



Disaster Recovery by Data Replication – some issues and pitfalls

A Management White Paper

The need for a Disaster Recovery Strategy

Defining the requirements

Architecting a solution

The processes

Managing the DR environment

Getting the business to buy into DR

Conclusions

By John M Hudson
Director of T360

April 2002

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited
This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

The need for a Disaster Recovery Strategy

Disaster recovery planning is at the forefront of every IT centric business. Every organisation should have a DR plan and recovery strategy, scaled to meet the business needs. It is not the exclusive domain of large and wealthy organisations with two sets of everything, each in a purpose-built centre complete with support teams.

According to Gartner Group, two out of five enterprises that experience a disaster go out of business within five years. Threats to the modern data centre are many – Internet crime, virus, power failure, human error, fire, flood to name but a few.

Disaster recovery is often achieved by having a secondary site that can be brought online. This site can also implement continued service, even through planned downtime. A second site needs current data and this may be achieved by replicating data.

It is important to put DR into perspective – there is a lot more to a disaster recovery plan than just having a fairly up to date copy of the data at the secondary site! A DR plan must have the inherent mechanism and procedures which, when invoked after the failure at a primary site, will restore service to the users in a known time against operational objectives.

Defining the requirements

The first step is to define operational objectives. This may become an iterative process as those objectives may not be realisable – the business case cannot support that level of investment or the objective might be simply unachievable. Disaster Recovery Service Levels (DRSL) establish the priorities for recovering IT services from a disaster and should be analysed on a per application basis. DRSL has two components:

Recovery Point Objective (RPO)

The maximum time before the event, for which the business can allow the data to be rolled back to the time between back-ups.

Recovery Time Objective (RTO)

The maximum time the business will tolerate a discontinuity of service. This can be minutes, as a cluster fails over from one server to another, or days whilst new equipment is supplied, and applications and data restored.

For each application both the RPO and RTO can differ, along with the impact of applications. For example:

File & Print server

Both the RPO and RTO can be long, less business critical and the weekly backup is probably sufficient.

R & D server

Data generated by the R&D lab is valuable and costly to regenerate, so RPO should be minimal, although the RTO can be less demanding.

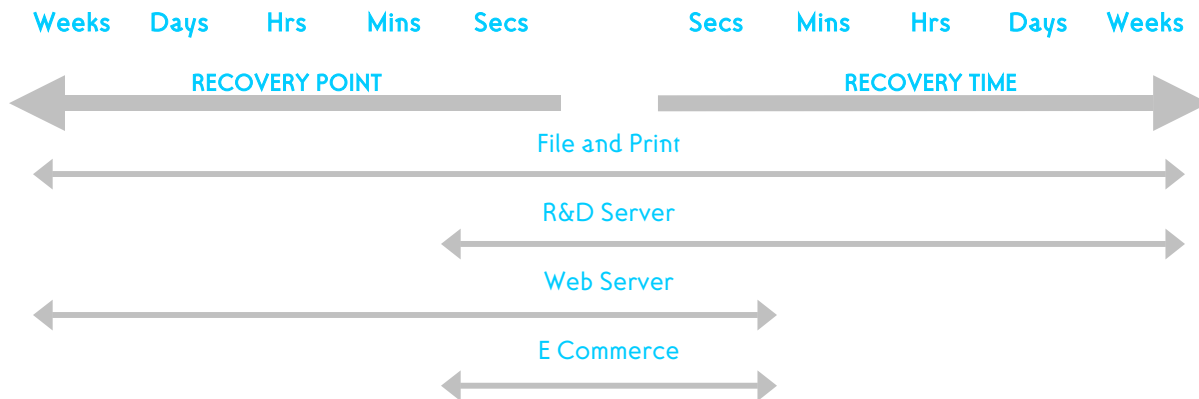
Web Server

The site itself changes infrequently and should be backed up as part of the change management process.

Disaster Recovery by Data Replication – some issues and pitfalls

E Commerce

Dynamic data and customer facing necessitate short RPO and RTO.



Threats to business IT availability can be grouped into five categories:

- **Local** Hardware or software failure
- **Logical** Software bug, virus, data corruption, accidental deletion, dropped table
- **Site** Natural – flood, earthquake etc
- **Third Party** Terrorism, sabotage etc, fire
- **Utility** Power failure, telecommunications

Each of these types has different implications

Local

If RTO is kind, a support contract with response time with the manufacturer could meet objectives – the real RTO has to include the reinstallation of the application and data. To reduce this threat and increase availability the systems can be design to be more robust. RAID storage, dual power supply, progressing to a warm standby or a cluster approach can be implemented.

Logical

This is the most likely threat for which a valid backup of all data is essential. Tape is the simplest and most cost effective backup medium but may struggle to meet either the RPO or RTO objectives. To improve RPO regular snapshots will provide a roll back mechanism. Online replication does not combat this threat: corrupted data, by whatever means, cannot be detected. To achieve a demanding RTO, these snapshot images must be transferred to the standby server. Logical threat requires a batch approach to data storage and recovery.

Site

A second site with equipment is required. The respective locations and functionality warrant close examination. It needs to be far enough away so that the same threat doesn't strike both operations – the same river burst its banks – served by the same power supply or substation. Too far and there are logistic and operational considerations.

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited

This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

Disaster Recovery by Data Replication – some issues and pitfalls

Choosing the secondary site must address:

- Purpose of site – hot, warm or cold standby
- Geographical – common natural effects or indeed its own risk
- Access – getting staff, data and other transferable essentials to it within the RTO
- Utilities – power supply and communications without dependency on the primary site

If the organisation is large enough, many of the above issues have less of an impact. Often two sites exist each with its own IT infrastructure. Duplication of equipment may not be required as both environments are already providing different IT services, and the fail over merely overlays multiple applications on one group of servers.

The table below sets out some typical scenarios as a framework for further discussion.

DR Classification	Acceptable data loss – RPO	Time to recover – RTO	Typical Implementation
a – Alpha	Less than 4 hours	Maximum 1 hour	One line replication onto warm standby
b – Beta	4 hours	4 to 12 hours	Regular batch replication
g – Gamma	24 hours	12 to 24 hours	Usual restore from off site tape
d - Delta	1 week	24 to 72 hours	Always restore from off site tape Delta

Even companies with a modest dependency on IT should have an effective two-site DR strategy. By incorporating a daily incremental backup into the process, a realisable RPO against local threats of 24 hours is attainable. Systems are either rebuilt from local data or in the case of a computer room fire, from the remotely stored tapes.

Gamma

By increasing the operational objectives, secondary servers must be available; the only component missing for fail over is data. Assuming that the distances are not too great, or there is a regular courier service between sites, daily backup tapes can be transported to the secondary site, restored and vaulted – don't forget to send over the backup catalogue too. This approach is ideal for batch-oriented businesses. There are some challenges that may preclude this approach:

- **Same site:** A single disaster could consume both the primary and secondary installation. Separate power supplies and communications might make this option acceptable – on cost justification against business impact.
- **Remote site:** The time taken to transport will delay the restore process. Backups can be sent to remote tape libraries but a dedicated fibre link can prove to be a prohibitive cost, especially outside metropolitan areas. A more cost effective alternative is fibre channel to ip, utilising WANs then back to fibre at the other end for connection to the tape library.
- **Data size:** As data grows the recovery time will also increase. Throwing more tape drives at the problem is expensive and will eventually reach a system bandwidth constraint imposed by the server or the storage. Realisable speeds of 10 to 12GB per hour for DLT and 20GB for LTO per tape drive mean that recovery times for 1TB plus will run into hours and that is on top of the courier time. What data is operationally essential? Can the services be prioritised to allow restoration to be phased?

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited
This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

Disaster Recovery by Data Replication – some issues and pitfalls

Replication by courier necessitates going back to older data, though if the failure occurs after restoration of backup data has been completed, the switchover can be completed quickly with the minimum of interruption of service to customers.

Beta

Some form of online replication is required to achieve a maximum 4 hours outage after a failure. Tape can be an option in limited circumstances, but backups must be completed online (impacting on service performance) and completed in less than 4 hours. Storage systems can provide a “snapshot” at intervals or create third party mirrors which can be broken off and backed up to another server to reduce the impact on service. A SAN allows these systems to be located in separate buildings, but this might not be considered a remote site.

Tape-based solutions have lower capital investment. However, the on-going costs associated with the management of the tapes cannot be ignored and the reliability of a fully automated process is greater than one dependent on people. Realistically, this approach is inappropriate when there is significant data and the operation is 7 by 24.

7 by 24 requires automated online replication. Both software and hardware interface to a wide range of WAN and fibre connections with their respective availability, performance and cost quotients. There are many factors that will impact on the respective data transfer rates.

Defining WAN bandwidth to accommodate bursts of data, generated at times of high activity, can be an unnecessary expense. For an RPO of 4 hours, short bursts can be addressed by additional cache to increase the replication buffer and thus facilitate a “catch up” process later. Transaction based applications, like databases, provide periodic snapshots for replication which will simulate a data burst. If there are several database instances, care needs to be taken to schedule the snapshots out as evenly as possible or better still, to periods of lower user activity.

To optimise WAN bandwidth, cache, resilience, performance, and ultimately cost, data transaction rates including snapshots should ideally be designed to be constant.

Alpha

Achieving this level of operational objective requires rigorous and detailed planning and is expensive and online replication is mandated. Storage is available with zero loss of data, but there is no safeguard in transit between sites.

Data replication can take place either synchronously or asynchronously. The better replication products support both types of transfers to be run concurrently, proving the most secure service to the highest priority application and improved performance for other elements of the application suite.

The following must be answered for each situation:

How much data can the business afford to loose? How far back in data terms will the business allow?

This determines the implementation method, which in turn starts to itemise the capital investment in servers, storage, databases etc.

How much performance degradation is acceptable during a disaster?

It is difficult to justify a business case that supports 100% server capacity on the secondary site for infrequent use for short periods, especially if it is targeted to provide a “needs must” service. Storage on demand, clustered servers with load balancing with applications shared across two active sites, will achieve a better return on investment.

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited
This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

Disaster Recovery by Data Replication – some issues and pitfalls

Review – Has this data and this application, the correct DR classification?

It is worth revisiting classifications of an application to establish the actual minimum required to sustain the business through a recovery, even if this means that reduced functionality is offered. The impact on investment and practical implementation may well be immense. For example, a reporting database could be deemed non-essential.

Architecting a solution

A mistake at this stage in the design will have significant economic and operational implications – design against now, the future and how to get there. The following questions cover areas that need to be addressed:

How do I set about sizing infrastructure to meet DR Service Level requirements?

This is the most demanding aspect of the project. There are many variables to be considered, from cache levels for data replication, through to application transaction rates and service levels on WAN links. Be wary of manufacturers' performance figures for hardware and software – achievable performance can be an order of magnitude, lower in real life situations. In practice, the actual performance of replication solutions is unpredictable. It is most unlikely that you can have high confidence in meeting any specific service level targets on day one, so performance logging, monitoring and modelling tools are essential. A simple change in an initial variable can have unpredictable effects further down the line; so empirical figures will be more use than theory.

How far to take “no single point of failure?”

A single point of failure can be costly and ideally should be avoided. Reducing the design from a “belt and braces” approach might be operationally acceptable and benefit from the capital savings. However, as with all aspects of DR, the rigour of a risk and impact analysis to the business is advisable to ensure any compromises are made knowingly.

How to size for the future?

A six-month period is classed as near term growth and, if the DR is part of a new application and service being offered, accurate predictions are difficult. What does the business case for the service predict? Investing in too much initial headroom will reduce the return on investment. Look at storage on demand, monitor the WAN traffic, provide end-to-end systems management, predict from real quantitative evidence and, therefore, be able to plan ahead. Change in these environments is resource hungry, disruptive and expensive.

Is hardware or software replication the best?

The two market leaders are Veritas, for software replication, and Hitachi Data Systems Lightning (9900 from HDS, SE9900 from Sun and XP512 from HP) for hardware replication.

The use of consolidated storage does not preclude a software approach; both have their merits and issues.

Operational aspects worthy of consideration are:

Does the number of servers requiring replication affect the decision?

For a few servers or small data requirement, investment in a large storage system such as Lightning could render the solution unviable. Conversely, a large number of servers could present an unacceptable administrative overhead. The cost implications associated with hardware replication are reduced if the applications are data centric, as savings can be made on the servers.

Disaster Recovery by Data Replication – some issues and pitfalls

With online replication, is it necessary to maintain a periodic backup policy?

It is essential to keep regular snapshots of data in line with the RPO for each application or element. Careful data management, including archiving, will reduce the activity, but a backup with the ability to roll back is the only insurance protection.

Should the infrastructure at each site be “over engineered”?

Elegant architecture will have ROI and aid business continuity. Clustering will increase application availability locally and, by strategic deployment, become instrumental in a program for server consolidation. Change management can be carried out in office hours and under fewer time constraints.

How can replication performance be increased?

Replication has a large computing burden on each end: on the server for software replication, or storage for hardware replication. These devices should not be under specified. Transactions are time stamped, checked and committed. When integrating into a fibre channel environment, extenders interface to routers, which in turn connect to the WAN. As the WAN will often be sized on “as small as possible” basis, bursts of activity and WAN downtime will result in queues. Sufficient cache is essential to provide enough buffering for these eventualities. If the sites are close enough (less than 100km), a direct fibre connection could be used although cost is a major issue. This removes the latency caused by channel extenders, routers, WAN downtime and the impact of sharing the link (in essence the solution becomes much more predictable).

The processes

What happens after the recovery process?

Whilst there are issues already outlined for sizing and architecting a solution to meet the requirements of replication in an online steady state, there is the often ignored consideration of returning to it after a recovery has been completed. An initial process of harmonising the data at both ends is accepted, but the issue is that after any fail over lengthy interruption of WAN, and possibly a scheduled downtime, the data must be synchronised. This can be accounted for in the initial design. The business is vulnerable throughout this period, so it is advantageous to minimise it. The WAN should not be sized on a “just adequate” basis – clustering, global and local can play a significant role here. When architecting the solution the method of approach to replication should be considered, with the objective to speed the return to the steady state.

When should DR be in place?

After testing and as the solution goes live this is not always possible. In an ideal world, the application is proven and removes the short-term need to upgrade, resulting in change management. Probably the most keenly reviewed aspect of replication is the quantities of data involved and their rate of change which in turn determines the pipe between the sites. An error here will be expensive – too large and it is wasteful, too small and the application would be forced to stall. Determining the transaction rate accurately is best assessed on a live instance of the application with real customers; if it is functionally incomplete or requires debugging, data structures may well change. Introducing upgrades to a live system that is replicated requires strategic planning and can take time to complete.

When should the DR architecture be tested?

Quarterly, and certainly after each major upgrade, whether that is hardware, operating system, application or middleware to ensure recovery is both functional and can be completed within the operational objectives. This is a resource intensive activity, but to discover an incomplete DR strategy triggered by an unscheduled failure defeats the original business case and could have serious financial implications.

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited
This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

Disaster Recovery by Data Replication – some issues and pitfalls

How should the recovery process be defined?

With absolute clarity and attention to detail. Nothing should be left to chance; there is an RTO to meet. The whole process requires documentation to allow a complete automaton-like approach to the process; there is no time to think or experiment, the systems guru could be unavailable, manuals take time to search and hand made notes are often ambiguous. The precise location of operating system, and application software CDs, must be guaranteed and the configurations known. Duplicate tapes and off site vaulting adds security and safety, but this extra level of complexity must be incorporated.

What happens after failure or change?

Downtime not only occurs as a result of a failure at the primary site. An upgrade to the operation system and hardware can result in operations being transferred to the secondary site, albeit controlled and planned. To complete the recovery process data on both sites must be re-synchronised. Throughout this time the business is vulnerable. Re-synchronising large amounts of data between two remote servers can take days.

One of the advantages of hardware replication is that data replication can be maintained whilst the servers are down. Re-synchronisation is, therefore, much less of an issue.

This consideration has most validity in rapidly changing environments, where the application is under regular review or when the server platform is less robust, such as non-Unix environments.

Managing the DR environment

With DR in place, does this security remove the need for system management?

No, it is more essential than ever. Monitoring allows pre-emptive maintenance, planned and scheduled. The correct toolkit should provide predicted growth based on quantitative evidence for accurate strategic investment. Long term performance monitoring also provides quantitative evidence for strategic change as operational demand develops.

How to deal with change?

These solutions are very complex, so change management must be planned after careful risk and impact analysis. Consider the impact on a live server requiring more capacity, additional hardware, operating system upgrade, or application enhancement. The problem is compounded because now there are two of everything, and a dynamic link connecting them. Consider the task of increasing a table size: this cannot be replicated until the change has been completed at the secondary as well. So the business is without DR during:

- Fail over to secondary
- Increase table size
- Re-synch data
- Fail-over to primary
- Increase table size
- Re-synch data

Ordinarily a DBA would complete this online, without interruption of service in moments, but in a DR scenario this is a major undertaking.

Disaster Recovery by Data Replication – some issues and pitfalls

The design of a DR solution must consider the process of change, whether online replication is used or a man on a bike with a tape in his pannier pedalling between sites. Better computing architectures offer 5 9's availability (under 5? minutes per year downtime), allowing operating systems to be upgraded online and storage to be added on demand. This functionality minimises the impact of change management. Often it is not the downtime itself that causes the problem, it is the re-synching which, when badly designed, will take several days, leaving the business vulnerable.

Getting the business to buy into DR

The economic implications to a business devoid of a DR strategy can be catastrophic, but equally it should be designed to meet the absolute operational objectives. Most organisations cannot warrant a sophisticated online replication approach but there are solutions that can be integrated and tailored to meet objectives and provide the ROI.

What happens if there is no plan, or it is inadequate?

There are the visible and invisible consequences. Visible are the lost revenues for the duration of the outage, but these are insignificant when matched against the consequential losses of providing a less than robust service. Customers move quickly, commercial loyalty is an obsolete concept.

How to strengthen the business case?

As a consequence of careful planning and architecting, IT infrastructure savings can be made. Most IT infrastructures have been built piecemeal over the years – a series of vertical applications supported by their own computing resource. Worse still is the architecture of more robust solutions. Typically, the predominance of host A / host B scenarios in corporate IT environments, based upon a large number of modest servers, all have significant capital and on going revenue implications. "Server consolidation through storage", another white paper by Tectonic 360 Ltd explores this topic in greater detail. However, it is worth noting that the business benefits discussed are magnified in a DR scenario.

What should the business look for in any proposed solution?

There are a number of considerations

- **Open systems:** A solution incorporating proprietary components, or those from marginal manufacturers adds levels of dependency on a single supplier. Not a good idea in a crisis.
- **Expertise:** There are many "gotchas" in DR. Specialists' experience will avoid bad calls whilst completing the DR strategy.
- **Own the problem:** Ultimately, it is your business and not the third party specialist, so ensure that there is a comprehensive hand over and training program so that destiny is in your hands.
- **Centralised management:** Common and centralised management brings a further business benefit to not only this solution but also the overall cost of running IT. Fewer valuable IT specialists are required to maintain infrastructure, releasing resource to enhance and push the business forward. Training budgets can be focused, and the development of a strategic alliance with a toolkit manufacturer will give greater scope for commercial negotiations to take place.

Conclusions

To successfully implement a disaster recovery program:

- IT must negotiate realistic DR service levels with its customers on an application-by-application basis.
- Hardware and software architectures must be defined with availability in mind.
- Set up a POC to evaluate the live dynamics: transaction rates, bursts etc. to allow accurate sizing for better investment.

Copyright 2002:

All information and ideas contained and presented in this white paper are the sole property of Tectonic 360 Limited
This White Paper formed the basis of a Paper presented to the Institute of Banking and Finance

Disaster Recovery by Data Replication – some issues and pitfalls

- Plan recovery after a fail over – minimise vulnerable time.
- Document the process in detail.
- Plan the process for change as part of the initial design.

T360

T360 is an innovative solutions provider that assists organisations to transition IT from being a necessary component to a strategic collaborator within the business.

'Aligning IT to the business is a Process not a Product'

Based upon best of breed ITIL solutions, T360 enable customers to achieve an end to end visualisation of their business services delivered by IT through the areas of:-

- Business Service Management
- Business Activity Monitoring
- Customer Experience Management

T360's philosophy is simple; we listen, understand and deliver. Our customers trust us because we provide innovation, expertise and commitment.