

Taking a Modern Approach to Data Protection for Web-Scale Infrastructure

The growing importance—and complexity—of data protection means old approaches no longer will get the job done in an era of exploding data volumes and ever-changing business requirements. It's time to reimagine and reengineer your IT infrastructure for a more efficient, affordable and manageable data protection framework.



Threats to data have never been greater or more sophisticated. From ransomware attacks and a dizzying array of unmanaged endpoints to simple human error, enterprises must do more to ensure the availability of their mission-critical data when and where it is required. The regulatory, legal, operational and financial implications of lost or compromised data are staggering.

As a result, more time and money are being poured into new data protection solutions. Data protection is expected to become a \$120 billion global market by 2022, growing by 16% on a compound annual basis since 2016.¹ And privacy initiatives such as the Global Data Protection Regulation (GDPR) are key drivers in making data protection a top priority for enterprise IT organizations and their business counterparts.

Unfortunately, there's a catch: Legacy backup point solutions are fragmented, inefficient, difficult to manage, and can't keep up with the stringent recovery point objectives (RPOs) and recovery time objectives (RTOs) associated with service-level agreements (SLAs) and user requirements. Legacy data protection systems also are hampered by the reality that most were designed and implemented for a far different age—years before the advent of cloud computing and the software-defined paradigm.

A new philosophy about how to deploy and manage data protection that handles essential functions such as backup, archiving and data recovery is needed. This is causing innovative IT leaders to reimagine and rearchitect their secondary data infrastructure to simplify enterprise data protection, maximize storage capacity, reduce total cost of ownership (TCO), provide near-instant recovery and natively integrate with the public cloud providers.

This paper looks at a new breed of modern, web-scale data protection solution—and examines how it makes data protection more manageable, more reliable and more affordable than legacy approaches.

The problems with Legacy Data Protection

Today's data protection threat landscape must account for and address a wide range of factors, from data breaches, configuration errors, and technology refresh cycles to more stringent SLAs and plain old user error. Although these challenges may not be new, the frequency and impact of such threats have increased dramatically. Additionally, organizations now need to meet much stricter SLAs and compliance requirements, which is exacerbated with rapidly growing data volumes.

Legacy data protection solutions, therefore, no longer are efficient or effective in dealing with rapidly evolving and high-impact threats. The result is a cascade of drawbacks and limitations, each with significant implications for the enterprise. For instance:

- **Highly fragmented data protection environments.** These include the existence of multiple point solutions for different data protection functions, such as backup, archiving and recovery, resulting in higher costs, more human intervention and potential gaps in protection schemes.
- **The need to manage multiple product vendors.** Having multiple vendors for different data protection point products (backup software, target storage, replication, media and proxy servers, and bolt-on cloud gateways) means multiple proprietary hardware solutions with their own software, user interface and maintenance model. This makes it harder to manage vendor relationships and the solutions themselves, difficult to ensure timely and smooth technology refresh, and support issues that may lead to confusion and finger pointing.
- **Storage inefficiency.** With fragmented infrastructure, multiple copies of the data sit in their own silos, plus dedupe and compression are limited to the node level.
- **Data protection processes and systems that are not optimized for modern architectures.** Traditional data protection solutions were designed for environments operating primarily with on-premises infrastructure, typically without adequate—or any—cloud integration capabilities.
- **Expensive insurance policy.** Once the data lands on the legacy backup solutions, it often sits idle on the target devices. The data cannot be made available to support other secondary workloads.
- **A need for forklift upgrades to accommodate evolving business requirements and new technologies and ensure that IT modernization keeps pace with digital transformation.** Forklift upgrades are expensive, time consuming and disruptive to the business. Additionally, these solutions are not designed to scale limitlessly to address exponential data growth stemming from multiple directions.

¹ "Data Protection Market by Component: Global Forecast to 2022," MarketsandMarkets, February 2018

- **Longer backup windows and slow recovery.** Business no longer is confined to traditional 9-to-5, five-days-a-week boundaries. Excessive downtime for backup, and to recover from unplanned service interruptions and data loss, has major financial, operational, legal, regulatory and brand impact.

How a modern, web-scale solution changes the game in data protection

So, if traditional approaches are inefficient and don't do nearly enough to ensure the very high levels of data protection required in today's demanding business environment, what kinds of changes are needed in order to change the rules of the game?

It is essential to architect your data protection environment with a modern, web-scale approach. This not only accommodates the massive increase in data—especially unstructured data—experienced by all enterprises, but also supports hyperscaled workloads that require high performance, resiliency, storage capacity, cost efficiency, rock-solid security, easy management and agility.

A modern, web-scale data protection infrastructure for today's demanding IT environment ideally must support:

- Elimination of multiple, disconnected hardware and software silos that have often sprung to accommodate different data protection functions.
- Tightly converged infrastructure design that integrates everything from storage hardware and software (backup, deduplication and compression) to cloud gateways and tape libraries.
- Faster backups for near-instant RTOs.
- Unlimited snapshots and clones for short RPOs—ideally as short as 5 minutes.
- Cloud-native, software-defined architecture that offers customers increased flexibility and choice.
- Scale-out architecture for improved performance at web scale.
- Non-disruptive online upgrades to eliminate planned downtime.
- Web-scale expandability with greater resiliency, higher storage efficiency and near-instant recovery.
- Improved storage capacity and performance utilization, which reduces the likelihood of overprovisioning.
- Cost-efficient commodity hardware, rather than proprietary, specialized infrastructure. Reduced Capex but also lower Opex from lower management costs, smaller footprint, and reduced power and cooling requirements—all of which result in lower TCO and faster time to value.
- Simplified management, reducing the need for extensive human tuning and monitoring of infrastructure performance, responsiveness and health.
- Expanded capabilities beyond backup to support other secondary workloads like files and objects, test/dev and analytics.

The Cohesity approach to modern, web-scale data protection

Cohesity, an innovator in developing and delivering efficient secondary storage solutions that ensure data protection for web-scale requirements, has developed a hyper-converged platform that eliminates data protection silos on a software-defined modern architecture for building an efficient infrastructure.

Cohesity DataProtect is an end-to-end data protection solution based on a hyper-converged, software-defined platform model. Designed with cloud in mind, DataProtect brings Google-like web-scale benefits to enterprises on a platform that spans from core, to the cloud, to the edge. Cohesity's DataProtect software provides:

- Instantaneous mass restore to unlimited virtual machines (VMs), to any point in time.
- Guaranteed data resiliency with strict consistency.
- Global storage efficiency from global variable-length dedupe.

- Non-disruptive upgrades.
- Easy integration with existing data center infrastructure.
- Optimization for the cloud and software-defined environments.
- Extensibility of the platform for other secondary workloads, including files and objects, test/dev and analytics.

Cohesity DataProtect breaks down the all-too-prevalent silos that house different data protection functions, including backup, archive and restore. Its hyper-converged infrastructure leverages low-cost, industry-standard hardware components and software-defined architecture for efficient scale-out, global deduplication and cost efficiency.

Cohesity DataProtect reduces management complexity with its unified infrastructure and automated policy management that allows administrators to specify unique SLA requirements such as RPO, retention, replication and archival in either the cloud or to tape libraries.

Because Cohesity DataProtect is designed with the cloud in mind, enterprise decision-makers have the confidence in knowing it works with all leading public cloud storage platforms, including Amazon Web Services, Microsoft Azure and Google Cloud.

Cohesity DataProtect software is a comprehensive solution that simplifies data protection across a single, unified, web-scale architecture. DataProtect is optimized for demanding workloads that require web-scale performance and high availability, providing sub-5-minute RPOs, instantaneous recoveries and tight backup windows.

For enterprises that have embraced modern software-defined data protection to improve agility and reduce costs, Cohesity DataProtect is an ideal fit because it delivers consistent, managed backups for all leading hypervisors, such as those from VMware, Microsoft, Citrix, Oracle and others.

Among its key features to enable modern, web-scale data protection are:

- Application-consistent backups for common operating systems and databases.
- Native protection of primary storage devices like NetApp, Pure and Nutanix.
- Google-like wildcard instant file-level search.
- VM, file and object-level recovery.
- Remote replication for disaster recovery and migrations.
- Role-based access control.
- Ransomware protection.
- Support tape archival.
- Cloud-based tiering, archive and replication.
- Software-based encryption of data at rest and in-flight.

Conclusion

As IT decision-makers look to ensure data integrity and availability in the face of rapidly evolving and emerging threats, decisions on data protection architecture become more central to the overall objective. Transitioning from inefficient legacy data protection point products to a modern, web-scale architecture helps organizations achieve that goal.

Modern, web-scale data protection offers improved manageability, scalability, storage and cost efficiency, as well as improved recovery times. By using hyper-converged solutions built on a software-defined architecture, this approach offers a new way to ensure reliable, secure and scalable data protection.

Cohesity's web-scale architecture is optimized for the demanding web-scale workloads that too often threaten the integrity and efficiency of legacy data protection systems. And Cohesity's DataProtect software simplifies the key data protection functions that until recently had morphed into multiple inefficiency silos.

Visit <https://www.cohesity.com/breakup/> to learn more.