

Remote Direct Data Placement
Working Group
Internet-Draft
Intended status: Informational
Expires: March 17, 2007

C. Bestler
R. Stewart
Editor
September 13, 2006

Stream Control Transmission Protocol (SCTP) Direct Data Placement (DDP)
Adaptation
[draft-ietf-rddp-sctp-07.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on March 17, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document specifies an Adaptation Layer to provide a Lower Layer Protocol (LLP) service for Direct Data Placement (DDP) using the Stream Control Transmission Protocol (SCTP).

Table of Contents

1.	Introduction	3
2.	Definitions	4
3.	Conventions	6
4.	Introduction	7
4.1.	Motivation	7
4.2.	Overview	7
5.	Data Formats	8
5.1.	Adaptation Layer Indicator	8
5.2.	Payload Data Chunks	8
5.2.1.	DDP Source Sequence Number (DDP-SSN)	9
5.2.2.	DDP Segment Chunk	9
5.2.3.	DDP Stream Session Control	10
6.	DDP Stream Sessions	11
6.1.	Sequencing	11
6.2.	Legal Sequence: Active/Passive Session Accepted	11
6.3.	Legal Sequence: Active/Passive Session Rejected	12
6.4.	Legal Sequence: Active/Passive Session Non-ULP Rejected	12
6.5.	ULP Specific Sequencing	13
6.6.	Other Sequencing Rules	13
7.	SCTP Endpoints	14
7.1.	Adaptation Layer Indication Restriction	14
7.2.	Multihoming Implications	14
8.	Number of Streams	15
9.	Fragmentation	16
10.	Sequenced Unordered Operation	17
11.	Procedures	18
11.1.	Association Initialization	18
11.2.	Chunk Bundling	18
11.3.	Association Termination	19
12.	IANA considerations	20
13.	Security Considerations	21
14.	Contributors	22
15.	Acknowledgments	23
16.	References	24
16.1.	Normative references	24
16.2.	Informative references	24
	Authors' Addresses	26
	Intellectual Property and Copyright Statements	27

1. Introduction

This document describes a method to adapt Direct Data Placement [[I-D.ietf-rddp-ddp](#)] to Stream Control Transmission Protocol (SCTP) [[RFC2960](#)].

Some implementations may include this adaptation layer within their SCTP implementations to obtain maximum performance but the behavior of SCTP will be unaffected. An SCTP Layer used solely by this adaptation layer is able to take certain optimizations based on the limited subset of SCTP capabilities used. In order to allow optimization for these implementations we specify the use of the new adaptation layer indication as defined in [[I-D.ietf-tsvwg-addip-sctp](#)]

2. Definitions

DDP - See Direct Data Placement Protocol.

DDP Endpoint - The logical sender/receiver of DDP Segments. An SCTP Stream pair is not assumed to have a DDP Endpoint. DDP Segments may only be sent once a DDP Endpoint has been assigned to an SCTP Stream pair by a local interface.

DDP Source Stream Sequence (DDP-SSN) - A stream specific sequence number assigned by the Adaptation Layer for each SCTP Data Chunk sent. This is the order that chunks were submitted to SCTP, no matter what order they are actually sent or received in.

DDP Segment - The smallest unit of data transfer for the DDP protocol. It includes a DDP Header and ULP Payload (if present). A DDP Segment should be sized to fit within the Lower Layer Protocol MULDPU.

DDP Segment Chunk - An SCTP Payload Data Chunk that encapsulates the DDP-SSN and a DDP Segment.

DDP Stream - a sequence of DDP Segments whose ordering is defined by the LLP. For SCTP, a DDP Stream maps directly to a bi-directional pair of SCTP streams with the same Stream IDs. Note that DDP has no ordering guarantees between DDP Streams.

DDP Stream Session - A single pairing of DDP Endpoints over a DDP Stream that lasts from a Initiation message through the Termination message(s).

DDP Stream Session Control Message - DDP Stream Session Control messages are used to control the association of the DDP Endpoint with the DDP Stream.

Direct Data Placement Protocol (DDP) - A wire protocol that supports Direct Data Placement by associating explicit memory buffer placement information with the LLP payload units.

Lower Layer Protocol (LLP) - In the context of DDP, the protocol layer beneath RDMA that provides a reliable transport service. The SCTP DDP adaption is one of the initially defined LLPs for DDP.

Protection Domain - A common local interface convention to control which Steering Tags (STags) are valid with which DDP Endpoints. Under this convention both the Steering Tag and DDP Endpoint are created within the context of a Protection Domain, and the

Steering Tag may only be enabled for DDP Endpoints created under the same Protection Domain.

RDMA - Remote Direct Memory Access.

RNIC - RDMA Network Interface Card.

SCTP association - A protocol relationship between two SCTP endpoints. An SCTP association supports multiple SCTP streams.

SCTP Data Chunk - An SCTP Chunk used to convey Payload Data. There can be multiple Chunks within each SCTP packet. Other Chunks are used to control the SCTP Association.

SCTP endpoint - The logical sender/receiver of SCTP packets. On a multi-homed host, an SCTP endpoint is represented to its peers as a combination of an SCTP port number and a set of eligible destination transport addresses to which SCTP packets can be sent.

SCTP Stream - A uni-directional logical channel established from one to another associated SCTP endpoint. There can be multiple SCTP Streams within each SCTP association. An SCTP Stream is used to form one direction of a DDP stream.

Transmission Sequence Number (TSN) - A 32-bit sequence number used internally by SCTP. One TSN is attached to each chunk containing user data to permit the receiving SCTP endpoint to acknowledge its receipt and detect duplicate deliveries.

Upper Layer Protocol (ULP) - In the context of RDMA protocol specifications, this is the layer using RDMA services. Typically this is an application or middleware. A primary goal of RDMA protocols is to enable direct transfer of payload to/from ULP buffers.

3. 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](#) [[RFC2119](#)].

4. Introduction

4.1. Motivation

This document specifies an Adaptation Layer which fulfills the requirements of a Lower Layer Protocol (LLP) for DDP using a specific subset of SCTP capabilities.

The defined protocol is intended to be implementable over existing SCTP stacks, while clearly defining what portions of SCTP are required to enable an implementation to be optimized specifically to support DDP.

4.2. Overview

The Adaptation Layer uses a pair of like-numbered SCTP Streams within an SCTP Association to provide a reliable DDP Stream between two DDP Endpoints. Except as specifically noted, each DDP Segment submitted by the DDP Layer is encoded as a single unordered SCTP Data Chunk. In addition to the DDP Segment the Data Chunk also contains a sequence number (DDP-SSN) that reflects the order that DDP submitted the segments in for that stream.

A DDP Stream Session is defined by DDP Stream Session Control Chunks that manage the state of the DDP Stream Session. These Chunks dynamically bind DDP Endpoints to the DDP Stream Session, and DDP Segment Chunks are used to reliably deliver DDP Segments with the session.

5. Data Formats

5.1. Adaptation Layer Indicator

The DDP/SCTP Adaptation Layer uses all streams within an SCTP association. An SCTP Association that has had the DDP Adaptation Indication negotiated will carry only SCTP Data Chunks as defined in this document.

It is presumed that the handling of incoming data chunks for DDP enabled associations is sufficiently different than for routine SCTP associations that it is undesirable to require support for mixing DDP and non-DDP streams in a single association. More than a single association is required if an application desires to utilize both DDP and non-DDP traffic with the same remote host.

We define a adaptation indication which MUST appear in the INIT or INIT-ACK with the following format as defined in [ADDIP-Draft] [[I-D.ietf-tsvwg-addip-sctp](#)]

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type =0xC006           | Length = Variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Adaptation Indication           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Adaptation Indication:

The following values are suggested for DDP in this document, but the final value will be assigned by IANA:

DDP - 0x00000001

5.2. Payload Data Chunks

The DDP SCTP adaptation uses two types of SCTP Payload Data Chunks, differentiated by the Payload Protocol Identifier:

DDP Segment Chunks are used to reliably deliver DDP Segments sent between DDP Endpoints.

DDP Stream Session Control Messages are used to establish and teardown DDP Stream Sessions, specifically by controlling the binding of DDP endpoints with SCTP streams.

Payload Protocol Identifier:

The following value are defined for DDP in this document, but the final values will be assigned by IANA:

```
DDP Segment Chunk          - 16
DDP Stream Session Control - 17
```

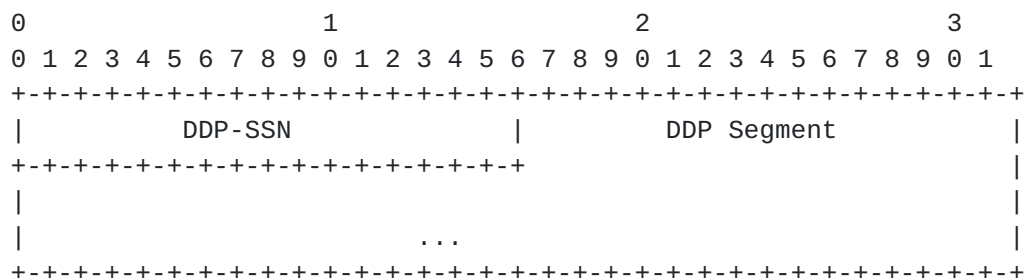
[5.2.1.](#) DDP Source Sequence Number (DDP-SSN)

All SCTP Payload Data Chunks used by this Adaptation layer include a DDP Source Sequence Number (DDP-SSN). The DDP-SSN tracks the sequence the messages were submitted to the SCTP layer for the SCTP stream in use. The DDP-SSN MUST have the same value that the SCTP Stream Sequence Number (SSN) would have been assigned had ordered SCTP Payload Data Chunks been used rather than unordered.

The rationale for specifying the DDP-SSN is as follows:

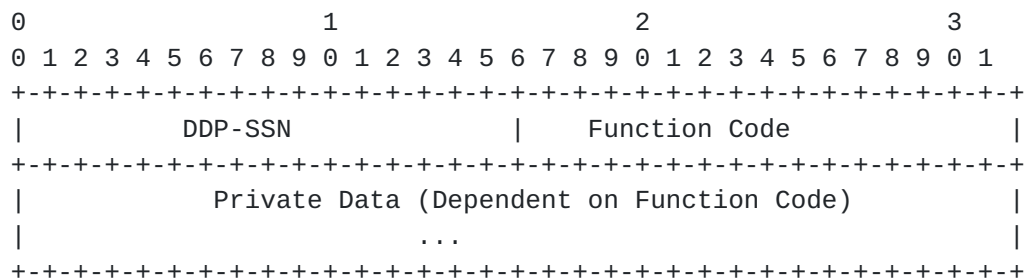
- o The SCTP Stream Sequence Number (SSN) is not suitable for this purpose, because all messages defined by this document use unordered Payload Data Chunks to ensure prompt delivery from the receiving SCTP layer.
- o The SCTP Transmission Sequence Number (TSN) is not suitable for determine the original order of Data Chunks within a stream. The sending SCTP layer is allowed to optimize the transmission sequence of unordered Data Chunks to encourage Chunk Bundling, or other purposes.

[5.2.2.](#) DDP Segment Chunk



DDP Segments are as defined in [[I-D.ietf-rddp-ddp](#)]. The DDP Segment Chunk serves the same purpose as the [[I-D.ietf-rddp-mpa](#)] Upper Layer PDU (MULPDU) in that it carries DDP Segments over a reliable protocol with added sequencing information.

5.2.3. DDP Stream Session Control



The following function code values are defined for DDP in this document:

DDP Stream Session Initiate	- 0x001
DDP Stream Session Accept	- 0x002
DDP Stream Session Reject	- 0x003
DDP Stream Session Terminate	- 0x004

ULP supplied Private Data MUST be included for DDP Stream Session Initiate, DDP Stream Session Accept and DDP Stream Session Reject messages. However, the ULP supplied Private DATA MAY be of zero length.

Private Data length MUST NOT exceed 512 bytes in any message.

Private Data MUST NOT be included in the DDP Stream Session Terminate message.

Received DDP Stream Session Control messages SHOULD be reported to the ULP. If reported, any supplied Private Data MUST be available for the ULP to examine.

The DDP/SCTP adaptation layer MAY limit the number of Session Initiate requests that it has submitted to the ULP. When a DDP Stream Session Initiate cannot be forwarded to the ULP due to such a limit the adaptation layer MUST respond with a DDP Stream Session Terminate message.

6. DDP Stream Sessions

A DDP Endpoint is the logical sender/receiver of DDP Segments. A DDP Stream connects two DDP Endpoints using a matched pair of SCTP Streams having the same SCTP Stream Identifiers.

A DDP Stream Session defines the sequence of Data Chunks exchanged between two DDP Endpoints over a DDP Stream that has a distinct beginning and end as defined in the following section. Data Chunks from one DDP Stream Session are never carried over to the next session. Each Data Chunk unambiguously belongs to exactly one session. The DDP-SSNs assigned to the Data Chunks for a session MUST NOT have any gaps.

The local interface MAY dynamically associate a DDP Endpoint with the DDP Stream based upon the initial exchanges of a DDP Session, and dynamically terminate that association at the session's end. Alternately a specialized local interface could simply statically map DDP Endpoints to DDP Streams.

Conventionally local interfaces for RDMA have deferred the selection of the DDP Endpoint until after the ULP decides to accept an RDMA connection request. But that is a local interface choice and not a wire protocol requirement.

A DDP Stream is associated with at most one Protection Domain during a single DDP Stream Session. On the passive side the association is typically deferred until the DDP Stream Session Accept message.

6.1. Sequencing

The DDP-SSN is reset to zero at the beginning of each DDP Stream Session.

The normative sequence for considering Payload Data Chunks within a given session is based upon each Data Chunk's DDP-SSN. When considered in this normative sequence, all sessions MUST conform to one the patterns defined in this section.

If the adaptation layer receives a Payload Data Chunk that conforms to none of the enumerated legal patterns the DDP Stream Session MUST be terminated.

6.2. Legal Sequence: Active/Passive Session Accepted

In this DDP Stream Session sequence one DDP Endpoint assumes the active role in requesting a DDP Stream Session, which the other side accepts.

Active Side sends a DDP Stream Session Initiate message.

Passive Side sends a DDP Stream Session Accept message.

Each side may then send zero or more DDP Segments with increasing DDP-SSNs, subject to any flow control imposed by other protocol layers.

The final User Data Chunk for each side MAY be a DDP Stream Terminate. At least one side MUST send a DDP Stream Terminate. Note that this would follow any RDMAP Terminate message, which to the Adaptation layer is simply another DDP Segment.

6.3. Legal Sequence: Active/Passive Session Rejected

DDP Stream Sessions allow each party to send a single non-payload message before the other end commits specific resources to the session. This allow each end to determine which resources are to be used, and how they are to be configured, or even if the session should be granted.

These decision MAY be influenced by the need to assign a specific Protection Domain, to determine how many RDMA Read Credits are required, or to determine now many receive operations the ULP should enable.

Because of these, or other, factors the Passive side MAY choose to reject a DDP Stream Session Request. This results in the following legal sequence:

Active Side sends a DDP Stream Session Initiate message.

Passive Side sends a DDP Stream Session Reject message.

An DDP Stream Session Reject message MUST NOT be sent unless the rejection is at the direction of the ULP.

6.4. Legal Sequence: Active/Passive Session Non-ULP Rejected

Acceptance or rejection of DDP Stream Session Initiate messages SHOULD be under the control of the ULP. This MAY require passing an event to the ULP. There MUST be a finite limit on the number of such requests that are pending a ULP decision. When more session requests are received the passive side MUST respond to the Initiate message with a DDP Stream Terminate Message.

6.5. ULP Specific Sequencing

An implementation MAY choose to support additional ULP specific sequences, but MUST NOT do so unless requested to do so by the ULP.

A defined ULP MUST be able to operate using only the defined mandatory session sequences. Any additional sequences must be used only for optional optimizations.

6.6. Other Sequencing Rules

A DDP Stream Session Control message MUST NOT be sent if it may be received before a prior DDP Stream Session Control message within the same DDP Stream Session.

An active side of a DDP Stream Session MUST NOT send a DDP Segment that might be received before the DDP Stream Session Initiate message.

This MAY be determined by SCTP acking of the Data Chunk used to carry the DDP Stream Session Initiate message, or by receipt of a responsive DDP Stream Session Control message.

A DDP Stream Identifier MUST NOT be re-used for another DDP Stream Session while any Data Chunk from a prior session might be outstanding.

7. SCTP Endpoints

7.1. Adaptation Layer Indication Restriction

The local interface **MUST** allow the ULP to specify an SCTP endpoint to use a specific Adaptation Indication. It **MAY** require the ULP to do so.

Once an endpoint decides on its acceptable Adaptation Indication(s), it **SHOULD** terminate all requests to establish an association with any different Adaptation Indication.

An SCTP implementation **MAY** choose to accept association requests for a given SCTP endpoint only until one association for the endpoint has been established. At that point it **MAY** choose to restrict all further associations for the same endpoint to use the same Adaptation Indication.

7.2. Multihoming Implications

SCTP allows an SCTP endpoint to be associated with multiple IP addresses, potentially representing different interface devices. Distribution of the logic for a single DDP stream across multiple input devices can be very undesirable, resulting in complex cache coherency challenges. Therefore the local interface **MAY** restrict DDP-enabled SCTP endpoints to a single IP address, or to a set of IP addresses that are all assigned to the same input device ("RNIC").

The default binding of a DDP enabled SCTP endpoint **SHOULD NOT** cover more than a single IP address unless doing so results in no additional bus traffic or duplication of memory registration resources. This will frequently result in a different default than for SCTP endpoints that are not DDP enabled.

Applications **MAY** choose to avoid using out-of-band methods for communicating the set of IP addresses used by an SCTP endpoint when there is potential confusion as to the intended scope of the SCTP endpoint. For example, assuming the SCTP endpoint consists of all IP Addresses advertised by DNS may work for a general purpose SCTP endpoint but not a DDP enabled one.

Even when multi-homing is supported, ULPs are cautioned that they **SHOULD NOT** use ULP control of the source address in attempt to load-balance a stream across multiple paths. A receiving DDP/SCTP implementation that chooses to support multi-homing **SHOULD** optimize its design on the assumption that multi-homing will be used for network fault tolerance, and not to load-balance between paths. This is consistent with recommended SCTP practices.

8. Number of Streams

DDP Streams are bidirectional. They are always composed by pairing the inbound and outbound SCTP streams with the same SCTP Stream Identifier.

The adaptation layer should request the maximum number of SCTP stream it will wish to use over the lifetime of the association. SCTP streams must still be bound to DDP Endpoints, and a DDP enabled SCTP association does not support ordered Data Chunks. Therefore the mere existence of an SCTP stream is unlikely to require significant supporting resources.

This mapping uses an SCTP association to carry one or more DDP Streams. Each DDP Stream will be mapped to a pair of SCTP streams with the same SCTP stream number. The adaptation MUST initialize all of its SCTP associations with the same number of inbound and outbound streams.

9. Fragmentation

A DDP/SCTP Receiver already deals with fragmentation at both the IP and DDP Layers. Therefore use of SCTP layer segmenting will be avoided for most cases.

As a Lower Layer Protocol (LLP) for DDP, the SCTP adaptation layer MUST inform the DDP layer of the maximum DDP Segment size that will be supported. This should be the largest value that can be supported without use of IP or SCTP fragmentation, or 516 bytes, whichever is larger.

A minimum of 516 bytes is required to allow a DDP Stream Session Control Message with 512 bytes of Private Data.

SCTP data chunk fragmentation MUST NOT be used unless the alternative is IP fragmentation.

The SCTP adaptation layer SHOULD set the maximum DDP Segment size below the theoretical maximum in order to allow bundling of Control Chunks in the same SCTP packet.

The SCTP adaptation layer MUST reject DDP Segments that are larger than the maximum size specified.

10. Sequenced Unordered Operation

The Adaptation layer MUST use the Unordered option on all Data Chunks (U Flag set to one). The SCTP Layer is expected to deliver unordered Data Chunks without delay.

Because DDP employs unordered SCTP delivery, the receiver MUST NOT rely upon the SCTP Transmission Sequence Number (TSN) to imply ordering of DDP Segments. The fact that the SCTP Data Chunk for a DDP Segment is prior the cumulative ack point does not guarantee that all prior DDP segments have been placed. The SCTP sender is not obligated to transmit unordered Data Chunks in the order presented.

The DDP-SSN can be used without special logic to determine the submission sequence when the maximum number of in-flight messages is less than 32768. This also applies if the sending SCTP accepts no more than 32767 Data Chunks for a single stream without assigning TSNS.

If SCTP does accept more than 32768 Data chunks for a single stream without assigning TSNS, the sending DDP must simply refrain from sending more than 32767 Data Chunks for a single stream without acknowledgment. Note that it MUST NOT rely upon ULP flow control for this purpose. Typical ULP flow control will deal exclusively with untagged messages, not with DDP segments.

The receiver MAY perform a validity check on received DDP-SSNs to ensure that any gap could be accounted for by unreceived Data Chunks. Implementations SHOULD NOT allocate resources on the assumption that DDP-SSNs are valid without first performing such a validity check. An invalid DDP-SSN MAY result in termination of the DDP Stream.

11. Procedures

11.1. Association Initialization

At the startup of an association, a DDP/SCTP Adaptation Layer MUST include an adaptation layer indication in its INIT or INIT-ACK (as defined in [Section 5.1](#). After the exchange of the initial first two SCTP chunks (INIT and INIT-ACK), an endpoint MUST verify and inspect the adaptation indication and compare it to the following table to determine proper action.

Indication type	Action
+++++	
NONE	This indicates that the peer DOES NOT support ANY DDP or RDMA adaptation and thus RDMA and DDP procedures MUST NOT be performed upon this association.
+++++	
DDP	This indicates that the peer DOES support the DDP/SCTP Adaptation Layer defined here.
+++++	
ANY-OTHER Indication	This indicates that the peer DOES NOT support the DDP adaptation and thus DDP procedures MUST NOT be performed upon this association.
+++++	

An implementation MAY require that all associations for a given SCTP endpoint be placed in the same mode.

The local interface MAY allow the ULP to accept only requests to establish an association in a specified mode.

11.2. Chunk Bundling

SCTP allows multiple Data Chunks to be bundled in a single SCTP packet. Data chunks containing DDP Segments with untagged messages SHOULD NOT be delayed to facilitate bundling. Data chunks containing DDP Segments with tagged messages will generally be full sized, and hence not subject to bundling. However partial size tagged messages MAY be delayed, as that they are frequently followed by a short untagged message.

11.3. Association Termination

Termination of an SCTP Association due to errors should be handled at the SCTP layer. The RDMAP defined RDMAP Terminate Message SHOULD NOT be sent on each DDP Stream when a determination has been made to terminate an SCTP association. Sending that message on each SCTP stream could severely delay the termination of the association.

The local interface SHOULD notify all consumers of DDP streams when the underlying SCTP stream has been terminated.

Other RDMAP defined Terminate Messages MUST be generated as specified when a DDP Stream is terminated. Note that with the SCTP mapping, termination of a DDP Stream does not mandate termination of the Association.

12. IANA considerations

This document defines a new SCTP Adaptation Layer Indication codepoint. [[I-D.ietf-tsvwg-addip-sctp](#)] will create the registry from which this codepoint is to be assigned. Any unallocated codepoint may be assigned. The value of one is suggested.

This document also defines two new SCTP Payload Protocol Identifier (PPIDs). [RFC2960](#) [[RFC2960](#)] creates the registry from which these identifiers are to be assigned. The Payload Protocol Identifiers allocated need to be unique, but have no other requirements. The following values are suggested:

DDP Segment Chunk	- 16
DDP Stream Session Control	- 17

13. Security Considerations

Any direct placement of memory could pose a significant security risk if adequate local controls are not provided. These threats are addressed in the appropriate DDP [[I-D.ietf-rddp-ddp](#)], RDMA [[I-D.ietf-rddp-rdmap](#)] or Security [[I-D.ietf-rddp-security](#)] drafts. This document does not add any additional security risks over those found in [RFC2960](#) [[RFC2960](#)].

The IPsec requirements for RDDP are based on the version of IPsec specified in [RFC 2401](#) [[RFC2401](#)] and related RFCs, as profiled by [RFC 3723](#) [[RFC3723](#)], despite the existence of a newer version of IPsec specified in [RFC 4301](#) [[RFC4301](#)] and related RFCs. One of the important early applications of the RDDP protocols is their use with iSCSI iSER [[I-D.ietf-ips-iser](#)]; RDDP's IPsec requirements follow those of IPsec in order to facilitate that usage by allowing a common profile of IPsec to be used with iSCSI and the RDDP protocols. In the future, [RFC 3723](#) may be updated to the newer version of IPsec, the IPsec security requirements of any such update should apply uniformly to iSCSI and the RDDP protocols.

14. Contributors

Many thanks to our contributors who have spent many hours reading and reviewing keeping us straight: Sukanta Ganguly an independent consultant, Vivek Kashyap of IBM, Jim Pinkerton of Microsoft, and Hemal Shah of Broadcom. Thanks for all your hard work.

15. Acknowledgments

The authors would like to thank the following people that have provided comments and input: Stephen Bailey, David Black, Douglas Otis, Allyn Romanow and Jim Williams.

16. References

16.1. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2960] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L., and V. Paxson, "Stream Control Transmission Protocol", [RFC 2960](#), October 2000.
- [RFC3723] Aboba, B., Tseng, J., Walker, J., Rangan, V., and F. Travostino, "Securing Block Storage Protocols over IP", [RFC 3723](#), April 2004.
- [I-D.ietf-rddp-ddp]
Shah, H., "Direct Data Placement over Reliable Transports", [draft-ietf-rddp-ddp-06](#) (work in progress), June 2006.
- [I-D.ietf-rddp-rdmap]
Recio, R., "A Remote Direct Memory Access Protocol Specification", [draft-ietf-rddp-rdmap-07](#) (work in progress), September 2006.
- [I-D.ietf-rddp-security]
Pinkerton, J., "DDP/RDMAP Security", [draft-ietf-rddp-security-10](#) (work in progress), June 2006.
- [I-D.ietf-tsvwg-addip-sctp]
Stewart, R., "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", [draft-ietf-tsvwg-addip-sctp-15](#) (work in progress), June 2006.

16.2. Informative references

- [I-D.ietf-rddp-mpa]
Culley, P., "Marker PDU Aligned Framing for TCP Specification", [draft-ietf-rddp-mpa-06](#) (work in progress), September 2006.
- [RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.

[I-D.ietf-ips-iser]

Ko, M., "iSCSI Extensions for RDMA Specification",
[draft-ietf-ips-iser-05](#) (work in progress), October 2005.

Authors' Addresses

Caitlin Bestler
Editor
Broadcom Corporation
16215 Alton Parkway
P.O. Box 57013
Irvine, CA 92619-7013
USA

Phone: 949-926-6383
Email: caitlinb@broadcom.com

Randall R. Stewart
Editor
Cisco Systems, Inc.
Forest Drive
Columbia, SC 29036
USA

Phone: +1-815-342-5222
Email: rrs@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

