ATI FirePro™ RG220

White Paper | Remote Graphics and the ATI FirePro™ RG220

Introduction

The ATI FirePro™ RG220 is a remote graphics solution that uses a technology that helps physically separate workstation users from their PCs while providing advanced graphics performance and the same work experience offered by discrete workstation graphics accelerators. This results in a work environment where users have a thin client system, displays, keyboard, and mouse on their desks, while the PC systems or servers containing the CPU, graphics solution, and data storage are located together in a remote location, such as a data center.

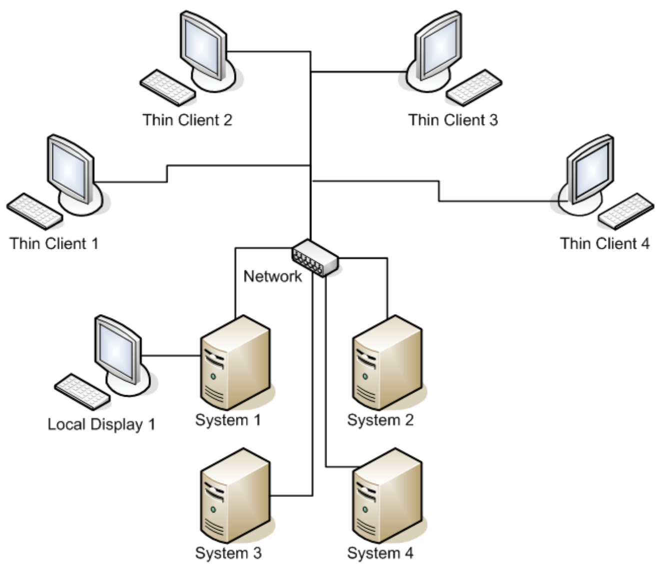

Figure 1: Typical Remote Graphics Setup

Note: This is not a virtualized system; users are connected to dedicated desktop PC or blade PC systems on a 1:1 basis and retain full discrete workstation graphics (and CPU) performance, subject to network bandwidth.

IT Has Challenges With Desktop Systems

Workstations that consist of individual desktop PCs may have several disadvantages from the perspective of a company and its IT department.

First and foremost, desktop systems can make it harder to safeguard high-value IP or data transactions. A PC on a desktop is vulnerable and accessible. With remote computing and graphics, on the other hand, users have all the power of a desktop PC but do not have physical access to the PC and its storage devices/outputs. Keeping the desktop systems in a central location also helps improve security against equipment theft, since they can be securely locked away at all times.

Desk-side support, such as patching and software updates, can be expensive and time-consuming for IT departments to provide, to say nothing of the productivity loss that can take place while a workstation is down. A central data center containing the PCs allows quicker and more convenient access for maintenance and troubleshooting, and also makes it easier for IT to migrate to new operating systems and install new applications.

Moving workstations (or deploying new ones) within a company can also be much simpler when the PC system itself never has to move or suffer the potential damage of moving.

All in all, remote computing and graphics offers a number of business solutions without requiring performance penalties.

The ATI FirePro™ RG220

The ATI FirePro™ RG220 is an integrated graphics accelerator plus display compression and IP transmission card. Its main function is to render all graphics data and then compress the dual display graphics data and output it out over a regular IP network to a remote thin client receiver system.

> ***Note:*** *"Local" here always refers to the host PC in which the graphics accelerator is installed, and "remote" refers to the thin client where the end user is located.*

The ATI FirePro RG220 uses the PC-over-IP® (PCoIP®) technology developed by Teradici® for compression and transmission and must be hooked up to a PCoIP thin client system that receives and decompresses the data. This thin client receiver system is sold separately. The ATI FirePro RG220 also compresses and transmits the keyboard and mouse data back and forth between the thin client and the host system. One of the key benefits of this technology is that the compression is offloaded from the CPU and has its own dedicated hardware. This allows for a much more robust remote PC experience.

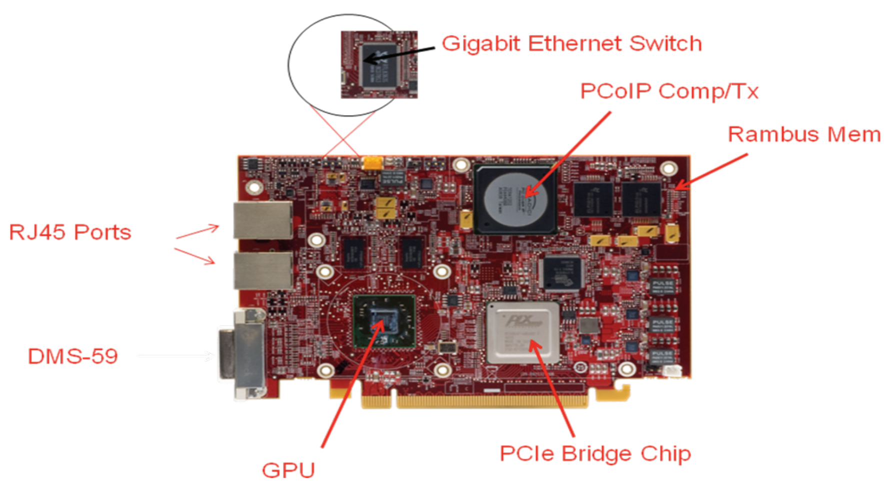

Figure 2: ATI FirePro™ RG220

Two ATI FirePro RG220 cards can be linked together to combine their display outputs for support for up to four displays, but this configuration must be connected to a quad thin client receiver box.

The ATI FirePro RG220 offers an integrated, single slot solution. There is only one device to maintain, stock, order, and control. The ATI FirePro RG220 uses less power in most cases than dual card solutions.

Each card has a unique MAC address. No special graphics drivers or special network equipment is required.

Output

The ATI FirePro™ RG220 has two RJ-45 connections: one **(1)** to receive input from another (primary) RG220 in a quad display configuration, and the other **(2)** to provide output to the thin client workstation (via the network) or another (secondary) RG220 in a quad display configuration. There is also a DMS-59 connection **(3)** for local dual DVI display output using an adapter cable.

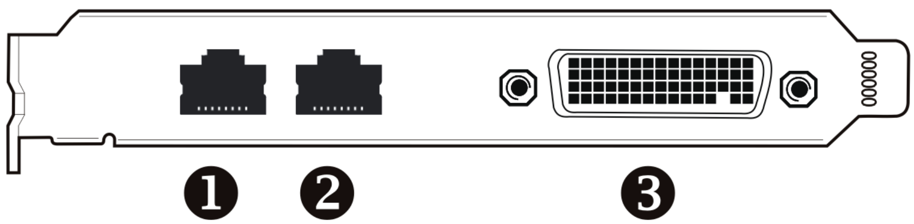

Figure 3: ATI FirePro RG220 Baseplate

Note: An ATI FirePro™ RG220 can act as either a primary or secondary graphics accelerator in a quad display configuration.

System Requirements

System Hardware	→ Intel® Pentium® 4/Xeon® or AMD Athlon™/AMD Opteron™ processor. → 512 MB of system memory; 1 GB or more recommended for best performance. → Motherboard with available x16 PCI Express® bus connection slot. → CD-ROM or DVD-ROM drive for installation software. → 450-watt or greater power supply recommended.
Operating System	→ 32- or 64-bit Microsoft® Windows® 7. → 32- or 64-bit Windows Vista®. → 32- or 64-bit Windows® XP Pro with Service Pack 1 (or higher).
Display	→ For local hosted output, only DVI displays are supported; VGA displays are not supported. → For remote thin clients, display output is determined by the output capabilities of the receiver box.

PC-over-IP Technology

PC-over-IP (PCoIP) is an innovative technology developed by Teradici that helps solve remote computing issues such as poor responsiveness, lack of media or graphics support, OS image management challenges, and peripheral interoperability. PC-over-IP technology enables remote graphics and desktop consolidation benefits to be realized.

TERADICI™

Figure 4: Teradici Logo

As with any desktop consolidation solution, PCoIP technology is dependent upon the network performance between the user and the centralized host computer. Unlike other solutions, however, PCoIP technology has been architected to optimize the transfer of user interface data and to enable the full-fidelity user experience that users have come to expect from traditional desktop PCs. Incorporated within this technology are algorithms that dynamically adapt the levels of compression to maximize the user experience. These and other advanced features help simplify the roll-out and management of PCoIP systems.

System Overview

PCoIP uses a two-chip solution that delivers the user's desktop over a standard IP network. The PCoIP Host Processor is used in the centralized host PC or workstation, while the PCoIP Portal Processor is used in a small, stateless desktop "portal" appliance.

The PCoIP Host Processor is tasked with encoding the complete desktop environment. The encoded desktop is then communicated in real time over an IP network (such as an enterprise network) to a PCoIP Portal Processor-enabled desktop.

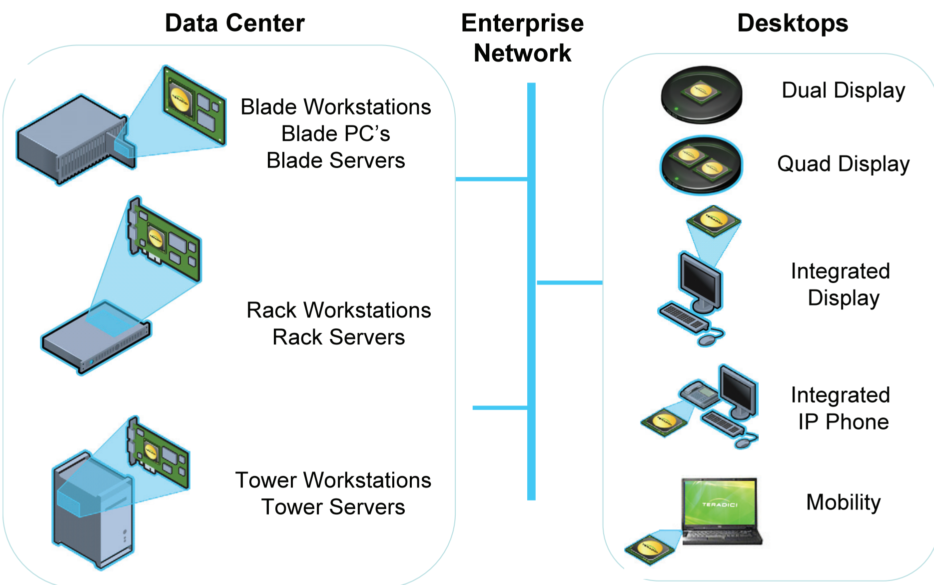

Figure 5: PCoIP System Options

PCoIP Host Processor

The PCoIP Host Processor receives standard Digital Visual Interface (DVI) display data from a PC's graphics processor. The PCoIP Host Processor also uses a PCIe® link for transparent bridging of HD audio devices and USB peripherals. Prior to network transmission, the PCoIP Host Processor compresses the video stream in real-time, optionally compresses the audio, and combines the bridged USB traffic. The combined data is AES encrypted at wire-speed and sent to the network using the on-chip integrated 10/100/1000 Ethernet MAC.

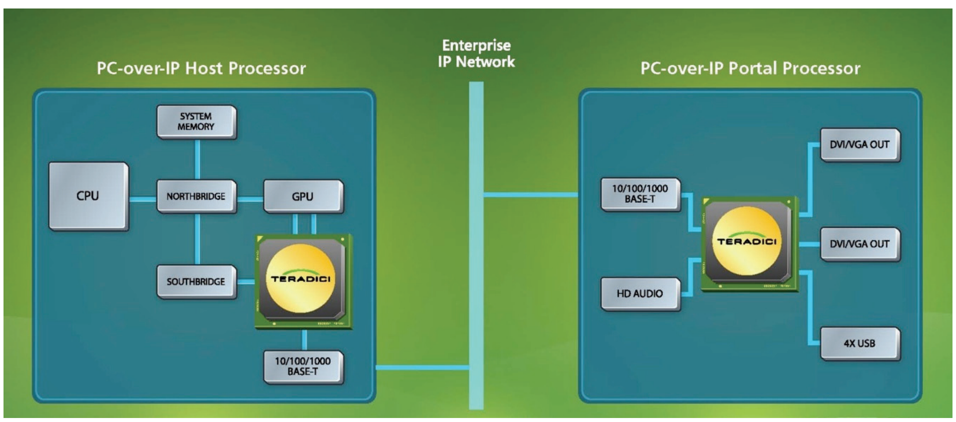

Figure 6: PCoIP Host and Portal Architecture

The PCoIP Host Processor employs sophisticated image compression techniques to help minimize desktop encoding system latency and reduce peak network bandwidth. The PCoIP image compression and transport algorithm is designed to provide an enhanced user experience at any available network bandwidth. The PCoIP Host Processor continuously monitors and adjusts the image compression rate to help ensure the optimal user experience for the available bandwidth. In all cases, static regions of the display image will be refined to be an exact replica of the image rendered on the host PC. This is a key requirement for visualization applications such as CAD or medical diagnostics.

The PCoIP Host Processor also adapts its traffic profile to fluctuations in available network resources – a key feature in large-scale PCoIP system deployments where high utilization and fair distribution of available network bandwidth is important.

The PCoIP Host Processor does not load the host PC in any way. The independent PCoIP Host Processor performs all the encoding and communication operations, freeing the host PC resources to perform the application and image rendering at full speed. Additionally, communication of the image is independent of the rendering, so unlike other remote computing protocols, the host PC and the application are not slowed down by a reduction in network bandwidth.

PCoIP Portal Processor

The Portal Processor is designed to be integrated into any user appliance, or portal, at the desktop. This appliance is called a "portal" because it is simply a secure window back to the centralized computer. The PCoIP Portal Processor decompresses and distributes the desktop-bound PCoIP packet stream to the video, audio, and USB interfaces. In the return path, the PCoIP Portal Processor combines audio and USB peripheral input data for transmission back to the PCoIP Host Processor.

PCoIP and the ATI FirePro™ RG220

ATI FirePro™ RG220 is fully compatible with any first or second generation Teradici-certified thin client solution.

Figure 7: PCoIP Certification Logo

It can be used with either a thin client desktop or thin client notebook, as well as either hardware or software Teradici clients.

Network Considerations

In a PCoIP system, only the display and peripheral data leaves the corporate datacenter to the local user network. This changes the resource profile required compared to the traditional model of PCs and workstations located at the user's desk.

PCoIP Technology Networking Capability Summary:

- PCoIP host and portal are IPv4 based (IPv6 capable) with static or DHCP IP assignment.
- PCoIP host/portal devices have integrated 10/100/1000 Ethernet.
- All PCoIP traffic is fully protected with IPSEC using wire-speed, hardware accelerated 128-bit AES encryption and authentication.
- PCoIP traffic is primarily point-to-point between host and the peer portal with the bandwidth typically dominated by downstream host-to-portal traffic
- PCoIP downstream bandwidth is determined primarily by user profile and screen resolution. Only screen changes are transferred so that a static display requires virtually no downstream network bandwidth.
- Upstream traffic from portal to host is dominated by USB input data.
- Progressive build is used to deliver an exact image of the rendered host PC display with a minimal network loading. Progressive build helps increase system responsiveness by quickly delivering an initial image to the desktop while still supporting a fully lossless display image transfer.
- The total network bandwidth used is automatically adjusted by PCoIP host and portal devices by dynamically changing the compression ratios used. Users can select their preference for image quality or image responsiveness.
- Integrated traffic shaping allows source bandwidth metering and supports a hard device bandwidth limit setting.

The following figure shows an example PCoIP deployment, and the flow of user and datacenter network traffic.

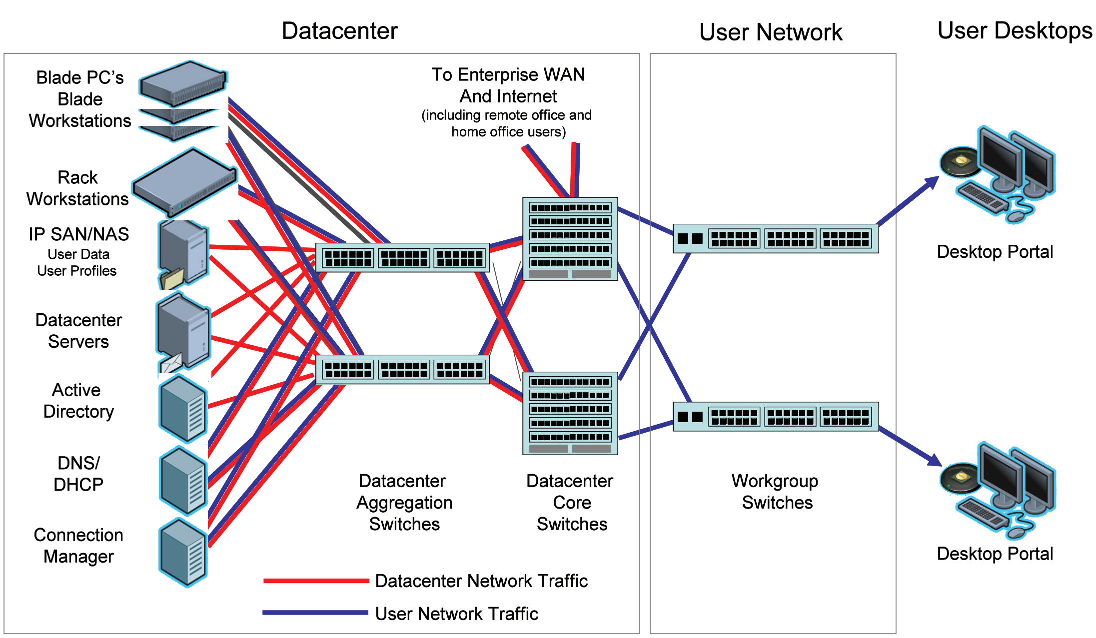

Figure 8: Enterprise Network Architecture

The PCoIP Host Processor and PCoIP Portal Processor both generate network traffic between the centralized datacenter and the user desktop. To ensure desktop responsiveness, this process of compression, transmission, and rebuilding the desktop occurs in a very short time, typically less than 10 ms.

PCoIP Network Traffic

PCoIP network traffic contains traffic for image data, USB and HD audio peripheral bridging, and system management. The traffic profile is dominated by the downstream (host to portal) compressed image data. USB and audio-generated traffic typically have significantly lower bandwidth requirements than the image data. There is a negligible amount of bandwidth required for system management traffic.

Imaging network bandwidth is only consumed during screen changes (opening/moving a window, high resolution video, etc.). Most office application scenarios are characterized by long periods with no screen changes, which result in long periods of virtually no network traffic. Overall bandwidth requirements are therefore based primarily on applications and usage scenarios, as well as on display resolution.

Network Recommendations for PCoIP Deployments

Today's enterprise networks have evolved to the point where a full-scale PCoIP deployment is readily achievable and can fit within existing switched 100Mb/s workgroup/datacenter infrastructures.

In a modern enterprise network, the available bandwidth is generally the same regardless of the direction from the datacenter. The dominant traffic in a PCoIP deployment is downstream (host to portal) so network sizing and requirements are specified on the link rates in the downstream direction.

The following procedure allows IT administrators to generate a network configuration that can enable an excellent user experience for all the users on the PCoIP system:

1. Segment the user base by applications and display resolution.
2. Determine the per-user bandwidth allocation for each user type.
3. Determine link planning bandwidth for each network link.
4. Set the device bandwidth target for each PCoIP Host.

Some key network bandwidth settings and considerations are defined below:

- **Planning Bandwidth** – defines the per-user bandwidth that will deliver the minimum desirable user experience, as defined by that user's profile.
- **Link Planning Bandwidth** – defines the bandwidth required for a certain network link.
- **Network Capacity** – is the capacity of the network to forward packets, described in bits per second.
- **Device Bandwidth Limit** – is a configuration parameter in PCoIP hosts and portals that is a hard upper bandwidth limit between the PCoIP hosts and portal pairs for downstream (host to portal) and upstream (portal to host) traffic. It is recommended that the default setting (disabled) be used in most environments.
- **Device Bandwidth Target** - is a configuration parameter in the PCoIP host and portal that defines the bandwidth that each attempts to maintain when the network is congested. This setting allows IT administrators to set a soft bandwidth floor for all users to improve the fair distribution of bandwidth when the network is congested.
- **Peak Bandwidth** - is the maximum bandwidth for a given user PCoIP host and portal pair.
- **Average Bandwidth** – is the average bandwidth for a given user PCoIP host and portal pair over a long period of time.

This procedure is a conservative estimate that guarantees a certain experience in the worst-case usage scenario where all active users are concurrently doing an operation that requires high imaging bandwidth (dragging windows or watching multimedia). This procedure is appropriate for deployments where the required amount of bandwidth is generally available, such as in LAN and campus networks. In cases where the required bandwidth is not available, the PCoIP host and portal pair will continue to operate by dynamically adjusting the network traffic generation to fit within the available network bandwidth while helping to minimize degradation in the user experience.

1. Segment the User Base

PCoIP network requirements depend highly on the usage scenario, ranging from plain-text data entry to 3D CAD. To understand the optimal network size and networking settings, IT administrators need to use their understanding of the user base. The following table shows common user types and associated profiles.

Table 1: Example User Segmentation to Optimize Network Resource Allocation

User Category	User Profile
Task Worker	Primarily text entry into forms.
Knowledge Worker/ Basic CAD	Uses standard office applications such as word processing, spreadsheets, and presentation tools, uses web, reads and writes emails, and simple CAD applications etc.
Performance User/ 3D Modeling	Similar to the Knowledge Worker user profile with the exception of occasional use of high-end visual applications, 3D modeling and may perform analysis on static images.
Video Editing	Similar to performance user, but requires consistent high-quality multimedia playback.
Extreme User	Discerning users of high-end visual applications such as 3D modeling, video editing or animation. Typically content has a higher resolution and the user performs technical analysis by dynamically manipulating images.

2. Determine Per-User Planning Bandwidth

Use the following table to determine the planning bandwidth per user. If a user has more than one display, use the column for the largest display.

Table 2: Planning Bandwidth Guidelines: User Category vs. Resolution

User category	1024x768	1280x1024	1680x1050	1600x1200	1920x1200
Task Worker	3 Mbps	5 Mbps	7 Mbps	6 Mbps	-
Knowledge Worker/Basic CAD	10 Mbps	15 Mbps	20 Mbps	22 Mbps	25 Mbps
Performance User / 3D Modeling	30 Mbps	45 Mbps	60 Mbps	65 Mbps	80 Mbps
Video Editing	-	60 Mbps	80 Mbps	85 Mbps	100 Mbps
Extreme Bandwidth Allocation	-	120 Mbps	160 Mbps	170 Mbps	200 Mbps

3. Determine the Link Planning Bandwidth

Conservatively, the network should be sized such that, when all users are active, each PCoIP connection is able to achieve the planning bandwidth allocated for each user. The planning bandwidth is the measure of the minimum quality a user will accept and feel they are receiving a true PC experience. For each link in the network, the network capacity required for a PCoIP implementation is calculated by summing the planning bandwidths of all the users sharing that network link, and adding 10% to allow for variations in the traffic profile and to ensure fair sharing of the network resources across multiple PCoIP Host/Portal pairs.

4. Set the Device Bandwidth Target

The Teradici PCoIP system provides an optional Device Bandwidth Target parameter to enable fair bandwidth sharing between many users sharing a constricted link. It is recommended to set the Device Bandwidth Target parameter on the PCoIP host to the planning bandwidth for the user.

VLAN and QoS Considerations

PCoIP technology was developed to work over standard IP networks without requiring special configurations. As such, virtual LAN (VLAN) and Quality of Service (QoS) settings are not required for system operation. In cases where the network is shared between PCoIP traffic and other data traffic, and the IT administrator wishes to guarantee a certain amount of bandwidth to PCoIP traffic, VLAN and QoS settings can be used within the network switches.

Network Latency Considerations

In general, to provide system responsiveness comparable to a traditional desktop, the round trip network latency should ideally be less than 30 ms (as part of the overall "perception budget"), which easily fits within a typical corporate or campus LAN infrastructures. Encoding, communication, and decoding of PCoIP data must be a small part of this "perception budget." The hardware advantages of PC-over-IP technology help mitigate the impact of high latency links to deliver an exceptional user experience.

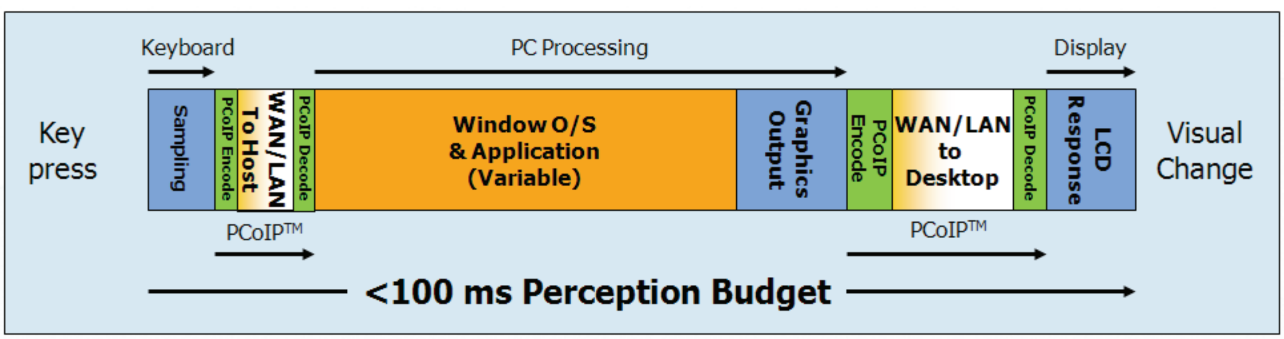

Figure 9: PCoIP as Part of the Latency Perception Budget

NAT Considerations

PC-over-IP traffic cannot undergo network address translation (NAT). The Host and destination IP addresses cannot be modified as the packets move between the host and the portal. To support NAT, PC-over-IP traffic can be tunneled through a VPN.