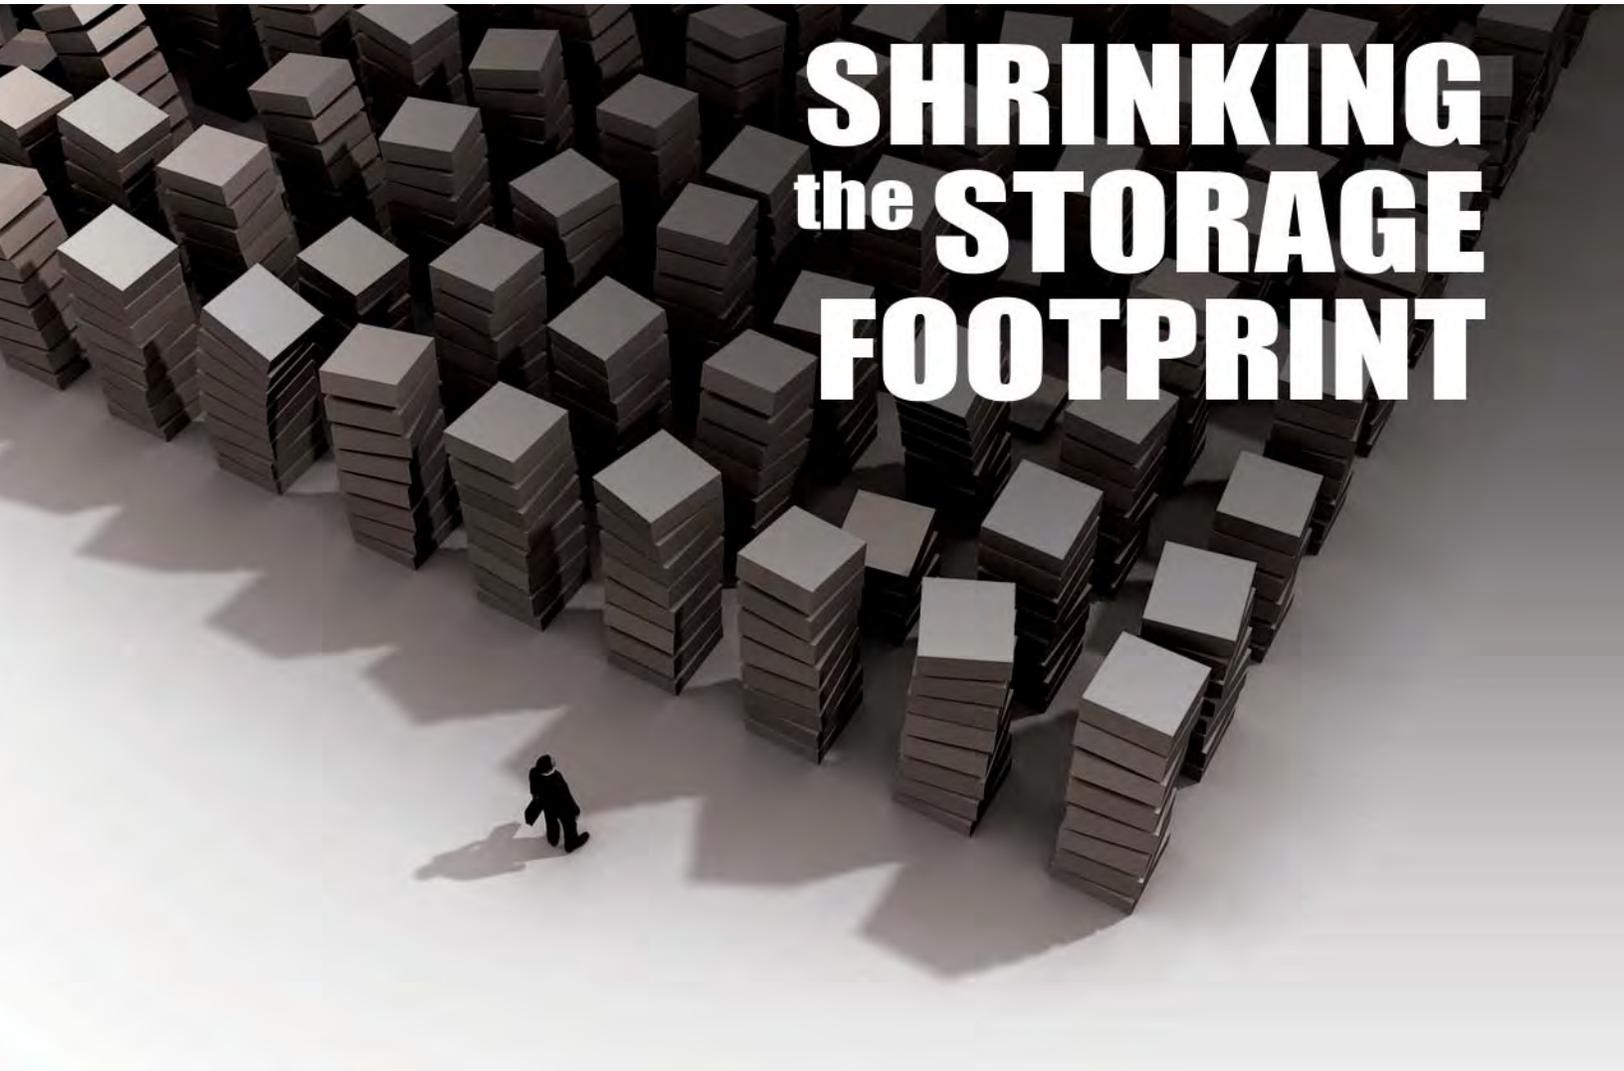


SigmaUptime

volume 15 number 4



SHRINKING the STORAGE FOOTPRINT

UPTIME

**Storage consolidation
with high-density
all-flash arrays
from HPE delivers
multitude of benefits.**

PRSR T STD
U.S. POSTAGE
PAID
Tulsa, OK
Permit No. 2146

Technology services to help
you streamline operations,
reduce costs and improve
business processes.

Companies today must align IT strategy with their corporate objectives, strategy and business model. Pivot has created a portfolio of operating companies and partners with a focus on helping you enhance and extend the capabilities of your technology assets.

Pivot provides technology services ranging from initial needs assessment and design, through procurement and implementation, to on-going support. As an adjunct to your IT team, we provide the resources that allow your team to offload some of the day-to-day operational challenges and focus on innovations that will drive business value and competitive advantage.

Contact us to learn more.



Contents

4 **Shrinking the Storage Footprint**

The growing storage infrastructure is creating a burden as it continually consumes more data center real estate. Hewlett Packard Enterprise is supporting new levels of data center consolidation with recent updates to its 3PAR StoreServ portfolio that take advantage of breakthroughs in 3D NAND flash technology.

8 **Divide and Conquer**

Microservice architecture is changing the way software is built. The idea is to break up rigid, monolithic applications into smaller, more manageable pieces that can be released rapidly for increased agility and scalability.

10 **Best of Both Clouds**

Hybrid clouds use a combination of public and private cloud services to meet the requirements of various workloads. The appeal of the hybrid cloud lies in its ability to provide cost savings and operational efficiency by allowing organizations to choose the optimal platform for each application and service.



4

Sigma Uptime

Copyright © 2016 CMS Special Interest Publications. All rights reserved.

Editorial Correspondence:

10221 East 61st Street, Tulsa, OK 74133

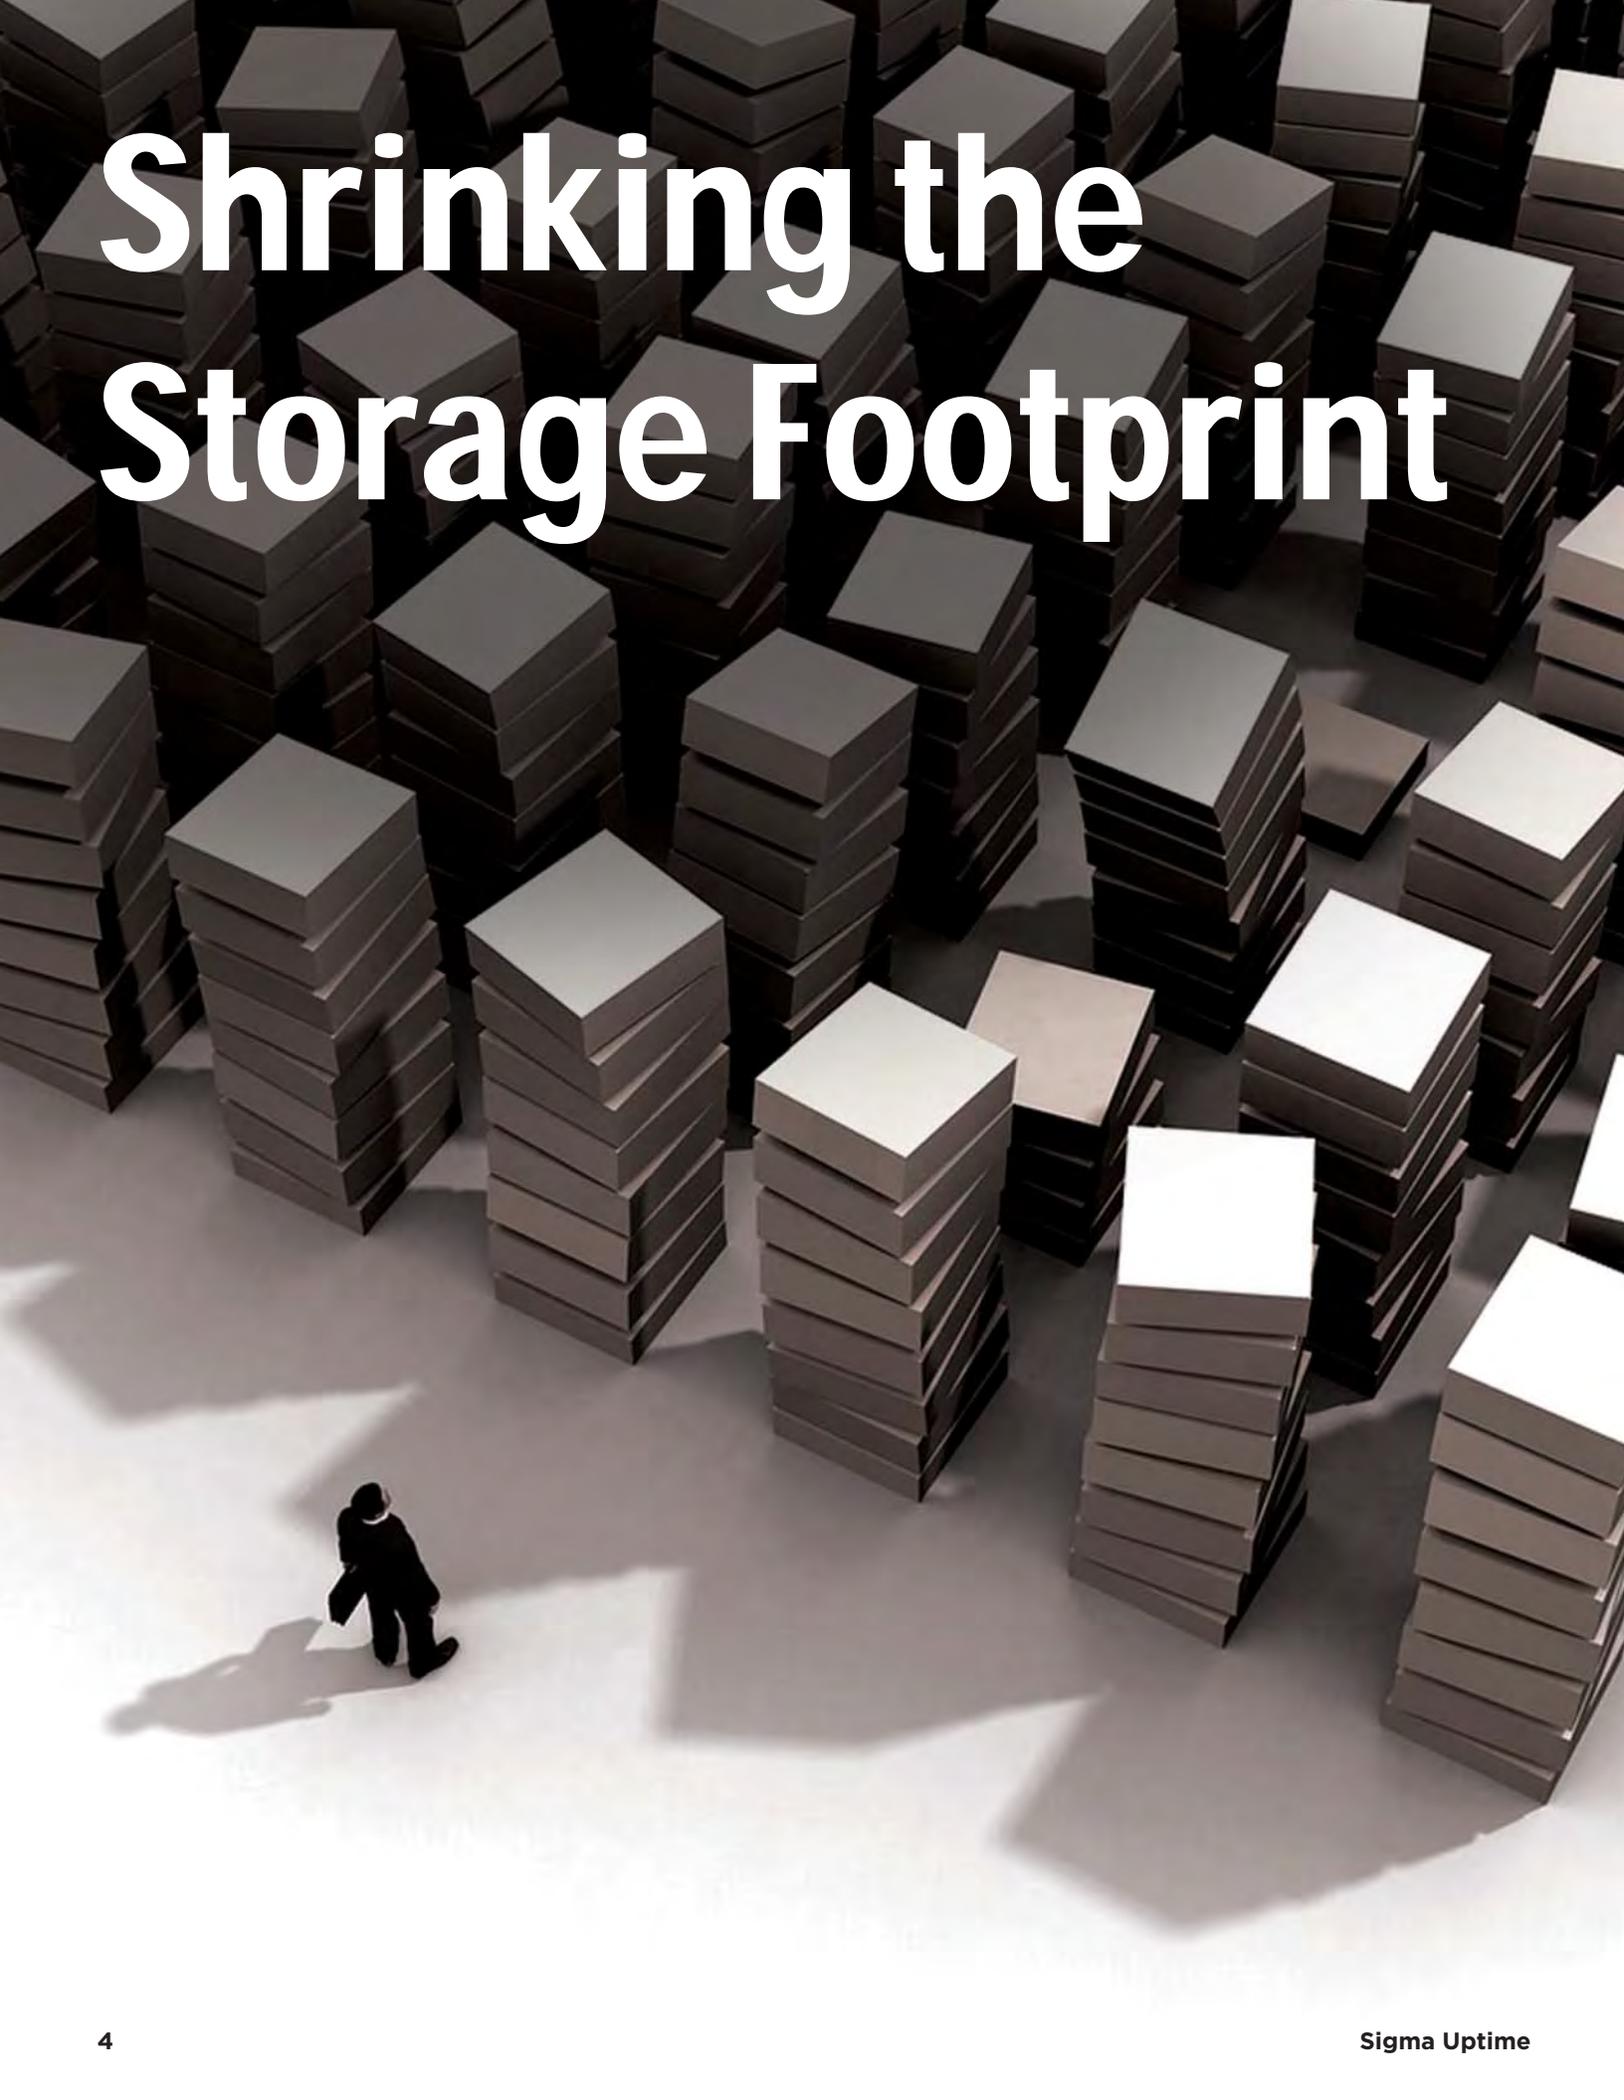
Phone (800) 726-7667 • Fax (918) 270-7134

Change of Address: Send corrected address label to the above address.

Some parts of this publication may be reprinted or reproduced in nonprofit or internal-use publications with advance written permission.

Sigma UPTIME is published bimonthly by CMS Special Interest Publications. Printed in the U.S.A. Product names may be trademarks of their respective companies.

Shrinking the Storage Footprint





Storage consolidation with high-density all-flash arrays from HPE delivers multitude of benefits.

From its utilitarian beginnings as a technology for recording transactions, data storage has evolved to a level of utmost significance and sophistication in the modern data center. In the age of big data, organizations are capturing ever-larger amounts of data from a variety of sources in order to gain insight, drive strategy and make real-time decisions.

However, this transformation comes with a cost. The growing storage infrastructure is creating a burden as it continually consumes more data center real estate. Increasingly, organizations are looking to reclaim that space, cut power costs and drive efficiency through the use of flash storage.

The benefits of flash-based solid-state drives (SSDs) are well understood by now. By eliminating the mechanical chokepoints of hard-disk drives (HDDs), flash storage dramatically improves data access times, improves application performance and streamlines storage management. But the benefits don't stop there.

“Early flash adopters were primarily looking to improve application performance, but now we see that there are benefits that extend all across the data center,” said Allan Doehler, Business Development Manager, Sigma Solutions. “Organizations are finding that flash enables them to consolidate their storage environment onto a more efficient, scalable platform that ultimately delivers significant cost savings.”

3D Density

Hewlett Packard Enterprise (HPE) is supporting new levels of data center consolidation with recent updates to its 3PAR StoreServ portfolio that take advantage of breakthroughs in 3D NAND flash technology. The company's new 7.68TB and 15.36TB 3D SSDs deliver 16 times more storage density over previous versions and 12 times more density than competitor products. In its most recent evalu-

continued on page 6

“These systems consume 85 percent less space than traditional high-end arrays and enable customers to consolidate multiple racks of legacy storage while improving quality of service.”

continued from page 5

ation of midrange storage arrays, Gartner gives StoreServ its highest score for the consolidation use case.

“These are the most dense and scalable all-flash arrays in the industry,” said Doehler. “A single system supports up to 24 petabytes of usable SSD capacity. These systems consume 85 percent less space than traditional high-end arrays and enable customers to consolidate multiple racks of legacy storage while improving quality of service.”

HPE says its 3PAR StoreServ arrays can reduce the storage footprint by 75 percent. Power savings is significant, as well. Per IOPS, flash uses around 600 times less energy than disk. Some data center operators have reported up to 90 percent reductions in power and cooling costs when replacing disk with flash.

The high storage densities in these new SSDs are enabled by 3D NAND, a type of flash memory in which the memory cells are stacked vertically in multiple layers. This architecture allows higher densities at a lower cost per bit than traditional planar NAND flash technology, which uses a single layer of memory cells. In fact, 3D technology has allowed flash to surpass HDDs in areal density, which is the measure of how much data can be stored on a given unit of physical space.

Enhanced Endurance, Administration

Two virtualization technologies — 3PAR Adaptive Sparing and 3PAR Express Layout — also contribute to the higher densities and improved management features that make StoreServ ideal for consolidating multiple racks of legacy high-end storage.

Adaptive Sparing is a patented 3PAR technology that optimizes capacity and extends drive endurance. Traditional array architectures dedicate spare drives for managing drive failures, but Adaptive Sparing unlocks this extra capacity to perform wear-leveling that extends the functional life of all SSDs. In the event of a drive failure, this extra capacity is automatically reclaimed to rebuild data from the failed drive.

Express Layout allows StoreServ controller nodes to share concurrent access to SSDs for higher throughput. HPE says allowing multiple controllers to drive input / output reduces management overhead by more than 10 percent. Concurrent access also allows controller nodes to use capacity from multiple drives to build logical capacity, making it easier for administrators to control where and how data is stored.

Flash further contributes to storage space saving by optimizing deduplication and compression — technologies that

have been popularized in the backup world but are not typically used in primary storage because they are so slow on disk. The StoreServ all-flash platform has thin capabilities built into both array hardware and software in order to deliver efficient data compaction technologies such as thin provisioning, block-level deduplication and space reclamation.

HPE flash gains performance advantages over disk with these tasks because StoreServ controllers feature dedicated chips that perform specific functions in silicon for extremely high performance and low latency. In contrast, in disk-based architectures, a general-purpose processor must be provided instructions and precompiled code in order to process data.

Data Protection Features

As flash expands across the data center, it also enables data protection and security at speeds not possible with legacy architectures. StoreServ flash arrays include built-in data availability and business continuity capabilities that reduce risk without performance penalty.

HPE 3PAR Remote Copy support for asynchronous streaming enables replication while optimizing latency, distance and recovery. Especially suited for the all-flash data center, asynchronous streaming eliminates the overhead of synchronous replication while providing a near-exact data copy with recovery point objectives measured in seconds.

HPE Recovery Manager Central (RMC) accelerates and simplifies backup by directly moving application snapshots from 3PAR flash to secondary protection storage up to 23 times faster than traditional storage backup. This not only reduces risk, but applying HPE StoreOnce deduplication to the process reduces capacity requirements by 95 percent or more.

“Unchecked data growth has pushed legacy storage environments to the edge of inefficiency,” said Doehler. “For years, companies have been cobbling together solutions that require more and more drives, yet result in unpredictable and poor performance, greater complexity, increased latency and higher risk.”

“As IT organizations are tasked with managing more and more data-intensive applications requiring a mix of storage attributes, flash has emerged as a dynamic solution. With these high-density StoreServ solutions, HPE is introducing innovations that make an all-flash data center an increasingly sensible alternative to traditional storage industry practices.”



Gain a competitive advantage

Flash storage so fast, scalable and affordable that it's almost unfair.

The HPE 3PAR StoreServ family of flash-optimized storage delivers performance advantages other midrange platforms simply cannot match. These lightning-fast arrays deliver 99.9999 percent uptime, all at a price point equal to traditional disk-based storage. Get rapid and automated provisioning, multi-tenant design, hardware-accelerated deduplication and sub-1ms latency — all within a single tier-1 storage architecture.

Contact your Sigma representative to learn more.



www.sigmasol.com
888.895.0495

SIGMA
A PIVOT COMPANY

Divide and Conquer?

Component-focused microservice architecture creates a shift in the way applications are developed.

As recently as two years ago, surveys showed that IT decision-makers overwhelmingly favored keeping core business applications and workloads running in the data center. There's been a seismic shift since then, however, as organizations find they can achieve significant cost, scalability and flexibility advantages by cloud-enabling essential apps.

An increasingly popular approach to application development known as microservice architecture is facilitating this movement. In this method, applications are built as collections of small, independent services that run their own distinct processes. The idea is to break up rigid, monolithic applications into smaller, more manageable pieces that can be released rapidly for increased agility and scalability.

“The microservice architecture is changing the way software is being built,” said Steve Anderson, founder of San Francisco-based venture-capital firm Baseline Ventures. “It enables the continuous delivery of applications in any language, across a host of on-premises and cloud environments ... giving enterprises the control and choice for where they run their workloads.”

In the traditional “waterfall” approach to application development, large development teams with dozens of software engineers gather all known requirements for an application, develop all elements and finally test the app before release. This is a reliable approach for producing monolithic applications such as business management software suites and transaction processing apps that require tightly integrated components, code modules and services.

However, this method has serious drawbacks in a cloud-enabled, web-scale world. Not only does it take a long time, but it forces the development team to go back to square one when errors are found. Even if just one small piece of the application needs revision, the whole app has to be compiled, tested and deployed again. That runs counter to the goal of cloud-enabling applications so that resources can be rapidly aggregated, orchestrated and scaled for use by global communities.

‘SOA Done Right’

By breaking apps into small components that perform specific functions, microservice architecture allows developers to update, change, add or remove services on the fly with minimal effect on any other parts of the application. Small teams of just a few developers focus on one piece of the app,



allowing them to perform continuous integration and deployment, and enabling a cycle of innovation of days or weeks rather than months or even years.

The idea of component-based application development is not really new. In fact, the microservice concept is closely related to the service-oriented architecture (SOA) approach in which applications are also broken down into components and services. In fact, microservices are often described as “SOA done right.”

There are, however, important distinctions. SOA is an attempt to get applications to interact and exchange data with other applications, independent of the underlying platform. As such, SOA is more of a strategic initiative involving multiple applications, systems and organizational units.

A microservice, meanwhile, is not designed to connect various applications, but rather to perform one specific function as part of a single application. This stripped-down functionality is reflected in the code base: An SOA service might comprise 1,000 lines of code compared to fewer than 300 for a microservice.

Additionally, microservices can be developed in any programming language and can communicate through a variety

of lightweight mechanisms such as HTTP-based RESTful APIs or simple text-based protocols such as STOMP. In an SOA environment, services rely on proprietary and code-heavy Enterprise Service Bus and Web Services specifications to communicate.

Netflix is frequently cited as among the first and most visible adopters of the microservice approach. The company was built on a unique DVD-by-mail business model supported by a monolithic Java application. Hundreds of engineers were involved in continually refining the application, and code updates were made every two weeks.

As streaming video gradually supplanted disc-based content, the company transitioned to a cloud-based environment. Now many small teams of engineers work on the end-to-end development of specific microservices that are aggregated and orchestrated to stream digital entertainment on a massive scale. Netflix operates one of the largest content-delivery networks in the world, serving nearly 70 million subscribers and accounting for nearly 40 percent of all online traffic in North America, according to a recent report from Sandvine.

Bring It Together

Amazon is another microservice pioneer, having moved to a cloud-based architecture to relieve huge server demands. The process of making a purchase on the company's website can involve hundreds of individual microservices — ranging from searching, image retrieval, pricing and shopping carts to billing, shipping, customer reviews and other product suggestions.

Of course, all of these services must be united in a way that produces a satisfactory user experience. That happens through the use of language-agnostic APIs, which expose each service's functionality and allow all services to communicate with each other. An API gateway orchestrates these interfaces, providing a single point of entry to the application and routing all requests to the appropriate service.

While there is no standard definition for microservices, this architectural trend is nevertheless gaining stream as organizations look to migrate monolithic, on-premises applications to the cloud in order to reduce complexity, improve agility and create maximum scalability. As with any emerging technology approach, organizations must plan carefully and do their due diligence before committing to this architecture. However, there is mounting evidence that this approach can streamline the deployment of applications that hold up to Internet-scale challenges.

“Microservices solve a lot of challenges, and that's why they're becoming the standard architecture both within and between applications,” said Dennis Callaghan, senior analyst of infrastructure software at 451 Research. “We anticipate accelerated adoption of microservices in enterprises this year.”



The Agile, Open and Secure Data Center

Cisco Application Centric Infrastructure (ACI) helps you create an automated IT environment that is fast, flexible, and responsive to the needs of your business. It accomplishes this using a software defined networking (SDN) policy model across networks, servers, storage, security and services.

ACI is ideally suited for microservices architectures. Network complexity is eliminated, so it no longer dictates application deployment or operation. Instead, application needs dictate IT deployment and operation.

Contact Sigma to learn more about ACI and microservice architectures.



www.sigmasol.com
888.895.0495



Best of Both Clouds

A hybrid cloud enables organizations to gain the agility of the public cloud while maintaining security and control.

Public or private cloud? That has been the conundrum organizations face as they seek to leverage the flexibility and scalability of the public cloud while addressing security, privacy, performance and availability concerns. The good news is that the cloud is not an “either-or” proposition. Increasingly, organizations are turning to the hybrid cloud model to gain the best of both worlds.

As the name implies, a hybrid cloud uses a combination of public and pri-

vate cloud services to meet the requirements of various workloads. For example, an organization could use a private cloud to maintain control over mission-critical applications and sensitive data, while leveraging the public cloud for data backup and archival, disaster recovery and application development and testing.

The benefits of the hybrid cloud are so compelling that it’s widely viewed as the logical end game for most organizations. The Rightscale 2016 State of the Cloud Report found that hybrid cloud

adoption increased from 58 percent to 71 percent year-over-year, with the average organization using six different clouds.

In a recent study by IDG Research Services, 83 percent of respondents said they currently use or plan to use a hybrid cloud environment, and 73 percent agree that the hybrid cloud model creates a path to digital transformation. The survey found that the hybrid cloud helps organizations increase IT agility by making it faster, easier and less expensive to implement technology initiatives.

By reducing IT costs, the hybrid cloud also enables greater investment in digital transformation, and the most aggressive hybrid cloud adopters are the most advanced at digital transformation. Of those surveyed, organizations with a significant number of hybrid cloud workloads are three times more likely than non-adopters to reach their digital business and infrastructure readiness goals.

Combining Forces

The public cloud offers attractive economic and operational benefits, allowing organizations to tap IT infrastructure, applications and services on a pay-per-use basis. However, many organizations are reluctant to place sensitive applications and data in the public cloud. In addition, accessing applications via the Internet brings latency and reliability concerns that make the public cloud unsuitable for many mission-critical workloads.

A private cloud gives organizations the agility of the cloud model in a dedicated, single-tenant environment. Whether deployed onsite or hosted, a private cloud gives organizations control over data center resources while accelerating service deployment and enhancing operational efficiencies.

However, the growing complexity of IT environments makes a hybrid cloud solution the best option in most cases. Organizations can retain control

where needed while relying upon cloud service providers for nearly limitless capacity without the associated infrastructure and management overhead.

It's important to note that a hybrid cloud isn't the same as using public and private cloud services simultaneously. In a true hybrid cloud, the public and private clouds are integrated, with a single management interface that makes it easy to move applications and services between the two environments. This eliminates the need to overprovision resources to account for spikes in demand — particularly beneficial for organizations that have highly dynamic workloads.

“Companies must take advantage of all of their resources to succeed and deliver services and apps in an ‘as-a-service’ model, regardless of whether they reside in a public or private cloud or on their existing infrastructure,” said Laura Sanders, IBM Global Technology Services General Manager for Systems Services. “Not surprisingly, clients are seeing great value from hybrid environments in driving business results and transformation. It's a positive balance of optimization with flexibility and agility.”

Management Is Key

A global study by the IBM Center for Applied Insights found that two-thirds of organizations implementing a hybrid cloud are gaining competitive advantage from their hybrid environments. Among hybrid cloud leaders, 90 percent report that the hybrid cloud is crucial to their business strategy and success.

Despite its potential benefits, a hybrid cloud doesn't always deliver ROI. In a survey from The Bunker, nearly two-thirds of CIOs

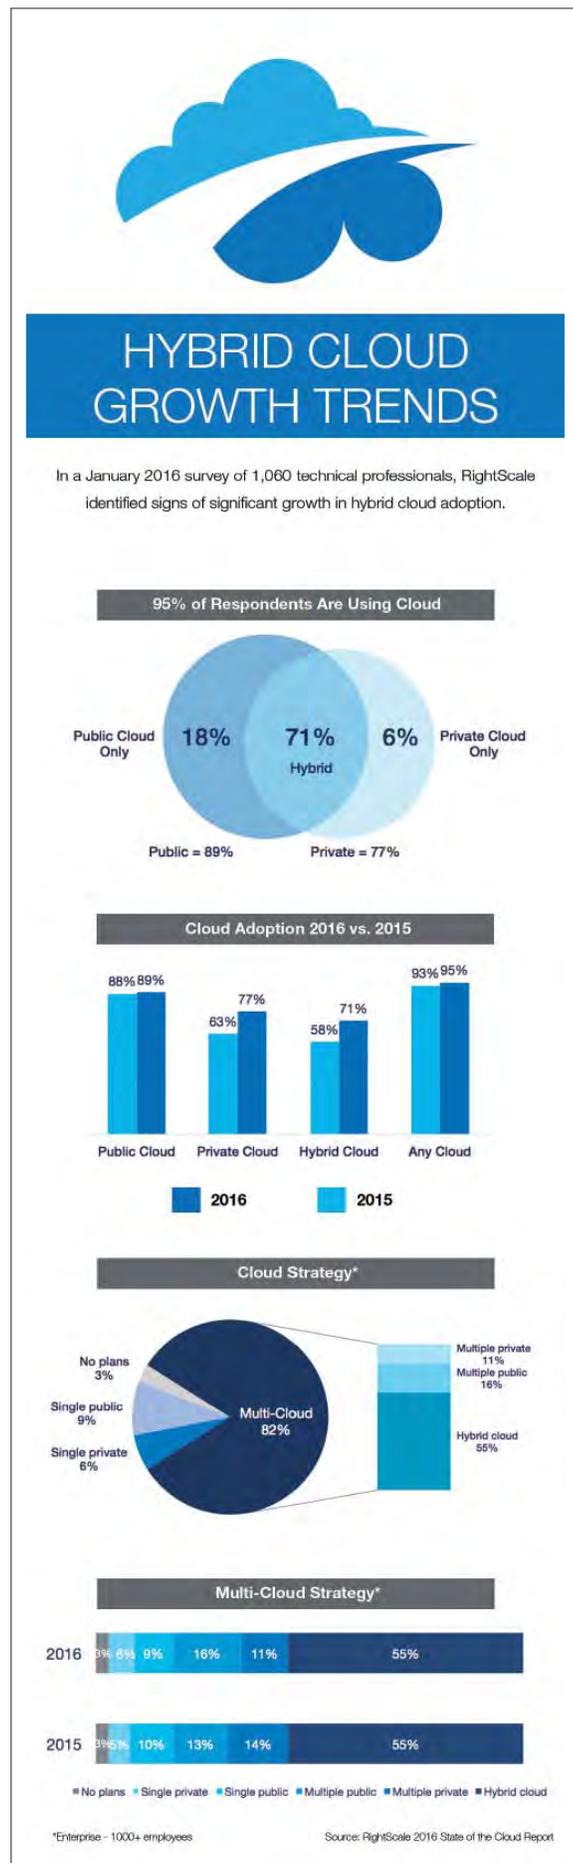
and IT decision-makers said their hybrid cloud deployments failed to meet expectations. Primary reasons were a lack of in-house expertise, poor advice, and a lack of integration of cloud and non-cloud resources.

Gaining the advantages of a private cloud requires a strategic approach. Experts at Taneja Group recommend that organizations begin with one or two primary use cases and define their performance, security and availability requirements. Only then can the organization select the appropriate public and private cloud services and develop best practices for implementation, configuration and ongoing management.

A hybrid cloud must be managed with as much rigor as traditional data center solutions. Organizations should take care not to allow a “shadow IT” approach to the hybrid cloud.

“With increasing cloud spending, many enterprises will have begun use of hybrid solutions without IT's direct involvement. With business unit managers increasingly buying their own cloud solutions, hybrid cloud architectures can proliferate, and sometimes not be aligned with the enterprise architecture,” said Chris Morris, vice president, Cloud and Services, IDC Asia/Pacific.

The appeal of the hybrid cloud lies in its ability to provide cost savings and operational efficiency by allowing organizations to choose the optimal platform for each application and service. But that's only the beginning. When properly deployed and managed, a hybrid cloud can help transform an organization by enabling IT to focus on strategic business value and innovation.



Focus on your business. We'll take care of the rest.

With a predictable, monthly cost structure, Sigma's comprehensive managed services reduce IT costs and risks. Contact us today to learn more.



www.sigmasol.com | 888.895.0495

