentials to perform a search. If this property is checked, the bind credentials provided in the namespace configuration will be used to perform a search in the LDAP directory server. If this property is unchecked, or bind credentials are not presented, the authenticated user credentials will be used for searching.

SSL certificate database

Specifies the name of the .zip file containing the certificate database files required by the directory server for establishing SSL connections. You can update this field after deployment.

Generic LDAP Directory Server

A Generic LDAP Directory Server user registry component represents an existing LDAP service that can be connected to a BI component. This LDAP service provides a user registry, also referred to as a namespace, for container-managed security.

You can provide the following attributes for each user registry:

Name A unique name that represents the name of the namespace. Specify this name for maintenance operations on a user registry.

Namespace ID

A value that identifies this namespace. This value is used internally at run time by the IBM Business Intelligence component and must be unique across all user registries.

Server and port number

The host name and the port number for the directory server, which is separated by a colon. For example, `ldapservname:389`. Instead of the host name, you can specify the IP address.

Base distinguished name

The base distinguished name of the LDAP server. For example, `o=acme,c=us`. These entries must exist in the LDAP server.

User lookup

The user lookup information that is used for binding to the LDAP directory server. For example, `(cn=${userID})`.

External identity mapping

The mapping that is used to construct a distinguished name (DN) or a search filter to locate a user in the LDAP directory server. If the application cannot find the environment variable or the DN in the LDAP server, it attempts to use the **User lookup** property to construct the DN.

If you specify `${environment("ENVIRONMENT_VARIABLE_NAME")}`, all instances of this string are replaced by the value of the environment variable that is provided by the web server. If the string does not begin with an open parenthesis, the result of the substitution is assumed to be the user's DN. For example, if you specify

`uid=${environment("REMOTE_USER")},ou=people, base DN`, where *base DN* is the Base Distinguished Name parameter value. If the value begins with an open parenthesis, the result of the substitution is assumed to be a search filter. For example, `(userPrincipalName=${environment("REMOTE_USER")})`. In this case, you must either enable anonymous access to the LDAP directory server or set the **Bind user DN** and **Password** properties.

Bind user DN

The user ID used for binding to the LDAP server when you search using the user lookup property, or when you perform operations that use external identity mapping.

Password

The password for the user ID used for binding to the LDAP server.

Use bind credentials for search

Specifies whether to use the bind credentials to perform a search. This property only affects users who do not use the external identity mapping. If this property is checked, the bind credentials provided in the namespace configuration will be used to perform a search in the LDAP directory server. If this property is unchecked, or bind credentials are not presented, the authenticated user credentials will be used for searching.

Unique identifier

The value that is used to uniquely identify objects that are stored in the LDAP directory server.

SSL certificate database

Specifies the name of the .zip file containing the certificate database files required by the directory server for establishing SSL connections. You can update this field after deployment.

The Generic LDAP Directory Server user registry component provides the following groups of extra attributes:

- For folder mappings:

Name The LDAP attribute that is used for the “name” property of a folder.

Object class

The name of the LDAP object class that is used to identify a folder.

Description

The LDAP attribute that is used for the “description” property of a folder.

- For group mappings:

Name The LDAP attribute that is used for the “name” property of a group.

Object class

The name of the LDAP object class that is used to identify a group.

Description

The LDAP attribute that is used for the “description” property of a group.

Member

The LDAP attribute that is used for the “member” property of a group.

- For account mappings:

Name The LDAP attribute that is used for the “name” property of an account.

Account object class

The name of the LDAP object class that is used to identify an account.

Description

The LDAP attribute that is used for the “description” property of an account.

Product locale

The LDAP attribute that is used for the “productLocale” property of an account.

Content locale

The LDAP attribute that is used for the “contentLocale” property of an account.

User name

The LDAP attribute that is used for the “userName” property of an account.

Password

The LDAP attribute that is used for the “password” property of an account.

Surname

The LDAP attribute that is used for the “surname” property of an account.

Given name

The LDAP attribute that is used for the “givenName” property of an account.

Postal address

The LDAP attribute that is used for the “postalAddress” property of an account.

Email The LDAP attribute that is used for the “email” property of an account.

Business phone

The LDAP attribute that is used for the “businessPhone” property of an account.

Home phone:

The LDAP attribute that is used for the “homePhone” property of an account.

Mobile phone

The LDAP attribute that is used for the “mobilePhone” property of an account.

Pager phone

The LDAP attribute that is used for the “pagerPhone” property of an account.

Fax/Phone

The LDAP attribute that is used for the “faxPhone” property of an account.

- For advanced authentication properties:

Advanced properties

The number of authentication properties that you want to set. Blank entries are ignored.

Name The name of the advanced property for the user registry. For example, `multitenancy.TenantPattern`.

Value The value that is associated with the advanced property.

- For custom LDAP properties:

Custom properties

The number of authentication properties that you want to set. Blank entries are ignored.

Name The name of the custom property. For example, dept.

Value The value that is associated with a custom property.

- For custom files:

Custom files

The name of the compressed file that contains custom files that are required by your Custom Java Provider user registry. When you deploy your pattern, the custom files are copied in the file structure that is specified in the compressed file.

Custom Java Provider

An Existing User Registry (Custom Java Provider) component represents a custom Java authentication provider that can be attached to an IBM Business Intelligence Pattern with BLU Acceleration component. The custom Java authentication provider provides a user registry, also referred to as a namespace, for container-managed security. This external resource must exist in your environment and be configured for use in authentication.

You can provide the following attributes for each user registry:

Name A unique name that represents the name of the namespace. Specify this name for maintenance operations on a user registry.

Namespace ID

A value that identifies this namespace. This value is used internally at run time by the IBM Business Intelligence component and must be unique across all user registries.

Java class name

The name of the Java class to authenticate users.

Selectable for authentication

When this option is selected, application users can select this user registry on the login page of the application.

The Custom Java Provider user registry component provides the following groups of extra attributes:

- For advanced authentication properties:

Advanced properties

The number of authentication properties that you want to set. Blank entries are ignored.

Name The name of the advanced property for the user registry. For example, multitenancy.TenantPattern.

Value The value that is associated with the advanced property.

- For custom files:

Custom files

The name of the compressed file that contains custom files that are

required by your Custom Java Provider user registry. When you deploy your pattern, the custom files are copied in the file structure that is specified in the compressed file.

Other components

Use these components for inbound and outbound TCP connections.

Connect In

This component is used to open the firewall for inbound TCP connections from a specified address or range of addresses, to a specified port in the target application component.

The following are the attributes for the component:

Name Specifies the name of the inbound connection.

Inbound Server(s) & Port(s)

Specifies the destination port on target servers for inbound connections from the application. Enter a port number or a range of port numbers. Separate multiple ports with a comma. If you specify port numbers only, the specified ports are open for any IP address. For example: 14121, 65432, 23411, 42111-42113.

If you want to open ports for specific IP addresses only, enter an IP address or netmask with the port number. You can specify IP addresses in IPv4 or IPv6 format. For example: 1.2.3.4:12311, 45221, 52341, 1.3.5.2/23:12311, [311:33:12:12]:1231.

Connect Out

This component is used to open the firewall for outbound TCP connections from a web or enterprise application to a specified host and port.

The following are the attributes for the component:

Name Specifies the name of the outbound connection.

Outbound Server(s) & Port(s)

Specifies the destination port on target servers for outbound connections from the application. Enter a port number or a range of port numbers. Separate multiple ports with a comma. If you specify port numbers only, the specified ports are open for any IP address. For example: 14121, 65432, 23411, 42111-42113.

If you want to open ports for specific IP addresses only, enter an IP address or netmask with the port number. You can specify IP addresses in IPv4 or IPv6 format. For example: 1.2.3.4:12311, 45221, 52341, 1.3.5.2/23:12311, [311:33:12:12]:1231.

Chapter 4. Planning the deployment

The information provided in this section will help you to plan a successful installation and be prepared with the configuration information that will be required during the installation process.

Use the following information to help you plan your deployment and collect the required information before you begin.

1. Plan your deployment by reviewing the “Prerequisites” on page 4 and Chapter 2, “Installation checklists,” on page 7.
2. Determine whether your application requires user authentication. By default, anonymous login is enabled if there is no user registry attached to the IBM Business Intelligence Pattern with BLU Acceleration component. With an existing authentication namespace, when you configure the pattern, use the administrative credentials that your component is pointing to, otherwise the deployment will fail.

Complete the following table if you are using an existing authentication namespace in your pattern deployment.

Table 4. Configuration information when using an existing authentication namespace

Information	Your value
Administrative namespace ID	
Administrative username	
Administrative password	

3. Review the pattern configuration options and collect the required information before starting your installation.
4. Determine which cloud group will contain the virtual application instance. Verify that the cloud group has the following characteristics:
 - • The cloud group has a running instance of the required shared services. For information about which shared services are required, see “Deploying shared services.”
 - • The cloud group has the required CPU, memory, disk, and IP addresses to run the application.
5. Determine whether you want to configure an RSA key pair to enable SSH key based access to the virtual machines. You can either provide your own externally-generated public key or use a key pair that is generated by the system.
6. Consider your security requirements and decide if you want to enable the Secure Socket Layer (SSL) protocol. If you enable SSL, it must be on port 636.

For information about installing and deploying the pattern, see Chapter 5, “Installing and deploying the pattern,” on page 27.

Deploying shared services

You must deploy some shared services before you can deploy IBM Business Intelligence Pattern with BLU Acceleration in your cloud group.

You must deploy the following services:

Database Performance Monitoring service

The Database Performance Monitoring service provides information on the health and availability of databases that are deployed in the IBM PureApplication System. This information provides early warning of degraded database performance. The database performance monitor also provides drill-down capability to determine the root cause of performance problems.

System Monitoring service

The System Monitoring service provides the monitoring infrastructure that allows monitoring agents to collect performance and availability information. You can use this information to troubleshoot your system and make business decisions about your hardware and service usage.

ELB proxy service

Elastic load balancing (ELB) is a shared service that provides routing and load balancing to deployed applications.

For information about how to deploy the previously listed services, see the Managing existing shared services topic in the IBM PureApplication System 1.1 information center (pic.dhe.ibm.com/infocenter/psappsys/v1r1m0/topic/com.ibm.puresystems.appsys.1700.doc/iwd/apt_existing_shareservices.html).

Reviewing security requirements

Using the built-in security features, you can configure an IBM Business Intelligence Pattern with BLU Acceleration virtual application instance for maximum security. Consider your environment before deciding which security features to implement.

Typical factors that influence the decision of how secure a system must be, include

- external exposure
Who are the users? Are there internal users as well as external or guest users?
- use of the public Internet
Is the system accessible by the Internet? Does a virtual private network (VPN) exist?
- data sensitivity
Departments such as human resources, finance, and accounting likely want the data protected the best way possible.

Enable SSL between Cognos BI components

You can use the Secure Sockets Layer (SSL) protocol for communications between IBM Cognos BI components and other servers.

When you create the pattern in Virtual Application Builder, you can enable SSL between Cognos BI components by checking the **Enable SSL** checkbox. When this option is checked, the pattern turns on SSL between the Elastic Load Balancing (ELB) service and Cognos BI, turns on SSL between all BI components, and generates application endpoints that use HTTPS. The default port number for HTTPS is 443.

Enabling SSL for use with an LDAP directory server

You can use the Secure Sockets Layer (SSL) protocol for communications with an LDAP directory server.

Before you begin

Enabling SSL requires an administrator with a good knowledge of the public key infrastructure (PKI) technology and SSL.

You must install a server certificate. For information about installing a server certificate, refer to the documentation for your directory server.

About this task

Secure LDAP protocol (LDAPS) encrypts the communication between the Access Manager component of Content Manager and the LDAP directory server. This prevents sensitive information in the directory server and the LDAP credentials from being sent as clear text.

To enable Secure LDAP protocol, you must

- install a server certificate that is either self-signed or signed by a certificate authority (CA) in your directory server
- create a certificate database to contain the certificates
- configure the pattern to use an LDAP namespace with SSL enabled

The server certificate must be a copy of either

- the trusted root certificate and all other certificates that make up the chain of trust for the directory server certificate.

The trusted root certificate is the certificate of the root certificate authority that signed the directory server certificate.

- the directory server certificate only

IBM Cognos works with both the cert8.db and cert7.db versions of the client certificate database. You must use the certutil tool from Netscape OpenSource toolkit NSS_3_11_4_RTM to create the certificate database. IBM Cognos does not accept other versions of cert8.db files, including those from the certutil tool that is provided with Microsoft Active Directory. The appropriate certutil tool is available from the FTP Web site at Mozilla.

You must also use version 4.6.7 of the NSPR library, which is available from the FTP Web site at Mozilla.

The default port for use with LDAPS is 636.

If you want to enable SSL in a deployed environment, you must back up your data, stop the deployed instance, change the setting, and then redeploy the pattern.

Procedure

1. Create a directory for the certificate database.

2. Create the certificate database by typing

```
certutil -N -d certificate_directory
```

where *certificate_directory* is the directory that you created in step 1.

This command creates a cert8.db file, a key3.db file, and a secmod.db file in the new directory.

3. Add the certificate authority (CA) certificate or the directory server certificate to the certificate database by typing the appropriate command for the type of certificate:

- For a CA certificate, type

```
certutil -A -n certificate_name -d certificate_directory -i CA.cert -t C,C,C
```
- For a directory server certificate, type

```
certutil -A -n certificate_name -d certificate_directory -i server_certificate.cert -t P
```

where *certificate_name* is an alias that you assign, such as the CA name or host name; and *server_certificate* is the prefix of the directory server certificate file.

4. Create a .zip file that contains all of the files in the certificate database directory.

What to do next

When building the pattern in Virtual Application Builder, ensure the user registry (namespace) configuration includes the following settings:

- **SSL certificate database** contains the name of the .zip file that contains the certificate database files.
- **Server and port number** specifies port 636.

Use IBM Cognos Application Firewall

IBM Cognos Application Firewall (CAF) supplements the existing IBM Cognos BI security infrastructure. By default, this supplemental security is enabled. To ensure that the IBM Cognos BI solution is secure, CAF should never be disabled in a production environment.

CAF protects the IBM Cognos BI components from processing malicious data. The most common forms of malicious data are buffer overflows and cross-site scripting attacks (XSS links), either through script injection in valid pages or redirection to other Web sites.

In a distributed environment, all CAF settings must be the same for all computers where IBM Cognos BI Application Tier Components are installed. If CAF is disabled on some computers and enabled on others, unexpected behavior and product errors may result. In IBM Business Intelligence Pattern with BLU Acceleration virtual applications, the application tier components are installed on the COGP virtual machine.

You can add host and domain names to the IBM Cognos list of valid names when building the pattern in the virtual application builder. For information about changing the valid domains in a deployed environment, see “Updating the list of valid domain names and host names” on page 51.

You can also track firewall activity by checking the log file, which contains rejected requests. By default, log messages are stored in the /c10_64/logs/cogserver.log file on the COGP virtual machine.

Techniques for Securing Access to IBM Cognos Connection

When you add an authentication provider to your IBM Business Intelligence Pattern with BLU Acceleration environment, all users in the directory have access to IBM Cognos Connection. To secure IBM Cognos BI, you must restrict this access.

The methods and best practices discussed in this section apply mostly to the LDAP authentication providers, Tivoli directory server and Active Directory.

Use whichever method applies to your organization:

- LDAP groups or roles

Use this method when the user accounts are located in different branches of the directory tree.

- LDAP OUs

Use this method when the user accounts are located in a specific Organizational Unit (OU) in the directory tree.

An Organizational Unit (OU) is a type of container in an LDAP directory structure. OU can contain user accounts, groups, roles, and other OUs.

Using LDAP Groups or Roles to Restrict Access

Not all users in your LDAP directory must use IBM Cognos BI. Grant only designated users access to IBM Cognos Connection. This can be done by creating an IBM Cognos BI-specific group or role in your directory server, adding the required users to its membership, and granting the group or role access to IBM Cognos Connection.

Whether you must create a group or a role depends on whether your authentication provider uses group membership or role membership as part of its user account information.

An alternative method of restricting access is based on using the LDAP organizational units (OUs).

The groups and roles used for this technique are created using your directory server software. For more information about creating groups or roles, see the directory server documentation.

When building the pattern in Virtual Application Builder, ensure that the following parameters are properly defined when configuring the LDAP namespace.

- **User lookup**

Configure the user lookup string for the selected LDAP provider to contain the attribute that will be used to authenticate against the `${userID}` variable. This variable takes the user name entered at logon and substitutes the variable with the value before passing the search string to the directory server. The distinguished name (DN) of the role must also be included in the string.

Here is an example of the lookup string using role based membership:

```
(&(uid=${userID})(nsrole=cn=Cognos,ou=people,dc=cognos,dc=com))
```

Here is an example of the lookup string using group membership:

```
(&(sAMAccountName=${userID})(memberOf=cn=ReportNet,ou=Groups,dc=cognos,dc=com))
```

- **External identity mapping**

Construct a string to locate a user in the LDAP directory server. At logon time, the environment variable `${environment("REMOTE_USER")}` in this string is replaced by the user name and then the string is passed to the directory server.

In the following example, the Web browser sets the environment variable `REMOTE_USER` that matches the user's `uid` attribute. Instead of substituting the hard-coded `sAMAccountName` value with `${userID}`, the environment variable is read from the browser session.

```
(&(sAMAccountName=${environment("REMOTE_USER")})(memberOf=cn=Cognos,cn=Groups,dc=cognos,dc=com))
```

Using LDAP OUs to Restrict Access

You can grant access to IBM Cognos Connection for a particular Organizational Unit (OU) or children of a particular OU in an LDAP directory. An OU usually represents a segment of an organization.

For this method to work, you must properly set up the **Base Distinguished Name** and **User lookup** properties for the LDAP authentication provider when building the pattern in Virtual Application Builder. By using different values for these properties, you can grant access for different OUs in your LDAP directory structure.

Consider the following directory tree:




Figure 1. Tree of organizational units of a fictitious company broken down by east and west

If users from only the East OU need access to IBM Cognos Connection, the values can be specified as listed in the following table.

Table 5. The Base Distinguished Name and User Lookup values for the East Organizational Unit

Property	Value
Base Distinguished Name	ou=East, ou=people, dc=abc, dc=com
User lookup	uid=\${userID}

If users from both East and West OUs require access, the values can be specified as listed in the following table.

Table 6. The Base Distinguished Name and User Lookup values for the East and West Organizational Units

Property	Value
Base Distinguished Name	ou=people, dc=abc, dc=com
User lookup	(uid=\${userID})

The parentheses () in the **User lookup** property are used as a filter that can search all OUs located under the specified Base DN. In the first example, only the East OU is searched for user accounts. In the second example, both the East and West OUs are searched.

However, in both of the above examples, groups are excluded from access to IBM Cognos Connection because they are located in a different branch of the directory tree than users. To include both the groups and users, the Base DN must be the root of the directory tree. The values would then be as listed in the following table.

Table 7. The Base Distinguished Name and User Lookup values for groups and users located at the root of the directory tree

Property	Value
Base Distinguished Name	dc=abc, dc=com
User lookup	(uid=\${userID})

As a result, all users in the directory have access to IBM Cognos Connection.

The last example shows that using OUs may not always be the most efficient way of securing access to IBM Cognos Connection. You can use this method if you want to grant access for all users in a particular OU. If you want to grant access only for specific users, you may want to consider creating a designated IBM Cognos BI group or role in your directory server, and granting this group or role access to IBM Cognos Connection.

Securing Data Source Signons

IBM Cognos BI uses data source signons to access the underlying data located in the reporting databases.

To protect against unwanted access to the data,

- grant execute permissions for users and groups that need the signons
Other types of permissions are not required.
- deny execute permissions explicitly for all users, groups, and roles that do not need the signons

This ensures that access is not permitted through an unknown membership.

For more information about data source signons, see the *IBM Cognos Administration and Security Guide*.

Setting the HTTPOnly parameter

Turning on the HTTPOnly parameter helps mitigate against malicious scripts attempting to steal a user's session identity.

Before you begin

Administrators who want to enable this attribute must ensure that users have a web browser that supports the HTTPOnly attribute.

About this task

CAM Passport identifies a user's session to the server. Setting the HTTPOnly parameter instructs a user's web browser to block scripts from being able to read or manipulate the CAM Passport cookie. You set this parameter in IBM Cognos Administration.

Procedure

1. Go to IBM Cognos Administration.
2. On the **Status** tab, click **System**.
3. In the **Scorecard** pane, from the **System** drop-down menu, click **Set properties**.
4. Click the **Settings** tab.
5. From the **Category** list, select **Environment**.

6. For the **HTTPOnly Cookie Support** parameter, select the check box in the **Value** column to turn HTTPOnly Cookie Support on. In the Default column, **No** indicates that the selected setting is different than the default value.
7. Click **OK**.

Preparing to add supplementary fonts

If your environment requires fonts that are not installed with IBM Cognos Business Intelligence 10.2.1, you can install a set of supplementary fonts. These additional fonts become automatically available when authoring reports.

You associate new fonts to your BI Pattern with BLU Acceleration virtual application by uploading a compressed file that contains the complete list of the additional fonts. These fonts are then associated with your virtual application instances. You can add font files at deployment time or later.

If you add fonts after your virtual application is deployed or if you refresh existing fonts, ensure that the compressed file contains all the fonts that were previously associated with the virtual application, even if the fonts have not changed.

Procedure

1. Create a compressed file with the additional TrueType font files. The supported file extensions are `ttf` or `ttc`.
2. From your system, use the Virtual Application console to upload the compressed file. Depending on whether your pattern is deployed, there are two ways to upload the compressed file:
 - If your BI virtual application instance is deployed, you can update the font files. For information about updating the fonts files, see “Updating fonts” on page 52.
 - If your BI virtual application is not deployed, you can add the compressed font file while creating the virtual application pattern.

Chapter 5. Installing and deploying the pattern

Review the steps required to install and deploy IBM Business Intelligence Pattern with BLU Acceleration.

Installation and deployment of the BI Pattern with BLU Acceleration consists of the following steps:

- Before deploying the pattern, ensure that you are prepared with the required information by reviewing the pattern configuration parameters.
- Download the command-line tool and use it to upload the pattern type.
- Enable the pattern to make it available for use.
- Create the Business Intelligence virtual application pattern that you want to deploy using Virtual Application Builder.
- Deploy the virtual application pattern to the cloud group where the shared services are running.
- Apply the latest fix packs in all environments.
- Access and test the BI virtual application.
- Determine a backup strategy.

For more information about post deployment tasks and maintenance operations, see Chapter 6, “Deployed virtual applications,” on page 41.

Downloading the command-line tool

The command-line tool is required to upload pattern types that are larger than 2 GB.

Before you begin

Before you download the command-line tool, ensure that the following prerequisites are met.

- Java Runtime Environment (JRE) V6 is installed on the machine on which you will run the command line tool.
- The PATH environment variable includes the path of Python.exe.
- A class path environment variable is not defined.
- JAVA_HOME points to the location where Java Runtime Environment (JRE) is installed.

About this task

The command-line tool can be downloaded from either the IBM PureApplication System or IBM Workload Deployer user interface. It is currently supported on Windows and Linux systems.

Procedure

1. Download the command-line tool from the user interface. From the Welcome page, click **Download tooling** > **Download command line tool**
2. Save the ZIP file to your local hard disk drive.

3. Extract the contents of the ZIP file to a directory on your hard disk. When extracted using IBM PureApplication System, a single directory called `pure.cli` is created. When extracted using IBM Workload Deployer, a single directory called `deployer.cli` is created.
4. Ensure that either the `JAVA_HOME` or the `PATH` environment variable is set to the location of your JRE.
5. Optional: If you are using Windows Server 2003 or 2008, perform this step. In the directory where you extracted the contents of the ZIP file, create a registry file in the `lib/<version>` with the following line:

```
python.os=nt
```

This causes Jython to bypass the normal operating system detection logic and treat the system as a Windows machine.

By default, the only subdirectory in the `lib` directory is a `<version>` subdirectory that matches the firmware level of the product from which the command-line tool was downloaded. If you use this command-line tool installation to communicate with different firmware levels, then there is one subdirectory under the `/lib` directory for each of these firmware levels. You must copy the registry file into each of these subdirectories, for example, `/lib/3.0.0.0-12345/registry`.

Results

Now that you have installed the command-line tool, you can verify the installation by running the shell script from the `install` directory.

If you are using IBM PureApplication System, run either **pure.bat** on Windows or **pure** on Linux from the `pure.cli/bin` directory.

If you are using IBM Workload Deployer, run either **deployer.bat** on Windows or **deployer** on Linux from the `deployer.cli` directory.

What to do next

Upload the pattern type using the command-line tool.

Uploading the pattern type

Use the command-line interface on a local machine to upload the IBM Business Intelligence Pattern with BLU Acceleration pattern type. This is required for all pattern types that are larger than 2 GB.

Before you begin

You must download and extract the command-line tool before you can use it to upload the pattern type.

About this task

The command-line tool can be run in either interactive mode or batch mode from a local machine running either a Windows or Linux operating system. This procedure covers the steps to run the command-line tool in interactive mode. For information about running the command-line tool in batch mode, see the IBM PureApplication System documentation.

Procedure

1. On the command line or shell, change the directory to the installation location of the command-line tool.
2. Use the following parameters to invoke the command-line tool in interactive mode.

-u <userID>

Specifies the user ID that is used to authenticate to the system. Use the same user ID and password that you use to log in to the system GUI.

-p <password>

Specifies the password that is used to authenticate to the system. Use the same user ID and password that you use to log in to the system GUI. You are not required to specify a password to access the help for the command-line tool.

-h <hostName>

Specifies the host name or the IP address of the system. If you specify this option, do not use the URL to access the Web interface. If this parameter is not specified, the command-line tool uses the value of the `host_name` environment variable to determine the host name.

-a Automatically accepts the certificates from the system. You can also format the parameter as `--acceptcert`.

Usage examples for IBM PureApplication System

In Windows:

```
.\pure -h <hostName> -u <userID> -p <password> -acceptcert
```

In Linux:

```
./pure -h <hostName> -u <userID> -p <password> --acceptcert
```

Usage examples for IBM Workload Deployer

In Windows:

```
.\deployer -h <hostName> -u <userID> -p <password> --acceptcert
```

In Linux:

```
./deployer -h <hostName> -u <userID> -p <password> --acceptcert
```

3. Upload the pattern type using the following command:
`deployer.patterntypes.create('<path>\<pattern_filename.tgz>')`

The time required to complete the upload varies on the file size and the network bandwidth.

4. After the pattern upload is complete, click **Cloud > Pattern Types** and search for the name of the .tgz file in the search box. Information about the pattern type appears in the pattern description pane.

What to do next

Once the pattern type has been uploaded, enable the pattern type by accepting the license agreement.

Enabling the pattern

To start using IBM Business Intelligence Pattern with BLU Acceleration, you must accept the license agreement and enable the pattern type.

Before you begin

Important: IBM Foundation Pattern is a prerequisite to using all other pattern types. The foundation pattern must be made available before you can accept the license agreement for the web application pattern type. For additional information about the IBM Foundation Pattern, see your IBM PureApplication System documentation.

About this task

The IBM Business Intelligence Pattern with BLU Acceleration license agreement encompasses a combination of product licensing. You can view the license agreement before and after you accept it.

Procedure

To accept the license agreement follow these steps:

1. Click the **Workload Console** tab on the Welcome page to open the workload console.
2. Click **Cloud > Pattern Types**.
3. In the list of pattern types, select the IBM Business Intelligence Pattern with BLU Acceleration pattern type.
4. In the **License Agreement** field, click **Accept**.
5. After reading the license agreement, click **Accept**. The license agreement changes to **Accepted**.
6. Click **Enable** to change the pattern type status to **Available**.

Results

You have accepted the license agreement and enabled the pattern type.

What to do next

You can now use the pattern to start building the application.

Creating the virtual application pattern

Creating a virtual application consists of using the Virtual Application Builder to define and create the virtual application templates and patterns, and then deploying the pattern in the cloud environment. A deployed virtual application pattern is a virtual application instance.

Before you begin

You must be granted access to patterns, access to create patterns, or have the **Workload resources administration** role with full permissions to complete this task.

About this task

When you create a virtual application, you first select a pattern type such as the IBM Business Intelligence Pattern with BLU Acceleration. You can then use Virtual Application Builder to create and deploy a virtual application pattern using the components associated with the selected pattern type.

For example, rather than installing, configuring, and creating a connection to a specific instance of a database, you can specify the need for a database and provide the associated database schema in your virtual application pattern. The database instance and the connection is then created for you by the virtual application pattern.

Note: When you start the Virtual Application Builder, it opens in a separate window. Depending on your browser settings, you might need to configure pop-up windows to open in a new tab so that the Virtual Application Builder displays correctly.

Procedure

1. In the **Workload Console**, click **Patterns > Virtual Applications**.
2. Under **Virtual Application Patterns**, select the IBM Business Intelligence Pattern with BLU Acceleration pattern type that you uploaded.

Tip: You can use the **Search** field to search for a pattern type. You can also select a pattern type group from the drop-down list to filter the list.

3. After the pattern type loads, click the **New** icon. The Virtual Application Builder opens in a new window.
4. Optionally, you can select a pattern template if you previously created one. See your IBM PureApplication System documentation for information about creating pattern templates.
5. Click **Start Building**. You have created a new virtual application pattern associated with the selected pattern type. The **Virtual Application Builder** opens in a new window where you can add components and policies.
6. On the Virtual Application properties pane, specify the following information:

Name The name of the virtual application pattern.

Description

(Optional) The description of the virtual application pattern.

Type Leave **Application** selected to create a virtual application pattern or select **Template** to create a virtual application pattern that is used as the basis for creating other virtual application patterns.

Lock option for plugin usage

Specify how this virtual application pattern is affected by upgrades to the pattern type or to IBM Foundation Pattern.

Unlock plugins

If the pattern type is upgraded use the latest versions of pattern type plug-ins. If IBM Foundation Pattern is upgraded, use the latest version.

Lock all plugins

Do not change the version of plug-ins or the version of the IBM Foundation Pattern associated with this virtual application pattern when an upgrade occurs.

Lock all plugins except Foundation plugins

If the pattern type is upgraded, do not change the version of the plug-ins associated with this virtual application pattern. If IBM Foundation Pattern is upgraded, use the latest version.

Note: You can view a list of which plug-ins are locked if you select **Lock all plugins** or **Lock all plugins except Foundation plugins**.

Click the **Source** tab in Virtual Application Builder. The application model source is displayed. Search for the element plugins to view the list.

7. Start building the virtual application by dragging the IBM Business Intelligence Pattern with BLU Acceleration component onto the canvas.

Note: If you previously selected a template, the pattern component might be on the canvas already.

8. Click the component, and on the Virtual Application properties pane, provide the configuration information as required. For more information about the configuration settings, see “IBM Business Intelligence Pattern with BLU Acceleration component” on page 9.
9. To add security to the deployment, add a user registry component to the canvas and provide the configuration information. For more information about the configuration settings, see “User registry components” on page 12.
10. Create connections between the pattern components as required.
 - a. To create the connections between the parts, hover over one of the objects until the blue circle turns orange.
 - b. Select the circle with the left mouse button, drag a connection to the second object until the object is highlighted, and release the mouse button.
11. Click **Save**.

Results

The virtual application pattern is created.

What to do next

After you create your virtual application pattern, you can deploy the pattern to create the virtual application instance.

Deploying virtual application patterns

After you create a virtual application, you can provision and deploy it to the cloud. You can deploy a virtual application multiple times and each deployment is a running *virtual application instance* on the cloud infrastructure. Use the workload console to complete this task.

Before you begin

Ensure that the virtual application pattern is configured with the required settings. After a virtual application pattern is deployed, changes to the virtual application pattern do not modify existing virtual application instances. You must stop the deployed virtual application instance before you can change it.

The Deployed LPAR host prefix must be unique for each virtual application deployment.

About this task

When a virtual application is deployed, the system allocates necessary resources, such as virtual machines and block storage on the cloud infrastructure, and deploys, configures, and starts the virtual application components in the cloud.

The time that it takes to deploy a virtual application depends on several factors, such as the size of the virtual application pattern parts and the interdependencies of parts in the application definition, network usage, storage usage, and the provisioning speed of the virtual machine on the cloud infrastructure.

Note: Connectivity issues with the DNS server can cause increased deployment times or failed deployments. The network administrator for the target network should check the routing tables of the DNS server to ensure that it can resolve the network address of the system.

You can add SSH key-based access to your workload virtual machine when you deploy the virtual application. This type of security provides better protection than password-based access.

Procedure

1. Click **Patterns > Virtual applications**.
2. Select the virtual application pattern that you want to deploy and click **Deploy** on the toolbar.
3. Specify the settings for the deployment:
 - a. Type a name for the deployed virtual application instance.
 - b. Select IPv4 or IPv6 in the **Filter by IP type** field.
 - c. Select the **Filter by environment profile type** from the menu.
 - d. Select the **Profile** from the menu.
 - e. Select the **Priority** from the menu.
 - f. Select the **Cloud group** from the menu.
 - g. Select the **IP group** from the menu.
4. Optionally, set up SSH access if you want to log in to the virtual application instance:
 - a. Expand the **Advanced** section to configure the advanced settings.
 - b. Use one of the following options to set the public key:
 - To use an existing SSH public key, open the public key file in a text editor and copy and paste it into the **SSH Key** field.

Important: Do not use **cat**, **less**, or **more** to copy and paste from a command shell. The copy and paste operation adds spaces to the key that prevent you from accessing the virtual machine.

- To generate a key automatically, click **Generate**. Click **Download** to save the private key file to a secure location. The default name is `id_rsa.txt`.

The system does not keep a copy of the private key. If you do not download the private key, you cannot access the virtual machine, unless you generate a new key pair. You can also copy and paste the public key into a text file to save the key. Then, you can reuse the same key pair for another deployment. When you have the private key, make sure that it has the correct permissions (`chmod 0400 id_rsa.txt`). By default, the SSH client does not use a private key file that provides open permission for all users.

5. Click **OK**.

When the virtual application is deployed, the virtual application instance is listed under the **Instances** section of the workload console. To view the virtual instance, click **Instances > Virtual Applications**.

- View the details of the deployed virtual application in the Virtual Application Instances pane.

The status values are listed in the following table:

Table 8. Possible status values for a deployed virtual application

Status	Deployment description	Virtual machine description
LAUNCHING	The virtual application is being deployed.	The virtual machine is being provisioned on the infrastructure cloud.
INSTALLING	Not applicable	The components of the virtual application are being provisioned on the virtual machine.
RUNNING	Resources are being provisioned on the infrastructure cloud.	The components of the virtual application are running on the virtual machine and can be accessed.
TERMINATING	The virtual application instance resources are being deleted.	The virtual machine is being deleted and its resources will be released.
TERMINATED	The virtual application instance resources are deleted. History files are retained.	The virtual machine is deleted and resources were released.
STOPPING	The virtual application instance is stopping.	The virtual machine is being stopped.
STOPPED	The virtual application instance is stopped. The virtual application can be made available again by starting the instance.	The virtual machine is stopped and it can be restarted.
FAILED	The deployment process could not be started due to either the application configuration or a failure that is occurring in the infrastructure cloud.	The virtual machine did not start successfully.
ERROR	An error occurred during deployment. Check the logs and determine the cause of the error before you redeploy the virtual application.	The virtual machine status.
NOT_READY	The virtual application instance is in maintenance mode.	The NOT_READY status does not apply to virtual machines.

You can also view the virtual machine role health status information. For example, a red check mark is displayed when the amount of processor becomes critical on the virtual machine.

Results

The virtual application instance is successfully deployed and started.

- To stop the virtual application instance, select the virtual application from the **Virtual Application Instances** panel, and click **Stop**.
- To start the virtual application instance again, select the virtual application from the **Virtual Application Instances** panel and click **Start**.

- To redeploy a virtual application, select the virtual application from the **Virtual Application Patterns** panel, and click the **Deploy** icon.
- To remove a stopped application, select it from the **Virtual Application Patterns** panel, and click **Delete**.

What to do next

After you deploy the virtual application, you can use the IP address of the virtual machines to access the virtual application endpoints. You can also view and monitor statistics for your deployed virtual machines, and download and view the log files.

Provisioned virtual machines

The deployed IBM Business Intelligence Pattern with BLU Acceleration provisions a minimum of four virtual machines.

You can log in to the virtual machines as `virtuser`, and type `sudo su - cognos` to switch to the `cognos` user. If you are unable to log in to the virtual machine, see “Cognos user cannot log in to virtual machines” on page 102.

<Hostname>-Primary virtual machine

This virtual application instance hosts the IBM DB2 database server.

It runs the following middleware service:

DB2 The endpoint information contains links that you can use to connect to the internal data store through JDBC. For example, the endpoint information might appear as shown in the following example:

AppUser endpoint:

```
jdbc:db2://172.20.190.215:50000/DATA:user=appuser;
password=VDznr4TD1tpHms;
```

AppDBA endpoint:

```
jdbc:db2://172.20.190.215:50000/DATA:user=appdba;
password=9FGP8K814dVyiXJ2;
```

Application users can use the `AppUser` endpoint to create a connection to the database. Database administrators can use the `AppDBA` endpoint to administer or work with databases.

The `appuser` and `appdba` accounts have database level access. If you want to change instance level settings on this virtual machine, you must log in as `virtuser`, and type `sudo su - db2inst1` to switch to the `db2inst1` user with instance level access.

<Hostname>CM1 virtual machine

This virtual application instance hosts the Content Manager server and IBM DB2 content store.

It runs the following middleware services:

CognosBICM1

The endpoint information contains a link that you can use to view information about the Content Manager, including whether it is in an

active or standby state. For example, the endpoint information might appear as shown in the following example:

IBM Cognos BI Content Manager URI: `http://BIwithBLUCM1:9300/p2pd/servlet`

where BIwithBLUCM1 is the virtual host name where Content Manager is hosted.

- DB2** The endpoint information contains a link that you can use to connect to the IBM DB2 content store for IBM Cognos BI through JDBC. The endpoint information might appear as shown in the following example:

Endpoint information

AppUser endpoint:

`jdbc:db2://172.20.190.233:50000/Content:user=appuser;
password=VDznr4TD1tpHms;`

AppDBA endpoint:

`jdbc:db2://172.20.190.233:50000/Content:user=appdba;
password=9FGP8K814dVyiXJ2;`

Application users can use the AppUser endpoint to create a connection to the database. Database administrators can use the AppDBA endpoint to administer or work with databases.

<Hostname>CM2 virtual machine

This virtual application instance hosts the standby Content Manager.

It runs the following middleware services:

CognosBICM2

The endpoint information contains a link that you can use to view information about the Content Manager, including whether it is in active or standby state. For example, the endpoint information might appear as shown in the following example:

IBM Cognos BI Content Manager URI: `http://BIwithBLUCM2:9300/p2pd/servlet`

where BIwithBLUCM2 is the virtual host name where Content Manager is hosted.

<Hostname>COGP virtual machine

This virtual application instance hosts the Cognos application tier components. It runs the following middleware services:

It runs the following middleware services:

CognosBIAS

The endpoint information contains a link to the IBM Cognos BI portal. It also provides information that you can use to configure other products to connect to your virtual application instance. For example, the endpoint information might appear as shown in the following example:

Application URI: `http://BIwithBLU/p2pd`

URI for external applications: `http://BIwithBLU/p2pd/servlet/dispatch`

Local dispatcher URI: `http://BIwithBLUCOGP:9300/p2pd/servlet/dispatch/ext`

To access the BI application Home page or Login page, use the **Application URI** endpoint.

When you configure other applications to connect to your virtual application instance, use the **URI for external applications** and the **Local dispatcher URI** information.

DownloadClient

The endpoint information contains links to the installation files for the IBM Business Intelligence Pattern with BLU Acceleration modeling tools and the IBM Cognos Samples models. For example, the endpoint information might appear as shown in the following example:

IBM Cognos Dynamic Query Analyzer installer:

http://BIwithBLUCOGP:9300/download/dqa_winx64h_10.2.5000.275_ml.tar.gz

IBM Cognos Cube Designer installer:

http://BIwithBLUCOGP:9300/download/dcubemodel_winx64h_10.2.5001.158_ml.tar.gz

IBM Cognos BI Samples model download

<http://BIwithBLUCOGP:9300/download/model.zip>

Documentation

The documentation endpoint provides links to the product documentation on the web.

BI Pattern with BLU Acceleration:

<http://pic.dhe.ibm.com/infocenter/cbiima/v1r0m0/index.jsp>

IBM DB2 v10.5

<http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp>

IBM PureApplication System v1.1

<http://pic.dhe.ibm.com/infocenter/psappsys/v1r1m0/index.jsp>

Viewing virtual machine details

Each virtual application instance consists of a set of virtual machines that represent a physical node in an application server environment. Use the workload console to view the details for each virtual machine.

Before you begin

You must be granted access to the virtual application instance to complete this task.

Procedure

1. Click **Instances > Virtual Applications**.
2. Select a virtual application instance.
3. Expand **Virtual machines** to view a list of the virtual machines in the virtual application instance. The number of virtual machines is dependent on the virtual application that was deployed. You can view the following details for each virtual machine.

Name This field displays the name of the virtual machine.

Public IP

This field displays the public IP address and fully qualified hostname of the virtual machine.

VM Status

This field shows the current status and health of the virtual machine.

Log This link opens the log file for the virtual machine in the Log Viewer.

Monitor

This link takes you to the monitoring page for the virtual machine.

Started on

This column shows when the virtual machine started.

Middleware Status

This shows the current status of the middleware. Click **Endpoint** to view the endpoint for the application. Click **Monitor** to open the monitoring page for the middleware.

Action

Click **View** to display the available actions for a virtual machine. If an action is not available, it is not active.

Stop If a virtual machine is running, then it can be stopped by using the stop action.

Start If a virtual machine is stopped, then it can be started by using this action.

Edit Configure the memory and CPU resources of the virtual machine.

Test connection

Use this action to test the connection between PureApplication System and the deployed virtual machine.

SSH To open an SSH connection to the virtual machine, supply the password or the private key file name. Then click **OK**.

Verifying the deployed BI application

After the virtual application is deployed, an IBM Business Intelligence Pattern with BLU Acceleration virtual application instance is created. After starting the application instance, you can access the Business Intelligence application by using the Virtual Application Instances dashboard in the workload console to determine the application endpoint URL.

About this task

After the virtual application is deployed, the BI instance generates middleware services that provision a minimum of four virtual machines; two virtual machines for the Content Manager Tier components, and one each for the Cognos BI Application Tier component, and the IBM DB2 server component. The virtual machines have endpoints that provide access to different components of your BI application. For a complete list of endpoints available on each of the virtual machines, see Chapter 3, “Configuring the pattern components,” on page 9.

Important: To access the virtual host name, you must register a DNS entry for the virtual host name.

Procedure

1. Click **Instances > Virtual applications** and select the virtual application instance that you deployed.
2. Expand each section of the **Middleware perspective** to view the endpoints that were created by the deployment.
3. Click each of the endpoints to verify the virtual application pattern deployment.

4. If the host name cannot be resolved by the client machine, update the host file on the local machine.

The location of the host file is platform specific.

In Windows:

```
c:\windows\system32\drivers\etc\hosts
```

In Linux:

```
/etc/hosts
```

- a. Expand the **Virtual machine perspective** to view the IP addresses of the virtual machines.
- b. Add each of the IP addresses to the host file on the local client machine.
- c. Click **Instances > Shared Services**.
- d. Select the **ELB Proxy Service** shared service and expand **Virtual machine perspective** to view the IP addresses for the ELB Proxy Service.
- e. Add the IP addresses for the **ELBInstance** virtual machines.
- f. Select the **System Monitoring - Shared** shared service and expand **Virtual machine perspective** to view the IP address for **Hub-TEMS**.
- g. Add the IP address for the **ITM-Hub-TEMS** virtual machine.
- h. Save and close the host file.

What to do next

Depending on your environment, there might be other post installation activities that you might want to perform. For more information, see “Post-deployment tasks” on page 41.

Chapter 6. Deployed virtual applications

After the virtual application is deployed, you can perform post-deployment tasks to ensure that your applications are performing well.

Post-deployment tasks

The following post-deployment administration tasks can help you to ensure that your applications run optimally.

Accessing the deployed BI application

After the pattern is deployed and the IBM Business Intelligence Pattern with BLU Acceleration virtual application instance is created, you can access the IBM Cognos Business Intelligence portal.

Before you begin

To access the virtual host name, register DNS entries for the IP addresses of the virtual machines. For more information, see “Verifying the deployed BI application” on page 38.

About this task

Use the **Virtual Application Instances** dashboard in the workload console to determine the endpoint and access the deployed IBM Cognos Business Intelligence portal.

Procedure

1. Click **Instances > Virtual applications** and select the virtual application instance that you deployed.
2. Click the **Endpoint** of the **CognosBIAS** middleware perspective.
3. Click the **Application URI** link.

Tip: If the host name cannot be resolved, check to ensure that the DNS entry is accurate.

Results

The IBM Cognos Connection portal opens in a browser window.

Identifying the active Content Manager

Content Manager is the IBM Cognos BI service that manages the storage of customer application data, including security, configuration data, models, metrics, report specifications, and report output. Content Manager is needed to publish packages, retrieve or store report specifications, manage scheduling information, and manage the Cognos namespace.

The IBM Business Intelligence Pattern with BLU Acceleration deployment installs Content Manager on two virtual machines, **<hostname>CM1** and **<hostname>CM2**. One

Content Manager is active and the other is on standby. Both the active and standby Content Manager use the same content store database, which is on the `<hostname>CM1` virtual machine.

About this task

Use the **Virtual Application Instance** dashboard in the workload console to manage the virtual application instance and determine which is the active Content Manager.

Procedure

1. Click **Instances > Virtual applications** and select the virtual application instance that you deployed.
2. In the middleware perspective, click the **Cognos BI Content Manager URI** endpoint for both the **CognosBICM1** and **CognosBICM2** roles. A window that identifies the state of each Content Manager as either Running or Running as standby is displayed.

Tip: If the host name cannot be resolved, check to ensure that the DNS entry is accurate.

Results

The state of the virtual machine that is hosting the active Content Manager is identified as Running.

Installing end user modeling tools

IBM Business Intelligence Pattern with BLU Acceleration provides end user modeling tools for working with IBM Cognos Dynamic Cubes.

IBM Cognos Cube Designer is the application that you use to model dimensional metadata and dynamic cubes. You use it to build dynamic cubes and publish them for use in the IBM Cognos studios.

IBM Cognos Dynamic Query Analyzer is a tool that provides graphical representations for the query logs produced by dynamic query mode queries. Cognos Dynamic Query Analyzer has Aggregate Advisor, a tool that analyzes dynamic cubes and suggests aggregates that, when implemented, improve cube performance.

The installation files for the end user modeling tools can be accessed from the **DownloadClient** endpoint, available from the middleware perspective of the deployed BI Pattern with BLU Acceleration instance.

The installation documentation for the modeling tools can be accessed from the **Documentation** endpoint, also available from the middleware perspective of the deployed BI Pattern with BLU Acceleration instance.

Installing translated product documentation

Some of the IBM Business Intelligence Pattern with BLU Acceleration product documentation is translated. To access the translated documentation, you must install it from IBM Cognos BI Supplementary Language Documentation.

Before you begin

Before installing the Supplementary Language Documentation, ensure that:

- BI Pattern with BLU Acceleration is deployed and running
- At least 220 MB of disk space is available to install supplementary language documentation
- Your software environment is supported

About this task

The IBM Supplementary Language Documentation disc contains all translated product documentation for IBM Cognos Business Intelligence. Not all documents on the IBM Supplementary Language Documentation disc are applicable to IBM Business Intelligence Pattern with BLU Acceleration.

Procedure

1. In the location where the Gateway component is installed, insert the IBM Supplementary Language Documentation disc or go to the directory where the installation files were downloaded and extracted.
 - On UNIX or Linux operating systems, mount the disc using Rock Ridge file extensions.
 - On Windows, the installation wizard starts automatically from the product disc.
2. To manually start the installation wizard, go to the operating system directory and do the following:
 - On Windows, if no Welcome page appears, double-click the `issetup.exe` file.
 - On UNIX or Linux, type
`./issetup`

Note: When you use the `issetup` command with XWindows, Japanese characters may be corrupted.

3. Follow the instructions in the installation wizard to copy the required files to the same location where you installed gateway components for IBM Cognos BI. Install in a directory that contains only ASCII characters in the path name. Some Web servers do not support non-ASCII characters in directory names. The supplementary languages documentation component is selected by default.

Configuring IBM Cognos BI products to connect to a BI virtual application

For Business Intelligence (BI) products, such as IBM Cognos Cube Designer and IBM Cognos Dynamic Query Analyzer, to connect to your virtual application instance, you must update their configurations with the location of your virtual application instance.

On each computer where IBM Cognos BI products are installed, you must update the **Gateway URI** and the **Dispatcher URI for external applications** properties using IBM Cognos Configuration.

Before you begin

Obtain the endpoint information for your virtual application instance. To obtain it, you can use the Virtual Application Instances dashboard and click the **Endpoint**

link of any virtual machines that host a BI application server. For example, assuming an application with a virtual host name of `ibmcognos`, the endpoint information might be as follows:

Endpoint information

Application URI: `http://ibmcognos/p2pd`

URI for external applications: `http://ibmcognos/p2pd/servlet/dispatch`

Local dispatcher URI: `http://9.24.207.162:9300/p2pd/servlet/dispatch/ext`

Procedure

1. On the client workstation, start IBM Cognos Configuration.
2. In the **Explorer** window, click **Environment**.
3. In the **Properties** window, update the following properties to point to your virtual application instance:
 - a. Update the **Gateway URI** property:
 - If SSL is enabled on a first-tier entry point, use the following syntax:
`https://virtual host name:443/p2pd/servlet/dispatch/ext`
For example: `https://ibmcognos:443/p2pd/servlet/dispatch/ext`.
 - If SSL is not enabled on a first-tier entry point, use the following syntax:
`http://virtual host name:80/p2pd/servlet/dispatch/ext`
For example: `http://ibmcognos:80/p2pd/servlet/dispatch/ext`.
 - b. Update the **Dispatcher URI for external applications** property with the information provided by the **Local dispatcher URI** endpoint value using the following syntax: `http://IP address of local dispatcher:9300/p2pd/servlet/dispatch`.
For example: `http://9.24.207.162:9300/p2pd/servlet/dispatch`.
4. From the **File** menu, click **Save**.

Importing or exporting BI deployment archives

You can use Business Intelligence (BI) deployment archives in your cloud environment.

First, you must copy the BI deployment archive to the virtual server hosting the active Content Manager. Then, a BI administrator can use IBM Cognos Administration to import or export the deployment archive into the BI application.

Before you begin

- You successfully deployed a BI Pattern with BLU Acceleration virtual application instance.
- You have access to the source environment where the deployment archives are stored.
- You know the IP address of the virtual machine that is hosting the active Content Manager. For information about identifying which virtual server hosts the active Content Manager, see “Identifying the active Content Manager” on page 41.

About this task

You must repeat this procedure whenever you redeploy an instance of the virtual application.

Procedure

1. You can log in to the active Content Manager server as `virtuser`, and type `sudo su - cognos` to switch to the `cognos` user.
If you cannot log in, the RSA key pair might be out of sync. For information about copying the `virtuser` key to the `/home` directory of the `cognos` user, see “Cognos user cannot log in to virtual machines” on page 102.
2. Find the deployment archive files in your source environment and copy them to the `<c10_installation>/deployment` directory in the active Content Manager server. For example, `/opt1/ibm/cognos/c10_64/deployment`.
3. Open IBM Cognos Administration and log in as an administrator.
4. Import or export the deployment archives. For information about working with deployment archives, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

Rendering images in HTML and PDF reports

For images to display in HTML and PDF reports in your deployed application, you must copy the images to the virtual application COGP server to make the images available.

Before you begin

- You successfully deployed a virtual application instance of the IBM Business Intelligence Pattern with BLU Acceleration.
- You have access to the source environment where the images are stored.
- You know the IP address of the virtual machine that is hosting the COGP server. You can use the workload console of your system to determine the IP address of the COGP server.

About this task

You must repeat this procedure whenever you redeploy an instance of the virtual application.

Procedure

1. You can log in to the COGP server as `virtuser`, and type `sudo su - cognos` to switch to the `cognos` user.
If you cannot log in, the RSA key pair might be out of sync. For information about copying the `virtuser` key to the `\home` directory of the `cognos` user, see “Cognos user cannot log in to virtual machines” on page 102.
2. Find the images in your source environment and copy them to the `/data/images` directory in the COGP server.
3. Add references to your images in your BI reports using the following syntax: `http://virtual_host/p2pd/images/image_name`. For example, to add an image called `logo.jpg`, you might specify `http://ibmcognos/p2pd/images/logo.jpg`.
For information about working with images, see “Insert an Image in a Report” in the *IBM Cognos Report Studio User Guide*.

Moving a BI virtual application to a production environment

After you test your BI virtual application, you can move it to a production environment.

Procedure

1. Complete one of the following tasks:
 - Perform a full content store deployment.
 - Perform a database backup of your existing content store and restore the content into a new content store database.

For information about backing up the content store, see “Determine a backup strategy” on page 48.

2. To enable the virtual host name to work for the application users, register a DNS entry for the virtual host name.
3. Update any content with URL links so that those links point to the new endpoint information. Test your changes.
4. For your application to use SSL on the first-tier proxy, make sure that SSL is configured on port 443 of the first-tier proxy.

Maintenance of BI virtual applications

Occasionally, you must apply maintenance activities such as password updates or software updates to virtual applications. These maintenance activities are called operations. Before you apply operations, an application must be in maintenance mode.

From your system, you can submit a maintenance operation for each administration task separately or you can combine multiple tasks into a single maintenance operation. Combining multiple operations reduces the downtime of your virtual applications. After a maintenance operation is applied, the system restarts the virtual application automatically. If the virtual application cannot be restarted, review the log files to determine the cause of the problem.

Updating IBM DB2 with BLU Acceleration settings

After initial load of the atomic facts, dimensions, and aggregate tables is complete, you must change some settings for IBM DB2 with BLU Acceleration before you run reports.

About this task

The following table shows the values of the settings before you perform this task and after you perform this task:

Table 9. IBM DB2 with BLU Acceleration settings for initial data load and incremental data loads

Settings for initial load	Settings for incremental load
40% buffer pool	25% buffer pool
40% SHEAPTHRES_SHR	50% SHEAPTHRES_SHR
$SORTHEAP = SHEAPTHRES_SHR / 5$	$SORTHEAP = SHEAPTHRES_SHR / 20$

Procedure

1. Set the application to **Maintenance** mode.
2. Click **Instances > Virtual Application**.
3. Select the Business Intelligence virtual application that you want to update.
4. Click the **Manage** icon. The Virtual Application Console window opens.

5. Click the **Operations** tab.
6. Click **TuneDB2**.
7. Click to expand the **Tune DB2** option.
8. Click **Submit**.

Viewing virtual application instance log files

You can view virtual application instance log files in the virtual application console.

Before you begin

All of the virtual machines in the deployed virtual application must be started before you view the logs.

Procedure

1. Click the **Workload Console** tab at the top of the Welcome page to open the workload console.
2. Click **Instances > Virtual Applications**.
3. Select a virtual application instance.
4. Click **Manage** on the toolbar.
5. Click the **Logging** tab.

Logs are organized by virtual machine name, and each name contains the name of the *role* associated with the virtual machine. A role within a virtual application refers to the middleware that is installed and configured. For example, a virtual machine named `ibmcognosCM1` contains `CM1`. This virtual machine has Content Manager installed and configured.

Under each virtual machine in the list, logs are organized into sections, such as operating system log, pattern type plug-in log, agent log, and logs for the type of middleware that is installed.

At the top of the list of virtual machines, you can specify a string to filter what files are displayed in the log viewer and then click **Go**. The filter can be prefixed by a tag that specifies one or more elements of the logs to be examined. The following tags are supported:

role: The role name. For example, `role:DB2` filters for role names that contain "DB2".

dir: The absolute path of the directory that contains the log. The filter searches the entire path for the specified file, but does not search within the log file name. For example, `dir:var` filters for directory paths that contain "var".

vm: The virtual machine name. For example, `vm:application` filters for virtual machine names that contain "application".

file: Log file name. For example, `file:trace` filters for file names that contain the word "trace".

The following usage rules apply to filters:

- If a tag is not included in a search string, the `file:` tag is assumed.
- Tag names are case-sensitive, but the filter string after the tag is not case-sensitive. For example, the string `"role:WAS"` matches a role name "WAS" as well as a role named "was", but the string `"ROLE:was"` does not match anything, because `ROLE` is not recognized as a valid tag.

- You can apply multiple filters. For example, `role:DB2,trace` filters for logs with a name that contains "trace", and that belongs to a role with a name that contains "DB2".
6. To search the text of a log, click the log file name in the list and then click the search icon on the toolbar and specify the search criteria. You can use the previous and next icons on the toolbar to search for the next or previous occurrence of the last specified search string.
 7. The following options are available to save logs:

Download all log files

When no filter is currently applied to the list of logs, or when the filter input box is cleared, you can click **Download All** to save an archived, compressed file, which contains all the logs on the virtual machine. If there are multiple virtual machines displayed in the log viewer, a separate archive file is returned for each virtual machine.

Download all files in a filtered list of logs

When you enter a string into the filter box to select a subset of logs, you can click **Download Filtered** to download all of the files that are displayed in the filtered list as a single archive file.

Download an individual log

To save a specific log file, click the log file name and click the **Save** icon on the toolbar.

Determine a backup strategy

We recommend that you determine a backup strategy to prevent the loss of your data in the event that it becomes damaged or stolen. IBM Business Intelligence Pattern with BLU Acceleration does not backup any of your data or configuration settings.

Because backing up consumes system resources, if IBM Cognos software is running while the database is backed up, its performance will be affected.

The following IBM DB2 databases must be backed up. For information about backing up the databases, see the Data Recovery topics in the IBM DB2 information center (pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp).

- IBM Cognos Business Intelligence content store
The content store database stores all application data, including security, configuration data, models, metrics, report specifications, and report output. The content store database, named Content, is stored in the `/home/db2inst1` directory on the `<hostname>CM1` virtual machine. The database owner is `db2inst1`.
- Internal data store
The internal data store is the analytic repository used for reporting and data analysis. The database, named DATA, is stored in the `/home/db2inst1` directory on the `<hostname>-Primary` virtual machine. The database owner is `db2inst1`.

In addition to the databases, ensure that your backup strategy includes the following:

- Report images stored in the `/data/images` directory on the `<hostname>COGP` virtual machine.
- Business Intelligence (BI) deployment archives stored in the `/deployment` directory on the virtual machine hosting the active content manager. For information about how to determine which virtual machine is hosting the active content manager, see "Identifying the active Content Manager" on page 41.

- Directories containing encryption and signing key settings. If the key stores are secured with passwords, ensure that you retain these passwords.

Diagnosing query and reporting problems

To diagnose query and reporting problems for a dynamic cube, you can increase the logging level for some logs of the cube, run reports, and then analyze the logs.

About this task

You can change the logging level for the following types of query logs that are related to dynamic queries and dynamic cubes:

- Aggregate cache (in-memory aggregates)
- In-database aggregate and aggregate awareness (in-database aggregate routing)
- SQL (database queries)
- Aggregate Advisor
- ROLAP query details

Changing the logging level of the previously listed query logs does not require a restart of the cube or the query service.

After you change the logging level for the cube, run the reports, and review the logs. The logs are stored in the *bi_install_location*\logs\XQE\xqe\log directory. The logs can help you to answer the following questions:

- What data was cached?
- Is the member cache up to date?
- Is the query pushed down from the database?
- Are all query joins correctly defined?
- What items undergo data type conversion?

Note: You can also change logging levels by changing query service properties in IBM Cognos Administration. However, if you change query service properties in Cognos Administration, you must restart the query service for your changes to take effect. For more information, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

In order of lowest to highest logging level, the valid logging values are no value, info, and trace.

Procedure

1. In a text editor, open the following file:

```
bi_install_location\configuration\xqe.diagnosticslogging.xml
```

where *bi_install_location* is the location where IBM Cognos Business Intelligence is installed.

2. To change the logging of the aggregate cache, edit the following line:

```
<eventGroup name="ROLAPCubes.AggregateCache" level="logging_value"/>
```

Tip: To enable the lowest logging level, set the value as shown in the following line:

```
<eventGroup name="ROLAPCubes.AggregateCache" level=/>
```

3. To change the logging of aggregate awareness, edit the following line:

```
<eventGroup name="ROLAPCubes.AggregateStrategy" level="logging_value"/>
```

4. To change the logging of SQL, edit the following line:


```
<component>
<component name="XQE">
<eventGroup name="QueryService.SQL" level="logging_value"/>
```
5. To change the logging of the Aggregate Advisor, edit the following line:


```
<eventGroup name="ROLAPAggregateAdvisor" level="logging_value"/>
```
6. To set the logging of ROLAP query and ROLAP query performance, edit the following line:


```
<eventGroup name="ROLAPQuery" level="trace"/>
<eventGroup name="ROLAPQuery.Performance" level="logging_value"/>
```
7. Save the xqe.diagnosticslogging.xml file.

Updating user credentials

You can update user names and passwords for the components used by the virtual application instances. In addition, you can update the credentials used for binding to the LDAP server.

Before you begin

Your application must be deployed, and all of your virtual machines, must be started before you can change the configuration.

Before you submit an operation, put your virtual application instance in maintenance mode until the operation is applied.

Procedure

1. Click **Instances > Virtual Applications**. The Virtual Application Instances palette displays.
2. Select the virtual application that you want to update.
3. Click the **Manage** icon.
4. Click the **Operations** tab.
5. From the **Operations** palette, select the Business Intelligence application (CognosBICM1).
6. Depending on the credentials that you want to update, expand the following sections:
 - For the credentials used to access the content store, the audit database, or the notification mail server, or to perform administrative tasks, go to the **Fundamental** section and expand the **Update configuration** section.
 - For the credentials used for binding to the LDAP server or to update the SSL certificate database, go to the **User registry** section and expand the **Update user registry binding credentials** section.
7. Type the new values.

Important:

- When you update the binding credentials for the user registry, you must specify the name of the existing user registry in the **User registry name** field. This value represents the namespace name.
 - When you update the administrative credentials, you must specify the namespace ID in the **Administrative namespace ID** field.
8. Click **Submit**.
 9. To monitor the progress of the operation, click the **Refresh** icon.

What to do next

After the values are applied to your virtual instance, the instance restarts automatically. To monitor the status of the instance, click **Instances > Virtual Applications**. From the Virtual Application Instances palette, select the instance to view detailed information about the status of the operation.

Important: If the credentials for the content store are not valid, the virtual instance cannot automatically restart.

If your virtual application instance is still in maintenance mode, click the **Resume** icon to exit maintenance mode.

Updating the list of valid domain names and host names

After you deploy an application, you can update the list of domain names or host names used by IBM Cognos Application Firewall (CAF) to validate requests.

Before you begin

Your application must be deployed, and all of your virtual machines must be started before you can change the configuration.

Procedure

1. Click **Instances > Virtual Application**. The Virtual Application Instances palette displays.
2. Select the virtual application that you want to update.
3. Click the **Manage** icon.
4. Click the **Operations** tab.
5. From the **Operations** palette, select the Business Intelligence application (CognosBICM1).
6. In the **Fundamental** section, expand the **Update configuration** section.
7. In the **Valid domains or hosts** field, type domain names and host names separated by a comma. For example: *.mycompany.com, hostname, hostname:80.
8. Click **Submit**.
9. To monitor the progress of the operation, click the **Refresh** icon.

What to do next

After the values are applied to your virtual instance, the instance restarts automatically. To monitor the status of the instance, click **Instances > Virtual Applications**. From the Virtual Application Instances palette, select the instance to view detailed information about the status of the operation.

Important: If the credentials for the content store are not valid, the virtual instance cannot automatically restart.

If your virtual application instance is still in maintenance mode, click the **Resume** icon to exit maintenance mode.

Updating notification mail server information

After you deploy an application, you can update the notification mail server information.

Before you begin

Your application must be deployed, and all of your virtual machines must be started before you can change the configuration.

Procedure

1. Click **Instances > Virtual Application**. The Virtual Application Instances palette displays.
2. Select the virtual application that you want to update.
3. Click the **Manage** icon.
4. Click the **Operations** tab.
5. From the **Operations** palette, select the Business Intelligence application (CognosBICM1).
6. In the **Fundamental** section, expand the **Update configuration** section.
7. Specify the new values for the notification mail server.
8. Click **Submit**.
9. To monitor the progress of the operation, click the **Refresh** icon.

What to do next

After the values are applied to your virtual instance, the instance restarts automatically. To monitor the status of the instance, click **Instances > Virtual Applications**. From the Virtual Application Instances palette, select the instance to view detailed information about the status of the operation.

Important: If the credentials for the content store are not valid, the virtual instance cannot automatically restart.

If your virtual application instance is still in maintenance mode, click the **Resume** icon to exit maintenance mode.

Updating fonts

After you deploy an application, you can add new fonts or refresh existing fonts by uploading a compressed file that contains all the supplementary font files associated with your application. You do not need to upload the fonts that are included by default with IBM Cognos Business Intelligence version.

Before you begin

- Your application must be deployed, and all of your virtual machines must be started before you can change the configuration.
- Prepare a compressed file that contains the TrueType font files. The compressed file must include all the supplementary font files (ttf or ttc format) required by your virtual application instance, even when these font files have not changed.
- The name of the compressed file does not have to match the name of the compressed file that was used in the original deployment.

Procedure

1. Click **Instances > Virtual Application**. The Virtual Application Instances palette displays.
2. Select the Business Intelligence virtual application that you want to update.
3. Click the **Manage** icon.

4. Click the **Operations** tab.
5. From the **Operations** palette, select the Business Intelligence application (CognosBICM1).
6. In the **Fundamental** section, expand the **Update configuration** section.
7. In the **Supplementary fonts** box specify the name of the compressed file that contains all the fonts to associate with your deployment.
8. Click **Submit**.

Updating the list of primary BlackBerry push servers

Cognos Mobile can use the BlackBerry MDS Connection Service capabilities to push reports to BlackBerry device users and to synchronize the reports. After you deploy an application, you can update the list of primary BlackBerry push servers.

Before you begin

- Your application must be deployed, and all of your virtual machines must be started before you can change the configuration.
- IBM Cognos Mobile must be installed.

Procedure

1. Click **Instances > Virtual Application**. The Virtual Application Instances palette displays.
2. Select the virtual application that you want to update.
3. Click the **Manage** icon.
4. Click the **Operations** tab.
5. From the **Operations** palette, select the Business Intelligence application (CognosBICM1).
6. In the **Fundamental** section, expand the **Update configuration** section.
7. Type the new values for the BlackBerry MDS servers.
8. Click **Submit**.
9. To monitor the progress of the operation, click the **Refresh** icon.

What to do next

After the values are applied to your virtual instance, the instance restarts automatically. To monitor the status of the instance, click **Instances > Virtual Applications**. From the Virtual Application Instances palette, select the instance to view detailed information about the status of the operation.

If your virtual application instance is still in maintenance mode, click the **Resume** icon to exit maintenance mode.

Performance monitoring

Monitor the performance of your deployed environment to validate the health and efficiency of your system and the individual workloads running on it.

Monitoring virtual application instances

Use reports and charting to monitor the status and performance of your deployed virtual application instances.

Before you begin

Your applications must be deployed and all of your virtual machines started before you can monitor results.

Procedure

1. Click the **Workload Console** tab at the top of the Welcome page to open the workload console.
2. Click **Instances > Virtual Applications**.
3. Select the virtual application instance. The virtual application instance details display to the right.
4. Click **Manage** on the toolbar.
 - Click **Monitoring > Virtual Machine** and select the virtual machine to display memory, processor, network and storage information.
 - Click **Monitoring > Middleware** and select the virtual machine containing the middleware that you want to monitor. The monitoring information that appears differs depending on what is installed on the virtual machine.

Monitoring virtual machines

Use reports and charting to monitor the status and performance of your deployed virtual machines.

Before you begin

Your applications must be deployed and all of your virtual machines started before you can monitor results.

Procedure

1. Click the **Workload Console** tab at the top of the Welcome page to open the workload console.
2. Click **Instances > Virtual Applications**.
3. Select the virtual application instance. The virtual application instance details display to the right.
4. Expand the virtual machine perspective to view the list of virtual machines.
 - To monitor the virtual machine, click **Monitor** in the VM Status column beside the virtual machine that you want to monitor.
 - To monitor a specific role on a virtual machine, click **Monitor** in the Middleware Status column beside the role that you want to monitor.

Monitoring databases

IBM PureApplication System provides a web console that you can use to isolate and analyze typical database performance problems.

You can use the database performance monitor to view a summary of the health of your databases and drill down for more details. For more information, see the Monitoring database performance with the database performance monitor (pic.dhe.ibm.com/infocenter/psappsyst/v1r1m0/topic/com.ibm.puresystems.appsys.1700.doc/OPM/com.ibm.datatools.perfmgmt.monitor.doc/p_monitor.html) topic in the IBM PureApplication System information center.

The Database Performance Monitoring service must be deployed in your cloud environment before you can use it to view information about the health of your databases. For more information, see “Deploying shared services” on page 19.

Monitoring dynamic cubes

You can use cube metrics that are available on the Status System page in Cognos Administration to monitor the database aggregate table hit rate.

After running reports for a cube, look at the cube metrics and if necessary, run Aggregate Advisor again. For more information about the Status System page in Cognos Administration, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

For techniques and scenarios for dynamic cube performance tuning, see the *Optimization and performance tuning* chapter in the *IBM Cognos Dynamic Cubes RedBook* (www.redbooks.ibm.com/abstracts/sg248064.html).

Chapter 7. Data preparation

Data preparation for IBM Business Intelligence Pattern with BLU Acceleration includes loading data into IBM DB2 with BLU Acceleration columnar-organized tables, optimizing dynamic cube performance, and configuring near real-time analysis.

Related concepts:

“ETL scenarios and best practices”

There are a number of scenarios and best practices when extracting, transforming, and loading (ETL) data for IBM DB2 with BLU Acceleration in columnar-organized tables.

“Near real-time analysis” on page 61

The Sample Outdoors sales model contains dynamic cubes that are configured to do near real-time analysis. Use the sample dynamic cubes as a guide to configure your model to do near real-time analysis.

Related information:

“Dynamic cube performance” on page 60

To optimize cube performance, your administrator can use Cognos Administration to monitor the metrics of the dynamic cubes and change the cube configuration. To further optimize performance, your system analyst can run reports that are a representative workload against the dynamic cube. Aggregate Advisor uses the resulting workload logs to generate recommendations for in-database aggregates.

ETL scenarios and best practices

There are a number of scenarios and best practices when extracting, transforming, and loading (ETL) data for IBM DB2 with BLU Acceleration in columnar-organized tables.

Creating and loading data into columnar tables

Loading data into column-organized tables is very similar to loading data into row-organized tables, with some exceptions.

For an introduction to column-organized tables, restrictions, configuration requirements, steps to create column-organized tables, examples, and scenarios, see the "Column-organized tables" topic in the IBM DB2 10.5 information center (pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.dboobj.doc/doc/c0060592.html).

For details on converting row-organized tables into column-organized tables, see the "db2convert - Convert row-organized tables into column-organized tables" topic in the IBM DB2 10.5 information center (pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp?topic=/com.ibm.db2.luw.admin.cmd.doc/doc/r0060728.html).

Loading atomic facts and dimensions

When loading atomic facts and dimensions into IBM DB2 with BLU Acceleration, there are three possibilities.

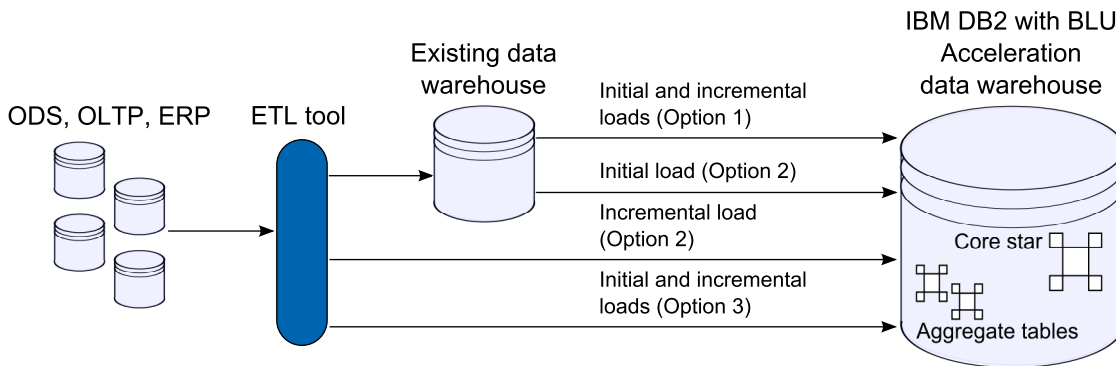


Figure 2. ETL options

Option 1

Continue to maintain an existing data warehouse, and push the data to the IBM DB2 with BLU Acceleration data warehouse for initial load and incremental loads. If you are using this option for incremental load, a change data capture field must exist in the atomic fact and dimension tables. If the change data capture field does not exist, then use option 2.

Option 2

Migrate initial data from an existing data warehouse to the IBM DB2 with BLU Acceleration data warehouse, and then push the incremental loads from the source to the IBM DB2 with BLU Acceleration data warehouse using the ETL tool. If you are using this option for incremental load, a change data capture field must exist in the source tables.

Option 3

Choose this option if you want a clean start, if you want new source columns added (updates will be costly), or to make sure that the data is in sync with the source. The time and effort to repeat initial load can be costly. If the initial load ETL from source to the IBM DB2 with BLU Acceleration data warehouse is large, then options 1 or 2 are preferable. If you have an ETL job to load your existing data warehouse, you can use it to load the IBM DB2 with BLU Acceleration data warehouse. If you are doing a new implementation in which you do not have an existing data warehouse, use this option.

If the schema is the same in an existing warehouse and the IBM DB2 with BLU Acceleration data warehouse, then ETL can be done with direct mapping. No transformation is required.

Create batch jobs to run initial and incremental load. You must create an execution history table that records the last successful run date for each of the tables. Incremental loading of new or updated records from the source or existing data warehouse are added to the target database by comparing the change data capture field in the source tables to the last successful run date.

After the initial data load of atomic fact tables, dimensions, and aggregate tables is complete, perform “Updating IBM DB2 with BLU Acceleration settings” on page 46.

Use stored procedures to load aggregate tables

You can use the stored procedures that are created by Aggregate Advisor in your own extract, transform, and load (ETL) tool to load data into in-database aggregate tables.

Aggregate Advisor generates in-memory and in-database recommendations. When you apply the in-database recommendations, Aggregate Advisor creates stored procedures to load aggregate tables. Aggregate Advisor can schedule an event to run the stored procedures. For more information about Aggregate Advisor and applying in-database recommendations, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

If you want to use your own ETL tool to load your aggregate tables, you can use the stored procedures that are created by Aggregate Advisor in the following ways:

Use stored procedures in your ETL tool

Connect the stored procedures to your ETL tool.

Build your own ETL job

Use the stored procedures as a source of information to build your own ETL job.

Stored procedures

A stored procedure does both initial load and incremental load for an aggregate table. If the aggregate table is empty, an initial load is done. If there is data in the aggregate table, an incremental load is done.

A stored procedure uses a change data capture field from the atomic fact to determine the new and changed transactions for an incremental load. For example, `CHANGED_DATE`. You must set this change data capture field in the Aggregate Advisor component of IBM Cognos Dynamic Query Analyzer in the following situations:

- The first time that you apply in-database recommendations.
- A new atomic cube has been added to the model.

For information about how to set the change data capture field, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

A stored procedure writes the execution status of the stored procedure to a metadata table.

The stored procedures are stored in the IBM DB2 with BLU Acceleration database. An atomic cube can have one or more associated aggregate cubes. Each aggregate cube is mapped to one aggregate table. The name of the aggregate cube is the same as the name of the aggregate table. There is one stored procedure for each aggregate table. A stored procedure is named *aggregate_cube_name_SP*. For example, to find the code of the stored procedure for aggregate cube `AGG_1785921059_SP`, run the following query:

```
SELECT TEXT FROM SYSCAT.ROUTINES WHERE ROUTINENAME='AGG_1785921059_SP'
```

Updating IBM DB2 with BLU Acceleration settings

After initial load of the atomic facts, dimensions, and aggregate tables is complete, you must change some settings for IBM DB2 with BLU Acceleration before you run reports.

About this task

The following table shows the values of the settings before you perform this task and after you perform this task:

Table 10. IBM DB2 with BLU Acceleration settings for initial data load and incremental data loads

Settings for initial load	Settings for incremental load
40% buffer pool	25% buffer pool
40% SHEAPTHRES_SHR	50% SHEAPTHRES_SHR
$\text{SORTHEAP} = \text{SHEAPTHRES_SHR} / 5$	$\text{SORTHEAP} = \text{SHEAPTHRES_SHR} / 20$

Procedure

1. Set the application to **Maintenance** mode.
2. Click **Instances > Virtual Application**.
3. Select the Business Intelligence virtual application that you want to update.
4. Click the **Manage** icon. The Virtual Application Console window opens.
5. Click the **Operations** tab.
6. Click **TuneDB2**.
7. Click to expand the **Tune DB2** option.
8. Click **Submit**.

Dynamic cube performance

To optimize cube performance, your administrator can use Cognos Administration to monitor the metrics of the dynamic cubes and change the cube configuration. To further optimize performance, your system analyst can run reports that are a representative workload against the dynamic cube. Aggregate Advisor uses the resulting workload logs to generate recommendations for in-database aggregates.

For information about optimizing a dynamic cube and changing the cube configuration, see the *IBM Cognos Dynamic Cubes User Guide*.

Aggregate Advisor can apply the in-database recommendations after they are generated. For more information about applying the in-database recommendations, see the *IBM Business Intelligence Pattern with BLU Acceleration Dynamic Query Analyzer User Guide*.

Enabling the result set cache

Enabling the result set cache for a dynamic cube can improve dynamic cube performance and as a result, report performance.

About this task

The result set cache contains the result of MDX queries that are run against a dynamic cube. It is important to allocate the correct amount of disk space for the result set cache because it resides on the file system.

The size of the result set cache depends on the number of named users on your system. Use the following table to determine the size of the result set cache to use.

Table 11. Result set cache size to use

Number of named users	Result set cache size
less than or equal to 100	1 GB
101 to 1000	10 GB
1001 to 5000	50 GB
greater than 5000	100 GB

You must set the result set cache for each cube in IBM Cognos Cube Designer.

Procedure

1. In IBM Cognos Administration, on the **Status** tab, click **Data Stores**.
2. In the **Scorecard** section, select the **Data Stores - (All)** view.
3. Click the dynamic cube that you want to modify, click the actions drop-down menu, and click **Set properties**.
4. Clear the **Disable result set cache** property.
5. Set the **Maximum amount of disk space to use for result set cache** to the applicable value for your system.
6. Restart the dynamic cube.

Near real-time analysis

The Sample Outdoors sales model contains dynamic cubes that are configured to do near real-time analysis. Use the sample dynamic cubes as a guide to configure your model to do near real-time analysis.

You must perform the following tasks for each fact in your model for which you want to do near real-time analysis:

1. Create a recent fact table to store recent transactions. You create this fact table based on an existing fact that contains all historical data except for the recent transactions.
2. Update your model to merge historical and recent transactions. You do this by creating a virtual cube in your model. The virtual cube merges historical fact data with the recent fact data into one cube.
3. Create a task to refresh the data cache of the recent cube whenever the recent fact table is updated. The virtual cube is refreshed at the same time.

Create a recent fact table to store recent transactions

The recent fact table uses the database logs to do change data capture (CDC) to record recent transaction changes.

Use the following guidelines to create a recent fact table in your data warehouse and update it with recent transactions:

- To add records to your recent fact table, use a log-based change data capture (CDC) process that is applicable for your source system. The CDC process you choose must track changes on the source system as they occur and apply them to the recent fact table. Set the CDC process to run as often as required to meet your business needs.
- You must have an extract, transform, and load (ETL) system for loading data to your data warehouse from your source system. Before you load them to your

data warehouse, you can write recent transactions to a flat file, message queue, or staging tables. Use your ETL system to load recent transactions to the recent fact table.

- If the schema on your source system is identical to the schema for your recent fact table, then you can copy the record layout from the source system and use your ETL system to write recent transactions directly to the recent fact table.
- If the data for your recent fact table comes from multiple source tables, then you must capture changes from each of those tables. Create fact rows by joining those tables and insert or update them in the recent fact table using your ETL system. If a table used in the join does not have a change record for the original record, then you must create a database jointo the historical fact to populate the missing columns.
- Configure an ETL job to run when transaction changes occur.
- Ensure that your recent fact table is truncated at the appropriate time.

The following transactions are examples of the types of transactions that your CDC methodology must process. The transactions represent sales of products and are all entered today.

Table 12. Examples of product sale transactions

Trans date	Product key	Customer key	Sales amount	Description
June 12, 2013	1	1000	\$100	A new sales transaction of \$100 for today entered in the OLTP system. Insert a record with a sales amount of \$100 into the recent fact table.
June 12, 2013	1	1000	\$50	The sales amount for a transaction that was created today (the record in the previous row) has changed to \$50. Update the sales amount of the record in the recent fact table to \$50.
June 12, 2013	1	1000	\$50	A transaction that was created today has been deleted from the OLTP system. Delete the record from the recent fact table.
May 10, 2013	5	20051	-\$75	The original sales amount for a historical transaction is \$300. The sales amount has changed to \$225 in the OLTP system. Insert a record into the recent fact table so that the net amount of the historical transaction plus today's transaction equals \$225. That is, $300 + -75 = 225$.
May 10, 2013	6	20034	-\$200	The sales amount for a historical transaction has been deleted from the OLTP system. The amount of the original sale is \$200. Insert a record into the recent fact table so that the net amount of historical plus today's transaction equals zero. That is, $200 + -200 = 0$.

Updating a model to merge historical and today's transactions


Use IBM Cognos Cube Designer to create a virtual cube in your model that merges historical and today's transactions.

About this task

This task refers to the `gosldw_sales` model and `RECENT_SALES` fact table that are in the Great Outdoors sales data sample. The sample files are used to explain how to update a model to merge historical and today's transactions. The steps in the procedure refer to the sample files to do the following things:

- Create a cube called `recent_sales` to model today's sales transactions. The measures in `recent_sales` will be based on a fact table called `RECENT_SALES`.
- Create and publish a virtual cube that merges historical sales data and today's sales transactions.

Procedure

1. Open Cognos Cube Designer.
2. In the **Project Explorer**, click the arrow beside Model, right-click the `gosldw_sales` model, and select **New > Cube**. Name the cube `recent_sales`.
3. Select all measures from `RECENT_SALES` and drag them to the `recent_sales` cube.
4. Select the dimensions Time, Promotions, Order Method, Products, and Retailers and drag them to the `recent_sales` cube.
5. Double-click the `recent_sales` cube.
6. Define the relationship between each dimension and the cube by clicking **Edit** after a dimension and specifying the join to a measure.
7. Right-click the `gosldw_sales` model, and select **New > Virtual Cube**. Name the virtual cube `all_sales`.
8. Select `gosldw_sales` and `recent_sales` as source cubes and click **OK**.
9. In the **Project Explorer**, double-click the `all_sales` virtual cube.
10. Verify that all dimensions are mapped correctly.
11. Verify that all measures are mapped correctly by clicking the Measures icon .
12. In the **Project Explorer**, right-click the `recent_sales` cube and select **Publish**.
13. In the **Publish** window, select the **Select all options** check box and click **OK**.
14. In the **Project Explorer**, right-click the `all_sales` cube and select **Publish**.
15. In the **Publish** window, select the **Select all options** check box and click **OK**.


Refreshing the data cache in the recent transaction cube

In IBM Cognos Administration, you can schedule a task that refreshes the data cache in the `recent_sales` cube at regular intervals.

About this task

When the data cache in the `recent_sales` cube is refreshed, the `all_sales` virtual cube is also refreshed.

Procedure

1. In IBM Cognos Connection, click **Launch > IBM Cognos Administration**.
2. On the **Configuration** tab, click **Content Administration**.
3. Click the **New Query service administration task** icon , and select **Dynamic cube**.
4. Type a name for the task and click **Next**.
5. In the **Operation** field, select **Refresh data cache**.

6. Under **Cubes**, select the **recent_sales** cube and then click **Next**.
7. Under **Action** select **Save and schedule**, and then click **Finish**.

Chapter 8. IBM Cognos Business Intelligence samples

The IBM Cognos Business Intelligence sample databases contain sales, marketing, and financial information for a fictional company named the Sample Outdoors Company that sells sporting equipment. The samples illustrate product features and technical and business best practices, and can be used for experimenting with and sharing report design techniques, and for troubleshooting.

The samples for IBM Business Intelligence Pattern with BLU Acceleration are partially deployed and configured when the pattern is deployed if you have selected the **Deploy IBM Cognos BI Samples** checkbox. You can find the samples in the **Public Folders** tab in IBM Cognos Connection, in a folder named **Samples_DQ**, **Samples_Dynamic_Cubes**, and **Samples_Mobile**. To make the samples fully functional, you must republish the sample dynamic cubes. For more information, see “Deploying sample dynamic cubes” on page 66.

The models and reports in the sample are optimized for dynamic query mode. For example, reports were updated to apply a specific sorting order and to specify an aggregation mode.

You can use the GO Sales sample to understand how to configure near real-time analysis for your model. For more information, see “Near real-time analysis” on page 61.

Importing the IBM Cognos Dynamic Cubes sample deployments

If you did not select the **Deploy IBM Cognos BI Samples** checkbox before deploying the pattern, follow this procedure.


About this task

The deployment archive (`model.zip`) that contains the dynamic cube sample models is available from the DownloadClient endpoint, labelled **IBM Cognos BI Samples model download**. This archive contains the package that is required for the IBM Cognos Dynamic Cubes samples.

Procedure

1. Download the zip file from the endpoint to the directory where your deployment archives are saved.
The default location is `c10_location/deployment`.
2. Connect to the IBM Cognos BI portal.
3. Open IBM Cognos Administration. On the **Welcome** page, click **Administer IBM Cognos content**.
4. On the **Configuration** tab, click **Content Administration**.

Important: To access this area in IBM Cognos Administration, you must have the required permissions for the **Administration tasks** secured feature.

5. On the toolbar, click the **New Import** button .
6. In the **New Import** wizard, in the **Deployment Archive** box, select the **IBM_Cognos_DynamicCube** archive, and click **Next**.

7. Type a unique name and an optional description, including a screen tip for the deployment archive, select the folder where you want to save it, and click **Next**.
8. In the **Public Folders Content** box, select the packages and folders that you want to import.
9. Select the options that you want and your conflict resolution choice for options that you select, and click **Next**.
10. In the **Specify the general options** page, select whether to include access permissions and references to external namespaces, specify who owns the entries after they are imported, and click **Next**.
11. Review the summary information, and click **Next**.
12. Choose how you want to import the sample deployment:
 - To run once now or later, click **Save and run once**.
Click **Finish**, specify the time and date for the run, then click **Run**. Review the run time and click **OK**.
 - To schedule at a recurring time, click **Save and schedule**.
Click **Finish**, and then select frequency and start and end dates. Click **OK**.

Tip: To temporarily disable the schedule, select the **Disable the schedule** check box.
 - To save without scheduling or running, click **Save only** and click **Finish**.
13. When the import is submitted, click **Finish**.

Dynamic cube samples

You can use the sample data to learn how to design and model dynamic cubes and use the data in reporting environments. You can safely use the samples instead of accessing corporate data.

Sample Models

The IBM Cognos Dynamic Cubes samples refer to the GOSALESDW database that contains sample data that other IBM Cognos products use. In IBM DB2 with BLU Acceleration, the dynamic cube uses a single schema from the database.

The `gosldw_sales` cube contains sales data with measures such as Quantity, Unit cost, and Revenue. It is based on the `SLS_SALES_FACT` table. Dimensions include Time, Products, Order Method, and others.

The `gosldw_target` cube contains sales target data, and has only one measure, Sales target. It is based on the `SLS_SALES_TARG_FACT` table. The dimensions are Time to month and Product type.

The `gosldw_sales_and_target` virtual cube includes measures from both of the other cubes in one cube. It has the three measures Quantity, Revenue, and Sales target. The dimensions are Products and Time.

Deploying sample dynamic cubes


To work with sample dynamic cubes in the IBM Cognos studios, use IBM Cognos Cube Designer to republish the sample dynamic cubes to deploy and make them available as data sources. The sample dynamic cubes are `gosldw_sales`, `gosldw_target`, and `gosldw_sales_and_target`.

Before you begin

The *great_outdoors_warehouse* data source connection must be set up before you deploy sample dynamic cubes.

If anonymous access is disabled, you must use the credentials that are associated with your account to publish cubes. Go to the **Personal** tab in the **Set preferences** window of the IBM Cognos Portal, and create your credentials before you proceed.

Procedure

1. From the **Start** menu, click **Programs > IBM Cognos 10 > IBM Cognos Cube Designer**.
2. Open the project that contains the dynamic cube that you want to deploy and publish.
 - a. From the toolbar, click **Open** .
 - b. Navigate to the folder that matches the data source that you are using. For example, DB2, and select `model.fmd`.
The default location for the sample files is `c10_location\webcontent\samples\models\great_outdoors_dynamiccube\datasource_type`
3. In the **Project Explorer** tree, expand the project and model.

Note: You may be prompted to log in to an IBM Cognos Business Intelligence Server.

4. Right-click the `gosldw_sales` dynamic cube, and select **Publish**.
5. To deploy the dynamic cube and configure the cube as a data source, in the Publish window, expand **Additional Options** and select the **Add the dynamic cube to the default dispatcher** check box.
6. To start the dynamic cube, select the **Start the dynamic cube** check box.
7. If anonymous access is disabled, make sure that the **Associate my account and signon with the cube datasource** check box is selected. For anonymous access, clear the check box.
8. Repeat steps 4 - 7 for the cube `gosldw_target`, and lastly for the virtual cube `gosldw_sales_and_target`.
9. Click **OK**.

Results

If the deployment and publish process is successful, no errors are reported. A message confirms that the cubes started successfully. You can now use the sample packages to create reports that rely on dynamic cube data sources. You can also run the sample reports that are available on the **Public Folders** tab in the portal.

Dynamic query mode sample reports

Sample models and reports that are optimized for the dynamic query mode are included with IBM Cognos Business Intelligence.

When installed and deployed, you can find the updated samples in the **Public Folders** tab in IBM Cognos Connection, in a folder named **Samples_DQ**. The updated reports were also renamed with the suffix `_DQ`.

The samples were modified slightly to benefit from the key improvements of the dynamic query mode. For example, reports were updated to apply a specific sorting order and to specify an aggregation mode.

To access the dynamic query mode samples, you must modify the data source connections to two sample data sources to enable JDBC connections and then import the updated samples deployment archive.

Cognos Mobile samples

The IBM Cognos Business Intelligence server installation includes a set of sample active reports for the iPad and a set of sample static reports optimized for mobile devices.

The iPad and static sample reports illustrate product features and technical and business best practices. Users can use them for experimenting with and sharing report design techniques and for troubleshooting.

iPad users can try out the interactive functionality of active reports. These reports let users compare different areas of their business to determine trends, for example, over time, by region, by departments or in combination, or compare business methods and statistics.

The iPad sample active reports include the following:

- **Employee Recruitment**
This active report compares the effectiveness of various employee recruitment methods for each department and country or region. It shows the organization names, positions filled, planned positions, and a bulleted chart of positions filled versus planned positions.
- **Customer Satisfaction**
This active report compares the number of returns by customers by order method and region. The report provides additional information about the order method with the highest number of returns. It also shows customer survey results for different regions.
- **Inventory turnover report**
This active report shows information about the regional product inventory turnover, based on two years of comparative data. The report provides key inventory metrics that a company might use to manage its inventory. You can drill down on each product category to view the detailed inventory information and the number of failed orders related to the inventory.
- **Financial report**
This active report shows current performance and changes in the financial position of an enterprise. This type of information is useful to all users who are involved in making business decisions. However, the Finance department is most likely to benefit from this information when implementing the checks and controls in the system to comply with legal, tax, and accounting regulations and requirements, and when providing advice about future directions, performance, and opportunities for the business.

Cognos Mobile sample active reports demonstrate the following product features:

- Interactive behavior between controls.
- Access to Details on Demand by leveraging drill-down functionality.
- Conditional palette and drill-down to details from a chart.

- Specific design for iPad gestures, such as swiping and scrolling.
- Particular user interface design, such as cover page and color palette.
- Different type of active report items, such as Deck, Tab Control, Chart, Buttons, Drop-down list, Iterator and Slider

The static sample reports let users try the following features:

- Drilling through from one IBM Cognos Mobile report to another.
- Using tree prompts to filter reports.
- Viewing different report types, such as charts, crosstab, and list reports.
- Using calculated values and showing the results as a percentage.

The Mobile sample static reports include:

- Running Total For Promotion
This report shows planned revenue by percentage for each product line. The report uses charts, crosstab reports, and list reports.
- Sales Revenue Expectation
This reports shows the goal for the percentage change in sales revenue for retailers. It uses a calculated value for forecasted revenue.
- Tree Prompt Retailers Set
This report shows the revenue for a set of retailers. The report is a drill-through target for the Sales Revenue Expectation report.

The Sample Outdoors Company

The Sample Outdoors Company samples illustrate product features and technical and business best practices.

You can also use them for experimenting with and sharing report design techniques and for troubleshooting. As you use the samples, you can connect to features in the product.

The Sample Outdoors Company, or GO Sales, or any variation of the Sample Outdoors name, is the name of a fictitious business operation whose sample data is used to develop sample applications for IBM and IBM customers. Its fictitious records include sample data for sales transactions, product distribution, finance, and human resources. Any resemblance to actual names, addresses, contact numbers, or transaction values, is coincidental. Unauthorized duplication is prohibited.

Samples outline

The samples consist of the following:

- Two databases that contain all corporate data, and the related sample models for query and analysis
- Sample cubes and the related models
- A metrics data source including associated metrics and a strategy map for the consolidated company, and a model for Metric extracts.
- Reports, queries, query templates, and workspaces

To run interactive reports, scripts are required. To see all the reports included in the samples packages, copy the files from the samples content installation into deployment folder and then import the deployments into the IBM Cognos Business Intelligence product.

Security

Samples are available to all users.

To implement security, see the IBM Cognos Business Intelligence *Administration and Security Guide*.

The Sample Outdoors Group of Companies

To make designing examples faster, especially financial examples, some general information about The Sample Outdoors Company is useful.

To look for samples that use particular product features, see the individual sample descriptions in this section.

Revenue for The Sample Outdoors Company comes from corporate stores and from franchise operations. The revenues are consolidated from the wholly-owned subsidiaries. There are six distinct organizations, each with its own departments and sales branches. Five of these are regionally-based companies.

The sixth company, GO Accessories:

- Has its own collection of products, differentiated from the other GO companies by brand, name, price, color and size.
- Sells from a single branch to all regions and retailers.
- Functions both as an operating company based in Geneva, and as a part owner of the three GO subsidiaries in Europe.

The diagram illustrates the consolidated corporate structure, including the percentage changes in ownership for GO Central Europe, and shows the reporting currency and GL prefix for each subsidiary.




Figure 3. Consolidated corporate structure

Each corporation has the same departmental structure and the same GL structure, shown in the table. Divisions may not report in the same currencies. For example, the Americas subsidiary reports in US dollars, but the Corporate division local currency is Canadian dollars, and the Operations division local currency is pesos.

Table 13. Departmental structure

Division (GL)	Department (GL)
Corporate (1700)	Sales (1720)
	Marketing (1750)
	IS&T (1760)
	Human Resources (1730)
	Finance (1740)
	Procurement (1710)
Operations (1800)	Production and Distribution (1820)
	Customer Service (1820)

Each corporation has a complete chart of accounts. Most of the accounts, such as those under non-personnel expenses, are at the department level, and contain only summary amounts. For example, although each marketing department has expenses, the cost is unspecified at the transaction level where marketing promotions occur.

Employees

The Sample Outdoors data contains a full list of employees in all divisions, departments, and locations.

Data is available for reports about bonuses (Global Bonus report) and sales commissions (Sales Commissions for Central Europe report), training (Employee Training by Year report), and performance reviews and employee satisfaction surveys (Employee Satisfaction 2012). If you use Metric Studio, sample metrics for human resources are also available.

In the GO Data Warehouse (analysis) package, groups of measures and the related dimensions are organized into folders. The employees are organized in hierarchies for region and manager, to make different kinds of aggregation easy to report on. Aggregation has been defined for the Employee Position Summary measures, so that Position count and Planned position count aggregate correctly at each level of time: monthly, quarterly, or yearly. For example, see the Planned Headcount report.

The employees are also listed in a sample LDIF file which could be used for any LDAP IBM product authentication including Tivoli. This authentication directory is necessary for IBM Cognos Planning samples. No other samples depend on security profiles.

Sales and marketing

Data about sales and marketing is available for all of the companies in the Sample Outdoors group.

GO Accessories has richer details to support analysis examples. For example, see the Revenue vs % Profit Margin by Product Brand analysis, based on the Sales and Marketing cube. Marketing and sales campaigns are tied to the Sample Outdoors regional companies.

Overall, the GO companies have experienced solid growth across most product lines (Sales Growth Year Over Year), in all regions (Revenue by GO Subsidiary 2011), because of factors like an increase in repeat business and new or improved products, such as the high margin sunglasses product line. In the product lines sold by the five regional companies (all but GO Accessories) promotions have had mixed success (Promotion Success by Campaign, Bundle and Quarter). If you use Metric Studio, this can also be seen in the sample metrics.

Customer surveys

The data also contains information from customer surveys. For example, the product line that includes bug spray, sun screen, and so on has not been successful (Product Satisfaction - Outdoor Protection 2011) and a source of retailer dissatisfaction may be the level of customer service rather than the returns (Customer Returns and Satisfaction). If you use Metric Studio, this information can also be monitored in metrics.

Sales outlets

Revenue from the corporate outlets is available at the transaction level. Revenue from the franchise outlets is available at the consolidated level only (Sales and Marketing cube). Metrics about retailers show that the number of new retail outlets has dropped over the time period covered by this data.

GO Accessories sells worldwide, and sells only accessories. Transaction data for GO Accessories is the primary source for analysis of product by brand, color and size. The other five subsidiaries in the group of companies are regional and sell all product lines for retailers in their region. For example, the report Top 10 Retailers in 2011 uses sparklines and list data to review revenues at the retailer level.

Sample Outdoors database, models, and packages

The Sample Outdoors models illustrate modeling techniques and support the samples.

The models are based on the GO data warehouse and the GO sales transactional database and are the basis for the sample reports and queries. Each model contains two packages for publishing analysis (dimensional) and query views of the data.

You must have access to Cube Designer to look at the sample models.

GO Data Warehouse

The GO Data Warehouse model is based on the database GOSALESDW. It contains data about human resources, sales and marketing, and finance, grouped into business areas. In the Database view, the three business areas are grouped into separate namespaces. The Database view contains a fourth namespace (GO Data) for the common information.

The Database view is very similar to the structure of the underlying database. All tables (database query subjects) are unchanged. This enables IBM Cognos BI to retrieve metadata directly from the package in most cases, instead of using a metadata call to the database. The following changes and additions have been made in the Database view:

- Joins have been added as necessary.
- To allow for aggregation at different levels of granularity, some model query subjects have been created. For example, see the relationships between Time and Sales or Sales fact.
- To allow single joins to be made between the lookup tables and each level in a dimension, lookup tables have been copied. For example, see the Products look up tables.

The Business view contains only model query subjects, with no joins. The following changes and additions have been made in the Business view:

- Calculations were added to the model query subjects. For example, the time dimension contains language calculations.
- Where the database has multiple hierarchies, new dimensions have been created to organize each hierarchy. For example, the employee hierarchies are organized into several categories, such as manager and region.

The GO Sales transactional database

The GO Sales model, is based on the GOSALES database, which is structured as a transactional database. It contains principally sales data.

The Database view is very similar to the underlying database structure. The following changes and additions have been made in the Database view:

- To make it possible to join the fact tables to the time dimension, model query subjects and multipart joins have been used.
- Other joins have been added as necessary.

The Business view contains only model query subjects, with no joins. The following changes and additions have been made in the Business view:

- Calculations were added to the model query subjects.
- Model query subjects that were created in the Database view to enable joins on the time dimension have been linked as reference shortcuts.
- Where the database has multiple hierarchies, new dimensions have been created to organize each hierarchy.
- Sales Staff is a subset of the slowly changing Employee dimension. There is no unique Employee key in GO Sales, so a filter retrieves the current record only. This model does not use historical data.

The samples packages

The Sample Outdoors samples include several packages. A brief description of each available package is provided.

Go Data Warehouse (analysis) is a dimensionally modeled view of the GOSALESDW database. This package can be used in all studios, including IBM Cognos Analysis Studio. Using this package you can drill up and down.

Sales and Marketing (cube) is an OLAP package, based on the sales_and_marketing.mdc cube.

Great Outdoor Sales (cube) is an OLAP package, based on the great_outdoors_sales_en.mdc cube.

Note: The OLAP packages, Great Outdoor Sales (cube) and Sales and Marketing (cube), are not multilingual. The IBM_Cognos_PowerCube.zip archive contains five versions of each package; one in English, French, German, Japanese and Chinese.

Sample Outdoors organization and schemas

Under the holding company GO Consolidated, there are six distinct sales organizations, each with their own departments and sales branches.

GO Consolidated consists of sales organizations for GO Americas, GO Asia Pacific, and GO Accessories. The GO Accessories sales organization consists of GO Northern Europe, GO Central Europe, and GO Southern Europe.

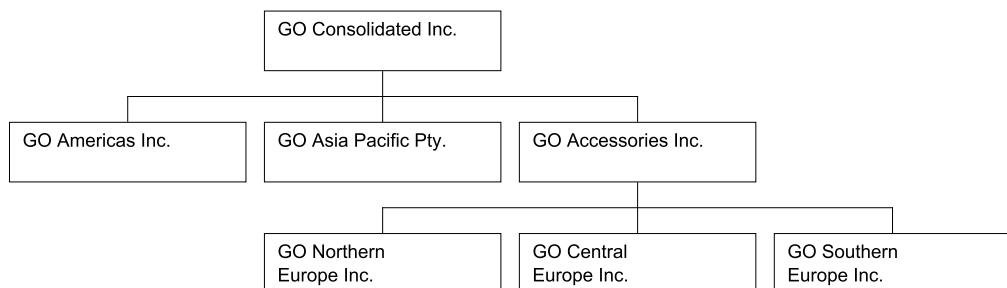


Figure 4. GO Organization

The Outdoors Companies

Five of the six companies are regionally-based companies, whose sales region is identified in the company name. These five companies sell the complete line of outdoor products including:

- Camping Equipment
- Golf Equipment
- Mountaineering Equipment
- Personal Accessories
- Outdoor Protection

The Accessory Company

GO Accessories Inc. operates differently than the five outdoor companies and stands apart in several other areas:

- GO Accessories has its own collection of accessories, differentiated from the outdoor companies by brand, name, price, color and size.
- GO Accessories sells from a single branch to all regions and retailers.
- GO Accessories functions both as an operating company based in Geneva, and as owner or part owner of three GO subsidiaries in Europe. This offers a complex structure that is useful in samples that show office-of-finance type reporting.

Unbalanced hierarchy

The ownership by GO Accessories of the European companies supports applications designed for office-of-finance reporting. Sample applications can include inter-company eliminations and consolidation up to the highest level node - GO Consolidated. In this structure, aggregations of the European organizations occur through GO Accessories.

The resulting unbalanced hierarchy has one additional level, as shown in the following table. At level 3 of the organization structure, GO Central Europe (GOCEU) is at the same level as the operations department for GO Americas (GOAMXOP).

Table 14. Sample GO Accessories Hierarchy

Level	Example (ID)	Example (data)
Org level 1	GOCON	GO CONSOLIDATED
Org level 2	GOAMX GOEUX	GO Americas GO Accessories
Org level 3	GOAMXOP GOCEU	GO Americas operations GO Central Europe
Org level 4	30 GOCEUOP	Sales branch 30 - Sao Paulo - BRA GO Central Europe operations
Org level 5	30 6	Sales branch 30 - Sao Paulo - BRA Sales branch 6 - Paris - FRA

Balanced hierarchy - view

You can use views to balance the organization hierarchy by side-stepping the relationship between GO Accessories and the European companies. Use a view when you want to create samples that show metrics and aggregations that are distinct for each corporate entity.

Gosales VIEW_BALANCED_ORG

```
SELECT TOP 100 PERCENT ORGANIZATION_1.ORGANIZATION_PARENT AS ORG_LEVEL1_CODE,
    gosaleshr.ORGANIZATION.ORGANIZATION_PARENT AS ORG_LEVEL2_CODE,
    gosaleshr.ORGANIZATION.ORGANIZATION_CODE
FROM    gosaleshr.ORGANIZATION ORGANIZATION_1 RIGHT OUTER JOIN
    gosaleshr.ORGANIZATION ON ORGANIZATION_1.ORGANIZATION_CODE =
    gosaleshr.ORGANIZATION.ORGANIZATION_PARENT
WHERE (gosaleshr.ORGANIZATION.ORGANIZATION_CODE BETWEEN N'006' AND
    N'8820')
ORDER BY ORGANIZATION_1.ORGANIZATION_PARENT,
    gosaleshr.ORGANIZATION.ORGANIZATION_PARENT DESC,
    gosaleshr.ORGANIZATION.ORGANIZATION_CODE
```

gosalesdw.VIEW_BALANCED_ORG

```
SELECT gosalesdw.GO_ORG_DIM.ORGANIZATION_KEY, GO_ORG_DIM_1.ORGANIZATION_PARENT AS
    ORG_LEVEL1_CODE, GO_ORG_NAME_LOOKUP_1.ORGANIZATION_NAME_EN AS ORG_LEVEL1_NAME,
    gosalesdw.GO_ORG_DIM.ORGANIZATION_PARENT AS ORG_LEVEL2_CODE,
    GO_ORG_NAME_LOOKUP_2.ORGANIZATION_NAME_EN AS ORG_LEVEL2_NAME,
    gosalesdw.GO_ORG_DIM.ORGANIZATION_CODE AS ORG_CODE,
    gosalesdw.GO_ORG_NAME_LOOKUP.ORGANIZATION_NAME_EN AS ORG_NAME
FROM gosalesdw.GO_ORG_NAME_LOOKUP GO_ORG_NAME_LOOKUP_2 INNER JOIN
    gosalesdw.GO_ORG_DIM INNER JOIN
    gosalesdw.GO_ORG_NAME_LOOKUP ON
    gosalesdw.GO_ORG_DIM.ORGANIZATION_CODE =
    gosalesdw.GO_ORG_NAME_LOOKUP.ORGANIZATION_CODE ON
    GO_ORG_NAME_LOOKUP_2.ORGANIZATION_CODE =
    gosalesdw.GO_ORG_DIM.ORGANIZATION_PARENT INNER JOIN
    gosalesdw.GO_ORG_NAME_LOOKUP GO_ORG_NAME_LOOKUP_1 INNER JOIN
    gosalesdw.GO_ORG_DIM GO_ORG_DIM_1 ON GO_ORG_NAME_LOOKUP_1.ORGANIZATION_CODE =
    GO_ORG_DIM_1.ORGANIZATION_PARENT ON
    gosalesdw.GO_ORG_DIM.ORGANIZATION_PARENT =
    GO_ORG_DIM_1.ORGANIZATION_CODE
WHERE (gosalesdw.GO_ORG_DIM.ORGANIZATION_CODE BETWEEN N'006' AND N'8820')
```

GO Departments

The Sample Outdoors companies are divided into operating and corporate departments. Each organization has the same structure.

Applications for financial and planning data use the department structure in the sample budgets and forecasts.

Sales branches under operations, and HR departments under corporate, are the richest sources of data for querying and analysis.

Table 15. Sample Outdoors Departments

Division (GL)	Department (GL)
Corporate	Sales (Corporate)
Corporate	Marketing (1750)
Corporate	IS&T

Table 15. Sample Outdoors Departments (continued)

Division (GL)	Department (GL)
Corporate	Human Resources
Corporate	Finance
Corporate	Procurement
Operations	Sales branches
Operations	Production and Distribution
Operations	Customer Service

GO Operations

Metrics such as revenue, quantities, and cost of goods, aggregate through levels from sales rep to branch to operations under each corporate banner.

Data from other departments ties into the sales data. For example, the marketing department implements sales campaigns.

Human Resources

Human resources data exists for metrics such as compensation, benefits, training and surveys. Employees are paid hourly, or receive salaries. Compensation may include commission and bonuses. Vacation and sick days are part of the benefits package.

Each Sample Outdoors organization is staffed by a variety of employees including personnel for marketing, finance, sales, as well as those in other departments. Employee history exists for employees that change positions or have new managers.

In the warehouse data (GOSALESDW), the employee history is contained in a slowly changing dimension, in which an employee code may be a repeating value and the employee key is unique.

In the transactional data (GOSALES), employee history requires filtering on a date range in order to assign transaction records to employees in the appropriate time period. Filtering also prohibits double-counting that can occur with repeating employee codes. For example, to capture an accurate history of an employee at the time of training, the date range is qualified.

```
WHERE (TRAINING_DETAILS.TRAINING_DATE BETWEEN
EMPLOYEE_HISTORY.RECORD_START_DATE AND EMPLOYEE_HISTORY.RECORD_END_DATE)
```

Alternatively, you can select active records by filtering out employee history, which returns only the current manager or job position of the employee.

```
WHERE (EMPLOYEE_HISTORY.RECORD_END_DATE IS NULL)
```

Warehouse Schema

Gosalesdw Schema

Table 16. Great Outdoors sales data warehouse schema

Table	Record count	Functional area
DIST_INVENTORY_FACT	53,837	Distribution

Table 16. Great Outdoors sales data warehouse schema (continued)

Table	Record count	Functional area
DIST_PRODUCT_FORECAST_FACT	129,096	Distribution
DIST_RETURN_REASON_DIM	5	Distribution
DIST_RETURNED_ITEMS_FACT	10,249	Distribution
EMP_EXPENSE_FACT	127,984	Distribution
EMP_EXPENSE_PLAN_FACT	30,150	Distribution
EMP_EXPENSE_TYPE_DIM	39	Distribution
FIN_ACCOUNT_DIM	242	Finance
FIN_FINANCE_FACT	164,132	Finance
FIN_SUBM_DIM	52	Finance
GO_BRANCH_DIM	29	Geography
GO_REGION_DIM	21	Geography
MRK_PROD_SURVEY_TARG_FACT	5,824	Marketing
MRK_PRODUCT_SURVEY_DIM	7	Marketing
MRK_PRODUCT_SURVEY_FACT	165,074	Marketing
MRK_PROMOTION_DIM	112	Marketing
MRK_PROMOTION_FACT	11,034	Marketing
MRK_PROMOTION_PLAN_FACT	8,652	Marketing
MRK_RTL_SURVEY_DIM	9	Marketing
MRK_RTL_SURVEY_FACT	22,508	Marketing
MRK_RTL_SURVEY_TARG_FACT	64	Marketing
EMP_POSITION_DIM	57	Organization
GO_ORG_DIM	123	Organization
EMP_EMPLOYEE_DIM	972	Personnel
EMP_POSITION_SUMMARY_FACT	15,050	Personnel
EMP_RANKING_DIM	5	Personnel
EMP_RANKING_FACT	1,897	Personnel
EMP_RECRUITMENT_DIM	14	Personnel
EMP_RECRUITMENT_FACT	416	Personnel
EMP_SUCCESSION_FACT	181	Personnel
EMP_SUCCESSION_STATUS_DIM	5	Personnel
EMP_SUMMARY_FACT	24,233	Personnel
EMP_SURVEY_FACT	5,725	Personnel
EMP_SURVEY_TARG_FACT	20	Personnel
EMP_SURVEY_TOPIC_DIM	5	Personnel
EMP_TRAINING_DIM	42	Personnel
EMP_TRAINING_FACT	4,465	Personnel
GO_SATISFACTION_DIM	5	Personnel
SLS_PRODUCT_DIM	274	Product
MRK_ACTIVITY_STATUS_DIM	2	Retailer
SLS_RTL_DIM	847	Retailer

Table 16. Great Outdoors sales data warehouse schema (continued)

Table	Record count	Functional area
SLS_ORDER_METHOD_DIM	7	Sales
SLS_SALES_FACT	446,023	Sales
SLS_SALES_ORDER_DIM	446,023	Sales
SLS_SALES_TARG_FACT	233,625	Sales
GO_TIME_DIM	1,465	Time
Xgorev	15	Database Admin

Gosalesdw Lookup tables

Table 17. Great Outdoors sales data warehouse Lookup

Table	Record count	Functional area
EMP_EXPENSE_UNIT_LOOKUP	3	Expenses
FIN_ACCOUNT_CLASS_LOOKUP	5	Finance
FIN_ACCOUNT_NAME_LOOKUP	242	Finance
FIN_ACCOUNT_TYPE_LOOKUP	4	Finance
FIN_SUBM_CURRENCY_LOOKUP	7	Finance
FIN_SUBM_TYPE_LOOKUP	3	Marketing
MRK_BUNDLE_GROUP_LOOKUP	15	Marketing
MRK_CAMPAIGN_LOOKUP	12	Organization
EMP_POSITION_LOOKUP	57	Organization
GO_ORG_NAME_LOOKUP	123	Personnel
EMP_TERMINATION_LOOKUP	6	Personnel
GO_GENDER_LOOKUP	2	Product
SLS_PRODUCT_BRAND_LOOKUP	28	Product
SLS_PRODUCT_COLOR_LOOKUP	27	Product
SLS_PRODUCT_LINE_LOOKUP	5	Product
SLS_PRODUCT_LOOKUP	274 per language	Product
SLS_PRODUCT_SIZE_LOOKUP	55	Product
SLS_PRODUCT_TYPE_LOOKUP	21	Product
GO_TIME_QUARTER_LOOKUP	20	Time

Transaction Schemas

Gosales Schema

Table 18. Great Outdoors sales schema

Table	Record count	Functional area
INVENTORY_LEVELS	53,837	Distribution
PRODUCT_FORECAST	129,096	Distribution
RETURN_REASON	5	Distribution
RETURNED_ITEM	10,249	Distribution

Table 18. Great Outdoors sales schema (continued)

Table	Record count	Functional area
BRANCH	29	Geography
COUNTRY	21	Geography
SALES_REGION	5	Geography
PRODUCT	274	Product
PRODUCT_BRAND	28	Product
PRODUCT_LINE	5	Product
PRODUCT_TYPE	21	Product
CONVERSION_RATE	624	Sales
EURO_CONVERSION	8	Sales
ORDER_DETAILS	446,023	Sales
ORDER_HEADER	53,256	Sales
ORDER_METHOD	7	Sales
SALES_TARGET	233,625	Sales
TIME_DIMENSION	1,465	Time
xgorev	16	Database admin

Gosales Lookup tables

Table 19. Great Outdoors sales Lookup tables

Table	Record Count	Functional Area
PRODUCT_COLOR_LOOKUP	27	Product
PRODUCT_NAME_LOOKUP	274 per language	Product
PRODUCT_SIZE_LOOKUP	55	Product
CURRENCY_LOOKUP	21	Sales
TIME_QUARTER_LOOKUP	20	Time

Gosaleshr Schema

Table 20. Great Outdoor sales human resources schema

Table	Record count	Functional area
EMPLOYEE_EXPENSE_DETAIL	127,997	Expenses
EMPLOYEE_EXPENSE_PLAN	37,317	Expenses
EMPLOYEE_SUMMARY	24,233	Expenses
EMPLOYEE_SURVEY_TOPIC	5	Expenses
EXPENSE_GROUP	10	Expenses
EXPENSE_TYPE	39	Expenses
EXPENSE_UNIT	3	Expenses
ORGANIZATION	123	Organization
POSITION_DEPARTMENT	445	Organization
POSITION_SUMMARY	15,050	Organization
EMPLOYEE	766	Personnel

Table 20. Great Outdoor sales human resources schema (continued)

Table	Record count	Functional area
EMPLOYEE_HISTORY	972	Personnel
EMPLOYEE_SURVEY_RESULTS	5,725	Personnel
EMPLOYEE_SURVERY_TARGETS	20	Personnel
RANKING	5	Personnel
RANKING_RESULTS	1,898	Personnel
RECRUITMENT	416	Personnel
RECRUITMENT_MEDIUM	14	Personnel
RECRUITMENT_TYPE	7	Personnel
SATISFACTION_INDEX	5	Personnel
SUCCESSION_DETAILS	182	Personnel
SUCCESSOR_STATUS	5	Personnel
TRAINING	42	Personnel

Gosaleshr Lookup tables

Table 21. Great Outdoors sales human resources Lookup tables

Table	Record count	Functional area
DEPARTMENT_LOOKUP	12	Organization
POSITION_LOOKUP	45	Organization
GENDER_LOOKUP	2	Personnel
TERMINATION_LOOKUP	6	Personnel
TRAINING_DETAILS	4,471	Personnel

Gosalesmr Schema

Table 22. Great Outdoors sales marketing schema

Table	Record count	Functional area
PRODUCT_SURVEY_RESULTS	165,074	Marketing
PRODUCT_SURVEY_TARGETS	5,824	Marketing
PRODUCT_SURVEY_TOPIC	7	Marketing
PROMOTION_BUNDLE_GROUP	15	Marketing
PROMOTION_CAMPAIGN	12	Marketing
PROMOTION_PLAN	8,652	Marketing
PROMOTIONS	112	Marketing
RETAILER_SURVEY_RESULTS	22,508	Marketing
RETAILER_SURVEY_TARGETS	64	Marketing
RETAILER_SURVEY_TOPIC	9	Marketing

Gosalesrt Schema

Table 23. Great Outdoors sales retailer

Table	Record count	Functional area
ACTIVITY_STATUS_LOOKUP	2	Retailer
RETAILER	562	Retailer
RETAILER_ACTIVITY	17,754	Retailer
RETAILER_CONTACT	847	Retailer
RETAILER_SITE	847	Retailer
RETAILER_SITE_MB	847	Retailer
RETAILER_TYPE	8	Retailer

Chapter 9. Fix packs and upgrades

Fix packs contain updates for specific products in IBM Business Intelligence Pattern with BLU Acceleration. An upgrade contains updates to the IBM Business Intelligence Pattern with BLU Acceleration pattern type.

You can apply the following types of fix packs and upgrades to your environment:

Business Intelligence Pattern with BLU Acceleration combined product fix pack

Contains product fixes that are critical to IBM Business Intelligence Pattern with BLU Acceleration. A combined product fix pack can contain fix packs for multiple products or a single product. For example, a combined fix pack can contain a fix for IBM DB2 and a fix for IBM Cognos Business Intelligence. A combined fix pack can also contain a fix for DB2 or a fix for Cognos BI.

When you apply a combined fix pack, the fixes are applied to the appropriate products in one maintenance operation.

DB2 fix pack

Contains an IBM DB2 fix. This type of fix pack is not commonly applicable for Business Intelligence Pattern with BLU Acceleration. If there is a DB2 fix pack that is applicable to Business Intelligence Pattern with BLU Acceleration, then it will be in a combined fix pack.

IBM Cognos Business Intelligence fix pack

Contains a Cognos Business Intelligence fix. This type of fix pack is not commonly applicable for Business Intelligence Pattern with BLU Acceleration. If there is a Cognos Business Intelligence fix pack that is applicable to Business Intelligence Pattern with BLU Acceleration, then it will be in a combined fix pack.

IBM Business Intelligence Pattern with BLU Acceleration upgrade

Contains improvements and fixes for the IBM Business Intelligence Pattern with BLU Acceleration pattern type.

Base operating system image fix packs

For information about how to update your base operating system virtual image, go to the IBM PureApplication System Information Center (pic.dhe.ibm.com/infocenter/psappsys/v1r1m0/index.jsp) and search for *updating base operating system virtual images*.

Applying a fix pack includes the following high-level steps:

- • Back up the IBM Cognos Business Intelligence content store and internal data store. For more information, see “Determine a backup strategy” on page 48.
- • Search for “Business Intelligence Pattern with BLU Acceleration” in Fix Central (www.ibm.com/support/fixcentral/) to get a list of all available fix packs.
- • Download the fix pack that you want to install to your computer.
- • Add the fix pack to the catalog.
- • Apply the fix pack to a role.

Combined product fix packs

A combined product fix pack contains product fixes that are critical to IBM Business Intelligence Pattern with BLU Acceleration.

A combined product fix pack can contain fix packs for multiple products or a single product. For example, a combined fix pack can contain a fix for IBM DB2 and a fix for IBM Cognos Business Intelligence. A combined fix pack can also contain a fix for DB2 or a fix for Cognos BI.

When you apply a combined fix pack, the fixes are applied to the appropriate products in one maintenance operation.

Adding combined or Cognos Business Intelligence fix packs to the catalog

You must add a combined or IBM Cognos Business Intelligence fix pack to the catalog before you can apply it to your environment.

Before you begin

Before you add fix packs to the catalog, ensure that the following prerequisites are met:

- The user name that you use to log in to IBM PureApplication System must have upload and download permissions.
- Download the fix pack from Fix Central to your computer.

Procedure

1. In the **Workload Console**, click **Catalog > Emergency Fixes**.
2. Click the **New** icon on the toolbar.
3. Type a name in the **Emergency fix name** field. This name is the identifier for the fix pack that you are adding to the catalog. Use a meaningful name that identifies the type of fix pack and software version.
4. Optional: In the **Description** field, enter the description for the fix pack.
5. Click **OK**.
6. Click **Browse**, select a file to upload, and click **Open**.
7. Click **Upload**. You are automatically assigned access to a fix pack when you create it.
8. If you want to allow other users to apply this fix to various instances, beside **Access granted to**, click **Add more**, and select a user name.
9. Select a severity level for the fix pack. This field is informational and identifies the urgency for applying this fix pack.
10. Under **Applicable to**, select the plug-ins to which the fix pack applies.

Tip: To see a list of available plug-ins, put your cursor in the **Plugins** field. A list of available plug-ins displays. If you type in the field, the plug-ins that match the text are listed.

What to do next

Apply the fix pack to your environment.

Adding combined or Cognos Business Intelligence fix packs to the catalog by using the command-line tool

Using the command-line tool is an alternative way to add a combined or IBM Cognos Business Intelligence fix pack to the catalog. If the fix pack is large, it can be faster to use the command-line tool to add the fix pack to the catalog.

Before you begin

Before you perform this task, ensure that the following prerequisites are met:

- The command-line tool is on your computer. For information about how to download the command-line tool, see “Downloading the command-line tool” on page 27.
- The fix pack that you want to add to the catalog is on the computer and in the directory where you run the command-line tool.

About this task

The command-line tool can be run in either interactive mode or batch mode from a local machine running either a Windows or Linux operating system. This procedure covers the steps to run the command-line tool in interactive mode. For information about running the command-line tool in batch mode, see the IBM PureApplication System documentation.

Procedure

1. On a command prompt or shell, change directory to the installation location of the command-line tool.
2. Invoke the command-line tool in interactive mode.
3. Create an entry in the catalog for the fix pack by typing the following command:

```
deployer.fixes.create({"name":"entry_name","description":  
"entry_description"})
```

where *entry_name* is the name of the catalog entry and *entry_description* is a description for the entry.

The following example command creates an entry in the catalog with a description of BIWithBLU fix pack for BI and DB2, and a name of BIWithBLU fix pack V1:

```
deployer.fixes.create({'name':'BIWithBLU fix pack V1',  
'description':'BIWithBLU fix pack for BI and DB2'})
```

4. Upload fix pack to the catalog, and associate the fix pack with the entry by typing the following command:

```
deployer.fixes.list({'name':'entry_name'})[0].archive.set  
("fix_path_file_path")
```

where *entry_name* is the name of the catalog entry and *fix_pack_file_path* is the path to the fix pack on your computer.

The following example command uploads the BIWithBLU_BIANDDB2V1.tar.gz fix pack and associates it to the BIWithBLU fix pack V1 entry in the catalog:

```
deployer.fixes.list({'name':'BIWithBLU fix pack V1'})[0].archive.set  
("/BIWithBLU_BIANDDB2V1.tar.gz")
```

5. In the **Workload Console**, click **Catalog > Emergency Fixes**.
6. Under **Emergency Fixes**, select the fix pack that you uploaded.

7. If you want to allow other users to apply this fix to various instances, beside **Access granted to**, click **Add more** and select a user name.
8. Select a severity level for the fix pack. This field is informational and identifies the urgency for applying this fix pack.
9. Under **Applicable to**, select the plug-ins to which the fix pack applies.

Tip: To see a list of available plug-ins, put your cursor in the **Plugins** field. A list of available plug-ins displays. If you type in the field, the plug-ins that match the text are listed.

Applying combined fix packs

A combined fix pack can contain fixes for multiple products or for a single product. The advantage of a combined fix pack is that all fixes that it contains are applied in one maintenance operation.

Before you begin

Before you perform this task, ensure that the following prerequisites are met:

- The combined fix pack that you want to apply is in the catalog.
- The user name that you use to log in to IBM PureApplication System must have upload and download permissions.

About this task

Use the following procedure for a combined fix pack that you found in Fix Central by searching for “Business Intelligence Pattern with BLU Acceleration”, even if the fix pack contains a single fix for IBM DB2 or a single fix for IBM Cognos Business Intelligence.

Procedure

1. In the **Workload Console**, click **Instances > Virtual Applications**.
2. In the drop-down list, select your version of IBM Business Intelligence Pattern with BLU Acceleration.
3. In the **Virtual Application Instances** list, select the virtual application instance that you want to update.
4. Click the **Maintain** icon.
5. Click the **Manage** icon. The Virtual Application Console opens.
6. Click the **Operations** tab, and from the Operations palette select **ProductPatches**.
7. Under **Product Maintenance**, expand **Install stack products updates**.
8. Click **Select** and select the fix pack from the list.
9. Type the **User Name** and **Password** that you use to log in to IBM PureApplication System and click **Submit**. The status of the operation is displayed under **Operation Execution Results**. The fix pack operation starts when the status changes to **Done**.
10. In the **Workload Console**, expand **Virtual machine perspective**. The fix pack has been applied when the **Middleware Status** of each of the roles is **Running**.
11. Click the **Resume** icon to exit maintenance mode.

Applying Cognos Business Intelligence fix packs

After you have added an IBM Cognos Business Intelligence fix pack to the catalog, you can apply it.

Before you begin

The IBM Cognos Business Intelligence fix pack that you want to apply must be in the catalog.

About this task

Use this procedure to apply an individual IBM Cognos BI fix pack. If you have a combined fix pack that contains a fix for IBM Cognos BI, then perform “Applying combined fix packs” on page 86.

Procedure

1. In the **Workload Console**, click **Instances > Virtual Applications**.
2. In the drop-down list, select a version of IBM Business Intelligence Pattern with BLU Acceleration.
3. In the **Virtual Application Instances** list, select the virtual application instance that you want to update.
4. Click the **Maintain** icon.
5. Click the **Manage** icon. The Virtual Application Console opens.
6. From the **Operations** palette, select the Business Intelligence application **CognosBICM1**.
7. Under **Product maintenance**, expand **Install IBM Cognos Business Intelligence Updates**.
8. Click **Select**, select the fix pack from the list, and click **Submit**. The status of the operation is displayed under **Operation Execution Results**. The fix pack operation starts when the status changes to **Done**.
9. In the **Workload Console**, expand **Virtual machine perspective**. The fix pack has been applied when the **Middleware Status** of each of the roles is **Running**.
10. Click the **Resume** icon to exit maintenance mode.

Uninstalling a Cognos Business Intelligence fix pack

To uninstall an IBM Cognos Business Intelligence fix pack, you must use the command line interface.

Before you begin

Back up your system and the content store.

About this task

The `\opt1\ibm\cognos\c10_64\uninstall\backup` directory contains folders for each fix pack that has been applied to your system. A backup folder contains copies of specific files that were modified by the fix pack. It does not contain the full content of the original directory.

You must perform this procedure on each virtual machine that has the fix pack applied.

Procedure

1. Log in to the virtual machine.
2. Type the following commands:

```
sudo su
su - cognos
```
3. Verify that the JAVA_HOME environment variable is pointing to the correct directory and, if necessary, set it by typing the following commands:

```
echo $JAVA_HOME
/opt1/ibm/cognos/c10_64/java/jre
```
4. Go to *cognos_install_location/bin64*.
5. Stop the IBM Cognos service by typing the following command:

```
./cogconfig.sh -stop
```
6. Go to the *\opt1\ibm\cognos\c10_64\uninstall\backup* directory. To determine the version to revert to, look at the *cmplst.txt* file in the backup folder for the fix pack that you want to revert.
7. Copy the contents of the backup folder and paste it to *\opt1\ibm\cognos\c10_64*. This action overwrites the files for the fix pack that you are uninstalling.

Tip: If you receive a copying error, type the following command to clear the cache and then try the copy again:

```
slibclean
```
8. Start the IBM Cognos service by typing the following command:

```
./cogconfig.sh
```

DB2 fix packs

An IBM DB2 fix pack contains an IBM DB2 fix.

This type of fix pack is not commonly applicable for Business Intelligence Pattern with BLU Acceleration. If there is a DB2 fix pack that is applicable to Business Intelligence Pattern with BLU Acceleration, then it will be in a combined fix pack.

Adding DB2 fix packs to the catalog

Use this procedure to add an IBM DB2 fix pack to the catalog that you have specifically searched for.

Before you begin

Before you add fix packs to the catalog, ensure that the following prerequisites are met:

- You have uploading and downloading permissions.
- You have downloaded the fix pack from Fix Central to your computer.

About this task

Use this procedure to add an individual IBM DB2 fix pack to the catalog. If you have a combined fix pack that contains a fix for IBM DB2, then perform “Adding combined or Cognos Business Intelligence fix packs to the catalog” on page 84.

Procedure

1. In the **Workload Console**, click **Catalog > DB2 Fix Packs**.
2. Click the **New** icon on the toolbar.
3. Type a name in the **DB2 Fix Pack Name** field. This name is the identifier for the fix pack that you are adding to the catalog. Use a meaningful name that identifies the type of fix pack and software version.
4. Optional: In the **Description** field, enter the description for the fix pack.
5. Click **Save**.
6. Click **Browse**, select a file to upload, and click **Open**.

What to do next

Apply the fix pack to your environment.

Adding DB2 fix packs to the catalog by using the command-line tool

Using the command-line tool to add an IBM DB2 fix pack to the catalog is an alternative to using IBM Workload Deployer. If the fix pack is large, it can be faster to use the command-line tool to add the fix pack to the catalog.

Before you begin

Before you perform this task, ensure that the following prerequisites are met:

- The command-line tool is on the computer you are working on. For information about how to download the command-line tool, see “Downloading the command-line tool” on page 27.
- The IBM DB2 fix pack that you want to add to the catalog must be on the computer in the directory where you run the command-line tool.

About this task

The command-line tool can be run in either interactive mode or batch mode from a local machine running either a Windows or Linux operating system. This procedure covers the steps to run the command-line tool in interactive mode. For information about running the command-line tool in batch mode, see the IBM PureApplication System documentation.

Procedure

1. On a command prompt or shell, change directory to the installation location of the command-line tool.
2. Invoke the command-line tool in interactive mode.
3. Create an entry in the catalog for the fix pack by typing the following command:

```
deployer.db2fixpacks.create({"name":"entry_name","description":  
"entry_description"})
```

where *entry_name* is the name of the catalog entry and *entry_description* is a description for the entry.

The following example command and example output creates an entry in the catalog with a description of db2 v105fp2, an ID of f552d174-89da-4799-a130-82fe2973d8bd, and a name of v105fp2:

```

deployer.db2fixpacks.create({"name":"v105fp2","description":"db2 v105fp2"})
{'description': 'db2 v105fp2', 'id': 'f552d174-89da-4799-a130-82fe2973d8bd',
 'name': 'v105fp2'}

```

4. Upload the fix pack to the catalog by typing the following command:

```

deployer.db2fixpacks.upload("fix_pack_file_path")

```

where *fix_pack_file_path* is the file path to the fix pack on your computer.

The following example command and example output uploads fix pack v10.5fp2_aix64_server_r.tar.gz to the catalog:

```

deployer.db2fixpacks.upload("/tmp/v10.5fp2_aix64_server_r.tar.gz")

```

```

{'file': 'v10.5fp2_aix64_server_r.tar.gz',
 'file_name': 'v10.5fp2_aix64_server_r.tar.gz',
 'fileName': 'v10.5fp2_aix64_server_r.tar.gz'}

```

5. Associate the IBM DB2 fix pack with the entry in the catalog by typing the following command:

```

deployer.db2fixpacks.get("fix_pack_ID").update("python_directory_object")

```

where *fix_pack_ID* is the ID of the entry in the catalog and *python_directory_object* is the name, description, ID, and name of the fix pack file.

For example,

```

deployer.db2fixpacks.get("f552d174-89da-4799-a130-82fe2973d8bd").update
({"name": "v105fp2", "description": "db2 v105fp2",
 "id": "f552d174-89da-4799-a130-82fe2973d8bd",
 "packageFile": "v10.5fp2_aix64_server_r.tar.gz"})

```

Applying DB2 fix packs

After you have added an IBM DB2 fix pack to the catalog, you can apply it.

Before you begin

Before you apply DB2 fix packs, ensure that the following prerequisites are met:

The DB2 fix pack that you want to apply must be in the catalog.

Back up your DB2 database. For information about how to back up your DB2 database, see the IBM DB2 information center (pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp).

About this task

Use this procedure to apply an individual IBM DB2 fix pack. If you have a combined fix pack that contains a fix for IBM DB2, then perform “Applying combined fix packs” on page 86.

You must apply the DB2 fix pack to both DB2 roles.

During the upgrade process the system is restarted and all application connections to the database are closed. The database is stopped during the upgrade process and restarted when the upgrade is complete.

Procedure

1. In the **Workload Console**, click **Instances > Virtual Applications**.

2. In the drop-down list, select a version of IBM Business Intelligence Pattern with BLU Acceleration.
3. In the **Virtual Application Instances** list, select the virtual application instance that you want to update.
4. Click the **Maintain** icon.
5. Click the **Manage** icon. The Virtual Application Console opens.
6. Click **Operations**, and select the Content Manager role. The Content Manager role ends with CM1.DB2.
7. Under **Fundamental**, expand **Apply a DB2 fix pack**.
8. In the **Fix pack** field, select the fix pack to apply from the list, and click **Submit**. The status of the operation is displayed under **Operation Execution Results**. The fix pack operation starts when the status changes to **Done**.
9. Under **Operations**, select the DB2 primary node role and perform the same steps to start the fix pack operation for it. The DB2 primary node role ends with Primary.DB2.
10. In the **Workload Console** tab, expand **Virtual machine perspective**. The fix pack has been applied when the **Middleware Status** of each of the roles is **Running**.
11. Click the **Resume** icon to exit maintenance mode.

Uninstalling a DB2 fix pack

You can revert to a previous fix pack or a major release of IBM DB2.

Before you begin

Before you uninstall a fix pack, you must remove or uninstall functionality that is not part of the fix pack or major release that you are reverting to. For example, before you revert to IBM DB2 Version 10.5 you must make the following changes to the database:

- Drop work action sets that were associated with a workload.
- Drop U0WTOTALTIME thresholds.

Procedure

1. Log in to the virtual machine and open the IBM DB2 command line interface.
2. Type the following command:

```
./installFixPack -f level -b DB2DIR
```

where *DB2DIR* is the installation location of the IBM DB2 product that you want to revert.

3. Associate the instance with the reverted IBM DB2 version. You can update all instances at one time by running the `thedb2iupdt -D` command from the new directory.

Base operating system image fix packs

You can update your base operating system virtual image by using the workload console.

For information about how to update your base operating system virtual image, go to the IBM PureApplication System Information Center (pic.dhe.ibm.com/infocenter/psappsys/v1r1m0/index.jsp) and search for *updating base operating system virtual images*.

Business Intelligence Pattern with BLU Acceleration upgrade

An IBM Business Intelligence Pattern with BLU Acceleration upgrade contains improvements and fixes for the IBM Business Intelligence Pattern with BLU Acceleration pattern type.

An upgrade can include updates to the following plug-ins:

- biima
- biima-ldap
- biimaelb

Upgrading Business Intelligence Pattern with BLU Acceleration

Upgrading IBM Business Intelligence Pattern with BLU Acceleration installs improvements and fixes to the IBM Business Intelligence Pattern with BLU Acceleration pattern type.

Before you begin

You must have the plug-in file that you want to upgrade to on your computer.

Procedure

1. In the **Workload Console**, click **Cloud > System Plug-ins**.
2. In the drop-down list, select a version of IBM Business Intelligence Pattern with BLU Acceleration.
3. Click the **New** icon on the toolbar.
4. Click **Browse**, select the plug-in to upload, and click **Open**.
5. Click **Instances > Virtual Application**.
6. In the drop-down list, select a version of IBM Business Intelligence Pattern with BLU Acceleration.
7. In the **Virtual Application Instances** list, select the virtual application instance that you want to upgrade.
8. Beside the **Pattern type** field, click **Check for upgrades**. All available updates are displayed.
9. Click **Update**. All available updates are applied. The upgrade is complete when the status icon of the virtual application instance is running. Review the updated deployment and click **Commit**. If the upgrade fails, the backup data is restored and the virtual application is returned to its previous state.

Appendix A. Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products.

For information about the commitment that IBM has to accessibility, see the IBM Accessibility Center(www.ibm.com/able).

Keyboard navigation

This product uses standard browser navigation keys.

The Virtual Application Builder is not keyboard accessible. If you need to work with virtual application patterns without using a mouse, you can use the command-line interface provided by your system or the REST API.

IBM Cognos software

In IBM Cognos Administration, you can enable system-wide settings to create accessible report output. For more information, see the *IBM Cognos Business Intelligence Administration and Security Guide*.

IBM Cognos HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Appendix B. Troubleshooting and Support

The troubleshooting information helps you understand, isolate, and resolve problems with IBM Business Intelligence Pattern with BLU Acceleration, and contains instructions for using the problem determination resources that are provided with your products.

Operations performed in BI Pattern with BLU Acceleration are recorded in various log files for tracking purposes. For information about accessing and viewing the logs, see the documentation for your system. Logs are organized by virtual machine name. They are grouped under categories such as DeploymentInlet, OS, IWD Agent, and usually also under the role of the virtual machine. For information about the Business Intelligence log files, review the *IBM Cognos Business Intelligence 10.2.1 Troubleshooting Guide*.

Troubleshoot a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM technical-support representative know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, which can then lead to a resolution of the problem.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is “What is the problem?” This question might seem straightforward; however, you can break it down into several focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is the problem a loop, hang, crash, performance degradation, or incorrect result?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

If one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and the hardware. Confirm that you are running within an environment that is supported; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for cases that are one-time occurrences. You can most easily develop a timeline by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as an upgrade or an installation of software or hardware?

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to occur?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

Problems that you can reproduce are often easier to solve. However, problems that you can reproduce can have a disadvantage. If the problem has a significant business impact, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation. Answer the following questions:

- Can the problem be re-created on a test system?
- Are multiple users or applications encountering the same type of problem?

- Can the problem be re-created by running a single command, a set of commands, or a particular application?

Searching knowledge bases

You can often find solutions to problems by searching the IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

About this task

You can find useful information by searching the information center for IBM Cognos, but sometimes you need to look beyond the information center to resolve problems.

Procedure

To search knowledge bases for information that you need, use one or more of the following approaches:

- Find the content that you need by using the IBM Support Portal (www.ibm.com/support/entry/portal/).
The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services. The IBM Support Portal lets you access the IBM electronic support portfolio from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution. Familiarize yourself with the IBM Support Portal by viewing the demo videos (www.ibm.com/blogs/SPNA/entry/the_ibm_support_portal_videos) about this tool. These videos introduce you to the IBM Support Portal, explore troubleshooting and other resources, and demonstrate how you can tailor the page by moving, adding, and deleting portlets.
- Search for content about IBM Cognos by using one of the following additional technical resources:
 - IBM Cognos BI APARs (problem reports) (www.ibm.com/search/csass/search?q=&ibm-go-btn.x=0&ibm-go-btn.y=0&ibm-go-btn=Go&sn=spe&filter=language:en+collection:dblue+ibmcontenttypeid:CT748&prod=U487531B48832U38::#q%253d%2526filter%253d%2526tsswrelease%253a841%2520%2526ibmcontenttypeid%253aCT748%2520%2526language%253aen%2520%2526collection%253adbblue%2526prod%253dU487531B48832U38%2526sn%253dspe%2526sortby%253d-moddate%2526o%253d0)
 - Searching technotes (www.ibm.com/support/docview.wss?uid=swg21373237).
 - IBM Cognos forums and communities (www.ibm.com/support/entry/portal/Forums_communities/Software/Information_Management/Cognos_Business_Intelligence_and_Financial_Performance_Management).
 - Analytics Center of Excellence community (www.analyticszone.com/homepage/web/displayAcePage.action)
- Search for content by using the IBM masthead search. You can use the IBM masthead search by typing your search string into the Search field on any ibm.com® page.
- Search for content by using any external search engine, such as Google, Yahoo, or Bing. If you use an external search engine, your results are more likely to include information that is outside the ibm.com domain. Sometimes you can find useful problem-solving information about IBM products in newsgroups, forums, and blogs that are not on ibm.com.

Tip: Include “IBM” and the name of the product in your search if you are looking for information about an IBM product.

Getting fixes

A product fix might be available to resolve your problem.

Procedure

To find and install fixes:

1. Determine which fix you need (Fix Central) (opens in new window) (www.ibm.com/support/fixcentral/)
2. Download the fix. Open the download document and follow the link in the “Download package” section.
3. Apply the fix by following the instructions in the “Installation Instructions” section of the download document.
4. Subscribe to receive weekly email notifications about fixes and other IBM Support information.

Contacting IBM Support

IBM Support provides access to a variety of IBM resources for help with software questions.

Before you begin

After trying to find your answer or solution by using other self-help options such as technotes, you can contact IBM Support. Before contacting IBM Support, your company must have an active IBM maintenance contract, and you must be authorized to submit problems to IBM. You should also have the following information at hand:

- Your customer identification number
- Your service request number, if it is an ongoing service request
- The phone number where you can be reached
- The version of the software you use
- The version of the operating environment you use
- A description of what you were doing when the problem occurred
- The exact wording of any error messages that display
- Any steps you took to attempt to solve the problem

For information about the types of available support, see the Support portfolio topic in the *Software Support Handbook* (opens in new window).

Procedure

Complete the following steps to contact IBM Support with a problem:

1. Define the problem, gather background information, and determine the severity of the problem. For more information, see the Getting IBM support (www14.software.ibm.com/webapp/set2/sas/f/handbook/getsupport.html) (opens in new window) topic in the *Software Support Handbook*.
2. Gather diagnostic information.
3. Submit the problem to IBM Support in one of the following ways:

- By using IBM Support Assistant (ISA): Use this feature to open, update, and view an Electronic Service Request with IBM. Any data that has been collected can be attached to the service request. This expedites the analysis and reduces the time to resolution.
- Online through the IBM Support Portal (www.ibm.com/software/support/) (opens in new window): You can open, update, and view all your Service Requests from the Service Request portlet on the Service Request page.
- By phone: For the phone number to call, see the Directory of worldwide contacts (www.ibm.com/planetwide/) (opens in new window) web page.

Results

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support website daily, so that other users who experience the same problem can benefit from the same resolution.

Exchanging information with IBM

To diagnose or identify a problem, you might need to provide IBM Support with data and information from your system.

In other cases, IBM Support might provide you with tools or utilities to use for problem determination.

Sending information to IBM Support

To reduce the time that it takes to resolve your problem, you can send trace and diagnostic information to IBM Support.

Procedure

1. Open a problem management record (PMR). You can use the IBM Support Assistant (www.ibm.com/software/support/isa) (opens in new window) or the IBM Service Request tool (www.ibm.com/support/servicerequest) (opens in new window).
2. Collect the diagnostic data that you need. Diagnostic data helps reduce the time that it takes to resolve your PMR. You can collect the diagnostic data manually or automatically.
3. Compress the files by using the TRSMAN or AMATERSE program. Download the free utility from the IBM to the IBM Cognos BI system and then install the utility using the TSO RECEIVE command.
4. Transfer the files to IBM. You can use one of the following methods to transfer the files to IBM:
 - The Service Request tool (www.ibm.com/support/servicerequest) (opens in new window)
 - Standard data upload methods: FTP, HTTP
 - Secure data upload methods: FTPS, SFTP, HTTPS
 - Email

If you are using an IBM Cognos product and you use ServiceLink / IBMLink to submit PMRs, you can send diagnostic data to IBM Support in an email or by using FTP.

All of these data exchange methods are explained on the IBM Support site (www-01.ibm.com/software/support/exchangeinfo.html) (opens in new window).

Receiving information from IBM Support

Occasionally an IBM technical-support representative might ask you to download diagnostic tools or other files. You can use FTP to download these files.

Before you begin

Ensure that your IBM technical-support representative provided you with the preferred server to use for downloading the files and the exact directory and file names to access.

Procedure

To download files from IBM Support:

1. Use FTP to connect to the site that your IBM technical-support representative provided and log in as anonymous. Use your email address as the password.
2. Change to the appropriate directory:
 - a. Change to the `/fromibm` directory.
`cd fromibm`
 - b. Change to the directory that your IBM technical-support representative provided.
`cd nameofdirectory`
3. Enable binary mode for your session.
`binary`
4. Use the `get` command to download the file that your IBM technical-support representative specified.
`get filename.extension`
5. End your FTP session.
`quit`

Subscribing to Support updates

To stay informed of important information about the IBM products that you use, you can subscribe to updates.

About this task

By subscribing to receive updates, you can receive important technical information and updates for specific Support tools and resources. You can subscribe to updates by using one of two approaches:

RSS feeds and social media subscriptions

The following RSS feeds and social media subscriptions are available for IBM Cognos BI:

- RSS feed for a developerWorks® forum (www.ibm.com/developerworks/data/community/#cognos) (opens in new window).
- RSS feed for the Support site for IBM Cognos Business Intelligence (www-947.ibm.com/systems/support/myfeed/xmlfeeder.wss?feeder.requid=feeder.create_public_feed&feeder.feedtype=RSS&feeder.maxfeed=25&OC=SS9RTN&feeder.subdefkey=swgimgmt&feeder.channel.title=Cognos%208)

%20Business%20Intelligence&feeder.channel.descr=The%20latest%20updates%20about%20Cognos%208%20Business%20Intelligence) (opens in new window)

For general information about RSS, including steps for getting started and a list of RSS-enabled IBM web pages, visit the IBM Software Support RSS feeds (www.ibm.com/software/support/rss/) (opens in new window) site.

My Notifications

With My Notifications, you can subscribe to Support updates for any IBM product. You can specify that you want to receive daily or weekly email announcements. You can specify what type of information you want to receive, such as publications, hints and tips, product flashes (also known as alerts), downloads, and drivers. My Notifications enables you to customize and categorize the products that you want to be informed about and the delivery methods that best suit your needs.

Procedure

To subscribe to Support updates:

1. Subscribe to the *Product* RSS feeds.
2. To subscribe to My Notifications, begin by going to the IBM Support Portal (www.ibm.com/software/support/) (opens in new window) and clicking **My Notifications** in the **Notifications** portlet.
3. If you have already registered for My support, sign in and skip to the next step. If you have not registered, click **Register now**. Complete the registration form using your email address as your IBM ID and click **Submit**.
4. Click **Edit profile**.
5. Click **Add products** and choose a product category; for example, **Software**.
6. In the second list, select a product segment; for example, **Data & Information Management**.
7. In the third list, select a product subsegment, for example, **Databases**.
8. Select the products that you want to receive updates for.
9. Click **Add products**.
10. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
11. Select **Please send these documents by weekly email**.
12. Update your email address as needed.
13. In the **Documents list**, select the product category; for example, **Software**.
14. Select the types of documents that you want to receive information for.
15. Click **Update**.

Results

Until you modify your RSS feeds and My Notifications preferences, you receive notifications of updates that you have requested. You can modify your preferences when needed (for example, if you stop using one product and begin using another product).

Troubleshoot IBM Business Intelligence Pattern with BLU Acceleration

This section provides solutions for problems you may encounter when using IBM Business Intelligence Pattern with BLU Acceleration.

Cognos user cannot log in to virtual machines

In some cases, you can no longer log in to the virtual machines as the **cognos** user. This happens when the public key of the **cognos** user account was removed or the public key does not match the private key associated with the virtual application.

To fix this problem, you can manually update the public key in the `authorized_keys` file for the **cognos** user account. When you update the key manually, it does not affect the availability of your BI application.

To manually update the public key file for the **cognos** user account, perform the following steps:

1. On the virtual machine, log in as **virtuser**.

For example, you might run a command similar to the following command:

```
ssh -i id_rsa virtuser@machine_ip
```

where

- *id_rsa* represents a file in the current directory with the public key file
 - *machine_ip* represents the IP address of the virtual machine
2. Gain root access by running the following command: `sudo su -`.
 3. Copy the `authorized_keys` file from the `/home/virtuser/.ssh/` directory to the `/home/cognos/.ssh/` directory.
 4. Verify that the file ownership of the public key is the **cognos** user. If not, change the file ownership so that it is owned by the **cognos** user.

Deployment fails when starting IBM Business Intelligence Pattern with BLU Acceleration

If your deployment fails and no error message is displayed, review the logs for the following error:

Unable to print or tune the deployment.

Check the Administrative Namespace ID or credentials.

Procedure

1. Click **Instances** and then click **Virtual Applications**.
2. Select the virtual application that contains the `trace.log` file you want to review.
3. Click **Manage**.
4. Click the **Logging** tab.
5. Expand the CM1 virtual machine.
6. Expand the IWD Agent `.../logs/CM1 server name` folders.
7. Click the `trace.log` file. This will populate the log contents in the viewing pane.
8. In the viewing pane, search for error: CM-REQ-4342. This will help confirm the cause of the issue.

Results

ERROR Unable to print or tune the deployment.

Check the Administrative Namespace ID or credentials:

```
CM-REQ-4342 An error occurred with the client.
```

This issue may occur when a User Registry Component has been populated on the Virtual Application Builder canvas, but the Administrative namespace ID, Administrative username and Administrative password have not been accurately populated on the BI Pattern with BLU Acceleration component.

This message is written to the IBM BI Pattern with BLU Acceleration (IWD Agent) trace.log associated with the CM1 server.

To resolve the problem, accurately enter the Administrative namespace ID as defined in the User Registry Component into the **Administrative credentials section** of the IBM Business Intelligence Pattern with BLU Acceleration component. By entering the Administrative namespace ID along with an Administrative username and password, the required authentication references can be made.

Notices

This information was developed for products and services offered worldwide.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
Attention: Licensing
3755 Riverside Dr.
Ottawa, ON
K1V 1B7
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

Trademarks

IBM, the IBM logo and ibm.com 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 IBM trademarks is available on the Web at “ Copyright and trademark information ” at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.



Index

A

- accessibility 93
- accessing
 - Cognos BI 41
- active Content Manager 42
- Active Directory 22
- Aggregate Advisor 59
- aggregate tables 59
- aggregating
 - samples and time dimensions 72
- authentication
 - SSL using LDAP 21
- authentication providers
 - Active Directory 22
 - Tivoli directory server 22

B

- backups 48
- Blackberry push servers
 - updating primary list 53

C

- catalog
 - adding Cognos BI fix packs to 84
 - adding combined fix packs to 84
 - adding DB2 fix packs to 88
- change data capture 59, 61
- checklist
 - installation 7
 - post-installation 8
 - pre-installation 7
- Cognos BI
 - accessing 41
- Cognos Business Intelligence fix pack 85, 87
 - adding to catalog 84
- Cognos Cube Designer 63
 - installing 42
- Cognos Dynamic Query Analyzer
 - installing 42
- combined fix pack
 - adding to catalog 84
- command-line tool 85, 89
 - downloading 27
- components
 - configuration settings 9
- configuration settings 9
- configuring
 - pattern components 9
- Content Manager
 - active state 42
 - standby 42
- creating the pattern 30

D

- data preparation 57

- data source signons
 - securing 25
- database monitoring
 - IBM PureApplication Systems 54
- databases
 - backing up 48
- DB2 fix pack 89, 91
 - adding to catalog 88
- deploying
 - Database Performance Monitoring service 19
 - elastic load balancing (ELB) shared service 19
 - System Monitoring service 19
 - virtual applications 32
- deployment
 - planning 19
 - steps 27
 - verification 38
- deployment archives
 - exporting 44
 - importing 44
 - importing samples 65
 - model.zip 65
- disk space 60
- documentation
 - translated documents 43
 - translations 43
- domain lists
 - changing 51
- dynamic cubes
 - deploying sample data 67
 - importing sample models 65
 - optimizing 60
 - performance monitoring 55
 - samples 66

E

- enable SSL
 - BI components 20
- error messages
 - CM-REQ-4342 102
 - Unable to print or tune the deployment 102
- ETL scenarios and best practices 57
 - Aggregate Advisor 59
 - creating and loading data into columnar tables 57
 - loading aggregate tables 59
 - loading atomic tables and dimensions 57
 - tuning after initial data load 46, 60

F

- fix packs
 - applying Cognos Business Intelligence 87
 - applying combined 86
 - applying DB2 90
 - types of
 - Cognos Business Intelligence 83
 - combined 83
 - DB2 83
 - uninstalling 87, 91

- fonts
 - adding 26
 - updating 52

G

- gosldw_sales sample dynamic cube 67
- gosldw_sales_and_target sample dynamic cube 67
- gosldw_target sample dynamic cube 67
- groups
 - securing IBM Cognos Connection 23

H

- hardware requirements 4
- HTTPOnly parameter 25

I

- IBM Cognos Application Firewall 22
- IBM Cognos Connection
 - securing 22, 23, 24
- IBM Cognos Cube Designer
 - deploying sample cube 67
- IBM DB2 with BLU Acceleration
 - tuning after initial data load 46, 60
- IBM Support
 - contacting 98
 - sending and receiving information 99
- IBM Tivoli Directory Server
 - user registry component 12
- images
 - rendering 45
- importation
 - samples 65
- in-database recommendations 59
- installation
 - checklist 7
 - checklists 7
 - steps 27
- installing
 - modeling tools 42
 - translated documentation 43

L

- LDAP
 - enabling SSL 21
 - namespaces 22
- LDAP Directory Server
 - user registry component 12
- Lifecycle Manager
 - connecting to BI Pattern applications configuration 43
- log files
 - troubleshooting 49, 95

M

- maintaining
 - domain lists 51
 - notification mail server 52
- Microsoft Active Directory
 - user registry component 12
- Mobile samples 68

- modeling tools
 - installing 42
- models
 - sample models and packages 73
 - samples 67
- monitoring
 - virtual application instances 54
 - virtual machines 54

N

- namespaces
 - Active Directory 22
 - LDAP 22
 - Tivoli directory server 22
 - user authentication sources 12
- near real-time analysis 61
- notification mail server
 - changes 52

O

- organizational units (OUs) 24
- other components
 - specifications 18

P

- packages
 - samples 67
- pattern
 - components 9
 - creating 30
 - deployed virtual machines 35
 - deployment 27
 - enabling for use 30
 - installation 27
 - post-deployment tasks 41
 - prerequisites 4
 - uploading 28
 - virtual machines 37
- pattern deployment
 - verifying 38
- performance
 - dynamic cubes 60
- performance monitoring 53
 - database monitoring 54
 - dynamic cubes 55
- planning the deployment 19
- plug-ins 83
- post-deployment tasks 41
- post-installation
 - checklist 8
- pre-installation
 - checklist 7
- preparing data 57
- prerequisites 4
- problem determination
 - exchanging information with IBM Support 99
- product
 - overview 1
 - Release Notes 5
 - software environments 5
- production environments
 - moving to 46

R

- recent fact tables 61
- refreshing data cache 63
- Release Notes 5
- reports
 - images 45
 - samples 67
- result set cache 60
- roles
 - securing IBM Cognos Connection 23
- RSS feeds
 - troubleshooting 100

S

- sample databases
 - importing 65
- Sample Outdoors Company 70
 - databases, models, and packages 73
 - samples 69
- samples 65, 67, 69
 - database, models, and packages 73
 - deploying dynamic cubes 67
 - dynamic cubes 66
 - employees 72
 - GO data warehouse 73
 - GO Sales transactional database 73
 - Mobile 68
 - models 66
 - packages 74
 - sales and marketing data 72
 - The Sample Outdoors Company 70
- secure LDAP communication 21
- securing
 - data source signons 25
 - IBM Cognos Connection 22, 23, 24
- security
 - HTTPOnly parameter 25
 - requirements 20
- software environments 5
- software requirements 4
- SSL
 - BI components 20
 - LDAP namespace 21
- stored procedures 59
- supplementary fonts
 - adding 26

T

- Tivoli directory server 22
- troubleshooting 95
 - contacting IBM Support 98
 - exchanging information with IBM Support 99
 - failed deployment 102
 - fixes
 - getting 98
 - getting fixes 98
 - identifying problems 95, 102
 - knowledge bases
 - searching for troubleshooting solutions 97
 - log files 49, 95

- troubleshooting (*continued*)
 - logging in to virtual machines 102
 - queries 49
 - reports 49
 - searching knowledge bases 97
 - subscribing to support 100
 - subscribing to updates 100
 - virtual machines 95
- tuning
 - IBM DB2 with BLU Acceleration settings 46, 60

U

- uninstalling
 - DB2 fix pack 91
- uninstalling fix packs
 - Cognos Business Intelligence 87
- updating
 - binding credentials to LDAP server 50
 - fonts 52
 - list of primary Blackberry push servers 53
 - passwords 50
 - user names 50
- upgradingBusiness Intelligence Pattern with BLU Acceleration 92
- uploading the pattern 28
- user authentication sources
 - user registry components 12
- user registry components
 - authentication sources 12
 - specifications 12

V

- verifying the deployment 38
- virtual application instance
 - Cognos BI 41
 - viewing details 37
- virtual application instances
 - monitoring 54
- virtual application pattern
 - creating 30
- virtual applications
 - deploying 32
 - logs 47
 - maintenance
 - overview 46
 - operations 46
 - post-deployment tasks 41
 - user authentication sources 12
- virtual cubes 63
- virtual machines
 - login errors 102
 - monitoring 54
 - provisioned 35
 - resolving host names 38
 - viewing 37

Z

- ZIP files
 - model.zip 65