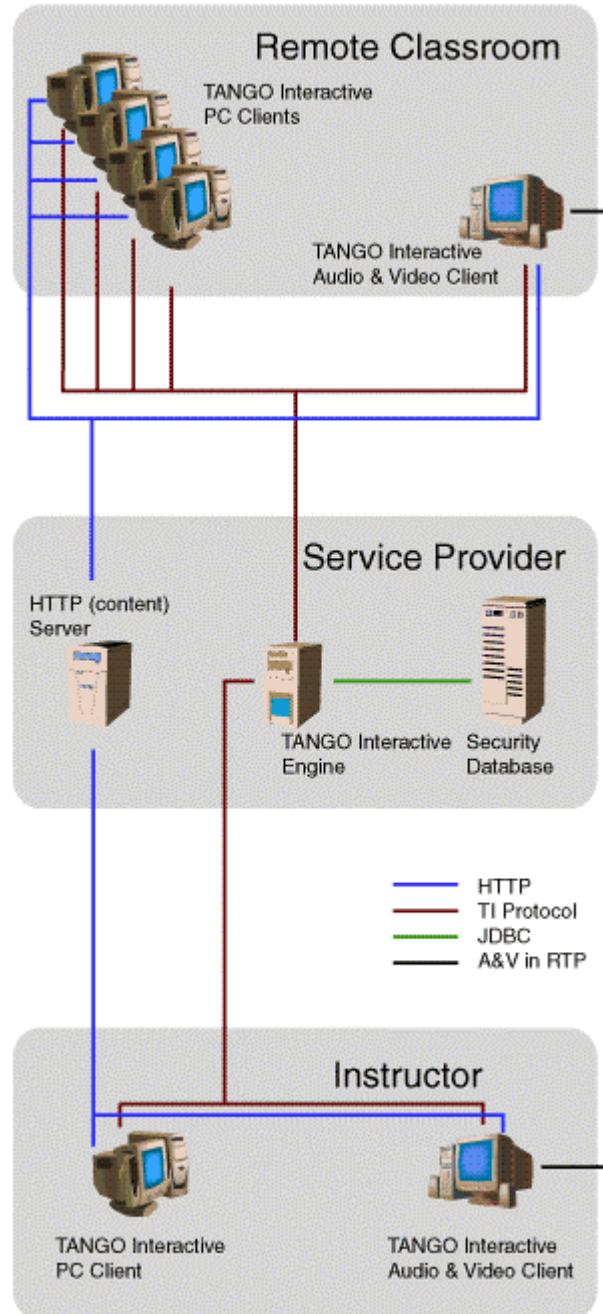


Typical arrangements of Virtual Classrooms

In practical applications of *Interactive Synchronous Distance Learning (ISDL)*, we usually deal with a small number of typical network topologies. The two most basic arrangements are: *Instructor to a Classroom* and *Instructor to a Group of Geographically Dispersed Students*. Arrangements that are more complex may include a combination of the two basic topologies and, possibly, a local classroom from which the instructor delivers the class to both local and remote students.

In the following, we will discuss the two basic arrangements in detail and provide general guidelines for setting up more complex distance learning topologies.

Figure 1 (right) shows a typical *Instructor to a Classroom* setup. In this topology, the instructor delivers the class to a group of students gathered in one computerized classroom at a remote location. There is no limit on the distance between the instructor and students locations. TANGO Interactive™ Virtual Classroom requires network connectivity between both locations, using either the Internet or a Virtual Private Network.



Multicast versus Unicast on today's Internet

Before we go any further with explanation of different network topologies for ISDL, it is necessary to explain the role and status of multicast transmission on today's packet switched networks. TANGO Interactive™ Virtual Classroom provides a voice over IP application. Actually, both audio and video are supported. These media streams are being sent from the instructor to students, and can be sent from students to the instructor as well.

Delivery of audio and video streams can be organized in one of two ways: either a separate stream is sent from the instructor's workstation to each student's using the primary IP address of the student's machine, or just one multimedia stream is sent from the professor's workstation to a class D (multicast) address to which all students' workstations listen. In theory, the later solution is far superior.

The problem is that, on today's Internet, the support for multicast is very incomplete. Many ISPs offer no multicast, and, even on campuses, the coverage is far from universal. It is always possible to join the MBone by setting up multicast tunnels, but this is a relatively complex operation, which significantly increases the cost of the customer's network support.

The multicast technology has been popular with large, technically savvy universities and with certain DoD agencies. Our recent studies show that even at the technically capable organizations the cost of maintaining the MBone connection is so high that many of the current multicast technology users switch to the less expensive unicast methodology. This transition is stimulated by the improving bandwidth and stability of the Internet backbone.

TANGO Interactive™ Virtual Classroom has been designed around the concept of inexpensive, commodity based technology. Consequently, given the problems with multicast availability, the default transmission method for BuenaVista, our video and Voice-over-IP application, is unicast. Multicast support is only available as a custom solution upon explicit request. This tradeoff imposes certain limitations on system scalability, and it has an impact on the recommended Virtual Classroom configuration.

1. "Instructor to a Classroom" setup

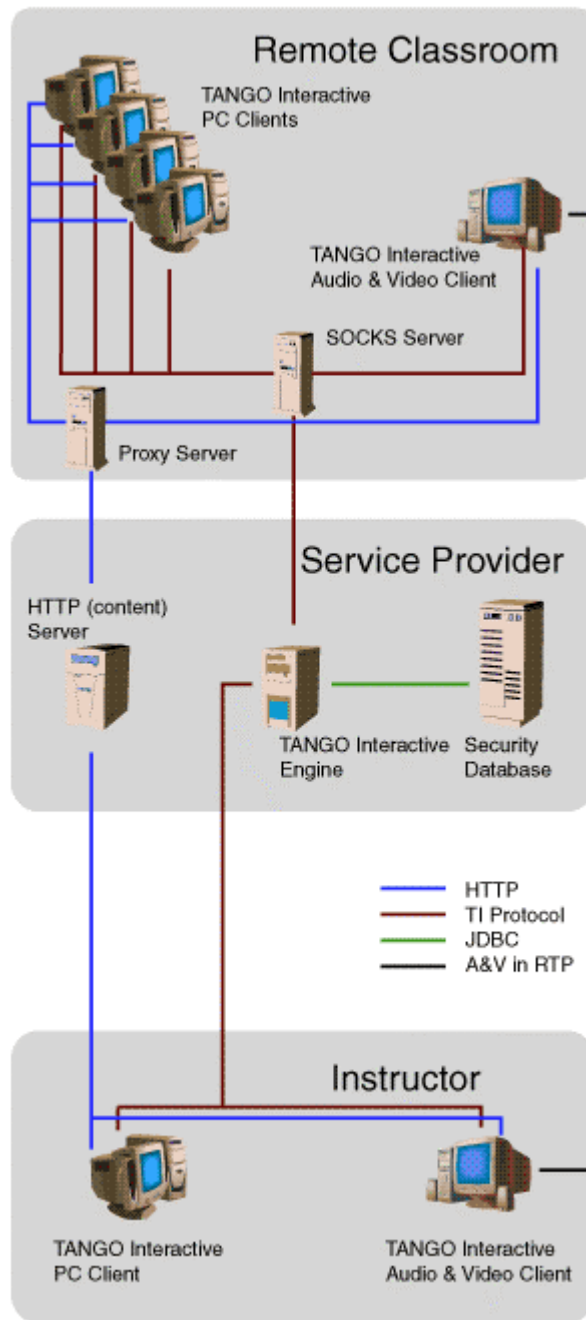
In the simplest case, *Instructor to a Classroom* setup involves the instructor workstation at one location and a number of PC or UNIX workstations installed in one physical classroom at a remote

location (see Figure 1 above). With such an arrangement, it is not necessary for students in the classroom to receive separate audio and video streams. A combination we have successfully tested at many locations consists of a dedicated audio and video workstation, connected to a large screen display and to an acoustic amplifier or a set of good quality amplified speakers. The large display is used to show the real-time video of the instructor, and the speakers are transmitting instructor's voice. To ensure interactivity, the instructor can activate the full track audio-video agent on each student's workstation, so that the students can freely ask questions using audio channel. If the student's workstation is equipped with video camera, student's video will be shown on the large screen too.

In situations where high system availability is critical, we recommend that at least one of the students' workstations is set up so that it can be quickly turned into the audio-video workstation. A simple switching device or a projector with double, selectable input is very helpful in such situations.

Recommended configurations for the audio & video and student workstations are discussed in the section [Configuration of the workstations.](#)

2. Classroom Behind a Firewall



A classroom used to receive ISDL classes may be a part of a corporate Intranet and be placed behind a firewall. TANGO Interactive™ supports such network configuration. To receive classes delivered from outside of the firewall, the receiving site must provide two proxy servers: a standard HTTP proxy (e.g., Netscape Proxy Server) and a SOCKS4 proxy.

Such configuration is displayed in Figure 2.

The HTTP proxy server channels all TI VC traffic, including Session Manager and application code download, as well as all HTTP data traffic created by TI VC application modules. Installation of the HTTP proxy server may be beneficial even if there is no firewall involved. A proxy server local to the classroom can significantly improve

system performance, especially on slow networks.

To channel the HTTP traffic through the proxy server, all browsers used in the TI Virtual Classroom must be pointed to the proxy server. The exact method of configuration depends on the browser and on the general site proxy configuration policy.

In addition to the HTTP network traffic, the TI VC system needs a connection to the TANGO Interactive™ class server. Since TI uses proprietary and application-dependent communication protocol, the system needs a SOCKS proxy server to traverse the firewall. At present, TANGO Interactive™ requires SOCKS4 protocol. Most SOCKS5 proxy servers support version 4 of the protocol as well.

TANGO Interactive™ system uses browser SOCKS setup and it does not need separate configuration. TI plug-in knows how to retrieve browser's SOCKS proxy protocol setup.

Current version of TANGO Interactive™ VC does not support SOCKS protocol for BuenaVista. We expect that this feature will be added in version 2.1. BuenaVista proxying will require a proxy server supporting SOCKS version 5.

Netscape Proxy Server supports both HTTP and SOCKS4/5 protocols. In general, TANGO Interactive™ VC should work with any HTTP and SOCKS proxy.

Special network configurations

Certain organizations use "IP translation" or "IP masquerading" on their internal networks. In essence, this solution involves using addresses from the IP class A 10.*.* range on the internal network. The machines are known via different addresses externally, and address translation is being done by the ISP.

As a rule, the translation supports TCP protocol, while all UDP protocols are being rejected by the gateway. Frequently, special measures must be taken to enable UDP-based applications for such networks.

In most cases, TANGO Interactive™ core system will work normally across such "pseudo-firewall," but, if there is no UDP support on the gateway, BuenaVista will not.

If enabling address translation for UDP is not possible, the customer can request a special version of BuenaVista which uses TCP for media transport. While this is technically non-optimal solution, it will enable BuenaVista on the networks hiding their IP addresses.

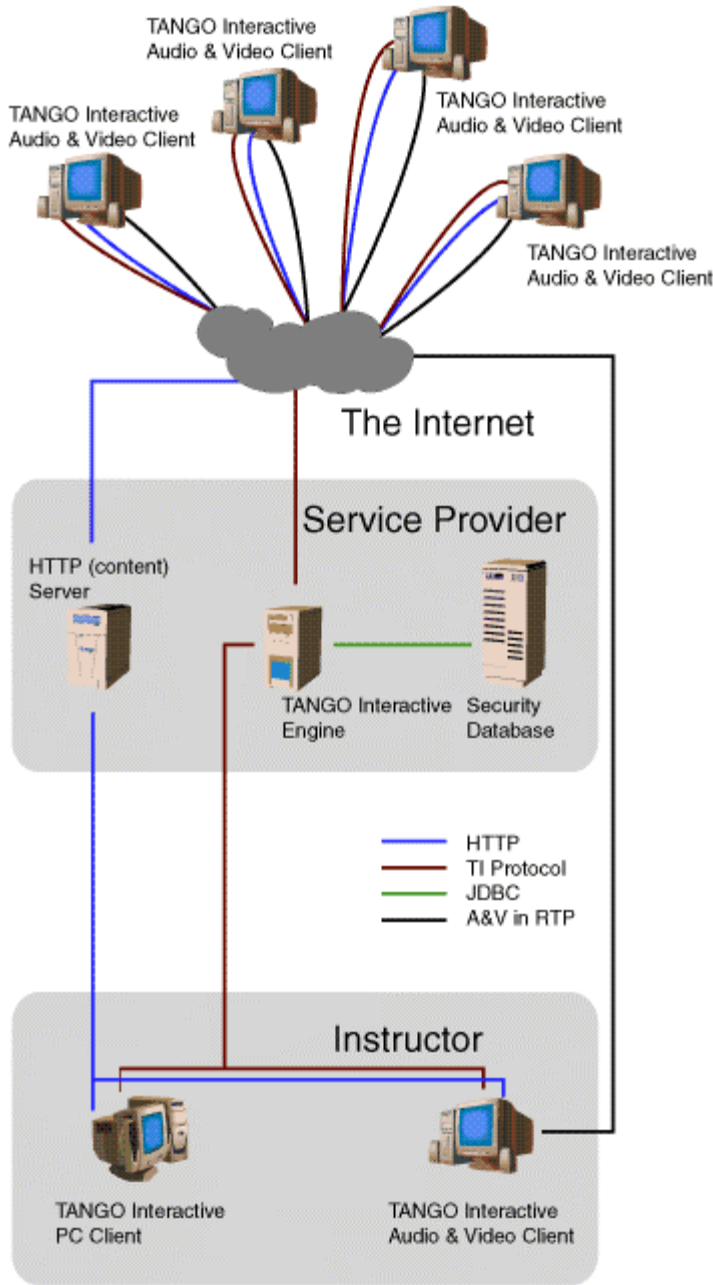
3. "Instructor to a Group of Geographically Dispersed Students" Setup

This setup is fundamentally different from the first one. The students, instead of being gathered in one of few classrooms, are

attending the class from individual workstations located at random places. They may access the class via a dial-up link or cable modem provided by ISP or via a corporate or campus link to the Internet.

The basic problem with this approach is scalability of audio & video. Obviously, audio (and, optionally, video) stream must be delivered to the workstation of each class participant. At present, use of multicast for such class arrangement is out of question: there is practically no support for multicast by the ISPs.

Geographically dispersed users, having no benefit of the technical support usually available in institutional virtual classrooms, impose special demands on the ISDL software and Virtual Classroom support personnel. Construction of multicast tunnels is usually beyond technical capabilities of such users.



There is a technical limit on the number of users that can be concurrently supported in such configuration. The unicast mode (see below) of BuenaVista creates performance and bandwidth

problems on the instructor's workstation. The practical limit for a number of class recipients in the "dispersed" setup is roughly equal 20 for strictly one-way audio. For fully interactive two-way audio, the maximal number of users is 8 to 10. These numbers can be a little higher if high-end workstations are used and the network bandwidth is very good (T1 range).

Connectivity requirements for dispersed users are discussed below in the section "[Protocols and bandwidth.](#)"

4. Mixed configurations

In real-life applications, we often encounter situations where all three of the network topologies discussed above are intermixed. TANGO Interactive™ VC permits such configurations. It is possible to have a number of Virtual Classrooms on-line, some of them connecting via firewall, along with a number of individual, "dispersed" users. TANGO Interactive™ VC has been successfully used in such situations. The legitimate question is, then, how many classrooms and individual students can TANGO Interactive™ VC concurrently support?

There is no exact answer to this question, as scalability limits depend on both hardware and software. Faster machines used as servers and clients will push the numbers higher. Remember that TANGO Interactive™ VC has been designed to support interactive synchronous distance learning. The number of students in the classroom is limited by the very nature of the process. Above certain limit, the interactive class becomes a broadcast.

The most stringent limitation is imposed by the audio & video agent. The system will support up to 25 audio & video channels. (BuenaVista uses one-to-many concept of audio & video transmission. This means that at any given time, there is only one speaker addressing entire class or breakout group. The instructor is in control of who can send audio and video). Using *Instructor to a Classroom* topology, or a mixture of different topologies, TANGO Interactive™ VC can easily support up to a hundred of concurrent students. For "dispersed users" topology, the number of concurrent users is imposed by BuenaVista scalability limits.

Protocols and bandwidth

The session setup and message passing protocol used internally by TANGO Interactive™ uses TCP as transport layer. Data transfer from the HTTP servers also uses TCP.

Bandwidth requirements for the TANGO Interactive™ protocol are negligible. The system has been successfully tested over wireless

9.6 kbps links, as well as over the marginal networks with packet loss as high as 60%. No problems other than slow response were observed.

Bandwidth requirements for the data transfer from HTTP servers are strictly application-dependent. Our courseware management system offers multiple versions of the class visuals, from large images to simplified HTML. It is possible to adjust the quality of visuals to network conditions to improve system performance. Another solution for slow networks is installation of a proxy server for HTTP protocol (see [Classroom Behind a Firewall](#)). This solution will ensure smooth operation of the TANGO Interactive™ framework even on extremely bad networks.

Audio and video over IP is somewhat more demanding. BuenaVista uses TCP for internal communication and UDP (more precisely, RTP) for media transport. Although BuenaVista can produce high-quality, high bandwidth data streams on LANs, it is also equipped to deliver audio and video over slow connections. The GSM audio nominally requires only 13 kbps per audio stream. Since BuenaVista suppresses transmission of silence (if properly set up), the effective audio stream for a typical class requires approximately 10 kbps. A useful video stream using our H.263 encoder can be tuned to approximately 12 kbps. *Video stream transmission is always optional.*

Does it mean that TANGO Interactive™ VC users can connect to classes via modems? The answer is yes, if they have a 56 kbps modem and if their ISP provides decent service. Keep in mind that a 56 kbps modem bandwidth is a myth. Typically, a "good" connection to an ISP will provide 41-44 kbps downstream, upstream bandwidth is always limited to 33.6 kbps. Assuming the best possible connection, it is usually possible to receive TANGO Interactive™ VC classes with audio only. Video reception on such links is not practical. The minimal home link bandwidth necessary to receive audio, video, visuals, and to participate in other TANGO Interactive™ VC sessions without interruptions of the audio stream is approximately equal to 80 kbps and hence requires an ISDN link or better (cable modem, ADSL, etc.).

Since the default transmission method used by BuenaVista is unicast, the outgoing link from the teacher site must support N times the average audio and video bandwidth for audio & video streams alone. In practice, the instructor site requires at least T1 connectivity.

Configuration of the Workstations

All recommendations in this section are supposed to provide the reader with reliable information on how to configure the system to obtain excellent performance. **These are not "minimal system requirements."**

Tip!

For minimal system requirements take a look at the box on the left hand side of the [TANGO Interactive™ Virtual Classroom Overview](#) page.

TIVC will run on much weaker workstations, including such outdated platforms as 166 MHz Pentium machines, but the performances will not be acceptable, except perhaps for a system demonstration. All configurations listed below ensure excellent system performance, necessary for delivery of high-quality ISDL classes.

Audio & video workstation

Platforms: At present, TANGO Interactive™ Virtual Classroom supports two platforms as audio & video workstations. One option is a PC with MS Windows NT operation system. Hardware specification for a PC based audio video platform is provided in the table below.

Another platform is an SGI O2 workstation running IRIX 6.5. A standard, vendor recommended, configuration of this platform is adequate for TIVC applications. The O2 workstations come, per default, with a camera, an A&V module, and 128 MB of RAM.

We discourage use of MS Windows 95 or 98 as the operating system for TANGO Interactive™ Virtual Classroom. The system will work, but its stability may be less than perfect. This is due to the poor performance of the Windows 95 and 98 under heavy load, caused by weak process management in the non-preemptive environment.

Operating System	MS Windows 2000, MS Windows NT (preferred) MS Windows 98 (acceptable)
CPU	Pentium III 450 MHz (or faster)
Memory	128 MB
Hard Drive	No special requirements
Graphics Board	Capable of displaying 65536 or more colors at 1152x864
Projector	Capable of displaying the same resolution at the graphics board
Network	Ethernet adapter 10/100 Mbps
Video Capture Board	One of: <ul style="list-style-type: none"> • Winnov PCI Videum, • Panasonic EggCam, • Cards based on Brooktree chipset (see Appendix at the end of this document). <p><i>Deprecated devices:</i></p> <ul style="list-style-type: none"> • <i>Parallel port cameras,</i> • <i>USB port cameras (if MS Windows NT used).</i>
Audio Board	Unnecessary if Winnov Videum used. Otherwise, Soundblaster 64 or compatible, <i>with full-duplex device driver.</i>
Microphone	Integrated with a headset, such as VXI Parrott 10-3 . <i>Use of any camera built-in microphone is unacceptable.</i>
Audio output	Amplifier or amplified speakers.
Echo canceller	Optional. Biamp System's VC 3000 Voicecrafter

Backup: In the classroom setup, two workstations with this configuration are recommended, one of them acting as a backup. A mix of an IRIX and a PC workstation is acceptable.

Video capture cards: Current market for the PC video capture cards is very crowded. All products are touted as “high-quality” videoconferencing devices. Unfortunately, the reality is very different. Only a handful of products is useful for professional applications. We strongly recommend that only certified solutions be used with TANGO Interactive™ Virtual Classroom. BuenaVista may or may not work with other cards and cameras. We cannot support or troubleshoot unstable and incompatible drivers coming with audio or video capture solutions not on our certified devices list. Use of such devices is entirely on the end-user risk.

Microphones: For good audio quality, special attention must be paid to the choice of microphones. Many cameras, even the certified ones, come with a built-in microphone. Built-in

microphones are unsuitable for multiparty conferencing. They are overly sensitive, omni-directional devices. They pick up and transmit ambient and background noise, dramatically lowering signal-to-noise ratio. Digitization of a poor quality audio signal tends to make things worse. One open camera mike in the conference can effectively disturb all conference participants. The only acceptable solution is a good quality, unidirectional, head-worn microphone, preferably equipped with ambient noise suppression via a built-in Digital Signal Processing (DSP) chip.

Headphones vs. speakers: For individual student workstations, and for the "dispersed" users receiving individual audio streams, no speakers should be used. Instead, the audio should be received via headphones. The best choice is a set of headphones with an integrated microphone.

Echo cancellation: A useful addition to the virtual classroom setup is an echo canceller. In the absence of such a device, the audio played by speakers can be picked up by the microphone and transmitted back to the sender. Since Voice-over-IP applications always have noticeable delay, related to various buffers in the system, the "echo" transmission is extremely annoying for the audio sender and disturbing to other receivers. If the audio stream is played back in the classroom using speakers, the classroom microphone must be turned off when not in use. This somewhat lowers the interactivity. An echo canceller (listed in the table above) helps to alleviate this problem. With an echo canceller installed, it is possible for each student to use the classroom microphone at any time to ask a question.

Audio equipment quality: Since good audio quality is critical for reception of the class, it's worth investing in better class audio components than the ones usually offered with computers.

Classroom video: With just one audio & video workstation in the classroom, instructor's video is shown via a projector. Many computer screen projectors only handle screen resolutions up to 800x600. The projector used in the virtual classroom should match the resolution of the audio & video workstation. Since this workstation can also be used to show certain specialized visuals, the suggested screen resolution is 1152x864 in Hi Color (65536 colors) or True Color (16.7 million colors) mode.

According to our experience, a video link showing the instructor greatly enhances class reception and is always praised by students.

Student's workstations

Operating System	MS Windows 2000, MS Windows NT (preferred) MS Windows 98 (acceptable)
CPU	Pentium II 400 MHz (or faster)
Memory	128 MB
Hard Drive	No special requirements
Graphics Board	Capable of displaying 65536 or more colors at 1152x864
Projector	Capable of displaying the same resolution at the graphics board
Network	Ethernet adapter 10/100 Mbps
Video Capture Board	One of: <ul style="list-style-type: none"> • Winnov PCI Videum, • Panasonic EggCam, • Cards based on Brooktree chipset (see Appendix at the end of this document). <p><i>Deprecated devices:</i></p> <ul style="list-style-type: none"> • <i>Parallel port cameras,</i> • <i>USB port cameras (if MS Windows NT used).</i>
Audio Board	Unnecessary if Winnov Videum used. Otherwise, Soundblaster 64 or compatible, <i>with full-duplex device driver.</i>
Microphone	Integrated with a headset, such as VXI Parrott 10-3 . <i>Use of any camera built-in microphone is unacceptable.</i>
Audio output	Headsets, such as VXI Parrott 10-3. <i>Speakers shouldn't be used.</i>

The table above lists the recommended configuration of a student's workstation in the Virtual Classroom. The preferred platform is PC-based.

Other supported platforms include Linux, IRIX 6.5, and Solaris 2.6. Note however that at present there is no audio & video support for Linux and Solaris. In general, system performance on UNIX machines isn't as good as on MS Windows.

For "dispersed students," we recommend identical configuration, with two exceptions:

- A Pentium III based machine is preferable if the student wishes to receive both audio and video. In general, workstations on slower network links should have faster

- CPUs, to make up for the more complex audio & video codecs used for higher media compression.
- A dial-up adapter or cable modem replaces the Ethernet card. For ISP link recommendation, consult "[Protocols and bandwidth](#)" section earlier in this document.

Instructor's workstation

Operating System	MS Windows 2000, MS Windows NT
CPU	Dual Pentium III 450 MHz (or faster)
Memory	256 MB
Hard Drive	No special requirements
Graphics Board	Capable of displaying 65536 or more colors at 1280x1024
Projector	Capable of displaying the same resolution at the graphics board
Network	Ethernet adapter 10/100 Mbps
Video Capture Board	<p>One of:</p> <ul style="list-style-type: none"> • Winnov PCI Videum, • Panasonic EggCam, • Cards based on Brooktree chipset (see Appendix at the end of this document). <p><i>Deprecated devices:</i></p> <ul style="list-style-type: none"> • <i>Parallel port cameras,</i> • <i>USB port cameras (if MS Windows NT used).</i>
Audio Board	Unnecessary if Winnov Videum used. Otherwise, Soundblaster 64 or compatible, <i>with full-duplex device driver.</i>
Microphone	Integrated with a headset, such as VXI Parrott 10-3 . <i>Use of any camera built-in microphone is unacceptable.</i>
Audio output	Headset, such as VXI Parrott 10-3
Echo canceller	Optional. Biamp System's VC 3000 Voicecrafter

Instructor's workstation is a critical part of the system. This workstation consolidates features needed in the audio & video and student machines. Further, the instructor often needs to concurrently use a variety of application modules that allow system and class monitoring. This workstation must handle a dozen or more processes on top of the usual complement of system applications. High performance of this workstation is the key to high performance of the entire TANGO Interactive™ VC system.

We strongly suggest that instructor's workstation run Windows NT or Windows 2000 operating system. Using MS Windows 98 or 95

may result in poor system performance, instability of the operating system and, in the worst case, system crash.

Due to the very high system load, we recommend that instructor workstation is a two-CPU machine. This configuration allows the operating system to allocate very CPU intensive video and audio encoding and transmission to one processor, leaving ample computing power to handle Java VM running the TANGO Interactive™ VC kernel.

Tip!

If a dual CPU machine is not available, we recommend that the instructor is provided with two computers: one handling audio & video, and another running the rest of the TIVC applications.

Another critical aspect of the instructor's machine is the high resolution display unit. We recommend use of the 21" monitors with highest possible resolution: 1280x1024, or 1600x1200 if available. Such resolutions allow the instructor to neatly arrange all needed software tools on the screen in one always visible layer.

All issues regarding audio gear quality discussed in the section "Audio & video Workstation" apply to the instructor's workstation as well. Echo canceller should be given serious consideration.

Software requirements

Each workstation participating in TANGO Interactive™ VC system should have the following software installed:

- MS Windows 2000 Professional; or MS Windows NT 4.0 with Service Pack 5 or newer; or MS Windows 98/95 (discouraged),
- Netscape Communicator 4.7 (version 4.6 is not supported as it has some serious implementation problems) or
- MS Internet Explorer 5 with a recent MS Java VM (available from <http://windowsupdate.microsoft.com>).

Some screen and application sharing features of TANGO Interactive™ Virtual Classroom require a VNC Server (available for free from [AT&T](#)) and MS NetMeeting. Refer to the software documentation for details.

Support issues

Delivery of ISDL classes requires technical support. The practice has proven that in the Virtual Classroom configuration there should be one technical support person per remote classroom location. The Instructor site should have a teaching assistant on-line. Service provider should expect to have a technical person per each set of three concurrent ISDL classes.

Tip!

TIVC provides out-of-band communication channels for support personnel. Support personnel workstations should conform to the specification of the student workstation.

TI VC Service Hosting

TI VC service can be hosted by a provider or by the end-user organization. Service provider site must provide two server machines: one for the TANGO Interactive™ Engine, and another for the Web server hosting TANGO Interactive™ VC downloadable software.

For small sites, these services can be collocated on one machine. For sites supporting multiple concurrent classrooms and extensive range of courseware, we recommend that these services run on separate servers.

TANGO Interactive™ Engine can run on MS Windows NT or UNIX workstations. The Engine requires Java runtime, version 1.1 or 1.2 (Java 2 platform). The CPU requirements imposed by the engine are minimal. It is possible to run a cluster of up to 32 server instances on one workstation.

An HTTP server is required as a TANGO Interactive™ VC application server. TANGO Interactive™ VC imposes no special requirements on the server. If the system is used in conjunction with [WebWisdom 2000](#) Courseware Management System and application server, our recommended choice for the HTTP server is Apache for both UNIX and MS Windows NT machines.

Appendix A

The following video capture cards use Brooktree/Rockwell Bt848 chipset:

- [Intel Smart Video Recorder III](#),
- Panasonic [Eggcam Kit](#) (includes a camera),
- Hauppauge Computer Works—[Wincast/TV](#),
- Diamond Multimedia—[SupraVideo Kit](#),
- miro—miroMediaTV,
- 3Com—Big Picture Kit (includes a camera & modem),
- [Zoom Telephonics](#)—VideoCam (includes video capture & modem board and camera)
- [Boca Research](#)—VideoCommunication Suite (includes video capture & modem board and camera).

It is usually possible to use the original Intel Smart Video Recorder III device driver with these cards.