

Network Working Group
Request For Comments: 1551
Obsoletes: RFC 1362
Category: Informational

M. Allen
Novell, Inc.
December 1993

Novell IPX Over Various WAN Media (IPXWAN)

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. 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. This document supercedes RFC 1362 and adds capability for Unnumbered RIP links, On-demand statically routed links and client to router connectivity.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 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 | 3 |
| 1.4 Operation Over Frame Relay | 3 |
| 1.5 Operation Over Other WAN Media | 3 |
| 2. Glossary Of Terms | 4 |
| 3. IPX WAN Protocol Description | 4 |
| 3.1 The Initial Negotiation | 5 |
| 3.2 Information Exchange | 9 |
| 3.3 NAK Packets | 10 |
| 4. Information Exchange Packet Formats | 10 |
| 4.1 Timer Request Packet | 11 |
| 4.2 Timer Response Packet | 14 |
| 4.3 Information Request Packet | 15 |
| 4.4 Information Response Packet | 18 |
| 5. Running Unnumbered RIP | 19 |
| 6. Workstation Connectivity | 19 |
| 7. On-demand, Statically Routed Links | 19 |
| 8. References | 21 |
| 9. Security Considerations | 21 |
| 10. Author's Address..... | 22 |

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.

The 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 negotiate down to a null set of 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. Options negotiated by the IPXWAN protocol MUST supercede any options negotiated by the IPXCP.

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 Permanent Virtual Circuits

To determine when a permanent virtual circuit (PVC) has become active or inactive, the router interacts periodically with either a private Frame Relay switch or a public Frame Relay network. The method used depends on the switch or service provider. Some support [7], section 61 others support [3], Annex D. Novell supports both methods.

When a router is restarted, IPXWAN exchanges over active Frame Relay PVCs (that is, PVCs that have remained active before and after restart) can begin immediately.

Each IPX packet is encapsulated in a Frame Relay frame sequence as defined in [3] without additional framing.

When a router detects that a Frame Relay PVC has transitioned from an inactive to an active state, link establishment is considered complete and IPXWAN exchange over this newly activated link begins.

When an active PVC becomes inactive, the router reflects the lost connectivity in its routing information database.

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.

For workstation (client) connectivity, it is useful if the client connection software is configured with a symbolic name reflecting the name of the client. This allows a router management utility to determine which connection connects with which client/router. If no name is configured, it is recommended that a default string such as "DIAL-IN-CLIENT" is used.

3. IPX WAN Protocol Description

After the underlying data link connection is established as described in the preceding media dependant description, the IPXWAN protocol is activated to exchange identities and determine certain operational characteristics of the link.

There are two steps in the IPXWAN operation:


- Negotiating master/slave role and choice of routing protocol.
The master/slave roles persist for the IPXWAN exchanges only;
- Information exchange of final router configuration.

After these steps are concluded, transmission of IPX routing packets begins - using the routing protocol negotiated - as well as

transmission of IPX data traffic.

3.1 The Initial Negotiation

The first exchange of packets decides the master/slave roles and the routing protocol to be used on the link and gauges the link delay for the routing metrics. The initial negotiation is the same for all protocols.



After link establishment, both sides of the link send Timer Request packets and start a timer waiting for a Timer Response. These Timer Requests are sent every 20 seconds until a response is received or a decision is made that the remote node is not responding. This could be after a predefined time (min. 60 seconds) or a number of retries (e.g., 16).

In composing the Timer Request, the router or workstation takes into consideration:

- Which types of routing protocols it supports;
- Whether it is prepared to assign a network address to the link;
- For workstations, whether they require the ability to specify their network/NIC address on a reconnect;

- Whether it is able to support IPX header compression [6].

For each routing protocol supported, place an option in the Timer Request packet. The Routing Type options should be added in the originator's order of preference with the most preferred option first.

Some of the newer (or modified) routing protocols do not have the requirement to allocate a network number on a WAN link. This type of routing protocol has the advantage of potentially simpler configuration as no network number pools are necessary for WAN links. However, these router implementations may still wish to interoperate with the older IPXWAN implementations which are able to allocate network numbers to the WAN link. In this case, the following method is used to force the older implementation to become the link master. It should be noted that a router implementation capable of supporting workstation dial-in MUST be able to supply AT LEAST ONE network number on which the workstation can reside.

If the router is prepared to assign an IPX network number to the link, it sends its primary network number in the Timer Request WNodeID field, and omits the Extended Node ID option. On the other hand, if the router is NOT prepared to assign an IPX network number to the link, it sets the Timer Request WNodeID field to zero, and includes its primary network number in an Extended Node ID option.

Workstations follow a similar, but slightly different set of rules for setting the WNodeID field. If this is the first time the workstation is connecting to the router, the workstation will set the WNodeID to zero indicating the router should be the link master and allocate a network number for the new link. In this case, the workstation will respond to the router's Timer Request and acknowledge only the Workstation Routing Type option. Note that a workstation does NOT include an Extended Node ID option in its timer request.

If the workstation is reconnecting a link after an earlier inactivity disconnect, it is necessary for the workstation to be able to specify its network, NIC address and "Router Name" field (so that file server connections can be maintained after the reconnect). In this case, the workstation will set its WNodeID field to FFFFFFFFh forcing itself to be the link master. In this case, the router will respond to the workstation's Timer Request with only the Workstation Router Type acknowledged.

Further packets in the IPXWAN exchange MUST use the correct WNodeID (workstations will always use zero).

On receiving a Timer Request packet, a router determines its role - master or slave - for the remainder of the IPXWAN exchanges. The master role does not denote special privileges, it merely means that the router is the requestor in the ensuing request/response exchanges. The decision is made as follows:

- a) If there is an Extended Node ID present (and we understand the option), this must be compared to our Primary Network Number. If we have the lower Primary Network Number, we MUST respond with a Timer Response and become the slave.
- b) If there is NO Extended Node ID, we must compare the WNodeID of the received Timer Request with our Primary Network Number and respond if we have a lower number.

Note: The Primary Network Number for a workstation when determining master/slave roles depends on whether the workstation requires itself to be the master of slave. It should compare the received WNodeID to that sent in its own Timer Request.

The numeric comparisons are done by considering each byte of the WNodeID or Extended Node ID fields as an unsigned integer, and the first byte as most significant.

The link slave responds to the Timer Request with a Timer Response. To do so, each option in the received Timer Request is parsed. If an option is not supported (or recognized), that option is rejected by changing the WAccept field to "NO" for that option.

When selecting the router type which will be used on the link, the first option in the Timer Request which can be supported should be accepted. All other router types should have the WAccept field set to "NO". A router MUST NOT accept workstation connectivity to a node which is another router.

Note: It is permitted for a router to support a numbered routing type, but not be able to assign the network number. In this case, that routing type can be selected only if the other router supports it and is able to assign the network number. This can be determined by the value of the received WNodeID field. If the router is unable to assign a network number to the link, it MUST support Unnumbered RIP and include this option in the Timer Requests.

If a router wishes to provide WAN Client access without supporting other WAN routing types, a potential problem arises since a router and WAN client would both just be sending a single Routing Type option indicating the use of WAN Client. The IPXWAN specification

does not allow a WAN workstation to connect to another WAN workstation. The method for detecting this is that the sent and received Timer Requests have a single Routing Type defined of WAN Client. To overcome this problem, IPXWAN defines that a router MUST NOT send a single Routing Type if that type is just WAN Client. The router MUST additionally include one (or more) of the defined routing types (like WAN RIP) with the WAccept option set to NO. This is so that a workstation may detect that this is actually a router sending the Timer Request and not just another workstation trying to call a workstation. The extra option will serve to be a counted Routing Type that will be ignored. If a workstation detects it is connecting to another workstation, it should disconnect the link.

Note that a router supporting a workstation will need to be able to supply AT LEAST one network number for workstations. All dial-in workstations could share the same network, and be assigned unique node numbers by the router, or each workstation could be assigned a different network number. This is a router specific implementation detail. Use of a single network for all clients is preferred, however, this does involve extra work by the router when dealing with broadcast frames. When the router is the link master and allocating NIC addresses on a single network, it should ALWAYS use a unique value - by incrementing the NIC address for each client connection. This allows a workstation which is reconnecting the ability to specify his old network and NIC address. It is unlikely with a 6 byte NIC address, that there will be wrap-around in the numbers that would cause a problem. Router Node Number allocation should follow a few simple rules. The six byte NIC address SHOULD have the first byte set to 2.

```

Byte #  +-1-----2-----3-----4-----5-----6--+
         | 02 | XX | XX | XX | XX | XX |
         +-----+

```

In an IEEE address space, this would represent a non-multicast, locally defined address. Node numbers of zero or -1 are not allowed.

If a slave determines it cannot support any of the supplied routing protocols in the received Timer Request, it MUST 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 4.3.

3.2 Information Exchange

After exchanging Timer Request packets, the link master and slave have been determined, and the Routing Protocol to be used on the link is negotiated. The link master is now responsible for sending an Information Request packet to the slave specifying the network number to be used on the new link (zero for unnumbered RIP and On Demand), the calculated transport time to be used in the routing metric, the Router Name (for management purposes), and for a workstation connection, the NIC address the workstation will be adopting. The NIC address option is a separate option added in the Information Request/Response for workstation connectivity. It is NOT present for router to router connections.

If a router receives an inappropriate Information Request from a workstation trying to set the common network number and NIC address the router MUST overwrite these values with preferred values. When the workstation receives the Information Response, it MUST note the new values. If the workstation is unable to adjust to the new values, it MUST issue a disconnect on the link. If a workstation is the link master (i.e., it is reconnecting), the router is additionally responsible for ensuring the "Router Name" field matches that of the original connection. If the values differ, the call should be disconnected.

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.

When Unnumbered RIP, On-demand or Workstation options are negotiated, Information Request packets are repeated every 20 seconds until a response is received. For the Numbered RIP links, the Information Request is NOT resent. Instead, the link is disconnected after a suitable delay (min. 60 seconds) - this requirement ensures interoperability with earlier versions of IPXWAN. When Information Requests are repeated, they should continue for a preconfigured time (min. 60 seconds) or a preconfigured number of retries (e.g., 16). Each retry uses an incremented sequence number.

3.3 NAK Packets

The IPXWAN protocol uses a NAK packet to indicate the received IPXWAN packet was not acceptable. A NAK packet is an exact copy of the received packet with the WPacketType field set to NAK. There are two anticipated uses of this packet.

- The received WPacketType is invalid or not recognized;
- A badly formed IPXWAN packet is received.

Returning a NAK packet allows the sender a chance to work out what was wrong. If the sender was unable to determine the problem, the call can then be disconnected.

The value of the NAK WPacketType is FFh.

4. Information Exchange Packet Formats

All IPX WAN protocol exchanges utilize 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 protocol exchange. IPX defines that a network number of zero (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. For a workstation, the node number will be explicitly assigned by the router and notified to the workstation in the Information Request packet.

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). Current Routing Types are:

| | |
|-------|--|
| 00 | Numbered RIP/SAP |
| 01 | NLSP (no RIP/SAP - defined in [8]) |
| 02 | Unnumbered RIP/SAP |
| 03 | On Demand, static routing (no RIP/SAP or NLSP) |
| 04 | Workstation (no RIP/SAP) |
| 05-FF | Currently undefined |

WOption Number assignment. These numbers only need to be assigned from Novell for the "Timer Request" and "Timer Response" packets.

Packet Types also need to be assigned by Novell. However, the options within these packets are dependant on the "Router Type" negotiated. WOption numbers in these packets are then defined by the vendor defining the Routing Type. The same packet format should still be maintained.

Router Type 01 will not be described in this document since it involves knowledge of the NLSP protocol to implement. Please refer to [8] for a complete specification of these NLSP IPXWAN exchanges and the NLSP protocol.

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 | xx | Number of options |
| WOption Number | xx | Option Identifier |
| WAccept Option | xx | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | xx xx | Number of following option bytes (Hi Lo) |
| WOption Data | nn | Option specific data |

Routing Type Option:

One or more of the following router type options should be included in a Timer Request packet. A router should ALWAYS include Routing Type zero (0) if full interoperability is required with an older implementation. The router types MUST be included in the senders order of preference. If a router receives a Timer Response with more than one Router Type having WAccept set to Yes, the link MUST be

disconnected.

| | | |
|------------------|-------|--|
| 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 | 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 | 01 | NLSP |
| 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 | 02 | Unnumbered RIP/SAP |
| 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 | 03 | On-demand, static Rting |
| 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 | 04 | Client - No RIP/SAP except on request |

Extended Node ID Option:

The extended node ID should only be included if the WNodeID field is set to zero AND the node constructing the packet is a router. Note that an older version of IPXWAN will just reject this option and automatically become the link master as the WNodeID is zero.

| | | |
|------------------|-------------|--|
| WOption Number | 04 | Extended Node ID |
| WAccept Option | 01 | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | 00 04 | Pad data length (Hi Lo) |
| WOption Data | xx xx xx xx | Real primary network # of this router (Hi-Lo) |

Header Compression Option:

Although more than one header compression option may be specified in a Timer Request packet, it is important that a MAXIMUM of ONE header compression option is accepted. If an implementation receives a Timer Response with more than one header compression option with the accept option set to Yes, the link MUST be disconnected. [Ref 6] defines the full Telebit Header Compression method.

| | | |
|------------------|-------|-------------------------|
| WOption Number | 80 | Header Compression |
| WAccept Option | 01 | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | 00 03 | Variable - at least 1 |
| WOption Data | 00 | 0 = Telebit Hdr Compr. |
| | xx | Compression Options |
| | xx | Compression Slots |

PAD Option:

The PAD option is used to fill the Timer Request up to the 576 byte limit. This field will be of variable length depending on the number of other options in the packet. 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 should be filled with a selection of totally random numbers to avoid compression modems or PPP data compression from destroying the link delay calculation. Note that this is different from the original RFC 1362 specification. This should not affect implementations. Implementations should not attempt to verify the contents of a PAD option.

| | | |
|------------------|----------------|-------------------------|
| WOption Number | FF | Pad option |
| WAccept Option | 01 | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | xx xx | Pad data length (Hi Lo) |
| WOption Data | Random numbers | (enough to fill packet) |

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.

As mentioned earlier, the WNodeID field may be set to zero for a router if it is unable to provide a network number for the link. If a router ONLY supports the Numbered RIP/SAP option, it MUST be capable of proving a network number for a WAN link.

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

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 | xx | Number of options |
| WOption Number | xx | Option Identifier |
| WAccept Option | xx | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | xx xx | Number of following option bytes (Hi Lo) |
| WOption Data | nn | Option specific data |

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

If the Timer Request packet contained more than one Router Type option and the "Slave" supports all the options, the "Slave" MUST set the WAccept Option to YES on the FIRST Router Type supported and NO to ALL other Router Types. This is the Router Type which is to be

adopted by both ends of the link. Information exchanges will then proceed by the link master based on the accepted Router Type.

This packet must contain the same sequence number as the received Timer Request. This packet should ONLY be sent by the router or workstation determining themselves to be the "Slave" of the link. (Workstations are ALWAYS the link slave).

Routers MUST set the WNodeID to their correct value when responding with the Timer Response. A value of zero must NOT be used.

4.3 Information Request Packet

| | | |
|----------------------------|-------------------|---|
| 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. |

Routers MUST set the WNodeID to their correct value when sending an Information Request. A value of zero must NOT be used.

A workstation should replace the Router Name with the configured name, or a constant descriptor string as described in section 2.

For a Workstation Information Request, an extra option is added which specifies the NIC address for the workstation. In this case, the number of options will be set to two (2).

| | | |
|------------------|-------------------|-------------------------|
| WOption Number | 05 | Define NIC Address |
| WAccept Option | 01 | 0=No,1=Yes,3=Not Applic |
| WOption Data Len | 00 06 | Option length (Hi Lo) |
| WOption Data | 02 xx xx xx xx xx | NIC Address for W/S |

Routers or workstations should not refuse to use a NIC address having a first byte with a value other than 02.

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
if (link_delay < 1)
{
    link_delay = 1;
}/*IF*/
// We are on a slow net, so add some biasing to help stop
// multiple workstation sessions timing out on the link
link_delay *= 6; /* Add the biasing for multiple sessions */
link_delay *= 55; /* Convert link delay to milliseconds */
```

If a higher resolution timer is available, better results may be obtained using the following algorithm:

```
conversion_factor = number of timer units in 1/18th second;
link_delay = ((end_time - start_time) * 6) / conversion_factor;
if (link_delay == 0)
{
    link_delay = 1;
}/*IF*/
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". Link Delay is specified in milli seconds.

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. If the negotiated Router Type was Unnumbered RIP, On-demand, or NLSP, the specified Common Net # will be zero.

This packet should contain a sequence number starting at zero. This packet should ONLY be sent by the router or workstation determining themselves to be the "Slave" of the link.

If extra options are included in this packet, they should be silently discarded. If the information option is missing, the link MUST be disconnected.

4.4 Information Response Packet

| | | |
|----------------------------|-------------------|--|
| 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 | Same as Info Request |
| 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 defined in section 2. |

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

A link slave will additionally respond with the received NIC address option as a confirmation of receipt. A workstation should replace the Router Name with the configured name, or a constant descriptor string as described in section 2. If the Information Request contained an inappropriate Common Net # or NIC address, the Information Response may set new values. The receiver of the Information Response is responsible for checking on the value and terminating the connection if the new values cannot be used.

Routers MUST set the WNodeID to their correct value when sending an Information Response. A value of zero must NOT be used.

5. Running Unnumbered RIP

Unnumbered RIP refers to the case where two WAN routers are communicating using the RIP protocol across a link with NO physical IPX network address. The premise for this ability is that there is no need to address a packet to anything on that WAN link. RIP and SAP run in exactly the same way as before, except the source and destination network numbers should be set to zero.

The advantage to running unnumbered RIP links is that it is not necessary to allocate/configure a pool of usable IPX network numbers which can be used on the WAN links. The other advantage is that when there is a large number of WAN links, it is not necessary to flood the network with an unnecessary set of extra RIP information.

6. Workstation Connectivity

Workstations MUST reside on a network and have a unique NIC address on that network to be individually addressable. However, workstations do not need to periodically receive RIP and SAP broadcasts as they play no part in the routing process. This allows routers to reduce background traffic on the workstation link by not sending any periodic RIP and SAP data. Note that it will not cause a problem if the RIP and SAP is sent. It will just slow down the workstation access times.

RIP and SAP information should ONLY be sent if the workstation makes a specific request for information - like a service or route request.

It should also be noted that if multiple workstations are attached to a single WAN workstation network (per router), broadcasts on that network - whether originated from a workstation or the router - MUST reach ALL other workstations. This will involve the router duplicating the packet to all WAN workstation connections.

7. On-demand, Statically Routed Links

On-demand, Static Routing serves two purposes. The "on-demand" part means that a router initiates communication to a destination only when there is data to be forwarded to that destination. "Initiating communication" includes making a datalink call (where necessary) and performing the IPXWAN exchange. A transient connection is closed after a period of inactivity.

The "static routing" part means that no routing information is sent over the link - no RIP, no SAP, and no NLSP. Instead, the router at each end is configured with the routes and services accessible through the link.

With on-demand, static routing, the called router must be able to identify the calling router so that statically configured routes and services can be attached to that connection. For example, with X.25 switched virtual circuits, the calling DTE address can be used; with PPP, the PPP authentication can be used; after IPXWAN has completed, the "Router Name" can be used; with a persistent datalink connection, the physical port identifier or a permanent virtual circuit identifier can be used. The choice of identifier is an implementation decision. Whatever value the called router uses is called a Remote System Identifier, or RSI. For PPP links, Novell uses PPP PAP or CHAP authentication to determine the caller.

A router implementing on-demand, static routing must maintain a database of RSIs, and lists describing the network numbers and services reachable through each RSI. These lists determine the reachability information it transmits to other routers in a routing area. Other routers treat each on-demand, static routing link as though it were permanently available.

The on-demand exchange has a slight variation on the IPXWAN protocol. The differences are as follows.

In the Timer Request, the calling router offers only the "On-demand, static routing" Routing Type. If the called router is capable of On-demand static routing, it offers "On-demand, static routing" in the Timer Request, along with any additional routing types it is willing to support on the link. The Master/Slave election and choice of routing type proceeds as described previously. If the Slave detects a mismatch in routing types, it disconnects the link.

For a persistent datalink (like X.25 PVCs), there may be no descenderable "link establishment" event. For such media, arrival of a Timer Request plays the role of detecting link establishment.

As with Unnumbered RIP, there is no network number assigned to the link. NLSP Packets are not sent on the link. Moreover, periodic RIP and SAP packets are not sent on the link. However, a router must respond to RIP and SAP queries received on the link.

8. References

- [1] Simpson, W., Editor, "The Point-to-Point Protocol (PPP) for the Transmission of Multi-protocol Datagrams over Point-to-Point Links", RFC 1548, Daydreamer, December 1993.
- [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 1490, Wellfleet Communications, Inc., Ascom Timeplex, Inc., July 1993.
- [4] Simpson, W., "The PPP Internetwork Packet Exchange Control Protocol (IPXCP)", RFC 1552, Daydreamer, December 1993.
- [5] Novell IPX Router Specification.
Novell Part Number 107-000029-001. This document may be retrieved via Anonymous FTP to SJF-LWP (130.57.11.140) under /sys/ftpguest/nw_shell/ipxdocs/ipxrout.zip.
- [6] Mathur, S., and M. Lewis, M., "Compressing IPX Headers Over WAN Media (CIPX)", RFC 1553, Telebit Corporation, December 1993.
- [7] ANSI, "Integrated Services Digital Network (ISDN) - Digital Subscriber Signalling System Number 1 (DSS1) - Signalling Specification for Frame Relay", ANSI T1.617-1991, June 1991
- [8] Novell NetWare Link Services Protocol (NLSP) Specification.
Novell part number 100-001708-002. Information on this document may be obtained by sending e-mail to dsink@novell.com.

9. Security Considerations

Security issues are not discussed in this memo.

10. Author's Address

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

EMail: mallen@novell.com

The working group can be contacted via the current chair:

Fred Baker
Advanced Computer Communications
315 Bollay Drive
Santa Barbara, California, 93111

EMail: fbaker@acc.com

