



# Pulse2011



## IBM Cloning Tool for z/OS

Presenter's Name – Amerigo Baldassarri

Presenter's Title – Technical Specialist

# Trademarks and disclaimers

© Copyright IBM Australia Limited 2011 ABN 79 000 024 733 © Copyright IBM Corporation 2011 All Rights Reserved.  
TRADEMARKS: IBM, the IBM logos, ibm.com, Smarter Planet and the planet icon are trademarks of IBM Corp registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.



# AGENDA

- Trends and directions
- Cloning terminology
- Use cases
- Options for copying data
- Clone DB2 systems
- 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



## 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



## Cloning Terminology

- A clone is an exact 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
- DB2 table and index space refresh automation
  - Refreshes specific table and index spaces
  - Lowest level is by data set



## Examples where you might want to have the ability to Clone 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 easy 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)



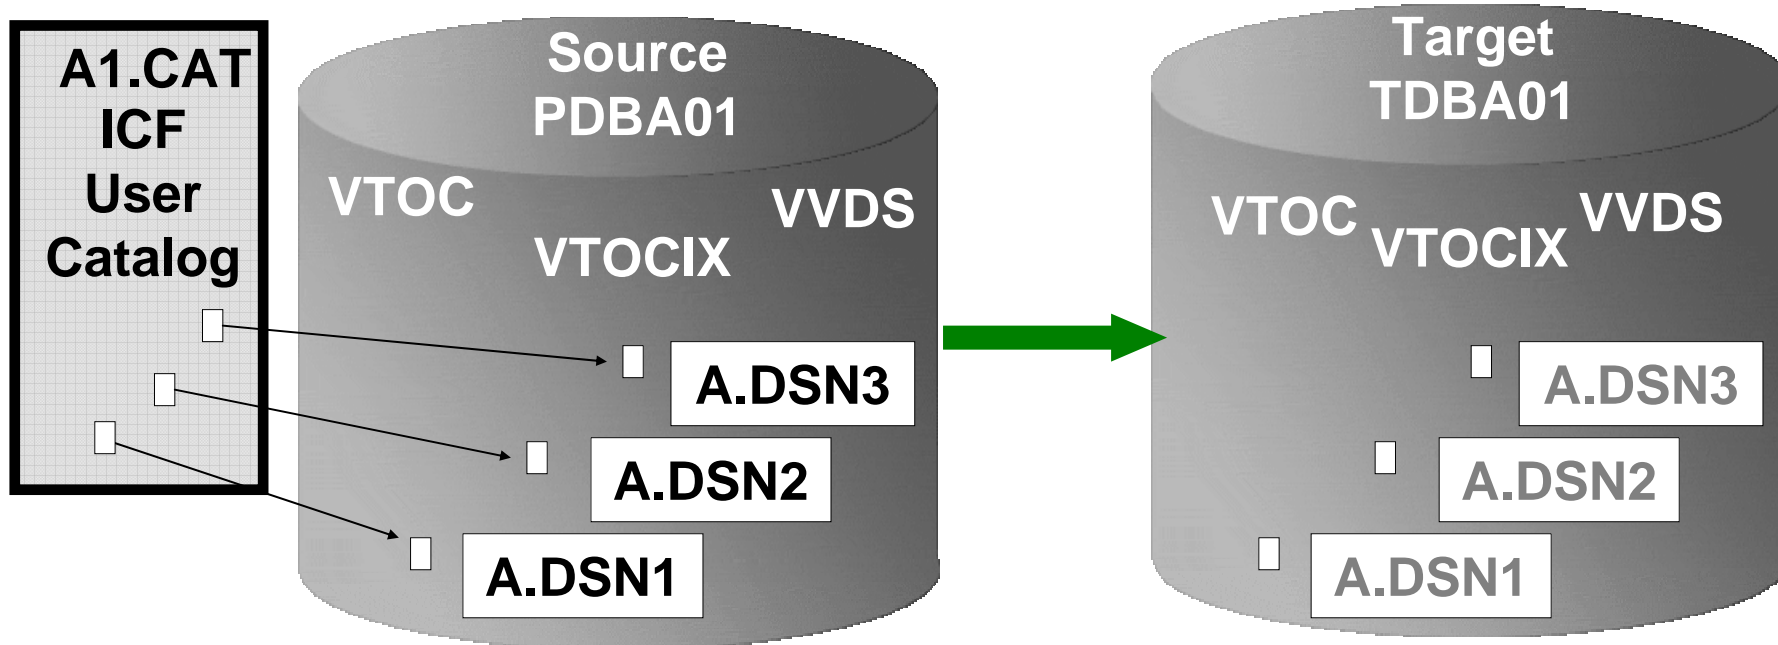
#### Result:

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
    - Adjusts 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



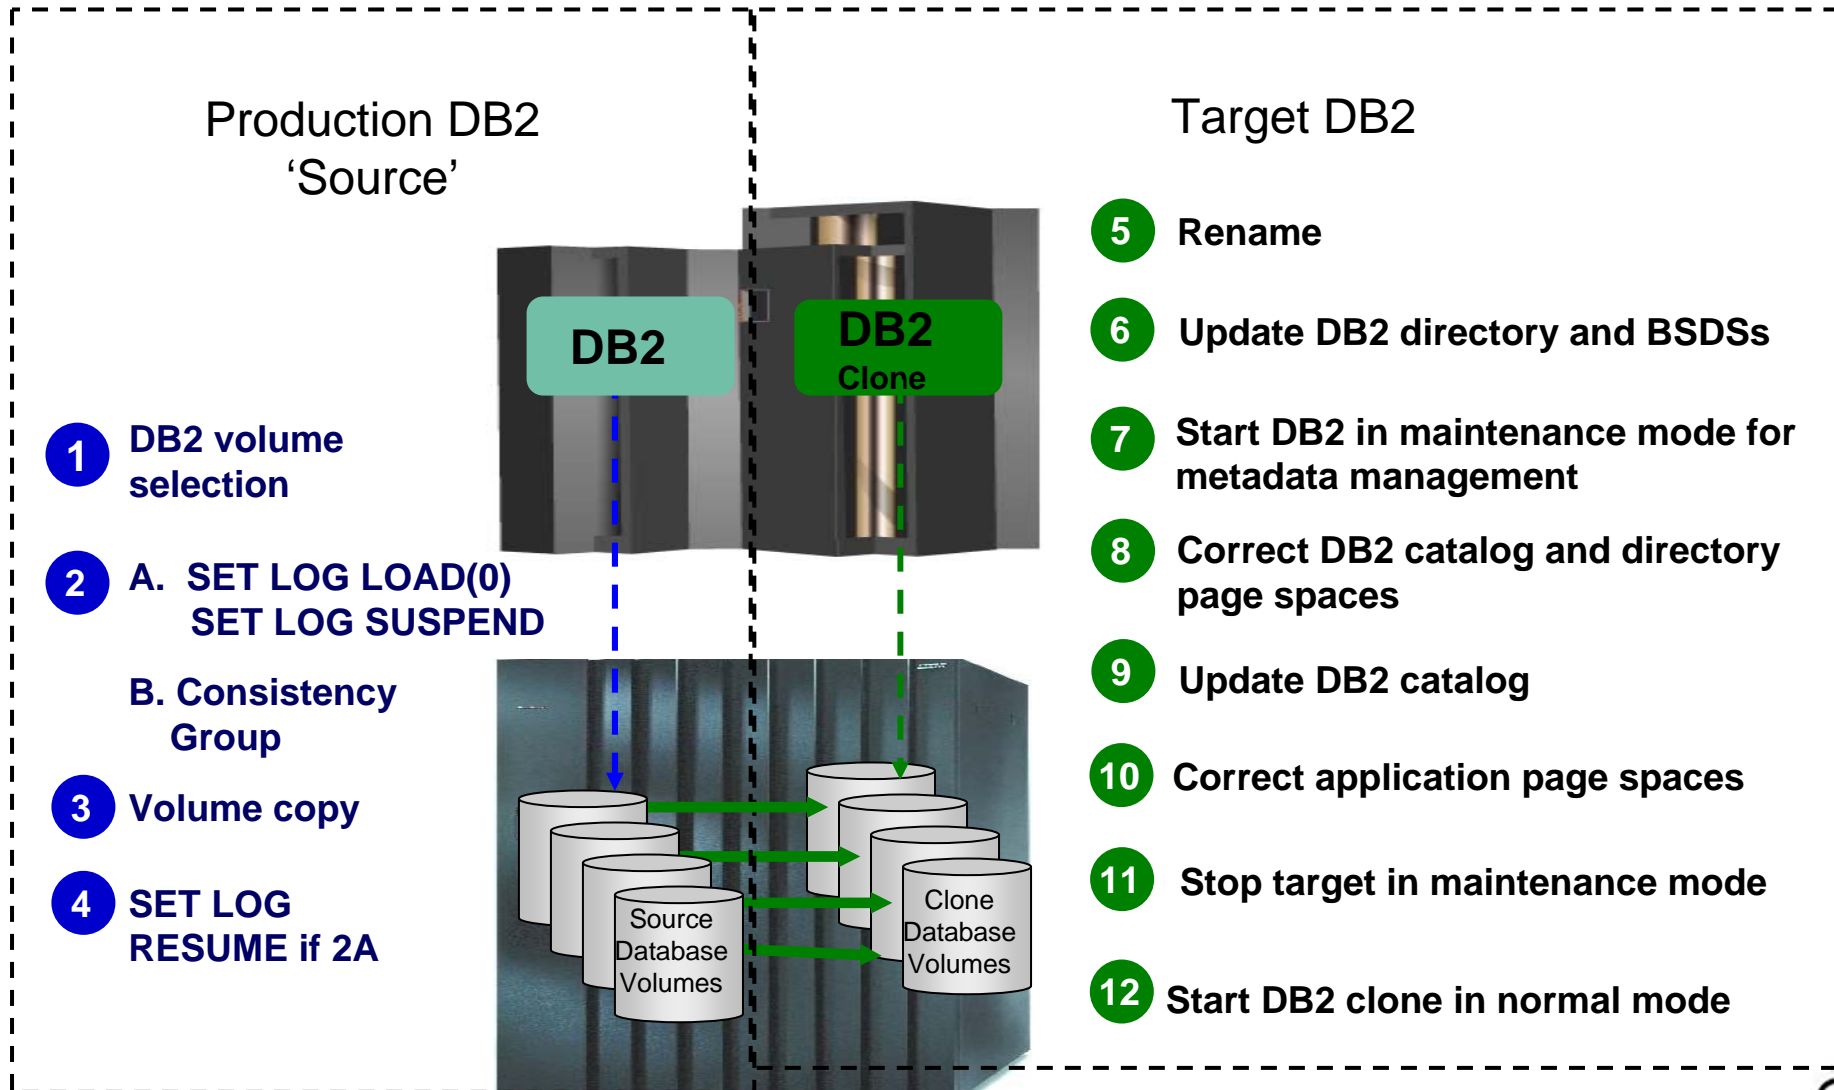


## Pre-DB2 System Cloning Steps

- **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



## Task Management! Who does what?

- **Cloning can be broken up into two distinct areas of responsibility**
  - Storage Functions
  - Database Administrative Functions

### STORAGE FUNCTIONS

SOURCE/TARGET Volume Selection and Preparation

COPY Step

RENAME Step

### DATABASE ADMINISTRATIVE FUNCTIONS

DB2 ONE TIME INITIAL SETUP eg ZPARMS and PROCLIB

DB2 SUBSEQUENT CONDITIONING STEPS eg VCAT, BOOTSTRAP



## 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



# Types of DB2 Support available

## ■ DB2 Support

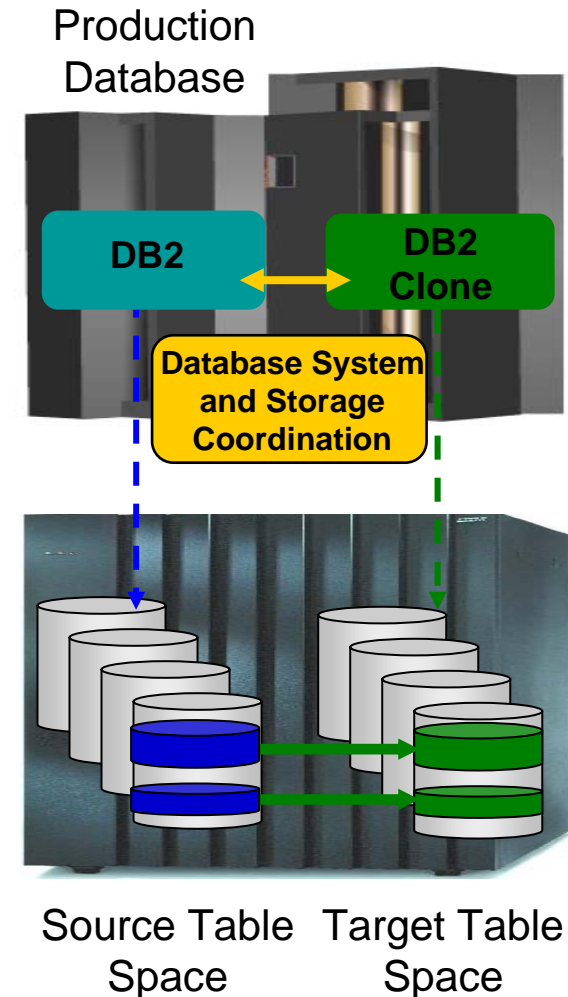
1. DB2 offline cloning DB2 Subsystem completely stopped
2. DB2 online cloning where a DB2 SET LOG SUSPEND command is issued to suspend transactions, or when a consistent copy process is used.
3. DB2 data sharing
4. DB2 data sharing with many to less members
5. DB2 data sharing to non-DB2 data sharing



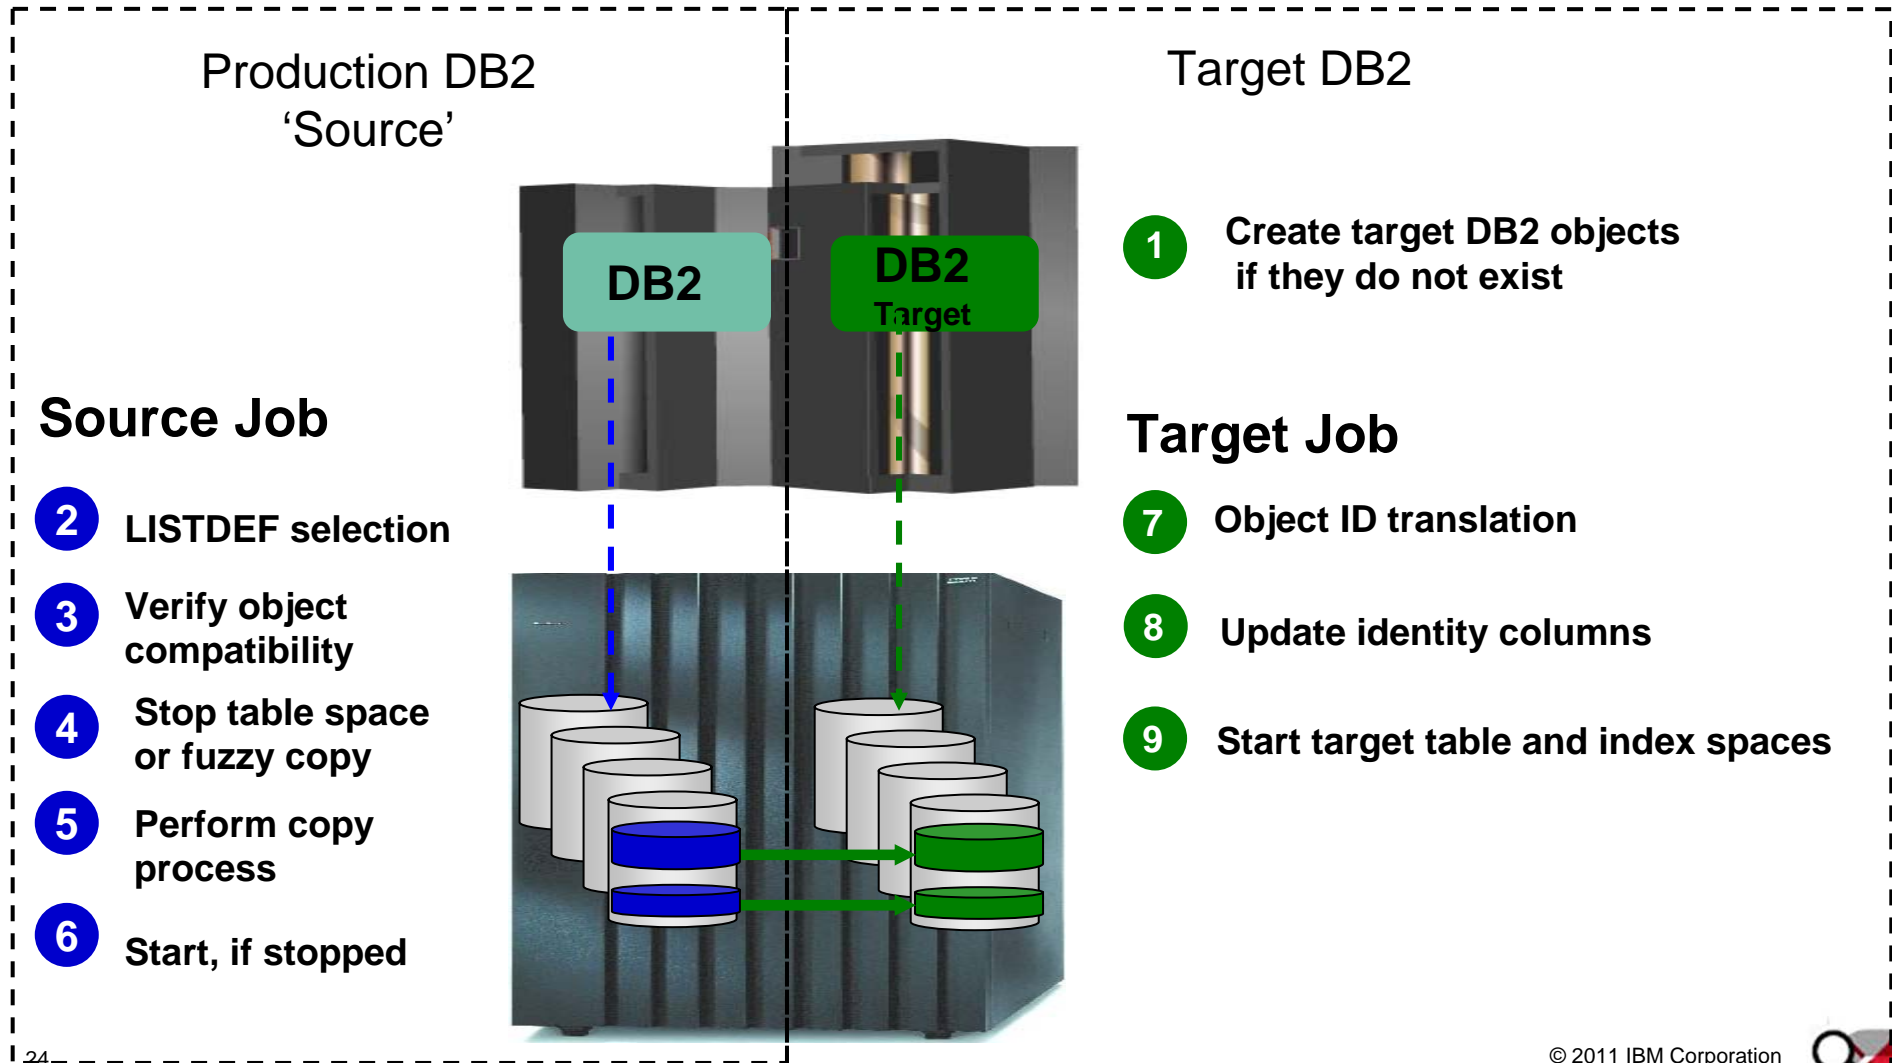
# Refresh DB2 Table and Index Spaces Using DB2 Cloning Tool

## 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 and target system meta-data**



# DB2 Table and Index Space Refresh Steps





# Compatibility

- Determines compatibility between the source and target DB2 objects:
  - Qualifiers
  - Object names
  - Object attributes
  - Column attributes
  - Object IDs
  - DSSIZE
  - Buffer pool size
  - CCSID
  - And others



# Data Masking

- Option to mask one or more columns during the table space refresh process
- The masking changes are made during OBID translation step based on masking rules that are enabled during the copy
- All referential integrity columns can have same masking function applied
  - Examples of fields that a user might change are US--- >Social Security numbers< --- credit card numbers, names and addresses



# Data Masking

- Types of data masking functions are provided, such as:
  - *PATTERN RULE* (*Sir | Mr*) *Bill*
  - *SCRAMBLE RULE* , *SCRAMBLE(FIELD)*
  - *USER EXIT RULE* , *FIELD = USER\_EXIT()*
  - *STATIC RULE* , *FIELD = CONSTANT VALUE*
  - *MASK RULE* , *FIELD = [a-z0-9]\*10*
  - *RANDOM RULE* , *FIELD = RAND(1, 100)*
  - *SEQUENCE RULE* , *SEQ(1, 1)*
  - *CURRENT DATE* , *CURRENT TIME* , *CURRENT TIMESTAMP RULES*
  - *CURRENT USER RULE*



## 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



# Q & A

