Continuous Access Storage Appliance

10

The HP OpenView Continuous Access Storage Appliance (CASA) solves a wide range of problems that may be encountered in enterprise storage environments. This chapter provides an overview of CASA as well as information about how to integrate CASA solutions into general HP StorageWorks Fibre Channel SAN installations. The following topics are discussed in this chapter:

- Overview of CASA
- How CASA Works
- CASA Features
- CASA Management
- Security Implications of CASA
- Supported Systems and Software
- Configuration Rules
- CASA Services
- Additional Information Sources

Note: The information in this chapter is specific to version 5.6.1 of CASA.

Overview of CASA

CASA provides data replication on SANs consisting of heterogeneous mixes of servers and RAID array storage devices. By making all available storage capacity accessible by all servers—with appropriate access controls as required—CASA helps you optimize the use of the server and storage systems in your installation. Data may be placed on the storage device that makes the most sense, regardless of server driver, host bus adapter, or operating system.

CASA supports data mirrors and point-in-time snapshots. These may be done between heterogeneous storage devices. By providing the opportunity for flexible data placement, CASA allows you to distribute redundant copies of data on the storage device most appropriate for each specific type of copy.

CASA can be used to migrate data between heterogeneous storage devices. By removing limits to where data is stored, CASA facilitates the retirement of legacy equipment and the addition of new equipment. If your data availability specifications change, CASA helps you adapt existing data to different availability configurations. As your data center requirements change to meet new business conditions, CASA provides the adaptable data placement and migration tools to optimize your new SAN configuration.

CASA provides an incremental approach to storage virtualization, because it works in conjunction with traditional SAN solutions. If only a subset of your data requires replication or migration, then adding CASA to your existing SAN won't disturb the rest of the data.

The features and supported configurations described here reflect CASA SANOS software version 5.6.1.

Figure 57 is an overview of a typical CASA configuration.

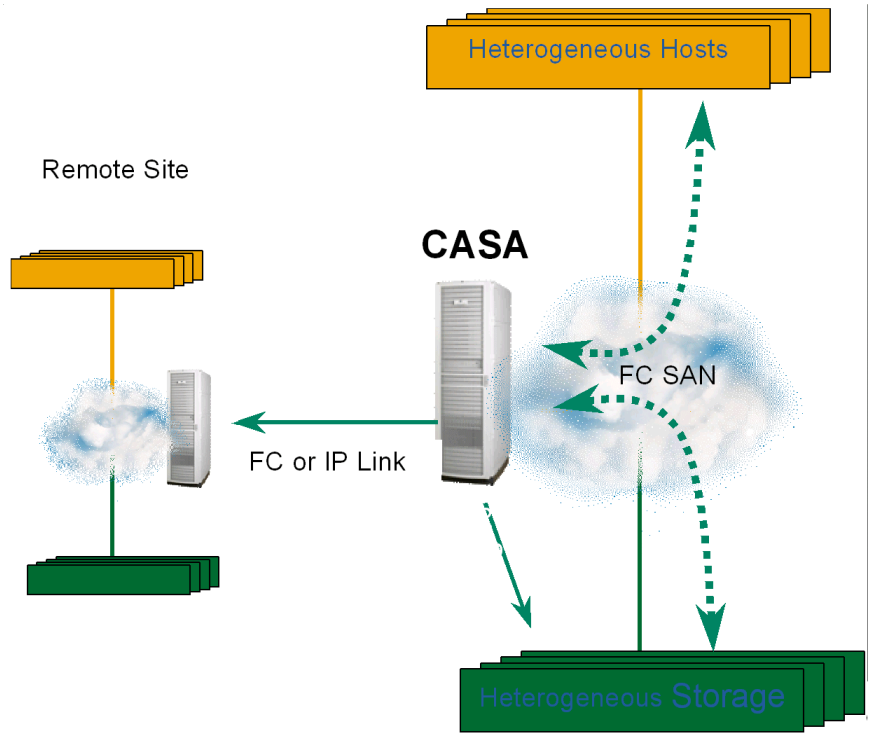

Figure 57: Typical CASA Deployment

How CASA Works

The Continuous Access Storage Appliance consists of the following components:

- Physical application servers and physical storage arrays
- A CASA appliance with two internal nodes and shared metadata storage
- The CASA utility software
- Management tools used to manage the operating characteristics of the CASA
- Optional racks to hold the appliance, servers, storage, and related equipment

Application servers are connected to CASA using traditional Fibre Channel (FC) Host Bus Adapters (HBAs), cables, and switches. RAID storage arrays are connected using FC cables and switches as required for the specific configuration. The detailed requirements for these components are discussed in this chapter.

Figure 58 shows a schematic of the internal architecture of the CASA. The target ports present virtual disks to the physical servers, while the initiator ports present virtual servers to the physical storage arrays. Shared storage is used within the appliance to store metadata information about the virtual devices. Gigabit Ethernet ports are used to connect the nodes in the appliance.

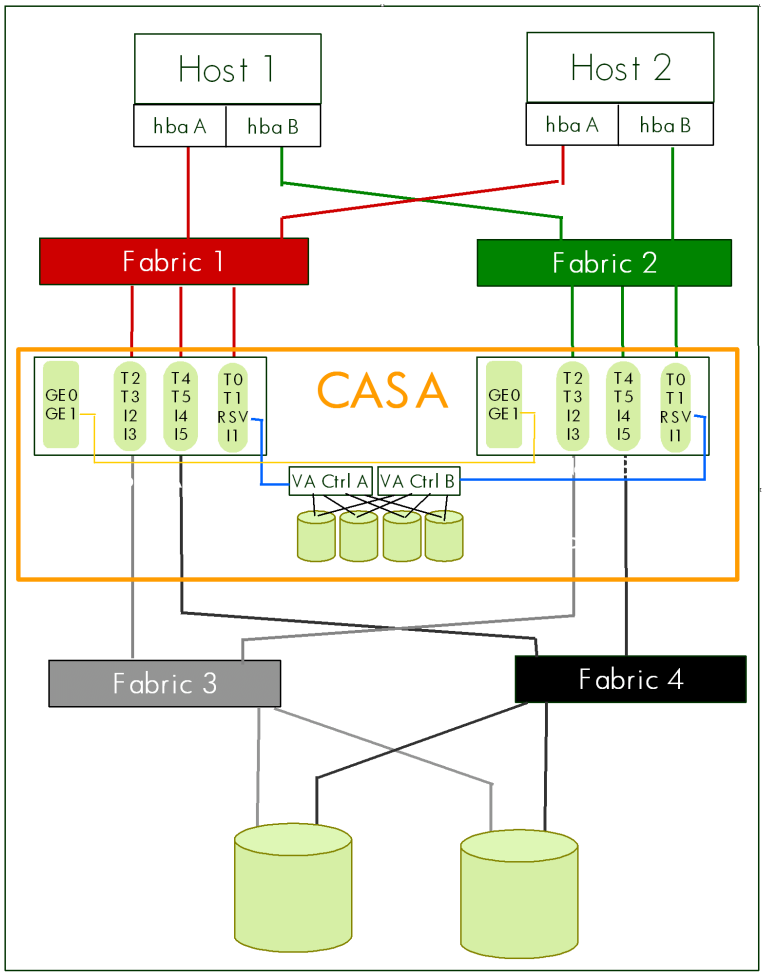

Figure 58: CASA Internal Architecture

The CASA software implements two kinds of virtual devices:

- The appliance presents virtual target logical units (LUNs) to the application servers.
- The appliance presents virtual initiators to the storage arrays.

The servers "see" the virtual storage devices provided by the appliance, but not the physical storage arrays in your SAN. Similarly, the storage arrays "see" the virtual servers provided by the appliance, but not the physical servers in your SAN.

This isolation of the physical servers from the physical arrays provides the opportunity for tremendous flexibility in the deployment of the servers and storage in the SAN, and is a core element of the storage infrastructure that supports the HP Adaptive Enterprise environment.

Changes to the array configuration can be made without the knowledge of the servers, and changes in the server configuration can be made without the knowledge of the arrays. For example, failures in disk arrays can be made completely transparent to the servers. The failure recovery mechanism in the array works with the virtual server in the appliance to manage the failure, but the appliance masks this activity from the servers on the SAN. Certain appliance failures are visible to the servers—in the same way that array failures are visible to the servers in a traditional SAN—but because all of the storage provided to the server is from the appliance, there is only one storage failure model that the server needs to handle. This means that a server may connect to storage capacity provided by a heterogeneous mix of storage array types, while only implementing a single storage failure handling model. This failure handling model is the one provided by the appliance.

Appliance Ports and Paths

Each CASA appliance is fully redundant, and includes two peer nodes, each with:

- 6 target ports (host-connect)
- 5 initiator ports (storage-connect)
- 1 dedicated initiator for shared metadata storage
- 3 Gigabit Ethernet ports

The CASA appliance has a total of 12 target ports and 10 initiator ports.

Shared metadata (data about the data) is stored on fully redundant dual controller local storage.

The CASA appliance has redundant paths:

- From hosts to CASA
- From CASA to storage

CASA Features

Using the method described above, the CASA provides the following capabilities:

Storage Pooling

Under the control of CASA, physical SAN storage is collected into a virtual capacity pool. Unused storage capacity (“stranded capacity”) can be allocated from the virtual capacity pool and then assigned to where it is needed. Mirroring and related replication technology can be applied to the pool in a flexible fashion, without disrupting the servers’ view of the virtual devices. This capability optimizes the utilization of existing storage capacity.

Local Data Replication

1. Data replication. Data in a given LUN may be mirrored to up to nine other LUNs. This adds to the reliability of the data stored in the physical arrays by protecting against array failure. Mirroring can be used to add flexibility to your backup process by allowing multiple copies of the data to be available at one time.
2. Data snapshot. The Vsnap feature creates a space efficient point-in-time image of a LUN. This capability can be used to create additional static copies (up to nine) of databases or other information. Typically, Vsnap uses only a fraction of the space that would be required for a full mirror of the LUN.

Remote Data Replication

1. For CASA Fibre Channel replication (FCP mirrors) or for Synchronous IP replication, CASA's must be within a campus configuration, typically within 40 kilometers.
2. A cascaded configuration supports 2 CASAs only.
3. Remote data replication. Mirror copies may be made (up to nine copies) between multiple CASAs, providing disaster tolerance and the ability to recover quickly from a site failure.
4. Remote cross mirroring between two appliances. The virtual storage capacity pool is distributed across the two sites, and servers at each site may have local and remote mirrors. This provides a fully disaster tolerant configuration that can withstand the failure of either site. For more information on this CASA application, refer to http://www.hp.com/products1/storage/products/virtualization_appliances/network/sv3000/infolibrary/CASA_CAMs.pdf
5. Synchronous and asynchronous mirroring. Both types of remote mirroring are supported. In the synchronous case, I/O operations issued by a server are not reported as complete until the remotely replicated operation has completed. In the asynchronous case, I/O operations are reported as complete when the local operation completes, which improves performance in cases where distance-related latency is undesirable. Both cases provide guaranteed write ordering technology, so that a disaster recovery operation will have a coherent data image with which to continue operation.
6. "N to 1" replication (for IP replication only.) Up to three sites can be mirrored back to a single central site. This may be used to support centralized backup of multiple sites. This feature allows the use of asynchronous replication to all of the the cascaded sites: No snapshot is required to handle multiple sites.

Figure 59 shows a cascaded configuration.

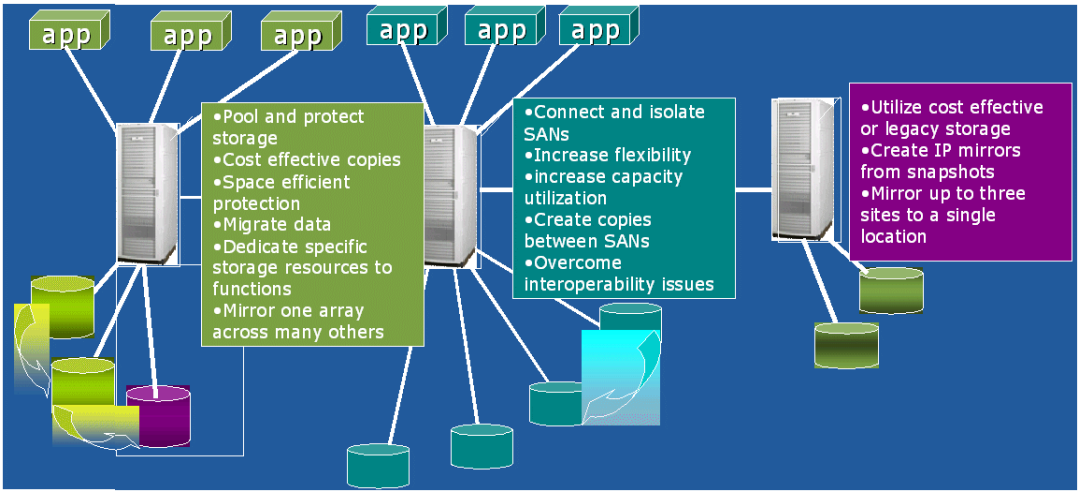

Figure 59: Cascaded CASA Configuration with Three Sites

IP/FCP Mirroring

The CASA appliance contains a pair of redundant nodes. These nodes work in tandem as peers to provide a high level of availability to the CASA system. "IP/FCP mirroring" is used to maintain coherence between the two nodes. A journal is maintained on a shared metadata disk array built into the CASA, making the appliance fully redundant. If one node fails, the other node has full access to the shared journals and can fully recover the data.

Mirroring availability continues under a variety of failure conditions:

- Local storage failure
- Remote storage failure
- Either IP link failure
- Either CASA node failure
- In the case of either IP link failure, SCSI requests are routed to the peer node
- If one node loses access to storage, SCSI requests are routed to the peer node

Heterogeneous Storage

In all of these cases, heterogeneous mixes of storage arrays are supported. CASA contains HBAs, HBA firmware, driver software, and path failover software appropriate for use with all supported combinations of storage array devices. Since the device characteristics are hidden from the application servers, heterogeneous mixtures of storage array devices may be used without requiring any changes or special configuration options in the application servers.

This powerful feature adds considerable flexibility to the HP StorageWorks SAN. Benefits include:

- Existing arrays may be mixed with new arrays to support data migration.
- Low cost arrays may be mixed with enterprise-class arrays to optimize cost.
- Migration from one type of array to another may be done without impact to the servers.
- Configuration changes required to support new storage requirements may be done without interfering with production work.

Many other important applications for this feature may be imagined without difficulty.

CASA Management

CASA environments are managed using the CASA Management Service (AMS), a centralized web-based user interface. The management service is used to configure and control all aspects of the CASA system. All CASA features are presented in a common fashion to make it easy to control both local and remote devices and the distributed virtual capacity pool. AMS implements a secure management interface for all CASA-related functions.

CASA Graphical User Interface

The CASA graphical user interface (GUI) supports remote management of CASA appliances. It provides navigation between multiple CASA nodes and appliances without having to login, and provides optional access to the command line interface (CLI) if needed. The GUI incorporates settable user privileges to provide selective access to management operations.

CASA Command Line Interface

The CASA command line interface (CLI) provides remote console based management of CASA appliances. It uses a UNIX shell-like interface that has scripting capability, the ability to process multiple requests from a single file, and flexible navigation between CASA nodes and appliances. The scripting capability allows a CASA to be managed by third party clients.

AMS Server

The AMS server software runs on the CASA, and is responsible for handling user management requests from management clients, either the GUI or the CLI.

The GUI or CLI sends an XML request to the management server, which in turn performs the required validation and translates the request to a command that is understood by the appliance. The XML handler is capable of processing management requests for the appliance backend engine and for B-Series switches.

The appliance processes the request, and then sends and appropriate response to the management server, which in turn creates an XML response message and sends it back to the requesting client.

Prior to forwarding any request to the appliance backend engine, the management server first authorizes the request with the security service, as discussed in the CASA security section, below.

Integration of CMS with OpenView SAM

OpenView Storage Area Manager (OpenView SAM) is used to handle SNMP traps generated by CASA. OpenView SAM 3.0 Suite DPIs are available for integration with Storage Node Manager, Storage Builder, Storage Optimizer, and Storage Accountant. These provide centralized discovery, mapping, performance planning and management, and billing capabilities.

Additional Information About CASA Management

Refer to the *HP OpenView Continuous Access Storage Appliance System Administrator's Guide* for additional information about managing the CASA system.

Security Implications of CASA

Traditional networked storage systems deliver a high level of security. In many cases this security is built into the SAN, because typical SANs are constrained to fairly small physical areas (such as a single machine room, single building, or single campus) and because SAN infrastructure components (such as Fibre Channel switches) incorporate various security control methods. In those cases where a SAN is extended beyond these limits, additional techniques (like encryption of data passed on extended links) must be used to maintain a suitable level of security.

Security Features

CASA systems achieve a level of security similar to that of traditional SAN systems by the use of strong access controls and redundancy.

- Passwords protect against intrusions through the management interface.
- Every CASA system is designed using a no-single-point-of-failure topology with redundant components and redundant meta-data storage.
- LUNs are mapped to hosts by unique worldwide name (WWN) to protect data from access by unauthorized servers. New hosts on the network have no access until LUNs are explicitly mapped.
- Hosts can have exclusive storage for independent applications or shared storage to enable failover for clustered applications
- Mapping is network-based, so no host software is required.

The security component provided by the CASA Management Service (AMS) provides ticket-based authentication and authorization for services on a CASA appliance. This is used by AMS to control access to the CASA appliance, B-Series switches, and its own administrative interface. It also provides an audit trail of authentication and authorization operations, as well as of its own administrative operations.

Security Services provided:

- Identification and Authentication:
  - Challenge-Response Mechanism. Password is never transmitted over the wire.
  - Encryption Based on Shared Knowledge of the Password and User ID.
  - 128 bit Encryption utilizing Blowfish
  - Result: Ticket is Granted
  - Tickets have a tunable timeout—Default is 8 hours.
  - Originator IP Address is contained within the encrypted portion of the Ticket.
  - Ticket must be passed with every request.
- Authorization:
  - Authorization request contains: Ticket, Originating Host, Comma separated list of requested operations
  - The requestor must have privileges required for ALL operations to allow any to be performed.
- User and Role Administration:
  - Add/Mod/Delete/Query Users
  - Add/Mod/Delete/Query Roles
  - List Roles for a User
  - List Privileges for a User

Architectural Advantages of this approach include:

- XML is a standard language that allows an open, human-readable protocol.
- The security service is usable by clients written in any language that can output XML on a socket connection.
- Implementation in Java enables platform independence.
- XML-based socket level protocol

The strong security features of AMS provide a high level of protection against intrusion through the management interfaces. In addition, if someone were to obtain unauthorized access to the appliance itself, a valid account and password are required to use the GUI or CLI even from the local machine.

Refer to Chapter 9, "SAN Security" for additional information.

Supported Systems and Software

CASA 5.6.1 supports the following Fibre Channel SAN switches, storage arrays, and server operating systems. Contact your Hewlett-Packard representative for information on specific supported models and version numbers.

Supported Fibre Channel SAN Switches

CASA supports the full line of HP StorageWorks SAN switches, as shown in the following tables.

Table 52: HP StorageWorks B-Series Product Line Switches

HP StorageWorks Switch Name		Firmware Version	Number of Ports
HP StorageWorks MSA SAN switch 2/8		3.1.1c	8
HP StorageWorks SAN Switch 2/8 EL, 2/8 Power Pak			8
HP StorageWorks SAN Switch 2/16, 2/16 EL, 2/16 Power Pak			16
HP StorageWorks SAN Switch 2/32, 2/32 Power Pak		4.1.2b	32
HP StorageWorks Core Switch 2/64, 2/64 Power Pak			64 (2 switches per chassis, for a total of 128 ports per chassis)
HP Switch Name	**Compaq StorageWorks Switch Name**		**Number of Ports**
HP Brocade 2400 (HP reseller)	Compaq StorageWorks SAN Switch 8	2.6.1c	8
N/A	Compaq StorageWorks SAN Switch 8-EL		8
HP Brocade 2800 (HP reseller)	Compaq StorageWorks SAN Switch 16		16
N/A	Compaq StorageWorks SAN Switch 16-EL		16
HP Surestore FC Switch 6164 (64 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/32 (64 ISL Ports)		32 (counts as 6 switches and 2 hops when applying configuration rules)
HP Surestore FC Switch 6164 (32 ISL Ports)	Compaq StorageWorks SAN Switch Integrated/64 (32 ISL Ports)		64 (counts as 6 switches and 2 hops when applying configuration rules)
HP Surestore FC 1Gb/2Gb Entry Switch 8B	N/A	3.1.1c	8
N/A	Compaq StorageWorks SAN Switch 2/8-EL		8
N/A	Compaq StorageWorks SAN Switch 2/16-EL		16
HP Surestore FC 1Gb/2Gb Switch 8B	N/A		8
HP Surestore FC 1Gb/2Gb Switch 16B	Compaq StorageWorks SAN Switch 2/16		16

Table 53: HP StorageWorks M-Series Product Line Switches

HP StorageWorks Switch Name		Firmware Version	Number of Ports
HP StorageWorks edge switch 2/12		05.05.00-12	4 to 12
HP StorageWorks edge switch 2/16		05.01.00-24	16
HP StorageWorks edge switch 2/24			8 to 24
HP StorageWorks edge switch 2/32			16 to 32
HP StorageWorks director 2/64			32 to 64
HP StorageWorks director 2/140			64 to 140
HP Switch Name	**Compaq Switch Name**		**Number of Ports**
N/A	McDATA ES-3016 (Compaq reseller)	05.01.00-24	16
N/A	McDATA ES-3032 (Compaq reseller)		32
McDATA ED-5000 (McDATA reseller)		04.00.00-16	32
HP Director FC-64	Compaq StorageWorks SAN Director 64	05.01.00-24	64

In addition to the switches listed in Table 52 and Table 53, CASA is also supported with the following Fibre Channel switch models (vendor branded):

Table 54: Brocade and McData Fibre Channel Switch Support for CASA-only SAN

Switch Brand	Switch Model	Firmware	Hub
Brocade	2400	2.6.1c	No
	2800	2.6.1c	No
	3200	3.1.1c	No
	3800	3.1.1c	No
	3900	4.1.2b	No
	12000	4.1.2b	No
McData	6064 (1 Gb directors)	04.01-02-4	No
	6140 (2 Gb directors)	05.01.00-24	No
	3216	05.01.00-24	No
	3232	05.01.00-24	No

Note: Table 54 lists switch vendor branded switch models supported by CASA only. For general non-CASA SAN configurations, refer to Chapter 3 for a list of supported HP-branded switch models.

Supported RAID Storage Arrays

- HP StorageWorks XP48, XP512, XP128, XP1024
- HP StorageWorks EVA v2, EMA12000, MA8000
- HP StorageWorks MSA1000
- HP StorageWorks va7400, va7410, va7100
- EMC Symmetrix 4 and 5
- EMC CLARiiON 4700 and 5700
- Hitachi 9200 and 9900
- Dell Powervault 650F

Supported Host Operating Systems

- Windows 2000, requires AutoPath version 2.0 for failover
- Windows NT 4.0, requires AutoPath version 1.05 for failover
- Solaris 2.6, 2.7, 2.8, requires VERITAS DMP for failover
- HP-UX K, L, R, and V class servers running HP-UX 10.20, 11, 11i, requires PVlinks for failover
- IBM AIX 4.3.3, requires AutoPath version 2.0 for failover
- Red Hat 7.1/Linux Kernel 2.4, requires Native Red Hat failover
- Novell NetWare 5.1, requires Native NetWare failover

Configuration Rules

CASA supports the full range of HP StorageWorks Fibre Channel SAN configurations as documented in this Guide. The following additional rules apply to all HP StorageWorks SAN installations that include CASA appliances.

Ask your HP representative for additional guidance on configuration rules.

Note: It is required that all host HBA ports are individually zoned to CASA target ports and that all CASA initiator ports are individually zoned to storage target ports.

Number of SAN Fabrics

For the purpose of availability, CASA installations normally use four separate Fibre Channel fabrics. Two fabrics are used to provide redundancy for the connection between the application servers and the appliances. Two additional fabrics are used for the connections between the appliances and the storage arrays. For installations where all the storage capacity is to be managed by CASA, this is the preferred configuration because it maximizes the availability of the entire system.

CASA may be used in installations where some of the storage capacity is managed by the CASA and some is directly connected[1] to application servers. In this case two fabrics are required. The failover functionality in the application servers, CASAs, and storage arrays makes this a no-single-point-of-failure configuration, however, there may be additional failover delay associated with the failure of one of the fabrics.

Number of CASAs

CASA is deployed with pairs of nodes in order to provide failover capability. A minimum CASA deployment has two nodes and is described as “one CASA.”

Multiple CASAs may be included in a SAN. Storage capacity is not shared between CASAs, except in those cases where replication is used. There is no specified limit to the number of CASAs that may be deployed in a single SAN, but in practice the connectivity limits of the SAN will restrict the number of CASAs.

Recommended SAN Topology

The recommended SAN topology for CASA deployments is core-edge (or director-edge) interconnection. Other topologies may not provide adequate port-to-port bandwidth.

CASA is supported in all HP StorageWorks SAN topologies.

Connection Rules

CASA requires a high-performance connection to the SAN for all of its ports. For this reason, the CASA should be connected directly to the core.

Application servers may be connected directly to the core or to edge switches, depending on the application workload.

Storage arrays may be connected to edge switches or directly to the core, depending on the workload requirements. In many cases the storage array will see a heavy workload and will need to be connected to the core.

Failover Software Rules

If all of the storage capacity in the SAN is under the management of CASA, the application servers must have failover software appropriate for the CASA. The physical storage devices are consolidated by CASA, so the failover software depends only on the CASA.

The following failover software must be installed on the application servers. Note that this failover management software is used when connecting to the CASA regardless of the storage arrays that are present in the configuration.

- AutoPath VA for Microsoft Windows and IBM AIX
- Veritas DMP for Sun Solaris
- PVLinks for HP-UX
- Native Linux
- Secure Path for Microsoft Windows

In order to handle the event of a path failure between the CASA and a storage array, the following failover software is used in the CASA:

- Native Active-Active (XP, VA, EMC)
- Secure Path (HSG80, HSV110,and MSA)
- ATF (for Clariion)

---

1. The connection may be a direct physical connection between the application and the storage array, if this is supported for the required server/array combination, or may be through an intermediate SAN. In this discussion “direct connection” includes both possibilities.

If some of the storage capacity is managed by the CASA and some is directly connected to application servers, then the application servers must have the appropriate failover software for the storage arrays to which they are connected. In these configurations the following issues should be considered.

- May require multiple flavors of failover software on the host (one for CASA storage, one or more for physical storage).
- Requires LUN mapping/masking on storage to allocate CASA LUNs and host accessible LUNs. Ensure that CASA LUNs can only be accessed by CASA by using an appropriate combination of LUN mapping, LUN masking, and zoning.

Example Configurations

Three example configurations are shown below. They cover the following cases:

- A Single CASA Manages all the Storage Arrays
- A Single CASA Manages a Subset of the Available Storage Arrays
- Multiple CASAs Manage the Storage Arrays

Similar configurations may be suitable for customer installations, depending on the specific requirements at hand.

Single CASA Manages all the Storage Arrays

Figure 60 shows a simple CASA configuration with four single-switch SAN fabrics. CASA Node C and CASA Node D are the redundant pair of nodes that make up “the CASA” in this illustration.

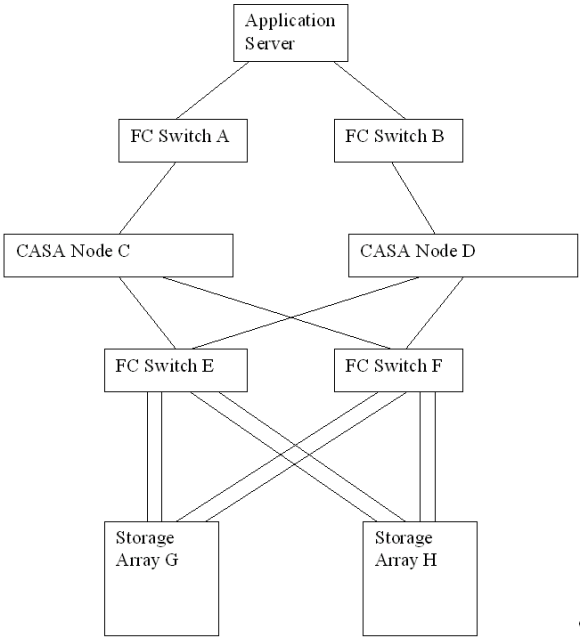

Figure 60: Single CASA Configuration

Single CASA Manages a Subset of the Available Storage Arrays

Figure 61 shows a more complex configuration where one CASA (pair of nodes) manages some of the storage capacity, while other storage capacity is connected directly to the application servers. Additionally, one of the storage devices, Storage Array H, is configured with LUN masking so that some of its capacity is connected to the servers and some to the CASA. Storage Array G is managed by the CASA while Storage Array I is connected directly to the application server.

In this case two fabrics are used because a connection between the application server and the storage arrays is needed. Each fabric has a core-edge topology, and the servers are connected to the edge switches as was discussed above. All of the CASA ports are connected directly to core switches.

Note that some of the connections to the storage arrays are not included on this drawing, for the purpose of simplification.

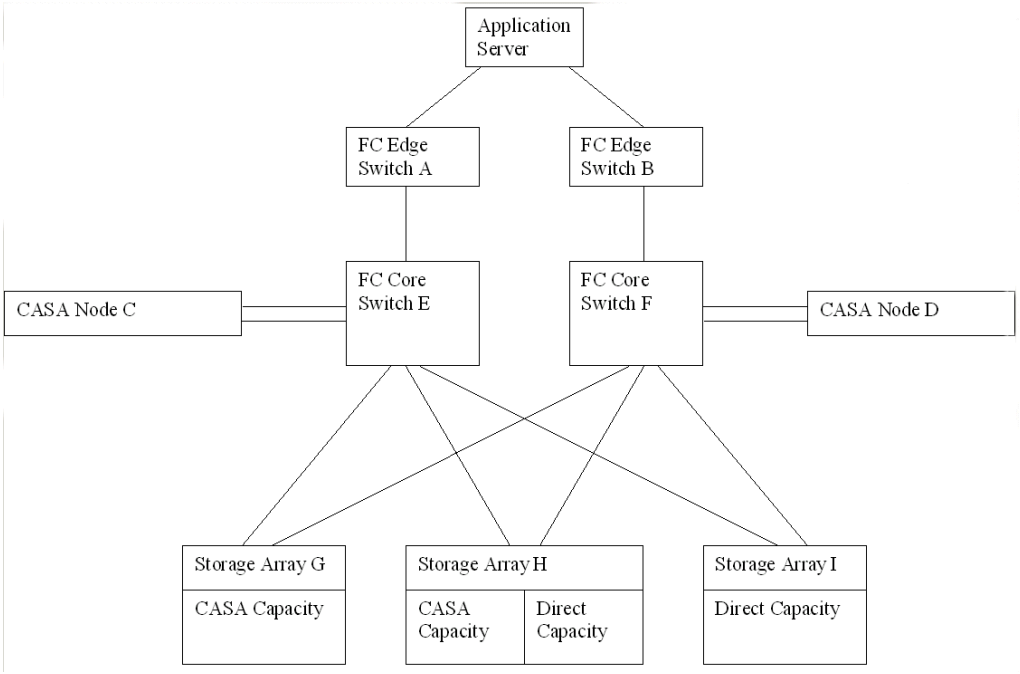

Figure 61: Single CASA Mixed with Non-CASA Storage

Multiple CASAs Manage the Storage Arrays

Figure 62 shows a larger configuration with multiple CASAs and multiple storage arrays.

Each CASA (node pair) controls a specific set of physical LUNs. The LUNs may be located on a single storage array, but in that case LUN masking (for example, Selective Storage Presentation on EVA arrays) must be used to isolate the LUNs.

CASA Nodes C and D are a pair, and have storage capacity assigned to them on storage arrays G and H. CASA Nodes J and K are a pair, and have storage capacity assigned to them on storage array G. Some of the capacity on storage array H, and all of the capacity on storage array I is assigned for direct access (through the SAN) by the application servers.

If one imagines an even larger configuration, and keeps in mind the requirement that all the CASA ports be connected directly to the fabric core, it is easy to see how the number of CASAs in a single installation is limited only by the number of ports on the fabric core.

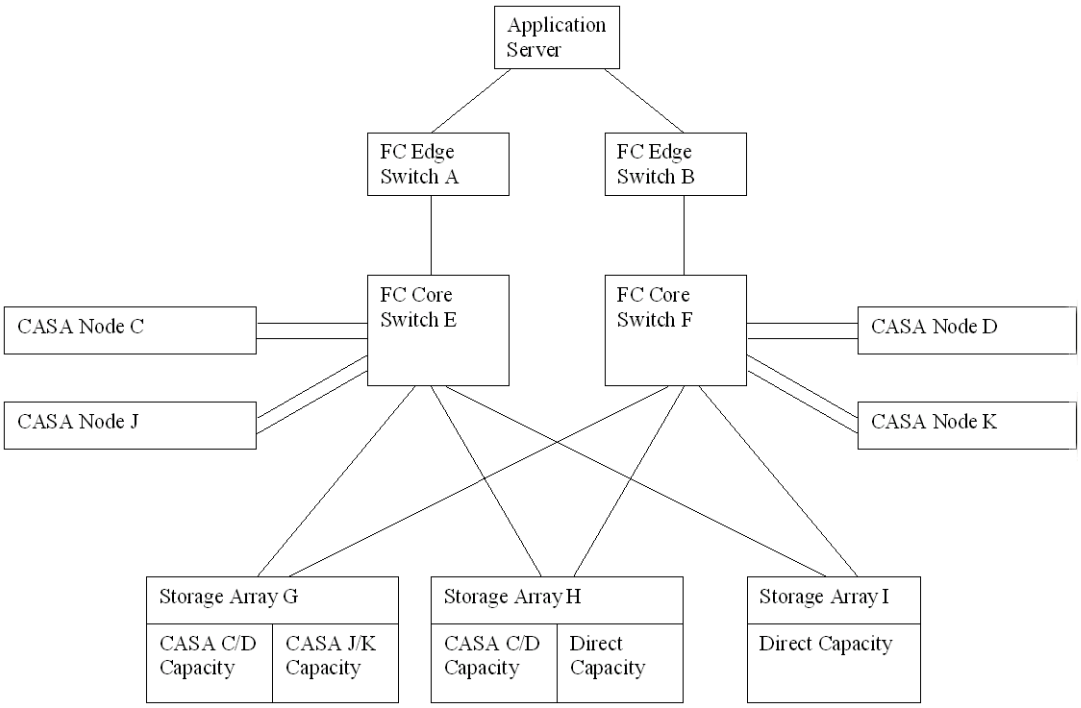

Figure 62: Multiples CASA Supporting Mix of Arrays

CASA Services

HP offers three levels of CASA Implementation Services. Simple (one operating system) and moderately complex (three operating systems, 12 hosts, and two Fibre Channel switches) installation services are available at fixed prices. Installation services for complex implementations are quoted after an assessment performed within the standard HP custom quoting process.

CASA Implementation Services include the following:

- Implementation planning and consulting for CASA.
- Hardware installation and configuration of CASA solution.
- Basic SANOS configuration including: password management, configuration of shared disk parameters, setting connections between nodes and shared disk, setting management and LAN network properties, and creating peer relationships between the nodes.
- Verify functionality of environment by performing: storage discovery, verification of host registration, creation of sample partitions and expansions, sample LUN mappings, and sample Fibre Channel Mirror and Fibre Channel Mirror configuration.
- Configure virtual LUNs and mappings per customer provided requirements and verify visibility of LUNs to the applicable nodes. The number of nodes verified is bounded by the complexity of the environment (simple, medium, or complex).
- Configuration of WAN IP addresses for each CASA and enablement of WAN interface.
- Configuration of IP Mirror relationship between source and target CASA units.
- Enablement of IP Mirror and creation of sample mirrored volume to verify functionality.
- Preparation of sample source and target LUN for Vsnap snapshot.
- Enablement of Vsnap snapshot, configuration of sample source and target LUNs to verify functionality.

The use of these services is highly recommended, particularly in complex environments.

Additional Information Sources

Refer to the CASA documentation at:

http://www.hp.com/products1/storage/products/virtualization_appliances/network/sv3000/index.html

for detailed information on supported RAID array storage devices.