Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 4, 2013

Murtaza S. Chiba
Alexander Clemm
Steven Medley
Joseph Salowey
Sudhir Thombare
Eshwar Yedavalli
Cisco Systems
October 1, 2012

Cisco Service Level Assurance Protocol
draft-cisco-sla-protocol-03

Abstract

Cisco's Service Level Assurance Protocol (Cisco's SLA Protocol) is a Performance Measurement protocol that has been widely deployed. The protocol is used to measure service level parameters such as network latency, delay variation, and packet/frame loss. This draft describes the Cisco SLA Protocol UDP measurement type to enable vendor interoperability.

Requirements Language

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 RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on April 4, 2013.

Copyright Notice

Table of Contents

1. Introduction 3
2. Protocol 3
2.1. Control Phase 5
2.1.1. Control Request 6
2.1.1.1. Command Header 7
2.1.1.2. CSLDs 9
2.1.2. Control Response Message 14
2.2. Measurement Phase 15
3. Implementation notes 19
4. Extensions 20
5. IANA Considerations 21
6. Security Considerations 23
6.1. Message Authentication 23
6.2. IPSec Considerations 23
6.2.1. Control Traffic 24
6.2.2. Measurement Traffic 24
6.3. Replay Protection 24
7. Terminology 25
8. Acknowledgements 25
9. References 25
9.1. Normative References 25
9.2. Informative References 26
Authors' Addresses 26

1. Introduction

Network performance measurements are becoming critical data points for administrators monitoring the health of the network. As Service Providers look to differentiate their offerings, performance measurement is increasingly becoming an important tool to monitor Service Level guarantees and, in general, to monitor the health of a network.

Performance metrics, both one-way and two-way, can be used for pre-deployment validation as well as for measuring in-band live network performance characteristics. It can be used to measure service levels in L2 and L3 networks as well as for applications running on top of L3. Performance measurements are gathered by analyzing actively generated synthetic request and response packets/frames. This is in contrast to passive measurements that analyze production traffic flowing through a particular network element.

There is a growing body of work on Performance Measurement standards that enable interoperability between different vendors network elements by describing common measurement protocols as well as metrics. IETF has actively developed standards on the subject and two such standards are One-Way Active Measurement Protocol(OWAMP) [RFC4656] and Two-Way Active Measurement Protocol (TWAMP) [RFC5357].

Cisco's SLA Protocol is another example of a performance measurement protocol that offers a rich set of measurement message types. The measurement types can be classified as those that test connectivity (ping like) by providing round trip or, one-way latency measures and those that provide a richer set of statistics including network jitter and packet/frame loss. Each type of active measurement exchanges mimic an actual protocol exchange.

Cisco's SLA Protocol UDP measurement message exchanges, as covered in this document to enable interoperability, simulates a UDP application and can be used to simulate either Voice or Video traffic that is encoded in RTP frames within UDP envelopes. The UDP measurement type message exchanges carry information that provide the ability to derive a robust set of statistics.

2. Protocol

The Cisco Service Level Assurance Protocol consists of two distinct phases, the Control phase and the Measurement phase. Each phase is comprised of exchange of information between a network element acting as the Sender and another element designated as the Responder.

The Control Phase is the first phase of message exchanges and forms the base protocol. This phase establishes the identity of the Sender and provides information for the Measurement Phase. A single message pair of Control Request and Control Response marks this phase. The Sender initiates a Control Request message that is acknowledged by the Responder with a Control Response message. The Control Request may be sent multiple times if a Control Response has not been received; the number of times the message is retried is configurable on the Sender element.

The Measurement Phase forms the second phase and is comprised of a sequence of Request/Response messages. These messages may be exchanged as often as required. Each Measurement Request message is acknowledged by the Responder with a Measurement Response Message.

The number and frequency with which messages are sent SHOULD be controlled by configuration on the Sender element, along with the waiting time for a Control Response.

The following sequence diagram depicts the message exchanges:

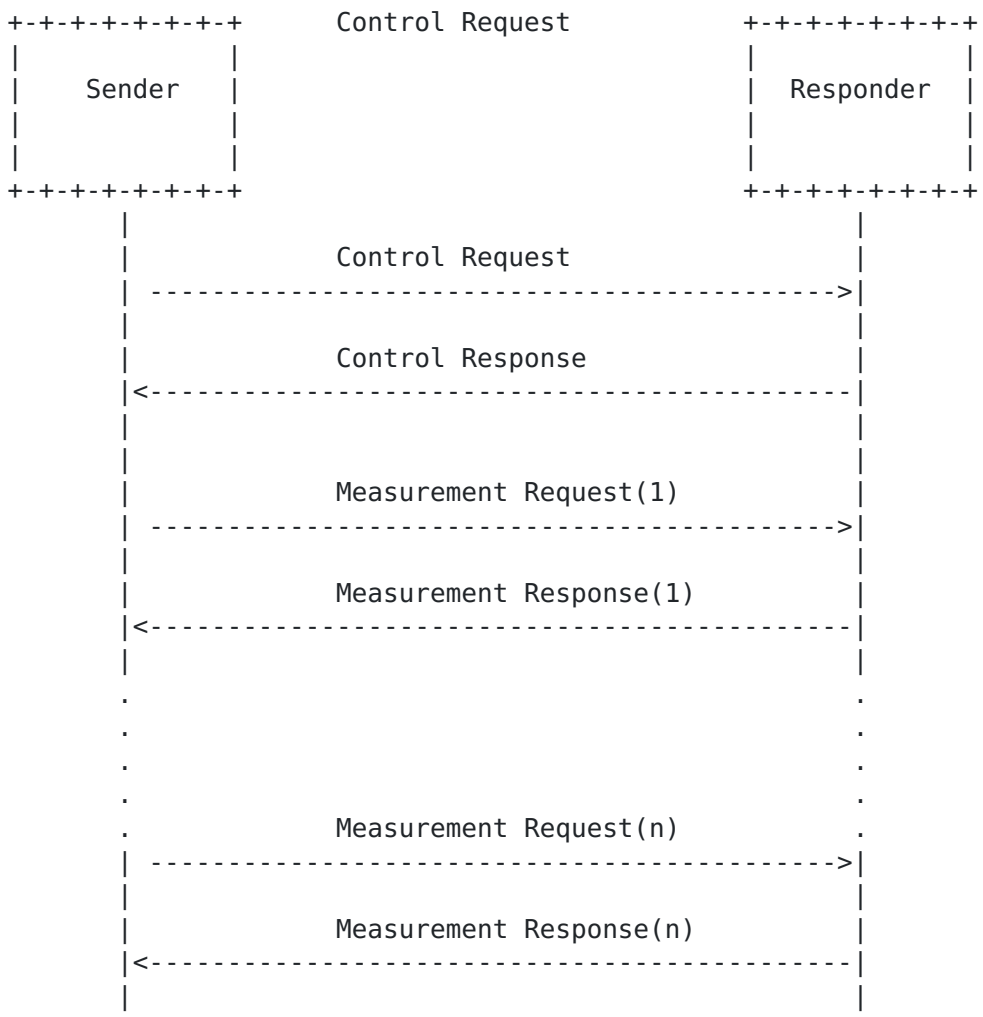

2.1. Control Phase

The Control Phase begins with the Sender sending a Control Request message to the Responder. The Control Request message is sent to UDP port 1167 on the Responder requesting a measurement phase UDP port be opened and, in addition, indicates the requested amount of time that the port needs to be opened for. The Responder replies by sending a Control Response with appropriate Status indicating Success when the sender identity is verified (if used) and the requested UDP port was successfully opened. In all other cases a non-zero Status is returned.

The sequence of exchanges is as indicated in the diagram.

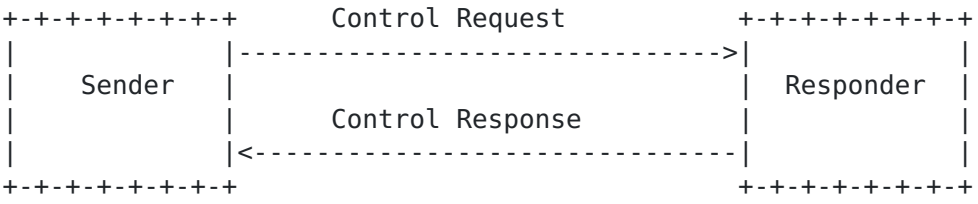

2.1.1. Control Request

The Control Request message consists of a Command Header followed by one or more Command, Status, Length and Data sections (henceforth known as CSLD). At the minimum, there SHOULD be at the least two CSLD sections, one of which is the authentication CSLD section and the other carries information for the Measurement Phase simulation type.

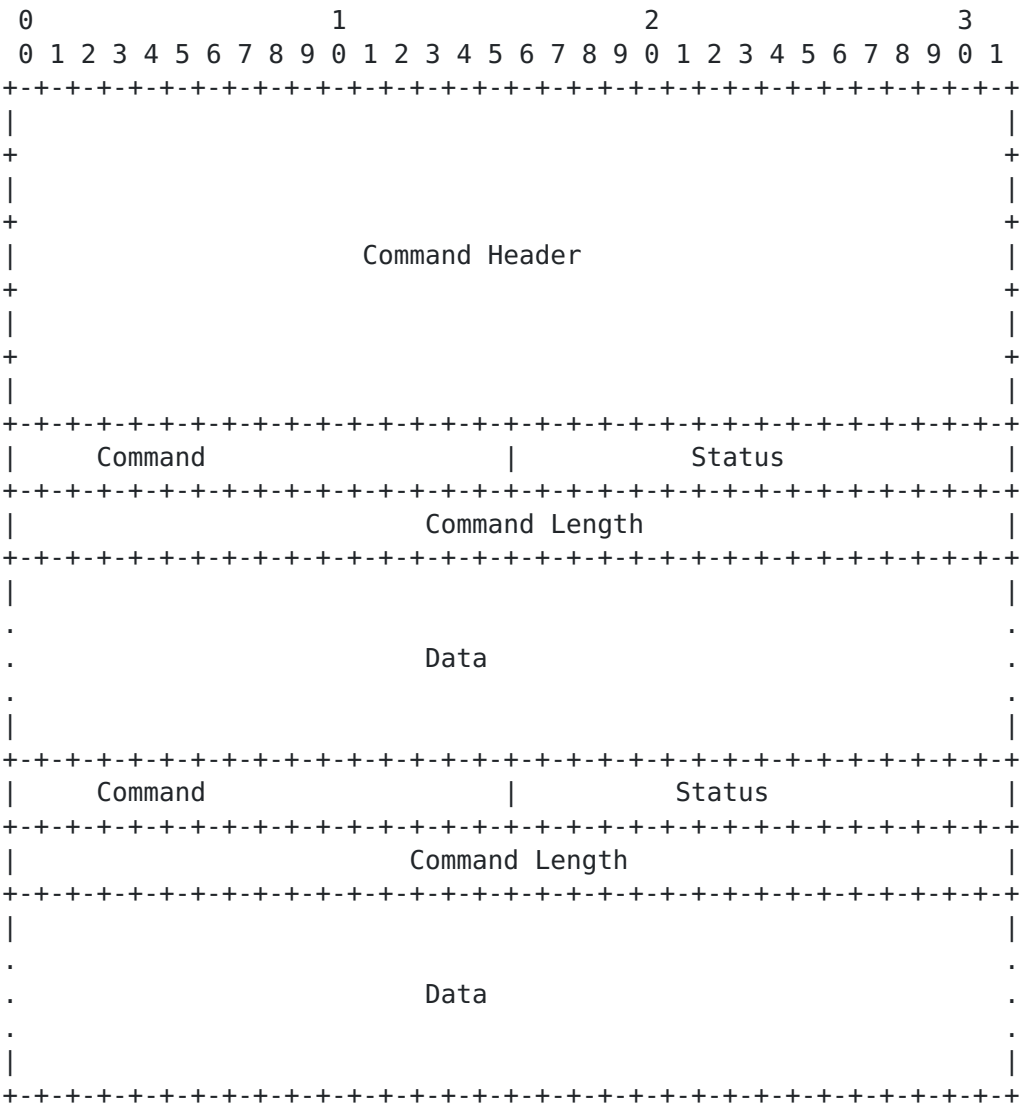

2.1.1.1. Command Header

The Command Header is the first section of the Control Request message and is depicted below:

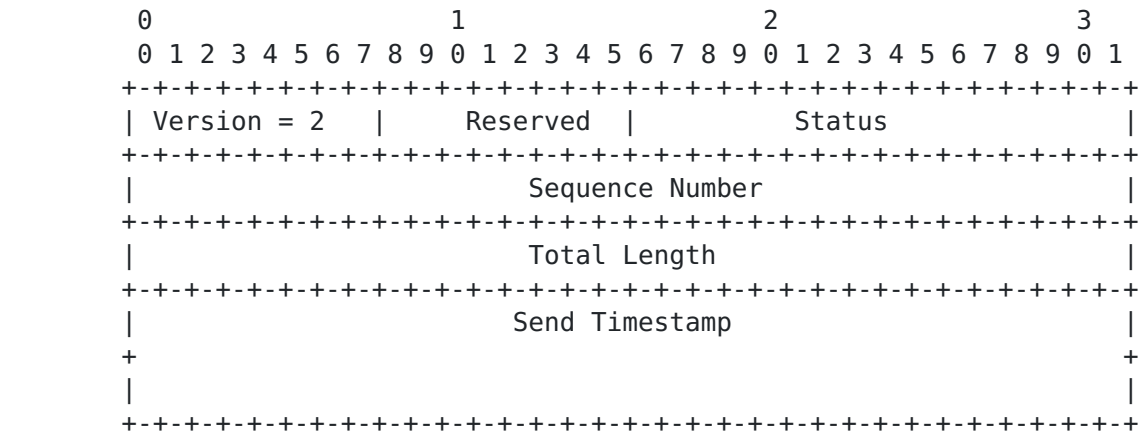

The Command Header fields hold the following meaning:

Field	Size (bits)	Usage
Version	8	Current version supported and is to be set to 2.
Reserved	8	Reserved field, MUST be set to 0
Status	16	Indicates success or failure for the entire message; not used for request and MUST be set to 0
Sequence Number	32	Used to map requests to responses. This is a monotonically increasing number. Implementations MAY reset the sequence number to 0 after a reboot, and SHOULD wrap around after all bits have been exceeded.
Total Length	32	Carries the total length of the control message in number of octets
Send Timestamp	64	This field is set to the time the command was submitted for transmission and is updated for a response. This field MAY be used when security is of concern in order to prevent replay attacks. SHOULD be updated for a response. When not being used it MUST be set to all 0's. The format is as given in RFC5905

The sequence number field MUST include a new number for each new request and is monotonically increasing. When the Control Request is to be retried, the sequence number MUST remain unchanged.

2.1.1.2. CSLDs

The two CSLDs to be included, in order, along with the Command Header are:

- The Authentication CSLD
- A Measurement Type CSLD

In this revision of the protocol, only a single Measurement Type CSLD has been defined, the UDP Measurement Type CSLD. For future extensions it is possible to add additional Measurement Type CSLDs. For more details please see the section on Extensions.

2.1.1.2.1. Authentication CSLD

The Authentication CSLD provides the message authentication and verifies the requester knows the shared-secret. The following is the format for the Authentication CSLD

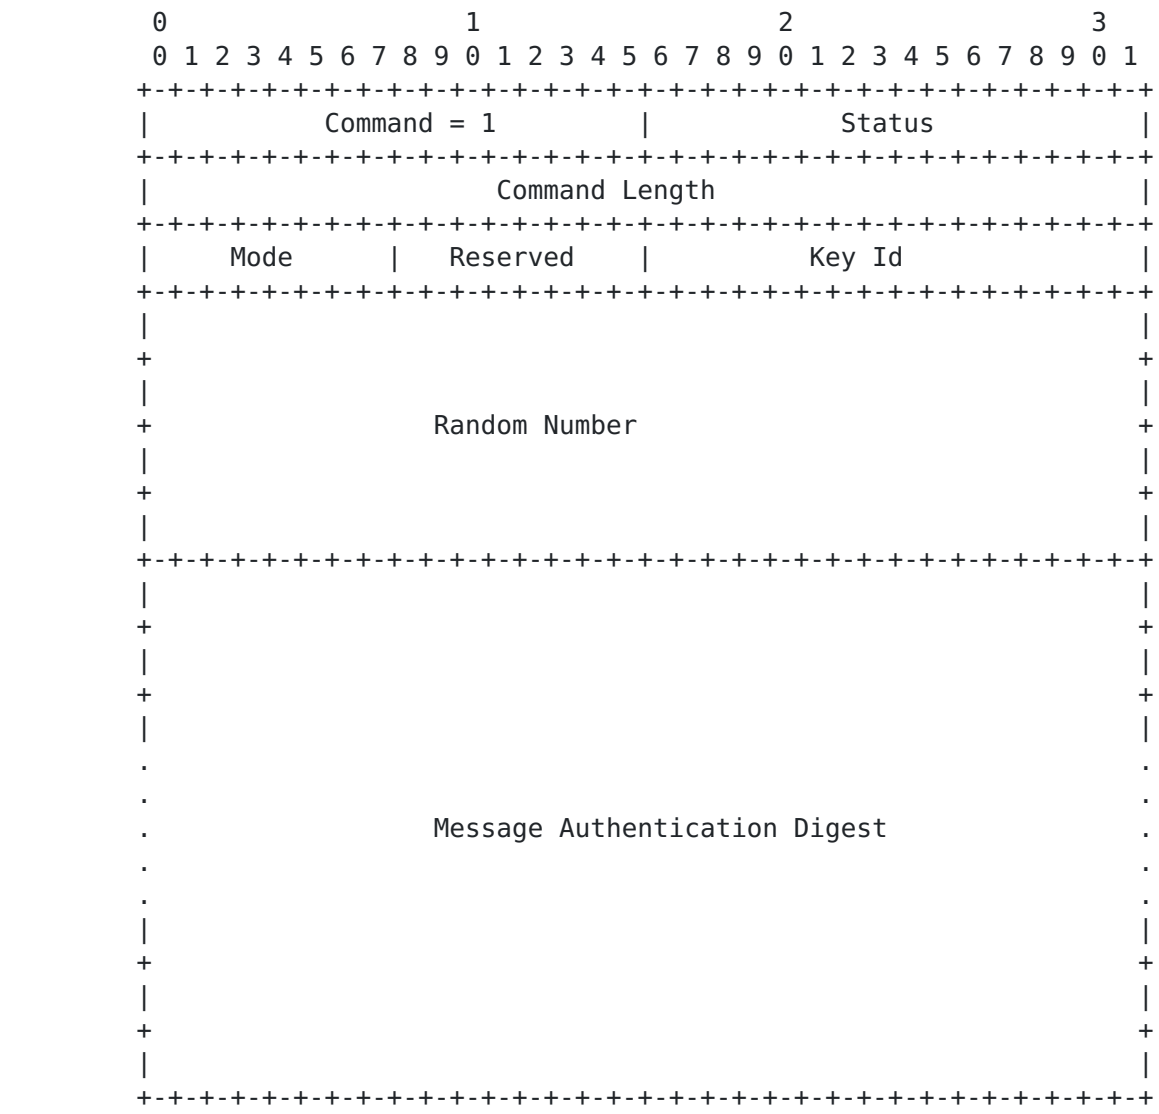

The fields for the Authentication CSLD have the following meaning

Field	Size (bits)	Description
Command	16	Indicates the CSLD is of type Authentication
---	---	---
Status	16	Not used for a request and MUST be set to 0
---	---	---
Command Length	16	Indicates the length of the CSLD
---	---	---
Mode	8	Indicates the type of authentication being used and is set as follows: 0 - No Authentication, 1 - SHA256 Authentication, 2 - HMAC-SHA-256
---	---	---
Reserved	8	This field is reserved for future extensions and MUST be set to 0
---	---	---
Key ID	16	Indicates the index number of the shared-secret to be used for authenticating the Control Request Message
---	---	---
Random Number	128	This field is to be unique over the shared secret life and is used to make it difficult to predict the shared secret via multiple packet captures. The value is reflected in a response message. This field MAY be used when security is of concern and is useful to prevent dictionary attacks. When not being used it should be set to all 0's
---	---	---
Message Authentication Digest	256	Contains the message authentication digest and is computed over the entire control packet including this field set to all 0s

2.1.1.2.2. UDP Measurement CSLD

The UDP Measurement CSLD indicates the Measurement Type to be used during the Measurement Phase and specifies the addresses and UDP port to be opened as well as the duration the port has to be kept open for the measurement phase. The format of the CSLD is as follows:

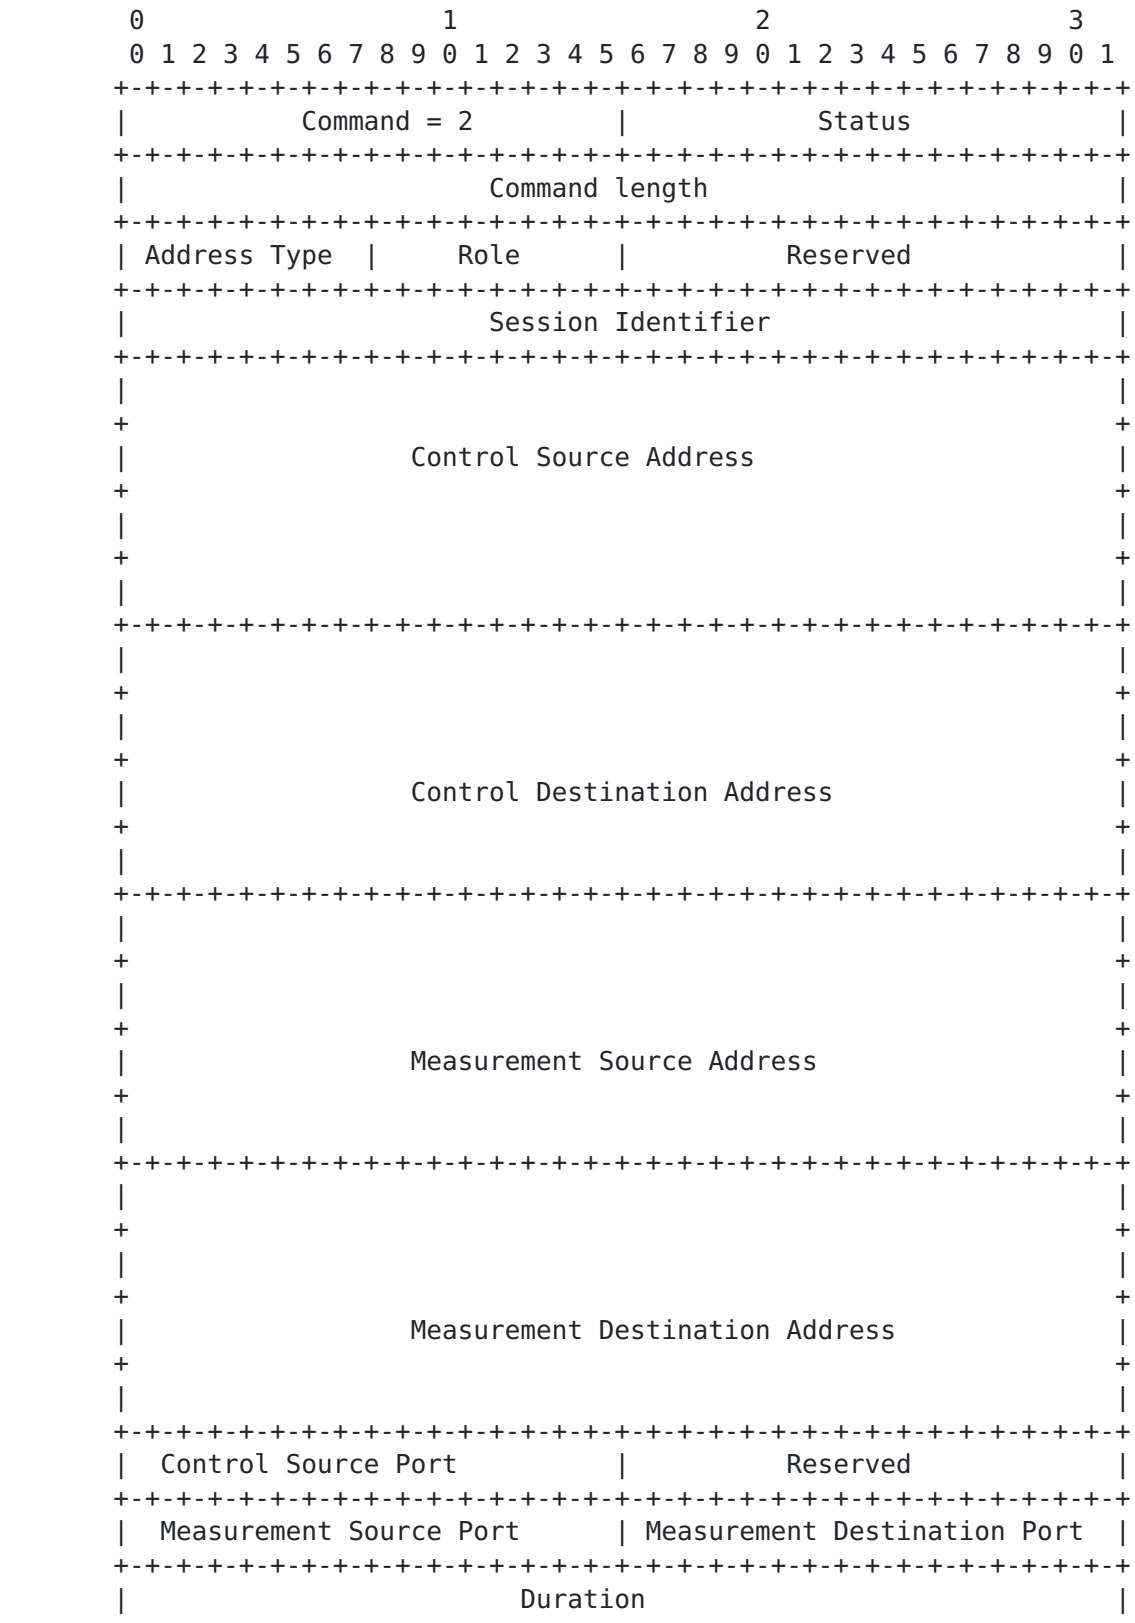
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Command = 2
Status
Command length
Address Type
Role
Reserved
Session Identifier
Control Source Address
Control Destination Address
Measurement Source Address
Measurement Destination Address
Control Source Port
Reserved
Measurement Source Port
Measurement Destination Port
Duration

Note: Duration is specified in seconds.

The fields in the UDP Measurement CSLD have the following meaning:

Field	Size (bits)	Definition
Command	16	Indicates that the CSLD is to simulate UDP traffic measurements.
---	---	---
Status	16	Not used for a request and MUST be set to 0
---	---	---
Command Length	32	Indicates the length of the CSLD
---	---	---
Address Type	8	Indicates the address type and is set to one of the values in the Cisco Service Level Performance Measurement Protocol - Address Registry. (1 - IPv4 addresses, 2 - IPv6 addresses)
---	---	---
Role	8	Indicates the role of the endpoint receiving the control message and is set as follows: 1 - Responder
---	---	---
Reserved	16	Reserved and MUST be set to 0
---	---	---
Session Identifier	32	Carries a session identifier that is a locally significant unique value to the originator of the message. MUST be 0 when not specified.
---	---	---
Control Source Address	128	Set to the address from which the Sender initiates control messages. For IPv4 addresses only the first 32 bits are filled and the remaining bits MUST be set to 0
---	---	---
Control Destination Address	128	Set to the address on the Responder where the control message will be sent. For IPv4 addresses only the first 32 bits are filled and the remaining bits MUST be set to 0
---	---	---

Measurement Source Address	128	Set to the address of the Sender from where the measurement packets will originate. For IPv4 addresses only the first 32 bits are filled and the remaining bits MUST be set to 0
Measurement Destination Address	128	Set to the address on the Responder towards which the measurement packets will be sent and is a way to identify an ingress interface on the Responder. For IPv4 addresses only the first 32 bits are filled and the remaining bits MUST be set to 0
Control Source Port	16	Indicates the port on the Sender from which Control message is sent. If unset the value should be derived from the incoming packet.
Reserved	16	Reserved Field, MUST be set to 0.
Measurement Source Port	16	Indicates the UDP Port on the Sender from which the measurement packets will be sent
Measurement Destination Port	16	Indicates the UDP Port on the Responder towards which the measurement packets will be sent
Duration	32	This is the duration in seconds the port needs to be kept open for accepting measurement phase messages. Measurement messages received after the duration MUST be ignored

Note: The source addresses are only indicative of identity of the originator and cannot be used as destination address for responses in a NAT environment.

2.1.2. Control Response Message

In response to the Control Request Message the network element designated the Responder sends back a Control Response Message that reflects the Command Header with an updated Status field and includes the two CSLD sections that also carry updated Status fields. Hence, the format is identical to the Control Request message as described above.

Following table shows the supported values of the Status fields:

Status Value	Description
0	Success
1	Fail - catch all
2	Authentication Failure
3	Format error - sent when any CSLD type is not recognized or any part of a CSLD has a value that is not recognized
4	Port in use - the UDP/TCP port is already being used by some other application and cannot be reserved
5+	Future extension and experimental values, please refer to Status Types Registry in the IANA Considerations section

The Command Header Status indicates Success only if all the CSLD sections have Status as Success. It is non-zero otherwise. Future extensions MAY extend these values as appropriate.

The Control Response message, besides the update of the Status fields, SHOULD also update the Sent Timestamp (if used) in the Command Header and the Message Authentication Digest in the Authentication CSLD. The Message Authentication Digest is computed in the same way as the Control Request message. The Random Number field MUST be reflected without modification. The Session Identifier MAY be updated to reflect a locally significant unique value, MUST be 0 if not specified.

2.2. Measurement Phase

Upon receiving the Control Response message with the Status set to Success, the second phase of the protocol, the Measurement Phase, is initiated. In all other cases when the Status is not success no measurement traffic is initiated. In the Measurement Phase the Sender sends a stream of measurement messages. The measurement message stream consists of packets/frames that are spaced a configured number of milliseconds apart.

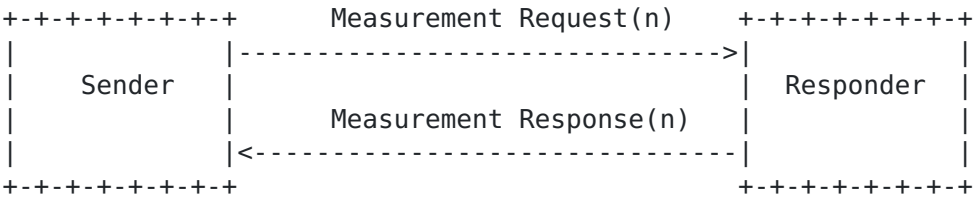

The format of the Measurement messages as defined by this document for UDP Measurements is as shown below and is the same for the exchange in both directions, that is the format is the same when sent from the Sender to the Responder and when sent back from the Responder to the Sender with the only difference being the update of those fields that are designated with the Responder prefix, all other fields MUST remain unchanged.

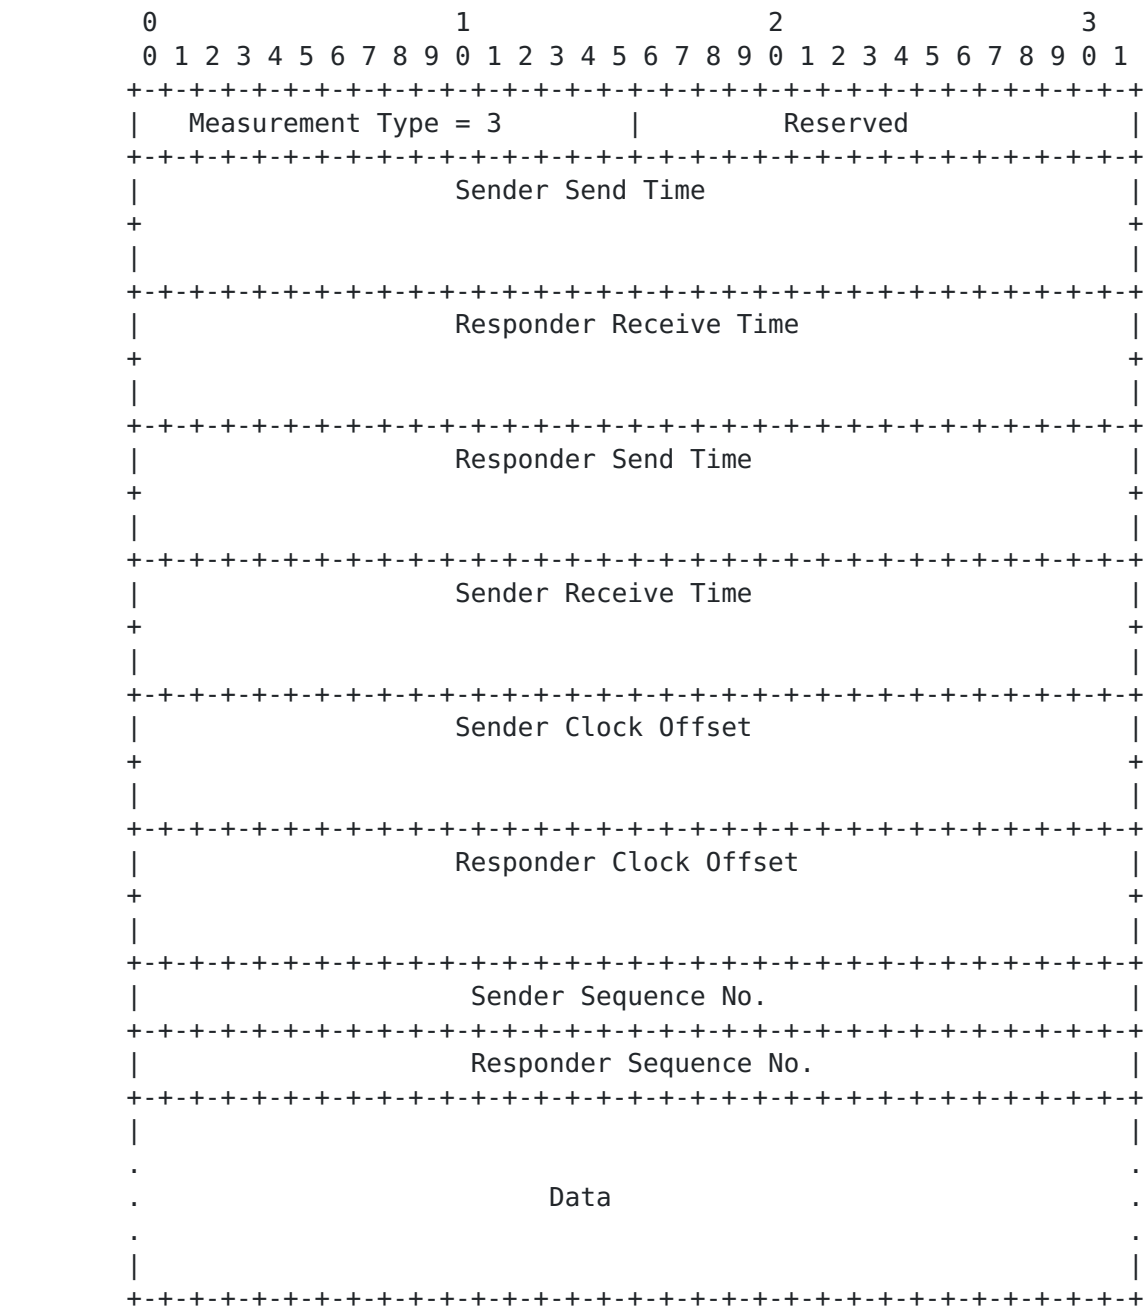

The fields for the UDP Measurement Request have the following meaning:

Field	Size (bits)	Description
Measurement Type	16	Carries the type of measurement being performed; 1 - Reserved, 2 - Reserved, 3 - UDP
---	---	---
Reserved	16	Reserved field and MUST be set to 0
---	---	---
Sender Send Time	64	Carries the timestamp when the measurement message was submitted for transmission by Sender
---	---	---
Responder Receive Time	64	Carries the timestamp when the measurement message was received by Responder
---	---	---
Responder Send Time	64	Carries the timestamp when the measurement message was submitted for transmission by the Responder. It MUST be 0 in the Sender to Responder direction
---	---	---
Sender Receive Time	64	Carries the timestamp when the Sender received the measurement message. It MUST be 0 in both directions on the wire and is filled on the Sender side as soon as the measurement message is received
---	---	---
Sender Clock Offset	64	Gives an estimate of the Sender clock skew measured in second and fractional seconds
---	---	---
Responder Clock Offset	64	Gives an estimate of the Responder clock skew measured in seconds and fractional seconds
---	---	---
Sender Sequence Number	32	The sequence number of the measurement message on the Sender side. This field is monotonically increasing and MAY wraparound
---	---	---
Responder Sequence Number	32	The sequence number of the measurement message on the Responder side. This field is monotonically increasing and MAY wraparound

| --------- | --------- | -------------------------- |
| Data | 32 bit aligned | This field is used to pad up to the configured request data size. The minimum requested data size SHOULD be 512 bytes and this field will be of length 512 minus the length of the previous fields |

Note: All timestamps have the default format as described in RFC 5905 [RFC5905] and is as follows: the first 32 bits represent the unsigned integer number of seconds elapsed since 0h on 1 January 1900; the next 32 bits represent the fractional part of a second thereof. The timestamp definition is also similar to RFC 4656 [RFC4656]

In addition, the timestamp format used can be as described for the low-order 64 bits of the IEEE 1588-2008 (1588v2) Precision Time Protocol timestamp format [IEEE1588]. This truncated format consists of a 32-bit seconds field followed by a 32-bit nanoseconds field, and is the same as the IEEE 1588v1 timestamp format. This timestamp definition is similar to the default timestamp as specified in RFC 6374 [RFC6374]

Implementations MUST use only one of the two formats. The chosen format is negotiated out-of-band between the endpoints or defaults to the format as defined in RFC 5905. [RFC5905]

3. Implementation notes

Responder implementations SHOULD support simultaneous measurements destined to a single port either from the same or a different Sender. For different measurement instances that originate from the same sender, there MUST be a clear method for the Responder to distinguish the traffic, for example per a unique 5-tuple of protocol, source address, source port, destination address and destination port.

A Control Request that is received for the same measurement request as identified by the 5-tuples, for instance, SHOULD result in the resetting of the duration timer as well as the Responder Sequence Number.

A Control Phase followed by the Measurement Phase can be repeated in order to have a continuous measurement over the entire time a device is alive.

The Random Number field in the Measurement packets is to be set to a random value in environments where security is a concern and is used

to prevent dictionary attacks. It MUST always be included, when not used it MUST be set to all 0s.

The Authentication CSLD MUST always be included. When the mode field is set to 0, the Random Number field and the Message Authentication Digest MUST both be set to all 0s. For the SHA256 authenticator mode the shared secret is prepended to the Control Message and the authentication algorithm is then run over the complete data including the shared secret. The SHA256 mode is included for ease of implementation and it is recommended to use the HMAC variant to afford better security.

If the UDP port indicated in the UDP Measurement CSLD is busy, the Responder MAY suggest an alternative port, the Status of the UDP Measurement CSLD MUST be set to Success in that case. The Sender MAY set a value of 0 in the field, in which case the Responder MAY choose to open a port and send that back along with the Status of Success. It should be noted that this behavior has security ramifications and the port needs to be chosen very carefully by the Responder.

The measurement stream typically consists of packets/frames with a periodic inter-packet distribution. The Sender need not wait for a Measurement Response packet to arrive before sending another Measurement Request packet, and in many cases it will not be possible in order to maintain the desired inter-packet distribution.

The default format for all timestamps is as specified in RFC 5905. [RFC5905]

All messages and all fields within a message are assumed to be in network order. In addition, all data fields are unsigned unless mentioned otherwise.

4. Extensions

This section describes how the protocol can be extended to allow for additional functionality, such as new types of measurements.

In order to allow for new types of measurements, additional Measurement Type CSLDs can be defined to be carried within the Control Request and Control Response messages in place of the UDP Measurement CSLD defined in this document . The meaning and precise format of such CSLD needs to be defined in a separate specification. Such a specification will also need to describe the appropriate formats for the messages in the Measurement Phase.

In addition, the protocol can be extended by adding support for new

values to registries defined in this document.

5. IANA Considerations

The following registries are needed for the extensibility of the protocol.

Cisco Service Level Assurance Protocol - Version Number Registry

Version	Description
1	Reserved
2	Version for protocol in this document
3 - 200	Available or future extensions
201 - 255	Experimental values for private use

The version number should only be changed when the structure of the Command Messages is different from the basic Command Header and CSLD structure described in this document.

Cisco Service Level Assurance Protocol - CSLD Command Registry

CSLD Type	Description
1	Authentication CSLD
2	UDP Measurements
3 - 52	Reserved
53 - 2047	Available for future extensions
2048+	Experimental values for private use

It is envisioned that future documents will provide their own measurement type number along with the format of the Data portion.

Cisco Service Level Assurance Protocol - Status Types Registry

Status	Description
0	Success
---	---
1	Fail - catch all
---	---
2	Authentication failure
---	---
3	Format error - sent when any CSLD type is not recognized or any part of a CSLD has a value that is not recognized
---	---
4	Port in use - the UDP/TCP port is already being used by some other application and cannot be reserved
---	---
5 - 4095	Available for future extensions
---	---
4096+	Experimental values for private use

Cisco Service Level Assurance Protocol - Authenticator Modes Registry

Mode	Description
0	No Authentication
1	SHA256
2	HMAC-SHA-256
3 - 200	Available for future extensions
201 - 255	Experimental values for private use

Cisco Service Level Assurance Protocol - Address Registry

Address Type	Description
1	IPv4
2	IPv6
3 - 200	Available for future extensions
201 - 255	Experimental values for private use

Cisco Service Level Assurance Protocol - Roles Registry

Role	Description
1	Sender
2	Responder
3 - 2047	Available for future extensions
2048+	Experimental values for private use

Cisco Service Level Assurance Protocol - Measurement Type Registry

Measurement Type	Description
1	Reserved
2	Reserved
3	UDP
4 - 52	Reserved
53-2047	Available for future extensions
2048+	Experimental values for private use

6. Security Considerations

6.1. Message Authentication

When the mode for the Authentication CSLD is set to 1, the Message Authentication Digest is generated using the SHA 256 algorithm and is to be calculated over the entire packet including the Message Authentication Digest field which MUST be set to all 0s.

When the mode for the Authentication CSLD is set to 2, the Message Authentication Digest is generated using the HMAC-SHA-256 as described in RFC 4868 [RFC4868]algorithm and is to be calculated over the entire packet including the Message Authentication Digest field which MUST be set to all 0s

When the mode field is set to 0, the Random Number field and the Message Authentication Digest MUST both be set to all 0s.

6.2. IPSec Considerations

It is RECOMMENDED that IPSec be employed to afford better security. IPSec provides enhanced privacy as well as an automated key distribution mechanism. The following recommendations are similar to RFC3579, Section 2 [RFC3579]

6.2.1. Control Traffic

For Senders implementing this specification, the IPSec policy would be "Initiate IPSec, from me to any, destination port UDP 1167". This causes the Sender to initiate IPSec when sending Control traffic to any Responder. If some Responders contacted by the Sender do not support IPSec, then a more granular policy will be required, such as "Initiate IPSec, from me to IPSec-Capable-Responder, destination port UDP 1167".

For Responders implementing this specification, the IPSec policy would be "Require IPSec, from any to me, destination port UDP 1167". This causes the Responder to require use of IPSec. If some Sender does not support IPSec, then a more granular policy will be required: "Require IPSec, from IPSec-Capable-Sender to me".

6.2.2. Measurement Traffic

As the Control Phase occurs before the Measurement Phase, it should be possible to build an IPSec Security Association once a successful Control Response is received.

For Senders implementing this specification, the IPSec policy would be "Initiate IPSec, from me to negotiated address, destination is negotiated port". This causes the Sender to initiate IPSec when sending Measurement traffic to the Responder. If some Responders contacted by the Sender do not support IPSec, then a more granular policy will be required, such as "Initiate IPSec, from me to IPSec-Capable-Responder, destination is negotiated port".

For Responders implementing this specification, the IPSec policy would be "Require IPSec, from negotiated address to me, destination is negotiated port". This causes the Responder to require use of IPSec. If some Sender does not support IPSec, then a more granular policy will be required: "Require IPSec, from IPSec-Capable-Sender to me, destination is negotiated port".

6.3. Replay Protection

For the Control Messages the originator of the message MAY choose to include a current value in the Sent Timestamp field indicating the time the message was submitted for transmission, it MUST be set to 0 otherwise. The receiver of the message MAY choose to validate if the timestamp is within an acceptable range. The Measurement Traffic described in this document contains a timestamp to indicate the sent time and hence no new field is required.

7. Terminology

Term	Description
Control Phase	A phase during which Control Request and Control Response is exchanged.
---	---
L2	OSI Data Link Layer
---	---
L3	OSI Network Layer
---	---
Measurement Phase	Active measurement phase that is marked by a sequence of Measurement Request and Measurement Response exchanges.
---	---
Metric	A particular characteristic of the network data traffic, for example latency, jitter, packet/frame loss
---	---
Responder	A network element that responds to a message
---	---
RTP	Real-time Transport Protocol
---	---
Sender	A network element that is the initiator of a message exchange
---	---
Service Level	This is the level of service that is agreed upon between the Provider and the Customer
---	---
UDP	User Datagram Protocol

8. Acknowledgements

The authors wish to acknowledge the contributions of several key people who contributed to the current form of the document. Hanlin Fang, David Wang, Anantha Ramaiah, Max Pritikin, and Malini Vijayamohan.

9. References

9.1. Normative References

[IEEE1588]
IEEE, "1588-2008 Standard for a Precision Clock

Synchronization Protocol for Networked Measurement and Control Systems", March 2008.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", RFC 4868, May 2007.

[RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

9.2. Informative References

[RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", RFC 3579, September 2003.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

Authors' Addresses

Murtaza S. Chiba
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: mchiba@cisco.com
URI:

Alexander Clemm
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: alex@cisco.com
URI:

Steven Medley
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: stmedley@cisco.com
URI:

Joseph Salowey
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: jsalowey@cisco.com
URI:

Sudhir Thombare
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: thombare@cisco.com
URI:

Eshwar Yedavalli
Cisco Systems
170 West Tasman Drive
San Jose, 95134
USA

Phone: 1-408-526-4000
Fax:
Email: eshwar@cisco.com
URI: