White Paper

FUJITSU Storage ETERNUS AF series and ETERNUS DX series Non-Stop Storage Reference Architecture Configuration Guide

Non-stop storage is a high-availability solution that combines ETERNUS SF products with the ETERNUS AF series or ETERNUS DX series. It provides an ideal system environment for enterprise systems that need to cater for future business growth and data accumulation.

Table of Contents

Preface

Non-stop storage is a high-availability solution for business continuity.
A non-stop operation is possible during problems and disasters with a cluster configuration created by combining FUJITSU Storage ETERNUS SF products with the FUJITSU Storage ETERNUS AF series or ETERNUS DX series. This document provides an overview of the non-stop storage solution and describes its configuration.

■Prerequisites

The product lineup and product information stated in this document are current as of November 2018.

■Target Readers

This document is intended for readers familiar with FUJITSU Storage ETERNUS SF products, the FUJITSU Storage ETERNUS AF series and ETERNUS DX series, VMware vSphere, and Windows Server.

■Target Models

This document targets the following storage system models.

- FUJITSU Storage ETERNUS AF series
- FUJITSU Storage ETERNUS DX S3 series and later*

*The ETERNUS DX60 S4/S3, ETERNUS DX100 S4/S3, and ETERNUS DX8100 S4/S3 are excluded.

■Abbreviations

This document uses the following abbreviations.

- FUJITSU Storage ETERNUS AF series	... ETERNUS AF series
- FUJITSU Storage ETERNUS DX series	... ETERNUS DX series
- FUJITSU Storage ETERNUS AF series and ETERNUS DX series	... ETERNUS AF/DX
- FUJITSU Storage ETERNUS SF Storage Cruiser	... ETERNUS SF Storage Cruiser
- FUJITSU Storage ETERNUS SF products	... ESF
- ETERNUS SF Manager	... ESF Manager
- ETERNUS SF MA	... MA
- Automated Storage Tiering function	... AST
- Automated Quality of Service	... Auto QoS
- VMware vSphere	... vSphere
- VMware ESXi	... ESXi
- VMware vCenter Server	... vCenter Server
- Microsoft Windows Server 2012 R2	... Windows Server
- Windows Server Failover Clustering	... WSFC
- Active Directory Server	... AD Server
- Domain Name System Server	... DNS Server
- Network Time Protocol Server	... NTP Server

■Related Document

Refer to the following document for information on how to set up the environment used in this document.

FUJITSU Storage ETERNUS AF series and ETERNUS DX series Non-Stop Storage Reference Architecture Configuration Procedures for Storage Cluster

1. Overview of Non-Stop Storage Solution

Redundancy has long been adopted for servers, with this also providing redundancy for any internal hard disks in the server. While configurations such as the shared disks in cluster systems or server virtualization have involved servers using external disk storage, redundancy for the storage systems themselves has relied on the server software and its use has remained limited.

In recent years, however, awareness of business continuity has been rising and virtualization involves running many virtual machines on the same physical storage system. For these reasons, redundancy is also required for general-purpose storage systems.

Non-stop storage means a storage environment that can maintain uninterrupted operation with high reliability and optimal performance over long periods of time, even in the event of physical failures, problems, or emergencies. Locating physical servers and storage systems in different buildings or at different sites provides an ideal environment for mission-critical systems because they can remain in operation during emergencies such as problems or power outages.

1.1. Features of Non-Stop Storage Solution

This solution provides redundancy for storage systems to provide an environment in which the customer's business can continue to operate even if physical damage occurs as a result of a problem or emergency by automatically switching the storage systems. Moreover, data consistency and the ability to recover are maintained even if data corruption is present after such a switchover because the backup data is automatically carried over, along with business data.

This solution includes Automated Storage Tiering (AST) that has the ability to automatically select a suitable storage device for the storage cost and the Automated Quality of Service (Auto QoS) function that automatically optimizes access to the storage area network (SAN) for linking servers with storage systems. This provides a reliable storage environment that can both maintain performance and reduce storage costs despite fluctuations in workload, business expansion, or the accumulation of data over time.

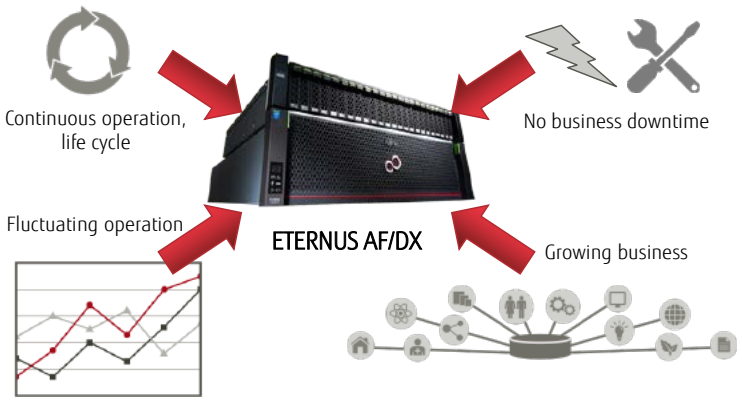

1.2. Suitable Systems for Non-Stop Storage

This solution is suitable for the following systems.

- Systems that require high reliability with uninterrupted 24-hour/365-day operation
  Systems that must remain in operation even if storage systems are unavailable due to equipment failure or building power failure, etc.
- Systems that demand stable storage performance
  Systems that are subject to variable loads of access to storage systems and virtualization environments that host various business systems
- Systems where workloads are expected to increase over time
  Systems on which storage capacity requirements and throughput continue to increase after installation due to business growth

The following shows application examples of this solution.

- Application to financial systems
  The solution provides non-stop operation with data redundancy across storage systems for systems that provide 24-hour/365-day services, such as Internet banking or ATMs, including during electrical maintenance inspections, storage failures, and disasters.
- Application to systems for managing medical imaging data
  The solution provides data redundancy at the storage system level for medical imaging data, the safe-keeping of which is a legal obligation, and reduces storage costs by automatically moving data for recovered patients to low-cost Nearline disks.

1.3. Protecting Investment and Optimizing Investment Cost

By using storage virtualization to start small and expand as required, this solution provides a storage environment that protects the customer's initial investment and can flexibly deal with business growth even for systems where future expansions are difficult to estimate.

Storage scalability typically means capacity expansion and throughput improvement.
Capacity expansion is achieved by adding drives. If more expansion is needed, additional drive enclosures can be added to provide space for more drives. If a large amount of capacity is required for reasons such as business growth, additional storage systems can be installed.
These three ways of adding capacity provide seamless capacity scalability in response to business growth.

In the case of IOPS performance improvement, the IOPS for a specific operation can be upgraded by using flash memory to provide a larger cache or by selecting high-speed drives (SSDs).
The cache size can be increased up to 800GB for the entry-level ETERNUS DX200 S4 storage system that is used in the example for this solution or up to 5.6TB for an enterprise storage system.
Furthermore, the addition of a small-capacity SSD can increase IOPS by using AST. This function automatically reallocates frequently accessed areas to SSDs to provide performance in the way that best suits the operation.

To improve the throughput per unit time, additional ETERNUS AF/DX storage systems can be installed without business downtime. Installing additional storage systems dramatically enhances the throughput.

In this way, the solution protects the customers' investment by making the most of their initial investment and by allowing additional investment to be made in small increments.

1.4. FUJITSU Storage ETERNUS SF Software for Non-Stop Storage Operation

The FUJITSU Storage ETERNUS SF storage software is used by this solution to provide non-stop operation and to handle the operation and management life cycle. With an extensive product range that extends from resource management of storage systems to the provision of additional functions such as replication, it flexibly supports administration of the storage environment and system-wide optimization.
The following sections describe the features of the ETERNUS SF products that provide powerful support for the reliable operation of the storage systems centered on the ETERNUS AF/DX.

1.4.1. ETERNUS SF Storage Cruiser

ETERNUS SF Storage Cruiser [*1] provides integrated management of the storage environment to support reliable operation of the entire system.
Used together with options such as AST [*2] and Auto QoS [*3], it provides cost reductions and more efficient use of resources.

The main functions are as follows.

- Simplified storage design and configuration
  A simple GUI is provided for easy configuration, from volume and virtual storage pool creation to allocation to servers.
  Users can start small with the minimum physical drive capacity and later add drives without interrupting operation.
- Visualization of storage resources
  A dashboard displays information of interest, enabling the status of the storage system and any problems to be identified early.
  Performance graphs provide the information for optimizing the I/O performance needed to satisfy business requirements.
  It provides integrated management of the components, including the storage system, Fibre Channel switches, physical servers, and virtual servers using Hyper-V, VMware, or other software.
- Use of AST to optimize investment
  AST is used in environments that combine drives of different types to automatically reallocate data to the most appropriate drive based on access frequency.
  It optimizes storage costs by allocating frequently accessed data to SSDs or other ultra-high-speed drives to maintain performance, and allocating less frequently accessed data to inexpensive Nearline disks with high capacity.
- Use of Auto QoS to allocate performance resources
  For systems where the storage system is being shared by a number of operations, Auto QoS prioritizes the I/Os of priority operations thereby ensuring that high priority operations operate reliably without being affected by fluctuations caused by I/O loads of other operations.
  This makes efficient use of resources by allowing business data with a wide range of different priorities to be combined.
  Response times can also be reduced by using automatic tuning together with AST.

  The ETERNUS AF/DX and ETERNUS SF Storage Cruiser together provide an ideal system for maintaining business continuity that can be used safely and reliably.

*1: Requires the ETERNUS SF Storage Cruiser license for each ETERNUS AF/DX.
*2: Requires the ETERNUS SF Storage Cruiser Optimization option for each ETERNUS AF/DX.
*3: Requires the ETERNUS SF Storage Cruiser QoS Management option for each ETERNUS AF/DX.

1.4.2. Storage Cluster

Storage Cluster *4 is an optional transparent failover feature of ETERNUS SF Storage Cruiser, which enables a non-stop storage infrastructure in the event of a single point of failure (SPOF) in a storage system.
To improve system availability, it automatically detects when the storage system recovers so as to restore redundancy as soon as possible.

The main functions are as follows.

- Data mirroring provided by storage functions
  The synchronous Advanced Copy function of the ETERNUS AF/DX maintains data consistency across different storage systems.
  Copying does not use server resources, thereby keeping any load on the operation to a minimum.
- Automated monitoring, switchover, and post-recovery switch-back
  When a failure is detected in the ETERNUS AF/DX, Storage Cluster switches the access destination automatically without operator intervention.
  Even when AST is running, data allocation is still optimized on the Secondary storage based on the access information on the Primary storage. This means that the same I/O performance is maintained after a switchover occurs and the operation can continue to function reliably.
  Post-recovery procedures are simplified by automatically detecting a data consistency restoration after the storage system recovers and the access destination switches back.
- Continuity of replication
  When replication is used on the Primary storage, the same copying is also performed on the Secondary storage at the same time.
  Because copying continues to function after switchover, replication continues and its high reliability is maintained.

  Using Storage Cluster provides a system without business downtime.
  Because automatic switchover does not require any operation, it helps minimize the operating costs of the customer's system.

*4: Requires the ETERNUS SF Storage Cruiser Storage Cluster option and the ETERNUS SF AdvancedCopy Manager license for each ETERNUS AF/DX.

1.4.3. ETERNUS SF MA

ETERNUS SF MA collects and archives performance data, and can output information such as automatic diagnosis results or long-term performance, not only on screen but also in the form of customizable reports.

The main functions are as follows.

- Automatic diagnosis
  Problems with performance can be pre-empted through automatic diagnosis performed using indicators determined based on expert know-how and (customizable) threshold values.
  Diagnosis results can be scheduled for automatic output in regular reports.
  Diagnosis results can also be presented as Baseline reports output in Excel format.
- Extensive performance data display and Navigation pane
  The relationship between the resources can be analyzed from the server side or the storage system side by tracking the resources in a tree structure.
  This provides intuitive operation of display and performance analysis, enabling bottlenecks or the scope of problems to be identified quickly.
  For example, ETERNUS SF MA interoperates with Microsoft Windows Server, VMware vCenter Server, and Hyper-V to undertake investigations that link the storage system and the IP network switches on the paths with performance problems for business servers by displaying the storage access performance from the physical servers and the virtual machines.
- Storage cost reduction by using compression and high-speed queries
  ETERNUS SF MA uses Fujitsu's highly efficient data compression for long-term storage of performance and event information.
  Performance analyses covering long periods of time can be performed quickly using in-memory processing.

The solution provides a storage system that can operate non-stop, 24 hours a day, 365 days a year.
The performance data display and automatic diagnosis functions of ETERNUS SF MA are useful to identify long-term performance trends and prevent performance problems arising due to business growth.
Use of the Navigation pane and reports of ETERNUS SF MA also reduces administrator workloads.

2. System Configuration

Non-stop storage systems use the transparent failover feature of Storage Cluster.
The system is configured with two ETERNUS AF/DX storage systems, the business servers, the Fibre Channel switches, and the management servers running ETERNUS SF products.

This chapter describes the following aspects of the non-stop storage system configuration.

2.1. Example of a Non-Stop Storage Configuration

Describes a redundant configuration for business environments as a basic configuration of non-stop storage systems.
The configuration uses two ETERNUS DX200 S4 storage systems.
The management server is virtually configured in a single physical server.

2.2. Example of a Non-Stop Storage Configuration (with Redundant Management Servers)

Describes a configuration that improves the availability of non-stop storage operation and monitoring.
The configuration uses two physical servers for the management servers and uses WSFC to make them redundant.

2.3. Hardware Layout

Describes a hardware layout chosen to facilitate a response to problems at installation sites or during disasters.

2.4. Workflow for System Configuration

Describes the sequence of steps for setting up the non-stop storage system.

2.1. Example of a Non-Stop Storage Configuration

2.1.1. Overview of System Configuration

The system uses VMware vSphere 6 for virtualization.
Storage Cluster is used to provide duplicate storage systems and vSphere HA is used to provide duplicate ESXi servers for the business servers.
Storage Cluster supports FC and iSCSI for connections between servers and storage systems, and connections between storage systems.
This configuration uses an FC.
Two Fibre Channel switches are also provided for redundancy.
A management server is used to manage the virtual environment and storage systems and to collect performance data.

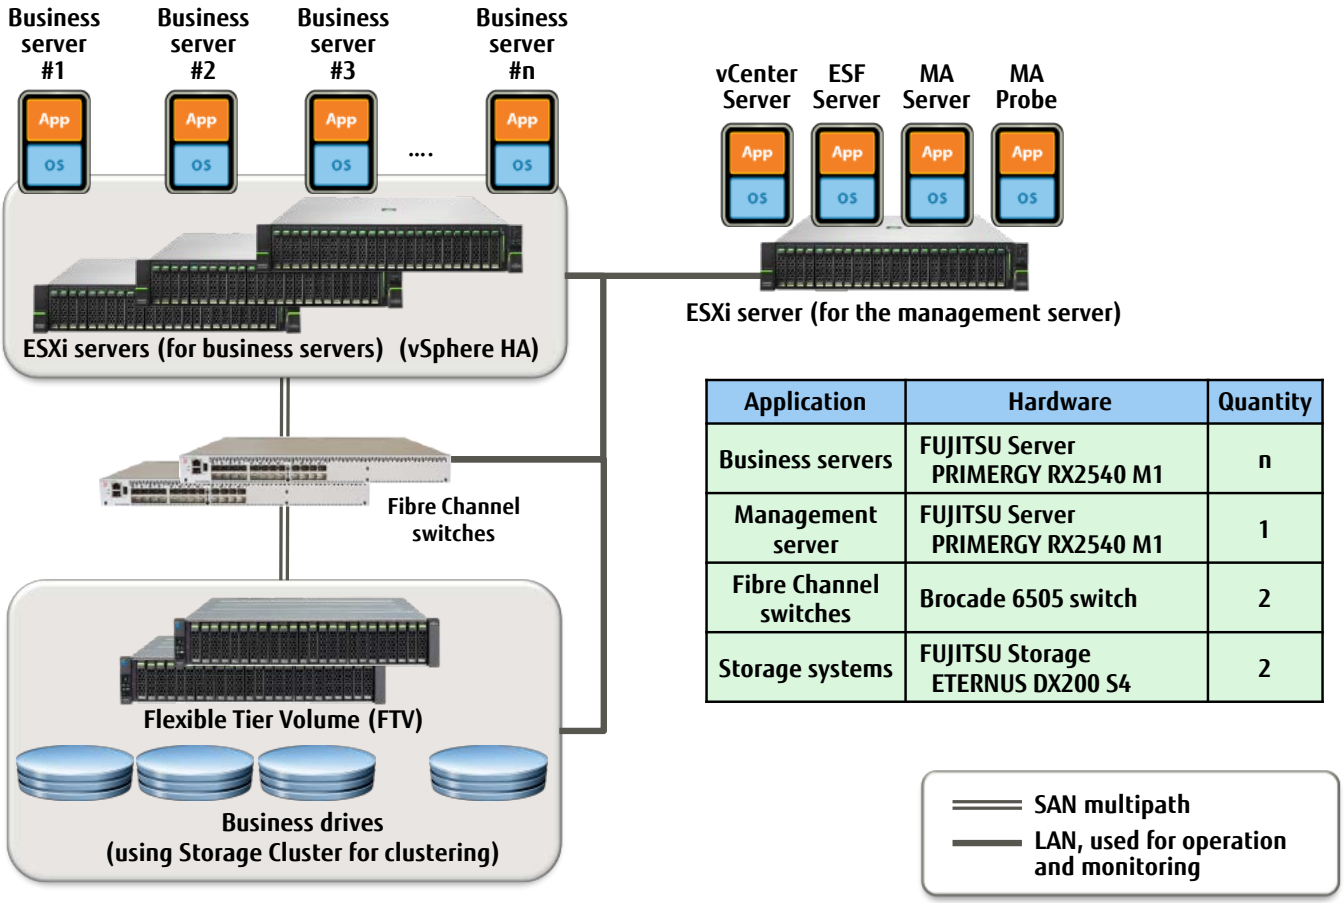

Application	Hardware	Quantity
Business servers	FUJITSU Server PRIMERGY RX2540 M1	n
Management server	FUJITSU Server PRIMERGY RX2540 M1	1
Fibre Channel switches	Brocade 6505 switch	2
Storage systems	FUJITSU Storage ETERNUS DX200 S4	2

Figure-1 System overview

Note that the servers listed below and network equipment such as the L2SW are not explained in this document.

NTP Server
The ESXi servers, vCenter Server, ESF Server, MA Server, MA Probe, Fibre Channel switches, and ETERNUS AF/DX use the NTP service to synchronize their clocks.

DNS Server and AD Server
Required by vCenter Server, the ESXi servers, and management servers to check FQDNs with each other.

2.1.2. Software Block Diagram

This section describes the software used to provide non-stop storage.

Business servers

VMware vSphere 6.0 is set up on ESXi servers and used for the hypervisor layer.
vSphere HA is used for the cluster layer, with the virtual servers able to run on any of the ESXi servers.
Either Windows or Linux can be used as the OS, depending on the requirements of the respective business servers.
Select the middleware that suits the operation based on the choice of OS.

Management server

As is the case with the business servers, VMware vSphere 6.0 is set up on an ESXi server and used for the hypervisor layer.
Because this example uses a single management server, vSphere HA is not used for the cluster layer.

ESF Server	: As this document assumes a Windows OS, the Windows versions of ETERNUS SF are used.
MA Server	: MA Server is set up as a directly connected appliance on a virtual server.
MA Probe	: MA Probe is set up as a directly connected appliance on a virtual server.
vCenter Server	: vCenter Server is set up as a directly connected appliance on a virtual server.

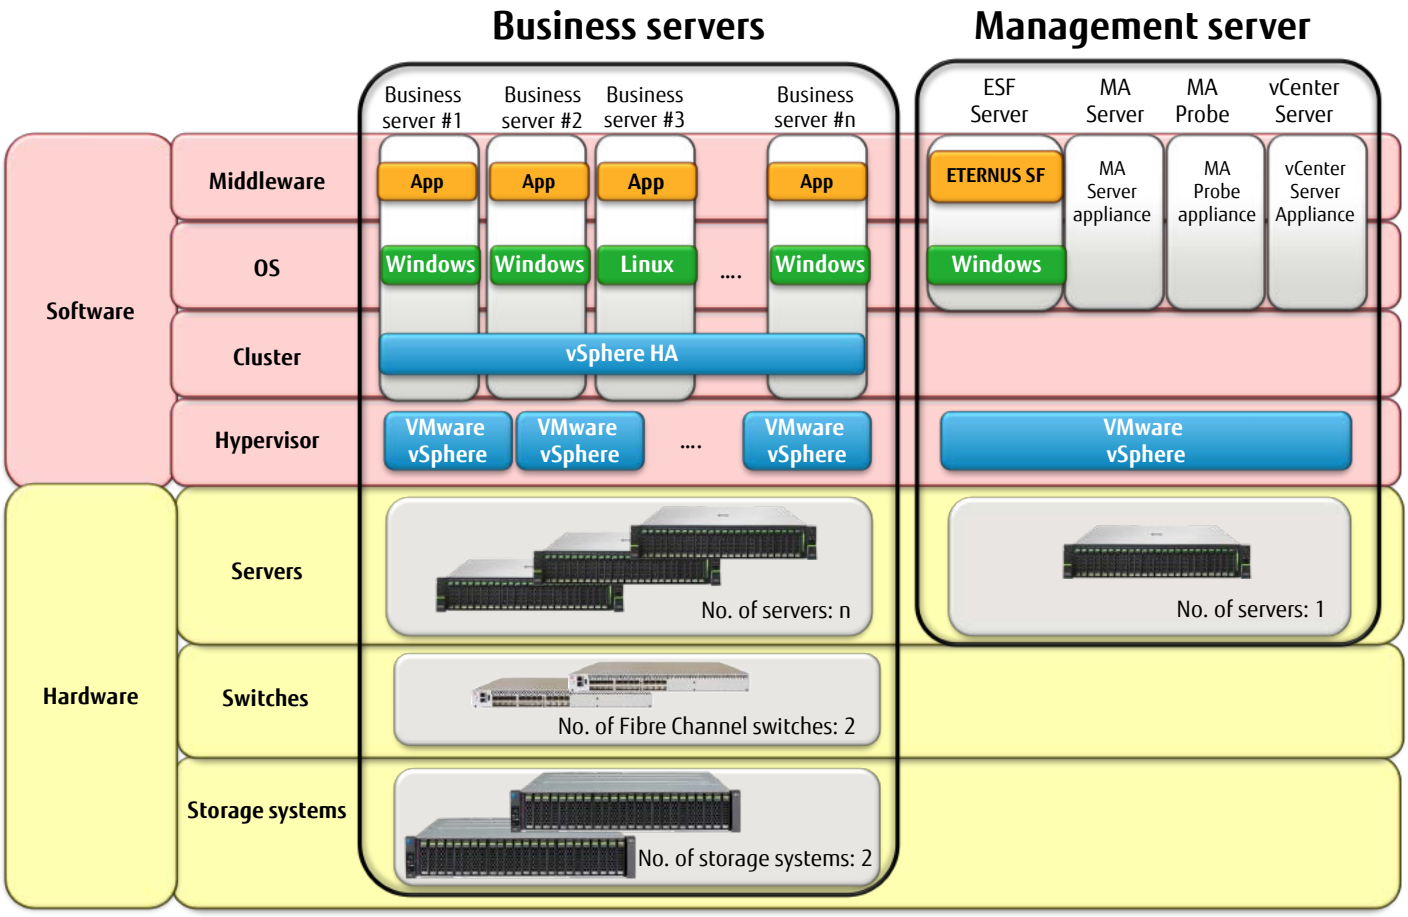

Figure-2 Software block diagram

2.1.3. Overview of Business Server Operation

This section describes how the business servers work with non-stop storage.
The storage system (ETERNUS DX200 S4) is clustered at the storage system layer. If a failure occurs on the Primary storage, the storage connected to the servers is switched over automatically by making the Secondary storage active.
AST is used to allocate data to the storage system in an efficient manner. Storage clustering is performed at the level of the volumes used for AST.
Because the business servers are virtualized, users are not aware of which ESXi server they are running on.
If a failure occurs on an ESXi server, its business servers are automatically shifted to another ESXi server and operation continues.

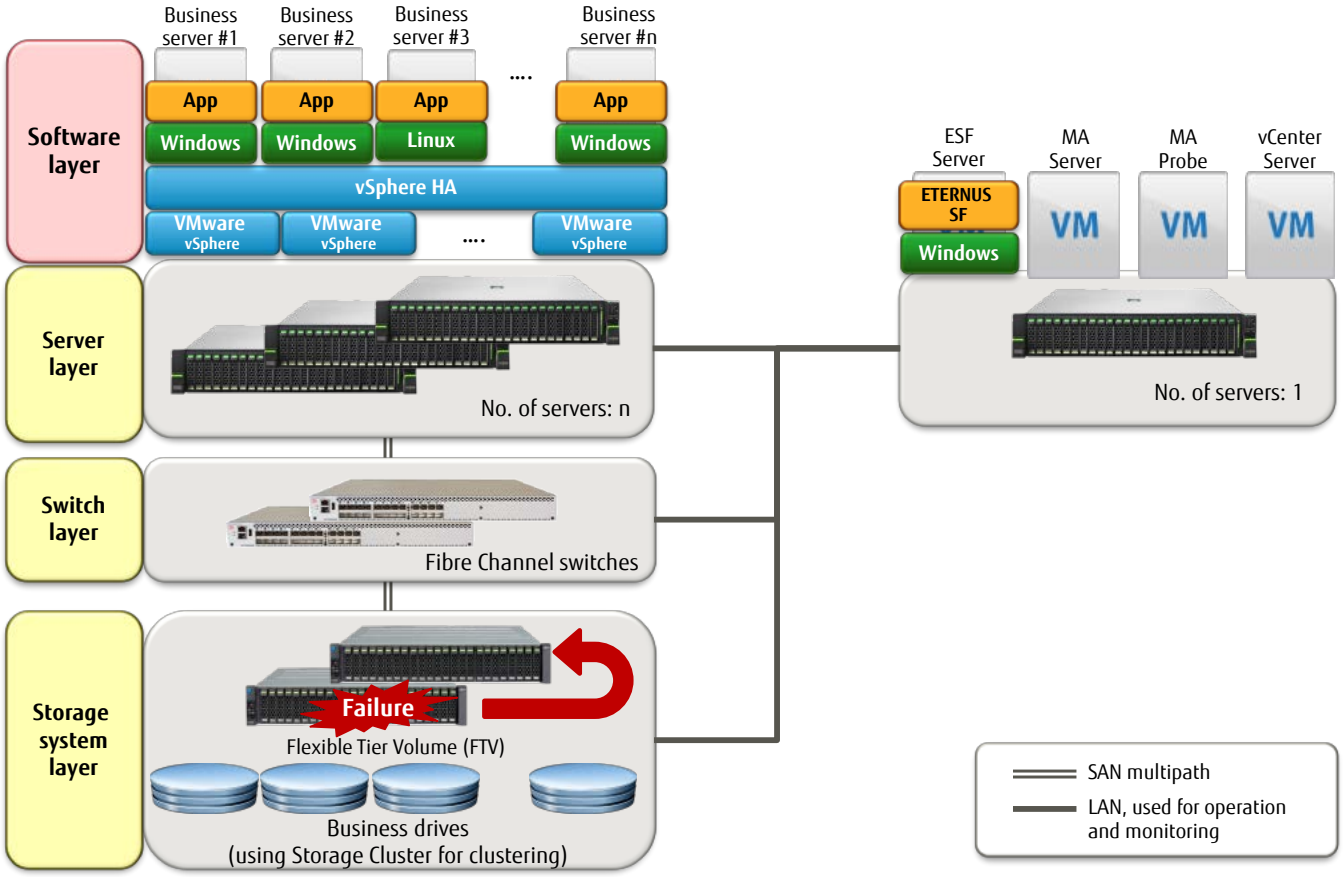

Figure-3 Overview of business operation

2.1.4. Overview of Monitoring Server Operation

ESF Server is used to manage the storage system (ETERNUS DX200 S4). It handles all aspects of storage operation, including configuration changes, AST settings, and Auto QoS policy settings.
MA Probe collects performance data on the storage systems and servers. It can collect performance data from a variety of sources, including ETERNUS Probe, VMware Probe, Linux Probe, and Windows Probe.
MA Server stores the performance data collected by MA Probe. It provides long-term archiving of the performance data, enabling analysis of performance over a period of years.
vCenter Server is used to configure and manage the virtual environments running on the ESXi servers. If a failure occurs on an ESXi server, its virtual servers are switched over to another ESXi server and automatically restarted.

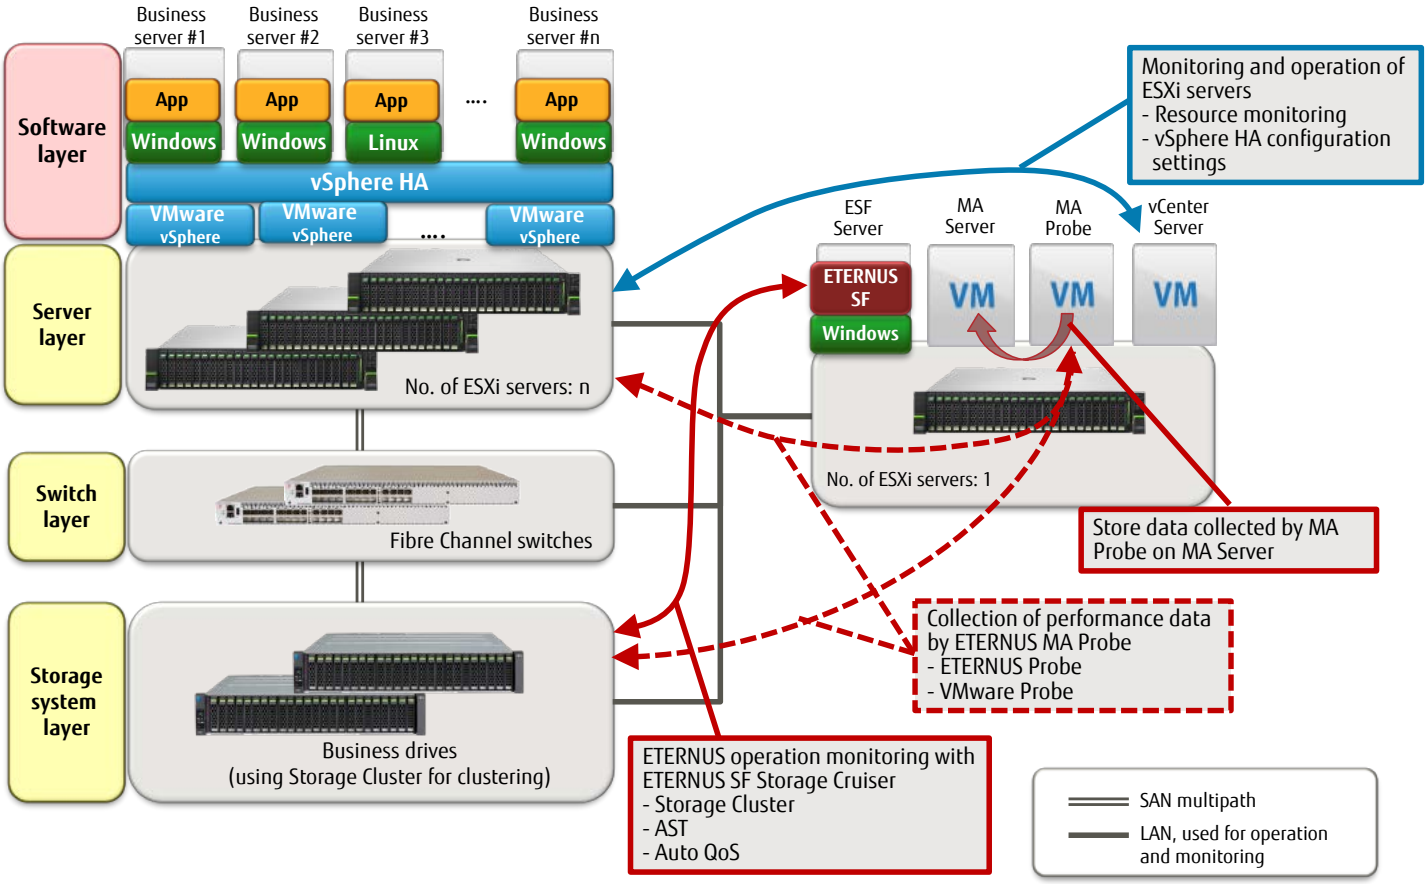

Figure-4 Overview of the monitoring operation

2.2. Example of a Non-Stop Storage Configuration (with Redundant Management Servers)

2.2.1. System Configuration Overview

This configuration is intended to improve the availability of the management servers in a non-stop storage system.
WSFC is used to configure a cluster of two physical servers in which ESF Manager handles storage management and the Storage Cluster controller handles storage monitoring.
The shared drives for the management servers are allocated in the storage systems.

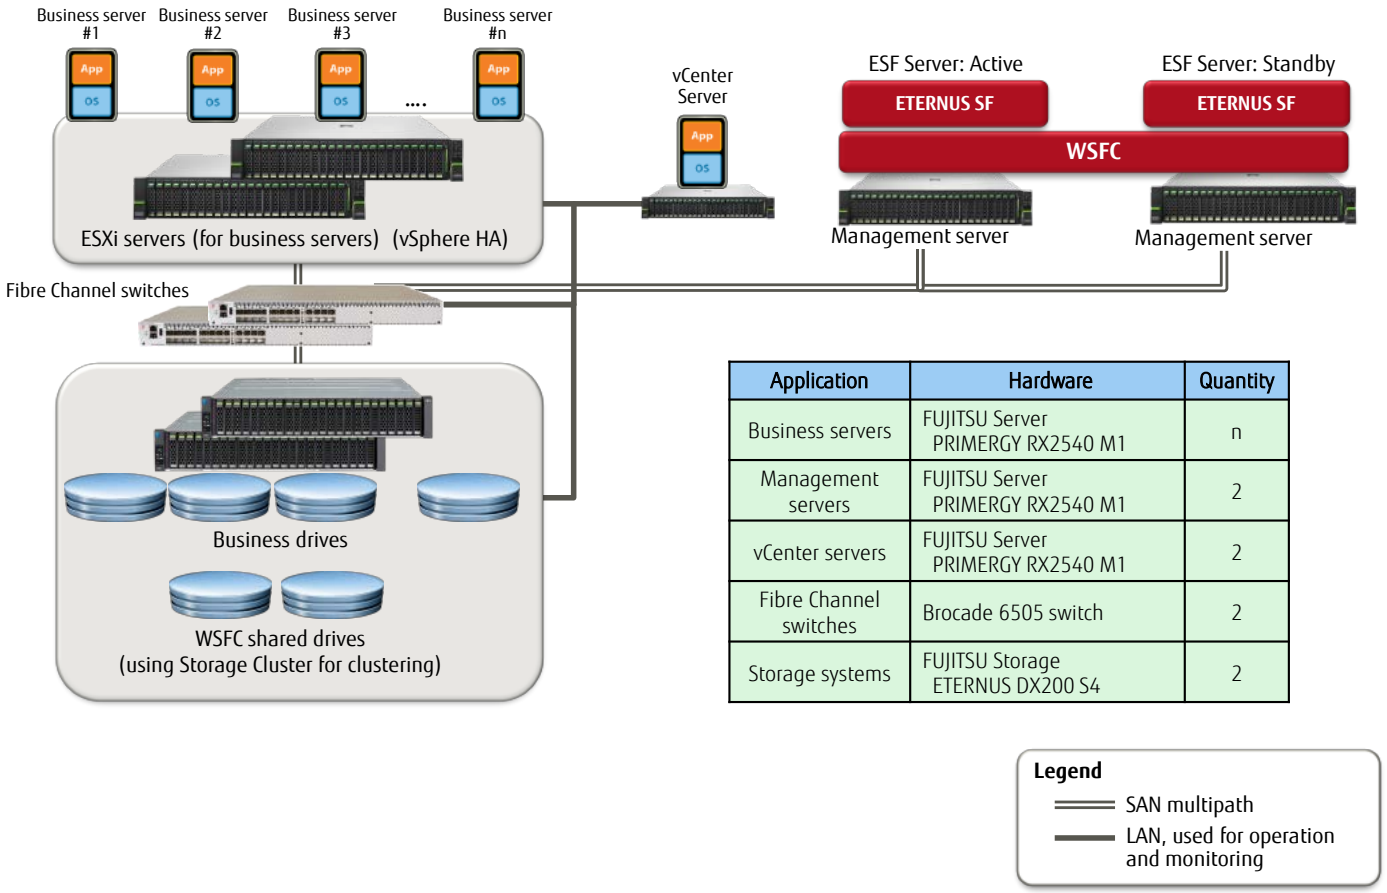

Application	Hardware	Quantity
Business servers	FUJITSU Server PRIMERGY RX2540 M1	n
Management servers	FUJITSU Server PRIMERGY RX2540 M1	2
vCenter servers	FUJITSU Server PRIMERGY RX2540 M1	2
Fibre Channel switches	Brocade 6505 switch	2
Storage systems	FUJITSU Storage ETERNUS DX200 S4	2

Figure-5 System overview for management servers clustered using WSFC

Refer to "2.1. Example of a Non-Stop Storage Configuration" for details of the business server configuration and operation.

AD Server is also required to use WSFC.

2.3. Hardware Layout

Protection against power supply problems or other disasters at the installation site requires that hardware units be physically installed in different sites.
Accordingly, install duplicate hardware units in different buildings, on different floors, or in different racks.
Whether management servers are also installed at different sites depends on whether redundancy is used.

The following figure shows a layout in which management server redundancy is not used.

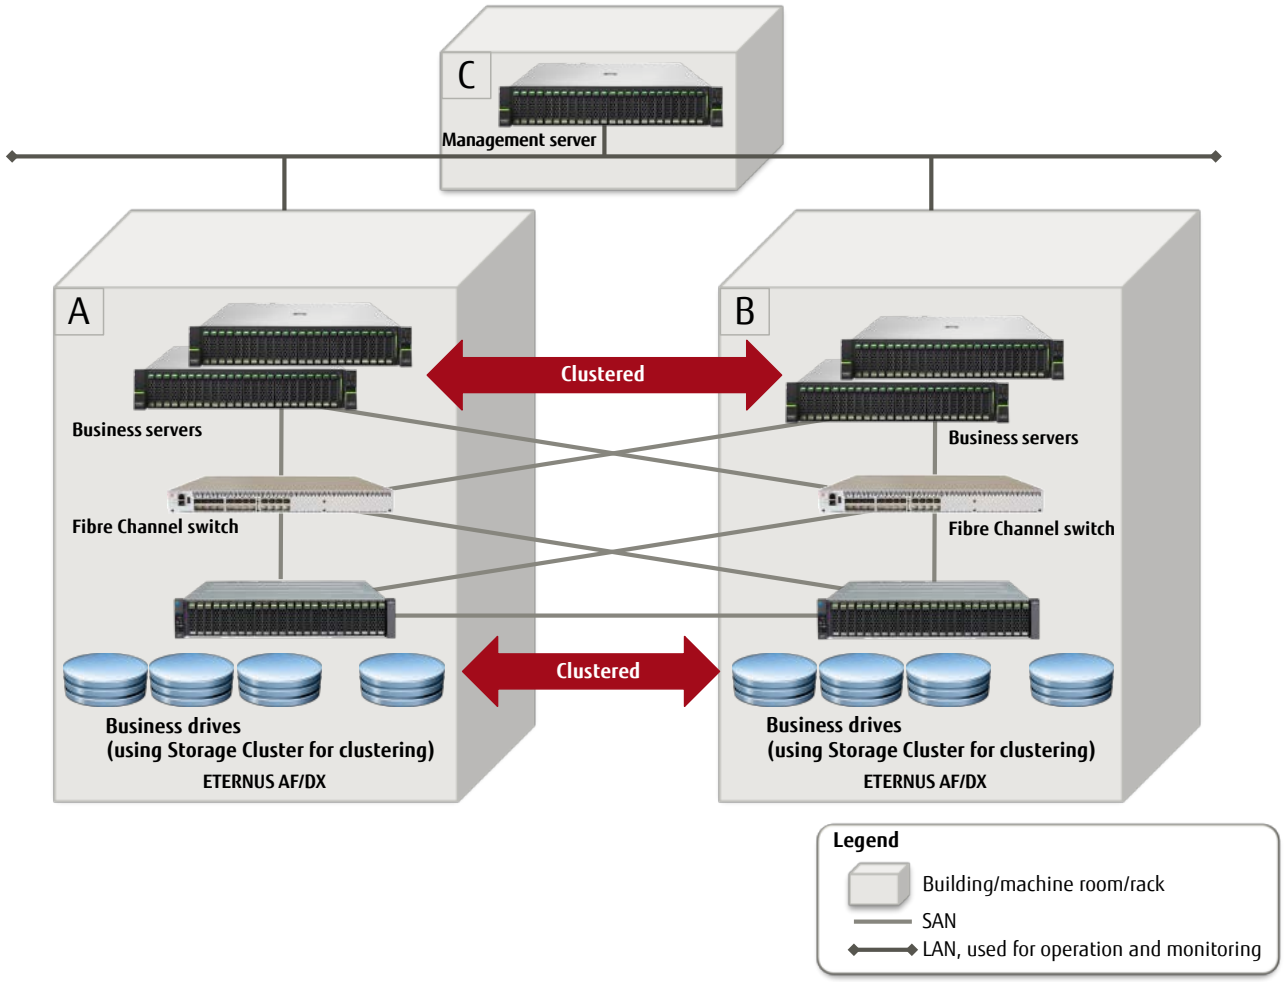

Figure-6 Layout in which management server redundancy is not used

The business servers, Fibre Channel switches, and storage systems are installed in different buildings, on different floors, or in different racks.
By clustering the business servers and the storage systems between different sites, business continuity can be maintained in the event of a problem that affects the entire site.
Install the management server at a different site to prevent it from being affected by the problems in the business servers and storage systems.

System availability can be further improved by providing redundancy for the management servers and by installing them in different sites.
The following figure shows a layout in which WSFC is used for management server clustering.

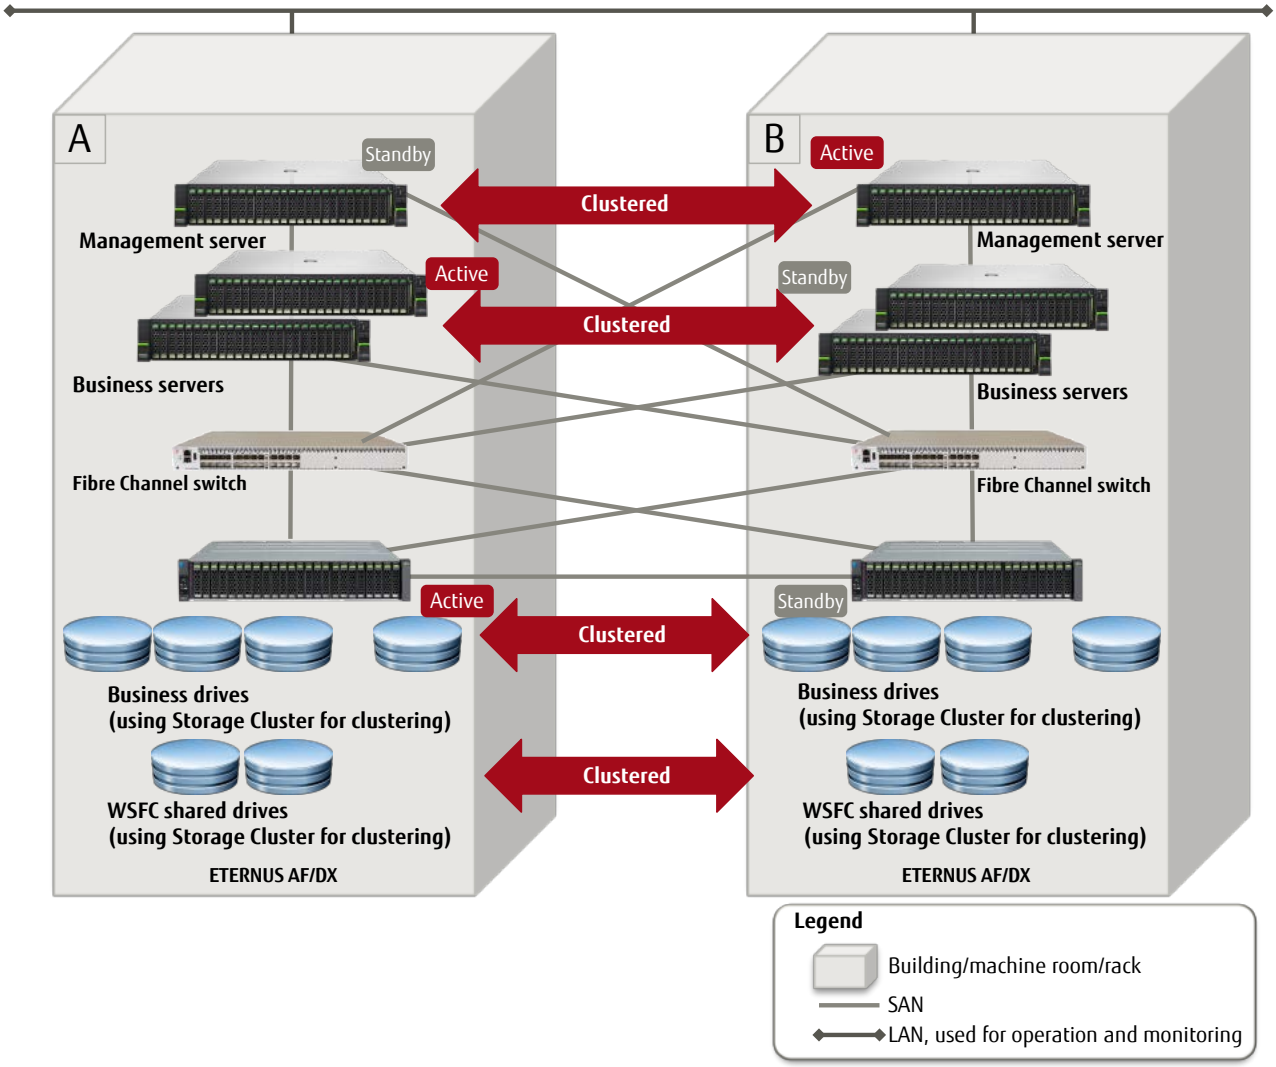

Figure-7 Layout in which WSFC is used for management server clustering

The management servers, business servers, Fibre Channel switches, and storage systems are installed in different buildings, on different floors, or in different racks.
Separate management servers, business servers, and storage systems are clustered between different sites. In addition, the active nodes of the business servers and storage system are located in a different site from the active node of the management server.
In the above example, site A is used for the active business servers and storage system and site B for the active management server.
This means that, in the event of a problem that affects the entirety of site A, monitoring continues to function, and the system continues to operate by switching over to site B.

2.4. Workflow for System Configuration

This section describes the sequence of steps for setting up the example non-stop storage system configurations.
When the management servers are configured as VMs on vSphere, the steps to be performed are as shown below (from "Advance preparation" to "Configuration of MA").
When the management servers are configured as a cluster using WSFC, the steps from "Configuration of Storage Cluster" onwards consist of "Preparation of the WSFC environment" to "Customization of management server operations".

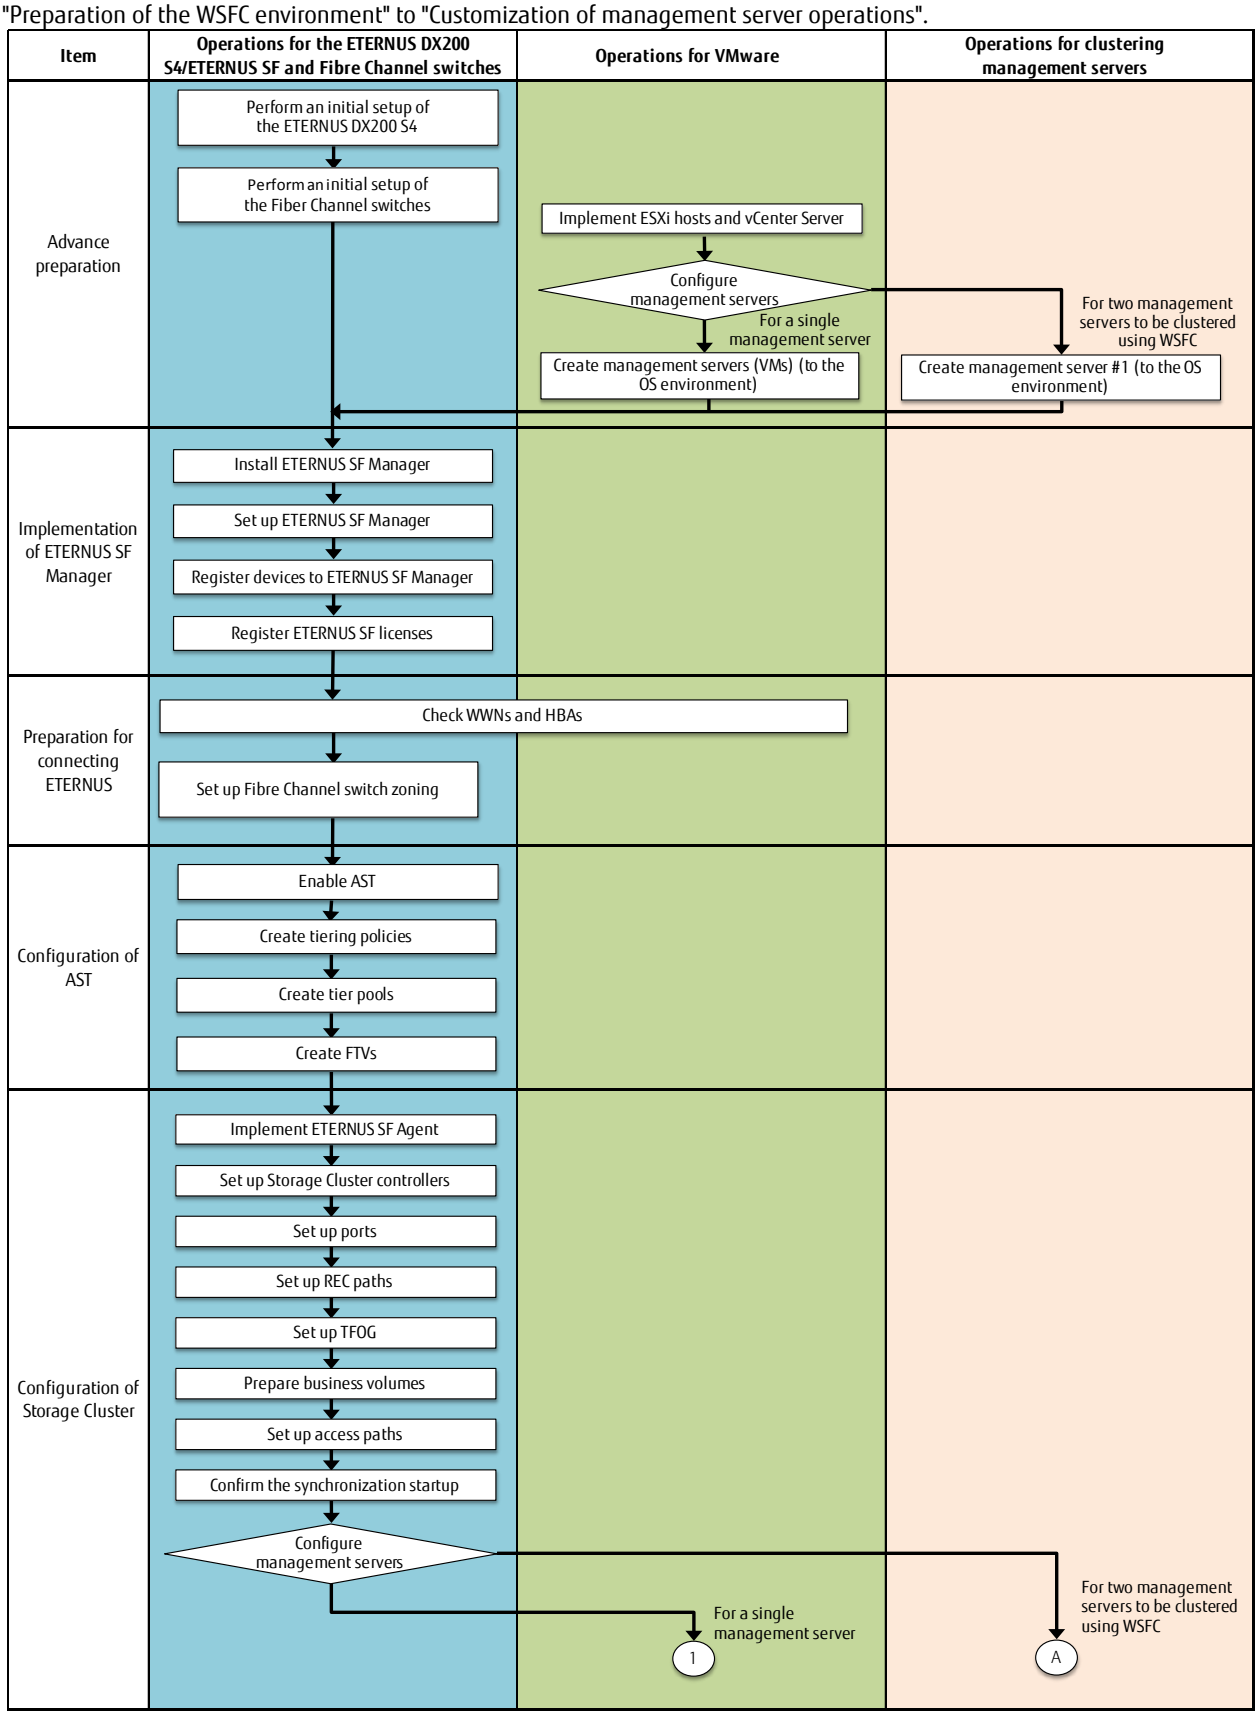

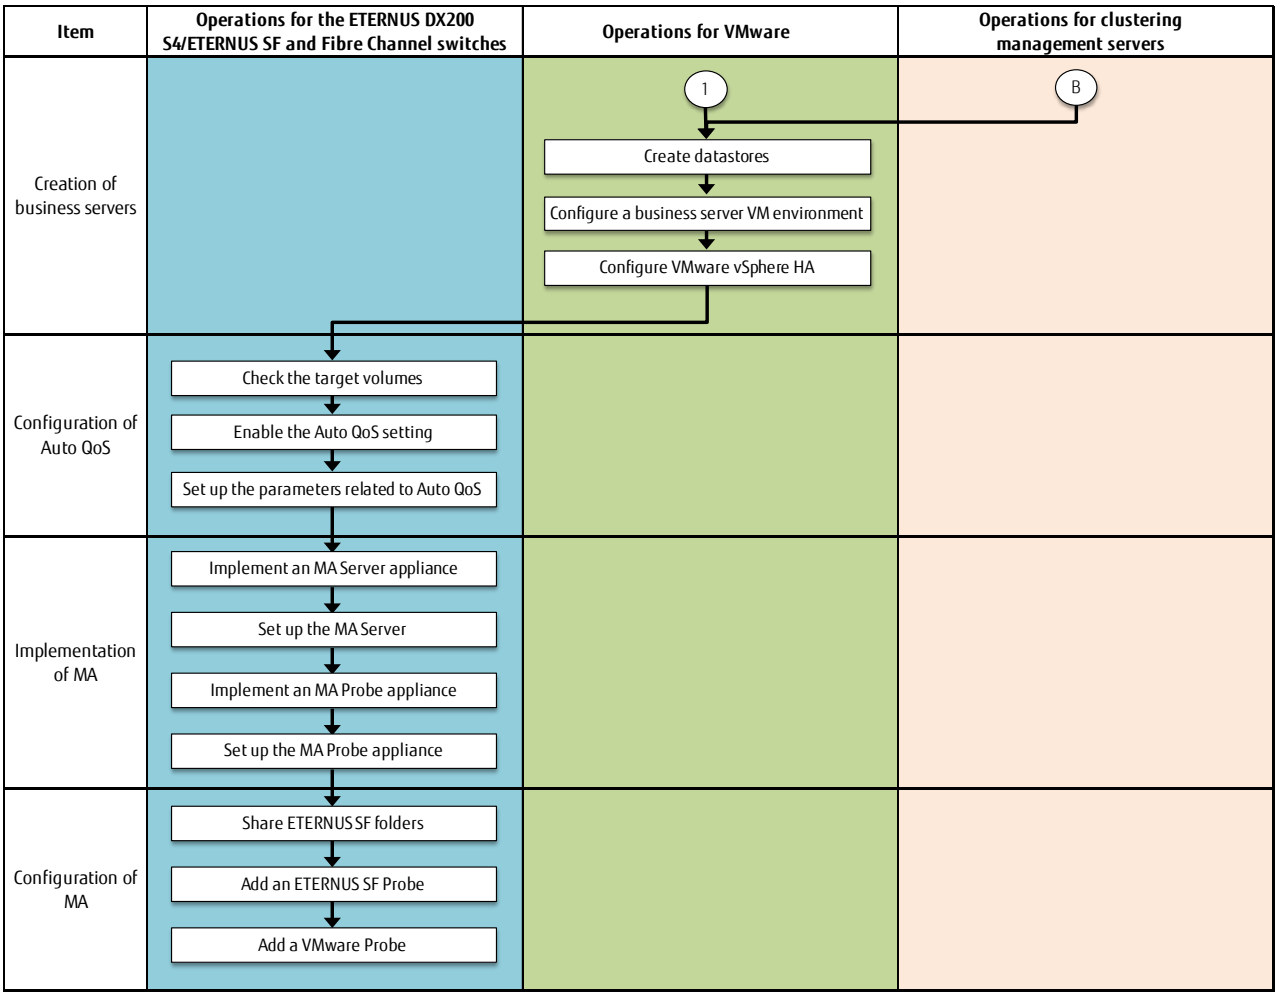
Item
Operations for the ETERNUS DX200 S4/ETERNUS SF and Fibre Channel switches
Operations for VMware
Operations for clustering management servers
Creation of business servers
1
B
Create datastores
Configure a business server VM environment
Configure VMware vSphere HA
Configuration of Auto QoS
Check the target volumes
Enable the Auto QoS setting
Set up the parameters related to Auto QoS
Implementation of MA
Implement an MA Server appliance
Set up the MA Server
Implement an MA Probe appliance
Set up the MA Probe appliance
Configuration of MA
Share ETERNUS SF folders
Add an ETERNUS SF Probe
Add a VMware Probe

When clustering the management servers using WSFC, perform the following steps in the workflow after Storage Cluster is configured in one of the management servers.

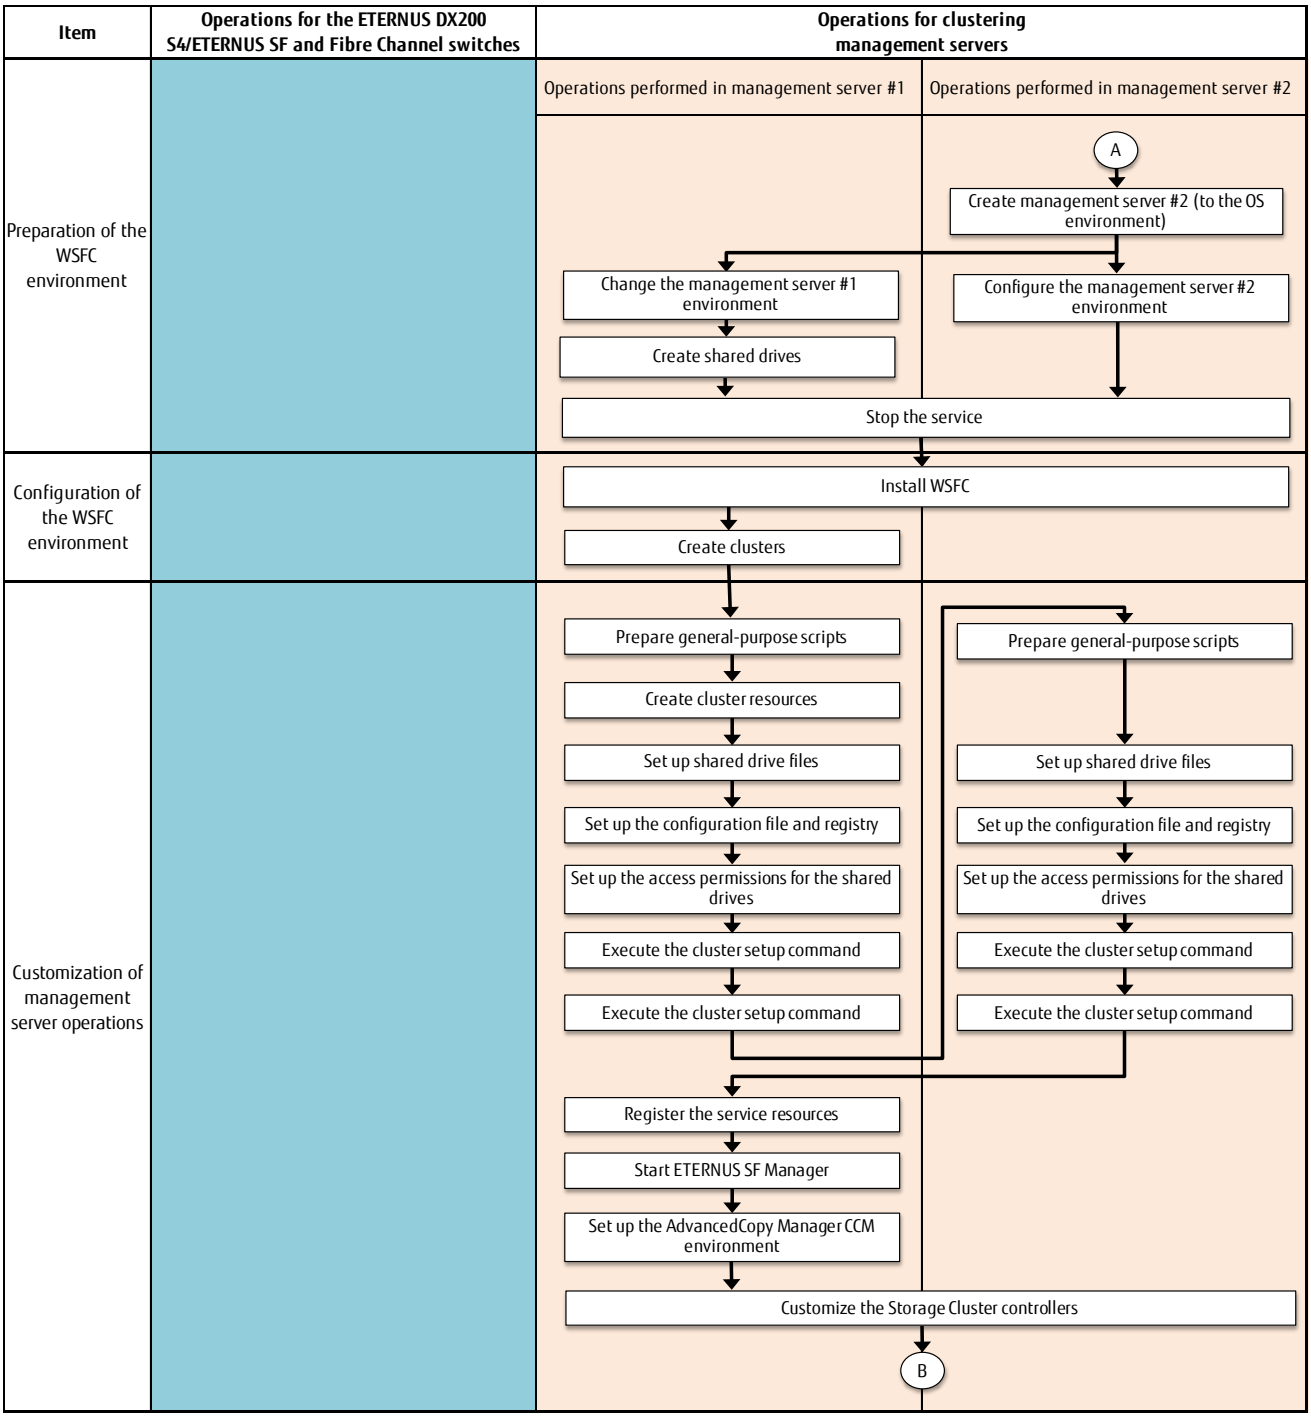

3. Conclusion

Non-stop storage is a solution for maximizing system availability with a cluster configuration created by combining ETERNUS SF products with the ETERNUS AF/DX.
This meets the need for operational continuity in virtual environments by providing a non-stop system with reliable 24-hour/365-day operation. The installation of hardware in different sites maintains operation in the event of a problem, power outage, or other disaster.

Contact
FUJITSU Limited
Website: http://www.fujitsu.com/eternus/