

Using VMware vFabric Postgres

vFabric Postgres 9.2.4

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-001174-00

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

Copyright © 2012, 2013 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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

Contents

Preface	5
1 VMware Customizations for PostgreSQL	7
vFabric Postgres Enhancements	7
Deploying vFabric Postgres	8
Passwords in vFabric Postgres	8
2 Installing vFabric Postgres	11
Installation Overview	11
System Requirements	12
Deploy the OVA File	13
Install vFabric Postgres Using RPM Files	14
3 vFabric Postgres Client Tools and Libraries	17
Overview of Tools and Libraries	17
Client Tool Packages and Drivers	18
Install the Client Tools Package	19
Add an x86 vFabric Postgres ODBC Data Source on Windows	20
Relink Your Application with vFabric Postgres libpq	20
4 Managing vFabric Postgres	23
Add a License Key	23
Migrate PostgreSQL Data from Earlier Versions Into vFabric Postgres 9.2	24
Migrate PostgreSQL Data Into vFabric Postgres	24
Restarting the vFabric Postgres Service	25
Connection to a vFabric Postgres Database	25
Accounts and Services	25
Using Perl and Python Language Extensions	26
Troubleshooting Guidelines	27
5 Using the Graphical User Interface	29
Deploy the Graphical User Interface	29
Access the Graphical User Interface	29
Database Entity Management	30
SQL Management	35
Index	37

Preface

Using VMware vFabric Postgres provides information about installing and using a VMware vFabric Postgres Standard Edition DBMS.

Intended Audience

This information is intended for anyone who wants to install or use a vFabric Postgres Standard Edition DBMS. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Related Publications

The vFabric Suite documentation has information about the components of the vFabric suite.

For information about managing vFabric Postgres databases, see the public PostgreSQL documentation at <http://www.postgresql.org/docs/>. Because vFabric Postgres is compatible with PostgreSQL, you can manage vFabric Postgres databases using the information in that documentation.

To access the current versions of VMware documentation, go to <http://www.vmware.com/support/pubs>.

VMware Customizations for PostgreSQL

1

VMware vFabric Postgres is an ACID-compliant, ANSI-SQL-compliant transactional, relational database management system that is designed for the virtual environment and optimized for vSphere. It is based on the PostgreSQL open-source relational database and is compatible with PostgreSQL.

vFabric Postgres databases are managed by a DBMS that consists of a server and a client. VMware supports two versions of vFabric Postgres, vFabric Postgres Standard Edition and vFabric Postgres Standard Edition for Data Director.

vFabric Postgres Standard Edition

Supports all standard PostgreSQL connection tools and interaction methods. Optionally allows you to use a GUI tool for database management.

vFabric Postgres Standard Edition for Data Director

Seamlessly integrates with Data Director and can be managed from the Data Director GUI. You can also use traditional database tools to interact with this version, but some modifications, for example to connection strings, might be required.

This chapter includes the following topics:

- [“vFabric Postgres Enhancements,”](#) on page 7
- [“Deploying vFabric Postgres,”](#) on page 8
- [“Passwords in vFabric Postgres,”](#) on page 8

vFabric Postgres Enhancements

VMware vFabric Postgres includes memory and performance features that are not available with an open source PostgreSQL DBMS.

Elastic Database Memory

Elastic database memory enables vFabric Postgres to run with graceful performance degradation under heavy over-commitment of memory. vFabric Postgres participates in memory resource management with the guest operating system and the vSphere hypervisor to achieve elastic database memory.

vFabric Postgres monitors requests for memory from the vSphere hypervisor and monitors swap activity within the guest operating system. When the hypervisor or the guest operating system needs more memory, the vFabric Postgres buffer manager shrinks the database buffer pool to make memory available. When more memory becomes available, the buffer manager increases the amount of memory dedicated to the buffer pool.

Elastic database memory is enabled by default in the virtual appliance, but is disabled if you use the RPM files. RPM installation is supported only for the standalone version of vFabric Postgres.

Ease of Tuning on the Appliance

If you deploy the vFabric Postgres appliance or the vFabric Postgres for Data Director appliance, associated vFabric Postgres databases have higher default values than standard PostgreSQL databases for many critical settings, including `shared_buffers`, `checkpoint_segments`, and `wal_buffers`. The higher default values improve out-of-box vFabric Postgres performance with a slight increase in disk space and memory requirements. The result is that users of an embedded vFabric Postgres database can more easily tune the database for their workload.

If you are using vFabric Postgres, and you use the RPM installation, these changes to default values are not made.

Checkpoint Tuning

vFabric Postgres improves on the Postgres algorithm to make the tuning more dynamic.

In I/O-constrained environments, periods of heavy checkpointing activity often alternate with periods of light checkpoint activity.

By default, vFabric Postgres performs dynamic tuning of checkpoint parameters so that rapid changes in available I/O bandwidth or changes in the database workload reduce the likelihood of database throughput oscillations.

Deploying vFabric Postgres

You can deploy a vFabric Postgres DBMS as a virtual appliance (OVA file) or by using RPM packages.

- You can deploy the virtual appliance to create a virtual machine with the operating system (SLES 11, SP 2 64-bit Linux), a vFabric Postgres server, and a vFabric Postgres client preinstalled. The appliance version of the vFabric Postgres database includes VMware virtualization technology.
- You can use RPM to deploy vFabric Postgres. Use RPM installation with a physical host, or create a virtual machine and install one of the supported operating system that are listed on the datasheet. Use `-ivh` commands to install the RPMs. You can use this method to install the vFabric Postgres server and client software.

Passwords in vFabric Postgres

You specify the passwords for vFabric Postgres users during deployment, and can change passwords for each user individually at a later time.

Passwords differ slightly depending on whether you deploy the OVF or perform an RPM installation.

OVA Deployment

After deployment of a vFabric Postgres OVA, three users are defined.

Table 1-1. vFabric Postgres Users for OVF Deployment

User Name	User Type
root	operating system user
postgres	operating system user
postgres	database user

You can specify a single password for all three users during deployment of the OVA. If you did not specify a password during deployment, you are prompted for a password when you access the virtual appliance console for the first time.

For automated deployments, you can use `ovftool --prop:Password=secret`.

RPM Installation

During RPM installation, the installer creates the following users.

Table 1-2. vFabric Postgres Users for RPM installation

User Name	User Type
postgres	operating system user
postgres	database user

With RPM installation, no initial password is set for the postgres operating system user or the postgres database user. The root user already exists before the RPM installation and its password is being set using Linux commands.

- To set the postgres operating system user password, log in as root and use the operating system mechanism for setting passwords.
- To set the postgres database user password, log in as the postgres operating system user and use the `alter` command.

```
/opt/vmware/vpostgres/current/bin/psql -c "alter user postgres with password 'your-password'"
```

Changing Passwords After Installation

After installation, you can change passwords as follows.

- Use the `/opt/aurora/sbin/set_password` command to change the password for all three users.
- Use the `passwd` command to change passwords individually for the system users.
- Use the following command to change the password for the postgres database user.

```
/opt/vmware/vpostgres/current/bin/psql -c "alter user postgres with password 'your-password'"
```


Installing vFabric Postgres

Before you install vFabric Postgres, review the requirements and the deployment or installation process.

This chapter includes the following topics:

- [“Installation Overview,”](#) on page 11
- [“System Requirements,”](#) on page 12
- [“Deploy the OVA File,”](#) on page 13
- [“Install vFabric Postgres Using RPM Files,”](#) on page 14

Installation Overview

The vFabric Postgres server and client software are distributed together. You can either deploy an Open Source Virtual Appliance (OVA) file or install a series of RPM packages.

Virtual Appliance Deployment Overview

The process of deploying a vFabric Postgres virtual appliance is similar on all the different supported virtualization platforms.

- 1 Install one of the VMware virtualization platforms such as vSphere 5.x, VMware Workstation 9.x, VMware Fusion 5.x, or VMware Player 5.x.

NOTE For a production system, only vSphere 5.x is supported.

- 2 Deploy the virtual appliance.
- 3 To manage the new DBMS, log in to the virtual appliance console and use the preinstalled psql tool or point your Web browser to https://your_vApp_IP:8443.
- 4 To manage the virtual appliance, log in to the virtual appliance console or point your Web browser to https://your_vApp_IP:5480.

RPM Deployment Overview

The process of installing the vFabric Postgres DBMS from RPM packages consists of the following high-level tasks.

NOTE This process is supported only for the standalone version. It is not supported for vFabric Postgres for Data Director.

- 1 Make sure the host or virtual machine that you want to use is running a supported operating systems and meets all the other requirements.

- 2 Download and install the client, server, and init RPM files.
 - a client package
 - b server package
 - c init package
- 3 Log in to the new DBMS using the client software.

System Requirements

You can deploy the virtual appliance and install the RPM packages on several operating systems.

Supported Platforms for OVA Deployment

For the virtual appliance (OVA), several virtualization platforms are supported during development, but support is more limited during production.

Development	While you develop your application and run tests, you can deploy the virtual appliance on the latest edition of any VMware virtualization platform, including VMware vSphere 5.x, VMware Workstation 9.x, VMware Fusion 5.x, or VMware Player 5.x.
Production	In a production environment, you must install vFabric Postgres on VMware vSphere 5.x.

Resource Requirements for RPM Installation

If you install the RPM files, you must have a physical host or a virtual machine that meets the following minimum requirements.

RAM	512 MB
CPUs or vCPUs	1 or more
Disk Space	12 GB

Resource Requirements for OVA Deployment

The virtual appliance requires 1GB of RAM. See the vSphere documentation for information on the required memory overhead.

Operating Systems

The vFabric Postgres server software is currently supported on the following operating systems.

Red Hat Linux	RHEL 6.2 (64 bit)
SUSE Linux	SLES 11 SP 1 (64 bit) or SLES 11 SP2 (64 bit)
Oracle Enterprise Linux	OEL 6

Database Clients

The vFabric Postgres product includes custom JDBC, ODBC, and LIBPQ drivers. Drivers and tools are available for the standalone vFabric Postgres Standard Edition and the vFabric Postgres Standard Edition for Data Director.

Database clients for Windows, Linux, and MAC OS X, both 32 bit and 64 bit, are included.

Many community PostgreSQL clients, such as Npgsql, and pycopg2 are also supported in both 32-bit and 64-bit configurations.

Deploy the OVA File

You can deploy the OVA file on vSphere 5.x for use during development or for production environments. In addition, can deploy the OVA file on VMware Workstation 9.x, VMware Fusion 5.x, or VMware Player 5.x for use during development.

This topic describes the process when you use the vSphere Web Client with vSphere. If you are using the vSphere Client, or one of the other VMware products, the process is similar but the prompts might differ slightly.

See [“Passwords in vFabric Postgres,”](#) on page 8 for an introduction to user accounts and passwords.

Prerequisites

Download the OVA file from the VMware download site.

NOTE Different OVA files exist for vFabric Postgres Standard Edition and vFabric Postgres Standard Edition for Data Director.

Procedure

- 1 Connect to a vCenter Server with the vSphere Web Client.
- 2 Right-click an inventory object that is a valid parent object of a virtual machine, such as a datacenter, folder, cluster, resource pool, or host and select **Deploy OVF Template**.
- 3 If prompted, download the client plug-in.
You have to close all browsers to download the plug-in.
- 4 Respond to the wizard prompts.

Screen	Action
Select Source	Specify the location of the OVA file.
Review Details	Review the OVA information.
Accept EULAS	Review and accept the license agreement.
Select name and folder	Specify the name and location for the virtual appliance.
Select storage	Select the storage for the virtual appliance. You can use the pull-down menu to change the disk format.
Setup networks	Map the networks used in the OVF template to networks in your inventory and select
Customize template	Specify the password that you want to use initially for the three users that the OVA file defines. A minimum of six characters is required.
Ready to Complete	Review the settings and click Finish to start deployment.

When deployment completes, the virtual appliance powers on.

- 5 (Optional) If you did not specify a password during deployment, specify it now. right-click the virtual appliance and select **Open Console** and enter the initial password for all user accounts.
 - a Click **Edit Settings**.
 - b Click the Options tab and select Properties under vApp Options.
 - c Enter the initial password in the text box on the right.

What to do next

You can now manage your vFabric Postgres environment.

- To manage the new DBMS, log in to the virtual appliance console and use the preinstalled `psql` tool or point your Web browser to `https://your_vApp_IP:8443`.
- To manage the virtual appliance, log in to the virtual appliance console or point your Web browser to `https://your_vApp_IP:5480`.

Install vFabric Postgres Using RPM Files

If you want to install vFabric Postgres on a new or existing virtual machine or on a physical host, you can use RPM installation process.

Different editions of the vFabric Postgres server software are supported on different operating systems. See the datasheet for information.

Prerequisites

- Create a new virtual machine running a supported operating system, or log in to a virtual machine where one of these operating systems is currently running. You can also install the RPMs on a physical host that runs one of the supported operating systems.
- Verify that you have access to the Internet to download the RPM packages.
- If you install 32-bit binaries on a 64-bit system, install compatibility libraries as well. On RHEL6, use `yum install glibc.i686 nss-softokn-freebl.i686`.

Procedure

- 1 Download at a minimum the following ZIP files from the VMware download site.
 - ZIP file for vFabric Postgres
 - ZIP file for the vFabric Postgres client tools and libraries

Optional components, 32-bit client RPMs, and client tools for Windows, Macintosh, ODBC, and JDBC are also available on the download site.

- 2 Unzip the ZIP file and install each of the RPM files using the `rpm -ivh` command, in the order shown below, or install all files at once with a single command.

```
>rpm -ivh V
VMware-vFabric-Postgres-client-version-XXXXXX.x86_64.rpm
VMware-vFabric-Postgres-version-XXXXXX.x86_64.rpm
VMware-vFabric-Postgres-server-init-version-XXXXXX.x86_64.rpm
```

After the files have been installed, a database instance is in the `/var/vmware/vpostgres/current/pgdata` directory. The user name for the database is `postgres`. A user with the name `postgres` has become available in your operating system.

- 3 Log in to the database with the 60 day license key and set the password manually.

You can set the password for the database user with username `postgres`, and for the operating system user with user name `postgres`. These two users are not the same.

- (Optional) To connect to the database from the local host as the database user with username postgres, set the password, as in the following example.

```
[root@rhel-6-64-esx41 ~]# /opt/vmware/vpostgres/current/bin/psql -U postgres
psql.bin (version)
Type "help" for help.
postgres=# alter user postgres password 'mypassword';
ALTER ROLE
postgres=#
```

- To set the password for the operating system user with user name postgres, log in as root and set the password as follows.

```
$passwd postgres
```

- If you want to you use the GUI, you can download a ZIP file that contains the vpgdbem.war file from the vFabric Postgres download site and move the file to the webapps directory of your Tomcat server.

You can then use the vpgdbem URI to access the GUI. For example, if your Tomcat server is installed for port 8080, you can access the GUI at <http://ipaddress:8080/vpgdbem>.

vFabric Postgres Client Tools and Libraries

3

You can download and install vFabric Postgres client tools for the operating system that you are using to manage and access vFabric Postgres databases. The command line front end to PostgreSQL, `psql`, is also included.

NOTE Different client tools and libraries are available for vFabric Postgres and for vFabric Postgres for Data Director. Go to the correct download location to download the tools and libraries you need.

This chapter includes the following topics:

- [“Overview of Tools and Libraries,”](#) on page 17
- [“Client Tool Packages and Drivers,”](#) on page 18
- [“Install the Client Tools Package,”](#) on page 19
- [“Add an x86 vFabric Postgres ODBC Data Source on Windows,”](#) on page 20
- [“Relink Your Application with vFabric Postgres libpq,”](#) on page 20

Overview of Tools and Libraries

The vFabric Postgres client tools are based on the Postgres client database tools and are customized for vFabric Postgres. The tools support common configuration commands. The libraries include several APIs and the ODBC driver for PostgreSQL.

Separate download packages are available for vFabric Postgres and for vFabric Postgres for Data Director.

Versions for Linux x86, 32 bit and 64 bit, for Windows x86, 32 bit and 64 bit, and for Mac-OS are available.

Linux The Linux RPM includes ODBC drivers for vFabric Postgres. The Linux ODBC driver requires unixODBC-2.3.1 or greater.

Windows The vFabric Postgres client tool installer package for Windows includes ODBC and JDBC drivers for vFabric Postgres

The vFabric Postgres client database tools that are included in the vFabric Postgres client tools packages are the same tools that are available as part of standard PostgreSQL and include `pg_dump`, `pg_restore`, and `psql`. See the postgresql.org Web site for details.

The vFabric Postgres client tools ship with the following libraries.

Table 3-1. vFabric Postgres Client Tool Libraries

Library	Description
libpq.so (Linux) or libpq.dll (Windows)	The C API to PostgreSQL. Libpq is the underlying engine for several PostgreSQL APIs such as those written for C++, Perl, Python, Tcl, and ECPG.
psqlodbcw.so (Linux) or psqlodbc35w.dll (Windows)	The ODBC driver for PostgreSQL.

The vFabric Postgres client tool libraries are customized for use with vFabric Postgres databases, but you can use the standard Postgres libraries. To ensure that you link with the vFabric Postgres libraries, do one of the following.

- If you want to keep the standard Postgres libraries on your system, ensure that your `LD_LIBRARY_PATH` environment variable specifies the location of the standard PostgreSQL libraries first.
- If you do not want to keep the standard PostgreSQL libraries, remove them and ensure that your `LD_LIBRARY_PATH` environment variable points to the location of the vFabric Postgres libraries on your system.

Client Tool Packages and Drivers

You can download client tool packages for Windows and Linux from the VMware download site. After you download the tools, you can use the drivers included in the packages.

Packages

If you plan to on writing code and on compiling an application to link with libpq, download both the client package and the development package.

You can download the client tool package for your platform from the VMware download site. Be sure to download the appropriate package for your environment.

vFabric Postgres <http://vmware.com/go/download-vfabric-postgres>

vFabric Postgres for Data Director <http://www.vmware.com/go/download-datadirector>

You can download tools and drivers for Windows, Linux, Java, or Macintosh.

Drivers

The vFabric Postgres client tools package includes a JDBC driver and an ODBC driver customized for vFabric Postgres. To connect to Data Director, use the vFabric Postgres JDBC or ODBC drivers, not the standard PostgreSQL drivers.

JDBC Driver After installation, you can find the JDBC driver in the following locations.

Microsoft Windows C:\Program Files\VMware\vPostgres\version\JDBC

Linux /opt/vmware/vpostgres/current/JDBC

The `Samples` directory contains a simple Java example and README file that show how to connect to Data Director using JDBC.

For example, if your application uses the JDBC driver to access a database, and you install the application as `/usr/local/lib/myapp.jar` and the PostgreSQL JDBC driver as `/usr/local/pgsql/share/java/postgresql.jar`, you run the application as follows.

```
export
CLASSPATH=/usr/local/lib/myapp.jar:/usr/local/pgsql/share/java/postgresql.jar:. java MyApp
```

ODBC Driver

The vFabric Postgres installation process installs the vFabric Postgres ODBC driver. You can verify the Windows ODBC driver installation as follows.

- 1 Select **Start > Administrative Tools > Data Sources (ODBC)**.
- 2 Click the **Drivers** tab.
- 3 Verify that the VMware vFabric Postgres ODBC driver appears in the list of installed ODBC drivers.

Install the Client Tools Package

You can install the vFabric Postgres client tools on Windows or Linux systems. The package includes drivers customized for vFabric Postgres. You can install only the base package, or install the development RPMs as well.

Prerequisites

- Download the package.
- If you are using vFabric Postgres for Data Director, verify that the Data Director ESXi host or virtual machine is running.

Procedure

- 1 Install the package.

Operating System	Installation Process
Linux	<p>Install the RPM files by running the following command.</p> <pre>rpm -ivh pathToClientRpms rpm -ivh pathToClientRpms</pre> <p><i>pathToClientRpms</i> is the full pathname of the RPM package location on your system. The default installed location is <code>/opt/vmware/vpostgres/<i>version</i></code>.</p> <p>Use <code>-Uvh</code> instead of <code>-ivh</code> if you perform an upgrade.</p>
Windows	<ol style="list-style-type: none"> a Double-click the installer to start the installer. b Accept the license agreement and confirm the install location. Installation proceeds. The default installed location is <code>\Program Files\VMware\vPostgres\<i>version</i></code>. If you install the x86 vFabric Postgres client tools on a Windows 64-bit system, the Windows installer places the client tools in <code>\Program Files (x86)\VMware\vPostgres\<i>version</i></code>.
Macintosh	<p>Run or rerun the installer. You can double-click the PKG file to start the installer GUI or install from the command line by running the following command.</p> <pre># sudo installer -pkg /path/to/VMware-vPostgres-client-....pkg -target /</pre>

- 2 Ensure that your `PATH` environment variable includes the location of the vFabric Postgres client tools, for example `C:\Program Files\VMware\vPostgres\version\bin`.

What to do next

If you install both the x86 and the 64-bit vFabric Postgres client tools on a 64-bit Windows system, see [“Add an x86 vFabric Postgres ODBC Data Source on Windows,”](#) on page 20.

If you are developing a custom application, relink with libpq. See [“Relink Your Application with vFabric Postgres libpq,”](#) on page 20.

Add an x86 vFabric Postgres ODBC Data Source on Windows

If you install both the x86 and the 64-bit vFabric Postgres client tools on the same 64-bit Windows system, you must explicitly add an x86 ODBC data source.

Prerequisites

Install the x86 and the 64-bit vFabric Postgres client tools.

Procedure

- 1 In Windows Explorer, go to C:\Windows\SysWOW64\.
- 2 Double-click `odbcad32.exe`.
- 3 Select the **System DNS** tab and click **Add**.
- 4 Select the VMware vFabric Postgres Unicode 32-bit data source.
- 5 Click **Finish**.

Relink Your Application with vFabric Postgres libpq

If you want to use an existing PostgreSQL application with vFabric Postgres, you can relink the application. Because vFabric Postgres `libpq.so` is dynamically linked with `libssl`, the static `ld` linker does not recognize the `rpath` of `$ORIGIN`. You can relink to specify the `rpath`.

Prerequisites

Install the vFabric Postgres client tools. You can relink without installing the development RPMs.

Procedure

- ◆ Relink with vFabric Postgres based on your operating system.

Operating System	Relinking Process						
Linux	<p>a Read <code>/opt/vmware/vpostgres/current/share/libpq-doc/README.vpostgres-libpq</code>.</p> <p>b Override the dynamic library search path by adding <code>/opt/vmware/vpostgres/current/lib-public</code> to <code>LD_LIBRARY_PATH</code>.</p> <pre># export LD_LIBRARY_PATH=/opt/vmware/vpostgres/current/lib-public # mypgapp</pre> <p>or</p> <p>c Relink using the vFabric Postgres libpq.</p> <pre># gcc -o t.t.c -L/opt/vmware/vpostgres/current/lib -Wl,'-rpath=/opt/vmware/vpostgres/current/lib' -lpq</pre>						
Windows	<p>Copy libpq and other libraries to the directory of the application binaries and relink.</p> <p>By default, the libraries and header files are in the following locations.</p> <table border="0"> <tr> <td>Development libraries</td> <td>C:\Program Files\VMware\vPostgres\version\dev</td> </tr> <tr> <td>libpgport.lib and libpq.lib libraries</td> <td>C:\Program Files\VMware\vPostgres\version\dev\lib</td> </tr> <tr> <td>libpq header files</td> <td>C:\Program Files\VMware\vPostgres\version\dev\include</td> </tr> </table>	Development libraries	C:\Program Files\VMware\vPostgres\version\dev	libpgport.lib and libpq.lib libraries	C:\Program Files\VMware\vPostgres\version\dev\lib	libpq header files	C:\Program Files\VMware\vPostgres\version\dev\include
Development libraries	C:\Program Files\VMware\vPostgres\version\dev						
libpgport.lib and libpq.lib libraries	C:\Program Files\VMware\vPostgres\version\dev\lib						
libpq header files	C:\Program Files\VMware\vPostgres\version\dev\include						
Mac OS X	<p>Perform one of the following tasks.</p> <ul style="list-style-type: none"> ■ Override the dynamic library search path by adding the <code>/opt/vmware/vpostgres/version/lib</code> to the <code>DYLD_LIBRARY_PATH</code> environment variable, as follows: <pre># export DYLD_LIBRARY_PATH=/opt/vmware/vpostgres/current/lib # mypgapp</pre> ■ Relink using the vFabric Postgres libpq library during compilation. Relinking requires the Xcode developer toolset. For example, to embed the full path of <code>libpq.dylib</code> in the executable binary <code>mypgapp</code>, run this command. <pre># gcc -o mypgapp mypgapp.c -L/opt/vmware/vpostgres/current/lib -lpq</pre> ■ Relink using the vFabric Postgres libpq after compilation. Relinking requires the Xcode developer toolset. <p>NOTE This changes the binary to use vPostgres libpq.</p> <pre># install_name_tool -change "/usr/lib/libpq.5.dylib" "/opt/vmware/vpostgres/current/lib/libpq.5.dylib" mypgapp</pre> <p>To confirm which library is linked, run this command.</p> <pre># otool -L mypgapp</pre> 						

Managing vFabric Postgres

After you have installed the vFabric Postgres DBMS and the client tools, you can perform a variety of management tasks.

NOTE If you are using the vFabric Data Director version of vFabric Postgres, you perform most of your management tasks from the Data Director GUI. This includes creating, monitoring, and managing vFabric Postgres databases, and managing license keys for production system. See *vFabric Data Director Administrator and User Guide*.

This chapter includes the following topics:

- [“Add a License Key,”](#) on page 23
- [“Migrate PostgreSQL Data from Earlier Versions Into vFabric Postgres 9.2,”](#) on page 24
- [“Migrate PostgreSQL Data Into vFabric Postgres,”](#) on page 24
- [“Restarting the vFabric Postgres Service,”](#) on page 25
- [“Connection to a vFabric Postgres Database,”](#) on page 25
- [“Accounts and Services,”](#) on page 25
- [“Using Perl and Python Language Extensions,”](#) on page 26
- [“Troubleshooting Guidelines,”](#) on page 27

Add a License Key

vFabric Postgres supports using a permanent license key. You create a file in the virtual machine, or you use the vFabric Suite license server.

Procedure

- 1 Log in as the root user.
- 2 Get the vFabric Postgres license key.

The license key consists of a sequence of numbers and letters plus the quantity and expiration, for example, XXXXX-XXXXX-XXXXX-XXXXX-XXXXX [quantity=1, expiration=Permanent].

- 3 Use your valid license key to create the `/etc/opt/vmware/vfabric/vf.vpg-serial-numbers.txt` file, as in the following example.

```
# echo "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX [quantity=1, expiration=Permanent ]"  
>> /etc/opt/vmware/vfabric/vf.vpg-serial-numbers.txt
```

- 4 Change the group and file permissions.

```
# chgrp vfabric /etc/opt/vmware/vfabric/vf.vpg-serial-numbers.txt
# chmod 644 /etc/opt/vmware/vfabric/vf.vpg-serial-numbers.txt
```

- 5 Restart the vFabric Postgres service and check the license key log at `/opt/vmware/vFabric/vf.vpg.log` to verify that a `LicenseActivatedEvent` happened.

If vFabric Postgres does not find this file, it will attempt to use the vFabric license server. If the license server does not respond, the vFabric Postgres server uses the remainder of the evaluation license if one exists.

See the vFabric Suite 5.1 documentation for more information about vFabric Suite licenses.

Migrate PostgreSQL Data from Earlier Versions Into vFabric Postgres 9.2

If you want to import PostgreSQL data or vFabric Postgres data from an earlier version, such as PostgreSQL 9.1 or vFabric Postgres 9.1, into vFabric Postgres 9.2, you can use a backup and restore process.

Procedure

- 1 Log in to the vFabric Postgres host.

- 2 Back up the existing database.

```
/opt/vmware/vpostgres/current/bin/pg_dumpall -c -h
ip_address_of_existing_db -U postgres -l postgres
```

- 3 Restore the database on vFabric Postgres 9.2.

```
/opt/vmware/vpostgres/current/bin/psql -f
/var/vmware/vpostgres/current/mydb-backup
```

- 4 Vacuum the database after restore to make sure there is no error.

```
/opt/vmware/vpostgres/current/bin/vacuumdb -a -z -U postgres
```

Migrate PostgreSQL Data Into vFabric Postgres

You can migrate an existing PostgreSQL PGDATA directory that is using PostgreSQL version 9.2 to vFabric Postgres by stopping the service, archiving the data, and copying the archived file.

Migrating PostgreSQL data from an earlier version of PostgreSQL requires a different process. See [“Migrate PostgreSQL Data from Earlier Versions Into vFabric Postgres 9.2,”](#) on page 24.

As part of the migration process, you stop and start the vFabric Postgres service. See [“Stop and Start the vFabric Postgres Service on the Virtual Appliance,”](#) on page 25 and [“Stop and Start the vFabric Postgres Service for RPM Installations,”](#) on page 25

Procedure

- 1 Stop the vFabric Postgres service

- 2 Archive the existing PGDATA directory using any archiving tool.

- 3 Copy the PGDATA archive to the vFabric Postgres.

- 4 Unarchive and replace the existing PGDATA directory, `/var/vmware/vpostgres/current/pgdata`.

- 5 Start the vFabric Postgres service.

Restarting the vFabric Postgres Service

You can stop and start the vFabric Postgres service on the virtual appliance or for an RPM installation. Restarting the service might be required if you change the configuration.

Stop and Start the vFabric Postgres Service on the Virtual Appliance

Stop and start the service on the virtual appliance if you deployed the OVF template and you want to change the database configuration.

If you deployed the vFabric Postgres OVF template, use the following commands to stop and then restart the service. For the appliance service, these commands also stop and start the VMware HA (high availability) monitor process that makes sure the database process is up and running.

Procedure

- 1 Change the configuration.
- 2 Stop the service.


```
$service aurora_mon stop
```
- 3 Restart the service.


```
$service aurora_mon start
```

Stop and Start the vFabric Postgres Service for RPM Installations

If you change the database configuration of your RPM installation, you can perform a configuration file reload or you can stop and start the vFabric Postgres service.

If you installed the vFabric Postgres service using RPM files, the service is running within your virtual machine. In contrast to the virtual appliance, VMware HA is not integrated with the service and is not affected.

For many configuration options, a configuration file reload with `SIGHUP` or `pg_ctl reload` is sufficient.

Procedure

- 1 Change the configuration.
- 2 Stop the service.


```
$service vpostgres stop
```
- 3 Restart the service.


```
$service vpostgres start
```

Connection to a vFabric Postgres Database

Connecting to a vFabric Postgres database that is not integrated with Data Director is the same as connecting to a standard PostgreSQL database.

Accounts and Services

When you deploy the vFabric Postgres DBMS, two users are created. When the vFabric Postgres server is running, it includes services that accept remote connections.

Accounts Created During OVA Deployment

When you deploy the vFabric Postgres OVA template, the following users are created.

Table 4-1. vFabric Postgres Users Created During OVA Deployment

User Name	User Type	
root	operating system user	The root user can log into the appliance from the guest console using the same random password as the postgres user. Remote ssh logins are disabled for root. Database access is also disabled for root.
postgres	operating system user	The postgres operating system user can start the database instance and act as the super user. This user can log in to the console and into a shell.
postgres	database user	The postgres user is a database administrator account. This user can log into the appliance from the guest console, log in remotely using ssh, or connect to the database service on port 5432. The OVF deployment wizard prompts you to specify the password for all three users. Use the <code>/opt/aurora/sbin/set_password</code> command to change the password for the postgres user.

Services that Accept Remote Connections

The following vFabric Postgres services accept remote connections to the virtual appliance by default.

Table 4-2. vFabric Postgres Services that Accept Remote Connections

Service	Port
Postgres service	5432
SSH service	22
VAMI (Virtual Appliance Management Infrastructure) Web Management UI. You can connect to port 5480 via https to update or reconfigure the appliance.	5480
VAMI SFCB broker	5488 and 5489
GUI HTTPS service	5432
GUI HTTP service, which redirects to HTTPS	8080

If you are performing the RPM installation, only port 5432 is relevant.

Using Perl and Python Language Extensions

You can use vFabric Postgres with the PL/Perl and PL/Python language extensions. You must make sure you are using the correct versions of the language and the operating system.

PL/Python and vFabric Postgres

The PL/Python vFabric Postgres extension is supported with the following Python versions and Linux distributions.

Python version

You must install the Python 2.6.x RPM on your system.

You cannot use the extension with earlier versions of Python (2.5.x) or with later versions of Python (2.7.x, 3.x).

Linux version

The RHEL 6, SLES 11 SP1, and SLES 11 SP2 distributions provide the Python 2.6 RPM and are supported for the PL/Python vFabric Postgres extension.

RHEL 5.x does not provide Python 2.6 RPM. The PL/Python vFabric Postgres extension is not supported on RHEL 5.x.

PL/Perl and vFabric Postgres

The PL/Perl vFabric Postgres extension is supported with the following Python versions and Linux distributions.

Perl version	You must install the Perl 5.10.x RPM on your system. You cannot use the extension with earlier versions of Perl (5.8.x) or with later versions of Perl (5.12.x, 5.14.x).
Linux version	The RHEL 6, SLES 11 SP1, and SLES 11 SP2 distributions provide Perl 5.10.x RPMs and are supported for the PL/Perl vFabric Postgres extension. RHEL 5.x does not provide the RPMs. The PL/Perl vFabric Postgres extension is not supported on RHEL 5.x.

On supported Linux distributions, the `VMware-vPostgres-server-extensions` RPM, which contains the PL/Perl extension, includes an `install-time` scriptlet that attempts to locate the `libperl.so` shared library on the system by looking in the following locations. The scriptlet looks for the Perl binary in the following locations.

- 1 `libperl.so` path as defined by the Perl binary, where the scriptlet looks for the Perl binary in the following locations.
 - a `PATH` variable
 - b RPM location database
 - c `/usr/bin`
- 2 `libperl.so` under `/usr/lib64`
- 3 `libperl.so` under `/usr/lib`

The scriptlet creates a soft link to `libperl.so` under `/opt/vmware/vpostgres/version/lib`. If the script cannot find `libperl.so` on the system, a warning is printed during RPM installation and the PL/Perl vFabric Postgres extension might not work properly.

Troubleshooting Guidelines

Use the options listed in this section to analyze connection or performance problems.

Client Cannot Connect

If your client cannot connect to the vFabric Postgres appliance or to a vFabric Postgres server installed using RPMs, follow these steps to troubleshoot the issue.

- 1 Ping the server IP from your client.
- 2 Verify that Postgres is running by running the following command on the command line.


```
ps ax | grep postgres
```
- 3 Try to connect a local PostgreSQL client to the vFabric Postgres server.
- 4 Review the logs in `/var/vmware/vpostgres/current/pgdata/pg_log`.

Database Transactions Per Second Less Than Expected

If the database transactions per seconds are less than expected, follow these steps to troubleshoot the issue.

- 1 Make sure your PGDATA VMDK is on a high-performance datastore.
- 2 Look for missing indexes in your SQL queries.

- 3 Analyze concurrent queries for conflicts.
- 4 Increase the number of vCPUs and/or memory.
- 5 As a last resort, turn off `synchronous_commit` in `var/vmware/vpostgres/current/pgdata/postgresql.conf` and restart the appliance. Monitor for performance changes. See the PostgreSQL documentation for details.

Using the Graphical User Interface

A graphical user interface for managing database entities and running SQL commands is available as part of vFabric Postgres. The interface is included with the virtual appliance, and can be installed as part of RPM installation.

The tool supports database entity management and SQL management.

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

This chapter includes the following topics:

- [“Deploy the Graphical User Interface,”](#) on page 29
- [“Access the Graphical User Interface,”](#) on page 29
- [“Database Entity Management,”](#) on page 30
- [“SQL Management,”](#) on page 35

Deploy the Graphical User Interface

If you are using the OVA file to deploy vFabric Postgres, the GUI is included. If you are using the RPM installation, you place a WAR file on the Tomcat server to make the GUI available.

Procedure

- 1 Download the ZIP archive that has a name like `VMware-vFabric-Postgres-dbem-9.2.2.0-XXXXXX.zip` from the VMware download site and extract the `vpgdbem.war` file from the ZIP archive.
- 2 Copy the `vpgdbem.war` file to the `webapps` directory of your Tomcat server.

You can access the GUI by using the Tomcat server address and the `vpgdbem` URI. For example, if your Tomcat server is installed for port 8080, you can access the GUI at the following location:

```
http://ipaddress:8080/vpgdbem
```

Access the Graphical User Interface

The graphical user interface allows you to view database entities, create schemas and tables, and perform other management tasks. You can use the GUI console to interact with the database by using SQL.

Prerequisites

- If you are using the virtual appliance (OVA), the GUI is preinstalled.
- If you used the RPM installation process, you deploy a WAR file on your Tomcat server and access the GUI from there. See [“Deploy the Graphical User Interface,”](#) on page 29.

Procedure

- 1 Access the GUI. from the vSphere Client or by using a URL.

Option	Process
vSphere Client	<ol style="list-style-type: none"> a Use a vSphere Client to connect to the vCenter Server that manages the host or cluster on which the virtual machine runs and click the Summary tab. b If you click the Available link next to Status, you are directed to the URL of the GUI.
URL	The blue virtual appliance console that appears after all setup scripts run displays the IP address to access the GUI. The address is <i>virtual_machine_IP</i> :8443. Access the GUI from your Web browser using this address.

You cannot access the GUI from the vSphere Web Client.

- 2 When prompted, give the login credentials for the postgres database user.

Database Entity Management

Administrators perform database entity management tasks to ensure the effective and efficient operation of databases.

You can manage database entities from the **Database** tab. Managing database entities includes vacuuming and analyzing databases, and creating, altering, dropping, and browsing database entities such as the following.

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

The left pane displays schema objects. The middle pane allows you to manage individual objects.

Create a Database

The vFabric Postgres GUI allows you to create a database and customize its attributes as part of creation.

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- 1 Right-click the host that is displayed in the left pane and select **Create Database**.

- Specify database information in the Create Database dialog.

Option	Description
Name	Name of the database (required)
Owner	Owner of the database. Always postgres for vFabric Postgres databases.
Encoding	Not currently supported.
Tablespace	Not currently supported.
Connection Limit	Number of connections. Defaults to -1, which is unlimited. Change this value only when you want to enforce that only a limited number of connections can be established with this database. Set this value lower than the maximum number of database server users. For example, if you have five databases and you want to make sure one of those databases does not have more than five users while the other databases can have as many users as the database server can handle, you use this property to limit the number of users for that database.
Comment	Comment about the purpose and characteristics of the database.

- Click **OK** to create the database.

View Database Entities

You can view and examine database entities such as schema, views, and so on from the GUI.

Procedure

- Double-click a database.
- If prompted, log in to the database
- Expand entity icons such as Catalogs and Schemas in the left pane and select individual items to view details.

What to do next

Manage database entities.

Vacuum Analyze a Database

You can use Vacuum Analyze to discover and reclaim storage occupied by dead tuples. Tuples you delete or that are made obsolete when you update the database remain in their table until you perform a Vacuum action.

You can perform Vacuum and Analyze on a database or a table.

- Use Vacuum periodically, especially on frequently updated tables, to keep the database performing well.
- Use Analyze to collect statistics about the contents of tables and to store the results.

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- Right-click your database in the left pane and select **Vacuum Analyze Database**.
- (Optional) Customize the operation.
 - To perform the Vacuum operation, uncheck the **Analyze** checkbox and click **OK**.
 - To perform the Analyze operation, uncheck the **Vacuum** checkbox and click **OK**.
 - To include the **Full** or **Freeze** operations with the Vacuum operation, check those checkboxes.

- 3 Click **OK** to perform the selected operations.

Create a Schema

After you create a database, you set up its entities, starting with the database schema.

Prerequisites

- Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.
- Create a database. See [“Create a Database,”](#) on page 30.

Procedure

- 1 Click the database in the left pane to display the Catalogs and Schemas icons.
- 2 Right-click **Schemas** and select **Create Schema**.
- 3 Enter the schema information.
- 4 Click **OK**.

What to do next

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

Create a Table for Schema Data

After you create a schema, you create tables to contain the schema's data.

Prerequisites

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

You created a database and a schema, and are in the **Console**.

Procedure

- 1 In the left pane, click the **Schemas** arrow to expand it.
- 2 Right-click the schema and select **Create > Table**.
- 3 Type the table name, fill factor, and comment.
- 4 Click **Next**.
- 5 Click **Add** to add a column.
 - a Type the column name, and select the column type..

Depending on the column type, you can specify a length or precision, a default value for the column, and add a comment.
 - b If users must enter a value for the column, select the **Not Null** check box.
 - c If the column is a primary key, select the **Primary Key** check box.
- 6 (Optional) In **Constraints**, select the type of constraint, **Foreign key**, **Unique**, or **Check**, that applies to the new column.

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

- a Click **Create**.
- b Enter the conditions for the constraint, and click **OK**.
- c Click **Next** to continue, or click **Finish** to create the table.

- 7 (Optional) In the **Auto Vacuum Settings**, select settings for removing stale data from your table, The default settings work well for most environments. For information about autovacuum, see the documentation on the Postgres.org site.
- 8 Click **Finish** to create the table.

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.

Prerequisites

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

Procedure

- 1 In the left pane, click the **Schemas** arrow to expand it.
- 2 Right-click the schema and select **Create > View**.
- 3 Enter the view properties.
 - a Type a unique name in the **Name** text box.
If the name is case-sensitive, select the **Case sensitive** check box.
 - b (Optional) To restrict who can modify the view, select an owner for the view definition from the drop-down menu.
 - 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"
```
- 4 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 33.

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.

Procedure

- 1 In the left pane, click the **Schemas** arrow to expand it.
- 2 Click the arrow next to the schema to expand it.
- 3 Select **Views** in the left pane.
All views under the schema appear in the list in the middle pane.

- 4 Right-click a view and select **Open**.
The view appears in the left pane.
- 5 Click the **View Data** tab to examine the data associated with the view.

Create a Constraint

Constraints let you reduce data entry errors by verifying data before inserting the data into a table.

You can create constraints when you create a table, or you can add them later by entering SQL fragments.

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- 1 Click a table to select it, and click the gear icon.
- 2 Select **Create > Constraint**.
- 3 Select a constraint to create.

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

- 4 Complete the dialog and click **OK**.

Example: Create a Check Constraint

A check constraint evaluates to a Boolean value. Use Check constraints to determine whether 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.

- 1 Click a table to select it.
- 2 Click the gear icon, and select **Create > Constraint**.
- 3 Select **Check Constraint**.
 - 1 Type a name for the constraint, such as *check_positive_price*.
 - 2 Enter the constraint in the **Check** text box.
 - 3 (Optional) Enter a comment that describes the constraint.
 - 4 Click **OK**.

Change the postgres Database User Password

You can change the password for the database user postgres from the console or from the vFabric Postgres GUI.

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- 1 In the vFabric Postgres GUI, click a database arrow to display catalogs, schemas, and db login users.
- 2 Click DB Login Users.
- 3 In the right pane, right-click the postgres user and select **Properties**.
- 4 Specify a new password, confirm the password, and click **OK**.

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

You can enter and run SQL queries from the vFabric Postgres GUI.

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- 1 Click **Enter SQL**.
- 2 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.
- 3 Click **Execute** to run the query.

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

View a Query Plan

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

Prerequisites

Connect to the vFabric Postgres GUI. See [“Access the Graphical User Interface,”](#) on page 29.

Procedure

- 1 Expand **Schemas** in the left pane.
- 2 Select a 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, runtime, and CPU cost.

What to do next

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

Index

A

accounts 25
administer SQL 35
analyze SQL query plan 35
aurora_mon start 25
aurora_mon stop 25
auto-vacuum 32

C

checkpoint 7
client tools 17
configure auto-vacuum 32
constraint creation 34
create check constraints 32
create column constraints 32
create constraints 34
create database 30
create database schema 32
create foreign key 32
create SQL queries 35
create unique constraints 32
create views 33

D

database entities, managing 30
database schema 32
deploy database 8

E

elastic database memory 7
enhancements 7

F

foreign key 32

G

graphical user interface 29

I

installation 11
installation overview 11
installing client tools 19

J

JDBC 25

L

libpq 20
libraries 17
linking 20
Linux packages 18
local license, installing 23

M

manage database entities 30
managing databases 23
migrate 24
modify SQL queries 35

O

ODBC data source 20
OVA, vSphere 5 deployment 13

P

packages 18
password, postgres user 34
passwords 8
pg_config 17
pg_dump 17
pg_restore 17
PL/Perl 26
PL/Python 26
postgres user 25
PostgreSQL 5
psql 17, 25

Q

query plan 35

R

relinking 20
remote connections 25
requirements 12
root user 25
RPM, GUI 29
RPM Files 14
RPM service, stop and start 25

S

schema 32
schema data, creating tables for 32

- SQL management **35**
- SQL query **35**
- SQL query management **35**
- SQL query plan **35**
- start vPostgres service **25**
- stop vFabric Postgres service **25**

T

- tables, creating for schema data **32**
- troubleshooting **27**

U

- user interface **29**

V

- vacuum analyze a database **31**
- view data, examining **33**
- view database entities **31**
- view SQL query plan **35**
- views, creating **33**
- vSphere 5 deployment **13**

W

- Windows packages **18**