12-9-2014

An Analysis of Storage Virtualization

Nicholas Costa

An Analysis of Storage Virtualization

By

Nicholas Costa

Thesis submitted in partial fulfillment of the requirements for the degree of
Masters of Science in Networking and Systems Administration

Rochester Institute of Technology

B. Thomas Golisano College of Computing and Information
Sciences

Department of Information Sciences & Technologies

December 9, 2014

Rochester Institute of Technology

B. Thomas Golisano College
of
Computing and Information Sciences

Master of Science in Networking and Systems Administration

Thesis Approval Form

Student Name: Nicholas Costa

Thesis Title: An Analysis of Storage Virtualization

Thesis Committee

Name	Signature	Date

Chair (Bill Stackpole wrsics@nssa.rit.edu)

Committee Member (Daryl Johnson dgjics@rit.edu)

Committee Member (Charlie Border cbbics@rit.edu)

Abstract

Investigating technologies and writing expansive documentation on their capabilities is like hitting a moving target. Technology is evolving, growing, and expanding what it can do each and every day. This makes it very difficult when trying to snap a line and investigate competing technologies. Storage virtualization is one of those moving targets. Large corporations develop software and hardware solutions that try to one up the competition by releasing firmware and patch updates to include their latest developments. Some of their latest innovations include differing RAID levels, virtualized storage, data compression, data deduplication, file deduplication, thin provisioning, new file system types, tiered storage, solid state disk, and software updates to coincide these technologies with their applicable hardware. Even data center environmental considerations like reusable energies, data center environmental characteristics, and geographic locations are being used by companies both small and large to reduce operating costs and limit environmental impacts. Companies are even moving to an entire cloud based setup to limit their environmental impact as it could be cost prohibited to maintain your own corporate infrastructure. The trifecta of integrating smart storage architectures to include storage virtualization technologies, reducing footprint to promote energy savings, and migrating to cloud based services will ensure a long-term sustainable storage subsystem.

Table of Contents

Table of Figures

Table of Tables

List of Terms and Abbreviations

AC – Air Conditioning

Amazon EC2 – Amazon Elastic Compute Cloud

BOM – Bill of Material

CAD – Computer Aided Design

CIFS – Common Internet File System

COTS – Commercial Off The Shelf

CPU – Central Processing Unit

CRAC – Computer Room Air Conditioner

CSP – Concentrated Solar Power

DAS – Direct Attached Storage

DMZ – Demilitarized Zone

EPA – Environmental Protection Agency

FCIP – Fiber Channel of Internet Protocol

FLOPS – Floating Point Operations per Second

GB – Gigabyte

GSR – Renewables Global Status Report

GW – Gigawatts

GWh – Gigawatt hours

HP – Hewlett Packard

HPC – High Performance Computing

HSM – Hierarchical Storage Management

I/O – Input/Output

IDE – Integrated Device Electronics

IE – Internet Explorer

IHS – Interstate Highway System

iSCSI – Internet Small Computer System Interface

IT – Information Technology

JBOD – Just a Bunch of Disks

JPEG – Joint Photographic Experts Group

KVM – Keyboard, Video, and Mouse

LAN – Local Area Network
LDAP – Lightweight Directory Access Protocol
LEED – Leadership in Energy and Environmental Design
LTO – Linear Tape-Open
LUN – Logical Unit Number
MB – Megabyte
MIPS – Millions of Instructions per Second
MIS – Management Information Systems
MP3 – Moving Picture Experts Group Layer-3 Audio
MPEG – Moving Picture Experts Group
MR Heads – Magneto Resistance Heads
MW – Megawatts
NAS – Network-Attached Storage
NFS – Network File System
NHS – National Highway System
PPA – Power Purchase Agreements
PV – Photovoltaic
PVU – Processor Value Unit
RAID – Redundant Array of Inexpensive Disks
RAM – Random Access Memory
REC – Renewable Energy Certificates
REN21 – Renewable Energy Policy Network for the 21st Century
RHEV – RedHat Enterprise Virtualization
RMA – Return Merchandise Authorization
SaaS – Software as a Service
SAN – Storage Area Network
SAS – Serial Attached SCSI
SD – Secure Digital
SDHC – Secure Digital High Capacity
SLA – Service Level Agreement
SMB – Server Message Block

SNIA – Storage Networking Industry Association

SNMP – Simple Network Management Protocol

SSD – Solid State Disks

SSO – Single Sign-On

TB – Terabyte

TP – Thin Provisioning

TSM – Tivoli Storage Manager

UNIX – Uniplexed Information and Computing System

USB – Universal Serial Bus

VMDK – Virtual Machine Disk

WAN – Wide Area Network

Acknowledgements

Thank you to all my teachers in life, to include RIT faculty and staff and mentors in my professional career.

Thank you to my family for their never ending love and support.

And a special thank you to my wife Erin Humphries for introducing me to a whole new level of drive and dedication in my career, in life, and in our lifelong partnership.

1 Introduction

Storage virtualization is the consolidation of homogeneous or heterogeneous physical disks, specifically to enable a single or set of virtualized partitions to be presented to compute nodes. Storage administrators utilize this technology in the hardware or software layers to improve disk utilization, throughput, backups, archival, recovery, and redundancy while disguising the actual complexity of the Storage Area Network (SAN). Some hybrid appliances also exist to combine the technologies of the hardware, such as perpendicular storage on the hard disks with the software layer application. An example of a software based technology would be file deduplication. Hybrid appliances use the best technologies of both hardware and software to be combined in a single solution. These solutions come in form factors and price points for all sized businesses and applications.

As multimedia becomes more and more prevalent with smartphones, tablets, and laptops at all-time record numbers, and with social media exploding with pictures and video on these devices, the backing storage system needs to be able to handle this increased storage capacity. Apps for phones and tablets are used to access multimedia content of high definition video and CD quality audio by downloading or streaming from a source. Cloud based storage, such as Google Drive, Apple iCloud, and Amazon Cloud, are all service based cloud storage applications to maintain your online multimedia purchases, documents, and apps. Keeping your content in a centralized location and letting the provider store, backups, and retain this content is cost efficient for the end user. This also consolidates data centers and allows your content to be efficiently stored, spread among many spinning disks for faster access and disk redundancy.

Maintaining these centralized storage environment in a cloud based storage subsystem for the end user presents an opportunity for the company to consolidate its storage into an energy efficient storage architecture using green technologies and environmental aspects. Because Internet throughput bandwidth is also increasing, this allows companies and cloud providers to build data centers in convenient locations where the environmental benefits decrease operating costs. This methodology is far more efficient than each user to maintain their own content on their device with their own individual backup solution, such as a USB external hard drive.

Section 1.1 will address the problem statement as to the need for smart virtualized storage subsystems and the environmental impact. The importance of this virtualized storage research, combined with the environmental concerns facing us today will be discussed in Section 1.2. Section 1.3 then provides a brief insight into businesses from small to large and how virtualized storage can benefit their company. A review of the current research is examined in Section 1.4. Summarizing thoughts and laying out the rest of the document is in the overview, Section 1.5.

1.1 Problem Statement

The problem that we face today is that the amount of storage required for both personal and business is growing exponentially. The days of text are long gone. Today's needs and desire is for multimedia which far outweigh the storage capacity compared to text in the past. Websites, apps, audio, video, and pictures have become commonplace with the expansion of personal computing into personal devices such as laptops, tablets, and smartphones. To accommodate these new storage requirements, the storage subsystems have to become faster, larger, and more accessible. Some examples of cloud based storage include Google Drive, Apple iCloud, and Amazon Cloud. These systems are used for storing email, calendar appointments, contacts, music, movies, and smartphone device backups. The cloud capability makes all your content available from your smartphone, tablet, and browser from any device connected to the Internet. This is ushering the World of digital bits and bytes into data proliferation (McArthur & Villars, 2005).

Data proliferation is a term concerned with the large number of files and amount of data stored by entities such as governments and businesses (wiseGEEK). The massive amount of data being generated exceeds the space and hardware advancements of today. Figure 1 below depicts statistics written by (Okwuke) showing the exponential growth of data with an estimated worldwide storage into the Exabyte range.

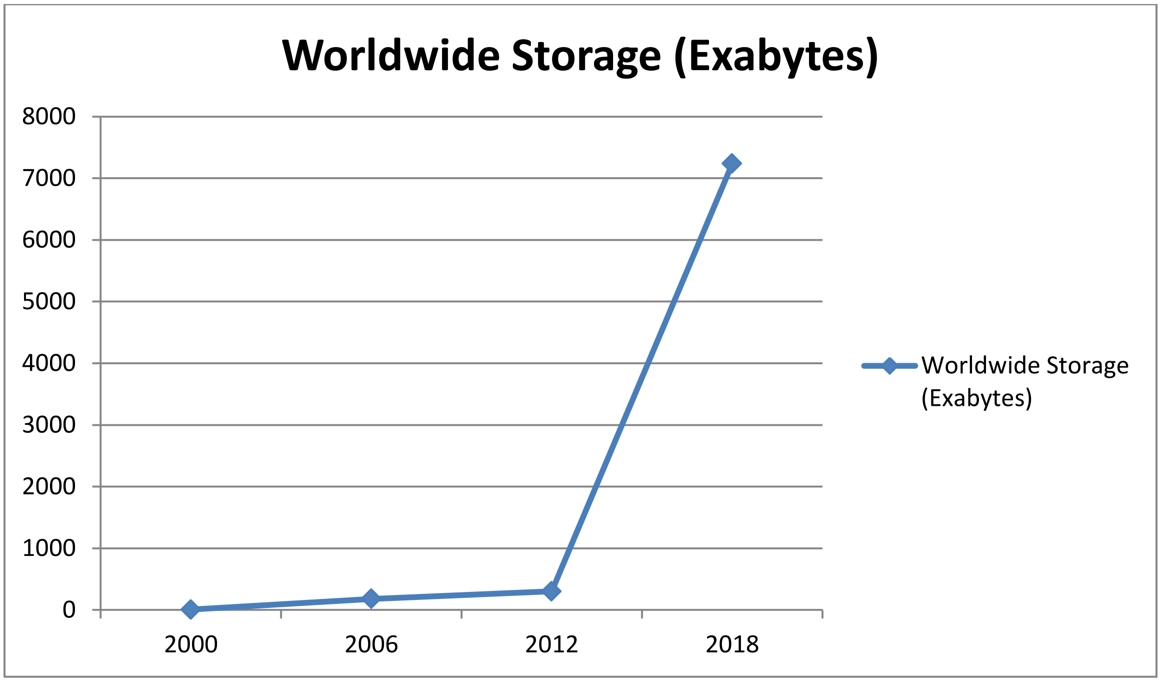

Figure 1 - Worldwide Storage in Exabytes

The main concern with data proliferation is that individual consumers or even households can purchase additional hard drive space to meet the requirements on their own. However, when data is presented into the Cloud or to businesses such as Facebook or Google, data proliferation may manifest much quicker. To combat this large influx in data, businesses have to resort to storage virtualization technologies and techniques to decrease the amount of physical disk, and potentially decrease the number of data centers to maintain a cost effective storage solution for the end users. Storage Virtualization provides some new technologies, but its main purpose is combining existing technologies for an integrated storage subsystem solution. These technologies are researched in detail in Section 3 – Storage Technologies Research.

Surrounding these integrated storage subsystems are data centers that power and cool the physical hardware. Part of the problem is the expense as a result of how we power these facilities. Energy costs are increasing, just as the need for additional storage, exacerbating the issue of more storage, more data centers, more power, and more cooling. Companies are switching to greener solutions to power and cool these facilities. And not just changing the technologies, but the geographical location as well. Section 4 – Green Energy Technologies Research will discuss the green technology factors associated with storage virtualization and the data centers that maintain the equipment.

1.2 Importance

The importance of this research is to highlight the environmental impacts that data centers are having on today's ailing world with overpopulation and consumption of natural resources. With the cost of energy today, it makes it very hard to have large data centers anywhere in the world. According to a 2008 whitepaper from Green Storage Initiative, about 2 percent of the total energy consumption in the United States is by data centers, equating to approximately $4 Billion annually (Clark & Yoder, 2008). This is equivalent to six 1,000 Megawatt power plants just to sustain the current operations.

In addition to the waste of old computer equipment and power consumption, throwing more hardware to scale up a data center is no longer a viable solution due to energy costs for both powering and cooling the data centers. In this paper, we will discuss in length technologies for both the storage solutions within the data centers as well as the electrical systems that power and cool these massive facilities.

The purpose of this paper is to make businesses of all sizes aware of their options for storage virtualization systems. Also, how these systems can help to benefit their energy costs. These technologies discussed are meant to revolutionize the old data centers as well as provide less invasive solutions to simply modify or compliment the current systems that exist today.

You cannot put a price on data. "43% of companies were immediately put out of business by a "major loss" of computer records, and another 51% permanently closed their doors within two years – leaving a mere 6% "survival" rate" (Gartner, 2010)". These are scary numbers. Providing a solution for long term storage resiliency mitigates the risk of data loss and potentially going out of business provided catastrophic failure of your company's storage subsystem.

1.3 Business Impacts of All Sizes

As part of the deliverables for this paper, I have put together sample architectures for three differing business constructs; small, medium, and large scale business keeping to a budget of $100,000, $1,000,000, and $10,000,000 respectively. These sample architectures of discussed in length in Section 8 – Appendices. The point of providing a deliverable based on different scale of business is to show that there are many ways to go about providing storage for your company of that particular size. There is no single correct solution but there are storage virtualization technologies that will help to mitigate data center footprint and storage subsystem costs. As you will see, the options are very different based on the scale of the business and storage needs.

1.4 Review of Current Research

Research regarding storage virtualization technologies and their impact on data centers for increasing storage functionality without adding additional spinning disks is prevalent in both academia and big business. They are all interested in reducing the data center footprint and power consumption while meeting their service level agreement (SLA) with their customers. In addition to being more efficient with the use of their storage architecture, energy efficiencies and considerations are being addressed related to data centers as a whole due to the increased costs in energy worldwide and the trend toward green computing.

(Verma, Koller, Useche, & Rangaswami, 2010) conducted extensive research into an energy proportional storage system through power-aware dynamic storage consolidation. Essentially they created a virtualization layer that optimizes energy consumption for dynamic I/O workload intensity. Their software, SRCMap, samples a subset of blocks from each data volume and replicates these onto other physical volumes. After the consolidation is complete, the physical disks spin down and redirect their workload to replicas on the active volumes. This software helps to minimize the power consumption of enterprise storage subsystems by maintaining only the minimum number of spinning disks required. Interestingly, because the data blocks were migrated to only the disks still spinning, access time was not impacts while power savings ranged from 35 to 60 percent, dependent upon the testing scenario.

(Clark & Yoder, 2008) released a whitepaper for SNIA (Storage Networking Industry Association) for their Green Storage Initiative. This best practices article addresses fundamental considerations regarding data center IT (Information Technology) operations and the energy required to support future data centers. The paper goes on to discuss 20 different best practices into how to decrease your footprint of your data center and thus reducing energy consumption. Some of these are technologies and some are examples of policy implementation with customers and end users. All technologies mentioned in this whitepaper are discussed in Section 3 where we will go into detail about the technologies related to storage virtualization and how they can benefit data centers of all sizes.

(Aikema, 2013) wrote his dissertation on assessing and reducing data center and financial costs. His paper goes very in depth into the complexity of computing and its relationship to energy and carbon emissions. The author is most concerned with the analysis of how data centers are impacting the environment and optimizing the energy consumption. This differs from my paper in that I'm focused on technologies in the data centers, most notably related to storage subsystems and how that can decrease energy consumption all together.

A very short but interesting article was discussing whether or not virtualization really does save energy. This was very basic and only harnessed virtual machines and not storage virtualization. The result of his analysis is potentially a six figure savings over 3 years but virtualizing just a handful of servers (Carl, 2011).

1.5 Overview

The remainder of this document is organized as follows. Section 2 presents the concept basics to understanding the research and background foundational information to accurately address concerns facing us today. Section 3 expands to an in-depth look at the technologies helping the storage virtualization. Sections 4, similar to Section 3, goes through an in depth look at green technologies and their linkages to data centers. Section 5 reviews cloud computing and its positive environmental impact. Section 6 presents the conclusions of the research. Section 7, 8, and 9 address future work, appendices, and references respectfully.

2 Background

2.1 The History of Storage

The history of storage dates back to the early 1950's where the only storage technology at the time was magnetic tape. The first device to incorporate this technology in the commercial world was IBM with its 305 RAMAC. This device could hold up to 3.75 Megabytes of data and consisted of fifty 28 inch disk platters (Dell, 2013).

It wasn't until 1976 when the first floppy disk was released. At this time, the storage could hold just 110 kilobytes with a double sided denser model coming out in 1978 that could hold 360 kilobytes. These disks were used to boot operating systems and applications separately. They were commonly referred to as "five and a quarter floppies" for their 5 and ¼ inch diameter.

In 1979, the next commercial device became available to purchase or lease from IBM. The storage unit was called the IBM 3370. Its storage capacity was just over 500 megabytes and could be purchased for $35,100 at the time. In today's dollars, that is nearly $200,000 per gigabyte.

The Seagate ST-506 took the concept the of 5.25" floppy disk and made it into a hard drive for personal computers. This was released in 1980 with a storage capacity of 5 megabytes.

The smaller floppy, commonly referred to as the diskette, was released in 1981. It took all the way until 1988 to outsell the 5.25 floppy disks, likely due to the expense in the personal computer market to change hardware platforms in such a short period of time. Its compact size was ideal for personal computing and businesses alike and maintaining the same storage capacity as the double layer 5.25" floppy of 360 kilobytes. The later revolution of the diskette came as a dual sided storage medium as well, increasing its storage capacity on its "high density" flavors to 1.44 megabytes 1987. This is where this storage technology began its decline.

Western Digital, a primary competitor of Seagate, came out with the 3.5” form factor hard drive named the Caviar. This disk had a rotation speed of 3595, which did not last long. The company just a year later introduced the 5400 RPM spinning disk speed, which was a standard for hard disks for many years, and is still used today in some portable computer applications due to its low power consumption. This hard disk was 40 megabytes of data storage at its release in 1990. This was the first device to utilize the IDE (Integrated Device Electronics) interface.

In 1991, there was a breakthrough by IBM with their 0663 Corsair hard drive. They started using magneto resistance heads (MR heads). The storage community had been using magnetic tapes since 1984. This new technology allowed for much greater storage density, thus greater storage capacity, jumping up this single drive to 1 gigabyte worth of storage.

At this point in time, technology for storage subsystems had really taken off. Just a year later, Seagate announced its Barracuda hard drive making it the fastest at the time with a transmit speed of 113 megabits per second. This was many firsts for hard drive storage technology; first drive to have a capacity of great than 2 gigabytes, first drive operating at 7200 RPM, and first multi-head hard drive that could use more than 1 head simultaneously by splitting the stream. By 2006, Seagate would be the first hard drive manufacturer to produce a hard drive costing less than $1 per gigabyte. Seagate would make the Barracuda line of hard drives all the way until 2013.

MMC (Multi Media Card) came into the marketplace by SanDisk in 1997 at the size of 4 megabytes. These cards established many form factors over the years, the primary one in 1999 called the SD (Secure Digital) cards. Today, SD cards can hold up to 128 gigabytes worth of data is their SDHC (SD High Capacity) formats and has been used in portable electronics such as cell phone and digital cameras. The SD card is by far the most prolific of the MMC types.

Western Digital released the Raptor in 2003 as the first 10,000 RPM hard drive. Storage sizes varied with the most popular being 36 gigabytes. This was the first hard drive with a SATA connection available to industry and personal computing alike.

SSD's (Solid State Disks) date back to the 1970's, but entered the market as a viable option for personal computer hard drives in the late 2000's. At that time, their downside was a very expensive price per gigabyte, about 10 times higher than the spinning disk alternative at the time. Today, SSD are prevalent everywhere. The cost continues to decrease, provides much higher transfer rates than spinning disk, has a significantly longer life span, does not degrade performance due to disk fragmentation, and uses less power for those portable devices using battery. This is most certainly the storage of the future for low power, low heat, and quick access storage medium. The power consumption of a SSD is roughly a third of what a spinning disk uses (Schmid & Roos, 2008).

Spinning disk will likely take many years to go away due to is low cost per terabyte and amount of storage in the same form factor as SSD (Hess, 2011). However, tiered storage technology will let a storage subsystem take advantage of different storage mediums and prioritize for speed based on access frequency.

2.2 The History of Data Centers

The history of data centers is quite simple considering their existence does not predate World War II. The time period for data centers really came alive in the 1980s with companies like IBM and Sun Microsystems producing mainframe computers. Mainframes were built to be reliable systems for transaction processing of commercial exchange of good, services, or money. Transactional processing is their primary goal and is measured by their input/output of performing thousands of concurrent transactions.

Supercomputers differ from that of mainframes because they are measured in FLOPS (Floating Point Operations Per Second), whereas mainframes are measured in MIPS

(Millions of Instructions Per Second). Supercomputers are constrained by calculation speed. Their primary goal is for very fast and complex calculations while pushing the limits of computational speed (ASPG, 2013).

During the mid-1990s, the supercomputers/computers were far outnumbering the purchasing orders of mainframe computers. Mainframes were being replaced by personal computers connected by a LAN (Local Area Network). This was due to the much easier programming projects that could be accomplished on personal computers connected to servers rather than mainframe based programming. In 1995, Paul Watz, Director of IT for Motorola was told to cut costs of their MIS (Management Information Systems) from 3.7% down to 1% (Woods, 2013). He was able to accomplish this by migrating away from 2 IBM mainframes to three distinct LANs consisting of 1,000 workstations each, while connecting to 30 centralized servers. This trend and fundamental concept of LAN based client to server architecture exploded into what became known as the Dot-Com Boom.

Despite the negative connotation with “Dot-Com Boom” regarding its implications to the rise and fall of Internet companies and startups in the stock market, there really was a paradigm shift in the way data centers and companies operated their IT infrastructure. The Internet was all of a sudden widely available to the general public and companies were trying to keep up with the demand of traffic and web site access requests.

Regardless of the types of systems being used, the data center was and is a fundamental component to a business to house, power, and cool its compute nodes and network infrastructure. Back during the inception of data centers, there were a few key components as to where they were geographically located:

1) Labor Availability and Quality – One key factor for the original data centers included the need for skilled engineers and administrators that were computer savvy to work on the complex set of servers and workstations, as well as network and storage hardware. The original concentration and most famous pool of

skilled engineers and scientists came from the major universities in the San Jose, California area, soon to be nicknamed Silicon Valley.

Dating back to the 1950’s, Silicon Valley was home to major technological corporations including Bell Telephone Laboratories and Fairchild Semiconductor (Gigerich, 2012). These companies were in proximity to Stanford University uniting academia and the commercial worlds boasting development of bleeding edge technology of the time. This area in California also obtained generous funding from the Defense Department and venture capital firms after the successful development of the microprocessor. At that time, this area was the quintessential geographic location for scientists, engineers, and administrators of technological developments.

2) Telecommunications Infrastructure Availability – The telecommunications infrastructure was a key driver to the location of data centers because connectivity was very scarce. Port based cities were identified as key locations because of underwater fiber optic connections, also known as submarine communication cables. This includes locations such as San Francisco (Silicon Valley), New York, and London. In 2001, the largest point to point connection in regards to bandwidth was from New York to London at 25 Gbps. The total US & Canada to Europe bandwidth at the time was only 56 Gbps. As of 2012, the throughput between US & Canada to Europe now exceeds 7 Tb/sec; over a 12,000% increase in 11 years (Lennett & Meinrath, 2009). Even within the United States, ground based high bandwidth connectivity was only availably in large metropolitan areas requiring the geographic location of data centers limited to these strategic sites.

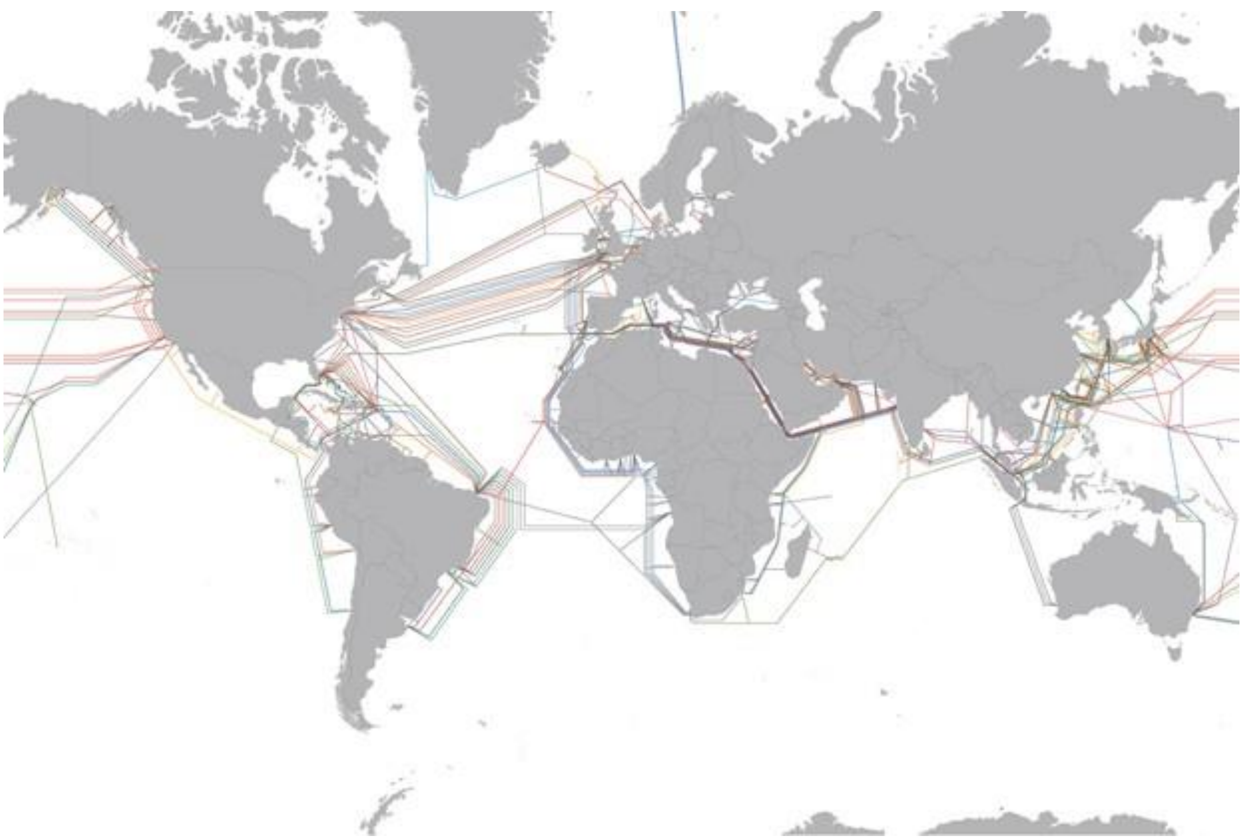

Figure 2 - Global Underwater Internet Connectivity (Wired)

3) Availability and Cost of Electricity – One obvious major proponent of a data center is powering and cooling the data center. This can result in massive electricity costs, however originally, finding locations that can support the amount of electricity needed was challenging. The maturity of the grid was easily overlooked. The concern around energy in the inception timeframe of the data center was most likely how to fuel the car and not how to power your computer. Powering the equipment and refrigeration are the two largest users of electricity in a data center. Being energy efficient was not an original concern, simply finding the power to run them were most definitely a challenge. Again, just as telecommunication infrastructure in high quality was scarce, also was the electricity. Major metropolitan hubs were again identified as key locations for data center sites.

Today, roughly 20 years later, the concerns are somewhat the same but more have made the list based on the environment that the world has evolved into. Climate change, population density, strategic telecommunications infrastructure, natural disasters, and taxes play a large role as to geographically where data centers are located today.

1) Labor Availability, Quality, and Cost – Whereas engineering personnel drove high costs due to geographical location years ago, availability of high speed computer systems and infrastructure access makes this now far less of a concern. The best and brightest around the world can be resourced and work remotely on servers and workstations around the globe. This can benefit companies by keeping a lower standard for salaries driving down cost because remote employees do not have to suffer the high cost of living in major metropolitan areas as they did years prior. Point number 2 below discusses improved telecommunications infrastructure availability, which helps to address number ones concern of labor availability, quality, and cost.

2) Telecommunications Infrastructure Availability – Today, this is now far less of a concern because high speed networks are prevalent everywhere. Data centers are no long dependent geographically based on previously limited high bandwidth connectivity locations. In bullet number 2 related to data centers of old, it was mentioned that San Francisco area, New York City, and London were ideal locations back in the day for data centers because of their telecommunication infrastructure availability. Today, major high bandwidth and low latency fiber optic connections exists both undersea to connect continents as well as underground terrestrial connections linking the intercontinental port based connections to inland cities and connectivity hubs. These connections traditionally follow the National Highway System (NHS) that comprises of approximately 163,000 miles of roadway, including over 46,000 miles of the Interstate Highway System (IHS). The NHS reaches nearly every part of the country. 90 percent of the United States population lives within 5 miles of an NHS roadway making the network of highways an ideal cost-effective and sustainable means to burry high speed fiber optics for remote connectivity. This also allows for easy access in case a repair of the system is needed. What this means for big businesses is that they can strategically place their data centers in accordance with numbers 4 and 5 below to limit the potential for natural disasters and decrease the cost of real estate and property taxes.

3) Availability and Cost of Electricity – Availability of electricity to power and cool data centers is no longer a concern in the Western world. In the United States, energy costs vary state to state only slightly due to the Federal Government subsidizing both fossil fuels and renewable energy. In Europe however, energy costs are much higher due to the EU’s strict and ambitious climate change policies. Many European companies that require high usage of energy are moving their facilities elsewhere, particularly the United States. There are however, a few places in Europe that are embracing their unique climate characteristics and making it into profitable energy conscientious businesses. To be more specific, Iceland is standing up data centers due to their 100 percent

renewable energy sources thanks to geothermal and hydroelectric power. This power is not only dirt but it also decreases the carbon footprint of a company utilizing these Icelandic data centers (Gilmer, 2011). Another component to Iceland that helps its business model for data centers is the cool temperate climate due to the Gulf Steam feeding the North Atlantic Current and into the Irminger and Norway currents. The northern latitude of the island is the reason for cool temperatures year round and circulated outside air does not require refrigeration also helping to keep costs down. One company embracing this new data center location is BMW. BMW migrated ten High Performance Computing (HPC) clusters into the Icelandic data center reducing its annual carbon emissions by 4,000 tons and the operating costs by up to 82 percent (Information Age, 2012). These systems are used for large computational required applications such as crash simulations, aerodynamic calculations, and Computer Aided Design (CAD). Because of the telecommunications infrastructure availability in the modern world, as we discussed previously, computers do not need to be collocated with the administrators and engineers of the system. BMW is a perfect example of a company offloading its CPU cycles to a country where the power is cheap, clean, and renewable. According to the U.S. Energy Information Administration statistics back in 2011, the electricity in Iceland is as low as 4 cents per kilowatt hour compared to 11 cents in the United States.

4) Natural Disasters – One topic that has bubbled to the top is the decision for data center geographic location based on frequency and type of natural disasters. Earthquakes, hurricanes, tornadoes, blizzards, floods, volcanic activity, and lightning storms are just a few considerations when discussing situational awareness for locating a data center. We discussed previously the positive implications for a data center in Iceland due to its renewable clean energy. On the contrary, if the systems you are deploying are deemed mission critical, think again due to seismic and volcanic activity which is very prone in this part of the world. We are able to forecast the weather but volcanic activity and earthquakes

can strike without notice. This makes California a difficult location for data centers due to its seismic activity (Latimer, 2011). There exists special seismic bracing for raised floors and rack units but the lack of Internet connectivity from your ISP and power outages due to seismic activity will bring down the data center before physical damage is the main culprit of downtime. It is best to avoid these areas with a high probability of natural disasters.

5) Real Estate Costs and Property Tax Rates – Another consideration for where to locate a data center are the real estate costs and the property taxes associated with that piece of property. This was hardly a consideration back in the day because the data centers needed to be placed where the telecommunications, engineers, and administrators of the systems were located. Now with remote access capability, the geographical location of the data center is not of as much concern and can be placed where the real estate prices and tax rates are conducive to a low cost piece of property. Colorado was recently acclaimed as the new and up-and-coming data center hot spot in the United States due to its moderate real estate costs and low property taxes when compared to coastal locations such as Silicon Valley and New York City. Google has a presence throughout Colorado as well as Seagate, Lockheed Martin, L3 Communications, Digital Globe, Ball Aerospace, and a new Microsoft facility.

6) Proximity to Major Airports – A close proximity to a major airport allows for connections to major markets and customers. Easy accessibility to fly coast to coast and international opens up possibilities with remote customers and international markets. This also provides the ability to send and receive equipment with a short turnaround time and lower shipping costs. This is for both sending Return Merchandise Authorization (RMA) equipment and receiving new items for the data center and business as a whole.

Despite all the pros and cons reviewed in this section related to the history of data centers, the transition to Cloud based services must also be addressed. Companies are farming out storage to cloud based services like Google Apps for Business @ $50 per user per year and other cloud services in order to avoid maintaining a data center all together. Today's data centers are shifting from an infrastructure, hardware, and software ownership model, toward a subscription and capacity on demand model, hence the cloud computing based services. This means storing data offsite into a highly redundant secure facility stripped across many disks in many locations. This could also mean paying for CPU cycles from a remote data center provider. This would mean not having to run and maintain your own infrastructure. Rather you would pay a data center company to spin up a virtual machine (VM) and use it for a short duration then shut it down. Depending on the use case, this could be far more cost effective than having your own hardware and supporting infrastructure running 24 hours a day, 7 days a week.

3 Storage Technologies Research

There are many different technologies that make up storage virtualization. Some are specific to a single hard or solid state disk and some are at the level of the entire storage subsystem. Within the storage subsystem, there can be many different parts that accumulate to a storage solution such as a RAID (Redundant Array of Inexpensive Disks) system, fiber channel or iSCSI (Internet Small Computer System Interface) connection medium, fiber channel switches, virtualization cache appliances, software application layer, NAS (Network-Attached Storage) filer heads, and magnetic tape disk. Not all these pieces are required to make up a storage virtualization solution but some combination of this equipment with the technologies listed in the following subsections of Section 3 will enable the implementation of storage virtualization. The intent of Section 3 is to dig deep into the intricacies of each technology followed by putting them to use in Section 8 as applicable for small, medium, and large sized businesses.

3.1 Thin Provisioning & Over-Allocation/Over-Subscription

Thin Provisioning (TP) is a method of optimizing the efficiency with which the available space is utilized in a SAN. TP operates by allocating flexible disk storage space among multiple users based on the minimum space required by each user at any given time (Rouse, whatis.com). Whether using a Server Message Block/Common Internet File System (SMB/CIFS) share for Windows or a Network File System (NFS) for Uniplexed Information and Computing System (UNIX), thin provisioning works all the same (Cason, 2014). It is typically implemented by either the Network-Attached Storage (NAS) filer head or the NFS/SMB/CIFS server itself.

If TP was implemented by a NAS filer head, then the hardware appliance would typically connect to the backing storage through fiber channel switches or directly to the storage subsystem itself. When creating a volume, you would have the option to implement TP. TP is the opposite of Fat Provisioning (FP), also known as thick provisioning, which instantly allocates 100 percent of the storage for the particular volume selected when being created even if none of the storage is being used. The figure below shows the capacity saved in the thin provisioning stack on the right as compared to thick or fat provisioning.

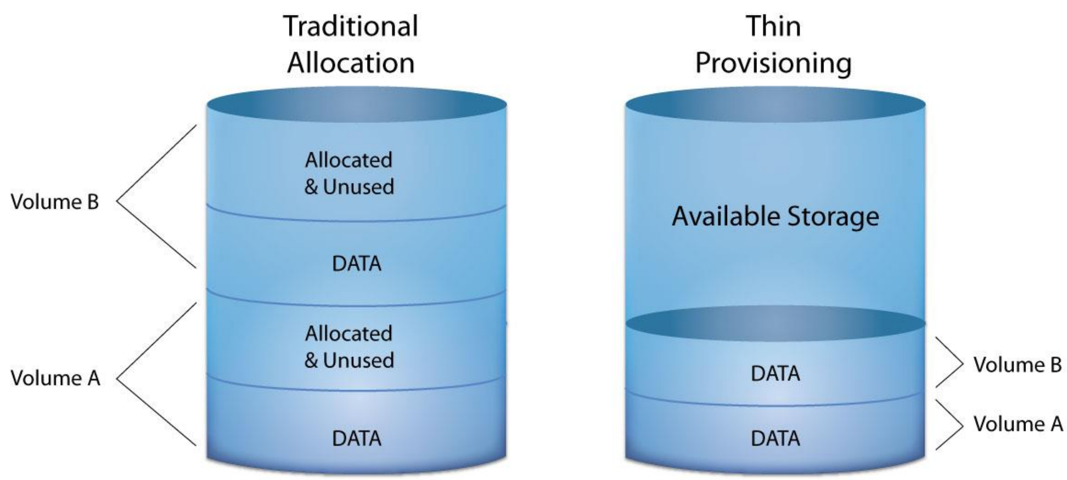

Figure 3 - Storage Savings in Thick versus Thin Provisioning (Thin-Provisioning)

Thin provisioning technology is commonly used in conjunction with over-allocation/over-subscription of the storage subsystem (Warden, 2006). A good use of thin provisioning would be for home directories. In this example, if a volume was created for 1 Terabyte (TB) of storage, then thin provisioning would show the NFS/SMB/CIFS clients 1 TB of free space with only a few Megabytes (MB) of storage being used on the storage array itself, primarily for the file system overhead. This differs from the old methodology of FP, which would instantly allocate 1 TB of the backend storage resulting in an inefficient use of the backing storage.

One advantage of TP is the flexibility of being able to allocate more storage than you have available on your backing storage subsystem. The disadvantage to this, going back to our home directory example, if you allocate 1 TB for home directories on a 500 Gigabyte (GB) storage array, this would be considered 100 percent over-committed or over-subscribed. Users can now exceed the capacity of your storage subsystem as soon as they consume over 500 GB of storage in the home directory volume. This could result in a read only file system or the entire system to go offline. Worst case, you end up with data corruptions because multiple concurrent writes were abruptly stopped when the backing storage array ran out of physical disk space.

To avoid this situation, there are many checkpoints for warning messages via email or SNMP (Simple Network Management Protocol) traps that can be identified for when you are approaching your over-allocation thin provisioning threshold. It is recommended that as soon as 75 percent of the backing storage subsystem is utilized, more physical disk is added to this storage array in order to not run the risk over hitting 100 percent utilization of the available physical disk.

Another advantage of TP is that an organization can purchase less storage capacity up front, thus deferring storage capacity upgrades in line with actual business usage. This would help to maintain lower operating costs for both electricity and raised floor space.

One disadvantage of thin provisioning is that during a large high speed data write to the share point, thus the backing disk, the file system will have to expand its sizing boundaries on the disk concurrently with the data transfer. This becomes a limiting factor to the transmission speed based on the speed at which the file system can expand. This should result in a negligible amount of disk read and write latency but should be evaluated in the specific implementation of thin provisioning and over-allocation on the storage array to meet the required business needs.

Another situation to avoid is what is known as “thin on thin”. This is when a volume is thinly provisioned in which its purpose is to NFS mount virtual machines. If those virtual machines are then thinly provisioned inside of a thinly provisioned shared file system, you run into the thin on thin paradigm. This would result in complicated math that can quickly be overlooked as a potential disaster that cannot be mitigated fast enough because of the high percentage of over-allocation.

StorageTek was the first to implement the thin provisioning over-allocation concept with their Iceberg product in 1992. This was revolutionary at the time but flawed by under reporting the amount of used space and did not cope well with snapshots. Iceberg was proprietary for IBM mainframe systems of the time. EMC became the dominant player in the mainframe market shortly after the Iceberg product was released, thus ending the reign of the IBM/StorageTek OEM system.

In the open computing world, DataCore Software Corporation released their Network Managed Volume (NMV) feature in their SANsymphony product in 2002. DataCore provided the capability to centralized storage monitoring, simplified storage allocations, and maximizes capacity utilization. SANsymphony is a middleware product that resides as software on a Windows based system between the storage arrays and edge servers utilizing the storage subsystems, either as iSCSI, raw LUNs (Logical Unit Number), or NFS shares. Additional details on how this product can be used and configured will be discussed in Section 8.2.3 for a medium sized business implementation.

Other large companies such as EMC Corporation and Hitachi Data Systems sell hardware storage platforms in which software that runs on Windows and Linux operating systems connect to the storage units for configuration. This software provides thin provisioning and over-allocation/over-subscription capabilities. The available configurations and capabilities are more or less the same from vendor to vendor.

VMware has thin provisioning capabilities as well for its virtual machines. The virtual machines local storage resides in large flat files called Virtual Machine Disks (VMDKs) (VMware, 2014). These files can be thinly provisioned. For example, a virtual machine can be allocated 100 GB of data storage, and after installing the operating system, if only 7 GB is used, then 7 GB is the size of the VMDK file. This file can also reside on shared storage as recommended by best practices in order to have disk redundancy underneath the virtual machines, contain additional physical disks for performance enhancements, and provide easy virtual machine migration. This paradigm you could run into is thin on thin as we discussed previously, if the VMDK is thinly provisioned as well as the storage subsystem that it resides on.

3.2 File Deduplication

File deduplication reduces the amount of storage needed for a given set of data where there are multiple copies of existing files. Pointers are implemented to allow a single copy of the file to exist on the backing storage subsystem. This technology is also known as intelligent compression and is important to note that this is implemented at the file system level (Stewart, 2013). Third party software vendors have created file deduplication software that runs on both Windows and UNIX to identify multiple instances of files and alerts the user to determine if they want to take action. This works well for pictures and documents but not at all scalable or automated for business or enterprises.

Like thin provisioning, this technology can be implemented at either the software layer or from a hardware appliance. The software layer would assume a storage virtualization software suite is being used, which is where the file deduplication will be instantiated. The hardware layer implies that the file deduplication is being done on the controllers of the storage array or on the NAS filer heads doing the data I/O for the storage subsystem.

Whether being implemented at the software layer or the hardware layer, there are two differing types of file deduplication: post-processed and in-line. Post-processed file deduplication means that the data is written and committed to disk. Later on in time, analysis of the data will be done to determine identical files, which will be deleted from the storage array and replaced with pointers thus saving critical disk space. The benefit to post-processed file deduplication is the insurance that there is no wait time for hash calculations and lookups to occur before storing the data thereby guaranteeing maximum write performance. The optimal usage of this technology would be implemented as a policy based rule set outside of business hours and completing immediately before nightly backups of the data stores occur. This benefits both the storage array as well as the backup system minimizing the amount of data to be stored and written to the backup system.

In-line file deduplication processes hash calculations in real time. If the software or hardware device detects a file that already exists in the storage subsystem, the file is never committed to disk, rather a pointer references the existing file. The benefit to in-line over post-processing is that if the data being transferred is rather large, the data is never written to disk, saving the time and storage space from ever being written to. However, the downside is the length of time to calculate the hash values and lookups into the volume. The usage of in-line versus post-processing file deduplication is heavily debated but is dependent upon the implementation of the storage volume.

A good candidate volume for file deduplication would be a common shared file system for Windows based systems. Typically in a business, there are Microsoft Office document types that have an existence in many personal folders stemming from the same parent shared volume. File deduplication would cut down the used storage on the volume or shared file system by removing the duplicates and instantiating pointers as we discussed previously. The following figure shows a six to one storage reduction by removing the other five 100 GB instances of the file.

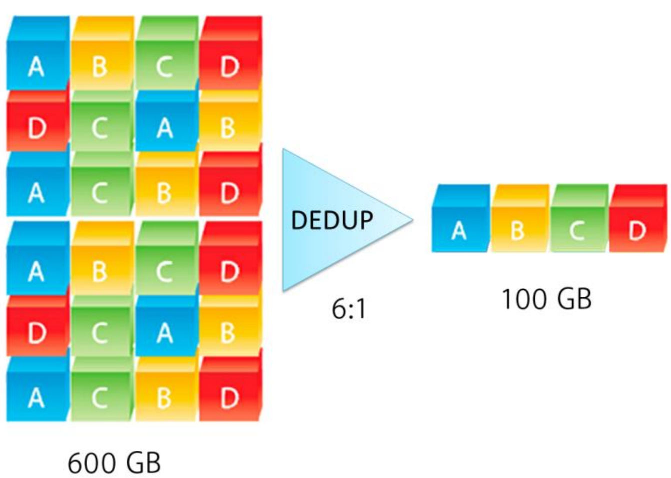

Figure 4 - File Deduplication Example File Set (SoftNAS Cloud Features)

A similar implementation of file deduplication is available on an email server application such as Microsoft Exchange or IBM Lotus Notes. Back in the inception of the email client, an email with a 1 MB attachment send to 100 people would take up 100 MB of storage on the email servers file system. Today, that same message would only take up 1 MB of storage and reference that file from the other 99 sent messages.

3.3 Block Deduplication

Block deduplication, also known as data deduplication, reduces the amount of storage needed for a given set of data where only one unique instance of data is retained on storage media (Rouse, 2010). Pointers are implemented to allow a single copy of the data block to exist on the backing storage subsystem. Block deduplication is very similar to file deduplication but can only be implemented at the hardware layer. Block deduplication is different than file deduplication because the hash comparison mechanism happens at the block level on the physical volume. This is far more efficient because unlike file deduplication, block deduplication can take 2 large files with 1 block difference and is only required to retain a copy of just the 1 differing block (Whitehouse, 2008). This means that since redundant data is referenced via pointers and is a virtual reference, data backups also benefit from block level deduplication. A real world example would be a large application code base in a version controlled system such as IBM Rational ClearCase. As the code matures and is updated into new builds, only the differing blocks of data will be required to have their own storage location. All the matching blocks of data will exist in a single location with pointers back to the common blocks. The following figure shows a progressive code set and eliminating commonality at the block level.

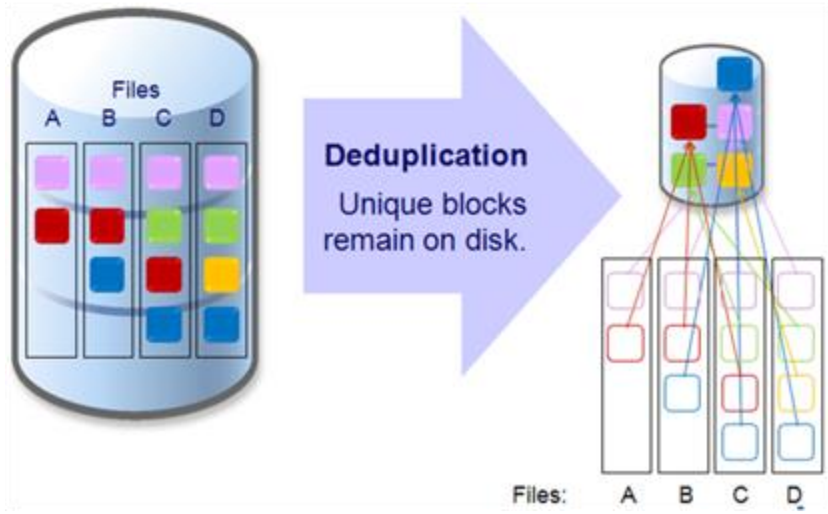

Figure 5 - Block Level Deduplication Example of Progressive Data Set (Data Protect Works, 2013)

If block level deduplication is going to be utilized in a storage subsystem, the block size of the volume being created on the storage array must be considered for the best compromise depending on what the volume is going to be used for. Using a larger block size guarantees maximum read/write performance as significant sequential data can be contained in a single block and this will minimize fragmentation. However, larger block sizes will decrease the amount of benefit that block level deduplication could provide because there would be far fewer identical blocks on the volume. On the flip side, small blocks on the storage volume could provide maximum benefits for block level deduplication for maximum storage savings but would suffer from significant disk fragmentation and poor read/write speeds. A best practices implementation is a compromise of block size based on the size of the file system. In the case of a file system of 1 TB or smaller, a 4,096 byte (aka 4K) block size provides the best efficiency for data deduplication (Symantec, 2012). In file systems greater than 1 TB in size, an 8K block size provides the best compromise of file system speed and data deduplication benefits.

As early as 2005, two companies, Data Domain and Avarmar, were filing patents for data deduplication technologies. Although they had differing algorithms, their results were very similar and the intent was the same. As of 2009, EMC now owns both Data

Domain and Avarmar. These algorithms are still used heavily today amongst all companies implementations of both file and data deduplication technologies.

3.4 Storage Replication/Mirroring

The first line of storage mirroring is the RAID levels of the storage subsystem. A storage subsystem could be as small as two drives in a workstation or server. A storage subsystem could also be as large as a full rack to multiple racks of spinning and/or solid state disks providing multi-terabyte physical and logical volumes.

Storage mirroring is very much addressed at the RAID levels within a workstation, server, or storage system and also defined as local or short distance operations. The most common mirroring is RAID 1 and is the mirror copy from one disk to another (Lubbers, et al., 2002). This is common in workstations as well as servers for the mirrored pair of disks where the operating system resides. Not only do you gain the peace of mind that you data resides over two physical disks, you also increase your read performance because there are two disks with mirrored copies of the data resulting in potentially twice the throughput. The following figure depicts that RAID 1 is an exact copy of disk 1 to disk 2 at the block level.

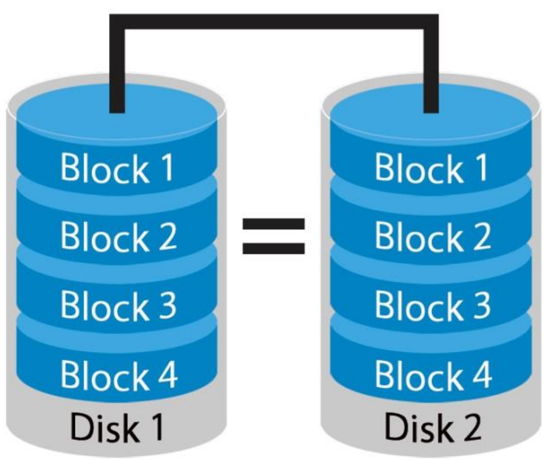

Figure 6 - RAID 1 - Disk Mirroring (Petemarovich Images, 2013)

The most common RAID level implementations in large scale data subsystems include

RAID 5 and RAID 6. These use the same technology but the RAID 6 includes one additional disk of parity. This provides the ability to lose one disk in a RAID 5 configuration or two disks in a RAID 6 configuration before eminent data loss is possible. This means RAID 6 will lose an entire disk of storage capacity at the physical volume layer but is a most attractive to taking a more risk adverse position. This is still in the realm of local or short distance operation. Because RAID 6 has an additional spinning disk compared to its RAID 5 counterpart, RAID 6 does offer a slight read and write advantage over RAID 5.

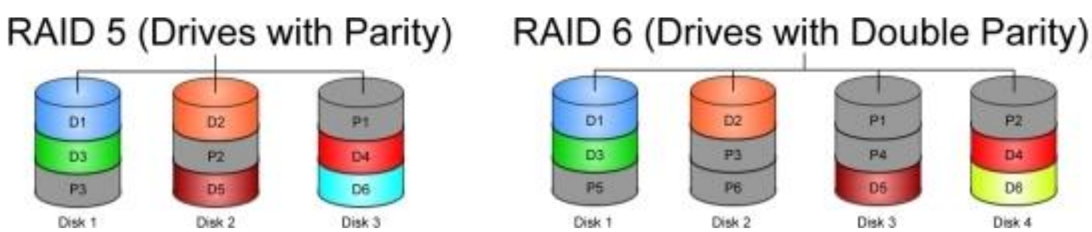

Figure 7 - RAID 5 & RAID 6 - Single versus Double Parity Disks (Net DNA)

When discussing storage replication, this refers to a far larger and more complex system. Typically, storage replication is the mirroring between two data centers independent storage subsystems. This type of storage replication is referred to active real-time and allows the fault tolerance of an entire downed data center. This is costly to provide an entire storage subsystem for the only purpose of maintaining an online copy in a hot standby scenario should one data center go down due to a power outage or maintenance window. This type of redundancy is primarily for mission critical systems that require a 100 percent uptime. Active storage replication is implemented by distributing updates of a block device to both storage subsystems that reside in separate data centers. This way, any writes to the file system level that resides on top of those blocks are independent to the mirroring happening at the block level underneath.

There are two types of storage replication: synchronous and asynchronous. Synchronous storage replication guarantees zero data loss by means of a write completing on both storage sub systems or not at all (Calder, et al., 2011). An acknowledgement by both local and remote storage elements will determine if the transaction is completed before proceeding with additional work. The result of this is overall decreased performance. An additional downside to a synchronous storage replication setup is that if either storage subsystem goes down or the medium connecting the two data centers is unavailable, the file systems shared between the data centers become read only. This stopping of all writes ensures zero data loss. Synchronous data replication is most common when running virtual machines on shared storage to ensure during an outage that the virtual machine can migrate without disk corruption.

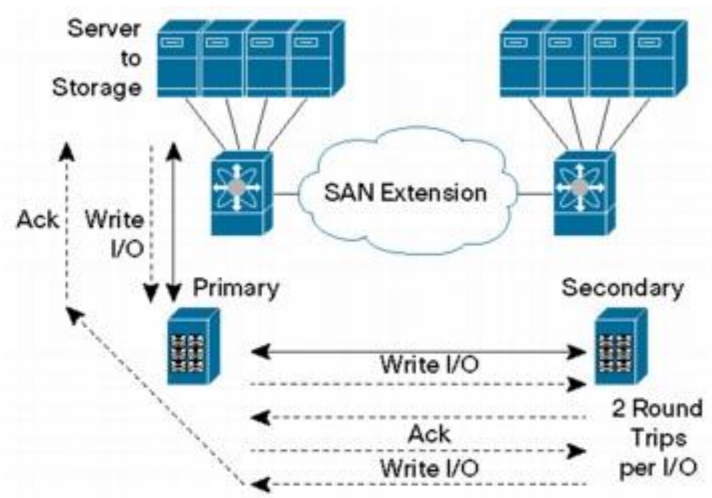

Figure 8 - Synchronous Data Center Transmissions (Cisco, 2006)

Asynchronous storage replication considers the write complete as soon as the local storage subsystem acknowledges the write. The remote storage element is updated with a small lag. This means performance is greatly increased but in the case of losing your local storage subsystem, the remote storage is not guaranteed to have the most current copy of data and may have the most recent data lost. Asynchronous data replication is most common with shared file systems such as home directories or large data stores. Asynchronous is also used for long distance data center replication due to the latency involved with long distances transmissions.

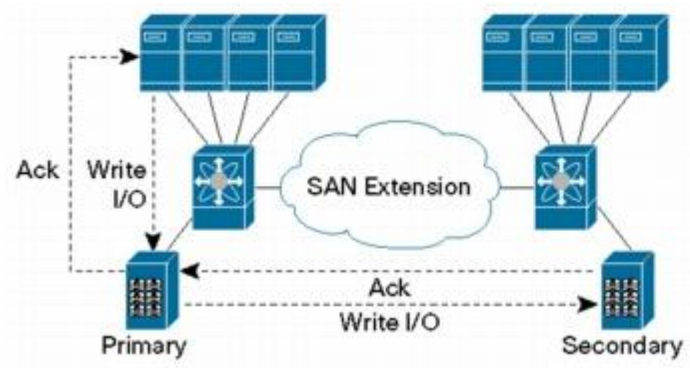

Figure 9 - Asynchronous Data Center Transmissions (Cisco, 2006)

In 2002, Hewlett Package (HP) filed and was granted the paten for “Storage area network, data replication and storage controller, and method for replicating data using virtualized volumes”. This paten was the foundation for the multi data center replication technologies. In 2011, Microsoft filed patens for synchronous and asynchronous data replication enhancing the original capabilities of simply mirroring blocks of data.

Applications of these storage mirroring and replication technologies exist by all the major storage corporations in both hardware and software based implementations. Larger storage single solutions by companies like EMC and Hitachi Data Systems will integrate these capabilities into the hardware whereas smaller installations will be software based such as DataCore’s SANsymphony product or Symantec Data Replicator.

3.5 Non-disruptive Data Migration

Data migration is the process of transferring data between storage types, formats, or systems. Data migration can be done intra-data center or inter-data center and can also be performed between homogeneous and heterogeneous data storage systems (Radhakrishnan & Singh, 2010). Regardless of the reasons or type of data migrations organizations are finding themselves having to perform, keeping this process the least disruptive as possible is often a lofty yet common goal. Downtime means disruption and expensive loss of business (Brocade, 2009).

One of the major benefits is that since the physical disk is abstracted from the host system, the host only knows about the logical disk (LUN). This means that data can be moved block to block while the host system continues to operate during the migration. After the migration occurs, the metadata is updated to ensure the host points to the new storage location. An example of this would be moving a virtual machine from one data centers storage system to another in order to perform maintenance in the primary data center. Another example would be migrating a share from slower SATA (Serial Advanced Technology Attachment) storage to faster SAS (Serial Attached SCSI) storage based on the needs and requirements of the enterprise. We will investigate the details of tiered storage in an upcoming section.

3.6 Compression

Data compression has been used for quite some time in storage systems. Compression can impose a performance penalty because of the encoding and decoding process when the data is accessed. Many file types of related to media such as JPEG, MPEG, and MP3 formats are already compressed thus not allowing additional compression. Text files, especially with significant white space, are great candidates for compression.

Almost all enterprise level backup software today uses some level of compression when taking data from a storage volume onto tape. As stated before, compression has performance penalties associated with it but is it acceptable levels due to the traditional timing for data backups to occur.

3.7 Disk Expansion and Shrinking on the Fly

The expansion and shrinking of a LUN is known as dynamic LUN expansion. This means that in the background, on the fly, you are able to expand and shrink the LUN size. This is a powerful tool provided by many of the modern storage arrays for all different sized business applications. The benefit to this functionality is to allow for the addition or removal of storage without having to take down the file system. Many business run into the paradigm depicted in the figure below, which is a direct relationship of storage usage as time goes on. The blue line represents the employees' usages of a particular storage share over time. The red line represents a 100 GB increase in the storage size every other month to increase the capacity. This is an example of where having the ability to dynamically modify the underlying storage system, and thus the file system size on the fly is a handy feature and storage virtualization technology. This also provides the opportunity for enforcement of storage management and the SLA while maintaining a high percentage of file system utilization.

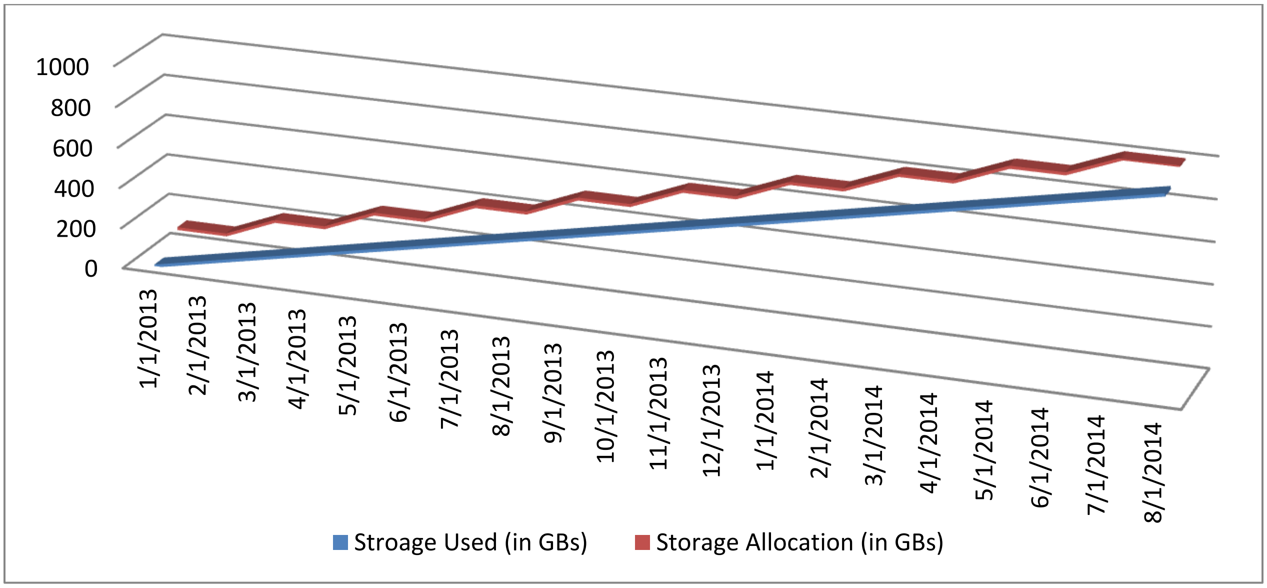

Figure 10 - Storage Subsystem Dynamic Allocation Increasing

This particular technology is very crucial to mission critical systems that do not allow downtime. In Section 3.5 we discussed the non-disruptive data migration capabilities that some modern storage virtualization applications provide. This capability of resizing LUNs and file systems on the fly goes hand in hand with non-disruptive data migrations.

3.8 Manual and Automated Tiered Storage

Tiered storage is a combination of different classes of storage systems and data migration tools that enables administrators to align the value of data to the value of the storage container in which it resides. Manual refers to conducting the data migrations, as previously discussed in Section 3.5, by hand. Conducting the data migrations automatically using data migration or storage virtualization software by setting up access policies denotes automated tiered storage.

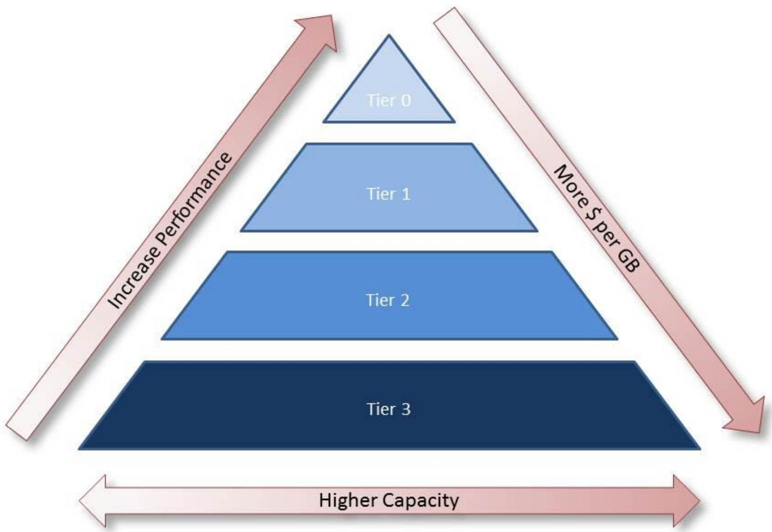

Figure 11 - Tiered Storage Sizing and Cost Relationship

In tiered storage, there are typically three levels of storage classes. Level 1, being the fastest read and write capabilities but usually the smallest tier in terms of terabytes, consisting in a RAID of SAS disks. Level 2 includes more storage capacity than level 1 but suffers from slower read and writes speeds contained in a RAID of SATA disks. Level 3 is the largest and slowest tier of storage. This includes the robotic tape library for long term storage with offsite mobility. An even larger more dynamic offsite storage capability that is taking over robotic tape backup solutions is the cloud. This is discussed in length in Section 5. The robotic tape solutions will never go away for classified or sensitive information backup solutions for the industries that require it.

Just recently, another tier of storage has been added, which is level 0. This refers to the SSD level of storage, even faster than SAS but again, smaller in terms of terabytes.

This is to be used sparingly due to the cost per gigabyte. Larger storage arrays enable enterprises to deploy differing storage tiers thus allowing for in-chassis data migrations.

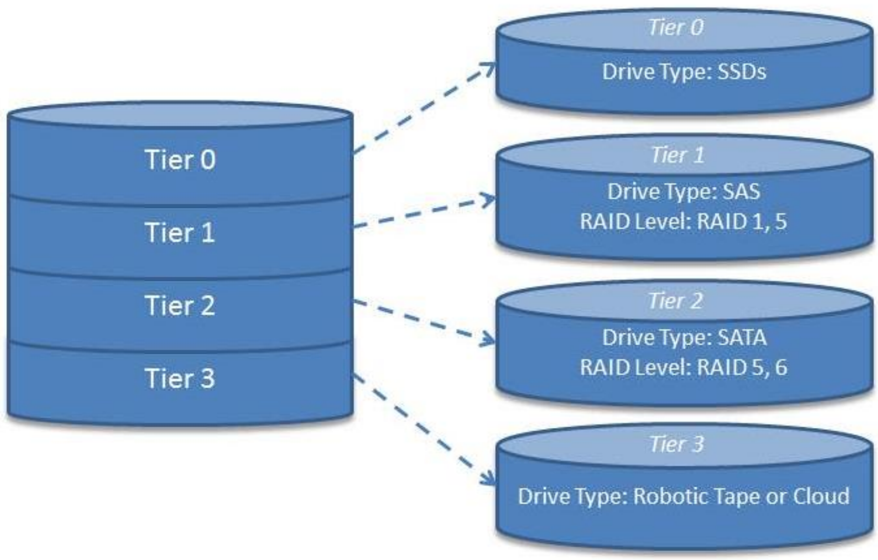

Figure 12 - Storage Tiers and Recommended Drive & RAID Levels

The automated tiered storage software can incorporate policy based migration of data based on a multitude of criteria. The primary criterion is based upon frequency of read, write, or access times. This is more granular than the manual tiering because the automation can go down to the file or block level and put specific files into differing tiers without manual intervention.

In a manual tiered storage setup, there exists the capability to specify the level of tier for a specific LUN or "tag" that is placed on the file system or data share. This is at a higher level and less granular than the automated solutions. For some smaller enterprises, a manual specification of a LUN or data share to a specific tier is sufficient.

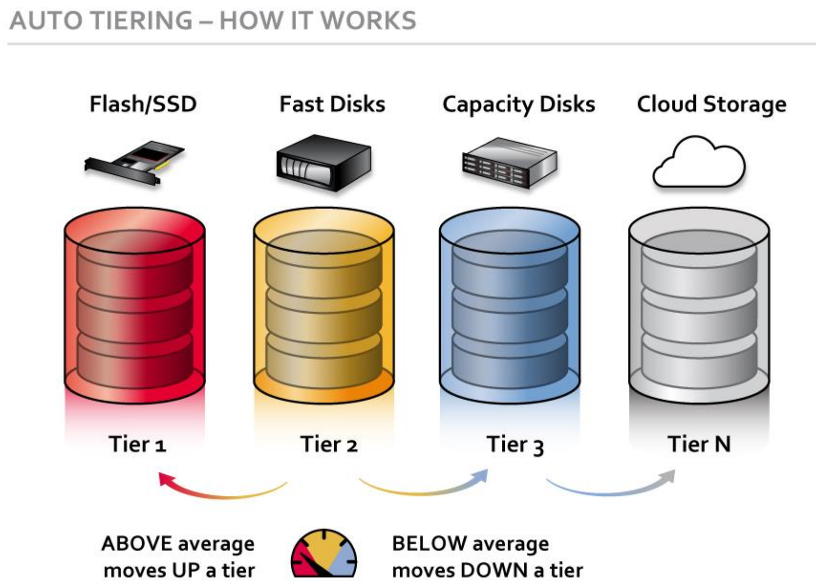

Figure 13 - DataCore SANsymphony Auto Tiering (DataCore Software)

Figure 13 depicts DataCore SANsymphony products ability to do auto tiering. Frequently access files and volumes are automatically moved to a higher tier of storage to optimize performance.

Another name for storage tiering is Hierarchical Storage Management (HSM). This was developed by IBM for their mainframe platforms and debuted on their 3850 mass storage platform in 1976.

3.9 Perpendicular Storage Recording

Perpendicular storage recording allows achieving of higher densities by aligning the poles of the magnetic elements, which represent bits, perpendicularly to the surface of the disk platter. This was first proven by a professor at the Tohoku University in Japan in 1976 but was not commercially implementing until 2005. Perpendicular storage recording delivers more than three times the storage density of traditional longitudinal recording. During its first implementation, the longitudinal approach resulted in 100 to 200 gigabits per square inch. Perpendicular storage recording allowed for 1,000 gigabits per square inch in 2010.

In addition to increased storage capacity, there was also an increase in performance. Data transfer rates increased up to 30 percent from longitudinal storage recording hard drives. This is due to the shorter distance that each bit represents on the disk.

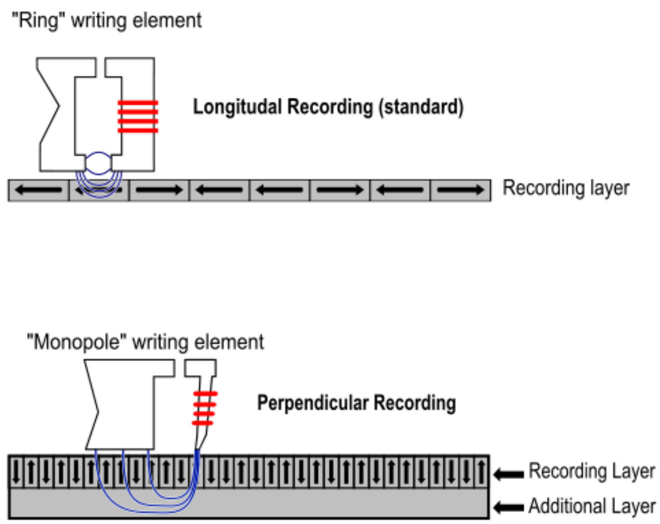

Figure 14 - Longitudinal Recording versus Perpendicular Recording (Cassioli, 2005)

How this relates to storage virtualization is that being able to fit more data onto a single hard drive means that less spinning disks are required to establish the same data requirements of the enterprise. Less spinning disks means smaller data center footprint and inevitably lower electricity usage.

3.10 Summary

The combination of the storage virtualization technologies mentioned in this section roll up into a consolidated scalable storage subsystem solution that meets the needs of the enterprise. Most solutions would not include all the technologies listed in this section, predominately due to cost. A subset of these technologies implemented would allow the enterprise reap the benefits of storage virtualization. In the Appendices of Section 8, we will see differing levels of storage virtualization technologies being implemented due to procurement cost and needs of the business.

Many of these technologies lend themselves to obtain greater success when used in conjunction. For example, combining storage replication with non-disruptive data migration is a natural fit. This is because if you have a multi data center setup that contains storage replication capabilities, then it would be expected that you can move LUNs and data shares between facilities without disruption. And many times, with Commercial off the Shelf (COTS) software or hardware appliances, these capabilities are presented and used together to compliment the capabilities of the entire storage subsystem suite of equipment.

4 Green Energy Technologies Research

Sustainable Energy is energy that can be kept up well into the future without harmful repercussions for future generations. Renewable Energy is energy that when used is not depleted. Examples of sustainable and renewable energies include geothermal, solar, wind, and hydroelectric sources.

In June 2004, delegates from 154 countries gathered in Germany for the world's first government-hosted international conference on renewable energy. The group that was created is called Renewable Energy Policy Network for the 21st Century (REN21). Now celebrating their tenth anniversary, this important group of men and women develop a yearly report titled Renewables Global Status Report (GSR). The 2014 document was released and contains over 200 pages of clear documented advancements in the uptake of renewables but also shows the long road ahead to achieve a fossil fuel free energy production worldwide. To date, fossil fuels and nuclear still make up 77.9 percent of the world's energy consumption as depicted below.

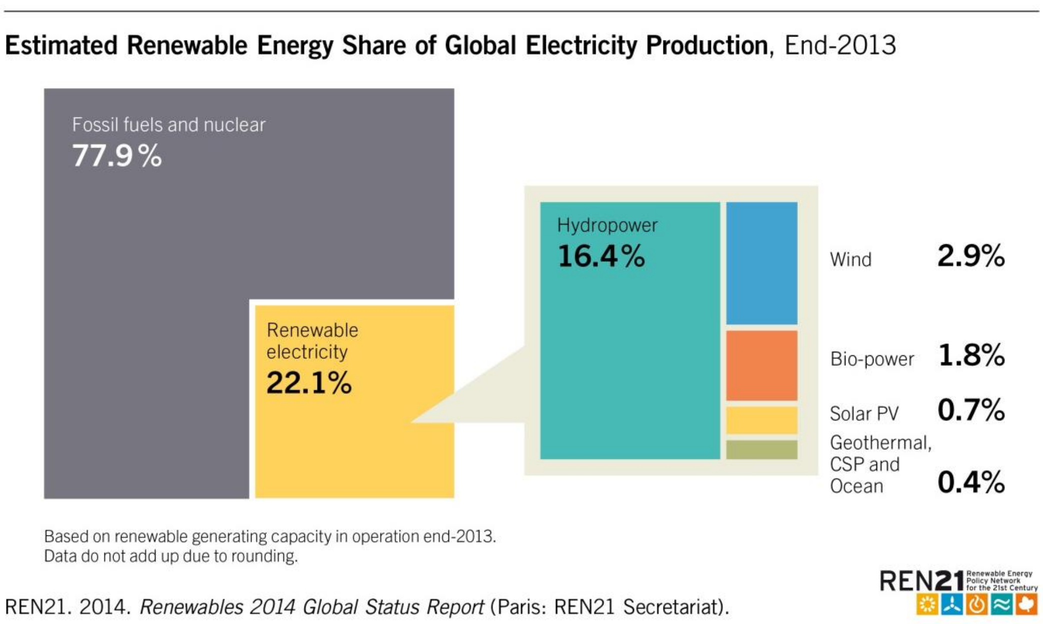

Figure 15 - REN21 2014 Global Electricity Production (Secretariat, 2014)

The executive summary of the document highlights the major countries and their contributions respective to the type of power. China was ranked first for their hydroelectric capacity, United States for geothermal, Germany for Solar photovoltaics, Spain for concentrating solar power, and Denmark for their wind power capacity.

The following subsections will go into detail about individual technologies that make each of the sustainable and renewable energies possible. These technologies will then be related back to how companies can benefit from these technologies and who are currently utilizing these systems to provide energy efficient data centers.

4.1 Geothermal Cooling

Geothermal cooling systems include an array of vertical holes drilled into the earth that house a closed-loop pipe system filled with water or coolant. The cool and constant temperature in the earth's surface allows the piping to serve as a heat exchanger (Well Drillings School). In the figure below is a simplistic view of a geothermal cooling system for home applications. In home applications, you can reverse the process at the flow center to heat the home in the winter. This example also elaborates to include a hot water tank to fully utilize the heating capabilities of this geothermal application.

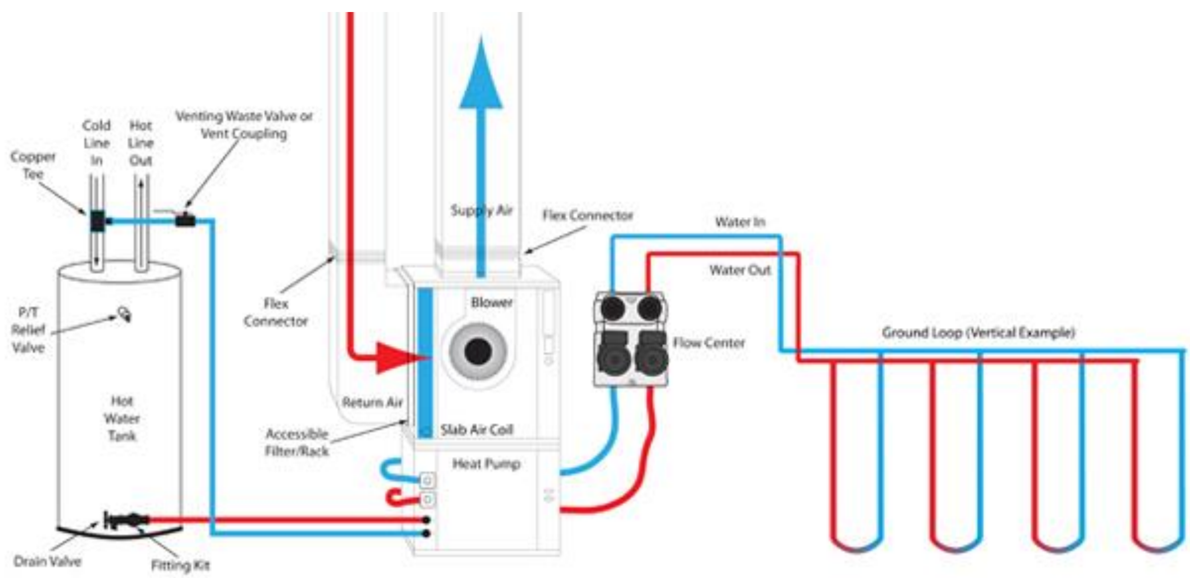

Figure 16 - Geothermal pictorial for home application (Geothermal Energy Systems, 2009)

A data center in Iowa City, Iowa at the American College Testing contains a 4,000 square-foot raised-floor data center cooled by geothermal. This is one of the first implementations of a geothermal cooling to be used in a data center.

The facility became the first data center in the United States to complete the LEED (Leadership in Energy and Environmental Design) Platinum certification. The LEED certification is a set of rating systems for the design, construction, operations, and maintenance of green buildings and facilities. There are four levels of certifications that can be obtained; Certified, Silver, Gold, and Platinum (Miller, 2009). The ratings are based on a point system based on criteria from window construction to power and cooling systems. It is a great honor to obtain this certificate, especially in a power hungry data center. And not to mention just any level of their certification, but their highest level of Platinum certified.

In Nebraska, a company called Prairie Bunkers invested in a 760-acre property near Hasting, Nebraska. The company has retrofitted dozens of World War II ammunition bunkers into data centers. They are hoping to capture the ultra-secure, risk adverse, and energy conscientious markets. Each bunker contains a geothermal cooling system drilled 500 feet into the earth (Prairie Bunkers Data Centers). This location could also obtain a LEED Platinum certification for a fully outfitted bunker data center. Reuse of existing structure, vegetation on three walls and roof, maximizing use of renewable energy in the area such as wind and hydro, and geothermal cooling provide ample justification for a Platinum certification.

Geothermal cooling dates back to the invention of the heat pump by Lord Kelvin in 1852 and was refined by Robert C. Webber in the 1940's when experimenting with the freezer in his basement. After touching the outlet pipe on the heat pump, he realized the amount of heat that was being discharged. The basic principal of the heat pump is to move heat from one place to another. In the case in Webber's basement, the heat extracted from the freezer and put into the air in his basement as waste. Being an inventor, he developed a longer exhaust pipe made of a copper coil and attached a fan to distribute the heat throughout the home.

It wasn't until the 1970's during the Arab oil embargo that the heat pump became most useful (Pyranak Design Group, 2004.). Dr. James Bose, a professor at Oklahoma State University began development on the latest era of geothermal systems. The International Ground Source Heat Pump Association was formed in Oklahoma and today resides on the campus of Oklahoma State University where Dr. Bose is the Executive Director.

The Environmental Protection Agency (EPA) has named geothermal "the most energy-efficient and environmentally sensitive of all space conditioning systems." The primary benefit applicable to data centers include cutting energy costs by 20 to 50 percent and reduced maintenance costs over traditional cooling systems. Because of the size and location of this system within a data center, the unit runs much quieter reducing noise within the data center This system also provide a more consistent temperature throughout the data center due to the constant flow of cool air as opposed to air conditioned units coming on and off as the thermostat regulates.

This technology is helping to increase efficiencies in a data center by eliminating the need to have power hungry CRACs (Computer Room Air Conditioner) and using a sustainable resource of cooling by tapping deep into the earth where the temperature will remain constant. There is also no place where a data center exists that this technology cannot be used.

4.2 Geothermal Energy

Geothermal energy is thermal energy generated and stored in the earth's crust which originated from the original formation of the planet and radioactive decay of minerals. The geothermal gradient is the difference in temperature between the core of the earth and its surface that drives a continuous conduction of thermal energy in the form of heat from the core to the surface. Geothermal power is cost effective, reliable, and sustainable. Historically it has been limited to areas near tectonic plate boundaries where the earth's heat was abundant near the surface but advances in technology has dramatically expanded the range and size of viable resources. The following figure depicts how this energy system works.

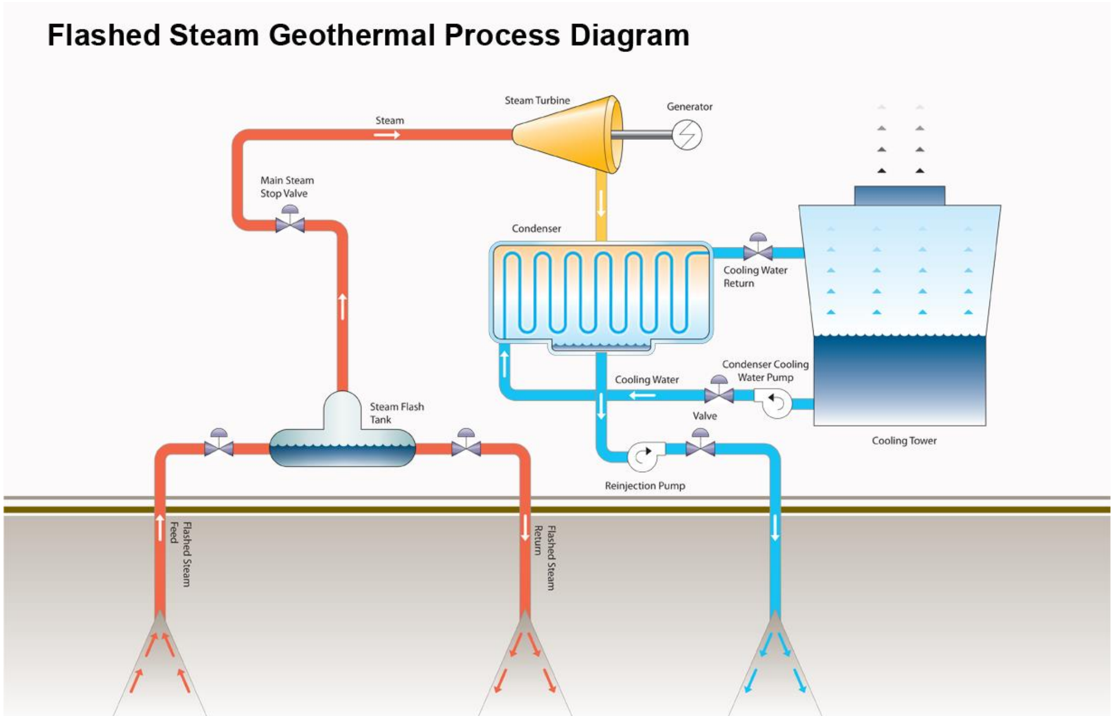

Figure 17 - Flashed Steam Geothermal Process Diagram (DiBello, 2013)

According to the 2014 Annual U.S. & Global Geothermal Power Production Report put out by the Geothermal Energy Association, the geothermal power industry continues to grow steadily at a rate of 4 to 5 percent per year. In total there are about 12,000 MW (Megawatts) of nameplate capacity today. If all current projects around the world finish on schedule, nearly 13,500 MW of capacity could be obtained by 2017 (Geothermal Energy Association, 2014). In terms of production, in 2014, approximately 7,500 GWh (Gigawatt Hours) can be attributed in one year to geothermal power product; enough to power Las Vegas for over 100 years. The numbers continue to grow as the power is sustainable and far cheaper than oil, nuclear, and natural gas sources of energy.

Despite the United States having the largest capacity of geothermal energy of any country, at just over 3,000 MW, this still makes up less than one percent of the nation's electricity production (Holm, Blodgett, Jennejohn, & Gawell, 2010). The west coast states are known for their geological activity. This activity has made California, Nevada, and Utah the three leading states for geothermal energy for production and ongoing projects. Iceland, on the other hand, an area of significant geological activity has a mere 575 MW of capacity but power 30 percent of the nation's production. A country like this can reap other benefits due to its northerly latitude as we discussed at length in Section 2.2. The cool climate in Iceland produces free natural air cooling reducing the use of electricity even further than United States locations.

How does it work turning heat into energy? A geothermal power plant pipes hot water from inside the earth through a sealed well heated by magma. This water is extremely hot and it put into pressurized chambers. The steam is evacuated from the reservoirs and piped to a large turbine in which the steam pressure spins the turbine to create energy. Water is then returned via another well to be reheated deep into the earth making this sustainable and renewable clean energy.

Around the same time as we discussed in the previous section about Lord Kevlin and his invention of the heat pump, Piero Ginori Conti, Prince of Trevignano in Italy, used steam from natural steam geysers to product electricity in 1904. By 1916, Piero developed his system to the point of distributing 2,750 kW (kilowatts) of electricity to the area surrounding the village. By 1960, Pacific Gas and Electric began operation of the first successful geothermal electric power plant in the United States. This technology became popular in Sweden as a result of the 1973 oil crisis growing steadily into the worldwide market.

The benefits of this technology should be obvious. Its clean, renewable, and sustainable nature makes it an attractive option for energy production. Unlike wind and solar, the heat in the earth is consistent and can be relied upon. One of the downsides is that there is the potential for an increased number of earthquakes due to the excessive drilling deep into the earth. Carbon emissions from nonrenewable sources may cause global warming thus melting the polar icecaps but the potential for catastrophic earthquakes due to digging deep into the earth's crust is a real concern. In December of 2006, Geothermal Explorers International managed to set off an earthquake in Basel, Switzerland measuring 3.4 on the Richter scale. The main quake was followed by 60 aftershocks in the weeks that followed.

Earthquakes typically occur around unstable areas such as volcanoes and fault lines in places like California and Iceland. These earthquakes are naturally occurring but when digging for geothermal power, introducing water into molten hot areas deep in the earth's crust cause excessive steam and pressure; enough to move bedrock in subterranean regions causing the rock to expand and contract. All energy production comes with its benefits and drawbacks. The key is to weigh the benefits with the drawbacks and pick the most viable solution for long term clean sustainable energy production.

Geothermal energy production technologies are helping to increase efficiencies in a data centers by providing clean sustainable energy to companies like BMW in the Icelandic data centers and Apple in Nevada data centers. These companies not only are using geothermal energy to power their data centers, they are collocated with the power sources. This concept is different than the traditional data centers that plug into the grid which services other commercial and residential power needs.

4.3 Solar Power

Solar power is the conversion of sunlight to electricity. This can be done with the more popular and well known photovoltaics (PV) or the less common concentrated solar power (CSP).

CSP is a system in which uses many lenses and mirrors on double axis stands to concentrate significant solar thermal energy onto a small area. This light is converted to heat which typically drives a steam turbine connected to an electrical power generator.

The global capacity for CSP is only 3.4 GW (Gigawatts) with Spain currently the world leader of with 2.2 GW of capacity. Despite the PV systems being more well-known due to their residential applications in calculators, watches, and rooftop implementations, the CSP systems are for large scale industrial usage. Below shows a figure of the CSP facility in Spain.

Figure 18 - Concentrated Solar Power Plant in Spain (RTCC, 2012)

In 1880, Auguste Mouchout from France was way ahead of his time when he was quoted saying: "Eventually industry will no longer find in Europe the resources to satisfy its prodigious expansion... Coal will undoubtedly be used up. What will industry do then?" It was around this time when Mouchout was demonstrating his parabolic trough to produce steam for the first solar steam engine. Years later, Frank Shuman from the United States developed much larger solar-powered devices. In 1913, Shuman finished a parabolic solar thermal energy station in Egypt used for irrigation. Professor Giovanni Francia designed the first commercial concentrated solar plant near Genoa, Italy in 1968 where the plant produced 1 MW of capacity using superheated steam.

In the previous section, we mentioned Apple using geothermal power for its Reno, Nevada data center. This clean energy usage will be coupled with its solar facility that Apple is developing with the local utility company NV Energy. This goal of this system is to contain a 20 MW capacity; more than enough to power the data center and the rest being fed into the local power grid. Apple is funding the project but giving NV Energy the opportunity to purchase some of the output directly or buy the whole facility outright. This is not new to Apple as they have similar solar farms already in operations in both their North Carolina data center locations.

This facility that Apple is building in Nevada is a large scale photovoltaic system consisting of hundreds of thousands of solar panels over 167 acres. Photovoltaic systems can be scaled down for home or small business use as well. The largest photovoltaic solar array is the Agua Caliente Solar Project is Southwest Arizona. This array has a capacity of 290 MW total nameplate capacity and contains more than 5 million solar panels. This solar farm is so large it can be seen from space. Considering the amount of sun produced in the desert, this is a prime location for solar farms such as Agua Caliente. It is also conveniently located near California where major corporations have large scale data centers powering their data processing subsystems and providing cloud based services for customers. Even closer to home are roof mounted solar arrays that reside with the data centers themselves. Googles headquarters in Mountain View, California has 1.8 MW of solar capacity all on the rooms of its buildings as seen below.

Figure 19 - Google Headquarters Solar System (Global Energy Network Institute, 2014)

The global operating capacity for the photovoltaic solar market hit the 100 GW milestone in 2012 and by the end of 2013 has a global total capacity for solar PV at 139 GW.

4.4 Wind Power

Wind power is the conversion of wind energy into a useful form of energy. In the context of providing clean renewable energy for data center usage, the wind power is converted into electricity from wind turbines. These turbines are connected to an electric power transmission network and connected back to the local power grid. Countries like Denmark are generating more than 25 percent of its electricity from wind. By the end of 2013, the wind energy production accounted for 2.9 percent of total worldwide electricity usage and growing rapidly. Over 200,000 wind turbines worldwide are operating with a nameplate capacity at over 318 GW; 26 times more than geothermal and over twice the amount of solar energy production.

Wind farm construction for onshore applications is far cheaper than conventional fossil fuel plants. Offshore wind is steadier and stronger than on land and have less visual impact, however construction and maintenance costs are considerably higher. All wind energy production results in no greenhouse gas emissions during operation. Below is an image of an offshore wind farm in the North Sea between the United Kingdom and The Netherlands.

Figure 20 - Offshore Wind Farm in the North Sea (Vattenfall, 2009)

Wind power history dates back thousands of years; ever since humans have put sails on boats to propel their vessels out to sea. The first windmill to source electricity was built in Scotland in 1887 by Professor James Blyth of Anderson's College in Glasgow. The wind turbine was made of wood and cloth was installed at the University to first power the lighting in the cottage. This was the first home in the world to have its electricity supplied by wind power. Back in the states, specifically Cleveland, Ohio in 1888, Charles Brush built a 60 foot tower with a 56 foot diameter rotor. This wind turbine was rated at 12 kW and was used to power 100 incandescent light bulbs in Brush's laboratory. It was years later before wind energy was deemed consistent enough for power generation due to the dependency on the weather and inconsistent power output.

Google may not have any clean energy sources collocated with their data centers but they do claim to use 50 percent less energy than typical data centers by raising the temperature to 80 degrees Fahrenheit and use outside air for cooling. That's not to say Google is not committed to cleaner energy promoting a better future. They have purchased 1,040 MW of wind energy in the form of a Power Purchase Agreements (PPA's). This is a long-term financial commitment, typically 20 years, to buy clean energy from a particular producer. Google has spread out their commitments amongst wind farms throughout the world to offset their electricity usage when plugged into the grid.

A Power Purchase Agreement works by providing developers with a solid commitment of financial support, and in exchange, the supporting company receives clean energy at competitive prices as well as the Renewable Energy Certificates (RECs) to reduce the companies' carbon footprint. This is Google's current vehicle for reducing their carbon emissions and helping by investing in the renewable energy market for a brighter future.

Moreover, Google has committed over $1 billion to renewable energy projects for large-scale wind and solar projects. These projects represent over 2 GW in total energy capacity. This is far more than what is needed for Google offices and data centers but is again showing its commitment to the industry and the need for renewable and sustainable clean energy initiatives.

4.5 Hydroelectric Power

Hydroelectricity is the term referring to electricity generated by hydropower. The earth's gravitation means water is falling or flowing in which is captured to spin turbines to create electricity. This form of clean energy is the most widely used form of renewable energy accounting for 16 percent of global electricity generation and continues to increase. By the end of 2013, hydropower encountered a milestone; 1,000 GW of capacity worldwide. This is greater than bio-power capacity, bio-power generation, geothermal, solar PV, solar CSP, and wind combined.

Hydropower has been used dating back to the 18th century in watermills to grind flour and perform other tasks. The Industrial Revolution drove out innovation helping to advance this technology even further. It wasn't until 1878 that water was used to generate electricity. This was developed in England by William George Armstrong. It was used to power a single lamp in his art gallery. In 1881 on the United States side of Niagara Falls, the first Edison hydroelectric power plant was operating in production. Monumental feats of engineering were to follow. The Hoover Dam was built in 1928 with an initial capacity of 1.3 GW that came online in 1936. Many dams were built over the years, each surpassing the previous Gigawatt capacity. Today's largest dam is the Three Gorges Dam in China with an operating capacity of 22.5 GW.

A hydroelectric dam works by a large scale filter at the intake point, a down sloping tunnel for the water which spins the turbine with extreme water pressure, and is released out the back of the dam gently down the river. At this point in time, Niagara Falls operates at only 25 percent of its full potential, meaning 75 percent of the water is diverted into the turbines for generating power. The capability exists to completely turn off Niagara Falls in order to produce maximum electricity if needed. Below is a pictorial view of how a hydroelectric dam works.

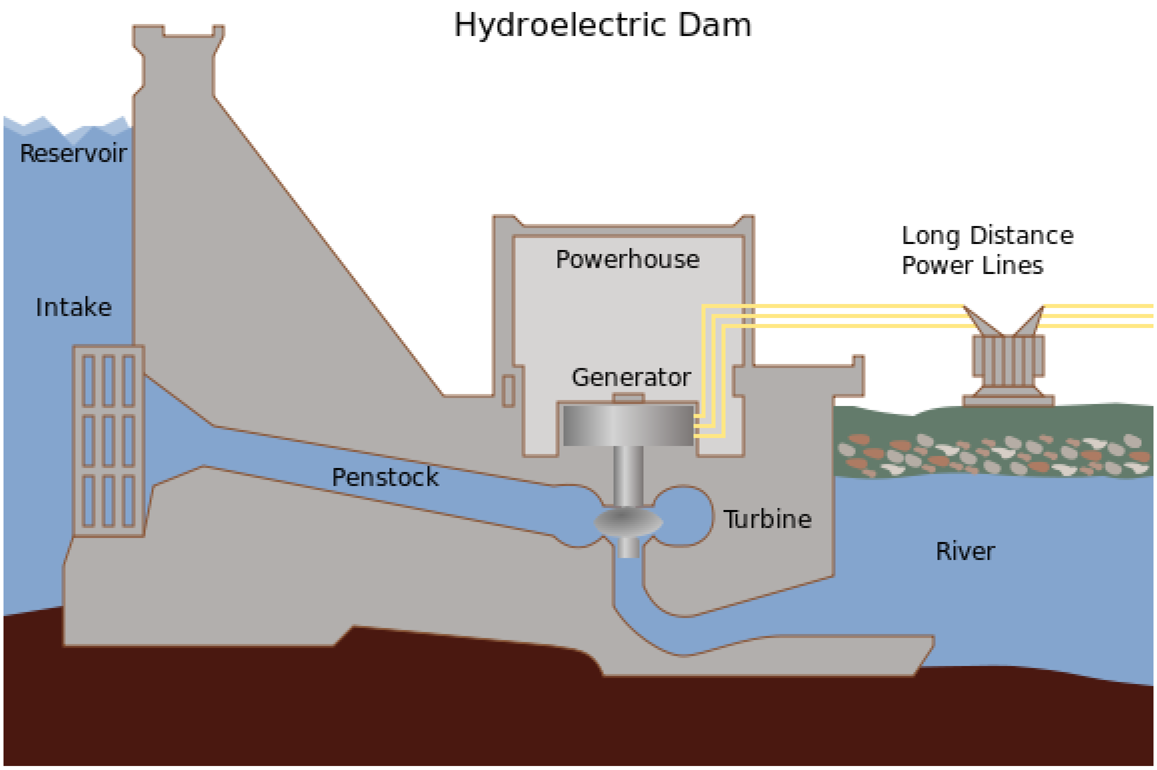

Figure 21 - Hydroelectric Dam (Tomia)

This is helping to increase efficiencies in data centers by companies such as Facebook. Facebook has a new data center in Sweden powered by 100 percent hydroelectric power. The dam doing the heavy lifting is just upstream from the town of Luleå, less than 70 miles from the edge of the Arctic Circle. Similar to Iceland, this data center uses fresh outside air to cool the data center rather than conventional refrigeration systems. The air is funneled in, filtered, and misted with humidity to ideal conditions then washed over the compute nodes. With the warmest time of the year in July, the average temperature is only 62 degrees Fahrenheit, making this location idea for running hot systems. Facebook has announced the Luleå data center runs one of the cleanest data centers in the world with a 1.04 to 1 ratio of watts of energy for power and cooling to watts for computing (BusinessGreen Staff, 2013). Traditional data centers run with a ratio of 3 to 1.

5 Cloud Services Technologies Research

Cloud service technologies refer to the remote capabilities that companies offer as a service either for free or paid. The Cloud is a network of remote servers and storage hosted around the world used to store, manage, and process data in place of local servers or personal computers. The term “The Cloud” has become a shorthand way to refer to cloud computing infrastructure and came from the cloud symbol that network engineers used on architecture diagrams to represent network segments unknown to them or out of their control. This term has been used even further in the marketing campaigns selling and promoting remote CPU cycles, software applications, and storage as a service.

Storage as a service is the primary capability being provided by companies today. This can be a free or paid service. A personal account with Google for example allocates 15 GB of data for free to be used between their Gmail email, Google+ (social networking) photos, and Google Drive (remote storage) services. To increase that storage capacity, Google charges $1.99 per month for 100 GB and incremental plans all the way up to $299.99 per month for 30 TB or storage. Google Apps is another service that increases your base storage to 30 GB and provides additional applications for $50 per user per year (Google).

Google is just one example of storage as a service in the Cloud for personal, small business, or educational use. Other companies are using a similar business model such as Apple. The Apple iCloud provides 5 GB of free storage with your Apple ID. This account is created with iTunes or when you purchase an iPhone. You can place purchased content into the iCloud and access it from any workstation, laptop, or mobile device. For additional storage, Apple charges a one-time fee of $20 for 10 GB more, bringing the iCloud total up to 15 GB (Apple, 2014). The functionality between Apple and Google is very similar but both great examples of storage as a service.

Services like Carbonite, Mozy, and Norton Online Backup exist as a different flavor of online storage. They come in the ability to conduct backups of personal workstations, laptops, and in some cases, mobile devices. These services typically run an agent on your system and backup on a schedule defined by the user. Some services offer mobile device backups and some offer unlimited storage. Most of these backup services are between $50 and $100 per year. This methodology is a different take on how personal computer backups have been done in the past. Not so many years ago, endless CDs, DVDs, and thumb drives made up the backup mechanisms for many home computers. Then the USB external hard drive was the next choice for backups due to the larger size and cheaper price per gigabyte than the other mediums were offering. Those devices and media are still susceptible to failure and negates the disaster recovery best practices of getting your data out of the same location where the primary data exists in case of fire or other total loss situation. Backing up to the Cloud provides this capability without having to purchase additional hardware or worry about finding a secure storage facility away from your home.

As we discussed in Section 2.2, BMW is using CPU (Central Processing Unit) cycles as a service by farming out to a data center in Iceland. This is helping to improve the company image with less carbon emissions and relies on a data center optimized for performance and efficiency, and 100 percent renewable and sustainable energy sources with natural cooling. The Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the Cloud. It is design to allow a user to boot a new server instance within minutes, scale the computing requirements as needed, and run the newly defined resource. Amazon EC2 allows you to pay for the server capacity (CPU, memory, and storage) for what you need in a fault tolerant efficient data center. For a business that needs high powered computing only some of the time, this is a great option to dynamically scale your virtual machine to what system requirements are needed for a specific time period and only pay for those compute cycles (Amazon).

Software as a Service (SaaS) is a software licensing model in which software is licensed on a subscription basis and hosted in the Cloud. The primary driver for this is to outsource the hardware and software maintenance and support to the SaaS provider.

This again helps to eliminate the need to hire on a systems administrator by subscribing to a capability rather than maintain it in house. An example of SaaS is GoToMeeting by Citrix. This system allows for collaboration over the Internet with a slew of online meeting tools. All you need to provide is a browser for your workstation or laptop. This solution also has an app for the mobile component. Based on the level of service your company subscribes to, you can host a video teleconference with collaboration tools with up to 100 people. And the best part is guaranteed uptime and no hardware or software maintenance costs. All content is outsourced to the destination company.

Cloud based services for storage is the way of the future for both personal use and small businesses. Larger business can create their own Cloud architecture by spreading out their storage subsystems to multiple data centers in differing geographic locations as we will discuss in Section 8.3. Centralized storage in a secure facility is a far better approach for backups and provides a central point for accessing your data from both personal computers and mobile devices. These locations will be assumed to provide the best equipment in the best of environments to optimize storage and CPU cycles for end users. Farming out CPU cycles is a better approach for businesses because centralized data centers built for the purpose of high capacity computing are optimized for efficiency and performance. Maintaining your own servers powered up 24 x 7 for a typical business week of 8 to 5 wastes significant energy. A cost benefit analysis is required for any business to determine the best approach for the needs of their particular enterprise.

6 Conclusions

With the proliferation of data by means of increased numbers of computing systems and cloud based services, it has never been more critical than now to implement proper storage virtualization techniques to bolster current storage architectures to maximize storage efficiencies. Internet speeds are increasing, short messages measured in bytes are being surpassed by multimedia, and what used to be computer based is now in mobile devices. Technology continues to evolve but is having trouble keeping up with the user base. Storage virtualization is crucial for increasing the storage size and efficiencies of the backend storage subsystem to handle this increase in storage requirements.

Also, another critical point to review is the need for sustainable and renewable sources of energy to maintain our data centers. With the increase in data means an increase in the size and number of data centers to meet the need of the industry. But when energy continues to be a major concern for reasons of environmental impact and cost, it becomes apparent that the methods in which companies and their associated data center traditionally conduct business will no longer be acceptable.

To tie this all together, one of most important examples of data proliferation and folding in energy contentiousness is Facebook. Founded in 2004, just 10 years later, the company has reported 1.23 billion users (Protalinski, 2014). This company has been a driver of the social media revolution and has harnessed the concepts of both data proliferation and sustainable energy sources. Facebook handles over 500 TB of new data every day requiring numerous data centers over multiple storage subsystems to constantly write data just to keep up, no doubt utilizing the most advanced storage virtualization technologies (Kern, 2012). In respect to their environmental conscience, as we discussed previously in Section 4.5, Facebook has deployed a data center in Sweden driven by 100 percent hydroelectric power (Vance, 2013).

Smart implementation of disk storage technologies combined with clean renewable and sustainable energy sources means long term low impact efficient storage subsystems. This is necessary for companies to remain competitive in the digital landscape of today.

7 Future Work

There are many areas in the field of storage virtualization and sustainable computing that can benefit from additional research and development.

In Section 3, there was extensive elaboration on existing storage virtualization technologies and how they complement each other to provide the best possible solution to decrease spinning disk, decrease data center footprints, and increase efficiencies for both power consumption and disk utilization. The next steps would include researching as to how the non-spinning disk storage technologies play into the data centers of tomorrow and their storage virtualization implementations. Also, as new technologies arise, determining how they play a role in storage virtualization and which current technologies are no long utilized will provide feedback as to where the industry is headed.

In addition to data centers that benefit from the local climate, like Iceland, which was discussed in Section 4.2, there are operating times daily, weekly, monthly, or annually in which additional efficiency benefits can be gained. Researching to understand these efficiencies for the best location and proper timing can further compliment these technologies to their fullest potential. Or perhaps there is a better geographical location if thinking decades into the future instead of years.

Is there merit to operating equipment for a shorter duration if the renewable energy fluctuates over time? These considerations would be applicable to locations in where the climate is fluctuating daily based on time of day, like in the desert for instance, or in latitudes in the extreme North or South to account for seasonal temperature and drastic day light swings. Researching the feasibility of a primary data center facility running in Iceland during the summer with a cold backup in the sever southern hemispheres such as South America, could provide the proper operating temperature of cool ocean regulated climates, combined with up to 20 hours of day light for a solar equipped data center. During the changing of the seasons, turn the cold backup into your primary data center to benefit from the solar while shutting down the northern data center for winter.

As far as future work for myself and personal development, I plan on obtaining the CompTIA Strata Green IT certificate (Whitney, 2010). This was designed for IT managers learning good practices and methods related to sustainable computing and green data centers as well as to why it’s important to incorporate them into an organization.

8 Appendices

8.1 Small Business Storage Virtualization Application ~$100,000

There is a multitude of vendors and products related to storage virtualization that even a small business can gain benefits from. The following subsections will provide a sample set of requirements and architecture artifacts to show how the design will meet the needs of the business. A sample Bill of Materials (BOM) for COTS software and hardware will be used to provide a sample IT system.

8.1.1 Requirements

Requirements define the needs of the business. Policies, procedures, and SLAs are put in place around the requirements to ensure these needs are met.

- Single office building with 50 employees with a mix of workstations and laptops of both UNIX and Windows operating environments.
- Budget of one hundred thousand dollars
- Nightly backups of the entire data center with a weekly offsite to a secure storage location of the entire data subsystem. Full backup every Saturday with incremental backups throughout the week. This provides full recovery capability. Backups only occur of data on the storage array and not on any individual workstations or laptops.
- Application servers, development servers, and test servers have access to the storage subsystem. Each of the three environments maintain their own shares in order to not disrupt business operations by means of segregated shares.
- UNIX based Lightweight Directory Access Protocol (LDAP) and Windows based Active Directory (AD) servers provide a single username and single password capability between the differing environments. Infrastructure services are divided between the primary and backup pair of each system instance.
- Home directories need to be 500 GB. Dev needs 2 TB of storage. Test needs 2 TB of storage. Infrastructure needs 500 GB of storage. Virtual machines supporting both the development and test systems take 1 TB of storage in total. Infrastructure systems are physical machines as part of best practices.

8.1.2 Hardware BOM

The first line item on the Hardware BOM is the server that will connect to the storage array. This type of external connection to a storage array is known as a Direct Attached Storage (DAS). The storage is attached with 6 Gbps SAS cables and presents the storage to the software running on the host system. The servers customized configuration is listed below:

Dell PowerEdge R320	PowerEdge R320, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(1) Intel Xeon E5-2430 2.5 GHz, 15 MB Cache, 6 Core
Memory	(4) 8 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 512 MB Cache
External Storage RAID Controller	PERC H810 RAID for External JBOD, 1 GB NV Cache
Hard Drives	(2) 200 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1 (Mirror)
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in	Intel X520 Dual Port 10 Gb SFP+
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise
Warranty & Service	3 Year Basic Hardware Warranty Repair, 5x10 Hardware Only
Total Cost	**$9,796.25**

Table 1 - Dell PowerEdge R320 Configuration

There are a multitude of options when it comes to rack mount server options from Dell. The R300 series product line offers hardware RAID controllers, as opposed to software RAID controllers. This provides a greater level of configuration management and better performance since there is dedicated hardware to run the local hard drives.

The two 200GB solid state drives will be configured in RAID 1. This mirrored pair of disks will run the RedHat Linux operating system for the R320 server. Single-Level Cell (SLC) flash memory is not commercial grade flash memory like Multi-Level Cell (MLC) flash memory. The SLC SSDs are meant for home use and are significantly cheaper but could degrade in performance and lifespan very quickly due to the I/O that servers produce. Therefore, the R320 was configured with MLC SSDs.

The iDRAC7 Enterprise component for system management was selected because this is not a huge cost upper but provides the capability to power on or power off the server remotely and gives secure access to a terminal window. This is a powerful tool for remote administration and keeps the management network traffic away from the in-band interface needed for storage access.

The second line item in the Hardware BOM is the storage array. This device will service all the storage needed for the small business unit by an external connect to the R320 host system described above.

Dell PowerVault MD1220	PowerVault MD1220, Rackmount, 24 Bay, Dual Power Supplies
Hard Drives	(4) 400 GB SSD SAS SLC 6 Gbps 2.5 Inch Hot-plug
Hard Drives	(10) 1.2 TB 10,000 RPM SAS 6 Gbps 2.5 Inch Hot-plug
Hard Drives	(10) Filler blank
Storage RAID Controller	(2) PERC H800 RAID Kit for External JBOD, 1 GB NV Cache
Cable	6 Gb Mini to Mini SAS Cable, Expansion Cable
Warranty & Service	3 Year Basic Hardware Warranty Repair, 5x10 Hardware Only
Total Cost	**$20,042.20**

Table 2 - Dell PowerVault MD1220 Configuration

This device can handle up to 24 drives in total of mixed drive types which allows for unique customization of the storage system. The internal RAID controller is a PERC H800, which can configure at RAID Levels 0, 1, 5, 6, 10, 50, and 60.

The four 400GB SSDs will be configured in a RAID 5 configuration. This will provide a raw configuration of 1.2TB (400GB x 4 drives, minus 1 drive for parity). This high speed optimized array will be used to host the storage for the operating system disks for the development and test systems. Each guest host virtual machine will have its own LUN for its hard disk carved up at the physical layer.

The ten 1.2TB SAS drives will be configured in a RAID 6 configuration. This will provide a raw configuration of 9.6TB (1.2TB x 10 drives, minus 2 drives for parity). This will be the slower of the 2 arrays, but still plenty fast due to the high speed spin rate of 10,000 RPM and 6 Gbps SAS drive interface. This array will be carved up into individual LUNs at the physical layer. This set of LUNs will serve home directories as well as infrastructure, development, and test data stores, totaling 5TB of storage.

The configuration of the storage to the host server, or controller, can be attached in many different configurations. The best configuration for this small business application is a unified-mode configuration where the storage enclosure is directly attached to the servers external RAID controller card with a SAS cable. Depicted below is the connection between the host controller (R320) and the storage array (MD1220). This also shows the daisy chain capability, with up to four storage trays for future expansion if needed:

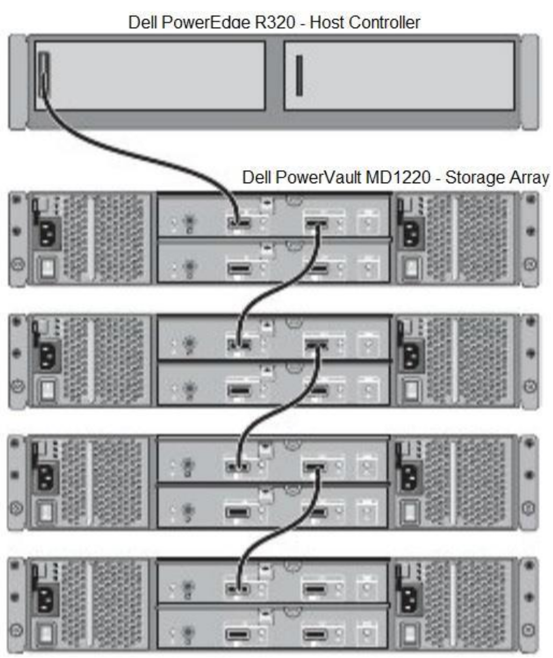

Figure 22 - Connecting the Host Controller to the Storage Array (Dell, 2010)

The Dell R700 series servers exist as a hypervisor to provide processing power and memory for the hypervisors residing on those physical systems. These systems are very robust with 32 cores and 128 GB of memory. Also purchased with these units is VMware. Each hypervisor will run VMware ESX and contain an Enterprise license so all features are activated. Locally on the system are two 200 GB MLC SSDs in RAID 1 to handle the installation of VMware ESX. Redundant power supplies were configured due to the criticality of these systems and the number of virtual machines they will be supporting. Below is the configuration table and cost:

Dell PowerEdge R715	(2) PowerEdge R715, Up to 6 Hot-plug 2.5 Inch Drives
Processor	(2) AMD Opteron 6378, 2.4 GHz, 16 core, 16 MB cache
Memory	(16) 8 GB RDIMM, 1600 MT/s
Storage RAID Controller	PERC H710, 512 MB Cache
Hard Drives	(2) 200 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in	Intel Gigabit NIC, Quad Port, Copper
Network Adapter Add-in	Broadcom 57810 Dual Port 10 Gb SR/SFP+
Operating System	VMware ESX 5.x
System Management	iDRAC7 Enterprise
Warranty & Service	5 Year ProSupport, Next Business Day Onsite (required with VMware)
Total Cost	**$34,479.08**

Table 3 - Development, Test, and Infrastructure Hypervisors

Additional hardware is required to complete the implementation of the small business storage solution. In the following table, equipment required to complete the design is addressed:

Black Box Rack	ClimateCab NEMA Server Cabinet with AC Unit, 42U
Body, Doors, Rails	12-gauge steel construction
Rack Units	42U
Capacity	2000 pounds
Total Cost	**$5,898.95**

Table 4 - Rack Assembly with AC Unit

A rack is required to house the servers, storage, and networking equipment. Since this is a small business, they will likely be able to fit all equipment into a single 42U rack. This can be place into a communications closet rather than a raised floor data center. The Air Conditioning (AC) unit that is attached will be more efficient to cool just the contents of the rack rather than cool an entire room.

The rack the equipment will also require a Keyboard, Video, and Mouse (KVM) setup to manage the systems locally. The solution selected can connect up to 16 devices and provides local authentication with username and password for security:

Console KVM	Tripp Lite B040-016-19
Ports	16 KVM systems
Monitor	19" 1366x768 resolution
Total Cost	**$1,539.99**

Table 5 - Tripp Lite KVM Console

In order to connect the host controller of the storage array to the rest of the development, test, and infrastructure environments, we need to provide network equipment capable of handling the throughput and communicate backbone for the enterprise needs. Below is a Dell N3048 rack mountable switch to facilitate the connectivity to the rest of the environments:

Dell Networking Switch	Dell N3048 Layer 3 Switch
Connectivity	48 x 1 GbE, 2 x 10 GbE SFP+
Rear Bay Uplink Module	2 x 10 GBase-T (RJ45)
Stacking Cables	(1) Stacking Cable
Total Cost	**$4,813.94**

Table 6 - Dell N3048 Layer 3 Switch

This switch has ample ports to connect the servers contained in the single rack of equipment to each other with ports to spare. The switch fabric capacity for this device is

212Gbps. This far exceeds the needs of the business at this point in time. There is a capability to stack these switches if needed in the future which is addressed in section 8.1.8.

Another option that was selected for future proofing this storage system solution was the selection of the Intel X520 Dual Port 10 Gb NIC card for the host controller. This option will be used immediately to uplink to the 10 Gb ports on the R320 to the N3048 switch. This will ensure that hosts connected at 1 Gb cannot saturate the entire in-band storage interface for the host controller and storage array.

8.1.3 Software BOM

The RedHat Linux server 6.X license for the R320 host controller was purchased with a 5 year license and 1 virtual guest. RedHat allows the procurement tiers of different virtual guests if you plan to host virtual machines on your RedHat server with a functionality called RHEV (RedHat Enterprise Virtualization). This feature will not be used in this environment because the R320 is the host controller for the storage and will not be running virtual machines. A bare metal installation of RedHat 6 will run on the RAID 1 SSDs. RedHat is an open source operating system, however the purchase will provide fixes and support if necessary by the local administrator of this sample small business.

Dell OpenManage comes with the MD1220 disk array and will be used to configure and monitor the disk array. It will be installed into the bare metal RedHat Linux instance on the R320 and accessed locally of through a web browser. It is supported by Internet Explorer (IE), Mozilla Firefox, and Google Chrome.

The Dell R320 running the Dell OpenManage storage software will also be the NFS host for this environment since this is a direct attached storage configuration. After the physical and virtual disks have been created as described in Section 8.1.7, the RedHat host will be able to configure its /etc/exports file and restart the nfs daemon. These shares are now available for use by the development and test virtual machines as well as the user community.

8.1.4 Services BOM

Because this is a small business, it doesn't make sense to add complexity by purchasing a backup solution that requires significant administrator intervention. For a business of this size, it is cost effective to purchase a backup service through one of the many options. The solution I have chosen is the Server Plus Plan by Carbonite. This service provides may features including an unlimited number of systems capped only by the number of terabytes in the plan that was purchased. In our case, our plan included 7TB of space to account for 100 percent utilization of all shared storage, virtual machines, and critical infrastructure systems such as Windows Active Directory and LDAP. Carbonite provides a live backup of Active Directory including registry and system state for a snapshot backup approach allowing full recovery from total data loss.

Many companies are reluctant to keep copies of its most critical files in the hands of a service provider. This situation is no different but Carbonite provides the peace of mind that you information is safe by encrypting the backed up data with Advanced Encryption Standard (AES) 128-bit encryption before it leaves the host system then transmitted to the offsite cloud using Secure Socket Layer (SSL) (Nadel, 2012).

Below is the single line item bill of materials for the small business services:

Carbonite Backup	Carbonite Server Plus Plan – 7TB
Total Cost	**$7,399.34 per year**

Table 7 - Carbonite Server Plus Plan Cloud Backup (Carbonite)

8.1.5 Total Cost of Solution

Here is the total cost of hardware, software, and services for a complete storage system solution for small businesses:

Line Items	Cost
Dell PowerEdge R320	$9,796.25
Dell PowerVault MD1220	$20,042.20
Dell R710 (2)	$34,479.08
Black Box Rack	$5,898.95
Console KVM	$1,539.99
Dell Networking Switch	$4,813.94
Carbonite Backup	$7,399.34 per year
Total Cost	**$83,969.75**

Table 8 - Total Cost of Procurement

8.1.6 Architecture and Interconnect Drawings

The next two diagrams display the rack layout and interconnect drawings. The rack layout shows both front and back of a single 42U rack. This is ample size for this relatively simple architecture with room for expansion in the future. The interconnect diagram shows both Ethernet and external SAS connection to the JBOD (Just a Bunch of Disks).

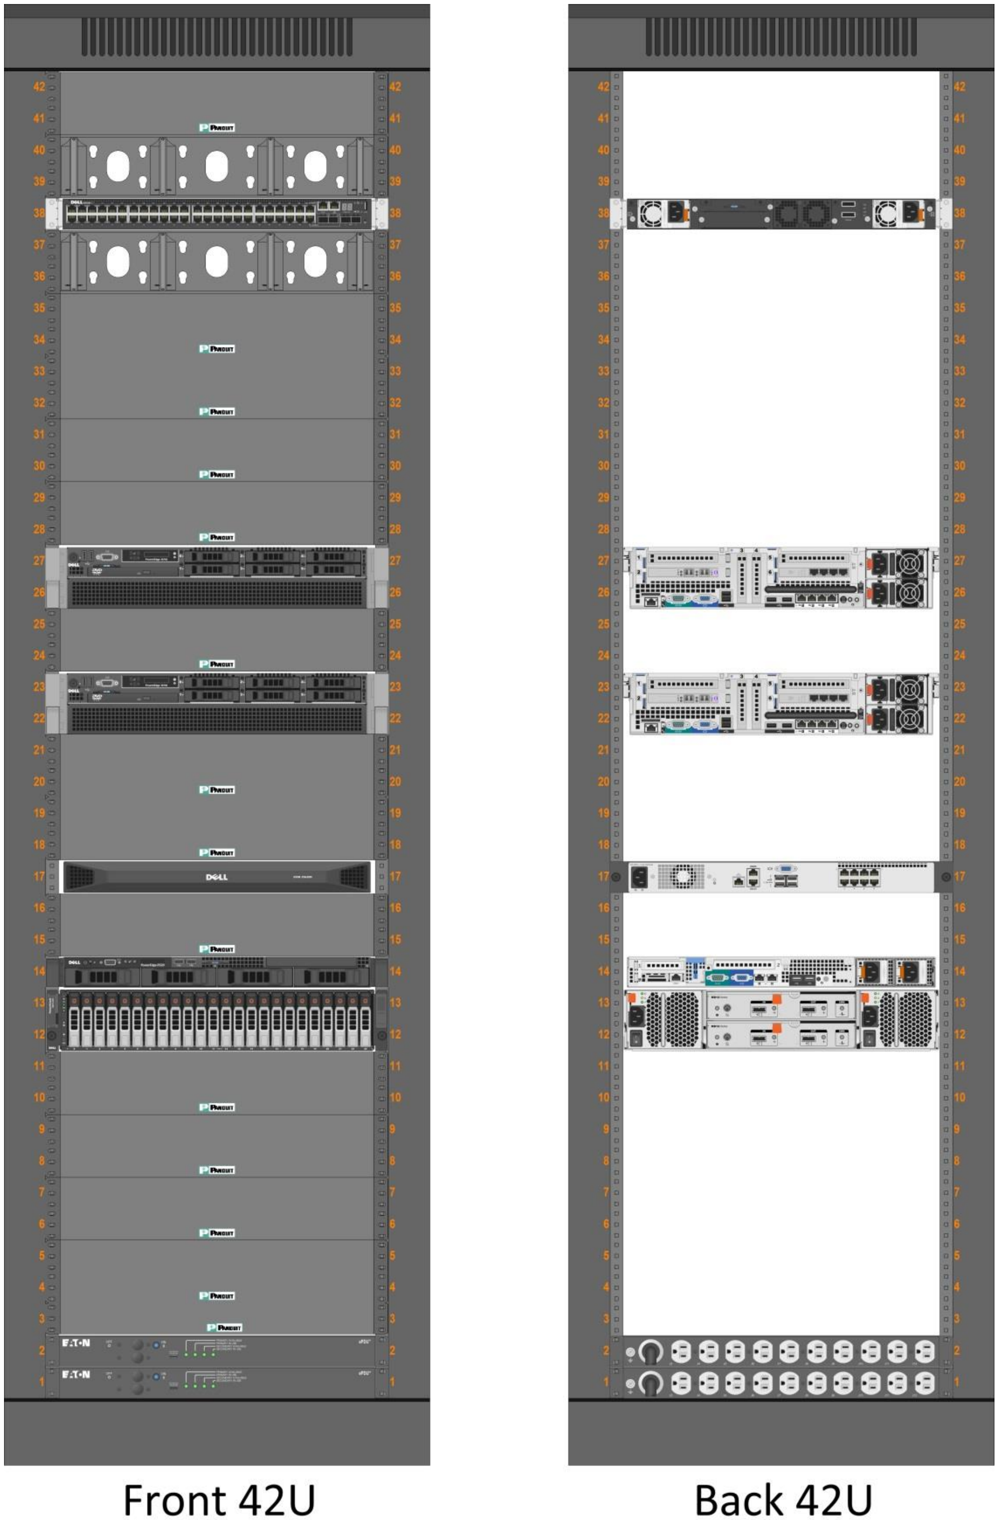

Figure 23 - Rack Layout

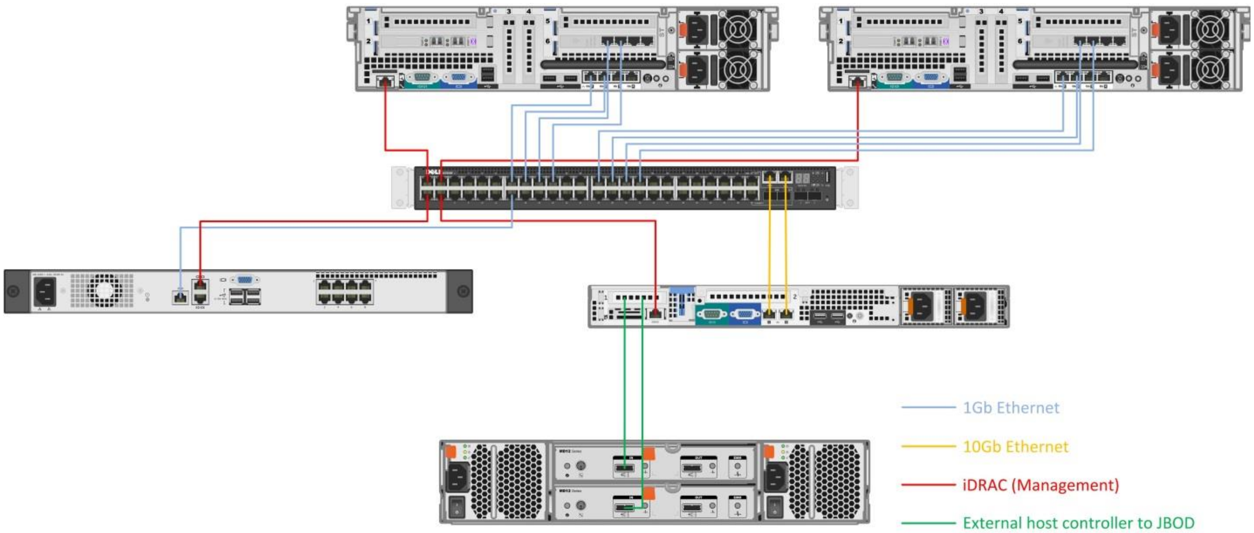

Figure 24 - Interconnect Diagram

8.1.7 Storage Virtualization Design

As one would expect, basic storage virtualization technologies exist with the MD1220 disk array. Relating this device back to Section 3 and the storage virtualization technologies discussed at length, this device incorporates thin provisioning with over-allocation, compression, disk expansion and shrinking on the fly, and perpendicular storage recording. Perpendicular storage recording is inherent on the hard drives themselves. Disk expansions and shrinking on the fly would be used if the storage requirements ever changes in the future and the storage administrator needed to make a change to the architecture. Compression and thin provisioning with over-allocation are not currently needed in this example due to the sizing requirements of the small business. However, in the future, these are great techniques to provide more storage while preserving the same physical hardware longer into the future. These technologies are easily configured using the Storage Manager component in the Dell OpenManager software bundled with the hardware.

The figure below depicts the Dell MD1220 physical disks, physical disk groups, and virtual disks. This view shows how the virtual disks are created on top of the physical disk groups, which are created on top of the physical disks. The disk array is carved up both physically and virtually to provide the storage needed detailed in the requirements Section 8.1.1

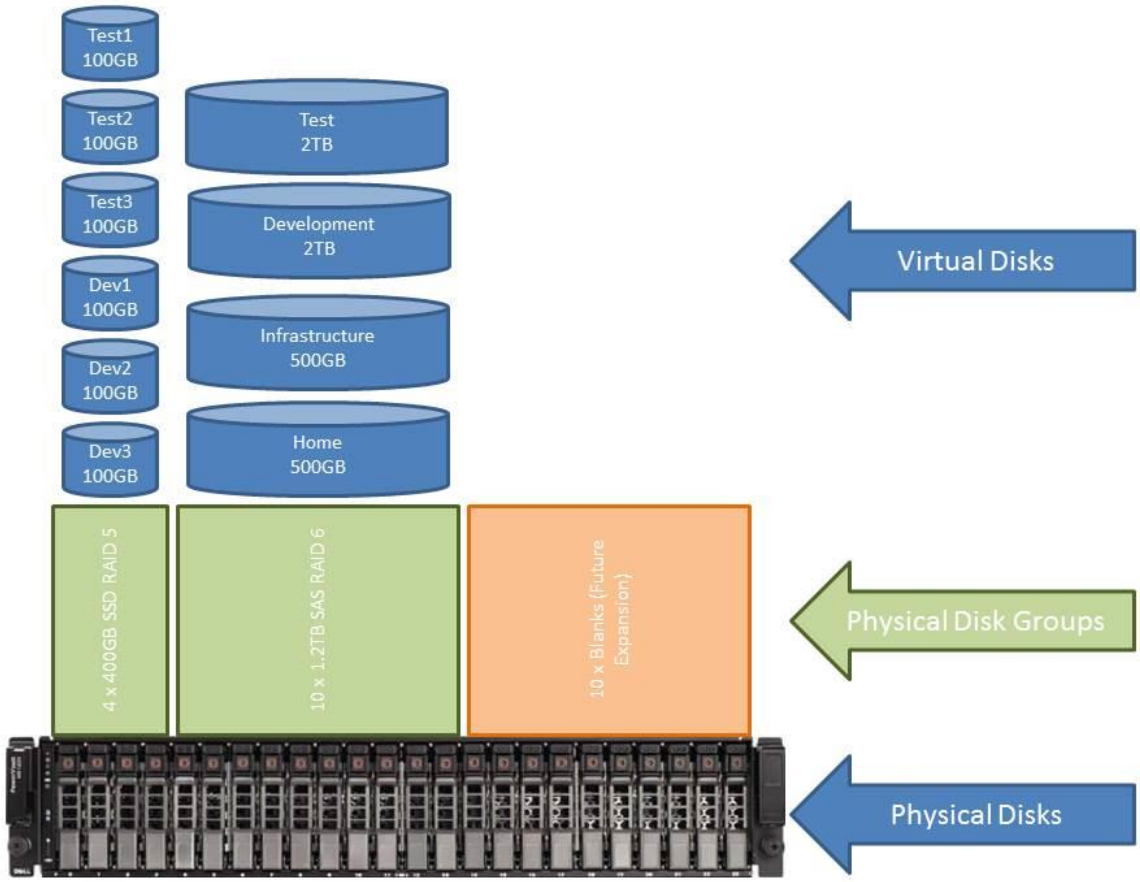

Figure 25 - Physical and Virtual Storage Allocations (SAN Storage Works)

8.1.8 Expandability & Future Proofing

When configuring the storage array, there was an option to add an additional hardware RAID controller. This option was selected with a PERC H800 RAID Kit for External JBOD with 1 GB of NV Cache. What this means is that there is an additional hardware controller that can connect to an external set of disks via the cable (also listed in the Hardware BOM) to interface between a future JBOD expansion. This option was relatively cheap but allows for an additional 192 disk devices to be configured in an attached array. The card itself is simply a PCI Express 2.0 x8 interface with a data transfer rate of 600 Mbps.

Because of the hot-plug capability of this particular server model, this means hard drives, either spinning disk or SSD, can be pulled and swapped on the fly while the system is running. For example, if a disk goes bad, due to the RAID configurations mentioned above, the bad disk can be removed and a new hard drive put into its place. The MD1220 will then rebuild the RAID itself and return to normal functionality. This feature can also be used with the ten blank slots in the MD1220. If the desire was to increase the storage space of either the RAID 5 of SSDs or RAID 6 of SAS drives, then the feature exists to procure additional drives and insert them into the storage array while running the system. The OpenManage software suite will recognize the new drive and manual intervention will add that device to the running cluster and immediately integrate the drive for additional storage to be allocated and increase read and write performance.

The N3000 series switch has a stackable feature for expansion needs. This single switch can be daisy chained with up to twelve total switches equating to a total of six hundred twenty four 1 GbE ports. This provides significant scalability for future proofing the network architecture.

8.1.9 Complexity

The initial setup and configuration of this storage system solution would require a fair amount of expertise. This would take a full time employee no longer than a month to get up and running. However, once configured and running properly, this solution would require very little intervention. The Carbonite backup system would run on schedule as configured with individual users having the capability to restore files through their web interface. The storage array was not configured to over-subscribe so there is no chance of a system crash due to overuse of the disk.

Small businesses are challenged in maintaining a full time employee for their IT systems. This is due to the fact that IT personnel are entirely overhead expenses and are not revenue generating. This makes it difficult to staff a full time IT person for such a small business. After the implementation period, a part time contractor employee knowledgeable in network, system, and storage administration would be the best person to employ. The system is not complex enough to staff individual administrators that are deep in knowledge of each domain. A basic knowledge of each domain would be sufficient for this example enterprise.

8.2 Medium Sized Business Storage Virtualization Application ~$1,000,000

As compared to the small business example in Section 8.1, a medium sized business has larger revenue, more employees, and ultimately more complex IT needs. This sample business will employ greater levels of storage virtualization due to the increased storage requirements and additional budget. A sample Bill of Materials (BOM) for COTS software and hardware will be used to provide a sample IT system.

8.2.1 Requirements

Requirements define the needs of the business. Policies, procedures, and SLAs are put in place around the requirements to ensure these needs are met.

- Single multistory office building with 500 employees with a mix of workstations and laptops of both UNIX and Windows operating environments.
- Budget of one million dollars
- Nightly backups of the entire data center with a weekly offsite to a secure storage location of the entire data subsystem. Full backup every Saturday with incremental backups throughout the week. This provides full recovery capability. Backups only occur of data on the storage array and not on any individual workstations or laptops.
- Application servers, development servers, and test servers have access to the storage subsystem. Each of the three environments maintain their own shares in order to not disrupt business operations by means of segregated shares.
- UNIX based Lightweight Directory Access Protocol (LDAP) and Windows based Active Directory (AD) servers provide a single username and single password capability between the differing environments. A Cisco ACS server, terminal servers, and shared KVM system in the lab environment also all take part in the single sign-on capability. Infrastructure services are divided between multiple physical systems as well as virtual machines residing on additional physical instances.
- Home directories need to be 10 TB. Dev needs 50 TB of storage. Test needs 50 TB of storage. Infrastructure needs 10 TB of storage. Virtual machines

supporting both the development and test systems take 50 TB of storage in total. Virtual machines are required to be on SSDs. Security requires its own share with root owned permissions of 50 TB. Infrastructure systems are a combination of both physical and virtual machines as part of best practices.

- A network Demilitarized Zone (DMZ) exists which provides a website and portal for its customers. Storage requirements for the DMZ are to be virtually segregated from the internal storage subsystem with at least 15 TB.

8.2.2 Hardware BOM

The first line item on the Hardware BOM is the server that will connect to the storage array to be used in the DMZ. This type of external connection to a storage array is known as a Direct Attached Storage (DAS). The storage is attached with 6 Gbps SAS cables and presents the storage to the software running on the host system. The customized configuration for the server is listed below:

Dell PowerEdge R320	PowerEdge R320, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(1) Intel Xeon E5-2470 v2 2.4 GHz, 25 MB Cache, 10 Core
Memory	(4) 16 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 512 MB Cache
External Storage RAID Controller	PERC H810 RAID for External JBOD, 1 GB NV Cache
Hard Drives	(2) 400 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1 (Mirror)
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in	Intel X520 Dual Port 10 Gb SFP+
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise
Warranty & Service	3 Year Basic Hardware Warranty Repair, 5x10 Hardware Only
Total Cost	**$14,595.41**

Table 9 - Dell PowerEdge R320 DMZ DAS Controller

There are a multitude of options when it comes to rack mount server options from Dell. The R300 series product line offers hardware RAID controllers, as opposed to software RAID controllers. This provides a greater level of configuration management and better performance since there is dedicated hardware to run the local hard drives. The next server in the lineup that provides a hardware RAID controller is the R700 series. This series is not cost effective for the application required.

This system is the same model that we choose for the small business application, however, this was configured with double the Random Access Memory (RAM), 10 core processor versus the 6 core, and twice the size SSDs of the MLC flavor.

The two 400 GB solid state drives will be configured in RAID 1. This mirrored pair of disks will run the RedHat Linux operating system for the R320 server. Single-Level Cell (SLC) flash memory is not commercial grade flash memory like Multi-Level Cell (MLC) flash memory. The SLC SSDs are meant for home use and are significantly cheaper but could degrade in performance and lifespan very quickly due to the I/O that servers produce. Therefore, the R320 was configured with MLC SSDs.

The iDRAC7 Enterprise component for system management was selected because this is not a huge cost upper but provides the capability to power on or power off the server remotely and gives secure access to a terminal window. This is a powerful tool for remote administration and keeps the management network traffic away from the in-band interface needed for storage access.

The second line item in the Hardware BOM is the storage array. This device will service all the storage needed for the medium sized business unit by an external connect to the R320 host system described above.

Dell PowerVault MD1220	PowerVault MD1220, Rackmount, 24 Bay, Dual Power Supplies
Hard Drives	(10) 400 GB SSD SAS SLC 6 Gbps 2.5 Inch Hot-plug
Hard Drives	(14) 1.2 TB 10K RPM SAS 6 Gbps 2.5 Inch Hot-plug
Storage RAID Controller	PERC H800 RAID, 1 GB NV Cache
Cable	6 Gb Mini to Mini SAS Cable, Expansion Cable
Warranty & Service	3 Year Pro Hardware Warranty, On-site Service
Total Cost	**$39,110.65**

Table 10 - Dell PowerVault MD1220 DMZ Storage

This device can handle up to 24 drives in total of mixed drive types which allows for unique customization of the storage system. The storage RAID controller is a PERC H810, which can configure at RAID Levels 0, 1, 5, 6, 10, 50, and 60.

The ten 400 GB SSDs will be configured in a RAID 5 configuration. This will provide a raw configuration of 3.6 TB (400 GB x 10 drives, minus 1 drive for parity). This high speed optimized array will be used to host the storage for the operating system disks for the virtual machines and any high speed database transactional file systems. Each guest host virtual machine will have its own LUN for its hard disk carved up at the physical layer.

The fourteen 1.2 TB SAS drives will be configured in a RAID 6 configuration. This will provide a raw configuration of 14.4 TB (1.2 TB x 14 drives, minus 2 drives for parity). This will be the slower of the 2 arrays, but still plenty fast due to the high speed spin rate of 10,000 RPM and 6 Gbps SAS drive interface. This array will be carved up into individual LUNs at the physical layer. This set of LUNs will serve a multitude of shared file systems within the DMZ. This enclave will require home directories, infrastructure and security partitions. The total storage of this MD1220 array will be 18 TB thus exceeding the requirement of 15 TB.

The configuration of the storage to the host server, or controller, can be attached in many different configurations. The best configuration for this DMZ web services application is a unified-mode configuration where the storage enclosure is directly attached to the server's external RAID controller card with a SAS cable. The connection between the host controller (R320) and the storage array (MD1220) is the same as depicted in Figure 22 in the Small Business Section 8.1.2. This also shows the daisy chain capability, with up to four storage trays for future expansion if needed.

This concludes the server and storage hardware needed for the DMZ portion of the architecture. The remaining portion of this section will address the physical infrastructure such as racks, KVM system, backup system, networking equipment, and additional storage and servers needed to support the development, test, and infrastructure systems internally.

The hypervisors for the medium sized business will look identical to the small business application. The exception is that there will be four in total to cover primary and backup infrastructure functionality, as well as one for development and one for test. Each server can support numerous virtual machines discussed earlier due to their 32 cores and 128 GB of memory. The odd number of the Dell R715 model indicates the AMD Opteron series processor. The reason this option was selected was due to the cost per core. The AMD series chips offer many more cores for the price as opposed to Intel thus making the AMD Opteron ideal for hypervisors.

Dell PowerEdge R715	(4) PowerEdge R715, Up to 6 Hot-plug 2.5 Inch Drives
Processor	(2) AMD Opteron 6378, 2.4 GHz, 16 core, 16 MB cache
Memory	(16) 8 GB RDIMM, 1600 MT/s
Storage RAID Controller	PERC H710, 512 MB Cache
Hard Drives	(2) 200 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1 (Mirror)
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in	Intel Gigabit NIC, Quad Port, Copper
Network Adapter Add-in	Broadcom 57810 Dual Port 10 Gb SR/SFP+
Operating System	VMware ESX 5.x
System Management	iDRAC7 Enterprise
Warranty & Service	5 Year ProSupport, Next Business Day Onsite (required with VMware)
Total Cost	**$68,958.16**

Table 11 - Development, Test, and Infrastructure Hypervisors

The next set of servers, two to be exact, are going to be used as the DataCore SANsymphony-V nodes. We discussed DataCore and their SANsymphony-V product at length in Section 3.1. In essence, the servers act as your storage controllers and provide SMB, NFS, and direct attached capabilities for all LUNs configured on the storage array. The hardware consists of dual 8 core Intel processors, 192 GB of memory and a RAID 10 with four MLC SSDs to host the operating system. The reason for the 192 GB of memory is that the SANsymphony-V product uses local memory on the nodes as cache, just as a storage controller would. Files that are accessed frequently and read/writes to the controller do not have to go back to reference the spinning disk or SSDs until later. This is a quick read/write access capability with the legwork being done after the fact.

These two storage controller nodes will run Windows 2012 R2 Datacenter Edition. Datacenter Edition is required because of the amount of memory and dual CPUs. Windows was selected because the DataCore product only runs on that platform at the moment but is to be ported to Linux for support in the near future. These nodes also contain two dual port 16 Gb fiber channel cards to connect to a Brocade fiber channel switch configured later in this section. The table below articulates the hardware components to support the SANsymphony-V product as a storage controller:

Dell PowerEdge R720	(2) PowerEdge R720, Up to 4 Hot-plug 3.5 Inch Drives
Processor	(2) Intel Xeon E5-2670 2.6 GHz, 20 MB cache, 8 core
Memory	(12) 16 GB RDIMM, 1866 MT/s
Storage RAID Controller	S110 Software RAID
Hard Drives	(4) 400 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1 (Mirror)
Network Adapter Add-in	Broadcom 5720 Quad Port 1 Gb
Network Adapter Add-in	Intel X540 Dual port 10 GBASE-T
Network Adapter Add-in	Qlogic 2662, Dual Port 16 Gb Fiber Channel
Operating System	Microsoft Windows Server 2012 R2 Datacenter Edition
System Management	iDRAC7 Enterprise
Warranty & Service	3 Year ProSupport, Next Business Day Onsite (required with VMware)
Total Cost	**$46,854.26**

Table 12 - Dell R720 DataCore Storage Controllers

The next items are the storage arrays themselves. Since we are using the DataCore SANsymphony-V product, we can use a single or variety of storage mediums and devices. The software virtualization product does not care. Coming up in Section 8.2.7, we will discuss at length the configuration applied to make the storage tiering possible without notice to the end user of performance degradation between shares. The following 2 tables will show both the fast and slow Dell PowerVault MD3660f storage arrays:

Dell PowerVault MD3660f	PowerVault MD3660f, Rackmount, 60 Bay, Dual Power Supplies
Hard Drives	(40) 4 TB 7200 RPM Near-Line SAS 6 Gbps 3.5 Inch Hot-plug
Hard Drives	(20) 1.2 TB 10,000 Near-Line RPM SAS 6 Gbps 2.5 Inch Hot-plug
Cable	(4) Multi-Mode Fiber Channel LC-LC 10 Meters
Warranty & Service	3 Year Pro Hardware Warranty, On-site Service
Total Cost	**$59,758.00**

Table 13 - Dell PowerVault MD3660f Slow Array

Dell PowerVault MD3660f	PowerVault MD3660f, Rackmount, 60 Bay, Dual Power Supplies
Hard Drives	(40) 400 GB SLC SSD 6 Gbps 2.5 Inch Hot-plug
Hard Drives	(20) 1.2 TB 10,000 Near-Line RPM SAS 6 Gbps 2.5 Inch Hot-plug
Cable	(4) Multi-Mode Fiber Channel LC-LC 10 Meters
Warranty & Service	3 Year Pro Hardware Warranty, On-site Service
Total Cost	**$133,994.95**

Table 14 - Dell PowerVault MD3660f Fast Array

The configuration of the Dell MD3660f Slow Array (Table 13) includes two RAID 6 configurations. The first RAID 6 will provide a raw configuration of 152 TB (4 TB x 40 drives, minus 2 drives for parity) of slow 7,200 RPM spinning disks. The second RAID 6 will provide twenty 1.2 TB SAS drives. This will provide a raw configuration of 21.6 TB (1.2 TB x 20 drives, minus 2 drives for parity) of high speed 10,000 RPM spinning disks.

The configuration of the Dell MD3660f Fast Array (Figure 29) includes a single RAID 6 configuration consisting of all 40 SSDs. This will provide a raw configuration of 15.2 TB

(400 GB x 40 drives, minus 2 drives for parity). The twenty 1.2 TB SAS drives will consist of a single RAID 6. This will provide a raw configuration of 21.6 TB (1.2 TB x 20 drives, minus 2 drives for parity) of high speed 10,000 RPM spinning disks.

When adding up the total storage, the slow spinning disk provides 152 TB of raw capacity. The fast spinning disk provides 43.2 TB of raw capacity, and 15.2 TB of raw solid state capacity. Raw capacity means RAID configured but no partitioning or file system has been put onto the storage. To recall the requirements, this does not quite meet the need from the perspective of terabyte for terabyte. However, coming up in Section 8.2.7, we will discuss the block level auto tiering done by the SANsymphony-V software which will meet requirements by virtue of storage virtualization technologies.

The last server required for the infrastructure is the backup server. Below is the configuration of a simple single CPU 1U server. The single CPU cuts down on licensing costs for the COTS product conducting the backups called Tivoli Storage Manager (TSM) from IBM. More on TSM in the next section - Software BOM section.

Dell PowerEdge R420	PowerEdge R420, Up to 4 Hot-plug 3.5 Inch Drives
Processor	(1) Intel Xeon E5-2450 v2 2.5 GHz, 20 MB Cache, 8 Core
Memory	(4) 16 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 512 MB Cache
Hard Drives	(2) 300 GB Solid State MLC Drives, SATA, 3 Gbps in RAID 1
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in 1	Intel X520 Dual Port 10 Gb SFP+
Network Adapter Add-in 2	Brocade 825 Dual Port 8 Gb Fiber Channel
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise
Warranty & Service	3 Year Pro Hardware Warranty, On-site Service
Total Cost	**$11,940.11**

Table 15 - Dell PowerEdge R420 Backup Server

Black Box Rack	(3) Elite Data Cabinet, 84”H (45U) x 30”W x 36”D

Body, Doors, Rails	12-gauge steel construction
Rack Units	45U
Lock	Key
Doors	Mesh
Sides	Solid
Fans	2 Fan Top Panel Module
Capacity	2000 pounds
Total Cost	**$7,002.00**

Table 16 - Rack Assembly

Racks are required to house the servers, storage, and networking equipment. This is a medium sized business, therefore needing multiple racks in a cooled room. In this scenario, the hardware can be accommodated into 3 racks. Due to the size, this can be place into a communications closet rather than a raised floor data center.

The rack the equipment will also require a Keyboard, Video, and Mouse (KVM) setup to manage the systems locally. The solution selected can connect up to 16 devices and provides local authentication with username and password for security:

Console KVM	Tripp Lite B040-016-19
Ports	16 KVM systems
Monitor	19” 1366x768 resolution
Total Cost	**$1,539.99**

Table 17 - Tripp Lite KVM Console

The networking equipment is needed to facilitate communications within the private enterprise, out to the Internet, and to support the DMZ. No routers are needed as the firewalls act as a gateway to the Wide Area Network (WAN). On the trusted side of the firewalls is the Local Area Network (LAN). Below is a simple topology using Juniper equipment to provide connectivity. The full detailed network topology will be shown in Section 8.2.6.

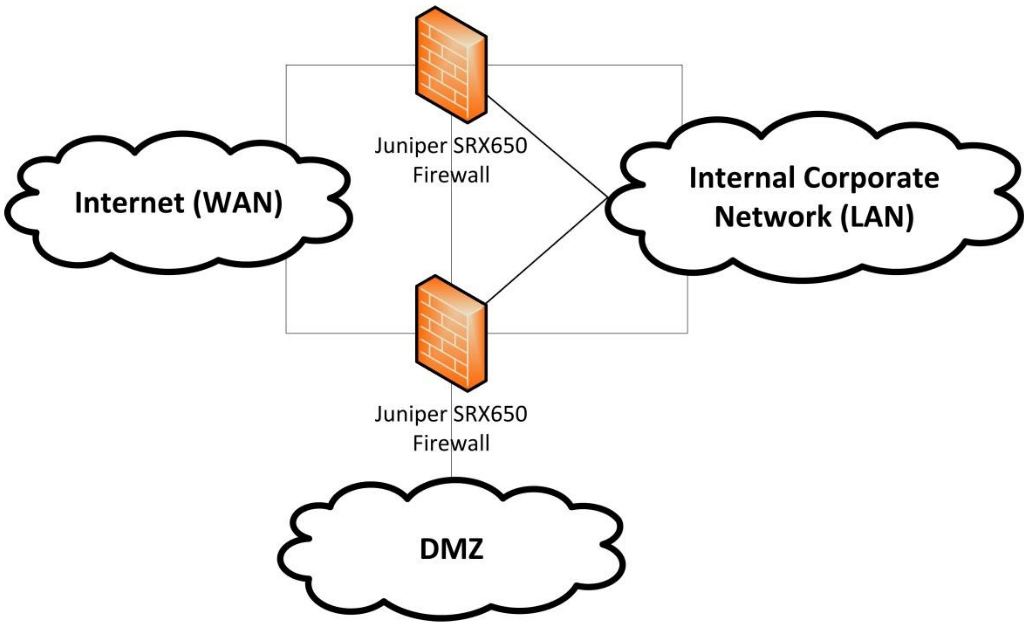

Figure 26 - Medium Business High Level Network Topology

Juniper Networks Firewall	Juniper Networks SRX650
Base Chassis	(2) SRX650 with 8 slots
Connectivity 1	(2) Quad T1/E1, 4-port with CSU/DSU
Connectivity 2	(4) 2-port 10 GbE SFP+ / 10 G Base-T Copper
Connectivity 3	(2) Ethernet Switch 16-port 10/100/1000
Support	Juniper Care Next Day Support
Renewable Protection	5-year subscription for AppSecure and IPS Updates
Total Cost	**$94,612.00**

Table 18 - Juniper SRX650 Firewalls

In order to complete the network design on the trusted side, the following switches are needed to allow connectivity for test, development, and infrastructure servers as well as each individual user station.

Juniper Networks Switch	Juniper Networks EX3300
Base Chassis	(12) EX3300 48-port 10/100/1000 Base-T with 4 SFP+ Uplink Ports
License Type	Enhanced Feature License (Supports Layer 3 Protocols)
Support	Juniper Care Next Day Support

Total Cost	**$84,371.52**

Table 19 - Juniper EX3300 Layer 3 Switches

The switches above have ample ports to connect the servers and workstations contained in the three racks of equipment to each other with ports to spare. The switch fabric capacity for this device is 80 Gbps. This far exceeds the needs of the business at this point in time. There is a capability to stack these switches which will be used in order to provide a simpler network but will be configured in detail at the software layer on the switch using Juniper's virtual chassis.

A pair of fiber channel switches to support fabric A and fabric B for redundancy will be required to connect the DataCore nodes, storage subsystem, backup server, and tape library. In the table below is a configuration of Brocade fiber channel switches:

Brocade	Brocade 6510
Base Chassis	(2) 24-port 16 Gbps Fiber Channel Switch
Warranty & Support	3 Years
Connectivity	(48) SFP+ 16 Gbps
Total Cost	**$31,855.98**

Table 20 - Brocade 6510 Fiber Channel Switches

The fiber channel switches supports a redundant interconnect architecture. This is needed for the DataCore nodes, storage subsystem, backup server, and tape library. The configuration of the devices above are full 3 year parts and maintenance, includes all SFP (Small Form-factor Pluggable) connections, and has a fabric backplane of 768 Gbps in a 1U form factor. The management of the switch is done though the management Ethernet port via HTTP (Hypertext Transport Protocol) through a compatible web browser. This allows for configuration of the device for zoning and port up/down for security purposes.

The last hardware device to discuss is the tape library. The device chosen is the IBM TS3200 rack mount tape library with an additional storage sled. This device can backup to a single tape at a time with a total of 40 tapes in the device. The TS3200 has fiber channel connections to provide LAN-free backup capabilities. More on LAN-free backup capabilities in the next section when discussing TSM.

IBM	Tape Library TS3200
Base Chassis	(2) LTO 6 Fiber Channel Drives
Tapes	(10) LTO 6 Data Cartridges 5-packs
Warranty & Support	3 Years
Connectivity	(2) 8 Gbps Fiber Channel Port
Total Cost	**$35,360.00**

Table 21 - IBM TS3200 Tape Library

As configured, this device uses Linear Tape-Open (LTO) 6 drives. The LTO 6 tape was released in 2012 and can hold up to 2.5 TB each uncompressed and up to 6.25 TB compressed. This means the tape library has a theoretical maximum capacity of 250 TB, exceeding the full storage requirements for this business unit.

8.2.3 Software BOM

The RedHat Linux server 6.X license for the R320 host controller was purchased with a 5 year license and 1 virtual guest. RedHat allows the procurement tiers of different virtual guests if you plan to host virtual machines on your RedHat server with a functionality called RHEV (RedHat Enterprise Virtualization). This feature will not be used in this environment because the R320 is the host controller for the storage and will not be running virtual machines. A bare metal installation of RedHat 6 will run on the RAID 1 SSDs. RedHat is an open source operating system, however the purchase will provide fixes and support if necessary by the local administrator of this sample small business.

The Microsoft Windows Server 2012 R2 Datacenter Edition was configured as part of the servers thus not requiring additional licensing purchasing. As noted before, Datacenter Edition was required due to the amount of RAM configured on the systems and dual CPUs.

For the DMZ enclave, Dell OpenManage comes with the MD1220 disk array and will be used to configure and monitor the disk array. It will be installed into the bare metal RedHat Linux instance on the R320 and accessed locally of through a web browser. It is supported by Internet Explorer (IE), Mozilla Firefox, and Google Chrome.

For the internal storage subsystem, Dell OpenManage will be used to monitor and configure RAID groups on the pair of Dell MD3360f disk arrays. Unlike the smaller MD1220 array in the DMZ, these storage systems do not require an external server for storage maintenance. A small Linux virtual machine will be used to host the Dell OpenManage software within the infrastructure hypervisors. The detailed configuration will happen within the DataCore SANsymphony-V product from the Windows Server DataCore nodes. This product also runs in a clustered node configuration using both nodes for balanced throughput and is expandable up to 24 nodes for additional balanced throughput and multi data center remote copy capability.

Due to the lack of accessibility to the website and not truly potential business for the company, I was unable to receive an actual quote. However, based on previous experience with the product, the following table depicts the options for the SANsymphony-V product with an estimated total cost:

DataCore	SANsymphony-V
Node Type	(2) Node Cluster License
Sizing	Up to 250 TB
Licensing Level	Enterprise, all features unlocked
Total Cost	**$200,000.00 (Estimated)**

Table 22 - DataCore SANsymphony-V Software Licensing

The backup software selected is IBM Tivoli Storage Manager. This software comes in many flavors but the version selected for the medium sized business is the LAN-free version of the software. This lowers CPU utilization on the backup server, decreases Ethernet network traffic, and provides greater bandwidth over fiber channel for a faster backup implementation. The difficult part about purchasing IBM products is to go through their intricate licensing model. IBM uses a Processor Value Unit (PVU) mechanism to price their software licensing. You have to use a chart and decipher the processor type and align the PVU count with the type and number of processors. Since we are not backing up any client systems, we only need to license the server for its PVU count and it is allowed to backup unlimited data.

IBM	Tivoli Storage Manager 6.x for Storage Area Networks
Node Type	Volume License for 250, 10 PVUs
Sizing	Unlimited
Warranty	1 Year Software Maintenance
Total Cost	**$36,323.00**

Table 23 - IBM Tivoli Storage Manager

8.2.4 Services BOM

Due to the size and nature of this particular medium sized business example, no services are required.

8.2.5 Total Cost of Solution

Here is the total cost of hardware, software, and services for a complete storage system solution for small businesses:

Line Items	Cost
Dell PowerEdge R320	$14,595.41
Dell PowerVault MD1220	$39,110.65

Dell R715	$68,958.16
Dell R720	$46,854.26
Dell MD3660f (slow)	$59,758.00
Dell MD3660f (fast)	$133,994.95
Dell R420	$11,940.11
Black Box Rack	$7,002.00
Console KVM	$1,539.99
Juniper Networks SRX650	$94,612.00
Juniper Networks EX3300	$84,371.52
Brocade 6510	$81,855.98
IBM TS3200	$35,360
DataCore SANsymphony-V	$200,000.00
IBM Tivoli Storage Manager	$36,323.00
Total Cost	**$916,276.03**

Table 24 - Total Cost of Procurement

8.2.6 Architecture and Interconnect Drawings

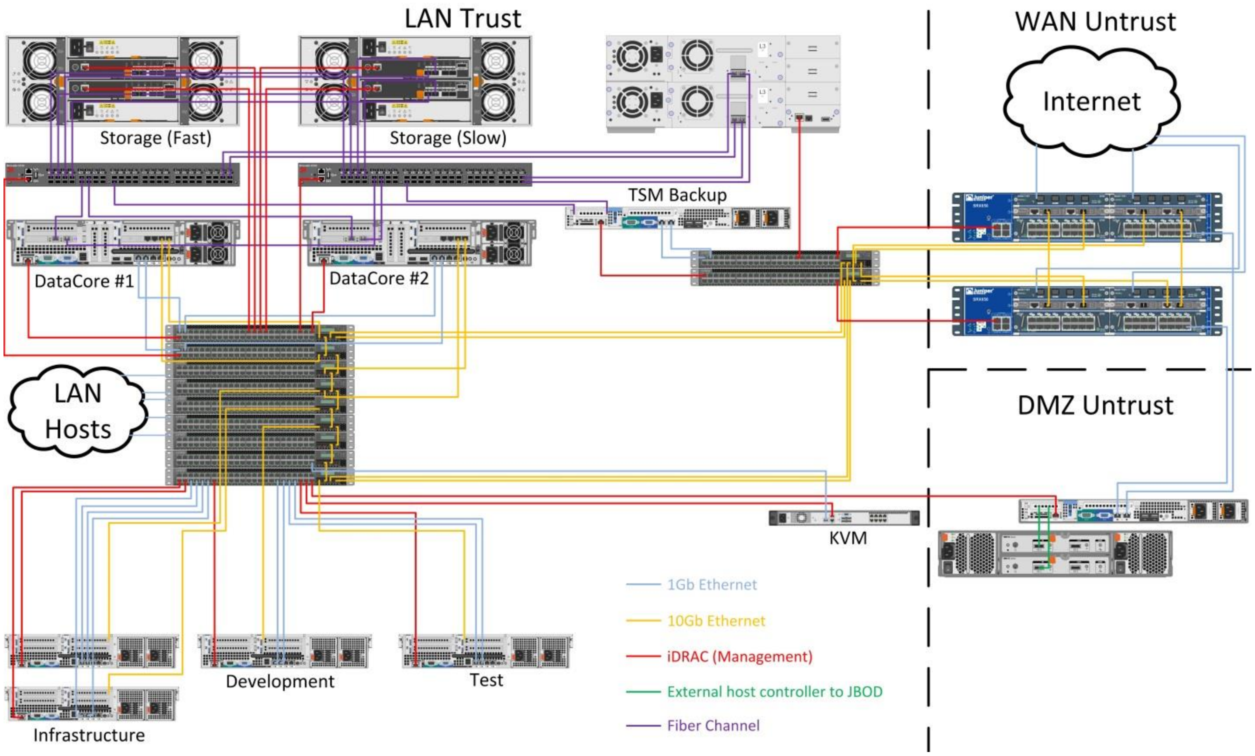

Figure 27 - Interconnect Diagram

Juniper networking equipment has the ability to configure logical networks or logical systems known as LSYS. Logical systems allow a single physical switch to be partitioned into many separate networks providing security boundaries and smaller broadcast domains. This technology will be used to divide the management network from the operational network to cut down on traffic to avoid performance degradation. To achieve the interconnect above, the EX3300 switches are trunked together to make one large logical switch via 10 Gigabit fiber trunk ports then subdivided into management and operational LSYS domains.

Another important architectural aspect to note is the DMZ. As we discussed before, this network is not trusted and therefore hangs off the Juniper firewalls. As you can see, the management interface on the R320 server in the DMZ connects back into the trusted environment. By not allowing routing between the maintenance and DMZ networks, this still supports a security boundary that meets best practices. This allows the server in the DMZ to be monitored and managed by the virtual machines on the infrastructure hypervisors.

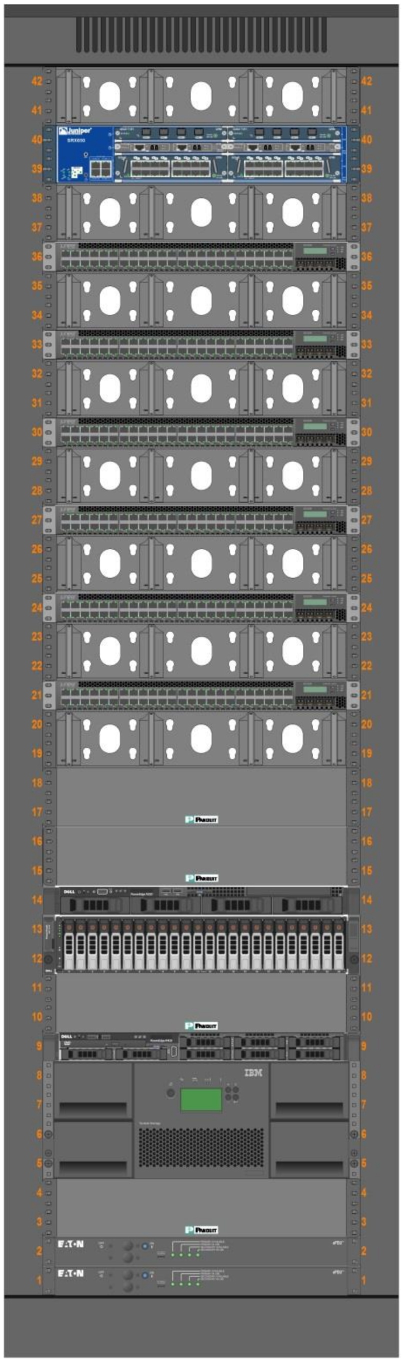

Rack 1

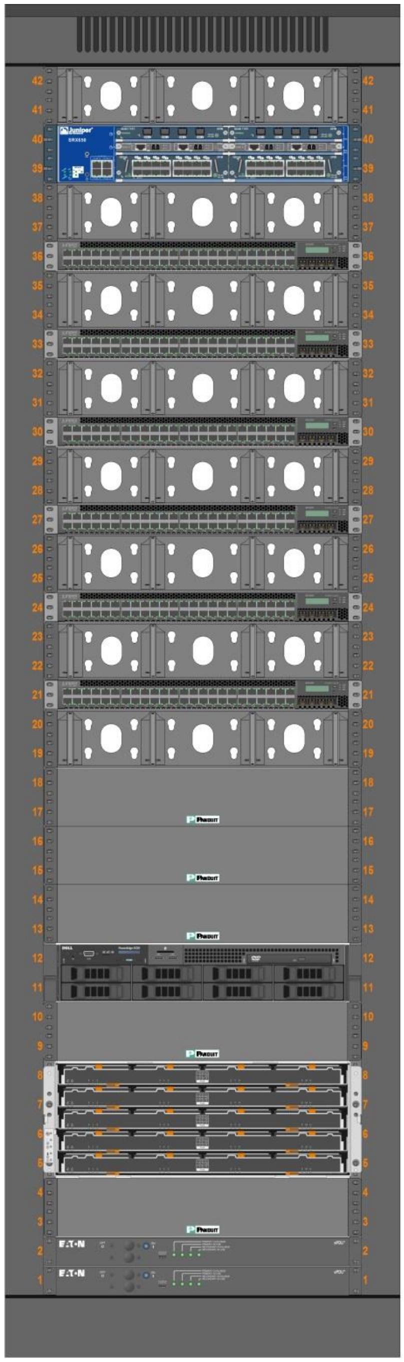

Rack 2

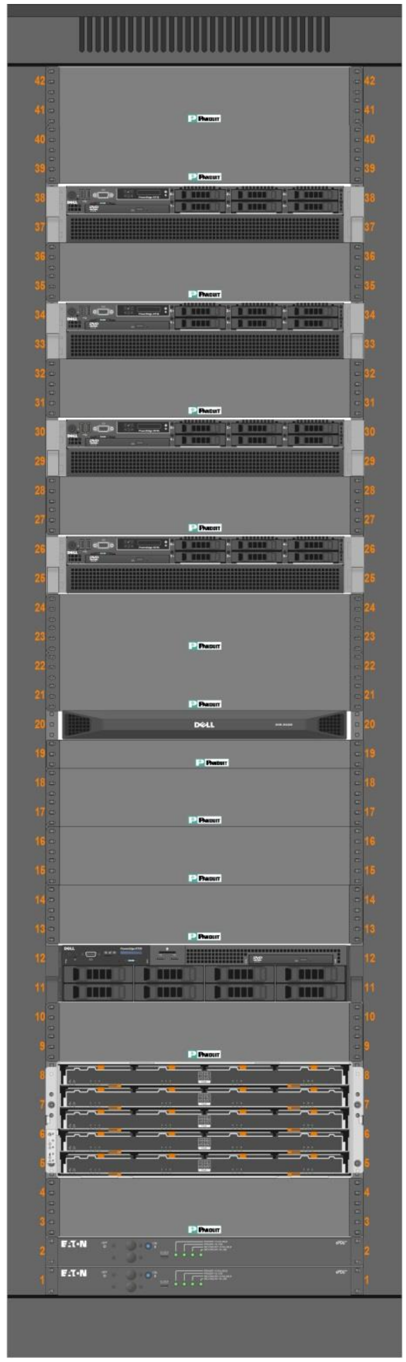

Rack 3

Figure 28 - Rack Layout – Front

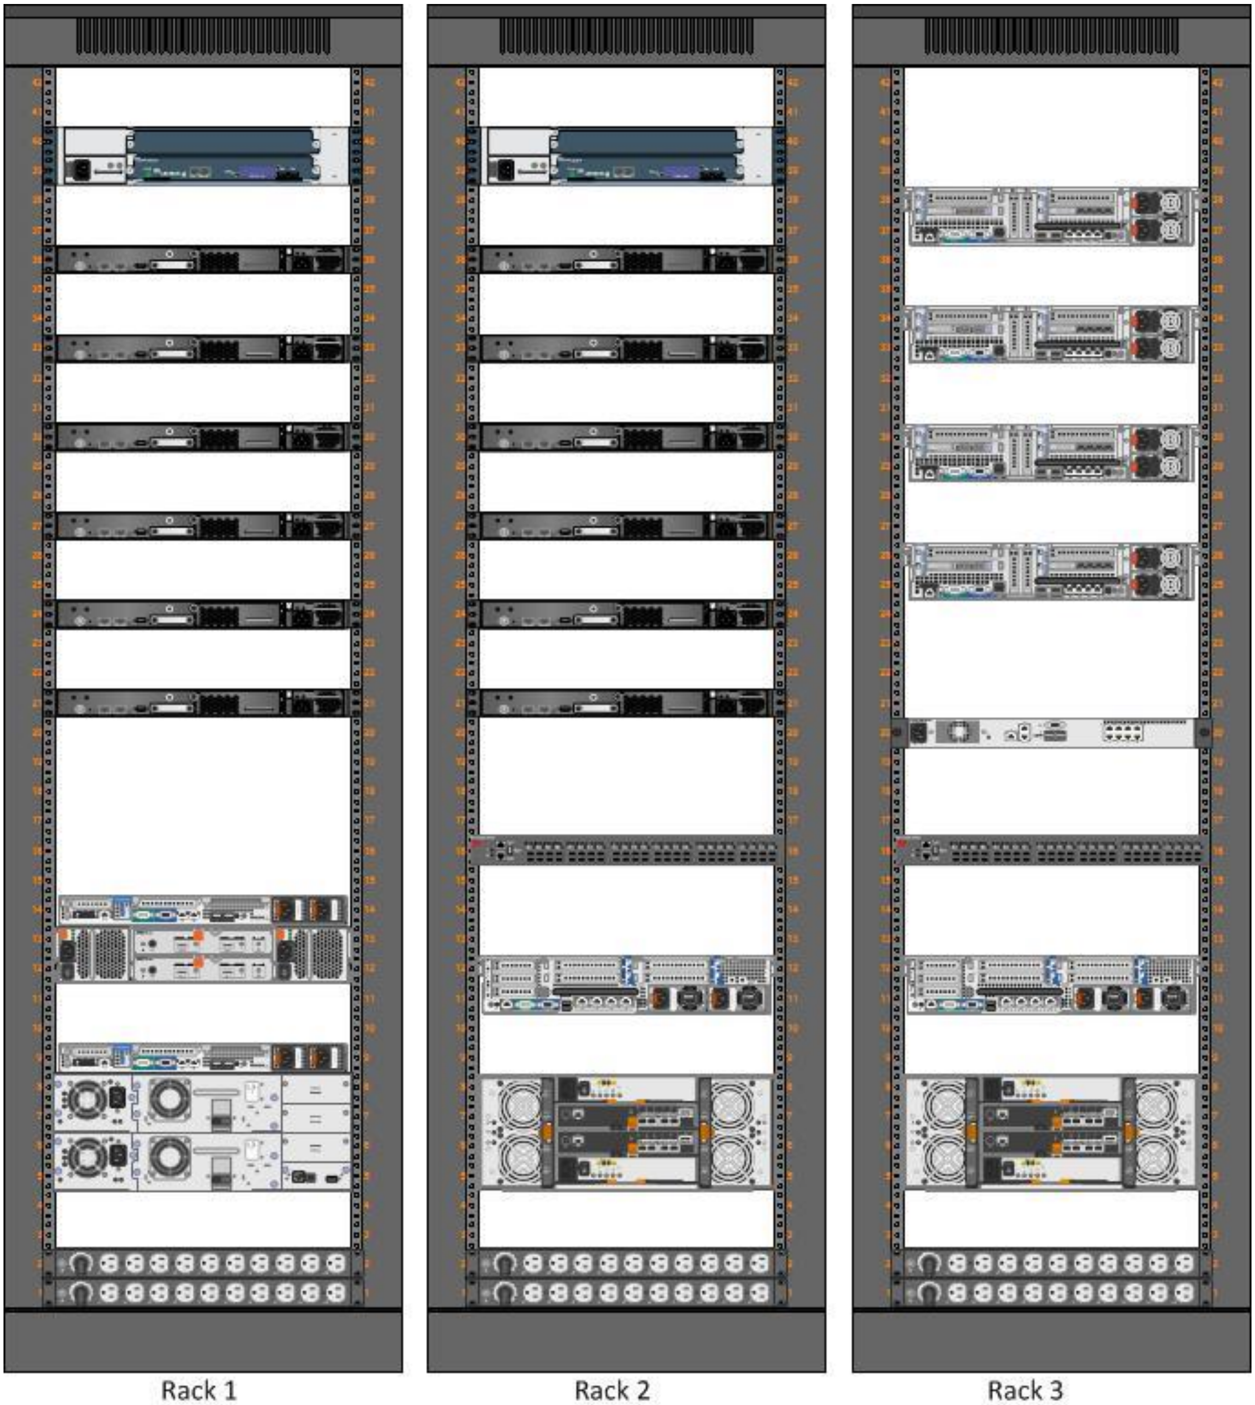

Figure 29 - Rack Layout - Back

8.2.7 Storage Virtualization Design

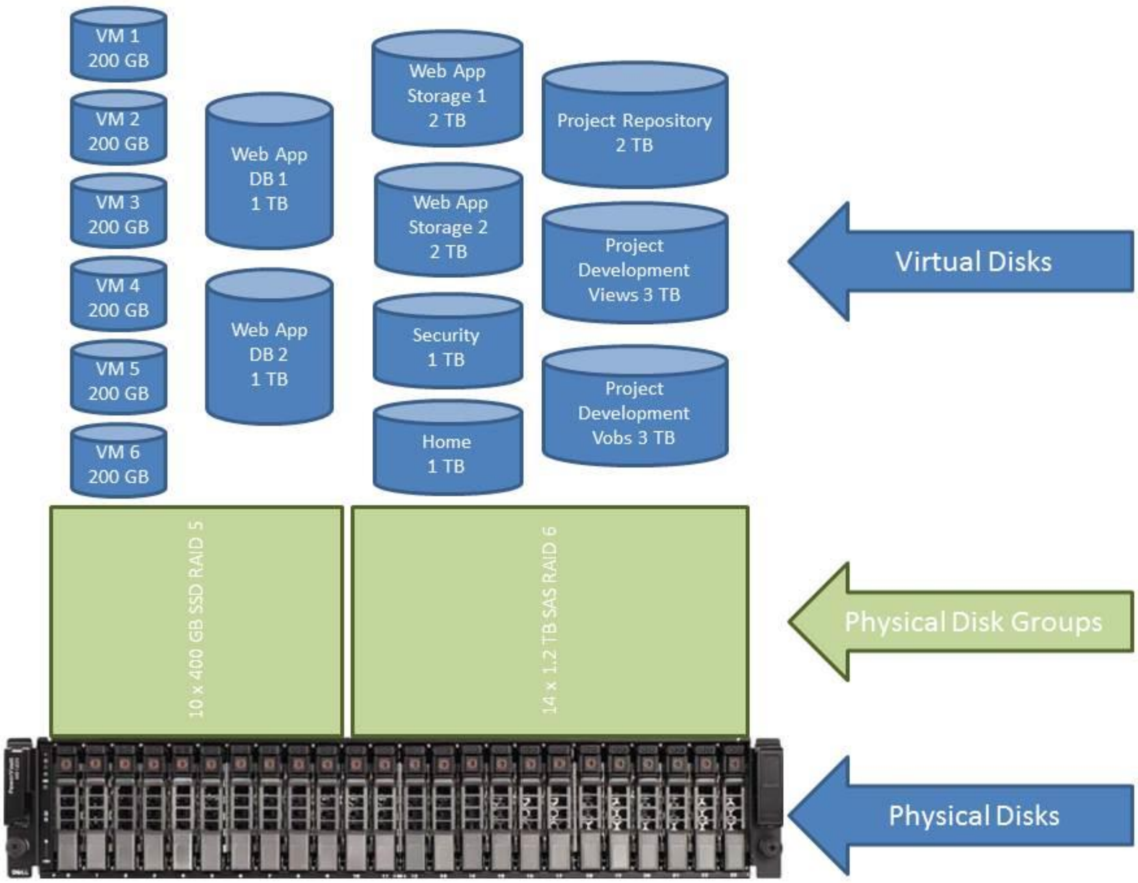

Figure 30 - DMZ - Physical and Virtual Storage Allocations (SAN Storage Works)

The storage in the DMZ will be carved up into many individual virtual disks to support the requirements. Since this device is the same used in the small business sample storage architecture, the configurations will be very similar. The storage virtualization technologies discussed at length in Section 3 such as thin provisioning with over-allocation, compression, disk expansion and shrinking on the fly, and perpendicular storage recording are all incorporated within the Dell MD1220 storage array. Perpendicular storage recording is inherent on the hard drives themselves. Disk expansions and shrinking on the fly would be used if the storage requirements ever changes in the future and the storage administrator needed to make a change to the architecture. Compression and thin provisioning with over-allocation are not currently needed in this example due to the sizing requirements of the DMZ. However, in the future, these are great techniques to provide more storage while preserving the same physical hardware longer into the future. These technologies are easily configured using the Storage Manager component in the Dell OpenManager software bundled with the hardware.

Figure 30 above depicts the Dell MD1220 physical disks, physical disk groups, and virtual disks. This view shows how the virtual disks are created on top of the physical disk groups, which are created on top of the physical disks. The disk array is carved up both physically and virtually to provide the storage needed detailed in the requirements Section 8.2.1.

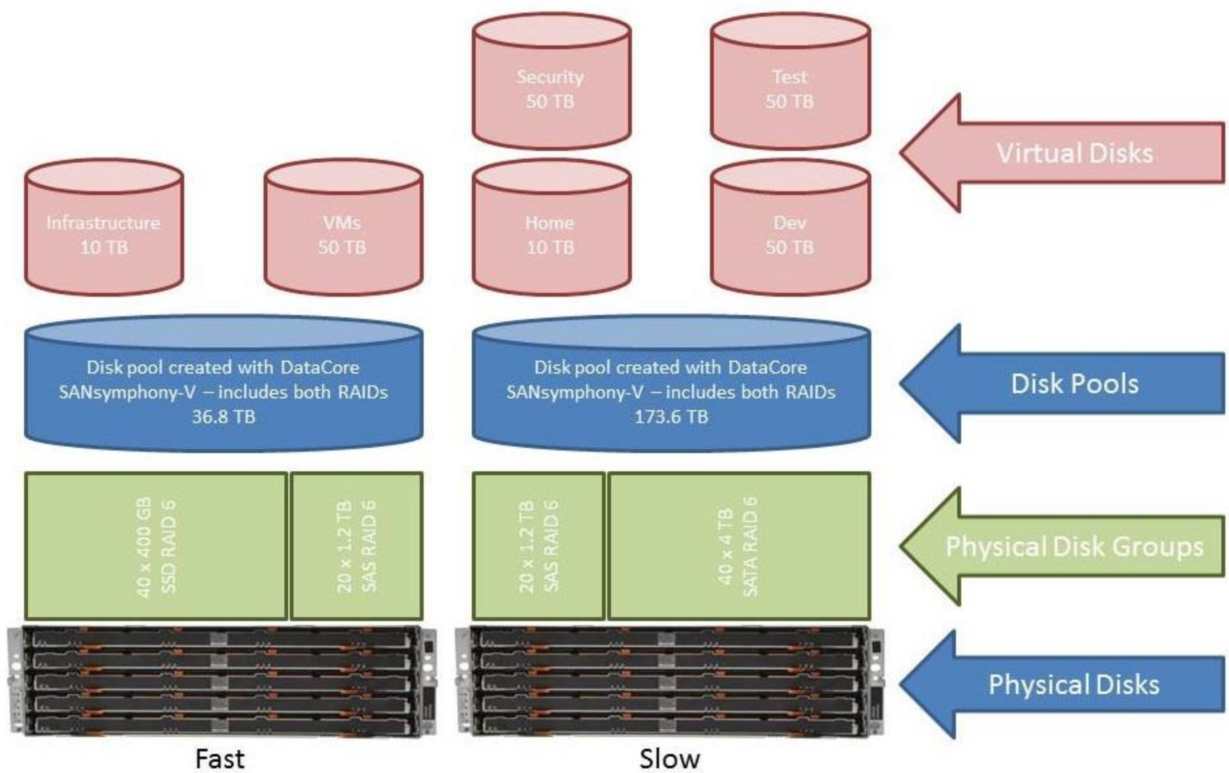

Figure 31 - Internal - Physical and Virtual Storage Allocations (EBI-Tech)

The internal multi storage array is a heterogeneous set of disks tied together with DataCore's storage virtualization solution to provide virtual disks within assorted medium options in the disk pools. In the Dell MD3660f storage array, each shelf contains 12 disks in a 4 by 3 configuration. Each storage array was configured to fill each disk slot with a drive and leave no spot empty resulting in a total of 120 disks of internal storage for a total of 210.4 TB of raw storage capacity.

The low level storage virtualization technologies comprised into this design include perpendicular storage recording on the spinning disks and RAID configurations within the storage arrays for fault tolerance. RAID 6 was used because the performance is similar to RAID 5 but provides and additional parity disk for added redundancy.

At the software layer, storage virtualization technologies are plentiful due to DataCore's SANsymphony-V COTS product. Auto tiering of storage is accomplished by nature of the software and how it is used to create disk pools. Above the physical disk group layer, which is done on the storage device itself, the software product defines a single or multiple disk pools within in a specific disk type or across heterogeneous disk groups. In our medium sized business sample architecture, a single disk pool across fast SSD and medium speed 10,000 RPM SAS drives was created. This disk pool within DataCore is allowed to create thinly provisioned virtual disks from the disk pool. Virtual disks can be granted a priority level which dictates which type of disks to migrate the virtual disk to, but in general, if all priorities are the same, DataCore has the ability to span its virtual disks at the block level across heterogeneous disk groups. This concept incorporates thin provisioning with over-allocation and automated tiered storage.

Automated tiered storage is a great feature to have, especially implementable at the block level. This helps to mitigate files that are read and written to infrequently from taking up crucial and expensive solid state disk. For example, if a virtual disk was created for an operating system on the fast storage array, blocks that are not frequently accessed may reside on the 10,000 RPM SAS spinning disk with no degradation to performance. However, files accessed frequently, are moved to the solid state disks even if not part of a virtual machine. This technology in this price point is hard to come by and provides the company with the opportunity to save money by not buying solid state disks to account for every gigabyte of storage underlying a virtual machines operating system, especially with thinly provisioning those hosts as well. Storage can go much further for a lot less when there is a software storage virtualization in between the compute nodes and backing storage subsystem.

As part the full storage virtualization solution, Windows Server 2012 contains data deduplication for its local and shared file systems. Since DataCore only runs on Microsoft Windows, the DataCore nodes end up providing the iSCSI and/or NFS shared file systems to the rest of the environment for use. Because the target shares are from a Windows based operating system, when the virtual disk is configured by DataCore, a simple rescan of the storage systems in Windows storage manager will display the new share. This share then needs to be initialized, partitioned, and then formatted before use. As part of that process, it would be extremely beneficial to install and enable the data deduplication feature within Windows Server 2012. This provides an additional layer of storage virtualization technology to minimize the disk usage on the backing storage subsystem.

The IBM TS3200 tape library uses compression to increase the storage capacity on a single LTO 6 tape. Uncompressed, an LTO 6 tape can hold 2.5 TB. Under optimal compression, a single LTO 6 tape can hold 6.25 TB of data (StorageNewletter.com, 2012). Also, TSM backup software implementation uses data deduplication when conducting backups to minimize tape usage and shortens backup durations.

8.2.8 Expandability & Future Proofing

When configuring the DMZ storage array, there was an option to add an additional hardware RAID controller. This option was selected with a PERC H800 RAID Kit for External JBOD with 1 GB of NV Cache. What this means is that there is an additional hardware controller that can connect to an external set of disks via the cable to interface between a future JBOD expansion. This option was relatively cheap but allows for an additional 192 disk devices to be configured in an attached array.

The SAN was configured with two 24-port fiber channel Brocade switches. Less than half these ports were used in the design providing room for expansion for additional storage or another DataCore server if throughput on the storage virtualization server becomes a bottleneck.

Table 11 shows the hardware BOM table for the Dell R715. These four servers are the VMware hosts for Infrastructure, Development, and Test environments. VMware is configured for booting the operating systems from the Dell MD3660f with the SSD's in a shared NFS file systems. The alternative to a simple NFS share with many VMDKs is to use the DataCore SANsymphony-V COTS product to show the operating system partitions straight through to the hypervisors as raw LUNs. The raw partitions are configurable per virtual machine on each of the hypervisors shown through as raw storage. When the hypervisor boots the virtual machine, the disk is unformatted space.

Non-disruptive data migration and mirroring is also possible with the storage virtualization solution chosen. DataCore has the ability to send its node into maintenance mode and allows a redundant copy to run from a separate data center. This storage can also be migrated between data centers and onto different storage arrays with no downtime. This technology is very implementable in this solution; however, no storage subsystems were identified as mirrored copies in differing data centers. This product can handle a mirrored copy on a storage array from a totally different manufacturer with different disk types and still function properly thanks to its clever ambiguous storage pools. The technology exists for future proofing the solution.

The Juniper EX3300 series switch has a stackable feature of up to 256 switches for expansion needs. This single switch is already chained with twelve total switches equating to a total of 576 1 GbE ports. This provides significant scalability for future proofing the network architecture if needed to go beyond twelve switches.

The IBM TS3200 can have their drives upgrade when the new technology is available for additional storage capacity and faster transfer mediums beyond LTO 6. Also available within the backup architecture is to add additional TSM servers to load balance the backups. Due to the CPU and memory intensive churning through data to determine the hash of every file to determine whether or not it needs to be backed up could delay backups into the following business day. At this point, it is advised to add an additional backup server to offset the load.

8.2.9 Complexity

The initial setup and configuration of the storage system solution would require a team of experts. The team of experts would need knowledge in SANs, medium sized networks with layer 3 switching, VMware, DataCore SANsymphony-V, Windows, RedHat, and tape libraries. Once the configuration is up, running, and configured properly, the solution would require minimum intervention. The single largest intervention into this system is to remember to offsite the LTO 6 tapes weekly to meet requirements. After that, the system only requires constant monitoring via the infrastructure stack of virtual machines for faulty VMs, storage allocations running low, and hardware/networking problems.

Medium sized business would be difficult to staff a team of engineers and administrators to man this size of a system. It would almost be beneficial to have an outside company conduct the fit-up and train a few individuals in each of the hardware and software implementations. This would save the cost of staffing an entire team for the integration phase then have too many administrators without enough to do on a day by day basis.

8.3 Large Business Storage Virtualization Application ~$10,000,000

As compared to the small and medium business examples in Section 8.1 and 8.2, a large sized business has larger revenue, more employees, and ultimately more complex IT needs. Unlike the small and medium business examples, a large business will likely necessitate a multisite data center requiring site replication and a high bandwidth low latency connection between sites. This sample business will employ greater levels of storage virtualization due to the increased storage requirements and additional budget. A sample Bill of Materials (BOM) for COTS software and hardware will be used to provide a sample IT system.

8.3.1 Requirements

Requirements define the needs of the business. Policies, procedures, and SLAs are put in place around the requirements to ensure these needs are met.

- Three multistory office buildings with 5,000 employees in total. The headquarters facility resides in Denver, Colorado which contains the data center attached to the multistory office building. In Sunnyvale, California and New York, New York a few racks of equipment are in the satellite office to support local operations, connectivity, and remote backup capabilities. The business supports a mix of workstations and laptops of both UNIX and Windows operating environments.
- Nightly backups of the entire data center with a weekly offsite to a secure storage location of the entire data subsystem. Full backup every Saturday with incremental backups throughout the week. This provides full recovery capability. Backups only occur of data on the storage array and not on any individual workstations or laptops. Offsite storage for backups is required.
- Application servers, development servers, and test servers have access to the storage subsystem. Each of the three environments maintains their own set of data stores in order to not disrupt business operations by means of segregated storage. For easy of manageability, the storage is collocated in a single data center with a separate storage subsystem for test, development, and

miscellaneous data stores. To keep the environments segregated, the shared data stores are logically separated from each other.

- UNIX based Lightweight Directory Access Protocol (LDAP) and Windows based Active Directory (AD) servers provide a single username and single password capability between the two differing environments. A Cisco ACS server, terminal servers, and shared KVM system in the lab environment also all take part in the single username authentication mechanism. Infrastructure services are divided between the primary and backup pair of each system instance.
- Home directories need to be 20 TB. Dev needs 150 TB of storage. Test needs 150 TB of storage. Infrastructure needs 20 TB of storage. Virtual machines supporting both the development and test systems take 70 TB of storage in total. Virtual machines are required to be on SSDs. Security requires its own share with root owned permissions of 20 TB. Infrastructure systems are physical machines as part of best practices.

8.3.2 Hardware BOM

The first line item on the Hardware BOM is the server that will connect to the storage array. This type of external connection to a storage array is known as a Direct Attached Storage (DAS). The storage is attached with 6 Gbps SAS cables and presents the storage to the software running on the host system. There are 3 Dell R720s configured in the table below; one per storage array type. In the next section, the configuration for the Dell MD3060e devices for SSD, SAS, and SATA will be defined. The servers customized configuration is listed below:

Dell PowerEdge R720	(3) PowerEdge R720, Up to 8 Hot-plug 2.5" Drives
Processor	(2) Intel Xeon E5-2690 3.0GHz, 25M Cache, 10 Core
Memory	(8) 8GB RDIMM, 1866MT/s
Internal Storage RAID Controller	PERC H710, 1GB Cache
External Storage RAID Controller	PERC H810 RAID for External JBOD, 1Gb NV Cache
Additional PCIe Card	(2) SAS HBA External Controller
Hard Drives	(4) 800GB Solid State MLC Drives, SAS, 6Gbps in RAID 10 (Mirror

	& Striped)
Network Adapter Onboard	Broadcom 5720 Quad Port 1Gb
Network Adapter Add-in	Intel X520 Dual Port 10Gb SFP+
Fiber Channel Adapter	(2) Qlogic 2662, Dual Port 16GB Fiber Channel HBA
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise with Vflash, 16GB SD Card
Warranty & Service	7 Year ProSupport Plus Warranty, Next Business Day Onsite Service
Total Cost	**$91,029.30**

Table 25 - Dell PowerEdge R720 Storage Controllers

There are a multitude of options when it comes to rack mount server options from Dell. The R700 series product line offers hardware RAID controllers, as opposed to software RAID controllers. This provides a greater level of configuration management and better performance since there is dedicated hardware to run the local hard drives.

The four 400GB solid state drives will be configured in RAID 10. This means that the hard drives will be configured to be striped for added performance followed by the striped pair to be in a mirrored pair. This is also known as RAID 1+0. These mirrored disks will run the RedHat Linux operating system for the R720 server. Single-Level Cell (SLC) flash memory is not commercial grade flash memory like Multi-Level Cell (MLC) flash memory. The SLC SSDs are meant for home use and are significantly cheaper but could degrade in performance and lifespan very quickly due to the I/O that servers produce. Therefore, the R720 was configured with MLC SSDs. The MLC drives, at the time of this configuration, were the same cost as the SLC. Also, the MLC drives selected for this configuration were larger than the SLC drives as faster (6Gbps versus 3Gbps).

The iDRAC7 Enterprise component for system management was selected because this is not a huge cost upper but provides the capability to power on or power off the server remotely and gives secure access to a terminal window. This is a powerful tool for remote administration and keeps the management network traffic away from the in-band interface needed for storage access. The different with this iDRAC versus the small and medium business option was that this iDRAC was configured with the Vflash with a 16GB SD card. This allows for KVM over IP and operates as emergency access over the out-of-band network interface.

The second line items in the Hardware BOM are the storage arrays. The Dell MD3060e storage JBODs can hold up to 60 drives each in the 4U form factor. Up to 4 of these JBODs can be daisy chained together. In this case, 3 trays were needed for SSD, 2 trays for SATA, and 1 tray for SAS to meet the storage requirements. These devices will service all the storage needed for the large business unit by an external connect to the R720 host systems described above. Three storage subsystems will be configured with one per storage type; SSD, SAS, and SATA.

Dell PowerVault MD3060e	(3) PowerVault MD3060e, Rackmount, 60 Bay, 4U Dense Expansion Enclosure, Dual Power Supplies
Hard Drives	(60) 400GB SSD SAS SLC 6Gbps 2.5in Hot-plug
Storage Connectivity	JBOD Connection to Dell R720 via 6Gbps HBA
Cable	(4) 6Gb Mini to Mini SAS Cable, Expansion Cable
Warranty & Service	5 Year ProSupport Plus Warranty, Next Business Day Onsite Service
Total Cost	**$500,309.07**

Table 26 - Dell PowerVault MD3060e Fast Configuration with Solid State Drives

This configuration of the Dell MD3060e Fast Array consists of three of the 4U modules which contain a total of 180 drives. Each drive is a 400GB SAS SLC SSD capable of 6Gbps. The RAID controller on the Dell R720 host server is capable of configuring the direct attached storage in RAID level 0, 1, 5, 6, 10, 50, and 60. The configuration of this array will include a single RAID 6 consisting of all 60 SSDs per 4U disk tray. This will provide a raw configuration of 23.2 TB (400 GB x 60 drives, minus 2 drives for parity) per 4U tray. This will provide a raw configuration of 69.6 TB for the three trays together. This raw configuration will be presented to the SANsymphony-V software which conducts the auto tiering of the storage discussed in Section 8.3.7.

Dell PowerVault MD3060e	PowerVault MD3060e, Rackmount, 60 Bay, 4U Dense Expansion Enclosure, Dual Power Supplies
Hard Drives	(60) 1.2TB 10,000 RPM SAS 6Gbps 2.5in Hot-plug
Storage Connectivity	JBOD Connection to Dell R720 via 6Gbps HBA
Cable	(4) 6Gb Mini to Mini SAS Cable, Expansion Cable
Warranty & Service	5 Year ProSupport Plus Warranty, Next Business Day Onsite Service
Total Cost	**$53,468.22**

Table 27 - Dell PowerVault MD3060e Medium Configuration with SAS Drives

This configuration of the Dell MD3060e Medium Array consists of a single 4U module which contains a total of 60 drives. Each drive is a 1.2TB 10,000 RPM SAS drive capable of 6Gbps. The RAID controller on the Dell R720 host server is capable of configuring the direct attached storage in RAID level 0, 1, 5, 6, 10, 50, and 60. The configuration of this array will include a single RAID 6 consisting of all 60 SAS disks. This will provide a raw configuration of 69.6 TB (1.2 TB x 60 drives, minus 2 drives for parity). This will provide a raw configuration of 69.6 TB for the three trays together. This raw configuration will be presented to the SANsymphony-V software which conducts the auto tiering of the storage discussed in Section 8.3.7.

Dell PowerVault MD3060e	(2) PowerVault MD3060e, Rackmount, 60 Bay, 4U Dense Expansion Enclosure, Dual Power Supplies
Hard Drives	(60) 4TB 7,200 RPM Near-Line SAS 6Gbps 3.5in Hot-plug
Storage Connectivity	JBOD Connection to Dell R720 via 6Gbps HBA
Cable	(4) 6Gb Mini to Mini SAS Cable, Expansion Cable
Warranty & Service	5 Year ProSupport Plus Warranty, Next Business Day Onsite Service
Total Cost	**$115,455.34**

Table 28 - Dell PowerVault MD3060e Slow Configuration with SATA Drives

This configuration of the Dell MD3060e Slow Array consists of a single 4U module which contains a total of 60 drives. Each drive is a 4TB 7,200 RPM Near-Line SAS drive capable of 6Gbps. The RAID controller on the Dell R720 host server is capable of configuring the direct attached storage in RAID level 0, 1, 5, 6, 10, 50, and 60. The configuration of this array will include a single RAID 6 consisting of all 60 Near-Line SAS drives per 4U disk tray. This will provide a raw configuration of 232 TB (4 TB x 60 drives, minus 2 drives for parity). This will provide a raw configuration of 464 TB for the two trays together. This raw configuration will be presented to the SANsymphony-V software which conducts the auto tiering of the storage discussed in Section 8.3.7.

The configuration of the storage to the host server, or controller, can be attached in many different configurations. The configuration for each storage subsystem is directly attached to the server's external RAID controller card with a SAS cable. The connection between the host controller (R720) and the storage array (MD3060e) is the same as depicted in Figure 22 in the Small Business Section 8.1.2. This also shows the daisy chain capability, with up to four storage trays for future expansion if needed.

The hypervisors for the large sized business are drastically different than those identified for the small and medium sized businesses. Primarily because the large business has more funding to spread out the services and virtualized hosts for the development, test, and infrastructure systems. Each Dell R820 server in the table below can support a significant number of virtual machines. Each system contains four 10 core Intel processors and 512 GB of memory. These 2U servers are optimized for a high number of CPU cycles in supporting multiple virtual machines and heavily storage traffic over the 10 Gbps Ethernet or 16 Gbps fiber channel interfaces. A total of 8 will be procured; 3 for development, 3 for test, and 2 for infrastructure services.

Dell PowerEdge R820	(8) PowerEdge R820, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(4) Intel Xeon E5-4640 v2, 2.2 GHz, 10 core, 20 MB cache
Memory	(32) 16 GB RDIMM, 1866 MT/s
Storage RAID Controller	PERC H710P, 1 GB Cache
Hard Drives	(2) 800 GB Solid State MLC Drives, SAS, 6 Gbps in RAID 1 (Mirror)
Network Adapter	Intel Ethernet I350 Quad Port 1Gb
Network Adapter Add-in	Broadcom 5719 Quad Port 1Gb

Network Adapter Add-in	(2) Intel X520 Dual Port 10Gb SFP+
Network Adapter Add-in	(2) Qlogic 2662, Dual Port 16Gb Fiber Channel HBA
Operating System	VMware ESX 5.x Embedded Image on Flash Media, vSphere Enterprise Plus 5.x, 4 CPU, 5 Year Subscrption
System Management	iDRAC7 Enterprise with Vflash 16GB SD Card
Warranty & Service	7 Year ProSupport, Next Business Day Onsite
Total Cost	**$491,462.08**

Table 29 - Dell PowerEdge R820 Development, Test, and Infrastructure Hypervisors

An additional feature on these Dell R820s above that were not previously configured for the small or medium sized businesses hypervisors is the VMware embedded flash media. This device is internal to the server and allows for an installation of VMware ESXi 5.x in this case being local to the server. This helps in that if this server was configured for hypervisors only running from centralized shared storage, no local hard drives would be required for this system. In the case of the large business implementation, a proper local hard drive storage controller with 2 SSDs was configured in case this server was ever repurposed for a physical host instead of a hypervisor.

The next set of servers, two to be exact, are going to be used as the DataCore SANsymphony-V nodes. We discussed DataCore and their SANsymphony-V product at length in Section 3.1. In essence, the servers act as your storage controllers and provide SMB, NFS, and direct attached capabilities for all LUNs configured on the storage array. The hardware consists of dual 8 core Intel processors, 384 GB of memory and a RAID 10 with four MLC SSDs to host the operating system. The reason for the 384 GB of memory is that the SANsymphony-V product uses local memory on the nodes as cache, just as a storage controller would. Files that are accessed frequently and read/writes to the controller do not have to go back to reference the spinning disk or SSDs until later. This is a quick read/write access capability with the legwork being done after the fact.

These two storage controller nodes will run Windows 2012 R2 Datacenter Edition. Datacenter Edition is required because of the amount of memory and dual CPUs.

Windows was selected because the DataCore product only runs on that platform at the moment but is to be ported to Linux for support in the near future. These nodes also contain two dual port 16 Gb fiber channel cards to connect to a Brocade fiber channel switch configured later in this section. The table below articulates the hardware components to support the SANsymphony-V product as a storage controller:

Dell PowerEdge R720	(2) PowerEdge R720, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(2) Intel Xeon E5-2667 v2 3.3 GHz, 25 MB cache, 8 core
Memory	(24) 16 GB RDIMM, 1866 MT/s
Storage RAID Controller	H710P 1GB NV Cache
Hard Drives	(4) 800 GB Solid State MLC Drives, SAS, 6 Gbps in RAID 10 (Mirror & Striped)
Network Adapter Add-in	Broadcom 5720 Quad Port 1 Gb
Network Adapter Add-in	Broadcom 5719 Quad Port 1 Gb
Network Adapter Add-in	Intel X520 Dual Port 10Gb SFP
Network Adapter Add-in	Qlogic 2662, Dual Port 16 Gb Fiber Channel
Operating System	Microsoft Windows Server 2012 R2 Datacenter Edition
System Management	iDRAC7 Enterprise with Vflash, 16GB SD Card
Warranty & Service	7 Year ProSupport, Next Business Day Onsite
Total Cost	**$69.029.46**

Table 30 - Dell PowerEdge R720 DataCore Storage Controllers

The backup server is a critical piece of the infrastructure. Below is the configuration of a simple single CPU 2U server. The single CPU cuts down on licensing costs for the COTS product conducting the backups called Tivoli Storage Manager (TSM) from IBM. The reason that 2 backup servers were purchased was in support of redundant backup systems. 1 will resides in the New York office and 1 in the Sunnyvale office. This will ensure that backups are redundant of the Denver data center as well as maintained off site from the primary storage subsystem.

Dell PowerEdge R720	(2) PowerEdge R720, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(2) Intel Xeon E5-2667 v2 3.3 GHz, 25 MB Cache, 8 Core

Memory	(4) 16 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 1 GB Cache
Hard Drives	(4) 800 GB Solid State MLC Drives, SAS, 6 Gbps in RAID 10
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in 1	Intel X520 Dual Port 10 Gb SFP+
Network Adapter Add-in 2	Qlogic 2662, Dual Port 16 Gb Fiber Channel HBA
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise with Vflash, 16 GB SD Card
Warranty & Service	7 Year Pro Hardware Warranty, On-site Service
Total Cost	**$58,109.64**

Table 31 - Dell PowerEdge R720 Backup Server

Each backup server is complimented with an IBM Tape Library TS3200 with an additional storage sled. This device can backup to a single tape at a time with a total of 40 tapes in the device. The TS3200 has fiber channel connections to provide LAN-free backup capabilities.

IBM	(2) Tape Library TS3200
Base Chassis	(2) LTO 6 Fiber Channel Drives
Tapes	(10) LTO 6 Data Cartridges 5-packs
Warranty & Support	3 Years
Connectivity	(2) 8 Gbps Fiber Channel Port
Total Cost	**$70,720.00**

Table 32 - IBM TS3200 Tape Library

As configured, this device uses Linear Tape-Open (LTO) 6 drives. The LTO 6 tape was released in 2012 and can hold up to 2.5 TB each uncompressed and up to 6.25 TB compressed. This means the tape library has a theoretical maximum capacity of 250 TB, exceeding the full storage requirements for this business unit.

The last server class hardware in the hardware BOM is the infrastructure systems. Best practices states that critical infrastructure systems such as your VMware vSphere server, Windows Active Directory servers, and RedHat LDAP servers should be physical installations so in a catastrophic failure, the critical infrastructure systems are easily recoverable or replaceable in their physical form instead of virtual machines. Below is a configuration for simple servers to handle local user authentication, DNS, and other infrastructure services localized to the facility they are supporting. 3 Windows Active Directory Domain Controllers provide DNS and user authentication for the Windows domain. These 3 servers, 1 per facility, are in a multi-master configuration. These Windows servers also host the VMware vCenter software for local and remote manageability of the virtual machines using the VMware infrastructure. The Windows servers also provide TACACS which is the network device authentication mechanism and it is tied into Windows AD natively. The RedHat servers are providing LDAP authentication as a multi-master configuration.

Dell PowerEdge R720	(3) PowerEdge R720, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(2) Intel Xeon E5-2667 v2 3.3 GHz, 25 MB Cache, 8 Core
Memory	(4) 16 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 1 GB Cache
Hard Drives	(4) 800 GB Solid State MLC Drives, SAS, 6 Gbps in RAID 10
Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in 1	Intel X520 Dual Port 10 Gb SFP+
Network Adapter Add-in 2	Qlogic 2662, Dual Port 16 Gb Fiber Channel HBA
Operating System	RedHat Linux 6.X, 5 Year Premium License
System Management	iDRAC7 Enterprise with Vflash, 16 GB SD Card
Warranty & Service	7 Year Pro Hardware Warranty, On-site Service
Total Cost	**$87,164.46**

Table 33 - Dell PowerEdge R720 RedHat Infrastructure Support Servers

Dell PowerEdge R720	(3) PowerEdge R720, Up to 8 Hot-plug 2.5 Inch Drives
Processor	(2) Intel Xeon E5-2667 v2 3.3 GHz, 25 MB Cache, 8 Core
Memory	(4) 16 GB RDIMM, 1600 MT/s
Internal Storage RAID Controller	PERC H710, 1 GB Cache
Hard Drives	(4) 800 GB Solid State MLC Drives, SAS, 6 Gbps in RAID 10

Network Adapter Onboard	Broadcom 5270 Dual Port 1 Gb
Network Adapter Add-in 1	Intel X520 Dual Port 10 Gb SFP+
Network Adapter Add-in 2	Qlogic 2662, Dual Port 16 Gb Fiber Channel HBA
Operating System	Windows Server 2012 R2 Data Center Edition, Unlimited VMs
System Management	iDRAC7 Enterprise with Vflash, 16 GB SD Card
Warranty & Service	7 Year Pro Hardware Warranty, On-site Service
Total Cost	**$87,638.76**

Table 34 - Dell PowerEdge R720 Windows Infrastructure Support Servers

In general, all the servers for the large business are significantly more expensive than the small and medium sized business examples. This is due to how these machines contain additional memory, additional cores in the CPUs, upgraded to SSD of the performance and enterprise level, added maintenance for the operating system, and prolonged warranty and compared to small and medium businesses.

Black Box Rack	(4) Elite Data Cabinet, 84"H (45U) x 30"W x 36"D
Body, Doors, Rails	12-gauge steel construction
Rack Units	45U
Lock	Key
Doors	Mesh
Sides	Solid
Fans	2 Fan Top Panel Module
Capacity	2000 pounds
Total Cost	**$9,336.00**

Table 35 - Rack Assembly for Denver Data Center

Black Box Rack	(2) ClimateCab NEMA Server Cabinet with AC Unit, 42U
Body, Doors, Rails	12-gauge steel construction
Rack Units	42U
Capacity	2000 pounds
Total Cost	**$11.797.90**

Table 36 - Rack Assembly with AC Unit for New York & Sunnyvale

For the remote facilities in New York and Sunnyvale, a single air conditioned rack will be used to hold and cool the equipment locally. These racks are defined in Table 36 above. In Table 35 above are the generic Black Box racks that will hold the equipment in the data center located in Denver in an air conditioned room.

Additional hardware is required to complete the implementation of the small business storage solution. In the following table, equipment required to complete the design is addressed:

The rack the equipment will also require a Keyboard, Video, and Mouse (KVM) setup to manage the systems locally. One per site will be procured for the satellite offices with one per rack, for a total of 4, for the Denver data center. The solution selected can connect up to 16 devices and provides local authentication with username and password for security:

Console KVM	(6) Tripp Lite B040-016-19
Ports	16 KVM systems
Monitor	19" 1366x768 resolution
Total Cost	**$9,239.94**

Table 37 - KVM Console

The networking equipment is needed to facilitate communications within the private enterprise and out to the Internet. No routers are needed as the firewalls act as a gateway to the Wide Area Network (WAN). On the trusted side of the firewalls is the Local Area Network (LAN). Below is a simple topology using Juniper equipment to provide connectivity. The full detailed network topology will be shown in Section 8.3.6.

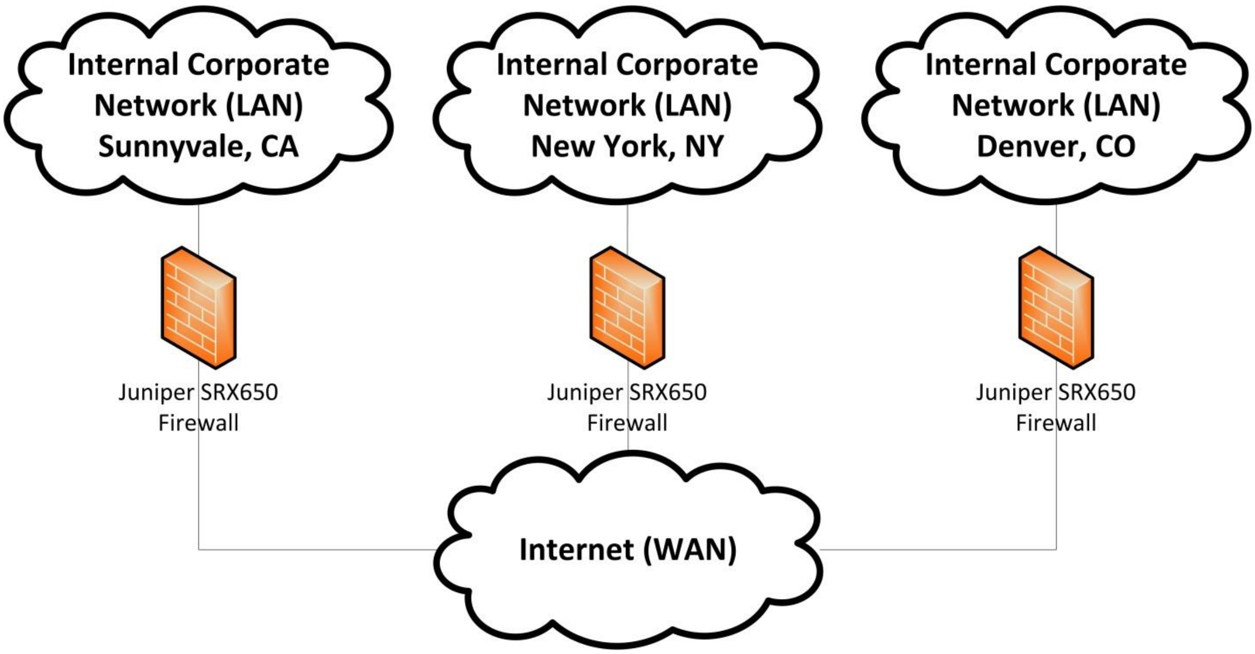

Figure 32 - Large Business High Level Network Topology

The Juniper SRX series firewalls provide a site-to-site VPN connection capability. This means each of the 3 firewalls know about each other and create an IPsec tunnel in which traffic flows from facility to facility as if it were all on the same LAN. With route-based VPNs, a policy does not specifically reference a VPN tunnel. Instead, the policy references a destination address. When the security device does a route lookup to find the interface through which it must send traffic to reach that address, it finds a route via a secure tunnel (ST) interface, which is bound to a specific VPN tunnel. The firewall devices are configured in the table below:

Juniper Networks Firewall	Juniper Networks SRX650
Base Chassis	(3) SRX650 with 8 slots
Connectivity 1	(2) Quad T1/E1, 4-port with CSU/DSU (SRX-GP-QUAD-T1-E1)
Connectivity 2	(6) 2-port 10 GbE SFP+ / 10 G Base-T Copper
Support	Juniper Care Next Day Support
Renewable Protection	5-year subscription for AppSecure and IPS Updates
Total Cost	**$148,724.00**

Table 38 - Juniper SRX650 Firewalls

In order to provide a high speed backbone for the Denver data center, pair of Juniper QFX5100 switches is implemented between the access layer switches and the Juniper firewall SRX650 edge node. This device has 96 ports with 10 Gbps SFP+ connectivity, however is capable of 1 or 10 Gbps and will be used by all the downstream switches to uplink to. Only a pair is needed for the Denver data center because the remote facilities do not have high bandwidth uplinks like the data center facility.

Juniper Networks Switch	Juniper Networks QFX5100-96S
Base Chassis	(2) QFX5100 96 Port Managed Switch
Connectivity 1	8-port 40 Gb QSFP+
Connectivity 2	Ethernet Switch 96-port 1 Gb SFP / 10 Gb SFP+
QSFP	(8) 40Gbps QSFP+ Transceiver
SFP+	(96) 10Gbps SFP+ Transceiver
Support	Juniper Care Next Day Support
Total Cost	**$95,396.00**

Table 39 - Juniper QFX5100 Switches

At the Denver facility, the switch chosen as the access layer switch for client workstation and laptop connectivity are the huge in size and the hugely capable Cisco 6513 layer 3 switches. Below is the configuration table showing a standard chassis with 2 supervisor units that contain redundant 10 Gb uplink ports and eleven 48 port gigabit Ethernet blades. Note that not all switches are displayed in the network architecture or rack diagrams in Section 8.3.6 as these devices will reside at the headquarters facility in Denver would be contained in a communications closet. The switch fabric capacity for each device is 720 Gbps exceeding the business needs of an access layer device. There exists the capability to stack these switches which will be discussed in the future proofing Section 8.3.8.

Cisco	(12) Cisco 6513 Layer 3 Switch
Supervisor Unit	(2) Supervisor Engine 2T 2 x 10GbE SFP+
Connectivity	(11) 48 x 10/100/1000 Gbps

Total Cost	**$1,121,158.56**

Table 40 – Cisco 6513 Layer 3 Switches for Client Connectivity

Each of the facilities will receive 3 of the Juniper EX3300 switches which will be used for both server and client connectivity for the small offices and just server connectivity for the data center in Denver. These switches are needed to provide ample copper Ethernet connectivity as well as 4 uplink ports that can be used for 1 Gbps or 10 Gbps fiber optics for uplink connectivity to the Juniper SRX650 firewalls or core switches.

Juniper Networks Switch	Juniper Networks EX3300
Base Chassis	(9) EX3300 48-port 10/100/1000 Base-T with 4 SFP+ Uplink Ports
License Type	Enhanced Feature License (Supports Layer 3 Protocols)
Support	Juniper Care Next Day Support
Total Cost	**$63,278.64**

Table 41 - Juniper EX3300 Layer 3 Switches

Due to the requirement to have the backups reside away from the data center, the Brocade 7800 devices are being installed to extend the fiber channel fabric over the WAN. Tivoli Storage Manager (TSM) is known for its LAN free backup capabilities. This means that the RedHat backup servers in New York and Sunnyvale will backup over the SAN fabric. This does require an extension of the fabric over the WAN to the east and west coast offices. Leveraging the next-generation Fiber Channel and advanced FCIP (Fiber Channel of Internet Protocol) technology, the Brocade 7800 creates a ring topology with all 3 sites to all LAN free backups to occur over the WAN to remote facilities shown below:

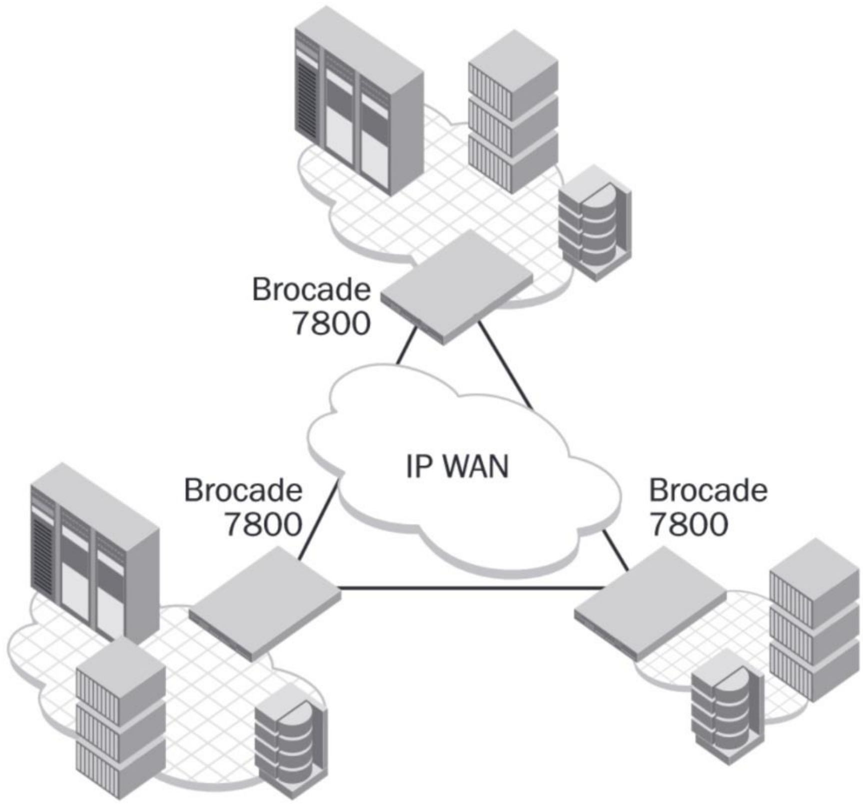

Figure 33 - FCIP Point to Point Architecture (Brocade, 2013)

Brocade FCIP Switch	(3) Brocade 7800 SAN Extender
Connectivity 1	16 x 8Gbps Fiber Channel
Connectivity 2	6 x 1Gbps Ethernet SFP
Management Port	1 x 1Gbps (RJ45)
Total Cost	**$37,656.99**

Table 42 – Brocade 6520 FCIP Switches for WAN Connectivity

A pair of fiber channel switches to support fabric A and fabric B for redundancy will be required to connect the DataCore nodes, storage subsystem, and uplink to the Brocade 7800s for FC over the WAN. In the table below is a configuration of Brocade 6520 fiber channel switches:

Brocade Fiber Channel Switch	(2) Brocade 6520 Fiber Channel Switch
Connectivity	96 x 16Gbps Fiber Channel (48 ports configured

	with expansion of up to 96 total ports)
Management Port	1 x 1Gbps (RJ45)
Total Cost	**$137,859.98**

Table 43 – Brocade 6520 FC Switches

8.3.3 Software BOM

The RedHat Linux server 6.X license for the R720s purchased for the storage controllers, backup servers, and infrastructure support servers were all configured with a 5 year premium license. This allows for extended vendor support as well as media and patch updates as the operating system continues to mature. RedHat allows the procurement tiers of different virtual guests if you plan to host virtual machines on your RedHat server with a functionality called RHEV (RedHat Enterprise Virtualization). This feature will not be used in this environment because the R720s will not be running virtual machines. A bare metal installation of RedHat 6 will run on the RAID 10 SSDs. RedHat is an open source operating system, however the purchase will provide fixes and support if necessary by the local administrator of this sample small business. Since RedHat is open source, the development, test, and infrastructure virtual machines will employ the free RedHat operating systems without the support contract.

The Microsoft Windows Server 2012 R2 Datacenter Edition was configured as part of the DataCore and Infrastructure servers. As noted before, Datacenter Edition was required due to the amount of RAM configured on the systems and dual and quad CPU platform configurations. Additional Windows licensing will be purchased to support the virtual machines in the development, test, and infrastructure environments.

Microsoft	Windows 2012 R2 Data Center Edition
Node Type	(20) Node Locked License
Total Cost	**$100,000.00**

Table 44 – Microsoft Windows Software Licensing

Dell OpenManage comes with the MD3060e disk arrays and will be used to configure and monitor the disk array. It will be installed into the bare metal RedHat Linux instance on the R720 and accessed locally of through a web browser. It is supported by Internet Explorer (IE), Mozilla Firefox, and Google Chrome. The detailed configuration will happen within the DataCore SANsymphony-V product from the Windows Server DataCore nodes. This product also runs in a clustered node configuration using both nodes for balanced throughput and is expandable up to 24 nodes for additional balanced throughput and multi data center remote copy capability.

Due to the lack of accessibility to the website and not truly potential business for the company, I was unable to receive an actual quote. However, based on previous experience with the product, the following table depicts the options for the SANsymphony-V product with an estimated total cost:

DataCore	SANsymphony-V
Node Type	(2) Node Cluster License
Sizing	Up to 500 TB
Licensing Level	Enterprise, all features unlocked
Total Cost	**$400,000.00 (Estimated)**

Table 45 - DataCore SANsymphony-V Software Licensing

The backup software selected is IBM Tivoli Storage Manager. This software comes in many flavors but the version selected for the medium sized business is the LAN-free version of the software. This lowers CPU utilization on the backup server, decreases Ethernet network traffic, and provides greater bandwidth over fiber channel for a faster backup implementation. The difficult part about purchasing IBM products is to go through their intricate licensing model. IBM uses a Processor Value Unit (PVU) mechanism to price their software licensing. You have to use a chart and decipher the processor type and align the PVU count with the type and number of processors. Since we are not backing up any client systems, we only need to license the server for its PVU count and it is allowed to backup unlimited data. Enough PVUs to cover 2 physical systems is required due to the 2 backup systems in the architecture; 1 east coast office and 1 west coast office for offsite redundancy.

IBM	Tivoli Storage Manager 6.x for Storage Area Networks
Node Type	(2) Volume License for 250, 10 PVUs
Sizing	Unlimited
Warranty	1 Year Software Maintenance
Total Cost	**$72,646.00**

Table 46 - IBM Tivoli Storage Manager

Because of the number of VMware ESXi hypervisor instances, it is advised to employ a VMware vCenter Operations Management Suite Enterprise Edition to manage all the virtual machine hypervisors. Since this product runs on Windows, it makes most sense to install 3 copies; one per instances of the Windows Infrastructure servers. This negates the potential circular dependency of having to manage a virtual machine instances with vCenter which resides in a virtual machine. This would be against best practices.

VMware	vCenter Operations Management Suite Enterprise
Node Type	(3) License for Windows
Total Cost	**$52,125.00**

Table 47 – Vmware vCenter Software Licensing

8.3.4 Services BOM

Due to the size and nature of this particular large sized business example, no services are required.

8.3.5 Total Cost of Solution

Here is the total cost of hardware, software, and services for a complete storage system solution for small businesses:

Line Items	Cost
R720 RHEL storage controller	$91,029.30
R720 RHEL Backup servers	$58,109.64
IBM Tape Libraries TS3200	$70,720.00
R720 RHEL Infra servers	$87,164.46
R720 DataCore Windows	$69,029.46
R720 Windows Infra servers	$87,638.76
R820 VMware hypervisors	$491,462.08
PowerVault MD3060e Fast	$500,309.07
PowerVault MD3060e Medium	$53,468.22
PowerVault MD3060e Slow	$115,455.34
Black Box Racks	$14,004.00
Racks with AC	$11.797.90
Rack mount KVMs	$9,239.94
Juniper SRX650	$148,724.00
Juniper QFX5100	$95,396.00
Cisco 6513 Access Switches	$1,121,158.56
Juniper EX3300	$63,278.64
Brocade 7800	$37,656.99
Brocade 6520	$137,859.98
Windows 2012 R2 Data Center Edition	$100,000.00
DataCore SANsymphony-V	$400,000.00
IBM TSM	$72,646.00
Vmware vCenter	$52,125.00
Total Cost	**$3,876,475.44**

Table 48 - Total Cost of Procurement

8.3.6 Architecture and Interconnect Drawings

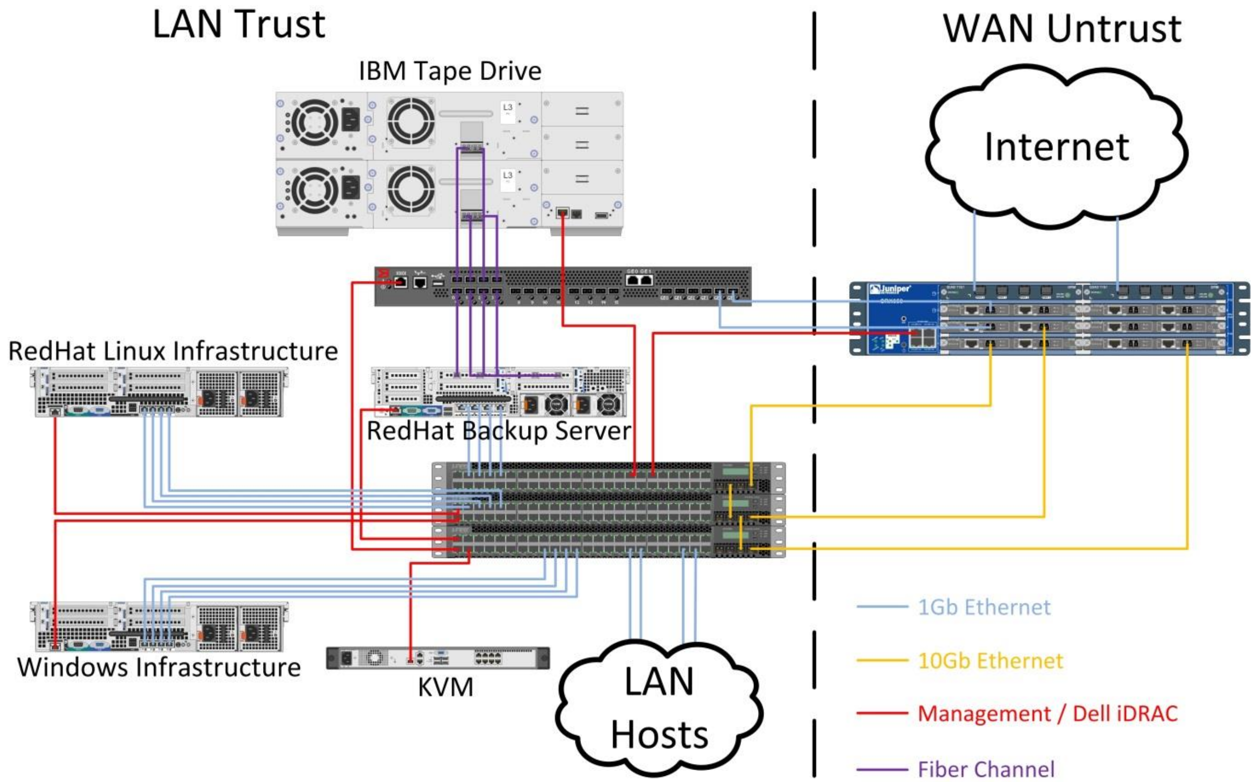

Figure 34 - Sunnyvale Office Interconnect Diagram

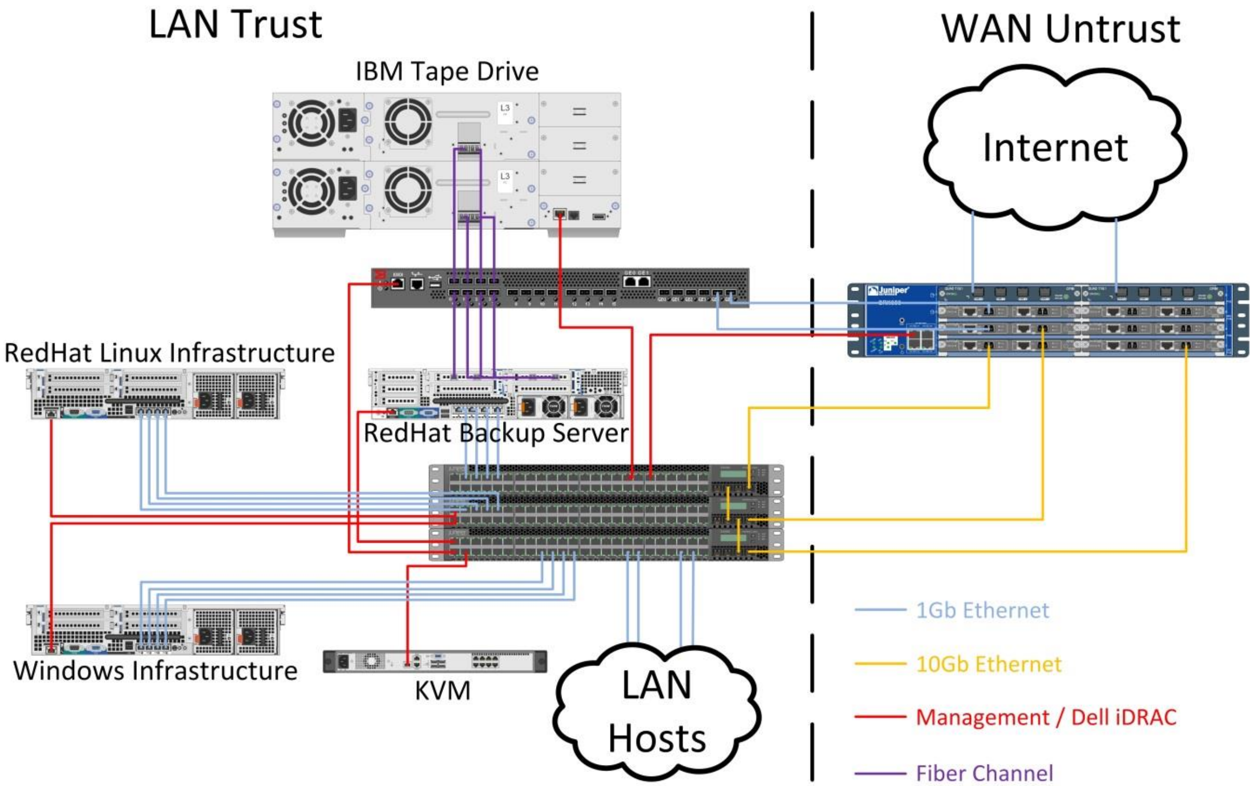

Figure 35 - New York Office Interconnect Diagram

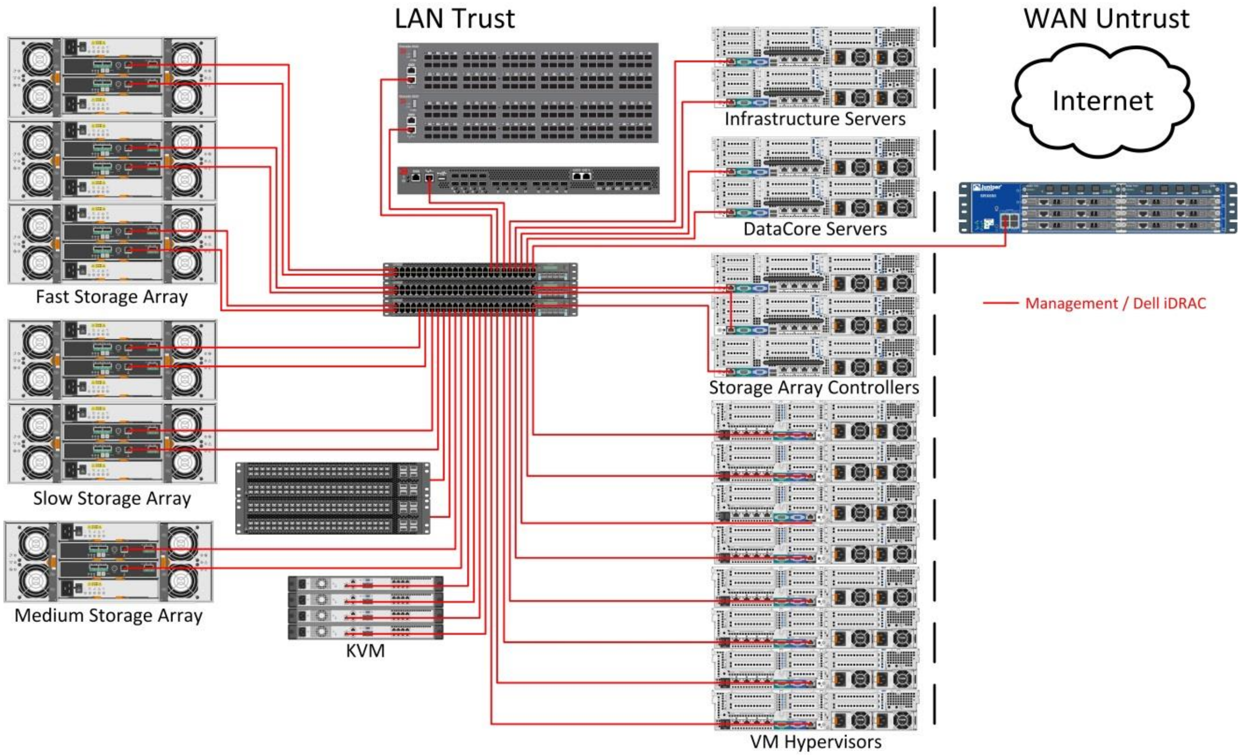

Figure 36 - Denver Data Center Interconnect Diagram for Management Network

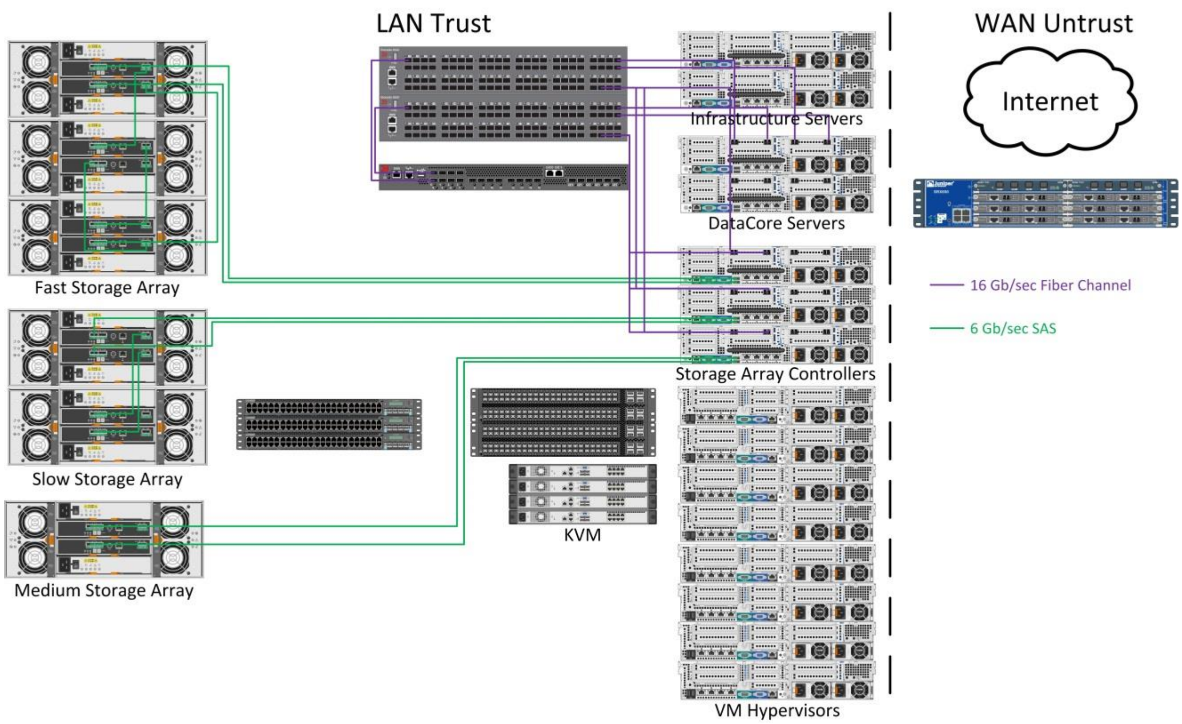

Figure 37 - Denver Data Center Interconnect Diagram for Storage Connections

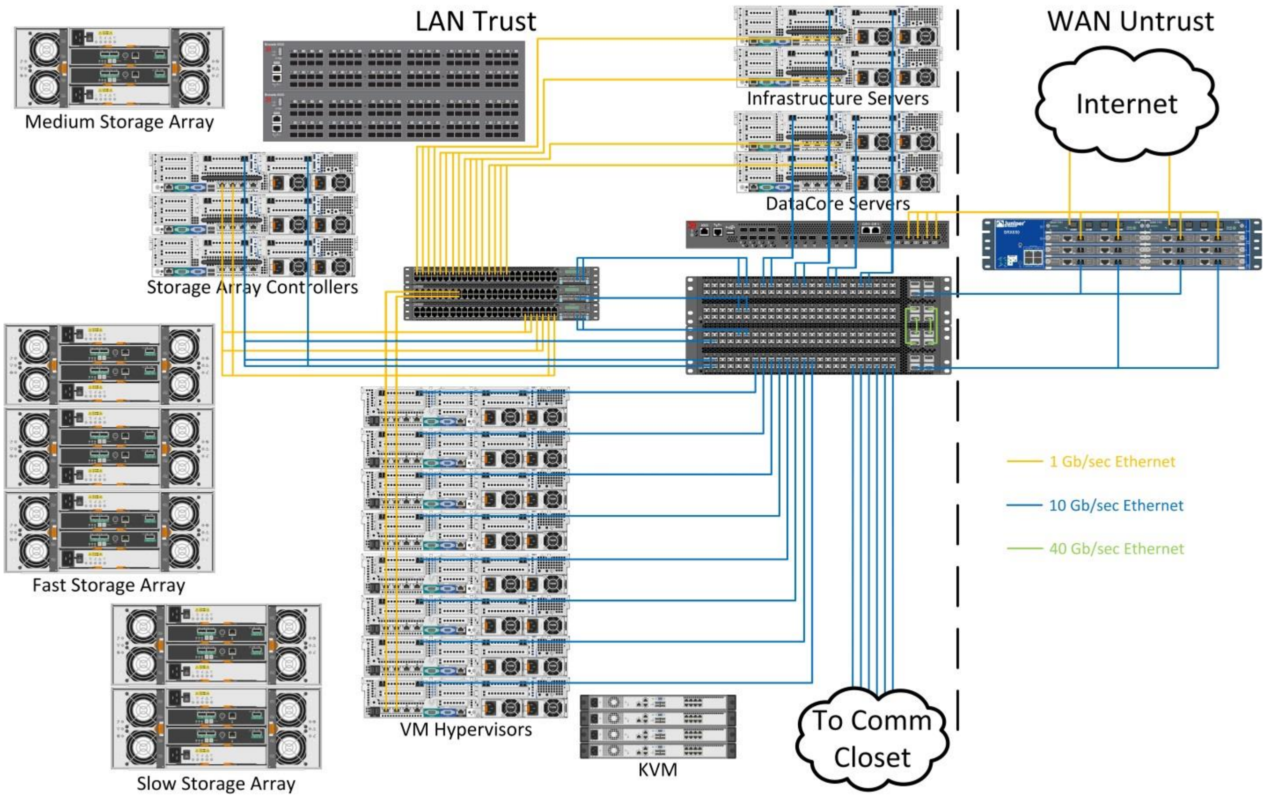

Figure 38 - Denver Data Center Interconnect Diagram for Ethernet Network

In order to avoid confusion with all connections on a single diagram, the three interconnect diagrams above culminate to depict the entire connectivity needed for the large business model. The interconnect diagrams above show Ethernet supporting management, Ethernet supporting both copper and fiber optic connections from 1 Gbps up to 40 Gbps, and a storage connections diagram showing SAS interfaces from the storage devices to the HBAs on the storage controllers and fiber channel connectivity.

Juniper networking equipment has the ability to configure logical networks or logical systems known as LSYS. Logical systems allow a single physical switch to be partitioned into many separate networks providing security boundaries and smaller broadcast domains. This technology will be used to divide the management network from the operational network to cut down on traffic to avoid performance degradation. To utilize this capability, you do not have to trunk your devices. You can simply have an

LSYS span multiple switches, similar to a VLAN. To achieve the interconnect above, the Juniper EX3300 switch, QFX5100 switch, and SRX650 firewall contain the knowledge of all LSYS configurations in the enterprise.

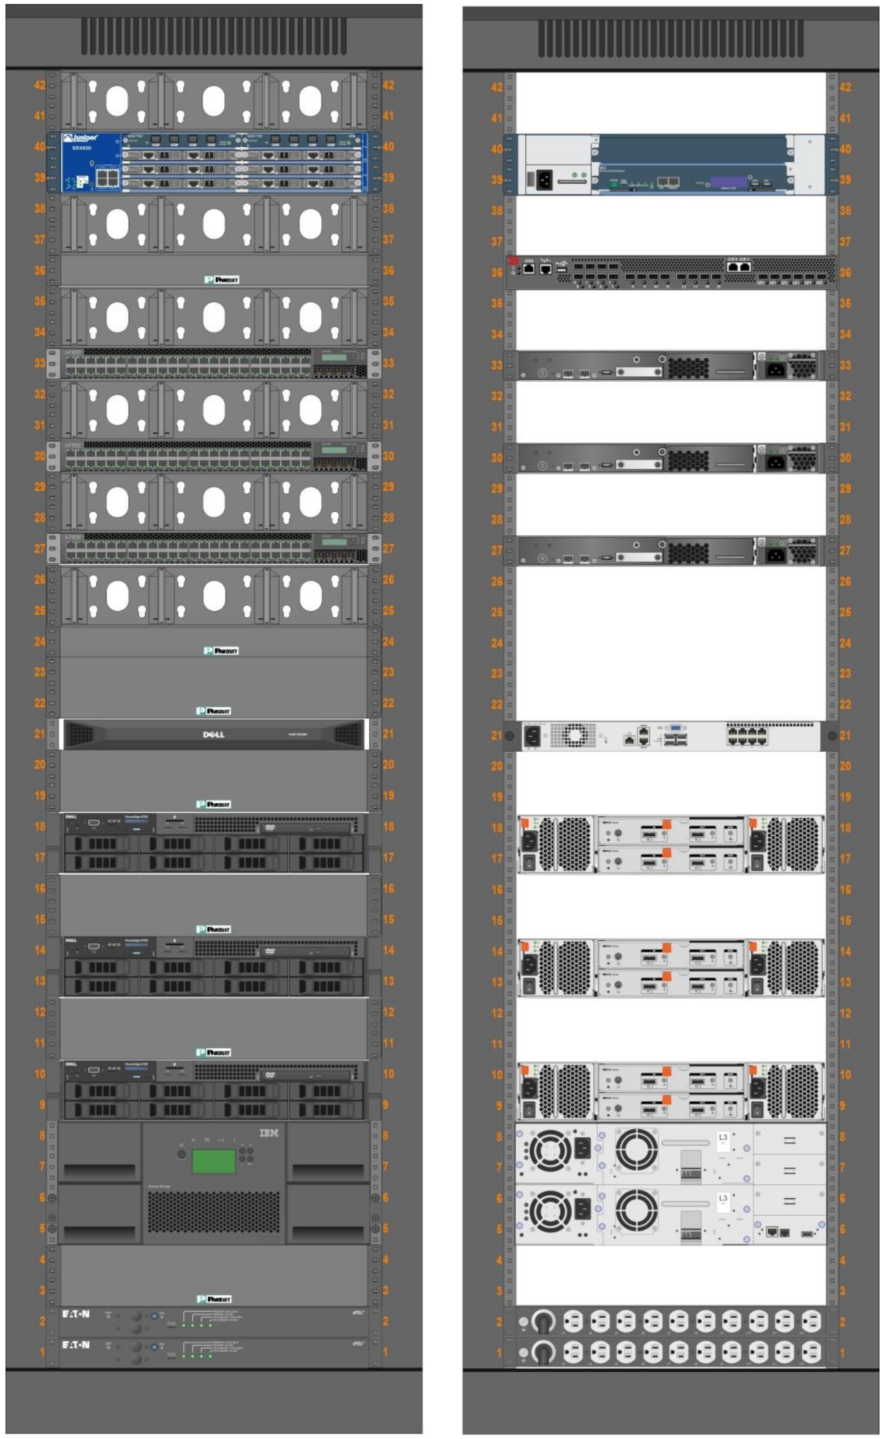

Rack 1 – Front – Sunnyvale, CA   Rack 1 – Back – Sunnyvale, CA

Figure 39 - Sunnyvale Rack Layout

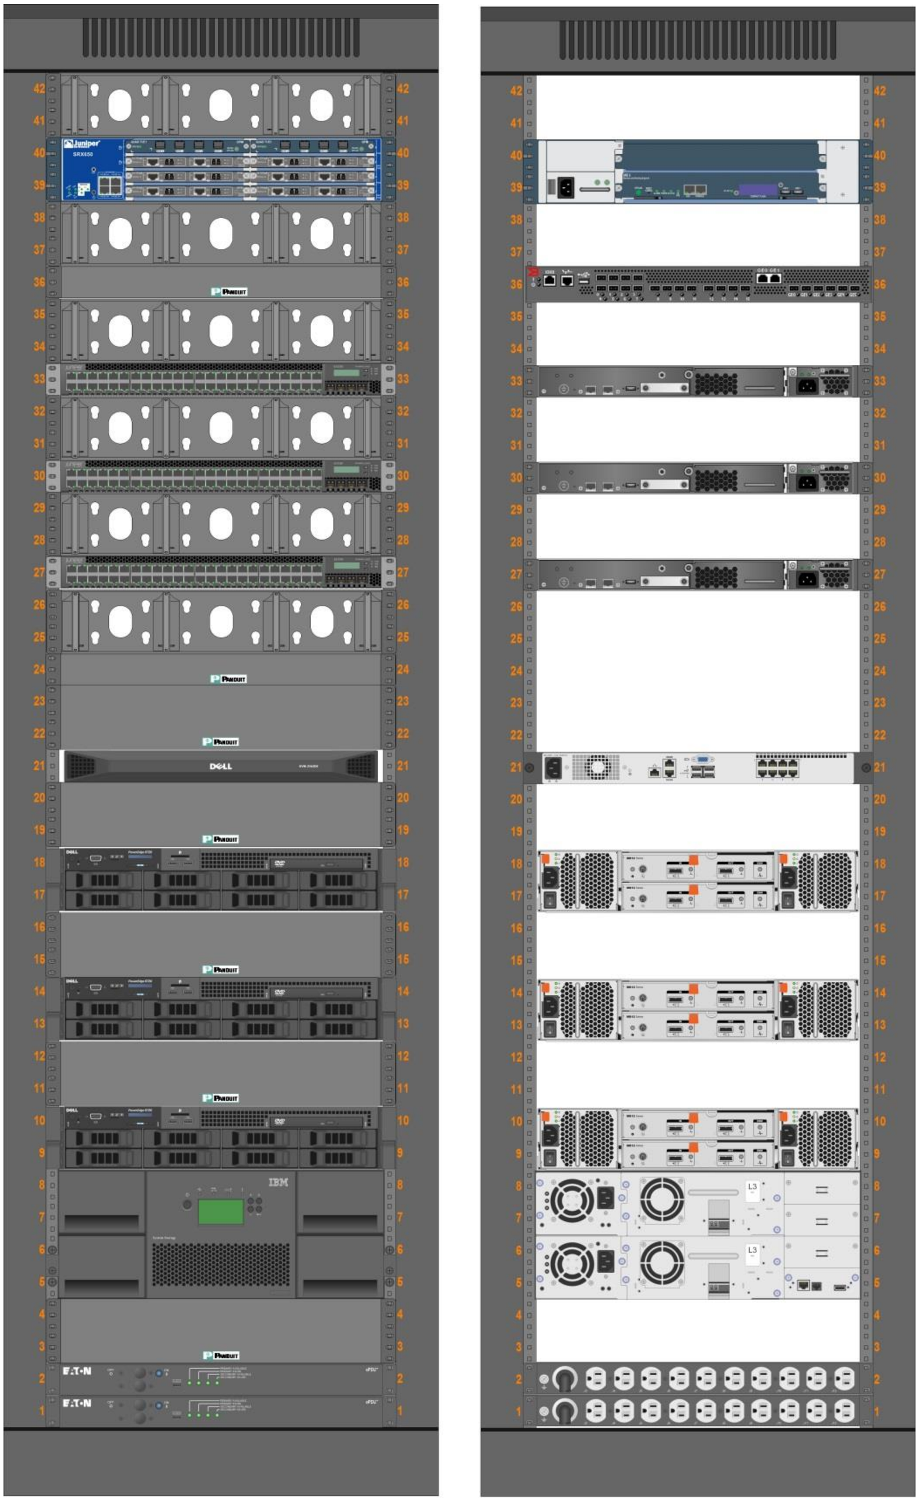

Rack 1 – Front – New York, NY

Rack 1 – Back – New York, NY

Figure 40 - New York Rack Layout

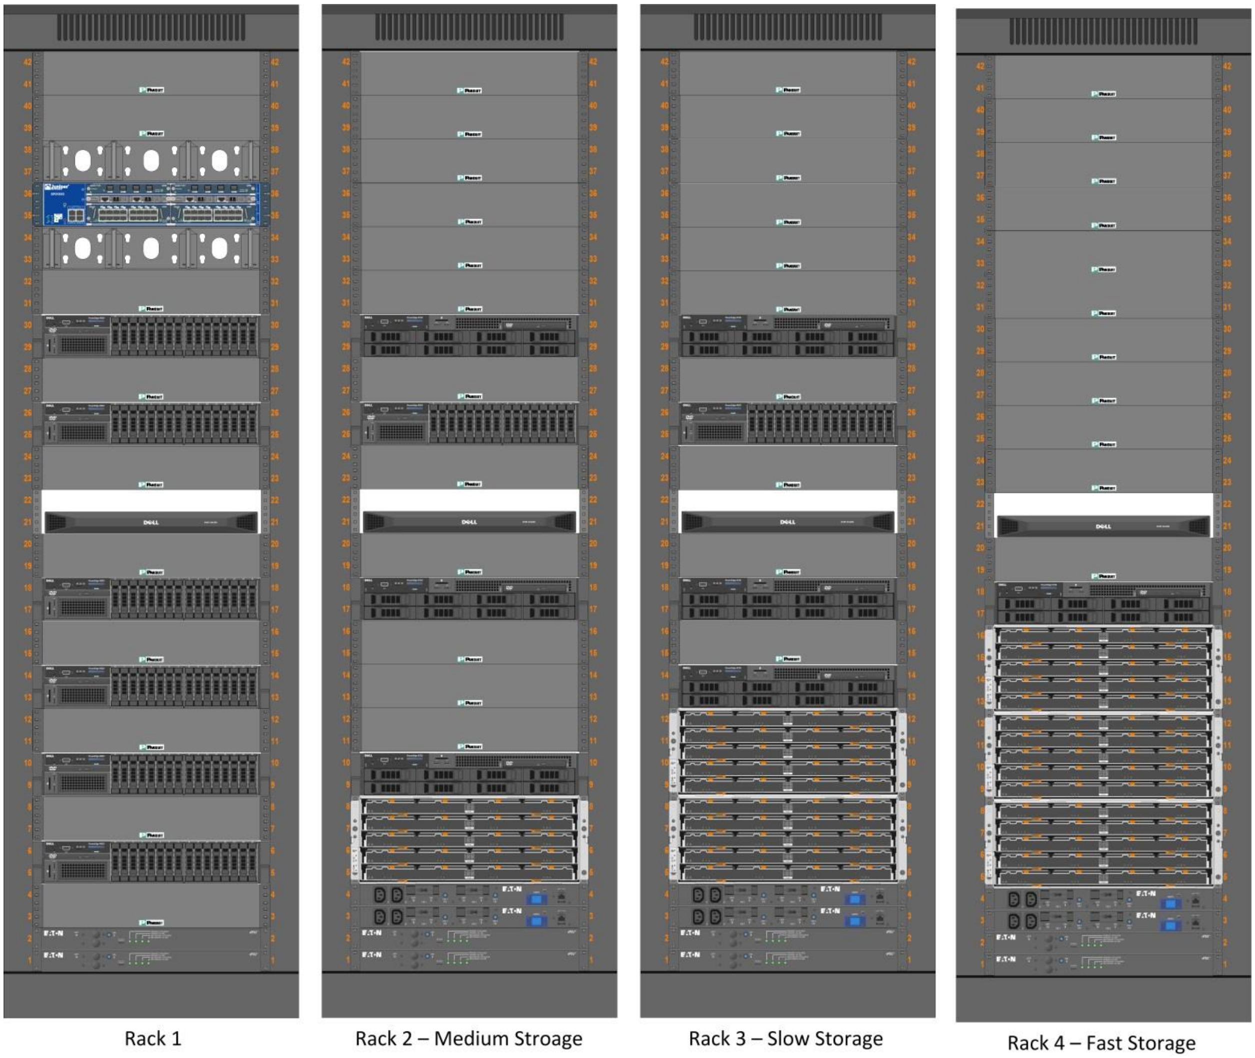

Figure 41 - Denver Rack Layout Front

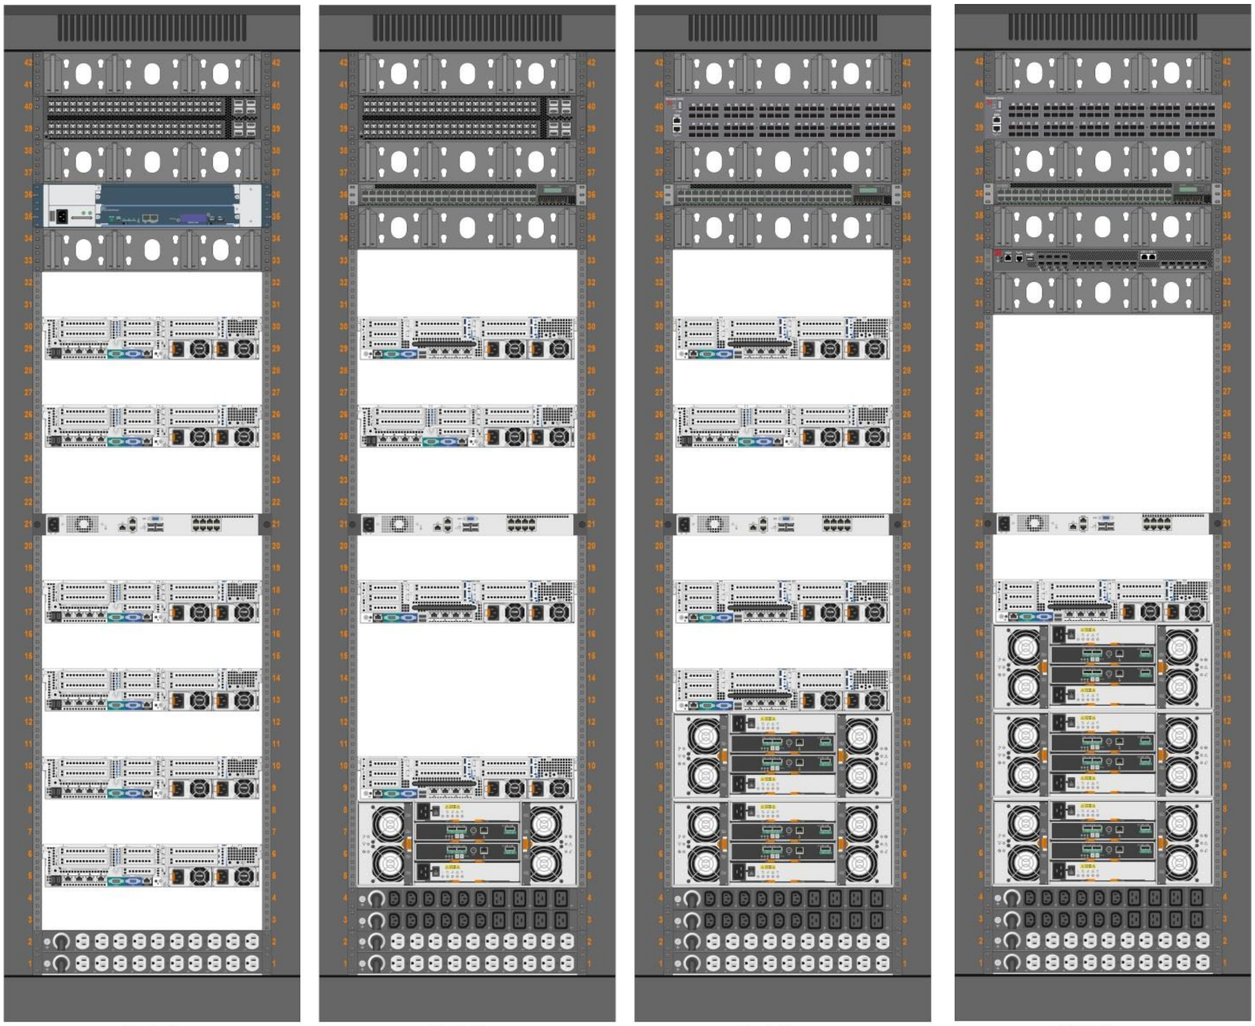

Figure 42 - Denver Rack Layout Back

8.3.7 Storage Virtualization Design

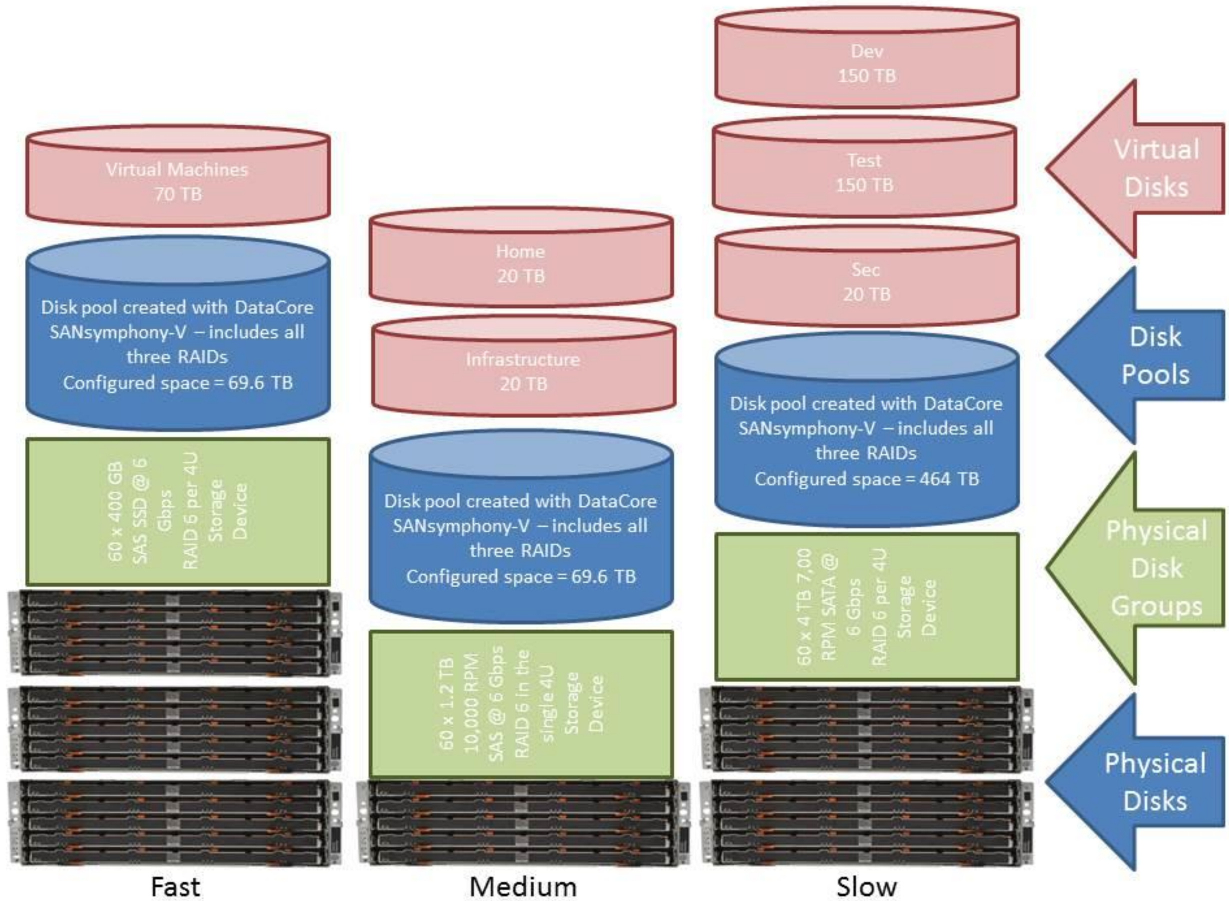

Figure 43 – Physical and Virtual Storage Allocations

The internal multi storage array is a heterogeneous set of disks tied together with DataCore's storage virtualization solution to provide virtual disks within assorted medium options in the disk pools. In the Dell MD3660e storage array, each shelf contains 12 disks in a 4 by 3 configuration. Each storage array was configured to fill each disk slot with a drive and leave no spot empty resulting in a total of 360 disks of internal storage for a total of 624 TB of raw storage capacity.

The low level storage virtualization technologies comprised into this design include perpendicular storage recording on the spinning disks and RAID configurations within the storage arrays for fault tolerance. RAID 6 was used because the performance is similar to RAID 5 but provides and additional parity disk for added redundancy.

At the software layer, storage virtualization technologies are plentiful due to DataCore's SANsymphony-V COTS product. Auto tiering of storage is accomplished by nature of the software and how it is used to create disk pools. Above the physical disk group layer, which is done on the storage device itself, the software product defines a single or multiple disk pools within in a specific disk type or across heterogeneous disk groups. In our large sized business sample architecture, a single disk pool across fast SSDs, medium speed 10,000 RPM SAS drives, and slower 7,200 RPM SATA drives was created. This disk pool within DataCore is allowed to create thinly provisioned virtual disks from the disk pool. Virtual disks can be granted a priority level which dictates which type of disks to migrate the virtual disk to, but in general, if all priorities are the same, DataCore has the ability to span its virtual disks at the block level across heterogeneous disk groups. This concept incorporates thin provisioning with over-allocation and automated tiered storage.

The capabilities of DataCore's SANsymphony product and IBM tape library devices are discussed at length in Section 8.2.7 as the products for both the medium and large scale business application were the same.

8.3.8 Expandability & Future Proofing

The SAN was configured with two 96-port 16 Gbps fiber channel Brocade switches. Only a few of these ports were used in the design providing room for expansion for additional storage or another DataCore server if throughput on the storage virtualization server becomes a bottleneck.

One of the key line items in the hardware BOM are the eight Dell R820 servers. These four servers are the VMware hosts for Infrastructure, Development, and Test environments. VMware is configured for booting the operating systems from the Dell MD3660e with the SSD's in a shared NFS file systems. The alternative to a simple NFS share with many VMDKs is to use the DataCore SANsymphony-V COTS product to show the operating system partitions straight through to the hypervisors as raw LUNs.

The raw partitions are configurable per virtual machine on each of the hypervisors shown through as raw storage. When the hypervisor boots the virtual machine, the disk is unformatted space.

Non-disruptive data migration and mirroring is also possible with the storage virtualization solution chosen. DataCore has the ability to send its node into maintenance mode and allows a redundant copy to run from a separate data center. This storage can also be migrated between data centers and onto different storage arrays with no downtime. This technology is very implementable in this solution; however, no storage subsystems were identified as mirrored copies in differing data centers. This product can handle a mirrored copy on a storage array from a totally different manufacturer with different disk types and still function properly thanks to its clever ambiguous storage pools. The technology exists for future proofing the solution.

The IBM TS3200 can have their drives upgrade when the new technology is available for additional storage capacity and faster transfer mediums beyond LTO 6. Also available within the backup architecture is to add additional TSM servers to load balance the backups. Due to the CPU and memory intensive churning through data to determine the hash of every file to determine whether or not it needs to be backed up could delay backups into the following business day. At this point, it is advised to add an additional backup server to offset the load.

As part of the storage architecture, Brocade 7800s were allocated for each of the 3 facilities. These are in place to facilitate a fiber channel SAN fabric that spans from coast to coast supplying the remote facilities the abilities to conduct backups offsite. This also positions the large business for future data growth by have the storage fabric available at all facilities through FC over IP transmitted over the WAN.

8.3.9 Complexity

The initial setup and configuration of the storage system solution would require a team of experts. The team of experts would need knowledge in SANs, large sized networks with layer 3 switching, VMware, DataCore SANsymphony-V, Windows, RedHat, and tape libraries. Once the configuration is up, running, and configured properly, the solution would require on staff expertise to manage the data center. Due to the complex nature of the tiered storage, numerous virtual machines, and remote elements of offsite back infrastructure, an IT administrator would be required at each of the 3 facilities. The data center requires the most expertise to include a network engineer, storage engineer, and system administrator. The system requires constant monitoring of the infrastructure stack looking for faulty virtual machines, storage allocations running low, and hardware and networking issues or faults.

8.4 Assumptions

In this section, there exists a bulleted list of assumptions regarding the pricing and configurations of the equipment in Section 8:

1.) Unless otherwise specific, all networking and SAN fabric equipment contained some flavor of SFP for connectivity.
2.) Equipment was configured and priced from a multitude of web site from July 2014 through November 2014. Prices may not reflect what future consumers would experience. It is always recommended to go through your respective sales representatives for respective manufacturers.

9 References

Aikema, D. (2013, April). *Optimizing Data Centre Energy and Environmental Costs.* Retrieved from http://theses.ucalgary.ca/bitstream/11023/664/2/ucalgary_2013_aikema_david.pdf

Amazon. (n.d.). *Amazon*. Retrieved from Amazon EC2: http://aws.amazon.com/ec2/

Apple. (2014, August 6). Retrieved from http://support.apple.com/kb/ht4874

ASPG. (2013, June 5). *Mainframes vs. Supercomputers*. Retrieved from Advanced Software Products Group, Inc.: http://aspg.com/mainframes-vs-supercomputers/#.U_EAPWOHPwM

Brocade. (2009). *Finding the Best Approach for Efficient, Non-Disruptive Data Migration.*

Brocade. (2013, May). *SAN Design and Best Practices (Version 23)*. Retrieved from http://www.brocade.com/downloads/documents/best_practice_guides/san-design-best-practices.pdf

BusinessGreen Staff. (2013, June 13). Retrieved from Green Biz: http://www.greenbiz.com/blog/2013/06/13/facebook-flicks-hydropowered-data-center-sweden

Calder, B. G., Nilakantan, N., Srivastav, S., Wu, J., Abbasi, A. R., Mainali, S., et al. (2011, May 23). *Patent No. US 20120303576 A1.*

Carbonite. (n.d.). Retrieved from Server Plans: http://www.carbonite.com/backup/pricing-plans/server-plans

Carl, S. (2011). *The Green Economgy Post*. Retrieved from Watts Up: Does Virtualization Really Save Energy?: http://greeneconomypost.com/virtualization-data-center-energy-efficiency-19009.htm

Cason, J. (2014, May 28). *Thoughts on Thin Provisioning*. Retrieved from RAID-ZERO.COM: http://raid-zero.com/2014/05/28/thoughts-on-thin-provisioning/

Cassioli, L. (2005, June 21). Retrieved from http://en.wikipedia.org/wiki/Perpendicular_recording#mediaviewer/File:Perpendicular_Recording_Diagram.svg

Cisco. (2006, October 26). Retrieved from http://www.cisco.com/image/jpg/en/us/guest/products/ps5479/c1244/ccmigration_09186a00801b97b3_09186a008076b044-1.jpg

Clark, T., & Yoder, A. (2008, October). *Best Practices for Energy Efficient Storage Operations.* Retrieved from http://www.snia.org/sites/default/files/GSI_Best_Practices_V1.0_FINAL.pdf

Data Protect Works. (2013, February 21). Retrieved from Progressive Deduplication: http://www.dataprotectworks.com/images/software/Progressive_Deduplication_sm.png

DataCore Software. (n.d.). *Auto-Tiering*. Retrieved from DataCore Software: http://international.datacore.com/Software/Features/List-of-Features/Auto-Tiering.aspx

Dell. (2010). Retrieved from Dell PowerVault MD1220 Technical Guidebook: http://www.dell.com/downloads/global/products/pvaul/en/storage-powervault-md12x0-technical-guidebook.pdf

Dell. (2013). *The History of Data Storage.*

DiBello, L. (2013, June 5). Retrieved from Geotech: http://www.google.com/imgres?imgurl=http%3A%2F%2Fwww.flowserve.com%2Ffiles%2FFiles%2FImages%2FIndustries%2Fstatic%2F020-I-Geothermal_Flashed_Steam_large.png&imgrefurl=http%3A%2F%2Ffahygeo.blogspot.com%2F&h=792&w=1224&tbnid=g0v7dnSTg2SOyM%3A&zoom=1&docid=5

EBI-Tech. (n.d.). Retrieved from http://www.etb-tech.com/powervault-md3660f-20-x-300gb-15k-sas.html#.U_F9BGOHPwM

Gartner. (2010, February 19). Gartner: only 6 percent of companies survive longer than two years after losing data. *Homeland Security News Wire*.

Geothermal Energy Association. (2014). *Annual U.S. & Global Geothermal Power Production Report.*

Geothermal Energy Systems. (2009). *Distribution of energy within a building*. Retrieved from http://www.africanecosystems.co.za/building%20distribution.html

Gigerich, L. (2012, April). Critical Factors: The Data Center Location Decision. *Area Development - Site and Facility Planning*.

Gilmer, E. (2011, October 17). *Can an IT snowball grow in Iceland?* . Retrieved from EE News: http://www.eenews.net/stories/1059955037/print

Global Energy Network Institute. (2014, June 3). Retrieved from Global Energy Network Institute: http://www.geni.org/globalenergy/library/articles-renewable-energy-transmission/graphics/googleplex.jpeg

Google. (n.d.). *Google Apps for Business*. Retrieved from http://www.google.com/intx/en/enterprise/apps/business/pricing.html

Hess, K. (2011, February 24). *10 Reasons Why SSDs Are Better Than Mechanical Disks*. Retrieved from ServerWatch - IT Business Edge: http://www.serverwatch.com/trends/article.php/3926131/10-Reasons-Why-SSDs-Are-Better-Than-Mechanical-Disks.htm

Holm, A., Blodgett, L., Jennejohn, D., & Gawell, K. (2010). *Geothermal Energy: International Market Update.*

Information Age. (2012, October 10). BMW moves into Icelandic data centre.

Kern, E. (2012, August 22). *Facebook is collecting your data — 500 terabytes a day*. Retrieved from Gigaom: http://gigaom.com/2012/08/22/facebook-is-collecting-your-data-500-terabytes-a-day/

Latimer, J. (2011, December 10). *Gigaom*. Retrieved from Top 5 places to build a new data center: http://gigaom.com/2011/12/10/latimer-where-to-build-data-center/

Lennett, B., & Meinrath, S. (2009, January). *New America Foundation*. Retrieved from Building a 21st Century Broadband Superhighway: http://www.newamerica.net/publications/policy/building_21st_century_broadband_superhighway

Lubbers, C., Elkington, S., Hess, R., Sicola, S. J., McCarty, J., Korgaonkar, A., et al. (2002, March 26). *Patent No. US 6880052 B2.*

McArthur, J. T., & Villars, R. L. (2005, September). *White Paper: Storage Virtualization: Analysis of the Benefits of Hitachi's Approach.* Retrieved from IDC - Analyze the Future: http://www.hds.com/assets/pdf/idc-storage-virtualization-analysis-of-the-benefits-of-hitachis-approach-wp.pdf

Miller, R. (2009, August 13). *Geothermal Data Center is LEED Platinum*. Retrieved from Data Center Knowledge: http://www.datacenterknowledge.com/archives/2009/08/13/geothermal-data-center-is-leed-platinum/

Nadel, B. (2012, February 6). *Computer World*. Retrieved from 5 online backup services keep your data safe: http://www.computerworld.com/s/article/9223805/5_online_backup_services_keep_your_data_safe

Net DNA. (n.d.). Retrieved from http://4f70kd3zi1axzrkd71h56h2e3e.wpengine.netdna-cdn.com/wp-content/uploads/2013/09/RAID-5-vs-RAID-6-Comparison.jpg

Okwuke, E. (n.d.). *Global installed data storage capacity to hit 7,235 exabytes by 2017 – Report.* Retrieved from http://dailyindependentnig.com/2013/10/global-installed-data-storage-capacity-to-hit-7235-exabytes-by-2017-report/

Petemarovich Images. (2013, November 24). Retrieved from http://www.petemarovichimages.com/wp-content/uploads/2013/11/RAID1.jpg

Prairie Bunkers Data Centers. (n.d.). Retrieved from http://prairiebunkers.com/sustainability-2/

Protalinski, E. (2014, January 29). Retrieved from Facebook passes 1.23 billion monthly active users, 945 million mobile users, and 757 million daily users: http://thenextweb.com/facebook/2014/01/29/facebook-passes-1-23-billion-monthly-active-users-945-million-mobile-users-757-million-daily-users/

Pyranak Design Group. (2004.). *The History of Geothermal Heating and Cooling*. Retrieved from http://www.pyranak.com/4greenbldg/Geothermal/History%20of%20geothermal.html

Radhakrishnan, P., & Singh, V. (2010, September 8). *Patent No. US 8583867 B1.*

Rouse, M. (2010, July). *whatis.com*. Retrieved from Data deduplication (Intelligent compression or single-instance storage): http://searchstorage.techtarget.com/definition/data-deduplication

Rouse, M. (n.d.). *whatis.com*. Retrieved from Thin Provisioning: http://searchstorage.techtarget.com/definition/thin-provisioning

RTCC. (2012, September 12). *Doha talks open door to Gulf-funded climate change technology*. Retrieved from Responding to Climate Change: http://www.rtcc.org/files/2012/05/masdar-funded-gemasolar-csp-in-spain-source-masdar-400px.jpg

SAN Storage Works. (n.d.). Retrieved from http://www.sanstorageworks.com/images/PowerVault/MD3600/MD3620f-front-array.png

Schmid, P., & Roos, A. (2008, June 26). *Tom's Hardware*. Retrieved from The SSD Power Consumption Hoax: http://www.tomshardware.com/reviews/ssd-hdd-battery,1955-14.html

Secretariat, P. R. (2014). *Estimated Renewable Energy Share of GLobal Electricity Production, End-2013*. Retrieved from http://www.ren21.net/Portals/0/Images/Figures/GSR2014/Figure_3_Electricity_oNr.jpg

SoftNAS Cloud Features. (n.d.). Retrieved from Inline Block Deduplication: https://www.softnas.com/wp/nas-storage/softnas-cloud-aws-nfs-cifs/features4/

Stewart, V. (2013, December 10). *Pure Storage*. Retrieved from Data Deduplication and Compression Means Double Savings - See more at: http://www.purestorage.com/blog/data-deduplication-and-compression-means-double-savings/#sthash.T7e494tz.dpuf: http://www.purestorage.com/blog/data-deduplication-and-compression-means-double-savings/

StorageNewletter.com. (2012, June 11). *Licensing Specs for LTO-6 Available in August, 2012*. Retrieved from Storage Newsletter: http://www.storagenewsletter.com/rubriques/tapes/licensing-specs-august-2012/

Symantec. (2012, February 3). *Symantec Enterprise*. Retrieved from About best practices for using the FileStore deduplication feature: http://www.symantec.com/business/support/index?page=content&id=HOWTO60443

Thin-Provisioning. (2013, May). Retrieved from Raxco: http://blog.raxco.com/wp-content/uploads/2013/05/Thin-Provisioning.jpg

Tomia. Hydroelectic Dam. *Wikimedia Commons.* Creative Commons Attribution 2.5.

Vance, A. (2013, October 4). *Inside the Arctic Circle, Where Your Facebook Data Lives*. Retrieved from Bloomberg Businessweek Technology: http://www.businessweek.com/articles/2013-10-03/facebooks-new-data-center-in-sweden-puts-the-heat-on-hardware-makers

Vattenfall. (2009, September 20). Retrieved from http://www.solaripedia.com/images/large/707.jpg

Verma, A., Koller, R., Useche, L., & Rangaswami, R. (2010). *SRCMap: Energy Proportional Storage using Dynamic Consolidation.* Retrieved from https://www.usenix.org/legacy/event/fast10/tech/full_papers/verma.pdf

VMware. (2014, August 1). *Knowledge Base - Thin Provisioned Disks*. Retrieved from http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1005418

Warden, T. (2006, November 17). *Storage Pooling, Thin Provisioning And Over Subscription*. Retrieved from Las Solanas Consulting: http://www.las-solanas.com/storage_virtualization/thin_provisioning.php

Well Drillings School. (n.d.). *Geothermal Heating & Cooling Systems.*

Whitehouse, L. (2008, September). *The pros and cons of file-leve vs. block-level data deduplication technology*. Retrieved from TechTarget: http://searchdatabackup.techtarget.com/tip/The-pros-and-cons-of-file-level-vs-block-level-data-deduplication-technology

Whitney, L. (2010, March 25). *CNET*. Retrieved from CompTIA offering Green IT exam: http://www.cnet.com/news/comptia-offering-green-it-exam/

Wired. (n.d.). Retrieved from Submarine Global Internet Connectivity: http://cdni.wired.co.uk/620x413/s_v/Submarine%202012%202_620x413.jpg wiseGEEK. (n.d.). *What Is Data Proliferation*. Retrieved from wiseGEEK: http://www.wisegeek.com/what-is-data-proliferation.htm

Woods, J. (2013, March 5). *Silicon Angle*. Retrieved from The evolution of the data center : Timeline from the Mainframe to the Cloud: http://siliconangle.com/blog/2014/03/05/the-evolution-of-the-data-center-timeline-from-the-mainframe-to-the-cloud-tc0114/