BCV4

IN FOCUS

Homogeneous System Copy? Who needs that?

Your company will need that as soon as it begins to deal with large scale application tests against production data. A subsystem clone always is the least costly method to create test environments containing live data, playgrounds for impending release changes, read-only query systems, and much more. We call that a pre-production environment or a golden copy. Another good reason for a “cloned” pre-production as a test data source: consistency is established by BCV4 at copying time; after that the consistency state of the cloned environment is “frozen” since no applications are active in the “golden copy”.

Every night a new production mirror is created, but those volumes are intended for emergency purposes and must not be touched. Can I use BCV4 anyway?

No problem at all. BCV4 considers the mirrored volumes as “intermediate copy” and propagates their contents to a third set of volumes. From this place the regular “make useable” process is then executed.

What is the fastest method to duplicate an entire DB2 subsystem?

A volume copy, carried out with BCV4. BCV4 makes DB2 and, optionally, IMS useable on these volumes, as new subsystems with different names.

Why is BCV4 that fast?

BCV4 has two important working areas: renames within the ICF Catalog and renames within the DB2 Catalog. For both tasks it does not use standard utilities but spezialized application programs. In addition, BCV4 parallelizes tasks wherever possible.

How can I make sure that job member names and number of job remains the same in each new clone execution, in order to avoid confusion for the job scheduler?

There is no need at all for special precautions. On each new clone execution BCV4 creates the same number of jobs, and the names of all job members remain the same.

Do we have to ask Sysprog to create the ICF Catalogs for the target system?

Among all other required steps, the fully automatic BCV4 process contains the creation of target ICF Catalogs.

How do I prepare the BSDS datasets for the target system?

You don’t have to deal with this. BCV4 creates a fully automatic job chain that contains each and every step of the cloning process.

I want to clone DB2 but there is a large amount of non-DB2 datasets with the same first level qualifier like DB2. They must not be lost when the system is cloned.

There is a processing option for BCV4 that allows it to exclude certain target volumes from being overwritten and to keep and recatalog their datasets, although the associated ICF Catalog is temporarily removed. For many other “everyday problems” like that, there is a resolving option in BCV4, for example regarding tape datasets, GDGs, migrated datasets, etc.

Is Flashcopy™ a prerequisite for BCV4?

It is not. In order to use BCV4 you neither need Flashcopy nor another fast copying tool. Only if you are planning to clone active DB2 subsystems is such a tool is required.

Is it possible to merely refresh DB2 subsystems using BCV4, rather than replacing them as a whole?

By definition, a clone is an absolutely identical copy of the source system. It is either replaced or stays as it is. For a high-performance refresh of DB2 subsystems we recommend to use our products BCV5 or BCV6.

Is it possible to use BCV4 with a job scheduler?

Yes, the JCL generated by BCV4 is compatible to schedulers. It also supports schedulers which require unique job member names across all scheduled applications. In addition, BCV4 has an automativ submit function for those cases where creating a job net appears too time-consuming.

Our production is more than DB2, of course. There is the IMS, and there is a flock of ancient VSAM clusters, too. While we are at cloning, then those have to be cloned, too. Is that feasible?

BCV4 supports clones which contain DB2 and IMS and collections of VSAM clusters.

The system to-be-cloned comprises of more than 400 volumes. Do I have to build this huge volume list by myself? And how can I avoid the trouble to keep that list current?

The only items to specify are the SMS storage group names of source and target system. BCV4 builds the volume list from this information. And it automatically refreshes the list before each new execution of the clone.

We cannot use volume copies. Our DB2 table spaces are spread unsystematically across all volumes. No chance to use BCV4, right?

Sure you can! BCV4 supports copies on dataset name level as well. Again, a fast copying tool is not prerequisite but may be useful.

We have model-3, model-9, and model-27 volumes. May this be an obstacle for the clone?

BCV4 takes care that volume pairs always contain matching volumes. However, volume copies between differently sized source and target volumes are forbidden by technical reasons. If your environment requires cloning between unlike volume types then a copy on dataset name level is to be considered.