NISHAN SYSTEMS

Internet Storage Name Service (iSNS) — A Technical Overview

For more information:

http://www.ietf.org/internet-drafts/draft-ietf-ips-isns-04.txt

Nishan Systems
3850 North First Street
San Jose, CA 95134
Tel 408-519-3700
Fax 408-519-3705
www.NishanSystems.com

Introduction

This paper outlines the technical details of Internet Storage Name Service, or iSNS. It is a standards-track document within the Internet Engineering Task Force (IETF) actively pursued within the IETF IP Storage Work Group, http://www.ietf.org/html.charters/ips-charter.html. The current draft (June 2001) of iSNS is:

http://www.ietf.org/internet-drafts/draft-ietf-ips-isns-04.txt

Authors of iSNS include:

Kevin Gibbons, Nishan Systems
Josh Tseng, Nishan Systems
Charles Monia, Nishan Systems
Franco Travostino, Nortel Networks
Ken Hirata, Vixel Corporation
Mark Bakke, Cisco Systems
Jim Hafner, IBM Research
Howard Hall, Pirus Networks

For questions or comments on this paper, please email mailto:tclark@NishanSystems.com .

Background on Storage and Storage Networks

As the repositories of information that sustain institutions and enterprises, storage devices are at the center of modern network designs. High-speed access and high availability of stored data are now critical for both internal business operations and external e-commerce applications. Storage devices, however, represent unique components in network design due to their specific requirements. Unlike hosts in a typical data communication network, storage devices do not initiate transactions. In the vocabulary of storage administration, storage devices are *targets*, responding to read and write requests from *initiators* such as servers or workstations. In direct SCSI-attached storage configurations, a server can easily discover the storage resources at its disposal by polling the SCSI bus to which disk arrays are attached. Since the server is the exclusive owner of its attached storage, it can be assumed that any devices it discovers are available for its use.

In a storage area network (SAN), by contrast, disk and tape resources may be dispersed across a complex network. This situation presents challenges for device discovery as well as device ownership. Initiators must be able to identify storage resources in the SAN and determine whether they have access to them. SANs therefore require a mechanism to facilitate device discovery and to assign potentially shared storage resources to specific initiators. iSNS facilitates device discovery in Fibre Channel storage area networks and in IP Storage Networks.

Discovery in Fibre Channel SANs

Fibre Channel SANs support two basic topologies: arbitrated loop and fabric. An arbitrated loop is a shared medium, analogous to a token ring LAN. Devices attached to the loop must arbitrate for access to the shared bandwidth before launching transactions, and only 126 end nodes are supported. Fabric topologies, in contrast, provide full bandwidth to each device and can support up to 15.5 million end nodes. The two topologies can be combined — for example, a loop segment can be attached to a fabric switch to allow communication between loop nodes and fabric-attached nodes.

Discovery in an arbitrated loop relies heavily on the limited number of devices that can be connected to a single loop. Following loop initialization, an initiator such as a server simply can poll through the 126 possible addresses to solicit responses from potential targets. Polling through the address space for 126 possible destinations may be inefficient, but occurs fairly quickly at gigabit speeds.

Within an arbitrated loop, targets that respond to an initiator's queries verify their presence on the loop and an initiator can thereafter establish sessions with each device and begin storage transactions. In some proprietary implementations, a positional map that is generated during loop initialization is used for target discovery. Each active device records its loop address in the positional map, and the map in turn can be used to create an address list to facilitate session establishment. While this feature, called positional mapping, shortens the device polling process, not all loop devices support it.

Device polling is not a viable solution for Fibre Channel fabrics, however, because there are over 15.5 million possible addresses. Fibre Channel standards therefore provide a name service definition that enables device discovery without walking through an enormous address space. Whereas in loop environments, the initiator must perform all the work of device discovery, in a fabric topology, this responsibility is shared between initiators and the Fibre Channel switches that compose the fabric topology.

A device attached to a fabric switch must first log onto the fabric to obtain a unique network address. The device must then register its presence in the fabric by logging on to the Simple Name Server (SNS) at a well-known address. The SNS maintained by the switch is a small database containing the permanent device identifier, fabric address, class of service parameters, and other information. Of special importance for device discovery, the SNS provides an entry for the type of upper-layer protocol supported, which for storage applications is serial SCSI-3. Since every target device registers with the SNS, an initiator simply can send a query to the SNS asking for its list of devices that support the serial SCSI-3 protocol. The address list returned from the SNS then becomes the polling list for the initiator, which can in turn send port logins to the listed targets and establish sessions with them.

In some cases, it may not be desirable for an initiator to discover all possible targets in the Fibre Channel fabric. An NT server, for example, probably should not be allowed to discover and establish a session with a Unix storage array as the NT server could potentially overwrite the

array's boot sector and render it useless to Unix. Given that the SNS reports only addressing and protocol support information, there is no means within the SNS itself to restrict discovery of target devices. In Fibre Channel fabric switches, segregation of devices is possible only through a technique called zoning. Zoning creates groups of devices authorized, either on the basis of port attachment or via WWNs, to communicate with each other. In certain implementations, zoning definitions override an initiator's SNS query, and thus the SNS reports only those targets that are in the same zone as the initiator. This leaves the initiator safely ignorant of other storage resources that may present conflicts and/or security breaches if reached by the 'wrong' initiator.

A third and related component of device discovery accommodates changes that may occur once devices have been discovered. The state change notification service is provided by a facility within the fabric that allows initiators to be notified if storage resources are removed or added to the network. State change notification (SCN) enables active responses to resource availability. If, for example, a new storage resource enters the fabric, initiators can be notified and quickly be allowed to establish sessions with it, if appropriate and allowed by its zoning permissions.

One of the problematic issues for Fibre Channel device discovery concerns scalability. Each fabric switch maintains its own SNS database. As more switches are added to a single fabric, they must be able to share SNS data so that an initiator anywhere on the network can discover viable targets. In addition, since zoning may be used to restrict the discovery process, zone information also must be exchanged. This scheme places an ever-increasing burden on switch resources as SANs scale from small departmental configurations to larger implementations. Network convergence time is adversely impacted by the greater latency of a large network as switches update each other with SNS, zoning, and SCN information.

Discovery in IP Storage Networks

While Fibre Channel has been forced to pioneer device discovery techniques where no precedents existed, IP Storage networks are able to draw from both IP networking applications such as Domain Name Server (DNS) as well as Fibre Channel applications including SNS, zoning, and state change notification. Nishan Systems authored the first comprehensive discovery proposal for IP Storage, Internet Storage Name Service (iSNS), which has been submitted as a draft for standardization consideration to the Internet Engineering Task Force (IETF). iSNS leverages the database objects of SNS as well as familiar DNS techniques to create a discovery mechanism that can be centralized or distributed, which allows it to be scalable.

IP Storage Protocols

Because IP Storage solutions can be based on several distinct protocols, iSNS provides support for a variety of implementations of block-based storage over IP. The three main protocols for IP Storage networks are Fibre Channel over IP (FCIP), Internet Fibre Channel Protocol (iFCP), and Internet SCSI (iSCSI). As shown in Figure 1, the FCIP, iFCP, and iSCSI protocols support serial SCSI-3 interfaces to the standard SCSI command set expected by the operating system and upper-layer applications. This interface allows conventional storage I/O to be performed over a high performance gigabit transport. Serial SCSI-3 transactions are carried over TCP/IP, although only iFCP and iSCSI leverage native TCP/IP for each storage end device. Each IP Storage protocol has unique requirements for discovery.

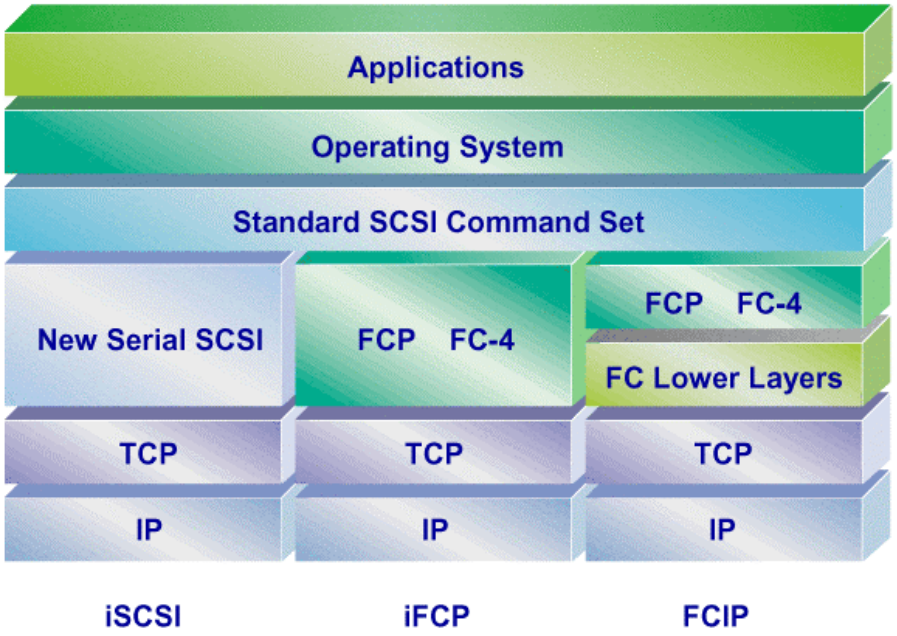

Figure 1: *iSCSI, iFCP, and FCIP Protocol Stacks*

FCIP is used to tunnel Fibre Channel traffic between two geographically separate Fibre Channel SANs. As shown in Figure 2, frames originating on one SAN are wrapped in IP packets and forwarded to the destination SAN. At the receiving end, the IP header is removed and native Fibre Channel frames are delivered to the fabric. A Fibre Channel fabric switch then makes the decision as to which end device the frame is intended. In terms of discovery, the only devices that have IP addresses are the FCIP gateways themselves. IP discovery is thus limited to the FCIP gateways, while Fibre Channel discovery and management is still required for the Fibre Channel storage end devices. Since FCIP tunneling requires both IP and Fibre Channel management applications, additional overhead is necessary for a tunneled solution. Due to the limited management requirements of FCIP gateways, FCIP is not included in the discovery and management services provided by iSNS.

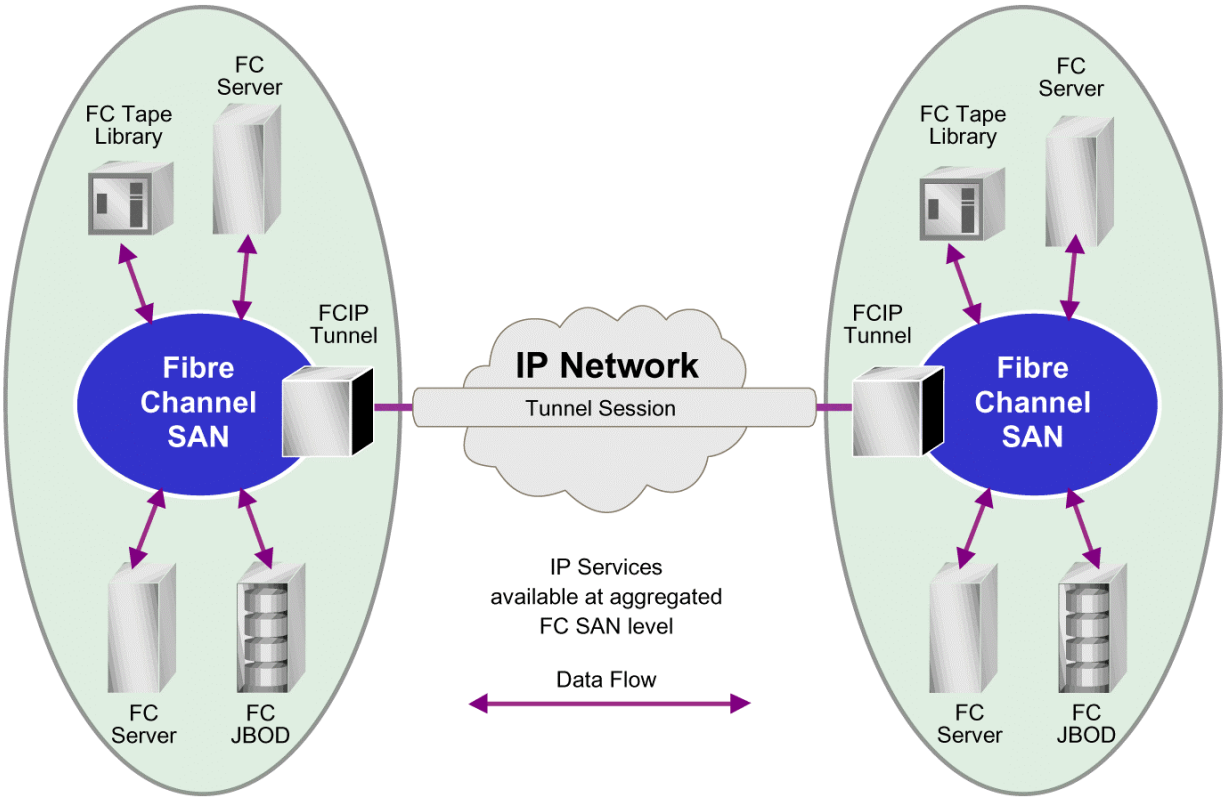

Figure 2: *Joining Two Fibre Channel SANs with an FCIP Tunnel*

While FCIP can be used to connect only Fibre Channel SANs, and not individual devices, iFCP can be used to connect individual Fibre Channel devices to create a native IP Storage network. As shown in Figure 3, storage devices can be dispersed throughout an IP network without being confined to Fibre Channel SANs. iFCP storage switches are a direct replacement for Fibre Channel switches, which implies that iFCP needs something comparable to SNS for end node discovery. This function is included in iSNS.

n addition to a 3-byte Fibre Channel address, an iFCP switch assigns a 4-byte IP address to each Fibre Channel end node. When a Fibre Channel device sends an SNS query, the request is intercepted and interpreted by iSNS. At the Fibre Channel layer, a list of appropriate target addresses are reported to the initiator, while the IP “storage aware” fabric handles the mapping of Fibre Channel addresses to their corresponding IP addresses to enable devices to be linked across an IP network. The list of iSNS entries for iFCP devices therefore includes standard Fibre Channel-specific objects such as port address and upper layer protocol support (e.g., SCSI-3), as well as IP-specific entries.

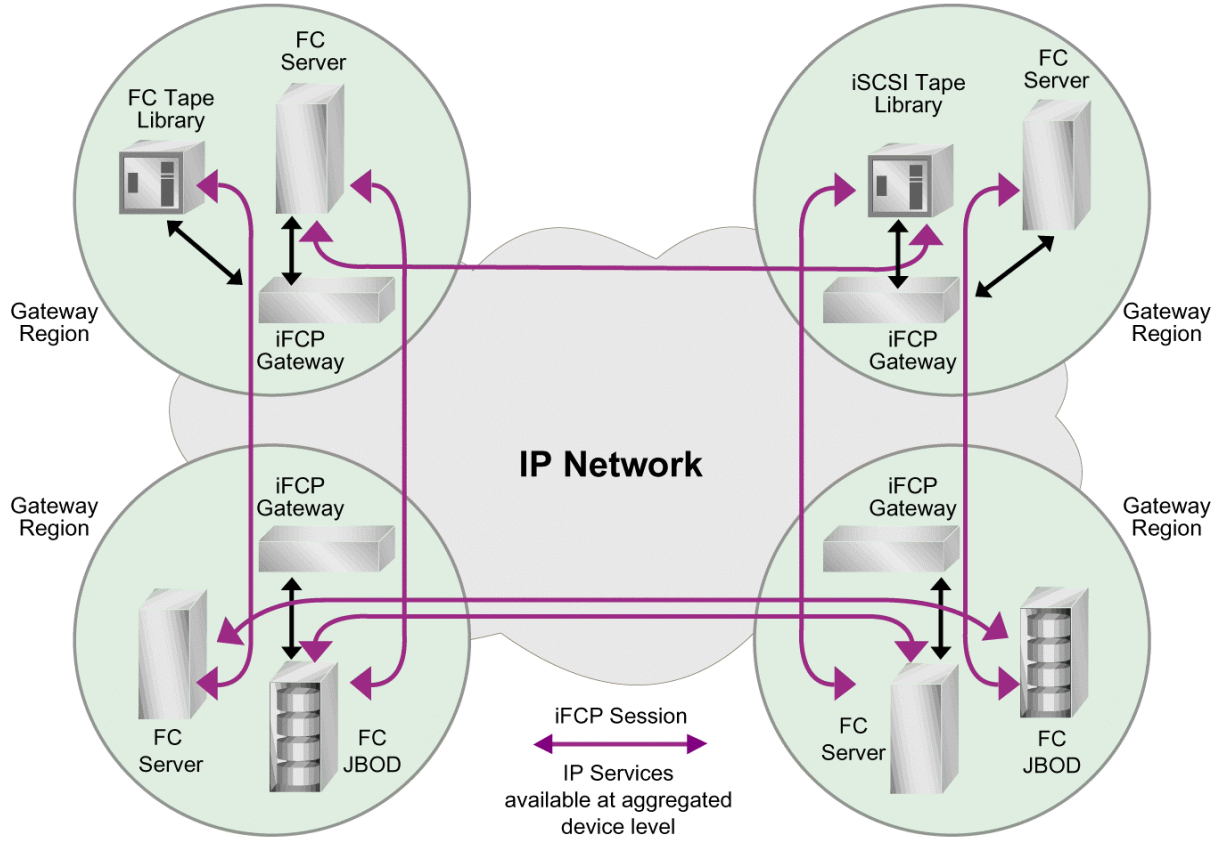

Figure 3: *An iFCP storage network*

iSCSI is a start-from-scratch reconstruction of a serial SCSI-3 protocol residing over the IP stack, and it assumes that both initiators and targets are native iSCSI devices. As shown in Figure 4, iSCSI end devices can be connected by conventional IP routers and switches. This network configuration has the advantage of native IP attachment, but does not accommodate existing Fibre Channel devices. For iSCSI implementations, a gateway is required to bring both iSCSI and Fibre Channel devices into the same network. For an iSCSI initiator to discover iSCSI targets, it needs to identify which devices in the network are storage resources and what IP addresses it needs to access them. A query to an iSNS server consequently returns a list of IP addresses that the initiator has permission to access.

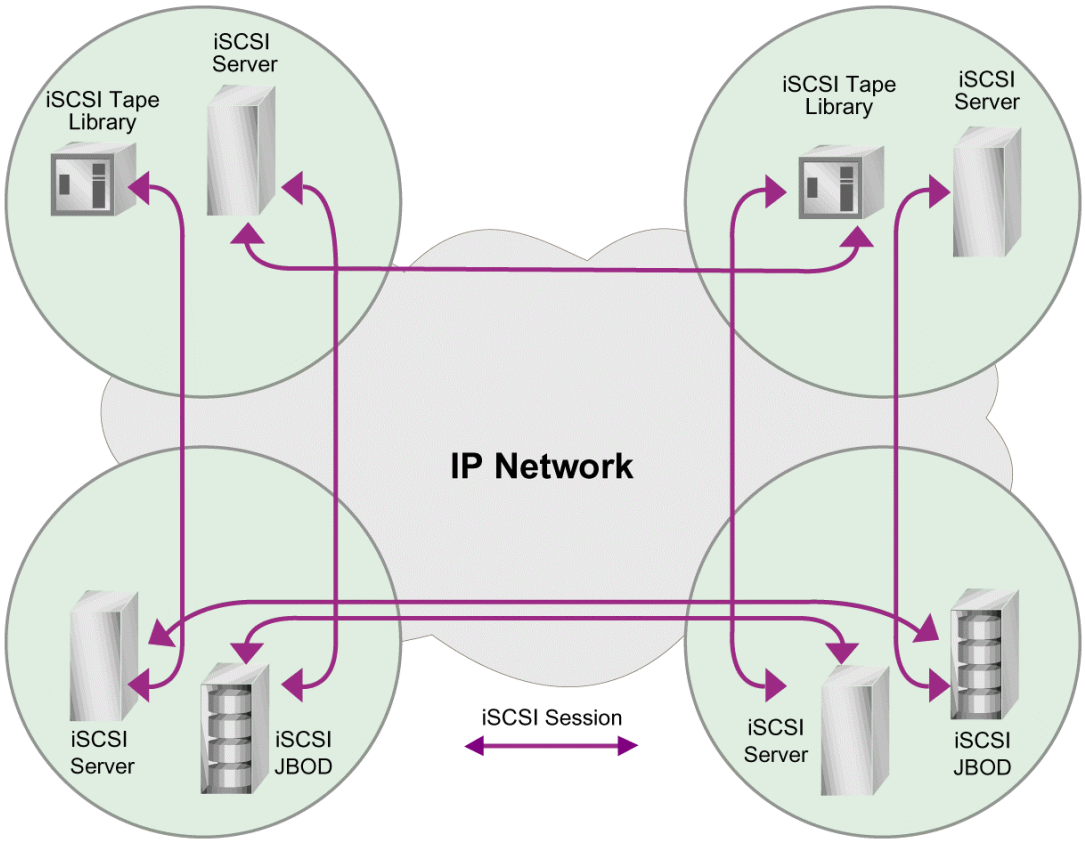

***Figure 4**: An iSCSI storage network*

An enterprise network may contain all three IP Storage protocol implementations, which presents an additional challenge for any discovery method. While FCIP will maintain Fibre Channel mechanisms for discovery, iSNS accommodates diversity across devices by including mechanisms for iSCSI and iFCP as well as future native IP Storage protocols that may emerge.

iSNS Features

iSNS is designed to be a lightweight discovery protocol that can be deployed in iSNS servers, IP Storage switches, and target devices. Features include facilities for registration, discovery, and management of IP Storage resources as well as zoning and state change management. The name registration service enables IP Storage devices to register their attributes and addresses in a manner analogous to Fibre Channel SNS. Initiators then can query the iSNS to identify potential targets. Zoning functionality is provided by Discovery Domains, which restrict the discovery of IP Storage targets to authorized functional groups. State change notification alerts iSNS clients to any change in status of a registered device or reconfiguration of the client's Discovery Domain.

Discovery Domains enable a device to participate in one or more zones. Like Fibre Channel zones, Discovery Domains must be manually administered, at least for the initial establishment of functional groups within the network. By default, a new device is isolated from the storage network until a management workstation assigns it to a specific Discovery Domain. This protective feature prevents inadvertent access by unauthorized initiators. Once a Discovery Domain has been configured for the device, state change notification is used to alert authorized initiators that a new resource has been added to the Domain.

iSNS also supports Discovery Domain Sets (DDS). Analogous to zone sets in Fibre Channel, a DDS can be used to quickly reconfigure an IP SAN for different application requirements. One DDS, for example, could include a tape resource in an NT Discovery Domain for one configuration, while an alternate Discovery Domain configuration could move the tape device into a Unix Discovery Domain.

As shown in Figure 5, the iSNS server can reside anywhere within the IP network, accessible by iFCP or iSCSI clients. One or more management workstations communicate with the iSNS server, either by iSNS protocol or Simple Network Management Protocol (SNMP).

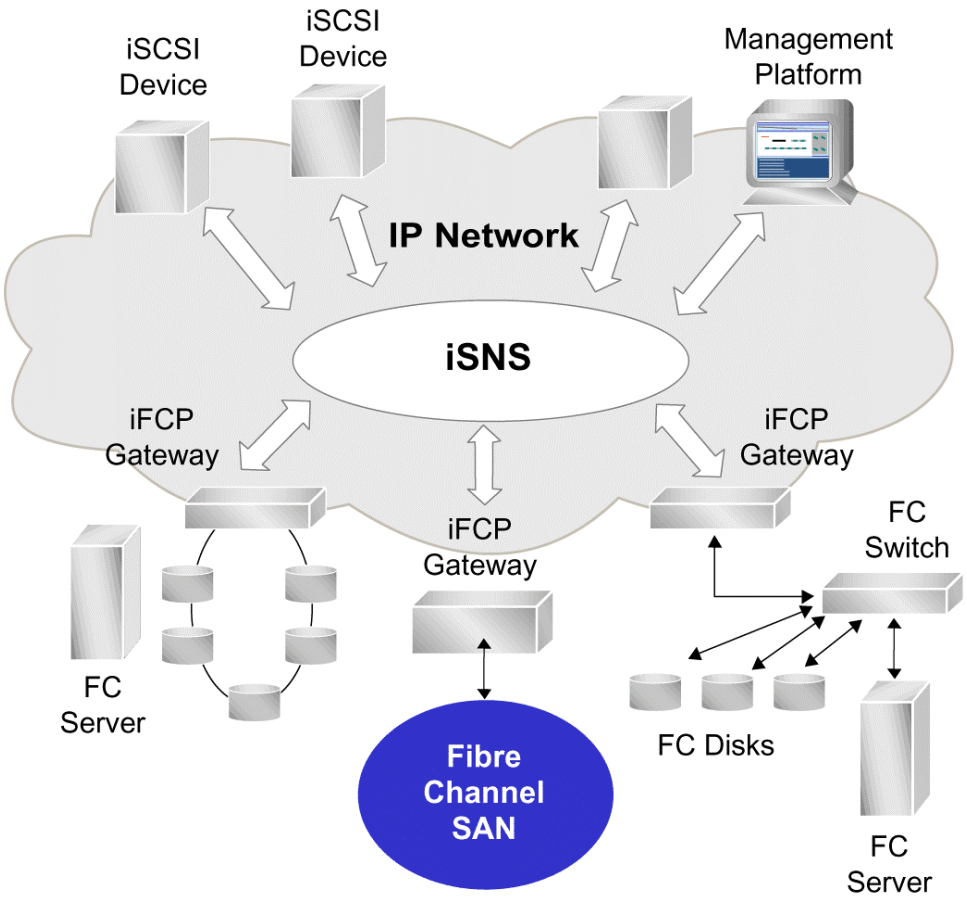

Figure 5: *An IP Storage network with iSNS server and clients*

Since iSNS provides a common resource for a variety of IP Storage types, each can register with and query the iSNS server for information appropriate to the functionality it supports. An iFCP gateway, for example, could be notified of a change of state of a peer gateway. The iSNS server provides the information database, but creative use of this information is product dependent. An iFCP storage switch, for example, could query the iSNS server for the existence of iSCSI storage targets and proxy additional entries that would make those resources available to iFCP initiators.

A final benefit of iSNS is that it facilitates the integration of existing FC devices with next-generation iSCSI networks. iSNS accomplishes this by providing a common device representation model for both iSCSI and FC devices. Because the iSNS server is capable of simultaneously storing information about both types of devices, a gateway bridging legacy Fibre Channel devices to iSCSI devices can use the iSNS to transparently map FC device name references (based upon World Wide Names) to the appropriate iSCSI device name alias. Similarly, the same gateway can use iSNS to map iSCSI device references to the equivalent Fibre Channel name alias.

iSNS Discovery Process

The first step in the iSNS discovery process is device registration. Depending on the IP Storage device type (iFCP or iSCSI), a device will register its attributes and address information to the iSNS server. The server thus builds a database of iSNS clients, which forms the raw material for assignment of Discovery Domains.

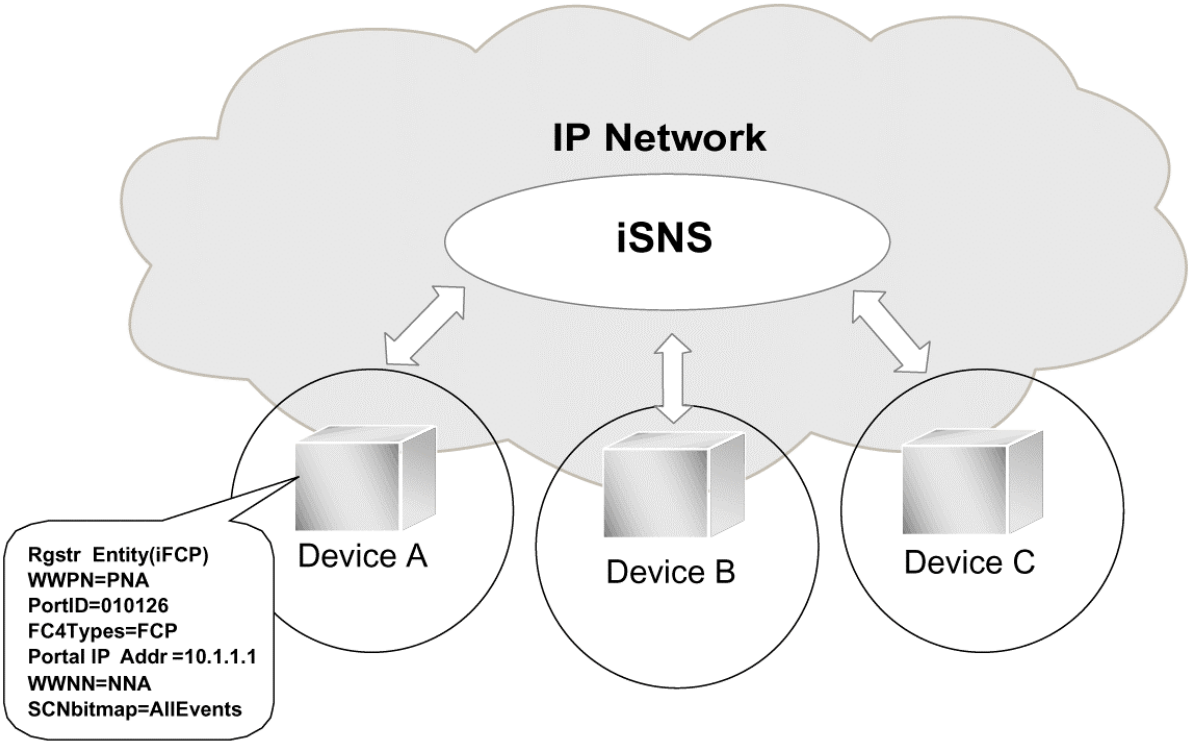

Figure 6: *iSNS Registration*

In Figure 6, devices A, B, and C have registered with the iSNS server. An example of the registration information for an iFCP device is shown for device A. This includes the entity type (iFCP), the device's 64-bit World Wide Port Name, a port ID number, the upper layer protocol support (in this case, FCP), the device's iFCP IP address, a 64-bit World Wide Node Name, and a bit map signifying what state changes this device should be alerted to (in this case, all events). Since these devices have just registered with the iSNS server, they have not yet been assigned to Discovery Domains. For initiators, this means that no storage resources are visible.

Once devices have registered with the iSNS server, zoning information is supplied by a management workstation. With the appropriate Discovery Domains defined, the iSNS server can notify the clients that a reconfiguration of the network has occurred. As shown in Figure 7, this is done via state change notification.

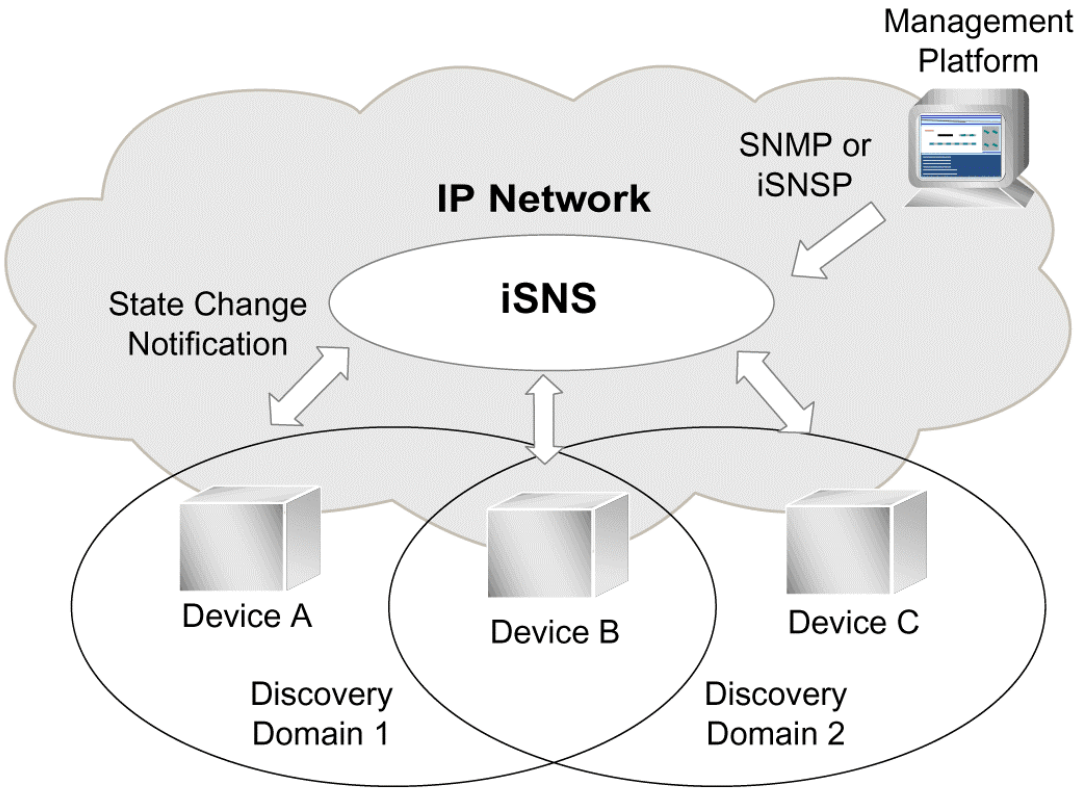

Figure 7: *Creation of Discovery Domains Following Device Registration*

In this example, two Discovery Domains (DD 1 and DD 2) have been created via the management workstation. These Domains group devices A and B into a common zone, and devices B and C into a zone. Since device B is a participant in two Discovery Domains, it will be able to discover both devices A and C. Device A, however, will only discover device B, or any additional resources that are subsequently added to DD 1.

The state change notification issued by the iSNS server will prompt any initiators to query the iSNS for available resources. An iSNS response to a query by device A, for example, would return the address and SCSI-3 capability of device B. Device A can then perform a login to device B and begin storage transactions.

This discovery process can scale from small departmental SANs to extended enterprise-class SANs that may span regional, national, or international boundaries. Except for the initial creation of Discovery Domains via management, the iSNS discovery process is automatic and requires no further administration. Network administrators, however, have the flexibility to reassign resources through reconfiguration of Discovery Domains and can verify network participation through the iSNS information base.

iSNS State Change Notification and Entity Status Inquiry

In the example above, state change notification was used to alert devices to changes in Discovery Domain configuration. As in Fibre Channel, an SCN is triggered by management instruction or by addition or removal of a device from the storage switch. SCN allows for active management of end nodes and enables iSNS to maintain updated information on device availability.

Since the storage switch is directly monitoring the insertion or removal of nodes on its ports, it is immediately aware of changes and can generate change notification. In IP Storage environments, however, a native iFCP or iSCSI device may not be directly attached to an SNS-aware switch, but to a standard Gigabit Ethernet switch. A means is therefore required to monitor state changes in devices that may be anywhere in an IP routed network. For iSNS, this is achieved through Entity Status Inquiry (ESI). The iSNS server polls a registered device at pre-established intervals to monitor its availability. If an ESI Response is not received after a number of retries, the device is de-registered from the iSNS server and, in the case of target devices, a state change notification will be issued to interested initiators. The iSNS server thus can maintain an updated database of active devices and actively report any changes throughout the IP Storage network.

iSNS Objects

iSNS database objects include structures that are broad enough to support a diversity of products and specific, when required, for detailed information on a device's attributes. At the top of the object hierarchy is the concept of a Network Entity. A Network Entity could be an iFCP gateway or an iSCSI gateway or device. The Network Entity will have one or more IP interfaces to the network, defined as Portals in iSNS.

The iFCP object model defines the iFCP gateway as well as iFCP storage devices it services. As shown in Figure 8, the Network Entity for iFCP can be an iFCP gateway with Fibre Channel loop or fabric attached storage devices. The iFCP storage switch represented by the Network Entity designation can have one or more Portals into the IP network, with unique IP address and TCP port numbers. The iFCP object for iSNS also includes a Storage Node, representing a disk controller or tape device that may have several Fibre Channel connections (Storage Ports) to the gateway. The Storage Node and Storage Port objects follow traditional Fibre Channel naming conventions, with both Node and Port World Wide Names. For discovery, each Storage Node that represents a disk or tape resource is identified as a target device.

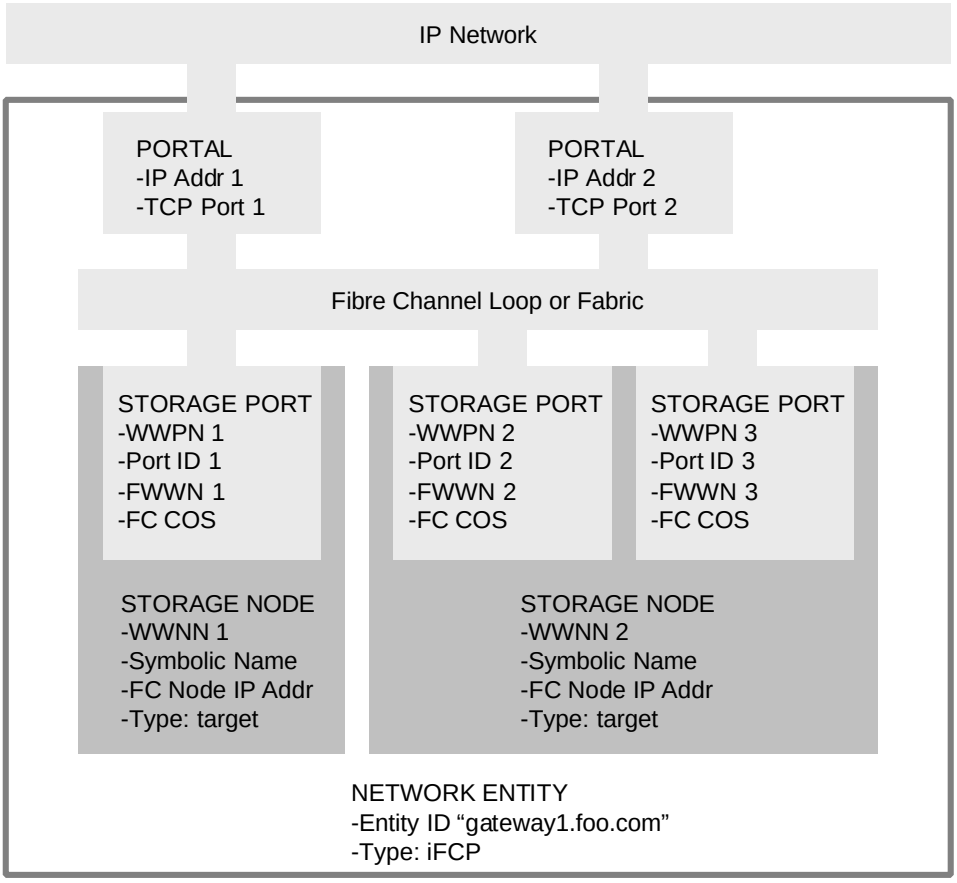

Figure 8: *The iSNS Object Model for iFCP*

For iSCSI, the iSNS object model includes the Network Entity, its Portal to the IP network, and a Storage Node object that specifies whether the iSCSI device is an initiator or target. In the case of iSCSI devices, the 64-bit identifier corresponding to a Fibre Channel World Wide Name is known as a World Wide Unique Identifier (WWUI). For discovery, an iSCSI initiator would register its own presence with the iSNS server and, after being notified of a state change of Discovery Domains, query the iSNS server for Storage Nodes with a Type of “target”.

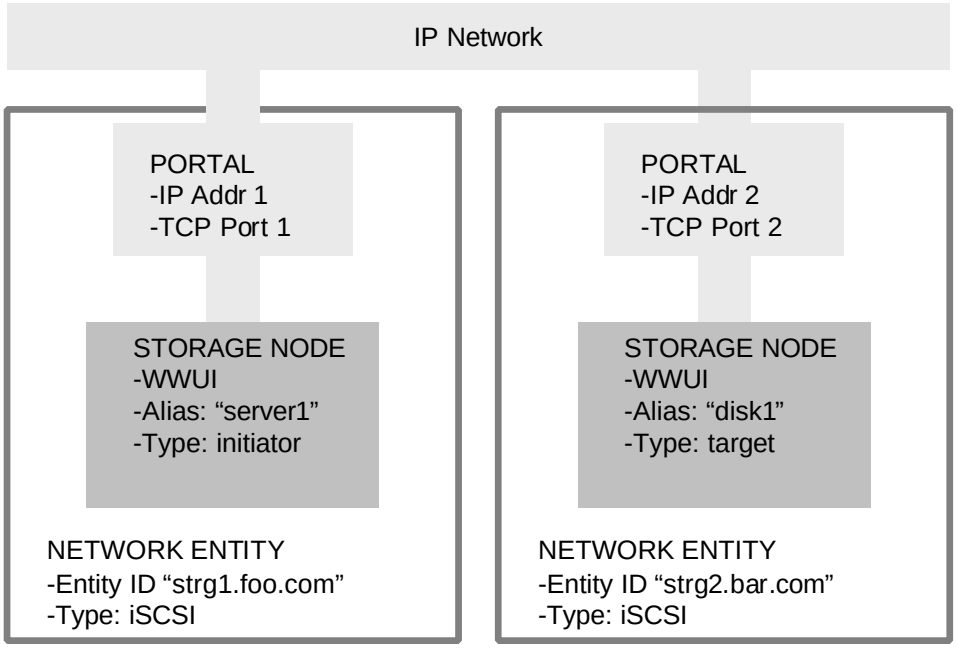

Figure 9: *The iSNS Object Model for iSCSI*

iSNS Security

In addition to restricting access between initiators and targets via Discovery Domains, it is also desirable to have higher-level authentication for storage resources. As the central repository of iSNS data for device discovery and Discovery Domain enforcement, the iSNS server is the logical place to host security services. As part of the registration process, for example, an IP Storage device could register its X.509 Public Key Certificate with the iSNS server. Once Discovery Domains are established, the iSNS server can distribute the appropriate Public Keys between devices in the same Domain. As shown in Figure 10, the exchange of Private Keys and digital signatures necessary for device authentication occurs during the login process. In this example, an iSCSI login between initiator and target includes Public and Private Key exchanges to establish secure communication between them.

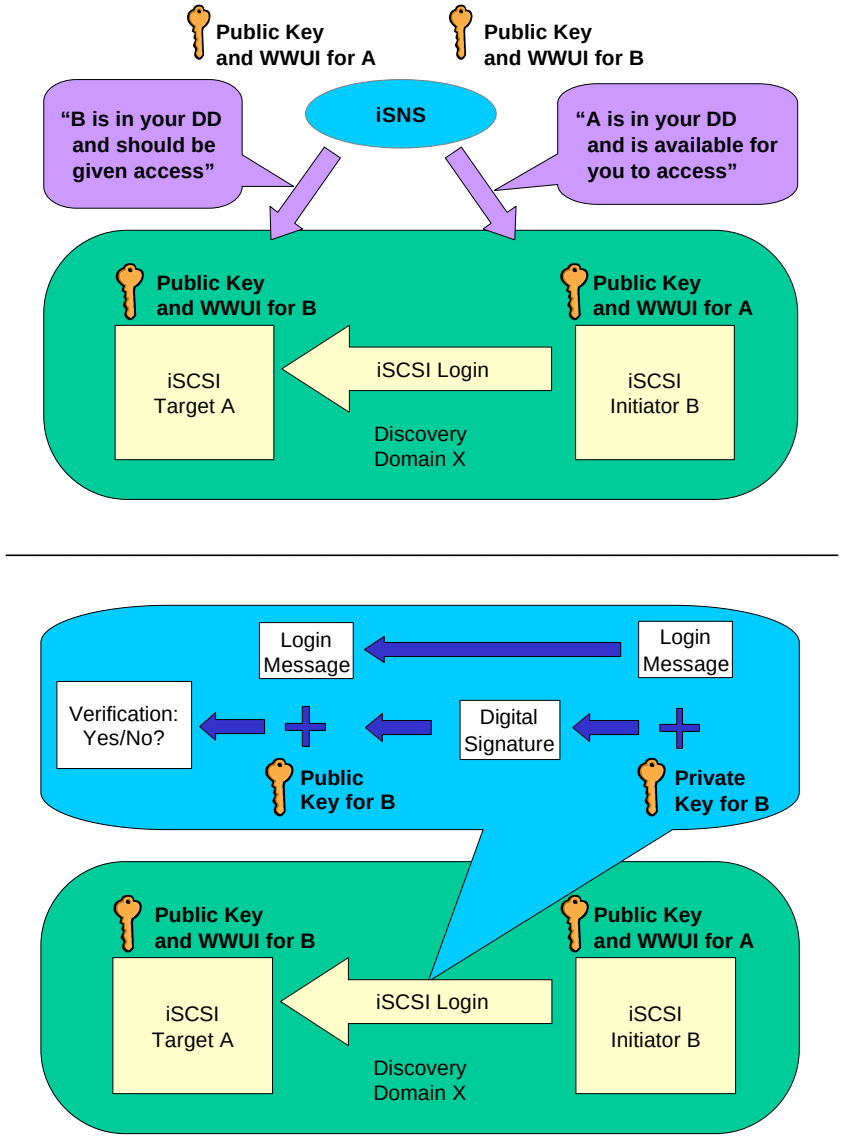

***Figure 10**: X.509 Public Key authentication via iSNS*

The advantage of Public Key Certificates over other security methods is scalability. While manual administration (e.g., Kerberos) may be suitable for small IP Storage networks, it is more convenient to leverage Public Key distribution from iSNS for enterprise-class storage networks.

Summary

iSNS provides a comprehensive discovery and management solution for a variety of IP Storage devices. As a lightweight protocol, iSNS functions can be embedded in an IP Storage switch, gateway, or router, or centralized in an iSNS server. iSNS minimizes the manual administration of an IP Storage network by automating the process of registration and state change notification. With Public Key Certificate distribution, iSNS also enables secure storage transactions between devices that may exist in a public or corporate IP network. Finally, iSNS facilitates the integration of next-generation iSCSI devices with existing Fibre Channel devices and fabrics.

As one of the original authors of the iSNS protocol, Nishan Systems is incorporating iSNS support in its product line, including the IPS 3000 Series IP Storage switch, IPS 1000 Series IP Storage gateway, and IPS 2000 Series Gigabit Ethernet and SCSI storage switch. iSNS support insures compatibility for mixed IP Storage environments and facilitates topology discovery for both IP Storage and legacy Fibre Channel devices.

Printed in the U.S.A. NSWP-06, August 2001.