



CDP on System i High Availability has Changed Forever

Barry Kirksey

Principal Manager

Lakeview Technology

IBM® System i™ Education Series

Get the Facts!

September 14, 2006



Today's Topics

- Defining Continuous Data Protection (CDP)
- How CDP works
- Benefits of CDP
- CDP and HA
- Introducing MIMIX CDP™





Defining Continuous Data Protection (CDP)



Get the Facts. Innovations in System i5.



What is Continuous Data Protection (CDP)?

Wikipedia

Continuous data protection (CDP), refers to backup of data by automatically **saving a copy of every change** made to that data, essentially capturing every version of the data that the user saves. It allows the user or administrator to **restore data to any point in time**.



Aberdeen

CDP **captures every data write** and transfers the new data immediately to a protection server disk. There is also the ability to recover to any point-in-time (APIT) by undoing the updates to **revert back to any time**.

Computerworld

CDP systems **automatically record every write** made to disk and time-stamp each of these I/O events, enabling you to **recover to any point in time** and restore the data to how it was before the corruption occurred.



Another Definition

CDP is a methodology that *continuously captures or tracks data modifications* and stores changes independent of the primary data, *enabling recovery points from any point in the past*. CDP systems may be block-, file- or application based and can provide fine granularities of restorable objects to highly variable recovery points.

- *Source: Storage Networking Industry Association - Data Management Forum - CDP Special Interest Group working definition.*



Let's Simplify Things

- Common Elements of CDP
 - Track all the changes
 - Recover to any point in time



- VCR or DVR for your data



Continuous Data Protection



- CDP is not a single technology
 - It is a Data Protection concept or process
 - Implemented in a variety of ways
- Blending of characteristics from backup, replication and snapshot technologies



Different Types of CDP

The World According to Gartner*

- True-CDP
 - Capture EVERY data write
 - Transfer immediately to secondary disk
 - Ability to recover to ANY point in time
- Near-CDP
 - Does not capture every write
 - Specific recovery points
 - May only transmit changes periodically

Gartner SOURCE: Dave Russell

"Continuous Data Protection: What It Is and Why It Is Needed"

June 8, 2006

CDP Architecture Overview



- Host Based
 - Agents communicate with a central server
 - Centrally store all recovery data
 - Well suited to a multi-server environment
- Network Based
 - Combination of dedicated appliance and agents
 - Journal all writes
 - Well suited to networked storage environments


STORAGE
Managing the information that drives the enterprise

SOURCE: Jerome Wendt :
"Nonstop Data Protection"
October 20, 2004

Components of Availability

RPO - Recovery Point Objective

How much data loss can you tolerate?



	+2.688
0	+5.000
1	+1.500
0	+1.125
0	+1.062



RTO - Recovery Time Objective

How long can you afford to not have systems access to data and applications?

How Does CDP Impact These Metrics?

RPO - Recovery Point Objective

CDP drives the RTO down
toward zero



	+2.688
0	+5.000
1	+1.500
0	+1.125
0	+1.062



RTO - Recovery Time Objective

CDP enables the recovery time to
be “variable”

CDP: Changing the Storage Landscape

- Stand-alone products generally becoming part of larger product suites
- Vendor Consolidation is already taking place
 - 7+ CDP vendor acquisitions in the last year
 - Multiple partnership announcements
 - Co-marketing
 - OEM/Private Label
- Deployments lag behind
 - Market penetration is still very low
 - Mostly impacting SMB



What is the market for CDP



- Most compelling for applications
 - that have a high rate of data change
 - where a loss of twenty-four hours of data could mortally wound the business
 - that take up a fairly large amount of storage
 - where downtime will have significant impact on the company

CDP Deployment Trends

Initial Critical Applications

- Electronic mail
 - ERP
 - CRM
 - “Files”
-
- Complementing existing data protection
 - Coordinated with existing solutions or practices
 - Replacement for specific RPO/RTO challenges





How CDP Works



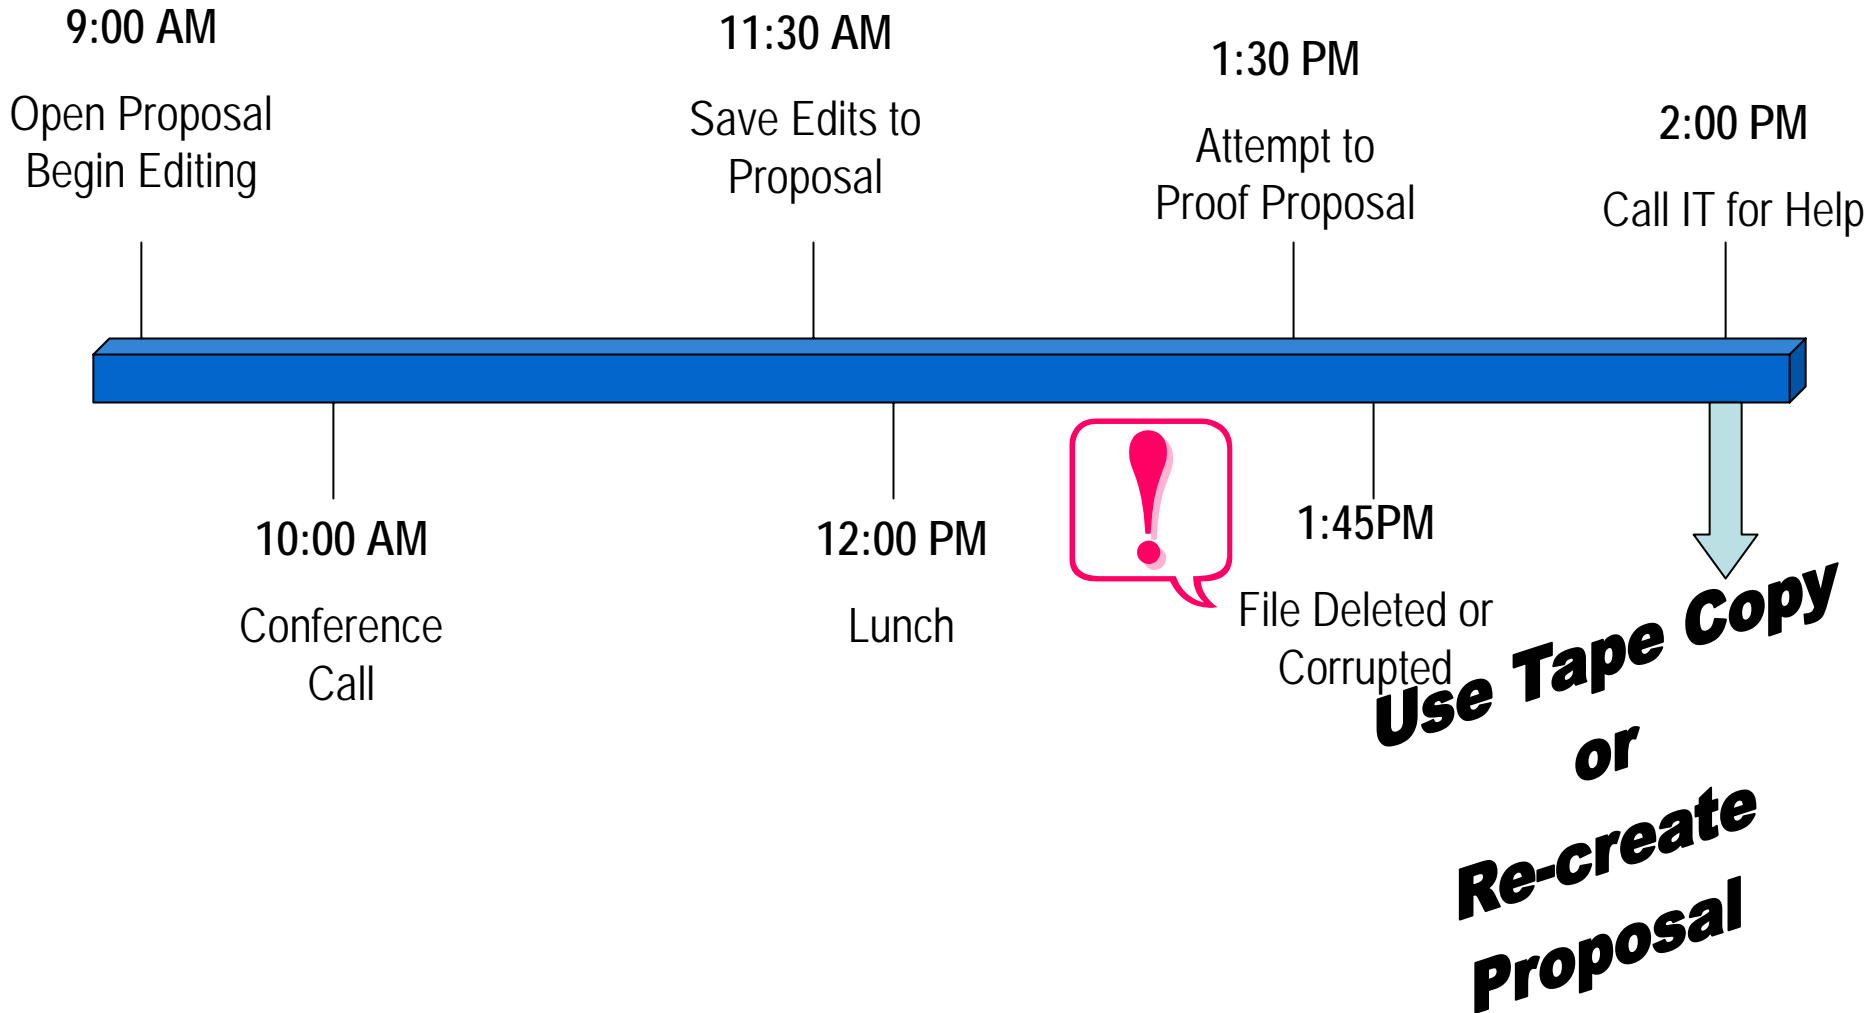
Get the Facts. Innovations in System i5.

Continuous Data Protection Basics


- Simple Concept
 - Capture every change
 - Recovery to any point in time
- Flexibility
 - File level
 - Volume level
 - Entire Server



Recovery without CDP



Recovery with CDP



**Use CDP to
restore copy
of your file
from 11:30**



Benefits of CDP



Benefits of CDP

- The ability to recover to virtually any point in time
 - Go back minutes, hours, days or even weeks
 - Flexibility of restoring defined volumes or single files
- Recovery Points are virtual
 - Created on demand when needed
 - Not predetermined or periodic
- Faster recovery
 - No tape based restore
 - Disk image can be used very quickly



More Benefits of CDP

- Reduce the impact of the “backup window”
- Reduce potential data loss
- Meet strict Recovery Point Objectives



Benefits of CDP over Snapshots

- Data changes are recorded continually, as opposed to halting an application's activity to create the data snap.
- Data changes are stored incrementally rather than storing numerous data images.





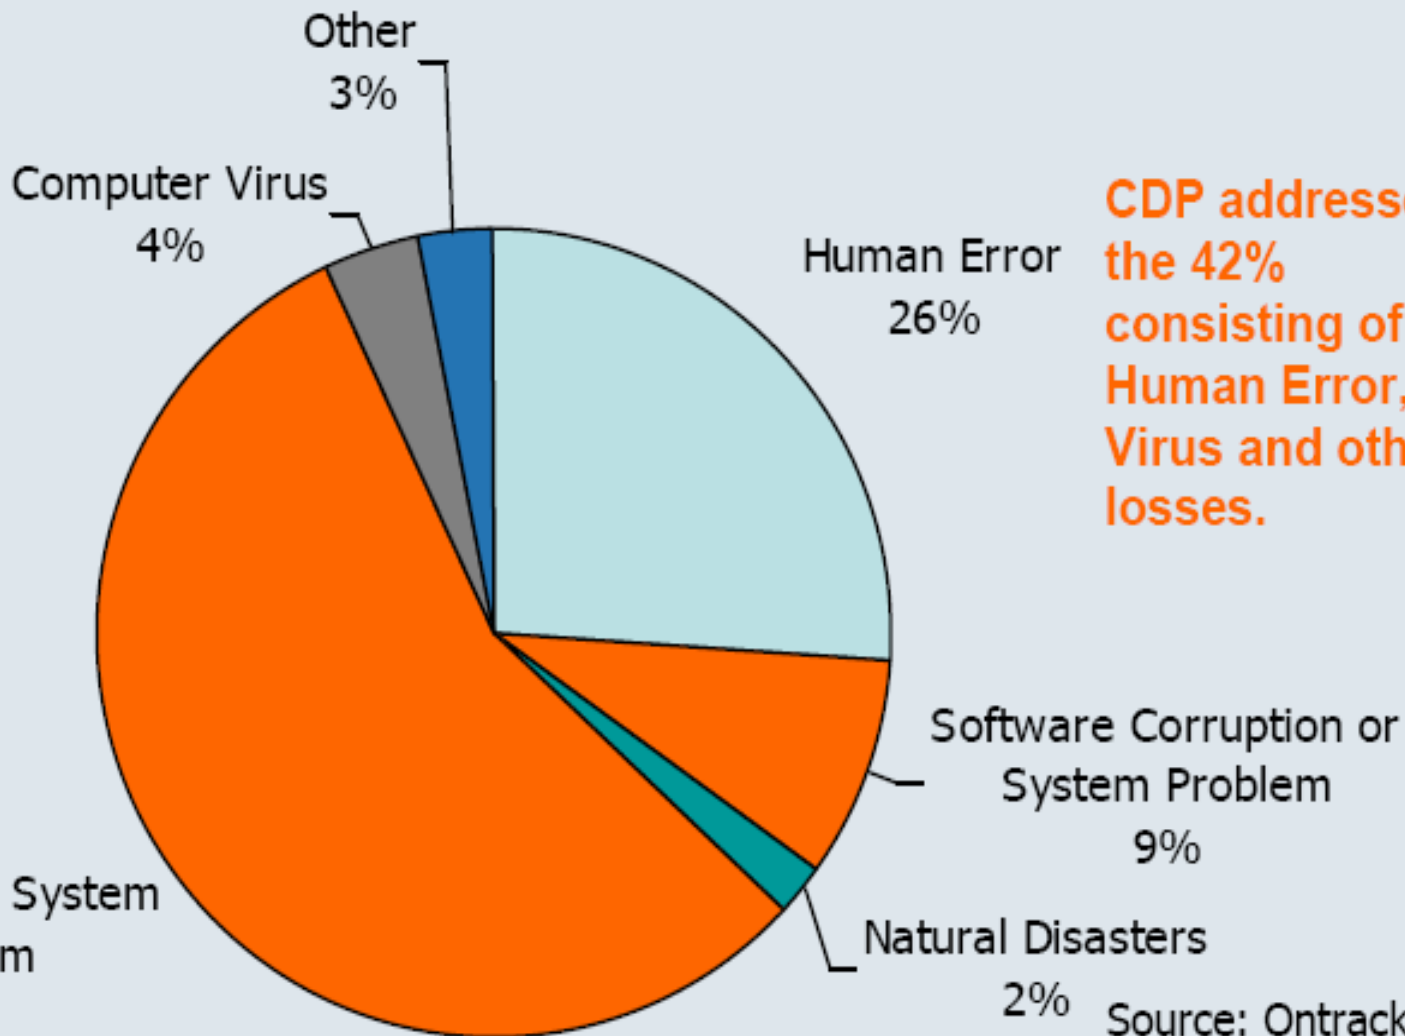
CDP and HA





Causes of Data Loss or Downtime

Hardware or physical protection is good & can recover 56% of the problems.



CDP addresses the 42% consisting of Human Error, Virus and other losses.

Hardware or System Problem
56%

Source: Ontrack, 2004



CDP & HA – Powerful Combination

- Business data is vulnerable
 - Accidental deletion
 - Malicious corruption
- High Availability Solutions may not offer protection
 - Replicate deletions
 - Replicate corrupted files

Gartner

By 2011, some form of CDP will be Deployed in 80% of Fortune 2000.

- Dave Russell, Research Director

Aberdeen Group

CDP has enormous potential as a strategic piece to future data protection policies.

- Sonia Lelii, Senior Research Analyst



Introducing MIMIX CDP[™]



Get the Facts. Innovations in System i5.



MIMIX CDP™ – A first for IBM System i

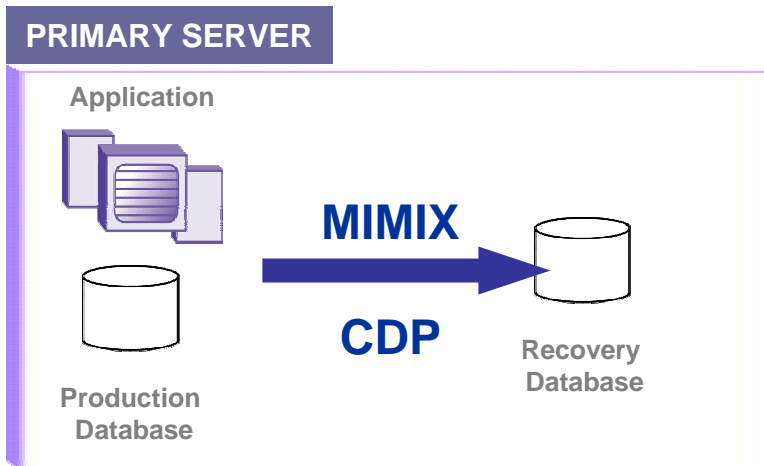
- New feature of MIMIX ha1 and MIMIX ha Lite
- No additional cost
- Integrated User Interface
- Flexible recovery options

MIMIX®
ha Lite™

MIMIX®
ha1™

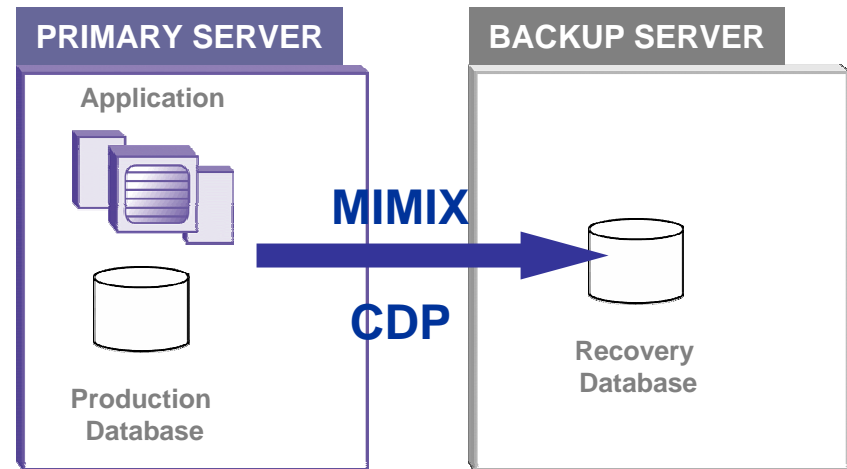
How MIMIX CDP Can Be Implemented

Single Server MIMIX CDP



Enables Recovery of Data on the Same Server

Dual Server MIMIX CDP Combined w/ HA / DR



Enables Recovery of Data on a Second Server and utilized data for HA / DR as needed



Frequently Asked Questions

FAQ

Q: Is MIMIX CDP a true-CDP or a near-CDP solution.

A: Gartner defines true-CDP solutions as those which enable any point-in-time recovery of data.

MIMIX CDP enables any point-in-time recovery of data and is a true CDP solution.



Frequently Asked Questions

FAQ

Q: Is MIMIX CDP a stand-alone product?

A: MIMIX CDP is a feature of MIMIX ha1 and MIMIX ha Lite.

MIMIX CDP is not being launched as a stand-alone product because our market research has indicated that customers want CDP to be integrated into their High Availability or Backup and Recovery Solutions.



Frequently Asked Questions

FAQ

Q: How large can a users set the Optimized Recovery Window ?

A: MIMIX CDP currently supports an Optimized Recovery Window of 60 days.



Frequently Asked Questions

FAQ

Q: Is there any additional hardware required to use MIMIX CDP?

A: MIMIX CDP will require additional DASD to store the MIMIX Recovery Log. The size of this log will depend on the size of the MIMIX Optimized Recovery Window and the amount of user changes that occur.

The impact of CPU is practically zero.

Summary

- CDP Definitions
- A Look at How CDP Works
- Benefits of Combining CDP and HA
- Introduction to MIMIX CDP™





Thank You

Barry Kirksey

Principal Manager

Lakeview Technology

kirkseyb@lakeviewtech.com

IBM[®] System i[™] Education Series

Get the Facts!

September 14, 2006



Get the Facts. Innovations in System i5.

Next Seminar

Date

September 14, 2006

Topic

“SOA on i - High Availability Makes IT Happen”

**First-generation**

- Requires software to be locally administered on the server.
- Products such as Veritas' Volume Replicator, Softek's Replicator or XOsoft's Enterprise Data Rewinder let you copy data with minimal impact to active application; can optionally configure checkpoints to allow secondary image to be restored or even configured to allow the application to run from the secondary data replica.
- Consider for applications that require direct administration, need to run on secondary data replica or simply require tested, proven and well-supported tools.

Second-generation

- Topio's Data Protection Suite and Mendocino Software's Realtime use a client/server architecture for centralized management and recovery of data.
- Agents communicate with central server.
- These approaches centrally store all of the recovery data using asynchronous IP to move the data from the clients to the main server.
- Consider for multiserver environments where restores need to be up to the minute and completed in minutes.

First-generation

- Products such as FalconStor's IPStor put the intelligence in the storage network with features found on many storage arrays.
- Enabling mirrors and snapshots between any vendors' arrays without need to deploy agents on the servers, allow users to unobtrusively deploy CDP solutions in networked storage environments.
- Consider when all storage can be placed under the control of this type of CDP application and recoveries can be done from the last snapshot or mirror.

Second-generation

- Products like Revivio's CPS-1000 and Alacritus Software's Chronospan.
- Run on switch blades or as network appliances, appear as LUNs to the individual hosts; can use host-based volume management software to allow mirrored writes to primary storage and to the LUNs presented by these products.
- Journal all writes and allow for recoveries from nearly any point in time.
- Consider for networked storage environments but roll out to entire environment slowly after extensive testing.