

Addressing the challenges of data protection

Mehran Hadipour explores some of the drawbacks of data protection technologies and provides a checklist of aspects to consider when choosing a solution.

Protecting high-value data and delivering 24x7 data protection and business continuity is of paramount importance to organisations throughout the world. Unfortunately, those organisations that have embarked on this mission have found considerable challenges along the way from the many infrastructure challenges associated with managing heterogeneous platforms, applications and data, to the challenges presented by the limiting and costly technology options available today.

INFRASTRUCTURE CHALLENGES

IT infrastructures usually include a myriad of server, storage, and application platforms. In addition, data and applications often span across distributed or clustered servers and storage. Supporting and protecting these heterogeneous platforms is a complex issue. Furthermore, as not all data is of equal value to an organisation, and as the value of data can change, determining how to most effectively protect this data is an ongoing problem.

Managing an end-to-end disaster recovery solution across an enterprise is currently an extremely complex challenge. Different storage platforms offer proprietary disaster recovery solutions, each with its own management challenges. Host based solutions can impact server performance and require another layer of data management. In addition, many disaster recovery solutions today also require additional infrastructure (like protocol converters) that in turn add yet another layer of complexity. And of course, organisations must deliver disaster recovery solutions without impacting the performance of key applications.

These many infrastructure challenges result in costly implementations that often do not address the complete disaster recovery needs of an organisation.

DATA REPLICATION CHALLENGES

Enterprises need a disaster recovery solution that delivers a reliable up-to-date remote copy of its mission critical data but will not result in performance degradation of the applications; it must be cost-effective and therefore must use minimal extra storage (an original and one copy should be enough), and must support the organisation's specific (and dynamic) availability requirements.

Data replication methods, from synchronous to asynchronous to point in time, have evolved over the years, in an attempt to address these dynamic needs of enterprises. Unfortunately, whereas each method offers advantages over the others, significant disadvantages are also present in all.

Synchronous replication addresses the very fundamental requirement for any effective disaster recovery solution of having an up-to-date remote copy of the data. With this replication method, every write transaction must be acknowledged from the remote site. This method ensures that an up to date copy of the primary site is maintained at a secondary site and that if a disaster occurs in the primary site, the secondary site will be consistent with the primary site. This works well for replication within a local SAN environment; however, extending this approach to transfer data over the WAN results in significant latency problems, high bandwidth costs and a

dramatic degradation in the performance of critical business applications. This can have a highly disruptive effect on business operations.

With asynchronous replication every write transaction is acknowledged locally and then added to a queue of writes waiting to be sent to the remote site. Although asynchronous replication does not reduce the bandwidth requirements associated with synchronous replication, it does reduce the latency problems. Unfortunately, however, for "write intensive" applications performance will eventually deteriorate to that of synchronous replication. Furthermore, with asynchronous replication, the copy at the secondary site is not necessarily up to date; as a result, in most disaster scenarios, data will be lost. Another key drawback of asynchronous replication is data inconsistency: in certain situations, even the most advanced solutions currently available are unable to maintain "write order fidelity" at the remote site and in the event of a disaster, no consistent copy will exist. Additionally, existing asynchronous solutions do not scale well and are either limited to one storage subsystem or one server.

With both synchronous and asynchronous replication, all modified data is transferred to the remote location. As a result, resource requirements, including storage and bandwidth, are high and costly. With snapshot replication, a consistent image of the changes made to the primary site (since the previous snapshot) is periodically transferred to the remote site, thus reducing the amount of transferred data. The advantages of this approach include lower bandwidth costs and minimal application degradation. However, in practice, existing solutions can be prohibitively expensive due to the cost of excessive storage requirements. In addition, snapshot replication provides limited protection in the event of a disaster; since the snapshot at the remote site will not be up to date, there could be significant data loss. Furthermore, existing solutions remain bandwidth intensive as they transfer data in an unreasonably large granularity.

Enterprises are thus challenged with the fact that although each replication method addresses important issues, none of them is ideal for the dynamic requirements of the organisation. It is clear that what is needed is a replication methodology that encompasses the advantages of the above methods but that eliminates the disadvantages; a replication methodology that can intelligently and dynamically select and utilise a replication method based on customer provided policies and on the point in time availability of network resources.

EXISTING DISASTER RECOVERY SOLUTIONS

Commonly used solutions such as off-site back up tapes do not provide up to date protection of data nor do they enable rapid recovery. The need to use communication lines for hot replication to a remote disaster recovery site is thus clear.

Current solutions, including volume mirroring, host based replication, storage based replication and database replication, are either limited in functionality, only work with a selected platforms, are expensive to implement, or both.

The industry is now seeing a new technology that moves the intelligence for data protection into the network (both SAN and LAN) and provides an intelligent universal solution for data protection for all of the storage and servers on the network.

TECHNOLOGY OPTIONS

There are several technology options available that offer some form of disaster recovery, and like the case of the replication methods discussed above, each one was designed to address the deficiencies of the other.

'Volume mirroring' creates an exact mirror of the original data and therefore demands extremely short distances. In addition, in order to reduce application degradation, it requires an extremely high-speed connection, resulting in high network costs. Furthermore, of course, it does not provide the distance required for effective disaster recovery.

In 'host based replication', the distance between the sites can be extended dramatically. However, since the replication software resides in each server, it takes valuable host cycles away from the application, possibly degrading application performance. This solution often requires a significant WAN bandwidth and also has a major impact on the local application performance. Furthermore, installation and setup of the replication software in each and every server can easily become a cumbersome and costly endeavour.

'Storage based replication' offers a host-independent solution, offloading the host from replication responsibilities. Many storage vendors offer their own proprietary solution and therefore only support that specific storage platform. This limitation results in undesired management complexity and cost.

'Database replication' is offered by many database vendors as a way to protect data within its control. As a result, only a portion of an organisation's data can be protected in this fashion and customers must use additional technologies to cover all other data types.

Which technology to use in order to deliver reliable disaster recovery is therefore a difficult challenge. With currently available options the decision always results in costly implementations that are not optimised to the organisation's dynamic needs. Clearly what is needed is a next generation replication technology that delivers maximum data protection with no data loss, that is host and storage platform independent, that can understand database linkages and dependencies, that will not increase management complexity, and that will adapt to an organization's changing needs – in a cost-effective manner.

THE REALITY OF DATA PROTECTION WITHOUT LIMITS

Next-generation networked-based architectures based on an intelligent data protection appliance connecting to the SAN and IP infrastructure provide data protection for all the storage and servers attached to the network.

These appliances experience no data loss by making an up to date copy of the data available at the remote site while combining this with very short recovery time in the event of a disaster. The solution intelligently recognises the differences between the local and WAN environments and utilises unique algorithmic-based technologies to combine the best features of each of the three existing replication approaches, while avoiding their disadvantages. It achieves this by dynamically adapting the replication approach to changes in traffic conditions due to the output load from the host application and as data is transferred from the local environment to the WAN.

A powerful, highly differentiating feature of the new generation approach is its ability to establish flexible replication policies based not just on the widely used technical parameters (e.g. maximum "write lag limit" between the primary and remote sites) adopted by other replication solutions but on criteria directly linked to business performance. For instance, the frequency with which data from a specific application is replicated can be set to reflect the relative business risk and cost to the company of lost data and/or application downtime when compared to data generated by other applications.

In the event of a disaster in which the primary storage system is temporarily disabled, a data replication appliance ensures rapid recovery with full data consistency and no data loss. This achieves the business continuity of a synchronous solution while, at the same time, it minimises application degradation and the bandwidth and storage costs associated with any one individual replication approach. In addition, 'multiple snapshot' techniques enable users to 'roll back' to a snapshot of the data as at various points prior to the time of the disaster as an added precaution against the risk of data corruption.

LAN

A replication solution should support multiple host/multiple storage system environments and integrates fully with all existing local replication and management solutions, thus allowing companies to leverage their existing storage infrastructure.

Other features to look for in data replication:

Universal data protection

Data protection for all open server and storage platforms on the network. The solution must remove the SAN distance limitation allowing DR site to be far apart. It should offer all possible replication policies in a single system (snapshot, asynchronous, and synchronous replication) without the need of edge connect or WDM devices using the standard shared or dedicated IP infrastructure.

Autonomous management

Make sure your replication system is capable of adjusting to WAN bandwidth and or application demand changes dynamically while enforcing the established policies for specific applications. The user should be capable of establishing different policies for each volume group and to enforce the established policies automatically. These systems allow different sets of policies dictated by the business need and criticality of the application and data enabling multiple service levels through the same infrastructure. Disasters, site or storage system failures are overcome through a rapid resynchronisation feature minimising down time.

Application aware compression

Look for unique agent technology that supports typical applications such as Oracle databases, detects the nature of the application and is capable of optimising the replication techniques while maintaining always-consistent remote copy.

Delta differentials

The system should maintain the write order fidelity and track and transmit only the changed bytes as opposed to writing the complete block of data multiple times. This saves bandwidth and improves performance.

Hot spot compression

Operating in the snapshot mode the system should track multiple WR requests against the same data blocks and only transmit the last WR while maintaining the consistency of the remote copy at all times.

SUMMARY

Data drives much of the value created by the enterprises today and data loss, whether due to human or equipment error, natural or artificial disaster, has business implications of enormous proportions. The need for key data to be 100 percent reliable, always accessible and fully up-to-date is clear for a growing number of enterprises. These conditions must be met at a cost that is affordable and without in any way hampering the operation of critical business applications.

Without an adequate disaster recovery solution in place, lost data and prolonged downtime could result in a loss of massive amounts of revenue and productivity, as well as customer trust and brand equity, which take years to build but just hours to destroy. With hourly downtime costs over \$6 million for some organisations (Gartner), the need for an effective disaster recovery solution is high up on the strategic agendas of the CIOs of leading international companies.

Mehran Hadipour is the vice president of marketing for Kashya, Inc.

www.kashya.com Hadipour has over 23 years of experience in senior and executive roles in the storage industry, including 19 years at IBM, and can be reached via e-mail at mehran@kashya.com