

Clearswift Secure Gateways

Virtual Deployment Guidelines

Issue 1.2

August 2020

Copyright

Version 1.2, August 2020

Published by Clearswift Ltd.

© 1995–2020 Clearswift Ltd.

All rights reserved.

All rights reserved. The intellectual property rights in the materials are the property of Clearswift Ltd and/or its licensors. The materials may not be reproduced or disseminated or transmitted in any form or by any means electronic, mechanical, photocopying, recording, or otherwise stored in any retrievable system or otherwise used in any manner whatsoever, in part or in whole, without the express permission of Clearswift Ltd.

The Clearswift Logo and Clearswift product names are trademarks of Clearswift Ltd. All other trademarks are the property of their respective owners. Clearswift Ltd. (registered number 3367495) is registered in Britain with registered offices at 1310 Waterside, Arlington Business Park, Theale, Reading, Berkshire RG7 4SA, England. Users should ensure that they comply with all national legislation regarding the export, import, and use of cryptography.

Clearswift reserves the right to change any part of this document at any time.

Click [here](#) to read Copyright, Trademark, and third party acknowledgments in full.

Contents

1	Introduction	4
2	Recommendations	6
2.1	Hardware Requirements.....	6
2.2	Virtual guest configuration	7
2.3	Gateway Installation and Configuration	8
2.4	Performance and Scalability.....	8
2.5	Run-time Optimisation and Management.....	9
	Conclusion.....	10

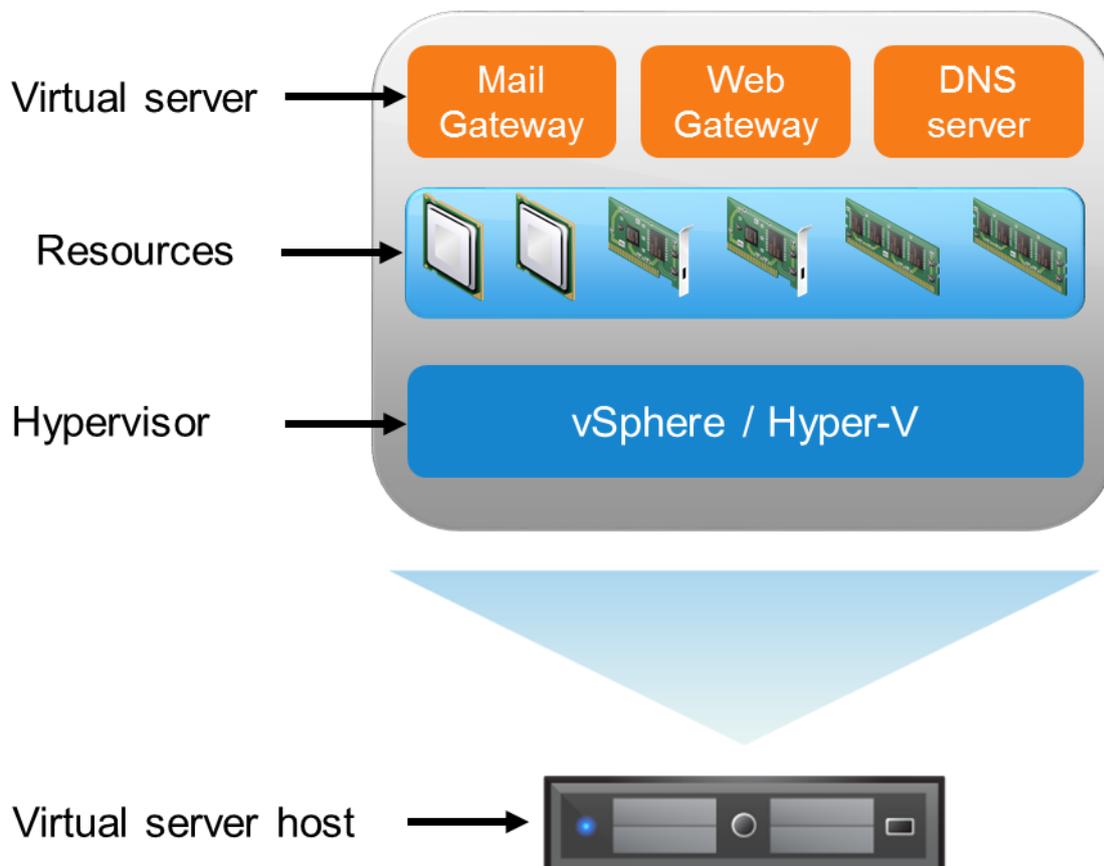
1 Introduction

The Clearswift Secure Web and Email Gateway family are multi-process, multi-threaded applications, which run as dedicated solutions on:

- Version 4.x - Red Hat Enterprise Linux 6
- Version 5.x – Red Hat Enterprise Linux 7

Our Gateway solutions make relatively high demands on all of the underlying system resources (CPU, Memory, Disc, NIC, etc.). This is 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 Gateways 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.

For clarity this document will use the following terminology:



Deploying a Clearswift Gateway as a 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.

Product	Version	Production	Evaluation	Demonstration
VMware vSphere	6	✓	✓	✓
VMware Workstation	13-14	✗	✓	✓
VMware Fusion	8	✗	✓	✓
Microsoft Hyper-V	2016 - 2019	✓	✓	✓

The following sections discuss some of the issues to be aware of and provide recommendations on how to plan and manage your Gateway VM installations.

2 Recommendations

2.1 Hardware Requirements

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

Gateway family	Resource Requirements			
	CPU	Memory	Disc I/O	Network
Web	High	High	Moderate	High
Email	High	Moderate	High	Moderate

Content-scanning is by nature a CPU-intensive process, so in almost all deployment scenarios both Gateways will typically run CPU-bound. CPU is thus the most critical of resources to overall performance. The Gateway's 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 a key system parameter and typically more important than overall processing power. In a Web family product, the system response time is also critical, so the processing power plays also an important role.

You can maximize the total number of vCPUs available to the VMs in several ways:

- By utilizing multi-core processors
- By enabling hyper-threading at the Hypervisor/BIOS level
- By using clustering or resource pooling from several deployments

Memory is the second most critical resource. The performance of a resource intensive application such as a Gateway will be rapidly compromised if physical memory is not sufficient for all VMs. Memory underpins access to all other resources, so the virtual machines must be configured according to the Gateways requirements.

Both the Email and Web Gateways family of products are very sensitive to the speed of the storage. In mid to heavy loaded systems this gets especially relevant, becoming a critical factor for the products to provide a proper response time. It must be assured that a performance similar to the expected in a physical one is provided.

2.2 Virtual guest configuration

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

- **Virtual disk sizing**

Gateway 4.x and later support virtual disk expansion, but it is recommended to create a pre-allocated disk when building the system. Be sure to follow the Clearswift supplied disk-sizing guidelines for the type of Gateway you are installing and to achieve optimum performance and ensure the disk is defragmented before installation.

- **Virtual CPUs**

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

- **Virtual switches**

If running several Gateway VMs as a peer group on a single virtual host server, consider using a single shared virtual switch 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 options of SCSI adaptors. It is advisable to research the best options for your chosen guest OS. For example, VMware recommends that you use the "LSI Logic virtual SCSI adapter" for all Linux-based GuestOS VMs rather than VMware BusLogic virtual adaptors.

- **Virtual NICs**

Whilst VMware defaults the network card to use VMXNET, it is recommended to use E1000 adaptors instead, as they are better supported by the drivers on the Gateways.

- **Gigabit Ethernet**

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

2.3 Gateway Installation and Configuration

- **Disk formatting**

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

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

- **NTP**

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

2.4 Performance and Scalability

A Clearswift Gateway installed and running on native hardware will always outperform a single VM running on the same hardware. On a high specification Virtual host server a single Gateway 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 Gateway 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 between 16 and 32 “virtual” processors.

- On native hardware the Gateway software (in particular the Linux kernel) may be able to directly utilize features and optimizations of the native underlying chip-set. 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 a virtual server host is ideally suited to running several Gateway VMs in a peered configuration. A peered group of VM Gateways is a highly efficient mode of operation, and additional VMs can easily be added to provide simple scalability. It should be noted that having **all** of the Gateway VMs running on the same virtual server host is a risk if the server itself fails.

For Gateways where the response time is critical, as it happens with the Web family Gateways, it is recommended to use resource reservation, to avoid the VMs the need to wait for available resources.

2.5 Run-time Optimisation and Management

You should make use of any virtualization drivers and management tools and features (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 Clearswift Gateways.

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 Gateway VMs to use fewer vCPUs.

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

Avoid running any Gateway VM with 100% Memory or CPU utilization. An overloaded VM creates additional overheads in the VM kernel, which in turn impacts the performance of all VMs. Instead, as Gateway throughput demands increase, add additional Gateway VMs in a peered, load-balanced configuration.

Conclusion

In summary, to run a Clearswift Secure Web or Email Gateway optimally as a VM:

- Use clusters and resource pooling, to maximize resource availability.
- Maximize available vCPUs by using multi-core processors and enabling hyper threading.
- When preparing VMs for Gateway use, follow the Clearswift guidelines for minimum disk and memory requirements.
- Configure Gateway VMs as multi-vCPU (min 2), but only if there is no risk of CPU contention.
- Avoid over-committing total physical memory.
- Configure Gateway VMs to use Gigabit Ethernet adapters and LSI Logic virtual SCSI adapters.
- Enable NTP (both at virtual server host level and Gateway level).
- In a multi-peer environment use a shared virtual switch to reduce network traffic.
- Do not overload Gateway VMs. Instead add additional Gateway VMs in a peered, load-balanced, configuration to give the scalability you need.
- Do Resource Reservation for Web Gateways in heavy loaded environments to ensure a consistent response time.