

EXECUTIVE SUMMARY

Storage Area Networks (SAN)

DEFINING STORAGE AREA NETWORK

A storage area network (SAN) attaches shared high performance disks to servers in such a way that the storage *appears* as directly attached to the server while actually contained in the shared SAN infrastructure. Storage is allocated to applications on an as-needed basis, dramatically increasing the rate of utilization by eliminating large amounts of unused disk space. The SAN model also enables storage space to be added dynamically, based on demand, without interruption to operations. All SAN components can be implemented with redundancy, eliminating numerous single points of failure inherent in the older model of tethering storage devices directly to the server they support.

SAN configurations offer numerous operational advantages over siloed (direct server connect) storage models, including improved reliability, performance, and manageability. A SAN advances the Disaster Recovery processes by decreasing time to recover data via superior data backup and recovery methodologies. Other benefits include the ability to take Point-in-Time “Snapshots” of data, enabling the creation of Test and Development copies of production data, central management of the storage environment and automated data de-duplication.

TRADITIONAL ARCHITECTURE

Historically, data centers have created “islands” of disk arrays directly attached to every server. Each island was dedicated to a server/application. Two significant issues are often seen in traditional storage architecture:

1. Underutilized disk space – In order to create reliable data storage, it is often necessary to allocate far more drive space than the application is likely to require.
2. Manual addition of space – When a server or application requires additional storage, dedicated disks are purchased, the application taken off line for a significant period of time, and the new disks are installed and configured. In some cases, the entire data set must be restored, extending down time.

SAN ARCHITECTURE

Breaks in the one-to-one relationship between servers and storage allow disk space to be shared and automatically allocated based on demand for storage. All servers connected to the SAN share a pool of high speed, highly available disk storage via a redundant network connection. Storage is allocated, configured, and tuned for optimum performance based on the differing requirements of each application.

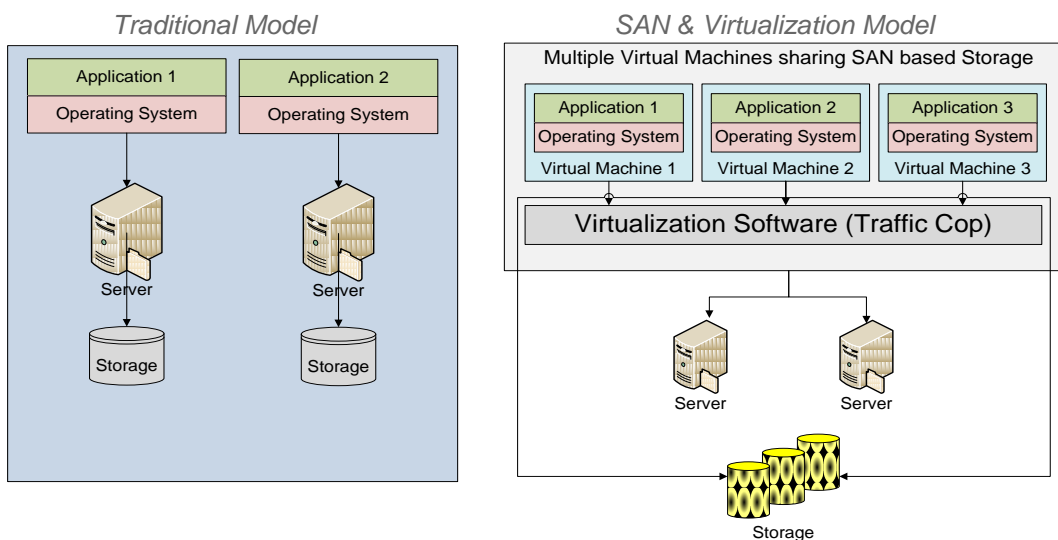


Figure 1 – SAN Architecture

ADDING COMPUTING CAPACITY AND SIMPLIFYING STORAGE MANAGEMENT

Capacity Management via the SAN storage model is highly automated, with software tools monitoring the need for additional storage and adding the required space automatically *without interrupting* production processes. Rather than requiring a technician to shut down the application (loss of business service), add and configure newly purchased disk space, and reload data to the upgraded server, the SAN toolset adds space once preset thresholds are reached, and does so without requiring technical administration or impacting business operations.

Shared storage simplifies data administration and adds flexibility – space is allocated based on demand proactively rather than the traditional reactive model.

The SAN model supports data de-duplication, automatically locating duplicate files and consolidating to a single copy available to authorized users. De-duplication reduces total storage volume by eliminating multiple copies of identical documents and increases efficiency and accuracy by ensuring staff is utilizing the single current version of any file or data set.

Many SAN implementations support multiple storage technologies in a unified environment. Data and applications requiring high performance access will be placed in storage areas comprised of the fastest drive technologies. Data and files requiring normal access will be stored on secure and redundant drive technologies with appropriate access comprised of less expensive drive technology.

Backup and Restore performance can be improved by directly connecting the SAN to the Backup infrastructure. Data destined for tape will no longer have to utilize general network resources, dramatically increasing backup performance while reducing traffic on the production network.

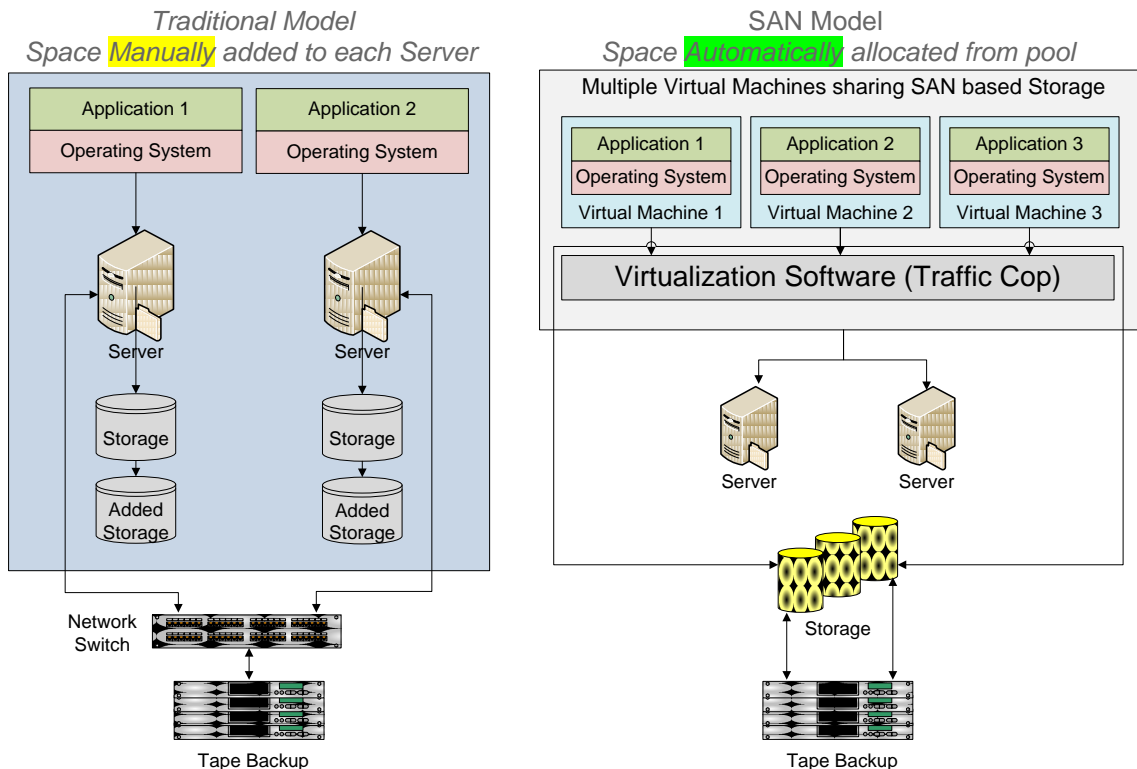


Figure 2 – Adding Capacity and Backup Efficiency

HIGH AVAILABILITY ARCHITECTURE

In historic, data siloed computing environments, when storage becomes unavailable due to an unplanned outage or planned maintenance, the application it supports is not available to users for the duration of the outage. In the SAN model, virtually all components and connections are redundant, which significantly reduces the risk of outage. The SAN can survive the loss of individual disk drives, a network connection, internal wiring, or a power supply without interrupting data availability and business process.

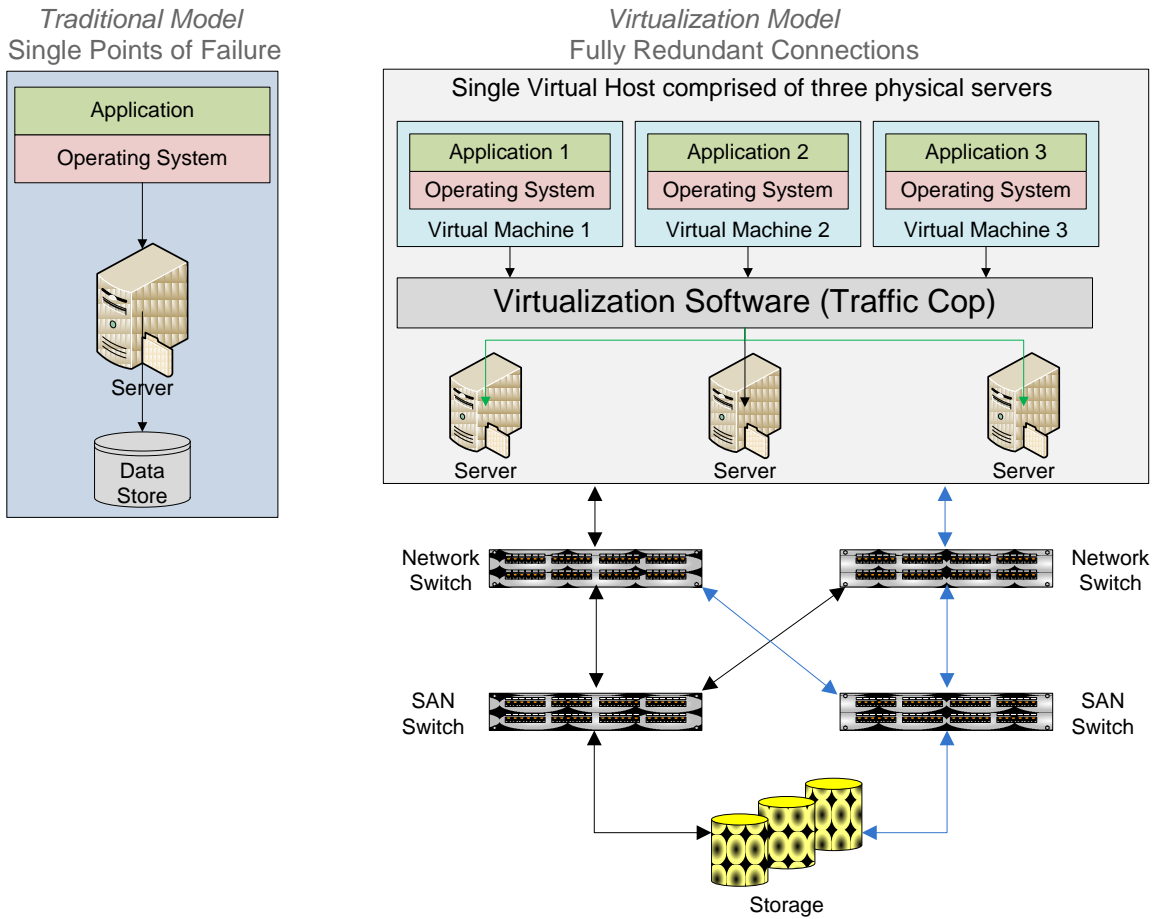


Figure 3 – High Availability

DISASTER RECOVERY

A Disaster Recovery model utilizing server-dedicated storage requires the Information Technology department to reload lost data from tape, a time consuming process that may result in an extended outage of the impacted systems.

The introduction of a SAN will reduce storage failure rates and leverage SAN-specific data management processes to recover data and applications far more rapidly, if necessary. As the application and data reside on the SAN rather than storage directly attached to a failed server, other (virtual) servers can assume operations while the primary server is repaired or replaced. The introduction of SAN specific technologies, such as the ability to take frequent Point-in-Time “Snapshots” of data, allows administrators to recover from an application failure far more quickly than from traditional tape restoration.

The Snapshot feature also enables copies of production data to be made, simplifying testing of both applications and the Disaster Recovery processes without interrupting production activities. There are numerous benefits of Snapshot copies of data including:

- Rapid Creation of test and development data sets
- Test and development are actual copies of production improving QA
- Facilitates testing of patches and updates without impact to production

Further improvement in the Disaster Recovery position can be realized by introducing a SAN at the Recovery site and enabling data replication. Data Replication ensures that the information stored on the Disaster Recovery SAN is kept synchronized with production in near real time. This model enables a seamless transition to the Disaster Recovery site during an event which prevents access to the production site, with minimal down time or lost data.

The introduction of SAN technology, in partnership with virtualization, will reduce Recovery Time and Recovery Point Objectives for key systems from hours/days to immediate/minutes, and the time to recover secondary systems from days/weeks to minutes/hours. SAN technology has been widely adopted by large enterprises for over a decade and by mid-sized businesses over the last five years as lower cost options have become available.