ANNEX U IP CLIENT (Optional)

This Annex specifies operation of IP (the Internet Protocol) over STANAG 5066, covering both IPv4 (RFC 791) and IPv6 (RFC 8200). This enables support of some IP applications over STANAG 5066.

U.1. Changes in This Edition

The core text of this annex is taken from Annex F Section F.12 of Edition 3.

The functional differences between this specification and Edition 3 are set out in Section U.9.

The key functional changes are:

1. Support of IPv6
2. Significant clarification and tidiness.

U.2. Model of Operation

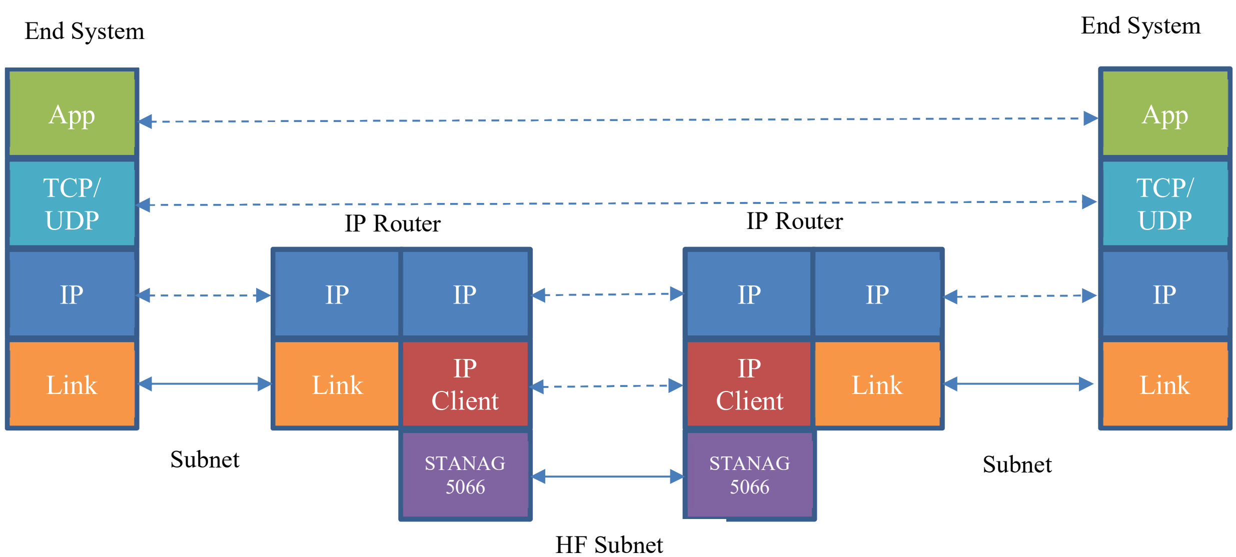

Figure U-1: IP Client HF Subnet Model

The general model of IP usage and IP Client is shown in Figure U-1. Applications communicate end to end, over end to end layer protocols such as TCP and UDP. IP is used to communicate over a series of one or more subnets. IP is a per-hop protocol with end to end implications. For any subnet, a link protocol appropriate to subnet technology is used to communicate over the subnet. IP Client, specified in this annex is used to communicate over an HF subnet using STANAG 5066. IP Client is a simple layer protocol to enable IP to be exchanged using STANAG 5066.

In most configurations, the IP Client will be logically associated with an IP Router, and not an application, end system or host. It is possible to provide a router function with IP Client co-resident on an end system.

It is also possible to configure and end system to directly use IP Client, and for the end system to have an IP address on the subnet.

U.3. Scope of Application

The model shown in Figure U-1 is general purpose, and in principle can be used to support any IP service running over an HF subnet using IP Client. In practice, the choice of IP service that can be usefully deployed over IP Client is limited. Under load, the interaction between TCP and an IP Client subnet is inefficient in many configurations, and for many deployments use of IP Client is not suitable for TCP and HTTP provision.

IP Client is suitable to provide a range of IP services, such as ICMP Ping, and some low volume protocols operating over UDP.

U.4. General Requirements

An IP client implementation **shall** be capable of sending and receiving encapsulated IP datagrams with unicast (i.e., point-to-point) IP addresses, using both ARQ- and non-ARQ-transmission modes in STANAG 5066.

An IP client **may** be capable of sending and receiving encapsulated IP datagrams with multicast (i.e., point-to-multipoint) IP addresses, using non-ARQ-transmission modes.

U.5. Encapsulation of the IP Datagram using STANAG 5066 Service

IP datagrams **shall**(1) be encapsulated within the U_DPU field of S_UNIDATA_REQUEST Service specified in Annex A, and delivered to clients at the destination node(s) within the U_DPU field of S_UNIDATA_INDICATION Service. There are no framing characters required or allowed.

The first byte of the header of the IP datagram **shall**(2) be aligned with the first byte of the U_PDU field within the primitive, and so on to the last byte of the IP datagram and U_PDU field.

The encoded bytes of an IP datagram submitted for transmission over the subnetwork **shall**(3) be bit-/byte- aligned with the octets in each U_PDU encapsulated in the S_UNIDATA_REQUEST primitive.

The least-significant bit (LSB) of each octet in the IP datagram **shall**(4) be aligned with the LSB of the U_PDU's octet.

IP_HDR (n bytes)	Byte[0]	-- IP_DATA (e.g., TCP, UDP) --	Byte[m-1]; m = k - n

Byte[0]	- - - U_PDU - - -	Byte[k-1]

Figure U-2: - Mapping of IP Datagram into an S_UNIDATA_REQUEST U_PDU

U.5.1. IPv4 and IPv6

Both IPv4 and IPv6 datagrams can be carried. The version of IP used can be determined from the first four bits of the first byte of the IP Header. Choice of IP version is transparent to this protocol, but IP Client implementations **may** need to control actions based on IP version. If needed, this version is straightforward to determine.

U.6. IP-Client Subnetwork Service Requirements

IP clients **shall**(1) bind to the HF Subnetwork at SAP ID 9.

IP support by the HF subnetwork **shall**(2) be configurable to use either ARQ or non-ARQ delivery services.

Selection of the delivery mode (established in the S_UNIDATA_REQUEST message), and priority should be chosen based on the type of IP address and other information in the IP datagram, such as the IP protocol and protocol-specific features such as port.

An IP client **may** set the subnetwork's default service requirements in the S_BIND_REQUEST message as a function of the most likely traffic that it expects to process.

An IP client **shall**(4) be capable of overriding the subnetwork's default Service Type requirements and dynamically setting the S_UNIDATA_REQUEST Delivery Mode for each IP Datagram submitted to the HF subnetwork.

U.6.1. IP Support using ARQ Service

HF subnetwork support using reliable point-to-point delivery between a pair of nodes is preferred for efficiency in the IP and higher-layer protocols, but in general cannot support IP-multicast protocols. [NB: The exceptional case is when an IP multicast address can be mapped to a STANAG 5066 unicast address, e.g., when tunnelling multicast traffic over an HF point-to-point link.]

The service definition for reliable-IP datagram delivery using the ARQ service **shall**(1) be as follows:

1. Transmission Mode = ARQ,
2. Delivery Confirmation = NODE DELIVERY,
3. Deliver in Order = IN-ORDER DELIVERY or AS-THEY-ARRIVE
4. Priority set according to Quality of Service as described in Section U.8.

U.6.2. IP Support using non-ARQ Service

If IP-multicast address groups are supported within the HF subnetwork environment (i.e., an application wishes to take advantage of the broadcast nature of the HF channel to support IP multicast), then the HF subnetwork **shall**(1) be configured in non-ARQ mode to support this requirement.

IP support using non-ARQ service **may** be used for IP unicast services.

For IP datagrams using non-ARQ service, the HF subnetwork service **shall**(2) be configured as follows:

1. Transmission Mode = non-ARQ,
2. Delivery Confirmation = none,
3. Deliver in Order = AS-THEY-ARRIVE.
4. Priority set according to Quality of Service as described in Section U.8.

The number of repeats for the D_DPUs in the service **may** be set to a value greater than one to provide some increased probability of receipt and reliability when using the subnetwork for IP multicast support.

U.6.3. Addressing using Independent Subnets (Router Architecture)

The IP Client Architecture, as shown in Figure U-1, interconnects IP subnetworks.

This is anticipated to be the most common mode of deployment. This means that the IP addressing is at a higher level than STANAG 5066 addressing. Two models of managing the address mapping are described. An IP Client implementation **may** implement one or both modes. IP users will be connected to independent subnets, which will access the IP Client through a router.

An IP Client product has two basic choices to support this:

1. It can operate by connecting to an IP router, using an implementation chosen communication mechanism.
2. It can provide IP router functionality as part of the IP Client product.

An IP Client implementation will be associated with an IP router on the local subnet. Communication between IP Client and the IP router is an implementation choice.

There are two choices for determining STANAG 5066 addresses to which an IP packet should be sent using IP Client.

U.6.3.1. **Subnet based mapping of STANAG 5066 Addresses**

The destination STANAG 5066 address is determined from the IP Subnet of the IP address, with an optional default mapping. This simple approach can be used in support of fixed IP routing, as it will be possible to specify a fixed mapping.

U.6.3.2. **Link based mapping of STANAG 5066 Address**

A second approach is to associate IP Client with a link to a remote router, and to associate the STANAG 5066 address of the peer with this link. This approach is preferable when dynamic IP routing is used between a pair of routers connected over HF, as it allows the mapping to STANAG 5066 addresses to not be affected by the IP routing.

U.6.4. Addressing using a Single Subnet (Bridged Architecture)

Another approach is to assign IP addresses on either side of an HF link to the same IP subnet. This leads to IP Client operating with a bridged architecture. In this mode of operation, IP client will need to operate as a bridge and map between the whole IP address and STANAG 5066 Address. This mode of operation will generally be impractical, but may be appropriate to support HF systems with a single node that does not have other IP connectivity.

U.6.5. IP Datagram Queueing

A key choice for IP Client implementation is handling STANAG 5066 flow control, when the load of arriving IP packets is greater than the HF subnet can immediately handle. A simple choice is to drop the IP packet when there is STANAG 5066 flow control. This is a simple approach, which follows the IP model of discarding packets on congestion and expecting the higher layers to adapt. Because of long HF delays, this discarding can lead to significant inefficiency, particularly when ARQ is used.

The alternate option is to queue arriving IP packets. This can lead to improved performance in some situations, but large queues building up with very long delays will lead to a different type of performance issues. Handling this is a key implementation choice.

U.7. Router Functions

There is a close relationship between IP Client and an associated router. Consideration needs to be taken as to where some functions are provided.

U.7.1. Segmentation and Reassembly

Any Unidata containing IP data **shall**(1) contain a complete IP datagram.

Segmentation and re-assembly, if used will generally be performed by the router. Segmentation and re-assembly **may** be performed by IP Client.

U.7.2. Internetwork Control Message Protocol

The IP Client and Router **shall** jointly ensure that various standard requirements are addressed between them:

1. That the Internet Path MTU discovery (PMTU) Protocol [RFC1191] for IP datagrams marked with the DON'T_FRAGMENT flag that also exceed the HF subnetwork MTU size.

2. That Internetwork Control Message Protocol (ICMP) is supported to provide MTU discovery.

U.8. Selecting Options to Provide Best Quality of Service

For each IP packet handled, IP Client can make a number of choices:

1. ARQ vs Non-ARQ.
2. Priority.
3. Queueing strategy, as described in Section U.6.5.
4. Discard/Filtering. IP Client may choose to discard certain IP packets.

Considerations for this choice can include general conditions, such as typical transfer speed and SNR. The most useful choices are based on each IP packet. Options that **may** be considered include:

1. Destination Address.
2. IP Protocol (e.g., ICMP, GRE, UDP, TCP etc)
3. For IP protocols with Ports, in particular TCP and UDP, the port used.

It is **recommended** to have this choice configurable. The following considerations may be helpful.

1. ARQ generally gives better performance than non-ARQ.
2. Some queuing generally leads to better performance, but very long queues need to be avoided.
3. Discarding traffic that is not explicitly allowed can be helpful.
4. Giving higher priority to control traffic such as ICMP is generally desirable.

U.9. Changes in Edition 4

General descriptive text is added, including notes on recommended scope of application of IP Client.

This annex is changed from mandatory to optional.

Support for IPv6 is added. This needed clarification of procedure, but no change of protocol.

Support of multicast is made optional. This can be non-trivial to support, and its benefit is unclear.

Removal of controls related to Differentiated Services and TOS. These are not relevant to modern IP networking and added unnecessary complexity.

Address handling rewritten, as Edition 3 text does not make practical sense. Addressing needs to be handled as subnet level.