



~

General description of the product

Date the 27th of January. 2022
Version of this document V3.2

Contact :
Bruno Taverne
Email: btaverne@m81.eu
Phone: +33 6 20 27 22 17

1. Product description

1.1 General presentation

The *Recover for i* product has a simple objective: to **reduce data loss to a few minutes in the event of a "crash" or "incident" on your database.**

Of course, the reason of the "incident" can be a "disaster" or a "crash" of several disks, which potentially generates the most important data loss (except if you have an HA solution). In such case (and without HA) the penalty is immediate: you have to restart from the backup of the previous day.

An "**incident**" on the Database can also be defined as anything that "corrupts" the data:

- deleted tables,
- loss of database integrity,
- damaged files,
- ...

due to a multitude of potential causes:

- accidental or malicious manipulation (or hacking)
- big bug in interactive or batch programs
- object damaged or corrupted data due to an unpredictable failure (program crash, maintenance drift, ...).

If essential tables are corrupted, after important and stressful efforts to try to restore the integrity of the DB, **it might become necessary to start again from the last backup and thus to lose from a few hours to several tens of hours of data!**

There are many other examples of "incidents" of all types ... Fortunately for the owners of IBM i the DB2 database has not yet known any virus ... until the first one occurs ...

The target of *Recover for i* is to export to an external server, distinct of your "IBM Power i", all the updates that are made on the Production Database (records creation, deletion, modification) and, in the event of an incident, to import and reapply these updates to the database restored from the last backup.

RPO (Recovery Point Objective) result = a few minutes instead of many hours!

1.2 Mode of operation

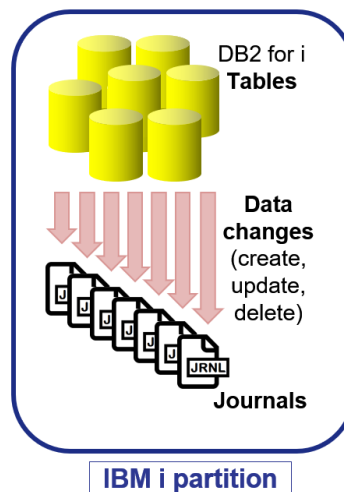
a) Journaling

The **Recover for i** product is based on the Journaling of the Production database.

Like all databases, "DB2 for i" records all the updates made to the data (creations, deletions, modifications) in database Logs called "Journals" on the IBM i. You certainly heard about the JOURNALING mechanism.

Logged data is stored in a "Journal Receiver" attached to a Journal.

**Recover for i is based on DB2
'native' Journaling protection**



IBM i system commands allow on-demand or by program to detach these Journal Receivers (replacing them immediately with new ones).

IBM i system commands allow you to use these Journals:

- to "rollback" the last transactions,
- to reapply from a given step (e.g. on a restored DB) the saved transactions.

This potential is already used by some ERP software to 'guarantee' their transactions.

Nota Bene : there are 2 limits / pitfalls to Standard Journaling :

- the Journals are stored on the disks of the production partition,
- the "rollback" of the last transactions is not always possible

Recover for i will use this IBM i Journaling mechanism improving and automating it.

With **Recover for i** it is possible

- to use existing Journals,
- to create new ones.

b) Journal Export

Thus **Recover for i** will outsource the changes made to the Database by detaching the Journal Receivers, and exporting them to a file server, external to the IBM i partition.

Those outsourcing operations are performed automatically, at regular intervals, for example every 15 minutes or even every 5 minutes. The frequency has to be setup.

The Journal Receivers are simply stored on the external file server for later (potential) use.

The file server must have some disk space to keep the Journals for some days. If needed, **Recover for i** can compress the Journals sent to the file server to reduce the space required. **Recover for i** manages also the cleaning process after a defined number of days, both on the IBM I and on the File server.

The external file server should support **FTP** or **NFS**, which is the standard protocol of most Windows, Linux or NAS file servers (or other appliances like the one of Synology called 'NAS DiskStation' available for a few hundred Euros).

Placing the file server far away from the IBM i server ensures that its contents is retained in the event of a "disaster".

c) Journals Import and Applying

When and How these "Journals" stored at regular intervals on an external file server will be used?

We will see it in the following pages:

- Situation 1: WITHOUT any HA solution
 - Case 1: *Disk crash*
 - Case 2: *complete Disaster*
 - Case 3: *Disaster* with "cold back-up" solution
 - Case 4: *Data Incident*

- Situation 2: "Data incident" WITH an HA solution

Note: By "HA solution"(High Availability solution) we mean one "hot backup solution" from the market (SW replication or disks HW mirroring).

Situation 1: you don't have any HA solution

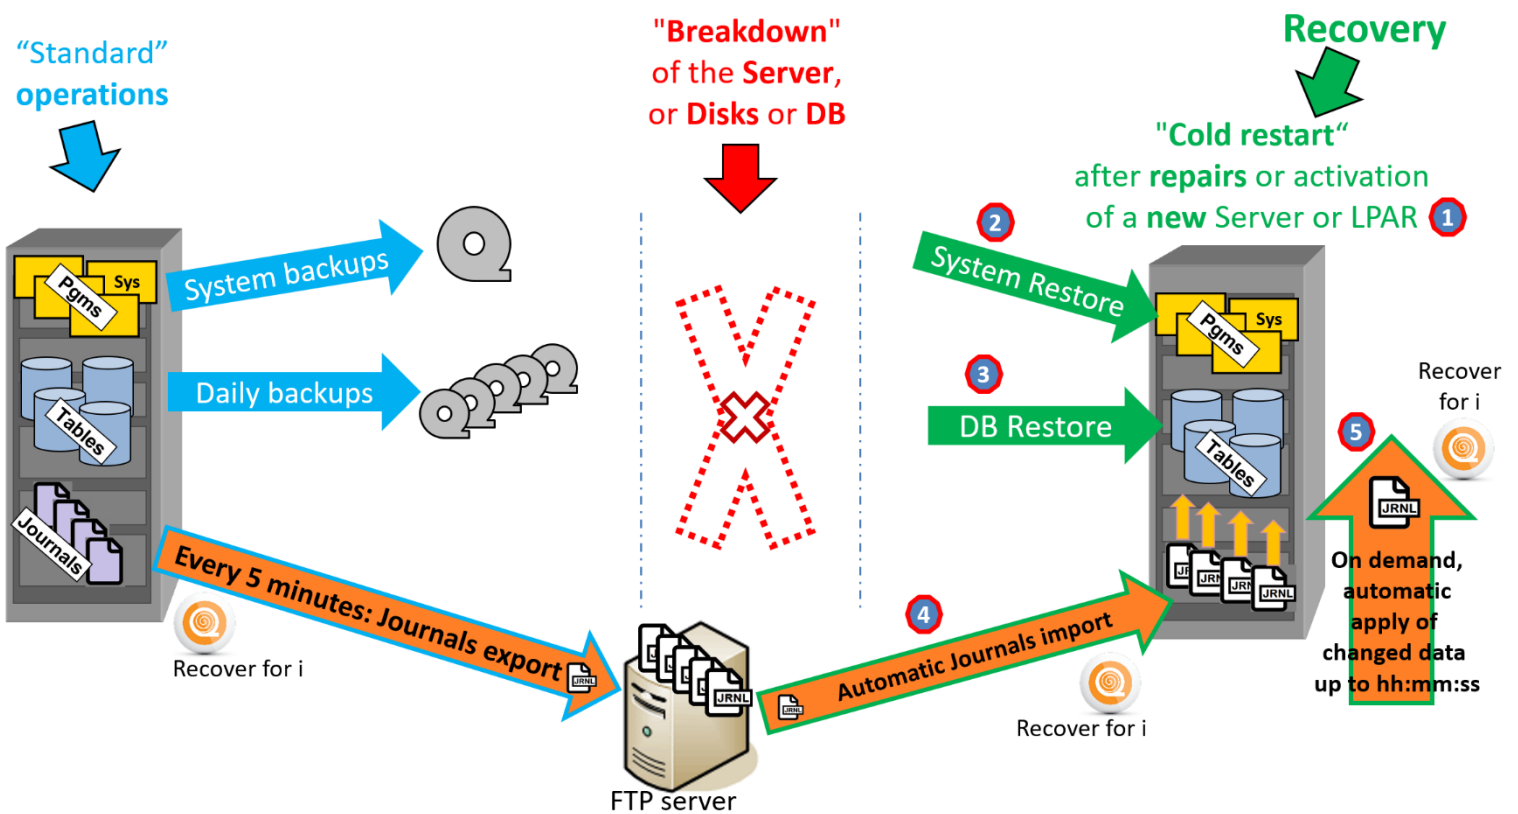
Case 1: disk crash

The database is lost ...

The steps for restarting are as follows:

- Repair of disks by the maintenance service ①
- Complete restoration of the Production partition:
 - o from a system backup tape (SAVE21) ②
 - o then from the last daily backup tape ③
- **Recover for i** is therefore also restored.
 - o With a 1st command, it will automatically
 - retrieve and identify on the file server the Journals created since the last backup,
 - import them on the IBM i partition ④
 - o with a 2nd command it will automatically
 - apply the changes (recorded in the Journals) to the database ⑤.

➔ The database will then be restored automatically to the state it was a few minutes before the disk crash!



Case 2: complete Disaster WITHOUT rescue solution

After a disaster, the server is "out" (or even the entire building or site).

In this case, we come back to the previous case: the same recovery steps are to be performed, except that before you have to supply an intact server ...

In this case, you understand easily the interest to externalise the tapes and the Journals...

➔ **The rescue server will be restarted with a DB as it was 5 minutes before the disaster!**

Case 3: Disaster with a "cold rescue" solution

This variant consists in having a "spare server" on another site or a contract for the provision of a backup server (or for an IBM i partition) with a Hosting firm or a Cloud Provider ; their offers are often called : **DR as a Service (DRaaS)** or **Rescue as a Service (RaaS)**.

In this case **Recover for i** can be configured to send the Journals to the Cloud Provider's file server. The extra cost of a few GB of storage compared to the other components of a DRaaS/RaaS offer will be very low, while the time saving result will be incomparable!

➔ **The rescue partition will be restarted with a DB as it was 5 minutes before the disaster!**

Case 4: Data incident without any HA solution

In case of a "data incident", the difference of process comes only from the following points:

- The Server is still up, as well as the Production partition
- The system has not to be restored from a 'save21' tape, only the last daily backup tape has to be restored on the Production partition
- *After that Recover for i* has imported the Journals on the IBM i partition, you will have to specify in the 2nd command to apply the changes to the entire restored database up to a few minutes before the "incident". (the IBM i administrator will have to determine when the incident took place).

➔ **Even in case of a "data incident", using *Recover for i* means that no more than a few minutes of data are lost!**

Situation 2: " Data incident " WITH an HA solution

If you have an HA solution:

- after a disk crash or a disaster you can find back your data on the remote HA site 😊 ,
- but after a "data incident*" the data loss/corrupted or spoiled data will be **immediately propagated and applied on the remote Standby Database !...**

And after hours of efforts, it might be potentially necessary to decide to restore the DB from the last backup **tape** losing several hours of data (and disappointed to have invested in an HA solution believing to be safe from this kind of misfortune 😞) ...

In case of a "data incident" WITH HA, the process is the same than in Case 4 above

➔ **In case of a "data incident" with HA, using *Recover for i* = only 5 minutes of data lost!**

⇒ ***Recover for i* is therefore an excellent supplement to an HA solution!!**

**corruption, deletion of tables, damaged files, loss of integrity, ...*

2 Complements:

2.1 *Warning :*

Recover for i is not a High Availability (HA) product. Its main differences compared to an HA product are the following:

- Updating of data in the target database is not performed in real time.
- The selection of the objects to be considered is made at library level. It is not possible to exclude certain files or objects.
- Only libraries and their contents are taken into account. No processing is done on IFS or DLOs.

During a phase with intense development activity, this product is not suitable; indeed:

- If the database structure is changed, only the data updated before the change will be taken into account.
- Journaling must never be stopped.
- The processing of programs or commands objects are only partially taken into account.

Although programs, system objects, user profiles, configuration data, audit Journal are not supported in **Recover for i**'s automation systems, it is possible to launch, for example, a daily outsourcing copy of these components to the FTP server by means of a command.

2.2 *Automation and monitoring*

All treatments are fully automatic.

All the operations described above are carried out without any impact on the users or on the process in progress.

Several commands allow you to view (and control) the status of the export, import and re-application of Journals.

A number of commands allows you to automatically monitor the good operations proceedings. These commands can be integrated into a monitoring tool such as Nagios for example (and of course they are integrated into our **Control for i** monitoring product).

2.3 Prerequisite

- IBM i (OS/400): Version 7.1 and above
- Storage: an external FTP or NFS server (including any appliance or NAS with built-in FTP or NFS protocol)
- Space required: experience shows that, in general, the Journals of a day needs a relatively small volume (but there may be exceptions).

2.4 Some références

35 LPARs are already running with **Recover for i**.



This document is not a contractual one. It is provided for information purposes only. M81 reserves the right to modify the specifications or the form of the products without notice and without engaging its responsibility.