Evault Software

IBM i Agent 7.2

User Guide

EVault®
from Carbonite

Revision: This manual has been updated for version 7.2.
Software version: 7.20 (March 2017)

Carbonite, Inc.
Two Avenue de Lafayette
Boston, MA 02111
www.evault.com

Acknowledgements: Two encryption methods, DES and TripleDES, include cryptographic software written by Eric Young. The Windows versions of these algorithms also include software written by Tim Hudson. Bruce Schneier designed Blowfish encryption.

The Evault Software Agent, Evault Software CentralControl, and Evault Software Director applications have the encryption option of AES (Advanced Encryption Standard). Advanced Encryption Standard algorithm (named Rijndael, pronounced “Rain Doll”) was developed by cryptographers Dr. Joan Daemen and Dr. Vincent Rijmen. This algorithm was chosen by the National Institute of Standards and Technology (NIST) of the U.S. Department of Commerce to be the new Federal Information Processing Standard (FIPS).

The Evault Software Agents and Evault Software Director applications also have the added security feature of an over the wire encryption method.

Contents

1 Introduction

The IBM i (formerly iSeries) Agent backs up data from IBM i (iSeries) systems across a local network or the Internet to a secure vault.

The Agent must be installed and running on each IBM i system that you want to back up, and each system must be connected to the network. The application runs a background program which, when configured, schedules and runs backups automatically.

You can configure, schedule and monitor the Agent using a 5250 terminal emulator Command Line Interface (CLI).

Note: You cannot manage the IBM i Agent using Portal or a CentralControl application.

2 Install the IBM i Agent

This section describes how to install or upgrade the IBM i Agent. The installation requires that you have the Agent for IBM i installation kit and a running IBM i system.

2.1 System Requirements

Before you can install or upgrade the IBM i Agent, the system must meet the following prerequisites:

- QSECOFR or equivalent user profile.
- If you are installing the Agent from a CD, a Vary-On optical device connected to CD-ROM.
- If you are installing the Agent from a save file, an ACTIVE Ethernet device connected to local network or internet.
- At least 1GB available system storage on *SYSBAS.
- The following licensed programs must be installed and licensed:

Program	Option	Description
5770SS1		Library QGPL
5770SS1		Library QUSRSYS
5770SS1	*BASE	IBM i
5770SS1	1	Extended Base Support
5770SS1	3	Extended Base Directory Support
5770TC1	*BASE	IBM TCP/IP Connectivity Utilities for i

2.2 Start the installation or upgrade from a CD

To start the installation or upgrade from a CD:

1. Download the IBM i Agent ISO file from your service provider's website to your computer.
2. Burn the ISO file to a CD. For instructions on creating an installation CD, refer to your CD writing software documentation.
3. Insert the installation CD into the IBM i server.
4. Open a 5250 terminal session. Log on to the IBM i system with the QSECOFR user profile or the equivalent.
5. Run the following command to open the IBM i Agent Setup Main Screen:

   ```
   LODRUN DEV(opticalDeviceName)
   ```

6. Install or upgrade the IBM i Agent from the CD. See Install the Agent or Upgrade the Agent.

2.3 Start the installation or upgrade from a save file

To start the installation or upgrade from a save file:

1. Download the IBM i Agent zip file from your service provider's website to your computer.
2. Extract the zip file to a temporary location on your computer. The extracted file is a save file.
3. Open a 5250 terminal session. Log on to the IBM i system with the QSECOFR user profile or an equivalent.
4. Run the `CRTLIB EVSAVE` command to create a temporary library.
5. Run the `CRTSAVF EVSAVE/IBMIAGENT` command to create a temporary save file:
6. Do the following to upload the save file to the IBM i system using FTP:
   a. On a Windows machine, open a command prompt.
   b. Run the `FTP` *<IBM i system IP address>* command to start an FTP session.
   c. Enter your user profile (QSECOFR or equivalent).
   d. Enter your password.
   e. Enter `BIN`
   f. Enter `LCD C:\`*<IBM i Agent save file local location>*
   g. Enter `PUT` *<local file name>* `EVSAVE/IBMIAGENT`
   h. Enter `quit`
7. On a 5250 terminal, do the following to restore objects from the save file:
   a. Run the following command to retrieve the saved library name:

      ```
      DSPSAVF FILE(EVSAVE/IBMIAGENT)
      ```

   b. Run the following command to restore all objects from the save file to the EVSAVE library:

      ```
      RSTOBJ OBJ(*ALL) SAVLIB(<saved library>) DEV(*SAVF)
      SAVF(EVSAVE/IBMIAGENT) MBROPT(*ALL) ALWOBJDIF(*ALL)
      RSTLIB(EVSAVE)
      ```

      Where *<saved library>* is the library name found in Step 7a.

   c. Run the `DSPJOBLOG` command to verify that objects were restored successfully:

      Note: If the objects were restored successfully, an "*XX* objects restore. 0 not restored to EVSAVE" message appears below the RSTOBJ command.

   d. Run the `CALL EVSAVE/ISSETUP` command to open the IBM i Agent Setup Main Screen.
8. Install or upgrade the IBM i Agent from the save file. See Install the Agent or Upgrade the Agent.

2.4 Install the Agent

Before installing the IBM i Agent, make sure that the following prerequisites are met on the system:

- The Agent user profile does not exist.
- The AGENT subsystem (*SBSD) does not exist.
- The QCTLSBSD system value is set to either QSYS/QBASE or QSYS/QCTL. If the control subsystem is not QBASE or QCTL, use the CHGSYSVAL command to change the QCTLSBSD system value to either QSYS/QBASE or QSYS/QCTL before installing the Agent. You can revert it back to the original value after the Agent is installed.

To install the Agent:

1. On the IBM i Agent Setup Main Screen, press **F6** to perform a fresh installation.

   Note: If a previous version of the IBM i Agent is installed on the system, you cannot perform a fresh installation. See Upgrade the Agent.

2. Complete the fields on the screen. The default product library is BUAGENT. The default product directory is /buagent.

   If the installation is successful, a confirmation message appears at the bottom of screen.

2.5 Upgrade the Agent

If you recently upgraded your operating system to IBM i V7R3 and a previous version of the IBM i Agent is installed, you must upgrade the Agent to version 7.20 before performing any backups or restores.

Before upgrading the Agent, be sure that the following prerequisites are met:

- The AGENT user profile is not actively logged on the system.
- No job is running in the AGENT subsystem.

To upgrade the Agent:

1. Log on to the IBM i system with the QSECOFR user profile or equivalent. You cannot log on as AGENT to upgrade the Agent.
2. Run the `CALL EVSAVE/ISSETUP` command to open IBM i Agent Setup Main Screen.
3. Enter option **6** to upgrade the agent.
4. Enter **YES** in the confirmation screen.

   If the upgrade is successful, the Agent version is refreshed with the upgraded version. A confirmation message appears at the bottom of screen.

2.6 Uninstall the Agent

Before uninstalling the Agent, make sure that the following prerequisites are met:

- The AGENT user profile is not actively logged on the system.
- No job is running in AGENT subsystem.

To uninstall the agent

1. Log on to the IBM i system with the QSECOFR user profile or equivalent. You cannot logon as AGENT to uninstall the Agent.
2. Run the `CALL EVSAVE/ISSETUP` command to open the IBM i Agent Setup Main Screen.
3. Enter option **4** to uninstall the Agent.
4. Enter **YES** in the confirmation screen.

   If the Agent is successfully uninstalled, a confirmation message appears at the bottom of screen.

2.7 Installed Agent objects and files

The following IBM i Agent objects and files are installed on the system:

- Objects in the BUAGENT library (where BUAGENT is the product library name). This can be verified by running the `WRKOBJ BUAGENT/*ALL` command.
- Files in the '/buagent' IFS system (where /buagent is the product directory name). This can be verified by running the `WRKLNK '/buagent'` command.
- AGENT user profile. This can be verified by running the `WRKUSRPRF AGENT` command.
- AGENT subsystem. This can be verified by running the `WRKSBSD BUAGENT/AGENT` command.

  After a successful installation, the AGENT subsystem is activated automatically. This can be verified by running the `WRKACTJOB` command.

2.8 License the IBM i Agent

The IBM i Agent has a 30-day trial Agent license when it is first installed. The expiry date for the license is displayed in the Expiry Date field of the Work with License screen.

To obtain a code for a full license, contact your service provider and provide the serial number and software group of the system where the IBM i Agent is installed. This information is displayed on the Work with License screen. You can then enter the license code. See Add an Agent license.

3 Get started with the IBM i Agent

3.1 Authority of user profile for running backups and restores

To run backups and restores, the user profile must have the following special authorities:

- *ALLOBJ special authority to open and read libraries and objects.
- *SAVSYS special authority for saving system state.

The Agent installation automatically creates an "AGENT" user profile with the minimum user class and required special authorities. We recommend using this user profile for agent and job management, backups and restores.

The default password for the AGENT user profile is AGENT. For security reasons, we recommend changing this password. You can change the password using the `CHGUSRPRF` command. We do not recommend changing any other user profile parameters.

3.2 Access the IBM i Agent main menu

When you sign on to the IBM i server with the AGENT user profile, you will be automatically directed to the ISAGENT menu.

If the initial menu parameter for your user profile is not set to ISAGENT, you can access the ISAGENT menu by running the `GO BUAGENT/ISAGENT` command (where BUAGENT is the name of the agent product library).

The following menu options are available from the IBM i Agent main menu:

Menu option	Description
1. Agent configuration	Configure vault settings, retention schemes, email notifications, licenses and other Agent settings.
2. Work with Jobs	Create, change and delete jobs, and view safeset and log information. You can also run backups, restores and synchronizations from this menu.
3. Work with schedules	Create, change, delete, enable and disable schedules for backup jobs.
4. Backup	Run an immediate (ad-hoc) backup.
5. Restore	Restore data from a backup.
9. Synchronize	Synchronize backup job information between Agent and the vault.
10. Auto Job Creation	A tool that helps create jobs for backing up the entire system.
90. Signoff	Sign off the current user profile.

4 Configure the Agent

After installing an IBM i Agent, you can configure settings for the Agent.

From the Agent Configuration menu, you can access screens for configuring vaults, retention schemes, email notifications, licenses and other Agent settings.

4.1 Create and manage vault configurations

Before an Agent can back up data to a vault, the Agent must have at least one vault configuration. A vault configuration includes a vault address, account and account credentials. To obtain this information, contact your service provider.

Note:

- An Agent can only be registered once to a vault using the same account.
- An Agent can have a maximum of 10 vault configurations.

4.1.1 Create vault configurations

To create a vault configuration:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **1 (Work with Vaults)**.
3. On the Work with Vault screen, press **F6**.
4. On the Add Vault Configuration screen, complete the following fields:

Field	Description
Vault Name	The name of a vault where the Agent can back up data. An Agent can have a maximum of 10 vault configurations. Each vault name for an Agent must be unique. Vault names can be a maximum of 15 characters and can include letters, numbers, and the following special characters: # @ _ and $. A vault name cannot begin with a number.
Account Name	The vault account name provided by your service provider. An account name can have a maximum of 40 characters.
User Name	The user name provided by your service provider.
Password	The password provided by your service provider.
Verify Password	The password provided by your service provider.
Network Address	The network address (IP address or DNS name) for the vault, provided by your service provider.
Port Number	The port that the IBM i Agent uses to communicate with the vault. The port is configured by your service provider. The default value is 2546.

Field	Description
Reconnect delay	The time in seconds that the Agent waits before attempting to reconnect to the vault after a communication or session failure. The default value is 30 seconds. The maximum value is 1800 seconds. If a reconnection is successful, the backup will continue without data loss. If a reconnection attempt is unsuccessful, the Agent waits this amount of time before the next attempt, until the overall timeout (specified by the Retry Timeout value) is reached.
Retry Timeout	The overall amount of time, in seconds, that the Agent tries to reconnect to the vault. The default value is 3600 seconds. When the retry timeout value is reached, the backup fails, an error message is added to the log file, and an email notification is sent (if email notifications are configured). *Note:* The Retry Timeout value must be greater than the Reconnect delay value.
Over the Wire Encryption	Select YES to encrypt your backup data as it goes from the IBM i system to the vault.

4.1.2 Change vault configurations

To change a vault configuration:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **1 (Work with Vaults)**.
3. On the Work with Vault screen, enter option **2** for the vault configuration that you want to change.
4. On the Change Vault Configuration screen, complete the following fields:

Field	Description
Vault Name	The name of a vault where the Agent can backup data. An Agent can have a maximum of 10 vault configurations. Each vault name for an Agent must be unique. Vault names can be a maximum of 15 characters and can include letters, numbers, and the following special characters: # @ _ and $. A vault name cannot begin with a number.
Account Name	The vault account name provided by your service provider. An account name can have a maximum of 40 characters.
User Name	The user name provided by your service provider.
Password	The password provided by your service provider.
Verify Password	The password provided by your service provider.
Network Address	The network address (IP address or DNS name) for the vault, provided by your service provider.
Port Number	The port that the IBM i Agent uses to communicate with the vault. The port is configured by your service provider. The default is 2546.

Field	Description
Reconnect delay	The time in seconds that the Agent waits before attempting to reconnect to the vault after a communication or session failure. The default value is 30 seconds. The maximum value is 1800 seconds. If a reconnection is successful, the backup will continue without data loss. If a reconnection attempt is unsuccessful, the Agent waits this amount of time before the next attempt, until the overall timeout (specified by the Retry Timeout value) is reached.
Retry Timeout	The overall time, in seconds, the Agent tries to reconnect to the vault. The default is 3600 seconds. When the retry timeout is reached, the backup fails, an error message is added to the log file, and an email notification is sent (if email notifications are configured). Note: The Retry Timeout value must be greater than the Reconnect delay value.
Over the Wire Encryption	Select YES to encrypt your backup data as it goes from the IBM i system to the vault.

4.1.3 Verify vault configurations

To verify a vault configuration:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration)**.
2. On the Agent Configuration menu, enter **1 (Work with Vaults)**.
3. On the Work with Vault screen, enter option **8** for the vault configuration that you want to verify.

   A message at the bottom of the screen indicates if the verification was successful or failed.

4.1.4 Delete vault configurations

You cannot delete a vault configuration if an existing job is configured with the vault.

When you delete a vault configuration on the Agent side, it does not delete the Agent registration on the vault. If a vault configuration has been deleted from an Agent, you can recover it by reregistering the Agent with the vault. See Reregister an Agent with a vault.

To delete a vault configuration:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **1 (Work with Vaults)**.
3. On the Work with Vault screen, enter option **4** for the vault configuration that you want to delete.
4. On the confirmation screen, press **ENTER**.

4.2 Reregister an Agent with a vault

*Y*ou can reregister an Agent as a computer that was backed up to a vault so you can restore data from the computer that was backed up.

Note: You cannot reregister an Agent across operating systems. Before reregistering an Agent, back up the '/buagent' directory.

To reregister an Agent with a vault:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **1 (Work with Vaults)**.
3. On the Work with Vault screen, press **F8**.
4. On the Re-register with Vault screen, enter the vault network address, account and user name.
5. After the Agent connects to the vault, you will be asked to select a registered computer to recover. Once you confirm, any local existing global and job configurations will be overwritten.

4.3 Create and manage retention schemes

A retention scheme specifies how many copies of a backup are stored on the vault, how many days a backup is kept online, and how many days it is held in the archive.

An Agent can have a maximum of 10 retention schemes.

There is always a minimum number of online copies and online days available, even if one is less than the other. For example, if you specify 7 online copies of a backup for 7 online days, there will always be 7 copies even if they are more than 7 days old. In addition, there will always be 7 online days, even if more days have passed.

The oldest safeset is deleted first. You cannot delete all safesets. The most recent safeset is always kept.

4.3.1 Create retention schemes

To create a retention scheme:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **2 (Work with Retention Schemes)**.
3. On the Work with Retention Scheme screen, press **F6**.

4. On the Add Retention Scheme screen, complete the following fields:

Field	Description
Retention Name	The name of a retention scheme. An Agent can have a maximum of 10 retention schemes. Each retention scheme name for an Agent must be unique. Retention scheme names can be a maximum of 32 characters and can include letters, numbers, and the following special characters: # @ _ and $. A retention scheme name cannot begin with a number.
Online days (1-9999)	The number of days a safeset is stored on the vault before it expires. When the expiry date is reached, the safeset is automatically deleted. There will always be at least the number of online copies (below), regardless of the setting for online days.
Online copies (1-999)	The minimum number of safesets to maintain online. When the minimum number of safesets is exceeded, the oldest safeset is deleted.
Archive Backup (*YES/*NO)	Enter *YES to archive your backup for a specific number of days. You can then enter a value from 365 to 9999.

4.3.2 Change retention schemes

To change a retention scheme:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **2 (Work with Retention Schemes)**.
3. On the Work with Retention Scheme screen, enter option **2** for the retention scheme that you want to change.
4. On the Change Retention Scheme screen, change values in some or all of the fields:

Field	Description
Online copies (1-999)	The minimum number of safesets to maintain online. When the minimum number of safesets is exceeded, the oldest safeset is deleted.
Online days (1-9999)	The number of days a safeset is stored on the vault before it expires. When the expiry date is reached, the safeset is automatically deleted. There will always be at least the number of online copies (below), regardless of the setting for online days.
Archive Backup (*YES/*NO)	Enter *YES to archive your backup for a specific number of days. You can then enter a value from 365 to 9999.

4.3.3 Delete retention scheme

To delete a retention scheme:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **2 (Work with Retention Schemes)**.
3. On the Work with Retention Scheme screen, enter option **4** for the retention scheme that you want to delete.
4. On the confirmation screen, press **ENTER**.

4.4 Configure email notifications

You can configure email notifications for an Agent, so that one or more recipients receive an email when a job fails or succeeds. You must configure email notification for all of an Agent's jobs. An email notification cannot be created for a specific job.

To configure email notifications:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **4 (Optional settings).**
3. On the Optional Settings screen, complete the following fields:

Field	Description
Email Notification	***NONE** – Do not send email notifications. ***ALL** – Send email notifications when a backup or restore succeeds or fails. ***SUCCESS** – Send email notifications when a backup or restore finishes successfully. ***FAIL** – Send email notifications when a backup or restore fails.
From Email Address	Email address from which email notifications will be sent.
Recipients	Email notification recipient email addresses.
SMTP Server	Network address of the outgoing mail server (SMTP) that will send the email.
User Name	SMTP username.
Password	SMTP password
Domain	SMTP domain.

4.5 Add an Agent license

The Agent is installed with a 30-day trial license. With this trial license, the agent has full functionality for30 days. Enter a code for a full license before the end of the 30-day trial period.

To obtain a license code, contact your service provider.

To add an Agent license:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **5 (Work with Agent Licenses)**.

   The Work with License screen shows the Agent version, license expiry date and server information. You service provider might need this information to provide you with a full license code.

3. Enter your license code in the **License Key** field.

   The format of the license key is thirty (including hyphens) upper-case characters in the format XXX-XXXXXXXX-XXXXXXXX-XXXXXXXX.

4.6 Configure backup/restore priority and bandwidth throttling

You can choose to use all the available network bandwidth for backups and restores, or you can restrict the amount of bandwidth to a specific value.

To configure backup/restore priority and bandwidth throttling:

1. On the IBM i Agent Main Menu, enter **1 (Agent configuration).**
2. On the Agent Configuration menu, enter **6 (Advanced configuration)**.
3. On the Advanced Configuration screen, complete the following fields

Field	Description
Backup/Restore Priority	The backup or restore priority. Available values are 1 to 9. The lower the number, the higher the priority.
Use all available Bandwidth	Enter ***YES** to use all available network bandwidth when creating a backup. This option is disabling bandwidth throttling. Enter ***NO** to customize the bandwidth used during a backup.
Limit Bandwidth usage to	The amount of bandwidth in kilobytes per second (kb/s) allocated to the backup.
All Day	Enter ***YES** to apply the bandwidth settings to an entire day. Enter ***NO** to specify the bandwidth settings to a specific hour of the day.
Start Hours	The hour when bandwidth throttling starts.
Start Minutes	The minutes when bandwidth throttling starts.
End Hours	The hour when bandwidth throttling ends.
End Minutes	The minutes when bandwidth throttling ends.
On the following Days	The days of week when bandwidth throttling is applied.

5 Create and manage backup jobs

Before you can back up data, you must create a job. There is no limit to the number of jobs you can create.

You can select one of these types of data to back up:

- *OBJ – Backs up libraries and objects in a native system.
- *IFS – Backs up folders and stream files in an Integrated File System.
- *SYS – Backs up system state data. System state data is critical to the recovery of the operating system. See About System State (*SYS) backup for more details.

A set of backup jobs can protect your entire system:

- One *SYS job
- Several *OBJ jobs to backup the QSYS system and user data. We recommend at least one *ALLUSR job and one *IBM job.

  To do a full backup if you configure iASP devices, you must set up two or more *ALLUSR jobs: one for the *SYSBAS ASP device and one for each iASP device.

- One or several *IFS jobs to backup the entire *IFS system.

At least one retention scheme must exist before you can create a backup job. See Create and manage retention schemes. When you create a backup job, the default retention scheme (DAILY) is applied to the job. To use a non-default retention scheme, you must create a job with the default retention scheme, and then change the job to assign a deferent retention scheme.

5.1 Create backup jobs

To create a backup job:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, press **F6**.

3. On the Add New Backup Job screen, complete the following fields and press **ENTER**:

Field	Description
Job name	The job name. Job names can be a maximum of 30 characters. You can use all alphabet letters and the numbers 0 to 9. Special characters #, @, _, and $ are allowed. Job names exceeding eight characters may appear truncated due to operating system limits. The operating system allows system job names to have up to ten characters, so longer Agent Job names will be truncated. The Agent uses a two-character suffix, namely _B, @B, #B, _R, #R, @R to indicate the type of process so that leaves eight characters for the system job name. If there is already another Agent job with the same eight-character abbreviation, then the eighth position will be substituted with a number and the number will be incrementally increased until a unique name is found. The process will continue for the 7th, 6th, and so on until a unique name is found.
Vault name	The name of the vault where the data is backed up. Use the F4 function key to get the list of available vaults.
Data type	Enter the type of data to backup in the job: • ***OBJ** – Backs up libraries and objects in a native system. • ***IFS** – Backs up folders and stream files in an Integrated File System. • ***SYS** – Backs up system state data. System state data is critical to the recovery of the operating system. See About System State (*SYS) backup for more information.

4. On the Add New Job screen, complete options for the job data type. See Object (*OBJ) backup job options or IFS (*IFS) backup job options.

5. Enter encryption information in the following fields:

Field	Description
Encryption	The type of encryption standard to use to protect data. Beginning with Agent version 7.20, AES256 is the only option for creating new backup jobs. If an existing job has no encryption or uses another encryption type (e.g., AES 128 bit, Blowfish, DES, Triple DES), which are now set to *ORIG, you can continue to encrypt the job using that type. However, if you change the encryption type for an existing job, you cannot change the encryption type back to the original type. Only AES 256 is available. *Note:* If you change encryption options for an existing job, it will force a reseed. The next backup will take longer than previous delta backups, and the amount of data stored on the vault will increase temporarily, depending on your retention settings.
Password	The encryption password. Maximum 31 characters. *Note:* You must remember your encryption password. Otherwise, you will not be able to recover your data.

6. To specify compression, deferral or other advanced options, press **F10** and then enter values in some or all of the following fields:

Field	Description
Compression	The compression level specifies the amount of compression, if any, for backup data. Compression levels optimize the data size against the speed of processing. In some cases, it might be better to use additional time and processing to compress the data before sending it. Available levels are: • ***NONE** – Do not compress data. • ***DEFAULT** – Balance CPU consumption against the size of backup data. • ***STANDARD** – Balances processing against backup size. • ***NORMAL** – Balances processing against backup size. • ***MINIMUM** – Minimizes processing, possibly at the expense of a larger backup size. • ***BETTER** – Minimizes backup size, possibly at the expense of extra processing. • ***MAXIMUM** – Always minimizes backup size, regardless of the amount of processing required.
Deferring	Define the behavior when a backup job reaches a specified amount of time. Enter ***NO** to allow the backup job to run without a time limit. Enter ***YES** and a number of minutes to stop backing up any new data after the specified amount of time, even if some data is not backed up. Changes to data that was previously backed up will be backed up, regardless of the backup time window. ***Note:*** When the backup window is expired, no more data is saved. However, the Agent will take some time to traverse all the rest of files to gather metadata for the next backup.
Quick File Scanning	Any data streams that have not changed since the last backup are skipped. This option reduces the amount of data read during the backup process. When ***NO** is specified, files are read in their entirety, which may result in longer backup times. A read is forced for the following object types, regardless of the Quick File Scanning setting: • FILE contains multiple members and size greater than OBJSIZ • FILE is under journaling • object size greater than OBJSIZ • OUTQ • object name contains special characters (",/,', etc)

Field	Description
Log Detail	The amount and type of backup information to include in the log files. Enter one of the following options: • ***FILE** – Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up. • ***NONE** – Do not generate log file. • ***SUMMARY** – Provides high-level information, including the vault and Agent version, and backup sizes. • ***DIRECTORY** – Provides less detail than the Files logging level. Provides information about folders that are backed up.
Threading Model	On a multi-CPU system, you can use one or more threads to improve backup and restore performance. Enter one of the following options: • ***DEFAULT** – Agent selected threading model is used. • ***SINGLE** – A single threading model is used. • ***COMBINED** – A combined threading model is used. • ***BLOCKPROCESSOR** – The block processor threading model is used with up to four processing threads. • ***MAXBLOCKPROCESSOR** – The block processor threading model is used with up to five processing threads.
Disable CRC	Disable CRC can improve the performance of backups and restores.

7. Press **ENTER**.

   If successful, the new job appears on the Work with Job screen. If a problem occurred, an error message appears at the bottom of the Work with Job screen.

5.1.1 Object (*OBJ) backup job options

Field	Description
ASP Device	The name of the auxiliary storage pool (ASP) device where the data is stored. These are the available options: • Name – The name of the ASP device. • *SYSBAS – The system ASP.

Field	Description
Include Objects	Objects to include in the backup. The list can include a maximum of 128 entries. Object – The name of the object to back up. • ***ALLUSR** – Saves all objects in the user libraries. QSYS becomes the only legal library name. If you specify more than one entries including *ALLUSR, all the other entries are removed. Refer to **LIB** option in SAVLIB command for more details. • ***IBM** – Saves all objects in system (IBM) libraries. QSYS becomes the only legal library name. If you specify more than one entries including *IBM, all the other entries are removed. Refer to **LIB** option in SAVLIB command for more details. • ***ALL** – Saves all objects in a specific library. • generic* – Generic names start with one or more valid characters, followed by the wildcard. Library – The library name of the object. • If *ALLUSR or *IBM is specified, the library name must be QSYS. • When a library name is specified in the Object field, and *LIB is specified in the Type field, the library name must be QSYS. A library cannot be a generic* name. Type – The type of objects to back up. Use *ALL for all types. *Note:* When user creates an *IBM job, 3 more include entries are automatically added to the job config: • All *CMD objects in QSYS library • All *MENU objects in QSYS library • All *PNLGRP objects in QSYS library
Exclude Objects	Objects to exclude from the backup. The list can include a maximum of 128 entries. Object – The name of the object to exclude from the backup. • *ALL • generic* – Generic names start with one or more valid characters, followed by the wildcard. • Library – The library name of exclude object. A library cannot be a generic* name. • Type – The type of objects to exclude. Use *ALL to exclude all types. Library – The library name of the object. Type – The type of objects to back up. Use *ALL for all types.
Recursive	When backing up a library, specifies whether objects inside the library are included.

Field	Description
Small object size (KB)	Specify the maximum size of object for using the **Save File Method**. The Save File Method save small object to a save file and backup the save file to vault. This method significantly increases restore performance for a job that contains many small objects, without reducing backup performance.
Save savefile data	Specify whether to save the entire save file or just save the save file header. • ***YES** – Save the entire save file. • ***NO** – Save the save file object but ignore its data. Refer to the SAVFDTA option in the SAVOBJ command for more details.
Spool data (V5R4 and up)	Specify whether to save spooled files inside output queue (*OUTQ). • ***ALL** – Save all spooled files inside output queue. • ***NONE** – Do not save spooled files inside output queue. Refer to the SPLFDTA option in the SAVOBJ command for more details.
Save active	This option is used to define check point processing while an object is in used. • ***NONE** – Object in use is not saved. • ***SYSDFN** – Object can be saved while they are in use. • ***TRIGGER** – (obsolete) Refer to the SAVACT option in the SAVOBJ command for more details.
Save active wait time	Specify the amount of time to wait for check point when object is in use. Refer to the SAVACTWAIT option in the SAVOBJ command for more details.
Update History	Specifies whether the save history information of each saved object is changed with the date, time, and location of this save operation. Refer to the UPDHST option in the SAVOBJ command for more details.

5.1.2 IFS (*IFS) backup job options

Field	Description
Include IFS stream files	The absolute path and file name pattern to include in the IFS backup. The list can include a maximum of 128 items.
Exclude IFS stream files	The absolute path and file name pattern to exclude from the IFS backup. The list can include a maximum of 128 items.
Recursive	When backing up a directory, specifies whether files and sub-directories inside the directory are included.
Suppress archive-bit process	Specifies whether the archive-bit of each saved file is altered.

5.1.3 About System State (*SYS) backups

A System State (*SYS) backup job includes the following items:

Item	Description
User Profiles and Authorization Lists	The same as **SAVSECDTA** command. It's temporarily saved in save file BUAGNET/SYSSTS1.
Configuration Objects	The same as **SAVCFG** command. It's temporarily saved in save file BUAGENT/SYSSTS2
System Values	All system values. Refer to **WRKSYSVAL *ALL** command. It's temporarily saved in user space BUAGENT/SYSSTS3.
TCP/IP configuration	Save the following objects from QUSRSYS library, temporarily to save file BUAGENT/SYSSTS4: QAOK*, QAOS*, QATOCIFC, QATOCLIFC, QASO*, QATM*, QATOC*, QAZSP*, QPTMPLPD, QTCP, QTCPASC, QTCPEBC, QTFTP, QTMFQ00021, QTMFTPS, QTMHHTTP, QTMSMTPS, QTMSTRACE, QTPPPOUTQ, QWEBADMIN, QZDATRC, QZHQTRC, QZMF*, QZRCTRC, QZSOTRC *Note:* These objects will be automatically excluded from *OBJ job.
Job schedule entries	Save all *JOBSCD objects from QUSRSYS library, temporarily to save file BUAGENT/SYSSTS5. *Note:* These objects will be automatically excluded from *OBJ job.

The following table shows differences between a *SYS job and the SAVSYS command:

Data type	*SYS job	SAVSYS command
Security Data (SAVSECDTA)	Y	Y
Configuration Objects (SAVCFG)	Y	Y
Licensed Internal Code (LIC)	(*)	Y
QSYS library	(*)	Y
System Values	Y	-
TCP configurations	Y	-
Job schedule entries	Y	-

*: this type of data is covered by an *IBM job. *SYS job does not contain this data.

5.2 Change backup jobs

To change a backup job:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.

2. On the Work with Job screen, enter option **2** for the backup job that you want to change.
3. On the Change Job screen, enter new values in the fields you want to change.

   For field descriptions, see Create a backup job, Object (*OBJ) backup job options and IFS (*IFS) backup job options.
4. To change compression, deferral or other advanced options, press **F10** and then enter values in some or all of the fields.

5.3 Delete backup jobs

You can delete a backup job when it is no longer needed. The delete operation removes the job configuration, metadata files and log files from your local Agent.

Deleting a job will not remove the job configuration, metadata files and existing backups from the vault. If a job is deleted from an agent, you can recover it by re-registering the Agent with the vault. See Re-register an Agent with a vault.

To delete a backup job:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, enter option **4** for the backup job that you want to delete.
3. On the Confirm for Delete Item screen, press **ENTER**.

5.4 Add a custom command to a backup job

You can attach a custom command to a backup job to run a customized operation before or after the backup runs.

Note: To avoid backup failure, a pre-backup command should be tested before it is attached to a backup job.

If a custom command is a batch job (using the SBMJOB command), the Agent considers the custom command to be successfully completed once the job is successfully submitted to the system. The user should check the custom command job log for the result.

To add a custom command to a backup job:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, enter option **3** for the backup job for attaching a custom command.
3. On the Work with Customer Command screen, enter one of the following values in the Command Class field:
   - ***PRE** – When the backup starts, the backup job will kick off the custom command and wait until it finishes. If the command fails, the backup will not start and an error message will appear in the backup log.

- ***POST** – When the backup is completed (all data has been committed to vault), the backup job will kick off the custom command and wait for it finish. If the command fails, the backup will still complete with a warning.

4. In the Command field, enter the custom command.

5.5 View a job's safesets

To view a job's safesets:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, enter option **5** for the job for which you want to view safesets.

   The Display Safeset screen shows the job's safesets that are stored in the vault. A sequential number is assigned to each safeset after every backup. Location, status and date and time the safeset is created are also displayed.

3. To view more information about a safeset, enter option **5** for the safeset.

   The Display Safeset Detail page shows the following information:

Field	Description
Job Name	The name of the job that created the safeset.
Catalog Number	The sequential number assigned to this safeset.
Location	The name and address of the vault where this safeset is kept.
Status	The safeset status. These options are available: • Online – You can use the safeset for a restore. • Work Area – The data is in transition to Online status. You can either wait for the system to change the status, or you can execute a Synchronize command • Archived – The safeset is stored off-line from the vault.
Backup Time	The date and time the safeset was created.
Backup Type	The type of backup.
Storage Size	The size of backup. These fields are available: • Original – The size of the original backed up data. • Deltized – The amount of deltized data (zero for the initial seed). • Compressed – The size of the data stored on the vault.
Retention	The retention plan for the safeset. These fields are available: • Days – The number of days this safeset is kept. • Copies – The number of backup copies kept. • Archived Days – The number of days an archive copy is kept.
Encrypted	Indicates if the backup data was encrypted.
Compressed	Indicates if the backup data was compressed.

Field	Description
Media Type	The type of media used to store the backup data.
Expiry	The date the safeset will be deleted.

5.6 View a job's log files

A log file is created for every backup, synchronize, or restore. You can view a log file to determine if the backup, synchronize or restore operation was successful, the date and time the event completed, and what objects were backed up or restored (when the *FILE log option is specified, see the **Log Detail** field).

To view a job's log files:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, enter option **6** for the job for which you want to view log files.

   The Display Object Links screen lists the job's logs. Backup logs are named with the safeset number (e.g., 00000001.LOG). Restore logs are named in RST*YYYYMMDD-HHMMSS*.LOG format. The SYNCH.LOG is the most current synchronization log file.
3. Enter option **5** for the log that you want to view.

5.7 View the Agent directory

The Agent directory is created during installation. The Agent uses this location to store Agent and job configuration files, as well as metadata and log files for jobs.

To view the Agent directory:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, press **F7** to open the Agent directory.

   Note: This is a shortcut of the **WRKLNK '/buagent'** command.

5.8 View job statuses

This screen shows all running agent jobs in the AGENT subsystem, as well as related jobs in the QSYSWRK subsystem.

To view job statuses:

1. On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**.
2. On the Work with Job screen, press **F10** to view the job statuses.

   Note: This is a shortcut of the **WRKACTJOB SBS(AGENT QSYSWRK)** command.

6 Schedule backups, synchronizations and custom commands

After creating a backup job, you can create one or more schedules for running the job automatically. You can also schedule synchronizations and pre-defined custom commands.

The Agent uses the built-in IBM i job scheduler to start backups, synchronizations or pre-defined custom commands. Once a job schedule entry is created, you can find the corresponding job entry in the WRKJOBSCDE command.

Important: Do not use the IBM Advanced Scheduler to schedule Agent jobs. The Agent job scheduler requires multi-threading. The IBM Advanced Scheduler uses the RCLRSC command, which is not compatible with multi-threaded applications.

To schedule a backup, synchronization or custom command:

1. On the IBM i Agent Main Menu, enter option **3 (Work with schedules)**.
2. On the Work with Schedule screen, press **F6**.
3. On the Work with Scheduler screen, complete the following fields, and press **ENTER**:

Field	Description
Command	Enter the type of operation to schedule: • ***BACKUP** – Schedule a backup. • ***SYNCH** – Schedule a job synchronization, where the Agent checks which safesets for the job are online and available for restore. • ***CUSTOM** – Schedule a custom command or user script to run.
Job Name	The name of a backup job. Press **F4** to retrieve a job list.

4. Enter values in any of the following fields that appear:

Field	Description
Retention Name	The name of a valid retention scheme. Press **F4** to view a list of retention schemes.
Command Cycle	The frequency that the schedule runs (see the Frequency field in the WRKJOBSCDE command). These options are available: • ***WEEKLY** • ***MONTHLY**
Minute/Hour	The time for the job to start.
Day of week	The day of the week for the job to run.
Date of month	The date of the month for the job to run.
Quick File Scanning	Any data streams that have not changed since the last backup are skipped.
Deferring	Define the behavior when a backup job reaches a specified amount of time.

6.1 Change schedules

Once a job schedule entry is changed, you can find the associated changes in the corresponding job entry in the WRKJOBSCDE command.

To change a schedule:

1. On the IBM i Agent Main Menu, enter option **3 (Work with schedules)**.
2. On the Work with Schedule screen, enter option **2** for the schedule that you want to change.
3. On the Work with Scheduler screen, enter values in any of the following fields that appear:

Field	Description
Retention Name	The name of a valid retention scheme. Press **F4** to retrieve a retention list.
Command Cycle	The frequency that the schedule runs (see the Frequency field in the WRKJOBSCDE command). These options are available: • ***WEEKLY** • ***MONTHLY**
Minute/Hour	The time for the job to start.
Day of week	The day of the week for the job to run.
Date of month	The date of the month for the job to run.
Quick File Scanning	Any data streams that have not changed since the last backup are skipped.
Deferring	Define the behavior when a backup job reaches a specified amount of time.

6.2 Delete schedules

To delete a schedule:

1. On the IBM i Agent Main Menu, enter option **3 (Work with schedules)**.
2. On the Work with Schedule screen, enter option **4** for the schedule that you want to delete.
3. On the Confirm for Delete Item screen, press **ENTER**.

   When a job schedule entry is deleted from the Work with Schedule screen, the corresponding job entry in the WRKJOBSCDE command will be removed.

6.3 Enable or disable schedules

To enable or disable a schedule:

1. On the IBM i Agent Main Menu, enter option **3 (Work with schedules)**.

2. On the Work with Schedule screen, do one of the following:
   - To disable a schedule, enter option **3** for the schedule.
   - To enable a schedule, enter option **6** for the schedule.

   When a schedule is enabled or disabled from Work with Schedule screen, the status of the corresponding job entry in the WRKJOBSCDE command will be changed to HLD (hold) or SCD (scheduled).

7 Run ad-hoc (unscheduled) backups

After a backup job is created, you can run the backup at any time, even if the job is scheduled to run at specific times.

To run an ad-hoc (unscheduled) backup:

1. Do one of the following:
   - On the IBM i Agent Main Menu, enter option **4 (Backup)**. On the Run Backup screen, enter the name of the job that you want to run. Press **F4** to retrieve a job list.
   - On the IBM i Agent Main Menu, enter option **2 (Work with Jobs)**. On the Work with Job screen, enter option **7** for the job that you want to run.
2. On the Run Backup screen, complete the following fields:

Field	Description
Retention Scheme	The name of a valid retention scheme. Press **F4** to retrieve retention list.
Quick File Scanning	Any data streams that have not changed since the last backup are skipped. See Common backup job options for more detail.
Disable Deferring	Define the behavior when a backup job reaches a specified amount of time. See Common backup job options for more detail.
Defer after	A maximum amount of time (in minutes) that the backup job can run. See Common backup job options for more detail.

7.1 Seeding and Re-seeding

When you run a backup job the first time, a "seed" backup is created on the vault. This seed backup contains all the data selected for backup in the job. Subsequent backups are much smaller and only include changes (deltas) that have occurred since the last backup.

If the encryption type or password has changed since the last backup, the next backup will automatically be a re-seed.

In the case of a re-seed, your backup will take longer to complete and a message about re-seeding is created in the log file.

7.2 Concurrent backup instances

A backup job cannot have two running instances at the same time. If you start a backup job when another backup job is running, the second job will not start until the first job finishes. The other instance will stay in SEMW (semaphore wait) status until the running backup finishes.

7.3 Save while active

An object backup (*OBJ) job uses the IBM i save-while-active function to save an object while it is in use by another job. This function allows you to back up your data without stopping active processes. For more details about the save-while-active function, see articles in the IBM Knowledge Center: http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_73/rzaiu/rzaiurzaiu300.htm

The Agent includes this function setting in an *OBJ job configuration. However, the Agent currently only supports the *SYSDFN (system-defined synchronization) option, which means that objects in the backup job might reach check points at different times even though they are in the same library. This occurs because the Agent saves each object individually. Using the *SYNLIB (full synchronization) or *LIB (library synchronization) options will take much longer for the backup job to complete check point processing.

To protect your critical data in consistent mode, we recommend you journal the objects, enable the save-while-active function in the backup job, and use another backup job to back up the journals (*JRN) and journal receivers (*JRNRCV).

If your critical data is not currently configured for journaling, refer to articles in the IBM Knowledge Center for instructions on setting up journaling: http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_73/rzaki/rzakikickoff.htm

Here is an example:

1. Create a single journal for each library to be backed up. This simplifies administration and adds the journal and receivers in the correct library.
2. Create the main job to back up all objects and data and specify two exclusions for *JRN and *JRNRCV object types. Specify *SYSDFN in the save active parameter.
3. Create a separate job for *JRN and *JRNRCV objects.

Running backups and restores in the proper sequence is very important. If the backup sequence is not correct and data is updated while the backup is running, some journal receivers might not contain all the data required for point-in-time recovery. If the restore sequence is not correct, some objects might not be journaled after the restore.

The proper backup sequence is:

1. Run the main backup job.
2. Run the *JRN and *JRNRCV job.

The proper restore sequence is:

1. Restore *JRN and *JRNRCV.
2. Restore other objects.
3. Apply the journal back to the point in time.

8 Synchronize jobs

When you synchronize a job, the Agent checks which safesets for the job are online and available for restore, and performs the following tasks on the Agent:

- Updates safeset status information.
- Recreate the delta file (DTA), if the most recent delta file is missing.
- (Optional) Resynchs catalog files that are not available on the Agent system.

To synchronize a job:

1. Do one of the following:
   - On the IBM i Agent Main Menu, enter option **9 (Synchronize)**. On the Synchronize screen, enter the name of the job you want to synchronize. Press **F4** for a job list.
   - On the IBM i Agent Main Menu, enter **2 (Work with Jobs)**. On the Work with Job screen, enter option **8** for the job you want to synchronize.
2. To resynchronize a missing catalog file for the job, press **F10** on the Synchronize screen. Enter ***YES** to resynch the missing catalog file.

9 Auto Job Creation

The Auto Job Creation tool helps users create jobs to protect the entire system while best using system resources and optimizing backup times.

You can only use automatic job creation once. To use this tool, do the following:

1. Install the Agent on the system.
2. Register the Agent to the vault.
3. Run the auto job creation.

9.1 Start the auto job creation

The first step of auto job creation is to scan all libraries and objects on the system, gather attributes and information, and create the database.

Note: Once the database is created, you cannot create it again.

To start the auto job creation:

1. On the IBM i Agent Main Menu, enter option **10 (Auto Job Creation)**.
2. On the AUTOJOB Menu, enter **1 (Create Database)**.
3. On the DSPJOBD Import/Generate screen, complete the following fields:

Field	Description
Import or Generate Data	Generate database for Auto Job Creation, or import it from existing database. • ***IMPORT** – specify an existing database. • ***GENERATE** – generate a new database.
Library	The library name of database file. This option is for *IMPORT only.
File	The database file name. This option is for *IMPORT only.

9.2 Library/Job Selection

Note: Once library selection is done, you cannot run it again.

For library/job selection:

1. On the IBM i Agent Main Menu, enter option **10 (Auto Job Creation)**.
2. On the AUTOJOB Menu, enter option **2 (Library/Job Selection)**.
3. On the Automatic Job Creation/Exclude Libraries screen:
   - Enter the library names you do not want to backup.

     Note: You cannot use wild card characters in the library names.

- Press **F6** to view and edit fields on the Select Libraries screen. An **X** mark in the table shows the library name selected by the job number. You can remove the **X** mark to remove a library from the job, or move to a different job.
- Press **F6** to create a new include/exclude list. Press **F2** to save configurations.
- Press **F11** to preview the job name and include/exclude number of libraries. You can modify the job name here. Press **F2** to save your changes.
- Before exit from the Select Libraries screen, press **F2** to save your configuration.
- Before exiting from the Automatic Job Creation/Exclude Libraries screen, press **F2** to save your configuration.

9.3 Retention Schedule

For retention schedule:

1. On the IBM i Agent Main Menu, enter option **10 (Auto Job Creation)**.
2. On the AUTOJOB Menu, enter option **3 (Retention Schedule)**.
3. Complete the following fields, and then press **F2** to save your changes.

Field	Description
Retention	The retention name from the Agent config (Global.vvc) file.
Period	The frequency that the job runs. See **Command Cycle** in Schedule options. These options are available: • *WEEKLY • *MONTHLY
Days/Month (MTWTFSS/1-31):	The day of week or the date of month for jobs to run. You can enter a value from Monday to Sunday, the dates from 1 to 31, or *First, *Last. For example, you enter _T_T__S to run a job on Tuesday, Thursday and Sunday. You must include spaces for days that will not be scheduled.
Time	The time the job runs in hours and minutes. The hour is in 24 hours format.
Quick Scan	Any data streams that have not changed since the last backup are skipped. See Common backup job options for more detail. Two options are available: • *YES • *NO
Defer Disable	Define the behavior when a backup job reaches a specified amount of time. See Common backup job options for more detail. Two options are available: • *NO • *YES
Defer Time	When deferring is enabled, specify the time in minutes for the backup window. See Common backup job options for more detail.

9.4 Create jobs

Note: Once jobs are created, you cannot run it again.

To create jobs:

1. On the IBM i Agent Main Menu, enter option **10 (Auto Job Creation)**.
2. On the AUTOJOB Menu, enter option **4 (Create Jobs)**.
3. On the Create Jobs screen, do the following:
   - Enter the vault name for jobs backup to. Press **F4** twice to get a list of registered vaults.
   - Enter the encryption password.
   - Press **F2** to save your changes.
   - Press **F11** to preview the job name and include/exclude number of libraries. You can modify the job name here. Press **F2** to save your changes.
   - To create jobs, press **F6**.

9.5 Schedule jobs

Note: Once schedule entries are created, you cannot run it again.

To schedule jobs:

1. On the IBM i Agent Main Menu, enter option **10 (Auto Job Creation)**.
2. On the AUTOJOB Menu, enter option **5 (Schedule Jobs)**.
3. On the Schedule Jobs screen, do the following:
   - Use an **X** mark in the table to define the job and scheduled retention. Each **X** mark creates a schedule entry.
   - Press **F2** to save your changes.
   - Press **F6** to create schedule entries.

10 Restore data

There are several scenarios for restoring IBM i data:

- To recover one or more objects or IFS stream files. You can restore them to their original location, overwriting any that are there, or restore them to a different location as desired.
- To restore data that was backed up on one system to another system.

  Note: This option is restricted by the operating system version of the target release (TGTRLS). You can only restore from a system if the operating system version is in the list of TGTRLS options on the restore system. (Refer to command SAVOBJ for TGTRLS options).
- To recover a complete system (i.e., perform a disaster recovery) when the original system has been lost. See Disaster Recovery.

  Note: In a disaster recovery scenario, the operating system of the restore system must be the same as the backup system.

To restore data:

1. Do one of the following:
   - On the IBM i Agent Main Menu, enter option **5 (Restore)**. On the Run Restore screen, enter the name of the job from which you want to restore data. Press **F4** to retrieve a job list.
   - On the IBM i Agent Main Menu, enter option **2 (Work with Jobs**). On the Work with Job screen, enter option **9** for the job from which you want to restore data.
2. On the Run Restore screen, complete the following fields:

Field	Description
Safeset No	The safeset number which is going to restore data from. • *LASTONLINE – restore from the most recent backup. • 1-99999999 – specified a safeset number
Data Type	The data type of this backup job. These are available options: • *OBJ - to restore libraries and objects in the native system. • *IFS - to restore folders and stream files in the Integrated File System. • *SYS - to restore system state data.

Field	Description
Allow object differences	Specifies whether differences are allowed between the saved objects and the restored objects. These are the available options: • *NONE • *ALL • *AUTL • *COMPATIBLE • *FILELVL • *OWNER • *PGP See the ALWOBJDIF option in the RSTOBJ command for each option specification.
Log detail	The amount and type of backup information to include in the log files. Choose from these options: • *FILE – Provides the most detailed information, and is typically used for troubleshooting. Provides information about files that are backed up. • *NONE – Do not generate log file. • *SUMMARY – Provides high-level information, including the vault and Agent version, and backup sizes. • *DIRECTORY – Provides less detail than the Files logging level. Provides information about folders that are backed up.

3. Complete other fields required for the job data type. See Object (*OBJ) job restore options, IFS (*IFS) job restore options or System State (*SYS) job restore options.

10.1 Object (*OBJ) job restore options

Field	Description
Include objects	The objects to include in the restore. The list can include a maximum of 128 items. You can configure these options: • Name – The name of the object to restore. Specify a valid object name or generic name to restore one or more objects, or specify *ALL to restore all objects in the specified library. • Library – The name of the library to restore. • Type – The object type to restore. Specify *ALL to restore all types. *Note: Type *LIB will only restore the library object.* • File member – The database file members to restore. Specify a valid member name or generic name to restore one or more members in the specified file. Specify *NONE to restore only the file object without any member. Specify *ALL to restore all file members. *Note: To restore a specific member, *FILE Type must be specified.*

Field	Description
Exclude objects	The objects to exclude from the restore. The list can include a maximum of 128 items. You can configure these options: • Name – The name of the object to exclude from the restore. Specify a valid object name or generic name to exclude one or more objects, or specify *ALL to exclude all objects in a library. • Library – The name of the library to exclude from the restore. • Type – The object type to exclude from the restore. Enter *ALL to exclude all types. • Member – (obsoleted) *Note:* When type is specified to *FILE, you cannot use generic object name. However, *ALL is allowed.
Data base member option	The database member restore option for database file exists on the system. These are the available options: • *ALL • *MATCH • *NEW • *OLD For more information, see the help text for the MBROPT option of the native operating system RSTOBJ command.
Spool file data	Restores spooled file data and attributes. These are the available options: • *NEW – Select this option to restore spooled file data saved with the output queue if it does not currently exist. • *NONE – Do not restore spooled file in output queue (*OUTQ) For more information, see the help text for the SPLFDTA option of the native operating system RSTOBJ command.
Restore Library	The library to restore to. You select the original source library, or a different one. • Name – Restore to a named library. • *SOURCE – Restore to original library. *Note:* The Agent will create the library if it does not exist, using default attributes and authorities.
Restore ASP Device	Restores data to an auxiliary storage pool (ASP) device. These are the available options: • Name – The name of the ASP device to which you want to restore data. • *SAVASPDEV – Data is restored to the same ASP device from which it was saved.
Restore ASP Number	Objects are restored to the auxiliary storage pool (ASP) from which they were saved or to the system ASP (ASP number 1) or to a basic user ASP (ASP numbers 2 through 32).

Field	Description
File Overwrite/Rename Obj	Specify the restore option if an object already exists. • *OVRWRT – Overwrite all existing objects without asking. • *PMTOVRWRT – Prompt for overwrite option on each existing object. • *NOOVRWRT – Do not overwrite if object exists. *Note:* When you choose the *PMTOVRWRT option, the main restore job will be pause on MSGW status when an existing object is detected, you should run the WRKACTJOB SBS(AGENT) command. Use option **7** (Display Message) to input the overwrite option.
Ignore security data	Turn on or off the granting of private authorities on an object after it is restored. • *YES – Ignore the security data and stop the GRTOBJAUT command after object is restored. • *NO – Issue the GRTOBJAUT command to grant private authorities after object is restored. When this option is used, an authorization list or user authorities that were added for an object after the backup will not be changed or removed when the object is restored. It is recommended that you enter *YES to improve restore performance. You can later run the RSTAUT command to set private authorities.
No. of Jobs for small objects	The number of jobs to run concurrently for restoring the objects backed up by the **Save File Method**. You can enter values from 1-8. The default is 4.

10.2 IFS (*IFS) job restore options

Field	Description
Include IFS stream files	The IFS stream files to restore. The list can include a maximum of 128 entries. • Absolute path – The absolute path for the stream files saved on the vault. • File – A name pattern for files to restore. • Recursive – Includes sub-directories in the restore.
Exclude IFS stream files	The IFS stream files to exclude from the restore. The list can include a maximum of 128 entries. • Absolute path – The absolute path for the stream files saved on the vault. • File – A name pattern for files to exclude from the restore. • Recursive – Excludes sub-directories in the restore.
Destination	The directory to restore. It must be an absolute path. • Name – Restore to a named directory. • *SOURCE – Restore to original directory. *Note:* The Agent will create the directory if it does not exist, using default attributes and authorities.

Field	Description
File Overwrite/Rename Options	Specify the restore option if an object already exists. • *OVRWRT – Overwrite all existing files without asking. • *PMTOVRWRT – Prompt for overwrite option on each existing file. • *NOOVRWRT – Do not overwrite if the file exists. • *RNMINC – Renames the restoring (incoming) file with a unique number (0001, 0002) appended to the file name if that file exists. • *RNMEXT – Renames the existing file with a unique number (0001, 0002) appended with the file name, and then restore the file with original name. *Note:* When you choose *PMTOVRWRT option, the main restore job will be pause on MSGW status when an existing object is detected, you should run **WRKACTJOB SBS(AGENT)** command, use option **7** (Display Message) to input the overwrite option.
Create sub-file	These are the available options: • *YES – Recreates the restored file structure like the backup structure. • *NO – Files are added to the top level, and no sub-files are created.

10.3 System State (*SYS) job restore options

Field	Description
System states	Specify the type of system data to restore. These are the available options: • *ALL – All system state data. • *CFG – Configuration Objects (**RSTCFG**) • *USRPRF – User profiles (**RSTUSRPRF**) • *SYSVAL – All system values. • *QUSRSYS – TCP/IP configuration files in QUSRSYS library. • *JOBSCD – Job schedule objects in QUSRSYS library.
Prompt RSTxxx command	These are the available options: • *YES – Agent will prompt **RSTCFG** and **RSTUSRPRF** commands for user to specify different options. • *NO – No prompting, agent completes each restore command with default options.

10.4 About System State (*SYS) restores

System State restores must be run on the console. The console device is usually DSP01. This is because restore user profiles (RSTUSRPRF) require the system to run in a restricted state (**ENDSYS** or **ENDSBS**

***ALL** commands). There are more restrictions applied by the RSTUSRPRF and RSTCFG commands. Refer to information about these commands for more details.

There are two steps for system state restore:

Step 1: Agent restore safe files and user space object from vault to local agent product library.

Step 2: put the system to restricted state, issue restore commands to restore each part of data.

On the RESTORE command, you can choose which part of the system state data to restore, or you can specify *ALL to restore all of them.

Note: The temporary save files remain in the agent product library, except system values, you can manually restore other parts of system state from save files.

Restoring TCP data may involve a network interruption. This occurs because the TCP configuration is being overwritten by the restore. You may need to review the TCP configuration after restore.

If you restore the system value, the system IPL might need to apply system value changes. Refer to each system value for more information.

For more *SYS restore information, please refer to Disaster Recovery.

11 Disaster Recovery

In the event of a system disaster where the source (IPL) disk is lost, you will need to recover or rebuild the system from the ground up.

Important: For detailed disaster recovery information, please contact your service provider.

11.1 Prerequisites

Before completing a disaster recovery, the system must have:

- Valid backups.

  You need valid backups (such as ***SYS**, ***ALLUSR**, ***IBM**, ***IFS**) to protect the entire source system. Valid safesets must be stored on a vault.

- A basic operating system.

  The basic operating system includes **Licensed Internal Code** (LIC), **57xxSS1** (option *BASE, Library QGPL and Library QUSRSYS). You can either install from IBM supplied media, or restore from a SAVSYS tape.

- Network access to a vault.

  You need to have the licensed program 57xxTC1 installed on the system. TCP/IP interface needs to be configured and enabled, and have access to the vault.

- An IBM i Agent.

  An up-to-date agent must be installed on the IBM i system.

11.2 Complete a disaster recovery

This section provides a high level view of how to recover an IBM i system. For detailed disaster recovery information, please contact Support.

To complete a disaster recovery:

1. Install the Licensed Internal Code (LIC) and the operating system (OS) from IBM supplied media or SAVSYS tape.
2. Restore the QGPL and QUSRSYS libraries from IBM-supplied media, or from SAVSYS tape if they are saved in SAVSYS tape.

   Note: We recommend backing up the QGPL and QUSRSYS libraries to the same SAVSYS tape for disaster recovery.

3. Install the following license programs from IBM supplied media:
   - 57xxSS1 option 1
   - 57xxSS1 option 3
   - 57xxTC1 option *base

   Where xx is the licensed program code for your OS level

4. (Optional) Load recent PTFs if QUSRSYS is installed from IBM supplied media.
5. Configure TCP/IP connectivity to the LAN/WAN.
6. Install the IBM i agent. If you install the agent on the original system, the original agent license key can be re-used. If you install the Agent on a new system, please contact your Service Provider for a new license key.
7. Re-register the agent from the vault to recover original agent configuration and jobs. If your job is encrypted, you need to re-enter the correct encryption password.
8. If the system is rebuilt from IBM supplied media, you need to restore System State (*SYS) job first.
9. Restore *IBM job
10. Restore *ALLUSR job.
11. Restore user data from other jobs.
12. Restore *IFS jobs.
13. Run the RSTAUT command to apply private authorities.
14. Restart the system.

12 Best Practices

12.1 Keep your system protection up-to-date

Before using the IBM i Agent to protect your data, we recommend that you have a completed full system (SAVE 21) tape backup of the original system.

When using the IBM i Agent to protect your system, we recommend that you periodically have a system data (SAVSYS) tape backup. SAVSYS saves the Licensed Internal Code (LIC) and the operating system, as well as security data (SAVSECDTA) and configuration objects (SAVCFG). This could help you quickly rebuild your operating system.

We also recommend saving the QGPL and QUSRSYS libraries to SAVSYS tape, typically when you a PTF is applied. This helps keep your PTF level up-to-date in a disaster recovery.

12.2 Carefully plan your backup jobs

Plan your backup jobs carefully to cover your entire system and critical user data. We recommend you at least have:

- One *SYS job to back up the system state
- One *IBM job to back up IBM supplied libraries.
- One or more separate *OBJ and *IFS jobs to back up critical user data.
- One or more separate *OBJ jobs to back up Journal (*JRN) and Journal Receiver (*JRNRCV)
- One *ALLUSR job to cover the rest of the user data, excluding data covered by other *OBJ jobs.
- One *IFS job to cover the rest of the IFS data, excluding data covered by other *IFS jobs.
- If you want to protect spooled files, create one or more *OBJ jobs for output queue (*OUTQ) only with SPLFDTA(*ALL), and exclude *OUT object from other jobs.

Also plan your backup windows and job scheduling carefully. Delta backups significantly reduce backup times, but other tips can maximize your system resources and avoid affecting business hours:

- Separate jobs can run in parallel, except for *SYS jobs. *SYS jobs run quickly, so run them exclusively. Consider running approximately 3 jobs per available processor.
- Keep the job size small. If a job is big, divide it. For *OBJ jobs, keep the total object count per job to 10,000 or less.
- Keep referenced objects in one job. If you have referenced PF and LF in different libraries, save them in a single job.
- Create separate jobs for volatile objects like databases and journal receivers, and for non-volatile objects like programs and service programs even if they reside in the same library.
- Schedule jobs to back up journals and journal receivers to start after jobs to back up databases are completed.

12.3 Journal critical databases and back them up

The IBM i Agent is not a replication solution; it is a data protection solution. If you need to protect your critical databases in a consistent state, we recommend the following:

- Journal the critical database files
- Back up the database files, with save-while-active.
- Back up journals and journal receivers, with save-while-active.

When you need to restore a databases, restore the journal and journal receiver first and then restore the database files. You can then move transactions backwards or forwards using journal entries.

12.4 Logical File (LF) backup and restore

We recommend you include associated physical files and logical files in one job, to keep them as close as possible.

The IBM i Agent only backs up the logical file header. No access path is saved in a safeset. When it is restored, the Agent submits a separate job in the AGENT subsystem to rebuild the access path.

Logical files are backed up and restored in the following way:

- All logical files included in the job are queued to the end of safeset. In the backup log, you should see all logical files are listed at the end.
- Logical files are restored at the last sequence of safeset.
- The access path is rebuilt.

12.5 Database Files

Because IBM limits access to deleted record information, deleted records will not be restored during a restore. Because some very old AS400 applications relied on relative record numbers for indexing, you may need to check with your application provider to see if you need to be concerned with relative record numbers and deleted records.

Note: IBM strongly discourages the use of relative records numbers in applications since it is a system controlled value and gets changed whenever a RGZPFM command (database defrag) is run. If relative record numbers are sensitive in your application, use the API option to back up database files.

12.6 IFS device, block, special character and socket files

IFS device, block, special character and socket files are typically created by the operating system install or operating system running programs. The Agent does not restore these objects and some errors will be listed in the log file when you attempt to restore these files.

13 EVault Customer Care

If you have a question about EVault that isn't covered in this guide, our frequently-updated Knowledge Base contains comprehensive information about EVault. The Knowledge Base is your first stop when searching for any EVault solutions you may need. We highly recommend searching here first for the quickest answers to your questions.

EVault Knowledge Base: http://support.carbonite.com/evault

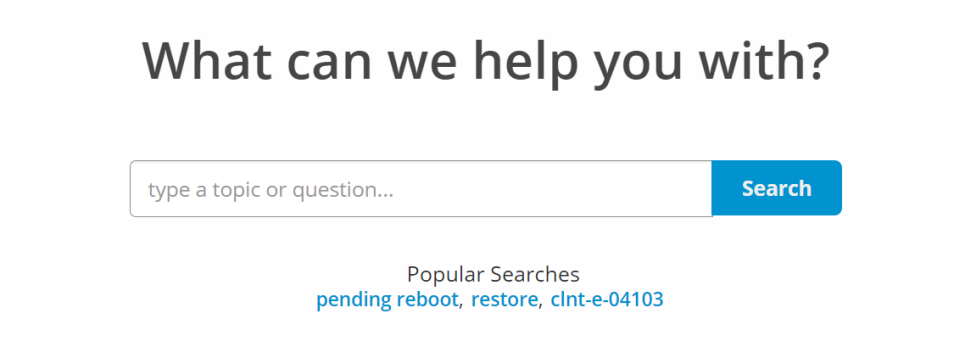

13.1 Contacting EVault

If you need live assistance from a qualified support agent, EVault Customer Care is here for you 24 hours a day, 7 days a week (excluding US holidays). Please feel free to get in touch with us, and we'll help out any way we can! You can find the contact information for EVault Customer Care in the EVault Knowledge Base: http://support.carbonite.com/evault

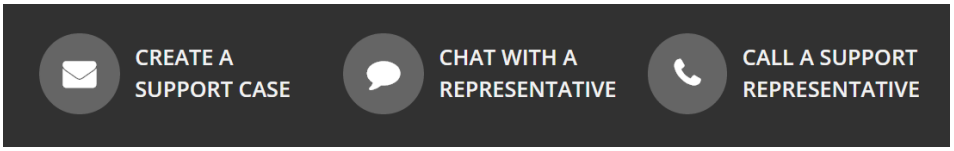

Tip: When contacting EVault Customer Care with a technical issue, please have both the program's log files and the store you are having difficulty with ready.

To gather log files, click **File** menu and choose *Open log folder*. Compress the contents of the folder in a .zip file and attach it to your support request.

If the log archive and/or mail store exceeds 10MB, you may not be able to send them as an email attachment. In that case, upload instructions will be provided to you upon request.