

SMART HYPERCONVERGENCE

INNOVATION IN THE DATA CENTER COMES IN WAVES

Innovation in the data center comes in waves. The client-server wave drove IT architectures through the 1990s into the early 2000s. The virtualization wave came next. Now hyperconvergence is gathering momentum and the implications are profound.

Modern applications have tested the limits of existing virtualized infrastructures. In response to

traditional infrastructure's inability to meet their needs, companies like Amazon, Google and Facebook evolved an approach to datacenter infrastructure that leveraged software as the mode for deploying functionality, including data and storage management and control. This approach is called hyperconvergence.

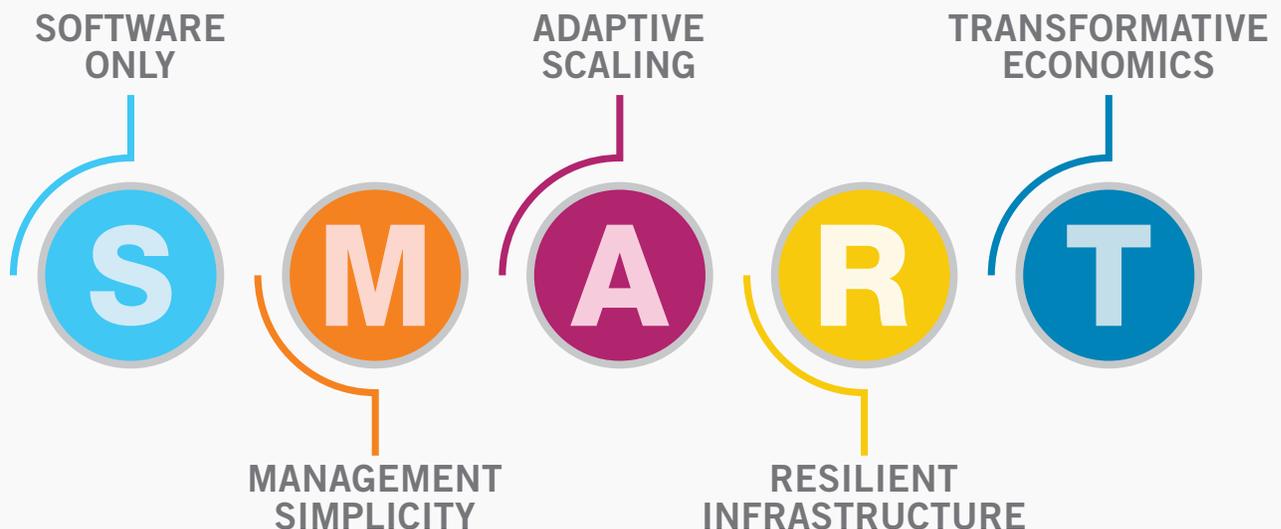
Springpath Data Platform delivers on the promise of this new approach by turning industry-standard servers and storage into dynamic pools of resources

that scale elastically and drive hyper-efficient utilization with self-monitoring and healing capabilities exposed through familiar and intuitive management interfaces. Springpath also eliminates storage silos that are based on proprietary hardware storage arrays or converged appliances.

This pure software solution approach to hyperconvergence enables customers to continue using standard servers of choice while delivering the most agile and lean IT hyperconverged infrastructure available today.



Five Key Ingredients of Successful Hyperconverged Solutions



SOFTWARE ONLY SOLUTION RUNNING ON X86 SERVERS OF CHOICE

In advanced hyperconverged solutions, customers are able to use their preferred server vendors and direct attached storage for all infrastructure needs. Hyperconvergence software is intelligent enough to allocate and utilize resources dynamically, based on application needs. Administrators simply provision workloads from the compute and storage pool spanning the full range of available processors, flash and spinning disk drives and the software makes it happen.

Springpath Data Platform is based on a distributed file system built from the ground up to deliver high performance, reduce operational complexity and simplify data management. Springpath eliminates the need for network storage, so there are no more volumes to provision or LUNs to configure. The Data Platform automatically utilizes all compute and storage resources available to the cluster to ensure optimal performance and resource utilization.

MANAGEMENT SIMPLICITY, NO LEARNING CURVES

Hyperconverged environments put standard, low cost hardware under simplified, efficient and robust, virtual-aware software management. Both IT staff resources and hardware resources are optimized favorably, impacting both CapEx and OpEx. This simple, efficient design enables IT to manage aggregated resources across individual nodes as a single converged system.

Springpath seamlessly extends vCenter and empowers VM administrators to manage their storage and data without having to learn yet another management tool. VMware snapshots and cloning capabilities are replaced by more scalable and performant Springpath native snapshots and clones. These familiar, integrated capabilities make a VM admin's life simpler, more productive and allow fewer admins to manage more VMs.

Clones and snapshot are taken at a VM level and hundreds of clones can be created and deleted in a matter of minutes. Whether admins require

multiple snapshots for backup or aim to stand up new environments for Test/Dev or VDI, all operations can be done in a few easy steps using the vSphere plug-in. Simple, efficient, familiar and easy.

ADAPTIVE SCALING OF COMPUTE, CACHING OR CAPACITY

Since software-defined hyperconvergence provides a unified resource pool across a range of capabilities, it is now possible to scale compute, caching or capacity resources just-in-time and in just the right increments. The ability to tune application performance by dialing up cores, cache or hard disk capacity is an admin's dream, and puts cost and SLA control in the hands of business users and IT staff.

Springpath allows IT staff to add servers without downtime and resources are made available immediately and automatically. Simply add resources to the cluster and the Springpath Data Platform automatically starts utilizing and rebalancing the available resources. You can add complete servers or just one tier of resources, depending on the needs in your environment.

At a more granular level, the Springpath Data Platform effectively allows IT staff control over scaling IO performance or storage capacity independently. Compute only nodes (e.g blades) can be added to a cluster at any time. When better storage IO performance is needed, caching capacity can be added; or, when data growth results in the need for more capacity, simply insert more hard drives to existing nodes.

RESILIENT INFRASTRUCTURE THAT IS SELF-HEALING

Modern software-defined hyperconverged solutions should be able to sustain data access in the face of inevitable hardware failures and provide proactive alerts so administrators can take timely action.

Springpath Data Platform is an enterprise-grade solution. It provides high availability by mirroring incoming

data, synchronously, to one or two other disks located in different servers. When a disk fails, a rebalancing engine uses the remaining copies to rebuild and redistribute data without incurring down time. In addition, IT administrators are notified about the failed components to ensure rapid resolution.

The platform proactively collects and sends health and usage metrics to Springpath in order to discover and address issues before they become critical. IT admins can log in at any time to access performance analytics and explore historical trends and insights.

TRANSFORMATIVE ECONOMICS AND UTILIZATION

Given the hyper-rapid pace of data growth, hyperconverged systems must be able to effectively utilize and manage consumption in order to keep costs under control. One of the primary benefits of a software-defined hyperconverged datacenter is the transformative economic savings that organizations can now achieve.

With Springpath's annual subscription based licensing, customers can easily start with a low-cost, low-risk deployment. As the environment grows, Springpath maximizes infrastructure resources with a robust set of data management, reduction and efficiency services, driving up to 5X data reduction. With a single resource pool, there are no more islands that can't interoperate and thus end up underutilized.

Space-efficient, pointer-based snapshots and clones are delivered simply and effectively without bogging down application performance. At the same time, inline de-duplication is leveraged across all media and inline compression is always on to achieve optimal utilization levels. The outcome is dramatically lower Total Cost of Ownership (TCO) compared to traditional infrastructure.

