REALIZING SOFTWARE-DEFINED STORAGE WITH EMC ViPR

EMC ViPR
Software-Defined Storage

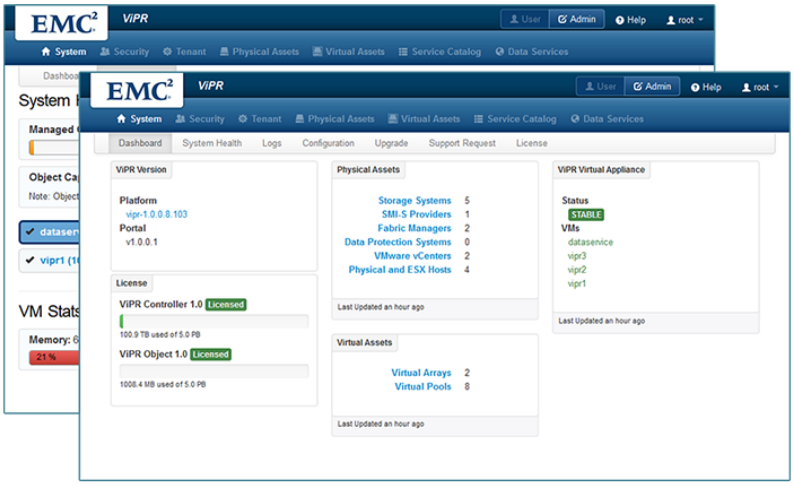

Storage automation
Provisioning storage with fewer steps and less time means potential cost savings.

Monitoring & reporting
Extend monitoring and reporting on storage health into enterprise management tools.

VMware integration
Execute storage automation within vSphere workflows, shortening deployment times.

Object storage access
Leverage cloud-based applications within your own infrastructure.

Meeting the storage requirements of a large enterprise can be an overwhelming task. Organizations spend a great deal of time keeping up with the dynamic storage needs of all the various groups they support. As storage needs grow, and new arrays from multiple vendors are put in place to meet the ever-growing demand for data storage, administrators may find themselves drowning in the procedures needed to maintain individual storage systems. Consistent monitoring and reporting across these arrays means manually extracting the data from the storage systems, correlating the data, and transforming it into a useable data set. With exploding data growth and the rise in virtualized infrastructure, storage consumption can occur at astounding rates, with frequent changes and additions required to meet demands. Organizational inefficiencies, such as complicated change-control processes and workflow hand-offs, introduce additional delays and make it difficult to provide the rapid responses modern customers have come to expect.

EMC ViPR, a software-defined storage solution, provides a streamlined, uniform storage-provisioning interface, as well as a common API to the storage layer regardless of the underlying storage infrastructure. A software-only product, ViPR integrates your existing storage resources into a single virtualized platform capable of making storage more agile and less complex. Storage automation removes the risk of human error, and allows organizations to execute storage delivery in a consistent, timely manner without

APRIL 2014
A PRINCIPLED TECHNOLOGIES TEST REPORT
Commissioned by EMC

having to wait for multiple hand-offs among functional groups, or for change-control approvals.

ViPR has two components:

1. ViPR Controller, which automates and centralizes storage management, configuration, provisioning, and delivers storage-as-a-service.
2. ViPR Services, which provides cloud-scale storage services that can be layered over ViPR-supported arrays to unlock new operations on data-in-place.

ViPR is not in the data path for block and file storage, which means performance charactistics of the underlying arrays, as well as their native capabilities, remain intact. ViPR abstracts data access, removing dependencies and limitations common to storage systems, which makes it possible to develop and deploy new services such as the ability to store objects on existing file storage systems.

Both in our own labs and on site at EMC, Principled Technologies worked closely with ViPR to learn about how it automates storage provisioning, consolidates monitoring and reporting, integrates with VMware, and provides cloud-compatible object stores that enable you to bring your cloud applications in-house to your datacenter. ViPR proved to be extremely effective in all of these areas and could be a very wise investment for enterprises seeking to add a software-defined storage solution to their infrastructure.

SIMPLE, EXTENSIBLE, AND OPEN

While the amount of data that enterprises are storing has grown exponentially, the number of administrators supporting this storage has not. As a result, organizations face an increasingly heavy management burden. EMC ViPR virtualizes the underlying storage and automates the associated provisioning processes, simplifying storage delivery throughout the organization. ViPR is a simple, extensible, and open software-defined storage platform that consolidates heterogeneous block, file, and object data storage into a single virtualized platform, providing consistency and ease-of-use for every storage array on the back-end. EMC ViPR has a growing catalog of supported storage arrays, including third-party storage arrays such as NetApp. EMC ViPR can integrate with enterprise monitoring tools to provide detailed metrics for the storage environment. ViPR plugins for VMware vSphere provide automation and monitoring capabilities to virtualization administrators. With object storage capabilities, EMC ViPR can provide compatible storage for cloud-based applications utilizing traditional file storage arrays.

In this report, we explore these various aspects of ViPR and its efficacy, efficiency, and usability as compared to existing physical storage solutions and discuss potential savings based on those findings.

DETAILED TEST CASES

About our approach

We wanted to see how the automated ViPR approach would differ from the manual process required without ViPR. To that end, we carried out a series of tasks typical of those that storage administrators execute on a regular basis. As we performed these tasks both with ViPR and manually, we recorded the number of steps and the amount of time required. Our steps include clicks and data entry items. In all of our self-service provisioning test cases, ViPR could execute a procedure to provision storage in fewer than five clicks.

In all of our self-service provisioning test cases, ViPR could execute a procedure to provision storage in fewer than five clicks.

Using ViPR's automated approach was very straightforward; our timer ran continuously from the time we started each task until we completed it. Using the various corresponding manual processes to carry out the tasks, however, was more complicated. These approaches required us to type commands, and we made occasional typographical errors. We also needed to engage in a certain amount of trial and error as we worked.

For example, if our initial manual attempts at zoning a host to an array failed, we had to drill down into the host bus adapter (HBA) of each system and determine whether the selected interface was the correct one, and whether we had correctly keyed in the zone entry in the switch. We corrected our mistakes each time and eventually documented a flawless manual methodology, which we used for our timed testing. For this reason, the time and number of steps that we provide for each manual task in this report reflect an unrealistically error-free experience—a best-case scenario. Under real-world conditions, the time and number of steps necessary for the manual processes would likely be considerably higher than those we report. Additionally, in many environments, the hosts, SAN switches, and arrays may be managed by different groups, and the hand-offs required to execute the tasks would introduce additional delays and the potential for human error. Because our tests were not managed within a change-control approval process, that additional time could not be directly quantified. Consequently, in real-world situations, the reduction in time and complexity that ViPR offers could be much greater than what we report here.

ViPR's automated approach reduced human error and provided for a smooth provisioning experience compared to executing tasks with a manual procedure.

Scenario

Many large enterprise datacenters include a variety of storage arrays the company has acquired over the years. These arrays are frequently from different vendors and may be managed by different groups within the organization. Because each

array has a different method for executing tasks, storage administrators must be careful to keep working documentation for each procedure and ensure these procedures transfer to new owners. When documentation is not fully maintained, the risk of errors increases and quality control can be seriously jeopardized. In addition, because of bureaucratic complexity, it can take weeks to request, review, and implement changes in storage provisioning, and the process can still be prone to human error.

Test environment

For our tests, we visited the labs at EMC, where we were provided a test environment with physical hosts connected by HBAs to Cisco MDS storage network switches and to a data network. These hosts contained a generic installation of Red Hat Enterprise Linux 6, Windows 2012, and VMware ESX 5.1. The storage switches connected the physical hosts to EMC block storage arrays—VMAX and VNX—configured with disks allocated into disk pools for provisioning. The data network presented access to EMC VNX file storage, as well as to EMC Isilon and NetApp file array simulators. We document this environment in our reporting.

Additionally, EMC provided a VMware-virtualized environment suitable for hosting the various infrastructure pieces required to host and support our ViPR installation, such as network services, DNS, and other various virtual appliances. We do not include this infrastructure in our test reporting because it falls outside the framework of our test environment. We installed the ViPR solution, using an OFV template, into this infrastructure environment.

We tested ViPR 1.0. At the conclusion of our testing, we performed an upgrade to ViPR 1.1 and found the upgrade process to be very easy. ViPR 1.1 differs from ViPR 1.0 in that it provides support for the ViPR Hadoop Distributed File System (HDFS) service, as well as the ability to provision EMC SRDF with supported block storage. The additional functionality provided by ViPR 1.1 fell outside the scope of our tests. Therefore, all the information and test results presented here for ViPR 1.0 should be the same for ViPR 1.1.

The ViPR interface provides the ability to toggle between admin and user functions with just a single click. Admin functions include tasks such as discovery and ingestion of physical assets such as arrays, fabric managers, and hosts; defining virtual assets such as storage pools for provisioning volumes or the storage networks used to provide storage to hosts; and configuration tasks such as setting up ViPR Services for object. The admin view can be used to modify the User Service Catalog to ensure users have access only to the provisioning functions allowable for that user. User functions are tasks that involve provisioning storage or performing tasks on associated hosts once storage has been provisioned. These self-service tasks automate all of the time-consuming and difficult processes that occur when a user makes a storage request.

See Appendix A for information on deploying ViPR into a pre-existing environment.

Storage automation

Identification, discovery, and registration

Before we begin, it is important to point out that ViPR supplements array-specific management interfaces such as EMC Unisphere. Storage administrators will still use storage element managers to set up the underlying components of their arrays, define disk pools, and perform other custom configuration tasks. The individual array management interfaces are still used to manage the physical aspects of each array—ViPR augments physical storage element managers with an abstraction layer to automate storage provisioning, relieving storage administrators of time-consuming, error-prone tasks. ViPR does not configure storage arrays; it consolidates and virtualizes storage provisioning, as well as monitoring and reporting, across the supported arrays connected to it.

ViPR supplements array-specific storage element managers, such as EMC Unisphere, by adding a layer of abstraction above the physical management of connected storage systems.

After the ViPR installation was complete, we began configuration by adding the various storage platforms to the system. After adding a storage array to ViPR, the ViPR system automatically obtains information directly from the array and determines the amount of storage in use , as well as the amount of storage available for provisioning. As shown in Figure 1, simply enter the Admin → Physical Assets section of ViPR and add the IP address of the storage array management controller to the storage components to be managed by ViPR.

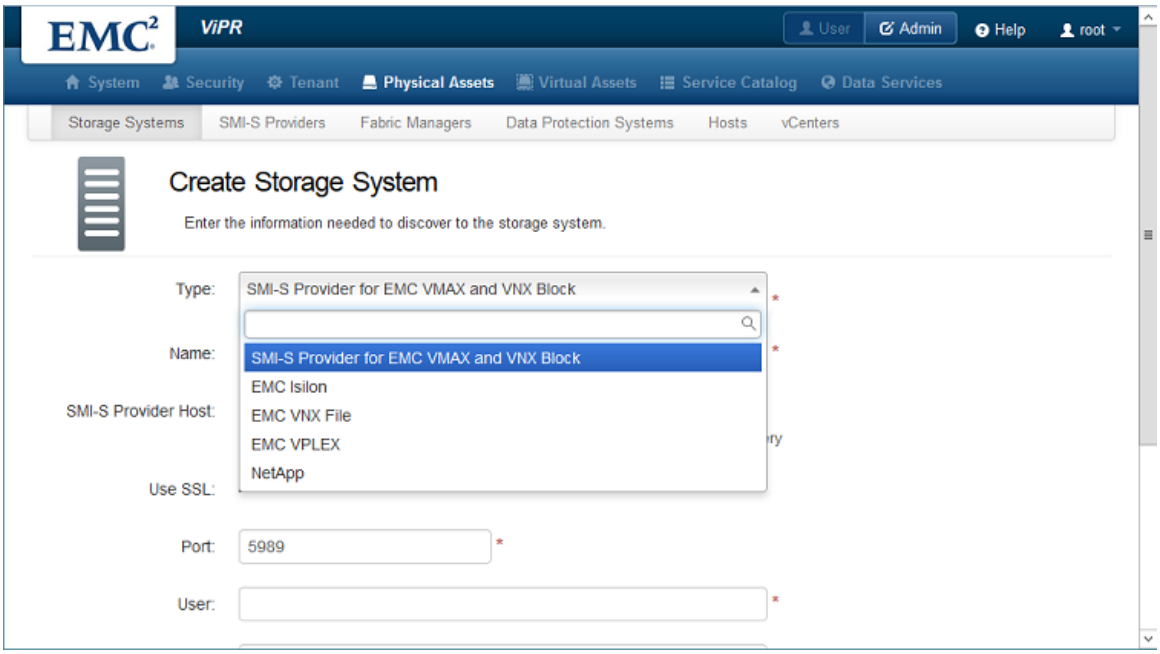

Figure 1: Adding storage systems to ViPR is as easy as entering a few configuration parameters and clicking Save.

A Storage Management Interface – Specification (SMI-S) provider is necessary to discover block storage. The SMI-S provider is a host directly connected to block storage arrays, which acts as a sort of proxy for gathering configuration information from an array. Connected arrays that cannot be discovered directly, such as VMAX, will provide gatekeeper LUNs to the SMI-S provider. The gatekeeper LUNs allow the SMI-S provider to obtain the information from the target storage arrays, which ViPR can then use for array registration. The gatekeeper LUNs must be configured using the VMAX array Unisphere instance—those tasks fall outside the scope of this report. Discovery of a VNX array by an SMI-S provider does not require gatekeeper LUNs.

As shown in Figure 2, instead of adding the array directly, enter the SMI-S provider section and define a new SMI-S provider. After adding the SMI-S provider hosts to the ViPR configuration, ViPR automatically identifies the connected storage, discovers its capabilities, and registers it for use. In our test environment, the ingestion of storage into ViPR took less than a minute.

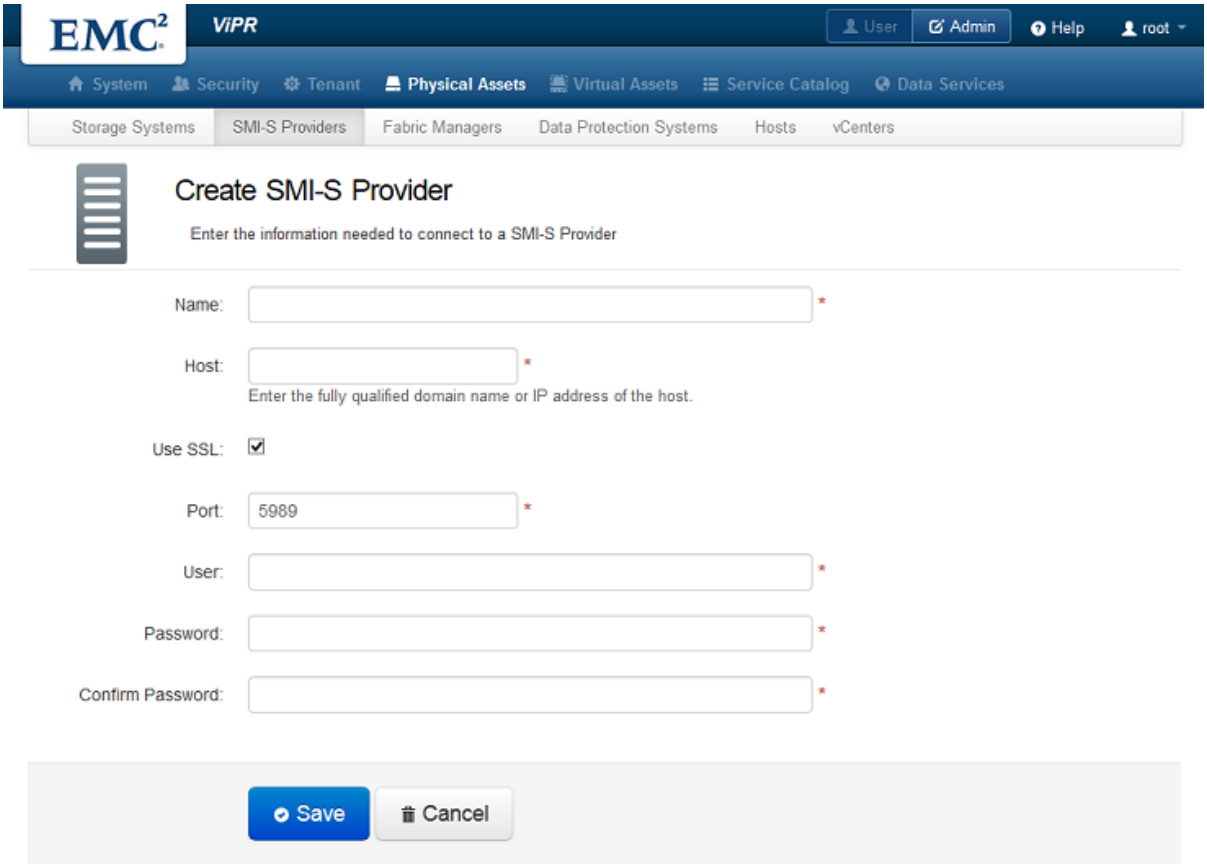

Figure 2: Defining a new SMI-S provider in ViPR.

Similarly, connecting file storage to ViPR was very easy. We simply entered the Admin → Physical Assets → Storage Systems section and put in the configuration details for the file storage array we were connecting. As shown in Figure 3, ViPR logs into the array and discovers its capabilities.

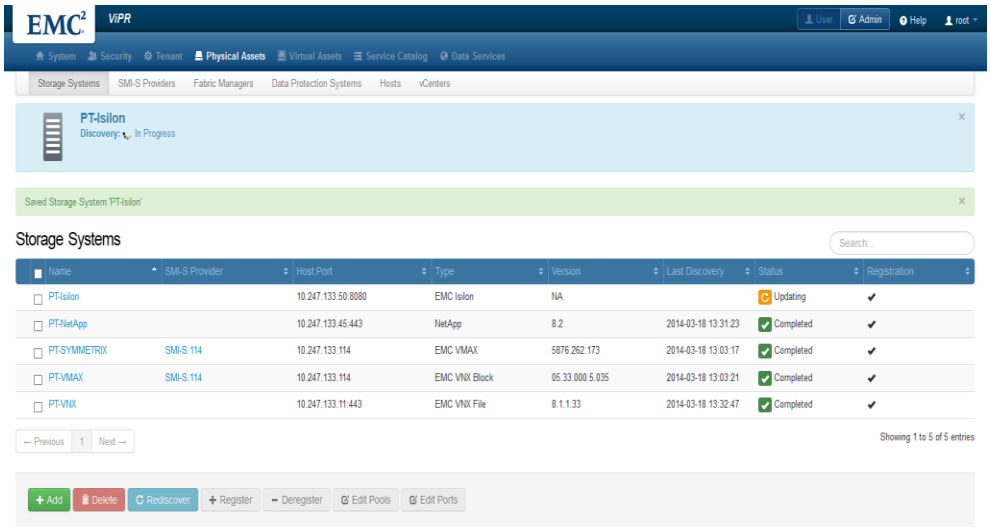

Figure 3: Connecting file storage to ViPR.

See Appendix B for information on adding block storage arrays, and Appendix C for details on how to add file storage arrays to ViPR.

Compared to manual methods, SAN zoning is easy with ViPR. Because ViPR can manage storage networks, an administrator can simply define hosts and fabric managers within ViPR and let the system do the work. ViPR logs into each system, identifies its hardware components, and makes them available within ViPR. The administrator only has to define the parameters of a given storage network and ViPR makes all necessary connections automatically. It is important to mention that configuring ViPR to handle the zoning is a one-time effort, while manually performed SAN zoning tasks would recur as changes, additions, or deletions necessitate. As Figure 4 shows, using ViPR to perform SAN zoning requires no time or administrative effort, saving us 7 minutes, 17 seconds over the idealized manual process each time we need to amend the network, not accounting for any additional time caused by distractions, human error, or retries.

EMC ViPR uses best practices to execute tasks in a consistent, repeatable manner.

See Appendix D for the methodology we used to perform this test.

Setting up SAN zoning	Number of admin steps	Admin time (mm:ss)
ViPR	0	00:00
Manual	24	07:17

Figure 4: Using ViPR for SAN Zoning completely eliminated associated administrative time and effort compared to the manual method.

Bottom-line impacts

Because setting up ViPR, just once, to perform your zoning tasks eliminates the need for you to perform manual procedures every time you make a change, using ViPR to manage your environment means you can significantly reduce your administrative effort and save measureable time—and associated staff costs.

Organizational efficiency

It is important to take a moment and consider the impact automation can have on an organization. In large environments, roles and responsibilities are commonly divided among specialized functional groups that manage smaller "silos" within the infrastructure. To organize workflow, organizations may implement IT Service Management (ITSM) to help coordinate changes due to requests, new releases, and response to incidents.

ViPR automation reduces the risk of human error and the need for lengthy management reviews, and can make your organization more nimble.

Because each group involved in a change must successfully execute its part of the change and then hand off the next steps to other groups, management is usually involved to approve changes and to understand potential risks to the overall environment. Human error can be a significant risk, and the approval process can introduce lengthy wait times while documentation and planning occur. Additionally, changes are typically performed during times that are least likely to affect customers, so that an organization can "roll back" a change that creates an unforeseen impact to the environment. All of these factors can impede the ability to fulfil a request quickly.

One of the major benefits of storage automation as implemented by ViPR is its ability to remove change control from the procedures. By pre-allocating storage, setting limits, and defining boundaries for a customer within a Tenant/Project, customers—or the administrators fulfilling their requests—can provision their own storage as needed, without impact to the rest of the environment. ViPR automation reduces the risk of human error and the need for lengthy management reviews, and can make your organization more nimble.

Provisioning

We tested the ability of a customer to use ViPR to provision block storage and prepare it on the target Windows host compared to performing those same tasks manually. We captured the time and steps required to perform those tasks using each method, and compared the results.

As Figure 5 shows, compared to the manual approach on both VMAX and VNX, using ViPR to provision block storage reduced the number of steps by up to 76.1 percent and reduced time by up to 86.5 percent.

See Appendix E for the methodology we used to perform this test.

Provision block storage	Number of admin steps	ViPR reduced steps by…	Admin time (mm:ss)	ViPR reduced time by…
ViPR	11		00:44	
Manual - VMAX	46	76.1%	05:27	86.5%
Manual - VNX	37	70.3%	02:49	74.0%

Figure 5: Time and number of steps to perform the block storage provisioning task with and without ViPR. Fewer steps and less time are better.

With ViPR, it takes less time and fewer steps to perform the provisioning test we executed. Because it can interact with a storage system, a fabric manager, and a host directly—and work across physical IT infrastructure from end-to-end, ViPR was able to automatically provision block storage, perform the necessary SAN zoning required to present the new storage to the host, and mount and format the new storage for use on a Windows host. The automation is so simple and complete that a non-administrative customer with access to the ViPR console could perform this function with no administrative assistance.

Compare that with the manual methods, which would require at least one storage administrator to provision the storage and perform the storage network changes. In larger organizations, one would expect to have the request accepted by a customer service specialist, a change control item created, management approval required, storage administrators to create the storage, network managers to perform the changes, host administrators to discover and format the storage, and potential delays during quality control and hand-off phases. In short, ViPR required less time and effort to execute a sequence of tasks than the manual process did; in a real-world situation, ViPR could certainly help circumvent unnecessary delays and reduce administrative burden.

In our tests, ViPR required less time and effort than manual processes.

Just as we did with SAN zoning automation, we can quantify storage automation savings in the steps required to provision storage and the time involved. Organizations can then apply their own financial model to calculate the potential savings in operational costs.

Provisioning file storage with ViPR was equally simple. One advantage to using ViPR to manage storage arrays is that the procedure for creating new storage is uniform across all storage types. In user mode, you simply use the service catalog to select the tasks you'd like to perform, fill out the required information (such as host name and the amount and types of storage you want), place the order for storage, and ViPR handles the rest. Figure 6 illustrates the process.

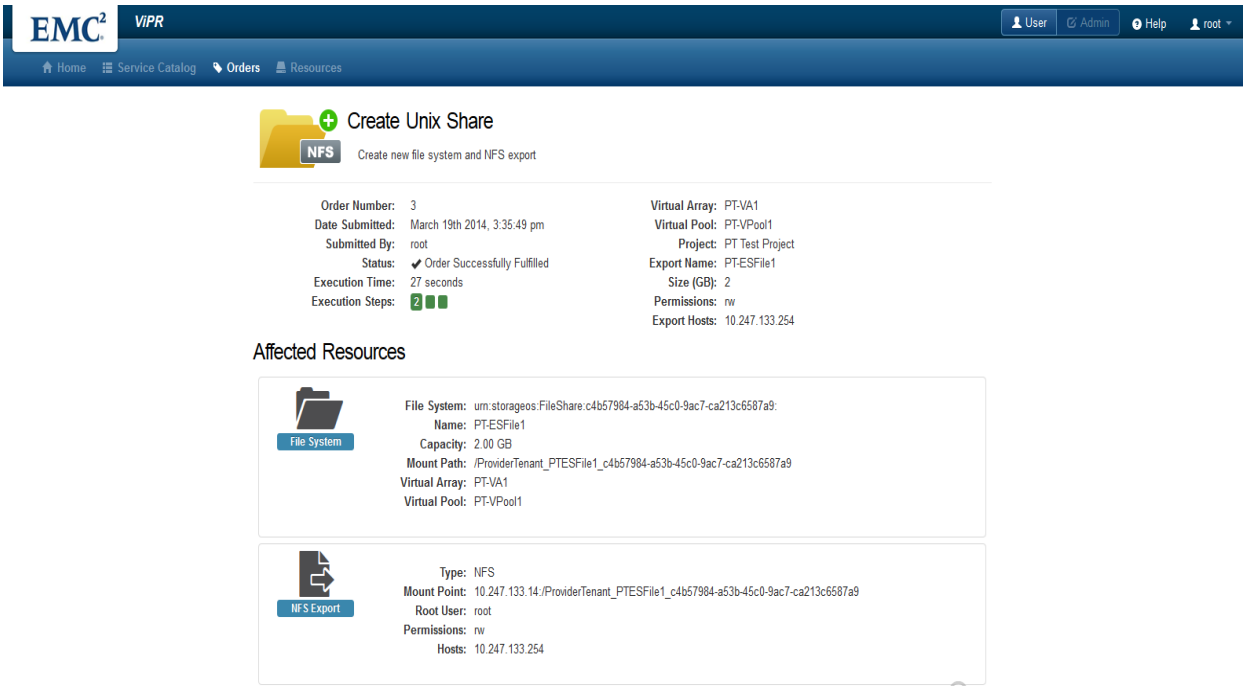

Figure 6: Provisioning file storage with ViPR.

EMC ViPR provides a single interface and method for provisioning heterogeneous storage.

In our tests, we compared the ability of a customer to create file storage (NFS) with ViPR compared with the manual procedure for three different storage systems. In addition to requiring less administrative time, we were able to use the same interface and method for each different system, an important factor in large organizations that utilize a wide variety of storage platforms.

Each of the platforms we tested had significantly different access methods and user interfaces. A user or administrator would have to become familiar with the capabilities and operation of each storage system to utilize the manual method effectively. ViPR provides a consistent user experience, regardless of the storage system—file or block—used to provision the new storage.

As Figure 7 shows, compared to the manual approach on VNX and Isilon, using ViPR to provision file storage reduced the number of steps by up to 45.0 percent and reduced time by up to 60.5 percent. When used with NetApp file storage, ViPR reduced the time to provision storage by almost half (45.6%).

Provision file storage	Number of admin steps	ViPR reduced steps by…	Admin time (mm:ss)	ViPR reduced time by…
ViPR	11		00:49	
Manual - VNX	12	8.3%	01:35	48.4%
Manual - Isilon	20	45.0%	02:04	60.5%
Manual - NetApp	15	26.7%	01:30	45.6%

Figure 7: Time and number of steps to perform file storage provisioning task with and without ViPR. Fewer steps and less time are better.

As with the other provisioning tasks we tested, the reductions in time and effort ViPR brings compared to manual methods can contribute to cost reductions in operating expenses. That's in addition to time saved by avoiding change management and needless workflow hand-offs.

See Appendix E for details on how to provision block storage, and Appendix F for procedures about how to provision file storage.

Integration with native array-based features

ViPR utilizes array-native features, such as FAST, and can allow users to order snapshots of volumes from a common user interface, regardless of array and data type, when supported by the underlying storage.

EMC ViPR, as a storage provisioning solution, can utilize the native capabilities of the connected storage arrays. When provisioning storage, ViPR will take advantage of the features presented by the disk pools allocated to it by array administrators. For example, if an array supports EMC Fully Automated Storage Tiering (FAST), and the array administrator presents a FAST-capable pool to ViPR for provisioning, any storage provisioned from that disk pool will take advantage of the FAST features.

Another example of how ViPR can leverage the native features of its managed arrays is its ability to protect data at the storage level. When virtual pools are created in ViPR, administrators can simply turn on data protection services and define the number of snapshots and continuous copies to allow for that pool. As shown in Figure 8, ViPR users can order a volume snapshot within the service catalog, provided the snapshot capability is available on the storage array housing the volume.

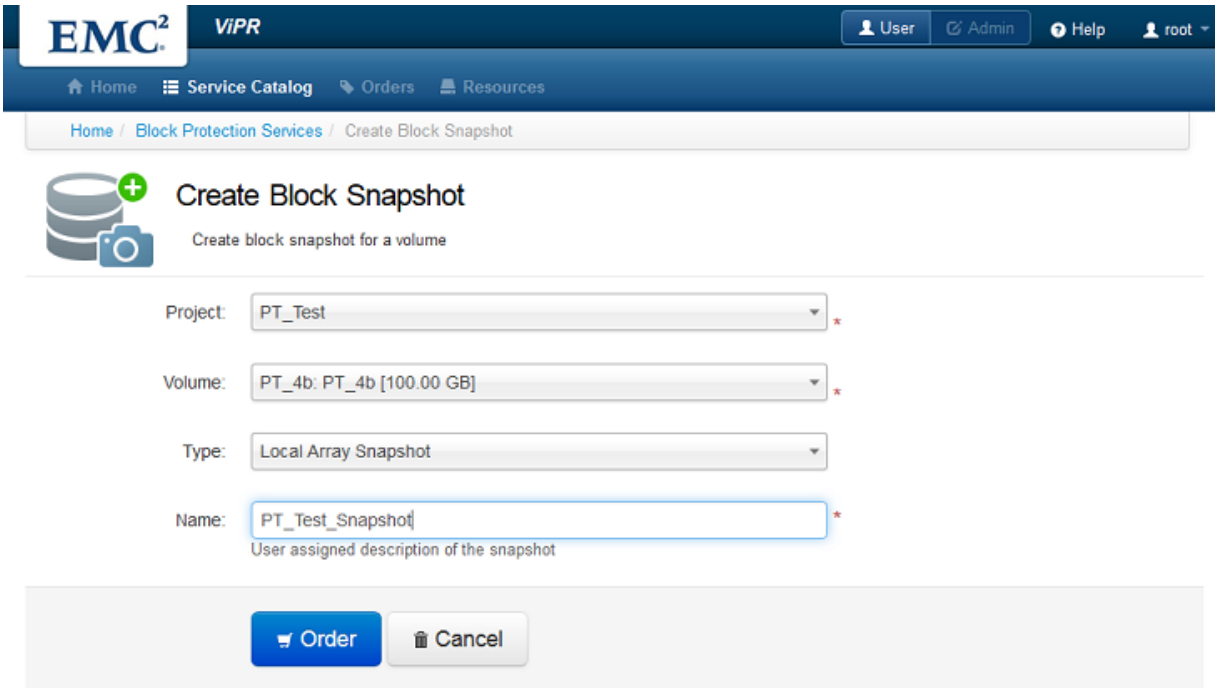

Figure 8: With ViPR, users can self-service provision snapshots of the data in their volumes.

Additional information

In addition to the storage automation we tested above, the EMC ViPR abstraction layers integrate and complement EMC VPLEX—a storage aggregation and virtualization engine that masks the back-end storage from front-end hosts and presents volumes and storage across the connected arrays. EMC VPLEX provides the ability to transparently move storage between arrays—even across geographically dispersed areas. EMC ViPR can leverage VPLEX like other physical storage arrays—ViPR provides the automation and orchestration for storage provisioning while VPLEX provides virtualization of the actual provided storage. For more information on EMC VPLEX and ViPR, see the following links:

- The Storage Chap blog: VPLEX or ViPR
- ViPR Community: What does ViPR do in a VPLEX environment?

ViPR can be configured to utilize EMC RecoverPoint for data protection. RecoverPoint is added as a Physical Asset in the data protection subsection, and can then be used as the data protection type when provisioning new ViPR storage. For more information about EMC RecoverPoint and ViPR, see the following link:

- ViPR Community: Procedure for adding data protection to ViPR

Monitoring and reporting

Monitoring

The ability for administrators to monitor the complete storage system is one of the most important aspects of managing a solution. Administrators use monitoring statistics to identify and resolve health issues, determine whether sufficient resources are available to meet the ongoing needs of an organization and its customers, and to help management make decisions about the timing for procuring new storage—and retiring existing storage. Additionally, many organizations utilize thin provisioning. Thin provisioning is a method to over-provision storage in the short term—allocating all of the storage a customer requests without having all of the physical storage on hand. This allows storage administrators to add additional physical resources to meet the actual usage demand as needed, rather than making costly investments up front and leaving large amounts of storage unused. Administrators need to be able to manage actual utilization compared to allocations in order to ensure they always have sufficient resources available to meet demand.

The ViPR dashboard view can immediately inform users of a potential system health issues. Shown in Figure 9, this view can be useful for verifying the service is stable and there are no licensing issues to address. The dashboard shows the Physical Assets managed by the ViPR system, providing a quick way to drill into the various managed components for additions, changes, or deletions.

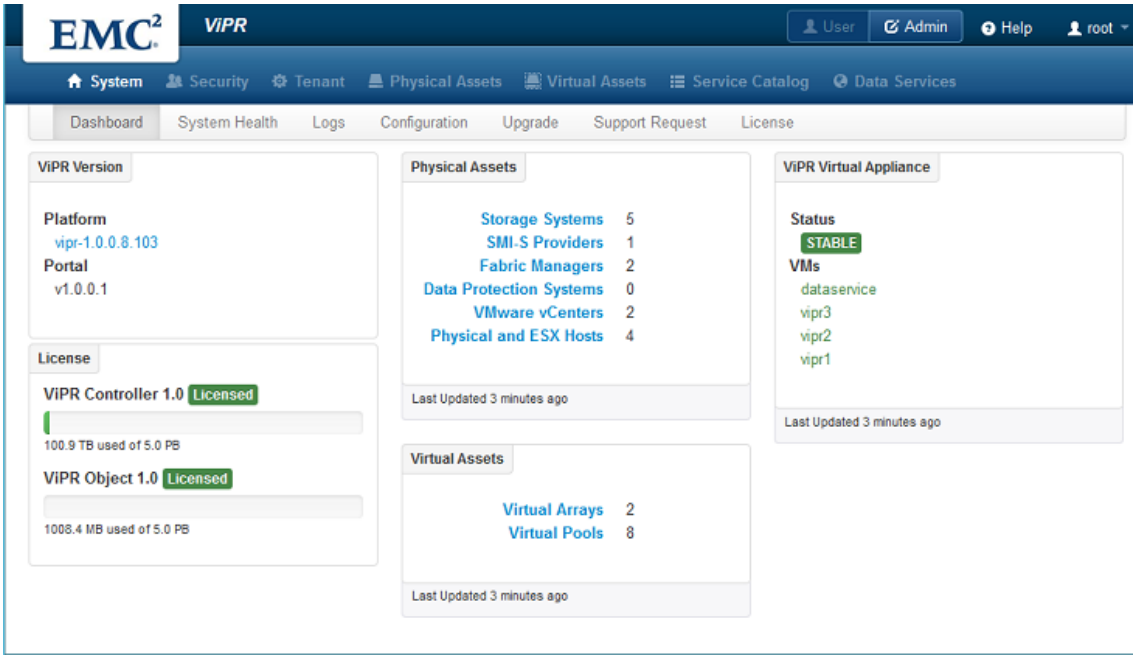

Figure 9: The ViPR dashboard provides a quick way to check on the general system health.

The dashboard is useful, but administrators have an even more detailed view within ViPR. As shown in Figure 10, System Health gives quick-glance view of the sub-components of the ViPR system. This view can give administrators a quick way to drill into the components and look at the stats associated with each member. Clicking each of the components changes the view to reflect the statistics of each node.

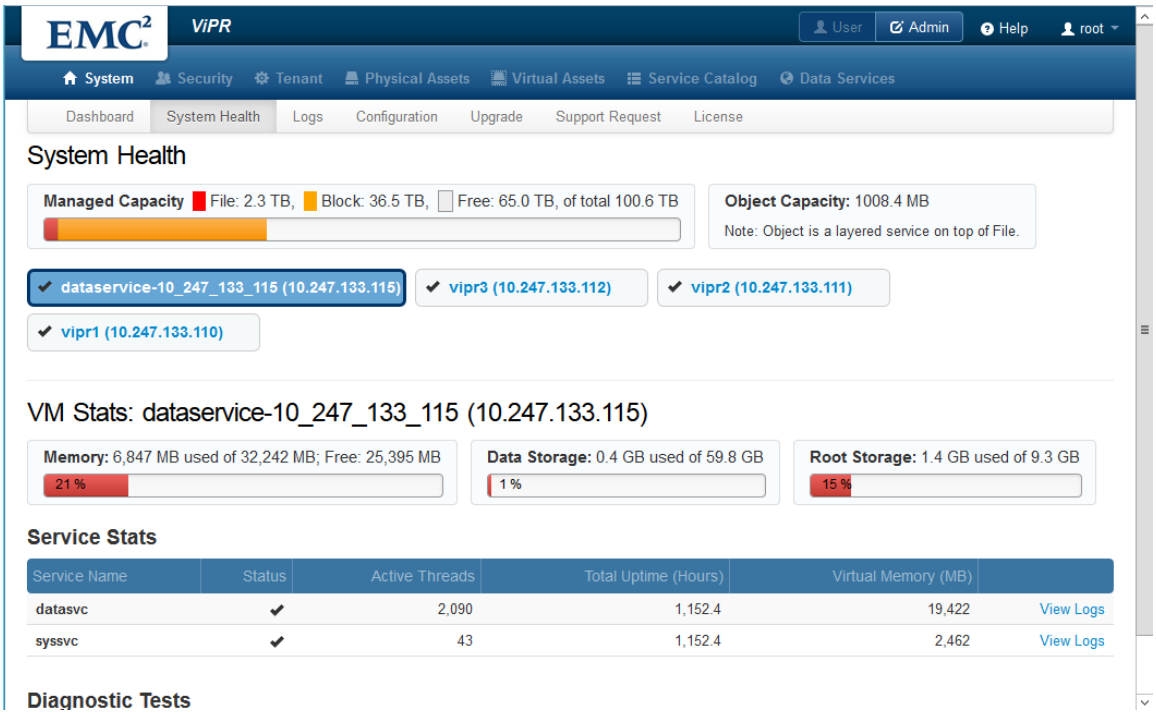

Figure 10: System Health gives a view of the statistics for each of the sub-systems within ViPR.

EMC ViPR integrates into existing management tools to extend views into ViPR supported storage.

The EMC ViPR solution is more than just the base installation. Organizations can take advantage of plugins, extensions, and management packs to add ViPR functionality into management tools already deployed in datacenters. For example, the ViPR monitoring and reporting SolutionPack delivered with ViPR allows administrators to get a complete in-depth picture of the health, capacity, and availability of the storage arrays attached to ViPR. Figure 11 shows an example of the ViPR monitoring views.

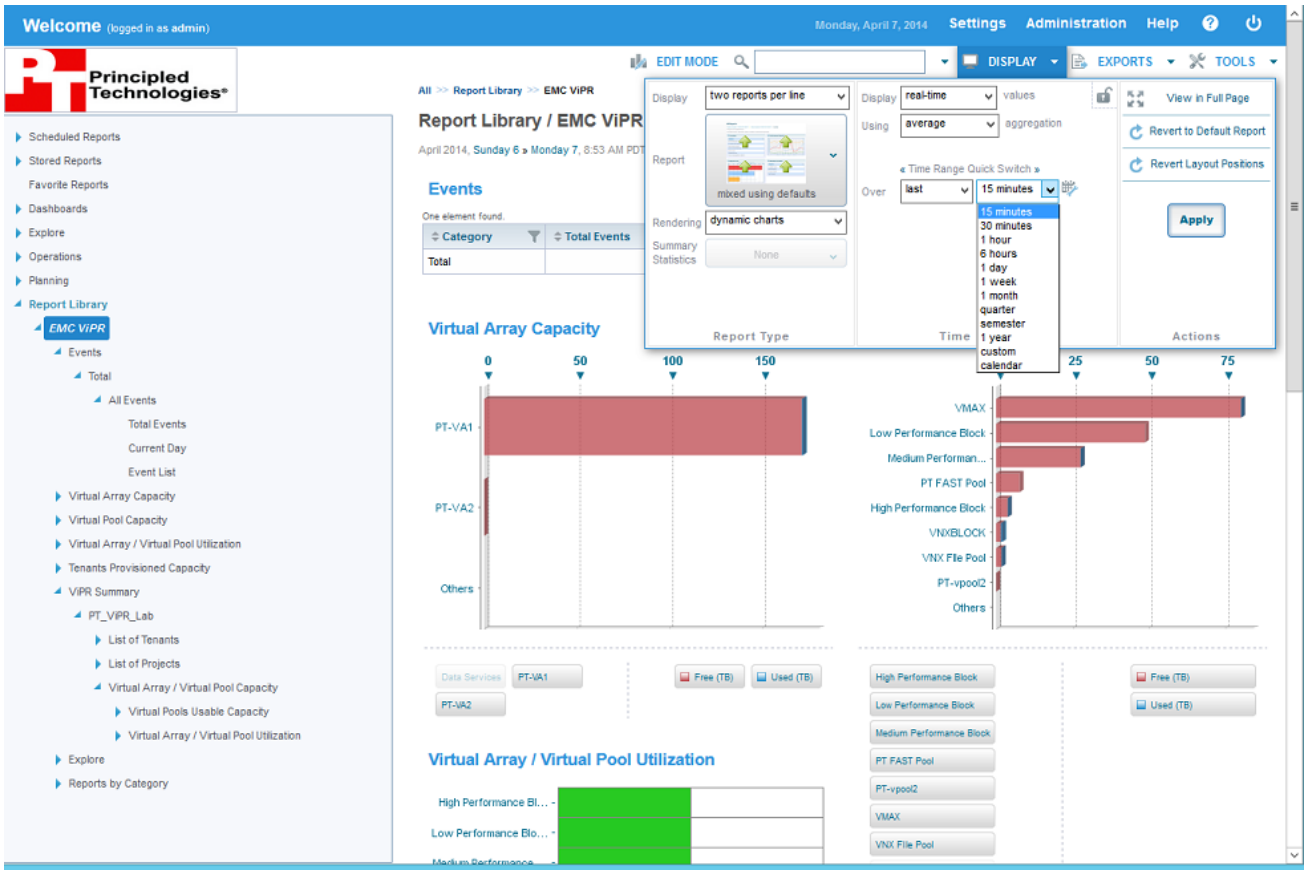

Figure 11: The ViPR SolutionPack allows administrators to view events, capacity, and system health information.

Reporting

Another advantage of using ViPR plugins in existing management applications is its ability to generate customized reports administrators can use to maintain documentation on their environments over time. Organizational management may use other specialized reports to guide purchasing and for developing strategies for meeting their customers' needs while expanding their customer base. Figure 12 shows how easy it is to export a report with the ViPR reporting mechanism. See Appendix L for a sample report.

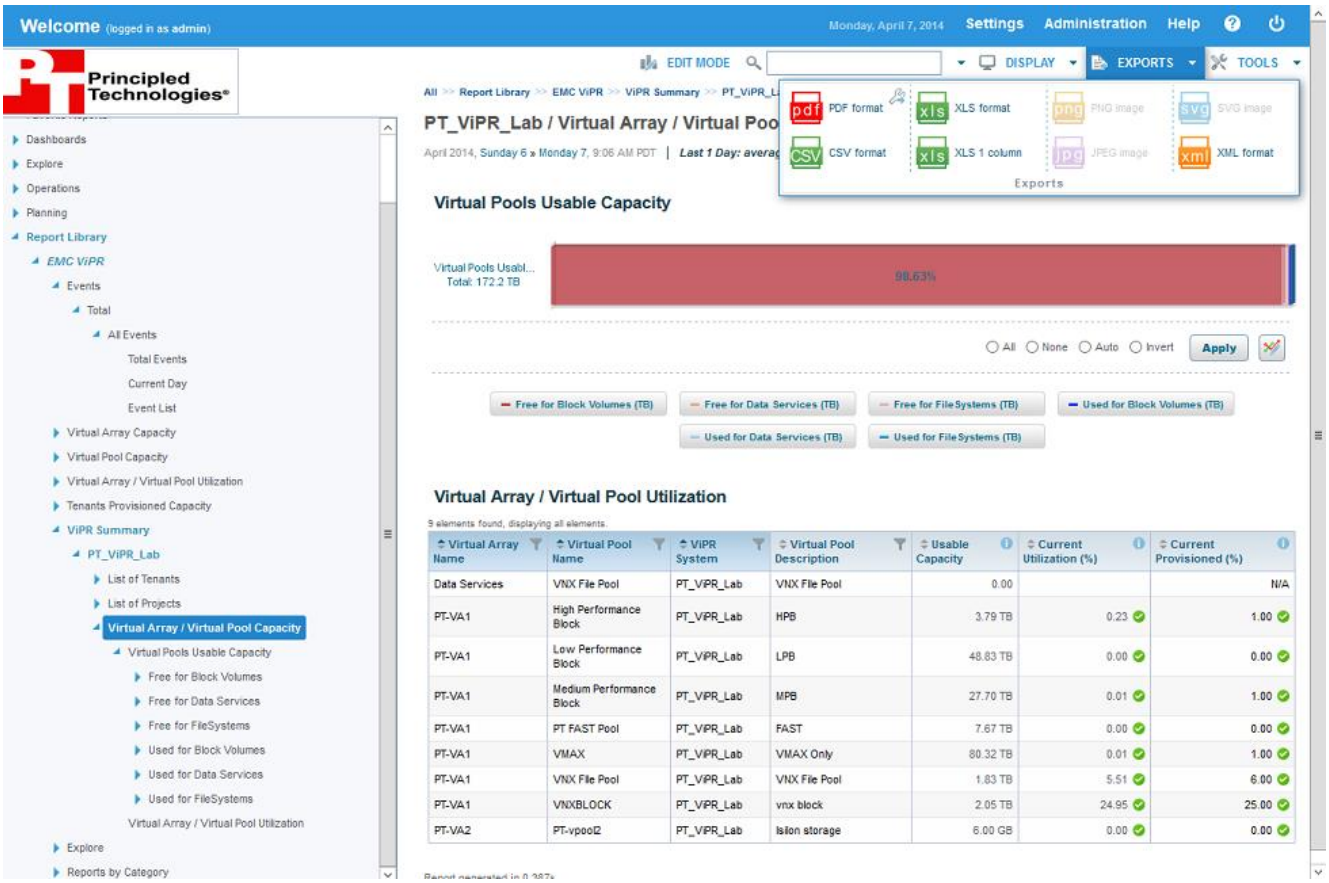

Figure 12: EMC ViPR with the ViPR SolutionPack makes it easy to create reports that enhance storage management.

Additional information

The EMC Storage Resource Management (SRM) Suite is a feature-rich application suite designed to allow storage administrators to "visualize, analyze, and optimize"[1] their storage environments. The ViPR Solution Pack bundled with the ViPR Controller provides some of these advanced enterprise features as part of the base ViPR package. The ViPR SolutionPack can be augmented by adding the SRM Suite, which provides enhanced capacity planning and management with consistency for the entire ViPR-supported and legacy storage environment. For more information about the SRM Suite and ViPR integrations, see the following links:

- Technical Documentation: EMC ViPR SolutionPack Installation and Configuration Guide
- Demonstration: EMC ViPR 1.1 and EMC SRM Suite 3.0

[1] www.emc.com/data-center-management/storage-resource-management.htm#!details

VMware integration

In addition to integration with monitoring applications, EMC extends the ViPR product for integration into virtual environments, such as VMware vSphere. From direct integration with vCenter to plugins and integration packs that enhance virtual environment monitoring, orchestration, and cloud deployments, EMC ViPR makes storage automation an excellent companion for virtualization administrators, and delivers on the promises of a software-defined datacenter. Organizations using EMC ViPR with VMware provide their customers the means to consume their storage through VMware. A VMware vSphere environment with ViPR integrations enables users to orchestrate workflows, present ViPR-supported storage, and monitor their software-defined storage solution along with the rest of their virtualized infrastructure through vCenter and vCloud tools.

VMware vCenter Orchestrator

In our test environment, we set up VMware vCenter Orchestrator (vCO) and imported the ViPR vCO plugin. With that simple integration, we were able to execute an orchestration workflow that automatically provisioned a new disk for an existing virtual machine (VM).

Figure 13 shows the workflow as tasks are performed by vCO. The power of vCO, when extended with the ViPR vCO plugin, means that administrators can automate many of their deployment tasks for both new and existing machines—reducing the risk of human error and decreasing the burden on both virtualization administrators and storage administrators.

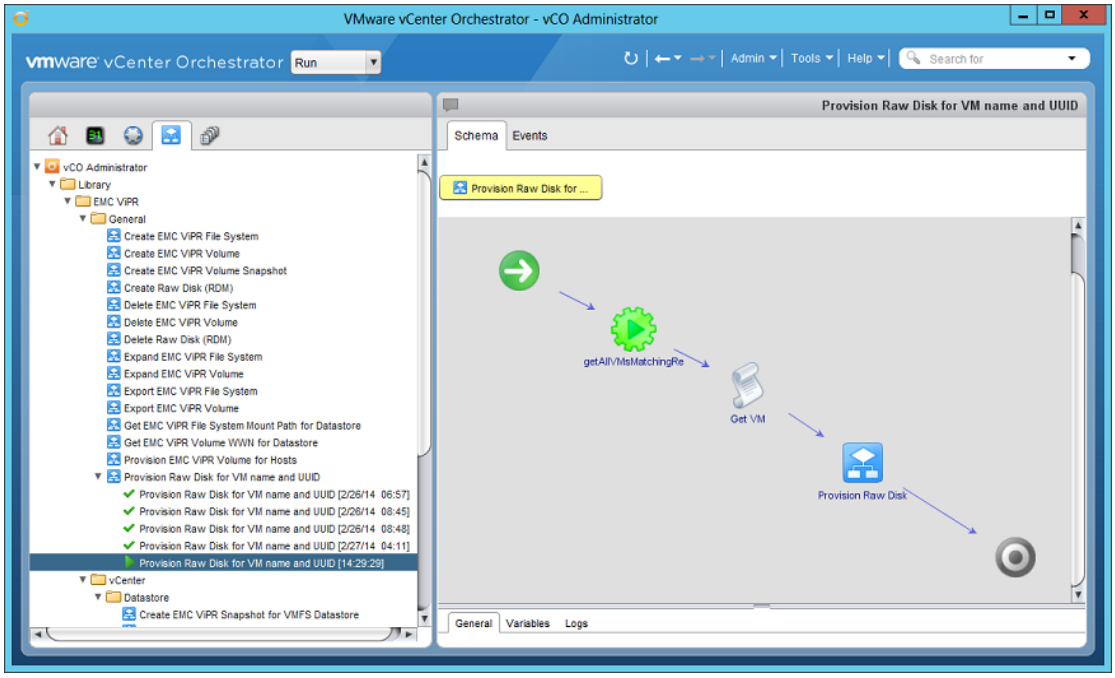

Figure 13: The EMC ViPR plugin for vCO enables VMware vCenter Orchestrator to perform automated storage provisioning tasks.

See Appendix G for details on how we installed vCO and executed storage provisioning workloads.

VMware vCloud Automation Center

For environments with rapid provisioning needs, such as virtualized cloud environments, EMC ViPR can be leveraged by VMware vCloud Automation Center (vCAC) to enable users to provision and deploy multiple VMs with customized storage needs. We set up vCAC in our test environment to show how easily end users can leverage the ViPR storage automation.

Organizations that leverage vCAC within their virtual environments will greatly benefit from the ability to perform repetitive tasks with minimal data entry. Figure 14 shows the limited information required to execute and deploy new VMs once an administrator has developed and enabled a blueprint that automates the workflow.

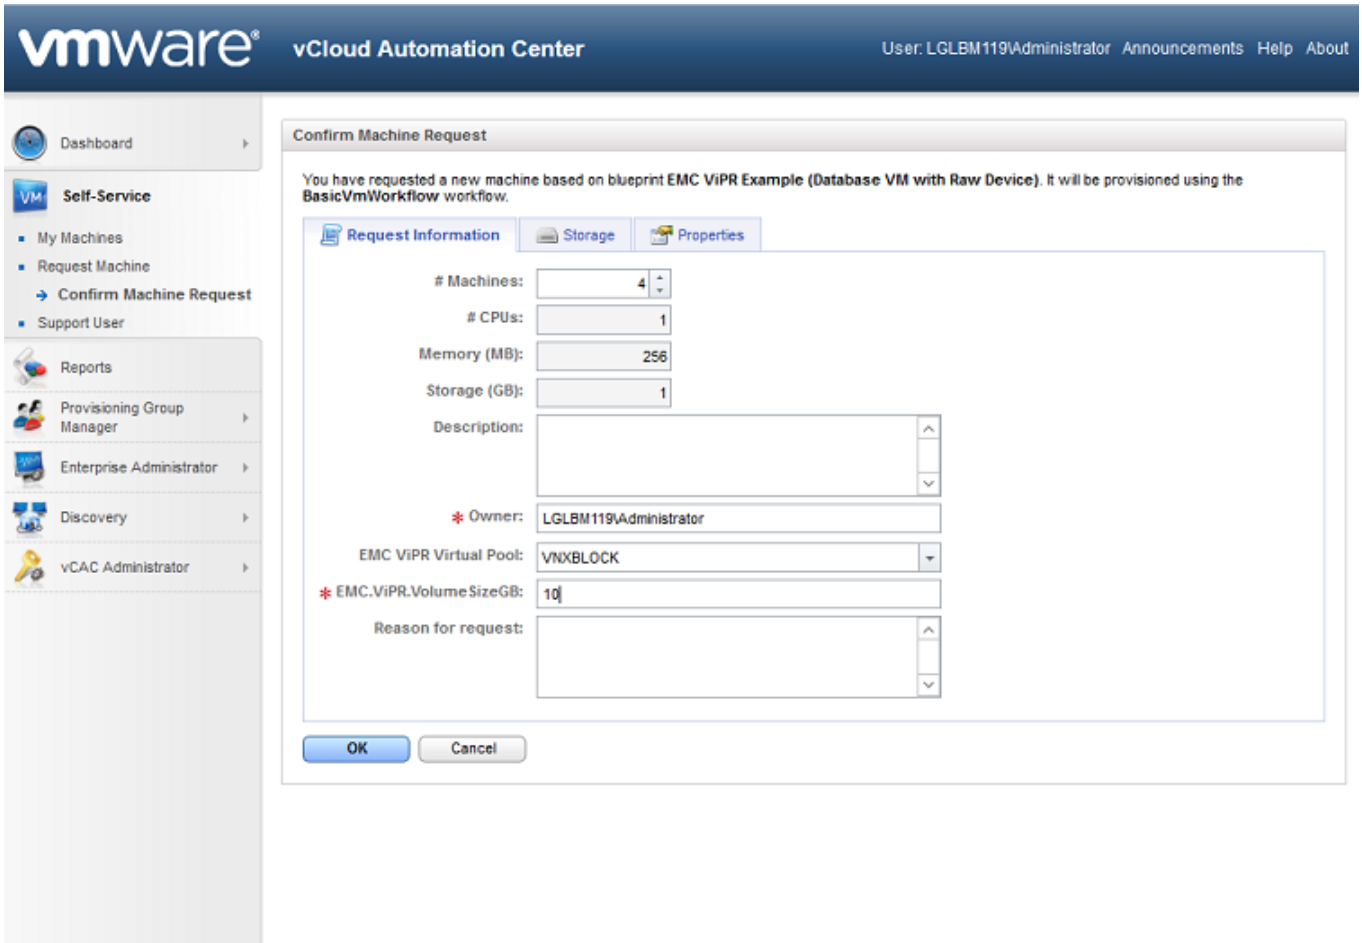

Figure 14: VMware vCloud Automation Center with ViPR integrations allows users to self-provision multiple new VMs with access to attached dedicated block storage.

Leveraging some of the same workflow we used in vCO, we were able to deploy multiple new VMs, each with its own block storage volume. Clicking OK submitted the requests, and new entries showed up in the Self-Service → My Machines section. Once completed, the new machines powered on and were available for use—each with its own attached block storage volumes. Figure 15 shows the machines as they are being provisioned.

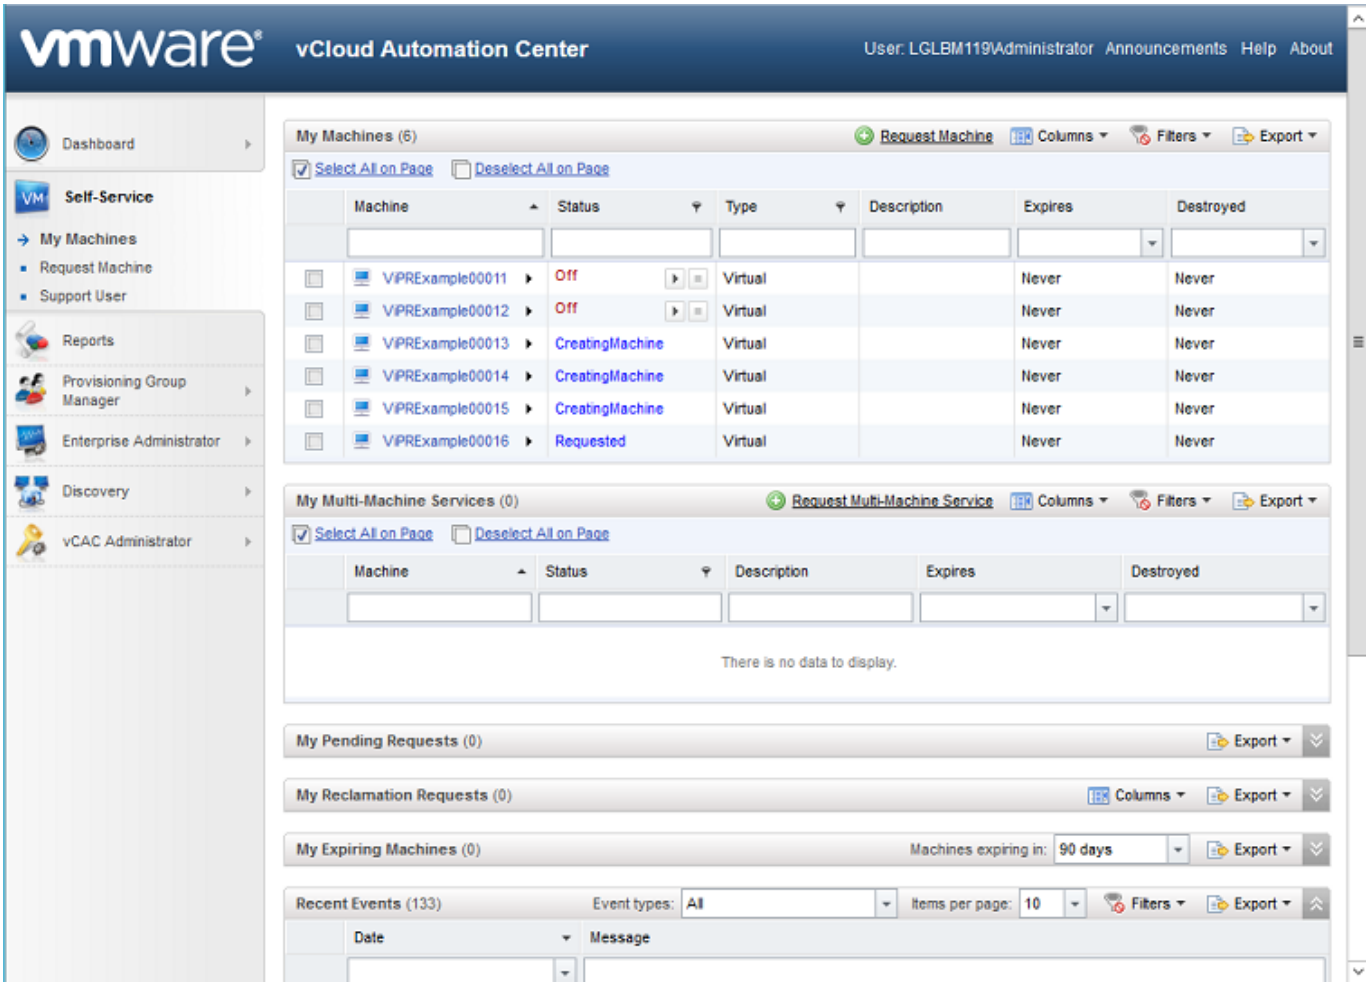

Figure 15: New virtual machines with ViPR-provisioned storage automatically appeared in our bucket of virtual machines.

It is worth taking a moment to go into greater detail about the automation that occurred in our vCAC demonstration, and how tightly all the software components and integrations worked together. In our example, our administrator ordered four new machines using a preconfigured blueprint in vCAC. The blueprint called for the creation of VMs, and then passed the UUID and hostname information to the workflow defined in vCO to provision and map Raw Disk storage to the newly created VMs. Further, this was done within vCAC Self-Service, which means the entire process was executed with neither a VMware vSphere administrator nor an enterprise storage administrator.

This translates to a major benefit of ViPR—users do not have to wait for VM provisioning, storage provisioning, workflow hand-offs, change approvals, or any of the other delays that are endemic to large enterprise organizations. The potential time savings can mean many things—faster development cycles, improved customer response times, and less temptation to spend money on a public cloud provider, which could lead to project cost overruns.

VMware vCenter Operations Manager

EMC ViPR Analytics Pack for vCenter Operations Management Suite extends vCOPS monitoring to include ViPR-supported storage.

VMware vCenter Operations Manager (vCOPs) is the monitoring solution vSphere administrators use to gather performance, health, and capacity metrics for their virtual environments. By integrating the EMC ViPR Analytics Pack for VMware vCenter Operations Management Suite, we were able to extend the vCOPs monitoring capabilities to include these same metrics for our ViPR-supported storage. In complex virtual environments, early warnings about potential problems with the underlying storage is a big advantage in that issues can be mitigated before they cause significant performance impact to customer virtual machines.

Figure 16 shows the EMC ViPR Performance view, which indicates heavy workloads, general health, capacity, and utilization of the subsystems managed within ViPR. These views can be modified to provide a customized dashboard for administrators who want to gather the most relevant data in a single glance.

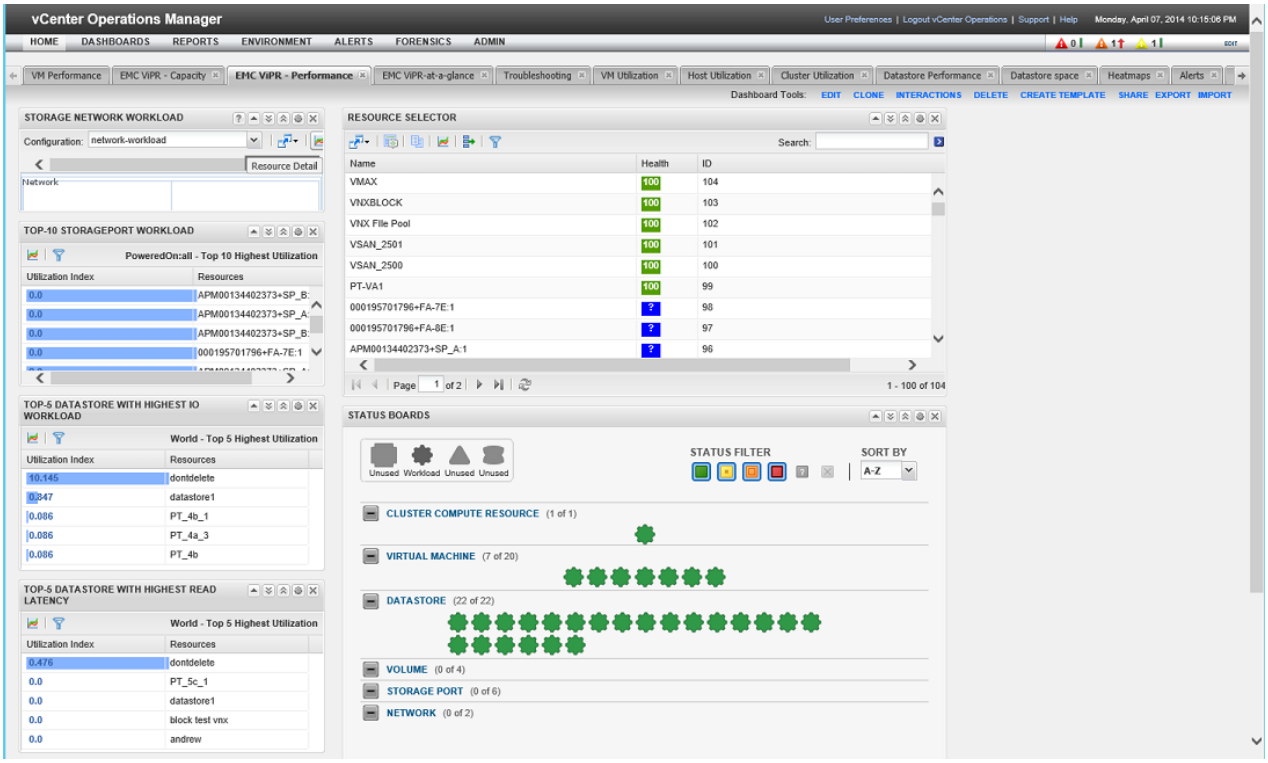

Figure 16: ViPR integrated vCenter Operations Manager provides health, capacity, and utilization data for ViPR-supported resources.

Manageability

Another aspect of ViPR VMware integration is the ability to use ViPR to create and mount volumes within vCenter to house VMs. A ViPR administrator simply adds the vCenter as a physical asset and its network interfaces to the appropriate storage networks. Then, within ViPR, a virtualization administrator can choose to both provision new storage **and** mount the new storage as a datastore. This eliminates the need for multiple hand-offs between the two different types of administrators and enables rapid provisioning within vCenter.

Let's briefly consider the alternative: the silo approach. We will numerate the potential hand-offs to give an example of the kinds of delays ViPR can help you avoid.

1. A customer makes a request to a VMware administrator stating they need one or more VMs of a certain size.
2. The VMware administrator realizes the environment does not currently have sufficient storage to meet the needs of the customer, and must generate a request for new storage.
3. Request analysts assign the first request to the storage administrator.
4. The storage administrator must determine which storage platform meets the needs of the customer making the request, based on the information contained within the service request. The storage administrator identifies the storage, and provisions disk pools and LUNs for the vSphere environment.
5. The storage administrator makes a new request to the storage network managers to add zoning rules to map the new storage to the target hosts.
6. Request analysts assign the new request to the storage network management group.
7. The storage network management group realizes this request constitutes a change for a production environment and generates a request for change.
8. The request for change is reviewed by the change-control advisory group, is approved, and is scheduled for a time when the risk of potential impact is lowest. The next available routine change window is two nights from now.
9. The change is executed by the storage network management group.
10. The change is reviewed by someone other than the executors to ensure no impact to the environment.
11. The storage network managers notify the storage administrator that the request has been fulfilled.
12. The storage administrator notifies the VMware administrator that their request was fulfilled.

13. The VMware administrator attempts to discover the new storage in vCenter, finds no new storage presented, and must open an incident ticket to troubleshoot the issue.

We will stop here and assume we have made the point. In a change-managed operational environment divided into functional silos, a simple change involves numerous hand-offs and delays that are built into the system even when everything goes according to plan. With VMware vSphere and EMC ViPR, datacenter virtualization is unified and provisioning can be implemented without lengthy delays and complicated change processes, allowing organizations to meet the needs of their customers as quickly as possible.

Figure 17 shows an example of the type of self-service offerings available to VMware administrators within the ViPR console.

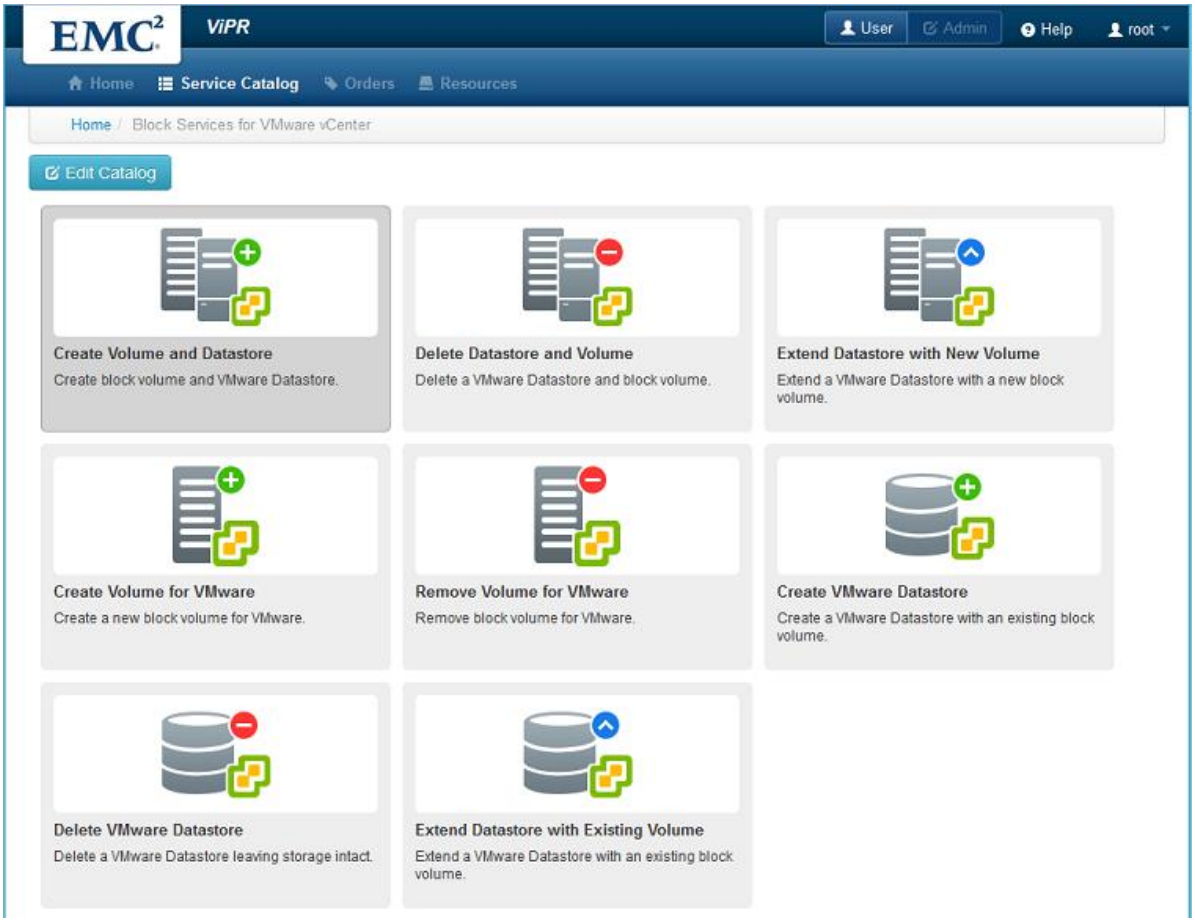

Figure 17: The service catalog provides the user self-service menu with options available to a virtualization administrator.

We compared ViPR's automated method with manually creating new storage, and making that storage ready for VM deployment. As Figure 18 shows, compared to the manual approach on both VMAX and VNX, using ViPR to provision block storage for VM reduced the number of steps by up to 54.0 percent and reduced time by up to 63.7 percent.

Provision block storage for VM	Number of admin steps	ViPR reduced steps by…	Admin time (mm:ss)	ViPR reduced time by…
ViPR	23		02:10	
Manual - VMAX	50	54.0%	05:58	63.7%
Manual - VNX	43	46.5%	04:00	45.8%

Figure 18: Time and number of steps to perform the block storage for VM provisioning task with and without ViPR. Fewer steps and less time are better.

Naturally, reductions in time and effort contribute to potential operational costs.

We performed the same tests on file storage systems. As Figure 19 shows, compared to the manual approach on VNX, Isilon, and NetApp, using ViPR to provision file storage for VM reduced the number of steps by up to 39.5 percent and reduced time by up to 31.2 percent.

Provision file storage for VM	Number of admin steps	ViPR reduced steps by…	Admin time (mm:ss)	ViPR reduced time by…
ViPR	23		02:19	
Manual - VNX	30	23.3%	03:17	29.4%
Manual - Isilon	38	39.5%	03:22	31.2%
Manual - NetApp	33	30.3%	03:11	27.2%

Figure 19: Time and number of steps to perform the file storage for VM provisioning task with and without ViPR. Fewer steps and less time are better.

Once again, it is important to note that we used the same procedure for all the storage types—we used the Service Catalog in user mode to provision both block and file storage. We provided the information ViPR required to create the requested storage, and allowed automation to perform the rest. Once ViPR indicated the storage was ready, we simply logged into vSphere and built the new VMs on the newly provisioned storage. Compare that with having to log in to the various storage arrays (each with a different interface), provision the types of storage needed, document the output, log in to vCenter, and manually discover the storage and create a new datastore—all before performing the simple task of creating a new VM.

Again, these reductions in time and effort translate into cost savings. Organizations need to apply the time savings we quantified to their own financial models, taking into account costs associated with a multi-step manual process and the inevitable downtime that will result from human error to arrive at operational cost savings.

See Appendix I for details on how we created datastores on block storage using ViPR and using manual methods. See Appendix J for details on how we performed those tasks on file storage.

Additional information

If VMware vSphere were the only virtualization platform supported by EMC ViPR, the features and integrations would still be impressive. However, in an effort to bring the benefits of software-defined storage to a wider audience, EMC ViPR has integrations for Microsoft® Hyper-V® and OpenStack® cloud environments.

For Microsoft Hyper-V environments, the EMC ViPR integration provides the ability to provision and mount new volumes to hypervisor clusters, expand those volumes, and delete them. Additionally, ViPR can present volumes directly to VMs within the environment and expand those volumes as storage is consumed. For more information about EMC ViPR integration with Microsoft Hyper-V, see the following links:

- ViPR Community: ViPR Add-in for Microsoft System Center Virtual Machine Manager (SCVMM) Overview
- Technical Documentation: EMC ViPR Add-in for Microsoft System Center Virtual Machine Manager

In OpenStack implementations, EMC ViPR supports the iSCSI and Fibre Channel storage connected to Cinder—a core OpenStack project related to block storage—and can create and host Swift-compatible object stores. Swift is the OpenStack project specifically for cloud-based object stores. For more information about EMC ViPR integration with OpenStack Cinder, see the following link:

- Rethink Storage blog: ViPR and OpenStack Integration: How It Works

With RESTful APIs for customers and other vendors to integrate their solutions with ViPR, EMC ViPR brings the potential for software-defined storage to just about any environment.

Object capabilities

As organizations look for ways to rapidly develop Web-based applications, cloud-based object storage is an attractive resource, because it is inexpensive, distributed, and accessible using normal Internet protocols. As developers build applications that utilize cloud-based storage, they may find difficulties in "porting" those applications to secured environments that use traditional storage models.

EMC ViPR provides organizations with an easy method of bringing their data and applications in house. ViPR can be used to create object data stores, protected with key-based authentication compatible with the storage models used by major cloud storage

providers. Within ViPR, storage administrators can create object data stores that overlay on file storage arrays, allowing traditional storage technology to be used with the latest in Web-based storage models. Figure 20 shows how easy it is to create new object-based data stores within ViPR.

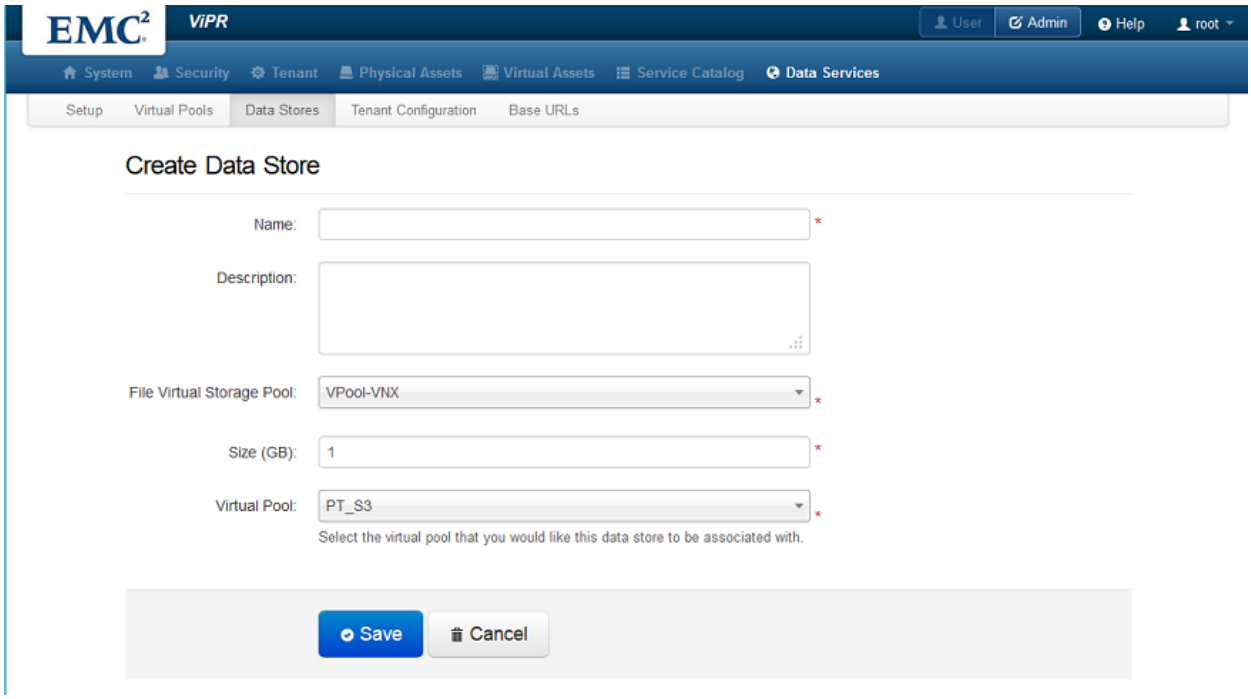

Figure 20: ViPR enables users to create object data stores using the same familiar interface and methods used to create block- and file-based volumes.

Once created, access to objects placed within the data stores requires keys to authenticate and retrieve objects. Web applications may store these keys and transmit them automatically in order to seamlessly access data. Unlike how it supports block and file, where it stays out of the data path, EMC ViPR places itself in the data path for object, creating a URL and key-based authentication model for objects stored on file storage arrays behind ViPR. Figure 21 shows an example set of keys that can be used to access data stores in ViPR.

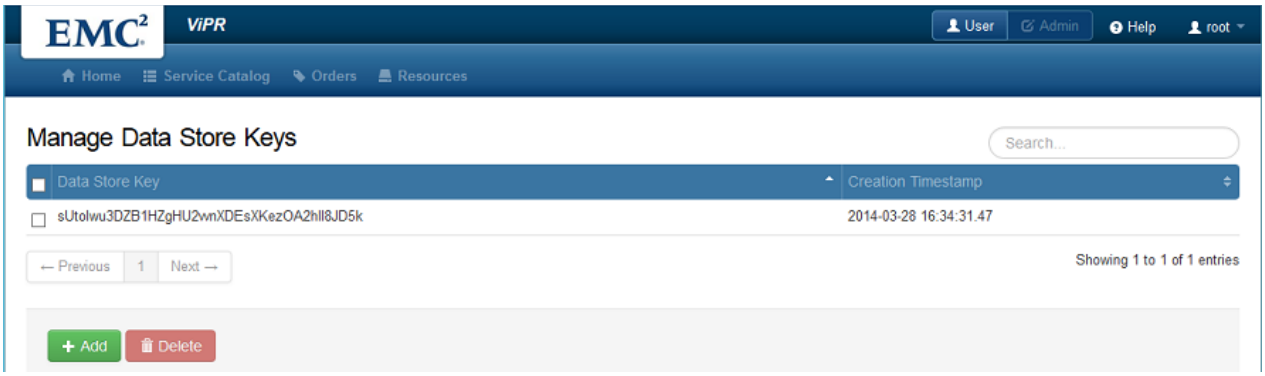

Figure 21: Data Store Keys are used to access the object data stores created in ViPR.

… your Web-based applications that utilize cloud storage take advantage of the data stores located on premises.

Once created, with just minimal reconfiguration, your Web-based applications that utilize cloud storage take advantage of the data stores located on premises. This allows you to bring on site and secure data that was previously stored in a public cloud. We tested this approach by using an S3-compatible browser to access object data stores housed within a ViPR system. Figure 22 shows how we could access objects contained in our ViPR object stores using a browser designed for cloud-based object store access.

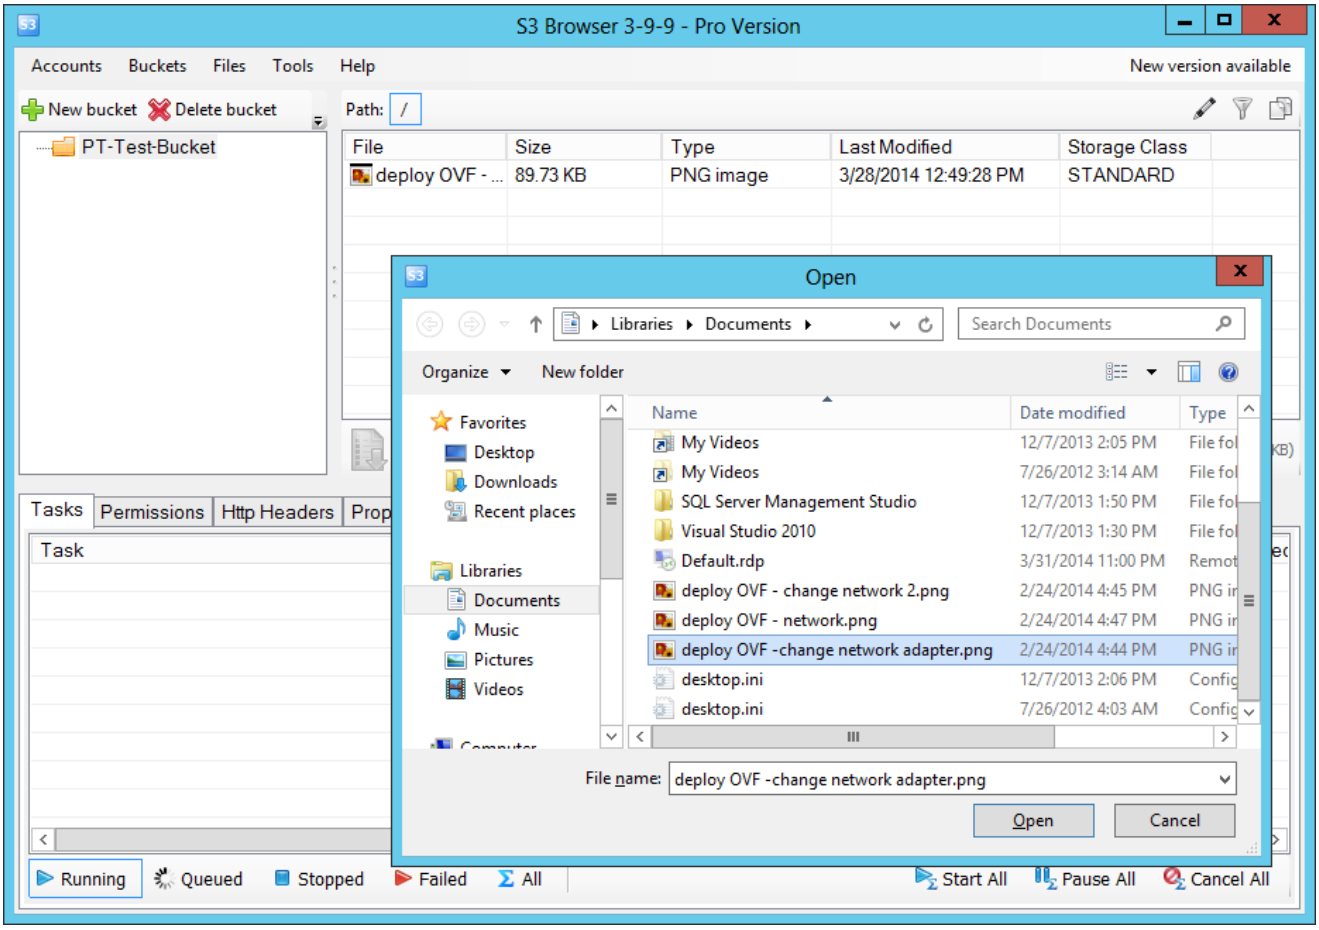

Figure 22: EMC ViPR can create S3 and Swift compatible object stores for use with apps configured to use cloud storage.

See Appendix K for details on how we configured Data Services and accessed data stores using a cloud-based application.

CONCLUSION

By automating routine yet complicated tasks, ViPR saves time and administrative effort and reduces the risk of human error. Built for the cloud, ViPR creates the ability for users to self-service provision storage from a diverse storage-array portfolio, increasing operational efficiency while enhancing organizational responsiveness and flexibility.

The benefits of ViPR are clearly demonstrated in its tight integrations with VMware. Virtual workflow automation, rapid provisioning, and enhanced monitoring of the ViPR environment could greatly enhance a virtualization administrator's ability to proactively manage their environment. Organizations that have come to appreciate the benefits of virtualized computing can now apply the same approach to storage and manage it through their VMware tools. EMC ViPR includes many plugins and extension packs allowing third-party monitoring and reporting tools to effectively provide detailed views of environmental conditions within a ViPR-supported system. Finally, EMC ViPR brings object storage capabilities to organizations by overlaying an object target on top of managed storage platforms.

An application that brought even one of these features to an organization would be a valuable addition to a software portfolio. With all of these features, and a growing catalog of supported storage platforms, it is easy to see how EMC ViPR can bring significant value to the datacenter.

For more information about EMC ViPR, including access to a no-charge fully-functional ViPR software download, see the following links:

- EMC ViPR Community
- EMC ViPR Download

APPENDIX A – INSTALLING AND CONFIGURING EMC VIPR 1.0

Deploying EMC ViPR

1. Connect to a vCenter using administrator credentials.
2. Select the cluster or host that will run the ViPR appliance.
3. Select File → Deploy OVM Template.
4. Browse to the location of the vipr-1.0.0.8.103-controller-XXX.ovf file. We selected vipr-1.0.0.8.103-controller-2+1.ovf. Click Open.
5. Click Next.
6. Review the template information, and click Next.
7. Click Accept to accept the terms of the license agreement, and click Next.
8. Provide a name for this ViPR installation or accept the default, and click Next.
9. Select the destination storage for the installation, and click Next.
10. Select Thin Provisioning, and click Next.
11. Select the correct destination network to attach to the ViPR appliance. We used the network designated as 192.168.1.X. Click Next.
12. Provide IP addresses within the network you selected for each of the ViPR servers in the appliance, plus one for client access, and the network mask and gateway. We selected 192.168.1.110, 192.168.1.111, 192.168.1.112, and 192.168.1.113. We used 255.255.255.0 as the network mask and 192.168.1.1 as the gateway.
13. Provide the IP addresses for your DNS and NTP servers, and click Next.
14. Review the installation information, and click Finish to deploy the ViPR vApp.
15. Select the newly deployed vApp and press the Play button located above the inventory list to power on all of the servers in the vApp.

Performing the initial configuration

1. Open a Web browser and connect to the client access address provided in the previous steps. We used 192.168.1.113.
2. Log on with the default root credentials (root / ChangeMe). For the purposes of our lab, we did not change the default password.
3. Enter the System Passwords account. We used `ChangeMe`. Click Next.
4. For ConnectEMC, change the transport setting to None.
5. Click Finish.
6. Scroll up, and click Save.
7. When prompted, click browse to locate the license file.
8. Locate the license file, and click Open.
9. Click Upload License file.
10. In the upper-right corner of the page, click Admin.
11. Click Tenant.
12. In Projects, click Add to create a project.
13. For Name, provide the name for a default project. We used `PT Test Project`
14. Click Save.

APPENDIX B – IDENTIFYING, REGISTERING, AND DISCOVERING BLOCK STORAGE IN VIPR

These procedures assume the back-end storage arrays are pre-configured and all VMAX gatekeeper LUNs are being presented to the SMI-S provider.

Setting up an SMI-S provider

1. Log in to support.emc.com and download the SMI-S provider for the operating system of a server connected to the SAN fabrics you wish to add to ViPR. Version 4.6.1.1 or higher is supported in ViPR 1.0. We downloaded SMI-S Provider 4.6.1.1 for SMI-S 1.5 for Linux.
2. Copy the file you downloaded to your host. We copied the file into a temp directory located in the home directory for our root user.
3. Open a terminal session.
4. Unzip the file. On our RHEL 5 Linux host, we entered the command `tar -xvf se7611-final-Linux-i386-SMI.tar.gz`
5. Install the SMI-S provider by typing `./se7611_install.sh -install`
6. Change the directory to /opt/emc/ECIM/ECOM/bin and execute `./ECOM -d`
7. Execute `./TestSmiProvider`
8. Accept all the default values to reach the command menu.
9. Type `disco` to discover VMAX storage attached to the SMI-S provider.
10. Type `addsys` and press Enter to add VNX block storage.
11. Type `y` to add a system. Type `1` and press Enter.
12. Enter the IP address of the VNX storage processor and press Enter.
13. Press Enter again.
14. Type `2` and press Enter.
15. Type `y` to find peer IP addresses and press Enter.
16. Enter the admin credentials for the VNX storage processor. We used `sysadmin`
17. Type `dv` to display the version information. This will include the arrays configured within the SMI-S provider.
18. Add a SAN Fabric Manager – see Appendix D for details.

Discovering block storage

1. Log in to the EMC ViPR console.
2. In the upper right of the page, click Admin.
3. Click on Physical Assets.
4. Under Physical Assets, click SMI-S Providers.
5. Click Add.
6. For Name, enter the name for the SMI-S provider you configured in the previous steps.
7. For Host, enter the IP address of the SMI-S provider you configured in the previous steps.
8. For User, enter `admin`
9. Enter `#1Password` for the Password and Confirm Password fields. Click Save.
10. Under Physical Assets, click Storage Systems. The SMI-S provider automatically presents the connected storage systems.

Assigning block storage to virtual pools

1. Log in to the EMC ViPR console.
2. In the upper right of the page, click Admin.
3. Under Virtual Assets, click Virtual Arrays.
4. Click Add to create a new virtual Array.
5. For Name, provide a name for the new virtual array. We used `PT-VA0`
6. Click Save.
7. Scroll down to the bottom of the Assign Networks to Virtual Array page, and select all the Fibre Channel networks available.
8. Click Assign Networks.
9. Under Virtual Assets, click Virtual Pools.
10. Click Add to create a new virtual pool.
11. Provide a name for the virtual pool. We used `PT-VPool0`
12. Check the box for PT-VA0 under Virtual Arrays. Wait for the storage to be located.
13. For Pool Assignment, use the pull-down menu to change the selection to Manual.
14. Clear the checkboxes from the storage pools you do not wish to have assigned to this virtual pool.
15. Scroll down, and click Save to define the virtual pool.

Enabling Data Protection services

1. Select Virtual Pools.
2. Click PT-VPool0.
3. Clear the checkbox for Expandable.
4. Scroll down to the Data Protection section.
5. Change the value for Maximum Native Snapshots to 1.
6. Change the value for Maximum Native Continuous Copies to 1.
7. Click Save.

APPENDIX C – IDENTIFYING, REGISTERING, AND DISCOVERING FILE STORAGE IN VIPR

Identifying, registering, and discovering file storage

1. Log in to the EMC ViPR console.
2. In the upper right of the page, click Admin.
3. Click on Physical Assets.
4. Select Storage Systems.
5. For Type, use the pull-down menu and select EMC VNX File.
6. For Name, type `PT-VNX-File`
7. Enter the IP address of the VNX control station. We used 192.168.1.11.
8. Accept the default port of 443. Provide the credentials of the administrator of the VNX file storage system to add. We used `nasadmin`
9. Scroll down and provide the SMI-S provider user credentials. We used `nasadmin`
10. Click Save to begin registration and discovery of VNX file storage for the identified array.
11. Click Virtual Assets.
12. Click the Add button to add a virtual array.
13. Provide a name for the virtual array. We used `PT-VA1`
14. For SAN Zoning, accept the default Automatic, and click Save.
15. Provide a name for the IP network associated with the virtual array. We used `PT-VA1-Net1`
16. Scroll down to the Add Ports to Network section, and check the box beside PT-VNX-File. Click the Add button.
17. Scroll down again to the Add Ports to Network section, and click the tab for Ports From hosts.
18. Check the boxes for all hosts that may have access to this virtual array. Click Add.
19. Click Save.

Assigning file storage to virtual pools

1. Under Virtual Assets, click Virtual Pool.
2. Click Add.
3. Provide a name for the virtual pool. We used `PT-VPool1`
4. Provide a description for the storage pool
5. For storage type, use the pull-down menu and select File.
6. For Provisioning Type, use the pull-down menu and select Thin.
7. For Virtual Arrays, check the box beside PT-VA1 to assign the virtual array to the storage pool.
8. Scroll down, and click Save.

Enabling Data Protection services

1. Under Virtual Assets, select Virtual Arrays.
2. Click the PT-VA1 virtual array.
3. Under Virtual Pools, select PT-VNX-FP.
4. Scroll down to Data Protection.
5. Change the value for Maximum Native Snapshots to 1.
6. Click Save.

APPENDIX D – SETTING UP SAN ZONING

Adding SAN Fabric Managers in ViPR

SAN zoning occurs automatically with ViPR. The steps below show how to add fabric managers as physical assets with ViPR. These are not zoning procedures.

1. Open a Web browser and connect to the EMC ViPR system.
2. Log in with administrative credentials. We used `root`
3. In the upper right of the page, click Admin.
4. Click on Physical Assets.
5. Under Physical Assets, click Fabric Managers.
6. Click Add.
7. For Type, select the storage switch to manage. We selected Cisco MDS.
8. For Name, provide a descriptive name for the switch. We used `MDS-A`
9. Enter the IP address of the storage switch.
10. Provide the administrative credentials for the switch. We used `admin`
11. Click Save.
12. Click Add.
13. For Type, select the storage switch to manage. We selected Cisco MDS.
14. For Name, provide a descriptive name for the switch. We used `MDS-B`
15. Enter the IP address of the storage switch.
16. Provide the administrative credentials for the switch. We used `admin`
17. Click Save.

Manual SAN Zoning

The following procedures are used to show how we manually performed SAN zoning. Unlike the automated process in ViPR, these manual procedures will have to be executed each time new storage or hosts are connected.

1. Log in to the storage system.
2. Obtain wwn of a “front-side” physical Fibre Channel port connected to the storage network.
3. Log on to the target server host.
4. Obtain the wwns of the Fibre Channel HBAs connected to the storage network.
5. Open a terminal session, connect via SSH, and log in to the “A-side” storage switch.
6. Type `configure` and press enter.
7. At (config) prompt, type `zone name PT_Test_A vsan 2500`
8. At (config-zone) prompt, type `member pwwn {wwn of host 'a' port}`
9. At (config-zone) prompt, type `member pwwn {wwn of storage front-side port 1}`
10. At (config-zone) prompt, type `member pwwn {wwn of storage front-side port 2}`
11. Type `exit` to return to the config prompt.
12. Type `exit` to return to the exec prompt.
13. Type `copy run start`
14. Type `exit` to end the session.
15. Open a terminal session, connect via SSH, and log in to the B-side storage switch.
16. Type `configure` and press enter.
17. At (config) prompt, type `zone name PT_Test_B vsan 2500`

18. At (config-zone) prompt, type `member pwwn {wwn of host 'B' port}`
19. At (config-zone) prompt, type `member pwwn {wwn of storage port 3}`
20. At (config-zone) prompt, type `member pwwn {wwn of storage port 4}`
21. Type `exit` to return to the config prompt.
22. Type `exit` to return to the exec prompt.
23. Type `copy run start`
24. Type `exit` to end the session.

APPENDIX E – PROVISIONING BLOCK STORAGE

Performing this task using ViPR

Self-Service Provisioning Block Storage

1. Log in to the EMC ViPR console.
2. Click Service catalog.
3. Click Block Service for Windows.
4. Click Create and Mount Volume.
5. For Windows Host, use the pull-down menu to select a Windows Host within the ViPR system.
6. For Virtual Array, use the pull-down menu to select a virtual array with block storage available. We selected PT-VA0.
7. For Virtual Pool, use the pull-down menu to select a virtual pool associated with the virtual array. We selected PT-VPool0.
8. For Project, use the pull-down menu and select the project for which you are granting access. We accepted the default project "PT Test Project."
9. For Name, provide a name for the Volume. We used `PT-ViPR-WB1`
10. For size, enter a value in GB for the size of the volume. We used `2` GB.
11. Click Order.

Additional Self-Service: Creating a Block Storage snapshot

1. Click Service catalog.
2. Click Create Block Snapshot.
3. For Volume, use the pull-down menu to select the volume you wish to snapshot.
4. For Type, use the pull-down menu and select Local Array Snapshot.
5. Give the snapshot a descriptive name. We used `Snap1`
6. Click Order.

Performing this task manually on VMAX

Manually creating volumes from a physical pool - VMAX

1. Using a Web browser, log in to EMC Unisphere for VMAX.
2. On the home screen, click the array.
3. Click Storage→Volumes.
4. In the far right, under Common Tasks, click Create volumes.
5. For volume type, select Regular.
6. For Configuration, select Standard.
7. For Disk Technology, select FC.
8. For Protection, select RAID-5 (3+1).
9. For number of volumes, enter `1`.
10. For volume capacity, select 2 GB.
11. Beside Add to Job List, click the pull-down menu and select Run Now.
12. Capture the name of the created volume, and click Close.

Creating storage groups - VMAX

1. Click Storage→Storage Groups.
2. In the far right, under Common Tasks, click Create a Storage Group.
3. Enter the Storage Group name. We used the name of the server plus its first WWN. Click Next.
4. For Volumes Type, use the pull-down menu and select Regular Volumes. Click Next.

5. For Disk Technology, select FC.
6. For Protection, select RAID-5 (3+1).
7. For number of volumes, enter 4.
8. For volume capacity, enter 20 GB. Click Next.
9. Review the information, and click Finish.
10. When storage group creation has completed, click Close.

Assigning hosts and volumes using VMAX

1. Click Hosts.
2. In the far right, under Common Tasks, click Create a new host.
3. Enter the name for the new host. We used the name of the server.
4. For initiator, enter the first WWN of the host, and click Add.
5. In the initiator field, enter the second WWN of the host, and click Add.
6. Click Next.
7. In the Provisioning Storage Window, use the pull-down menu for Provision By, and select Use an existing Storage group.
8. Scroll down in the list below storage group, and find the group you created in the previous steps. Click Next.
9. Under Port Group Definition, for select Ports, select the FC ports you want to use for this assignment. Click Next.
10. Click Finish.
11. When completed, click Close.

Mount the volume on the target host

1. Log in to the Windows target host.
2. Open the computer management MMC.
3. Expand Storage and select Disk Management.
4. Right-click Disk Management and select Rescan Disks.
5. Disk management displays the new disks as offline. Right-click one of the disks and select Online.
6. Right-click the now online disk, and select Initialize.
7. Click OK to accept MBR and write a signature to the disk.
8. Right-click the Basic disk, and select New Simple Volume.
9. Click Next in the New Simple Volume Wizard.
10. Click Next to accept the defaults for size.
11. Click Next to accept the defaults for drive letter assignment.
12. Click Next to accept the defaults to format as NTFS.
13. Click Finish to complete volume creation.

Performing this task manually on VNX

Manually creating volumes from a physical pool - VNX

1. Using a Web browser, log in to the EMC Unisphere instance for the VNX array.
2. From the pull-down menu at the top of the dashboard, select the array to manage.
3. Click Storage→LUNs.
4. At the bottom of the page, click Create.
5. On the LUN creation window under LUN Properties, clear the checkbox for Thin.
6. For user capacity, enter 20 GB.
7. Under LUN Name, select Name and provide a new name for the LUN. We used PTBlockTest_VNX
8. Click Apply.
9. Click Yes to confirm LUN creation operation.

10. Click OK to confirm successful completion.
11. Click Cancel to close the window.

Creating storage groups - VNX

1. Click Hosts→Storage Groups.
2. Under Storage Group Name, click the Create button.
3. Provide a name for the new Storage Group. Click Ok.
4. Click Yes to confirm Storage Group Creation operation.
5. Click Yes to add hosts and LUNs to the storage group.

Assigning hosts and volumes using VNX

1. On the LUNs tab of the Storage Group Properties page, expand SPA.
2. Locate and select the LUN you created in the previous step, and click Add.
3. Click on the Hosts tab.
4. On the left panel, locate and select the unassigned host you wish to add to the storage group.
5. Click the button containing the right-pointing arrow to move the host to the right-panel.
6. Click OK.
7. Click Yes to confirm adding the host and LUN to the storage group.
8. Click OK to confirm successful completion.

Mounting the volume on the target host

1. Log in to the Windows target host.
2. Open the computer management MMC.
3. Expand Storage and select Disk Management.
4. Right-click Disk Management and select Rescan Disks.
5. Disk management displays the new disks as offline. Right-click one of the disks and select Online.
6. Right-click the now online disk, and select Initialize.
7. Click OK to accept MBR and write a signature to the disk.
8. Right-click the Basic disk, and select New Simple Volume.
9. Click Next in the New Simple Volume Wizard.
10. Click Next to accept the defaults for size.
11. Click Next to accept the defaults for drive letter assignment.
12. Click Next to accept the defaults to format as NTFS.
13. Click Finish to complete volume creation.

APPENDIX F – PROVISIONING FILE STORAGE

Performing this task using ViPR

Self-Service Provisioning File Storage

1. Log in to the EMC ViPR console.
2. Click Service catalog.
3. Click File Storage Services.
4. Click Create UNIX Share.
5. For Virtual Array, use the pull-down menu to select PT-VA1.
6. For Virtual Pool, use the pull-down menu to select PT-VPool1.
7. For Project, use the pull-down menu to select the project for which you are granting access. We accepted the default project "PT Test Project."
8. For Export Name, provide a descriptive name. We used `PT-EXFile1`
9. For size, enter a value in GB for the size of the Share. We used `2`.
10. For Export Hosts, enter the IP address of a server with access to the share. The host must have been already been provided with access to the virtual array by the administrator.
11. Click Order.

Self-Service Creating a File Storage snapshot

1. Click Service catalog.
2. Click File Protection services.
3. Click Create File Snapshot.
4. For File System, use the pull-down menu to select a file system to snapshot. We selected PTEXFile1.
5. For Name, provide a descriptive name for the snapshot. We used `FileSnap1`
6. Click Order.

Performing this task manually on VNX

1. Using a Web browser, connect to the IP address for the VNX Control Station. Change the scope to local and login as nasadmin.
2. Use the pull-down menu beside Dashboard, and select the array.
3. Select Storage→File System.
4. Click Create.
5. In the Create File System window, provide a File System Name. We used `PT_VNX_File_Test`
6. For Storage Capacity, enter a value in GB for the size. We used `2` GB.
7. Click OK.
8. Select Storage→Shared Folders→NFS.
9. Click Create.
10. For File System, select the File system you just created (PT_VNX_File_Test).
11. For Read/Write Hosts, enter the IP address of a server with access to the share. The IP address should be on a network segment with access to the data mover.
12. Click OK.

Performing this task manually on Isilon

1. Open a Web browser and connect to the IP address of the Isilon Array.
2. Log in to OneFS as root.
3. Click the File System Management tab.
4. Click File System Explorer
5. Select the /ifs folder under directories.
6. In the right panel, click Add Directory.
7. In the Directory Name field, provide a new name for the directory. We used `PTIsilonTest`
8. Click Submit.
9. Select the Protocols tab.
10. Click UNIX Sharing (NFS).
11. Click the link for Add an Export.
12. For Description, enter information describing the purpose of the export.
13. For Clients, add the IP address of the host you want to have access to the NFS share.
14. For Directory, click the Browse button to locate the directory for export.
15. Select the directory you just created (PTIsilonTest). Click Select.
16. For User/Group Mappings, use the pull-down menu to select Use custom.
17. Under Map to user Credentials, use the pull-down menu for Map these users and select All users.
18. Select Specific username and enter root.
19. Clear the checkbox for Group.
20. Click Save.

Performing this task manually on NetApp

1. Open the NetApp OnCommand System Manager.
2. Select the NetApp storage system you want to connect to, and click Login.
3. Enter root credentials, and click Sign in.
4. Expand Storage, and select Volumes.
5. Click Create.
6. For Name, enter the name of the volume you wish to create. We used `PT_NA_Vol`
7. For Aggregate, accept the default, or click the choose button to select the aggregate that will house the volume.
8. For Size, enter a value in GB for the size of the volume. We used `2` GB.
9. Click Create.
10. In the left pane, under storage, select Exports. The newly created volume is displayed with the Export path. Select the Export.
11. Under Client Permissions for Export, select the security policy, and click Edit.
12. In the Edit Export Rule Window under Security Flavor, clear the checkbox for UNIX.
13. Under Client Permissions, select All hosts, and click Edit.
14. Under client, enter the IP address of the host you want to have access to the export. Click Save.
15. Under Anonymous Access, select Grant root access to all hosts. Click Modify.

APPENDIX G – SETTING UP VCENTER ORCHESTRATOR (VCO)

Deploying vCO

1. On the virtual center, select a host in your cluster.
2. In the top menu, select File → Deploy an OVF template.
3. Browse to the location of the vco appliance OVF file.
4. Select vCO_VA-5.1.0.0-81795_OVF10.ovf, and click Open.
5. Click Next.
6. Review the OVF template details, and click Next.
7. Click the Accept button to accept the license agreement terms, and click Next.
8. Provide a name for the vCO appliance. Click next.
9. Select the destination datastore for the appliance, and click Next.
10. Select Thin Provision, and click Next.
11. Provide the default gateway, DNS, and network IP address for the appliance. Click Next.
12. Check the box for Power on after deployment, and click Next.
13. Open a virtual KVM connection to the vCO appliance.
14. Click on the KVM session and press Enter to Login.
15. Provide the root credentials `root` and `vmware` to access the command line interface.
16. Type `yast2` and press enter.
17. Using the arrow keys, select Network Devices→Network Settings. Press Enter.
18. Using the Tab key, select Hostname/DNS and press Enter.
19. Change the hostname to something meaningful.
20. Provide the domain name for the network.
21. Clear the checkbox for Change Hostname via DHCP.
22. Provide the name of at least one name server.
23. Add the network domain to the Domain Search list. Press F10 for OK.
24. Press Enter to select Network Settings.
25. Tab to Routing and press Enter.
26. Enter the default gateway for the network. Press F10 for OK.
27. Press Enter to select Network Settings.
28. In the Overview section, use the arrow keys to select the network adapter from the list. Press F4 to edit.
29. Use the arrow keys to select Statically assigned IP address. Press spacebar to select.
30. Provide the IP address and subnet mask for the vCO server.
31. Press F10 for Next.
32. Press F10 for OK.
33. Press F9 to quit.
34. Open a Web browser. Use https to connect to the IP address of the vCO. Use port 8283 (example: https://192.168.1.123:8283).
35. Enter the default credentials `vmware` with the password `vmware`. When prompted, change the password to something unique.
36. Select Plug-ins from the menu on the left of the browser window.
37. Scroll down and check the box for vCenter Server 5.1.0.446. Click Apply Changes.
38. Select Startup Options from the menu on the left of the browser window.
39. Click the link for Restart vCO Configuration Server.

40. Enter the username `vmware` and the password for that account. Press Login.
41. Select Network from the menu on the left of the browser window.
42. Select the tab for SSL Trust Manager.
43. Enter the vCenter server address to import a certificate. Example: https://192.168.1.118:7444. Click Import.
44. Review the certificate information, and click the Import link beneath it.
45. Click the Network tab.
46. Use the pull-down menu for IP address to select the IP address assigned to your vCO server.
47. Click Apply changes.
48. Select vCenter Server (5.1.0) from the menu on the left of the browser window.
49. Select the tab for New vCenter Server Host.
50. Enter the IP address for a vCenter server.
51. Enter the vCenter Admin user credentials. We used `root`
52. Click Apply changes.
53. Select Authentication from the menu on the left of the browser window.
54. Use the pull-down menu for Authentication mode and select LDAP authentication
55. Use the pull-down menu for LDAP client and select OpenLdap.
56. Click Apply changes.
57. If prompted by the browser click Remember Password.

Integrating ViPR with vCO

1. Download the ViPR 1.0 plug-in for VMware vCenter Orchestrator.
2. Click Plugins.
3. Scroll down to Install new plug-in. Click the button with the magnifier glass icon.
4. Browse to the location of the ViPR plugin. Select EMC-ViPR-vCO-Plugin-1.0.0.7.41.dar, and click Open.
5. Click Upload and Install.
6. In the left menu, click EMC ViPR Plugin (1.0.0).
7. Enter the IP address of the ViPR instance. (Example: 192.168.1.113).
8. Enter the ViPR username and password. We used the `root` account.
9. Click Verify Connection.
10. Click Apply changes.
11. Reboot the vCenter Orchestrator server.

Executing a ViPR integrated workflow – Provision Raw Disk for VM name and UUID

1. Open a Web browser. Enter the IP address of the vCO server and press Enter (example: 192.168.1.123).
2. Click the link for Start Orchestrator Client.
3. Select the correct application for launching java applications, and click OK. (example: Java™ Web Start Launcher).
4. Provide the appropriate user credentials for vCO. We used `vcoadmin`. Click Login.
5. At the top of the left panel, click the blue workflow icon.
6. Expand vCO Administrator→Library→EMC ViPR→General.
7. Select Provision Raw Disk for VM name and UUID.
8. In the top of the right panel, click the green triangle to start the workflow.
9. Enter the name of the virtual machine targeted for new storage.
10. Enter the UUID of the virtual machine targeted for new storage.
11. Enter the EMC ViPR storage volume name for the new storage.

12. Define the size of the new disk in GB as 1. Click Next.
13. Select the EMC ViPR virtual array that will house the new storage.
14. Select the EMC ViPR virtual pool name.
15. Click Submit. The workflow will display a graphical representation of the tasks as they execute. The final icon will change to a green bull's eye. A check will appear beside the workflow in the left pane when the task is complete. Verify the changes to the target VM by checking the configuration in vCenter.

APPENDIX H – SETTING UP VCLOUD AUTOMATION CENTER 5.1

Deploying vCAC 5.1

Preparing the host

1. Build Windows 2008 R2 Virtual machine. Information about required IIS modules and other Windows components, such as .NET Framework and PowerShell can be found at www.vmware.com/pdf/vcac-51-installation-guide.pdf.
2. Launch SQL Server 2008 R2 Express Installer.
3. Select new Installation or add features to an existing installation.
4. Select New installation or add shared features, and click Next.
5. Check the box for I accept the license terms, and click Next.
6. Check the box for Database Engine Services, and click Next.
7. Select Default Instance, and click Next.
8. Accept the defaults, and click Next.
9. Select Windows authentication mode, ensure the administrator account is specified as a SQL Server administrator, and click Next.
10. Accept the defaults, and click Next.
11. Click Finish to close the installer.

Installing vCAC Manager

1. Browse for the vCAC installer program. Double-click DCAC-Manager-Setup.exe.
2. In the VMware vCloud® Automation Center™ Setup Wizard, click Next.
3. Check the box for I accept the terms in the license agreement, and click Next.
4. Click browse to locate the filename of your license file. Select the license file, and click Open.
5. Click Next to begin installation.
6. Select the drop-down for Database and select Entire feature will be installed on local hard drive.
7. Click Next to continue.
8. Select the pull-down menu for the application Web site and choose Default Web Site.
9. Select HTTP as the binding protocol, and click Next.
10. Click Next to accept the database defaults.
11. Click Test Connection to validate database engine availability. Click Ok.
12. Accept all defaults, and click Next.
13. Select File based XML store, and click Next.
14. Enter the FQDN for the vCAC Web site Hostname.
15. Enter the SMTP server details and the From address. Click Next.
16. Enter the account credentials that will run the vCAC service. We used `LOCALHOST\administrator`. Click Next.
17. Enter the FQDN for the Model Manager Web Service Hostname. We used the FQDN of the local host. Click Next.
18. Enter the username and password for the vCAC Web portal. We used `LOCALHOST\administrator`. Click Next.
19. Click Install.
20. When installation has completed, click Finish to exit the installer.

Installing vCAC Distributed Execution Manager (DEM) orchestrator role

1. Browse for the vCAC DEM setup wizard. Double-click DCAC-Dem-Setup.exe.
2. Click Next to begin installation.
3. Check the box for I accept the terms in the license agreement, and click Next.

4. Enter the DEM instance name. We used *hostname*-dem-orch
5. Enter a DEM description. We used orchestrator DEM
6. Select the Orchestrator Role.
7. Clear the checkbox for Use HTTPS. Click Next.
8. Click Next to begin installation.
9. Enter the FQDN and port number for the vCAC server Manager Service. We used FQDN:9003
10. Enter the FQDN and port number for the vCAC server Model Manager Web Service. We used FQDN:80
11. Enter the credentials for the Model Manager user. We used LOCALHOST\Administrator. Click Next.
12. Enter the credentials for the DEM user. We used LOCALHOST\Administrator. Click Next.
13. Click Install.
14. Click Finish to exit the installer.

Installing vCAC Distributed Execution Manager (DEM) worker role

1. Log in with administrator credentials.
2. Browse for the vCAC DEM setup wizard. Double-click on DCAC-Dem-Setup.exe.
3. Click Next.
4. Check the box for I accept the terms of the license agreement, and click Next.
5. Enter the DEM Instance Name. We used *localhost*-dem-work
6. Enter the description for the DEM.
7. Select Worker Role.
8. Clear the checkbox for Use HTTPS. Click Next.
9. Click Next.
10. Enter the Manager service Hostname and port. We used FQDN:9003
11. Enter the Model Manager Web service hostname and port. We used FQDN:80
12. Enter the Model Manager credentials. We used LOCALHOST\Adminstrator. Click Next.
13. Enter the DEM credentials to use. We used LOCALHOST\Administrator. Click Next.
14. Click Install.
15. Click Finish to exit the installer.

Creating an Endpoint in vCAC for vSphere

1. Open a browser window and connect to the vCAC Web interface. (Example: http://192.168.1.119/dcac)
2. Select vCAC Administrator → Endpoints.
3. Click New Endpoint.
4. Enter the name of the vCenter to define as an endpoint.
5. Provide the IP address for the vCenter server.
6. Click the button beside credentials to select the credentials to use for this endpoint.
7. Click New Credentials.
8. Enter the name used to identify this credential set. We used vCenter credentials
9. Enter the vCenter credentials for the new endpoint. We used root. Click OK.
10. If prompted by the browser, click Remember password.
11. Select the newly created credentials, and click OK.
12. Click OK to define the new endpoint.

Deploying the vCAC Agent to the vSphere server

1. Browse for the vCAC Agent setup wizard. Double-click DCAC-Agent-Setup.exe.
2. Click Next.
3. Check the box for I accept the terms in the License Agreement, and click Next.

4. Enter the name of the target vCenter for the Agent Name.
5. Clear the checkbox for use HTTPS.
6. Enter the IP address and port of the vCAC Server. (Example: 192.168.1.119:80)
7. Enter the IP address and port for the model manager Web service (Example: vcac.local.test:80)
8. Click Next.
9. Select vSphere Agent. Click Next.
10. Enter the credentials for the vCAC user. We used `LOCALHOST\administrator`. Click Next.
11. Enter the credentials for the model manager user. We used `LOCALHOST\administrator`. Click Next.
12. Enter the name of the endpoint you created in the previous steps. Click Next.
13. Click Install.
14. Click Finish to exit the installer.

Creating an Endpoint in vCAC for vCO

1. Select vCAC Administrator→Endpoints.
2. Click New Endpoint.
3. Enter the name of the vCO server to define as an endpoint.
4. Provide the IP address for the vCO server.
5. Click the button beside credentials to select the credentials to use for this endpoint.
6. Click New Credentials.
7. Enter name used to identify this credential set. We used `vCO credentials`
8. Enter the credentials for the vCO server. We used `vcoadmin`
9. Click OK.
10. If prompted by the browser, click Remember Password.
11. Select the vCO credentials you just created, and click OK.
12. Under Custom Properties, click the plus sign to create a New Property.
13. For the Property Name, use `VMware.VCenterOrchestrator.Priority`
14. For the Value, use 1
15. Click OK to define the new endpoint.

Creating an Enterprise Group

1. Click vCAC Administrator → Enterprise Groups.
2. Click the link located in the upper right of the page for New Enterprise Group.
3. Provide the name for the new Enterprise Group. We used `PT Enterprise Group`
4. Enter the username for the Enterprise Administrator. We used `LOCALHOST\Administrator`
5. Click OK.
6. Click PT Enterprise Group → Edit.
7. Check the box for the cluster resources displayed under compute resources.
8. Click OK.

Creating a Provisioning Group

1. Click Enterprise Administrator → Provisioning Groups.
2. Click New Provisioning Group.
3. Provide a name for the provisioning group. We used `PTPG-1`
4. Enter an email address for the group manager.
5. Click the button beside Default Machine Prefix to select or create a machine prefix.
6. Click New Machine Prefix.
7. Enter the machine prefix. We used `ViPRExample`
8. Provide the number of digits to include in the machine name. We used `5`.

9. Set the Next Number to 1.
10. Click OK.
11. Click the button beside Default Machine Prefix to select the newly created prefix.
12. Select ViPRExample. Click OK.
13. For the Active Directory Container, enter `cn=computers`
14. For each role, enter `LOCALHOST\administrator` as the user.
15. Click OK.

Creating a Reservation Policy

1. Click Enterprise Administrator→Reservation Policies.
2. Click New Reservation Policy.
3. Provide a name for the reservation policy. We used `PTRP1`
4. Click the check to save the policy.

Creating a Reservation

1. Click Enterprise Administrator→Reservations.
2. Click New Reservation→Virtual.
3. On the Reservation Information tab, for Compute resource use the pull down menu to select a resource.
4. Accept the default name for the reservation.
5. For Provisioning group, use the pull down menu to select a group.
6. For Reservation policy, use the pull down menu to select a policy.
7. Leave the Machine quota blank.
8. For Priority, select `1`.
9. Click on the Resources tab.
10. For Memory, select `64` GB.
11. For Storage, check the box beside the storage path you want to use for deploying VMs.
12. On the same line, set the value for This Reservation Reserved to `100`.
13. On the same line, set the value for Priority to `1`.
14. For Network, check the box beside the network you want to assign to your VMs.
15. Click OK.

Creating a Blueprint

1. Select Provisioning Group Manager → Blueprints.
2. Click New Blueprint → Virtual.
3. Provide a Name for the blueprint. We used `PTBP1`
4. For Group, use the pull down menu to select a group.
5. For Reservation policy, use the pull down menu to select a policy.
6. Check the box for Enabled.
7. Check the box for Master (copyable).
8. Set the Archive value to `2`.
9. Click the Build Information tab.
10. For Platform type, use the pull down menu to select vSphere (vCenter).
11. For Provisioning workflow, use the pull down menu to select BasicVmWorkflow.
12. In Machine Resources, set the # CPUs value to `1`.
13. Set the Memory (MB) value to `256`.
14. For Volumes, click Add Volume.
15. For Capacity (GB), change the value to `1`.

16. For Drive Letter/Mount Path, enter /data.
17. Click the check to accept the values.
18. Click the Properties Tab.
19. In Custom properties, click New Property.
20. For the property name, enter `VMware.VirtualCenter.OperatingSystem`
21. For the value, enter `sles64Guest`
22. Click the check to accept the values.
23. Click OK to complete the new blueprint.

Creating a Self-Service Machine

1. Select Self-Service→Request Machines.
2. Select the PTBP1 blueprint to request a machine.
3. Click OK.
4. Click Self-Service→My Machines to view the status of the creation job.

Integrating ViPR with vCAC

1. Download the ViPR 1.0 Enablement Kit for vCloud Automation Center from support.emc.com.
2. Download the VMware vCloud Automation Center 5.1 – Development Kit and associated license files from www.vmware.com.
3. Copy EMCViPREnablementKitforvCAC-1.0.zip, vCAC-51-DK-Installation.zip, and the Development Kit license files to a directory on the vCAC server.
4. Extract the contents of EMCViPREnablementKitforvCAC-1.0.zip.
5. Extract the contents of vCAC-51-DK-Installation.zip.
6. Double-click \vCAC-51-DK-Installation\Setups\CDK-Setup.exe.
7. Click Next.
8. Check the box for I accept the terms in the License Agreement. Click Next.
9. Click Browse to locate the license files downloaded for the vCAC Development Kit.
10. Select CDKlicense_5.1.XML. Click Open.
11. Click Next to validate the license.
12. Accept the installation defaults, and click Next.
13. Clear the checkbox for Use HTTPS.
14. Enter the IP address and port for the Model Manager Web Service. (Example 192.168.1.119:80)
15. Enter the credentials for the model manager. We used `LOCALHOST\administrator`. Click Next.
16. Click Install.
17. Click Finish to close the installer.

Installing EMC ViPR workflow

1. Open extracted EMCViPREnablementKitforvCAC-1.0 and copy EMC.ViPR.VMwithRDM.Example.XAML and External-EMC.ViPR.VMwithRDM.Example.XML to C:\Program Files (x86)\DynamicOps\DCAC Server\ExternalWorkflows\xmldb.
2. Copy EMC.ViPR.VMwithRDM.Example.XAML and External-EMC.ViPR.VMwithRDM.Example.XML to C:\Program Files (x86)\DynamicOps\Design Center.
3. Restart the VMware vCloud Automation Center service from the Windows Services mmc.
4. Open a PowerShell command prompt.
5. Change directories to C:\Program Files (x86)\DynamicOps\Design Center\.
6. Type `./CloudUtil.exe Workflow-Install –f EMC.ViPR.VMwithRDM.Example.xaml –n EMC.ViPR.VMwithRDM.Example`

Configuring vCAC to execute the ViPR workflow

1. Open the vCAC Web console.
2. In the left menu, select Enterprise Administrator→Property Dictionary.
3. Click Add Property Definition.
4. Enter `EMC.ViPR.VirtualPool` for the Name.
5. Enter `EMC ViPR Virtual Pool` as the Display Name.
6. Select DropDownList as the Control Type.
7. Check the box for Required.
8. Click the check to the far left of the data fields to accept the entries.
9. Click Edit in the far right under Property Attributes.
10. Click Add Property Attribute.
11. Use the pull-down menu under Type and select ValueList.
12. For the Name, enter `ViPR.VirtualPoolList`
13. For Value, enter the names of virtual pools configured in ViPR separated with commas.
14. Click the check at the far left of the data fields to accept the entries.
15. Click OK.
16. Select Enterprise Administrator→Build Profiles.
17. Click New Build Profile.
18. For Name, enter `EMC ViPR Provisioning`
19. In Custom Properties, click New Property.
20. For Name, enter `EMC.ViPR.VirtualPool`
21. For Value, enter the virtual pool that will be default. We used `VNXBLOCK`
22. Click the green check to the left of the fields to accept the values.
23. Click New Property.
24. For Name, enter `EMC.ViPR.VolumeNamePrefix`
25. For Value, enter a prefix for the ViPR created volumes. We used `ViPRvol_`
26. Click the green check to the left of the fields to accept the values.
27. Click New Property.
28. For Name, enter `EMC.ViPR.VolumeSizeGB`
29. For Value, enter `1`
30. Click the green check to the left the fields to accept the values.
31. Click OK to save the new build profile.

Creating a ViPR specific blueprint

1. In the left menu, select Provisioning Group Manager→Blueprints.
2. Click New Blueprint.
3. Click Copy from Existing Blueprint.
4. Select PTBP1.
5. For Name, enter `EMC ViPR Example (VM with Raw Device)`
6. Click on the Properties tab.
7. In the Build Profiles section, check the box for EMC ViPR Provisioning.
8. In the Custom properties section, click New Property.
9. For Name, enter `EMC.ViPR.vCenterRDMExampleTrigger`
10. For Value, enter `True`
11. Click the check to the left of the fields to accept the entries.

12. Click OK to save the ViPR blueprint.

Executing a ViPR integrated workflow – Creating a VM with Raw Disk Mapping

1. In the left menu, select Self-Service → Request Machines.
2. Under Request Machine, click EMC ViPR Example (VM with Raw Device).
3. Accept the defaults, and click OK.
4. Validate the job status by clicking Self-Service → My Machines or by viewing the configuration settings of the newly created VM in vCenter.

APPENDIX I – PROVISIONING BLOCK STORAGE FOR A VM

Performing this task using ViPR

Creating a Volume and a Datastore in vCenter

1. Log in to the EMC ViPR console.
2. Click Service Catalog.
3. Click Block Services for VMware vCenter.
4. Click Create Volume and Datastore.
5. For Name, provide a name for the datastore. We used `PT_ViPR_Block_vSphere`
6. For Datacenter, use the pull-down menu and select the datacenter you want to provision storage to. We selected PT ViPR Test.
7. For ESX host, use the pull-down menu to select a cluster member. We selected 192.168.1.252.
8. Select the virtual array you want to use for block storage. We selected PT-VA0.
9. Select the virtual pool you want to use for block storage. We selected PT-VPool0.
10. For Name, provide a description of the volume.
11. For Size, enter a volume size in GB. We entered `50` GB.
12. Click Order.

Creating a new VM on the new block storage

1. Log in to the vCenter.
2. Select the cluster. Right click, and choose New Virtual Machine.
3. Accept the typical configuration, and click Next.
4. Give a name for the virtual machine, and click Next.
5. Choose a host within the cluster, and click Next.
6. Scroll down the list of available volumes until you find the volume you created in previous steps. Our volume name was PT_ViPR_Block_vSphere. Click Next.
7. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next
8. Click Next.
9. Select the Virtual Disk size. We selected 20 GB.
10. Select Thin Provision, and click Next.
11. Click Finish to complete the virtual machine creation.

Performing this task manually on VMAX

1. Log in to EMC Unisphere for VMAX.
2. On the home screen, click on the array.
3. Click Storage→Volumes.
4. In the far right, under Common Tasks, click Create volumes.
5. For volume type, select Regular.
6. For Configuration, select Standard.
7. For Disk Technology, select FC.
8. For Protection, select RAID-5 (3+1).
9. For number of volumes, enter `1`.
10. For volume capacity, select 100 GB.
11. Beside Add to Job List, click the pull-down menu and select Run Now.
12. Capture the name of the created volume, and click Close.

13. Click Storage→Storage Groups.
14. In the far right, under Common Tasks, click Create a Storage Group.
15. Enter the Storage Group name. We used `PT_VMAX_ESX`. Click Next.
16. For Volumes Type, use the pull-down menu and select Regular Volumes. Click Next.
17. For Disk Technology, select FC.
18. For Protection, select RAID-5 (3+1).
19. For number of volumes, enter 1.
20. For volume capacity, enter 100 GB. Click Next.
21. Review the information, and click Finish.
22. When storage group creation has completed, click Launch Provision Storage Wizard.
23. In the Provision Storage window, for Host, use the scrolling list to select the ESX cluster initiator group.
24. For Provision By, use the pull down menu and select Use an existing Storage Group.
25. Scroll down in the list below storage group, and find the group you created in the previous steps. Click Next.
26. Under Port Group Definition, for select Ports, select the FC ports you want to use for this assignment. Click Next.
27. Click Finish.
28. When completed, click Close.
29. Log in to the vCenter.
30. Select a cluster member.
31. Click the Configuration tab.
32. Under Hardware, click Storage.
33. Click Add Storage.
34. Click Next.
35. Select the EMC Fiber Channel Disk in the list with the capacity of 100 GB. Click Next.
36. Select VMFS-5, and click Next.
37. Review the disk layout information, and click Next.
38. For datastore name, use `PT_VMAX_ESX`. Click Next.
39. Use the maximum space available, and click Next.
40. Click Finish.
41. Select the vSphere cluster. Right click, and choose New Virtual Machine.
42. Accept the typical configuration, and click Next.
43. Give a name for the virtual machine, and click Next.
44. Choose a host within the cluster, and click Next.
45. Scroll down the list of available volumes until you find the volume named PT_VMAX_ESX. Click Next.
46. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
47. Click Next.
48. Select the Virtual Disk size. We selected 20 GB.
49. Select Thin Provision, and click Next.
50. Click Finish to complete the virtual machine creation.

Performing this task manually on VNX

1. Using a Web browser, log in to the EMC Unisphere instance for the VNX array.
2. From the pull-down menu at the top of the dashboard, select the array to manage.
3. Click Storage→LUNs.
4. At the bottom of the page, click Create.

5. On the LUN creation window under LUN Properties, clear the checkbox for Thin.
6. For user capacity, enter 100 GB.
7. Under LUN Name, select Name and provide a new name for the LUN. We used PT_VNX_ESX
8. Click Apply.
9. Click Yes to confirm LUN creation operation.
10. Click OK to confirm successful completion.
11. Click Cancel to close the window.
12. Click Hosts→Host List
13. Locate the ESX host in the host list and double-click it.
14. Click the storage tab on the Host Properties window. At the bottom, determine the storage group it is a member of. Click Cancel.
15. Click hosts→Storage Groups
16. Locate the Storage Group for the ESX host and double-click it.
17. In the storage groups properties page, click the LUNs tab.
18. On the LUNs tab of the Storage Group Properties page, expand SP A and SP B.
19. Locate and select the LUN you created in the previous steps, and click Add.
20. Click Yes to confirm adding the host and LUN to the storage group.
21. Click OK to confirm successful completion.
22. Log in to the vCenter.
23. Select a cluster member.
24. Click the Configuration tab.
25. Under Hardware, click Storage.
26. Click Add Storage.
27. Click Next.
28. Select the DGC Fiber Channel Disk in the list with the capacity of 100 GB. Click Next.
29. Select VMFS-5, and click Next.
30. Review the disk layout information, and click Next.
31. For datastore name, use PT_VNX_ESX. Click Next.
32. Use the maximum space available, and click Next.
33. Click Finish.
34. Select the vSphere cluster. Right click, and choose New Virtual Machine.
35. Accept the typical configuration, and click Next.
36. Give a name for the virtual machine, and click Next.
37. Choose a host within the cluster, and click Next.
38. Scroll down the list of available volumes until you find the volume named PT_VNX_ESX. Click Next.
39. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
40. Click Next.
41. Select the Virtual Disk size. We selected 20 GB.
42. Select Thin Provision, and click Next.
43. Click Finish to complete the virtual machine creation.

APPENDIX J – PROVISIONING FILE STORAGE FOR A VM

Performing this task using ViPR

Creating a Volume and Datastore in vCenter

1. Log in to the EMC ViPR console.
2. Click Service Catalog.
3. Click File Services for VMware vCenter.
4. Click Create FileSystem and NFS Datastore.
5. For Name, provide a name for the datastore. We used `PT_ViPR_File_vSphere`
6. For Datacenter, use the pull-down menu and select the datacenter you want to provision storage to. We selected PT ViPR Test.
7. For ESX host, use the pull-down menu to select a cluster member. We selected 192.168.1.252.
8. Select the virtual array you want to use for block storage. We selected PT-VA1.
9. Select the virtual pool you want to use for block storage. We selected PT-VPool1.
10. For Export name, provide a name used for the NFS mount. We used `PT_ViPR_File_vSphere_export`
11. For Size, enter a volume size in GB. We entered `50` GB.
12. Click Order.

Creating a new VM on the new file storage

1. Log in to the vCenter.
2. Select the cluster. Right click, and choose New Virtual Machine.
3. Accept the typical configuration, and click Next.
4. Give a name for the virtual machine, and click Next.
5. Choose a host within the cluster, and click Next.
6. Scroll down the list of available volumes until you find the volume you created in previous steps. Our volume name was PT_ViPR_File_vSphere. Click Next.
7. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
8. Click Next.
9. Select the Virtual Disk size. We selected 20 GB.
10. Select Thin Provision, and click Next.
11. Click Finish to complete the virtual machine creation.

Performing this task manually on VNX

1. Using a Web browser, connect to the IP address for the VNX Control Station. Change the scope to local and log in as nasadmin.
2. Use the pull-down menu beside Dashboard, and select the array.
3. Select Storage→Storage Configuration→File Systems.
4. Click Create.
5. In the Create File System window, provide a File System Name. We used `PTVNXESX_NFS`
6. For Storage Capacity, enter a value in GB for the size. We used `100` GB.
7. Click OK.
8. Select Storage→Shared Folders→NFS.
9. Click Create.
10. For File System, select the File system you just created (PTVNXESX1).

11. For Read/Write Hosts, enter the IP addresses of servers with access to the share. The IP addresses should be on a network segment with access to the data mover.
12. Click OK.

Creating a new datastore and create a new VM on the new datastore

1. Log in to the vCenter.
2. Select a host that is a member of the cluster.
3. Click Add Storage.
4. Select Network File System. Click Next.
5. For Server, enter the IP address of the data mover defined on your storage array.
6. For Folder, enter the path to the newly created NFS share.
7. For Datastore name, enter the name displayed in vCenter for this datastore. We entered `PTVNXESX_NFS`
8. Click Finish.
9. Select the cluster. Right click, and choose New Virtual Machine.
10. Accept the typical configuration, and click Next.
11. Give a name for the virtual machine, and click Next.
12. Choose a host within the cluster, and click Next.
13. Scroll down the list of available volumes until you find the volume you created in previous steps. Our volume name was PTVNXESX_NFS. Click Next.
14. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
15. Click Next.
16. Select the Virtual Disk size. We selected 20 GB.
17. Select Thin Provision, and click Next.
18. Click Finish to complete the virtual machine creation.

Performing this task manually on Isilon

1. Open a Web browser and connect to the IP address of the Isilon Array.
2. Log in to OneFS as root.
3. Click the File System Management tab.
4. Click File System Explorer.
5. Select the /ifs folder under directories.
6. In the right panel, click Add Directory.
7. In the Directory Name field, provide a new name for the directory. We used `PTIsilonESX_NFS`
8. Click Submit.
9. Select the Protocols tab.
10. Click UNIX Sharing (NFS).
11. Click the link for Add an Export.
12. For Description, enter information describing the purpose of the export.
13. For Clients, add the IP address of the hosts you want to have access to the NFS share.
14. For Directory, click the Browse button to locate the directory for export.
15. Select the directory you just created (PTIsilonESX_NFS). Click Select.
16. For User/Group Mappings, use the pull-down menu to select Use custom.
17. Under Map to user Credentials, use the pull-down menu for Map these users and select All users.
18. Select Specific username and enter root.
19. Clear the checkbox for Group.

20. Click Save.

Creating a new datastore and create a new VM on the new datastore

1. Log in to the vCenter
2. Select a host that is a member of the cluster.
3. Click Add Storage.
4. Select Network File System. Click Next.
5. For Server, enter the IP address of the Isilon Storage.
6. For Folder, enter the path to the newly created NFS share.
7. For Datastore name, enter the name displayed in vCenter for this datastore. We entered `PTIsilonESX_NFS`
8. Click Finish.
9. Select the cluster. Right click, and choose New Virtual Machine.
10. Accept the typical configuration, and click Next.
11. Give a name for the virtual machine, and click Next.
12. Choose a host within the cluster, and click Next.
13. Scroll down the list of available volumes until you find the volume you created in previous steps. Our volume name was PTIsilonESX_NFS. Click Next.
14. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
15. Click Next.
16. Select the Virtual Disk size. We selected 20 GB.
17. Select Thin Provision, and click Next.
18. Click Finish to complete the virtual machine creation.

Performing this task manually on NetApp

1. Open the NetApp OnCommand System Manager.
2. Select the NetApp storage system you want to connect to, and click Login.
3. Enter root credentials, and click Sign in.
4. Expand Storage, and select Volumes.
5. Click Create.
6. For Name, enter the name of the volume you wish to create. We used `PTNAESX_NFS`
7. For Aggregate, accept the default, or click the choose button to select the aggregate that will house the volume.
8. For Size, enter a value in GB for the size of the volume. We used `25` GB.
9. Click Create.
10. In the left pane, under storage, select Exports. The newly created volume is displayed with the Export path. Select the Export.
11. Under Client Permissions for Export, select the security policy, and click Edit.
12. In the Edit Export Rule Window under Security Flavor, clear the checkbox for UNIX.
13. Under Client Permissions, select All hosts, and click Edit.
14. Under client, enter the IP address of the host you want to have access to the export. Click Save.
15. Under Anonymous Access, select Grant root access to all hosts. Click Modify.
16. Log in to the vCenter.
17. Select a host that is a member of the cluster.
18. Click Add Storage.
19. Select Network File System. Click Next.
20. For Server, enter the IP address of the NetApp Filer.

21. For Folder, enter the path to the newly created NFS share.
22. For Datastore name, enter the name displayed in vCenter for this datastore. We entered PTNAESX_NFS
23. Click Finish.

Creating a new datastore and create a new VM on the new datastore

1. Select the cluster. Right click, and choose New Virtual Machine.
2. Accept the typical configuration, and click Next.
3. Give a name for the virtual machine, and click Next.
4. Choose a host within the cluster, and click Next.
5. Scroll down the list of available volumes until you find the volume you created in previous steps. Our volume name was PTNAESX_NFS. Click Next.
6. Select the operating system for the guest VM. We selected Linux. We accepted the Red Hat Enterprise Linux 6 (64-bit) version. Click Next.
7. Click Next.
8. Select the Virtual Disk size. We selected 20 GB.
9. Select Thin Provision, and click Next.
10. Click Finish to complete the virtual machine creation.

APPENDIX K – CONFIGURING OBJECT STORE

Setting up Data Services

1. Log in to the EMC ViPR console.
2. Click Admin in the upper right of the page.
3. Click Data Services.
4. Select the IP network for the file storage. We used `PT-VA1-Net1`
5. Click Save. The following steps execute the instructions to complete the data services setup.
6. At the top menu bar, click System.
7. Click Configuration.
8. In the network section, click on the Data Service IP addresses field and enter an IP address for the data services node. We used 192.168.1.115. Click Save. Click OK to confirm reboot.
9. Download the config.iso as directed by the instructions in step 2.
10. Copy the config.iso file to the same directory as the ViPR dataservice OVF.

Deploying the data service node

1. Connect to the vCenter, and click File->Deploy OVF template.
2. Click browse to locate the vipr-1.0.0.8.103-datasservice.ovf file. Click Open.
3. Click Next.
4. Review the template details, and click Next.
5. Click Accept to accept the terms of the license agreement. Click Next.
6. Provide a name for the new data services node. We used `PTDS_115`. Click Next.
7. Select the destination storage for the virtual machine. Click Next.
8. Select Thin Provision, and click Next.
9. Select the destination network for mapping the data services node from the pull-down menu under Destination Networks. We used the network designated 192.168.1.X. Click Next.
10. Provide the IP address, netmask, and gateway for the data services node. We used 192.168.1.115 for the address, 255.255.255.0 for the netmask, and 192.168.1.1 for the gateway. Click Next.
11. Review the installation summary, and click Finish to begin installation.
12. Power on the VM after installation is complete.

Configuring S3 object store

1. In the EMC ViPR Web console, click Data Services.
2. Under Data Service, click Virtual Pools.
3. Click Add to create a virtual pool.
4. Enter the name of the new virtual pool. We used `PT_S3`
5. Provide the optional description, and click Save.
6. Click Data Stores.
7. Click Add to create a new data store.
8. Enter the name of the new data store. We used `PT_S3_DS1`
9. Enter a value for the size of the pool. We used 20 GB. Click Save.
10. Click Tenant Configuration.
11. Provide a name for the tenant namespace. We used `Tenant1`
12. Select the default Virtual Pool for this tenant. We selected PT_S3.
13. Select the default project for this tenant. We selected PT Test Project. Click Save.

14. In the upper right of the screen, use the pull-down menu beside root and select Manage Data Store Keys.
15. Click Add to create a new data store key.
16. Right-click and copy the string of characters under Data Store Key. Use this to authenticate to the S3 compatible storage.

Testing the S3 object store

1. Open the S3 browser.
2. Select Accounts→Add New Account…
3. In the Add New Account window, for Account Name, provide a name for the new account. We used `PT_ViPR`
4. For Access Key ID, enter root.
5. For Secret Access key, paste the character string you copied from Data Store Key into the field.
6. Click the link for Advanced.
7. In the Advanced account properties window, check the box for Use Amazon S3 compatible storage.
8. Enter the IP address and port of the ViPR Data Services node you just created. We used `192.168.1.115:9021`. Click Close.
9. Click Add new account.
10. Click Save.
11. In the S3 Browser window, click New bucket.
12. In the Create New Bucket window, for Bucket name provide a name for the bucket. We used `PT-Test-Bucket`
13. Click Create new bucket.
14. Select the bucket, and click Upload.
15. Browse to select an object to upload to the object storage. We selected a graphics file. Click Open.
16. The file uploads from your system to the object store.
17. Click on the file in the S3 Browser, and click Preview. A view of the object displays in the bottom panel.
18. The URL provided is the address S3 compatible applications will use to access the object.
19. Set up the S3 browser on another client. Connect to the S3 compatible storage.
20. The client retrieves the available buckets automatically. Select PT-Test-Bucket.
21. The client displays the available objects. Select the graphics file object. Click Download.
22. Browse for the folder you wish to use as the download target. We selected desktop. Click OK.
23. The client saves the object on your desktop. Close the S3 Browser.

APPENDIX L – SAMPLE ViPR REPORT

All >> Report Library >> EMC ViPR >> ViPR Summary >> PT_ViPR_Lab >> Virtual Array / Virtual Pool Capacity

PT_ViPR_Lab / Virtual Array / Virtual Pool Capacity

April 2014, Sunday 6 » Monday 7, 9:06 AM PDT | *Last 1 Day*

Virtual Pools Usable Capacity

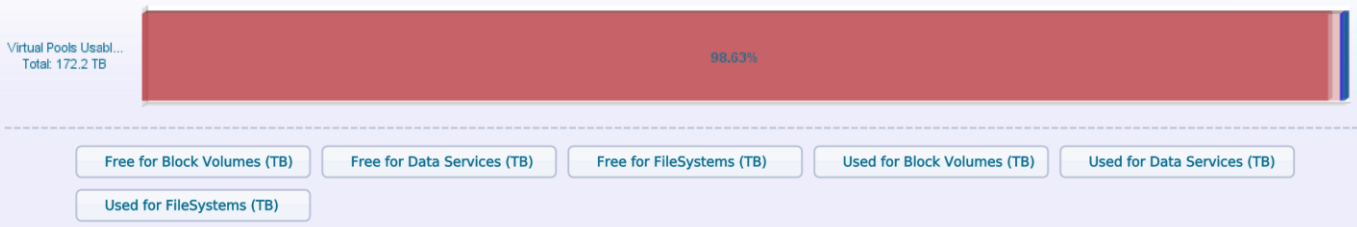

Virtual Array / Virtual Pool Utilization

9 elements found, displaying all elements.

Virtual Array Name	Virtual Pool Name	ViPR System	Virtual Pool Description	Usable Capacity	Current Utilization (%)	Current Provisioned (%)
Data Services	VNX File Pool	PT_ViPR_Lab	VNX FIle Pool	0.00		N/A
PT-VA1	High Performance Block	PT_ViPR_Lab	HPB	3.79 TB	0.23	1.00
PT-VA1	Low Performance Block	PT_ViPR_Lab	LPB	48.83 TB	0.00	0.00
PT-VA1	Medium Performance Block	PT_ViPR_Lab	MPB	27.70 TB	0.01	1.00
PT-VA1	PT FAST Pool	PT_ViPR_Lab	FAST	7.67 TB	0.00	0.00
PT-VA1	VMAX	PT_ViPR_Lab	VMAX Only	80.32 TB	0.01	1.00
PT-VA1	VNX FIle Pool	PT_ViPR_Lab	VNX File Pool	1.83 TB	5.51	6.00
PT-VA1	VNXBLOCK	PT_ViPR_Lab	vnx block	2.05 TB	24.95	25.00
PT-VA2	PT-vpool2	PT_ViPR_Lab	Isilon storage	6.00 GB	0.00	0.00

APPENDIX M – SOFTWARE VERSIONS

Software	Version
VMware vCenter Orchestrator	5.1.0 (build 2725)
VMware vCloud Automation Center	5.1.1 (build 55)
EMC ViPR	1.0.0.8.103
EMC Host Interface	1.0.0.0.174
SRM Suite	3.0
Isilon	b.7.0.2.3.r.vga
NetAppSim	8.2.0GA
EMC Solutions integration service	OVF10

Figure 23: Software version numbers.

APPENDIX N – HARDWARE DETAILS

System	Cisco UCS C220 M3
General	
Number of processor packages	2
Number of cores per processor	4
Number of hardware threads per core	1
CPU	
Vendor	Intel®
Name	Xeon®
Model number	E5-2609
Socket type	LGA 2011
Core frequency (GHz)	2.4
Bus frequency	6.4 GT/s
L1 cache	32 + 32 KB (per core)
L2 cache	256 KB (per core)
L3 cache	10 MB
Platform	
Vendor and model number	Cisco® UCSC C220 M3
BIOS name and version	C220M3.1.5.3b.0.082020130601
BIOS settings	Defaults
Memory module(s)	
Total RAM in system (GB)	128
Speed (MHz)	1,600
Size (GB)	8
Number of RAM module(s)	16
Chip organization	Double-sided
Operating system	
Name	VMware ESXi 5.1.0
Build number	799733
Language	English
RAID controller	
Vendor and model number	Emulex LightPulse LPe12002

System	Cisco UCS C220 M3
Ethernet adapters	
Vendor and model number	Intel I350-T2 Dual-port 1Gb NIC
Type	PCI-e

Figure 24: Configuration details for our test server.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.