

~

General description of the product

Date the 7th of April 2020
Version of this document V3.02

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

Recover for i

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 an "accident" on your database.**

We define "**accident**" on the Database as anything that "corrupts" the data:

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

due to a multitude of potential causes:

- accidental or malicious manipulation
- bugs in interactive or batch programs
- object damaged due to an outright failure
- Corrupted data due to a program crash, or due to an upgrade or a maintenance operation that goes wrong, ...
- cyber crime: viruses, cryptovirus, hacking, ...

Of course, if you don't have an High Availability (HA) solution that replicates data on a Standby Database, the most important data loss is potentially the one generated by a disaster or a crash of several disks.

Indeed, there, the penalty is immediate: you have to start again from the backup of the day before.

But also in the case of the corruption of a single important table... Often, after important and stressful efforts to try to restore the integrity of the DB, it becomes 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 examples of "accidents" of all types ... And fortunately for the owners of IBM i the DB2 database has not yet officially known any virus ... until the first one occurs ...

The idea 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 (record creation, deletion, modification) and, in the event of an accident, 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!

Combined with our *Flash for i* product, *Recover for i* allows you to dramatically improve restart times (RTO) and granularity of data recovery (see Chapter 2).

Recover for i

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" historizes all the updates made to the data (create, delete, modify) in database "Logs" called on IBM i "Journals". In the IBM i environment, everyone has heard about the JOURNALING mechanism.

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


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 on the DB, from a given step, the saved changes.

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

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



It is possible with **Recover for i** to use existing Journals, but also, if necessary, **Recover for i** can take care of creating new ones.

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

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 storage server, external to the IBM i partition. Those outsourcing operations are performed automatically, at regular intervals, for example every 15 or even 5 minutes. The frequency can be setup. The Journal Receivers are simply stored on the external storage server for later use.

The external storage server must have some disk space to keep the Journals for one day (or even 2 or 3 days). If needed, with **Recover for i** it is possible to compress the Journals sent to the storage server to reduce the space required.

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

The fact that the FTP/NFS server is external to the IBM i partition ensures that its content is retained in the event of a "hardware crash" or a "data accident".

c) Journals Import and Applying

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

Case 1: Disk crash WITHOUT backup solution

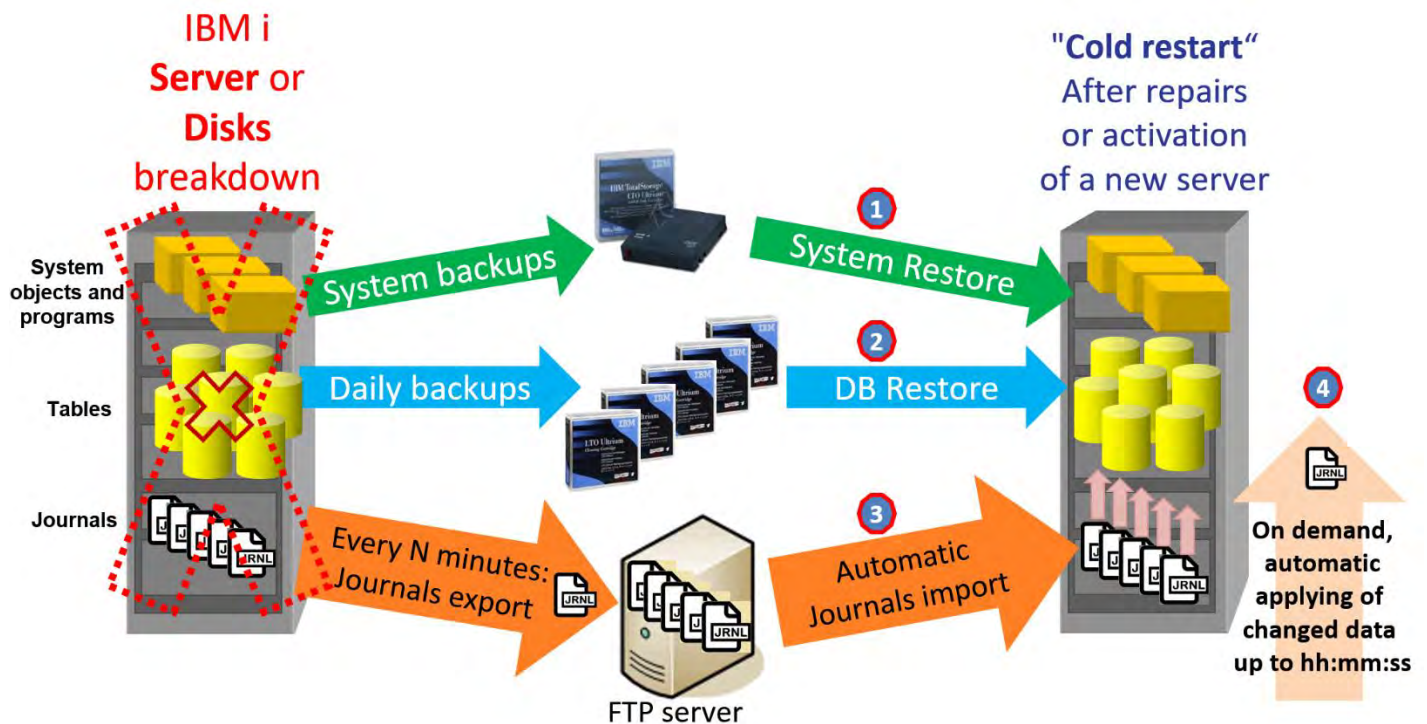
The database is lost ...

The steps for restarting are as follows:

- Repair of disks by the after-sales service
- Complete restoration of the Production partition:
 - o from a system backup (SAVE21) ①
 - o then from the last daily backup tape ②
- **Recover for i** is therefore also restored.
 - o It will automatically retrieve and identify the Journal Receivers on the storage server,
 - o import them on the IBM i partition ③
 - o and apply the changes (recorded in the Journals) to the database ④.

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

Recover for i



Variant 1bis: complete Disaster WITHOUT HA solution

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

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

In this case, we come back to the previous case n°1: the same recovery steps are to be performed, except that upstream you have to supply an intact server ...

In the event of a disaster, outsourcing the FTP file server appears to be better ...

Variant 1ter: Disaster with "cold" back-up solution

This variant consists in having a "spare server" on another site or a contract for the provision of a backup server (or an IBM i partition) from a hoster or a Cloud Provider (= PRA as a Service or **Rescue as a Service** / RaaS).

In the latter case **Recover for i** can be configured to send the Journals to a Cloud Provider's file server. The extra cost of a few GB of storage compared to the basic RaaS offer will be low, while the result will be incomparable!

➔ The recovery partition will be restarted with the database as it was a few minutes before the disaster!

Recover for i

Case 2: " Data accident " WITH or WITHOUT HA solution

If you have an HA solution, after a disaster you can find your data on the backup site. Indeed, they are replicated on the stand-by database ("software" solution) or mirrored on the mirror disks ("hardware" solution).

But when it is a "data accident" (reminder: corruption, deletion of tables, damaged files, loss of integrity, cryptovirus, ...) then WITH or WITHOUT HA solution, the result is the same: **the HA solution immediately propagates the "data accident" on the Standby Database !...**

And after hours of efforts, it will be potentially necessary, HA or not, to decide to start again from the last backup losing several hours of data (and to have invested a lot of money in a HA solution believing to be safe from this kind of misfortune 😞) ...

In the case of a "data accident" (WITH or WITHOUT HA solution), the process is as follows:

- The Server is still up, as well as the Production partition:
 - o restoring the last daily backup tape to the Production partition (or to another partition or to another server)
- **Recover for i :**
 - o will automatically retrieve and identify Journal receivers on the FTP storage server,
 - o import them to the IBM i partition,
- The administrator will determine when the accident took place because it is often better not to apply changes until the exact moment of the accident ...
- **Recover for i :**
 - o Will automatically apply the changes to the entire restored database up to a few minutes before the "crash".

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

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

Recover for i

2 'Recover for i' associated with 'Flash for i'.

Flash for i is another M81 product, available since 4 years and running on 150 IBM i partitions (see www.m81.eu/flash-for-i/).

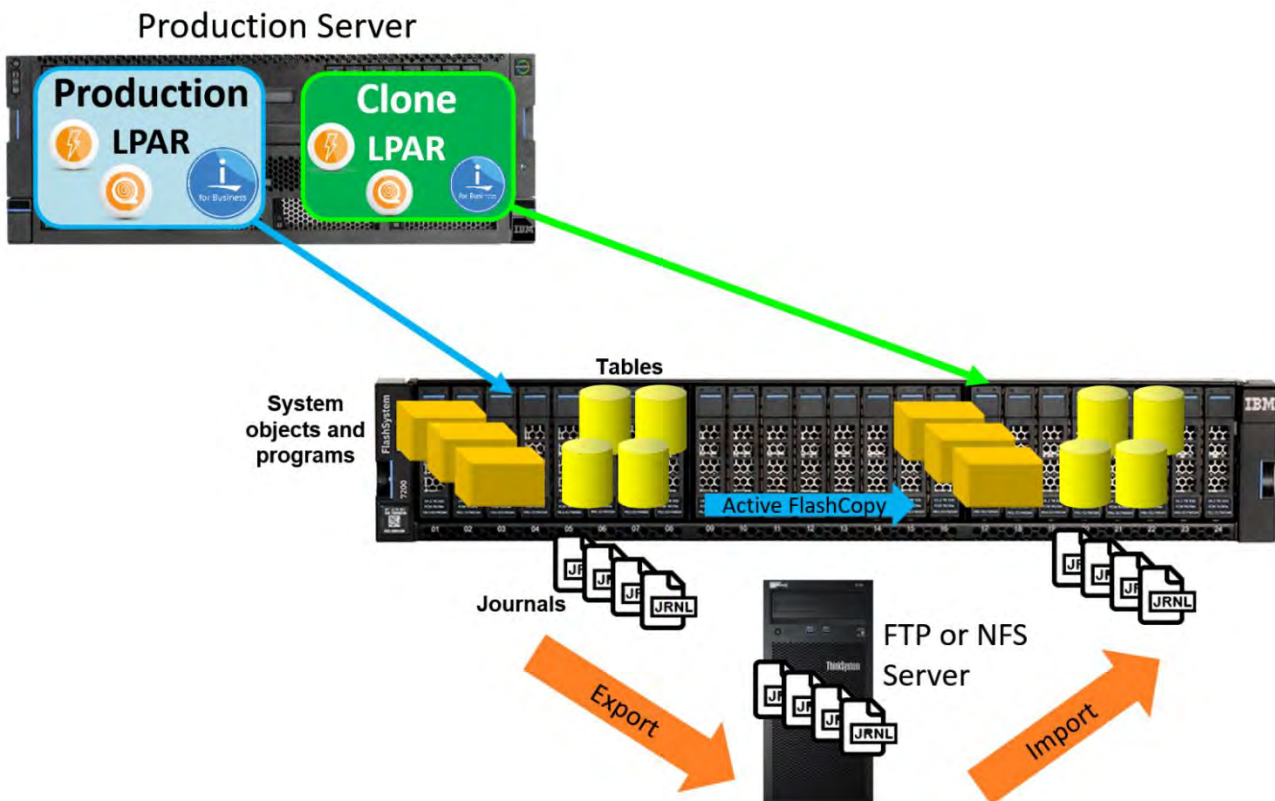
When *Recover for i* is used in conjunction with *Flash for i* then automation, schedule and granularity are improved.

2.1 Operating process

The automated operations are as follows:

- **Flash for i** (on the 'Prod' partition):
 - o manages the creation of a 'clone' of the production discs
 - o boots a 'clone' partition from the 'clone' discs.
 Note: this 'clone' partition can for example be used to make backups without stopping users.
- **Recover for i** (on the 'Prod' partition):
 - o manages the detachment of Journals from the production database every N minutes
 - o sends and stores Journals on the FTP server
- **Recover for i** (on the 'Clone' partition):
 - o searches and detects Journals on the FTP server
 - o imports and restores them.

⇒ The Journals are therefore present on the 'Clone' partition but 'dormant'.



Recover for i

In case of "data accident" on the Production database (always WITHOUT or WITH a HA solution):

- The administrator decides the exact time until **Recover for i** will reapply on the BD the changes (from the previously repatriated Journals).
- Through a single command (on the 'Clone' partition):
 - o **Recover for i** will automatically apply the changed data
 - on the 'Clone' database (which is in the status at the time of the creation of the last clone),
 - until the desired time (hh:mm:ss).

→ **A complete, upstanding Database, with data dating from a few minutes before the "accident" is therefore available and operational on the 'CLONE' partition.**

To return to PRODUCTION with a Non-corrupted DB, administrators have two solutions:

- A. Manually copy the updated tables from the 'Clone' partition to the 'Prod' partition.
- B. Boot the Prod partition using the 'Clones' disks (containing the updated Database). This last technique is called by IBM "Reverse FlashCopy" and is documented by M81.

→ Thus, even in case of a "data accident" the duo **Recover for i** and **Flash for i** allows:

- to restart very **quickly**
- with a complete Non-corrupted Database,
 - o whether processing 100% of the database at once...
 - o or just one by one, the few **libraries** or **tables** that have been corrupted.

2.2 Variant: operations carried out on the HA remote site

With a backup solution in place (HA Software or Hardware)

⇒ *The product duo "Flash for i" and "Recover for i" can be used on the backup partition.*

The objective remains the same: to minimize the potential loss of data in the event of an "accident" on the PROD database (such as table deletion, loss of integrity, basic corruption, etc...) knowing that the incident is instantly reproduced on the backup database by the HA solution in place.

⇒ *The interest lies in the fact that two process can be carried out in parallel:*

- o *attempt to restore the integrity of the Production DB on the local Production partition*
- o *perform the "recovery" operations of the DB on the remote HA configuration*

With **Flash for i** and **Recover for i** used on the **backup** partition, all the operations described in the previous chapter (§22.1 can be carried out and allow a reboot.

- **quickly**
- on the **backup site**
- with **only a few minutes of data loss!**

Recover for i

3 Complements:

3.1 *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.

3.2 *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** product).

Recover for i

3.3 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 takes up a relatively small volume (but there may be exceptions).

3.4 Some références



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.