



## **IVI-6.1: IVI High-Speed LAN Instrument Protocol (HiSLIP)**

June 8, 2010  
Revision 1.0

## **Important Information**

---

The IVI-6.1: High-Speed LAN Instrument Protocol Specification is authored by the IVI Foundation member companies. For a vendor membership roster list, please visit the IVI Foundation web site at [www.ivifoundation.org](http://www.ivifoundation.org).

The IVI Foundation wants to receive your comments on this specification. You can contact the Foundation through the web site at [www.ivifoundation.org](http://www.ivifoundation.org).


### **Warranty**

The IVI Foundation and its member companies make no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The IVI Foundation and its member companies shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

### **Trademarks**

Product and company names listed are trademarks or trade names of their respective companies.

No investigation has been made of common-law trademark rights in any work.



---

<b>1</b>	<b>Overview of the IVI HiSLIP Specification .....</b>	<b>6</b>
1.1	IVI HiSLIP Overview .....	6
1.2	References .....	6
1.3	Definitions of Terms and Acronyms .....	6
<b>2</b>	<b>HiSLIP Protocol Description .....</b>	<b>8</b>
2.1	HiSLIP Port Assignment .....	8
2.2	Message Header .....	8
2.3	Summary of HiSLIP Messages .....	9
2.4	Numeric Values of Message Type codes .....	13
2.5	Locking Mechanism .....	14
	2.5.1 Server Behavior When Client Does Not Have a Lock .....	14
2.6	HiSLIP Buffer Sizes .....	16
<b>3</b>	<b>Overlapped and Synchronized Modes .....</b>	<b>17</b>
3.1	Synchronized Mode .....	18
	3.1.1 Synchronized Mode Server Requirements .....	18
	3.1.2 Synchronized Mode Client Requirements .....	18
3.2	Overlapped mode .....	20
	3.2.1 Overlap Mode Server Requirements .....	20
	3.2.2 Overlap Mode Client Requirements .....	20
<b>4</b>	<b>HiSLIP Transactions .....</b>	<b>21</b>
4.1	Initialization Transaction .....	22
4.2	Fatal Error Detection and Synchronization Recovery .....	25
4.3	Error Notification Transaction .....	26
4.4	Data Transfer Messages .....	27
4.5	Lock Transaction .....	29
	4.5.1 Unlock Considerations .....	31
4.6	Lock Info Transaction .....	32
4.7	Remote Local Transaction .....	33
4.8	Trigger Message .....	36
4.9	Vendor Defined Transactions .....	37
4.10	Maximum Message Size Transaction .....	38
4.11	Interrupted Transaction .....	39
4.12	Device Clear Transaction .....	40
	4.12.1 Feature Negotiation .....	41
4.13	Service Request .....	42
4.14	Status Query Transaction .....	43
	4.14.1 MAV Generation in Synchronized Mode .....	44

4.14.2	MAV Generation in Overlapped Mode .....	44
4.14.3	Implementation Note.....	44

<b>A.</b>	<b>Analysis of Interrupted Conditions.....</b>	<b>46</b>
A.1	Slow Client.....	46
A.2	Fast Client .....	47
A.3	Intermediate Timing .....	49

## IVI HiSLIP Revision History

---

This section is an overview of the revision history of the IVI HiSLIP specification.

**Table 1. IVI HiSLIP Class Specification Revisions**

<b>Status</b>	<b>Action</b>
Revision 1.0	First version of specification.

# 1 Overview of the IVI HiSLIP Specification

---

HiSLIP (High Speed LAN Instrument Protocol) is a protocol for TCP-based instrument control that provides the capabilities of conventional test and measurement protocols with minimal impact to performance. The HiSLIP protocol includes:

- Device clear
- Instrument status reporting with message available calculation per IEEE Std 488.2
- Instrument remote/local status control
- Instrument locking
- Service Request from the instrument to the client
- End message
- Message exchange protocol interrupted error detection

## 1.1 IVI HiSLIP Overview

HiSLIP creates two TCP connections to the same server port<sup>1</sup> referred to as the synchronous channel and asynchronous channel. HiSLIP sends packetized messages between the client and server on both channels.

The synchronous channel carries normal bi-directional ASCII command traffic (such as SCPI) and synchronous GPIB-like meta-messages (such as END and trigger). Generally, both the client and server can leave messages in the synchronous buffers and execute them in a synchronous fashion.

The asynchronous channel carries GPIB-like meta-messages that need to be handled independently of the data path (such as device clear and service request). Generally, both the client and server need to treat asynchronous messages as higher priority and act on them before messages from the synchronous channel,

The HiSLIP protocol permits multiple virtual instruments at a single port at a given IP address. When the connection is initialized, the client specifies a sub-address that designates the specific virtual instrument to be associated with this connection. The protocol does not associate any aspects of the connections to multiple virtual instruments.

## 1.2 References

Several other documents and specifications are related to this specification. These other related documents are the following:

VXIPNP	VXI <i>plug&amp;play</i> VISA Specification defines the API between the client application and the client-side HiSLIP implementation. This specification is available from the IVI Foundation web site at <a href="http://www.ivifoundation.org">www.ivifoundation.org</a> .
VXI-11.1, 11.2, and 11.3	These standards define the VXI-11 protocol which is the primary predecessor to HiSLIP.
IEEE 488.2	IEEE Std 488.2 defines the interrupted protocol requirements as well as the appropriate server behavior for several of the GPIB messages.

## 1.3 Definitions of Terms and Acronyms

This section defines terms and acronyms that are specific to the HiSLIP protocol:

RMT	From IEEE Std 488.2: Response Message Terminator. RMT is the new-line accompanied by END sent from the server to the client at the end of a response. Note that with HiSLIP this is implied by the <i>DataEND</i> message.
-----	--

---

<sup>1</sup> The IANA assigned port for HiSLIP is 4880

END	From IEEE Std 488.2: END is a protocol generated indication of the end of a message. It is not indicated with an 8-bit value in the data stream. This message is provided by HiSLIP.
eom	From IEEE Std 488.2: end-of-message. eom is the termination character of a message to the server. The eom is either: new-line, END, or a new-line accompanied by an END. For the purposes of HiSLIP, eom is implicit after group execute trigger.
interrupted	From IEEE Std 488.2: A protocol error indicating that a server received an input message (either a command or query) before the client has fully accepted the response of the preceding message.
HiSLIP	High Speed LAN Instrument Protocol defined in this specification.
MAV	From IEEE Std 488.2: A bit indicating that there is a message available from the server.

## 2 HiSLIP Protocol Description

---

Both the synchronous and asynchronous channels send all command and data information in a fixed packet format. A complete packet is referred to as a message.

### 2.1 HiSLIP Port Assignment

All HiSLIP clients and servers shall use the IANA assigned port number of 4880.

### 2.2 Message Header

The messages consist of a header followed by a counted payload. However, the payload count is frequently zero.

**Table 2 HiSLIP Message Header Format**

Field	Octets	Field Offset
Prologue (ASCII "HS")	2	0
Message Type	1	2
Control Code	1	3
Message Parameter	4	4
Payload Length	8	8
Data	Payload Length	16

Table 2 defines the header used for HiSLIP messages. The fields are:

Prologue	A pattern to facilitate HiSLIP devices detecting when they receive an ill-formed message or are out of sync. The value shall be ASCII 'HS' encoded as a 16 bit value. 'H' is in the most significant network order position and 'S' is in the second byte.
Message Type	This field identifies this message. See Table 3, <i>HiSLIP Messages</i> , for a description of the HiSLIP messages. See Table 4, <i>Message Type Value Definitions</i> , <b>Error! Reference source not found.</b> for the numeric values of each message type.
Control Code	This 8-bit field is a general parameter for the message. If the field is not defined for a message, 0 shall be sent.
MessageParameter	This 32-bit field has various uses in different messages. If this field is not defined for a message, 0 shall be sent.
Payload Length	This field indicates the length in octets of the payload data contained in the message. This field is an unsigned 64-bit integer. The maximum data transfer size may be limited by the implementation, see section 4.10, <i>Maximum Message Size Transaction</i> for details. If the message type does not use a payload, the length shall be set to zero.

All Hi-SLIP fields are marshaled onto the network in network order (big endian). That is, most significant byte first. Where the specification calls for an ASCII string as the payload the payload length shall refer to the length of the number of characters in the string and extended ASCII (8-bit) is implied. A trailing NUL character shall not be sent or accounted for in the length.



## **2.3 Summary of HiSLIP Messages**

Table 3 summarizes the HiSLIP messages.

**Table 3 HiSLIP Messages**

Sender	Channel	Message Type (1 byte)	Control Code (1 byte)	Message Parameter (4 bytes)	Payload (counted field)
C	S	<i>Initialize</i>		UpperWord : Client protocol version LowerWord : Client-vendorID	sub-address inASCII, may be of zero length
S	S	<i>InitializeResponse</i>	Bit 0 : 1 Prefer Overlap : 0 Prefer Synchronized	UpperWord : Server Protocol version LowerWord : SessionID	--
E	E	<i>FatalError</i>	ErrorCode (see Table 7)	--	Error Message in ASCII, may be of zero length
E	E	<i>Error</i>	ErrorCode (see Table 9)	--	Error Message in ASCII, may be of zero length
C	S	<i>Data</i>	Bit 0 : 0 RMT was not delivered Bit 0 : 1 RMT was delivered	MessageID (this msg)	Counted data
S	S		--	MessageID – usage depends on overlapped or synchronized mode	Counted data
C	S	<i>DataEND</i>	Bit 0 : 0 RMT was not delivered Bit 0 : 1 RMT was delivered	MessageID (this msg)	Counted data
S	S		--	MessageID – usage depends on overlapped or synchronized mode	Counted data
C	S	<i>AsyncLock</i>	1 – Request	Timeout (in ms)	LockString in ASCII, may be of zero length.
			0 – Release	MessageID of last sent message	--
S	A	<i>AsyncLockResponse</i>	<i>In response to request:</i> 0 – Failure 1 – Success 3 - Error	--	--
			<i>In response to release:</i> 1 - Success exclusive 2 - Success shared 3 - Error	--	--
C	A	<i>AsyncLockInfo</i>	--	--	--

S	A	<i>AsyncLockInfoResponse</i>	0 – No exclusive lock granted 1 – Exclusive lock granted	Number of HiSLIP clients holding locks when <i>AsyncLockInfo</i> was processed	--
C	A	<i>AsyncRemoteLocalControl</i>	0 – Disable remote 1 – Enable remote 2 – Disable remote and go to local 3 – Enable Remote and go to remote 4 – Enable remote and lock out local 5 – Enable remote, go to remote, and set local lockout 6 – go to local without changing REN or lockout state	MessageID of last sent message	--
S	A	<i>AsyncRemoteLocalResponse</i>	--	--	--
C	A	<i>AsyncDeviceClear</i>	--	--	--
S	A	<i>AsyncDeviceClearAcknowledge</i>	Feature-bitmap	--	--
C	S	<i>DeviceClearComplete</i>	Feature-bitmap	--	--
S	S	<i>DeviceClearAcknowledge</i>	Feature-bitmap	--	--
C	S	<i>Trigger</i>	0 – RMT was not delivered 1 – RMT was delivered	MessageID (this message)	--
S	S	<i>Interrupted</i>	--	MessageID	--
S	A	<i>AsyncInterrupted</i>	--	MessageID	
C	A	<i>AsyncMaximumMessageSize</i>	--	--	8-byte size – note that the payload length is always 8 and the count is in the payload.
S	A	<i>AsyncMaximumMessageSizeResponse</i>	--	--	8-byte size – note that the payload length is always 8 and the count is in the payload.
C	A	<i>AsyncInitialize</i>	--	SessionID	--
S	A	<i>AsyncInitializeResponse</i>	--	Server-vendorID	--
S	A	<i>AsyncServiceRequest</i>	Server status	--	--
C	A	<i>AsyncStatusQuery</i>	0 – RMT was not delivered 1 – RMT was delivered	MessageID of last sent message	--
S	A	<i>AsyncStatusResponse</i>	Server status response	--	--
E	E	<i>VendorSpecific</i>	Arbitrary	Arbitrary	Data

In Table 3 :

In the Sender column :

S indicates

Server generated message

C indicates

Client generated message

E indicates

A message that may be generated by either the client or server

In the channel column :

S indicates

Synchronous channel message

A indicates

Asynchronous channel message

E indicates

A message that may be send on either the synchronous or asynchronous channel

## 2.4 Numeric Values of Message Type codes

The Table 4 defines the numeric values of the Message Type codes.

**Table 4 Message Type Value Definitions**

Designation	Channel	Numeric Value (decimal)
<i>Initialize</i>	Synchronous	0
<i>InitializeResponse</i>	Synchronous	1
<i>FatalError</i>	Synchronous	2
<i>Error</i>	Synchronous	3
<i>AsyncLock</i>	Asynchronous	4
<i>AsyncLockResponse</i>	Asynchronous	5
<i>Data</i>	Synchronous	6
<i>DataEnd</i>	Synchronous	7
<i>DeviceClearComplete</i>	Synchronous	8
<i>DeviceClearAcknowledge</i>	Asynchronous	9
<i>AsyncRemoteLocalControl</i>	Asynchronous	10
<i>AsyncRemoteLocalResponse</i>	Asynchronous	11
<i>Trigger</i>	Synchronous	12
<i>Interrupted</i>	Synchronous	13
<i>AsyncInterrupted</i>	Asynchronous	14
<i>AsyncMaximumMessageSize</i>	Asynchronous	15
<i>AsyncMaximumMessageSizeResponse</i>	Asynchronous	16
<i>AsyncInitialize</i>	Asynchronous	17
<i>AsyncInitializeResponse</i>	Asynchronous	18
<i>AsyncDeviceClear</i>	Asynchronous	19
<i>AsyncServiceRequest</i>	Asynchronous	20
<i>AsyncStatusQuery</i>	Asynchronous	21
<i>AsyncStatusResponse</i>	Asynchronous	22
<i>AsyncDeviceClearAcknowledge</i>	Asynchronous	23
<i>AsyncLockInfo</i>	Asynchronous	24
<i>AsyncLockInfoResponse</i>	Asynchronous	25
reserved for future use		26-127
<i>VendorSpecific</i>	Either	128-255 inclusive

## 2.5 Locking Mechanism

The HiSLIP protocol supports the VISA locking mechanism<sup>2</sup> including exclusive and shared locks.

After the initialization transaction, all subsequent HiSLIP messages from the client to the server are subject to access control governed by the locking transaction.

Section 4.5 describes how one or more clients can obtain a lock. If the server is locked, only the client or clients that have been granted the lock are permitted access to the server. Clients that do not have the lock behave as described in section 2.5.1, *Server Behavior When Client Does Not Have a Lock*.

If a client is holding an exclusive lock, the server guarantees that only that client has received an exclusive lock. If no client is holding an exclusive lock, then multiple clients may be granted a shared lock. If multiple clients hold a shared lock, each is permitted access to the server. If the server has no outstanding locks, then any client is permitted access to the server.

If any client is holding a shared lock, only a client holding a shared lock shall be granted an exclusive lock. However, if no client is holding a shared lock, a client requesting an exclusive lock shall be granted that lock.

If a server receives a *Lock* message requesting a lock from a client that is not holding the lock, it will wait for the timeout time indicated in that message for the lock to become available.

If a HiSLIP connection is closed any locks assigned to the corresponding client are immediately released by the server.

Nothing in this section should be taken to require that a server is precluded from implementing other security mechanisms that may result in it refusing to grant access to any client.

The behavior of a server that supports connections other than HiSLIP is beyond the scope of this specification. However, it is appropriate for a server to manage a single lock across several connection styles, and therefore, a HiSLIP client may be refused a lock although no other HiSLIP client is holding a lock. Also, a HiSLIP lock may impact other connection styles.

### 2.5.1 Server Behavior When Client Does Not Have a Lock

This section describes how servers behave towards clients that do not have the lock when some other client has been granted either an exclusive or shared lock.

Exchanges with a client that does not have a lock while another client has been granted a lock are handled as follows:

- The *AsyncDeviceClear* transaction (including all protocol messages) is completed immediately per section 4.12, *Device Clear Transaction*. Note that only the channel upon which the *AsyncDeviceClear* was sent is impacted by this transaction. In this circumstance, device clear is not permitted to impact any of the server functions other than those associated with this HiSLIP communication session.
- Any transactions that are in-process are completed normally. Thus, the following may be sent:
  - *AsyncLockResponse*
  - *AsyncInterrupted*
  - *AsyncMaximumMessageSizeResponse*
  - Appropriate *VendorSpecific* messages (subject to the vendor-specific definition of these messages)

---

<sup>2</sup> See VXIPNP per section 1.2, *References*

- *AsyncStatusQuery*, *AsyncMaximumMessageSize* and *AsyncLockInfo* transactions are completed.
- *AsyncRemoteLocalControl* transactions are completed. However, the client shall not wait for locks on other clients to be released. If another client is holding a lock, the server shall only act on the remote/local request after the lock is released. The behavior is device dependent.
- *AsyncServiceRequest* messages are sent normally.
- All synchronous messages are left in the input buffer. Normal TCP behavior to prevent buffer overflows takes place. *Data*, *DataEND*, *Trigger*, *Error*, *FatalError*, and *VendorSpecific* messages are left in the input buffer.
- Data traffic from the server to the client (*Data* and *DataEND* messages) are sent normally. However, since incoming synchronous messages are blocked, only responses to operations that were begun before the lock was granted will be generated.

## 2.6 HiSLIP Buffer Sizes

HiSLIP provides the *AsyncMaximumMessageSize* transaction to inform the client and server of the largest message they are permitted to send to the other device on the synchronous channel. Clients should initiate this transaction as part of the initialization process (shortly after the HiSLIP initialization) in order to guarantee that messages do not overflow the other devices buffers.

Each device shall also have a buffer suitable to accept any asynchronous message. The buffer size for asynchronous messages are dominated by the messages with a variable payload. That is:

- *Error* and *FatalError* which include a variable length string in addition to the 16-byte message header
- *AsyncLock* message (sent to the server) that includes the variable length lock string.

The VISA<sup>3</sup> limitation for these strings is 256 characters, yielding a maximum message size of 272 bytes.

If either the server or client detect the arrival of a message that is too large to properly handle it shall send an *Error* message with the appropriate HiSLIP defined error value.

---

<sup>3</sup> VISA rule 3.6.13



### 3 Overlapped and Synchronized Modes

---

In order to maintain compatibility with GPIB, VXI-11 and USB-TMC instruments, the HiSLIP protocol supports two different operating modes:

- Overlap mode      In overlap mode input and output data and trigger messages are arbitrarily buffered between the client and server. For instance, a series of independent query messages can be sent to the server without regard to when they complete. The responses from each will be returned in the order the queries were sent. Thus multiple query operations may be initiated and conducted by the server independent of the rate at which the client consumes the responses.
- Synchronized mode      In Synchronized mode, the client is required to read the result of each query message before sending another query<sup>4</sup>. If the client fails to do so, the protocol generates the interrupted protocol error and the response from the preceding query is cleared by the protocol.

All HiSLIP clients shall support both synchronized and overlapped mode. HiSLIP servers shall support either synchronized or overlapped mode or both.

The following sections describe the implementation of these modes.

Note that the calculation of message available (MAV) and the *AsyncStatusQuery* transaction differ between the overlapped and synchronized modes also. See section 4.14, *Status Query Transaction* for details.

---

<sup>4</sup> Per the IEEE 488.2 definition of a response message.

### 3.1 Synchronized Mode

Synchronized mode closely mimics the requirements of the IEEE Std 488.2 message exchange protocol to detect the interrupted error.

#### 3.1.1 Synchronized Mode Server Requirements

HiSLIP servers shall implement the following:

1. When the server application layer (nominally an instrument parser) requests that HiSLIP send a response message terminator the server shall verify that no data is in the server input queue. If there is data in the input queue, the server shall declare an interrupted error.

To declare an interrupted error, the server shall:

- Use the server error reporting mechanism to report the interrupted error within the server.
  - Clear the response message just received from the server application layer, and any other messages buffered to be sent to the client.
  - Send the *Interrupted* transaction to the client, including both the *Interrupted* and *AsyncInterrupted* messages.
2. When receiving *Data*, *DataEND*, *Trigger*, or *StatusQuery* verify the state of the RMT-delivered flag.

The server shall maintain a flag that indicates RMT-expected. RMT-expected shall be set *true* when the server sends a *DataEND* message (that is, when it sends an RMT).

The RMT-expected bit shall be cleared when the server receives *AsyncStatusQuery* with RMT-delivered flag set to true.

When *Data*, *DataEND*, or *Trigger* are received, if RMT-expected and RMT-delivered are both either true or false the RMT-expected bit shall be cleared.

When *Data*, *DataEND*, or *Trigger* are received, if RMT-expected and RMT-delivered are different, the server shall declare an interrupted error. The server shall use the server error reporting mechanism to report the interrupted error. No indication of this interrupted error is sent to the client by the HiSLIP protocol.

When servers send the *DataEND* message, they shall set the MessageID field to the MessageID of the client message that contained the eom that generated this response.

When servers send the *Data* message, they may set the MessageID field to the MessageID of the client message that contained the eom that generated this response so long as that eom is at the end of the identified message. In some circumstance (for instance, if the eom is not at the end of the message), the server might not be able to provide the MessageID of the message ending in the eom. In these circumstances, the server shall set the MessageID to 0xffff ffff.

After interrupted error processing is complete, the server resumes normal operation.

#### 3.1.2 Synchronized Mode Client Requirements

HiSLIP clients shall implement the following:

1. When receiving *DataEND* (that is an RMT) verify that the MessageID indicated in the *DataEND* message is the MessageID that the client sent to the server with the most recent *Data*, *DataEND* or *Trigger* message.

If the MessageIDs do not match, the client shall clear any *Data* responses already buffered and discard the offending *DataEND* message.

2. When receiving *Data* messages if the MessageID is not 0xffff ffff, then verify that the MessageID indicated in the *Data* message is the MessageID that the client sent to the server with the most recent *Data*, *DataEND* or *Trigger* message.

If the MessageIDs do not match, the client shall clear any *Data* responses already buffered and discard the offending *Data* message.

3. When the client sends *Data*, *DataEND* or *Trigger* if there are any whole or partial server messages that have been validated per rules 1 and 2 and buffered they shall be cleared.
4. When the client receives *Interrupted* or *AsyncInterrupted* it shall clear any whole or partial server messages that have been validated per rules 1 and 2.

If the client initially detects *AsyncInterrupted* it shall also discard any further *Data* or *DataEND* messages from the server until *Interrupted* is encountered.

If the client detects *Interrupted* before it detects *AsyncInterrupted*, the client shall not send any further messages until *AsyncInterrupted* is received.

Clients shall maintain a MessageID count that is initially set to 0xffff ff00. When clients send *Data*, *DataEND* or *Trigger* messages, they shall set the message parameter field of the message header to the current MessageID and increment the MessageID by two in an unsigned 32-bit sense (permitting wrap-around).

The MessageID is reset to 0xffff ff00 after device clear, and when the connection is initialized.

After interrupted error processing is complete, the client resumes normal operation.

## 3.2 **Overlapped mode**

In overlapped mode commands and responses are buffered by the client and server and I/O operations are permitted to overlap.

No special processing is required in the server or client other than buffering inbound messages until the respective application layer requires them. Buffers are only cleared by a device clear.

### 3.2.1 **Overlap Mode Server Requirements**

HiSLIP overlap mode servers maintain a MessageID and use it as follows:

1. The MessageID shall be reset to 0xffff ff00 after device clear or initialization.
2. When the server sends *Data* or *DataEND* messages it shall place the MessageID into the message parameter and increment it by two in an unsigned 32-bit fashion (permitting wrap-around).

### 3.2.2 **Overlap Mode Client Requirements**

HiSLIP clients shall implement the following:

1. In overlap mode, when sending *AsyncStatusQuery* the client shall place the MessageID of the most recent message that has been entirely delivered to the client in the message parameter.

Clients shall maintain a MessageID count that is initially set to 0xffff ff00. When clients send *Data*, *DataEND* or *Trigger* messages, they shall set the MessageID field of the message header to the current MessageID and increment the MessageID by two in an unsigned 32-bit sense (permitting wrap-around).

The MessageID is reset after device clear, and when the connection is initialized. In overlap mode, the MessageID is only used for locking.

## 4 HiSLIP Transactions

---

The following sections describe the HiSLIP protocol transactions.

In the following sections angle brackets (<>) are used to separate the various fields of the message. This is always expressed as four fields, the fourth field represents the payload and a count is always implied. <0> in the payload field indicates a count of zero and no payload.

## 4.1 Initialization Transaction

The purpose of the initialization transaction is to establish the HiSLIP connection between the client and server. This requires opening a synchronous and an asynchronous channel on the same server port and associating the two together. The two are associated through a session ID that is provided to the client by the server in response to the *Initialize* message.

Server's shall support multiple simultaneous clients initializing.

**Table 5 Initialization Transaction**

Step	Initiator	Message content	Action
0	Server	none	Server passively opens TCPserver socket on the IANA assigned port.
1	Client	Opens the synchronous TCP connection (TCP SYN message)	Client does an active TCP open, the server continues to wait for additional connections
2	Client	<Initialize><0><upper:client-protocol-version : lower :client-vendorID><sub-address>	Client starts the initialization by identifying the vendor, specifying the subaddress, and advertising the protocol version it supports.
3	Server	<InitializeResponse><overlap-mode><upper :server-protocol-version : lower :SessionID><0>	Server responds with its protocol version. The two go to the lower of server version and client version.  The server also provides the SessionID to send with <i>AsyncInitialize</i> .  Note that several <i>FatalError</i> messages are appropriate at this time.
5	Client	Opens the asynchronous TCP connection (TCP SYN message)	Client opens second connection for the asynchronous channel on same server port
6	Client	<AsyncInitialize><0><SessionID><0>	The client sends SessionID to associate this TCP session with the established HiSLIP synchronous channel.
7	Server	<AsyncInitializeResponse><0><server-vendorID><0>	Server acknowledges initialize and provides the vendor ID.  The HiSLIP connection is ready for use.

The following are the fields in the *Initialize* client message:

client-protocol-version This identifies the highest version of the HiSLIP specification that the client implements. Per the IVI standards requirements, HiSLIP specification versions are of the form <major>.<minor>. The major specification revision, expressed as a binary 8-bit integer is the first byte of the client version. The minor number expressed as a binary 8-bit integer is the second byte of the client version.

The client version is sent in the most significant 16-bits (big endian sense) of the 32-bit message parameter.

client-vendorID This identifies the vendor of the HiSLIP protocol on the client. This is the two-character vendor abbreviation from the *VXIplug&play* specification VPP-9. These abbreviations are assigned free of charge by the IVI Foundation<sup>5</sup>.

The client vendorID is sent in the least significant 16-bits (big endian sense) of the 32-bit message parameter.

sub-address This field corresponds to the VISA LAN device name. It identifies a particular device managed by this server. It is in the payload field and therefore includes a 64-bit count. The count is followed by the appropriate length ASCII sub-address. For instance: "device2". The maximum length for this field is 256 characters.

If the sub-address is null (zero length) the initialize opens the default (perhaps only) device at this IP address.

The following are the fields in the *InitializeResponse* server message:

server-protocol-version This identifies the highest version of the HiSLIP specification that the server implements. It is expressed the same as the client-version field in the *Initialize* client message.

The server version is sent in the most significant 16-bits (big endian sense) of the 32-bit message parameter.

SessionID This is used to associate the synchronous and asynchronous connections and must be provided by the client in the *InitializeAsync* message. This associates the two TCP connections into a single HiSLIP connection.

The client vendorID is sent in the least significant 16-bits (big endian sense) of the 32-bit message parameter.

overlap-mode The server uses this field to indicate if it initially in overlapped or synchronous mode. 0 indicates synchronous, 1 indicates overlapped.

The following is the field in the *AsyncInitialize* message:

SessionID This is the session ID provided by the server in the *InitializeResponse* message. It associates the synchronous and asynchronous connections. This may be discarded by the client after this message.

The following is the field in the *AsyncInitializeResponse* message:

server-vendorID This identifies the vendor of the server. This is the two-character vendor abbreviation from the *VXIplug&play* specification VPP-9. These abbreviations are assigned free of charge by the IVI Foundation.<sup>6</sup>

After the initialization sequence, the client and server will both use the highest protocol revision supported by both devices (that is, the smallest of the two exchanged versions). Note that all HiSLIP devices must support earlier protocol versions.

---

<sup>5</sup> Contact the IVI Foundation ([admin@ivifoundation.org](mailto:admin@ivifoundation.org)) to register a new vendor ID (also known as a vendor prefix). Vendors do not need to join the IVI Foundation to obtain a defined two-character abbreviation.

<sup>6</sup> Ibid

Clients that require exclusive access to the server must immediately follow the initialize transaction with an appropriate *Lock* transaction. If the *Lock* operation fails then the client can close the connection. Servers shall not automatically grant a lock to new clients, and the connection may be opened when another client is holding the lock.

If the client closes the connection after receiving the *InitializeResponse*, the server should not declare an error as this is a legitimate way for a client to validate the presence and version of a server.



## 4.2 Fatal Error Detection and Synchronization Recovery

**Table 6 Synchronous Fatal Error Message**

<i>Initiator</i>	<b>Message</b>	<b>Data Consumer</b>
<i>Either client or server</i>	<FatalError><ErrorCode><0><length><message>	Accept data and handle appropriately
<i>Initiator</i>	Close the connection	If initiator is the client, it may re-open the connection per 4.1

At any point, the client or server may encounter a non-recoverable error situation. For instance, the prologue may be incorrect. If either device detects an error condition that is likely to cause the two devices to lose synchronization it shall send the *FatalError* message on the synchronous channel and the asynchronous channel with appropriate diagnostic information.

The IVI Foundation defines the error codes listed in Table 7. Error codes from 128-255 inclusive are device defined.

The payload shall be of the specified length and contain a human readable error description expressed in ASCII. A length of zero with no description is legal.

If the error is detected by the client, after sending the *FatalError* messages it shall close the HiSLIP connection and may attempt to re-establish the connection (that is, close both synchronous and asynchronous connections and re-establish the connection per section 4.1, *Initialization Transaction*).

If the error is detected by the server, after sending the *FatalError* messages, it shall close the HiSLIP connection. The client may re-establish the connection. However the SessionID for the new session will not necessarily relate to the previous SessionID. Note that locks will not be retained and must be re-acquired.

**Table 7 HiSLIP Defined Fatal Error Codes**

<i>Error Code</i>	<b>Message</b>
<i>0</i>	Unidentified error
<i>1</i>	Poorly formed message header
<i>2</i>	Attempt to use connection without both channels established
<i>3</i>	Invalid Initialization Sequence
<i>4</i>	Server refused connection due to maximum number of clients exceeded
<i>5-127</i>	Reserved for HiSLIP extensions
<i>128-255</i>	Device defined errors

### 4.3 Error Notification Transaction

**Table 8 Synchronous Error Notification Transaction**

<b>Initiator</b>	<b>Message</b>	<b>Data Consumer</b>
<i>Either client or server</i>	<Error><ErrorCode><0><length><message>	Accept data and handle appropriately
<i>Initiator</i>	No further action	No further action

If either the client or server receive a message that it is unable to process but that does not cause it to lose synchronization with the sender it shall discard the errant message and any payload associated with it, then reply with the *Error* message.

The *Error* message shall be sent on whichever connection (synchronous or asynchronous) that the errant message arrived on.

The payload shall be of the specified length and contain a human readable error description expressed in ASCII. A length of zero with no description is legal.

The IVI Foundation defines the error codes listed in Table 9, error codes from 128-255 inclusive are device defined.

After sending the *Error* message, the device shall return to normal processing.

For example, the *Error* message should be sent in reply to unrecognized vendor specific messages or unsupported MessageIDs or control codes.

**Table 9 HiSLIP Defined Error Codes (non-fatal)**

<b>Error Code</b>	<b>Message</b>
<i>0</i>	Unidentified error
<i>1</i>	Unrecognized Message Type
<i>2</i>	Unrecognized control code
<i>3</i>	Unrecognized Vendor Defined Message
<i>4</i>	Message too large
<i>5-127</i>	Reserved for HiSLIP extensions
<i>128-255</i>	Device defined errors

#### 4.4 DataTransfer Messages

Table 10 Data Transfer Messages from Client to Server

<b>Initiator</b>	<b>Message</b>	<b>Data Consumer</b>
<i>client</i>	<Data><RMT-delivered><MessageID><length><data>	<p>Accept data and use it appropriately</p> <p>RMT-delivered is 1 if this is the first <i>Data</i>, <i>DataEND</i> or <i>Trigger</i> message since the client delivered RMT to the application layer.</p> <p>The client increments the MessageID with each <i>Data</i>, <i>DataEND</i> or <i>Trigger</i> message sent.</p>
<i>client</i>	<DataEND><RMT-delivered><MessageID><length><data>	<p>Accept data and use it appropriately. Final data byte has an accompanying END.</p> <p>RMT-delivered is 1 if this is the first <i>Data</i>, <i>DataEND</i> or <i>Trigger</i> message since the client delivered RMT to the application layer.</p> <p>The client increments the MessageID with each <i>Data</i>, <i>DataEND</i> or <i>Trigger</i> message sent.</p>

**Table 11 Data Transfer Messages from Server to Client**

<b>Initiator</b>	<b>Message</b>	<b>Data Consumer</b>
<i>server</i>	<Data><0><MessageID><length><data>	Accept data and handle appropriately  The MessageID in synchronized mode is the ID of the message containing the RMT that generated this response or 0xffff ffff. The MessageID in overlapped mode is an ID that is incremented with each data transfer message sent.
<i>server</i>	<DataEND><0><MessageID><length><data>	Accept data and handle appropriately. Final data byte has an accompanying END.  The MessageID in synchronized mode is the ID of the message containing the RMT that generated this response. The MessageID in overlapped mode is an ID that is incremented with each data transfer message sent.

Either the server or the client is permitted to initiate a data transfer at any time.

For client originated message:

- RMT-delivered RMT-delivered is 1 if this is the first *Data*, *DataEND*, *Trigger*, or *StatusQuery* message since the HiSLIP client delivered an RMT to the client application layer.
- MessageID MessageID identifies this message so that response data from the server can indicate the message that generated it. For generation of the MessageID, see sections 3.1.2 (*Synchronized Mode Client Requirements*) and 3.2.2 (*Overlap Mode Client Requirements*).

For server originated messages:

- MessageID MessageID in synchronized mode identifies the client message responsible for generating this response. In overlapped mode, the MessageID is a continuously incrementing count that assists in MAV generation. See section 3.1.1(*Synchronized Mode Server Requirements*)and 3.2.1(*Overlap Mode Server Requirements*), and section 4.14.2, *MAV Generation in Overlapped Mode*

The *DataEND* message indicates that the END message should be processed with the final data byte.

These messages are not acknowledged.

## 4.5 Lock Transaction

**Table 12 Lock Transaction – Requesting a Lock**

Step	Sender	Message	Action
1	Client	<AsyncLock><1=request><timeout><LockString length><LockString>	Request lock, wait up to timeout milliseconds for it to become available.  LockString is an ASCII string indicating shared lock identification. A zero length string indicates an exclusive lock request.
2	Server	<AsyncLockResponse><0=failure, 1=success, 3=error><0><0>	Response indicates if the lock was successful.

**Table 13 Lock Transaction – Releasing a Lock**

Step	Sender	Message	Action
1	Client	<AsyncLock><0 =release><MessageID><0>	Release lock
2	Server	<AsyncLockResponse><1=success exclusive, 2=success shared, 3=error><0><0>	Response indicates the type of lock released if any

The *AsyncLock* client message is used to request or release a lock as described in Table 14.

The following are the fields in the *AsyncLock* transaction:

LockString	an ASCII string that identifies this lock
MessageID	designates the last message to be completed before the release takes place.
Timeout	The <i>AsyncLock</i> request passes a 32-bit timeout in the MessageID field. This is the amount of time in milliseconds the client is willing to wait for the lock to grant. If the lock is not available in this amount of time, the <i>AsyncLockResponse</i> will fail and return failure. A timeout of 0 indicates that the server should only grant the lock if it is available immediately.

The server shall always reply with an *AsyncLockResponse* per Table 14.

**Table 14 Lock request/release operation descriptions**

<b>Lock Control Code</b>	<b>LockResponse Control Code</b>	<b>Description</b>
0 (release)	1 (success)	Release of exclusive lock was granted.
	2 (success)	Release of shared lock was granted
0 (release)	3 (error)	Invalid attempt to release a lock that was not acquired.
1 (request)	0 (fail)	Lock was requested but not granted (timeout expired)
1 (request)	1 (success)	The lock was requested and granted
1 (request)	3 (error)	Invalid (redundant) request that is, requesting a lock already granted

A null LockString in the *AsyncLock* LockString with a Control Code of 1 indicates a lock request for an exclusive lock.

A non-null LockString in the *AsyncLock* LockString with a Control Code of 1 indicates a request for a shared lock.

An *AsyncLock* with the Control Code set to 0 indicates a request to release the lock.

HiSLIP servers shall respond to lock requests as shown in Table 15 Lock BehaviorTable 15. Note that the ‘lock state’ described in this table is the lock state across all active HiSLIP sessions.

After the Initialization and Device Clear transaction, the client shall use the MessageID = 0xffffefe (0xfffff00-2) in the *AsyncLock* (release) message.

**Table 15 Lock Behavior**

Initial State	Lock request	Client	New State
Unlocked (initial state)	Shared lock	any client → success	Shared Locked
	Exclusive lock	any client → success	Exclusive Locked
	Release	any client → error	Unlocked
Exclusive locked	Shared Lock	holder of exclusive lock → error	Exclusive Locked
		other client → fails after lock timeout	Exclusive Locked
	Exclusive lock	client not holding the lock → fails after lock timeout	Exclusive Locked
		holder of exclusive lock → error	Exclusive Locked
	Release	holder of exclusive → succeeds	Unlocked
		client without exclusive lock → error	Exclusive locked
Shared locked	Shared lock	holder of shared lock → error	Shared Locked
		other client with right key → succeeds	Shared Locked
		other client with wrong key → fails after lock timeout	Shared Locked
	Exclusive lock	holder of shared lock → succeeds	Both Locks
		client not holding shared lock → fails after lock timeout	Shared Locked
	Release	holder of the shared lock when 2 or more are holding shared lock → succeeds	Shared Locked
		the only remaining holder of the shared lock → Succeeds	Unlocked
		client holding no locks → error	Shared Locked
Both locks	Shared Lock	holder of shared lock → error	Both locks
		other client with right key → succeeds	Both locks
		other client with wrong key → fails after lock timeout	Both locks
	Exclusive lock	client not holding the exclusive lock → fails after timeout	Both locks
		client holding the exclusive lock → error	Both locks
	Release	holder of exclusive and shared lock → succeeds	Shared locked
		other client holding a shared lock → succeeds	Both Locks
		client holding no locks → error	Both locks

HiSLIP servers shall:

1. Go to the Unlocked state when the first connection is initialized
2. Release all locks assigned to a client when that client connection closes

#### 4.5.1 Unlock Considerations

In an unlock operation, the MessageID in the message parameter designates the last *Data*, *DataEND* or *Trigger* message to be completed before the lock is released. The client is only allowed to specify messages that were transmitted before the *AsyncLock* operation. The *AsyncLock* transaction will not complete the unlock until that message is complete.

The *AsyncLock* unlock operation can only be aborted by device clear. The normal device clear behavior will abandon any pending transactions the unlock operation may be waiting for. Note that this requires the server respond in a timely fashion to an *AsyncDeviceClear* message while waiting for an *AsyncLock* unlock operation to complete. When an *AsyncLock* unlock operation is abandoned by a device clear, the lock shall be released per the pending unlock operation and no confirmation sent to the client.

## 4.6 Lock Info Transaction

Table 16 Lock Info Transaction

Step	Sender	Message	Action
1	Client	<AsyncLockInfo><0><0><0>	Request from the client for lock information.
2	Server	<AsyncLockInfoResponse><exclusive-locks-granted><locks-granted><0>	The server returns information regarding locks it has granted.

The following are the fields in the *AsyncLockInfoResponse*:

exclusive-locks-granted    1 if an exclusive lock has been granted and 0 otherwise.

locks-granted                the number of clients that were holding locks when *AsyncLockInfo* was processed. A client holding both a shared and exclusive lock is counted only once.

The *LockInfo* transaction is used by the client to determine how many other clients are connected and how many locks have been granted. These values are sampled values from the server and may be inaccurate since other clients may be simultaneous releasing and requesting locks or connections. However, the locks-granted and clients-connected values shall be self-consistent at some point in time when the *LockInfo* was processed by the server.

This transaction is processed regardless of whether the client currently holds a lock.



## 4.7 Remote Local Transaction

HiSLIP supports GPIB-like remote/local control. The purpose of remote/local is to:

- Prevent front panel input from interfering with remote operations
- Permit front panel local key to re-enable the front panel input
- Provide a way to lockout the local key<sup>7</sup> when the controller needs exclusive access

**Table 17 RemoteLocal Control Transaction**

Step	Sender	Message	Action
1	Client	<AsyncRemoteLocalControl><request><MessageID><0>	Request remote local operation
2	Server	<AsyncRemoteLocalResponse><0><0><0>	Confirm remote/local request

request                      The values of the request field are shown in Table 18, *Remote Local Control Transactions*

MessageID                    designates the MessageID of the most recent *Data*, *DataEND*, or *Trigger* message sent by the client.

The server is permitted to act on the *AsyncRemoteLocalControl* immediately, or wait until after the preceding operations have been acted on by the server.

In some conditions, TCP may deliver the remote/local requests before it delivers a preceding *Data/DataEND* or *Trigger* message generated by the client. The server should consult the MessageID provided with the *AsyncRemoteLocalControl*. If this MessageID is not equal to the MessageID of the last received *Data*, *DataEND*, or *Trigger* message then the remote/local request should be deferred until after the message (on synchronous channel) with the designated MessageID was processed.

The server shall send the *AsyncRemoteLocalResponse* after any server defined actions are complete, however it shall not wait for locks granted to other clients to be released. Although the *AsyncRemoteLocalResponse* is sent immediately if another client is holding a lock, the server shall only act on the remote/local request after the lock is released. The behavior is device dependent.

Three logical variables maintained by the server dictate its behavior:

Remote	Controls if front panel input is accepted. Note that remote input is always accepted. If Remote is true, front panel input is not accepted, with the exception of the local key. If the local key is pressed and LocalLockout is set to false, Remote is set false so that subsequent front panel input is accepted. When the connection is initialized Remote is false.
RemoteEnable	Mimics the GPIB REN line, but is maintained by the individual server. When the connection is initialized RemoteEnable is true.
LocalLockout	If true, the front panel local key has no affect. If false, the front panel local key sets Remote to false. When the connection is initialized LocalLockout is false.

<sup>7</sup> The local key, as defined by IEEE 488, is a key on the instrument an operator can use to gain front panel access to the instrument when front panel access has been automatically disabled by the 488 protocol.

If RemoteEnable is true and new data or control information arrives via the Hi-SLIP protocol, Remote is set to true. Specifically, any of the following messages on the synchronous channel set remote true, provided that a deferred implication is processed:

- Data
- DataEND
- Trigger

Or any of the following messages on the asynchronous channel:

- AsyncStatusQuery
- AsyncDeviceClear
- AyncLock

Servers are permitted to take a device specific action for *VendorSpecific* messages.

*RemoteLocal* HiSLIP messages set these state variables as described in Table 18. In that table, T indicates the variable is set, F indicates the variable is cleared, and nc indicates the variable is not changed.

The remote/local control codes correspond to the parameters to the VISA viGpibControlREN<sup>8</sup> function call. The behavior is chosen to emulate the behavior of a GPIB device.<sup>9</sup>

After the Initialization and Device Clear transaction, the client shall use the MessageID = 0xffffefe (0xfffff00-2) in the *AsyncRemoteLocalControl* message.

<sup>8</sup> The VISA specification (vpp43, Table 6.5.1) specifies the following:

Mode	Action Description
VI_GPIB_REN_DEASSERT	Deassert REN line.
VI_GPIB_REN_ASSERT	Assert REN line.
VI_GPIB_REN_DEASSERT_GTL	Send the Go To Local command (GTL) to this device and deassert REN line.
VI_GPIB_REN_ASSERT_ADDRESS	Assert REN line and address this device.
VI_GPIB_REN_ASSERT_LLO	Send LLO to any devices that are addressed to listen.
VI_GPIB_REN_ASSERT_ADDRESS_LLO	Address this device and send it LLO, putting it in RWLS.
VI_GPIB_REN_ADDRESS_GTL	Send the Go To Local command (GTL) to this device.

<sup>9</sup> The VISA API provides general control of GPIB that is not necessary for a Hi-SLIP client. Practical Hi-SLIP applications can be handled by using just three values for the mode: VI\_GPIB\_REN\_DEASSERT which will always place the instrument in local, VI\_GPIB\_REN\_ASSERT\_ADDRESS\_LLO which will always put the instrument into remote with local-lockout, and VI\_GPIB\_REN\_ASSERT\_ADDRESS which will place the instrument into remote, but enable the front panel local key (with automatic transitions back to remote when remote data is received). Unfortunately, the names of these modes are not very mnemonic.

**Table 18 Remote Local Control Transactions**

Control Code (request)	Corresponding VISA mode from viGpibControlREN	Behavior		
		RemoteEnable	LocalLockout	Remote
0 – Disable remote	VI_GPIB_REN_DEASSERT	F	F	F
1 – Enable remote	VI_GPIB_REN_ASSERT	T	nc	nc
2 – Disable remote and go to local	VI_GPIB_REN_DEASSERT_GTL	F	F	F
3 – Enable remote and go to remote	VI_GPIB_REN_ASSERT_ADDRESS	T	nc	T
4 – Enable remote and lock out local	VI_GPIB_REN_ASSERT_LLO	T	T	nc
5 – Enable remote, go to remote, and set local lockout	VI_GPIB_REN_ASSERT_ADDRESS_LLO	T	T	T
6 – go to local without changing state of remote enable	VI_GPIB_REN_ADDRESS_GTL	nc	nc	F

If multiple clients make changes the behavior shall be the same as if a single client made all the requests serially in whatever order the requests are handled by the server.

On closing the connection, the remote local behavior is defined by the server.

## 4.8 Trigger Message

Table 19 Trigger Message

Step	Sender	Message	Action
1	Client	<Trigger><RMT-delivered><MessageID><0>	Initiate a trigger

The trigger message is used to emulate a GPIB Group Execute Trigger. This message shall have the same instrument semantics as GPIB Group Execute Trigger.

The fields in the *Trigger* message are:

RMT-delivered	RMT-delivered is 1 if this is the first <i>Data</i> , <i>DataEND</i> , <i>Trigger</i> , or <i>StatusQuery</i> message since the HiSLIP client delivered an RMT to the client application layer.
MessageID	MessageID identifies this message so that response data from the server can indicate the message that generated it. For generation of the MessageID, see sections 3.1.2 ( <i>Synchronized Mode Client Requirements</i> ) and 3.2.2 ( <i>Overlap Mode Client Requirements</i> ).

## 4.9 Vendor Defined Transactions

Table 20 Vendor Defined Transaction

Step	Sender	Message	Action
1	Either	<VendorDefined><arbitrary><arbitrary ><length><payload>	Vendor defined
2		Response – if unrecognized non-fatal error, if recognized vendor defined.	

*VendorDefined* messages may be used arbitrarily by vendors on either the synchronous or asynchronous channels. Clients or servers that do not recognize *VendorDefined* messages shall ignore the message including the number of subsequent data bytes.

Devices or Servers receiving VendorDefined commands they do not support shall respond with an *Error* message on the same channel the Vendor defined message arrived on specifying “Unrecognized Vendor Defined Message”.

## 4.10 Maximum Message Size Transaction

Table 21 Maximum Message Size Transaction

Step	Sender	Message	Action
1	Client	<AsyncMaximumMessageSize><0><0><8><8-byte size>	The server sets the maximum message size it will send to the client to the specified value
2	Server	<AsyncMaximumMessageSizeResponse><0><0><8><8-byte size>	The client sets the maximum message size it will send to the server to the specified value

The *AsyncMaximumMessageSize* transaction is used to inform the client and server of the maximum message size they are permitted to send to the other on the synchronous channel. This is especially important for small devices that may be unable to handle large messages.

The *AsyncMaximumMessageSize* transaction is initiated by the client. Neither clients nor servers are obligated to accept a particular message size beyond what is necessary during initialization. Therefore it is prudent for clients to initiate this transaction as part of initialization to inform the server of its message size limitations and determine the server limitations.

The specified message sizes only apply to the synchronous channel.

The 8-byte buffer size is sent in network order as a 64-bit integer.

Servers shall keep independent client message sizes for each HiSLIP connection.

## 4.11 Interrupted Transaction

Table 22 Interrupted Transaction

Step	Sender	Message	Action
1	Server	<AsyncInterrupted><0><MessageID><0>	Clear buffered messages.
2	Server	<Interrupted><0><MessageID><0>	Clear buffered messages.

The interrupted transaction is sent from the server to the client when the server detects an interrupted protocol error. The client shall clear any buffered *Data*, *DataEND*, or *Trigger* messages from the server and ignore any subsequent *Data*, *DataEND*, or *Trigger* messages until it has received both the synchronous *Interrupted* message and asynchronous *AsyncInterrupted* messages arrive.

The MessageID field indicates the MessageID of the *Data*, *DataEND*, or *Trigger* message that interrupted the server response.

For use of the interrupted transaction, see section 3 (*Overlapped and Synchronized Modes*).

## 4.12 Device Clear Transaction

Device clear clears the communication channel.

**Table 23 Device Clear Complete Transaction**

Step	Sender	Message content	Action
1	Client	<AsyncDeviceClear><0><0><0>	
----	Client	complete messages underway and abandon any pending messages	Abandon pending messages and wait for in-process synchronous messages to complete
2	Server	<AsyncDeviceClearAcknowledge><featurePreference><0><0>	The client shall wait for this acknowledgement before additional processing.
3	Client	<DeviceClearComplete><featureRequest><0><0>	Indicate to server that synchronous channel is cleared out.
4	Server	NA	Upon receipt of the sync or async clear messages abandon any operations in progress.
5	Server	NA	Disregard input messages until the <i>DeviceClearComplete</i> message is found. But continue to require well-formed messages.
6	Server	<DeviceClearAcknowledge><featureSetting><0><0>	Client and server each resume normal operation.

To send a device clear, the client will:

1. Finish sending any partially sent messages on either channel.
2. Send the *AsyncDeviceClear* message on the asynchronous channel.
3. If the protocol was amidst any of the following transactions permit them to complete if the server responds before sending *DeviceClearAcknowledge*
  - a. *Lock (client waiting for AsyncLockResponse)*  
If *DeviceClearAcknowledge* arrives from the server before these other operations are acknowledged, the client HiSLIP shall assume the operations were not completed.
4. Clear messages on the synchronous and asynchronous channels with the exception of *FatalError*, *DeviceClearAcknowledge*, and *AsyncDeviceClearAcknowledge*.
5. Wait for the *AsyncDeviceClearAcknowledge* message.
6. Send the *DeviceClearComplete* message on the synchronous channel indicating to the server that no further messages will be sent to it.
7. Wait for the server to respond with *DeviceClearAcknowledge* on the synchronous channel.
8. *The MessageID is reset to 0xffff ff00.*
9. Resume normal operation.

When the server receives the asynchronous *AsyncDeviceClear* message, it shall:

1. Complete any partially complete transactions involving the asynchronous channel without waiting for timeouts.
2. Send *AsyncDeviceClearAcknowledge*.
3. Finish sending any partially sent messages to the client. Complete with the normal behaviors, without waiting for any timeouts.
4. Abandon any buffered unsent transactions.
5. Clear any well-formed messages received from the client on the synchronous channel.



6. Accept and ignore subsequent synchronous messages until it finds the synchronous *DeviceClearComplete* message.
7. Send *DeviceClearAcknowledge* message back to the client (after the above steps are complete and it has received the device clear complete message).
8. Resume normal operation.

If at any time during device clear management either the client or server encounter a poorly formed message they shall send a *FatalError* message and do the *FatalError* processing.

If at any time during the device clear management the client or server determines that the other device is not responding in a timely fashion, it shall send a *FatalError* message and do the *FatalError* processing. The determination of an appropriate time may vary with the application, 40 to 120 seconds are reasonable values.

#### 4.12.1 Feature Negotiation

During device clear, the features listed in Table 24 are negotiated between the client and server. The features are specified through a feature bitmap that is sent in the control code of three different messages.

The feature negotiation occurs in three steps:

1. The server proposes values that it prefers with the *AsyncDeviceClearAcknowledge* message.
2. The client indicates values that it requests in the *DeviceClearComplete* message.
3. The server indicates the values that both client and server will use in the *DeviceClearAcknowledge* message.

The server shall identify the default values it prefers for the features in the *AsyncDeviceClearAcknowledge* message. Servers shall support any such capabilities that it requests.

The server shall accept the value proposed by the client in the *DeviceClearComplete* message if it is capable of supporting them.

The client shall use the values specified by the server in the *DeviceClearAcknowledge* message.

Note that the *InitializeResponse* has the server preferred feature specified in it.

**Table 24 Features negotiated during device clear**

Control Code Bit Position	Name	Meaning
0	Overlapped	False- Synchronized mode True - Overlapped mode

### 4.13 Service Request

Table 25 Service Request

Step	Sender	Message content	Action
1	Server	<AsyncServiceRequest><status><0><0>	Client Initiated request for service

The server requests service by sending the *AsyncServiceRequest* message.

The control code contains the 8-bit status register.

No values, including rqs (request service), are cleared in the status register. Note that since no values are cleared, the client must do a *AsyncStatusQuery* to clear the rqs bit and enable additional *AsyncServiceRequest* messages.

This message is not acknowledged.

## 4.14 Status Query Transaction

**Table 26 Status Message**

Step	Sender	Message content	Action
1	Client	<AsyncStatusQuery><RMT-delivered><MessageID><0>	Client Initiates a request for status
2	Server	<AsyncStatusResponse><status><0><0>	Status information is sent back in the control code field

The status query provides a 8-bit status response from the server that corresponds to the VISA viReadSTB operation. The status query is initiated by the client and sent on the asynchronous channel.

The calculation of the message available bit (MAV) of the status response differs for overlapped and synchronized modes and requires the client to provide a different Message ID.

The following are the fields of the *AsyncStatusQuery* client message:

RMT-delivered	RMT-delivered is 1 if this is the first <i>Data</i> , <i>DataEND</i> , <i>Trigger</i> or <i>AsyncStatusRequest</i> message since the client delivered RMT to the application layer. Note that RMT-delivered is only reported once.
MessageID	In synchronized mode, this field contains the MessageID of the most recent <i>Data</i> , <i>DataEND</i> , or <i>Trigger</i> message sent by the client.  In overlapped mode, this field contains the MessageID of the most recent <i>Data</i> or <i>DataEND</i> message delivered to the client application layer.

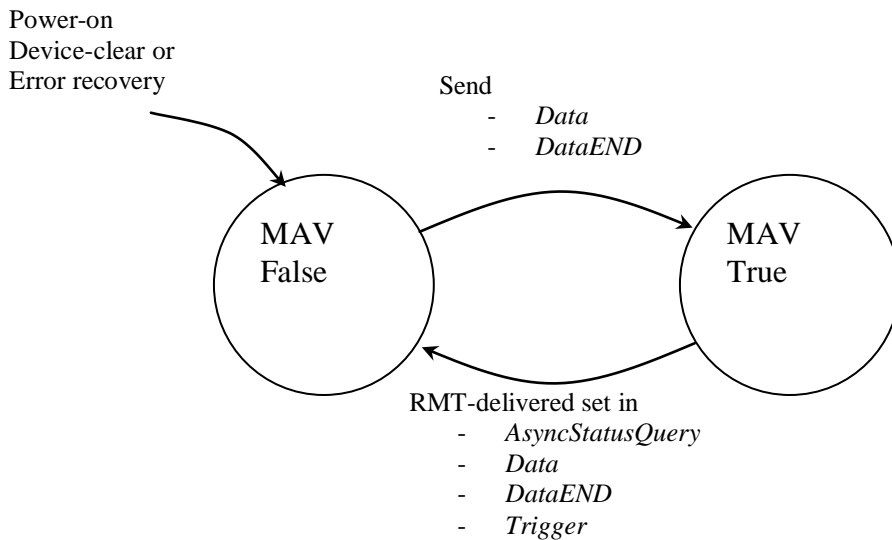
The following are the fields of the *AsyncStatusResponse* server message:

status	This field contains an 8-bit status response from the server.
--------	---

When clients send *StatusQuery* messages, they shall set the message header to the MessageID field of the message header of the most recently sent *Data*, *DataEND*, or *Trigger* message. See section 3, *Overlapped and Synchronized Modes* for the server construction of the status response.

After the Initialization and Device Clear transactions, the client shall use the MessageID = 0xffffefe (0xfffff00-2) in the *AsyncStatusQuery* message.

### 4.14.1 MAV Generation in Synchronized Mode



**Figure 1 MAV Generation in Synchronized Mode**

HiSLIP asynchronously reads the status back from the server using the *StatusQuery* message on the asynchronous channel. However, IEEE Std 488.2 requires servers to include a MAV (Message Available) bit in position 4 that indicates if data is available from the server. The RMT-delivered indication sent from the client to the server assists the server in determining the correct value for MAV.

Per IEEE Std 488.2 the MAV bit shall be sent in bit 4 (zero-based) of the status response from the server.

MAV shall be set true when the server sends the first *Data* or *DataEND* of a response.

MAV shall be set false when RMT-delivered is indicated by the client in either the *Data*, *DataEND*, *Trigger* or *AsyncStatusQuery* message.

Figure 1 shows how the MAV shall be calculated.

Note that new-reason-for-service is only asserted on the transitions between these states. Therefore, a *AsyncServiceRequest* is only generated for MAV once per message.

### 4.14.2 MAV Generation in Overlapped Mode

In overlapped mode, the server shall compare the MessageID specified in the *AsyncStatusQuery* to the current MessageID counter. If any messages have not been fully delivered to the client application, MAV shall be set true.

### 4.14.3 Implementation Note

In some conditions, TCP may deliver a *AsyncStatusQuery* before it delivers a preceding *Data/DataEND* or *Trigger* message generated by the client.

In synchronized mode, the server should consult the MessageID provided with the AsyncStatusQuery. If this MessageID is not equal to the MessageID of the last received *Data*, *DataEND*, or *Trigger* message then MAV shall be set false. Note that in this special case the server either has no data for delivery or will be interrupted with the next pending synchronous message. In the case where MAV is reported as false because of a pending interrupted error, it is not necessary for the server to detect and report the error in this status response however the interrupted error will be reported subsequently with normal interrupted processing.

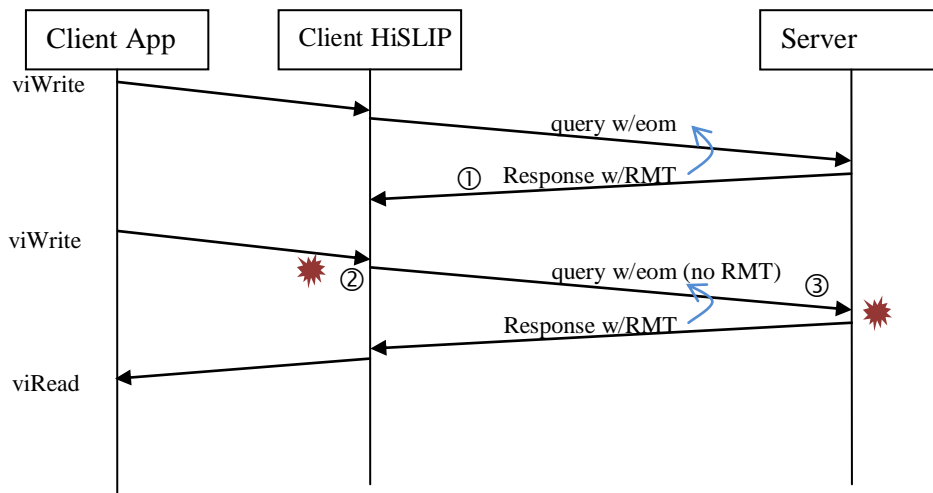
In overlapped mode, the provided MessageID strictly indicates the availability of new data. Therefore, client should never presume that the absence of message available indicates additional data will not be made available later.

## A. Analysis of Interrupted Conditions

This is an informative appendix and is not part of the HiSLIP standard requirements.

The following transaction diagrams describe HiSLIP behavior in synchronized mode with various interrupted conditions.

### A.1 Slow Client

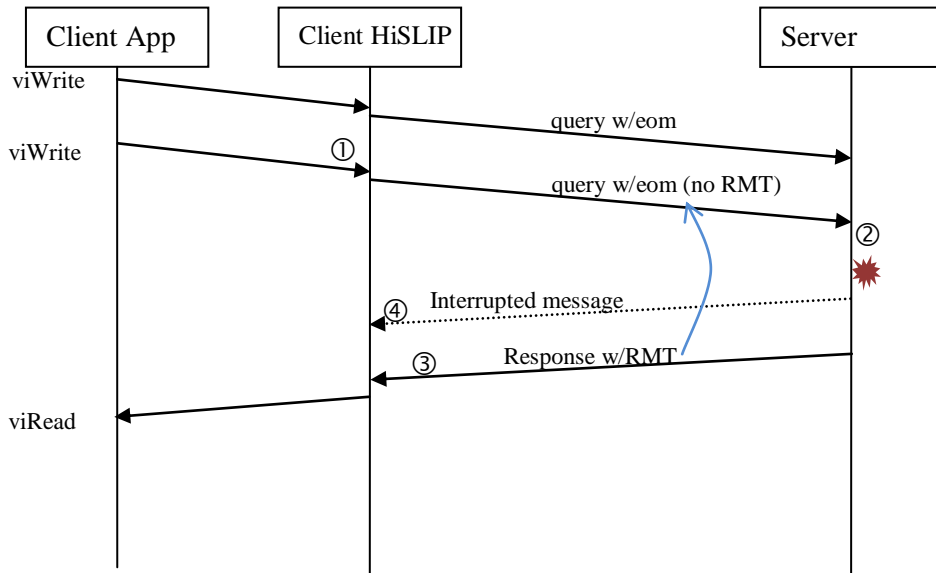


**Figure 2 Interrupted error with slow client**

Detailed explanation:

1. At this point the response indicates the eom from the preceding message.
2. Note that the client probably will not get a chance to execute until the client application calls viWrite. However, the client HiSLIP has the opportunity to take care of input processing before attempting to send the second write message. At this point, the client detects the error based on section 3.1.2 rule 3 and clears the first response from its buffer. (this error detection is essential at this point to ensure that the buffered response is not provided to a subsequent viRead). The client then sends the second query normally.
3. Note that the second query indicates that the RMT was NOT delivered to the application layer. Therefore the server will also detect the error based on 3.1.1 rule 2. Since the last action by the server was to send RMT, no error handling is necessary other than reporting the error.

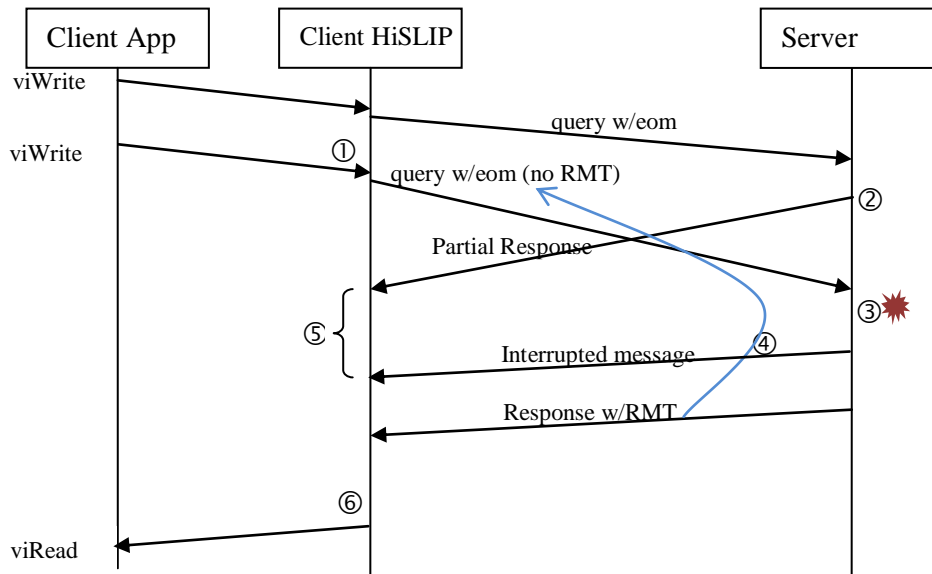
## A.2 Fast Client



**Figure 3 Interrupted error with fast client**

Detailed explanation:

1. Note that from the perspective of the client HiSLIP, there may not be a problem, since the data being written may be commands.
2. Per section 3.1.1 rule 1 The server detects interrupted because it has a complete response (with RMT) and the input buffer is not empty. The response to the first query is never sent.
3. The response to the second query is sent normally.
4. NOTE – although not required here, the server sends an interrupted message (as shown in *Figure 4 Interrupted error with fast client and partial response*).



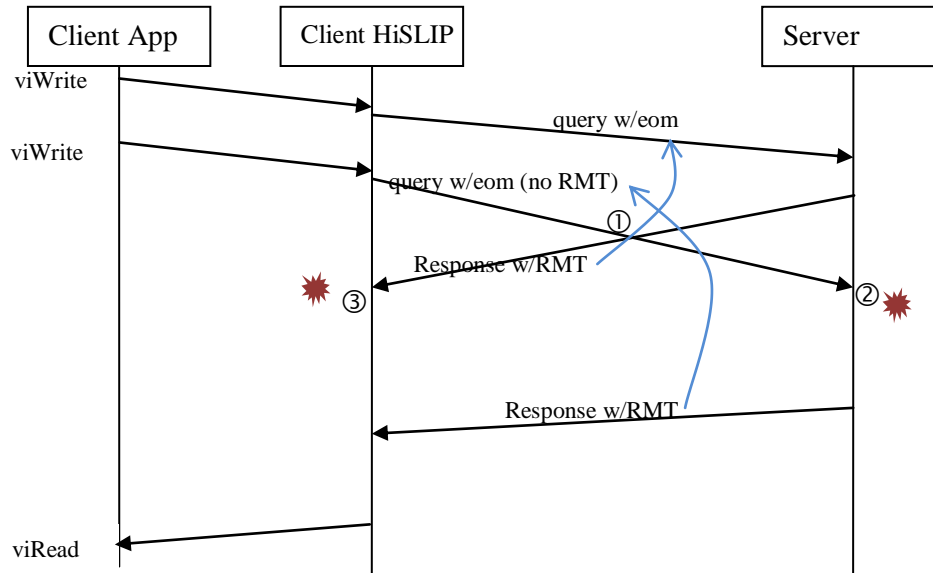
**Figure 4 Interrupted error with fast client and partial response**

Detailed explanation:

1. Note that from the perspective of the client HiSLIP, there may not be a problem, since the data being written may be commands.
2. Per IEEE Std 488.2, the server chooses to send a partial response (without RMT) to the client. IEEE Std 488.2 permits delivering this to the client, but the RMT corresponding to the first query must not be delivered.
3. Per section 3.1.1 rule 1, the server detects interrupted because it has a complete response (with RMT) and the input buffer is not empty (note the server is still completing processing on the first query). The final portions of the response to the first query are not sent.
4. Server informs the client that the partial response should be cleared if not already delivered.
5. The partial response is only delivered to the client if the client application attempts to read before the *Interrupted* or *AsyncInterrupted* message arrives and is detected by the client protocol. In this illustration, HiSLIP client will not deliver the partial response.
6. The response to the second query is sent normally.



### A.3 Intermediate Timing



**Figure 5 Race condition: first response and second query pass in-flight**

Detailed explanation:

1. Second query and first response cross in-flight. Note that it is essential to have some error detection on the client in this case, otherwise this errant response would be delivered.
2. Based on section 3.1.1 rule 2, the server detects the error and reports it. Since it has already launched the first response it takes no additional action.
3. Based on section 3.1.2 rule 1, the client detects the stale response and clears it without offering any to the client app.