



## Deployment Guide v3.0

### SECTIONS

1. COMPONENT OVERVIEW
2. STANDARD CONFIGURATION OVERVIEW
3. DR CONFIGURATION OVERVIEW
4. ENHANCED REPLICATION
5. OFFSITE VAULTING
6. OTHER CONSIDERATIONS



# Component Overview

This guide is intended to provide you with an understanding of the basic deployment scenarios for Alike Standard and DR Editions. Covered in this document will be the installation options and best practices for Enhanced Backup, Enhanced Replication and Offsite Vaulting jobs. For more in depth information on any particular setting or topic, please refer to the Alike Admin Guide.

As an overview, Alike v3.0 consists of the following three components:

1. Alike Services and Manager (UI)
  2. Alike Data Store (ADS)
  3. Alike Backup Delegate (ABD – only required for XenServer)
1. The Alike services (BackupScheduler and DataVault) perform all of the job scheduling, database maintenance, data validation and other housekeeping tasks. These services may be installed on any Windows machine, either physical or virtual. Because these services do not perform any of the data deduplication or processing, they generally require only modest resources. The system running the Alike services does not need to be high end, as it will not significantly contribute to Alike's backup performance, however, higher speed disks for the Alike databases can help increase the performance of validation times.
  2. The Alike Data Store (ADS) holds all of Alike's critical deduplicated backup data. As such, it will require sufficient storage space to hold all the desired backup versions. The ADS can be located on any CIFS compliant storage. If desired, the ADS can even be shared locally from a path on the Alike Server itself (where the services reside). Often however, it is preferable to locate the ADS directly on a NAS or storage server for performance and fault tolerance purposes. If the Data Store is local to the Alike Server, you may want to refer to our knowledge-base article on tuning CIFS performance for windows.
  3. The Alike Backup Delegate (ABD) and Q-Hybrid are the two types of “processors” in Alike. They both perform the exact same function (process deduplication, compression and encryption), but work very differently under the hood.

## ABDs

The ABD is only available on XenServer platforms, and provides the “enhanced” methods of execution. They are run as tiny virtual appliances, and are provisioned by Alike from a template only when needed, and are disposed of when their work is complete. The ABD template is a tiny (128MB) virtual appliance image that Alike deploys and manages in your XenServer pool. The ABD template must reside on shared storage, which must be accessible by all hosts in a resource pool. For environments with multiple pools, one template image must be deployed to the shared storage of each pool. In single XenServer environments, you may deploy your ABD image to local storage. ABDs are dynamically provisioned by Alike to perform an Enhanced backup, replicate or restore job, and are subsequently removed when no longer needed. Alike will provision one ABD for each job (sometimes two for ER jobs-- one on either end).



## **Q-Hybrid**

Q-Hybrid jobs leverage Alike's tiny, dynamic Q-Hybrid agent, that is transferred to the guest VM at the time of backup or replication, executed, then removed when complete. No additional configuration or installation is required for the guests. The Q-Hybrid agent is therefore not installed ahead of time. Q-Hybrid does require that the Alike server has network connectivity, WMI access and access to the target VM's default administrative share (admin\$).

Q-Hybrid is optional in an XenServer environment, but mandatory for Hyper-V. Q-Hybrid offers a number of benefits over the ABD approach including NTFS optimization, full VSS "quiesced" support, and iSCSI drive protection.



# Standard Configuration Overview

Alike offers high performance backup and replication features using our ABD and Q-Hybrid technologies. Below are the typical deployment options available for Alike.

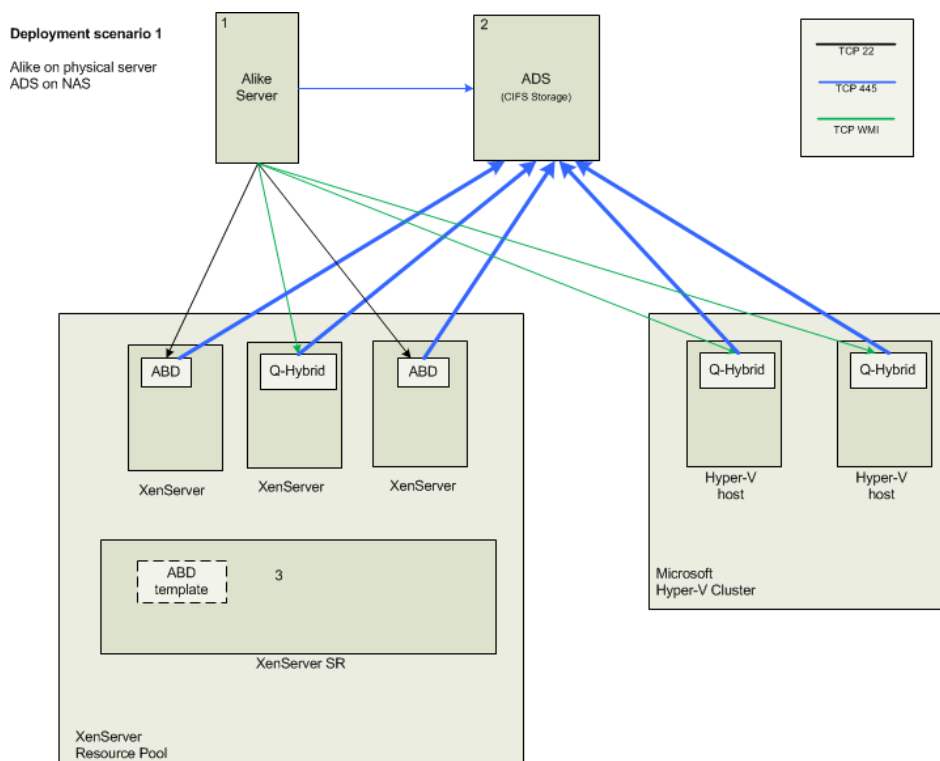
## Configuration I

Alike Services: Physical Host

ADS: NAS CIFS Share

In this configuration, the Alike services are installed on a physical Windows machine, separate from the virtualized environment. Additionally, the ADS is located on a CIFS share residing on a separate NAS device. This configuration provides the highest degree of flexibility and scalability. It also offers several additional benefits.

First, the Alike services are not part of the virtualized environment they are protecting. This offers insulation from virtualization failures, and affords the fastest possible recovery time since the Alike installation would not need to be recovered in the event of a hypervisor/SAN failure.



This also provides scalability benefits, as the ADS is isolated from the Alike services and the two do not effect each other. Additionally, the storage system housing the ADS can be scaled separately to meet the backup demands.



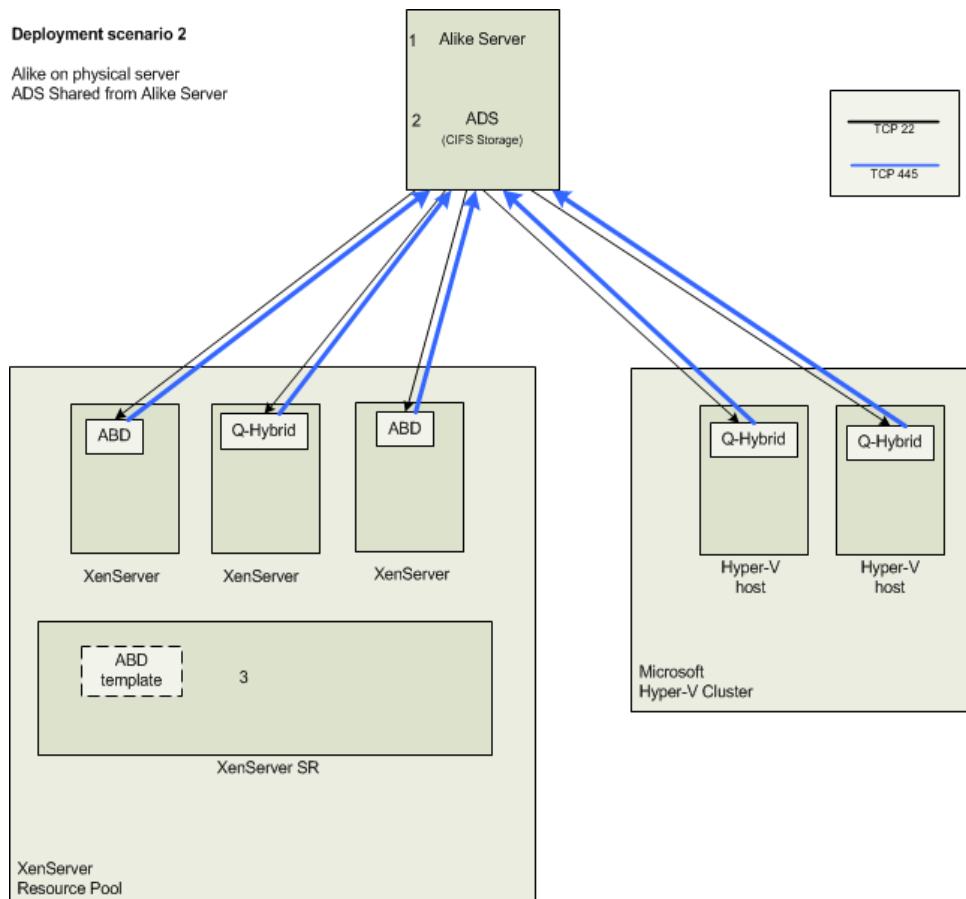
## Configuration II

**Alike Services:** Physical or Virtual Windows computer

**ADS:** Local to Alike Server, Windows Share

For less demanding installations, the Alike services and ADS can be installed on the same machine. This machine can be either physical or virtual. This can be a simple and convenient method as it reduces the number of required systems to single host. Although the scalability might be less than providing a physically separate ADS, this configuration should be sufficient to handle most small to medium sized environments.

This configuration can be ideal for customers looking to only use Enhanced Replication, since ER jobs do not require any additional storage. In this case, Alike can be installed on a very small Windows virtual machine (even a workstation OS) with no significant local storage. For further information on Enhanced Replication requirements and recommendations, please see the section on Enhanced Replication below.



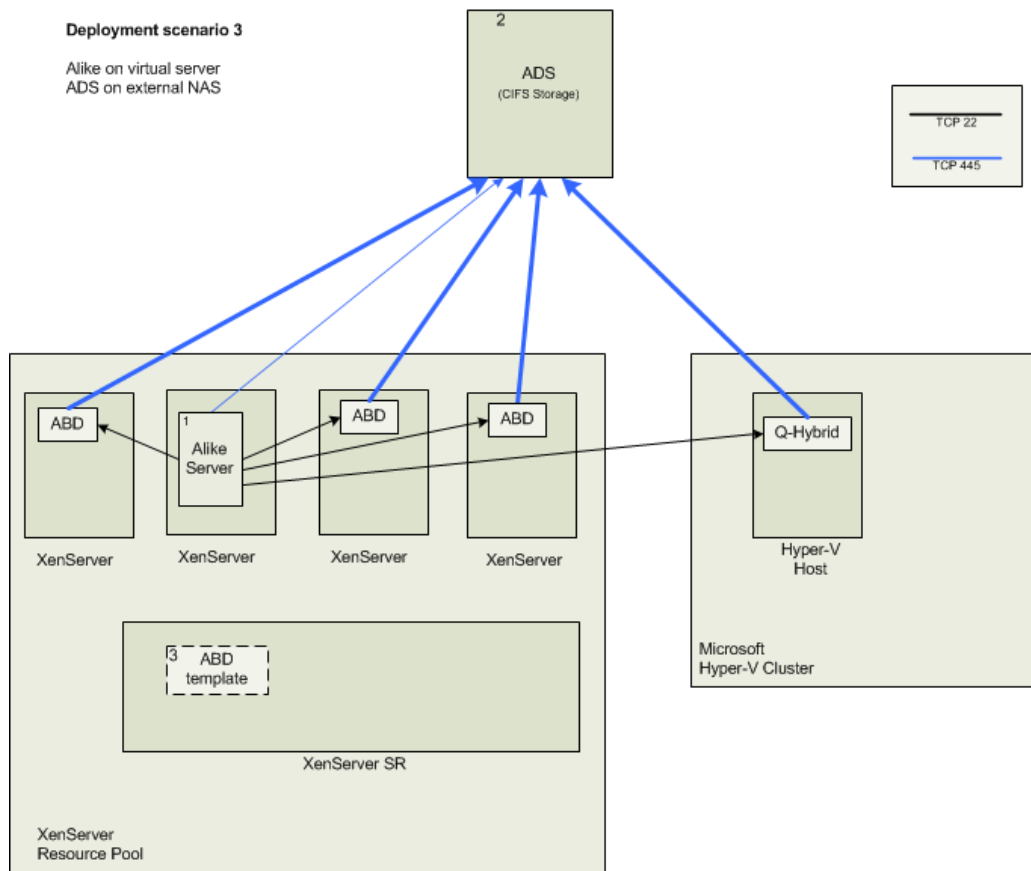
### Configuration III

#### Alike Services: Virtual Machine

#### ADS: NAS CIFS Share

The following configuration provides many of the performance and scalability benefits of configuration 1, and the convenience of running the Alike Services on a virtual machine. Since the most I/O intensive stage of any job is during data processing and transmission, the most critical component for I/O performance is the ADS. By offloading the ADS to a high performance storage system, you can gain a high degree of scalability, even while the Alike server runs as a virtual machine.

The down side of running the Alike services in a virtual machine is that, in the event of a total virtual system failure, Alike would need to be recovered/installed to another system to begin restores. For more information about recovering in such a situation, please refer to the Admin Guide under Recovery Installation.



# DR Configuration Overview

## Enhanced Replication

Included in Alike DR Edition, Enhanced Replication (ER) is a powerful way to provide disaster recovery services for your XenServer or Hyper-V environments. Built on the same ABD and Q-Hybrid technologies that power Enhanced Backups, ER leverages the Alike's deduplication engine to transfer only block level changes to your destination.

While very straightforward to configure and use, there are several options to consider when planning your replication schedule. It will also help to understand the basics of how ER works to optimize your performance and configure for your networking requirements.

Enhanced Replication requires no significant storage to execute, so it can be deployed without the need for a large ADS. ER jobs leverage snapshots on the source host, and send the delta data directly to the target host, via the ABDs or Q-Hybrid agent. No intermediary backup is required. Of course, if your ER jobs are configured with one-pass backup and replication, Alike will need sufficient storage for the backups.

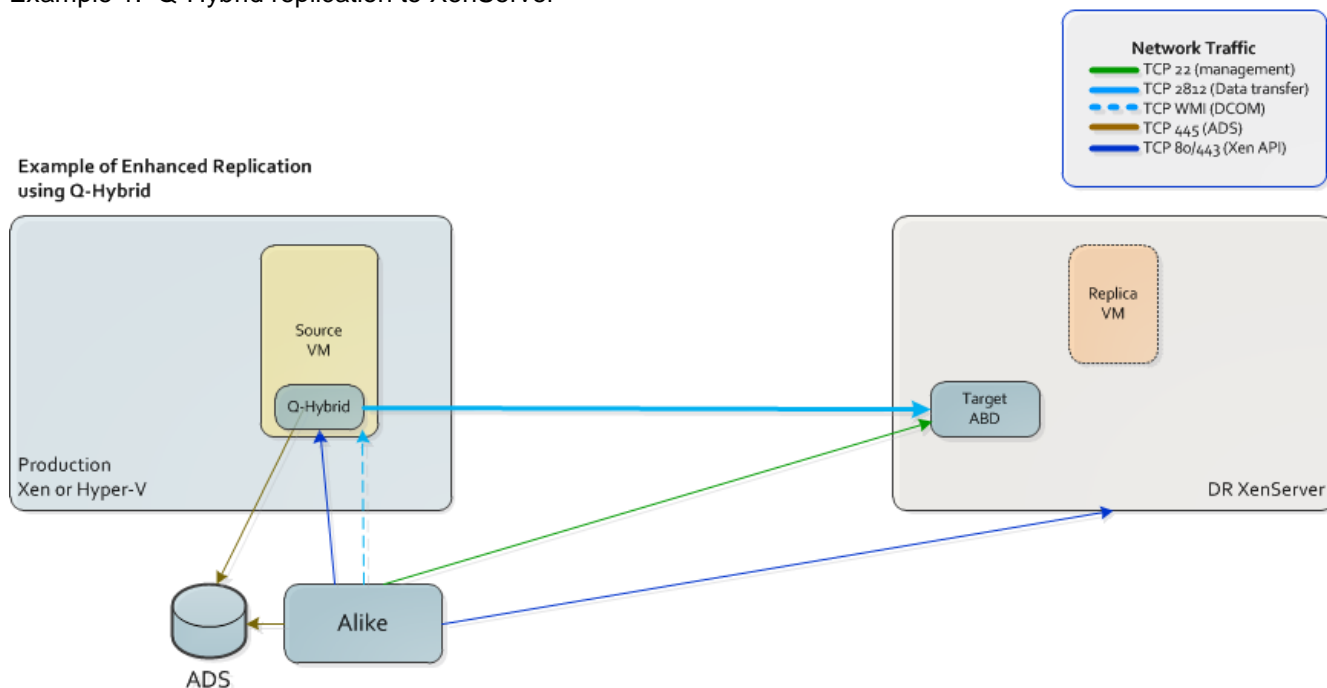
Replication jobs can either use ABDs or Q-Hybrid or both, as their replication method, depending upon your configuration. Both options will be discussed below.



## Q-Hybrid Based Replication

When dealing with Hyper-V hosts as the source or destination for a replication job, Alike will always use Q-Hybrid as the replication engine. Likewise, when replicating to a XenServer host, Alike will always require an ABD for the target side. However, you may choose Q-Hybrid or ABD for the replication engine when both sides are XenServer.

Example 1. Q-Hybrid replication to XenServer

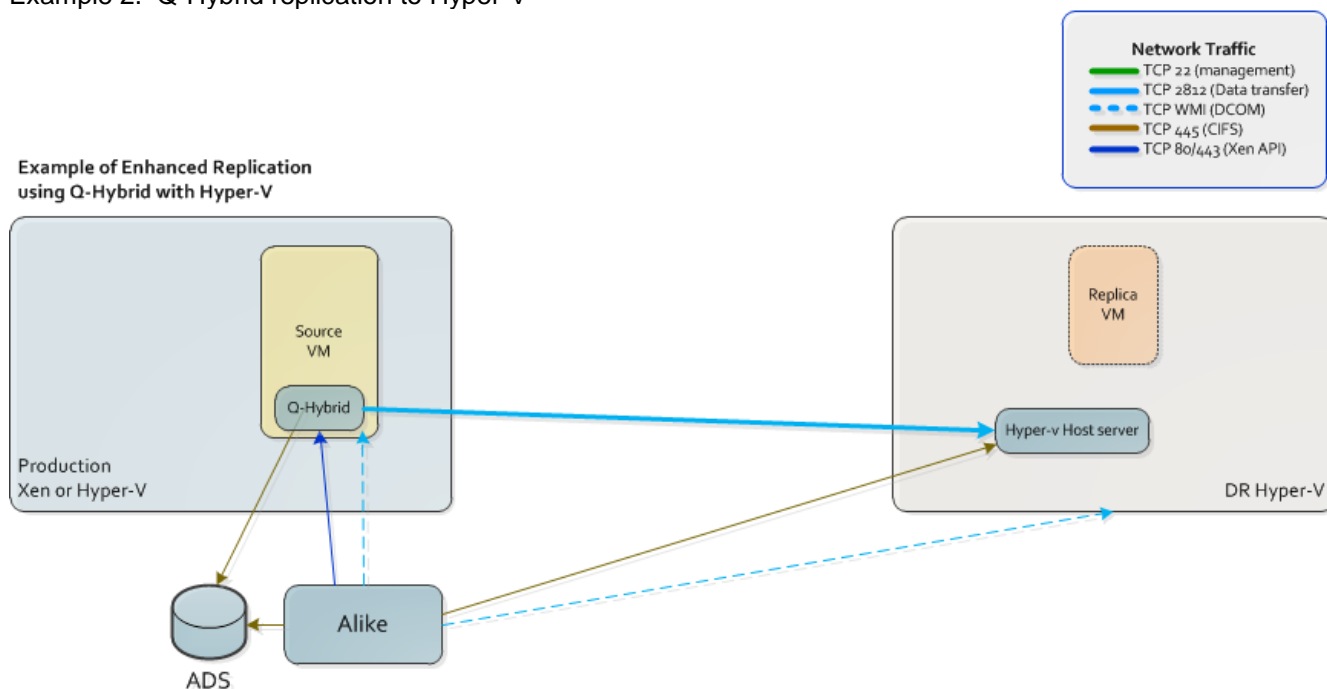


In the above example, Q-Hybrid is being used as the replication method, and the target host is a XenServer. In this configuration, Alike will deploy 1 Q-Hybrid agent to the source VM, and 1 ABD to the target to perform the replication work.





## Example 2. Q-Hybrid replication to Hyper-V



In the second example above, Q-Hybrid is being used as the replication method, but Hyper-V is the target host type. In this scenario, Alike will not utilize any ABDs for the replication job. Instead, it will dynamically deploy the Q-Hybrid agent to the source VM, as well as dynamically deploy the replication components to the Hyper-V host, for the duration of the job.

## ABD Based Replication

ER jobs may use one or two ABDs to perform their work, depending on the location of the VDIs for the VM being replicated. If a single ABD is able to see both the source and destination VDIs, then only 1 ABD will be provisioned. This is the case when replicating within a resource pool, to or from shared storage. However, if a single ABD is not able to see both source and destination disks, a second ABD will be provisioned to handle to destination. This is common when replicating between local SRs, or across resource pools.

## Enhanced Replication Networking Requirements

In a single ABD ER job, the networking requirements are relatively simple. The Alike Services must be able to access the ABD via TCP 22 on the assigned IP address. Once Alike automatically configures the ABD for the job, no significant amounts data will traverse your networks. During the course of the job, the ABD will access the ADS occasionally to record data, so access to the ADS on TCP 445 is required. The actual replication work, and data transfer, will be done entirely between your SRs from within the single ABD.

For ER jobs that require 2 ABDs, additional networking considerations must be made. This includes all ER jobs that traverse a WAN or VPN. Please note, these considerations and requirements change depending on the direction of the selected data 'flow' (push vs. pull). For the purpose of illustration,



the ABD deployed to the source VM will be called the 'source ABD', and the ABD deployed to the destination VM will be the 'target ABD'.

### Push replication flow

By default, ER uses the 'push' data flow. In this case, the source ABD should be 'closer' to the ADS, as it will be used to communicate with the ADS. The target ABD does not need to communicate with the ADS, and will not initiate any networking communication. Therefore, the target ABD can be deployed in a remote location to support the ER job. During the course of the ER job, the source ABD will connect to the target on either its assigned IP, or alternate IP to perform the data transfers as needed.

ER required ports:

Alike Server → Source ABD on TCP 22

Alike Server → Target ABD on TCP 22

Source ABD → Target ABD on TCP 22 & 2812 (default)

Source ABD → ADS on TCP 445

See figure 1 below for a typical “push” replication example network flow. Note the thicker lines indicate heavier network traffic, and thin lines indicate light or management network traffic. In this configuration, the source ABD is local to the ADS, so all communications have a very low latency, allowing Alike to perform at its peak.




Figure 1. Ideal “push” replication configuration



Figure 2 demonstrates an inefficient “push” replication configuration. In this case, the Alike installation and ADS reside on the far end of the WAN (in this case close to the target VM). Note that the source ABD must connect over a (presumably) slow WAN connection to communicate with the ADS, which can substantially slow the replication job as latency rises. In this situation a “pull” replication flow would be significantly more efficient.




Figure 2. Inefficient “push” replication configuration



## Pull replication flow (ABD only)

When the data flow is set to 'pull', Alike will essentially reverse the roles of the ABDs. In a pull job the target ABD should be 'closer' to the ADS, and will require networking connectivity to the ADS. Additionally the target ABD will initiate connectivity to the source ABD on its assigned or alternate IP address. In this configuration, Alike allows you to 'pull' VMs from a remote location to the network running Alike. This option is very useful for DR service offerings.

ER required ports:

Alike Server → Source ABD on TCP 22

Alike Server → Target ABD on TCP 22

Target ABD → Source ABD on TCP 22 & 2812 (default)

Target ABD → ADS on TCP 445

Figure 3 below demonstrates a properly configured “pull” replication flow. In this situation, the target ABD is near to the ADS, and therefore has a low latency. This configuration is ideal for situations where Alike is installed and running in the DR datacenter, commonly seen with MSP (managed service provider) solutions. In this scenario, no Alike software needs to be installed in the production environment (except for the ABD template).




Figure 3. Ideal “pull” replication configuration.



Figure 4 shows an incorrect “pull” replication flow configuration. This configuration is less than ideal, because the Alike installation and ADS are remote to the target ABD, which must communicate with it frequently. Latency over a WAN connection in this configuration will impact the replication job's performance.




Figure 4. Inefficient “pull” replication configuration.

For full details about ER job settings and configuration, please refer to the Alike Admin Guide.



# Offsite Vaulting

Offsite Vaulting (OSV) is offered with Alike DR as a means to efficiently transfer some or all of your backup data to a secondary location. When necessary due to disaster, audits or otherwise, you can install a temporary Offsite Console, or attach another existing Alike installation to your Offsite DataStore, and restore directly with either file level recovery, or full VM restores. Offsite Vaulting fully leverages Alike's global block level deduplication, and only transfers post deduplicated, compressed data over the network, minimizing bandwidth requirements. During normal operation, only a CIFS compliant network share is required in the offsite location.

To configure OSV, you must simply configure your main installation of Alike as required, then configure the Offsite CIFS information in the Offsite tab of the settings screen. The 'onsite' installation of Alike must be able to access your 'offsite DataStore' CIFS share via TCP port 445. For instructions on basic installation and configuration of Alike, please refer to the Alike DR Getting Started Guide.

The diagram below provides a simple installation example of Alike with Offsite vaulting configured.




Figure 5. Example Offsite Vaulting configuration.

Please note that OSV does not require that the secondary location have XenServer in place during normal, backup operations. It is perfectly viable to vault to a standalone/headless CIFS share and only provision a Xen host when necessary for recovery.

As can be seen in figure 5, the networking requirements for OSV are very simple. The Onsite installation of Alike will attempt to connect to the Offsite DataStore on the default CIFS port (TCP 445).



## Other Considerations

### ***Concurrency and scalability with ABDs***

Alike dynamically provisions ABDs for each Enhanced job from the template uploaded to your shared storage. This means that each ABD provisioned does not actually consume any additional storage in your SR to exist, as they all point back to the original base image.

Each ABD can simultaneously process all of the drives of a single VM at once. In order to concurrently backup multiple VMs, you must create separate jobs. For example, if you had 8 VMs to process, you could split them into 2 jobs with 4 VMs in each, and Alike would provision 2 ABDs, and process 2 VMs at a time. Or, similarly, if you configured 4 jobs with 2 VMs each, Alike would provision 4 ABDs and process 4 VMs simultaneously.

Please be aware that each ABD requires 1 IP address to operate. You can enter as many IP addresses as you like in the Alike Manager->Settings->Networking tab. You must have enough IP addresses configured to support the number of concurrent jobs you wish to run. Additionally, the IP addresses assigned to ABDs are grouped by Xen Pool. So if you have multiple pools that Alike is protecting, you must enter IP addresses in for each pool.

Also be advised that the IP addresses assigned to the ABDs must be able to access your ADS CIFS (TCP 445) share directly. The IP address/virtual network that you assign to your ABDs determine which network the backup data will traverse. You may segment your backup traffic by selecting a specific VLAN for your ABDs.

