



International DB2 Users Group



## Big Data Disaster Recovery Performance

1221 - F04 - May 13, 2014 - 4:30 PM - 05:30 PM

David Beulke  
Dave@davebeulke.com Twitter:  
@DBeulke  
www.davebeulke.com/blog



International DB2 Users Group



2

### Big Data Disaster Recovery- Abstract

Big Data Disaster Recovery can be a big problem. In this session you will learn the latest techniques and design architectures to provide the best Big Data disaster recovery performance. The various hardware and software techniques will be discussed highlighting the Flash Copy and replication procedures.

Performance and system availability have always been paramount and today's Big Data systems make your design architectures, hardware and software choices critical. Understand the architecture choices hardware and software settings that will help you avoid problems and keep your Big Data application available. Learn from real world experience and hear the facts so you avoid the issues with your Big Data disaster recovery performance.

- Understand the architecture choices hardware and software settings that will help you avoid problems and keep your Big Data application viable.
- Learn from real world experience and hear the facts so you avoid the issues with your Big Data disaster recovery performance.

DaveBeulke.Com



**IDUG**  
Leading the DB2 User  
Community since 1988

## International DB2 Users Group


**Dave@davebeulke.com**

- Member of the inaugural IBM DB2 Information Champions
- One of 45 IBM DB2 Gold Consultant Worldwide
- Past President of International DB2 Users Group - IDUG
- Best speaker at CMG conference & former TDWI instructor
- Former Co-Author of certification tests
  - DB2 DBA Certification test
  - IBM Business Intelligence certification test
- Former Columnist for IBM Data Management Magazine

**Weekly Performance Tips:**  
[www.DaveBeulke.com](http://www.DaveBeulke.com)


<b>Consulting</b>	<b>Teaching Educational Seminars</b>
• CPU Demand Reduction Guaranteed!	• DB2 Version 11 Transition
• DB2 Performance Review	• DB2 Performance for Java Developers
• DW & Database Design Review	• Data Warehousing Designs for Performance
• Security Audit & Compliance	• How to Do a Performance Review
• DB2 11 Migration Assistance	• Data Studio and pureQuery

- Extensive experience in VLDB databases, DW design and performance
  - Working with DB2 on z/OS since V1.2
  - Working with DB2 on LUW since OS/2 Extended Edition
  - Designed/implemented first data warehouse in 1988 for E.F. Hutton
  - Working with Java for Syspedia since 2001 – Syspedia - Find, understand and integrate your data faster!



**IDUG**  
Leading the DB2 User  
Community since 1988

## International DB2 Users Group



### DR Status

**76%** of enterprises have declared a disaster or experienced a major business disruption –

- Power Failures 42%
- Natural Disasters 33%
- IT Hardware Failures 31%


April 3, 2008 Building the case for Disaster Recovery Spending – Forrester

**87%** of enterprises have indicated that improving DR capabilities was critical

- Cost of Downtime 56%
- Improving Mission Critical Availability 52%
- Requirements to Stay Online 24/7 48%
- Increased Risk 44%

June 5, 2008 The Forrester Wave: Disaster Recovery Services Providers Q2 2008 -Forrester

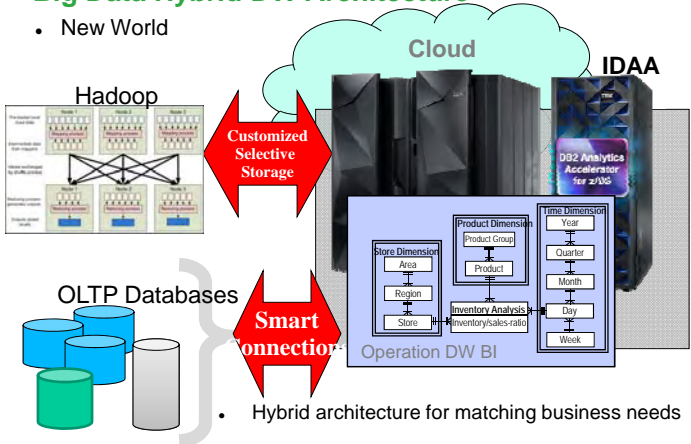
DaveBeulke.Com


**IDUG**  
Leading the DB2 User Community since 1988

 International DB2 Users Group #IDUG

## Big Data Hybrid DW Architecture


- New World



The diagram illustrates a hybrid data warehouse architecture. On the left, 'Hadoop' and 'OLTP Databases' are shown. Red arrows labeled 'Customized Selective Storage' and 'Smart connection' point from these components to a central 'Cloud' environment. The cloud contains server racks and a component labeled 'IDAA DB2 Analytics Accelerator for z/OS'. Below the cloud, a 'Dimensional Model' is shown with 'Store Dimension' (Area, Region, Store) and 'Time Dimension' (Year, Quarter, Month, Day, Week). A 'Product Dimension' (Product Group, Product) is also shown, with 'Inventory Analysis' and 'Inventory/sales-ratio' metrics. The bottom part of the diagram is labeled 'Operation DW BI'.

- Hybrid architecture for matching business needs


DaveBeulke.Com


**IDUG**  
Leading the DB2 User Community since 1988

 International DB2 Users Group #IDUG

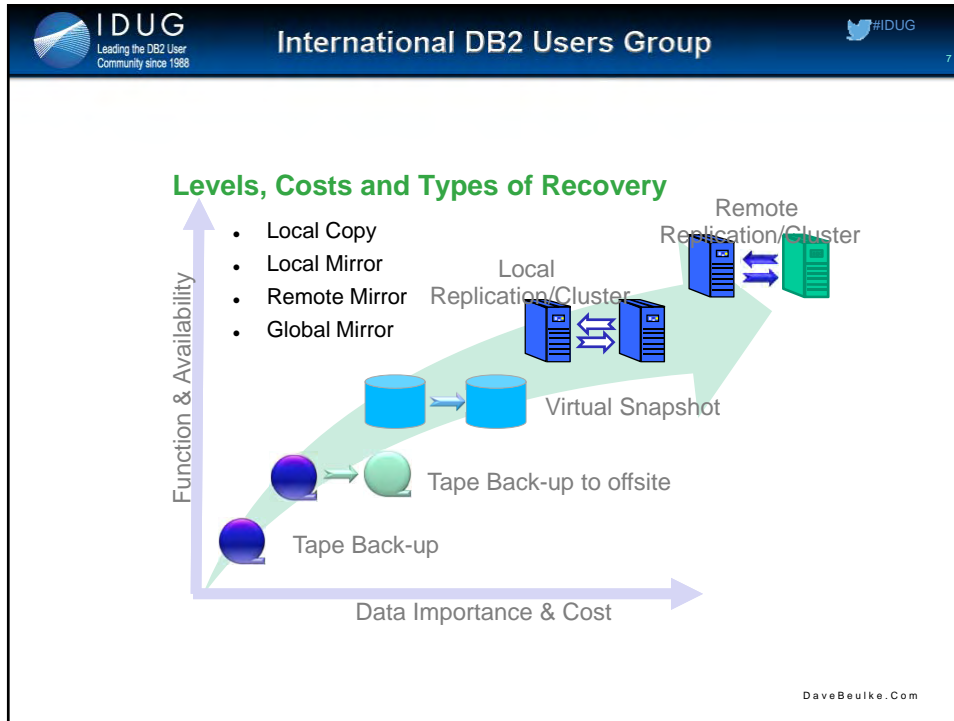
## Two Basic types of DR

- DR by duplication
  - Duplicate Cluster or Mirrored VM
  - Performance ≠ Near zero time?
    - Recovery Time Objective (RTO), Recovery Point Objective (RPO)
  - Very Costly
    - 80+% cost of current environment
      - Licensing/Capacity issues
    - Most need remote location
- DR by backup
  - Tape or Image backup
  - Performance slower - days or ??
  - Costs are cheaper but at what price
    - Server/Image synchronization/rebuild time



The diagram shows a blue triangle representing a trade-off between Cost, Risk, and Performance. The text 'Pick 2' is written below the triangle, indicating that only two of these three factors can be optimized simultaneously.

DaveBeulke.Com



International DB2 Users Group

#IDUG

8


### Linux, UNIX & Window systems are bullet proof

- DB2 HADR provides state-of-art recovery
  - Options from replication to no-downtime ever!
  - Production "hot" server → Read-only standby "cold" server →
  - Hot → Hot -Full HADR for full production
    - Ready to take over at any moment without missing any transaction

Remote Replication/Cluster

Check out any of the DB2 HADR sessions for details

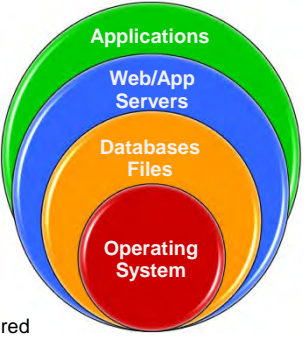
DaveBeulke.Com


**IDUG**  
Leading the DB2 User Community since 1988


 International DB2 Users Group #IDUG

## Tier strategies for different types of data

- Measuring the RTO
- Server(s) environment
  - OS/DBMS/Data/Configurations
- VM Replication
  - Capacity Requirements
  - CPU/Memory/Storage
  - Consolidated DR environment
- Big Data: Back to fundamentals
  - How critical/fast/protected/recovered
  - Time based or content based criteria



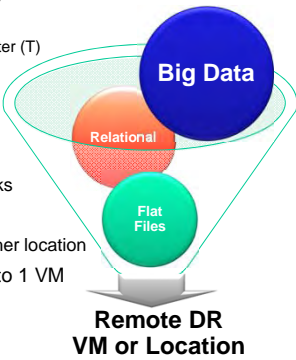
DaveBeulke.Com


**IDUG**  
Leading the DB2 User Community since 1988


 International DB2 Users Group #IDUG

## Big Data provides special RTO challenges

- RTO – Recovery Time Objective – do the numbers
  - The formula for calculating optimal RTO
  - Maximal RTO = minutes(T, S/B)
    - Processing time at the backup data center (T)
    - Resulting data set size (S)
    - Effective available bandwidth (B)
- Big Data - No data compression
  - WAN accelerators can help but big bucks
- Bandwidth considerations
  - Daily ingest rates to other DR box or other location
- Security – challenging to consolidate into 1 VM
  - Cross Realm Kerberos is a challenge




DaveBeulke.Com


**IDUG**  
Leading the DB2 User  
Community since 1988


 #IDUG 11

## International DB2 Users Group

### Big Data – same problems just bigger

- DR problems are constant, big or small
  - Use standard tools helps
  - Business Impact Analysis is still required
- Problem is scope, size (amount/occurrences)
- Use Standard Practices-no overhaul necessary only modifications?
  - Business/Process Impact(s)
    - Of course it is **all** Mission Critical 
  - RTOs and RPOs – How many parts within your Big Data?
    - How many parent process dependencies within each part
    - Quantitative Impact – Financial/social/situational
- Sub-optimal solutions/Alternatives
  - New technology ready? – Cloud and VMs
  - Use 80/20 what is truly critical? Where is the business value?

DaveBeulke.Com

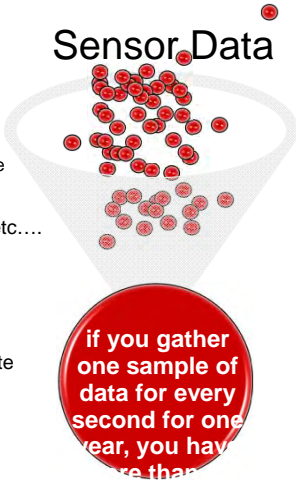

**IDUG**  
Leading the DB2 User  
Community since 1988

 #IDUG 12

## International DB2 Users Group

### Level of Security

- Technology bandwidth –
  - Regulatory drivers
  - Use it for DR and
    - Drive business growth competitive
- Security impacts
  - PII, HIPPA, Masking, Encryption etc....
- Framework for business continuity
  - Physical to VM
  - VM to Physical
  - VM to VM
  - Logical sync point - Local or remote
  - File or transaction
  - Requirements/Money/Technology



DaveBeulke.Com

**IDUG** Leading the DB2 User Community since 1988 #IDUG 13

### International DB2 Users Group

## Big Data Hybrid DW Architecture

- New World

The diagram illustrates a hybrid data warehouse architecture. On the left, 'Hadoop' and 'OLTP Databases' are shown. A red double-headed arrow labeled 'Customized Selective Storage' connects Hadoop to a central 'Cloud' component. Another red double-headed arrow labeled 'Smart Connection' connects OLTP Databases to the Cloud. The Cloud contains server racks and a box labeled 'IDAA DB2 Analytics Accelerator for z/OS'. Below the Cloud is a 'Operation DW BI' diagram with various dimensions: Store Dimension (Area, Region, Store), Product Dimension (Product Group, Product), and Time Dimension (Year, Quarter, Month, Day, Week). A 'Inventory Analysis' box is also present.

- Hybrid architecture for matching business needs

DaveBeulke.Com

**IDUG** Leading the DB2 User Community since 1988 #IDUG 14

### International DB2 Users Group

## Hadoop's Disaster Recovery - special needs

- Hadoop Three Sections
  - Data, System & Configuration
- Biggest is input data
  - Replicated data to 2 or 3 nodes
  - Best: backup as it comes in
  - Disk, Node, Rack, Site failures
    - Standard DR
- System/Application(s)
  - Frequent regular backups
- Configuration
  - Frequent regular backups

The diagram shows the Hadoop disaster recovery process. At the top, 'Input Backup' is shown with arrows pointing to a data source. A blue arrow points down to three nodes: Node 1, Node 2, and Node 3. Each node has a 'Mapping process' box. Below the mapping processes, a network of arrows shows data being distributed to 'Reducing process' boxes on each node. Each reducing process box is connected to a blue box representing the final output or storage.

DaveBeulke.Com

**IDUG**  
Leading the DB2 User  
Community since 1988

**International DB2 Users Group**

#IDUG

## Hadoop's Disaster Recovery - special needs

- Community is working on HDFS Snapshot capabilities
  - MapRfs – Amazon
    - Provides snapshots
  - Namenode single point
    - Dual nodes heartbeat syn
- Map Reduce Output(s)
  - Standard DR Backups
  - Transform to Archive
  - Standard reports/files

**Best backup Hadoop:  
Do only Config, Inputs, Outputs**

DaveBeulke.Com

**IDUG**  
Leading the DB2 User  
Community since 1988

**International DB2 Users Group**


#IDUG

## Big Data is a different story

- Bottom Line: Business Continuation - 4 "W"s
  - What systems are synchronized with Big Data
  - When were they synchronized?
  - Where are the data source going to be recovered?
  - Why is this - the best point? the correct data amount?
- These are the keys to all DR strategies
  - Big Data just makes it more intense!

DaveBeulke.Com



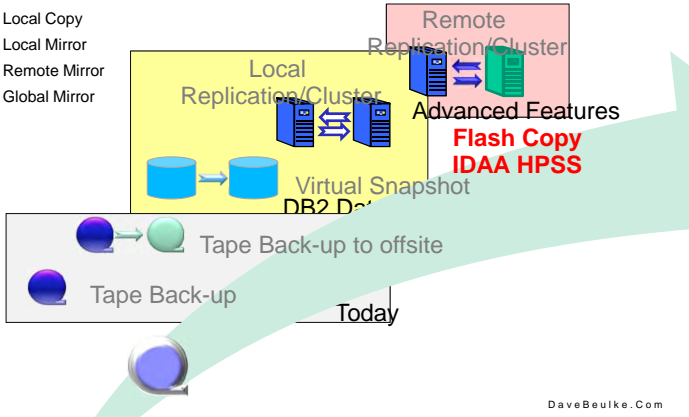

**IDUG**  
Leading the DB2 User Community since 1988

 #IDUG 17

## International DB2 Users Group


### DB2 z/OS – the BEST easiest Big Data solution

- As flexible as small LUW systems
  - Just more complex, transactions, interfaces and harder to sync
- Local Copy
- Local Mirror
- Remote Mirror
- Global Mirror



The diagram illustrates a data replication architecture. It shows a 'Local Replication/Cluster' box containing two server icons and a 'Remote Replication/Cluster' box containing two server icons. A double-headed arrow connects them. Below the local cluster is a 'Virtual Snapshot' box with a cylinder icon. A large green arrow points from the 'Today' box (containing 'Tape Back-up' and 'Tape Back-up to offsite') towards the 'Advanced Features' box (containing 'Flash Copy', 'IDAA', and 'HPSS').

DaveBeulke.Com


**IDUG**  
Leading the DB2 User Community since 1988


 #IDUG 17


## International DB2 Users Group

### IDAA HPSS and DB2 Flash Copy

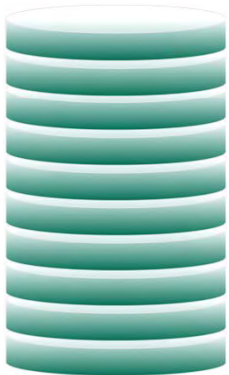
- DB2 IDAA – HPSS
  - DB2 subsystems
  - High Perf. Super Saver
  - IDAA accelerator(s)
- DB2 IDAA HPSS
  - Eliminates extra data copies
  - New HPSS – many features
  - Archive Data into IDAA
  - Protect/prevent data issues
  - Data to multiple accelerators
- Protected DR

- Flash Copy
  - Snapshot, TimeFinder, Clone
- Many different options
  - DB2 system replication
  - Partial system replication
  - TS IX Level Replication
  - Object replication
- DB2 Cloning Tool
  - Helps with all setup aspects
  - Templates – Masks
  - Log

 **IDUG**  
Leading the DB2 User  
Community since 1988


**International DB2 Users Group**  19


## DB2 IDAA High Performance Super Saver-HPSS



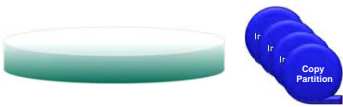
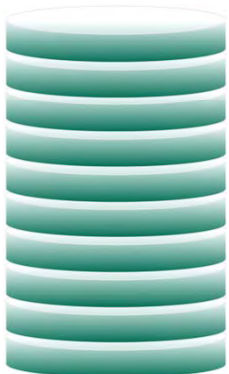
- Big Data Partitioned Tables
- Migrate partition(s) to IDAA
  - One or All partitions
- New "PRO" TS part status
- Image Copy status to protect backup
- Saves the space in primary DB2

DaveBeulke.Com

 **IDUG**  
Leading the DB2 User  
Community since 1988

**International DB2 Users Group**  20

## DB2 IDAA High Performance Super Saver-HPSS



- Up to 4 image copies of data
- "PRO" status
  - Prevents SQL data changes

DaveBeulke.Com

**IDUG**  
Leading the DB2 User  
Community since 1988

International DB2 Users Group

#IDUG

21

Push it into multiple IDAAs - Eliminate the I/O

Protected TS Status - PRO

DB2 Analytics Accelerator for z/OS

DB2 Analytics Accelerator for z/OS

Flexible Templates

DaveBeulke.Com

**IDUG**  
Leading the DB2 User  
Community since 1988

International DB2 Users Group

#IDUG

22


DB2 HPSS IDAA

DB2

Copy Partition


DB2 Analytics Accelerator for z/OS

DaveBeulke.Com




**IDUG**  
Leading the DB2 User  
Community since 1988

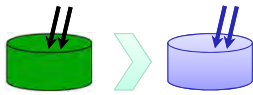
International DB2 Users Group

 #IDUG  
23


## Other options - Many quick Cloning DR options

- Complete DB2 Subsystem Clone
  - DB2 Subsystem
  - System Level - Backup System
- Partial DB2 Subsystem Clone -PSSC
  - Selected storage volumes
- Skeleton DB2 Subsystem Clone
  - DB2 directory, catalog, BSDS, and active logs only
  - Populated by datasets
- DB2 Individual Table Space Clones
  - By dataset/object definition






DaveBeulke.Com



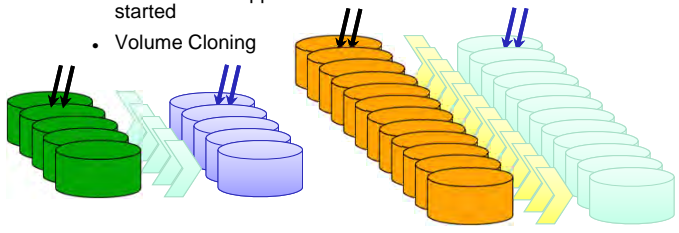
**IDUG**  
Leading the DB2 User  
Community since 1988

International DB2 Users Group


 #IDUG  
24

## Cloning is based on Storage Management

- Storage group management
  - Leveraging Storage technology
- **Offline Mode**
  - DB2 BACKUP SYSTEM – DB2 Utility Command
    - Using DB2 BACKUP SYSTEM is stopped and started
    - Volume Cloning



DaveBeulke.Com




**IDUG**  
Leading the DB2 User  
Community since 1988

**International DB2 Users Group**


#IDUG

## DR Subsystem Cloning Options

- Online Mode - non disruptive
  - DB2 BACKUP SYSTEM is online
    - using DB2 Log Suspend and Resume
  - Via consistent
    - IBM FlashCopy,
    - STK SnapShot
    - EMC TimeFinder
  - Via consistent split or break mirror.
    - Can be same LPAR
    - Must be a ***different DB2***



DaveBeulke.Com



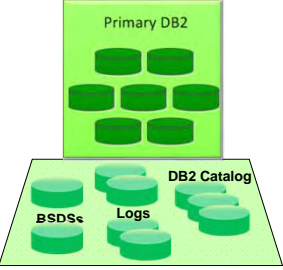
**IDUG**  
Leading the DB2 User  
Community since 1988

**International DB2 Users Group**

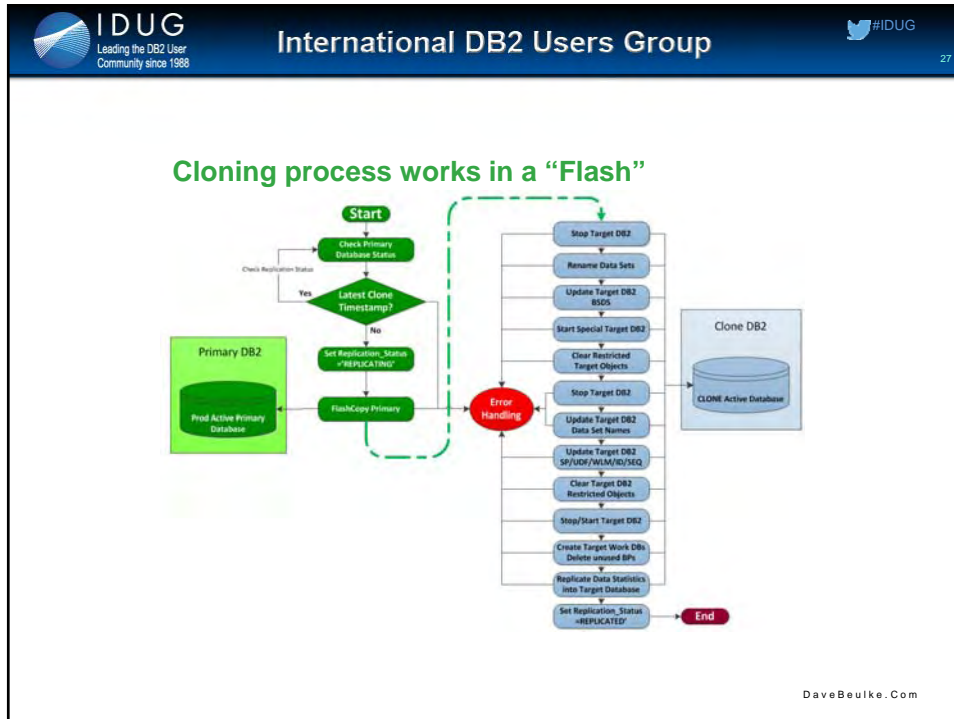
#IDUG

## Flash Copy DR DB2 System Components

- Clones DB2 System Objects
  - DB2 directory
  - DB2 catalog
  - BSDSs and Active logs
- Updates the DB2 internals
  - Cloned DB2 subsystem can be accessed from the same z/OS LPAR
- Cloning renames and catalogs the data sets on the cloned volumes, fixes the volume internals



DaveBeulke.Com



DaveBeulke.Com

- DR Cloning Flexibility Options**
- Target can be different configuration than original
    - DB2 data-sharing groups to non data-sharing DB2 subsystems
  
  - TS cloning provides segmentation of important data
    - Only do the TS that are absolutely needed
  
  - Customizable Flexibility Options
    - Target DB2 can be configured many different ways
      - Reporting Subsystem
      - Archive system

DaveBeulke.Com

**IDUG** Leading the DB2 User Community since 1988 **International DB2 Users Group** #IDUG 29

### Special Considerations

- Identity columns
- PBG
- LOBS
- Table space reordered row format
- LONGVAR vs. VARCHAR
- Just VARCHAR
- Objects created using DEFINE NO . .
- Clone tables


DaveBeulke.Com


**IDUG** Leading the DB2 User Community since 1988 **International DB2 Users Group** #IDUG 30

### Summary

- Performance RTO is based on amount of I/O
- Petabytes can be done
  - Just a matter of the technology and time
  - Bandwidth/Channels/Parallel/PAVs
- RPO – Quiet Time?
  - Not Quiet – Slow
  - System/Appl Sync
- Minimize I/Os - Data
  - Compress
  - Eliminate
  - Don't duplicate data

DaveBeulke.Com



**International DB2 Users Group**  31

### Resources -

- **Manuals**
  - IBM DB2 Cloning Tool – User Guide Version 3.1
  - IBM System Storage DS8000: Command-Line Interface User’s Guide, GC53-1127
  - IBM System Storage DS8000: Architecture and Implementation, SG24-8886
- **Redbook –**
  - Managing DB2 z/OS Utilities with DB2 Tools Solution Packs – SG24-8046
  - IBM System Storage DS8000 Copy Services for IBM System z SG24-6787-06

DaveBeulke.Com



**International DB2 Users Group** 

## Big Data Disaster Recovery Performance

1221 - F04 – May 13, 2014

David Beulke  
Dave@davebeulke.com  
www.davebeulke.com/blog



*Please fill out your session evaluation before leaving!*

