

White Paper

Enterprise Software Defined Block and Object Storage

- Built for Continuous Availability -



Table of Contents	
Preface.....	3
1. Background.....	4
2. Availability by Design	5
3. Advancing to Continuous Availability Lifecycle Management (CALM)	5
4. ETERNUS DSP's CALM, How it Works	6
5. CALM Best Practices	8
6. High Availability Basics.....	8
6.1. Redundancy	9
6.2. Remediation	9
6.3. Degraded Time	9
6.4. Degraded Behavior	10
7. Software Upgrades.....	11
8. Summary and Conclusions	11

Preface

This document describes the basic architecture of the ETERNUS DSP system and provides the features of the enterprise storage system by explaining the product specifications, benefits, and tradeoffs.

The product lineup and product information stated in this document are current as of January 2020.

■ Intended Audience

This document targets the following audience:

- Those who are considering installation or replacement of storage systems
- Those who are proposing installation or replacement of storage systems

■ Applicable Model

This document covers the following model.

- FUJITSU Storage ETERNUS Data Services Platform Software

■ Naming Conventions

The following abbreviations are used in this document.

- ETERNUS DSP FUJITSU Storage ETERNUS Data Services Platform Software
- Web GUI Web GUI for ETERNUS DSP
- SSD Solid State Drive

1. Background

The legacy approach to data availability, often called high availability, is no longer sufficient in ensuring a business can extract maximum value of data over its lifecycle. Legacy approaches to high availability (HA) mandate rigorous adherence to configuration, maintenance and repair of the storage systems, and in return data will be available without disruption 99.9999% of the time. Data will be available for 100+ years if you perfectly adhere to rules set out when you installed and configured the storage system. Highly available, just don't change anything, ever.

The problem with this approach is that there are many types of data, each with unique lifecycles that are very dynamic in support of extracting value from the data. In the Data Era, over its lifecycle, data must be continuously available while the storage system entrusted with the data must simultaneously support business velocity, operational agility, technology advancement and access visibility.

ETERNUS DSP has advanced the concept of high availability via the more modern concept of **Continuous Availability Lifecycle Management** or CALM. CALM support continuous data availability while simultaneously enabling two critical business capabilities:

1. The transition of **data** through a business-value lifecycle where, at any point in time, the data can be non- disruptively optimized to match the value it brings to the business.
2. The evolution of the storage **infrastructure** through a business-value lifecycle where, at any point in time, the infrastructure can be non-disruptively optimized to match the value it brings to the business.

Distilled, data and infrastructure can be simultaneously and independently optimized, to maximize business value while providing continuous data availability.

CALM embraces velocity and agility without compromising data availability. CALM is achieved by design, not just by observation.

In this paper we will:

- Describe how we achieved CALM
- How CALM advances the industry's view on availability
- The benefits of CALM
- CALM best practices
- How ETERNUS DSP meets the legacy standard of high availability


2. Availability by Design

Building a continuously available, large scale system for the most demanding enterprise application was our founding principal, it has been a very deliberate process that continues to this day. Each day presents opportunities to learn how to improve and manage the delicate balance of conflicting requirements in the modern data center.

ETERNUS DSP set out from the very beginning to utilize a closed loop continuous improvement process consisting of six steps:

- Architecture
- Design
- Implementation
- Test
- Observation
- Analytics

It is important to point out that in the initial implementation, the steps, from architecture through design, implementation, test, observation and analytics, might be thought of as sequential. Once the product began to ship, these steps formed an ongoing process, each potentially affecting the others. Information gathered during testing, telemetry observed in production, and analytics derived from the installed base, continue to refine and optimize the architecture, design and implementation. In sum, each of these steps is an important element in delivering continuous availability, but collectively they form the essential process for delivering continuous availability and balancing many other, often conflicting, requirements for enterprise software defined storage.



Critical to ETERNUS DSP's process for continuous availability is the assumption that needs change over time and as such the system must be able to change in response to changing needs. Built for change, whether initiated by the customer, informed through installed base analytics or by adoption of new technology, ETERNUS DSP's approach ensures customers get both availability, while adapting for competitive advantage. This need for availability during change is the genesis of CALM.

3. Advancing to Continuous Availability Lifecycle Management (CALM)

The definition of High Availability of storage has not kept pace with the way data is used to extract business value. CALM has come into existence to address a business need. Historically data was associated with a single application throughout the data lifecycle. In the modern data-centric data center the lifecycle results in data being associated with many applications: real-time analytics, machine learning training and batch analytics are modern examples in addition to existing uses business processes, reporting, data protection etc.

A key aspect of the data lifecycle is that there is a long tail to the business value of data. This long tail mandates that the data and the infrastructure be decoupled so that as the infrastructure on which the data is currently stored becomes obsolete the data can be non-disruptively moved to modern infrastructure that is commensurate with the data's business value.

Finally, the need for businesses to move quickly to create and sustain competitive advantage, termed business velocity, mandates that businesses make decisions with less planning. The consequence of quick decision making is the need to refine after the decision is implemented. For the storage infrastructure, enabling business velocity also means that the way data is stored, even early in the lifecycle, may not be the best match to the data's business value. The storage infrastructure must support changing the cost, performance, protection and distribution in all phases of the data lifecycle.


Continuous Availability Lifecycle Management (CALM) is a major advancement in how storage can contribute business value when more and more value is coming from effectively exploiting data assets. The goal of CALM is to enable the extraction of value from data by decoupling the data lifecycle from the infrastructure lifecycle while providing continuous data availability from the infrastructure. CALM help both IT operations and business stakeholders, those extracting value from data, to remain calm by enabling independence of IT operations and data value lifecycle.

4. ETERNUS DSP's CALM, How it Works

ETERNUS DSP's innovative architecture simultaneously enables business velocity without regret and infrastructure optimization without compromise. In this section the holistically engineered set of mechanisms that enable CALM will be described from the outside in.

Like all storage systems, ETERNUS DSP's architecture consists of a control plane, data plane, and the physical infrastructure. Where ETERNUS DSP innovates is in how these architectural components form a closed loop.

The control plane continuously analyzes and adjusts the operations against the objectives of both the infrastructure and data stakeholders.



In ETERNUS DSP's architecture, the roles of the key components are as follows:

- **Control Plane** - optimize replication, distribution and management of data based on the current policies specified by the tenants. Utilize the telemetry from the infrastructure and cloud based analytics to optimize, data placement, resource utilization and respond to external requests to manage the system (create, remove, scale,...)
- **Data Plane** - Efficiently (i.e. with performance) respond to data access requests via the supported protocols and access data as specified by the policies provided by the control plane. Distribute data utilizing the time- based coherency protocol.
- **Commodity Infrastructure, Extent Store** - Respond to read/write requests from the data plane. Optimize accesses for media types, CPU and memory resources. Perform data management functionality (encrypt, compress, dedup, snapshot) as specified in the metadata provided by the control plane and communicated by the data plane. Collect telemetry on the I/O and the infrastructure and communicate to the control plane.

In the ETERNUS DSP architecture the control plane plays a larger role by not only providing management interfaces, it is responsible for continually assessing and balancing the policies associated with the data and the consumption of resources. This closed loop process is a key enabler to delivering CALM.

In the ETERNUS DSP software there are a number of key mechanisms that enable CALM. The following list of mechanisms, and the contribution to CALM, is presented in the approximate order they relate to in the closed loop architecture starting with the way users interact with the system.

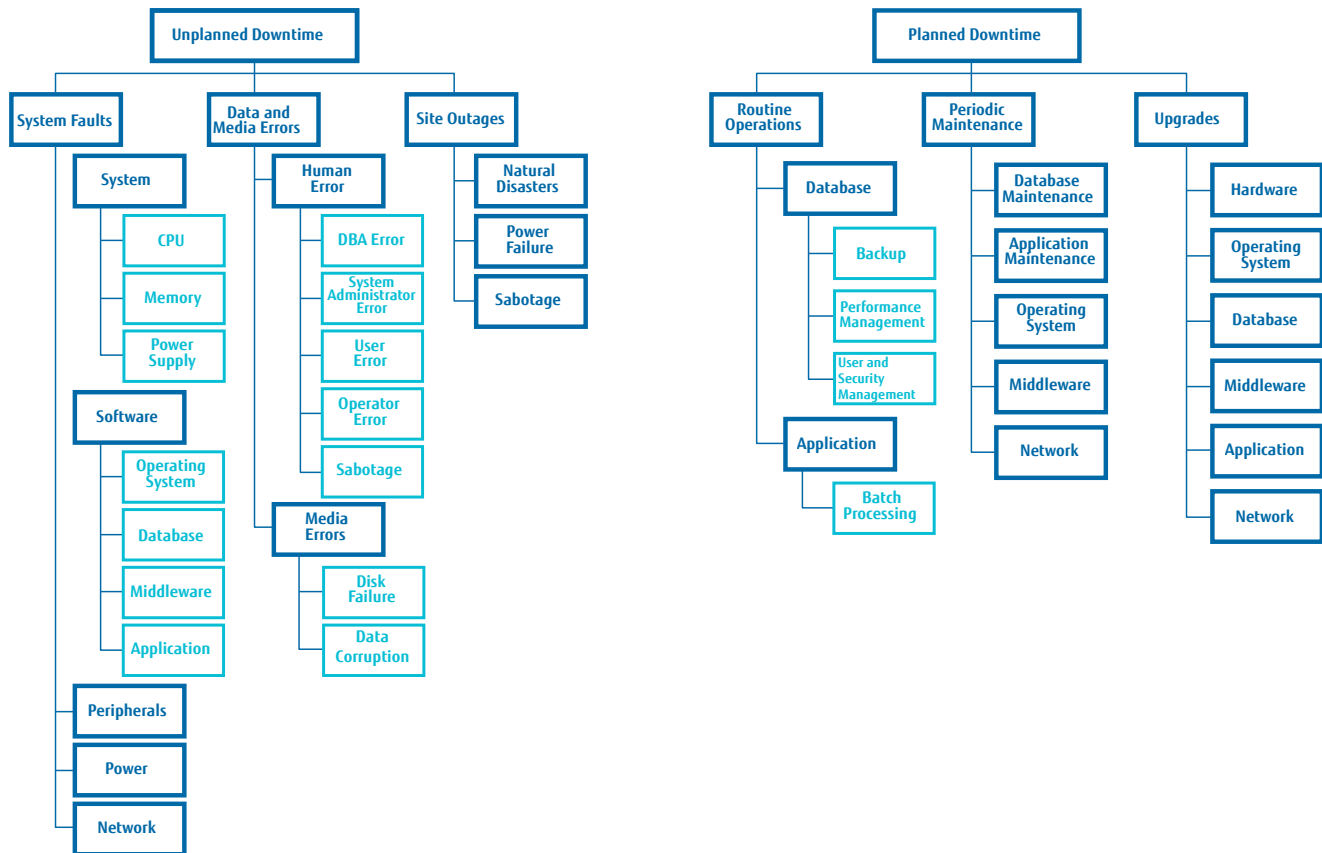
- **Application Templates with Ecosystem integration** - The preferred method of creating and managing storage is via application templates programmed via the control plane. When an application needs storage the templates describe the properties of the storage including one or more volumes and the policies for data management, replication, fault domains, performance and tenancy. The templates are integrated with orchestration ecosystems (e.g. Kubernetes, OpenStack,...) so all storage can be provisioned as part of the application provisioning. The benefit of this approach is that it raises the level of abstraction for configuring and managing storage. This reduces the probability of misconfiguration and enables application owners to provision and manage storage based on their application objectives without having to be storage experts.
- **API First** - ETERNUS DSP support several different types of interfaces including ReSTful programmable interfaces, Command Line and Graphical User interfaces. The benefit of an API first approach is that all actions manifesting from the interfaces resolve to a single set of APIs within the ETERNUS DSP software to insure consistent behavior across interface types.
- **Applications Instances with Policy Based Automation** - In the control plane, storage for an application is managed as an application instance with associated policies to a service level objective (SLO). The control plane monitors performance against the SLO and automatically adjust to deliver to the SLO. The SLO can be changed at any time via the templates and / or APIs and the control plane will automatically adjust in response to changes. The benefit of this approach is that it lets the user specify their intent without having to configure the system to achieve the intent. The ability to specify intent and change the intent increases the level of interaction with the system, reducing the instances of human error-based issues.
- **Current / Future Maps** - ETERNUS DSP utilizes a map-based approach to data placement in contrast to systems that use a hash based (calculation) approach. ETERNUS DSP extends this concept to utilize two maps, a current map that describes the steady-state placement of data and a future map that describes forward looking data placement when the system encounters a need to change. Changes to the maps occur in response to changes to SLOs, to better meet existing SLO's, failure remediation, growing and shrinking the system. The benefit of this mapping approach is that the system can change data placement at any time. Furthermore, this pervasive mechanism provides an elegant and well tested methodology for dealing with change in contrast to other systems that are difficult or risky to change or execute little tested code paths for error handling.
- **Scale Out, Shared Nothing with Asymmetric Nodes** - ETERNUS DSP is utilizing a modern scaling approach where nodes can be added incrementally as needed. ETERNUS DSP extends this concept by allowing nodes with different capabilities, CPU, memory, storage, networking, to be added. The control plane will automatically recognize and take advantage of whatever resources are added. The benefit of this approach becomes apparent when coupled with the way the control plane manages via SLO and can provide new placements maps at any time. The asymmetry can be exploited to adopt new technology, improve performance of running applications, add capacity in anticipation of applications enhancements or data growth and to take advantage of cost saving opportunities. No forklift upgrades.

- **Time-based Coherency for Data Distribution** - As writes are received into the data plane from applications, the data is time-stamped and distributed to the cluster nodes in accordance with the current and future placement maps provided by the control plane. Utilizing time enables cluster nodes to guarantee ordering and versioning of data with minimal communication in comparison to systems that use locking protocols. The benefit of this approach is a very low latency data path with guaranteed accuracy and forward progress. Utilizing time for synchronization also will allow any node to serve any data, true active / active data access across a cluster that may be widely distributed. This approach enables cluster nodes to be distributed across fault domains for better large scope failure resiliency without performance compromise.
- **Floating Data and Metadata** - In the ETERNUS DSP system neither data or metadata is tethered to any node in the cluster. Each node can perform any cluster function and all nodes have a copy of all cluster metadata. Access to any data in the cluster can move among the nodes and the data, based on the associated policies can be stored on any node and migrated at any time. The benefit of this approach is that nodes can be added, serviced, upgraded or removed and any services performed by the node and any data stored can be moved as needed automatically and without error prone user intervention.
- **Stateless Extent Store** - The extent store is the software component of the ETERNUS DSP system responsible for persisting and managing the data on the physical infrastructure. Any node in the cluster that stores data runs this software component but the instances of this software component on each data node in the cluster are stateless relative to one another. This independence enables local optimizations for the type of physical infrastructure but more importantly the independent operation enables linear scaling without increased communication. Separation of concerns is a critical element in achieving CALM at large scale when other systems have increasing communication and risk as the scale increases.
- **Writes to Non-Volatile Memory** - Data nodes have non-volatile memory that is used to accept data written into the system and respond to the application without having to store the data to the storage devices. The benefit of non-volatile memory is that write performance is similar across different types of nodes and very low write latency does not require hyper-optimization of the data path, a historical source of storage system- based data corruption and loss.
- **Data Management and Log Structured Writes** - When data is destaged from non-volatile memory, data management functions such as snapshots, deduplication, compression and encryption are performed as background operations. Data is then placed into logs sized to optimize storage device performance and longevity. Additional data integrity information is also added. The benefit of this approach is that the coalescing of writes reduces stress and increases longevity of the underlying storage media and the additional data integrity information allows remediation of silent data corruptions from the media. Reducing device failure rates and detecting silent data corruption are material contributors to high availability.
- **Telemetry Gathering** - As data flows through the system and various management activities occur, telemetry data regarding resource usage statistics, activity logging, performance statistics and physical telemetry such as temperature, fan speed etc. are gathered. The benefit of telemetry is that it can be used to assess the operation of the system both locally and globally and having this information can be utilized in the future for purposes not defined when the information was saved.
- **Cloud Based Analytics and Artificial Intelligence** - One use of the telemetry data is in the cloud-based analytics and user portal. The telemetry data is analyzed and visualized for the user. This information can be utilized to better understand how the system is operating for each tenant, how the system is operating relative to service level objectives and how the system is performance relative to the ETERNUS DSP installed base. The benefit of this approach is to allow users to visualize and manage their system(s) and pinpoint how the system can be optimized to better serve the data or how the policies surrounding the data can be tuned to better utilize the system resources.
- **Closing the Loop** - The telemetry information communicated to the control plane closes the loop to form the self-optimizing intelligent and automated ETERNUS DSP system. Starting with the application templates that capture the intent of the applications storage needs, this gets converted into internal policies that drive service level objectives. The SLO's enable the control plane to create current maps that specify data distribution and when change occurs, future maps to plot out the future state of the system. Embracing asymmetry in the scale-out cluster allows for a variety of application needs as well as the adoption of new technology and the retirement of aging infrastructure. The time-based coherency protocol accelerates performance through minimizing cluster communications while ensuring data correctness even when the cluster is distributed across multiple fault domains. Floating data and metadata ensure no single node in the system becomes a risk while the stateless extent stores allow the system to provide linear scaling while enabling per-node optimization. Non-volatile memory insure that write latencies are low and enables data management functions to be performed without impacting performance. Log structured writes ensure maximum performance and longevity from the underlying storage media along with an extra measure of data integrity checking to prevent silent data corruption. Telemetry gathering insure both the cluster control plane can monitor the operation of the system and optimize to meet the SLO while also enabling cloud-based analytics to guide the user to optimal operation of both the infrastructure and the data. All of these mechanisms form a robust, performant, self-optimizing and continually available system to enable customers to extract value from their data while reducing the burden of managing the infrastructure. This is the delivery of Continuous Availability Lifecycle Management - CALM.

5. CALM Best Practices

This section would describe best practices in system implementation and management as they relate to continuous availability.

6. High Availability Basics



A subset of continuous availability is what historically would be called high availability. As a basic requirement all vendors are expected to meet, high availability itself is no longer a distinguishing capability in the market. That said, different products may go about achieving high availability in different ways and with different tradeoffs. Some are expensive, some intrusive, while others have large impacts in performance... not all systems deliver the basics in the same way.

Availability is one of the many "ilities". Reliability, manageability, usability, affordability, supportability, scalability, serviceability... there is a long list of "ilities" that often are in tension with one another. For the purpose of this paper we obviously will center on availability, and when appropriate other "ilities" may get mentioned if they contribute to or conflict with delivering continuous data availability.

ETERNUS DSP achieves the requisite number of "9's," the percentage of time data is available to applications, using some common techniques along with some key innovations.

Nearly all primary storage systems, including ETERNUS DSP, have a simple prioritization: Data Integrity, Data Availability and then all their Functionality (performance, efficiency, data management...). In this hierarchy Data Availability is the second priority behind data integrity. The reason this is mentioned is that ETERNUS DSP also has innovations in enhancing data integrity, ensuring data availability.

Many PhD thesis have been written on high availability but with more than 30 years working on enterprise storage including more than 20 different product architecture, the author has distilled High Availability down to four key factors:

- **Redundancy** - provisioning of additional resources to allow resource failures to occur while keeping data available.
- **Remediation** - the detection of failures and execution of processes to restore availability to an acceptable, ideally pre-failure, state.
- **Degraded Time** - the duration from the failure occurrence until remediation is complete.
- **Degraded Behavior** - Changes in externally visible behavior during degraded time.

Put simply, when a failure occurs remediation should minimize degraded time and degraded behavior without excessive redundancy.

Following is how ETERNUS DSP delivers high availability utilizing the four key factors:

6.1. Redundancy

ETERNUS DSP implements a scale-out shared-nothing cluster architecture and does fully distributed data placement via replication. Customers can choose the number of replicas from 1-5 with increasing replicas improving availability as well as performance and other attributes. Replicas can be specified to be distributed across fault domains to manage failure scope. Full distribution of data ensures that replicas are placed across as many nodes as possible. This distribution also minimizes the scope of impact of failures. ETERNUS DSP requires redundant networking and best practices specify separate networking for initiator and cluster traffic. ETERNUS DSP can utilize a wide variety of standard x86 servers, currently all of which have redundant and hot swappable power and cooling as well as hot pluggable storage devices.

6.2. Remediation

When failure(s) occur the processes for recovery will vary depending on the nature of the failure(s). In all cases the user will be notified of the failure and any corrective action.

- **Remediation of Failed Storage Devices** - In the event a storage device that holds user data fails, the ETERNUS DSP software will automatically create new placement maps for the data associated with the failed drives to all nodes in the cluster. Any data that has been written to the thinly provisioned volumes will be redistributed on other drives in the system restoring the system to its previous redundancy level.
- **Remediation of Failed Servers Hosting The ETERNUS DSP Software** - In the event a server hosting the ETERNUS DSP software fails, any responsibility for serving IO will be shifted to surviving cluster nodes. Due to the amount of data held by a server and the resulting long redistribution time, it is recommended that the server be repaired and brought back into service. Once back in service, any changes to data stored on the restored server will be merged back in when the server is restored. In the event the server will not be brought back into service, the ETERNUS DSP software will operate in a manner similar to a failed drive when the surviving nodes will be provided new placement maps and data will be reconstructed on the surviving nodes to bring the data back into its previous redundancy levels.
- **Remediation of ETERNUS DSP Software Failures** - In the event of a software failure, if the scope of the failure is limited to a micro-service then the micro-service will be restarted, and operation will continue. If the software failure is larger in scope, impacting the server, responsibility for serving IO along with any cluster wide responsibilities (i.e. Control Plane master), will be shifted to surviving nodes and the node will be rebooted. If rebooting is successful, the node will be joined back into the cluster. If a reboot is not successful, the node will be deemed to have failed.
- **Remediation of Failed Network Infrastructure** - ETERNUS DSP utilizes standard networking as the interconnect to the hosts that are accessing data as well as for intra and inter cluster communications. ETERNUS DSP depends on a reliable network for availability but the network is not part of the ETERNUS DSP system. Alerts will occur if networking anomalies are detected that impact the ETERNUS DSP software and it is expected that the network will be brought back into operation in compliance with ETERNUS DSP networking requirements.
- **Remediation of Infrastructure Sub-Components** - Some failures, more specifically, fans, power supplies and other components where the failed component is a redundant sub-component of a server or switch, are less urgent to repair as subsequent failures will not cause an outage but may result in the failure of the larger component where the redundant sub-component is installed. ETERNUS DSP is designed to handle failures of the servers and switches so sub-components in these components should be replaced at a reasonable service opportunity but have no effect on user visible behavior and the components will operate indefinitely with a single failed sub-component.

Remediation of failures at a high level is conceptually straightforward. Actual implementations and where ETERNUS DSP differentiates itself will be covered in other sections of this paper. Suffice it to say that a general pattern of detect, notify, and re-mediate exists in all products.

6.3. Degraded Time

The time a product spends in a degraded state, i.e. during remediation, is the largest variable in achieving high availability. During remediation the system is vulnerable to additional failures resulting in a loss of availability or a loss of data. This section will describe key capabilities that reduce Degraded Time.

Full Mesh Rebuild for Storage Device Failure Recovery: In storage systems the most common long-running remediation process is the recovery from a failed storage device. With the increase in capacity of storage devices, device failures is the largest factor in achieving high availability.

ETERNUS DSP fully exploits the scale-out nature of the architecture to bring available resources to bear in accelerating remediation. When a storage device failure occurs the ETERNUS DSP control plane will build new placement maps for redistributing the data across the cluster. These new placement maps tell all nodes where data will be redistributed to restore the data to the previous level of redundancy. Redistribution is then delegated to all nodes in the cluster, termed a full mesh rebuild, and each node rebuilds a portion of the failed devices data.

Full mesh rebuild accelerates device failure recovery in the following ways:

- Placement maps are thinly provisioned so only data that has been written is rebuilt.
- All nodes participate in recovery, accelerating rebuilds commensurate with the size of the cluster. This contrasts with many other products that have a single coordinator of the rebuild that then becomes a chokepoint.
- All suitable devices are targets for the rebuilt data. This contrasts with many other products that may rebuild to a single device or a single node, which then becomes the chokepoint of the rebuild.
- Replication enables simple copying of data, a single I/O, to enable redistribution. This is in contrast to many other products where erasure coding can create up to 10X or more I/Os to rebuild the same data.

Full mesh rebuild, can reduce rebuild time by orders of magnitude with a commensurate benefit to reducing data availability risk associated with additional failures during degraded time.

Any Cluster Node Can Run Any ETERNUS DSP Micro-Service - In an ETERNUS DSP cluster all nodes have all the node level metadata for the entire cluster. This provides both redundancy of the metadata but more importantly, in terms of degraded time, a server or node failure results in other nodes quickly taking over the responsibility for a degraded node.

Fast Reboot - ETERNUS DSP is built utilizing a custom Linux kernel. The benefit to this is that unneeded services can be eliminated as well as being able to create a kernel that support extremely fast reboots compared to off-the-shelf distribution. This fast reboot enables much short-degraded time, orders of magnitude compared to off-the-shelf distributions, so that nodes can be brought booted or rebooted very quickly for error recovery, cluster additions and software upgrades.

6.4. Degraded Behavior

One of the largest problems plaguing the storage industry over the years is that when a failure happens, the behavior of the storage system during the remediation process and time can cause failures elsewhere in the system. Long failovers, large performance impacts, or inaccessibility of management interface as examples of behaviors that can cause upstream applications and operating systems to consider the storage system to have failed. This upstream failure often occurs even though the storage system is available, at least from its own perspective.

ETERNUS DSP is designed to minimize the impact during remediation in the following ways:

- Scale-out is implemented such that the performance and resiliency impact is commensurate with the proportion of the resources that are degraded. For example, in a 10-node cluster, if one node is degraded, 90% of the cluster resources are still available.
- During device recovery the full mesh rebuild distributes the work of rebuilding the data. This is in comparison to alternative approaches where a single node, controller or device becomes the bottleneck and services provided by that node can be severely impacted.
- Because any node can run any service, reassignment of responsibilities within the cluster occur both quickly and seamlessly in regard to any external behavior associated with the service.
- Floating IP's decouple IP addresses from node. This enables both data and management services to migrate among nodes in the cluster without requiring network reconfiguration.
- Many failures are automatically handled to restore redundancy to previous level. Sometimes called lazy repair, this enables component replacements to be scheduled at a time it is convenient to administrators. This contrasts with systems where failed components must be replaced in order to restore the system to an acceptable state.
- Distributing data across all suitable nodes and devices in a cluster means that when a failure occurs in a node or device there are multiple other paths from which the data can be stored and retrieved. This is in contrast to dual-redundant systems where in the event of a failure, I/O is concentrated to the surviving node (in the event of a node failure) or a node that is responsible for rebuilding data for a failed device may have a large, externally visible impact on I/O rates and response times.
- ETERNUS DSP uses variable resource consumption during remediation. ETERNUS DSP monitors the level of activity in the nodes and cluster and will throttle remediation activities to ensure applications are being served.
- During remediation, if resource consumption is high, the system will defer resource intensive data management functions, such as compression and deduplication. When resource consumption returns to an acceptable state, these data management functions will then be performed.

A key aspect of enterprise storage is predictable behavior during remediation. Historically enterprise storage systems achieved this through massive redundancy with the resulting high costs. ETERNUS DSP's scale-out architecture, key innovation such as full mesh rebuilds, along with careful attention to resource utilization, can deliver enterprise storage capabilities during both normal and remediation processes.

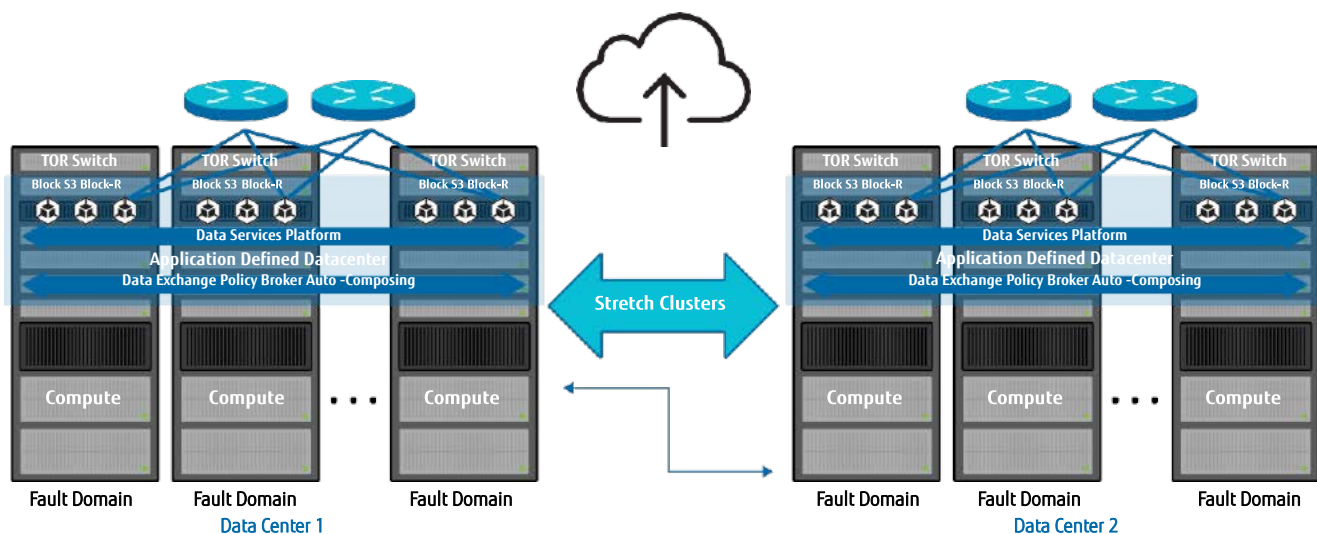
7. Software Upgrades

Another attribute of enterprise storage system is the ability to upgrade any necessary software components in the system while the system and the data remain available. To meet this requirement ETERNUS DSP implements rolling upgrades. Upgrades include all ETERNUS DSP software components as well as any software / firmware for the infrastructure. No user intervention is required for upgrades, all nodes will be upgraded as a single operation.

ETERNUS DSP's approach to high availability is ideal: remediation is fast, changes in behavior or minimal and all of this is achieved with minimal redundancy.

8. Summary and Conclusions

The power of software defined, shared nothing ETERNUS DSP multi-node clusters enables customers the ability to easily architect for continuous data availability. Using multiple replica counts, fault domains, stretch clustering and L2/L3 networking, allows the infrastructure to withstand loss of multiple racks, or geo-locations, with no loss of data availability.



FLAT NETWORK

- All hardware is available "as-a-service"
- Composable & scalable resource pools
- Eliminate complex and costly migration

FLAT IP NAMESPACE

- Policy defined L3 virtualization
- Floating nodes and services (ports)
- Flexible data placement & live data mobility

CONTINUOUS AVAILABILITY

- ETERNUS DSP survives MULTIPLE full rack failures to deliver 24x7/365 lights-out service
- Snapshot S3 (and other public clouds)

Unlike legacy, dual controller, scale-up arrays, ETERNUS DSP scale-out, policy and application driven architecture delivers significant innovation to enable customers the ability to achieve a high performing and continuously available data infrastructure. It limits impacts and risks to performance during hardware failures, lets you avoid costly downtime for upgrades, maintenance and the most troublesome for legacy arrays, migration. Simply continue to add heterogeneous media and server nodes to your ETERNUS DSP cluster, to increase performance, capacity AND availability. Simple, with no additional software or appliances to purchase. ETERNUS DSP CALM truly delivers the broadest architectural model for achieving a continuous available data infrastructure at significantly better economics.

Copyright Information

Copyright 2020 FUJITSU LIMITED. All rights reserved.

No part of this document covered by copyright may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system - without prior written permission of the copyright owner.

Software derived from copyrighted Fujitsu material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY FUJITSU "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL FUJITSU BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Fujitsu reserves the right to change any products described herein at any time, and without notice. Fujitsu assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Fujitsu. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Fujitsu.

Contact

FUJITSU LIMITED

Website: <https://www.fujitsu.com/eternus/>

■ Trademarks

Third-party trademark information related to this product is available at:

<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

■ Disclaimer

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. FUJITSU LIMITED is not responsible for any damage or indemnity that might be caused by the content in this document.