



WHITE PAPER

Converged Infrastructure vs. Hyper Converged Infrastructure

Simplifying Application Placement Across Public, Private, and Hybrid Multi-Clouds

Introduction

The title of this paper might suggest having to pit Converged (CI) against Hyper Converged (HCI), but quite the contrary. In the next few pages we will show you that your application and user needs can thrive with both CI and HCI, as long as they can communicate and share across a data fabric.

Overview

Before the cloud arrived, diverse enterprise applications existed for decades—housed in IT data centers. As mainframe computers gradually gave way to open systems servers based on UNIX/Linux and Microsoft operating systems, application placement typically consisted of mission-critical applications running on robust clustered high availability (HA) systems. Less critical applications were commonly found on standalone servers.

As servers, and the applications they supported, evolved into essential ingredients of business, proper placement of applications became more sophisticated. Fast CPUs with ample of DRAM and quick-spinning HDDs were favored over complex clustered HA systems within critical applications such as OLTP, MRP, and ERP databases. To ensure that these applications always remained available, synchronous mirroring and remote disaster recovery sites became the norm. This practice guaranteed recovery time objectives (RTOs) within seconds and recovery point objectives (RPOs) within a handful of database transactions.

Other applications became important to businesses, but not necessarily critical to operations. Applications such as email and web services fell into this category. To reduce budgets, IT organizations often deployed “server farms” with row upon row of low-cost servers to support these applications. Less-expensive, slower-spinning HDDs with high capacities were paired with these servers to further reduce costs.

All other applications and data became relegated to commodity equipment—that is, the least expensive servers and tape-based storage systems available. Legacy applications, backups, and various archival copies kept for legal compliance reasons fell into this category.

Despite constant changes in technologies, companies continued the same approach. They placed critical applications on the more reliable and faster equipment and placed noncritical applications on progressively slower, more cost efficient equipment.

The Need to Rethink Application Placement

As we entered the 2000s, several changes caused a major shift in enterprise strategies. First, the entire concept of application placement came into question. Were email applications mission critical or business critical? Likewise, what was the business impact of web servers going down for a few days, or even a few hours? As society came to depend more on digital information, and businesses relied more on IT services, the entire world of

application placement was turned on its head. It seemed that every application was becoming mission critical!

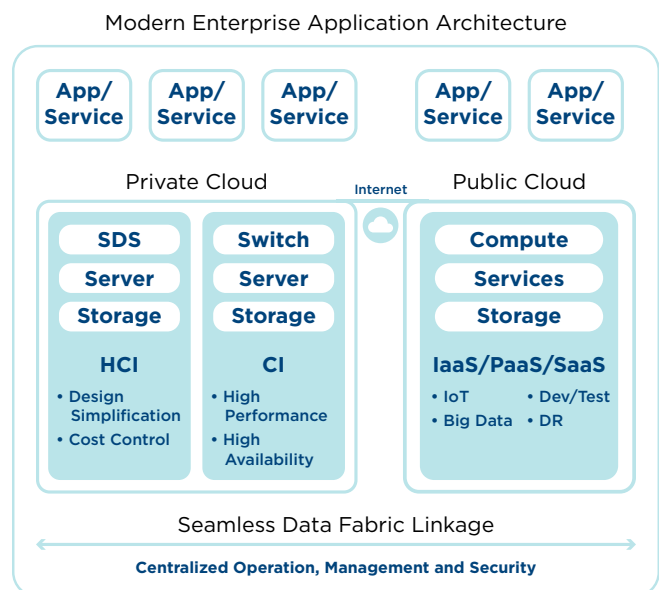
Second, new technologies appeared on the horizon—namely, virtual machines (VMs), cloud computing, and flash storage. Exciting new technologies, yes—but often these technologies were ahead of the curve when it came to application placement, and IT suffered some growing pains as a result.

A mission-critical OLTP application, for example, was blazingly fast running on an all-solid-state-drive (SSD) storage array, but enterprise reliability and disaster recovery mechanisms were far from proven. Or, a critical business application could be quickly developed and rolled out in the cloud using web-based tools. But when it came time to put the application into production, IT organizations often found it difficult to provide the bandwidth and security required by the application.

The result was a siloed approach to application placement. Mission-critical applications ran on dedicated flash-based storage systems, web-developed applications remained hosted by public cloud service providers, and legacy applications remained on existing UNIX and Windows servers, with data served by traditional storage arrays. As IT would find, this configuration was not optimal and led to higher costs and management headaches. A better solution was needed. In the remainder of this paper, we'll discuss the ingredients of a modern application placement environment: one that takes advantage of the best in today's technologies.

Modern Application Placement

As shown in the following diagram, the choices available to IT architects today are quite different than they were just a few years ago. Instead of monolithic, application-centric servers, IT prefers the resource-sharing capabilities of public and private clouds. Driven by the development of VMs, these clouds allocate compute, memory, and storage resources as needed.



In this way, resources can be consolidated and operated at much higher utilization rates, allowing IT to operate with much more agility and efficiency.

These cloud-based IT architectures require a different approach to application placement. Unlike the days of everything on premises, this model requires that applications and data be able to seamlessly move from platform to platform and cloud to cloud as application configurations shift with dynamic business priorities.

Private Cloud Application Strategies

As the number of applications managed in an enterprise has grown, there has been a strong trend to standardize them in VMs and containers. The aim has been to simplify management and make applications compatible with the cloud.

However, applications with special infrastructure needs often require different platforms than standardized applications to accommodate higher performance, extremely low latency, high availability, or other specific infrastructure needs.

In private clouds, two platforms have emerged to address this need: high-performance/high-availability converged infrastructure (CI) systems, and simplified, scalable hyper converged infrastructure (HCI) systems. These systems can be thought of as the old mission-critical/business-critical application platforms with one important difference: both must operate in a dynamic environment.

Because of the diversity of application requirements, both platforms are necessary—but application data must be able to flow smoothly between platforms as priorities shift. This flow is the nature of IT today, and the application environment must adjust accordingly.

Public Cloud Application Strategies

Public clouds, once considered an experiment for enterprise organizations attempting to reduce IT costs, have become a key ingredient in transforming businesses and gaining a competitive advantage. As Gartner states in its article “Is Your IT Team Prepared for Public Cloud?”:

Public cloud services are an inextricable part of how enterprises are responding to the digital economy. They are adopting public cloud services to gain agility, enter a new market or try a new idea without too much capital investment. They may also adopt these services for the opportunity to create a business model based on application programming interfaces (APIs).

Gartner further states that spending on public cloud services is growing more than five times faster than overall IT spending. However, are applications “born” in a public cloud destined to stay in that cloud? With a modern application placement strategy, the answer is: No! As performance and security requirements change, public cloud applications should be able to shift seamlessly from off premises to on premises as needed.

Any modern IT environment must recognize and accommodate this need.

The Data Fabric Effect on Application Placement

Organizations must adjust to the realities of modern IT and properly support application placement. They need a method for controlling and managing data between on-premises private clouds and the many endpoints in the public cloud.

Fundamentally, a data fabric is a way to manage data, both on premises and in the cloud, using a common structure and architecture. A data fabric provides efficient data transport, software-defined management, and a consistent data format, allowing data to move more easily among clouds. This approach is what allows hybrid clouds to operate in enterprise application environments.

With data portability enabled by a connected data fabric, application data becomes fluid. This fluidity leads to better utilization of resources. Mature applications, for example, often take up data center space, power, and the resources of a skilled IT staff. A data fabric enables you to selectively move applications from your private cloud to a public cloud so that internal IT resources and mindshare can focus on innovation.

Summary

A data fabric makes modern application placement work. It seamlessly connects different data management environments across disparate clouds, making them cohesive, integrated, and whole. The NetApp® Data Fabric allows IT organizations to embrace private and public clouds on their terms. Simply stated, the Data Fabric enabled by NetApp allows businesses to accelerate capabilities while also reducing costs.

Having a fabric means that data is free to move dynamically across all private and public cloud resources. Businesses realize greater efficiencies by pairing workload requirements with cloud economic models in real time, and without disruption. NetApp’s approach to application placement facilitates the expert connection of cloud resources, with a highly efficient transport between systems and clouds and a single purview of data management over a seamless data fabric.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

WP-7286-1118