XMPP over HF Radio: Measurements and Analysis

Steve Kille
CEO – Isode Ltd

HFIA – San Diego, January 2011

Overview

- Gave talk on theory of XMPP over HF at February 2009 HFIA
- Today: What we’ve built; measurements; and analysis
- What is XMPP and why you want to run it over HF
- General purpose enhancements and operation over Iridium satellite (2400 bits/sec)
- Additional functionality for HF Radio (STANAG 5066 support)
- Measurements
- STANAG 5066 considerations for XMPP and IM

About Isode

- Software product company based in Hampton, UK.
- Server applications for general Military/Government use and operation over constrained links including HF and Satcom.
  - STANAG 4406 formal military messaging.
  - SMTP Messaging.
  - Directory Replication and File Transfer by Email
  - XMPP (M-Link & M-Link Edge Products)
- Run over both IP and STANAG 5066.

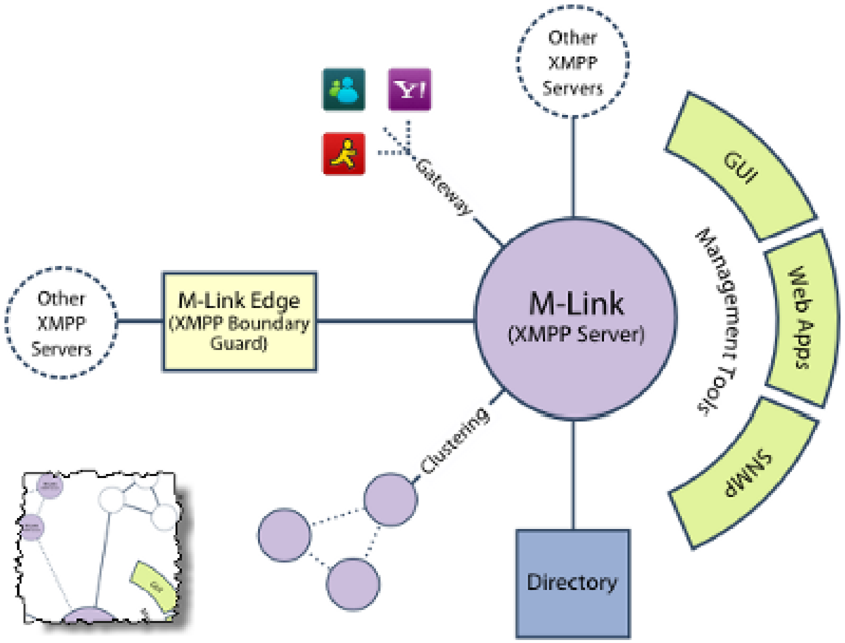

XMPP

- Internet Standard eXtensible Messaging and Presence Protocol
- Open Standard led by XSF (XMPP Standards Foundation)
- 1:1 Chat and MUC (Multi-User Chat) are key visible user services
- XMPP is also a building block for other services that need an open real time messaging infrastructure

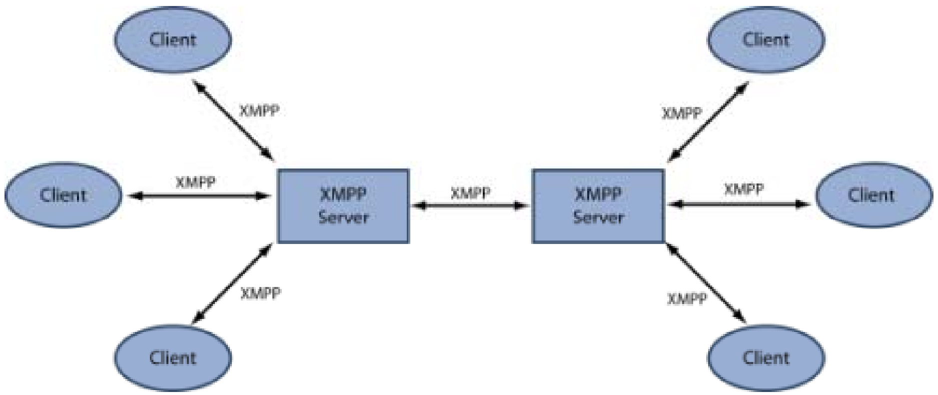

XMPP For Military Use

- XMPP is being adopted for Military Use
- Strategic direction for DISA and NATO
- Extensive operational use in ISAF
- Military XMPP clients such as Transverse (shown), JChat and SAFEchat
- Important for real time decision making and applications such as time sensitive targeting

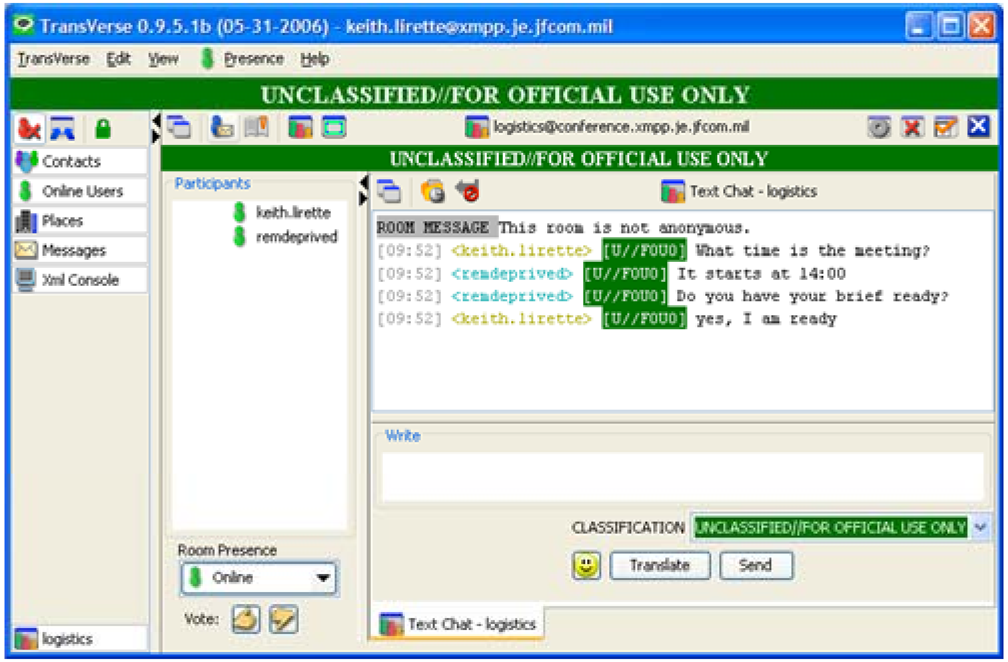

Military Deployment over Constrained Nets

- XMPP needs to work over constrained links, because they are used:
  - Slow SATCOM (e.g., Iridium)
  - Radio Links
  - HF
- XMPP seen as a fall-back technology, when videoconferencing and other technologies will not work
- Talk looks at how to achieve this

Optimizing XMPP for Constrained Links

- Key Approaches:
  - Server to Server Architecture
  - Not sending messages
  - Removing data from messages
  - MUC Optimizations
  - General purpose protocol optimization
  - HF protocol optimization

Why Server to Server

- When dealing with constrained links, there is significant benefit to handling the links between a pair of servers
  - Isolates client from poor network performance
  - Allows any standard client to be used (by avoiding building special protocols into clients)
  - Enables useful server caching (especially in support of multiple clients)
- For small systems, may be better to think of "server" as a "proxy" supporting one or two clients
- Isode XMPP approach is server to server
  - Peering controls allow independent settings for each peer

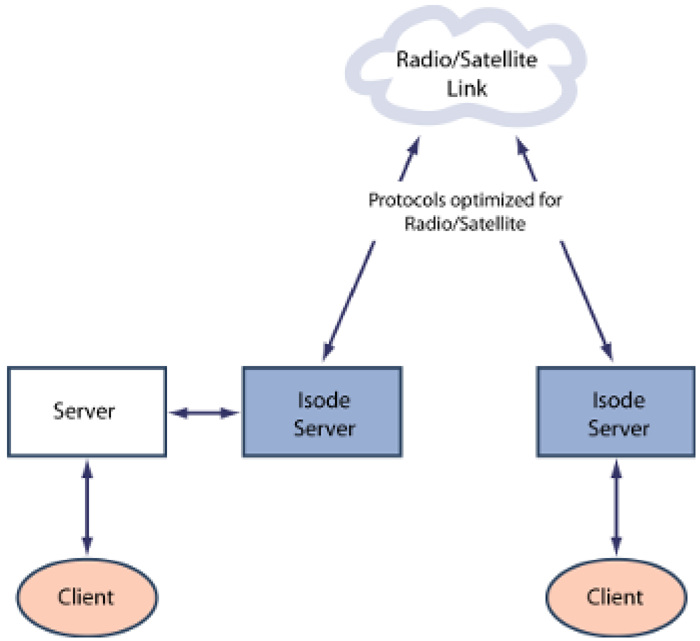

Not Sending Messages

- An easy way to optimize network use
- Some messages, such as “User Fred is Typing” (Chat State Notifications), while handy on fast networks, are best not sent over slow links
- An extreme approach (perhaps for very slow nets) would be to send messages only, and not presence information

Filtering Messages

- XMPP Messages are extensible, and XMPP clients and applications make use of this to “add value”
- User and Presence messages end up having elements that are not “core functionality” (e.g., HTML version of the message)
- Stripping these extra fields is sensible for slow links

Multi-User Chat Optimizations

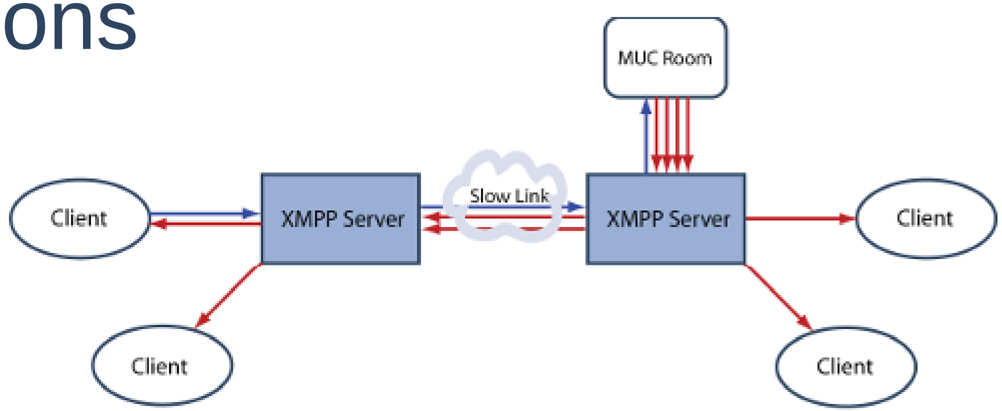

- Multi-hop (shown above)
- Disconnected working (e.g., so that ships and shore can continue MUC)
- Optimizing presence updates for room members (shared in MUC)
- An important topic – for another talk
  - Federated MUC (FMUC) is anticipated solution

Standard Protocol between XMPP Servers

- Referred to as XMPP S2S
- Runs over two TCP connections
  - New standard (BIDI (XEP-0288) will reduce to one); M-Link will be one of the first implementations
- Handshaking at start-up; Fully asynchronous after that
- Includes compression (DEFLATE)
- Good performance (for constrained networks) once connection is established
- Numbers in following slides from “M-Link & XMPP Performance Measurements over Satcom and Constrained IP Networks”
  - http://www.isode.com/whitepapers/xmpp-performance-constrained.html

XMPP S2S Performance on Established Link

- Efficient network utilization (especially for a protocol not designed for bulk data transfer)
- About 30 bytes XMPP protocol overhead (significantly less than IP overhead)
- No latency overhead (no handshaking)

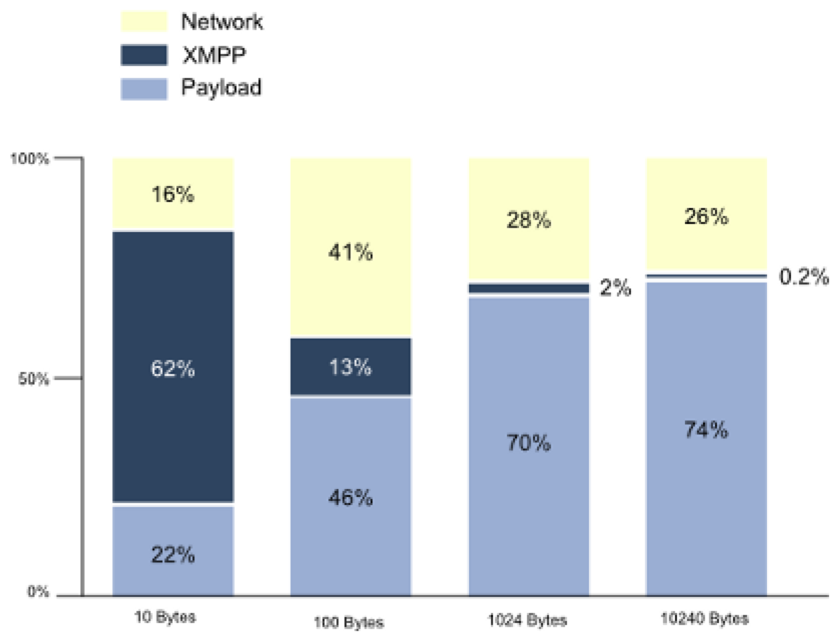

XMPP S2S Start-up Costs

- Standard XMPP S2S has the following approximate overheads
  - 20 handshakes (10 if BIDI is used)
  - 10 kByte data exchanged
- For Iridium Satellite (2400 bits/sec, 1 second latency) this gives
  - 40 seconds delay due to latency; plus
  - 36 seconds delay due to data transfer
- This is clearly not acceptable

Optimized XMPP S2S

- Isode protocol for use as an alternative to Standard S2S
  - Configurable in M-Link for use on peer to peer basis
- Key features:
  - Reduce amount of data exchanged (by using peer configuration to hold configuration data and minimize what is negotiated)
  - Single TCP connection
  - Remove all handshaking (zero application handshake) so that only a single TCP handshake remains
  - Like standard S2S once connection established
- Overhead of start-up on simulated Iridium link (2400 bits/sec 1 second latency) reduced from 75.3 seconds (Standard S2S) to 4.1 seconds (Optimized S2S)

Can you use Optimized S2S over TCP over HF?

- Yes
- For example using the IP Client Mapping of STANAG 5066
- However, it does not generally work very well
- Talk will look at the solution optimized for HF, and then compare performance with this architecture

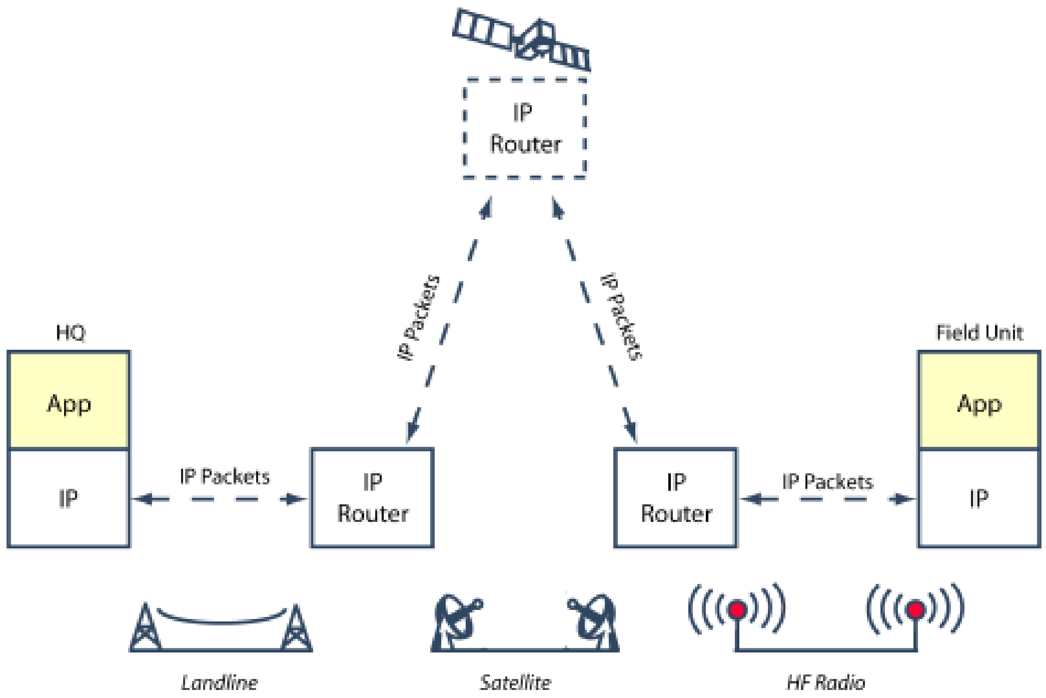

Optimized S2S over STANAG 5066

- Isode M-Link provides a direct mapping of Optimized S2S onto STANAG 5066
- Uses STANAG 5066 RCOP (Reliable Connection Oriented Protocol) to transfer compressed XMPP Messages
  - Simple data transfer
- Uses ARQ, with acknowledgements back to the XMPP server (important for dealing efficiently with STANAG 5066 timeouts)
- Typically will be used with 3G waveform (STANAG 4539)

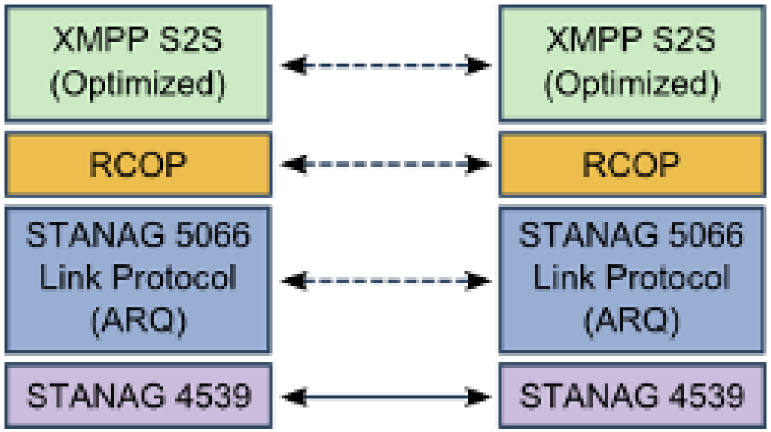

Optimized S2S over S5066 Deployment Architecture

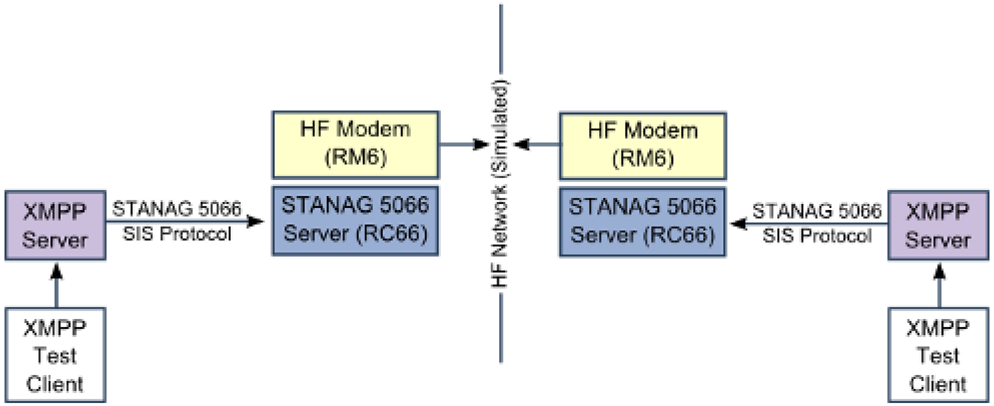

- Diagram shows how Optimized S2S over STANAG 5066 would be deployed
  - Also the test architecture for measurements made
- XMPP Servers talk STANAG 5066 SIS Protocol to STANAG 5066 servers (typically co-located with the HF Modem)
- As well as data transfer, STANAG 5066 provides:
  - Multiplexing (so modem can be shared with other applications, such as messaging)
  - Flow control (key to optimize link utilization)

Isode's HF Radio Test Setup

- Uses RapidM RM6 Military Grade HF Modems
- RapidM RC66 STANAG 5066 Servers
- STANAG 4539 3G Waveform
- Audio link replaces radios:
  - Equivalent to "perfect" HF Radio

- Measurements in "M-Link & XMPP Performance Measurements over HF Radio using STANAG 5066 and IP"
  - http://www.isode.com/whitepapers/xmpp-performance-hf-radio.html

Optimized S2S over S5066: Ping Test

Link Speed (bits/sec)	First Round Trip (secs)	Average Subsequent Round Trip (secs)
75	125.5	29.5
1200	18.0	11.7
9600	21.1	13.9

- Ping Test: application sends and echoes short XMPP message
- For 1200 and 9600 (mid and top speed HF) dominated by S5066 + Modem performance
  - First round trip is slower, primarily due to S5066 soft link establishment
  - Impact of user and XMPP data negligible
- For 75 (lowest HF speed) data transfer has significant effect
  - 1 kbyte startup data. This can usually be avoided as peer agreements can be very long lived without data transfer
  - Data transfer accounts for 6 seconds of subsequent round trips

Optimized S2S over S5066: Payload @ 1200 bps

Payload Size (bytes)	First Round Trip (secs)	Average Subsequent Round Trip (secs)
0	18.0	11.7
10	18.8	13.4
50	18.2	12.9
100	18.2	12.8
200	19.7	14.2
512	20.8	13.8
1024	32.2	15.7
2048	52.0	22.6

- Round trips (ping test) minimally affected by payload size up to 512 bytes
  - So STANAG 5066 limited round trip times might be seen by users exchanging short messages

Optimized S2S over S5066: Throughput

Payload	Measured	Utilization
10	695	29%
100	1070	45%
1024	1720	73%
10240	1970	79%

Line Speed	Measured	Utilization
75	100	68%
1200	1720	73%
9600	14800	78%

- Tests send a sequence of packets with a payload
- Although XMPP is not designed for bulk data, the results show good network utilization

Optimized S2S over S5066: Summary

- Efficient use of the underlying STANAG 5066 Service
- Good support of the XMPP application service

IP vs Direct Application over S5066: Ping Test

Link Speed (bits/sec)/Round Trip Type	STANAG 5066	Optimized S2S	XMPP S2S
75 First	125.5	-	-
75 Average	29.5	-	-
1200 First	18.0	39.1	330
1200 Average	11.7	12.9	12.7
9600 First	21.1	41.7	Not Measured
9600 Average	13.9	14.9	Not Measured

- TCP will not work at slow HF speeds
- Stable “ping” operation has similar values
- Start-up significantly slower when using IP, even for Optimized S2S

IP vs Direct Application over S5066: Throughput

Link Speed (bits/sec)	STANAG 5066	IP
75	68%	-
1200	73%	66%
9600	78%	11%

- Throughput levels over STANAG 5066 are "immediate"
- Over IP took 3 minutes to reach 66% level and 1 hour to reach 11%
  - In practice "data transfer" will be for shorter periods, so utilization will be lower
- Key problem is TCP Window size (which controls rate) and this opens very slowly due to long turnaround times

IP vs Direct Application over S5066: Summary

- Operation over IP
  - Can work reasonably when load is light and after connections are established
    - Unlike email, which almost always works badly over IP
  - Connection establishment slow
  - Bulk transfer poor
  - Does not work at lower HF speeds
  - Will have poor co-existence with other applications, particularly if load is high
  - Will deal badly with errors
  - Will deal badly with speed changes
- Direct operation over STANAG 5066 will generally work well
  - Noting constraints of underlying medium

The Big Problem with running applications over HF

- Long Turnaround time
  - Typically 5-30 seconds
  - A consequence of HF simplex nature and other characteristics
  - Compounded by interleavers (which are often key to optimizing HF performance)
- Problem made worse by other HF characteristics
  - Low speed
  - Highly variable speeds
  - High error rates

Why Application direct over S5066 Works Well

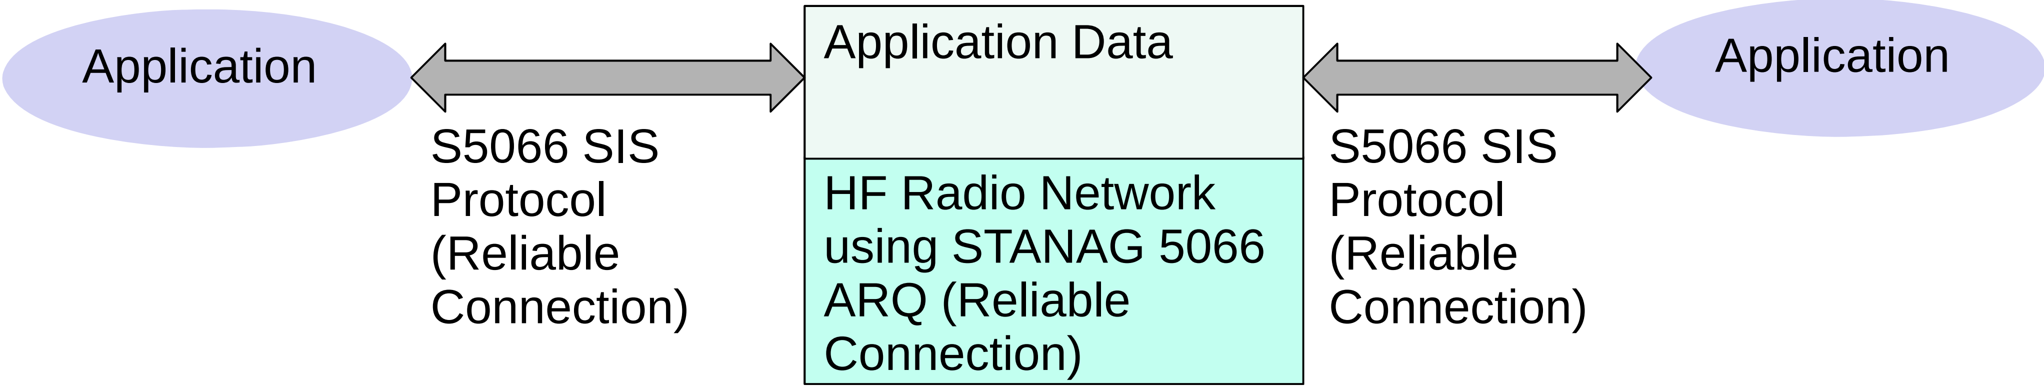

- HF Subsystem with ARQ provides reliable transfer
- STANAG 5066 SIS Protocol connection with application is reliable
- End to end reliable transfer of data
- STANAG 5066 SIS Flow Control enables efficient use of the pipe, irrespective of radio errors, delays, and change of speed

Why Application over IP over S5066 Works Badly

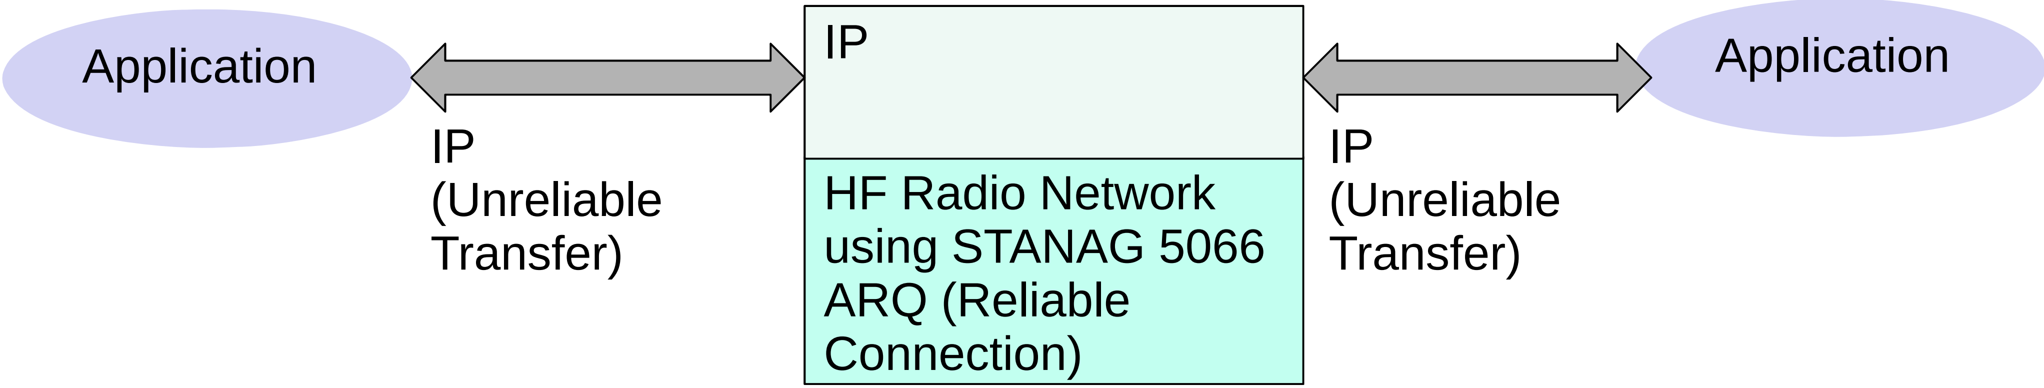

- IP is unreliable datagram protocol
- If an application does not receive data, it cannot tell if it is delayed or lost (retransmit vs wait)
  - This is a generic problem, not just TCP
- Application has no flow control, so transmission rate is hard to determine
  - Need to avoid data loss (from sending too fast) or failing to “fill the pipe”
  - Particular problem if radio errors, change of speed, or sharing with other applications

Delayed Ping & STANAG 5066

- We performed a modified ping test, where the return message is delayed
  - The intention was to simulate a chat user thinking about a response before sending
- For delays up to 4 secs, there was no change on round trip time
  - The delay was to some extent "absorbed"
- For delays of 5 seconds the round trip increased dramatically (50 seconds). Reasons:
  - Soft-link was terminated
  - Delay caused by:
    - Time to terminate soft link
    - Time to re-establish soft link
    - Delays due to RapidM collision avoidance algorithm

STANAG 5066 Link Control for XMPP & IM

- Can the “rapid soft link restart” problem be avoided?
- For a system with multiple nodes: “No”
  - Trying to guess traffic patterns is a losing strategy, so you should release the link early, and just accept the hit when traffic timing is “unfortunate”
- For two node systems: “To some extent”
  - If the system has two nodes, the XMPP (or other IM protocol) should hold the soft link open for a reasonable period, so that normal chat does not hit this problem
  - Need to “change ends” so that link will pick up both replies and more messages from the original sender
- For XMPP over STANAG 5066, the XMPP server should do its own soft link management

Other potential approaches to XMPP over S5066

- Non-ARQ mapping
  - Can enable EMCON (Radio Silence operation)
  - Could do application level acks, but for peer to peer, probably best to use ARQ
- Broadcast (also non-ARQ)
  - With a broadcast HF network and MUC distribution, broadcast support would be good
  - STANAG 5066 ed 3 would be essential, as ad hoc collision avoidance introduces too many delay

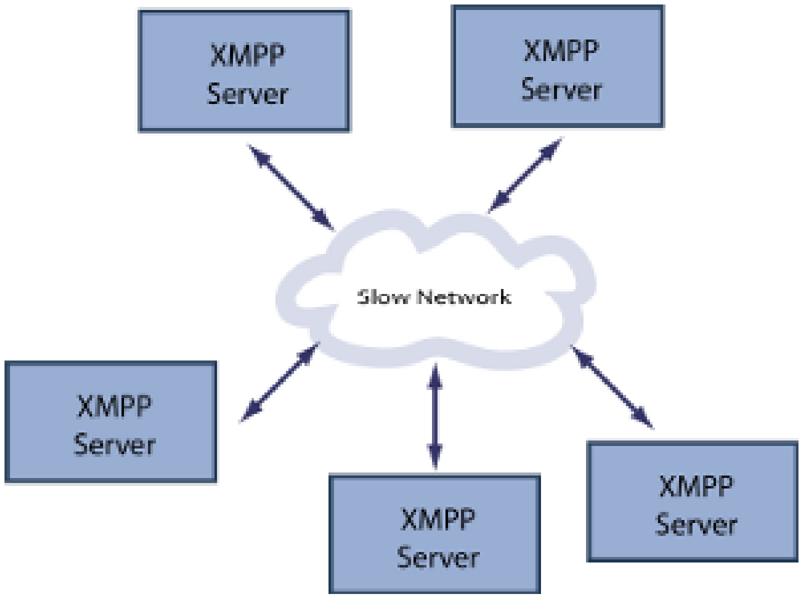

Questions?

- steve.kille@isode.com
- Presentation on Isode website
  - www.isode.com