PATHFINDER REPORT

The Emerging Role of Object Storage in Compliance with Regulatory Data Requirements

DECEMBER 2017

COMMISSIONED BY

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
1411 Broadway
New York NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
(Ground floor)
30, Artillery Lane
London, E1 7LS, UK
P +44 (0) 207 426 1050

BOSTON
75-101 Federal Street
5th Floor
Boston, MA 02110
Phone: +1 617.598.7200
Fax: +1 617.357.7495

EXECUTIVE SUMMARY

Digital transformation aligns IT innovation with business strategy and gains insight from big data to transform customer services, continuously improve business processes, invent new business models, and create new businesses. IT leaders have an opportunity to create value through digital transformation, but with opportunity comes responsibility – and regulatory requirements to protect big data collected from customers, employees and partners, as well as to protect transactional data.

Data is increasing in volume and comes from a variety of internal and external applications that support an increasingly mobile workforce and the Internet of Things (IoT). Dealing with the rapid growth of data volume continues to be a significant pain point for organizations. According to 451 Research's Voice of the Enterprise annual storage studies, the increasing amount of data capacity continues to surpass other storage challenges, such as capacity planning, costs, performance and regulatory compliance. But meeting compliance requirements is a growing concern as data volumes increase and regulators focus on data collection, analytics and reporting.

Managing user-created unstructured data is one of the top drivers for data growth in organizations, before storage requirements for relational databases, backup and archive, analytics, email messaging applications, and even data retention for information governance purposes. Unstructured data is difficult to identify, manage and retain to meet the regulatory requirements in many industries such as banking and finance, healthcare, and insurance, and to protect data to meet global data privacy laws.

Figure 1: Security Concern Driven by Compliance

What is the most important aspect of compliance that's driven your concern over the last 90 days?

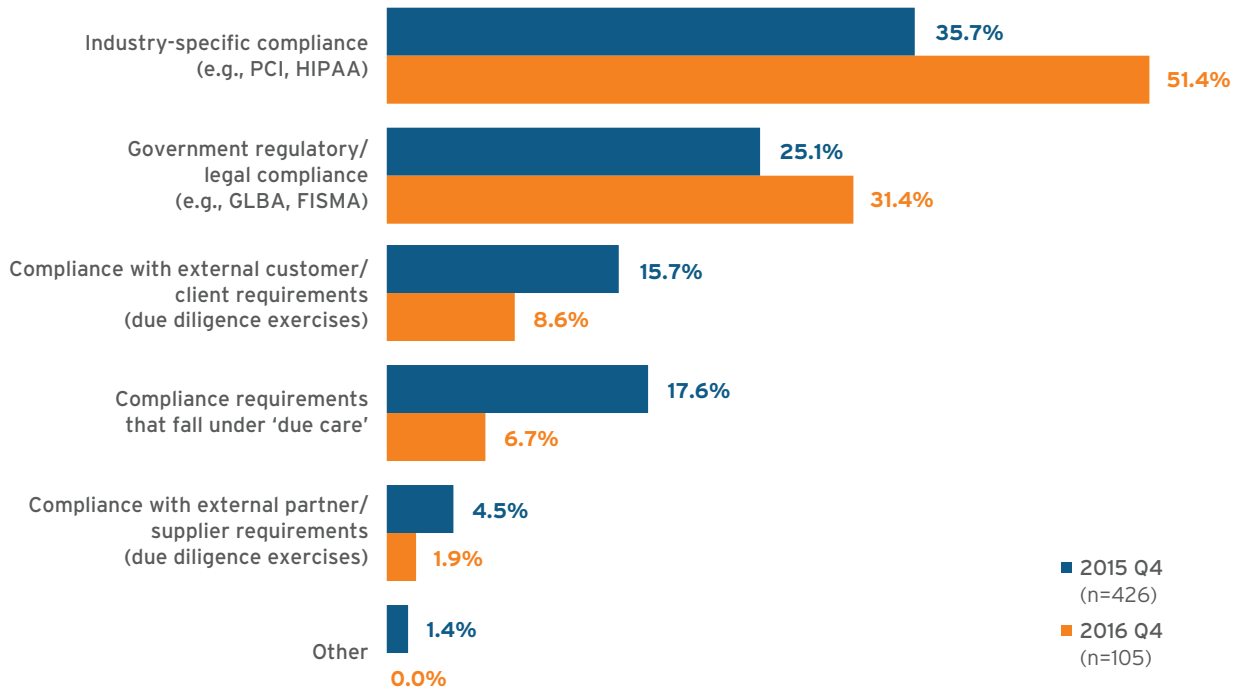

Source: 451 Research, Voice of the Enterprise: Information Security, Budgets and Outlook 2016

The top compliance concerns for IT decision-makers in 2015 and 2016 were industry-specific compliance, such as the Payment Card Industry Data Security Standard (PCI DSS), and legal and regulatory compliance, such as the Gramm-Leach-Bliley Act. Regulatory compliance challenges force organizations to retain specific data for extended periods, preventing storage professionals from deleting old data to make room for the new. The legal significance of copying data and keeping backups requires organizations to manage data growth using storage-reduction technologies, such as compression and deduplication and compression.

Figure 2: Current and Future Usage of Cloud Storage Services

Q. Which of the following public cloud and SaaS-related storage services and capabilities does your organization utilize today? And in two years? (N=460)

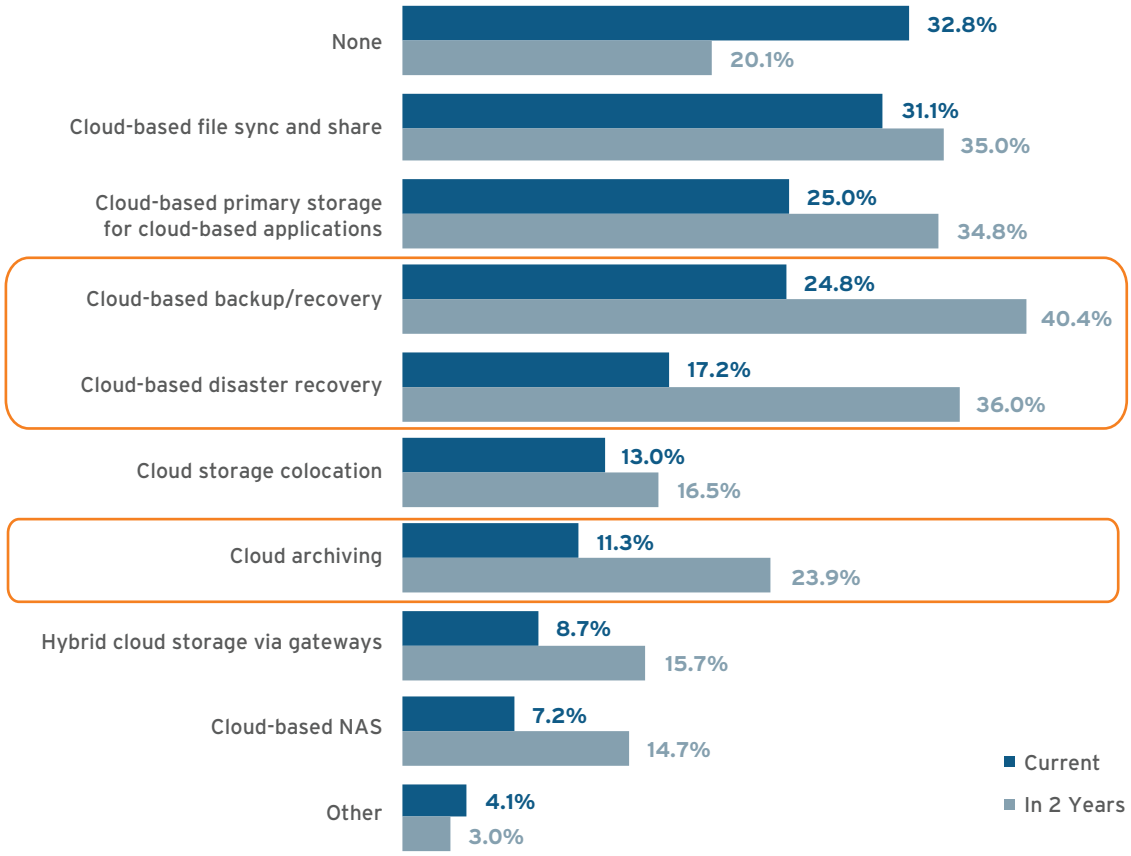

Source: 451 Research, Voice of the Enterprise: Storage, Q2 2016

Cloud storage adoption for archiving, backup and disaster recovery will more than double over a two-year period due to object storage technology's linear scalability, enhanced management capabilities, and support for petabyte-size files and trillions of individual objects. But there remains a significant market for on-premises object storage by companies adopting a hybrid, multi-cloud model, where object storage is an ideal platform for content repositories, tier 2 and tier 3 applications, and regulatory compliance to identify, categorize and apply consistent and persistent policies to unstructured data.

This report discusses the legal and regulatory compliance challenges, and the benefits of object storage for legal and regulatory compliance.

Legal and Regulatory Compliance Challenges

Complying with laws and regulations in various industries, such as financial services, healthcare and insurance, is data-intensive. There are regulations for data residency, data integrity, data access, accounting and reporting, and there are fines or penalties for non-compliance and data breaches, which can also result in reputational damage and lost revenue. Organizations must ensure security controls are persistent, correctly implemented and built into hardware and software used to collect and store data to meet regulatory requirements in regulated industries and comply with increasingly data-intensive regulations to guard against bribery and corruption, safeguard payment card data, and protect personally identifiable information in the global marketplace.

BANKING AND FINANCIAL SERVICES

The Securities and Exchange Act of 1934 requires registered broker-dealers to create (Rule 17a-3) and maintain (Rule 17a-4) specific records in an easily accessible manner and make and keep comprehensive records of securities transactions and their securities business. These record-keeping requirements allow the Securities and Exchange Commission (SEC), self-regulatory organizations such as the Financial Industry Regulatory Authority (FINRA), and state securities regulators to conduct examinations of broker-dealers to protect investors. Record retention is the primary means of monitoring compliance with banking and securities laws, including anti-fraud provisions and financial responsibility standards.

The length of time broker-dealers must keep records depends on the record type. For example, firms must maintain blotters containing all purchases and sales of securities for at least six years. They must keep copies of confirmations three years. Broker-dealers must retain business communications for three or six years, and during the first two years, records must be readily accessible.

When electronic media are used to retain records, the SEC requires the data to be in a non-rewriteable, non-erasable format (write once, read many – WORM). The system of record must document record sources, support an audit trail to validate the storage media recording process, serialize original and duplicate units of storage media, and time-date records for retention. The SEC states that compliant storage systems assign permanent, default retention periods for stored records without an allotted retention period. The system must also maintain records under legal requirements, such as a legal hold or subpoena, for the required periods.

FINRA regulates firms and professionals selling securities in the US and the US securities markets and enforces federal securities rules and laws, as well as the Municipal Securities Rulemaking Board controls. Among other regulations, FINRA Rule 4511 requires member firms to keep all records without a prescribed record-retention period for six years after an account is closed, and Rule 4513 requires regulated companies to keep customer complaints for four years.

Strengthened by the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank Act), the Commodity Futures Trading Commission (CFTC) oversees the $400 trillion-plus swaps market. The Dodd-Frank Act requires public reporting of outstanding and future swaps. Besides the swap reporting requirements, the Dodd-Frank Act requires each swap dealer and major swap participant to maintain trading records, including recorded communications such as emails, instant messaging and telephone recordings. Records must be kept by all swap parties during the swap and for five years following its termination. The files must be readily accessible throughout the life of the swap and for two years following its completion.

The Federal Deposit Insurance Corporation (FDIC) requires banks and other financial institutions to maintain customer identification programs to prevent money laundering. Under the Bank Secrecy Act (BSA), banks must keep a detailed history of checking accounts for at least five years after the information is obtained and retain records of large international transactions and fund transfers. The BSA also requires banks to keep for five years suspicious activity reports filed with the Financial Crimes Enforcement Network and original business records and other evidence of alleged fraud. Although banks are not required to keep separate systems of record for BSA requirements, the institutions can keep the information in any form, including electronic and microfilm, but it must be readily accessible upon request.

An electronic storage system that prevents overwriting, erasing or otherwise altering a record during a retention period through integrated hardware and software control codes can satisfy SEC and other recordkeeping rules for institutions regulated by the FINRA, CFTC and the FDIC. Such a system will also help meet regulations in other industries, such as healthcare, insurance and retail, and help manage global data privacy risks.

HEALTHCARE AND INSURANCE

State laws govern the retention of medical records, so the time frame varies from 7-10 or more years. However, the administrative simplification rules of the Health Insurance Portability and Accountability Act (HIPAA) require covered entities, such as physicians and hospitals that bill Medicare, to retain required documentation for six years from its effective date. HIPAA preempts state laws if the retention period is shorter than six years, but if longer, then state law determines the retention period.

The Centers for Medicare & Medicaid Services (CMS) require healthcare providers that submit cost reports to retain original forms for at least five years. CMS also needs Medicare managed care program providers to maintain records for 10 years. Medical providers and suppliers should keep medical files for each patient who is a Medicare beneficiary. According to Medicare, the data must be accurately written, accessible, and retained in a system using author identification and record maintenance that ensures the integrity and security of the documents.

State laws also regulate insurance carriers, but the Department of Labor Fiduciary Rule makes capturing new administrative data types a high priority for insurance carriers and distributors. Agents need to track and capture communications with clients to show that they acted in the best interest of investors when making retirement account recommendations. Otherwise, carriers and agents can be the subject of individual lawsuits for conflict of interest and imprudent investment advice. Automating data capture and storage is the easiest way to show compliance and ensure extensive evidence of a careful process in advising clients without regard to compensation.

RETAIL: PAYMENT CARD DATA SECURITY STANDARDS

The PCI DSS enhances cardholder data security and encourages the adoption of consistent data security measures. PCI DSS provides organizations with technical and operational requirements designed to protect cardholder data (CHD). The standards apply to all entities that store, process or transmit CHD and sensitive authentication data, including merchants, payment processors, card issuers and service providers. Security requirements apply to all system components in a cardholder data environment, which includes all processes and technology that store CHD.

The PCI DSS has 12 requirements, which include directives to protect stored CHD and track and monitor all access to it. CHD contains primary account numbers, cardholder names, card expiration dates and service codes. Sensitive authentication data includes full track data (magnetic-stripe data and its equivalent on a chip), the three- or four-digit values printed on the front or back of payment cards, and personal identification number blocks. After cardholder authentication, the PCI DSS permits entities to store primary account numbers, cardholder names, service codes and expiration dates, but primary account numbers must be rendered unreadable.

PCI DSS requires data minimization: keep stored CHD to a minimum by executing data retention and disposal policies and perform a quarterly manual or automatic process to identify and securely delete stored CHD that exceeds defined retention periods. Other conditions include classifying media so the sensitivity of the data can be determined, and retaining audit trail history for at least one year, with a minimum of three months immediately available (e.g., online, archived or restored from backup) for analysis.

Protection methods to comply with the PCI DSS include encryption, truncation, masking and hashing. Potential risk-mitigation methods should also be considered to protect stored data, including minimization – not storing CHD unless necessary and, even then, for only the period required by law or business need.

GLOBAL DATA PRIVACY LAWS

Laws that regulate privacy involve the application of fair information practices, or privacy by design, fostered by the Organization for Economic Cooperation and Development Guidelines on the Protection of Privacy and Trans-Border Flows of Personal Data (1980, amended in 2013). Privacy by design restricts the collection, use and dissemination of identifiable personal data, assures data reliability for its intended purpose, and prevents data from misuse while ensuring the security, accuracy, and validity of data collected and stored.

Country-specific laws incorporate privacy by design – e.g., the Argentine Data Protection Regulations (ADPR), the Japanese Act on the Protection of Personal Information (APPI), and the European Union General Data Protection Regulation (GDPR), which goes into effect in May 2018. The ADPR, APPI and GDPR require data collectors and processors to assess the risks of collecting personal data from customers, employees and partners and determine appropriate protections to store and process the data, as well as ensure its accuracy, accessibility, confidentiality and validity.

The GDPR applies to the processing of personal data from a business activity in the EU without regard to whether the processing occurs inside or outside the EU. The regulation gives EU residents more control of their data. Individual powers include the ability to prohibit data processing beyond its specified purpose for collection, the right to be forgotten, and the ability to withdraw consent to the collection and use of personal data. Controllers and processors must implement organizational and technical measures to ensure data security appropriate to the risks of the stored data. The measures include 'pseudonymizing' or encrypting the data to ensure confidentiality of the data and access to it in the event of a data breach.

Recordkeeping for data processing under the GDPR requires organizations with 250 or more employees to retain the full details of processing activities. The features include the purpose of processing, data subjects' consent to the processing, and the categories of personal data and the recipients to whom the personal data have been or will be disclosed. Other recordkeeping rules require documenting the time limits for erasure of different content types, and a record of the organizational and technical security measures applied to stored data.

RISK OF FINES, PENALITIES

Regulatory requirements add quantifiable risks to improperly collecting, processing and using personal data. The cost of failing to comply with recordkeeping and global privacy requirements and, by implication, associative technology controls and protections, is high. For example, failing to comply with the GDPR can cost up to €20m in fines, or up to 4% of the total annual revenue of the preceding financial year, and failure to comply with the APPI can result in criminal prosecution and sentencing. Besides regulatory fines and penalties, companies face the costs of defending litigation from individual and classes of data subjects, mandatory breach notifications to customers, pervasive government monitoring for prescribed periods, and a loss of customer confidence and revenue.

Penalties for violating the document preservation requirements include monetary penalties and can subject violators to criminal and civil enforcement actions. In 2016, FINRA *fined 12 firms* for $14.4m for deficiencies in preserving broker-dealer and customer records in a format that prevents alteration. In March 2014, FINRA *settled with Barclays Capital* for $3.75m, alleging that the firm violated, among other rules, SEC Rule 17a-4 for failing to preserve electronic records in a WORM format.

Businesses should adopt risk-based information governance principles to collect, retain and use personal data that digitally transforms their businesses to customer-centric operations. Organizations need methods and mechanisms to store personally identifiable information that complies with regulatory requirements to continue doing business.

Object Storage Helps Organizations Comply with Legal and Regulatory Requirements

The increasingly massive volume and variety of data stored in organizations that needs to meet regulatory requirements is pushing existing block and file-based storage services, such as network-attached storage (NAS) and storage-area network (SAN) systems, to their limits. These proprietary hardware and software platforms provide data protection and system management features that span arrays of hard disks, but the platforms have limited capabilities to identify and present necessary information about the content and nature of data collected to fulfill legal and regulatory requirements.

NAS and SAN file systems have a limited framework to locate files – they use a hierarchical system based on device, folder and subfolder, and simple check-box attributes to identify metadata, such as creation and last modification date – and to fulfill indexing and archiving requirements. Besides the lack of useful information about content, NAS and SAN file systems do not scale to extensive data applications and conditions that require billions of files over multiple locations – like the use cases for regulatory compliance.

Rather than identifying data as a location on disk, object storage identifies data as an individual object with a unique identifier, and separately maintains configurable metadata about the object to classify, index, manage and retain the data. Any data, including unstructured and multimedia files, can be captured, preserved, archived and accessed in an object store. Regulatory mandates continue to require organizations to keep increasingly massive datasets to prove compliance. Modern regulations are making the case for using object storage technology to store, manage and preserve organizational data more compelling.

Object storage features to identify, manage and control access to content include content security, media and interface efficiency, and compatibility.

Object storage security features include:

- Locking down data with erasure-code-based data protection to apply retention policies and secure objects from deletion until a retention period, legal hold, subpoena or other legal requirement has expired.
- Single copy with code-based erasure protection provides compliance-enabled data functionality without the need for replication or multiple copies.
- Applying default retention policies to stored content, and configuring variable, user-defined retention periods for specific content.
- Monitoring and documenting content access and use with audit logs.

The media and interface efficiency features of object storage include:

- Scaling storage across multiple nodes and geographic regions with global name-space capabilities.
- Supporting non-disruptive data movement when migrating to newer infrastructure.
- Reducing friction and minimizing errors with interfaces that can be directly programmed and controlled by applications.

Object storage compatibility features include:

- Standardizing access for compliance using modern user interfaces and application programming interfaces (APIs).
- Running object storage systems on standard servers and hardware-agnostic storage infrastructure.

Conclusion

Organizations should turn to object-based storage platforms to protect sensitive data and meet stringent legal and regulatory compliance requirements. Each object includes data, metadata and a globally unique identifier. The object storage model can provide adequate information and customizable controls to automate data-retention practices, dynamically scale over large, multi-node systems, and economically migrate objects between storage tiers. Unlike file and block storage, object storage has no limits to the number and size of objects in storage, and supports programmable application interfaces and data management capabilities such as data distribution and replication at the object level, self-healing, and erasure-code-based data protection.

CONTENT FURNISHED BY IBM

IBM Cloud Object Storage is a breakthrough platform that helps solve unstructured data challenges for companies worldwide. It is available as a service in the IBM Cloud or deployed as an integrated appliance or capacity based software on your hardware in your data center(s). Relied upon by some of the world's largest repositories, Cloud Object Storage turns storage challenges into a business advantage. It does this by reducing storage costs while reliably supporting both traditional and emerging cloud-born workloads for enterprise mobile, social, analytics and cognitive computing. It is designed to provide scalability, simplicity, security and storage economics for an overall lower total cost of ownership (TCO).

Serving as a scalable content or backup repository, active archive or as cloud-native storage, the IBM® Cloud Object Storage System™ is a private cloud for your on-premises environment. Customers can start small or in a configuration that matches their needs. In all cases the storage is always on and never goes down even as you upgrade, expand capacity or experience failures, including an entire site. This is a system that lasts as long as or longer than your data, which saves cost and valuable time for your staff.

Security and encryption is built into the overall design and process of the storage system including encrypting all data at rest (DAR). With IBM's compliance-enabled vaults, customers can lock down data with write-once read-many (WORM) functionality and meet regulatory compliance requirements like those found for SEC17a-4(f). Data can be kept safe and secure but still remains securely accessible via the standardized S3 API interface.