FUJITSU ETERNUS AF/DX
Optimization Features
White Paper

Automated storage tiering and automated quality of service

Contents

Management summary and general remarks

Automated storage tiering (AST) refers to the ability of the storage array to move chunks of data between different disk types and RAID levels to meet the right balance between performance and space usage thus avoiding so-called hot spots. Frequently accessed data can be moved to high speed drives such as SSDs and less frequently accessed data to cost-effective disks with large capacities.

Quality of service automation (Automated QoS) ensures that particular applications always get a certain, predefined performance level. Adjusting the bandwidth and performing an automatic tuning of the I/O performance makes sure the required response time per application will be achieved.

Combining both optimizing features helps administrators balancing between performance, capacity and cost and to overcome peak loads with just some mouse clicks.

This white paper elaborates how automated QoS and AST are implemented in Fujitsu's ETERNUS storage systems. The concepts are explained in general and enriched with best practices.

Introduction

The amount of data to be retained and managed is rapidly increasing, even though much of the data is rarely or never accessed again. Proliferating capacity needs go hand in hand with higher service level requirements, while enterprise IT budgets are shrinking. Two basic solutions are thinkable: The first one is about moving rarely accessed data to lower cost tiers built from low-cost slowly spinning disk drives and place the data which is needed by mission and business-critical applications with highest service level requirements on the fastest storage media available. The second approach looks after application priorities. By prioritizing data access and dynamically managing any I/O conflict, high performance can be guaranteed for high-priority applications. At the same time capacity is used more efficiently, thus increasing storage utilization without sacrificing performance.

So far so good – but these valid approaches have some pitfalls. Data must be qualified, access frequency and service levels like response times or batch runtime must be measured and evaluated to decide which data has to be stored at a given time in a certain tier or which application needs to change its priority. These facts have been the main drivers for implementing automated storage tiering and automated quality of service concepts in external storage arrays.

Rarely accessed data does not need to be stored on expensive high performance disk drives but should be moved to a lower-cost tier consisting of less expensive disk drives. Without automation moving this data is an expensive and time-consuming task. Administrators must collect and analyze access data to decide which data may be moved to a lower-cost tier, doing this several times a week or a day depending on the current or predicted application workload.

The automated storage tiering function is defined by policies and allows changing data locations dynamically corresponding to the performance status of the data.

An array based quality of service option just limits the IOPS for specific volumes in a static way and requires a lot of expertise and continuous tuning to find the optimum settings. To ease these tasks automated quality of service management (automated QoS) lets administrators set priorities based on performance requirements much more easily and dynamically adjusts the I/O bandwidth along with the result of continuous performance monitoring.

This feature makes it easier for the user to assign I/O priorities. Furthermore, the automatic tuning ensures that the targets are more accurately achieved, resulting in better service level fulfillment.

For both options administrators are supported in tasks of performance estimation, layout design and relocation of data according to performance and cost needs.

All of the above prerequisites and trade-offs have been taken into consideration when implementing the AST and automated QoS functionality into ETERNUS storage systems. Following the family concept of offering uniform management and same functionality for all members of the family, the features are available from entry system to high-end models.

Figure 1 shows the environment for the ETERNUS DX optimization options, which include automated storage tiering and automated quality of service.

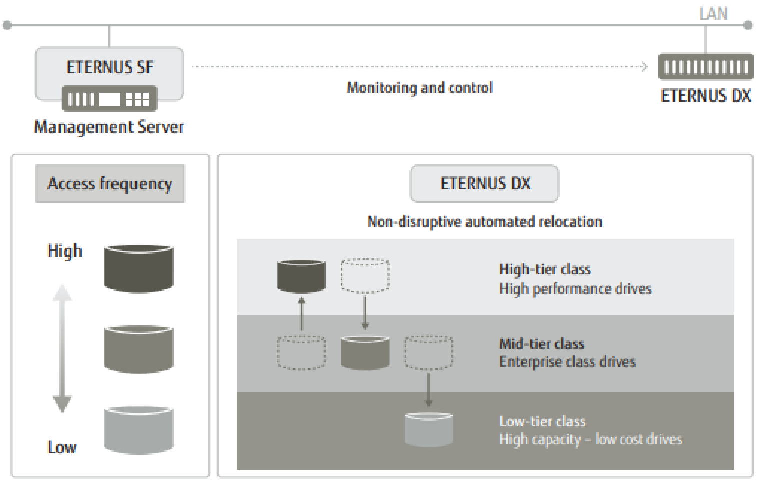
LAN
ETERNUS SF
Management Server
Monitoring and control
ETERNUS DX
Access frequency
High
Low
ETERNUS DX
Non-disruptive automated relocation
High-tier class
High performance drives
Mid-tier class
Enterprise class drives
Low-tier class
High capacity – low cost drives

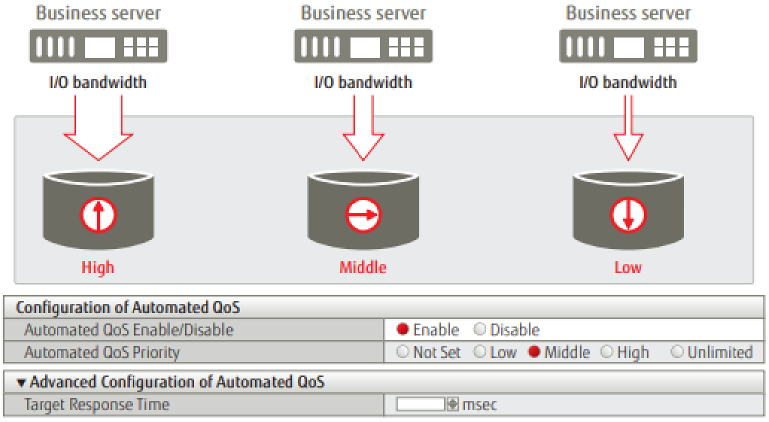
Business server
I/O bandwidth
Business server
I/O bandwidth
Business server
I/O bandwidth
High
Middle
Low
Configuration of Automated QoS
Automated QoS Enable/Disable
Enable
Disable
Automated QoS Priority
Not Set
Low
Middle
High
Unlimited
Advanced Configuration of Automated QoS
Target Response Time
msec

Figure 1

AST basic definitions

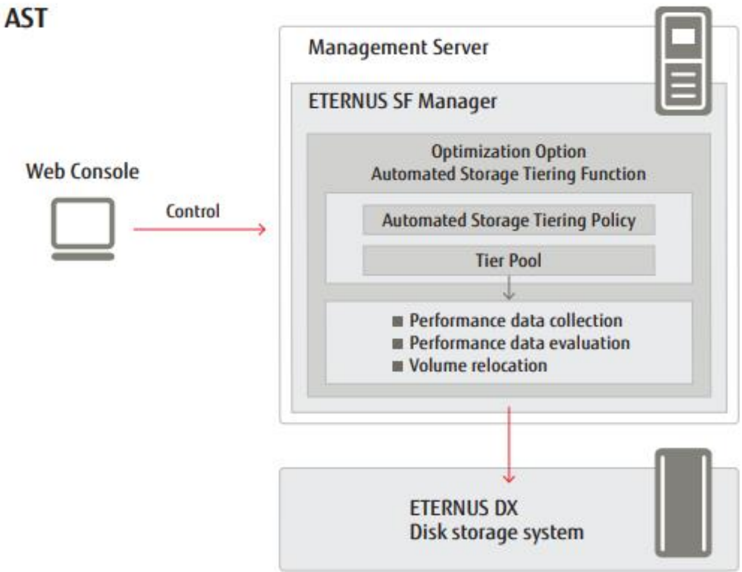

Figure 2

Prerequisites and licenses

Required software

The automated storage tiering feature is controlled by the ETERNUS SF storage management software suite, which is delivered with any ETERNUS DX storage array. ETERNUS SF can be installed on either a Windows, RHEL or Solaris host as well as on a virtual machine provided by VMware or HyperV.

Required software

The function is enabled by an optional license called ETERNUS SF storage cruiser optimization option in addition to the ETERNUS SF storage cruiser standard license; it cannot be setup with the ETERNUS SF storage cruiser basic license or the free-of-charge ETERNUS SF express. These licenses must be activated for each ETERNUS DX system regardless of installed disk capacity.

In addition, the hardware-based thin provisioning license must be registered on the storage array itself.

AST is not relevant for ETERNUS all-flash storage systems, because only one flash-tier is available.

The automated storage tiering implemented in ETERNUS DX distinguishes three types of so-called tiering objects which are defined as follows:

Automated tiering policy – defines when, how and under which conditions the relocation of data is executed. The tiering policy is the central part of the automated storage tiering functionality. The baseline for relocation is the IOPS values measured on the sub-LUNs, either as peak values or average values within an evaluation interval.

Flexible tier pool – a flexible tier pool consists of two or three tiering sub-pools, which are storage areas of thin provisioned RAID groups. In case three sub-pools are chosen, these reflect the low, middle and high tiers with regard to performance or cost per GB. The flexible tier pool is bound to one dedicated tier policy – when choosing a 2-tier policy the middle sub-pool will be omitted.

Flexible tier volume – flexible tier volumes are volumes which are created in a flexible tier pool and are the entities which are presented to the hosts like any other volume via the common mechanisms of mapping and defining LUN affinities

Setup of tiering objects

Tiering objects are a group consisting of tiering policies, tier pools and tier volumes which all must be properly configured to enable the AST feature.

Tiering policies

The first step of implementing automated storage tiering is the setup of at least one tiering policy, which defines when and how data relocation is triggered. The system constantly measures how many IOPS are executed on the sub-LUNs of the flexible tier volume. The measurement method can either be related to the peak value or to the average value within an evaluation interval.

The evaluation interval can be set either on an hourly or a daily base. Hourly measurement spans 4, 6, 8 or 12 hours, after which the evaluation process starts over again. The daily based measurements span from 1 day to 31 days with increments of 1.

The tiering policy also defines the threshold values for triggering the relocation of data from one sub-pool to another and allows timely limitation of the evaluation and relocation process itself.

If the interval is set to an hourly base, the relocation process starts immediately after the completion of measurement and analysis.

In case of daily based measurement the administrator can define a measurement period within a day to limit for example measurement to business hours. The daily policy also allows to define the starting time of the relocation process to execute measurement in periods of low system activity.

Flexible tier pools

The next step is to set up the flexible tier pools that are used by the policy. A flexible tier pool consists of two or three tiering sub-pools which are storage areas of thin provisioned RAID groups.

The three sub-pools are nominated as low, middle and high with regard to the performance or cost of the chosen disk types or RAID levels. Classically, in a three-tier environment the high sub-pool is created from fast SSDs, the middle sub-pool is created from SAS disks, while the low sub-pool consists of slower high capacity nearline SAS disks.

The creation of sub-pools is much more flexible with Fujitsu's implementation of Automated Storage Tiering. It is also possible to create a two-tier pool by omitting the middle sub-pool and it is also possible to not only map different physical disk types to different sub-pools, but e.g. also the same disk types with different RAID configurations. Thus, for example, the higher sub-pool can be created out of a RAID1 group of 15k rpm SAS disks while the lower sub-pool is made of a RAID5 group of 15k or 10k rpm SAS disks.

Flexible tier volumes

Flexible tier volumes are generated within a tiering pool. They are the entities which are presented to the hosts via the common mechanisms of mapping and defining affinities. Flexible tier volumes are thin provisioned volumes which consist of sub-LUNs (chunks) with a size of 252 MB. These are the smallest entities which are moved between the sub-pools of the tier pool

The process of creating the flexible tier volumes allows assigning the tier sub-pool for the initial location of sub-LUNs before any performance monitoring and analysis has started.

Tiering process

The tiering process itself is done in a three-step cycle:

- **Step 1:** IOPS numbers of the sub-LUNs are monitored and collected for the duration of the evaluation interval defined in the tiering policy. The measurement method itself can be either peak IOPS or average IOPS
- **Step 2:** At the end of the collection period the IOPS numbers are analyzed and compared with the thresholds defined in the tiering policy.
- **Step3:** If IOPS numbers of a sub-LUN exceed the threshold, the sub-LUN is relocated to a higher tier; if IOPS reach the lower threshold, it is moved to a lower tier.

The tiering mode defines the execution mode for the tiering process. Beside automatic, it can be performed: semi- automatic as well as manual. These options define which of the three steps are executed automatically or have to be triggered by the administrator.

In automatic mode all three steps are executed automatically, monitoring and collecting data, evaluating the data and relocating the data if appropriate.

In contrast, relocation is not performed automatically in semi-automatic mode and thus this third step has to be triggered by the administrator. So a human decision is responsible for starting or skipping the relocation.

The manual mode only does step one automatically, which is the performance data collection. The administrator can confirm the evaluation results and can manually select the appropriate volumes for relocation as needed.

The parameter [Expire Time] in the policy setup defines the maximum duration for the evaluation and relocation phase. It can span from 1 to 24 hours. The relocation will be forcibly stopped, whether finished or not, when the expire time is up. The value can also be set to 0, giving unlimited execution time for the evaluation and relocation phase. In this case, if the data evaluation and relocation process takes longer than the next collection interval, the next data evaluation and relocation phase will be skipped

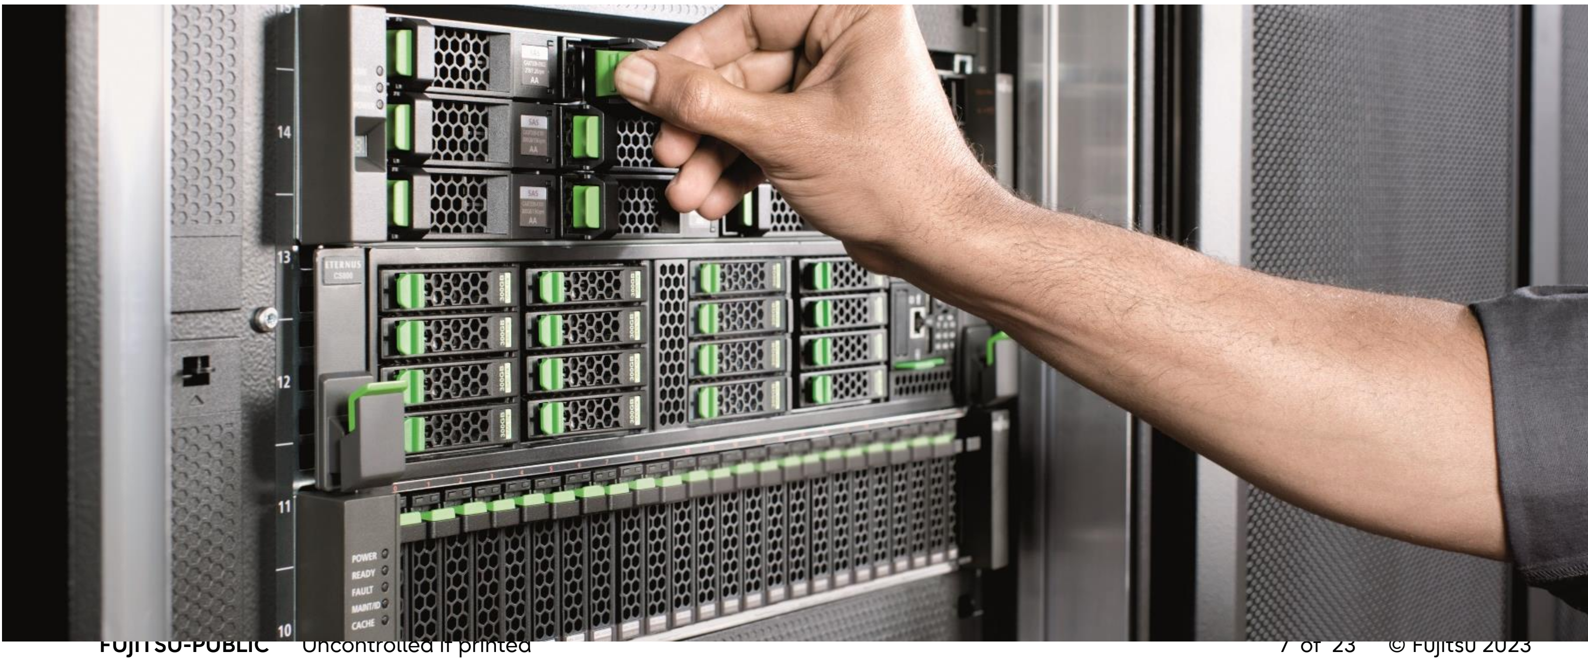

Example for daily based measurement and execution of automated tiering

Figure 3 shows a screenshot of setting up a tier policy for automated execution on a daily base.

[Execute Mode] is set to "Auto" and the measurement method [Evaluation Method] is based on average IOPS [Type] for a duration of three days [Interval] in which the evaluation time starts at 7:00 h and ends at 18:00 h each day [Evaluated Timeframe]. Relocation is triggered to start at 3:00 h [Execute Time] and will be executed for a maximum time of 2 hours [Expire Time].

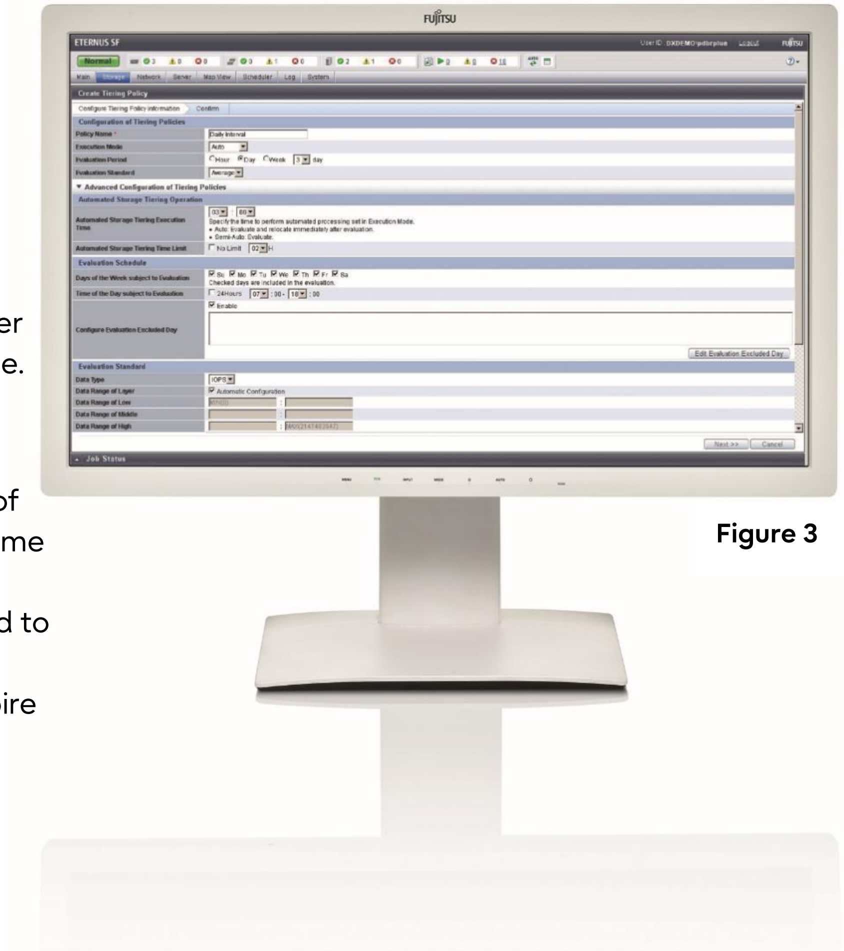

Figure 3

So let's see in Figure 4 what happens exactly when this policy is executed. Whenever the administrator starts the AST policy for the first time, the data collection will be started the same or next day at 7:00 h. After the full interval of three days, collecting data every day between, 7:00 h and 18:00 h data evaluation is completed on the data collected during the interval of these three days. And if the thresholds are matched, the relocation is started at 3:00 h the next day and will be executed for a maximum of 2 hours, while the next 3-day evaluation period starts over again at 7:00 h.

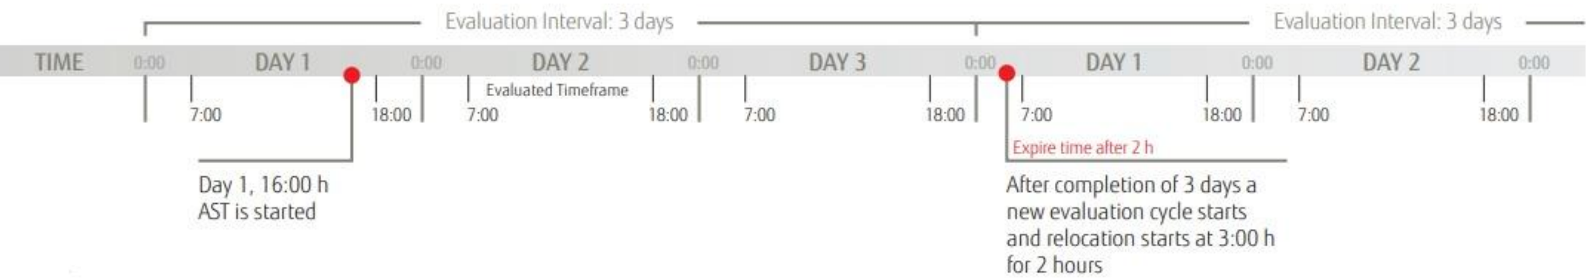

Figure 4

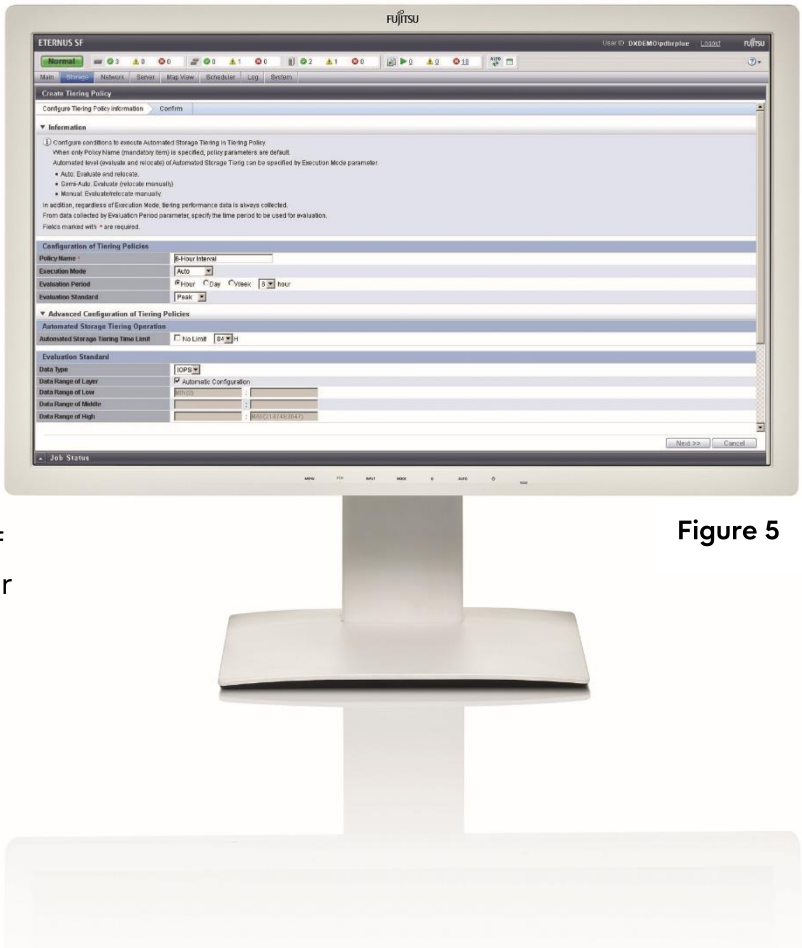

Figure 5

Example for hourly based measurement and execution of automated tiering

Figure 5 shows a screenshot of setting up a tier policy for automated execution on an hourly base.

[Execute Mode] is again set to "Auto" and the measurement method [Evaluation Method] is now based on peak IOPS [Type] for a duration of 6 hours [Interval]. Relocation will be executed for a maximum time of 4 hours [Expire Time].

Administrators must keep in mind that if shorter intervals are chosen in terms of hours instead of days, the ending of a collection interval is the starting point for the evaluation and relocation interval as well as the starting point for the next data collection interval. This is a continuous process and thus it is not required to manually set the [Evaluated Timeframe] and the [Execution Time].

Figure 6 shows what happens exactly when this policy is executed. If the AST process is started by the administrator on the current day at 15:00 h, the first interval has been started in between the regular starting times. Then at 18:00 h the evaluation and relocation phase is started for the data collected before, but only for the first 3 hours from 15:00 h to 18:00 h. Possible data relocation for this interval is skipped, because the available data is for less than 75 percent of the desired interval length (3 hours of 6 hours).

At the same time, at 18:00 h, the next data collection starts as well, for which the data evaluation and relocation would start 6 hours later, at 0:00 h the next day. In our example the choice of an interval size of 6 hours divides a day into 4 sections with starting times at 0:00 h, 6:00 h, 12:00 h and 18:00 h

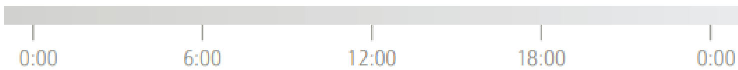

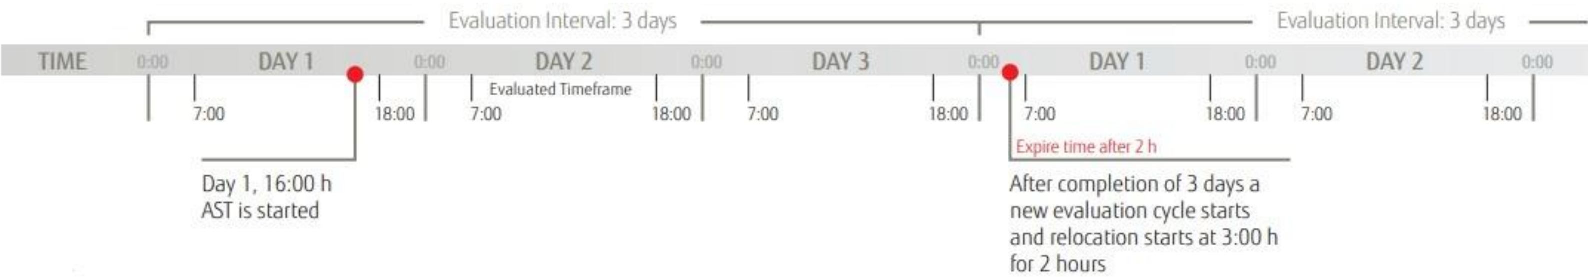

Figure 6

Flexible tier volumes, Sub-LUNs and IOPS

To understand the mechanisms of Automated Storage Tiering it is important to know that the IOPS thresholds are separately defined for each of the sub-LUNs or chunks. These IOPS numbers may be very much different from the traditional understanding of IOPS numbers, which count the total IOPS between a host and an array or between an HBA and a volume.

This is because the IO traffic between a server and a tier volume can be very distributed over the sub-LUNs in many ways, while often changing in unpredictable patterns.

The measurements and relocation are done on a sub-LUN basis. An ETERNUS DX S3 sub-LUN is 252 MB. The IOPS rate has to be seen on one chunk and not in sum for the whole LUN. Thus the IOPS on all sub-LUNs may be below the limit for relocation to the high tier; even if the IOPS in total are higher than the threshold.

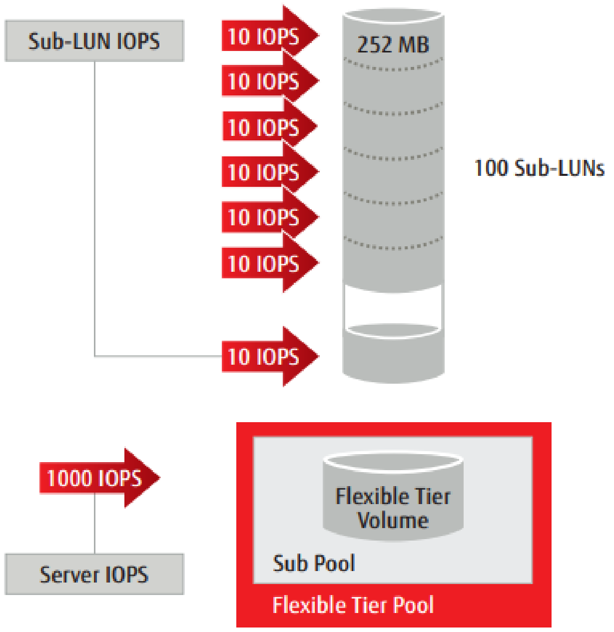

Figure 7: Flexible tier volume with balanced IOPS on Sub-LUNs

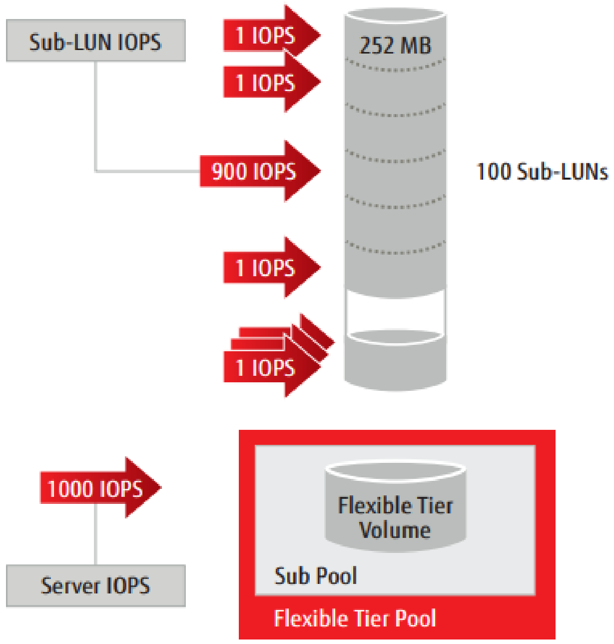

Figure 8: Flexible tier volume with unbalanced Sub-LUN IOPS

In Figure 7 we see an ideal picture where 1,000 server IOPS are equally processed in a flexible tier volume consisting of 100 sub-LUNs. The example given earlier in the setup of the tier policy in this case shows that all sub-LUNs will stay in the low tier, as the peak value does not exceed the lower threshold.

In contrast, Figure 8 shows that the majority of the 1,000 server IOPS is processed to one sub-LUN only. After the collection evaluation period this sub-LUN will be the only one to be migrated to the high tier, as its value by far exceeds the high treshold value.

Best practices at high level

Before configuring automated tiering configurations administrators have to clearly analyze and understand the data and application structure. This document provides some high level hints. For more detailed design and best practices a variety of documentation is available. Here are the basic points administrators should consider when building an automated tiering infrastructure:

Analysis of the data structure

Categorize data

Applications may have different types of data with different performance requirements. Planned data LUNs have to be categorized accordingly. It is crucial to have a clear understanding of the IOPS requirements for data and applications.

Define rough performance needs

Categorize data to possible tier classes, which can then be dedicated to possible storage areas.

Group data in tier pools

Data, which may fit together in performance requirements over time, can be put into one tier pool. Decide which data can be in the same tier pool. Data of different applications with the same SLA requirements can be grouped in one tier pool or be divided in different tier pools.

Define the sub-pools of your tier pool

Following a performance estimation of the data some LUNs will certainly need high speed storage areas as SSDs, some others are seen in the medium performance range and a third part with slow performance requirements. Decide if a specific tier pool is to consist of two or of three sub-pools and keep in mind that two sub-pools can be upgraded to three sub-pools, but sub-pools cannot be removed.

Define these LUNs as flexible tier volumes and initially place them in the appropriate sub-pool, where probably most of the data will stay

Design of the tiering policies

Once the rough design on data is completed, the tiering policy to apply can be defined. Some parameters have to be set for the definition of the tiering policy: the execution mode and the execution interval in compliance with the performance needs of your applications.

Choose the execution mode

Applications with unsteady performance requirements and moving hot spots may tend to more relocation and the relocation may be counterproductive, because today's and tomorrow's hot spots may differ. The application may also have cycles in using its data, which also should be taken into account. For example batch jobs or backups from the productive volumes may provoke IOPS peaks in the night, whereas during daytime processing other parts of the volume may have performance peaks. This is the reason why the automatic execution mode may not be applicable for some applications. In this case semi-automatic or manual execution mode might be more appropriate.

Define the execution interval

The execution interval should be chosen in such a way that the intervals divide the time into parts, which reflect the performance windows of the application well, and allow a reaction with data relocation to benefit later timeframes. The interval should be short enough to enable the application to benefit from relocation, but also long enough to collect the right sample of the application.

Define the execute time

When defining an execution interval of days, the execute time of the evaluation and relocation process can be chosen. The time to start the execution should be set to a time where the storage is not heavily used and has sufficient free resources to relocate the data.

Monitoring and optimizing the tiering process

When automated tiering is in production mode, administrators should periodically monitor the history of the tier pool and analyze the usage of the sub-pools by the flexible tier volume. In case of insufficient performance or resources the administrator can then increase the size of the sub-pools and/or manage individual flexible tier volumes by setting sub-pool quotas and/or the initial sub-pool allocation for new data. The space usage of a sub-pool, as well as the allocated ratio in subpools per flexible tier volumes can be displayed on the ETERNUS SF Web Console. This helps administrators to understand the system status and to plan capacity upgrades more effectively.

Specifics of the automated stprage tiering implementation in ETERNUS DX

Fujitsu's implementation of AST was carried out with the aim of offering maximum flexibility and incorporates some special features which are not common in standard AST implementations that are available today.

- The ETERNUS DX AST feature is available regardless of the class of the storage system from the entry model ETERNUS DX100 S3 up to the enterprise models ETERNUS DX8700 S3 and DX8900 S3, offering the benefits of Automated Storage Tiering at an affordable cost to customers in small and medium-sized businesses.
- Licenses are priced at controller level so that customers have no restrictions in terms of the capacity used.
- In general, the fastest tier is built of expensive and fast solid state disks (SSD), while the middle tier is made of fast spinning SAS disk drives, with cheap and big nearline SAS drives being used for the lower tier. The configuration alternatives of AST in ETERNUS DX offer much more flexibility than these standard implementations. Both the configuration of tier policies with either two or three tiers is an option. In addition, there is no limitation with regard to a rigid mapping of tier classes to dedicated disk types.
- In Fujitsu's implementation disk drives and RAID levels can be freely assigned to any tier which gives storage administrators even more possibilities to finely balance the ratio between capacity, speed, and cost. Figure 9 shows that very flexible service levels can thus be assigned to applications.

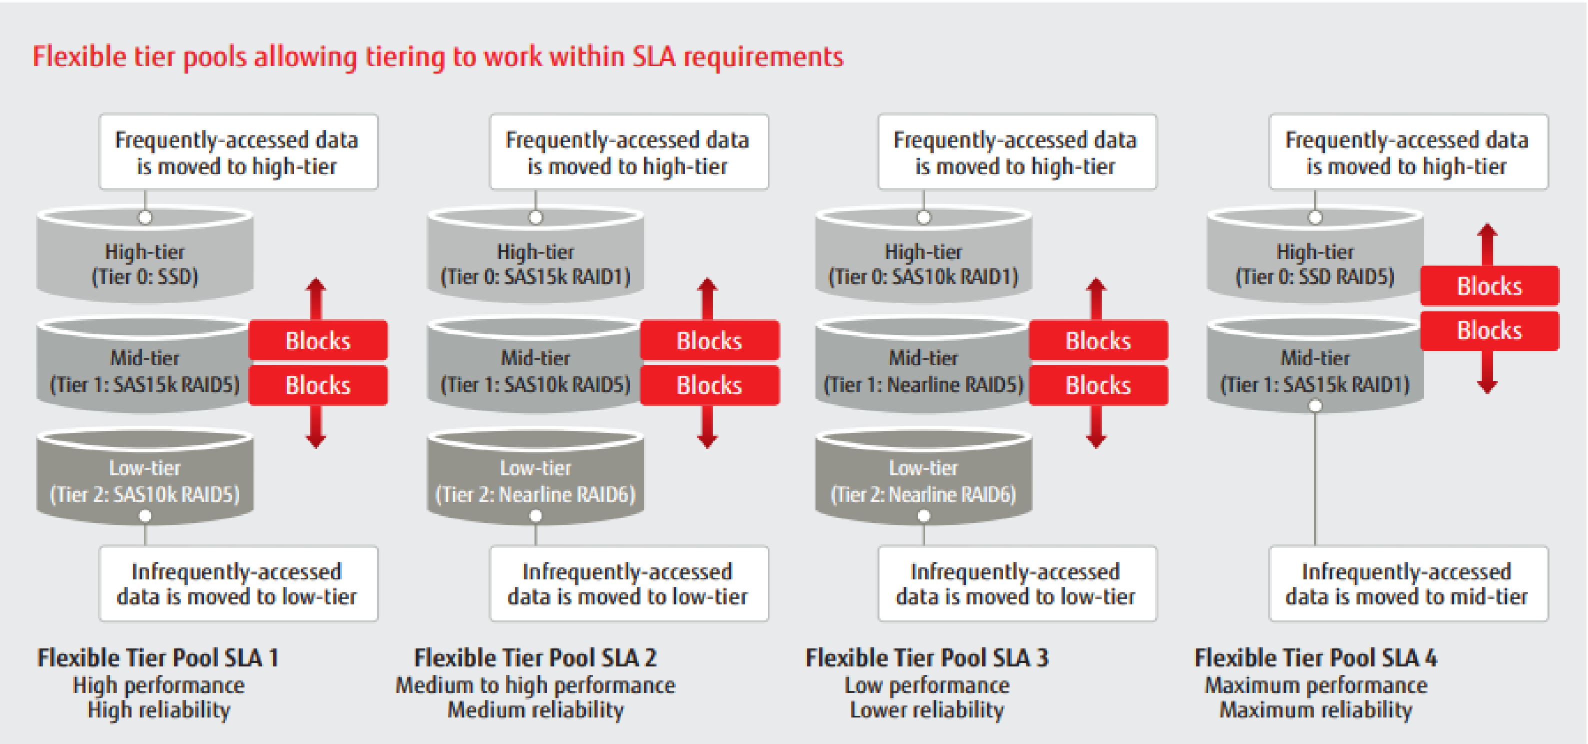

Figure 9

Automated QoS basic definitions

If different applications share the same storage platform, potential problems may occur like:

- Workloads with I/O and cache conflicts, such as online transaction processing (OLTP) and data warehousing
- Tiered storage access restrictions, such as development and production applications
- Peak processing demands for critical applications versus maintenance activities.

The ETERNUS AF/DX **automated quality of service** feature ensures that multiple applications can run on a single storage platform without affecting each other.

Performance limits are set for each connected server according to its priority. By prioritizing data access and dynamically managing any I/O conflict, high performance can be guaranteed for high-priority applications, and at the same time capacity is used more efficiently. The QoS policies allow the user to specify the expected I/O patterns of each application (random, sequential, read or write-based, and mixed).

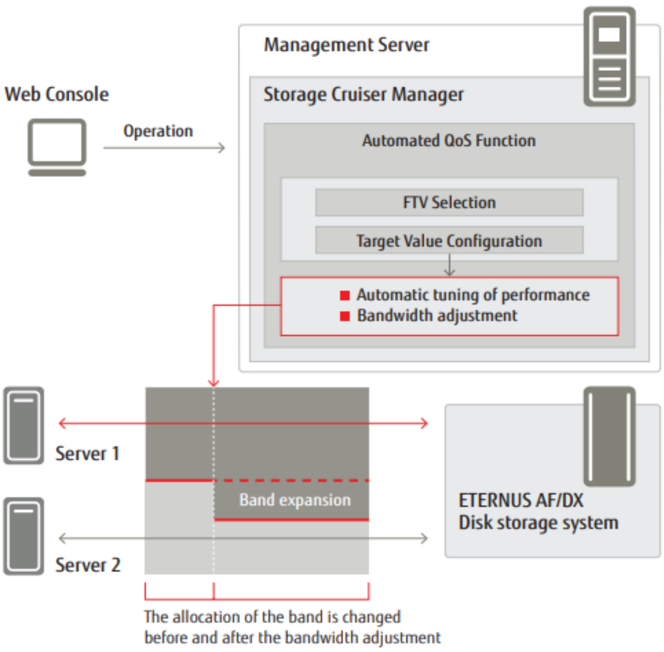

Figure 10

To define the Automated QoS targets two parameters can be set:

- **Automated QoS priority:** defines a service level for the volume representing its priority. Volume priority can be set as low, middle, high or unlimited. By setting priority to a volume, performance for the volume can be relatively adjusted in the other volumes.
- **Target response time:** defines the performance per volume as the average of the read + write response time of this volume. Performance for the volume is adjusted automatically according to the defined target response time.

Each target value can be set per flexible tier volume (FTV). Automated QoS limits the bandwidth of volumes of lower priorities as defined. Volumes of higher priorities are granted greater bandwidth and achieve better response times close to the target.

Other than the array-based quality of service option, which just limits the IOPS for specific volumes in a static way and requires a lot of expertise and continuous tuning to find the optimum settings, the ETERNUS AF/DX automated quality of service feature dynamically adjusts the I/O bandwidth along with the result of continuous performance monitoring. To facilitate these tasks the ETERNUS SF automated quality of service management option lets administrators set priorities based on performance requirements much more easily.

Automated QoS not only helps to avoid potential problems for shared storage access but resolves these issues and enables the consolidation of multiple applications with different I/O performance requirements in a single storage system.

Prerequisites and licenses

Required software

The automated QoS feature is controlled by the ETERNUS SF storage management software suite, which is delivered with any ETERNUS AF/DX storage array. ETERNUS SF can be installed on either Windows, RHEL or Solaris hosts as well as on virtual machines provided by VMware or HyperV.

Required licenses

ETERNUS AF

ETERNUS SF storage cruiser quality of service management option as well as the ETERNUS SF storage cruiser standard license are included in the free of charge All-in FlashPack software package coming with every ETERNUS AF system.

ETERNUS DX

The function is released by an optional license called ETERNUS SF storage cruiser quality of service management option as well as by the ETERNUS SF storage cruiser standard license; it cannot be set up with the ETERNUS SF storage cruiser basic license or the free-of-charge ETERNUS SF express. The licenses are activated for each ETERNUS DX system – regardless of installed disk capacity.

In addition, the hardware-based thin provisioning license must be registered on the storage array itself.

Setup of automated QoS

To run automatic tuning either Automated QoS priority or target response time has to be defined in order to control the priority of applications or the bandwidth they are allowed to use.

Setting an Automated QoS priority adjusts the priority of a volume relatively to the other volumes. One of the following values can be set to define the priority of business applications:

- **Unlimited:** Allows the volume to use the available bandwidth as much as possible, without limitation.
- **High:** Highest performance is required for the volume.
- **Middle:** High performance is required, but bandwidth limitation can be accepted if resources have to be shared with another volume with priority “High”.
- **Low:** High performance is not required. The volume tolerates bandwidth limitations to other volumes. If “Low” is selected, the Automated QoS function adjusts the other “High” priority volumes while considering the performance of the volumes that were selected with “Low”.

An alternative to the predefined priorities is the option to configure an average target response time for a volume. This allows even greater flexibility, because response times can be set according to the business needs within the following boundaries:

- 0 (msec): The configuration does not limit the bandwidth and uses it as much as possible.
- 1-10000 (msec): Target response time is set.

Tuning process

After setting automated QoS priority and target response times, QoS automation controls the system and allows volumes of higher priorities to use a greater bandwidth and brings the actual measured values closer to the target response times.

Tuning of the response times starts as soon as automated QoS is enabled. Its result is influenced by the following parameters:

- Automated QoS priority
- Current I/O performance
- Performance of disks

The next step is a check of the current value against the target response time. Bandwidth adjustment is performed if the measured response time of the volume does not match the expected targeted performance.

After achieving the performance targets, automated QoS continues to monitor the performance and makes adjustments when necessary.

However, automatic tuning cannot always guarantee achievement of the target value that was set, for example if multiple volumes with high priority are sharing the same resource.

If the performance target is not attained, the following actions can be taken:

- Check all settings to make sure that automated QoS is ready to be executed properly
- Lower the priority of a volume that shares the same resources
- Review the volume configuration to change the tuning parameters.

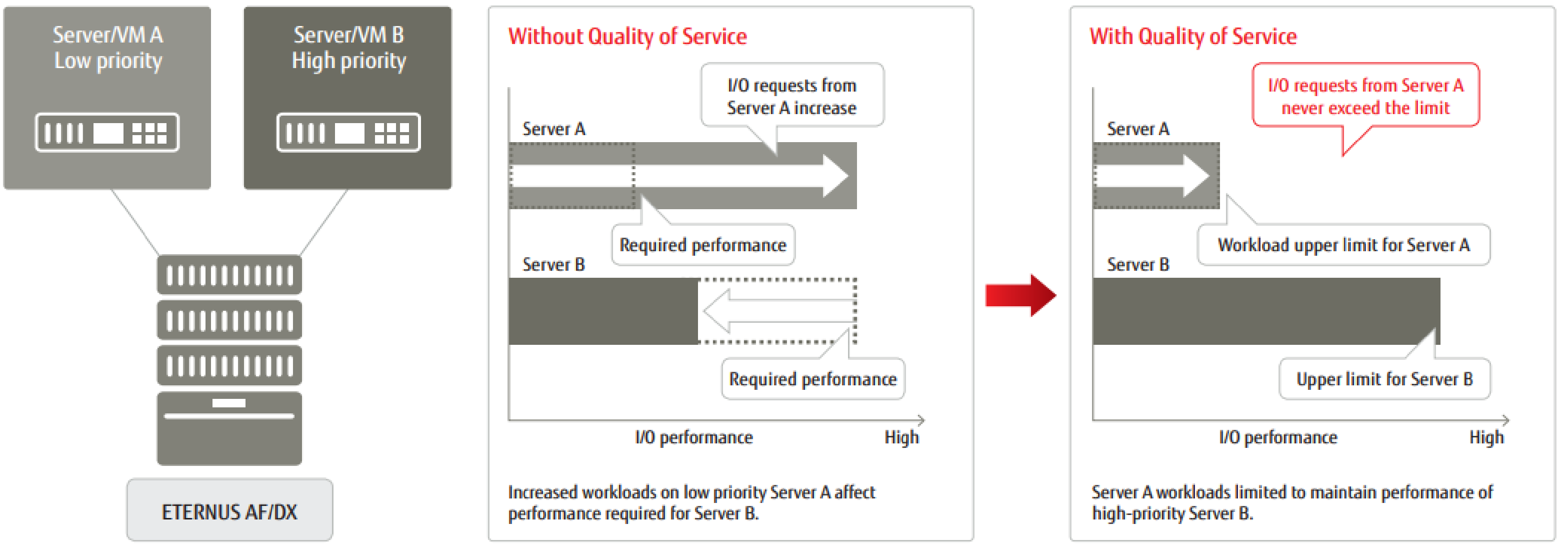

Figure 9

Example 1:

Bandwidth adjustment with automated QoS priority

The following example shows the adjustments that are performed with automatic tuning depending on the automated QoS priority setting.

- The three volumes (Vol#1 - Vol #3) show different automated QoS priority settings: Vol#1 is set to "High priority", Vol#2 is set to "Middle priority", and Vol#3 is set to "low priority".
- Figure 12 shows what happens when the automatic tuning is started at 15:00 h. Based on the measured performance of the high priority volume Vol#1 a calculation of the target response times for Vol#2 and Vol#3 is made. Then the calculated response time is compared to the measured performance for Vol#2 and Vol#3.
- If the results of the measured performance exceed the target performance, the bandwidth is narrowed by one level. If the measured performance is lower than the target performance, the bandwidth is widened by one level.
- At 15:10 h the three volumes are running with adjusted bandwidth reflecting the importance and performance demands of the corresponding applications.

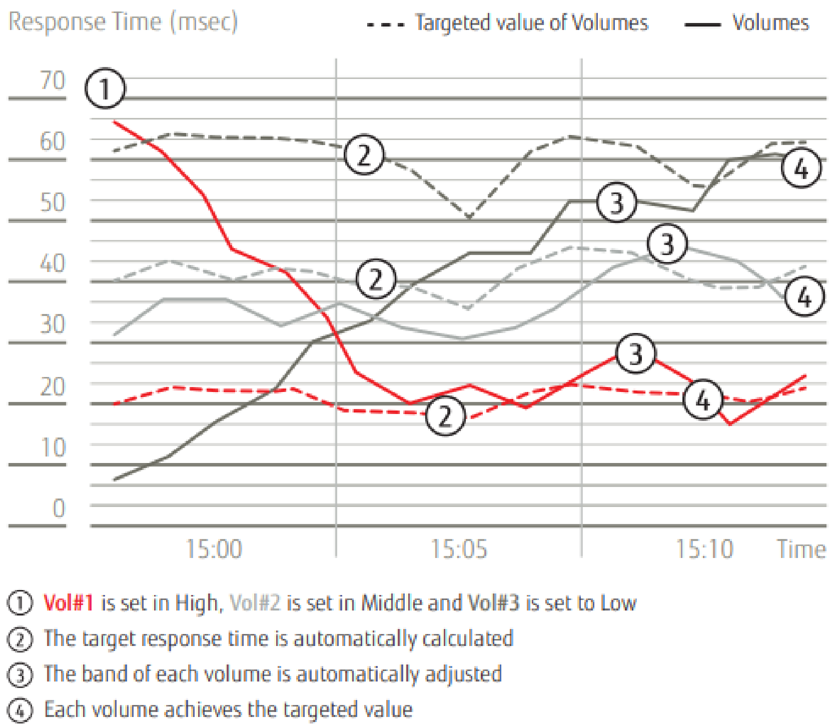

Figure 12

Example 2:

Bandwidth adjustment with target response time

The following example in Figure 13 shows adjustments that are performed with automatic tuning depending on the target response time settings.

- The three volumes need different response times: Vol#1 needs the fastest response times and Vol#3 the lowest ones.
- When automatic tuning is started at 15:00 h, the I/O performance of the business applications is checked using the web console performance graph screen. The measured response time of Vol#1 shows 50 milliseconds, slower than the required response time of 30 milliseconds. At this time there is no need to set the target response time for the other volumes (Vol#2, Vol#3) as well.
- To enhance the I/O performance of Vol#1 the Automated QoS function performs adjustments. It widens the bandwidth range for Vol#1 and at the same time narrows the bandwidth for Vol#2 and Vol#3 to approach the response time of the target value.
- At 15:10 h the target response time of Vol#1 is achieved and the I/O performance of the business applications running on it increased.

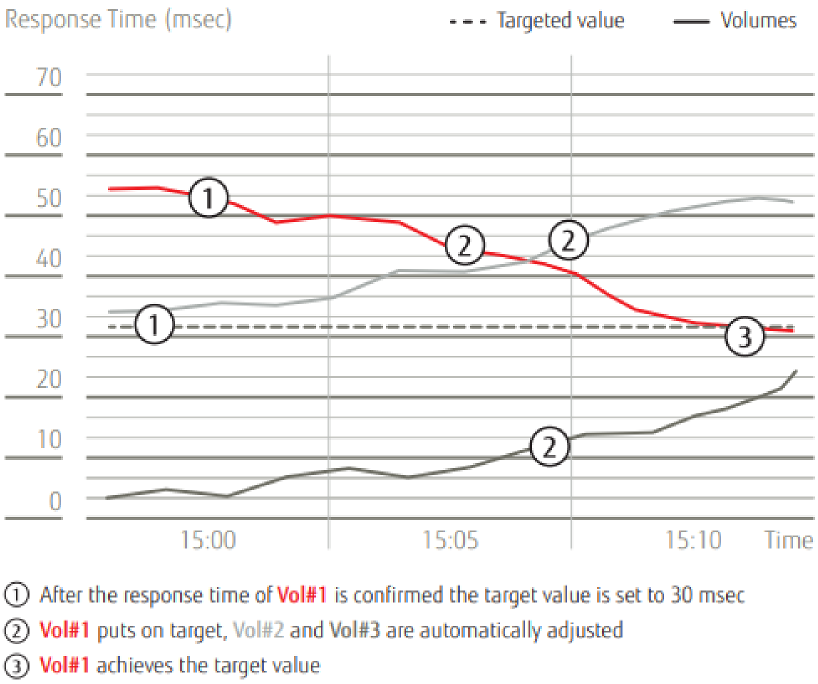

Figure 13

Best practice at high level

Before configuring automated QoS administrators have to clearly analyze and understand the data and application structure. This document provides some high level hints. For more detailed design and best practices a variety of documentation is available. Here are the basic points administrators should consider when building an Automated QoS infrastructure:

Define rough performance needs

Applications may have different types of data with different performance requirements. It is crucial to have a clear understanding of the performance requirements and the corresponding demands for the storage system

Size disk resources

Depending on the performance estimation some applications will certainly need high speed storage areas as SSDs, some others are seen in the medium performance range and a third part with slow performance requirements. Decide in favor of sizing of the different storage pools.

Decide in favor of tiering pool configuring

Automated QoS needs a tiering pool in the background. This can be a dedicated tiering pool reserved for Automated QoS or the tiering pool of AST. To get the best results out of Automated QoS a combination with AST is recommended.

Categorize data

As applications are of different importance for the company, so are the data. Data of business-critical core applications need very fast access, while other data accept lower access. Depending on the importance of the data, volumes have to be categorized. Categorize the volumes according to the importance of the data on them.

Assign priorities

Priorities have to be assigned to the volumes according to the categories. Start using the fixed priority model first. This model is easier to use and shows in most cases the required results. Reduce the usage of the target response time for special requirements.

Check results

When Automated QoS is in production mode, administrators should constantly monitor the performance values and check them against the settings.

Please note: Automated QoS becomes more effective over time and sometimes it needs some time to reach the performance targets.

Specifics of the automated QoS implementation in ETERNUS AF/DX

The ETERNUS AF automated QoS feature is available for all models. Licenses come free of charge with every ETERNUS AF system and can be used without further activation.

The ETERNUS DX automated QoS feature is available regardless of the class of the storage system from the entry model ETERNUS DX100 S3 up to the enterprise models ETERNUS DX8700 S3 and DX8900 S3, offering the benefits of Automated QoS at an affordable cost to customers in small, medium-sized and enterprise businesses.

Licenses for ETERNUS DX are priced at controller level so that customers have no restrictions in terms of the capacity used.

The Automated QoS feature can not be used in combination with deduplication/compression.

Automated QoS and AST

The AST function of ETERNUS DX can be further improved by using automated QoS over a flexible tier pool.

If the target response time for a volume cannot be met through the automated QoS function alone, the hot blocks of volume are automatically moved to a faster tier by the automated storage tiering function. The quota shares of a volume in the different storage tiers are automatically adjusted to achieve the desired response time. An efficient and automated method to get the best possible performance for the business applications is provided by integrating automated QoS and AST.

The following picture (Figure 15) shows how AST can affect the response time of a volume which is already tuned by automated QoS. Without AST the volume cannot achieve its performance targets and runs above the required target response time. After integrating AST relocation to a faster tier takes place automatically and after some time a significant performance improvement can be seen.

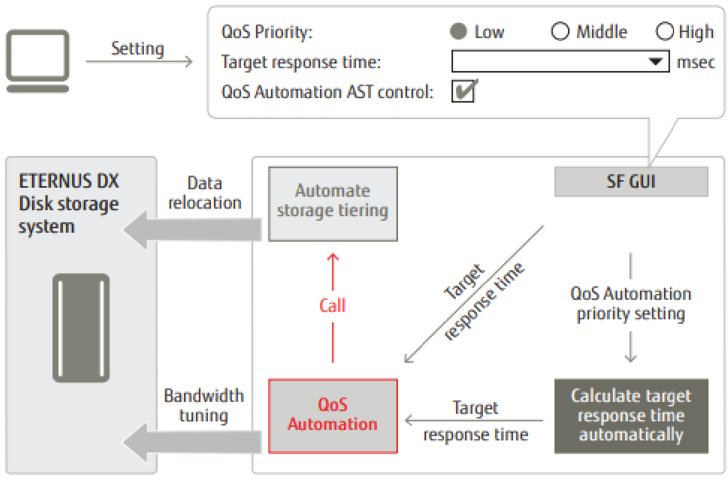

Figure 14

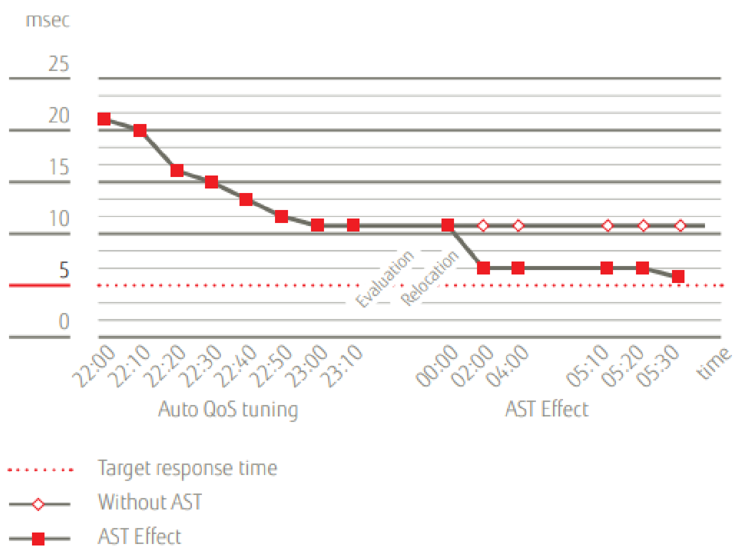

Figure 15

Conclusion

Compared with traditional setups with optimization features, customers can achieve highest performance at the same or even lower cost.

Many verifications show that the I/O response time can be improved significantly with Automated Storage Tiering and Automated quality of service while the cost for disk drive purchase and the footprint of the disk storage system can be cut. Fast SSDs can be used more efficiently and the automatic control of SLAs guarantees the fulfilment of customer performance requirements. In addition, low-power SSDs reduce the power consumption of the whole system. This also results in cost savings for power and cooling.

With an easy to understand and easy to use management interface administrators can intuitively create the complete subset of ETERNUS optimization features. They provide relief from time-consuming tasks like data analysis, manual movement of data to other tiers or manual adjustments of application priorities, while still offering the choice of semi-automatic or manual operation.

In summary, this makes the optimization features of ETERNUS AF and ETERNUS DX both a highly flexible and affordable solution.

.

FUJITSU ETERNUS AF/DX

Optimization Features

White Paper

For more information on ETERNUS optimization feature:

fujitsu.com/eternus

Contact

Fujitsu Technology Solutions GmbH
Mies-van-der-Rohe-Strasse 8
D-80807 Munich
www.fujitsu.com