2

Getting Started

The first steps in getting started with cloud computing involve data, security, topology, and Linux considerations.

Data Considerations

Databases and data sources can be co-located in the cloud with your cloud application. Or, they can be located on-premises behind your firewall, but with a secure connection to your cloud. Cognos Business Intelligence has two classes of data sources: 1) the content store and metric store database and 2) the query databases and other data sources. Figure 2.1 depicts these aspects of the Cognos 8 tiered architecture.

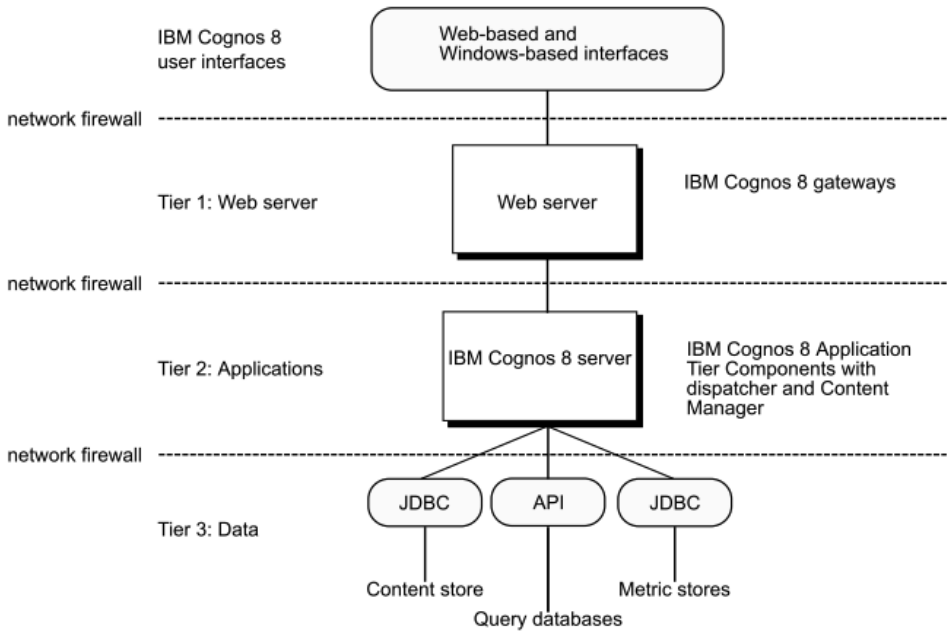

Figure 2.1: The Cognos 8 tiered architecture and data sources.

The *content store* is a relational database that contains data that Cognos BI needs, including report specifications, published models (and the packages that contain them), connection information for data sources, information about your users, and information about scheduling and bursting reports. The *metric store* is the equivalent of the content store for Metric Studio (an optional component of Cognos BI). It contains content for metric packages and other Metric Studio settings, such as user preferences. If you are not using Metric Studio, you do not need a metric store.

The *query databases* are relational databases that can be accessed through Cognos BI. They provide the data for its reports and analyses, through a JDBC or Virtual View Manager connection.

The *data sources* include all relational databases. Other, less common data sources can be accessed through Cognos BI, as well. These are not relational databases; they are things such as dimensional cubes and files.

For best performance, the content store and metric store databases should be as close as possible to the application on the network. Close proximity on the network minimizes latency between the Cognos BI Server components and the content and metric store databases. Ideally, therefore, these databases should be

in the IBM Cloud environment—either in a separate virtual machine instance or with your application tier components in the same instance.

For your query databases, consider the specifics of your intended workload and scenario, as outlined in Table 2.1. Perhaps, for example, your reports require high performance or rapid querying of your data, and this data can easily be moved to the cloud, too. At the same time, privacy or security concerns are not a priority. In this case, you can realize significant cost savings by creating the query databases in the IBM Cloud.

In other situations, high performance or rapid data querying is not a priority, and there is a large amount of data that is difficult or expensive to move. Privacy, security, or other legal reasons may require you to maintain the data within your corporate firewalls. In this case, the query databases should be kept on-premises, within the network bounded by your firewalls. These databases can then be accessed from the IBM Cloud through a secure network connection. This configuration is sometimes referred to as a "hybrid cloud" environment because it is a mix of cloud instances and traditional behind-the-firewall instances.

In some circumstances, your query database is already in the cloud (e.g., Salesforce data). In this case, the security and latency challenges associated with the query data are not new to a cloud solution. Such "cloud-born" query data sources are ideal candidates for leaving in the cloud.

Table 2.1: Considerations for Locating Cognos Query Databases

Query Database in the Cloud	Query Database On-Premises
Workloads require high performance (e.g., rapid queries or large amounts of data).	Workloads have acceptable query performance over a network connection.
New or existing data is easily moved to the cloud (e.g., test data) or is "cloud-born" (e.g., Salesforce data).	A large amount of data exists, or the data is difficult to move to the cloud.
An acceptable level of privacy/security comfort exists around the location of data (e.g., public or non-sensitive data).	Privacy, security, or legal reasons require data to remain on-premises.

Another combination is also worth mentioning. In some situations, database replication can be used to copy an on-premises database instance to a database instance in the cloud. For example, one of IBM DB2®'s various replication alternatives might be an attractive option for you, provided you leverage a secured connection for the transaction.

File-based data sources, such as dimensional cubes and other files, are usually amenable to synchronization or transport to the cloud in a directory that is

accessible to the Cognos application instance. They can also be synchronized to an IBM Cloud storage instance that appears as a mounted directory to your instance.

Security Provider Considerations

Many Cognos applications authenticate users through third-party security tools such as LDAP or Active Directory. Such third-party authentication sources are typically used to create groups of users and to restrict content access to certain user groups. If your workload includes such a requirement, your cloud topology will also need to include your authentication source, as Figure 2.2 illustrates.

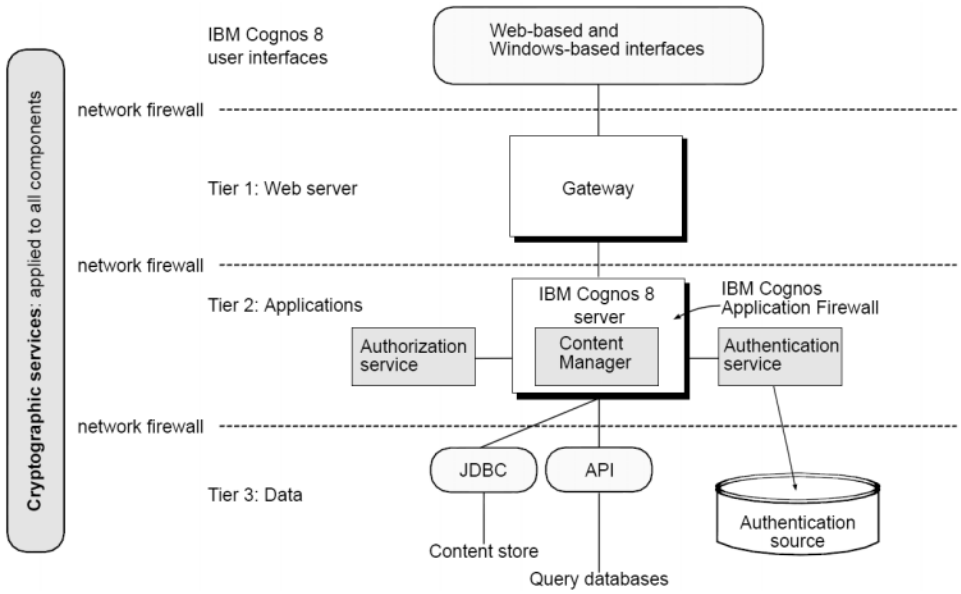

Figure 2.2: The Cognos tiered architecture with authentication.

In addition to data considerations, you will need to consider the location of your authentication source in this scenario (Table 2.2). You can co-locate your authentication source in the cloud, as you did with your Cognos application, or you can keep it on-premises behind your firewall with a secure connection to the cloud.

Table 2.2: Determining the Placement of the Authentication Source	
Security Provider in the Cloud	**Security Provider On-Premises**
Workloads require the best network latency or highest performance.	Workloads have acceptable query performance over a network connection.
User data is new or easily moved to the cloud (e.g., test-user data).	The user directory is too large or difficult to move to the cloud (as in the case of password synchronization).
An acceptable level of privacy/security comfort exists around the location of user data.	Privacy, security, or legal reasons require you to maintain user information on-premises.

Perhaps you're providing IT resources to a new or small set of users, such as a development or test/QC group. If the user data can be easily moved to the cloud (because it is small or static enough to be managed via SFTP or database replication) and there are no security or privacy concerns, then installing the authentication source in the cloud alongside or within your application cloud instance is recommended. On the other hand, a secured network connection between the on-premises authentication tool and the Cognos application can be used if any of the following are true:

- The authentication source needs to leverage an existing on-premises installation.
- The data cannot be easily replicated or moved.
- Privacy, security, or performance concerns exist.

Designing and Testing Your Topology

As you design and refine your topology, start simply and avoid unnecessary complexity. Satisfy your requirements, but always keep the number of cloud instances in your topology as low as possible. Adding instances is easy in the event that you encounter a future need based on additional requirements or load, so it's usually preferable to underestimate your requirements initially.

Starting simply also applies when it comes to the number of unique cloud images. For example, it is easier to manage a single DB2 database image that customizes itself on startup than to create five different query database images. Or, rather than have a Cognos content store image and a Cognos report server image, it might be simpler to have a generic Cognos BI image, with which you start only the services required for the instance upon startup.

The process of designing and then testing your topology is very iterative, likely consisting of these steps:

1. Design/refine your topology.
2. Create/customize the required instances.
3. Install/configure the instances.
4. Save image snapshots.
5. Test for functionality and performance.
6. Repeat.

If you start simply and keep your topology as small as possible, each of these steps becomes as straightforward as possible. We offer other tips in the pages that follow to help reduce the number of unique images required for your solution.

Embracing Linux®

The IBM Cloud gives you the flexibility to choose from several different operating systems on your instances, including Microsoft Windows and several versions of Linux. There are advantages to opting for one of the Linux-based distributions:

- The underlying infrastructure of the IBM Cloud is based on Linux technologies, including virtualization based on Xen.
- The licensing costs associated with Linux are lower than for Windows-based operating systems.
- Many cloud-appropriate applications and technologies—and the associated user communities and knowledge bases—are already focused on Linux environments.

Other things being equal, embracing Linux can thus be a very efficient and cost-effective part of your cloud strategy. Within your cloud design, deployment, and operations teams, there is also probably a rich storehouse of Linux experience and skills, which you can leverage to ensure the success of your project.

3

Installation and Configuration

This chapter describes the steps required to install and configure IBM Cognos Business Intelligence, either Version 8 or Version 10, into a single IBM Smart Business Development and Test Cloud image. When finished, you will have created a private image on the IBM Cloud. You can use this image to create and deploy instances of a fully configured and operational Cognos 8 or Cognos 10 instance, which can immediately be used for your analytics applications.

The techniques and information used to create this single image of Cognos 8 or Cognos 10 will form the basis for extending out to multiple image topologies. In addition, while the steps assume IBM DB2 as the database, IBM WebSphere® Application Server as the application server, and IBM HTTP Server as the Web server, the information should be useful to whatever configuration you are ultimately aiming for.

Set Up the Windows Client

While the server instances will be in the IBM Cloud, several tools allow easy access and management of the cloud instances from a PC running Microsoft Windows. Installing these tools is a one-time step; after that, you can use them to connect to any instance created in the IBM Cloud.

Install SSH Client (PuTTY)

You need a Secure Shell (SSH) client to log into your newly created instance and interact with it. If you are using a Windows machine, we recommend the PuTTY freeware SSH client. (The latest PuTTY version as of June 2010 is 0.60.)

Install WinSCP

You also need an SSH file transfer program to move software and data to your cloud instances. For Windows machines, we recommend WinSCP freeware. (The latest WinSCP version as of August 2010 is 4.2.8.)

Install X-Windows

To run X-Windows programs from the cloud, such as the Cognos Configuration tool, and have them appear on your machine, you will need an X-Windows client. On Windows machines, you can choose to install a commercial product, such as OpenText's Exceed, or opt to use one of several freeware clients, such as Xming.

If you are using Xming, follow these steps:

1. Download and install Xming, following the instructions available from the download site. (The latest Xming version as of August 2010 is 7.5.0.24.)
2. Launch Xming, and confirm that the Xming icon appears in the system tray.

Set Up and Configure the Cloud Instance

This section summarizes the steps involved to create your machine instance in the cloud. More detailed documentation is available online on the IBM Smart Business Development and Test Cloud documentation areas and asset catalog.

Create the Cloud Instance

To create the cloud instance, take these steps:

1. Browse to *http://www.ibm.com/cloud/enterprise*.
2. Register for an account, if required.
3. Sign in with your user ID and password.
4. In the Control Panel, click the **Add Instance** button.

5. In the "Add Instance" image selection that follows, scroll down and select "Red Hat Enterprise Linux 5.4 (32-bit)" to choose a Red Hat Enterprise Linux 5.4 base operating system instance. (For a 64-bit install, select "Red Hat Enterprise Linux 5.4 64-bit.") Click **Next**.
6. Specify your desired instance name and size type (Medium or Large), and then click **Add Key**.
7. Specify the name of the key to generate, and click **Continue**.
8. Ensure you download the private key file and save it. Do this by clicking **Save** or by creating the key manually. To create the key manually:
   a. Choose to open the file during download. Your key will appear in a separate Web page, with a **BEGIN RSA PRIVATE KEY** line at the top of the file and an **END RSA PRIVATE KEY** line at the bottom.
   b. Copy and paste the page contents (including the **BEGIN** and **END** lines) into a new file using a text editor.
9. Save the file with a name such as **ibmcloud_your.name@ibm.com_rsa**. Select your new key.
10. Verify overall selections for the new instance, and click **Next**.
11. Read and agree to the terms and conditions. Click **Submit**.
12. The Add Instance dialog's "Submitting" dialog appears. When the Add Instance dialog reports "Success," click **Close**.
13. The Control Panel should now list the new instance with a status of "Active," as shown in Figure 3.1.

Figure 3.1: The Control Panel showing an active instance.

Modify Security Permissions

In a single-image environment, you want to ensure that the permissions are correctly set up so that the **idcuser** account is able to install and run all the

software. We describe this process below. Keep in mind that there will be accounts with least privileges accessing the instance. Because this is the case, consider revoking root privileges after installation, or create a separate user account just to run Cognos 8 or Cognos 10.

1. Change the group for **idcuser** to **root** by setting the value of the fourth field to 0 (zero):

```
$ sudo vi /etc/passwd
idcuser:x:500:0:idcuser:/home/idcuser:/bin/bash
```

2. Prepare the target root installation folder and set group-level write permissions:

```
$ sudo mkdir -p /opt/ibm
$ sudo chmod -Rf g+w ibm
```

This step is important for the Cognos installation program to succeed.

Enable X11 Forwarding on Your Cloud Instance

X11 Forwarding is the part of X-Windows that permits messages from the cloud instance's Windows manager to be "forwarded" to your Windows client. By default, your instance should be set up for X11 Forwarding, but we will verify that here.

1. As a root user, ensure that the ports required for X11 forwarding are enabled:

```
$ sudo /sbin/iptables -L
Chain INPUT (policy DROP)
target    prot opt source       destination
ACCEPT    all  --  anywhere     anywhere     state RELATED,ESTABLISHED
ACCEPT    tcp  --  anywhere     anywhere     tcp dpt:ssh
Chain FORWARD (policy DROP)
target    prot opt source       destination
Chain OUTPUT (policy ACCEPT)
target    prot opt source       destination
```

2. Confirm that all policy values are "ACCEPT." If this is not the case (as shown above), you will need to update the values accordingly in the main **iptables** configuration file, **/etc/sysconf/iptables**, and then restart the **iptables** service:

```
$ sudo vi /etc/sysconfig/iptables
```

The file should look something like this:

```
# Generated by iptables-save v1.3.5 on Thu Mar  4 22:07:58 2010
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p tcp -m tcp --dport 22 -j ACCEPT
COMMIT
# Completed on Thu Mar  4 22:07:58 2010
```

3. Restart the service:

```
$ sudo /sbin/service iptables stop
$ sudo /sbin/service iptables start
```

Install X11 Support in Your Cloud Instance

In addition to the X-Windows client, you'll need to install Motif on the cloud instance to run X-Windows programs such as the Cognos Configuration tool on the cloud.

1. Create a temporary directory, such as **/home/idcuser/installs/centos**.
2. Add the CentOS repository to your system by creating the file **Centos.repo** in that directory, or create it on your local machine and use WinSCP to transfer the file to that directory (this step might not be required if your RedHat image is activated and connected to the RedHat repositories; in that case, skip to step 4):

```
[base]
name=CentOS-$releasever - Base
mirrorlist=http://mirrorlist.centos.org/?release=5.5&arch=x86_64&re
  po=os
#baseurl=http://mirror.centos.org/centos/$releasever/os/$basearch/
gpgcheck=1
gpgkey=http://mirror.centos.org/centos/RPM-GPG-KEY-CentOS-5
protect=1
[updates]
name=CentOS-$releasever - Updates
mirrorlist=http://mirrorlist.centos.org/?release=5.5&arch=x86_64&re
  po=updates
#baseurl=http://mirror.centos.org/centos/$releasever/
  updates/$basearch/
gpgcheck=1
gpgkey=http://mirror.centos.org/centos/RPM-GPG-KEY-CentOS-5
protect=1
```

3. Update your system to recognize the CentOS repository:

```
$ sudo cp /home/idcuser/installs/centos/Centos.repo /etc/yum.
  repos.d/
```

4. Download OpenMotif and associated X-Windows libraries:

```
$ sudo yum install openmotif.i386 xorg-x11-xauth libXtst xterm
$ sudo ln -s /usr/lib/libXm.so.4 /usr/lib/libXm.so.3
```

5. Edit **/etc/ssh/sshd_config** and ensure that the following lines are uncommented:

```
X11Forwarding yes
X11UseLocalHost yes
AllowTcpForwarding yes
```

6. If changes to **sshd_config** were required, restart the **sshd** service:

```
$ sudo /sbin/service sshd restart
```

Configure the Windows Client

Now that you have created the cloud instance, follow the steps described below to connect the Windows client to it. You can repeat these steps for all the IBM Cloud instances that are available to you.

Configure PuTTY

For each cloud instance, create a PuTTY session that lets you connect to it:

1. PuTTY requires a "ppk" version of the private key previously downloaded from the IBM Cloud. This step is required only once for each keyfile you have from the IBM Cloud.
   a. Launch **puttygen.exe**.
   b. Click **Load**, and choose the keyfile (e.g., **ibmcloud_your.name@ibm.com_rsa**). Click **OK** when you see the dialog box shown in Figure 3.2.

Figure 3.2: A successful conversion of the private key using PuTTYgen.

   c. Click "Save private key," optionally providing a passphrase and specifying a **.ppk** extension (e.g., **ibm_cloud_your.name@ibm.com.ppk**).

2. Launch PuTTY, and specify these settings:
   a. In the Session section, specify the IP address of your newly created instance.
   b. In the Data section, specify **idcuser** as the username.
   c. In the Auth section, select "Allow attempted changes of user name," and specify the path to the private keyfile you saved in Step 2 (**ibm_cloud_your.name@ibm.com.ppk**). Figure 3.3 shows these settings.