ANNEX R ROUTING SUBLAYER (Optional)

The Routing Sublayer is an optional sublayer above Channel Access Sublayer (CAS) and below Subnet Interface Sublayer (SIS). The Routing Sublayer supports transfers of data that needs to traverse one subnet twice or needs to traverse multiple subnets. It also supports systems with multiple independent subnets. The primary goal of Routing Sublayer is to support operation over Wireless Ring Token Protocol specified in STANAG 5066 Annex L when there is partial connectivity.

R.1. Routing Sublayer Overview

The terms Subnet is used for a set of nodes connected to an HF (or other) channel communicating using STANAG 5066. SIS (Subnet Interface Service) was originally conceived as a service to access a single subnet/channel. The Routing Sublayer does not change this service, but enables the SIS service to be provided across a concatenation of Subnets.

The STANAG 5066 service defines operation over a single Subnet. This specification adds a routing layer so that the STANAG 5066 SIS service can be provided across multiple Subnets and support repeated transmission over a single Subnet. The repeated transmission is needed to provide data communication between all nodes when Wireless Ring Token Protocol is used with partially connected topology. Routing Sublayer also enables a SIS client to communicate with peers on different Subnets, with the choice of Subnet transparent to the client. Finally, the Routing sublayer enables operation over a set of concatenated Subnets.

R.2. Goals of the Routing Sublayer

R.2.1. **Wireless Token Ring Protocol Support**

Wireless Token Ring Protocol (WRTP) as specified in STANAG 5066 Annex L, defines a framework for supporting communication between nodes sharing a channel where the nodes cannot communicate directly. It specifies a MAC level service that provides information on node connectivity to the higher layers of STANAG 5066.
The routing sublayer defined in this specification makes use of this connectivity information to enable users of the STANAG 5066 service to communicate data to nodes that are not directly connected.

R.2.2. **Routing across Multiple Subnets**

STANAG 5066 defines a service with access to a single channel (typically accessed using a single modem) and service to nodes connected by that channel. A system may be connected to multiple channels with independent STANAG 5066 services. This routing sublayer enables transparent interconnection of such services. Examples as to where this might be useful.

1. A ship with two HF connections: Skywave reach back to shore; WTRP surface wave communication to other ships. This routing protocol would enable them to interconnect.
2. A shore station with skywave access to several mobile units, each with a dedicated channel and STANAG 5066 service. The routing sublayer would enable STANAG 5066 communication between the mobile units.
3. A MANET (Mobile Ad hoc Network), with a mix of HF links and line of sight links (e.g., UHF).

R.2.3. **Transparent Client Access to multiple Subnets**

Where a system has multiple STANAG 5066 Subnets, without Routing Sublayer, there would be a need for each Subnet to have its own STANAG 5066 server. An application using STANAG 5066 would need to connect to the "right" subnet for a given peer. Operating with the Routing Sublayer enables a SIS client to address STANAG 5066 nodes on any connected Subnet through a single SIS interface.

R.3. Alternate Approaches

There are two "obvious" alternatives to the routing sublayer defined here, which are discussed below. Both are highly desirable approaches for some deployments, but they cannot address all scenarios.

R.3.1. **Application Relay**

When transferring application data over HF, it is highly desirable to proceed one hop at a time. This enables communication to be tuned for one network, without introducing relay. This can be done, and is recommended, whenever routing configuration is stable.

However, where routing can change, for example with WTRP, application relay is not viable, because an application relay connecting over SIS will not have access to the dynamically changing WTRP configuration.

Application Relay is useful for some applications such as messaging and XMPP, but other applications such as Web Browsing do not usefully support application relay.

R.3.2. **IP Relay**

Use of IP is a common approach used to build MANETs over multiple subnetworks. In many cases, this is a good and flexible choice. However, HF subnetworks have high and variable latency, which leads to problems deploying these applications over HF. Use of relay over multiple networks would compound these problems. This IP approach will not be generally viable for HF, particularly "bulk" applications using TCP. Research supporting this assertion is provided in the Isode White Paper, Measuring and Analysing STANAG 5066 F.12 IP Client[1]. Another Isode White Paper, Measuring and Analysing HF-PEP for TCP communication and Web Browsing over HF[2] shows

[1] https://www.isode.com/whitepapers/measuring-stanag5066-f12-ip-client.html
[2] https://www.isode.com/whitepapers/hf-pep-measurements.html

that an approach operating directly over STANAG 5066 can address these issues. This solution can provide good performance using Routing Sublayer.

R.4. Routing Layer Architecture

This section provides a recap of the STANAG 5066 Ed3 architecture, and shows how the Routing Sublayer (RS) extends this architecture.

R.4.1. **STANAG 5066 Ed3**

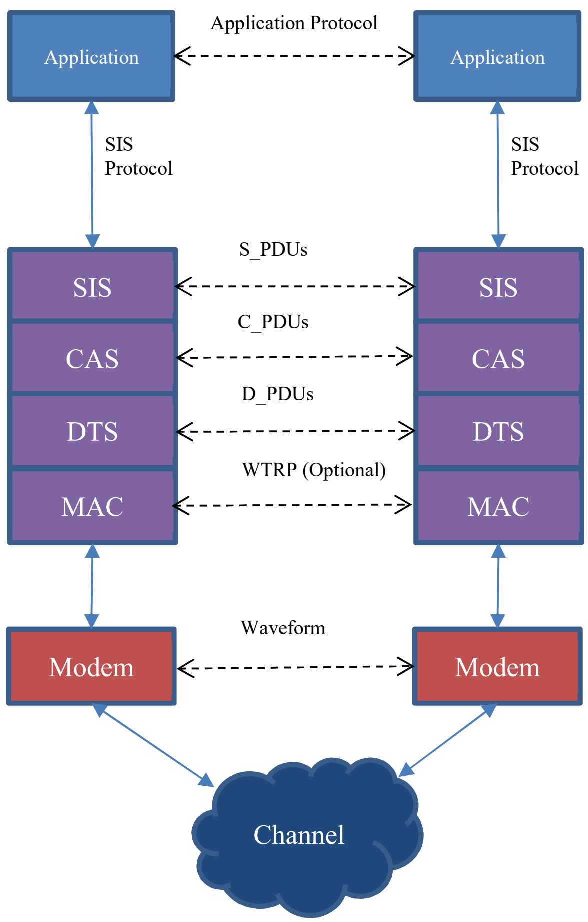

Figure R-1 STANAG 5066 Ed3 Architecture

The STANAG 5066 Ed3 protocol architecture is illustrated above. Notes on this architecture:

1. Modems communicate with a waveform, such as STANAG 5069. There will be other hardware needed to connect a modem to an HF Channel (or to another channel such as UHF).
2. STANAG 5066 Applications communicate with STANAG 5066 servers using the SIS protocol.
3. Applications running over STANAG 5066 define their own end to end application protocols, such as COSS used for ACP 127 messaging.
4. STANAG 5066 is defined as four layers, with PDUs defined at each layer that (except for MAC) are transferred by the protocol layer below. These layers are:
    a. Subnet Interface Sublayer (SIS) which communicates S_PDUs.
    b. Channel Access Sublayer (CAS) which communicates C_PDUs.
    c. Data Transfer Sublayer (DTS) which communicates D_PDUs.
    d. Media Access Control (MAC) sublayer. The communication depends on the MAC layer choice:
        i. The MAC layer is optional and is not required on a channel with just two nodes.
        ii. Annex K (CSMA). The MAC layer is procedure only, and no protocol is exchanged.
        iii. WTRP. This exchanges protocol, but uses special D_PDUs (rather than having MAC layer PDUs).
        iv. TDMA. There is a placeholder (Annex M) for this, but there is no specification.

R.4.2. **Routing with Multiple Independent Subnets**

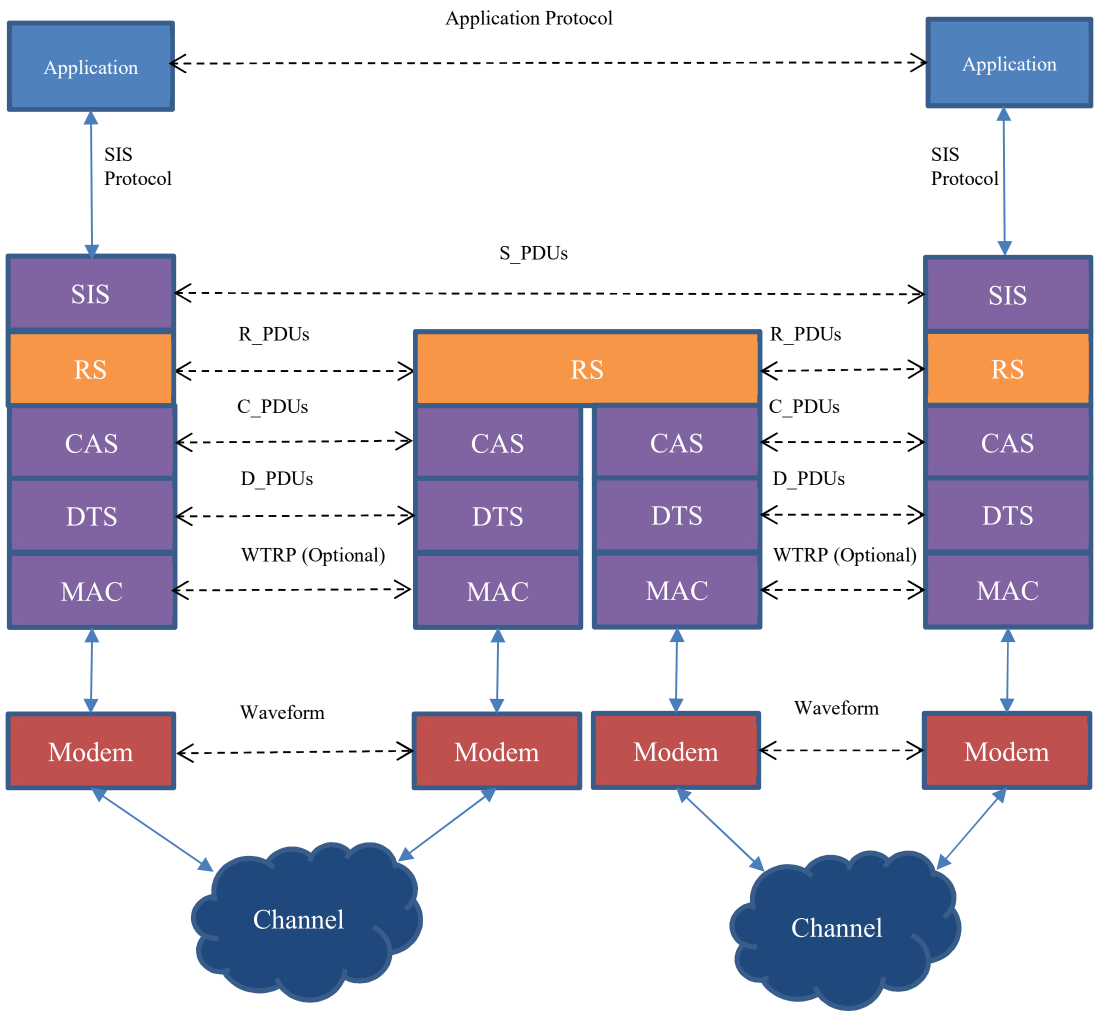

Figure R-2 Routing with Multiple Independent Subnets

Routing Sublayer (RS) is added as a STANAG 5066 layer between SIS and CAS, with exchange of R_PDUs. The SIS protocol is end to end, whereas CAS and the layers below are single hop.

The above diagram illustrates how the routing sublayer is used with two independent channels. If the LHS node in the above diagram sends data to the RHS node which is not connected, the routing sublayer communicates to the intermediate node the destination address in the Routing Sublayer protocol. This enables the middle node to route traffic between the two nodes. Note that this routing is completely transparent to the application.

NOTE: Many STANAG 5066 servers have been developed as modular products, and it may be desirable to provide the routing layer as a separate module, so that it can make use of multiple modular STANAG 5066 servers. It would be clean to use SIS for this purpose, but unfortunately SIS does not provide sufficient information and it is desirable to orient SIS to client access. The separation described here could be achieved by a new protocol to access the CAS service, which might be vendor-specific or standardized as a part of a future edition of STANAG 5066.

R.4.3. **Routing with Single Channel**

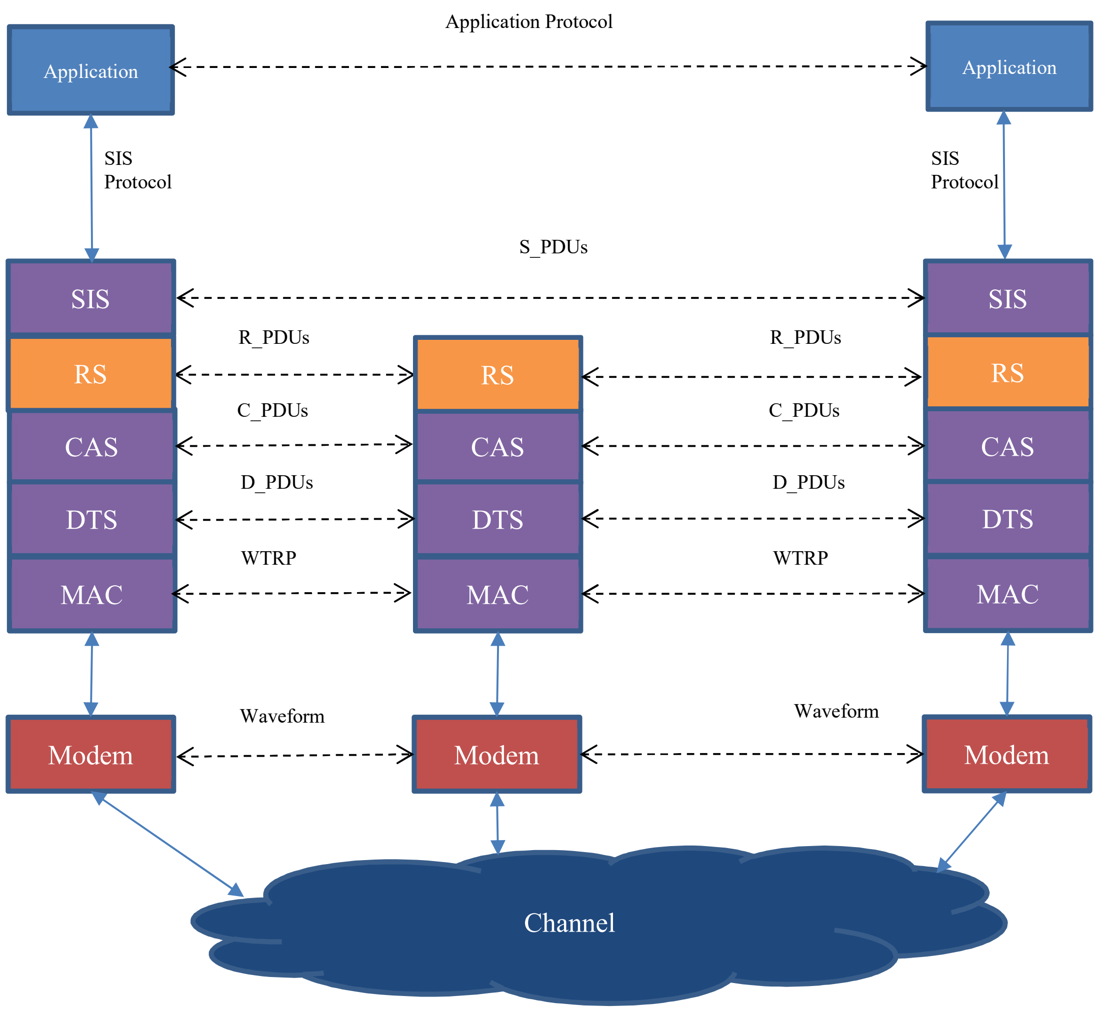

Figure R-3 – Routing with Single Channel

The above diagram shows how the routing sublayer works with a shared channel. WTRP is the only MAC layer that supports such a channel, so this description is given for WTRP, although the routing sublayer would work for a different MAC layer with the same characteristics.

Although all three modems are connected to the same channel, the diagram above relates to a scenario where the middle node can communicate with both of the end nodes, but traffic does not propagate between the end nodes. WTRP supports this configuration and will ensure that only one node transmits onto the channel at any one time.

All transmissions onto the channel are broadcast at the physical layer. When the LHS node transmits data to the RHS node, the address of the RHS node is included in the routing sublayer protocol, and the lower layers send traffic to the middle node. The middle node receives this message which is handled by the routing sublayer, which will send the data through the same stack but addressed to the RHS node. Both RHS and LHS nodes will hear data transmitted by the middle node, but will only handle traffic addressed to them.

The information to control this routing is provided by the WTRP layer specified in Annex L. This defines in Section L.2.9.2 the *next hop table*, which contains the necessary information.

R.4.4. **Transparent Client Access to Multiple Subnets**

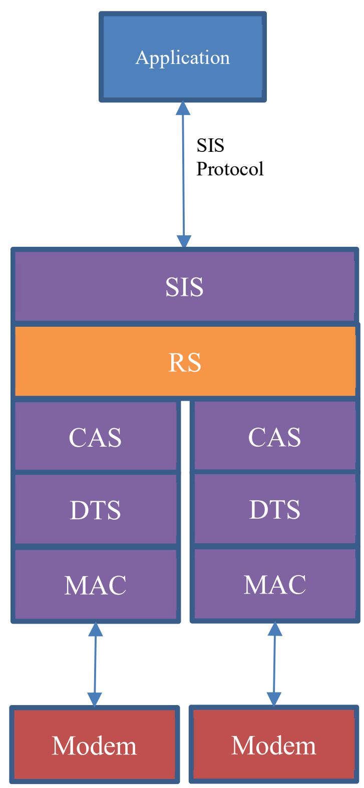

Figure R-4 – Client Access to Multiple Subnets

When a system is connected to two or more modems through independent STANAG 5066 stacks, the routing sublayer (RS) enables SIS provision to an application to transparently access the lower layers based on STANAG 5066 address of the destination. The application does not need to be aware of which modem and stack is used.

R.5. Routing Layer PDUs

The routing sublayer uses five R_PDUs. These use a single leading byte to identify the type of the PDU.

R.5.1. **Direct R_PDU**

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 0				Not used			
1	S_PDU							
n								

Figure R-5 – Direct R_PDU

The Direct R_PDU is used when the next hop is the final destination for the data. There is a single byte header of Type=0, and the rest of the data is the encapsulated S_PDU.

This R_PDU is for cases where the Routing Sublayer is not needed. It is a minimal PDU that needs to be present to provide a coherent layer service.

R.5.2. **Indirect R_PDU**

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 1				Not used			
1	MSB			Final Destination Address				
2								
3								
4								LSB
6	MSB			Original Source Address				
7								
8								
9								LSB
10	S_PDU							

n		

Figure R-6 Indirect R_PDU

The Indirect R_PDU is used when the next hop is not the final destination. The Indirect R_PDU is type 1, and includes the STANAG 5066 address of the destination. This will enable CAS and layers below to send the S_PDU to the next hop. It also includes the original source address, to enable the final SIS service to report this address to the recipient.

R.5.3. General Broadcast R_PDU

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 2				Not used			
1	S_PDU							
n								

Figure R-7 – Broadcast R_PDU

The General Broadcast R_PDU is used when data is being broadcast (or multicast), to all the nodes on a single subnet This is always used for an initial broadcast/multicast transmission. There is a single byte header of Type=2, and the rest of the data is the encapsulated S_PDU.

R.5.4. **Selective Broadcast R_PDU**

	MSB							LSB
	7	6	5	4	3	2	1	0
0	Type = 3				Not used			
1	MSB			Original Sender				
2								
3								
4								LSB
5	MSB			Destination Count=d				LSB
6	MSB							
				Destination Address				
								LSB
9								
10	MSB							
				Destination Address				
								LSB
13								
5 + 4*d	S_PDU							
n								

Figure R-8 – Multicast R_PDU

The Selective Broadcast R_PDU is used when a broadcast/multicast is intended for a specific set or nodes only. This is used when a broadcast/multicast message being relayed or for initial broadcast/multicast when there are multiple subnets connected. The STANAG 5066 address of the node that originally sent the broadcast is included, as this will not be available at on relay. Then there is a list of nodes to which the message needs to be sent. This is to ensure that a message is only re-broadcast when needed.

The number of destination STANAG 5066 addresses is encoded as Destination Count. This allows the list of destination addresses to be determined.

R.5.5. **Routing Update R_PDU**

	MSB 7	6	5	4	3	2	1	LSB 0
0	Type = 4				Not used			
1	Routing update information							
n								

Figure R-9 – Routing Update R_PDU

The Routing Update R_PDU is used to share routing information between nodes on different subnetworks. This is type 2. Routing update is a sequence of five byte information blocks. The Routing Update R_PDU is transmitted without any additional data, so the lengths of this PDU and number of blocks can be determined from the lower layers. Routing updates are only sent directly (single hop) so that the source of the Routing Update can be determined from the lower layers.

Each five byte block is encoded as follows.

	MSB 7	6	5	4	3	2	1	LSB 0
0	MSB			Reported Address				
1								
2								
3								LSB
4	Hop Count							

Figure R-10 – Block Encoding of Routing Update R_PDU

The first four bytes of each routing update information block is the STANAG 5066 address of the node being reported on. The fifth byte is a Hop Count, which indicates the number of routing hops to this node from the node that is sending the routing update. Where the node is directly reachable from the reporting node, hop count is set to 1.

R.6. Procedures of Operation

This section describes how the routing sublayer works.

R.6.1. **CAS Layer Interface**

The interface to CAS layer is modified as follows:

1. Where STANAG 5066 Annex B uses S_PDU, R_PDU is used when the routing sublayer is used.
2. Other parameters to CAS layer are specified in the following procedures.

R.6.2. **Determining Routing**

The routing sublayer needs to know the full set of addressable STANAG 5066 nodes and for each node it has to determine the next hop. This information can be obtained by four mechanisms:

1. Configuration. For example when a subnet is set up with a specified set of STANAG 5066 nodes. This could be a fixed configuration or a configuration updated dynamically by an external process.
2. From Wireless Token Ring Protocol, as specified in Annex L, which supports topologies where nodes are not directly connected. Information on this topology is passed upwards using an Annex J M_ service primitive. This is specified in Annex L as the next hop table which gives a list of all nodes in the ring, and indicates which nodes are directly connected. For nodes that are not directly connected it indicates a preferred data relay node, which is the directly connected node that is recommended to use for data relay. Annex L makes the next hop table available to the routing sublayer.
3. By discovery. There is no standardized mechanism to achieve this. Isode has defined an open protocol to achieve this in “HF Discovery, Ping and Traffic Load” (S5066-APP2).
4. Using the Routing Update procedures of this specification, which enables discovery of nodes which can be routed to indirectly over subnets that are connected indirectly.

NOTE: When used with WTRP, Annex L will determine and provide next hop information based on link quality, which will dynamically adapt to provide the best route. This information does not need to be shared with the routing protocol. The routing mechanism for concatenated networks is based on simple connectivity. It may be desirable, in a future version of this protocol, to enhance routing update to provide link quality information, which can give better routing in complex networks.

R.6.3. **Submission Procedure**

The SIS layer will provide an S_PDU and other parameters to be passed down to the CAS layer.

1. If the target address is a broadcast address, there are three distinct scenarios:
    a. If the S_PDU is provided by the SIS layer and there is a single subnet, encapsulate the S_PDU in a General Broadcast R_PDU. Then pass the R_PDU and submission parameters to CAS layer.
    b. If the S_PDU is provided by the SIS layer and there are multiple subnets for each of the connected subnets determine the list of destination

addresses for each subnet, so that all target addresses are reached by exactly on subnet. Then encapsulate the S_PDU in a Selective Broadcast S_PDU, with the Original Sender set to the local node address and the destination addresses included.

c. If the S_PDU is provided by relay (i.e., when it has been delivered from CAS layer and routing has determined the need for relay), there will be an Original Sender and a list of destination addresses to handle. Determine which subnet will handle each address, and for each subnet submit a Selective Broadcast R_PDU with the set of addresses determined to be on that subnet

2. If the target address is a single node for which the local node does not have routing information, the SIS request is rejected with the C_UNIDATA_REJECT owith reason “Address Not Routable”.

3. If the target address is a single node which can be directly reached on a connected subnet, then encapsulate the S_PDU in a Direct R_PDU and pass to CAS layer with the provided parameters.

4. If the target address is a single node which can be reached indirectly, then encapsulate the S_PDU in an Indirect R_PDU, with the final destination address set to the target address and original source address set to the source address. Then pass to the CAS layer the R_PDU with SIS submission parameters with the target address replaced by the next hop address.

NOTE: because relayed PDUs will inherently have higher latency, it **may** be desirable to transmit them ahead of non-relayed PDUs.

R.6.4. **Reception Procedure**

The CAS layer will provide received R_PDUs to the routing sublayer .Handling of Routing Update R_PDUs is described in the next section. The different PDUs are handled as follows:

1. The S_PDU from Direct R_PDUs are passed directly to the SIS layer.

2. When an Indirect R_PDU is received, the final destination is examined. If it is local, the S_PDU and associated parameters are passed up to SIS layer, with final destination address and original sender address (so that relay is transparent to the receiver). Otherwise, the S_PDU is extracted. Then Submission Procedure is followed using the S_PDU, service parameters associated with the received R_PDU and the target address being the destination address from the Indirect R_PDU. This process is termed relaying.

3. When a Broadcast R_PDU is received over WTRP, the Relay Responsible List of nodes provided by the MAC layer is considered. For each of these addresses,

determine the subnet to be used for relay. Use multicast submission procedure for this list of addresses, with the broadcast sender used as original sender.

4. When a Broadcast R_PDU is received over a subnet that is not WTRP, as list of nodes is generated from the broadcast/multicast address, with the local node and original sender removed, For each of these addresses, determine the subnet to be used for relay. Use multicast submission procedure for this list of addresses, with the broadcast sender used as original sender.

5. When a Multicast R_PDU is received, the following procedure is followed.

    a. If the original sender is the local node, the PDU is discarded and processing stops.
    b. The S_PDU is delivered locally following standard procedures.
    c. Consider the list of destination addresses included in the Multicast R_PDU. If the subnet is WTRP, consider the Relay Responsible List, and eliminate all addresses that are not in this list. If there are no addresses remaining, processing stops.
    d. Pass the PDU for relay using the multicast submission procedure with this list of addresses and the original sender.

In the event of any local failure of a directly submitted R_PDU, error information from the CAS layer ispassed up to the SIS layer.

In the event of any failure of a relayed R_PDU the error must be handled by the routing sublayer. TTL will be associated with any received R_PDU. In the event of TTL expiry, the R_PDU is discarded. In the event of any other error, the R_PDU is resubmitted. The routing calculation must be repeated for the destination, as the preferred next hop may have changed since the original submission.

R.6.5. **Updating Routing**

An node connected to multiple subnets **shall** use the Routing Update R_PDU to communicate routing information to each directly connected peer node. Routing information is provided on the nodes that the local node can reach, but which are not connected to the subnet on which the Routing Update is being sent. Routing information that needs to be shared will only arise either when a node is connected to multiple Subnets, and relates either to nodes on the other subnet or routing derived from received routing updates. This information is shared at intervals or whenever routing information changes. Note that routing update is not needed on a single WTRP network.

On reception of a Routing Update R_PDU, the node must update its local configuration of reachable nodes.
The default approach to sharing Routing Update R_PDUs on a subnet is to use ARQ communication to each node on the subnet. This will ensure that full routing information is shared about all reachable nodes to all nodes.

For an Annex K (CSMA) network or a WTRP network where all nodes are directly reachable, Routing Update R_PDUs may be shared using non-ARQ broadcast.