hp
storage
software

november 2002

technical white
paper

hp OpenView continuous access storage appliance

data mirroring, data replication, and snapshots in and between heterogeneous SANs

table of contents

executive summary

To address the challenge of storing, managing, and keeping available increasing volumes of information without significant increases in related costs, many organizations are implementing storage area networks (SANs). A SAN separates the storage systems from the application servers and has the potential to significantly improve the quality of service to storage consumers (applications) while reducing storage management costs. The HP OpenView Continuous Access Storage Appliance (CASA) is part of the HP portfolio of storage products that provides a SAN storage infrastructure to bring real business benefits.

CASA is a SAN infrastructure device that adds data mirroring, data replication, and snapshot functionality to a SAN fabric. Additionally it simplifies, centralizes, and adds flexibility to storage capacity management and allocation. It can be added to an existing SAN or used to create a new SAN, and its wide operating system and storage system support allow the SAN to be built from heterogeneous best-of-breed components. This functionality is delivered with high-performance software that uses the unique HP storage virtualization technology.

CASA is an evolution of the StorageApps SANLink, HP Surestore SANlink, and HP sv3000 appliances that have sold over 700 licenses worldwide. It is built from the high-performance software preinstalled onto industry-standard hardware that is optimized for performance and fault tolerance. CASA is delivered in a standard HP rack system or as a rack-ready product for installation into a customer's rack.

introducing the hp OpenView continuous access storage appliance (CASA)

HP CASA continues the development and evolution of HP SAN storage appliances. The first generation was launched in 2000 as the StorageApps SANLink. The HP SANlink and most recently, the HP sv3000 appliances have followed this. Over 700 licenses have been sold worldwide for these appliances. CASA builds on the strengths of the previous generation appliances and adds a number of valuable new features. CASA is backwards compatible with the HP sv3000.

The value of CASA is derived from the unique HP software that leverages storage virtualization technology to deliver data replication and other services across heterogeneous environments. The software powering CASA is SAN.OS 5.5.

For added reliability, optimized performance and easy installation, the software is preinstalled onto industry-standard hardware. This is delivered in a standard HP rack system or as a rack-ready product for installation into a customer's rack. The CASA hardware is shown in **figure 1**. To further reduce the complexities of installation, CASA is backed with a range of professional installation services from HP.

figure 1. The HP OpenView Continuous Access Storage Appliance is delivered as an integrated solution. These are front and rear views of the industry-standard hardware onto which the software is preinstalled.

installation

CASA can be used to add functionality to an existing SAN fabric, used to build a new SAN fabric, or dedicated for use as a replication engine for a storage system. The appliance provides four ports for storage system connections and 10 ports for host connections, which in the recommended high-availability redundant configuration allows two storage system connections and five host connections. To increase the port count, for both storage system and host connections, Fibre Channel switches are required.

When an existing SAN fabric is upgraded with CASA, the storage administrator can choose to route some or all of the I/O paths through CASA. Once it's installed, hosts continue to access the same data that they used before the upgrade.

Once the physical connections are made, CASA rescans connected storage controllers to recognize presented LUNs. No formatting is done on the connected storage systems; all existing data remains in place and accessible. A registration utility is run on the connected hosts to register their details with CASA. No host software is required, so there is no impact on hosts' processing power.

CASA GUI

CASA has a browser-based management graphical user interface (GUI) that is used for most day-to-day management and monitoring tasks. Use of a secure logon and password allows for remote administration via a Web browser. For initial setup and scripting to automate repeated tasks, there is a keyboard-video-mouse console included, which the administrator can use with the command line interface (CLI).

For organizations that implement CASA remote data replication and for those that have multiple CASAs, these can all be managed from a single GUI as illustrated in **figure 2**. It is a simple task to register multiple CASAs to be managed via a single GUI with a single logon and password. The CASA administrator can also set up visitor logons with only read access.

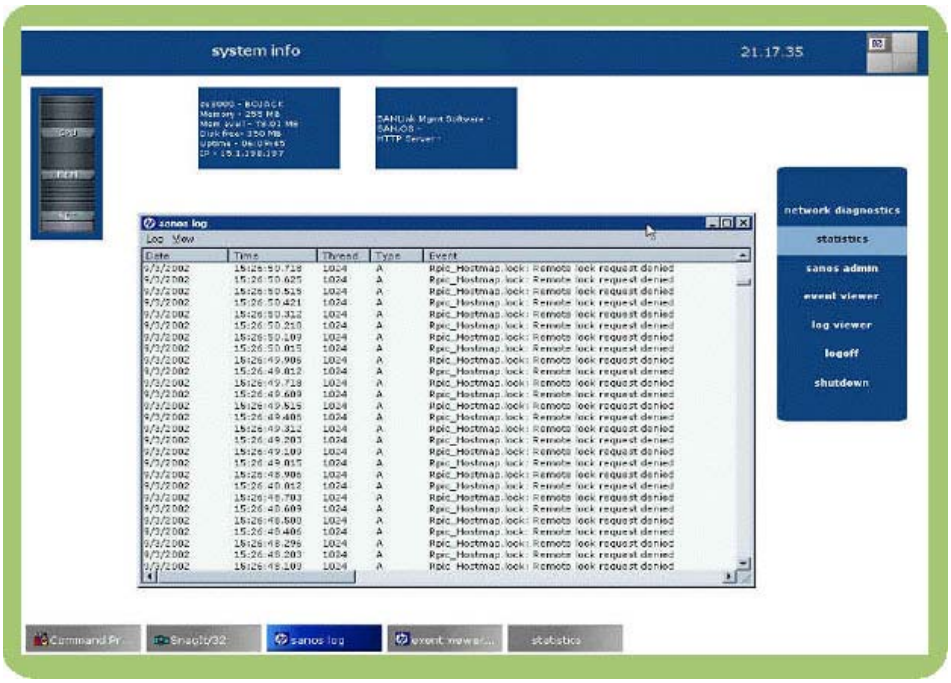

figure 2. CASA can be managed remotely using the secure Web-based graphical interface.

storage management

Once the LUNs on the connected storage systems have been scanned they are managed as CASA virtual disks. The storage virtualization technology allows these virtual disks to be partitioned to make multiple, smaller virtual disks or expanded (concatenated) to make larger virtual disks. To the storage administrator, a virtual disk is managed the same regardless of the storage system that provides the physical capacity. The storage virtualization technology also allows a range of data replication services at the virtual disk level; they are described below.

Managing all the connected storage capacity as virtual disks is very efficient. It allows one storage administrator to manage the same capacity that may have previously required two or three administrators. Consider the case when the storage systems connected to the CASA SAN are from three different vendors—say, HP, EMC, and LSI. In this case, each storage system has a different management program that has to be mastered in order to reconfigure storage and allocate it to hosts. Migrating data from one storage system to another is also a disruptive task that requires considerable planning.

CASA allows storage to be reconfigured and allocated from a single interface. Migrating data between connected storage systems is a painless and nondisruptive task. Additionally, the capacity from all three storage systems can be managed as a consolidated resource or "pool" rather than three islands, as shown schematically in **figure 3**.

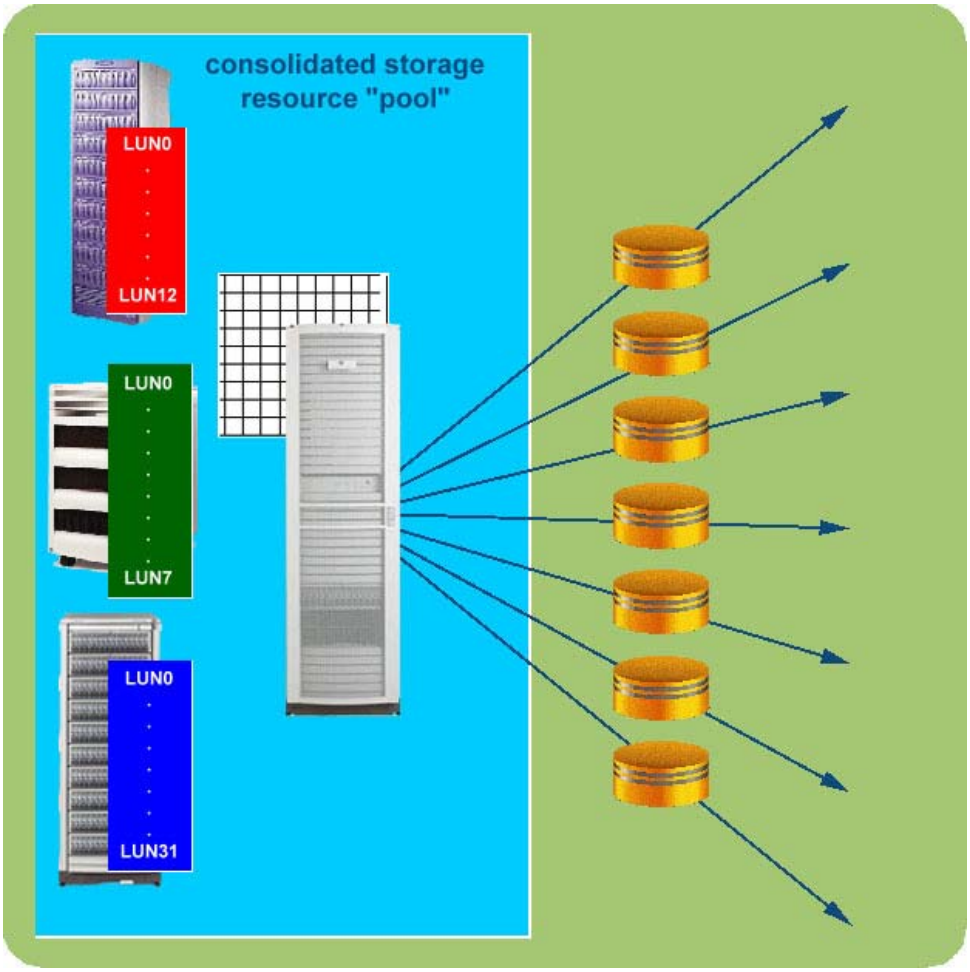

figure 3. Storage capacity connected to CASA is managed as a consolidated resource or "pool." From this "pool" the storage administrator creates virtual disks that are securely mapped to hosts.

true heterogeneous connectivity

CASA provides connectivity to the most widely used host platforms and the leading vendors' storage systems. This speeds up the commissioning of the SAN by removing the requirement for interoperability testing, and it also increases design flexibility. For the storage administrator, existing storage systems and hosts are readily connected to the SAN. More importantly, future purchases of storage systems can be made based on what meets the requirements best, rather than what matches the existing storage systems. This allows best-of-breed storage systems to be added to the SAN and extends the useful life of existing storage systems.

Connectivity testing is an ongoing process in which HP invests considerable resources. Everything on the CASA-supported list has been thoroughly tested. For the latest listing see www.hp.com.

CASA data replication options

The single most important task of a storage administrator is to safeguard an organization's data and ensure its availability. The HP OpenView Continuous Access Storage Appliance makes this task, which is often time-consuming, faster and simpler. But what is unique about CASA is its truly heterogeneous data replication capability and the potential for significantly reducing storage costs.

Within the SAN, or an extended SAN, the CASA data mirroring feature allows fully synchronous data copies to be made. For near-instant point-in-time images, the CASA Vsnap feature is used. To replicate data over longer distances, CASA has an IP gateway that allows appliances to be attached to and connected by a LAN/MAN/WAN without the uses of a storage router.

Each of these data replication options allows the source and target data to be stored on a different storage system. In addition to giving the storage administrator greater flexibility, all of these options can help to greatly reduce the total cost of storage.

local data mirroring

CASA can mirror a virtual disk to any other virtual disk of a similar or larger size. The mirror operates synchronously over Fibre Channel Protocol (FCP). For added protection, it is possible to mirror a source virtual disk to multiple target virtual disks (1-to-n mirroring).

A new feature for the CASA is the capability to extend the SAN fabric by cascading two appliances as illustrated in **figure 4**.

A Fibre Channel interconnect is used so the two CASAs can be several kilometers apart. Each CASA can have connected hosts and storage. The cascading feature allows mirrors to be created between the storage connected to each appliance. It also allows hosts connected to CASA A to access data on storage connected to CASA B and vice versa. This makes mirroring between different buildings or sites simple to implement, and it protects data against local disasters and outages.

The cascading feature also allows bi-directional mirroring, so each CASA can act as a recovery site for the other. It also means that there is full recoverability from a failure in or the total loss of either appliance.

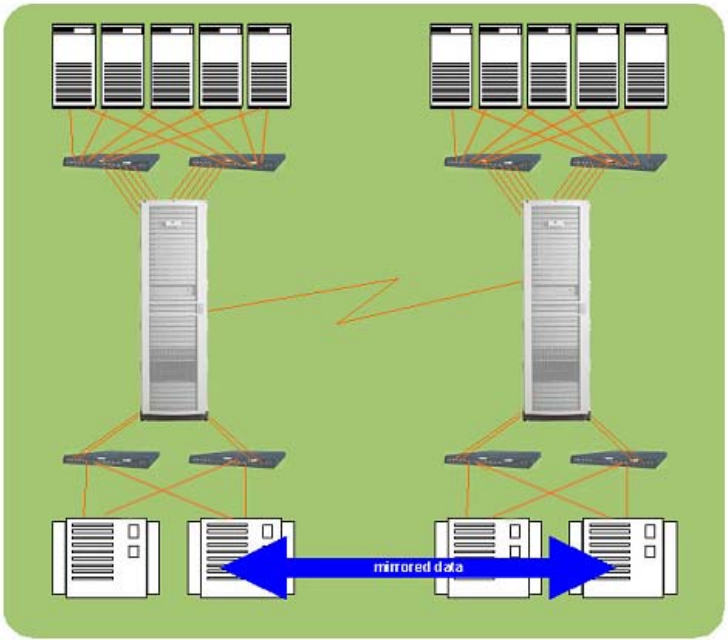

figure 4. The new cascading appliances feature allows the SAN to stretch over a several kilometers to improve resilience to disasters.

Vsnap (snapshot)

The Vsnap (snapshot) feature of CASA creates a near-instant, exact, virtual copy of a specified-source virtual disk at the time the Vsnap is executed. The CASA Vsnap image is a virtual copy of the data in the source virtual disk that can be viewed and used in the same way as the source virtual disk. An advantage over some "snapshot" or "instant image" implementations is that the CASA Vsnap image can be read from and written to.

Unlike the CASA local mirror option, the source virtual disk and the Vsnap copy are maintained independently of each other using a copy-on-write process. Before a modification is made to data on the source virtual disk, the unchanged data is copied so that the Vsnap copy remains an exact replica of the data at the time the Vsnap was executed.

A Vsnap copy can be integrated into the backup process to back up online data at a particular point in time without impacting applications' or users' access time to that data. Another common usage is application development, where the Vsnap copy allows test files and databases to be quickly and frequently made available for testing. Decision support applications require data that they can "slice and dice." A Vsnap copy can be used to make data available faster than by making a full physical copy of the required data, and it also does not put the production data at risk.

The CASA Vsnap feature is very space-efficient. Because the Vsnap copy contains a block-based map of the source virtual disk plus physical copies of only those blocks of data on the source virtual disk that have changed since the Vsnap copy was made, its physical storage requirement is less than the source virtual disk. CASA takes advantage of the reduced physical storage requirements of the Vsnap copies by storing multiple Vsnap copies in a single virtual disk. This virtual disk is known as the backing store. The physical size of the backing store may need to be only 10%–20% that of the source virtual disks. This is illustrated in **figure 5**. All the Vsnap copies in the backing store are managed as regular virtual disks.

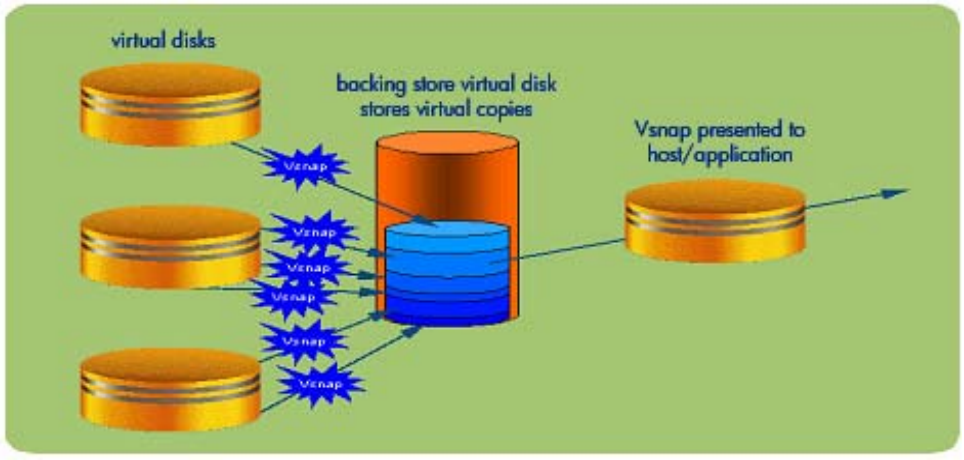

figure 5. The Vsnap feature creates near-instant virtual data copies that require much less physical storage capacity than the source virtual disks. These virtual copies can be mapped to hosts and used just like regular virtual disks.

long-distance data mirroring

For business continuity and data protection, mirroring data to a remote location is a widely used method. Of the many long-distance data mirroring solutions available today, none is as cost-effective as the CASA solution.

This long-distance data mirroring solution requires a CASA at both the local and the remote location as shown in **figure 6**. A Gigabit Ethernet (GbE) port in each CASA provides an IP gateway from the SAN that is used to connect the CASAs to an IP-based network. This can be a dedicated network, a shared network, or even the Internet. Each CASA has two GbE ports available for long-distance mirroring, so for resiliency two independent paths can be specified between the local and the remote CASA.

figure 6. The CASA allows data to be mirrored synchronously or asynchronously between geographically remote locations using an IP-based network.

Any virtual disk on the local CASA can be mirrored to any equal or larger sized virtual disk on the remote CASA. Mirroring is also possible from the remote appliance to the local CASA. Each mirror pair created can be specified as an asynchronous or synchronous mirror. The choice depends upon the storage administrator's requirement for data currency at the remote location versus I/O latency at the local storage system.

Specifying synchronous mirroring creates a real-time copy of the source virtual disk at the remote location. Writes are only acknowledged to the application after they have been committed to storage that is connected to the local CASA and to the storage connected to the remote CASA. This adds latency to the writing application's I/O. As long as there is a connection between the two CASAs, a consistent image of the virtual disk is guaranteed at both locations. Specifying asynchronous mirroring creates a copy of the virtual disk at the remote location that trails the source virtual disk at the local site. This does not add latency to the writing application's I/O.

The performance and reliability of the connecting network(s) must be considered when specifying asynchronous or synchronous mirroring. To use asynchronous IP mirroring, a high-performance, probably dedicated network is required. For synchronous mirroring, the connection needs to be reliable because the host application will see I/O failure if the connection(s) is down. Using asynchronous mirroring, writes are queued on a local journal before being transmitted to the remote storage. The journal file can be sized, up to 1 TB, to compensate for low-bandwidth links or link outages. Writes are dispatched from the journal to the remote CASA in the same order that they were committed locally. Time consistency is guaranteed and allows a recoverable image of data on the remote side.

Compared to storage system controller-based solutions for remote data mirroring, using CASA has several advantages. The main advantage is the significantly reduced cost. Additionally, CASA can mirror data between unlike storage systems from multiple vendors. However, it must be remembered that the CASA remote mirroring solution does not provide automatic failover and failback in and following site outages. If this type of nonstop access to data is required, a solution such as HP StorageWorks Continuous Access XP is required.

The reduced cost of the CASA remote data mirroring solution compared to storage controller-based solutions means that many more businesses and organizations can afford to implement remote data mirroring. Many organizations that currently rely on local tape backups that are transported offsite for disaster recovery and business continuance can now enhance their data availability and disaster recovery options using CASA remote data mirroring.

As well as being implemented for business continuity and disaster recovery, CASA facilitates backing up data on geographically dispersed SANs to a central location. Data managed by multiple CASAs can be mirrored to a single CASA. This n-to-1 mirroring feature allows a single appliance to be a target for up to three source CASAs. In this way, up to four locations can have their backup consolidated to a single location.

combining data replication options

The data replication options described above can be combined to add resiliency to the storage infrastructure and to take processing requirements off production storage systems. Some examples of this are:

- By using cascaded CASAs, local synchronous data mirrors can be created between storage systems in different buildings on a campus. One building houses the production storage systems; the other houses the storage used for holding the data mirrors. The secondary storage systems can be mirrored to storage systems at a remote site by using the IP-based mirroring functionality. This provides two levels of mirroring for a high level of data protection while maintaining the speed and easy recovery of synchronous mirrors.
- A CASA virtual disk can be a source for a local (FCP-based) mirror and a long-distance (IP-based) mirror. This allows data to be mirrored locally and remotely for added protection.
- Using either the FCP-based or IP-based data mirroring, a virtual disk can be mirrored to a secondary storage system. The target virtual disk can be used as a source for Vsnaps. These Vsnaps can be used for online data recovery as well as for archiving data to tape.

usage scenarios

scenario 1: storing database copies off the production storage system (figure 7)

Almost all organizations rely on databases for their key business processes. Many of these databases are classified as mission critical and consequently are protected using online recovery solutions that are integrated with high-end storage systems such as HP Continuous Access XP or EMC Symmetrix SRDF.

In addition to keeping database copies for recovery purposes, they are also kept for testing, development, running queries, and archiving. Often three, four, five, or even seven database copies are maintained. In cases where these frequently static copies are currently stored on the same storage system as the production database, it is beneficial to move the copies to other, lower-specification storage systems. Adding CASA to the storage environment gives the storage administrator the flexibility to do this.

Typically, the primary database and one or more mirrors are maintained on the production storage system. One of these mirrors is presented to CASA, which allows copies to be created and stored on more cost-effective storage systems. This enables organizations to move the copies for test, development, and queries to less-expensive storage, which has an immediate effect of reducing the cost of storage. This also enables the production database to have more of the I/Os and bandwidth of the high-end storage system.

A space-saving alternative to physical database copies is to generate virtual copies using the CASA Vsnap feature. A Vsnap produces the copy almost instantly and is very space efficient, typically taking only 10%–20% of the physical space of source data. With high-end storage systems costing more than $0.10/MB, using Vsnap copies can save considerable amounts of money.

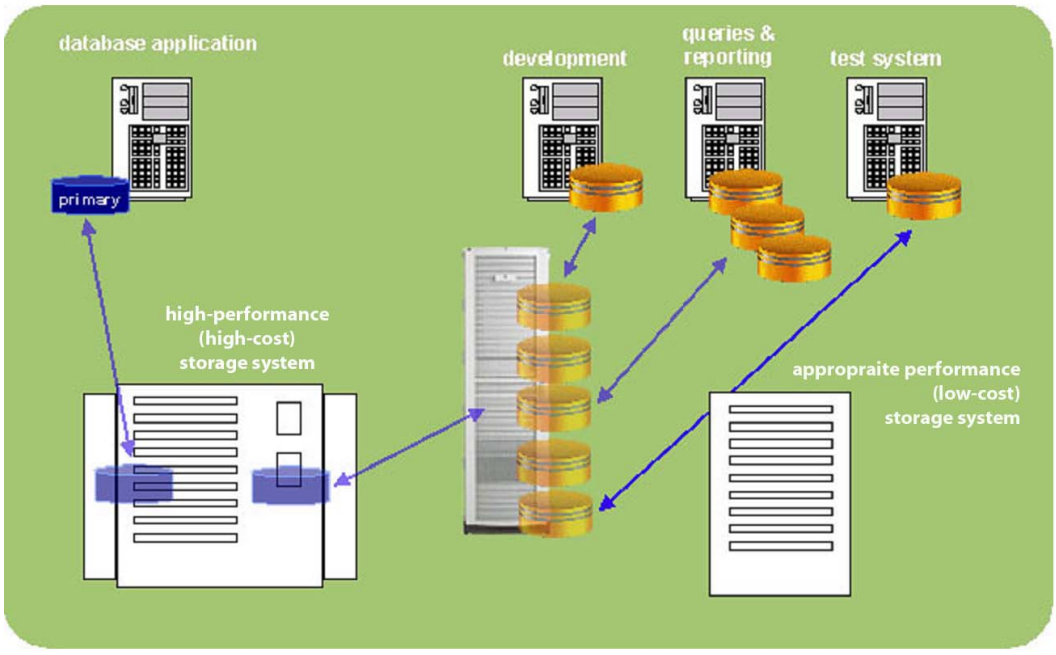

figure 7. CASA allows the primary databases and database copies to be stored off the production storage system, which reduces the cost of storage, reduces space usage (with Vsnaps), and improves database application performance.

scenario 2: online and offsite data backup

Cost is the reason why many more organizations do not have an online recovery solution for their data. Even those organizations that do have an online recovery plan usually only have this for a subset of their data. The expense of high-end storage and integrated disaster recovery solutions such as HP Continuous Access XP or EMC SRDF is only justified for mission-critical, as opposed to business-critical, data.

Tape backup is a proven and reliable solution for data protection that is significantly less expensive than online recovery models. However, despite large improvements in throughput in recent years, recovery from tape can take many hours plus time to bring the media from offsite storage.

CASA facilitates a backup solution that allows faster recovery than tape-based solutions and less data loss, in the event of a catastrophic failure, than nightly tape backups. It is also significantly less expensive than storage system based online recovery models.

Consider an organization that has a head office with several branch offices. Each location has a backup regime that requires a nightly backup to tape. Media is stored off-site. This organization cannot justify moving to a storage system based online recovery solution, even for the head office, but wants to speed up recovery in the event of data loss and reduce its exposure to the possibility of data loss.

Connecting the storage systems to the hosts via CASA adds the capability to move data to an online recovery model. Data to be protected is presented to CASA. Using the Vsnap function, point-in-time images are made of the data at regular intervals. These are held online in case recovery is needed from user error, viruses, or software errors. These Vsnaps are also mapped to a host running a backup application so data can be periodically archived to tape to protect the organization's data against hardware failure.

This implementation can be done in stages with increasing amounts of data accessed via CASA. Further efficiency improvements can be gained by using the remote data mirroring function of CASA as illustrated in **figure 8**. By installing CASA in branch offices, data can be replicated over an IP network to another branch office or the head office. In addition to protecting data by moving it offsite, this also allows the tape backup task to be consolidated to a single location; one CASA can act as a remote mirroring target for up to three CASA source appliances, allowing a single remote site to be allocated as the failover site. A typical data protection plan with this setup would be to mirror data from the branch offices to the head office, then to create Vsnaps of the mirror targets. From these Vsnaps, the storage administrator can run backups to a central tape library.

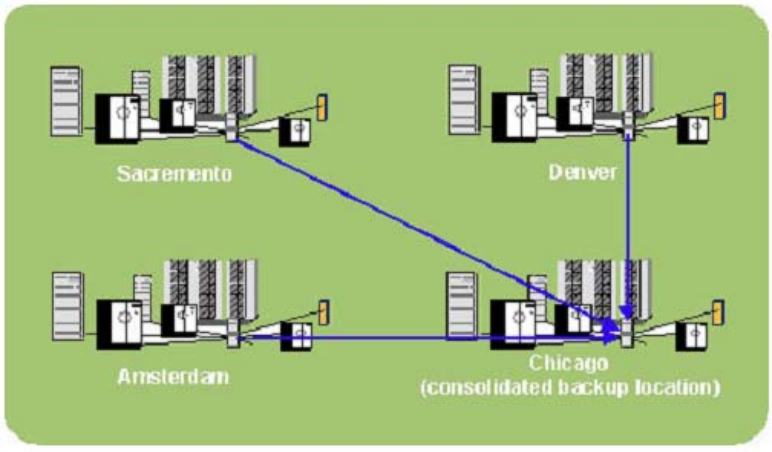

figure 8. Consolidating backup to a single location is made possible using the n-to-1 mirroring capability of CASA.

scenario 3: migrating data to a new storage system

Data migration is a task that can take storage administrators a considerable length of time and lead to data unavailability. The task of data migration has added complexity when the source and target storage systems are different models or from different vendors as shown in **figure 9**.

CASA provides a simple way to migrate data between like and unlike storage systems. The source and target storage systems need to be connected to CASA and the necessary LUNs presented to CASA. The local FCP-based mirroring feature is used to make a copy of the source data on the target storage system. After mirror set synchronization is complete, the CASA virtual disk presented to the host is changed to the virtual disk containing the copied data. The mirror is then broken. This can be done without impacting host applications.

In addition to allowing new storage systems to be integrated quickly, CASA provides a nonintrusive solution for environments with changing availability requirements. Data can also be migrated between storage systems to enable higher levels of availability and protection.

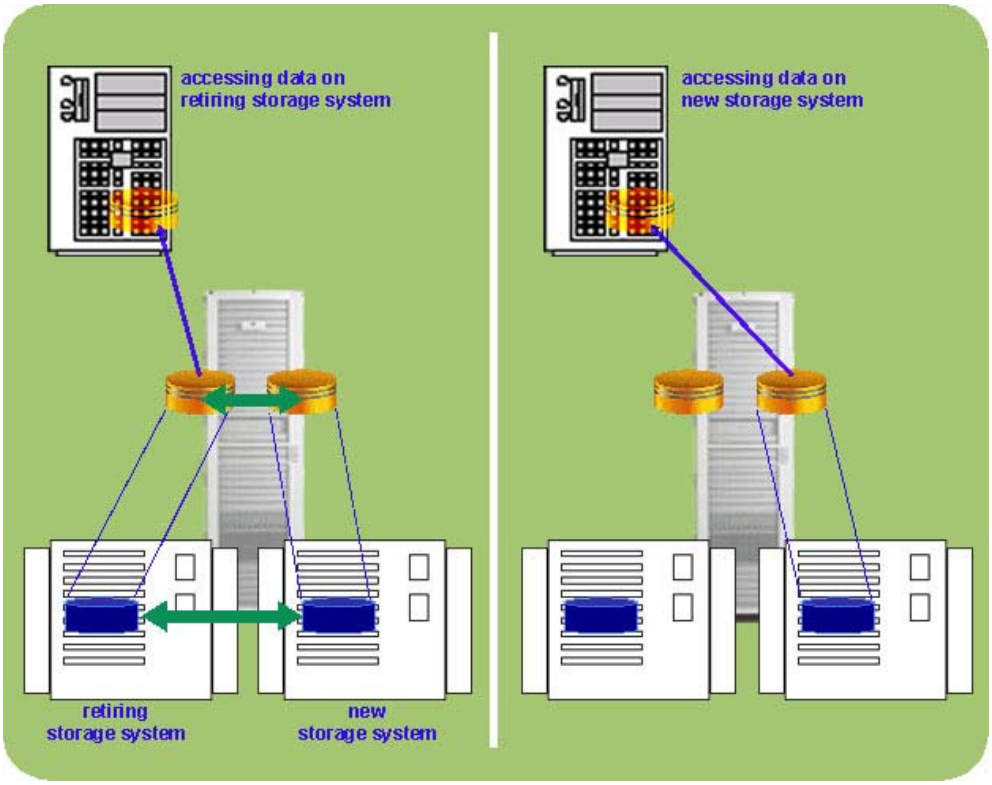

figure 9. Data migration is simplified using CASA, greatly reducing the expense and planning required for this task.

data security

The HP OpenView Continuous Access Storage Appliance is located in the center of the SAN fabric. All file and application data as well as control data for the storage managed by CASA passes through it. This facilitates tight security, with CASA managing a host's access to virtual disks.

CASA allows the storage administrator to centrally manage the access to data on multiple heterogeneous storage systems. The security database is stored on the appliance for all connected hosts and storage devices.

A host added to the SAN cannot access any data on LUNs assigned to CASA without the administrator first setting permission from CASA. In effect, all virtual disks are masked from hosts until the storage administrator maps them to a host. This allows a more granular access policy than is possible using only fabric switch security features such as zoning.

boot from SAN

The ability to boot a host from the SAN is useful for recovery following server hardware failure. Storing the boot files and operating system, configuration details, and applications on a CASA virtual disk allows a new host to be connected and booted from the SAN. This is a much faster recovery than a tape-based recovery or reloading the operating system and applications from the original media.

It is also possible to reduce platform size by configuring each host with its own CASA virtual disk as a boot device. In addition, this boot device can be mirrored to a remote site, so in the event of a local storage system failure, the host can be mapped to the remote boot device.

performance and scalability

All file and application data I/O and control data passes through CASA. It is a pass-through appliance similar to a Fibre Channel switch. The sustained performance is up to 80,000 I/Os per second (IOPS) with a data throughput of up to 800 MB/s. To achieve this performance, CASA is built using best-of-breed industry-standard hardware. The SAN.OS program, which provides the intelligence and functionality of CASA, is a multi-threaded application that introduces minimum latency into the I/O path.

Without any Fibre switches, up to two storage systems and five hosts can be connected to CASA in a high-availability redundant configuration. In a nonredundant configuration the number of hosts and storage connections is doubled. In a CASA SAN with additional switches, up to approximately 20 hosts and four midrange storage systems or a single enterprise-class storage system is the recommended maximum configuration. This obviously is I/O profile dependent and depends upon the type of host application and the size of the storage system. As a general rule, approximately 150 TB of heterogeneous capacity is a recommended maximum for a single CASA. The new capability of cascading CASAs doubles the recommended maximum hosts and storage system connections in the extended SAN. Up to 4,096 LUNs from the connected storage systems can be managed by CASA.

product information

The HP OpenView Continuous Access Storage Appliance (CASA) is available direct from HP or through a limited number of qualified HP partners. Options are available to extend the functionality and extend the warranty of the base product to meet exact customer requirements. Base warranty is one-year, same-day service during business hours (08:00 to 17:00). Installation services must be ordered at the same time as the product. The level of installation service depends upon the size and complexity of the environment.

A7540AZ	HP OpenView Continuous Access Storage Appliance (factory racked)
A7540A	HP OpenView Continuous Access Storage Appliance (rack ready)
U2418A	installation services – simple environment
U2419A	installation services – intermediate environment
U2420A	installation services – complex environment
A7541A	add IP data replication functionality
A7542A	add Point-In-Time Image functionality
A7543A	add heterogeneous (non-HP) storage connectivity
A7544A	add IP data replication and Point-in-Time Image functionalities
H4405A #8JC	warranty upgrade to 1-year, 24 x 7
H4403Y #8JC	warranty upgrade to 3-year, same-day, during business hours
H4405Y #8JC	warranty upgrade to 3-year, 24 x 7

for more information

www.hp.com/go/san

11/02