

Teradici Session Planning Guide

For access to the full Teradici product documentation visit [Teradici Support](#).

About PCoIP Session Planning

PCoIP is a display protocol that encodes a complete desktop or workload, which is then displayed through a PCoIP client device over a standard IP network. PCoIP technology uses advanced display compression to provide end users with cloud-based virtual computers such as GPU-enabled virtual workstations or standard virtual desktops as a progressive alternative to a local deployment model. It also supports many of the services available to physical machines, including keyboard, mouse, USB, multiple monitors, printers, audio devices, as well as custom options.

The core technology comprises the PCoIP protocol which compresses, encrypts, and transmits only pixels to a broad range of software clients, mobile clients, and stateless PCoIP Zero Clients, providing a highly secure data exchange.

The image is rendered on the host to provide the framework in which the host can transmit only the pixels across the network without being concerned about the applications or responses from the client. In best case scenarios, zero clients are used to receive the pixels and decode them, essentially eliminating display latency. The PCoIP protocol is configured to enable the display representation rendered by the virtual workstation to be exactly reproduced at the endpoint. This is referred to as lossless reproduction. This is critical, particularly in instances such as medical diagnostics, geospatial analysis, and media production, where the image itself contains important visual information. PCoIP protocol uses the User Datagram Protocol (UDP) which is much better suited for streaming media and in high-intensive graphic situations.

The PCoIP protocol provides a real-time delivery of a rich user desktop experience in virtual desktop and remote workstation environments. To ensure a responsive desktop, the PCoIP protocol must be deployed across a properly architected virtual desktop network infrastructure.

The configuration information and checklist helps users meet the said infrastructure's requirements, such as for bandwidth, quality of service, latency, jitter, and packet loss.

PCoIP Protocol Overview

The PCoIP protocol incorporates the following features:

- It uses UDP packets similar to other real-time protocols, such as VoIP and video conferencing.
- It is inherently reliable and applies intelligence to lost packet re-transmission based on the information type, for example, the PCoIP protocol guarantees the delivery of all USB information irrespective of network packet loss whereas the protocol is selective with compressed audio or pixels to avoid out-of-date re-transmission. PCoIP offers dramatic user experience advantages over traditional methods such as TCP, especially on high latency networks.
- The PCoIP protocol has a dramatic performance advantage over protocols such as TCP on high latency networks as data lost on the network does not need to be re-transmitted.
- It performs traffic shaping on the PCoIP traffic streams and will dynamically adapt image or audio quality depending on the available network resources.
- It fairly shares the available network resources with other PCoIP sessions on the network.
- It has WAN optimizations already built into the protocol.

For more detailed information, please see techsupport.teradici.com.

Who Should Read This Guide?

This guide provides capacity planning, optimization, and troubleshooting information for system administrators preparing networks for PCoIP traffic or preparing workloads for remote access using the PCoIP protocol. It acts as a troubleshooting and network planning guide.

In this guide, you will learn about:

- Network Requirements
- Planning your Network Capacity
- PCoIP Session Variables
- Sample Network and Bandwidth Settings
- Evaluating Workload Performance

Who Uses PCoIP Technology?

Teradici PCoIP technology is used in a wide range of industries, including government, education, financial services, healthcare, oil and gas, automotive, media and entertainment, architecture, engineering and construction, manufacturing, and design. For information on specific industry applications, check out the [case studies](#) featured on the Teradici website.

Network and System Requirements

The following sections outline the important network requirements and guidelines to take into consideration when setting up a PCoIP session. This section covers the following areas:

- [Packet Management](#)
- [WAN Guidelines](#)
- [WRED and VLAN Layer Configuration](#)
- [Network Port Guidelines](#)
- [Latency Guidelines](#)
- [VPN Guidelines](#)
- [Session Establishment](#)

General Network Guidelines

- Consider quality of service options such as Class-Based Weighted Fair Queuing (CBWFQ) or Low Latency Queuing (LLQ) on switch uplinks and on Layer 3 WAN/LAN links.
- Mark and classify PCoIP traffic the same as real-time interactive traffic according to your quality of service marking scheme. (namely, below VoIP RTP but above all other traffic). This is necessary for the real-time responsiveness of the protocol.

Packet Management

- If using DSCP markings, PCoIP traffic should be marked to DSC AF41 or AF31. This ensures low drop probability inside each queue if weighted random early drop (WRED) must be configured for the queue servicing the PCoIP protocol. The choice of which DSCP value to use is influenced by the presence of possible video and/or VoIP control packets.

End-to-End priority mapping

Not all switches support the same number of priority queues. Work with service providers to ensure proper end-to-end priority mapping.

- Avoid using LLQ for PCoIP packets on links that do not carry VoIP and have availability greater than 1.544 Mbps. Consider the 33% LLQ rule, which limits the amount of strict priority queuing configured on an interface to no more than $\frac{1}{3}$ of the link's capacity. The strict priority queue should only be considered if performance is suffering and there are many different types of traffic competing with PCoIP.
- Avoid adjusting the maximum transmission unit on low bandwidth links. PCoIP protocol packets should not be fragmented. It may be difficult to guarantee high-quality conversations with both VoIP and PCoIP on links with less than 1.544 Mbps of bandwidth.
- Consider tuning the hardware transmit ring to '1' to ensure that software queuing takes place if LLQ is not possible, and PCoIP or VoIP are experiencing high jitter.

Large pack serialization

Large pack serialization can sometimes cause high amounts of jitter. This should not be done in most cases as proper CBWFQ usage will allow for acceptable guaranteed session quality.

- Increase the queue depth settings in the PCoIP queue if tail drops are experienced. If you are near the maximum recommended queue depths, consider optimizing PCoIP for lower bandwidth or increasing the link bandwidth.
- Packet loss should be zero for properly configured LAN/WAN deployments. Packet loss within a single PCoIP session should target less than 0.1%. Users typically notice performance degradation if the session packet loss is greater than 0.1%, although higher loss may be tolerated.
- PCoIP packets that arrive excessively out of order may be considered as lost packets by the PCoIP protocol. Avoid packet re-ordering in the network. This will show as packet loss in the PCoIP session logs, but not in network device logs.
- Ensure that PCoIP packets are not fragmented at any point in the network path.
- Ensure that the maximum transmission unit in network devices is not below the PCoIP packet maximum transmission unit size. Defaults are 1200 or 1300 bytes for PCoIP software (depending on the vendor), and 1400 bytes when connecting PCoIP zero clients to PCoIP remote

workstation cards. Increase router maximum transition unit before reducing PCoIP packet maximum transition unit, as lower PCoIP protocol maximum transition unit can impact desktop performance. Keep in mind that network devices may add additional encapsulation and increase the PCoIP packet size.

- Do not use per-packet load balancing for any load balancing decisions along the path of traffic, including but not limited to EIGRP load balancing, static route load balancing, and MPLS load balancing.

Out of order packets

Out of order packets adversely affect the quality of the PCoIP protocol.

- For load balancers, ensure affinity (or related) is set to '1'. Ensure that the same Source Address/Destination Address is sent on the same path.
- Ensure that small packets are not prioritized over larger packets. This can cause PCoIP packet reordering, as small PCoIP packets jump ahead of larger ones.

WAN Guidelines

- Ensure that your classification and quality of service schemes inter-operate with your WAN carrier's quality of services schemes. This is especially applicable to MPLS networks.

MPLS networks

Most WAN carriers only offer three or four different classes of traffic on MPLS networks.

- Configure WAN optimization devices to bypass PCoIP packets, unless the devices explicitly support the PCoIP protocol. Some WAN optimization products can impact PCoIP packets, causing increased latency and packet loss, as well as packet reordering.

WRED and VLAN Layer Configuration

- Configure WRED in the path of all PCoIP conversations. On Cisco routers this is the 'random-detect' command. The PCoIP protocol incorporates rate limiting and flow control mechanisms optimized for virtual workloads such as VDI or cloud-based graphics applications.

NDP Protocol

The Neighbor Discovery Protocol (NDP) is a protocol in the Internet protocol suite used with IPv6. Unlike most NDP traffic, PCoIP protocol benefits from WRED mechanics. When an WRED algorithm causes packet loss PCoIP protocol adapts to network congestion with minimal impact on the user experience. If a Tail Drop scheme is deployed, large bursts of PCoIP data may be suddenly dropped by the network before PCoIP protocol has time to adapt to congestion.

- Confirm that the network interface is not configured for WRED if you have selected WRED for the service policy on that interface. Configuring WRED on the physical interface overrides all other quality of service queuing configurations.
- Consider segmenting PCoIP traffic via Layer 2 VLAN and/or class of service types at the access layer of your network.
- Only use Layer 2 quality of service class of service prioritization if there is noted congestion at the access layer or between the access and aggregation (distribution) layer. Consider adding Layer 2 uplink bandwidth before applying Layer 2 quality of service, if possible.
- Avoid the use of AutoQoS features at the Layer 2 layer for devices that do not explicitly support AutoQoS for PCoIP packets, as this may result in WRED being applied at the switchport layer through the use of Shared/Shaped Round Robin (SRR) queues. When using the AutoQoS feature, SRR queues are automatically configured on many access layer platforms. By default, these enforce WRED for all but trunked packets marked with class of service 5 (generally VoIP packets from a hardphone). Often PCoIP packets are treated as scavenger class traffic, which can negatively impact desktop performance.
- Avoid traffic shaping unless absolutely necessary. Shaping works to smooth traffic bursts and achieve a defined committed access rate (CAR) by buffering packets. Traffic shaping increases PCoIP packet latency and can impact user experience. If necessary, consider traffic policing as an alternative.
- Ensure that a full-duplex end-to-end network is used.

Monitoring older switches

Older switches may incorrectly default to half-duplex when connected to a link with auto-negotiation. In this case, explicitly set the switch link to full-duplex.

Network Port Guidelines

- Ensure that network ports are open for the PCoIP protocol and virtual desktops. For details, see the Teradici Knowledge Base article [What are the required TCP/UDP ports for PCoIP technology? \(KB 15134-114\)](#).
- Ensure that PortFast is enabled on all network ports that have PCoIP end points connected to them.

PortFast mode on an IP phone

If an IP phone is connected between the client and the switch, you may need to set a different PortFast mode because of the internal switch inside the phone. This ensures that the port is immediately configured to forward traffic in the event of a spanning tree recalculation.

- Ensure intrusion protection services (IPS) in network devices and/or laptop/desktop software have been disabled or configured to allow the PCoIP protocol and other virtual desktop network ports. IPS can block some or all network ports and/or throttle bandwidth for the PCoIP protocol.
- Avoid gaps in network connectivity. PCoIP sessions will disconnect after 30 seconds of loss in traffic in either network direction or the PCoIP port (4172 UDP). Intrusion protection services (IPS) or intrusion detection services (IDS) should be disabled, or configured to allow (4172 UDP).

Latency Guidelines

- Ensure that the round trip network latency is within specification. Excessive latency will impact desktop performance.
- Latency should be less than 250 ms round trip for virtual desktops, Cloud Access Software and PCoIP remote workstation cards.
- Ensure the latency variation is less than 30 ms.

Latency variation for a remote workstation implementation

For a remote workstation implementation, ensure that the latency variation is limited to approx. 1 frame when playing a 30 fps (HD) video.

VPN Guidelines

- PCoIP can be used with a VPN. The VPN must be UDP compatible. VPNs using SSL with a TCP transport layer are not supported due to the error correction of TCP. There are limitations around the performance with TCP VPN and PCoIP UDP.
- If a VPN is used, confirm that UDP traffic is supported (IPsec, or DTLS-enabled SSL solutions). Do not route PCoIP traffic through TCP-based SSL tunnels.
- Avoid VPN overhead. If possible, consider a VPN-less secure remote access solution that supports the PCoIP protocol.
- Use QoS Pre-Classify if CBWFQ or LLQ is necessary on the outgoing interface of the VPN device. This may not be available on many platforms or in many designs.
- Confirm the VMware ESXi virtual switch traffic shaper is turned off.
- Determine other protocol traffic that exists on the network, especially other high priority traffic that could impede PCoIP forwarding.
- Determine the network characteristics that are key for a successful real-time protocol deployment, including latency, jitter (latency variation), and packet loss.
- Optimize networks for virtual desktop connections. For details, see the following topics [What are the recommended host card bandwidth settings? \(KB 15134-242\)](#) and [How do I optimize the Windows desktop experience settings to reduce PCoIP session bandwidth? \(KB 15134-880\)](#) on the Teradici support site.

Session Establishment

For troubleshooting tips, FAQs and specific documentation around PCoIP Session Establishment, see the following KB article <https://help.teradici.com/s/article/4529>. This article includes guidelines, troubleshooting checklists as well as links to the PCoIP connection instructions found in the various Teradici component guides.

Network Capacity Planning

The following section outlines important bandwidth requirements to consider when you are negotiating network capacity.

Network Capacity Planning Guidelines

- Ensure sufficient bandwidth is allocated to PCoIP traffic. Plan for an average of 150–250 kbps for standard office applications with Windows Experience settings optimized. Or, plan for 80–150 kbps when optimizing the Windows Experience settings and also optimizing PCoIP session variables.
- Consider the bandwidth required for audio input/output and USB traffic. PCoIP Zero Client bandwidth limit will limit USB traffic per PCoIP session.
- Plan a minimum of 1 Mbps per simultaneous user watching a 480p video window. More bandwidth may be required depending on video resolution/user quality requirements.
- Ensure that network bandwidth analysis includes following network loading guidelines (for example, peak bandwidth less than 80% of theoretical link capacity.)
- Ensure that there is sufficient bandwidth headroom for bursts of PCoIP protocol traffic.
- Plan for a minimum bandwidth headroom of 500 kbps to 1 Mbps. The actual peak bandwidth required will depend on the user and applications (for example, the number of pixels changing, the level and complexity of compression, and the display frame rate).
- Plan for adequate peak capacity for Remote Workstation Card users. Many workstation applications require a higher peak bandwidth that varies per user and imaging workload. It also depends on how many users are sharing the link.
- Over-subscription analysis (optional) can enable efficient link sizing while maintaining a reasonable expectation that peak bandwidth is available per session.
- Consider bandwidth reduction options when operating in a known constrained network environment.
- Ensure you optimize Windows desktop experience settings, PCoIP session variables and display resolution and/or number of displays.

Network Capacity Planning Guidelines

- PCoIP Ultra with GPU optimizations configured for 4:2:0 chroma sub-sampling offers additional bandwidth savings for deployments comprising NVIDIA NVENC GPUs. For more information on Session Planning for PCoIP Ultra, see [here](#).

Workload Analysis and Use Case Classification

It is recommended that use cases are evaluated in view of user experience demands as well as network bandwidth considerations. The following sections provide metrics in relation to bandwidth and user experience ratings, user types, network bandwidth and frame rates.

Use Case Classification

PCoIP technology provides a range of capabilities suitable for a variety of user types, including:

- **Task Workers:** Task workers typically use applications involving simple screen data and text entry, such as call centers or companies that deliver online certification tests.
- **Knowledge Workers:** Knowledge workers use enterprise-level productivity suites, interactive CD quality audio communication tools, consume high-definition video and may have some lightweight 3D applications.
- **Artists/Designers:** Artists and designers perform creative tasks that need high-quality, graphic-intensive 3D rendering for high definition imaging, video, and animation work.
- **Scientists/Engineers:** Scientists and engineers use compute farms and High Performance Computing (HPC) platforms to extract high resolution visual information and animations from complex data sets.

Bandwidth Consumption

Bandwidth Consumption for Moderate Workloads

The following estimates assume 1920 x 1080 displays. Higher resolutions may increase your actual bandwidth requirements. The following table correlates bandwidth consumption according to user type:

Bandwidth Consumption by User Type

Example User Type	Average Bandwidth	Video Content	Peak Bandwidth
Task Worker	70-100 kbps	No Video	0.5-1 Mbps
Knowledge Worker	100-500 kbps	5-10 Mbps	5-10 Mbps
Artists/Designers	10 Mbps	Included	30 fps workload: Typ. 30-60 Mbps 60 fps workload: Typ. 60-120 Mbps

Bandwidth Consumption for Graphic Intensive Workloads

For users such as artists and designers that require graphic-intensive 3D rendering or engineers and scientists that use HPC platforms. Network administrators should consider the peak bandwidth allocation needed to sustain the highest image quality. If insufficient bandwidth is available, the image quality is automatically adjusted accordingly.



Peak bandwidth calculation

Peak bandwidth demand can usually be estimated at 0.5 - 2.0 bits-per-pixel compression ration when using standard SSE 4.2 PCoIP encoding. CPU optimizations using AVX2 or GPU optimizations using NVENC in 444 chroma sub-sampling mode. When using NVENC with 420 chroma sub-sampling, peak bandwidth can typically be estimated at 0.1-0.5 bits-per-pixel. In unusual use cases, where it is critical to sustain detailed screen content (e.g. CAD textures and fine lines) at both very high quality levels (e.g. Q90) and maximum frame rate, the peak bandwidth consumption may need to be adjusted to support a compression ratio of 3-4 bpp.

Average bandwidth vs peak bandwidth

While use case dependent, the average bandwidth is typically less than 10% of the peak demand, considering minimal consumption during periods of display inactivity. Bandwidth provisioning of a shared network should be based on a sum of averages, plus an empirical peak allocation, rather than a strict sum of peak values from the table alone.


Image quality considerations

Peak and average bandwidth consumption can be intentionally constrained by using non-default quality settings as discussed in the PCoIP Session Variable section. For example, video content bandwidth may be lowered to less than 0.25 bpp if the Maximum Initial Image Quality is adjusted from a default 80 to a value of 70.

Display Frame Rate Performance

Teradici PCoIP technology has been widely adopted by industries that demand the highest possible image quality with multi-monitor displays such as the CAD, finance, medical, animation, and film industries involving ultra-high resolution display and rendering operations.

Example Target Frame Rates supported by the PCoIP Protocol



Display Resolution	Single	Dual
Zero Client		
1920 x 1080 FHD	Up to 60	Up to 60
2560 x 1600 WQXGA	Up to 48	Up to 30
PCoIP Ultra Software Client*		
1920 x 1080 FHD	Up to 60	Up to 60
2560 x 1600 WQXGA	Up to 60	Up to 30
3840 x 2160 4K/UHD	Up to 30	Up to 30

- Teradici Cloud Access Software running PCoIP Ultra protocol enhancements delivers up to 60 fps dependent on host and client CPU specifications. Both host and client require support for the AVX2 CPU extensions to take advantage of the enhancements. Refer to the Cloud Access Software Administrator Guide for platform requirements.

User Experience Ratings and Bandwidth Considerations

For your convenience, Teradici has established five user profiles, each with a unique combination of PCoIP session variables. The session variables may be configured on a specific user profile basis per Sample Policies by User Profile or via the Group Policy Editor (Windows) or session configuration file (Linux) as described in the PCoIP Session Configuration. The profiles, their level of experience and the recommended configurations, are described next.

Profile A represents the best graphics experience, default for PCoIP Remote Workstation Cards.

Profile B represents a standard experience, default for PCoIP Standard and Graphics Agents.

Profile C represents bandwidth-optimized experience for knowledge workers operating in constrained network scenarios.

Profile D represents bandwidth-constrained experience for task workers operating in constrained network scenarios.

Profile E represents maximum bandwidth-constrained experience suitable for task workers only - for example 10 small screen users sharing a T1 link.

The following table shows the bandwidth and user experience for the different user types, and the worker profiles they correlate to:

Bandwidth and User Experience ratings

User Experience Ratings and Bandwidth Considerations

Worker Type	Bandwidth	User Experience	Use Cases
Profile A	Highest (LAN)	Best	Artists/Designers/Scientists/Engineers
Profile B	High (LAN/WAN)	Great	Artists/Designers/Scientists/Engineers
Profile C	Optimized	Good	Knowledge Workers
Profile D	Constrained	Good	Knowledge Workers/Task Workers
Profile E	Lowest	Limited	Task Workers

PCoIP Session Configuration

This section lists the configuration parameters available for Windows or Linux host environments.

Configure Build-to-Lossless Feature

The *Turn off Build-to-Lossless* feature setting specifies whether to disable the build-to-lossless feature of the PCoIP protocol, which is disabled by default.

The build-to-lossless feature is disabled by default for Cloud Access Software Agents in order to provide network bandwidth savings. However, in environments that require high levels of image quality and/or lossless image reproduction, the build-to-lossless feature should be enabled.

Configure the Maximum PCoIP Session Bandwidth

This setting specifies the maximum bandwidth, in kbps, in a PCoIP session. The bandwidth includes all imaging, audio, virtual channel, USB, and control PCoIP traffic.

Set this value based on the overall capacity of the link to which your endpoint is connected, taking into consideration the number of expected concurrent PCoIP sessions.

The limit set can vary across usage scenarios, such as the following:

- Single user on a link (that is, home user on an Internet connection). Set the maximum PCoIP session bandwidth to 90%.
- The maximum PCoIP session bandwidth should not be set to a value lower than 300 kbps, for example Profile E in the Sample Policies by User Profile table.
- The Sample Policies by User Profile table shows maximum PCoIP session bandwidth settings for different user profiles.

Configure the PCoIP Session Maximum Transmission Unit

This setting specifies the maximum transmission unit (MTU) size for UDP packets for a PCoIP session.

Configure the PCoIP Session Bandwidth Floor

The MTU size includes IP and UDP packet headers. TCP uses the standard MTU discovery mechanism to set MTU and is not affected by this setting. The maximum MTU size is 1500 bytes. The minimum MTU size is 500 bytes. The default value is 1200 bytes.

Typically, you do not have to change the MTU size. Change this value if you have an unusual network setup that causes PCoIP packet fragmentation.

This setting applies to the agent and client. If the two endpoints have different MTU size settings, the lowest size is used.

If this setting is disabled or not configured, the client uses the default value in the negotiation with the agent.

Configure the PCoIP Session Bandwidth Floor

This setting specifies a lower limit, in kbps, for the bandwidth that is reserved by the PCoIP session. This setting configures the minimum expected bandwidth transmission rate for the endpoint. Note that the session will use less bandwidth than the floor if there is less data that needs to be sent. This setting is useful for two primary scenarios:

- For networks with persistent low levels of packet loss like wireless networks, this setting causes the PCoIP bandwidth manager to ignore this packet loss and maintain the bandwidth level indicated in this setting.
- The initial responsiveness of a session can be improved by this setting since the session does not have to discover how much bandwidth is available.

Make sure that you do not over-subscribe the total reserved bandwidth for all endpoints. Make sure that the sum of bandwidth floors for all connections in your configuration does not exceed the network capability:

*$(n \text{ users}) * (\text{minimum bandwidth}) \leq 90\% \text{ of available link bandwidth}$*

The default value is 0, which means that no minimum bandwidth is reserved. When this setting is disabled or not configured, no minimum bandwidth is reserved. This setting applies to the agent and client, but the setting only affects the endpoint on which it is configured.

Bandwidth Floor Configuration Examples

Configure the UDP Port to which the PCoIP Server Binds and Listens

User Type	Link Capacity [Mbps]	# Users	Packet Loss	Example Bandwidth Floor Configuration	Corresponding Bandwidth Floor Policy [kbps]
Artist/ Designer	100	1	0.1%	10% of Link Bandwidth	10000
	100	5	0.1%	20% of Link Bandwidth*	20000
	100	1	1%+	50% of Link Bandwidth	50000
Knowledge Worker	10	1	0.1%	10% of Link Bandwidth	1000
	10	5	0.1%	20% of Link Bandwidth*	2000
	10	1	1%+	50% of Link Bandwidth	5000
Task Worker	1	1	0.1%	50% of Link Bandwidth	500
	1	5	0.1%	20% of Link Bandwidth*	200
	1	1	1%+	90% of Link Bandwidth	900

*The sum total of bandwidth floor values across should not exceed the available Link Bandwidth or the number of users.

Configure the UDP Port to which the PCoIP Server Binds and Listens

This setting specifies the UDP agent port bound to by software PCoIP hosts.

Configure the PCoIP Session Audio Bandwidth Limit

The UDP port value specifies the base UDP port that the agent attempts to bind to. The UDP port range value determines how many additional ports to try if the base port is not available. The port range must be between 0 and 10.

This setting applies to the agent only.

Configure the PCoIP Session Audio Bandwidth Limit

This setting specifies the maximum bandwidth that can be used for audio (sound playback) in a PCoIP session.

The audio processing monitors the bandwidth used for audio. The processing selects the audio compression algorithm that provides the best audio possible, given the current bandwidth utilization. If a bandwidth limit is set, the processing reduces quality by changing the compression algorithm selection until the bandwidth limit is reached. If minimum quality audio cannot be provided within the bandwidth limit specified, audio is disabled.

This setting applies to the agent only. You must enable audio on both endpoints before this setting has any effect. In addition, this setting has no effect on USB audio.

When this setting is modified during an active PCoIP session, it will take effect immediately.

Configure PCoIP Image Quality Levels

Image quality control PCoIP image compression levels during periods of network congestion. The Minimum Image Quality, Maximum Initial Image Quality, and Maximum Frame Rate values interoperate to provide fine control in network-bandwidth constrained environments. When dynamic content such as video is displayed, consecutive frames are presented between the configurable Minimum and Maximum Initial Quality values, the quality of each frame determined by bandwidth availability. If there is insufficient network bandwidth, a frame is discarded. Static content is initially presented at a value no higher than the Maximum Initial Quality value, then increased to perceptually lossless or numerically lossless quality according to the Build-to-Lossless configuration.

- **Minimum Image Quality** Use the Minimum Image Quality value to balance image quality and frame rate for limited-bandwidth scenarios. You can specify a value between 30 and 100. A lower value allows higher frame-rates, but with a potentially lower quality display of dynamic image content. A higher value provides higher image quality, but with potentially lower frame

Configuration for Windows

rates when network bandwidth is constrained. When network bandwidth is not constrained, PCoIP maintains maximum quality regardless of this value. The default value is 40 for Cloud Access Software and 50 for Remote Workstation Cards.

- **Maximum Initial Image Quality** Use the Maximum Initial Image Quality value to reduce the network bandwidth peaks generated by PCoIP by limiting the initial quality of the changed regions of the display image. You can specify a value between 30 and 100. The default value is 80 for Cloud Access and 90 for Remote Workstation Card. A lower value reduces the image quality of content changes and decreases peak bandwidth demand. A higher value increases the image quality of content changes and increases peak bandwidth requirements. Unchanged regions of the image progressively build to a lossless (perfect) quality regardless of this value. A value of 80 or lower best utilizes the available bandwidth. The default value is 80 for Cloud Access Software and 90 for Remote Workstation Cards.
- **Maximum Frame Rate** Use the Maximum Frame Rate value to manage the average bandwidth consumed per user by limiting the number of screen updates per second. You can specify a value between 1 and 60 frames per second. A higher value provides smoother frame transitions important for some graphics applications but may use more bandwidth. A lower value uses less bandwidth but may increase the jitter rate of applications with a high source frame rate. The default value is 60 for both Cloud Access Software and Remote Workstation Cards.

Configuration for Windows

This section lists the configuration parameters in the Windows Local GPO editor. The PCoIP settings will not appear in the GPO editor until the Administrative Template File is imported. If you have not done so, consult your respective administrators's guides on how to import the Administrative Template File.

Configuration for Linux

The Agent for Linux uses a single configuration file which is only writable by the administrator.

For information on how to set up your session configuration, consult the respective [agent guides for Linux](#).

Windows and Application Optimizations

Teradici recommends the following Windows and application optimizations for the preceding sample profiles.

Windows Performance Options Optimization

For Virtual Desktop Infrastructure (VDI) use case in which high-server consolidation ratio's are advantageous, Microsoft offers a host of suggested Windows 10 optimizations: <https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds-vdi-recommendations>

Furthermore, third parties such as Login VSI, provide tools and guidelines for additional Windows 10 optimizations: <https://www.loginvsi.com/blog/520-the-ultimate-windows-10-tuning-template-for-any-vdi-environment>


Optimizations related to Windows 10 Visual Effects

The performance Options dialog is accessible from **Control Panel>System and Security>System>Advanced System Settings**. For profiles **D and E**, Teradici recommends that you disable Show window contents while dragging.

Adobe Reader Optimization

For Adobe Reader, go to **Edit>Preferences>Page Display** and set the *Smooth Text* setting to **For Monitor**, as shown below

Windows and Application Optimizations



Adobe Premiere Optimizations

If Adobe Premiere is experiencing unusual lag or problematic audio/video synchronization behavior, you can adjust Windows for best performance from **Control Panel>System>Adjust the Appearance and Performance of Windows>Adjust for best performance**.

Remote Access to 3D Rendering Software

The user interactivity of creative applications such as visual effects or architectural rendering software is contingent on reserved CPU availability to service peripheral devices and refreshing the GUI in a responsive manner. Applications with built-in or plug-in render capabilities are prone to poor user interactivity and low frame rates if the render engine consumes disproportionate CPU resources.

[Teradici Cloud Access Software](#) requires a small proportion of high priority CPU access to enable a responsive remote user experience. In many cases this may be ensured by adjusting render settings within the creative application software itself as many popular render engines provide CPU utilization controls via processor affinity, thread count or thread parameters.

In cases where remote user interactivity is adversely impacted by software applications using disproportionate CPU resources, performance may be remedied by one or more of the following actions:

Linux Environments

Linux systems spawn the `pcoip-server` process within scope of the `system.slice` while render processes are typically spawned in the `user.slice` (user session context), it is making it difficult to schedule relative priorities between `pcoip-server` and conflicting render processes.

- If available, adjust render engine settings to allow CPU resources for other tasks
- Use Linux `cgroups` to re-allocate shares in favor of `system.slice`, for example by allotting `system.slice` twice as many `cpu.shares` as the `cpu.slice`. It is highly recommended to test thoroughly in advance of deployment because share re-allocation may have unanticipated side effects such as enabling a system daemon to starve other user applications of CPU resources.

Read more on Linux `cgroups`: <https://www.golinuxcloud.com/beginners-guide-cgroups-slice-examples-tutorial/>

Windows Environments

- If available, adjust the render engine settings to allow CPU resources for other tasks.
- Manually set processor affinity for the application software, allowing at least one core for the pcoip agent (pcoip_server.exe). Note that this approach will not work if the application software spawns render tasks that override Windows processor affinity settings.
- Deploy third party real-time CPU optimization software.

Read more on Windows Processor Affinity: <https://www.thewindowsclub.com/processor-affinity-windows>

Enhanced Audio/Video (A/V) Synchronization

This client side setting which is available on PCoIP Software Clients, reduces the delay between the audio and video channels. Enhanced A/V Sync offers improved lip sync while also reducing the number of dropped video frames during movie playback. When enabled, a small lag in user interaction and responsiveness is introduced and the maximum frame rate may be reduced. Enhanced A/V Sync is enabled on a per-display basis, so you can dedicate individual displays to playback without impacting responsiveness on the others. For further information on turning on and using the Enhanced A/V Sync on the PCoIP Software Clients, see the following guides:

- [PCoIP Software Client for Windows](#)
- [PCoIP Software Client for macOS](#)
- [PCoIP Software Client for Linux](#)

Some video editing applications such as Adobe Premiere, provide user settings that adjust A/V synchronization. In the case of Adobe Premiere, audio latency can be adjusted in the audio hardware preferences page. For more information on this, consult the [Adobe Premiere Pro User Guide](#).

Wacom Tablet Support

Teradici Cloud Access Software supports Wacom tablets in two configurations: bridged, where peripheral data is sent to the desktop for processing, and locally terminated, where peripheral data is processed locally at a PCoIP Software Client or a PCoIP Tera2 Zero Client. Some of the advantages of using Cloud Access Software in conjunction with Wacom tablets include:

- The ability to interact with your Wacom tablet in a remote environment as if using a desktide workstation.
- Increase the flexibility and footprint of your VFX or video editorial pipelines by centralizing your workstations and providing WAN access for remote artists or satellite studios.
- It dramatically increases your deployment options by enabling your artists with public cloud remote workstations.

The Benefits of Local Termination

While Wacom tablets offer reasonable interactivity in bridged mode at latencies up to 25ms, locally termination greatly improves the responsiveness and tolerance for higher-latency (including 25ms and higher) network connections. High interactivity is very important for artists using Wacom tablets. Local termination presents an immediate local cursor to assist hand-eye coordination in addition to bridging parameters such as pen position, pressure and angle.

For information on which Wacom tablet models are supported with local termination on a PCoIP Software Client and PCoIP Tera2 Zero Client, see [Locally Terminated Wacom Tablets - Windows](#) and [Locally Terminated Wacom Tablets - Linux](#).

Working-from-Home Support

Teradici Cloud Access Software enables remote access to Windows or Linux based remote workstations including:

- Physical workstations or standalone computers (either desktide or centralized).
- Virtual workstations on VMware ESXi, KVM or Nutanix AHV Hypervisors.
- Non-graphics virtual desktops on VMware ESXi, Redhat KVM or Nutanix AHV Hypervisors.
- Remote workstations that reside on the public cloud.

Teradici Cloud Access Software is ideally positioned to manage and provide solutions to your pressing corporate work-from-home demands. For an in-depth view of our work-from-home offerings, please see our [Work-from-Home Rapid Response Guide](#).

This guide outlines:

- [Work-from-Home options for Standalone Computers](#).
- [Work-from-Home options with Remote Workstation Cards](#).
- [Work-from-Home options with Cloud Access Software](#).
- [Work-from-Home options for VMware Horizon](#).
- [Performance Tips for Work-from-Home Use Cases](#).

Sample Network and Bandwidth Settings

Example Configurations

This section lists examples of configuration and bandwidth settings.

User Profiles by Worker Type

The following table assists administrators in classifying their workers based on application usage. Once a category is determined, a suitable profile can be assigned to the worker using the table below.

Recommendations for Different Remote Workstation Desktop Categories

Visual Desktop Worker Categories	Data/Text Entry	Office & Web	Embedded Web page	Half Screen	Full Screen	3D Apps	Occasional Video Usage	Frequent Video Usage
Task Workers	✓							
Basic Productivity Workers	✓	✓					✓	
Advanced Productivity Workers	✓	✓	✓				✓	
Knowledge Workers	✓	✓	✓	✓				✓
Artists/Designers	✓	✓	✓	✓	✓	✓		✓

The following table shows which User Profiles are applicable to each type of User Type.

User Profiles by Worker Type

Remote Workstation Worker Categories and Sample Profiles

User Type	A	B	C	D	E
Task Workers		✓	✓	✓	✓
Basic Productivity Workers		✓	✓	✓	
Advanced Productivity Workers		✓	✓		
Knowledge Workers		✓	✓		
Artists/Designers	✓	✓			

You can use these user profiles as guidelines when planning your workload requirements in specific user scenarios.

Use profiles and guidelines as a starting point

These user profiles and provided guidelines are only intended to provide you with a starting point to plan your deployment. User perceptions can vary greatly, and more or less bandwidth may be needed to satisfy your user performance requirements based on the operating system and specific applications being used. The information provided here should not replace real-world testing in your own environment with actual workloads.

Sample Policies by User Profile

The following table outlines sample policies by user profile.

Group Policy Editor Configuration	A	B	C	D	E
Maximum PCoIP Session Bandwidth (kbps)	900000	900000	4000	1200	300
Enable Build-to-Lossless feature	Yes	No	No	No	No
Minimum Image Quality	50	40	40	30	30
Maximum Initial Image Quality	90	80	70	70	70
Maximum Frame Rate (fps)	60	30	16	8	4
Session Audio Bandwidth Limit (kbps)	256	256	48	48	0-32
Windows and application optimizations	N/A	N/A	Yes	Yes	Yes

The above table deals with the profiles A thru E which are defined in [User Experience Rating and Bandwidth Considerations](#). The table shows recommended group policies for user profiles A thru E for users with HD or 2K display configurations. Profile A or B users with 4K/UHD display configurations should use the settings described in [Session Planning for PCoIP Ultra Enhancements](#).

Cisco Router Configuration Example

The following example contains marking and Class-Based Weighted Fair Queuing (CBWFQ) with Low Latency Queuing (LLQ) for VoIP. SIP traffic is not treated. The example assumes a LAN Ethernet interface and a WAN Serial T1 interface. Quality of service is configured to guarantee the following:

- Strict priority for four G.729 VoIP calls marked as EF.
- Reserved bandwidth for two 'task worker' PCoIP sessions marked as AF41 (500 kbps minimum peak bandwidth, limited ability for over-subscription).
- The default class gets all the remaining bandwidth and is fair queued.

Sample Cisco router configuration settings:

```
!match PCoIP packets
```

```
access-list 100 permit tcp any any eq 4172
```

```
access-list 100 permit udp any any eq 4172
```

```
``bash class-map match-all VOIP-IN match ip rtp 16384 16383 class-map match-all PCOIP-IN
match access-group 100
```

```
```bash
class-map match-all VOIP-OUT
match ip dscp EF
class-map match-all PCOIP-OUT
match ip dscp AF41
```

```
policy-map ETH-IN
class VOIP-IN
set ip dscp EF
class PCOIP-IN
set ip dscp AF41
```

## Cisco Router Configuration Example

```
policy-map SERIAL-OUT
class VOIP-OUT priority 128
class PCOIP-OUT
bandwidth 1000
class class-default
fair-queue
```

```
interface Serial 0/1
bandwidth 1544
no fair-queue
service-policy output SERIAL-OUT
```

```
!trust dscp markings coming into this router from across the WAN
```

```
!do this if you need Layer 2 COS QoS and have a DSCP-COS map defined or set COS
on e0/1
```

```
mls qos trust dscp
```

```
interface Ethernet 0/1
service-policy input ETH-IN
```

# CPU and GPU Optimizations

The CPU and GPU optimizations offered by PCoIP Ultra protocol enhancements are targeted at desktops and workflows where high performance and interactivity are required. This section provides video playback benchmarks to help you plan your PCoIP Ultra deployment.

## CPU and GPU Optimizations

PCoIP Ultra offers the choice between CPU and GPU optimizations. The CPU optimization mode targets up to dual 4K/UHD display resolutions, leveraging the CPU's AVX2 instruction set to provide efficient content encoding at very high color accuracy, including the option for Build-to-Lossless image quality. The GPU optimization modes targets up to dual 2560x1600 display resolutions, leveraging the NVIDIA NVENC hardware encoder if present, which provides CPU relief for applications such as interactive rendering where CPU clock cycles may be premium. PCoIP Policies can be used to configure NVENC for either YUV 4:4:4 or YUV 4:2:0 chroma sub-sampling schemes.

## PCoIP Ultra Options

The following table outlines the CPU and GPU optimizations that work with specific system requirements:

## PCoIP Ultra Options

Requirement	CPU Optimizations - AVX2 RGB 4:4:4	GPU Optimizations - NVENC YUV 4:4:4	GPU Optimizations - NVENC YUV 4:2:0
Highest Image Quality Desktop (Up to 25x16 at 60 fps or 4K/UHD @ 30 fps)	✓		
CPU Relief e.g. VFX Render Applications (Up to 25x16 at 60 fps, 4K/UHD not recommended)		✓	
Bandwidth Preservation (e.g. constrained WAN)			✓
Highest frame rate from a low-cost thin client			✓

# Host System Requirements

To take advantage of PCoIP Ultra, you need:

- A **PCoIP agent** (any type), 19.11.0 or later
- A **PCoIP Software Client** (any type), 19.11.0 or later

 **PCoIP Tera2 Zero Clients do not support PCoIP Ultra**

PCoIP Ultra is only available through PCoIP Software Clients.

- **CPU optimization** requires CPU support for the AVX2 instruction set on both the agent and client machines.
- **GPU optimization** requires an NVIDIA graphics card that supports NVENC on the agent machine.

For more information the requirements for PCoIP Ultra, see:

- [PCoIP Graphics Agent for Windows](#)
- [PCoIP Graphics Agent for Linux](#)
- [PCoIP Standard Agent for Windows](#)
- [PCoIP Standard Agent for Linux](#)

# Client Device Requirements

The following client device requirements are necessary to utilize PCoIP Ultra:

- **PCoIP Ultra with CPU optimizations enabled:** Intel 8<sup>th</sup> Generation Core i5/i7 processors @3.0 GHz or higher base frequency with dual memory banks are required to achieve 4K30 FPS. An example of this includes Asus VivoMini VC66C (Intel Core i7).
- **PCoIP Ultra with GPU optimizations enabled:** PCoIP Ultra with GPU optimizations uses software-based H.264 decoding, supporting resolutions up to 2560x1600 at 30 FPS. Teradici recommends Intel 8<sup>th</sup> Generation Core i5/i7 processors @ 3.0 GHz or higher base frequency with dual memory banks.

# PCoIP Ultra Compatibility Matrix

The following table outlines PCoIP Host and PCoIP Client configurations across Windows, Linux and Mac OS for PCoIP Ultra.

PCoIP Host Configuration	PCoIP Client Configuration		
	Windows	Linux	Mac OS
PCoIP Graphics Agent for Windows or Linux: Nvidia NVENC or AMD RapidFire supported	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 software decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 hardware decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 software decoder.
PCoIP Graphics Agent for Windows: Nvidia GeForce graphics	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.
PCoIP Standard Agents for Windows or Linux	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.	<b>PCoIP Ultra:</b> CPU Offload with PCoIP AVX2 decoder.

## PCoIP Mobile and PCoIP Zero Client Configuration

The PCoIP Graphics Agents for Windows or Linux, PCoIP Standard Agents for Windows or Linux and PCoIP Remote Workstation Card with the PCoIP Mobile Client does not support PCoIP Ultra. It is recommended to use the PCoIP ARM decoder.

The PCoIP Graphics Agents for Windows or Linux, PCoIP Standard Agents for Windows or Linux and PCoIP Remote Workstation Card with the PCoIP Zero Client does not support PCoIP Ultra. It is recommended to use the TERA2 decoder.

# Example Configurations with Benchmarks

The following is an example of a set configuration with specific benchmarks set:

**Host Platform:** Dell R7910 (Xeon Haswell E5-2637v3 3.5 GHz) Windows 10

**Hypervisor:** VMware ESXi 6.7

**Virtual Machine:** Windows 10, 8 GB RAM

**Graphics:** NVIDIA M60, GRID 8.0

**Network:** 1 Gbps LAN

**Client Endpoint:** Asus VivoMini VC66C, Windows 10

**Display:** 4K/UHD 3840 x 2160

**PCoIP Agent:** Version 19.08

**PCoIP Client:** Version 19.08

**Video Content:** Big Buck Bunny 1080p 24fps (Opening Scene)

## Network Bandwidth Consumption

The network bandwidth used in conjunction with PCoIP Ultra AVX2 CPU optimizations are dependent on the configured image quality and size of the video display surface. The table below shows the average bandwidth consumption (in Mbps) for the opening scene of Big Buck Bunny at native 1080p resolution and 1080p upscaled to full-screen (4K/UHD) playback at PCoIP Maximum Initial Image Quality settings of Q70 – Q100.

	Q70	Q80*	Q90	Q100
Native 1080p Playback	16	27	40	61
4K/UHD Playback	27	47	68	107

## Host CPU Utilization

\*Q80 is the default

On constrained networks, the PCoIP protocol dynamically adjusts to the available network bandwidth. In cases where bandwidth needs to be proactively constrained, refer to [PCoIP settings section](#) for policies that enforce bandwidth and/or frame rate limits.

## Host CPU Utilization

The table below shows host virtual machine average CPU utilization as measured by the VMware ESXi Performance Monitor at default image quality settings for various virtual machine pCPU core allocations.

	2 pCPU	4 pCPU	6 pCPU	8 pCPU
Native 1080p Playback	20%	12%	6%	5%
4K/UHD Playback	41%	24%	12%	8%

For troubleshooting information around implementing the PCoIP Ultra protocol enhancements, see the knowledge base article: [Troubleshooting PCoIP Ultra](#).