



KNOWLEDGE SHARING 2011 BOOK OF ABSTRACTS

An EMC Proven Professional Publication

EMC²

TABLE OF CONTENTS

FIRST-PLACE KNOWLEDGE SHARING ARTICLE

Data Migration—Best Way to Do It?

Randeep Singh, HCL Comnet

2

SECOND-PLACE KNOWLEDGE SHARING ARTICLE

Data Warehouse Fundamentals for Storage Professionals— What You Need to Know

Bruce Yellin, EMC

3

THIRD-PLACE KNOWLEDGE SHARING ARTICLE

Practical Guide to Challenges Faced with Centralization of Records

Sonali Bhavsar, Microlink LLC

4

BEST OF SAN ASSESSMENT STRATEGY

3D SAN Audit—The Most Effective Methodology to Assess Your Infrastructure

Vasily Pantyukhin, EMC

4

BEST OF DATA PROTECTION

Data Protection in a Nutshell—A Tiered Approach

Luiz Pissinatti, EMC

5

BEST OF VDI

Addressing the VDI Storage Bottleneck

Vinayak Sivanand, Tata Consulting Services Limited

6

ARCHIVE

Synergistic Opportunity Between DiskXtender and PowerPath

Vinay Parthasarathy, EMC

6

BACKUP AND RECOVERY

A System Administrator's Approach to Monitoring Avamar

Kelly Mills, EMC

7

NetWorker Best Practices with Automation

Ajit Kumar Rath, EMC

8

How to Start Your Disaster Recovery in this “Cloudy” Landscape

Roy Mikes, Mondriaan Zorggroep

9

NetWorker Data Migration—EDL to Physical Tape

Gurpreet Dhillon, EMC

10

CELERRA AND CLARIIION

Celerra SnapSure Best Practice Insight of its Space Management and Performance

Frank Zhao, EMC

Kevin Xu, EMC

11

Integrating and Troubleshooting VMware with EMC CLARiiON: Deep Dive

Jason L. Gates, Presidio Networked Solutions

11

Celerra Disk Reconfiguration and Advanced Management

James Ingle, Alexander Open Systems

12

CLOUD

Deadly Trio Integration

Anuj Sharma, EMC

12

Why Cloud Computing?

Bashanta Phukon, EMC

14

Riding the Clouds—Best Practices in Data Center Transformation into the Next Century

Paul Brant, EMC

15

Optimizing IT Services using Cloud, Data Center, and Adherence to Compliance Standards <i>Amrith Raj Radhakrishnan, Tata Consultancy Services Limited</i>	16
CONTENT MANAGEMENT	
A Guide to Enterprise Content Management <i>Mouli Ponnupandy, Keane, an NTT DATA Company</i>	17
Global Scale Documentum Implementation—Design Considerations <i>Sameer Patil, Tata Consultancy Services Limited</i>	18
SAN	
EMC SAN Copy Migration—A Complete Experience <i>Dinesh Pahari, Dimension Data Australia Pty Ltd.</i>	19
Host integration with EMC Storage Arrays <i>Upanshu Singhal, EMC</i> <i>Shelesh Chopra, EMC</i>	20
STORAGE	
Using the Building Block Approach for Scaling Enterprise Storage <i>Mark Rogov, EMC</i>	21
Entering Yottabyte Era using Enterprise Flash Drive Technology <i>Milan K Mithbaokar, Dell Inc.</i>	21
Architectural Plans for Enabling a Dynamic Data Center for Today and in the Future <i>Puneet Goyal, HCL Comnet</i>	22
Information-in-Motion (IiM) <i>Lalit Mohan, Dell Global B.V</i>	22
Storage Capacity Management: Simplicity is Efficiency <i>Olumide Ademidun, ING Bank Canada</i>	23
STORAGE MANAGEMENT	
Automating EMC IONIX ControlCenter Host Agent Deployment and Maintenance Process <i>Mehmet Ozdemir, EMC</i>	23
STORAGE NETWORKING	
Fibre Channel over Ethernet for Beginners <i>Bruce Yellin, EMC</i>	24
Convergenomics of FCoE <i>Akshat Sawhney, EMC</i> <i>Sahil Gupta, Ace Data Devices</i>	25
SYMMETRIX	
Using SE Configchecker for Checking Software Compatible Pre-requisites for Symmetrix <i>Parthasarathi Ilangovan, EMC</i>	26
VIRTUALIZATION	
Discovering VIO Relationships for EMC Arrays <i>Ewa Nowatkowska-Glowacz, EMC</i>	26
VPLEX: Interesting Use Cases <i>Sandeep Choudhary, Tech Mahindra</i> <i>Puneet Goyal, HCL Comnet</i>	27
Storage and Virtualization Technologies in Information Technology Education <i>Dr. Timur Mirzoev, Georgia Southern University</i>	27



Knowledge Sharing Winners 2010 Awards (left to right): Sejal Joshi, John Bowling, Ben Dingley, John Shubeck, Denis Guyadeen, Paul Brant, Rajaneesh Sasidharan, Mohammed Hashim, Ken Guest

THANK YOU!

For the fifth consecutive year, we are pleased to recognize our EMC® Proven™ Professional Knowledge Sharing authors. This year's Book of Abstracts demonstrates how the Knowledge Sharing program has grown into a powerful platform for sharing ideas, expertise, unique deployments, and best practices among IT infrastructure professionals. Articles from our contributing Knowledge Sharing authors have been downloaded more than 212,000 times, underscoring the power of the knowledge sharing concept. View our full library of Knowledge Sharing articles, and articles from the 2011 competition, published monthly, at: <http://education.EMC.com/KnowledgeSharing>.

Our continuing success is built on the foundation of committed professionals who participate, contribute, and share. Through the Knowledge Sharing program, your industry-leading expertise can reach the wider audience of IT professionals and further strengthen the thought leadership position of all EMC Proven Professionals.

We thank all those who participated in the 2011 Knowledge Sharing competition.

Tom Clancy
Vice President
EMC Education Services

Alok Shrivastava
Senior Director
EMC Education Services

FIRST-PLACE KNOWLEDGE SHARING ARTICLE

DATA MIGRATION—BEST WAY TO DO IT?

Randeep Singh, HCL Comnet

Relocating a data center can be a complex endeavor. Because of the high risk and low margin for errors, it can be one of the most challenging IT projects any organization will ever undertake. One of the toughest activities during data center relocation is data migration because that is where your whole business lies. Ask any storage professional for a list of his/her biggest challenges and data migration will be at, or near, the top of the list. Often, data migration is required during technology refresh initiatives where newer storage systems replace legacy systems.

Many IT projects today involve some kind of corrective measure at the platform level, whether integrating application silos, upgrading packaged applications, consolidating redundant IT systems, or migrating data and applications from old to modern platforms. These project types are related because they all involve moving data from database to database or application to application. Furthermore, they all require similar data modeling and data integration skills, and users apply one or more project types together. Among these, data migration is prominent because it is both a discrete project and a common component of other data-movement projects.

For IT managers, data migration has become one of the most routine and challenging facts of life. Workload balancing, technology refresh, server and storage consolidation, data center relocation, data classification, and mergers/acquisitions are all drivers to move data from one storage device to another on a regular basis. You would think that any operation performed routinely would become easy, but it is not so. Migration—moving data from one device to another (especially as the capacity of storage devices continues to grow) and then redirecting all I/O to the new device—is an inherently disruptive process. With the increase in the percentage of mission-critical data and the proportionate increase in data availability demands, downtime—with its huge impact on a company's financial bottom line—becomes unacceptable. In addition, business, technical, and operational requirements impose challenging restrictions on the migration process itself. Resource demands (staff, CPU cycles, and bandwidth) and risks (application downtime, performance impact to production environments, technical incompatibilities, and data corruption/loss) make migration one of IT's biggest challenges.

Since the majority of storage systems purchased by customers are used to store existing rather than new data, getting these new systems production-ready requires that data be copied/moved from the old system being replaced to the new system being deployed. As the services arm of one of the industry's leading storage companies, EMC Professional Services has developed a core competency in data migration, fine-tuning its migration methodology over time with each migration services engagement delivered. Whether the migration is a same vendor system upgrade (for example, EMC CLARiiON® to EMC CLARiiON) or a vendor take-out (for example, other vendor to EMC), and whether the migration is performed by internal IT or an external services provider, the migration methodology is the same.

This article explores the following:

1. Key Challenges During Data Migration Activity
2. Business Impact
3. Data Migration Key Drivers
4. Data Migration Methodology
5. Data Migration Types

6. Using Virtualization when Migrating Data
7. Data Migration Plan, Phases, and Process
8. Replication Core of Data Migration Process
9. Best Practices for Data Migration
10. Taking Pain out of Data Migration
11. Data Migration Activity—Customer Scenario

In summary, data migration is a routine part of IT operations in today's business environment. Even so, it often causes major disruptions as a result of downtime or application performance problems, and it can severely impact budgets. To prevent these problems, organizations need a consistent and reliable methodology that enables them to plan, design, migrate, and validate the migration. Further, they need migration software that supports their specific migration requirements, including operating systems, storage platforms, and performance. In addition, migration products that maintain continuous data availability during the migration, without affecting performance, are desirable.

SECOND-PLACE KNOWLEDGE SHARING ARTICLE

DATA WAREHOUSE FUNDAMENTALS FOR STORAGE PROFESSIONALS—WHAT YOU NEED TO KNOW

Bruce Yellin, EMC

When organizations build a data warehouse, they do so to answer a number of strategic questions that greatly affect their future as a company. These questions involve such important areas as growing market share and revenue, improving profitability, and investing in their future. The warehouse ties together all sorts of daily and historical internal information along with external data into a living, breathing system that can be mined for informational insight into the business.

In these economically trying times, a great deal of importance is placed on the data warehouse. As a result, warehouse teams face considerable risk trying to keep these huge and highly visible systems available and running at peak performance 24/7. That is where you, the storage professional, comes in. The warehouse group needs your expertise with the design, implementation, and support of essential systems, even though you may not be familiar with data warehouse concepts and terminology.

This Knowledge Sharing article will demystify data warehousing and turn you from a warehouse beginner to a contributing member of a decision support group as you leverage your current storage expertise and newly acquired insight. Using a step by step, easy to follow approach, you will be introduced to the concepts behind data warehousing and data marts. Case studies are offered, including how one company's warehouse helped them increase sales of beer and diapers. You will review design concepts, learn how a warehouse accepts new data, and how other systems feed off of it. Critical warehouse storage topics such as backup, maintenance, virtualization and the cloud, NAS versus SAN, reliability, availability, scalability, performance, and disaster recovery/business continuity are also discussed.

THIRD-PLACE KNOWLEDGE SHARING ARTICLE

PRACTICAL GUIDE TO CHALLENGES FACED WITH CENTRALIZATION OF RECORDS

Sonali Bhavsar, Microlink LLC

As mergers and acquisitions have become the norm in today's business environment, there is an increased focus on governance and compliance to identify relevant data within the existing repositories. What should you do to meet business, governance, risk management, and technology needs to identify a centralized access point for maintaining and controlling siloed information stored by various vendors in their proprietary formats? Should you decommission the existing systems completely and transition to a newer, single system which would not only provide access to existing repositories but also provide more functionalities? Or should you continue to support the legacy systems until a budget is assigned for migration? Is federation of a solution by a single product the ultimate game changer? What is the approach to understand legacy repositories and their content? How do you determine if there is a scalable need to support newer forms of data such as social media and mobile communications, and what are the best practices that you should follow?

Companies are constantly faced with the challenge of how to integrate official records under centralized taxonomy rules, support one entry point for search, and maintain the record repository for audit as well as eDiscovery. How do you develop a successful program to address these basic issues? Should you just install a search appliance which serves search and eDiscovery needs? The answer goes beyond choosing a third-party tool or selecting a vendor to provide a solution. This Knowledge Sharing article provides best practices along with a step-by-step procedure covering strategies you can adopt to centralize records across your company.

The article explores the challenges faced by technology decision makers in looking at all aspects of records stored in disparate repositories spread across multiple vendors (for example: Autonomy, IBM FileNet, Email Archiving—in-house or vendor specific, EMC Documentum®, etc.) and suggests best practices that can be incorporated prior to making a decision for decommissioning or adding new repositories.

BEST OF SAN ASSESSMENT STRATEGY

3D SAN AUDIT—THE MOST EFFECTIVE METHODOLOGY TO ASSESS YOUR INFRASTRUCTURE

Vasily Pantyukhin, EMC

Proactive administration is the key to effective storage management. It provides many advantages with respect to reducing OPEX and saving time. The main requirement for proactive management is to understand what happens in your storage infrastructure. It isn't just monitoring of HW and SW current state. Rather, it is a regular, deep analysis of all aspects, which helps find hidden or potential problems and predict an evolution of infrastructure as a whole.

One of the most important subsystems of a storage complex is a storage area network (SAN). That's why an accurate SAN audit must be a regular, essential storage administration procedure.

From actual assessments in finance, telco, and transport companies, a very effective methodology is offered to examine a SAN infrastructure. Employing a multidimensional approach helps facilitate a detailed investigation of the most important factors:

- Architecture
- Physical state
- Fault-tolerance
- Configuration
- Management
- Performance
- Operations
- Security

The methodology described in this Knowledge Sharing article allows one to visualize maturity levels of SAN management areas, and develop recommendations for their optimization.

This article is particularly useful for architects and administrators responsible for SAN development and management.

BEST OF DATA PROTECTION

DATA PROTECTION IN A NUTSHELL—A TIERED APPROACH

Luiz Pissinatti, EMC

There are two main reasons people do backups: first, to enable recovery from failures of any kind with minimum data loss, regardless of the nature of the problem; second, to provide data for auditing or compliance reasons—whether from a couple of months ago or even years later.

These two objectives have very different characteristics and requirements, and are often underestimated by some CIOs who try to solve both of them with a single solution. The single-solution approach might have worked well so far, but with the huge data growth predicted by marketing research companies (such as Gartner and IDC), it might not be enough to meet increasingly challenging service levels demanded by new business applications.

This Knowledge Sharing article describes how to build a whiteboard discussion explaining backup mechanics that show the evolution of traditional backups and their challenges, as well as what new technologies exist today that can also be used.

It will explain how multiple solutions such as replication, Continuous Data Protection (CDP), backup-to-disk, and data deduplication can be combined to augment traditional backup strategies, delivering a tiered approach to data protection that is capable of meeting multiple SLAs within the enterprise.

This whiteboard works as a customer knowledge-level settlement tool that enables you to assertively position multiple EMC solutions such as NetWorker®, Avamar®, and Data Domain®, along with DPA, RecoverPoint, Replication Manager, and storage products, and also consulting and implementation services.

BEST OF VDI

ADDRESSING THE VDI STORAGE BOTTLENECK

Vinayak Sivanand, Tata Consulting Services Limited

Many organizations have deployed Virtual Desktop Infrastructure (VDI) with a conventional storage infrastructure only to find that the cost of scaling storage for user desktops is cost-prohibitive, and end-users are reluctant to accept the slower desktop experience. Earlier operating systems were designed to operate with low-latency dedicated local disks for every desktop. However, in a VDI environment, Microsoft® Windows® operating system tasks place a heavy load on storage as each user, application, and desktop compete for limited input/output capacity (measured in input/output per second—IOPS). Without adequate storage IOPS, virtual desktops take longer to boot, applications respond slowly, and leave users frustrated.

Organizations that try to address the problem by adding more hardware, networking, or storage, find that VDI quickly uses up the additional capacity without solving the problem. In order for VDI to scale to thousands of users, IT organizations need to take a new approach to storage for VDI that dramatically increases performance, lowers storage costs, and enables centralized desktop image management.

ARCHIVE

SYNERGISTIC OPPORTUNITY BETWEEN DISKXTENDER AND POWERPATH

Vinay Parthasarathy, EMC

By combining two time-tested technologies that operate in the path management and archival space, we can build great synergies and prove yet again that the value of combined products is more than the sum of individual values. Here we see how we can leverage two individual technologies: EMC DiskXtender® and EMC PowerPath®. Before we start, it is best to gain a quick understanding of each of these solutions and set the stage for synergies.

DiskXtender extends the primary volume by transparently archiving files that do not necessarily have the highest business value as of a particular time. It does this via policies created by the administrator called move rules, retain rules, purge rules, etc. This sets the stage for Information Lifecycle Management (ILM), a buzzword that we are familiar with. This solution can also ensure proper file deletion after a certain period or retain files for a preset time. DiskXtender operates as a file system filter driver which builds an index of files on the primary volume, and schedules moves and purging based on the policies. Operating as a file system filter, DiskXtender has good knowledge of the typical applications that create file traffic. This file or application knowledge is of particular interest to us in this Knowledge Sharing article.

PowerPath is time-tested industry-standard multipathing software that can reside on almost any host that is connected to a storage array, typically an EMC array. PowerPath sees every I/O that gets written and read from the devices it manages. It performs load balancing of I/O between several paths to the device, facilitates automatic failover if a path fails, and provides high application availability.

One highly desirable feature is “Application Performance Monitoring,” and for this to occur, PowerPath needs application awareness. We can surely leverage this from DiskXtender, which contains knowledge about the applications pumping I/O.

Now comes the synergy part. What does this application knowledge do for PowerPath? It enables us to facilitate some helpful functionality. We could perhaps build in “analytics” capabilities so that PowerPath can understand the I/O pattern in a particular environment, and be able to distinguish “Business-Critical Applications” against those that are “Routine” such as antivirus and scheduled applications like backups. We could also add a new feature called “Bandwidth Tiering” wherein an administrator can allocate a percentage of I/O bandwidth to Business-Critical Applications.

What does DiskXtender get in return? The fact that PowerPath talks to the arrays can be leveraged by DiskXtender by providing some kind of Fully Automated Storage Tiering (FAST) at the software level. DiskXtender has knowledge of those applications that are business critical and those that need high performance. This data is typically used to configure gold and silver storage policies on a FAST array. With this knowledge, we can directly translate the archival policies to program and use FAST technology on newer generation arrays without additional administrative burden. Other interesting use cases involve PowerPath capabilities such as “Migration Enabler” and “Data @ Rest Encryption” which could be explored while archiving.

Both of the additional features that this article mentions are significant value additions and innovative in approach. They could potentially be distinguishing features against the commodity multipathing software from OS vendors for PowerPath, as well as a great value addition in some kind of “End-to-End FAST” for DiskXtender. Proof-of-Concept can be possible with the great depth of engineering talent that exists in both product groups.

BACKUP AND RECOVERY

A SYSTEM ADMINISTRATOR’S APPROACH TO MONITORING AVAMAR

Kelly Mills, EMC

Do you lie awake at night wondering if your mission-critical servers are being backed up? Do you have service level agreements tied to the uptime of your backup environment and the ability to perform timely restores? Are costs preventing you from building the monitoring and alerting system you desire? Do you need to monitor a single or even multiple Avamar systems spread across multiple sites?

Depending on your role in the backup service, it could be of utmost importance to know when your backup infrastructure is down or experiencing problems. It is also important that those problems are identified and alerts generated in a timely manner, with no pause until acknowledged. Realtime monitoring provides a proactive approach that detects faults and allows quicker resolution to maintain uptime. The alerts need to be available in many forms to reach the intended audience. Avamar provides a host of options for alerting the user when the application has issues, but what do you do when the application is down or incapable of providing those alerts? The system needs to take into account the Avamar application itself but also be aware of the environment in which it is running.

A monitoring and alerting system must be robust and, at the same time, easy to manage and configure. The system must not be so cumbersome that it becomes a burden to use. It should provide the basic functions for adding new nodes to monitor, processes to check, and people to contact when things go wrong. The system should facilitate as much automation as possible. It should be centralized, providing a secure single place for modifications and access. You might also want to know things about your Avamar environment that are not necessarily urgent; for

example, current disk utilization. A central monitoring system can provide this functionality in a single location for a single point of view.

The aim of this Knowledge Sharing article is to offer System Administrator and backup operations expertise to provide ideas for addressing the issues mentioned in this abstract. It will lay out a foundation for monitoring Avamar that includes points to consider and real-world experience. The article will describe some of the built-in monitoring functions of Avamar and how to use those to your advantage. Also included is a sample setup to get you up and running in minimal time, with little or no cost, using industry trusted tools.

NETWORKER BEST PRACTICES WITH AUTOMATION

Ajit Kumar Rath, EMC

In the world of EMC NetWorker, it becomes more challenging and difficult for a NetWorker Administrator to handle a large enterprise backup environment where the number of NetWorker servers, storage nodes, and clients increases day by day. This requires a NetWorker Administrator's proper planning, designing, and implementing efforts while configuring. It becomes increasingly difficult for NetWorker Administrators to manage such huge environments manually using the existing NetWorker features. It may also require additional NetWorker Administrators to manage huge enterprise backup environments.

This Knowledge Sharing article describes how we can use Automation with NetWorker to reduce manual efforts and save the time of NetWorker Administrators, and apply best practices to manage such huge enterprise backup environments in a much more organized way. This will also help in managing such environments with fewer NetWorker Administrators.

Take the example of an environment consisting of different data zones having more than 50 NetWorker servers, hundreds of storage nodes, and thousands of clients. Though NetWorker has features for monitoring and configuring schedules, groups, clients, and media management, it becomes more difficult for NetWorker Administrators to manually configure these features in such an environment.

Configuring groups, clients, and schedules for more than 10 servers and 50 clients takes time and effort if done manually using NMC or nwadmin. Similarly, ejecting, inserting, and recycling more than 50 tapes takes time and manual effort.

Using Automation with the help of Perl scripting in such a backup environment not only saves time, effort, and costs but also improves scalability. NetWorker Administrators who are required to configure and monitor such enterprise backup environments with best practices are the target audience for this article.

The following best practices are discussed:

- Commissioning and Decommissioning Backup Clients using Automation
- Media Management using Automation
- Groups and Schedules Management using Automation
- Client Failures Troubleshooting using Automation
- NetWorker Server Maintenance using Automation

This Knowledge Sharing article will empower resident engineers to better leverage EMC NetWorker service and support.

HOW TO START YOUR DISASTER RECOVERY IN THIS “CLOUDY” LANDSCAPE

Roy Mikes, Mondriaan Zorggroep

Let's start with a simple quote, “Information is the organization's most important asset.”

Data is created by the applications and is processed to become information. Information is undoubtedly the most important asset for an organization. Does this make sense? Absolutely. The digital footprint for each person on this planet is growing. It does not matter whether we as individuals or as a corporation store data, it has to be protected. For some, photos are just as important as a company's ERP system. Understanding this, storage vendors put a lot of energy into managing this information.

From a disaster recovery (DR) perspective, the world is divided into two types of businesses; those that have DR plans and those that don't. Of these two types, if a disaster strikes, which do you think will survive? When disaster strikes, organizations without DR plans have an extremely difficult road ahead. If the business has highly time-sensitive critical business processes, that business is almost certain to fail. If a disaster hits an organization without a DR plan, that organization has very little chance of recovery—and, at that point, it's certainly too late to begin planning. Even organizations that do have DR plans may still have a difficult time when a disaster strikes. They may have to put in considerable effort to recover time-sensitive critical business functions. However, those with a DR plan have a fighting chance at survival.

Does your organization have a disaster recovery plan today? If not, how many critical, time-sensitive business processes does your organization have? Many organizations think they have a DR plan. They think they have some procedures, and that is all it takes. True, you need procedures, but you also need to be sure that you actually can failover. How do you manage that? Testing live may do more damage than knowing you can. Are you certain of every organization change? Many organizational infrastructures change per hour. Try to fit in your DR plan when changing that fast. Where does that leave you? Good question. Probably, when you test your failover, you do it once per year, maybe twice, or even each quarter. How much do you think has changed since the last time you performed your failover. Thus, this is a considerable challenge.

Luckily, there are many techniques and solutions, such as “clouds,” where DR plans are probably already well-organized, or VMware® Site Recovery Manager (SRM), which can help you with your failover. VMware SRM is a business continuity and disaster recovery solution that helps you plan, test, and execute a scheduled migration or emergency failover of data center services from one site to another. However, the most beautiful part of SRM is that you can test a plan without doing it live. Does this mean that you can failover anytime without doing damage to the infrastructure environment? Yes. Virtualization and cloud computing (both public and private) can make DR implementations easier. However, there is a lot to consider before taking this road.

At minimum, consider the following items as you set up a disaster recovery program:

- How does it correspond to your business continuity?
- What is the maximum downtime you can accept? What is your Recovery Time Objective (RTO)?
- How much data loss can you accept? What is your Recovery Plan Objective (RPO)?

- What kind of recovery plan do you need for a partial disaster, compared to a complete disaster?
- How do you classify which data/applications are mission- or business-critical or part of the DR?
- How do you classify your computing resources?

Depending on the nature of your business, good disaster recovery is achieved by designing a process which enables your operations to continue to work, perhaps from a different location; with different equipment; or from home; making full use of technology to achieve a near seamless transition that is all but invisible to your customers and suppliers. Insurance can mitigate the cost of recovery, but without a disaster recovery plan that gets you back up and running, your business could still be at risk. Indeed, more than 70 percent of businesses that don't have a DR plan will fail within two years of suffering a disaster.

NETWORKER DATA MIGRATION—EDL TO PHYSICAL TAPE

Gurpreet Dhillon, EMC

Performance is vital to business applications. We have often seen Fortune 500 companies struggling with performance to get their data ready for offsite storage. There could be a number of reasons for this: budget cuts, the lack of a next-generation backup and recovery solution, an aging backup and recovery infrastructure, poor capacity/storage planning, and even the inability to put a cap on an ever growing demand for data backups.

This Knowledge Sharing article provides backup administrators, professional services, and support services for EMC Disk Library (EDL) and NetWorker environments with guidelines on how to migrate data from the EDL to a physical tape library, transparent to the NetWorker server. This article will benefit EMC customers who are interested in migrating data offsite from EDL to a physical tape library for offsite storage—without NetWorker performance being impacted in the whole process. Most importantly, this process provides NetWorker with data availability even when NetWorker is NOT used for this data migration.

This Knowledge Sharing article will help individuals leverage EDL features in NetWorker-based backup and recovery environments where coherent working of multiple products in an architected solution is imperative. By taking advantage of multiple product functionality in a backup and recovery infrastructure, we provide a solution that is transparent to NetWorker for data migration.

CELERRA AND CLARIION

CELERRA SNAPSURE BEST PRACTICE INSIGHT OF ITS SPACE MANAGEMENT AND PERFORMANCE

Frank Zhao, EMC

Kevin Xu, EMC

EMC Celerra® SnapSure™, a pointer-based data protection feature in Celerra, enables users to create and manage checkpoints for data protection/business continuity. Due to its fast creation, prompt recovery, and economic space usage, SnapSure acts as the fundamental and enabling disaster recovery (DR) technique for Celerra Replicator™. It is also integrated with Celerra NDMP backup to assure backup data consistency.

Checkpoint and checkpoint-based replication are widely used by Celerra customers. However, their “copy on first modify” principle and the need to reside on a separate

volume (a.k.a., SavVol) bring additional challenges to performance tuning and space management—two important aspects for a data protection solution. A typical scenario is that there are dozens of user checkpoints, some Replicator internal checkpoints, and some NDMP automatic checkpoints residing on a single SavVol. As a result, an appropriate design for storage space usage and scheme become critical to satisfy both local data protection and remote replication requirements.

On the other hand, SnapSure is not an independent DR solution. Instead it is derived and always bound with PFS, which makes it difficult to guarantee consistent performance. If performance is not well-tuned, it would cause negative impact not only on checkpoint but also on PFS, Celerra Replicator, and backup, and may eventually impair the service level agreement (SLA). Actually, space management and performance tuning are pervasive throughout the solution design, implementation, operation, and troubleshooting stages.

This Knowledge Sharing article provides insight into SnapSure space management and performance rather than a general introduction of SnapSure. In particular, for the solution design stage, this article explains how to distinguish data access patterns and identify data protection requirements, as well as work out the storage configuration and evaluate its performance changing curve during business operations. For the implementation stage, this article explains how to construct corresponding tests to verify the solution through some helpful utilities. In addition, this article provides some guidelines and tips for performance tuning and analysis of the impact on PFS, Celerra Replicator, and NDMP.

The target audience of this Knowledge Sharing article includes Celerra solution consultants, implementation engineers, system administrators, and technical support engineers.

INTEGRATING AND TROUBLESHOOTING VMWARE WITH EMC CLARIIION: DEEP DIVE

Jason L. Gates, Presidio Networked Solutions

Many in the EMC Proven Professional community have questions about best practices and the advanced skills required to troubleshoot problems at the Systems Engineer level when it involves VMware and CLARiiON in one solution/configuration. Problems do arise and resolving these problems requires an engineer and/or administrator with a good understanding of the VMware hosts and expected behavior of the CLARiiON at the SCSI and Fibre Channel layers.

What kinds of problems can arise? Examples include: LUN trespassing, slow performance of VMFS, loss of connectivity, iSCSI configuration, and SCSI reservation issues. What happens when these type of problems occur? The storage administrator points fingers at the VMware server farm and the VMware engineer points fingers at the storage subsystem.

As a VMware Certified Professional (VCP) and CLARiiON Storage Admin expert, I have implemented and troubleshot numerous VMware ESX® and CLARiiON configurations. Some common customer questions addressed include: “What is the best path policy?” “What invokes a trespass?” “How does VMware native multipath behave?” “What causes SCSI Reservations and what is the best way to resolve them?” “What are the best practices to configure iSCSI with the CLARiiON?” “How should my LUNs be carved out based on Virtual Machine needs?” Also, I look to provide detailed information on what logs and support files can be used to assist from both VMware and CLARiiON environments.

This Knowledge Sharing article will take an in-depth look into the questions mentioned above. It is offered to shed some technical light on the kinds of situations that many in our industry encounter on CLARiiON—the world’s number one mid-range storage subsystem—and VMware’s leading virtualization software.

CELERRA DISK RECONFIGURATION AND ADVANCED MANAGEMENT

James Ingle, Alexander Open Systems

When it comes to connectivity and data management, the EMC Celerra Unified Storage platform is a very flexible storage solution. Due to the tight integration of EMC Celerra with EMC CLARiiON, storage administrators are able to provide a scalable storage environment for SMB to enterprise environments. While the setup and basic administration of the EMC Celerra is quick and simple, reconfiguring back-end storage and following best practices takes some consideration.

It is extremely easy to add disk and expand storage with the Celerra, but what happens when you change directions with your storage design and wish to reconfigure the CLARiiON storage behind the Celerra datamovers? There is a specific process that must be followed using only the command line interface (CLI). There are caveats that must be considered when deleting the underlying volumes, disks, LUNs, and RAID groups that make up the Celerra storage.

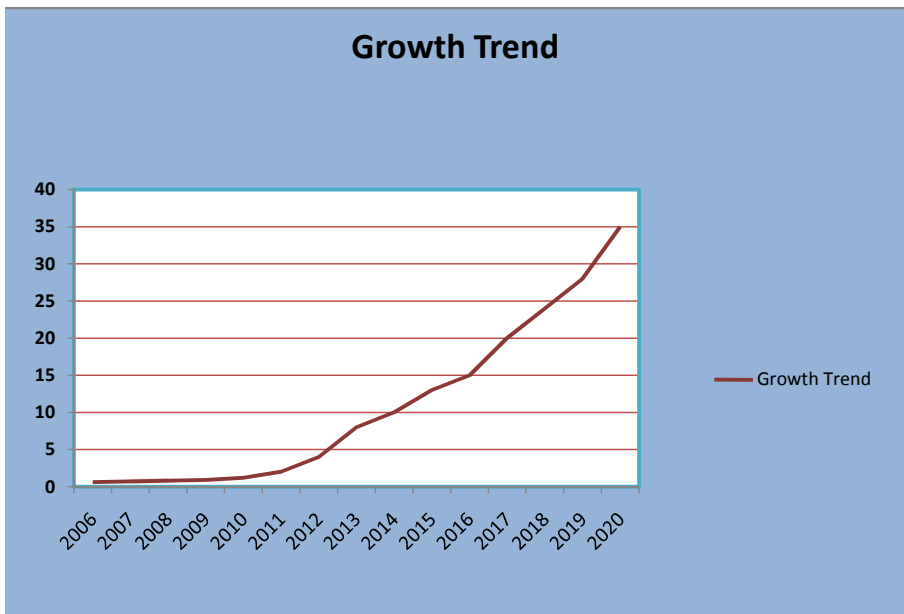
I have been working with EMC Celerra for many years now, and have performed this task several times. The process is straightforward once you understand it, but until now it took several different documents and white papers to find the necessary commands and tasks that had to be performed. In this Knowledge Sharing article, I break down the process and list the steps within an easy to follow methodology. Having this information in a single document will enable storage administrators to confidently perform this task on their own.

CLOUD

DEADLY TRIO INTEGRATION

Anuj Sharma, EMC

While the global economy went through a recession in 2008-2009, there was no recession in data growth. According to IDC’s Digital Universe study, it was estimated that the worldwide volume of digital data grew by 62 percent between 2008 and 2009 to nearly 800,000 petabytes (PB). Further, IDC claims that this “Digital Universe” will grow to 1.2 million PB, or 1.2 zettabytes (ZB) in 2010 and reach 35 ZB by 2020. Other results from the 2010 study include the amount of containers that will hold all of the digital data, and the changing sizes of the data files. The study estimates there will be 25 quintillion (that’s 25,000,000,000,000,000) information containers—packets, files, images, records, and signals—that will hold the Digital Universe data by 2020. Additionally, the files are getting smaller with the growth of embedded systems, including smart grids, smart cities, and logistic item tracking.



X axis: Years **Y axis:** Zettabytes

This clearly indicates that we will be dealing with huge amounts of critical data in the times ahead. The word ‘critical’ itself demands that business continuity procedures be in place in case of data loss. As we all are aware, many companies ceased to exist after 9/11 as all the data was lost, and there were no business continuity practices in place. This reinforces the fact that data is the most vital asset for any enterprise—whether it’s a FMCG, IT, banking, or business in any other sector. Data loss for any organization can lead to huge financial losses as well as loss of reputation in the market. As soon as an enterprise goes live into production stage, it is critical to have business continuity procedures already in place.

As data grows daily at a rapid pace, so too do backup challenges. Per the IDC study, of the 35 ZB of data expected by 2020, nearly 75 percent will be copies of original data. This represents an opportunity for cost reductions using advanced data services such as compression and data deduplication.

This Knowledge Sharing article will focus on the integration of NetWorker, Avamar, and Data Domain for designing, implementing, and administering a robust deduplication-enabled backup infrastructure, and will benefit anyone who is involved in designing, implementing, and administering backup infrastructure. This article will also touch upon what makes NetWorker, Avamar, and Data Domain stand tall among the competitors, and how customers can benefit from the features, achieve high return on investments (ROI), and lower total cost of ownership (TCO). The following aspects will be covered in the article:

- Factors to be Weighed while Designing the Backup Solution
- NetWorker Integration with Avamar
- NetWorker Integration with Data Domain with the DDBoost Software Component
- NetWorker Integration with Atmos™ Cloud
- NetWorker Integration with VMware
- Positioning of Data Domain and Avamar in the Backup Infrastructure
- Implementation of the Disaster Recovery Infrastructure
- NetWorker Implementation and Administration Best Practices
- Avamar Implementation and Administration Best Practices
- Data Domain Implementation and Administration Best Practices
- Disaster Recovery Procedures
- Other Practices That Will Help Deploy and Administer a Robust Backup Infrastructure

WHY CLOUD COMPUTING?

Bashanta Phukon, EMC

Wikipedia says “Cloud computing is Internet-based computing, whereby shared resources, software, and information are provided to computers and other devices on demand, like the electricity grid.” So what does this mean? Most companies will provide users with software that resides on their servers in the so called “cloud” or the “grid.” Users can access the software or information from anywhere and on any computer attached to the Internet. In other words, software would cease to be just a tangible product that is installed on the computer. Software is evolving into a service that can be accessed over the Internet. Examples of cloud computing can be seen with services such as Gmail, Google Docs, Office Live, and other online platforms.

Cloud computing is not a technology; rather, it is a way in which we can use technology to decrease IT overhead (cost-wise, in theory). Cloud computing is essentially letting someone else (outsourcing) provide the hardware and infrastructure required to run aspects of an IT setup (infrastructure can be connected over the Internet). Simply put, cloud computing is a business model that companies can adopt.

Frequently, “virtualization” is sometimes confused with the cloud. However, virtualization isn’t something used only by the cloud; rather, it is a technology that many organizations can leverage to maximize IT hardware utilization. It therefore comes with a host of benefits for an organization, some of which include: reduced administration and hardware costs, and lower electricity bills. It can get confusing and downright cloudy when talking about technology and concepts, and how to implement technology. However, virtualization and cloud computing are very different things and making a clear distinction between the two makes it much easier to talk to the business about the benefits of a “concept” or a technology.

How can cloud computing benefit the enterprise? The cloud would be an advantage for companies that have no prior investments in infrastructure, and have a steady cash flow. With debt financing for startups almost disappearing, a utility model is an attractive option. For large organizations with existing infrastructure, there is an up-front expense in laying out an infrastructure to support a model that allows elastic deployment of applications. Virtualization, automation, application rewrites, platform consolidation, and evangelizing will take up a lot of time and money before any cost benefit can be seen. In these scenarios, “Cost benefit can be taken as a fruit to be savored later.” By leveraging the cloud, new ideas can come to market earlier since the infrastructure needed to validate the viability of these ideas is available with a few mouse clicks.

Over the past couple of years software providers have been moving more and more applications to the “cloud.” It is the wave of the future in terms of delivering software as a service. Indeed, the case for cloud computing is quite appealing. Many companies are looking for alternatives for employees, such as a work from home option, thereby increasing the number of employees making up a mobile workforce.

This shift toward cloud computing will enable businesses to save money while minimizing their impact on the environment. The reduction in application costs and technical support will easily accommodate the trend toward a mobile workforce that will utilize the computer grid. Everyone will benefit from the increased availability and affordability of applications that were previously beyond reach due to cost, installation complexity, and maintenance. At the end of the day, the successful operational model is one that will enable the business to benefit from cloud computing solutions.

RIDING THE CLOUDS—BEST PRACTICES IN DATA CENTER TRANSFORMATION INTO THE NEXT CENTURY

Paul Brant, EMC

There is a digital crisis. The digital information universe is growing at an exponential pace. Some consider it a never-ending deluge. How do you find information? How do you handle unstructured data? How do you add structure? What are containers and how does this concept address these issues? How do you handle object affinity, manage distribution of assets, and delve into the storage ecosystem? How do we tackle all of these challenges, and what will the next-generation data center look like? This digital crisis requires that we re-think established truths and transform business and technical processes as soon as possible.

Faced with retaining and preserving huge amounts of digital information for very long periods, both large and small organizations are at the front edge of a troubling crisis. Digital information is actually easier to lose than its paper or film equivalent. It is one thing to manage a domain of digital records that an archivist can personally guard and shepherd, but it is quite another to meet the archival challenges of today's enterprise data center.

These data centers can be characterized as environments with petabytes of distributed information and high data growth rates. They encompass many facilities and departments with uncoordinated responsibilities and requirements, as well as a lack of business-level budgets and focus on archives and recently created information. All of these operating challenges are compounded by high risk—risk of failure and fines from legal discovery, compliance requirements, or security threats. Add the risk of losing information that may be of great value to the organization and the picture looks daunting.

The digital crisis will be exacerbated by time. In 10 years, 50 years, 200 years, which applications will still be around? What computer and storage system will be able to read old information, provided that it is not corrupted by then? Finding a single piece of content, and all the linked objects that contain associated content, amid trillions of distributed information objects could become a costly and complex adventure.

The problems are already huge and here is the dilemma. Many standards and best practices exist today that document managing, creating, utilizing, and preserving digital information. Yet, none of them address the core problems caused by inadequacies and inefficiencies in the supporting information storage infrastructure.

There is good news though. With the advent of new business, process, and technology practices, there is hope. The private and public cloud is a burgeoning approach foreshadowing what the next-generation data center will look like. Virtualization, compliance, security, long-term retention, geographic transparency, and global reach of computing resources, as well as many more aspects to this challenge will be discussed in an effort to find a way to grasp these issues and solve them.

Today's data centers and IT infrastructures face unprecedented challenges. Many are experiencing a capacity crisis as they reach the limits of older facilities and legacy, siloed infrastructures. Space is tight. Technology modernization is overdue. Energy costs are high. As a result, many companies still spend up to 60 to 70 percent of their IT budgets on operations and maintenance, instead of innovation.

Meanwhile, customers are demanding wider access to information, transactions, and services. Pressure may be growing to evolve your IT organization from cost center to strategic business enabler, providing a true service-based infrastructure. Moreover, you must address all of this in the face of an uncertain near term, not to mention an extended future.

What are the best practices to thrive despite unpredictability and exponential information growth? With a more elastic, resilient, and optimized next-generation data center built on the attributes discussed, you will be ready for whatever comes next.

This Knowledge Sharing article will describe how to ride the “cloud” and utilize what this technology can afford, as well as offer best practices that will align with the most important goal; creating a next-generation data center that addresses the business challenges of today and tomorrow through business and technology transformation.

OPTIMIZING IT SERVICES USING CLOUD, DATA CENTER, AND ADHERENCE TO COMPLIANCE STANDARDS

Amrith Raj Radhakrishnan, Tata Consultancy Services Limited

Today, data drives the world. The success of business depends on keeping data secure and accessible so that it can be put to work as business information. In the last few years, the IT industry has seen a lot of changes. New technologies are blooming like never before. The most recent, cloud computing, is now a reality, and evolutions in various hardware, software platforms, and the data center have been seen in the IT industry. Businesses are focusing on excelling in their core business rather than focusing on IT. This is only possible if IT is agile, robust, and reliable.

For businesses, IT shouldn't be a matter of concern but rather, a solution to its problems. However, finding a solution to a major IT problem is a big challenge. New technologies and options that have emerged recently contribute to increasing complexity and confusion. Multiple options and solutions make it difficult to determine what to select and what not to. Companies are unclear what to choose to enable IT to transform their businesses. Should a business build, manage, and support its own IT infrastructure? Should it host its hardware resources with a data center provider in the form of co-location? Should it use resources from a cloud service provider or even outsource IT and services to a third-party firm? The confusion increases as we apply each solution.

Each solution has its advantages and disadvantages. A good feature of a solution might be excellent for one type of business, but might turn out to be inappropriate for another. With the features cloud computing provides, every company will be keen to harness its power. For example, it is widely known that data centers are one of the highest consumers of electricity. The need to reduce carbon emissions is also rising due to ecological compliance. New solutions are needed not just to comply with the policies and regulations but also to reduce the rising electricity bills of the data center.

So, how do we decide what is good for an enterprise with regard to running the business effectively? How do we choose an IT solution to solve business problems without causing a chain-reaction of complex issues after implementing a wrong one? Will adding more servers and storage solve the problem? How do we govern our IT infrastructure? How do we apply existing service management tools? How do we transform an existing IT infrastructure which is currently suffering from problems into a long lasting, strong, robust setup? How do we get insight into new technologies that can transform IT for the business? These questions and more will

be answered in this Knowledge Sharing article. Included will be a multi-dimensional view of current IT problems and solutions.

In summary, the article will explain how IT can deliver more services using cloud computing, the data center, servers, I/O virtualization, converged infrastructure, enterprise storage, proper governance and management, and more. The article offers a view from the business with respect to IT, and provides a detailed analysis explaining applications across various types of enterprises. The objective is that IT should deliver the best to a business by leveraging the cloud, the data center, and compliance.

CONTENT MANAGEMENT

A GUIDE TO ENTERPRISE CONTENT MANAGEMENT

Mouli Ponnupandy, Keane, an NTT DATA Company

Art is beautiful only when it reaches its audience. Likewise, technology is useful only when it achieves intended business benefits. Customers making technology investments often need to rely on industry reports, competitive analysis, vendor capability, and product datasheets before arriving at an informed choice of product or technology. These data sources are full of buzzwords and industry jargon which customers need to cut through to really understand what they want to buy, and if the investment choice is good.

This Knowledge Sharing article intends to present enterprise content management from an end user perspective, and aims to educate the reader on constituents of content management. It offers insight into document capture, document management, digital asset management, web content management, web experience management, enterprise collaboration, records management, document output management, business process management, and enterprise search. It explores how features of different functional components overlap and interconnect. This helps customers identify what functional components are essential to meet the needs of content management in their organization. Business and technology experts like us need to understand enterprise content management from a buyer's perspective to effectively service customer needs.

As a consultant, when I meet with customers, the first question they seem to need an answer to is, "What are data, metadata, and content?" Computers are designed to process elemental data. Humans are inefficient at processing data but adept in processing content. To make it simple, computers are analytical—breaking things into parts to process; humans are contextual—putting things together to understand the big picture. Both speak different languages. So for us to communicate effectively, we need a common language—metadata—which in basic terms is data about data, or let's say, context to the content. Structured metadata is stored systematically in a database table which interacts with unstructured content according to business rules. Computers manage metadata while effectively delivering content to us in a usable format.

Content Management aims to manage content through the entire lifecycle—from creation, to delivery, and finally, disposition by associating it with appropriate metadata. The primary purpose of an enterprise Content Management System (CMS) is to Consume, Process, and Deliver content in a usable form, and to manage any integration with existing business applications.

Consume—There are three specific ways in which a user can bring electronic content into the CMS. Create: the most direct and probably the easiest way is to create the electronic content from within the CMS. Capture: there are certain business

processes which require handling of paper-based content objects. The paper-based documents are scanned and introduced into the CMS as electronic content. Import: a good amount of content files reside in various legacy and business applications as silos. The electronic content in these disparate systems are imported or migrated into the CMS.

Process—Content by itself does not add any value to end users. It needs to be processed in some form for it to be fit for use. **Classify**: the preliminary processing adds method to the madness by way of classification. Content consumed is organized in the form that best suits business needs. **Collaborate**: people from across boundaries need to share business content to get the job done. Collaboration facilitates processing of content per the desired business outcome. **Manage**: the system should manage both content and people who access content. Authentication takes care of data security and authorization takes care of people security.

Deliver—The content is processed in a neutral environment so that multiple output formats and channels are supported by single source content. The idea is to have single source content. **Personalization**: content is then delivered either in an intermediate format which allows the end user to personalize presentation, or delivery channel of the content based on conditional processing (metadata). The idea is to reuse, recycle, and repurpose content as there are multiple users and delivery channels for enterprise content. **Compliance**: content (most, not all) becomes a record when it reaches the end of its active lifecycle, and it needs to be preserved either due to internal policies or external regulations. An ideal compliance initiative should have provision for both retention and archival. Retention aims at holding on to a record for a fixed period of time and then ensuring that it is purged from the system, effectively bringing record life to its end. Archival aims to move non-essential records to low-cost storage devices, either automatically based on business rules or manually driven by administrative needs.

GLOBAL SCALE DOCUMENTUM IMPLEMENTATION—DESIGN CONSIDERATIONS

Sameer Patil, Tata Consultancy Services Limited

Companies with a world-wide presence face specific challenges when implementing applications that must cater to the requirements of all colleagues globally. Document management systems are no exception. This Knowledge Sharing article describes a solution that can be used by architects while addressing global requirements, and explains how the EMC Documentum product mix can be used to meet these requirements.

As a solution architect, I have worked on global implementation of a Documentum system for an insurance brokerage firm. This article defines the challenges faced while designing repository, application integration, and migration to meet global requirements.

Some of the challenges faced while implementing Documentum are summarized below:

- Implementing country-specific requirements for security, taxonomy, and object model
- Meeting local standards for network, database, LDAP, and coding
- Implementing country-specific policies with respect to storage and retention periods
- Addressing different language requirements
- Implementing country-specific functional requirements

- Facilitating compatibility with local applications interacting with the document management system
- Migrating data from existing local repositories

To overcome these challenges, we began by classifying essential out-of-the-box feature, configuration, and customization requirements.

For requirements needing customization, we tried to define the configuration part which made it reusable across different geographies. We helped the client rationalize the object model in order to meet country-specific requirements, using different solution mixes based on the given requirements. We also educated the client about the extent to which language requirements could be met.

While implementing the application integration requirement, we helped the client map to single sign-on in the application. We minimized the impact of existing integration by providing a common framework with an extension of application-specific code for rationalization of the metadata, protocol, and application-specific functionality.

During migration from different repositories, we engaged with the client on:

- Mapping metadata from multiple systems to the new system
- Identifying error cases long before the production move
- Meeting stringent timelines while migrating a massive amount of data

SAN

EMC SAN COPY MIGRATION—A COMPLETE EXPERIENCE

Dinesh Pahari, Dimension Data Australia Pty Ltd.

EMC has been a market leader in developing tools that simplify the SAN administrator's work life. Due to the dynamic nature of IT environments, it is imperative that businesses evolve and adapt to proven new technologies in the market. This not only helps businesses fully utilize their IT investment but also enables them to become leaders in a very competitive market.

From medium-size businesses to large enterprises, storage needs nearly double every year, if not more. Therefore, businesses have a big decision to make when they invest in their SAN environment. So, what do the businesses do? This is a big question, and EMC has been very successful over the last decade in providing the best answers with many proven solutions. One of those solutions is to facilitate migration of legacy storage onto a new storage system—without the hassle of losing data or extensive downtime to the production environment.

EMC SAN Copy™ is among the many solutions that EMC has put forward for SAN migration. SAN Copy has been in the market for many years and, therefore, is a proven product for high-speed data mobility, migration, and protection between EMC CLARiiON networked storage and other storage systems that are qualified by EMC. These systems include EMC CLARiiON, EMC Symmetrix®, IBM, HDS, and Sun storage arrays.

This Knowledge Sharing article focuses on the process and procedure involved in SAN Copy migration between two CLARiiON arrays. This is based on the process I followed that led to a successful SAN Copy migration of multiple hosts from a CLARiiON array CX400 to a CX4-480.

Below are the basic steps for migration which are discussed in this article.

- Basic requirements and environmental updates (EMC PowerPath®; OS, etc.)
- Creating SAN Copy sessions
- Host preparation
- Starting SAN Copy sessions
- Monitoring SAN Copy sessions
- Completing the migration
- Troubleshooting

HOST INTEGRATION WITH EMC STORAGE ARRAYS

Upanshu Singhal, EMC

Shelesh Chopra, EMC

“Do you have any documents which can help me configure my host with an EMC array?” The usual answers given to this frequently asked question include, “Check Powerlink,” “Register for EMC array training”, or “I will help you.” The same question is asked every time a new engineer joins the team, as they typically need to develop or test their software with EMC storage arrays, or need some larger storage space. Even if someone finds a document to configure, for instance, Microsoft Windows, the next request will be about Solaris or other operating systems, which bring their own nuances. Perhaps you have configured an array with a physical host, and now you have a couple of virtual machines which need EMC array configuration. The engineer ends up going through multiple trainings or piles of lengthy documents, or other members spend significant time to help them. Many times engineers do not have time to go through the complete documents, instead opting to browse the text to help them quickly satisfy their needs.

We have used EMC arrays for our software development for the past 8 years, have gone through several training sessions, and have hands-on experience with EMC CLARiiON, EMC Symmetrix, multiple operating systems, and physical to virtual machines.

Providing a single document with an introduction to storage area network (SAN) components, and a step-by-step procedure on host integration, will help bring engineers (development, QA, and support) up to speed quickly and enable them to be productive in a timely fashion.

This Knowledge Sharing article will provide introductions to:

- SAN hardware components
- SAN software components
- Configuring third-party hardware (e.g., Host Bus Adapter “HBA” and its software)
- Different types of EMC arrays and their use in a customer environment
- A step-by-step guide on host integration using EMC array GUI and command line interfaces used for storage provisioning
- Configuring EMC array replication technologies
- Configuring Windows or non-Windows hosts with EMC arrays
- Configuring physical and virtual machines or VMware ESX servers, and much more

Note: This article will not cover SAN theory, SAN topology, storage array details, or Fiber Channel switches.

Benefits of this document:

- Single point of reference to configure the host with EMC storage
- Quick turnaround time and smooth transition

- Self-learning with very minimal effort and time required from others
- The document is expandable. Any new learning will be consolidated here, no need for everybody to change/modify/update/recreate their own documents
- The document can easily be integrated into the Procedure Generator Tools and can be made available via the web

Engineering personnel are the intended audience for this Knowledge Sharing article but it can also be published for end users as a “Quick User Guide” for rapid host configuration with EMC arrays.

STORAGE

USING THE BUILDING BLOCK APPROACH FOR SCALING ENTERPRISE STORAGE

Mark Rogov, EMC

Does your company have to comply with stringent security regulations? Do you have to provide storage for both super I/O-hungry databases and slow WAN-connected Windows users? How does one achieve maximum value with a minimal price tag?

This Knowledge Sharing article explores the challenges that a large government contractor faced in its effort to optimize its storage vision. It offers insight into balancing a dwindling budget with growing capacity demands, while offering better scalability and regulations compliance. It explores issues such as secure data cleanup and modular architecture, virtualization, and the disappearing ratio between usable and raw capacity.

At the end of this article, we present a large government contractor’s solution that utilizes the building block approach to answer all internal and external goals and requirements, while using the latest cutting-edge technologies.

ENTERING YOTTABYTE ERA USING ENTERPRISE FLASH DRIVE TECHNOLOGY

Milan K Mithbaokar, Dell Inc.

Enterprise Flash Drives (EFDs) have entered the ever-expanding Digital Universe market and will soon lead us toward the yottabyte era. The yottabyte is a unit of information storage equal to one septillion (10²⁴) bytes. The unit symbol for the yottabyte is YB.

EFDs, also known as Solid State Drives (SSDs), contain no moving parts and therefore are not constrained by seek time or rotational latency. This absence of moving parts dramatically improves the ability of EFDs to deliver a significant amount of IOPS with very low response times. The EFDs can produce approximately 30 times more IOPS than traditional Fibre Channel (FC) HDDs. In addition, because there are no mechanical components, EFDs consume up to 98 percent less energy per I/O than FC HDDs.

EFD technology opens new doors for organizations that need to make decisions at speeds that were previously deemed impossible. Organizations now have the option of going with Hybrid Storage models (using SSD and FC Drives), a current trend in the storage industry.

The introduction of SSD technology and the Hybrid Storage model into the IT infrastructure business world will offer countless new avenues for success to company leaders who can harness its blinding speed and vast data analysis potential to achieve their own business goals.

This Knowledge Sharing article will explore the challenges facing SSD technology and how it will lead us to the yottabyte era.

ARCHITECTURAL PLANS FOR ENABLING A DYNAMIC DATA CENTER FOR TODAY AND IN THE FUTURE

Puneet Goyal, HCL Comnet

Data centers are at a crossroad. Massive data growth, tough economic conditions, and physical data center limitations (power, heat, and space) exert extreme pressure on IT infrastructures. Finding ways to take cost, complexity, and associated business risk out of the data center without sacrificing service levels has become a major objective for almost every enterprise. In theory, consolidation and virtualization make it possible to change the way data centers are architected, built, and managed, but how do you start your improvement journey and avoid disruptions and added costs to your ongoing operations?

In today's resource-constrained world, success calls for streamlining the inefficiencies of your storage infrastructure and processes.

Is your underlying technology prohibiting your IT organization from effectively managing costs and responding quickly to the demands of business? This Knowledge Sharing article outlines EMC architectural plans for enabling a dynamic data center for today and in the future. We will discuss:

- **EMC Unified Storage architecture enabling today's dynamic data center—** Elements of the EMC Unified Storage architecture are described in detail. This includes true multiprotocol support, a single management interface, integrated data protection, support for multiple tiers of storage (primary, secondary, and archive/compliance), quality of service, and the ability to act as a front-end for legacy storage systems.
- **Scale-out storage and the dynamic data center of the future—** Describes the requirements for scale-out storage in the enterprise in terms of resiliency, scalability, manageability, and other factors. Also discussed is the current state of EMC scale-out deployments as well as specific operational and business advantages that the EMC approach delivers to the enterprise data center.
- **Building a dynamic data center—** Discusses the steps required to begin building your next-generation data center. It describes the elements of the new dynamic data center infrastructure as well as specific cost, time-to-market, and risk mitigation advantages that the EMC approach delivers to the enterprise data center.

INFORMATION-IN-MOTION (IIM)

Lalit Mohan, Dell Global B.V

Information is always in motion. This has attracted considerable interest in the industry. Understanding the mechanics of this motion enables the enterprise to estimate the evolving value of each piece of information. This empowerment has its rewards, as each piece may be matched with the storage infrastructure component most suited in functional and financial characteristics.

This Knowledge Sharing article traces the journey of information as it traverses a typical enterprise, capturing and attempting to quantify the short- and long-term flows and functionalities that typically exist in any storage infrastructure. Even though every enterprise and its information storage footprint is unique, it may be decomposed into a number of common component categories (e.g., transactional, messaging, batch processing activities, and so forth).

Readers may benefit from this article by understanding the enterprise framework, molding it for their use, and reaping the rewards of matching each information piece with the appropriately suited storage infrastructure components.

STORAGE CAPACITY MANAGEMENT: SIMPLICITY IS EFFICIENCY

Olumide Ademidun, ING Bank Canada

For years now, I have designed some of our storage infrastructure architecture around EMC Centera® which suddenly has become an important and integral part of our storage environment as more applications and data are being migrated to the device pool. However, a problem that I immediately encountered was the limited tools for planning, monitoring, capacity management, and forecasting.

With a thin storage budget, I realized that it would be difficult to convince my management to invest in the high-end EMC ControlCenter® which would provide a solution to my problem. With this in mind, I started building and customizing my own tools using a simple Microsoft Excel® spreadsheet.

In this Knowledge Sharing article, I describe how this spreadsheet has grown to become an indispensable tool for performing one of our most important daily tasks; planning, managing, and reporting storage usage in our environment.

This simple yet powerful tool enables me to:

- Track individual node capacity, when it was upgraded, and the type of fault tolerance the node uses
- Track cluster capacity, application pool allocation, percentage usage, and when to re-order
- Proactively allocate more storage to the application pool based on total storage available
- Translate information to graphical form (useful as supporting evidence for budget requests)

STORAGE MANAGEMENT

AUTOMATING EMC IONIX CONTROLCENTER HOST AGENT DEPLOYMENT AND MAINTENANCE PROCESS

Mehmet Ozdemir, EMC

As strange and funny as it is to state that “you can be successful where there is no information,” it is equally funny to mention that “one is bound to failure where information is vast and available.” Information should be regarded as “a living fact” and it should be recalled that information is valuable and useful only as long as it is managed properly and kept up-to-date. Sole existence of information is not sufficient. Incorrect, missing, and most importantly out-of-date information may cause loss of labor, time, and inevitably, money.

One of EMC’s core products, EMC Ionix™ ControlCenter (ECC) contains a vast amount of information and is well worth the investment. ECC is used to manage a large SAN cloud, as well as for scaling and reporting functionality. Storage administrators can use the ECC as a tool to simplify their daily work life. However, the process of keeping the ECC information up-to-date and available is critical for optimum use. It is achieved through implementation and proper maintenance of software agents in the SAN environment by the storage administrators.

It is widely known that new sales campaigns are introduced frequently due to stiff competition, especially at large telcos. New applications and servers are introduced to the environment all the time. In such dynamic environments, storage administrators may witness that new servers are added to the ECC environment, that the new servers are not defined in the ECC yet, or that some of the older host agents may not work anymore. The storage administrators have to track down the

responsible server administrator, question why the agents stopped working, or request the installation of new agents on the servers.

Even though storage administrators may share routine installation procedures with server administrators, they may end up having to wait for hours or sometimes, days due to busy schedules. In some environments the administrators may routinely change jobs or go on vacation, resulting in situations where the root password may not even be known. Just as the situation is resolved, the cycle of new and updated servers begins again. Soon enough, a wonderful product such as ECC that could make life easier may seem like an unattainable utopia and part of an annoying infinite cycle by the users.

Turkcell, the largest telco in Turkey, has been experiencing this problem for many years. A lot of effort has been wasted in an upstream battle. Last year, as the onsite EMC engineer, I devised a flexible and easily adaptable solution using UNIX scripting that addresses this problem and keeps the information provided by ECC up-to-date by combining native tools from UNIX and EMC.

In summary, this script suite is copied remotely to all servers over the HP OpenView framework. It detects all UNIX servers connected to the SAN, enabling ECC discovery, remotely installs and configures all necessary software, and starts up the services. If the applications are already installed, it checks the services and starts only those that are required. When errors are encountered, the log files are collected and copied over to a centrally managed location, accessible by the storage administrators. At the same time, it keeps track of all activity in a daily log on a separate remote server. If a host is connected to an EMC storage device, the script collects an "emcgrab" (EMC's utility to analyze host side issues and interoperability) and routinely emails all of its activities.

Using this script, the storage administrators at Turkcell no longer need to worry about installing and maintaining EMC host agents when they allocate EMC or non-EMC storage to a host. They can focus on more important tasks rather than worrying if the server will appear on the EMC console or not. The later sections of this Knowledge Sharing article contain details about the methods used and the value it contributed to the ECC environment at Turkcell.

STORAGE NETWORKING

FIBRE CHANNEL OVER ETHERNET FOR BEGINNERS

Bruce Yellin, EMC

The data center of tomorrow is envisioned to reduce costs, have greater flexibility, and be friendlier to the planet. That dream requires a unified network capable of carrying all data traffic through a single cable. Converged networks have been tried before and were met with limited success. A new approach called Fibre Channel over Ethernet (FCoE) is about to change all of that.

FCoE represents a quantum leap in network technology that supports Ethernet and Fibre Channel storage networks simultaneously over the same cable. FCoE comes to market at a time when data center bandwidth and storage resources are in great demand due in part to the popularity of multi-core processors, virtualization, cloud computing, networked storage, and I/O convergence.

Written for the beginner, as well as for those versed in a single network construct, this Knowledge Sharing article begins with basic concepts and along the way explores the practicality, costs, and a step-by-step approach of how it works. Topics covered include LAN, SAN, voice and video using a single cable, implementation in phases, and how FCoE saves money, manpower, and is future-proof.

CONVERGENOMICS OF FCOE

Akshat Sawhney, EMC

Sahil Gupta, Ace Data Devices

Organizations are in the process of transforming their data center IT infrastructure environments into those that are: more robust and ecological, in line with changing business requirements, focused on eliminating old monolithic legacy operations, and entering into the world of virtualization and cloud. Fibre Channel over Ethernet (FCoE), an industry standard protocol (ANSI T11) that maps Fibre Channel (FC) storage traffic over an enhanced 10 Gigabit (10Gb) Ethernet network, is a key technology that will enable that transformation.

FCoE has experienced a fast growth cycle moving from concept to beta and now to fourth-generation product technology within a few quarters. With this unified network protocol, it is now possible to carry both LAN and SAN traffic over a single Ethernet network. FCoE extends Fibre Channel traffic onto a lossless 10 Gigabit Ethernet fabric, converging LAN and SAN I/O onto one set of cables. This simplified infrastructure reduces the cost of management by joining both worlds (i.e., SAN and LAN). It also significantly lowers overall infrastructure requirements by reducing the number of switches, cables, and cards to deliver great value in terms of total cost of ownership (TCO).

While IT professionals will be happy to see the benefits from a reduction in power, fewer cables, and simplified management, questions remain. What about inter-compatibility? What about investments in existing infrastructure? The simple answer is FCoE fits with the existing IT infrastructure. Implementing FCoE in an IT environment does not require changes to the core network. Support for FCoE traffic will typically require an upgrade to one or more edge switches, such as embedded blade and top-of-rack switches, but does not affect the core switching equipment or topology. Moreover, this switch upgrade is only an incremental addition to the upgrade to 10 Gigabit Ethernet from Gigabit Ethernet that many organizations will be undertaking in the next year or two.

Top level CIOs and CTOs managing their company's IT infrastructure will be happy to learn the definition of a new word "convergeconomics," which talks in terms of graphs, percentages, and numbers. For example, FCoE delivers 28 percent savings on switches, 42 percent savings on power and cooling, and 80 percent saving on cables. Implementation personnel will be pleased to know that FCoE works on an incremental deployment model. In short, it causes minimum disruption in existing environments while ensuring maximum benefits.

FCoE has been designed from the ground up for maximum flexibility, including the capability to coexist with existing infrastructure, management topologies, and networking technologies. Backing from industry leaders such as EMC, Cisco, IBM, HP, NetApp, and Intel helps ensure the continuing innovation that is pushing FCoE into the networking environments of companies of all sizes and across all industries.

This Knowledge Sharing article explores the challenges faced by data center customers, and provides insight into navigating a course of action to provide budget relief while offering better services to internal and external customers. The article goes into detail about the fundamentals by which FCoE works, security concerns, its implementation, and the role it plays in virtualization, cloud computing, and much more.

SYMMETRIX

USING SE CONFIGCHECKER FOR CHECKING SOFTWARE COMPATIBLE PRE-REQUISITES FOR SYMMETRIX

Parthasarathi Ilangovan, EMC

The core value of a company like EMC is to not only provide powerful cutting-edge hardware and software but also to deliver complete targeted solutions to help businesses address specific enterprise needs.

The wealth of products brings with it a range of support issues involving unsupported configurations, invalid configurations, obsolete versions, and so on. While there has been every attempt to document information in as user friendly a way as possible, it is still a Herculean task to deliver this information to a customer as an easy, elegant, and complete solution.

SE Configchecker is a product designed to resolve problems by checking for all supported matrix for almost all major EMC Symmetrix products in terms of pre-requisites, unsupported configurations, recommendations, device settings, firmware checks, third-party software compatibilities, user permissions, and so forth, and deliver a consolidated report to the user with a list of success/failure results. The product is available across all major platforms, and has an integrated update feature to seamlessly pull new test checks to the customer site from the EMC back end.

The goal of this Knowledge Sharing article is to provide an overview (and much needed visibility) into SE Configchecker's features and usage for the target customer demographic including, but not limited to, direct customers, regular users, customer service teams, and pre-sales teams. The full potential of SE Configchecker is to deliver on automatic checking for standard EMC-supported configurations for each product, resulting in early detection of issues. This frees the support engineers to focus on more targeted expertise needs, while also facilitating a better customer experience for the end user.

VIRTUALIZATION

DISCOVERING VIO RELATIONSHIPS FOR EMC ARRAYS

Ewa Nowatkowska-Glowacz, EMC

This Knowledge Sharing article describes ways of discovering VIO relationships for IBM p-based systems connected to EMC Symmetrix arrays. The Virtual I/O Server (VIOS) provides virtual storage to client logical partitions. It allows a physical adapter with attached disks on the VIOS partition to be shared by one or more Virtual I/O Client (VIOC) partitions.

This type of configuration with use of VIOS/VIOC relationship is becoming increasingly popular among enterprise customers where IBM servers are installed. In large migration projects, discovering and documenting relationships between servers and their clients is a time-consuming task.

Currently, there are no tools available which could perform (from the client point of view) volume or device mapping between servers and clients. The identification method shown in this article could be useful for verifying current configuration for migration purposes or can even be used for daily reporting. The article cites a real implementation project where the EMC team worked out methodology and scripts needed to discover current relationships to perform migration of VIOS resources between arrays.

The configuration description and commands' outputs shown in this article are based on documentation of a real environment, and help provide an understanding of VIO configuration challenges. Sample scripts for reports automation for large configurations are shown at the end of the article.

This Knowledge Sharing article helps the reader to understand and implement ways of discovering configuration behind a VIOS server in a customer's environment where VIO configuration is involved.

VPLEX: INTERESTING USE CASES

Sandeep Choudhary, Tech Mahindra

Puneet Goyal, HCL Comnet

Cloud, cloud computing, cloud environment, cloud service, cloud center, and on and on! As field engineers, when visiting data centers at customer sites, the conversation usually comes around to the cloud facility. Comparisons are made between NetApp V series, HDS USP-V, and others. But when it comes to EMC VPLEX™, all seem to be in agreement; everyone sees that EMC has taken things one step further with VPLEX. In our view, EMC VPLEX jumps several steps ahead.

With VPLEX, you can share storage across sites using distributed caching algorithms making active/active data centers a reality. Currently available across synchronous distances (100km) and for local site use cases, this product does all the things EMC Invista® does, plus cache coherency; plus across distance; plus with a scale out cluster model; plus with simplified setup and management—without performing the virtualization at the SAN switch level (thus not requiring stretched fabrics between sites); plus doing it with an appliance model, plus...well, you get the point.

The main reason for this technology is to further enable cloud-based IT architectures by dynamically moving storage workloads across data centers (or even within different arrays in the same data center) seamlessly, without any application downtime or interaction. One can also make a case for easy tech refreshes since virtualizing the storage volume decouples it from the physical array, allowing you to play the “storage vMotion®” game with the back-end array on which the data actually resides. This is very similar to how VMware decouples the application/OS workload from the physical server.

Those are all pretty standard use cases for storage virtualization, but how about some others? In this Knowledge Sharing article we will discuss some of the cases in detail covering the following topics:

- An alternative to array-based replication
- In the future, CDP using VPLEX
- Cache is king
- What is the use of SRM if you have a VPLEX?
- When cloud brings the rain

STORAGE AND VIRTUALIZATION TECHNOLOGIES IN INFORMATION TECHNOLOGY EDUCATION

Dr. Timur Mirzoev, Georgia Southern University

This Knowledge Sharing article describes the implementation and use of virtualization and storage technologies in Information Technology education. Universities' involvement with academic programs such as EMC Academic Alliance and VMware IT Academy allows the educational experience to become much more than theory and presentations. This article presents many aspects of the utilization,

deployment, and administration of physical resources that allow for an advanced learning experience.

The workforce of America is initially shaped by the educational experience of students. This process has a tremendous impact on the professional success of future IT workers. Today, the gap between IT education and commercial training is closing rapidly—technology leaders such as EMC and VMware support academic programs that disseminate commercial training materials to students free of charge. Students have an opportunity to receive certifications such as EMC Proven Professional and VMware Certified Professional before they graduate. This practice of presenting real-world IT knowledge to students is a sign of significant progress for companies interested in hiring storage and virtualization professionals.

Our IT program is both an EMC Academic Alliance partner and an International VMware IT Academy Center. There were several investments and donations of commercial-grade equipment that boosted our IT program. Today, over 30 x86 servers and several open-source storage servers support the infrastructure of over 350 virtual machines, and 6 TB of storage—all for educational purposes. Currently, VMware vSphere® 4.1, Openfiler, and Dell iSCSI systems are utilized.

Students with various permissions and privileges use commercial equipment and learn to become future storage and data center administrators. High Availability, Distributed Power Management, Storage vMotion, and many other technologies are used in our laboratory to support our students' learning experiences.

It took three years to reach the commercial level of sophistication for our IT program to become the leading school in storage and virtualization education. Readers of this Knowledge Sharing article will gain valuable information about the educational experience of our students as well as learn about advanced virtualization, storage, and cloud computing technologies.

CONTACT US

Contact Us by Phone or Email

Online: <http://education.EMC.com>

Email: EdServices@EMC.com

Phone: 1-888-EMC-TRNG(888-362-8764)

International:

APEducation@EMC.com

+61 2 9922 7888 (ANZ)

+65 6333 6200 (South Asia)

EMEA_Education@EMC.com

+44 208 758 6080 (UK)

+49 6196 4728 666 (Germany)

GCCEducation@EMC.com

+86 10 8438 6593 (Greater China)

India_Education@EMC.com

+91-80-6737-5064 (India)

Japan_Education@EMC.com

+81 44 520 9830 (Japan)

Korea_Education@EMC.com

+82 22125 7750 (Korea)

EMC², EMC, EMC Proven, Atmos, Avamar, Centera, Celerra, Celerra Replicator, CLARiiON, ControlCenter Data Domain, DiskXtender, Documentum, Invista, Ionix, NetWorker, PowerPath, SAN Copy, SnapSure, Symmetrix, VPLEX, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. VMware, ESX, vMotion, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners. © Copyright 2011 EMC Corporation. All rights reserved. Published in the USA. 4/11 Customer Profile H2771.8

EMC Corporation

Hopkinton, Massachusetts 01748-9103

1-508-435-1000 In North America 1-866-464-7381

www.EMC.com

EMC²[®]