



SECURE Web Gateway v4 Sizing Guide

Technical Guide

Version 06

19/04/2016

Copyright

Version 1.0, April, 2016

Published by Clearswift Ltd.

© 1995–2016 Clearswift Ltd.

All rights reserved.

The materials contained herein are the sole property of Clearswift Ltd unless otherwise stated. The property of Clearswift 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.

Information in this document may contain references to fictional persons, companies, products and events for illustrative purposes. Any similarities to real persons, companies, products and events are coincidental and Clearswift shall not be liable for any loss suffered as a result of such similarities.

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.

Contents

1	Introduction	4
2	Overview	4
2.1	Concurrent connections	4
2.2	Sustained Bandwidth	5
2.3	Inspection policy	6
2.4	Additional considerations	6
3	Sizing guidelines	7
3.1	Hardware sizing	7
3.2	Virtual environments	7
3.3	Gateway Reporter	8
3.4	Inspection policy considerations	8
3.4.1	Lexical Analysis	9
3.4.2	Database Optimization	10
4	Sizing Examples	10
4.1	Marketing company – 2000 users – 100 Mbps Internet connection	11
4.2	Standard company – 2000 users – 100 Mbps Internet connection	11

1 Introduction

Web traffic can highly vary between organizations with similar number of users. This guide will help you understand which the demands of your users are to do a proper sizing of your web security platform.

The guide applies to v4.3 and higher of the SECURE Web Gateway.

2 Overview

When sizing a web security platform there are many different parameters to take into account. For example, based on the traffic profile of a company, the data sizes types can be completely different.

However, there are three key metrics to take into account when sizing a web security platform:

- Concurrent connections
- Sustained bandwidth
- Inspection policy

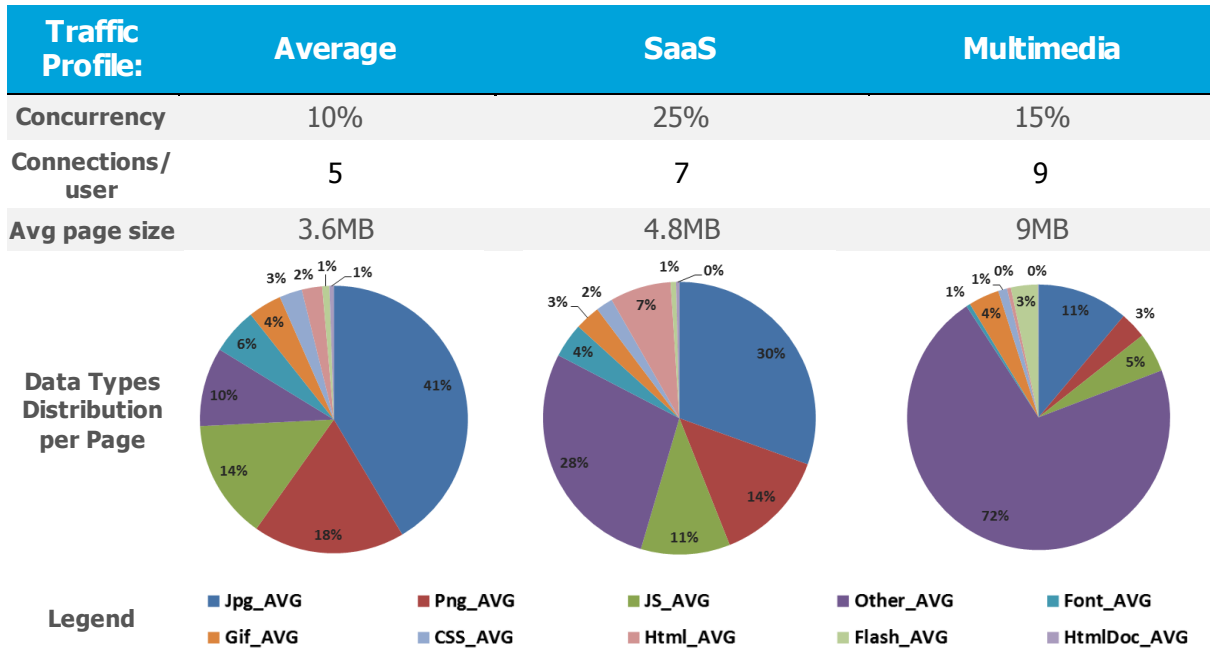
Each of them has an impact in the total performance and must be taken into account when sizing the platform.

2.1 Concurrent connections

The number of concurrent connections is usually one of the hardest values to get, even more if there is no existing web security platform. They depend on the concurrency, but also in the number of connections per connected user, both of which are highly based on the traffic profile of the organization and users' desktop configuration.

Different traffic profiles also affect the average page size and how data types are split across the traffic. This needs to be taken into account in terms of the impact when doing text analysis, as organizations with a multimedia profile might require higher bandwidth, but less performance to enforce data loss prevention policies.

Based on what has been observed on our customer base, the peak values for these parameters are summarized in the following table:



So a 1500 users organization that does heavy use of multimedia content will approximately have $1500 \times 15\% \times 9 = 2025$ concurrent connections. Whilst a similarly sized company with an average use of the Internet will only have 750 concurrent connections.

These values are approximate and must be taken as a reference rather than a definitive value. It must also be taken into consideration that the bigger the organization, the lower the concurrency tends to be.

2.2 Sustained Bandwidth

Regarding the sustained bandwidth, some considerations must be taken into account at peak times:

- Bandwidth usage is typically around 80-90% of the total available bandwidth
- HTTP/S bandwidth can highly vary, but is typically around 60% of the total bandwidth usage of an organization
- Outbound HTTP/S traffic is less than 10% of the total web browsing traffic

So depending on which type of inspection needs to be done, the bandwidth to analyze can highly vary. For example, an average company with a 100 Mbps connection to the Internet will typically require to support a 51 Mbps HTTP/S bandwidth at peak times. But if only outbound traffic is to be inspected to enforce

DLP policies, still the same traffic will traverse the gateway but just about 5 Mbps traffic will require to be inspected.

These figures should only be taken as a reference and real traffic values must be used instead to provide an accurate platform sizing.

2.3 Inspection policy

Organizations have different information security requirements. Some might require inspecting traffic bi-directionally and doing deep analysis whilst others might only need to inspect outbound traffic.

Within the inspection policy, performing lexical expression analysis has a bigger impact than looking for specific data types. And using regular expressions within lexical analysis has a bigger impact than using standard expressions.

Similarly, doing HTTPS inspection has a bigger impact in CPU usage than not doing it. Applying per user policies requires users to be authenticated, which heavily relies on the responsiveness of the directory service used for that purpose. From a resource perspective, it will impact memory, CPU and network usage.

2.4 Additional considerations

The SECURE Web Gateway can run on Clearswift provided servers, standard Intel servers¹ as well as on VMware. Regardless of the platform used, the expected performance and available resources are similar.

It is a common mistake to provide less resources or less performant ones to a virtual platform. While an architecture change like having a bigger number of instances might be suitable, the platform should be able to provide the required performance.

¹ Please refer to Red Hat Enterprise Linux 6 Catalogue for a list of certified Red Hat platforms for v4 Gateways:
<https://access.redhat.com/ecosystem>

3 Sizing guidelines

3.1 Hardware sizing

Clearswift provides physical appliances with the Clearswift SECURE Web Gateway preinstalled on them. Their performance and specifications should be used as a guidance in order to size a hardware platform:

Server Specification	Sustained Bandwidth	Peak Bandwidth	Peak Connections
(A) 1 x Intel G3430 (Dual-core 3.30GHz), 4GB RAM, 500GB SATA@7200	25 Mbps	30 Mbps	280
(B) 1 x Intel E3-1240 v3 (Quad-core 3.40GHz), 4GB RAM, 500GB SATA@7200	60 Mbps	75 Mbps	700
(C) 2 x Intel E5-2609 v3 (Hexa-core 1.9GHz), 8GB RAM, 3x300GB SAS@15000	80 Mbps	110 Mbps	1000

Important:

- The above figures are based on HTTP traffic, using a 200 Mbps Internet pipe with off-box reporting enabled and the proxy cache disabled.
- When the proxy cache is enabled, an SSD drive MUST be used. In this case the bandwidth will be lower than shown above.

In order to select the appropriate server, both the bandwidth and peak connections should be taken into account and be met or exceeded by the selected configuration.

It must be noted that a storage performance similar to the one provided by the above specifications is critical for the SECURE Web Gateway to achieve the above figures.

If the bandwidth or number of connections required are more than the limits delivered by servers specified above, multiple servers can be used to meet the requirements.

The peak columns indicate the maximum values obtainable for short durations.

3.2 Virtual environments

The Clearswift SECURE Web Gateway can be deployed as a virtual gateway running on VMware. The performance requirements are still the same, so it must be assured that the platform provides adequate resources.

The SECURE Web Gateway is also very sensitive to waiting times for resources. That means that any delay in obtaining access to resources because of the overhead of

sharing resources with other virtual machines, will result on a poor browsing experience for users.

In order to avoid some of the problems usually found in virtual environments, the following settings are recommended:

- Use VMXNET3 network adapters
- Install VMware tools for 64-bit Red Hat Enterprise Linux 6
- Use resource reservation to ensure the required resources are always available for the SECURE Web Gateway
- A bigger number of smaller VMs provide more flexibility to allocate the required resources to run the Gateways
- Ensure the VM gets a fast enough response, as this is one of the main causes of low performance for Gateways running on virtual environments

Please remember that regardless of whether the SECURE Web Gateway runs on VMware or physical hardware, the performance requirements stay the same.

3.3 Gateway Reporter

The server specification for the Gateway Reporter is determined by the amount of storage required. Storage is calculated as the product of the number of days audit data is retained and the number of transactions audited across all Gateways.

The retention period, current database size and average number of daily transactions processed during the previous seven days are all displayed under System > System settings > Report Data Settings.

Each transaction stored requires approximately 600 bytes of disk space.

Using the above you can estimate the disk space required. For example,

270,500 transactions per day kept for 60 days will require:

$270,500 \text{ transactions} * 60 \text{ days} * 600 \text{ bytes} = 9,738\text{MB or } 9.7\text{GB of disk space}$

3.4 Inspection policy considerations

Once deployed, there are some policy components and system configurations that can place additional processing demand on the SECURE Web Gateway, affecting performance. The following section highlights these areas and provides guidance on best practice.

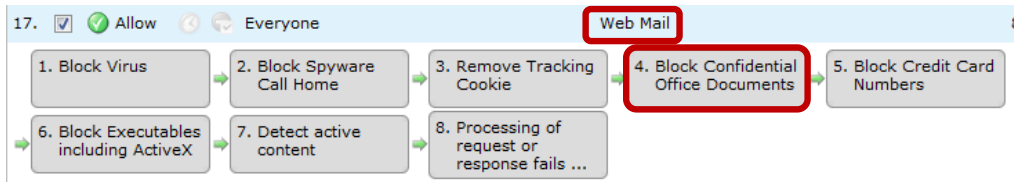
3.4.1 Lexical Analysis

The lexical analysis content rule is very powerful and can be used to identify key words and phrases within web content and file attachments. This rule also allows complex regular expressions capable of identifying patterns within the text – e.g. customer reference numbers – to be defined. Regular expression processing requires more CPU power than searching for simple keywords such as 'Top Secret'.

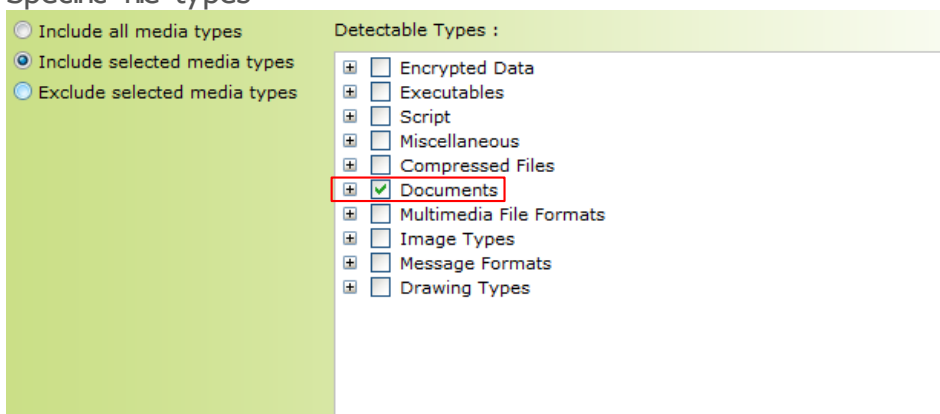
The SECURE Web Gateway allows the textual searching to be targeted at particular parts of the web transfer rather than searching all the web content. By being more specific about site type, file type, location within documents and desired search direction, processing overheads and risk of identifying false positives can be reduced. For example, you only need to search outbound web traffic for sensitive phrases related to confidential business information.

To reduce performance overheads associated with textual searching, consider how you can limit the areas searched to:

- Particular types of sites and documents



- Specific file types



- Web page or document content, URL, HTTP header or even the header, footer and properties of the document.

Note: Selecting 'HTTP header' and/or 'Request URL' is rarely needed. Searching every HTTP header and every URL for a phrase will impact on performance, therefore only select these after careful consideration.

Lexical Expression

Scan the following parts of the HTTP conversation :

HTTP Headers
 Request URL
 Content

Using the expression list and trigger conditions below :

Expression list : Confidential Material

Threshold : 10

Trigger : Greater than or equal

Document options (where applicable) :

Scan body
 Scan header and footer
 Scan properties

- Direction – although inbound inspection might be required, data only leaks out.

Direction To Apply

Where the item was detected

leaving the company (uploading)
 either leaving or entering the company
 leaving the company (uploading)
 entering the company (downloading)

3.4.2 Database Optimization

There are two aspects to database optimization:

1. Rebuilding the database indexes:

By default the index rebuilding is performed weekly, on Saturday at 21.00 hours. This day and time has been selected because it's out of hours and therefore doesn't impact the performance of the web proxy.

2. Shrinking the database:

Database shrinking means releasing redundant disk space occupied by deleted rows in the database. This option should not be enabled unless explicitly instructed to do so by Clearswift Customer Support.

4 Sizing Examples

The following sizing examples use the guidelines provided in previous sections to choose the appropriate hardware platform and instances for the organization requirements.

4.1 Marketing company – 2000 users – 100 Mbps Internet connection

Being a Marketing company, they make heavy use of online multimedia content, so fall under the Multimedia profile. Their Internet link is 100 Mbps, and the utilization is 80%. Out of this traffic, 65% is caused by browsing traffic.

According to the above requirements, this company will require:

- **Connections:** $2000 \text{ users} * 15\% \text{ concurrency} * 9 \text{ conn/user} = \mathbf{2700}$
- **Bandwidth:** $100 \text{ Mbps} * 80\% * 65\% = \mathbf{52 \text{ Mbps}}$

The required sustained bandwidth could be supported by one (C) type server. However, the number of connections is over the supported ones for the same server, so 3 x (C) servers would be required.

4.2 Standard company – 2000 users – 100 Mbps Internet connection

This company doesn't have any special activity that continuously requires Internet and follows a Standard profile. Their Internet link is 100 Mbps, with a utilization of 95%. A big part of this traffic is caused by a WAN connection to a regional office, leaving 55% of this traffic used by browsing activity.

The requirements for this company are as follows:

- **Connections:** $2000 \text{ users} * 10\% \text{ concurrency} * 5 \text{ conn/user} = \mathbf{1000}$
- **Bandwidth:** $100 \text{ Mbps} * 95\% * 55\% = \mathbf{52.25 \text{ Mbps}}$

Both the number of connections and the required bandwidth are in line with what a (C) server can provide, so one of them is enough to deal with their traffic requirements.