A Wanova White Paper

Wanova Mirage™

Cloud Desktop Management

Architecture

By Issy Ben-Shaul

CTO, Wanova

910 Campisi Way, Suite 2C

Campbell, CA 95008

408-201-9420

wanova.com

January, 2012

CONTENTS

INTRODUCTION

Wanova Mirage™ software is a cloud desktop management solution that combines the manageability benefits of centralization with the flexibility and user experience of local execution.

Mirage clones the desktop image in the private data center for enhanced manageability and protection, retains a copy of the image at the end point for local execution, and optimizes the synchronization between the local and central copies, making the solution optimal for environments with distributed and laptop users

This paper outlines the full architecture and typical use cases for Wanova Mirage software.

CHALLENGES WITH DESKTOP VIRTUALIZATION FOR DISTRIBUTED ENTERPRISES

IT administrators regularly try to simplify the complex and expensive task of managing, supporting and ensuring continuity for desktops in the enterprise. This task is even more complex for distributed organizations with many branch offices and an increasingly mobile workforce. Desktop Virtualization promises to simplify the manageability of desktops through centralization and consolidation of desktop-related management tasks. However, there are many limitations when considering this technology for distributed enterprises, especially in environments with a large number of laptop users and knowledge workers. .

There are two main kinds of desktop virtualization. The most common method is server-hosted desktop virtualization or Virtual Desktop Infrastructure (VDI) . By hosting the desktops in the data center on virtual machines and allowing users to connect to their desktop from anywhere using a remote desktop protocol, VDI promises centralized image management and enhanced data protection. VDI is effective for stateless task workers that are always tethered to the LAN. However, VDI doesn't work for every user. Examples include: laptop users who must be able to work offline and from multiple networks (including high latency ones), distributed users who connect to the data center over WAN who cannot tolerate severe performance degradation created by network latency or low bandwidth, and power users who need to run demanding multimedia applications.

An organization must also consider infrastructure costs. VDI requires significant computing and storage capacity to accommodate the execution of user's desktops in the data center, while the end point resources are often underutilized.

Laptops, especially those used by knowledge workers, offer unique management challenges: They are often used offline or outside the corporate network, making it harder for IT to ensure compliance with software updates or resolves issues when they arise. Further, laptops are much more prone to loss or damage, and require special disaster recovery methods as they are often remote from the data center. For many users, their laptop is their lifeline: worker productivity depends on the ability to quickly recover a laptop back to normal operation.

Client-hosted desktop virtualization tries to address the needs of mobile/distributed users by executing the desktop on a virtual machine (VM) at the end point, in some cases, combining this with central management of the desktop images. This lets users work offline with predictable user experience independent of bandwidth and latency constraints. While a hypervisor-based solution may make sense for a "Bring Your Own Computer" model, it can create performance, security and support problems. If the host OS is not operating properly, then the VM will not operate either, and if the host OS is managed, then IT ends up managing and licensing two images, requiring more – not less – work. Recent client-hosted solutions offer a Type-1 hypervisor, eliminating the need for a fully-provisioned host OS. However, Type-1 hypervisors introduce hardware compatibility issues, which are more severe with laptops due to the myriad of peripherals and device drivers. Finally, regardless of the type of hypervisor being used, client-hosted solutions lack the centralization benefits of VDI: the desktop instances are not centralized, and therefore they are less controlled, managed, and protected.

CLOUD DESKTOP MANAGEMENT

COMBINING THE CENTRALIZATION AND MANAGEMENT BENEFITS OF VDI WITH THE NATIVE PERFORMANCE OF A LAPTOP.

Wanova introduces Cloud Desktop Management, an innovative architecture that combines the benefits of centralization with the flexibility to use a fully-personalized laptop and work offline.

The general concept behind the architecture is as follows:

Centralized Desktops – When the Mirage Client is installed on an end point, the full file system of each desktop – including the operating system, corporate applications, user-installed applications, user data and settings – is cloned to the data center[1] into an object called a Centralized Virtual Desktop, or CVD. The CVD object represents a logical desktop and can be stored and managed centrally throughout its lifetime. Furthermore, it can be restored to or re-assigned to a different physical or virtual device as the execution container of the CVD.

Local Execution on the Desktop – While the primary desktop image resides in the data center in the form of a CVD, end users execute a copy of their CVD on their PC, giving them native PC performance, the flexibility to install applications and work offline. The Mirage client is embedded in the operating system and does not require a hypervisor, simplifying deployment and ongoing management.

---

[1] Data center and cloud are used interchangeably throughout the rest of this article

 |

Bi-directional Synchronization – The CVD acts as a two-way cache: when a user connects to the network, changes made by IT to the centrally-managed part of the CVD propagate to the end point, and changes made by end users to the personal part of the CVD propagate to the data center. In order to accommodate the transfer of large amounts of data over WAN links, Wanova employs advanced optimization technologies including desktop streaming, network and storage deduplication and compression.

Image Management and Layering – Wanova provides a toolset for managing CVDs in the data center. At a high level, IT managers define a Core Image, which includes all elements that are controlled/managed by IT, and a set of operations for maintaining the Core Image, including provisioning, updates, and enforcement. The Core Image is decoupled from the personalized and "non-managed" layer, which is nevertheless backed up and part of the CVD, hence it can be restored or re-assigned to a different device.

The Mirage administrator can apply a Core Image to a collection of CVDs, based on various criteria including Active Directory attributes. From then on, updates are made to the Core Image and then applied automatically to the relevant collection of CVDs.

CVD Layering – Wanova defines six layers, 3 in the core image, and 3 in the "personalized" section:

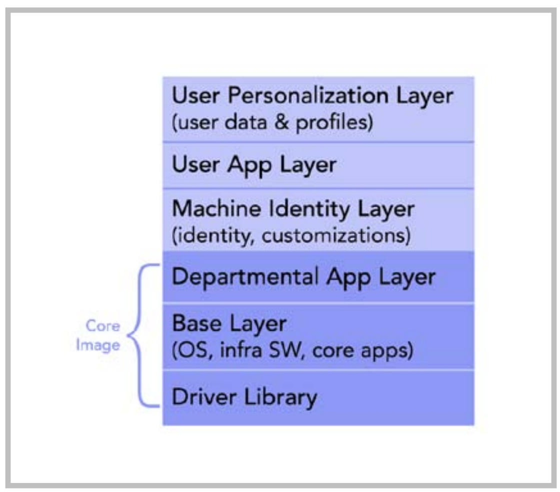

Base Layer: Contains the base Operating System, core infrastructure software (e.g., anti-virus and VPN) that is shared by all users of that Base Layer, and core applications (e.g., Office). Typically, the base layer is built and maintained on a virtual-machine and is HW agnostic.

Driver Library: New in 3.0, this layer extracts drivers from the base layer and thus allows IT to maintain a hardware-agnostic base-layer without inflating it with a superset of all drivers that need to be managed in all base-layers. Wanova offers tools to manage, organize and deliver the proper drivers to end points based on hardware matching rules.

Department Applications Layer (new and experimental in 3.0): This layer allows IT to define managed apps per department, which further reduces the number of base layers that need to be maintained, and reduces the places where apps need to be managed.

Machine Identity Layer: Includes the machine ID, domain name, OS license and other machine-related settings. In 3.0, when a CVD is migrated to a new device, it can be configured to either retain the original machine ID, or overwrite it with the ID of the previous machine.

User Applications Layer: Applications that are not part of the Core Image, including user-installed applications and applications that are delivered by external tools to the end points (e.g., application catalog) – defined automatically by virtue of not being part of the Core Image.

User Personalization Layer: User files, data, settings and other personalization – defined by policy.

The advantage of layering is that it allows IT to easily update and maintain the integrity of the Core Image for compliance and PC stability, while still allowing users to retain their personalization. For example, an Admin can issue an “Enforce Core Image” operation that validates the end point Core Image and repairs it as needed, while keeping all personalized apps and machine state intact. This is in contrast to other desktop virtualization methods, where an image update wipes out all user personalization on the “system drive” and the user has to re-install his applications and rebuild his environment, a tedious task for both IT and the end user.

WANOVA MIRAGE™ ARCHITECTURE

The Wanova Mirage™ architecture is shown in Figure 2, and includes the following components:

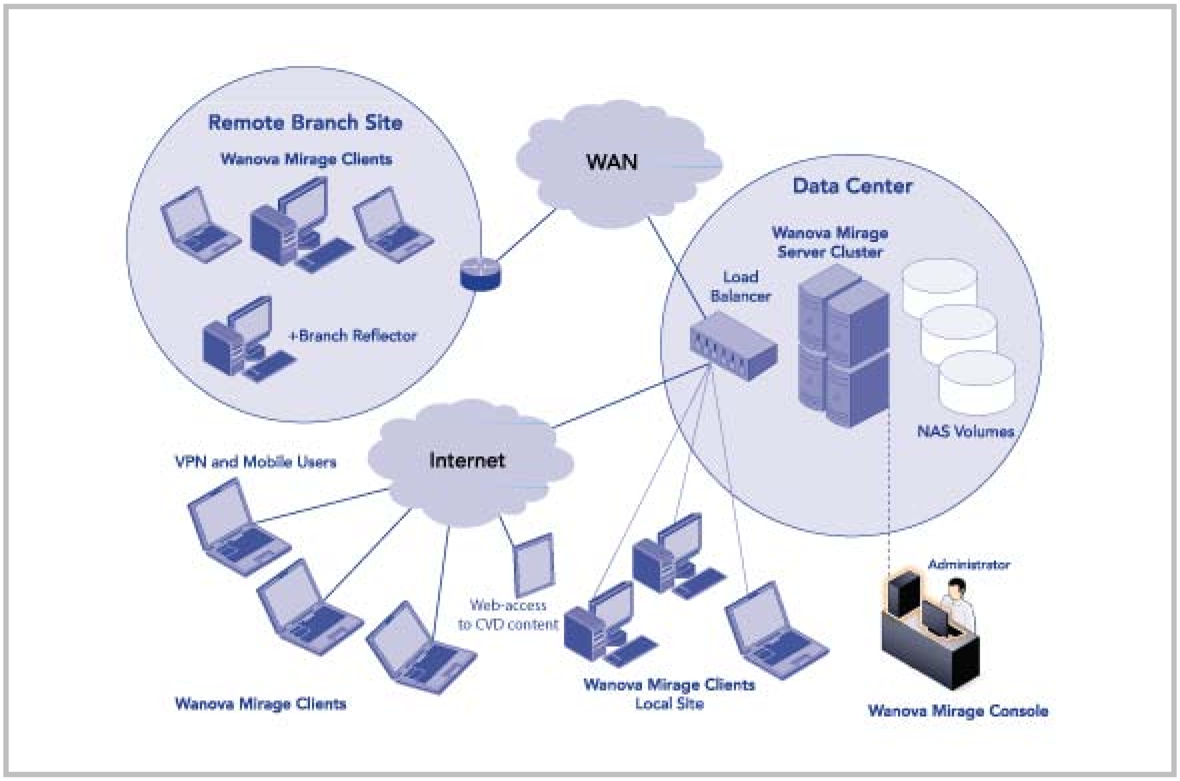

Figure 2: Wanova Mirage™ Architecture

WANOVA MIRAGE™ SERVER

The primary function of the Mirage™ Server is to manage the repository of CVDs and Core Images, and coordinate interaction with Mirage clients.

Mirage Server software can be installed on one or more (physical or virtual) servers running Windows 2008 R2 server OS. Because Mirage Server does not execute actual desktop workloads, each server can handle far more clients than a typical VDI host. Specifically, a standard server with 2 Quad core CPUs and 16GB of RAM can typically host 20-40 virtual desktops in a standard VDI environment. The same type of server would support 1500 – 2000 Wanova Mirage CVDs.

Clustering and Failover:
The Mirage server is stateless, meaning that all system state information is stored in a shared external data repository. This design enables both scalability and high availability, since a client can connect to any available server. Mirage supports clustering either using Microsoft's Network Load Balancing (NLB) service (part of Windows 2008 R2), or it can leverage any existing investment in third party load-balancers. This clustered architecture, combined with the fact that clients continue to operate even if the server is not available, makes it possible for administrators to easily add or remove server nodes on the fly. Finally, Wanova provides a unified management console for the entire deployment, where administrators can address cluster-level configuration and reporting, as well as drill-down to individual servers as needed.

Storage
Mirage uses two storage sub-systems: an SQL database (MS-SQL) to store meta-data for management purposes, and a single-instance-store (SIS) to store CVD data. The SIS can leverage either DAS or SAN (for single-server deployments) or NAS for multi-node, clustered deployments.

Mirage supports dynamic management of storage volumes, making it possible for an administrator to move CVDs between volumes (e.g., for storage balancing or as part of tiered-storage architecture with varying SLAs).

Mirage Storage Optimizations
Wanova Mirage implements a Single Instance Store (SIS), meaning that only a single copy of each file is stored in the data center regardless of how many instances of this file exist within user desktops and across multiple snapshots. Additionally, Mirage performs sub-file deduplication for large and frequently changing files (such as Outlook PST files). Finally, the storage deduplication system is uniquely integrated with the network deduplication system via a shared index database for a highly optimized solution.

Snapshots
Mirage maintains snapshots for each CVD at a frequency determined by the Mirage administrator. The snapshots can be used for two purposes:

1. **Self-service file restore** – allows users to restore deleted, corrupted, or overwritten files from any snapshot using a simple application invoked from the Mirage icon in the system tray.
2. **System-restore** – Administrators can restore the system elements of a CVD (or optionally include user-data) to return a desktop quickly into a last known

configuration state, providing a network-based restore point to recover from any corruption or malfunction of the end point.

Because Mirage transfers and stores CVD snapshots optimally on the network and storage, respectively, it is common for IT to set the snapshot upload interval to once an hour, ensuring near continuous data protection. Snapshots are aggregated in the data center, and once a day they get “collapsed” into a daily snapshot. Similarly, daily snapshots are collapsed to a weekly snapshot and weeklies into monthly, providing long term data and system recovery.

Security

Wanova Mirage employs several security measures to ensure the safety and privacy of the data at rest and in transit. At rest, Wanova supports seamlessly most third-party disk encryption solutions, and encourages their use. The Mirage client-server protocol runs over SSL using standard TCP/IP. Finally, all user access to the central CVD is protected by AD authentication,

Portal for Universal Web-Access

Wanova 3.0 offers end users access to the centralized files of their personal CVD through a standard Web-based portal provided by the Wanova Server. When a user registers a CVD, he or she enters credentials, which are used to bind the CVD with a specific authenticated user. From then on, the user can access her personal CVD from any device, including tablets and smart phones, via a Web-based interface. The Wanova server authenticates the user (via AD) and then connects the user to her CVDs for direct access to her files.

Mirage™ Client

The Mirage Client’s primary function is to allow end points to run a copy of their centrally-stored desktop image. The Mirage Client is not a hypervisor – nor does it require one. The client installs into the base OS and turns it into “self-managed”. The Mirage Client has three main roles. First, it serves as the end point of the transfer and optimization framework of Mirage. This includes the ability to stream down images and act as a cache (fetch missing content on demand), as well as scan and upload snapshots to the server. Second, it performs the “image swapping” (pivot) functionality that enables IT the “zero-touch” capability to swap images and restore old images “in place” onto the end points, not requiring external boot media or PXE server. Third, it act as a “merge calculator” for merging the core-image and personal layers.

MIRAGE™ MANAGEMENT SERVER

The Mirage Management Server is a service that is responsible for scheduling and carrying out management tasks, as well as monitoring and configuring system activities. The Management Server is controlled by the Mirage Management Console, which is the main interface for the administrator, and provides a complete set of tools for image and desktop management, reporting, and system status. The system offers a consolidated dashboard to provide at-a-glance monitoring of the entire deployment, with the ability to drill down into clusters or individual devices as needed. The dashboard monitors:

- Overall system status and health
- Capacity and available licenses
- Progress of image updates and deployments
- State of data protection to identify end point-side changes that have yet to be synchronized
- Compliance of end points with Core Image and the progress of invoked operations

Figure 3 depicts the systems dashboard and control pane.

 |

Figure 3: Mirage Management Console Dashboard

MIRAGE BRANCH REFLECTOR

For enterprises with many branch offices, distributing significant software updates or large images (e.g., in the case of a service pack upgrade, or a full OS upgrade such as XP to Windows 7 migration) to each of the clients in the branch can be prohibitive – both in terms of bandwidth requirements and completion times. Wanova Mirage offers a simple, yet powerful capability that significantly simplifies software distribution, without requiring dedicated proxy hardware at each branch. With a mouse click, IT can designate any desktop as a Branch Reflector. In case of a Core Image update, an image is first downloaded to the Branch Reflector, and then all other end points whose CVD is linked to that specific Core Image can leverage the Branch Reflector to fetch the updated image locally. Hence, the image is downloaded only once over the WAN, and then peer PCs can pull the elements they need to recompose the image directly from the Branch Reflector using the LAN.

MIRAGE NETWORK OPTIMIZATIONS

Wanova Mirage was designed for users who connect over networks with low-bandwidth and high latency. As such, it employs a number of technologies to optimize the transfer of large amounts of data over the network.

GLOBAL DATA REDUCTION

Mirage recognizes that many of the user and system files are identical or similar across users. This inherently large redundancy exists across different generations of the same desktop, as well as across multiple desktops, at both file level, as well as sub-file variable-sized "chunk" level. To optimize data uploads from the clients to the server, Wanova employs an advanced network deduplication technology that eliminates this redundancy across all users of the Mirage deployment, dramatically reducing the amount of transferred data.

This deduplication functions across the entire desktop, including OS, applications and data. For example, assume three thousand end points have the same copy of a Powerpoint file. Once the first copy has been transferred, Mirage will recognize that the file already exists in the Server index database, and will not transfer the file again from any of the other end points. Furthermore, when users make changes to this file and create their own local versions, only these minimal changes are uploaded, leveraging a global "chunk" database in the server. Finally, the same kind of technology is applied when downloading images from the data center to end points, using an index generated by the Mirage Client based on the contents of the local disk at the end point.

Desktop Streaming over the WAN

In addition to network and storage efficiencies, Wanova provides a unique desktop streaming technology that enables unprecedented speed when delivering a whole operational desktop environment to remote end points. There are several instances where this capability can be used. For example, if a laptop is lost or broken and the end user acquires a new PC, IT can quickly move their CVD to the new device. Another use case is if IT wants to restore a PC to a previous state via a snapshot.

With desktop streaming, Mirage automatically identifies the minimum number of files that are required to get a PC to boot, which is much smaller than the entire CVD. The minimal set is optimally transferred to the end point and is then used to boot the system. Once the system is up and running, the remainder of user data and files are streamed in the background, ordered by their usage frequency. That is, files and applications that are used more often by the user will be streamed before those that are rarely used. Finally, if a user or application requests a file before it has been transferred, Mirage fetches the file from the server on-demand with high priority. The bottom line is that desktop enables administrators to get a user back to productivity with their full, personalized desktop in minutes instead of hours or days – even over a WAN.

Deploying and Using Wanova Mirage™ Software

System Deployment and End point Centralization

Wanova Mirage is typically deployed with the following simple sequence of steps.

1. Install the Mirage Management Server, Management Console and Mirage Server software (a process that typically takes between 5 to 20 minutes, provided that the server and storage sub-systems are configured and available.) The server can be hosted on either an on-premises data center or on a public cloud.

2. Install the Client on the desired end points. The Client is a small, 2MB standard MSI executable that can be installed on any Windows OS in about 2 minutes. Once the Client is installed, it signals to the administrator it is ready to be cloned. The IT administrator then verifies whether the user is entitled to be managed by Wanova and if so, initiates the "CVD centralization" phase.

3. The desktop file system contents are uploaded optimally to the data center and stored in the CVD repository. When finished, the end point is considered "protected".

4. The next phase is for the IT administrator to construct a Core Image and assign it to a collection of CVDs that match the criteria for that Base Image. The centrally-stored CVD is updated first, regardless of whether the end point is connected to the network. When that user connects their PC to the network, the Core Image propagates down to the copy on that end point.

5. The Core Image is staged in the end point, and then the user receives a prompt to reboot into the managed image. Post reboot, the old image is replaced with the managed new image, and is considered "managed". Any updates made by IT to the Core Image will follow the same process, synchronizing with each end point. Similarly, any changes made by the user will propagate back to the central CVD.

Core Image Management and Updates

The basic flow of Core Image updates is depicted in Figure 4.

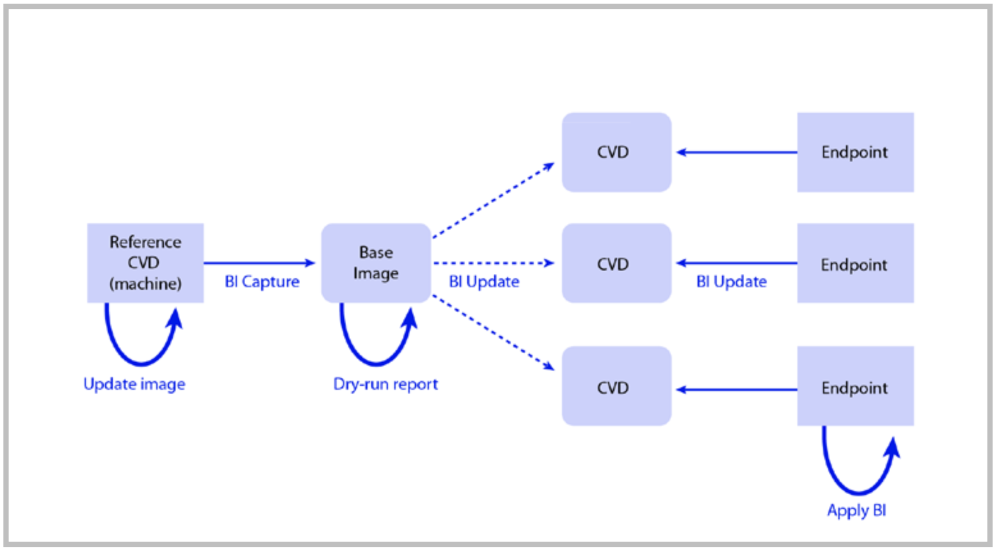

Figure 4: Core Image Update Flow

- The Core Image is initially created following simple guidelines. When IT wants to update the Core Image, the admin uses a special type of CVD called a

Reference CVD, and assigns it to a machine (typically a VM) to perform the necessary changes.

- When the Core Image is ready for deployment, the admin issues the "Core Image capture" operation (a one click operation), which extracts the Core Image from the reference CVD and generates a new version.
- The administrator can optionally run a "dry-run" report on the new Core Image, which conducts a "what-if" analysis of the impact of the Core Image update on the CVDs.
- When IT is ready to deploy the Core Image, the admin invokes a Core Image update operation. This operation transfers the Core Image to the specified end points (possibly using a Branch Reflector to eliminate extraneous transfers over WAN).
- Once complete, the Mirage Client notifies the user about the availability of a new Core Image, and users, at their discretion, apply the update locally, merging the Core Image updates with the personal layers. In cases where an administrator does not want to retain the personalized applications, the 'Update' operation can be applied with a "cleanup" option, which removes any user-installed applications from the CVD and reverts it into a "pristine" state.

Core Image Compliance

Mirage allows IT administrators to enforce a Core Image, which brings an end point back into compliance with the approved Core Image. The primary purpose of this operation is to repair a PC when IT believes a corruption or problem has occurred in a Core Image element. Using the Enforce Core Image command, a malfunctioning PC can be repaired quickly, without requiring a tier-2 helpdesk resource, or a painful end point re-image. Another common use of this command is to periodically ensure that the end points comply with the contents of the corporate Core Image, improving stability for each end point. The flow is quite simple – the administrator selects the target CVDs and invokes the Enforce Core Image operation on them.

Desktop Recovery and Continuity

Restore a CVD from a Snapshot

In addition to enforcing or updating the Core Image component of the CVD, Mirage allows administrators to restore the entire CVD including user data, personalized applications and settings from a stored snapshot. For example, if the CVD was corrupted in the user-data layer or user-installed applications, or when an administrator wants to revert a CVD into its last known working configuration. The flow is simple – the desired snapshot is selected, and a restore operation is invoked. In cases where the end point cannot even boot prior to the restore, Mirage offers a "recovery" boot media with an embedded Mirage client that can be used to boot the end point and restore the CVD from the Mirage server.

Another variant of restore allows users to restore only the system portion of the CVD, namely the Core Image and "machine" layer, but retain the user-data layer as it exists on the end point. This is useful when IT wants to repair a PC to a known good state, but the user has files or data that were created since that point in time.

Finally, Wanova also offers a self-service file-level restore service, which allows the user to select individual files that were deleted or mistakenly modified, and restore them from any of the snapshots.

Universal Access to CVD files via the Portal

Recall the use case where a user loses or damages his laptop, or is simply without access to it, e.g. left at home and working with a tablet. As a first measure, IT can provide immediate access to the CVD files using the Web-based portal. From any Web-browser the user connects to the portal, authenticates with AD credentials, and is presented with CVDs that are registered with that user. From then on, the user can browse his files and download them for access.

Re-assign CVD to a Central VM and to a new End point

The Web access method is useful when only files are required. However, when the real machine is needed, e.g., to run applications and use the machine as a gateway to the enterprise apps etc., a different access method is needed. Wanova Mirage offers universal access to the CVD by using the unique re-assignment operation, applied from physical onto virtual machine in the data center. Once re-assigned to the VM, the user has full access to the full "desktop" on a VM via any standard remote desktop protocol, similar to a remote access to VDI.

It is important to denote that when the end user regains access to his old or new laptop , the CVD can be re-assigned back to the physical device and the end user continue to work, with all the work that was done on the VM being kept and in-synch.

HARDWARE REFRESH

Migrating users to a new PC, whether it is part of a scheduled refresh cycle, or because of loss or breakage of a laptop, is time consuming, expensive, and usually results in lost productivity while the end user 're-personalizes' their desktop. These tasks can include re-installing applications that weren't centrally provisioned, recovering settings and preferences, and restoring user data.

With Mirage, the CVD already resides in the data center. The administrator makes sure the end user has connected to the network and synchronized any changes they may have made with their centrally-stored CVD. The administrator then assigns a Core Image to the new PC – this new hardware can be from the same hardware family of the old desktop (e.g., Dell, HP, Lenovo, etc.), in which case the same Core Image is selected, or it can be a completely different hardware family, in which case a new Core Image is selected.

Using the Management Console, an administrator assigns this Core Image to a CVD (or a collection of CVDs). If the new Core Image is for a different hardware family, Mirage dynamically swaps the drivers and hardware-specific services required for the new hardware. The CVD can then be assigned and then pipelined to the new end point. When the new end point reboots, it will instantiate the whole CVD, so the user's desktop will look and feel exactly as it did before, complete with their files, preferences and both IT-provisioned and user-installed applications, all in their original place.

XP TO WINDOWS 7 MIGRATION

Using a process that is similar to 'Re-Base', Wanova Mirage dramatically simplifies the process of migrating PCs from Windows XP to Windows 7, even if the end user is remote or mobile. The administrator can initiate the migration for a large number of users via a simple Wizard in the Mirage Management Console.

The process works as follows: The administrator builds a Windows 7 image, including core applications that need to be delivered. Once ready, she starts the Wizard, which binds the image with a collection of users/end points, performs validations, and gets ready for deployment. Once ready, the image is deployed efficiently to end points (leveraging Branch Reflectors wherever possible). Each end point then waits for the image to be staged, and then performs an image replacement (which can be manually controlled) that reboots the system into Windows 7. Post migration, the system restores user files that were under XP into their proper location in the Windows 7 folder structure, and performs few other tasks such as re-joining the machine to the domain etc.

Note: While Mirage typically preserves user installed applications when enforcing a Core Image, the Mirage Windows 7 migration utility will only preserve user data and settings. IT can still deploy Windows 7 applications that they want to manage centrally as part of the Core Image, but user-specific applications will need to be installed via other means.

To summarize, the following elements differentiate Wanova for Windows 7 migration:

- Zero-touch – no need to send an IT person to remote sites, migration is fully automated even in locations without PXE server or "distribution points"
- WAN optimized, without dedicated infrastructure. Including de-duplication and branch-reflector technology.
- Minimal downtime to the user, typically 30 min only.
- A full copy of the XP image is centralized, allowing the user to revert back to XP or restore user-data, or even re-assign the XP image onto a VM to access XP apps.
- Centrally monitored and configured.

SUMMARY

For enterprises that support distributed mobile laptops and desktops, Wanova introduces a new architecture –Cloud Desktop Management – which brings the manageability and centralization benefits of desktop virtualization to mobile and distributed users, without degrading the end user experience. Wanova's many advantages stem from a new architecture that integrates innovations across the client, server, network and image management. Administrators get powerful centralized management and control, efficient network utilization, and the performance that knowledge workers expect.

