MSP

Your ICT Solution Integrator

Levels of Data Resilience and Solutions

Source: IBM Redpaper Data Resilience Solutions

MSP

Your ICT Solution Integrator

Business Continuity Problem Categories

Business continuity is the capability of a business to withstand outages and to operate important services normally and without interruption in accordance with predefined service-level agreements.

Your ICT Solution Integrator

Business Continuity Problem Categories

The Business Continuity solution must address:

- The data
- The operational environment
- The applications
- The application hosting environment
- The end-user interface

All of the above must be available to deliver a good and complete business continuity solution

MSP
Your ICT Solution Integrator

Business Continuity Problem Categories

Business Continuity includes:

- *High Availability (HA)*: The ability to withstand all outages (planned, unplanned and disasters) and to provide continuous processing for all important applications
- *Disaster Recovery (DR)*: In case of a disaster, a set of resources, plans, services and procedures is used to recover important applications and to resume normal operations for these applications at a remote site.

The main *differentiators* of an HA solution over a DR solution are:

- More demanding recovery time objectives (seconds to minutes)
- More demanding recovery point objectives (zero end user disruption)
- Fully automated failover to a backup system

→ CLUSTER TECHNOLOGY !!

A cluster is a collection of interconnected complete systems, used as a single, unified computing resource.

Business Continuity Problem Categories

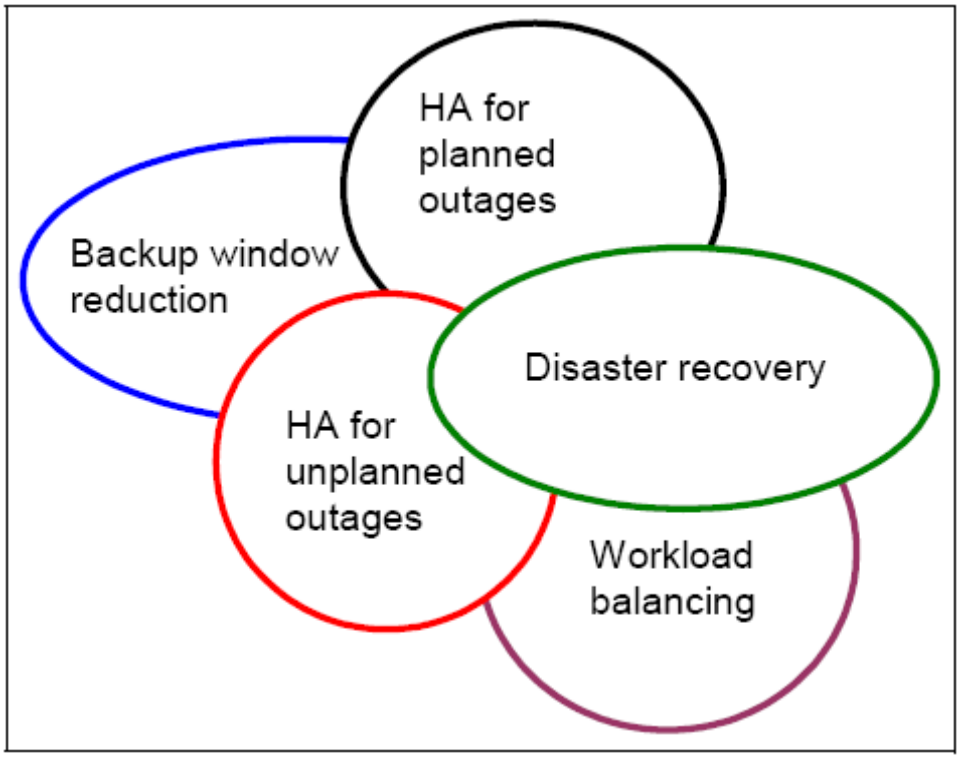

The above mentioned problem sets are not mutually independent.

Business Continuity Problem Categories (cont.)

- Backup Window Reduction
  - Reduce or eliminate non-productive time to do regular backups
  - Although daily outage for backups may be available, the length of the outage exceeds the available window
- Availability for planned outage events
  - Ensure that production can continue even if outage is required for planned maintenance (f.e. Hardware service/upgrade, software service/upgrade, ...)
  - Business operations demand that applications and data must continue to be available during time of planned maintenance
- Recovery from disaster related outage events
  - Requirement to ensure that the business is not affected by an extended outage due to a disaster (local system outages and site disasters)
  - Typically applications and data are relocated to a backup system, meaning IT services are unavailable 12 to 48 hours.
  - Disaster recovery may involve some amount of manual processing → assumed to be rarely needed

Business Continuity Problem Categories (cont.)

- High Availability for unplanned outage events
  - Ensure that critical system resources such as applications are available during both unplanned and planned outages
  - Business operations cannot tolerate extended outages while a disaster recovery site is brought online and dictate that failover processing must be fully automated
- Workload balancing
  - Multiple workloads are hosted in a multiple server environment
  - In case of system overload, application instances can be relocated to another server which has available capacity

MSP
Your ICT Solution Integrator

Business Continuity Technologies

There are several categories of technology solutions, however a single technology is not capable of addressing all needs of all customers.

4 categories of technology solutions:

- Backup window reduction
- Workload balancing
- Application resilience (applies to HA & DR)
- Data resilience (applies to HA & DR)

Both application resilience and data resilience are needed to achieve HA!

Business Continuity Technologies (cont.)

Backup window reduction

Reducing or eliminating the backup window by decreasing the time to perform the backup or decreasing the amount of data backup up.

Techniques:

- Improved tape technologies
- Parallel saves
- Saving to non-removable media
- Data archiving
- Saving only changed objects
- Saving to a second system
- Save while active

Business Continuity Technologies (cont.)

Workload balancing

Involves moving work to available resources.

Examples:

- Front end routers
- Multiple application servers
- Distributed, multi-part application
- Controlled application switchover

MSP
Your ICT Solution Integrator

Business Continuity Technologies (cont.)

Application resilience

Application resilience can be classified according to the following categories:

- No application recovery
- Automatic application restart and manual repositioning within applications
- Automatic application restart and semi-automatic recovery
- Automatic application restart and automatic recovery to last transaction boundary
- Full application resilience with automatic restart and transparent failover

Important remark: any of these application resilience mechanisms can be combined with the data resilience mechanisms described in the following section to provide a complete solution.

Business Continuity Technologies (cont.)

Data resilience

A number of data resilience technologies are available to address the data resilience requirements:

- Logical replication
- Switchable device
- Cross-site mirroring (XSM)
- ESS PPRC used in conjunction with iSeries Copy Services for ESS toolkit

Detailed attributes of data resilience technologies

Logical Replication

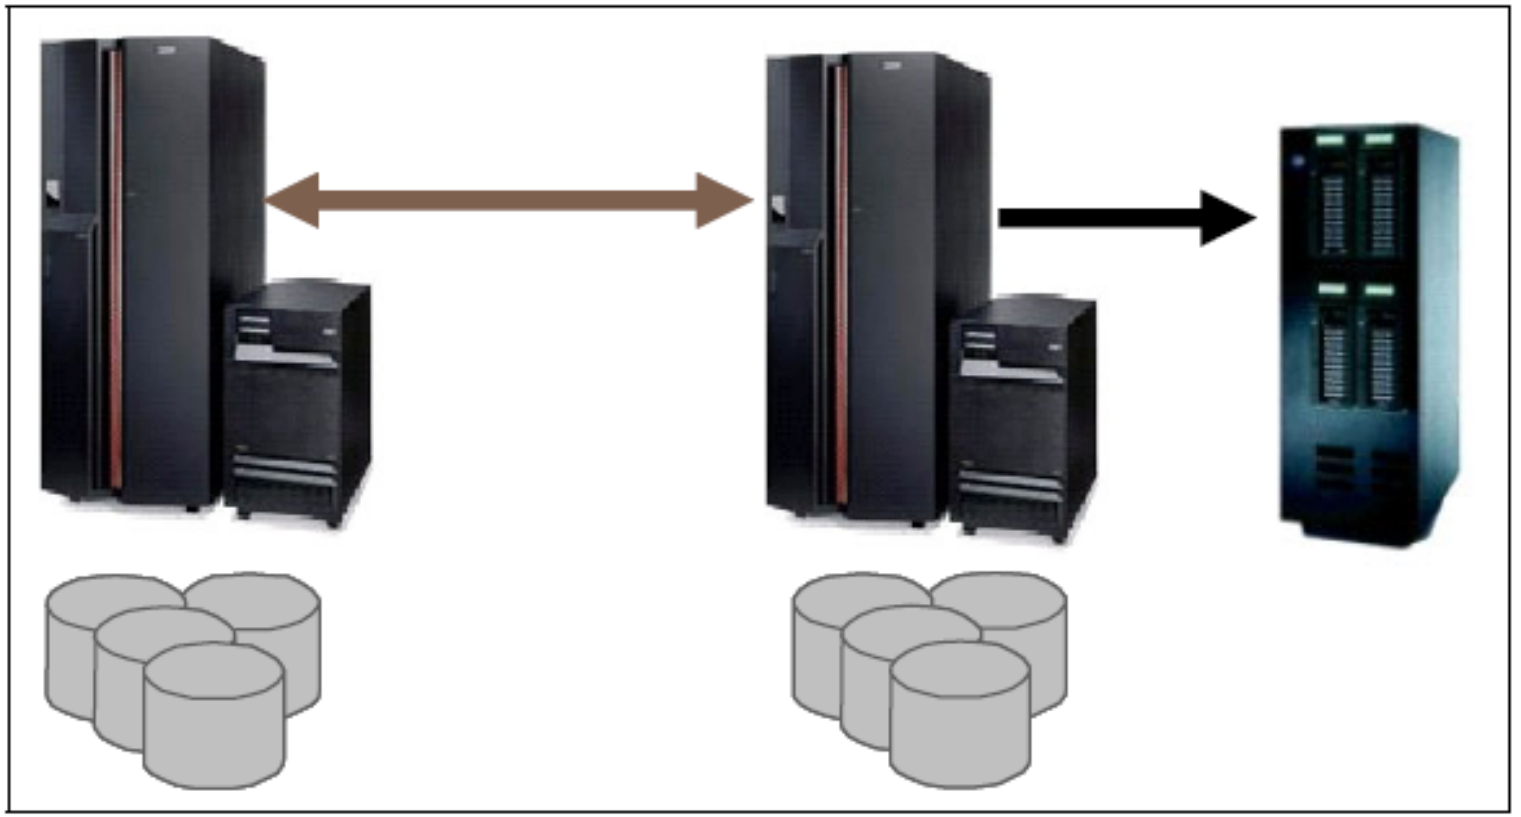

Detailed attributes of data resilience technologies (cont.)

Logical Replication (cont.)

Solution Characteristics

- Typically deployed via an HABP solution package
- Replication executed via software methods on objects
- Changes to the objects are replicated to a backup copy
- Replication near real time (simultaneous)
- Typically, if object is journaled (such as file) replication is handled at record level
- For other objects that are not journaled to a database journal, replication is handled at object level (entire object is replicated)
- Additional auditing capabilities
- Observe replication status in real time
- Automatically add newly created objects to those being replicated
- Possibility to replicate only a subset of objects in a given library or directory
- Possibility to use synchronous/asynchronous remote journaling!!

Detailed attributes of data resilience technologies (cont.)

Logical Replication (cont.)

Solution Benefits

- You can rapidly activate your production environment on the backup server via a role-swap operation
- Backup database file is 'live', meaning it can be accessed in real time for backup operations or other read-only application types (f.e. Building reports)
- Minimal recovery is needed when switching over to the backup copy

Solution Limitations / Drawbacks

- Complexity that can be involved with setting up and maintaining the environment
- Danger of undisciplined modification of the live copies of objects residing on the backup server
- HABP's have designed a large number of tools to simplify the management aspects and perform periodic data validation !!!
- Latency of the replication process, i.e. Amount of lag time between the time at which changes are made on the source system and the time at which those changes become available on the backup system → Synchronous remote journaling can mitigate this to a large extent

Detailed attributes of data resilience technologies (cont.)

Remote Journaling

Characteristics

- Function in OS/400 operating system that allows you to establish journals and journal receivers on a target eServer iSeries system and associate them with specific journals and journal receivers on a source iSeries system
- Remote journal function can replicate journal entries from source to the journals and journal receivers on the target system
- Supports both synchronous and asynchronous modes of operations

Detailed attributes of data resilience technologies (cont.)

Remote Journaling

Benefits

- Remote journal function was added by IBM to the iSeries within the licensed internal code layer of OS/400 → provides greater iSeries integration and efficiency than HABP's send and receive processes
- Lower CPU consumption on the source machine by shifting processing required to receive the journal entries from the source to the target system (asynchronous delivery)
- Eliminates the need to buffer journal entries to a temporary area before transmitting them from the source machine to the target machine → less disk writes and greater DASD efficiency on the source system
- Implemented in micro code, thus significant improvement of the replication performance of journal entries and allowing database images to be sent to the target system in real-time (synchronous delivery, meaning journal entries are guaranteed to be in main storage on the target system prior to control being returned to the application on the source machine)
- Allows for journal receiver save and restore operations to be moved to the target system → resource utilization on source system can be reduced

Detailed attributes of data resilience technologies (cont.)

Remote Journaling

Limitations

- No filtering of journal entries possible on source system → issue in case of limited communications bandwidth

Detailed attributes of data resilience technologies (cont.)

Switchable device

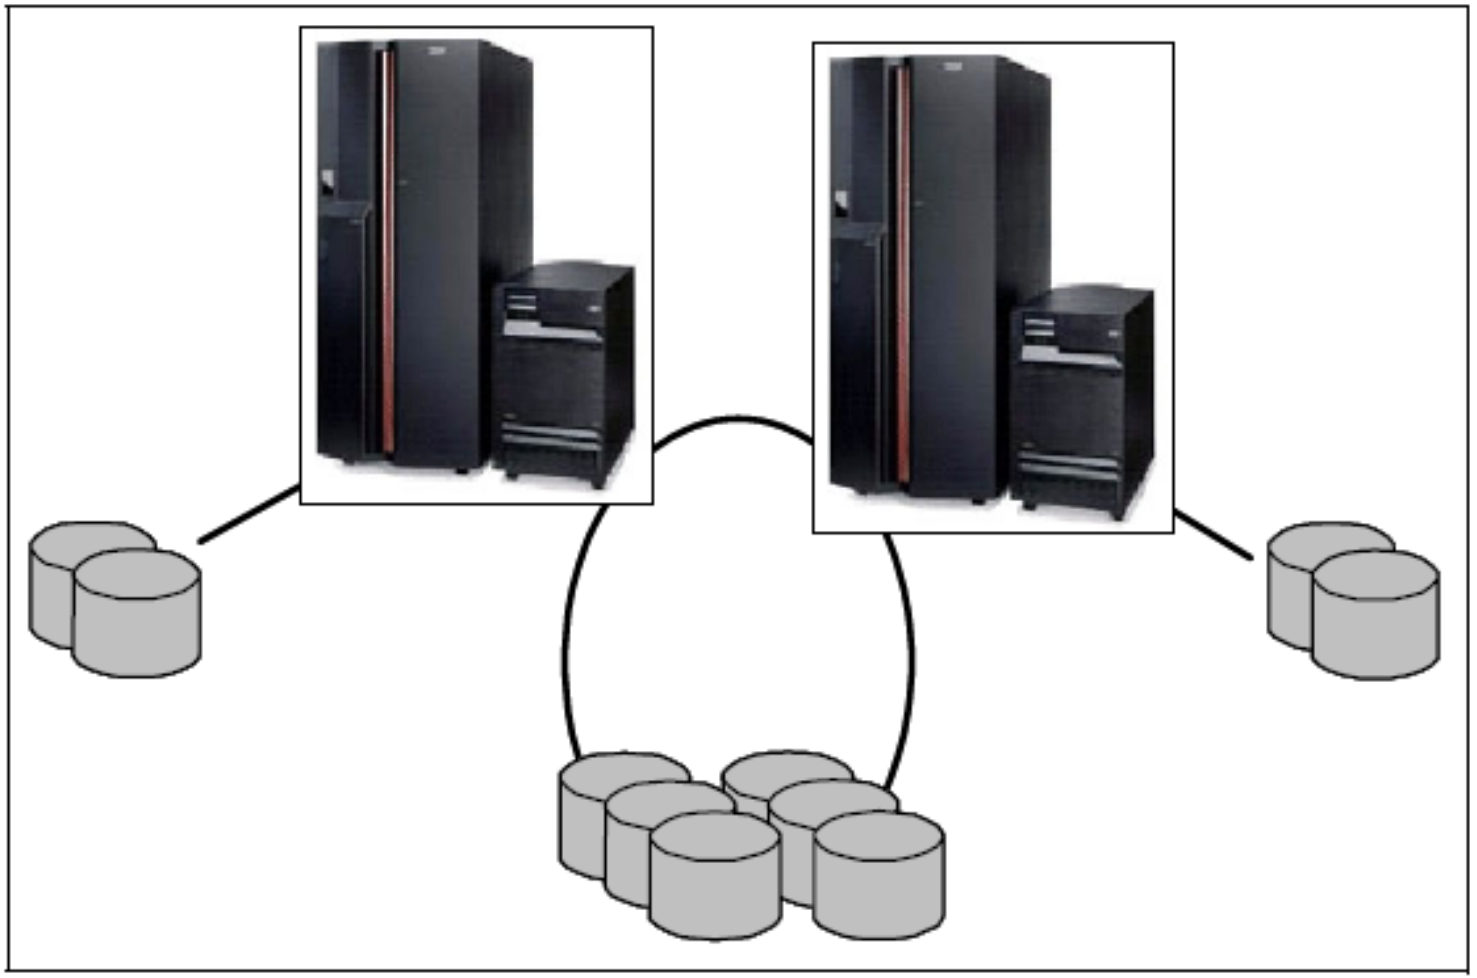

Detailed attributes of data resilience technologies (cont.)

Switchable device (cont.)

Solution Characteristics

- iASP (independent auxiliary storage pools) supports both directory objects (IFS) and library objects
- Is provided as part of i5/OS option 41, High Availability Switchable resources
- A logical solution as opposed to a purely mechanical-switching solution
- Architecture is deployed within the operating system as a special class of ASP that is independent of a particular host system
- Switching an iASP from one system to another involves less processing time than a full initial program load (IPL)
- iASP is placed into a switchable tower that is attached to two servers (or partitions via a high-speed link loop)

Detailed attributes of data resilience technologies (cont.)

Switchable device (cont.)

Solution Benefits

- Operational simplicity: single copy of data is always current (no synchronization to other copy)
- No data loss due to asynchronous transmission
- Minimal performance overhead
- Role swapping or switching is relatively straight forward (vary on iASP)
- Zero-transmission latency

Solution Limitations / Drawbacks

- Only one logical copy of the data in the iASP (Single point of failure)
- Data cannot be concurrently accessed from both hosts (no read access or no diverted backup)
- Certain object types cannot be stored in iASP (such as configuration objects) → other mechanisms required to replicate these objects to other system
- Distance limits in HSL loop technology
- Outages associated with certain hardware upgrades (iASP cannot be brought online to a down level system)
- Considerations about database restrictions on cross-iASP relationships such as Joins and referential integrity rules (up-front database design and analysis essential)

Detailed attributes of data resilience technologies (cont.)

Cross-site mirroring

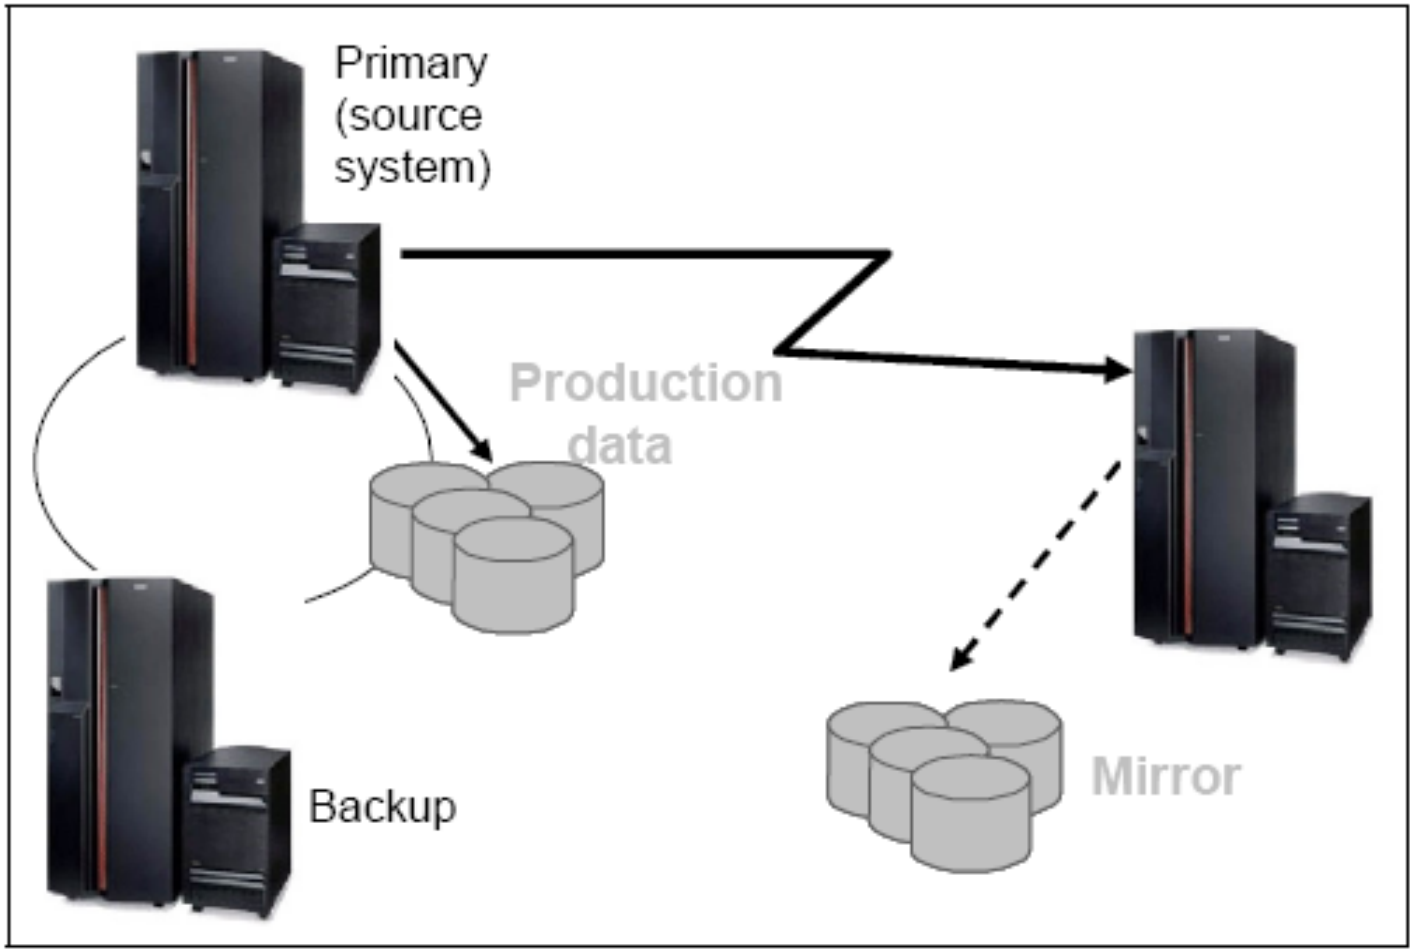

Detailed attributes of data resilience technologies (cont.)

Cross-site mirroring (cont.)

Solution Characteristics

- Involves mirroring of iASP data via i5/OS storage management to a second and perhaps remote server over a communications fabric
- Included in option 41 of i5/OS Version 5 Release 3
- Enables switching or automatic failover to a mirrored copy of the iASP in addition to locally switching the iASP between systems
- Addresses the single point of failure issue of basic switchable device structure
- Provides a means to develop a remote mirrored copy of your iASP data via *geographic mirroring*

Detailed attributes of data resilience technologies (cont.)

Cross-site mirroring (cont.)

Solution Benefits

- Same as basic switchable device solution with added advantage of providing DR to a second copy at increased distance
- Operational simplicity → all data placed in production copy of iASP (incl. Journal receivers) is mirrored to a second iASP on a second, perhaps remote system
- Switching operations essentially the same as switchable device solution, but you can also switch to the mirror copy of the iASP
- Solution easy to deploy and to operate
- XSM also provides real-time replication support for hosted integrated environments such as Microsoft Windows and Linux → generally not possible through journaling-based logical replication

Solution Limitations / Drawbacks

- Potential performance impacts in certain workload environments (use of synchronous communications → consider bandwidth, distance and latency limitations)
- Concurrent operations cannot access the mirror copy of the iASP
- Detaching an reattaching the mirrored copy or loss of all communication paths for an extended period of time implies a synchronization of the primary and backup iASP copies → system is exposed while synchronization is occurring
- Certain object types cannot be stored in iASP (such as configuration objects) → other mechanisms required to replicate these objects to other system

Detailed attributes of data resilience technologies (cont.)

IBM TotalStorage Enterprise Storage Server PPRC used with the iSeries Copy Services for ESS toolkit

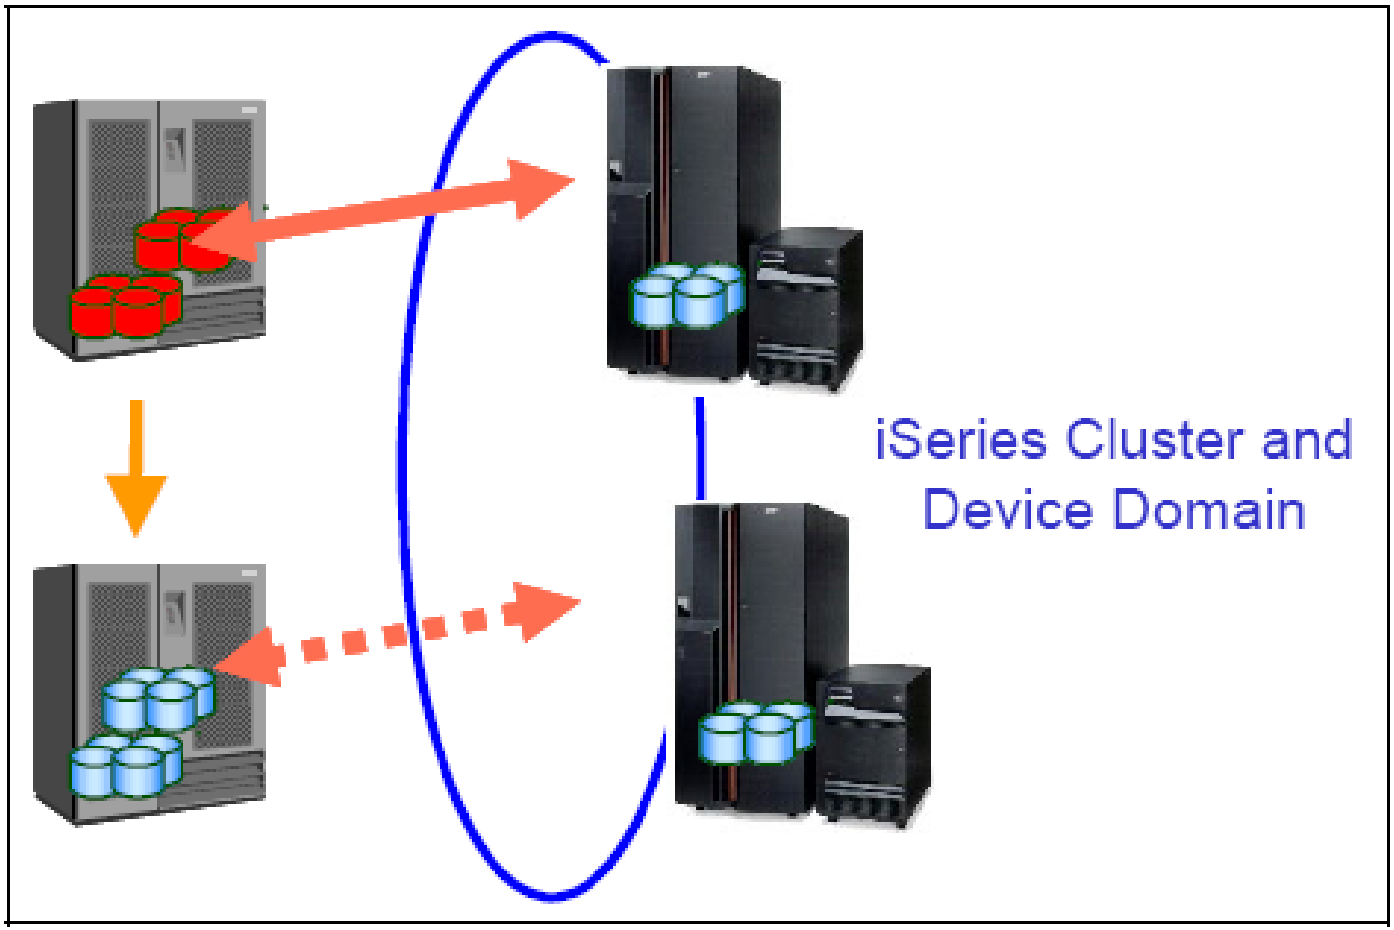

Detailed attributes of data resilience technologies (cont.)

IBM TotalStorage Enterprise Storage Server PPRC used with the iSeries Copy Services for ESS toolkit (cont.)

Solution Characteristics

- Involves the replication of data at the storage controller level to a second storage server using ESS copy services
- PPRC generates a second copy of the iASP on another ESS
- Toolkit provides a set of functions to combine PPRC, iASP and i5/OS cluster services for coordinated switchover and failover processing through a cluster resource group

Detailed attributes of data resilience technologies (cont.)

IBM TotalStorage Enterprise Storage Server PPRC used with the iSeries Copy Services for ESS toolkit (cont.)

Solution Benefits

- Remote copy function and coordinated switching operations → good data resiliency capability if replication is done synchronously
- Toolkit enables you to attach the second copy to a backup server without an IPL
- No load source recovery is involved in the operations

Solution Limitations / Drawbacks

- Switchover processing (mostly automated) requires some manual intervention to coordinate actions between i5/OS and the ESS
- If the switchover is unscheduled, an extra IPL on the failed system is required before it can accept the iASP again
- Solution isn't defined as an HA solution because of the required manual processing
- If two ESS are connected synchronously, you must also be aware of the distance limitations associated with transmission times

Remark: a solution of ESS PPRC without iASP and the Toolkit involves a long recovery period due to load source recovery processing and long IPL recovery steps.

Detailed attributes of data resilience technologies (cont.)

Business continuity requirement	Data resilience technologies			
	Logical replication	Switched disk	XSM	ESS toolkit for PPRC
Backup window reduction*		N/A	N/A	N/A
Planned maintenance				
Recovery for disaster outage		N/A		
HA for unplanned outage				N/A
Workload balancing				N/A

MSP
Your ICT Solution Integrator

Technology Comparison

	Logical replication	Switchable IASPs	XSM with geographic mirroring	ESS PPRC with IASP and iTC Toolkit
Primary use	HA (including DR)	HA (no DR)	HA (including DR)	DR
Characteristic of replication mechanism	► Object based replication ► Changes at record or object level based on data and audit journal ► Logical copy of object-level changes for selected objects	► No replication ► One copy of data that is switchable between systems	► Page-level replication as controlled by operating system based on storage management writes ► Logical copy since physical DASD configurations can differ	► Sector level replication of all pages written to disk ► Physical copy of an IASP based on disk I/O (cache based)
Recovery time considerations	► Apply lag plus replication switchover overhead ► Journal settings ► No IPL required ► Minutes	► IASP vary on ► SMAPP or journal settings ► No IPL required ► Minutes	► IASP vary on ► SMAPP or Journal settings ► No IPL required ► Minutes	► Quiesce time ► Manual steps plus vary on ► SMAPP/Journal settings ► IPL sometimes required before use backup again ► Tens of minutes
Recovery point considerations	► Transaction boundary with commitment control ► Mixed, audit and data journal ► Data or objects sent to target are recovered ► Any changes not transmitted are lost (zero data loss with synch remote journal)	► Transaction boundary with commitment control ► Last data written to IASP ► Objects not in IASP	► Transaction boundary with commitment control ► Last data written to IASP ► Objects not in IASP	► Quiesce point for breaking PPRC ► Transaction boundary with commitment control ► Last data written to disk (some automation and protection from mistakes)
Ordering of changes	► Based on journal receiver content and HABP ability to synchronize changes from data and audit journals	► Ordering preserved	► Ordering at system level ► Ordering preserved across ASP group	► Ordering at controller level ► Preserved at LUN set level for sync PPRC ► No order for asynch until 2.4.0
Concurrent access	► Typically read only, possibly shared data ► Always some lag time in data currency ► Remote Journal helps	► No concurrent access since no copy of data	► No, requires resynchronization ► Second copy current	► No concurrent access ► Copy current with synch PPRC; incoherent with asynch PPRC
Geo dispersion	Virtually unlimited	Limited (250 M)	Virtually unlimited	Virtually unlimited

Technology Comparison (cont.)

	Logical replication	Switchable IASPs	XSM with geographic mirroring	ESS PPRC with IASP and iTC Toolkit
Number of backup systems	1<= n <127 (or BP max)	n=1 (with switchable towers)	1<= n <=3 (2 or 3 with switchable towers)	1<=n<=2 (2 with cascading PPRC)
Number of data copies	127 (or BP max)	None	1	2
Cost factors	▸ Any DASD configuration ▸ HABP software ▸ Bandwidth ▸ Duplicate disks	▸ Switchable tower (or IOP) ▸ i5/OS Option 41	▸ Any (flexible) DASD configuration ▸ i5/OS Option 41 ▸ Bandwidth ▸ Duplicate disks	▸ Ext DASD (2 x Shark) ▸ Bandwidth ▸ i5/OS Option 41 ▸ PPRC ▸ Toolkit ▸ Duplicate disks
End user disruption	▸ Replication overhead ▸ Can automatically restart application	▸ Can automatically restart application	▸ Geographic mirroring overhead ▸ Can automatically restart application	▸ PPRC and toolkit overhead ▸ Semi automatic application restart
Outage coverage	Planned, unplanned, disaster, save window	Planned, unplanned	Planned, unplanned, disaster	Some planned outages, disaster
Cluster control	Yes	Yes	Yes	Yes – of switchable devices
Risks	▸ Loss of in flight data ▸ Mismatch of data levels for various objects ▸ Monitoring logical object replication environment	▸ Disk subsystem is single point of failure, therefore no protection against catastrophic disk failure	▸ Asynch case: Loss of copy for some double failure situations; OK if can quiesce and vary-off mirror copy ▸ Resynch may yield lengthy unprotected condition (especially with only two systems)	▸ IPL on backup systems in some situations ▸ Somewhat complex ▸ Never use asynch PPRC, unless using the new Global Mirror option

Your ICT Solution Integrator

General Guidelines when considering a data resilience technology

Consider ***Logical replication*** when:

- Two or more copies of the data are needed
- Backup window reduction is needed
- Concurrent access to the second data copy is needed
- Only a selection of objects in a library or directory needs to be replicated
- You have available IT staff that can monitor the state of the replication environment
- Geographic dispersion between copies is needed
- Already another solution using logical object replication has been deployed
- You need a solution that has no special hardware configuration requirements
- Failover and switchover times should not exceed tens of minutes
- Transaction level integrity is important for all journaled objects

Your ICT Solution Integrator

General Guidelines when considering a data resilience technology (cont.)

Consider ***Switchable iASP's*** when:

- Only one copy of the data with hardware protection satisfies you requirements
- There is a need of a low cost and low maintenance solution
- Disaster Recovery is not needed → only coverage for planned and certain unplanned events is required
- Source and target system are at the same site
- Failover and switchover times should not exceed ten minutes and do not depend on transaction volumes
- Transaction-level integrity is important for all objects
- The highest throughput environment is needed
- Your environment calls for multiple, independent databases that can be moved between systems

General Guidelines when considering a data resilience technology (cont.)

Consider ***Cross-site mirroring*** when:

- You want a system-generated second copy of the data (iASP level)
- You need two copies of data, but no concurrent access of second copy
- You desire a low cost and low maintenance solution, but you also need disaster recovery
- Geographic dispersion between copies is needed
- You want consistent failover and switchover times within minutes and that do not depend on transaction volumes
- Transaction-level integrity is important for all objects
- The second copy that is not available during resynchronization fits within your service level objectives

Your ICT Solution Integrator

General Guidelines when considering a data resilience technology (cont.)

Consider ***IBM TotalStorage Enterprise Storage Server (ESS) peer-to-peer remote copy (PPRC) with iASP and Toolkit*** when:

- You need a storage-based solution for DR (multi platform)
- You do not need complete High Availability, but you do want to cover DR and some planned outages for critical application data
- Recovery times of one hour or more are acceptable (times can vary)
- You want two copies of the data, but concurrent access to the second copy is not required
- Geographic dispersion between the copies is needed
- Transaction-level integrity is important for all objects

General Guidelines when considering a data resilience technology (cont.)

Consider ***a combination solution*** when:

- No single solution meets all of your business continuity requirements

Your ICT Solution Integrator



MSP n.v.
Your ICT Solution Integrator
Veldkant 7 • 2550 Kontich • Belgium
Tel. +32 (0)3 454 27 28 • Fax. +32 (0)3 454 27 88
info@msp.be • www.msp.be • BTW BE 0870.749.501