

EXECUTIVE SUMMARY

Data Replication

DEFINING DATA REPLICATION

Data replication is a backup in real time that duplicates all data and selected application images stored on the Production Storage Area Network (SAN) to a remote SAN at the Disaster Recovery site. The availability of a high speed data line between the Production and Disaster Recovery sites ensures that all changes and transactions are synchronized between the two storage devices. This process guarantees that the Disaster Recovery site will be current should a business continuity event trigger a failover to the Disaster Recovery site and infrastructure.

Utilizing SAN to SAN replication enables an automated and almost instantaneous transition to a DR environment, eliminating the need for IT staff to reload applications and data from secondary storage (tape and disk), a time-consuming process during which impacted applications will be unavailable to staff and customers.

When the primary site is returned to the normal operational state, the data replication can be reversed and, upon completion of the resynchronisation, processing is to be switched back to the primary site, enabling the business to continue with minimal disruption.

A significant advantage of the replication model is the reduction in overall Wide Area Network (WAN) traffic when compared to traditional data migration strategies. SAN based replication transmits only the changed portion of a record or file over relatively high cost WAN bandwidth. For example, when a telephone number is updated in a customer record, only the updated data is sent over the WAN to the DR site. In traditional transfer models, at best, the entire customer record is transferred, and in many cases, the entire customer database would be copied from site to site.

TRADITIONAL ARCHITECTURE

Without SAN-based replication maintaining an accurate and current copy of production data at the disaster recovery site, it is difficult if not impossible to create an automated and rapid transition to the DR data center. The two primary methodologies utilized in traditional disaster recovery data transfer models are:

- Periodic copying of full data sets from production to disaster recovery locations
- Periodic restore from tape of full backups taken at the production site.

Both of these options fail to keep data at the disaster recovery site synchronized with production. Should a transition of production processing to the disaster recovery SAN be required, the DR SAN data will be up to 24 hours out of date. Transactions occurring between the nightly data migrations from production to DR will, in many cases, be gone, resulting in lost productivity as staff identifies and reenters the missing transactions or recreates files.

In the traditional model, failing over to the Disaster Recovery site will take many hours and require considerable labor - data and applications must be reloaded and transactions recreated before production processing can resume at the DR site. The same process will be required to bring the production site back on line, incurring similar down time and labor to ensure the production site data set is current before the transition can occur.

DATA AND APPLICATION REPLICATION – SERVER AND STORAGE VIRTUALIZATION

Server and Storage Virtualization solutions enable the automated transfer of information and operating environments between physically separated data centers. Applications and data are replicated continuously from the production data center to the disaster recovery site, ensuring that current data is immediately available during a disaster recovery event. These capabilities dramatically reduce the time and effort required to bring the disaster recovery site on line.

Implementing this model requires the appropriate SAN and replication software at both sites, and a high speed data line between the sites. The expense of this connectivity can be greatly reduced by placing both sites at colocation centers owned and operated by a single colocation service provider and utilizing the very high bandwidth, low cost connectivity between the two sites.

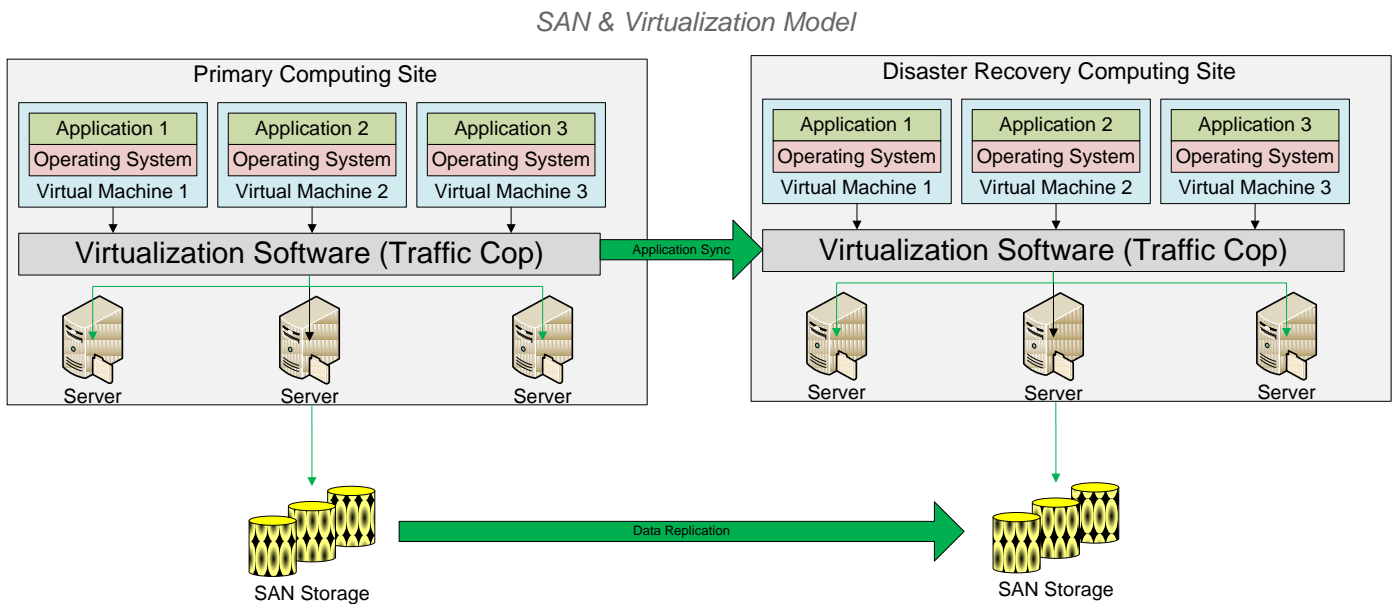


Figure 1 – Data Replication Architecture

DISASTER RECOVERY

SAN Replication and Virtual Server control applications have monitors that continually poll the production environment's availability and efficiency. When preset minimum processing thresholds are not achieved, or when the monitors fail to contact the critical components, the SAN and Virtual Server environment failover processes are automatically triggered. Production processing is transitioned to the disaster recovery site, with customers, partners, and staff transparently redirected.

Although users may see a short interruption in service, few if any transactions will be lost. Typically, only incomplete transactions, those that are currently being updated when the event occurs, will need to be recreated.

The capabilities of SAN based data replication have significant positive impact on Recovery Time and Recovery Point Objectives.

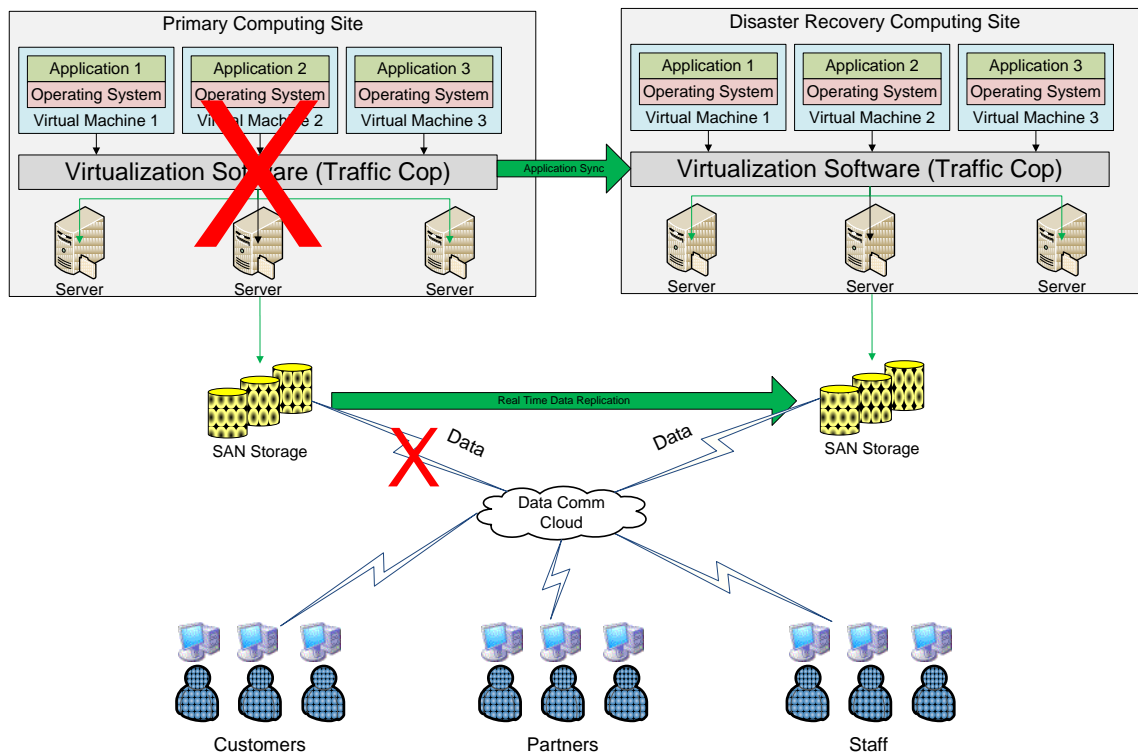


Figure 2 – Automated Fail Over to DR SAN with full copy of Production data