



89 Fifth Avenue, 7th Floor
New York, NY 10003
www.TheEdison.com
212.367.7400

White Paper

Oracle ZFS Storage Appliance Comparative Management Costs Study



Printed in the United States of America.

Copyright © 2010 Edison Group, Inc. New York. Edison Group offers no warranty either expressed or implied on the information contained herein and shall be held harmless for errors resulting from its use.

All products are trademarks of their respective owners.

First Publication: August 2010

Produced by: Kalicharan Rakam, Senior SME; Barry Cohen, Editor-in-Chief; Kelly Strand Anderson, SME; Craig Norris, Editor



Table of Contents

Executive Summary	1
Introduction	3
Objective	3
Audience	3
Contents of this Report	3
Background	4
Methodology	6
CMCS Methodology	6
Laboratory Configuration	6
Test Procedures	7
Test Results	10
Comparative Management Cost Savings	10
Qualitative Analysis	12
Quantitative Analysis	18
Data Protection	23
Maintenance and Configuration Changes	24
Monitoring	25
Conclusions	40
Appendix	41
The Methodology Defined	41
The Study	41
Tasks	41
Complexity Calculation Formula	43



Executive Summary

The drumbeat for expanding storage capacity continues for organizations of all sizes. Today, the pressure is to add storage capacity for databases. Tomorrow, users are clamoring for more space for their documents. Next week, storing video files becomes critical to the training department. Then R&D needs some place to store data for analysis from the production monitoring system. Every day all the business units are demanding new virtual machines to host their applications. The message is, more, more, more: now, now, now: block storage, file storage, streaming storage: don't you ever lose anything we store. Ever.

Storage systems dedicated to providing block storage for databases over fibre channel do not excel at supporting CIFS file systems. NAS filers, ideal for storing user productivity software created files, usually do not provide block storage for databases. Line-of-Business applications running on virtual machines are not budgeted to connect to expensive and complex fibre channel SANS.

One way of addressing these disparate and conflicting demands is by deploying Unified Storage systems. These systems support the popular storage network protocols; Fibre Channel, iSCSI, NAS. They also provide both block and file-level storage architectures. Since these platforms provide such a broad range of functionality, they are inherently complex. A sure way to lower the ownership costs for these systems is to temper that complexity by having easy to use management software.

Edison Group has developed a methodology for comparing management systems called Comparative Management Cost. Using this methodology, Edison subject matter experts compared the management software for the NetApp FAS family of storage systems, Data ONTAP, with Integrated Software Suite of Oracle's Sun ZFS Storage Appliance.

The results of our testing show overall that a storage administrator can perform a series of day-to-day tasks in 36 percent less time with 38 percent less complexity on the Sun ZFS Storage Appliance than is required for the NetApp FAS filer. These savings can be perceived as providing a Full Time Equivalent operating cost savings of almost \$27,000 per year. These savings in time also mean that administrators can provision storage for new applications or virtual machines more quickly and at lower cost. The vastly superior monitoring tools provided with the Sun ZFS Storage Appliance can greatly reduce the time required for troubleshooting (our tests showed a 44 percent advantage), which can translate into better customer service, and a higher end-user productivity due to a greatly reduced potential for downtime. Reducing potential downtime also



translates into better maintenance of revenue streams derived from the applications using the storage.

For organizations in need of flexible, multipurpose storage systems, Edison strongly recommends consideration of Sun ZFS Storage Appliances. The several models available can meet a wide range of capacity and reliability requirements while lowering operating costs through greater ease of management.



Introduction

Objective

The goal of this paper is to provide the reader with an understanding of the cost of ownership and the business value advantages of choosing Oracle's ZFS Storage Appliance family of Unified Storage devices. In order to illustrate these advantages, Edison compares its ease of management for day-to-day tasks, with an emphasis on troubleshooting, with that of the NetApp FAS family of storage arrays.

Audience

The audience for this white paper is IT and business decision makers who are well informed about current technologies.

Contents of this Report

- Executive Summary – a brief review of the context for the research and Edison's findings.
- Introduction – this section.
- Background – a more thorough discussion of the context in which Unified Storage systems have come to exist and the management challenges inherent in the technology.
- Methodology – an overview of the Edison CMCS Methodology, the laboratory configuration used for testing, and the specific procedures used in this study.
- Test Results – this section includes both a qualitative analysis that compares the testers experience with the software from the two vendors and a quantitative analysis that describes each set of tests and the results of that testing.
- Conclusions – this section provides a final summary of findings and recommendations from Edison subject matter experts.
- Appendix – a more detailed explanation of the Edison CMCS methodology.



Background

The rapid growth of digital data requires the near continuous growth of storage capacity in organizations of all sizes. This rapid growth is adding to Information Technology budget constraints in many organizations, as other IT investments are deferred and head count is restricted to compensate for acquisitions needed to accommodate capacity growth. The discussion that follows briefly describes storage technologies for capacity growth. It also describes the need to lessen the complexity of managing these technologies as a means for controlling IT costs while supporting IT goals. Among these goals is greater business agility enabling IT to focus on revenue-generating applications of technology through more efficient IT management.

Today, storage capacity growth is usually provided through some form of network storage. The currently predominant storage network protocols include Network Attached Storage (NAS) where dedicated storage servers, alone or pooled, provide users and servers with file storage functionality. Another popular storage network protocol is Fibre Channel. In a Fibre Channel storage network, storage arrays¹ and servers (hosts) are connected over a dedicated fiber optic network. A third, more recently emerging storage network protocol is iSCSI (Internet Small Computer System Interface). In an iSCSI network, hosts and arrays are connected over an Ethernet Network using the TCP/IP protocol. Though iSCSI can run over the same physical network that the hosts and client computers use, most iSCSI implementations use a dedicated Ethernet network, sometimes even utilizing specialized iSCSI network interface cards (NICs) in the hosts. A new storage network protocol, Fibre Channel over Ethernet (FCoE), has emerged that enables the Fibre Channel protocols to run over an Ethernet network. Intended to utilize 10 GB Ethernet technologies, FCoE is expected to simplify physical network provisioning in large data centers.

Choosing from the many platforms offered largely depends upon the amount and types of data that must be stored, the current and expected rates of data growth, and the nature of the applications creating and accessing the data. The focus of this white paper is on those organizations that have heterogeneous storage requirements. These include

¹ A note on terminology: In this white paper, storage devices are described using several labels. An *Array* generally refers to the hardware containing the disk drives. The *Controller* is the computer system that provides the intelligence to operate the storage device and to which the drives and hosts are attached. *Array* and *System* are used interchangeably to label the combination of disks and controllers. *Platform* is used to describe both the physical storage devices and the operating system and management software that enables the physical storage to be utilized by hosts.



file services (the underlying services that enable file servers to share data) and block-level storage (both the means by which data is stored to disk and the preferred storage system for most enterprise database systems). On a storage network, file services are generally delivered via NAS devices while block storage is delivered over FC or iSCSI protocols. For many organizations, all three protocols are in use, with the choice based upon application and cost considerations.

An ideal storage platform for these organizations is one that can be accessed using any storage network protocol. These platforms are usually labeled *Unified Storage* or *Network Unified Storage*. Since these storage platforms provide both block-level and file system storage access over NAS, FC, iSCSI, and other storage network protocols, management of these devices is inherently complex.

Controlling this complexity is critical if organizations are to have flexible, dynamic IT infrastructures. This is especially the case as the number of storage devices and their capacity continue to grow.

Edison set out to compare the manageability of two Unified Storage platforms, one from Oracle and the other from NetApp, in order to demonstrate how better systems management efficiency can both lower costs and enable greater flexibility. The remainder of this report describes the research methodology and the results of testing that Edison used to quantify the relative efficiency of the two platforms. We also attempt, through depiction of how several tasks were performed, to provide the reader with a qualitative perspective on the testing.



Methodology

Edison's team of subject matter experts deployed storage arrays from NetApp and Oracle in a lab within its New York City facility. The subject matter experts then performed a series of management tasks, capturing the steps and time required using the Edison Group Comparative Management Costs Study (CMCS) Methodology. Details about the methodology as well as the lab configuration and the testing process follow below.

CMCS Methodology

Edison developed the CMCS methodology to enable an objective comparison of the manageability ease of use for enterprise technology products. The CMCS methodology provides an objective metric for comparing Information Technology management platforms. More subjective metrics, such as those derived from interviews with groups of testers in typical user experience testing, are excellent for gaining an understanding of the preferences of the testers at the time of testing, but the results can vary significantly for different groups of testers and different test conditions. Objective metrics based upon mouse clicks cannot adequately address issues such as the difference between a Wizard-based versus a non-Wizard based interface. The Edison methodology attempts to address these concerns by establishing a set of rules that can account for complexity and different interface design philosophies. A full description of the methodology can be found in the appendix of this report.

Laboratory Configuration

Edison's test lab was configured with arrays from NetApp and Oracle. **Table 1** lists the storage array models, hardware features, and the licensed software installed in the lab.

Table 1: Lab Configuration

NetApp
Hardware
Filer Model FAS2040 as a single enclosure and drive shelf
12 500 GB SATA Disks
Licensed Software
Data ONTAP

NetApp
NFS, iSCSI, ASIS (Dedupe) NearStore, HTTP, SyncMirror, OPS- DFM Management Server, SnapDrive for Windows, SnapRestore, SnapVault ONTAP, FlexClone, MultiStore, NearStore. Additional available software was neither licensed, nor used.
Oracle
Hardware
ZFS Storage 7310, configured as a cluster
2 J4400 Drive Enclosures
24 1 TB SATA Disk Drives
Software
Oracle Open Storage Software

While a review of acquisition costs is not part of this study, it is important to note that the Oracle array software license is available at no extra charge and includes a wide range of capabilities for each of which a NetApp customer must separately purchase and renew licenses. In addition to any potential cost savings, not needing to acquire a license to use an unanticipated feature saves considerable time and administrative complexity.

Test Procedures

This study focuses on regularly performed administrator tasks. Edison divided the tasks into categories that relate to the tasks typically performed by storage administrators on a daily basis²:

- Provisioning – Provisioning includes all the tasks or operations performed after installation in order to attach an array to servers and store data. These include global tasks such as creating a storage pool (Oracle) or aggregate (NetApp), configuring network ports for connectivity to hosts, and the creation LUNS so the array can actually be put to use.
- Data Protection – Data protection covers such array features as the creation of snapshots and clones.

² The tasks listed for the categories are illustrative not inclusive. Not all of the tasks listed were evaluated nor are all of the possible tasks for a category listed. In addition, there in reality, these categories are not exclusive: tasks listed for one category may be justifiably perceived as belonging to a different category.



- Maintenance and Configuration Changes – Maintenance and Configuration Changes include time spent performing maintenance tasks such as adding additional capacity, cleaning up unneeded snapshots, clones, and so forth².
- Monitoring and Troubleshooting – Monitoring an array consists of identifying performance and operational parameters that affect uptime and performance, setting alerts, and generating reports on status and system condition. Troubleshooting consists of identifying the causes of issues identified through monitoring or from other sources such as server administrators or end users concerned about performance or data availability. Only troubleshooting tasks were evaluated in this study.
- Overhead or Other – these tasks include such tasks as budget or planning meetings, training, and other activities not directly related to any specific storage array. These activities were not evaluated in this study.

Edison’s subject matter experts identified tasks under each category and matched them between the two platforms compared in this study. Since the vendors use either different nomenclature for the same feature or the same nomenclature for different features, this was a very important first step.

Administrators rarely perform specific tasks in isolation. Instead, the administrator usually has a goal that requires the performance of a combination of several tasks or operations. To emulate this approach, Edison further categorized the task list to reflect a goal-oriented approach. This also simplified comparisons where the architectural differences between the two platforms would make task-to-task comparisons difficult or even impossible. The table (Table 2, Page 8) shows the list of tasks and goals, organized by category.

Table 2: Task Categories

Provisioning	
Baseline Provisioning: Tasks that need to be performed prior to creating usable storage on any storage system	
Preparing Storage	This series of tasks includes the creation of a storage pool (Oracle) or aggregate (NetApp) from the disks in the array.
Preparing Connectivity	For this study, this series of tasks includes all the operations required for setting up iSCSI initiators and targets on the array. Configuration of other network protocols was not evaluated in this study.

Provisioning	
Configuring Usable Storage	This category includes the creation of accessible storage (defined as a Volume or FlexVol on the NetApp filer and as a Project ³ on the Oracle array) and the creation of a LUN for attachment by a host.
Data Protection	
Creating Snapshots	These tasks include the creation and scheduling of snapshots.
Creating Clones	These tasks include the creation of clones.
Replication	These are tasks related to replicating data onto remote storage arrays. Replication tasks were not performed during this study.
Maintenance and Configuration Changes	
Removal of unused storage objects	This includes the deletion of old snapshots or clones, deletion of LUNs, etc.
Monitoring and Troubleshooting	
Monitoring and Troubleshooting	This includes the performance of a series of tasks related to the identification of problems on the arrays.

³ In a technical sense, a NetApp Volume and an Oracle Project are very different things. Volumes on a NetApp filer (whether Traditional or Flexible – go to www.netapp.com to find out more) reserve space on the aggregate or drive pool. An Oracle Project performs the same configuration role as a NetApp volume, but no actual physical capacity (beyond an insignificant 50 K at configuration) is reserved.



Test Results

This section provides a summary plus Edison's analysis of the test results for the study and for each task category. It includes a description of the different approaches to storage management taken by Oracle and NetApp, a tabulation of the complexity and time required to perform the various tasks, and detailed depictions of several of the tasks. These detailed depictions are intended to provide readers with a qualitative understanding of the quantitative data.

Managing complex systems is inherently complex. Edison believes that the management tools an administrator uses can be a critical factor in dealing with that multifold complexity. In large part, the nature of the management tools a vendor provides is deeply tied to both the product's history and the vendor's approach to systems. The two storage platforms are examples of how a product's history and the vendor's approach affect interface design and function.

Edison believes that well designed Graphical User Interfaces (or GUIs) can greatly improve the quality of system management as compared to the exclusive use of a command line interface (CLI) or scripting. This is especially true for the less skilled staff usually tasked with day-to-day management.

Edison does not oppose the use the CLI or scripts, but we believe that, for the majority of day-to-day tasks, a GUI can reduce training and task-oriented support costs and prevent operator error, both significant factors in administrative efficiency and thus cost. In fact, the ideal GUI should enable the use of scripts by providing an interface for running and creating scripts through learning or similar capabilities.

Comparative Management Cost Savings

The Edison CMCS analysis assigns a complexity metric and measures and compares the time required to perform different tasks. The savings in time delivered by a less complex management platform can be given a financial value. In most real-world business environments, management costs will far outweigh vendor charged licensing and support costs throughout the life of the product. With this in mind, we estimated the annual cost savings a business might expect due to storage administrator time savings that result from one product being easier to administer and operate than the other.



The overall results of Edison’s evaluation showed that managing the Sun ZFS Storage Appliance required 36 percent less time with 38 percent less complexity than was required for the NetApp Filer. A description of the testing and the results can be found in the sections that follow.

In order to compute cost savings Edison used Storage Administrator salary figures gathered from a query of salaries made at Simplyhired.com⁴. Edison searched for the salaries seven metropolitan areas of various sizes and costs of living. Though different choices would affect the results, Edison believes the average salary is reasonable.

Table 3: Average Storage Administrator Salaries

Average Storage Administrator Salaries	
Storage Administrator in NY, NY	\$87,000
Storage Administrator in Boston, MA	\$84,000
Storage Administrator in Chicago, IL	\$76,000
Storage Administrator in Memphis, TN	\$64,000
Storage Administrator in Houston, TX	\$71,000
Storage Administrator in San Francisco, CA	\$77,000
Storage Administrator in Birmingham, AL	\$65,000
Average	\$74,857.14

The average of the results of this search (Table 3, Page 11), when multiplied by the average time savings show cost savings:

$$\text{Median Storage Administrator Salary} * (\text{Storage Administrator Time Savings} = 74,857.14 * 36\% = \$26,948.57^5)$$

By substituting its own costs, an organization can determine its own potential savings. Edison recognizes that these savings are not going to be recognized directly. Instead, they need to be interpreted within the context of overall storage administrator duties and storage group or IT budgets. Within this framework, savings in time of over 35 percent can be equated as lessening the need for new hires when additional storage arrays are acquired or by enabling a better allocation of resources to mission critical tasks.

⁴ This and similar queries can be performed at <http://www.simplyhired.com/a/salary/home>.

⁵ See the Results Summary in Table 4, on Page 17.



It is not practical to attach direct financial costs to complexity. However, it should be apparent that management tools that are more complex require more user training, can cause user error simply due to complexity, and otherwise affect productivity.

Qualitative Analysis

NetApp Filer

At its inception, NetApp set out to change the data storage industry through the introduction of its innovative *Filers*. These network-attached storage devices were essentially file servers without the computational overhead that general-purpose servers required. Over the years, as competitors emerged and newly evolved technologies were incorporated that were complementary to the original Filer design, the simple, single-purpose Filer became a much more complex solution. When the Filers first appeared, management tasks were straightforward and the environments in which the devices were installed were almost exclusively managed from the command line. The same forces that led to the incorporation of additional technologies as well as the success of Microsoft Windows servers and GUI-based management systems have resulted in the evolution of GUI ⁶ management software for NetApp's filers as well.

NetApp currently offers two GUI-based management consoles. Data ONTAP is a web-based platform (or BUI) that runs on an embedded server within each filer. System Manager is a Microsoft Management Console plug-in application that adds the ability to manage multiple filers from the same console. In addition, NetApp Operations Manager is a server-based monitoring and reporting console that is accessed through a browser interface.⁷ NetApp also has a complete CLI for managing the filer.

A review of the documentation accessible from the Data ONTAP main menu as well as in many of the blogs and other reference materials available on the NetApp website reveals a preference for the Command Line. This preference is also evident in the structure of the ONTAP Browser User Interface.⁷ The navigation bar contains a combination of feature- or task-specific links (Figure 1, Page 13). Some of these links are to Wizards that direct the user through a series of the feature or task specific links in order to accomplish a task. In the illustration, the LUNs navigation branch has been expanded, revealing a LUN Wizard plus various other links for such activities as adding or managing LUNs, plus several sub-branches each with its own set of links.

⁶ In this white paper, GUI is used for the entire class of graphical user interfaces including both browser based and freestanding programs. The term, Browser User Interface or BUI, is used for those programs whose graphical interface is delivered via a web browser.

⁷ Edison used the browser-based NetApp Data ONTAP and NetApp Operations Manager in this study.

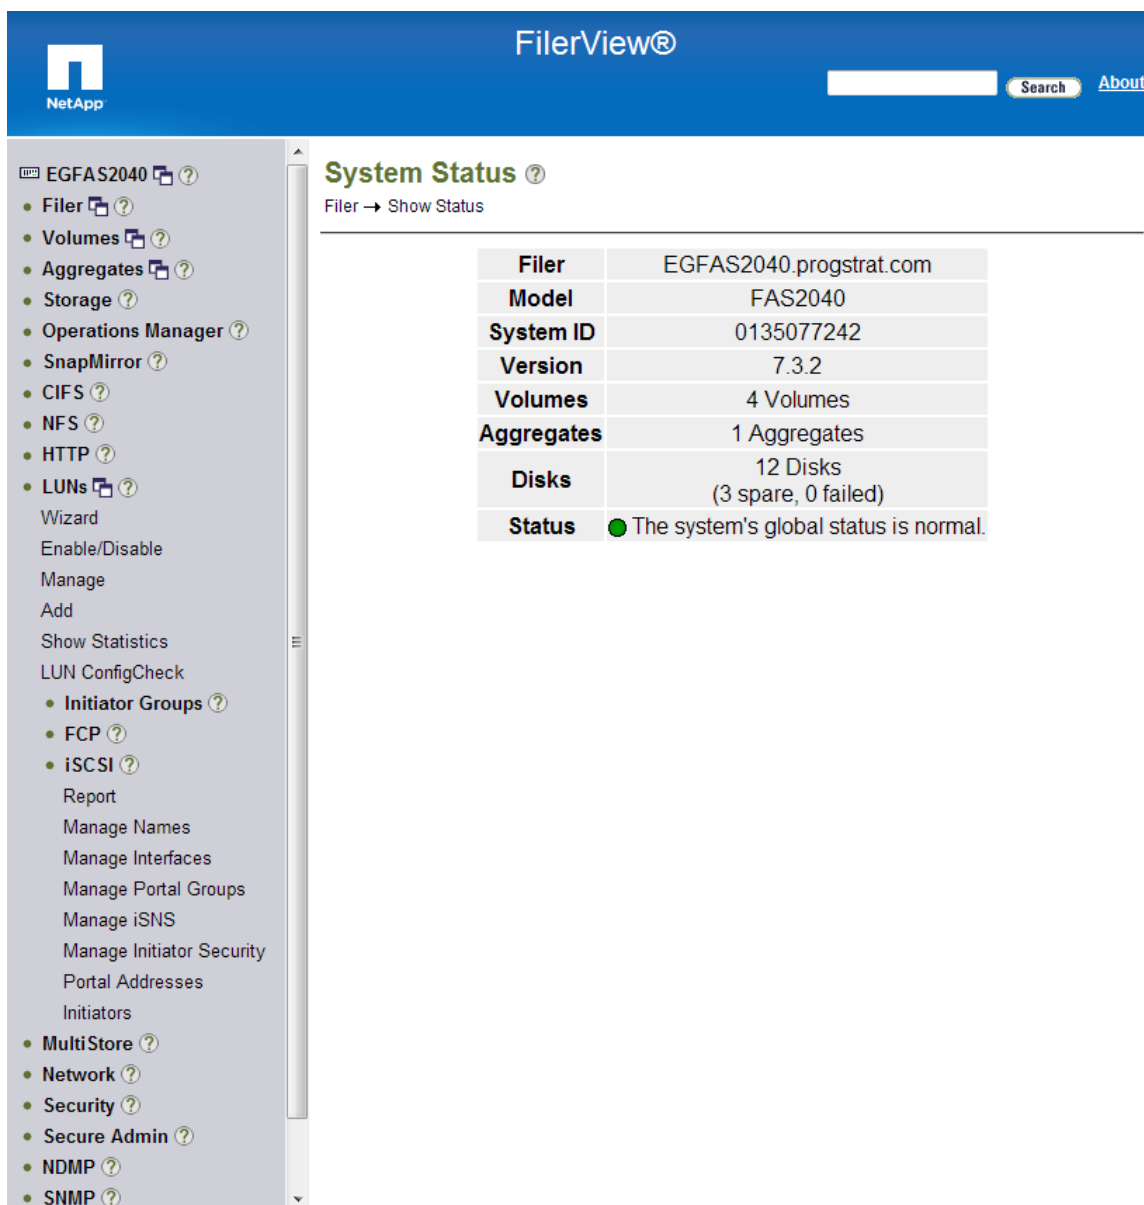


Figure 1: NetApp FilerView Home Page

This highly granular approach, including the Wizards, mirrors the menu-oriented structure common in complex CLI systems. An experienced storage administrator will know what most of the links refer to and can usually select the items to be configured with little difficulty. However, a less experienced administrator will need to develop, through practice, their own documentation in order to perform the series of tasks needed to accomplish various management goals, as it is unclear from the BUI alone which prerequisite tasks are required and in what order they must be performed.



NetApp provides good contextual Help, but ad hoc administrative efforts are not readily rewarded with appropriate results if the user makes the wrong initial selection.

Oracle Sun ZFS Storage Appliance

Sun Microsystems, now a part of Oracle, has long been in the forefront of UNIX-based workstations and servers. As part of its complete product offering, Sun has offered a range of storage systems, acquiring StorageTek in 2005 to fill out its offerings. Also in 2005, Sun integrated ZFS (originally the Zetabyte File System) into its Solaris operating system. ZFS provides Solaris with features that were quickly recognized as ideal for the creation of what, for Sun, was a new class of storage devices. These features include:

- High Storage Capacities — 2^{58} zettabytes where each ZB is 2^{70} bytes: a great deal of storage capacity
- Integration of concepts of filesystems and volume management into a single system
- Snapshots
- Copy-on-write clones
- Continuous integrity checking and automatic repair
- RAID-Z⁸
- Native NFSv4 ACLs – providing security as well as protocol support to take advantage of clustered server deployments, including the ability to provide scalable parallel access to files distributed among multiple servers (pNFS extension).

Another key feature provided by the Solaris operating system is DTrace. DTrace is a comprehensive, dynamic tracing framework developed by Sun Microsystems for troubleshooting problems on production systems in real time. It can be used to get a global view of a running system via such metrics as the amount of memory, CPU utilization, file systems, and network resources being used by active processes. On a Sun ZFS Storage Appliance, DTrace is used for the Status Dashboard and for the Analytics (Figure 2, Page 15) of the system. The troubleshooting section of this white paper is based upon utilization of the analytics functionality of DTrace.

⁸ Sun's ZFS implements an integrated redundancy scheme similar to RAID 5, which is called *RAID-Z*. RAID-Z avoids the RAID 5 "write hole" by its copy-on-write policy: rather than overwriting old data with new data, it writes new data to a new location and then atomically overwrites the pointer to the old data. It avoids the need for read-modify-write operations for small writes by only performing full-stripe writes; small blocks are mirrored instead of parity protected, which is possible because the file system is aware of the underlying storage structure and can allocate extra space if necessary. There is also *RAID-Z DP* which doubles the parity structure to achieve results similar to RAID 6, enabling the ability to sustain up to two drive failures without losing data., Triple-parity RAID was added in July 2009.

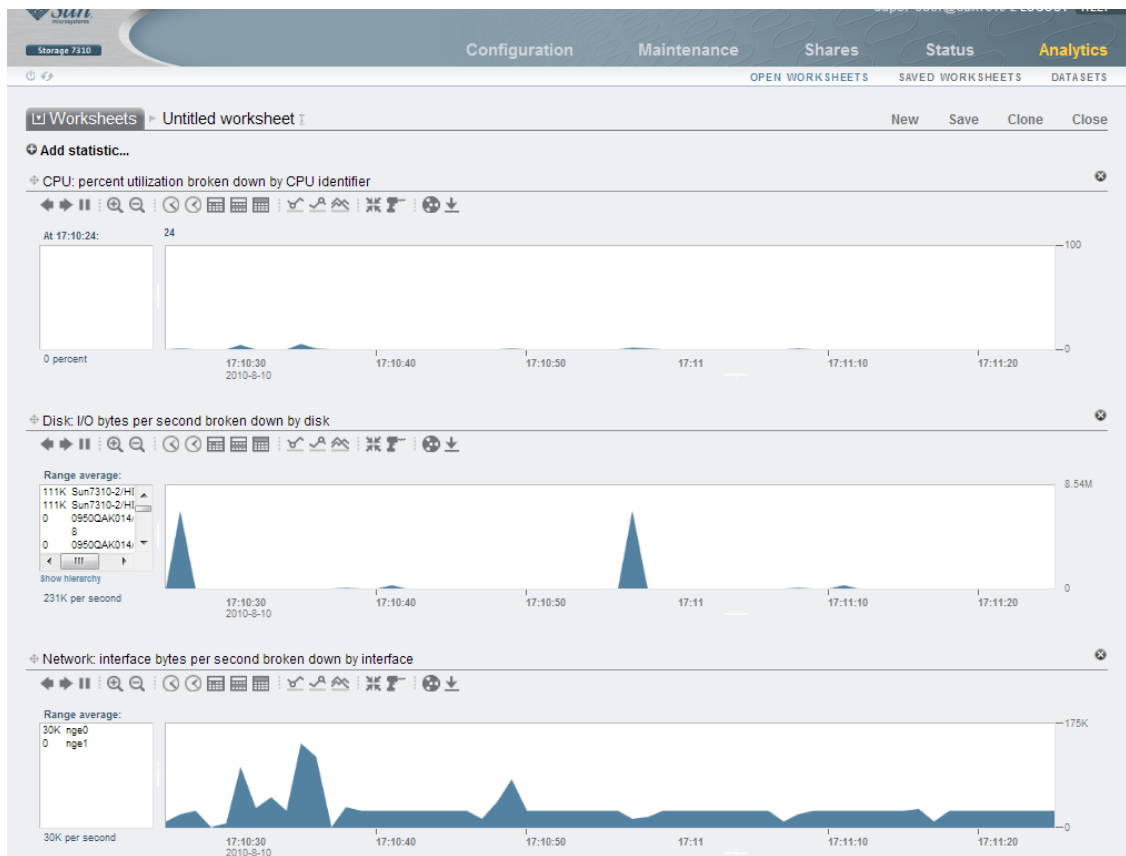


Figure 2: Oracle Integrated Software Suite - Analytics

Oracle had an inherent advantage in designing the Integrated Software Suite for Unified Storage Systems, the administration console for the Sun ZFS Storage Appliance: it did not have a legacy management platform to redesign. The Integrated Software Suite is a browser-based system that runs on the appliance in single or clustered configurations. It does not provide management of multiple appliances. Two key differentiators for the Integrated Software Suite as compared to NetApp Data ONTAP are:

- **Status Dashboard home page:** showing live graphic indicators of system performance and the activity-oriented layout. (Figure 3, Page 16)
- **Analytics page:** the location where administrators can access the wide range of monitoring tools provided through DTrace.

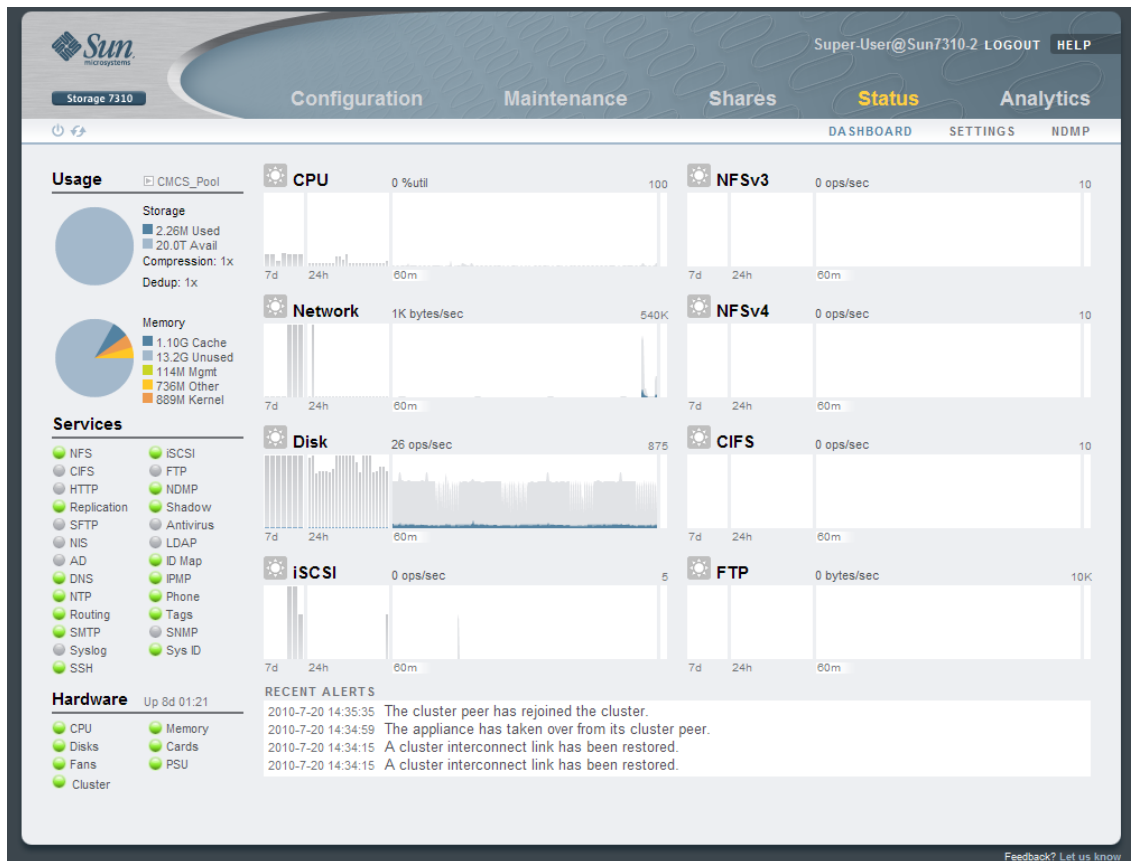


Figure 3: Oracle Integrated Software Suite - Home Page

The Status Dashboard will be covered in more detail later in the discussion of troubleshooting tasks. For current purposes, it should be obvious to most readers that the Oracle dashboard offers a more informative view than that offered by NetApp. The activity-oriented layout deserves additional discussion here. Rather than focusing on specific hardware or software features, the Oracle interface presents, in a horizontal, tab-like design, important tasks that an administrator needs to perform: Configuration, Maintenance, Shares, Status and Analytics. Each of these “tabs” provides a selection of activity tabs appropriate to the tasks. Clicking on any of these activities tabs presents the administrator with a graphic representation of status and click-through to the tools for performing or managing the activity.

Instead of using Wizards or a procedure-based configuration model, Oracle utilizes an object-oriented approach that provides for direct viewing and manipulation of all of the settings applicable for a specific function or feature. If an activity consists of several optional items, each of these items has its own sub-page. For example, the following two illustrations show the Network Configuration page and the configuration settings for one of the interfaces. (Figure 4: Oracle Integrated Software Suite - Network

Configuration, Page 17, and Figure 5: Oracle Integrated Software Suite - Network Interface Settings, Page 18) Though an extensive Help system is available for each page in the BUI, the pages each have an explanation for use that is generally adequate for navigation and task performance. Assuming a new administrator was familiar with the nomenclatures, such as those used on the Network Configuration tab, selection and configuration is straightforward. For administrators in need of clarification, the Help system provides both basic information on the terms and technology and the methodology for each page.

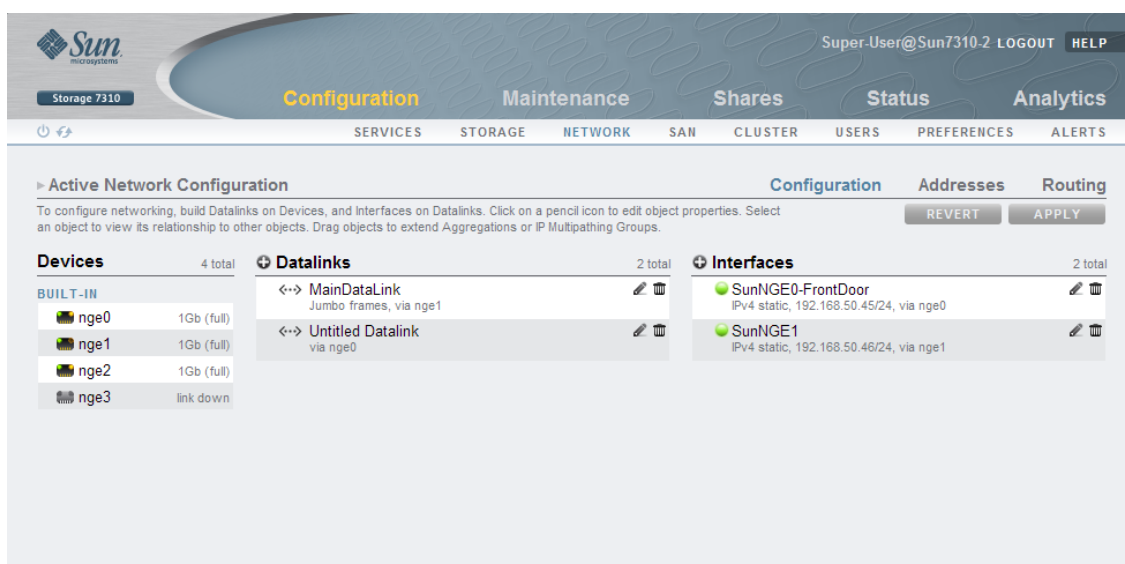


Figure 4: Oracle Integrated Software Suite - Network Configuration

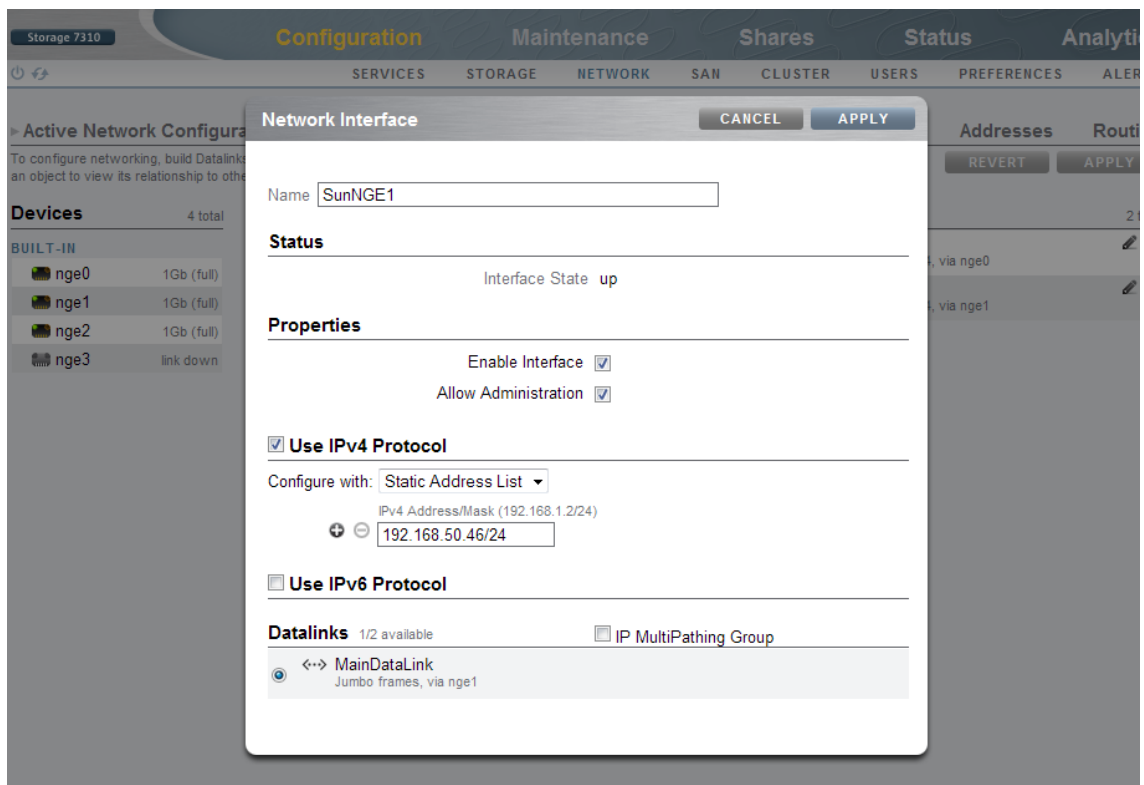


Figure 5: Oracle Integrated Software Suite - Network Interface Settings

Edison’s CMCS Methodology enables comparing the Wizard-oriented approach, where creating or changing configuration settings requires navigation of a series of Wizard pages, to an object-oriented approach, where all the options for creating or changing configuration are accessible on a single page. It does this by defining a Step as the commitment of a change that affects the state of the system, and an Increment as the number of items that need to be selected, whether on multiple pages in a Wizard or on the single object page, to perform the task.

Quantitative Analysis

Summary

A summary of Edison’s findings can be found in Table 4, Page 19. The results show that, in most task areas evaluated, managing the Sun ZFS Storage Appliance required an overall average of 36 percent less time with 38 percent less complexity. The one area where there was near equal performance was in the area of Data Protection where the only feature tested was creation of snapshots. The 5 percent time difference was insignificant and there was no difference in complexity. This is not surprising since



snapshot capabilities are inherent within the underlying technologies of both storage platforms.

The areas where the Sun ZFS Storage Appliance had its greatest advantage over NetApp were in Provisioning, especially in for configuring usable storage and Monitoring. As will be seen in the detailed results below, the Oracle Integrated Software Suite offers a much simpler approach to provisioning than that offered by NetApp Data ONTAP. In addition, its analytic tools provide a single, extremely versatile interface for meeting most any monitoring requirement as opposed to the multiplicity of software tools available and required for NetApp.

Table 4: Summary of Results

Task Areas	Time	Steps
Total Baseline Provisioning		
Oracle Advantage (Oracle - NetApp)	-15	-1
Oracle Advantage % (Oracle/NetApp-1)	-18%	-33%
Total Configure Usable Storage		
Oracle Advantage (Oracle - NetApp)	-56	-2
Oracle Advantage % (Oracle/NetApp-1)	-45%	-40%
Total Provisioning		
Oracle Advantage (Oracle - NetApp)	-71	-3
Oracle Advantage % (Oracle/NetApp-1)	-34%	-38%
Total Data Protection		
Oracle Advantage (Oracle - NetApp)	1	0
Oracle Advantage % (Oracle/NetApp-1)	5%	0%
Total Maintenance & Configuration		
Oracle Advantage (Oracle - NetApp)	-63.942	-1
Oracle Advantage % (Oracle/NetApp-1)	-31%	-11%
Total Monitoring		
Oracle Advantage (Oracle - NetApp)	-111	-12
Oracle Advantage % (Oracle/NetApp-1)	-44%	-50%
Total		
Oracle Advantage (Oracle - NetApp)	-244.942	-16
Oracle Advantage % (Oracle/NetApp-1)	-36%	-38%



Details

The results below are presented for the tasks evaluated, organized by the four task categories previously described: Provisioning, Data Protection, Maintenance and Configuration Changes, and Monitoring. Each section consists of a description of the task or tasks evaluated, a table showing the time (in seconds), the number of steps required for performing the task, and a brief analysis of the results.

Provisioning

Provisioning begins after the storage appliances have been installed in the rack and connected to the networks and power sources, and the management software has been initialized. Edison did not compare these tasks because they are usually performed only once in the life of the system. For this study, Provisioning has been divided into two categories: Baseline Provisioning are those tasks that are prerequisite for other tasks. Creating Usable Storage consists of the creation of LUNs and their attachment or presentation to hosts.

There are two sub categories of Baseline Provisioning tasks. Preparing Connectivity consists of identifying the appliance interfaces to the network and identifying initiators (the hosts) for presentation. Both the appliance interfaces and initiators are organized into groups for convenience when performing other configuration tasks. The other sub category is Preparing Storage. This task results in the virtualization of the physical disks in the appliance so that the combined storage capacity can be provisioned for use by the hosts.

Storage preparation on a NetApp filer requires the identification and selection, through a Wizard of the disks to be combined into an aggregate. On the Oracle Appliance, the disks are configured in pools. The options for the two platforms are similar, as the underlying hardware is essentially the same: hard disks on shelves.

As can be seen in Table 5, Page 21, Edison's testing shows that Baseline Provisioning on an Oracle array requires 28 percent less time and is 20 percent less complex than is required by the NetApp filer.



Table 5: Baseline Provisioning

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Provisioning				
Preparing Storage				
Create Storage Pool/Aggregates	43	1	59	1
Preparing Connectivity				
Create Targets	16	1	17	1
Create Target Groups	30	1	59	1
Create Initiator Groups	25	1	24	2
Total Baseline Provisioning	114	4	159	5
Oracle Advantage (Oracle - NetApp) in seconds	-45	-1		
Percent Oracle Advantage (Oracle/NetApp-1) Oracle Advantage (Oracle/NetApp-1)	-28%	-20%		

Configuring Usable Storage is the creation of LUNS and their presentation to hosts. This is an area where there are key differences between the two platforms. In order to create a LUN NetApp requires setting aside actual physical capacity in the form of Volumes with defined parameters. Only then, can capacity be set aside as a LUN and attached to a host. Though the theoretical limit is high (500 per node, 1,000 per cluster), the number of possible volumes on the array is limited by the storage capacity of the array.

In order to create a LUN on the Oracle appliance, the administrator first creates a Project. A Project is actually a set of policies that can be inherited by any LUNs (or Filesystems for NAS hosts) created within the project. These properties cover a range of possible settings, such as LUN size, the enabling of data deduplication and compression, and so forth. A Project does not utilize any appreciable storage capacity — in our testing, a project consumed about 50 KB, which on a 20 TB appliance is insignificant.

Once a NetApp filer volume has been created, LUNs can be created as needed using the LUN Wizard. This task requires knowing the Path to the LUN (either by having written it down or by going outside the Wizard and referring to the Manage Volumes page in ONTAP), entering the size and LUN protocol (operating system to support), disabling space reservation (not recommended), knowing what LUN IDs are available, and



entering a new ID. On the Oracle appliance, creating a LUN can be as simple as clicking the Add LUN button, entering a name for the LUN, and accepting the defaults previously set for the parent project.

As can be seen, creating several LUNs using the BUI is a much simpler process on an Oracle appliance than on a NetApp appliance.

Table 6: Provisioning Usable Storage

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Provisioning				
Configuring Usable Storage				
Project/Volume	22	1	44	2
LUN	22	1	45	1
Map/Add LUN	24	1	35	2
Total Configure Usable Storage	68	3	124	5
Oracle Advantage (Oracle - NetApp) in seconds	-56	-2		
Percent Oracle Advantage (Oracle/NetApp-1)	-45%	-40%		

Table 6: Provisioning Usable Storage shows that the Oracle appliance requires 45 percent less time and is 40 percent less complex than is required for NetApp.

Table 7: Total Provisioning

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Total Provisioning	182	7	283	10
Oracle Advantage (Oracle - NetApp) in seconds	-101	-3		
Percent Oracle Advantage (Oracle/NetApp-1)	-36%	-30%		



Overall, the Sun ZFS Storage Appliance required 56 percent less time and was 30 percent less complex to manage than the NetApp Filer. These results are especially significant in organizations where administrators are frequently adding LUNS for a rapidly expanding pool of virtual servers.

Data Protection

Data protection on a Storage Appliance includes such features as multiple controllers, multiple I/O paths, RAI, remote replication, and the creation of snapshots and clones. Multiple controllers and I/O paths are dependent upon the physical attributes of the appliance and are generally configured as part of installation. For this study, RAID settings are included in the creation of aggregates or pools. Replication was not tested because the Edison lab only had one appliance from each vendor. That leaves the creation of clones and snapshots.

A clone is defined the same way on both platforms — a writable copy of existing data on the appliance managed in the same manner as the source data — but is architecturally very different on the two platforms.⁹ Because of these differences, Edison did not compare clone creation in this study.

A snapshot is defined as a read-only copy of data on the appliance. On NetApp, snapshots are created at the volume level. On Oracle, snapshots can be created at the Project, Filesystem, or LUN level. On both platforms, snapshots require only minimal storage capacity at creation, only utilizing storage space when they are used as a data source presented to a host.

On the NetApp filer, snapshots are enabled and scheduled by default; therefore, manually creating a snapshot is a very straightforward process, merely requiring navigation to the Add Snapshot page (under Volumes), selecting the volume upon which the snapshot will be made, giving the snapshot a name, and clicking Add.

On the Sun ZFS Storage Appliance, the administrator must navigate to the Project upon which the snapshot will be taken before a snapshot can be added. Once navigation is complete, the Add Snapshot button is clicked, and a snapshot name is assigned applied by clicking Apply.

⁹ Called a FlexClone Volume on a NetApp filer, a FlexClone is the same size as its source volume and can only be created within the same aggregate as within the parent aggregate. On an Oracle Storage Appliance, a clone is created from a snapshot made at the Project, Filesystem, or LUN level. Like a snapshot of a Filesystem, a clone does not consume any additional space until the data within it changes.



Table 8: Data Protection - Snapshots

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Create and Schedule Snapshots	21	1	20	1
Total Data Protection Replication	21	1	20	1
Oracle Advantage (Oracle - NetApp)	1	0		
Percent Oracle Advantage (Oracle/NetApp-1)	5%	0%		

As the results in Table 8 on Page 24 show, creating a snapshot using either appliance is a simple and quick process. Though the process is five percent faster on NetApp, the actual time difference is insignificant.

Maintenance and Configuration Changes

The Maintenance and Configuration Changes tasks evaluated in this study consist of such activities as deleting outdated snapshots, managing and expanding LUNs, destroying inactive LUNs, managing snapshot schedules, and reconfiguring or decommissioning storage pools or aggregates. Other tasks that fit under the topic but that were not evaluated include removing and replacing failed drives, adding additional drives, and so forth.

Both vendors provide access to these maintenance tasks through the same interfaces that are used for configuration. The Sun ZFS Storage Appliance adds a Maintenance tab, the features of which were not evaluated for this study. The Maintenance tab provides access to hardware monitors for the controllers and drive shelves, a System view that provides access to installing software updates or support bundles as well as backing up and restoring configuration files. The Maintenance tab also includes a tab for monitoring persistent problems on the system that require replacement or software updates. System logs and workflows are also available under the Maintenance tab. NetApp FilerView does not offer a directly comparable feature, though similar information is available under the Filer menu on the navigation bar or under the Network and Storage menus.

As can be seen in Table 9, Page 25, both vendors have made these tasks simple to perform, requiring only a single step for each task. The need to either utilize a Wizard or perform such activities as taking a volume offline, even if only requiring a single button click and confirmation before deleting that volume, makes these tasks take 35 percent longer on NetApp than on Oracle.



Table 9: Maintenance & Change Configuration

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Maintenance and Change Configuration				
Unconfig Pools or Aggregates	20	1	30	1
Destroy Project or Volume	10	1	25	1
Expand LUN	20	1	30	1
Destroy LUNs	16	1	24	1
Destroy Snapshots	16	1	30	1
Manage Snapshot Schedules	24	1	18	1
Manage LUN properties (Additional Replication, Thin Provisioned & Data Compression)	15	1	30	1
Total Configure Usable Storage	121	7	187	7
Oracle Advantage (Oracle - NetApp) in seconds	-66	0		
Percent Oracle Advantage (Oracle/NetApp-1)	-35%	0%		

Monitoring

In an ideal world, storage administrators spend their time monitoring stable systems, occasionally adding a shelf of drives and creating new LUNs. In the real world, storage administrators also have to deal with problems. Storage system problems include the obvious — failed drives, network adapters, and so forth — as well as the less obvious issues of application performance and other uses of the data being stored.

Edison has come to believe that the quality of monitoring and troubleshooting tools provided by IT solutions can be a critical criterion for selecting the solution. This is especially true for those solutions targeting the SME and midrange business market where it is unlikely that enterprise-wide management solutions are in place. Because of this, monitoring and troubleshooting tasks were given special attention in this study. Before we look at the results of Edison’s evaluation of monitoring and troubleshooting tasks, we have provided an overview of the tools that NetApp and Oracle provide for the purpose.

In keeping with the evolution of the industry and the NetApp product line, the company provides many different monitoring tools. The table below (Table 10, Page 26) lists just a handful of the tools that Edison investigated while determining what to monitor and troubleshoot. In addition to these BUI and command line tools, NetApp also offers, through separate license, a DataFabric Manager (DFM) server and Management Console, which provides administrators with a set of tools for performing such tasks as data backup protection, space management, resource provisioning, data migration and performance tuning on multiple NetApp systems from a single interface.¹⁰

Table 10: NetApp - Sample listing of monitoring tools

Task Areas	Tool Used
CPU	
CPU Percent Utilization	Operations Manager
Disks	
Disk: I/O bytes	Operations Manager
Disk: I/O operations	Operations Manager
Disk: WAFL logical I/O bytes	Sys stat
Disk: WAFL logical I/O operations	Sys stat
Network	
Network: Device bytes	Operations Manager
Protocol	
Protocol: SMB operations	Perfmon
Protocol: Fibre Channel bytes	Operations Manager
Protocol: iSCSI bytes	Operations Manager
Protocol: NFSv2 operations	Operations Manager
Protocol: NFSv3 operations	Operations Manager
Misc	
Cache	Sys Stat

For the Sun ZFS Storage Appliance, Oracle offers essentially the same range of monitoring and troubleshooting capabilities as NetApp from a single source: DTrace. The table (Table 11, Page 27) shows a listing of many of the possible monitoring counters available, all powered by DTrace and accessible on the management console Analytics tab.

¹⁰ The NetApp DataFabric Manager and its applications - Protection Manager and Provisioning Manager – were not evaluated for this study. While utilizing DFM would have enabled Edison to evaluate additional NetApp monitoring and troubleshooting features, adding an additional set of tools would have further disadvantaged NetApp within the parameters of the study.

Table 11: Oracle - Sample listing of monitoring tools

Task Areas	Tool Used
CPU	
Percent utilization by CPU mode	DTrace
Percent utilization by CPU identifier	DTrace
Percent utilization by CPU Process identifier	DTrace
Cache	
Cache: ARC accesses	DTrace
Cache: L2ARC I/O bytes	DTrace
Cache: L2ARC accesses	DTrace
Cache: ARC evicted bytes	DTrace
Cache: ARC size	DTrace
Cache: ARC target size	DTrace
Cache: DNLC accesses	DTrace
Cache: DNLC entries	DTrace
Cache: L2ARC errors	DTrace
Cache: L2ARC size	DTrace
Disks	
Disk: Disks	DTrace
Disk: I/O bytes	DTrace
Disk: I/O operations	DTrace
Disk: ZFS logical I/O bytes	DTrace
Disk: ZFS logical I/O operations	DTrace
Network	
Network: Device bytes	DTrace
Network: Interface bytes	DTrace
Protocol	
Protocol: SMB operations	DTrace
Protocol: Fibre Channel bytes	DTrace
Protocol: Fibre Channel operations	DTrace
Protocol: FTP bytes	DTrace
Protocol: HTTP/WebDAV requests	DTrace
Protocol: iSCSI bytes	DTrace
Protocol: iSCSI operations	DTrace
Protocol: NFSv2 operations	DTrace
Protocol: NFSv3 operations	DTrace
Protocol: NFSv4 operations	DTrace
Protocol: SFTP bytes	DTrace

Task Areas	Tool Used
Protocol: SRP bytes	DTrace
Protocol: SRP operations	DTrace
Misc	
Memory: Dynamic memory usage	DTrace
System: NSCD backend requests	DTrace
System: NSCD operations	DTrace

Monitoring with NetApp

NetApp provides two main BUI based applications for monitoring and troubleshooting the status of a storage appliance. The first is the simple Filer At-A-Glance, which is accessible from the Data ONTAP home page and from the FilerView Real Time Status menu (See Figure 6, page 28) Filer At-A-Glance provides five simple graphs for viewing capacity usage by volume, file ops/second, CPU utilization and network input and output in KB/sec. By clicking the capacity usage graph an administrator can drill down to view the usage by filesystem or LUN. (Figure 7, Page 29).

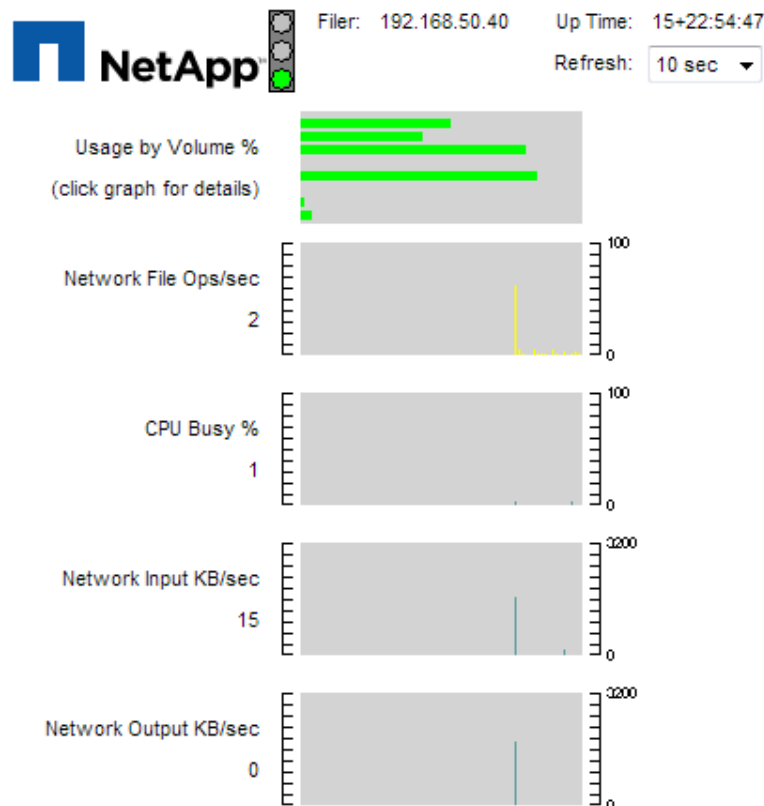


Figure 6: NetApp Filer-at-a-Glance



Filer: 192.168.50.40











Filesystem	kbytes	used	avail	%used	
/vol/EGiSCSI/	281245984	226842752	54403232	81%	
/vol/EGiSCSI/.snapshot	70311496	91351996	0	130%	
/vol/Test_Vol1/	83886080	21014344	62871736	25%	
/vol/Test_Vol1/.snapshot	20971520	1088	20970432	0%	
/vol/Lab_TestClone/	83886080	3808	83882272	0%	
/vol/Lab_TestClone/.snapshot	20971520	1072	20970448	0%	
/vol/Sun_TestVol1/	41943040	31527140	10415900	75%	
/vol/Sun_TestVol1/.snapshot	10485760	924	10484836	0%	
/vol/SunFlexClone/	41943040	1952	41941088	0%	
/vol/SunFlexClone/.snapshot	10485760	776	10484984	0%	

Figure 7: NetApp - Filesystems Status

The FilerView Real Time Status menu provides access to two additional monitoring tools:

- Health Monitor: provides a simple browser pop up with a colored icon (Green, Yellow, Red) and a text status message. This simple tool can be kept open on an administrator workstation to provide instant indication of storage array status.
- Performance Meter: provides another pop up with an Ops/Sec chart plus a traffic light icon to indicate status at a glance.

The FilerView home page defaults to a System Status view that includes basic configuration information and as well as the same status indicator as provided by the Health Monitor.

These simple tools provide adequate alerts to system status, but they provide no access, beyond the capacity usage drill-down mentioned above, to more detailed system status data. In order to obtain this data, NetApp provides Operations Manager, the second BUI based application. NetApp Operations manager is separately licensed and must be installed on a server (Linux, Solaris, or Windows). For reasons that should be obvious, the Operations Manager database should not be stored on a filer being monitored. NetApp Operations Manager can monitor multiple NetApp filers, providing a centralized platform for managing all of an organization's NetApp storage systems.

In Edison's view, NetApp Operations Manager is an excellent, fully featured monitoring and reporting platform; with emphasis on the reporting. Possibly, because it is designed to access multiple NetApp devices and must do so over a network, Operations Manager does not show live data. Rather, it polls the devices at scheduled intervals. (Polling schedules can be set for the many attributes being monitored individually.)¹¹

The Operations Manager home page (Figure 8, Page 30) shows a summary of the status of the NetApp devices being managed. It includes a line graph with data on various, user-selectable, monitors. Historic data shown on this chart can be viewed by the day, week, month, or quarter. In addition, the user can select from several Report views.

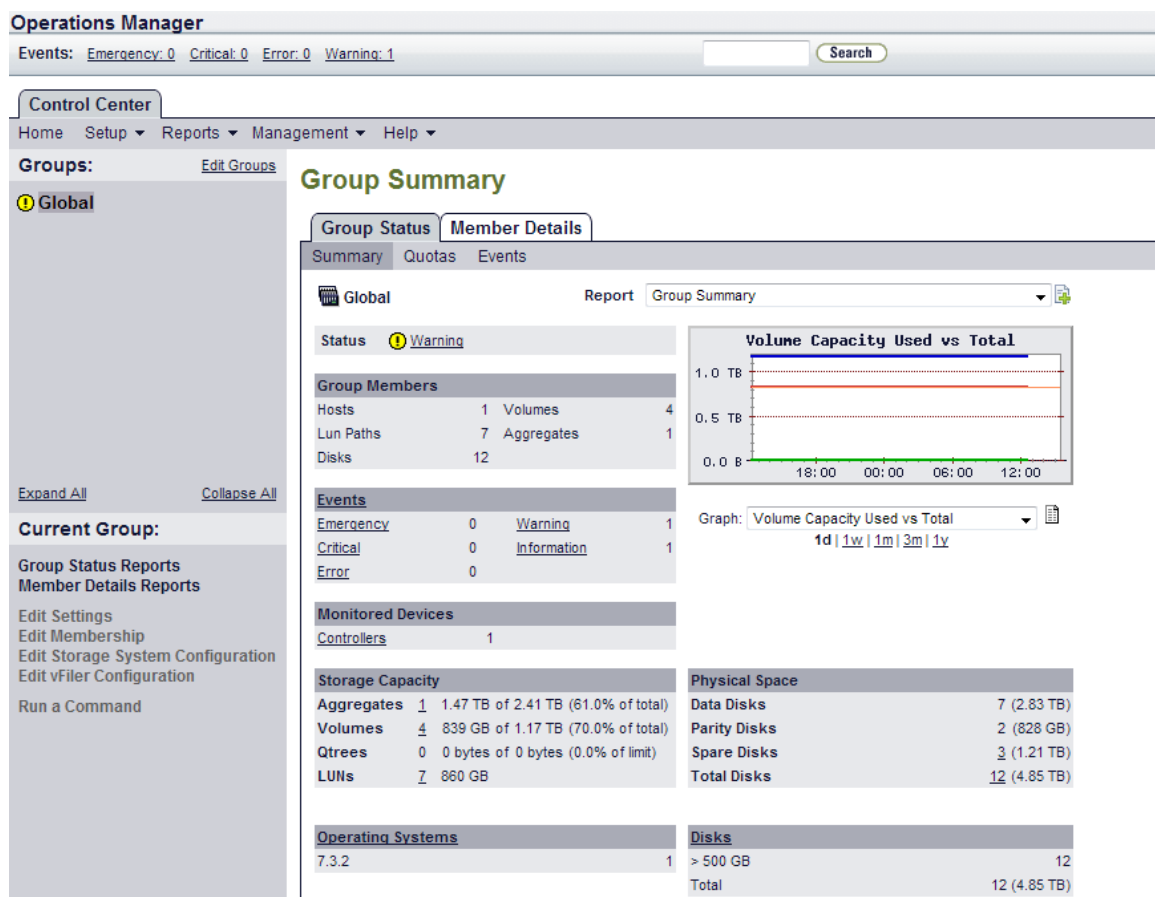


Figure 8: NetApp Operations Manager - Home Page

¹¹ Changing polling schedules requires accessing the Options page and navigating to the Monitoring page. There, the polling intervals can be set. Settings are entered manually: 5 minutes, 1 day, etc.

Beyond the home page, the most important area in Operations Manager is Reports. As can be seen in Figure 9, Page 31, a very large number of standard report categories exist, each with numerous reports from which to choose. Many of these reports are text reports, often with links for drilling down to access additional details. For example, clicking on a Storage Server from the table in NetApp Operations Manager - Volume Space Guarantees Report, Page 32, brings you to the home page view.

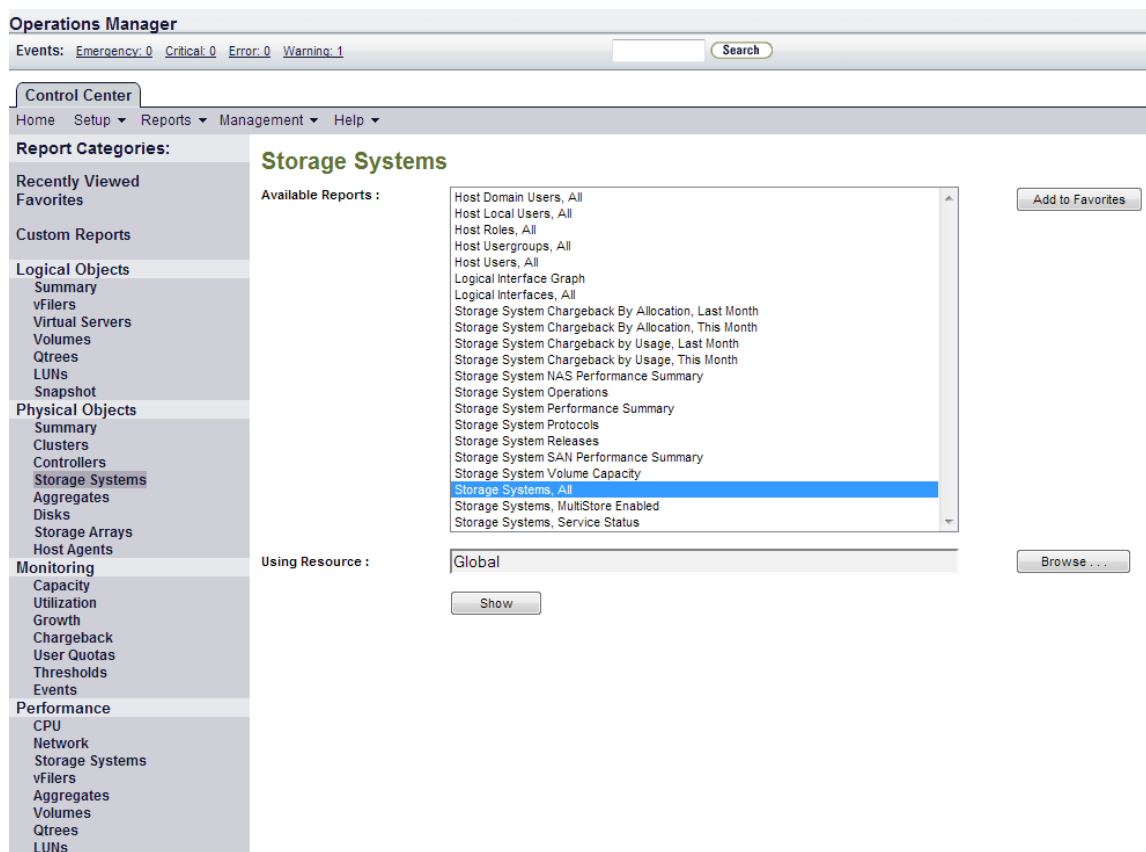


Figure 9: NetApp Operations Manager - Storage Systems Reports

Operations Manager can be a very useful tool, providing administrators with valuable insight into historic usage patterns and the state of their NetApp devices. What Operations Manager cannot do is provide immediate feedback on changes made to a device during troubleshooting. As will be seen in Edison's testing results, this is a far from efficient approach to monitoring and troubleshooting.

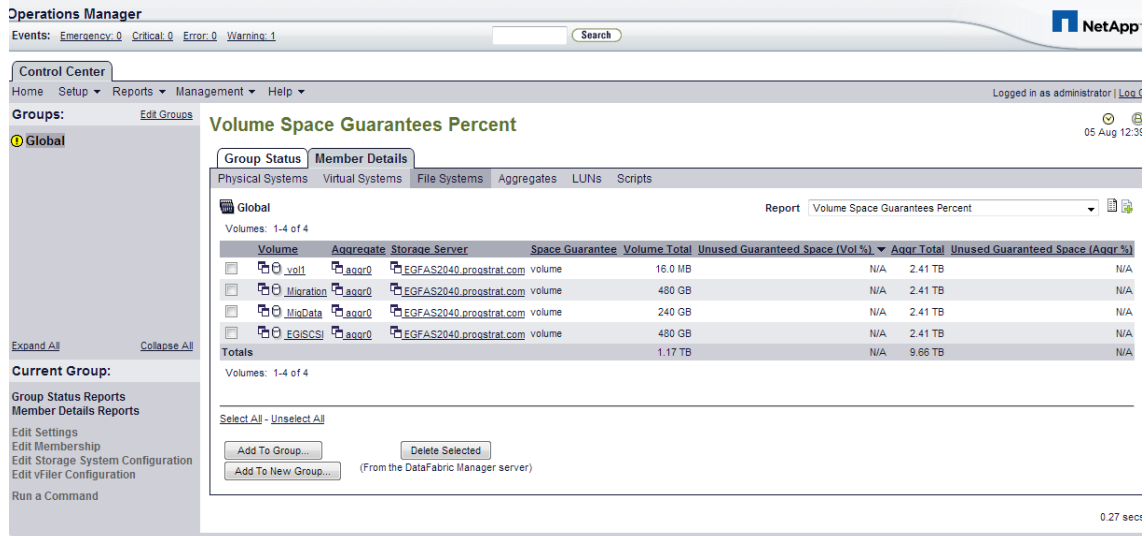


Figure 10: NetApp Operations Manager - Volume Space Guarantees Report

NetApp provides additional troubleshooting support on its Support website. Besides the usual knowledge base, “create a ticket,” and call support offerings, the NetApp support site includes AutoSupport. NetApp AutoSupport is where data is sent from a filer to a NetApp server from which NetApp engineers can provide technical support on demand (depending upon support contracts, of course) and reports on the system and its components (see Figure 11, page 32).

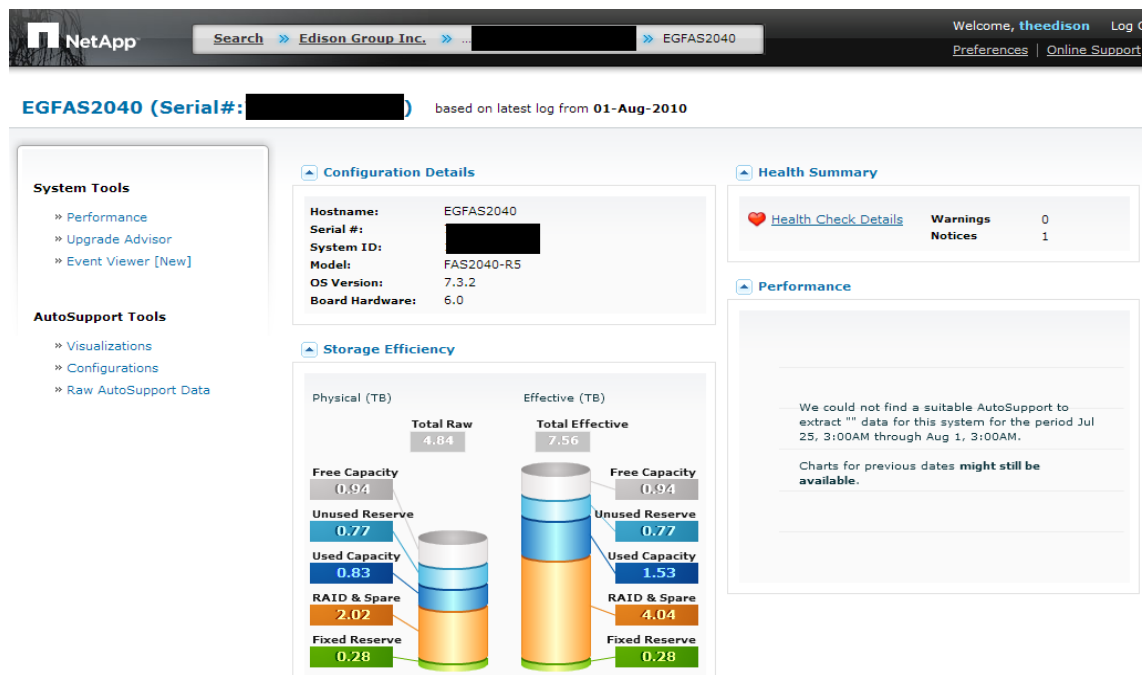


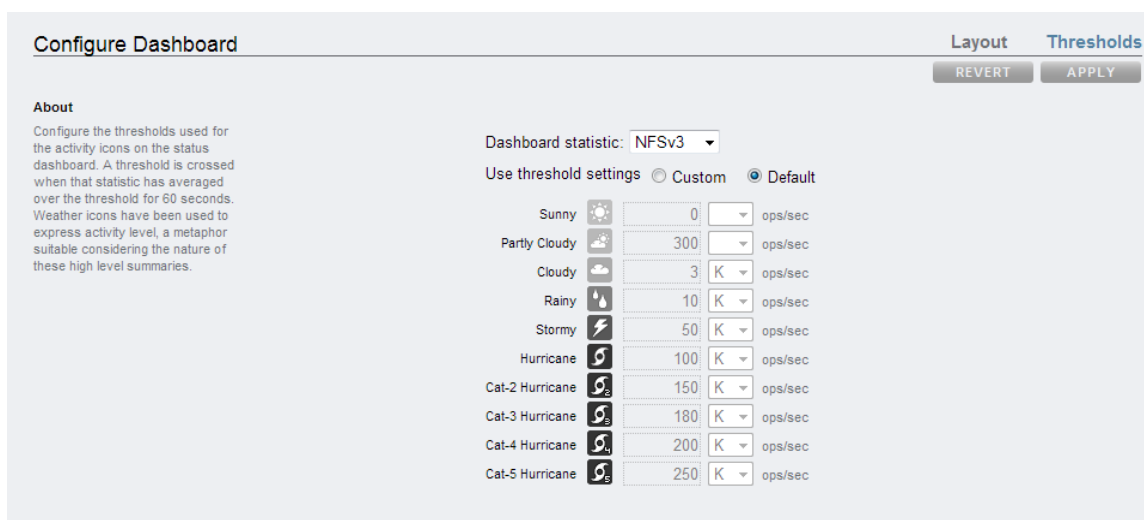
Figure 11: NetApp AutoSupport Dashboard

Monitoring with Oracle

Oracle takes a very different approach from NetApp in respect to monitoring and troubleshooting. DTrace, Solaris' built-in, dynamic tracing tool, is tightly integrated into the management console. DTrace functionality is visible on the default Status Dashboard Management home page (Figure 3, Page 16), is used on several of the configuration settings pages, and is the engine for the Analytics tab.

The Status Dashboard provides administrators with an overview of almost all aspects of appliance status. It contains live histograms of CPU, Network Disk, and five other key performance metrics (user-selectable from a list of fourteen options). It also provides pie chart views of storage and memory usage, a listing of services and their status (enabled or disabled plus health) and the status of key hardware components. There is also a listing of the most recent alerts.

The histograms can each have an icon that can indicate health at a glance. Thresholds for each statistic can be set by the user. (See Figure 12, Page 33) Clicking the Storage Usage pie chart takes you to the Shares All Projects view for a detailed view of storage use. Clicking the Memory Usage Pie Chart brings you to an Analytics Worksheet that shows a live view of dynamic memory usage broken down by application name. (Figure 13, Page 34) Clicking a Service brings you to the configuration page for that service. Clicking on Hardware icons bring you to the Maintenance Tab for Hardware from which you can drill down to check the live status of hardware components.



Weather Icon	Threshold	Unit
Sunny	0	ops/sec
Partly Cloudy	300	ops/sec
Cloudy	3	K ops/sec
Rainy	10	K ops/sec
Stormy	50	K ops/sec
Hurricane	100	ops/sec
Cat-2 Hurricane	150	K ops/sec
Cat-3 Hurricane	180	K ops/sec
Cat-4 Hurricane	200	K ops/sec
Cat-5 Hurricane	250	K ops/sec

Figure 12: Oracle Status Dashboard - Thresholds

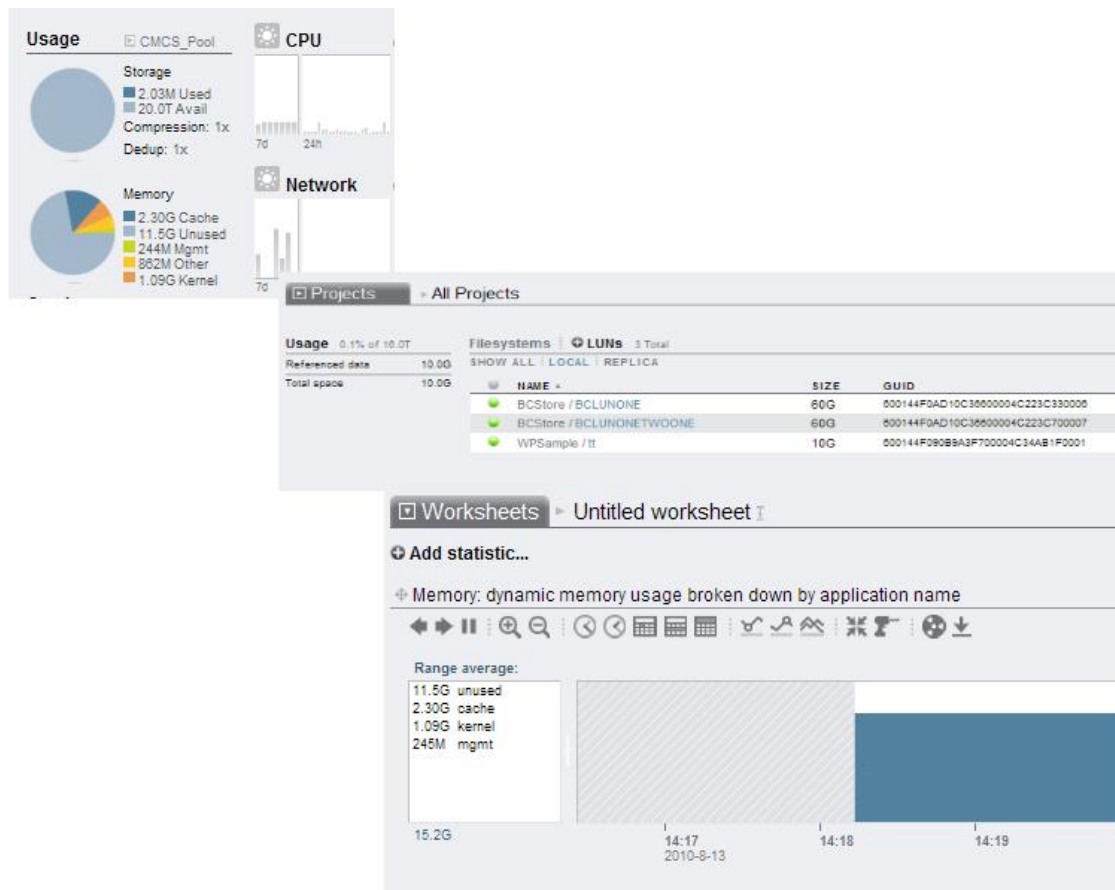


Figure 13: Oracle Status Dashboard - Usage pie graphs with drill down graphs

The Analytics Tab is the heart of monitoring and troubleshooting on the Sun ZFS Storage Appliance. Over thirty datasets of performance or operational data, cached and on disk, are constantly monitored by default. Datasets are automatically created when statistics are viewed, but are not saved to disk for future viewing unless they are archived.¹²

To use DTrace Analytics an administrator adds statistics from the list of datasets and views the graphs on a worksheet. Worksheets can be saved for future use or cloned to simplify the creation of variations on a theme. In addition to being able to view live and historic data, the dataset can also be exported for analysis or reporting using external software.

¹² Archiving sets a statistic to be permanently opened and archived to disk creating a permanent dataset for review.

DTrace Analytics works on the principle of drill-down analysis. Starting at a high level, an administrator can drill down to focus on finer details of a statistic. For example, an administrator can choose to monitor Disk I/O in bytes per second. By clicking the Drill Icon in the graph's tool bar the administrator can view I/O bytes per second broken down by type of operation or by disk. Selecting one of the operations, such as *write*, the administrator can drill down to view I/O bytes per second of type "write," broken down by disk. (Figure 14, Page 35)

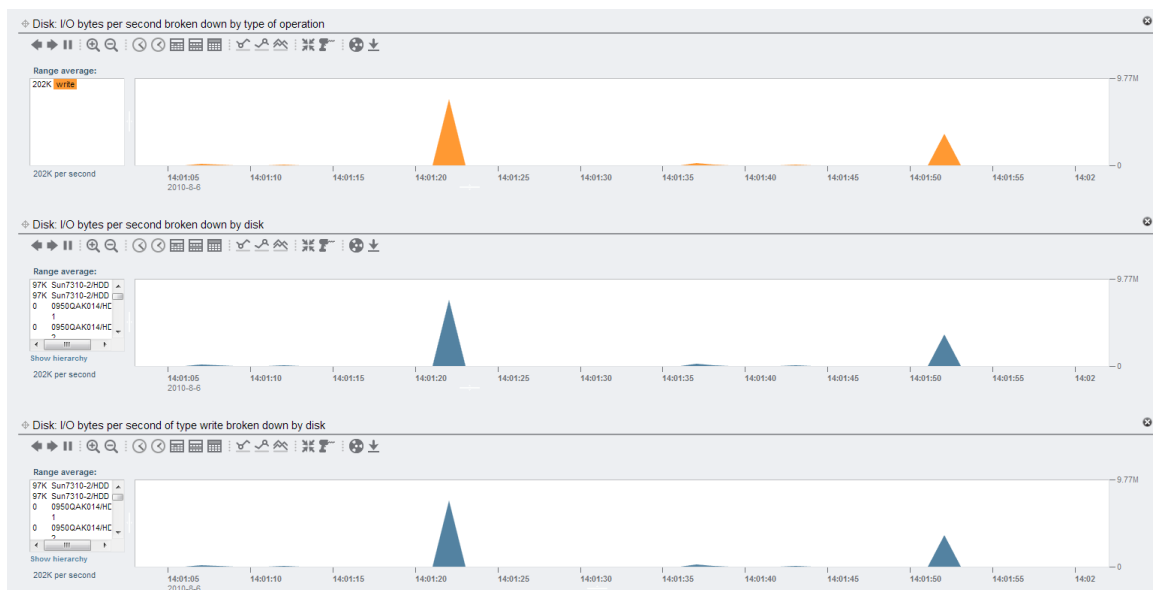


Figure 14: Oracle Integrated Software Suite: Disk I/O Drill Downs

The Sun ZFS Storage Appliance has a feature similar to NetApp's AutoSupport called Phone Home. Phone Home sends fault reports that can generate support cases. It also sends daily heartbeat messages to Oracle which can trigger notifications to the designated support contact (via e-mail, etc.) if the system fails to send a heartbeat for too long.

System configuration messages are also sent describing current hardware and software configuration data, used for statistical purposes. Not only does Oracle utilize these data for support, but also to identify hardware and software that may need to be updated, sending software updates, replacement hardware, and even field technicians when problems are identified.¹³

¹³ These support responses are dependent upon the nature and status of support contracts.



The Phone Home facility links to the Oracle/Sun support center, SunSolve, which provides access to the usual knowledge base and other support facilities. Included is an array of diagnostic and management tools that can streamline various operations activities.

Performing Monitoring Tasks

In its evaluation of monitoring and troubleshooting, Edison separated monitoring tasks from troubleshooting and only tracked the selection of a monitoring tool and the viewing of it. This was done because the actual performance of troubleshooting tasks could include:

1. Tweaking a configuration setting from the management console.
2. Checking and changing cables, interface cards and other system components.
3. Checking and changing switch, host, or even application settings.

Other than for the first option, performing these tasks not only does not involve the management consoles (the focus of this study) but are also essentially the same regardless of the storage appliance.

What the results table (**Table 12**, Page 38) does not show is how the success or failure of the troubleshooting effort is shown to the administrator. This is where the difference in approach between NetApp and Oracle is most clearly revealed. As an example, an application administrator has discovered that one of the LUNs being used to store application data is responding slowly.

In order to determine the cause and troubleshoot the problem, the NetApp administrator does the following:

- Login to NetApp FilerView
- Login to NetApp Operations Manager
- Verify status in FilerView by clicking to Filer at a Glance, or
- Verify filer health status in Operations Manager, drilling down through the Physical Systems page of Storage Controller Details
- Run LUN and other system reports to identify probable cause of the problem
- Perform the appropriate troubleshooting tasks for each probable cause
 - Validate the status of each system component in FilerView, checking configurations, etc.



- Check the status of any external devices to which the appliance is connected, including the cabling and switches.
- Verify troubleshooting success by refreshing monitoring data in Operations Manager and/or
- View status in FilerView
- Clear any alerts in the events report: alerts can be acknowledged, leaving them in the report for future reference, or deleted.

In order to determine the cause and troubleshoot the alert on the Oracle appliance, the Oracle storage administrator does the following:

- Visit BUI Dashboard home page and check charts and status indicators
- Drill down through the graph reflecting the issue to identify the source of the problem; initiator, target, project, LUN etc. Drilling down can continue from most of these filters.
 - If needed, check status of any external devices to which the appliance is connected, including the cabling and switches.
- For our purposes, these checks identify and fix the problem.
- View status on the Dashboard.

As can be seen, the performance of the complete troubleshooting activity can be much more complex and time-consuming on the NetApp filer than on the Sun ZFS Storage Appliance.

The data shows the results of Edison's testing for several categories of appliance component monitoring counters. The monitors selected for each category are listed by platform. The comparisons chosen do not come close to exhausting the monitoring options available on both platforms. The items chosen were those that were the most similar in function and information provided. As we have demonstrated above, accessing the monitors is only one aspect of the troubleshooting effort. What the results show is that not only is the overall troubleshooting process less complex on a Sun ZFS Storage Appliance, but also that the time required and the complexity of identifying, selecting, and viewing fresh monitor data is less, as well.



Table 12: Monitoring

Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Monitoring				
CPU				
CPU Percent Utilization (Oracle: CPU Raw Data)	12	1	15	2
Totals CPU	12	1	15	2
Disks				
NetApp: Physical Systems - Disks Throughput Blocks/Sec Oracle: Disk I/O operations per second by disk	12	1	15	2
NetApp: Physical Systems - Disks Read or Write Operations/Sec Oracle: Disk I/O Operations per second by disk	12	1	30	2
Totals Disks	24	2	45	4
Network				
NetApp: Controller Interface Graph Oracle: Interface Bytes/Second by interface	12	1	15	2
NetApp: Controller Interfaces in/out table Oracle: Device Bytes/Second by device	12	1	15	2
Totals Network	24	2	30	4
Protocols (iSCSI as example, other protocols follow same procedures.)				
NetApp: iSCSI Reads/Sec, Writes/Sec Operations/Sec by LUN Oracle: Choose from: iSCSI bytes or iSCSI operations by either: raw statistic, initiator, target, project, LUN, client, type of operation, command, latency, offset, size, client	12	1	15	2
Totals Protocols	12	1	15	2



Task Areas	Oracle		NetApp	
	Time in Seconds	Steps	Time in Seconds	Steps
Total Monitoring	72	6	105	12
Oracle Advantage (Oracle - NetApp) in seconds	-33	-6		
Percent Oracle Advantage (Oracle/NetApp-1)	-31%	-0.5		



Conclusions

Edison subject matter experts, using Edison's CMCS methodology, compared the management software provided with unified storage arrays from Oracle and NetApp. Tests were performed using NetApp Data ONTAP software running on a NetApp FAS2040 Filer and Oracle Integrated Software Suite on an Oracle 7310 ZFS Appliance.

The results of our testing show that managing the Sun ZFS Storage Appliance required 36 percent less time and was 38 percent less complex. This translates into lower operating expenses due to saved time, quicker provisioning of new applications (which is especially important for virtualized servers), and quicker resolution of problems through superior monitoring and troubleshooting tools.

Based upon these findings, Edison recommends that organizations looking for versatile, highly scalable, high performance storage solutions consider the Sun ZFS Storage Appliance as their standard unified storage platform.



Appendix

The Methodology Defined

Edison Comparative Manageability Cost Studies methodology is a product manageability cost evaluation process, whereby the products in question are compared against a set of task-oriented objective and subjective metrics in order to derive an accurate set of analytical results. The outcome of this study determines the Comparative Management Cost (CMC) incurred by managing and operating the products in a production environment. The methodology employed to conduct this comparison consists of the following elements.

The Study

The study is the baseline checklist of standard administration tasks routinely performed, quantitatively and qualitatively compared in order to determine, on a task-by-task basis, which product is superior. This is measured primarily in terms of ease of administration and secondarily (for certain tasks only) in terms of system speed of execution — the wall clock time it takes for the system in question to complete a job once it has been submitted by an administrator. The function of this study is to apply a set of quantitative metrics, developed by Edison Group, to a list of tasks typically regarded as qualitative in nature, in order to derive a meaningful set of CMCS statistics that can reveal the real difference in management costs for the two products in question.

Tasks

Edison defines a task as a complete logical activity, composed of one or more steps, all of which effect a significant alteration on the state of the device or software program that accomplishes a specific work goal. Each task is measured for time and complexity. Time and complexity, as used in the study, are defined as follows:

Time

Defined as the amount of time it takes to perform a given task. For certain (asynchronous) tasks, when a job can be run in the background so that the administrator can use the time for accomplishing other tasks, time is measured strictly in terms of the time it takes the administrator to perform the steps to configure, initiate, and submit a given task.



For other (synchronous) tasks in the study that demand the administrator's full attention and prevent the accomplishment of other tasks (as in performing a hot recovery operation on a live database), time is measured to include both the time it takes for an administrator to configure/execute the task in question as well as the time it takes the system to complete the task. All time metrics are measured in wall clock time.

Complexity

Complexity is measured using a proprietary metric devised by Edison Group: it is the number of system-affecting steps it takes to complete a given task, where a step is defined as a task component that effects a change of state to the system under test.

Because not all steps have the same inherent complexity, each step is further broken down into increments to account for the difference. An increment is a decision point that the user must make to complete a step. Increments are technically defined as a part of a step that will have a measurable effect on the state or execution path of that step in the task process, but which in and of itself does not effect a change upon the underlying system state until the step being executed is complete. For example, selecting Basic vs. Advanced Install with the installation Wizard is an increment and not a step.

- Complexity is then measured in terms of number of steps, but taking into account the following factors:
- The number of increments it takes to complete each step.
- Whether or not instrumentation for a given step is GUI-based or requires the use of a command line/scripting interface.
- Whether or not the task requires a context switch between multiple interfaces in order to be completed. If a context switch exists, then additional steps will be added to the total step count for a given task.

The above factors affect the complexity calculation as follows:

The primary measure is steps. If a step has many increments, it is considered several steps. The metric allows each step five increments, and thereafter we add steps for each additional five increments rounded up. So if a step has between 0–5 increments, it remains unchanged; if it has between 6–10 increments, it is increased by one; between 11–15 increments, it is increased by two; and so on. We decided to do this because, while increments are secondary to steps in determining complexity, they do modify the relative complexity of a given step in the course of completing a task. In other words, steps with a low number of increments are simple, and steps with a high number of increments are complex.

The other modifiers (instrumentation and context switching) may occur very infrequently in the products under review, but are a significant enough factor that we needed to account for them in some meaningful way in order to generate a measure of complexity that accurately reflects our experience of using the products.

Regarding instrumentation, if an operation could be executed entirely within a GUI interface, then the complexity/step value for that task would remain unmodified. If, on the other hand, a step required the use of a command line interface, this would increase the step count. For a simple single-line command operation, the step count was increased by one, whereas if the operation required the user to write a script, the step value was increased by two or more, depending on how much work was required to write the script in question.

Lastly, we come to the matter of context switching. If a context switch was encountered during the course of completing a given task, then two or more steps were added to the step count for that task. The possible addition of more than two steps was allowed for as a judgment call on the part of the subject matter expert performing the task under consideration. The reason tasks containing context switches may be penalized is that we regard it as inherently more complex to understand the dependencies of relating and performing a single operation in two different environments, as opposed to performing a similarly complex task in a well-integrated environment where all the operations can be accomplished in one place.

The following is the complexity formula utilized throughout these studies.

Complexity Calculation Formula

Complexity is defined as the number of computed steps it takes to complete a given task. The formula Edison typically uses to compute complexity for each of the tasks in this study is as follows:

For every five increments contained in a step, we increase the step value by one. For example:

If a step has 0–5 increments, step value = step + 0,

If a step has 6–10 increments, step value = step + 1

If a step has 11–15 increments, step value = step + 2

... and so on.



The type of instrumentation offered to perform a given task modifies the task's complexity.

If a task can be performed completely with a GUI, then step count = step count + 0.

If a task requires the use of a command line interface, then step count is modified as follows:

If the command line operation consists of a single-line command, then step count = step count + 1.

If the command line operation requires writing a script, then step count = step count + 2 or more steps, depending on a subjective assessment of the complexity of the script.

If a task requires a context switch between different environments, then step count = step count + 2 or more steps, depending on a subjective assessment of the complexity of the context switch.