Replify Quick Start Guide

Version 6.4

Contents

1 Overview

This document assumes that you have already installed a Replify Accelerator Virtual Appliance and wish to configure this to optimize traffic. If not, please refer to the Replify Accelerator installation guide.

All configuration will take place using the Web UI of the Virtual Appliance. The default login credentials for this are:

- Username: **admin**
- Password: **default**

If you have installed the Virtual Appliance directly on your own Debian server, the username and password will be any local Linux user who is in either the **replify** or the **root** system groups.

The following basic tasks are required to get Replify Accelerator up and running in every deployment.

1. License the system
2. Add at least one application server
3. Connect the Virtual Appliance to a remote Virtual Appliance or Accelerator Client

However, many more configuration options are available and can provide additional optimization in certain circumstances. Please read the User Guide or contact Replify Support for further details.

2 Licensing

A license is required for every Replify Accelerator deployment. A licence key is provided by Replify Support and can be used to licence a system or to update the licence on an existing system. The licence can easily be upgraded without impacting users on the system.

There are two mechanisms to license a Virtual Appliance. These are described below.

2.1 License Directly

You can license a Virtual Appliance (VA) directly using a licence key provided by Replify Support. If you re-deploy a Virtual Appliance with a production licence, you will need to request an updated licence key.

You can enter the licence key on the **Licensing** page under the System section of the Web UI.

2.2 License via a Replify Enterprise Manager

If there are multiple Virtual Appliances in a deployment, they can be licensed using a Replify Enterprise Manager (REM). This allows a single licence key to be used and means that Virtual Appliances can be moved/re-deployed etc without requiring a new licence key.

A licence key can be added to the REM on the **Licensing** Page under the System section of the Web UI.

A VA can be informed that it should be connected to a REM using the **Replify Enterprise Manager** page under the Configuration section of the Web UI.

3 Connecting Remote Accelerator Components

For Replify Accelerator to be able to accelerate traffic across a network link, there needs to be a Virtual Appliance on the side of the link that has access to the Application Servers to be optimized and either a Virtual Appliance or an Accelerator Client at the remote side. Details of both deployment models are described below.

Note that both deployment models are equally valid and a mixture of both can be used in a single deployment. The choice of which to use is normally determined by the customer's environment and which is easier to deploy.

3.1 Replify Peered Virtual Appliances

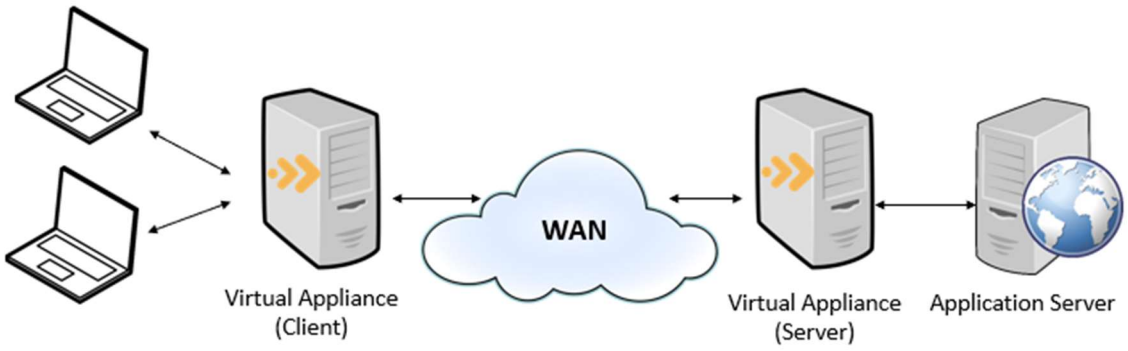

The peered appliance deployment model can be used when there are several devices requiring optimization at a single location. In this model a Virtual Appliance is deployed at both ends of the network link. Traffic from the site that is remote to the Application Servers needs to be routed via the local Virtual Appliance using one of the following mechanisms.

- Install a Replify Accelerator Client on each end user device and connect this to the client-side VA. (Note Client Location Awareness should be set to "All Clients are Local" on the **Client Location Awareness** screen in the Configuration Section of the Web UI.
- Install the Virtual Appliance inline on your network as a bridge.
- Configure the Virtual Appliance as a WCCP Client
- Use the Virtual Appliance as a default gateway
- Use a static route or Policy Based Routing

Contact Replify Support for more information on each of these approaches.

3.2 Replify Accelerator Client and a Virtual Appliance

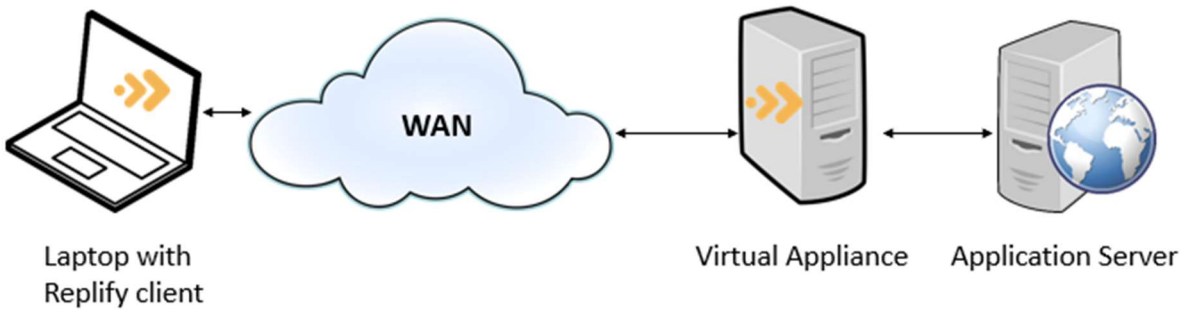

In this deployment, clients are installed on individual end user machines and these are configured to connect to either a Virtual Appliance or to a REM. This can be useful when clients are on mobile devices such as laptops or where the infrastructure isn't available to install a peered appliance.

Configuration can be completed using the Client's User Interface. Additionally, the following mechanisms can be used to configure the Client on Microsoft Windows environments:

- Installing with a pre-configured VA/REM using MSI parameters
- Pre-configuring the Client MSI installer to automatically connect to a VA or REM.
- Use Microsoft Windows Group Policy to configure the REM that the Client will connect to.

Contact Replify Support for more information about each of these mechanisms.

4 Application Servers

Replify Accelerator does not accelerate any traffic unless explicitly instructed to. Traffic is only accelerated if it matches an **Application Server** that has been defined in the **Application Servers** page in the configuration section of the web UI.

An Application Server has an IP address associated with it. This can either be a single IP address or a subnet in CIDR notation, for example 10.0.0.0/8

For an application server to be accelerated effectively, there should be a much faster connection between the Virtual Appliance and the Application Server than there is between the VA and it's associated remote VA or Replify Accelerator Client.

If using a peered VA deployment, the Application Server should **only** be added to the VA that is on the same side of the WAN as the Application Server.

Note that if you choose the Transparency or SSL Optimization options for the Application Server, you should carefully follow the instructions shown in the popup dialog that is displayed on the UI. Failure to do so may cause data connections to be blocked.

4.1 Application Servers that use SSL

One of the reasons SSL is used on a network is to ensure that traffic cannot be intercepted by a "man in the middle". This means that for Replify Accelerator's optimization algorithms to work on SSL traffic, it needs to be configured so that it is a trusted "man in the middle". There are two ways of doing this.

4.1.1 Provide Replify Accelerator with the application server certificate

If the server's SSL certificate and private key are available, these can be uploaded to the VA on the SSL Management page. This means that when Replify Accelerator decrypts the traffic to apply its optimization algorithms, it can re-encrypt it for the client using a certificate that the client deems acceptable.

4.1.2 Configure end user devices to trust the VA as a CA Authority.

If the option to use an auto-generated certificate is selected, the Virtual Appliance will generate its own certificate for encryption/decryption purposes, however this will be rejected by client browsers because they don't trust the validity of these certificates.

To circumvent this, each end user device that will be receiving traffic undergoing SSL optimization will need to add the Virtual Appliance's CA certificate to their list of trusted local Certificate Authorities. This can be downloaded from the **SSL Settings** page on the Configuration section of the web UI.

This is required in all deployment scenarios that require SSL optimization using auto-generated certificates.

5 TCP Optimization

If using Replify Accelerator on a network link that has high levels of congestion, high latency or packet loss, it is worth enabling some TCP optimizations.

5.1 WAN Connection Pooling

This option should be enabled on high latency connections to improve the response time for applications. WAN connection pooling maintains a pool of TCP connections between Replify Accelerator endpoints that can be re-used. This means that each request across the link can take place without requiring a TCP handshake.

WAN Connection Pooling can be enabled on the **Settings** page on the Configuration section of the Virtual Appliance UI.

5.2 Use Alternative Congestion Control Algorithm

Using an alternative congestion control algorithm can improve throughput on congested connections or on connections with high latency or packet loss.

An alternative congestion control algorithm can be configured on the **Settings** page of the Configuration section of the Virtual Appliance UI when using Replify Accelerator Clients or on the **Peered Accelerator Appliances** UI when using peered appliances.

It is recommended that **BBR** is used on higher bandwidth networks that are either congested or suffer from packet loss.

It is recommended that **Hybla** is used on satellite networks or other high latency networks.

Note that the best algorithm to use is very dependent on individual network conditions and therefore testing should be carried out with each algorithm to ensure that there is no performance degradation compared to using the default **Cubic** algorithm.

6 Testing Replify Accelerator Functionality

When Replify Accelerator is configured to accelerate traffic, the amount of optimization that is achieved can be determined using the web user interface.

6.1 Offload and Performance Improvement

The primary metric that is used by Replify Accelerator to measure the amount of optimization obtained is “WAN Offload”. This is a measure of the reduction in data that has been sent across the WAN due to Replify Accelerator being used. For example, if you accelerated 1GB of data and the amount of data that was transmitted between Replify devices across the WAN was 200MB, this would be described as an offload of 80%.

The Virtual Appliance can measure this offload easily by comparing the amount of data it receives from the LAN and the amount that it transmits across the WAN.

The performance improvement that is seen by an end user is difficult to determine by the Virtual Appliance. This usually correlates with the WAN offload measured by Replify Accelerator but may differ.

When Replify Accelerator’s TCP optimization features are enabled, this may improve the response time on links with high latency or packet loss, however they will not affect offload. These features can provide benefits even when the amount of offload is 0%

There may be scenarios where the bottleneck is not network related. For example, an under-resourced application server or malware running on the end user device.

Ultimately the only way of accessing the benefit of Replify Accelerator is to test it with a realistic user scenario using realistic network conditions.

6.2 Measuring offload

The offload is prominently displayed on the top right-hand corner of the Virtual Appliance Web UI, however there are several reports that can give you more detail about the optimization that is occurring.

6.2.1 Accelerator Clients

This report allows you to see the offload achieved for each client.

6.2.2 Live traffic

The live traffic graph shows a real-time view of the optimization that is taking place. The yellow graph shows the bandwidth being used by the WAN interface and the red graph shows the bandwidth being used by the LAN interface. On a system with high offload the red graph will be using much more bandwidth than the yellow graph.

6.2.3 Optimized Sessions

This report allows you to see the offload delivered for individual TCP connections.

6.2.4 Bandwidth Savings

This allows you to get an overall view of the amount of bandwidth saved for different protocols.

7 Adding Extra Disk Space for the Cache

Replify Accelerator maintains a cache of data that has been previously seen. This is used to reduce the amount of data traversing the WAN.

The amount of space used by the cache is configurable and the Virtual Machines supplied by Replify have enough disk space for most POCs or smaller deployments, however for many deployments extra disk space will be required.

7.1 Disk space requirements

Each client or peered Virtual Appliance will have a cache associated with it and the size of each cache is configurable. The required disk space should be at least 5GB greater than the total cache size. For example, if there are ten clients with a 256MB cache and one peered VA with a 10GB cache, the total cache size will be 12.56GB and the required disk space should be at least 17.56GB.

For more details on resource requirements please contact Replify Support.

7.2 Available disk space on Replify Virtual Machines

The VMWare and Hyper-V Virtual Appliance images provided by Replify have approximately 14GB of space available for the cache and the KVM/QEMU images have 48GB.

If this is not sufficient, an extra disk should be added to these machines for the purpose of storing cache data. **Note, extending the size of the current virtual disk is not supported by Replify.**

When the disk is added, this space can be made available either by rebooting the Virtual Machine or by logging in to the console and running:

```
add-disk-for-cache
```

7.3 Ensuring sufficient space on native Linux installations

If installing natively on Debian or on a cloud platform or other Hypervisor type, the cache will be stored at **/replify/db**.

If an extra disk, LVM volume etc is required to store the cache, this should be mounted at **/replify/db** and permissions should be set on this so that the **replify** Linux user can read and write to this disk.

8 TLS Certificate Management

Several TLS certificates are generated whenever a Replify Accelerator Virtual Appliance or Enterprise Manager is installed for the first time. These will be generated with default values and will be unique for each instance.

These certificates can be replaced or customised as necessary

8.1 Certificate Used to Access HTTPS Web UI

When a user accesses the web GUI over HTTPS, the server certificate that is generated at installation will be presented. This will have two Subject Alternative Names. One is the 'hostname' and one is the 'fully qualified domain name' of the machine. It will also be signed by the VA's CA certificate, so by default will not be trusted by most browsers.

Note that, if using Dynamic SSL functionality, your browser will need to trust this CA certificate, and therefore will trust the corresponding VA server certificate.

The server certificate has a lifetime of thirteen months. This is because most browsers mark certificates with a longer lifetime as invalid.

8.1.1 Using Your Own Certificate

If you want to use an alternative certificate instead, you can copy your own certificate and key in PEM format to the VA and use this by running the following commands:

```
replify-ctl set-configuration-value certfile PATH_TO_CERTFILE string
replify-ctl set-configuration-value keyfile PATH_TO_KEYFILE string
```

Note that the VA service should be restarted with '**replify-ctl restart'** before this takes effect.

8.1.2 Regenerating the Default Server Certificate

To regenerate the certificate that is used by the VA or REM, you can run

```
replify-ctl generate-new-server-certificate
```

This will ask you to specify each of the attributes that make up the certificate. The important values for HTTPS access are the '**Subject Alternative Name'** and the '**Common Name'**. These should be set to the hostname that you are using to access the VA web UI in your browser. For example, 'va.acme.com' or '*.acme.com'.

8.2 CA Certificate

Each Replify Accelerator Virtual Appliance and Enterprise Manager has a CA certificate associated with it. This is used for two purposes: to generate the server certificate; and on a Virtual Appliance, to generate dynamic certificates used for SSL optimization. The CA certificate has a lifetime of fifteen years.

8.2.1 Using your own CA certificate for dynamic SSL

To use your own CA , run the following commands

```
replify-ctl set-configuration-value replify_ca_cert PATH_TO_CERTFILE string
replify-ctl set-configuration-value replify_ca_key PATH_TO_KEYFILE string
```

8.2.2 Regenerating the default CA certificate

To regenerate the CA certificate that is used by the VA or REM, you can run

```
replify-ctl generate-new-ca-certificate
```

This will ask you to specify each of the attributes that make up the certificate.