Migrating to a new version of IMS with the help of IBM IMS Cloning Tool

Document version 1.1

Rosemary Shay, Rocket Software

Ron Bisceglia, Rocket Software

CONTENTS

LIST OF FIGURES

Executive summary

Migrating to a new version of IMS is required to stay current with technology. The journey is a formidable effort with high visibility to critical applications. The challenges of migration include refreshing test environments quickly to enable testing with current data, and maintaining exact replicas of your production environments to facilitate system testing that accurately reflects production characteristics and workloads. Because IBM's IMS Cloning Tool can quickly refresh data and test environments to significantly reduce your risk of impacting critical applications during the migration process, it helps you avoid surprises and delays.

Traditional methods for creating test environments typically require downtime and are lengthy, resource-intensive, manual processes. Creating a new IMS environment using traditional processes requires highly skilled resources to create new IMS PROCLIB and JOB members, create and run IMS SYSGENs, create new RECONs, and create new MDA (MVS Dynamic Allocation) members just to set up the environment. Additionally, to populate the databases, an UNLOAD/RELOAD process or an IMAGE COPY/RECOVERY process must be created and run for every IMS object in the source environment. Whichever method you use, cloning your databases is very labor intensive, requires a lot of host resources, and may require an outage on your production databases to eliminate the risk of pointer errors on the target database.

IMS Cloning Tool automates, simplifies, and speeds up the cloning process without affecting production IMS application performance or availability, allowing you to test with current data in your test environments; thus lowering your risk of affecting sensitive IMS business applications when you implement a new version of IMS into your production environment. IMS Cloning Tool is the foundation for cloning production IMS systems and databases. Automation is provided to perform meta-data management functions, allowing the data to be used by a cloned IMS system on the same or another shared disk System z logical partition (LPAR). It supports both data sharing and non-data sharing IMS environments.

IMS Cloning Tool leverages storage-based fast-replication to copy the IMS system and its data. Fast-replication allows database data to be replicated nearly instantaneously reducing copy time, minimizing outages needed do to the copying of data, and offloading the host resources needed to copy the data to the storage processor thus saving host CPU and I/O resources.

IMS Cloning Tool provides the automation required to leverage fast-replication and other data copy processes to refresh test environments or create new test environments very quickly. It performs volume reconditioning and data set renaming operations to allow the replicated data to be used by a cloned IMS system on the same or another shared disk LPAR.

This document will describe three types of cloning operations and how they can be used throughout the process of migrating to a new version of IMS to reduce the time and resources required and reduce the risk associated with implementing a new IMS version.

Applying the New Version of IMS

An IMS System Skeleton clone is often created to test the apply process. It is a very small environment that contains all of the system information used to define and bring up a cloned IMS environment, including database and application definitions, but does not contain any database data. This environment is well suited for testing your entire apply process and the nuances involved in activating functionality of a new IMS version.

How IMS Cloning Tool creates an IMS System Skeleton clone

The following figure depicts the processing steps that IMS Cloning Tool performs to clone an IMS System with the apply process:

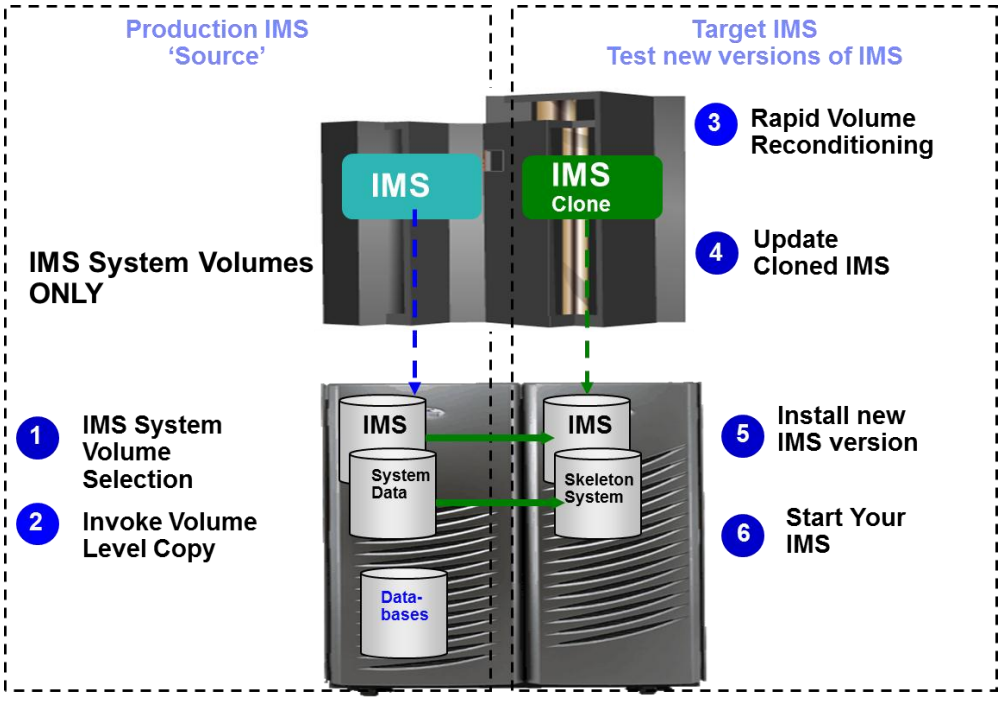

Figure 1: Creating an IMS System Skeleton Clone with IMS Cloning Tool

Process:

1. The user identifies the storage volumes that make up the production IMS system using specific VOLSER IDs, VOLSER masks, or SMS storage groups.
2. IMS Cloning Tool invokes DFSMSdss Copy to issue IBM FlashCopy or SnapShot commands, or EMC Timefinder/Clone commands to copy the data and back up the ICF catalogs that point to the data sets on the volumes being copied.

   To achieve data consistency during the copy process while IMS is up, there are two options:

   - Option A: IMS Cloning Tool internally suspends the IMS log updates for the milliseconds it takes to perform the copy. This process is called, "IMS SUSPEND" and the process ensures all production volume data is consistent.

   - Option B: IMS Cloning Tool invokes a storage-based consistency function whereby all the volumes in a consistency group are frozen at the same time for the milliseconds it takes to perform the copy, thus ensuring a consistent copy of the production environment.

3. IMS Cloning Tool performs Rapid Volume Reconditioning:
    a. If the target volumes are not already online, the volume internal identifier of each target volume is re-labeled so that they can be brought online to the same z/OS system, or a different z/OS system, without volume label conflicts.
    b. The VTOC and VVDS on the target volumes are renamed to match the target VOLSER names. The data set related entries in the VTOC and VVDS are updated to reflect new data set names to be used by the new, or target, IMS system.
    c. The data sets on the cloned volumes are renamed and re-cataloged. Data sets are renamed to eliminate duplicate data set names from a z/OS catalog perspective and to allow the new data set names to be integrated into the cloned IMS system. The cloned and renamed IMS data sets can then be used on the same LPAR and integrated into a different IMS system using the new data set names.
4. IMS Cloning Tool updates the cloned IMS system to reflect the new data set names, VOLSERs, IMSID, etc. The cloned IMS system updates include:
    a. RECON data sets – the data set names, IMS subsystem IDs, and VOLSERs are updated in the following RECON records: header record, database data set records, online log records, and back-out records.
    b. Optionally, the following RECON records are updated: image copy records, change accumulation records, System Log Data Set (SLDS) records, and Recovery Log Data Set (RLDS) records if they were on volumes that were cloned.
    c. IMS.PROCLIB, IMS.JOBS, and user JCL library members are updated with the new values for IMSID, VOLSERS, and data set names
    d. MDA (MVS Dynamic Allocation) members for databases or system data sets, such as the RECON, Online Log Data Sets (OLDS), or Write-Ahead Data Sets (WADS) are updated to reflect the new data set names.
    e. Dynamic Resource Definitions (DRD) data sets are updated if they are used.
    f. The default IMS SSID is updated in the online and batch nucleuses in IMS.SDFSRESL.
    g. The repository data sets, if they are used, are updated.
    h. If IMS data sharing is involved, each additional IMS data sharing member is updated.
5. Test your process for applying a new version of IMS.
6. Start your newly versioned IMS.

Testing the apply process in your IMS System Skeleton clone

With a better understanding of what an IMS System Skeleton clone can be used for, let's review the general steps in the apply process that can be tested in your IMS System Skeleton clone. Keep in mind if there are any failures in these steps, you do not have to go through a back-out process to start over. Simply rerun the cloning steps in order to recreate your IMS System Skeleton environment.

Procedure:

1. **Install prerequisite software and maintenance.** Apply prerequisite updates for tools, vendor, and home-grown software. Also apply any maintenance that has been determined critical to be included in the new IMS version.

2. **Apply coexistence maintenance to other IMS systems.** In this step you apply the coexistence APARs/PTFs to existing systems. You may have several related IMS subsystems in your IMS System Skeleton clone environment.
3. **Install new version of IMS.** Apply the new IMS version to the IMS System Skeleton clone via SMPE, CBPDO, ServerPac, or a home-grown process.
4. **Evaluate and reassemble IMS exit routines**. Any IMS exit routines you may have, such as the RECON I/O exit, randomizers, sparse secondary index routines, etc., can be reassembled with the new IMS MACLIBs.
5. **Run System Definition.** The IMS System Skeleton clone already has all the system definitions from the source IMS system. However the migration process dictates a regeneration of system definitions with the new version of IMS software. A copy of the GENDECKS will need to be created for the cloned system. You may need to change the SVCs and SDFSRESL data set names to the cloned IMS names. You will now be able to run your sysgen process for the new, or cloned, IMS system.
6. **Verify RECON compatibility.** When migrating to a new version of IMS, an area of particular concern is the RECONs. The recoverability of your IMS system depends on these data sets, yet records within these data sets are often changed from one version of IMS to another. In an IMS System Skeleton clone you will have a full copy of the RECONs from your source system. So when you update your RECONs, you can test fallback processing, test coexistence compatibility and change your MINVERS. Using an IMS System Skeleton clone will allow you to run your tests with a RECON full of Image Copy records, Logs, Change Accumulation and Allocate records, database, database data set, and partition definitions, etc. improving the likelihood of finding problems before a production implementation.
7. **Generate updated ACBs.** You can run the ACB generation process associated by most IMS version upgrades to determine if there are any problems that need to be addressed and also to create the new ACB library members prior to a production implementation.
8. **Review your cloned IMS.PROCLIB and IMS.JOBS data sets**. After creating an IMS System Skeleton clone it is recommended to review your cloned IMS.PROCLIB and IMS.JOBS data sets. These are replicas of the source libraries and they will have all the same configuration options and components. Take the time to evaluate what components and configuration options you need in your IMS System Skeleton clone and adjust appropriately. After the IMS.PROCLIB and IMS.JOBS members are initially configured for the new or target IMS system, they should not have to be changed for future cloning iterations.
9. **Start the cloned IMS system.** The cloned IMS is now ready to be started using the new IMS version.

Troubleshooting:

If the cloned IMS does not start successfully, check the following:

- Are you missing any modules or exits?
- Did your RECON upgrade correctly?
- Did the vendor and home-grown software generate errors?
- Did you skip a step?

When you have identified what may have prevented the new, or cloned, IMS system to start successfully, you can simply re-clone your IMS System and try again. The process is

one easily repeatable job. No back-out is needed, as the Cloning Tool automatically cleans up between cloning cycles.

Activating IMS V12 features in your IMS System Skeleton clone

Once you have successfully applied the new version of IMS to your IMS System Skeleton clone, you may want to activate some of the new features of IMS V12, including:

- Dynamic Buffer Pools.
- Dynamic Resource Definitions (DRD). You can test the new IMPORT OPTION(UPDATE) command and the new IMS Repository, or the steps to implement Dynamic Resource Definitions. This is an excellent environment to try DRD and the various implementations.
- Member Online Change NAMEONLY option.
- Migration to Log Striping and Log Buffers in 64-bit Storage. There are several steps to each of these migration processes that can be tested in the IMS System Skeleton clone.
- Extended Address Volume (EAV) for non-VSAM data sets.
- RECON MINVERS. You will have your full source RECONs as you execute this upgrade.
- IMS 11 SPE User Exit Enhancements (APAR PM04456/PTF UK67199), are now included in IMS V12.

Cloned environments are so easy to create they become disposable. If you want to change your process, try a new configuration, determine the impact of not upgrading a vendor's software, or manually clean up a RECON, just run the one job and re-clone your IMS System Skeleton environment.

Testing Business Applications on the New Version of IMS

IMS Cloning Tool also provides the ability to refresh specific IMS databases from a production IMS system. Having the ability to refresh specific databases from a production IMS system can save time and resources compared to cloning a complete IMS system.

IMS Cloning Tool provides facilities to refresh databases within or across IMS systems. The target IMS database definition may exist on the target IMS system before databases are refreshed, or optionally, IMS Cloning Tool will use the source characteristics to dynamically define new target database definitions, ACBLIB members, MDALIB members and DBRC definitions. IMS Cloning Tool assures the compatibility of the database definitions on the source and target IMS systems prior to initiating the data replication process. This will prevent potential errors when accessing the refreshed databases that could be time consuming to detect and resolve. IMS Cloning Tool can refresh IMS databases while they are online creating fuzzy database clones. The IMS Cloning Tool Log Apply option can then be used to make the cloned databases transactionally consistent. Optionally IMS Cloning Tool can either issue an IMS DBR, DBD or Quiesce command against the selected source databases and indexes to create a consistency point for transactionally consistent cloned databases.

IMS Cloning Tool can use storage-based data set level fast-replication to replicate production IMS data to a target IMS system. Target IMS data sets are renamed on the refreshed copy during the data set fast-replication process to associate them to a target

IMS system. Target IMS meta-data is updated to reflect the refreshed databases. Refreshing IMS databases using storage-based fast-replication takes only seconds to replicate the data and just a few minutes to complete the entire database refresh process.

If you are starting with a small IMS System Skeleton clone, you can selectively pick databases and just clone them into the IMS System Skeleton clone as needed. Because the IMS System Skeleton cloning process cloned all of your definitions from your source environment, the database definitions will not only exist but they will also be in sync with the source database definitions.

How IMS Cloning Tool refreshes databases

The following figure depicts the processing steps that IMS Cloning Tool performs to refresh databases from a production IMS system to an IMS System Skeleton clone or to an existing IMS environment:

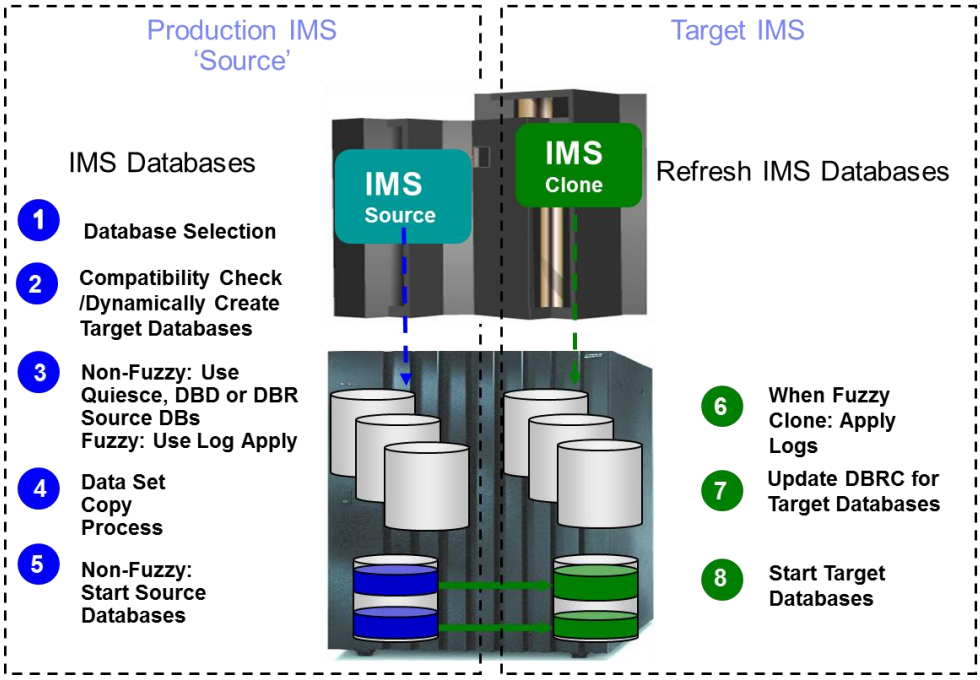

Figure 2: Refreshing databases from a production IMS system into an IMS System Skeleton clone or existing IMS environment

Process:

1. The user selects the databases to be refreshed on the source IMS system by database name or database mask. IMS Cloning Tool determines all indexes associated with the databases to be refreshed and also determines all associated data set names for the specified databases.
2. IMS Cloning Tool performs checks to ensure that the characteristics of the source and target IMS databases are compatible. Some characteristics that are checked include database type, access method, number of segments, data set groups, blocksize, randomizer parameters, number of partitions, partition selection routine, and KEYSTRNG. If the target environment is an IMS System Skeleton clone of the environment of the source database, all database characteristics will be compatible. Optionally, IMS Cloning Tool will use the source characteristics to dynamically define new target database definitions, ACBLIB members, MDALIB members and DBRC definitions.

3. Source objects can be copied while they are still online with the potential to create a fuzzy copy. If a fuzzy copy is created, IMS Cloning Tool offers an option to apply logs to the cloned databases to create transitionally consistent database clones. Optionally, IMS Cloning Tool issues an IMS DBR, DBD or Quiesce command against the selected source databases and indexes to create a consistency point for the source databases.
4. IMS Cloning Tool invokes DFSMSdss Copy to issue IBM FlashCopy or SnapShot commands, or EMC Timefinder/Snap commands to copy the data sets representing the databases to be refreshed.
5. If databases were left online during the cloning process this step is skipped. For databases that were Quiesced, DBDed or DBRed, IMS Cloning Tool automatically executes the commands to resume access to or start the source IMS databases.
6. If the source databases are kept online during the cloning process and update transactions were in process at the time of the cloning process, a fuzzy cloned database is created. IMS Cloning Tool can then optionally apply logs from the source environment to the cloned databases to create transitionally consistent database clones.
7. If the databases are registered to Database Recovery Control (DBRC), IMS Cloning Tool updates the target IMS RECONs to prevent invalid recovery processes, and sets the image copy recommended flag on for the refreshed databases and indexes. For HALDBs, IMS Cloning Tool will update the partition and reorg numbers to match the source databases.
8. IMS Cloning Tool starts target IMS databases, and makes them available.

NOTE: IMS Cloning Tool has a data masking feature that provides you with the opportunity to protect sensitive data, (for example, credit card numbers, Social Security numbers, names, and addresses) by masking data while refreshing your IMS databases. IMS Cloning Tool automates a wide variety of masking options against source data creating a masked unload file to be loaded into the target source database.

Testing with refreshed databases in your IMS System Skeleton clone

If you have already created an IMS System Skeleton clone, you can now start refreshing databases into the cloned environment. Start by cloning your IVP databases and completing the IVP validation steps. If you relinked any sparse secondary index routines or randomizers, clone the associated databases and run those tests. Next, you may want to focus on specific applications. You could test applications that are used by several systems such as security or common table reference systems, applications that have a history of having issues during past migrations, or applications that are business critical. It is easy to clone a group of databases using a wild card in the database name keyword or selecting several databases by name to be cloned in a single IMS Cloning Tool job step.

The IMS System Skeleton environment can also be used as a starting point to test new functionality. You can select databases based on their physical attributes to be refreshed into the cloned environment. You do not have to worry about corrupting the data during testing since you can always refresh the database by cloning. Below is just a short list of IMS V12 features that can be tested in this environment.

- Full Function databases:
  - Dynamic Database Buffering. Clone an IMS database that you have a long-running job to run against. Test database buffering maintenance tasks while the database is allocated.

  - Extended Address Volume (EAV) for non-VSAM data sets. Refresh an OSAM database to test this functionality.

- Fast Path (DEDB) Databases:
  - Buffer Manager 64 bit Enhancements. Refresh DEDB databases and analyze the effects of this enhancement on any potential impact on other workload on z/OS as well as FP transaction performance.
  - Secondary Index processing. If you have applications using DEDBs that could benefit from secondary indexes, you could validate the work to implement DEDB secondary indexing as well as any performance impact. If you already have a fast path database with a secondary index via another product, you can refresh the database and do performance comparisons to test this new feature available in IMS V12.

All of these IMS V12 features can be tested in a tightly controlled, high quality, small environment that is a clone of the production environment. The IMS Cloning Tool database refresh feature allows for easy, repeatable cloning at the database level.

Evaluating the Impact of a Production Implementation

The IMS Full System clone is a replica of the entire IMS environment: all the system data sets, system definitions, databases and data. This type of test environment had eluded the industry in the past because of the time, resources and cost to produce. Now with the IMS Cloning Tool leveraging fast-replication, rapid volume reconditioning, and automating adjustments to the cloned IMS, this test environment can become a reality. Customers who cloned their entire production system have gone from a several day process to an average of 30 minutes to complete the entire cloning process.

How IMS Cloning Tool clones a full IMS system

The following figure depicts the processing steps that IMS Cloning Tool performs to clone a full IMS system and its databases:

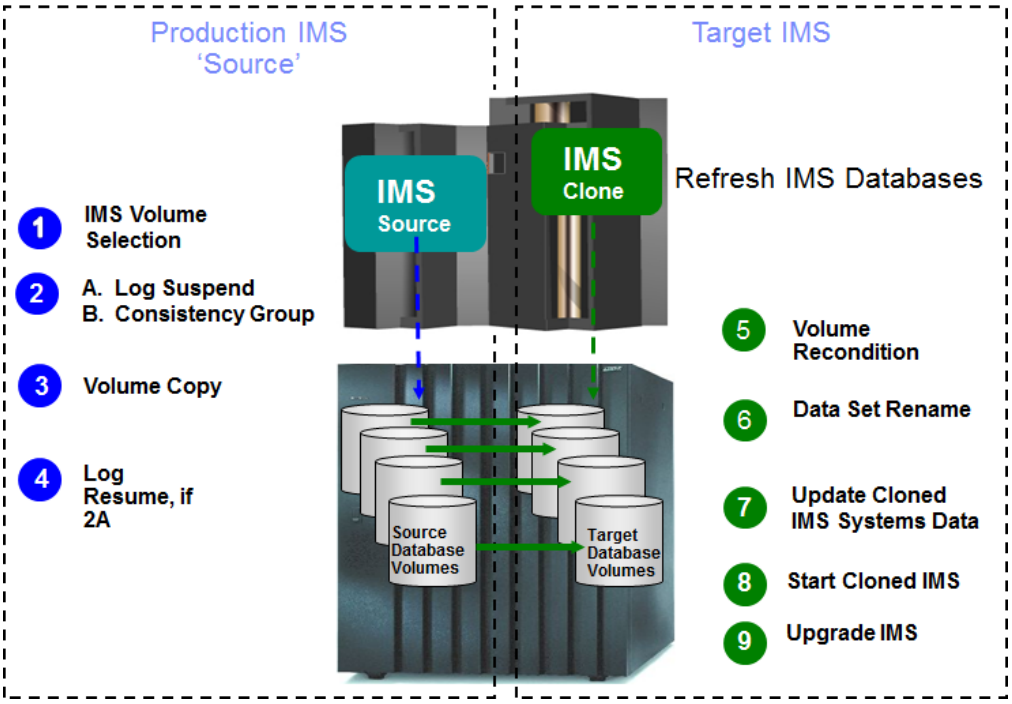

Figure 3: Cloning an IMS system and its databases with IMS Cloning Tool

Process:

1. The user identifies the storage volumes that make up the production IMS system and associated databases using specific VOLSER IDs, VOLSER masks, or SMS storage groups.
2. To create a state where the data being copied is consistent, meaning that IMS could be restarted from that point and any inflight updates could be backed out successfully while IMS is active, there are two options:

   - Option A: IMS Cloning Tool internally suspends the IMS log updates for the milliseconds it takes to perform the copy. This process is called "IMS SUSPEND" and the process ensures all production IMS data is consistent.

   - Option B: IMS Cloning Tool invokes a storage-based consistency function whereby all the volumes in a consistency group are frozen at the same time for the milliseconds it takes to perform the copy, thus ensuring a consistent copy of the production environment.
3. IMS Cloning Tool invokes DFSMSdss Copy to issue IBM FlashCopy or SnapShot commands, or EMC Timefinder/Clone commands to copy the data and back up the ICF catalogs that point to the data sets on the volumes being copied.
4. If you chose option A in step 2 (using IMS SUSPEND) above, IMS Cloning Tool internally resumes the IMS log updates immediately after the copy is performed (which takes seconds or less). This process is called "IMS RESUME".
5. If the target volumes are not already online, IMS Cloning Tool re-labels the volume internal identifier of each target volume so that they can be brought online to the same z/OS system, or a different z/OS system, without volume label conflicts.
6. The VTOC and VVDS on the target volumes are renamed to match the target VOLSER names. The data sets on the cloned volumes are renamed and re-cataloged to a new high level qualifier. Data sets are renamed to eliminate duplicate data set names from a z/OS catalog perspective and to allow the new data set names to be integrated into

the cloned IMS system. The cloned and renamed IMS data sets can then be used on the same LPAR and integrated into a different IMS system using the new data set names.

7. The cloned IMS system is updated to reflect the new data set names, VOLSERs, IMSID, etc. The cloned IMS system updates include:
    a. RECON Data Sets – The data set names, IMS subsystem IDs, and VOLSERs are updated in the following RECON records: header record, database data set records, online log records, and back-out records.
    b. Optionally, the following RECON records are updated: image copy records, change accumulation records, System Log Data Set (SLDS) records, and Recovery Log Data Set (RLDS) records if they were on volumes that were cloned.
    c. IMS.PROCLIB, IMS.JOBS, and user JCL library members are updated with the new values for IMSID, VOLSERS, and data set names.
    d. MDA (MVS Dynamic Allocation) members for databases or system data sets, such as the RECON, Online Log Data Sets (OLDS), or Write-Ahead Data Sets (WADS) are updated to reflect the new data set names.
    e. Dynamic Resource Definitions (DRD) data sets, if they are used, are updated.
    f. The default IMS SSID in the online and batch nucleuses in IMS.SDFSRESL is updated.
    g. The repository data sets, if they are used, are updated.
    h. If IMS data sharing is involved, each additional IMS data sharing member is updated.
8. The cloned IMS system is started and made available to users.
9. Upgrade the cloned IMS system to the new IMS version you wish to test with. If any IMS issues are uncovered due to the upgrade, simply correct the errors during business hours, without fear of an emergency back off at some early morning hour. Apply any necessary IMS version patches and continue testing. Migrate the new version of IMS to your production IMS with confidence

Testing in an IMS Full System clone

You can test everything in an IMS Full System clone: a complete dry run of the installation process, the new IMS software, vendor software, home-grown software, complete application testing, and complete new function testing.

An area this environment uniquely supports is load and performance testing. It is a safe, optionally temporary, full volume, complete environment. The IMS Full System clone allows detailed repeatable performance testing both before and after the IMS version upgrade. It is an environment in which you can determine the return on investment (ROI) of a new version or individual feature of that version of IMS. This information is normally very hard to measure yet very valuable in making decisions about which new features to use and when they should be used.

If you are having throughput issues, or if you are looking for throughput improvements and you have Queue Control Facility (QCF) and IMS Performance Analyzer (IMS PA), you can use QCF to capture an interval of work from your production environment. Because you have a clone of the entire environment, you can play back the work load in the cloned environment and take your time to analyze it. IMS PA can be used to show a performance base line from your current production environment and compare that to the same workload run in a production replica with the new IMS version. You can re-clone your environment or specific databases, tweak them for possible performance improvements,

re-run your QCF captured interval of work, and perform the analysis. Repeat the process as often as need to get the optimum performance from your IMS systems.

The IMS Full System clone has many uses outside of the migration process. A Full System clone can be used to offload business reporting, quickly create, duplicate and refresh training, testing, and maintenance environments, and since it is an exact replica of your production environment, you can debug issues that are hard to recreate from your production environment.

Summary

Leveraging these different types of cloning operations can provide cost effective, quality testing to help your organization smooth the IMS migration path and reduce your risk of impacting critical applications.

IMS Cloning Tool can create an exact clone of your production IMS system, production IMS databases or entire production IMS environments fast and effectively. Testing IMS migrations using current data and current environments lowers the risk of affecting sensitive IMS business applications when you go live.

IBM United States of America

Produced in the United States of America

IBM Director of Licensing
IBM Corporation
4205 South Miami Boulevard
Research Triangle Park, NC 27709 U.S.A.