

HP Anyware Session Planning Guide

23.12

Table of Contents

PCoIP Session Planning Guide	5
About the PCoIP Protocol	5
Who Uses PCoIP Technology?	6
About PCoIP Session Planning	7
Network and System Requirements	8
General Network Guidelines	8
Packet Management	8
WAN Guidelines	10
WRED and VLAN Layer Configuration	10
Network Port Guidelines	12
Latency Guidelines	12
VPN Guidelines	13
VPN Alternatives	14
Session Establishment	14
Cisco Router Configuration Example	14
Workload Characterization and Network Planning	16
Use Case Classification	16
Bandwidth Consumption Estimates	17
Peak Estimates for Graphic Intensive Workloads	18
Peak Bandwidth Determination	18
Average Bandwidth Determination	18
Image Quality Considerations	18
Capacity Planning Tips	19
PCoIP Session Configuration	21
Session Settings	21
Build-to-Lossless	21

Session Maximum Bandwidth	22
Session Maximum Transmission Unit (MTU)	22
Session Bandwidth Floor	22
UDP Listening Port	24
Session Audio Bandwidth Limit	25
PCoIP Image Quality Levels	25
PCoIP Performance Optimization	27
User Experience Profiles	27
Display Frame Rate Performance	29
PCoIP Ultra Performance	29
Anyware Agent Performance	30
Anyware Software Client Performance	31
Application and OS Optimization	32
Windows Performance Optimizations	32
Windows Performance Options Optimization	32
Improving Performance on Windows Remote Desktops	32
Optimizations related to Windows 10 Visual Effects	37
Adobe Reader Optimization	38
Adobe Premiere Optimizations	38
Remote Access to 3D Rendering Software	38
Linux Environments	39
Windows Environments	39
Enhanced A/V Sync	40
Wacom Tablet Support	40
The Benefits of Local Termination	41
Webcam Support	41
Requirements	41
Notes and Limitations	42
Setup	42

Work-From-Home Support	42
Session Planning with PCoIP Ultra	43
CPU and GPU Offload	43
PCoIP Ultra Auto-Offload	44
Benefits of Auto-Offload	44
PCoIP Ultra System Requirements	44
Client Device Requirements	45
PCoIP Ultra Compatibility Matrix	46
PCoIP Ultra Benchmarks	47
Network Bandwidth Consumption	47
Host CPU Utilization	48

PCoIP Session Planning 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.

The material in this guide is based on the supported features and requirements of **HP Anyware 2023.12**. Earlier versions of HP Anyware, and its components, may not provide the level of support documented here.

In this guide, you will learn about:

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

For complete product and component documentation, visit our [support site](#).

About the PCoIP Protocol

PCoIP is a display protocol that encodes a complete desktop or workload, which is then displayed through a Anyware 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 Anyware 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 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, see help.teradici.com.

Who Uses PCoIP Technology?

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 our website.

About PCoIP Session Planning

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 information and data in this guide will help you best work with your infrastructure's requirements, such as for bandwidth, quality of service, latency, jitter, and packet loss.

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 transition 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 transition unit in network devices is not below the PCoIP packet maximum transition unit size. Defaults are 1200 or 1300 bytes for PCoIP software (depending on the vendor), and 1400 bytes when connecting Anyware zero clients to 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 Knowledge Base article [FAQ - What are the required TCP/UDP ports for PCoIP technology?](#)
- Ensure that PortFast is enabled on all network ports that have Anyware 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, HP Anyware and 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

The following are a list of VPN guidelines that we encourage users to follow:

- The PCoIP protocol provides end-to-end AES-256 encryption of all traffic. When used in conjunction with a security gateway, corporate security compliance objectives can generally be achieved without the need for a VPN. However, PCoIP may be used with UDP-compatible VPN solutions. TCP-based VPNs are not supported.
- If a VPN is required, confirm that UDP traffic is supported (IPsec, or DTLS-enabled SSL solutions). Do not route PCoIP traffic through TCP-based SSL tunnels.
- 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.
- PCoIP does not support packet fragmentation so avoid using a VPN that restricts the MTU for packets transiting the VPN.
- Confirm the VMware ESXi virtual switch traffic shaper is turned off.
- Determine the nature of other protocol traffic that exists on the network, especially other high priority traffic that could impede PCoIP forwarding.
- Ensure the network meets the requirements for real-time protocol deployments, including latency, jitter (latency variation), and packet loss within specified PCoIP limits.
- Optimize your network for virtual desktop connections. For details, see the following KB articles:
 - [What are the recommended host card bandwidth settings? \(KB 1422\)](#)
 - [How do I optimize the Windows desktop experience settings to reduce PCoIP session bandwidth? \(KB 1359\)](#)
- When deploying ZScaler Private Access (ZPA), follow ZScaler recommendations for UC traffic and configure ZPA to bypass IP addresses configured for PCoIP connections. Sending PCoIP traffic to ZPA interferes with PCoIP network adaptation and increases latency. Consider a VPN alternative described below.

Why avoid TCP-based VPN or TCP Display Protocols?

TCP-based protocols (including TCP-based VPNs) provide guaranteed packet delivery. All dropped packets are retransmitted by network layer protocols regardless of their attributes; this increases interactive latency and more seriously, can cause a phenomenon termed *congestion collapse* when multiple connections are streamed over a shared high latency network. Additionally, PCoIP uses packet loss statistics as a feedback parameter for intelligent network bandwidth adaptation, if the

VPN masks the true network capability, bandwidth management is impeded, and the user experience is compromised.

VPN Alternatives

HP Anyware and the Anyware Connector offer highly secure encrypted remote access to enterprise desktops without the need for VPN technology. Cloud Access Manager includes multi-factor Authentication (MFA), secure connection management and a secure network address translation (NAT) gateway allowing public IP addresses for PCoIP traffic. Consult the HP Anyware Architecture Guide for further details on:

- [Managed connections for WAN users connection on-premises](#)
- [Managed Connections for Public Cloud Workstations](#)
- [Managed Connections for Multi-cloud Workstations](#)

Session Establishment

For troubleshooting tips, FAQs and specific documentation around PCoIP Session Establishment, see [FAQ - PCoIP Session Establishment](#) in our knowledge base. This article includes guidelines, troubleshooting checklists, and links to the connection instructions found in component guides.

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

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

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

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
```

Workload Characterization and Network Planning

Prior to production deployment, remote desktop workloads should be assessed in terms of user experience expectations, network bandwidth demands, and performance considerations related to both the host platform (i.e. physical or virtualized desktop) and the Anyware client endpoint. The following section provides a simple classification of user roles and typical workloads which may be used as a guideline on user expectations and assist with bandwidth consumption estimates.

Use Case Classification

PCoIP technology provides a range of capabilities suitable for a variety of user roles, 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 occasional HD video and may have some lightweight 3D applications.
- **CAD or Creative:** Designers and artists including architects, drafters, video editors, game developers and animators perform tasks involving high-quality, graphic-intensive content. Similarly, Scientists and Engineers use compute farms and High Performance Computing (HPC) platforms to extract high resolution visual information from complex data sets.
- **Review:** Broadcast video Quality Control (QC) and post-production finishing teams perform review tasks on color-accurate production grade content with extremely low tolerance for video defects

such as stutter or image artifacts. Similarly, Financial Traders, Analysts and some Healthcare Professionals rely on highly responsive systems that free of visible image artifacts.

	Task Worker	Knowledge Worker	CAD and Creative	Review
Data/Text Entry	✓	✓	✓	✓
Typical per-app window size	Full screen text	Partial screen	Up to full screen	Up to full screen
Office and Web		✓	✓	✓
Video Playback		Occasional	✓	✓
3D Applications		Light	✓	✓
Production quality content playback				✓

Bandwidth Consumption Estimates

The following table provides network bandwidth approximations for different user roles based on a 1920 x 1080 display topology.

Bandwidth Consumption by User Role

User Role	PCoIP Mode	Average Bandwidth	Peak Bandwidth
Task Worker	Default	< 100 kbps	< 1 Mbps
Knowledge Worker	Default	< 1 Mbps	< 10 Mbps
CAD or Creative (30 fps workloads)	PCoIP Ultra Auto-Offload (YUV 4:2:0 Example)	< 5 Mbps	< 50 Mbps
CAD or Creative (60 fps workloads)	PCoIP Ultra Auto-Offload (YUV 4:2:0 Example)	< 10 Mbps	< 100 Mbps
Review	PCoIP Ultra Auto-Offload (YUV 4:4:4 Q90 Example)	< 25 Mbps	< 250 Mbps

NOTES

- Bandwidth values provided are broad approximations only. PCoIP bandwidth depends on protocol configuration, user role, workflow and content characteristics. Higher display resolutions will

increase bandwidth requirements for video or graphics content as illustrated in the [PCoIP Ultra Benchmark Examples](#).

- We recommend that you establish a proof-of-concept (PoC) to characterize your anticipated bandwidth needs before proceeding with a production deployment.

Peak Estimates for Graphic Intensive Workloads

For graphic-intensive workloads, network administrators should consider both average network consumption and the peak bandwidth to sustain high image quality during high frame rate activities such as video playback. PCoIP automatically adjusts image quality and frame rate if insufficient network bandwidth is available.

Peak Bandwidth Determination

Peak bandwidth demand can typically be estimated at 0.5 - 2.0 bits-per-pixel (bpp) compression ratio when using PCoIP default or PCoIP Ultra CPU Offload modes. When using PCoIP GPU Offload or Auto-Offload with 4:2:0 chroma sub-sampling, peak bandwidth can be between 0.1 – 1.0 bpp, dependent on configured image quality and content complexity. In Creative Review cases where high image quality settings (e.g., Q90) are used to preserve detailed screen content such as fine lines and textural detail, a peak bandwidth allocation supporting a compression ratio of 3 - 4 bpp may be needed.

Average Bandwidth Determination

While use case dependent, the average bandwidth is typically less than 10% of the peak demand, considering that PCoIP consumes minimal bandwidth 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 associated with individual connections.

Image Quality Considerations

Peak and average bandwidth consumption can be forcibly constrained by using a session bandwidth limit or non-default quality settings as discussed in the PCoIP Session Variable section. For example,

by adjusting the Maximum Initial Image Quality from Q80 to Q70 when using PCoIP Ultra Auto-Offload, bandwidth savings of approx. 50% may be achieved during video playback.

Capacity Planning Tips

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

- Ensure sufficient bandwidth is allocated to PCoIP traffic by determining appropriate [User Experience Profiles](#) for your user roles, and setting [PCoIP Session Parameters](#) accordingly. For Knowledge Workers using Windows applications, plan for an average of 150–250 kbps for standard office applications.
- Consider any additional bandwidth required for audio input/output and USB traffic. Configure the Anyware Zero Client bandwidth limit parameter to limit USB traffic. Configure a PCoIP Session Audio Bandwidth Limit if needed.
- Plan for 1-2 Mbps per simultaneous user watching a 480p window or 10-13 Mbps for a 1080p video window at default image quality settings. Refer to [PCoIP Ultra benchmarks](#) for additional estimates based on protocol modes and quality settings.
- Ensure that network bandwidth analysis includes observing total network capacity guidelines (for example, peak bandwidth should be 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 configuring a Maximum PCoIP Session Bandwidth Limit when operating in a known constrained network environment. Additional bandwidth savings can be achieved by [optimizing Windows desktop experience settings](#), reducing the display resolution or limiting the number of connected displays.

- PCoIP Ultra in GPU Offload mode configured for 4:2:0 chroma sub-sampling offers additional bandwidth savings for deployments offering hardware video encoder support (e.g. NVIDIA NVENC or macOS hardware). For more information, see [PCoIP Ultra Session Planning](#).

PCoIP Session Configuration

This section lists Anyware agent configuration parameters that directly affect PCoIP session characteristics; there are additional settings that govern other session behaviors.

For configuration instructions specific to your host agent type and operating system, including settings not described here, consult the following product documentation.

Configuration Guides:

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

Session Settings

The following parameters directly affect PCoIP session characteristics and performance. Specific instructions for setting each of these are provided in the documentation listed above.

Build-to-Lossless

This setting specifies whether the protocol will use its *build-to-lossless* feature. This feature is disabled by default in order to conserve bandwidth. When disabled, images may never reach a truly lossless state.

Enable this feature in environments that require images and desktop content to be built to a lossless state, at the expense of increased bandwidth usage.

Session Maximum Bandwidth

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

The default setting is `9000000` kbps.

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. For example, for a single-user VDI configuration (a single PCoIP session) that connects through a 4Mbit/s internet connection, set this value to 4Mbit (or 10% less than this value to leave some allowance for other network traffic).

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.

Session Maximum Transmission Unit (MTU)

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

The default setting is `1200` bytes.

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 MTU can be specified in a range between 500 and 1500 bytes.

The MTU size does not normally need to be changed; you may consider changing it 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.

Session Bandwidth Floor

This setting specifies a lower limit, in kbps, for the bandwidth that is reserved by the PCoIP session.

The default value is 0, so there is no bandwidth floor reserved unless you set one.

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:

- In networks with persistent low-level packet loss, like wireless networks, this setting can be set to a level above the background loss; this will cause the Anyware agent to ignore it, 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 users) *(minimum bandwidth) <= 90% of available link bandwidth*

This setting applies to the agent and client, but the setting only affects the endpoint on which it is configured.

Bandwidth Floor Configuration Examples

User Type	Link Capacity [Mbps]	# Users	Packet Loss	Example Bandwidth Floor Configuration	Corresponding Bandwidth Floor Policy [kbps]
CAD or Creative	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.

UDP Listening Port

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

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.

Session Audio Bandwidth Limit

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

The default setting is 512 kilobits per second (kbps).

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.

PCoIP Image Quality Levels

There are a number of settings within this category, which 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, and the quality of each frame is 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 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 HP Anyware 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 HP Anyware 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 HP Anyware and Remote Workstation Cards.

PCoIP Performance Optimization

If fine tuning of user experience is needed, PCoIP session parameters impacting frame rate, image quality and bandwidth may be tuned via the Group Policy Editor (Windows) or session configuration file (Linux or macOS) as described in the respective Anyware Agent Administrator Guides.

User Experience Profiles

To assist tuning, five user experience profiles and recommended session parameters are provided below as basis for protocol optimization on different networks:

User Experience Profiles

Profile	User Experience	Bandwidth	Network	User Roles
Profile A	Best	Highest	LAN	Review/CAD and Creative
Profile B	Great	Moderate	LAN/WAN	CAD and Creative/Knowledge Worker
Profile C	Good	Optimized	WAN	Knowledge Worker
Profile D	Good	Constrained	WAN	Knowledge/Task Worker
Profile E	Limited	Lowest	Shared WAN	Task Workers with small screens

The table shown next provides recommendations for each of the 5 identified user profiles:

Session Parameters for User Profiles

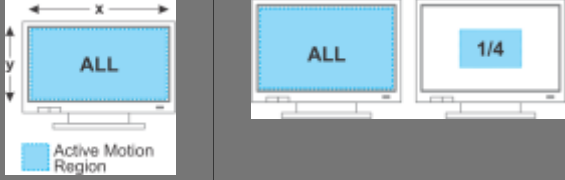
Configurable Session Parameter	A	B	C	D	E
PCoIP Protocol Mode (Graphics Agent)	CPU Offload	Auto-Offload	Auto-Offload	Auto-Offload	N/A
PCoIP Protocol Mode (Standard Agent)	Default	Default	Default	Default	Default
Chroma sub-sampling (Graphics Agent only)	YUV 4:4:4*	YUV 4:4:4*	YUV 4:2:0	YUV 4:2:0	N/A
Maximum PCoIP Session Bandwidth (kbps)	900000*	900000*	10000	2400	600
Enable Build-to-Lossless (BTL) Feature	BTL Enabled	BTL Disabled*	BTL Disabled*	BTL Disabled*	BTL Disabled*
Minimum Image Quality	50	40*	40*	30	30
Maximum Initial Image Quality	90	80*	80*	70	70
Maximum Frame Rate (fps)	60*	60*	30	16	8
Session Audio Bandwidth Limit (kbps)	Disabled*	Disabled*	Disabled*	48	0-32
Windows Visual Effects Optimizations	Best Appearance*	Best Appearance*	Best Performance	Best Performance	Best Performance

*Default Setting

Display Frame Rate Performance

PCoIP technology has been widely adopted by industries that demand the highest possible image quality and frame rate using high resolution displays in multi-monitor configurations. For CAD, visual effects (VFX), video editorial, animation and other graphics-intensive use cases, user satisfaction requires that the PCoIP protocol does not diminish the frame rate associated with the source content compared to local workstation experience. The table below provides example PCoIP frame rates under ideal network conditions.

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 1440 QHD	Up to 60	Up to 30
Software Client*		
1920 x 1080 FHD	Up to 60	Up to 60
2560 x 1440 QHD	Up to 60	Up to 60
3840 x 2560 UHD	Up to 60	Up to 60

(*) HP Anyware configured for PCoIP Ultra protocol enhancements delivers up to 60 fps dependent on host and client CPU specifications and available network bandwidth. PCoIP Ultra CPU Offload requires host and client support for the AVX2 CPU extensions to take advantage of the enhancements. PCoIP Ultra GPU Offload and Auto-Offload require hardware video encoder support on the host system. Refer to the CAS Administrator Guide for specific platform requirements.

PCoIP Ultra Performance

The following section outlines the Anyware Agent and Anyware Software Client performance with PCoIP Ultra modes enabled.

Anyware Agent Performance

PCoIP Ultra CPU Offload

When configured for PCoIP Ultra CPU Offload mode, the host system requires sufficient CPU and GPU resources to support both the host application software (e.g. a video editor tasked with rendering image frames) and CPU cycles required by the Anyware Agent software needed to encode the image frames at the desired frame rate without dropping frames. For best results using 4K/UHD display resolution in conjunction with PCoIP Ultra CPU Offload, at least 4 CPU cores should be provisioned in addition to cores needed for the user applications. For best results using FHD resolution, at least 2 CPU cores are required. Consider setting processor affinity for the PCoIP Server process to partition PCoIP encoding from the application workload if the user application demands high CPU usage.

PCoIP Ultra GPU Offload

When PCoIP Ultra GPU Offload is configured, the host application software utilizes CPU/GPU resources to render image frames, but PCoIP encoding is offloaded to the video hardware encoder of the GPU which frees CPU resources. Refer to the [PCoIP Ultra knowledge article](#) for additional information on GPU configurations.

PCoIP Ultra Auto-Offload

When PCoIP Ultra Auto Offload is configured, the host application software utilizes CPU/GPU resources to render image frames, and PCoIP encoding is offloaded to the video hardware encoder of the GPU whenever the pixel rate exceeds the PCoIP Ultra Offload MMPS threshold defined in the PCoIP session variables. Because Auto-Offload mode only requires CPU resources during low pixel rate encoding periods, Auto Offload mode has approximately the same overall CPU impact as GPU Offload.

Benchmarks

Refer to the [Host CPU Utilization Benchmarks](#) for estimates of PCoIP CPU utilization for different core counts and display resolutions.

Anyware Software Client Performance

PCoIP Ultra CPU Offload

To achieve high pixel rates at high display resolutions (e.g. 4K/UHD at 30 fps or higher), a suitably performing endpoint with AVX2 support is required. For example, at least an Intel 8th (or newer) Generation Core i5/i7 processors with dual memory banks is required to support CAD, Creative or Review workloads at 4K/UHD. An example of this includes NUC10i5FNH (Intel Core i5-10210U processor).

PCoIP Ultra GPU Offload and Auto-Offload

PCoIP Ultra GPU Offload or Auto Offload should be used in conjunction with a Anyware client with H.264 hardware-based decoder support for highest performance, especially when using 4K/UHD displays. We recommend Intel 8th Generation Core i5/i7 processors (or newer). Software clients can achieve higher pixel rates and performance when agents are configured for YUV 4:2:0 to enable hardware decode on supported hardware.

Refer to the [PCoIP Ultra knowledge article](#) for additional information on Anyware Client configuration.

Application and OS Optimization

Windows Performance Optimizations

We recommend 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>

Improving Performance on Windows Remote Desktops

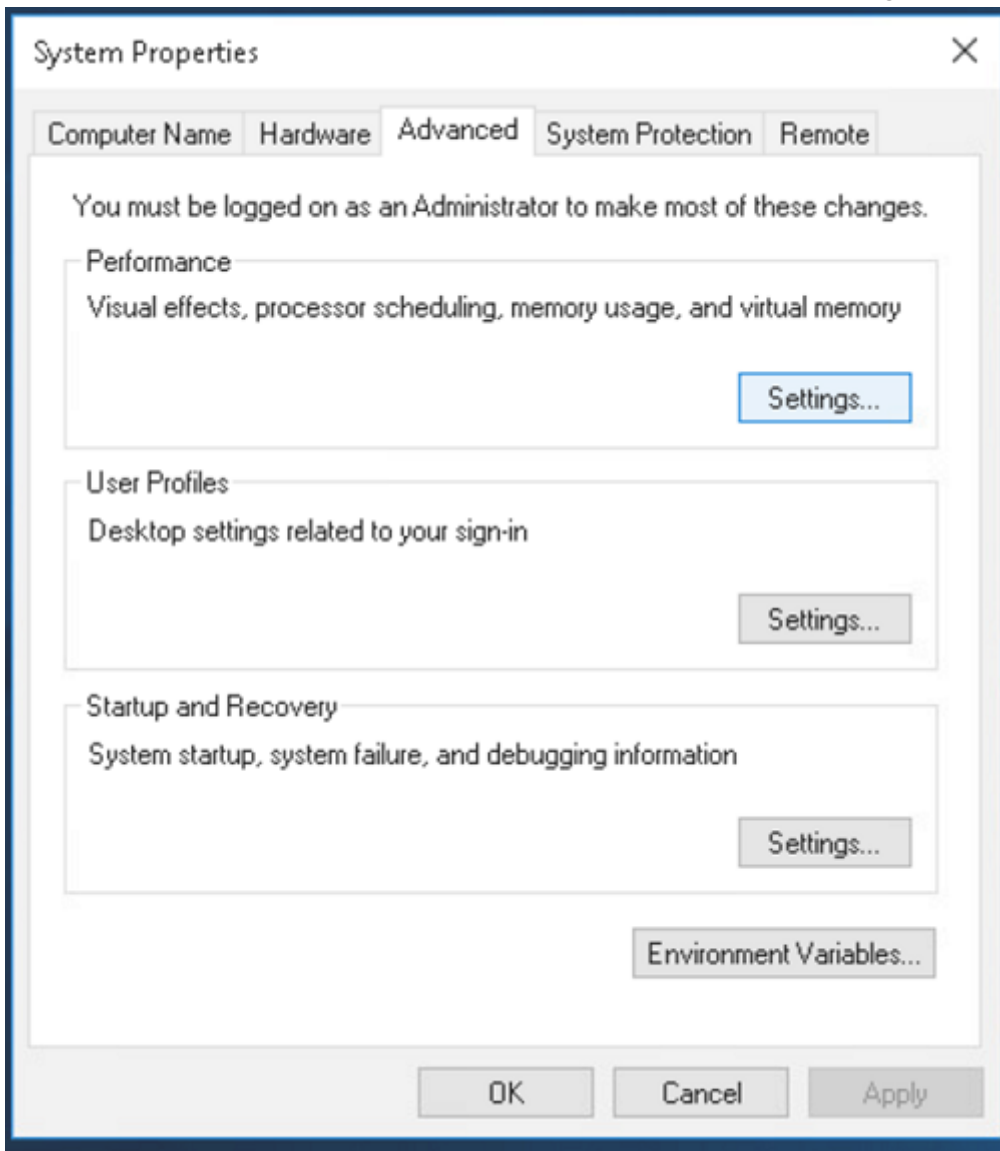
We recommend making updates to the following settings to improve performance for Windows 10 desktops:

- Visual Effects
- Verify Processor Scheduling
- System Page File


The steps involved in configuring these settings are outlined in the section below. Please ensure you save any work, and close all programs before making these changes.

1. Navigate to *View advance system settings* within the Windows system Control panel.

2. From the Performance box within the Advanced tab, select **Settings**.

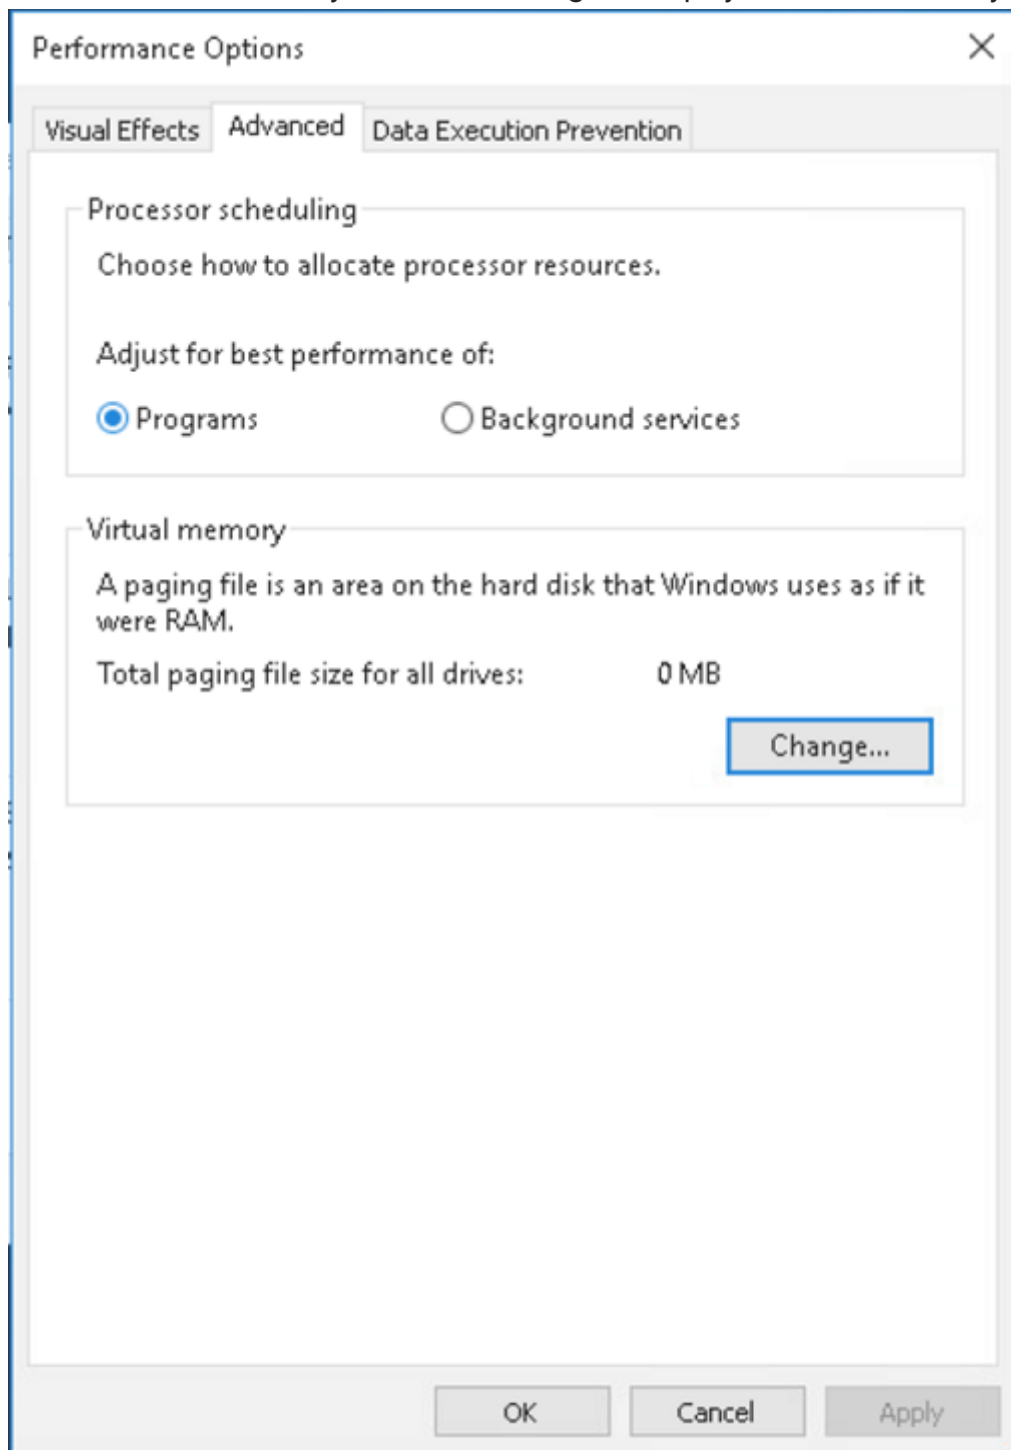


3. From the Performance Options page select **Adjust for best performance** and click **Apply**.



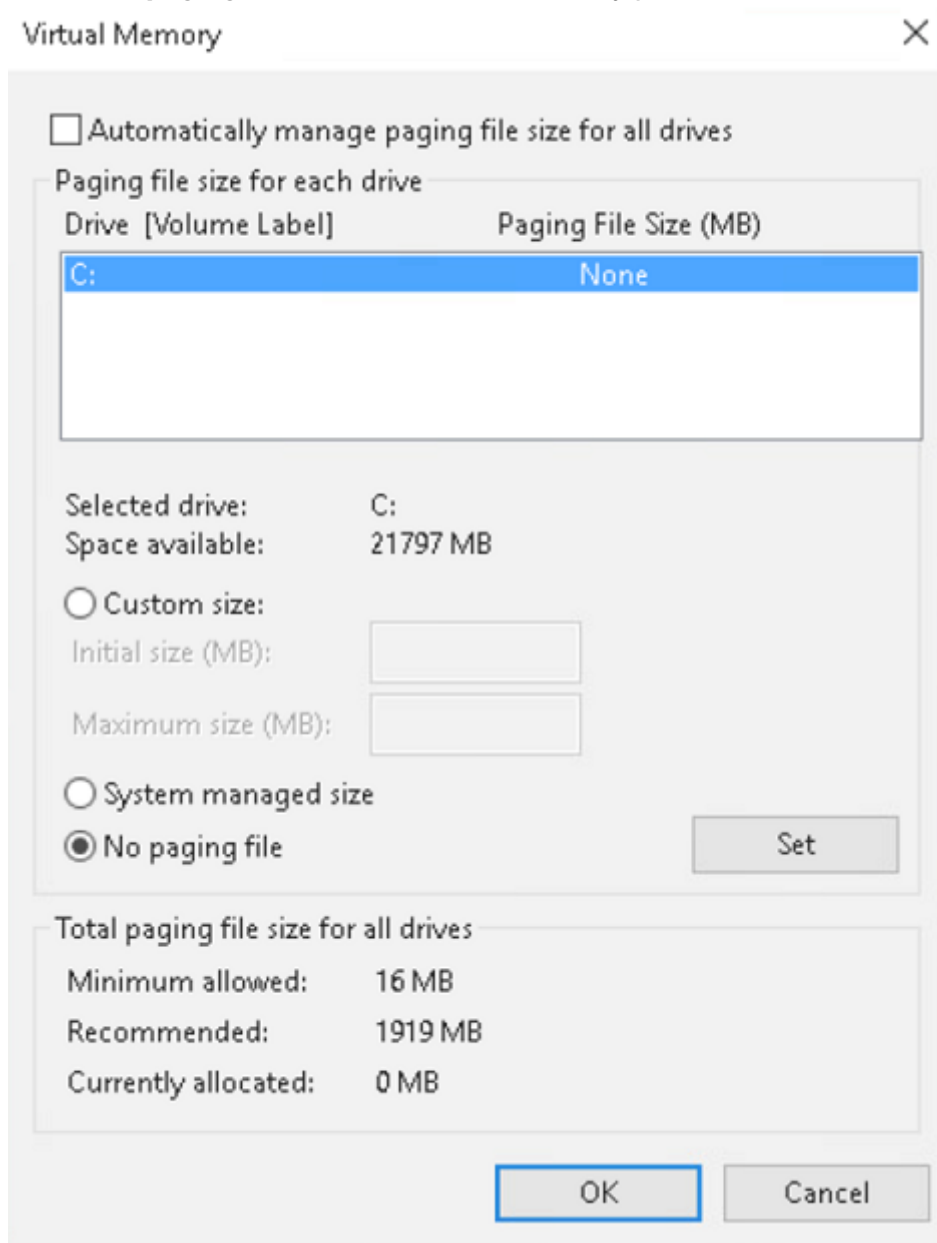
4. From the Performance Options page select the Advanced tab and click **Programs** is selected within the Processor scheduling box.

5. From the Virtual memory box select **Change** to display the Virtual Memory panel.



6. From the Virtual Memory panel, uncheck the option to **Automatically manage paging file size for all drives**.

7. Select **No paging file** from the Virtual Memory panel, as outlined in the image below.

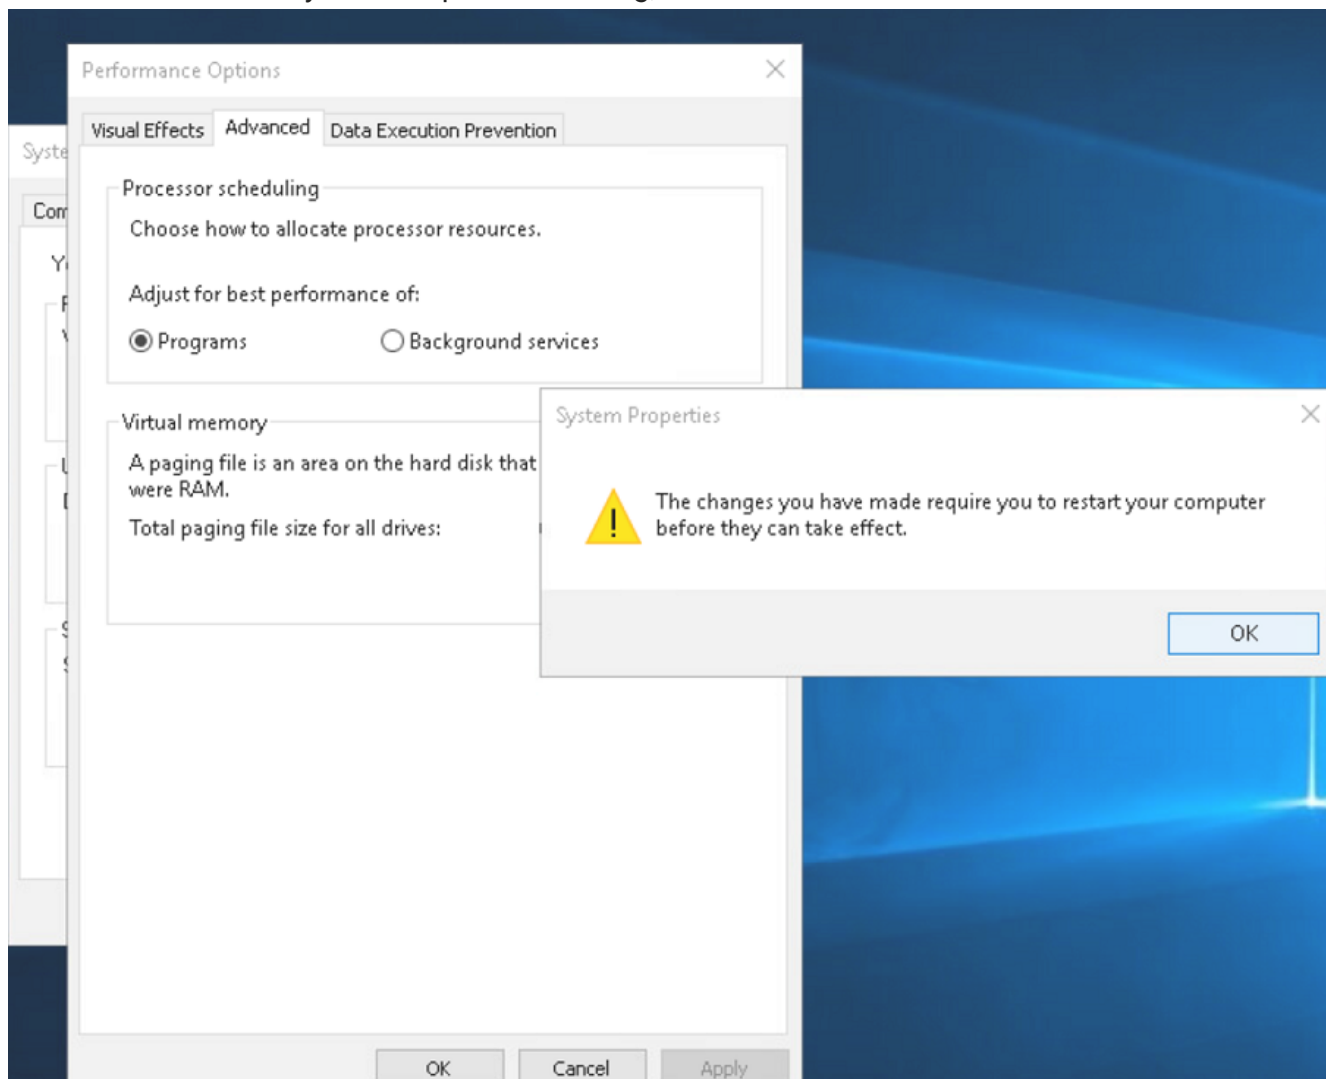


8. From the Virtual Memory panel, click **Set**.

9. You will see a System Properties warning, click **Yes**.

10. From the Virtual Memory panel, click **OK**.

11. You will see another System Properties warning, click **OK**.



12. Return to the Performance Options page and to the Advanced tab and click **Apply** and then **OK**.

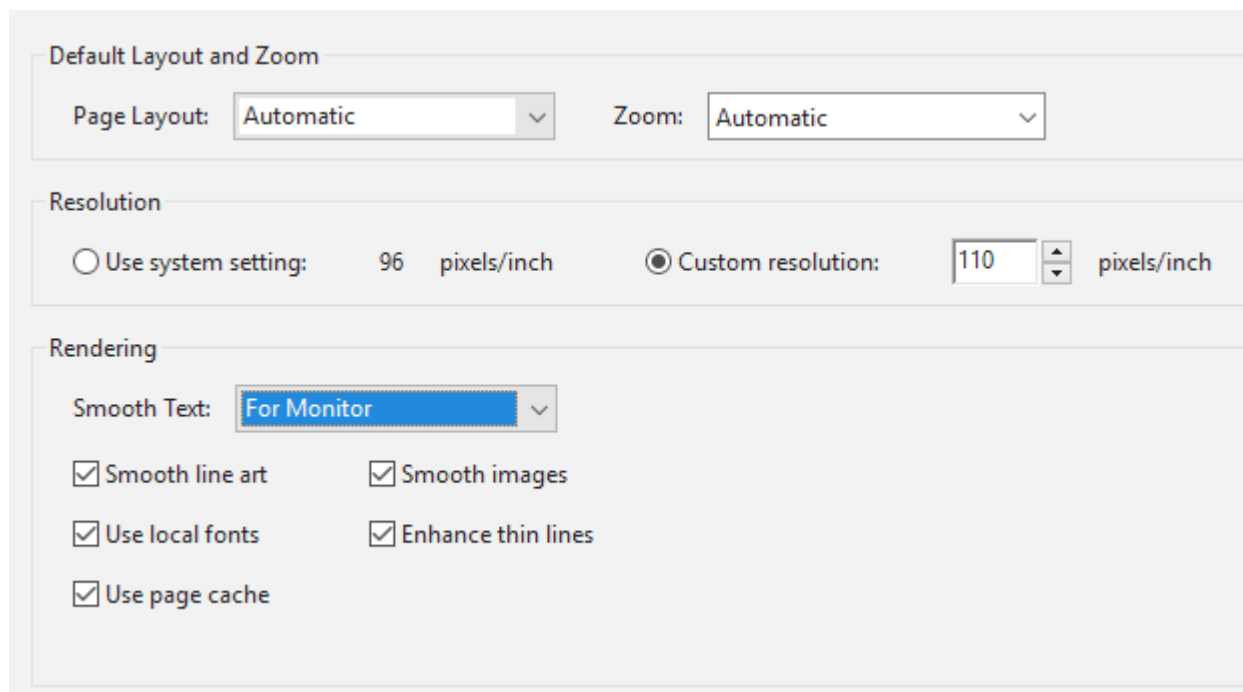
13. Restart your system for these changes to take effect.

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 [C, D and E](#), we recommend that you disable the 'Show window contents while dragging' policy. Additional performance optimizations are gained by adjusting Windows Visual Effects for best performance as described below.

Adobe Reader Optimization

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



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.

[HP Anyware](#) 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 Anyware 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 A/V Sync

This feature reduces the delay between the audio and video channels. It is enabled in the Anyware Client application. 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 Anyware Software Clients, see the following guides:

- [Anyware Software Client for Windows](#)
- [Anyware Software Client for macOS](#)
- [Anyware 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

HP Anyware 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 Anyware Software Client or a Anyware Tera2 Zero Client. Some of the advantages of using HP Anyware in conjunction with Wacom tablets include:

- The ability to interact with your Wacom tablet in a remote environment as if using a deskside 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 Anyware Software Client and Anyware Tera2 Zero Client, see [Locally Terminated Wacom Tablets - Windows](#), [Locally Terminated Wacom Tablets - Linux](#) and [Locally Terminated Wacom Tablets - macOS](#).

Webcam Support

The Anyware Software Client for Windows supports certain Webcam models. You must have a Anyware Software Client 21.03 or higher and Anyware Agent 21.03 or higher installed to enable these Webcam models. For detailed information on which models are supported and the performance metrics associated with these models, see [here](#). This KB article also details steps on how to test and verify other Webcam models.

Webcams may consume significant network bandwidth in the upstream direction from the endpoint with the webcam to the remote desktop. Depending on the webcam model and resolution, the bandwidth requirement to achieve ~ 30 frames per second may be anywhere between 17 Mbps and 40 Mbps. Note that if network bandwidth availability is limited, the frame rate of the webcam is scaled back accordingly.

This feature is enabled by default.

Requirements

Webcam support in HP Anyware requires the following:

- Anyware Software Client for Windows or Anyware Software Client for Linux, 21.07+
- Anyware Standard Agent for Windows or Anyware Graphics Agent for Windows, 21.03+
- USB-attached webcam

Notes and Limitations

- Webcams must be connected via USB. Webcams that are not USB, such as embedded laptop webcams, are not supported.
- Linux agents are not supported.
- Anyware Software Client for macOS is not supported.

Setup

On the Anyware Software Client, connect the webcam as described in the following sections:

- Anyware Client for Windows: [USB Bridging of Webcams](#).
- Anyware Client for Linux: [USB Bridging of Webcams](#).

Work-From-Home Support

HP Anyware 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.

HP Anyware 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 HP Anyware](#).
- [Work-from-Home options for VMware Horizon](#).
- [Performance Tips for Work-from-Home Use Cases](#).

Session Planning with PCoIP Ultra

PCoIP Ultra offers three modes that optimize system performance dependent on use case demands:

- CPU Offload for high resolution content at highest image quality.
- GPU Offload for CPU relief.
- Auto Offload provides both build-to-lossless image quality and optimum bandwidth and CPU savings for dynamic content.

This section provides video playback benchmarks to help you plan your PCoIP Ultra deployment.

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

Requirement	CPU Offload - AVX2 RGB 8 8 8	GPU Offload - NVENC YUV 4:4:4	GPU Offload - NVENC YUV 4:2:0
Highest Image Quality Desktop (Up to 25x16 at 60 fps or 4K/UHD @ 30fps)	✓		
CPU Relief e.g. VFX Render Applications (Up to 25x16 at 60 fps, 4K/UHD)			✓
Bandwidth Preservation (e.g. constrained WAN)			✓
Highest frame rate from a low-cost thin client			✓

CPU and GPU Offload

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 Auto-Offload

The auto-offload feature with PCoIP Ultra enables users to allow PCoIP Ultra to select the best protocol, whether that is CPU or GPU, based on display rate change. CPU Offload is used by default to provide the best image fidelity, GPU Offload is used during periods of high display activity to provide improved frame rates and bandwidth optimization. This setting is only effective if the remote host and client endpoints are capable of both CPU and GPU Offload. You can select this option when you are enabling PCoIP Ultra.

For information on how to do this, see [PCoIP Ultra - Windows](#), and [PCoIP Ultra - Linux](#).

Benefits of Auto-Offload

Using the auto-offload feature ensures that the best codec is used and enables a solid quality vs efficiency balance. Depending on the amount of display change PCoIP Ultra CPU or PCoIP Ultra GP will be used. If there is a large amount of change picked up the protocol will transition to PCoIP Ultra GPU to allow for higher frame rates and lower bandwidth. Auto-offload also ensures:

- No per-workload admin tuning required.
- No burst stealing of CPU or GPU cycles.
- No transitional image artifacts.
- No aggressive network consumption.

PCoIP Ultra System Requirements

To take advantage of PCoIP Ultra, you need:

- A **Anyware Agent** (any type), 19.11.0 or later
- A **Anyware Software Client** (any type), 19.11.0 or later

⚠ Anyware Tera2 Zero Clients do not support PCoIP Ultra

PCoIP Ultra is only available through Anyware Software Clients.

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

For more information the requirements for PCoIP Ultra, see:

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

Client Device Requirements

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

- **PCoIP Ultra with CPU Offload enabled:** Intel 8th Generation Core i5/i7 processors with dual memory banks are required to achieve 4K30 FPS. An example of this includes NUC10i5FNH (Intel Corei5-10210U processor).
- **PCoIP Ultra with GPU Offload enabled:** On Linux and Windows clients, hardware decoding is available for GPU Offload when YUV 4:2:0 chroma-subsampling is enabled. Software-based H.264 decoding is used in YUV 4:4:4 mode. We recommend Intel 8th Generation Core i5/i7 processors or newer with dual memory banks.

For more information on the Technical Requirements required for PCoIP Ultra, see [here](#).

PCoIP Ultra Compatibility Matrix

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

	Windows Client	Linux Client	macOS Client
Anyware Graphics Agent for Windows or Linux <i>Nvidia NVENC or AMD RapidFire</i>	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 software decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 hardware decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder or GPU Offload with H.264 software decoder.
Anyware Graphics Agent for Windows <i>Nvidia GeForce graphics</i>	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.
Anyware Standard Agents for Windows or Linux	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.	PCoIP Ultra: CPU Offload with PCoIP AVX2 decoder.

Important: Anyware Tera2 Zero Clients are not supported

PCoIP Ultra is supported only by Anyware Software Clients for Windows, Linux, and macOS.

PCoIP Ultra Benchmarks

The following benchmarks give some idea of the bandwidth consumption used with PCoIP Ultra sessions. The benchmarks were generated using the following system configuration:

	Description
Host Platform	Supermicro SYS1019GP-TT (Xeon Gold 6248 2.5 GHz, 20 cores)
Hypervisor	VMware ESXi 6.7
Virtual Machine	Windows 10, 96 GB RAM, grid_rtx6000_12q, HP Anyware 21.07.4
Graphics	NVIDIA RTX 6000, GRID 11.1
Network	1 Gbps LAN
Client Endpoint	Intel NUC10FNHi5, Windows 10
Display	4K/UHD 3840 x 2160
Anyware Agent	Version 21.07.4
Anyware Client	Version 21.07
Video Content	Big Buck Bunny 1080p 24fps (Opening 2 minutes)

Network Bandwidth Consumption

The network bandwidth used in conjunction with PCoIP Ultra depends on several factors, including PCoIP Ultra mode (CPU Offload, GPU Offload or Auto-Offload modes), display resolution, and configured image quality policy.

This table shows the average bandwidth consumption for the opening scene of Big Buck Bunny with the default PCoIP Image Quality setting (Q80) and at the following resolutions: native 1080p, 1080p scaled to 1440p, and full-screen 2560p (4K/UHD) playback.

Bandwidth Consumption in Mbps

	Auto-Offload* (YUV 4:2:0)	Auto-Offload* (YUV 4:4:4)	CPU Offload
1080p	13.0	19.0	34.4
1440p	17.3	27.5	43.1
2560p	25.0	43.8	61.4

* PCoIP Ultra GPU-Offload mode consumes comparable bandwidth to PCoIP Ultra Auto-Offload mode.

On constrained networks, the PCoIP protocol dynamically adjusts image quality and frame rate to suit the available network bandwidth. In cases where it is preferred to proactively constrain the bandwidth, as might be required to limit service provider network charges, PCoIP image quality and frame rate can be capped. Refer to [PCoIP settings section](#) for policies that enforce bandwidth and/or frame rate limits.

The following table provides an example of network bandwidth savings and consumption using PCoIP Ultra Auto-Offload at different image quality settings. These rates are at .

Bandwidth Consumption using Auto-Offload (in Mbps)

	Q60	Q70	Q80	Q90
1080p Auto-Offload (YUV 4:2:0)	3.7	6.6	13.0	34.8

Host CPU Utilization

PCoIP Ultra Auto Offload mode offers the highest host CPU efficiency by leveraging NVIDIA NVENC technology for display processing when the encoded pixel rate exceeds a programmed threshold.

The table below shows host virtual machine average CPU utilization when using PCoIP Ultra Auto-Offload and PCoIP Ultra CPU-Offload as measured by the VMware ESXi Performance Monitor at default image quality settings for various virtual machine vCPU allocations.

CPU Utilization with Auto-Offload and CPU-Offload

	4 vCPU	8 vCPU	16 vCPU	24 vCPU
1080p (Auto-Offload)	17%	10%	4%	2%
4K/UHD (Auto-Offload)	19%	11%	4%	3%
1080p (CPU-Offload)	38%	18%	9%	2%
4K/UHD (CPU-Offload)	58%	28%	14%	3%

PCoIP Hyper-Threading

PCoIP Ultra CPU Offload does not distinguish between real and virtual cores. PCoIP will take advantage of hyper-threading but N hyper-threaded cores do not provide the same performance as the same number of N physical cores.

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