Internet Draft
<draft-ietf-ips-iscsi-slp-02.txt>
Expires May 2002

Mark Bakke
Cisco

Joe Czap
Jim Hafner
John Hufferd
Kaladhar Voruganti
IBM

Howard Hall
Pirus

Jack Harwood
EMC

Yaron Klein
Sanrad

Marjorie Krueger
HP

Lawrence Lamers
San Valley Systems

Todd Sperry
Adaptec

Joshua Tseng
Nishan

November 2001

Finding iSCSI Targets and Name Servers Using SLP

Status of this Memo

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

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 management services, along with the SLP service type templates that describe the services they provide.

1. Acknowledgements

This draft was produced by the iSCSI Naming and Discovery team, including Joe Czap, Jim Hafner, John Hufferd, and Kaladhar Voruganti (IBM), Howard Hall (Pirus), Jack Harwood (EMC), Yaron Klein (Sanrad), Marjorie Krueger (HP), 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. Two of these methods can be accomplished without the use of SLP:

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

- Each address providing targets can be configured on the initiator; iSCSI provides a mechanism by which the initiator can query the address for a list of targets.

The above 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 management servers.

This draft is a work in progress. Searching for the string "WORK" in this document should find anything that is not considered to be complete. The following items are still open:

- Need to add RFC 3082 interaction. An initiator that is already up and running must be notified within a reasonable amount of time when a new target becomes available to it. This may be due to a storage device booting, a network interface being added to the device, a new target being created on the device, or the initiator being added to the access-list of an existing device. Work is under way to determine the best way to do this, either using the experimental RFC 3082 or some modification thereof. Note that it is a non-goal for SLP to notify an initiator when a target or one of its service URLs is no longer accessible; the initiator will find this out soon enough if it cares to attempt access to the target. Note that RFC 3082 takes care of a device booting, adding a new interface or target (and hence, a service URL), but not the access-list change.

- Add comments about lifetime of URLs and how it is used. URLs are registered with a finite lifetime. If the lifetime is too long, a lot of stale URLs may hang around; if it is too short, SLP participants will spend too much time re-registering the same old URLs. There is a definite recommendation by the SLP folks to stick with the default; I have to go look it up to see what it is.

- SLP can be set up to use either Unicast or Multicast. Add a discussion on when to use each.

- Storage Name Service or Storage Management Service? Need to settle on a generic name for things like this.

The following modifications have been made since draft-01:

- Removed the mgmt-ipaddress attribute from the template; if FQDN is not available, the IP address may be returned in its place as a

dotted-decimal string.

- Added example for finding targets that will allow access to any initiator.

- Updated Security Considerations to reference the IP storage security draft.

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 Management Server	An addressible entity that provides management services that benefit an iSCSI environment. "Storage management server" is used as a generic term, rather than a specific protocol or service.

5. Using SLP for iSCSI Service Discovery

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.

iSCSI targets are registered with SLP as a set of service URLs, one for each address on which the target may be accessed. Initiators discover these targets using SLP service requests. Targets that do not directly support SLP, or are under the control of a management service, may be registered by a proxy service agent as part of the software providing this service.

iSCSI entities may also use SLP to discover higher-level management services where needed.

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 management 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.

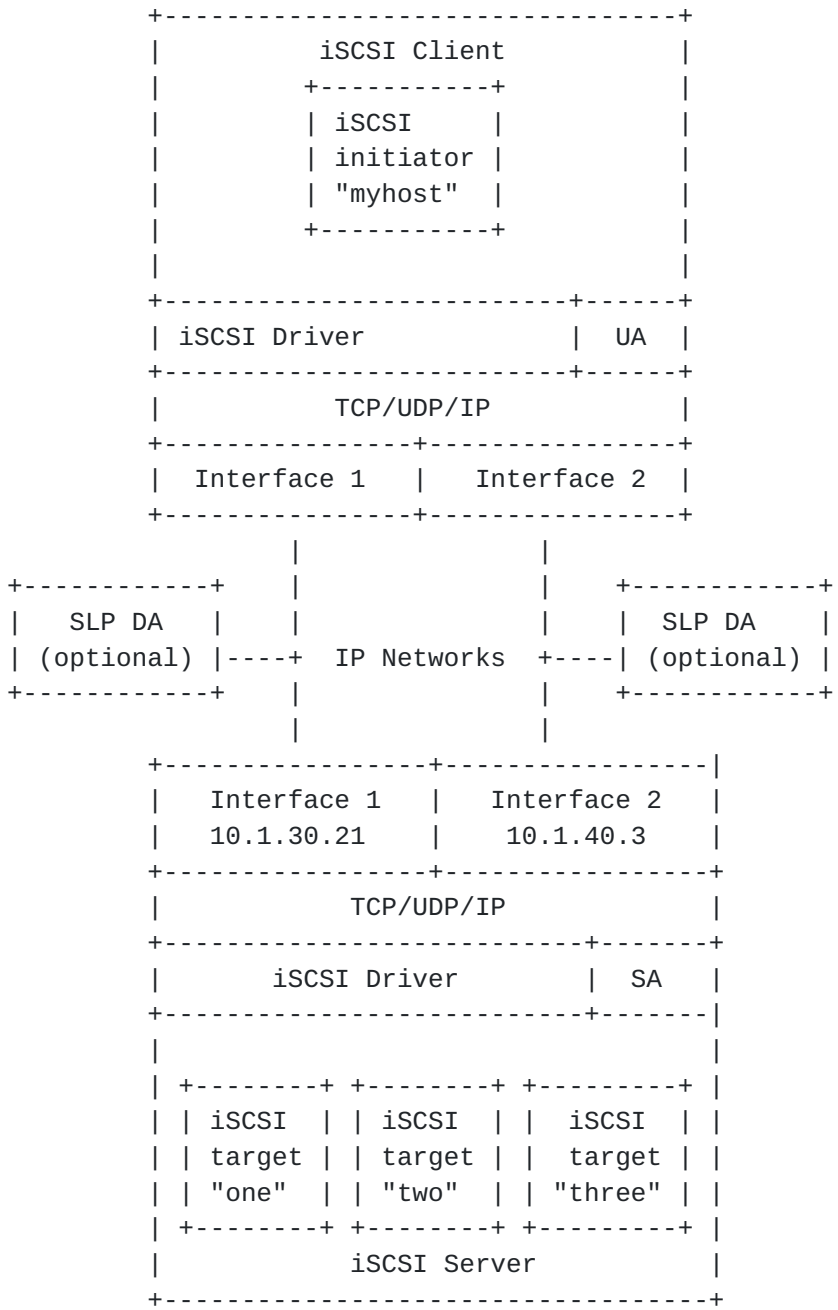

In the above drawing, the iSCSI server has three iSCSI targets that the client could discover, named "one", "two" and "three". The iSCSI client has an iSCSI initiator with the name "myhost". The iSCSI client may use the initiator name in its SLP Service Requests as a filter to discover only targets that are configured to accept iSCSI connections from "myhost".

Each iSCSI target and initiator has a unique name, called an 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", and "three" are used for the targets; the initiator uses the name "myhost". An actual iSCSI name would incorporate 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, would have two service 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 iSCSI name. An iSCSI server must register each of its individual targets at each of its network addresses.

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.

If the server in the above drawing is listening at TCP port 5003 for both network addresses, the service URLs registered would be:

- 10.1.30.21:5003/one

- 10.1.30.21:5003/two

- 10.1.30.21:5003/three

- 10.1.40.3:5003/one

- 10.1.40.3:5003/two

- 10.1.40.3:5003/three

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. If a UA is configured with the address of the SA, it may avoid multicast and contact an SA using unicast. 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.

As information contained in the iSCSI target template may exceed common network datagram sizes, the SLP implementation for both UAs and SAs supporting this template MUST implement SLP over TCP.

In some networks, the use of multicast for discovery purposes is either unavailable or not allowed. Such networks include public or service-provider networks that are placed in between an iSCSI client and server; these are probably most common between two iSCSI gateways, one at a storage service provider site, and one at a customer site.

In these networks, an initiator may, instead or in addition to its DA configuration, allow the addresses of one or more SAs to be configured. The initiator would then make unicast SLP service requests directly to these SAs, without the use of multicast to first discover them.

This functionality is well within the scope of the current SLP protocol. However, it does have two consequences for implementors:

- A service-agent responding to requests for iSCSI targets MUST implement SLP over TCP; UDP only is not enough.

- An initiator configured to make direct, unicast requests to an SA will have to add this to the SLP API, if it is following the service location API defined in [RFC2614].

5.2. Discovering Storage Management Services using SLP

Storage management servers can be built to manage and control access to 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 both clients and servers, in addition to the target discovery already being performed.

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

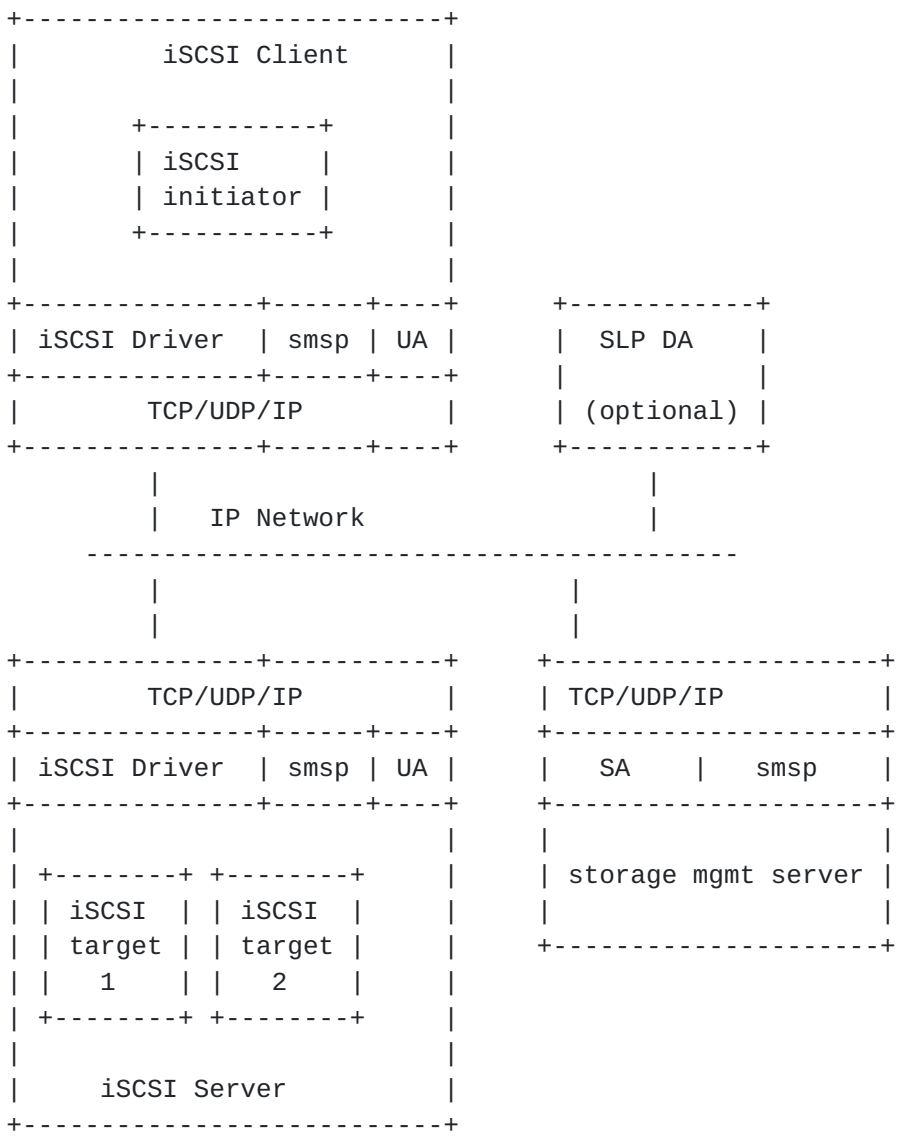

Note the difference between the storage management server model and the previously-defined target discovery model. When target discovery was used, the iSCSI Server implemented an SA, to be discovered by the initiator's UA. In the storage management server model, the iSCSI clients and servers both implement UAs, and the management server implements the SA.

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

The storage management server constructs a service advertisement of the type "service:iscsi:sms" 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 management service.

Targets that support a storage management service implement a UA in addition to the SA. A target may alternatively just implement the UA, and allow the storage management service to advertise its targets appropriately by providing an SA and registering the appropriate service:iscsi:target registrations on the target's behalf; the target device would not have to advertise its own targets. This has no impact on the initiator.

This allows the initiators' discovery of targets to be completely interoperable regardless of which storage management service is used, or whether one is used at all, or whether the target registrations are provided directly by the target or by the management service.

5.3. NAT and NAPT Considerations

Since SLP provides IP address and TCP port information within its payload, the addresses an SA or DA advertise may not be the same as those a UA must use if a Network Address(/Port) Translation (NAT/NAPT) device is present between the UA and the SA. This may result in the UA discovering address information that is unusable. Here are a few recommendations to handle this:

- Use a fully-qualified domain name instead of IP address in service URLs and in the mgmt-entity attribute.

- Stick with the default, IANA-assigned iSCSI TCP port number in service URLs, wherever possible.

- If advertising service URLs through a NAT/NAPT device, and the FQDN, IP address, or TCP port will be translated, the NAT/NAPT device can provide an SLP proxy capability to do the translation.

5.4. Implementation Considerations

This section will answer common questions for those who are not too familiar with SLP.

Where are the templates used? By the implementor; don't need to be installed in a DA (not like a MIB).

Who makes use of the templates?

- Implementor of iSCSI host drivers / adapters / devices
- Network Administrator (DHCP and DA)
- Storage Administrator (DA and SA)

Integrating SLP DA or SA within a storage management server

When to use multicast and/or unicast

Using DHCP to bootstrap SLP discovery

6. iSCSI SLP Templates

Three templates are provided: an iSCSI target template, a management 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 perhaps management services) wishing to discover targets in this way will generally use one of the following queries:

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

```
Service: service:iscsi:target
Scope:   initiator-scope-list
Query:   (iscsi-name=iqn.2001-04.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-scope-list
Query:   (access-list=iqn.1998-03.com.os.hostid.045A7B)
```

3. Find all of the iSCSI Target Names that may allow access to any initiator:

```
Service: service:iscsi:target
Scope:   initiator-scope-list
Query:   (access-list=iscsi)
```

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

```
Service: service:iscsi:target
Scope:   initiator-scope-list
Query:   (boot-list=iqn.1998-03.com.os.hostid.045A7B)
```

5. In addition, a management service may wish to discover all targets:

```
Service: service:iscsi:target
Scope:   management-server-scope-list
Query:   <empty-string>
```

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 must be the iSCSI
     ; name of the target being registered.
     ; A device representing multiple targets must individually
     ; register each target/address combination with SLP.
     ;
     ; Example:
     ;   service:iscsi:target://10.1.3.40:5003/iqn.2001-04.com.acme.sn.45678

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

   portal-group = integer
   # The iSCSI portal group tag for this address.  Addresses sharing
   # the same iscsi-name and portal-group tag can be used within the
   # same iSCSI session.  Portal groups are described in [ISCSI].

   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

   mgmt-entity = string O
```

```
   # The fully qualified domain name, or IP address in dotted-decimal
   # notation, of the management interface of the entity containing
   # this target.
   #
   # WORK - Should this be a URL?
   #  snmp://10.1.1.1
   #  http://mydisk.ssp.com:1080/
   #  telnet://mydisk.ssp.com

   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 default 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 Storage Management Service Templates

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

iSCSI clients and servers wishing to discover storage management services using SLP will usually search for them by the protocol(s)

they support:

```
Service: service:iscsi:sms
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:sms

template-version=0.1

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

template-url-syntax=
  url-path   = ; The URL of the management 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"  - This management service supports the use of the iSNS
#           protocol for access management, health monitoring, and
#           discovery management services.  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 management server is discovered, authentication and authorization are handled by the iSCSI protocol, or by the management 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.

7.1. IPsec Integration

Although SLPv2 security provides authentication, it does not provide confidentiality.

The use of IPsec and IKE for SLPv2 is discussed in [IPS-SEC], and is a work in progress. It will be discussed further here in a subsequent draft revision.

8. Summary

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

[RFC2614] J. Kempf, E. Guttman. An API for Service Location RFC 2614, June 1999.

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

[RFC3082] J. Kempf, J Goldschmidt. Notification and Subscription for SLP. RFC 3082, March 2001.

[ISCSI] J. Satran, et. al. "iSCSI", draft-ietf-ips-iscsi-08.txt, September 2001.

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

[NDT] K. Voruganti, et. al. "iSCSI Naming and Discovery", draft-ietf-ips-iscsi-name-disc-03, July 2001.

[ISNS] J. Tseng, et. al. "Internet Storage Name Service", draft-ietf-ips-isns-05, November 2001.

[BOOT] P. Sarkar, D. Missimer, C. Sapuntzakis. "A Standard for Bootstrapping Clients using the iSCSI Protocol", draft-ietf-ips-iscsi-boot-03, August 2001.

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

[IPS-SEC] B. Aboba, et. al., "Securing iSCSI, iFCP, and FCIP", draft-ietf-ips-security-04, October 2001.

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

English.

Acknowledgement

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