Technical Report

Infrastructure as a Service with Fujitsu ServerView Resource Orchestrator and NetApp Storage

Ernst Berberich, Fujitsu
Holger Kordowske, Fujitsu
Christian Winter, NetApp
Dr. Adolf Hohl, NetApp
June 2011 | TR-3934

ABSTRACT

This document highlights infrastructure-as-a-service (IaaS) use cases realized by the Fujitsu ServerView Resource Orchestrator (ROR) management software leveraging NetApp® storage. The focus of this document is on the storage integration and the enabled use cases.

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

This technical report outlines the interplay of Fujitsu ServerView ROR and NetApp storage in a cloud-based environment. It briefly explains fundamental concepts of Fujitsu ServerView ROR and how it can leverage unique properties of a NetApp storage infrastructure.

With ServerView ROR, Fujitsu offers resource orchestration software that enables IT organizations to deliver IT infrastructure resources faster while lowering the barriers for users to gain access to IT resources. Automation of provisioning processes helps to streamline infrastructure administration workflows with a positive impact on operational data center costs. A logical server (L-server) concept hides the complexity of the underlying physical infrastructure and provides users with an interface from which they can request server resources easily and efficiently. As one of the major prerequisites for operation in enterprise data centers, Fujitsu ServerView ROR offers operation security supporting role-based administration for multi-tenant environments. Moreover, users can choose from various high-availability options to make sure of business continuity, ranging from failover solutions for single physical and virtual servers to protection against complete site failures.

Fujitsu ServerView ROR supports a wide variety of use cases. The following section highlights the use cases that fall within the scope of this technical report.

2 USE CASES

- **Automated physical and virtual server provisioning.** The automated provisioning process includes storage and network resources.
- **Data center consolidation.** Fujitsu ServerView ROR provides resource pooling with different service qualities to allow mastering very different workloads within the same infrastructure.
- **Infrastructure blackboxing.** Fujitsu Server ROR abstracts the definition of servers by introducing an L-server concept. An L-server is a logical representation of a server, which can be either a physical or virtual server. The users have a simple and intuitive view to the infrastructure. The complexity of the underlying infrastructure is blackboxed and invisible to the user.
- **Role-based administration in multi-tenant environments.** IT infrastructure can be shared by different users. A multi-tenant concept allows sharing IT resources even if they have competing interests.
- **Built-in disaster recovery (DR) and failover.** Fujitsu ServerView ROR comes with built-in functionality to address failover scenarios and DR.

3 ARCHITECTURE

This section gives an overview of the major concepts of Fujitsu ServerView ROR. Each subsection provides a detailed discussion. Figure 1 provides an overview of the architecture.

Figure 1) Overview of resource pools, their assignments, and roles (graphic provided by Fujitsu).

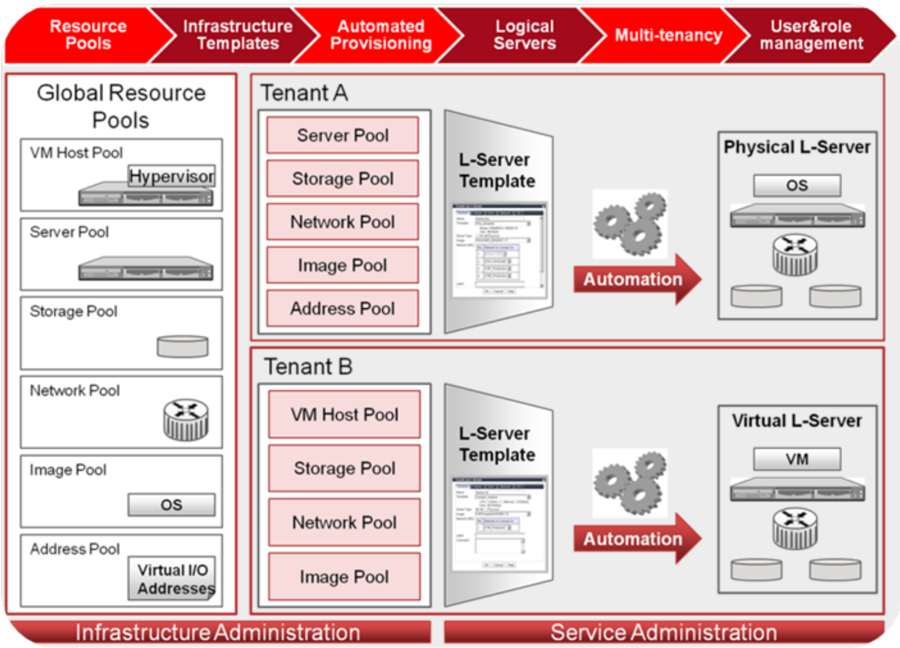

3.1 RESOURCE POOLS

A central element in Fujitsu ServerView ROR is the concept of resource pools. Pools group resources with similar attributes. Rather than consuming the resources of a specific infrastructure component, resources are consumed from pools. This simplifies the user experience because it hides physical borders of an infrastructure. Pool priorities define and control the order in which the resources are being used for provisioning.

Fujitsu ServerView ROR manages the following types of resource pools:

- **Server pool.** Bare-metal servers used to provision physical L-servers or spare servers
- **Virtual machine (VM) hosts pool.** Physical servers with installed hypervisor software to run virtual L-server
- **Network pool.** IP addresses and virtual LAN (VLAN) definitions, teaming
- **Address pool.** Virtual storage area network (SAN) and LAN addresses (World Wide Port Name [WWPN]/media access control [MAC])
- **Image pool.** Cloning master images and VM templates
- **Storage pool.** RAID groups, logical unit numbers (LUNs), virtual storage registered to VM management software

  Multiple pools with different characterizations and service levels can be defined to serve multiple service and application demands.

3.2 L-SERVER LIFECYCLE MANAGEMENT

Fujitsu ServerView ROR manages the complete lifecycle of L-servers. The activation of an L-server starts a provisioning process that automatically selects, allocates, and configures the requested computing, storage, and network resources from the available resource pools. For example, creation and mapping of LUNs, configuration of virtual I/O addresses and VLANs, OS image cloning, and personalization and the setup of a high-availability (HA) configuration are typical automated steps in a physical server provisioning process. If resources are no longer needed, they can be released to the resource pools for further use. Figure 2 illustrates sample L-servers that take their resources from pools.

Figure 2) Sample L-servers taking their IT resources from pools (graphic provided by Fujitsu).

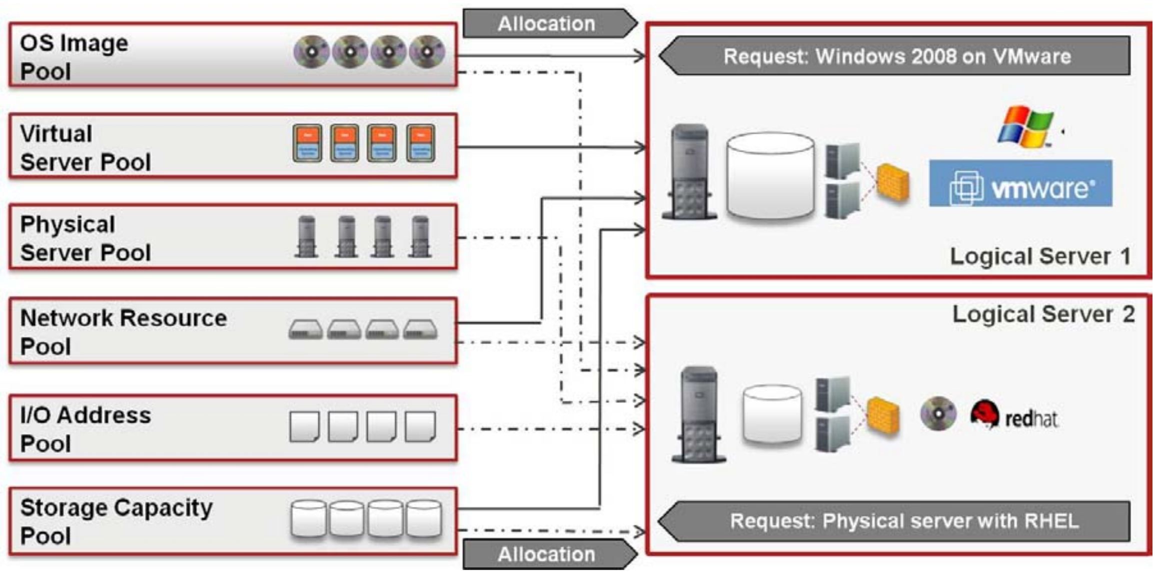

While it is possible to set all required parameters manually during an L-server definition, a template is the most efficient option to create an L-server. Templates are predefined and pretested L-server definitions, which are provided by infrastructure administrators in advance. Figure 3 is a screenshot of an L-server provisioned using a template.

Figure 3) Provisioning an L-server using templates (graphic provided by Fujitsu).

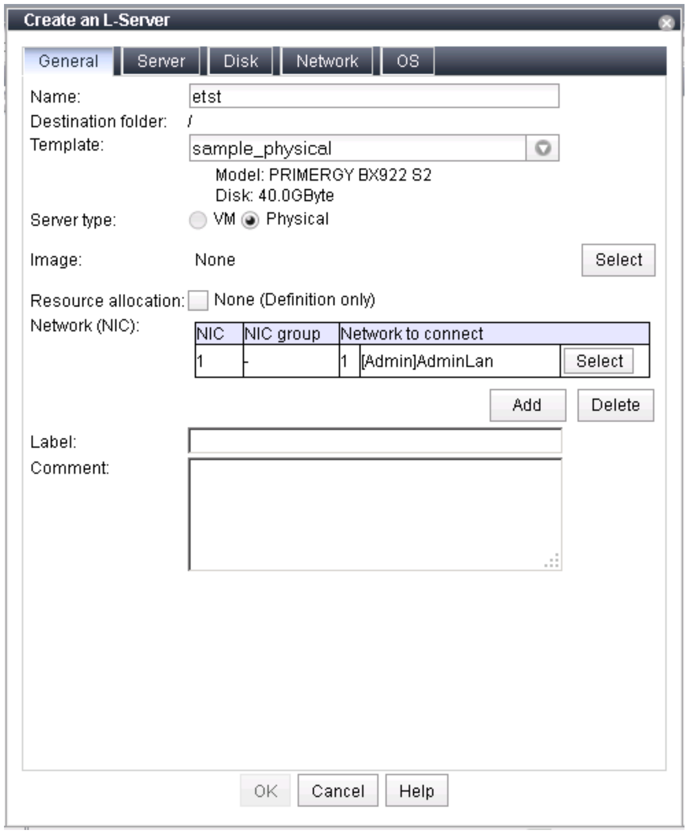

3.3 SEPARATION OF INFRASTRUCTURE AND SERVICE ADMINISTRATION

Resource abstraction based on L-servers enables IT organizations to separate service and infrastructure administrator roles. Service administration can now concentrate on service management. For example, service administrators no longer need to be involved in tasks such as the roll-out, operation, and maintenance of applications. They simply request infrastructure resources according to their service needs.

Also, infrastructure administrators for servers, network, and storage can now concentrate on delivering resources into the respective resource pools. Each expert can fully leverage his or her skills. Figure 4 shows the difference between infrastructure administration and service administration roles.

Figure 4) Infrastructure administration roles versus service administration roles (graphic provided by Fujitsu).

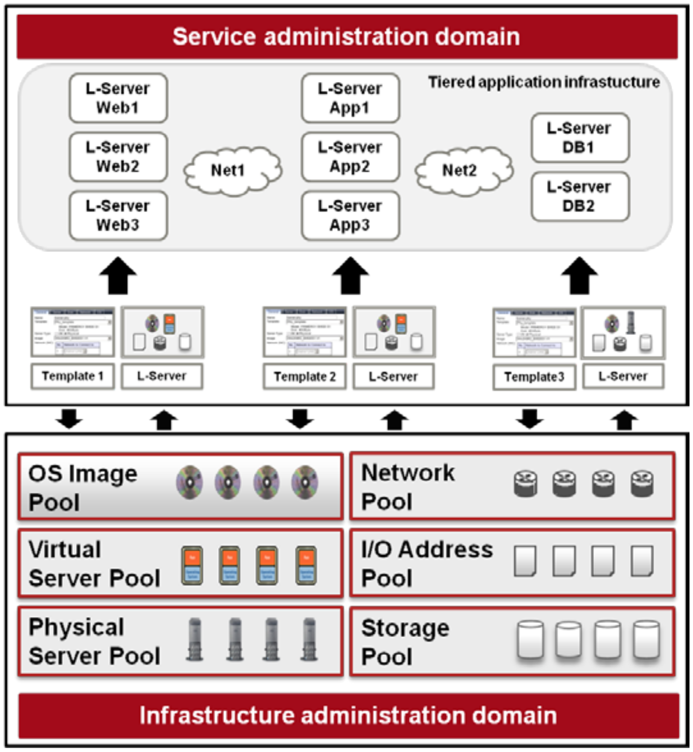

3.4 BUILT-IN HIGH AVAILABILITY

Fujitsu ServerView ROR has various HA options to facilitate business continuity, ranging from failover solutions for single physical or virtual servers up to support for complete blade chassis or storage system failover. DR in case of complete site failures is integrated.

The use of NetApp MetroCluster™ storage controllers complements these scenarios from a storage perspective. MetroCluster is a solution that offers combined array-based clustering with synchronous mirroring to deliver continuous availability and zero data loss, whether deployed in a data center, a building or campus, or across a metropolitan area. MetroCluster provides automatic transparent failover for all single-component failures. For complete site failures, an automated single command failover can be integrated in Fujitsu ServerView ROR to initiate the sequence.

Figure 5 shows a MetroCluster configuration spanning two sites.

Figure 5) NetApp MetroCluster spanning two sites.

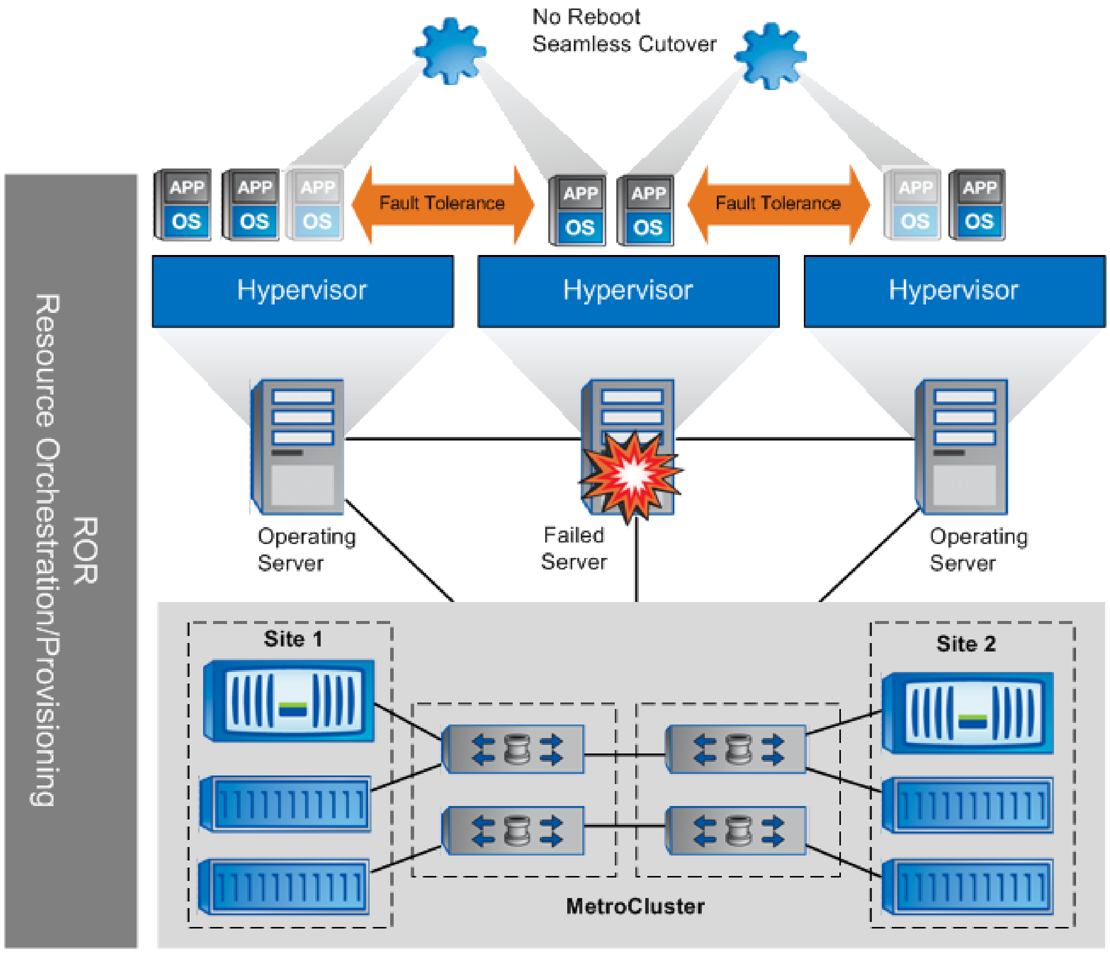

4 STORAGE INTEGRATION

This section describes how NetApp controllers are integrated with a Fujitsu ServerView ROR environment. First, outline the necessary steps to discover the resources of the NetApp storage controller within Fujitsu ServerView ROR. Second, outline how to define separate service classes of storage within a NetApp storage controller.

From an interconnectivity point of view, the storage integration of NetApp storage controller with Fujitsu ServerView ROR is built on Fibre Channel, which allows the NetApp storage to expose LUNs to the infrastructure. When Fujitsu ServerView ROR triggers the creation of a LUN, a volume with the requested LUN is provided by the NetApp controller.

4.1 ADDING NETAPP STORAGE TO FUJITSU SERVERVIEW ROR

Extending storage pools in Fujitsu ServerView ROR translates to assigning a storage controller, or more specifically, its aggregates, to the pool. Figure 6 shows the integration of NetApp storage with Fujitsu ServerView ROR. The aggregate `aggr1` is assigned to the storage pool. While assigning a NetApp storage controller, its credentials have to be configured in Fujitsu ServerView ROR. If multiple storage pools with different storage service levels are to be implemented, as in the following section, an assignment of controller to storage pool has to be performed.

Figure 6) Assignment of a NetApp storage controller (graphic provided by Fujitsu).

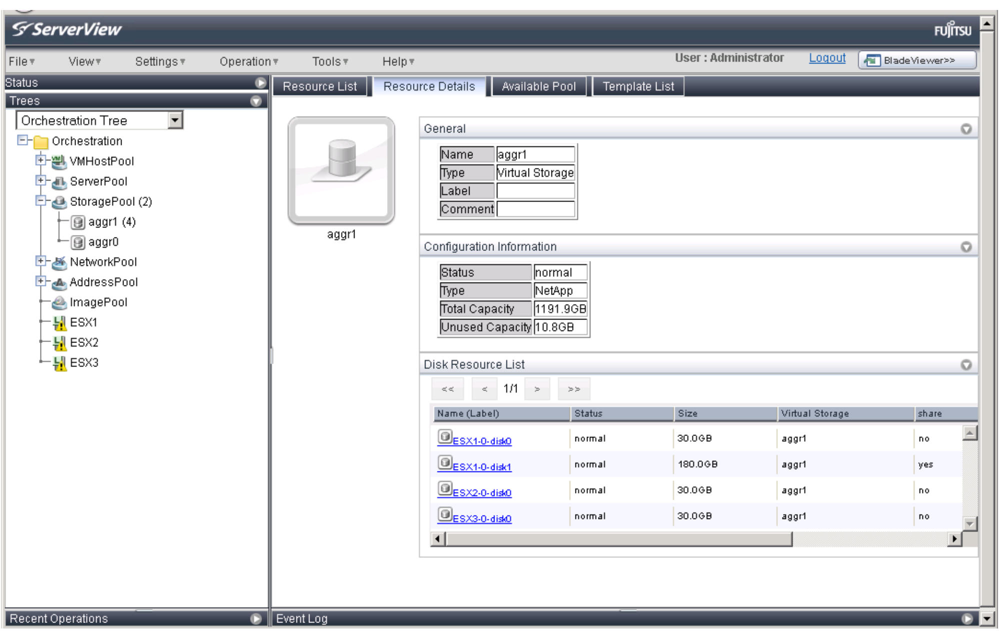

4.2 IMPLEMENTING STORAGE POOLS WITH DIFFERENT SERVICE LEVELS

One facet of the NetApp unified storage paradigm is the ability to incorporate different drive and caching technologies in one system, but without having the necessity to use different means for data protection, resiliency, availability, and management. It is possible to protect, through NetApp replication software, highest performance mission-critical storage using high-capacity and cost-effective SATA technology, combined with NetApp's superior storage efficiency mechanisms.

From a business continuity perspective, NetApp storage systems can consist of:

- One single storage controller (lower SLAs, backup to disk target, archive)
- A local HA paired configuration
- A geographically dispersed HA paired configuration, combined with synchronous data replication (up to 100km), called NetApp MetroCluster

Supporting different drive technologies, intelligent flash-based caching, and storage efficiency technologies allows automated storage tiering to realize multiple storage classes in a single platform. This decreases management overhead with a single management environment and provides greater flexibility to meet the needs of the design-for-service IT model.

This approach offers the following significant advantages over fixed tiers of storage:

- A single storage platform simplifies management and automation.
- You can apply the desired backup and DR policies, or any policy, to any class of service (COS).
- The class of service can be changed without moving data, significantly reducing governance costs.
- The motivation to oversubscribe the top tiers of storage goes away because users can easily change COS if they get it wrong. Other services, such as backup and DR policies, can be added without disruption.
- You only have to procure one type of storage, making it easy to standardize on particular platforms and disk types, which simplifies procurement and provides economies of scale.
- The amount of each capacity available in each class isn't static. If one class becomes more popular, you can shift resources to it without disrupting operations.

Most service providers and IT organizations need only three or four performance levels. It's tempting to describe dozens of virtual tiers, but, in practice, automated storage tiering allows you to use the infrastructure to meet the needs of most workloads on demand. Delivering this level of flexibility is the point of the automated storage tiering.

For a better understanding, NetApp exemplarily designs a couple of storage classes to be used by Fujitsu ServerView ROR. They appear as StoragePool_COS_A and so on in Figure 7.

Figure 7) Implementing storage pools with different characteristics (graphic provided by Fujitsu).

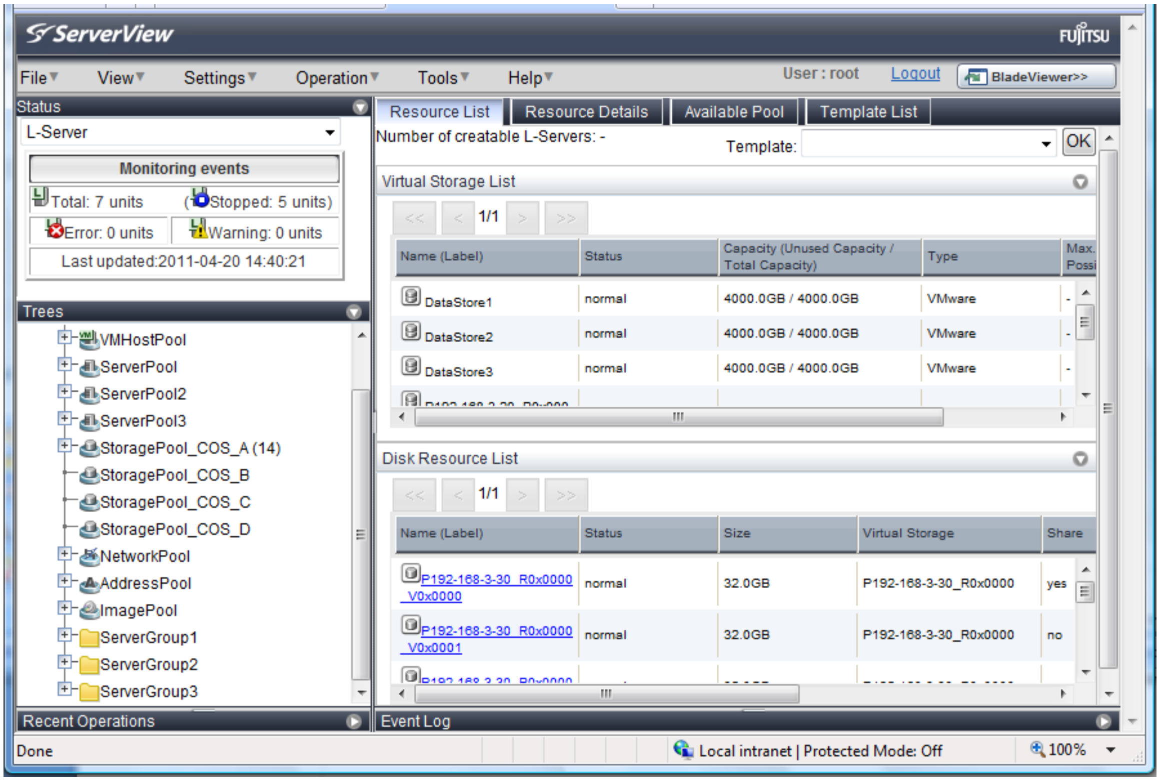

The classes for this sample are defined as follows:

- **COS A.** Highest performance storage; mission critical: highest recovery point objective (RPO)/recovery time objective (RTO) (seconds to minutes):
  - Based on SAS/FC disk drives and Flash Cache (formerly PAM II)
  - Based on fabric-attached storage (FAS) MetroCluster (out-of-the-box HA and DR)
- **COS B.** High-performance storage; business critical: high RPO/RTO (minutes to hours):
  - Based on SAS/SATA disk drives with Flash Cache
  - NetApp FAS systems, asynchronous data replication to DR site possible
- **COS C.** Medium performance storage; standard RPO/RTO (hours to days):
  - Based on SATA disk drives with Flash Cache and FlexShare®
  - NetApp FAS systems, asynchronous data replication to DR site possible
- **COS D.** Low-performance, high-density storage; standard RPO/RTO:
  - Based on SATA disk drives, enhanced storage efficiency enabled (deduplication and compression), or FlexShare
  - NetApp FAS systems, backup and recovery on demand

Figure 8 shows one platform serving different service classes.

Figure 8) Serving different service classes by one platform.

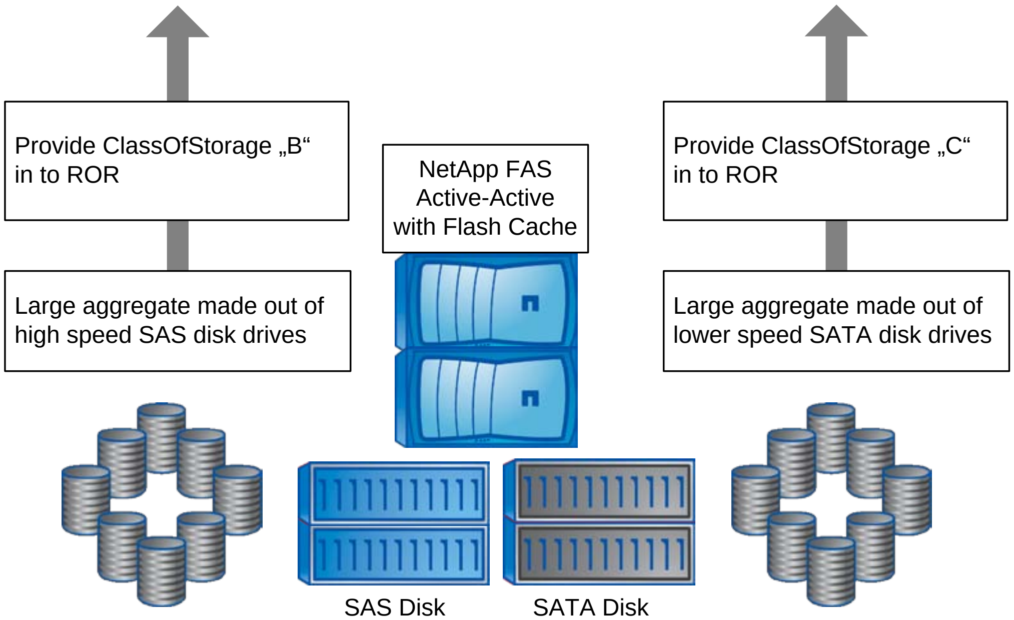

Figure 8 shows a sample implementation of what NetApp called COS B and C. There is an active-active FAS system connected to some significant capacity of both SAS- and SATA-based disk drives. In the sample design, the SAS portion, Class B, is managed primarily by one controller. The SATA portion, Class C, is managed by the other controller. This provides a workload distribution on both controllers, but in case of some failure in one system, the other system can serve both workloads without disruption of service.

More flexibility, but still with consolidation and efficiency in mind, can be achieved by introducing three different performance levels—COS B, COS C, and COS D—out of one system. Figure 9 shows differentiating services classes through FlexShare workload prioritization.

Figure 9) Differentiating service classes through FlexShare workload prioritization.

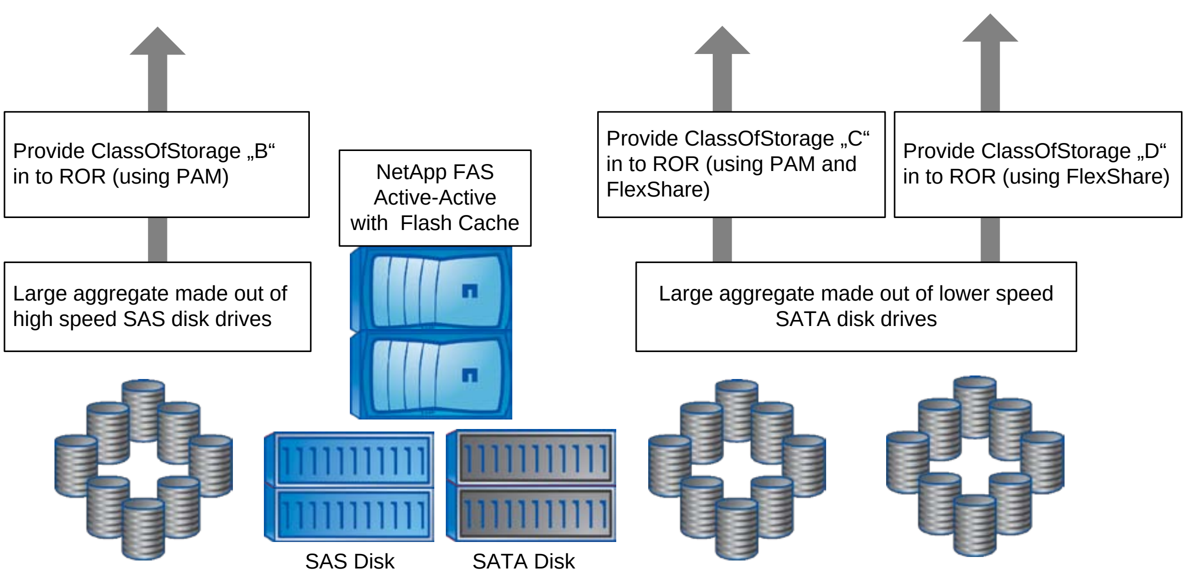

Using the option that Fujitsu ServerView allows to run postprocessing tasks after creating and attaching the LUN storage resource to the L-server, we can apply best practice settings specific to NetApp on the volumes or LUNs. One example of a best practice is to gain maximum efficiency results using deduplication.

The following list presents a possible use case:

- Messaging application, for example, Microsoft® Exchange 2010, is implemented using the SAS-based storage resources (COS B). E-mail is critical in most companies and needs to run on performing platforms, encompassing both read and write I/O.
- A collaboration application, for example, Microsoft SharePoint®, is implemented using the SAS-based storage resources (COS B). Due to the nature of user activity and workload characteristics, we can assume a read-write ratio of about 70:30. Therefore, retrieval load is well accommodated by NetApp Flash Cache.
- Some archival application has storage from COS C assigned. This application is not performance critical and uses storage system resources as available. This dynamic is achieved by assigning process prioritization using NetApp FlexShare.

5 CONCLUSION

This technical report describes the integration of Fujitsu ServerView ROR and NetApp storage to realize IaaS. The report focuses on showing how the combined HA capabilities and the seamless integration of different storage classes provided through the NetApp unified storage platform can offer improved business continuity across servers and storage, along with more service level–based management of storage resources.