



Rocket® Mainstar Clone and Rename for IMS

Reduce downtime significantly

Simplify copying IMS subsystems and databases

Reduce costs through automation

Improve productivity

Increase data availability

Support business resiliency initiatives

Clone entire IMS subsystems or select databases quickly and easily

Quickly Provide Access to Cloned Data

Rocket Mainstar Clone and Rename for IMS (ICR) provides greater data access by simplifying and automating the process of cloning entire IMS subsystems.

Many applications have critical requirements for quick and accurate cloning of entire IMS subsystems or select databases, due to the need for production or development copies, recovery initiatives, query-only access, reporting, or testing. Using current techniques, data cloning can require the better part of a day. With the ever-increasing demands for usable clones, you need an easy and automated data cloning process with minimal downtime.

Simplify the Cloning Process

Every hardware vendor and some software vendors use fast copy products to replicate data. However, the challenge of accessing the replicated data persists. Producing a usable clone on the same or shared image can be a manual and cumbersome process. As the demand for usable clones continues to escalate, the roles of IMS Database Administrators (DBAs) and Storage Administrators are becoming more complicated. With ICR, you can provide access to cloned data quickly.

The Data Access Dilemma

You make replicas, but how can you easily access and use the cloned data?

While fast copy tools can create replicas very quickly, the process does pose challenges:

- ❖ If the source volume label is retained, the target volume is varied offline by the system, making the data unusable for any online requirements.
- ❖ If the target volume label is retained, the target volume stays online. Because the data set names are replicas of the source data set names and are already cataloged, the data set names must be changed.
- ❖ Internal data, such as VTOC, VTOCIX, and VWDS, are replicas of the source volume names, and must be renamed before use.
- ❖ IMS control data sets reflect the source data set names, so unless the IMS control data sets are updated, the cloned IMS system will not be able to access the target data sets.

So, how do you access the data?

The Solution

ICR automates the manual tasks associated with cloning IMS subsystems and select databases by renaming the data sets, volumes, volume internals, and IMS internals; by providing access automatically; and by eliminating the need for a separate image.

Because the ICR processes are extremely fast, your replicated data is available in the shortest amount of time possible.

Refresh IMS Subsystems and Databases in Minutes

ICR can refresh IMS subsystems and databases in minutes through the use of fast copy tools and automation. ICR utilizes fast copy technology to

minimize the time IMS is offline. Therefore, an IMS clone can be created within minutes.

Automate the Cloning Process – Fast Copy Automation

When you're using IBM® FlashCopy®, EMC TimeFinder/Clone, EMC TimeFinder/Snap, or StorageTek SnapShot, ICR initiates volume copies by invoking DFSMSdss. ICR automatically pairs the volumes by manufacturer-specific requirements, SMS and non-SMS status, and by device size to eliminate failures.

Additionally, ICR can optionally re-label the target volumes and vary them online.

Support for More Efficient Storage Compatibility

You can use incremental FlashCopy when performing volume-level copies during subsystem cloning. By only copying the changed tracks, incremental FlashCopy reduces the amount of data that the storage processor must copy in the background processor when a volume copy request is issued.

You can also use space-efficient-capable target volumes when performing volume level FlashCopy during subsystem cloning. Space-efficient FlashCopy reduces the amount of storage that must be occupied by the target volumes by the storage processor when a volume copy request is issued.

IMS Clone and Rename includes new PPRC Preserve Mirror options when FlashCopy is performed and the target volume is a PPRC primary volume. Preserve Mirror reduces the amount of data that must be transferred from the primary storage controller to the secondary storage controller in order to get the target volume back into sync with its PPRC secondary volume.

Updating Cloned IMS to Reflect New Names

Whether you have cloned an entire IMS subsystem and its associated databases, or just individual databases, ICR automates the changes that need to occur on the target IMS to enable the use of the data you have just cloned.

When cloning an IMS subsystem, the RECONS, Dynamic Allocation library (MDA), and PROCLIB members are updated. When cloning select databases, the recovery information for the cloned databases is updated, automating the entire process.

Maintain Data Integrity

Storage Based Consistency to Provide Application Availability

Utilizing the consistency function of FlashCopy or EMC TimeFinder/Clone for volume-level fast-replication allows all of the volumes that are being copied to be copied at a consistent point in time. When the hardware consistency function of FlashCopy or EMC TimeFinder/Clone is used, you do not need to suspend IMS updating when performing the copy process, allowing the application to be more available during the cloning process.

Apply Logs After Refreshing a Database in Order to Make the Database Consistent

You can refresh databases without taking the source databases offline, and you can create target databases that do not have any in-flight units of work. By reading the archive logs from the source IMS environment after a fuzzy copy of the source

database has been made, and then applying them to the copy or target database, all databases refreshed in the same IMSDBREFRESH command are updated to the same point of consistency.

Catalog Integrity to Ensure Rename Success

Some of the data required to rename data sets, such as PATH entries and GDG base records, only exists in the ICF catalog, making it imperative to capture the information from the point in time the replica is made. ICR creates backups of the source ICF catalog that reflect the status of the source volume data sets at the time the point-in-time (PIT) copy was created. Alternatively, you can read the source ICF catalogs off of the target volumes if they were also included in the PIT copy.

When the source ICF catalogs are backed up, ICR uses the same fast catalog backup logic that was developed in Rocket Mainstar Catalog RecoveryPlus.

Integrity Checks to Eliminate Orphaned Data Sets

When copying volumes, you don't want to miss pieces of multi-volume data sets, VSAM spheres, active generation data sets, or data set name aliases. If not copied, these become orphaned data sets that cannot be used. ICR provides complete data integrity checking.

Preserve Sensitive Data with Masking

You can mask or scramble sensitive data such as credit card numbers, names, and addresses when refreshing databases. This ensures the protection of data when cloning across different environments.

Support the DB QUIESCE Command for Refreshing Databases

You can perform a database refresh and stop the source database using DB QUIESCE support.

Extensive DFSMS Class Assignments for Enhanced Flexibility

ICR provides extensive DFSMS class assignments for target data sets selected during the cloning process. SMS class information can be specified to:

- ❖ Accept the SMS values copied from the data set names of the source volume.
- ❖ Assign specific data classes, management classes, or storage classes.
- ❖ Re-enter SMS ACS routines.

Easy to Use

Dynamically Define New Databases on the Target

Avoid manual intervention by optionally creating the database objects on the target if needed. Defining the new database on the target makes the cloning process more manageable and streamlined.

Specify Volumes by Volume Mask or SMS Storage Group

If you have data that spans volumes, ICR ensures that all related volumes are processed together. This guarantees that all necessary pieces of data are copied, and prevents data from being orphaned. When using FlashCopy or SnapShot, volumes can be specified by volser, volser masks, SMS storage groups, or any combination to minimize setup errors.

Rename More than Just the High-Level Qualifier

Target volume data sets are renamed based on the RENAME masks specified.

- ❖ Include changes to one or more qualifiers of a data set name.
- ❖ Introduce or eliminate entire qualifiers.

A SAFE option allows the RENAME step to be rerun if rename masks are incorrectly specified, without losing the PIT copy.

Remove Catalog Entries from Previous Clones

ICR can catalog target volume data sets to either a populated ICF catalog or an ICF catalog that is initially empty except for the aliases for the renamed target data sets. ICR provides support to clean out any data set entries from a previous cloning execution.

If each copy process is repetitive, each cycle of the process is likely to leave orphaned data sets. Orphaned data sets are created when a catalog entry exists for a data set that does not exist. These orphans are data sets (copied from a previous execution) that exist in the target ICF catalogs, but are not presently on the source volumes, and therefore will not be copied, leaving catalog entries without associated data sets.

ICR identifies all data sets cataloged in the previous run and cleans out everything written to the ICF catalog on the previous run, leaving all other catalog entries intact.

Delete Target Database Data Sets from Previous IMSDBREFRESH

You can delete the target database data sets from a previous IMSDBREFRESH to free up the space allocated by the target database data sets after they are no longer needed.

Validate Through Simulation

The ICR simulate feature performs all of the normal cloning functions, except the actual initiation of copies or the updating of any IMS data sets. ICR will:

- ❖ Verify syntax, match source to target volumes, and display DFSMSdss COPY FULL commands if using FlashCopy or SnapShot, without initiating the copy operations.
- ❖ Perform masking comparisons to both ICF catalog and volume records to detect any potential problems such as:
 - Incomplete VSAM sphere renames
 - Multi-volume data sets not wholly contained in the volume list
 - Catalog entries that match a mask and have one or more volumes that were not copied
- ❖ Detect database data sets on volumes that were not included in the copy process.
- ❖ Detect any data sets needed for recovery that were not included in the copy process.
- ❖ Verify that the required IMS control information is available for the databases being copied.

High Level Features and Benefits

Automatic Pairing ❖ Provides automatic pairing of volume characteristics. ❖ Increase productivity through automation.

Flexible Cloning Options ❖ Allows FlashCopy or SnapShot by volser masks, entire storage groups, or any combination to eliminate the requirement for individual volume specification. ❖ Simplify

Validity Checking ❖ Automatically validates the integrity of the data before it is cloned so that all required data is encompassed. ❖ Reduce the risk of error.

Rename Volume Conflicts ❖ Quickly fixes volume conflicts so you can use data from the same LPAR. ❖ Improve data access.

Extended Rename Capability ❖ Renames the target data sets and allows you to change, add, or delete any qualifier. ❖ Tailor data cloning to meet the unique needs of your environment.

Extensive SMS Options ❖ Enables you to determine how the SMS class constructs will be applied to your cloned data sets to ensure they are managed correctly. ❖ Facilitate data management.

IMS Cloning Support ❖ Helps you clone IMS subsystems or databases in minutes - without requiring a separate image. ❖ Facilitate data management.

IMS Data Sharing Support ❖ Supports either IMS data sharing or from data sharing to non-data sharing. ❖ Clone IMS data, regardless of data sharing status.

IMS Internals Support ❖ Updates the appropriate target IMS control data sets eliminating the need for a lengthy, manual cloning process. ❖ Improve productivity.

System Requirements

Operating System Environment: Any z/OS version; and OS/390® Version 2 Release 8 or later

 www.rocketsoftware.com

 info@rocketsoftware.com

 twitter.com/rocket

 www.youtube.com/rocketsource

 www.linkedin.com/company/rocket-software

 plus.google.com/u/0/104109093105646534918

