VIRTUE
COMMUNICATIONS

Web Conferencing
Technical White Paper

Published: April 2004

Table of Contents

1. Architecture Overview

Virtue Web Conferencing (VWC) offers a multi-tier web application with a sophisticated server component and a browser-based thin client. The system was designed for easy integration with other products and works together with all usual standard components of a typical IT infrastructure.

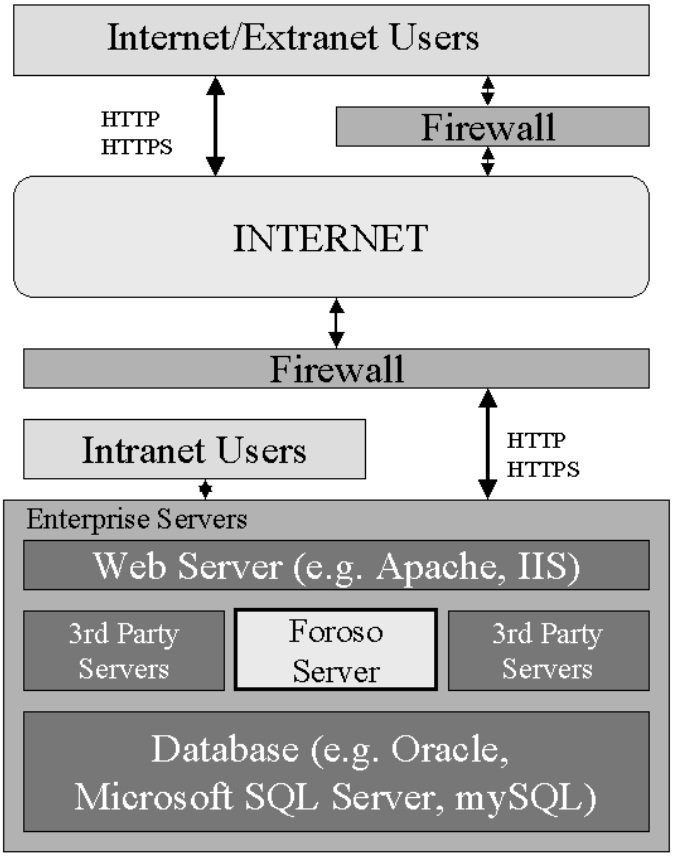

Client Platforms:
Windows 95, 98, 2000, NT, XP
Mac, Linux, Solaris, OS/2
IE 5.0+, NS 4.7+

Required Bandwidth:
28.8 modem for basic features
33.6 modem for VoIP
ISDN for application sharing

Server Platforms:
Windows 2000, NT, XP
Linux
Solaris

Required Bandwidth:
see chapter „Scalability“

Required Environment:
Web Server
SQL Database
PHP

The VWC server is a component of the enterprise server environment. The underlying database technology is an SQL server, such as Oracle, Microsoft SQL Server, or my SQL. The service is provided through standard web servers, such as Apache or IIS. VWC co-exists and integrates with other complementary server products.

VWC does not require installation of any client software beyond a standard browser with Java support. This eliminates a lot of potential problems that come with installation, version updates, and varying configurations at the end users computer. Virtue designed its product in a way that minimizes browser requirements. Thus, a wide range of operating systems and browser versions can be supported.

As for network protocols, VWC relies only on http and https. There is no need to open up any specific ports, to configure firewalls or proxy servers. The server automatically detects the best way to connect to the client and reacts in an appropriate way.

2. Feature Explanations

2.1 Basic Features

For this document, we summarize all features except VoIP and application sharing under the term “basic features”. This includes the general synchronization of the environment, as well as interactive features such as text chat, up-to-date participant list, various annotation features, and others.

The main concept of VWC’s architecture is a central server that keeps all connected clients up-to-date. Content and all other information get uploaded to this server prior or during the meeting and downloaded during the meeting.

2.2 Voice over IP

VWC follows the basic concept for VoIP, that there is always one active speaker and all others can listen. Active speakers can be switched very quickly, thus ensuring a pleasant communication process. An additional advantage of this concept is that it is always very clear who is currently speaking.

VWC’s concept is not optimized for concurrent speaking, especially of many participants. This is partly due to our belief, that the “one speaker at a time” model is more efficient than the model of many concurrent speakers. Since VoIP always has a certain network delay due to unavoidable latency problems, concurrent speaking with VoIP is even more confusing than on phone conferences.

Virtue’s priorities are on very high quality and robustness. Robustness relates to low bandwidth consumption, as bandwidth bottlenecks are the most typical sources for problems. Another main requirement was that VoIP must not require an explicit client software installation, as almost all similar products need it.

2.3 Application Sharing

VWC offers the capability to share an application that runs on the computer of one participant with all other users. This doesn’t even require installation of this application at the computers of the participants who share the application with the application owner. Thus, for example a Mac user could watch the application of a Windows application owner.

Application sharing is a different way to load new content than downloading shared content from the central server. Generally, the fastest option is to share just a brief URL. Then, each participant downloads just this URL and each participant’s browser downloads the respective content directly from the respective content server. The second fastest option is to upload content (e.g. Powerpoint, Excel or images) to the server and enable all participants to download it from there.

These options are ideal if the presenter just wants to display content to all participants, without modifying the content. However, if the content is a shared URL, each participant can follow links in the content at his own discretion, until the presenter synchronizes all again. Similar, if the shared content is an Excel file, each participant can modify the content in his local copy.

With application sharing, the presenter can ensure that each single activity, including mouse movements and menu selections, are immediately visible to all participants.

The basic approach for VWC’s application sharing is to “take photos” of the application screen of the application owner and transfer them to all participants for display at their screens. As the typical data volume of such a “photo” is around 1 MB, the frequency of updates is massively dependent on the available bandwidth. With an ISDN line, an update would be possible roughly every 1-2 minutes. Obviously, this would not be an acceptable solution for most application sharing scenarios.

Thus, compression of these “photos” is of very high importance. VWC achieves an average compression rate of over 98 %. This means, that an ISDN user gets an update about every 1-2 seconds, which is absolutely adequate for almost all scenarios (with the most notable exception of a video stream).

3. Bandwidth Requirements

3.1 Client Bandwidth Requirements

Users participate in sessions through their standard Web browsers. The central server distributes all information to all participants and receives certain information from participants. Different features have different load patterns that will now be described in more detail.

3.1.1 Basic Features without VoIP and without Application Sharing

There is a certain minimum traffic that happens continuously between the user PC and the central server. This ensures that all participants are synchronized, receive and send chat messages, see and create annotations, get up-to-date information about who is in the meeting room, as well as support for other basic features. The resulting traffic is so little, that it is irrelevant.

The only noteworthy exception is the change to a new content page. Then, the PC of each participant needs to automatically download this new content. Depending on the size of the content, this can take just a moment or several seconds. If the content is big, such as a high-resolution image or even a video, then it can take quite a while. The behavior is identical to the behavior of a plain browser. If an image takes five seconds to download from a web site, then it also takes five seconds to download into the meeting room.

The resulting pattern is a very low basic load and peaks whenever new content gets loaded.

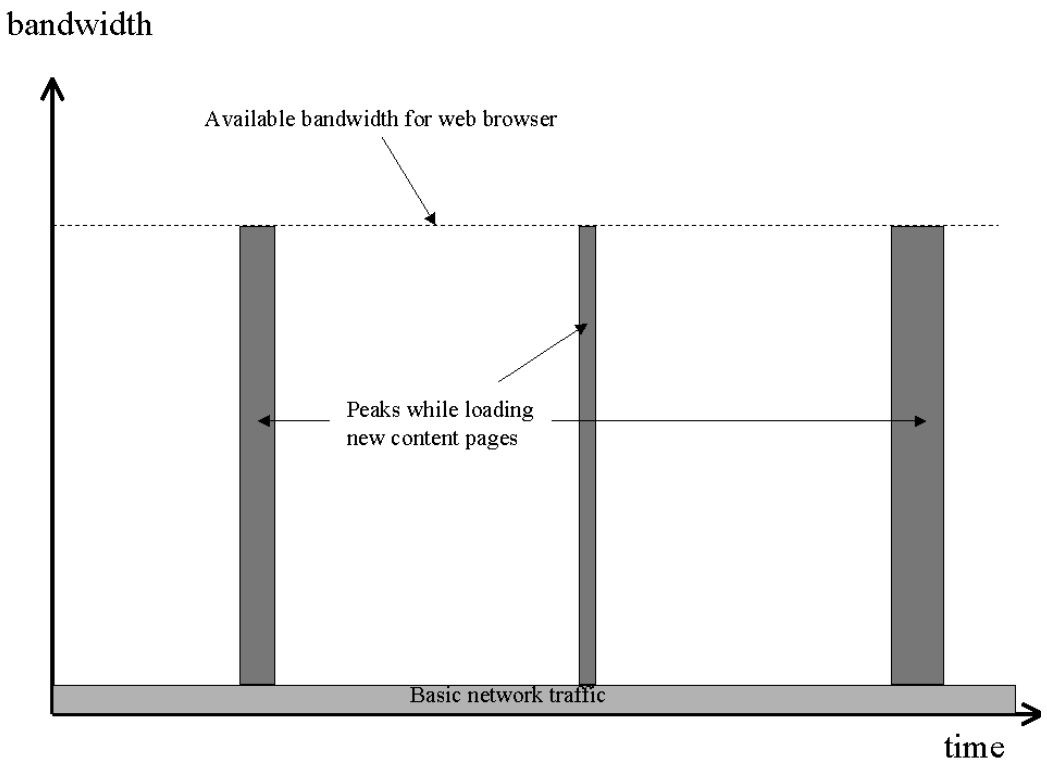

The above pattern is typical for an ISDN line and content such as Powerpoint slides. Most of the time, the ISDN line will not really be busy. If a participants has just a slow modem connection, the peak times increase in duration. If there is more bandwidth, such as DSL, the peaks are very short and often not even notable to the user.

3.1.2 Voice over IP

As outlined above, VWC follows the basic concept for VoIP, that there is always one active speaker and all others can listen.

Thus, with activated VoIP, each participant either receives the voice stream from the active speaker or sends a voice stream if he is the active speaker. If nobody speaks (even during relatively short breaks), then there is no voice stream traffic at all. The traffic from VoIP is on top of the basic traffic. Content pages are loaded as before and cause the same peaks. If combined with VoIP, there is less bandwidth available than without VoIP and it takes longer until new content gets loaded. Therefore it is recommended that the speaker pauses a moment until the new content is visible to everyone.

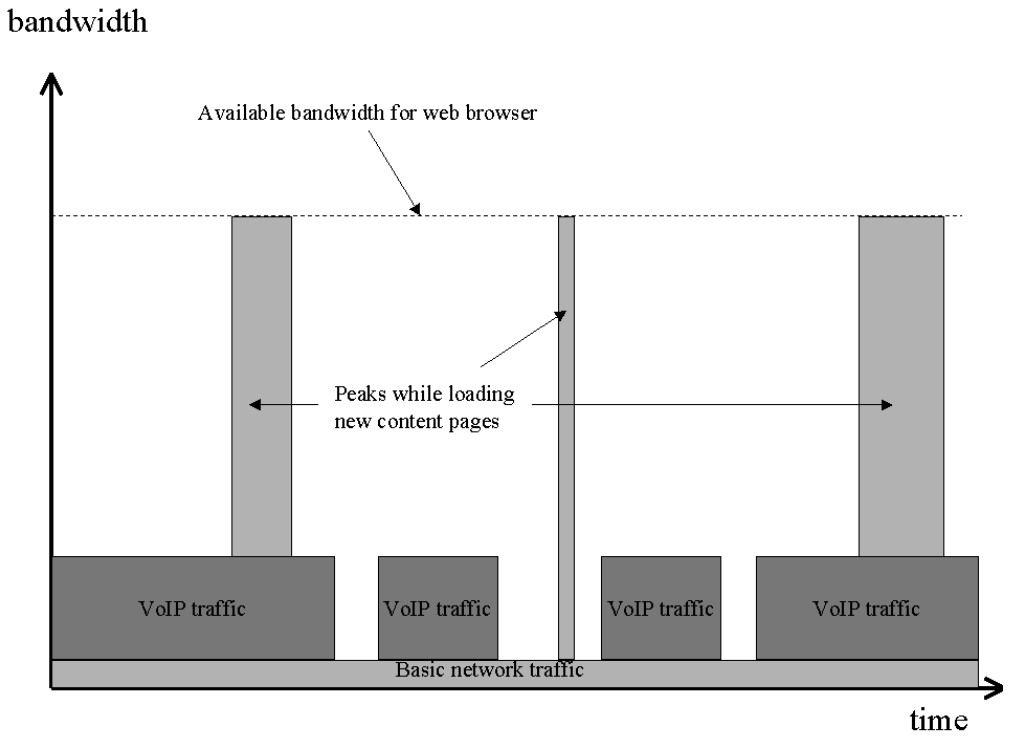

The above pattern is typical for an ISDN connection. VoIP requires only a fraction of a full ISDN and leaves enough bandwidth for content loading. With modems, content loading takes longer if done in parallel with speaking. And if the available bandwidth gets below the capacity of a 33.6 modem, VoIP may be delayed or even some words may be dropped.

With a DSL connection, VoIP is using up only a very small fraction and leaves plenty of bandwidth for very quick content loading.

3.1.3 Application Sharing

When application sharing is in use, the presenter selects an application on his computer and makes it available to all participants. The respective pattern is similar to downloading new content pages. The main difference is that application sharing typically means continuous updates to the content, as the presenter moves his mouse, selects menu entries and objects, changes content by filling out forms or typing text. Thus, this modified content has to be distributed continuously to all participants.

VWC does various optimizations in order to minimize network traffic. This includes compression techniques as well as smart algorithms that determine just the changes and thus further reduce the transferred data volume. Thus, the pattern is characterized by a sequence of full updates and delta updates, on top of the basic traffic and VoIP traffic. There are no peaks resulting from content page loading, as this is not relevant during application sharing.

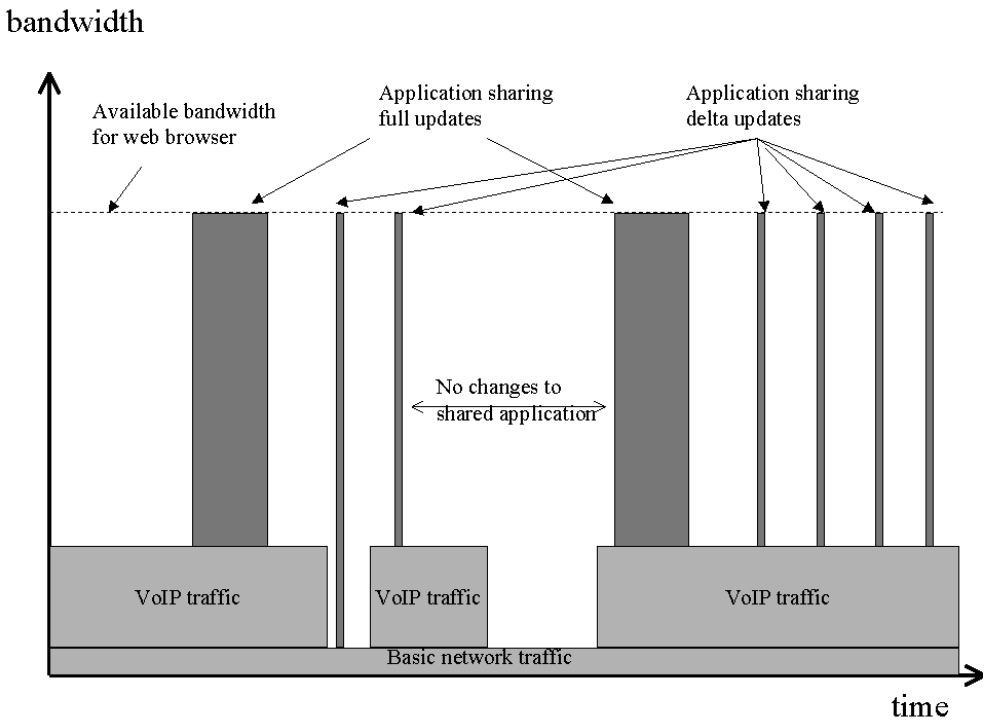

The above image shows a typical pattern where application sharing and VoIP are used in parallel on a DSL connection. Only the initial start of the application and significant changes to the application display max out the bandwidth, especially if someone speaks in parallel. But most of the time, there is plenty of free bandwidth available. Below is the same scenario with an ISDN connection. Then, most of the bandwidth is occupied and large changes might take more than just a few seconds. If the connection bandwidth drops further, e.g. to the capacity of a 33.6 modem, then application sharing will not follow all changes anymore and VoIP might get troubled.

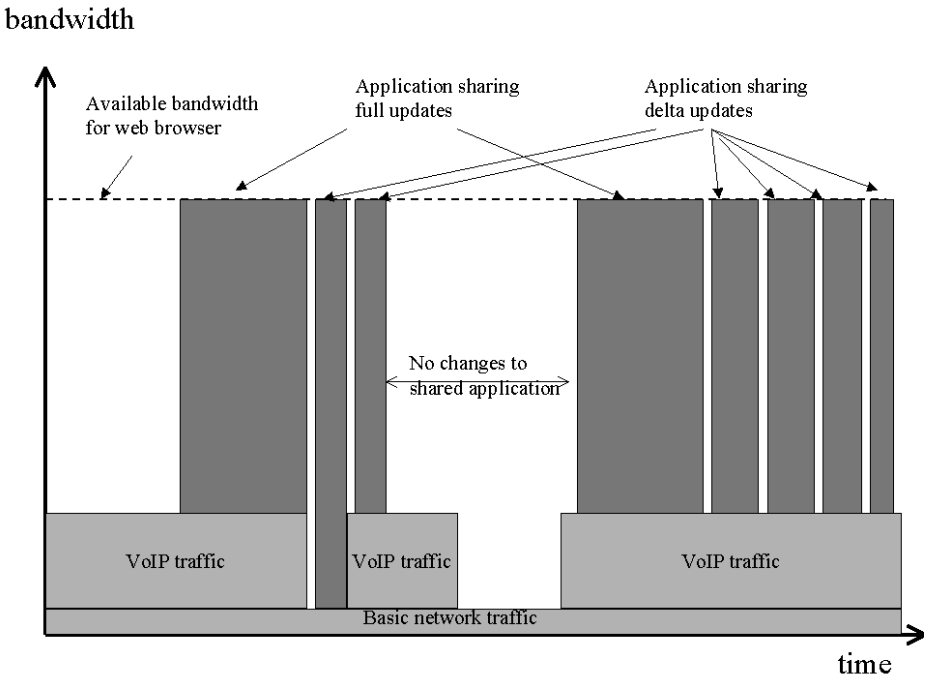

3.2 Server Bandwidth Requirements

The patterns on the server are basically identical with client patterns. The only difference is the overall bandwidth that is required. While each participant has one connection to the central server, the server has one connection with each single participant.

The three main factors for appropriately sizing the server environment are

- the number of concurrent users (i.e. are there just a few or hundreds of concurrent users)
- user behavior (i.e. how actively do they collaborate)
- content type and size (i.e. is it just simple data like text and images or multimedia streams)

The number of concurrent users can typically be determined early on and should be predictable in most usage scenarios.

User behavior is also quite predictable. A session that lasts 30 minutes will most likely use 8 to 30 pages of content where participants spend between one and five minutes on each page.

The most typical content comes from Powerpoint presentations, office documents, web editors, and images. The size of these data types is usually between 20 and 100 KB. Large high resolution images should be avoided unless the quality is essential for the collaboration task. Audio can easily exceed some MBs and video even tens or hundreds of MBs. Clearly, this has major implications on bandwidth and CPU requirements.

3.2.1 Server bandwidth requirements for data conferencing

During the early stage of a typical usage patterns, this doesn't have to be much of a concern. A typical backend environment with server hardware such as a SUN UltraSPARC or an Intel Pentium III can easily support one hundred concurrent users. Only if you have really high load in excess of one hundred concurrent users, you have to plan for more server hardware.

Before we go into more detailed bandwidth calculations, we summarize the typical bandwidth packages that are available to corporations. Please note that the scale is not linear. For instance, a T3 has almost 7 times more capacity than a T2.

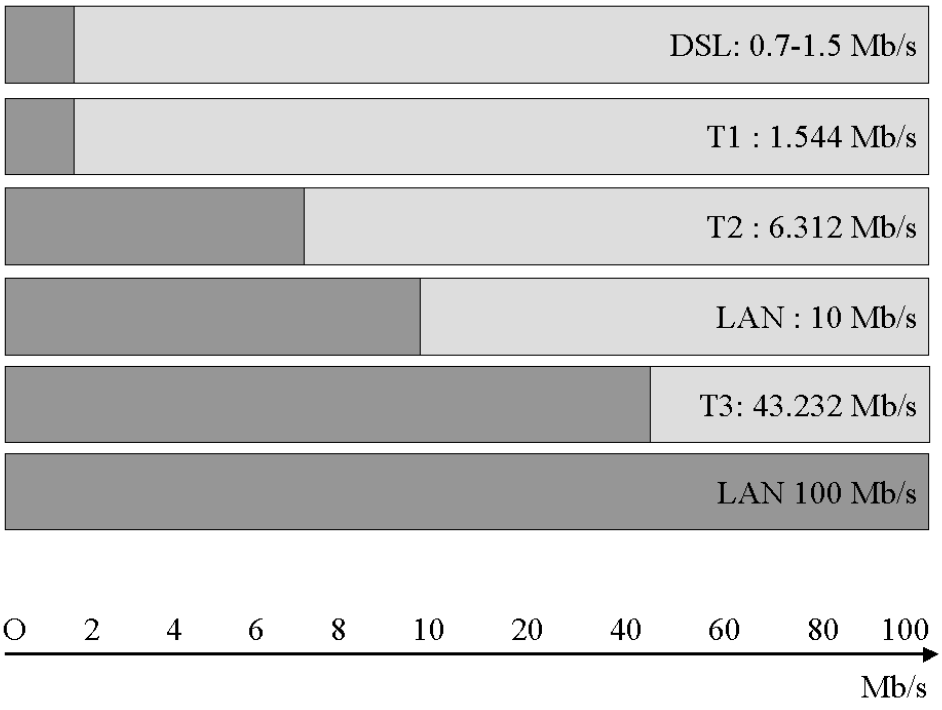

Next, we want to summarize the typical data volumes that are used in collaborative conferencing applications. The most typical formats are:

- Powerpoint slide: 30-60 KB
- Word page: 25-40 KB
- Excel sheet: 20-60 KB
- PDF page: 50-100 KB
- Image: 20-100 KB
- Web page: 20-80 KB
- Survey page: 10-20 KB
- Agenda page: 20 KB

Most content pages are between 20 and 100 KB in size and an average of 50 KB per page is probably a conservative assumption.

If we now make some assumptions about how often a new page gets displayed and how actively the participants interact, then we can calculate the average value for the system and determine, how much server bandwidth is needed for how many concurrent users. The result for Virtue's data conferencing system is an average of 12 kb/s and per user. However, more important that these mathematical results are the typical usage patterns that occur when the system is in real use. The most important factor are the peaks that happen when new pages are displayed. Then, all users have to wait for the new page and the server has to distribute it as quickly as possible. This is the most challenging situation and thus we want to analyze it in more detail.

We want to analyze five different load scenarios to come to useful recommendations for server bandwidth.

1) 1S10: One session with ten concurrent users
2) 1S100: One session with hundred concurrent users
3) 10S10: Ten sessions with ten concurrent users in each
4) 1S1000: One session with thousand concurrent users
5) 100S10: Hundred sessions with ten concurrent users in each

1S10: 1 session with 10 participants

10S10: 10 sessions with 10 participants in each

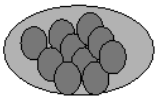

1S100: 1 session with 100 participants

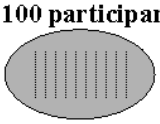

100S10: 100 sessions with 10 participants in each

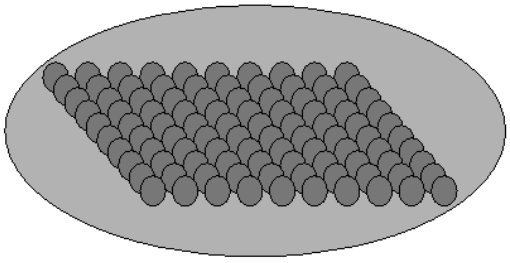

1S1000: 1 session with 1000 participants

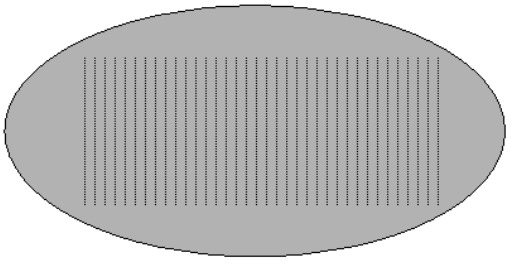

Another important assumption is the maximum waiting time that is tolerated by participants. We'll use ten seconds as our assumption here for the maximum time that a new page display must not exceed. Please note that participants will still see the current page and thus usually not notice when the ten seconds start. Thus, even a longer time might be quite acceptable. If instead of ten seconds only half the time or twice the time should be used as maximum, then all numbers have to be adjusted accordingly.

1S10: A new page with 50 KB gets displayed to all 10 participants. That means 500KB (or 4.000 Kb) to be downloaded. If the download has to be completed in 10 seconds, then 400Kb/s have to be provided from the server. A T1 line could deal with that volume in less than 3 seconds. Or it could handle about 40 participants in 10 seconds.

1S100: The same scenario for 100 concurrent users means that 4.000 kb/s have to be provided by the server. A T1 line would require almost 30 seconds for the download. A T2 line would only need less than 7 seconds.

10S10: The big difference here is that not all 100 online users would share the same peak. Instead, each group would have their own peaks as with 1S10. Two or at most three peaks might overlap, but more would be a rare exception. Thus, the 10S10 scenarios can be calculated by tripling the 1S10 calculation. That means that 1.200 Kb/s have to be provided for the maximum peaks. This means that a T1 link can serve all 100 users in max. 10 seconds.

1S1000: This is an extreme situation where 1,000 users all participate in one single huge session. This could only be a large-scale information presentation session (e.g. a company meeting or a public presentation). Here, 40,000 kb/s have to be provided by the server. This requires a T3 line in order to meet the 10 second maximum requirement.

100S10: Here we have again 1,000 users, but now with 100 groups with 10 users in each. This is the basic 1S10 scenario again, but with a lot of peak overlap. If we assume that on average there are 200 seconds between new page displays, then we come to 5 parallel peak periods on average. But since the peak periods will not be distributed so evenly, we conservatively assume 15 parallel peak sessions as the worst case scenario. This means that 6,000 kb/s have to be served, which almost uses up a T2 line.

3.2.2 Adding VoIP

VoIP with VWC adds another 14 kb/s for each user. However, this is a steady stream and there are no peaks to consider. If we add the additional 14 kb/s to each scenario, then we get:

1S10: Instead of 400 kb/s, we now have 540 kb/s. A T1 line is quite enough and could almost handle three times as many users.

1S100: Here we get to 5,400 kB/s, which means that a T2 is still easily sufficient.

10S10: For this scenario, we get to 2,600 kb/s. This requires 1,7 T1 lines.

1S1000: The result here is 54,000 kB/s. This would require either 1,25 T3 lines or a good LAN. Keep in mind: this means 1,000 people listening to voice in real-time.

100S10: This scenario adds up to 20,000 kb/s, which can be served by 3,2 T2 lines or 0,5 T3 lines.

3.2.3 Adding application sharing (plus VoIP)

There are a variety of details in application sharing that impact the bandwidth usage patterns. For example, if there are no activities in the shared application, then there is no network traffic. This happens a lot, as the presenter will show something and then explain to the participants what it means. If there are small changes to the screen (e.g. filling in a cell in Excel or opening a menu), then VWC does not send a complete update screen, but only the incremental changes to the display. VWC also reduces the color depth, which is hardly a difference for most office applications, but might change the appearance of a high resolution image quite a bit. Thus, typical business and office applications are very well handled by VWC's application sharing, whereas applications that require a high color depth or a high update rate (like videos or games) are at a disadvantage.

On the server side, a combined peak per user of 44 Kb/s for application sharing and 14 Kb/s for VoIP has to be planned for. The basic traffic of changing pages is here not relevant, as application sharing replaces it completely.

1S10: We get to a maximum of 580 Kb/s. That means that a T1, that could deal with up to 30 users, is easily enough.

1S100 and 10S10: Here the result is a maximum of 5,800 kb/s and requires almost a T2 line.

1S1000 and 100S10: Here we get to 58,000 kb/s. That's about 1.35 T3 lines.

3.2.4 Summary Table

Scenario	Data only	Data + VoIP	Data + VoIP + App Sharing
1S10: 1 session with 10 users	0.4 Mb/s	0.54 Mb/s	0.58 Mb/s
1S100: 1 session with 100 users	4 Mb/s	5.4 Mb/s	5.8 Mb/s
10S10: 10 sessions with 10 users each	1.2 Mb/s	2.5 Mb/s	5.8 Mb/s
1S1000: 1 session with 1,000 users	40 Mb/s	54 Mb/s	58 Mb/s
100S10: 100 sessions with 10 users each	6 Mb/s	20 Mb/s	58 Mb/s

4. Hardware Requirements

4.1 Client Hardware Requirements

4.1.1 Client requirements for basic data conferencing

Any PC will be fine, if it is sufficient to work on the web in general. Data conferencing has no significantly higher requirements than plain browsing. A 28.8 modem will be OK for regular data conferencing.

4.1.2 Client requirements for VoIP and audio streaming

Obviously, you'll need a sound card and speaker if you want to hear something. And, if you want to speak with VoIP, you'll need a microphone or a headset. VoIP doesn't require the latest PC hardware, but it should be at least at the level of a typical Pentium II machine. As for connectivity, an IDSN connection is recommended both for speakers and listeners, although a steady connection through a 33.6 modem is sufficient.

4.1.3 Client requirements for application sharing

A Pentium II machine is here also sufficient. However, an ISDN connection is recommended for frequent updates of the shared application. Applications that require a consistent high update rate need more bandwidth then ISDN.

4.1.4 Client requirements for video streaming and video conferencing

Video conferencing will be described in a separate paragraph at the end of this document, as it is not a standard component of Virtue's offering. This is partly because video conferencing has much higher infrastructure requirements, i.e. very fast PCs, additional equipment, and much higher bandwidth. Video streaming is similar to audio streaming, but with a higher data transfer volume.

4.2 Server Hardware Requirements

The minimum hardware requirement for the server is a 500MHz PIII server with 256 MB RAM and a 1.5 Mb/s connection (T1).

For 100 meeting participants using VoIP there should be at least a 500MHz PIII server with 512 MB RAM and a 6 Mb/s connection (T2).

More information about issues related to server installation can be found in the Virtue Web Conferencing Installation Guide.

5. Security

VWC supports a number of mechanisms to ensure top-level security. Security can be categorized into network security, user security, and content security.

5.1 Network security

By enabling https as protocol, all content and interaction is using 128-bit SSL encryption for data transfer. This means that even if someone intercepts the transmission, all data is completely unreadable and secure. Typically, the server runs behind a firewall that protects all systems.

5.2 User security

Each user has a unique name and personal password. This information is stored on the central server and is sent only by encrypted transmission. Users can change their passwords. If users leave a session, they have to re-enter their password.

5.3 Content security

An effective roles and object access right system ensures that users can only access the information that they are supposed to. Sessions can be protected with additional passwords. If a user has no access right for certain information, then this information is not even displayed for him. All irrelevant information is filtered out at the server.

5.4 Control

All data access is logged at the server and all unauthorized access can be analyzed.

6. Accessibility

One of the most unique attributes of the VWC system is its excellent accessibility. Accessibility is a key success factor, as professional applications cannot afford to lock our users. And every technical problem negatively impacts end user acceptance. Not to mention that it causes additional support cost and wasted time.

6.1 No client installation

Experience shows that many users have difficulties to install software and to get it to work properly. Often the basic installation fails due to unexpected configurations or simply because the installation was not executed correctly by non-technical users. Additional problems arise when new product versions are released and different users run different versions that might cause incompatibility problems.

VWC requires no client installation at all. VoIP and active application sharing require that the browser has activated its in-built Java support. Everything is loaded from the server during the session.

6.2 No firewall problems

Many real-time products suffer under firewall problems. Often, these products try to resolve the situation by switching between various access options. In a way, these products try to trick the firewall and to get access through an open hole. But since there are so many firewall configurations and administrators continuously try to shut out access that is not using the approved mechanisms, it becomes a great race between real-time products and firewall configurations.

Virtue designed its product in a way that causes no firewall problems at all. All access uses only the standard http protocol. There is no need to open up ports or to configure proxies.

6.3 Low browser requirements

Many users configure their browsers in a restrictive manner. Thus, some browsers would not accept a plugin installation and some do not even allow for Java applets to be loaded.

VWC has very low browser requirements and thus works even in very restrictive settings.

6.4 Support for many operating systems and browser versions

In large corporations, employees often do not work with the latest browser versions. This might be due to the fact that an IT department has to approve the officially authorized software and thus an older version is the only officially available choice. Or users simply do not care for the latest version, because they feel they don't really need it or they are not even aware of its existence. Thus, there is a wide variety of operating system and browser version combinations. Unfortunately, these are often behaving quite differently, especially if a product is using the latest features of a new browser version.

VWC supports a wide variety of operating system and browser versions. All features are available for Windows systems with a Microsoft Internet Explorer 5.5 or higher. Other browsers or operating systems restrict the display of some document formats, like Word or Excel, and the usage of cards on top of URL and plugin content (e.g. Adobe Acrobat PDF files). Voice and applications can be recorded only on Windows systems, but they can be played on all systems.

The following operating systems and browser combinations have been successfully tested:

Windows	**Mac PowerPC**	**Linux**	**OS/2**
Internet Explorer 5.0	Internet Explorer 5.0	Netscape 4.7	Netscape 4.61
Internet Explorer 5.5	Internet Explorer 5.1	Netscape 6.2	
Internet Explorer 6.0	Netscape 4.		
Netscape 4.7	Netscape 4.7		
Netscape 6.2			
Netscape 7.1			

Internet Explorer 4.0, Netscape 3 and Opera browsers are not supported.

6.5 Low bandwidth requirements

Many end users still do not have broadband Internet access. And even intranet installations are often overloaded if many employees are active. Thus it is important that new applications carefully deal with bandwidth and do not generate huge peaks or continuous high load on the network.

VWC adds very little load beyond the shared content and minimizes the number of connections that are necessary. All data is routed through the central server. Thus, additional active users only add linear load instead of exponential increases.

6.6 Ease of use

Accessibility is not just a technical attribute, but is also determined by the simplicity of the product itself. Many powerful products are very demanding on users and even after extended training they only master a small fraction of all possibilities.

Virtue's product was designed for absolute ease of use. Users can participate in online sessions without any prior training and 15 minutes of introduction are enough to understand even all advanced features.

6.7 Video Conferencing

Note: Virtue's product does not include video conferencing capabilities. But it is possible to integrate third-party video conferencing products.

There are three main reasons that cause Virtue not to provide in-built video conferencing.

6.7.1 Video conferencing means greatly increased complexity

With today's state of the technology, there is no way to deliver video conferencing without client software installation, firewall and network problems, and significantly increased complexity for users at the interface. Additionally, video hardware (e.g. cameras) introduce a variety of operating system dependencies and setup problems.

While data and audio conferencing can be offered without any hassle and be used as easily as email or phone, video conferencing is much more complex.

6.7.2 Video conferencing means greatly increased bandwidth requirements

Video is somewhat complex to assess, as there are several choices to be made. Required bandwidth is massively influenced by the size of the video image, and the quality. Quality is determined by the compression rate (how good is the image) and the update rate (how many images are displayed per second). Typically, video conferencing requires between 200kb/s and 1,000 kb/s per user. Please not that this means neither full screen nor TV quality video.

The implication is that even small and not very fluent video requires significant bandwidth both at the users end and even more at the server. A T1 line can deal with a maximum of 5-7 users. Even a T2 line can support only 10-20 users. Large groups require a dedicated broadband network.

6.7.3 Video conferencing is much more specialized than data and audio conferencing

All business scenarios require data to be shared between participants. And there is always a need for communication. Thus, data and audio conferencing are mandatory requirements. Video conferencing, however, is only required in certain scenarios, where participants need to watch their counterparts continuously. Given the limitations of today's video conferencing, the effect of a live meeting can rarely be achieved when it comes to facial expressions and body language. Thus, it is often clearly better to use the screen for data display.

6.7.4 Video conferencing is an option

All this said, Virtue offers integration of third party video conferencing as an option. There are certainly scenarios where video conferencing is essential and Virtue can meet these needs.

7. Contact details

Virtue Communications
Parchetwiesen 33
D-82362 Wilheim / Muncih

Telephone : +49 (700) 3676 7600
Facsimile : +49 (700) 3676 7600
Email : info@foroso.com
Web: www.virtue-communications.com