Voice-over-IP and Video Conferencing under Linux

Telephonitis

Convenient Internet-based phone and video conferencing facilities were the reserve of Windows or Mac users for a long while. Now, Linux users can leverage the power of GnomeMeeting, a feature-rich program with an intuitive interface. BY KILIAN KRAUSE AND CHRISTIAN STRAUF

Phone and video conferences have become an every-day tool for many users, allowing them to communicate with relatives or colleagues. GnomeMeeting runs on the major desktop environments, GNOME and KDE, but also supports more exotic window managers. It co-operates with other **clients** to support platform-independent communication.

No Trick

GnomeMeeting requires a full-duplex soundcard, that is a card that can record and play back at the same time. Modern soundcards, or integrated soundchips on modern motherboards, will typically provide this. The *Advanced Linux Sound Architecture (ALSA) [3]*, which has become the standard for many distributions, and is the official choice of the Linux kernel developers as of kernel 2.6, also has no trouble supporting this function. The OSS drivers used by kernel 2.4 only support this with the *SoundBlaster Live!* card.

For video transmissions you will also need a webcam or some other device that supports *Video4Linux* (v4l) [7], for example a TV card with a video camera attached. We also recommend a headset to avoid echoes when using your microphone. USB headsets are recommended to avoid the loss of signal quality that typically occurs with analog connectors.

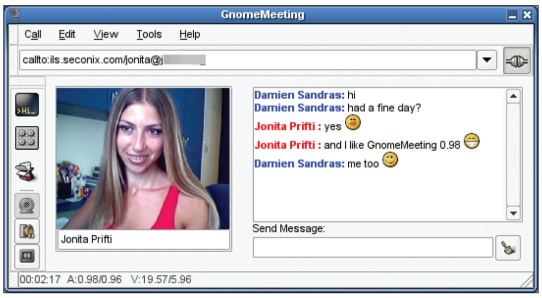

Figure 1: GnomeMeeting at work.

If you choose your audio and video **codecs** carefully, the defining factor will not be the speed of your Internet connection, but rather a low latency connection that allows a constant data flow. Useful results can be achieved with a 56k modem or an ISDN connection, although this will mean some limitations to video facilities. Of course, this will not leave you with a great deal of bandwidth to surf the Web, or download files, at the same time.

If you only require audio transfer facilities, you can use a low-footprint codec that requires only 8 Kilobits/sec. (kbps) – that is an eighth of the bandwidth provided by ISDN – per audio stream, where a connection will typically comprise one incoming and one outgoing stream.

GnomeMeeting does require a few GNOME libraries [1], but it does not need a fully-fledged GNOME installation. If your distribution does not pre-install them, you will need to install the basic GNOME libraries (*gnome-libs*), the OpenH323 library [5] (Suse, Red Hat: *openh323*, Mandrake, Debian: *libopenh323*) for the **H.323** protocol, and the *Portable Windows Library* (Suse, Red Hat: *pwlib*, Mandrake: *libpwlib*, Debian: *libpt*). These packages are available at [2].

Setting Up

There is one more requirement before you can start using your Internet phone: an IP address which is accessible directly across the Internet. An internal network that uses a **NAT** router to access the public network, will not have this. The router would need native support for H.323, and this is something that typical hardware routers, and the *h323* Linux kernel module do not fully support.

You can work around this problem by enabling **port forwarding for H.323** on your router. This means forwarding TCP 1720 and TCP 30000 through 30010 such as **UDP** 5000 through 5003, and if you use a *Gatekeeper* also UDP 5010 through 5013. Also, disable NAT for H.323 to prevent a collision between the port forwarding and NAT rules.

After setting up forwarding, you can enable *NAT IP Address Translation* in your GnomeMeeting preferences.

If the person you call has image and sound, but you do not, your firewall or router configuration may be to blame. This type of redirection will not typically work for other H.323 clients as they use other ports.

Figure 2: The Configuration Druid saying hello.

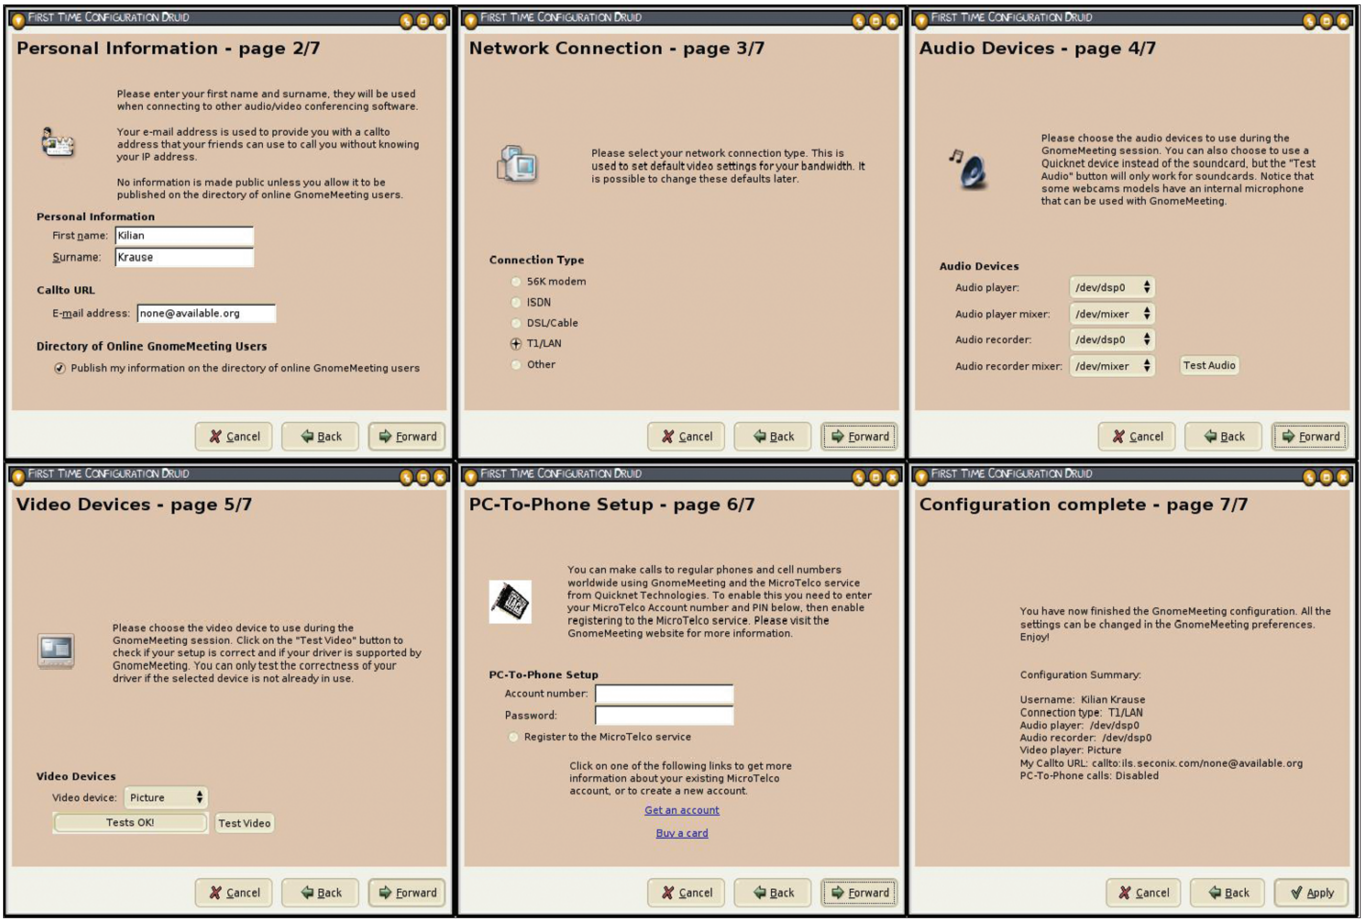

Figure 3: The druid helps you find the right settings.

When you launch the program for the first time, a *Configuration Druid* appears. This is a graphical tool that helps you set up a basic configuration (see Figure 2).

Add the appropriate details in the windows that appear; these include a directory server, which we will be looking at later, details on your Internet connection, your audio, video and phone hardware (see Figure 3). You can typically accept the defaults suggested by the druid, and perform a few tests to check whether your hardware is working. To modify these settings later, call up the configuration program by selecting *Edit / Preferences* in the menu.

Making a Call

There are many ways to call a business partner. The easiest way is to enter the network name, or IP address, of your business partner's target machine in the address line (see Figure 4). You can then click the *Connect* button to open a connection. This assumes that the target client is running and accepting incoming connections. **IPv6** addresses need to be enclosed in square brackets (see Figure 5). Specify the protocol name (*h323:@*) before the address to ensure that GnomeMeeting will understand the URL.

When a call comes in, the VoIP application asks you if you want to accept the call. Should you choose to do so, it then goes on to negotiate the connection.

Forgotten the exact target address? Don't worry – you can ask an **ILS** directory server. When you set up GnomeMeeting, the program suggests that you register your data at *ils.seconix.com*, or at *ils.ipv6.seconix.com* for IPv6. These servers by the author of GnomeMeeting, Damien Sandras, are a useful way of getting in touch with other GnomeMeeting, or even NetMeeting, users.

The above servers run filters to block pornographic content, so you can rely on clean entries. In contrast to other directories, they do not store the internal IP of

GLOSSARY

Client: *A computer or an application that utilizes the services provided by a server, or communicates with another client.*

Codec: *Short for "Coder/Decoder": A filter that manipulates data, for example, by compressing sender-side and decompressing receiver-side.*

H.323: *A protocol that handles VoIP and video conferencing. VoIP applications use this protocol to negotiate audio and video codecs, and network ports.*

NAT: *"Network Address Translation" distributes incoming packets from the Internet to the appropriate computers on an Intranet. DSL routers use this technique to allow multiple machines to share a single DSL connection.*

Port forwarding for H.323: *A NAT router redirects any packets it receives for a specific port to a specific machine on the internal network. As H.323 uses a* **UDP***-based protocol for packet transmissions, the NAT router cannot route packets directly back to the calling machine.*

UDP: *The "User Datagram Protocol" is based on the Internet Protocol (IP), just like TCP, however, it does not provide reliability by checking whether packets are actually reaching their final target.*

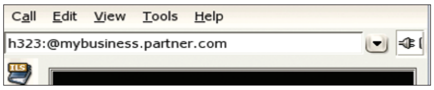

Figure 4: URL line with hostname.

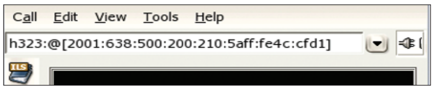

Figure 5: GnomeMeeting can handle IPv6.

your computer if you use NAT, but only the public IP, thus allowing other clients to find you behind the NAT router.

You can use the ILS browser (see Figure 6) to search the directory for potential contacts. The browser is located in the address book which you can access via *Tools / Address Book*, directly via the address book button, or by clicking on the taskbar icon with the center mouse button. You can double-click a contact to call that contact. To populate your private address book, you can either add entries manually, or drag & drop entries found on the ILS server.

If you do not need the address book, simply enter the contact data in the address line. The URL comprises the protocol (*callto*), the ILS server, and the email address: *callto:ils.seconix.com/user@email.adresse.net*

Gatekeeper

Many network administrators set up a H.323 gatekeeper [4, 10] at the interface between their Intranet and the Internet, or as a switchboard for internal H.323 connections. Among other things, this allows for alias and password-based authentication, and conferencing with contacts behind firewalls.

Some providers, like Microtelco [9], use gatekeeper software to redirect calls to PST networks (*PC2Phone*). To use a normal phone to call another across the Internet, you also need a *Quicknet* card, or a *Creative VoiceBlaster*.

Then tell GnomeMeeting your provider's gatekeeper address, and connect your phone up to the Quicknet card. When you pick up the phone, you will hear a normal dialing tone. This is extremely useful if you have Internet access but no phone, or if phone calls are more expensive than Internet access charges – as is the case in many hotels. Quicknet cards use a commercial *G.723.1* codec to provide excellent voice quality despite low bandwidth.

Configure GnomeMeeting to use a gatekeeper by selecting *Edit / Settings / H.323 Settings / Gatekeeper Settings*, and enter the gatekeeper server name, your login data, and your alias (see Figure 7).

To call a contact that uses a gatekeeper called *gatekeeper.example.com*, and is logged on as *joeuser*, type the following in the GnomeMeeting URL line: *h323:joeuser@gatekeeper.example.com.* It is important to supply the username as multiple users are typically logged on to the gatekeeper, and can only be distinguished by their names.

During Your Call

After you have opened a connection, GnomeMeeting can provide you with a wide range of statistics on the quality of that connection (see Figure 8). To enable this display, select *View / Control Panel / Statistics* below the video window.

The connection statistics reveal the percentage of lost and delayed packets, the **latency**, and the size of the **jitter** buffer. If the connection is poor, these statistics may explain why – because too many packets are being dropped, for example. Also, there are graphs that show the bandwidth of the incoming and outgoing video and audio streams. The graphs tell you whether your codecs use variable bitrates. In one case, the bitrate drops almost to zero if one of the callers is not speaking, and if there is little or no image data to transfer.

Besides statistics, there are also tabs for the *Dialpad*, *Audio Settings* and *Video Settings*. The *Dialpad* is used to dial numbers in PST networks via a PC2Phone hardware. The controls in the *Audio* tab allow you to set the output volume and microphone input level during a conversation. The *Video* tab allows you to adjust the brightness and contrast of the image.

The H.323 protocol has native features that allow you to mute, forward, or hold a call, or simply freeze the image. When you mute a call, the protocol actually interrupts the audio stream, thus allowing you to save bandwidth on a poor connection, and continue the conversation without having to re-connect.

Call History

The call history is accessible via a shortcut [Ctrl-H] or the menu *Tools / Calls History*. It stores the time, name, and network path for incoming and outgoing calls, the call duration, and the client at the opposite end. You can also see how the conversation was terminated.

The general history (*Tools / Generic History*) records additional call details. A successful call to a partner who uses NetMeeting would be recorded as shown in the History box.

If you cannot negotiate a successful connection with NetMeeting users, one of the following conditions may apply:

- NetMeeting is not using the *MS-GSM* codec. You can use the *instcodec.exe* tool on the GnomeMeeting homepage [2] to enable this codec for NetMeeting. Changes will take effect after restarting NetMeeting.

Figure 6: Querying the ILS server for other GnomeMeeting users.

- NetMeeting is running behind a firewall and/or a NAT router. As NetMeeting switches UDP ports, this issue cannot be resolved by port forwarding. Instead, a router with H.323 support is required. The *nmproxy* [6] program, which maps the dynamic ports to static ports, can be useful in this case.
- In video conferencing scenarios, NetMeeting is known to have non-reproducible video reception issues; there is no known fix for this.
- On slow links (modem or ISDN), the sound and image quality received by NetMeeting is poor. The *H.261* codec that GnomeMeeting uses for video transmissions does not achieve the same compression rate as the commercially licensed *H.263* codec used by NetMeeting. You should do without video in this case.

Table 1 provides an overview of the connection types between GnomeMeeting and other clients.

Conferencing

If you want to use GnomeMeeting for conferences with multiple participants, you will need additional software. There are a few free projects and also commercial packets available.

OpenMCU by the *OpenH323* project [5] is a free variant. This program runs on almost any Unix-style platform, and on Windows. The conference participants call the Multipoint Control Unit (MCU) which then sets up the conference connections. OpenMCU allows up to four people to participate in a conference. The MCU transmits all the video images to all participants as a single video stream (see Figure 9). The audio stream is filtered to allow each participant to hear the other participants, but not themselves. An echo mode allows

History

```
01 23:36:44 h323:1.2.3.4 calling
02 23:36:45 Opening video device USB Camera with driver V4L
03 23:36:46 Video device USB Camera, opened channel 0
04 23:36:48 Opened Plantronics Headset for recording with plugin ALSA
05 23:36:48 Started transmission of MS-GSM{sw}
06 23:36:48 Started transmission of H.261-QCIF
07 23:36:48 Connected to Kilian Krause, using Microsoft(r) NetMeeting(r)
   3.0      181/21324
08 23:36:48 Opened Plantronics Headset for playing with plugin ALSA
09 23:36:48 Started reception of MS-GSM{sw}
10 23:36:49 Updated information on the users directory ils.seconix.com.
```

GLOSSARY

IPv6: *"Internet Protocol version 6" by the Internet Engineering Task Force (IETF) will replace the current standard IPv4 some day. It users longer IP addresses that are written in hexadecimal.*

ILS: *The "Internet Locator Service" is used for entering and querying contact information. It is often used as a phone book for VoIP or video conferencing.*

Latency: *This value indicates the time that a packet with voice and video data takes to encode at the source, cross the wire, and decode at the target end. The smaller this value is, the smaller the delay for the two callers.*

Jitter: *An indicator of the quality of the traffic flow across the Internet connection between the sender and the receiver. It is calculated as the difference between the times that individual voice or video packets take for the trip from the sender to the receiver. If the jitter is too large, it will be impossible to hold a conversation. To compensate for this, GnomeMeeting uses a Jitter buffer that caches data and thus smoothes out any deviations. This also avoids gaps when packets are dropped.*

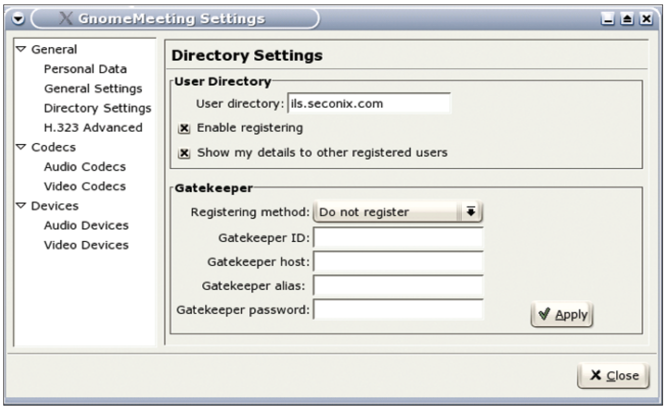

Figure 7: Gatekeeper settings.

Table 1: Connectivity with third-party clients

Client	Protocol	Notes
OpenH323-based clients (OpenPhone, OhPhone, OpenMCU, CPhone, OpenAM, OpenIVR, Asterisk, ISDN2H323)	H.323	unrestricted use
Polycom MCU	H.323	unrestricted use (Polycom update may be required)
Cisco IP phones	H.323	unrestricted use
CuSeeMe	H.323/SIP	can be used with H.323
MSN Messenger, Nero SIPPS	SIP	not usable, SIP use planned as of GnomeMeeting version 1.00
Y! (Yahoo Messenger), Kazaa Skype	proprietary	protocol not disclosed, not usable

you to hear yourself for testing. OpenMCU is suitable for occasional use, but crashes tend to inhibit regular use.

Polycom MCU [11] by Trinity provides more stability. From the user's viewpoint, the features still apply.

As a MCU generates an individual datastream for each conference participant, it makes sense to use this on servers with a high bandwidth; use on DSL lines is only possible in theory.

An MCU can define virtual rooms to avoid collisions between various groups. You can select a room via the URL line: *callto:roomname@mcu.server.remote-host.com*. As the MCU itself does not provide any authentication facilities, a combination that uses a gatekeeper to log users on is recommended. This ensures that only authorized users will have access to specific rooms, preventing interns from crashing a meeting of executive officers.

Forwarding

"I'm out for the afternoon, but I still want to take my calls." This kind of statement motivated the GnomeMeeting programmers to develop a feature that can forward calls to another H.323 endpoint. This could be an answering machine which has voice or video capabilities (IVR) or even another H.323 client.

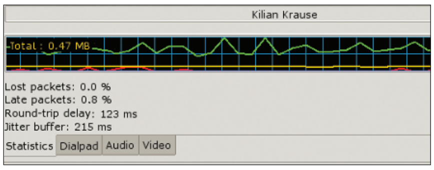

Figure 8: Connection statistics.

You can choose to forward incoming calls when a line is busy due to a previous call, or if no-one answers the call within a certain time. These settings are available in *Edit / Preferences / General / H.323 Advanced / H.323 Call Forwarding*.

In the Pipeline

The current developer version has removed a few restrictions that applied to the stable version. For example, it can use ALSA to communicate directly with the soundcard without an OSS emulation. Also, system notifications no longer depend on the GNOME *ESound* sound system. This means that users of other desktops, such as KDE, can hear the phone ring when a call comes in.

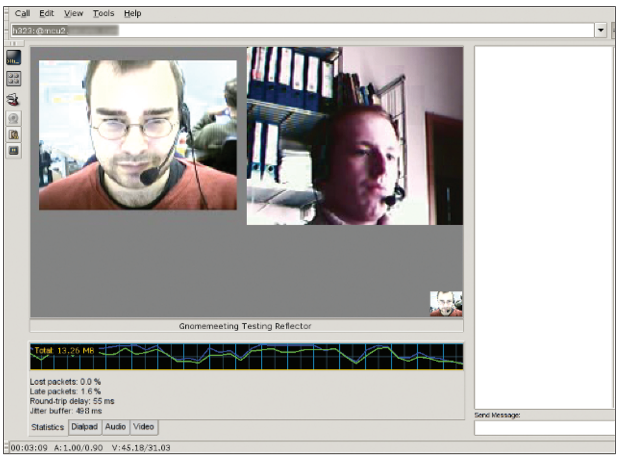

Figure 9: MCUs allow for conferences with multiple participants.

In future it should be possible to add on software plug-ins for third-party (commercial) codecs. The right codec will allow video conferencing over a modem connection. Additionally, support for the *Session Initiation Protocol* (SIP) has been in the pipeline for a long while now. This would allow GnomeMeeting to exchange calls with MSN Messenger users. The GnomeMeeting developers have decided to postpone this goal until after the version 1.00 release.

There will also be a function for stopping, starting and scaling the video stream.

If you do run into problems with GnomeMeeting, despite the intuitive interface, the FAQ on the website at [2] and the mailing list at [8] both provide additional help. ■

INFO

[1] GNOME: *http://www.gnome.org/*

[2] GnomeMeeting: *http://www.gnomemeeting.org/*

[3] ALSA: *http://www.alsa-project.org/*

[4] H.323 gatekeeper (International Engineering Consortium): *http://www.iec.org/online/tutorials/h323/topic06.html*

[5] OpenH323: *http://www.openh323.org/*

[6] nmproxy: *http://www.cryogenic.net/nmproxy.html*

[7] Video for Linux resources: *http://www.exploits.org/v4l/*

[8] GnomeMeeting Mailing-List: *http://mail.gnome.org/mailman/listinfo/gnomemeeting-list*

[9] Microtelco Site: *http://www.linuxjack.com/*

[10] GNU gatekeeper: *http://www.gnugk.org/*

[11] Polycom MCU: *http://www.trinityvideo.net/products/polycommcu.htm*