IP Storage Working Group
Internet Draft
Document: draft-ietf-ips-iscsi-boot-10.txt
Category: Standards Track

Prasenjit Sarkar
IBM
Duncan Missimer
Brocade
Constantin Sapuntzakis
Stanford University
19 June 2003

Bootstrapping Clients using the iSCSI Protocol

Status of this Memo

Abstract

iSCSI is a proposed transport protocol for SCSI that operates on top of TCP. This memo describes a standard mechanism to enable clients to bootstrap themselves using the iSCSI protocol. The goal of this standard is to enable iSCSI boot clients to obtain the information to open an iSCSI session with the iSCSI boot server.

1. Introduction

The Small Computer Systems Interface (SCSI) is a popular family of protocols for communicating with I/O devices, especially storage devices. SCSI can be characterized as a request/response messaging protocol with a standard architecture and componentized command sets for different device classes.

iSCSI is a proposed transport protocol for SCSI that operates on top of TCP. The role of iSCSI is necessitated by the evolution of the system interconnect from a shared bus to a switched network. IP networks meet the architectural and performance requirements of transporting SCSI, paving the way for the iSCSI protocol.

Many diskless clients sometimes bootstrap off remote SCSI devices. Such diskless entities are lightweight, space-efficient and power-conserving, and are increasingly popular in various environments.

This memo describes a standard mechanism to enable clients to bootstrap themselves using the iSCSI protocol. The goal of this standard is to enable iSCSI boot clients to obtain the information to open an iSCSI session with the iSCSI boot server. It is possible that all the information is not available at the very outset, so the memo describes steps to obtain the information required to bootstrap clients off an iSCSI boot server.

2. Requirements

1. There must be no restriction of network topology between the iSCSI boot client and the boot server other than those in effect for establishing the iSCSI session. Consequently, it is possible for an iSCSI boot client to boot from an iSCSI boot server behind gateways or firewalls as long as it is possible to establish an iSCSI session between the client and the server.

2. The following represents the minimum information required for an iSCSI boot client to contact an iSCSI boot server: (a) the client's IP address (IPv6 or IPv4); (b) the server's iSCSI Target Name; and (c) mandatory iSCSI initiator capability.

The above assumes that the default LUN for the boot process is 0 and the default port for the iSCSI boot server is the well-known iSCSI port. However, both may be overridden at the time of configuration.

Additional information may be required at each stage of the boot process.

3. It is possible for the iSCSI boot client to have none of the above information or capability on starting.

4. The client should be able to complete boot without user intervention (for boots that occur during an unattended power-up). However, there should be a mechanism for the user to input values so as to bypass stages of the boot protocol.

5. Additional protocol software (for example, DHCP) may be necessary

if the minimum information required for an iSCSI session is not provided.

3. Related Work

The Reverse Address Resolution Protocol (RARP) [Finlayson84] through the extensions defined in the Dynamic RARP (DRARP)) [Brownell96] explicitly addresses the problem of network address discovery, and includes an automatic IP address assignment mechanism. The Trivial File Transfer Protocol (TFTP) [Sollins81] provides for transport of a boot image from a boot server. BOOTP [Croft85, Reynolds93, Wimer93] is a transport mechanism for a collection of configuration information. BOOTP is also extensible, and official extensions have been defined for several configuration parameters. DHCPv4 [Droms97, Droms93] and DHCPv6 [Droms02] are standards for hosts to be dynamically configured in an IP network. The Service Location Protocol (SLP) provides for location of higher level services [Guttman99].

4. Software stage

Some iSCSI boot clients may lack the resources to boot up with the mandatory iSCSI initiator capability. Such boot clients may choose to obtain iSCSI initiator software from a boot server. Currently, there are many established protocols that allow such a service to enable clients to load software images. For example, BOOTP and DHCP servers have the capability to provide software images on requests from boot clients.

It is to be noted that this document does not recommend any of the above protocols, and the final decision of which boot protocol is to be used to load iSCSI initiator software is left to the discretion of the implementor.

5. DHCP stage

In order to use an iSCSI boot server, the following pieces of information are required for an ISCSI boot client.

- The IP address of the iSCSI boot client (IPv4 or IPv6)

- The IP transport endpoint for the iSCSI Target Port for the iSCSI boot server. If the transport is TCP, for example, this has to resolve to an IP address and a TCP port number. TCP is currently the only transport approved for iSCSI.

- The eight-byte LUN structure identifying the Logical Unit within

the iSCSI boot server.

At boot time, all or none of this information may be stored in the iSCSI boot client. This section describes techniques for obtaining the required information via the DHCP stage. Otherwise, if the iSCSI boot client has all the information, the boot client may proceed directly to the Boot stage.

An iSCSI boot client which does not know its IP address at power-on may acquire its IP address via DHCP. An iSCSI boot client which is capable of using both DHCPv6 and DHCPv4 should first attempt to use DHCPv6 to obtain its IP address, falling back on DHCPv4 in the event of failure.

Unless otherwise specified here, DHCP fields such as the client ID and gateway information are used in an identical way as applications other than iSCSI do.

A DHCP server (v4 or v6) MAY instruct an iSCSI client how to reach its boot device. This is done using the variable length DHCP option named Root Path. The use of the option field is reserved for iSCSI boot use by prefacing the string with "iscsi:".

The option field consists of an UTF-8 [Yergeau98] string. The string MUST contain only alphanumberic characters, "." , ":" and "-"; no other characters are permissible. The string has the following composition:

"iscsi:"<servername>":"<protocol>":"<port>":"<LUN>":"<targetname>

The fields "servername", "port", "protocol" and "LUN" are OPTIONAL and should be left blank if there are no corresponding values. The "targetname" field is not optional and MUST be provided.

The "servername" is the name of iSCSI server and contains either a valid domain name, a literal IPv4 address, or a literal IPv6 address.

If the "servername" field contains a literal IPv4 address, the IPv4 address MUST be in standard dotted decimal notation as defined in Section 2.1 of RFC 1123 [Braden89].

If the "servername" field contains an IPv6 address, the address MUST be represented in the IPv6 address format x.x.x.x.x.x.x.x where the 'x's are the hexadecimal values of the eight 16-bit pieces of the address. Note that this format representation is specific to iSCSI boot.

If the "servername" is a domain name, the name MUST be a fully

qualified domain name (FQDN) and should abide by the rules specified in Sections 3.1 and 3.5 of RFC 1034 [Mockaopetris87] and the reply from the host configuration server should contain the Domain Name Server Option [Alexander93]. It must also be pointed out that the use of DNS for address translation in enterprise environments must contain adequate levels of fault tolerance and security.

If the "servername" field contains 4 decimal components, the "servername" is assumed to be an IPv4 address. If there are more than 4 decimal components or if there is a hexadecimal component, the the "servername" is assumed to be an IPv6 address. If the least significant (rightmost) component is an approved domain extension, then the "servername" field is assumed to be a domain name. If the "servername" field is left blank, then no default value is assumed in its place.

The "protocol" field is the decimal representation of the IANA-approved string for the trasport protocol to be used for iSCSI. If the protocol field is left bank, the default value is assumed to be "6" for TCP. The transport protocol MUST have been approved for use in iSCSI; currently, the only approved protocol is TCP.

The "port" is the decimal representation of the port on which the iSCSI boot server is listening. If not specified, the port defaults to the well-known iSCSI port.

The "LUN" field is a hexadecimal representation of the LU number. If the LUN field is blank, then LUN 0 is assumed. If the LUN field is not blank, the representation MUST be divided into four groups of four hexadecimal digits, separated by "-". Digits above 9 may be either lower or upper case. An example of such a representation would be 4752-3A4F-6b7e-2F99. For the sake of brevity, at most three leading zero ("0") digits MAY be omitted in any group of hexadecimal digits. Thus, the "LUN" representation 6734-9-156f-127 is equivalent to 6734-0009-156f-0127. Furthermore, trailing groups containing only the "0" digit MAY be omitted along with the preceding "-". So, the "LUN" representation 4186-9 is equivalent to 4186-0009-0000-0000. Other concise representations of the LUN field MUST NOT be used.

Note that SCSI targets are allowed to present different LU numberings for different SCSI initiators, so that to our knowledge nothing precludes a SCSI target from exporting several different LUs to several different SCSI initiators as their respective LUN 0s.

The "targetname" field is an iSCSI Name that is defined by the iSCSI standard [Satran02] to uniquely identify an iSCSI target.

If the "servername" field is provided by DHCP, then that field is

used in conjunction with other associated fields to contact the boot server in the Boot stage (Section 7). However, if the "servername" field is not provided, then the "targetname" field is then used in the Discovery Service stage in conjunction with other associated fields. (Section 6).

6. Discovery Service stage

This stage is required if the DHCP server (v4 or v6) is unaware of any iSCSI boot servers or if the DHCP server is unable to provide the minimum information required to connect to the iSCSI boot server other than the targetname.

The Discovery Service may be based on the SLP protocol [Guttman99, Bakke02] and is an instantiation of the SLP Service or Directory Agent. Alternatively, the Discovery Service may be based on the iSNS protocol [Tseng03] and is an instantiation of the iSNS Server.

The iSCSI boot client may have obtained the targetname of the iSCSI boot server in the DHCP stage (Section 5). In that case, the iSCSI boot client queries the SLP Discovery Service using query string 1 of the iSCSI Target Concrete Service Type Template as specified in Section 6.2 of the iSCSI SLP interaction document [Bakke02] to resolve the targetname to an IP address and port number. Alternatively, the iSCSI boot client may query the iSNS Discovery Service with a Device Attribute Query with the targetname as the query parameter [Tseng03]. Once this is obtained, the iSCSI boot client proceeds to the Boot stage (Section 7).

It is possible that the port number obtained from the Discovery Service may conflict with the one obtained from the DHCP service. In such a case, the implementor has the option to try both port numbers in the Boot stage.

If the iSCSI boot client does not have any targetname information, the iSCSI boot client then may query the SLP Discovery Service with query string 4 of the iSCSI Target Concrete Service Type Template as specified in Section 6.2 of the iSCSI SLP interaction document [Bakke02]. In response to this query, the SLP Discovery Service provides the boot client with a list of iSCSI boot servers the boot client is allowed to access. Alternatively, the iSCSI boot client can query the iSNS Discovery Service to verify if the targets in particular Discovery Domain are bootable [Tseng03].

If the list of iSCSI boot servers is empty, subsequent actions are left to the discretion of the implementor. Otherwise, the iSCSI boot client may contact any iSCSI boot server in the list. Moreover, the

order in which iSCSI boot servers are contacted is also left to the discretion of the implementor.

7. Boot stage

Once the iSCSI boot client has obtained the minimum information to open an iSCSI session with the iSCSI boot server, the actual booting process can start.

The actual sequence of iSCSI commands needed to complete the boot process is left to the implementor. This was done because of varying requirements from different vendors and equipment, making it difficult to specify a common subset of the iSCSI standard that would be acceptable to everybody.

The iSCSI session established for boot may be taken over by the booted software in the iSCSI boot client.

8. Security Considerations

The security discussion is centered around securing the communication involved in the iSCSI boot process.

However, the issue of applying credentials to a boot image loaded through the iSCSI boot mechanism is outside the scope of this document. One key difference between the iSCSI boot mechanism and BOOTP-based image loading is the fact that the identity of a boot image may not be known when the Boot stage starts. The identity of certain boot images and their locations are known only after examining the contents of a boot disk exposed by the iSCSI boot service. Furthermore, images themselves may recursively load other images based on both hardware configurations and user input. Consequently, a practical way to verify loaded boot images is to make sure that each image loading software verifies the image to be loaded using a mechanism of their choice.

The considerations involved in designing a security architecture for the iSCSI boot process include configuration, deployment and provisioning issues apart from typical security considerations. Enabling iSCSI boot creates a critical operational dependence on an external system with obvious security implications, and thus administrator awareness of such enablement is extremely important. Therefore, iSCSI boot SHOULD NOT be enabled, or put high in the boot order, without an explicit administrative action.

The software stage SHOULD NOT be involved in a secure iSCSI boot

process as this would add the additional complexity of trying to secure the process of loading the software necessary to run the later stages of iSCSI boot. Authentication and integrity protection of downloaded boot software has proven to be difficult and complex due to administrative issues and limitations of the BIOS environment. It is therefore assumed that all the necessary software is resident on the iSCSI boot client.

If the DHCP stage is implemented, the iSCSI boot client SHOULD implement the DHCP authentication ([Droms01], [Droms02] for IPv6). In this case an administration interface SHOULD be provided for the configuration of the DHCP authentication credentials, both when the network interface is on the motherboard, and when it is removable. Note that DHCP authentication ([Droms01],[Droms02] for IPv6) is focused on intra-domain authentication, which is assumed to be enough for iSCSI boot scenarios. In the context of the secure iSCSI boot process, the reply from the DHCP server in the DHCP stage SHOULD include the serverName in IPv4 (or IPv6) format to avoid reliance on a DNS server (for resolving names) or a Discovery Service entity (to look up targetnames). This reduces the number of entities involved in the secure iSCSI boot process.

If the Discovery Service stage is implemented using SLP, the iSCSI boot client SHOULD provide IPsec support (OPTIONAL to use) for the SLP protocol, as defined in [Bakke02] and [Aboba03]. If the Discovery Service stage is implemented using iSNS, the iSCSI boot client SHOULD provide IPsec support (OPTIONAL to use) for the iSNS protocol, as defined in [Tseng03] and [Aboba03]. When iSNS or SLP are used to distribute security policy or configuration information, at a minimum, per-packet data origin authentication, integrity and replay protection SHOULD be used to protect the discovery protocol.

For the final communication between the iSCSI boot client and the iSCSI boot server in the Boot stage, IPsec and in-band authentication SHOULD be implemented according to the guidelines in the main iSCSI draft [Satran02] and [Aboba03]. Due to memory constraints, it is expected that iSCSI boot clients will only support the pre-shared key authentication in IKE. Where the host IP address is assigned dynamically, IKE main mode SHOULD NOT be used, as explained in [Satran02] and [Aboba03]. Regardless of the way parameters in previous stages (DHCP, SLP, iSNS) were obtained (securely or not), the iSCSI boot session is vulnerable as any iSCSI session (see [Satran02] and [Aboba03] for iSCSI security threats). Therefore security for this session SHOULD be configured and used according to [Satran02] and [Aboba03] guidelines.

Another point to be noted is that if a boot image inherits an iSCSI session from a previously loaded boot image, the boot image also

inherits the security properties of the iSCSI session.

Acknowledgments

We wish to thank John Hufferd for taking the initiative to form the iSCSI boot team. We also wish to thank Doug Otis, Julian Satran, Bernard Aboba, David Robinson, Mark Bakke, Ofer Biran and Mallikarjun Chadalapaka for helpful suggestions and pointers regarding the draft document.

Normative References

[Aboba03] Aboba, S. et al. "Securing Block Storage Protocols over IP", Work in Progress, January 2003.

[Alexander93] Alexander, S., and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, Lachman Technology, Inc., Bucknell University, October 1993.

[Bakke02] Bakke, M., et al. "Finding iSCSI Targets and Name Servers using SLP", Work in Progress, March 2002.

[Braden89] Braden, R., "Requirements for Internet Hosts - Application and Support", RFC 1123, October 1989.

[Bradner96] Bradner, S., "The Internet Standards Process -- Revision 3", RFC 2026, October 1996.

[Bradner97] Bradner, S. "Key Words for use in RFCs to indicate Requirement Levels", RFC 2119, Harvard University, March 1997.

[Croft85] Croft, B., and J. Gilmore, "Bootstrap Protocol (BOOTP)", RFC 951, Stanford and SUN Microsystems, September 1985.

[Droms93] Droms, D., "Interoperation between DHCP and BOOTP" RFC 1534, Bucknell University, October 1993.

[Droms97] R. Droms, "Dynamic Host Configuration Protocol", RFC 2131, Bucknell University, March 1997.

[Droms01] Droms, R., Arbaugh, W., "Authentication for DHCP Messages", RFC 3118, June 2001.

[Droms02] Droma, R., Bound, J., Volz, B., Lemon, T., Perkins, C.,

Carney, M., "Dynamic Host Configuration Protocol for IPv6", Work in Progress, November 2002.

[Guttman99] Guttman, E., Perkins, C., Verizades, J., Day, M., "Service Location Protocol v2", RFC 2608, June 1999.

[Mockaopetris87] Mockaopertis, P., "Domain Names - Concepts and Facilities", RFC 1034, November 1987.

[Reynolds93] Reynolds, J., "BOOTP Vendor Information Extensions", RFC 1497,
USC/Information Sciences Institute, August 1993.

[Satran02] Satran, J. et al., "iSCSI", Work in Progress, September 2002.

[Tseng03] Tseng, J., Gibbons, K., Travostino, F., Du Laney, C., Souza, J., "Internet Storage Name Service", Work in Progress, June 2003.

[Yergeau98] Yergeau, F., "UTF-8: A Transformation Format for ISO-10646", RFC 2279, January 1998.

[Wimer93] Wimer, W., "Clarifications and Extensions for the Bootstrap Protocol", RFC 1532, Carnegie Mellon University, October 1993.

Informative References

[Brownell96] Brownell, D, "Dynamic RARP extensions for Automatic Network Adress
Acquisition", RFC 1931, SUN Microsystems, April 1996.

[Finlayson84] Finlayson, R., Mann, T., Mogul, J., and M. Theimer, "A Reverse
Address Resolution Protocol", RFC 903, Stanford, June 1984.

[Sollins81] Sollins, K., "The TFTP Protocol (Revision 2)", RFC 783, NIC,
June 1981.

Authors' Addresses

Prasenjit Sarkar
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120, USA
Phone: +1 408 927 1417

Email: psarkar@almaden.ibm.com

Duncan Missimer
Brocade Communication Systems
1745 Technology Drive,
San Jose, CA 95110, USA
Email: dmissime@brocade.com

Constantine Sapuntzakis
Stanford University
353 Serra Hall #406
Stanford, CA 94306, USA
Phone: +1 650 520 0205
Email: csapuntz@stanford.edu

Full Copyright Statement

Intellectual Property Rights

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology

described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat."

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.