

# Running a Successful Proof of Concept



1.1 Approved

26<sup>th</sup> July 2011

© Replify 2011

## Table of Contents

1. You're not alone.....	3
2. Introduction .....	3
3. Is Replify going to help you? .....	3
4. Want to do a quick and simple test? .....	4
5. Who and What? .....	4
6. Routing your traffic .....	5
7. Understand your network.....	5
8. Not using a VPN?.....	6
9. Sizing and configuring the system .....	6
10. Success Criteria .....	7
11. Gotchas .....	7
12. Not seeing any benefits?.....	8
The VAs are under-resourced. ....	8
The traffic isn't going through Replify. ....	8
The traffic cannot be optimized.....	8
The client isn't aware of its location .....	8
Transparency is required .....	8
13. Quick Start Steps.....	9
11. Replify Installation Checklist .....	9

# Running a Successful Proof of Concept

---

## 1. You're not alone

If you follow the advice in this document you'll likely have a straightforward proof of concept. You're not alone though – our pre-sales and support teams are there to help, and if the guidance doesn't cover your use case, or you still have issues after following the steps, give us a call please.

## 2. Introduction

So, you're considering evaluating the Replify Accelerator product. This probably means that you're a distributed enterprise with one or more remote branches, and you almost certainly have people that work remotely from the office – perhaps home workers, or people who travel on behalf of the business. There's also a fair chance you're one of the rapidly growing number of organizations who are adopting cloud services in some form. And presumably you're experiencing one or more of the following problems:

1. Users complaining that the company's applications are too slow to use remotely.
2. Users having trouble accessing the corporate web and other collaboration sites.
3. Users having difficulty in upload or downloading documents, drawing, designs, databases, software updates etc.
4. Users complaining that at certain times of the day they have reduced connectivity because something else e.g. email, is clogging up the connection.
5. Unacceptably high bandwidth costs because you've tried to fix the problems above by buying more bandwidth. And worse, it may not be making much difference.

It is very likely that Replify Accelerator can help with these problems, and perhaps eliminates some altogether, but you want to satisfy yourself that this is the case, and that's the right thing to do. So keep reading.

## 3. Is Replify going to help you?

Replify can work with almost any kind of data carried over a TCP connection. It understands some application protocols specifically and delivers specific optimizations. Those protocols are:

- HTTP and HTTPS
- CIFS (the Windows file transfer protocol)
- MAPI and MAPI over HTTPS
- FTP
- The backup protocols used by DoubleTake and Dell EqualLogic

We're adding more protocols, but even if you're using another protocol not on our list, Replify can still apply the technique of de-duplication which eliminates the sending of blocks of data which have already traversed the connection in either direction. So sites that are visited repeatedly, documents and other artifacts which are exchanged (even after modification), back-ups and other transfers which work with similar data on each operation, will all result in far less data being sent over the network.

# Running a Successful Proof of Concept

---

The final technique which Replify Accelerator applies is compression – and for many types of data this can reduce data volumes by 50% or more.

For most organizations the techniques above will apply to most of the troublesome traffic on their network, but to be clear, there are connection and data types which Replify does not optimize, or accelerates to a lesser degree. Specifically:

- UDP – Accelerator does not optimize UDP traffic types such as VOIP
- RDP – Accelerator does not intercept and accelerate the screen drawing component of RDP traffic such as Citrix, but it can help greatly with associated file transfers and print jobs.
- Video and Audio – most video and audio formats are already highly compressed and hence the compression engine with Replify may only decrease file sizes by a small additional amount. We do de-duplicate though, so if the same media is passing back and forth we'll provide great offload.

This is a good time to discuss offload and acceleration. Replify Accelerator can dramatically reduce the amount of data being sent across the network. It can also shorten the length of time that it takes for a user to carry out an operation such as copy a file or access a web page. Offload and Acceleration often go hand-in-hand and you get them together, but as the speed of the connection goes up, and the latency goes down, there is less and less room for Replify Accelerator to increase the speed of the interaction so the acceleration becomes less noticeable. The offload is constant though and this is almost always still valuable.

Replify Accelerator, when deployed on a suitably sized platform, will provide acceleration on connections up to around 100Mbps, but the benefits beyond that are confined to offload. The greatest benefits will be seen on connections on slower connections (20Mbps or lower say) with latencies of 50ms or greater.

## 4. Want to do a quick and simple test?

In that case, contact our sales team and we'll give you remote access to a Replify virtual appliance running in the Amazon cloud. All you need to do is to install a client on a PC and then you can:

- Try copying files of any kind you like up and down to a file share
- Access a test SharePoint server and copy files up and down from shared repositories

This gives you a very quick and easy way to observe the levels of offload and acceleration you can get. That may be all you need to understand the benefits that Replify would give remote workers in your organization. If not, read on.

## 5. Who and What?

So what problem are you trying to solve? Are you:

- 1) Trying to improve the performance of a link between two or more locations,
- 2) Trying to make remote users more productive (and happy)
- 3) Both

# Running a Successful Proof of Concept

---

If your focus is solely on sites, then you probably want to deploy two Virtual Appliances for your initial trial: one in the Data Center (or cloud) and one in the test site. If the site is very small (less than say 10 users) you may just want to install an Accelerator Client on each machine and just have a VA (and the REM) in the data center.

If your focus is on remote or roaming users then you only need the VA (and the REM) in the data center(s) and an Accelerator Client on each test machine. Make sure those machines are Windows XP, Vista or 7, and patched with the latest service pack. Windows 7 is preferred – see the section on Gotchas.

Having decided on the problem you're trying to solve, choose which applications you want to look at first. Is it site-to-site disaster recovery (e.g. Dell EqualLogic or DoubleTake), is it collaborating users perhaps centered on SharePoint, is it emailing large attachments?

Don't boil the ocean– pick a subset of your total problem applications and start with that. It's very easy to add new services (i.e. applications and data sources) to the system later.

## 6. Routing your traffic

WAN optimization relies on intercepting the traffic to work its magic which means you need to arrange for the traffic between the users and the servers to be re-routed. There are several ways to do this, but it's very important to make the right choice up front. These are your options:

- 1) Static routing. This isn't very flexible but it's quick and simple and works for a small number of client machines when you have installed a local VA. Just add the static route to the hosts table. But beware; in some environments we've seen routers ignore the static routes.
- 2) WCCP (Web Cache Communication Protocol). This is a Cisco defined protocol to handle re-routing of application traffic to local cache servers and it works very well with Replify. But... WCCP is a complex beast and unless you're already familiar with it, this is not the time to start learning it.
- 3) Accelerator Client. Even if the client PCs in question are never leaving a site with a VA in it, you can install an Accelerator client on them and it will transparently re-route the traffic to a local VA. Simple. And if the PC in question does roam around, even better because the client is smart and will know when it's not located in a branch with a local VA, and it will provide direct WAN optimization to remote apps all by itself when it's not.

## 7. Understand your network

When PoCs get into difficulty it's usually because the traffic isn't flowing correctly through the Accelerator. Firewalls are a frequent obstacle – they need to pass traffic on the ports that Replify is using. We've made it simple for you by allowing you to define just two ports for Replify to use – one for regular tcp traffic and one for encrypted (SSL) traffic. You can use any ports you like, including port 80 which is probably already open for web traffic.

# Running a Successful Proof of Concept

---

If you want to use Replify with SSL traffic, that's not a problem, but you'll need to install your SSL certificates on to the Data Center virtual appliance. That's an easy thing to do, and perfectly safe as long as you have placed your VA into the same secure environment as the application server itself.

If you have a very "smart" firewall, which uses DPI to check that the traffic on a port such as port 80 is really HTTP, then it's not going to be happy with Replify using that port. It's very straightforward to verify that the Replify Accelerator clients can see the Virtual Appliance, that peered Virtual Appliances can talk to one another, and that the Data Center (or Cloud) VA can talk to the application servers. Always start by verifying the connectivity and then begin user testing.

## 8. Not using a VPN?

Typically, you'll be using a VPN between remote clients/sites and your data center, but you may have reasons for not doing that. In that case, you have to take account of the Network Address Translation going on in your gateways. Remote Clients or VAs need to be configured with the *external* IP address of the REM, and the REM needs to know the external address of the VAs in order to communicate them to any Client or remote VA connecting to it. The easiest way to do this is to determine the external address of the REM (by configuring your gateway router) and then use this address with your VAs, *including* any VAs in the same site as the REM. Because the local VAs have the external REM address, their request to connect to the REM will bounce off the gateway and arrive at the REM with the VA's external address as the originating IP address. The REM will therefore store the correct VA address to share with clients and remote VAs.

## 9. Sizing and configuring the system

There's no point in installing Replify Accelerator on a small server, configuring it to intercept all the tcp traffic going over a 1Gbps link, and then standing back to see what happens. The system won't cope well. Much better to start by picking a small subset of your application traffic – perhaps a single troublesome files share, or a poorly performing web application. The Service configuration screens in Accelerator make it very easy to choose individual application servers, and add more as and when you need to.

Replify is deployed as a virtual appliance, so the resources that are allocated to this can be increased as more applications are added to the deployment. We support VMware, Hyper-V, Xen Server and KVM virtualization platforms. You can run on bare metal too, but if you're planning on doing that you should have a word with us first.

The minimum configuration we recommend for a VA is:

- 2 virtual CPUs
- 2GB RAM
- 4 GB disk space

It will support around 20 users on a connection up to 10Mbps.

# Running a Successful Proof of Concept

---

Disk speed is important; slow, or contended, disks will limit the data rates which the VA can handle. A disk speed of 7600 rpm is suggested as a minimum. RAID striping is even better, and if you've got HDDs, well lucky you.

The high level steps for deploying a solution for a POC are as simple as:

1. Deploy the Appliance on your virtualization platform
2. License it
3. Configure one (or more) application servers
4. Deploy a peered VA if it is a multi-site scenario, establish a peering relationship and configure client location awareness (If using Accelerator clients).
5. Install the client
6. Configure the client with a REM address

The Replify user guide at <http://www.replify.com/guide> will give more details on each of these steps or you can contact [support@replify.com](mailto:support@replify.com) for more help.

The choice of client cache size is important in a PoC focused on individual users. Most PCs these days have well over 100GB of hard drive, so it's reasonable to allow 500MB or more of cache space. If you're copying large files, for example 2GB video files, then make sure the cache is at least as large as the size of file you want to work with. If it's not then the cache will "wrap" on each transfer, overwriting the cached data, and de-duplication just isn't possible.

Note that each client cache needs to have a corresponding cache of the same size on the virtual appliance, so you need to take this into account when determining cache sizes. In general, the more disk space that can be allocated, the better.

## 10. Success Criteria

What are you hoping to see? Offload? Acceleration? Both? Be clear at the outset what you need from the product. On a low latency link end-users may not see much increase in responsiveness and if you've set high expectations you could be disappointed. But you may still get great offload, and that could mean more users on a connection, or using a smaller and cheaper connection.

So upfront quantify the level of offload that you need, and that it's realistic to expect. For back-up traffic, halving the volume of data and the back-up time is a huge improvement, so 50% or so is a realistic goal. For designers bouncing large drawings back and forth, offload of 80% or more is entirely achievable.

Maybe you have a problem that file transfers frequently fail and have to be re-tried. Well the de-duplication engine in Replify means that failed transfers should pick up from where they left off, so that might be the benefit you need to see.

## 11. Gotchas

Try as we might, we can't make Replify Accelerator work with every other piece of software out there. On Windows XP and Windows Vista, our client uses a component called DNE to intercept the

# Running a Successful Proof of Concept

---

tcp traffic. Other applications also use DNE and some of them are unfriendly. We have problems with versions of SonicWall and CheckPoint VPNs which can't be fixed on our side. We may be able to advise on reconfiguring those applications so they can coexist with Replify, so give us a call. On Windows 7 we use a different approach, called WFP, and that works well with everything we've come across so far.

## 12. Not seeing any benefits?

So you've followed all my great advice, but you're not seeing any of the expected benefits. What's going wrong? Well assuming that you're not trying to achieve the impossible and accelerate uncompressible UDP traffic on a 10GB link, there are a number of possible reasons:

### The VAs are under-resourced.

If this is the case then there will be tell-tale signs in the management screens – high RAM usage, dropped writes, cache misses, timeouts. Take a look at the info screen to check resource usage and uptime.

### The traffic isn't going through Replify.

This could be a routing problem, a firewall problem or a client interception problem.

Follow these steps to verify that you have connectivity.

- If you have reconfigured traffic to be routed via the Accelerator, does looking at the output of a Traceroute command to the application server have the VA listed as a hop?
- Do your firewall logs show traffic being blocked?

### The traffic cannot be optimized

Some software vendors have realized that their software will be used over a WAN and have designed it to work well in this scenario through the use of compression techniques; alternatively they may encrypt the data going across the WAN so that it cannot easily be inspected by devices such as Replify. In these scenarios, a simple configuration change to the application may be all that is required, or Replify may need to have the SSL certificates installed.

You might have enabled SMB signing in your network; it's intended to prevent man-in-the-middle attacks but it's not very useful and slows your network down, and... we don't support it.

### The client isn't aware of its location

The Replify client can be made aware of its location so that it knows not to optimize traffic when in the data center, sent traffic to the local VA if it is in a branch office or send it directly to the data center VA when out on the road. However for this to work correctly each VA needs to know what subnets are considered local.

### Transparency is required

It could also be that you need to enable "transparency" in the Virtual Appliance to hide the re-routing from the applications – but don't turn it on without reading the user guides and understanding what you're getting into.

## 13. Quick Start Steps

### Pre-requisites;

In the main/data center site: an existing VMware/Hyper-V Host to run two virtual machines;

- 1 x Replify Enterprise Manager (1 x vCPU, 512GB RAM, 5GB storage space)
- 1 x Replify Virtual Appliance (1 x vCPU, 2GB RAM, 40GB storage space), this disk space will increase by 30GB per remote connected site

Branch/remote sites:

- Either Replify Client software deployed/installed to all PC's. Or
- 1 x Replify Virtual Appliance (1 x vCPU, 2GB RAM, 40GB storage space)

Here are the steps to start your proof of concept:

1. Prepare/build host servers
2. Register for download links on Replify Website ([www.replify.com/download/](http://www.replify.com/download/))
3. You will then receive an email with the download links
4. At this point are ready to build/import the Enterprise Manager and Virtual Appliance
5. Assign static IP addresses to new Virtual Appliances
6. Configure disk allocation on Virtual Appliance
7. Deploy Virtual Appliance onto branch/remote servers
8. Configure Replify application acceleration and test site to site connections
9. Configure local client to talk to local Virtual Appliance
10. Carryout testing of optimized solution/access to profiles etc.

## 11. Replify Installation Checklist

Scenario 1 – Central Virtual Appliance and Mobile Clients

### Replify Enterprise Manager

1. Replify Enterprise Manager deployed, static IP address configured and able to log into the Web site to access configuration and reports etc.
2. Appropriate hypervisor tools installed in guest operating system
3. Trial license obtained and added to Enterprise Manager Configuration
4. Re-set root password

### Replify Virtual Appliance

1. Replify Virtual Appliance deployed, static IP address configured and able to log into the Web site to access configuration and reports etc.
2. Additional disk size calculated and added (also need to run 'add-new-disks' within the console)
3. Check/verify that the VA is being licensed and controlled by the Enterprise Manager
4. Configure the default initial cache size, clients can be increased up to 2.78GB
5. Increase the cache size of any existing client or peered VA's

## Running a Successful Proof of Concept

---

6. Add application server (either individual servers or subnet range, if using subnet which includes addresses of the Replify Enterprise Manager and Virtual Appliance make sure to add exclusions for these appliances.
7. Check that client location awareness is configured correctly, in this specific scenario it must be set to 'All Clients are Remote'
8. Re-set root password

### **Replify Client**

1. Check that client is installed and the Replify service 'Replify Accelerator Client' appears and is started
2. Check that client is connected to Replify Enterprise Manager
3. Check that the Virtual Appliance is showing in the client under Edit Appliances
4. Check that the Appliance is 'connected', i.e. green light and that it is displaying a list of accelerated applications (i.e. the ones added above)
5. Open a new session to one of the application servers and then open Connection Spy and verify that the traffic is being intercepted
6. Login to the Virtual Appliance and verify that the client appears in 'Accelerator Clients' with open connections and that optimization is occurring