

VMWARE HORIZON 7 INSTANT-CLONE DESKTOPS

VMware Horizon 7

Table of Contents

Introduction	3
JMP – Next-Generation Desktop and Application Delivery Platform	3
Instant-Clone Use Cases	3
Cloning Technology	4
Full Clones	4
Linked Clones	4
Instant Clones	4
How Instant-Clone Desktop Pools Are Created	6
Publishing the Master Image	6
Provisioning Instant-Clone Desktops	10
What Happened to These Linked-Clone Settings?	12
Benefits of Using Instant Clones	13
Simple Setup	13
Easy Management of OS and Software Updates	13
Smaller Load on vCenter	14
Simplified Desktop Pool Management	14
Reduced Storage Costs	15
Instant-Clone Restrictions	16
About the Authors and Contributor	17
Additional Resources	17

Introduction

VMware Horizon® 7 remotely delivers a virtual Windows desktop experience to a user's endpoint device from centralized VMware vSphere® enterprise servers. In comparison to physical desktops, virtual desktops improve IT responsiveness, and the centralized delivery enhances application and data security.

The new VMware [Instant Clone Technology](#) included in the View component of the Horizon 7 Enterprise Edition improves and accelerates the process of creating cloned virtual desktops over the previous View Composer [linked-clone technology](#). In addition, instant clones require less storage and less expense to manage and update because the desktop is deleted when the user logs out, and a new desktop is created using the latest master image.

For Just-in-Time Desktops, you can combine instant clones with [VMware App Volumes™](#) and [VMware User Environment Manager™](#) to create disposable desktops that retain user customizations, personas, and user-installed apps from session to session, even though the cloned desktop is destroyed when the user logs out. Users experience a stateful desktop, while the enterprise realizes the economy of stateless desktops. Just-in-Time Desktops are part of the JMP desktop and application delivery platform, a feature of the Horizon 7 Enterprise Edition.

This white paper is for View administrators who are considering using instant clones. It provides a technical explanation of Instant Clone Technology and its benefits and limitations. This paper is intended both for readers who are familiar with the View Composer linked-clone technology and for readers who are new to the topic.

JMP – Next-Generation Desktop and Application Delivery Platform

JMP (pronounced *jump*) represents capabilities in VMware Horizon 7 Enterprise Edition that deliver Just-in-Time Desktops and Apps in a flexible, fast, and personalized manner. JMP is composed of the following VMware technologies:

- VMware Instant Clone Technology for fast desktop and RDSH provisioning
- VMware App Volumes for real-time application delivery
- VMware User Environment Manager for contextual policy management

JMP allows components of a desktop or RDSH server to be decoupled and managed independently in a centralized manner, yet reconstituted on demand to deliver a personalized user workspace when needed. JMP is supported with both on-premises and cloud-based Horizon 7 deployments, providing a unified and consistent management platform regardless of your deployment topology. The JMP approach provides several key benefits, including simplified desktop and RDSH image management, faster delivery and maintenance of applications, and elimination of the need to manage “full persistent” desktops.

Instant-Clone Use Cases

If you want to use instant clones, you create one or more pools of instant clones. Like View Composer linked-clone pools, instant-clone pools have been tested to support up to 2,000 desktops in a pool.

Instant clones fit several use cases in enterprise environments:

- For task workers, where everyone in the pool requires the same applications, and a specific desktop need not be dedicated to a specific user
- Where you want to take advantage of the benefits of cloned desktops, and provide personalization with other VMware products such as App Volumes and User Environment Manager
- Wherever you previously used View Composer floating desktop pools that used a naming pattern and where 3D rendering was not required

Cloning Technology

A clone is a copy of a master virtual machine with a unique identity of its own, including a MAC address, UUID, and other system information. Instant clones represent the newest generation of cloning technology, after full clones and View Composer linked clones.

Full Clones

A full clone is an independent copy of a virtual machine (VM). It shares nothing with its master VM, and it operates entirely separately from the master VM used to create it.

Linked Clones

A linked clone uses significantly less storage space than a full clone because it accesses software on shared virtual disks. Because of this sharing mechanism, a linked clone must always have access to the disk used for cloning.

To make a linked clone, you take a snapshot of the master VM and then the cloning process creates a replica VM to use for cloning. The linked clone shares virtual disks with the replica VM. The differential—the bits of software that are unique to the linked clone—is stored in a *diff disk* or *redo disk*. This arrangement allows the linked clone to occupy a smaller amount of physical disk space than the master VM but still access the software installed on the shared virtual disks. You can create hundreds of linked diff disks from one replica, reducing the total storage space required.

Instant Clones

Like a linked clone, an instant clone shares virtual disks with the replica VM after the linked clone is created. The process of creating instant clones differs from that used for linked clones in the following way: The cloning process creates a running parent VM from the replica VM. At creation time, the instant clone shares the memory of the running parent VM from which it is created.

Instant clones use copy-on-write for memory and disk management. Instant clones are based on a *running parent VM*, derived from a master VM. (For terminology, see Table 1.) When an instant clone is created from a running parent VM, any reads of unchanged information come from the already existing running parent VM. However, any changes made to the instant clone are written to a *delta disk*, not to the running parent VM. This strategy preserves security and isolation between the instant clones by ensuring that

- Each instant clone is immediately accessible.
- Changes do not affect the shared data and memory of the running parent VM on which all other instant clones are based. Sharing the memory of a running parent VM at creation time enables instant clones to be created within a few seconds and instantly powered on. An instant clone requires no boot time when the cloning process is finished.

Because an instant clone can be created so quickly, it does not need to persist after a user logs out. Instead, the instant clone is deleted when the user logs out. Depending on the number of spare VMs configured for the desktop pool, a new instant clone might be created immediately after a used instant clone is deleted. In this manner, users get a newly created desktop whenever they log in. If the master image—the VM snapshot used to create the pool—has been updated since the last login, the user gets the new image.

Note: The instant clone is deleted when the user logs out, not necessarily when the user disconnects. If the user disconnects the session, the virtual desktop remains, unless the administrator has configured the user to be automatically logged out after disconnecting.

The following table summarizes the types of VMs used or created during the instant-cloning process.

TERM	DEFINITION
Master VM	The VM that you create and configure as the model for the virtual desktops you plan to deploy.
Master VM snapshot	A reproduction of a master VM exactly as it was when you took the snapshot. The snapshot includes the state of the data on all VM disks.
Internal template VM	A linked clone of the master VM, based on the master VM snapshot. This VM is used for the AD domain join.
Replica VM	A thin-provisioned full clone of the internal template VM. The View Storage Accelerator feature uses a content-based read cache digest of this VM.
Running parent VM	A linked clone of the replica VM, based on a snapshot the system takes of the replica VM.
Instant clone	A desktop VM created from the memory and disk of the running parent VM. After the instant clone is created, it shares the read disks of the replica VM, exactly like a linked clone.

Table 1: Instant-Clone Terminology

Figure 1 shows the process of creating the various types of VMs that ultimately produce an instant clone.

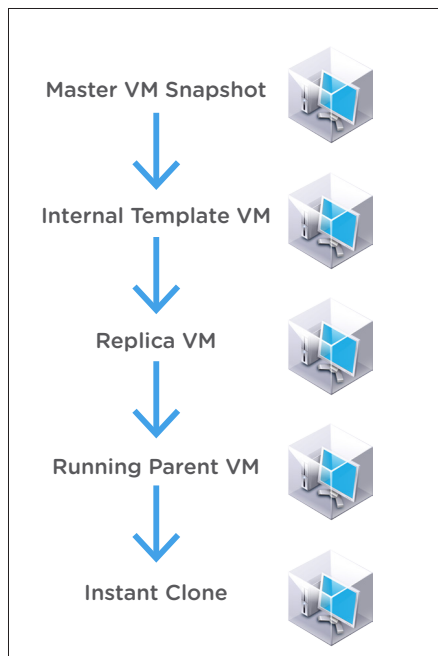


Figure 1: Instant-Clone Sequence

How Instant-Clone Desktop Pools Are Created

Creating an instant-clone desktop pool is a two-part process:

- Publishing—also called priming—the master image
- Provisioning the desktops in the pool

Publishing the master image can take from 7 to 40 minutes, depending on the type of storage you are using. Provisioning the desktops takes only 1 or 2 seconds per desktop. You can perform these tasks at separate times, so that the provisioning process occurs either at a scheduled time or immediately after the publishing process is complete.

Publishing the Master Image

The Add Desktop Pool wizard in Horizon Administrator guides you through the process of publishing the master image—that is, using the specified VM snapshot and creating all the required types of internal VMs, including the running parent VMs. Completing the Add Desktop Pool wizard for instant clones is similar to adding any type of pool in View, except there are fewer settings to configure.

Publishing the master image means completing the process to create running parent VMs so that the system is ready to instantly clone desktops during the second part of the pool-creation process.

Figure 2 depicts the VMs that get created as part of this process and shows an example of where each VM might be stored. In this first example, the VMs are stored on separate datastores, as in a tiered storage model.

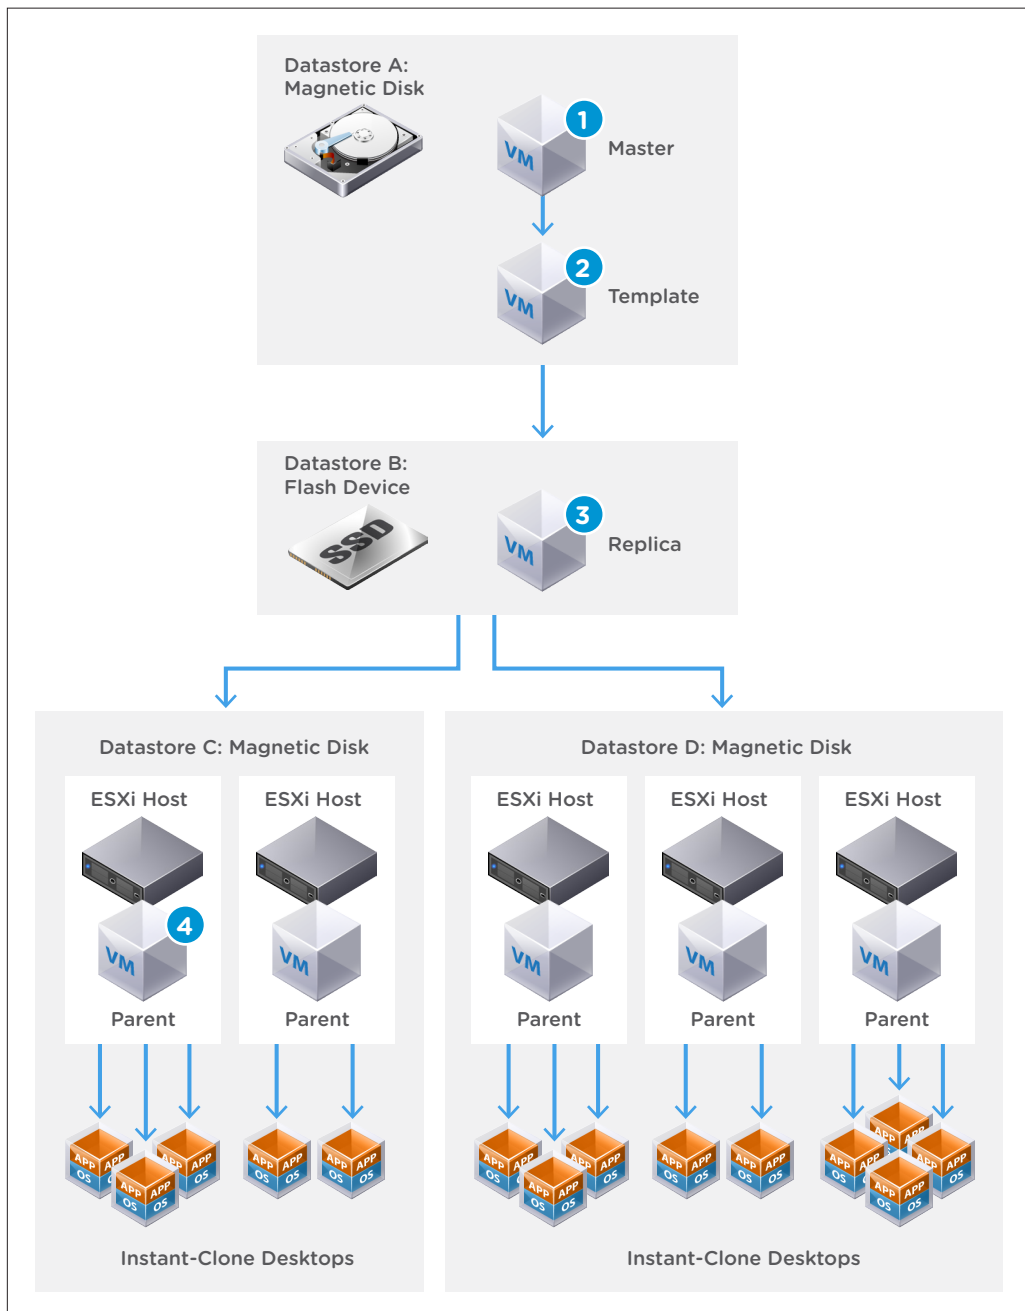


Figure 2: Creating Instant-Clone Desktops and Storing Replicas on Separate Datastores for Tiered Storage

The following steps describe the publishing process shown in Figure 2.

1. In the Add Desktop Pool wizard, you select a snapshot of a master VM to use for the instant-clone pool.

After you complete the wizard, the instant-clone engine performs the rest of the steps.

2. The instant-clone engine uses the master VM snapshot that you specified to create one internal template VM on the same datastore as the master VM. This internal template VM is linked to the master VM. The internal template VM is used to create the replica VMs.

The system performs a domain join on this internal template VM, which ensures that all the proper Windows registry keys and settings are correctly populated. This process involves a reboot.

Using this internal template VM for the domain-join process means that the process is performed only once on one VM. Without the internal template VM, the domain-join process would have to be performed on all the replica VMs that are created during the next step of the process.

3. One or more replica VMs are created from the internal template VM. The replica VM is a thin-provisioned full clone of the internal template VM. The replica VM shares a read disk with the instant-clone VMs after they are created. A content-based read cache (CBRC) digest is created. (For more information about this digest file, see the vSphere [blog about View Storage Accelerator](#).)

While completing the Add Desktop Pool wizard, you can choose to place the replica VM on a different datastore from the instant-clone VMs in the desktop pool, just as you can for View Composer pools. For example, in a tiered storage scenario such as that shown in Figure 2, you can store the replica VM on a solid-state, disk-backed datastore. Solid-state disks have low storage capacity and high-read performance, typically supporting 20,000 IOPS.

If you choose to have replica VMs on the same datastores as your instant clones, one replica VM is created per datastore, as shown in Figure 3.

4. The system takes a snapshot of each replica VM and uses it to create one running parent VM per VMware ESXi™ host per datastore. The running parent VM's memory and disks are used to create the instant clones.

Note: Although the running parent is used to create the instant clone, after the clone is created, the clone is linked to the replica VM and not to the running parent VM. You can delete the running parent VM without affecting the instant clone.

Important: If you use the same master VM snapshot for multiple pools, those pools use the same running parent VMs to create the pools. If additional pools use additional datastores or ESXi hosts, additional running parent VMs are created for the ESXi hosts and datastores.

In Figure 3, replica VMs are placed on the same datastores as the instant clones. One replica VM is created for each datastore.

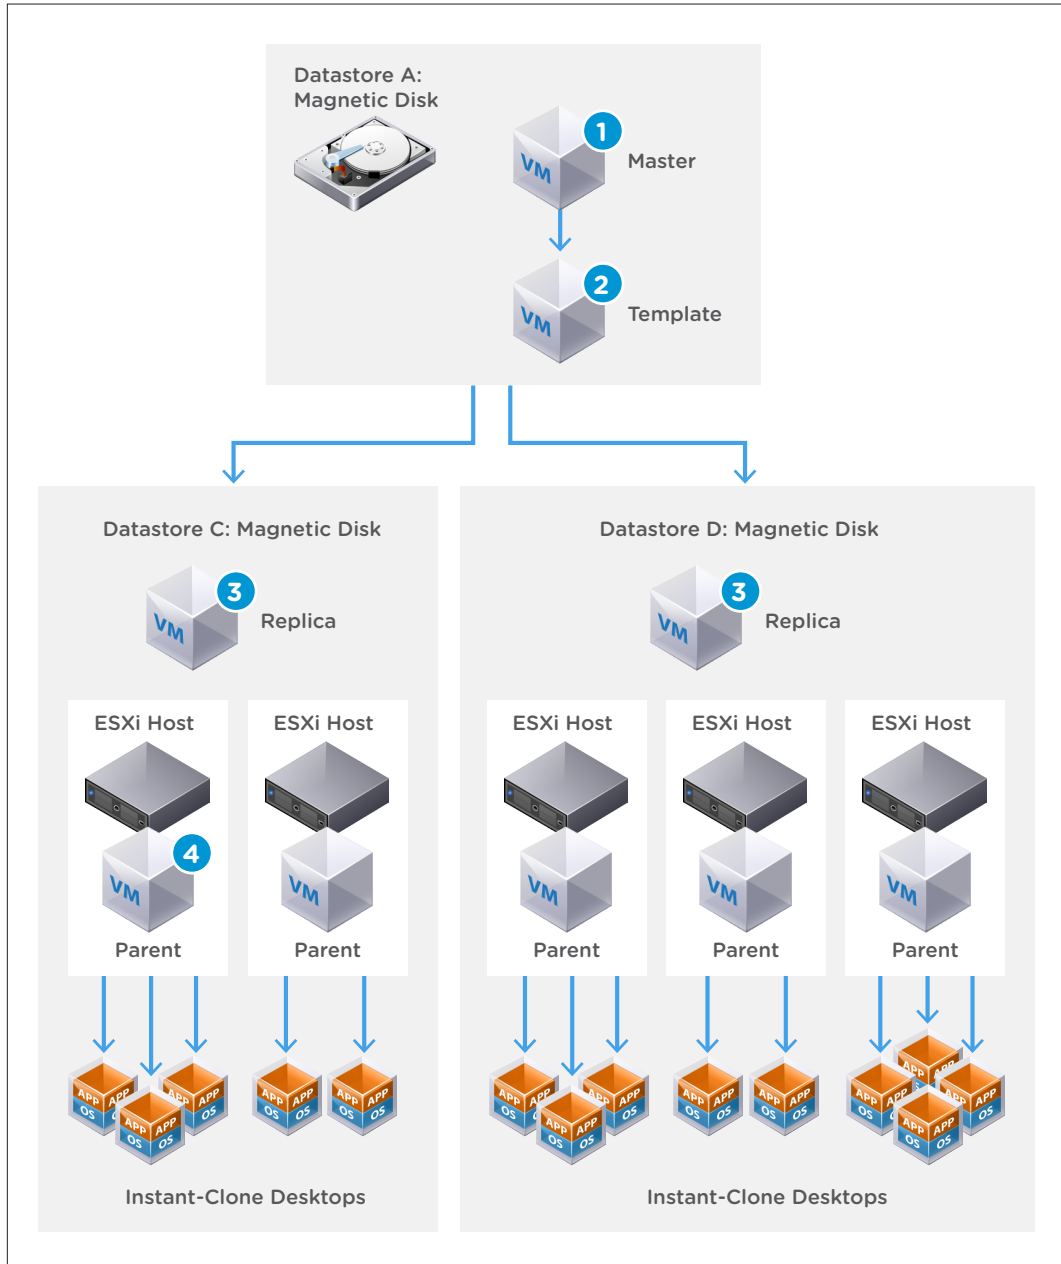


Figure 3: Placing Replica VMs on the Same Datastores as Instant Clones

Provisioning Instant-Clone Desktops

After the master image is published and the running parent VMs are created, the provisioning of instant clones can begin. The instant-clone engine performs the following tasks to create instant clones:

1. The engine brings the running parent VM to a *quiescent*, or quiet, state and then forks it using the vSphere vmFork technology. The forking process is like creating two similar branches of development so that disk and memory can be shared.
2. The engine customizes each forked instant clone. This *ClonePrep* process performs the following customization tasks, all without requiring a reboot:
 - Gives the VM a unique MAC address
 - Changes the Active Directory password
 - Joins the machine to the Active Directory domain
 This domain join does not require a reboot because the associated internal template VM was already joined to the domain and rebooted during the publishing process described earlier.
 - Activates the Microsoft license

You can use scripts for the ClonePrep process so that one script runs immediately after a clone is created and another script can run before the clone is powered off. These scripts can invoke any process that can be created with the Windows CreateProcess API, such as **cmd**, **vbscript**, **exe**, and batch-file processes.

The provisioning process does not require power operations, and the clones are forked from a running parent VM, so the process takes only a couple of seconds. Figure 4 shows the steps in the provisioning process for instant clones as compared with the View Composer provisioning process for linked clones.

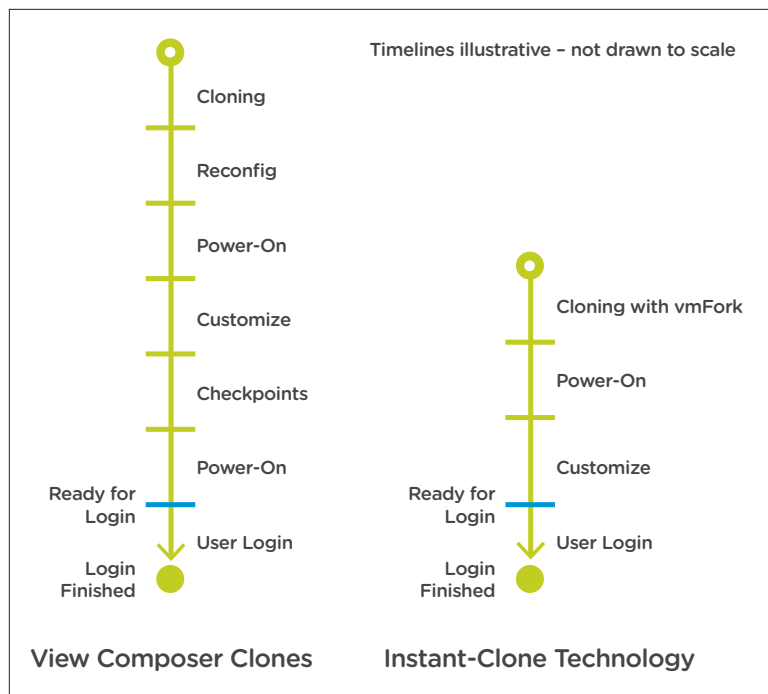


Figure 4: View Composer Process for Creating Linked Clones Compared to the Process for Creating Instant Clones

Figure 5 shows the entire process of publishing the image and provisioning the instant-clone desktops. The publishing process, which comprises the first 10 of the 12 steps in the diagram, can be completed ahead of time so that when it is time to create the pool, desktops appear virtually instantly, just in time for users to log in.

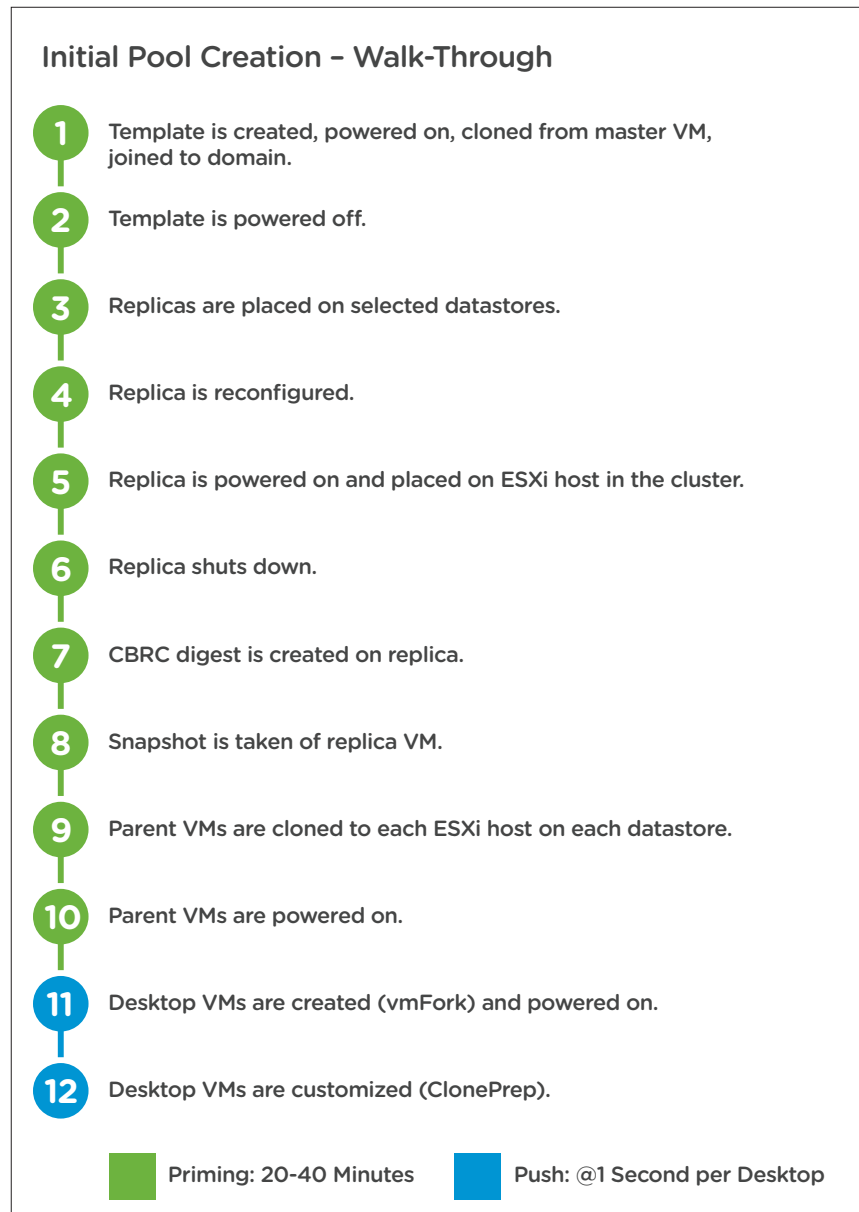


Figure 5: Instant Clone Process from Publishing to Provisioning

What Happened to These Linked-Clone Settings?

If you are already using View Composer linked clones and are planning to replace them with instant-clone pools, you will find that the wizard for creating instant-clone pools requires less information. The following View Composer settings are not needed for instant-clone pools.

- **Delete or refresh machine on logout** – An instant clone is always deleted whenever a user logs out. An instant clone with the newest image is ready when the user logs in again.
- **Remote machine power policy** – Instant clones are always powered on. If an instant clone is powered off, the user is logged out. When a user logs out, the virtual desktop is deleted and recreated.
- **Allow users to reset their machines** – To reset the machine, the user is logged out. When a user logs out, the virtual desktop is deleted. If a virtual desktop freezes, the administrator can use a recover operation, which deletes the current VM and creates a new one.
- **Redirect Windows profile to a persistent disk** – To accomplish the same effect, use App Volumes.
- **Disposable file redirection** – With instant clones, the virtual desktop is deleted when the user logs out.
- **Select separate datastores for persistent and OS disks** – Deciding whether to use separate datastores is not necessary because with instant clones, there are no persistent disks. You can use App Volumes instead of using persistent and OS disks.
- **Minimum number of ready (provisioned) machines during View Composer maintenance operations** – With instant clones, you do not use recompose operations, which typically occur during scheduled maintenance windows. Instant clones are provisioned so rapidly that maintenance operations are not required.
- **Reclaim VM disk space** – Reclaiming disk space is not needed because the virtual desktop is deleted when a user logs out. VM disk space does not have much time to grow. The space-efficient disk format (SE sparse), with its wipe and shrink process, is not needed.

These settings are automatically enabled:

- **Use View Storage Accelerator** – View Storage Accelerator is automatically enabled. Clone-level CBRC is no longer needed, so you do not need to specify the level of CBRC. Master VMs and replicas still use CBRC, and the CBRC digest is calculated automatically.
- **Transparent page sharing scope** – Transparent page sharing is automatically enabled at the VM level.

Benefits of Using Instant Clones

Instant clones are easy to implement and manage and offer significant savings by reducing storage costs and support costs. Instant Clone Technology does not require a database or a separate dedicated server, so overall support costs for the virtual desktop infrastructure (VDI) are reduced, as is the complexity of future infrastructure upgrades.

You can also use VMware vSphere Distributed Resource Scheduler™, VMware vSphere High Availability (HA), and VMware vSphere vMotion® with instant clones.

Simple Setup

Instant Clone Technology is built into the View server and agent components and does not require a separate server or database:

- The instant-clone engine is a Java app that runs on the Tomcat server. One View Connection Server 7.0 (or later) instance in a replicated group runs the engine. If the server fails, another one takes over.
- The instant-clone information is stored in VMware vCenter Server™ (vSphere Version 6 Update 1 or later) rather than in a separate database.
- Horizon Agent 7.0 (or later) with the instant-clone component has code for instant clones and the instant-clone customization engine.

Important: When you install the agent, you can select either the instant-clone component or the View Composer linked-clone component. You cannot install both components in the same VM.

Easy Management of OS and Software Updates

To manage OS patches and software updates with instant clones, you use the push-image operation. You make changes to the master VM and take a snapshot or use a snapshot from a different master VM. You can recreate the pool from any snapshot of any master VM.

You can then either push the new image out immediately or schedule it for a more convenient time. The process for publishing a new image takes from 7 to 40 minutes (depending upon the storage used) and occurs when you complete the Push Image wizard page. The instant clones are then created within seconds at the time you specify. You can force users to log out or wait until the next time they log out to get the new image.

The push-image operation for instant clones achieves the same goal as the recompose operation for View Composer linked clones. However, the recompose operation is a much slower process and requires you to plan for desktop maintenance windows to perform the operation at off-peak hours. Because the provisioning of instant clones is significantly faster than that of View Composer linked clones, instant clones eliminate the need for desktop maintenance windows.

Unlike linked clones, instant clones do not need to be recomposed, refreshed, or rebalanced. When a user logs out of the desktop, that desktop gets deleted and recreated. This approach to desktop deletion and recreation staggers the patching operation across desktops, eliminates boot storms, reduces storage IOPS, and creates less of a load on the vCenter Server.

Smaller Load on vCenter

Three types of VMware vCenter™ calls are made during the View Composer cloning process: cloning, power cycle, and reconfiguration. Table 2 contrasts the number and types of calls required by View Composer and instant clones.

VCENTER CALL TYPE	LINKED CLONES	INSTANT CLONES
Cloning	1 clone call	1 vmFork call
Power cycle	2 power-cycle calls	0
Reconfiguration	3-4 reconfiguration calls	0

Table 2: Loads on vCenter for View Composer Linked Clones Compared to Instant Clones

Although vmFork calls do place a higher load on vCenter than linked-clone calls, the instant-clone process eliminates the power-cycle and reconfiguration calls. The result is a significantly lower overall load on vCenter. And fewer calls per task means fewer opportunities for errors to occur.

Simplified Desktop Pool Management

With instant clones, many management tasks are automated, and architectural elements are eliminated. For example, a separate database and dedicated server are not used. The following tasks are automated:

- If you edit a pool to add or remove datastores, rebalancing the virtual desktops occurs when a new clone is created, such as when a user logs out or when you increase the size of the pool.
 - Note:** Instant-clone pools can be on the same ESXi host as View Composer linked-clone pools and can share the same datastores.
- Outdated images are cleaned up. If a pool is deleted, or when a new image is pushed, the old image is unpublished, including the running parent VMs, the internal template VM, and the replica VM or VMs. If you cancel a push operation, the primed image is rolled back and deleted.
- Automatic recovery scenarios are built in. VMware HA can power on desktops if a vSphere host fails. The instant-clone engine powers on running parent VMs if they are powered off for some reason.
- If the View Connection Server host controlling the instant-clone engine goes down, the instant-clone engine moves to another instance in the cluster.

You can easily monitor the following operations in Horizon Administrator:

- Current image, pending image, current state, and current operation are displayed on the Summary tab of the Pool Details page, along with a starting time stamp for the push operation, which is useful for troubleshooting.
- The Inventory tab of the Pool Details page identifies which desktops are using which image during a push.
- Tasks to switch users to a new image appear on the Tasks tab of the Pool Details page.

Reduced Storage Costs

Instant clones share the virtual disk of a replica VM and therefore consume less storage than full VMs. In addition, because instant clones are created at a high speed, you do not need to provision a large number of desktops up front or have many spare desktops. You can provision on demand and size desktop pools based on the number of concurrent users. For this reason and others, when compared with View Composer linked clones, instant clones can reduce the amount of hardware you need. Other reasons include

- When a user logs out, the desktop VM is deleted and a new desktop is created, automatically controlling VM growth. You do not need to refresh the desktop pool.
- View Storage Accelerator is automatically enabled. This feature uses the CBRC feature in vSphere hosts to cache VM disk data.
- Transparent page sharing is automatically enabled.
- VMFS, VMware vSAN™, and NFS disks are supported.

To help you determine storage requirements, Table 3 lists the various types of VMs that get created as part of the publishing process. The disk space usage shown is for general reference only. These numbers are not intended to provide sizing guidance for every environment. In this example, a 60-GB VM is used, but because the desktops are thin-provisioned, only 2.5 GB per desktop is used to start with.

NUMBER OF OBJECTS	VM TYPE	DISK SPACE USED
1	Master	60.0 GB
1	Internal template	55.0 MB
1 per datastore	Replica	21.0 GB
1 per ESXi host	Running parent	2.5 GB (primarily swap)
1 per desktop	Instant-clone desktop	2.5 GB (growth on use)

Table 3: Disk Usage for a Pool of Windows 7 Instant-Clone Desktops

Note: If multiple pools use the same master VM snapshot, the publishing process is performed only once. For example, if four infrastructure VMs (master, internal template, replica, running parent) are created during the publishing of one image, and you create 10 pools from that image, the same four infrastructure VMs are used for all 10 pools, as long as the pools use the same datastores and ESXi hosts.

Instant-Clone Restrictions

For Horizon 7.0, the instant-clone feature has the following limitations:

- Only Windows 7 and Windows 10 guest operating systems are supported.
- The 3D Renderer feature, which includes hardware acceleration for graphics such as vGPU, is not available.
- You can have up to two monitors. The maximum display resolution of each monitor is 2560 x 1600 pixels.
- Remote Desktop Session (RDS) host farms cannot be composed of instant clones. You must use full clones or View Composer linked clones instead.
- You cannot use Sysprep with instant clones, so all instant clones in a pool share the same system ID (SID). Some legacy applications might require a unique SID.
- Instant clones cannot reuse pre-existing computer accounts in Active Directory.
- Virtual Volumes and VAAI (vStorage APIs for Array Integration) native NFS snapshots are not supported.
- Local ESXi datastores are not supported.
- IPv6 is not supported.
- You cannot assign multiple network labels to a pool, therefore the pool cannot use multiple VLANs.
- View Persona Management is not available.
- You cannot clone automated pools of instant clones.

The workaround for these limitations is to use View Composer linked clones. Check the [release notes](#) for the most up-to-date information and to find out whether these limitations have changed.

About the Authors and Contributor

Caroline Arakelian, Senior Technical Marketing Manager, End-User-Computing Technical-Marketing Center of Excellence, VMware, and Jim Yanik, EUC Architect, EUC Technical-Marketing Center of Excellence, VMware, wrote this paper.

Fred Schimscheimer, Staff Engineer in the End-User-Computing Office of the CTO, VMware, contributed to the content of this paper.

To comment on this paper, contact the VMware End-User-Computing Technical-Marketing Center of Excellence at euc_tech_content_feedback@vmware.com.

Additional Resources

[VMware Instant Clone Technology for Just-In-Time Desktop Delivery in Horizon 7 Enterprise Edition](#) blog post

[Setting Up Desktop and Application Pools in View](#)

[View Administration](#)

[View Installation](#)



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2017 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: VMW-TWP-HOR7INSTANTCLONEDKTP-USLTR-20170210-WEB