VMware vFabric Data Director Administrator and User Guide

vFabric Data Director 1.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-000709-01

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About VMware vFabric Data Director

The *VMware vFabric Data Director Administrator and User Guide* provides information about administering VMware® vFabric Data Director. Administration tasks include creating organizations and database groups, managing users and roles, resource allocation, database and backup configuration, cloning databases, safeguarding data, and monitoring your system.

The Data Director software solution provides on-site self-service database provisioning and automation to database administrators and application developers, including the following.

- Self-service database creation and resource allocation.
- Flexible, policy-based resource management.
- Resource isolation within organizations and within databases.
- Security policy implementation through role-based access control.
- Delegating and granting customizable roles and privileges to specify users' allowed actions.

Self-service database lifecycle management enables application developers to create new databases, manage schemas, configure backups, perform restores, clone databases for testing and development, scale up database sizes, and decommission databases. Customizable database configuration and backup templates enable administrators to control database parameters and enforce resource allocation policies, and provide application developers with simplified database creation and resource allocation.

Intended Audience

This document is for administrators and application developers.

- System administrators use this document to learn about managing and monitoring a Data Director environment. System administrators create organizations, allocate resources to them, and perform other high-level tasks.
- Organization administrators use this document to learn about managing and monitoring database groups and databases. Organization administrators can use and customize database templates, can assign resources, and can monitor their organization.
- Application developers use this document to learn about creating, managing and monitoring databases.

Updated Information

This *vFabric Data Director Administrator and User Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *vFabric Data Director Administrator and User Guide*.

Revision	Description
EN-000709-01	■ The topic Chapter 1, "VMware vFabric Data Director and vFabric Postgres Overview," on page 11 clarifies information about user management modes. ■ The topic Chapter 8, "Cloning Databases," on page 57 clarifies clone point information. ■ The topic Chapter 12, "Managing Licenses," on page 103 clarifies information about license types and database usage types. ■ Minor revisions.
EN-000709-00	Initial release.

VMware vFabric Data Director and vFabric Postgres Overview 1

The VMware vFabric Data Director and vFabric Postgres software solutions enable you to provide on-site self-service database provisioning and automation to database administrators and application developers.

This chapter includes the following topics:

VMware vFabric Data Director System Architecture

The Data Director architecture provides database as a service (DBaaS) to application developers with security and resource isolation as well as flexible, policy-based resource management and role-based access control for system administrators. Data Director is optimized for VMware vSphere.

At the system level, Data Director supports flexible, policy-based resource management and provides resource isolation between organizations and databases. As a Data Director system administrator, you can implement security policies through role-based access control, controlling users' allowed actions with customizable roles and privileges that you delegate and grant as required.

Within organizations, Data Director offers self-service database lifecycle management for VMware vFabric Postgres (vPostgres) databases. You control database parameters with customizable database configuration and database backup templates. These templates simplify database creation and provisioning for application developers. Developers can create databases and allocate resources for them, manage schemas, set up backups and perform restores, clone databases for testing and development, scale database sizes up, and decommission databases without assistance.

The vPostgres database is based on the open source Postgres database, an ACID-compliant, ANSI-SQL-compliant transactional relational database. The database optimized for vSphere and is compatible with Postgres client tools and drivers.

VMware vFabric Data Director Components

VMware vFabric Data Director consists of the database lifecycle management platform and vFabric Postgres (vPostgres).

The Data Director hierarchy has the following levels.

- System (the database lifecycle management platform)

- Organizations
- Database groups
- Databases (vPostgres databases)

System administrators perform management tasks at the system level, which is the top level of the hierarchy. To edit system-level settings you must have system privileges, but having system privileges does not allow you to make changes to the other levels.

A system can contain multiple organizations. An organization can contain multiple database groups. A database group can contain multiple databases. You cannot create database groups at the system level. They can exist only within organizations. Databases can exist only within database groups.

The following figure shows the Data Director system hierarchy.

Figure 1-1. Data Director System Hierarchy

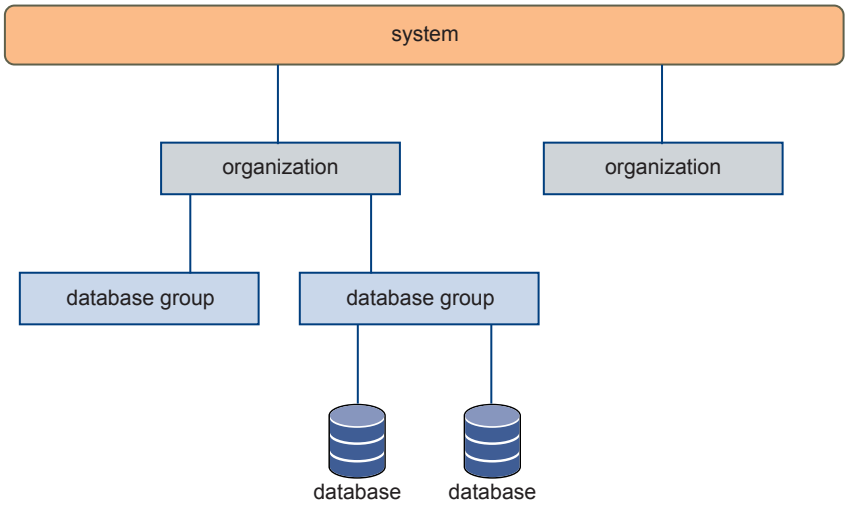

System administrators manage Data Director resources at the system, organization, and database group levels. System administrators create resource bundles from vSphere resource pools (CPU and memory resources) and networking and storage resources, and allocate one or more resource bundles to each organization. Organization administrators assign resources from their resource bundles to database groups for consumption by databases.

Data Director User Management Modes

Data Director user management modes control how users are assigned and managed among different organizations. Data Director has two user management modes: Global mode (for enterprises) and By Organization mode (for service providers). Global user management mode is the default.

By Organization user management mode has the following characteristics.

- Organizations are set up as separate, isolated enterprises with no visibility into other organizations.
- The Data Director system user list is not visible to organizations.
- No organization can see another organization's user list.
- Organization administrators send email to invite users to join their organization, or users can navigate to the Data Director application URL and click an email link to request access to an organization.

Global user management mode has the following characteristics.

- Organizations are set up as separate departments, business units, or groups within one enterprise, such as a corporation's HR and Finance departments.

- All Data Director users are visible to all organizations within Data Director.
- Organization administrators contact users to invite them to the organization or grant access directly from the system user list.

You configure the Data Director user management modes during installation. User management mode cannot be changed. In both Global and By Organization user management modes, organization administrators must grant users access to their organization.

About Data Director Administration

Data Director system administrators perform Data Director administration on the system level. Organization administrators perform Data Director administration on the organization level.

You create the initial Data Director system administrator account during Data Director setup. That system administrator creates other Data Director users, including other system administrators and organization administrators, and performs administration tasks at the system level.

By default, users do not have roles or permissions and cannot access any organizations. Organization administrators assign roles and permissions to users and grant them access to specific organizations.

System administrators perform system-level operations for Data Director or for an entire organization. System administrators perform the following tasks.

Table 1-1. System-level Operations

Operation Type	Examples
Resource management operations	■ Creating resource bundles. ■ Assigning resource bundles to organizations.
User and organization management operations	■ Creating users. ■ Creating organizations. ■ Creating organization administrators. ■ Designating existing users as organization administrators.

Organization administrators perform organization-level operations within their organizations. Organization administrators perform the following tasks.

Table 1-2. Organization-level Operations

Operation Type	Examples
Resource management operations	■ Creating database groups. ■ Creating database configuration templates. ■ Creating database backup templates. ■ Allocating resources to database groups within the organization.
User management operations	■ Creating organization users. ■ Granting organization access to existing Data Director users. ■ Assigning organization roles to users in the organization. ■ Creating organization roles and granting roles to organization user. ■ Defining organization permissions and granting permissions to organization users.

By default, Data Director system administrators do not have access to organizations. Organization administrators have access only to their own organization, can create organization users, and can grant access to existing Data Director users.

Data Director system administrators can create users, but only organization administrators can grant those users access to organizations.

vFabric Postgres Databases

Data Director provides self-service database provisioning and automation with vFabric Postgres (vPostgres). vPostgres is built on the open source Postgres database. It is compatible with pSQL and the PostgreSQL tools and client drivers. vPostgres databases are fully ACID and ANSI SQL-compliant. The ACID properties (Atomicity, Consistency, Isolation, and Durability) guarantee that database transactions are processed reliably.

Database administrators and application developers administer databases within their organizations. Database administration includes the following tasks.

- Creating databases and allocating resources to them.
- Cloning databases.
- Managing database users, roles, privileges, and permissions.
- Maintenance, including backups, restores, and removing old and unused data.
- Scaling databases up.
- Monitoring database usage and performance.
- Monitoring database alarms.
- Decommissioning databases.

See the *vFabric Postgres Standard Edition User Guide* for information about the Postgres database features for Data Director.

Managing Data Director Resources 2

System administrators manage CPU, memory, storage, and networking resources for different organizations. Organization administrators manage resources for database groups and for databases.

This chapter includes the following topics:

Resource Management Overview

System administrators allocate resources to organizations. These virtual resources come directly from the physical resources of the cluster on which Data Director runs. Organization administrators assign organization resources to database groups and databases.

A vSphere cluster consists of several ESXi hosts that provide the physical CPU and memory resources for the databases managed by Data Director. As part of installation, you create the cluster and enable vSphere High Availability (HA) and vSphere Distributed Resource Management (DRS) for the cluster. Data Director can take advantage of the vSphere HA and vSphere DRS functionality because Data Director runs on top of the cluster. See the *vSphere Availability* and the *vSphere Resource Management* documentation for details.

A Data Director resource bundle includes CPU, memory, storage, and networking resources. The CPU and memory resources come from a resource pool in the vSphere cluster. The storage and networking resources are assigned to Data Director during installation or at a later time. Data Director includes a set of VLANs to carry different types of network traffic.

When system administrators create an organization, they can assign virtual resources to the organization as resource bundles. When organization administrators create a database group, they assign virtual resources to the database group. These virtual resources are backed by the physical resources of one or more clusters. vSphere clusters provide failover protection and support efficient use of physical resources.

System administrators can assign resources when they create an organization (see "Create an Organization," on page 38) or assign resources to an existing organization (see "Assign a Resource Bundle to an Organization," on page 22). Organization administrators can assign resources when they create a database group or assign resources to existing database groups.

To help you specify the resources associated with a database template, Data Director includes a calculator that computes the optimum resource configuration based on the anticipated usage patterns. When you create databases from the template, the specified resources are allocated.

Resource Bundles and Resource Pools

A resource bundle is a set of compatible IT resources for provisioning databases. A resource bundle includes CPU and memory resources as vSphere resource pools, and storage and networking resources.

To assign the appropriate amount of resources to each organization, system administrators create resource bundles and assign them to organizations. System administrators specify a resource pool and storage and networking resources when they create a resource bundle.

Resource Pool

All CPU and memory resources of a resource bundle come from a vSphere resource pool that is created in the vSphere Client with reservation equal to limit. See "Create a Resource Pool," on page 20.

Storage Resources

Each resource bundle includes storage resources for data and storage resources for backup. The storage resources must be visible to all hosts that use the resource bundle.

DB Access Networks

DB Access Networks provide communication for databases. A DB Access Network corresponds to a vSphere port group. Each network must be visible to all hosts that use the resource bundle. DHCP is required.

Selecting one or more DB Access Networks allows you to isolate different database groups from one another, for example, to isolate a QA database group from a Production database group. When no DB Access Networks have been assigned in the environment, select the network that is mapped to the Web Console Network. Do not select internal networks for DB Access Network traffic.

The following figure shows how Data Director resources come from vSphere resource pools, datastores, and port groups. When administrators create a resource bundle, the resources are always coming from the underlying vSphere environment.

Figure 2-1. Resources in vSphere and Data Director

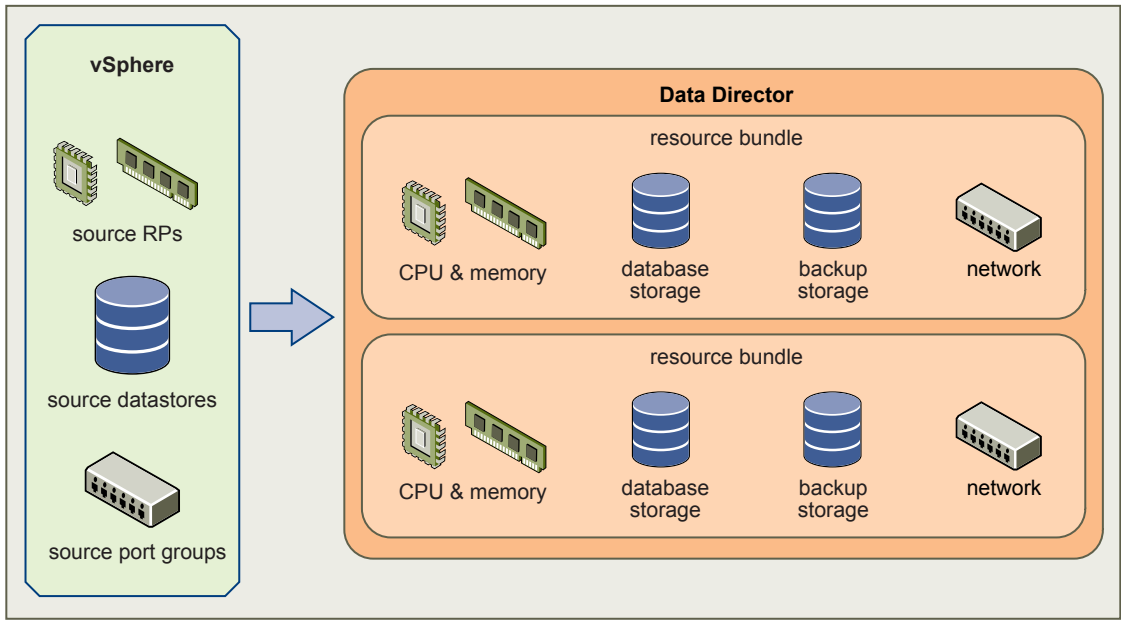

"Resource Assignment," on page 17 explains how resource assignment differs for the different levels of the hierarchy.

Resource Assignment

Resource assignment differs for organizations, database groups, and databases.

Resource Assignment for Organizations

System administrators can assign multiple resource bundles to each organization. Organization administrators allocate the resource bundles to database groups. When databases are created, they can only draw on the resources assigned to the database group. This resource isolation guarantees that different organizations and different database groups have control over their resources.

Resource Assignment for Database Groups

When you create a database group, you assign a resource bundle that specifies the resources for that group. You cannot assign more than one resource bundle to one database group. Multiple database groups can share one resource bundle.

When you assign a resource bundle to a database group, you can specify how to allocate each resource.

- CPU priority or reservation.
- Memory priority or reservation.
- Storage allocation.
- A network for the database group. You cannot divide the network. You can select only one network during database group creation even if several networks are associated with the resource bundle.

If you do not explicitly specify the CPU or memory allocation, Data Director sets the reservation to zero but sets expandable reservations to true. If expandable reservations is set to true, the CPU or memory can expand beyond the specified value.

Resource Assignment for Databases

A database consumes the resources assigned to its database group.

- You can specify the number of virtual CPUs, the memory size, and CPU and memory priority for each database that you create.
- You cannot specify storage allocation. All databases consume the data and the backup storage allocated to their parent database group.
- Each database uses the network assigned to the database group.

vSphere Resource Pools and Data Director

A vSphere resource pool is a logical abstraction for flexible management of CPU and memory resources. You add CPU and memory resources to Data Director resource bundles by adding a vSphere resource pool to the bundle.

CAUTION Data Director can only use resource pools if the corresponding cluster is enabled for DRS and HA. Do not disable DRS. If you do, Data Director can no longer use the resource pools even if you reenable DRS. See "Resource Bundles Become Unusable Because DRS Is Disabled," on page 116.

Resource pools allow you to group available CPU and memory resources. You can allocate resources explicitly, or use the resource pool share mechanism. You can hierarchically partition available CPU and memory resources by grouping resource pools into hierarchies. You can then allow different organization access to different resource pools. For example, a QA department might need large amounts of CPU and memory for running tests while the marketing department might require smaller amounts.

Data Director expects you to group the hosts that provide the CPU and memory resources into clusters. Each cluster owns the resources of all hosts. You can create one or more resource pools for the cluster, which has an invisible root resource pool. Each resource pool owns some of the cluster's resources. If necessary, you can create child resource pools. Child resource pools represent successively smaller amounts of CPU and memory.

How you allocate CPU and memory resources to database groups differs from how you allocate those resources to databases.

Creating Resource Pools

You create resource pools by using a vSphere Client connected to a vCenter Server system. Specify the following resource pool settings to ensure that Data Director always receives all of its allocated resources and does not have different amounts of CPU and memory available if the cluster is experiencing a light or a heavy load.

NOTE If you do not configure your resource pool with these settings, problems with resource bundle creation or other Data Director tasks might result. The primary problem is that resource pools with incorrect settins do not appear in the list of available resource pools when you create a resource bundle.

Set the Limit equal to the Reservation.	If the system never allocates more resources than you reserved, you do not experience resource fluctuations.
Set Expandable Reservation to checked or unchecked.	If the system does not attempt to allocate more resources than you reserved, you do not experience resource fluctuations.
Set Unlimited to unchecked.	Data Director requires this setting to avoid that a resource bundle takes more than its share.

After you create the resource pool, you create resource bundles. Each resource bundle uses one resource pool.

See "Create a Resource Pool," on page 20 and "Create a Resource Bundle," on page 21.

Allocating CPU and Memory Resources to Database Groups

When you create a database group and set its CPU and memory allocation, Data Director creates a child resource pool of the resource pool you select. Data Director configures the resource pool with the allocation you specify. Having a different resource pool for each database group isolates the database group's allocation and makes different groups independent.

- If you specify the CPU and memory allocation, Data Director uses the following settings for the resource pool it creates.
  - Reservation is set to the value you specify.
  - Expandable reservation is set to False.
  - Limit is set to unlimited.
- If you do not specify CPU or memory allocation, Data Director uses the following settings for the resource pool it creates.
  - Reservation is set to 0.
  - Expandable reservation is set to True, allowing the database group to consume resources as they are available.
  - Limit is set to unlimited.

Allocating CPU and Memory Resources to Databases

In the Data Director environment, a database is a virtual machine that consumes resources from the database group. You can specify the CPU and memory allocation for the database. Data Director always sets the limit to unlimited.

Viewing Resource Information

Data Director system administrators can view resource usage information for an organization from the Data Director **Manage & Monitor** tab.

When you log in to Data Director as a system administrator, you can view information about the resource usage of the different database groups and about the resource bundle or resource bundles that are being used by each database group.

- The **Organizations** pane allows you to manage organizations. You can view organization information, assign and unassign resource bundles, delete the organization, and view the organization's properties.
- The **Resource Bundles** pane allows you to view all resource bundles currently created for this instance of Data Director. You can display either allocation information or vCenter Server Object information.
  - You can click on an item in the heading, such as **Organization**, to re-sort the table based on that column. Right-click any resource bundle name and choose **Properties** to see detailed information about each resource bundle.
  - If you select **vCenter Server Objects**, Data Director displays the names of resource pools, datastores, and networks that you see in the vSphere Client UI.
- The **Datastore Usage** pane shows datastore usage for the main datastore and the backup datastore. You can see how resource bundles map to datastores and examine storage allocation information for each datastore.

See Chapter 11, "Monitoring the Data Director Environment," on page 89 for details on using the monitoring interface.

Monitor Resource Usage

System administrators can view usage information for resource bundles and datastores and can reassign resource bundles from the **Manage & Monitor** tab.

The focus of this task is on monitoring, not on changing current settings.

Prerequisites

- Log in to Data Director as a user with system administrator privileges.
- Verify that one or more organizations exist in your environment.
- Verify that resource bundles and datastores have been assigned to the organizations.

Procedure

1 In Data Director, click the **System** tab, and click the **Manage & Monitor** tab.

The Organizations panel displays resource allocation information about each organization.

2 Click one of the columns, for example **Total Memory**, to reorder the rows of the table.

3 Click one of the organizations to display resource bundle information for the selected organization.

4 Click **Resource Bundles** to display the **Resource Bundles** pane.

5 Click **Datastore Usage** to display information about available datastores, their capacity, and the allocated and unallocated storage for each.

6 Click one of the datastores to display the associated resource bundles and their storage allocation.

What to do next

You can change the resource bundle information by clicking the **Actions** icon and selecting **Properties**. If properties are dimmed, you do not have permissions to change them.

Create a Resource Pool

You allocate CPU and memory resources to Data Director by creating one or more resource pools from a vSphere Client connected to a vCenter Server system. From the Data Director user interface, you can then assign the resources from those resource pools to database groups and databases.

Before you create the resource pools, you must prepare a cluster. Enable the cluster for HA and DRS, and add all Data Director hosts to the cluster. See the *vFabric Data Director Installation Guide* for information.

Prerequisites

- Connect to the vCenter Server system by using a vSphere Client. You cannot create resource pools if the client is connected directly to a host.
- Verify that you have permissions sufficient to create a resource pool.
- Choose a location for the resource pool. Data Director cannot use resource pools at the vApp top level.
- See the *vSphere Resource Management* documentation for information about resource pools.

Procedure

1 In the vSphere Client, select **Home > Inventory > Hosts and Clusters**.

2 Select the cluster to which all Data Director hosts have been assigned.

3 Specify the settings in the following table for the resource pool.

Option	Description
Name	Name of the resource pool.
CPU Shares	Do not specify CPU shares. Instead, specify the CPU reservation.
CPU Reservation	CPU resources to allocate to this resource pool.
Expandable Reservation	Checked or unchecked.
CPU Limit	Maximum CPU resources available to this resource pool. Set Limit to be equal to CPU Reservation.
Unlimited	Unchecked.
Memory Shares	Do not specify memory shares. Instead, specify a memory reservation.
Memory Reservation	Memory resources to allocate to this resource pool.
Expandable Reservation	Checked or Unchecked.
Memory Limit	Maximum memory resources available to this resource pool. Set Limit to be equal to Memory Reservation.
Unlimited	Unchecked.

After the resource pool is set up, you can point to the resource pool when you create the Data Director resource bundle.

What to do next

Create a resource bundle. See "Create a Resource Bundle," on page 21.

Create a Resource Bundle

Resource bundles allow you to bundle CPU, memory, storage, and networking resources. You create resource bundles using the Data Director user interface.

When you create a resource bundle, the wizard displays only resource pools with a parent cluster that meets the following requirements.

- vSphere DRS and vSphere HA are enabled.
- VM Monitoring is set to **VM and Application Monitoring**.
- VM Restart Priority is not disabled for any of the virtual machines.
- Host monitoring is enabled.

See "Perform Advanced Cluster Configuration," on page 22 for details on recommended settings.

Prerequisites

- Create a resource pool to use for allocating CPU and memory resources. See "Create a Resource Pool," on page 20.
- Decide on the storage resources that you want to include in the resource bundle. Plan for storage resources for database storage and resources for backup storage.
- Decide on the networking resources that you want to include in the resource bundle. The resource bundle's networking resources are used for the public network for databases in an organization.

NOTE If you do not configure your resource pool with these settings, problems with resource bundle creation or other Data Director tasks might result.

Procedure

1 Log in to Data Director with system administrator privileges.

2 Select **System**, and click **Manage & Monitor**.

3 Click **Resource Bundles** in the left pane.

4 Click the plus (+) icon.

5 Specify the following information about the resource bundle in the wizard.

Wizard screen	Action
Name and Description	Type a name and optional description and click **Next**.
CPU and Memory	Select the resource pool from which you want to assign CPU and memory resources and click **Next**.
Storage	Click **Edit** to select a datastore, and allocate the number of GB to use with Data Director, or allocate all unallocated space. Repeat the process for backup storage. NOTE Do not select a datastore that is in a datastore cluster.
Networks	Select the networks that you want to have available to this resource bundle. These networks provide the public network for the organization's databases. Resource bundles must use a database network when available.

What to do next

System administrators can allocate the resource bundles to organizations, and organization administrators can assign resources to their database groups.

Assign a Resource Bundle to an Organization

System administrators can assign a resource bundle to an organization when they create an organization. You can also assign a resource bundle to an organization at a later time.

Prerequisites

Log in to Data Director as a system administrator or a user who can assign resource bundles to organizations.

Procedure

1 Click the **Manage & Monitor** tab and, click **Organizations**.

2 Right-click the organization that you want to assign a resource bundle to, and select **Assign Resource Bundle**.

3 Select the resource bundle that you want to assign from the list of resource bundles and click **OK**.

What to do next

You can create one or more database groups and databases. See "Create a Database," on page 53 and "Create a Database Group," on page 43.

Perform Advanced Cluster Configuration

During installation, you configure the Data Director cluster with vSphere DRS and vSphere HA enabled, and with certain monitoring settings. You can later edit the Data Director cluster configuration to change the monitoring sensitivity for virtual machines.

As part of the installation process, you configure the Data Director cluster. See the *vFabric Data Director Installation Guide*. After installation, you can customize the cluster to work in your environment. See the *vSphere Availability* documentation and the *vSphere Resource Management* documentation for background information.

Not all changes that you can make to a vSphere cluster are compatible with Data Director. You must make sure that the cluster settings remain compatible with Data Director. Data Director checks the following settings.

- DRS must be enabled. DRS automation level can be any of the supported options. Partially automated works best with Data Director in most situations.
- Admission control must be enabled.

If cluster settings are not compatible with Data Director, and if you create a resource pool in the cluster, you cannot import the resource pool into a Data Director resource bundle.

If you change cluster settings from Data Director compatible to Data director incompatible, Data Director displays alerts but does not revert the settings. You must revert the settings to make the cluster compatible again.

CAUTION Do not disable DRS because you lose all resource pools. Reenabling DRS does not resolve the issue. See "Resource Bundles Become Unusable Because DRS Is Disabled," on page 116.

If you customize the HA settings for a virtual machine, and if those settings are not compatible with Data Director, an alert appears. You are responsible to make the cluster compatible again.

Prerequisites

Verify that you have log-in privileges and privileges for cluster modification for the vCenter Server system on which the Data Director cluster runs.

Procedure

1 Log in to a vSphere Client that is connected to the vCenter Server on which the Data Director cluster runs.

2 Right-click the cluster and click **Edit Settings**.

3 Click **VM Monitoring**.

4 Select the **Custom** check box and specify custom settings.

The following are the lowest acceptable settings, values can be higher.

Option	Description
Failure interval	30 seconds
Minimum uptime	120 seconds
Maximum Per-VM resets	3
Maximum resets time window	Within 1 hour

5 Click **OK**.

Managing Users and Roles 3

User management controls the users that can log in to Data Director and what they can see and do after they log in.

This chapter includes the following topics:

User Management Overview

System and organization administrators use a combination of user logins, privileges, permissions, and roles (role-based access control) to manage Data Director users. Role-based access control provides management of users and the tasks that they can perform on objects. You can grant and revoke roles and permissions at the system level, on organizations, and on database groups, databases, and templates within organizations.

Roles are sets of permissions required to perform particular jobs. Jobs are sets of tasks that a user with a particular role is responsible for performing, such as the set of tasks that are the responsibility of a database administrator. System and organization administrators define roles as part of defining security policies, and grant the roles to users. To change the permissions and tasks associated with a particular job, the system or organization administrator updates the role settings. The updated settings take effect for all users associated with the role.

- To add a user to a job, the system or organization administrator grants the role to the user.
- To remove a user from a job, the system or organization administrator revokes the role from the user. Changes are effective immediately.

Roles apply only to the organization in which they are created. For example, an organization administrator creates a database administrator role that includes permission to add and remove database users, start and stop databases, and perform backups for a specific database in that organization. Users that are granted the database administrator role in that organization can perform database administrator tasks only within that organization.

Organization administrators usually manage role and permission assignments for their organizations. However, any user that has the permission to grant and revoke permissions on an object can grant all permissions on that object to any user or any role. Organization administrators can also grant permissions directly to users.

Each user's login account is unique in the system. Managing access, roles, and permissions for each user is based on their user login account. The organization administrator can grant users access to one or more organizations. Within those organizations, each user can be granted multiple roles and permissions.

Users who cannot view or access certain objects or cannot perform certain operations were not granted the permissions to do so.

The following figure illustrates the scope of users and roles in Data Director.

Figure 3-1. Scope of users and roles in Data Director

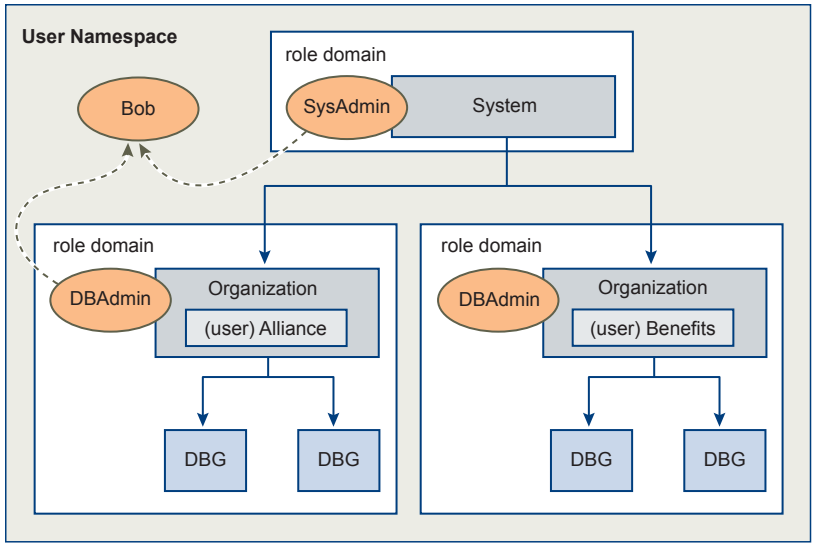

In the figure, user Bob is logged in to Data Director and has been granted access to the system and to the organization Alliance. Bob is also granted the SysAdmin role at the system level, and the DBAdmin role in the organization Alliance. Bob's SysAdmin role applies to the system level. The SysAdmin role does not propagate to any organizations. The role DBAdmin in organization Alliance and the role DBAdmin in organization Benefits are separate roles that apply only within their organizations. Bob has the DBAdmin role in the Alliance organization but does not have access to the Benefits organization.

Authenticating Users

User authentication is based on user login and password.

User login accounts and credentials are unique in Data Director. This enables managing credentials, roles, permissions, and privileges for each user based on the user login account.

Create users and passwords in the following ways.

- A system or organization administrator creates the user account and assigns a password.
- A user registers for a Data Director account and specifies a password as part of the registration request.

Data Director encrypts the password and stores it with the user information. When the user logs in, that user's credentials are stored in an HTTP session. Data Director uses the credentials to validate that the user is authorized to view organization objects (database groups and databases) and to perform tasks.

Role-Based Access Control

Role-based access control enables system and organization administrators to control user access to Data Director and to control what users can do after they log in. To implement role-based access control, system and organization administrators associate (or revoke) privileges, permissions, and roles with (or from) user login accounts.

Users

User logins (users) are unique accounts that enable users to access Data Director. They include a password and identifying information such as name, email address, and phone number. Because user login accounts are unique, system and organization administrators can control each user's access and actions by granting or revoking privileges, permissions, and roles to or from the user's login account.

Users can be active or inactive. Inactive users cannot log in.

Privileges

Privileges control all actions in Data Director. They define the allowable actions within an organization. Privileges apply to particular types of Data Director objects. For example, you can apply the **Stop Database** privilege to organizations, database groups, and databases and apply the **Create Database** privilege to organizations and database groups. Privileges by themselves are not associated with specific objects within an organization.

Permissions

Permissions associate a user and privilege pair with an object in Data Director. Examples are granting a user permission to start or stop a specific database, to modify an organization's backup templates, or to create other users in an organization.

You can grant permissions to users by assigning a role to a user, or by granting permissions directly to the user.

Roles

Roles are collections of permissions that can be associated with or granted to users. Roles provide a convenient way to package all the permissions required to perform a job, such as that of database administrator. Roles apply only to the entity in which they are created. If you create a role at the system level, it applies only to the system. If you create a role in an organization, it applies only to the organization. Organizations have no visibility into each others' roles. If two organizations in the same Data Director data cloud each have a role that has the same name, those roles are distinct within each organization.

One user can have multiple roles within an organization. Users can have access to multiple organizations and can have multiple roles in each organization.

A user can have different roles for different objects. For example, if you have two database groups in your organization, DBG1 and DBG2, you can grant the Database Admin role to a particular user on DBG1 and grant that user the DB User role on DBG2. These assignments might allow the user to perform administrative tasks in DBG1, but not in DBG2.

Predefined Roles

Data Director provides the predefined roles of system administrator, user administrator, and organization administrator. Predefined roles provide a starting point for administering Data Director users and roles and for defining custom roles. You can also create custom roles.

Organization administrator role

Organization adminstrators manage their organizations. They control which users can access the organizations, how users request access to the organizations, and what those users can see and do within the organization. This role has all privileges on the organization for which it is created. Organization administrators invite users to join the organization, grant access, roles, and permissions to users in the organization, create database groups, and can create databases. You can choose to create an administrator user when you create a new organization, or you can select an existing user as the new organization administrator.

Organization administrators perform all user management tasks within their organizations, including the following.

- Add users to organizations, database groups, and databases.
- Modify user settings.
- Remove users from organizations, database groups, and databases.
- Create roles.
- Grant privileges and permissions to roles and to individual users.
- View users, roles, and permissions granted to users and roles.

Organization administrators can view, grant, and revoke privileges on all objects within their organizations, including database groups, databases, and templates. Privileges include **Create Database Groups** and **Modify Database Configuration Templates**.

System administrator role

System administrators operate Data Director. The first system administrator user is created during Data Director installation. This role has all system-level privileges, including managing resources for the system and for organizations. System administrators can see, grant, and revoke permissions at the system level. The first system administrator configures Data Director, creates other system administrators and system-level users, and creates initial organizations. System administrators manage users at the system level. By default they do not have access to organizations unless an organization administrator grants access to them.

User administrator role

The User administrator role manages users at the system level, including creating, editing settings for, and deleting system users.

Privileges

Privileges define the allowable actions on objects in vFabric Data Director. You associate privileges with a user login and a Data Director object to define permissions.

For example, the **Start and Stop Database** privilege indicates that in general, Data Director users can start and stop databases. But the privilege by itself does not indicate which users can start and stop databases, or the databases that they can start and stop. To provide context, you associate the privilege with a user login and a Data Director object. The combination of privilege, user login, and Data Director object is a permission. You can group related permissions into roles to package all the permissions required to perform a job, such as that of database administrator.

System

System privileges relate to Data Director management, such as **Manage Resources** and **Manage System Settings**. These privileges apply only to the system. System privileges do not propagate to organizations.

Organizations

Privileges on organizations relate to organization management, such as **Manage Organization Settings** and **Manage Registration**. Organization privileges apply only to organizations. They do not propagate beyond organization boundaries.

Database group

Privileges on database groups relate to database group management, such as **Create Databases** and **Import Backups**. Database group privileges apply only within the organization and to the organization's database groups.

Organization administrators and users with database group management privileges grant and revoke privileges on database groups, and enable users to access a database group by adding the database group to the user's account.

Databases

Privileges on databases relate to database management, such as **Start and Stop Database** and **Edit Database Info**. Database privileges apply only to databases, database groups, and organizations. If a database-related privilege is on a database group, that privilege applies to all databases within that database group. If the database-related privilege is on an organization, it applies to every database group and database in the organization.

Organization administrators and users with database management privileges grant and revoke these privileges and permissions on databases. To gain access to databases, the databases must be added to a user's account.

Database configuration and database backup templates

Privileges on templates relate to template management, such as edit template and view and user template. Edit template applies only to the organization. View and user template applies to individual templates or to the organization. If a template privilege is on an organization, it applies to all templates within that organization.

Organization administrators and users with template management privileges grant and revoke template privileges and permissions. To gain access to templates, the templates must be added to a user's account.

Propagation of Permissions and Roles

How permissions and roles propagate through an organization depends on where and on what types of objects they are granted. Understanding how permissions and roles propagate can help you to assign them to users appropriately.

Permission and role propagation stops at the organization boundary. Permissions granted within an organization propagate only within that organization. Permissions granted at the system level do not propagate to organizations.

Permissions (and their associated privileges) that apply to an organization are inherited by that organization's database groups and databases. Users or roles can have permissions on specific database groups, and those permissions propagate to databases within the database groups.

Roles apply only to the organization in which they are defined. If a role is defined at the system level, it applies only to the system and is not visible to organizations. If a role is defined within an organization, it applies only to that organization and is not visible to the system or to other organizations.

You can grant permissions and roles on objects within an organization, such as on a database group, on a database, or on a template. For example, granting the Start/Stop Database permission on a database group means that the user or role has the Start/Stop Database permission on all databases within that database group. If a user is granted the Start/Stop Database permission on a database group, that user can start and stop any databases within that database group. However, permissions that apply only to certain types of objects do not propagate to other objects. For example, granting the database group permission Create Database on a database is meaningless.

Organization Privileges and Permissions

Organization administrators grant privileges and permissions to users and roles in their organizations. Those privileges and permissions propagate to database groups and databases in the organization.

You can grant the following types of privileges and permissions to users and roles on organizations.

- User and permission management, such as manage roles and registration and grant/revoke permissions.
- Organization management, such as manage organization settings, database configuration and backup templates, and import databases.
- Database group management, such as manage database groups, create databases, and import backups.
- Database management, such as edit database information, resource, and backup settings, modify database users, upgrade databases.
- Database operations, such as enable/disable databases, delete databases, start and stop databases, and restart databases.
- Database backup and recovery, such as create and delete snapshots, create and delete external backups, clone databases, and recover databases.
- Templates, such as use templates.
- View and monitor, such as viewing reports and monitoring resource usage.

Add Users to Your Organization

Users can self-register to login to Data Director, but cannot access Data Director organizations, database groups, or databases until organization administrators grant access to them. You must add the users to your organization to grant them access.

Prerequisites

- Verify that you have Manage Registration permission for the organization.
- Verify that the system setting Allow Public Registration is on.

Procedure

1 Log in as an organization administrator.

2 Click the **Administration** tab, expand Users and Roles, and click **Users**.

3 Click the plus (+) icon.

4 Complete the user information in the Credentials and Contact Information sections.

5 Grant roles and permissions now or choose to grant roles and permissions later.

6 Click **OK**.

If the Email Validation system setting is on, users receive an activation email that contains a link that they click to activate their account. The new users' status is Pending and the users cannot log in until they activate the account.

The new user appears in the Users list.

Add Roles to an Organization

Roles enable you to group the permissions required to perform tasks associated with a job, such as the job of database administrator. You can then grant the role to users rather than granting individual permissions needed for each task. You can add custom roles to your organization and grant them to the users who are responsible for performing particular jobs.

Prerequisites

- You are logged in to Data Director.
- You have the OrgAdmin role with permissions on all objects in the organization, or permissions for the organization in which to create the role.
- You have grant and revoke permissions on objects.

Procedure

1 Click the **Administration** tab.

2 Expand Users and Roles and click **Roles.**

The OrgAdmin role appears in the list.

3 Click the plus (+) icon.

4 Type a name for the role.

5 (Optional) Enter a description

6 Right-click **Status**.

- Select **Enable** to activate the role.
- Select **Disable** to deactivate the role.

7 In the Permissions section, select the permissions to grant to this role.

You can grant permissions to the role on the organization, database groups within the organization, databases within the organization's database groups, and on organization templates.

8 Click **OK**.

The new role appears in the Roles list.

What to do next

Grant this role to organization users.

Create other roles and grant permissions to them.

Grant a Permission to a User

If a user requires only limited privileges in your organization, you can grant just those privileges to the user instead of granting a role to that user.

Prerequisites

You are logged in to a Data Director organization as an organization administrator.

Procedure

1 Click the **Administration** tab, then click **Users**.

2 Select a user name.

3 Use one of the following methods to access the Edit Permissions window.

- Select the user name, click the gear icon, and select **Edit Direct User Permissions**.
- Right-click the user name and select **Edit Direct User Permissions**.
- Left-click the user name, select **Grant direct user permissions now**, then click **Edit.**

4 Grant privileges to the user.

- To grant a category of privileges to the user, click the **All privileges** check box.
- To grant a specific privilege to the user, click the privilege's check box.

5 Click **OK**.

What to do next

Use the Edit Permissions window to grant the user access to database groups, databases, and templates within the organization.

Modify Organization Security Settings

Organization security settings determine whether your organization allows open registration or users must be invited to register, and whether or not the system administrator can access your organization. You can change the security settings at any time.

Prerequisites

Log in as organization administrator or as a user with the **Manage Organization Settings** permission.

Procedure

1 Click the **Administration** tab.

2 Click **Settings**, then click **Security**.

3 Choose one of the following **Allow public registration** settings.

Setting	Description
No	User registration is by invitation only.
Yes	Users can see the organization and register themselves.

4 Choose one of the following **Allow System Administrator to log into Org** settings.

Setting	Description
No	Do not allow the system addministrator to log into the organization.
Yes	Allow the system administrator to log into the organization.

5 Click **Apply** to accept the settings.

Managing Organizations 4

The basic component of Data Director is the organization. Data Director system administrators create organizations, assign the initial organization administrator, and allocate resources to the organization.

This chapter includes the following topics:

Organization Structure

The structure of organizations depends on the operating mode: Global mode or By Organization mode.

Global Mode In Global mode, all users in the Data Director system are visible to all organizations. Global mode is best for operating Data Director for a single enterprise in which organizations represent business units or departments within the enterprise. Organization administrators can see the global user list and grant access to any user to their organization.

By Organization Mode In By Organization mode, Data Director operates as a service and each organization is a distinct enterprise. Organizations are not visible to each other in By Organization mode. Each organization has its own distinct user list that is not visible to any other organization. Users must either send a request to register to an organization and be approved by the organization administrator, or the organization administrator can invite a user to join the organization.

Organizations contain one or more database groups (DBGs) that in turn contain one or more databases, as shown in the following figure.

Figure 4-1. Data Director Organization Structure

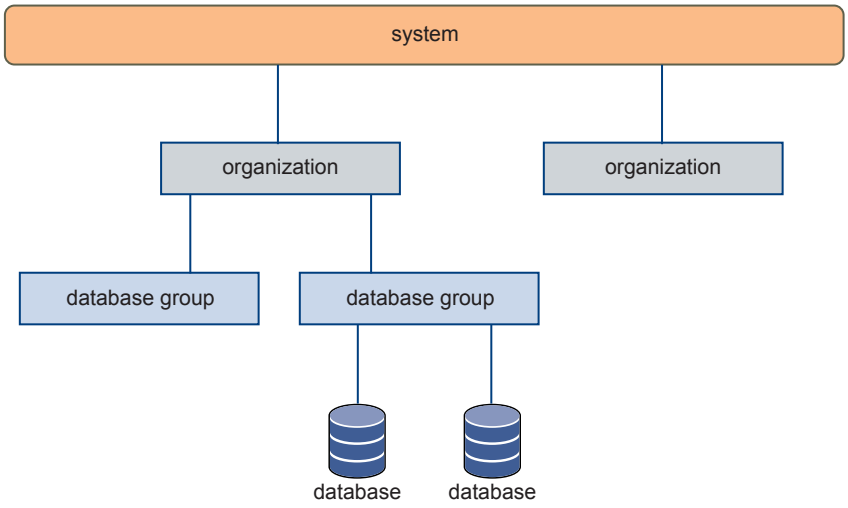

Organization names must be unique within Data Director. Organizations cannot be nested.

Organization roles, policies, and templates apply only within that organization. Resources allocated to an organization are reserved for that organization and cannot be shared among multiple organizations, whether in Global or By Organization mode. This restriction enhances security and ensures resource isolation among organizations.

See Chapter 2, "Managing Data Director Resources," on page 15 for details about resource management in Data Director.

Operating Organizations

Organization operations include system-level tasks such as creating and assigning resources to organizations, and organization-level tasks such as managing organization users, defining and granting roles, and creating database groups.

System administrators perform system-level organization tasks such as the following.

- Create an organization. See "Create an Organization," on page 38.
- View all organizations within Data Director
- Create organization administrators
- Allocate resources to organizations
- Revoke resource bundles from existing organizations
- Implement user authorization and authentication rules (security policies)
- Edit organization properties such as the organization name and description
- Delete disabled organizations

By default, system administrators cannot access organizations. Organization administrators can grant access to system administrators by modifying a security setting for their organization.

Organization administrators perform organization-level, day-to-day tasks such as the following.

- Manage organization users, roles, privileges, and permissions
- Create other organization administrators
- Grant access to the organization to existing users
- Allocate organization resources to database groups

- Implement organization security and backup policies
- Define roles
- Define database configuration and database backup templates
- Monitor organization performance, resource usage, and alarms

Managing Resources For Organizations

Organizations get their resources from vSphere resource pools and from networking and storage resources. These resources are allocated to the organization by Data Director system administrators.

Organizations manage resource bundles on behalf of their database groups and databases. Resource bundles are composed of vSphere resource pools (CPU and memory), storage, and networking resources, and provide the resources used to provision databases.

Resource pools initially created in vSphere are allocated to the Data Director system, where Data Director system administrators use them to create resource bundles. System administrators allocate resource bundles to organizations, and organization administrators can then assign resources to their database groups.

Figure 4-2. Resource Bundles, Organizations, and Database Groups

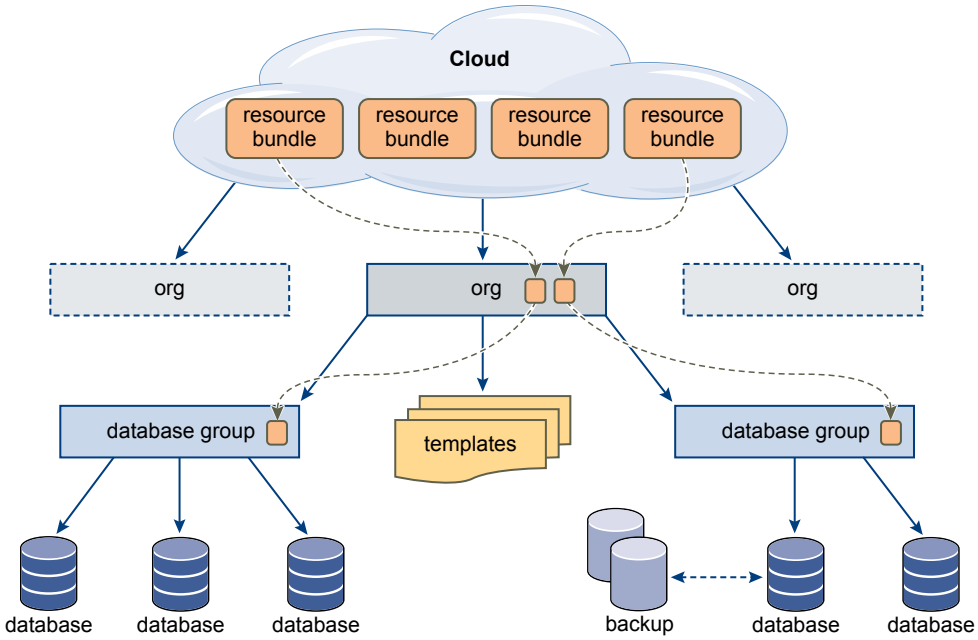

One or more resource bundles can be assigned to an organization, but a resource bundle cannot be shared across organizations. This restriction provides resource isolation and enhances security. Organizations do not compete for available resources and do not have access to each others' CPU, memory, storage, and network resources.

Storage resources are the datastores and allocation amounts for database data and backups. Network resources are the network or networks that are available to the resource bundle and that provide the network(s) for vFabric RelationalDB databases. Data Director system administrators can set up separate networks to provide database isolation.

Organization administrators can subdivide resource bundles across several database groups within their organization.

Databases draw their resources from their parent database groups, which draw their resources from their parent organizations. Organizations draw their resources from the Data Director system.

Managing Organization Users

Organization administrators control user access, roles, permissions, and privileges within their organizations.

Organization administrators control which users can access their organizations and what those users can do. Only organization administrators can grant access to their organizations and assign roles to users within their organizations.

Users can belong to multiple organizations and can be granted multiple roles within those organizations in either By Organization or Global mode systems.

- In a By Organization system, each organization has a distinct user list that is not visible to other organizations. To join an organization, users send a request to the organization administrator, or the organization administrator can invite a user to join.
- In a Global system, the user list for the system is visible to all organizations. All users belong to all organizations. Organization administrators grant roles to users to enable them to perform tasks in the organization.

Organization administrators can grant any roles defined within their organizations to organization users. In By Organization mode, the user must be on the organization's user list.

Organization administrators control what users can do in their organizations by defining roles, privileges, and permissions within their organizations, then granting them to organization users. Roles are specific to the organization in which they are created and are not visible to other organizations.

See Chapter 3, “Managing Users and Roles,” on page 25.

Create an Organization

The Data Director system administrator creates organizations to allow organization administrators independent management of their database groups and databases.

Prerequisites

- Resource bundle(s) must be created and available for allocation.
- You are logged in as a Data Director system administrator.

Procedure

1 With System selected, click **Manage & Monitor**.

2 Click **Organizations** in the left pane.

3 Click the plus (+) icon to create an organization and specify the organization information in the wizard.

Wizard screen	Action
Name and Description	Specify a name and optional description and click **Next**.
Organization Administrator	To create a new organization administrator user, perform the following tasks. a Click **Create a new user**. b Specify the user name, password, first and last name, and optionally, phone number. c Click **Next**. To use an existing user, perform the following tasks. a Click **Choose an existing user**. b Select the user from the list. c Click **Next**.
Resource Bundles	You can assign resource bundles at any time after creating the organization. To skip the assign resource bundles step, click **Assign resource bundles later**. To select a resource bundle now, click **Choose an existing resource bundle** and select a resource bundle from the list. Click **Finish**.

The new organization appears in the Organizations list.

What to do next

Create resource bundles and assign them to the organization. See "Create a Resource Bundle," on page 21.

Managing Database Groups 5

Database groups contain sets of databases within organizations. Database groups allow organization administrators to provide the resources for operating and provisioning databases and to apply access and authorization rules (security policies) to those databases. Grouping databases enables subdivision of resources from the organization's allocated resources.

This chapter includes the following topics:

Database Group Management Overview

Organization administrators create database groups to enable efficient management of databases and database templates. Administrators also allocate the resources required to provision, operate, and control database groups.

The databases within a database group are usually related. For example, in Global user management mode, where organizations represent business units in a single enterprise, database groups can group databases for departments within the business unit. In By Organization user management mode, where each organization represents a unique enterprise, database groups can group databases for business units within that enterprise.

Each database group can contain one or more databases. Databases must reside in one database group and cannot be divided among database groups.

Database groups must reside in one organization and cannot be nested.

The following figure shows the relationship between organizations and database groups.

Figure 5-1. Database Groups in the Data Director Architecture

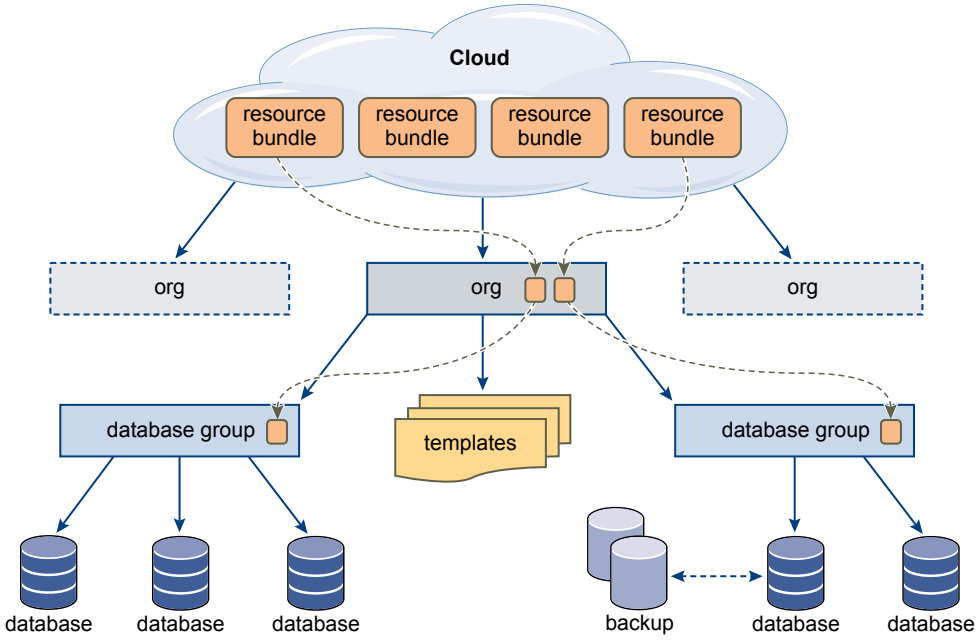

Managing Resources for Database Groups

Database groups require CPU, memory, storage, and networking resources to enable database operation, provisioning, and backup. To provide database groups with the required resources, organization administrators create resource bundles and allocate those resource bundles to their database groups.

Resource bundles consist of CPU, memory, storage, and networking resources. An administrator can share a resource bundle among multiple database groups. The administrator allocates part of the resource bundle to each database group. An administrator can also assign a resource bundle exclusively to one database group.

Organization administrators assign resources when they create database groups and can add or expand resources as required. Each database group has exclusive use of its assigned resources to ensure resource isolation. Resource isolation ensures that database groups and the databases that they contain do not compete for resources or have visibility into the resources of other organizations. When organization administrators create database groups, they optionally specify how much unused CPU and memory to reserve for the database groups. The administrator also assigns the database group's priority for distribution of unreserved resources. The priority options are high, medium, or low.

Because administrators allocate resources to organizations and then assign resources to that organization's database groups, each RelationalDB database must be contained within one database group. You cannot split databases among database groups, and you cannot move a database to a different database group after the database is created.

Use the following guidelines to estimate the resources that you need for a database group.

- Calculate the storage allocation based on the expected number of databases that the database group will contain, the amount of storage allocated for each of those databases, and room for growth:

  ```
  (number of DBs) X (storage for those DBs) + (room for growth)
  ```

- Determine the size of the backup storage allocation to support the external backups for each database in the database group plus the Point-in-Time Recovery allocation for each database.

Database Groups and Security

Role-based access control and direct user permissions form the security policies that determine which users can access particular database groups and the actions that the users can perform. Database groups inherit security policies from their organizations.

Organization administrators define the security policies for their organization, including user roles, permissions, and privileges.

For example, an organization administrator creates a user role with permissions on database groups. These permissions include create database, take database snapshots, and start or stop database. Those roles and their associated permissions apply to each database group within the organization, and to each database within each database group.

Chapter 3, "Managing Users and Roles," on page 25 discusses the Data Director security model and explains how you can use roles for fine-grained permission management.

Create a Database Group

Database groups contain sets of databases within an organization. Database groups enable grouping related databases and provide efficient use of resources needed to provision and operate databases.

Prerequisites

- Verify that at least one resource bundle is allocated to the database group's organization. See "Create a Resource Bundle," on page 21 if no resource bundle is available.
- You must be logged in as an organization administrator or have permissions to create or modify database groups.

Procedure

1 Click the **Manage & Monitor** tab.

2 Click the **Database Groups** tab.

3 Click the plus (+) icon to create a database group.

4 Specify the following information in the Create Database Group wizard.

Wizard screen	Action	
Name and Description	Type a name and optional description and click **Next**.	
Resource Bundle	Select a resource bundle from the list and click **Next**.	
Resources	Specify the resources for this database group.	
	Network	Select the network from the drop-down menu.
	CPU & Memory	■ Assign the priority (High, Medium, or Low). ■ (Optional) Check the **Reserve resources for this database group** check box and enter the reservation amounts for CPU and memory.
	Storage	Enter the amount of database and backup storage in the **Database storage** and **Backup storage** fields.

5 Click **Finish**.

The new database group appears in the database group list.

What to do next

Click the database group name to open the database group and view and edit its properties.

Managing Database Templates 6

Data Director database templates allow organization administrators to standardize databases and their backup policies. Templates also impose limits on resource consumption. Database administrators can create and back up databases consistently by using templates and can create, clone, and customize templates.

This chapter includes the following topics:

Introduction to Database Templates

Data Director includes database templates to help administrators streamline resource allocation and standardize database setup and backup setup. Templates help database administrators to quickly provision a database and to select a backup process.

Data Director supports database configuration templates and backup templates. Included with Data Director are several optimized templates. When you create an organization, Data Director copies the system-defined templates to the new organization. Organization administrators can modify the organization-specific template instances or configure new templates.

You can create database templates and publish them immediately or publish them later. When a template is not published, it can be viewed or managed, but cannot be used for provisioning or other purposes.

Database Configuration Templates

Database configuration templates define the computing and storage resources for creating a database, the database settings, and the high availability settings. Each template defines resource settings and database settings.

Resource Settings

When you create a template, you can specify the number of virtual CPUs, memory size, and recommended database storage allocation. You can enable high availability for the template and all corresponding databases. You can also choose the CPU and memory priority, which affects the allocation of resources for all databases in the database group. The levels (high, medium, and low) give certain databases higher priority than other databases in the same database group. The **CPU reservation** and **Memory reservation** fields let you explicitly reserve resources for each database you create from the template.

	If you make changes to a template, databases that are already created from the template are not affected.
Database Settings	Specify connection, memory, IO, WAL, checkpoint, logging, and other information. When you create a template, the wizard includes defaults for each value. The wizard also includes a database settings calculator. The calculator prompts for usage information and changes the default based on that information and the resource settings you specified on the first wizard pane.

You can create different templates for different situations. For example, you can define a configuration template for engineering with a small memory size and have high availability disabled. The configuration template for QA can be defined with a larger memory size and with high availability enabled.

Backup Templates

Backup templates define backup settings for databases. You can associate a backup template with a database when you create the database, or you can associate a backup template with a database at a later time. See "Select a Database Backup Template," on page 81.

You can use one of the predefined backup templates for consistency across your organization. See "Backup Template Settings," on page 80.

You can also clone and customize an existing template and associate the custom template with your database. You can customize frequency, start time, and retention for snapshots and for external backup. You can also enable and customize point-in-time recovery, and you can specify a backup label. See "Create a Backup Template," on page 48.

Create a Database Configuration Template

You can create a database configuration template by cloning a template or by configuring a new template. In both cases, you can specify the resource settings and the database settings for the template.

Only organization administrators or users with **Manage Database Configuration Templates** or **Manage Backup Templates** privileges can create, edit, and delete templates.

Prerequisites

Log in to Data Director as an organization administrator or as an administrator with privileges to create and modify templates.

Procedure

1 Click the **Administration** tab.

2 Click **Templates**, and click **Database Configuration Templates**.

3 Create a template or clone a template.

Creation Method	Action
New template	Click the green plus icon above the menu bar on the left.
Cloning	Right-click an existing template and choose **Clone**.

4 In the Create Database Configuration Template wizard, type a name and description.

5 Specify whether you want to publish the template, and click **Next**.

When a template is not published, you can view or manage it, but you cannot use it for provisioning or other purposes.

6 Enter resource settings for the template and click **Next**.

Option	Description
vCPUs	Number of virtual CPUs the database virtual machine will use.
High availability	Select **Enable** to protect the database with vSphere High Availability. See the *vSphere Availability* documentation.
Memory size	Amount of memory the database virtual machine will use.
Recommended database storage allocation	Recommended storage allocation for this database.
CPU and memory priority	Select **Automatic** to allow the vCenter Server system to allocate CPU and memory to the virtual machine. If you select another value, the CPU priority determines how unreserved CPU and memory resources are assigned to this database as compared to other databases in this database group.
Explicitly reserve resources for databases created by this template	If checked, you can reserve resources for running databases. Reservations guarantee that the database has the specified amount of CPU and memory available.
CPU reservation	Number of MHz to reserve for this database.
Memory reservation	Number of MB to reserve for this database.

7 Specify the database settings by accepting the default settings, by using the Database Settings Calculator, or by entering values explicitly, and click **Finish**.

Modify a Database Configuration Template

If the requirements for resources or other aspects of your environment change, you can modify existing database configuration templates. Databases that you create from the new template will use the new settings.

Prerequisites

Log in to Data Director as an organization administrator or as an administrator with privileges to modify a configuration template.

Procedure

1 Click the **Administration** tab.

2 Click **Templates**, and click **Default Database Configuration Templates**.

3 Right-click the template that you want to modify, and perform one of the supported actions.

Action	Description
Clone	Creates a copy of this template. When you clone a template, the Create Database Configuration Template wizard appears and you can configure the resource settings and database settings for the clone.
Delete	Deletes the selected template.
Unpublish	Disables provisioning and other capabilities for this template. When a template is not published, it can be viewed or managed, but cannot be used for provisioning or other purposes.
Edit Permissions	Allows you to specify who can use this template, and what each user can do. You can change the permissions for an existing user, remove an existing user, and add a role. Users who can create a database from the template do not automatically have permissions to modify the template.
Properties	Allows you to modify the settings that you specified when you created the template. See "Create a Database Configuration Template," on page 46 for a discussion of the properties you can change.

You can create databases with the new settings from the modified template. Databases you previously created from the template do not change.

Create a Backup Template

Backup templates include frequently used backup settings. You can use one of the existing templates, clone and customize a template, or create a template. You can then associate the backup template with a database that you create.

The system-defined backup templates use recommended settings for different situations.

Prerequisites

Log in to Data Director as an organization administrator or as an administrator with privileges to create and modify a backup template. See "Backup Template Settings," on page 80 for information about system-defined templates.

Procedure

1 Click the **Administration** tab.

2 Click **Templates**, and click **Backup Templates**.

3 Create a template or clone a template.

Creation Method	Action
New template	Click the green plus sign above the menu bar on the left.
Cloning	Right-click one of the existing templates and select **Clone**.

4 In the Backup Template wizard, type a name and description for the template and click **Next**.

5 In the Backup Settings panel, specify the snapshot settings.

Option	Action
Frequency	Select one of the options from the menu. Select **Never** if you do not want backups for databases that use this backup template.
Start Time	Select **Automatic** to allow the system to control the start time, or enter a start time. The system initiates a backup within two hours of the target start time, depending on system load.
Retention	Select one of the available prespecified durations.

6 Specify the external backup settings.

Option	Action
Frequency	Select one of the options from the menu. Select **Never** if you do not want backups for databases that use this backup template.
Start Time	Select **Automatic** to allow the system to control the start time, or enter a start time. The system initiates a backup within two hours of the target start time, depending on system load.
Retention	Enter the number of hours or the number of copies to retain.

7 Specify the general backup settings.

Option	Action
Point in time recovery	Enable or disable point-in-time recovery. Point-in-time recovery ensures that each database transaction is recorded and enables you to revert the database to any point in time within a certain time range. The start time for point-in-time recovery is right after point-in-time recovery is enabled, when the system creates a baseline backup or snapshot. You cannot remove the baseline backup. If you do, the start time for point-in-time recovery changes. The time range for point-in-time recovery is from the time of your oldest automatic backup to the present. The oldest backup can be an external backup or a snapshot. Backups with extended retention are not supported as oldest backups. Point-in-time recovery consumes space in the backup storage area. Depending on database load and retention lengths, this feature might require a significant amount of storage.
Backup label	Type the first part of the name of the backup. Defaults to *user-specified label-date_and_time-dbname*. For your database named db1, if you entered testbackup as the label and the backup starts at 12:30:45 on May 30, 2011, the full name is `testbackup-2011-05-30-12-30-45-db1`. If you do not specify a label, the system uses the date, time, and database name.

8 Click **Finish** to complete creating the template.

What to do next

You can assign the template to databases.

Modify a Backup Template

If the requirements for backups in your environment change, you can modify existing backup templates.

Prerequisites

Log in to Data Director as an organization administrator or as an administrator with privileges to modify a backup template.

Procedure

1 Click the **Administration** tab.

2 Click **Templates**, and click **Backup Templates**.

3 Right-click the template that you want to modify and perform one of the supported actions.

Action	Description
Clone	Creates a copy of this template. When you clone a template, the Create Backup Template wizard appears and you can configure the backup settings for the clone.
Delete	Deletes the selected template.
Unpublish	Disables provisioning and other capabilities for this template. When a template is not published, it can be viewed or managed, but cannot be used for provisioning or other purpose.

Action	Description
Edit Permissions	Allows you to change the permissions for an existing user, to remove an existing user, and to add a role.
Properties	Allows you to modify the settings you specified when you created the template.

You can create databases with the new settings from the modified template. Databases you previously created from the template do not change.

Managing Databases 7

Database administrators and application developers manage databases from creation to decommissioning and deletion. The tasks include managing database entities and data, scheduling backups, and performing recovery.

This chapter includes the following topics:

Database Lifecycle

The database lifecycle includes database creation and resource allocation, managing the database schema and data, performing backup and recovery tasks, and decommissioning databases. Database administrators and application developers perform the database lifecycle tasks.

Create

Create and allocate resources to a new database using database configuration templates. Database templates specify sets of database parameters, including resource limits. Application developers can perform do-it-yourself database creation using the templates. See "Requirements for Creating Databases," on page 52.

Administrators can grant permissions to their users that enable the users to create databases from templates, but do not allow them to modify the templates or change the default resource allocations. This restriction provides resource limit enforcement and allows administrators to retain control of resource and security policies. See Chapter 6, "Managing Database Templates," on page 45.

Manage schema

Manage the database schema and add data. You can create tables, designate primary and foreign keys and indexes, and create views, sequences, triggers, and other database entities.

Backup and restore

Safeguard your data by taking regular backups and testing your backups. See Chapter 10, "Safeguarding Data," on page 77.

Clone

Ensure access to consistent, yet isolated databases by cloning the database for specific purposes such as development or quality assurance. See Chapter 8, "Cloning Databases," on page 57.

Scale up

Dynamically increase the database size as required, during the development, test, and production phases.

Monitor performance and usage	Use the Data Director user interface to monitor tasks and events. See Chapter 11, "Monitoring the Data Director Environment," on page 89.
Stop and restart the database	Stop and restart, for example, to perform maintenance tasks.
Decommission the database	Disable and then delete databases. Free up the resources when they are no longer needed.

Every database requires a database owner account that can perform all schema management operations. This account is specific to the database and cannot log in to Data Director. You can add database owner accounts after database creation. Data Director users must log in with their database-specific credentials to view the database, its entities, and its data or to perform database management tasks.

Database administrators and application developers can manage databases only if they have appropriate permissions and roles granted to them by the organization administrator. The permissions and roles must be granted on the database group or on the database, and apply only within the organization in which they are granted.

Requirements for Creating Databases

You must have certain permissions to create databases, and you must calculate the storage needed for database and related data.

Permissions Required for Creating Databases

To create databases, you need **Create Databases** permission on the database group that will contain the database and **Use Template** permission on at least one database configuration template.

It is useful to have the following permissions on the database group and on the database.

- Create snapshots.
- Create external backups.
- Delete snapshots and manage their retention time.
- Clone the database.
- Recover the database from a backup or snapshot.

The organization administrator can create a role with these permissions and assign users in the organization to the role.

Calculating Database Storage Allocation

During the database creation process, you specify database storage allocation, point-in-time recovery storage allocation, and the database group for the database. The database group provides the CPU, memory, storage, and network resources required to run the database. The storage and point-in-time recovery allocations specify how much of the database group's resources to use for this database.

When you calculate the amount of storage to allocate to the database, proceed as follows.

- Estimate how much data will be stored in the database.
- Consider the number of users and average expected number of transactions in a particular time period and include room for growth.
- If you plan to enable point-in-time recovery, calculate additional storage to accommodate the point-in-time recovery write-ahead logs (WALs). The size of the allocation depends on the expected volume of transactions on the database.

Database storage allocation is for the database data only. It does not include overhead for the operating system, database software, swap space, or snapshots. You must have enough resources available to cover both the database allocation and to cover any overhead. Even if the database group has enough free space for creating a database, database creation does not complete if you do not have enough resources for the overhead.

Database Creator Permissions

After database creation finishes, the following permissions on the new database are granted to the database creator.

Edit Database Info	Enables the database creator to edit database properties such as the name, description, and size of the database.
Modify Database Users	Enables the database creator to add or modify database users for this database. Database users are granted full permission on this database.
Restart Database	Enables the database creator to add or modify database users for this database. The database users are granted full permission on this database.
View Database	Enables the database creator to view the database.

Create a Database

When you need a database for a new application, you can create it with a database configuration template. The template is configured to allocate resources to the database.

Prerequisites

- Verify that you have access to the organization and database group in which to create the new database.
- Verify that you have **Create Databases** permission on the database group in which you create the new database.
- Verify that you have **Use Template** permission on at least one database configuration template.

Procedure

1 Navigate to the organization and to the database group in which to create the database.

2 Click the **Manage & Monitor** tab.

3 Click the **Databases** tab.

4 Click the plus (+) icon and specify the following information in the **Create Database** dialog box.

Text Box	Action
Name and Description	Enter a name and optionally, a description of the database.
Owner Account	Enter an owner account for the database. Each database requires an owner who can perform all schema management operations. The owner account is specific to the database and cannot log in to Data Director.
Password	Enter and confirm the owner account password.
DB config template	Select a database configuration template from the drop-down menu. Database templates determine the initial size of databases and other parameters, such as resources the database can consume. The create database process calculates the default storage allocation values based on the resources defined in the database template.
Storage allocation	Select the storage allocation for this database. The minimum is 1GB. The recommended (and default) value is 5GB.
Backup template	(Optional) Select a backup template from the drop-down menu. You can select a backup template for specific purposes, such as development, or select no backups.

Text Box	Action
Point-in-Time Recovery allocation	Select the number of gigabytes to allocate for point-in-time recovery operations. The minimum is 1GB.
Encoding	Select the encoding for the database from the drop-down menu. The default is UTF8.
Tags	(Optional) Select one or more tags for this database. Use tags to filter the list of databases that you view in an organization's **Databases** tab, for example, all your customer relationship databases can have a tag called CRM.
Database group	Select the database group in which to create this database. Databases must reside in one and only one database group.

5 Click **OK**.

The database appears in the **Databases List** with a status of Creating. The process can take a few minutes. The status changes to Running when creation finishes successfully.

The following permissions are granted to the database creator after database creation finishes.

- **Edit Database Info**
- **Modify Database Users**
- **Restart Database**
- **View Databases**

What to do next

You can load the database data and use the database.

Using Tags

Tags are text labels that users create and associate with databases. Users can create tags on any databases that are visible to them. Tags provide a simple way to search for databases in a particular database group or organization.

Users can see only the tags that they create.

Tags enable filtering the list of databases that appear in an organization's **Databases** tab. For example, a user can create a tag called HR and associate the tag with all the HR databases in an organization. When that user views the **Databases** tab, filtering on the HR tag displays only the databases with that tag.

You can associate a tag with a database during database creation. See "Create a Database," on page 53. You can also associate tags with an existing database.

Create a Tag

Tags provide a simple way to search for databases in a database group or organization.

Procedure

1 Log in to Data Director as an organization administrator.

2 Click the **Manage & Monitor** tab.

3 Click the **Tags** tab.

4 Click the plus (+) icon.

5 Type the name of the tag in the **Create Tag** dialog box and click **OK**.

Associate a Tag with an Existing Database

Tags support searches for databases. You can associate a tag with a database to help with searches for databases.

Procedure

1 Log in to Data Director as an organization administrator or as a user with **Edit Database Info** permission.

2 Click the **Manage & Monitor** tab.

3 Click the **Databases** tab.

4 Right-click a database to display the **Actions** menu and select **Properties**.

5 Click the **Basic** tab, select a tag in the **Tags** field, and click **Edit**.

6 Click the check box for the tag or tags to associate with the database and click **OK**.

Cloning Databases 8

In Data Director, cloning a database is easy to configure and to perform. You have a choice of cloning operations that include full database clone, linked database clone, and schema only clone. You can customize the clone's database settings and backup settings during clone creation.

This chapter includes the following topics:

Clone Types

Data Director allows you to clone databases and supports several clone types with different contents, storage requirements, and performance characteristics.

You can create different types of clones depending on whether you need schema only, schema and data, full database clones, or clones that take advantage of the linked clone technology. You can also create a clone that includes neither schema nor data but includes the database settings and backup settings.

Schema Only Clone and Schema and Data Clone

If you create a schema only clone, none of the data in your database is cloned. If you create a schema and data clone, the complete set of schema and data is included in the clone. In that case, you might have to run a script over the clone to remove confidential data.

You can also clone only the configuration. In that case, the clone includes neither the schema nor the data.

Full Database Clones and Linked Database Clones

When administrators clone a database they can choose a full database clone or a linked database clone.

Full Database Clones A full database clone is a complete copy of the source database. Full clones allow you to isolate the source and the clone. The isolation might be useful, for example, if the source database cannot tolerate any performance degradation. Creating a full clone is typically more time consuming than creating a linked clone.

Figure 8-1. Full Database Clone

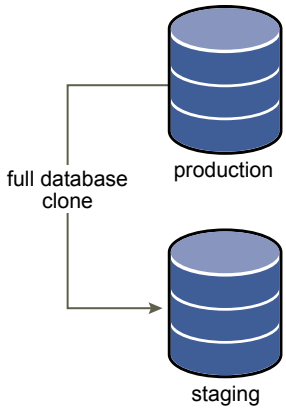

Linked Clones

Linked clones are two or more databases that share storage. The linked clone technology supports efficient sharing of duplicate data. Linked clones use delta disk backings. A delta disk backing is a virtual disk file that is located on top of a standard virtual disk backing file. When one of the databases writes to disk, the data is written to that database's delta disk. When one of the databases reads from disk, the read process first checks the delta disk. If the data is not in the delta disk, the database retrieves the information from the parent disk.

You can create a linked clone from a snapshot or from the current running point but not from an earlier backup or from a specific point in the past. If you use linked clones, the clone and the source database cannot change data disk size.

Figure 8-2. Linked Database Clone

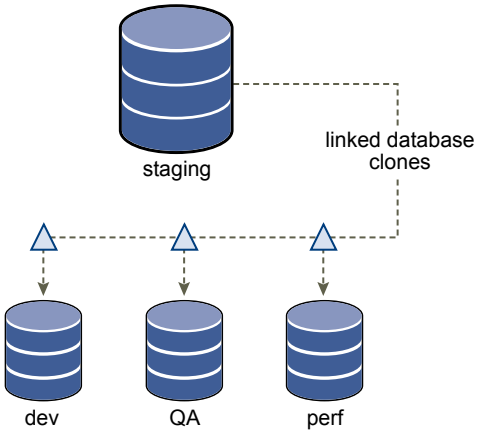

Choosing a Full Clone or a Linked Clone

To choose the clone type most appropriate for your situation, consider these points:

- Full clones require a longer time to create than linked clones.
- Linked clones are much faster to create.
- Linked clones do not support storage isolation. Having several linked clones can affect the performance of the source database and the performance of the linked clones.

VMware best practice is to first create a full clone of a production database to use it as a staging clone. Next, you create linked clones of the production system full clone, the staging system in the illustration. In this scenario, potential performance degradation affects only the staging system clone and not the production database.

Figure 8-3. Using a Full Clone as a Staging Clone

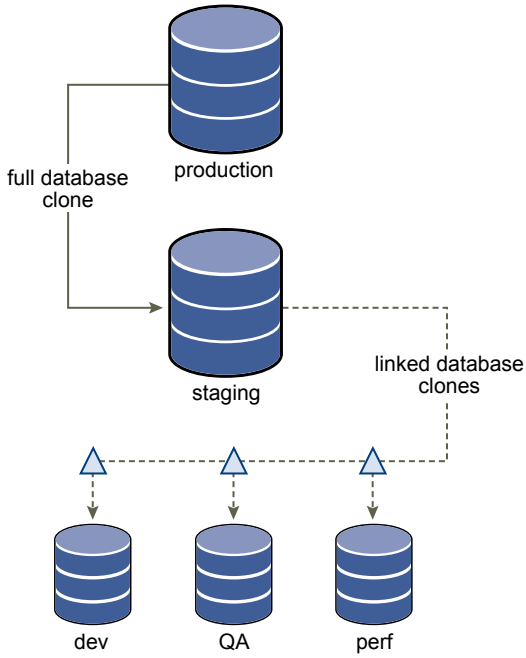

Cloning Customizations

You can customize a clone when you create it. You can specify new database configuration settings and backup settings for the clone, choose the clone point, set an immediate backup, and set an expiration date for the clone.

When you clone an existing database, you can customize the clone to suit your needs. For example, start with a staging database that requires frequent backups, a sizable storage allocation, and point-in-time recovery. For developers, you can create a clone of the staging database that uses a development backup template.

The cloning process also allows you to choose the clone point, which is the point in time at which the clone is created from the source database. You have the following choices.

- Clone the current state of the source database.
- Clone the source database as it was at a certain point in time. You can specify the date and time for this operation.
- Clone one of the backups of the source database.

Clone a Database

You clone a database to create an exact copy that you can use for testing or other purposes. When you start the database cloning process, the Clone Database wizard prompts you for clone type and clone customization information.

Prerequisites

Log in to Data Director as an administrator or as a user with **Clone Database** privileges.

Procedure

1 Configure the Clone Type on page 60

The clone type allows you to specify which data are cloned. You can also choose custom configuration settings and backup settings for the clone.

2 Configure the Database Settings of the Clone on page 61

When you clone a database, you can customize database attributes for the clone as part of clone creation.

3 Configure the Clone Point Settings on page 61

Data Director allows you to clone the current state of the database, clone a specific point in time, or clone a backup of the database.

4 Configure Recovery Point and Expiration for the Clone on page 62

You can specify an expiration date and time for a clone if you anticipate that the clone will no longer be useful at a certain date and time. Data Director displays an alarm when the clone expires. You can configure clone expiration to also send an email to the administrator when the clone expires. Expiration does not mean that Data Director stops or deletes the clone.

Configure the Clone Type

The clone type allows you to specify which data are cloned. You can also choose custom configuration settings and backup settings for the clone.

As part of the cloning process, you decide whether to create a full clone or a linked clone. See "Clone Types," on page 57 for a discussion of the differences.

Prerequisites

Log in to Data Director as an administrator or as a user with **Clone Database** privileges.

Procedure

1 Right-click an existing database and select **Clone** to start the Clone Database wizard.

2 In the Clone Type panel, select the data to clone and the type of cloning process to use.

Option	Description
Clone schema and data	Clones both the database schema and all data. If you clone both schema and data, you can select a full clone or a linked clone.
Clone schema only	Clones the database schema. Does not clone the data.
Do not clone schema or data	Clones only the database settings and database backup settings. Does not clone the schema, and does not clone the data.

3 Select the database configuration settings to use for the clone.

Option	Description
Clone from source database	Uses the current database settings of the source database. Click **Customize Settings** to modify the source settings.
Choose a new template	Allows you to select one of the available database templates. Select **Advanced selection** to override one or more values of the existing template.

4 Select the backup settings to use for the clone.

Option	Description
Clone from source database	Uses the backup settings of the source database. Click **Customize Settings** to modify the source settings.
Choose a new template	Allows you to select one of the available templates. Select **Advanced selection** to override one or more values of the existing template.

5 Click **Next** to continue to the Database Settings panel.

What to do next

Specify the database settings information, or click **Back** to return to the Clone Type panel for modifications.

Configure the Database Settings of the Clone

When you clone a database, you can customize database attributes for the clone as part of clone creation.

Prerequisites

Complete the **Clone Type** panel of the Clone Database wizard described in "Configure the Clone Type," on page 60.

Procedure

1 In the Database Settings panel, specify a name for the clone database, or leave the default.

2 Type a description of the clone database.

3 If you specified **Do not clone schema or data** earlier, specify storage allocation or leave the default.

   a Estimate the storage required for the data you expect to store in the database and enter that value for storage allocation, or leave the default.

   Database storage allocation is for the database data and does not include storage for operating system, database software, swap allocation, or snapshots. The default is based on the size of the source database.

   b Specify the point-in-time-recovery storage allocation, which depends on the volume of transactions that you expect.

   The point-in-time-recovery allocation stores the WAL (write-ahead logs).

   For other clone types, you cannot change storage allocation.

4 Select tags for the clone.

   Tags support a mechanism for finding multiple databases that share the same tag. See "Using Tags," on page 54.

5 If you are creating a full clone or you are cloning the database settings, choose the database group to add the clone to.

   You cannot change the database group for linked clones.

6 Click **Next** to continue to the Clone Point panel.

What to do next

Specify the settings in the Clone Point panel, or click **Back** to return to the Clone Type panel for modifications.

Configure the Clone Point Settings

Data Director allows you to clone the current state of the database, clone a specific point in time, or clone a backup of the database.

The available options for the clone point depend on the type of clone. For a clone of only database settings, clone points are not supported.

The following table summarizes the available clone point options.

Table 8-1. Clone Point Options

Clone Type	Now	Specific Point In Time (Snapshot)	From Backup
Full	Yes	Yes	Yes
Linked	Yes	Yes	No
Schema only	Yes	No	No

If you create a schema only clone, the database must be powered on. If you clone from a backup, see Chapter 10, "Safeguarding Data," on page 77.

Prerequisites

Complete the Clone Type and Database Settings panels of the Clone Database wizard.

Procedure

1 In the Clone Point Settings panel, specify the point from which to create the clone.

Option	Description	
Now	Clones the current state of the database.	
Specific point in time	Clones the source database at a point in time based on the point-in-time recovery settings for your database.	
Select a backup	Clones the backup that you select. You can select the backup based on its attributes.	
	Type	Snapshot or External.
	Retention	How long to keep the backup.
	Backup label	A unique label for the backup.

2 Click **Next** to continue to the Options panel.

What to do next

Specify the recovery point and expiration for the clone, or click **Back** to return to the Database Settings panel for modifications.

Configure Recovery Point and Expiration for the Clone

You can specify an expiration date and time for a clone if you anticipate that the clone will no longer be useful at a certain date and time. Data Director displays an alarm when the clone expires. You can configure clone expiration to also send an email to the administrator when the clone expires. Expiration does not mean that Data Director stops or deletes the clone.

Prerequisites

Complete the Clone Type, Database Settings, and Clone Point panels of the Clone Database wizard.

Procedure

1 (Optional) Specify an expiration time and behavior.

When a clone expires, Data Director displays an alarm. You can also set up Data Director to send an email to the administrator. The administrator can then delete the clone.

2 Click **Finish** to complete clone setup.

Data Director creates a clone of the current database using the settings you specify.

What to do next

Monitor the creation progress in the task bar on the right, or check the database list for the database group to verify that the clone is created.

Managing Database Entities 9

Managing vFabric Postgres database entities includes managing schemas and tables, administering the database, and performing SQL management tasks.

To manage databases, you need database management and database operations permissions on the database. Your organization administrator can create a role that grants the necessary permissions, and grant that role to you.

See "Requirements for Creating Databases," on page 52 and "Create a Database," on page 53 for information about creating vPostgres databases.

The following are database management and database operations-related privileges.

- **Edit Database Info**
- **Edit Database Settings**
- **Edit Database Resource Settings**
- **Edit Database Backup Settings**
- **Enable, Disable, and Delete Database**
- **Start and Stop Database**
- **Restart Database**
- **Modify Database Users**
- **Upgrade Databases**

Permissions that you have on the organization apply to all database groups and databases in the organization. Permissions on a database group apply to all databases in the database group.

Data Director supports the following types of management tasks.

Database Entity Management

Database entity management includes creating, replacing, updating, and deleting database entities. These database entities include schemas, tables, views, indexes, functions, sequences, triggers, constraints, and users.

Database Administration

Database administration includes the following tasks.

- Monitoring resource usage and database status.
- Analyzing usage statistics and performance tuning.
- Cleaning up stale data (vacuuming old data).
- Migrating data to vPostgres databases.

- Backing up databases.

SQL Management SQL management tasks include SQL profiling, query plan analysis, running ad-hoc queries or SQL scripts.

After you log in to an organization in Data Director, click the **Manage & Monitor** tab to view your database groups and databases. You see a list of database groups to which you have access and the list of databases within that database group. To view a database and its entities, double-click the database in the middle pane. The database appears under the database group in the left pane. You can expand the database to view its entities and the objects those entities contain, such as schemas and the number of tables, views, and other objects.

You can perform administrative and SQL management tasks.

- Monitoring database status and resource usage.
- Alarms.
- Status of tasks and events.
- Viewing reports.
- Checking the database logs.
- Entering SQL queries.
- Obtaining the JDBC connection string.

To manage a specific database entity or object, click the entity or object in the left pane. The entity or object appears in the middle pane and you can perform management tasks.

This chapter includes the following topics:

Database Entity Management

You can manage database entities such as databases, schemas, users, and roles from the Data Director **Manage & Monitor** tab.

Managing database entities includes creating, altering, dropping, and browsing database entities such as the following.

- Schemas
- Tables
- Views
- Columns
- Indexes
- Sequences
- Constraints (primary, foreign, and unique key)
- Users

Click any of the schema objects in the left pane to view a list of the objects in the middle pane. You can manage individual objects from the middle pane.

Create a Schema

After you create a database, you set up its entities, starting with the database schema. You create vPostgres database schemas from the Data Director **Manage & Monitor** tab.

Prerequisites

- Verify that a database exists in which you can create schema.
- Log in to Data Director as a user with database privileges.

Procedure

1 Click the **Manage & Monitor** tab.

2 Click your database group.

3 Double-click your database to select it.

4 Right-click the database name in the left pane, and select **Create > Schema**.

5 If a login prompt appears, enter the database login.

The database login is different from your Data Director login and might be, for example, the database owner login name and password. This login is for database security purposes.

6 Enter the schema information.

Data Director creates the database schema.

What to do next

Create schema entities such as tables, triggers, users, and so on.

Create a Table

After you create a schema, you create tables to contain the schema's data. Create tables from the Data Director **Manage & Monitor** tab.

You are a DBA or application developer setting up a database.

Prerequisites

You created a database and a schema.

Procedure

1 Log in to Data Director as a user with database privileges and click the **Manage & Monitor** tab.

2 In the left pane, click the arrow next to your database name to expand it.

3 Right-click the schema and select **Create > Table**.

The following tabs appear.

- **Basic Settings**
- **Columns**
- **Constraints**
- **Auto Vacuum Settings**

4 In the **Basic Settings** tab, enter information such as the table name, whether the name is case-sensitive, and fill factor.

5 You can now configure columns or exit the wizard.

- Click **Next** to configure columns.
- Click **Finish** to create the table.

6 (Optional) In the **Columns** tab, click **Add** to add a column, and complete the column fields.

a Type the column name.

b Select the column type.

Depending on the column type, you can specify a length or precision, and you can specify a default value for the column.

c If users must enter a value for the column, select the **Not Null** check box.

d If the column is a primary key, select the **Primary Key**check box.

e Click **Next** to continue, or click **Finish** to create the column.

7 (Optional) In the **Constraints** tab, specify constraints that apply to the new column.

a Use the drop-down menu to select the type of constraint: **Foreign key**, **Unique**, or **Check**.

You can create foreign key constraints only if the schema has more than one table.

b Click **Create** to specify the constraint.

c Enter the conditions for the constraint, and click **OK**.

Data Director creates the constraint.

d Click **Next** to continue, or click **Finish** to create the column.

8 (Optional) In the **Auto Vacuum Settings** tab, specify settings for removing stale data from your table,

The default settings work well for most environments. For information about auto vacuum, see the documentation on the Postgres.org site.

9 Click **Finish** to finish table creation.

Data Director creates the table.

Create a View

A view is a subset of related table data. For example, if you have a table that contains the locations of all corporate offices throughout the world, you can create a view of all the offices in Europe, in California, or Brazil. You create views from the **Views** tab in the schema pane.

Prerequisites

Verify that the table on which to create the view exists.

Procedure

1 Log in to Data Director as an organization administrator or user with database privileges and select a database.

2 Click the **Views** tab.

3 Click the plus icon to create a view.

4 Enter the view properties.

   a Enter a unique name in the **Name** field. If the name is case-sensitive, select the **Case sensitive** check box.

   b (Optional) To restrict who can modify the view, use the drop-down menu to select an owner for the view definition.

   c Enter a SQL query to define the view.

   For example, if you are creating a view of your office_locations table named China Offices, you might enter a query similar to the following to select all the office locations in China.

```
select office_name, addr1, addr2, addr3 from office_locations where country="China"
```

5 Click **OK**.

   The view appears in the left pane under the **Views** icon.

What to do next

Examine the data in the view. See "Examine View Data," on page 69.

Examine View Data

A view is a subset of related table data. After you create a view, you can examine the data in the view.

Prerequisites

Verify that a view is available. See "Create a View," on page 68.

Procedure

1 Log in to Data Director as an organization administrator or user with database privileges.

2 In the **Manage & Monitor** tab, navigate to your schema, and click to select it.

3 Click the **Views** tab.

4 Double-click a view to select it.

5 Click the **View Data** tab.

Create a Constraint

Constraints enable you to reduce data entry errors by verifying data before inserting it into a table. Define constraints on schema tables and columns from the Manage & Monitor tab.

You can create constraints when you create a table, or you can add them later. You can create the following types of constraints.

- Check constraint. Limits the values or value range that can be inserted in a column.
- Unique constraint. Ensures that a column or set of columns is unique.
- Primary key constraint. Uniquely identifies each row in a table. There can be only one primary key per table.
- Foreign key constraint. Points to a primary key in another table.

You enter SQL fragments to define a constraint.

Prerequisites

- You are logged in to your organization as an organization administrator or user with database privileges.
- The table on which to create the constraint already exists.

- You expanded the schema in the left pane, then selected Tables.
- The middle pane shows the Tables pane.

Procedure

1 In the left navigation pane, expand the schema that contains the table for which you want to create a constraint.

2 Click Tables. The Tables page appears in the middle pane with a list of the schema's tables.

3 Click the table to select it, then click the gear icon. The Actions drop-down menu appears.

4 Select Create > Constraint. The Constraint drop-down menu appears.

5 Click the type of constraint you want to create. A Constraint dialog appears.

6 Complete the dialog as appropriate for the constraint you're creating, then click OK.

Data Director creates the constraint.

Example: Create a Check Constraint

Check constraints evaluate to a Boolean value. Use check constraints to check whether or not a value entered for a column meets a specific truth-type requirement. For example, suppose that you create a column that must be a positive integer, such as a product price. You can create a check constraint to return TRUE when the product price is greater than 0, and to return FALSE when the product price is less than 0. The check constraint ensures that if a user tries to enter a negative product price, the data entry operation fails with a SQL error.

Enter a check constraint as follows.

1 In the **Constraints** tab, select **Check** from the drop-down menu, and click **Create**.

2 Enter a name for the constraint, such as check_positive_price, in the **Name** field.

3 Enter the constraint in the **Check** field. The constraint should be a simple equation; you do not have to enter SQL. For example, `product_price > 0`.

4 Optionally enter a comment that describes the constraint.

5 Click **OK**. Data Director creates the constraint.

Database Administration

Database administration involves performing routine maintenance for vPostgres databases to ensure efficient use of resources and to achieve optimum database performance. Users with appropriate roles and permissions perform administrative tasks from the Data Director user interface.

Database administration tasks include the following.

- Manage database properties to tune database performance.
- Monitor database statistics such as resource utilization and database performance.
- Manage database backup and restore operations.
- Re-index data as databases change.
- Cluster table data according to an index.
- Recover unused space from tables and indexes (vacuum).

See Chapter 10, “Safeguarding Data,” on page 77 for information about backing up and restoring data. See Chapter 6, “Managing Database Templates,” on page 45 for information about managing database configuration templates and database backup templates.

Manage Database Properties

DBAs and application developers with appropriate privileges manage database properties such as storage allocation, database name, and backup and configuration templates, and manage settings such as auto-vacuum, write-ahead log (WAL), and checkpoint frequency.

Database properties and settings control how the database operates. You can manage and adjust certain settings, such as resource allocation, database connection limits, and whether and how often to perform automatic tasks such as checkpoints and data vacuuming. You can view, but cannot change, database properties such as the UUID and connection string.

Right-click the database name and select **Properties**. The dialog box contains the following tabs.

Basic — The **Basic** tab contains basic information about the database, including its name, JDBC connection string, UUID, storage and point-in-time recovery allocation, version, and tags. You cannot change the database name, JDBC connection string, or the UUID. You can add a new dbowner account from the **Basic** tab.

DB Configuration — The **DB Configuration** tab shows the database configuration settings and their values, such as the number of vCPUs, CPU memory, priority, and reservation, connection limits, IO and WAL settings, and checkpoint and auto-vacuum settings. If you have appropriate privileges, you adjust the database configuration settings from the **DB Configuration** tab.

Backup — The **Backup** tab shows the database backup settings and their values. When you create a database, you choose a backup template. If you have appropriate privileges, you can adjust or override the backup settings from the **Backup** tab.

View or Change Basic Database Properties

Basic database properties include the database name, its UUID, JDBC connection string, storage allocation, its version, and database owner account. You can view and change the values for these properties.

Prerequisites

- You are logged in to your organization as a DBA or application developer with appropriate privileges on the database.
- The database is running.

Procedure

1 Click the organization's **Manage & Monitor** tab and select the database group.

2 Right-click the database name, and select **Properties**.

3 Click the **Basic** tab and view the properties, using the scroll bar as necessary.

You cannot change the database name, UUID, or JDBC connection string.

4 To change the storage allocation for the database or point-in-time recovery, enter new values in the **Basic** tab text boxes.

5 Add or change tags for the database.

a Click **Edit** next to the **Tags** text box.

b Enter tags, one per line.

c Click **OK**.

6 To add a new database owner account, select the **New owner account** checkbox, and click **Edit**.
   a Type the database owner account name and password .
   b Type the password in the **Confirm password** field.
   c Click **OK** to accept the new owner account.
7 Click **OK** to accept your changes.

Upgrade a Database

Upgrade your database to benefit from enhancements to vPostgres.

Prerequisites

- You are logged in to your organization as a DBA or application developer with appropriate privileges on the database.
- The database is running.

Procedure

1 Click the organization's **Manage & Monitor** tab and select the database group.
2 Right-click the database name, and select **Properties**.
3 Click the **Basic** tab.
4 Click **Edit** in the **Version** text box.
   If your database is already the latest version, the Upgrade text boxes are dim.
5 Select the **Upgrade to latest version** check box.
6 Specify a start time.
   Start now is the default.
7 Click **OK**.

The upgrade proceeds.

Change Database Configuration Settings

When your database expands or when usage patterns change, you can adjust database configuration settings to improve performance, provide more storage, and so on.

Prerequisites

- You are logged in to your organization as a DBA or application developer with appropriate privileges on the database.
- The database is running.

Procedure

1 In the organization's **Manage & Monitor** pane, select the database group.
2 Right-click the database name, and select **Properties**.
3 Click the **DB Configuration** tab and view the current settings, using the scroll bar as necessary.

4 Click **Edit**.

The top section of the dialog box shows the database configuration templates. Some settings have drop-down menus from which you select new values, while others allow you to enter new values in text fields. A restart icon marks the settings that require restarting the database to take effect.

- To change resource settings, click the **Resource Settings** tab, and select the check box in the **Override** column.
- To change database settings, click the **Database Settings** tab, and select the check box in the **Override** column.

  You can use the Database Calculator to calculate new settings based on the current database usage and requirements.

5 When you finish, click **OK** to close the **DB Configuration** dialog box.

6 Click **OK** to accept your changes.

7 If you changed settings that require a database restart, right-click the database name, and select **Power > Restart**.

Change Database Backup Settings

As your database usage or backup requirements change, you can adjust your database's backup settings to suit current usage patterns and database recovery requirements.

Prerequisites

- You are logged in to your organization as a DBA or application developer with appropriate privileges on the database.
- The database is running.

Procedure

1 In the organization's **Manage & Monitor** pane, select the database group.

2 Right-click the database name, and select **Properties**.

3 Click the **Backup** tab and view the current backup settings, using the scroll bar as necessary.

4 Click **Edit**.

The top section of the **Backup Settings** dialog box lists the available database backup templates. The bottom section shows the template settings and their current values.

5 Click the backup template that you want to adjust.

6 To adjust the backup template settings, select the check box in the **Override** column and adjust the settings.

7 When you finish, click **OK** to accept your changes.

8 Click **OK** to close the **Properties** dialog box.

Monitor Database Group and Database Statistics

Monitoring database statistics helps you to ensure that your databases run efficiently. The statistics enable you to identify and troubleshoot problem areas, such as running low on resources, that might affect the ability to meet service goals. You monitor database statistics by viewing resource usage and performance data in the **Manage & Monitor** tab.

Data Director keeps statistics for database groups and for the databases within those groups. You must have appropriate permissions on the organization, database group, or database to monitor statistics, run reports, or assign permissions.

Monitor Database Group Statistics

You can view resource usage, allocation, alarms, tasks and events, get reports, view permissions, and view a statistics breakdown for databases within the database group. Monitoring database group statistics helps to ensure that your database groups run efficiently, and lets you identify and troubleshoot problems that can affect performance, such as resource availability. Monitor database group statistics from the **Manage & Monitor** tab.

Prerequisites

- You are logged in to the organization as an organization administrator or user with appropriate permissions on the database group you want to monitor.
- You are in your organization's **Manage & Monitor** tab.

Procedure

1 Click the database group in the left pane.

2 Click the appropriate tab to view the database group statistics you want.

   - In the **Dashboard** tab, view graphical representations of resource usage statistics for the database group.
   - In the **Alarms** tab, view triggered alarms.
   - In the **Tasks & Events** tab, view tasks, events, and their status.
   - In the **Permissions** tab, view roles, assign roles, or grant permissions.
   - In the **Reports** tab, view database statistics reports or revise settings such as sampling intervals.

3 Click the **Databases** tab.

4 Click the database to select it.

5 Click the drop-down menu next to **View**, and select the statistics you want.

Resource Usage	Includes how much memory is in use, how much storage the database's data, backups, and point-in-time recovery log use, the amount of storage available to the database, any tags the database has, and the current database status.
Resource Allocation	(Default) Includes the amount of storage allocated to this database for data and point-in-time recovery, any tags in use, and the current database status.
Performance Statistics	Includes average read/write times and transactions per second.
Version	Includes whether upgrades or patches are needed and the vPostgres and Data Director versions currently running.

Monitor Database Statistics

You can view statistics and access your database activity logs from the **Databases** tab.

Prerequisites

- You are logged in to the organization as an organization administrator or user with appropriate permissions on the database you want to monitor.
- You are in your organization's **Manage & Monitor tab**.

Procedure

1 Click the database group that contains the database you want to monitor.

2 Click the **Databases** tab, and click the database you want.

3 Click the appropriate tab to view database-specific statistics.

- In the **Backup & Recovery** tab, monitor backups and perform recovery operations.
- In the **Alarms** tab, view warnings and alarms, such as health check and application status alarms.
- In the **Tasks & Events** tab, view tasks, events, and their status.
- In the **Reports** tab, view statistics reports or revise settings such as sampling intervals.
- In the **Logs** tab, view the database activity log. You can also download the log from this tab.
- In the **Permissions** tab, view roles, assign roles, or grant permissions.

SQL Management

Managing SQL includes developing and testing SQL queries and monitoring and tuning query performance. You must have appropriate permissions on the schema and database to develop and manage SQL queries. You can manage SQL from the schema page.

Enter and Run a SQL Query

Create and modify SQL queries.

Prerequisites

You are logged in to Data Director as a user with appropriate privileges on the database or schema.

Procedure

1 Click the **Manage & Monitor** tab and select the schema for which you want to manage SQL queries.

2 Click **Enter SQL**.

The SQL dialog box has the following sections.

- Entry pane.
- Output pane. Allows you to examine query output, examine the query's actual run time and cost, view any output messages, and examine query run history.

3 Enter a query in the entry pane.

You can type or modify a SQL query, test the query, and analyze the query's execution plan before running it.

- Type the query in the entry pane.
- Click **Open** to open a SQL script file.

4 Click **Execute** to run the query.

If the query runs successfully, data appears in the **Output** pane.

If the query does not run successfully, Data Director shows the error message in the **Messages** tab.

To cancel the query, click **Cancel**.

View a Query Plan

Viewing a SQL query execution plan enables you to analyze query run time and cost to ensure that your queries run as efficiently as possible.

Prerequisites

- You are logged in to Data Director as a user with appropriate privileges on the database or schema.
- You know how to enter and run a SQL query. See "Enter and Run a SQL Query," on page 75.

Procedure

1 Click the **Manage & Monitor** tab and select the schema for which you want to view a query pane.

2 Select the schema, and click **Enter SQL**.

3 Enter a SQL query in the entry pane, or click **Open** to open a SQL script file.

4 Click **Execute** to run the query.

5 Click **Explain** to view the query plan, run time, and CPU cost.

What to do next

Adjust the SQL query, rerun, and reexamine the query plan as necessary to tune performance.

Safeguarding Data 10

Data Director provides several options for managing backups and recovering databases.

Taking regular backups of your databases is essential to safeguarding your data. Data Director tracks and stores changes for each database on a virtual disk associated with that database. Back up your database to capture the changes, preserves the database, and enables recovering the database and restoring its data after a failure. You can also restore the database to its state at a particular time and replay changes to troubleshoot a problem.

Data Director offers the following features for safeguarding data:

- Manual and automated external and snapshot backups.
- Database recovery from external and snapshot backups.
- Point-in-time recovery

You can define backup retention time and storage allocation. You can use one of the predefined backup templates or create custom backup templates to ensure consistent backups of your databases and to enforce resource limitations.

This chapter includes the following topics:

Backup Strategies

Backup strategies center on your business requirements for protecting your data. Database backup strategies vary according to business requirements and the database environment, such as production, development, or QA.

For example, for a production database with a high transaction volume and business rules that require the highest possible database resiliency, you might define the following backups:

- Take full external backups twice a day.
- Take database snapshots every hour.
- Enable point-in-time recovery to keep a continuous log of all transactions as they occur on the running database.
- Retain your full backups for a month or more.

If your business rules state that you must preserve every transaction, you can specify that the database must shut down if the point-in-time recovery's write-ahead log runs out of space. For a development database where data loss is not a concern, you might take full external backups every week with daily snapshot backups and point-in-time recovery disabled.

You can initiate backups manually (one-time backups) or automatically (recurring backups). Backup methods are snapshots and external (full database backup). You can enable point-in-time recovery. Depending on your business rules, you can set up automated backups and use a combination of backup methods to safeguard data.

You set up automated backups by attaching a database backup template to your database. Backup templates contain backup and recovery settings. You can select a database backup template during database creation or attach a template at a later time. You can also attach a different backup template at any time. If you have sufficient privileges, you can modify the template settings or create custom backup templates. The backup process picks up the latest settings the next time it runs. The modified settings do not affect backups that are in progress.

Using database backup templates ensures that you can take consistent database backups, meet recovery goals, and enforce your business rules. Data Director provides preconfigured database backup templates that provide a range of backup and recovery settings. If you are not sure how much storage your backups will require, start with the most conservative settings. You cannot decrease the backup storage allocation, but you can increase it. Monitor the database activity and the backup size until you have a good idea of the workload and backup space needed, and then adjust the storage amount.

Backup Types

You manage backups and recover data using Data Director snapshot backups, external backups, and point-in-time recovery.

- External Backups on page 79

  External backups are full copies of the database saved to a datastore separate from the database. This section describes the pros and cons of using external backups.

- Snapshot Backups on page 79

  Snapshot backups capture the changes to the database after the snapshot is taken. Snapshots initially use less storage than external backup files and take just a few minutes regardless of database size.

- Point-In-Time Recovery on page 79

  If point-in-time recovery (PITR) is enabled, a write-ahead log (WAL) continuously records every change made to the database while the database is running. In the event of a failure, you can replay the WAL to restore the database to its state at a point in time within the retention period of the database backups.

External Backups

External backups are full copies of the database saved to a datastore separate from the database. This section describes the pros and cons of using external backups.

External backups use about the same amount of storage as the database itself. Because they reside on a separate disk from the database, external backups provide resiliency and benefits such as the following.

- External backups protect against data loss due to failure of the primary data storage device.
- External backup storage is more cost effective than using the primary data storage for backups.
- You can extend the data disk as needed.

The following are points to consider about using external backups.

- External backups can take a long time. Large amounts of data must be copied across devices.
- Each backup uses the full size of the data disk on the backup storage device.

Snapshot Backups

Snapshot backups capture the changes to the database after the snapshot is taken. Snapshots initially use less storage than external backup files and take just a few minutes regardless of database size.

Snapshot backups are stored in files called delta files or delta disks on the same data store as the database.

The following are points to consider about using Snapshot backups.

- Because snapshots reside on the same data store as the database, they do not protect against data loss due to failure of the data storage.
- As the database changes, the changes require more and more space on the virtual disk. That space is generally more expensive than backup storage.
- The recovery process from snapshots is not faster than the recovery process from an external backup.
- If you have snapshots, you cannot extend the data disk.

Point-In-Time Recovery

If point-in-time recovery (PITR) is enabled, a write-ahead log (WAL) continuously records every change made to the database while the database is running. In the event of a failure, you can replay the WAL to restore the database to its state at a point in time within the retention period of the database backups.

The WAL logs are archived and are subject to a retention period that you set. The time range for point-in-time recovery is from the time of your oldest backup to the present. The oldest backup can be an external backup or a snapshot.

By default, PITR is disabled. If you enable PITR, consider the following points.

- Because every change to the database is recorded, PITR requires additional storage. Depending on how large your database is and how many transactions occur during the WAL archive retention time, the amount of storage needed can be large.
- PITR has a performance impact on the database and on Data Director as a whole. The impact depends on the size of the database and the volume of database activity.

Start with a conservative storage allocation. You cannot decrease the storage allocation, but you can increase it. Monitor the size of the PITR logs until you understand the workload and storage needed, and adjust the storage amount.

You can specify whether to suspend the database or automatically increase the log retention period if PITR runs out of space.

When you enable PITR, Data Director creates a baseline external backup. The default retention period is **forever**. You can change the baseline backup's retention period from the database Properties dialog box's **Backup** tab.

Backup Template Settings

Data Director backup templates contain backup settings that use a combination of methods to safeguard data, provide consistent database backups, and enforce limits on resource consumption. You can use the default backup template settings or adjust the settings to suit your business requirements.

Each database backup template contains settings for snapshot backups, external backups, and point-in-time recovery.

External Backup Settings

Frequency How often to take backups. Settings are every 12 hours, daily, weekly, monthly, or never.

Start time Automatic means the system controls the backup start time. If you specify a start time, each external backup will be initiated within two hours of the target start time depending on system load.

Retention How long to keep the external backup. Retention time settings are 1 day, 1 week, 2 weeks, 1 month, 6 months, or 1 year.

Snapshot Backup Settings

Frequency How often to take backups. Settings are every 4, 8, 12, or 24 hours, or never.

Start time Automatic means the system controls the backup start time. If you specify a start time, each snapshot backup will be initiated within 10 minutes of the target start time depending on system load.

Retention Select how long to keep the snapshot backup, or select how many copies of snapshot backups to keep.

- You can retain snapshot backups for 4, 8, 12, 24, or 48 hours.
- The number of copies of snapshot backups that you can keep varies according to the Frequency setting. You can keep from one copy up to as many copies of snapshot backups as are taken in a 24-hour period.

Table 10-1. Snapshot Backup Copies to Keep per Backup Frequency Setting

Take Snapshot Backups Every....	Copies to Keep
4 hours	1-12 copies
8 hours	1-6 copies
12 hours	1-4 copies
24 hours	1-2 copies

Point-In-Time Recovery (PITR) Settings

Enabled or disabled — Enable point-in-time recovery to continuously record each change to the database in a write-ahead log (WAL) while the database is running. In the event of a failure, you can replay the WAL to restore the database to its state at a point in time within the retention period of the database backups.

Recommended point-in-time recovery storage allocation — The recommended storage amount is based on the database size and storage allocation. You can accept the recommendation or enter a different amount.

If storage runs out — Select whether to suspend the database or adjust the point-in-time WAL retention period.

Backup Label

The backup label can be any text that helps you identify the backup. The format is *backup label*-yyyy:*mm:dd:hh:mm:ss-dbname*.

Preconfigured Backup Templates

The preconfigured backup templates enable you to standardize your database backups and enforce resource limitations. Organization administrators and organization users with sufficient privileges can modify the templates' default settings or create custom backup templates.

For more information, see Chapter 6, "Managing Database Templates," on page 45.

Data Director includes the following preconfigured backup templates. Each template has system-controlled start times.

Disabled. — No backups are taken.

Development. — Schedules snapshot backups every 24 hours and external backups each week. PITR is disabled.

Auto. — Schedules snapshot backups every 12 hours with retention period of 24 hours and external backups each day. External backup retention time is 1 month. PITR is enabled. If PITR storage runs out, the available PITR timeline is adjusted and the oldest archived WAL segments are deleted automatically.

Standard. — Same settings as Auto. If PITR storage runs out, the database is suspended.

Maximum. — Schedules snapshots every 4 hours with retention period of 48 hours, and external backups taken every 12 hours with retention time of 1 month. PITR is enabled. If PITR storage runs out, the database is suspended.

Select a Database Backup Template

You can associate a backup template with your database as part of database creation, or you can select a backup template later. Databases must be associated with a backup template to enable scheduling regular backups.

See "Preconfigured Backup Templates," on page 81 and "Backup Template Settings," on page 80 for information on the settings in the preconfigured templates.

Prerequisites

Log in to your organization as a user with at least the following privileges.

- **Use Templates**

- **Create Snapshots**
- **Create External Backups**

Procedure

1 Navigate to your database's Properties window.

   a In your organization, click the **Manage & Monitor** tab.

   b Select your database group, and click the down arrow to display the list of databases.

   c Right-click your database name and select **Properties**.

2 Click the **Backup** tab, and click **Edit**.

3 Click the name of the backup template to associate it with your database.

What to do next

Schedule regular database backups in the Details - Current Backup Configuration pane.

Schedule Regular Database Backups

To set up an automated schedule of backups of your database, you can specify the backup settings in the database's Properties window. You protect your data when you set up an automated database backup schedule.

Prerequisites

Verify that your database is associated with a database backup template. See "Select a Database Backup Template," on page 81.

Log in to your organization as a user with at least the following privileges.

- **Use Templates**
- **Create Snapshots**
- **Create External Backups**
- **Use Template**

Procedure

1 Navigate to your database's Properties window.

   a Click the **Manage & Monitor** tab.

   b Select your database group, and click the down arrow to display the list of databases.

   c Right-click your database name, and select **Properties**.

2 Click the **Backup** tab and click **Edit**.

3 To view the backup templates' configuration settings, click the backup template's name in the Backup Templates pane.

   The backup template configuration settings appear in the Details pane, and **<Current Backup Configuration>** is replaced with the name of the template. If you have the **Manage Backup Templates** privilege, you can override the template settings.

   a Select the check box in the **Override** column.

   b Adjust the settings for each type of backup.

c (Optional) Enable or disable point-in-time recovery.

d (Optional) Specify a backup label.

4 Review your backup settings and click **OK** to confirm.

What to do next

To view a list of backups and the status of each backup, open the database by double-clicking the database name in the **Manage & Monitor** tab. Click the **Backup & Recovery** tab in the middle pane. The list of database backups appears in the Backup List section.

To review the backup schedule, click the **Backup Schedule** link.

To take a one-time manual backup, right-click the database name in the navigation pane and select **Take Manual Backup**.

Recover a Database

Your ability to recover databases depends on scheduling regular backups. You can recover databases from backups taken using Data Director or from external backups taken using utilities such as VMware Virtual Data Recovery (VDR).

Regularly scheduled backups ensure that you can recover your databases and restore your data in the event of system failure or data corruption. See "Schedule Regular Database Backups," on page 82.

The recovered database is a full copy that is independent from any previously taken snapshot backup or clone database. A side-effect of this process is that you have a database that can be resized. Databases with snapshot backups or linked clones cannot be resized.

Prerequisites

Log in to Data Director as a user with appropriate privileges.

- **Create Snapshots**
- **Create External Backups**
- **Delete Snapshots**
- **Delete External Backups**
- **Recover**

Procedure

1 In Data Director, select the organization and click the **Manage & Monitor** tab.

2 In the navigation pane, select your database group, and click the down arrow to show the list of databases.

3 Right-click your database name and select **Recover**.

4 In the Recover dialog, select recovery options, depending on your setup.

5 Click **OK** to start the recovery process.

The database is unavailable while the recovery operation is in progress.

Data Director takes a complete backup of your restored database after database recovery finishes. This post-restore backup becomes your baseline backup.

Import Backups

If the retention period of a backup set has expired, the backup set is no longer in the Data Director backup storage archive. Data Director has no record of such a backup set and does not recognize it. To use an expired backup set, you must import it into Data Director and associate it with a database.

For example, suppose that you archive backup sets to tape just before they expire as part of your disaster recovery policies. You can later restore the archived backup sets from tape to your active system. Use the import backups feature to locate and import the backup sets, and use the imported backup sets to restore your database. The imported backup sets have a retention policy of Forever.

You can use the import database backup feature with a VMware backup solution such as VDR to implement an extended backup and restore solution. See "Use VMware Data Recovery for Backups," on page 84.

Prerequisites

Log in to Data Director as a user with the appropriate privileges.

- **Create Snapshots**
- **Create External Backups**
- **Delete Snapshots**
- **Delete External Backups**
- **Recover**

Procedure

1. In Data Director, click the organization's tab and in the navigation pane, expand the relevant database group to view the databases.
2. Right-click the database for which to import the backup, and select **Import Backups**.
3. Select the backup to import from the list and click **OK**.

The imported backup appears in the list of database backups when the import finishes.

Use VMware Data Recovery for Backups

Data Director packages external database backups as virtual machines. You can use any VMware virtual machine backup technology as an extended backup solution. One of the choices for backups is the extended backup process supported by the VMware Data Recovery (VDR) appliance.

VDR is available as an appliance for VMware vSphere and operates on vSphere resources. See the *vSphere Documentation Center for VMware Data Recovery* documentation.

You can configure VDR to backup the database group's entire Backup resource pool so that VDR backs up newly created backups. When you add a new database group, manually configure VDR to backup the new database group's backup resource pool.

Procedure

1. Install and Connect to the VMware Data Recovery (VDR) Appliance on page 85

   The VDR appliance is an optional appliance that may not be installed in your vSphere system. Verify the VDR appliance installation, install the appliance if necessary, and then connect to VDR.

2. Take External Backups with VMware Data Recovery on page 85

   You can use VMware Data Recovery (VDR) to take external backups of your Data Director virtual machines.

3 Restore a VMware Data Recovery Backup on page 86

Before you can import a VMware Data Recovery (VDR) backup, you must restore the backup in vSphere Client.

4 Import VMware Data Recovery Backups on page 86

After you restore a VMware Data Recovery backup, you can import that backup into Data Director.

Install and Connect to the VMware Data Recovery (VDR) Appliance

The VDR appliance is an optional appliance that may not be installed in your vSphere system. Verify the VDR appliance installation, install the appliance if necessary, and then connect to VDR.

You plan to take external backups of the Data Director virtual machines using VDR.

Procedure

1 Log in to vSphere Client as an administrator.

2 Verify that VDR is installed.

a Click **Home**.

b Check the Solutions and Applications section for the **VMware Data Recovery** icon.

If VDR is not installed, follow the instructions in the vSphere Documentation Center. You must install the client plugin and the appliance.

3 Connect to the VDR appliance.

a Click **Home**.

b In the Solutions and Applications section, click the **VMware Data Recovery** icon.

c In the Welcome page, click **Connect** to connect to VDR.

Take External Backups with VMware Data Recovery

You can use VMware Data Recovery (VDR) to take external backups of your Data Director virtual machines.

Prerequisites

Log in to the vSphere Client with administrator privileges and connect to the vCenter Server system where you installed VDR.

Procedure

1 On the VMware Data Recovery main page, click the **VDR Backup** tab and click **New**.

2 Enter a unique name for the backup in the **Name** text box and click **Next**.

3 Select the backup datastore associated with the database group's resource bundle.

4 Select a backup storage location for your backup from the list of storage devices and click **Next**

5 In the calender, select or deselect days and hours during which the backup can run and click **Next**.

6 Select the retention period for the backup and click **Next**.

7 Review the backup settings and click **Finish** to start the backup.

The backup begins. The process can take some time to complete. When the process finishes, you can see the backup virtual machine in your database group's backup resource pool.

Restore a VMware Data Recovery Backup

Before you can import a VMware Data Recovery (VDR) backup, you must restore the backup in vSphere Client.

Prerequisites

Log in to the vSphere Client with administrator privileges and connect to the vCenter Server system where you installed VDR.

Procedure

1 In vSphere Client, connect to VDR and click the **VDR Restore** tab.

2 Click the **Restore** link.

3 Select the database backup to restore.

   a Expand your database group's resource pool.

   b Expand the Backup resource pool.

   c Select the check box next to the backup to be restored.

   d Click **Next**.

4 Click through the inventory list to select the location for the restored backup (the datastore of the virtual machine and the `data.vmdk` file that you want to restore), and click **Next**.

5 Review the restore settings, and click **Restore**.

The restore begins. The process can take some time to complete. When the process finishes, you can see the restored backup virtual machine in the vSphere inventory.

Import VMware Data Recovery Backups

After you restore a VMware Data Recovery backup, you can import that backup into Data Director.

Prerequisites

Restore a VDR backup.

Procedure

1 In vSphere Client, power off the backup virtual machine.

2 Log into Data Director as an organization user with database backup and restore privileges for the database that you want to restore.

3 In the **Manage & Monitor** tab, expand the database group, then select the database to restore.

4 Right-click the database name and select **Import Backups**.

5 Select the backup to import.

Data Director imports the backup.

Database End of Life and Backups

When you decommission and delete a database, you decide whether to retain its backup files. The decision is based on your site's policies and whether you might need the database in the future.

When you delete a database, you can retain all external backups. The backups expire at the end of the normal retention period. It is good practice to take a final backup of a database and specify the final backup's retention period before you delete a database. If you retain the external backups, the snapshots and the executable instance of the database are deleted. If the deleted database had point-in-time recovery enabled, all the archived write-ahead log (WAL) segments are deleted as well. This means that the only way to recover the database is by using the external backups. You cannot recover the database using snapshots or point-in-time recovery.

If you do not retain the external backups, the database and its associated backups, snapshots, and WALs are deleted. In addition, the database resources are released, and the database cannot be restored.

Monitoring the Data Director Environment 11

System administrators can examine current resource usage, monitor events and alarms, view and download reports about their environment, and create diagnostic packages for individual databases and for the system itself. Organization administrators can examine how different database groups and databases use resources, and can view and monitor events and alarms for their organization.

This chapter includes the following topics:

Explore Monitoring Customization and Filtering

Customize your monitoring setup to find information quickly. Some of the customization and filtering tasks are the same for both system administrators and organization administrators.

You can explore how to optimize screen areas and how you can quickly find information by using filters. Filtering is not supported in all panels.

Prerequisites

Log in to Data Director, as the system administrator or as an organization administrator.

Procedure

1 Click the **Manage & Monitor** tab and click **Reports**.

2 Click **Summary Reports** or **Time Interval Reports** to display only one type of report and view the icons above the filter options.

Icon	Description
	The gear-shaped **Action** icon lets you choose an action. Available actions differ for different panels. For example, the **Reports** panel lets you download all reports or to download a selected report.
	The blue **Reload** icon lets you redisplay the current page.

3 Explore the Filter text box in the right corner.

a Type a search term to search the current items.

For example, search all tasks for the term Delete.

The search includes the currently displayed items and the list of available items.

b To filter which columns are searched, click the down-triangle, deselect checked boxes for columns you do not want to search, and click **OK**.

When you perform the next search, only the selected columns are searched.

4 Explore how to view a panel.

Action	Action
To reduce a panel	Click the down-facing triangle to the left of the panel name to reduce the panel to its title.
To expand and shrink a panel	Click the expand icon ▭ to expand a panel, and click the shrink icon ▭ to return the panel to its original size and position.
To close a panel	Click the X icon (✕).
To open a closed panel	Click the **Customize** button in the dashboard's top right corner to select the panel's check box.

5 Explore the sidebar.

The sidebar can include a panel for tasks, a panel for alarms and, in certain contexts, an SQL window. Task and alarm information is limited to the most recent tasks and alarms. You can collapse the side bar with the right-facing triangles ⏩ and expand it with the left-facing triangles ⏪.

What to do next

System administrators can see "Monitoring for System Administrators," on page 90. Organization administrators can see "Monitoring for Organization Administrators," on page 95.

Monitoring for System Administrators

The vFabric Data Director interface includes a set of monitoring and diagnostic tools for system administrators. Administrators can see system health information; explore how organizations use resources; view system-level events, alarms, and reports; and generate diagnostic packages for individual databases and for the system itself.

The information Data Director displays for system administrators differs from the information available to organization administrators. System administrators can use the following resources for monitoring overall system health.

Dashboard Tab

In the **Dashboard** tab, system administrators can see system health information and a system overview. Overview information includes the number of users, number of resource bundles, and number of organizations. Administrators can also see the total CPU and memory capacity and the total database storage and backup storage allocation for this instance of vFabric Data Director. The

	dashboard includes information about the top five consumers of resources. Administrators can customize the **Dashboard** to change the sampling and to add or remove the information displayed. See "Explore System Dashboard Customizations," on page 91.
Manage & Monitor Tab	In the **Manage & Monitor** tab, system administrators can view alarms and define new alarms, view tasks and events, and configure, display, and download reports. System administrators can display reports for the organization or for a resource bundle and filter the sampling interval, time range, and other fields. See "Explore Monitoring Customizations for System Administrators," on page 92.
Administration Tab	In the **Administration** tab, system administrators can create a diagnostic package for one or more databases and for the system itself. Diagnostics packages are for use by VMware Support.
Tasks and Alarms Side Bar	The tasks and alarms side bar, in the right panel of the main page by default, displays recent tasks and alarms. The **Manage & Monitor** tab includes more details about tasks and alarms.

Explore System Dashboard Customizations

Exploring system dashboard customizations allows you to see available options. You can customize the dashboard to suit your needs.

The system administrator dashboard differs from the organization administrator dashboard. See "Explore Organization Administrator Dashboard Customization," on page 96 if you are an organization administrator.

Prerequisites

Log in to Data Director as a user with system administrator privileges.

Procedure

1 Click the **Dashboard** tab.

2 Click the link for **Organization Stats** or **Resource Bundle Stats**.

- Click **Organization Stats** to evaluate resource usage for the top five organizations.
- Click **Resource Bundle Stats** to evaluate resource usage for the top five resource bundles.

3 Review the Overview panel.

The panel displays information about all items in the system, about the total CPU and memory capacity, and about database storage and backup storage allocation. You can click an object to display it.

4 Review the System Health panel below the Overview panel.

The System Health panel gives you access to the status of the different servers, systems, and networks.

- A green icon indicates no problems exist.
- A yellow icon warns of potential problems.
- A red icon indicates that a problem exists. vSphere services such as HA remedy the problem but certain tasks cannot be performed while the icon is red.

5 View the Top 5 CPU Usage panel.

You can customize the view to show information for 24 hours, 3 days, or 1 week.

6 To see the organization or resource bundle for which the information is displayed, click the name of the organization or resource bundle.

7 Change the sampling by choosing from the **Sampling** drop-down menu.

Option	Description
Average	Average value for the sampling period.
Minimum	Minimum value for the sampling period.
Maximum	Maximum value for the sampling period.

What to do next

Customize the dashboard to meet your needs.

Explore Monitoring Customizations for System Administrators

Explore monitoring customizations to learn about available options. You can then customize the **Manage & Monitor** tab.

The system administrator **Manage & Monitor** tab differs from the organization administrator **Manage & Monitor** tab. See "Explore Monitoring Customizations for Organization Administrators," on page 97 if you are an organization administrator.

Prerequisites

Log in to vFabric Data Director with system administrator privileges.

Procedure

1 Click the **Manage & Monitor** tab in the Data Director client.

2 Click **Organizations** to display all organizations and the resource allocation for each organization.

You can click an organization name to display details about the organization.

3 Click **Alarms** in the panel on the left, and select either the **Triggered alarms** tab or the **Definitions** tab.

Option	Description
Triggered alarms	Includes alarm severity, description, definition, and the target object. To acknowledge an alarm, right-click the **Acknowledged By** text box and select **Acknowledge**. If more than one system administrator monitors Data Director, the acknowledgement mechanism allows them to communicate that someone knows about the alarm.
Definitions	Includes the trigger type (event or performance), the alarm name and alarm trigger, where the alarm was define, the object the alarm is monitoring, who defined the alarm, and the alarm status. You can create custom alarms from this tab. See "Create a Custom Alarm," on page 100.

4 Click **Tasks & Events** in the panel on the left, and click either **Tasks** or **Events**.

Option	Description
Tasks	Tasks are scheduled system activities requested by the system or a user, for example, Create database and Repair database. A task can succeed or fail. The **Tasks** tab includes information about the target and the user who initiated the task. Information such as related events appear in the **Details** panel below the Tasks table. A task can have no related event or one or more related events. Click the event in the **Details** panel to view the event in the **Tasks & Events** tab. For certain tasks, a right-button menu allows you to cancel or retry the task.
Events	Events are records of user actions or system actions. For example, the system logs when a user logs in to Data Director, or when a database is repaired. Events can be of type info, error, or warning. Use **Hide Info Events** to limit the choices. You can filter the **Events** pane to show only some of the columns. Information such as related tasks appear in the **Details** panel below the **Events** table. An event can have no related tasks, or one related tasks. Click the task in the **Details** panel to view the task in the **Tasks & Events** pane.

5 Click **Reports**, and click **Summary Reports** or **Time Interval Reports**.

6 Use the filter options below the **Action and Reload** icon to customize what the system displays.

The customization steps depend on the report type.

Report type	Action
Summary Report	a Select **Resource bundle** or **Organization** from the **Type** drop-down menu and click **Browse**. b Select the specific resource bundles or organizations for which you want to generate a report. By default, no objects are selected. c Select a time range, or click **Customize** to configure a custom time range. d Select **Compute & Network** or **Storage** to focus the report. e Click the **Filter** button to filter the report. f Click the **Action** icon and select **Download** to download the report.
Interval Reports	a Select **Resource bundle** or **Organization** from the **Type** drop-down menu and click **Browse**. b Select the resource bundles or organizations for which you want to generate a report. c Select a sampling interval. d Select a time range or click **Customize** to configure a custom time range. e Select **Compute & Network** or **Storage** to focus the report. f Select the sampling mechanism from the **Sampling** drop-down menu. g Click the **Filter** button to filter the report. h Click the **Action** icon and select **Download** to download the report.

7 Click **Resource Bundles** or **Datastore Usage** to display all resource panels or all data stores and corresponding usage information.

What to do next

Customize the **Manage & Monitor** options, or customize and download reports.

Create, Download, and Delete Diagnostics Packages

Diagnostic packages are sometimes requested by VMware Support to help resolve a problem. System administrators can create diagnostics packages, download them for analysis, and delete them to save storage space.

Prerequisites

Log in to Data Director with system administrator privileges.

Procedure

1 Click the **Administration** tab and click **Diagnostics**.

By default, the diagnostics page is empty. Diagnostic packages are created on demand.

2 Click the green plus sign and select **Create**.

3 Provide information about the diagnostics package and click **OK**.

Option	Action
Database	Click **Add** and select the database for which you want to generate a diagnostics package.
Include system diagnostics	Select the check box to include system diagnostics. You can select just the system diagnostics without selecting a database.
Time range	Select a time range from the drop-down menu.

Data Director creates the package displays it.

4 To download a package, select **Download** from the **Actions** menu, and specify the download location when prompted.

5 To delete one or more packages, select the package or packages and select **Delete** from the **Actions** menu.

What to do next

Send the diagnostic package to VMware Support for analysis.

Understanding Cluster Alarms

The vSphere Cluster on which Data Director is installed must meet several configuration requirements. If the requirements are not met, or if a compatible cluster is modified to no longer be compatible, Data Director displays one or more alarms.

When you create a resource bundle, you can use resource pools only if the cluster in which you create the resource pools is compatible with Data Director. See the *Data Director Installation* documentation for initial cluster setup. After installation, you can customize the cluster. However, only certain customizations are compatible with Data Director. See *vFabric Data Director Installation Guide*.

Data Director requires the following settings and generates an alarm if you change them.

- vSphere DRS and vSphere HA are enabled.

CAUTION If you disable vSphere DRS, all resource pools in your environment become unusable. You must recreate the resource pools.

- Host monitoring is enabled
- **VM Monitoring** is set to **Virtual Machine and vApp Monitoring**.
- **Default VM Restart Priority** is not disabled.

Data Director also generates an alarm if a vSphere administrator makes the following changes to the cluster.

- Admission control for the cluster is disabled.
- The cluster's default VM monitoring settings are changed to be too low.
  - Heartbeat failure time less than 30 seconds.
  - Minimum uptime less than 120 seconds.
  - Maximum number of resets less than 3.
  - Time window for maximum number of resets less than 3600 seconds (1 hour).

If you encounter a cluster-related alarm, contact the vSphere system administrator and show them this information to resolve the problem.

Monitoring for Organization Administrators

The vFabric Data Director interface allows organization administrators to view CPU and memory utilization across the organization and to view database storage breakdown and utilization. Administrators can also monitor the databases and database groups, see events and alarms, create alarms, and generate and download reports.

The information that Data Director displays for organization administrators differs from the information available to system administrators. The panel on the left displays a hierarchy.

- Organizations are the top level.
- Expand an organization to display its database groups.
- Expand a database group to display its databases.

When you select an item in the hierarchy, the right panel displays information about it if you have permission to view the information.

Organization administrators use the following tabs and panels to monitor the organization.

Dashboard Tab

In the **Dashboard** tab, organization administrators can customize the resource usage information displayed to them. Administrators close any of the panels, and click **Customize** to include the panel in the dashboard again. See "Explore Monitoring Customizations for Organization Administrators," on page 97.

Manage and Monitor Tab

In the **Manage & Monitor** tab, organization administrators select tabs to manage and monitor parts of the organization.

Table 11-1. Manage and Monitor Tab

Tab	Description
Databases tab	View existing databases and their attribute and status. Create databases.
Database Groups tab	View existing database groups and their attributes. Create database groups.
Alarms tab	Includes a list of all alarms triggered so far. The system displays alarms for certain events. Administrators can create custom alarms, which are then included in the **Alarms** panes. See "Create a Custom Alarm," on page 100.

Table 11-1. Manage and Monitor Tab (Continued)

Tab	Description
Tasks and Events tab	Allows organization administrators to display information about all tasks and information about all events. You can display events of only a certain type, and you can filter tasks and events to drill down to information you really need. See "Explore Monitoring Customizations for Organization Administrators," on page 97.
Tags tab	Users can create tags and use them to tag the databases. Tags categorize databases and make search easier.
Reports tab	Allows organization administrators to customize the reports pane and to create and download custom reports.
Permissions tab	Allows organization administrators to view currently defined users and roles and the privileges granted to a selected role.

Which alarms, events, and tasks the system displays depends on the current selection in the left panel. For example, with a database group selected, clicking the **Events** tab displays events in that database group and its databases.

Tasks and Alarms Side Bar

The Tasks and Alarms side bar, in the right panel of each tab, displays recent tasks and alarms. "Explore Monitoring Customization and Filtering," on page 89 explains how you can collapse and expand the side bar.

Explore Organization Administrator Dashboard Customization

Exploring organization administrator dashboard customizations allows you to see available options. You can then customize the dashboard to suit your needs.

The organization administrator dashboard differs from the system administrator dashboard. See "Explore System Dashboard Customizations," on page 91 if you are a system administrator.

Organization administrators can use the main dashboard to monitor the organization and its database groups. Organization administrators can use the Database dashboard available from the **Manage & Monitor** tab to monitor databases. See "Explore Database Monitoring," on page 99.

Prerequisites

Log in to Data Director as a user with organization administrator privileges.

Procedure

1 Click the **Dashboard** tab.

2 Explore the Resource Bundles panel.

The panel displays information about each resource bundle that has been assigned to the organization and lets you see CPU and memory allocation and reservation, database and backup storage, and the number of database groups that use the resource bundle.

3 Explore one of the panels that includes lines or histogram bars, for example the **%CPU Utilization** panel or the **Database Storage Breakdown** panel.

   a Customize the view to show information for 1 hour, 24 hours, 3 days, or 1 week.

   b Move the cursor over a histogram bar or a line to view information, and click the item for details.

What to do next

Customize the dashboard to meet your needs.

Explore Monitoring Customizations for Organization Administrators

Organization administrators can customize the monitoring pane to view information relevant for their current needs.

Prerequisites

Log in to Data Director with organization administrator privileges.

Procedure

1 Click the **Manage & Monitor** tab in the Data Director client.

2 In the left panel, select the item you want to monitor.

Option	Description
Selecting a database group	Displays resource information for that database group.
Selecting a database	Displays resource information for that database.

3 Click the **Databases** tab and select **Database list** or **Statistic breakdown**.

Option	Description
Database list	Displays information about each database in the organization.
Statistic breakdown	Displays the top resource or top performance database. You can customize the view .

You can select a database and add it to **Favorites**, but you will be prompted for the owner user name and password to make changes to the database.

4 Click the **Dashboard** tab to display the Organization Resource Usage dashboard.

You can customize the dashboard by clicking the **Customize** button, or customize individual panels on the dashboard. By default, the following information is included.

Panel	Description
Resource Bundles	Displays the databases, associated resource bundles, currently allocated CPU and memory reservations, and currently allocated and free storage. This panel allows administrators to evaluate whether they have additional resources to allocate to a new or existing database group or database.
Recent Alarms	Displays recent system alarms and user-defined alarms.
CPU Utilization, Memory Utilization	Displays the CPU and memory utilization, allowing the administrator to see usage and usage patterns. The view can be set to display the last 1 hour, 24 hours, 3 days, or 1 week.

Panel	Description
Database Storage Breakdown, Backup Storage Breakdown	Pie charts that show the current state of storage and backup storage, including storage that is allocated, storage that is used, and storage that is allocated but not used. Placing the cursor inside a field of the pie chart displays information about that field.
Database Storage Utilization, Backup Storage Utilization	Charts that show storage utilization over the selected amount of time (1 hour, 24 hours, 3 days, or 1 week). Placing the cursor over a line displays information about that line.

5 Click **Tasks & Events**, and click either **Tasks** or **Events**.

Option	Description
Tasks	Tasks are scheduled system activities requested by the system or a user, for example, Create database and Repair database. A task can succeed or fail. The **Tasks** tab includes information about the target and the user who initiated the task. You can filter the **Tasks** panel to show only some of the columns.
Events	Events are records of user actions or system actions. For example, the system logs when a user logs in to Data Director, or when a database is repaired. Events can be of type info, error, or warning. Use the **Type** drop-down menu to display the type of event you are interested in. You can filter the **Events** column to show only some of the columns.

6 Click **Reports**, click **Summary Reports**, and customize the pane by using the filter options.

a Select **Database** or **Database Group** from the **Type** drop-down menu and click **Browse** to select a database or database group.

By default, no objects are selected. You can select more than one object.

b Select a time range or click **Customize** to configure a custom time range.

c Select **Compute & Network** or **Storage** to focus the report on networking or storage information.

d Select the sampling mechanism from the **Sampling** drop-down menu.

e Click the **Filter** button to filter the report.

Click the **Action** icon and select **Download** to download the report.

Customizing the pane does not customize the report itself.

7 Click **Reports**, click **Time Interval Reports**, and customize the pane by using the filter options.

a Select **Organization**, **Database Group**, or **Database** from the **Type** drop-down menu and click **Browse** to select the object you want to generate a report for.

You can select more than one object.

b Select a sampling interval.

c Select a time range or click **Customize** to configure a custom time range.

d Select **Compute & Network** or **Storage** to focus the report on networking or storage information.

e Click the **Filter** button to filter the report.

f Click the **Action** icon and select **Download** to download the report.

8 Click **Resource Bundles** or **Datastore Usage** to display all resource panels or all data stores and corresponding usage information.

What to do next

To monitor specific databases, see "Explore Database Monitoring," on page 99.

Explore Database Monitoring

The main organization dashboard allows administrators and other privileged users to monitor the organization and its database groups. Administrators can also monitor databases from the **Manage & Monitor** tab.

Database information provided by Data Director allows administrators to check whether a database is in use, check on the backup status of the database, see errors and alarms, and check resource allocation. Database administrators might find that backups or other tasks do not finish and can alert the organization administrator, who can allocate more resources. Organization administrators can allocate or remove resources, schedule backups, and perform other database-specific tasks.

Review tasks that administrators can perform on databases that cannot be performed on database groups and organization. For a general exploration of managing and monitoring for organization administrators, see "Explore Monitoring Customizations for Organization Administrators," on page 97.

Prerequisites

Log in to Data Director with monitor privileges on the database. You do not need login privileges for a database to monitor the database. You do need database login privileges to view information in the **Processes and Locks** tab, to see uptime information, and to see schemas, tables, DB login users, and user data size in the overview.

Procedure

1 Click the **Manage & Monitor** tab.

2 Open the organization and the database group, and select a database.

The right panel displays database information.

3 Click the **Dashboard** tab to examine the Overview panel.

The information includes details about the database contents, external backups and snapshots, recent alarms, and any clones that administrators might have created for the database.

4 Click **Resource Usage** to view and customize resource usage information.

5 Click **Database Stats** to display database information such as transactions per second, total IO per second, and so on.

By default, the information is updated in real time. You can change individual panels to use a less frequent update cycle.

6 Click **Processes and Locks** to view database processes and database locks.

You can show details about processes and locks and kill processes if you have permissions to do so.

NOTE You are prompted for database credentials if you select **Processes and Locks**.

---

7 Click **Custom**, and click the **Customize** button to create a custom dashboard for this database.

8 Click the **Logs** tab to display the postgresql log for the database.

You can show the next or previous 100 lines, search the log, and add a filter to search only the specified columns in the log. Filters can be added only if the database is running.

9 Click the **Action** icon and choose **Download File**.

10 Use the **Alarms** tab to view alarms and the **Tasks & Events** tabs to view tasks and events.

The information in the tabs is more detailed than the information in the side bar.

What to do next

Manage databases as discussed in Chapter 7, "Managing Databases," on page 51.

Working with Alarms

Data Director displays system-defined alarms to system administrators and organization administrators. Data Director also allows administrators to create custom alarms and to delete and disable alarms.

Create a Custom Alarm

Custom alarms allow you to display information in the Alarms panel or to send email if certain conditions are met. For each alarm, you can specify a name and description and a trigger.

You can create a custom alarm for the items below the selected item in the hierarchy. For example, if you create an alarm at the organization level, you can monitor database groups and databases. If you create an alarm at the database group level, you can monitor databases.

Prerequisites

Verify that you have permissions to create alarms for the object in which you create the alarm.

Procedure

1 Click the **Manage & Monitor** tab, click **Alarms** in the left pane, and click **Definitions**, and click the plus sign to start the Create Alarm Definition wizard.

2 Type a name and description.

3 (Optional) Change the status to Disabled so that the alarm is not enable immediately and click **Next**.

4 Select the trigger, which is defined by the following options, and click **Next**.

Option	Description
Object type	The object to monitor. The alarm is triggered when trigger conditions on the monitored object are reached.
Trigger type	Select **Performance** to trigger an alarm when the object moves beyond a specified warning or critical threshold. Select **Event** to trigger an alarm when a system event occurs.
Trigger	Select from the available options. Options differ depending on the trigger type (Performance or Event).
Severity	Select **Warning** to have a yellow warning icon associated with the alarm. Select **Critical** to have a red critical icon associated with the alarm. The icon appears in all displays.
Condition	For performance triggers, specifies whether the alarm is triggered when the value is below the current threshold or above the current threshold. For example, you might want a warning if Aborted transactions per second is more than (above) a specified number, or if Committed transactions per second is less than (below) a specified number.
Warning threshold	Threshold at which you want warning actions to take place. You can specify different actions for warning and critical threshold problems.
Repetition frequency	Available only for performance alarms. When the condition that triggers the alarm remains true, a second alarm is generated based on the repetition frequency. If alarm actions such as sending an email are specified, the actions are performed again.

5 Select **Send email** to send an email when the alarm is triggered.

The email is sent only if the SMTP parameters were set correctly during Data Director setup.

6 Click **Finish** to complete definition of the alarm.

The alarm appears in the Alarms panel when it is triggered even if you leave **Do nothing** selected. If you selected **Send email**, an email is sent as well.

Example: Custom Alarm that Monitors Resource Bundles

The following example illustrates how you can create a custom alarm that monitors resource bundles. The alarm sends an email when free space is below a certain threshold.

Log in as the system administrator and

1 Select the **System** tab.

2 Start the Create Alarm Definition wizard.

   a Click the **Manage & Monitor** tab.

   b Click **Alarms** in the left pane.

   c Click **Definitions**.

   d Click the plus sign.

3 Type a name and description and click **Next**.

4 Select the trigger.

Field	Value
Object Type	Resource Bundle
Trigger Type	Performance
Trigger	Database storage usage percentage
Condition	Above
Warning Threshold	80%
Critical Threshold	90%

When you complete the alarm, the result is a yellow (warning) alarm in the Alarms pane when free space drops below 20% and a red (critical) alarm when free space drops below 10%.

What to do next

You can disable or delete alarms. See "Delete or Disable a Alarm," on page 101.

Delete or Disable a Alarm

Administrators can delete or disable an alarm if they do not find it useful. You can delete and disable both system-defined alarms and custom alarms. If you disable an alarm, you can enable it again. If you delete an alarm, it is permanently removed from the system.

Prerequisites

Log in to Data Director as a user with permissions to delete alarms at the level where you want to delete them. A user who has monitor privileges on the object the alarm is monitoring can update or delete the alarm.

Procedure

1 Open the Alarm Definitions table.

- If you are a system administrator, click the **Manage and Monitor** tab, click **Alarms**, and click **Definitions**.
- If you are an organization administrator, click the organization or click the resource group that you want to delete and alarm for, click the **Alarms** tab, and click **Definitions**.

2 Disable or delete an alarm.

- To delete an alarm, right-click that alarm and select **Delete**.
- To disable an alarm, right-click that alarm and select **Disable**.

Managing Licenses 12

Data Director offers evaluation and permanent product licenses. System administrators have fine-grained control of licenses and license assignment using the Data Director Administration tab's License pane.

This chapter includes the following topics:

License Management Overview

System administrators can manage Data Director product licenses from the Licensing pane of the Data Director **Administration** tab.

Evaluation Evaluation licenses let you use Data Director for a limited period of time at no cost. The evaluation product is fully functional, but support is not available.

Permanent Permanent licenses provide full product functionality and never expire. Support and Subscription (SnS) licenses are required for all permanent software licenses.

When you buy permanent licenses, they replace any evaluation licenses you might have. After you upgrade to permanent licenses, only your permanent licenses appear in the **Licensing** pane. After you add a permanent license, you can no longer add evaluation licenses.

The following are the vFabric Postgres (vPostgres) database usage types.

Non-Production Use Includes internal development, quality assurance, proof of concept, or other testing purposes.

Production Use Includes use of vPostgres databases in any manner other than Non-Production Use.

You can change the vPostgres database usage type at any time. You can create databases at will as long as resources are available.

You cannot remove the last remaining permanent or evaluation license. Data Director provides vPostgres Non-Production Use licenses by default. The Non-Production Use licenses cannot be removed.

Purchasing Support

You can optionally purchase vFabric Developer Support for Production Use and Non-Production Use vPostgres databases.

SnS licenses are required for Data Director. SnS licenses are also required for vPostgres: Production Use databases.

- Basic SnS licenses provide weekday support for test, development, and non-critical deployments.
- Production SnS licenses provide focused, 24-hour support for production environments.

See the VMware Support Offerings Web site for details about the SnS licensing options. Manage your SnS licenses through standard VMware support processes.

For details about your licensing arrangement, contact your VMware representative.

Roles and License Management Tasks

The license management tasks that you can perform depend on your role, as shown in the following table.

Table 12-1. License tasks

Task	System administrator	Organization administrator
View licenses	Yes	Yes
Add licenses	Yes	No
Remove Licenses	Yes	No
Change database usage type	Yes	Yes

System administrators and organization administrators have licensing privileges by default, and can grant the **View and Manage Licenses** privilege to users. For example, an organization administrator can grant the **View and Manage Licenses** privilege to a database user. The database user can then view database license information and change the database usage type.

Organization administrators can view only their organization's license information.

Only users with the system administration role can add and remove licenses.

Counting Data Director Licenses

You count Data Director licenses according to the number of database virtual machines (DBVMs) and vPostgres databases in use. If you use more licenses than you have purchased, you can purchase additional licenses or change how licenses are used.

- Data Director requires one license per database virtual machine, regardless of the number of vCPUs, and regardless of whether the virtual machine is powered on or off.
- Backup database virtual machines do not count towards your license total.

vPostgres requires one license per virtual machine, up to two vCPUs per license. For example, a four-vCPU vPostgres virtual machine requires two licenses.

The following table shows the number of licenses required given the number of vCPUs for a single vPostgres virtual machine.

Table 12-2. Licenses Required for a Single vPostgres Virtual Machine

Number of vCPUs for a Virtual Machine	Number of Licenses
1-2	1
3-4	2
5-6	3
7-8	4

You can view the total number of vPostgres virtual machines and vCPUs in the **Licensing** pane of the Data Director **Administration** tab.

About Evaluation Licenses

Evaluation licenses offer full use of Data Director and vPostgres databases at no cost for a limited period of time (usually 90 days).

When you use the evaluation version of Data Director, a message appears when you log in that shows how many days remain in the evaluation period. When the evaluation period expires the following functionality becomes unavailable.

- Backup
- Clone database
- Create database
- Import database
- Repair database
- Restart database
- Restore database
- Start database

You can upgrade evaluation licenses to permanent ones. When you purchase permanent licenses, VMware issues one permanent license key per SKU. As a Data Director system administrator, you add the permanent license key(s) in the **Manage & Monitor** tab's **Licensing** pane. See "Add License Keys," on page 105.

Adding permanent licenses upgrades your evaluation licenses to permanent ones.

- Permanent licenses replace the evaluation licenses.
- Only permanent licenses appear in the license list.

Add License Keys

You can use Data Director only if enough licenses are available. System administrators add license keys in the **Licensing** pane of the **Administration** tab.

Prerequisites

- Obtain license keys for your Data Director products from your VMware representative.
- Log in to Data Director as a system administrator.

Procedure

1 Click the **Administration** tab.

2 In the left pane, expand Settings and click **Licensing**.

3 In the **License Keys** section, click the plus (+) icon.

4 Enter product license keys in the **License keys** text box (one per line), and click **Add License Keys**.

5 (Optional) Enter a label for your license keys in the **Optional label for license keys** text box.

6 Click **OK**.

The licenses appear in the license key list.

View License Information

You can view information about your Data Director and vPostgres licenses in the **Licensing** pane of the **Administration** tab. The information helps you determine whether you need additional licenses if you increase the size of your Data Director installation.

You can monitor product license usage and assignments. What you view depends on whether you are a system administrator, organization administrator, or organization user with the **View and Manage Licenses** privilege.

Procedure

1 Log in to Data Director and click the **Administration** tab.

2 In the left pane, expand **Settings** and click **Licensing**.

3 View the license information.

User	Description
System administrator	If you are a system administrator, view license information as follows. ■ Click the **Licensing** tab to view your product license and license key information. ■ Click the **Usage** tab to view database license usage.
Other users	If you are an organization administrator or a user with the **View and Manage Licenses** privilege, view license information as follows. ■ View your license usage in the **Summary** section of the **Licensing** pane. ■ View your databases and their usage types in the **Databases** section of the **Licensing** pane.

Change the Database Usage Type

You can designate databases for Production Use or Non-Production Use. The default database usage type is Non-Production Use. You can change the database usage type at any time from the **Licensing** pane of your organization's **Administration** tab.

- Use Non-Production Use databases for application development and testing purposes.
- Use Production Use databases for real-time, production applications.

Prerequisites

- Log in to Data Director as an organization administrator or as a user with manage licensing privileges.

Procedure

1 Click the **Administration** tab.

2 In the left pane, expand **Settings** and click **Licensing**.

3 In the **Databases** section, select the database that you want to change.

You can select multiple databases to change their usage types in one operation.

4 Right-click, and select the usage type.

5 Click **Yes** to confirm the change.

The updates appear in the Type column and the Summary pane.

Remove License Keys

To reallocate licenses in the Data Director environment, system administrators can remove license keys in the Licensing pane of the **Administration** tab.

Procedure

1 Log in to Data Director as a system administrator and click the **Administration** tab.

2 In the left pane, expand **Settings** and click **Licensing**.

3 In the **License Keys** section, click the license you want to remove.

4 Right-click the license, and click **Remove**.

5 Click **Yes** to remove the license.

The license no longer appears in the license keys list.

Reconfiguring Data Director Networks 13

During installation, you set up the networks that carry the different types of Data Director network traffic. You can reconfigure Data Director networks, for example, to improve throughput or provide better isolation for certain types of traffic, using vSphere Client. You can reconfigure certain network adapter settings in Data Director.

Data Director has the following types of network traffic.

- Web Console Network
- vCenter Network
- DB Name Service Network
- Internal Network
- DB Access Network

After you change network configuration using vSphere Client, you can confirm the changes in the Data Director **Network Setup** page. You might have to reenter your network settings after changing the mapping for the Web Console network, Web Console network adapter, and vCenter network. You might have to do this if you use static IP addressing. Network configuration information that you might have to reenter includes the netmask, gateway, and DNS Server information.

This chapter includes the following topics:

Change the vCenter IP Address

At certain times in the vCenter Server lifecycle, administrators have to change the vCenter Server IP address. If the IP address change is necessary, you must first power off the Data Director vApp. If you update the vCenter Server IP address while Data Director is running, Data Director cannot communicate with the vCenter Server system.

Because of an issue reported in the release notes, you must power off the Data Director vApp and also remove the vCenter Extension service and add it back.

Procedure

1 Log in to vSphere Client as an administrator.

2 Right-click the Data Director vApp and select **Power Off**.

3 Select **Administration > vCenter Server Settings**.

4 Click **Runtime Settings**.

5 Update the IP address in the **Managed IP Address** field.

6 Remove the vCenter Extension Service and add it back.

   a Right-click the Management Server virtual machine, select **Edit Settings**, and click the **vServices** tab.

   b Select **vCenter Extension Installation** and click **Edit**.

   c Select **<No Provider>** as the provider and click **OK** twice to exit.

   d Select the Management Server virtual machine again.

   e Click **Edit Settings** and select the **vServices** tab.

   f Select **vCenter Estension vService** as the provider and click **OK** twice to exit.

7 Power on the Data Director vApp.

The Data Director vApp communicates with the vCenter Server system using the updated vCenter Server IP address.

Reconfigure the Web Console Network Mapping or Network Adapter

During installation, you configure the Web Console Network mapping and the Web Console Network adapter. You can later reconfigure the Web Console Network by editing the network settings of the Management Server virtual machine in vSphere Client.

After you reconfigure the network mapping or adapter in the vSphere Client, you might have to reenter the network information from the Data Director UI. Reentering the settings for the Web Console network, Web Console network adapter, and vCenter network, is usually necessary if you use static IP addresses. Network configuration information that you might have to reenter includes the netmask, gateway, and DNS Server information.

Prerequisites

Review the information about network settings in Data Director in the *vFabric Data Director Installation Guide*.

Procedure

1 Log in to the vSphere Client as an administrator and select **Inventory > Hosts and Clusters**.

2 Select and expand the Data Director vApp, right-click the Management Server virtual machine, and select **Power > Power Off**.

3 Right-click the Management Server and select **Edit Settings**.

4 Reconfigure the vSphere network for the Web Console Network.

   a Select the **Hardware** tab, and click **Network adapter 1**.

   b In the drop-down menu, select the appropriate network in the **Network label** text box and click **OK**.

5 Reconfigure the Web Console Network Adapter.

   a Select the **Options** tab.

   b Select **vApp Options > Properties**.

c Change the settings for FQDN, static IP address, netmask, DNS Server 1, or DNS Server 2.

d Click **OK**.

6 Right-click the Management Server virtual machine and select **Power > Power On**.

What to do next

Verify the settings from the Data Director UI. See "Verify Network Settings in Data Director," on page 113.

Reconfigure the vCenter Network Mapping

You can reconfigure the network mapping for the vCenter Network by editing the settings for the Management Server and DB Name Server virtual machines in vSphere Client.

Prerequisites

Review the information about network settings in Data Director in the *vFabric Data Director Installation Guide*.

Procedure

1 Log in to the vSphere Client as an administrator and select **Inventory > Hosts and Clusters**.

2 Select and expand the Data Director vApp, and right-click the Management Server virtual machine.

3 Select **Power > Power Off**.

4 Right-click the Management Server and select **Edit Settings**.

5 Reconfigure the vSphere network for the vCenter Network.

a Select the **Hardware** tab, and click **Network adapter**.

b In the drop-down menu in the **Network label** field, select the appropriate network and click **OK**.

6 Right-click the Management Server virtual machine and click **Power > Power On**.

7 Repeat the process for the DB Name Server virtual machine.

What to do next

Verify the changes in Data Director. See "Verify Network Settings in Data Director," on page 113.

Reconfigure the vCenter Network Adapter Settings

During installation, you specify the vCenter Network Adapter settings. You can later reconfigure the vCenter Network Adapter settings by editing the settings in the Data Director UI.

Prerequisites

Review the information about network settings in Data Director in the *vFabric Data Director Installation Guide*.

Procedure

1 Log in to Data Director as a system administrator, click the **System** tab, and select **Administration**.

2 Click the **Network Setup** tab, and click **Edit Network Setup**.

3 If necessary, enable or disable DHCP or Static IP.

a Click **Edit**.

b Toggle the **DHCP** check box to enable or disable DHCP.

c Toggle the **Static IP**check box to enable or disable static IP, and enter the netmask if **Static IP** is enabled.

d Click **OK**.

4 In the **Network Adapters** section, edit the network adapter text boxes, and click **Next**.

5 Click **Next** again.

6 Click **Finish** to confirm and save the changes.

The changes take effect immediately.

7 Verify the IP address change in the vSphere Client.

a Log in to vSphere Client as an administrator.

b Select the Management Server virtual machine.

c Click the **Summary** tab, and examine the **IP Addresses** line.

d Click **View all** to check all IP addresses.

Reconfigure the DB Name Service Network or DB Name Service Network Adapter

During installation, you configure the DB Name Service Network and the DB Name Service Network Adapter. You can later reconfigure the DB Name Service Network or the DB Name Service Network Adapter by editing the settings in the Data Director UI.

Prerequisites

Review the information about network settings in Data Director in the *vFabric Data Director Installation Guide*.

Procedure

1 Log in to Data Director as a system administrator, click the **System** tab, and select **Administration**.

2 Click the **Network Setup** tab, and click **Edit Network Setup**.

3 Click **Next** to skip vCenter Network setup.

4 To reconfigure the DB Name Service Network mapping, select the appropriate network mapping in the **DB Name Service Network** drop-down menu, and click **Next**.

5 (Optional) Click **Finish** if you are not reconfiguring the DB Name Service Network Adapter or modify the adapter.

a To enable or disable DHCP or Static IP addressing, click **Edit** next to the network name.

b Toggle the **DHCP** check box to enable or disable DHCP.

c Toggle the **Static IP** check box to enable or disable static IP, and enter the netmask in the **Netmask** text box if **Static IP** is enabled.

d Click **OK**.

e In the **DB Name Server - DB Name Service Network Adapter** section, select DHCP or Static IP.

If you select Static IP, enter the IP address and FQDN.

6 Click **Finish** to commit your changes.

Reconfigure the Internal Network or Internal Network Adapter Mapping

You can reconfigure the Internal Network or the Internal Network Adapter by editing the settings in the Data Director UI. You can change Internal Network and Internal Network Adapter settings only when all database virtual machines are powered down.

Prerequisites

Review the information about network settings in Data Director in the *vFabric Data Director Installation Guide*.

Procedure

1 Log in to vSphere Client as an administrator and power down all the database virtual machines (DBVMs).

2 Log in to Data Director as a system administrator.

3 Click the **System** tab, and select **Administration**.

4 Click the **Network Setup** tab, and click **Edit Network Setup**.

5 Click **Next** to skip vCenter Network setup.

6 To reconfigure the Internal Network mapping, select the appropriate network mapping in the **Internal Network** drop-down menu, and click **Next**.

7 To enable or disable DHCP or Static IP addressing for the network, click **Edit** next to the network name.

   a Toggle the **DHCP** check box to enable or disable DHCP.

   b Toggle the **Static IP** check box to enable or disable static IP.

   c If Static IP is enabled, enter the netmask in the **Netmask** text box.

   d Click **OK**.

8 (Optional) If you are not reconfiguring the Management Server - Internal Network Adapter, click **Finish**, or modify the adapter settings.

   a Toggle the **DHCP** check box to enable or disable DHCP.

   b Toggle the **Static IP** check box to enable or disable static IP.

   c If you select Static IP, enter the IP address and FQDN.

9 Click **Finish** to commit your changes.

10 Log in to vSphere Client as an administrator, power on the DBVMs, and verify the network settings.

   a Right-click the DB Name Server virtual machine, and select **Edit Settings**.

   b Click **Network Adapter 2**.

   The Network Adapter 2 Network label should match the change you made in the **Data Director Network Setup** tab.

Verify Network Settings in Data Director

After you reconfigure the network mapping or adapter in the vSphere Client, verify the changes in Data Director. You might have to reenter the network information in the Data Director UI if vCenter Server settings and Data Director settings do not match.

Reentering the settings for the Web Console network, Web Console network adapter, and vCenter network, is usually necessary if you use static IP addresses. Network configuration information that you might have to reenter includes the netmask, gateway, and DNS Server information.

Prerequisites

Procedure

1 Log in to Data Director as system administrator and ignore any warning that might appear.

2 Click the **Administration** tab, and click **Network Setup**.

3 Click **Edit Network Setup** and verify that the changes you made in vSphere Client are reflected in Data Director.

4 In the **Network Settings** wizard, check the network settings and reenter them if necessary.

5 Click **Next** and click **Finish**.

6 Click **Test Network Setup** to verify the network is in sync with the vCenter Server system.

Data Director Troubleshooting 14

Troubleshooting information helps you when you encounter problems with your Data Director environment. See the *Release Notes* for discussions of known issues and for corresponding workarounds.

This chapter includes the following topics:

vCenter Server Stops Responding

A Lost vCenter Session alarm appears and vCenter Server is unavailable.

Problem

An alarm in the Data Director interface states the following:

```
Lost VCenter Session
```

In the System Health panel, vCenter Connectivity is associated with a red failure icon.

When you check the vCenter Server system event log, you see the following message:

```
The transaction log for database 'VIM_VCDB' is full. To find out
why space in the log cannot be reused, see the log_reuse_wait_desc
column in sys.databases
For more information, see Help and Support Center at http://url.
```

Cause

The SQL server that manages the vCenter Server database provides capabilities for recovering the database to any point in time since the last full backup. If no backups exist, the transaction log grows indefinitely or reaches the maximum allowed limit and stops. See VMware Knowledge Base article 1003980 for background information and solution.

Solution

Configure regular backups for the vCenter Server database or use simple recovery mode (no point-in-time recovery).

Resource Bundles Become Unusable Because DRS Is Disabled

When a vSphere administrator disables DRS for the cluster that Data Director uses, all resource pools become unavailable.

Problem

If a vSphere administrators disables DRS functionality for the Data Director cluster from the vSphere Client, all resource pools become unavailable and the installation is unusable. Resource bundles, database groups, and storage and backup storage no longer work properly.

Renabling the cluster for DRS does not resolve the issue. No automated way for recovering from the situation exists.

Cause

vCenter Server supports resource pools only for clusters that have DRS enabled.

Solution

1 Reenable DRS for the cluster.

2 Delete all resource bundles and database groups that use the affected cluster.

3 Recreate the resource pools in vSphere, and recreate the resource bundles and database groups in Data Director.

4 Import each database virtual machine into Data Director.

Missing Resource Pool

When you create a resource bundle, the resource pool you want to use is not included in the list of resource pools.

Problem

When you create a resource bundle you assign a resource pool that encapsulates the CPU and memory resources. The resource pool you want to use is not available in the list of resource pools to choose from.

Cause

Data Director includes only resource pools that are compatible with Data Director resource bundles. A resource pool must meet a set of criteria to be included in the list. If the criteria are not met, the resource pool is not included in the list of available resource pools.

The resource pool must meet the following requirements:

- Resource pool is not already in use by Data Director.
- CPU limit and CPU reservation must be the same.
- Memory limit and memory reservation must be the same.
- CPU limit is expandable.
- Memory limit is expandable.
- The resource pool has no child resource pools.
- The resource pool has no child virtual machines.
- The cluster configuration for the resource pool is compatible with Data Director.

Solution

1 To display the resource pool, uncheck **Show only compatible clusters and RPs**.

If the resource pool is displayed but is dimmed, it is incompatible.

2 In Data Director, verify that the resource pool is not in use by Data Director by checking the resource bundle list.

3 In the vSphere Client, verify that the resource pool meets the requirements.

- Resource pool settings are correct (limit equal to reservation, expandable checked, and so on.
- The resource pool is empty.
- The parent resource pool uses compatible vSphere HA and vSphere DRS settings. See "Perform Advanced Cluster Configuration," on page 22.

4 Ensure that the resource pool is included in a cluster that Data Director can access.

Data Director can access all clusters in the same vCenter Server data center but cannot access clusters in a different data center.

Index

O

P

Q

R

S

T