

# Teradici APEX™ 2800 Server Offload Card Functional Specification

TER1201001-2



Teradici Corporation  
#101-4621 Canada Way, Burnaby, BC V5G 4X8 Canada  
p +1 604 451 5800 f +1 604 451 5818  
www.teradici.com



---

The information contained in this document represents the current view of Teradici Corporation as of the date of publication. Because Teradici must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Teradici, and Teradici cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. TERADICI MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Teradici Corporation.

Teradici may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Teradici, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property. Visit [Teradici.com/pat](http://Teradici.com/pat) for additional information.

© 2012 Teradici Corporation. All rights reserved.

Teradici, PC-over-IP, and PCoIP are registered trademarks of Teradici Corporation.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

## Revision History

Version	Date	Description
2	Feb 24, 2012	Included support for ESXi 5.0
1	Jan 31, 2012	Initial release

## Contents

Contents .....	5
Table of Figures .....	6
Definitions .....	7
1 Introduction .....	8
1.1 Benefits .....	8
2 Product Overview .....	9
3 System Requirements .....	10
3.1 Virtual Desktop Server Platform .....	10
3.2 Client .....	10
3.3 Network .....	10
4 Functional Overview .....	11
4.1 Solution Components .....	11
4.1.1 APEX 2800 Card Hardware Characteristics .....	12
4.2 APEX 2800 Software Packaging and Delivery .....	14
4.2.1 APEX 2800 Card VIB Package (ESXi Driver) .....	14
4.2.2 APEX 2800 VM Software Package .....	15
4.3 About the Display Offload Feature .....	15
4.3.1 Display Rendering .....	16
4.3.2 Display Imaging Activity Monitoring .....	16
4.3.3 Display Offload Decision .....	16
4.3.4 Display Offload Priority .....	17
4.3.5 Display Offload Canceling .....	17
4.3.6 Display PCoIP Encoding/Compression .....	17
4.3.7 Per-Pixel Display Image Decomposition .....	18
4.3.8 Pixel-based Imaging Motion Compensation .....	18
4.4 Networking .....	18
4.5 USB .....	19
4.6 Audio .....	19
4.7 VMware vMotion Support .....	19

## Table of Figures

Figure 7-1: APEX 2800 Card Solution Software Modules.....	12
Figure 7-2: APEX 2800 Server Offload Card.....	13
Figure 7-3: APEX 2800 Server Offload Card Mechanical Drawing.....	14
Figure 7-4: VM Hardware Acceleration and Priority Settings.....	16

## Definitions

Definition	Description
APEX	Accelerated PCoIP EXperience
CPU	Central Processing Unit
ECC	Error Correction Code
FPS	Frames Per Second
GPU	Graphics Processing Unit
GUI	Graphical User Interface presented by the client On Screen Display when not operating in a PC-over-IP session
Mpps	Mega pixels per second
OS	Operating System
OSD	On Screen Display
PCIe	Peripheral Component Interconnect Express
PC-over-IP <sup>®</sup>	Personal Computer over Internet Protocol
PCoIP <sup>®</sup>	Personal Computer over Internet Protocol (PC-over-IP)
PCoIP Zero Client	Desktop side of PC-over-IP system (that is, client). For example, PCoIP zero client or PCoIP integrated display
PCoIP Host	Host side of PC-over-IP system
QoS	Quality of Service
vCPU	Virtual CPU within a virtual machine
VDI	Virtual Desktop Infrastructure
VIB	VMware vSphere Installation Bundle
VM	Virtual Machine
VPN	Virtual Private Network
Zero Client	See PCoIP zero client

# 1 Introduction

The Teradici APEX™ 2800 server offload card provides hardware-accelerated PCoIP® image encoding for server-hosted VMware View virtual desktops (often referred to as Virtual Desktop Infrastructure or VDI). Available as a standard PCI Express (PCIe) expansion card for industry-standard servers, the APEX 2800 card monitors the graphical demands of VDI displays and automatically offloads the PCoIP encoding tasks of the most demanding displays from the vCPUs to the APEX 2800 card.

As demands change, the card seamlessly and automatically shifts between hardware encoding on the APEX 2800 card and software encoding on the virtual desktops' vCPU(s). This lets the APEX 2800 card support a large number of virtual desktops while providing optimal use of resources. The APEX 2800 server offload card can offload 64 displays running at a maximum resolution of 1920 x 1200 and can monitor up to four displays per virtual machine.

**Note:** Each virtual machine display is monitored separately. The number of displays offloaded may vary based on screen resolution and display activity.

## 1.1 Benefits

- Protected end-user experience from pixel-encoding CPU peaks
- Increased VDI consolidation ratios per server
- Full integration into VMware View and managed through View Administrator
- Simple install and setup

For more details, see <http://www.teradici.com/pcoip/pcoip-products/pcoip-server-offload-card.php>.

## 2 Product Overview

The Teradici APEX 2800 server offload card has the following features and capabilities:

- Offloads PCoIP protocol encoding tasks of virtual desktop displays for server CPU
- Works with VMware View
- Support for 64 displays: Simultaneous PCoIP protocol encoding for offloaded displays
- Support for display processing up to 300 Mpps (Mega pixels per second)
- Dynamic imaging resource management:
  - Ensures the most active 64 displays are offloaded
  - Maximizes the card efficiency
- Per-pixel image decomposition: Ensures the right compression algorithm is used for the image content
- Pixel-based image motion compensation:
  - Optimizes network bandwidth utilization
  - Operating System (OS) and application independent
- PCoIP protocol network bandwidth management: Automatically adjusts image quality based on available network resources
- Card system memory
  - Two GB DDR3 to support 64 displays at a maximum resolution of 1920 x 1200
  - Error Correction Code (ECC) memory
- Card LED indicators
  - Power LED – solid green indicates stable power
  - Heartbeat LED – blinks twice per second for normal operation

## 3 System Requirements

The APEX 2800 card can be used in virtual desktop environments that meet the following system requirements.

### 3.1 Virtual Desktop Server Platform

- Server platform with an available PCI Express (PCIe) x8, full height, half length, compatible card slot
- ESXi 4.1 Update 1 build 348481, 4.1 Update 2 build 502767 or ESXi 5.0
- VMware View 4.6 or View 5.0
- Supported virtual machine operating systems:
  - Windows 7 32-bit/64-bit
  - Windows Vista 32-bit/64-bit
  - Windows XP
- Maximum of two APEX 2800 cards per server for 128 displays being offloaded

To get the latest drivers, register at <http://techsupport.teradici.com>

### 3.2 Client

- PCoIP zero clients with firmware 3.5 or later
- VMware View 4.6 clients or later
- Displays with a resolution of 1920 x 1200 or less (displays with a higher resolution setting are not eligible for offload with release 1.0)

### 3.3 Network

The PCoIP protocol provides a real-time delivery of a rich user desktop experience. To ensure a responsive desktop, the PCoIP protocol must be deployed across a properly architected virtual desktop network infrastructure that meets bandwidth, QoS, latency and packet-loss requirements. For more details, see the *PCoIP Protocol Virtual Desktop Network Design Checklist* TER1105004 document at <http://techsupport.teradici.com>.

## 4 Functional Overview

The APEX 2800 card is a host server add-in card that offloads PCoIP protocol processing for VMware View virtual machine (VM) servers, reducing the server CPU utilization. The APEX 2800 card can be used in standard rack/ tower servers or in blade server “side cars”. From a high level, the APEX 2800 card:

- Monitors VM display(s) imaging and PCoIP encoding load
- Provides dynamic offload selection to determine which displays to offload based on priority and imaging load
- Seamlessly switches between PCoIP software and APEX 2800 hardware encoding
- Transfers display pixels from the VM to the APEX 2800 card
- Simultaneously encodes hardware accelerated PCoIP protocol for all displays
- Transfers PCoIP protocol encoded pixels from the APEX 2800 card to the VM
- Delivers PCoIP protocol data to the soft client for decode
- Manages APEX 2800 card imaging resources
- Determines which displays require offload canceling (returns these displays to PCoIP software encoding)

**Note:** This only applies if more than 64 displays are being addressed by the server.

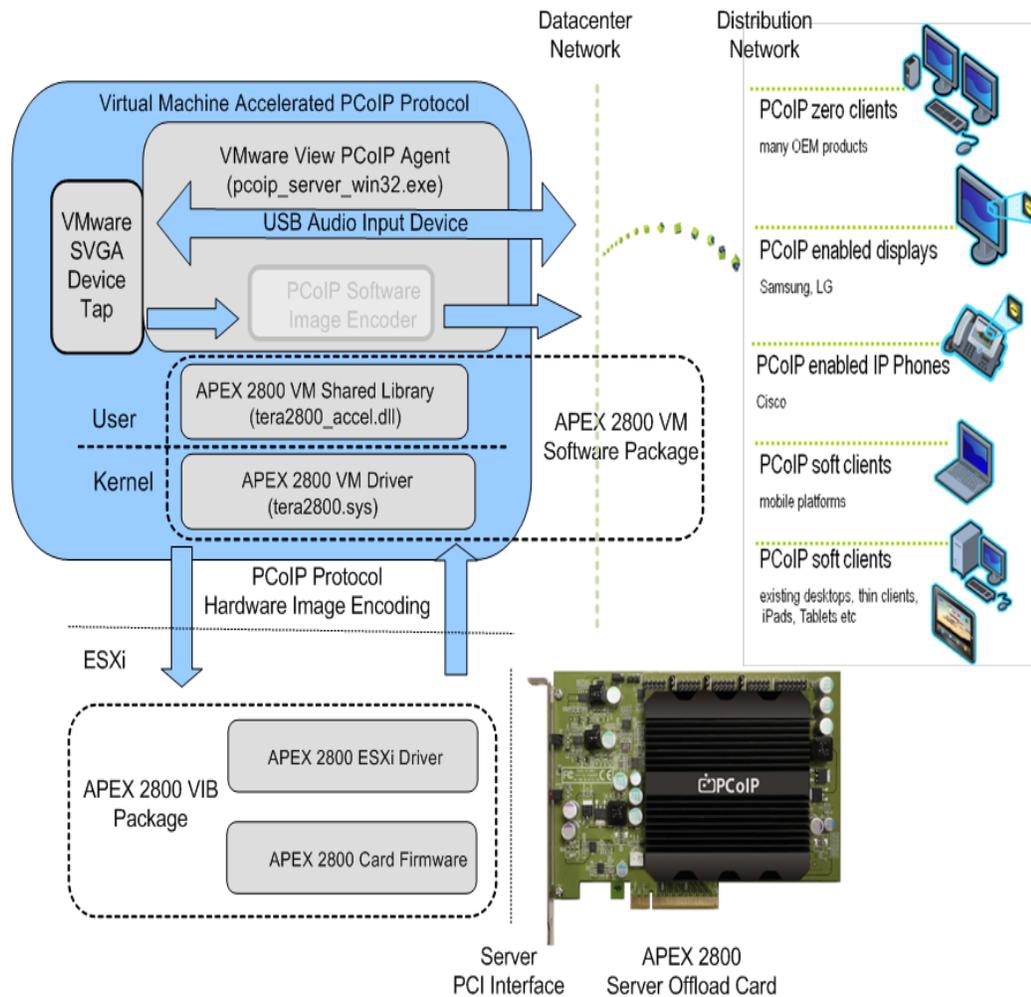
- Dynamically adapts PCoIP encoding based on available network resources uniquely for each PCoIP session
- Display offload imaging capabilities:
  - Per-pixel image decomposition
  - Motion compensation
  - Multiple display compression
  - Dynamic quality adjustment
  - VMware View vMotion

The APEX 2800 card works with an ESXi driver as well as the.dll library and driver in each VM to offload the desktop PCoIP protocol encoding. For selected displays, the display pixels are transferred to the APEX 2800 card for encoding. The encoded and compressed pixels are then sent back to the VM for transmission to the client.

### 4.1 Solution Components

There are three major solution components for the APEX 2800 server offload card as shown in the following figure:

- APEX 2800 PCIe card - The Card
- APEX 2800 ESXi drivers
- APEX 2800 VM software package



**Figure 7-1: APEX 2800 Card Solution Software Modules**

#### 4.1.1 APEX 2800 Card Hardware Characteristics

- Teradici TERA2800 PCoIP processor
- PCIe form factor server add-in card
- Plugs into any PCIe x8 slot
- PCIe x8 Gen 1.1
- Full height, half-length card (length 6.6", width 4.376", height 0.57")
- Power consumption less than 20W (15W typical)
- Passive cooling
- RoHS-6 compliant (including EU directive compliant)
- On-board card and device temperature monitoring
- Temperature range
  - Operational: 0°C to 55°C
  - Storage: -20°C to 70°C

- Humidity
  - Relative (non-condensing): 10% to 90%
  - Storage: 5% to 95%
- Hardware certifications include:
  - Product Safety Approvals:
    - UL for US
    - CSA for Canada
    - IEC 60950-1
  - EMI and EMC Approvals:
    - FCC Class A for USA
    - CE for EU
    - VCCI Class A for Japan (also DENAN and MIC)
    - C TICK/CISPR Class A for Australia and New Zealand
    - BSMI/CNS for Taiwan
    - KCC for Korea
    - IC for Canada



Figure 7-2: APEX 2800 Server Offload Card

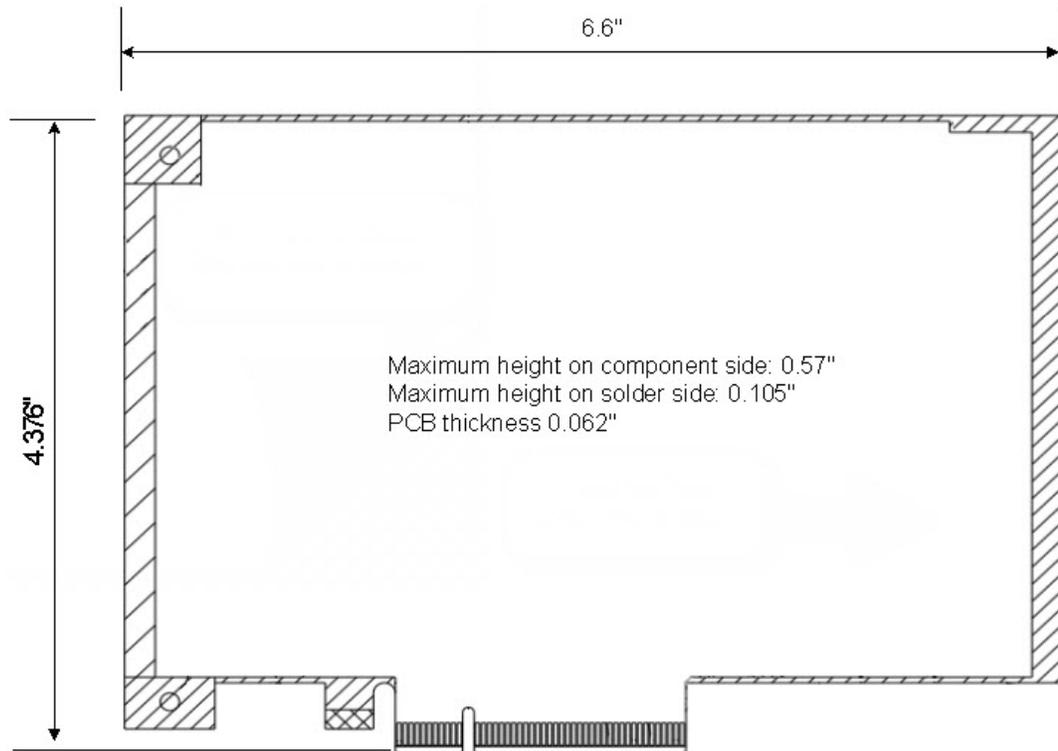


Figure 7-3: APEX 2800 Server Offload Card Mechanical Drawing

## 4.2 APEX 2800 Software Packaging and Delivery

The APEX 2800 ESXi driver is shipped within a single VMware vSphere Installation Bundle (VIB) package. The VIB is available from the Teradici Support website at <http://techsupport.teradici.com/>. System administrators install it using the 'esxupdate' utility.

**Note:** The VIB for ESXi 5.0 package level is VMware Supported.

The APEX 2800 VM Shared Library and APEX 2800 VM driver are packaged and shipped as an .exe file. You must install this file on each VM of the VDI pool that you want to make eligible for offload by the APEX 2800 card.

For installation details, see the TER1109003 *APEX 2800 Card Administrator's Guide*.

### 4.2.1 APEX 2800 Card VIB Package (ESXi Driver)

The APEX 2800 card firmware enables the card features and functionality. The firmware image is packaged within the APEX 2800 ESXi driver, and is conditionally written to the card's flash memory when the server boots up. If the current card firmware matches the firmware packaged with the ESXi driver, then no update takes place and the APEX 2800 card boots as normal. If the firmware is a different version, the most recent firmware is automatically loaded into the card's flash memory. The APEX 2800 card completes the boot process.

**Note:** The firmware version is only checked and updated during the installation of the ESXi driver. To upload a new firmware version you must go through the ESXi driver installation process.

The APEX 2800 card's ESXi driver contains the device drivers for the card functions and PCIe interfaces:

- APEX 2800 card loads, boots, and resets the card. It also communicates with multiple APEX 2800 cards in the server if present.
- Virtual machine management interface to provide individual VM's access to the APEX 2800 card's offload resources (via the APEX 2800 .dll in the VM).
- Display management interface to control access to the APEX 2800 card's display offload resources and to optimize system performance.

#### 4.2.2 APEX 2800 VM Software Package

The APEX 2800 solution also includes software for each virtual desktop VM consisting of a driver and a shared library that provides an API to the standard Teradici PCoIP software encoder engine in the VM to let the APEX 2800 monitor and offload display surfaces.

The APEX 2800 VM driver is a kernel-level device driver, which provides kernel-level memory management capabilities to the APEX 2800 shared library.

Specifically, the driver:

- Lets the PCoIP Agent in the VM request hardware offloading by the APEX 2800 card.
- Sends the display pixels from the display frame buffer to the APEX 2800 card.
- Receives the PCoIP encoded display pixels from the APEX 2800 card.

### 4.3 About the Display Offload Feature

To offload or cancel the offload for a VM, the APEX 2800 card:

1. Dynamically monitors the display imaging activity as rendered by the vCPU or a GPU.
2. Decides to offload a VM based on the following criteria:
  - a. Are offload card resources available? (That is, are there more than 64 displays to be offloaded on the server?)
  - b. What is the VM offload priority setting (highest first)?
  - c. Is the VM bandwidth above the set threshold? (View 5.0 clients only)
4. Makes the offload canceling decision based on the following criteria:
  - a. What is the VM offload priority setting (lowest first)?
  - b. What is the display imaging activity level?
  - c. Is the VM available bandwidth below the set threshold? (View 5.0 clients only)
5. If the offloading criteria are met, transfers display pixels from the VM to the APEX 2800 card.
6. Processes the image:

- Per-pixel display image decomposition
  - Motion compensation
  - PCoIP protocol encoding
7. Transfers PCoIP encoded data from the APEX 2800 card to the VM network interface.
  8. Transfers PCoIP packets to the client for decoding.

### 4.3.1 Display Rendering

The APEX 2800 card is not involved in the VM display rendering. Rendering is handled within the VM by VMware's virtual graphics driver or physical GPU (that is, direct mapped GPU). While the APEX 2800 card is not used in the display rendering process, the use of the APEX 2800 card in a server frees up CPU cycles that the VMware graphics driver can use for additional rendering load (that is, video decode, 3D or Aero Glass support). Also, direct-mapped or virtual graphics cards can be used for hardware accelerated display rendering along with the APEX 2800 card to maximize virtual desktop performance.

### 4.3.2 Display Imaging Activity Monitoring

The APEX 2800 card monitors the display activity for all VM's that have the .exe file installed and that have the **Hardware Acceleration** setting enabled through View Administrator. This lets the APEX 2800 card offload the most active displays and maximize the card benefit. Enabling the hardware acceleration and setting the priority for each VM offload is done on a per-display basis to maximize the card offload benefit. By default, VMware View desktops have PCoIP hardware acceleration enabled as shown in the following figure.

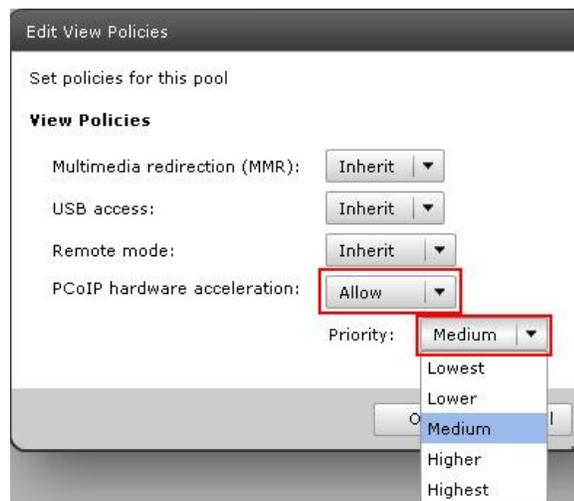


Figure 7-4: VM Hardware Acceleration and Priority Settings

### 4.3.3 Display Offload Decision

Virtual desktop displays are offloaded until the APEX 2800 card resources are fully assigned. The PCoIP software encodes any remaining displays. This provides scalability by allowing more desktop displays in a virtual desktop server than APEX 2800 card resources.

For example, a server may support 96 virtual desktops with dual display (192 displays) but only 128 displays can be offloaded via two APEX 2800 cards in the server.

#### **4.3.4 Display Offload Priority**

The APEX 2800 card supports VM offload priority to allow the assignment of card resources across different user pools or groups (for example, executive versus departmental desktop pools). Priority is handled as a strict priority across five levels. With strict priority, displays are offloaded from the highest priority first then to the lowest priority. All displays for a given VM priority are offloaded prior to servicing displays from the lower priorities. Servicing continues from the highest priority to the lowest priority until the card resources are fully assigned (that is, 64 displays are offloaded). Displays that are not offloaded continue to use PCoIP software encoding to maximize the number of VM's per server.

For VMware View 5.0 clients, the APEX 2800 card also considers the virtual desktop session bandwidth, which must be above a threshold before displays are eligible for offload. VM's that are below a minimum threshold for network bandwidth utilization are not offloaded as their offload has less benefit. Also, having these VM's continue to use the PCoIP software compression ensures that they can take advantage of software optimizations for low bandwidth operation.

#### **4.3.5 Display Offload Canceling**

The APEX 2800 card actively manages the imaging and encoding resources on the card. A change in the demand for the APEX 2800 card resources may require display offload canceling to stay within the available resources. This could involve additional high-priority displays going above the imaging or bandwidth threshold and becoming eligible for offloading. When the demand for imaging resources goes above what the card can support, display offloading is canceled until the demand fits within the available card resources.

Display offload canceling is handled at the lowest priority first for displays that have the lowest imaging activity. This process also follows strict priority, so the lowest priority displays have offloading canceled before canceling the offloading for displays at higher priorities. Displays that have the offloading canceled are switched to use the PCoIP software encoding without impacting the end-user experience.

#### **4.3.6 Display PCoIP Encoding/Compression**

The APEX 2800 card contains two internal PCoIP processing engines that together provide hardware acceleration of PCoIP encoding and image compression for up to 300 Mpps. The APEX 2800 card supports more sophisticated image compression algorithms than can be efficiently supported by software compression. This provides offload scalability benefit by freeing CPU cycles in all cases. It also provides improved image quality and network bandwidth savings in some cases.

Specifically, the APEX 2800 uses the full scope of PCoIP technology to detect, compress, and send changes in pixels to clients in the most effective manner:

- Image decomposition: Analyzes the pixels for a display to determine that optimum compression algorithm to use for each pixel.
- Progressive image refinement: Provides an efficient use of bandwidth while ensuring an exact copy of the source desktop imaging.
- Motion compensation of display regions that are scrolling or "dragged" across the display: Reduces network bandwidth utilization while improving desktop quality under certain workloads.

#### **4.3.7 Per-Pixel Display Image Decomposition**

For optimum image quality and bandwidth utilization, different image types in a PC display require different compression algorithms. Some compression algorithms are best for video and animation, others for natural picture content, and yet others for text and computer generated graphics. The PCoIP protocol uses sophisticated image decomposition algorithms to determine the optimal compression algorithm to use for each region of the display. The APEX 28000 card uses similar image decomposition algorithms to other Teradici hardware host processors.

Image decomposition in the APEX 2800 card is performed at the pixel level allowing a finer granularity and better selection of the imaging encoding algorithms to use on portions of the display screens as compared to software PCoIP encoders used by VMware View when the APEX 2800 card is not present. This is most beneficial when interacting with a zero client because zero clients can perform at-speed mask decoding at the pixel level, whereas block-based decomposition is used on soft clients to reduce the client CPU taxation.

For dynamic bandwidth optimizations, the APEX 2800 card supports the same quality/frame-rate trade-off techniques as the Teradici hardware host processors to ensure the best user experience for the given network bandwidth. These settings change dynamically over time as the network bandwidth varies.

#### **4.3.8 Pixel-based Imaging Motion Compensation**

The APEX 2800 card supports enhanced motion compensation to further reduce network bandwidth use. It detects blocks of the screen that have moved relative to their previous position usually due to scrolling of a page or dragging of an object across the screen. Because this detection is done based on the raw pixels on the display, it is independent of the application or operating system unlike other software-based motion compensation schemes which do not consistently work for all applications. When motion is detected, the moved pixels are simply copied by the client without resending them across the network, reducing the network bandwidth use.

### **4.4 Networking**

The APEX 2800 card does not need any specific network configuration. It works on any network optimized for the PCoIP protocol.

Key network requirements include:

- Full duplex, switched network
- Sufficient average bandwidth and minimum burst headroom for users sharing the network link
- Less than 0.1% packet loss to allow for proper virtual desktop delivery (that is, a real-time application)
- Latency for round-trip communication is 250 ms or less

For more information about PCoIP networking requirements and recommendations, see the following documents on the [Teradici Support Site \(techsupport.teradici.com\)](http://techsupport.teradici.com):

- TER1105004 *Network Design Checklist for Virtual Desktop deployments using PCoIP Protocol*
- TER1007002 *PCoIP Zero Client to VMware View 4 LAN and WAN Network Optimization Guidelines*

## 4.5 USB

The APEX 2800 card does not handle USB. The VMware View Agent handles USB redirection in the virtual machine.

## 4.6 Audio

The APEX 2800 card does not handle audio. It is handled by the VMware View Agent in the virtual machine.

**Note:** You can use the Teradici Virtual Audio driver for analog audio input and output. The audio driver in VMware View 4.6 or later only provides analog output.

## 4.7 VMware vMotion Support

During a vMotion operation, displays currently being processed by the APEX 2800 card seamlessly revert to the PCoIP software encoding. Once the vMotion event completes, the PCoIP agent requests the display processing be offloaded on the destination server. If there is an APEX 2800 card present on the destination server and resources are available, the end user resumes the hardware accelerated PCoIP session. Again, switching from hardware encoding to software encoding and back to hardware is seamless to the end user.

**Table 7-1: vMotion Migration Offloading Support Matrix**

Desktop VM vMotion Migration		Destination Server	
		Server without APEX 2800	Server with APEX 2800
Source Server	Server without APEX 2800	Displays are not offloaded	Displays are eligible for offload depending on loading and VM offload priority
	Server with APEX 2800	Displays are not offloaded	Displays are eligible for offload depending on loading and VM offload priority