

# FUJITSU Storage ETERNUS AF/DX

## ETERNUS Storage Cluster

### White Paper

**Introduction to the concept, technology and setup of the ETERNUS AF/DX transparent failover solution for optimal business continuity**

---

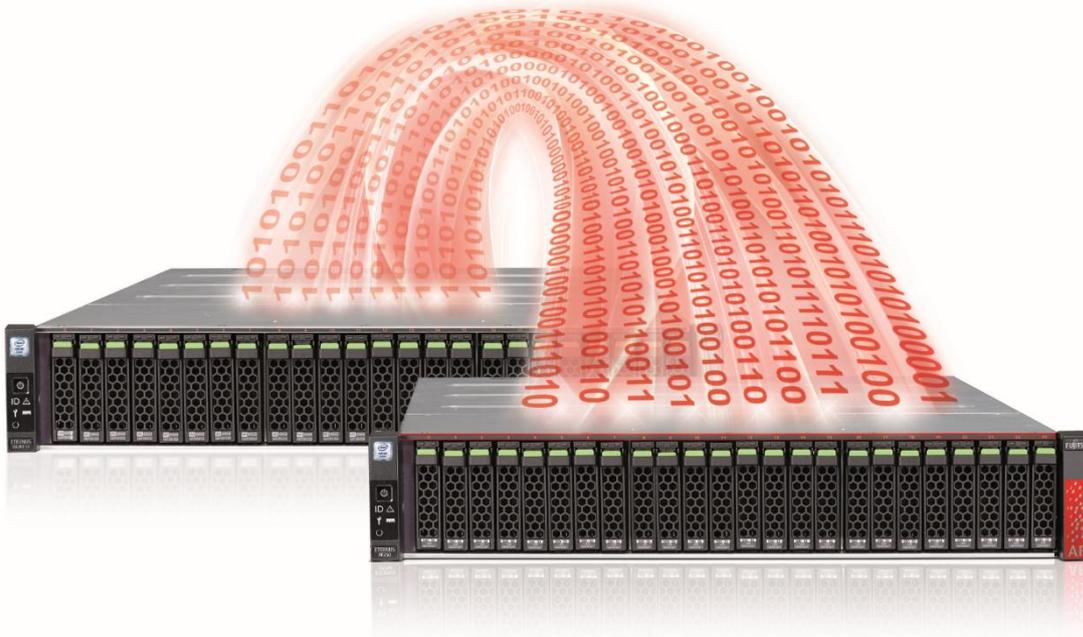
#### **Content**

Management summary .....	2
Business continuity, sever virtualization and storage.....	3
Basic architecture .....	4
Storage cluster features and benefits .....	5
Storage cluster in detail.....	7
Storage cluster in virtualized server environments.....	11
Conclusion.....	12

## Management summary

As data is one of the most important assets in any company, storage systems play a vital role in the IT infrastructures of all enterprises. IT administrators need to ensure that mission-critical data is continuously accessible, ensuring round-the-clock business continuity despite any system or site failures, as well as unplanned or planned outages. Fujitsu's ETERNUS online storage systems, with its Business-centric Storage approach, address this issue with a unique transparent application and server failover solution – storage cluster. storage cluster is the high-availability feature of the ETERNUS AF and ETERNUS DX S3 storage arrays. Data is synchronously mirrored between two interlinked storage systems. If the primary

(active) system should fail, all primary host connections are switched instantly to the secondary (standby) system. This failover is transparent for both servers and applications, and it ensures uninterrupted operations. In addition, the failover can be executed in both directions and between different ETERNUS AF all-flash models and ETERNUS DX, thus supporting non-stop operations very efficiently. This advanced feature helps fulfill service levels and delivers predictable operation for businesscritical applications – particularly in virtualized server environments. Storage cluster is integrated and controlled by the ETERNUS SF Management Suite and features simple setup and administration functionalities.

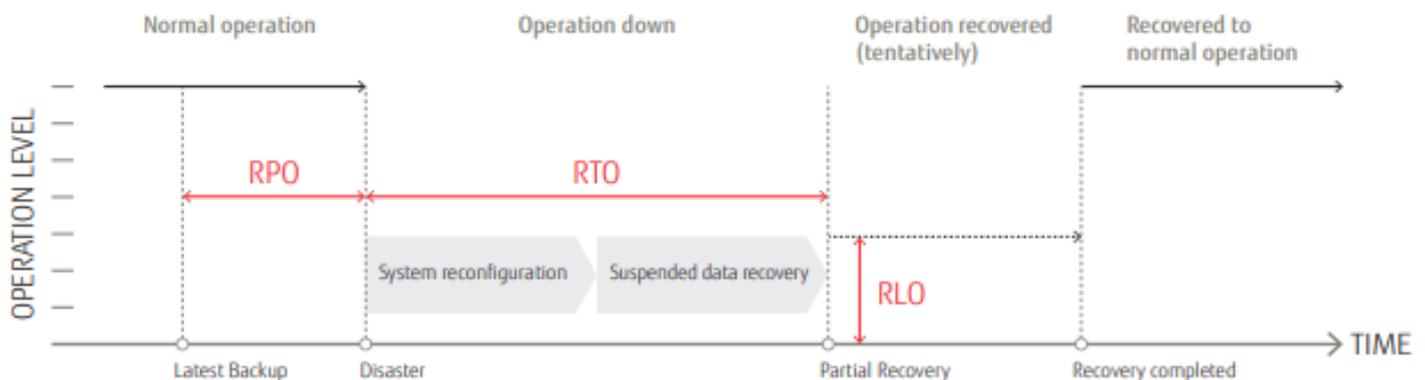


# Business continuity, sever virtualization and storage

A business continuity plan is built on the 3R fundamental objectives: RLO, RTO and RPO.

All objectives should be set for each business application based on

- The Recovery Level Objective (RLO) is defined according to the priority of the business application – it defines the granularity with which the recovered data must be enabled – i.e., whether it is necessary to recover a whole database, a web application, a site collection, a site, a list or library, or just an item.
  - Based on the RLO, acceptable values are defined for both Recovery Point Objective (RPO) and Recovery Time Objective (RTO)
- the expected loss for the business with the objective
  - the cost of achieving the objective



## RTO (Recovery Time Objective)

Metric for the amount of time it takes to recover from a data loss event, and how long it takes to return to service.

## RPO (Recovery Point Objective)

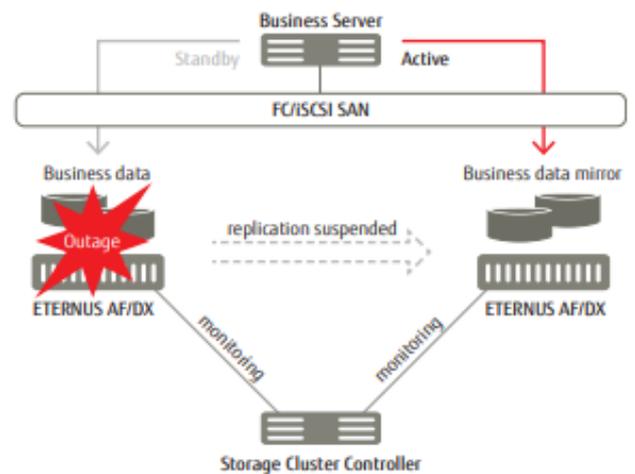
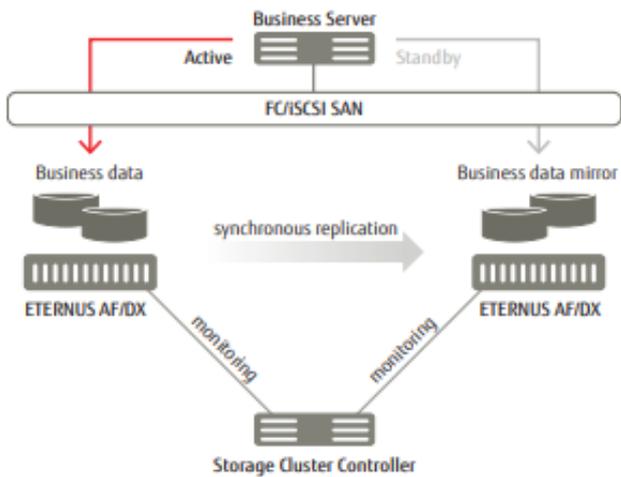
Metric for the amount of data at risk of being lost during a recovery.

## RLO (Recovery Level Objective)

Granularity level of the recovered data.

## Basic architecture

The concept of storage cluster is based on the deployment of a secondary storage system and a supervising instance – the storage cluster controller.



Storage cluster can be deployed in Fibre Channel (FC) and iSCSI environments. Thus organizations can leverage existing skills, because Ethernet is common standard and fast enough for most organizations.

As long as the primary storage system is running, data is transferred from it to the secondary system via a synchronous replication function. The storage cluster controller continuously checks the status of the primary storage. If a failure is detected, it runs the failover logic, and the primary storage information (e.g., LUN ID/WWN, or with iSCSI: IP addresses/port settings) is shifted over to the secondary storage in order to transparently identify the volume via the server I/O. Hence, operations run smoothly and ensure business continuity.

Due to the restrictions on synchronous mirroring between the storage systems, storage cluster can be deployed in building, campus and metro environments. For regional or global business continuity and disaster recovery scenarios, ETERNUS AF/DX can be deployed with asynchronous remote mirroring functionality resulting in increased RPO and RTO, and without automated and transparent failover.

Please note: In environments where iSCSI configurations are used for the host connection and the copy path, switching storage systems for a failover or failback requires approximately 30 to 120 seconds which is more time than for Fibre Channel (FC) configurations. Therefore, unlike FC configurations, a failover might not be performed transparently and the business server may be aware of the operation. For iSCSI configurations, please refer to "FUJITSU Storage ETERNUS DX, ETERNUS AF – storage cluster Function Guide".

## Storage cluster features and benefits

storage cluster technology is the foundation for a non-disruptive business continuity solution. By transferring data in real time and switching storage resources automatically and transparently for servers and applications, data is protected 100 percent and the business process is safe from any interruption.

- **Range of application:** Storage cluster is suitable for planned downtime use cases and different types of disaster on regional-level, particularly device failures, floor-level or building disasters. For wide-level disasters a disaster recovery solution is suitable



- Data consistency is managed by the storage systems without the need for specific volume management software. This frees up server resources while the secondary storage is securely synced with consistent data sets.
- Data is replicated along with the automated storage tiering (applies only to clusters with ETERNUS DX storage systems) layout and auto QoS settings.
- Failover to the standby resources is triggered automatically – there is no need for any human intervention.
- For planned downtime (e.g. for upgrades), manual failover, failover test or restoration after failover (failback), the failover operation can also be triggered from the user interface – it is simple and secure.
- Seamless server affinity – no change in access path, drive letter or mount point information after the failover. Volumes continue to be accessed transparently from the business application. No workload for volume reallocation or remount is needed or generated.
- After the recovery of the failed system, storage cluster automatically executes the failback to the initial state. System operation is always stable and secure.

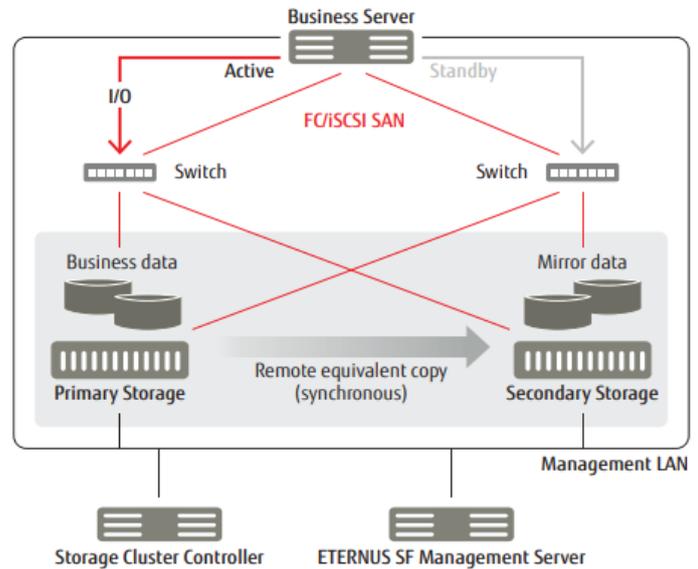
## Limitations

- Maximum distance between the sites is limited by the roundtrip time for the synchronous replication which must not exceed 10 milliseconds.
- To avoid misbehavior in case of complete site outages it is recommended to deploy the storage cluster controller and the ETERNUS SF Management Server on an independent third site.
- A redundant multipath configuration with at least one path to each storage controller is recommended.
- Storage cluster controller can be run as virtual machine without needing dedicated hardware.

# Storage cluster in detail

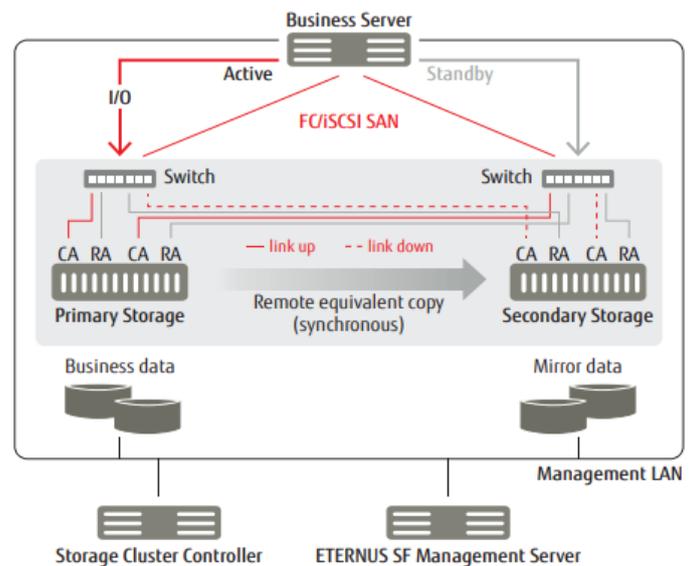
## Primary and secondary storage

The primary and secondary ETERNUS AF/DX storage array should preferably be located in different fire compartments – even better in different buildings or in metropolitan dispersion. storage cluster is set up using Transparent Failover Volumes (TFOV) which are part of a special copy group – the TFO group. Layout of the TFOVs and TFO groups are identical in the primary and secondary storage, even including the configuration settings for automated storage tiering, snapshots, etc. TFOVs are synchronously replicated from the primary to the secondary array.



## Channel adapter (CA) ports

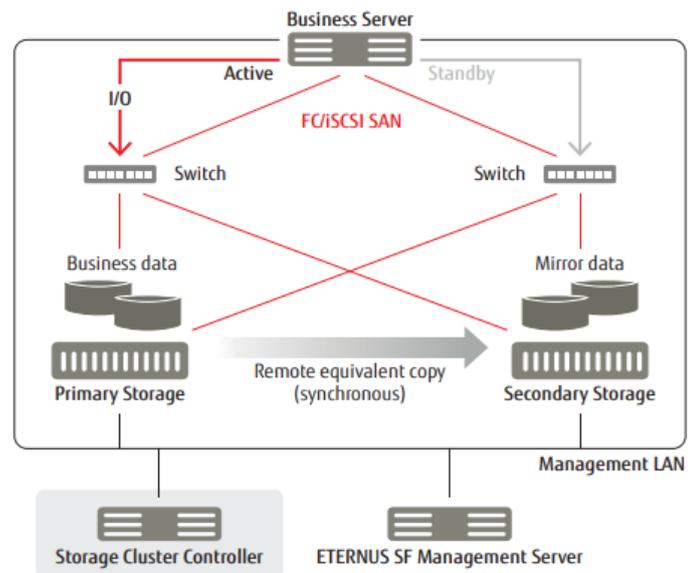
Primary and secondary storage have paired CA ports. Paired means that the ports on both sites have the same identity regarding WWN/ WWPN or IP address. Under normal conditions the CA port on the primary site is in "link up" status, and the port on the secondary site is in "link down" status, so all server I/O is processed to the primary storage. The CA port states, as well as the REC state, are controlled by the storage cluster feature.



### Storage cluster controller

The storage cluster controller is a server or virtual machine including an agent to connect with the ETERNUS SF management software. It monitors the health of primary and secondary storage in order to detect outages of the active system.

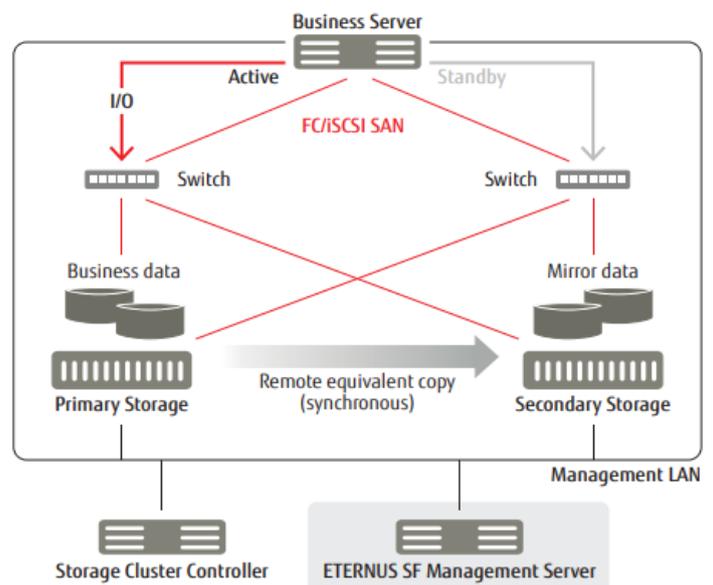
The storage cluster controller triggers the automatic failover in this scenario only and is not involved in cases of administrator-triggered manual failover or in cases of automated failover caused by RAID failures.



### ETERNUS SF management server

ETERNUS SF management is the prerequisite for setting up the storage cluster configuration with regard to TFO groups, TFOVs, copy groups and REC pairs. It also executes the failover and failback operations, either triggered by the storage cluster controller in cases of automatic failover, or manually by an operator. It also executes the automated failover in cases of RAID failures on the primary array.

ETERNUS SF and the storage cluster controller can be installed on the same physical or virtual server.



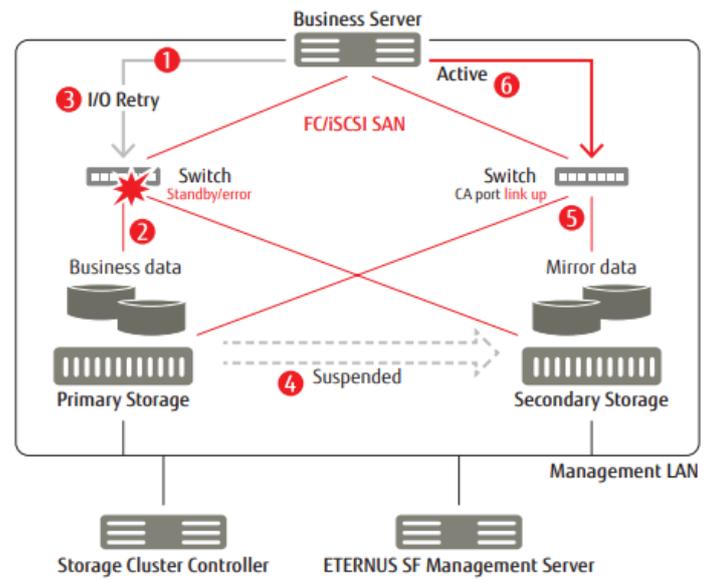
### Failover mechanism

Storage cluster reroutes I/O access from one array to the other as seen in this figure

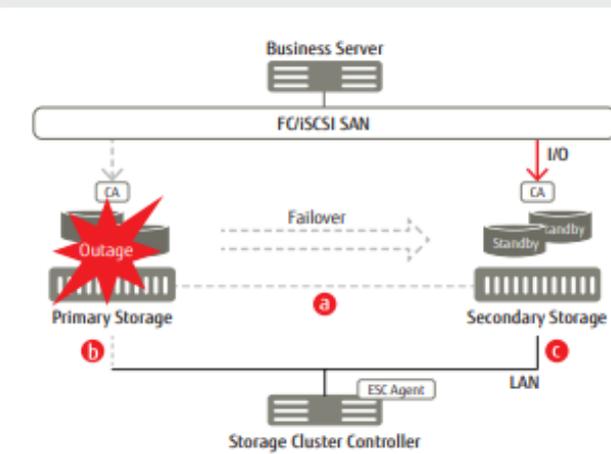
If an outage occurs the failover sequence is executed as follows:

- 1 The server sends I/O requests to the primary storage.
- 2 The primary CA port does not respond, the storage cluster controller detects the primary ETERNUS is unreachable and reports it to ETERNUS SF.
- 3 The server retries the I/O after a preset time-out
- 4 ETERNUS SF suspends the remote mirroring (REC) session, the replicated data becomes the actual business data.
- 5 The CA port on the secondary array is activated (link up) with the same identity (WWN/WWPN or IP address) as the primary CA port.
- 6 The server I/O is processed by the secondary storage before the retry time-out is exceeded. The application continues running without any restrictions

Such an automatic failover is typically completed within less than three seconds, which is sufficient for most applications to keep on running smoothly.



### Conditions for Automatic Failover



	Link status			Array Status	
	(a)	(b)	(c)	Primary	Secondary
1	OK	OK	OK	Normal	Normal
2	X	OK	OK	Normal	Normal
3	OK	X	OK		
4	OK	OK	X		
5	X	X	OK	Down	Normal
6	X	OK	X	Normal	Down
7	OK	X	X	Normal	Normal
8	X	X	X	Down	Down

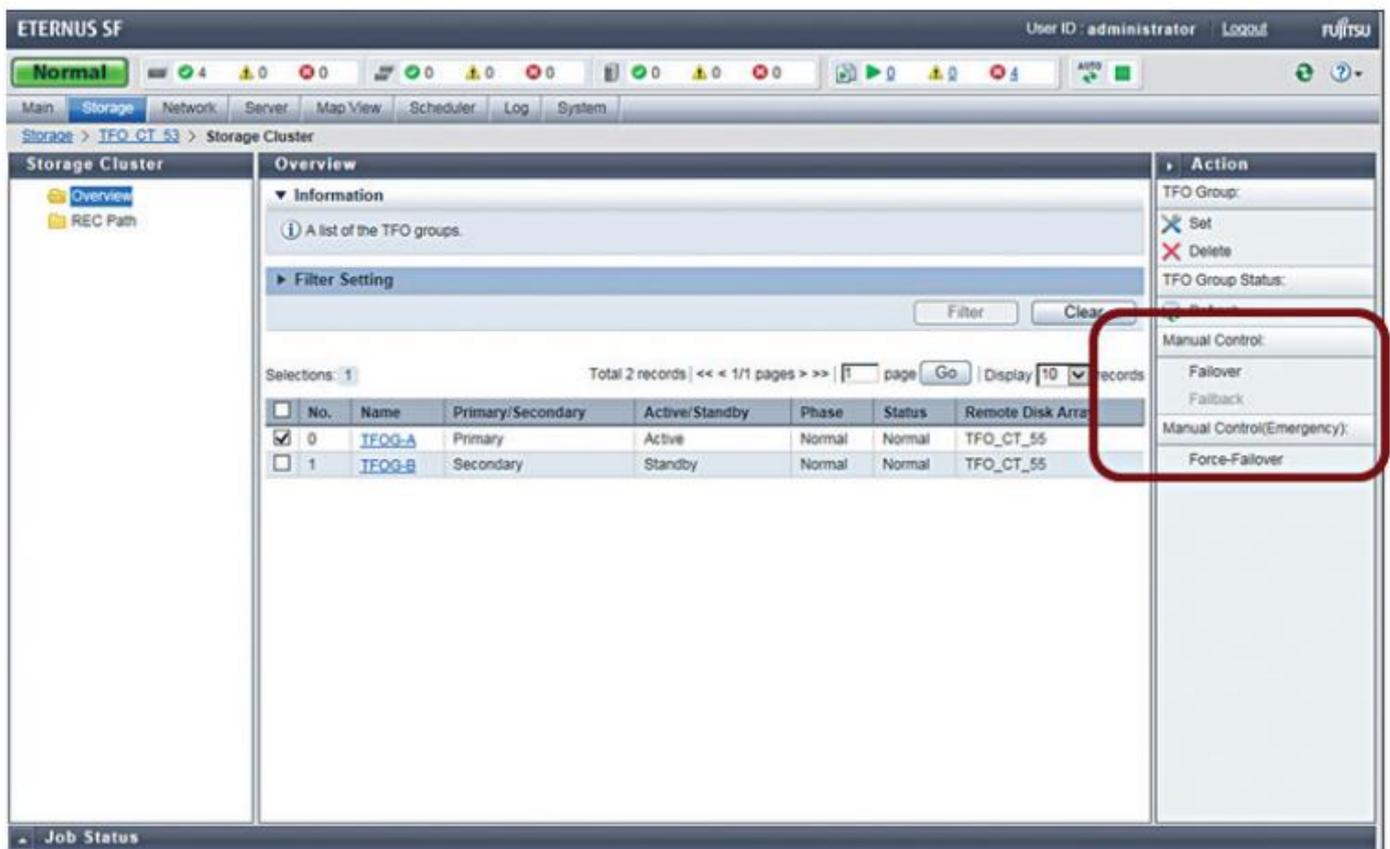
Automatic failover is only triggered if both the primary storage health supervision from Storage Cluster Controller (Link b) and the

remote replication (Link a) are down (5). In all other cases a failover is either not necessary or not possible (8)

## Failover and failback

Storage cluster can handle different failover types and scenarios:

- Auto failover:** Failover is triggered automatically when the primary storage becomes unreachable or a RAID group becomes unavailable. This mode ensures business continuity in cases of unpredictable failures or a disaster at the primary site.
- Manual failover:** Failover is triggered manually from the ETERNUS SF user interface by stopping access on the primary storage and activating the secondary storage. This mode ensures business continuity when planned downtime is required on the primary site, e.g., for maintenance, disruptive upgrades or planned power shutdowns. It can also be used for general testing of the failover mechanism.
- Force failover:** Failover is triggered manually from the ETERNUS SF user interface by activating the secondary storage regardless of the status of the primary storage. This mode ensures business continuity in cases of emergency when the primary storage is unreachable, and for any reason the auto failover cannot be executed.
- Auto failback:** The failover back from the secondary site to the primary site is triggered automatically under these conditions: The storage cluster controller confirms that all systems are operative, the REC session is established, and the business data and mirror data are consistent.
- Manual failback:** The failover back from the secondary site to the primary site is triggered manually via the ETERNUS SF user interface. This mode resets normal operation manually. The conditions of auto failback apply as well.

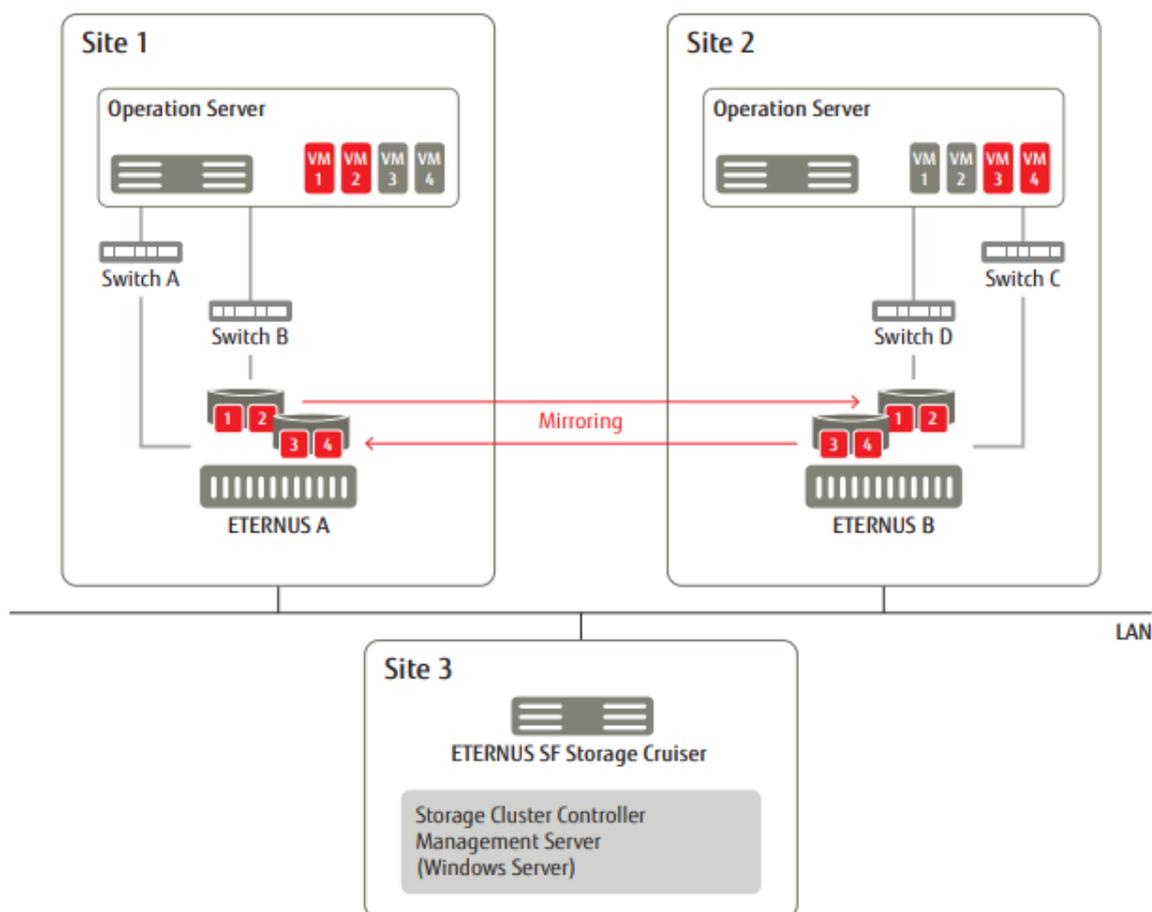


# Storage cluster in virtualized server environments

Thanks to bidirectional mirroring in combination with VMware’s HA/FT functionalities, storage cluster helps overcome even complete site outages in virtualized multisite server configurations. It provides instantaneous, non-disruptive failover in the event of server or site failures, delivering protection from even the slightest lapse, disruption or data loss.

Storage cluster configuration is based on remotely replicated TFOVs which can be freely configured and paired. Thus on one site both primary and secondary TFO groups and TFOVs can be configured – connected via linked-up (active) or passive (linked-down) CA ports, respectively.

## Bidirectional mirroring in multisite configurations



Under normal conditions, VM1 and VM2 at site 1 and VM1 and VM2 at site 2 are connected to the respective ETERNUS AF/DX arrays located at the same site – this is called the active (primary) site. All data is synchronously replicated to the standby (secondary) site, while in this example ETERNUS B is secondary for VM1 and 2 and ETERNUS A is secondary for VM 3 and 4, respectively.

The ports on both sites have the same identity regarding WWPN. As the port on the primary site is in “link-up” status and the port on the secondary site is in “link-down” status, all server inputs/outputs are processed to the primary storage.

# Conclusion

ETERNUS storage cluster provides zero-downtime, zero-data-loss continuous availability for any application, without the cost or complexity of traditional solutions like additional storage virtualization appliances or instances.

- Eliminates expensive downtime or data loss due to storage or site failures.
- Copes with both unplanned and planned downtime.
- Provides continuous service to any application, regardless of operating system.
- Provides easy setup through an intuitive administrative interface.

## Advantages – at a glance

### Feature

### Advantages

Flexible infrastructure options

In-house, campus-wide or metro-wide deployments; flexible use in mixed environments: iSCSI/FC and ETERNUS AF/DX as well as between different system generations and sizes

Ease of use

Simple wizard-based setup

Manual control

Administrator-initiated failover supports business continuity in cases of planned downtime like disruptive upgrades

Automated control

Protection against hardware, network or environmental failures to prevent unplanned downtime

Transparent Failover

Ensures business continuity for any kind of outages

# FUJITSU Storage ETERNUS AF/DX

## ETERNUS Storage Cluster

### White Paper

For more information on ETERNUS storage:

[fujitsu.com/eternus](https://fujitsu.com/eternus)

#### Contact

Fujitsu Technology Solutions GmbH  
Mies-van-der-Rohe-Strasse 8  
D-80807 Munich  
[www.fujitsu.com](http://www.fujitsu.com)

© Fujitsu 2023. All rights reserved. Fujitsu and Fujitsu logo are trademarks of Fujitsu Limited registered in many jurisdictions worldwide. Other product, service and company names mentioned herein may be trademarks of Fujitsu or other companies. This document is current as of the initial date of publication and subject to be changed by Fujitsu without notice. This material is provided for information purposes only and Fujitsu assumes no liability related to its use.