



IBM System i5™

IBM System i5 Disaster Recovery and/or “High”-Availability

IBS

Dirk Vande Walle

*IBS Technology and Services
Xavier De Cocklaan, 70
B-9830 Sint-Martens-Laetem
Tel. +32 9 280 21 33*

*IBM System i SE
IBM Certified*

Email: Dirk.Vande.Walle@ibsbe.be

IBS

Danny Simon

*IBS Technology and Services
Lewensesteenweg, 643
B-1930 Zaventem
Tel. +32 2 723 91 24*

*IBM System i SE
IBM Certified*

Email: Danny.Simon@ibsbe.be

Simplify your IT.

Who's who?

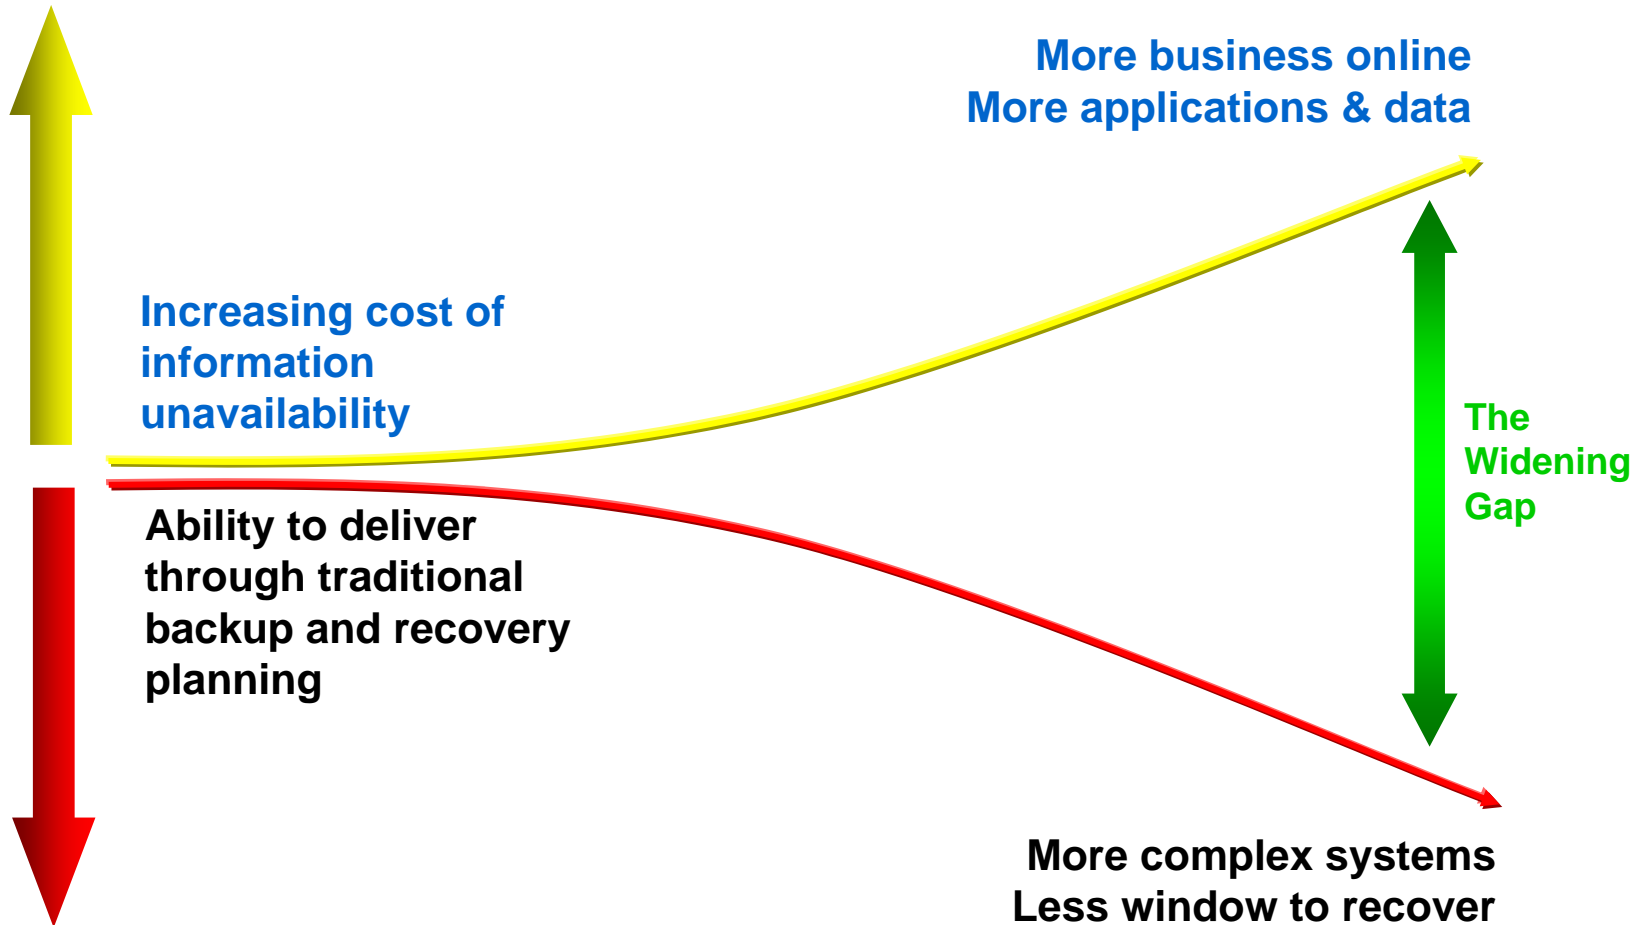


Dirk Vande Walle



Danny Simon

The Availability Business Challenge



Requires continuous information availability – by design

Agenda

- Define Your Business Availability Goals
- Single Server Availability
- Multi Server High Availability
- Application Availability
- Summary



Define Your Business Availability Goals

Impact of “Non”-Availability

- Lost business
 - Impact varies between companies
 - Can be permanently lost business
- Image and publicity
- Fines and penalties
 - Government and industry fines, service level agreements
- Staff costs
 - Pay for inactivity, overtime, recovery actions
- Impact on business decisions
 - Some business intelligence information must be real time

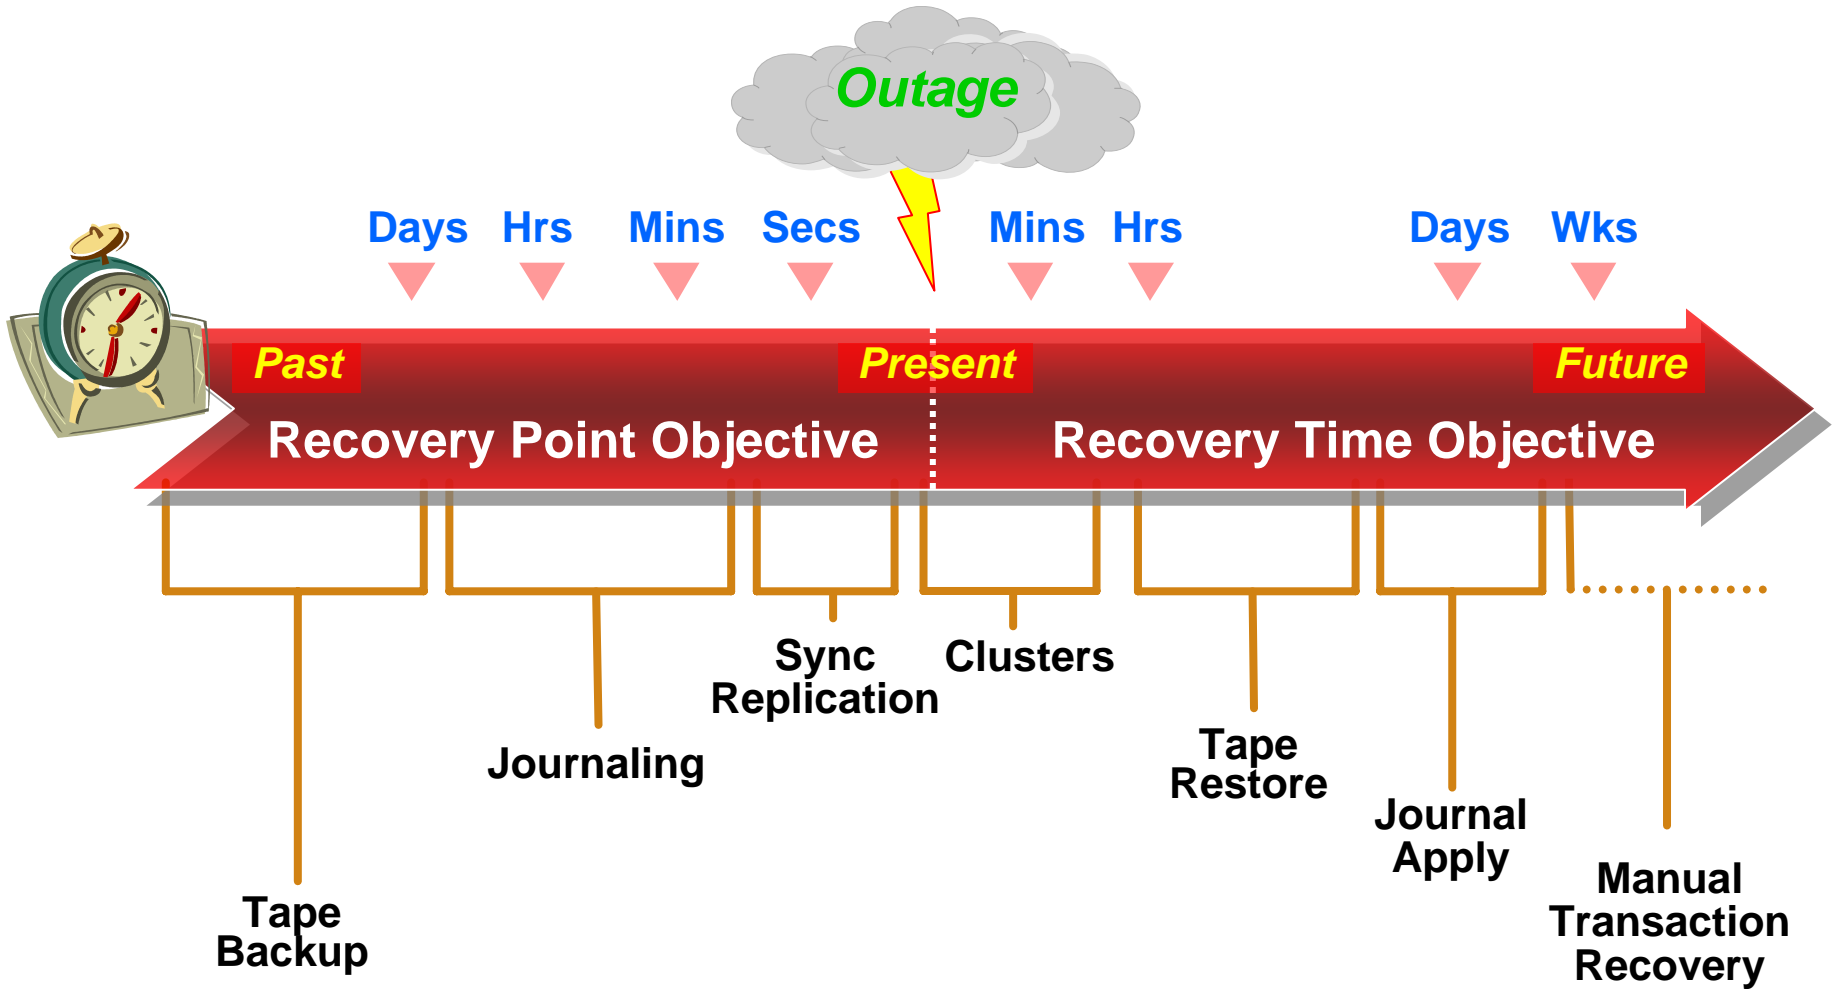


Business Availability Definitions

- **Service Level Agreements (SLA)** establish measurable targets of performance and availability
 - Provides a common understanding of the nature, level, and quality of service required
 - Provides priorities for availability management and investments
 - Eliminates assumptions, aligns business requirements with IT services
- **Recovery Objectives** for applications are key business measurements required to assess availability technologies and design solutions
 - *RPO*: Recovery Point Objective
 - *RTO*: Recovery Time Objective



Meeting Recovery Objectives



Analyze and Define Your Business Availability Goals

Fully identify the set of availability problems that you are attempting to address . . .

1. Backup window reduction
2. Availability for planned outage events
3. Availability for unplanned outage events
4. Recovery from disaster related outage events
5. Workload balancing

- . for each application,
 - for each partition,
 - and for each operating system

Identify All Critical Application Requirements (example)

What Are Your Business Applications Availability Requirements?

| Application | OS | Disaster Recovery | Fault Tolerant | Fault Resilient | High Availability | Highly Reliable | Base | Notes |
|---------------|---------|-------------------|----------------|-----------------|-------------------|-----------------|------|----------------------------------|
| HR | AIX | Desired | | Desired | Actual | | | |
| Financials | i5/OS | Actual | | Desired | | Actual | | |
| Replenishment | i5/OS | Desired | Desired | | Actual | | | Effects i5 business environments |
| Warehouse | i5/OS | Actual | | Desired | | Actual | | |
| Distribution | AIX | Desired | | Desired | | Actual | | Failover in test |
| LDAP | Windows | | Desired | | | Actual | | |
| Sales Audit | i5/OS | Desired | | Desired | Actual | | | |

Example Analysis

Disaster Recovery - A horizontal availability feature that is applicable to any of the AECs. It provides for remote backup of the information system and makes it safe from disasters such as fire, flood, hurricane, power outage, vandalism, or an act of terrorism.

Fault Tolerant - *Continuous* computing with any failure transparent to the user. No degradation in service and continuous 7X24 operation.

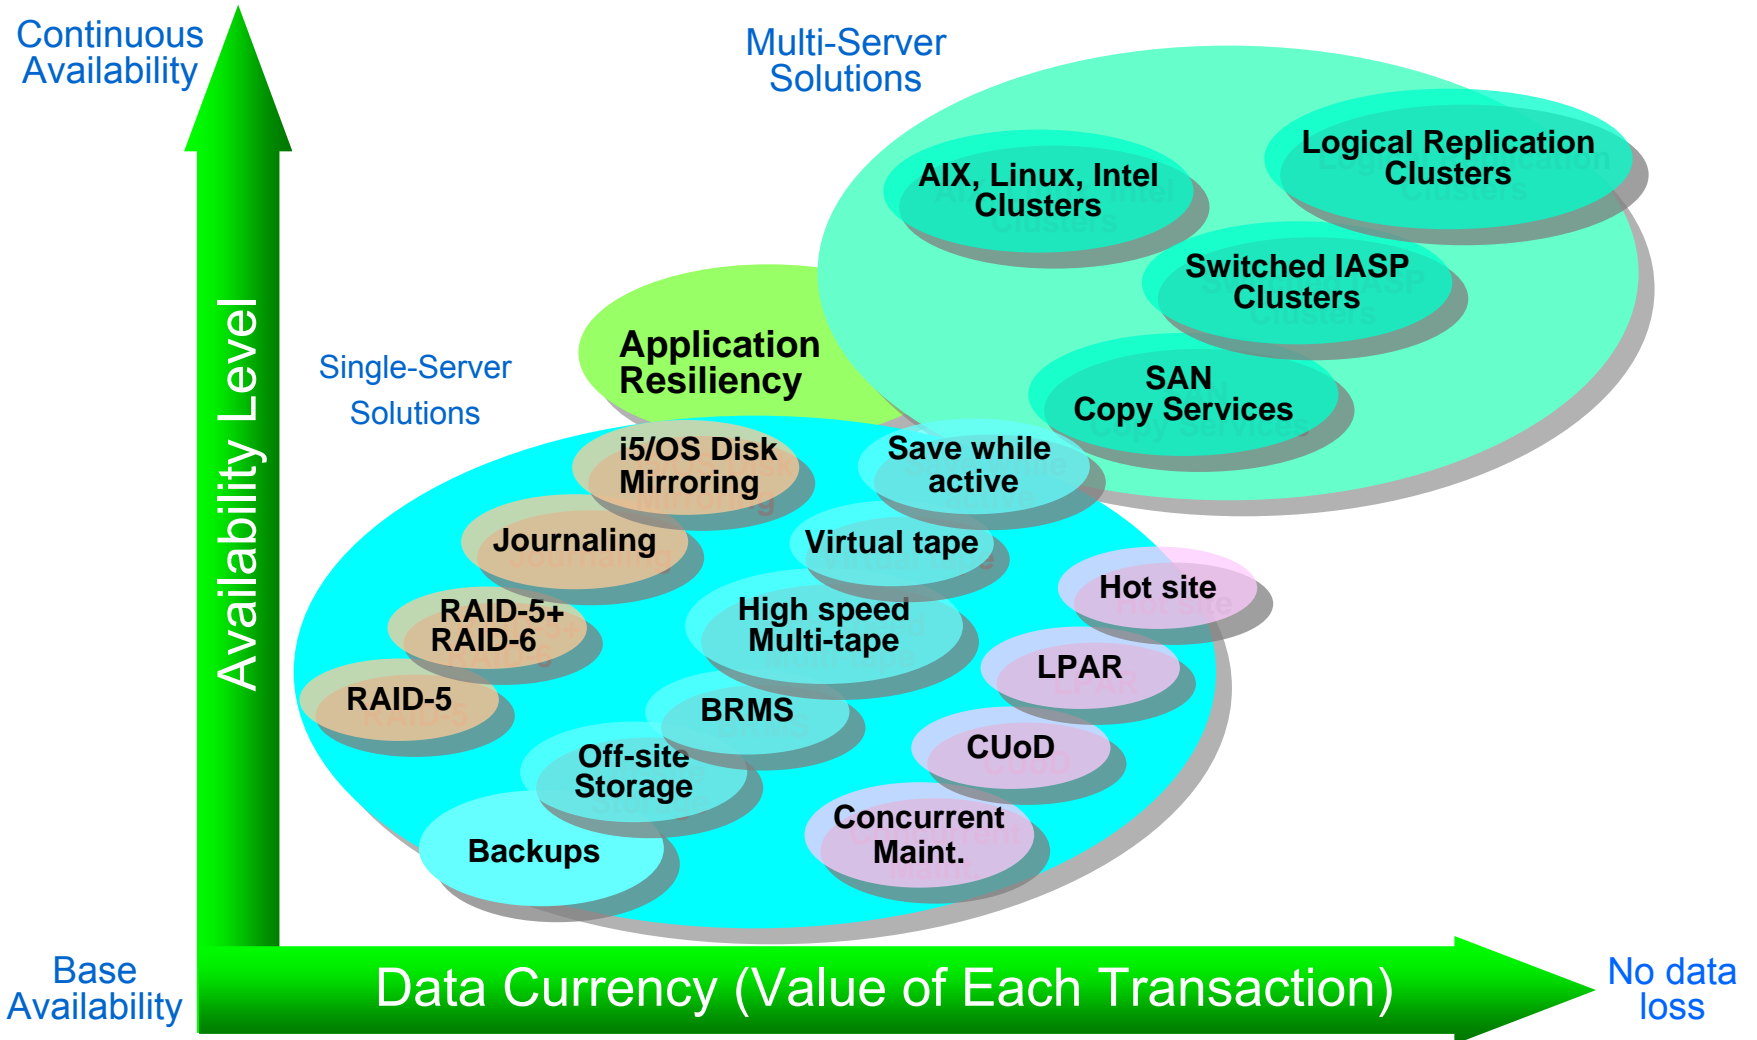
Fault Resilient - Business functions require *uninterrupted computing services during essential or designated time periods*. Users stay online at all times, however current transactions may require restarting. Users performance may experience degradation.

High Availability - Business functions allow *interrupted computing services*, but users can quickly recover from interruptions. They may have to rerun some transactions, and they may experience some performance degradation.

Highly Reliable - Business functions can be interrupted as long as the data integrity is insured. To the user work stops and uncontrolled shutdown occurs.

Base - Business functions that can be interrupted, and the integrity of the data is not essential. To the user work stops and uncontrolled shutdown occurs. Data may be lost or corrupted.

iSeries Availability Building Block Technologies



Single Server Availability

RAS - Reliability, Availability and Serviceability

■ Hardware Service

- *Automatic Failure Notification*
- *Hardware concurrent maintenance*
- *Hot plug cards*
- *IOP reset*
- *Predictive Failure Analysis & Reporting*
- *Dual Hardware Management Console*

■ Storage

- *RAID 5, 5+,6, 10, 51*
- *SAN Multipath*
- *Concurrent disk rebalance and migration*

■ Virtualization

- *Dynamic resource allocation*
- *Workload isolation with LPAR*

■ Power and Cooling

- *Redundant power supplies*
- *Dual power cords*
- *Dedicated UPS interface for monitoring*
- *Redundant cooling*

■ CPU and Memory

- *Chip Kill*
- *Error detection and correction*
- *Memory Scrubbing*
- *Array repair at IPL with redundant array bits*
- *Redundant service processors for 570 and 595*
- *CPU Gard*
- *CPU Sparing*

■ Networking

- *Virtual IP*
- *Route fault tolerance*
- *Load balancing*

■ Application and Data Resiliency

- *Object journaling*
- *Commitment Control*
- *DB re-org while active*
- *Application cluster enablement*

System i5 Storage Protection Choices

System i5 RAID-5 basic

- *In use since 1993*
- *Single copy of cache*

System i5 RAID-6

- *New - dual parity protection*
- *Single copy of cache*

System i5 RAID-5 with Redundant Cache

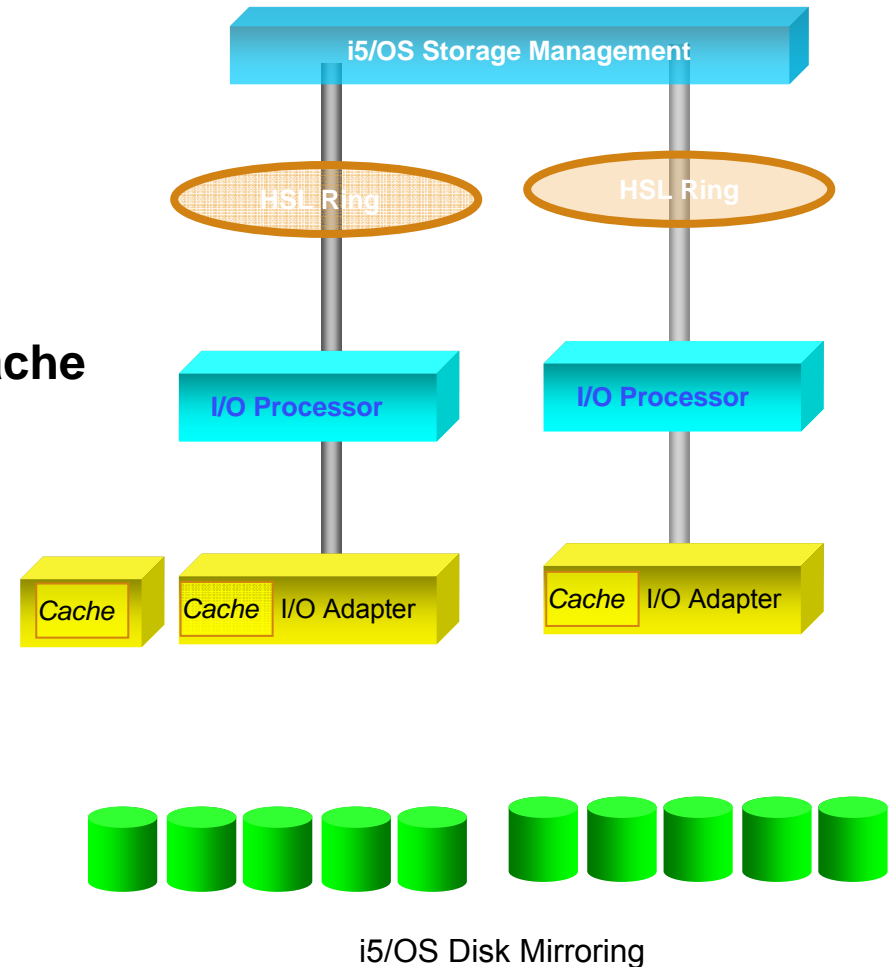
- *Modern RAID-5 solution*
- *Redundant cache*

i5/OS Disk Mirroring

- *Most robust solution*
- *Full I/O path, cache, and disk redundancy*

IBM TotalStorage DS8000

- *Redundant cache*
- *Hot spare*
- *Multipath fiber option*



Storage protection options

| Disk protection choice | Level of protection | What is protected |
|--|-----------------------------|--|
| RAID-5 | Basic | Single disk unit failure within a parity set. |
| RAID-6 | Basic + | Multiple disk unit failures within a parity set. |
| RAID with Auxiliary Write Cache | Intermediate | Same protection as RAID choice, plus reduced risk of write cache data loss. Outage occurs with disk IOA failure. |
| External storage with multipath | Intermediate to Best | Redundancy within storage subsystem minimizes outage and data loss due to multiple component failure in subsystem. Multipath minimizes outage due to IOA or path failure. |
| i5/OS disk mirroring | Best | Potentially multiple hardware components including HSL, IO tower, IOP, IOA, and disk with no outage required during component failure. |

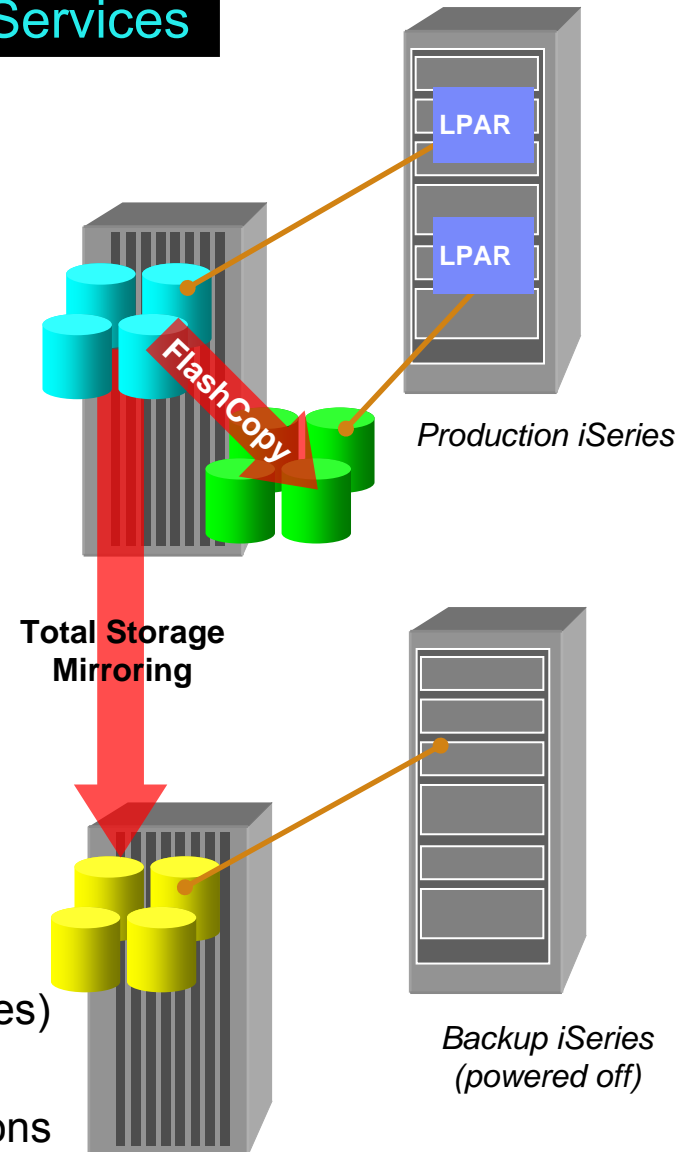
IBM Total Storage

Basic Copy Services

- **Metro Mirror** – continuous copy, synchronous
- **Global Mirror** – continuous copy, asynchronous
- **FlashCopy** – point in time snapshot

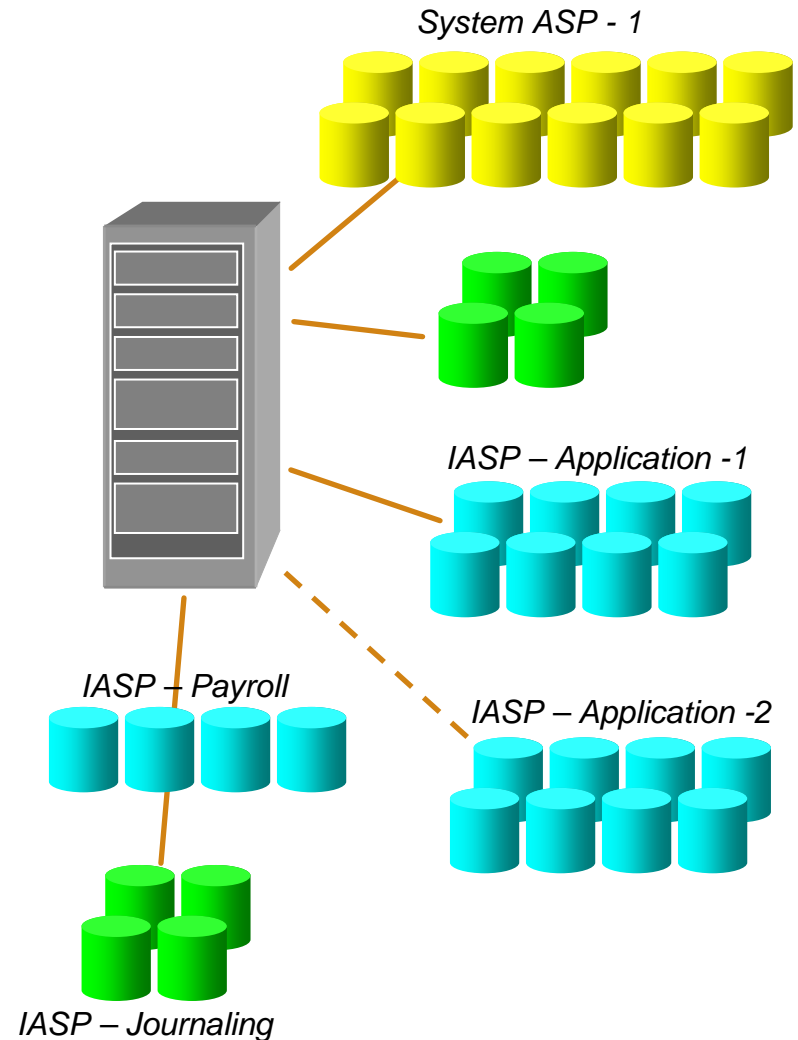
- **Benefits:**
 - Multiple platforms supported
 - Copy simplicity
 - Full system backups

- **Considerations:**
 - Available on *i5 models* only - managed via *HMC*
 - **Second copy cannot be used** for most planned outages
 - Source server power off **may be required**
 - **Target IPL** always required
 - **Requires journaling/SMAPP** to protect DB
 - More copies of disk than other solutions (2-5 copies)
 - New management complexities
 - Application and network performance considerations



Independent Auxiliary Storage Pools (IASP)

- **New storage design option - V5R2 and up**
 - Enhanced granularity and availability
 - IASP pool files support (V5R3)
- **Benefits:**
 - Maintains iSeries ease-of-use, plus increased availability:
 - IASP outage isolation
 - Isolate maintenance by IASP
 - Building block for Switch Disk
- **Clusters and TotalStorage Copy Services enablement building block**
- **Considerations:**
 - New storage and *application designs*, and new planning processes



Overview of factors influencing save performance



Target device technology

+




Connection technology

+




Bus and HSL speed

+




configuration

+



CPU and memory

+



Data Type (size and quantity)

=

GB per hour saved or restored

Save and Restore Planning


- Consider ***not only faster*** tape drives, but also the use of ***multiple tape drives***
- Improved with i5/OS V5R4 ***Virtual Tape***
- Use ***offline backups*** for the shortest window of all
- Understand your application/DB recovery requirements and expectations (does your ***DB require a quiesce*** before backup?)
- Manage complexity: use ***BRMS*** to manage both your backups and the latest backup and recovery technologies available
- Choose the ***simplest strategy*** that will meet your goals
- Ensure your operations teams have the appropriate ***skills***
- The only way to ensure you have a good backup strategy is to ***test your recovery***

Multi Server High Availability

System i5 Availability Solutions

IBM Offers Solution Choices To Fit Business Availability Needs

Logical Replication



Integrated or SAN

- Widely deployed
- Very flexible
- Rapid switch or failover
- Concurrent access

Basic Copy Services



- FlashCopy, Metro Mirror, Global Mirror
- Available on i5 models only
- Simple, but limited to outages addressed

Switch IASP Cluster



Integrated or SAN


- Simple availability
- Single copy of data
- Cluster building blocks for other solutions

Advanced Copy Services



- Increased function
- Switched IASPs plus DS Copy Services
- More complex, costs
- All iSeries models

Cross Site Mirroring



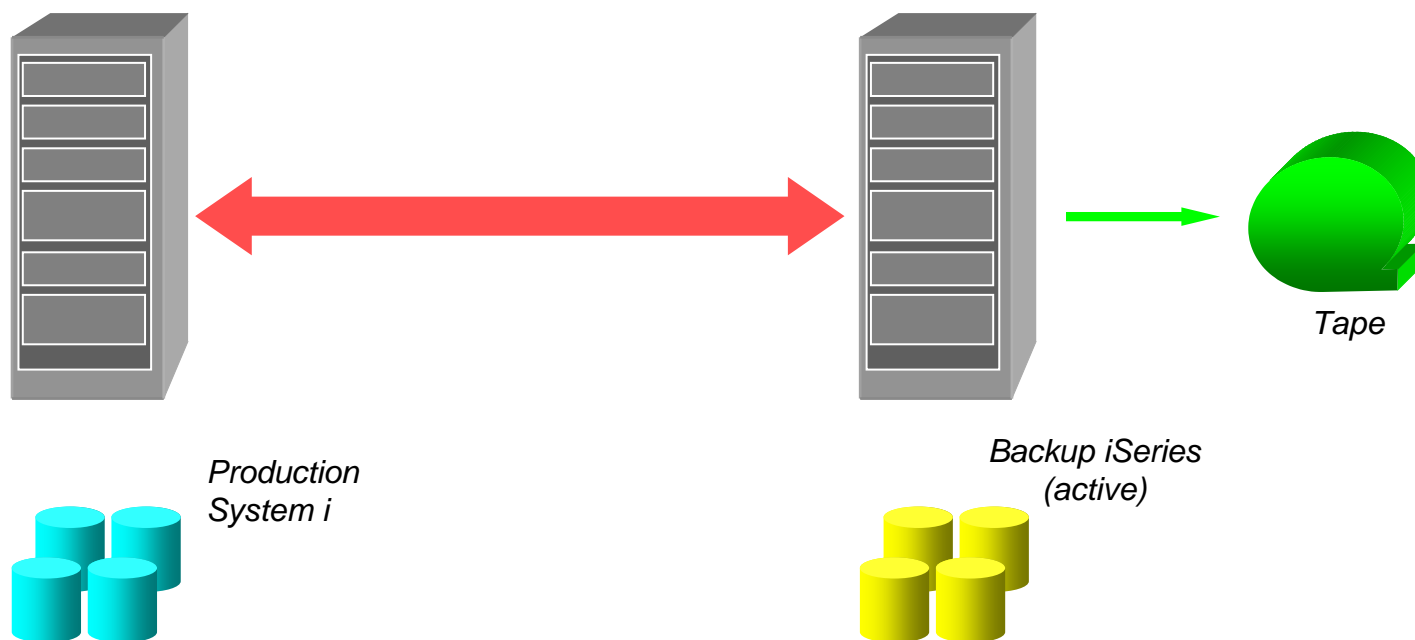
Integrated or SAN

- i5/OS managed copy of IASP
- DR and basic availability needs
- Enhancements with i5/OS V5R4

High Availability and Disaster Recovery is much more than just a copy of data or disk. RTO and RPO must be known, and the proper availability solution used, to meet those needs.

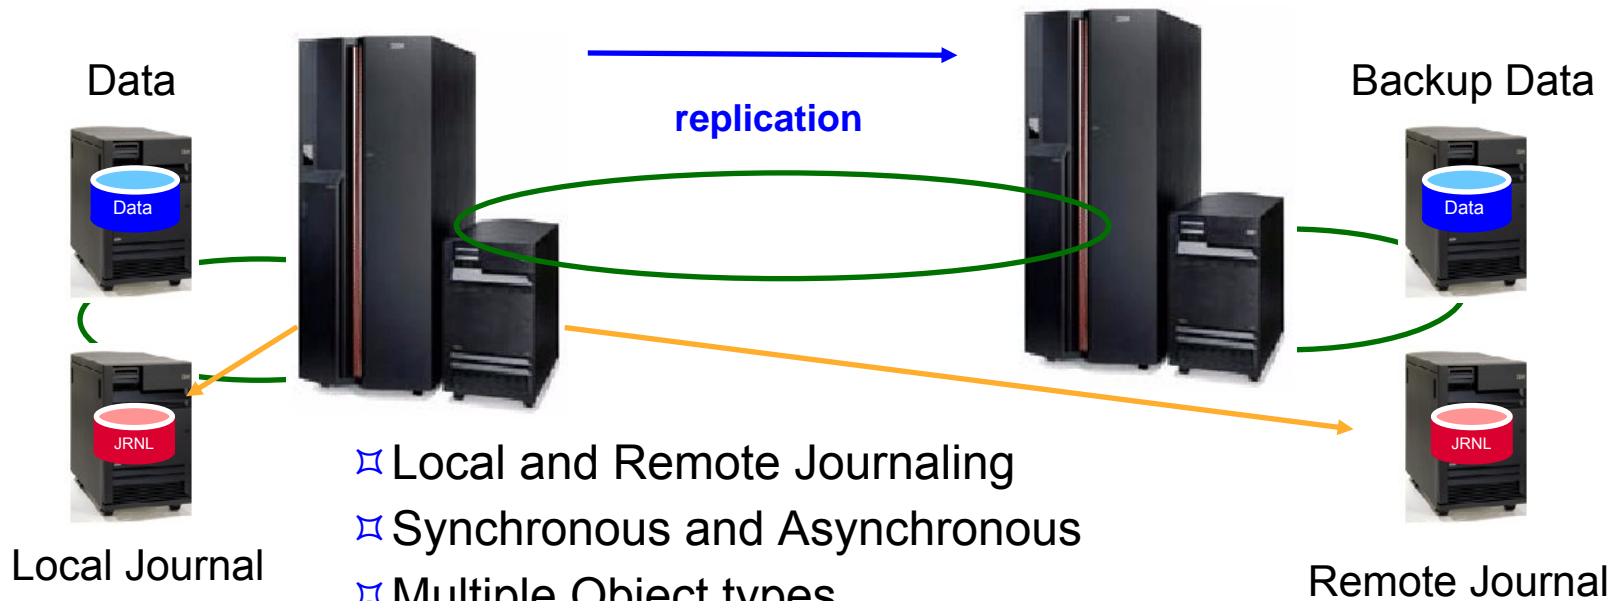
Logical Replication

- Most widely deployed data resiliency topology for IBM System i
 - Deployed via an HA Business Partner solution package
 - Or only basic functionality with remote journaling (not recommended)
- Replication done on data and object levels (file, member, record, data area, program, etc.) near real-time



i5/OS Journaling

Journaling ensures data integrity, by copying all object changes to a special place called the journal receiver



- ✧ Local and Remote Journaling
- ✧ Synchronous and Asynchronous
- ✧ Multiple Object types
 - ★ Database, IFS, Data Areas, Data Queues
- ✧ OS/400 Optional Feature 42
 - ★ Batch Journal cache
 - ★ Journal Standby Mode
- ✧ Journal Minimal data

When using remote journaling, the data is written first to the remote journal, then to the local table.

Logical Replication

■ Benefits:

– Costs

- Only **two copies** of disk required
- Very **flexible** (any hardware or disk, various OS release, most ISV software, etc.).
- Performed at the lowest unit of change for the object, e.g. **record level** for database
i5/OS Remote Journaling provides efficient, reliable transport mechanism
- HA Edition and Capacity Backup Edition available

– Geographical dispersion

- Replication on a data and object level (file, data area, program, etc.) **near real-time**

– Concurrent access

- Replicated data can be **concurrently accessed** for backups or other read-only apps

– RTO/RPO

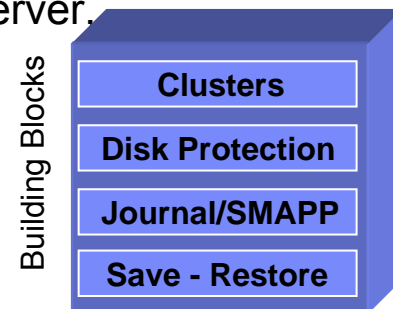

- Rapid switch or failover to production environment on backup server.
no IPL required

– Outage type coverage

- Covers all outage types

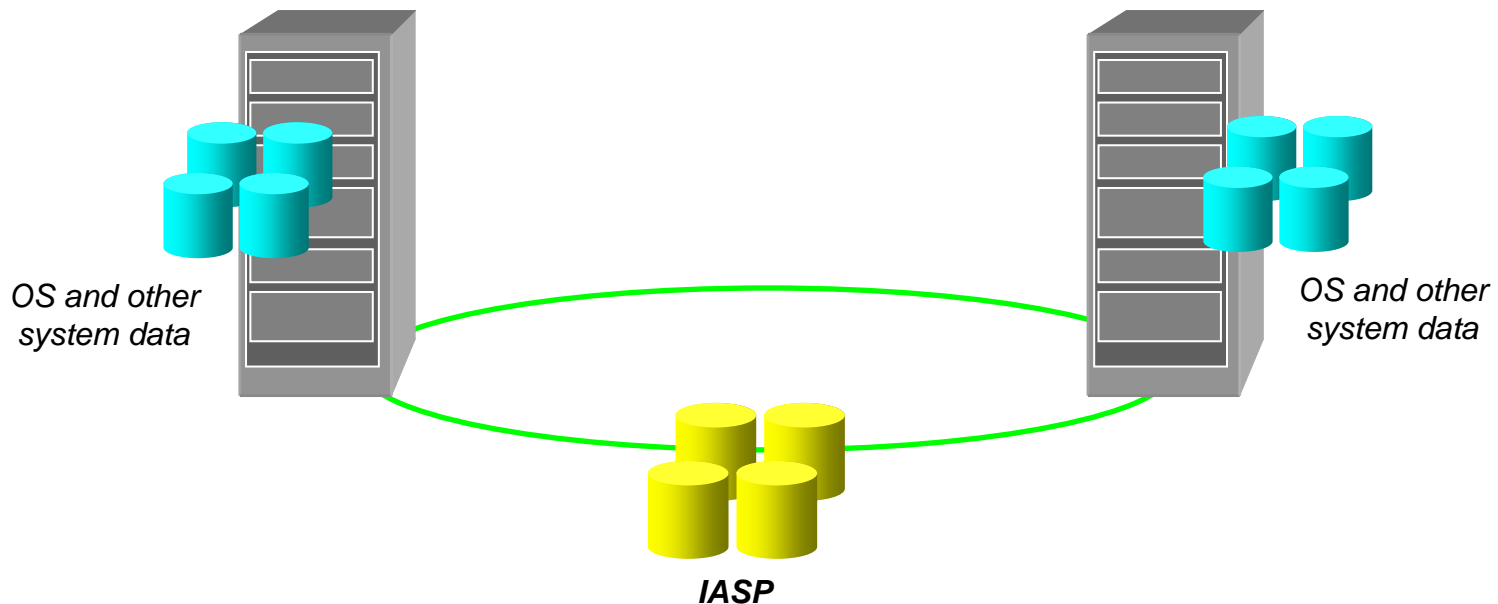
■ Considerations:

- Application and operations can be more complex to setup and maintain



Switchable Independent ASP (IASP)

- Single copy of data is maintained
- Data in IASP (or Disk Pool) is switched to backup system during outage



Switchable IASPs

- **Independent Auxiliary Storage Pools (IASPs)**


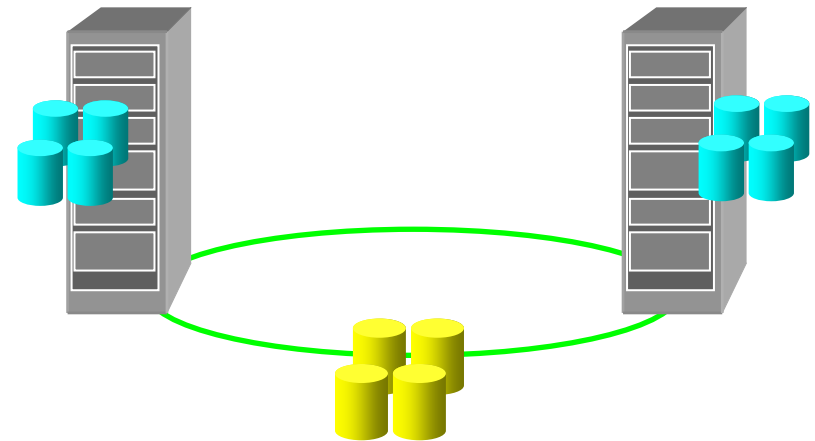
- Switch disk pool(s) from one server/LPAR to another

- **Benefits:**

- **Simplicity** and low cost
- Data is always **current** (no copy to synchronize)
- Minimal in-flight data to lose
- **Minimal performance** overhead
- Supports integrated Windows, Linux and AIX environments

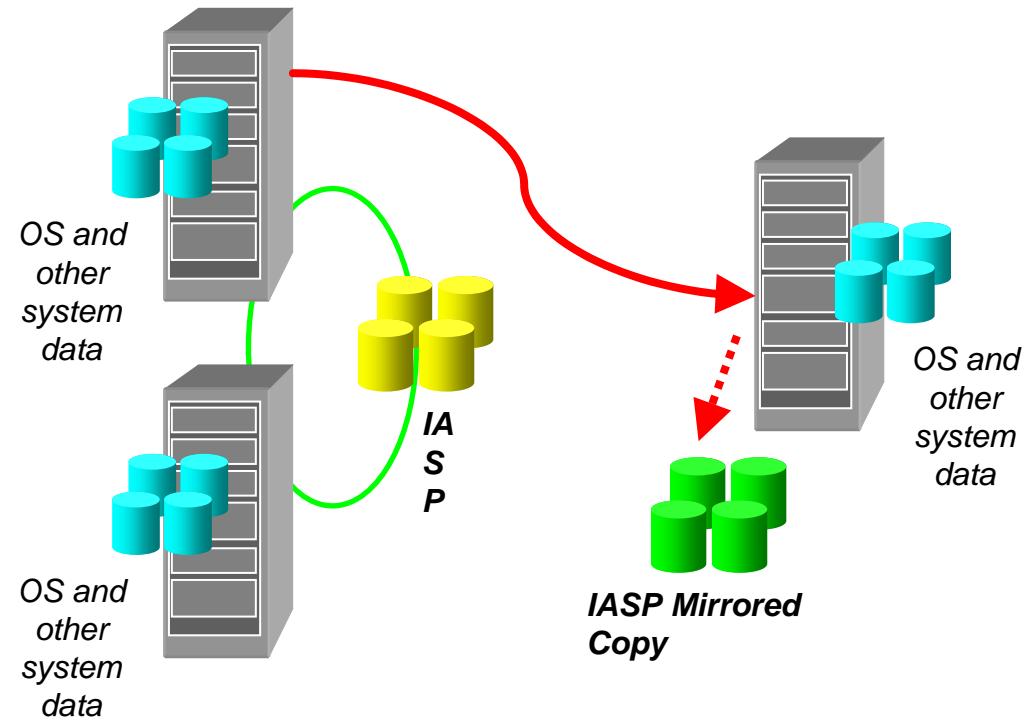
- **Considerations:**

- Design and setup of disk configuration, data and application
- **Single copy of data**
 - No concurrent access from both hosts
 - Cannot mask outages from tape backups
- HW restrictions (**limited distance** and concurrent maintenance)



Cross-Site Mirroring (XSM) with Geographic Mirroring

- Second copy of data in an IASP is generated logically identical to first
- Changes to production IASP replicated to second copy of IASP thru another system
- Operating system storage management based replication solution





Cross-Site Mirroring (XSM) with Geographic Mirroring

- **Mirroring of IASP data via OS/400 storage management to a second server**
 - Enables switching or automatic failover to mirrored copy of IASP

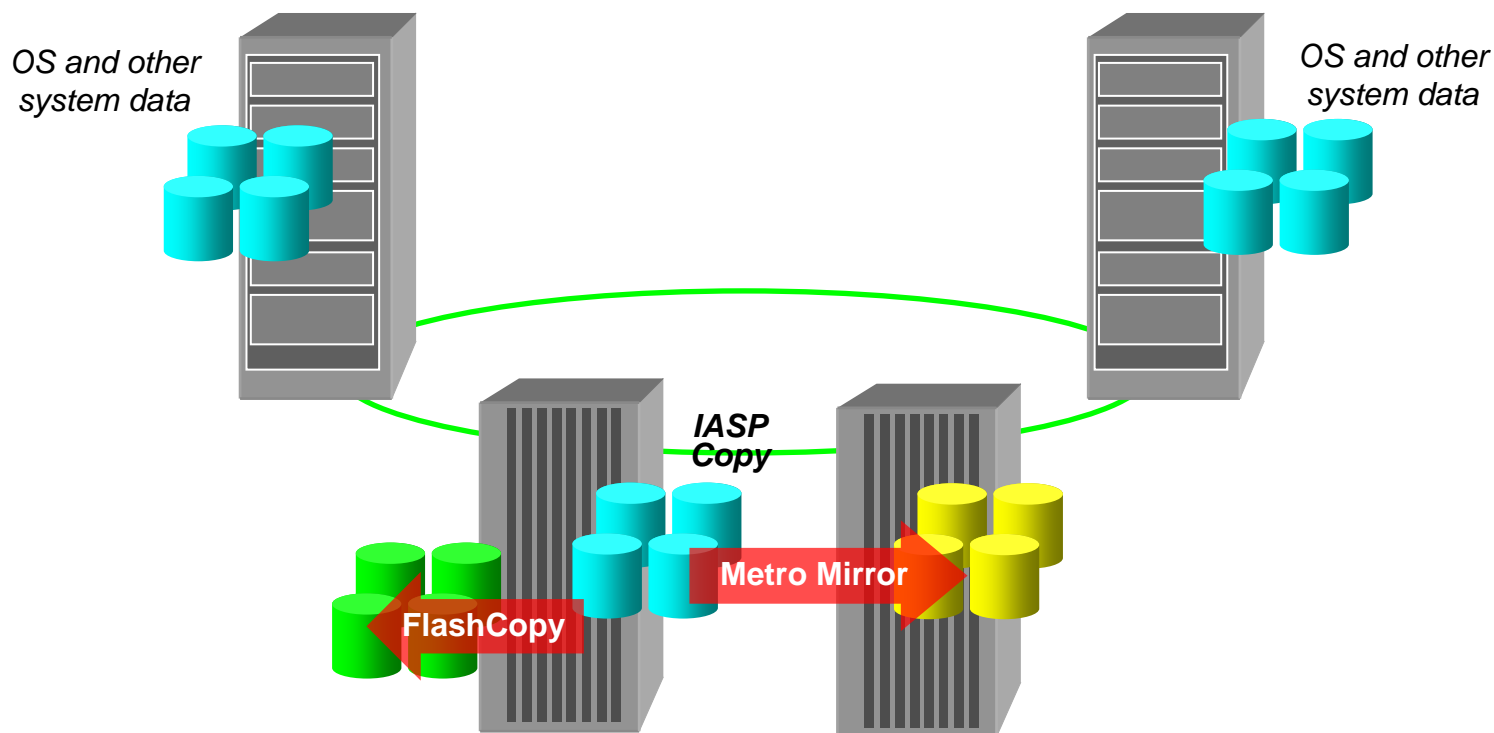
- **Benefits:**
 - Same as switched IASP
 - **Two copies** of IASP data
 - Lower costs - included in Option 41 of i5/OS V5R3
 - Can be local or remote
 - **Ease of deployment** and operation
 - Enhanced in i5/OS V5R4

- **Considerations:**
 - Performance impacts of synchronous operation, distance, bandwidth, latency
 - Mirror copy **cannot be concurrently accessed**
 - **Impractical to** detach mirrored copy to **do backups** to tape
 - Full data **re-synchronization** may be required for any persistent transmission interruption
 - Currently recommend at least a three system configuration
 - Can be done in combination with replication software



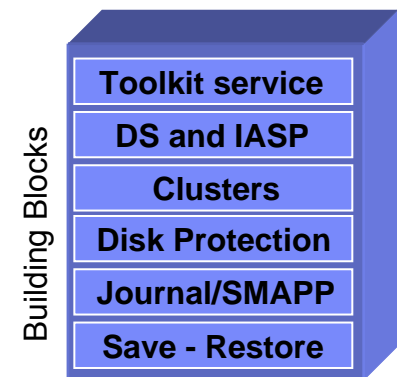

iSeries Copy Services toolkit for IBM TotalStorage

- Metro Mirror generates a second IASP copy identical to first
- Changes to production IASP replicated to second copy of IASP
- Cluster and SAN combined replication solution



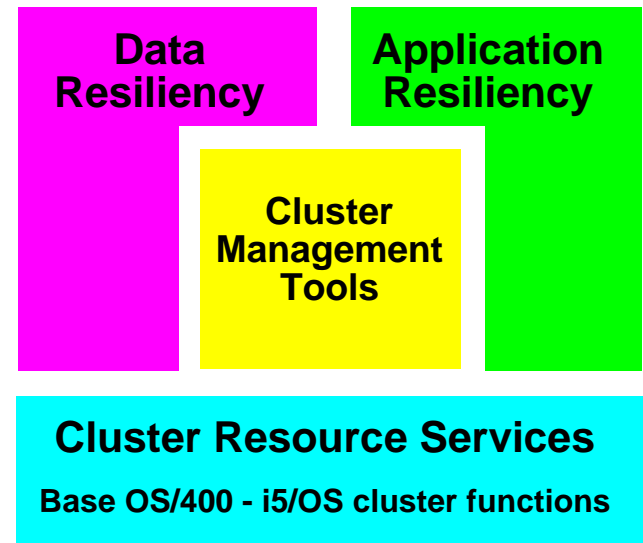
iSeries Copy Services toolkit for IBM TotalStorage

- **Replication of IASP data to a second DS8000 using Global Mirror**
 - Enabled via the “*iSeries Copy Services for IBM TotalStorage*” service offering
 - Combines ESS/DS Metro Mirror with IASP and OS/400 cluster services
- **Benefits:**
 - Remote copy and coordinated switching without an IPL
 - Combine with FlashCopy for backup window reduction
 - Does not eliminate need to quiesce for backup
 - Unlike basic SAN copy services, backup server remains active and can mask some unplanned outages
 - Any iSeries 8xx or i5 model
- **Considerations:**
 - Performance and distance limitations with distance
 - May require additional copies of disk to address all outages
 - Requires tools and services to deploy



Clustering Architecture

- A property of the OS
- **Enables switchover & failover automation**
- Provides the logical connections between **resilient data groups**
- Enables a **resilient application** to be switched, activated and repositioned to a defined state
- Enables automatic sequencing of events that bring the user, application, and data to a coherent production state




- Heart beating
- IP address takeover
- Reliable, ordered cluster messaging
- Switchover administration
- Distributed activities

Rochester Benchmark Center

Performance Benchmarks
High Availability feasibility testing
Cluster Lab
Consulting

Development Interface



Application Availability

Application Resiliency Building Blocks

The pyramid of end user happiness

- **Full application resilience:**

- Automatic restart and increased failover transparency

- **Automatic application restart and automatic recovery:**

- To last transaction boundary

- **Automatic application restart and semiautomatic recovery:**

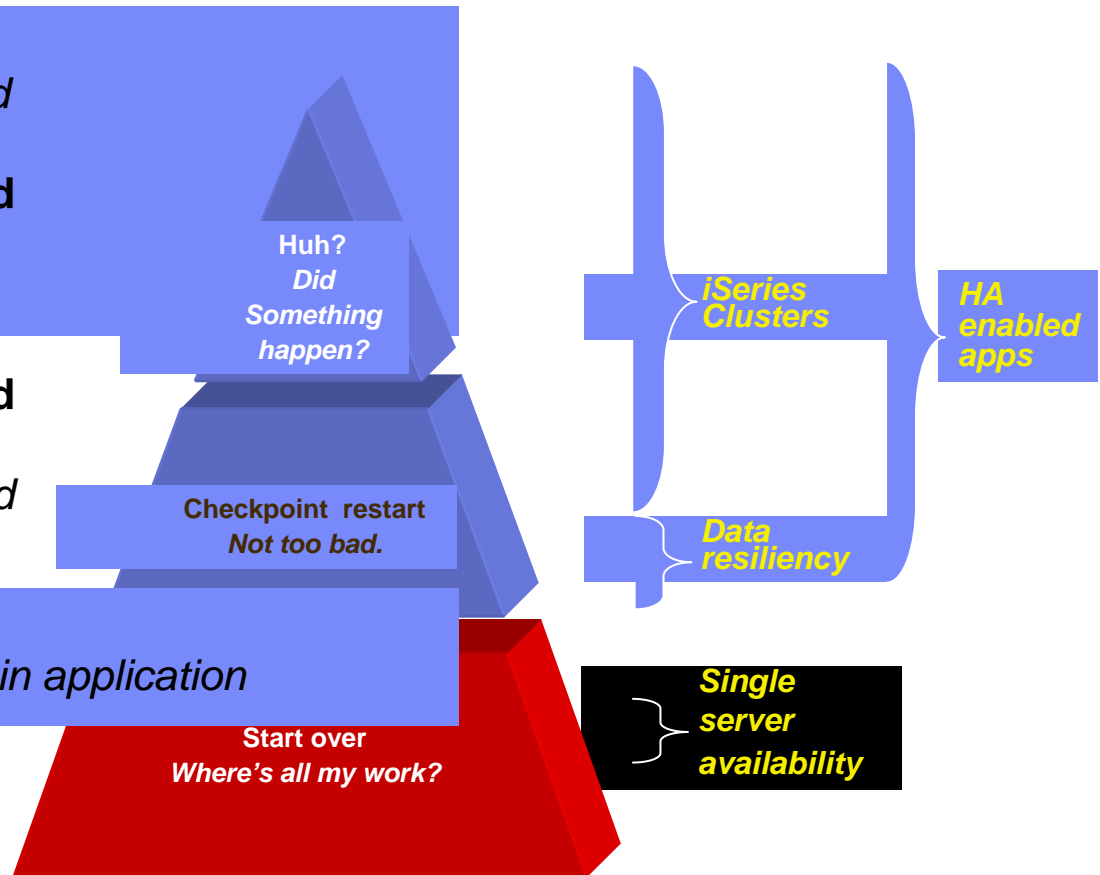
- User returned to last architected application “restart” point

- **Application restart:**

- User manually repositions within application

- **No application recovery:**

- Manual restart with resilient data



Summary

Assess and Define

You must first assess and define the requirements
... before designing a solution:



- Define your Service Level Agreements (*SLA*)
 - You will likely not meet availability expectations if business expectations are not known or current



- Define your Recovery Objectives
 - Recovery Time Objective (*RTO*)
 - Recovery Point Objective (*RPO*)



- Define the outages that must be masked or eliminated
 - Planned outage
 - Unplanned outages
 - Disaster Recovery

Consider All Major Requirements and Factors

| <i>Business Availability Criteria and Requirements</i> |
|--|
| 1. Up time requirements |
| 2. Recovery point objective (RPO) |
| 3. Recovery time objective (RTO) |
| 4. Resilience requirements |
| 5. Outage type coverage |
| 6. Concurrent access |
| 7. Geographic dispersion |
| 8. Tolerance for end user disruption |
| 9. Cost factors |
| 10. Downtime window availability |
| 11. Services and support |

| <i>Priority</i> | | | | |
|-----------------|---|-------------|---|---|
| <i>Less</i> | | <i>More</i> | | |
| 1 | 2 | 3 | 4 | 5 |
| | | | X | |
| | | | X | |
| | | | | X |
| | X | | | |
| | | X | | |
| | X | | | |
| | | X | | |
| | | | X | |
| | | X | | |
| | | | X | |

Which solution do I have to take ? 1/2

- When to consider Logical Replication?
 - Need **two or more copies** of the data
 - Want some level of **concurrent access** to second data copy
 - Need **backup** window reduction
 - Monitoring the state of the replication environment can be done by your IT staff
 - Geographic dispersion between copies is needed
 - Already have solution deployed using **logical object replication**
 - Need a solution that has no special hardware configuration requirements
 - Failover / switchover times should not exceed 10's of minutes
 - **Transaction level integrity** is important for all journaled objects
- When to consider Switchable IASPs ?
 - **Single copy** of data meets requirements; addressed exposure to disk subsys failures
 - Need a very simple, **low cost**, low maintenance solution
 - **No need for DR** solution
 - Source and target system will be at the **same site**
 - Want consistent fail/switchover times within minutes independent of transaction volumes
 - Need transaction level integrity for all objects; no loss of in-flight data
 - Need highest throughput environment




Which solution do I have to take ? 2/2

- When to consider Cross-Site Mirroring ?
 - Want a **system-generated second** copy of the data (at an IASP level)
 - Need two copies of data, but **do not need concurrent access** to second copy
 - Want relatively **low cost, low maintenance** solution, but also need disaster recovery
 - Need geographic dispersion between copies; distance does not impact performance goals
 - Want consistent **fail/switchover times within minutes** independent of transaction volumes
 - Need transaction level integrity for all objects; no loss of in-flight data

- When to consider copy services (Global/MetroMirror and FlashCopy)?
 - Want **storage based solution for DR**; especially if multiple platforms are involved
 - **Don't need complete HA**; want to cover DR and some planned outages for crit app data
 - **Recovery times** of one hour or more are acceptable (actual recovery times may be less)
 - Need two copies of data, but **do not need concurrent access to second copy**
 - Need geographic dispersion between copies; distance does not impact performance goals
 - Need transaction level integrity for all objects; no loss of in-flight data

- When to consider a combined solution?
 - **When no single solution meets all of your business continuity requirements**

Next Steps



Get business commitment

Determine what the business requires

Define the IT requirements

Assess and select the technologies to match requirements

Obtain the skills required to design, implement and maintain

Design the availability solution

Implement the availability solution

Keep the solution up to date



References

- ITSO Redpaper:
 - i5/OS High Availability Clusters: Data Resilience Solutions (redp0888)
 - www.redbooks.ibm.com
- RAS White Paper: IBM POWER5 Processor-based Servers - A Highly Available Design for Business-Critical Applications
 - www.iseries.ibm.com/ha
- System i5 High Availability and Clusters Website:
 - <http://www.ibm.com/servers/eserver/series/ha/>

References

- Redbook: iSeries in a Storage Area Network (SG24-6220)
 - Tape Chapter – www.redbooks.ibm.com
- iSeries Performance Capabilities Guide
 - <http://www.ibm.com/servers/eserver/series/perfmgmt/resource.html>
- Info Center Article: "Backing up the IFS - Experience Report"
 - <http://publib.boulder.ibm.com/infocenter/series/v5r3/ic2924/info/experience/ifsystm.pdf>
- Techdocs Article: How to Shrink your iSeries Backup Window - a Step by Step Guide

Ask your IBM rep or Business Partner to get you a copy from

 - <http://w3.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100600>
 - <http://partners.boulder.ibm.com/src/atmastr.nsf/WebIndex/WP100600>
- Magazine Article: System i Parallel Save/Restore
 - June 2006 Issue of COMMON Connect magazine
- Magazine Article: Virtual Tape
 - July 2006 Issue of IBM eServer System i magazine
- Redbook: i5/OS V5R4 Virtual Tape (SG24-7164)
 - www.redbooks.ibm.com
- System i Tape Info APARs
 - [http://www-912.ibm.com/n_dir/nas4apar.nsf/\\$\\$Search?openform](http://www-912.ibm.com/n_dir/nas4apar.nsf/$$Search?openform)
Search for II13513, II12621, II13904

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: i5/OS, i5/OS logo, AIX, AIX 5L, BladeCenter, Blue Gene, DB2, e-business logo, eServer, IBM, IBM Logo, Infoprint, IntelliStation, iSeries, pSeries, OpenPower, POWER5, POWER5+, Power Architecture, TotalStorage, Websphere, xSeries, z/OS, zSeries

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided "AS IS" without warranty of any kind.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US 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.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.