



NetApp White Paper

Architectural Considerations for Archive and Compliance Solutions

Chris Cummings, NetApp; Thomas Savage, NetApp
October 2008 | WP-7055-1008

DATA CENTER PROVEN SOLUTIONS

When it comes to archive and compliance, IT directors have a choice to make: implement independent archive and compliance architectures for individual applications as needed, or implement an archive and compliance architecture that dramatically simplifies traditional architectures and consolidates e-mail, file, database, ERP, and ECM data on a single platform. This paper investigates the architectural considerations for archive and compliance solutions and makes the case for a comprehensive and unified architecture for archive and compliance solutions.

TABLE OF CONTENTS

- 1 INTRODUCTION3**
- 2 RATIONALE FOR ARCHIVE AND COMPLIANCE3**
 - OPERATIONAL EFFICIENCY3
 - COMPLIANCE4
- 3 REQUIRED CAPABILITIES6**
 - MAPPING REQUIREMENTS TO ARCHIVE AND COMPLIANCE SOLUTION TYPES.....8
- 4 THE IDEAL ARCHITECTURE.....9**
 - TYPICAL APPROACH: ARCHIVE AND COMPLIANCE9
 - TYPICAL APPROACH: PRIMARY AND SECONDARY STORAGE9
 - SINGLE-PURPOSE SOLUTION INEFFICIENCIES9
 - COMPREHENSIVE, UNIFIED FRAMEWORK.....10
 - AN EXEMPLARY ARCHIVE AND COMPLIANCE SOLUTION ARCHITECTURE11
- 5 NETAPP ARCHIVE AND COMPLIANCE SOLUTION 12**
 - NETAPP STORAGE PLATFORM12
 - NETAPP EXTENDED CAPABILITIES12
 - NETAPP SOLUTION SUITES13
- 6 CONCLUSION 17**

1 INTRODUCTION

Today's organizations are facing exponential growth in the amount of data that they are required to store and access. It is estimated that over the next five years, digital archive capacity will grow nearly 1000% (or 10 times), from 9.2 exabytes (EB) to 90 EB of the capacity required today¹.

Along with this explosive increase in the amount of data being stored, the scrutiny surrounding the security and proper use of that data is also growing. A majority of this data is unstructured (upwards of 70%²), such as everyday office documents stored on network shares, project data and design files, which makes it harder to manage the data growth. Together, these trends are forcing data archive and compliance to the top of the list of business challenges for many of today's organizations.

In this paper, we will discuss the challenges of archiving mission-critical data and meeting compliance standards, propose an ideal architecture for archive and compliance solutions, and describe the NetApp solutions that fit into such an architecture.

We begin with a discussion of the challenges of archiving data. What was once exclusively an IT issue has now grown into a challenge for the entire organization. Operational efficiency, regulatory compliance, and long-term data retention requirements are forcing organizations to move from simple backup storage systems to more sophisticated archive systems. Customers are creating or overhauling their retention policies to ensure that their policies align with corporate goals and regulatory requirements in the face of legal discovery and increased regulatory oversight.

Then we'll look at the ideal architecture for archive and compliance solutions. How would you design architecture to simplify archive and compliance implementations? We'll take a close look at the requirements for the ideal architecture and make a case for disk-based archive and compliance solutions over other media-based solutions. Finally, we'll describe the NetApp approach, archive and compliance platform and solutions for e-mail, file, SAP, Oracle and ECM archive and compliance.

2 RATIONALE FOR ARCHIVE AND COMPLIANCE

When it comes to defining "archive" within the IT industry there is no clear agreement. For some people, it is simply a process of moving inactive data off their systems to enable their systems to maintain a certain level of performance. For others, archiving is adherence to strict retention policy that prescribes storing permanent copies of data, knowing where they are stored, and classification and search requirements for access to that data. But in today's business world, neither of these approaches is enough. You must take a more sophisticated approach to archiving—one that entails long-term retention, management of inactive data to meet compliance, data protection and security requirements and finally storage optimization.

No matter how you define it, there are three main reasons to archive: operational efficiency, compliance, and long-term information retention.

OPERATIONAL EFFICIENCY

Leading market research firms, such as The Enterprise Strategy Group, estimate that each year, the amount of data that medium and large-size businesses need to manage grows by 90 percent. As the amount of data grows, the cost of storing that data also increases, as does the complexity of managing all that data. We call this the unvirtuous cycle: increasing amounts of data require not just more primary storage, but more storage for disaster recovery and backup. Additional storage increases the complexity of storage and data management and lowers system performance and, in some cases, availability. The more complex the system, the more complex the backup, making it less likely that backups will be completed within the increasingly shrinking windows. Without backup, data is left unprotected.

¹ ESG Research Report 2007 Digital Archiving Survey

² *ibid*

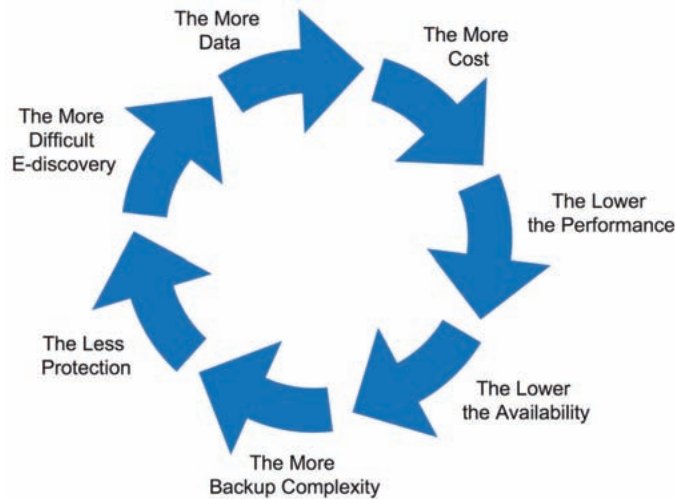


Figure 1) Unvirtuous cycle.

By relieving primary storage systems of data that does not have strict performance requirements or is inactive, archive solutions can significantly improve your operational efficiency in several ways: reducing the cost of primary storage, increasing application performance and in some cases availability, and relieving the burden on disaster recovery and backup systems and processes.

COMPLIANCE

Compliance is not a one-time exercise—it is an ongoing effort. New laws and regulations such as Sarbanes-Oxley and HIPAA require organizations to store more data for longer periods of time than ever before. But government regulations aren't the only forces driving compliance efforts. Internal requirements also demand that companies focus more of their time and resources on compliance. We split compliance into two categories: externally driven (governmental laws and regulations) and internally driven (internal business requirements). Each category faces its own challenges and failure risks.

EXTERNALLY DRIVEN COMPLIANCE

As more and more business operations are recorded and stored digitally, the thicket of laws and regulations governing businesses and data has become denser, and the consequences for failing to comply with these regulations have become more severe. For example, in the United States, failure to keep a customer's personal information secure may result in jail time and fines up to \$500,000. Some Sarbanes-Oxley legislation may bring penalties of up to 20 years in prison and fines of up to \$5 million for failures to comply with certain data retention policies. Management in Financial Instruments Directive (MiFID) requires retention of specific client data and communications for 3 years with penalties of \$\$\$ for failure to comply. Perhaps even more severe than fines and potential imprisonment are the loss of shareholder trust and the ensuing PR disaster of public naming and shaming.

The following table provides examples of regulations organizations around the world are required to meet.

Table 1) Global compliance regulation examples.

Regulation	Regulatory Agency	Industries Impacted	Requirements
SEC Rule 17a-4	Securities and Exchange Commission (SEC)	Financial services	Type and length of data retention Storage media requirements
EU Data Retention Directive	European Union	All industries with communication services	Mandatory data retention periods
21CFR (Code of Federal Regulations) Part 11	Food and Drug Administration	Pharmaceuticals and medical device manufacturers	Data security, integrity, auditability
Basel Capital Accord (or Basel II)	G-10 nations	All industries	Measure and report operational risk
U.S. Federal Rules of Civil Procedure (FRCP)	U.S. Federal Courts	All industries	Legal hold to preserve data unaltered Rapid production of unaltered data and records

INTERNALLY DRIVEN COMPLIANCE

In addition to government compliance regulations, organizations also have their own internal policies and procedures. In order to mitigate risk and control their own IT, organizations must establish their own internal compliance standards. Most internal compliance standards fall into one of three categories: records management, litigation support, and privacy and security.

Records are one of the most critical elements of any organization. Therefore, many organizations develop a set of policies governing the care and management of its records. These policies typically dictate the format in which the records are stored, the media on which they are stored, how long the data needs to be retained, and when and how the data should be expunged. While many organizations have adopted enterprise content management (ECM) applications that often help with records management, the vast majority of enterprise data remains outside these systems and therefore remains unmanaged. You may have come across this issue when trying to address request for enforcement of a corporate policy on managing intellectual property or search for data used in an human resource action.

While businesses must be compliant with regulations, it is also imperative that they be prepared for legal actions, including litigation. One of the greatest game-changing regulations to impact organizations across the board centers around legal discovery and production of evidence relevant to litigation. In the US, the Federal Rules of Civil Procedure (FRCP) have forced companies to rethink their need and strategy for archiving data. We see many other countries starting to adopt similar rules in their judicial systems or formal regulations to manage production and retention of electronic stored information for legal proceedings.

Archives are viewed as a key component of an eDiscovery solution. Preparing the vast information stores within a company for possible litigation can be daunting. But should a dispute arise, having archival information organized and accessible may preempt the spiraling costs of electronic disclosure and may help organizations avoid severe financial penalties for not providing requested records within required time periods. In 2006, a leading financial services firm was hit with a \$10 million penalty for failing to properly maintain records that pertained to a legal case in which it was involved. In another case, a company was fined \$15 million for failing to produce e-mail records on time. In a well-publicized study by duPont, their legal counsel found that over a 4 year period, the company would have saved, \$12 million with proper use of an archive to enforce a reasonable retention policy.³

Security and privacy are other types of internal compliance policy that organizations must adhere to. There are two basic types of security: people security and data security. People security helps keep unauthorized users from accessing systems and data. It requires authentication to prove the person trying to access the

³ James Michalowicz, Manager, Legal Services, DuPont Corporation, presented at Cohasset's Managing Electronic Records Conference, September, 2002

data is who he or she claims to be, authorization to assure that the person requesting the data is authorized to do so, and auditing to be absolutely sure that the authentication and authorization methods are working. Data security technologies such as firewalls, intrusion prevention systems, and virtual private networks (VPNs) work in combination with people-security methods and are most effective in protecting against external attacks, such as phishing. However, they are not very effective in protecting data from internal attacks, an increasingly common type of computer crime. Encryption, the last wall of privacy protection within an enterprise, provides effective means of protection against breaches as well as loss of media, such as tape. A 2005 survey by the U.S. Federal Bureau of Investigation and the Computer Security Institute found that 56 percent of organizations that responded to the survey had experienced internal security breaches that year.

LONG-TERM DATA RETENTION

In the past, archives had short shelf lives. For instance, at one time, businesses working with the German government were required to retain all e-mail communications for seven years. Now we are seeing retention periods of 100 years or more. For example, U.S. regulations require all medical records to be retained for 30 years after a person's death. And with today's life expectancy rates, this could mean retaining records for well over 100 years. Another good example would be a construction engineering firm with an internal policy to keep all project documentation for 100 years. Their buildings have a life expectancy of at least 100 years, and in order to control legal risk, they retain projects records for at least as long as the buildings are standing. A similar requirement in the auto industry requires long-term retention of design documents. As the amount of data to be analyzed and managed continues its rapid growth, this type of data management will become even more complex and costly to maintain.

3 REQUIRED CAPABILITIES

Operational efficiency, compliance, and long-term retention each has a set of required capabilities that must be met by the archive and compliance solution.

Access controls. Many regulatory compliance mandates require organizations to have controls in place to keep private data private. Access controls are used to govern user access and permissions to data and the network itself, helping to ensure data security.

Backup options. Historically, organizations have used tape for backup and more than likely will continue to do so. But now, lower-cost disks are opening up more, and better, options that eliminate some of the inherent pitfalls of tape—its fragile nature, ability to easily get lost or stolen, slow backup times and pitiful slow recovery times.

Disaster recovery. Disaster recovery planning is not just for primary data and applications. Even though the data is held in an application considered secondary or tertiary to hourly business operations, the data held in an archive did not lose its value or relevance as a business record or information when it moved from the primary application into an archive. It didn't change. Under some risk management approaches, protection of this information from disaster is just as critical (if not more critical) than when it was held in the primary application.

Choice of network storage protocol. There are two classes of network storage approaches: storage area networks (SANs) and network-attached storage (NAS), each with two different types of protocols. SANs are typically relied upon for highest performance. Fibre Channel SAN has been the standard protocol choice, but increasingly, IP SAN (iSCSI) is being used in enterprise deployments. NAS has primarily been used for file sharing for UNIX® file system (NFS) or Windows® file system (CIFS) protocols, but over the years gained presence in applications enterprise-wide. Storage options play two critical roles. First, if you consider that basic architecture of an archive application, it does not benefit from monolithic storage. To optimize the archive infrastructure, you must support storage requirements for the database, index and content/image archive resident in every archive application. Higher performing components (database and index) call for FC SAN or iSCSI and content/image archives call for "cheap and deep" benefits of NAS with SATA. Second, all protocols should be available to utilize in archive and compliance applications. Standard network protocols are essential especially for long-term retention so that companies are not tied to proprietary forms of communication to their archival and compliance systems.

Data classification and search. Many (most, in fact) enterprises lack a clear understanding of their data profile. Data indexing and classification can help such enterprises understand and manage their data. Data indexing is needed to efficiently search for information and fulfill compliance and litigation-driven discovery requests. Data classification is needed to effectively define data management.

Data integrity assurance. Data integrity assurance ensures that data is maintained identically throughout all operations, including transfer, storage, and retrieval. Assuring data integrity is an important factor in meeting many compliance regulations. This is particularly true for enterprises in the financial industry, where many SEC regulations require that data from all customer transactions be stored unaltered from its original state. Checksum algorithms, which execute a redundancy check on data, are a common form of data integrity assurance.

Access to data in the event of disk failure. Organizations need continuous access to data—even in the event of a disk failure. Redundant arrays of independent disks (RAID) and redundant arrays of independent nodes (RAIN) mirror data, ensuring continuous data access, even if a disk fails.

Deduplication. Deduplication is the process of removing unwanted duplicate copies of data. This helps reduce the amount of data to be stored, reducing the overall cost of storing and managing the data. Deduplication can be split into two categories: file-level and block-level. File level ensures that a file is stored only once in a defined set of storage and exact duplicates are removed. Block-level goes one step further by comparing the blocks of data that are written on the disk and eliminating redundant disks. Two similar but different 20-slide PowerPoints may have all information the exact same except for a few words. Block-level Deduplication would reduce the storage required down to nearly that of the single 20-slide PowerPoint. With the amount of data organizations are required to store, storage efficiency is an important factor in the overall cost of an archive system.

Storage and access in native data format. Many compliance regulations require organizations to store data for long periods of time—in some cases, a lifetime. When this data needs to be accessed (for example, for litigation support), you need a fast and easy way to locate and retrieve the information you need in its native format. Maintaining native format is critical because you have no idea what application(s) will be used 100 years from now (even 25 years from now) to access this information. Saving data in its native format will dramatically improve your ability to access the data in the future.

Ease of application integration. Application integration for archive and compliance has two key parts. First, there needs to be a means of policy-based migration from the primary tier to the secondary tier. Second, your users need the ability to quickly and easily locate and access information. Transparency is the most common mechanism for enabling your users to retrieve their data. For instance, with Exchange, users should be able to access archived e-mail files directly from Outlook, eliminating the need for drive mapping and network sharing. Other approaches address the challenges of saving stubs for decades and provide processes to enable users to readily and intuitively locate data using common mechanisms such as parallel directories and search

Seamless migration. Migration is the act of moving data from one storage tier to another to free up space on primary disks. With seamless migration, the data moves smoothly from tier to tier without affecting system performance. Further movement of data across dissimilar platforms is important, especially for large organizations with many remote offices.

Audit logging. Audit logging is critical to any organization's regulation compliance strategy. Audit log data can provide a complete record of access, activity, and configuration changes for applications, servers, and network devices. It can also provide auditors with the information required to validate security policy enforcement and proper segregation of duties.

Performance and Agility. Performance for an archive is not limited to simple rate of ingestion, but other key services that help you realize the benefits of an archive. As archives have more data pushed into them (for example growth in number and size of emails per person) and more content types, they must easily scale for volume and performance. But, also support services such as classification, search and export of data for audits and discovery. Agility is critical to meeting business demands of new retention policies and regulatory changes.

Selective encryption. Selective encryption makes it easy to quickly secure data. Selectively encrypting data at any granularity (at a record/file level or at an entire volume or storage system level) helps reduce overall system complexity; reduces redundancy; and provides a secure, efficient way to store sensitive data.

Selective data locking. The ability to lock data to protect it from being altered is very important to organizations that are required to produce unaltered data to meet regulatory compliance mandates. Furthermore, you can improve risk management profile by locking down static data to prevent accidental or malicious alteration or deletion. Ideally, organizations need to be able to selectively lock data at any granularity – one record, a group of records, or all records generated by an application. In combination with selective encryption, even individuals with access privileges could not alter or delete data.

Storage efficiency. You can easily be overwhelmed by the thought of stories terabytes (even petabytes) of data for years, decades, even “indefinitely”. Somehow you must store this ever growing pool of data as it swells from a pond to an ocean. Storage efficiency technologies such as thin provisioning, deduplication and efficient snapshots to name a few ensure that you can keep your acquisition, power, cooling and real estate costs down.

MAPPING REQUIREMENTS TO ARCHIVE AND COMPLIANCE SOLUTION TYPES

Most organizations view archive and compliance as two separate initiatives and deploy individual solutions for each. As you can see from the table below, many of the requirements for archive and compliance solutions overlap. Each time a single-purpose solution is deployed, many of the same capabilities are repeated. It makes much better sense to consider them together in a single architecture with a single solution that offers all of these capabilities together.

Table 2) Archive and compliance required capabilities.

Requirement	Operational Efficiency	Compliance	Long-Term Retention
Access Controls	■	■	■
Access to Data in Event of Disk Failure	■	■	■
Backup Options	■	■	■
Disaster Recovery	■	■	■
Choice of Network Storage Protocol	■	■	■
Data Classification and Search	■	■	■
Data Integrity Assurance	■	■	■
Storage Efficiency	■	■	■
Storage and Access in Native Data Format	■	■	■
Ease of Application Integration	■	■	■
Seamless Migration	■		■
Audit Logging		■	■
Performance and Agility	■	■	■
Selective encryption		■	■
Selective data locking		■	■

4 THE IDEAL ARCHITECTURE

As the amount of data and the number of regulatory compliance mandates grow, so will the complexity of managing this data, maintaining operational efficiency, and meeting regulatory compliance requirements. Adding new systems to address each storage or compliance need only adds to the complexity of data management. The ideal architecture would address and reduce this complexity; it would be a single, comprehensive, scalable, and unified architecture that leverages the same platform and infrastructure for both archive and compliance initiatives. Taking it one step further, the ideal architecture would also combine primary and secondary storage on the same infrastructure, enabling you to aggregate all archive and compliance initiatives on the same architecture as primary and secondary storage. Furthermore, you can simplify to complexity of the archive application itself with a unified architecture to address the architecture complexity of the application.

TYPICAL APPROACH: ARCHIVE AND COMPLIANCE

Many organizations are faced with the challenge of meeting numerous rules and regulations. Each time a new archive policy or compliance regulation is released, organizations must create new teams to evaluate and address it. This often causes a duplication of effort and sometimes results in the creation of conflicting policies.

TYPICAL APPROACH: PRIMARY AND SECONDARY STORAGE

Like archive and compliance, many organizations also implement separate primary and secondary storage solutions, which can introduce a great deal of complexity into the infrastructure. It is often much more difficult to offer application integration and seamless migration when primary and secondary storage are separated.

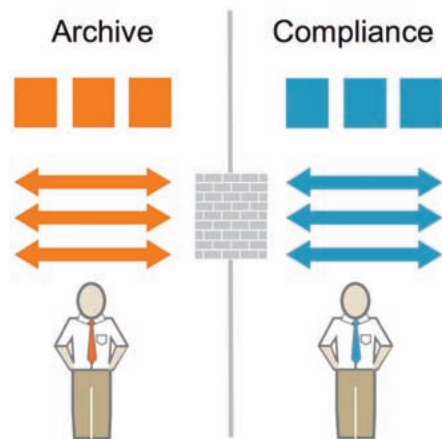


Figure 2) Addressing archive and compliance regulations individually.

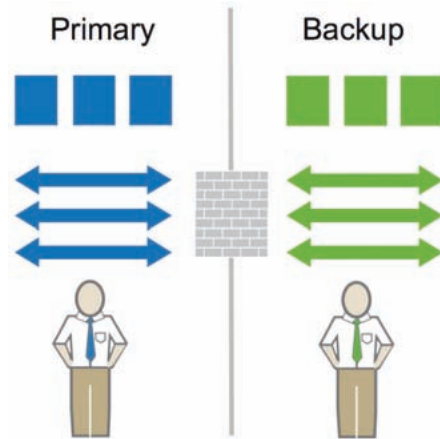


Figure 3) Addressing primary and secondary storage individually.

SINGLE-PURPOSE SOLUTION INEFFICIENCIES

A single-purpose approach is inefficient for many reasons. First, a single-purpose solution is sometimes unable to act seamlessly with the rest of the infrastructure, especially when it comes to data migration and application integration. Single-purpose solutions may also lack the ability to maintain a cohesive set of data. Multiple copies of the same data may reside in multiple places, making it more difficult to maintain and manage. In addition, single-purpose solutions are often inflexible, unable to easily adapt to process changes. It is also typically more expensive to deploy a single-purpose solution because they typically require additional investments in hardware, software, time, and training that is applicable no where else.

COMPREHENSIVE, UNIFIED FRAMEWORK

Ideally, organizations should implement a comprehensive, unified framework to address the requirements of multiple archive and compliance regulations together. This approach could simplify the process of meeting regulatory compliance and internal policy regulations by enabling external compliance regulations and internal policies to be mapped to a framework and addressed as a whole. New regulations and policies could easily be mapped to the existing framework, so only changing requirements would need to be addressed at any given time. To top it off, a combined archive and compliance solution would leverage the same infrastructure for primary and secondary storage, lowering total cost of ownership and delivering consistent performance. This type of framework would offer the flexibility to choose the best deployment model and utilize existing investments, processes, and skill sets for a high performance, cost-effective archive and compliance solution that serves short-term and long-term needs. Finally, these approaches treat the archive application as a monolithic data pump. But archive applications incorporate databases and indexes in addition to the content/image archive and require a more tailored approach from the archive storage solution.

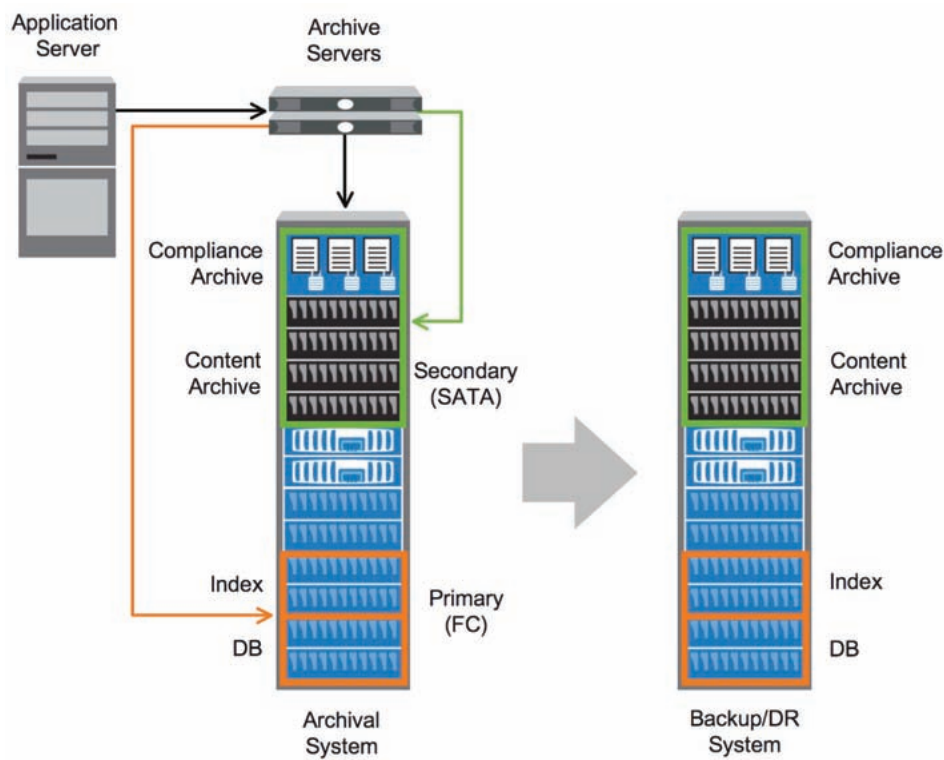


Figure 4) Unified architecture with logical single arrays.

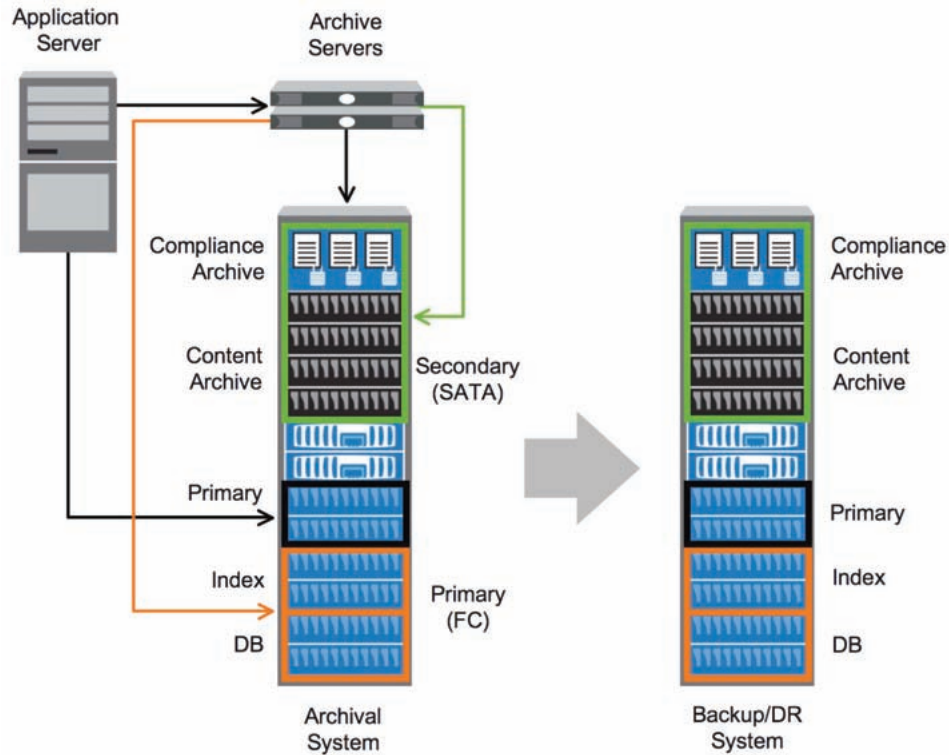


Figure 5) Unified architecture with single physical array.

AN EXEMPLARY ARCHIVE AND COMPLIANCE SOLUTION ARCHITECTURE

An exemplary archive and compliance solution architecture would be built with the future in mind. Offering a choice of SANs or NAS storage means organizations could easily add less-expensive SATA drives to meet the growing need for additional secondary storage. Combining archive and compliance capabilities into one unified framework means organizations could more easily and cost-effectively address new government or internal compliance policies as they arise. But not only would archive and compliance capabilities be combined on the same platform and infrastructure, primary and secondary storage would also be combined to provide more-efficient application integration, dramatically simplify your archive architecture, radically improve archive data protection, and enable more seamless data migration. Built-in compliance migration functionality would enable the creation of an audit trail to track data moved from one tier to another and help ensure that the data is retained for the amount of time necessary to meet compliance regulations.

Finally, the ideal unified framework would be built on a disk-based archive architecture to facilitate the key capabilities that are critical to the architecture: data classification, data security, and data discovery. Along with the platform and infrastructure components of this ideal architecture, other issues would also need to be addressed. For example, organizations are concerned about the amount of power used to maintain long-term compliance archives. The architecture would need to provide functionality, such as the ability to deduplicate and compress data, to reduce power consumption. Organizations are also concerned about how much physical space will be required to store its data for long periods of time, requiring the architecture to provide the ability to store large amounts of data in a small amount of space.

5 NETAPP ARCHIVE AND COMPLIANCE SOLUTION

With its unified platform, extended capabilities, and best-of-breed solution suites, the NetApp archive and compliance solution brings the ideal architecture to life.

NETAPP STORAGE PLATFORM

Other vendors require that you buy separate storage technologies and hardware platforms if you need primary (typically FC SAN) and secondary storage (typically NAS). Once you do this, you must make a choice: do you configure your secondary storage for archive, or do you configure it for compliance? If you require both, you will need to buy multiple systems.

What sets NetApp apart is the ability for all its storage systems to provide simultaneous support for all protocols—Fibre Channel SAN, IP SAN, NFS, and CIFS—in combination with Fibre Channel, SAS or SATA disk with the ability to selectively utilize additional capabilities, including data permanence and encryption. The result is that organizations can standardize on a single architecture for all archive and compliance initiatives that can be tailored to meet the needs for scale, cost, and performance. Fundamentally, the same architecture that delivers very high performance and reliability for enterprise-class applications, including Microsoft® Exchange, Oracle and SAP, can be utilized to deliver low-cost, massively scalable archival and compliance initiatives for the same applications. Having a single set of software, tools and processes simplifies the complex world of enterprise data management. A single process for activities such as installation, provisioning, mirroring, backup, and upgrading is used throughout the entire NetApp storage product line, lowering administrative costs and making it far easier to deploy new capabilities across all tiers of storage. Unifying storage and data management software and processes significantly reduces the complexity of data ownership, enables the ability to adapt to changing business conditions without interruption, and results in lower TCO.

NETAPP EXTENDED CAPABILITIES

NetApp offers a single, complete product infrastructure that seamlessly integrates with key enterprise applications that are driving the current generation of mission-critical enterprise data for archive and compliance. With NetApp, archive and compliance functions such as data classification and migration, data permanence, data security, and data discovery all integrate with the core platform, making NetApp a truly comprehensive, unified solution with extended capabilities that enable organizations to lock, encrypt, classify, and search data. The ability to selectively apply these capabilities to any data set provides organizations with the ability to meet any particular data management requirement.




Figure 6) NetApp archive and compliance solution architecture.

NETAPP SOLUTION SUITES

NetApp offers a portfolio of open-standard e-mail, file, SAP, and ECM solutions for archive and compliance, each of which is becoming increasingly critical to the enterprise. NetApp solutions are fully integrated, tested, and validated with leading archive and compliance partners, including EMC Documentum, FileNet, OpenText, SAP, Oracle, Symantec, CommVault and Autonomy.

Table 3) NetApp E-mail archival solution.

<p>Customer Needs</p>	<p>Increase performance and availability of e-mail systems Improve e-mail system protection and backup Increase control and visibility over e-mail and attachments Provide more-efficient, less-costly e-discovery capabilities Meet regulatory compliance mandates Increase storage efficiency of archive</p>
<p>Solution Components</p>	<p>Recommended Core storage to support e-mail archive: FC or SAS for archive indexes and databases; SATA for messages and attachments Integration with e-mail archive application to manage policies and migration processes for mailbox management and journaling, including Symantec® Enterprise Vault® for Exchange, Autonomy EAS for Lotus, as well as other innovative vendors, such as CommVault, Quest and Mimosa</p> <p>Optional SnapLock® for data retention and WORM requirements NetApp DataFort for security and encryption</p>
<p>Value Statement</p>	<p>NetApp dramatically reduces the cost and complexity of e-mail archival. The unique NetApp single-box e-mail archive storage solution combines high-performance SAN for e-mail archive indexes and databases and cost-effective SATA for message and attachment storage. NetApp FlexVol® technology makes things easy to manage and grow. RAID-DP® enables data availability, even with dual-drive failures in the same RAID group. Cascading Snapshot copies and SnapMirror improve recovery point and recovery time objectives. Extended NetApp capabilities enable organizations to lock, encrypt, classify, and search e-mail for regulatory and litigation support requirements. Finally, NetApp systems exhibit stellar performance, either on day one or in two years after being loaded up with data.</p>

Table 4) NetApp file archival solution.

<p>Customer Needs</p>	<p>Reduce cost of managing unstructured data, such as office documents, contracts, product information, project plans, and proposals</p> <p>Increase visibility into data, including the ability to distinguish between important and unimportant files and enable better accounting and service levels</p> <p>Increase data security of confidential customer, employee, and corporate data stores</p> <p>Provide more-efficient, less-costly e-discovery capabilities</p> <p>Meet regulatory and internal mandates for information security and retention</p>
<p>Solution Components</p>	<p>Recommended</p> <p>Secondary storage for archived information, with complete protection via RAID-DP and replication, utilizing NetApp SnapMirror®</p> <p>NetApp IS1200 for data classification, migration, and discovery</p> <p>Integration with key archival partners for data classification and migration</p> <ul style="list-style-type: none"> • Symantec Enterprise Vault • CommVault Simpana • Arkivio <hr/> <p>Optional</p> <p>SnapLock for complete data retention</p> <p>Decru DataFort for complete data security and encryption</p> <p>NetApp IS1200 for data classification, migration, and discovery</p> <p>Integrated archive application partners</p> <p>Data Assessment Service</p>
<p>Value Statement</p>	<p>NetApp simplifies file archival and enables companies to move to a tiered storage model to lower the cost of unstructured content storage. NetApp file archive solutions enable companies to profile their information, determine archive policies, and migrate applicable data to secondary storage, either through the IS1200 or in concert with application partners. NetApp further enables companies to selectively lock and encrypt information to meet key external and internal mandates.</p>

Table 5) NetApp SAP archival solution.

Customer Needs	<p>Increase SAP performance and scalability Increase user productivity Reduce backup times Access archived data and supporting documents Meet regulatory compliance mandates Move off optical storage media</p>
Solution Components	<p>Recommended Secondary (ATA disk) storage for archives</p>
	<p>Optional SnapLock to enforce the retention periods NetApp DataFort to encrypt sensitive data before sending off-site Partner SAP archive application FileNet: Application Connector for SAP (ACSAP) Open Text (IXOS): Livelink ECM Data Archiving for SAP PBS ContentLink</p>
Value Statement	<p>NetApp offers a high-performance, unified storage solution that can archive both SAP databases and content. NetApp RAID-DP offers the best protection against disk failure without sacrificing performance or usable capacity. And NetApp offers the SnapLock option to lock down archived data. All of these advantages mean customers will enjoy faster SAP performance and have a more-efficient archive solution that is easier to deploy and easier to manage and offers a lower TCO.</p>

Table 6) NetApp Oracle archival solution.

Customer Needs	<p>Reduce storage acquisition and management costs Increase user productivity Reduce backup times Meet regulatory compliance mandates and records management requirements Increase database performance and scalability</p>
Solution Components	<p>Recommended Primary (FC SAN) storage for high-performance databases Secondary (ATA disk) storage for archives Oracle ILM Assistant Oracle partitions software SnapLock to enforce the retention periods</p>
	<p>Optional PLSQL Plugin for extensibility to customer applications SnapManager for Oracle SnapMirror for disaster recovery</p>
Value Statement	<p>NetApp offers a high-performance, unified storage solution that can create a tiered storage environment for any application built on Oracle database. NetApp RAID-DP offers the best protection against disk failure without sacrificing performance or usable capacity. And NetApp offers the SnapLock option to lock down archived data. All of these advantages mean customers will enjoy faster Oracle database performance and have a more-efficient archive solution that is easier to deploy and easier to manage and offers a lower TCO.</p>

Table 7) NetApp ECM solution.

Customer Needs	<p>Fast, reliable, affordable storage for terabytes of content/images and ECM indices/databases</p> <p>Replace old, slow, unreliable optical storage</p> <p>Prevent tampering or early deletion at the storage level to enhance compliance</p> <p>Reduce backup and replication complexities</p> <p>Protect sensitive data</p>
Solution Components	<p>Recommended</p> <p>Secondary (ATA disk) and RAID-DP for content/documents</p> <p>Primary (FC disk) and SAN/iSCSI for ECM app database</p>
	<p>Optional</p> <p>SnapLock to enforce retention periods, fully integrated with FileNet and EMC Documentum</p> <p>Decru to encrypt sensitive data</p>
Value Statement	<p>NetApp offers a high-performance, unified storage solution that can hold both ECM documents and the index databases. NetApp RAID-DP offers the best protection against disk failure without sacrificing performance or usable capacity, while NetApp SnapLock helps lock down archived data to prevent tampering or early deletion, without requiring deployment of a separate storage architecture for compliance. All of these advantages give customers a more-scalable, more-efficient ECM storage solution that is easier to deploy and easier to manage and offers lower TCO.</p>

6 CONCLUSION

Increasing amounts of data and more-stringent government regulations have made archive and compliance top business challenges for many organizations. Many times, organizations treat archive and compliance as two separate issues, which they address with different solutions, creating a less-flexible, more-expensive, and more-complex environment. The ideal archive and compliance solution would include a comprehensive, unified framework that includes an infrastructure of products and solutions that enable customers to deliver on multiple archive and compliance initiatives while running on the same systems. This unified approach would leverage the same infrastructure for archive and compliance initiatives as it would for primary storage applications, reducing cost, simplifying management, and delivering high performance. NetApp archive and compliance solutions are built on a unified architecture, providing a single-box solution to a complicated challenge. NetApp tiered storage and storage services are the foundation of NetApp products. Software solutions from NetApp and our partners support overall enterprise archive applications and enable enterprise-wide discovery of information across all archive initiatives. All NetApp solutions are created with flexibility, integration, and the future in mind. As a leading innovator in networked storage systems and storage and data management, NetApp is dedicated to providing a simpler, unified approach to data archive and compliance. We want to help organizations quickly and cost-effectively reduce the complexity of meeting compliance regulations with solutions that work for today and easily adapt to the future.

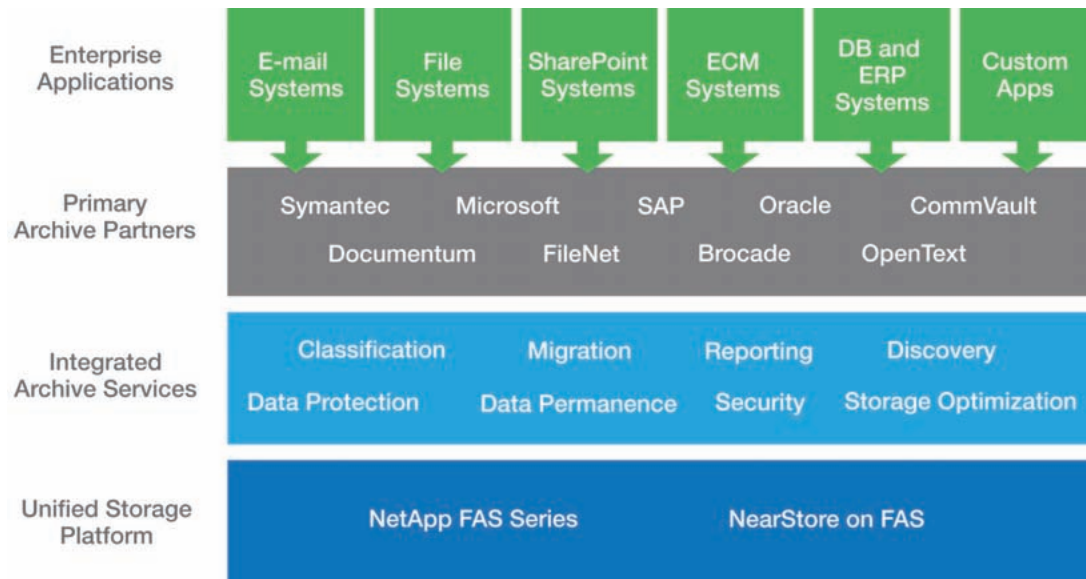


Figure 7) NetApp unified archive and compliance platform and solutions.