Voice-over-IP and Asterisk

Jaco Kroon

jaco@uls.co.za

Ultimate Linux Solutions

What is Voice-over-IP?

Voice-over-IP (or just VoIP) is a generic name used for any method used to carry voice data over an IP based data link. Generally when people refer to VoIP they mean SIP (which is one possible protocol for carrying voice using IP).

Analog vs VoIP

- ISDN makes use of 8-bit samples at a rate of 8kHz.

Analog vs VoIP

- ISDN makes use of 8-bit samples at a rate of 8kHz.
- Analog thus is considered to have a dedicated 64Kbps switched circuit, VoIP generally makes use of contended bandwidth.

Analog vs VoIP

- ISDN makes use of 8-bit samples at a rate of 8kHz.
- Analog thus is considered to have a dedicated 64Kbps switched circuit, VoIP generally makes use of contended bandwidth.
- Analog should thus have lower latency.

Analog vs VoIP

- ISDN makes use of 8-bit samples at a rate of 8kHz.
- Analog thus is considered to have a dedicated 64Kbps switched circuit, VoIP generally makes use of contended bandwidth.
- Analog should thus have lower latency.
- Analog lines performs no compression, and thus doesn't lose any voice data – supposedly resulting in better quality.

Analog vs VoIP (continued)

- VoIP can (in theory) use **any** sample size and rate (scales to available bandwidth).

Analog vs VoIP (continued)

- VoIP can (in theory) use **any** sample size and rate (scales to available bandwidth).
- VoIP makes use of (usually) codebook based compression codecs, resulting in insane bitrates for voice data (up to 200 % datagram overhead!).

Analog vs VoIP (continued)

- VoIP can (in theory) use **any** sample size and rate (scales to available bandwidth).
- VoIP makes use of (usually) codebook based compression codecs, resulting in insane bitrates for voice data (up to 200 % datagram overhead!).
- Re-uses your ethernet cables, or can even go wireless!

Analog vs VoIP (continued)

- VoIP can (in theory) use **any** sample size and rate (scales to available bandwidth).
- VoIP makes use of (usually) codebook based compression codecs, resulting in insane bitrates for voice data (up to 200 % datagram overhead!).
- Re-uses your ethernet cables, or can even go wireless!
- Phone can be anywhere in the world where an IP connection is available (virtual office).

GSM

GSM is used by cellular networks and whilst it also carries voice over a packetized network it does not use IP as the underlying datagram system. It does, however, need to consider many of the same constraints, but it’s a fully controlled network and thus should (in theory at least) be more stable.

VoIP protocols/implementations

- Skype

VoIP protocols/implementations

- Skype
- Session Initiation Protocol (SIP)

VoIP protocols/implementations

- Skype
- Session Initiation Protocol (SIP)
- Inter-Asterisk eXchange v2 (IAX/2)

VoIP protocols/implementations

- Skype
- Session Initiation Protocol (SIP)
- Inter-Asterisk eXchange v2 (IAX/2)
- Media Gateway Control Protocol (MGCP)

VoIP protocols/implementations

- Skype
- Session Initiation Protocol (SIP)
- Inter-Asterisk eXchange v2 (IAX/2)
- Media Gateway Control Protocol (MGCP)
- H.323

VoIP Process Overview

- Obtain “number” to dial.

VoIP Process Overview

- Obtain “number” to dial.
- Determine “how” to “route” the call (typically pass it on to SIP proxy).

VoIP Process Overview

- Obtain “number” to dial.
- Determine “how” to “route” the call (typically pass it on to SIP proxy).
- Establish RTP (in the case of SIP) parameters (codecs + endpoint IP+ports).

VoIP Process Overview

- Obtain “number” to dial.
- Determine “how” to “route” the call (typically pass it on to SIP proxy).
- Establish RTP (in the case of SIP) parameters (codecs + endpoint IP+ports).
- Send+Received packetized voice.

VoIP Process Overview

- Obtain “number” to dial.
- Determine “how” to “route” the call (typically pass it on to SIP proxy).
- Establish RTP (in the case of SIP) parameters (codecs + endpoint IP+ports).
- Send+Received packetized voice.
- Terminate call.

Voice (de)packetization

- Digitize some audio (typically 20ms).

Voice (de)packetization

- Digitize some audio (typically 20ms).
- Encode audio data.

Voice (de)packetization

- Digitize some audio (typically 20ms).
- Encode audio data.
- Slap RTP (12 bytes) + UDP (8 bytes) + IP (20 bytes) headers on (total 40 bytes).

Voice (de)packetization

- Digitize some audio (typically 20ms).
- Encode audio data.
- Slap RTP (12 bytes) + UDP (8 bytes) + IP (20 bytes) headers on (total 40 bytes).
- Transmit packet.

VoIP CODECs

CODEC	bitrate	payload	RTP overhead
G.711[a]	64 kb/s	160	25 %
G.723-1	5.3/6.3 kb/s	23/27	173/148 %
iLBC	15 kb/s	38	105 %
GSM	13 kb/s	33	121 %
G.729	8 kb/s	20	200 %
speex	VBR: 2.15 - 44.2		

[a] ulaw/alaw or PCMu/PCMa

SIP

- “de-facto” standard (everybody uses it)

SIP

- “de-facto” standard (everybody uses it)
- XML based

SIP

- “de-facto” standard (everybody uses it)
- XML based
- media stream separate from control stream

SIP

- “de-facto” standard (everybody uses it)
- XML based
- media stream separate from control stream
- Theoretical limit of 10 calls on 4Mbps ADSL line (ATM cell size)

SIP

- “de-facto” standard (everybody uses it)
- XML based
- media stream separate from control stream
- Theoretical limit of 10 calls on 4Mbps ADSL line (ATM cell size)
- Multiple UDP ports - difficult to firewall + NAT problems.

IAX/2

- Entirely open protocol

IAX/2

- Entirely open protocol
- Binary (not easily extendible)

IAX/2

- Entirely open protocol
- Binary (not easily extendible)
- More “wire” efficient

IAX/2

- Entirely open protocol
- Binary (not easily extendible)
- More “wire” efficient
- Generally a three-times higher call concurrency than SIP

IAX/2

- Entirely open protocol
- Binary (not easily extendible)
- More “wire” efficient
- Generally a three-times higher call concurrency than SIP
- Up to 44 calls on DSL.

IAX/2

- Entirely open protocol
- Binary (not easily extendible)
- More “wire” efficient
- Generally a three-times higher call concurrency than SIP
- Up to 44 calls on DSL.
- Trunking introduces jitter (need trunktimestamps to counter)

IAX/2

- Entirely open protocol
- Binary (not easily extendible)
- More “wire” efficient
- Generally a three-times higher call concurrency than SIP
- Up to 44 calls on DSL.
- Trunking introduces jitter (need trunktimestamps to counter)
- Single UDP port (easy to firewall).

VoIP Quality Killers

- TCP vs UDP.

VoIP Quality Killers

- TCP vs UDP.
- Jitter (public enemy number one).

VoIP Quality Killers

- TCP vs UDP.
- Jitter (public enemy number one).
- Analog echo.

VoIP Quality Killers

- TCP vs UDP.
- Jitter (public enemy number one).
- Analog echo.
- Packet loss (up to about 5 % is OK, up to 10 % may be bearable).

VoIP Quality Killers

- TCP vs UDP.
- Jitter (public enemy number one).
- Analog echo.
- Packet loss (up to about 5 % is OK, up to 10 % may be bearable).
- Voice Activity Detection (VAD) and Comfort Noise Generation (CNG).

TCP vs UDP

- TCP congestion control causes unpredictable re-transmit.

TCP vs UDP

- TCP congestion control causes unpredictable re-transmit.
- It’s better to just lose a packet and skip over 20ms of audio than to wait for that piece of audio.

Jitter

Definition: Variance in latency.

Jitter

Definition: Variance in latency.

- Packets takes different durations to reach destination.

Jitter

Definition: Variance in latency.

- Packets takes different durations to reach destination.
- Packets arrive out-of-order.

Jitter

Definition: Variance in latency.

- Packets takes different durations to reach destination.
- Packets arrive out-of-order.
- Packets arrive after they already had to be “played back”.

Jitter

Definition: Variance in latency.

- Packets takes different durations to reach destination.
- Packets arrive out-of-order.
- Packets arrive after they already had to be “played back”.
- Jitter Buffers: Adaptive vs Fixed (**USE THEM**).

Why analog ECHO is bad

- On analog the delay is short enough (and in fact comforting).

Why analog ECHO is bad

- On analog the delay is short enough (and in fact comforting).
- Due to packetization process this delay becomes audiable and annoying.

Why analog ECHO is bad

- On analog the delay is short enough (and in fact comforting).
- Due to packetization process this delay becomes audiable and annoying.
- Need to perform echo cancellation.

ECHO solutions

- HARDWARE: Doesn’t require (much) CPU.

ECHO solutions

- HARDWARE: Doesn’t require (much) CPU.
- SOFTWARE: Requires (potentially) lots of CPU.

ECHO solutions

- HARDWARE: Doesn’t require (much) CPU.
- SOFTWARE: Requires (potentially) lots of CPU.
  - KB1, JPAH, SEC(2) - untested (apparently not so good).

ECHO solutions

- HARDWARE: Doesn’t require (much) CPU.
- SOFTWARE: Requires (potentially) lots of CPU.
  - KB1, JPAH, SEC(2) - untested (apparently not so good).
  - MG2: Performs equivalently to hardware.

ECHO solutions

- HARDWARE: Doesn’t require (much) CPU.
- SOFTWARE: Requires (potentially) lots of CPU.
  - KB1, JPAH, SEC(2) - untested (apparently not so good).
  - MG2: Performs equivalently to hardware.
  - Oslec: Makes everything else looks silly.

Faxes over VoIP (FoIP)

The problem: Compressed codecs makes use of codebooks, destroying the frequencies that faxes/modems rely on.

[a] http://jkroon.blogs.uls.co.za/it/voip/fax2voip2em

Faxes over VoIP (FoIP)

The problem: Compressed codecs makes use of codebooks, destroying the frequencies that faxes/modems rely on.
Solutions:

- Use G.711

[a] http://jkroon.blogs.uls.co.za/it/voip/fax2voip2em

Faxes over VoIP (FoIP)

The problem: Compressed codecs makes use of codebooks, destroying the frequencies that faxes/modems rely on.
Solutions:

- Use G.711
- Get T.38 working (asterisk 1.6?)

[a] http://jkroon.blogs.uls.co.za/it/voip/fax2voip2em

Faxes over VoIP (FoIP)

The problem: Compressed codecs makes use of codebooks, destroying the frequencies that faxes/modems rely on.
Solutions:

- Use G.711
- Get T.38 working (asterisk 1.6?)
- Fax2VoIP2email2fax[a]

[a] http://jkroon.blogs.uls.co.za/it/voip/fax2voip2em

What Asterisk is (not)

- It’s NOT a PABX

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:
  - A PABX (FreePBX is an example).

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:
  - A PABX (FreePBX is an example).
  - A billing platform.

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:
  - A PABX (FreePBX is an example).
  - A billing platform.
  - A telephony switch.

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:
  - A PABX (FreePBX is an example).
  - A billing platform.
  - A telephony switch.
  - A call centre.

What Asterisk is (not)

- It’s NOT a PABX
- FreePBX is NOT Asterisk.
- Asterisk is a set of tools (framework) that can be used to implement:
  - A PABX (FreePBX is an example).
  - A billing platform.
  - A telephony switch.
  - A call centre.
  - A telephony media platform.

Asterisk “features”

- Customizable (programmable).

Asterisk “features”

- Customizable (programmable).
- Voicemail (very flexible).

Asterisk “features”

- Customizable (programmable).
- Voicemail (very flexible).
- Call Recording.

Asterisk “features”

- Customizable (programmable).
- Voicemail (very flexible).
- Call Recording.
- Music-on-hold.

Asterisk “features”

- Customizable (programmable).
- Voicemail (very flexible).
- Call Recording.
- Music-on-hold.
- Call Queues.

Asterisk “features”

- Customizable (programmable).
- Voicemail (very flexible).
- Call Recording.
- Music-on-hold.
- Call Queues.
- Call Origination.

Asterisk “features” (more ...)

- Call “Parking”.

Asterisk “features” (more ...)

- Call “Parking”.
- Call Detail Record (CDR) generation.

Asterisk “features” (more ...)

- Call “Parking”.
- Call Detail Record (CDR) generation.
- Call Conferencing.

Asterisk “features” (more ...)

- Call “Parking”.
- Call Detail Record (CDR) generation.
- Call Conferencing.
- Supports many channel types (SIP, IAX/2, H.323, MGCP, BlueTooth, gtalk, jingle, analog/ISDN).

Asterisk “features” (more ...)

- Call “Parking”.
- Call Detail Record (CDR) generation.
- Call Conferencing.
- Supports many channel types (SIP, IAX/2, H.323, MGCP, BlueTooth, gtalk, jingle, analog/ISDN).
- Too many to mention here ...

Asterisk Advantages

- It’s dirt cheap for the features.

Asterisk Advantages

- It’s dirt cheap for the features.
- If it’s not there yet, you can add it.

Asterisk Advantages

- It’s dirt cheap for the features.
- If it’s not there yet, you can add it.
- If it’s broken, you can (probably) fix it.

Asterisk Advantages

- It’s dirt cheap for the features.
- If it’s not there yet, you can add it.
- If it’s broken, you can (probably) fix it.
- Usually GOOD support (Just don’t mention FreePBX, Trixbox, Elastix or PBX in a flash in IRC and you’ll be fine).

Asterisk Disadvantages

- Hardware is not so cheap (but still comparable with equivalent analog equipment).

Asterisk Disadvantages

- Hardware is not so cheap (but still comparable with equivalent analog equipment).
- Not nearly a sufficient number of knowledgable people in South Africa (never mind “experts”).

Asterisk Disadvantages

- Hardware is not so cheap (but still comparable with equivalent analog equipment).
- Not nearly a sufficient number of knowledgable people in South Africa (never mind “experts”).
- If you’re not sufficiently knowledgable it’s HARD to fix asterisk.

Asterisk Disadvantages

- Hardware is not so cheap (but still comparable with equivalent analog equipment).
- Not nearly a sufficient number of knowledgable people in South Africa (never mind “experts”).
- If you’re not sufficiently knowledgable it’s HARD to fix asterisk.
- It’s probably best to pay someone to implement at least the base system for you.

FreePBX

- “Well” supported (wide user base).

FreePBX

- “Well” supported (wide user base).
- Doesn’t do so well transcoding wize ...

FreePBX

- “Well” supported (wide user base).
- Doesn’t do so well transcoding wize ...
- Generally install and leave (no touchy).

FreePBX

- “Well” supported (wide user base).
- Doesn’t do so well transcoding wize ...
- Generally install and leave (no touchy).
- WAY TOO MANY OPTIONS.

FreePBX

- “Well” supported (wide user base).
- Doesn’t do so well transcoding wize ...
- Generally install and leave (no touchy).
- WAY TOO MANY OPTIONS.
- G.729 pass-through is non-existent.

FreePBX

- “Well” supported (wide user base).
- Doesn’t do so well transcoding wize ...
- Generally install and leave (no touchy).
- WAY TOO MANY OPTIONS.
- G.729 pass-through is non-existent.
- Their way or the high-way.

Asterisk Distro's

- TrixBOX CE

Asterisk Distro's

- TrixBOX CE
- Elastix

Asterisk Distro's

- TrixBOX CE
- Elastix
- PBX in a flash

Asterisk Distro's

- TrixBOX CE
- Elastix
- PBX in a flash
- All FreePBX based (probably others too).

VoIP over the Internet

- Shared data/voice connections is BAD.

VoIP over the Internet

- Shared data/voice connections is BAD.
- Jitter **Usually** not a problem.

VoIP over the Internet

- Shared data/voice connections is BAD.
- Jitter **Usually** not a problem.
- Packet-loss can be an issue (even on DSL we sometimes see up to 20 % loss).

VoIP over the Internet

- Shared data/voice connections is BAD.
- Jitter **Usually** not a problem.
- Packet-loss can be an issue (even on DSL we sometimes see up to 20 % loss).
- Latency on non-DSL connections (albeit we have clients using it quite happily over 3G with latencies up to 180ms).

VoIP over the Internet (2)

- At sustained high concurrencies (>40 calls) diginet to your provider becomes worth it.

VoIP over the Internet (2)

- At sustained high concurrencies (>40 calls) diginet to your provider becomes worth it.
- Providers generally use Premi-Cells to get sub-R1.30 prices on cellulars, and the dencity of the premis are generally too high (microwave interference - bad quality).

VoIP over the Internet (2)

- At sustained high concurrencies (>40 calls) diginet to your provider becomes worth it.
- Providers generally use Premi-Cells to get sub-R1.30 prices on cellulars, and the dencity of the premis are generally too high (microwave interference - bad quality).
- Rather get “CLI” (more expensive, often in the range of R1.45 but significantly better quality).

VoIP over the Internet (2)

- At sustained high concurrencies (>40 calls) diginet to your provider becomes worth it.
- Providers generally use Premi-Cells to get sub-R1.30 prices on cellulars, and the dencity of the premis are generally too high (microwave interference - bad quality).
- Rather get “CLI” (more expensive, often in the range of R1.45 but significantly better quality).
- Make 100 % sure your provider uses G.729 - and ONLY G.729.

Q&A

Questions?
Suggestions?
Thoughts?

Thanks!