To: T10 Technical Committee
From: Greg Pellegrino (Greg.Pellegrino@compaq.com)
and Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)
Date: 16 January 2001
Subject: SRP InfiniBand™ annex

Revision History
Revision 0: 5 Jan 2001 first revision
Revision 1: 11 Jan 2001 with updates from the SRP meeting.

Related Documents
T10/srp-r02 – SCSI over RDMA revision 2
InfiniBand Architecture Volume 1 – General Specifications, Release 1.0
InfiniBand Architecture Volume 2 – Physical Specifications, Release 1.0
InfiniBand Architecture Volume 3 – Application of InfiniBand, Release 0.9: (draft)

T10/99-245r9 Access Controls (by Jim Hafner)
T10/00-261r0 Discussion of editorial changes to Access Controls (by Jim Hafner)
T10/00-287r1 TransportIDs for Access Controls (by Jim Hafner)
T10/00-381r0 Three minor modifications to Access Controls (by Jim Hafner)
T10/00-425r0 Long Identifiers in SPC-3, SAM-2, SBC-2 and other XOR issues (by Jim Hafner)
T10/01-026r0 SPC-3 Access Control conflicts due to TransportIDs (by Rob Elliott)

T11/FC-GS-3 Revision 6.42 (Section 6.1.2.3 Platform Object)
T11/99-697v0 Management Server Platform Extension (by Duane Baldwin) (source of FC-GS-3)

Overview
This proposes topics and text for an InfiniBand annex for the SCSI over RDMA (SRP) standard.

The goal is to identify all optional InfiniBand features that must be implemented to ensure useful, interoperable SRP devices. InfiniBand Volume 3 will describe how boot devices, a subset of SRP devices, are specifically identified and handles ("Storage Boot Wire Protocol").

All text outside [brackets] is part of the suggested text.

Suggested text.

Annex A (normative)
SRP for InfiniBand™ Architecture

[footnote] InfiniBand is a trademark and service mark of the InfiniBand Trade Association.

A.1 Related documents

InfiniBand™ Architecture Volume 1 – General Specifications. Release 1.0. InfiniBand^{SM} Trade Association.
InfiniBand™ Architecture Volume 2 – Physical Specifications. Release 1.0. InfiniBand^{SM} Trade Association. [needed for Chassis and Module discussion]
InfiniBand™ Architecture Volume 3 – Application of InfiniBand. Release 0.9 (draft). InfiniBand^{SM} Trade Association.

IETF RFC 2460 – Internet Protocol, Version 6. S. Deering and R. Hinden. Internet Engineering Task Force. [needed if GID format is referenced]

[rules: do not cross-reference section numbers in other documents. Can refer to section headings.]

A.2 Glossary

[Items may be added to the main glossary or defined in the annex. Since they are only used in the annex, I suggest defining them here.]

A.2.1 Overview

See the InfiniBand Architecture Volume 1 glossary for full definitions of InfiniBand terms.

A.2.2 Acronyms

CA: Channel adapter
GID: Global ID
GSA: General service agent
HCA: Host channel adapter
IBA: InfiniBand™ architecture
IOC: IO controller
IPv6: Internet Protocol version 6
LID: Local ID
MAD: Management datagram
QP: Queue pair
QPN: Queue pair number
ROM: Read only memory
SMA: Subnet management agent
SRP: SCSI over RDMA protocol
TCA: Target channel adapter

A.2.3 Definitions

A.2.3.1 Channel adapter (CA): Device that terminates a link and executes transport-level functions. One of a Host channel adapter or Target channel adapter.

A.2.3.2 Communications manager: The software, hardware, or combination of the two that supports the communication management mechanisms and protocols used to establish and release InfiniBand connections.

A.2.3.3 Global ID (GID): A port address used for directing packets between subnets. A GID is a valid 128-bit IPv6 address. Source and destination GIDs are carried in an optional global routing header.

A.2.3.4 General Service Interface: An interface providing management services other than subnet manager. Queue Pair 1.

A.2.3.5 Host channel adapter (HCA): A channel adapter used to support processor nodes.

A.2.3.6 IO unit: One or more IO controllers attached to the fabric through a single TCA.

A.2.3.7 IO controller: The part of an IO unit that provides IO services.

A.2.3.8 IPv6 Address: A 128-bit address constructed in accordance with IETF RFC 2460 for IPv6.

A.2.3.9 Local ID (LID): A port address used for directing packets within a subnet. Source and Destination LIDs are carried in every packet header.

A.2.3.10 Management datagram (MAD): A packet used for communication to manage an InfiniBand network.

A.2.3.11 Node: A channel adapter.

A.2.3.12 Port: Location on a channel adapter to which a link connects.

A.2.3.13 Queue pair (QP): An interface used for communication.

A.2.3.14 Queue pair number (QPN): A value that identifies a queue pair within a channel adapter.

A.2.3.15 Service ID: A value that allows a Communication Manager to associate an incoming connection request with the entity providing the service.

A.2.3.16 Subnet: A set of InfiniBand ports and links that have a common Subnet ID and are managed by a common Subnet Manager.

A.2.3.17 Subnet manager: Entity that configures and controls a subnet.

A.2.3.18 Subnet management agent (SMA): An entity present in every channel adapter that processes management datagrams from a subnet manager.

A.2.3.19 Target channel adapter (TCA): A channel adapter used to support I/O units.

A.3 Overview

This annex describes how SCSI over RDMA protocol (SRP) is used over InfiniBand Architecture (IBA), a transport that provides the necessary RDMA semantics.

An IBA processor node contains one or more host channel adapters (HCAs), each containing one or more ports (see Figure 1).

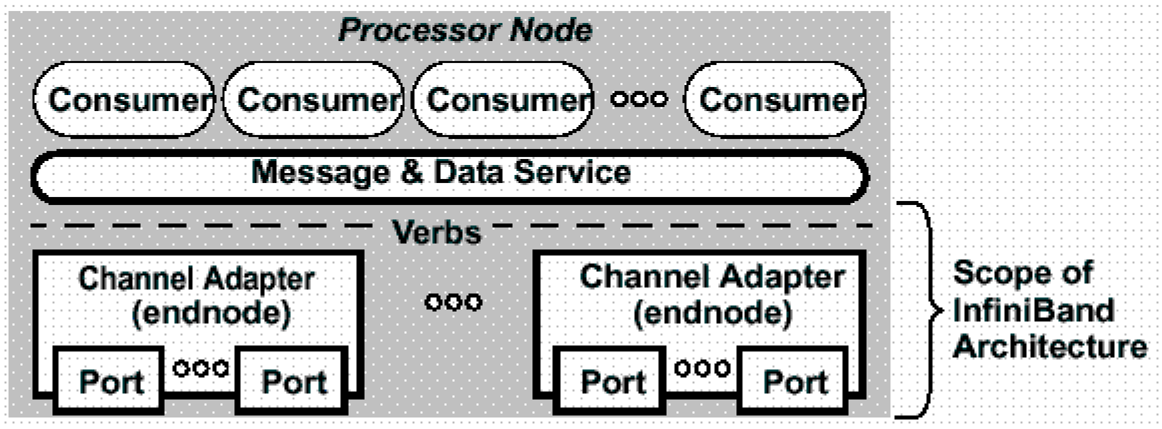

Figure 1. Processor node (from InfiniBand Architecture specification).

[Editor's note: taken from IBA Figure 17. Figure 171 might be a better source.]

An IBA IO unit contains a target channel adapter (TCA) with one or more ports, a subnet management agent (SMA), general service agents (GSAs), and one or more IO controllers (IOCs) (see Figure 2). Each IO controller may advertise multiple service IDs.

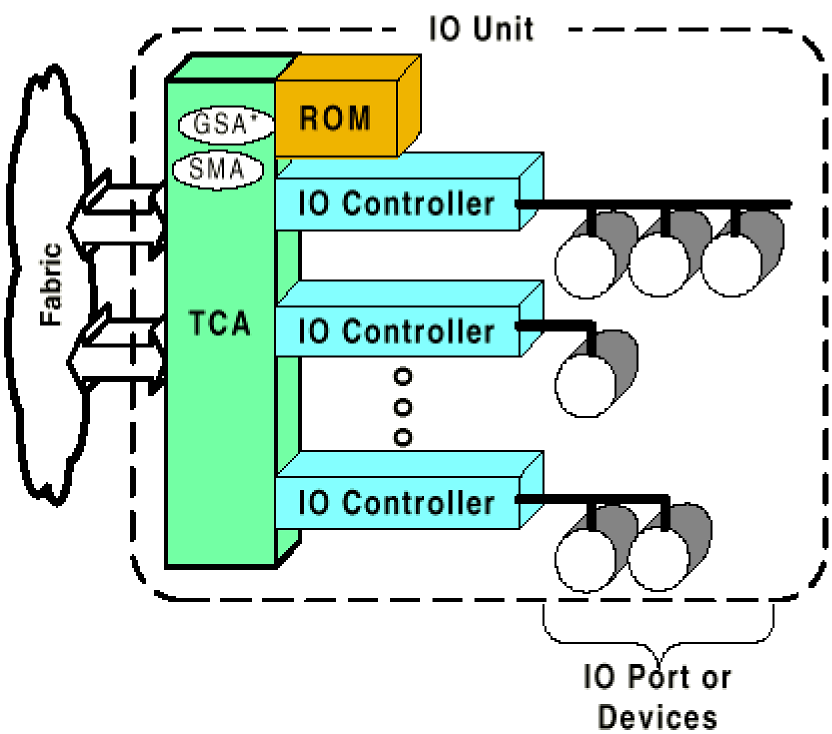

Figure 2. IO unit (from InfiniBand Architecture specification).

[Editor's note: taken from IBA Figure 35. Figure 171 might be a better source – it doesn't show a ROM.]

[Editor's note: Here is a picture adding in module and chassis concepts from Volume 2 that may also be useful:

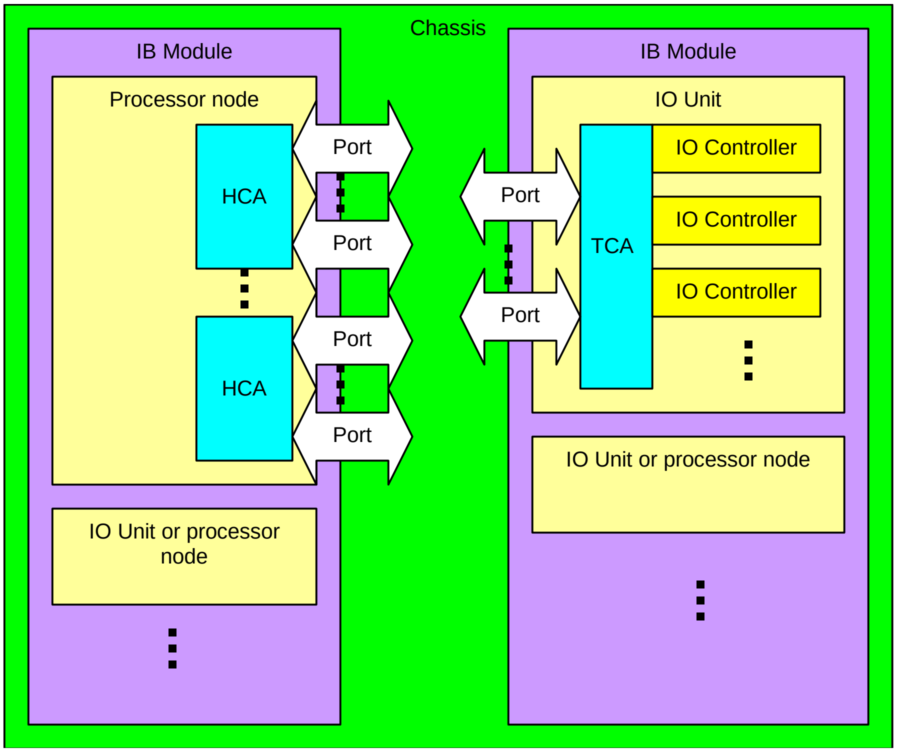

Figure 3. InfiniBand Architecture objects.

]

Each port has a globally unique identifier called a port GUID. Each channel adapter has a node GUID (which may be shared between multiple ports). Each IO controller has an IOC GUID.
[Editor's option:]
Each port is assigned a local ID (LID) or a range of LIDs by the subnet manager. Each port has one or more global IDs (GID). Each GID is globally unique, formed in part from the port GUID. The subnet manager provides GUID to GID/LID resolution.

An IBA module contains one or more xCA, switches, or routers. An IBA chassis contains one or more IBA modules. Each IBA module may contain a module GUID. Each IBA chassis may contain a chassis GUID.

This annex defines additional GUIDs called an initiator platform GUID and target platform GUID. When used in a processor node, either of these GUIDs may be shared among multiple channel adapters and is based on either the chassis GUID or a port GUID. When used in an IO controller, the GUID is based on the IO controller GUID.

Table 1 summarizes the IBA names and addressable entities relevant to SRP.

Table 1. IBA names and addressable entities.

Name	Scope of uniqueness	Description
Port GUID	worldwide	EUI-64 for each port within a channel adapter
Node GUID	worldwide	EUI-64 for each channel adapter (IO unit = TCA)

IO controller GUID	worldwide	EUI-64 for each IO controller in an IO unit
Chassis GUID	worldwide	EUI-64 for each chassis
Module GUID	worldwide	EUI-64 for each module within a chassis
Initiator platform GUID	worldwide	(new) 128 bits for each SRP initiator processor node-based initiator: created from the chassis GUID or a port GUID IO unit-based initiator: created from the IO controller GUID
Target platform GUID	worldwide	(new) 128 bits for each SRP target processor node-based target: created from the chassis GUID or a port GUID IO unit-based target: created from the IO controller GUID
LID	subnet	assigned by subnet manager to each port
GID	worldwide	assigned by subnet manager; subnet prefix plus the port GUID

[Editor's note: end of InfiniBand overview… now discuss how SRP maps into it]

An SRP initiator in a processor node uses a set of HCAs that share an initiator platform GUID.
An SRP initiator in an IO unit is an IO controller.

[Editor's note: if an SRP initiator is implemented in an IO controller, the IO controller cannot share [T]CAs and still follow the IO unit model. It can act as a single-CA processor node, though.]

An SRP target is an IBA IO controller. An IBA IO unit may contain one or more SRP targets.

[Editor's note: if an SRP target is implemented by a host using an HCA, it has to support inbound DevMgt MADs and look like an IO controller inside an IO unit. Unlike the initiator, defining SRP target for processor nodes does not describe how to find them.]

Each queue pair (QP) is a separate I_T nexus.

A.4 Communication Management

A.4.1 Overview

Communications Managers on each InfiniBand device manage InfiniBand connections using MADs over the General Service Interface on each system. SRP devices shall use the active/passive (client/server) connection establishment protocol. The processor node or IO unit containing the SRP target shall act as the server and the processor node or IO unit containing the SRP initiator shall act as the client.

A.4.2 Establishing a connection

[Editor's note: it is beyond the scope of SRP to define a Configuration Manager/name server like device to help an initiator find IO unit LID/GIDs worth querying for SRP services. Assuming that such devices have been found, a connection is established as described below.]

To discover how to establish an InfiniBand connection to an SRP target, the client retries the IOUnit and ServiceEntries attributes from an IO unit using a DevMgtGet MAD. This provides the list of IOCs, their supported protocols, and their associated Service IDs. The client filters the list to find a Service ID for "SRP.T10.NCITS".

[Editor's note: Ed Gardner suggests "SRP.T10.NCITS". IBTA TWG needs to assign this string. Boot-capable devices (with Boot ROMs) will advertise "SBWP.IBTA" as well.]

To establish an InfiniBand connection, the client places the Service ID in a Request (REQ) message. The server associates the request with the appropriate SRP target. If it accepts the connection request, the server returns a queue pair number (QPN) in a Response (REP) message. The client replies with a Ready To Use (RTU) message.

[Editors node: Two options for PrivateData are presented below. They are not exclusive; both could be used.]

[Editors Option 1:]
The PrivateData field of the REQ message may [shall?] include an SRP_LOGIN_REQ IU. The SRP target may choose to refuse the connection based on parsing this login message. If it accepts the login, the SRP target shall return an SRP_LOGIN_RSP IU in the PrivateData field of the REP MAD. This creates an SRP connection at the same time the InfiniBand connection is created.

[Editors Option 2:]
The PrivateData field of the REQ message includes a 128-bit initiator platform GUID. This uniquely identifies an SRP initiator and may be shared across multiple channel adapters. The PrivateData payload shall be constructed as shown in Table 2.

The REP message includes a 128-bit target platform GUID. This uniquely identifies an SRP target and may be shared across multiple channel adapters. The PrivateData payload shall be constructed as shown in Table 3.

In an IO controller, an initiator or target platform GUID shall be constructed by concatenating the IO controller GUID to a unique 64-bit suffix. In a processor node, an initiator or platform GUID shall be constructed by concatenating a unique 64-bit suffix to either a chassis GUID or a port GUID accessible to the platform.

[Coordinate this definition with the SNIA Fibre Channel Management work group's definition of platform ID generation algorithm.]

Table 2. Initiator Platform GUID

Byte Bit	7	6	5	4	3	2	2	0
0	TYPE (02H)							
1	RESERVED							
2	RESERVED							
3	RESERVED							
4	(MSB)							
	INITIATOR PLATFORM GUID MOST SIGNIFICANT 64 BITS							
11	(FROM CHASSIS GUID, PORT GUID, OR IO CONTROLLER GUID)							
12								
	INITIATOR PLATFORM GUID LEAST SIGNIFICANT 64 BITS							
19								(LSB)

Table 3. Target Platform GUID

Byte Bit	7	6	5	4	3	2	2	0
0	TYPE (03H)							
1	RESERVED							
2	RESERVED							
3	RESERVED							
4	(MSB)							
	TARGET PLATFORM GUID MOST SIGNIFICANT 64 BITS							
11	(FROM CHASSIS GUID, PORT GUID, OR IO CONTROLLER GUID)							

12	TARGET PLATFORM GUID LEAST SIGNIFICANT 64 BITS
19	(LSB)

If Alternate Path Migration (APM) is supported, the initiator and/or target may include a second address in its REQ or REP packet that may be used in case of failure of the path to the original address. This address includes an alternate LID and GID.

A.4.3 Releasing a connection

To release both an InfiniBand connection and any associated SRP connection, either the client or server sends a Disconnect Request (DREQ) MAD. The other device replies with a Disconnect Reply (DREP). This terminates the InfiniBand connection and any SRP connection that was still established.

[Editor's option 1: if an SRP_LOGOUT_REQ IU is added, it could be sent as PrivateData with DREQ. If SRP_LOGOUT_RSP is also added, it could be sent as PrivateData in DREP.]

[Editor's note: be careful of the difference between InfiniBand connection and SRP connection. FCP has a table regarding login/logout effects that might be needed here.]

A.5 Initiator-specific requirements

SRP Initiators shall support the Send and RDMA Write inbound transport functions. An SRP initiator shall not issue any data-out commands unless it supports the inbound RDMA Read transport function. [is the latter sentence too pedantic?]

The QP associated with an SRP target shall accept RDMA Write requests. If the SRP initiator issues data-out commands, the QP shall also accept RDMA Read requests.

SRP initiators are not required to check inbound RDMA transfer lengths (either to verify that the maximum transfer size of 2^31 was not exceeded or to verify that the sum of the packet payloads adds up to the requested transfer size).

A.6 Target-specific requirements

SRP targets shall support the Send inbound transport function. An SRP target shall not issue either the RDMA Write with Immediate Data or ATOMIC transport functions. SRP targets shall not use RDMA transfer lengths that exceed the maximum transfer size of 2^31.

SRP targets shall ensure that the sum of the packet payloads adds up to the requested transfer size.

The QP associated with an SRP target shall not accept RDMA requests.

SRP targets are not required to assume trivial subnet manager responsibilities.

SRP targets shall implement the DevMgt class of general management services.

SRP targets shall include device management IOUnit attributes as described in Table 4.

Table 4. IOUnit attributes for SRP devices

Field	Description
Change ID	As defined in InfiniBand Volume 1.
Max Controllers	At least one.
Option ROM	Vendor-specific [required for SBWP devices]
Controller List	At least one IO controller must be present.

SRP targets shall include the device management IOControllerProfile attributes as described in Table 5.

Table 5. IOControllerProfile attributes for SRP devices

Field	Description
[IO controller] GUID	The GUID shall be the same as that reported for LUN 0 in INQUIRY VPD Page 83h Type 2h (EUI-64) or Type 3h (FC) NAA=IEEE registered or NAA=IEEE registered extended
Device ID	vendor-specific [PCI compatible format?]
Vendor ID	vendor-specific [PCI compatible format?]
Device Version	vendor-specific [PCI compatible format?]
Subsystem Vendor ID	vendor-specific [PCI compatible format?]
Subsystem [Device] ID	vendor-specific [PCI compatible format?]
IO Class	TBD assigned by AWG
IO Subclass	TBD assigned by AWG
Protocol	Set to the value assigned to SRP by AWG.
Protocol Version	SRP devices shall use 0x0000. 0x0001 is reserved for future SRP-2 devices.
Service Connections	at least one
Initiators Supported	vendor-specific
Send Message Depth	vendor-specific
RDMA Transfer Size	vendor-specific
Controller Operations Capability Mask	These bits shall be set to one: 0: ST (Send Messages to IOCs) 1: SF (Send Messages from IOCs) 5: WF (RDMA Write Requests from IOCs) This bit shall be set to 1 by SRP targets supporting data-out commands: 3: RF (RDMA Read Requests from IOCs) These bits shall be set to zero: 2: RT (RDMA Read Requests to IOCs) 4: WT (RDMA Write Requests to IOCs) 6: AT (Atomic Operations to IOCs) 7: AF (Atomic Operations from IOCs)
Controller Services Capability Mask	This bit is set for SRP devices with boot support (see InfiniBand volume 3) 1: SBWP Storage Boot Wire Protocol
Service Entries	Set to the number of services entries (at least one).
ID String	Vendor-specific

The SRP target IO controller shall register with its Communications Manager a Service Name string of “SRP.T10.NCITS”. This string is assigned an IO SERVICE ID type service ID by the Communications Manager.

[Editor’s note: If we want to allow multiple targets within a single IO controller, we need a Service ID range assigned by IBTA. If one target per IO controller suffices, one Service ID suffices. Cris Simpson will push this thru the IBTA TWG.]

[Editor’s note (duplicate from earlier): Ed Gardner suggests “SRP.T10.NCITS”. IBTA needs to assign this string. Boot-capable devices (with Boot ROMs) will advertise “SBWP.IBTA” as well.]

SRP targets shall include the Service Name/Service ID pair in the device management ServiceEntries attribute pair at described in Table 6.

Table 6. ServiceEntries attribute pair for SRP devices

Field	Length	Description
ServiceName_n	320	"SRP.T10.NCITS"
ServiceID_n	64	Assigned by Communications Manager

[Editor's note: Boot devices are identified this way: An IO controller indicates that it supports the SRP storage protocol for booting by setting the SBWP bit in the IO Controller Profile's Controller Services Capability Mask. See IBA volume 1 release 1.0 Table 222 in section 16.3.3.4. It also sets IOUnitInfo.OptionROM to 1 and provides a ROM image for booting.]

[Editor's note: IBTA AWG may recommend a similar bit for all SRP devices, or leave discovery to the Service ID strings.]

A.7 Common initiator and target requirements

Targets and initiators shall support the Reliable Connection transport service type.

[Editor's note: Ed Gardner would like to define RD too. Need to add a Service ID range for RD, add EE context to the initiator identifier, and make other changes to support it. Ed will write a proposal or defer it until SRP-2.]

End-to-end flow control is not required in either the SRP initiator or target.

[Editor's note: Are Solicited Events needed? Ed thinks maybe... defer until someone asks for it.]

[Editor's note: Should automatic failover support (path migration) be mandatory? We could require the target support responding to automatic path migration requests (Alternate Path Response APR) but not require it to generate them (Load Alternate Path LAP) if it loses contact with the initiator. This way the initiator can force a failover if it wants. This feature would be optional to the initiator. This decision cannot be made until initiator identification is finalized.]

Any Service Level may be negotiated and used. Any Virtual Lanes may be negotiated and used.
[Editor's note: The SL for an SRP connection is chosen by the initiator. The target shall accept any SLs. Should this non-requirement be removed? The issue draws many questions so I suggest mentioning it here.]

SRP devices shall not set the Fence indicator on any transfers.
[Editor's note: The only case in Volume 1 section 10.8.3 that might break is an RDMA READ followed by a SEND before the RDMA READ completes. This should never occur as the device must wait for data transfer for the RDMA before issuing another SEND (e.g. for an SRP_LOGOUT_REQ IU).]

[Editor's note: There seems to be little value in listing these non-requirements:

Path MTU size support is device-specific. [all SRP message fit within the minimum size]

SRP devices shall not use multicast operations. [implied by not using unreliable datagrams]

SRP devices are not required to support more than two P_KEYs per port.

SRP devices are not required to support the Bad P_Key Trap and P_Key Violations Counter. [Volume 1 section 10.9.3-4]

SRP devices are not required to check M_Keys for reads or maintain M_Keys through power loss [Volume 1 section 10.9.9]

SRP devices are not required to implement any minimum RNR (resource not ready) NAK Timer field values.

Performance management features are all optional. [Volume 1 Section 16.1.3.2]

]

A.8 Initiator and target identifiers

The initiator identifier used by a target and recorded by a target for persistent reservations shall be the initiator LID, GID, and QPN. The initiator port GUID shall serve as the world wide identification also maintained by persistent reservations, used to track initiator identifier changes.

[Editor's note: if Reliable Datagrams were supported, then the initiator's EE Context would need to be retained along with the QPN.]

[Editor's note: Since the Port GUID is a worldwide unique name, the name persists through power loss.]

[Editor's note: It would be beneficial to use an identifier of wider scope than the Port GUID like the Node GUID or Platform GUID. A Node GUID would be unaffected by alternate path migration, LID changes, and queue pair number selection. The Platform GUID would be unaffected by CA swaps as well. However, this issue also applies to Fibre Channel and other transports. Adding this feature to Persistent Reservations in SPC-3 for all transports will be proposed separately.]
If Alternate Path Migration (APM) is supported and the initiator identifies a second address in its REQ packet during connection establishment, the target shall also record as part of an initiator identifier the second address. If requested by the initiator, the target shall start using that address in place of the original address.

The target identifier used by an initiator shall be the target LID, GID, and QPN.

[Editor's note: if Reliable Datagrams were supported, then the target's EE Context would need to be retained along with the QPN.]

[Editor's note: need to define the target port identifier also used during persistent reservations?]

A.9 Extended Copy target descriptor

[Editor's note: this goes in SPC-3 rather than in the protocol standard.]

[Editor's note: There are 3 FC target descriptors:
a) port WWN
b) port ID (D_ID, the 24-bit FC address)
c) port ID (D_ID) with port WWN checking.
None of them are node-based or platform-based.

The SRP equivalents are roughly:
a) Port GUID (8 bytes)
b) LID/GID (2 + 16 bytes = 18 bytes)

c) LID/GID with Port GUID checking (2 + 16 + 8 = 26 bytes)

With the 1 byte TYPE field and 8 byte LUN fields added, only the first fits in the 24-bit target descriptor and is documented here. The REPORT/CHANGE ALIASES method must be used to create longer descriptors – see T10/00-425 and below.

Should SRP move to higher levels and use Node GUID (8 bytes), IO Controller GUID (8 bytes), or Target Platform GUID (16 bytes)?

Should the Service ID be specified or should a search for a Service Name “SRP.T10.NCITS” be assumed?]

The target descriptor used by EXTENDED COPY shall have this format:

Table 7. Target descriptor for SCSI over RDMA (for InfiniBand) – Port GUID

Byte Bit	7	6	5	4	3	2	2	0
0	TYPE (05H)							
1	RESERVED							
2	RESERVED							
3	RESERVED							
4	LOGICAL UNIT NUMBER							
..								
11								
12	RESERVED							
..								
15								
16	TARGET PORT GUID							
…								
23								

The copy manager shall query the subnet manager to map the target Port GUID into a GID and/or LID.

A.10 REPORT ALIASES and CHANGE ALIASES Identifier/Address format

[Editor’s note: this may go in SPC-3 rather than in the protocol standard.]

[Editor’s note: based on Jim Hafner’s T10/00-425r0, which is new and subject to change.]

The identifier/address formats in Table 8 are used by the REPORT ALIASES and CHANGE ALIASES commands to identify SRP targets.

[Editor’s note: for Alternate Path Migration support, should a second LID/GID address field always be included?]

Table 8. Identifier/Address formats for SRP (for InfiniBand)

Protocol Identifier	Protocol Description	Type	Type Description	Format/Length
04h	SRP-n	00h	LID/GID	2 byte LID + 16 byte GID
		10h	LID/GID with Port GUID checking	2-byte LID + 16 byte GID + 8 byte GUID
		11h	LID/GID with Node GUID checking	2-byte LID + 16 byte GID + 8 byte GUID

		12h	LID/GID with IOC GUID checking	2-byte LID + 16 byte GID + 8 byte GUID
		13h	LID/GID with Platform GUID checking	2-byte LID + 16 byte GID + 16 byte GUID
		20h	Port GUID	8 byte GUID
		21h	Node GUID	8 byte GUID
		22h	IOC GUID	8 byte GUID
		23h	Platform GUID	16 byte GUID

A.11 Access Controls TransportID

[Editor's note: this goes in SPC-3 rather than in the protocol standard.]

The TransportID used by a target to identify initiators for access controls shall have the format in Table 9:

Table 9. TransportID for SRP (for InfiniBand)

Byte Bit	7	6	5	4	3	2	2	0
0	TYPE (02H)							
1	RESERVED							
2	RESERVED							
3	RESERVED							
4	RESERVED							
..	RESERVED							
7	RESERVED							
8	[EDITOR'S OPTION:] INITIATOR PLATFORM GUID MOST SIGNIFICANT 64 BITS							
…								
15								
16	[EDITOR'S OPTION:] INITIATOR PORT GUID [EDITOR'S OPTION:] INITIATOR NODE GUID [EDITOR'S OPTION:] INITIATOR PLATFORM GUID LEAST SIGNIFICANT 64 BITS							
…								
23								

[Editor's note: Fibre Channel's definition was originally the initiator Port and/or Node WWN. Dropping the Port WWN is discussed in T10/01-026.

Using the Node GUID (aka CA GUID) makes it immune to alternate path migration, LID changes, and queue pair number assignment variation. It is not immune to CA swaps. A Platform GUID would be immune to CA swaps.]

[Editor's note: It is important not to create multiple levels of transport IDs that may conflict with each other. The Initiator Platform GUID, if accepted, would work best.]

<u>Additional notes</u>

This text is not part of the proposed annex.

[Additional background on initiator identifier:]
[Excerpt from SPC-2 revision 18c:

> The device server shall preserve the following information for each registration across any reset, and if the APTPL capability is enabled, across any power off period:
> a) Initiator identifier;
> b) reservation key; and
> c) when supported by the protocol, the initiator port's world wide identification.

The device server shall preserve the following reservation information across any reset, and if the APTPL capability is enabled, across any power off period:
a) Initiator identifier;
b) reservation key;
c) scope;
d) type; and
e) when supported by the protocol, the initiator port's world wide identification.

For those protocols for which the initiator port's world wide identification is available to the device server the initiator port's world wide identification shall be used to determine if the initiator identifier has changed. This determination shall be made at any time the target detects that the configuration of the system may have changed. If the initiator identifier changed, the device server shall assign the new initiator identifier to the existing registration and reservation of the initiator port having the same world wide identification.

]

[Additional background on FC worldwide names:]
[Excerpt from FCP-2 revision 5:

5.3 Use of World Wide Names

As specified in FC-FS, each Fibre Channel node and each Fibre Channel port shall have a Worldwide_Name, a unique name using one of the formats defined by FC-FS and its extensions. Each target and its associated logical units has knowledge of the Port_Name of each initiator through the Fibre Channel login process. As a result, the relationship between address identifier of the initiator and a persistent reservation for a logical unit may be adjusted as defined in SPC-2 during those reconfiguration events that may change the address identifier of the initiator. If a target receives a PRLI or a PLOGI from an initiator FCP_Port with a previously known Worldwide_Name but with a changed initiator identifier, the device server shall assign the new initiator identifier to the existing registration and reservation to the initiator port having the same Worldwide_Name.

Each logical unit shall be able to present a Worldwide_Name through the INQUIRY command vital product data device identification page as defined by SPC-2. For devices compliant with this standard and having a LUN 0, the Worldwide_Name of the logical unit having a LUN of 0 may be the same as the Node_Name of the target. The Worldwide_Name for the port shall be different from the Worldwide_Name for the node.

]