



Hewlett Packard
Enterprise

HPE Connected Backup

Software Version: 8.8.7

Product Overview

Document Release Date: July 2017

Software Release Date: July 2017

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent software updates, go to <https://softwaresupport.hpe.com/patches>.

To verify that you are using the most recent edition of a document, go to <https://softwaresupport.hpe.com/manuals>.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support Online web site at <https://softwaresupport.hpe.com>.

This web site provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Access product documentation
- Manage support contracts

- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

To find more information about access levels, go to <https://softwaresupport.hpe.com/web/softwaresupport/access-levels>.

About this PDF version of online Help

This document is a PDF version of the online Help.

This PDF file is provided so you can easily print multiple topics or read the online Help.

Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online Help.

Contents

Chapter 1: Products and Features	7
Connected Backup Products	7
Licensed Product	7
Subscription Service	8
Key Features	8
Reliable Data Backup	8
File Retrieval	8
32-bit and 64-bit Support	8
Deployment Options	9
Support for Wise Transform	9
Agents for Disk Images	9
Digitally Signed Agent Installation Programs	9
A Graphic Interface for Setting Up Agents on Client Computers	9
An Efficient Way to Upgrade Agents	10
Branding Options	10
Options to Efficiently Back Up E-Mail for PC Agent Accounts	10
Support for Encrypted Files, Metadata, and Attributes	10
FDCC Compliance	11
Chapter 2: Product Architecture	12
Product Components	12
Agents	14
File Selection	14
File Compression and Encryption	14
File Transmission	15
Data Center	15
Data Center Servers	15
Data Center Databases	15
Data Center Configurations	16
Stand-alone Configurations	17
Mirrored Configurations	18
Clustered Configurations	20
Non-Mirrored Cluster Configurations	22

Data Center Management Console	24
Archive	24
Life Cycle of an Archive	24
Sizing and Splitting Archive	25
Support Center	26
Account Management Website	26
MyRoam	26
Chapter 3: Management Applications and Utilities	27
Compactor	27
DataBundler	28
Data Center Utilities	28
Data Center Toolkit	28
Chapter 4: Connected Backup Processes	30
File Backup	30
File Retrieval	31
File Management	32
Base and Delta Files	32
Expiration Dates	32
Rebasing	33
Compaction	33
Chapter 5: Connected Backup Deployment	35
Deploy the Data Center	35
Prepare Your Environment	35
Install and Enable MyRoam	36
Gather Information that You Need to Provide During Installation	36
Install the Data Center Software	36
Back Up Certificate Files	37
Verify the Data Center Installation	37
Install Support Center and Account Management Website	37
Use Support Center to Create Agent Communities	37

Use Support Center to Create Agent Configurations	38
Create Agent Rule Sets	39
Deploy the Account Management Website	39
Deploy Agents	39
Index	41
Send documentation feedback	44

Chapter 1: Products and Features

This chapter describes the Connected Backup products and their key features.

- [Connected Backup Products, below](#)
- [Key Features, on the next page](#)

Connected Backup Products

Connected Backup is software that backs up files on users' computers and stores the backed up files in a centralized location. If users lose files on their computers, they can retrieve the lost files from the stored backups.

You can purchase the following types of Connected Backup products:

- A licensed product for users who manage their data storage facility.
- A subscription service for users who back up files to a storage facility that another company manages

You can purchase the following Connected Backup solutions:

- Connected Backup for Mac
- Connected Backup for PC

Licensed Product

Customers who have the licensed product install file storage software on their hardware, install software on users' computers, and use a Web-based application to manage accounts for the computers that host Connected Backup software.

If you use the licensed Connected Backup product, refer to the following documentation:

- *Installing the Data Center*
- *Administering the Data Center*
- *Installing Agents* for your operating system
- *Administering Agents* for your operating system
- *Account Management Web Services Development*

Subscription Service

Customers who subscribe to the Connected Backup service use a storage system that a third party manages. They do not install the storage facility software. However, they install software on users' computers and use a Web-based application to manage accounts for the computers that host Connected Backup software.

If you use the subscription service, refer to *Installing Agents* guide for your operating system and *Administering Agents* for your operating system.

Key Features

The following topics describes the key features of the Connected Backup product.

Reliable Data Backup

Connected Backup uses software, the Agent, that you install on client computers and a centralized storage system, the Data Center, to back up files, operating systems, and applications.

A graphic interface on users' computers lets users review the set of files that they or an administrator specify for backup. If the Agent settings allow, users can edit the files in their backup set.

Administrators can configure the product so that backups start on a scheduled basis. Alternatively, users can initiate backups.

Connected Backup includes Agents for Mac, and Agents for PC.

For more information about how the Agent backs up files to the Data Center, see [File Backup, on page 30](#). Also refer to *Administering Agents* guide for your operating system.

File Retrieval

Users can retrieve one or more backed-up files and restore lost or corrupted files to their computers. Users can retrieve files from the Agent graphic interface or a command-line interface.

For more information about how the Agent retrieves files from the Data Center, see [File Retrieval, on page 31](#). Also refer to *Administering Agents* guide for your operating system.

32-bit and 64-bit Support

You can install the Connected Backup Agent on computers with a 32-bit platform or a 64-bit platform. The installation process is the same regardless of the platform.

You can retrieve backed up files across bit platforms. For example, you can retrieve files from a computer that has a 32-bit platform to a computer that has a 64-bit platform.

Administrators can also create or upgrade a Connected Backup Data Center using 32-bit or 64-bit Windows platforms.

Deployment Options

To enable the Agent to back up files on a computer, you must deploy the Agent on each computer from which you want to back up files. Connected Backup offers several ways to deploy the software. For more information, see [Deploy Agents, on page 39](#). Also refer to *Installing Agents* guide for your operating system.

Support for Wise Transform

You can use the Wise Transform tool instead of the Connected Backup command-line interface to set parameters in the Agent installation file that you deploy to users' computers. For information about how to use transforms, refer to *Installing PC Agents* guide.

Agents for Disk Images

You can install an Agent that does not back up files until a user or administrator activates it. Install this type of Agent when you do not want the Agent to back up files on the installation computer, but you want to include it in a disk image of a computer system that you want to clone to other computers. After users receive a computer that has this Agent, users or a script must activate and register the Agent.

For more information about Agents for disk images, refer to *Installing Agents* guide for your operating system.

Digitally Signed Agent Installation Programs

You can use certificates that you obtain from a trusted certificate authority to digitally sign the Agent Setup file. When you sign the Agent Setup file, users who install the Agent receive a message at the start of the installation process to ensure them that the file comes from a trustworthy source. For more information, refer to Support Center Help.

A Graphic Interface for Setting Up Agents on Client Computers

The Agent Startup Wizard is a graphic interface that, if enabled, helps you configure and register new Agent accounts and recover accounts. The Agent Startup Wizard opens when you start the Agent Setup program for the first time.

For more information about the Agent Startup Wizard, refer to *Installing Agents* guide for your operating system.

An Efficient Way to Upgrade Agents

You can use the Central Administration feature to upgrade Connected Backup Agents on client computers to a new version the next time that they connect to the Data Center. When an Agent connects to the Data Center, the Data Center downloads and installs the new Agent version. For more information, refer to *Installing Agents* guide for your operating system. Also refer to Support Center Help.

Branding Options

The branding feature lets you customize the Connected Backup product to use your company logo. For information about the components that you can brand, refer to *Administering Agents* guide for your operating system. For instructions on how to brand the product, refer to Support Center Help.

Options to Efficiently Back Up E-Mail for PC Agent Accounts

The Connected Backup solution provides the following ways to back up Microsoft Outlook Personal Folder files (*.pst files) and Lotus Notes database files (*.nsf files):

Standard method	Backs up e-mail messages and attachments as a single file. There is no check for duplicate attachments in files backed up by other Agents.
Connected EmailOptimizer method	Analyzes e-mail messages and attachments. Checks for duplicate attachments in files backed up by other Agents, and then backs up only one instance of duplicate attachments. During subsequent backups, the Agent backs up only information that has changed since the last backup, and not the entire attachment. The Connected EmailOptimizer feature is an optional, separately licensed component available for PC Agents only.

NOTE:

The Connected EmailOptimizer is not available for Mac Agents.

Support for Encrypted Files, Metadata, and Attributes

The Agent can back up and retrieve files with many different attributes, including the following types of files:

- Encrypted files.
- Multistream files.

- Files that are open and in use by software other than the Agent when a backup occurs.
- Reparse points.
- Resource forks.
- Security descriptors and ACLs.
- Sparse files.

For more information about the file types the Agent can back up and retrieve, refer to *Administering Agents* guide for your operating system.

FDCC Compliance

As of version 8.4, Connected Backup complies with Federal Desktop Core Configuration (FDCC) 1.2 requirements.

NOTE:

As per the FDCC 1.1/1.2 release, Group Policy Settings do not need to change to run Connected Backup components in FDCC-enabled environments.

For more information about preinstallation requirements for Support Center access in an FDCC environment, refer to *Administering Agents* guide for your operating system.

Chapter 2: Product Architecture

This chapter describes the components that comprise the Connected Backup product and how they work together to back up and retrieve files. It also describes the configuration types that you can deploy.

- [Product Components, below](#)
- [Agents](#) , on page 14
- [Data Center](#), on page 15
- [Archive](#), on page 24
- [Support Center](#), on page 26
- [Account Management Website](#), on page 26
- [MyRoam](#), on page 26

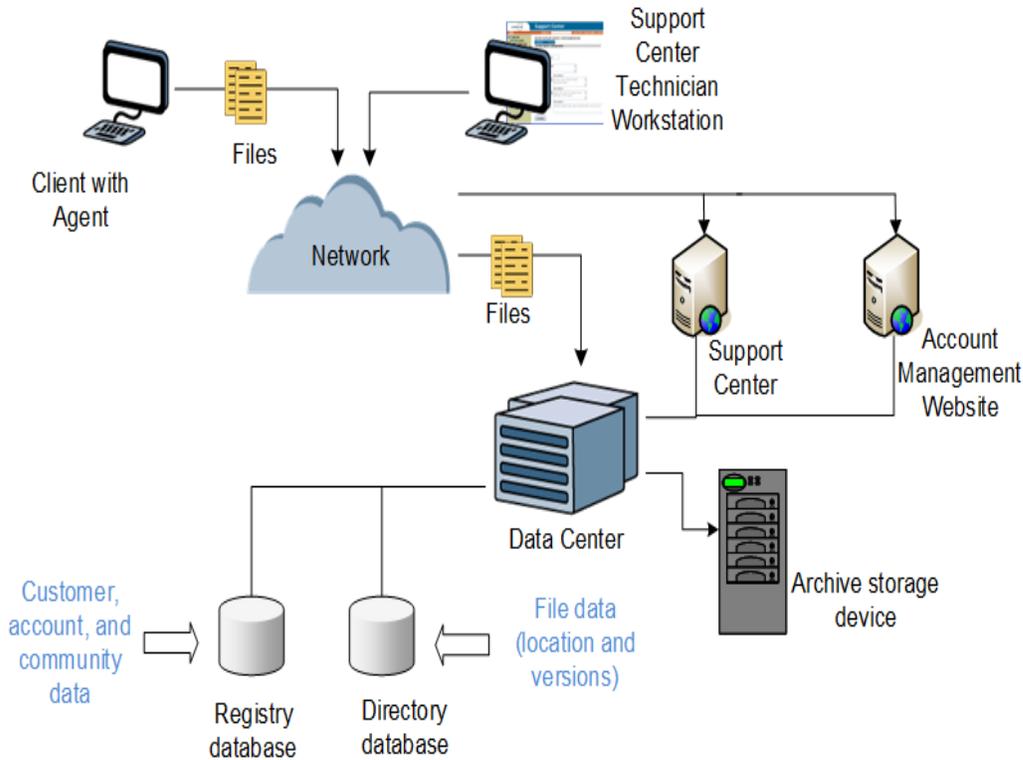
Product Components

Connected Backup consists of the following major components:

Agents	<p>The application that runs on client computers and communicates with the Data Center to back up and retrieve files. The Agent service runs in the background to perform backups. The Agent user interface lets users start backups manually, retrieve files, and view the backup history. For more information about the Agent user interface, refer to the Agent Help.</p> <p>Each client computer that hosts an Agent has one account with the Data Center.</p>
Data Center	<p>The hardware (one or more servers) and software that stores and manages backed-up files. The Data Center supports the following types of Agents:</p> <p>PC Agents</p> <p>Mac Agents</p> <p>The Data Center Management Console is the graphic interface that lets you manage the Data Center.</p> <p>For more information about the Data Center, see Data Center Databases, on page 15.</p>
Connected Backup databases	<p>Databases that the Data Center uses to store backed up data and information about Agent accounts.</p>
Storage devices	<p>Devices that are part of a Data Center that is configured to store files from the Data Center disks to another device.</p>

Support Center	<p>A Web-based application that lets you manage and configure Agents. Although you can install Support Center on the same computer as the Data Center computer, the best practice is to install it on a computer that does not host the Data Center server.</p>
Account Management Website	<p>A Web-based application that lets users manage their Agent accounts. For example, users can update their profile information and request backed up files on media.</p> <p>Although you can install the Account Management Website on the same computer as the Data Center computer, the best practice is to install it on a computer that does not host the Data Center.</p> <p>For more information about the Account Management Website, see Account Management Website, on page 26.</p>
MyRoam application	<p>A Web-based application that lets users retrieve backed up files from a computer that does not host an Agent. The MyRoam application is part of the Account Management Website. It is on the same server as the Account Management Website. For more information, see MyRoam, on page 26.</p>

The following figure shows how the Connected Backup components interact:



Agents

The Agent is software on users' computers that contacts the Data Center to back up and retrieve files.

The Connected Backup software supports the following types of Agents:

- PC Agents that run on Windows PC operating systems
- Mac Agents that run on Apple Macintosh operating systems

Each computer that hosts an Agent has an account in the Data Center. If your computer fails, or if you get a new computer, you can transfer the data associated with your account (such as backup settings and existing backups) to a different computer. An account can exist on only one computer at a time.

Agents perform the following actions:

- Specifies files on the user's computer to back up
- Compresses and encrypts files
- Sends files to the Data Center

Agents initiate all activity with the Data Center. The Data Center does not contact the Agent.

File Selection

The Agent uses rules to determine the location and types of files to select for backup. When a backup session begins, the Agent scans the files selected for backup. If a file is new or modified, the Agent backs it up. For more information about how the Agent uses rules to select files for backup, refer to the Agent Help for your operating system.

After the Agent scans a computer's hard drives for files that are eligible for backup, it contacts the Data Center to initiate a backup session. The Agent also prepares files for backup and notifies the Data Center if you deleted previously backed-up after the last backup session. To retrieve files, the Agent lets you issue requests to the Data Center and download the files.

After the Agent scans the computer's hard drive, the Agent determines which files to back up and processes the files for transmission to the Data Center. For the first system backup, the Agent sends entire files to establish a base for each file. For subsequent backups, it sends only the changes to the file since the last good backup. This change is called a delta.

File Compression and Encryption

Before the Agent sends files to the Data Center, it uses the ZLIB compression library to compress each file. ZLIB is a compression algorithm that reduces the size of a file. The Agent also encrypts each compressed file.

File Transmission

The Agent sends the compressed, encrypted files to the Agent's primary server on the Data Center. To transmit the files, the Agent uses TCP/IP over a local area network (LAN), wide area network (WAN), or a dial-up modem connection. In response to the transmission, the Data Center creates a new file, called an archive, and begins to receive backup information from the Agent. For more information about archives, see [Archive, on page 24](#).

Data Center

A Data Center consists of the hardware (one or more servers) and software that stores and manages backed up files. This section describes the components of the Data Center.

Data Center Servers

The Data Center has the following servers:

Backup server	Receives and stores backup data from Agents, and lets Agents retrieve backed-up data upon request
Web server	Runs Support Center and the Account Management Website (which includes the MyRoam application)

Data Center Databases

When you install the Data Center software, it creates the SQL databases that it uses to store information about files.

The following table describes each database and how the Data Center uses it:

Database	Description
Registry database	The Data Center uses the Registry database to perform the following tasks: <ul style="list-style-type: none">• Record account activity (for example, the creation of new Agent accounts or changes to account information).• Generate reports within Support Center.• Store user information.• In mirrored configurations, obtain information about the mirror. For information about mirrored configurations, see Mirrored Configurations, on page 18.
Directory database	The Data Center uses the Directory database to store the following information:

Database	Description
	<ul style="list-style-type: none"> • Archives. • User files in a given archive. • Locations of each archive on disk or on an archive storage device. • Information about replication and Compactor. For information about Compactor, see Compactor, on page 27. <p>The Data Center uses this database when a user tries to retrieve data.</p>

If you host your Data Center and you run it as a stand-alone server, the Data Center backs up the databases during the automatic daily and weekly maintenance. If a disaster occurs, you can use the backed-up database files and transaction logs to reconstruct the Data Center. For more information about these procedures, refer to *Administering the Data Center* guide.

If you host your own Data Center and you run it on a mirrored or clustered configuration, you can rebuild the databases from the mirror. For more information about the procedures that you can use to rebuild the databases, refer to *Data Center Disaster Recovery* guide.

During the course of daily activity, the Data Center replicates database tables between the mirrored servers. Several tables store information that is unique to each server. The Data Center also replicates these tables. For example, the Data Center replicates the ArchiveeClip table on Server A to a table named MirrorArchiveeClip on Server B. If Server A database fails, you can copy the database tables from Server B. This process is faster and more reliable than if you restore from database backups. For more information about the ReplicationServer service, refer to *Administering the Data Center* guide.

Data Center Configurations

You can set up the Data Center in the following configurations:

Stand-alone server	A Data Center with only one server.
Mirrored pair	A pair of Data Center servers that contain identical data. Two identical servers protects against data loss if a disaster affects either server in the pair. It also enables Agents to connect to a secondary server if the primary server is not available.
Clustered Data Center	Two or more pairs of mirrored servers.
Non-	A non-mirrored cluster with a single Registry Master Server and multiple Directory

Mirrored Cluster	Servers.
------------------	----------

Stand-alone Configurations

Stand-alone configurations have the following components:

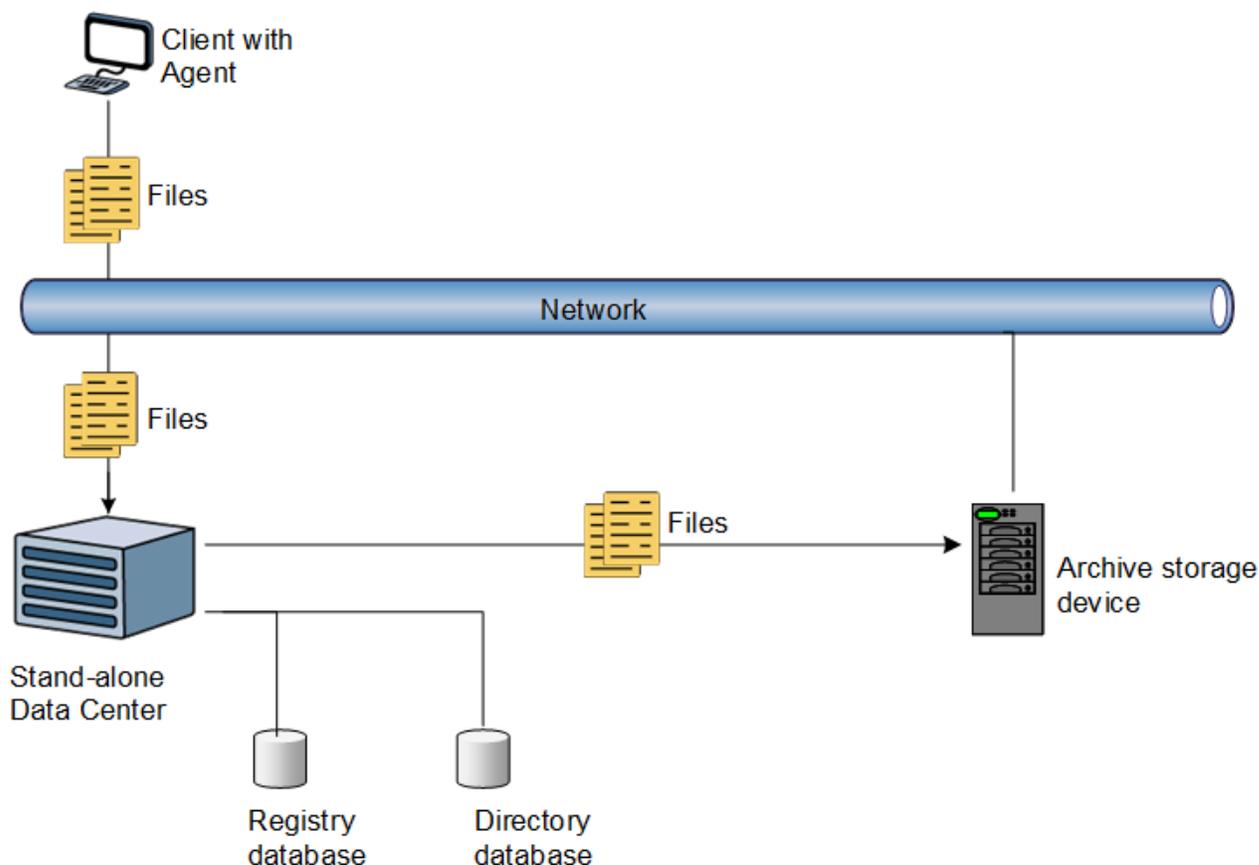
- A Data Center server that processes requests for new account, backups, and retrievals.
- A Web server that hosts Support Center and optionally, Account Management Website and the MyRoam application
- A storage device that stores files from the Data Center

The Data Center server stores backed up data on local disks. In stand-alone configurations, the data on the local disk is the only copy of the data. If a local disk fails, you lose the data. To protect your data more securely, you can use any of the following additional storage methods:

Mirrored configuration	A configuration of two identical Data Centers that provide redundant storage of data and failover backup and retrieval service. For information about mirrored configurations, see Mirrored Configurations, on the next page .
Redundant Array of Independent Disks (RAID)	A device that uses two or more disk drives. If one drive fails, the other takes over.

If the Data Center server is unavailable (for example, during maintenance), the Agent must wait for it to become available before the Agent can back up or retrieve files. In a mirrored configuration, if one server is unavailable, the Agent connects to the mirror server. In a mirrored environment, users remain unaware of server downtime.

The following figure shows a stand-alone configuration that uses an attached archive storage device:



Mirrored Configurations

Mirrored configurations use two identical Data Centers, ideally at physically separate sites. Each Data Center uses the same software, stores the same data, and services the same Agent communities. The two mirrored servers are redundant peers. Neither server is dominant over the other. Each server can perform any backup or retrieval operation.

When a Data Center server creates a new user account, the Agent assigns one of the mirrored Data Center servers as the primary Data Center for the account. An Agent typically contacts its primary Data Center. If the primary Data Center is unavailable, the Agent contacts the mirror instead. The alternating assignment of Data Centers as primary to individual accounts facilitates load balancing. This ensures that an approximately equal number of sessions connect with each Data Centers.

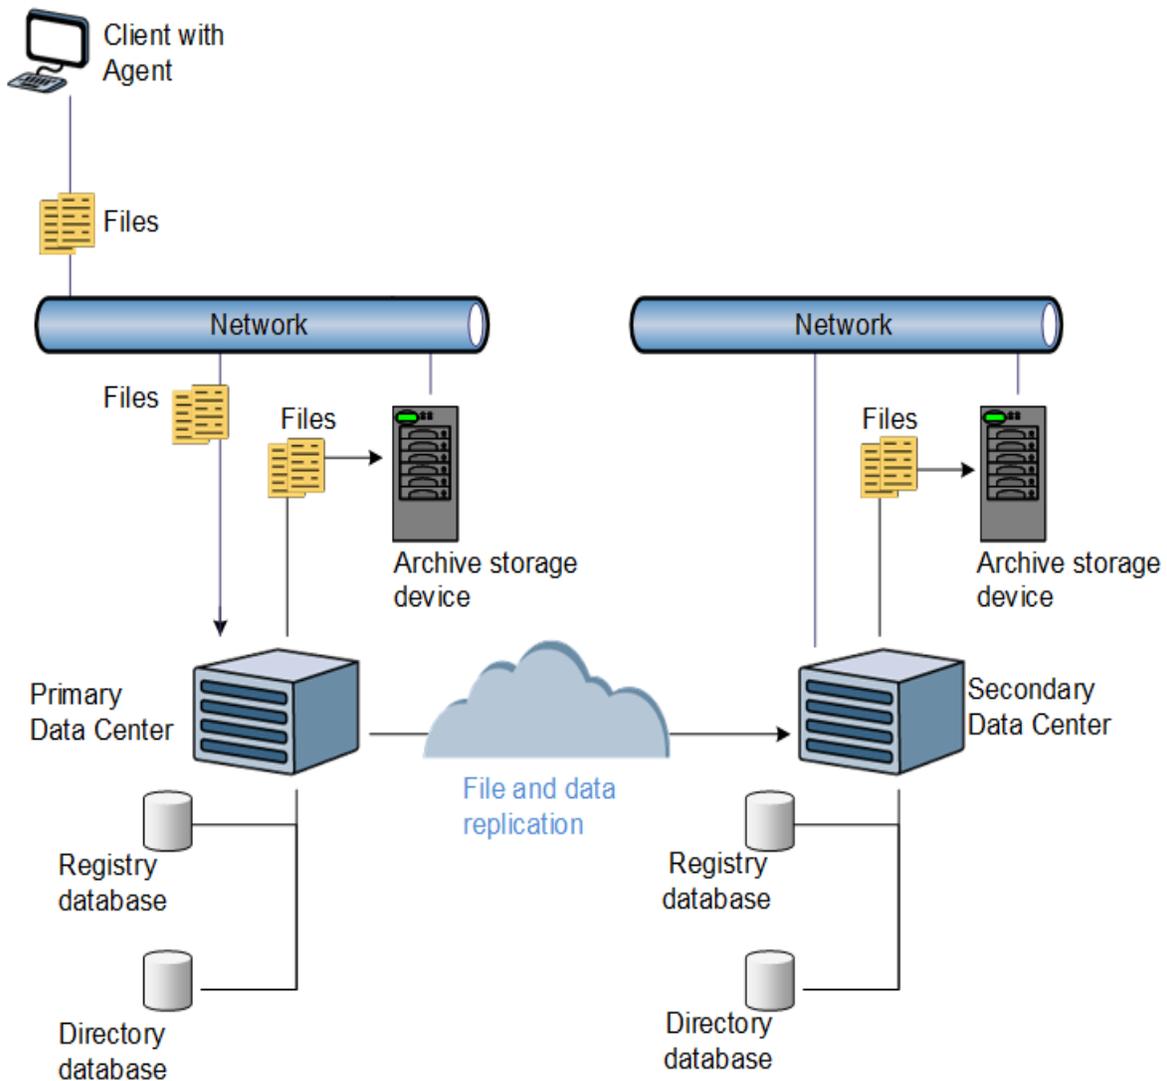
Either the primary or mirror Data Center can create new accounts or modify information in existing accounts. Any change to the Registry database duplicates across mirrored Data Centers so the Registry databases on the two mirrored servers have identical content. For more information about the Registry database, as well as other Data Center databases, see [Data Center Databases, on page 15](#).

When the Agent performs a backup, one Data Center receives the backup set, stores the data as an archive on its disk, and records it in the Directory database. To protect the data against loss, the Data

Center software replicates the data to the mirror server, and then records it into the mirror's Directory database. Every archive is stored redundantly on both Data Centers, initially on the disk and then on the server's archive storage device (if you use one). Therefore, either Data Center server can retrieve any backed up file, even if the backup originally went to the other Data Center server.

A Data Center can receive an archive when the mirror server is down (or while communication between the mirrored servers is down) and the mirror cannot replicate the archive immediately. When this occurs, the replication is deferred. The Data Center performs all deferred replications when the mirror resumes operation. This process to re-establish equivalency of data between the mirrors is called recovery of replication or resynchronization.

The following figure shows a mirrored server configuration:



The figure depicts a configuration that uses the highest level of data protection — a mirrored Redundant Array of Independent (or Inexpensive) Disks (RAID) with an attached storage device. Each side of the mirror is in a different location. Although the figure shows Agents for only the Primary Data Center, an equal number of Agents use the Secondary Data Center as their Primary Data Center. Both Data Center servers process requests over the network for new accounts and accept Agent backup and retrieval requests.

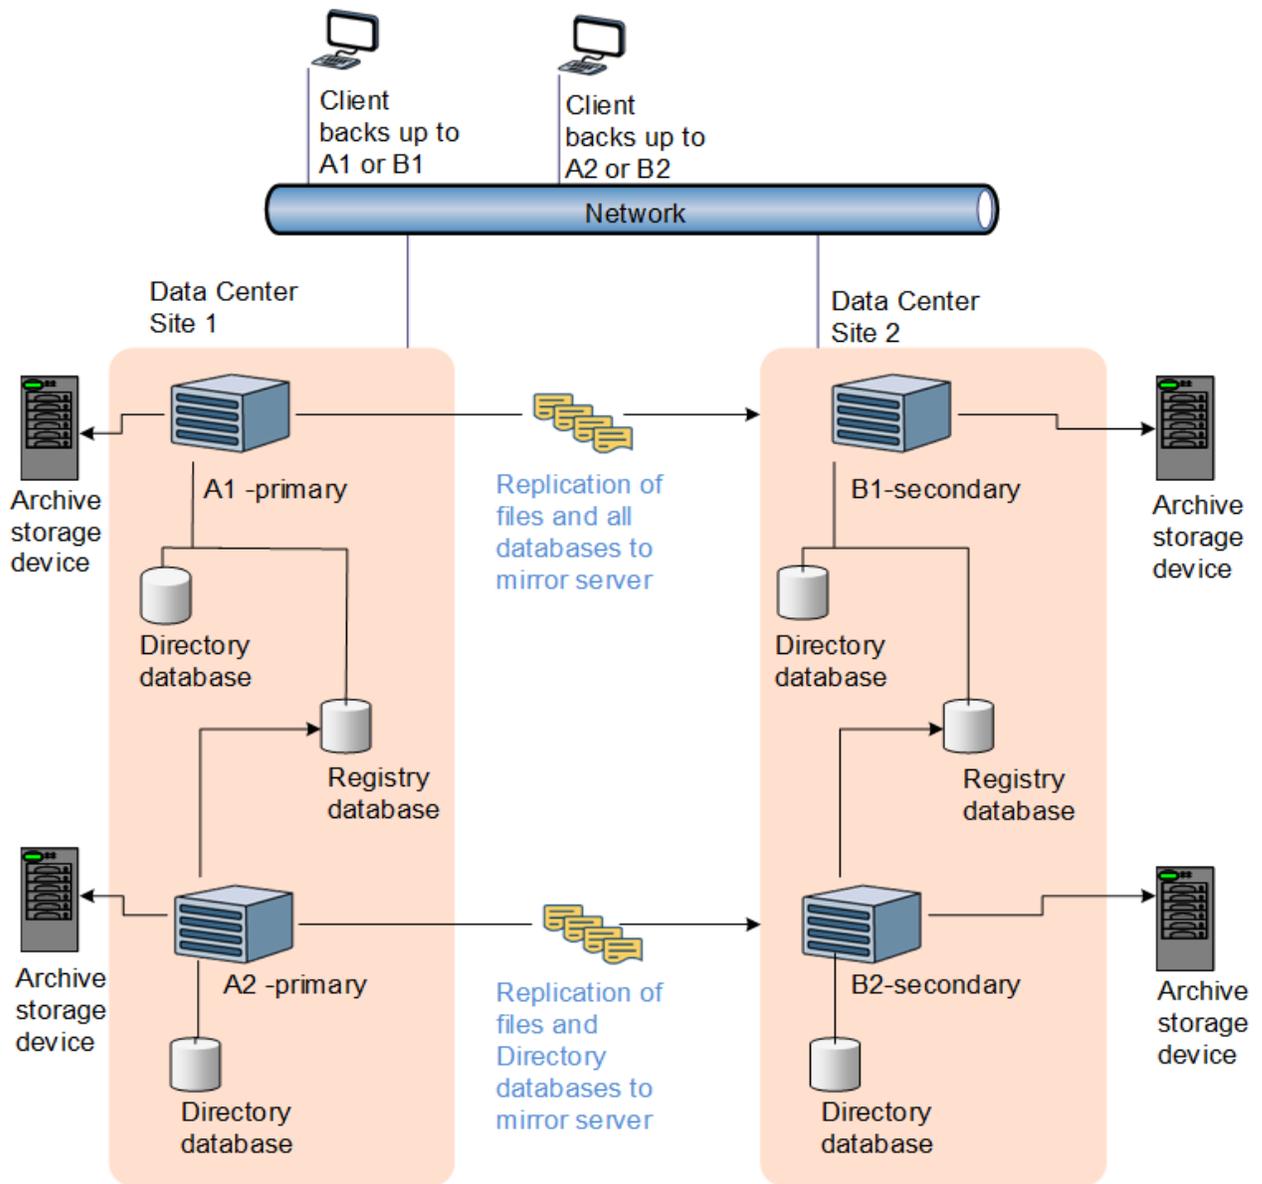
Clustered Configurations

For Data Centers that serve a large number of users or receive a large amount of data, a mirrored configuration might not have sufficient capacity to handle the backup and retrieval activity. In this situation, you might want to set up clusters at each Data Center site.

To balance loads, clustered Data Centers distribute backup and retrieval requests to more than one pair of mirrored servers. The mirrored servers in the cluster provide load balancing, redundant storage of data, and fail-safe availability of file backup and retrieval services.

Clustered configurations have two or more pairs of Data Center servers. Each server in a pair resides on a Data Center at a different location. Each mirrored pair of Data Center servers provides backup services to a subset of the user community and failover for each other. The group of Data Center servers is a **cluster**. The Data Center servers in the cluster share the same Registry database in SQL, but each mirrored pair of servers has its own Directory database.

The following figure shows Data Center servers A1 and B1, which form a mirrored pair, with archive storage devices for extended storage of backed up data:



In this configuration, the Data Center servers A1 and B1 share the same set of users and maintain identical Directory databases. The same is true for Data Center servers A2 and B2, which are mirrors of each other. Servers A1 and B1 serve a subset of their user community, while servers A2 and B2 serve the other half. Servers A2 and B2 use the Registry databases that reside on A1 and B1 respectively.

In this configuration, each individual user account is assigned to a primary server at the time the account is created. Thereafter, the user can back up and retrieve data to the primary server and retrieve from there or the server's mirror. For example, if an account is assigned to server A2, it attempts to back up to A2, but will back up to B2 if A2 is unavailable. It never backs up to A1 or B1.

Although this figure shows only the servers at Data Center site 1 as primary Data Centers, an equal number of Agents use the servers at Data Center site 2 as their primary Data Center.

Non-Mirrored Cluster Configurations

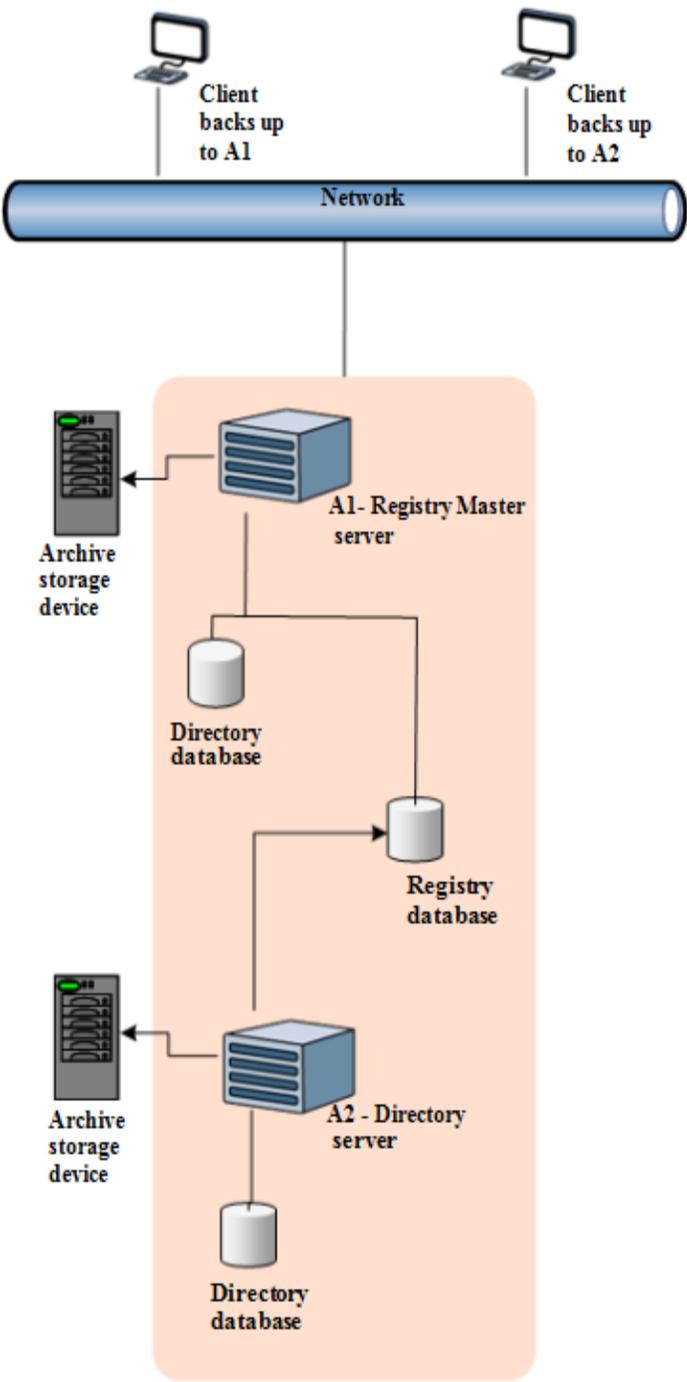
The non-mirrored cluster configuration is an extension to the concept of a Stand-alone Data Center server. This configuration contains one registry master server and multiple directory servers where directory servers are installed individually instead of pairs.

To provide geographical affinity to a single server, a non-mirrored cluster configuration can be setup. Based on the number of agent requests, you can subsequently add more directory servers to the site.

For load balancing, the Data Center assigns group of agents to a specific server in the cluster to manage the backup or retrieval requests. Each server in a cluster is associated with unique ID. Therefore, the agents attempt to connect only to the assigned server.

If the Data Center server to which the agent is assigned is unavailable (for example, during maintenance), the Agent must wait for the Data Center to become available before the Agent can back up or retrieve the files.

The following figure shows the Data Center servers, A1- Registry Master server and A2-Directory Server which forms a non-mirrored cluster, with archive storage devices to store the backed up data.



In this configuration, Server A2 uses the registry database of the Registry Master server. A subset of agents are assigned to the Server A1 and the remaining agents to the Server A2. If the registry database on the Registry Master server goes down, for example, services hang or other failure, this impacts all the servers (A1 and A2) in the cluster and the servers cannot process the agents request until the registry database is available.

Data Center Management Console

The Data Center Management Console (DCMC) is a snap-in for the Microsoft Management Console (MMC). DCMC lets system administrators monitor and control Data Center operations such as the Data Center services and Compactor. For more information about DCMC, refer to DCMC Help. For information about the Data Center services and Compactor, see [Management Applications and Utilities, on page 27](#).

Archive

An archive is a file in the Data Center that contains the compressed and encrypted files that an Agent transmits to the Data Center during one backup. The size limit of each archive is 90MB. This size limit optimizes performance of data recovery and preserve Data Center disk space.

The data in a single archive originates from one account. The header for the archive file includes the account number for identification in mirrored or clustered Data Center environments. The Data Center software replicates archives to each Data Center server's mirror.

Archive files have .arc file extensions. Connected Backup writes these files to a Customers directory on a Data Center server volume. The contents of a single archive always originate from the same account. In disk-only Data Centers that contain multiple volumes, archives from one account can reside on different volumes. Writing archives to the volume with the most available space evenly distributes disk consumption.

The following table describes the basic elements in an archive:

Element	Description
Archive header	The archive header contains information about the archive file. This information includes the name of the file, the account from which the data originated, and the current version of the Agent in use when the Data Center created the archive.
File header	The file header or headers contain detailed information about the files in the archive. It identifies whether the file is a baseline or delta, the current revision level (if it is a delta), and its path and file names. When you delete a file from the client or exclude a file from the backup list, the Agent transmits just a file header that signals to the Data Center that the user deleted or excluded the file.
File data	The file data in an archive consists of the files in their compressed and encrypted form.

Life Cycle of an Archive

The life cycle of an archive depends on the configuration of the Data Center where it resides. From the first transmission of account data to the Data Center, the life cycle of an archive proceeds as follows:

1. The Data Center BackupServer service distributes the client data into one or more archives on the volume that has the greatest amount of available disk space.
2. The Data Center IndexServer service indexes the archives and their individual contents in the Directory database to track each revision of each backed up file.
3. In mirrored or clustered Data Centers, the ReplicationServer service replicates the archives to the mirror (if applicable).
4. The Compactor service routinely checks each account's archives for expired revisions. It then freshens (repackages) the archives as needed.

For more information about the Data Center services that archive, index, and replicate data, refer to *Administering the Data Center* guide.

Sizing and Splitting Archive

When an Agent transmits account backup data to the Data Center, the BackupServer service groups it into archives. During this process, BackupServer regulates the size of the archives and creates new archives for the client data. The BackupServer creates a new archive under the following conditions:

- The archive reaches 90 MB in size.
- The number of files in the archive totals 2,000.
- The next file to come from the client is greater than 225 MB.

The BackupServer service creates a new archive as needed to efficiently group the data into the smallest possible units. For example, files A, B, and C together represent 40 MB, but file D equals 60 MB, then BackupServer combines A, B, and C together in one archive and creates a new archive for file D.

Archives can include information from many individual files. Baselines of new files (that is, backup information for a file's first backup) and deltas of older files (that is, backup information for a file previously backed up) can reside together in the same archive. However, backup information for a single file exists in multiple archives. This is because new archives of account data are created during every backup session. For example, a large file of 10 MB that undergoes significant revisions between each of five backup sessions have its backups stored in five archives. The first time the Agent backs up a file to the Data Center, the Data Center writes a baseline of the file to one archive. For the four subsequent revisions, the Data Center writes each file deltas to four additional archives.

NOTE:

The data in a single archive comes from one backup session from a single account.

For more information, see [File Management, on page 32](#).

Support Center

Support Center is a Web-based management application that lets you create and manage Agent configurations and registered accounts. For information about how to use Support Center, refer to Support Center Help and *Administering Agents* guide for your operating system.

Account Management Website

The Account Management Website is a Web-based application that lets users perform the following tasks:

- Register a new account and download the Agent software.
- View account information.
- Modify profile information such as e-mail address and password.
- Recover an account.
- Use the MyRoam application to retrieve files.
- For PC Agents, order images of backed-up files on different types of media. You cannot use media images with Mac Agents.

To access to the Account Management Website, you can enter the URL to the Website in a Web browser, or select a link in the Agent user interface. The URL for the Account Management Website is case-sensitive. Technicians can select a link in Support Center to gain access to users' account information on the Account Management Website. When technicians gain access to accounts in this manner, they can gain access to only the features and information available to the user.

MyRoam

The MyRoam application is an optional, Web-based application that lets users retrieve backed-up files using a Web browser instead of the Agent user interface. The MyRoam application runs on the same secure Web server as the Account Management Website.

To gain access to the MyRoam application, users must log on to the Account Management Website.

Chapter 3: Management Applications and Utilities

This chapter describes the Connected Backup applications and utilities that you can use to manage files.

- [Compactor, below](#)
- [DataBundler, on the next page](#)
- [Data Center Utilities , on the next page](#)
- [Data Center Toolkit, on the next page](#)

Compactor

Compactor is a Data Center service that removes data from the Data Center. Using the Compactor service is beneficial for the following reasons:

- Reduces overall storage requirement for the Data Center.
- Improves Agent file retrieval performance.
- Reduces the size of the databases.
- Improves data integrity.

Compactor runs on all Data Center configurations. However, it runs differently on a stand-alone server and non-mirrored cluster configuration than it does in a mirrored configuration.

The following table describes how Compactor runs on each type of Data Center server configuration:

Data Center Configuration Type	How Compactor runs
Stand-alone	Runs on one server.
Mirrored	Runs on two servers. However, only the primary server controls the workload of the compaction process.
Clustered	Runs on one primary server for every mirrored pair in the cluster. For example, a clustered Data Center with three mirrored pairs has three primary servers. You can check the status of the primary server(s) in the Compactor view of DCMC. For instructions, see DCMC Help.
Non-Mirrored Cluster	Runs on each server in the cluster.

For information about the compaction process, see [Compaction, on page 33](#).

DataBundler

The DataBundler application creates images that contain backed-up data for individual PC Agent accounts. If you are a licensed customer, you use your own software and hardware to burn the DataBundler images onto a CD or DVD. If you are a subscription customer, you request the images on CD or DVD from the organization that manages your Data Centers.

Users can recover Agent accounts or retrieve files from the images that DataBundler creates. Images are useful when you need to retrieve files but a client system is unable to connect to a Data Center, or if you need to retrieve a large number of files and the client system uses a slow connection to the Data Center. You cannot use Mac Agents to retrieve files from images.

When the DataBundler application receives a request for an account image, it interacts with the Data Center to gather all files for a selected account and assemble them into a compressed account image. The account image includes a copy of the Agent, enabling users to recover their accounts or retrieve some or all files through the familiar Agent user interface. The DataBundler application also creates shipping and disk labels you can print using your own software and hardware.

Technicians use Support Center to request an image on media for PC Agents. Users open the Account Management Website to request images on media.

For additional information about the DataBundler application and how to order images on media, refer to the following documentation:

- If you host your Data Center, for installation instructions, refer to *Installing the Data Center* guide and *Administering the Data Center* guide.
- For more information on how to use DataBundler, refer to DataBundler Help.
- For information about how to order account images on media, refer to Support Center Help and Account Management Website Help

Data Center Utilities

The Data Center installation program installs a number of utilities for Data Center maintenance. If you do not host a Data Center, you do not need these tools.

The following table describes each tool. For more information about the Data Center management utilities, refer to *Administering the Data Center*.

Data Center Toolkit

The Data Center Toolkit is a collection of Data Center utilities that you can use on a computer other than the Data Center server (for example, to monitor the status of the BackupServer service remotely).

You can install the Data Center Toolkit on as many computers as you like.

The Data Center Toolkit includes the following utilities:

- Data Center Management Console (DCMC)
- DC Message Logger
- Event log message support files
- Disk Status
- Tdate Converter
- dsping

For more information about how to install the Data Center Toolkit, as well as about each utility, refer to *Administering the Data Center*.

Chapter 4: Connected Backup Processes

This chapter describes the processes that Connected Backup uses to back up, retrieve, restore, and manage files.

- [File Backup, below](#)
- [File Retrieval, on the next page](#)
- [File Management, on page 32](#)
- [Compaction, on page 33](#)

File Backup

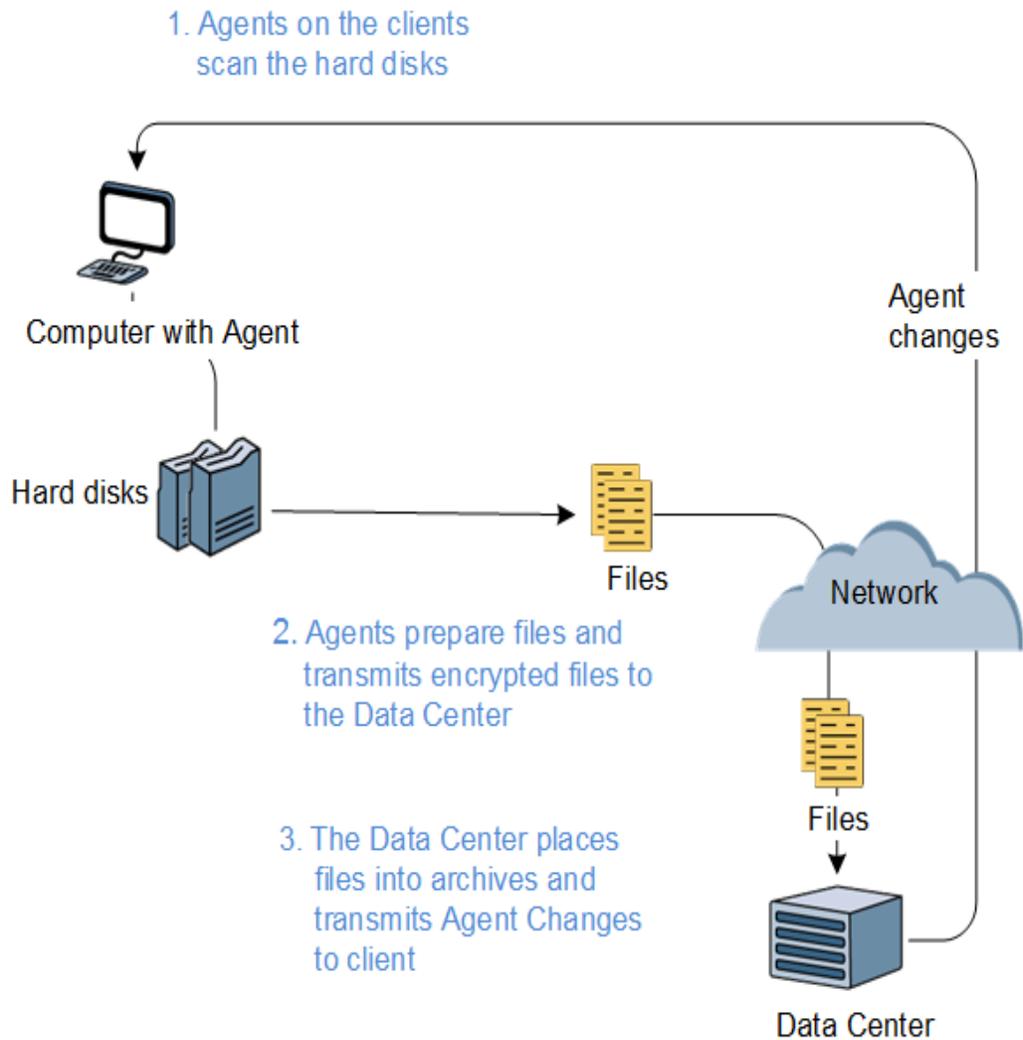
The Agent backs up new files and files that have changed since your previous backup. When the Agent backs up the changed files, it backs up only the changed portions of the files. When the Agent sends files to the Data Center, the Data Center compares the files to those backed up by other users. Duplicate files are not sent to the Data Center, using SendOnce technology.

You can use the Agent to review the set of files specified for backup and, if your Agent settings allow, edit the files in your backup set.

You can configure the Agent so that backups start on a scheduled basis, whether the Agent is open or closed. You can also use the Agent graphic interface or command-line interface to initiate a backup.

For more information about file backup, refer to *Administering Agents* for your operating system.

The following figure shows the backup process:

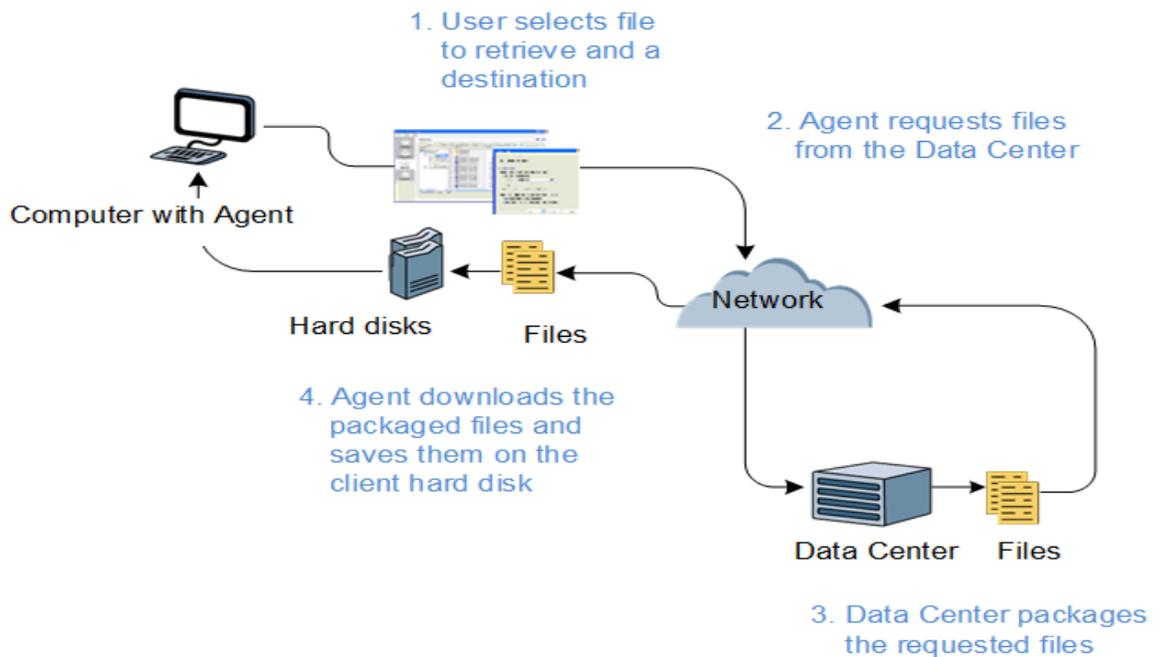


Agents synchronize with the Data Center before and after a backup. During synchronization, the Agent receives configuration changes, messages, and upgrades, and it uploads Agent history logs.

File Retrieval

You can use the Agent to retrieve one or more backed-up files from the Data Center and restore files that have been lost or corrupted on your client computer. To retrieve files, you can use the Agent graphic interface or the command-line interface. You retrieve files through a network connection to the Data Center. If you have PC Agents, you also can retrieve files from an image stored on a CD or DVD.

The following figure shows an example of the Retrieve process.



For more information about the Retrieve feature, refer to *Administering Agents* for your operating system.

File Management

This section describes how the Data Center handles files during backup.

Base and Delta Files

During the backup process, the Agent maintains a complete copy of the base file, and then captures the changes that occur between backups in delta files. The delta files are not complete pictures of the original file. They are only the changes to the file that occurred since last backup. After several backups, one file can accumulate many delta files, and reconstitution of that file for a retrieval request requires the Agent to get the baseline and all the accumulated deltas.

Expiration Dates

At some point, the original base file and some of its deltas expire. The expiration date depends on the expiration settings that you set for the Data Center. For example, some organizations, such as lawyers or medical offices, might turn off expiration dates for legal reasons.

Compactor, a Data Center service, searches the archives for expired files. Compactor evaluates each file. When Compactor determines that the base file or the base file and several delta files are expired, it

creates a new baseline or rebases the file. For more information about Compactor, see [Compaction](#), below.

Rebasing

Rebasing combines the original base file and expired deltas to create a new base. Compactor must rebase a file before its expired base and deltas can be deleted. It extracts the base file and deltas from the archive or archives and rolls them up into a new file. After Compactor creates the new base file, it deletes the expired files and then repackages the remaining unexpired files into one or more new archives.

A new file is created and backed up from the account on March 17th (3/17). Subsequent modifications and backups occur over the next two days and then cease for that file.

archive1	base1	created and backed up on 3/17
archive2	delta1	modifications backed up on 3/18
archive3	delta2	modifications backed up on 3/19

The next time Compactor reviews this account, it rolls up the data in base1, delta1, and delta2 to create a new baseline, base2:

archive1	base1	expired on 5/17
archive2	delta1	expired on 5/18
archive3	delta2	expired on 5/19
archive6	base2	rolled up file created and archived on 5/31

Compaction

The compaction process removes data from the Data Center. To remove data, Compactor, a Data Center service, performs the following process:

1. Selects accounts.
2. Analyzes and repairs the system.
3. Marks files as expired.
4. Repackages archives.
5. Deletes expired archives and database entries.
6. Informs the Agent of changes.

These tasks represent a standard compaction cycle on a mirrored Data Center. For more information about each task, refer to *Administering the Data Center*. For more information about Compactor, see [Compactor, on page 27](#).

Chapter 5: Connected Backup Deployment

This chapter describes, in general, the tasks that you must complete to deploy Connected Backup. The chapter presents the tasks in the order in which you must perform them.

- [Deploy the Data Center, below](#)
- [Install Support Center and Account Management Website, on page 37](#)
- [Use Support Center to Create Agent Communities, on page 37](#)
- [Use Support Center to Create Agent Configurations, on page 38](#)
- [Create Agent Rule Sets, on page 39](#)
- [Deploy the Account Management Website, on page 39](#)
- [Deploy Agents, on page 39](#)

Deploy the Data Center

If you want to host your Data Center, you must install the Data Center software. This section describes, in general, the tasks that you must perform to install the Data Center. For detailed information, refer to *Installing the Data Center*.

Prepare Your Environment

To prepare your environment

- Determine which Data Center configuration best suits your organization.

For information about types of Data Center configurations, see [Data Center Configurations, on page 16](#).

- Determine which licensing option best suits your organization.

You must license every Data Center. You can purchase licenses for PC accounts, Mac accounts, or any combination of these accounts. Before you install the Data Center software, you must obtain a permanent license. However, if you do not have a license when you install a Data Center, the Data Center Setup program creates a temporary license that expires in 30 days.

Your Data Center license gives you access to optional features. Before you purchase your licenses, decide which features you want to use. The license software also tracks the number of active users and servers on the Data Center and warns you when your contracted license agreement nears capacity.

- Review the following requirements:
 - Data Center server requirements.
 - Storage solutions requirements.

- Network requirements.
- Security requirements for your Data Center.

For information about these requirements, refer to *Installing the Data Center*.

- Install and configure the required Microsoft software on the computer that will host the Data Center software.

The Data Center requires the following Microsoft software:

- Microsoft Windows Server
- Microsoft SQL Server

For a list of required software, refer to *Connected Backup Requirements Matrix*. For instructions about how to configure the software, refer to *Installing the Data Center*.

- Prepare the servers that will host Support Center and Account Management Website.

The tasks that you must perform to prepare servers to host Support Center and Account Management Website depend on your Data Center configuration. For more information, refer to *Installing the Data Center*.

Install and Enable MyRoam

MyRoam lets users retrieve backed-up files to computers that do not host the Agent. After users select the files that they want to retrieve, the Data Center creates a ZIP archive file that contains the selected files. Users download this file and extract it to their computer.

To use MyRoam, users must log on to the Account Management Website. For more information about how to install and enable MyRoam, refer to *Installing the Data Center*.

Gather Information that You Need to Provide During Installation

Before you install the Data Center software, complete the worksheets in Worksheets and Checklists in *Installing the Data Center*. Be sure to store the Data Center information in a secure location.

Install the Data Center Software

After you prepare you environment and install MyRoam, you can install the Data Center software. Use the information you recorded in “Appendix A: Worksheets and Checklists” in *Installing the Data Center*. As you perform the installation, complete “Appendix B: Data Center Installation Worksheets.”

The Data Center installation process creates a Master Encryption Key (MEK). Store this key in a safe place. You might need the MEK if you reinstall the Data Center software.

Back Up Certificate Files

After you install the Data Center, back up the certificate files that the software uses. Because these files do not change, you have to back them up only once. For a list of certificate files that you have to back up, refer to *Installing the Data Center*.

Verify the Data Center Installation

To make sure that the Data Center installation functions, verify the following components operate properly:

- BackupServer
- Mirror site replication (if you use a mirrored or clustered configuration)

For more information about how to verify the Data Center installation, refer to *Installing the Data Center*. The guide also contains information about how to troubleshoot problems with the installation.

Install Support Center and Account Management Website

Support Center is a Web-based management application that lets you create and manage Agent configurations and registered accounts.

Account Management Website is a Web-based application that allows users to perform the following tasks:

- Register a new Agent account and download the Agent software.
- View Agent account information.
- Modify profile information such as e-mail address and password.
- Recover an Agent account.
- Use the MyRoam application to retrieve files.
- For PC Agents, order images of backed-up files on different types of media.

You must install the Account Management Website component on a separate server from the Data Center.

Use Support Center to Create Agent Communities

A community is a group of Agent accounts that share common characteristics, such as configuration settings, the geographic location of the clients, or similar bandwidth throttling requirements. You can use communities to run reports, edit settings, upgrade Agent configurations, and manage accounts.

Each community that you create in Support Center appears as a node in the left pane of the Support Center interface. The node contains subnodes that represent the accounts, technicians, reports, configurations, and subcommunities in the community. You can create as many communities and subcommunities as necessary to organize your accounts.

For more information, refer to *Administering Agents* for your operating system. Also refer to Support Center Help.

Use Support Center to Create Agent Configurations

You create Agent configurations in Support Center to specify the features that are in the Agents that you deploy. An Agent configuration contains the following components:

Profile and Website settings	Determine the information and features that are available to users in the Account Management Website. They also determine the information that is in users' profiles, and the fields that are in the Agent Startup Wizard.
Components list	<p>Support Center compiles Agent configurations from a set of components that each specify some of the Agent's features. Agent configurations can contain the following components:</p> <ul style="list-style-type: none">• Agent versions• Agent settings• Agent rule sets <p>For more information, refer to <i>Administering Agents</i> for your operating system. Also refer to Support Center Help.</p>
Default Agent configuration	<p>The Support Center installation creates a default Agent configuration (with a name similar to Default PC version Configuration or Default Mac version Configuration). The default configuration resides at the root level of the tree in Support Center and is available to all your communities. You might want to create your own configuration with a unique version, settings, and rules. However, you can use the default configuration as a reference when you create your own configurations.</p> <p>The default configurations contain the following components:</p> <ul style="list-style-type: none">• Default Agent Version• Default Agent Settings• Default Agent Rule Sets• Default Agent Profile and Website Settings <p>For more information, refer to <i>Administering Agents</i> for your operating system. Also refer to Support Center Help.</p>

Create Agent Rule Sets

The Agent uses rules to determine which files to include or exclude from backups and recoveries. Rule types determine the order in which the Agent processes rules. For more information about rules, refer to *Administering Agents* for your operating system. Also refer to Support Center Help.

Deploy the Account Management Website

The Account Management Website lets users download Agent configurations, manage their account information, retrieve files (if you enable MyRoam), or, for PC Agents, order backed-up files on media such as CDs or DVDs when they work from a computer that does not host an Agent. For information about how to deploy the Account Management Website, refer to *Installing the Data Center*.

Deploy Agents

To deploy Agents, you or end users install them on client computers. Choose from the following installation methods:

User installation	Users register the Agent, and then download the Agent Setup file from the Account Management Website. The installation process prompts users for set up information.
Distributed user installation	<p>Administrators download the Agent Setup file from Support Center and distribute it to users. The installation process prompts users for set up information. After users install the Agent, they must register it. If you enable the Agent Startup Wizard, the user experiences registration as a continuation of the installation process.</p> <p>Use this type of installation when you want users who do not have access to the Account Management Website to install the Agent themselves.</p>
Disk image installation	<p>Administrators install an inactive Agent that does not back up files until it is activated by a user. Install this type of Agent when you want only to include the Agent in a disk image of a computer that you want to clone to other computers.</p> <p>After users receive a computer that has an inactive Agent, they must activate and register the Agent. If you enable the</p>

	Agent Startup Wizard, the registration program runs the first time the user opens the Agent.
Hands-free installation	Administrators download the Agent Setup file from Support Center and use the installation command-line interface to install the Agent on multiple computers simultaneously without user participation.

For more information about the installation methods, refer to *Installing Agents* for your operating system.

Index

3

32-bit support 8

6

64-bit support 8

A

Account Management Website 26

definition of 13

additional storage methods 17

Agent Startup Wizard 9

Agents

central upgrade 10

definition of 12

deploying 39

description 14

for disk images 9

selecting files for backup 14

types of 14

architecture diagram 13

archive storage devices

definition of 12

archives

and expiration rules 32

sizing and splitting 25

splitting 25

attributes 10

B

backup

encrypted files 10

backup server 15

base files 14, 32

branding 10

C

cluster configurations 20

diagram 20

communities

assigned server pair 21

on clustered Data Center 21

compaction 33

Compactor 33

description 27

expiring contents of archives 32

components of Connected Backup 12

compression 14

configurations

cluster 20

mirrored 18

stand-alone 17

Connected Backup

components 12

databases, definition of 12

D

Data Center

archives 24

archives, lifecycle of 24

configurations 16

databases 15

definition of 12

load balancing 18

resynchronization 19

servers 15

toolkit 28

utilities 28

databases 12, 15

directory 15

registry 15

replication of 16

DataBundler, description 28

delta files 14, 32

deploying Agents 39

diagrams

architecture 13

digital signature 9

Directory database 15

disk image

Agent 9

- installation 39
- distributed user installation 39

E

- e-mail backup 10
- EmailOptimizer 10
- encrypted files 10
- expiration of files 32

F

- features
 - Agent deployment options 9
 - Agent Startup Wizard 9
 - Agents for disk images 9
 - backup 8
 - branding 10
 - central upgrade for Agents 10
 - cross platform support 8
 - digitally signed installation program 9
 - e-mail backup 10
 - retrieve 8
 - support for encrypted files, metadata, attributes 10
 - support for transforms 9

- file management process 32

files

- backup process 30
- base 32
- delta 32
- expiration 32
- expiration, and archives 32
- rebasing 33
- retrieval process 31

H

- hands-free installation 40

I

- installation
 - digitally signed 9
- installation options, Agents 39

K

- key features 8

L

- licensed product 7
- load balancing 18

M

- metadata 10
- mirrored configurations 18
 - definition of 17
 - diagram 19
- MyRoam 26
 - definition of 13

N

- network protocol used for backups 15

P

- processes
 - compaction 33
 - file backup 30
 - file retrieval 31
 - files
 - management 32

R

- RAID, definition of 17
- rebasing 33
- Registry database 15
 - in clustered Data Center 21
- replication
 - databases 16
- resynchronization 19

S

- server, account assignments in clustered Data Center 21
- servers
 - backup 15
 - Web 15

stand-alone configuration 17

 diagram 17

 limitations 17

Startup Wizard 9

subscription service 8

support

 cross platform 8

 transforms 9

Support Center 26

 definition of 13

T

TCP/IP 15

tool kit 28

transforms 9

transmission of files to archives 15

types of Agents 14

U

user installation 39

W

Web server 15

Wizard, Startup 9

Z

ZLIB compression library 14

Send documentation feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Product Overview (HPE Connected Backup 8.8.7)

Add your feedback to the email and click **Send**.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

We appreciate your feedback!