

Nuance[®]
Dragon[®]
Professional
Anywhere

Audio Routing Solutions
in Virtualized Environments



Table of contents

Overview	3
System configurations	4
Single-hop virtualization	4
Double-hop virtualization	5
Available solutions	7
Nuance virtual extensions	7
Vendor native audio channels	8
USB redirection	9
Recommendations	10
Zero clients	10
Single-hop configuration (thin or thick clients)	10
Double-hop configuration (thin or thick clients)	10
Recommendations overview	11
References	12

Overview

For accurate speech recognition, Dragon Professional Anywhere must capture clear audio from the user's device. To achieve this in a virtualized environment, it requires the following considerations:

- The client workstation, where the user is dictating, is not running a locally-installed instance of Dragon Professional Anywhere or the target application. The applications are running on a virtualization server or virtual desktop. The workstation displays a bitmap representation of the applications or virtual desktop via a receiver application.
- Dragon Professional Anywhere must be installed on the same server or virtual desktop as the target application to access the application's text fields.
- Button controls and high-quality audio must be routed from the workstation where the PowerMic is connected to the virtualization layer where Dragon Professional Anywhere and the target application are running.

This guide summarizes the types of virtualized environments at customer sites and the available Nuance solutions. It helps Nuance customer-facing personnel to recommend the best solution for each customer's configuration.

Note: Many sites install Dragon Professional Anywhere locally in addition to a virtual configuration; local installations are not the primary focus of this document.

System configurations

Dragon Professional Anywhere and the target application can be hosted on a virtualization server (for example, Citrix XenApp) or included in a virtual desktop image running on a virtualization server (for example, Citrix XenDesktop).

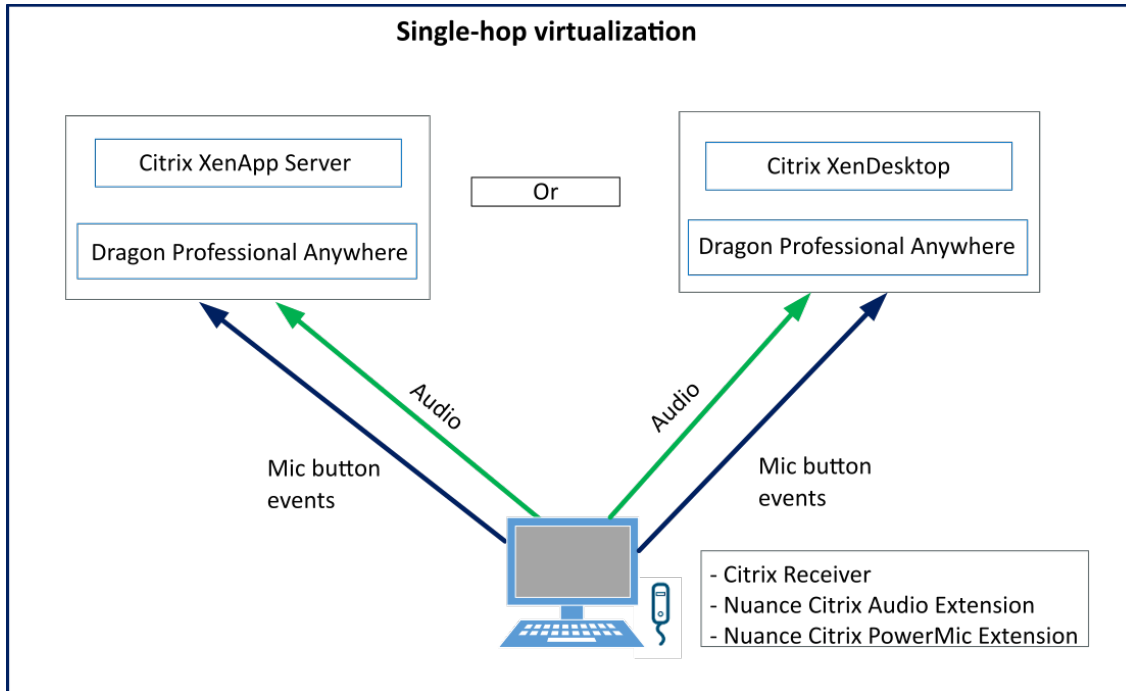
The workstation can be a thick client running a full Windows version, a thin client running Linux or Windows Embedded or a zero client with no operating system.

A site's configuration might include a virtualization vendor for which Nuance provides custom channels (Citrix, VMware); this affects the choice of solution for the site.

A customer might use more than one of these configurations. Therefore, more than one Nuance solution might be needed to fulfill a site's requirements.

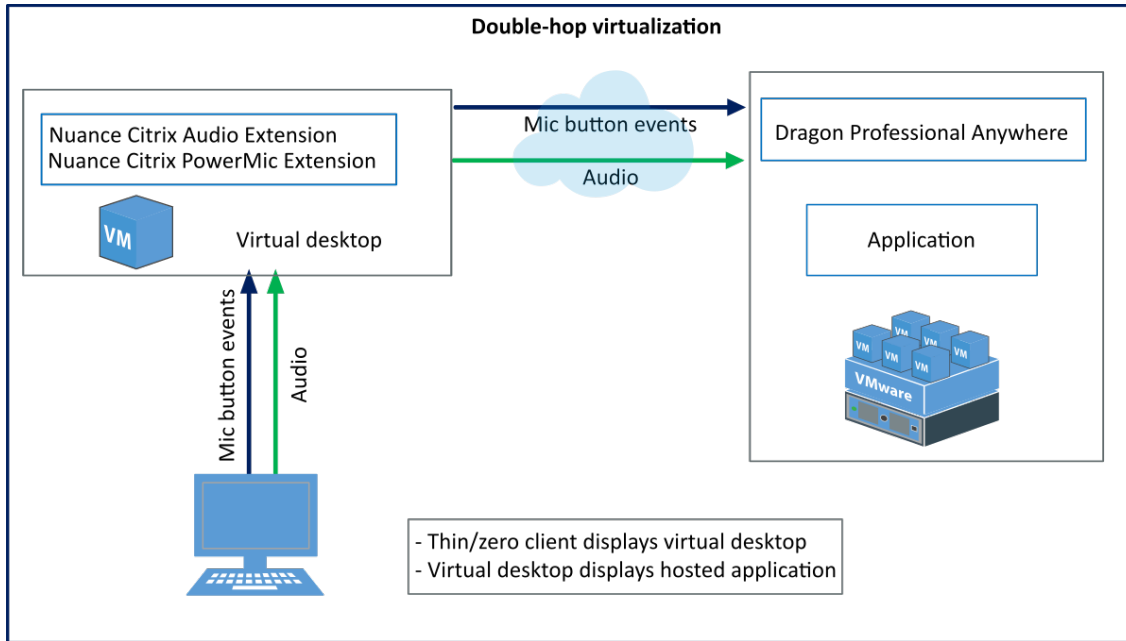
Single-hop virtualization

Dragon Professional Anywhere and the target application are hosted on a virtualization server or included in a virtual desktop. The applications/virtual desktop are streamed to the user's workstation. Audio and microphone buttons must be routed from the workstation to the virtualization server/virtual desktop where Dragon Professional Anywhere and the target application are installed.



Double-hop virtualization

A double-hop configuration is most common when a site accesses an application hosted by a vendor and also uses a desktop virtualization solution such as XenDesktop. In this configuration, Dragon Professional Anywhere is hosted with the target application. Audio and microphone buttons must be either routed through two virtualization layers or delivered via an alternative route that bypasses the intermediate virtualization layer.



For more information on double-hop configuration for Citrix and VMware environments, see *Nuance - Double Hop Configuration for Citrix Environments* in the Docs & Guidelines folder of the Nuance virtual extensions package.

Available solutions

Nuance virtual extensions

The Nuance Citrix and VMware audio extensions and PowerMic extensions provide custom audio and microphone button channels for the following products:

- Citrix XenApp and XenDesktop/Citrix Virtual Apps and Desktops
- VMware Horizon View
- Microsoft Windows 8.1 and 10 and Windows Embedded operating systems
- Linux thin clients (for more information on thin clients supporting device splitting, contact your vendor)
- The following speech recognition applications:

Dragon Professional Anywhere

Applications based on Dragon Medical SpeechKit (.NET and COM editions)

Audio extensions

The Nuance Citrix/VMware audio extensions reduce the bandwidth required to transfer audio data from the client workstation to the server/virtual desktop where Dragon Professional Anywhere is installed. Native audio channels can require up to 1.4 Mbit/s bandwidth between the client workstation and the hosted application. The Nuance audio extensions require 27 kbit/s. This enables Nuance to offer compatibility with the most common virtualization solutions while at the same time delivering industry-leading speech recognition accuracy.

The Nuance audio extension is only installed on the client workstation (client end point).

PowerMic extensions

Microphone buttons, sliders and other controls must be routed to Dragon Professional Anywhere separately from audio. To enable this for the PowerMic II/III in a Citrix XenApp/XenDesktop or VMware system, the Nuance PowerMic extensions provide a custom channel for PowerMic button controls.

The Nuance PowerMic extension is only installed on the client workstation (client end point).

Vendor native audio channels

Important: To use PowerMic buttons, you need to configure device splitting (available for Citrix and VMware environments, Windows client end points and some Linux-based clients). For more information on thin clients supporting device splitting, contact your vendor.

Citrix HDX

Citrix HDX audio channel (Citrix optimized virtual channel) redirects audio data to the virtual desktop at a significantly lower bandwidth than regular USB redirection.

VMware RTAV

VMware Horizon View Real Time Audio and Video (RTAV) redirects audio and video data to the virtual desktop at a significantly lower bandwidth than regular USB redirection.

Limitations

If you use native audio channels instead of the Nuance audio extensions, you might experience the following limitations:

- The bandwidth increases from 28 kbit/s to 150 kbit/s or higher.
- The microphone input volume cannot be changed on the client end point; this might lead to audio gaps if the volume is not configured correctly.
- Automatic gain control (AGC) and voice activity detection (VAD) are not available.
- Controlling the client end point is not possible (for example, disabling the screen saver during recording or disabling the buttons/stopping recording when the screen is locked).
- In standby mode, the audio data will be streamed continuously from the client end point to the server or virtual desktop. When recording is on, the audio data will also be streamed even if the user is not speaking.

- If Citrix HDX is used in combination with USB device splitting, each time a user roams to another client end point, they must manually select the HID device they want to redirect to be able to use the microphone buttons.

USB redirection

Citrix XenApp/XenDesktop 7.15+ and VMware View 7.6+ can redirect the USB device from the client workstation to the server/virtual desktop. PowerMic II/III bandwidth consumption via USB redirection is approximately 440 kbit/s in Citrix environments and approximately 1 Mbit/s in other virtual environments. There can be limitations depending on the virtualization platform and version used by a site. For example, VMware Remote Desktop Services Host (RDSH) does not support USB redirection on the server side.

Note: USB redirection requires high bandwidth and is not optimized or reliable; it is preferred to use the [Nuance virtual extensions](#), which are optimized for audio data.

Use native audio channels and device splitting in the following cases:

- When using a PowerMic II/III with a zero client or with a thin client not supported by the Nuance virtual extensions.
- When using a double-hop configuration.

Recommendations

Note: More than one Nuance solution might be needed to fulfill a given site's requirements.

Zero clients

Alternative solution: [Vendor native audio channel](#).

- Should only be considered if the site needs to use PowerMic or composite devices from other microphone vendors.
- Does not need additional software to be installed on the client workstation.
- The audio is compressed to optimize network bandwidth and encrypted for security (if enabled).

Remarks

- Nuance virtual extensions cannot be installed on zero clients.

Single-hop configuration (thin or thick clients)

Recommended solutions:

- [Nuance virtual extensions](#) if the site uses a supported virtualization vendor.
- Vendor native audio channel if the virtualization vendor is not supported by the Nuance virtual extensions.

Double-hop configuration (thin or thick clients)

Recommended solutions: Combination of vendor native audio channel and Nuance virtual extensions.

Recommendations overview

Technology	Bandwidth	Recommended for
Nuance virtual extensions	27 kbit/s	Microsoft Windows and Windows Embedded operating systems Igel thin clients eLux-based thin clients
Citrix HDX	Approximately 150 kbit/s	Microsoft Windows operating systems Thin and zero clients Note: We recommend setting the audio quality option to High for optimal speech recognition quality.
VMware RTAV	> 200 kbit/s (depends on the VMware settings)	Microsoft Windows operating systems Thin and zero clients
Citrix USB audio	500 – 1000 kbit/s	Microsoft Windows operating systems Thin and zero clients
VMware USB audio	500 – 1000 kbit/s	Microsoft Windows operating systems Thin and zero clients

References

For more information, see the following:

- *Nuance - Deployment and Configuration for Citrix Environments*
- *Nuance - Deployment and Configuration for VMware Horizon View Environments*
- *Nuance - Deployment and Configuration for Microsoft RDS Environments*
- [IGEL Knowledge Base](#)