Network Working Group
Request for Comments: 4174
Category: Standards Track

C. Monia
Consultant
J. Tseng
Riverbed Technology
K. Gibbons
McDATA Corporation

September 2005

The IPv4 Dynamic Host Configuration Protocol (DHCP) Option for the Internet Storage Name Service

Status of This Memo

Copyright Notice

Abstract

This document describes the Dynamic Host Configuration Protocol (DHCP) option to allow Internet Storage Name Service (iSNS) clients to discover the location of the iSNS server automatically through the use of DHCP for IPv4. iSNS provides discovery and management capabilities for Internet SCSI (iSCSI) and Internet Fibre Channel Protocol (iFCP) storage devices in an enterprise-scale IP storage network. iSNS provides intelligent storage management services comparable to those found in Fibre Channel networks, allowing a commodity IP network to function in a similar capacity to that of a storage area network.

Table of Contents

1. Introduction

The Dynamic Host Configuration Protocol for IPv4 provides a framework for passing configuration information to hosts. Its usefulness extends to hosts and devices using the iSCSI and iFCP protocols to connect to block level storage assets over a TCP/IP network.

The iSNS Protocol provides a framework for automated discovery, management, and configuration of iSCSI and iFCP devices on a TCP/IP network. It provides functionality similar to that found on Fibre Channel networks, except that iSNS works within the context of an IP network. iSNS thereby provides the requisite storage intelligence to IP networks that are standard on existing Fibre Channel networks.

Existing DHCP options cannot be used to find iSNS servers for the following reasons:

a) iSNS functionality is distinctly different from other protocols using DHCP options. Specifically, iSNS provides a significant superset of capabilities compared to typical name resolution protocols such as DNS. It is designed to support client devices that allow themselves to be configured and managed from a central iSNS server.

b) iSNS requires a DHCP option format that provides more than the location of the iSNS server. The DHCP option has to specify the subset of iSNS services that may be actively used by the iSNS client.

The DHCP option number for iSNS is used by iSCSI and iFCP devices to discover the location and role of the iSNS server. The DHCP option number assigned for iSNS by IANA is 83.

1.1. Conventions Used in This Document

iSNS refers to the Internet Storage Name Service framework, which consists of the storage network model and associated services.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

All frame formats are in big-endian network byte order. RESERVED fields SHOULD be set to zero.

This document uses the following terms:

"iSNS Client" - iSNS clients are processes resident in iSCSI and iFCP devices that initiate transactions with the iSNS server using the iSNS Protocol.

"iSNS Server" - The iSNS server responds to iSNS protocol query and registration messages and initiates asynchronous notification messages. The iSNS server stores information registered by iSNS clients.

"iSCSI (Internet SCSI)" - iSCSI is an encapsulation of SCSI for a new generation of storage devices interconnected with TCP/IP.

"iFCP (Internet Fibre Channel Protocol)" - iFCP is a gateway-to-gateway protocol designed to interconnect existing Fibre Channel devices using TCP/IP. iFCP maps the Fibre Channel transport and fabric services to TCP/IP.

2. iSNS Option for DHCP

This option specifies the location of the primary and backup iSNS servers and the iSNS services available to an iSNS client.

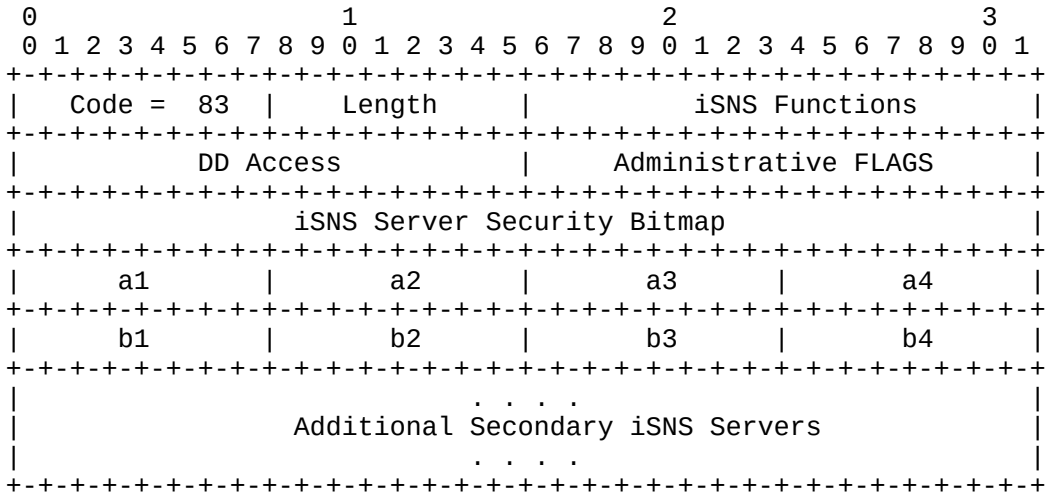

Figure 1. iSNS Server Option

The iSNS Option specifies a list of IP addresses used by iSNS servers. The option contains the following parameters:

Length: The number of bytes that follow the Length field.

iSNS Functions: A bitmapped field defining the functions supported by the iSNS servers. The format of this field is described in section 2.1.

Discovery Domain Access: A bit field indicating the types of iSNS clients that are allowed to modify Discovery Domains. The field contents are described in section 2.2.

Administrative Flags field: Contains the administrative settings for the iSNS servers discovered through the DHCP query. The contents of this field are described in section 2.3.

iSNS Server Security Bitmap: Contains the iSNS server security settings specified in section 2.4.

a1...a4: Depending on the setting of the Heartbeat bit in the Administrative Flags field (see section 2.3), this field contains either the IP address from which the iSNS heartbeat originates (see [iSNS]) or the IP address of the primary iSNS server.

b1...b4: Depending on the setting of Heartbeat bit in the Administrative Flags field (see section 2.3), this field contains either the IP address of the primary iSNS server or that of a secondary iSNS server.

Additional Secondary iSNS Servers: Each set of four octets specifies the IP address of a secondary iSNS server.

The Code field through IP address field a1...a4 MUST be present in every response to the iSNS query; therefore the Length field has a minimum value of 14.

If the Heartbeat bit is set in the Administrative Flags field (see section 2.3), then b1...b4 MUST also be present. In this case, the minimum value of the Length field is 18.

The inclusion of Additional Secondary iSNS Servers in the response MUST be indicated by increasing the Length field accordingly.

2.1. iSNS Functions Field

The iSNS Functions Field defines the iSNS server's operational role (i.e., how the iSNS server is to be used). The iSNS server's role can be as basic as providing simple discovery information, or as significant as providing IKE/IPSec security policies and certificates for the use of iSCSI and iFCP devices. The format of the iSNS Functions field is shown in Figure 2.

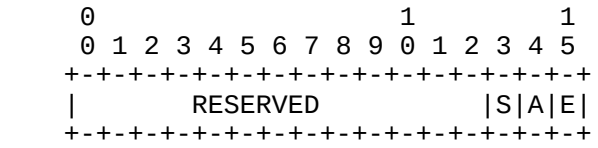

Figure 2. iSNS Functions Field

Bit Field	Significance
15	Function Fields Enabled
14	DD-Based Authorization
13	Security Policy Distribution

The following are iSNS Functions Field definitions:

Function Fields Enabled: Specifies the validity of the remaining iSNS Function fields. If it is set to one, then the contents of all other iSNS Function fields are valid. If it is set to zero, then the contents of all other iSNS Function fields MUST be ignored.

DD-based Authorization: Indicates whether devices in a common Discovery Domain (DD) are implicitly authorized to access one another. Although Discovery Domains control the scope of device discovery, they do not necessarily indicate whether a domain member is authorized to access discovered devices. If this bit is set to one, then devices in a common Discovery Domain are automatically allowed access to each other (if successfully authenticated). If this bit is set to zero, then access authorization is not implied by domain membership and must be explicitly performed by each device. In either case, devices not in a common discovery domain are not allowed to access each other.

Security Policy Distribution: Indicates whether the iSNS client is to download and use the security policy configuration stored in the iSNS server. If it is set to one, then the policy is stored in the iSNS server and must be used by the iSNS client for its own security policy. If it is set to zero, then the iSNS client must obtain its security policy configuration by other means.

2.2. Discovery Domain Access Field

The format of the DD Access bit field is shown in Figure 3.

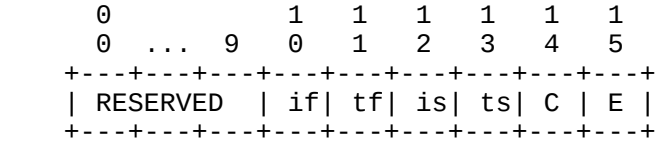

Figure 3. Discovery Domain Access Field

Bit Field	Significance
15	Enabled
14	Control Node
13	iSCSI Target
12	iSCSI Initiator
11	iFCP Target Port
10	iFCP Initiator Port

The following are Discovery Domain Access Field definitions:

Enabled: Specifies the validity of the remaining DD Access bit field. If it is set to one, then the contents of the remainder of the DD Access field are valid. If it is set to zero, then the contents of the remainder of this field MUST be ignored.

Control Node: Specifies whether the iSNS server allows Discovery Domains to be added, modified, or deleted by means of Control Nodes. If it is set to one, then Control Nodes are allowed to modify the Discovery Domain configuration. If it is set to zero, then Control Nodes are not allowed to modify Discovery Domain configurations.

iSCSI Target, iSCSI Initiator, iFCP Target Port, iFCP Initiator Port: Determine whether the respective registered iSNS client (determined by iSCSI Node Type or iFCP Port Role) is allowed to add, delete, or modify Discovery Domains. If they are set to one, then modification by the specified client type is allowed. If they are set to zero, then modification by the specified client type is not allowed.

(A node may implement multiple node types.)

2.3. Administrative Flags Field

The format of the Administrative Flags bit field is shown in Figure 4.

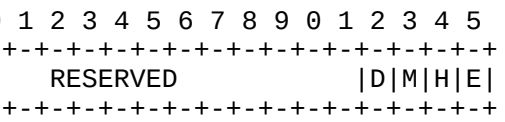

Figure 4. Administrative Flags

Bit Field	Significance
15	Enabled
14	Heartbeat
13	Management SCNs
12	Default Discovery Domain

The following are Administrative Flags Field definitions:

Enabled: Specifies the validity of the remainder of the Administrative Flags field. If it is set to one, then the contents of the remaining Administrative Flags are valid. If it is set to zero, then the remaining contents MUST be ignored, indicating that iSNS administrative settings are obtained through means other than DHCP.

Heartbeat: Indicates whether the first IP address is the multicast address to which the iSNS heartbeat message is sent. If it is set to one, then a1-a4 contains the heartbeat multicast address and b1-b4 contains the IP address of the

primary iSNS server, followed by the IP address(es) of any backup servers (see Figure 1). If it is set to zero, then a1-a4 contain the IP address of the primary iSNS server, followed by the IP address(es) of any backup servers.

Management SCNs: Indicates whether control nodes are authorized to register for receiving Management State Change Notifications (SCNs). Management SCNs are a special class of State Change Notification whose scope is the entire iSNS database. If this bit is set to one, then control nodes are authorized to register for receiving Management SCNs. If it is set to zero, then control nodes are not authorized to receive Management SCNs (although they may receive normal SCNs).

Default Discovery Domain: Indicates whether a newly registered device that is not explicitly placed into a Discovery Domain (DD) and Discovery Domain Set (DDS) should be automatically placed into a default DD and DDS. If it is set to one, then a default DD shall contain all devices in the iSNS database that have not been explicitly placed into a DD by an iSNS client. If it is set to zero, then devices not explicitly placed into a DD are not members of any DD.

2.4. iSNS Server Security Bitmap

The format of the iSNS server security Bitmap field is shown in Figure 5. If valid, this field communicates to the DHCP client the security settings that are required to communicate with the indicated iSNS server.

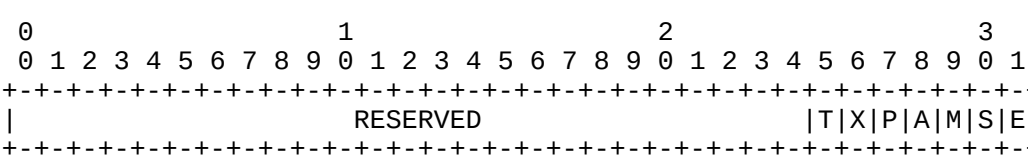

Figure 5. iSNS Server Security Bitmap

Bit Field	Significance
31	Enabled
30	IKE/IPSec
29	Main Mode
28	Aggressive Mode
27	PFS
26	Transport Mode
25	Tunnel Mode

The following are iSNS Server Security Bitmap definitions:

Enabled: Specifies the validity of the remainder of the iSNS server security bitmap. If it is set to one, then the contents of the remainder of the field are valid. If it is set to zero, then the contents of the rest of the field are undefined and MUST be ignored.

IKE/IPSec: 1 = IKE/IPSec enabled; 0 = IKE/IPSec disabled.

Main Mode: 1 = Main Mode enabled; 0 = Main Mode disabled.

Aggressive Mode: 1 = Aggressive Mode enabled;
0 = Aggressive Mode disabled.

PFS: 1 = PFS enabled; 0 = PFS disabled.

Transport Mode: 1 = Transport Mode preferred; 0 = No preference.

Tunnel Mode: 1 = Tunnel Mode preferred; 0 = No preference.

If IKE/IPSec is disabled, this indicates that the Internet Key Exchange (IKE) Protocol is not available to configure IPSec keys for iSNS sessions to this iSNS server. It does not necessarily preclude other key exchange methods (e.g., manual keying) from establishing an IPSec security association for the iSNS session.

If IKE/IPsec is enabled, then for each of the bit pairs <Main Mode, Aggressive Mode> and <Transport Mode, Tunnel Mode>, one of the two bits MUST be set to 1, and the other MUST be set to 0.

3. Security Considerations

For protecting the iSNS option, the DHCP Authentication security option as specified in [RFC3118] may present a problem due to the limited implementation and deployment of the DHCP authentication

option. The IPsec security mechanisms for iSNS itself are specified in [iSNS] to provide confidentiality when sensitive information is distributed via iSNS. See the Security Considerations section of [iSNS] for details and specific requirements for implementation of IPsec.

In addition, [iSNS] describes an authentication block that provides message integrity for multicast or broadcast iSNS messages (i.e., for heartbeat/discovery messages only). See [RFC3723] for further discussion of security for these protocols.

If no sensitive information, as described in [iSNS], is being distributed via iSNS, and an Entity is discovered via iSNS, authentication and authorization are handled by the IP Storage protocols whose endpoints are discovered via iSNS; specifically, iFCP [iFCP] and iSCSI [RFC3720]. It is the responsibility of the providers of these services to ensure that an inappropriately advertised or discovered service does not compromise their security.

When no DHCP security is used, there is a risk of distribution of false discovery information (e.g., via the iSNS DHCP option identifying a false iSNS server that distributes the false discovery information). The primary countermeasure for this risk is authentication by the IP storage protocols discovered through iSNS. When this risk is a significant concern, IPsec SAs SHOULD be used (as specified in RFC 3723). For example, if an attacker uses DHCP and iSNS to distribute discovery information that falsely identifies an iSCSI endpoint, that endpoint will lack the credentials necessary to complete IKE authentication successfully, and therefore will be prevented from falsely sending or receiving iSCSI traffic. When this risk of false discovery information is a significant concern and IPsec is implemented for iSNS, IPsec SAs SHOULD also be used for iSNS traffic to prevent use of a false iSNS server; this is more robust than relying only on the IP Storage protocols to detect false discovery information.

When IPsec is implemented for iSNS, there is a risk of a denial-of-service attack based on repeated use of false discovery information that will cause initiation of IKE negotiation. The countermeasures for this are administrative configuration of each iSNS Entity to limit the peers it is willing to communicate with (i.e., by IP address range and/or DNS domain), and maintenance of a negative authentication cache to avoid repeatedly contacting an iSNS Entity that fails to authenticate. These three measures (i.e., IP address range limits, DNS domain limits, negative authentication cache) MUST be implemented for iSNS entities when this DHCP option is used. An analogous argument applies to the IP storage protocols that can be discovered via iSNS as discussed in RFC 3723.

In addition, use of the techniques described in [RFC2827] and [RFC3833] may also be relevant to reduce denial-of-service attacks.

4. IANA Considerations

In accordance with the policy defined in [DHCP], IANA has assigned a value of 83 for this option.

There are no other IANA-assigned values defined by this specification.

5. Normative References

[DHCP] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[iSNS] Tseng, J., Gibbons, K., Travostino, F., Du Laney, C., and J. Souza, "Internet Storage Name Service (iSNS)", RFC 4171, September 2005.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

[RFC3720] Satran, J., Meth, K., Sapuntzakis, C., Chadalapaka, M., and E. Zeidner, "Internet Small Computer Systems Interface (iSCSI)", RFC 3720, April 2004.

[RFC3723] Aboba, B., Tseng, J., Walker, J., Rangan, V., and F. Travostino, "Securing Block Storage Protocols over IP", RFC 3723, April 2004.

6. Informative References

[iFCP] Monia, C., Mullendore, R., Travostino, F., Jeong, W., and M. Edwards, "iFCP - A Protocol for Internet Fibre Channel Storage Networking", RFC 4172, September 2005.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", RFC 3833, August 2004.

Authors' Addresses

Kevin Gibbons
McDATA Corporation
4555 Great America Parkway
Santa Clara, CA 95054-1208

Phone: (408) 567-5765
EMail: kevin.gibbons@mcdata.com

Charles Monia
7553 Morevern Circle
San Jose, CA 95135

EMail: charles_monia@yahoo.com

Josh Tseng
Riverbed Technology
501 2nd Street, Suite 410
San Francisco, CA 94107

Phone: (650)274-2109
EMail: joshtseng@yahoo.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.