IT@Intel White Paper

Intel IT
IT Best Practices
Cloud Computing and IT Efficiency
July 2011

Implementing Cloud Storage Metrics to Improve IT Efficiency and Capacity Management

Sudip Chahal
Principal Engineer, Intel IT

Krishnamurthy Anandarao
Storage Capacity Planner, Intel IT

Chris Peters
Director, Intel IT

Shane Healy
Enterprise Architect, Intel IT

Nigel Wayman
Storage Engineer, Intel IT

Steve Owen
Storage Engineer, Intel IT

Executive Overview

Intel IT is transitioning to an enterprise private cloud that will support our office and enterprise computing applications. This cloud is based on virtualized infrastructure accessing shared pools of storage through storage area networks (SANs).

Over the past year, we have developed and begun refining a storage metrics methodology that could span both our new private cloud and our traditional enterprise environment.

As we migrate to the private cloud, we seek to optimize the efficiency of the SAN environment by increasing storage capacity utilization. Demand for storage capacity proliferates by about 35 percent a year, and capacity purchases represent a significant and growing business cost. By increasing storage utilization, we can curb growth and contain cost.

We are evaluating a wide array of technologies and approaches to improve SAN efficiency and utilization over time, including thin provisioning. Because many customers use only a small portion of their storage allocation, thin provisioning technology allows IT to overallocate capacity and thereby increase utilization of storage resources. However, overallocation also requires improved capacity and risk management so that we have adequate capacity to meet customer requirements.

We recognized a need for enterprise-wide storage capacity metrics that accurately measure storage utilization and allocation across our SAN environment. Over the past year, we have developed and begun refining a storage metrics methodology that could span both our new private cloud and our traditional enterprise environment.

These metrics focus on three related areas: efficiency, capacity management, and risk management. We are building business intelligence and reporting capabilities based on these metrics. Our storage capacity metrics are designed to:

- Establish a clear link between cost (the purchase of raw capacity) and value (the use of that capacity to store internal customer data).
- Reflect efficiency gains due to new technology adoption and storage efficiency approaches such as thin provisioning and data de-duplication.
- Establish operational thresholds, based on allocation and utilization levels, that can alert us when we need to add or reallocate capacity before we jeopardize customer service levels.
- Remain uniform over time, providing us with a consistent view of efficiency across supplier product lines and technology generations.

We are starting to use these new metrics to analyze and compare the efficiency of our private cloud at multiple levels, from individual storage pools to a global view.

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

Custom reports provide an efficiency overview for senior IT managers, help data center managers plan capacity purchases, and enable operations engineers to respond to day-to-day capacity requests.

Potential future enhancements include customized risk management thresholds based on a predictive algorithm that can consider factors such as historical growth in customer storage requirements. We also anticipate being able to apply our metrics across our traditional office and enterprise network attached storage (NAS) environment.

Additionally, we have identified a need to add storage performance health metrics. As we scale our storage environment and increase utilization, performance health metrics will be essential to understand application and system requirements, identify potential bottlenecks, and assist in successful deployment. Accordingly, we plan to add performance health metrics into our capacity management and efficiency views.

BUSINESS CHALLENGE

Intel IT is undertaking a major transition to an enterprise private cloud that will support our office and enterprise computing applications. This phased, multiyear initiative is designed to enable greater agility and efficiency.

We are building this multitenant environment on virtualized infrastructure as a service (IaaS). This infrastructure is based on clusters of Intel® Xeon® processor-based servers accessing shared pools of storage through storage area networks (SANs), as shown in Figure 1.

As we transition to the private cloud, SAN storage is becoming an increasingly important element of Intel's IT storage infrastructure. Currently, Intel has a total about 25 PB of global storage; Intel's office and enterprise environment accounts for about 8 PB of this. Within this office-enterprise environment, SAN storage is used to support our enterprise private cloud as well as mission-critical and other applications.

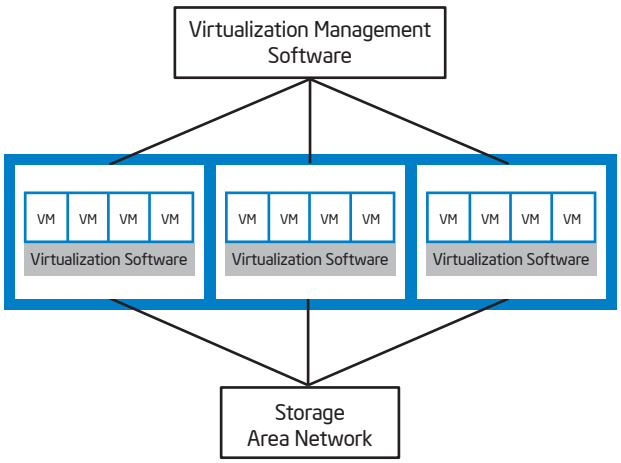

Figure 1. In Intel IT's private cloud infrastructure, servers access shared storage pools through a storage area network.

The cost of SAN storage is becoming a significant factor in overall IT costs. This is because SAN storage is more expensive technology than locally attached storage and because a larger proportion of our storage is becoming SAN based. Key factors in SAN capacity growth include:

- **Server virtualization.** To build the foundation for our private cloud, we began accelerating the pace of server virtualization in late 2009. This drives incremental demand for SAN storage, because data that was previously stored on physical servers' local drives is now migrated to the shared SAN environment. We more than tripled the rate of virtualization in 2010; by the end of the year, about 42 percent of servers were virtualized, and our goal is to increase this to 75 percent within our office-enterprise environment.
- **Rapid growth in storage demand.** We are experiencing a 35 percent annual growth rate in raw data requirements due to trends such as increasing use of video and other graphical information, legal retention requirements, larger databases for operational and business intelligence use, and backup needs.

THE VALUE OF OPTIMIZING SAN CAPACITY UTILIZATION

With our increasing use of SAN storage, we saw the need for strategies and tools that would enable us to measure and optimize efficiency by maximizing SAN capacity utilization. Increasing utilization can reduce cost by helping to curb growth in the amount of capacity we need to purchase.

Due to the large size of the Intel IT environment, even small efficiency improvements can result in significant financial savings. This is because capacity purchases are the largest factor in Intel IT storage costs; other costs such as operating costs, power and cooling, SAN switches, and cabling represent a much smaller percentage of overall total cost of ownership (TCO).

THE NEED FOR ENTERPRISE-WIDE METRICS

To manage capacity and improve efficiency, we need to be able to accurately measure capacity utilization across our SAN environment. To date, this has been challenging for several reasons:

- Storage suppliers often use different terminology and methodologies for measuring capacity utilization. It is difficult to combine information from multiple suppliers into a single view that reflects Intel IT business requirements.
- Existing metrics do not accurately reflect efficiency improvements achieved using new technologies such as thin provisioning.

Due to the complexity of the environment and the lack of a standard method for measuring storage utilization, we have found inconsistencies in the way we measure efficiency across different groups and teams within Intel IT.

We recognized the need for a single set of metrics that provides standard measures of utilization across our diverse and constantly changing technology base.

STRATEGY FOCUS AREAS

These enterprise-wide metrics must account for three interrelated focus areas within our storage management strategy: efficiency, capacity management, and risk management. The relationship between these three areas and their importance to different groups within Intel IT are described below and shown in Figure 2.

Efficiency

Clearly, our goal is to achieve high storage capacity utilization without impacting our internal customers or increasing IT operational costs. Storage efficiency metrics provide visibility into how efficiently our storage capacity is being used to store customer data. Efficiency metrics are particularly interesting to Intel IT executives and senior managers, who view storage efficiency as a key IT indicator.

Capacity Management

As we increase efficiency and utilization, capacity management becomes increasingly important. We need to be able to analyze utilization growth rates and accurately anticipate and determine when we need to purchase more capacity. This enables us to make purchases in the most timely and cost-effective way.

Capacity management-related metrics are used by a variety of people within Intel IT. These include the data center managers, capacity planners, and operations engineers who are responsible for making sure there is enough SAN capacity to meet demand.

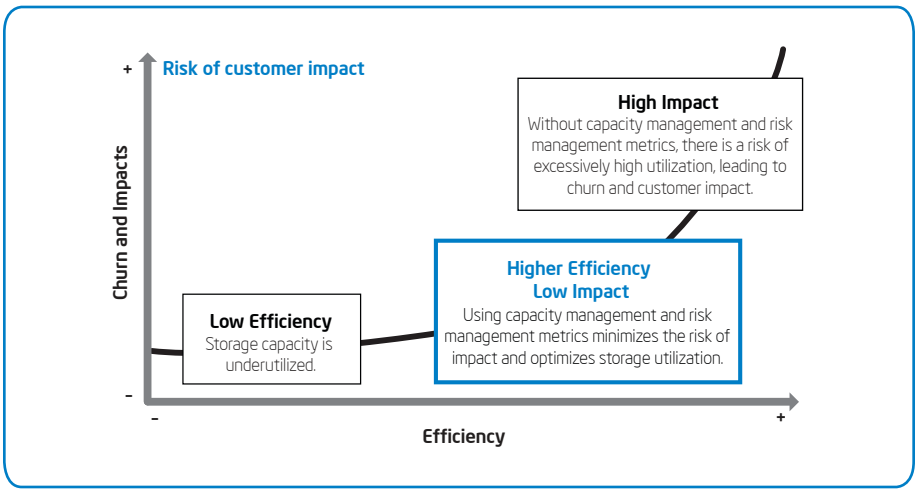

Figure 2. To maximize storage efficiency, enterprise-wide metrics must focus on three interrelated areas within our storage management strategy: efficiency, capacity management, and risk management.

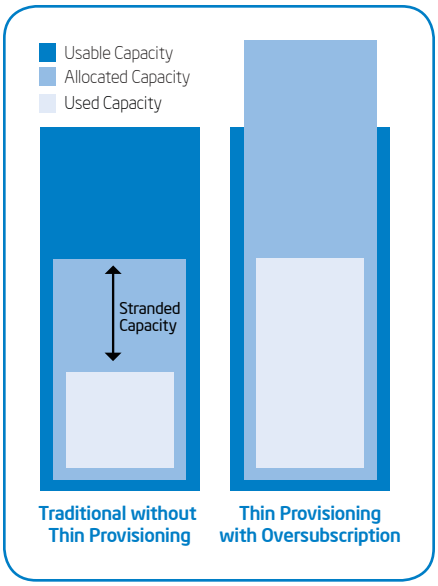

Figure 3. Thin provisioning enables us to oversubscribe (overallocate) capacity based on the assumption that business groups ask for more capacity than they immediately need. Oversubscription effectively increases the capacity that we can allocate to our internal customers, without requiring additional raw storage.

Risk Management

As we drive storage efficiency and utilization to higher levels, there is an increasing risk of customer impact. In an extreme case, running out of capacity can result in an unacceptable disruption to business applications. Very high storage utilization can also impact application performance.

There is also the potential for increased IT operational costs. To avoid customer impact, Intel IT operations engineers may devote a considerable amount of time to manually migrating data from heavily utilized to less-utilized pools

These potential issues mean that as we improve efficiency, we also need to develop metrics that enable us to manage risk and have capacity available when needed. These include capacity utilization and allocation thresholds that can be used to trigger alerts before utilization reaches unacceptably high levels.

Risk management metrics are of primary interest to Intel IT storage operations engineers, who are responsible for making capacity available to support day-to-day business needs.

IMPROVING STORAGE EFFICIENCY WITH THIN PROVISIONING

We are applying a number of approaches to reduce cost and curb storage growth by utilizing capacity more efficiently. These include thin provisioning, reclaiming unused storage, and sharing storage infrastructure across multiple projects. Together, these approaches saved Intel a total of USD 15.8 million in 2010.

Thin provisioning is a key technology, especially within our private cloud. The concept is analogous to the way that airlines overbook flights. As shown in Figure 3, we can oversubscribe, or overallocate, capacity to provide customers with adequate headroom, based on the assumption that they typically ask for more capacity than they need. The effect of oversubscription is to increase capacity utilization and enable us to avoid purchasing more capacity than customers actually need.

Because we are oversubscribing available capacity, our storage capacity metrics must be able to measure the efficiency benefits achieved with thin provisioning and help us manage capacity and risk.

We are widely deploying thin provisioning to increase efficiency across our private cloud storage environment. We began using thin provisioning in 2010 and plan a full rollout in 2011.

We are also exploring the use of data de-duplication, which recognizes when multiple instances of the same data are stored by different customers. The technology consolidates the multiple data sets into a single copy, freeing space that can then be allocated for use.

STORAGE METRICS

We formed a cross-functional team to create a new storage metrics framework for use across Intel's environment. The team included members from Intel IT architecture, engineering, and operations groups, and from Intel business groups.

At the core of this framework is a single set of capacity utilization metrics that we can apply across our entire SAN environment to support efficiency, capacity management, and risk management. In the future, we also anticipate being able to apply this approach across our traditional office and enterprise network attached storage (NAS) environment.

We designed our metrics so that we can aggregate data to create different views that enable us to analyze storage use and efficiency across the entire storage environment. We can analyze the metrics data by data center, storage tier, and storage frame. In the future, our goal is to show by individual customers.

Key Storage Metrics and Concepts

Our metrics are designed to support our three primary strategy focus areas: efficiency, capacity management, and risk management.

EFFICIENCY METRICS

We defined three core metrics that measure the storage efficiency of different aspects of the environment. We believe that these metrics represent a minimum set necessary to measure the efficiency of storage capacity use and thereby help contain cost. The metrics are orthogonal, or independent: Ultimately, to maximize efficiency throughout the environment, we need to optimize all three.

- Slot Utilization. The ratio of storage frame slots that are populated with drives compared to the total available storage frame slots.
- Overall Storage Efficiency. The ratio of customer stored data compared to the raw storage capacity.
- Low-cost Storage Percentage. The ratio of customer data stored on our low-cost storage tier compared to the total customer data stored.

Implementation of these metrics across the private cloud is a significant undertaking. Therefore, we are taking a phased implementation approach, focusing initially on the Overall Storage Efficiency metric.

Intel IT Storage Tiers

Intel IT organizes storage area network (SAN) storage into tiers that differ based on their performance, cost, reliability, and capacity. Each tier consists of storage frames containing drives. Our high-end tier provides the greatest performance and reliability, but the cost can be 2x to 4x higher per gigabyte as compared to lower tiered storage. Due to constant technology advances, drive specifications change frequently and are not listed here.

The primary tiers include:

- **H1.** Highly reliable, scalable high-end frames utilizing high-performance Fibre Channel (FC) drives
- **M1.** Mid-range frame with high-performance FC drives
- **M2.** Mid-range frame with lower cost high-capacity SATA drives

Our private cloud currently uses the M1 and M2 tiers. The H1 tier is used for mission-critical applications such as enterprise resource planning software, databases, and other software that has not yet been virtualized and implemented in our private cloud.

We need to balance use of the different tiers to maximize efficiency while maintaining adequate capacity at each tier.

Slot Utilization Percentage

Frames are expensive and occupy costly data center capacity. Therefore, efficient utilization of frame capacity is an important factor in minimizing overall storage cost.

Our measure of frame capacity utilization is the Slot Utilization Percentage metric. Each frame includes a manufacturer-defined maximum number of slots. Intel IT storage operations groups may restrict the number of these slots that are available for use to provide good performance for Intel workloads. As shown in Figure 4, we define Slot Utilization Percentage as the percentage of these available slots that are populated with drives.

Overall Storage Efficiency

The goal of our storage strategy is to cost-effectively store data required by our internal customers. Therefore, we defined a metric that measures how efficiently storage capacity is utilized for storing customer data. This metric, Overall Storage Efficiency, is defined as the ratio of stored customer data (Used Capacity) to raw storage capacity (Raw Capacity), as shown in Figure 5. Improving Overall Storage Efficiency is a primary focus of our storage strategy.

Challenges of measuring Used Capacity: data duplication and orphaned data. Measuring Used Capacity presents some complexity for IT organizations because of data duplication (multiple instances of business data stored by one or more customers) and inefficient management of orphaned data (stored data that is no longer accessed by customers). The challenge is that the customer's view of the amount of data stored—the Used Capacity—may differ from the storage frame view. As we evaluate data de-duplication technology, we are initially focused only on the storage frame view of Used Capacity as a reasonable approximation for customer stored data.

Low-cost Storage Percentage

This represents the percentage of customer data that is stored on the lowest-cost tier, which is currently our M2 tier (see sidebar). Intel IT is planning storage tier initiatives with the goal of moving less critical business data to lower cost tiers while continuing to support customer needs.

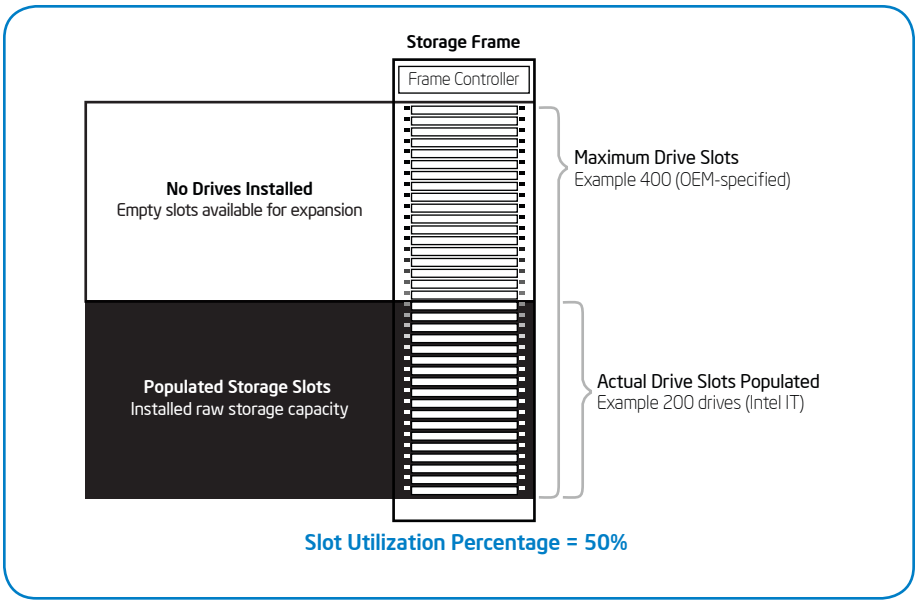

Figure 4. Slot Utilization Percentage metric.

CAPACITY MANAGEMENT AND RISK MANAGEMENT METRICS

Some storage efficiency techniques, such as thin provisioning, increase utilization but also increase the risk of having insufficient capacity to meet growing and changing business needs. If we oversubscribe capacity (allocate more than is actually available), it becomes even more important to carefully monitor how much of this capacity is being used, how quickly usage is growing, how much more capacity we can safely allocate, and when we need to add or rebalance capacity.

Our key capacity management and risk management metrics are designed to support these requirements. The metrics measure allocation and utilization within each storage pool, but can also be aggregated to provide summary views at the frame and data center level. The metrics measure allocation and utilization relative to Usable Capacity, not Raw Capacity, since this is what matters to the IT operations team and our customers. The metrics are shown in Figure 6 and described below.

To minimize risk, we need to be alerted when allocation or utilization levels reach predefined thresholds. We have established several initial policy-based thresholds, as described below. We expect to refine these policies as we gain more experience in production. For example, we anticipate being able to incorporate better forecasting of customer utilization behavior and trends in storage demand growth, enabling us to include "time to full" projections. We also anticipate including performance health monitoring.

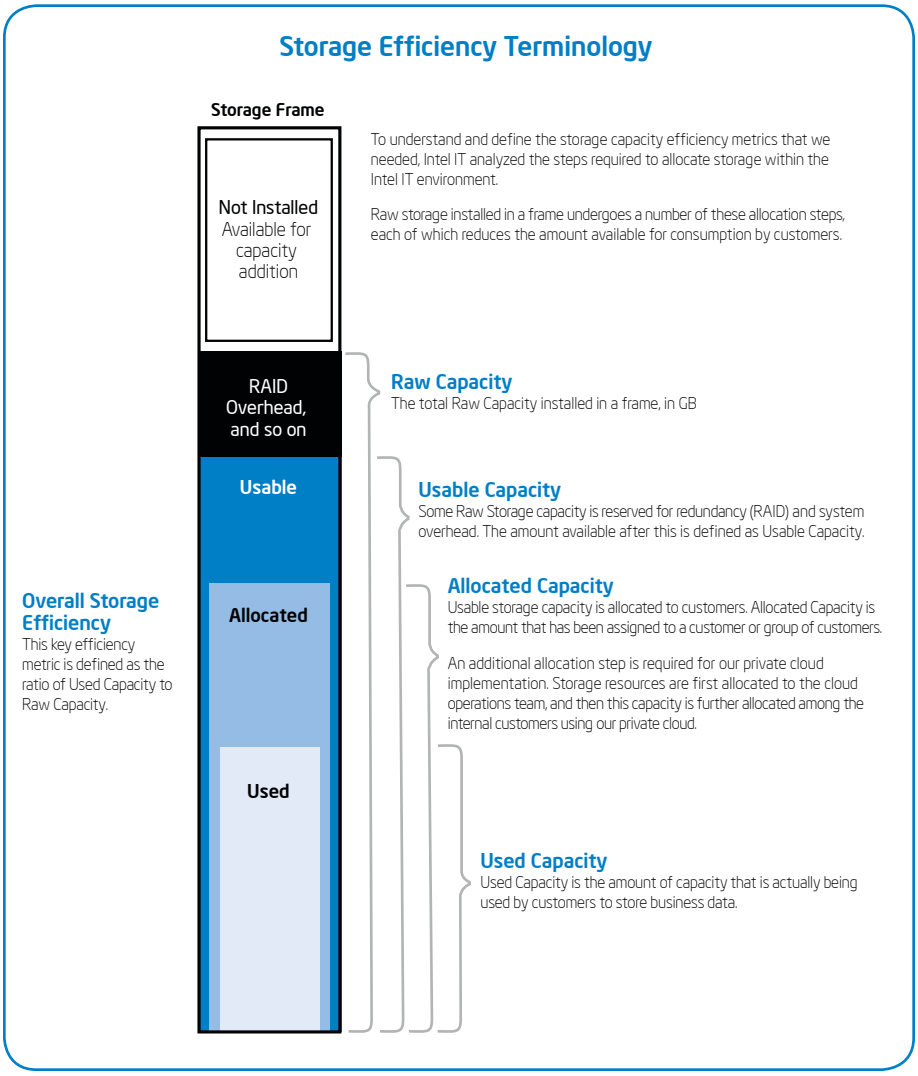

Figure 5. Overall Storage Efficiency and other key metrics.

Used Percentage

This is the percentage of Usable Capacity that is used to store customer data. It is defined as the ratio of Used Capacity to Usable Capacity. By definition, this metric can never exceed 100 percent. We currently define two operational thresholds to manage risk as storage pools reach high utilization levels.

- **Close pool threshold.** When a pool's Used Percentage reaches 50 percent, the pool is considered closed. No further capacity can be allocated from this pool; only organic growth of existing customer storage workloads is allowed.
- **Rebalance pool threshold.** When the Used Percentage for a pool reaches 70 percent, storage capacity must be reassigned to other pools to avoid the risk of customer impact.

Allocation Percentage

This is the percentage of Usable Capacity that has been allocated to customers. With thin provisioning, this metric can exceed 100 percent. We have established an allocation limit of 150 percent. Once a pool's Allocation Percentage reaches this level, IT operations will not allocate more capacity, regardless of how much of the capacity is actually being used by customers.

Customer Utilization of Allocated Capacity

This represents how much specific customers have used the storage capacity that has been allocated to them. Ideally, all customers would consume 100 percent of their allocated capacity. However, most customers use only a small percentage of their allocation. Higher levels of storage pool oversubscription are possible in pools with lower customer utilization. We often refer to this metric simply as Customer Utilization.

Allocation Headroom

Ultimately, the operations team needs to know whether they can allocate more storage from a pool in response to customer demand and, if so, how much storage can be allocated. The Allocation Headroom metric is designed to enable IT storage operations engineers to quickly and easily determine this.

Allocation Headroom is defined as the additional storage capacity, in GB, that can be allocated from a pool. It is calculated based on several factors, including Allocated Percentage, Used Percentage relative to the thresholds described above, and Customer Utilization of Allocated Capacity. We define policies that alert operations engineers when a pool's Allocation Headroom has decreased to a predetermined low level.

ANALYZING EFFICIENCY: BUSINESS INTELLIGENCE AND REPORTING

We have begun building business intelligence and reporting capabilities based on our metrics. Our goal is to use these capabilities to analyze and manage SAN storage across our office and enterprise private cloud.

We created a storage resource management (SRM) tool that automatically gathers storage capacity-related data across our SAN environment. Using the SRM tool, we can pull data from individual storage pools or aggregate data by frame or by data center. This allows managers to compare efficiency across different data centers, frames, and pools.

As of May 2011, we have established regular on-demand reporting across 56 pools in 15 data centers, representing over 1.37 PB of raw storage within our private cloud.

We use the information gathered by the SRM tool to create customized reports for different purposes and users within Intel IT.

For example, executives may receive a high-level overview of average efficiency and other summary data. Engineers responsible for day-to-day operations need reports that are at a more granular level and to alert them to potential problems such as overallocated storage pools.

Some example reports are shown below. These show actual Intel IT operational data; we have changed some identifying details to protect Intel proprietary information. We are applying thin provisioning to increase efficiency within these data centers.

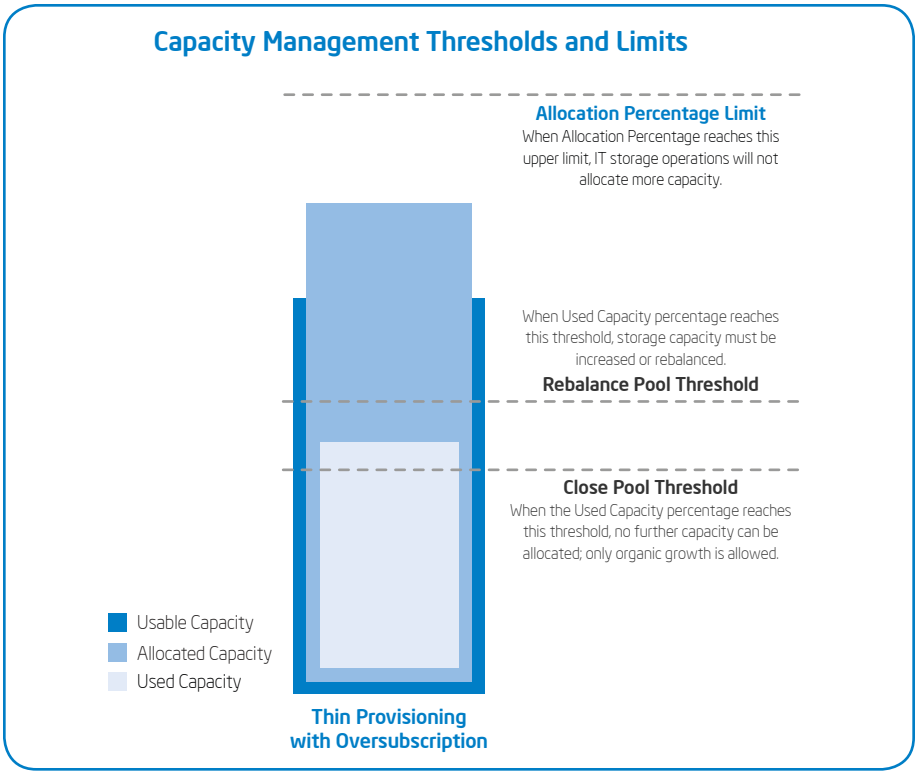

Figure 6. Intel IT capacity management and risk management metrics applied to a storage pool with thin provisioning.

EXECUTIVE STORAGE EFFICIENCY REPORT

This report, shown in Table 1, includes the Overall Storage Efficiency metric and also provides a capacity overview for Intel IT executives and senior managers at the data center level.

DATA CENTER MANAGER REPORT

Table 2 shows a report designed to help Intel IT data center managers administer storage capacity. These managers need to know how much usable capacity exists, how that capacity is being allocated and used, and how much headroom is available in each storage tier.

This information is critical for planning purchases of more storage capacity based on established IT operations policies and customer service-level agreements (SLAs).

The report illustrates how our metrics and policies are used to create triggers that alert our data center managers to the need for action. In the example below, the available Allocation Headroom for Fiber Channel (FC) storage at several data centers has decreased to low levels, triggering status alerts that indicate the need for more capacity.

While this report provides a useful overview at data center and storage tier levels, data center capacity managers typically also need information that helps them manage capacity at the level of individual frames.

A typical scenario is illustrated in the example of a frame-level report shown in Table 3. Overall, Data Center A appears to have adequate FC headroom available and is therefore considered open for new storage allocations.

However, two frames provide this FC capacity, and an examination of the metrics for each frame reveals that while Frame 1 has headroom for further allocation, Frame 2 does not. This is because the Used Percentage in Frame 2 has exceeded our 50 percent threshold. Frame 2 is therefore closed to further allocations, even though the data center as a whole remains open.

This frame-level report tells data center managers that they need to take action

Table 1. Executive Storage Efficiency Report

Data Center	Raw Capacity (GB)	Used Capacity (GB)	Overall Efficiency Percentage
Data Center 1	114,174	13,999	12.26%
Data Center 2	14,762	2,266	15.35%
Data Center 3	40,260	13,328	33.10%
Data Center 4	106,296	28,766	27.06%

Table 2. Data Center Manager Report

Data Center	Storage Type	Usable Capacity	Allocated Capacity (GB)	Allocation Percentage	Allocation Headroom (GB)	Used Capacity (GB)	Used Percentage	Pool Status
Data Center 1	Fibre Channel	14,980	19,271	128.64%	1,337	6,962	46.48%	**Closed: Add Capacity**
Data Center 1	SATA	71,428	72,722	101.81%	34,419	7,037	9.85%	**Open**
Data Center 2	Fibre Channel	11,770	16,464	139.88%	1,191	2,266	19.25%	**Closed: Add Capacity**
Data Center 3	Fibre Channel	32,101	33,903	105.61%	3,744	13,328	41.52%	**Closed: Add Capacity**
Data Center 3	SATA	43,395	45,136	91.38%	7,289	20,049	40.59%	**Closed: Add Capacity**
Data Center 4	Fibre Channel	32,100	32,456	101.11%	15,694	8,717	27.16%	**Open**

Table 3. Frame-level Data Center Report

Data Center	Storage Type	Frame	Usable Capacity (GB)	Allocated Capacity (GB)	Allocation Headroom (GB)	Used Capacity (GB)	Used Percentage	Pool Status
Data Center A	Fibre Channel	All Frames	59,908	64,809	21,317	18,990	31.70%	**Open**
Data Center A	Fibre Channel	Frame 1	50,278	52,887	21,317	13,621	27.09%	Open
Data Center A	Fibre Channel	Frame 2	9,630	11,922	0	5,369	55.75%	Closed: Add Capacity

by working with the operations team to rebalance storage utilization among the existing frames or by acquiring more storage capacity. If they identify that more storage capacity is required, they can examine slot availability to determine whether there are available slots in existing frames or whether new frames need to be purchased.

POOL STATUS REPORT

Operations engineers responsible for day-to-day storage management need more granular information at the storage pool level. This enables them to respond to requests for more storage by determining which pools have enough capacity.

Table 4 shows a pool status report that indicates, based on our usage and allocation thresholds, whether individual pools are open or closed. The report also shows each pool's available Allocation Headroom.

In this example, FC Thin Pool 1 for Data Center 1 is closed because two thresholds have been exceeded: our 150 percent Allocation Percentage limit and our 50 percent Used Percentage close pool threshold. Two pools in Data Center 2 are also closed because the Used Percentage close pool threshold has been exceeded. The other pools remain open.

CONCLUSION AND NEXT STEPS

Storage capacity utilization metrics enable us to improve the efficiency of our private cloud by providing a framework we can use to analyze and manage SAN storage use across our diverse IT environment. Metrics also provide information required by Intel IT for capacity management and risk management.

We find that the reports and metrics help our storage operations engineers and storage capacity managers work together to maximize storage use efficiency. The reports help them identify the best course of action. For example, when a pool's utilization or allocation reach high levels, they can determine whether they need to purchase new capacity or can simply rebalance storage across existing pools, frames, or data centers.

We have established on-demand reporting across 56 pools in 15 data centers, representing more than 1.37 PB of raw storage within our private cloud. We continue to increase the reporting coverage across our private cloud, and, in the future, we also anticipate being able to apply this approach across our traditional office and enterprise NAS environment.

We continue to enhance our metrics. For example, we have developed an adaptive algorithm that can create customized allocation headroom estimates for each pool based on factors such as the rate at which a customer's capacity utilization is increasing.

Additionally, we identified a need to add storage performance health metrics. As we scale our storage environment and increase utilization, performance health metrics will be essential to understand application and system requirements, identify potential bottlenecks, and aid in successful deployment. Accordingly, we plan to add performance health metrics into our capacity management and efficiency views. These performance metrics could include concepts such as response time, queue length, and storage processor utilization.

In 2011, we plan a full roll-out of thin provisioning, and we are also exploring the use of other efficiency technologies such as data de-duplication. Our metrics will enable us to measure the efficiency improvements resulting from the implementation of these technologies.

Table 4. Pool Status Report

Data Center	Storage Type	Frame	Pool	Usable Capacity (GB)	Allocation Percentage	Allocation Headroom	Used Capacity (GB)	Used Percentage	Pool Status
Data Center 1	Fibre Channel	Frame 1	Thin Pool 1	10,700	154.57%	0	6,695	62.57%	**Closed: Add Capacity**
Data Center 1	SATA	Frame 1	Thin Pool 3	21,953	96.60%	161	7,891	35.94%	**Open**
Data Center 1	Fibre Channel	Frame 1	Thin Pool 4	4,240	136.97%	2,264	761	17.95%	**Open**
other pools not shown									
Data Center 2	Fibre Channel	Frame 2	Thin Pool 8	9,630	95.70%	0	5,369	55.75%	**Closed: Add Capacity**
Data Center 2	Fibre Channel	Frame 3	Thin Pool 9	9,630	101.54%	0	5,809	60.32%	**Closed: Add Capacity**
Data Center 2	Fibre Channel	Frame 3	Thin Pool 10	10,700	103.64%	5,810	99	0.93%	**Open**
other pools not shown									

CONTRIBUTORS

Husni Bahra
Asoka Diggs
Jason Herrick
Joshua J. Jin
Jeff McLeod
Guillermo Rueda
Terry Yoshii

ACRONYMS

FC	Fibre Channel
IaaS	infrastructure as a service
NAS	network-attached storage
SAN	storage area network
SLA	service-level agreement
SRM	storage resource management
TCO	total cost of ownership
VM	virtual machine

For more information on Intel IT best practices, visit www.intel.com/it.

Printed in USA
0711/CPP/KC/PDF

Please Recycle
325586-001US