TECHNOLOGY OVERVIEW

CISCO MDS 9000 FAMILY SANTap SERVICE—ENABLING INTELLIGENT FABRIC APPLICATIONS

Intelligent fabric applications are storage features that reduce the total cost of data management by taking into consideration the value of different types of data at different times, as well as the associated requirements for data accessibility, performance, availability, and protection. The intelligent features that make up this approach to storage networking specifically address customer challenges related to storage provisioning, data migration and replication, backup and recovery, storage utilization, and increasing storage costs.

The Cisco® MDS 9000 SANTap Service allows customers to deploy third-party appliance-based storage applications without compromising the integrity, availability, and performance of primary I/O. Cisco SANTap provides a reliable copy of storage write operations, enabling applications to provide data continuity, data protection, online data migration, storage performance, and SLA monitoring, without the traditional drawbacks of deploying devices in-band within the data path or out-of-band in conjunction with host-based software agents.

Cisco SANTap is enabled through the Cisco Storage Services Module (SSM) line card, which can be inserted into any Cisco MDS 9500 Series or Cisco MDS 9200 Series multilayer intelligent storage switch.

THE NEED FOR CISCO SANTAP

The Cisco MDS 9000 family SANTap Service allows the deployment of third-party application appliances without any of the traditional disadvantages associated with connecting appliances to SANs. Cisco SANTap:

- Does not compromise SAN performance, integrity, or availability.
- Nondisruptively enables fabric-based storage applications.
- Provides flexibility to choose the storage applications or appliances that satisfy business and operational needs.
- Protects existing investments in storage arrays via a software-based solution.

Figure 1 and Table 1 provide a comparison of in-band, out-of-band, and SANTap-based features.

Figure 1. Comparison of Out-of-Band, In-Band, and SANTap-Based Connectivity Options for Storage Application Appliances

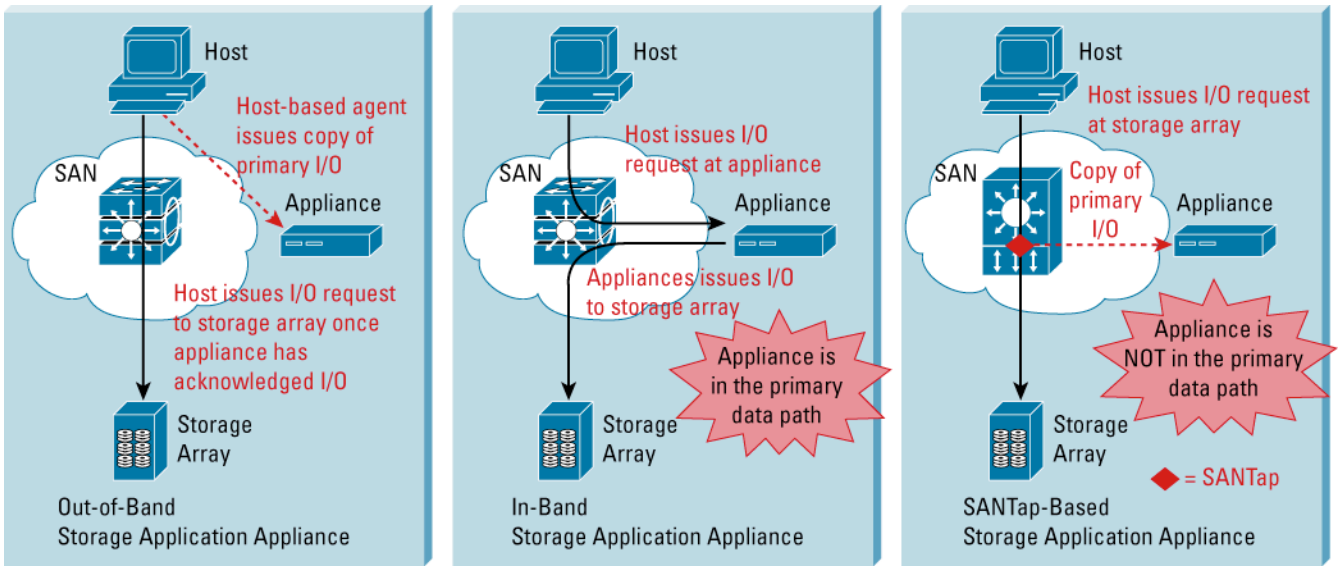

Table 1. Comparison of In-Band, Out-of-Band, and SANTap-Based Functions

Functional Comparison	Out-of-Band	In-Band	SANTap-Based
Requires software installed on host (consuming CPU, memory, and I/O)	Yes	No	No
Potentially compromises I/O performance	Yes*	Yes	No
Potentially compromises I/O integrity/availability by introducing a single point of failure	No	Yes	No
Can be deployed nondisruptively	Yes	No	Yes
Supports heterogeneous storage environments	Yes	Yes	Yes

* Host effectively issues double the write I/O

Using Cisco SANTap to deploy intelligent fabric applications has the following advantages:

- **Smooth insertion and provisioning of storage applications**
  - SANTap eliminates the service disruption caused by inserting appliances in-band.
  - SANTap reduces or eliminates host-side agents.
- **No disruption of the primary I/O from the server to the storage array**
  - SANTap eliminates the risk of an appliance or host-side agent affecting the availability and performance of deployed storage solutions.
- **Deployment flexibility and investment protection**
  - SANTap-enabled applications can be provided to all of the existing servers in the SAN, regardless of their operating systems. Deploying SANTap allows customers to get more out of their existing storage and server infrastructures.
  - Multiple best-of-breed storage applications can be concurrently added to servers and storage.

- **Enables on-demand storage services**
  - SANTap-enabled applications can be provisioned on demand without any application downtime for any server or storage appliance connected to any port of a storage network.
  - SANTap reduces implementation risk by enabling gradual introduction of storage applications in a customer environment. Customers can use a SANTap-enabled application in parallel with their existing applications in production.
- **Eliminates the scalability limitations for intelligent fabric applications by eliminating performance bottlenecks**
  - SANTap is enabled through custom intelligent Small Computer System Interface (SCSI) processing application-specific integrated circuits (ASICs) from Cisco Systems. A single SSM line card provides 320,000 I/O operations per second (IOPS) performance and 20 Gbps throughput. Moreover, SANTap has been implemented in a distributed architecture that enables multiple SSMs on a storage network to provide the SANTap services. Customers are no longer constrained by the performance limitations of host CPU cycles and in-band appliances.
  - SANTap can distribute workload to multiple application servers based on type of application and host/target combinations.

SANTap Protocol

The SANTap protocol operates between a Cisco MDS 9000 family multilayer intelligent switch and a SANTap-enabled storage application appliance. It contains commands for appliances to communicate housekeeping tasks to the switch (to commence SANTap service, for example) and to query the status of the SANTap service; it also contains data commands from the switch to provide a copy of data to the appliance. In addition, the SANTap protocol includes mechanisms for dealing with errors and outages with either the storage application appliance or the ports/connectivity between the appliance and the switch.

All SANTap communication is based on industry-standard SCSI commands running over Fibre Channel (SCSI-FCP). The SANTap service registers as both an initiator (host) and a target device (storage array) in the Fibre Channel name server. Communication between the SANTap service and the storage application appliance fits into three classes:

1. Control messages from the storage application appliance to the SANTap service.
2. Control messages from the SANTap service to the storage appliance.
3. Data traffic (reliable writes) mirrored from a host issuing a write to a storage array.

The first two classes of communication are messages or notifications between the devices to control various aspects of the SANTap service. Since the SANTap service appears as both a standard SCSI initiator and target, SCSI write operations are used between the SANTap service and the storage application appliance to convey control messages.

The third class of communication contains copies of any write I/O traffic between a host and a storage array. Copying of data traffic commences once the storage application appliance has registered itself with the SANTap service and has requested the service to start. SANTap guarantees that the mirrored write I/O traffic is an exact copy of write I/O operations issued by the host.

Communication between SANTap and the storage application appliance starts with the appliance registering for the SANTap service. The appliance discovers the SANTap service when it logs into the Fibre Channel fabric and queries the name server. Once discovered, the appliance will issue Port Login (PLOGI) and Process Login (PRLI) commands, followed by the standard SCSI device discovery process. The SANTap service will respond to a SCSI inquiry with vendor information set to “CISCO MDS” and product identification set to “MDS9000 SANTAP CVT”. The storage application appliance initializes the SANTap protocol via a control message requesting a copy of all write operations from a given initiator to a given target. The SANTap service is activated when the appliance issues a command to commence the copy. These steps are shown in Figure 2.

Figure 2. Steps for a Storage Application Appliance to Register for the SANTap Service

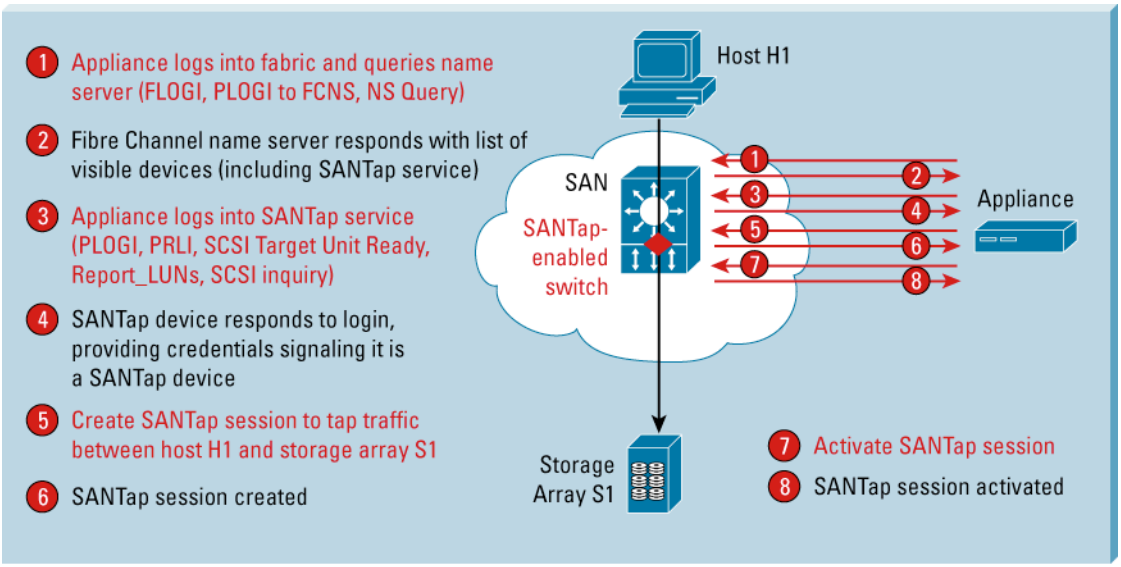

Once a Cisco SANTap service is operational, write operations are intercepted and delivered to the appliance-based storage application in parallel to the I/O being delivered to the target device (Figure 3). In addition to write operations (SCSI opcodes 0x0A, 0x2A, 0x8A, 0xAA), applications can request notification for other SCSI operations such as WRITE SAME, FORMAT UNIT, and RESERVE/RELEASE. If the appliance fails to acknowledge the I/O, the primary I/O path is unaffected and I/O between the initiator and target continues normally with SANTap operating in error recovery mode.

Figure 3. Cisco SANTap Service Reliable Write Operations

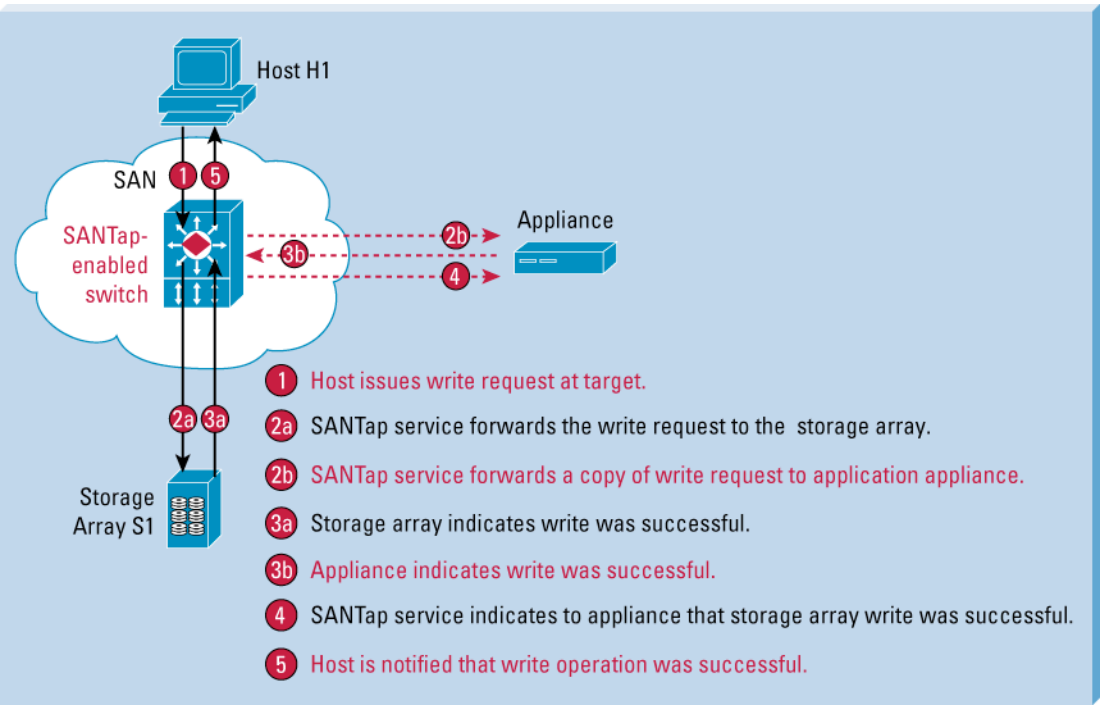

During session creation, the appliance-based storage applications can instruct the SANTap service to also record a log of write operations. The error recovery log enables rapid recovery in the event of either Fibre Channel port failure or appliance failure. If the appliance is offline or otherwise unable to acknowledge write operations, the SANTap service will maintain a record of I/O requests in the Appliance Recovery Log (ARL). Upon the appliance becoming available again, this log can be used to enable rapid resyncronization with the changed data.

In addition to providing appliance-based storage applications with a reliable copy of write I/O operations, appliances can issue control messages to the SANTap service to:

- **'Quiesce'** I/O from the initiator to the target. Quiescing temporarily pauses I/O from the host, queueing up requests until an 'unquiesce' operation is performed. Quiescing provides a synchronization point, allowing for operations such as point-in-time snapshots and switchover between log files or backup arrays.
- **'Redirect'** the primary I/O from the initiator to the appliance-based storage application only. This means that I/O (reads and writes) is not sent to the target at all, but is instead redirected to the storage appliance.

The variety of run-time protocol configuration options provides flexibility in the types of applications that can take advantage of Cisco SANTap. The Cisco SANTap protocol provides the necessary network-assisted intelligence for data continuity and data protection applications (point-in-time snapshots, synchronous replication, and asynchronous replication, for example), as well as the capability to enable a whole new breed of online data migration and storage performance monitoring applications.

SANTap Performance

The Cisco SANTap service is one of the many intelligent network services enabled through the SSM line card, which can be inserted into any modular switch within the Cisco MDS 9000 family. Each SSM contains multiple embedded processors, which provides a distributed architecture capable of providing inline SCSI support for up to 320,000 IOPS and in excess of 20 Gbps of throughput per module. Multiple SSMs may be deployed in a chassis for higher aggregate performance, and multiple SSMs can be distributed across multiple chassis.

Each SSM contains 32 Fibre Channel front-panel ports. In the case of a transparent mode SANTap service, either the initiator or target must be connected to one of these ports. In the case of proxy modes, SANTap service traffic will flow through the SSM, consuming some of the 20 Gbps throughput, regardless of what other devices may be connected to the ports. Compared to traditional appliance-based storage application deployments, a SANTap-enabled application can offer higher levels of performance—it resides outside the primary data path and only receives copies of write operations.

SANTap Error Recovery Services

The Cisco SANTap service provides several error recovery services to permit rapid recovery in the event of appliance or port failure. As part of establishing a SANTap service, the appliance-based storage application may elect for the SANTap service to provide one of three possible types of recovery logs. The primary purpose of these logs is to record the write operations from the initiator to a target while SANTap communication to the appliance is unavailable.

The types of error recovery logs are:

- **Appliance recovery log (ARL)—**Used for fast recovery in the event of appliance failure. When an appliance comes back online, the appliance could query the ARL to identify which blocks have changed since the appliance failure.
- **Pending write log (PWL)—**Used for fast recovery in the event of SANTap port failure. The PWL enables the appliance to isolate the outstanding I/O's, at the time of the SANTap failure, and synchronize with what was actually written or not written to the storage array.
- **Circular log**—Used for fast recovery by a standby appliance in the event of a primary appliance failure. This log ensures that the standby appliance is synchronized with all of the I/O's that were committed by the active appliance.

The ARL and the PWL keep a record of what write operations have been performed. This record can be stored in the form of a bitmap (dividing the total storage size into multiple regions, then setting a bit whenever there is a write that modifies a region) or in the form of a list of logical blocks that have been modified. These logs can be used for an appliance to rapidly "sync up" with storage changes while the appliance was unavailable.

The ARL and circular log can be queried and cleared by the appliance at any time through specific control messages. Querying and clearing the log in conjunction with the quiesce command allows SANTap-enabled appliances to gather a definitive list of storage blocks that were modified while the appliance was unavailable.

In addition to providing error recovery logs, the Cisco SANTap service provides two target devices that can be used for error recovery. The Cisco SANTap Appliance Virtual Target (AVT) enhances usability of the solution by masking the appliance's identity to be that of the host (providing a means of avoiding configuration changes on the storage array to grant the appliance access). This enables the appliance to access the storage target directly, enabling the appliance to synchronize with the storage. For application appliances that expose a point-in-time copy of data, a virtual LUN needs to be exposed to the host, in the same VSAN as that of the host. The Cisco SANTap Recovery Virtual Target (RVT) can be used to provide connectivity from the host for performing I/O against the virtual LUN exposed by the appliance.

SANTap Security

Cisco SANTap services are compatible with the management and fabric/target access security mechanisms that Cisco MDS 9000 family switches offer. The SANTap service may be deployed in conjunction with the following security features:

Fibre Channel Zoning

Zoning is the security mechanism within Fibre Channel used to restrict communication between devices within the same Fibre Channel fabric. Since all SANTap service communication between the switch and the appliance is based on standard SCSI/SCSI-FCP, both the SANTap service and the storage application appliance must be configured in a common Fibre Channel zone to provide connectivity. Using zoning, it is possible to limit what appliances are capable of establishing SANTap services.

LUN Zoning and Read-Only Zones

Cisco MDS 9000 family switches can provide more detailed zoning than is generally available today. Based on deep frame inspection, hard zoning within Cisco MDS 9000 family switches can restrict access to explicit logical unit numbers (LUNs) within a storage array and can even restrict write SCSI I/O operations, enforcing read-only access. LUN zoning and read-only zones can be used to help ensure that application appliances do not write to the primary volumes of the storage arrays.

Virtual SANs

Virtual SANs (VSANs) can be used to create multiple logical SANs over a common physical infrastructure. Each VSAN runs its own set of fabric services, providing for absolute partitioning between virtual fabrics. VSANs can be used to achieve higher security and greater stability in Fibre Channel fabrics by providing isolation among devices that are physically connected to the same set of switches. With Cisco SANTap, an appliance can be connected in its own VSAN, separate from the initiators and targets that it is receiving traffic from.

Port Security

Port security can be used to limit access to the Fibre Channel fabric based on the device identity attributes. Port security prevents unauthorized access to a switch port by binding specific worldwide names (WWNs) as having access to one or more given switch. When port security is enabled, all devices connected to a switch must be in the port security database and must be listed in the database as bound to a given port. Port security can be used to lock specific authorized appliances to specific switch ports.

Roles-Based Access Control

Roles-Based Access Control (RBAC) enables different users to have different roles, responsibilities, management capabilities, and restrictions. RBAC can be used to separate administrators for SANTap services from other storage administrators. User role configuration may either be configured locally on the switch or stored centrally and distributed to the switch as part of authentication and authorization via either RADIUS or TACACS+.

FC-SP DH-CHAP

Fibre Channel Security Protocol (FC-SP) Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) can be used to help ensure data integrity and authentication for device communication. Authentication is based on DH-CHAP, and can be performed locally in the switch or remotely through a centralized RADIUS or TACACS+ server.

FC-SP DH-CHAP provides absolute protection against WWN spoofing on a compromised port, even when physical security of the switch has been compromised and a rogue device has been installed on the same physical switch port. FC-SP DH-CHAP can be used to help ensure that only trusted appliances are entitled to enable a SANTap service.

SANTap Deployments

A Cisco SANTap service can operate in many modes, allowing a great deal of flexibility in deployment options (Figure 4).

- **Transparent mode** enables the Cisco SANTap service to be deployed nonintrusively into an existing SAN environment between an initiator and a target with no outage or loss of SAN service. As its name implies, transparent mode can be deployed without any changes to the Fibre Channel addresses (FC_IDs) of either the initiator or the target. Transparent mode can only be used when either the initiator or target is directly attached to a Fibre Channel port on the SSM. The initiator and target must be in the same VSAN; the appliance may be in the same VSAN or in a different VSAN.

- **Proxy mode 1** is required when neither the initiator device nor target device is directly connected to a Fibre Channel port on the SSM. In proxy mode 1, the SANTap service will use FC_ID translation to enable SANTap to perform its function. In proxy mode 1, initiator devices do not communicate directly with target devices—instead, they communicate with a virtual target within the SANTap service, which proxies communication to the real target. With proxy mode 1, the only requirement is for the initiator and target to be within the same fabric (VSAN), but either or both may be connected to any switch within the fabric. The primary disadvantage with deploying Cisco SANTap with proxy mode 1 is that since all communication is translated between the initiator and target, "storage bindings" on the host may need to be configured. Likewise, any LUN mapping, masking, or security on the storage array may need to be reconfigured to the FC_ID presented by the SANTap service.
- **Proxy mode 2** aims to simplify the deployment and alleviate the primary disadvantage associated with proxy mode 1. In proxy mode 2, the initiator and target are configured in different fabrics (VSANs), with the Cisco SANTap service providing proxied connectivity between the two devices such that they can communicate as if they were in the same fabric. Since the two devices are in separate VSANs, there is no need for any FC_ID translation and the SANTap service can present the target to the initiator (in the initiator's VSAN) with the same FC_ID the target has in its own VSAN. Proxy mode 2 allows maximum flexibility—neither the initiator nor the target needs to be directly attached to a Fibre Channel port on the SSM, or even attached to the same switch.

Regardless of the mode used for the SANTap service, there are no restrictions on connectivity for the storage application appliance itself; it may be connected via any switch in the fabric and may reside in any VSAN.

Figure 4. Cisco SANTap Modes

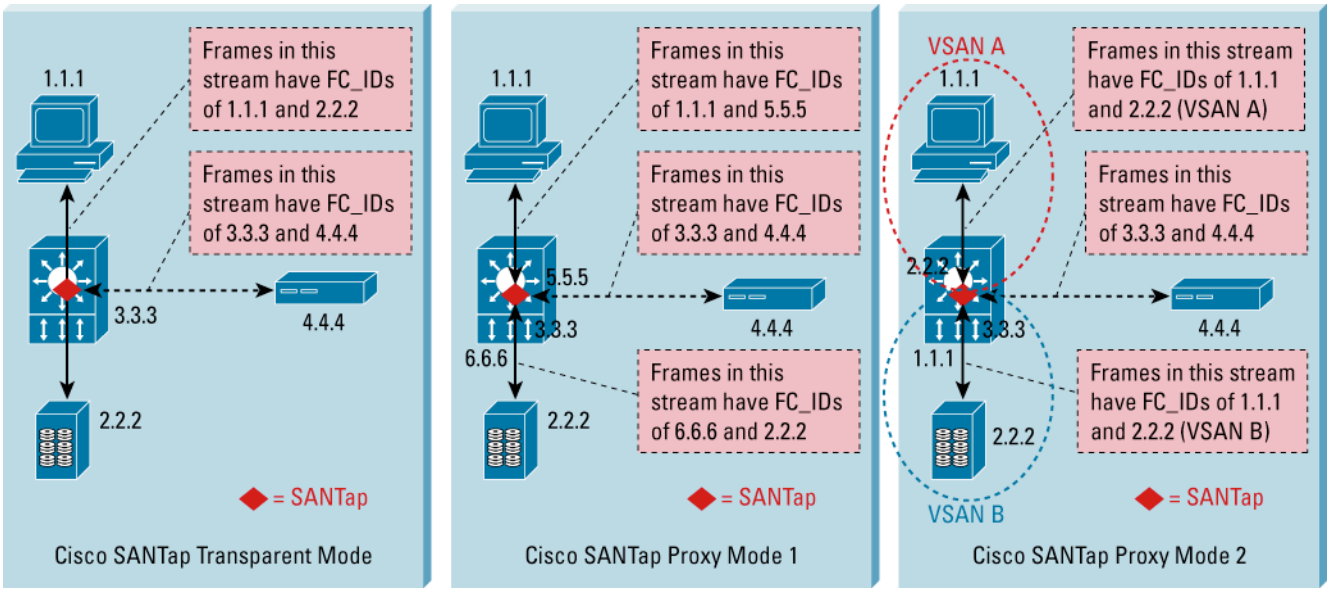

Both the Cisco SANTap service and SANTap-enabled appliances can be deployed seamlessly into an existing SAN since both can be provisioned and enabled without any impact to existing traffic flows. Furthermore, SANTap-enabled appliances can provision and enable the Cisco SANTap service through the SANTap control protocol, substantially simplifies SANTap management since all service provisioning can be performed through a single appliance console.

The Cisco SANTap service can be deployed as a high-availability solution for mission-critical deployments. It may be used in conjunction with multipathing or redundant fabric deployments, even where initiators are performing active/active multipathing. This type of deployment is shown in Figure 5.

Figure 5. Cisco SANTap High-Availability Deployment

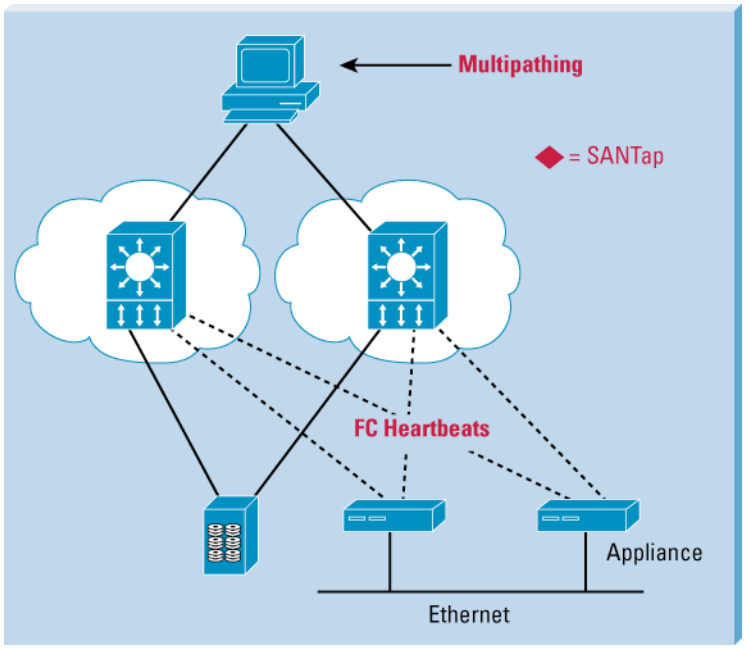

SUMMARY

The Cisco MDS 9000 family SANTap Service allows the deployment of third-party application appliances without any of the traditional disadvantages associated with connecting the appliance to the SAN. Cisco SANTap:

- Does not compromise SAN performance, integrity, or availability.
- Nondisruptively enables fabric-based storage applications.
- Eliminates the service disruption caused by inserting appliances in-band.
- Reduces or eliminates host-side agents.
- Provides flexibility to choose the storage applications or appliances that satisfy business and operational needs.
- Protects existing investments in storage arrays via a software-based solution.
- Allows storage services to be added to any server/storage device in the network without any rewiring.
- Enables on-demand storage services.
- Allows multiple appliance-based storage services to be concurrently added to servers and storage.
- Reduces implementation risk by enabling gradual introduction of services for staging.
- Enables commodity-appliance-based storage applications to scale.
- Allows distribution of workload to multiple appliances based on application and source/target combinations.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204181.k_ETMG_DB_2.05

Printed in the USA