White Paper
ETERNUS AF series All-Flash Arrays Best Suited for Databases

Typically, complicated RAID designs are required for hard disk arrays to secure high-speed processing of databases. However, all-flash arrays can consolidate a large number of databases without complex RAID designs. This document verifies the benefits of consolidating databases by using an example of a consolidation in a SQL Server 2016 configuration that runs on Microsoft Windows Server 2016.

Table of Contents

List of Figures

List of Tables

Preface

The data handled in our information society continues to grow in volume, detail, and length of storage on a daily basis. Just as databases are essential to utilizing this huge volume of data, storage systems are indispensable for managing and archiving all of this data.
Most of the data in databases is stored in a storage system. This makes the performance of the database highly dependent on that of the storage system.
Conventionally, each database system requires a dedicated storage system. However, with all-flash arrays, multiple databases can be consolidated in a single storage system.
This document describes database consolidation in detail including the points to consider for consolidating databases.
The verification results described in this document are current as of January 2017.

The product lineup and product information stated in this document are current as of December 2017.

■ Target Readers

This document targets the following readers.

- Those who have designed and constructed infrastructure but are not familiar with database design
- Those who have designed databases but are not familiar with infrastructure design and construction

■ Applicable Models

This document targets the following storage system model.

- FUJITSU Storage ETERNUS AF250 S2/AF650 S2 and ETERNUS AF250/AF650

■ Naming Conventions

The following abbreviations are used in this document.

- ETERNUS AF series .... FUJITSU Storage ETERNUS AF250 S2/AF650 S2 and ETERNUS AF250/AF650 (All-Flash Arrays)
- ETERNUS AF650 .... FUJITSU Storage ETERNUS AF650
- ETERNUS AF250 .... FUJITSU Storage ETERNUS AF250
- PRIMERGY .... FUJITSU Server PRIMERGY
- SSD .... Solid State Drive
- TPP .... Thin Provisioning Pool
- TPV .... Thin Provisioning Volume

■ Terminology

The following terms are used in this document.

- Database buffer cache .. Memory areas in the server memory that are separately allocated to each database
- Cache hit .... Required data for database processing is stored in the database buffer cache and can be immediately read from the database buffer cache
- Cache miss .... Required data for database processing is not stored in the database buffer cache so the disk, such as HDDs and SSDs, is accessed

1 Realizing High-Speed Processing Best Suited for Databases

Databases need to manage large volumes of data and access disks frequently. Therefore, the processing speed of the entire database depends on the performance of the storage system as well as the performance of the server. General all-flash arrays do not allow SSDs to deliver their full performance potential because of the deduplication/compression function. However, Fujitsu's ETERNUS AF series can simultaneously handle the areas that use and the areas that do not use the Deduplication/Compression function. By using the areas that do not use the Deduplication/Compression function for databases, the ETERNUS AF series can maximize the performance of SSDs. The ETERNUS AF series is best suited for databases because it supports high-speed database processing and can also run multiple databases at the same time.

In general, databases perform high-speed processing by using the server memory. However, the moment a transaction is committed, the storage system is loaded with a huge amount of data access. This access to the storage system is the most time consuming process during a database processing. All-flash arrays handle the access using the flash memory to deliver high-speed processing.

All-flash arrays can deliver exceptionally high performance even after multiple databases are consolidated in the storage system. However, to ensure smooth operation after consolidating, several points must be considered before consolidating existing databases.

The following chapters provide the points to consider for smooth operation after consolidating.

2 Prior to Database Consolidation

Databases may not provide expected responses by consolidation alone. The factors for unsatisfactory responses are extremely difficult to identify once databases have been consolidated.
This chapter describes the points to consider before performing a database consolidation.

2.1 Database Characteristics

Databases are required to process data at high speed. For high-speed processing, servers are installed with a large amount of memory and the database buffer cache is allocated with a large memory size to increase the cache hit rate. If the database buffer cache is insufficient, the cache hit rate decreases and the data in the storage system is frequently accessed. This causes a delay to the entire database processing.

Databases synchronize the data between the memory and the disks after a certain number of processes are performed. At that time, a large number of disk I/Os are issued simultaneously. In order to deal with these large number of disk I/Os at the same time, the storage system must have high performance. If the performance of the storage system is low, the server has to wait for the disk I/Os to be processed. As a result, the entire database processing is delayed.

In addition to the server performance, the performance of the storage system directly affects the processing speed of the entire database that manages large volumes of data and issues a large number of disk I/Os.

2.2 Characteristics of Database Resources Depending on the Usage

Databases operate using resources that include the server's CPU and memory as well as storage systems. The use of these resources must be balanced to achieve efficient, high-speed operations. The following table shows the operation speed for various combinations of resources.

CPUs with low processing power are not suitable for databases. Even if the processing power is high, the database response deteriorates at an accelerated pace if the memory capacity is insufficient and the storage response is low.

CPU processing power	Memory capacity (Database buffer cache size)	Storage response	Database operation	Operation speed
High (High-speed processing is achieved)	Large (Cache hit rate is high and high-speed processing is achieved)	Fast (High-speed processing is achieved even for cache miss data)	High-speed operation is achieved	High-speed
		Slow (I/O load for cache miss data becomes a bottleneck)	Processing is delayed	
	Small (Cache miss increases and processing is delayed)	Fast (High-speed processing is achieved even for cache miss data)	Processing is delayed (However, the delay is minimized because of the ultrahigh-speed storage response of the all-flash arrays)	
		Slow (I/O load for cache miss data becomes a bottleneck)	Processing is delayed (Increased cache miss + I/O load bottleneck)	Low-speed
Low			Not suitable for databases	

Table 2-1 Operation characteristics depending on resources

As shown above, the minimum requirements that must be sufficiently secured for database consolidation are CPU, memory, and storage system performance.

3 Checks for Consolidating Databases

Databases operate using the CPU and memory of the server and the storage resources. The amount of resources and the balance between them must be considered to maintain high-speed operation after consolidating. The points to consider are the CPU, memory, and storage system.

To secure enough post-consolidation resources, the pre-consolidation conditions must be checked.

This chapter explains what, when, and how to check the conditions for securing post-consolidation resources.

3.1 Items to Check

Information about the CPU, memory, and storage system must be obtained before consolidating because after consolidating, a small number of servers or storage systems will handle environments that are run on several different servers.

For the CPU, check the utilization, the clock frequency, and the number of cores.

Constantly high CPU utilization causes process delays. To avoid this situation after consolidating, check the time period when the system load is high as well as the CPU utilization during the period.

In general, around 50% is an appropriate rate. When the utilization is 75% or more, it is close to the limit. When the utilization is 20% or less, consolidation of the servers should be considered because the servers have extra power.

The CPU clock frequency and the number of CPU cores must also be checked before consolidating because they are required for estimating the post-consolidation CPU utilization. The time period to obtain the information is irrelevant because the CPU clock frequency and the number of CPU cores are hardware-specific information.

For the memory, check the available capacity or usage rate, the physical capacity, and the size of the database buffer cache.

High memory usage rate or insufficient available memory causes process delays. To avoid this situation after consolidating, check the time period when the system load is high as well as the available memory capacity during the period.

The physical memory capacity and the size of the database buffer cache must also be checked because they are required for estimating the post-consolidation memory capacity. The time period to obtain the information is irrelevant because the physical memory capacity and the size of the database buffer cache do not change dynamically.

The available memory capacity indicates the physical memory capacity that is currently available for use.

If the value calculated by (available memory capacity ÷ physical memory capacity)*100 is under 10% at any given time, consider adding memory.

For the storage system, check the IOPs, the response time, and the database disk usage.

High IOPs or late response times during processing cause process delays. To avoid this situation after consolidating, check the time period when the system load is high as well as the IOPs and the response time during the period.

In general, if the response time exceeds 20ms, the process delay is significant.

To estimate the post-consolidation performance, check the IOPs and the response time of the post-consolidation storage system.

If the total amount of disk space used in each pre-consolidated database exceeds the capacity of disks in the post-consolidation storage system, consolidation is not possible. Therefore, the disk usage and the additional amount must be checked for each database.

3.2 Timing and Period for Checks

For the values of the CPU utilization, the available memory capacity, the IOPs, and the response time, know when and how long the load becomes high for each system.

Check the information during the following time period.

- For systems with constant load, one business hour in the morning and afternoon
- For systems where daily, weekly, or monthly processes are required, the time period when those processes are performed (including approximately two hours before and after those processes)

3.3 Procedure for Checking Items in Pre-Consolidated Environments

3.3.1 Windows Environments

Check the resources in the Windows environments as follows.

3.3.1.1 CPU Clock Frequency and Number of CPU Cores

This section provides an explanation using Windows Server 2012 R2 as an example. Start Task Manager to check the items.

1. Right-click an empty area on the taskbar and select [Task Manager] to start Task Manager.
   If the details are not displayed, click "More details" on the bottom-left side of the screen to display them.
2. Click the [Performance] tab and select "CPU" to display the CPU clock frequency and the number of cores.

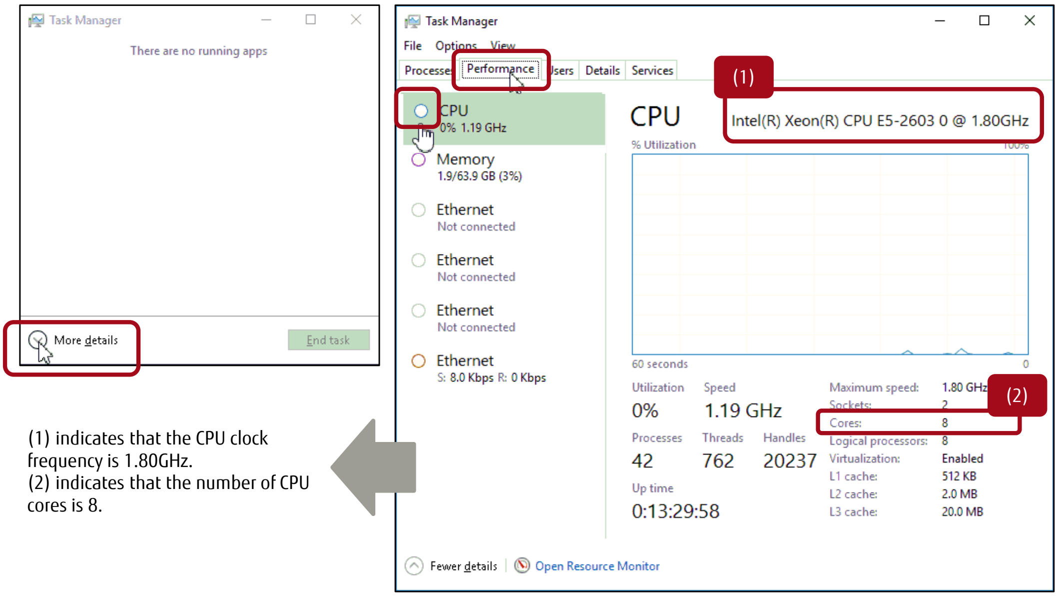

Figure 3-1 Checking the CPU clock frequency and the number of CPU cores

3.3.1.2 Physical Memory Capacity

Click the [Performance] tab to check the physical memory capacity in the same manner as the CPU clock frequency and the number of CPU cores.

Then, select "Memory" to display the physical memory capacity.

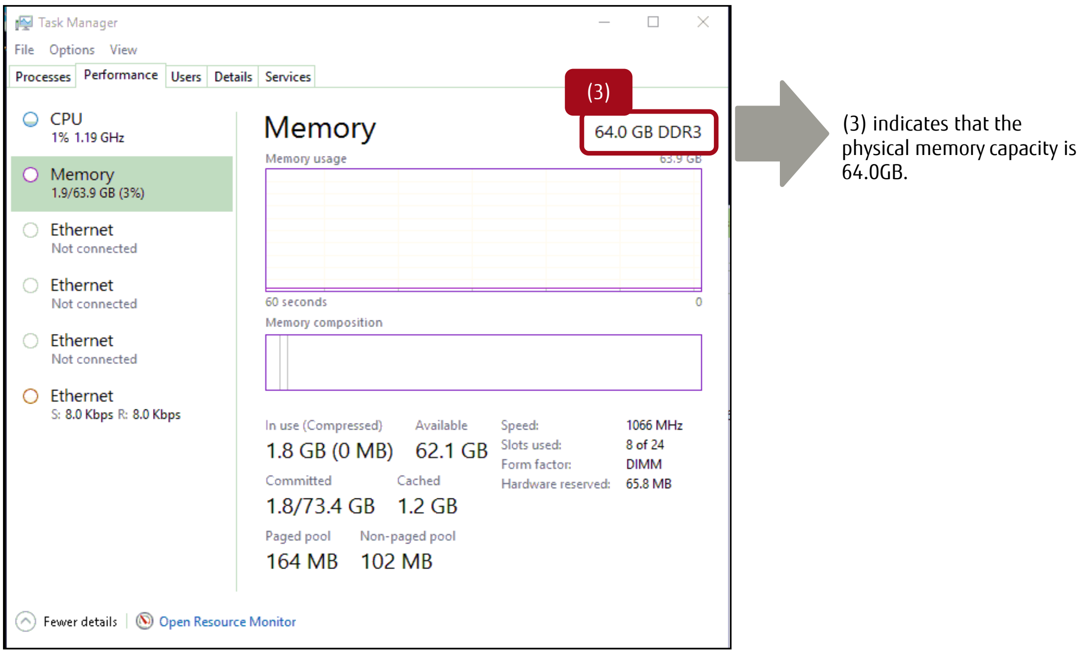

Figure 3-2 Checking the physical memory capacity

3.3.1.3 Database Buffer Cache Size

This section provides an explanation using Microsoft SQL Server 2014 as an example. Check the cache size using "Microsoft SQL Server Management Studio".

1. Start "SQL Server 2014 Management Studio" from the Windows start menu.
2. Right-click the server name and select [Properties] to open the [Server Properties] screen.
3. Select "Memory" from "Select a page" in the left pane to display the size of the database buffer cache.

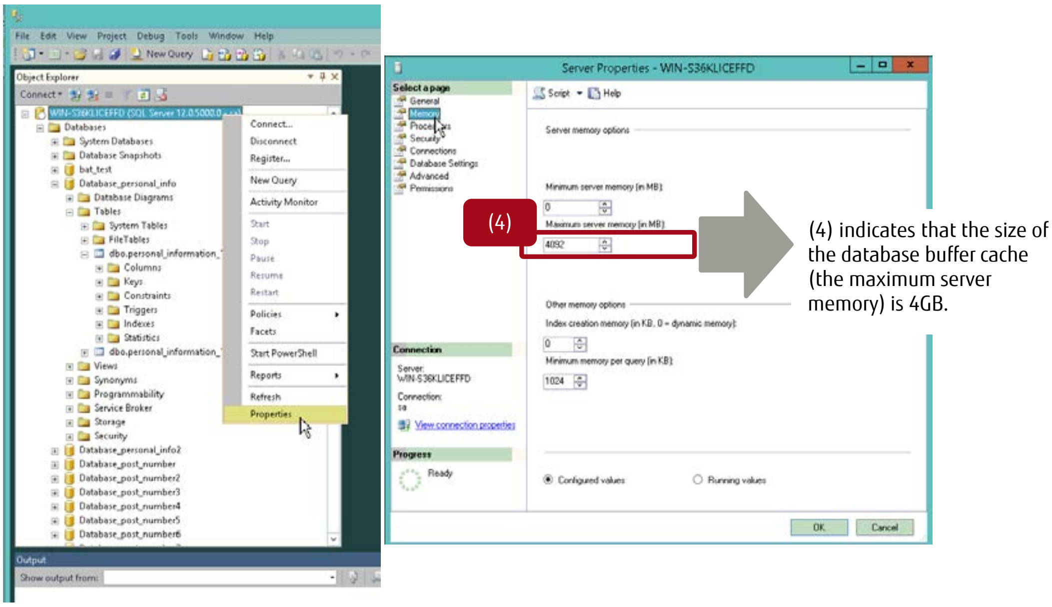

Figure 3-3 Checking the size of the database buffer cache

3.3.1.4 Database Disk Usage

This section provides an explanation using Microsoft SQL Server 2014 as an example. Check the disk usage using "Microsoft SQL Server Management Studio".

1. Start "SQL Server 2014 Management Studio" from the Windows start menu.
2. Right-click the database name and select "Reports" - "Standard Reports" - "Disk Usage" to open the [Disk Usage] screen.
   Check the value for "Total Space Reserved". This value includes spaces in the data areas and the log areas.

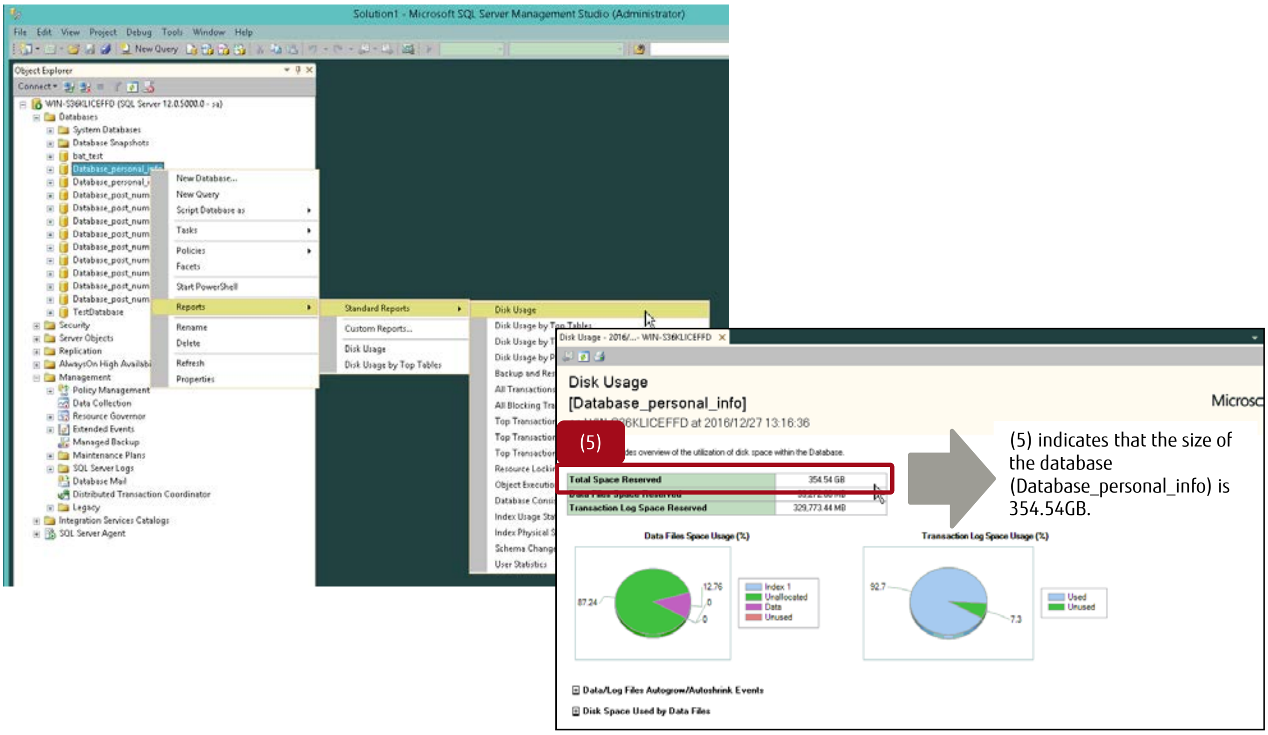

Figure 3-4 Checking the disk usage in the database

3.3.1.5 CPU **Utilization**, Available Memory Capacity, and Response Time

Use the "logman" commands of Windows to obtain the information during the period described in "3.2 Timing and the Period to Check ".
This information is called "performance log". Each data included in the performance log is called "counter".
The CPU utilization, available memory capacity, and response time can be checked by defining them as counters.

1. Defining the counters

   1-1. Preparing a text file

Acquisition purpose	Value to define
Checking the CPU utilization	“\Processor(*)\% Processor Time”
Checking the available memory capacity	“\Memory\Available MBytes”
Checking the response time	“\PhysicalDisk(*)\Avg. Disk sec/Read” “\PhysicalDisk(*)\Avg. Disk sec/Write” “\PhysicalDisk(*)\Avg. Disk sec/Transfer” Enter the disk number and the drive letter in (*). [explained below]

Table 3-1 Defining counters (acquisition purpose and defining values)

For the values to de defined for checking the response time, define counters for each disk where the database is stored.
The following example indicates that C:\ is the target disk and the disk number of C:\ is "0".

```
c:\tmp>type counters.txt
"\Processor(*)\% Processor Time"
"\Memory\Available MBytes"
"\PhysicalDisk(0 C:)\Avg. Disk sec/Read"
"\PhysicalDisk(0 C:)\Avg. Disk sec/Write"
"\PhysicalDisk(0 C:)\Avg. Disk sec/Transfer"

c:\tmp>
```

Figure 3-5 Text file that describes the counters

Refer to Performance Monitor to check the relationship between the drive letter and the disk number that are used for checking the response time. Use extreme caution when checking because the combination of the drive letter and the disk number differs depending on the environment.
To start Performance Monitor, select [Administrative Tools] → [Performance Monitor] from the start menu.
After the Performance Monitor screen appears, select [Performance Monitor] on the left of the screen. Select "Properties" on the right side of the screen where the graph of the current activity status is displayed.

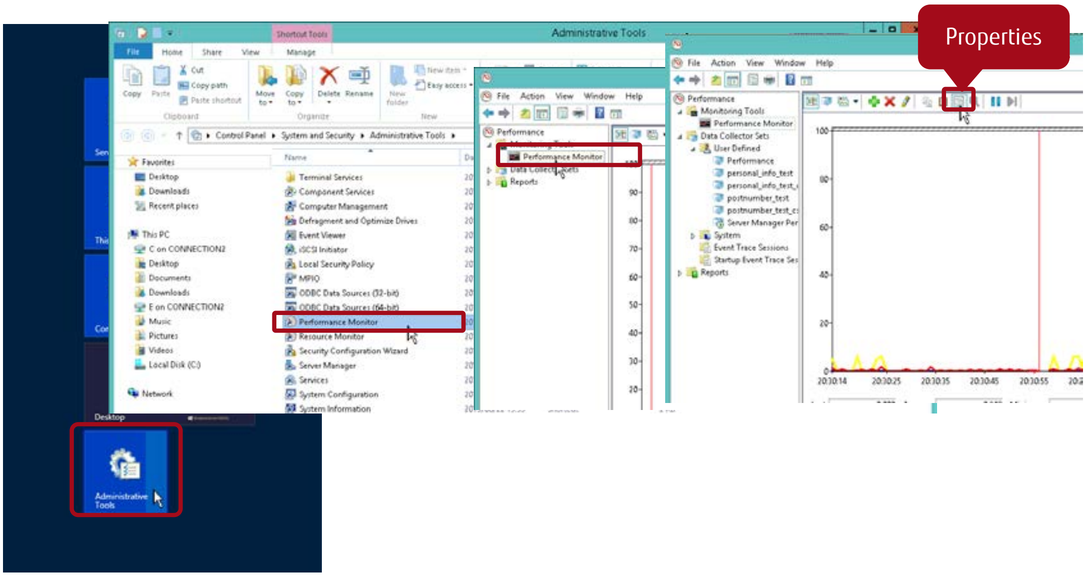

Figure 3-6 Opening Performance Monitor Properties

In the [Performance Monitor Properties] screen, check that the [Data] tab is displayed and click [Add]. Check the relationship between the drive letter and the disk number of the target disk in [Instances of selected object].

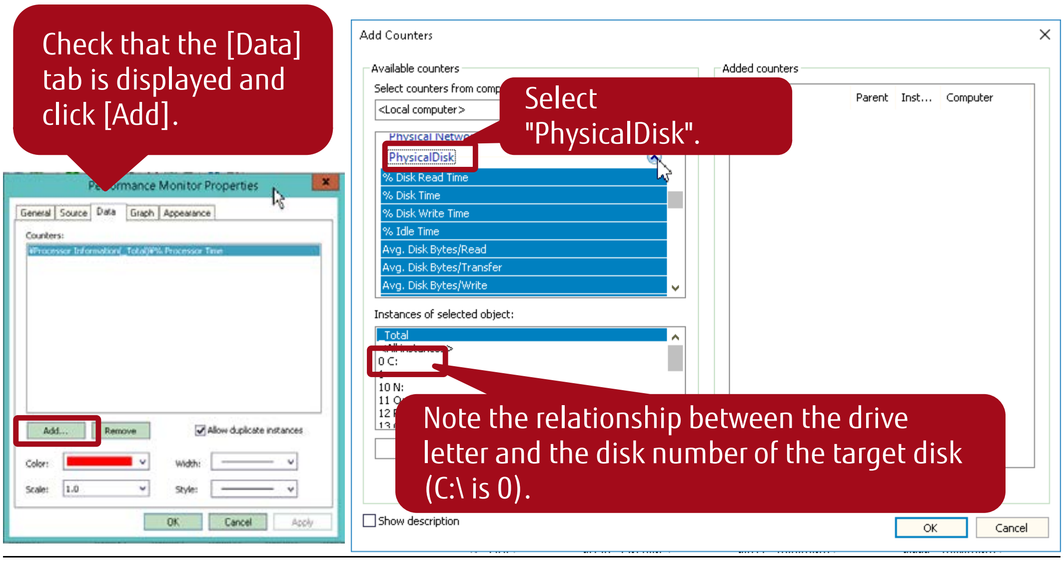

Figure 3-7 Checking the disk number

After taking a note, click [Cancel] to close the opened screen because the Performance Monitor screen is opened only for reference. Enter the noted values in the text file (counters.txt) as described in "Figure 3-5 Text file that describes the counter".

1-2. Defining the counters

To define the counters on Performance Monitor, execute the "logman" commands using the DOS command prompt.

The newly-defined names are "personal_info_test" and "personal_info_test_csv", and the file path is "C:\sqltest\test_personal_info\0001" in the following example.

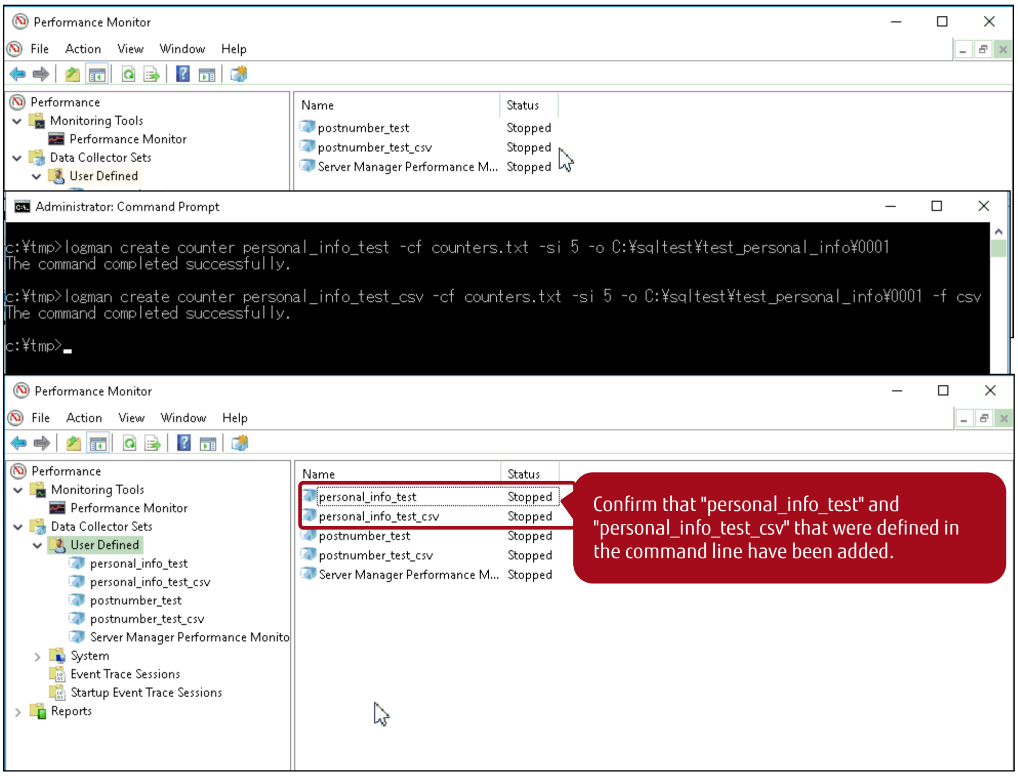

Figure 3-8 Definitions for Performance Monitor

Execute the following "logman" commands using the DOS command prompt.

(a) "logman create counter personal_info_test -cf counters.txt -si 5 -o C:\sqltest\test_personal_info\0001"

(b) "logman create counter personal_info_test_csv -cf counters.txt -si 5 -o C:\sqltest\test_personal_info\0001 -f csv"

(a) indicates that the information is obtained to the file arbitrarily named "personal_info_test" in the format where data can be seen in the Performance Monitor screen.

(b) indicates that the information is obtained to the file arbitrarily named "personal_info_test_csv" in the csv format.

Both (a) and (b) indicate that the information is obtained in five second intervals, the counter definition file is "counters.txt", and the output file is "C:\sqltest\test_personal_info\0001".

After executing the commands, confirm that "personal_info_test" and "personal_info_test_csv" that were defined in the command line have been added in the Performance Monitor screen.

2. Starting the obtainment of the performance log

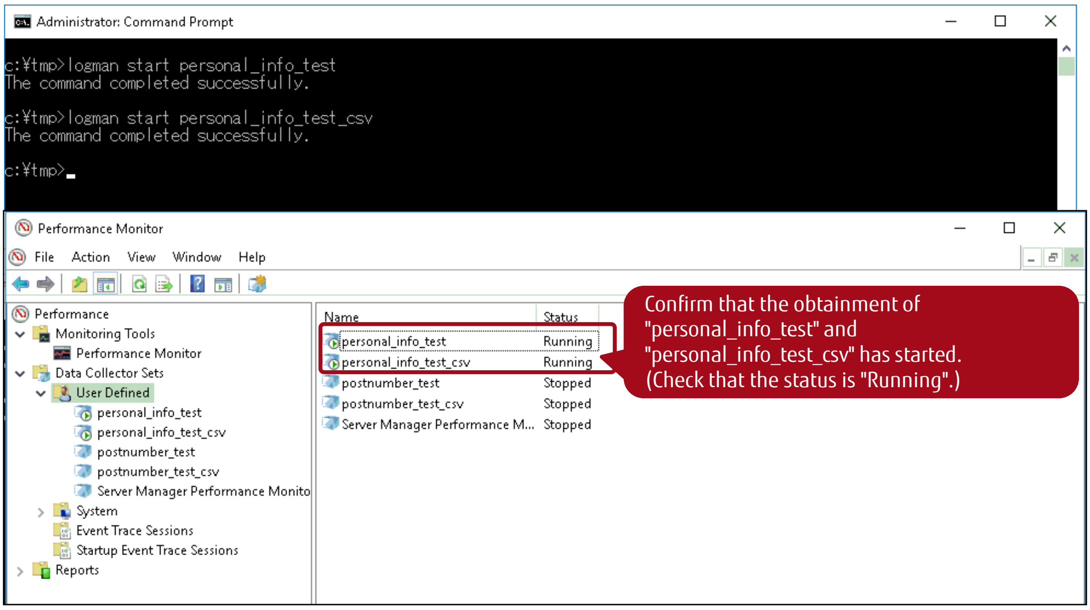

Figure 3-9 Starting the obtainment of the performance log

To start obtaining the performance log, execute the following "logman" commands using the DOS command prompt.
(a) "logman start personal_info_test"
(b) "logman start personal_info_test_csv".

After executing the commands, confirm that "personal_info_test" and "personal_info_test_csv" that were defined in the command line have been started in the Performance Monitor screen.

3. Stopping the obtainment of the performance log

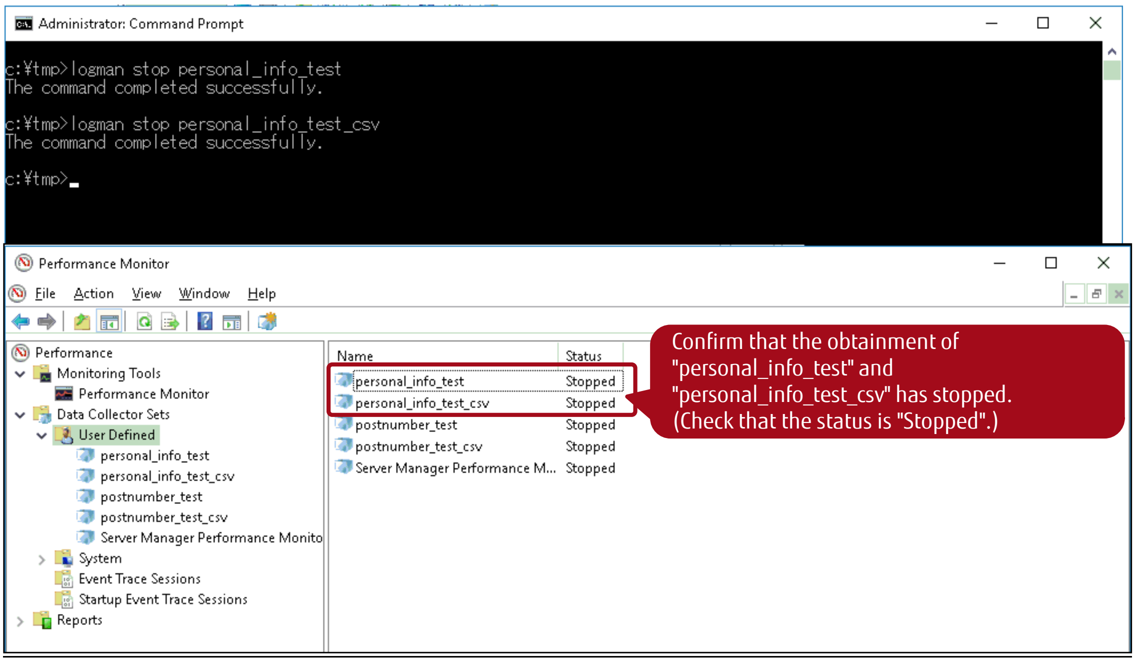

Figure 3-10 Stopping the obtainment of the performance log

After the performance log has been obtained, execute the following "logman" commands using the DOS command prompt to stop obtaining the performance log.
(a) "logman stop personal_info_test"
(b) "logman stop personal_info_test_csv".

After executing the commands, confirm that "personal_info_test" and "personal_info_test_csv" that were defined in the command line have been stopped in the Performance Monitor screen.

4. Checking the performance log information

The blg file is created when the performance log is obtained. Double-click the blg file to start Performance Monitor and then check the performance log information on the opened Performance Monitor screen.

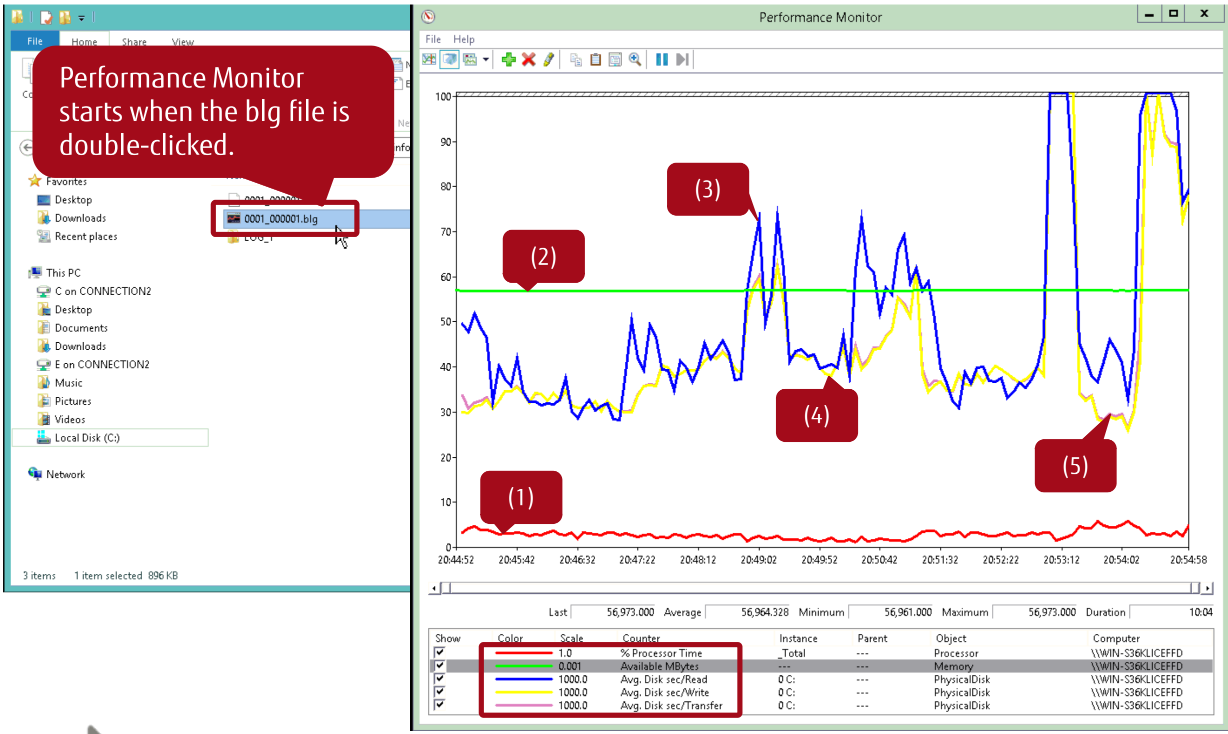

(1) % Processor Time : CPU utilization
Indicates that the CPU utilization is approximately 5% (5/1 = 5 because the value in the graph is 5 with a scale of 1.0).

(2) Avg. Disk sec/Transfer : available memory capacity
Indicates that the available memory capacity is approximately 57GB (58/0.001 = 58000MB ≒ 57GB because the value in the graph is 58 with a scale of 0.001).

(3) Avg. Disk sec/Read : response time (Read)
(4) Avg. Disk sec/Write : response time (Write)
(5) Avg. Disk sec/Transfer : response time (Transfer)
Indicate that the response time is approximately 30ms when it is low (30/1000 = 0.03 seconds because the value in the graph is 30 with a scale of 1000.0), and 100ms or more when it is high (100/1000 = 0.1 seconds because the value in the graph is 100 with a scale of 1000.0).

Figure 3-11 Checking the obtained information using Performance Monitor

To check the obtained information, double-click the "0001_000001.blg" file in "C:\sqltest\test_personal_info\" and start Performance Monitor or refer to the csv file.

3.3.2 Linux Environments

Check the resources in Linux environments as follows. The following sections provide explanations using Red Hat as an example.

3.3.2.1 CPU Clock Frequency

Execute "cat /proc/cpuinfo | grep Hz" from the command line to check the CPU clock frequency.

```
(1) linux ~]# cat /proc/cpuinfo | grep Hz
    name    : Intel(R) Xeon(R) CPU          X7542  @ 2.67GHz
cpu MHz        : 2661.000
model name    : Intel(R) Xeon(R) CPU          X7542  @ 2.67GHz
cpu MHz        : 2661.000
model name    : Intel(R) Xeon(R) CPU          X7542  @ 2.67GHz
cpu MHz        : 2661.000
model name    : Intel(R) Xeon(R) CPU          X7542  @ 2.67GHz
cpu MHz        : 2661.000
[root@linux ~]#
```

(1) Indicates that the CPU clock frequency is 2661MHz (2.67GHz).

Figure 3-12 Checking the CPU clock frequency

3.3.2.2 Number of CPU Cores

Execute "cat /proc/cpuinfo | grep "cpu cores"" from the command line to check the number of CPU cores.

```
[root@linux ~]# cat /proc/cpuinfo | grep "cpu cores"
cpu cores                                  : 6
(2)   es                                   : 6
      es                                   : 6
cpu cores                                  : 6
     :                                     :
cpu cores                                  : 6
cpu cores                                  : 6
cpu cores                                  : 6
cpu cores                                  : 6
[root@linux ~]#
```

(2) indicates that the number of CPU cores is 6.

Figure 3-13 Checking the number of CPU cores

3.3.2.3 CPU Utilization

Execute "sar -u -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" from the command line to check the CPU utilization.

```
[root@linux ~]# sar -u -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05
Linux 3.10.0-327.18.2.  (3)  _64 (sv-openstack.localdomain) 12/05/2016 _x86_64_(24 CPU)

09:00:01 PM   CPU   %user   %nice   %system   %iowait   %steal   %idle
09:10:01 PM   all   1.14    0.00    0.49      0.13      0.00     98.25
09:20:01 PM   all   1.14    0.00    0.49      0.12      0.00     98.26
09:30:01 PM   all   1.12    0.00    0.49      0.12      0.00     98.26
09:40:01 PM   all   1.13    0.00    0.49      0.13      0.00     98.25
09:50:01 PM   all   1.12    0.00    0.49      0.12      0.00     98.27
Average:      all   1.13    0.00    0.49      0.12      0.00     98.26
[root@linux ~]#
```

(3) indicates that the CPU utilization (avg.) is 1.13%.

"sar -u -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" indicates that the CPU usage from 21:00 to 22:00 on the 5th is output.
"%user" indicates the CPU utilization of the users.

Figure 3-14 Checking the CPU utilization

3.3.2.4 Physical Memory Capacity

Execute "cat /proc/meminfo | grep MemTotal " from the command line to check the physical memory capacity.

```
[root@linux ~]# cat /proc/meminfo | grep MemTotal
MemTotal:          65641640 kB
[root@linux ~]#
```

(4)

(4) indicates that the memory capacity is 65641640kB (≒ 64GB).

Figure 3-15 Checking the physical memory capacity

3.3.2.5 Memory Usage Rate

Execute "sar -r -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" from the command line to check the memory usage rate.

```
[root@linux ~]# sar -r -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05
Linux 3.10.0-327.18.2.el7.x86_64 (sv-openstack.localdomain)    12/05/2016    _x86_64_   (24 CPU)

09:00:01 PM kbmemfree kbmemused  %memused kbbuffers  kbcached  kbcommit   %commit  kbactive   kbinact   kbdirty
09:10:01 PM  36639564  29002076     44.18      1860   4727256  61069080     61.56  24742852   2199260        28
09:20:01 PM  36629136  29012504     44.20      1860   4735944  61124560     61.62  24744244   2207452        52
09:30:01 PM  36629464  29012176     44.20      1860   4736444  61078604     61.57  24749344   2207452        56
09:40:01 PM  36625924  29015716     44.20      1860   4736948  61133924     61.63  24749448   2207464        12
09:50:01 PM  36618788  29022852     44.21      1860   4737452  61078112     61.57  24749988   2207484         4
Average:     36628575  29013065     44.20      1860   4734809  61096856     61.59  24747175   2205822        30
[root@linux ~]#
```

(5)

(5) indicates that the memory usage rate (avg.) is 44.2%.

"sar -r -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" indicates that the memory usage from 21:00 to 22:00 on the 5th is output. "%memused" indicates the memory usage rate.

Figure 3-16 Checking the memory usage rate

3.3.2.6 Response Time

Execute "sar -p -d -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" from the command line to check the response time.

```
[root@linux ~]# sar -p -d -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05
Linux 3.10.0-327.18.2.el7.x86_64 (sv-openstack.localdomain)    12/05/2016    _x86_64_   (24 CPU)

09:00:01 PM     DEV     tps  rd_sec/s  wr_sec/s  avgrq-sz  avgqu-sz     await   svctm   %util
09:10:01 PM     sda    9.92      0.00    304.70     30.73      0.16     15.91    4.39    4.36
09:10:01 PM     sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00
09:20:01 PM     sda    9.85      0.00    321.31     32.62      0.16     16.02    4.30    4.24
09:20:01 PM     sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00
09:30:01 PM     sda    9.69      0.00    287.93     29.73      0.16     16.16    4.59    4.44
09:30:01 PM     sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00

09:30:01 PM     DEV     tps  rd_sec/s  wr_sec/s  avgrq-sz  avgqu-sz     await   svctm   %util
09:40:01 PM     sda   10.71      0.00    291.62     27.24      0.15     14.02    4.81    5.15
09:40:01 PM     sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00
09:50:01 PM     sda    9.91      0.00    272.71     27.52      0.15     14.79    4.74    4.69
09:50:01 PM     sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00
Average:        sda   10.01      0.00    295.66     29.53      0.15     15.36    4.57    4.58
Average:        sdb    0.00      0.00      0.00      0.00      0.00      0.00    0.00    0.00
[root@linux ~]#
```

(6)

(6) indicates that the sda response time (avg.) is 15.36ms.

"sar -p -d -s 21:00:00 -e 22:00:00 -f /var/log/sa/sa05" indicates that the response time from 21:00 to 22:00 on the 5th is output. "await" indicates the response time (ms).

Figure 3-17 Checking the response time

3.3.3 Storage System Environment

Use the tools provided with the storage system (such as ETERNUS SF Storage Cruiser) to check the IOPs and the response times before consolidating.

Obtain the volume-based performance information (IOPs and response times) from the tool output.

Present Time : 16:18:31
Interval Time(sec) : 30

Time	No	Name	Type	(1) I/O Count (IOPS) Read	Write	RD+WRT	Throughput (MB/s) Read	Write	RD+WRT	(2) Response Time (msec) Read	Write	RD+WRT
10:51:10	3	TPV_3	TPV	0	1	1	0	1	1	-	1	1
10:51:10	8	DBtest_0	TPV	24	37	60	5	3	8	2	1	1
10:51:10	9	DBtest_1	TPV	50	941	991	16	14	30	2	1	1
10:51:10	10	DBtest_2	TPV	100	4140	4240	41	48	89	2	1	1
10:51:10	11	DBtest_3	TPV	139	6462	6601	59	73	131	2	1	1
10:51:10	12	DBtest_4	TPV	553	8061	8613	121	85	206	1	1	1
:	:		:	:	:	:	:	:	:	:	:	:
11:09:10	3	TPV_3	TPV	5445	7312	12756	53	62	115	1	1	1
11:09:10	8	DBtest_0	TPV	1185	2401	3585	26	25	50	1	1	1
11:09:10	9	DBtest_1	TPV	2555	6200	8754	44	58	102	1	1	1
11:09:10	10	DBtest_2	TPV	4126	6753	10879	49	62	111	1	1	1
11:09:10	11	DBtest_3	TPV	1263	2285	3547	91	26	116	1	1	1
11:09:10	12	DBtest_4	TPV	6760	6850	13609	124	60	183	1	1	1
:	:		:	:	:	:	:	:	:	:	:	:
Average	3	TPV_3	TPV	2179	3793	5972	51	38	88	1	1	1
Average	8	DBtest_0	TPV	2130	3918	6047	51	39	90	1	1	1
Average	9	DBtest_1	TPV	2266	4493	6758	58	45	103	1	1	1
Average	10	DBtest_2	TPV	2270	4466	6736	58	45	103	1	1	1
Average	11	DBtest_3	TPV	2157	3910	6067	53	40	93	1	1	1
Average	12	DBtest_4	TPV	1810	3949	5758	57	49	105	1	1	1

The following information can be checked for each volume.
(1) IOPs
(2) response times

Figure 3-18 Checking the storage system performance (IOPs and response times of the pre-consolidated database)

3.4 Procedure for Checking Items in Post-Consolidated Environments

This section provides the method to calculate the post-consolidation CPU utilization, memory capacity, and storage system performance using the results in "3.3 Procedure for Checking Items in Pre-Consolidated Environments".

If the post-consolidation model is determined, check that each resource is sufficient. If the resources are insufficient, consider the procurement of additional resources. If the post-consolidation model is not determined, select the model that satisfies the requirements of each resource.

3.4.1 Calculation Method for the CPU Utilization

1. Consideration method for the load and time period

The post-consolidation CPU utilization must be estimated using the values at the time period when the sum of the CPU utilization in all the pre-consolidated environments is the highest.

Example: If the following two environments are consolidated, the values from 18:00 to 0:00 are used because the sum of the CPU utilization in environments A and B (15% + 30% = 45%) is the highest.

	CPU clock frequency	Number of CPU cores	CPU utilization 0:00 – 6:00	6:00 – 12:00	12:00 – 18:00	18:00 – 00:00
Environment A before consolidating	1.8GHz	8	5%	15%	25%	15%
Environment B before consolidating	2.0GHz	4	5%	5%	5%	30%

Table 3-2 Checking the peak time (CPU)

2. Calculating the coefficient

Calculate the coefficient using the CPU clock frequency, the number of CPU cores, and the CPU utilization during the period described in Section 1 for the pre-consolidated environments.

Coefficient = the sum of each pre-consolidated environment's [CPU clock frequency × number of CPU cores × CPU utilization (%)/100]

Example: If the following two environments are consolidated, the coefficient is 4.56.

	CPU clock frequency	Number of CPU cores	CPU utilization
Environment A before consolidating	1.8GHz	8	15%
Environment B before consolidating	2.0GHz	4	30%

Environment A before consolidation: 1.8 × 8 ×15/100 = 2.16 (A)

Environment B before consolidation: 2.0 × 4 × 30/100 = 2.4 (B)

Coefficient: 2.16 (A) + 2.4 (B) = 4.56

Table 3-3 Calculating the coefficient (CPU)

3. Guideline of the CPU utilization when the post-consolidation model is determined

The post-consolidation CPU utilization is calculated using the coefficient shown in Section 2, the post-consolidation CPU clock frequency (GHz), and the number of post-consolidation CPU cores as shown below.

CPU utilization = coefficient ÷ (post-consolidation CPU clock frequency × number of post-consolidation CPU cores) × 100

Example: The guideline of the CPU utilization is 23.75% when the coefficient is 4.56, the post-consolidation CPU clock frequency is 3.2GHz, and the number of post-consolidation CPU cores is 6.

The guideline of the CPU utilization: 4.56 ÷ (3.2 × 6) × 100 = 23.75

4. Guideline for selecting CPUs when the post-consolidation model is not determined

To calculate the resources required in the post-consolidated environment, use the coefficient shown in Section 2 and the target CPU utilization.

CPU clock frequency × number of CPU cores = coefficient ÷ target CPU utilization/100

Example: Select the model so that "CPU clock frequency × number of CPU cores" is 11.4 or more when the coefficient is 4.56 and the target CPU utilization is 40%.

The guideline of "CPU clock frequency × number of CPU cores": 4.56 ÷ 40/100 = 11.4

With 2.3GHz CPU clock frequency and 6 CPU cores, the result of 13.8 satisfies the requirement of 11.4 or more.

With 3.4GHz CPU clock frequency and 4 CPU cores, the result of 20.4 satisfies the requirement of 11.4 or more.

Therefore, the above combinations are possibilities.

3.4.2 Calculation Method for the Size of the Database Buffer Cache

The size of the post-consolidated database buffer cache can be calculated using the size of the pre-consolidated database buffer cache as shown below. For the safety coefficient, 20% is added. This is because processes may be delayed in the consolidated environment if the size of the database buffer cache is insufficient.

Required size of the database buffer cache ≥ (the sum of each pre-consolidated environment's size of the database buffer cache) × 1.2 (safety coefficient)

Example: If the following two environments are consolidated, the required database buffer cache size is 19.2GB or more.

	Size of the database buffer cache
Environment A before consolidating	4,092MB
Environment B before consolidating	12,276MB

Required size of the database buffer cache: (4,092 + 12,276) × 1.2 = 19,641.6 → capacity of 19.2GB or more

Table 3-4 Calculating the size of the database buffer cache

3.4.3 Calculation Method for the Physical Memory Capacity

1. Consideration method for the load and time period

The post-consolidation physical memory capacity must be estimated using the values at the time period when the sum of the available memory capacities in all the pre-consolidated environments is the lowest or the sum of the memory usage rates is the highest.

Check the pre-consolidation available memory capacity in Windows environments or the pre-consolidation memory usage rate in Linux environments.

Example: If the following two environments are consolidated in Windows environments, the values from 6:00 to 12:00 are used because the sum of the available memory capacities in environments A and B (48GB + 96GB = 144GB) is the lowest.

	Physical memory capacity	Available memory capacity			
		00:00	6:00 12:00	18:00	00:00
Environment A before consolidating	64GB	60.8GB	48GB	60.8GB	60.8GB
Environment B before consolidating	128GB	121.6GB	96GB	108.8GB	121.6GB

Table 3-5 Checking the peak time (physical memory capacity in Windows environments)

Example: If the following two environments are consolidated in Linux environments, the values from 6:00 to 12:00 are used because the sum of the memory usage rates in environments A and B (25% + 25% = 50%) is the highest.

	Physical memory capacity	Memory usage rate			
		00:00	6:00 12:00	18:00	00:00
Environment A before consolidating	64GB	5%	25%	5%	5%
Environment B before consolidating	128GB	5%	25%	15%	5%

Table 3-6 Checking the peak time (physical memory capacity in Linux environments)

2. Calculating the memory usage

In Windows environments, calculate the memory usage using the available memory capacity during the period described in Section 1 and the physical memory capacity before consolidating. In Linux environments, calculate it using the memory usage rate in the period described in Section 1 and the physical memory capacity before consolidating.

The guideline of the memory usage for Windows:

the sum of each pre-consolidated environment's (physical memory capacity - available memory capacity)

Example: 48GB = (64GB [A] - 48GB [A]) + (128GB [B] - 96GB [B])

The guideline of the memory usage for Linux:

the sum of each pre-consolidated environment's (physical memory capacity × memory usage rate/100)

Example: 48GB = (64GB [A] × 25% / 100 [A]) + (128GB [B] × 25% / 100 [B])

3. Guideline of the memory usage rate when the post-consolidation physical memory capacity is determined

Calculate the guideline of the post-consolidation memory usage rate using the guideline of the post-consolidation memory usage described in Section 2 and the post-consolidation physical memory capacity.

Memory usage rate = (the guideline of the post-consolidation memory usage ÷ post-consolidation physical memory capacity) × 100

Example: When the guideline of the post-consolidation memory usage is 48GB and the post-consolidation physical memory capacity is 256GB, the guideline of the memory usage rate is 18.75%.

The guideline of the memory usage rate: (48 ÷ 256) × 100 = 18.75

4. Guideline for obtaining the required physical memory capacity when the post-consolidation physical memory capacity is not determined

Calculate the guideline of the post-consolidation physical memory capacity using the post-consolidation memory usage described in Section 2 and the target memory usage rate.

Post-consolidation physical memory capacity = the guideline of the post-consolidation memory usage ÷ the target memory usage rate/100

Example: When the guideline of the post-consolidation memory usage is 48GB and the target memory usage rate is 20%, the physical memory capacity must be 240GB or more.

The guideline of physical memory capacity: 48GB ÷ 20/100 = 240GB

3.4.4 Calculation Method for the Storage System Performance

1. Checking the performance of the post-consolidation storage system

Check the IOPs and the response time from the relative catalogs.

Example: For the ETERNUS AF250, the maximum IOPs is 400,000 IOPs and the response time is average 0.5ms.
For the ETERNUS AF650, the maximum IOPs is 640,000 IOPs and the response time is average 0.5ms.

2. Consideration method for the load and time period

The performance of the post-consolidation storage system must be estimated using the values at the time period when the sum of the IOPs in all the pre-consolidated environments is the highest.

Example: If the following three environments are consolidated, the values from 18:00 to 0:00 are used because the sum of the IOPs in environments A and B (2000 + 3200 + 2400 = 7600) is the highest.

	IOPs 0:00	6:00	12:00	18:00	00:00
Volume 1 in environment A before consolidating	400	400	1200	2000	
Volume 1 in environment B before consolidating	2400	400	200	3200	
Volume 2 in environment B before consolidating	400	1400	3200	2400	

Table 3-7 Checking the peak time (IOPs)

3. Calculating post-consolidation IOPs

Calculate the post-consolidation IOPs using the pre-consolidation IOPs during the period described in Section 2.

Guideline of post-consolidation IOPs = the sum of each pre-consolidated environment's IOPs

Example: 7600 = 2000 (Volume 1 in A) + 3200 (Volume 1 in B) + 2400 (Volume 2 in B)

4. Determining the availability of consolidation

Determine the availability of consolidation using the post-consolidation storage system performance described in Section 1 and the post-consolidation IOPs described in Section 3.

If the storage system performance value is greater than or equal to the post-consolidation IOPs, consolidation is possible.

Example: When the post-consolidation storage system performance is 400,000 and the post-consolidation IOPs is 7,600, consolidation is possible because the storage system performance is much higher than the post-consolidation IOPs.

3.4.5 Calculation Method for Storage System Disk Capacities

Check if the sum of the disk usage for the pre-consolidated database can be secured for the storage system in the post-consolidated environment.

Required disk capacity ≥ the sum of the pre-consolidated environment's [disk usage for the database]

Example: If the following two environments are consolidated, a disk capacity of 460GB or more is required.

	Disk usage for the database
Environment A before consolidating	100GB
Environment B before consolidating	360GB

Disk capacity: (100 + 360) = 460 → capacity of 460GB or more

Table 3-8 Calculating the disk capacity

* For systems where the database size is expected to increase, consider additional increases.

4 Verifying Database Consolidation

This chapter verifies database consolidation based on "3 Checks for Consolidating Databases".

The following shows the verification environment.

- Pre-consolidation: Database data is stored in the internal disks (HDDs) of each server and each server processes SQL statements.
- Post-consolidation: Database data is consolidated to an ETERNUS AF650 and servers are also consolidated to a single PRIMERGY server that processes the SQL statements simultaneously.

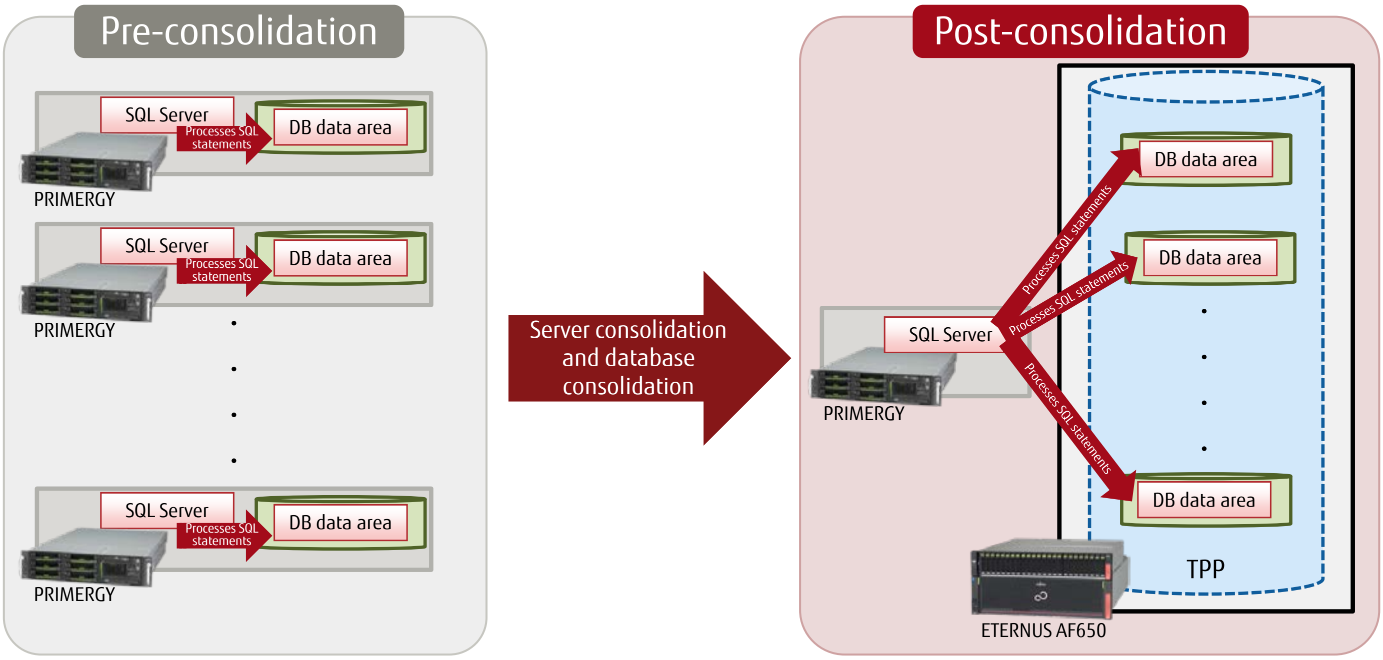

Figure 4-1 Consolidation overview

The following shows the details of the verification environment.

Pre-consolidation configuration (common to each server)

Server	: PRIMERGY (CPU: [1.8GHz, 4 cores] × 2)
OS	: Microsoft Windows Server 2012 R2 Standard Edition (Desktop Experience)
Database	: Microsoft SQL Server 2014
Disk device	: Internal HDDs
Database volume	: Single volume defined in a disk space (RAID0(2) × 1)

Post-consolidation configuration

Server	: PRIMERGY (CPU: [2.3GHz, 6 cores] × 1)
OS	: Microsoft Windows Server 2016 Standard Edition (Desktop Experience)
Database	: Microsoft SQL Server 2016
Disk device	: ETERNUS AF650 SSDs
Database volumes	: Same number of TPVs as environments to be consolidated are defined in a TPP (RAID5(3+1) × 1)

A PRIMERGY server meeting the specification requirements calculated in "3.4 Procedure for Checking Items in Post-Consolidated Environments" is used in the post-consolidation configuration to satisfy the higher processing power for consolidation. For the storage, the ETERNUS AF650 replaces the internal HDDs.

4.1 Server Consolidation (CPU Utilization)

1. Calculating the pre-consolidation coefficient

Perform the procedure described in "3.3.1 Windows Environments" during the peak time of the database. Then, calculate the coefficient by using the CPU clock frequency, the number of CPU cores, and the CPU utilization (%) that are obtained in the procedure.

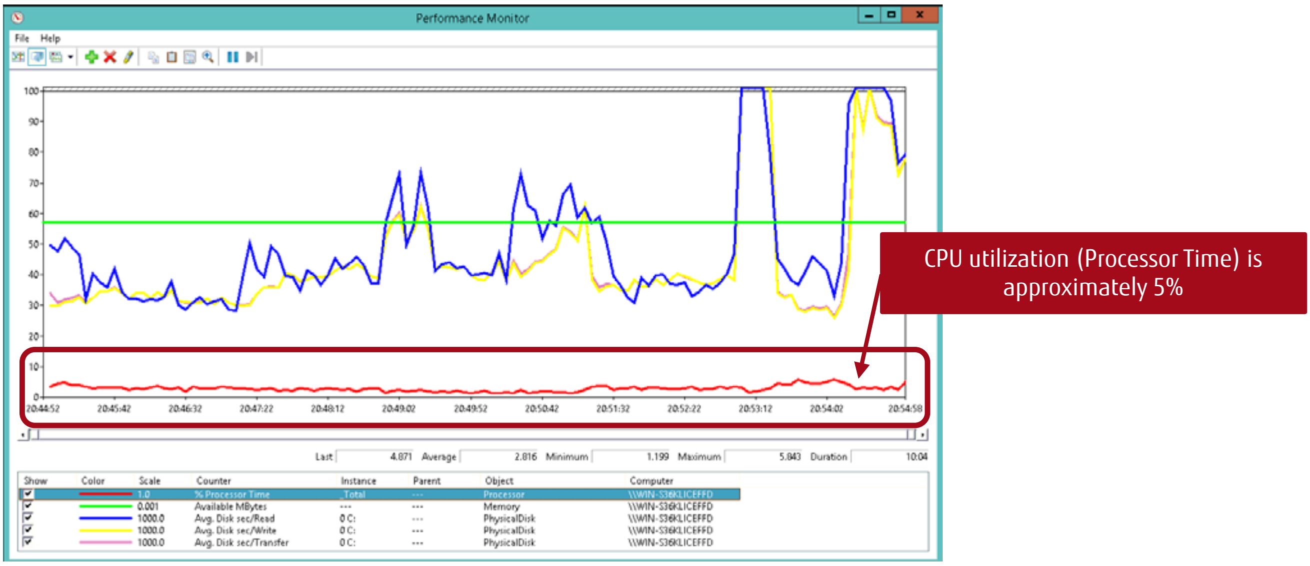

Figure 4-2 CPU utilization before consolidating

The CPU utilization was approximately 5% in the pre-consolidated environment for this test.

By using the CPU clock frequency (1.8GHz) and the number of cores (8) in the pre-consolidated environment, a coefficient of 0.72 is calculated as follows:

Coefficient for a single environment before consolidating $1.8 \times 8 \times 0.05 = 0.72$

To consolidate three of the above pre-consolidated environment and ten of the above pre-consolidated environment, the coefficients are calculated as follows:

Coefficient = the sum of each pre-consolidated environment's [CPU clock frequency × number of CPU cores × CPU utilization (%)]

Coefficient for consolidating three environments $0.72 \times 3 = 2.16$

Coefficient for consolidating ten environments $0.72 \times 10 = 7.2$

2. Calculating and verifying the post-consolidation CPU utilization

   Estimate the post-consolidation CPU utilization as follows by using the coefficient calculated in Section 1, the post-consolidation CPU clock frequency (GHz), and the number of post-consolidation CPU cores.

   CPU utilization = coefficient ÷ (post-consolidation CPU clock frequency × number of post-consolidation CPU cores)

   For a three environment consolidation, the CPU utilization is calculated using the CPU clock frequency (2.3GHz) and the number of cores (6) as shown below.

   2.16 ÷ (2.3 × 6) = 0.157 The estimated CPU utilization is approximately 15.7%.

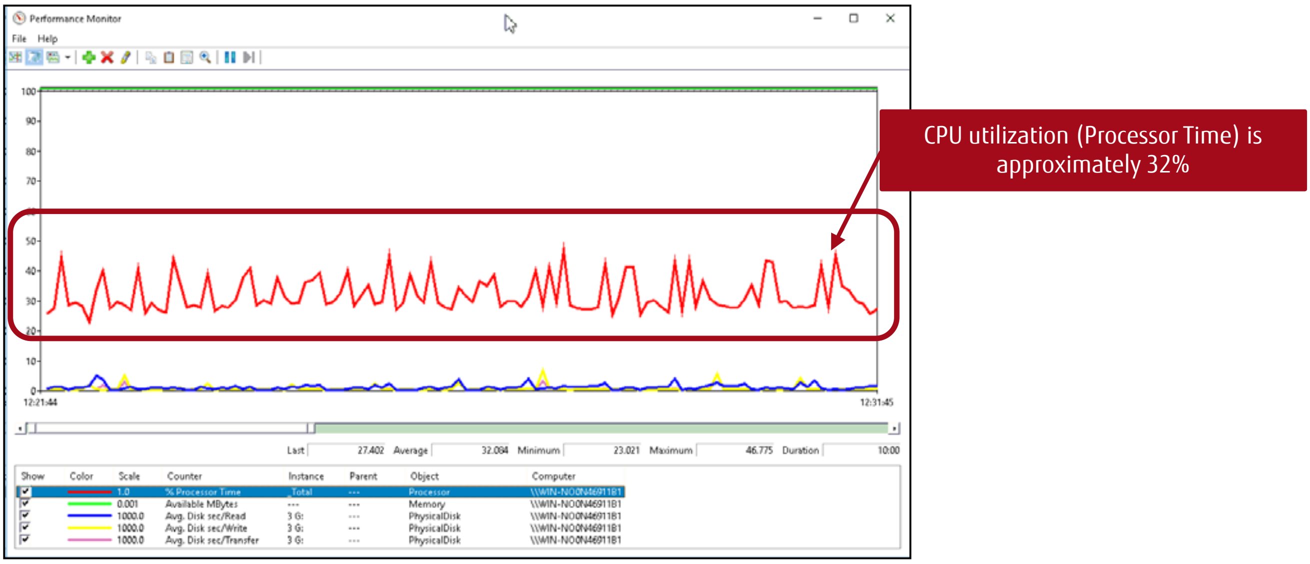

Figure 4-3 CPU utilization after consolidating three environments (results)

   The estimated CPU utilization was 15.7% but the test results showed that the CPU utilization was approximately 32%.

   For a ten environment consolidation, the CPU utilization is calculated using the CPU clock frequency (2.3GHz) and the number of cores (6) as shown below.

   7.2 ÷ (2.3 × 6) = 0.522 The estimated CPU utilization is approximately 52.2%.

Figure 4-4 CPU utilization after consolidating ten environments (results)

   The estimated CPU utilization was 52.2% but the test results showed that the CPU utilization was approximately 60%.

The test results show that even though there is a gap between the estimated value and the test results when three environments are consolidated, the impact on the system is small because the CPU utilization is approximately 30%. However, when ten environments are consolidated, the results are roughly within the range of the estimated value.

4.2 Sever Consolidation (Memory Capacity)

1. Obtaining the memory capacity in the pre-consolidated environment

Perform the procedure described in "3.3.1 Windows Environments" during the peak time of the database to obtain the database buffer cache size, the physical memory capacity, and the available memory capacity in the pre-consolidated environment.

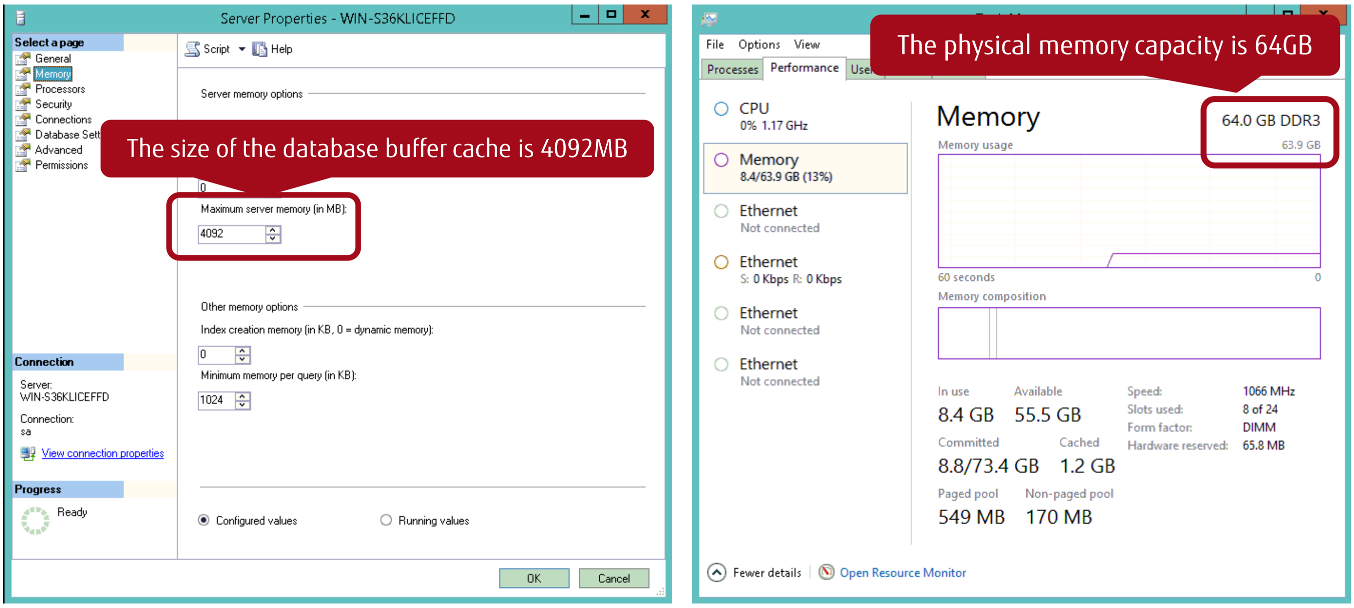

Figure 4-5 Database buffer cache size and physical memory capacity before consolidating

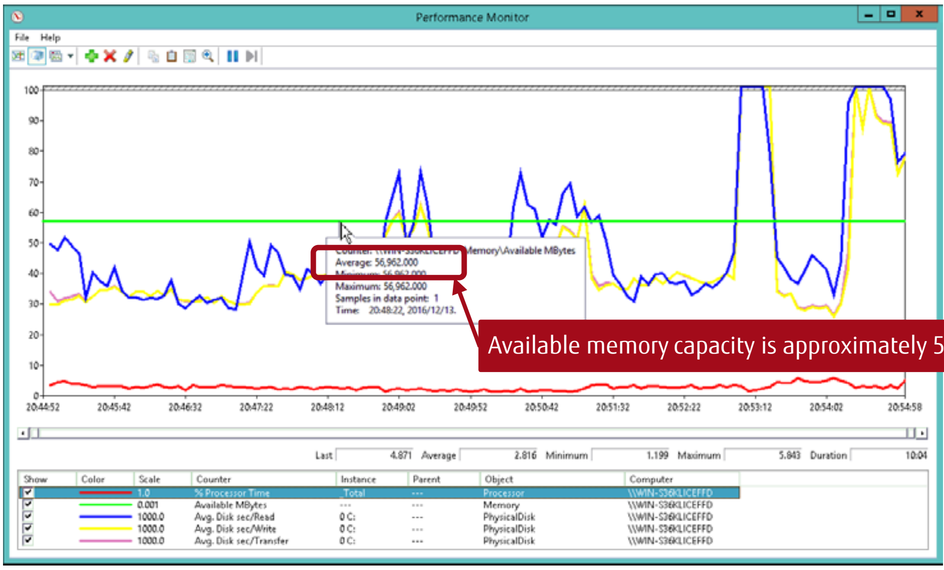

Figure 4-6 Available memory capacity before consolidating

Database buffer cache size before consolidating = 4GB (4092MB)
Physical memory capacity = 64GB
Available memory capacity = 57GB

2. Calculating and verifying the post-consolidation memory capacity

Calculate the post-consolidation database buffer cache size and the post-consideration physical memory capacity by referring to "3.4.2 Calculation Method for the Size of the Database Buffer Cache" and "3.4.3 Calculation Method for the Physical Memory Capacity".

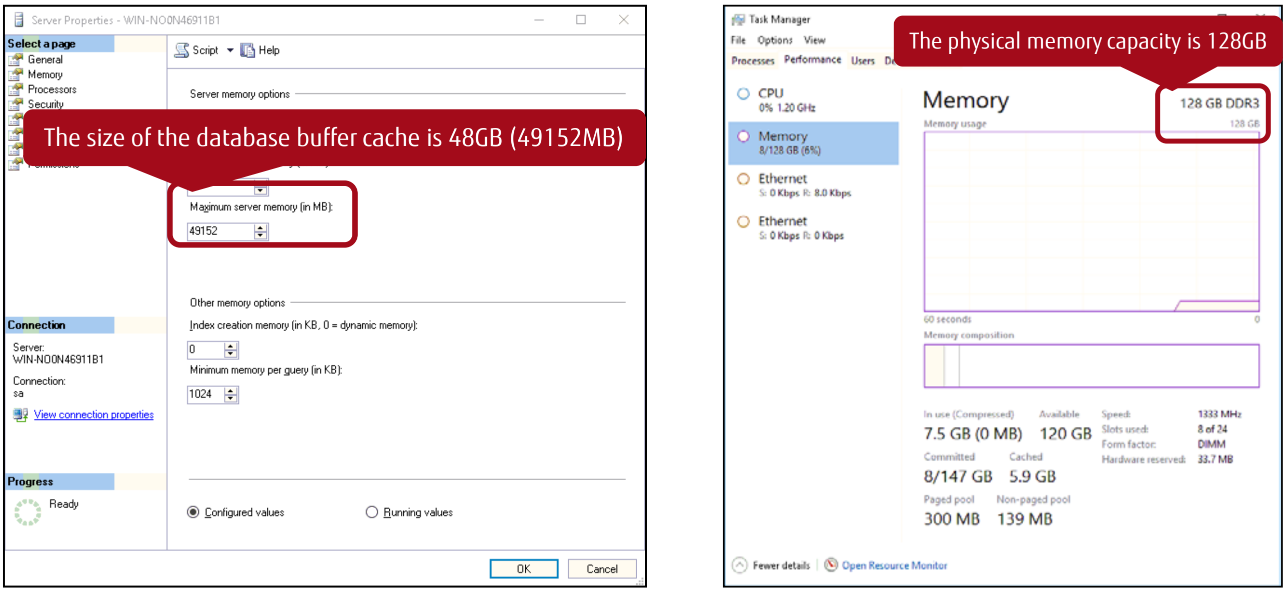

Figure 4-7 Database buffer cache size and physical memory capacity after consolidating

The guideline of the database buffer cache size ≥ (the sum of each pre-consolidated environment's size of the database buffer cache) × 1.2 (safety coefficient)

For consolidating three environments 4GB × 3 × 1.2 = 14.4GB

For consolidating ten environments 4GB × 10 × 1.2 = 48GB

The test was conducted with the assumption that the size of the database buffer cache was 48GB after consolidating.

The guideline of the physical memory usage = the sum of each pre-consolidated environment's (physical memory capacity - available memory capacity)

For consolidating three environments (64GB - 57GB) × 3 = 21GB

For consolidating ten environments (64GB - 57GB) × 10 = 70GB

The test was conducted with the assumption that the physical memory usage was 128GB after consolidating.

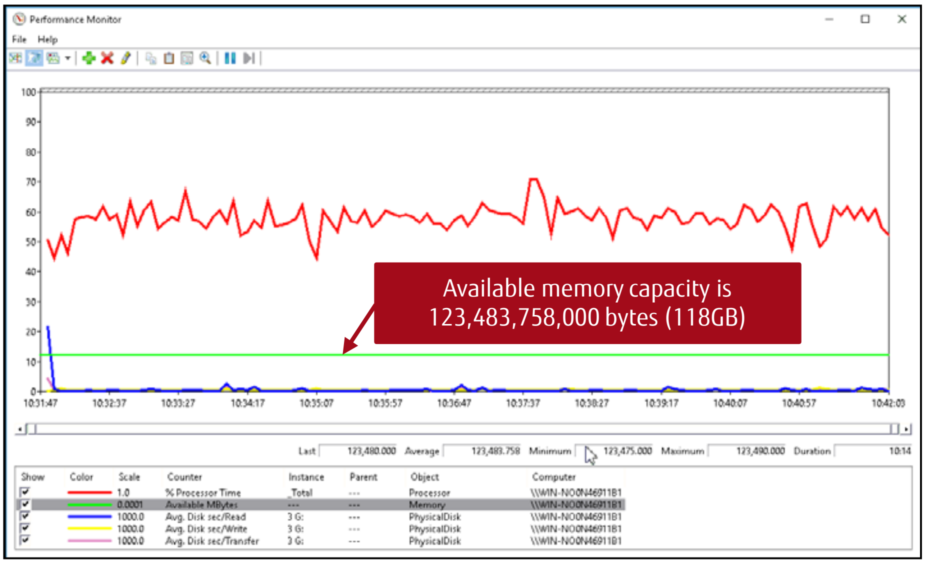

Figure 4-8 Available memory capacity after consolidating ten environments (results)

The test results are within the range of the estimated values.

4.3 Storage Consolidation

1. Comparing disk response times

The disk response time is 40ms in the pre-consolidated database where RAID0(2) is configured with internal disks and Microsoft SQL statements are processed on Windows Server 2012. The disk response time in actual databases rarely exceeds 20ms but a higher load was applied in this test.

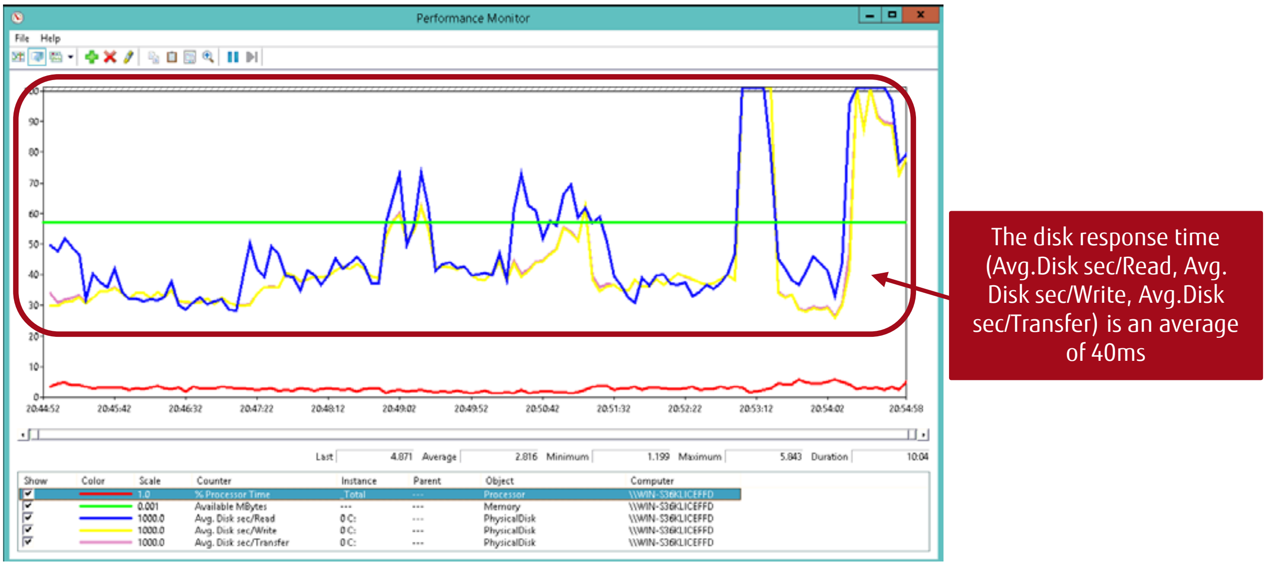

Figure 4-9 Disk response time before consolidating

Ave. Disk sec/Read, Ave. Disk sec/Write, and Ave. Disk sec/Transfer show varying values but fluctuate around an average of 40ms.

After consolidating, RAID5(3+1) is configured with the SSDs in the ETERNUS AF650, which is an external storage.

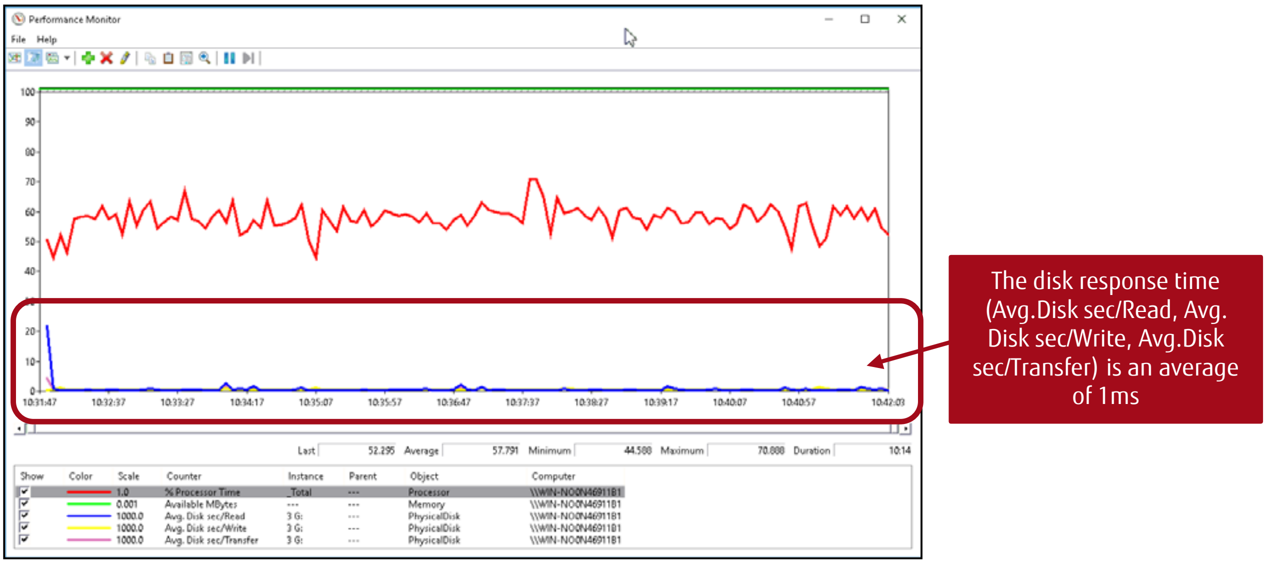

Figure 4-10 Disk response time after consolidating ten environments (results)

The test was conducted with the assumption that the ten-environment consolidation processed the SQL statements simultaneously for each of the ten databases that had the same pre-consolidated database configuration. The disk response time was approximately 1ms even when ten environments were consolidated.

	Disk	RAID	Number of consolidation environments	Disk response time
Pre-consolidated environment	Internal	RAID0(2)	1	Approximately 40ms
Post-consolidated environment	ETERNUS AF650	RAID5(3+1)	10	Approximately 1ms

Table 4-1 Comparison of disk response times

2. Comparing SQL statement processing times

Approximately 10 minutes were required to process the SQL statements in an HDD based database. However, by switching the HDDs with the ETERNUS AF650, the time required to process the SQL statements in a single execution was significantly reduced.

	Pre-consolidation	Post-consolidation	
	Single environment (Unconsolidated)	Three environments	Ten environments
SQL statement processing time	Approximately 10 minutes	Approximately 2 minutes	Approximately 5 minutes
Proportion of the post-consolidation processing time against the pre-consolidation processing time (per process)	1	1/15	1/20

Table 4-2 Comparison of SQL statement processing times

Compared to the pre-consolidation time, the processing time was reduced to approximately 1/15 for the three environment consolidation and 1/20 for the 10 environment consolidation per process.

The disk usage was reduced and the CPU utilization was improved by using the ETERNUS AF650 to process SQL statements that required high disk usage.

As long as server resources are sufficient, consolidating databases to an all-flash array is highly beneficial. Conventional databases have storage in database servers and when multiple databases access the storage, they have to compete for storage resources. However, with all-flash arrays, even if they are accessed from multiple database servers, storage resources are still sufficient.

3. Post-consolidation disk capacity

Ten 500GB volumes were created in the ETERNUS AF650 and a database (with a disk usage capacity of around 350GB) was stored to each of the volumes.

The ETERNUS AF series can consolidate a large number of databases. Consider the ETERNUS AF250 or the ETERNUS AF650 depending on the number of installable disks necessary.

5 Conclusion

Maintaining server resources is critical for database consolidation. The relationship between the CPU clock frequency and number of cores, the size of the database buffer cache, the installation memory size, and the disk capacity can be easily estimated by referring to this document.

The ETERNUS AF series allows SSDs to deliver their full performance potential by simultaneously handling the areas that use and the areas that do not use the Deduplication/Compression function. It can dramatically accelerate the processing speed of mission-critical systems that include databases by eliminating disk bottlenecks.

Instead of redesigning and tuning a database for a database replacement, the highest performance can be easily obtained by installing the ETERNUS AF series and simply maintaining the server performance by securing server resources.

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