

DHC Working Group
INTERNET DRAFT
Expires: March 2004
Internet Draft
Document: <draft-ietf-dhc-isnsoption-09.txt>
Category: Standards Track

Charles Monia
Josh Tseng
Kevin Gibbons
Nishan Systems
September 2003

The IPv4 DHCP Option for the Internet Storage Name Service

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of [RFC2026].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Comments

Comments should be sent to the DHCP mailing list (dhcwg@ietf.org) or to the authors.

Table of Contents

Status of this Memo.....	1
Comments.....	1
Abstract.....	3
Conventions used in this document.....	3
1. Introduction	3
2. iSNS Option for DHCP	4
2.1 iSNS Functions Field	5
2.2 Discovery Domain Access Field	7
2.3 Administrative Flags Field	8
2.4 iSNS Server Security Bitmap	9
3. Security Considerations	10
4. IANA Considerations	11
5. Normative References	11
6. Non-Normative References	11
7. Author's Addresses	12
Full Copyright Statement.....	13

Abstract

This document describes the DHCP option to allow Internet Storage Name Service (iSNS) clients to automatically discover the location of the iSNS server 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 as a storage area network.

Conventions used in this document

iSNS refers to the Internet Storage Name Service framework consisting 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.

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 needs 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 <<TBD>>.

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.

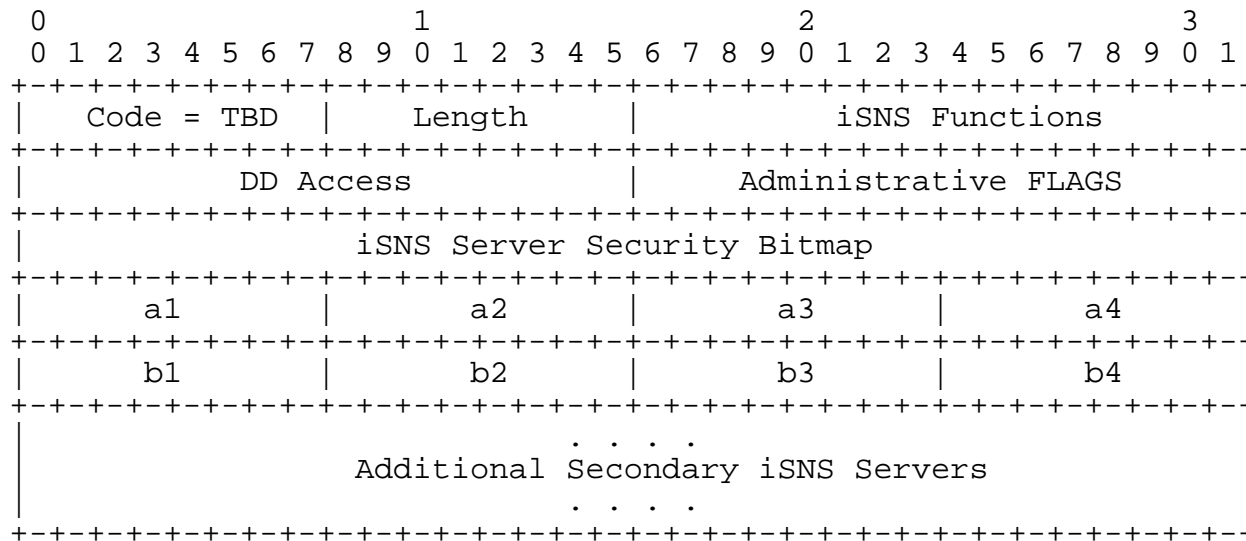


Figure 1 -- iSNS Server Option

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 a secondary iSNS server.

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

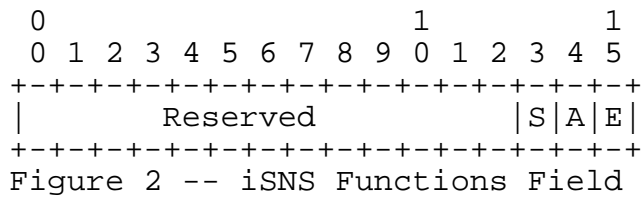
The Code field through IP address field a1...a4 MUST be present in every response to the iSNS query, hence the Length field has a minimum value of 14.

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:



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

iSNS Functions Field definitions:

Function Fields Enabled:	This bit specifies the validity of the remaining iSNS Function fields. If set to one, then the contents of all other iSNS Function fields are valid. If set to zero, then the contents of all other iSNS Function fields MUST be ignored.
DD-based Authorization:	Indicates whether or not 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 or not 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 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 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:

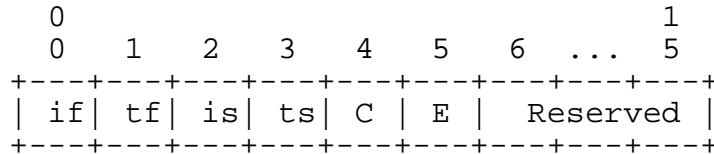


Figure 3 -- Discovery Domain Access Field

Bit field	Significance
5	Enabled
4	Control Node
3	iSCSI Target
2	iSCSI Initiator
1	iFCP Target Port
0	iFCP Initiator Port

Discovery Domain Access Field Definitions:

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

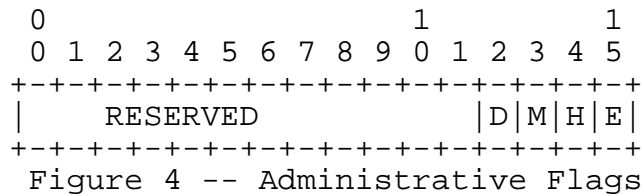
Control Node: Specifies whether the iSNS server allows Discovery Domains to be added, modified or deleted by means of Control Nodes. If set to one, then Control Nodes are allowed to modify the Discovery Domain configuration. If set to zero, then Control Nodes are not allowed to modify Discovery Domain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

iSCSI Target, iSCSI Initiator, iFCP Target Port, iFCP Initiator Port: These bits 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 set to one, then modification by the specified client type is allowed. If 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:



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

Administrative Flags Field definitions:

- Enabled:** Specifies the validity of the remainder of the Administrative Flags field. If set to one, then the contents of the remaining Administrative Flags are valid. If 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 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 set to zero, then a1-a4 contains 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 to receive Management State Change Notifications (SCN's). Management SCN's are a special class of State Change Notification whose scope is the entire iSNS database. If set to one, then control nodes are authorized to register to receive Management SCN's. If set to zero, then

control nodes are not authorized to receive Management SCN's (although they may receive normal SCN's).

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 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 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.

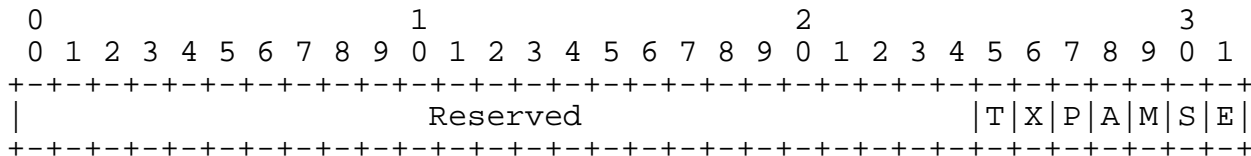


Figure 5 -- iSNS Server Security Bitmap

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

iSNS Server Security Bitmap definitions:

Enabled	This bit specifies the validity of the remainder of the iSNS server security bitmap. If set to one, then the contents of the remainder of the field are valid. If 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, an implementation SHALL enable:

a) One of Main Mode or Aggressive Mode but not both and

b) One of Transport Mode or Tunnel Mode but not both.

3. Security Considerations

[RFC3118] should be consulted to determine the requirements for additional security measures to authenticate the iSNS option message received by the DHCP client. ~~If necessary, the authentication option described in [RFC3118] should be utilized.~~

With regard to security considerations specific to the use of this DHCP option for iSNS server discovery, exposure to a "man-in-the-middle" attack by a hostile entity modifying or replacing the original iSNS option message should be considered a potential security exposure. If the authentication option in [RFC3118] is not implemented, then an attacker may trick the iSNS client into connecting into rogue iSNS servers.

Therefore, if [RFC3118] is implemented in the client and server, it MUST be used. If the authentication option for DHCP is not implemented and it is determined that the potential exists for a "man-in-the-middle" attack, then the DHCP option message for iSNS should not be utilized.

iSNS security considerations are discussed in [iSNS] and [SEC-IPS].

4. IANA Considerations

In accordance with the policy defined in [DHCP], IANA has assigned a value of TBD 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, Bucknell University, March 1997.

[iSNS] Tseng, J. et al., "iSNS - Internet Storage Name Service", Internet draft (work in progress), draft-ietf-ips-isns-12.txt, August 2002

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996

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

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

6. Non-Normative References

[iFCP] Monia, C., et al., "iFCP - A Protocol for Internet Fibre Channel Storage Networking", Internet draft (work in progress), draft-ietf-ips-ifcp-13.txt, May 2002

[iSCSI] Satran, J., et al., "iSCSI", Internet draft (work in progress), draft-ietf-ips-iSCSI-15.txt, August 2002

[SEC-IPS] Aboba, B., et al., "Securing IP Block Storage Protocols", draft-ietf-ips-security-14.txt, June 2002

7. Author's Addresses

Kevin Gibbons,
Charles Monia,
Josh Tseng

Nishan Systems
3850 North First Street
San Jose, CA 95134-1702
Phone: (408) 519-3700
Email: cmonia@nishansystems.com
jtseng@nishansystems.com
kgibbons@nishansystems.com

Full Copyright Statement

"Copyright (C) The Internet Society September 2003. All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

