
GIOP SCTP Protocol Mapping

Initial Submission

Submitted by

**Highlander Engineering Inc.
Nokia**

Supported by

Data Connection, Ltd.

*April 2, 2001
OMG Document orbos/2001-04-02*

Copyright © 2001 Highlander Engineering Inc.
Copyright © 2001 INPRISE Corporation.
Copyright © 2001 Nokia.
Copyright © 2001 Data Connection, Ltd.

All Rights Reserved.

The companies and organisations listed above hereby grant a royalty-free license to the Object Management Group (OMG) for world-wide distribution of this document or any derivative works thereof, so long as the OMG reproduces the copyright notices and the below paragraphs on all distributed copies.

The material in this document is provided for evaluation by the OMG. Submission of this document does not represent a commitment to implement any portion of this specification in the products of the above said companies and organisations.

While the information in the publication is believed to be accurate, the companies and organisations listed above 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 companies and organisations listed above shall not be liable for errors contained herein or for incidental or consequent damages in connection with the furnishing, performance or use of this material. The information contained in this document is subject to change without notice.

This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright here on may be reproduced or used in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without permission of the copyright owners. All copies of this document must include the copyright and other information contained on this page.

The copyright owners grant the Object Management Group permission to reproduce and use the information contained in this document. The copyright owners grant a limited waiver of the copyright on this document to members of the Object Management Group so that they can each reproduce up to 50 copies of this document for their internal use as part of the OMG evaluation process.

RIGHTS RESERVED LEGEND. Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c)(1)(ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013.

All trademarks acknowledged.

Contact Persons:

Jon Currey
Highlander Engineering Inc.
208 East Pine Street
Lakeland, FL 33801
USA
Phone: +1 863 686 7767
Fax: +1 863 687 7767
Email: jon@highlander.com

Dr. Kimmo Raatikainen
Nokia Research Center
P.O. Box 407
FIN-00045 NOKIA GROUP
Finland
Phone: +358 7180 36275
Fax: +358 7180 36308
Email: kimmo.raatikainen@nokia.com

Jon Berger
Data Connection Ltd.
100 Church Street
Enfield
Middlesex
EN2 6BQ
UK
Phone: +44 20 8366 1177
Fax: +44 20 8363 1039
Email: jrb@dataconnection.com

This document specifies an architecture and interfaces to support a GIOP Mapping over the SCTP Protocol as requested in the RFP orbos/2000-12-18

1.1 Submission Overview

The submission is organized as follows.

In the current chapter we describe our design rationale and make our statement of proof of concept. We also address each mandatory and optional requirement in the RFP as well as issues to be discussed.

In Chapter 2 we present the Architectural Framework. In Chapter 3 each of the key concepts, that is SCTP IOP and SCTP IOR, is described in details.

The associated IDL specifications are provided in Chapter 3.

1.2 Design Rationale

The basic design principles have been client-side and server-side ORB application transparency, and interoperability.

The RFP requested a mapping of GIOP onto SCTP using IPv6, and the IOR profile for this mapping. This submission addresses this requirement, as well as the optional requirement of using SCTP with IPv4 addresses, by making the “host address” portion of the IOR a string.

1.3 Proof of Concept

The design is affected by experiences of the submitters in implementing IOP in VisiBroker as well as the pluggable protocol interface which allows non-TCP protocols to be used by the ORB.

1.4 Mandatory Requirement

GIOP mapping and IOR profile for SCTP using IPv6 .

Responses to this RFP shall provide a mapping of GIOP onto SCTP using IPv6

Responses to this RFP shall define an IOR profile for SCTP using IPv6

This Proposal meets this requirement by specifying host addresses as strings.

1.5 Optional Requirements

GIOP mapping and IOR profile for SCTP using IPv4 .

Responses to this RFP shall provide a mapping of GIOP onto SCTP using IPv4

Responses to this RFP shall define an IOR profile for SCTP using IPv4

This Proposal meets this requirement by specifying host addresses as strings.

GIOP mapping and IOR profile for SCTP using IPv4 .

Responses to this RFP may define an interoperability bridge between the IOP and the IOP defined in the response.

This Proposal does not address this optional requirement.

1.6 Issues to be discussed

Proposals shall discuss the interworking issues between IPv6 and IPv4.

It is not clear to the submitters, what the RFP is requesting here since interworking issues between IPv6 and IPv4 seem to be an IETF issue.

1.7 References

IETF, Stream Control Transmission Protocol, RFC 2960, October 2000.

The key concepts in this specification are: SCIOP and the SCIOR. This proposal will use TCP/IP as the baseline protocol to describe the differences introduced by the SCTP protocol.

The **SCIOP** is the concrete mapping of GIOP over SCTP, and the usage of SCTP associations, and streams, analogous to IIOP's usage of TCP/IP connections.

The **SCIOR** is an object reference profile analogous to the IIOP 1.2 IOR profile. It identifies the host address(es) and the port on which the target object resides.

The SCTP protocol is ideally suited as a protocol onto which GIOP may be mapped directly: Quoting from OMG/00-11-07 Chapter/Section 15.5:

15.5 GIOP Message Transport

The GIOP is designed to be implementable on a wide range of transport protocols. The GIOP definition makes the following assumptions regarding transport behavior:

- *The transport is connection-oriented. GIOP uses connections to define the scope and extent of request IDs.*

SCTP is connection oriented like TCP. SCTP refers to a connection as an SCTP Association. Within an SCTP Association, there may be multiple Streams (independent unidirectional data streams).

- *The transport is reliable. Specifically, the transport guarantees that bytes are delivered in the order they are sent, at most once, and that some positive acknowledgment of delivery is available.*

SCTP fulfills this requirement.

- *The transport can be viewed as a byte stream. No arbitrary message size limitations, fragmentation, or alignments are enforced.*

While SCTP is not a byte stream, it is actually a reliable packet based protocol, (i.e. a reliable, connection oriented UDP), within SCTP, there may be a message size limitation based on the MTU. While SCTP implementation must support receiving and reassembly of fragmented user data messages, it is optional that an implementation support fragmentation of user data messages sent. This can be resolved by setting the GIOP Fragment size within the sending GIOP engine. There are no arbitrary alignment requirements in SCTP.

- *The transport provides some reasonable notification of disorderly connection loss. If the peer process aborts, the peer host crashes, or network connectivity is lost, a connection owner should receive some notification of this condition.*

SCTP fulfills this requirement.

- *The transport's model for initiating connections can be mapped onto the general connection model of TCP/IP. Specifically, an agent (described herein as a server) publishes a known network address in an IOR, which is used by the client when initiating a connection.*

SCTP fulfills this requirement.

The server does not actively initiate connections, but is prepared to accept requests to connect (i.e., it listens for connections in TCP/IP terms). Another agent that knows the address (called a client) can attempt to initiate connections by sending connect requests to the address. The listening server may accept the request, forming a new, unique connection with the client, or it may reject the request (e.g., due to lack of resources). Once a connection is open, either side may close the connection. (See Section 15.5.1, "Connection Management," on page 15-46 for semantic issues related to connection closure.) A candidate transport might not directly support this specific connection model; it is only necessary that the transport's model can be mapped onto this view.

SCTP can be utilized in this fashion.

The above notwithstanding, SCTP does pose some issues to be resolved to be used as an effective protocol for making CORBA invocations:

SCTP Endpoints: There are several differences between TCP/IP endpoints and SCTP endpoints worthy of note:

TCP/IP endpoints are specified by a host/port pair. SCTP endpoints are specified by a list of host addresses, and a single port. This is used by SCTP to provide a low-level fault tolerance by allowing SCTP associations to exist across multiple interfaces between the communicating hosts. The server-side SCTP endpoint will specify which host-addresses (which physical interfaces on that host) may be used to communicate with that endpoint. Likewise the client-side SCTP endpoint similarly may choose which of its physical interfaces will be used when creating an association with the server.

SCTP Associations & Streams: SCTP associations are analogous to TCP/IP connections with the following caveat: A TCP/IP connection represents a single bi-directional stream, while a SCTP association may contain multiple uni-directional streams. In TCP/IP the application writes data to the TCP/IP endpoint (socket). In SCTP, the application writes data to an SCTP-endpoint, SCTP-association ID, and Stream-ID. To summarize, once a SCTP association is established, each stream contained within it, is effectively connected as well.

The creator of an SCTP endpoint specifies the maximum number of Streams which that endpoint will accept and use (outbound-streams, and inbound-streams). The Client-side initiator of an SCTP association requests the number of streams it would like to use, and the lesser of the two values of the server-side inbound-streams and the client-side outbound-streams will be used.

SCTP Inter-ORB Protocol (SCIOP)

Specific APIs for libraries supporting SCTP/IP vary, so this discussion is limited to an abstract view of SCTP/IP and management of its connections. The mapping of GIOP message transfer to SCTP/IP associations (and streams within associations) is called the SCTP Inter-ORB Protocol (SCIOP). SCIOP is based on GIOP 1.2. and may support all components found in IIOP 1.2. It is the submitters' intent that the SCIOP will track and use the same version number as IIOP.

SCTP Association Usage:

Agents that are capable of accepting object requests or providing locations for objects (i.e., servers) publish SCTP addresses in IORs, as described below. A SCTP address consists of a sequence of IP host addresses, (each address may be represented by a host name, IPV4 address, or IPv6 address), and a SCTP port number. Servers must listen for association connection requests.

A client needing an object's services must initiate a connection with the SCTP endpoint specified in the IOR, with an association connect request.

The listening server may accept or reject the connection request. In general, servers should accept association connection requests if possible, but ORBs are free to establish any desired policy for association connection acceptance (e.g., to enforce fairness or optimize resource usage).

Once an association connection is accepted, all streams contained within that association are deemed to be also connected and each stream represents an individual GIOP connection. Note that per the IETF SCTP specification (RFC2960), streams represent uni-directional flows, and the `STREAM_ID` used during a send operation indicates the target (receiver's) Stream to which a datagram will be sent. An SCTP user: Upper Layer protocol (ULP) can use the streams as a bi-directional flow by using the same stream-id on which a message was received to send the reply messages. This bi-directional usage of Stream_IDs shall be used by SCIOP. The client may send

Request, LocateRequest, CancelRequest, or CloseConnection messages by writing to the SCTP Association_ID, and target Stream_ID for the contained GIOP connections.

The server may send **Reply, LocateReply, and CloseConnection** messages by writing to its SCTP Association_ID, and target Stream_ID. If BiDirectional GIOP is in use, the client may also send **Reply** and **LocateReply** messages.

After receiving a **CloseConnection** message, an ORB must close the GIOP connection associated with that Stream_ID. Any subsequent GIOP messages should be assumed to be within a new GIOP connection for this Stream as if a new association was formed. This will give the same behavior as an individual TCP/IP connection for that Stream_ID. After sending a **CloseConnection** message on all Streams contained within an association, an ORB may close the SCTP association immediately, or may delay closing the association until it receives an indication that the other side has closed the association.

3.1 SCIOP IOR Profiles:

SCIOP IOR profiles, identifying individual objects accessible through the SCTP Inter-ORB Protocol, have the following form:

```

module SCIOP { //
  struct Version {
    octet major;
    octet minor;
  };

  struct ProfileBody_1_2 { //
    Version SCIOP_version;
    sequence <string> host_addresses;
    unsigned short port;
    unsigned short max_streams;
    sequence <octet> object_key;
    sequence <IOP::TaggedComponent> components;
  };

```

SCIOP Profile version number:

- Indicates the SCIOP protocol version.
- Major number can stay the same if the new changes are backward compatible.
- Clients with lower minor version can attempt to invoke objects with higher minor version number by using only the information defined in the lower minor version protocol (ignore the extra information).

An instance of one of these structure types is marshaled into an encapsulation octet stream. This encapsulation (a **sequence <octet>**) becomes the **profile_data** member of the **SCIOP::TaggedProfile** structure representing the SCIOP profile in an IOR, and the tag has the value **TAG_SCTP_IOP** (to be assigned by OMG).

The version number published in the Tag SCTP IOP Profile body signals the highest GIOP minor version number that the server supports at the time of publication of the IOR.

No order of use is prescribed in the case where more than one TAG SCTP IOP Profile is present in an IOR.

The members of **SCIOP::ProfileBody_1_2** are defined as follows:

- **sciop_version** describes the version of SCIOP that the agent at the specified address is prepared to receive. When an agent generates SCIOP profiles specifying a particular version, it must be able to accept messages complying with the specified version or any previous minor version (i.e., any smaller version number). The major version number of this specification is 1; the only minor version number defined to date is 2. Compliant ORBs must generate version 1.2 profiles, and must accept any profile with a major version of 1, regardless of the minor version number. If the minor version number is 2, the encapsulation is fully described by the **ProfileBody_1_2** structure. If the minor version number is greater than 2, then the length of the encapsulated profile may exceed the total size of components defined in this specification for profiles with minor version number 1 or 2. ORBs that support only version 1.2 SCIOP profiles must ignore, but preserve, any data in the profile that occurs after the components member, for SCIOP profiles with minor version greater than 1.2

Note – As of version 1.2 of GIOP and SCIOP and minor versions beyond, the minor version in the **TAG_SCTP_IOP** profile signals the highest minor revision of GIOP supported by the server at the time of publication of the IOR.

- **host_addresses** identifies a list of Internet host addresses supported by the server's association, to which a client will establish an SCTP association on which GIOP messages for the specified object may be sent. In order to promote a very large (Internet-wide) scope for the object reference, if an address within this list is a host name this will typically be the fully qualified domain name of the host, rather than an unqualified (or partially qualified) name. However, per Internet standards, the individual hostaddress strings may also contain a host address expressed in standard "dotted decimal" form (e.g., "192.231.79.52").
- **port** contains the SCTP port number (at all of the specified host addresses) where the target agent is listening for association connection requests. The agent must be ready to process SCIOP messages on associations established at this port.
- **max_streams** contains the number of individual inbound and outbound data streams within the SCTP association for this port number. The Client ORB may use less than **max_streams** when establishing an association.

- **object_key** is an opaque value supplied by the agent producing the IOR. This value will be used in request messages to identify the object to which the request is directed. An agent that generates an object key value must be able to map the value unambiguously onto the corresponding object when routing requests internally.
- **components** is a sequence of **TaggedComponent**, which contains additional information that may be used in making invocations on the object described by this profile. More research is needed to determine if all **TaggedComponents** that apply to IOP 1.2 are applicable to SCIOP, to support ORB services such as security, and to support other GIOPs, ESIOPs, and proprietary protocols.

3.2 *SCIOP SCTP Association and Stream usage:*

As mentioned above, the Client ORB will establish an SCTP Association to the Server address contained in the SCIOP IOR Profile. Note that this address may be a list of host addresses if the Server's SCTP association is available via multiple physical interfaces, and a single port number. The Client will create this association with the desired number of inbound and outbound streams which is less than or equal to the **max_streams** parameter contained in the IOR. Because SCIOP uses the streams as bi-directional data flows, the number of inbound & outbound streams must be equal. Upon successful establishment of the association, any of the streams may be used for GIOP message traffic. An association should only be closed after sending a GIOP **CloseConnection** on all streams within the association. Failure to do so will constitute an ungraceful teardown of the associated GIOP connections by the connected peer.

SCTP Association Restart: Unlike TCP/IP, if a host connected via an association disappears and then restarts (for example, it has rebooted without the peer's SCTP engine being aware of its departure), then the GIOP connections previously associated with the "old" association must be destroyed. Then a "new" association will be established w/ new GIOP connections created on the peer receiving the restart notification.

