incipient

Deploying iNSP Software in Cisco® MDS 9000 Based Storage Area Networks

White Paper

April 2006

Contents

Introduction

This paper will help you understand the benefits companies will achieve by implementing the Incipient Network Storage Platform (iNSP) software suite in a Cisco MDS 9000 based SAN environment.

> *iNSP software is the industry's first and only switch-resident enterprise-class storage virtualization solution on the market today*

iNSP leverages Cisco's industry leading MDS 9000 SAN switches and includes storage-related services such as non-disruptive data migration, automated storage provisioning and point-in-time copy services. iNSP runs natively on Cisco MDS 9000 Storage Services Modules (SSMs) that are designed to work with intelligent fabric applications. iNSP does not require any external appliances or servers and works with any Cisco MDS 9500 Director or 9200 fabric switch. Cisco SSMs are designed specifically to support intelligent fabric applications such as iNSP software. This provides investment protection for Cisco MDS customers as well as rapid access to storage services based applications.

With SAN storage rapidly growing in capacity and complexity each year, the need to automate SAN management tasks, increase utilization, and reduce storage and storage management costs will become a requirement in most enterprise environments.

Understanding SAN Challenges and iNSP's Impact on SAN Management

Widespread adoption of SAN technology has proven both a blessing and a burden to data center administrators everywhere. One of the benefits of implementing a SAN is to consolidate distributed storage resources and centralize administrative tasks. But as the SAN fabric grew in port counts, storage, and number of applications, so did the management complexity.

For example, a single storage array once bound tightly to a small number of servers may now be utilized by hosts across the entire SAN. Managing the array itself along with access to the array, can involve exponentially more components and dependencies. When you multiply the number of storage arrays and hosts present in the fabric, even simple administrative tasks such as allocating new capacity becomes a challenge requiring careful planning.

The SAN has also blurred the lines between storage and host management. Although a server connected to a SAN can potentially be granted access to any number of storage array volumes in the SAN, doing so involves managing many components in the environment. For example, careful consideration must be taken when associating servers to storage such as making changes in the fabric when provisioning new volumes. In this example, the host must be configured with the correct drivers and firmware which vary greatly according to storage vendor, array version, host bus adapter, and host operating system. Fabric changes may be required too. Many storage vendors seem to take the complexity of these dependencies for granted, as part of the cost of deploying a SAN.

iNSP software was developed to unleash the full potential of the SAN. It transforms the SAN from a complex set of point-to-point connections into an intelligent, adaptive network. By treating storage as virtual **pools** of capacity, iNSP software can remove the burden of directly managing the relationship between hosts and storage, while intelligent policies ensure that reliability and performance requirements are met automatically. iNSP software provides storage services that promote the data volume from being merely a statically allocated LUN to an independent entity that can be moved, copied, and fluidly managed within the network. The iNSP Data Mobility feature set, for example, makes data migration tasks almost effortless.

The full feature set and power of iNSP software goes far beyond the scope of this white paper. Although, it will touch upon a number of key features that demonstrates how iNSP can simplify the deployment and management of intelligent SANs within a Cisco MDS SAN switch environment. A complete description of specific iNSP software and solutions are available in other white papers at Incipient.com, including topics such as:

- iNSP Administrator
- iNSP Data Mobility
- iNSP Copy Services
- iNSP Volume Management

iNSP Software Components

iNSP software has two main components: iNSP and iNSP Administrator. iNSP runs natively on a clustered set of Cisco MDS 9000 Storage Services Modules, also know as intelligent SAN switches (see diagram below). The Cisco SSM is a 32 Port Fibre channel switch blade and works with any Cisco MDS 9500 Director or 9200 fabric switch. The Cisco SSM blade has all of the connectivity capabilities of a traditional non-intelligent SAN switch, plus it contains multiple specialized processors optimized to deliver highly efficient storage virtualization services. Cisco MDS switches provides hardware optimized primitiveness that allows I/O (Input/Output) to be routed to specific physical locations. iNSP software leverages these primitives to provide volume virtualization and other storage services like data migration and copy services.

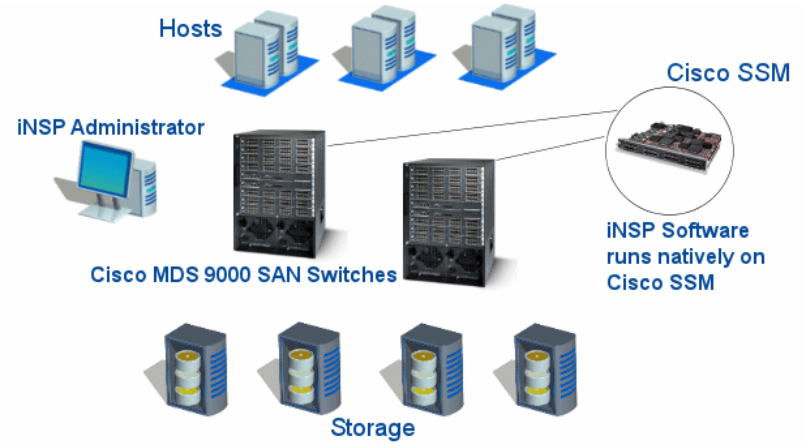

iNSP Administrator resides on a separate server and uses LAN connectivity to monitor and manage iNSP. Its capabilities go far beyond that of a storage array element manager, though, offering the tools to construct an intelligent infrastructure around an existing environment. iNSP Administrator can probe the SAN to discover detailed information about the fabric topology, from zone configurations to attached storage arrays and hosts. With its **Deep Discovery** capability, it can examine storage arrays to obtain information about even the physical layout of storage array volumes, determining the physical disks and RAID configuration. Rather than simply presenting this daunting set of facts to the administrator, it uses this information in conjunction with policy-based planning to make intelligent decisions when provisioning and providing other storage services..

iNSP software in a Cisco MDS based fabric relaxes the dependencies between storage and server. iNSP storage virtualization software acts as an abstraction layer, shielding host operating systems from the complexities of the SAN, and freeing the storage administrator to treat storage arrays as simple containers for their all-important data. iNSP Administrator provides the intelligence to guide that abstraction layer, configuring and placing data in appropriate and available containers. Within a SAN environment, storage administrators allocate capacity from storage arrays into **storage pools**, each with its own performance, reliability, or user-defined attributes. Storage administrators also define **volume classes**: templates which detail the performance, reliability, and layout requirements of a particular user-defined type of data volume. iNSP Administrator, drawing upon policies carefully designed to match industry best practices, matches storage pools and volumes classes to ensure that the SAN configuration that is deployed meets the SAN configuration that is required.

For a full description of the features available when using iNSP Administrator, please see the white paper, iNSP Administrator's Role in SAN Management, available at Incipient.com.

Gradual Deployment of iNSP in a Cisco MDS Based Fabric Environment

When designing new technology for the data center, coexistence and complementing existing infrastructures is the key to protecting existing IT investments. The iNSP software suite was designed to coexist with existing SAN deployments and complement existing SAN management practices and technologies. iNSP can be deployed gradually and without massive infrastructure changes, software upgrades, or firmware upgrades. Adoption of the features and solutions can proceed at the pace decided by the data center administrator and at the convenience of the data center staff.

Every feature within iNSP, from the abstraction of storage virtualization to the detailed investigation of deep discovery, was developed to be gradually added to a SAN environment. No storage volume or storage array, regardless of whether it is attached to iNSP, is *required* to be virtualized. Capacity can be added and managed by the iNSP software suite gradually, volume by volume or array by array. Likewise, existing data and volume layouts can be imported into the system intact, and will remain that way until the iNSP Administrator is given explicit instructions by the user.

iNSP Software and the Switch Platform

The core of the iNSP software is designed to provide a complete storage virtualization environment that is resident within the SAN.

The Intelligent Switch Platform

The Cisco SSM running within the Cisco MDS chassis offers a number of ports with specialized capabilities, each connected within a hierarchy of processing units to deliver a unique set of intelligent features. At the level closest to the port, processors can direct an incoming I/O to be handled according to a small set of mapping and redirecting primitives, at line speed and with little more overhead than an "unintelligent" port. These ports and processors are collectively known as the "**fast path**[1]", since data and I/O commands pass through with little overhead.

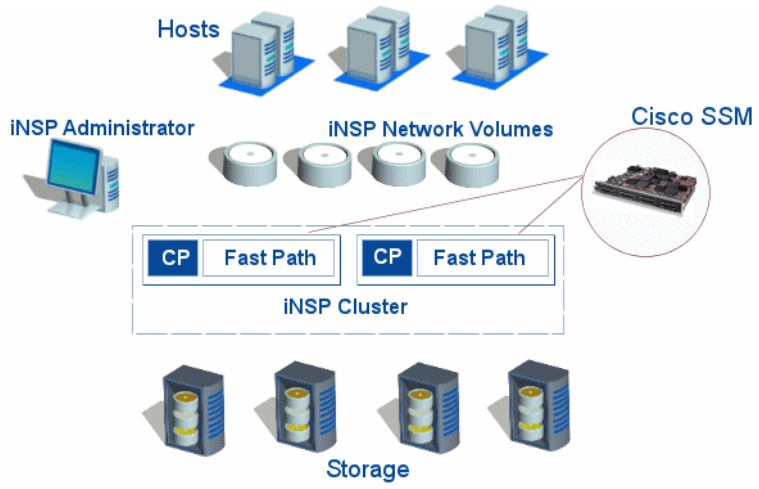

iNSP is deployed on what is called the "**control path**" (CP) of the intelligent switch (see diagram). From there, it monitors all of the fast path ports, configuring and coordinating their simple primitives. From its vantage point, complex storage virtualization layouts are just a matter of directing the correct I/O to the correct physical storage connected to the intelligent ports. Fast path processing ensures that the configurations it creates are handled efficiently. In the extremely rare case that a fast path processor is unable to correctly complete an I/O it receives, iNSP will receive a notification on the control path, and act to resolve it.

> A simple example is a SCSI Inquiry command directed to a virtualized storage volume. The fast path processor will correctly note that this command should be handled by the iNSP software, and will pass the command structure through to the control path. From that point, iNSP will quickly construct the correct response and return it to the fast path.

[1] Fast path is sometimes also termed "data path".

Intelligent switches running iNSP software are configured into clusters for reliability and performance. An iNSP Cluster requires a minimum of two Cisco SSM switch blades. These blades are fully integrated into the fabric, and are able to access storage or hosts attached to any other non-intelligent fibre channel switch blades, given the correct fabric zoning.

Supported Environments

Incipient and its partners conduct rigorous, ongoing testing of the iNSP software suite in a wide variety of hardware and software environments. Staff routinely tests interactions with new and old products from all of the major storage, switch, and operating system vendors, in configurations that range from simple and straightforward to highly complex. iNSP features and layouts are combined to create environments that routinely put far more stress on the software and the hardware than a typical customer will ever encounter, all with the goal of ensuring reliability and quality.

Incipient maintains a support matrix which describes the latest results of our testing. It provides a complete and detailed list of arrays, operating systems, and host bus adapters (HBAs) that have passed ours and our partners testing, as well as firmware versions, driver versions, and multipathing driver versions. For the most current support matrix, please visit Incipient.com.

Deploying iNSP into an Existing Environment

Deploying iNSP into an existing Cisco SAN environment does involve a number of change points. While some of them, such as storage and host configurations, can be done gradually, others, like most network changes, do need to be done up-front. Any product that introduces new hardware as well as new data paths should be treated with the proper level of preparation and planning.

Network Configuration

The placement of iNSP into an existing Cisco SAN, as with any network configuration changes, requires careful planning and design decisions. The first stage in a deployment is locating the appropriate network location in which to insert the intelligent switches. Two slots must be available: either both in a single chassis or for greater fault-tolerance, one in each of two chassis.

Once the hardware is deployed, a decision must be made about the desired SAN topology. An intelligent switch running iNSP is fully compatible with level-2 switching technology, and may be anywhere in the topology and still offer virtualization services. But virtualization doesn't mean that you can forego good SAN practices, and topology still "needs to make sense".

If you deploy a core-edge design, with a core made up of highly available directors, it makes sense to use the intelligent switch in the core of your network. Depending on SAN port count, the number of arrays, and the number of hosts, it may be necessary to use additional layer-2 switches to increase the host- or storage port count, resulting in a 3 tier core-edge design.

Storage Configuration

Once the iNSP software has been installed and configured on the Cisco SSM blades within the MDS 9000 Chassis, iNSP is ready to begin offering storage services. The next step is to grant iNSP access to some actual, physical storage.

Before iNSP is granted access to a storage array volume, arrays exposed through the fabric to hosts retain all of their familiar properties and identities, regardless of whether a connected port is intelligent or not. For any given array or even storage array volume, deployment of iNSP doesn't have to change that. But if a data center administrator does want to take advantage of any of the powerful iNSP features, at least some storage must be placed under their control.

Allocating storage array capacity to be managed by iNSP is simply a matter of exposing storage array volumes to the iNSP nodes. These exposures can be performed using the storage array element manager or any existing SAN management utilities with that capability, including iNSP Administrator itself. From that point on, the allocated capacity becomes the responsibility of iNSP and iNSP Administrator. iNSP Administrator, under strict guidelines by the user, automates the layout and provisioning of that storage. iNSP uses those capacity layouts to provide new, virtualized volumes, called **Network-Volumes**, which can be exposed to the hosts.

While it is sufficient for raw capacity from storage array volumes to be simply added to iNSP storage pools, storage array volumes which contain existing data must be handled differently. These array volumes are **imported** into iNSP as encapsulated volumes, with their layouts left completely intact. They can be immediately mapped to hosts as Network-Volumes, and iNSP will manage host access to them. While configured thusly, any iNSP solution can be used with them except iNSP Network Volume Management. However, data contained in an encapsulated volume *can* be migrated off of that array volume transparently to the host (see the iNSP Network Data Mobility for SAN Management white paper), and the capacity it used can then be re-allocated to the iNSP Administrator storage pools.

Another component of storage configuration that must be performed during initial storage reconfiguration is iNSP metadata management. Information about the array volumes that iNSP is managing is stored within that storage array, and a small amount of capacity must be set aside on each array for that purpose. By placing the metadata describing Network-Volumes utilizing a storage array on that same array, iNSP ensures that access to those volumes will never be affected by the independent failure of any other information store.

Host Configuration

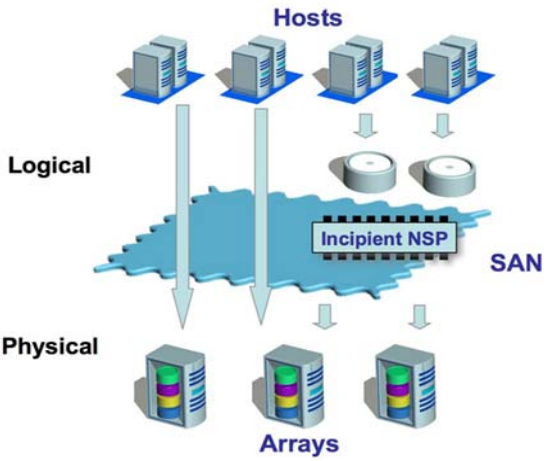

Once iNSP is installed and provided with some storage capacity, it is ready to begin presenting Network-Volumes to hosts. As was the case with storage arrays, hosts can begin to use iNSP gradually. Hosts can still use the SAN to access their old configurations and original storage arrays, and the mere presence of iNSP does not require any server changes. However, until a host is given access to an iNSP Network-Volume, it still must be managed using the original, complex processes.

Once the data center administrator decides to either present a new Network-Volume to a host, or to decouple a host from its existing storage array volume and present it as an encapsulated Network-Volume instead, some reconfiguration is necessary. A key point, though, is that this reconfiguration is the *last* time that this volume will need to be reconfigured on that host. An iNSP Network-Volume completely masks any further changes that need to be made to the underlying physical storage; whether the data is migrated to a new array or just has its layout modified for a performance boost, the host will never see a difference.

Using iNSP Administrator

For a full description of the features available when using iNSP Administrator, please see the white paper, iNSP Administrator's Role in SAN Management. But as mentioned above, certain features are integral to the iNSP environment, and serve to showcase the power of the iNSP solution. In particular, the creation of storage pools gives data center administrators the ability to treat arrays as raw capacity, without concerning themselves with the specifics of provisioning. Likewise, the creation of volume classes lets the administrator define a set of requirements around the data, rather than trying to match configuration to array capabilities.

While storage pools and volume classes are extremely flexible and customizable, a typical design example might involve the creation of three storage pools to represent particular customers or levels of service:

- "Bronze"
    - Intended for non-critical data, and composed of capacity from storage arrays that are considered low-end, and therefore less expensive and reliable.
- "Silver"
    - Intended for business-critical data, allocated from mid-range arrays that are moderately more reliable and expensive than the Bronze pool.
- "Gold"
    - Intended for mission-critical data on high-end arrays that are significantly more expensive and reliable.

Some of the typical volume classes might be:

- "File server"
    - With very basic requirements: Storage-based RAID-5 for performance, multipathing, and drawn from the Bronze pool.
- "e-mail"
    - Stricter requirements: Network-based RAID-0 for performance, composed of physical array volumes with RAID-5 configurations, multipathing, and drawn from the Silver pool.
- "Database logs"
    - Even stricter requirements: Network-based RAID-0 for performance, composed of physical array volumes with RAID-1 configurations, multipathing, and drawn from the Gold pool.
- "Database tables"
    - High-reliability, high-capacity requirements: Network-based RAID-0 for performance, composed of RAID-5 volumes drawn from the Gold pool.

Once these storage pools and volume classes are created, provisioning becomes a simple matter of telling the iNSP Administrator what size of what class of volume should be allocated to which host.

Using iNSP in Disaster Recovery Environments

The same principles that make incremental deployment of iNSP possible also allow its deployment into disaster recovery environments. Array-based services such as EMC SRDF, IBM PPRC, or HP CA can successfully co-exist with iNSP environments, but as with any disaster recovery solution, some care must be taken during planning and configuration. Please see the white paper iNSP Network Volume Management and Array-based Services at Incipient.com for detailed discussion of the steps and issues to be taken into consideration.

Conclusion

The iNSP software suite delivers a bold new set of tools and solutions for storage and SAN management, offering data center administrators the opportunity to simplify storage administration and shrink costs for managing SANs within their data centers:

- **iNSP Administrator** offers a single pane of glass into a SAN, allowing a comprehensive view of the entire SAN and deep discovery of all SAN components,
- **iNSP Network Volume Management** allows both existing and new storage arrays to be managed as anonymous pools of capacity, provisioned according to the data's performance and reliability requirements,
- **iNSP Network Data Mobility** solutions gives the administrator control and flexibility when making SAN changes, masking them from the host while efficiently and rapidly moving data within the network,
- **iNSP Network Copy Services** offer new ways to manage volume copies cost-effectively, utilizing capacity that is appropriate for the actual use of the copy, on the same array or in another storage pool in the SAN.

All of these solutions are designed to be deployed gradually and at the pace determined by the data center administrator. The iNSP software suite will function non-disruptively with existing data, arrays, and host infrastructure, alongside existing data center processes.

Deployment of iNSP offers the best of all possible solutions for data centers seeking solutions to the complexities introduced by storage area networks: investment protection for existing infrastructure, and comprehensive utilities and services to use when planning for the future.

Incipient Storage Glossary

array - A physical storage device (subsystem) with multiple hard drives and RAID functionality.

array controller - A processor in a storage array that provides volume management and other storage services. A single storage array usually contains multiple controllers.

array volume - A SCSI logical unit presented by a storage array.

bound array volume - An array volume which has been assigned to management by the iNSP Administrator.

consistency group - A collection of Network-Volumes that are treated as a single unit for snapshot operations, ensuring referential data integrity across multiple volumes. Applications such as large databases or email servers may span multiple volumes, and the use of consistency groups guarantees that snapshots of all volumes represent a single, atomic point in time.

data mobility - A feature of iNSP that allows the assignment of data to physical storage to be dynamically modified without host disruption. See also *migration, virtualization*.

Deep Discovery - An iNSP Administrator operation which queries attached storage arrays for detailed information about the physical layout of array volumes, including their RAID properties and mappings to physical hard drives.

discovery - The process of obtaining information about the physical and logical devices connected to a SAN, either through in-band (fabric-based) or out-of-band (TCP/IP-based) queries.

fabric - One or more interconnected Fibre Channel switches that can physically transmit data between any two N_Ports on any of the switches.

Fibre Channel (FC) - Fibre Channel (FC) is a technology for networking computer devices. Fibre Channel is primarily used for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives. See also *storage area network (SAN)*.

host - A computer that hosts applications and may initiate I/O commands to a target.

host cluster - A set of two or more hosts configured to cooperatively share SCSI volumes.

iNSP Cluster - A set of two or more iNSP nodes that act together as a single unit to provide storage services on the network.

Information Lifecycle Management - A process which comprehensively and dynamically manages the assignment of storage resources to data, based on attributes such as the criticality of the data, its reference and update frequency, regulatory requirements, and retrieval costs.

intelligent FC switch - An intelligent FC switch is a Layer-3 switch that is bundled with a programmable virtual storage router to enable highly efficient network-based storage services.

migration - The process of moving data from one set of physical extents to another while the Network-Volume remains online.

mirror - A Network-Volume which maintains at least one complete, dynamic copy of the data stored within its logical extents.

mirroring (RAID-1) - A method for data redundancy that maintains two or more identical, dynamic, real-time

copies of data on separate disks. Mirroring is transparent to the host, which perceives a single volume.

Network-Volume - A virtual SCSI logical unit created by the iNSP, constructed from extents of one or more array volumes.

node - One instance of the iNSP software residing in the network (e.g., on an intelligent FC switch processor). Nodes are always deployed in sets of two. Each node is allowed to be in only one cluster.

policy - An enforceable set of rules to restrict and ensure accountability, a particular sequence of actions, or correct allocation and use of storage resources.

pool - A collection of bound array volumes that provide capacity that can be allocated for a specified class of Network-Volumes. For policy control, each user can be given access to one or more pools, and pools may be defined based on a number of different storage attributes.

provisioning - The process of assigning storage to data, to optimize performance, cost, and reliability requirements.

snapshot - A point-in-time copy service that duplicates the contents of a source Network-Volume to a target Network-Volume.

storage area network (SAN) - A dedicated network for storage operations, whose primary purpose is to transfer data between computer systems and storage devices. SANs provide the ability to share a single storage array across many servers or applications.

storage array - An intelligent storage device whose software provides functionality such as RAID protection, volume management, and replication. Also called a *storage subsystem*.

switch - A physical device that selects a path for sending a unit of data to its next destination.

synchronize snapshot - Update the relationship between a source Network-Volume and its target Network-Volume(s), bringing the target's copy of data up to the current point in time.

target - A SCSI device that that receives I/O requests from another device (initiator).

tiered storage - The assignment of data to different types of storage media, based on cost, reliability, and performance requirements.

virtualization - The insertion of an abstraction layer between hosts and storage which enables advanced, network-based functionality, such as: network-based volume management, pooling of storage from multiple SAN-attached storage devices from various vendors, and advanced network-based copy services.

virtualizer - The network component responsible for virtualization.

volume - A virtual disk (logical volume) created by a storage-, network-, or host-based volume manager. A volume may be composed of one or more physical extents, but to the host application, a volume appears as a single logical extent.

volume class - User-defined provisioning template that specifies the RAID attributes, number of paths, min-max volume size, and the storage pools from which resources are selected.

Additional definitions can be found on the Incipient Web Site: www.incipient.com

Doc Number: NSPM-WP-042307

Incipient, Inc.
230 Third Avenue
Waltham, MA 02451

781-906-3501 main phone
781-906-2503 fax
www.incipient.com