Internet Draft
<draft-ietf-ips-iscsi-slp-00>
Expires October 2001

Mark Bakke
Cisco

Joe Czap
IBM

Jim Hafner
IBM

Howard Hall
Pirus

Jack Harwood
EMC

John Hufferd
IBM

Yaron Klein
Sanrad

Lawrence Lamers
San Valley Systems

Todd Sperry
Adaptec

Joshua Tseng
Nishan

Kaladhar Voruganti
IBM

April 2001

Finding iSCSI Targets and Name Servers Using SLP

Status of this Memo

Copyright Notice

Abstract

The iSCSI protocol provides a way for hosts to access SCSI devices over an IP network. This document defines the use of the Service Location Protocol (SLP) by iSCSI hosts, devices, and name services, along with the SLP service type templates that describe the services they provide.

1. Acknowledgements

This draft was produced as a companion document for the iSCSI Naming and Discovery team, including Joe Czap, Jim Hafner, John Hufferd, and Kaladhar Voruganti (IBM), Howard Hall (Pirus), Jack Hardwood (EMC), Yaron Klein (Sanrad), Lawrence Lamers (San Valley), Todd Sperry (Adaptec), and Joshua Tseng (Nishan). Thanks also to Julian Satran (IBM) for suggesting the use of SLP for iSCSI discovery, and to Matt Peterson (Caldera) and James Kempf (Sun) for reviewing the document from an SLP perspective.

2. Introduction

iSCSI [iSCSI] is a protocol used to transport SCSI [SAM2] commands, data, and status across an IP network. This protocol is connection-oriented, and is currently defined over TCP. iSCSI uses a client-server relationship. The client end of the connection is an initiator, and sends SCSI commands; the server end of the connection is called a target, and receives and executes the commands.

There are several methods an iSCSI initiator can use to find the

targets to which it should connect. The discovery of iSCSI targets can be configured on the initiator in several ways:

- Each target and its address can be statically configured on the initiator.

- Each address providing targets can be configured on the initiator; the initiator can query the address for a list of targets.

- A storage name server address can be configured on the initiator; the initiator can use the storage name servers's protocol it provides to obtain a list of targets.

These methods are further defined in "iSCSI Naming and Discovery Requirements" [NDT].

Each of the above methods requires a small amount of configuration to be done on each initiator. The ability to discover targets and name services without having to configure initiators is a desirable feature. The Service Location Protocol (SLP) [SLP] is an IETF standards track protocol that provides several features that will simplify locating iSCSI services. This document describes how SLP can be used in iSCSI environments to discover targets, addresses providing targets, and storage name servers.

3. Notation Conventions

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].

4. Terminology

Here are some definitions that may aid readers that are unfamiliar with either SLP, SCSI, or iSCSI. Some of these definitions have been reproduced from [RFC2608] and "Finding an RSIP Server with SLP" [RSIP].

User Agent (UA)   A process working on the client's behalf to establish contact with some service. The UA retrieves service information from the Service Agents or Directory Agents.

Service Agent (SA)   A process working on behalf of one or more services to advertise the services and their

capabilites.

Directory Agent (DA) A process which collects service advertisements. There can only be one DA present per given host.

Scope A named set of services, typically making up a logical administrative group.

Service Advertisement A URL, attributes, and a lifetime (indicating how long the advertisement is valid), providing service access information and capabilities description for a particular service.

Initiator A logical entity, typically within a host, that sends SCSI commands to targets to be executed. An initiator is usually present in the form of a device driver.

Target A logical entity, typically within a storage controller or gateway, that receives SCSI commands from an initiator and executes them. A target includes one or more Logical Units (LUs); each LU is a SCSI device, such as a disk or tape drive.

iSCSI Name A UTF-8 character string which serves as a unique identifier for iSCSI initiators and targets. Its format and usage is further defined in [NDT].

iSCSI Client A logical entity, typically a host, which includes at least one iSCSI Initiator.

iSCSI Server A logical entity, typically a storage controller or gateway, which includes at least one iSCSI Target.

Storage Name Server An addressible entity that provides one of several discovery and management services that benefit an iSCSI environment. Storage Name Server is used as a generic term, and is not necessarily equivalent to that used within Fibre Channel.

5. Using SLP for iSCSI Service Discovery

At least two entities are involved in iSCSI discovery. The end result is that an iSCSI initiator (e.g. a host) discovers iSCSI targets, usually provided by storage controllers or gateways. An iSCSI initiator may either discover these targets directly, using SLP, or may opt to discover them through a storage name service. One name service protocol currently under development is iSNS [ISNS]. In the first case, only SLP is needed. In the second, SLP may be used to discover a name server, which can be used to provide additional capabilities beyond simple target discovery.

This section first describes the use of SLP for discovery of targets by iSCSI initiators, and then describes the use of SLP to discover storage name servers.

This document assumes that SLPv2 will be used when discovering iSCSI-related services; no attempt is made to include support for SLPv1.

5.1. Discovering iSCSI Targets using SLP

The following diagram shows the relationship between iSCSI clients, servers, initiators, and targets. An iSCSI client includes at least one iSCSI initiator, and an SLP user agent (UA). An iSCSI server includes at least one iSCSI target, and an SLP service agent (SA). Some entities, such as extended copy engines, include both initiators and targets. These include both an SA, for its targets to be discovered, and a UA, for its intiator(s) to discover other targets.

```
              +---------------------------------+
              |      iSCSI Client               |
              |                                 |
              |     +-----------+               |
              |     | iSCSI     |               |
              |     | initiator |               |
              |     +-----------+               |
              |                                 |
              +--------------------------+------+
              | iSCSI Driver             |  UA  |
              +--------------------------+------+
              |           TCP/UDP/IP            |
              +----------------+----------------+
              |  Interface 1   |   Interface 2  |
              +----------------+----------------+
                       |               |
     +------------+    |               |    +------------+
     |   SLP DA   |    |               |    |  SLP DA    |
     |            |----+  IP Networks  +----|            |
     | (optional) |    |               |    | (optional) |
     +------------+    |               |    +------------+
                       |               |
              +-----------------+-----------------|
              |   Interface 1   |   Interface 2   |
              |   10.1.30.21    |    10.1.40.3    |
              +-----------------+-----------------+
              |            TCP/UDP/IP             |
              +---------------------------+-------+
              |       iSCSI Driver        |  SA   |
              +---------------------------+-------|
              |                                   |
              | +--------+ +--------+ +---------+ |
              | | iSCSI  | | iSCSI  | |  iSCSI  | |
              | | target | | target | |  target | |
              | | "one"  | | "two"  | | "iscsi" | |
              | +--------+ +--------+ +---------+ |
              |                                   |
              |            iSCSI Server           |
              +-----------------------------------+
```

In the above drawing, the iSCSI server has three iSCSI targets that the client could discover. The first two, targets "one" and "two", are actual targets that can support SCSI commands. The third, "iscsi", is the canonical target. An initiator can log in to the canonical target, and use an iSCSI command called "SendTargets" to obtain a list of the other targets within this iSCSI server.

Each of the iSCSI targets has a unique name, called an iSCSI Node Name, or simply iSCSI Name. This identifier is the same regardless of the network path (through adapter cards, networks, interfaces on the storage device) over which the target is discovered and accessed. For this example, the iSCSI names "one" and "two" are used. A real iSCSI name incorporates more structure, including a naming authority, and is not described here.

Each of the iSCSI targets in the drawing can appear at two addresses, since two network interfaces are present. Each target, then, would have two URLs.

An iSCSI target URL consists of its fully qualified host name or IP address, the TCP port on which it is listening, and its world- wide unique identifier. If the server in this drawing is listening at TCP port 3000 for both network addresses, the service URLs present would be:

- 10.1.30.21:3000/one

- 10.1.30.21:3000/two

- 10.1.30.21:3000/iscsi

- 10.1.40.3:3000/one

- 10.1.40.3:3000/two

- 10.1.40.3:3000/iscsi

An iSCSI server has two options when registering targets with the service location protocol:

- Register the canonical target "iscsi" at each of its network addresses. The initiator can use this target to discover the others. This would create two service URLs in the example.

- The server can register the individual targets, "one" and "two", at each of its network addresses. This would create four service URLs in the example.

Although a target could do both types of registration, it would probably not be helpful.

The iSCSI server constructs a service advertisement of the type "service:iscsi:target" for each of the service URLs it wishes to register. The advertisement contains a lifetime, along with other attributes which are defined in the service template.

The remainder of the discovery procedure is identical to that used by any client/server pair implementing SLP:

1. If an SLP DA is found, the SA contacts the DA and registers the advertisement. If no DA is found, the SA maintains the advertisement itself, answering multicast UA queries directly.

2. When the iSCSI initiator requires contact information for an iSCSI target, the UA either contacts the DA using unicast or the SA using multicast. The UA includes a query based on the attributes to indicate the characteristics of the target(s) it requires.

3. Once the UA has the host name or address of the iSCSI server as well as the port number and iSCSI Target Name, it can begin the normal iSCSI login to the target.

5.2. Discovering Storage Name Services using SLP

Storage Name Servers can be built to perform discovery of targets in a variety of ways. They can also provide extended services beyond discovery, which could include storage allocation and management. None of these services are defined here; the intent of this document is to allow these services to be discovered by clients.

The following drawing shows an iSCSI client, an iSCSI server, and a storage name server. To simplify the drawing, the second IP network is not shown, but is assumed to exist. The storage name server would use its own protocol (snsp) to provide capabilities to iSCSI clients and servers; these clients and servers could both use SLP to discover the storage name server.

```
      +---------------------------+
      |         iSCSI Client      |
      |                           |
      |       +-----------+       |
      |       | iSCSI     |       |
      |       | initiator |       |
      |       +-----------+       |
      |                           |
      +---------------+------+----+      +------------+
      | iSCSI Driver   | snsp | UA |      |  SLP DA    |
      +---------------+------+----+      |            |
      |        TCP/UDP/IP         |      | (optional) |
      +---------------+------+----+      +------------+
               |                               |
               |   IP Network                  |
           ------------------------------------------
               |                          |
               |                          |
      +---------------+-----------+     +---------------------+
      |        TCP/UDP/IP         |     | TCP/UDP/IP          |
      +---------------+------+----+     +---------------------+
      | iSCSI Driver   | snsp | UA |     |   SA    |   snsp    |
      +---------------+------+----+     +---------------------+
      |                           |     |                     |
      | +--------+ +--------+     |     | storage name server |
      | | iSCSI  | | iSCSI  |     |     |                     |
      | | target | | target |     |     +---------------------+
      | |   1    | |   2    |     |
      | +--------+ +--------+     |
      |                           |
      |     iSCSI Server          |
      +---------------------------+
```

Note the difference between the storage name server model and the previously-defined target discovery model. When target discovery was used, the iSCSI Server implemented an SA, to be discovered by the initiator's UA. In the storage name server model, the iSCSI clients and servers both implement UAs, and the name server implements the SA. To be discovered by an initiator in this model, a target first find the storage name server, and then register itself with the name server using the name server's protocol.

A storage name server's URL contains the domain name or IP address and TCP port. No other information is required.

The iSCSI server constructs a service advertisement of the type "service:iscsi:target" for each of the addresses at which it appears. The advertisement contains the URL, a lifetime, along with other attributes which are defined in the service template.

The remainder of the discovery procedure is identical to that used to discover iSCSI targets, except that both initiators and targets would normally be "clients" of the storage name service.

5.3. Supporting Discovery of Targets and Name Services

An initiator should support, at a minimum, discovery of targets using SLP. If an initiator also supports a storage name service, the initiator should use SLP to discover the storage name service, and perform target discovery through the name service. An initiator may discover other targets directly even when using a name service, or may discover targets via multiple name services.

For example, the name service may be used to discover targets at a remote location, and SLP may used to discover local targets that are not part of the service.

Targets that support a particular storage name service can be configured to be discovered through that name service. A target configured in this way should not answer direct SLP discovery requests from initiators; they should effectively disable their service:iscsi:target SA functionality while under the control of a name service. This prevents targets from being duplicated between SLP and the name service. A target should also not be configured to support more than one name service at the same time.

5.4. Interoperating Between Initiators, Targets, and Name Services

Many initiators will support the basic discovery of targets using SLP, but will not always support a particular storage name service in use. A storage name server may provide the capability to support these initiators by forming service advertisements for use through its SA on behalf of the targets it has discovered. This would allow an initiator to discover targets in a storage name service environment, without having direct support for the storage name service protocol in use.

Targets not supporting the storage name service protocol in use can be discovered separately from the name service.

6. iSCSI SLP Templates

Three templates are provided: an iSCSI target template, a name service template, and an abstract template to encapsulate the two.

6.1. The iSCSI Abstract Service Type Template

This template defines the abstract service "service:iscsi". It is used as a top-level service to encapsulate all other iSCSI-related services.

Name of submitter: Mark Bakke
Language of service template: en
Security Considerations:
  See the security considerations of the concrete service types.

Template Text:

```
-------------------------template begins here-----------------------
template-type=iscsi

template-version=0.1

template-description=
  This is an abstract service type.  The purpose of the iscsi
  service type is to encompass all of the services used to support
  the iSCSI protocol.

template-url-syntax=
  url-path=  ;  Depends on the concrete service type.

--------------------------template ends here-----------------------
```

6.2. The iSCSI Target Concrete Service Type Template

This template defines the service "service:iscsi:target". An entity containing iSCSI targets that wishes them discovered via SLP would register each of them, with each of their addresses, as this service type.

Initiators and name services wishing to discover targets in this way will generally use one of the following query strings:

1. Find a specific target, given its iSCSI Target Name:

```
Service: service:iscsi:target
Scope:   initiator-or-name-server-scope-list
Query:   (iscsi-name=fqn.com.acme.sn.456)
```

2. Find all of the iSCSI Target Names that may allow access to a given initiator:

```
Service: service:iscsi:target
Scope:   initiator-or-name-server-scope-list
Query:   (access-list=fqn.com.os.hostid.045A7B)
```

3. In addition, a name service may wish to discover all targets, and assume responsibility for them. It may issue a simple query for all of the targets:

```
Service: service:iscsi:target
Scope:   name-server-scope-list
Query:   none
```

4. Find the iSCSI Target Names from which the given initiator is allowed to boot:

```
Service: service:iscsi:target
Scope:   name-server-scope-list
Query:   (boot-list=fqn.com.os.hostid.045A7B)
```

More details on booting from an iSCSI target are defined in [BOOT].

Name of submitter: Mark Bakke
Language of service template: en
Security Considerations:
  See later section.

Template Text:

```
   -------------------------template begins here-----------------------
   template-type=iscsi:target

   template-version=0.1

   template-description=
     This is concrete service type.  The iscsi:target service type is used
     to register individual target addresses to be discovered by others.
     UAs will generally search for these by including one of the following:
     - the iSCSI target name
     - the iSCSI initiator name (must be in the access-list of the target)
     - the service URL

   template-url-syntax=
     url-path   =  ipaddr [ : tcpport ] / iscsi-name
     ipaddr     =  DNS host name or ip address
     tcpport    =  decimal tcp port number
     iscsi-name =  iSCSI target name
     ; The iscsi-name part of the URL is required and may be either the iSCSI
     ; name of the target being registered, or the canonical name "iscsi". If
an
     ; initiator discovers a canonical iSCSI name, it should log in to that
     ; target, and issue the iSCSI SendTargets command to discover additional
     ; targets.  A device representing multiple targets may then either
     ; register each of them with SLP, or just register a single "iscsi"
     ; target, which will be used to discover the remainder of the
     ; targets.
     ;
     ; Examples:
     ;   service:iscsi:target://10.1.3.40:4000/fqn.com.acme.sn.45678
     ;   service:iscsi:target://mystorage.mycompany.com/iscsi

   iscsi-name = string
   # The iSCSI Name of this target.
   # This must match the iscsi-name in the url-path.

   transports = string M L
   tcp
     # This is a list of transport protocols that the registered
     # entity supports.  iSCSI is currently supported over TCP,
     # but it is anticipated that it could be supported over other
     # transports, such as SCTP, in the future.
   tcp

   entity = string O
   # Normally the FQDN of the management interface of the entity
   # containing this target.
   mgmt-ipaddr = string O
   # The IP address of the management interface appropriate for SNMP,
   # web-based, or telnet management of the entity containing this
   # target.

   alias = string O
   # The alias string contains a descriptive name of the target.

   access-list = string M
   # A list of iSCSI Initiator Names that can access this target.
   # Normal iSCSI names will be 50 characters or less; max length is 255.
   # Normally, only one or a few values will be in the list.
   # Using the equivalence search on this will evaluate to "true"
   # if any one of the items in this list matches the query.
   # If this list contains the canonical name "iscsi", any initiator
   # is allowed to access this target.

   boot-list = string M O
   # A list of iSCSI Initiator Names that can boot from this target.
   # This list works precisely like the access-list attribute.  A name
appearing
   # in this list must either appear in the access-list, or the
   # access-list must contain the initiator name "iscsi".  Otherwise, an
   # initiator will be unable to find its boot target.
   # If boot-list contains the name "iscsi", any host can boot from it,
   # but I am not sure if this is useful to anyone.
   # If this attribute is not registered, this target is not "bootable".
   #
   # Note that the LUN the host boots from is not specified here; a
   # host will generally attempt to boot from LUN 0.
   #
   # It is quite possible that other attributes will need to be defined
   # here for booting as well.

   --------------------------template ends here------------------------
```

6.3. iSCSI Name Service Templates

This template defines the service "service:iscsi:name-service". An entity supporting one or more iSCSI name service protocols may register itself with SLP as this service type.

Initiators wishing to discover name services using SLP will usually search for them by the protocol(s) they support:

```
     Service: service:iscsi:name-service
     Scope:   initiator-scope-list
     Query:   (protocols=isns)
```

Name of submitter: Mark Bakke
Language of service template: en
Security Considerations:
  See later section.

Template Text:

```
-------------------------template begins here-----------------------
template-type=iscsi:name-service

template-version=0.1

template-description=
  This is a concrete service type.  The iscsi:name-service service type
  provides the capability for entities supporting iSCSI to discover
  appropriate name services.

template-url-syntax=
  url-path   = The URL of the name service.  Defined in RFC 2608.

protocols = string M L
# The list of protocols supported by this name service.  This
# list may be expanded in the future.  There is no default.
#
# "isns"  - The name service supports the use of the iSNS protocol
#           to locate and register targets, and provide further
#           information on them.  This protocol is defined in [ISNS].
isns

--------------------------template ends here------------------------
```

7. Security Considerations

Service type templates provide information that is used to interpret information obtained by clients through SLP. If the iSCSI templates are modified or if false templates are distributed, iSCSI targets and name servers may not correctly register themselves, or iSCSI clients may not be able to interpret service information.

SLP provides an authentication mechanism for UAs to assure that service advertisments only come from trusted SAs. [RFC2608] If trust is an issue, particularly with respect to the information sought by the client about IPSEC and IKE support, then SLP authentication should be enabled in the network.

Once a target or name server is discovered, authentication and authorization are handled by the iSCSI protocol, or by the name server's protocol. It is the responsibility of the providers of

these services to ensure that an inappropriately advertised or discovered service does not compromise their security.

8. Summary

This document describes how SLP can be used by iSCSI initiators to find iSCSI targets and name servers. Service type templates for iSCSI targets and name servers are presented.

9. References

[RFC2608] E. Guttman, C. Perkins, J. Veizades, M. Day. Service Location Protocol, version 2 RFC 2608, July, 1999.

[RFC2609] E. Guttman, C. Perkins, J. Kempf. Service Templates and service: Schemes RFC 2609, July, 1999.

[RFC2119] S. Bradner. Key Words for Use in RFCs to Indicate Requirement Levels. RFC 2119, March 1997.

[ISCSI] J. Satran, et. al. "iSCSI", draft-ietf-ips-iscsi-04.txt, February 2001.

[SAM2] ANSI T10. "SCSI Architectural Model 2", March 2000.

[NDT] K. Voruganti, et. al. "iSCSI Naming and Discovery Requirements", draft-ietf-ips-iscsi-name-disc-01, April 2001.

[ISNS] J. Tseng, et. al. "Internet Storage Name Service", draft-ietf-ips-isns-00, January 2001.

[BOOT] P. Sarkar, D. Missimer, C. Sapuntzakis. "A Standard for Bootstrapping Clients using the iSCSI Protocol", draft-ietf-ips-iscsi-boot-02, February 2001.

[RSIP] Kempf, J., Montenegro, G., "Finding an RSIP Server with SLP", draft-ietf-nat-rsip-slp-00, February 2000.

Author's Address:

Mark Bakke
Cisco Systems, Inc.
6450 Wedgwood Road
Maple Grove, MN

USA 55311

Voice: +1 763-398-1000
E-Mail: mbakke@cisco.com

Full Copyright Statement

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.

Acknowledgement

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