

White Paper

Capacity Reduction Concept Using the Deduplication/Compression Function of FUJITSU Storage ETERNUS AF series and ETERNUS DX series

The Deduplication/Compression function of the FUJITSU Storage ETERNUS AF series and the ETERNUS DX series provides storage efficiency and reduces the required number of drives. This document describes the effects of both deduplication and compression, and provides a concept for capacity reduction.



Table of Contents	
Preface	4
1 Deduplication/Compression	5
1.1 Deduplication Mechanism	5
1.2 Compression Mechanism	5
1.3 Deduplication/Compression of the ETERNUS AF/DX	6
1.3.1 Deduplication/Compression Setup Patterns	7
1.3.2 Calculating the Space Reduction Rate	9
2 Fluctuations of the Space Reduction Rate after Data Updates	11
2.1 Fluctuations of the Space Reduction Rate with Deduplication	11
2.2 Fluctuations of the Space Reduction Rate with Compression	12
2.3 Space Reduction Rate Examples for Windows System Areas	13
2.4 Space Reduction Rate Examples for System Areas of Red Hat Enterprise Linux	16
3 Estimating the Size Reduced with Deduplication/Compression	18
3.1 Estimating the Required Physical Size for Windows Server 2016	19
3.2 Reduced Size for System Areas of 1000 Virtual Machines Created by Cloning	22
3.3 Reduced Size for Data Areas	24
3.4 Summary of Reduced Size	24
4 Costs Associated with SSDs	25
5 Summary	25

List of Figures

Figure 1-1 Effects of deduplication5
 Figure 1-2 Effects of compression5
 Figure 1-3 Range of effects for deduplication/compression6
 Figure 1-4 Space reduction with deduplication and compression7
 Figure 1-5 Space reduction with deduplication or compression8
 Figure 1-6 Original size, reduced size, and space reduction rate.....9
 Figure 1-7 Space reduction rate when the deduplication setting is enabled9
 Figure 1-8 Space reduction rate when the compression setting is enabled10
 Figure 1-9 Space reduction rate when the deduplication and compression settings are enabled10
 Figure 2-1 Fluctuations of the space reduction rate with deduplication11
 Figure 2-2 Fluctuations of the space reduction rate with compression12
 Figure 2-3 Space reduction rate with deduplication for Windows Server 201613
 Figure 2-4 Space reduction rate with compression for Windows Server 201613
 Figure 2-5 Space reduction rate with deduplication and compression for Windows Server 201614
 Figure 2-6 Space reduction rate with deduplication for Red Hat Enterprise Linux 7.x16
 Figure 2-7 Space reduction rate with compression for Red Hat Enterprise Linux 7.x16
 Figure 2-8 Space reduction rate with deduplication and compression for Red Hat Enterprise Linux 7.x17
 Figure 3-1 Safety factor18
 Figure 3-2 Reduced size with deduplication19
 Figure 3-3 Reduced size with compression20
 Figure 3-4 Reduced size with deduplication and compression21
 Figure 3-5 Reduced size with deduplication (VM clone)22
 Figure 3-6 Reduced size with compression (VM clone)23
 Figure 3-7 Reduced size with deduplication and compression (VM clone)24

List of Tables

Table 2-1 Number of files based on the extension for Windows Server 201615
 Table 2-2 File size based on the extension for Windows Server 201615
 Table 2-3 File size based on text formats for Windows Server 201615

Preface

The amount of data handled by systems keeps increasing on a daily basis in our growing information society. In this circumstance, storage systems are indispensable for archiving the increasing data. However, adding HDDs and SSDs to deal with the growth is costly. Efficient measures to store data and reduce costs will be necessary even more in the future.

The FUJITSU Storage ETERNUS AF series All-Flash Arrays and ETERNUS DX series Hybrid Storage Systems provide the Deduplication/Compression function to save the storage space.

This document describes the Deduplication/Compression function and then provides a concept for storage space reduction by showing the difference between the Deduplication function and the Compression function. The verification results described in this document are current as of April 2017.

The product lineup and product information stated in this document are current as of November 2019.

■ Target Readers

This document targets the following readers.

- Readers who are looking to install or replace storage systems or to make a proposal to install or replace storage systems

■ Assumption for Readers

The target readers are assumed to have the following knowledge.

- General concepts of file systems and storage systems
- Concepts for the deduplication/compression functions of the storage systems

■ Target Models

This document targets the following storage system models.

- FUJITSU Storage ETERNUS AF250 S3/S2 and AF650 S3/S2
- FUJITSU Storage ETERNUS DX200 S5/S4, DX500 S5/S4, DX600 S5/S4, DX900 S5, and DX8900 S4^{*1 *2}

^{*1} To use the Deduplication/Compression function with the ETERNUS DX500 S4/DX600 S4, the Memory Extension must be installed.

^{*2} Only the Compression function is supported by the ETERNUS DX900 S5 and the ETERNUS DX8900 S4.

To use the Compression function, each Controller Enclosure (CE) must be installed with 768GB or more of system memory.

■ Abbreviations

The following abbreviations are used in this document.

- FUJITSU Storage ETERNUS AF S3/S2 series All-Flash Arrays ETERNUS AF series
- FUJITSU Storage ETERNUS DX S5/S4 series Hybrid Storage Systems ETERNUS DX series
- ETERNUS AF series and ETERNUS DX series ETERNUS AF/DX
- Web GUI for the ETERNUS AF/DX ETERNUS Web GUI
- FUJITSU Storage ETERNUS SF Storage Cruiser ETERNUS SF Storage Cruiser
- Solid State Drive SSD
- Thin Provisioning Pool TPP
- Thin Provisioning Volume TPV

■ Terminology

The following terms are used in this document.

- Drive An HDD or SSD that is used in the storage system to store data
- Deduplication/Compression A function that can perform deduplication, compression, or both deduplication and compression

1 Deduplication/Compression

Recent storage systems are equipped with deduplication/compression functions to save storage space.

Deduplication is a function that automatically analyzes and eliminates duplicate data. This function minimizes the number of times data is written and prolongs the life of SSDs.

Compression is a function that reduces the size of data by converting it to different data while maintaining its nature.

Both the deduplication function and the compression function reduce data size. However, their mechanisms are different from each other.

Fujitsu's ETERNUS AF/DX provides deduplication, compression, or both deduplication and compression.

Because their mechanisms differ from each other, the effects of the Deduplication function should be examined separately from the effects of the Compression function. This approach can determine the size to be reduced and save on extra expenditures.

The following sections describe the mechanisms of the Deduplication function and the Compression function, and provide a concept of space reduction by taking the difference between these functions into consideration.

1.1 Deduplication Mechanism

Deduplication divides newly written data into data blocks and compares them to existing data blocks. If identical data blocks are not found, the new data blocks are stored in the drive.

If identical data blocks are found, they are recognized as duplicates.

For duplicate data block recognition, an association is performed between the newly written data and the data blocks that have already been written to the physical area.

As a result, storage space can be saved because duplicate blocks are not saved in the drive.

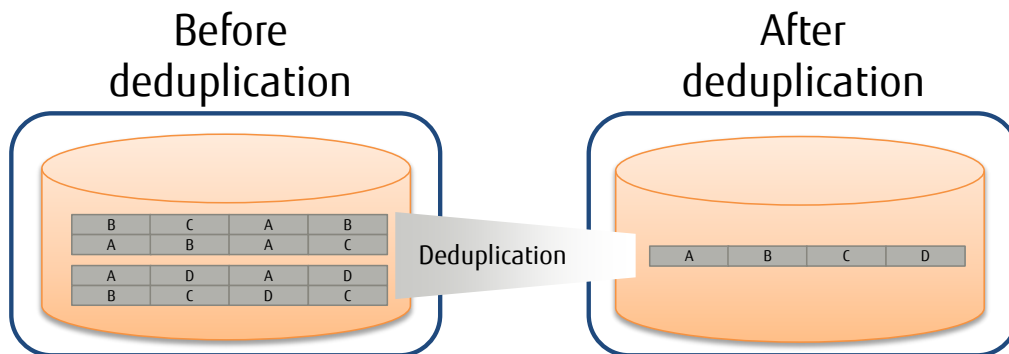


Figure 1-1 Effects of deduplication

1.2 Compression Mechanism

Compression works in a similar manner as the basic ZIP compression format. This function analyzes newly written data for duplication and repetition based on the bit order and then converts the data on a per-block basis to a compressed format that can be restored to the original bit order.

By reducing the size of the data that is to be saved, compression provides storage space efficiency.




Figure 1-2 Effects of compression

1.3 Deduplication/Compression of the ETERNUS AF/DX

The ETERNUS AF/DX uses the "inline" method to reduce the total size of data to be written. This method detects and removes duplicates on a per-block basis in the controller and compresses the data before writing it to a drive. Because the inline method does not require a temporary area, implementation and operation of this function are easy to design and the initial investment cost can be reduced. In addition, reducing the data size to be written also leads to prolonging the life of the SSDs.

Deduplication/Compression is effective within the range of a TPP. Data is processed in the memory allocated for each TPP. The Deduplication function is not performed between TPPs.

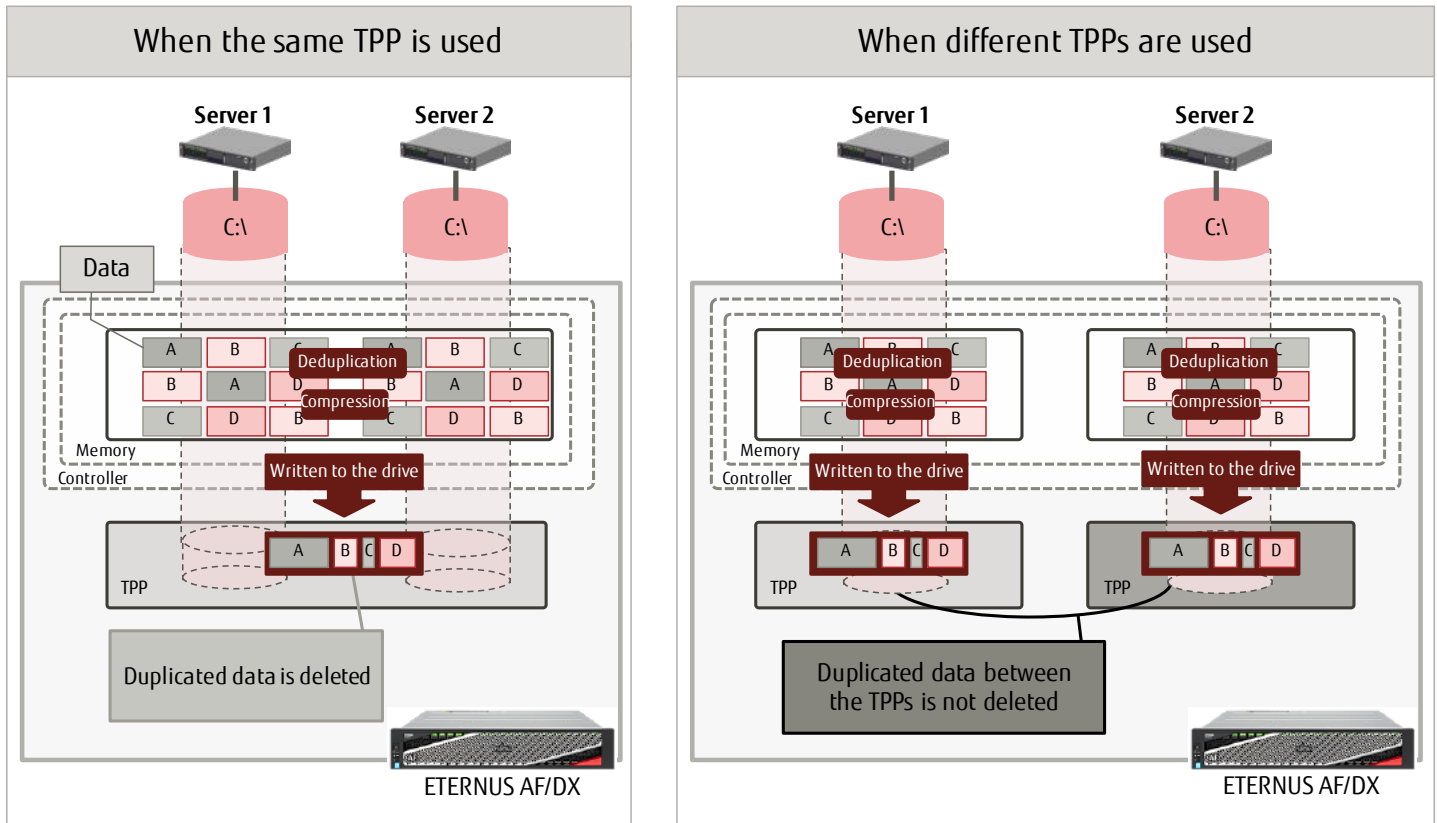


Figure 1-3 Range of effects for deduplication/compression

1.3.1 Deduplication/Compression Setup Patterns

Deduplication, Compression, or both Deduplication and Compression can be enabled for each TPP in the ETERNUS AF/DX. This setting is performed while a TPP is being created and cannot be changed once it is set.

To save data to TPPs, Thin Provisioning Volumes (TPV) must be defined in a TPP and be recognized by the server OS.

Deduplication/compression is enabled or disabled for each TPV according to the mode that is set for the relevant TPP.

By configuring a suitable setting for the TPP according to the data to be stored, the operational efficiency of the storage system can be improved.

When both deduplication and compression are enabled, the written data is deduplicated before being compressed.

The following figure compares the size of the data written to the drives when Deduplication/Compression is enabled and when disabled. It shows that the deduplication and compression processes reduce the size of the data to be written to the drives and save storage space.

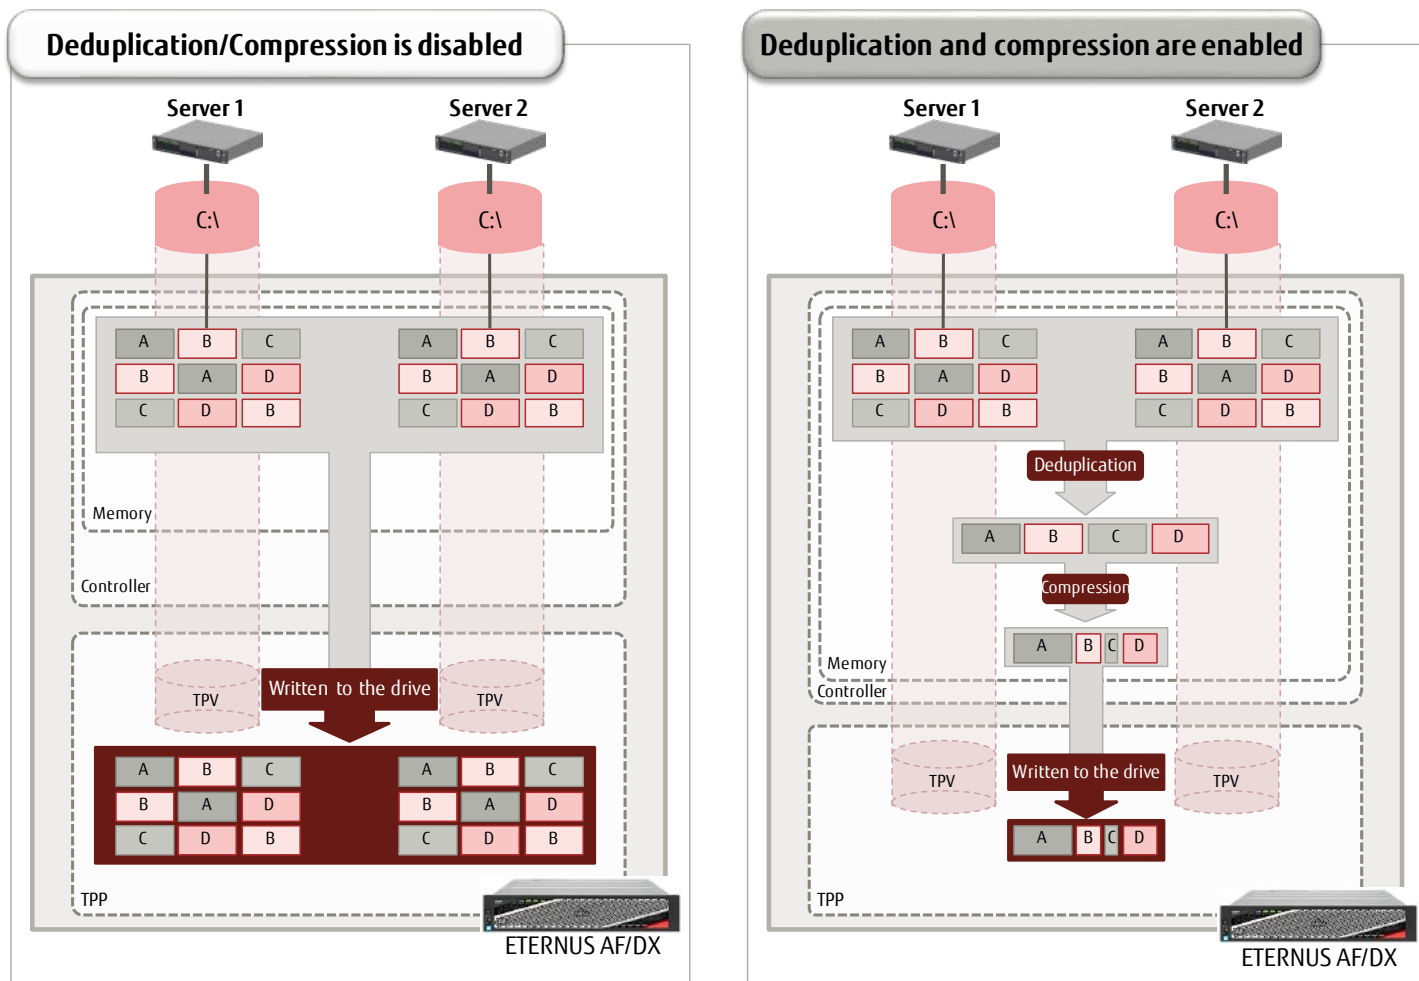


Figure 1-4 Space reduction with deduplication and compression

The following figure illustrates how the data to be written to the drives is reduced when only deduplication is enabled or when only compression is enabled. It shows that the size of the data to be written to the drives is reduced and storage space is saved.

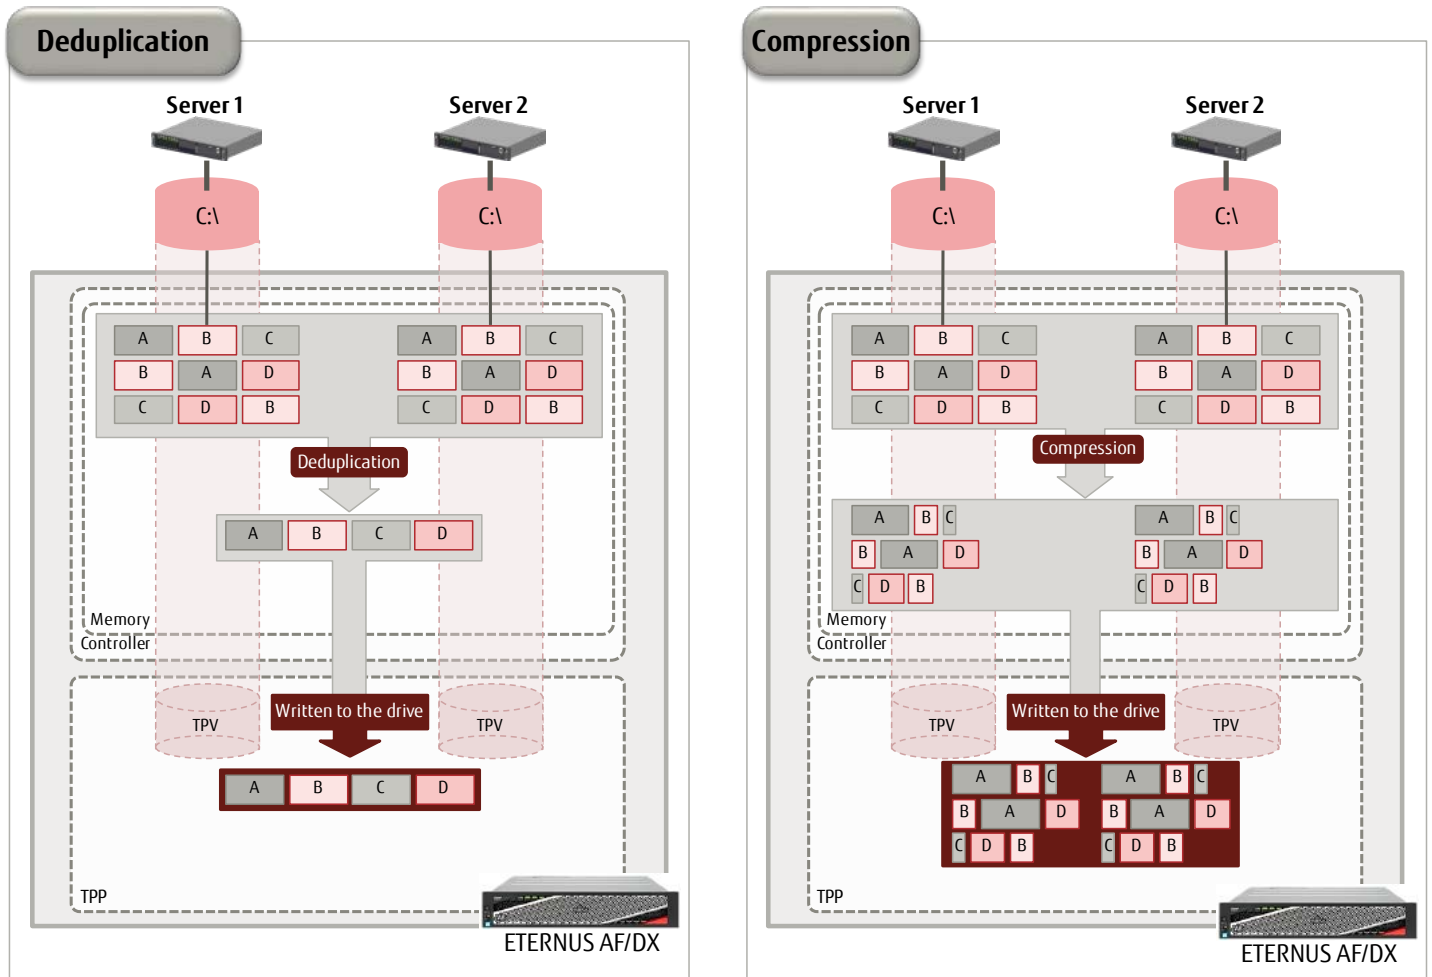


Figure 1-5 Space reduction with deduplication or compression

1.3.2 Calculating the Space Reduction Rate

The space reduction rate is displayed as a percentage for each TPP in the ETERNUS AF/DX to show the amount reduced with Deduplication/Compression.

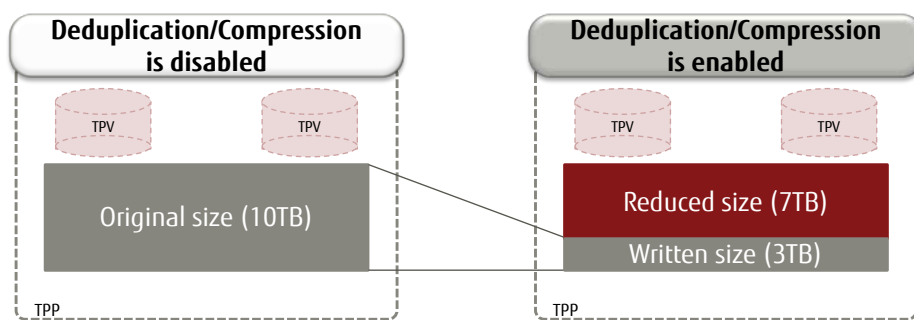
This section describes the space reduction rate.

The space reduction rate can be checked from "Capacity Reduction Rate" in the [Thin Provisioning Pool] screen of ETERNUS Web GUI or from "Thin Provisioning Pool: Deduplication/Compression Capacity Reduction Rate Display" of ETERNUS SF Storage Cruiser.

The space reduction rate indicates the percentage with which the size is reduced using the Deduplication/Compression function when compared to the size that is written when the Deduplication/Compression function is disabled.

The size that is written when Deduplication/Compression is disabled is hereinafter referred to as "original size", and the reduced size by this function is referred to as "reduced size".

$$\text{Space reduction rate} = (\text{reduced size}) \div (\text{original size}) \times 100$$

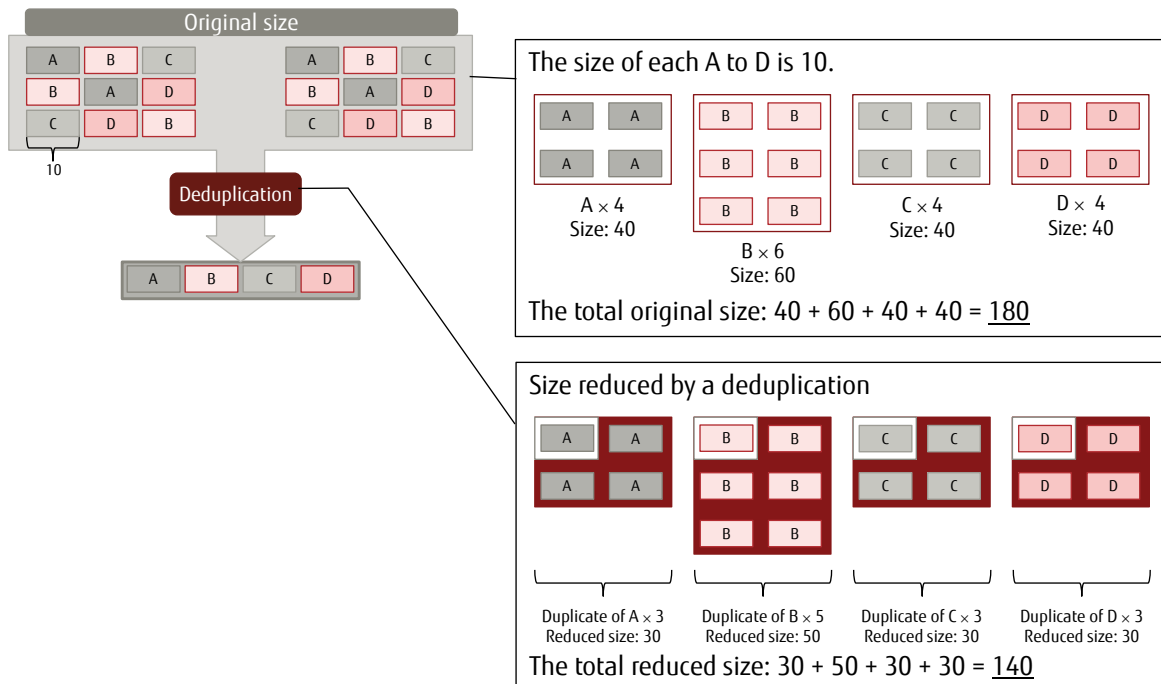


$$\text{Space reduction rate} = 7\text{TB} \div 10\text{TB} \times 100 = 70\%$$

Figure 1-6 Original size, reduced size, and space reduction rate

The concept for the space reduction rate is provided below according to the deduplication/compression setup patterns.

When the deduplication setting is enabled:



$$\text{Space reduction rate} = (\text{reduced size}) \div (\text{original size}) \times 100$$

$$140 \div 180 \times 100 \approx 78 \quad 78\%$$

Figure 1-7 Space reduction rate when the deduplication setting is enabled

When the compression setting is enabled:

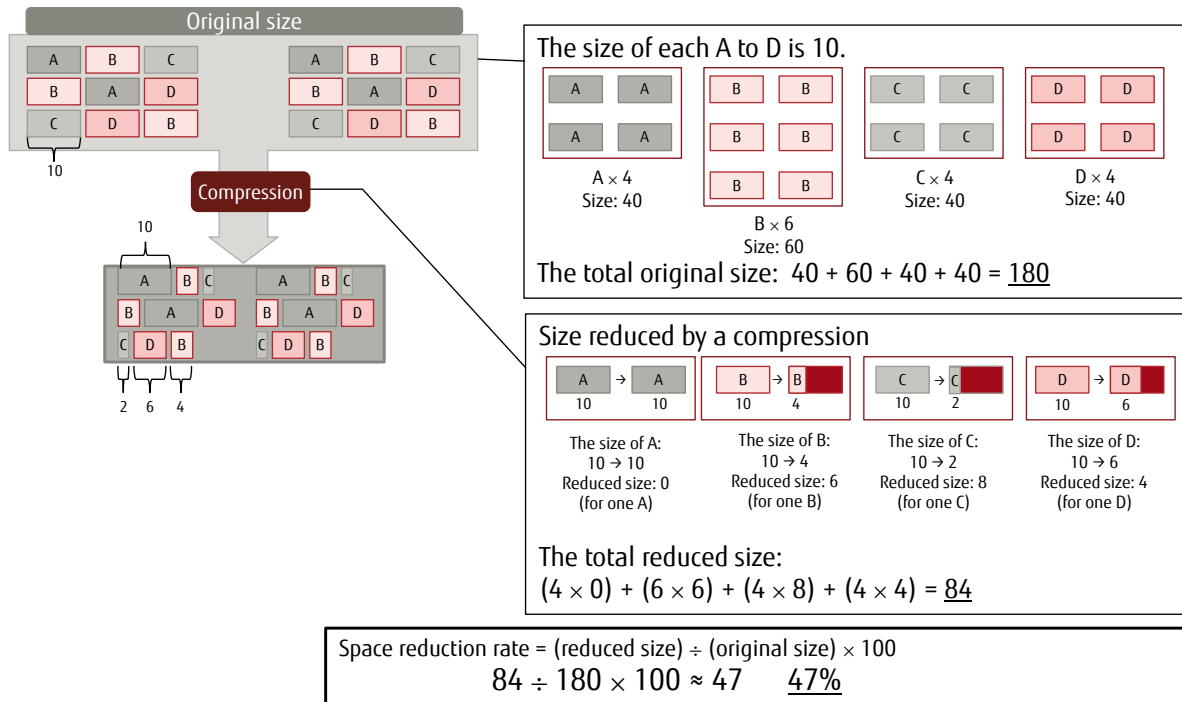


Figure 1-8 Space reduction rate when the compression setting is enabled

When both the deduplication setting and the compression setting are enabled:

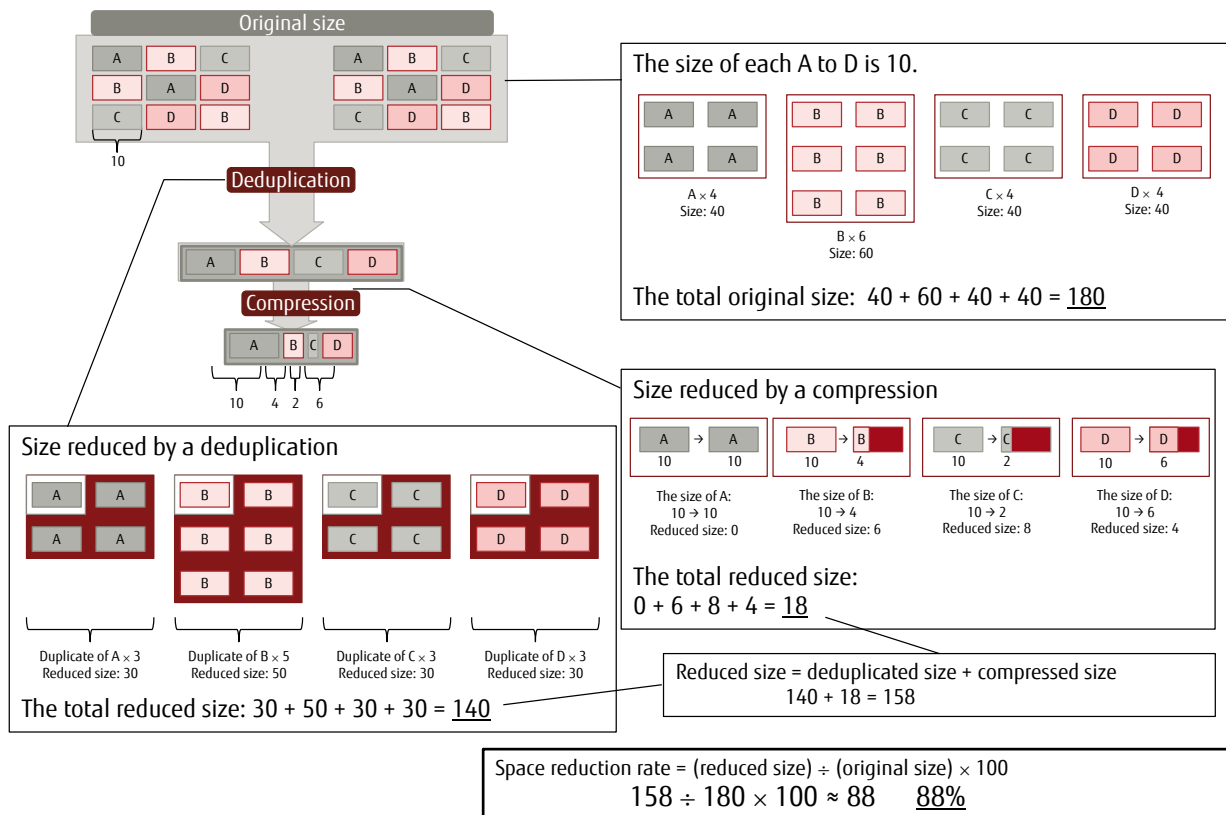


Figure 1-9 Space reduction rate when the deduplication and compression settings are enabled

2 Fluctuations of the Space Reduction Rate after Data Updates

The space reduction rate changes according to the update order or the update targets. In addition, the way the rate changes is different between the Deduplication and Compression functions because they have different mechanisms.

This chapter describes how the space reduction rate is changed by the Deduplication function and the Compression function after data is updated.

2.1 Fluctuations of the Space Reduction Rate with Deduplication

The Deduplication function is mainly applied to areas that contain large amounts of duplicate data created by file backups or virtual machine (VM) clones. In this case, the space reduction rate can be very high when deduplication is performed for the first time.

However, updating this kind of data separately leads to less duplicates.

For example, 10 cloned VMs have a large amount of duplicates but after a patch is applied in five of the guest OS VMs, the files updated by the patch do not have duplicates.

This way, because data updates may greatly influence the space reduction rate, predicting the fluctuation is not easy for areas with a high space reduction rate.

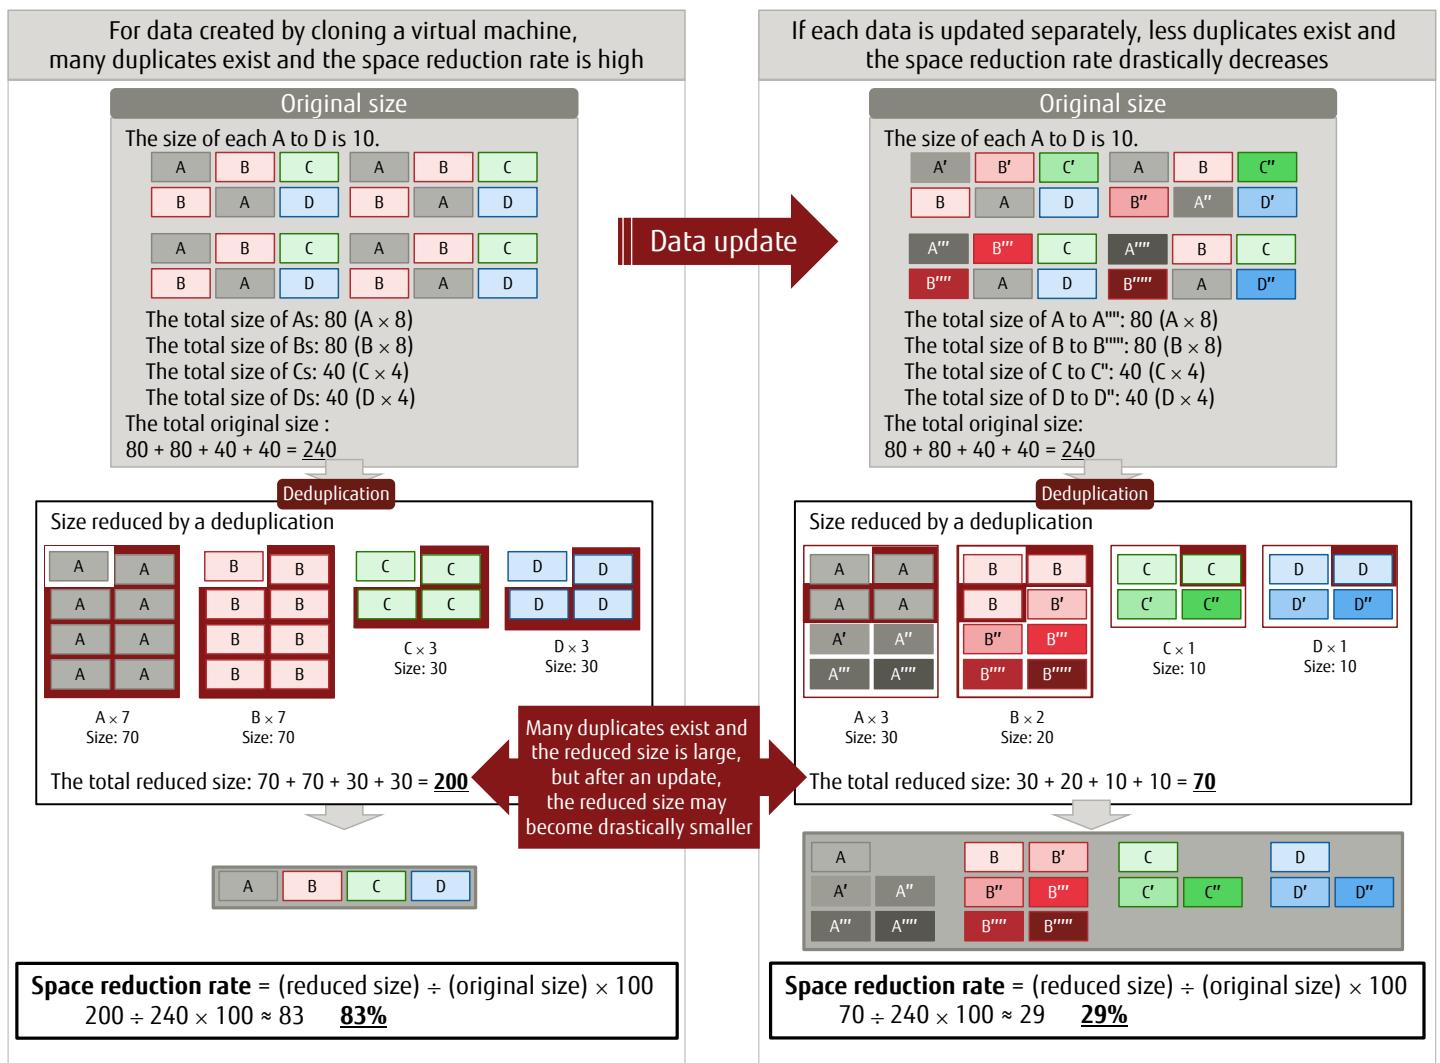


Figure 2-1 Fluctuations of the space reduction rate with deduplication

2.2 Fluctuations of the Space Reduction Rate with Compression

For compression, fluctuations of the space reduction rate depend on the type of data stored. The space reduction rate is high for text data and uncompressed data. On the other hand, it is very low for already compressed data and program files.

Increasing or reducing the number of characters in a text file does not affect the space reduction rate. Even if data that was created by copying or by cloning a VM is updated separately, predicting fluctuations is easy unless the file configuration drastically changes.

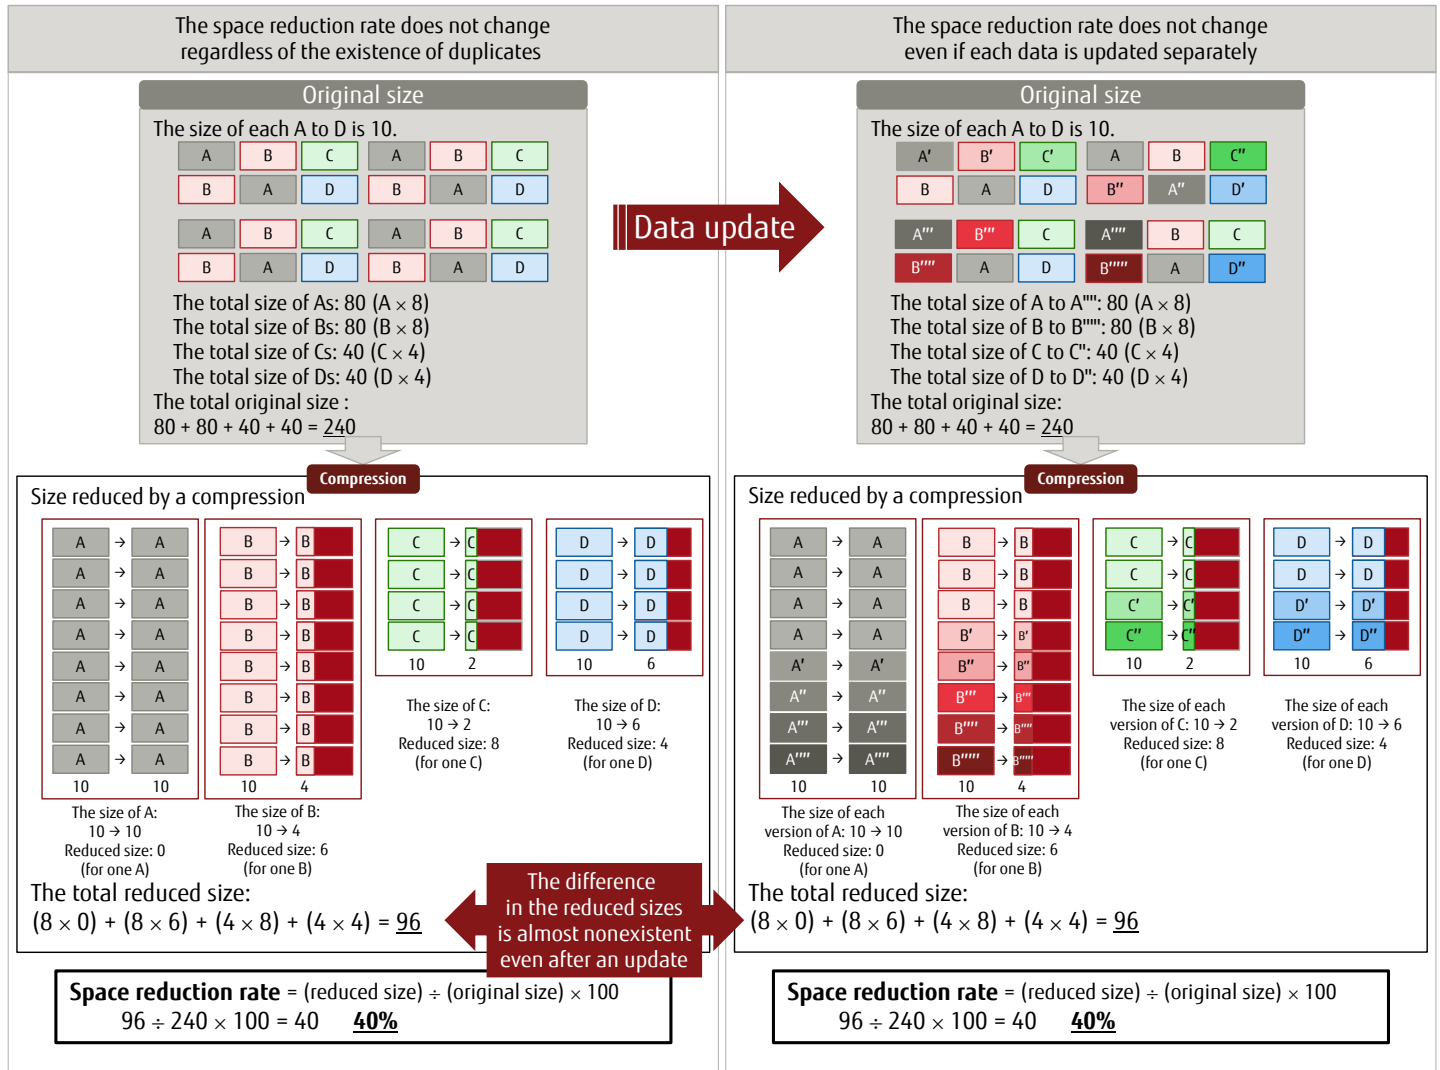


Figure 2-2 Fluctuations of the space reduction rate with compression

2.3 Space Reduction Rate Examples for Windows System Areas

This section shows examples of the space reduction rate obtained when the Deduplication/Compression function is applied to the system area where the Windows Server 2016 OS is installed.

Only the OS is installed in the 1TB TPV created in each TPP to compare the three Deduplication/Compression setup patterns.

[For TPPs with deduplication enabled]




Figure 2-3 Space reduction rate with deduplication for Windows Server 2016

When Windows Server 2016 is installed in the TPV with deduplication enabled, the space reduction rate is 2%.

[For TPPs with compression enabled]




Figure 2-4 Space reduction rate with compression for Windows Server 2016

When Windows Server 2016 is installed in the TPV with compression enabled, the space reduction rate is 11%.

[For TPPs with deduplication and compression enabled]

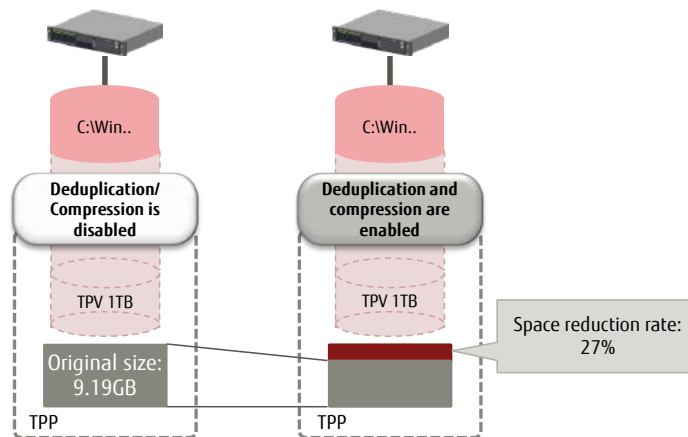


Figure 2-5 Space reduction rate with deduplication and compression for Windows Server 2016

When Windows Server 2016 is installed in the TPV with deduplication and compression enabled, the space reduction rate is 27%.

For Windows Server 2016, the space reduction rate with deduplication is 2%. This indicates that little or no duplicates exist. The space reduction rate with compression is 11%. This indicates that compression is slightly effective. However, with deduplication and compression, the space reduction rate is 27%.

To help understand how the Compression function is effective, this section provides the file configuration of the system area that was used for the verification test.

The total number of files in Windows Server 2016 is 104,986 and 59.2% of those files consists of four file formats (or extensions). Because the fifth most common format accounts for only 3.7% of the total, most of the files consist disproportionately of the four formats. The total size of these four formats accounts for 26.1% or approximately one fourth of the entire system area.

The number of .manifest files (text files) is large (24,044). However, the total size is only 0.1% of the entire system area. Because the space reduction rate is up to 0.1%, compression has little effect for this file type.

The following table shows the number of common files based on the extension.

Extension	Number of files	Rate for the number of files (Total number of files: 104,986)	Total file size (Mbyte)	Size rate (Total file size: 33,833MB)
manifest	24,044	22.9%	19	0.1%
dll	17,665	16.8%	8,351	24.7%
cat	11,065	10.5%	188	0.6%
mui	9,413	9.0%	223	0.7%
mum	3,890	3.7%	8	0.0%

Table 2-1 Number of files based on the extension for Windows Server 2016

When file sizes are compared, .sys and .dll files account for 83.6% or approximately four fifths of the entire capacity.

Because a single system area with no cloned VMs was used for the verification test, deduplication had little effect.

The capacity was reduced mainly due to the compression function. Generally, .sys and .dll files keep their sizes even after compression is performed because they are mainly comprised of object code. However, when the compression function is applied to the system area, a certain amount of space is reduced.

The following table shows the top five total file sizes based on the extension.

Extension	Number of files	Rate for the number of files (Total number of files: 104,986)	Total file size (Mbyte)	Size rate (Total file size: 33,833MB)
sys	1,312	1.2%	19,920	58.9%
dll	17,665	16.8%	8,351	24.7%
exe	2,521	2.4%	558	1.6%
lex	311	0.3%	542	1.6%
ttc	40	0.0%	509	1.5%

Table 2-2 File size based on the extension for Windows Server 2016

The following table shows the rates for text format files in the Windows Server 2016 system area.

* Text format files are displayed with "text" (for example "ASCII text") when the Linux "file" command is executed. This includes txt, htm, and xml files.

File format	Number of files	Rate for the number of files (Total number of files: 104,986)	Total file size (Mbyte)	Size rate (Total file size: 33,833MB)
Text format	23,021	21.9%	412	1.2%
Others	81,965	78.1%	33,421	98.8%

Table 2-3 File size based on text formats for Windows Server 2016

2.4 Space Reduction Rate Examples for System Areas of Red Hat Enterprise Linux

This section shows examples of the space reduction rate obtained when the Deduplication/Compression function is applied to the system area where the Red Hat Enterprise Linux 7.x OS is installed.

Only the OS is installed in a 1TB TPV created in each TPP to compare the three Deduplication/Compression setup patterns.

[For TPPs with deduplication enabled]




Figure 2-6 Space reduction rate with deduplication for Red Hat Enterprise Linux 7.x

When Red Hat Enterprise Linux 7.x is installed in the TPV with deduplication enabled, the space reduction rate is 30%.

[For TPPs with compression enabled]




Figure 2-7 Space reduction rate with compression for Red Hat Enterprise Linux 7.x

When Red Hat Enterprise Linux 7.x is installed in the TPV with compression enabled, the space reduction rate is 33%.

[For TPPs with deduplication and compression enabled]

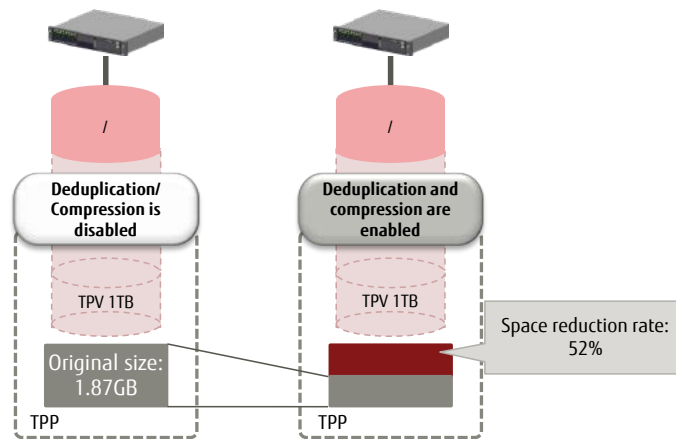


Figure 2-8 Space reduction rate with deduplication and compression for Red Hat Enterprise Linux 7.x

When Red Hat Enterprise Linux 7.x is installed in the TPV with deduplication and compression enabled, the space reduction rate is 52%.

For Red Hat Enterprise Linux 7.x, the space reduction rate with deduplication is 30%. This indicates that duplicates exist. The space reduction rate with compression is 33%. This indicates that compression is also effective. The space reduction rate with deduplication and compression is 52%. This indicates that half the data size can be reduced.

3 Estimating the Size Reduced with Deduplication/Compression

This chapter describes a method to consider how much size can be reduced from the original size using a single system area with OS example and a system area example of 1000 VMs that are created by cloning.

The term "safety factor", which is used in this manual, is described below.

When system designers take the effects of the Deduplication/Compression function into consideration to determine the required number of drives, they usually allow for some wiggle room. The rate of the wiggle room is referred to as "safety factor".

$$\text{Safety factor} = \text{size predicted to fluctuate} \div \text{reduced size} \times 100$$

Half the 100GB capacity (= 50GB) is assumed to be reduced by deduplication.

If a size of 30GB is assumed to fluctuate by operations, the safety factor is 60% (30GB ÷ 50GB × 100).

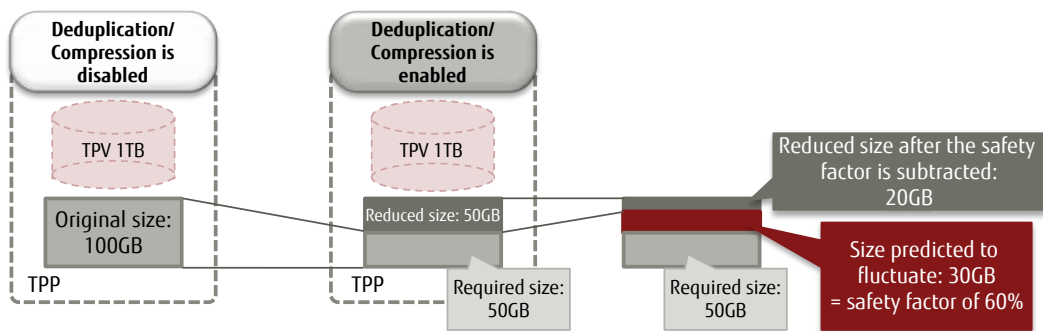


Figure 3-1 Safety factor

If the size is not assumed to fluctuate due to operations (and no extra size is required), the safety factor is 0% (0GB ÷ 50GB × 100).

If no reduction can be expected due to fluctuations by operations, the safety factor is 100% (50GB ÷ 50GB × 100).

3.1 Estimating the Required Physical Size for Windows Server 2016

This section describes a method for estimating the required storage size of a single system area installed with Windows Server 2016. The safety factors used in the following sections are example values. Obtain an appropriate value according to the data property and the frequency of data updates.

Windows Server 2016 is installed in a 1TB TPV created in a TPP.

[For TPPs with deduplication enabled]

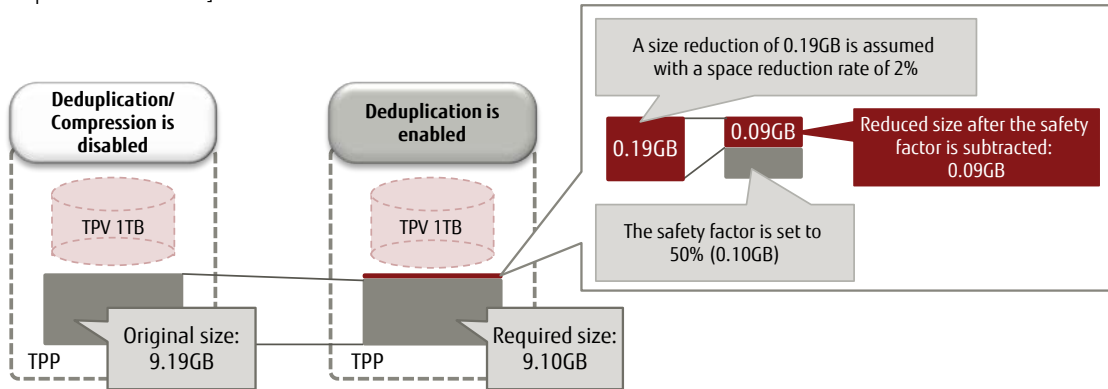


Figure 3-2 Reduced size with deduplication

Space reduction rate with deduplication: 2% (from the verification example in Chapter 2)
 Safety factor: 50% (example)

Calculation: Reduced size after the safety factor is subtracted =

$$\text{original size} \times \text{space reduction rate with deduplication} \div 100 \times (1 - \text{safety factor} \div 100)$$

$$9.19\text{GB} \times 2 \div 100 \times (1 - 50 \div 100) = 0.09\text{GB}$$

Required size after deduplication = original size - reduced size

$$9.19\text{B} - 0.09\text{GB} = 9.10\text{GB}$$

Reduced size: 0.09GB

Required size after deduplication: 9.10GB

Conclusion: A storage size of 9.19GB is required when the deduplication setting is disabled. Enabling this setting reduces the required storage size to 9.10GB.

If the reduced size is small, using the original size is an option by assuming that the size cannot be reduced.

[For TPPs with compression enabled]

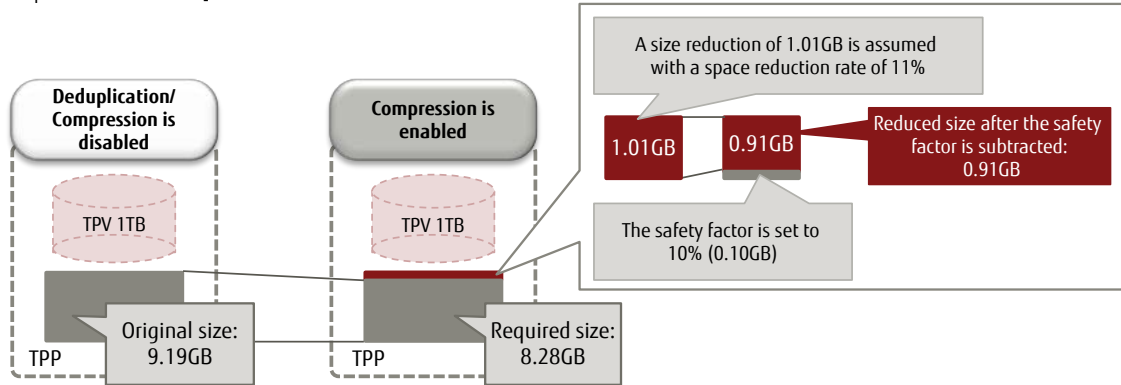


Figure 3-3 Reduced size with compression

Space reduction rate with compression: 11% (from the verification example in Chapter 2)
 Safety factor: 10% (example)

Calculation: Reduced size after the safety factor is subtracted =

$$\text{original size} \times \text{space reduction rate with compression} \div 100 \times (1 - \text{safety factor} \div 100)$$

$$9.19\text{GB} \times 11 \div 100 \times (1 - 10 \div 100) = 0.91\text{GB}$$

Required size after compression = original size - reduced size

$$9.19\text{B} - 0.91\text{GB} = 8.28\text{GB}$$

Reduced size: 0.91GB

Required size after compression: 8.28GB

Conclusion: A storage size of 9.19GB is required when the compression setting is disabled. Enabling this setting reduces the required storage size to 8.28GB.

[For TPPs with deduplication and compression enabled]
 When both the deduplication and compression settings are enabled, deduplication is performed first.
 The reduced size with compression is calculated by using the required size after deduplication.

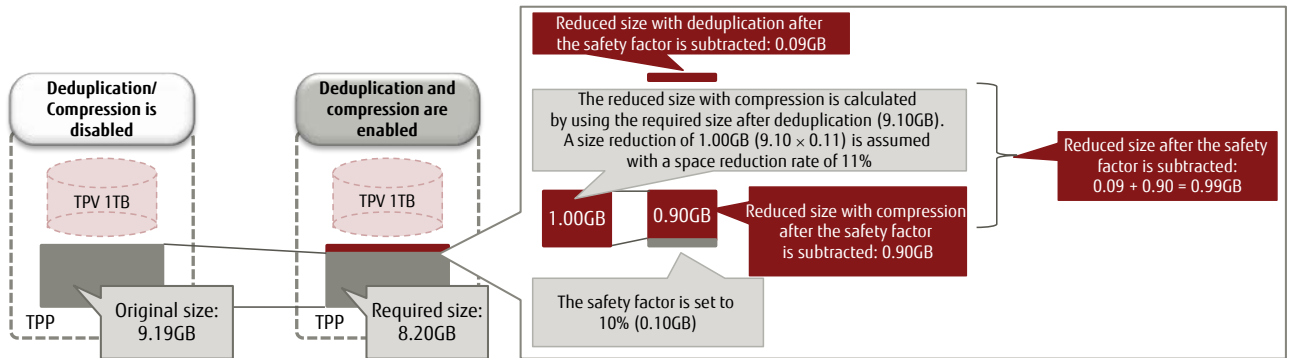


Figure 3-4 Reduced size with deduplication and compression

Required size after deduplication: 9.10GB (from the calculated result in the [For TPPs with deduplication enabled] section)
 Space reduction rate with compression: 11% (from the verification example in Chapter 2)
 Safety factor: 10% (example)

Calculation: Reduced size after the safety factor is subtracted =
 reduced size with deduplication +
 reduced size with compression [required size after deduplication × space reduction rate with
 compression ÷ 100 × (1 - safety factor ÷ 100)]
 $0.09\text{GB} + [9.10\text{GB} \times 11 \div 100 \times (1 - 10 \div 100)] = 0.99\text{GB}$

Required size after deduplication and compression = original size - reduced size
 $9.19\text{GB} - 0.99\text{GB} = 8.20\text{GB}$

Reduced size: 0.99GB

Required size after deduplication and compression: 8.20GB

Conclusion: A storage size of 9.19GB is required when the deduplication and compression settings are disabled. Enabling these settings reduces the required storage size to 8.20GB.

3.2 Reduced Size for System Areas of 1000 Virtual Machines Created by Cloning

This section describes a method for estimating the required storage size of the system areas for 1000 VMs when Windows Server 2016 is installed in the system areas. Windows Server 2016 is installed in a 100GB TPV created in a TPP and the TPV is used as the clone source to create 999 VM clones in the TPP.

[For TPPs with deduplication enabled]

Some effects by deduplication can be expected for VM clones. However, because 1000 VMs individually update the data, the safety factor must be set considering the space reduction rate that may decrease due to operations.

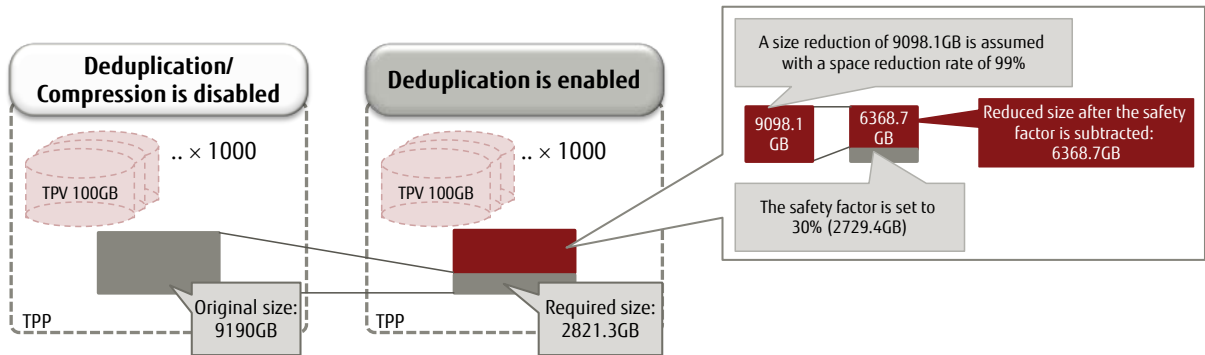


Figure 3-5 Reduced size with deduplication (VM clone)

Space reduction rate with deduplication: 99% (Duplicates exist for the number of clones.)
 Safety factor: 30% (example)

$$\begin{aligned} \text{Calculation: Reduced size after the safety factor is subtracted} = & \\ & (\text{original size for a VM}) \times \text{number of VMs} \times \\ & \text{space reduction rate with deduplication} \div 100 \times (1 - \text{safety factor} \div 100) \\ & 9.19\text{GB} \times 1000 \times 99 \div 100 \times (1 - 30 \div 100) = 6368.7\text{GB} \end{aligned}$$

$$\begin{aligned} \text{Required size after deduplication} = & (\text{original size for a VM}) \times \text{number of VMs} - \text{reduced size} \\ & 9.19\text{GB} \times 1000 - 6368.7\text{GB} = 2821.3\text{GB} \end{aligned}$$

Reduced size: 6368.7GB

Required size after deduplication: 2821.3GB

Conclusion: A storage size of 9190GB is required when the deduplication setting is disabled. Enabling this setting reduces the required storage size to 2821.3GB.

[For TPPs with compression enabled]

For the system areas of 1000 VMs, set the safety factor low because the space reduction rate with compression rarely fluctuates. If the number of VMs increases, simply multiply by the number of VMs because the compression rate does not change.

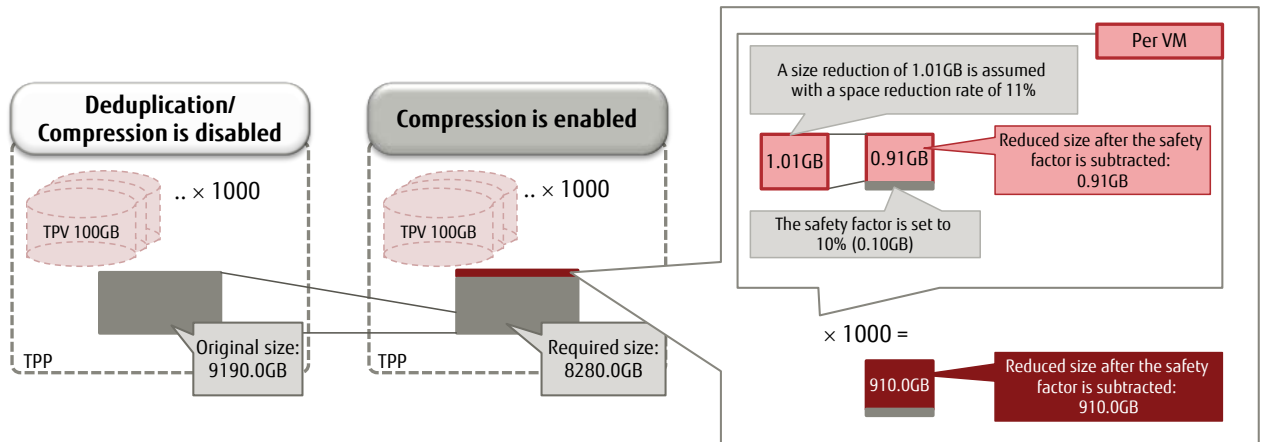


Figure 3-6 Reduced size with compression (VM clone)

Reduced size for a VM: 0.91GB (from the value in [For TPPs with compression enabled] of section 3.1)

Calculation: Reduced size after the safety factor is subtracted = reduced size for a VM × number of VMs
 $0.91\text{GB} \times 1000 = 910.0\text{GB}$

Required size after compression = (original size for a VM) × number of VMs - reduced size
 $9.19\text{GB} \times 1000 - 910.0\text{GB} = 8280.0\text{GB}$

Reduced size: 910.0GB

Required size after compression: 8280.0GB

Conclusion: A storage size of 9190.0GB is required when the compression setting is disabled. Enabling this setting reduces the required storage size to 8280.0GB.

[For TPPs with deduplication and compression enabled]

The required size can be calculated using a reduced size with deduplication for 1000 VMs + a reduced size with compression for a VM because deduplication is performed first.

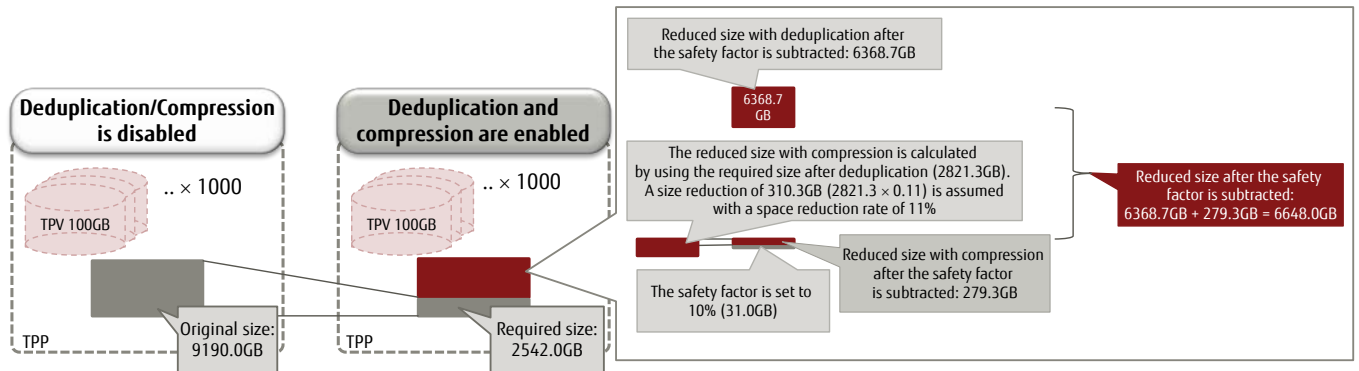


Figure 3-7 Reduced size with deduplication and compression (VM clone)

Required size after deduplication: 2821.3GB (from the calculated result in the [For TPPs with deduplication enabled] section)

Space reduction rate with compression: 11%

Safety factor: 10% (example)

Calculation: Reduced size after the safety factor is subtracted =
 reduced size with deduplication +
 reduced size with compression [required size after deduplication ×
 space reduction rate with compression ÷ 100 × (1 - safety factor ÷ 100)]
 $6368.7\text{GB} + [2821.3\text{GB} \times 11 \div 100 \times (1 - 10 \div 100)] = 6648.0\text{GB}$

Required size after deduplication and compression =
 (original size for a VM) × number of VMs - reduced size
 $9.19\text{GB} \times 1000 - 6648.0\text{GB} = 2542.0\text{GB}$

Reduced size: 6648.0GB

Required size after deduplication and compression: 2542.0GB

Conclusion: A storage size of 9190.0GB is required when the deduplication and compression settings are disabled. Enabling these settings reduces the required storage size to 2542.0GB.

3.3 Reduced Size for Data Areas

While system areas store systems such as the OS, data areas store only data. Data areas contain data generated or updated by application programs. The effectiveness of Deduplication/Compression depends on the program.

Whether the Compression function can reduce the size of data areas can be easily tested using a basic compression function such as ZIP. If the examples in this document apply, use them as a reference.

If the examples do not apply or the effects cannot be assessed, the assumption is that the size cannot be reduced. In this case, using Deduplication/Compression is worth considering because in some cases a certain amount of data can be reduced such as when .sys and .dll files are compressed in Windows.

3.4 Summary of Reduced Size

The safety factor can be set low for the Compression function because a space reduction rate with compression rarely fluctuates. Using fewer drives for a backup saves storage space.

If a high space reduction rate is expected with the Deduplication function such as when clones are used, the space reduction rate increases as the number of servers and VMs increases. When the number of servers and VMs is large, the number of required drives can be drastically reduced. This significantly reduces the overall cost as well.

4 Costs Associated with SSDs

Even though SSDs are more expensive than HDDs, the difference in costs becomes small by a space reduction with deduplication and compression. The savings can be seen in the reduced number of required drives and running costs due to smaller installation space and power consumption. For this reason, the total cost is not much different from when HDDs are used.

5 Summary

The ETERNUS AF/DX allows users to set deduplication and compression separately. Selecting the best setup pattern provides efficient use of the limited storage space in the drives. Choose the ETERNUS AF/DX, the best suited for cost reduction.

Contact
FUJITSU Limited
Website: <https://www.fujitsu.com/eternus/>

■ Trademarks
Microsoft®, Windows, and other Microsoft's product names are registered trademarks or trademarks of Microsoft Corporation in the United States, and other countries. Names and logos related to Red Hat are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries. The product names and company names in this document are registered trademarks or trademarks of their respective companies. VMware and VMware product names are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Trademark symbols such as (R) and (TM) may be omitted from system names and product names in this document.

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