

IBM DB2 Cloning Tool for z/OS

Amerigo Baldassarri – Rocket Software – Technical Specialist

27 August 2013



Agenda


- Trends and directions
- Cloning terminology
- Use cases
- Options for copying data
- Clone DB2 systems
- Clone from a system level backup
- Refresh DB2 table and index spaces

Database and Storage Administration Trends and Directions

- Large DB2 and IMS systems require high availability
 - Fast and non-intrusive backup and cloning facilities are required
 - Fast recovery capabilities are required to minimize downtime and promote high availability
 - Most backup, recovery and cloning solutions do not leverage storage-based fast-replication facilities
- Storage-based fast-replication facilities are under-utilized
 - Tend to be used by storage organizations
 - Tend not to be used by database administrators (DBAs)
- Storage aware database products
 - Allow DBAs to use fast-replication in a safe and transparent manner
 - Provide fast and non-intrusive backup and cloning operations
 - Simplify recovery operations and reduces recovery time
 - Simplify disaster recovery procedures

Cloning Terminology

- A clone is an exact but independent replica
 - Clone a DB2 system by volume
 - Clone a table space by data set
- DB2 system cloning and table space refresh
 - The act of replicating the data, making the replica accessible, and then using the replica in lieu of the original data
- DB2 system cloning automation
 - Clones a complete DB2 system including all its databases **by VOLUME**
- DB2 table and index space refresh automation
 - Refreshes specific table and index spaces **by data set**
 - Lowest level is by data set



Cloning Terminology

Fast-replication or Snap

The capability of a storage system which uses its processor and memory to give the appearance of an instantaneous copy of a volume or a data set, Database, Table Space, system.

Slow copy

Any copy utility that uses the z/OS system's host CPU to copy a volume or a data set, Database, table space, system.

Mirror

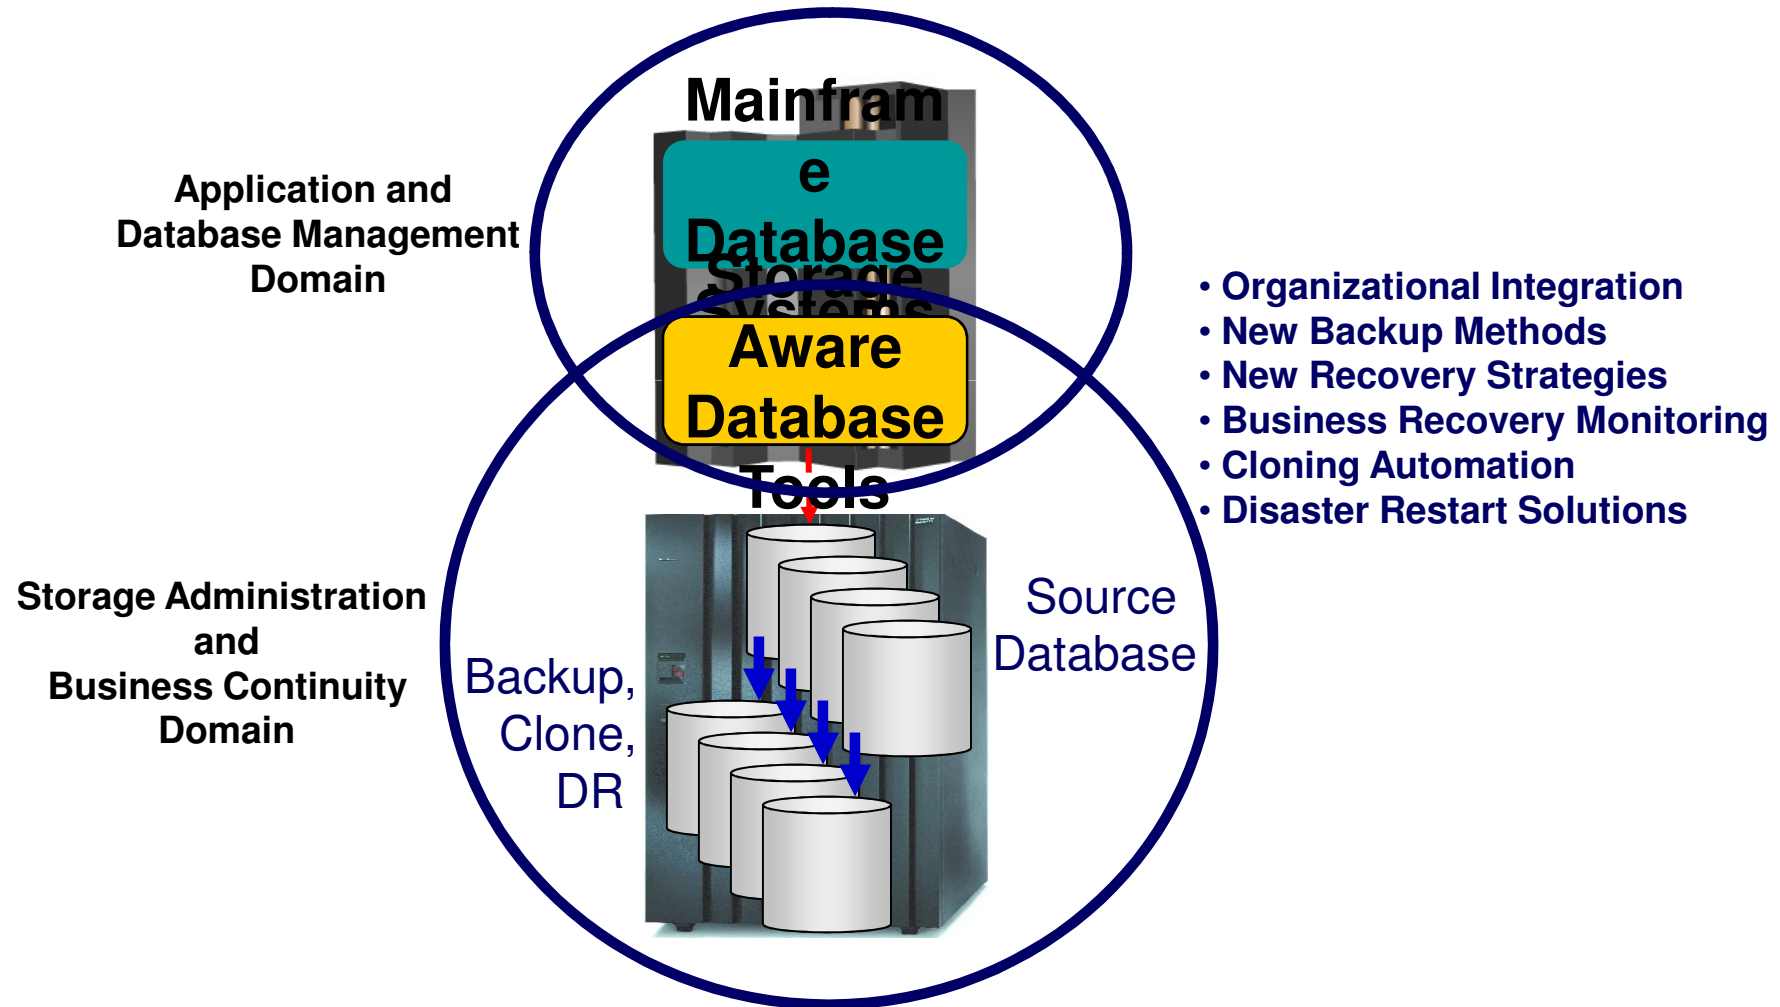
A copy procedure that establishes a relationship between a source and target volume such that all changes on the source are "mirrored" on the target volume.

Smoothing the migration path

Cloning DB2 is not a new concept

- There are 2 basic traditional methods
 - The first method – Resource intensive:
 - Creates a new DB2 (the clone)
 - Creates the objects from a source DB2 on the new DB2 via DDL
 - Unloads the tables of the source DB2
 - Loads the table rows of the source DB2 onto the tables of the clone
 - The second method – Requires an isolated LPAR:
 - Dump the data sets of a source DB2 to tape or some movable media
 - On an **isolated** LPAR restore the data sets of the source DB2
 - Start the clone DB2
- A third and non-traditional method is DB2 Cloning tool
 - Safer
 - Faster
 - Not resource intensive
 - Can clone on the same or on a shared LPAR
 - Uses storage technology which you already own

Database and Storage Integration



Database and Storage Integration Operational Advantages

- Reduce backup, recovery, and cloning administration costs
- **Reduce host CPU and I/O resource utilization**
- **Perform backups and create clone copies instantly**
- Fast restore and parallel recovery reduces recovery time
- Simplify disaster recovery operations and procedures
- DBMS and storage-based fast-replication integration
 - Leverage storage processors and fast-replication investments
 - IBM, EMC, HDS, STK
 - Expose fast-replication capabilities to the DBAs ***safely and transparently*** using “***storage-aware***” database utilities
- Provide a sophisticated infrastructure and metadata to manage the DBMS and storage processor coordination

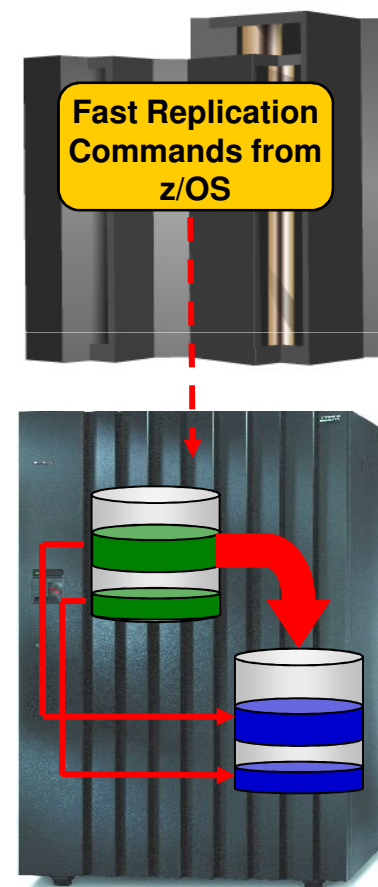
Use Cases for Cloning DB2 Systems

- Offload business reporting to a DB2 clone
 - Improve production performance
- Create or refresh test, development, or quality assurance environments
 - Copy SAP interrelated data
 - Copy PeopleSoft interrelated data
 - Stage data-warehouse loads
- Aid in problem determination
 - Troubleshoot a production problem
 - Test new releases of DB2
 - Apply maintenance and verify integrity before applying to production
- Clone from a system level backup

Fast Replication Data Copy Options

Fast copy processes offloaded to the storage processor
No host CPU or I/O resources

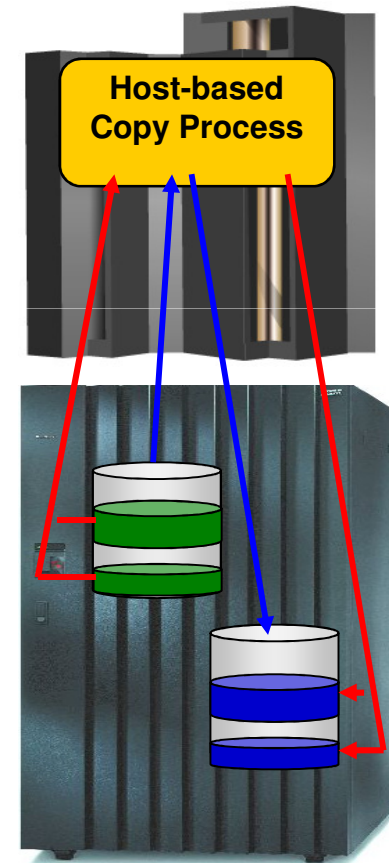
- Volume based fast replication options for DB2 system cloning
 - FlashCopy (IBM,EMC,HDS)
 - SnapShot (IBM,STK)
 - TimeFinder/Clone Volume Snap (EMC)
 - TimeFinder/Snap (EMC)
 - Mirror processes
 - PPRC (IBM,EMC,HDS)
 - TimeFinder/Mirror, SRDF (EMC)
 - ShadowImage HUR (HDS)
- Data set based fast replication options for DB2 table space refresh
 - Data Set FlashCopy (IBM,EMC,HDS)
 - Data set SnapShot (IBM,STK)
 - TimeFinder/Clone Data set Snap (EMC)



Host Based Data Copy Options

Data copy processes use host based CPU and I/O facilities
Slower than storage-based fast replication

- Volume copy options for DB2 system cloning
 - TDMF (IBM)
 - FDRPAS (Innovation Data Processing)
 - DFSMSDss (IBM)
 - FDR (Innovation Data Processing)
- Data set copy options for DB2 table space refresh
 - Any traditional data set copy processes



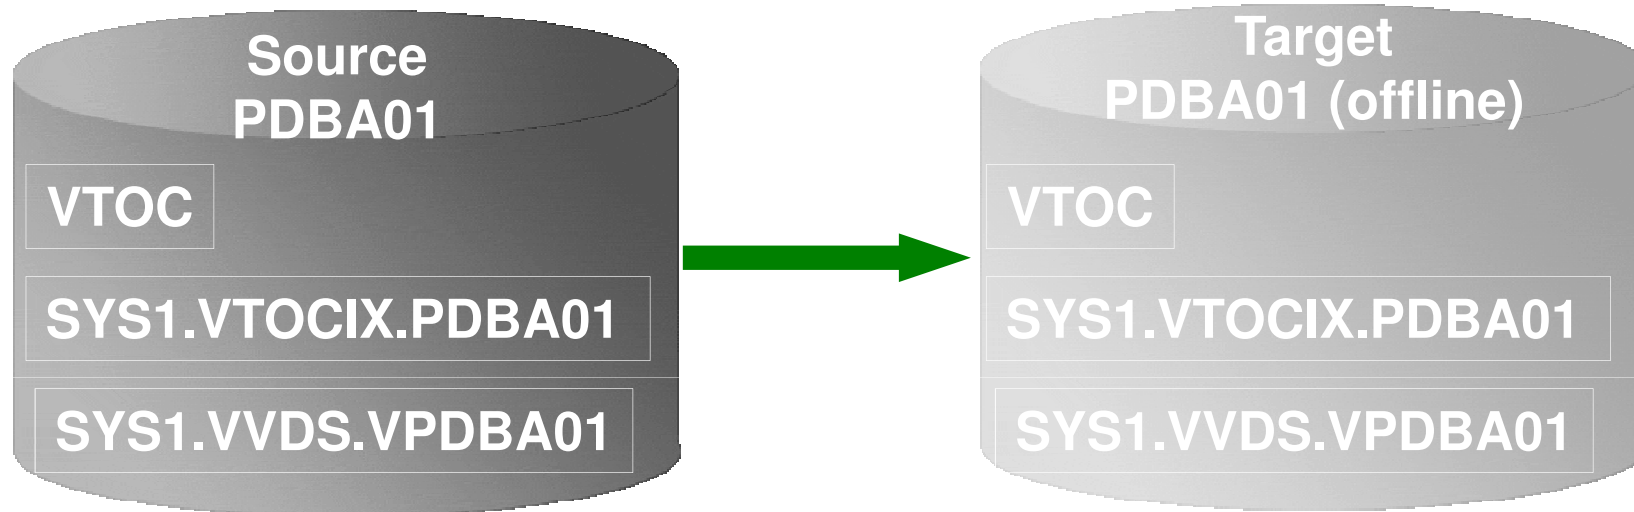
Challenges to Data Access On the Same or Shared LPAR

- DB2 system cloning is best done using volume replication
- Volume data is replicated fast and easily but...
 - How do you access the cloned data?
- **Inherent Problems:**
 - VOLSERS may have the same volume names as the source
 - Volume VTOC, VTOCIX and VVDS may be the same or different
 - Data set names are the same on source and target volumes
 - **If you want to access the data from a common LPAR, how do you access the data?**



Challenges to Data Access

On the Same or Shared LPAR - Volume ID Issues

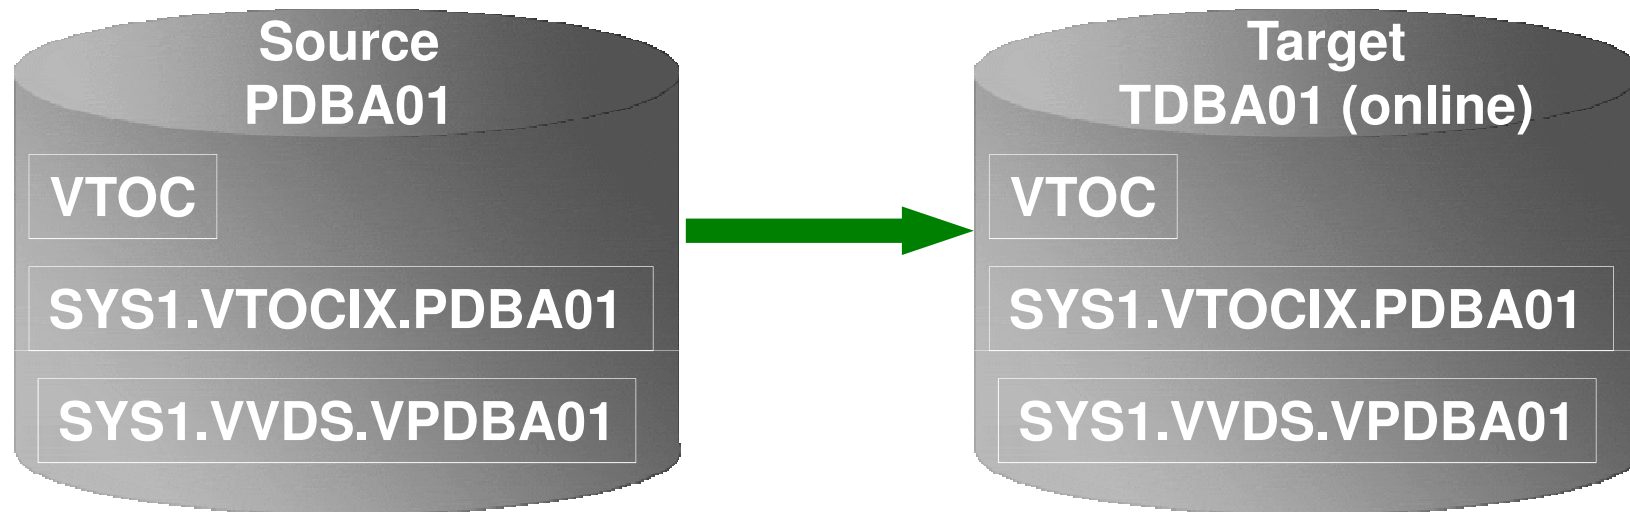


Result:

1. Source PDBA01 keeps its original VOLSER on the target volume
2. VTOC, VTOCIX, and VVDS *reflect* PDBA01
3. The target volume is offline because it has the same VOLSER as the source. It could be brought online on another non-sharing system to use the clone
4. Target volume can't be used on the same or shared disk LPAR without re-labeling the volume.

Challenges to Data Access

On the Same or Shared LPAR - Volume ID Issues (2)

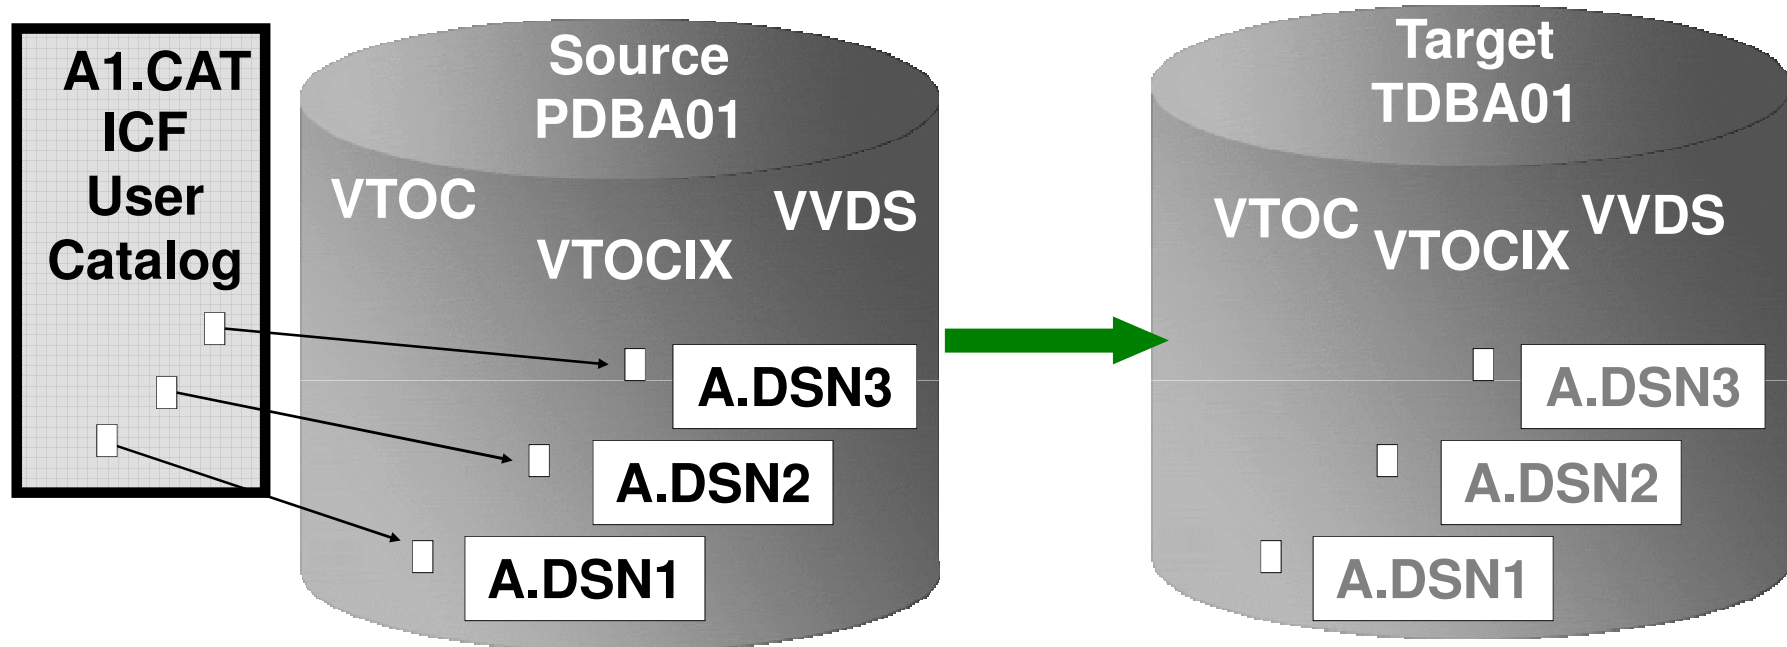


Result:

1. Source PDBA01 becomes new VOLSER TDBA01 on the target volume
2. VTOC, VTOCIX, and VVDS *reflect* PDBA01

Challenges to Data Access

Data Set Name and Cataloging Issues



Result:

1. Data sets on the volume are copied, but keep their original name
2. Only the source data sets are cataloged; even if the catalog is on the cloned volumes, it isn't connected to the system's master catalog

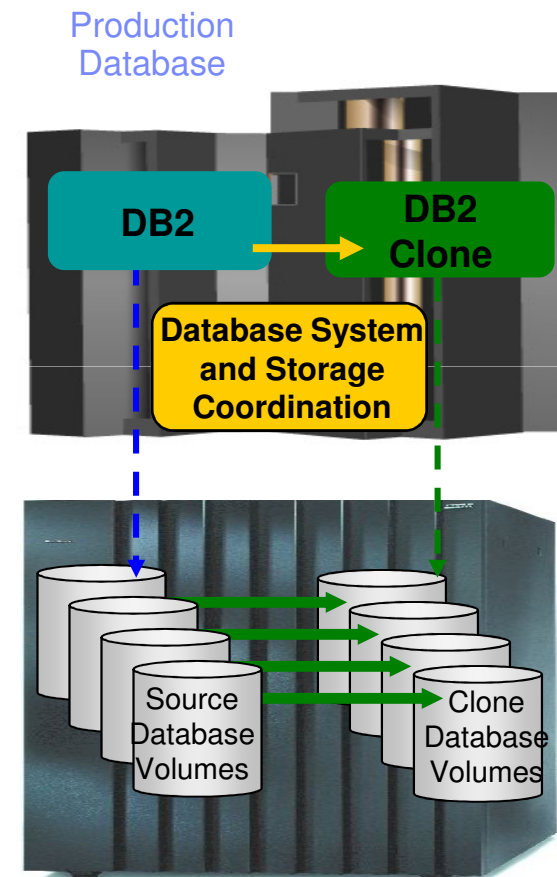
DB2 Cloning Tool Provides Data Access On the Same or Shared LPAR

- Solution – DB2 Cloning Tool
 - the *Key to Data Access*
 - Provides *access to data sets* on target volumes
 - Renames the VTOC, VTOCIX, and VVDS to match the target volume
 - Renames and catalogs all data sets to a new HLQ
 - Conditions the target DB2 system to accommodate and accept the cloned data
 - Solution to enable the use of the data sets on the target volumes on the same image
 - No more complicated cloning
 - No more requirement for a separate image

Clone DB2 Systems Using DB2 Cloning Tool

DB2 System Cloning Automation

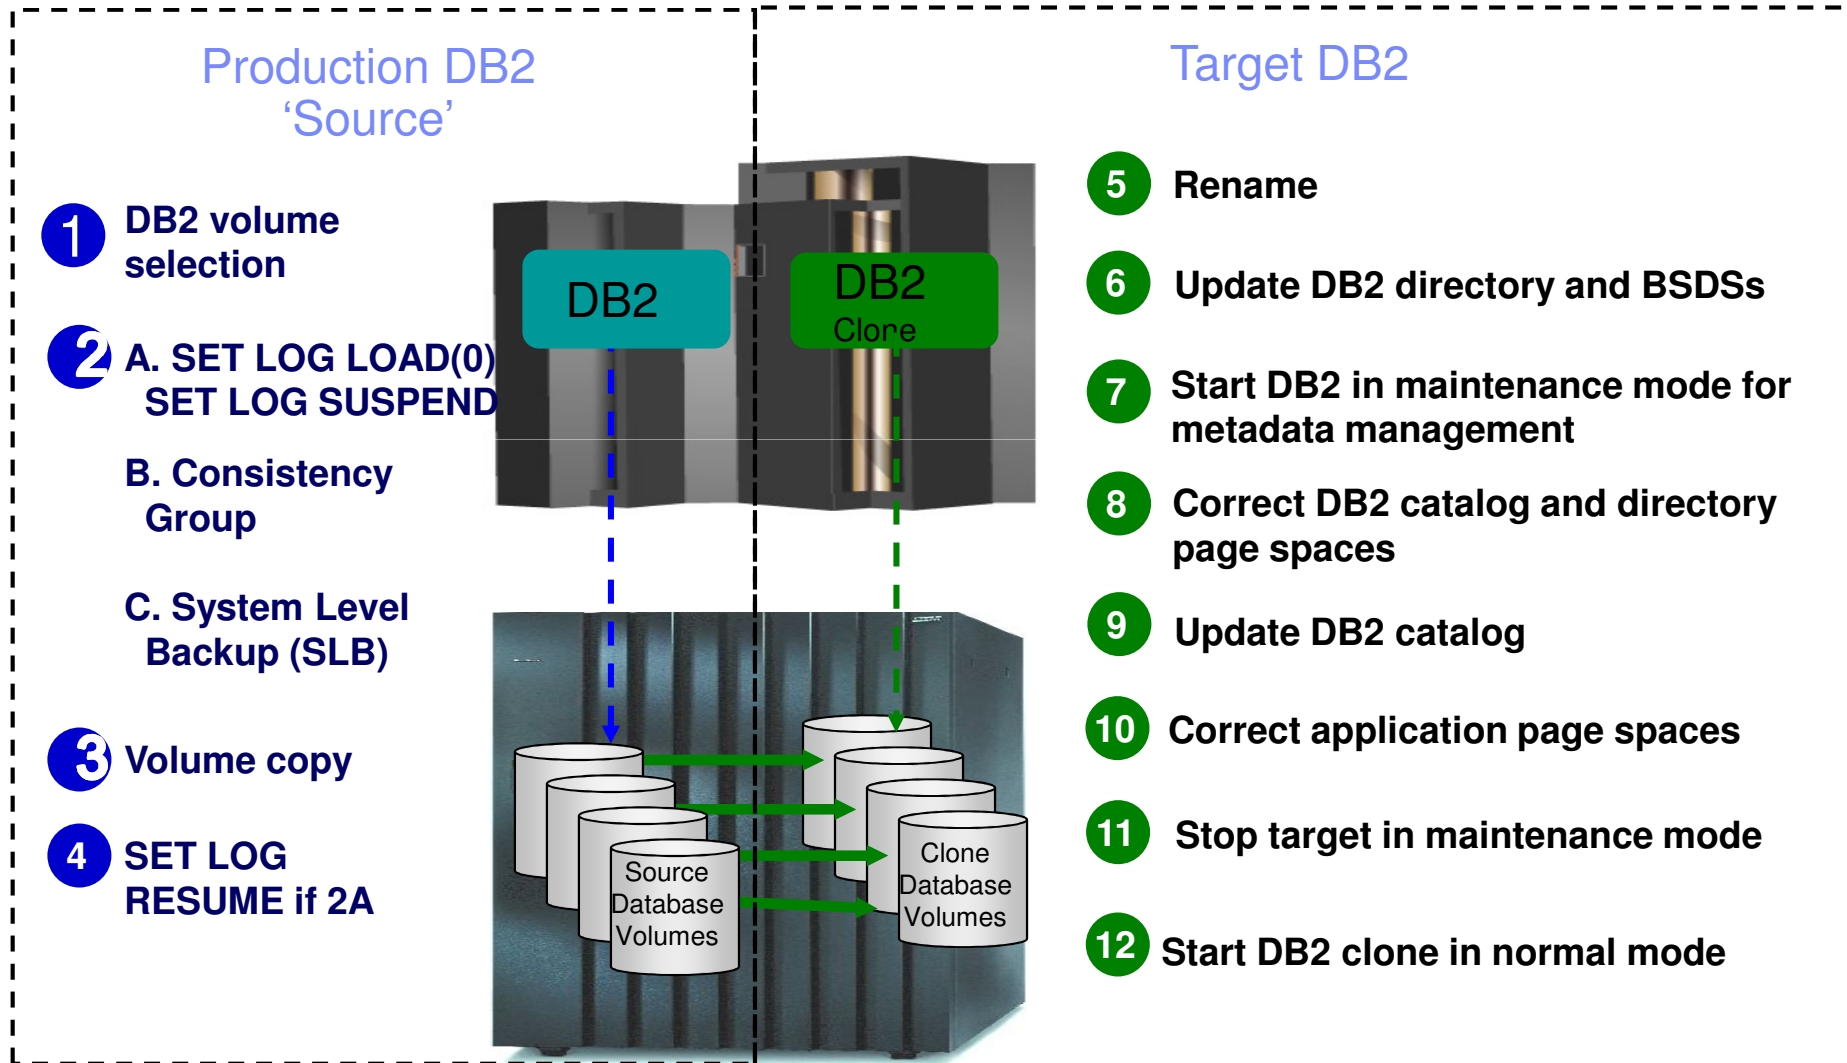
- Performs automated cloning of DB2 systems
- DB2 data copied using storage-based dataset fast-replication
 - Data can be cloned while online or offline
- Performs rapid volume reconditioning and data set renaming on cloned volumes to solve the data access challenges
 - Target volumes retain their target volume label
 - Renames the VTOC, VTOCIX, and VVDS to match the target volume
 - Renames and catalogs all data sets to a new HLQ
- Adjusts target DB2 system to accommodate and accept the cloned data
 - DB2 catalog, directory, BSDS, active / archive Log
 - Makes data accessible on the same or shared LPAR



Pre-DB2 System Cloning Steps

- Set up the target DB2 system definition before performing DB2 cloning automation
- Set up the special DB2 system maintenance mode ZPARM before performing DB2 cloning automation
 - Allows the DB2 catalog to be updated
 - Starts the target DB2 in DEFER mode to prevent back out processing

DB2 System Cloning Steps



Status of In-Flight Transactions

- Status of In-Flight Transactions
 - An online cloning solution often results in transactions in flight
 - These in-flight transactions, cloned to the target system result in the same target system action that would happen on the source system if it were to have died at that same time and then restarted
 - The target restart is essentially an emergency restart of a failed system

Target DB2 System (Clone) Updates

- DB2 directory updates (*step 6 in previous diagram*)
 - The VCATNAME
 - Optionally, the DB2 storage group names
- BSDSs updates (*step 6 in previous diagram*)
 - The DB2 catalog name
 - The 'active' log data set names
 - Optionally, the ARCHIVE data set names and volume serial numbers
 - Optionally updates the target DB2 BSDS's DDF parameters
- DB2 Catalog updates by SQL (*step 9 in previous diagram*)
 - The DB2 VCATNAME name

A DB2 stored procedure that works with the DB2 Administrative Task Scheduler to clone an entire DB2 subsystem

- Stored procedure is an application interface (API) to DB2 Cloning Tool
- The stored procedure will:
 - Generate the necessary jobs to do the subsystem cloning
 - Schedule the jobs in the DB2 Administrative Task Scheduler and cause the jobs to be submitted
 - Monitor the execution of the jobs
- The stored procedure will return to the caller, when the requested cloning has ended, either in success or failure
- It can be called from a platform other than z/OS to clone an entire z/OS DB2 subsystem
- End users can initiate their own re-cloning if desired

DB2 Support

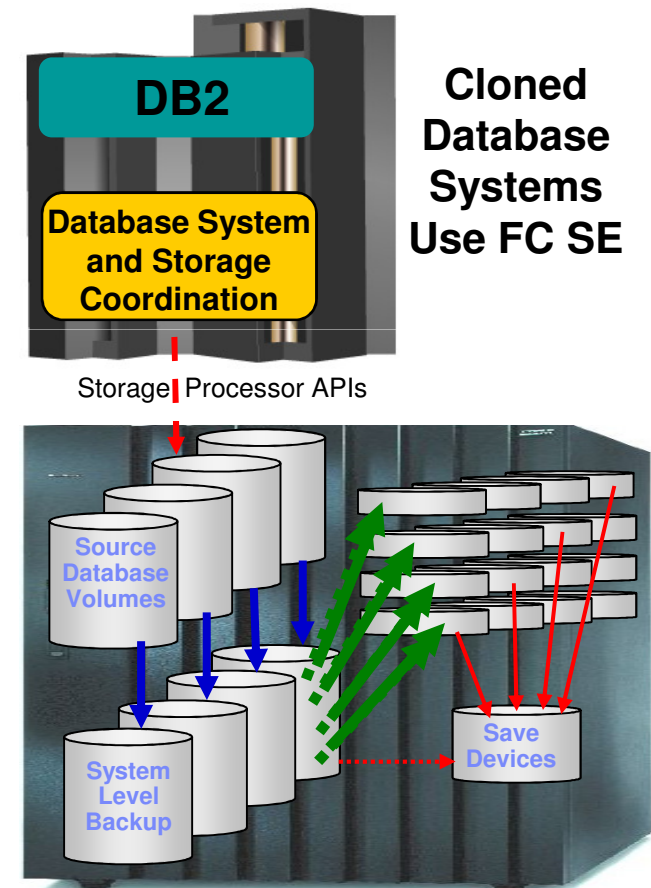
- DB2 Support
 - DB2 offline
 - DB2 online
 - DB2 data sharing
 - DB2 data sharing with many to less members
 - DB2 data sharing to non-DB2 data sharing

Introduction to System Level Backups

- System Level Backups (SLB) are performed by:
 - DB2 Recovery Expert for z/OS
 - DB2 BACKUP SYSTEM utility
- System Level Backup Overview and Advantages
 - Back up entire DB2 systems near instantaneously with no impact to the business
 - Provide fast restore and parallel recovery to reduce recovery time and complexity
 - Transform disaster recovery into a disaster restart process reducing RTO
 - Significantly reduce costs by using less CPU and I/O resources
 - Utilize one backup for multiple purposes to reduce storage costs
 - Performs backup validation to achieve successful recoveries
 - Eliminate backup windows and extend batch processing windows by removing image copies
 - Eliminate contention on production objects by creating image copies from system backups
 - Simplify and enhance DB2 BACKUP SYSTEM implementations
 - Perform DB2 system clone operations from a system level backup

Clone from a System Level Backup without using Additional Storage

- Full system-level backup created using full volume fast-replication
- DB2 system clone uses SLB backup volumes as source
 - DB2 system clone can use Space Efficient FlashCopy
- SLB volumes are used to service I/O for database clone access
 - Clone writes (few) go to the save pool
 - SLB writes (none) go to the save pool
- Storage-aware database tools provides infrastructure and metadata to manage database and storage processor coordination



Q & A



Use Cases for Refreshing DB2 Table and Index Spaces

Offload business reporting to a DB2 clone

Improve production performance

Refresh data in a test, development or quality assurance environment

Development refresh

To stage data-warehouse loads

To aid in problem determination

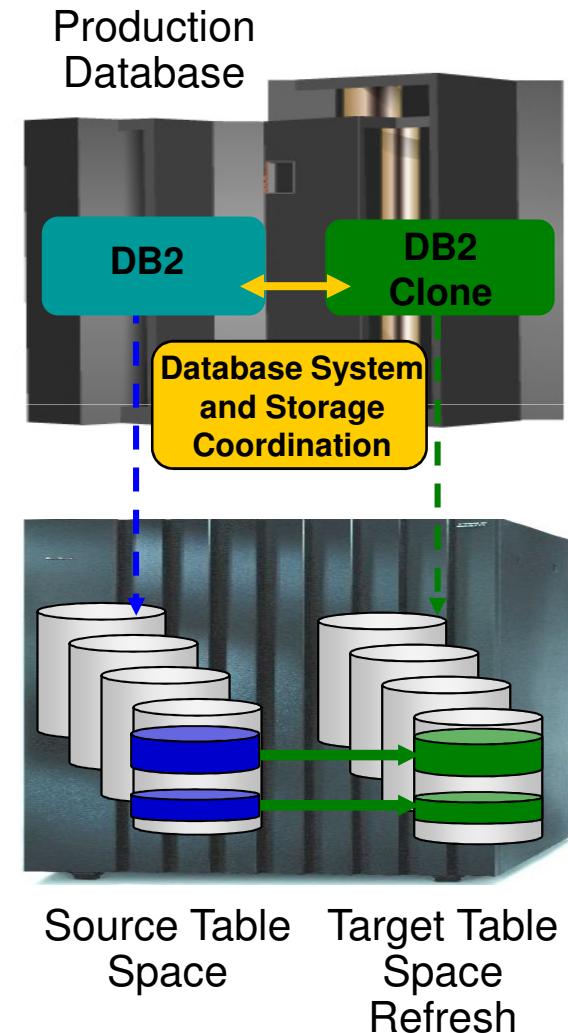
Troubleshoot a production problem

Refresh table spaces into a previously cloned DB2 skeleton

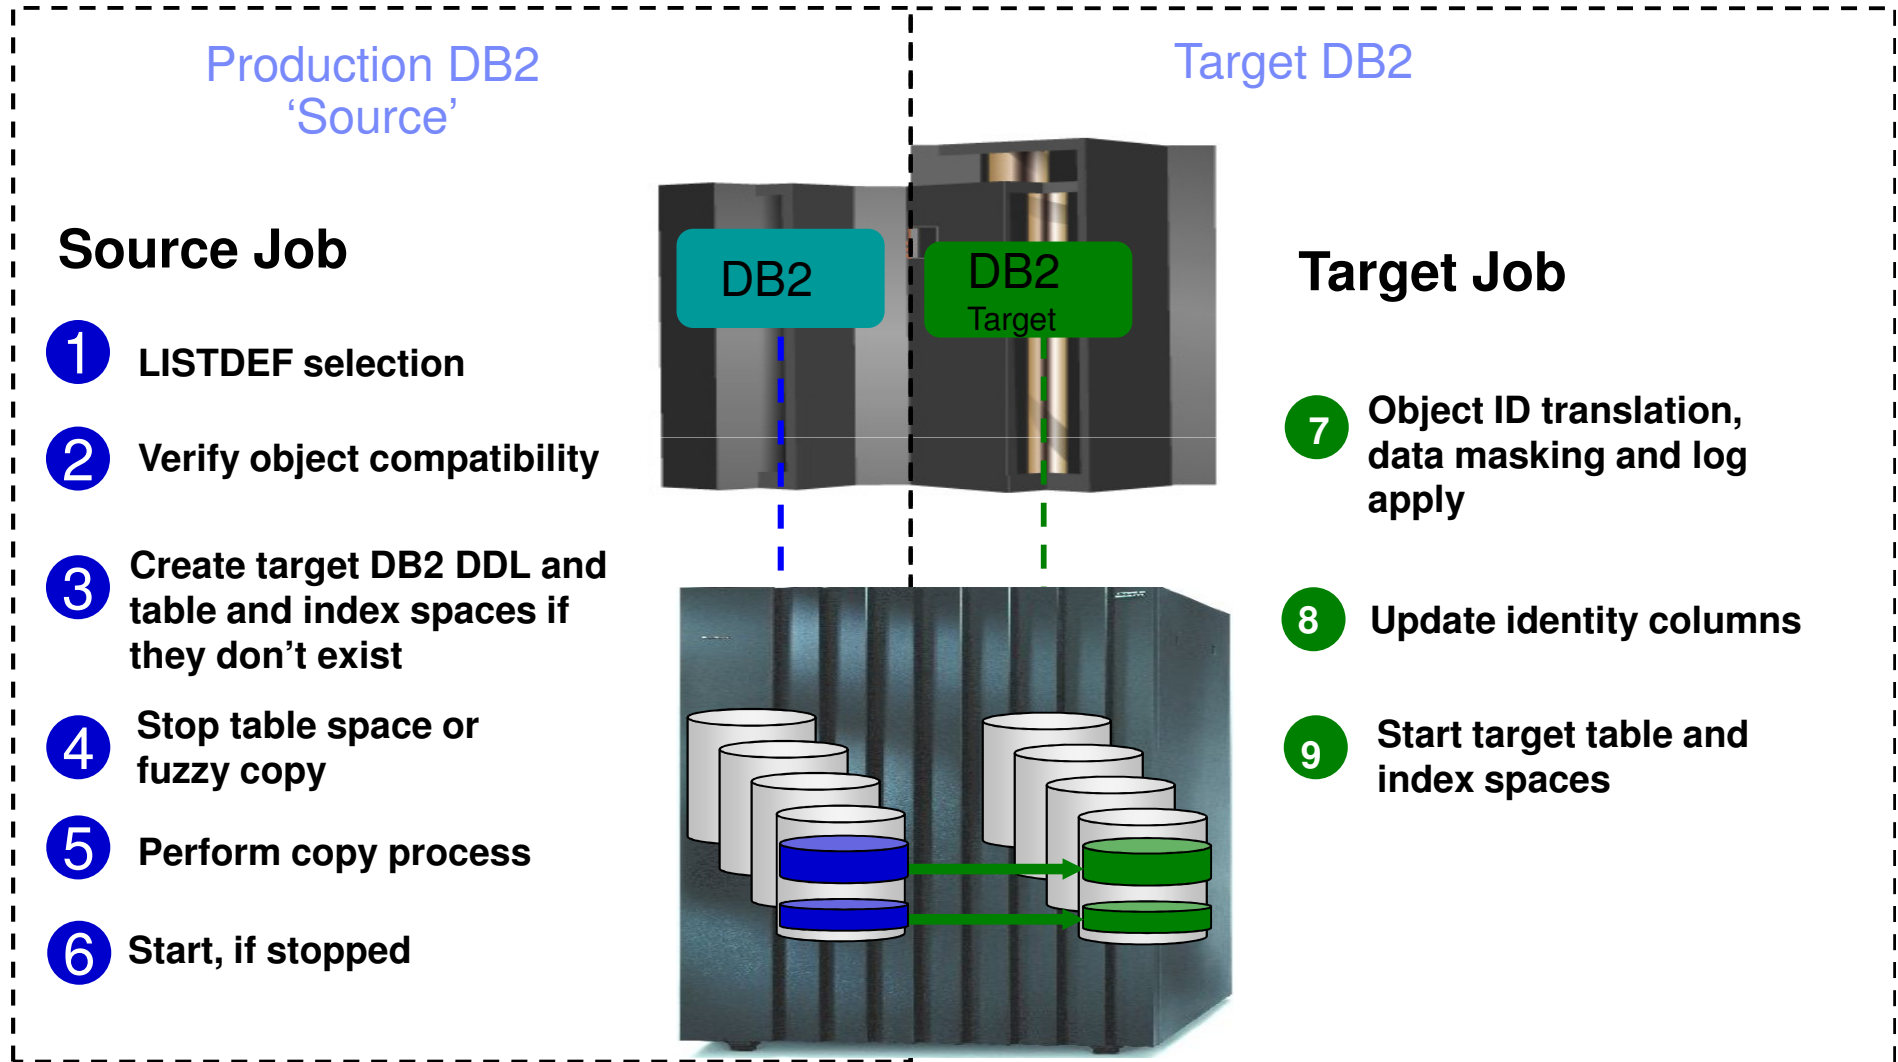
Refresh DB2 Table and Index Spaces Using DB2 Cloning Tool

Table and Index Space Refresh Automation

- Performs automated DB2 table and index space refresh operations
 - DB2 RI relationships, LOBS, and Identity columns
 - XML on DB2 V9 or greater
- Verifies source target database compatibility
- DB2 data copied using storage-based dataset fast-replication
 - Data can be cloned while online or offline
 - Slow copy mechanism can be used
- Performs object ID translations, data masking, log apply, and target system meta-data



DB2 Table and Index Space Refresh Steps



Target locations that DB2 Cloning Tool supports

- Within the same DB2 system or to another DB2 system
- To the same database name or to another database
- To the same table space name, or to another table space name
- To the same Creator ID or to another Creator ID

DB2 Cloning Tool does the Catalog Research for you

- Determines the source and target data sets
- Determines the source and target object IDs
- Determines compatibility
- Creates the XLATE parameters
- Tracks source DB2 extents

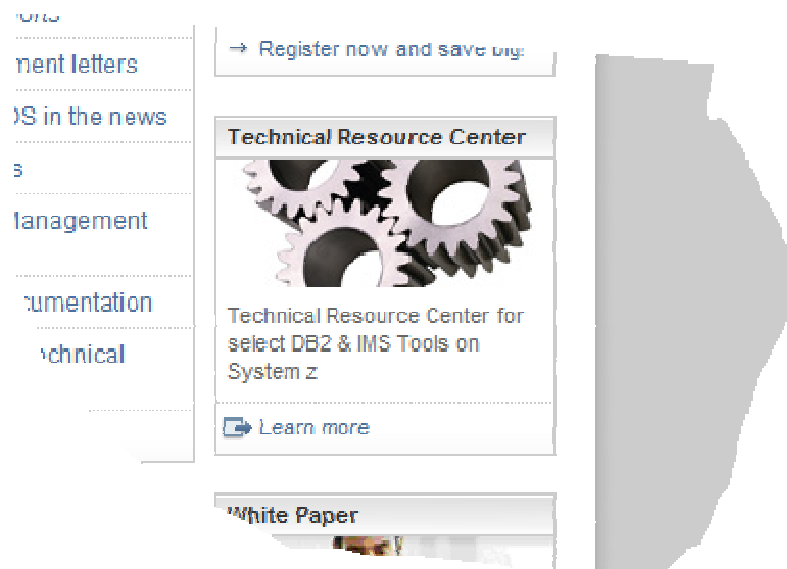
DB2 Cloning Tool Session Summarization

- DB2 Cloning Tool simplifies and automates database administration tasks
 - Integrates and coordinates database and storage activities
- Promote high availability solutions
 - DB2 Cloning Tool clones DB2 systems fast and effectively
 - DB2 Cloning Tool provides fast and effective refresh operations
- Integrates storage-based fast replication
 - Provides fast and non-intrusive cloning operations
- Offloads copy to the storage processor
 - Reduces host CPU and I/O resource utilization



The DB2 and IMS Tools Technical Resource Center

- Focused education on select DB2 and IMS Tools
 - Access to product FAQs, video downloads, interactive webinars, and subject matter experts
 - Find on the DB2 Tools and IMS Tools Home Pages



Q & A

