

White Paper

FUJITSU Storage ETERNUS CS8000 and Commvault Software

With the combination of FUJITSU Storage ETERNUS CS8000 and Commvault Software, Fujitsu provides a comprehensive solution that helps IT departments solve data protection challenges by combining backup and archiving as well as the related information management in an efficient and cost-effective way.



Content	
1. Introduction	2
2. Challenges in Modern Data Protection	2
3. Components of a Unified Data Protection Environment	2
3.1 FUJITSU Storage ETERNUS CS8000	4
3.2 Commvault Software	5
4. Use Cases	6
4.1 Backup in Open Systems Environments	6
4.1.1 Flexible, Scalable and Streamlined Workflows	6
4.1.2 Protecting Data in Virtualized Environments	7
4.1.3 Enhanced Reliability	8
4.2 Archiving and second-tier file storage	9
4.2.1 Improving Storage Tiering for Increased Efficiency	9
4.2.2 Archive Protection	10
4.3 Consolidation of data protection environments	10
4.3.1 Consolidation, Simplification and Optimization	11
4.3.2 Implementing Modern Information Management	11
5. Conclusion	13
6. Index of Figures	13



1. Introduction

According to analysts, at least 30 percent of all businesses plan to implement new backup solutions in the medium term. In fact, the number of those willing to trade their old solutions for a new and functional one might be even higher. This is because most companies still run multiple separate solutions for backup, recovery, archiving as well as the related data and storage resource management tasks in parallel. Unfortunately, these point solutions are no longer suited for today's usage scenarios dealing with data explosion, higher availability levels or stricter statutory regulations and compliance standards. Worse, applying independent hardware systems has induced and still induces considerable manual integration efforts. Furthermore, due to manual management efforts, errors can never be ruled out. In short, although operating costs keep growing companies still run the risk of storage, server, and application failures.

Therefore, the vast majority of IT managers are looking for contemporary, flexible solutions for backup, archiving and the related data management that combine as many features as possible and yet can be efficiently administered. The new platform should be technically up to date, help reduce operating costs and be easily and flexibly adaptable to future requirements. This ensures an early ROI as well as protection of necessary investments. Against this backdrop, Fujitsu and Commvault have expanded their strategic and technological alliance and developed a set of integrated data protection and data management solutions in which sophisticated data protection appliances and advanced software functions merge into a consistent whole.

The [White Paper: Fujitsu Storage ETERNUS and Commvault Software](#) focuses on integrated solutions for various backup scenarios where Commvault's Software and Data Platform is used together with Fujitsu storage ETERNUS CS appliances, ETERNUS AF all-flash and ETERNUS DX hybrid series, and/or ETERNUS LT tape storage systems.

This white paper focuses on backup and archiving scenarios where Commvault Software is used together with Fujitsu Storage ETERNUS CS8000, the datacenter storage solution for backup and archive automation.




2. Challenges in Modern Data Protection

The growing importance of IT results in a double dilemma for both business departments and data centers. On the one hand, they are faced with exploding amounts of information and tightening SLA requirements, on the other hand with shrinking budgets. With regard to backup, recovery and archiving, this means that more data have to be processed simultaneously while errors must be eliminated – all without additional efforts and expenses. At the same time, business-critical data must remain permanently available without backend processes slowing down access, and organizations need to make sure that centralized backups and archiving encompass all parts of their IT infrastructure, that way prohibiting the loss of vital information. Trends like virtualization, digitalization, mobile and cloud computing have further sharpened this predicament.

However, virtualization and cloud computing have some very positive effects. Probably the most important is that after years of neglect, enterprises once again realize how much they depend on a functioning storage infrastructure and efficient data management. Consequently, they are also more willing to tackle some or all of the serious challenges that IT departments are faced with, namely:

- Consolidation, simplification and optimization of storage environments
- Establishing flexible, scalable and streamlined workflows
- Protecting data in virtualized environments
- Improving storage tiering for increased efficiency
- Archive protection and protection against cyberattacks
- Enhancing reliability
- Implementing modern information and IT service management

In this white paper, we will discuss use cases to explain how ETERNUS CS8000 and Commvault Software together will solve these pain points.



3. Components of a Unified Data Protection Environment

In a typical productive environment, many critical data, the lifeblood of a business, can be stored anywhere. These data are backed up on a regular basis, using a backup application and a hardware grid consisting of several nodes (media servers) and a master console. The main advantage of this architecture is its flexibility and scalability: in case you need to back up more data or accelerate the process, you simply add and/or reallocate media servers and bandwidth.

Unfortunately, the same flexibility and scalability do not exist at the backend, namely within the backup storage infrastructure. Media servers send their data to various hardware repositories, depending on importance, frequency of use, etc. These repositories are usually encapsulated in disk appliances or tape libraries that offer only limited scalability. This may work fine so long as only one running backup storage system is sufficient. However, the challenge starts as soon as more than one of these systems is involved. Unlike the media servers and their master server, the backup storage systems are isolated from each other

and can merely serve as data silos with limited capabilities. This lack of flexibility turns moving data between repositories – which is necessary to free up storage space – into cumbersome, time-consuming manual work, leading to increased complexity and costs, quite the opposite of what an efficient organization wants. Separated from the backup storage systems, there are the archive storage systems. As if all that wasn't bad enough, a high percent of all data residing on expensive primary (online) storage systems is "inactive," meaning that it hasn't been touched in weeks or even months. To call this "inefficient" does not quite hit the nail on the head; in fact, it is a waste of valuable disk capacity and money. Clearly, offloading inactive files to another storage tier could help to drastically cut daily, weekly and monthly data transfers of backup copies and result in faster, more efficient backups. Again, now this would often require lots of costly manual work.

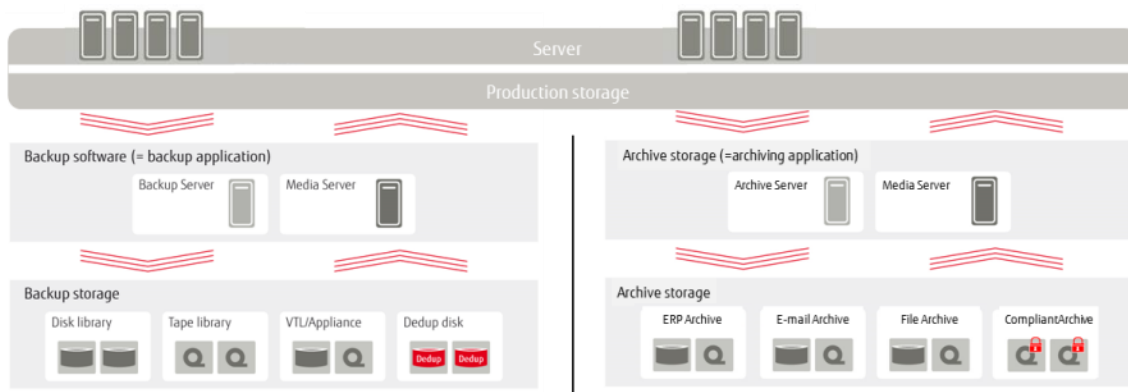


Figure 1: Typical data protection environment

Figure 1 shows a simplified model of the data protection environment typically found in today's enterprises: Company X runs a selection of mainframe, UNIX, and x86 servers at its main data center. Each of these architectures requires its own backup solution, with separate target systems (tape libraries or disk arrays), separate software, and separate networking equipment. The same setup is repeated for archiving.

The potential for improvement is obvious, particularly in mid-range and enterprise environments where the amount of data that must be stored, backed up, and archived is in the tera- to petabyte range, or where availability requirements are very high. Here, every limitation hits twice as hard, and manual administration efforts are (usually) not an option. Consequently, IT departments and data centers need a platform that provides best-in-class automation and centralization capabilities.

Fujitsu and Commvault have approached the task from different angles, resulting in a unique combined solution. By combining ETERNUS CS8000 – the datacenter storage solution for backup and archive automation from Fujitsu - with the backup and archive software from Commvault data protection environments will benefit from a solution that meets these complex recurring challenges in a cost-efficient way. The combined solution replaces outdated and ultimately dysfunctional methods. The idea is to consolidate data protection environments and to simplify backup, recovery and archiving processes as much as possible. Primarily, backup is about recovery and the risk of losing – or worse: not even capturing – important data has to be eliminated; costs for data protection and related data management need to be reduced. Ultimately, the goal is to create a system that aligns with customers' business needs.



3.1 FUJITSU Storage ETERNUS CS8000

The ETERNUS CS8000 is a datacenter storage solution for backup and archive automation with superior functionality that already solves many problems on its own. Unlike several competing products, it does not merely consist of simple pieces of hardware with firmware attached; instead, each customer gets an intelligent solution tailored to his or her individual business needs.

ETERNUS CS8000 provides Virtual Tape Library (VTL) interfaces for storing block-based tape volumes and a Network Attached Storage (NAS) interface for storing file-based data structures. The ETERNUS CS VTL subsystem is designed to store and retrieve sequential block-based tape volume data within a Storage Area Network (SAN). To the backup software, it appears exactly like one or more physical tape libraries. The ETERNUS CS NAS subsystem (Virtual Network Storage) provides within a TCP/IP network a NAS based file data archive with a virtual extremely high capacity.

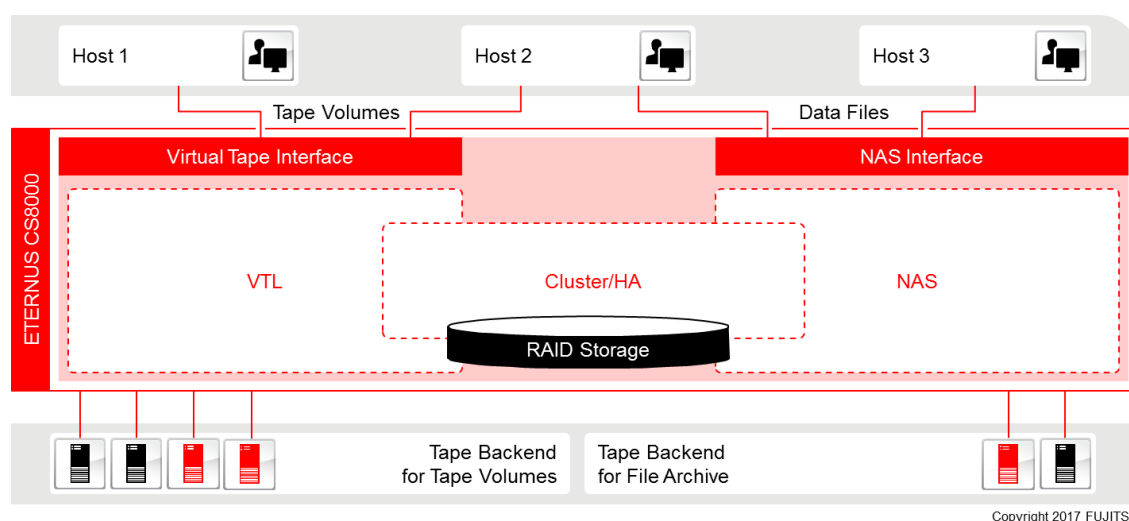


Figure 2: Fujitsu Storage ETERNUS CS8000

Within the ETERNUS CS8000 solution, the ETERNUS CS VTL and the ETERNUS CS NAS operate as two autonomous and distinct subsystems that rely on one single, unified, highly available cluster architecture (see Figure 2). However, ETERNUS CS8000 can also be modified to run only one subsystem – VTL or NAS – according to individual requirements.

The tight integration of hardware components within the modular, highly scalable grid architecture offers a disaster-resilient, datacenter solution for backup and archiving purposes. Based on a highly flexible frontend/backend design, the hosts are directly connected to the ETERNUS CS VTL or ETERNUS CS NAS frontend interfaces. Completely decoupled from the hosts, the ETERNUS CS8000 not only controls backup and archive data within its internal RAID storage, but also the migration of this data to tape or disk storage devices at the backend. In addition, optional deduplication of backup data on disk reduces disk capacity requirements. Provisioning of resources like frontend connectivity, appliance performance or aggregate storage capacity is freely configurable.

The virtualization subsystems eliminate dependencies between the data accessed via specific host interfaces and the target storage devices where the files are physically stored. One of the major benefits of abstracting hosts from the actual storage is the ability to use independent and centralized, device-based hierarchical storage management (HSM) while maintaining concurrent I/O access. Device-based HSM allows for automated determination of storage locations for backup and archive data across different types of storage devices and media. The non-disruptive data movement according performance, capacity and cost requirements embraces state of the art technologies like tape, disk, compression, and deduplication.

The ETERNUS CS8000 solution offers several advantages:

- Host-independent unified backup and archive data management on disk and tape
- Industry standard hardware components within one single scale-out architecture
- Disaster resiliency and failover between sites based on a unified cluster technology
- Common subsystem control with central appliance control units
- Common user interface and service functions

For more detailed configuration options, please see the datasheets at www.fujitsu.com/eternus-cs8000.



3.2 Commvault Software

Commvault’s Data Platform and Solutions Portfolio is uniquely positioned to help customers activate their data by accelerating the transformation from legacy data management to a modern data environment to unlock critical business insight and drive new value from customers’ technology investments. The Commvault solutions portfolio is comprised of an industry-leading product offering in the areas of data protection and recovery, cloud, virtualization, archive, and file sync and share that addresses evolving market trends, growth areas, and customer use cases. The Commvault Data Platform is open and standards-based, helping customers better leverage data, improve IT operations, and enabling third-party innovation.

The Commvault Data platform consists of four product pillars:






 Commvault Complete Backup & Recovery	 Commvault HyperScale™	 Commvault Orchestrate™	 Commvault Activate™
Includes everything you need to conduct backup, recovery, and archiving activities, enable operational reporting, and perform hardware snapshot management, all in one complete solution.	Software leverages the cost and scale efficiencies of the cloud on-premises using a software-defined scale-out architecture	Simplifies disaster recovery, development and testing, and workload migration through automation and orchestration	Extracts data insights for better data governance and business outcomes
Complete can be extended with each module to address customer scale and complexity challenges ----->			

Figure 3: Commvault Data Platform

With a flexible, yet comprehensive data management platform, you can quickly deliver actionable insights from data stored across your IT infrastructure. Commvault Software protect, manage and access a broad range of data types across primary infrastructure, cloud and secondary storage locations. You can quickly test a new customer application, support a new overseas office, and have the confidence to protect your organization from a ransomware attack. Commvault’s data protection solution supports all major operating systems, applications, and databases on virtual and physical servers, NAS shares, cloud-based infrastructures, and mobile devices.

For more information, please refer to www.fujitsu.com/emeia/commvault or www.commvault.com



4. Use Cases

Now it's time to look at a number of real world scenarios and practical use cases that demonstrate how working with ETERNUS CS8000 and Commvault Software really makes a difference and which underpin the solution of the above mentioned pain points. The use cases outlined in the next chapter will show you what the technical improvements mean from a business perspective.



4.1 Backup in Open Systems Environments

Figure 4 depicts in how backups for open systems work in a redesigned data protection environment with the ETERNUS CS8000 deployed.

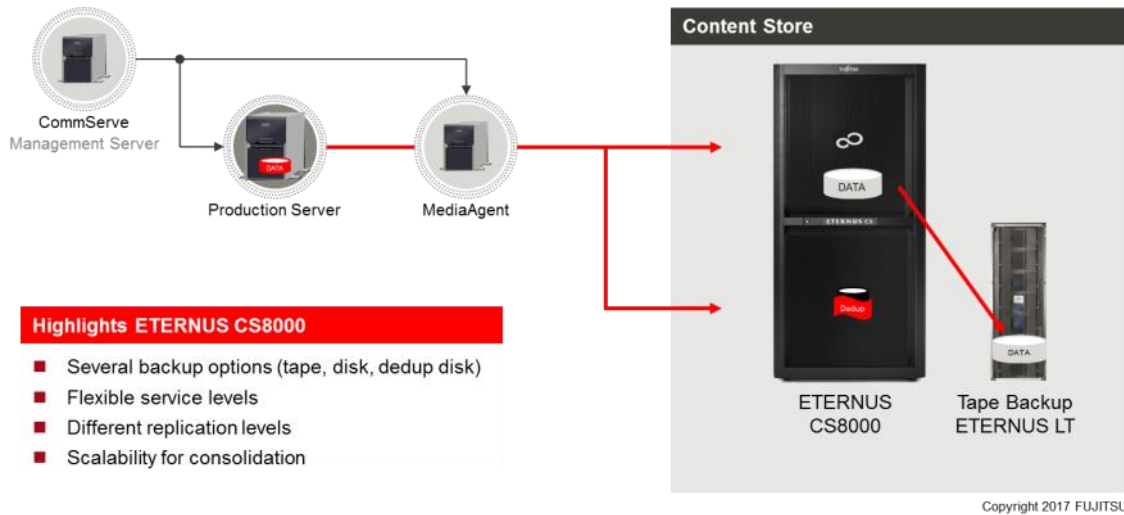


Figure 4: Backup with ETERNUS CS8000 and Commvault Software

As you can see, the entire backup process is launched by the CommServe component of the Commvault Software (top left). The CommServe component is a piece of server software that controls the MediaAgent, another server system that transmits data from the production server to the ETERNUS CS8000. The data protection appliance then sends the data to the different categories of backend storage according to predetermined backup policies. In this particular setup, production data goes to the internal RAID system, consisting of high-speed disk arrays, where they are kept for a user-defined time to enable fast, complete restores of active physical or virtual systems. Data for long-term, nearline and offsite retention is stored in an internal online cache and additionally sent to a tape library attached at the ETERNUS CS8000 back-end (red arrow). As an alternative, backup data can also be stored to a deduplication store within the ETERNUS CS8000.



4.1.1 Flexible, Scalable and Streamlined Workflows

Quite frequently, one major issue directly derives from the heterogeneity of open systems environments. As many data protection and management solutions have been customized to fit specific server hardware and operating systems, they tend to be extremely static – meaning that they rely on fixed backup paths and expect predefined targets. This results in a lack of flexibility, scalability, and adaptability to new technologies and media types. The combined Fujitsu-Commvault solution addresses these problems in an elegant way, as both the backup software and the backup appliance work in accordance with the same principle – using a grid of nodes managed by a master instance.

On the software side, this means that growing backup workloads can easily be managed by adding new media agents while management is done by a single master component, the CommServe.

Commvault Software also provides a number of additional scalability features such as

- **Workflow automation:** Workflow capabilities in Commvault Software let you automate repetitive or highly complex data management tasks. Admins can select from an extensive catalog of typical tasks (such as “Register New Clients”), or create their own custom workflows using the intuitive GUI. With these advanced tools, admins may automate business tasks by bringing together sets of individual tasks in a specific order or decision tree.
- **Virtual Machine Auto-Protection:** Commvault Software provides industry-leading capabilities for policy-based auto-protection of virtual machines to ensure no VM goes unprotected. Auto-protection rules leverage deep hypervisor integration to deliver great flexibility. Several discovery rules can be combined, mixed and customized to include or exclude VMs based on affinity rules for data stores, resource pools, virtual data centers (VDCs), VM names, and a host of other options.

The unique ETERNUS CS8000 architecture follows a similar principle. No matter how much capacity or performance is required, Commvault Software continues to see a single backup repository. But this approach calls for extreme scalability, which the solution can easily provide. Customers may add processor nodes to increase performance or add disk storage components, tape drives and tape libraries to increase capacity. Still, the ETERNUS CS8000 at any time only represents a single backup repository that can itself easily adapt to growing workloads by adding new grid elements.

In terms of capacity, this means that one single ETERNUS CS8000 scales from 14 terabytes internal disk capacity up to 96 petabytes. All data received via the VTL interface are automatically compressed before being saved to disk. Up to 4.8 Petabyte of disk capacity can be used as integrated deduplication disk store. In addition, the ETERNUS CS8000 can handle exabytes of tape capacity. To ensure adequate access times to this enormous storage pool, customers may add further server nodes to improve performance. That way, they may achieve sustained data transfer rates of up to 150 TB per hour.

For technical details please refer to ['Data sheet ETERNUS CS8000'](#).

As a result, the combination of Commvault Software and ETERNUS CS8000 forms a backup solution with practically unlimited scalability and a central administration console that allows for a lean, standardized management.

The advantages of a backup with ETERNUS CS8000 and Commvault Software are as follows:

- The IT department and administrators benefit from the immense scalability of the ETERNUS CS8000 as well as from an unprecedented freedom of choice. Virtually all characteristics of the backup and recovery process can be properly specified, including target systems, backup technologies and paths, number and depth of replication levels and service options for restores. For example, administrators may determine exactly what happens to a given data set– i.e. whether and when it is moved to disk, or to dedup disk, or to an inexpensive tape, or to multiple targets. Furthermore, they may also define the number of local and remote data copies that must be kept. As a result, customers will be able to meet even the tightest SLAs and obligations regarding backup/recovery speed, disaster resilience, and costs.
- In addition, companies also profit from the most comprehensive set of data protection and recovery capabilities available in a single appliance, which includes features such as multiple copy functions, asynchronous and cascaded replication, remote replication to tape systems, virtual tape copies for emergency restores, and self-healing copy management.

4.1.2 Protecting Data in Virtualized Environments



The past decade has seen an unprecedented boost of server and storage virtualization in open systems environments. Driven by budget restraints, new environmental laws, server and storage consolidation efforts and the simple need for replacing outdated hardware, many enterprises have successfully deployed the technology, mainly based on implementations of VMware vSphere and Microsoft Hyper-V. With virtualization, most data centers achieved their goal of reducing costs while preserving or even improving their service quality, thereby avoiding a trap that had seemed inescapable before. This success, however, had its price.

In its early stages, virtualization inherently complicates backup and recovery. As with physical servers, the main challenge for administrators is to create consistent backups that can be restored quickly and swiftly in order to repair crashed systems or applications. In highly virtualized enterprise environments, the problem is further amplified because VMs permanently write to the data store of physical servers, thereby changing the file system content.

Finally, the task of backing up entire virtualization farms with hundreds or thousands of VMs in a short period of time requires a new, modern approach. An adequate solution needs to reflect the change and to provide methods for continuous or near-continuous data protection along with the ability to leverage storage snapshots. The problem is well-known; and backup specialists as well as the virtualization market leaders offer tools to address it.

Unfortunately, the bulk of these solutions suffer from technical shortcomings. Several programs are platform-dependent, meaning they only work with a limited choice of server-storage combinations. Others are designed for enterprises of a specific size, which means that companies have to invest in new software as they expand – including additional training. Moreover, the selection of backup and recovery methods and paths is often limited. Finally, although state-of-the-art products generally integrate snapshot technology for fast restores, their effectiveness is frequently hampered by hardware limitations.

Our combined solution sidesteps these shortcomings and offer most of the qualities that virtualization requires:

- Because of their underlying architecture and unified approach, it works with all open systems platforms, regardless of server and storage hardware and operating systems, as well as with all applications found in these environments.
- It automatically detects virtual machines using protection rules such as “backup all VMs residing on data store X”
- Administrators can pick the backup method best suited for the job at hand and either can create drive images to preserve system configurations or store data at file or block level. In terms of backup types, all systems support full, incremental, differential, and synthetic backups. Supported backup paths include tape, disk, bare metal, and snapshots.
- The software facilitates hypervisor migrations by converting VMs created in e.g. Hyper-V to VMware format and vice versa.

- Recovery options are equally varied and reflect backup methods, types and paths. Administrators can achieve multiple levels of granularity and restore entire servers, single virtual machines, or specific files – if necessary, all from one single image level backup.
- It automates the creation of snapshots at significant points in time (e.g., before software patches are installed) to ensure data and application consistency. As a result, administrators can choose between multiple point-in-time recovery options in case of a failure.
- Thanks to technologies like agent-less backup of VMs and built-in global deduplication, the software layer uses less bandwidth and storage space than competing solutions.

With their comprehensive feature set, the combination of ETERNUS CS8000 and Commvault Software allows for seamless and consistent data protection of all physical and virtual servers and for fast, reliable and efficient restores – everything needed to ensure business continuity in emergencies.



4.1.3 Enhanced Reliability

As data is the most valuable asset of a company, it is of paramount importance that the Fujitsu-Commvault solution guarantees data availability and reliability at any time. To achieve this, we built it around the concept of redundancy, implemented at different levels.

Originally developed as the first data protection appliance to consolidate backup data of open system and mainframe environments, ETERNUS CS8000 is designed as a system with no single point of failure, to fulfill even mainframe reliability levels. Hardware redundancy means that for every single primary component of an ETERNUS CS8000 – hard disks, processors, RAM, tape drives, RAID controllers, PSU etc. – there is a reserve building block that can take over in case of a malfunction. The risk of a system failure is thus substantially minimized, if not excluded.

However, there is more to redundancy than fitting a system with duplicate components for emergency cases. That is why at the second level we aimed to implement redundancy of information as well as of data protection processes in various ways. Thanks to running disk-to-disk-to-tape (D2D2T) backups, the ETERNUS CS8000 data protection appliance is able to generate and manage multiple copies of data on local disk or tape, from where they may be easily recovered in case an application fails. In addition, copies can be written to remote tape libraries located thousands of miles away. This provides an extremely easy and economical way to introduce disaster resilience. In addition backup to tape offers offline backup which is essential to thwart ransomware and malware attacks.

Third, features like remote and asynchronous replication between several ETERNUS CS8000 systems ensure that up-to-date copies of production data and recent backups are always available for speedy and simple restores. Crucial information remains accessible and serviceable at all times, thanks to a built-in functional layer that affords redundancy of data.

Finally, the ETERNUS CS8000 split-site configuration with cache mirror ranks reliability to new heights. In this setup, two identical halves of an ETERNUS CS8000 work as one logical unit that is distributed across two remote locations. One site will continue to deliver full data availability even if the other suffers from power outages, blazes or other catastrophes. Here, internal system redundancy turns into the highest level of disaster resilience.

On the software side the Commvault Software provides several redundancy and failover features:

- **Dynamic Drive Sharing (DDS)**
With Dynamic Drive Sharing (DDS) activated, all libraries, drives or ETERNUS CS8000 can be shared by the MediaAgents that have access to it. One or more MediaAgents can be configured as failover candidates to ensure that all jobs using the storage target devices are performed without interruption in the event of a failure in any one MediaAgent.
- **Auxiliary Copy**
The Auxiliary Copy operation creates additional standby copies of data. These standby copies are referred to as secondary copies. Primary and secondary copies may be stored on different media and in different libraries. If the primary copy becomes inoperative due to failure or network malfunctions, administrators can promote a synchronous secondary copy to take on the role of a primary copy. Thus, applications can continue to run like before, and IT departments may carry out repairs without interrupting backup and restore operations.
- **GridStor® (Alternate Data Paths)**
A data path is the combination of MediaAgent and storage pool to perform a backup operation. All storage data has a default data path. The administrator can define alternate data paths providing the following advantages:
 - Automatic switch-over to an alternate data path when one of the components in the default data path is not available.
 - Usage of available libraries and drives in the event of failure or non-availability of these resources

- Load-balancing between alternate data paths, which provides evenly-distributed backup operations among available resources

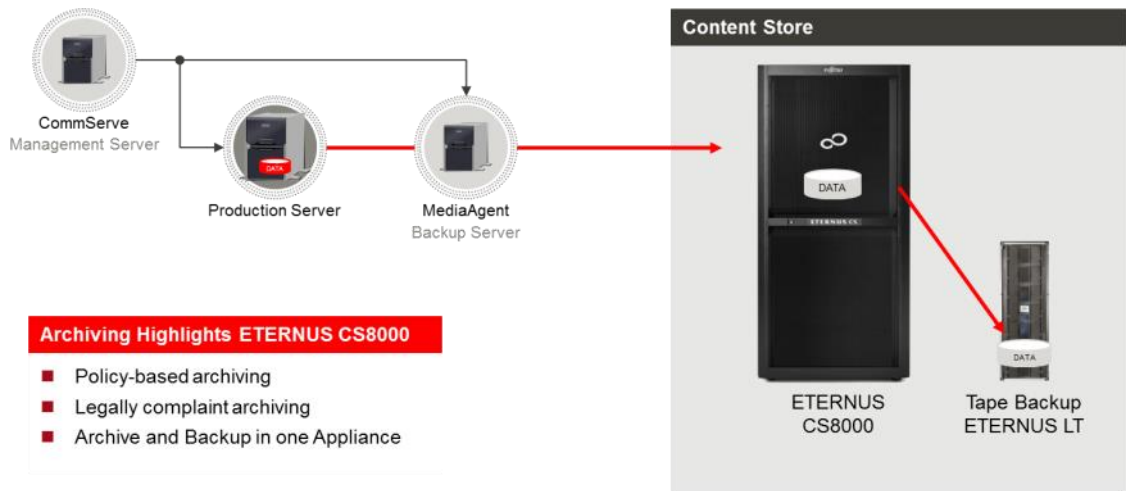
■ Disaster Recovery of the CommCell

CommServe functionality provides centralized management and administration for the entire CommCell group. Hence, CommServe system failures result in the entire CommCell becoming non-functional. When this happens, admins can switch to a standby CommServe or temporary CommServe to quickly resume CommCell operations. Disaster recovery procedures help to shift CommCell operations from the primary to a standby CommServe system with minimal data loss and downtime.



4.2 Archiving and second-tier file storage

Figure 5 shows how the restructured archiving process differs from its predecessor.



Copyright 2017 FUJITSU

Figure 5: Archiving with ETERNUS CS8000 and Commvault Software

The archiving process depicted is nearly identical with the backup process except for one main difference. The archiving module of the Commvault Software prompts the Media Agent to write data to a particular share – e.g. in ETERNUS CS8000 NAS share – where the data is kept for a predefined retention period. Archiving rules and policies are set within the ETERNUS CS8000; if necessary, data can be written to multiple targets for different retention periods using the dual-/triple-save feature. This makes archiving more targeted and efficient. The IT department can define rules and policies for legally and standards-compliant archiving in the same way it define rules for backups, and so implements an automated, swift and efficient process. As outlined in the earlier Virtualization section, IT companies can now run the archiving process simultaneously with backup, analysis, and deduplication. Moreover, there is no need for extra hardware, as archives can reside on the same systems as nearline backups.



4.2.1 Improving Storage Tiering for Increased Efficiency

The basic idea of a backup is to create copies of the original data from the production systems with the purpose of being able to recover data when necessary. These copies are then used to get staff back to work following application failures or hardware outages – or to get entire businesses back up and running after floods, fires and other disasters. Backups are focused on constantly changing business information, which means they are generally kept for the short term and often overwritten; say monthly, when full backups are due. This makes them a poor choice for retaining data for compliance reasons. By moving old or non-critical data from production storage to second-tier file storage, administrators can increase the usable storage capacity on systems that hold mission-critical data. The reduction of data volume stored in productive environments results in faster backup and recovery, and data in second-tier file storage can be managed with more ease at a different service level.

Archives, on the other hand, are not copies of production data, but rather the primary version of a piece of data, which is often inactive or non-changing. When data stops changing, it is best to move it to an archive, where it lives outside the backup window but can still be accessed. Archives do not focus on recovering an application or business data, but on information retrieval – usually at the level of a file, e-mail, or other individual piece of content. Archives are typically used for long-term retention of information, which makes them the best choice for managing data with regulatory requirements.

Traditionally, companies have treated backup/recovery and archiving as separate processes that run on separate systems. Over the years, this has led to the formation of two strictly distinct infrastructures, each of which maintains its own set of rules, tool chest, and upgrade cycles. While initially this parallelism made sense, it turned into a weakness as storage infrastructures evolved; nowadays, it means that one and the same data set is touched multiple times, by different administrators and with different software, to achieve what is essentially the same goal – being moved from production systems through backups to archives during their lifecycle. This is a waste of time, workforce, and money. To stop it, modern storage solutions must eliminate inefficient data silos, consolidate managed data, and automate retention and tiering according to the policies customers define. Plus, they must provide a global, intelligent index and multiple retrieval methods so that users can easily access the information they need when they need it, whether it's for recovery, eDiscovery or data mining.

The ETERNUS CS8000 with Commvault Software delivers all the required functionalities. Archiving agents move data from primary to secondary storage to optimize storage space, and retain and find relevant data on any disk, tape, or cloud storage. As a result, all traditional backup and archiving processes occur in a single operation that moves all data to secondary storage, where it serves as both backup and archive copy. On the primary storage system(s), only stubs are retained that point to the data moved as part of backup. The archived data is available for quick and easy retrieval. Administrators may retain, store, classify, and access information according to its business, compliance, or evidentiary value, with one method of access and preservation across all ESI (Electronically Stored Information).



4.2.2 Archive Protection

Since an archive is not a backup that gets overwritten in due time, it needs to be protected, too. With the ETERNUS CS8000, that's fairly easy, as it can store data in both local and remote tape libraries. This is an uncomplicated way to introduce disaster resiliency. What's more, the data protection appliance also natively supports asynchronous replication to a remote ETERNUS CS8000 system for the archive/files store, so that no extra replication software is needed. This is particularly useful for long-distance replication scenarios where data centers are thousands of miles apart (cf. Chapter 4.1.3).

The core element of the most disaster-resilient architecture is a logical ETERNUS CS8000 system deployed across two geographically separate sites. The internal infrastructure is thereby extended to a second site which can be 100 km away. The associated applications from the production environment still only see one logical setup, without ever being aware of the geographical location. Data can be written to or read from both sites. This dramatically increases the availability of data, even in worst-case scenarios.

Often, data that is stored in an archive is not part of the backup process, leaving it unprotected. This can be a serious problem, as data moved away from the backup-protected a productive environment is at risk. To solve this issue, ETERNUS CS8000 includes an integrated, automated backup function for archive and second-tier file storage data that reside on the NAS subsystem. This provides lots of benefits, such as increased availability levels, protecting organizations against user errors and viruses, and meeting specific SLA requirements. Compared with external archive backup solutions, which are very complex and use up a large amount of resources, the solution integrated into the ETERNUS CS8000 reduces management efforts and costs to a minimum.

The ETERNUS CS8000 brings robustness and resiliency to the archive environment, at the same time it simplifies operations by avoiding complex policies and administration tasks.



4.3 Consolidation of data protection environments

To better understand the situation, let's first look at a typical data protection environment. Traditionally datacenters have two dedicated infrastructures in place – one for backup and one for archiving (cf. Figure 1: Typical data protection environment). Both run on their own hardware (disk arrays, dedup systems, and tape libraries) and require separate management tools. This is the standard case for any open systems environment. Depending on the size and type of organization, further complication arises if a company still operates mainframe computers that demand yet another dedicated infrastructure – as is common, for instance, in banks, insurances or the pharmaceutical industry. Additional problems confront quite a few companies that use their backups as archives or, worse, forego archiving altogether in breach of compliance regulations. And, as noted before, we need to remember that companies worldwide are facing exponential data growth while CIOs have to ensure they maintain or, better still, reduce the costs of running IT infrastructures. But what does this mean with regard to backup and archiving?

As a result of working with isolated point solutions, companies face several serious problems:

- Backup and archiving costs at its main data center are already very high due to the complex setup – and will explode in the medium term, because an increase in customers has triggered exponential data growth.
- The explosive data growth has caused backups to violate backup window restrictions; restores are equally time-consuming. The resulting delays negatively impact the company's level performance by slowing down business and production processes.

- Transferring data from backup to the archiving part of the infrastructure is accompanied by manual efforts that cause a delay in optimal usage of the infrastructure. "Cold" data blocks highly available, expensive backup storage for longer than needed before being transferred to cost-efficient second-tier or archiving platforms.



4.3.1 Consolidation, Simplification and Optimization

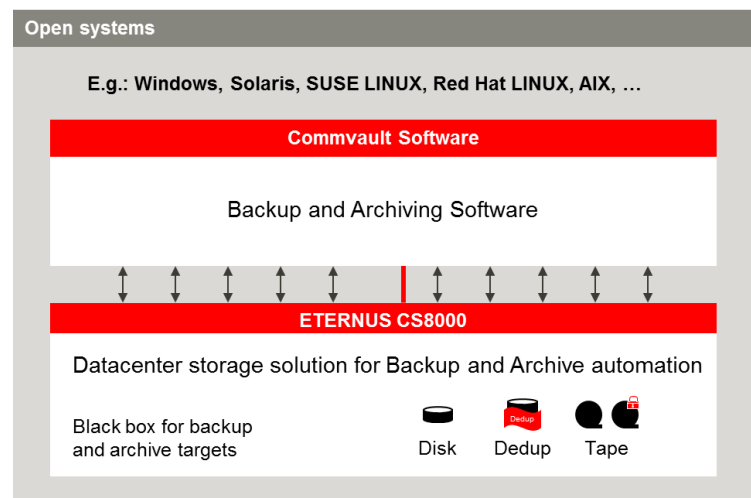
As mentioned before, one of the longest-standing nightmares for storage administrators, IT leaders and CIOs is being forced to deal with the effects of permanent exponential data growth. Until about ten years ago, the standard method was to add new storage systems whenever existing ones showed a lack in capacity or performance. Today that's no longer a viable solution, for various reasons. One reason is IT budget cuts and freezes. Companies of all sizes hesitate to invest in new hardware, especially if the investment seems to yield no immediate, visible results. Another, equally important reason is that this method is no longer technically appropriate: Surveys show that the annual data growth rate is currently at 50% on average; some industries even report a 100%, meaning that their data pools double each year. Estimates predict that by 2020 the total amount of data stored worldwide will reach 44 zettabytes – 35 times the amount of 2010. And given current and future trends – from Big Data to the burgeoning Internet of Things – the surge will not stop anytime soon: analysts reckon that by the end of the decade, about 200 billion smart devices will produce an as-yet-inconceivable flash flood of information.

Against this backdrop, it's becoming ever harder for IT departments and storage admins to fulfill their tasks. Still their employers expect them to square the circle and not only control data growth, but also somehow continue to provide the same or better backup, recovery, archiving and data management services. This is where the combination of Fujitsu ETERNUS CS8000 and Commvault Software comes into play: The solution has exactly the capabilities that are needed to contain the ongoing data increase and storage system sprawl, simplify backup and archiving through automation and rules-based management, and tune up entire storage environments to fit future demands. More specifically, they enable IT departments to drastically reduce the number of target systems required to back up and archive data from mainframe and open systems environments. They also implement an abstraction layer between open systems servers and backup applications, so that administrators may further reduce the amount of open system-specific targets. Due to its inherent multi-platform support, the ETERNUS CS8000 / Commvault Software combination swiftly integrates into existing data protection and data management processes. Thus it provides one single, streamlined method to incorporate data from all open systems environments into company-wide backup and archiving as efficiently and securely as possible.



4.3.2 Implementing Modern Information Management

One of the key tasks of IT departments is to deliver the right data to the right person or application at the right time. To reach this goal, all relevant information has to be managed from the moment it enters an organization until it can be retired. Powerful backup solutions provide a basis for this Information Lifecycle Management (ILM). But despite this close relationship – including content indexing, policy-based retention/archiving, search, and export to secondary storage – data protection and ILM typically remain isolated from one another. The solution of ETERNUS CS8000 and Commvault Software overcomes this separation by storing all data, regardless of origin, in a single virtual pool that is accessible from all applications, IT systems and locations. Thus, valuable information can be swiftly retrieved for multiple purposes such as OLAP or eDiscovery. The same concept is applied to hardware-related tasks like Storage Resource Management.



Copyright 2017 FUJITSU

Figure 6: Modern data protection environment

Figure 6 shows how deploying ETERNUS CS8000 data protection appliances with Commvault Software help simplify backup and archiving at the main data center. In this new topology, the ETERNUS appliance virtualizes the different target systems (disk, tape and dedup disk) and consolidates them into one extremely scalable, unified data protection platform. The radically simplified setup results in a number of key advantages:

- The unified platform provides a central storage pool for backups and archives from all types of hosts (mainframes, UNIX systems, and x86 servers) that can be managed as a single entity – as opposed to the point solutions from before that required separate management tools and processes. Administrators only need to define a set of policies for the ETERNUS CS8000 system (e.g. to determine whether data is stored on disk or tape), and all mainframe backup applications or, in the case of open systems, Commvault Software will cooperate nicely.
- With all disk and tape storage capacities being combined into one huge, scalable pool, previously underutilized systems may be easily retired, helping IT departments to substantially save on power consumption, cooling, and floor space. In addition, since all systems may be easily expanded by adding processor nodes, disks or tape drives, necessary hardware expenditures are kept to a minimum.
- The ETERNUS CS8000 handles all data management operations automatically. This applies to all types of data copies as well as to all types of “data handling operations,” including processes as diverse as tape migration and refresh, replication, mirroring and deduplication. Again, the overall results are large-scale cost savings.

Commvault Software enables progressive enterprises to rethink data management across the organization addressing the following:

- Protect your data.
 - Commvault keeps it simple: all customer data is protected and secure, whether in a physical, virtual or cloud environment, whether on desktops, servers or endpoints.
 - That gives customers the confidence and freedom to focus on what *really* matters: their business.
- Access your data.
 - Only Commvault safeguards all customer data – backups, snapshots, archives, etc. – in virtual native format, with full indexing.
 - That means customers know everything about that data: where it is, what it is, who owns it, when it was protected and much, much more. That means Commvault customers are guaranteed the most immediate availability in the industry.
 - No more static copies or dark data – Commvault customers get back data that is actionable and ready to put to maximum strategic use.
- Be compliant with your data.
 - Commvault solutions enable customers to produce, retrieve and review all discoverable information, on demand. No waiting.
 - Leveraging powerful indexed search across multiple data types with a single platform not only saves money and time, it minimizes risk and exposure.
- Share your data.
 - People share things. For enterprises that can be a problem, as employees use consumer tools to share company information. While user collaboration and productivity are served, security and data protection become huge issues.
 - Commvault brings control back into enterprise data management with secure file sharing and anytime/anywhere data access, powered by enterprise-grade backup. Users get powerful sharing and access, IT gets comprehensive backup and security for all enterprise data, and the company benefits from powerful analytics and reporting to help optimize data use and value.



5. Conclusion

By deploying combined ETERNUS CS8000 and Commvault Software solutions, customers can build reliable and highly efficient data protection infrastructures that comply with legal and professional standards and cut expenses at the same time. The integrated scenarios described in this paper only offer a glimpse at what companies can do to eliminate time-consuming, cumbersome processes and replace them with their flexible, agile counterparts. In short, a solution consisting of Commvault Software combined with the ETERNUS CS8000 solution enables IT departments to successfully handle future data growth, retain control over their information assets at any time, and own infrastructures that remain reliable for years to come.

The many advantages and savings opportunities over point solution scenarios include these select examples:

- **Reduction of administrative burdens and costs** through automation and centralization of core functions
- **Extreme scalability** – maximum performance can be achieved with minimal administrative costs
- **Utilization of synergies** between backup, archiving and second-tier file storage to reduce daily backup efforts
- **Investment protection** through better utilization of backend infrastructures
- **Simple planning and implementation** of storage levels and hierarchies, guaranteed data availability according to company guidelines
- **Transparent, policy-based archiving** and indexing of application data
- **Central classification, comment and search functions** for archived content
- **Stable, reliable performance** even in case of server, storage or network failures through **excellent error correction**
- **Significant acceleration of backups** and reduction of backup times even in Big Data environments
- **Easy integration** of heterogeneous storage hardware and virtual storage

6. Index of Figures

Figure 1: Typical data protection environment.....	3
Figure 2: Fujitsu Storage ETERNUS CS8000.....	4
Figure 3: Commvault Data Platform.....	5
Figure 4: Backup with ETERNUS CS8000 and Commvault Software	6
Figure 5: Archiving with ETERNUS CS8000 and Commvault Software	9
Figure 6: Modern data protection environment	11

Contact

FUJITSU Technology Solutions GmbH
Address: Mies-van-der-Rohe-Str. 8, D- 80807 Munich
E-mail: storage-pm@ts.fujitsu.com
Website: www.fujitsu.com/emeia
2018-12-05 WW EN

© Copyright 2018 FUJITSU LIMITED. All rights reserved, including intellectual property rights. Fujitsu, the Fujitsu logo, are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner. Commvault and the Commvault logo are trademarks or registered trademarks owned by Commvault Systems, Inc. in the United States and/or worldwide.