

Network Working Group
Request for Comments: 1538
Category: Informational

W. Behl
McDATA Corporation
B. Sterling
McDATA Corporation
W. Teskey
I/O Concepts
October 1993

Advanced SNA/IP : A Simple SNA Transport Protocol

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This RFC provides information for the Internet community about a method for establishing and maintaining SNA sessions over an IP internet. While the issues discussed may not be directly relevant to the research problems of the Internet, they may be interesting to a number of researchers and implementors. Any questions or comments relative to the contents of this RFC may be sent to the following Internet address: snaip@mcddata.com.

Table of Contents

1. Introduction.....	2
2. Motivation and Rationale.....	2
3. SNA/IP Protocol Specification.....	3
3.1 Glossary.....	3
3.2 Conventions and Assumptions.....	3
3.3 The Protocol.....	3
3.3.1 Connection Establishment.....	3
3.3.2 Data Transfer.....	5
3.3.3 Connection Termination and Loss.....	6
3.3.4 Session Data Flow.....	7
3.3.5 State Transition Table for the Initiating Node.....	8
4. LLC to SNA/IP Conversion.....	8
5. Performance.....	8
6. VTAM Definition.....	9
7. Acknowledgments.....	9
8. References.....	9
9. Security Considerations.....	10
10. Authors' Addresses.....	10
11. Disclaimer.....	10

1. Introduction

Advanced SNA/IP suggests a method for the transmission of SNA session data over an IP network. This memo documents the SNA/IP protocol as implemented in the McDATA LinkMaster(R) 6200 Network Gateway, McDATA LinkMaster(R) 7100 Network Controller, and I/O Concepts X-Direct TN3270 Server.

Advanced SNA/IP differs from other protocols designed to enable routing of SNA session traffic over an IP network. SNA/IP was originally designed for implementation in peripheral network nodes like SNA gateways and downstream nodes (DSNs). It is the authors' view, however, that SNA/IP could also be implemented in intermediate network nodes like routers as the base for an LLC to IP subnet gateway or data link switch function.

2. Motivation and Rationale

The token-ring media access control (MAC) protocol 802.5 and logical link control (LLC) protocol 802.2 were the first set of LAN protocols used to provide a reliable and connection-oriented data link service for SNA sessions in a LAN environment.

McDATA's experience with transporting SNA over 802.5 networks led to an 802.3/802.2 (Ethernet) based variation. As prospective customers were introduced to these Ethernet products, the question of routability arose. Network administrators, accustomed to working with Ethernet networks and the IP-based protocols, required an IP routable solution. McDATA's "SNA over Ethernet" products were bridgeable, but were not routable.

SNA sessions require a reliable and connection-oriented data link. TCP running over IP provides a reliable and connection-oriented transport service and has the added benefit of being routable. It seemed the UDP and TCP protocols could be used in place of 802.2 Type I and Type II levels of service used in traditional SNA token-ring implementations. Advanced SNA/IP was created as a result of these observations.

3. SNA/IP Protocol Specification

3.1. Glossary

Data Link Switching (DLSw) - This is best described as a routing protocol used for the conversion of LLC-based SNA sessions to an IP form. The initial version of the DLSw protocol is documented in the informational RFC 1434 [1].

Downstream Node (DSN) - An SNA Physical Unit (PU) type 2.0 or 2.1 device connected to the SNA network via a LAN (802.5, 802.3, etc.) as opposed to an SDLC, X.25, or channel connection.

SNA Gateway - A device that provides a data link control (DLC) conversion function for SNA PU type 5 (host) devices and LAN-attached DSNs.

Subnet SNA Gateway - A device connected to both a traditional SNA token-ring segment and an IP network that performs local termination of the LLC connections, a mapping function of source address to destination IP address, and a conversion (switching) function of LLC to IP.

3.2. Conventions and Assumptions

Frame formats are shown starting with the IP header. Other headers will, of course, appear in the actual frames sent, but these headers, and the numbers of them, will vary across MAC types.

It is assumed the reader is familiar with both the standard SNA protocol (to the extent it applies to SNA Gateway and DSN functions) and the base set of TCP/IP protocols. Where practical, the reader is asked to refer to appropriate SNA and TCP/IP documentation.

3.3. The Protocol

Conceptually, there are three phases to the Advanced SNA/IP protocol: the Connection Establishment phase, the Data Transfer phase, and the Connection Termination phase.

3.3.1. Connection Establishment

Connection Establishment involves the exchange of logical XID packets between the connecting end nodes and culminates in the establishment of a TCP connection. This process is similar to the IBM-specified Test, XID, SABME and UA exchange used to establish a Type II 802.2 connection for SNA traffic [2]. In place of the 802.2 Type I messages, SNA/IP defines the following set of UDP datagrams:

Logical Null XID

Use: Sent by an initiating node (such as a DSN) when the connection to another SNA node is desired.

The Logical Null XID communicates the sending node's desire to negotiate connection parameters. Once those parameters are established, the Logical Null XID communicates the sender's TCP port to which a connection is to be made.

Format:

```
-----
| IP Header | UDP Header | 0xBF |
-----
```

Source IP address:	The IP address of the initiating node.
Destination IP address:	The IP address of the partner SNA node.
Source UDP Port:	Must match the TCP port number to be used in the eventual TCP connection.
Destination UDP Port:	A known port on the partner node that expects SNA/IP datagrams.

XID Request

Use: Sent in response to a Logical Null XID and requests the receiving node to send a Logical SNA XID datagram.

Format:

```
-----
| IP Header | UDP Header | 0xBF |
-----
```

The source and destination IP and UDP port numbers follow, logically, from those provided in the Logical Null XID datagram.

The format of the XID Request and Logical Null XID are the same. The two types are distinguished by the roles assumed by the two nodes. In current implementations, the DSN initiates the XID exchange by sending the Logical Null XID. The SNA Gateway responds with the XID request.

Logical SNA XID

Use: Sent in response to an XID Request and in the context of SNA XID negotiation.

Format:

```
-----
| IP Header  |  UDP Header  | 0xBF | SNA XID data  |
-----
```

For PU 2.0 nodes, the SNA XID data consists of a Format 0 XID [3].

For PU 2.1 nodes, the SNA XID data consists of a Format 3 XID [3].

A typical Connection Establishment data flow appears below.

Node 1	Node 2
Logical Null XID ----->	
	<----- XID Request
Logical SNA XID ----->	
	<----- TCP SYN
TCP SYN ACK ----->	
	<----- TCP ACK

Note: The source UDP port of the Logical Null XID equals the destination TCP port of the TCP SYN segment.

Retries of the Logical Null XID by the initiating node should occur periodically until an XID Request is received in reply. The frequency of the retries is left up to the implementor. The lower bound on the retry timer should be more than the expected round trip time for a packet on the network.

3.3.2. Data Transfer

There are no special packets defined for the Data Transfer phase. Once the TCP connection is established, SNA Request Units (RUs) may be exchanged between the two end nodes. The SNA session data appears as TCP segment data. The only added SNA/IP requirement is that each SNA message consisting of a Transmission Header (TH), Request/Response Header (RH) and an optional Request/Response Request Unit (RU) be preceded by a two octet length field. Examples of Data

Transfer frames are shown below.

```

-----
| IP Header | TCP Header | SNA Msg 1 len | SNA Msg 1 |
-----

-----
| IP Header | TCP Header | SNA Msg 1 cont'd ->
-----

-----
| SNA Msg 2 len | SNA Msg 2 |
-----

```

The length field is passed in big endian format. 0 is a valid length value.

The format of the SNA Message pieces are as defined by SNA [3].

Reliable and sequential delivery of data is provided by the TCP protocol [5,6].

3.3.3. Connection Termination and Loss

Either SNA node may, at any time, terminate the logical SNA connection by issuing a TCP-level FIN segment. Dictates of the TCP protocol apply to this termination process [5,6].

A connection is also terminated, though not as cleanly, if a TCP Reset segment is sent by either SNA node.

Once a connection is terminated, a new connection may be established by the process outlined in the Connection Establishment section. For reconnections made to the LinkMaster 6200 gateway, the same UDP source port must be used by the initiating node. This implies that the same TCP port is used. This requirement stems from the fact the gateway may not always be aware that a TCP connection has been terminated. This would happen if the DSN became disabled prior to sending a FIN or Reset segment. Under these circumstances, SNA host resources remain allocated and a reconnection from a DSN, which the host believes to already be in session, is not allowed. By requiring the DSN to use the same port when reestablishing a connection, the LinkMaster 6200 is able to recognize when a reset of the host connection is required.

3.3.4. Complete Session Data Flow

```

Node 1                                     Node 2

Logical Null XID ----->
(UDP Datagram)
Logical Null XID ----->
(UDP Datagram)
<----- XID Request
          (UDP Datagram)
Logical SNA XID ----->
(UDP Datagram)
<----- TCP SYN
          (TCP Message)
TCP SYN ACK ----->
(TCP Message)
<----- TCP SYN
          (TCP Message)

***** Connection Established *****

<----- SNA ACTPU
          (TCP Message)
SNA ACTPU Response ----->
(TCP Message)
<----- SNA ACTLU
          (TCP Message)
SNA ACTLU Response ----->
(TCP Message)
.
.
.
<----- TCP FIN
          (TCP Message)
TCP FIN ACK ----->
(TCP Message)
<----- TCP ACK
          (TCP Message)

***** Connection Closed *****

Logical Null XID ----->
(UDP Datagram)
.
.
.
.

```

3.3.5. State Transition Table for the Initiating Node

Given State	Transition State			
	No Conn	Null XID Sent	SNA XID Sent	Conn Estb
No Connection		Internal Act. Stimulus ---> Sends 1st Null XID		
Null XID Sent		Internal Timer Event ----> Resend Null XID	XID Request Received ----> Sends SNA XID	
SNA XID Sent		Internal Timer Event ----> Resend Null XID	SNA XID Received ----> Send SNA XID	Indication that TCP connection is estb.
Connection Established	Indica- tion that TCP conn term.			SNA Session Data

A gateway state transition table is not provided here because the state transitions are dependent on the nature of the SNA host interface (3172 Channel Protocol, 3174 Channel Protocol, SDLC, etc.).

4. LLC to SNA/IP Conversion

The use of Advanced SNA/IP to convert conventional token ring-based SNA traffic to a routable form is both conceivable and practical. While interesting, a discussion of this application falls outside the context of this RFC. Very briefly, it can be said that an SNA/IP-based "subnet SNA gateway" application could do many of the things being discussed in the context of the DLSw specification [1].

5. Performance

The performance of SNA sessions running over an SNA/IP connection will be affected by the bandwidth available on the network and by how much traffic is on the network. SNA/IP is poised to take full advantage of the prioritization and class of service enhancements promised in the next generation of IP. Today, SNA/IP can take

advantage of router packet prioritization schemes based on port number. SNA/IP also leaves intact the standard SNA class of service prioritization protocol.

Performance measures taken at McDATA comparing the throughput of SNA/IP and LLC across a single token-ring segment showed approximately a 15 percent decrease in the maximum transactions per hour (1500 bytes to the DSN, 50 bytes out to the host) for SNA/IP. This decrease is well within the expected levels given the added processing requirements of TCP/IP over LLC in the LinkMaster 6200 and LinkMaster 7100 operating environments.

6. VTAM Definition

The host VTAM definition of SNA/IP downstream nodes is dependent on the gateway implementation. Downstream nodes may appear as switched major nodes connected to an XCA or as downstream nodes connected to a PU 2.0 controller [4].

7. Acknowledgments

The authors wish to acknowledge that the definition of SNA/IP was a collaborative effort involving many individuals ranging from customers to sales and marketing personnel to engineers. Particular thanks go to David Beal, Steve Cartwright, Tracey Floming, Audrey McEwen, Mark Platte, Paul Schroeder, Chuck Weil, and Marty Wright, who all played key roles in the development and testing of this protocol and also in the editing of this RFC.

8. References

- [1] Dixon, R., and D. Kushi, "Data Link Switching: Switch-to-Switch Protocol", RFC 1434, IBM, March 1993.
- [2] "Token-Ring Network Architecture Reference", IBM document #SC30-3374-02.
- [3] "Systems Network Architecture Formats", IBM document #GA27-3136-12.
- [4] "VTAM Resource Definition Reference", IBM document #SC31-6438-1.
- [5] Comer, D., "Internetworking with TCP/IP Volume I", Prentice Hall 1991.
- [6] Postel, J., "Transmission Control Protocol - DARPA Internet Program Protocol Specification", STD 7, RFC 793, USC/Information Sciences Institute, September 1981.

9. Security Considerations

This RFC does not address issues of security. SNA level security procedures and protocols apply when SNA/IP is used as the transport.

10. Authors' Addresses

Wilfred Behl
310 Interlocken Parkway
Broomfield, Colorado 80021

Phone: 303-460-4142
Email: wil@mcddata.com

Barbara Sterling
310 Interlocken Parkway
Broomfield, Colorado 80021

Phone: 303-460-4211
Email: bjs@mcddata.com

William Teskey
2125 112th Ave. North East
Suite 303
Bellevue, WA 98004

Phone: 206-450-0650
Email: wct@ioc-sea.com

Note: Any questions or comments relative to the contents of this RFC should be sent to snaip@mcddata.com. This address will be used to coordinate the handling of responses.

11. Disclaimer

McDATA, the McDATA logo, and LinkMaster are registered trademarks of McDATA Corporation. All other product names and identifications are trademarks of their respective manufacturers, who are not affiliated with McDATA Corporation.

