NASA/WVU Software IV & V Facility
Software Research Laboratory
Technical Report Series

NASA-IVV-95-003
WVU-SRL-95-003
WVU-SCS-TR-95-23
CERC-TR-RN-95-009

Reliable Multicast Protocol Specifications Packet Formats

by John R. Callahan, Todd Montgomery, and Brian Whetten

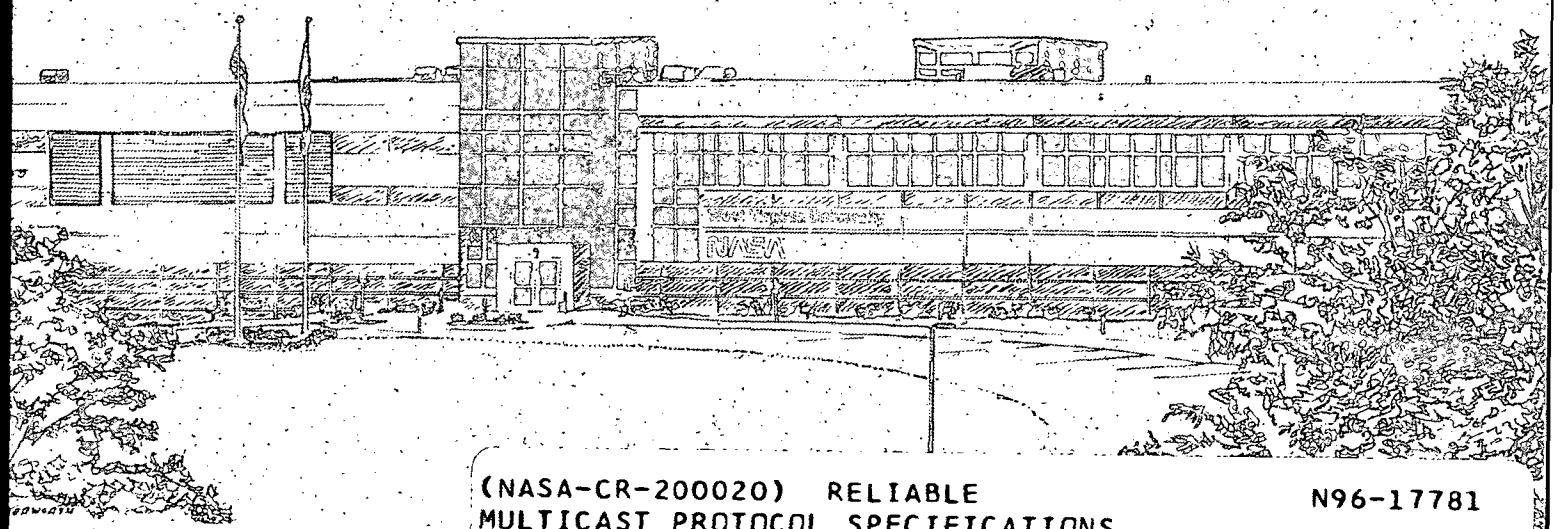

National Aeronautics and Space Administration

West Virginia University

Todd Montgomery
Brian Whetten
John R. Callahan
5 October 1995

The Reliable Multicast Protocol Specification
Protocol Packet Formats

RMP Library Version: 1.3b (1.3 Beta)

This appendix contains the packet formats for the RMP packet types. Each field of the packet is shown and briefly discussed. The packet formats assume a 32 bit IP Addressing scheme.

RMP Fixed Header

The RMP Fixed Header is included at the beginning of all RMP packets. It is the secondary means that packets bound for different token rings are filtered. The first being the IP Multicast Address and port of the token ring. Following the fixed header is one of the additional headers shown below or one or more options then the additional header.

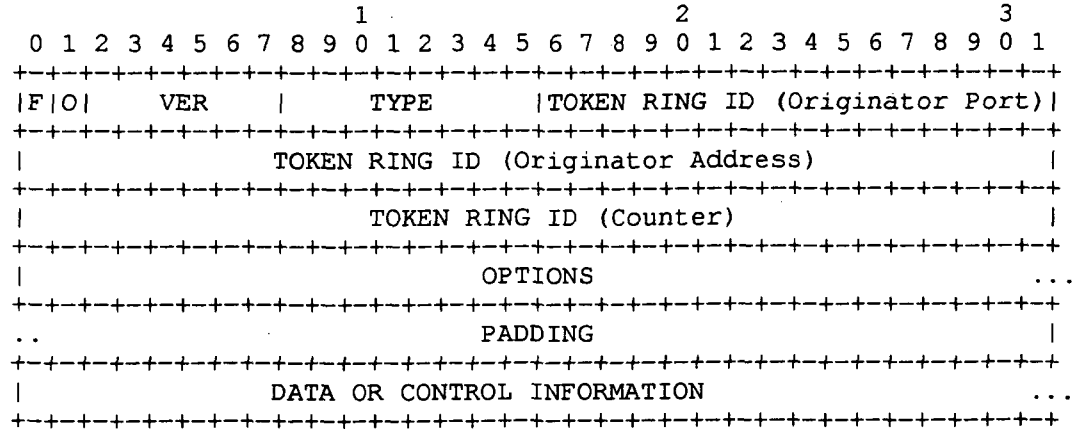

F(FWD): Specifies whether the packets should be multicast to the token ring IP Multicast address and port or not. This field is used to unicast a packet from a non-multicast capable member to a multicast capable member that may forward the packet to the token ring.

O(Options): Specifies if any options are present in this packet. If no options are present then this field is set to 0. If options do exist then this field is set to 1.

VER: Specifies the version of the protocol. The current version is 3. Any packet received with an invalid version should be rejected.

TYPE: Specifies the type of RMP packet. The valid types are:

Value	Packet Type	Abbreviation
0	Reserved	
1	Data	Data
2	ACK	ACK
3	Confirm Token Pass	Conf
4	NACK	NACK
5	New List	NL
6	List Change Request	LCR
7	Recovery Start	RecStart
8	Recovery Vote	RecVote
9	Recovery ACK New List	RecACKNL
10	Recovery Abort	RecAbort
11	Non Member Data	NMD
12	Non Member ACK	NMA
13	Ping	Ping
14	Ping Response	PingResp
15-255	Reserved	

TOKEN RING ID: Specifies the current Token List ID of the token ring.

OPTIONS: Specifies the location of any options present for the packet.

RMP Options

RMP Options take a general form for their first 32 bits. That form is:

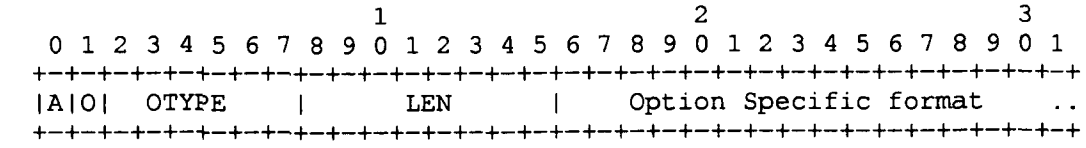

A(Action): Specifies what is to be done if the implementation does not support this option. The defined actions are: Field is 0, then skip over the option and continue processing packet, Field is 1, then discard the packet. Defaults are given for each option. However,

implementations should allow this to be customized.

O(Options): Specifies if more options exist after this option. If this field is set to 0, then this is the last option and the information after the option is the RMP additional header. If this field is set to 1, then another option follows this one.

LEN: Specifies the length of this option in 32-bit words. A value of 0 is invalid.

OTYPE: Specifies the type of option this is. The predefined values are:

OTYPE	Option
0	Null Option
1	32-bit Checksum of Packet
2	Authentcation Header
3	Security Header
4-63	Reserved for Future Use

Option Specific Format: The remaining LEN-1 words of the packet contain information specific to that option.

Defined RMP Options

Null Option

This option has the following format:

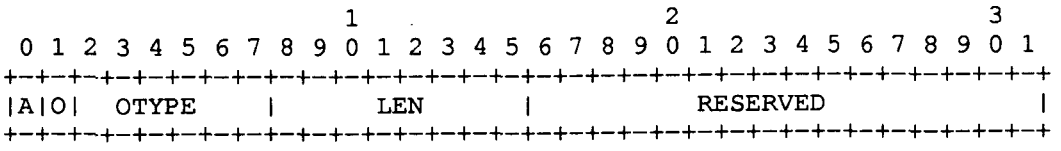

A(Action): Always set to 0 (skip).

OTYPE: Set to 0.

LEN: Set to 1.

RESERVED: Zeroed when sent, ignored when received.

32-bit Checksum Option

This option provides a means of checking for transmission error if RMP is run at lower IP layers. In most cases, this option need not be

used if the transport level already supports error detection.

This option is the last option to be filled when a packet is sent. In addition, this option is the second field checked when receiving a packet. The first field to be checked is the VER field of the Fixed Header.

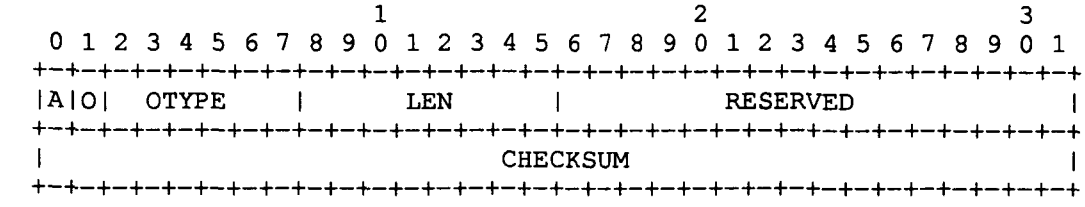

A(Action): Default is 0 (skip) if running over UDP, 1 (discard) if running on raw IP.

OTYPE: Set to 1.

LEN: Set to 2.

RESERVED: Zeroed when sent, ignored when received.

CHECKSUM: 32-bit checksum of entire RMP packet.

Authentication Header Option

RMP does not have special operations designed to support security considerations. These issues are best dealt with at higher level or lower level than RMP. However, RMP does provide mechanisms so that these issues can be effectively dealt with. This option contains information used to provide authentication, integrity, and if used with certain signature algorithms, non-repudication. This option uses the IPv6 concept of Security Associations, and allows RMP to use those values as it would normally use them in IPv6 applications.

This option and the Security Header option are handled as follows: The option is filled in prior to the packet being sent and, on reception, the option is handled after the check of the Token Ring ID of the Fixed Header and prior to checking the RMP additional header. The preferred handling of both of these options is to allow the application to do the actual authentication and encryption/decryption operations and provide, to the RMP implementation, information pertaining to what to do with the packet (such as drop it or send it).

The Authentication Header option format is given below.

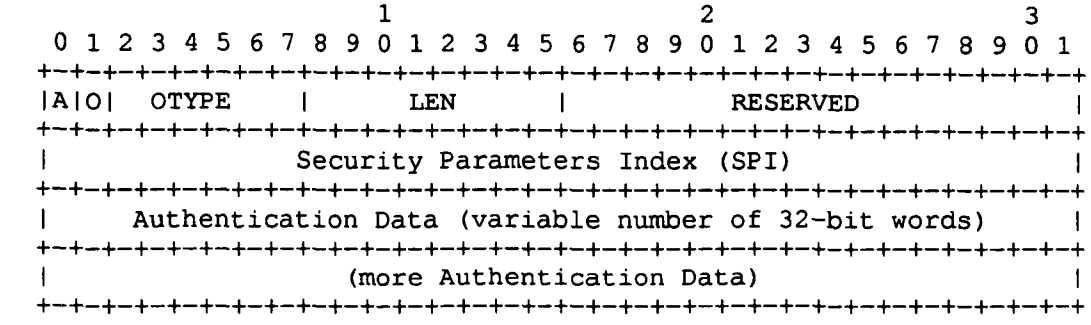

A(Action): Default is 0 (skip).

OTYPE: Set to 2.

LEN: Specifies the length (in 32-bit-words) of this option including the Authentication Data.

RESERVED: Zeroed when sent, ignored when received.

SPI: Specifies the Security Parameters Index value that indicates along with the group and Token Ring ID what kind of algorithm and key(s) should be used for this packet.

Authentication Data: Specifies the actual message digest data generated by authenticating all fields of the RMP packet. For MD5 this field is 128-bits in length (or specifying a LEN of 6). For the Secure Hash Signature (SHS), this field is 160-bits in length (or specifying a LEN of 7).

Security Header Option

Everything after this option is encrypted in a manner dicated by the Security Assocation of this option.

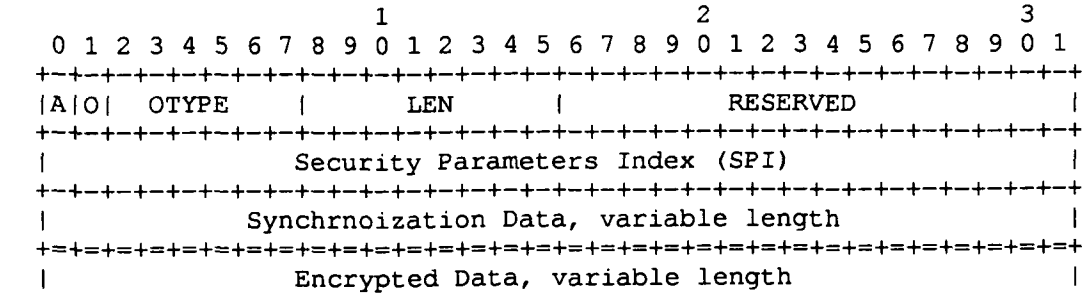

A(Action): Default is 1 (discard).

OTYPE: Set to 3.

LEN: Specifies the length of this option in 32-bit words.

RESERVED: Zeroed when sent, ignored when received.

SPI: Specifies the Security Parameters Index value that indicates along with the group and Token Ring ID what kind of algorithm and key(s) should be used for this packet.

Synchronization Data: Specifies any initialization needed for the decryption algorithm on the remaining data of this packet.

Encrypted Data: All data after this option is encrypted.

RMP Additional Headers

RMP Data Header

The Data Packet type holds data sent from one member of the token ring. Following the fixed header is an additional header, the RMP Data Header. The format for this additional header is shown below. Following the data header is the actual message data.

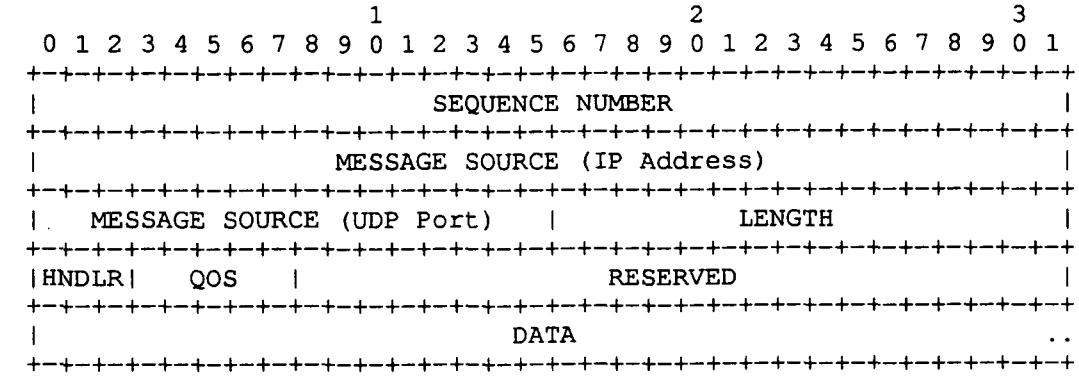

HNDLR(Handler): The handler lock number, if any, for the data packet. The valid values are:

Handler Value	Required Handler
0	None
1-6	The process, if any, that holds the handler with the same number
7	The process, if any, that holds the highest priority handler. 1 is the highest priority and 6 is the lowest

QOS: Specifies the desired QoS of the data packet. The semantics of the delivery of the packet are discussed on age{chap3:qos}. The valid values for the QoS field are:

QoS Value	QoS Desired
0	Reserved
1	Unreliable
2	Unordered
3	Source Ordered
4	Causally Ordered (Optional)
5	Totally Ordered
6-29	K Resilient, K set to (QoS-5)
30	Majority Resilient
31	Totally Resilient

SEQUENCE NUMBER: Each Data packet source stamps each Data packet it sends, except Data packets with QoS of Unreliable, with a sequence number. Each source also keeps a sequence number counter for each token ring it may be a member of. Because sequence numbers are 32 bits in length, they have a valid range of 0 to 2^32-1. Because this range is finite, all arithmetic and comparison with these numbers must be modulo 2^32. RMP requires that no more than 2^31 packets can be created over a period equal to the maximum lifetime for a datagram packet in the network. This condition holds true when IP is used as the datagram service.

MESSAGE SOURCE: Specifies the RMP Process ID of the message source.

LENGTH: Specifies the size of the data field in octets.

RESERVED: Unused field, zeroed when sent, and ignored when received.

DATA: The data to be delivered.

Control Packets

Control Packets are packets that contain information vital to RMP normal operation and maintenance of membership views. The packet types that are classified as Control Packets are ACK Packets, Confirm Packets, NACK Packets, New List Packets, and List Change Request Packets. Any retransmissions of these packet types due to a NACK for that packet do not require modification of the packet header to show the current state. The original packet in its entirety is retransmitted. The control information for each packet follows directly after the fixed header for the packet.

ACK Packet

ACK Packets transfer the token to a new token site and impose ordering on one or more Data Packets and/or Non-Member Data Packets.

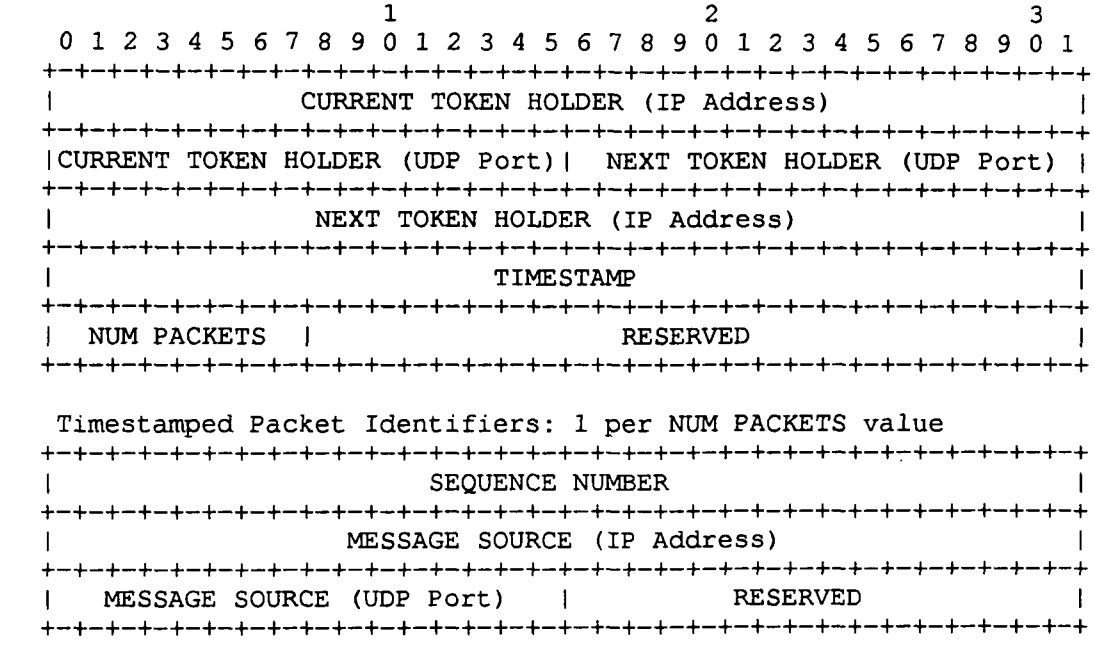

Timestamped Packet Identifiers: 1 per NUM PACKETS value

CURRENT TOKEN HOLDER: Specifies the RMP Process ID of the member sending the ACK packet.

TIMESTAMP: Specifies the timestamp of the ACK packet. Timestamps have a range of 0 to 2^32 and therefore all arithmetic and comparisons on them must be modulo 2^32.

NEXT TOKEN HOLDER: Specifies the RMP Process ID for the member that will become the next token site.

RESERVED: Unused field, zeroed when sent, and ignored when received.

NUM PACKETS: Specifies how many Data packets and Non Member Data packets are timestamped and acknowledged by this ACK packet.

Timestamped Packet Identifiers: There are NUM PACKETS sets of these identifiers that follow after the ACK header. Each identifier contains the same information as the Data packet of Non-Member Data packet it timestamps. The implied timestamps of the Data and Non member Data packets follow monotonically from the ACK timestamp field and are in the order shown in the ACK packet.

Confirm Token Pass Packet

The Confirm Token Pass Packet provides a positive acknowledgment for a token site passing the token.

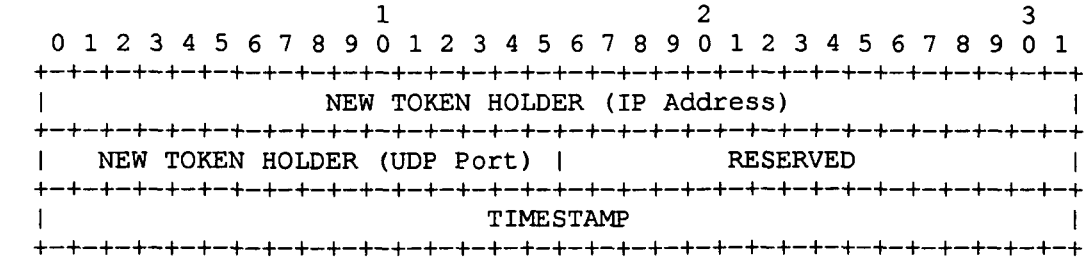

NEW TOKEN HOLDER: Specifies the RMP Process ID of the member that accepts the token and is sending the Confirm Token Pass packet.

RESERVED: Unused field, zeroed when sent, and ignored when received.

TIMESTAMP: Specifies the timestamp of the ACK packet or New List packet that passed the token to the member.

NACK Packet

The NACK Packet is used to request retransmission of lost packets. A NACK can be for a specific timestamp or it may be for a particular Data, Non-Member Data packet, or a List Change Request using the Data Source and Data Sequence Number. The RMP algorithms use timestamp NACKs, however, implementations may alternatively attempt to maximize NACK utility by using Source NACKs as well.

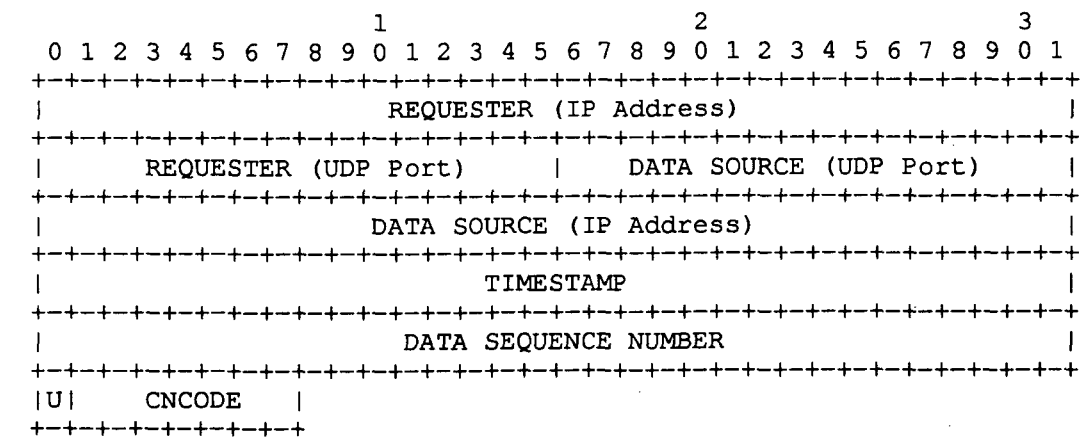

REQUESTER: Specifies the RMP Process ID for the member requesting the missing packets.

CNCODE(Congestion Code): Specifies the reason for the NACK. The valid values are:

Value	Semantic
0	Reserved
1	Buffer Overrun
2	Probable network congestion
3-128	Reserved

TIMESTAMP: Specifies timestamp of the missing packet requested.

U(Unicast Response): Specifies that the response to this NACK is to be unicast to requester in addition to being multicast to group.

DATA SOURCE: Specifies the RMP Process ID of the source of the requested Data, Non-Member Data, or List Change Request packet.

DATA SEQUENCE NUMBER: Specifies the sequence number of the requested Data, Non-Member Data, or List Change Request packet.

New List Packet

New List Packets contain a new membership view for the group and passes the token. The format of the packet is shown below.

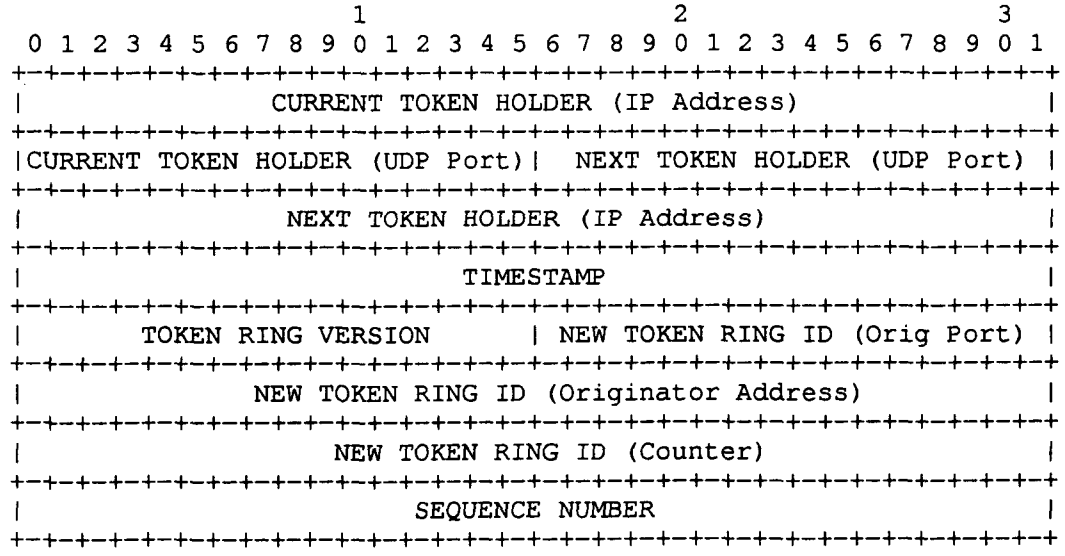

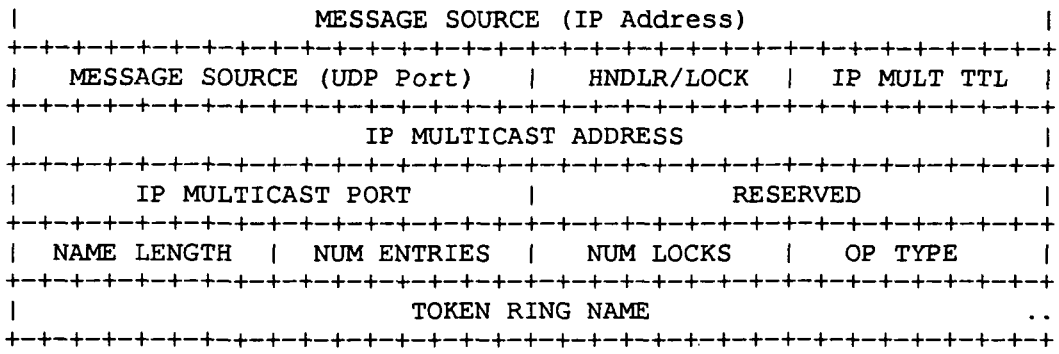

Token List Entries: 1 per NUMBER OF ENTRIES value

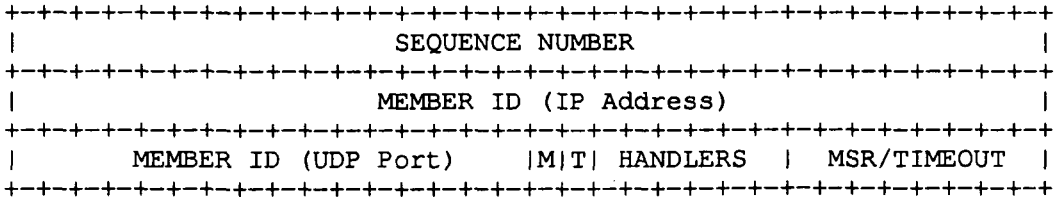

Extra Locks: 1 per NUMBER OF EXTRA LOCKS value

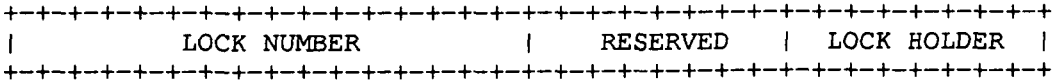

CURRENT TOKEN HOLDER: Specifies the RMP Process ID of the member sending the New List packet and passing the token.

TIMESTAMP: Specifies the timestamp of the New List packet.

NEXT TOKEN HOLDER: Specifies the RMP Process ID of the member that the token is being passed to.

NEW TOKEN RING ID: Specifies the new Token List ID for the token ring. This ID will be used as the TOKEN RING ID for each packet following this New List packet. The NEW TOKEN RING ID has a Originator IP Address and Originator Port that is equal to the RMP Process ID of the member that generates the New List packet, the CURRENT TOKEN HOLDER. The Counter is the value of the old TOKEN RING ID Counter + 1. This Counter has a range of 0 to 2^32-1. Thus all arithmetic and comparison operations on these values must be done modulo 2^32.

NUMBER OF ENTRIES: Specifies the number of Token List members, or entries.

NUMBER OF EXTRA LOCKS: Specifies the number of locks being held by members of the group. This does not count handler locks.

MESSAGE SOURCE: Specifies the RMP Process ID for the source of the List Change Request packet that this New List packet corresponds to.

SEQUENCE NUMBER: Specifies the sequence number of the List Change Request packet that this New List packet corresponds to.

TOKEN LIST VERSION: Specifies the version of the new token list.

IP MULTICAST ADDRESS: Specifies the IP Multicast address used by the token ring. This is used to notify members who join through a unicast List Change Request packet sent to a well known member of the token ring what the IP Multicast address of the ring is.

NAME LENGTH: Specifies the Token Ring Name length in bytes.

IP MULT TTL: Specifies the IP Multicast TTL value for the token ring.

IP MULTICAST UDP PORT: Specifies the IP Multicast UDP Port used by the token ring.

HNDLR/LOCK: Specifies the handler value that the List Change Request packet corresponding the the New List packet was to be performed on. This field is only used in operation types 3-6.

OP TYPE: Specifies the type of operation that was performed on the token list. This field, in effect, details why the New List packet was generated. It also may report errors. The valid operation types are:

Value	Operation Type Description
0	Reserved
1	Requesting Member added to Token List
2	Requesting Member removed from Token List
3	Requesting Member received a Handler Lock
4	Requesting Member released a Handler Lock
5	Requesting Member was denied a Handler Lock
6	Requesting Member attempted to release a Handler Lock it did not hold
7	Reformation occurred and was successful
8	Reformation occurred with possible atomicity violations
9	Failed Reformation due to partition criteria violations (An Invalid List was created)
10	Change TRID Request granted
11	Change TRID Request denied

12	Change TTL Request granted
13	Change TTL Request denied
14	Requesting Member Add Rejected
15	Requesting Member changed Minimum Size Requirement
16	Requesting Member added to Token List by Proxy
17	Requesting Member removed from Token List by Proxy
18-255	Reserved

TOKEN RING NAME: Specifies the null-terminated ASCII name for this token ring. This name must end on a word boundary, which may necessitate that 1-3 extra octets of padding be included after the trailing zero of the name. These octets must be set to zero.

Token List Entries:

Each entry contains information on each member of the token list. There are NUMBER OF ENTRIES of these entries. When this list is committed to an application, the RMP Process examines the list and caches these values for later use in generating ACKs, examining sequence numbers, etc. When the process sees an entry corresponding to itself, it then marks the entry following it as the site it will pass the token to when necessary. If the process is the last entry in the list, then the process marks the first entry in the list as the member it will pass the token to.

M(Multicast Capable): Specifies whether the member is multicast capable or not. Each member of the token list that is not multicast capable requires that each other member in the list unicast each packet that it sends to the list to these members as well as sending the packet to the IP Multicast group. Non-Members of the token ring have this field set to zero.

T(Token Ring Member): Specifies whether the RMP Process is a member of the token ring or not. Non-Members periodically get flushed from the list. The inclusion of Non-Members into the token list is used to bring new token list members up to date with respect to recent Multi-RPC activities.

HANDLERS: Specifies the Handler mask for the member. Each bit position represents a handler lock value. Non-Members of the token ring have this field set to zero.

SEQUENCE NUMBER: Specifies the next sequence number from the RMP Process that is to be delivered. This is used for Non-Members as well.

MEMBER ID: Specifies the RMP Process ID for the token list entry. This is used for Non-Members as well.

MSR/TIMEOUT: Specifies either: (1) the members vote for the minimum partition size that is allowed to form after a failure, or (2) the time in seconds until the non-member times out and is removed from the list.

Extra Locks:

The token ring has 255 locks that members may request and release. Each one is assured to be mutually exclusive, only one member may be in possession of it at one time. The first six locks (1-6) are handler locks. Locks 7-255 are extra locks. These locks have semantics that are totally dependent on what the application desires to use them for. Each lock that is being used is represented by a LOCK NUMBER, LOCK HOLDER tuple in the New List Packet.

LOCK NUMBER: Specifies the lock number. Valid range is 7-255.

RESERVED: Unused, zeroed when sent, and ignored when received.

LOCK HOLDER: Specifies the index in the token list of the member that holds the lock. The first member in the list is denoted 1.

List Change Request Packet

Changes to the membership view are requested by a List Change Request Packet. The format for the packet is shown below.

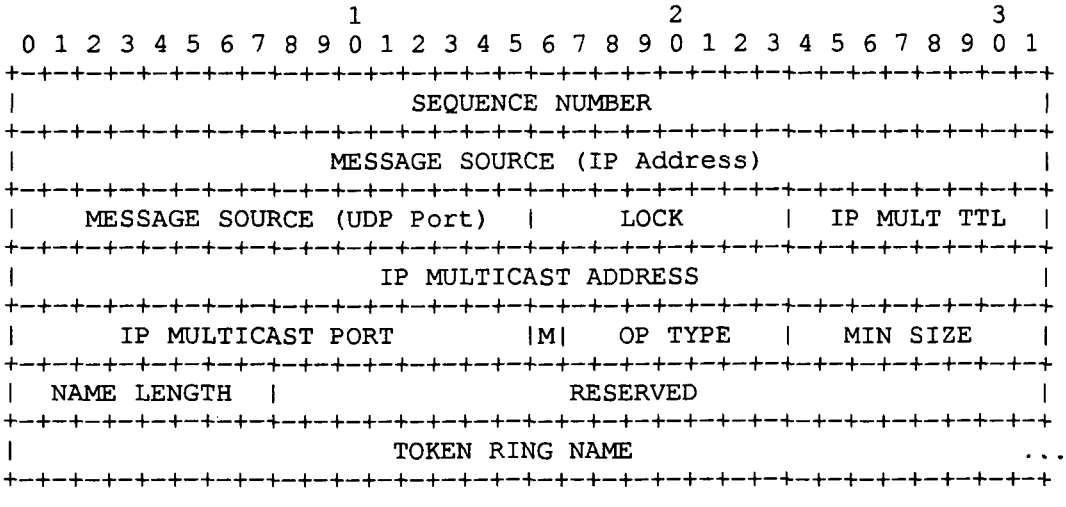

IP MULTICAST ADDRESS: Specifies the IP Multicast address used by the token ring.

IP MULT TTL: Specifies the IP Multicast TTL used by the token ring.

IP MULTICAST UDP PORT: Specifies the IP Multicast UDP Port used by the token ring.

M(Multicast Capable): Specifies whether the process in the List Change Request is multicast capable or not.

LOCK: Specifies the lock value associated with the List Change Request packet. This is used only for OP TYPE of values 3-4.

OP TYPE(Operation Type): Specifies the operation type desired by the List Change Request packet. The valid values for this field are:

Value	Operation Type Description
0	Reserved
1	Request Add Member to Token List
2	Request Remove Member from Token List
3	Request Handler Lock
4	Release Handler Lock
5	Request Change of TRID
6	Request Change of group TTL
7	Request Change of Member Minimum Size Requirement
8	Request Add Member to Token List by Proxy
9	Request Remove Member from Token List by Proxy
10	Request New List
11-15	Reserved

MESSAGE SOURCE: Specifies the RMP Process ID for the process sending the List Change Request packet.

SEQUENCE NUMBER: Specifies the sequence number of the List Change Request packet.

NAME LENGTH: Specifies the length of the TOKEN RING NAME field. This is done exactly as it is done in New List packets.

MIN SIZE: The RMP process' vote for the minimum partition size of the token ring.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Failure Recovery Packets

Failure Recovery Packets contain information relevant to the Reformation Extension of RMP operation. Each Failure Recovery packet type contains a fixed header. The Token List ID for the fixed header is

the last known Token List ID before the failure was detected. This ID does not change until the New List is committed after reformation has finished.

Recovery Start Packet

The Recovery Start Packet is sent by the failure detecting site to the members of the token ring.

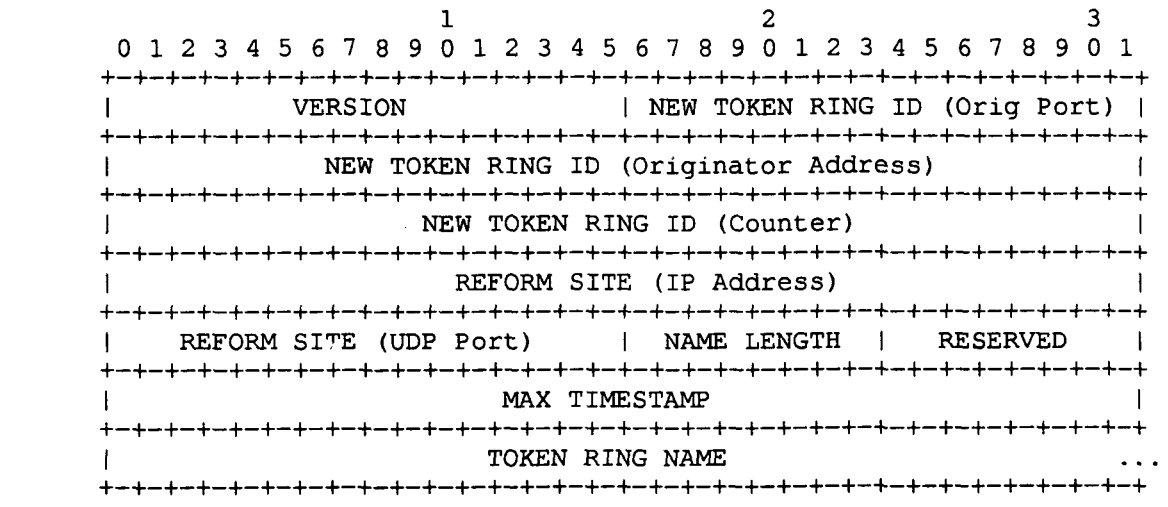

VERSION: Specifies the token ring version of this current reformation. The valid range of VERSION is 0 to 2^16-1. And all arithmetic operations and comparisons must be modulo 2^16.

NEW TOKEN RING ID: Specifies the Token List ID if the reformation succeeds.

REFORM SITE: Specifies the RMP Process ID for the site that is initiating and controlling the reformation, the reform site.

MAX TIMESTAMP: Specifies the SynchTSP for the reformation.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Recovery Vote Packet

Upon receiving a Recovery Start Packet from a member of the token ring, the other sites in the token ring send Recovery Vote Packets to the reform site.

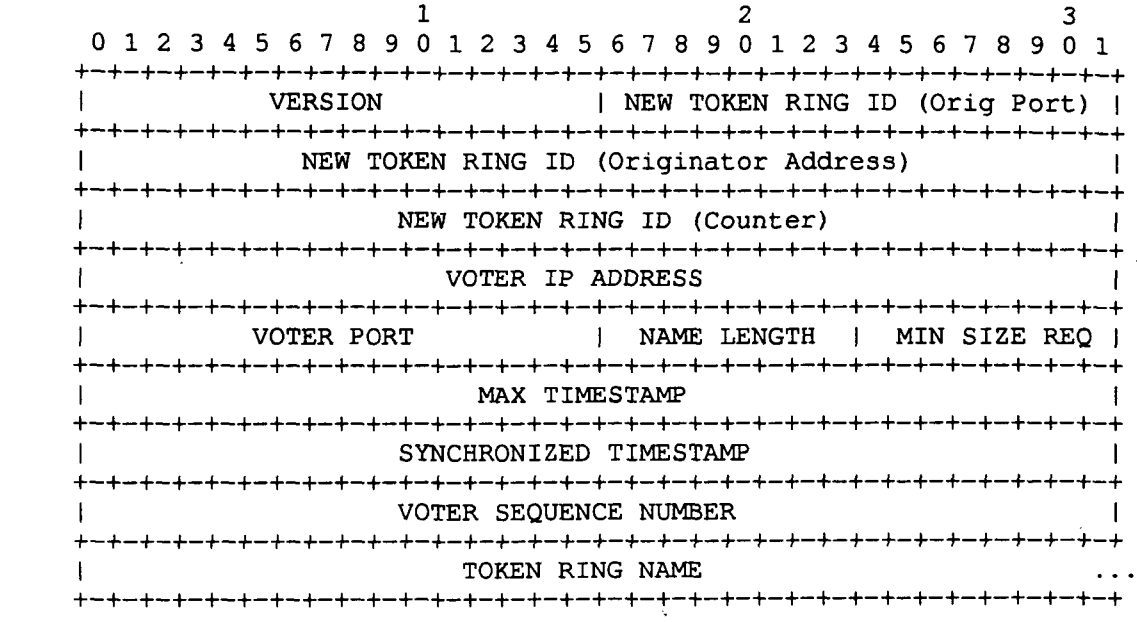

VERSION: Specifies the token ring version of the current reformation. This has the same semantics of the Recovery Vote packet field.

TOKEN RING ID: Specifies the Token List ID for the ring if the reformation succeeds.

VOTER IP ADDRESS: Specifies the IP Address of the site sending the vote.

VOTER PORT: Specifies the UDP Port of the site sending the vote.

VOTER SEQUENCE NUMBER: Specifies the last acknowledged sequence number of the voting site.

MAX TIMESTAMP: Specifies the SynchTSP for the current reformation. If the process sending the Recovery Vote packet has a higher timestamped packet than the Recover Start packet field, then the process sends this new value.

SYNCHRONIZED TIMESTAMP: Specifies up to what timestamp the process sending the Recovery Vote packet is synchronized up to.

VOTER SEQUENCE NUMBER: Specifies the next delivered sequence number for the voter after Reformation completes.

MIN SIZE REQ: Specifies the voters minimum size requirement for this

reformation.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Recovery ACK New List Packet

The Recovery ACK New List Packet is sent to the reform site to signify that the sending site received a New List Packet.

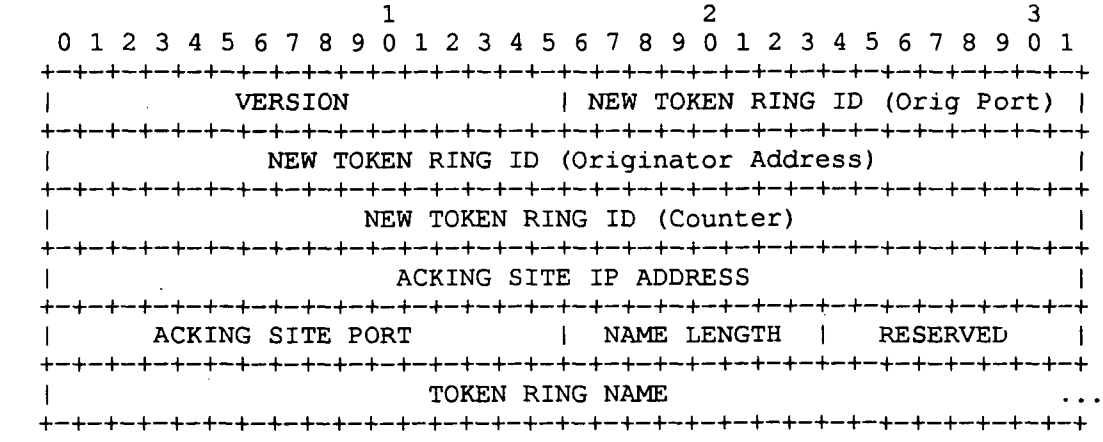

VERSION: Specifies the token ring version for the current reformation. This has the same semantics as the Recovery Start packet field.

NEW TOKEN RING ID: Specifies the Token List ID for the ring the current reformation succeeds.

ACKING SITE IP ADDRESS: Specifies the IP Address of the site acknowledging the New List.

ACKING SITE PORT: Specifies the UDP Port of the site acknowledging the New List.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Recovery Abort Packet

Upon detection of a failure during reformation, a Recovery Abort Packet is sent to signify that the current reformation is to be aborted and a new reformation initiated.

1 2 3

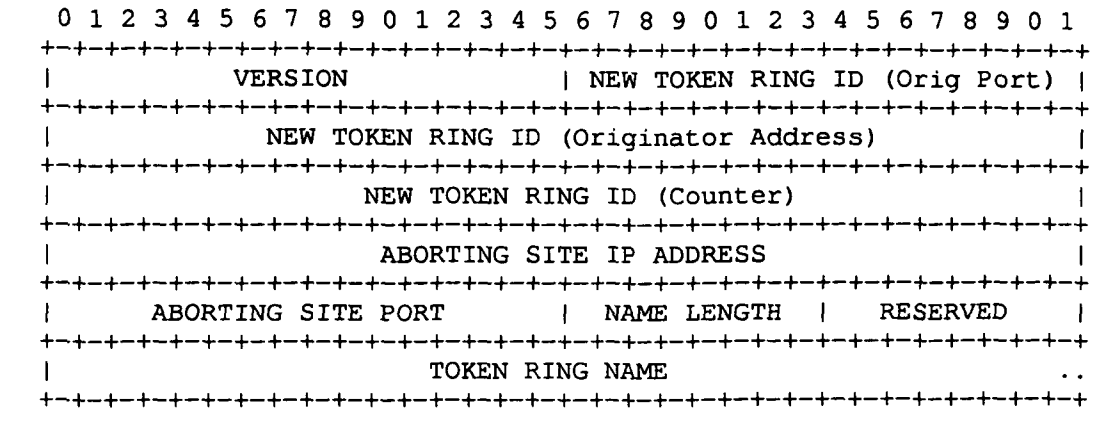

VERSION: Specifies the token ring version for the current reformation. This has the same semantics as the Recovery Start packet field.

NEW TOKEN RING ID: Specifies the Token List ID of the failed reformation.

ABORTING SITE IP ADDRESS: Specifies the IP Address of the site sending the Recovery Abort.

ABORTING SITE PORT: Specifies the UDP Port of the site sending the Recovery Abort.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Non-Member Packets

Non Member Packets are sent from a Non-Member of the token ring to the token ring to be ordered and optionally replied to by a member of the ring. Each Non-Member Packet contains a fixed header. The Token Ring ID for the fixed header is set to zero unless the Non-Member happens to know the current Token Ring ID.

Non-Member Data Packet

The Non-Member Data Packet is analogous to the Data Packet but is used exclusively by Non-Members of the token ring.

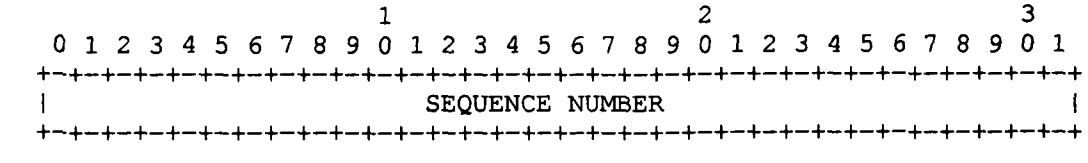

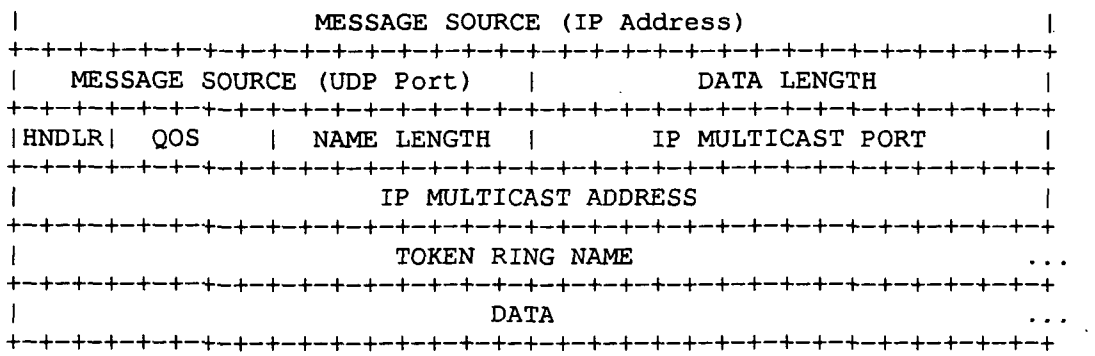

HNDLR(Handler): Specifies the handler value, if any, for the Non Member Data packet. This has the same semantics as the Data packet HNDLR field.

QOS: Specifies the desired QoS for the Non Member Data packet. This has the same semantics as the Data packet QoS of the same value.

SEQUENCE NUMBER: Specifies the sequence number of the Non Member Data packet. This has the same semantics as the Data packet sequence numbers.

IP MULTICAST ADDRESS: Specifies the IP Multicast Address this packet is destined for.

IP MULTICAST PORT: Specifies the IP Multicast Port this packet is destined for.

MESSAGE SOURCE: Specifies the RMP Process ID for the process sending the Non Member Data packet.

NAME LENGTH: Specifies the length of the TOKEN RING NAME field. This has the same semantics as the New List packet field.

DATA LENGTH: Specifies the size of the data field in octets.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. This has the same semantics as the New List packet field.

DATA: The data to be delivered.

Non-Member ACK Packet

Non-Member ACK Packets are sent from a member of the token ring to the Non-Member as a means of notification that a corresponding set of Non-Member Data Packets have been received and ordered. Optionally,

the sender may attach reply information at the end of the Non-Member ACK Packet.

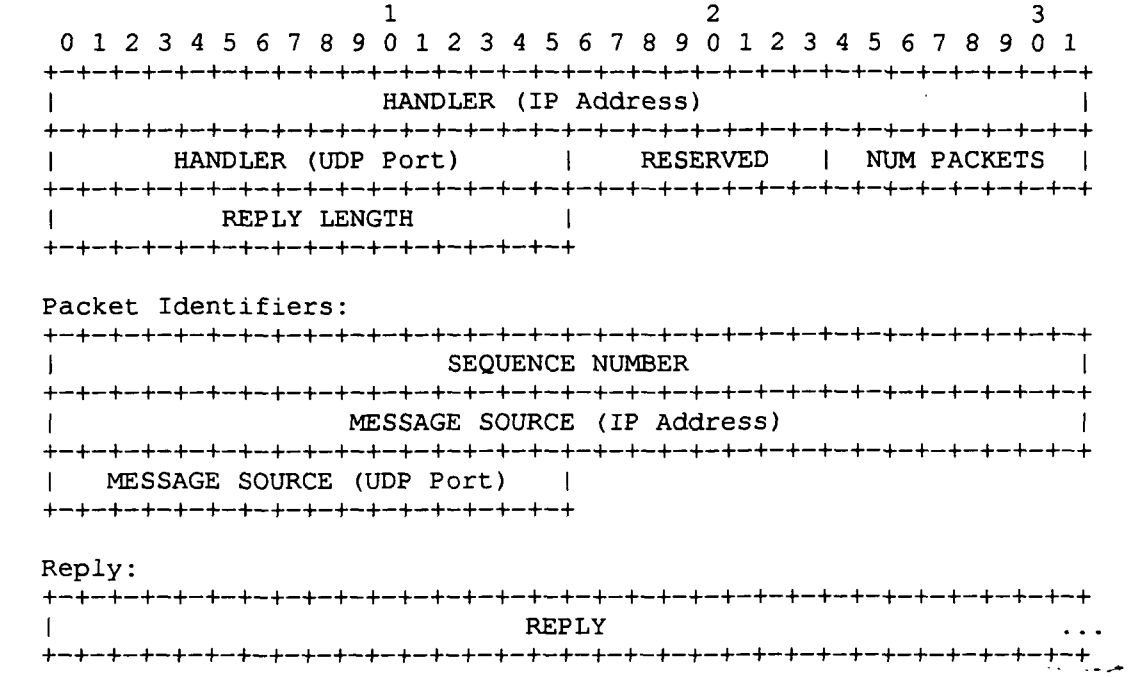

HANDLER: Specifies the RMP Process ID for the member of the token ring that generated the Non Member ACK packet.

NUM PACKETS: Specifies how many Non Member Data packets this Non Member ACK packet provides acknowledgments for.

REPLY LENGTH: Specifies the length of the REPLY field in octets.

Packet Identifiers: Each packet identifier represents a Non Member Data packet that is acknowledged by this Non Member ACK packet. The fields of the packet identifier correspond with the Non Member Data packet fields.

REPLY: The data of the reply. (Optional).

Ping Response Packet

A Ping Response packet is sent in acknowledgement of each Ping packet. These packets provide information to diagnostic tools as well as provide RMP Listeners with membership information.

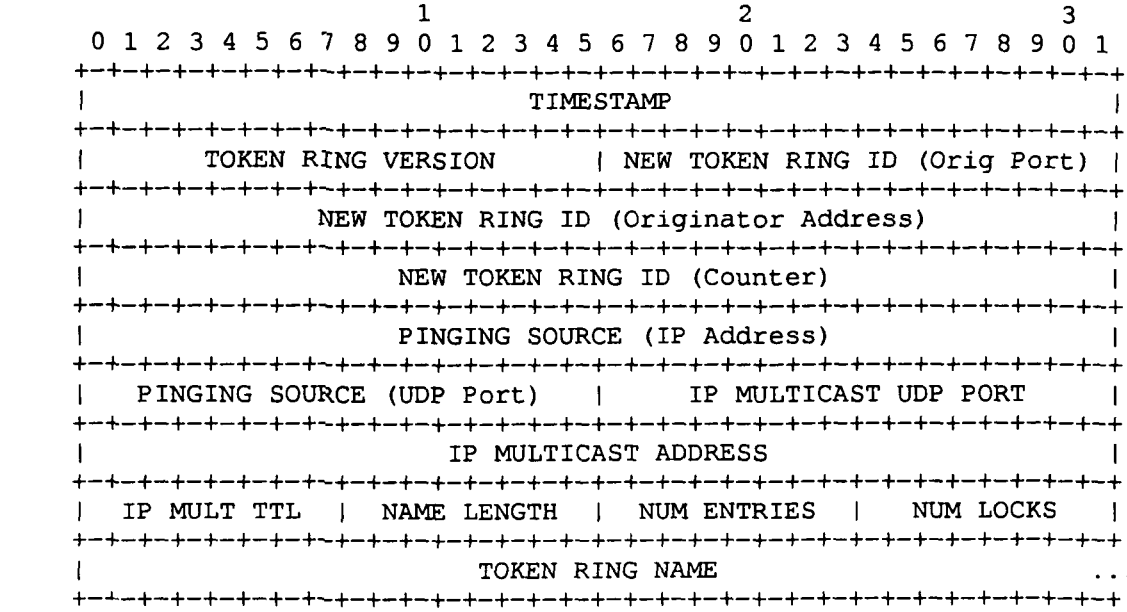

Token List Entries: 1 per NUMBER OF ENTRIES value

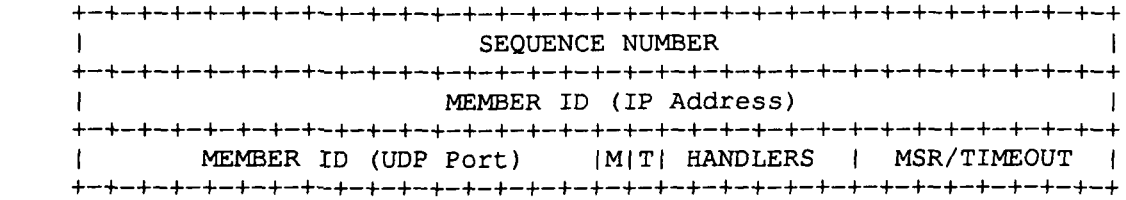

Extra Locks: 1 per NUMBER OF EXTRA LOCKS value

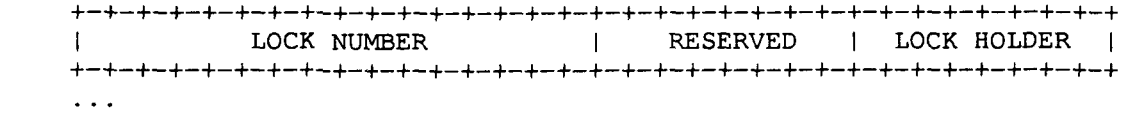

TIMESTAMP: Specifies the timestamp that was last sent to the group by the last token site.

TOKEN RING ID: This is the current Token Ring ID in use by the group.

NUMBER OF ENTRIES: Specifies the number of Token List members, or entries.

NUMBER OF EXTRA LOCKS: Specifies the number of locks being held by

members of the group. This does not count handler locks.

TOKEN RING VERSION: Specifies the version of the new token list.

IP MULTICAST ADDRESS: Specifies the IP Multicast address used by the token ring.

NAME LENGTH: Specifies the Token Ring Name length in bytes.

IP MULT TTL: Specifies the IP Multicast TTL value for the token ring.

IP MULTICAST UDP PORT: Specifies the IP Multicast UDP Port used by the token ring.

TOKEN RING NAME: Specifies the null-terminated ASCII name for this token ring. This name must end on a word boundary, which may necessitate that 1-3 extra octets of padding be included after the trailing zero of the name. These octets must be set to zero.

Token List Entries:

Each entry contains information on each member of the token list. There are NUMBER OF ENTRIES of these entries.

M(Multicast Capable): Specifies whether the member is multicast capable or not. Each member of the token list that is not multicast capable requires that each other member in the list unicast each packet that it sends to the list to these members as well as sending the packet to the IP Multicast group. Non-Members of the token ring have this field set to zero.

T(Token Ring Member): Specifies whether the RMP Process is a member of the token ring or not. Non-Members periodically get flushed from the list. The inclusion of Non-Members into the token list is used to bring new token list members up to date with respect to recent Multi-RPC activities.

HANDLERS: Specifies the Handler mask for the member. Each bit position represents a handler lock value. Non-Members of the token ring have this field set to zero.

SEQUENCE NUMBER: Specifies the next sequence number from the RMP Process that is to be delivered. This is used for Non-Members as well.

MEMBER ID: Specifies the RMP Process ID for the token list entry. This is used for Non-Members as well.

MSR/TIMEOUT: As with New List packets.

Extra Locks:

The token ring has 255 locks that members may request and release. Each one is assured to be mutually exclusive, only one member may be in possession of it at one time. The first six locks (1-6) are handler locks. Locks 7-255 are extra locks. These locks have semantics that are totally dependent on what the application desires to use them for. Each lock that is being used is represented by a LOCK NUMBER, LOCK HOLDER tuple.

LOCK NUMBER: Specifies the lock number. Valid range is 7-255.

RESERVD: Unused, zeroed when sent, and ignored when received.

LOCK HOLDER: Specifies the index in the token list of the member that holds the lock. The first member in the list is denoted 1.

Ping Packet

Ping packets are used to get information about a groups membership.

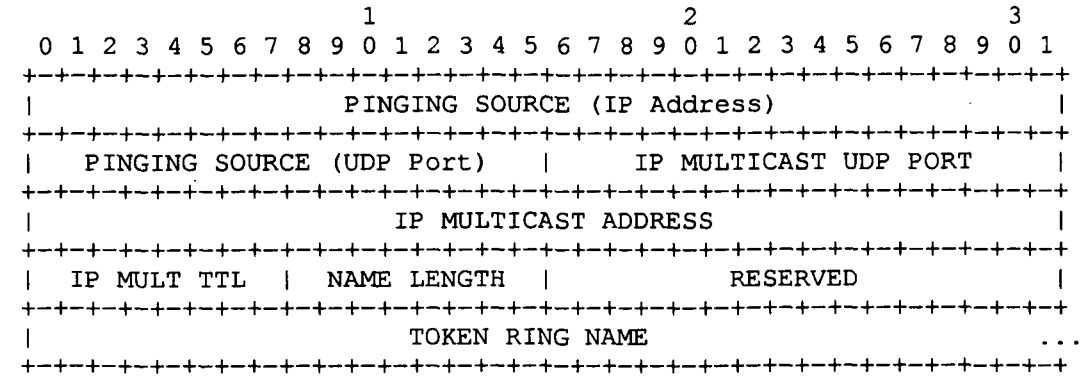

IP MULTICAST ADDRESS: Specifies the IP Multicast address used by the token ring.

IP MULT TTL: Specifies the IP Multicast TTL used by the token ring.

IP MULTICAST UDP PORT: Specifies the IP Multicast UDP Port used by the token ring.

PINGING SOURCE: Specifies the RMP Process ID for the process sending the Ping packet.

NAME LENGTH: Specifies the length of the TOKEN RING NAME field. This is done exactly as it is done in New List packets.

TOKEN RING NAME: Specifies the null-terminated ASCII name for the token ring. It is handled exactly as is done in New List packets.

Authors' Addresses:

Brian Whetten
University of California Berkeley
Berkeley, CA
Email: whetten@tenet.cs.berkeley.edu

Todd Montgomery
West Virginia University
Morgantown, WV
Email: tmont@cerc.wvu.edu

John R. Callahan
West Virginia University
Morgantown, WV
Email: callahan@cerc.wvu.edu

304 367-8348 ❑ FAX 304 367-8211 ❑ 100 University Drive ❑ Fairmont WV 26554