

EMC RecoverPoint Support for Cisco MDS 9000 SANTap Service: Intelligent Fabric-Based Data Replication

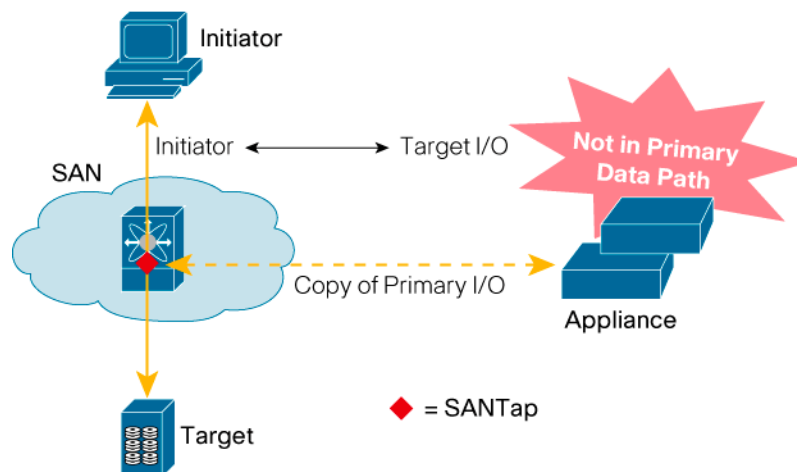
Intelligent storage networks promise a new heterogeneous, flexible, and high-performance platform for hosting storage applications. EMC has delivered on this promise by using the Cisco[®] MDS 9000 Family SANTap Service, offering a fully heterogeneous fabric-based data replication solution without the need of a host agent. The EMC RecoverPoint data protection appliance replicates data in any storage area network (SAN) attached array, local or remote across any distance, to any other array, without any of the traditional downsides of deploying devices in-band (within the data path) or out-of-band, in conjunction with host-based software agents. RecoverPoint reduces the total cost of data protection and management by taking into consideration the value of different types of data and the associated requirements for data accessibility, performance, availability, and protection.

RecoverPoint's intelligent features specifically address key customer concerns related to data protection and replication, recovery, storage use, and high cost of bandwidth and storage management. Using the Cisco MDS 9000 Family SANTap Service, RecoverPoint can be deployed out-of-band (outside the data path) and without a host agent, because Cisco SANTap provides a reliable copy of storage write operations. RecoverPoint is a modular architecture that offers solutions within the local data center environment for Continuous Data Protection (RecoverPoint CDP) and for Continuous Remote Replication (RecoverPoint CRR). Cisco SANTap functions are activated through the Cisco MDS 9000 Family Storage Services Module (SSM) line card, which can be inserted into any modular switch within the Cisco MDS 9000 family of multilayer intelligent storage switches.

Cisco SANTap—Optimized Architecture for Data Replication

Storage applications can be delivered on the host or initiator, on the storage or target, or inside a SAN-attached appliance. A purpose-built appliance that is designed to provide a specific storage application such as replication offers the benefit of support for heterogeneous storage while minimizing the effect on hosts, applications, and array performance. Long-distance data replication is uniquely suited for this architecture, because hosts and storage are relieved from the burden of managing the data movements across the WAN and dealing with complex data-consistency concerns inherent in cross-site data replication. An appliance placed at the SAN and WAN junction will have to be connected either in-band (between host and storage) or out-of-band, with the aid of a host agent. Both approaches can have obvious drawbacks. The in-band approach can compromise SAN performance, integrity, and availability, and deployment can be very disruptive. The out-of-band deployment removes these drawbacks, but requires installation of a driver in the host, potentially adding complexity to the solution's deployment. Cisco SANTap eliminates the need for host agents, allowing a simplified and agent-free implementation of RecoverPoint. This out-of-band appliance approach offers an optimized architecture for data replication that protects existing investments in storage arrays with no host footprint and delivers high-performance bidirectional data replication between SANs across any distance, without any effect on host and array performance. With ample processing, and using native SAN and WAN interfaces, the RecoverPoint appliance efficiently replicates data across multiple (homogeneous or heterogeneous) arrays without the need for protocol converters or edge connection devices (Figure 1).

Figure 1. Cisco SANTap Connectivity for EMC RecoverPoint



Cisco SANTap Benefits for Data Protection

- Transparent deployment of the RecoverPoint data protection appliance—RecoverPoint can be deployed transparently without any disruption. The Cisco SANTap architecture eliminates the need for host-side agents, further simplifying the deployment.
- No disruption of the primary I/O from the server to the storage array—RecoverPoint design takes advantage of Cisco SANTap to eliminate the risk of affecting the availability and performance of deployed applications.

- Deployment flexibility and investment protection—RecoverPoint data protection services can be provided to all existing servers and storage in the SAN, irrespective of the operating systems or array type, allowing customers to get more out of existing storage and server infrastructure investments.
- On-demand storage services—RecoverPoint data protection services can be provisioned on demand without any application downtime for any server or storage connected to any port of the storage network.
- Unlimited scalability and no performance bottlenecks — A single Cisco SSM line card provides 500,000 I/O operations per second (IOPS) performance and 20-Gbps throughput. Moreover, Cisco SANTap is implemented in a distributed architecture that allows multiple Cisco SSMs on a storage network to provide the Cisco SANTap services. These features, coupled with the linear scalability of data replication appliances, mean that customers are no longer constrained by the performance limitations of host CPU cycles and in-band appliances.

RecoverPoint Fabric-Based Protection Benefits

- Universal enterprise data protection—EMC RecoverPoint with the Cisco SSM SANTap solution is an end-to-end solution for Continuous Data Protection (CDP) and Continuous Remote Replication (CRR) across heterogeneous server and storage platforms, providing complete data protection for the entire enterprise. Because storage systems at the primary and secondary sites do not have to be the same, there is flexibility to deploy lower-cost storage (for example, at the secondary site) or to use existing storage. With the Cisco SSM SANTap architecture, RecoverPoint can be used with virtually every open-system-based environment. No changes are required on the storage arrays, hosts, or SAN configuration.
- Guaranteed data consistency—RecoverPoint guarantees a consistent replica of business-critical data in the event of any failure or disaster. With RecoverPoint, consistency is maintained at all times, even through rolling disasters or during resynchronization.
- Continuous Data Protection—RecoverPoint CDP utilizes CDP technology that provides application recovery to any point in time. With RecoverPoint CDP, each I/O boundary is captured as a time-stamped transaction that is recorded within the recovery journal.
- Intelligent application bookmarks—RecoverPoint uses intelligent application bookmarks that enable administrators to recovery to application specific points in time. These bookmarks enable an administrator to ensure consistent data recovery, minimizing potential data loss due to logical corruption.
- Intelligent use of bandwidth—RecoverPoint CRR employs intelligent bandwidth reduction technologies to deliver outstanding bandwidth savings. As a result, the system provides superior protection for the available bandwidth, while dramatically reducing WAN costs, particularly over long distances. Data reduction is achieved through application-aware and storage-aware algorithmic techniques that conserve bandwidth to an extent not possible with traditional compression technologies.

- Policy-based data replication—RecoverPoint CRR offers a full spectrum of replication modes, from synchronous to asynchronous, small-aperture snapshot, and point-in-time. The replication process is managed automatically, with strict adherence to user-defined policies that are tied to desired business objectives. The system adapts its replication dynamically for each application according to these policies, based on the available bandwidth and the application workload, greatly simplifying data and disaster recovery management for complex and heterogeneous environments. For example, a replication policy to minimize lag may be specified for a crucial application, in which case the system uses all available bandwidth to minimize the lag between the primary and secondary sites. Alternatively, a policy to minimize bandwidth may be specified for less-crucial data, causing the system to use as little bandwidth as possible, while maintaining the lag within a specified upper limit.
- Recovery to any point in time—RecoverPoint CRR efficiently maintains a snapshot history to allow convenient rollback to any point in time, providing quick and effective recovery from a disaster. It supports multiple transactional-consistent snapshots at the remote site, allowing reliable recovery in database environments. Frequent, small-aperture snapshots (seconds apart) are used to minimize the risk of data loss due to data corruption.
- Data processing on the target volume—Both RecoverPoint modules (CDP and CRR) support target Side Processing (TSP) which allows direct read/write access to the replicated image in real-time. There is no requirement to first make an additional copy or clone of the volume. The system supports robust failover and failback capabilities, reducing management and operating costs.
- Intuitive management—RecoverPoint offers multiple management features that make implementation and operation extremely easy. The system, consisting of any number of appliances, is centrally managed through a single intuitive GUI or through a command-line interface (CLI), both of which are accessible from a secure Internet connection. RecoverPoint provides automated discovery of logical unit numbers (LUNs), provides the use of Simple Network Management Protocol (SNMP) for integration with standard enterprise network and system management applications, and offers call-home capability that proactively reports the system status in the event of a failure.

RecoverPoint Data Protection Using Cisco SANTap Protocols

The RecoverPoint solution and the Cisco MDS 9000 family of multilayer intelligent switches use the Cisco SANTap protocol to allow the copying of host write operations to the storage array to a RecoverPoint appliance at the source site. For remote replication, these write processes are then efficiently transferred using TCP/IP over a WAN connection to a RecoverPoint appliance at the target side that manages write operations to the target replication volumes. RecoverPoint uses the Cisco SANTap protocol to communicate housekeeping tasks to the switch (for example, to commence Cisco SANTap service) and query the status of the Cisco SANTap service, as well as to send data commands from the switch to provide a copy of data to RecoverPoint. In addition, the Cisco SANTap protocol includes mechanisms for handling errors, failures, and outages in the storage array or the ports and connectivity between the appliance and the switch.

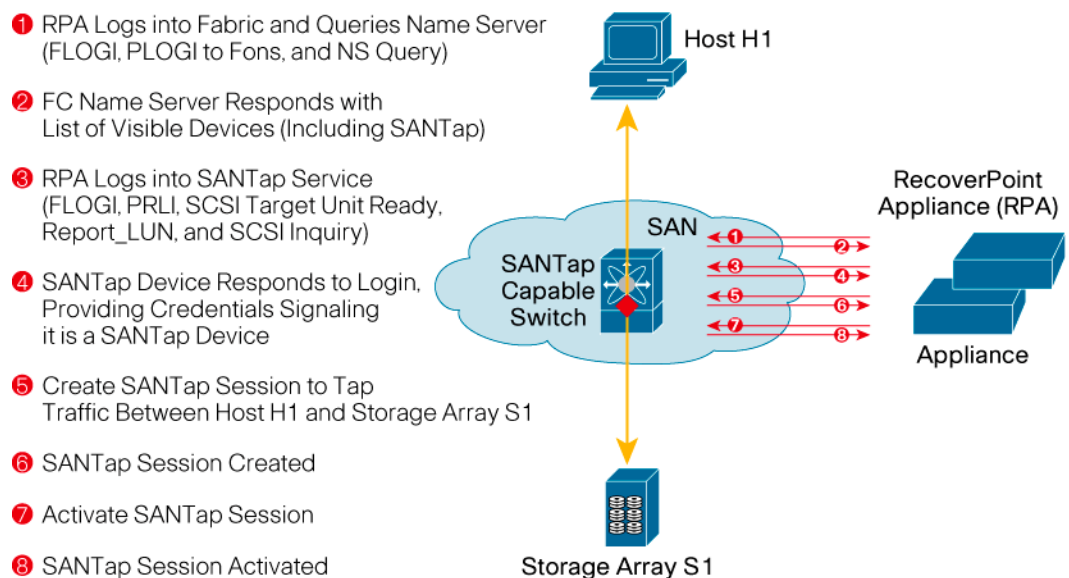
All Cisco SANTap protocol communications are based on industry-standard Small Computer System Interface (SCSI) commands running over Fibre Channel (SCSI-FCP). The Cisco SANTap service registers as both an initiator and a target device in the Fibre Channel name server. Communications between the Cisco SANTap service and RecoverPoint fit into three classes:

- Control messages from RecoverPoint to the Cisco SANTap service
- Control messages from the Cisco SANTap service to RecoverPoint
- Data traffic (reliable write operations) mirrored from a host sending a write process to a storage array

The first two classes of communications are messages or notifications between the devices to control various aspects of the Cisco SANTap service. Since the Cisco SANTap service and RecoverPoint each appear as a standard SCSI initiator and a target, SCSI write operations are used between the Cisco SANTap service and RecoverPoint to convey control messages.

The third class of communications contains copies of any write I/O traffic between a host and a storage array. Copying of data traffic commences after RecoverPoint registers itself with the Cisco SANTap service and requests the service to start. Cisco SANTap guarantees that the mirrored write I/O traffic is an exact copy of write I/O operations sent by the host. Communication between Cisco SANTap and RecoverPoint starts with the appliance registering for the Cisco SANTap service. The appliance discovers the Cisco SANTap service when it logs into the Fibre Channel fabric and queries the name server. After discovery, RecoverPoint will send Port Login (PLOGI) and Process Login (PRLI) commands, followed by the standard SCSI device-discovery process. The Cisco SANTap service will respond to a SCSI inquiry with vendor information set to "Cisco MDS" and the product identification set to "Cisco MDS 9000 SANTap CVT." RecoverPoint initializes the Cisco SANTap protocol through a control message requesting a copy of all write operations from a given initiator (host) to a given target (storage array). The Cisco SANTap service is activated when RecoverPoint sends a command to commence the copy operation. These steps are shown in Figure 2.

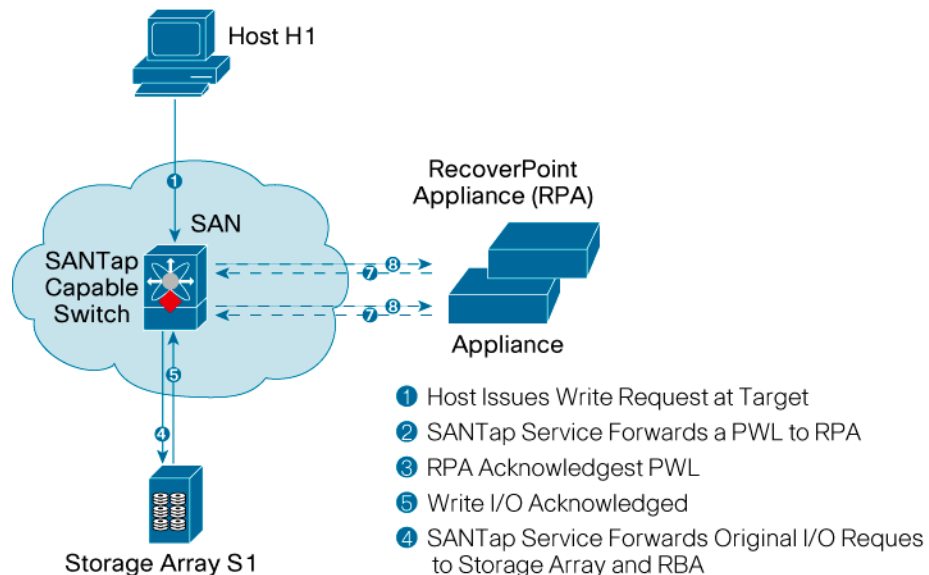
Figure 2. Steps for a Storage Application Appliance to Register for Cisco SANTap Service



After a Cisco SANTap service is operational, write operations are intercepted and delivered in a parallel manner to RecoverPoint and the storage array. Initially, Cisco SANTap will send a pending write log (PWL) to RecoverPoint. The PWL is a short SCSI command (several bytes) consisting only of the write operation's metadata. After the PWL is acknowledged by RecoverPoint, the Cisco SANTap service will simultaneously perform a write I/O operation to both RecoverPoint and the target device (storage array). RecoverPoint will then acknowledge the write I/O (see

Figure 3). Because the actual data is not written in the PWL, the latency introduced by this stage is kept to a minimum, yet this process helps ensure that no storage writes go undetected by RecoverPoint. Note that if RecoverPoint fails to acknowledge the I/O, the primary I/O path is unaffected, and I/O between the initiator and target continue as normal, with Cisco SANTap operating in error-recovery mode.

Figure 3. Cisco SANTap Service and RecoverPoint Parallel Write Options



Cisco SANTap Performance

The Cisco SANTap service is one of the many intelligent network services provided through the Cisco SSM line card that can be inserted into any modular switch within the Cisco MDS 9000 family of multilayer intelligent storage switches. Each Cisco SSM contains multiple embedded processors, providing a distributed architecture capable of providing inline SCSI support for up to 500,000 IOPS and in excess of 20 Gbps (full duplex) of throughput per module. Multiple Cisco SSMs can be deployed in a chassis for higher aggregate performance, and multiple Cisco SSMs can be distributed across multiple chassis. Each Cisco SSM contains 32 Fibre Channel front-panel ports. RecoverPoint uses a Cisco SANTap deployment mode called proxy mode 2 that allows any Fibre channel port to be used in replication (see “RecoverPoint and Cisco SANTap Deployment” later in this document).

RecoverPoint and Cisco SANTap Error Recovery Services

The RecoverPoint and Cisco SANTap solution provides a number of error recovery capabilities to allow rapid recovery in the event of various failure scenarios. These capabilities are supported through RecoverPoint’s internal recovery mechanisms in conjunction with error recovery logs provided by Cisco SANTap. RecoverPoint uses the following types of error recovery logs:

- Appliance Recovery Log (ARL)—Used for detection of host write traffic that does not reach RecoverPoint
- Pending Write Log (PWL)—Used for identification of write traffic that reached RecoverPoint but did not reach the storage array

- Circular Log (CL)—Used for fast recovery of host write traffic not marked by RecoverPoint because of a failure prior to marking

Both the ARL and PWL keep a record of the write operations that have been performed. This record can be stored either as a bitmap (dividing the total storage size into multiple regions and then setting a bit any time a write operation modifies a region) or as a list of logical blocks that have been modified. The CL keeps a complete log of write I/O operations performed by the host machines. In some failure conditions, host write operations will continue to be written to storage without being copied to RecoverPoint. After the failure ends, RecoverPoint will query the ARL to detect whether host write operations were performed during the failure condition. If needed, RecoverPoint can then rapidly resynchronize replication volumes between the source and target site. RecoverPoint uses markers to persistently log on a SAN volume the blocks that were modified in the storage array. In the case of a WAN failure or other failures at the local site, these markers are used to determine which blocks need to be resynchronized. In case of a failure in a RecoverPoint appliance before marking was persistently stored, RecoverPoint can recover this information from the Cisco SANTap service CL. The PWL is used by RecoverPoint to detect write operations that were not acknowledged by the storage array but did reach RecoverPoint. In this case, RecoverPoint resynchronizes the affected volumes.

RecoverPoint and Cisco SANTap Security

RecoverPoint and the Cisco SANTap services are compatible with all the management and fabric and target access security mechanisms offered on all members of the Cisco MDS 9000 family of multilayer intelligent switches. Using these security services, it is possible to deploy the RecoverPoint replication solution while maintaining high security and high service in the SAN fabric. The RecoverPoint and Cisco SANTap services can 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. Because all Cisco SANTap service communication between the switch and the appliance is based on standard SCSI and SCSI-FCP, both the Cisco SANTap service and RecoverPoint must be configured in a common Fibre Channel zone to provide connectivity. Because RecoverPoint requires access to the replicated storage (such access is used, for example, during replication image synchronization), an additional zone should be defined with RecoverPoint and the replicated storage arrays. Zoning therefore allows control of the RecoverPoint appliances that have access to Cisco SANTap services, limiting RecoverPoint access to only replicated storage arrays.
- Virtual SANs (VSANs)—VSANs can be used to create multiple logical SANs over a common physical infrastructure. Each VSAN runs its own set of fabric services, providing 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. In the case of the Cisco SANTap service and RecoverPoint, a VSAN containing replicated storage arrays and RecoverPoint appliances is created.
- 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 World Wide Names (WWNs) to access to one or more specific switches. When port security is active, 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 RecoverPoint appliances to specific switch ports.

- Role-based access control (RBAC)—RBAC provides the capability for different users to have different roles and responsibilities, management capabilities, and restrictions. RBAC can be used to separate administrators for Cisco SANTap services from other storage administrators. User role configuration can be either configured locally on the switch or stored centrally and distributed to the switch as part of authentication and authorization through RADIUS or TACACS+.
- Fibre Channel Security Protocol (FC-SP) and Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP)—FC-SP and DHCHAP can be used to help ensure data integrity (tamper-proof data) and authentication (nonrepudiation) for device communication. Authentication is based on CHAP with DH extensions. Authentication can be performed locally in the switch or remotely through a centralized RADIUS or TACACS+ server. FC-SP and DHCHAP 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 and DHCHAP can be used to help ensure that only trusted RecoverPoint appliances are allowed to establish a SANTap service.
- LUN zoning and read-only zones—Cisco MDS 9000 Switches can provide far more fine-grained zoning than is generally available today. Based on deep-frame inspection, hard zoning within the Cisco MDS 9000 family of switches can restrict access to explicit 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 source-site RecoverPoint does not perform write operations on the source site volumes of the storage arrays.

RecoverPoint and Cisco SANTap Deployment

The RecoverPoint solution operates Cisco SANTap services in proxy mode 2. This mode aims to simplify the deployment and support replication services on all ports available in the fabric: both Cisco SSM Fibre Channel ports and ordinary Fibre Channel ports. Furthermore, replication services can be used on Fibre Channel ports that are not attached to the same switch. Using proxy mode 2, replication services are transparently added to the fabric while maintaining all WWNs originally used, greatly simplifying deployment and configuration (Figure 4).

To allow this advanced feature, the initiator (host) and target (storage array) are configured to different fabrics (VSANs), with the Cisco SANTap service providing proxy connectivity between the two so that the two devices can communicate as if they were in the same fabric.

The RecoverPoint appliances are configured to the target's VSAN and are not accessible to initiators. Because of the proxy connection, Fibre Channel ID (FC ID) translation is not needed. The initiator (in the initiator's VSAN) uses the same FC ID that the actual target has in its own VSAN. There are no restrictions on connectivity for the RecoverPoint appliances; they can be connected through any switch in the fabric.

Figure 4. Cisco SANTap Modes

