


White Paper

FUJITSU Storage ETERNUS CS200c Powered by Commvault®

Data protection is an essential yet laborious and time-consuming task that every organization must perform. Identifying the right hardware and software, and managing the data protection environment, can tie up IT teams. The integrated backup appliance ETERNUS CS200c powered by the Commvault Data Platform® helps resolve these challenges and saves customers time and money.



Content

Introduction	2
ETERNUS CS200c – Overview	3
ETERNUS CS200c – Licensing	4
ETERNUS CS200c Software Features	5
Backup and Recovery	5
Basic Endpoint Backup	5
Deduplication	6
Encryption	8
Replication	8
Archiving	8
Snapshot support (Commvault IntelliSnap®)	10
Snapshot replication support for ETERNUS Storage Arrays	10
Tape support	12
Cloud Connection	12
Conclusion	13

Introduction

Data protection is an essential yet laborious and time-consuming task that every organization must perform. Organizations capture, store and retain more types of data for several periods of time than ever before. Identifying the right hardware and software, and managing the data protection environment, can tie up IT teams. Integrated backup appliances address these challenges. Market analysts expect the PBBA (Purpose-Built Backup Appliance) market to grow over the months ahead as customers continue to opt for appliances that meet their growing data protection and recovery challenges.

The FUJITSU Storage ETERNUS CS data protection appliances offer outstanding data protection capabilities by aligning storage resources with business priorities to deliver business-centric storage. The ETERNUS CS portfolio offers data protection solutions for all company sizes and needs: ETERNUS CS200c is an integrated backup appliance providing a complete solution without the need for additional backup software. ETERNUS CS800 is a deduplication appliance designed for cost-efficient backup to disk in small and mid-sized environments. The ETERNUS CS8050 NAS system is a reliable, cost-efficient storage appliance for archiving and second-tier file storage. The unified data protection appliance ETERNUS CS8000 supports upper mid-sized and large enterprises for the complete consolidation of backup and archiving infrastructures of open systems and mainframes (for more details, visit the [Fujitsu Internet site](#)).

Commvault's Software enhances Fujitsu's "[Business-Centric Storage](#)" strategy by providing a single, modular, and scalable architecture with robust data structures. The Commvault Software is a perfect match and logical extension for Fujitsu Storage ETERNUS hardware, resulting from a long and close collaboration between both companies that led to a strategic partnership in 2011. Since then, Fujitsu has augmented its storage systems with Commvault Software as well as related maintenance and management services.

Both customers and partners benefit from this approach because it increases flexibility and helps keep operational expenditures and the cost of growth under control.

This white paper describes the features of the integrated backup appliance ETERNUS CS200c Powered by Commvault. Each environment is unique. Several licensing options and appliance models provide the solution that offers a custom fit for your data protection needs.

ETERNUS CS200c – Overview

Building Blocks

The Fujitsu Storage ETERNUS CS200c Powered by Commvault integrates correctly-sized hardware and data protection software into a complete backup and archiving solution. The appliance consists of all required software and hardware components within a Commvault Software CommCell. The building blocks are:

- Management server, in Commvault nomenclature the so-called CommServe
- Backup server with installed MediaAgent
- Intelligent Data Agent (= client agents) for the clients
- Build-in software deduplication
- Raid-protected online hard disk storage and/or flash storage
- Tape attachment
- Cloud connection

The all-in-one appliance perfectly combines these blocks. The system is preconfigured and preinstalled to enable a fast and hassle-free setup of a comprehensive backup and archiving environment.

Industry-leading Commvault Software is perfectly integrated with powerful Fujitsu system technology, providing excellent features for business efficiency and continuity. The solution addresses the challenges in protecting business applications, virtualized, converged and hyper-converged environments and facilitates compliance with regulations.

The integrated Commvault Software provides a single platform and index for your entire environment, including backup, deduplication, archiving, disaster recovery, replication, snapshot management, and more to radically simplify data management across its lifecycle.

Fujitsu Storage ETERNUS LT tape libraries can be attached for long-term retention, archiving, and second-tier backup providing offline backup to thwart cyberattacks.

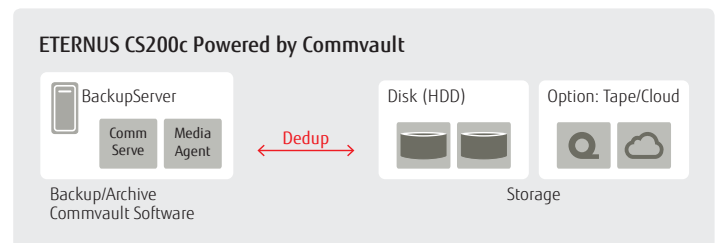


Figure 1: Fujitsu ETERNUS CS200c – architecture



ETERNUS CS200c – Licensing

The ETERNUS CS200c is powered by the industry-leading Commvault Software. The enterprise software is tightly coupled with best-in-class Fujitsu system technology. There are various licensing options available depending on the specific customer environment:

ETERNUS CS200c with software license for physical and virtual environments

The ETERNUS CS200c combines all the advantages of an integrated appliance with the flexibility of a dedicated backup software license. The appliance comes with pre-installed software and you can freely select out of all available Commvault software license flavors, among them capacity (BET, FET), pure VM infrastructure (number of VMs or sockets) or per endpoint/mailbox user.

The software provides modern deduplication, replication, backup, archiving and recovery capabilities to better protect ever-growing data volumes, to efficiently manage information assets, and to quickly find, recover and access data.

The ETERNUS CS200c supports application-aware backup for Windows, Linux/Unix, VMware, Hyper-V, Oracle, and SAP applications. Furthermore, the software delivers native support for more than 40 cloud storage providers across public and private clouds.

ETERNUS CS200c without a License

The ETERNUS CS200c without a license is the best choice for existing Commvault environments. The system is preconfigured and preinstalled for quick and easy start-up.

It is an easy way to quickly deploy new Commvault media agents in existing environments.

Back-end Terabyte versus Front-end Terabyte

The ETERNUS CS200c capacity is metered in Back-end Terabytes (BET). The main difference between BET and Front-end Terabyte (FET) licensing is:

- **Back-end Terabytes (BET)** are the amount of backup data stored on the disk storage and/or appliance storage (typically after deduplication). Applies to ALL used backend capacity, including secondary copies.
- **Front-end Terabytes (FET)** are the total amount of data being backed up before deduplication. Data is measured as the current largest full (or synthetic full) backup performed.

ETERNUS CS200c Software Features

Backup and Recovery

Traditional Backup and Recovery

Figure 2 shows a typical backup infrastructure as it can be found in most companies, regardless of revenue or number of employees. Whenever a backup is initiated, application servers (that store information on internal hard drives, dedicated disk arrays or networked storage) send copies of production data to so-called backup or media servers, which in turn pass it on to storage target systems at the backend, like tape, dedup disk, virtual tape library or target-based backup appliances. For restores, that process is simply reversed.

Backup and Recovery using ETERNUS CS200c

General

When using an integrated backup appliance like ETERNUS CS200c, all of these processes take place within the appliance. You can automate global protection and retention policies from a single, centralized management console.

The ETERNUS CS200c offers seamless and efficient backup and restore of data and information in your enterprise from any operating system, database, and application.

File system backups are supported for all major operating systems and include inherent file system capabilities. Application protection uses application-aware backup agents to provide consistent point-in-time protection for application data. Application-aware backups capture the state of application data at the time of the backup, including data in

memory and pending transactions, making it easier to restore the application for continued use. An application-aware backup for a virtual machine quiesces the application, flushes its memory, and completes pending writes, then creates a software snapshot of application data. After the software snapshot is created, the application resumes. Granular application protection facilitates flexible data protection strategies and simplified recovery methods.

Database protection also uses application aware agents to provide a simplified end-to-end backup solution for database environments of any size.

For more information including a list of supported applications, databases and hypervisors see: [BOL Data Protection and Recovery Agents](#)

Virtualization

The ETERNUS CS200c also supports data protection in virtual environments, regardless of whether they are traditional, converged or hyper-converged infrastructures. Virtualization demands a data management solution that is aware of dynamic workloads, consolidated resources and cloud-based computing models. The Commvault Software lets you virtualize even the most demanding applications, leveraging deep integration in the virtual infrastructure to deliver advanced data management capabilities. It optimizes the recovery and retention of files, virtual machines and virtualized applications. Policy-based auto-protection of virtual machines ensures that no VM will ever be at risk.

For more information about backup of virtual environments, see [Virtualization](#).

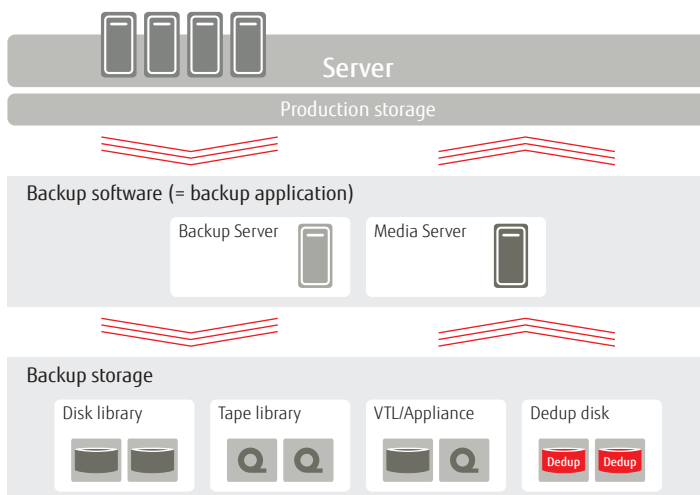


Figure 2: Typical backup infrastructure

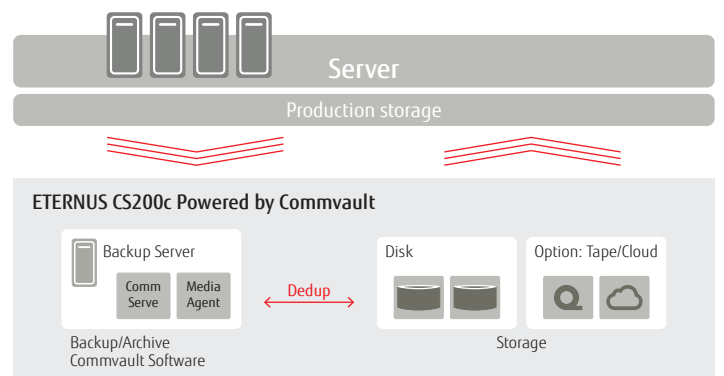


Figure 3: Backup infrastructure with ETERNUS CS200c

Basic Endpoint Backup

Commvault provides the capability to backup data from desktops and laptops to the ETERNUS CS200c. Basic Endpoint Backup features secure protection for business-critical data on laptops and desktops through source-side deduplication, policy-based scheduling, and intelligent bandwidth throttling.

Automated discovery and data access from mobile devices reduce administration efforts.

More information at: www.commvault.com/endpoint

Deduplication

Deduplication – Overview

At a high level, deduplication is a process which compares the blocks in an incoming data stream to the blocks that have already been written to a target. If a redundant block is recognized, a reference to that block is stored in a database, but the actual data is not written to the storage target. Deduplication provides an efficient method to transmit and store data by identifying and eliminating duplicate blocks of data during backups. All data types from Windows, Linux, UNIX operating systems and multiple platforms can be deduplicated when data is copied to secondary storage like integrated backup appliances. Deduplication optimizes the use of storage media by eliminating duplicate blocks of data and reduces network traffic by sending only unique data during backup operations.

The deduplication feature of ETERNUS CS200c can be enabled at the source (client server) or the target (ETERNUS CS200c appliance). Source-side deduplication minimizes the time needed to protect remote data and minimizes WAN requirements. Source-side deduplication conducts the comparison of the changed blocks on the server itself (= source) with a lightweight process, reducing the amount of data that will be

streamed from a client prior to it being transmitted across the network. The comparison is conducted across an environment by using a global deduplication technique. This results in a significant reduction of the backup process and associated window, as well as maximizing the efficiency of the network, LAN or WAN, that the data traverses.

Target-side deduplication (means deduplication on the ETERNUS CS200c) provides a second filter that maximizes the deduplication across clients before data is finally stored on the appliance disk storage.

With the ETERNUS CS200c Powered by Commvault, you can scale with a grid architecture that allows for load balancing and failover to the shared Deduplication Database (DDB). The functionality of DASH Copy (DASH = Deduplication Accelerate Streaming Hash) allows users to take advantage of the powerful Auxiliary Copy (AuxCopy) process. It copies data from one appliance to another in its deduplicated form factor. The feature maximizes storage savings and provides integrated replication of data in a deduplicated state to reduce network traffic. See the [replication section](#) for more detail.

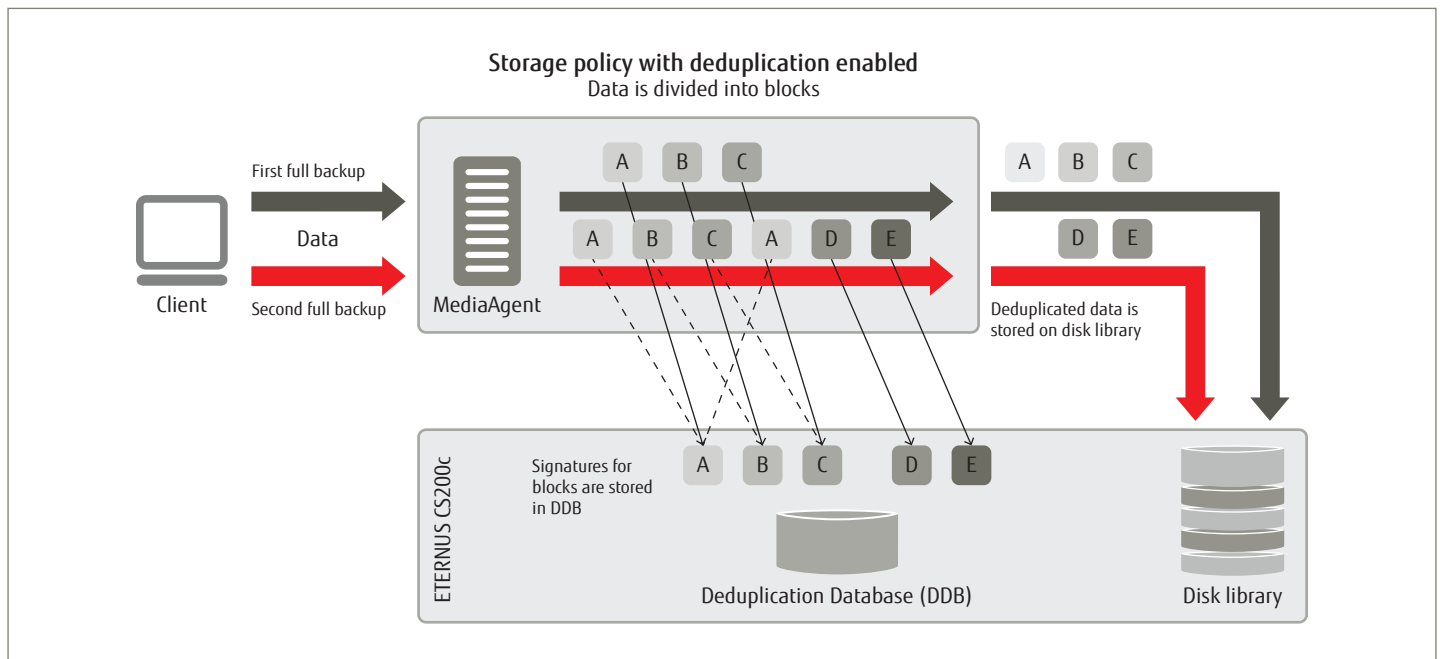


Figure 4: Deduplication

Deduplication process

Deduplication works as follows:

- A block of data is read from the source and a signature for the block of data is generated using a hash algorithm. Signatures are unique for each data block.
- The signature is compared against a database of existing signatures for data blocks that are already on the destination storage (i.e. ETERNUS CS200c). The database that contains the signatures is called the Deduplication Database (DDB).
- If the signature already exists, the DDB records that an existing data block is used again on the destination storage. The associated MediaAgent writes the index information and the duplicate data block is discarded.
- If the signature does not exist, the new signature is added to the DDB. The MediaAgent within the appliance writes both the index information and the data block to the destination storage within the appliance.

During the deduplication process:

- Data blocks can be compressed (default) and/or encrypted (optional).
- Data block compression, signature generation and encryption are performed in that order on the source or destination host.
- Signature comparison is done on the MediaAgent of the appliance. For performance benefits, a locally cached set of signatures on the

source host can be used for the comparison. If a signature does not exist in the local cache set, it will be sent on to the MediaAgent for comparison.

- An object (file, message, document, and so on) written to the destination storage may contain one or many data blocks. These blocks might be distributed on the destination storage. An index that is maintained by a MediaAgent tracks the location of the data blocks. This index allows the blocks to be reassembled so that the object can be restored or copied to other locations. The DDB is not involved in the restore process.

Block length

The standard deduplication algorithm uses a fixed block length. The administrator defines the block length depending on user data.

Additionally the administrator can change the fixed block length algorithm to a variable block length. But this method requires high system performance.

For more details see:

Commvault Online Books: [Deduplication](#)

Encryption

General

Data encryption provides the ability to encrypt data both for transmission over non-secure networks and for storage on media. The flexibility of key management schemes makes data encryption useful in a wide variety of configurations. Commvault Software provides hardware encryption (by encrypting data on tape drives) and software encryption. The ETERNUS CS200c uses the software encryption method.

Software encryption

Software encryption encrypts the data during the backup job data replication job, and auxiliary copy operations (encrypts backup data while being copied to secondary copies). All of the encryption is symmetric cryptography (the same key is used to encrypt and decrypt), so there is no need for a certificate or certificate authority. The decryption of the encrypted data occurs at the client during restore or on the MediaAgent during the synthetic full or during the auxiliary copy process.

Software encryption can be specified and configured at the following levels:

- Client (for backups)
Encryption on a client protects data during data protection and recovery operations. Users can select which encryption cipher to use and where keys are stored – in the CommServe database or optionally on the media itself.
- Subclient (for backups)
Encryption on a subclient allows users to select if and where encryption should be performed for the subclient data during backups. The encryption settings on a specific subclient or instance properties can be disabled.
- Replication Set
Encryption on a replication set allows you to protect replicated data as it transits the network. It encrypts data on the source, encrypts replicated data across the network to the destination system, and decrypts on the destination system. By default, data encryption applies to all of its replication pairs.
- Storage Policy Copy
The encryption on a storage policy copy encrypts data during data protection operations before writing it to the media. Data encryption keys are generated per storage policy copy of the archive file and stored in the CommServe database. If there are multiple copies in a storage policy, the same archive files in each copy get a different encryption key. Individual archive files, however, will have different encryption keys.

Algorithms

Several Data Encryption algorithms like Blowfish, GOST, AES, 3-DES and key lengths are supported and approved by the Federal Information Processing Standard (FIPS).

For more information see: [Data Encryption](#)

Replication

The ETERNUS CS200c supports replication with the DASH (Deduplication Accelerate Streaming Hash) Copy feature.

DASH Copy

The DASH Copy is an ideal solution for off-site copies to the recovery site. It is a fast method of copying data by transferring only the changed data to a secondary copy (second disk target or, in our case, to a second ETERNUS CS200c appliance). DASH Copy leverages a unique approach to managing deduplication that allows users to generate multiple copies of data while remaining in deduplicated format. Think of it as WAN-optimized, deduplication-aware replication of data. Customers can use DASH Copy to rapidly generate off-site copies of data for disaster recovery without having to rely on expensive array-based replication. DASH Copy is also great for remote and distributed office deployments where users may want to keep a local, deduplicated copy of data at the remote site while also quickly and efficiently sending a secondary copy back to the central site for retention purposes. DASH Copy uses network bandwidth efficiently and minimizes the use of storage resources.




Figure 5: DASH Copy

Archiving

The ETERNUS CS200c features file archiving based on the OnePass™ feature.

OnePass™ is a key software component of Commvault Software and ETERNUS CS200c. The feature drastically cuts the number of redundant data protection and data management processes. Usually, all backup, archiving and storage management programs collect data separately and place them in separate storage targets. By contrast, OnePass collects, classifies and stores data on the ETERNUS CS200c in a single run.

The OnePass feature converges backup, archive and reporting with a single scan of the data to simplify administration, reduce time and resources by 50 percent and to eliminate complexity for file data in physical and virtual environments.

The software features numerous automation options that simplify archiving considerably. For example, administrators can define simple sets of rules that specify which messages must be archived and where.

The unified index gives a detailed overview of the data in the central repository and thus ensures access to random subsets of data (backups, snapshots, archives, replicated DR copies, etc.).

Based on message journaling and full-text indexing, integrated reporting and analysis tools provide comprehensive information regarding employment and utilization of the ETERNUS CS200c appliance that further help to improve service quality. These features ensure that company management and legal departments have immediate access to crucial file data in emergency situations.

Depending on the individual scenario, an archive may either reside on the same ETERNUS CS200c or on a second ETERNUS CS200c which can be specified as an archive appliance. Another option is to transfer the data from the ETERNUS CS200c to the optional adaptable tape storage system like ETERNUS LT (if the focus is mainly on long-term retention). Fujitsu Storage ETERNUS LT tape libraries are capable of reading and creating WORM media in order to prevent the subsequent, illicit alteration of source material and thus warrant legally secure archiving.

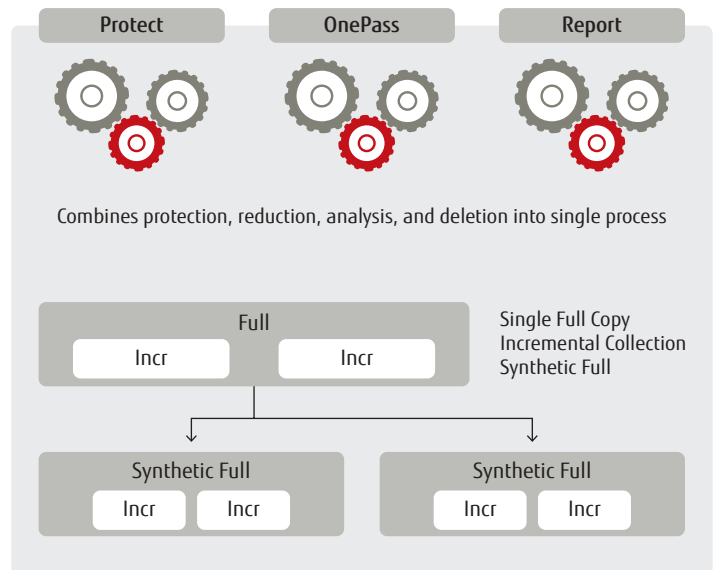


Figure 6: OnePass Feature

Backup versus Archive

A backup is not an archive. If you are still using backup storage as an archive, let's take this opportunity to differentiate between the two:

- **Backups** are secondary copies of active production information used when a recovery copy is needed to get an end user back to work or, in the case of a disaster, to get the business back up and running.
- **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 or is no longer frequently used, it is best to move it to an archive, where it resides outside the backup window, but can still be accessed.

Since backups are focused on constantly changing business information, backups are generally short term and often overwritten, say monthly, when full backups are made. Thus backup is a poor choice for retaining data for compliance purposes.

Archives do not focus on "recovering" an application or business data, but allow for 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 based on regulatory requirements.

The answer to the question of whether you need both backup and archive is likely "yes" – you need both:

- backup for disaster recovery
- archive for long-term data retention

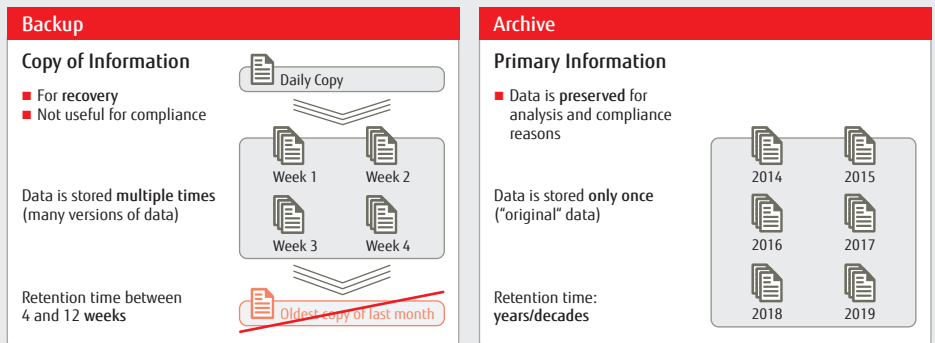


Figure 7: The difference between backup and archive

Snapshot support (Commvault IntelliSnap®)

The snapshot capability of ETERNUS CS200c automates snapshot management and application-aware recovery across the broadest selection of hardware arrays in the industry, including the Fujitsu Storage ETERNUS AF all-flash and ETERNUS DX hybrid storage arrays, without custom scripting. The feature enables faster and more efficient, granular and full recovery for applications and virtual server environments.

Snapshot technology is the fastest and most efficient way to protect and recover data and applications. But integrating snapshots and applications, and managing the process, can be very challenging. Commvault pioneered integrated snapshot and backup management across multiple disk vendors with its IntelliSnap technology. With IntelliSnap you can maximize the value of your snapshot technology while dramatically reducing management overhead and complexity.

The IntelliSnap technology automates the creation of application-aware hardware snapshot copies across a multi-vendor storage environment. The snapshot data is cataloged, which simplifies the

recovery of individual files without the need for a collection of scripts and disparate snapshot, backup, and recovery tools.

IntelliSnap supports 95 percent of all storage arrays as well as most business-critical applications, including DB2, Lotus Notes, Exchange and software from SAP and Oracle. Moreover, it cooperates with VMware's and Microsoft's virtualization platforms and various file systems for Windows, Linux and UNIX. Even more important, the module serves as a central hub between the applications and the array-specific tools, governs snapshot generation and automates snapshot management. Put another way, it serves to fully integrate snapshots into the rest of the backup process.

For details about Commvault IntelliSnap see:
[Online books Snapshot Management](#)
[IntelliSnap support Matrix](#)

Snapshot replication support for ETERNUS Storage Arrays

Commvault Software IntelliSnap supports hardware snapshots of ETERNUS all-flash and hybrid storage arrays. In addition, IntelliSnap features array-based snapshot replication for ETERNUS arrays to create more frequent and more current disaster recovery copies – with full application awareness and granular recovery capabilities. The replication feature creates a replica copy on a remote array, in addition to the primary snapshot copy on the local array. When you run a snapshot backup operation, the snapshot is created on the local array and on the remote array.

The key features of snapshot replication are:

- The primary copy and the replica copy are independent of each other, and can have separate retention policies.
- The replica copy is crash-consistent.
- Inline replication is performed only if the sub-client on which you run the snapshot operation is associated with the replica copy.

For more details go to books online:
[Inline Replication for ETERNUS Storage Arrays](#)

The setup illustrated below shows the snapshot feature of ETERNUS CS200c supporting hardware snapshots of the Fujitsu ETERNUS DX Disk Storage family. This is how the snapshot backup works:

- The iDataAgent module (integrated in ETERNUS CS200c) quiesces applications, databases and file systems on the production server and sets it to backup mode.
- File systems, applications, and virtual machines are now in a consistent state so that a snapshot of the associated storage arrays can be created. Administrators can manage all snapshots using the ETERNUS CS200c IntelliSnap feature which requires no additional array management software.
- Afterwards, the snapshot is transferred to the MediaAgent (integrated in ETERNUS CS200c), which acts as a proxy host and hands over selected snapshots to the ETERNUS CS200c integrated disk storage.

- The MediaAgent writes a catalog of the files in the snapshot in the same way it would during a regular backup. The catalog can be used for targeted information retrieval and granular restores of specific information. In addition, the MediaAgent can also perform “offline mining” operations that let users view and recover individual Exchange messages, documents or Active Directory objects.
- During office hours, snapshots can be taken on a regular basis in order to generate as many restore points as possible and meet relevant SLAs (SLA: Service Level Agreement). To prevent the loss of snapshot data (or minimize its impact), IntelliSnap automatically creates rule-based backups of selected snapshots, which are then transferred to the ETERNUS CS200c appliance.
- Since this backup copy is created by a proxy host, the process doesn't affect the production system. Consequently, administrators may induce hourly snapshots and back up the last one of the day.

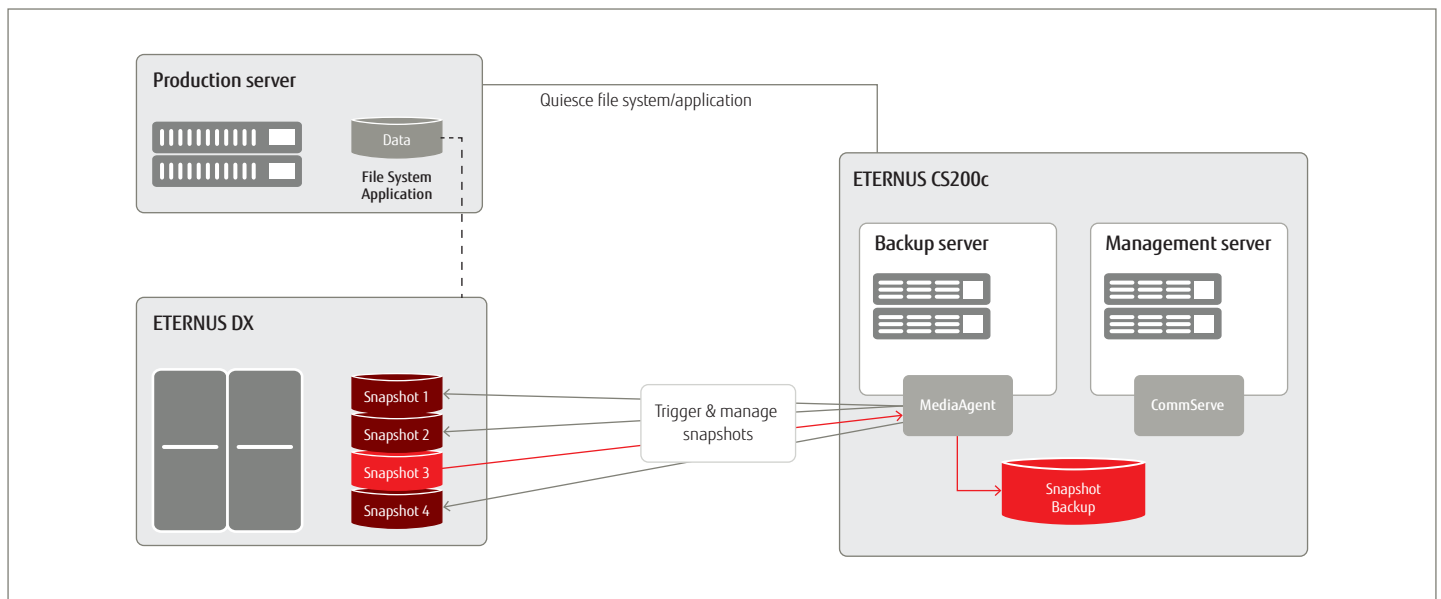


Figure 8: Snapshot-based backup with ETERNUS DX and ETERNUS CS200c with IntelliSnap

The ETERNUS CS200c IntelliSnap option also includes the IntelliSnap API, which opens up two alternate paths to data recovery from snapshot-based backups besides browsing the catalog:

- **Full Revert** – The entire snapshot is restored using the appropriate hardware features. Users should exercise extreme caution, since all files on a LUN or volume are reverted to an earlier state. IT departments planning to use this feature must configure each application carefully and ensure that there is only one database running on each volume. Full Revert is by far the fastest restore method and thus particularly suited to meet stringent Recovery Time Objective (RTO) requirements.
- **Granular Restore** – The second option is to mount the snapshot and copy back individual files or folders to the production system. The restores can either be performed over the LAN or independently (LAN-free). To use this option, the snapshot must be mounted on the production server.

Tape support

The tape attachment allocates a tape drive to the appliance providing second tier backup to tape. There is no limit to the data moved to tape. You can make as many copies as necessary and the tape capacity does not count against the disk capacity licensed in the appliance. Allocating only one drive within the Fujitsu Storage ETERNUS LT tape library to the appliance is possible. The Appliance Software is still aware of the data written on tape.

Choose the tape attachment to copy or move not frequently used, cold data from the appliance disk storage to ETERNUS LT. Tape storage provides offline backup to thwart cyberattacks and is ideal for long-term retention, archiving, and second-tier backup. Further advantages of backup-to-tape lie in the significantly lower storage costs per gigabyte and greater durability of the storage media – under appropriate conditions, tapes can protect data for decades. In addition, they can store information in WORM mode (for “write once, read many,” indicating protection against subsequent changes), are easily transferred to a second secure location and consume less power than hard disk drives, which draw current even in idle mode.

Cloud connection

Commvault Software allows customers to store their data with a cloud storage provider. The cloud storage architecture can be private, public and hybrid. The cloud connection feature of ETERNUS CS200c can be used to move backup data to a cloud data center, thus releasing licensed capacity on the ETERNUS CS200c appliance (similar to the tape attachment). The feature delivers native support for clouds supporting a RESTful interface. This currently applies to 40+ cloud storage platforms, among them Microsoft Azure, AWS, Google, Oracle and VMware clouds. The capacity for data to be transmitted to the cloud frees up the licensed capacity of an ETERNUS CS200c. Additional cost may result from using provider data storage capacity.

More information about Commvault Software and the cloud:
www.commvault.com/cloud
Online Books: Cloud Storage

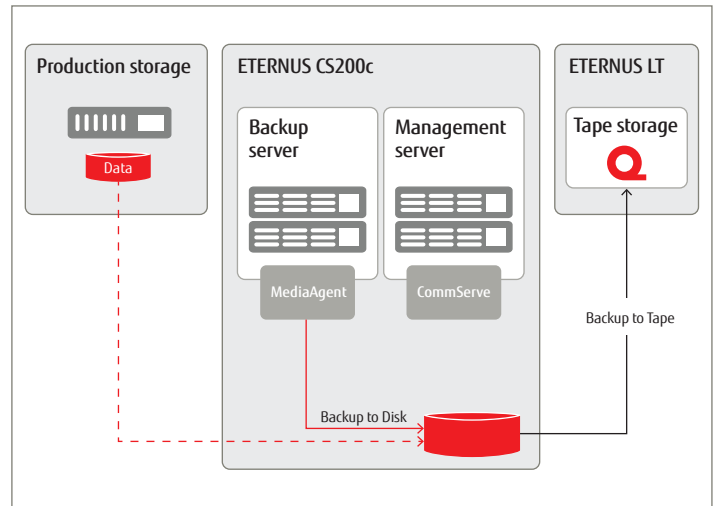


Figure 9: ETERNUS CS200c with tape attachment




Figure 10: ETERNUS CS200c with cloud connection

Conclusion

The ETERNUS CS200c is an enterprise-class backup and archiving solution with significantly reduced complexity.

By combining Fujitsu's system technology with Commvault's intelligent information management software, organizations can dramatically reduce deployment efforts, hardware and operating costs.

The ETERNUS CS200c appliance uses the mainstream functionality of Commvault Software as described in this white paper. Each environment is unique. Several licensing options and appliance models provide the solution that precisely addresses your data protection needs.

www.fujitsu.com/emeia/eternus-cs200c

▶ POWERED BY



Published by

Fujitsu Limited

Copyright 2019 FUJITSU LIMITED

All rights reserved, including intellectual property rights. Technical data subject to modifications 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. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries. Commvault and the Commvault logo are trademarks or registered trademarks owned by Commvault Systems, Inc. in the United States and/or worldwide.