

Clearswift Appliances and VMware

Technical FAQs

PLANNING AND DEPLOYMENT GUIDELINES

Introduction

Clearswift Email and Web Appliances are multi-process, multi-threaded applications, which run as dedicated solutions on a proprietary Clearswift hardened Linux OS.

Our appliance solutions make relatively high demands on all the underlying system resources (CPU, Memory, Disc, NIC). Simply because they were originally designed to run on dedicated hardware and not necessarily be sensitive to fluctuations in the availability of resources. For this reason the appliances make little or no attempt to throttle usage of system resources under high loads, and assume the underlying OS and hardware will self-govern access to key resources while continuing to operate efficiently at maximum capacity.

Deploying either a Email or Web Appliance as a VMware/ESX virtual machine (VM) is a recognized and supported deployment option, but care needs to be taken to ensure performance (its own or other VMs) are not unduly affected by the high demands placed upon resources.

Supported VMware platforms	Deployment types supported			
	Version	Production	Evaluation	Demonstration
VMware ESX	3.x	✓	✓	✓
VMware ESXi	3.x	✓	✓	✓
VMware Server	2.x	*	✓	✓

The following sections discuss some of the issues to be aware of and provide recommendations on how to plan and manage you Email or Web Appliance VM installations.

RECOMENDATIONS

Hardware requirements

As mentioned, both appliances make relatively high demands on underlying ESX resources. See the table below as a rough guide to overall resource demand, depending on the type of appliance being deployed:

Appliance type	Resource requirements			
	CPU	Memory	Disc I/O	Network
Web	High	High	Moderate to Low	High to Moderate
Email / Edge	High	High to Moderate	High to Moderate	Moderate

Content-scanning is by nature a CPU-intensive process, so in almost all deployment scenarios both appliances will typically run CPU-bound. CPU is thus the most critical of resources to overall

performance. The appliances demands on CPU are exacerbated by constant network activity, which in a virtual-environment equates to additional load on the CPU.

So the total number of virtual-CPU's available is typically more important than overall processing power. You can maximize the total number of vCPU's available to the VMs in several ways:

- By utilizing dual or quad core processors
- By enabling hyper-threading at the ESX BIOS level
- By using DRS clustering to pool resources from several ESXs.

Memory is the second most critical resource. The performance of a resource intensive application such as an appliance will be rapidly compromised if physical memory is not sufficient for all VMs. Memory underpins access to all other resources.

VMware configuration

When creating and configuring VMs for appliances, consider the following points:

Virtual disc sizing

Appliance VMs do not support virtual disc expansion, so it is important to allocate enough disc-space when the VM is first created. Be sure to follow the Clearswift supplied disc-sizing guidelines for the type of Appliance you are installing and to achieve optimum performance ensure the disk is defragmented before installation.

Virtual CPUs

Because of its multi-process, multi-threading design, an appliance performs best when they are running as a multi-vCPU VM. It is recommended that appliance VMs are configured with as many vCPU's as possible, whilst avoiding CPU contention.

Virtual switches

If running several appliance VMs, as a peer group, on a single ESX server - consider using a single shared vSwitch to allow inter-peer traffic to flow between VMs without hitting the physical network. This will speed up various management functions, most notably multi-peer reporting, and at the same time reduce the overall network load. Effectively you are trading network bandwidth for CPU cycles.

Virtual SCSI adapter

In common with many Linux Guest-OS's, there are known to be potential issues with the VMware BusLogic virtual SCSI adapter. VMware recommends that you use the "LSI Logic virtual SCSI adapter" for all Linux-based Guest-OS VMs.

Gigabit Ethernet

Gigabit Ethernet adaptors are recommended, but be aware that these will make higher CPU demands.

Appliance Installation and Configuration

Disc formatting

If a large network-based virtual-disc has been configured for the VM deployment, please be aware this may take significantly longer to format compared with installing an appliance on native hardware.

IMPORTANT: Do not be tempted to interrupt an appliance boot process or install wizard if it appears to stall. Allow plenty of time for the install process to complete.

TECH FAQ

NTP

Configuring an appliance to use NTP is essential. Both MIMESweeper Appliances run as many interdependent processes and are particularly sensitive to lost clock ticks from the underlying VM kernel. Clock drift on the appliances has been observed, even on moderately loaded ESX servers. The only way to control this currently is to configure the appliances to use an NTP server. This should be configured when completing the initial appliances set-up wizard or later through the appliance UI. Assuming the ESX server is already using an NTP server, the same server can be used for the appliance.

In the future, time-synchronization using VM-tools will also be an option.

Performance & Scalability

A Clearswift appliance installed and running on native hardware will always out-perform a single ESX VM running on the same hardware. On a high specification ESX server a single appliance VM will typically run at 20-50% of normal operational throughput, regardless of how many other VMs are running. There are several reasons for this, but the key ones are:

- On native hardware the appliance application/Linux-OS can utilize all cores of all available processors. On a modern high-specification server, with several quad-core processors, this might amount to 8 or 12 'virtual' processors. Running as a Guest-OS under ESX however, the appliance VM is limited to 4 virtual CPUs.
- On native hardware the appliance software (in particular the Linux kernel) may be able to directly utilize features and optimizations of the native underlying chip-set. In contrast the VMKernel emulates a standard chip-set that is well established but not so fully featured. Some or all of the benefits of the native chip-set will be lost in the emulation layer.

To counter this, it should be noted that VMware/ESX is ideally suited to running several appliance VMs in a peered configuration. A peered group of VM appliances is a highly efficient mode of operation, and additional VMs can easily be added to provide simple scalability.

Run-time optimization and management

Make use of any advanced VMware management tools and features (DRS-clustering, etc) at your disposal to help manage resources. Remember, the more CPU and memory resources in your cluster the higher the CPU and memory "reservations" you can define and the more your virtual machines are insulated from competition.

Be careful not to over commit the total physical memory, which will result in the VM kernel using its own swap-space. This will significantly impact the performance of all VMs, especially resource-intensive VMs such as the MIMESweeper appliances.

Monitor VM resources using the performance graphs available in VI Client. Check in particular the "CPU ready" graph, which will highlight CPU contention. If CPU contention becomes an issue, reconfigure your appliance VMs to use fewer vCPUs.

If clock-drift is observed confirm that NTP is enabled for both ESX and appliance. If NTP fails to keep clocks aligned, the load on the individual VM is possibly too high, see next point.

Avoid running any appliance VM with 100% Memory or CPU utilization. An overloaded VM creates additional overheads in the VMkernel, which in turn impacts the performance of all VMs. Instead, as

appliance throughput demands increase, add additional appliance VMs in a peered, load-balanced configuration.

VMware Tools

Clearswift Email and Web Appliance support for VM-tools installation is planned for the future. VM-Tools will enable better integration of the appliance Guest-OS into the VMware infrastructure. The main advantages will be:

- Improved visibility of guest-OS from within VirtualCenter
- Improved memory management and time synchronization

CONCLUSION

In summary, to run a MIMESweeper Email or Web Appliance optimally as a VM:

- Use ESX DRS clusters and resource pooling, to maximize resource availability.
- Maximize available vCPUs by using multi-core processors and enabling hyper-threading.
- When preparing VMs for appliance use, follow the Clearswift guidelines for minimum disc and memory requirements.
- Configure appliance VMs as multi-vCPU (2 or 4), but only if there is no risk of CPU contention.
- Avoid over-committing total physical memory.
- Configure appliance VMs to use Gigabit Ethernet adapters and LSI Logic virtual SCSI adapters.
- Enable NTP (both at ESX level and appliance level).
- In a multi-peer environment use a shared vSwitch to reduce network traffic.
- Do not overload appliance VMs. Instead add additional appliance VMs in a peered, load-balanced, configuration to give the scalability you need.