

Novell IPX Over Various WAN Media (IPXWAN)

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 document describes how Novell IPX operates over various WAN media. Specifically, it describes the common "IPX WAN" protocol Novell uses to exchange necessary router to router information prior to exchanging standard IPX routing information and traffic over WAN datalinks.

Table of Contents

1. Introduction	1
1.1. Operation Over PPP	2
1.2. Operation Over X.25 Switched Virtual Circuits	2
1.3. Operation Over X.25 Permanent Virtual Circuits	2
1.4. Operation Over Frame Relay	3
1.5. Operation Over Other WAN Media	3
2. Glossary Of Terms	3
3. IPX WAN Protocol Description	4
4. Information Exchange Packet Formats	5
4.1. Timer Request Packet	6
4.2. Timer Response Packet	8
4.3. Information Request Packet	10
4.4. Information Response Packet	12
5. References	12
6. Security Considerations	13
7. Author's Address.....	13

1. Introduction

This document describes how Novell IPX operates over various WAN media. It is strongly motivated by a desire for IPX to treat ALL wide area links in the same manner. Sections 3 and 4 describe this common "IPX WAN" protocol.

IPX WAN protocol operation begins immediately after link establishment. While IPX is a connectionless datagram protocol, WANs are often connection-oriented. Different WANs have different methods of link establishment. The subsections of section 1 of this document describe what link establishment means to IPX for different media. They also describe other WAN-media-dependent aspects of IPX operation, such as protocol identification, frame encapsulation, and link tear down.

1.1 Operation Over PPP

IPX uses PPP [1] when operating over point-to-point synchronous and asynchronous networks.

With PPP, link establishment means the IPX NCP [4] reaches the Open state. NetWare IPX will reject all NCP options, and uses normal frame encapsulation as defined by PPP. The IPXWAN protocol MUST NOT occur until the IPX NCP reaches the Open state.

PPP allows either side of a connection to stop forwarding IPX if one end sends an IPXCP or an LCP Terminate-Request. When a router detects this, it will immediately reflect the lost connectivity in its routing information database instead of naturally aging it out.

1.2 Operation over X.25 Switched Virtual Circuits

With X.25, link establishment means successfully opening an X.25 virtual circuit. As specified in RFC-1356, "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode" [2], the protocol identifier 0x800000008137 is used in the X.25 Call User Data field of the Call Request frame, and indicates that the virtual circuit will be devoted to IPX.

Furthermore, each IPX packet is encapsulated directly in X.25 data frame sequences without additional framing.

Either side of the virtual circuit may close it, thereby tearing down the IPX link. When a router detects this, it will immediately reflect the lost connectivity in its routing information database instead of naturally aging it out.

1.3 Operation over X.25 Permanent Virtual Circuits

The nature of X.25 PVC's is that no call request is made. When the router is informed that X.25 Layer 2 is up, the router should assume that link establishment is complete.

Each IPX packet is encapsulated in an X.25 data frame sequence without additional framing. Novell IPX assumes a particular X.25 permanent circuit is devoted to the use of IPX.

If a router receives a layer 2 error condition (e.g., X.25 Restart), it should reflect lost connectivity for the permanent circuits in its routing information database and re-perform the necessary steps to obtain a full IPX connection.

1.4 Operation over Frame Relay

Novell conforms to RFC-1294, "Multiprotocol Interconnect over Frame Relay" [3] for frame relay service and packet encapsulation. Currently, Novell has not stabilized the method for treating frame relay connections - whether they treat the connections as LANs or WANs.

1.5 Operation over other WAN media

Additional WAN media will be added here as specifications are developed.

2. Glossary Of Terms

Primary Network Number:

Every IPX WAN router has a "primary network number". This is an IPX network number unique to the entire internet. This number will be a permanently assigned network number for the router. Those readers familiar with NetWare 3.x servers should realize that this is the "Internal" network number.

Router Name:

Every IPX WAN router must have a "Router Name". This is a symbolic name given to the router. Its purpose is to allow routers to know who they are connected to after link establishment - particularly for network management purposes. A symbolic name conveys more information to an operator than a set of numbers. The symbolic name should be between 1 and 47 characters in length containing the characters 'A' through 'Z', underscore (_), hyphen (-) and "at" sign (@). The string of characters should be followed by a null character (byte of zero) and padded to 48 characters using the null character. Those readers familiar with NetWare 3.x servers should realize that the file server name is the Router Name.

3. IPX WAN Protocol Description

IPX WAN links have the concept of a LINK MASTER and a LINK SLAVE. This relationship is decided upon based on information contained within the first IPX packets transferred across the WAN link.

After link establishment, both sides of the link send "Timer Request" packets and start a timer waiting for a "Timer Response". These "Timer Request" packets are sent every 20 seconds until a response is received or a time-out occurs trying to initialize a connection (the timer is restarted each time a new "Timer Request" is sent). The time-out should be configurable, and is normally about one minute. This is directly dependent on the call setup time for the connection. If a time-out occurs, the router issues a disconnect on the offending connection and optionally attempts to retry the connection.

When a "Timer Request" is received, the router with the lowest primary network number MUST respond with a "Timer Response" packet - and become the "Slave" of the link. If the "Slave" determines that it cannot support any of the Routing Types included in the "Timer Request" packet, the "Slave" should issue a disconnect on the connection being established. The "Master" of the link (determined when a "Timer Response" packet is received) is responsible for defining the network number that is to be used as a common network number for the new WAN link, and for calculating the RIP transport time that will be advertized to other RIP routers for the new link. This is calculated by stopping the timer which was started when a "Timer Request" was initiated and applying the algorithm in section 4.2.

To allow this, both sides of the link MUST have an adequate pool of WAN network numbers (unique within the internetwork) available to be assigned to the link when the call is fully completed. The "Master" of the link MUST then select a network number and construct an "Information Request" packet containing the calculated link delay, the common network number, and its own router name. On receiving this packet, the "Slave" MUST turn the packet around, overwrite the router name and node identifier and send an "Information Response".

After the "Master" has received the "Information Response" and the "Slave" has received the "Information Request", standard IPX RIP and SAP packets are transferred across the WAN link, as currently defined for LAN links. The "IPX Router Specification" [5] contains information describing the Novell RIP/SAP protocol for third party developers.

Note that the "Information Request" and "Information Response" packets are specific to the "Routing Type"=0 information exchanges.

With this routing type, no retransmission is made of any of the Information packets. If a response has not been received within the predefined time-out period, a disconnect is issued on the link, and the link can optionally be attempted later.

If a router detects an error for which no suitable protocol response exists (e.g., unable to allocate a network number), the link should be terminated according to the relevant media specification.

Under certain circumstances, particularly on X.25 permanent circuits, it is only possible to detect the remote router went away when it comes back up again. In this case, one side of the link receives a Timer Request packet when IPX is in a fully connected state. The side receiving the Timer Request MUST realize that a problem occurred, and revert to the IPX link establishment phase. Furthermore, the routing information learned from this connection should be immediately discarded.

4. Information Exchange Packet Formats

All IPX WAN information exchange packets conform to the standard Novell IPX packet format. The packets use the IPX defined packet type 04 defining a Packet Exchange Packet. The socket number 0x9004 is a Novell reserved socket number for exclusive use with IPX WAN information exchange. IPX defines that a network number of 0 is interpreted as being a local network of unknown number that requires no routing. This feature is of use to us in transferring these packets before the common network number is exchanged. Some routers need to know a "Node Number" (or MAC address) for each node on a link. Node numbers will be formed from the "WNode ID" field. The node number will be the 4 bytes of WNode ID followed by 2 bytes of zero.

Router Type number assignment. Other vendors IPX routing protocols can make use of the IPXWAN protocol definition by obtaining Router Types from Novell. This document will then include the new Router Types (with the references to vendor protocol description documents).

WOption Number assignment. These numbers only need to be assigned from Novell for the "Timer Request" and "Timer Response" packets. Other packet types (e.g., the "Information Request" packets, are dependent on the "Router Type" negotiated and can contain any (vendor defined) packet type other than 0 or 1. WOption numbers in these packets are then defined by the vendor defining the Routing Type. The same packet format should still be maintained.

4.1 Timer Request Packet

Checksum	FF FF	Always FFFF
Packet Length	02 40	Max IPX size (576 bytes Hi Lo order)
Trans Control	00	Hops traversed
Packet Type	04	Packet Exchange Packet
Dest Net #	00 00 00 00	Local Network
Dest Node #	FF FF FF FF FF FF	Broadcast
Dest Socket #	90 04	Reserved WAN socket
Source Net #	00 00 00 00	Local Network
Source Node #	00 00 00 00 00 00	Set to zero
Source Socket #	90 04	Reserved WAN socket
WIdentifier	57 41 53 4D	Confidence identifier
WPacket Type	00	Timer Request
WNode ID	xx xx xx xx	Primary Net # of sending router (Hi Lo order)
WSequence #	xx	Sequence start at 0
WNum Options	02	2 Options follow
WOption Number	00	Define Routing Type
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	00 01	Option length (Hi Lo)
WOption Data	00	IPX RIP/SAP Routing
WOption Number	FF	Pad option
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	02 0E	Pad data length (Hi Lo)
WOption Data	00->FF's	Repeated sequence of 00 through FF's.

Note:

Timer Request packets will always be 576 bytes. However, there should be no assumption made about the number of options specified in this packet.

After link establishment, Timer Request packets are sent by both sides of the link. Each end starts their sequence number at zero. Subsequent retries (every 20 seconds) will increment the value of this sequence number. Only a Timer Response packet with a sequence number matching the last sent sequence number will be acted upon.

When receiving this packet, the WNode ID should be compared to the receiver's Primary Network #. If the WNode ID is larger than the receiver's Primary Network #, a Timer Response packet should be sent, and the receiver should become the link "Slave".

Packets received on the reserved socket number not having the WIdentifier set to the hexadecimal values noted above should be discarded.

Routing Type Option:

A routing type of zero (0) is the minimum interoperability requirement (as defined by this document). A router ready to send a Timer Response (and receiving a routing type of zero) MUST respond with a routing type of zero. A router ready to send a Timer Response (and receiving routing types not matching a supported value) SHOULD respond with a Routing Type of zero indicating support for the minimum common protocol.

Note that multiple Routing Type Options can be included in the Timer Request packet if the router supports multiple routing methods for IPX. The included Router Types MUST include and support this type zero option.

Accept Option (for Routing Type and PAD options):

This field MUST be set to YES if the option is supported, and NO if an option is not supported. A Timer Response MUST respond with ONLY one Router Type set to YES.

PAD Option:

This option will normally be the last entry in the packet. Its sole purpose is to fill the IPX packet to be 576 bytes. The pad option data will contain a repeating sequence of zero's through 0xFF's. This should stop compression modems from collapsing the packet and destroying the link delay calculation.

Currently Assigned WOption Numbers (Timer Request Packet):

Routing Type Option = 0x00; Option Length = 0001

Current option data values:

0 Novell RIP/SAP routing with network
 number assigned to the link.

PAD Type Option = 0xFF; Option Length = Variable

Compression Option = 0x80; Option Length = Variable
 (length dependent on compression type)

Current option data values:

Byte 1 Compression type

0 = Telebit compression (length=3) [6]

Telebit Byte 2 - Compression options

Telebit Byte 3 - Number compression slots

4.2. Timer Response Packet

Checksum	FF FF	Always FFFF
Packet Length	02 40	Max IPX size (576 bytes Hi Lo order)
Trans Control	00	Hops traversed
Packet Type	04	Packet Exchange Packet
Dest Net #	00 00 00 00	Local Network
Dest Node #	FF FF FF FF FF FF	Broadcast
Dest Socket #	90 04	Reserved WAN socket
Source Net #	00 00 00 00	Local Network
Source Node #	00 00 00 00 00 00	Set to zero
Source Socket #	90 04	Reserved WAN socket
WIdentifier	57 41 53 4D	Confidence identifier
WPacket Type	01	Timer Response
WNode ID	xx xx xx xx	Primary Net # of sending router (Hi Lo order)
WSequence #	xx	Same as Timer Request received
WNum Options	02	2 Options follow
WOption Number	00	Define Routing Type
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	00 01	Option length (Hi Lo)
WOption Data	00	IPX RIP/SAP Routing (Minimum interoperating requirement). Others may be defined by at a later date by Novell
WOption Number	FF	Pad option
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	02 0E	Pad data length (Hi Lo)
WOption Data	00->FF's	Repeated sequence of 00 through FF's to stop compression modems doing any compression for link delay calc.

The responses contained within this packet are as described in section 4.1. Any unknown options or not supported options from the Timer Request should have the WAccept Option set to NO.

If the Timer Request packet contained more than one Router Type option and the "Slave" supports all the options, the "Slave" should set the WAccept Option to NO on all Router Types except the Routing

Type which is to be adopted. The "Master" of the link will then adopt the routing option specified with the YES setting and complete further information exchanges according to that routing standard.

This packet should contain the same sequence number as the received Timer Request. This packet should ONLY be sent by the router determining themselves to be the "Slave" of the link.

Currently Assigned WOption Numbers (Timer Response Packet):

Routing Type Option = 0x00; Option Length = 0001

Current option data values:

0 Novell RIP/SAP routing with network
number assigned to the link.

PAD Type Option = 0xFF; Option Length = Variable

Compression Option = 0x80; Option Length = Variable
(length dependant on compression type)

Current option data values:

Byte 1 Compression type

0 = Telebit compression (length=3) [6]

Telebit Byte 2 - Compression options

Telebit Byte 3 - Number compression slots

4.3. RIP/SAP Information Request Packet (Router Type=0 Only)

Checksum	FF FF	Always FFFF
Packet Length	00 63	Size of header+data (Hi Lo order)
Trans Control	00	Hops traversed
Packet Type	04	Packet Exchange Packet
Dest Net #	00 00 00 00	Local Network
Dest Node #	FF FF FF FF FF FF	Broadcast
Dest Socket #	90 04	Reserved WAN socket
Source Net #	00 00 00 00	Local Network
Source Node #	00 00 00 00 00 00	Set to zero
Source Socket #	90 04	Reserved WAN socket
WIdentifier	57 41 53 4D	Confidence identifier
WPacket Type	02	Information Request
WNode ID	xx xx xx xx	Primary Net # of sending router (Hi Lo order)
WSequence #	00	Sequence start at 0
WNum Options	01	1 Option to follow
WOption Number	01	Define IPX RIP/SAP info exchange
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	00 36	Option length (Hi Lo)
WOption Data		
Link Delay	xx xx	Hi Lo link delay in milli seconds (see below for calculation)
Common Net #	xx xx xx xx	Hi Lo Common Network #
Router Name	xx (x 48 decimal)	Router name - as defined in section 2.

Calculation of link delay is performed as follows:

```
// Start_time is a time stamp when Timer Request sent out
// End_time is a time stamp when a Timer Response is
// received.
link_delay = end_time - start_time; // 1/18th second
// We are on a slow net, so add some biasing to help stop
// multiple workstation sessions timing out on the link
if (link_delay < 1)
{
    link_delay = 1;
}/*IF*/
link_delay *= 6;    // Add the biasing
link_delay *= 55;   // Convert link delay to milliseconds
```

The "Link Delay" is used as the network transport time when advertized in the IPX RIP packet tuple for the network entry "Common Net #". For a consistent network, a common link delay is required at both ends of the link and is calculated by the link "Master".

The Common Net # is supplied by the link "Master". This number must be unique in the connected internetwork. Each WAN call requires a separate number.

Currently only a single option is defined for the "Information Request" packet for Routing Type=0.

4.4. RIP/SAP Information Response Packet (Router Type=0 Only)

Checksum	FF FF	Always FFFF
Packet Length	00 63	Size of header+data (Hi Lo Order)
Trans Control	00	Hops traversed
Packet Type	04	Packet Exchange Packet
Dest Net #	00 00 00 00	Local Network
Dest Node #	FF FF FF FF FF FF	Broadcast
Dest Socket #	90 04	Reserved WAN socket
Source Net #	00 00 00 00	Local Network
Source Node #	00 00 00 00 00 00	Set to zero
Source Socket #	90 04	Reserved WAN socket
WIdentifier	57 41 53 4D	Confidence identifier
WPacket Type	03	Information Response
WNode ID	xx xx xx xx	Primary Net # of sending router (Hi Lo order)
WSequence #	00	Sequence start at 0
WNum Options	01	1 Option to follow
WOption Number	01	Define IPX RIP/SAP info exchange
WAccept Option	01	0=No,1=Yes,3=Not Applic
WOption Data Len	00 36	Option length (Hi Lo)
WOption Data		
Link Delay	xx xx	Hi Lo link delay (as received in Info Requ)
Common Net #	xx xx xx xx	Hi Lo Common Network # (as received in Info request)
Router Name	xx (x 48 decimal)	Router name - as defned in section 2.

The responses contained within this packet are as described in section 4.3.

5. References

- [1] Simpson, W., "The Point-to-Point Protocol (PPP) for the Transmission of Multi-protocol Datagrams over Point-to-Point Links", RFC 1331, May 1992.
- [2] Malis, A., Robinson, D., and R. Ullman, "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode", RFC 1356, August 1992.

- [3] Bradley, T., Brown, C., and A. Malis, "Multiprotocol Interconnect over Frame Relay", RFC 1294, January 1992.
- [4] Simpson, W., "The PPP Internetwork Packet Exchange Control Protocol (IPXCP) Compromise Version", Work in Progress.
- [5] Novell IPX Router Specification. Novell Part Number 107-000029-001. (Note: Currently, this document is only available as part of a Novell developers program as part of an SDK. Novell Labs, Provo (UT) should be able to provide more information on this document.)
- [6] Lewis, M., Telebit Corp. "IPX Header Compression based on Van Jacobson Header Compression for TCP/IP", Work in Progress, contact: (mlewis@telebit.com).

6. Security Considerations

Security issues are not discussed in this memo.

7. Author's Address

Michael Allen
Novell, Inc.
2180 Fortune Drive
San Jose, CA 95131

EMail: MALLLEN@NOVELL.COM

Chair's Address:

The working group can be contacted via the current chair:

Brian Lloyd
Lloyd & Associates
3420 Sudbury Road
Cameron Park, California 95682

EMail: brian@ray.lloyd.com

Phone: (916) 676-1147