
Overview of "Allow Multiple Instances of Bulk Import"

- **Product:** Voyager
 - **Product Version:** 8.2.0
 - **Relevant for Installation Type:** Multi-Tenant Direct, Dedicated-Direct, Local, TotalCare
-

Question

Overview of "Allow Multiple Instances of Bulk Import."

Answer

Support for multiple instances of Bulk Import (using the -M flag to run concurrent bulkimport sessions) was introduced in Voyager 8.2.1. It applies to sites with single databases and sites with [multiple databases](#) (where /m1/voyager is shared).

The -M flag essentially forces the operator to intentionally run multiple imports at the same time. Prior to Voyager 8.2.1 there was no way to "enforce" Support's long-standing recommendation to limit imports to one at a time. In other words, it is a safeguard to prevent an operator from accidentally running concurrent bulk imports.

Be aware that running multiple instances of Bulk Import [can cause the import to run more slowly](#) than if you were doing one single Bulk Import at a time. It can also have a significant impact on server performance in general.

Note

A good best practice is to start each new process *at least two minutes* after the one previous to avoid overloading the system.

Be careful to not run concurrent Bulk Imports that are importing the same records. Such imports should be run sequentially. Using the -M flag *assumes you have confirmed* there is no data overlap. If there is data overlap, data corruption can occur.

It is a known behavior that running multiple instances of bulk import can lead to the [inability to retrieve records using limits](#).

Additional Information

See the [Technical User's Guide](#), Bulk Import chapter.

A reclamation project is a good example of a use case for the -M flag because *there is no chance of data overlap*.

Some customers wonder how many concurrent imports can they run. It is difficult to predict this and provide a good answer because environments vary greatly from installation to installation, and while some installations may have ample resources to support multiple imports, others may not. The old adage: "your mileage may vary" applies here.

One suggested course of action is to do benchmark testing in the training database:

1. Import one file
2. Import two files simultaneously (always separate processes by at least two minutes)
3. Import three, four or more files simultaneously (the expectation here is that you will see a larger impact on import performance (imports will get slower) at or around this point, and server performance in general may be impacted).

This should provide a fair estimate of how similar imports will perform in production, what the boundaries are for your particular environment, and provide an opportunity to adjust any import plans before your production environment is impacted.

-
- **Article last edited:** 14-Sept-2019