# Enable Consumption-Based Data Management

Storage Switzerland, LLC

VEEAM

Hewlett Packard Enterprise

*The data protection process is full of challenges. One of the biggest is the economic unpredictability typical to the process. IT needs to upgrade backup infrastructures to keep pace with growing primary storage capacities and increased demand for more rapid recoveries while meeting the long-term retention requirements of recent regulations. These upgrades are almost always a surprise, and IT personnel rarely include them in advanced budget planning.*

# CHAPTER 1:
# The Impact of Unpredictable Data Protection

Most production IT projects are planned months, if not years, in advance of deployment. However, advanced planning of infrastructure upgrades for the data protection architecture seldom occurs. The problem is that updates to the backup infrastructure can often be expensive and because they are not budgeted, the organization is forced to scramble to allocate funds toward the project. Reallocation of funds means that other projects are delayed or not rolled out to their original scale. The lack of planning also forces the organization to pay extra by not spending the appropriate time finding more cost-effective solutions.

The lack of predictable data protection spend cycles means that sometimes funds just can't be allocated and IT is forced to cobble a solution (workaround) together to provide some means of protection. Finding funds is especially problematic when IT budgets are flat or declining. While creative, workarounds increase risk. The workaround may not provide frequent enough protection nor provide rapid enough recovery.

There is also a ripple effect of stretching the data protection process to protect a new application or environment in that it may leave incumbent environments more exposed to data loss or slower than originally promised recovery times.

The lack of budget planning also forces organizations to continue to use legacy data protection platforms that don't have advanced abilities regarding backup performance, rapid recovery, cloud support, and long-term data protection. These legacy platforms often require dedicated administration which increases the total cost of ownership.

Even if the organization decides to switch data protection platforms, it may not be able to since costs typically increase temporarily as the new solution is brought in and the old solution is phased out. For a time, the organization needs to run both solutions, and during that time, costs are effectively doubled.

In addition to increasing requirements like more frequent backups, more rapid recovery and increased data retention, the unprecedented growth of production data forces organizations to add additional secondary storage capacity. Secondary storage is typically bought dozens (if not hundreds) of terabytes at a time, which means a significant upfront cost when the current secondary storage system runs out of capacity.

—

While data protection planning and budgeting are essential, the reality is that most organizations won't or can't plan for "the spend." To be fair, they do the best they can, planning and budgeting for production spending. A new method of acquiring IT solutions is required, but it has to be a solution that doesn't force IT to move everything to the cloud.

Consumption-based IT is the answer. It enables IT to keep its resources on-premises, but the organization only pays for the resources that are actually in use. It eliminates the high costs and low utilization associated with upgrades. Finally, it provides active capacity management that regularly tracks IT resource utilization and adjusts IT spending accordingly.

# CHAPTER 2: Why Data Protection-as-a-Service is Unpredictable

Economic unpredictability is the most significant challenge associated with data protection. The pool of data that must be protected continues to grow exponentially, and meanwhile, recoveries must be nearly instantaneous.

Our previous chapter discussed the challenges inherent in meeting these unpredictable data protection requirements with a more traditional, fully on-premises IT infrastructure approach. The unpredictable spend cycles cause IT to scramble to allocate budget, often forcing a workaround to be cobbled together that may not provide sufficient levels of protection. The business may be forced to rely on legacy technologies, and pieces of the environment may be left exposed to data loss or slow recovery times. Not to mention, the business may also incur the cost and headache of running multiple point solutions at a time during the recovery process.

The pain points associated with investing in and managing a fully on-premises disaster recovery environment incentivize many IT planners to look to cloud-based data protection-as-a-service (DPaaS), including disaster recovery-as-a-service (DRaaS). One of the DPaaS goals is to smooth out the costs associated with data protection. However, DPaaS brings its own set of unpredictability and challenges.

DPaaS enables the business to defray upfront costs. However, the advent of stricter compliance regulations and increasing utilization of data and analytics to fuel new business outcomes require more data to be retained for longer periods of time. As capacity requirements increase and as data is stored over

months and years, this cost model can become inhibitive, far surpassing the cost that would be required to store that data on premises. Meanwhile, egress fees accumulated when the business goes to recover data are unpredictable and may be very sizeable since cloud service providers charge a per-gigabyte fee for data to be removed from their cloud storage and placed onto another infrastructure. Adding to this cost structure is the fact that any on-premises hardware that is required to run primary storage workloads must still be purchased up-front, as opposed to when it is consumed. Analyzing cloud utilization and costs remains a developing art as most organizations have fuzzy visibility at best into the cost structure of their cloud services over time.

Typically, DPaaS does not coexist with on-premises data protection infrastructure so the organization will typically go a cloud-only route. As a result, these costs can scale quickly and significantly. This also puts the organization at risk since DPaaS solutions tend to focus heavily on software innovation and as a result, might not take advantage of hardware-level innovations. Additionally, many of the vendors in this space are new startups. The customer is relying solely on a vendor, which might not be in business in five, ten or fifteen years, for critical business continuity.

In addition to unpredictable costs, the recovery time itself is also unpredictable in a DPaaS model. Recovery time will vary depending on the compute cycles that the cloud service provider has available to apply to the recovery task (many other organizations might also be recovering data at the same time), how much compute is required to complete the recovery, and the cloud provider's network latency.

Bringing consumption-based payment models to on-premises infrastructure is a compelling answer to the economic unpredictability of data protection. This approach enables customers to avoid heavy upfront capex investment in on-premises infrastructure, like in a DPaaS model, and introduces granular usage metering and scalability (both up and down) to enable customers to pay for the resources that are actually in use.

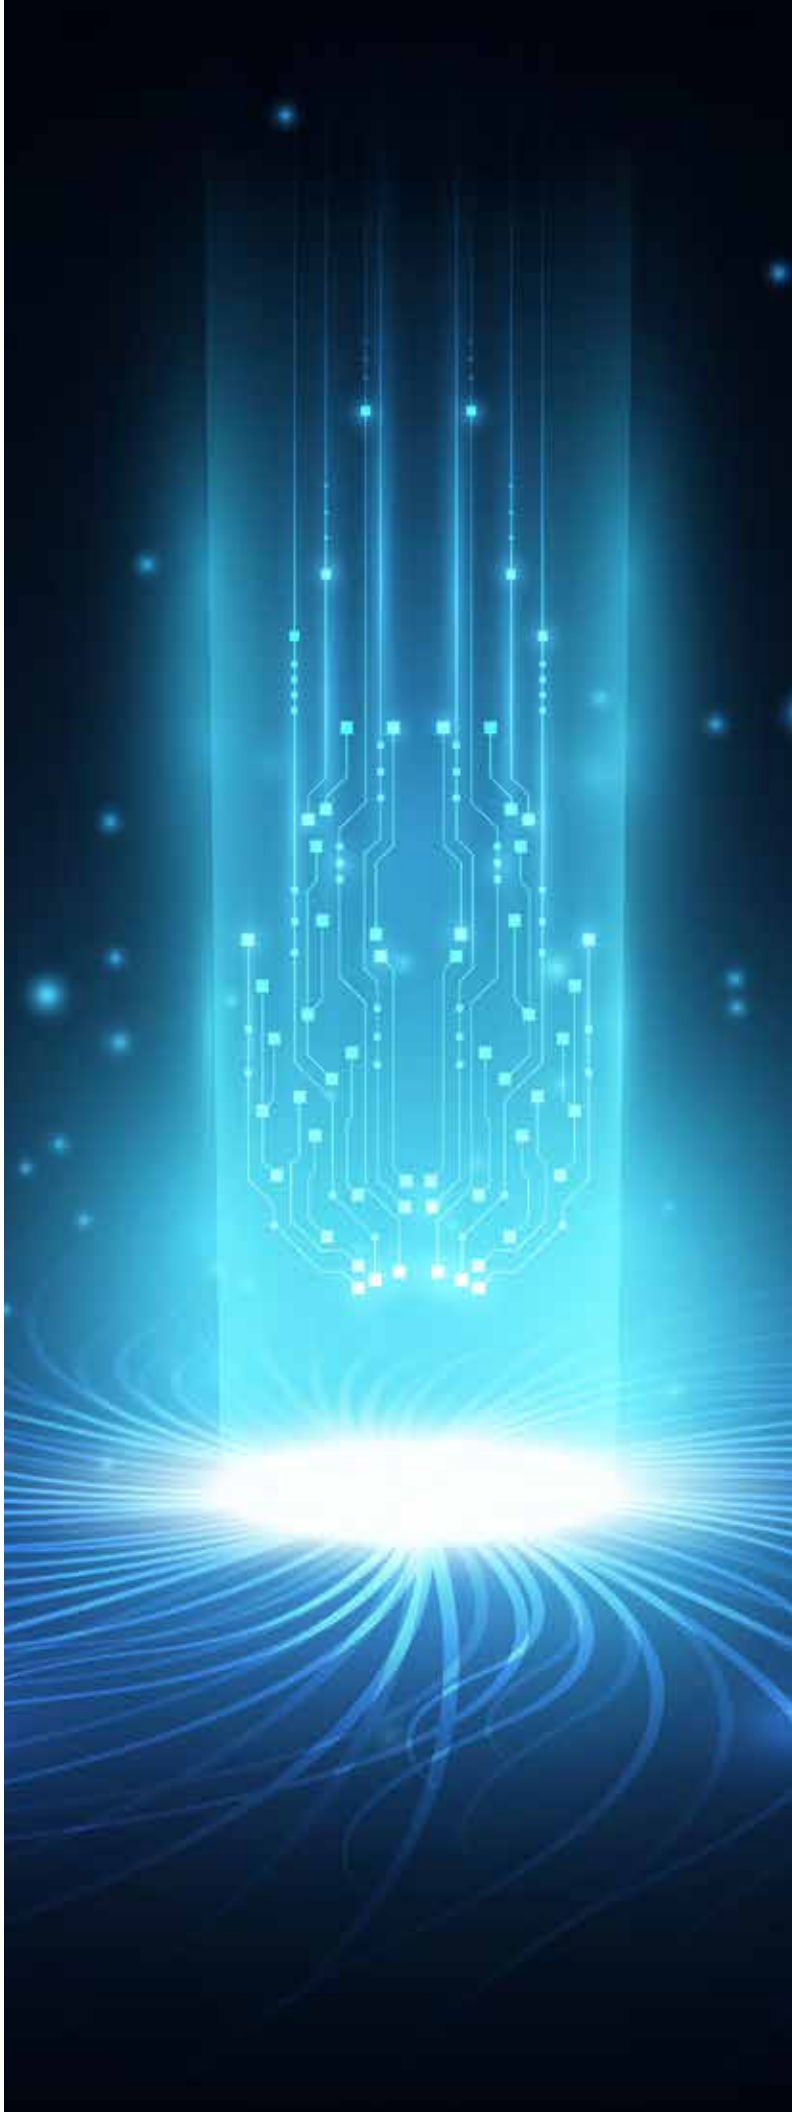

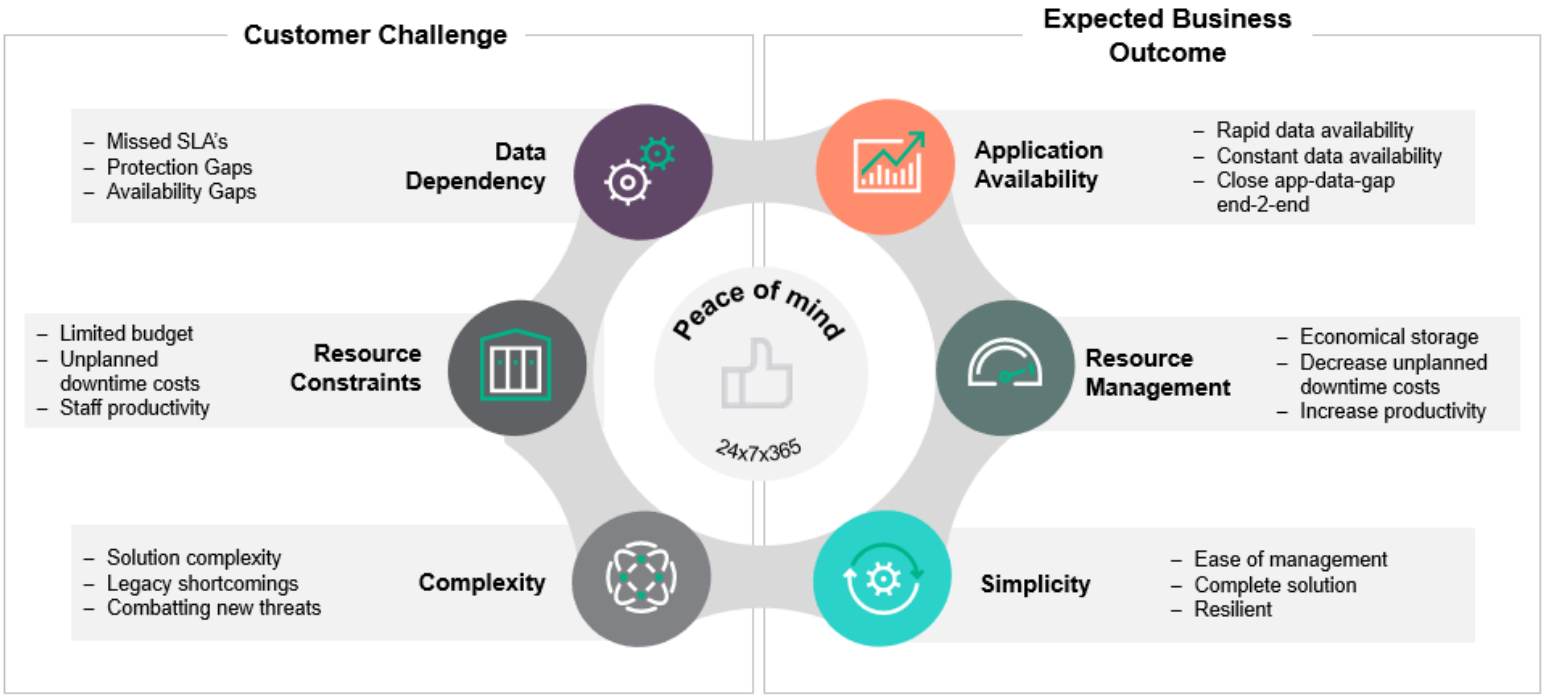


# CHAPTER 3: Consumption-based Data Protection Brings Predictability

Bringing predictability to the data protection process is not a lost cause. The problem is that IT lacks a purchasing model to make data protection predictable and a software/hardware stack that works well with that model. As discussed in the last installment, the cloud model requires too dramatic of a change. IT needs something that simulates the cloud business model but works both on and off-premises. Consumption-based IT provides a compelling alternative to the cloud-only business model for organizations looking to bring predictability to data protection.

## What is Consumption-based IT?

Vendors have tried to re-create cloud models in the past. They would supply the organization with IT hardware and software and the organization could pay for it on a periodic basis. In most cases these solutions were actually leases. There was almost always a minimum period of time the equipment needed to be kept, generally there were restrictive return policies, and payments were typically fixed even if all resources were not fully utilized.

Consumption-based IT is different than a lease. The vendor owns the hardware and software it supplies to the customer, so that the customer does not need to worry about buying more or returning unused hardware. These transactions happen automatically as part of the business arrangement. Furthermore, the nature of consumption means that the monthly cost will go up or down based on usage.

A key element to the consumption-based model is active capacity management, which includes metering and monitoring of hardware and software resources to make sure that capacity is always available ahead of demand, and to make sure that customers are only paying for what they use. In addition, consumption analytics provide visibility into usage and costs and help organizations forecast future capacity needs.

## Impact of Consumption IT on Data Protection

The result of a consumption model is that the organization benefits from the Opex-like business model of the cloud but still maintains an on-premises infrastructure footprint. The organization has its resources automatically downsized or scaled up as needed, and it pays only for the IT resources that were used during the billing period.

From a data protection perspective, consumption IT automatically grafts the purchase of software modules and capacity licenses as well as additional storage capacity into the business model. If, for example, the organization decides to move data to the cloud for long-term storage, the customer's costs for on-premises storage resources will scale down in alignment with their usage (subject to any mutually agreed upon minimum commitments). The business is no longer hit with surprise upgrades or capacity expansions; rather, a gradual increase of data capacity or utilization of a faster-performing tier of storage infrastructure is reflected gradually as an operating expense.

## The Benefit of Consumption IT for the Business

The business benefits from consumption-based IT in three ways. First, infrastructure management and planning are greatly simplified. IT no longer has to perform the task of understanding and adjusting its resource utilization. The consumption service provider does that for them. Second, consumption-based IT provides business agility. An organization can change directions almost instantly and the consumption service provider responds with the appropriate addition or removal of IT resources. Finally, the organization sees an improved time to market. It no longer has to make upfront investments in IT hardware and software that go unused while the project ramps up. The organization only pays for what is actually being used.

# CHAPTER 4:
# Veeam and HPE Consumption-based Data Management

Consumption-based IT enables the customer to obtain the pay-as-you-go cost effectiveness of cloud services while retaining an on-premises footprint. It also dramatically simplifies IT infrastructure management and planning. Data protection administrators in particular can benefit from this method of acquiring IT resources.

Hewlett-Packard Enterprise (HPE) and Veeam® Software are collaborating to offer Veeam's core offering, the Veeam Availability Platform™, through HPE's GreenLake Flex Capacity consumption-based IT procurement model.

Veeam Availability Platform manages and protects virtualized, physical and cloud-based primary and secondary storage resources from the centralized Veeam ONE console. Core capabilities include backup, recovery, replication and failover.

- With its Availability Platform, Veeam aims to facilitate recovery point objectives (RPOs) of less than 15 minutes. To do so, it offers plug-ins for enterprise applications including SAP HANA, alongside other capabilities such as its Backup from Storage Snapshots feature, which enables backups and replicas to be created from snapshots as often as every 15 minutes with minimal (or no) impact on the production environment.

- Veeam Cloud Tier enables automatic tiering of "colder" data to cloud-delivered object storage services including Amazon S3 and IBM Cloud Object Storage. This enhances the elasticity of backup storage infrastructure and can mitigate the amount of on-premises infrastructure that must be purchased and managed (thus saving on costs).

- Embracing the cloud also plays an important role when it comes to Veeam's replication and recovery capabilities. For example, the Veeam Cloud Mobility

capability enables two-step migration and recovery to AWS, Azure and Azure Stack. The Veeam Cloud Connect capability provides flexibility of choice for customers that desire to work with a different service provider. Veeam utilizes built-in WAN Acceleration to overcome the latency inherent in the cloud and applies encryption at the source as well as when data is in flight and at rest, to enhance security (a typical concern when offloading backup and disaster recovery to the cloud).

Through the HPE GreenLake consumption-based IT delivery model, storage managers have the option of deploying Veeam's Availability Platform on HPE's StoreOnce storage systems or HPE's Apollo servers. They can also, depending on their workload's capacity and performance requirements, use a mixture of the two. Usage is billed per GB or TB of storage capacity consumed, and for the Veeam Software, per sockets and agents used. The companies benefit from HPE's unique consumption model from a lifecycle asset management perspective, as well as a complementary and robust global footprint.

## StorageSwiss Take

Organizations large and small are facing business-threatening gaps to data availability due to complex and expensive legacy technologies and procurement models. With their new consumption-based offering, HPE and Veeam are making data protection available where previously it was price-inhibitive because of high upfront investments in hardware and software, as well as the cost over time of solution upgrades and administrative staff. Meanwhile, the solution stands to make data protection more cost-effective and easier to manage for enterprises with existing implementations. The solution is robust enough to offer comprehensive enterprise support, while also being flexible enough to scale down to address smaller companies' needs or to introduce specific functionalities into large enterprise environments. To learn more about consumption-based IT and how it enables a more predictable data protection infrastructure watch our on-demand webinar, "Consumption-Based Data Management Providing Peace of Mind."

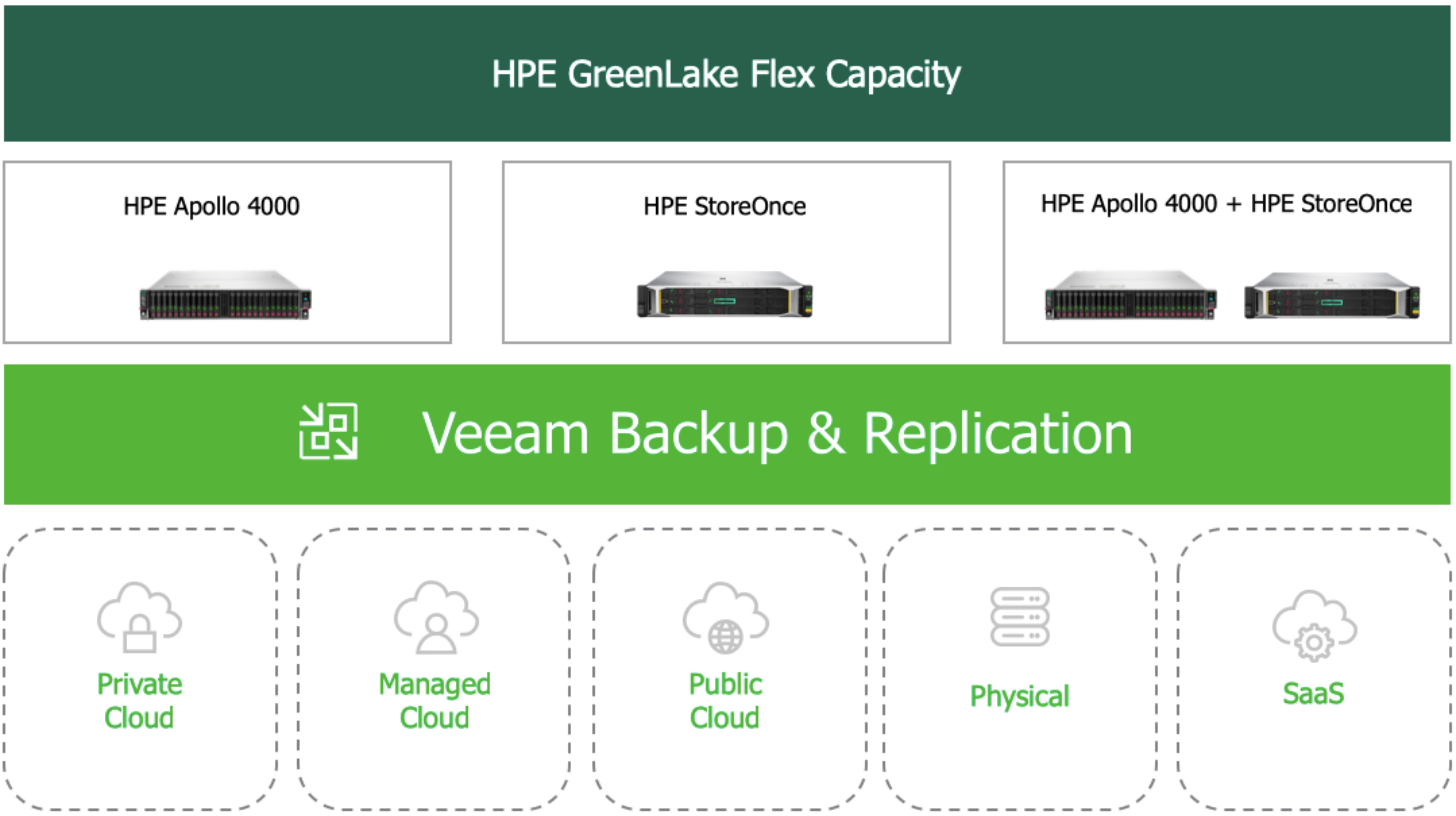

Storage Switzerland, LLC

## The Firm

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.

VEEAM

Hewlett Packard Enterprise

## About Our Partner

Veeam is the global leader in Intelligent Data Management. Veeam Availability Platform is the most complete solution to help customers on the journey to automating data management and ensuring the Availability of data. The combination of HPE (Hewlett Packard Enterprise) + Veeam helps customers optimize their data Availability solutions to deliver maximum application uptime, improve resource management and reduce complexity. Integration of Veeam Availability Platform with key HPE Storage Systems and flexible-capacity procurement models helps customers derive new insights, simplify IT, lower costs and increase business agility through Intelligent Data Management. To learn more on HPE and Veeam intelligent data management solutions, visit: https://go.veeam.com/hpe-veeam-digital-hub.

## The Analyst

George Crump is the founder of Storage Switzerland, the leading storage analyst firm focused on the subjects of big data, solid state storage, virtualization, cloud computing and data protection. He is widely recognized for his articles, white papers, and videos on such current approaches as all-flash arrays, deduplication, SSDs, software-defined storage, backup appliances, and storage networking. He has over 25 years of experience designing storage solutions for data centers across the U.S.