

Virtual Machine Infrastructure Overview

Wednesday, February 3, 2016

Author: Christopher Dudley,
CTO
info@terainsights.com

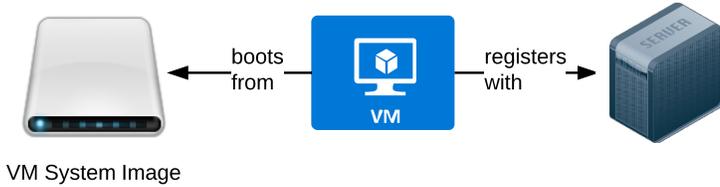
Virtual Machine Lifecycle

Stage 1: Operating System Image Creation



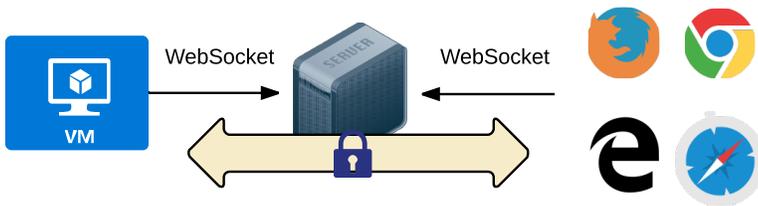
The VM's Operating system image is a thin, "copy-on-write" volume that keeps track of any differences between the VM's current state and the original Brick image. This ensures that every VM starts with the exact same initial state, and that VMs do not interfere with each other.

Stage 2: VM Creation



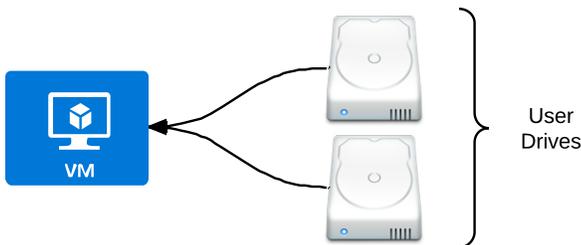
A new Libvirt domain is started with the image created in Stage 1 as the main disk. When basic initialization is complete, the VM sends a message to the server to notify it that it is ready. This message also contains some information about the VM, such as its IP address and its public key. In return, the server gives the VM the public key of the user that created it.

Stage 3: WebSocket Connection Establishment



The VM and browser both open WebSocket connections to the server, which joins the connections together. Using the public keys that were exchanged in the previous step, the VM and the browser establish an encrypted channel over the WebSocket connection. This encrypted channel is used by the browser to issue commands to the VM.

Stage 4: Drive Attachment

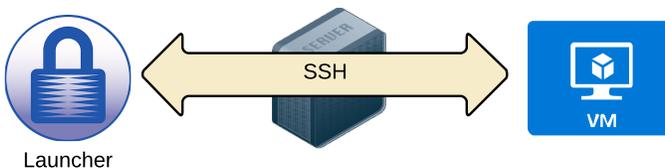


With the WebSocket connection established, any drives the user had selected for use with the VM are attached. The necessary encryption keys needed to unlock the encrypted drives are transmitted over the WebSocket connection.



At this point, the VM is considered "ready", and non-graphical VMs are ready for use immediately.

Stage 5: SSH Proxy (Graphical VMs only)

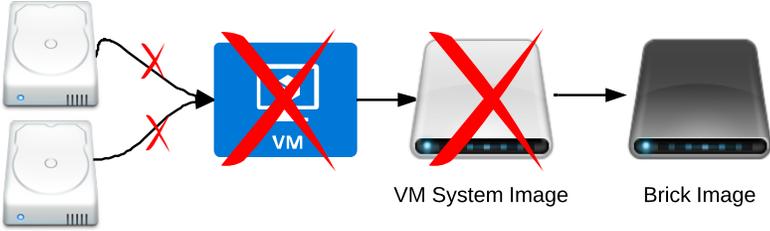


The server acts as a proxy for the SSH bridge between the launcher and the VM. This allows the VM to securely provide arbitrary services to the client. When establishing the SSH connection, the client and VM exchange keys through the encrypted WebSocket channel.



VM stays up until destroyed by the user or an administrator.

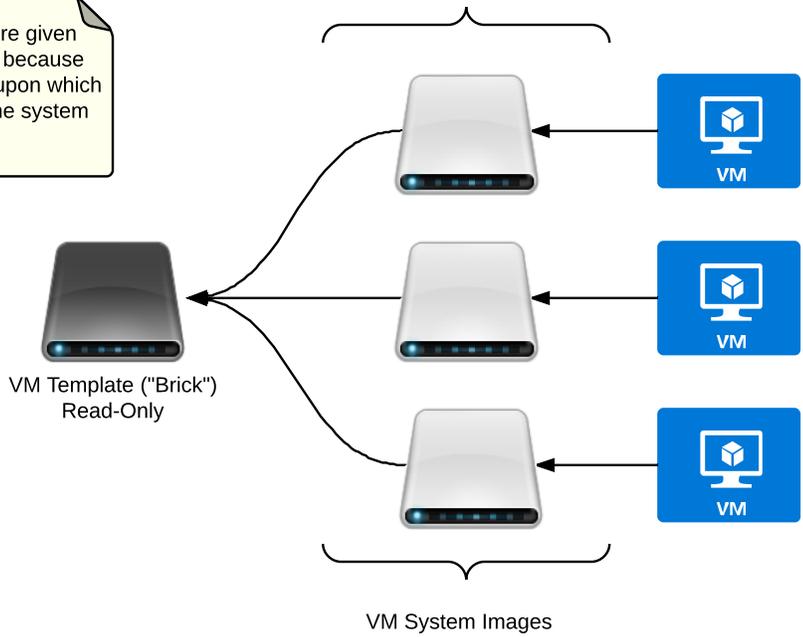
Stage 6: Shutdown and Deletion



All drives are detached from the VM, and then the VM is shut down. Once the VM has shut down, the copy-on-write system image it was using is deleted.

Virtual Machine Templates ("Bricks")

These templates were given the name "bricks" because they are the bricks upon which all of the VMs in the system are built.



Thin images on top of the brick. Contain only differences between the VM's current state and the underlying image using what is called "copy-on-write".

"Copy-on-Write" Semantics

Brick Image	System Image	Virtual Machine
C:\SomeFile.txt		C:\SomeFile.txt
	Read →	
C:\SomeFile.txt		C:\SomeFile.txt
	← Write	
C:\SomeFile.txt	C:\SomeFile.txt	C:\SomeFile.txt
	← Read or Write	
C:\SomeFile.txt	C:\SomeFile.txt	C:\SomeFile.txt

VM Loads file from disk into memory. File isn't present in the system image, so it is served from the underlying brick.

VM modifies the file

VM writes the modified file back to disk. The system image intercepts the write and stores the modified file.

Future reads AND writes to that file access only the system image. The data on the underlying brick is never changed.

Virtual Machine File System Structure

