# Asynchronous Replication

## Introduction

A disaster recovery site is crucial for companies which care about their data and it's A MUST for companies which are under regulatory requirements. On-site Redundancy (RAID, Synchronous Mirroring, etc.) does not guarantee data protection from hurricanes, floods and earthquakes that could destroy a whole primary datacenter.

## Problem

**High cost, performance impact** and **compromised data integrity** are typical issues for a conventional Disaster Recovery (DR) site implementation.

A traditional replication solution presents the following **cost-related problems**:

- Typical DR replication solutions require similar or even identical hardware set on the both primary and remote sites. That significantly increases both capital and operational expenditure because there is twice as much hardware to purchase.

- WAN throughput requirements increase if the data reduction technologies are not used on the primary site, since all the data changes made on the primary storage will be sent to the DR site. The demands for the WAN data link rise, resulting in the requirement for very high performance and expensive WAN connections.

Enabled Disaster Recovery **replication impacts performance** (network, CPU, hard disks, etc.) for the following reasons:

- Storage Array routine tasks and the recently enabled DR replication compete for system resources reducing performance. Running DR replication over the critical path may be a significant bottleneck in the performance and reliability of the entire system.

- Replicated data usually gets rebuffered on the primary site and this impacts the storage sub-system, causing a decline in performance and over-provisioning of disk space.

**Compromised data integrity.** If the data on the primary site gets corrupted for any reason (malware, human factor, etc.), then it will be corrupted on the DR site, since errors are also replicated. There would be no possibility of rolling back to the original state with consistent data since the data on both sites would be damaged.



*Typical DR solution*

## Solution

Minimal **StarWind Virtual SAN** DR setup is just one virtual machine on the DR site since it is a pure software solution, which is **extremely cost-effective**. It works on any hypervisor that runs on existing general-purpose hardware and in private or public Cloud (e.g. Amazon or Azure). This solution will not affect performance of the whole system since the DR site does not require high performance.

StarWind implements replication as asynchronous, putting it to the background and using snapshots as the source. As a result, **StarWind eliminates the performance and space over-provisioning problems** typically encountered with traditional DR replication solutions.

**Overload of the WAN data link is eliminated** by combining StarWind space reduction technologies — deduplication, compression, snapshots, change block tracking, etc. As a result there is a significant decrease in the amount of actual data transferred. Due to the reduction in WAN requirements there is no need to buy high-cost data links.

**Data Integrity.** StarWind snapshots technology restores to the latest integral point of data even if corruption occurred on the primary and DR sites.
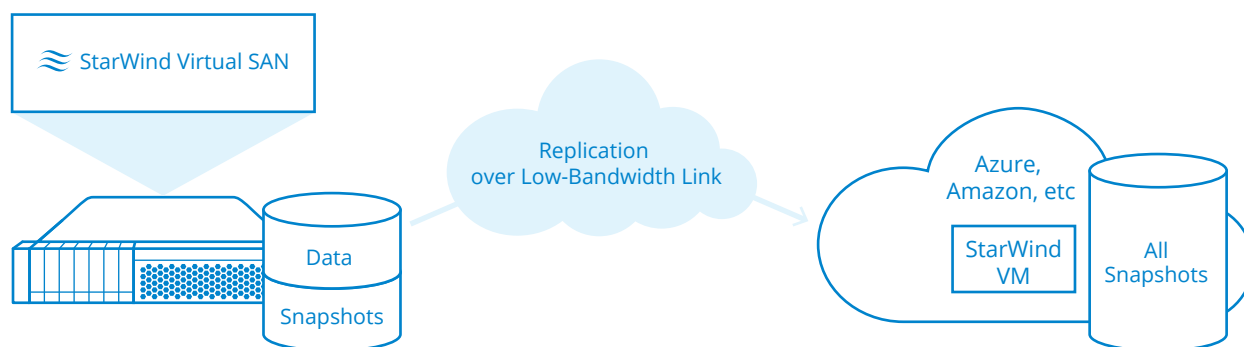


*StarWind Asynchronous Replication*

## Conclusion

The combination of StarWind data-compression and replication technologies eliminate the issues of high cost, performance impacting and data integrity  typical with conventional DR implementations. Allowing installation to take place in the cloud VM for DR purposes and moving the replication and data compression to the background, it prevents the common pitfalls found in traditional DR solutions.

In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms".

To learn more, visit **www.starwind.com**