

White Paper

# Converged

## Introduction

Contemporary datacenter infrastructure has adopted virtualization as a staple. To keep up with the growing market demands, storage systems, being integral to virtualization's proper work, evolve into hyperconverged infrastructure (HCI). Nonetheless, hyperconverged and converged configurations both exist and coexist still.


Hyperconverged doesn't always mean better. It's just that either of the scenarios tends to focus on slightly different tasks. Converged Infrastructure may be the way to go if you will require growing storage and compute resources independently of each other.

## Problem

Most often than not, businesses will need to increase their storage and compute resources separate from each other. The hyperconverged infrastructure route won't necessarily allow introducing such changes easily. Notably, higher compute and storage hardware parameters will have to be introduced simultaneously when scaling a hyperconverged cluster. Indeed, virtualized environments may expand in various dimensions:

- Storage resources (IOPS, capacity)
- Compute resources (RAM, CPU, etc.)

However, in a hyperconverged infrastructure scenario, every virtualization host is running both storage and compute utilizing the same hardware at the same time. Therefore, if you decide to increase one of the dimensions, you will have to increase the other one, too. This boils down to the fact that you may not actually have the need to increase the other dimension at all. As a result, you incur costs by buying additional hardware that wasn't essential to the initial problem that you had.




It's wasteful to upgrade compute and storage simultaneously every time.

## Solution

StarWind has been designed to perfectly fit either of the scenarios or any mix-up that you may have. Its features allow scaling storage and compute independently, with various levers regardless of each other, while still running all the systems without any downtime. This means that, with StarWind, you can add precisely what you need instead of buying hardware to sit there and wait for its future use. Such a flexible system can be easily adjusted to fit a specific task better.

When using StarWind in a converged scenario where clusters run on compute power mostly, you just need to add more RAM to the hypervisor server or add a new node, letting the performance and storage resources remain the same. Similarly, if your infrastructure is running out of storage IOPS or TBs, you can upgrade that particular dimension by adding a new storage node, without manipulating compute resources. The power of choice is yours.



Cost-efficiency lies in the freedom to scale your infrastructure how you see fit.

## Conclusion

You can choose to run StarWind as ready-to-go nodes (StarWind Storage Appliance) injected into an existing IT environment, or as Software-Defined Storage (StarWind VSAN) on top of existing hardware. The ability to run StarWind independently at the storage layer without involving compute provides for significantly reduced CapEx and OpEx thanks to acquired freedom of adding only what you find essential into your storage infrastructure. This way, you can tailor your system hardware to fit specific needs and tasks better and have a more granular approach to your business' IT.

## Contacts



### Phone number

+1 617 449 77 17  
+44 2037 691 857  
+34 629 03 07 17



### Fax

+1 617 507 58 45



### Voice Mail

+1 866 790 26 46



### Customer Support Portal

[www.starwind.com/support](http://www.starwind.com/support)

### Support Forum

[www.starwind.com/forums](http://www.starwind.com/forums)

### Sales

[sales@starwind.com](mailto:sales@starwind.com)

### General Information

[info@starwind.com](mailto:info@starwind.com)

To learn more, visit [www.starwind.com](http://www.starwind.com)