



EMC and Mainstar Solutions for SAP Systems Using DB2 for z/OS and EMC Symmetrix V-Max

Applied Technology

Abstract

This white paper describes SAP on z/OS usage characteristics and the operational challenges they present. SAP backup, recovery, disaster recovery, and cloning automation solutions are described, showing how EMC and Mainstar storage integrated products can simplify SAP on z/OS administration challenges. Examples show how EMC and Mainstar products perform and simplify SAP backup, recovery, disaster recovery, and cloning operations. This white paper provides EMC and Mainstar SAP solution support validation for SAP environments using EMC[®] Symmetrix[®] V-Max[™].

April 2009

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h5985.1

Table of Contents

Executive summary	4
Introduction	4
About the authors.....	4
Audience	4
SAP systems on z/OS usage characteristics	5
Backup and recovery challenges.....	5
Application and data cloning challenges.....	6
Solutions for SAP systems using DB2 for z/OS.....	7
Rocket Backup and Recovery for DB2	7
Creating a DB2 system backup.....	10
Performing DB2 system and application recovery	12
Implementing a tape-based DB2 disaster restart methodology	14
Creating and refreshing SAP test environments.....	16
Creating a test SAP application environment from a production system	16
Creating a test SAP environment from a system backup	18
Refreshing DB2 table and index spaces in a test SAP system	19
DB2 solution validation on EMC Symmetrix V-Max	21
Conclusion	22

Executive summary

SAP is a business-critical application that is used to integrate and automate diverse business management processes. Many SAP implementations use DB2 for z/OS as the database server to manage SAP application data. Typically, DB2 for z/OS is chosen as a database server due to its robust database management capabilities, high performance potential, high availability characteristics, and ability to scale to fit application requirements.

Dedicated DB2 for z/OS systems are used to support SAP applications. These DB2 systems have large numbers of correlated DB2 tables, table spaces, indexes, and index spaces that are used by the SAP application. SAP implementations can require 50,000 to 80,000 DB2 objects to be managed collectively within a given DB2 system to support SAP applications. Some DB2 objects may be dynamically created and deleted while sensitive SAP application data relationships, which are inherent in its databases, are maintained by the SAP application and by the DB2 system collectively. The dynamic interrelation of many DB2 objects requires all data to be managed as a unit and requires that the data be backed up and recovered as a unit otherwise data relationships become broken, jeopardizing the integrity of the SAP business applications.

SAP is a critical business application where high availability solutions are required to reduce the business effect caused by normal database maintenance as well as to minimize application outages caused by infrastructure failures. Fast backup, recover, and disaster recovery facilities are required to minimize business downtime during normal application operations and to reduce the impact of logical errors or the impact of a catastrophic local site failure. Low impact and fast DB2 system cloning facilities are required to reduce the administration and operational impact of creating test SAP application environments.

Introduction

This white paper describes SAP usage characteristics by laying out the challenges of backup and recovery, and application and data cloning for SAP systems running on z/OS. It details how EMC and Mainstar products simplify DB2 backup, recovery, disaster recovery, and cloning functions for SAP on z/OS environments.

The products used to achieve the results were EMC® TimeFinder®, Rocket Backup and Recovery for DB2, Mainstar Volume Clone and Rename, and Mainstar Fast Table Space Refresh for DB2. Examples are provided to show how backup and recovery operations are performed using TimeFinder integration. DB2 cloning examples are provided to show how TimeFinder can be used during the SAP cloning process in order to create and refresh test SAP landscapes.

About the authors

This paper was written by Rocket Software Inc. and EMC personnel. Ron Hauptert is a Senior Technologist at Rocket Software and responsible for creating storage integrated database and application management products and solutions. He has worked with various mainframe and open systems database technologies and implementation for over 30 years. Paul Pendle works in EMC Symmetrix® Partner Engineering and is responsible for IBM database solutions at EMC. Dawn Johnsey works in EMC Symmetrix Partner Engineering and is responsible for SAP integrated solutions at EMC. Mainstar is a subsidiary of Rocket Software.

Audience

The intended audience for this white paper is anyone working as a SAP basis administrator, DB2 system and database administrators, storage management personnel, disaster recovery personnel, and IT management.

SAP systems on z/OS usage characteristics

Backup and recovery challenges

Traditional DB2 for z/OS backup and recovery procedures do not work effectively for SAP systems using z/OS. The large number of integrated DB2 objects used to support SAP applications require strict backup and recovery coordination to ensure all data relationships are maintained through a recovery process. Failure to recover all related objects to exactly the same point results in compromised business data in the database server, which in turn negatively impacts the SAP business application.

A DB2 system backup and recovery methodology can simplify DB2 backup and recovery procedures and ensure that dynamically managed DB2 objects are recovered and all data relationships are preserved during a recovery process. A DB2 system backup and recovery methodology is one where the entire DB2 system is backed up as a unit and the backup can be used to recover the DB2 system as a unit. All data relationships are implicitly maintained during the recovery process, guaranteeing that data integrity is preserved for the SAP application. A DB2 system backup and recovery methodology is much faster to perform than using traditional DB2 image copy approaches and the same DB2 system backup and recovery procedures can be used for local site recovery as well as for offsite disaster recovery purposes.

Most DB2 system backup solutions use some form of storage-based fast-replication that is integrated with DB2 backup and recovery facilities. These backup and recovery solutions leverage modern storage processor capabilities and fast-replication products to perform backup and restore operations on behalf of the DB2 system. EMC fast-replication products used to support a DB2 system backup approach include: EMC TimeFinder/Mirror, EMC TimeFinder/Clone, EMC TimeFinder/Snap, and EMC Compatible Native Flash for Mainframe.

The advantages of using a DB2 system backup methodology that uses storage-based fast-replication include the following:

- Full DB2 system backups can be completed in seconds or less.
- There is little or no impact to applications while the backup is performed.
- Point-in-time restartable backup copies of an entire DB2 system are created.
- No host CPU and I/O resources are used to create the backup.
- Incremental fast-copy facilities can be used to reduce backup storage requirements and to reduce the storage processor resources used in the copy process.
- System backups can be used for multiple purposes:
 - Local site system recovery
 - Local site object recovery
 - As a source for creating a backup used for offsite disaster recovery
 - As a source for performing DB2 system cloning operations
- System backups can be archived using an independent tape copy process to reduce disk storage utilization while preserving the backups created in multiple backup cycles.

The advantages of using a DB2 system backup to perform a DB2 system recovery operation include the following:

- The entire DB2 system can be restored from disk instantaneously when storage-based fast-replication is used to restore the data.
- Parallel recovery can be performed. That is, DB2 log apply functions can be performed in parallel to the data restoration process, thus reducing overall DB2 system recovery time.
- The DB2 recovery process is done by processing the DB2 log one time to recover the entire DB2 system.
- The DB2 system is recovered as a unit so all data relationships are preserved during the restore and recovery processes, thus ensuring the SAP application's data integrity is preserved.

The advantages of using a DB2 system backup for offsite disaster recovery include the following:

- A DB2 system backup can be easily copied to tape and transported to a disaster recovery site.
- Traditional DB2 disaster recovery procedures are streamlined by using a tape-based disaster restart methodology. A disaster restart methodology is one where tapes containing a system backup are loaded at a disaster recovery site and DB2 is restarted. Recovery procedures are implicitly performed during the normal DB2 restart operation. Traditional DB2 disaster recovery procedures are not used. The normal DB2 restart process using a system backup transforms the system backup data state into a transactionally consistent data state. The DB2 system is ready to accept new application work after the restart process is complete. A disaster restart methodology is a fast and effective way to resume SAP processing at a disaster recovery site and minimize recovery time objectives and SAP application downtime.

Application and data cloning challenges

SAP application environments and their supporting data infrastructures are cloned for various application and operational purposes. Some SAP implementations that use DB2 for z/OS clone their supporting DB2 for z/OS system multiple times. Some reasons to clone SAP application environments and their supporting DB2 systems include:

- Creation of SAP application test environments
- Prototyping new SAP application functionality
- Testing new SAP release migrations before the migrations are performed on production systems
- Testing new DB2 for z/OS release and migration procedures
- Testing new system management and operational procedures

Cloning DB2 systems that support SAP application environments can take a long time when using host-based DASD volume or dataset copy procedures. Furthermore, when cloning a DB2 system in the same LPAR or in a different LPAR with shared disk, all dataset names must be changed so they can be integrated into the cloned DB2 system. The dataset rename process can be long when 50-80,000 datasets need to be renamed during the cloning process. Typically, cloning methodologies that use host-based volume or datasets copy procedures can take days or weeks to perform the DB2 system cloning operation. The steps required to clone a DB2 for z/OS system include the following:

- A data consistency mechanism must be chosen and used so that data relationships are preserved during the copy process and remain intact when the cloned DB2 system is started. Many DB2 system cloning methodologies require the SAP application to be down and its supporting DB2 system stopped to get a consistent DB2 clone copy.
- All data must be in a consistent state and copied from a source DB2 system to a target DB2 system.
- Dataset names must be changed during the copy process so the target datasets can be accommodated into the target DB2 system.
- The target DB2 system must be adjusted to accommodate the copied datasets and source DB2 system's recovery structures.
- Target SAP system identifiers must be adjusted to associate the cloned SAP application environment with its cloned DB2 system.

DB2 system cloning methodologies that leverage EMC Symmetrix V-Max™, Symmetrix DMX™, or Symmetrix TimeFinder facilities copy the data quickly. The copy process can be performed while the DB2 system is running. These cloning methodologies use TimeFinder to copy all volumes that make up a DB2 system to a target set of volumes. The volume copy process is fully contained and managed within the EMC storage processor. Once the DB2 volumes are copied, procedures are employed to condition the volumes and to rename all copied DB2 datasets so the volumes and their new dataset names can be accommodated into a target DB2 system. TimeFinder family products that are used to support a storage-based DB2 system cloning methodology include TimeFinder/Mirror and TimeFinder/Clone. The advantages of using a storage-based DB2 system cloning methodology include the following:

-
- DB2 systems can be copied instantaneously and while SAP applications are running.
 - There is little or no impact to the SAP application while the DB2 system is being copied.
 - Application data consistency is maintained using a DB2 Set Log Suspend or by using EMC consistency technology, which will eliminate application downtime.
 - The DB2 system cloning process is quick and easy to perform.
 - No host CPU and I/O resources are used for the data copy process.
 - Volume reconditioning and dataset renaming facilities must be employed to allow the volumes to be brought online and the datasets to be integrated into a cloned DB2 system.

Solutions for SAP systems using DB2 for z/OS

EMC and Mainstar offer products that simplify DB2 administration functions for SAP on z/OS environments. These products can be used independently or together to create DB2 system backup, recovery, disaster recovery, and DB2 cloning solutions on behalf of the SAP application. Three products that integrate EMC TimeFinder facilities and can be used to create SAP on z/OS solutions include:

- Rocket Backup and Recovery for DB2, which can be purchased from EMC through the EMC Select partner product program
- Volume Clone and Rename, which can be purchased from Mainstar
- Fast Table Refresh for DB2, which can be purchased from Mainstar

Rocket Backup and Recovery for DB2

Rocket Backup and Recovery for DB2 (RBR) is a DB2 for z/OS system backup and recovery solution designed specifically to simplify and speed up backup and recovery operations for SAP application environments that use the mainframe as their foundation. It is tightly integrated with EMC TimeFinder and it speeds up DB2 backup and recovery processing by using TimeFinder/Mirror, TimeFinder/Clone, or TimeFinder/Snap facilities. RBR automates storage management functions for DBAs so they can leverage Symmetrix V-Max, Symmetrix DMX, and Symmetrix fast replication without having to create or use storage-specific commands, scripts, or JCL.

RBR provides a fast and easy-to-use implementation of a DB2 system backup and recovery methodology. It reduces backup windows by leveraging TimeFinder fast-replication products such that backups of multi-terabyte databases can be performed in seconds. It simplifies backup and recovery methodologies by allowing both full-system and object-level recoveries to be performed from a common system backup. It provides DB2 system backup and recovery support for complex applications like SAP, where all of the application's data must be backed up, restored, and recovered as a unit.

RBR features include:

- **DB2 System Backup Configuration and Management.** These facilities discover the DB2 system and recommend layout and configuration changes such that the DB2 system is set up appropriately to accommodate a system backup and recovery methodology. Optionally, RBR will move datasets and modify the DB2 system configuration to accommodate a system backup and restore methodology.
- **DB2 System Backup and Recovery.** RBR uses profiles to define system backup and restore procedures that allow users to back up and restore DB2 systems as a unit. Object recovery profiles are used to define DB2 object and application recovery strategies using a system backup. The backup and restore procedures leverage storage-based fast replication to perform the backup and restore functions.
- **DB2 Object-Level Recovery.** Object-level recovery allows RBR users to recover DB2 objects or groups of related objects that represent applications from a system backup. Object-level recovery leverages storage-based dataset fast-replication facilities. The use of EMC Dataset Snap facilities allows object recovery to be performed in parallel to the object restore process, thus significantly reducing the overall recovery time. Object recoveries that traditionally have taken many hours can be performed in minutes or seconds using RBR.

-
- **Metadata Repository.** RBR provides a comprehensive metadata repository to record backup information like backup time, backup type, log byte addresses, volumes used for the backup, etc. Reports can be generated to monitor backup methods and operations, storage volume usage, system backup volume usage, archived backups, etc.
 - **Backup Validation.** RBR provides extensive backup validation to ensure the backup contains all DB2 files and catalog structures required for a successful recovery.
 - **EMC TimeFinder Support.** RBR has extensive EMC TimeFinder integration. It supports TimeFinder/Mirror, TimeFinder/Clone, and EMC Dataset Snap facilities. EMC virtual device support will be available in a pending release. Mixed vendor storage environments are supported when using EMC's Compatible Native Flash for Mainframe as the fast-replication method. Backups of DB2 and non-DB2 data can be accommodated by including the non-DB2 data volumes in a DB2 System Backup profile and using an appropriate TimeFinder fast-replication function along with associated EMC consistency technology to create a consistent backup.
 - **Tape Offload Support.** RBR for DB2 provides tape offload support to automate copying a backup from disk to tape. Backups created on disk can be copied to tape using DFSMSdss or FDR so the backup disk volume pool can be reused. A subsequent DB2 system restore operation will restore the backup from disk or tape depending on system backup availability and recovery scope. RBR will use the most appropriate backup for object-level recovery and will restore the object from disk, tape, or a previous image copy depending on which backup provides the most expedient recovery process.
 - **Disaster Recovery.** RBR provides disaster recovery support by transforming traditional DB2 disaster recovery procedures into a tape-based disaster restart methodology. System backups can be tagged for offsite transport to a disaster recovery site during the offload process. The tape-based disaster restart methodology loads the system backup tapes and restarts DB2 at the disaster recovery site. The DB2 restart process transforms the system backup into a transactionally consistent DB2 system that is ready to accept work. Optionally, DB2 logs generated after the system backup can be copied and subsequently applied at the disaster recovery site during the DB2 restart process. Using RBR to implement a DB2 tape-based disaster restart methodology simplifies disaster recovery procedures and reduces recovery time objectives.
 - **DB2 Version Support.** RBR supports DB2 versions 7, 8, and 9 using data-sharing and non-data-sharing modes of operation.

RBR navigation is done using a sophisticated ISPF interface. The menu interface provides quick and easy access to its backup and recovery functions. Figure 1 shows the main ISPF panel for RBR.

```
RBR$MAIN V2R1 ----- Rocket System Backup and Recovery for DB2 on z/OS -----
Option ===> _____

User: PDBOBA - RBR

-----

B. System Backup Profiles

R. System Restore and Offload

O. Object Recovery Profiles

D. Disaster Recovery

M. DB2 Subsystem Analysis and Configuration

S. Product Setup

X. Exit

Enter END command to return to ISPF.
```

Figure 1. RBR for DB2 main ISPF panel interface

A functional description of each menu option is provided below.

- **Product Setup** – Allows users to establish ISPF profile defaults and to register DB2 systems they want to back up and restore. Default specifications can be provided for system backup profiles. The default specifications will be used each time a new backup profile is created. Some default backup profile specifications include full or data-only backups, the number of online and offline backup generations to maintain, the type of EMC TimeFinder fast-replication method to use, and whether EMC consistency technology or a DB2 Set log Suspend command is used to enforce data consistency during a backup operation.
- **DB2 Subsystem Analysis and Configuration** – This option invokes the DB2 system configuration discovery facilities. This facility discovers the DB2 system and displays its current DASD configuration. It also recommends DASD layout and configuration changes so that the DB2 system is set up appropriately to accommodate a DB2 system backup and restore methodology. Optionally, ISPF screens are available to assist moving datasets and modify the DB2 system so system backup and restore functions are easily accommodated.
- **System Backup Profiles** – Backup Profiles allows users to define specific backup profiles that can be executed to perform a backup of a specified DB2 system. It allows users to specify the type of backup to be performed and the association of target volumes to be used for the backup. RBR automates the target volume selection to allow easy use of EMC TimeFinder features.
- **System Restore and Offload** – System restore allows users to restore and recover the DB2 system. DB2 systems can be recovered to a specified backup, to a time stamp, or to a specified RBA value. The restore process leverages EMC TimeFinder/BCV or TimeFinder/Clone functions to restore the backup and DB2 recovery facilities are invoked to recover the DB2 system. The RBR recovery management infrastructure coordinates the overall recovery process. System backups can be archived or offloaded to tape using this selection.
- **Object Recovery Profiles** – Object recovery profiles allow users to recover objects or groups of related objects as a unit. Object recovery profiles can be used to coordinate the recovery of all objects representing a given application. RBR will choose the most appropriate backup to restore and it will

apply log records in parallel to the restoration process if possible. Recovery profiles can be created proactively to simplify and speed up application recovery scenarios.

- **Disaster Recovery** - Disaster recovery profiles allow users to define specific disaster recovery options into a saved profile. The profiles can be executed at the local site to ready all the necessary resources to restore a DB2 system at a remote site. Users can optionally direct RBR to make copies of local archive logs to be shipped to the disaster recovery site. All the JCL and job cards needed to perform the DB2 restore and recovery at the remote site are generated and placed in a partitioned dataset that can be written to tape or to remote DASD for transport to a disaster recovery site. Users need only submit these jobs at the disaster recovery site to perform the DB2 restore and recovery operation.
- **Exit** – This option is used to exit the RBR product.

Creating a DB2 system backup

Implementing a system backup methodology requires the DB2 system and storage volume mapping to be identified and DB2 recovery structures to be isolated from DB2 catalog, directory, and application data structures. ICF User Catalog placement must be coordinated with the data they represent. The DB2 discovery and system backup configuration process is done using the DB2 Subsystem Analysis and Configuration panels. Figure 2 depicts a DB2 system that is appropriately configured to use a system backup methodology.

SAP Application Environment Using DB2 for z/OS

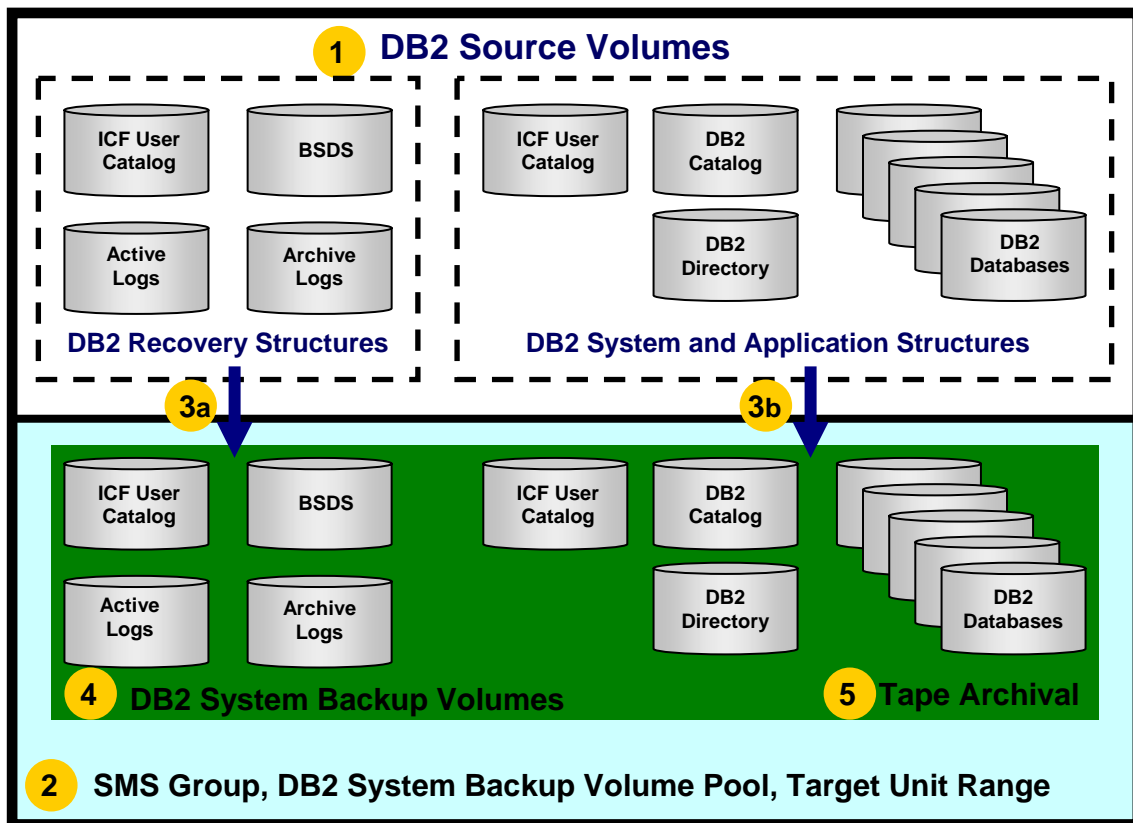


Figure 2. DB2 system configuration for a system backup methodology

The following steps, referred to in Figure 2, are performed by RBR to implement a system backup methodology.

1. The DB2 system is discovered, analyzed, and configured to accommodate a system backup methodology.

2. A System Backup Profile is created. During the profile creation process, DASD volumes used for the system backup are selected. The target backup volumes can be selected from a range of target units, a DB2 System Backup volume pool, or a SMS Storage Group.
3. A system backup is created using the information provided in the backup profile to drive an appropriate TimeFinder process. Backups executed using the “full” option, perform a backup of all volumes of the DB2 system (3a, 3b). Backups processed using the “data only” option back up only the data volumes (3b).
4. The backup is recorded in the RBR metadata repository. Recorded information includes backup type, time, associated RBA values, target volumes used, and dataset name mappings for use in object-level recoveries.
5. Online disk volumes can be archived to tape automatically after each system backup or on demand using the System Restore and Offload ISPF menus.

Figure 3 shows RBR creating a system backup for an SAP environment using DB2 for z/OS. RBR performs the system backup process using information specified in System Backup Profiles and using DASD volume relationships specified in the Backup Volume Selection (step 3) in Figure 3.

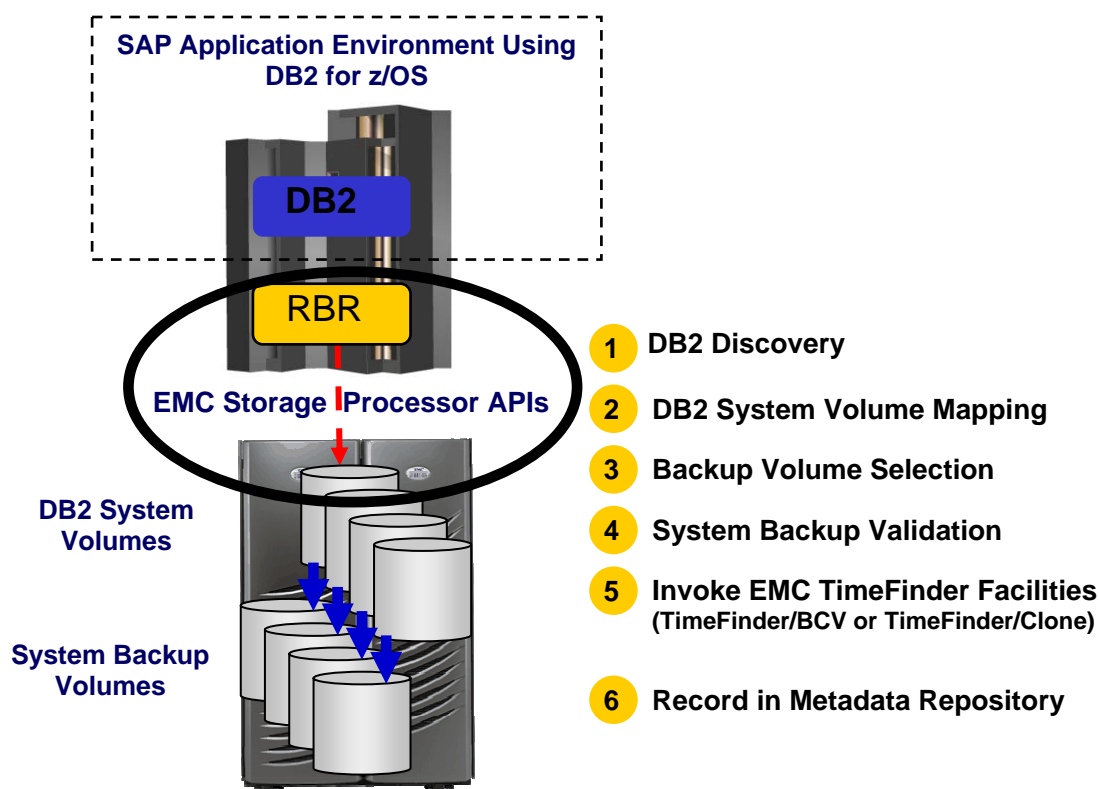


Figure 3. Creating a system backup for an SAP environment using DB2 for z/OS

The following processing steps, referred to in Figure 3, are used to perform a DB2 system backup operation using RBR:

1. The DB2 system is discovered and analyzed. Source volumes are identified and dataset and catalog placement is analyzed to determine if the DB2 layout can accommodate a system backup methodology. Optionally, RBR will move DB2 datasets and catalogs appropriately so that the DB2 system can accommodate a system backup and recovery methodology.

-
2. All source DB2 volumes are identified and analyzed to determine if non-DB2 data resides on the volumes. Warning messages are issued when non-DB2 data is included in a DB2 system backup.
 3. Source DB2 volumes are mapped to target volumes that will contain the system backup. The source-to-target volume mapping is done during the backup profile creation process.
 4. RBR performs system backup validation checks before each backup to ensure the backup is complete and can be used for a successful restore operation.
 5. RBR invokes the selected TimeFinder function. The backup is performed in the Symmetrix V-Max, Symmetrix DMX, or Symmetrix storage processor without using host CPU and I/O resources. The system backup is completed from a DB2 system and application perspective when the TimeFinder command is issued. Typically, full DB2 system backups complete in seconds. Data consistency functions are employed to ensure the backup data state can be used for recovery purposes. Data consistency is ensured using DB2 Set Log Suspend or an appropriate EMC consistency technology function. Using EMC consistency technology eliminates application downtime.
 6. When the backup is complete, information about the backup is recorded in the RBR metadata repository. The metadata information includes an inventory of where each DB2 pageset resides on the backup volumes. This information is used to restore individual datasets when performing a DB2 object recovery operation.

Performing DB2 system and application recovery

RBR automates DB2 system or object-level recovery from a system backup. System recovery is performed using the System Restore and Offload ISPF panels and object recovery is performed using the Object Recovery Profiles panel interface shown in Figure 1 on page 9.

System recovery can be a full or data-only recovery. When a full system recovery is performed RBR will restore all data and log volumes and no DB2 logs will be applied to the restored volumes. An example of a full system recovery is shown in Figure 2 by performing the reverse copy of steps 3a and 3b. This is shown in steps 3a and 3b in Figure 2 on page 10. Performing a full system recovery is analogous to performing a “recover to copy” when using image copies for recovery.

Data-only recovery directs RBR to restore only the data volumes from the system backup (step 3b in Figure 2). This leaves the DB2 log datasets unaffected and available to use for DB2’s log apply phase of recovery. RBR then executes a DB2 Restore System utility with the “LOG ONLY” keyword. This utility applies log records for the entire DB2 system with one pass (or reading) of the log files. This process can take place while the storage processor is restoring the data volumes in the background. The recovery process can use a RBA, LRSN, or time stamp to establish the end of the recovery process. Figure 4 depicts the use of a system backup to recover an SAP application utilizing DB2 for z/OS.

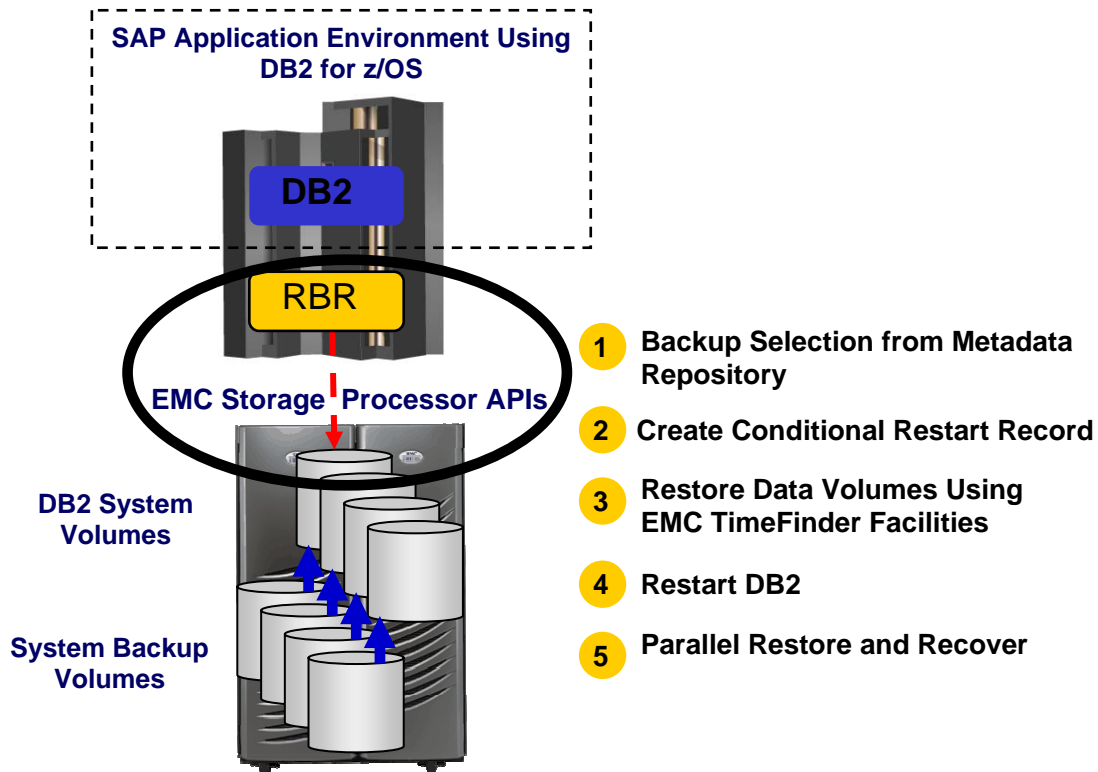


Figure 4. RBR used to recover an SAP on z/OS application environment

The following processing steps, referred to in Figure 4, describe a DB2 system recovery process using a RBR for DB2-generated system backup:

1. The appropriate system backup metadata is selected from the metadata repository using the recovery criteria specified in the Restore System Display from the System Restore and Offload panels.
2. A DB2 Conditional Restart record is created specifying at what point to stop the DB2 log apply recovery process. The information used to generate the conditional restart record is specified in the Restore Options panel.
3. RBR for DB2 invokes an appropriate TimeFinder function to perform the system restore process. The restore operation is performed in the Symmetrix V-Max, Symmetrix DMX, or Symmetrix storage processor without using host CPU and I/O resources.
4. The DB2 conditional restart and subsequent log apply process begins after the volume restore process has started.
5. The DB2 restart and roll forward log apply process is performed in parallel with the volume restore process to minimize overall DB2 recovery time.

DB2 object-level recovery is done by creating an Object Recovery Profile. Object Recovery Profiles describe the methods and options used to recover objects or groups of related objects representing applications. Recovery profiles are created and stored in the metadata repository and can be recalled for use when an application recovery is required. The Object Recovery Profile specifies a recovery point, which could include any of the following:

- Any specific RBA or LRSN
- Last application quiesce point

-
- Last backup copy
 - To current (end of log)

The Object Recovery Profile also describes which recovery resources to use by specifying whether to recover from disk, an offloaded tape, an image copy, or all available resources.

System and object recoveries can restore data and perform DB2 recovery operations in parallel when recovering from a system backup that resides on disk. RBR will invoke an appropriate EMC volume level TimeFinder operation or a dataset snap operation in the storage processor to restore data. While the data is flowing from the backup volumes to the source DB2 system to restore the data, DB2 logs are applied to roll the table or index space forward in parallel with the data restoration process. RBR's parallel restore and recovery process significantly reduces overall recovery time and increases application availability.

Implementing a tape-based DB2 disaster restart methodology

A tape-based disaster restart methodology is one where a restartable DB2 system is captured on disk and transferred to a disaster recovery site using tape as a transport mechanism. Traditional image copy based DB2 disaster recovery procedures are not used at the disaster recovery site. Instead, tapes containing a restartable DB2 system are loaded on to disk and DB2 is restarted at the disaster recovery site. The disaster recovery exercise is complete when the DB2 restart process completes.

RBR creates a restartable DB2 system while creating a system backup. RBR coordinates a DB2 Set Log Suspend function or an EMC consistency function with an appropriate TimeFinder/Mirror or TimeFinder/Clone operation to create a system backup that is dependent-write consistent. A dependent-write consistent data state is identical to a DB2 system that has been exposed to a power failure. When a DB2 system is restarted using a system backup, the dependent-write consistent data state that is inherent in the system backup is transformed to a transactionally consistent data state by the DB2 restart process. Once the restart process is complete, the DB2 system is available for SAP applications.

RBR has facilities to pre-process archive logs as they are copied to tape for disaster recovery purposes. The pre-processing creates a partitioned dataset with the information required to build a new BSDS and DB2 Conditional Restart Record at the disaster recovery site. The disaster recovery procedures exercised at the disaster site restore the DB2 system using the last offsite system backup, restore the rebuilt BSDS, replace the DB2 Conditional Restart Record, and then perform a Restore System Log Only operation. Figure 5 shows the steps required to transform traditional DB2 disaster recovery procedures into a tape-based disaster restart solution.

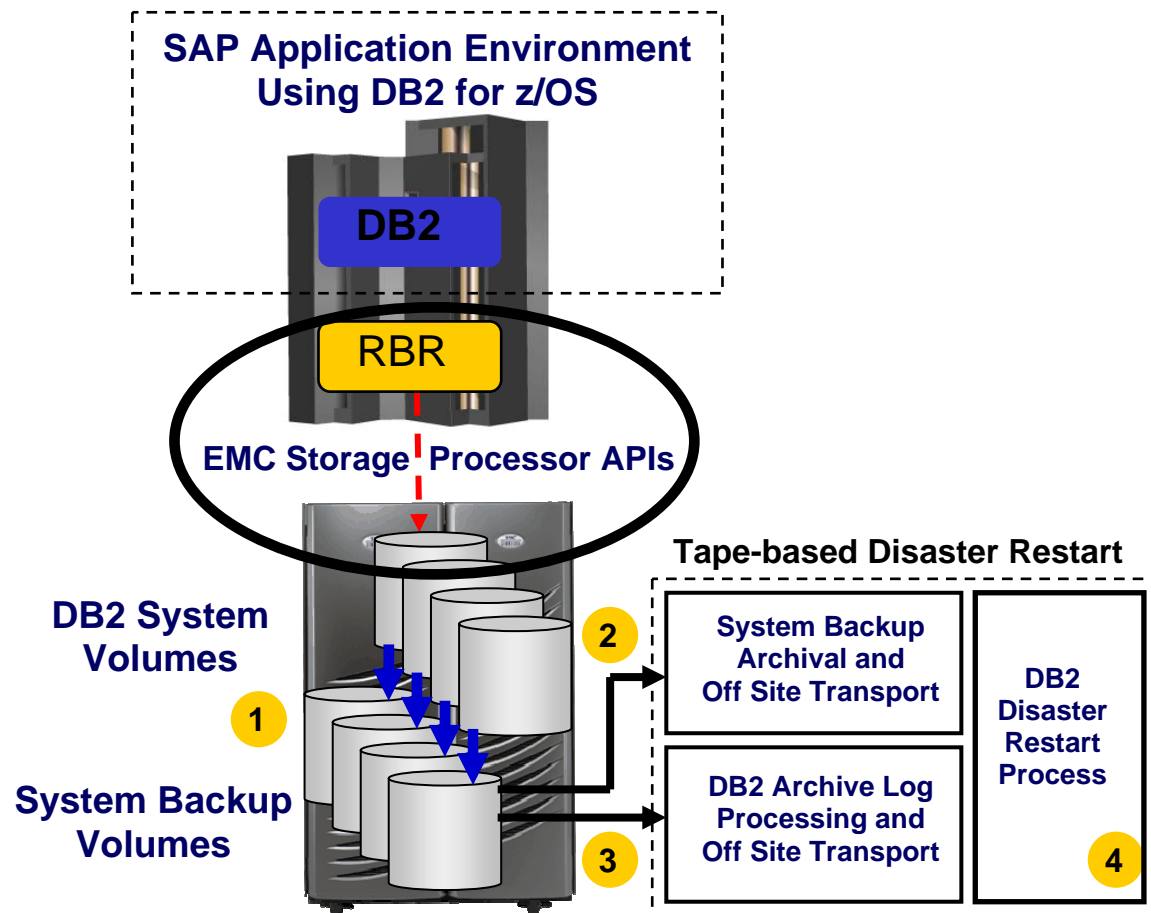


Figure 5. RBR used to implement a tape-based disaster restart methodology

The following processing steps, referred to in Figure 5, are used to implement a DB2 tape-based disaster restart solution:

1. RBR creates a system backup that has a dependent-write consistent data state.
2. The system backup is archived to one or more tape copies and one of the archive tape copies is targeted to be transported to the disaster recovery site. Tape archival options and offsite specification are provided in the System Restore and Offload panels (see Figure 1).
3. A Disaster Recovery Profile is created that holds all the options for executing a disaster recovery procedure at a remote site.
4. When the Disaster Recovery Profile is executed, DB2 archive logs are optionally copied to tape or an existing archive log is tagged to be taken offsite for disaster recovery processing. Information on how to perform the disaster recovery operation and associated recovery job JCL to perform the disaster recovery process is placed in a partitioned dataset (PDS) to be sent to the disaster recovery site. The recovery jobs perform the following functions at the disaster recovery site:
 - Restore the RBR metadata repository.
 - Restore all the volumes of the DB2 system from the last offsite system backup available at the disaster recovery site.
 - Create a new BSDS that has been conditioned for the disaster recovery site and place a DB2 system point-in-time restart record in the BSDS.

-
- Optionally copy the tape-based archive logs to DASD to speed the recovery process.
 - Perform a roll-forward process to bring DB2 to a transactionally consistent data state to the point of the last available DB2 log dataset at the disaster recovery site.
 - DB2 is ready to accept new work after the restart process is complete

Creating and refreshing SAP test environments

The Volume Clone and Rename (VCR) product is the foundation for cloning production SAP systems into test SAP application environments. It automates the DB2 for z/OS cloning process such that test SAP application environments can use data that is cloned from a production SAP system. VCR can use EMC TimeFinder to copy the DB2 data quickly and effectively without using host CPU or I/O resources and it performs fast and effective volume reconditioning and dataset naming operations to allow the data to be used by a cloned DB2 system on the same or another shared disk LPAR. The DB2 cloning automation process can be performed while the production SAP system is running or stopped and it supports data sharing and non-data sharing DB2 environments.

Fast Table Refresh for DB2 (FTR) is a selectable option for the VCR product. FTR provides facilities to clone or refresh table or index spaces within or across DB2 systems. It can be used to refresh DB2 table and index spaces from a production DB2 system into a test SAP environment. FTR is particularly useful when only certain DB2 objects need to be refreshed from a production SAP system into a test SAP environment. FTR can save time and processing resources by refreshing only those objects that need to be cloned instead of cloning a complete DB2 system again.

Creating a test SAP application environment from a production system

Typically, test SAP application environments are created using data from a production SAP system. VCR facilitates SAP test environment creation by automating the production SAP DB2 system cloning process. The target DB2 system must be defined prior to using the system cloning automation process. Figure 6 shows the steps VCR uses to clone a production SAP application environment to create a test SAP application environment. In Figure 6, processing steps shown in a blue background are performed on the production data while processing steps shown in a green background are performed on the cloned data.

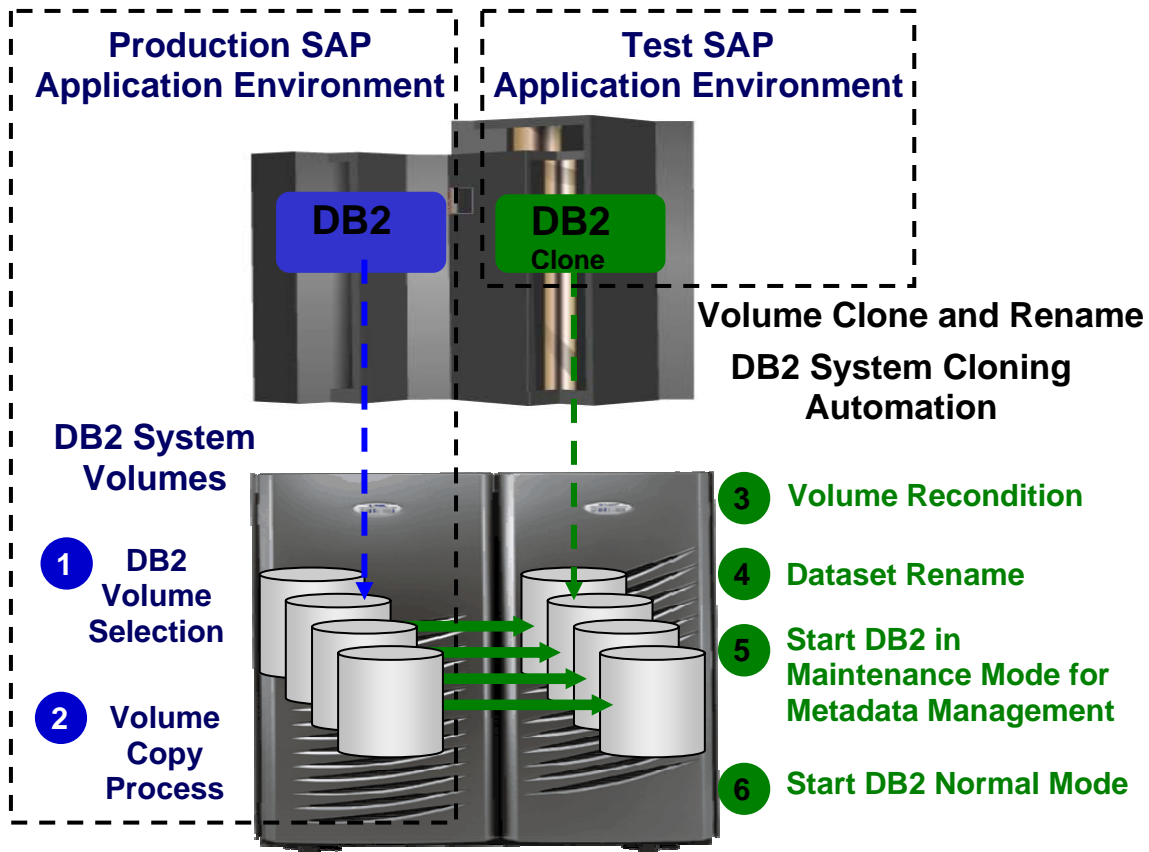


Figure 6. Creating a test SAP application environment from a production system using volume clone and rename

1. The DASD volumes that make up the production SAP DB2 system are identified by using specific VOLSER IDs, a VOLSER mask, or a SMS storage group.
2. An appropriate TimeFinder fast-replication process is performed before the VCR cloning automation is invoked and a list of copied storage volumes is passed to VCR for use in later processing steps.
3. Volume internal identifiers are relabeled so they can be brought online to the same or different z/OS system, without volume label conflicts, if they are not already online.
4. The datasets on the cloned volumes are renamed and recataloged to a new high level qualifier. Datasets are renamed to eliminate duplicate dataset names from a z/OS catalog perspective and to allow the new dataset names to be integrated into the cloned DB2 system. The cloned DB2 datasets can then be used on the same LPAR but integrated into a different DB2 system using different dataset names.
5. The cloned DB2 system is restarted in MAINTENANCE mode to allow the DB2 system catalog to be updated by the cloning automation process. When DB2 is started in MAINTENANCE mode, in-flight transactions are deferred and will not be backed out. The cloned DB2 system's directory, catalog, and BSDS are updated. The cloned data created during the copy process reflects DB2 production data and it must be changed to reflect the data in a cloned DB2 system such as the new dataset names, VOLSERs, and STOGRPs. Once all the metadata management processing is complete, in-flight transactions are resolved and the cloned DB2 system is stopped.
6. The cloned DB2 system is restarted using normal ZPARM parameters and is made available to users.

Creating a test SAP environment from a system backup

Test SAP environments can be created using a RBR generated system backup as input to VCR's DB2 cloning automation process described previously. When cloning DB2 from a RBR-generated system backup, VCR determines the volumes to be copied by accessing RBR's metadata repository. Once the system backup volumes are determined, they are passed to VCR to perform the DB2 cloning process described previously. Figure 7 depicts cloning a DB2 system from a RBR for DB2-generated system backup. Steps shown with a gold background color are performed by RBR for DB2, steps shown in blue are performed by the fast-replication process, and steps shown in green are performed by VCR.

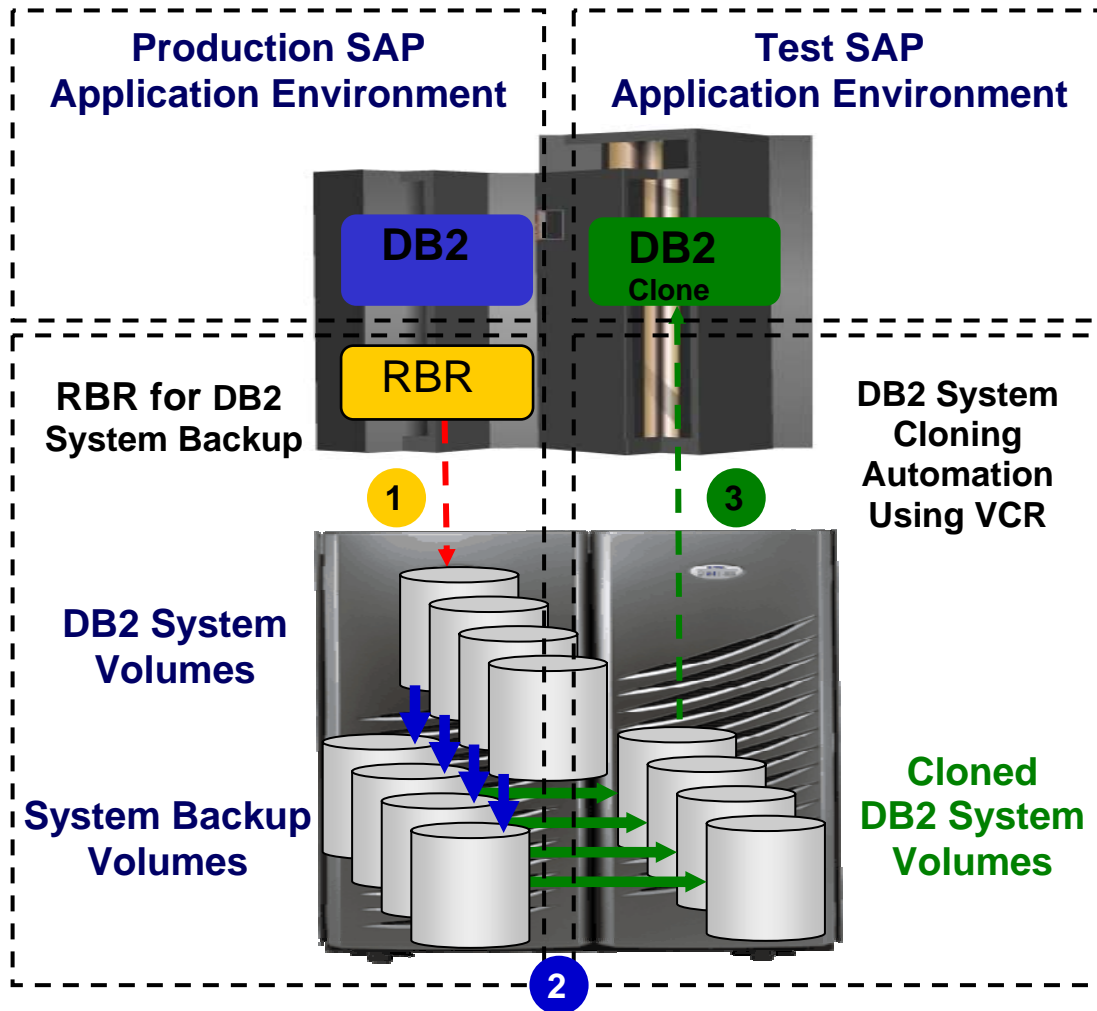


Figure 7. Cloning a DB2 system from a RBR for DB2-generated system backup

1. RBR generates a system backup using the procedures described in the section “Creating a DB2 system backup” on page 10.
2. VCR determines the volumes to be included in the DB2 cloning process by accessing the RBR metadata repository. VCR uses a specific DB2 identifier and an available system backup to retrieve a dataset containing corresponding source DB2 volumes and their associated backup volumes that are to be used for the cloning process. The backup volume dataset is used by VCR to copy the backup volumes to target volumes using TimeFinder facilities, recondition the volumes, and rename their datasets for the cloning process.

-
3. VCR performs the DB2 cloning automation described in the section “Creating a test SAP application environment from a production system” on page 16.

Refreshing DB2 table and index spaces in a test SAP system

Sometimes test SAP application environments require the ability to refresh particular DB2 objects from a production SAP system. Having the ability to refresh particular objects in a test SAP environment from a production SAP system can save time and resources over having to clone a complete DB2 system.

Fast Table Space Refresh for DB2 (FTR) provides facilities to clone or refresh table and index spaces within or across DB2 systems. The table and index space objects can be cloned or refreshed within and across databases and the creator ID of the cloned objects can be the same or different from that of the source objects. The target DB2 table and index space objects must exist on the target DB2 system before objects can be cloned or refreshed and they can be identified by FTR using the SIMULATE command or when the cloning process is executed. FTR can clone or refresh DB2 objects while they are stopped or while they are running to create fuzzy copies.

Figure 8 depicts refreshing table and index space objects from a production SAP system to a cloned SAP environment. All steps, except step 4, are performed by VCR’s Fast Table Space Refresh for DB2 option. Step 4 is executed using a FTR generated dataset list and is performed outside of FTR proper.

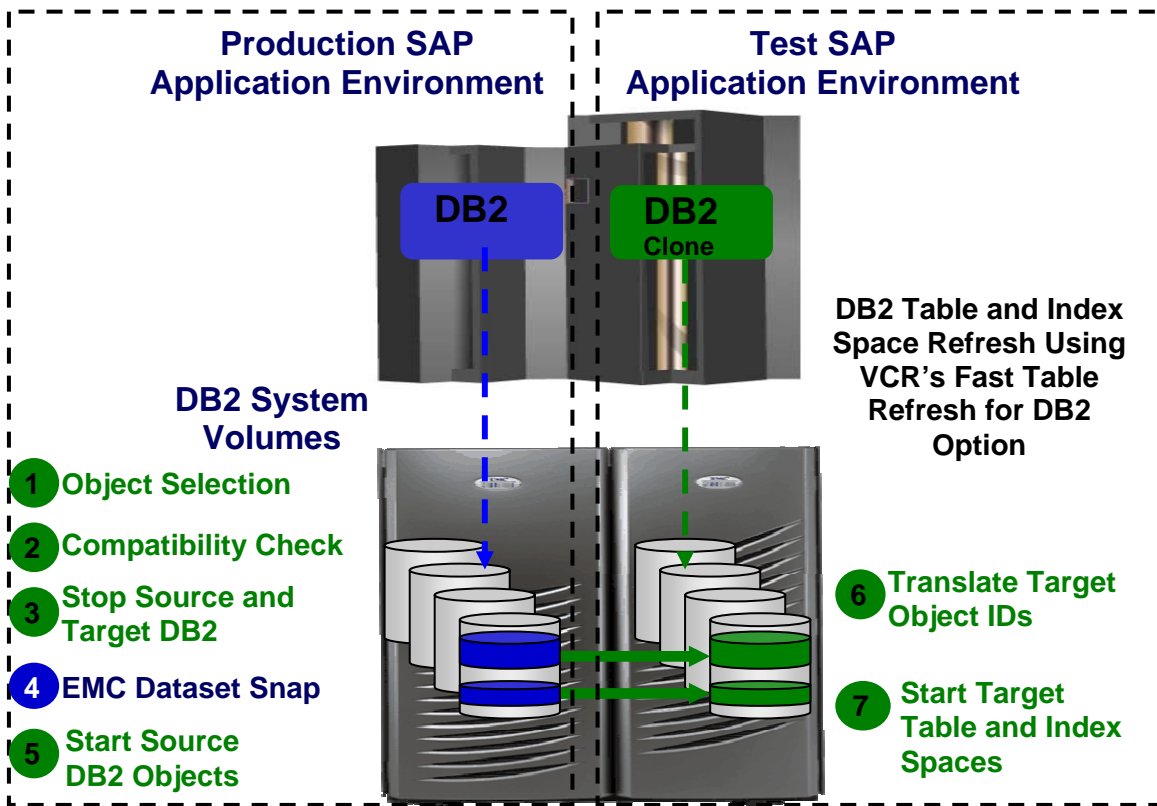


Figure 8. Refreshing DB2 table and index spaces from a production SAP system

1. The source table and index spaces to be refreshed or cloned are selected.
2. Checks are performed to ensure the characteristics of the source and target DB2 objects are compatible. Some characteristics that are checked include object names and qualifiers, column attributes, objects IDs, page size, and buffer pool utilization.
3. The source and target table and index space objects are stopped. Optionally, the source objects can be copied while they are running to create a fuzzy copy. Care must be taken when using the fuzzy copy option as transactional integrity is not guaranteed on the cloned copy.
4. Datasets representing the cloned or refreshed objects are copied. For SAP systems using TimeFinder, FTR optionally produces files describing source and target dataset information along with necessary control parameters so users can create their own dataset snap job streams.
5. The source DB2 objects are started unless a fuzzy copy was specified in step 2.
6. Object IDs are copied during the copy process in step 4. The object ID translation step changes the object IDs in the data pages to match those in the target DB2 catalog.
7. Target DB2 objects are started and made available to the SAP test environment.

DB2 solution validation on EMC Symmetrix V-Max

The EMC Symmetrix V-Max Series with Enginuity™ is a new entry to the Symmetrix family product line. Built on the strategy of simple, intelligent, modular storage, it incorporates a new scalable fabric interconnect design that allows the storage array to seamlessly grow from an entry-level configuration into the world's largest storage system. Symmetrix V-Max provides improved performance and scalability for demanding enterprise storage environments while maintaining support for EMC's broad portfolio of platform software offerings.

The Enginuity operating environment for Symmetrix version 5874 is a new, feature-rich Enginuity release supporting the Symmetrix V-Max storage arrays. With the release of Enginuity 5874, Symmetrix V-Max systems deliver new software capabilities that improve capacity utilization, ease of use, business continuity, and security.

Rocket Backup and Recovery for DB2 (RBR), Mainstar Volume Clone and Rename (VCR), and Mainstar Fast Table Space Refresh for DB2 (FTR) have been tested and validated using Symmetrix V-Max. The following SAP and DB2 management functions have been performed with the products described in this white paper using Symmetrix V-Max.

- RBR - DB2 system backup using EMC TimeFinder/BCV, TimeFinder/Clone, and TimeFinder/Snap
- RBR - DB2 system backup using EMC Compatible FlashCopy
- RBR - DB2 system backup using DB2 Backup System
- RBR - DB2 system restore and recover using EMC TimeFinder/BCV, TimeFinder/Clone, and TimeFinder/Snap
- RBR - DB2 system restore and recover using EMC Compatible FlashCopy
- RBR - DB2 system restore and recover using DB2 Restore System
- VCR - DB2 system cloning using EMC TimeFinder/Clone
- VCR - DB2 system cloning using EMC Compatible FlashCopy
- FTR - Fast Table Space Refresh using EMC Dataset Snap

Conclusion

SAP implementations that use DB2 for z/OS tend to require high performance and high availability operations in support of their SAP applications. Backup operations must not impact SAP applications. Recovery operations must be swift, streamlined, and easy to perform. SAP cloning operations must create test and quality assurance environments in a simple, transparent, and automated manner.

EMC and Mainstar have teamed to offer storage integrated database products and SAP solutions that simplify DB2 backup, recovery, disaster recovery, and cloning operations for SAP environments using z/OS. These products provide tight storage integration that simplifies DB2 backup processes and speeds SAP recovery. DB2 cloning automation provides an easy way to create and refresh test SAP landscapes.

The solutions discussed in this paper are based on the following EMC and Mainstar products: Rocket Backup and Recovery for DB2, Mainstar Volume Clone and Rename, and Mainstar Fast Table Space Refresh for DB2. All of these products are tightly integrated with EMC TimeFinder and can be used independently or collectively to simplify the management of SAP applications using DB2 for z/OS. All products and solutions referenced in this paper have been validated and approved for use with EMC Symmetrix V-Max.

About EMC

EMC Corporation is the world's leading developer and provider of information infrastructure technology and solutions that enable organizations of all sizes to transform the way they compete and create value from their information. Information about EMC's products and services can be found at www.EMC.com.

About Mainstar

Mainstar has focused on IBM System z software infrastructure, storage management, and disaster recovery for over 30 years. Mainstar is dedicated to providing software and services for z/OS customers. Seven of the top 10 Fortune 500 companies turn to Mainstar for System z products and expertise. Information about Mainstar's products and services can be found at www.mainstar.com.

Mainstar is a wholly owned subsidiary of Rocket Software.