

VPLEX Management Student Guide

Education Services
July 2015

VPLEX Management



© Copyright 2015 EMC Corporation. All rights reserved.

Support Contact: [Education Services](#)

1

Welcome to VPLEX Management.

Copyright ©2015 EMC Corporation. All Rights Reserved. Published in the USA. EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license. The trademarks, logos, and service marks (collectively "Trademarks") appearing in this publication are the property of EMC Corporation and other parties. Nothing contained in this publication should be construed as granting any license or right to use any Trademark without the prior written permission of the party that owns the Trademark.

EMC, EMC² AccessAnywhere Access Logix, AdvantEdge, AlphaStor, AppSync ApplicationXtender, ArchiveXtender, Atmos, Authentica, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, Avamar, Bus-Tech, Captiva, Catalog Solution, C-Clip, Celerra, Celerra Replicator, Centera, CenterStage, CentraStar, EMC CertTracker, CIO Connect, ClaimPack, ClaimsEditor, Claralert, cLARIiON, ClientPak, CloudArray, Codebook Correlation Technology, Common Information Model, Compuset, Compute Anywhere, Configuration Intelligence, Configuresoft, Connectrix, Constellation Computing, EMC ControlCenter, CopyCross, CopyPoint, CX, DataBridge, Data Protection Suite, Data Protection Advisor, DBClassify, DD Boost, Dantz, DatabaseXtender, Data Domain, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, DLS ECO, Document Sciences, Documentum, DR Anywhere, ECS, eInput, E-Lab, Elastic Cloud Storage, EmailXaminer, EmailXtender, EMC Centera, EMC ControlCenter, EMC LifeLine, EMCTV, Enginuity, EPFM, eRoom, Event Explorer, FAST, FarPoint, FirstPass, FLARE, FormWare, Geosynchrony, Global File Virtualization, Graphic Visualization, Greenplum, HighRoad, HomeBase, Illuminator, InfoArchive, InfoMover, Infoscape, Infra, InputAccel, InputAccel Express, Invista, Ionix, ISIS, Kazeon, EMC LifeLine, Mainframe Appliance for Storage, Mainframe Data Library, Max Retriever, MCx, MediaStor, Metro, MetroPoint, MirrorView, Multi-Band Deduplication, Navisphere, Netstorage, NetWorker, nLayers, EMC OnCourse, OnAlert, OpenScale, Petrocloud, PixTools, Powerlink, PowerPath, PowerSnap, ProSphere, ProtectEverywhere, ProtectPoint, EMC Proven, EMC Proven Professional, QuickScan, RAPIDPath, EMC RecoverPoint, Rainfinity, RepliCare, RepliStor, ResourcePak, Retrospect, RSA, the RSA logo, SafeLine, SAN Advisor, SAN Copy, SAN Manager, ScaleIO Smarts, EMC Snap, SnapImage, SnapSure, SnapView, SourceOne, SRDF, EMC Storage Administrator, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix, Symmetrix DMX, Symmetrix VMAX, TimeFinder, TwinStrata, UltraFlex, UltraPoint, UltraScale, Unisphere, Universal Data Consistency, Vblock, Velocity, Viewlets, ViPR, Virtual Matrix, Virtual Matrix Architecture, Virtual Provisioning, Virtualize Everything, Compromise Nothing, Virtuent, VMAX, VMAXe, VNX, VNXe, Voyence, VPLEX, VSAM-Assist, VSAM I/O PLUS, VSET, VSPX, Watch4net, WebXtender, xPression, xPresso, Xtrem, XtremCache, XtremSF, XtremSW, XtremIO, YottaYotta, Zero-Friction Enterprise Storage.

Revision Date: July 2015

Revision Number: MR-7CN-VPLEXMGMT.5.4.1.2

Course Overview

Description

This course provides the learner with an in-depth understanding of how to manage, monitor and maintain a functional VPLEX system. Operational aspects of a VPLEX Local and a VPLEX Metro, using both the VPLEX Management Console GUI and the VPLEX CLI, are covered. The emphasis will be on provisioning virtualized storage from VPLEX to hosts, and performing operations on live virtual volumes such as mobility and mirroring. Integration of VPLEX into existing data centers with live production data, and interoperability considerations, are included. This course also covers routine monitoring of VPLEX as well as an introduction to VPLEX and RecoverPoint. This course covers VPLEX version 5.4.

Audience

This course is intended for any person who presently or plans to manage a VPLEX solution.

Objectives

Upon completion of this course, you should be able to:

- Describe common VPLEX terms, configuration options, hardware and software architecture
- Describe high level data flow within a VPLEX system.
- Create Virtual Volumes and provision them to Hosts
- Create and Manage VPLEX Distributed Devices
- Manage VPLEX using Unisphere and CLI
- Perform volume expansion or migration
- Protect VPLEX Devices with RecoverPoint
- Perform Monitoring in VPLEX including batched mobility.



This course covers implementing a VPLEX System, managing it, and performing basic maintenance and monitoring.

Agenda

Modules

Labs

Day 1	• Course Overview • VPLEX Product and Technology Details • VPLEX Storage Provisioning	• VPLEX Cluster Health Check
Day 2	• VPLEX Storage Provisioning (cont) • VPLEX Distributed Devices • Volume Management and Protection	• Provisioning with VPLEX • Creating Distributed Devices
Day 3	• Volume Management and Protection (cont) • VPLEX Monitoring	• Mobility with VPLEX • Protection with RecoverPoint • VPLEX System Monitoring

Here we see the Agenda for this 3-day class; please take a moment to go over it.

This Course is Part of Learning Path for:



- **Implementation Engineer (EMCSA) Track**
EMC Proven Professional Certification
 - **E20-562 VPLEX Specialist Exam for Storage Administrators**
 - See: <https://edu.corp.emc.com/gs/certification/default.aspx>

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Course Introduction



4

This course is part of the learning path for VPLEX Specialist Exam for Storage Administrators Specialist EMC Proven Professional certification. For details on the EMC Proven Professional program, please see: <https://edu.corp.emc.com/gs/certification/default.aspx>.

Module: VPLEX Product and Technology Details

Upon completion of this module, you should be able to:

- List the ways VPLEX fits into the virtualized data center strategy
- Describe VPLEX architecture and product capabilities
- List the different VPLEX system and volume management options
- Describe the different VPLEX management options



This module explores VPLEX architecture and product capabilities. We will introduce some of the use cases for VPLEX along with some of the management options.

Lesson: VPLEX Product Overview


This lesson covers the following topics:

- VPLEX Product Overview
- VPLEX Use Cases
- VPLEX Data Protection



This lesson covers an overview of the VPLEX product.

The Journey to Platform 2.5



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


EMC²

3

For years, users have relied on “physical storage” to meet their information needs. New and evolving changes, such as virtualization and the adoption of Private Cloud computing, have placed new demands on how storage and information is managed. To meet these new requirements, storage must evolve to deliver capabilities that move information from being a physical element to being a virtualized resource – a resource that is: fully automated, integrated within the infrastructure, consumed on demand, cost effective and efficient, always on, and secure. The technology enablers needed to deliver these qualities combine unique EMC capabilities such as FAST, Federation, and storage virtualization. The result is a next generation Private Cloud infrastructure that allows users to:

- Move thousands of VMs over thousands of miles
- Batch process in low-cost energy locations
- Enable boundary-less workload balancing and relocation
- Aggregate large data centers
- Deliver “24 x forever” – and run or recover applications without ever having to restart.

Journey to the Private Cloud



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


4

When EMC discusses the Private Cloud, it is describing an infrastructure strategy that enables optimized resource use. This means you're optimized for energy, power, and cost savings. You can scale up and out simply and apply automated policies. You can guarantee greater availability and access for your production environment—significantly reducing or eliminating downtime.

Introduction to the EMC VPLEX Family

VPLEX is Designed to Provide:

- Data mobility
 - Within the same data center
 - Across a campus
 - Non-disruptive migration
 - Storage pooling
- High Data Availability
 - Within local and across campus IT centers
- Collaboration
 - Real-time data collaboration for Big Data application



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

5

VPLEX federates data that is located on heterogeneous storage arrays to create flexible, distributed and highly available data centers.


VPLEX addresses three primary IT needs:

- Mobility: VPLEX moves applications and data between different storage installations:
 - Within the same data center or across a campus (VPLEX Local)
 - Within a geographical region (VPLEX Metro)
- Availability: VPLEX creates high-availability storage infrastructure across these same varied geographies with a very high level of resiliency.
- Collaboration: VPLEX provides efficient real-time data collaboration over distance for Big Data applications.

EMC VPLEX

AccessAnywhere Enabled by Distributed Virtual Volumes

- Provides shared storage, concurrently Read And Write accessible, with the same storage identifiers, across sites
- Supported by nearly all open systems application clustering technologies and FC storage arrays
- Simplifies high availability and disaster recovery design and operations
- Enables 100% storage availability



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

6

EMC VPLEX enables Active/Active distributed datacenters.

VPLEX is a storage virtualization solution. There are many virtualization solutions in the marketplace today, but the key differentiator is that VPLEX was designed from concept to enable multi-site storage virtualization.

VPLEX is a hardware and software solution that can be used to create an abstraction layer between hosts and storage arrays. From a storage perspective VPLEX looks like a host, and you can present Physical LUNs from the storage array to VPLEX. From a host perspective VPLEX looks like a storage array, and you present virtual LUNs from VPLEX to the host. By breaking the link between the host and the physical storage array it enables storage virtualization use cases like non-disruptive data mobility across arrays.

In this context, it is possible to take on another datacenter, with another physical storage array and a completely separate, resilient VPLEX cluster. The VPLEX Clusters can then be federated. Federation is the process of creating a communication path between the VPLEX clusters. This can be done over FC or IP, however the RTT between the DCs must be less than 5ms and there must be enough bandwidth to sustain synchronous replication. You do not need to merge fabrics between sites.

Once federated it is possible to create a distributed virtual volume. This aspect of the distributed virtual volume is a bit like a RAID 1 mirror; reads and writes can occur at both sites and VPLEX consistently mirrors the data across storage arrays. One advantage of this architecture is that it enables 100% storage availability. In the event of a problem with an array, it is not required to fail application from site A to site B. Read and writes continue to be served from either VPLEX cache or the remote array.

The key consideration in regard to distributed virtual volumes is that they have exactly the same storage and LUN identifiers across both sites and are concurrently RW accessible. Why is that important..


VPLEX Product Family

EMC VPLEX Local



- Within a Data Center
- Single VPLEX cluster

EMC VPLEX Metro



- AccessAnywhere at Synchronous distances
- No more than 5ms RTT

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

7


The VPLEX product family is composed of VPLEX Local and Metro system configurations.

VPLEX Local provides seamless, non-disruptive data mobility and the ability to manage and mirror data between multiple heterogeneous arrays from a single interface within a data center. VPLEX Local consists of a single VPLEX cluster. It contains a next-generation architecture that allows increased availability, simplified management, and improved utilization and availability across multiple arrays.

VPLEX Metro enables active/active, block level access to data between two sites within synchronous distances. The distance is limited not only by physical distance but also by host and application requirements. Depending on the application, VPLEX clusters should be installed with inter-cluster links that can support not more than 5ms round trip time (RTT). The combination of virtual storage with VPLEX Metro and virtual servers enables the transparent movement of virtual machines and storage across synchronous distances. This technology provides improved utilization and availability across heterogeneous arrays and multiple sites.

VPLEX Use Cases

Online Migration between Arrays



- Benefits

- Non-disruptive to host applications
- Simplification of migration process
- Uniform migration process between heterogeneous arrays

- Use-cases

- Lease roll-over / tech refresh
- Load balancing between arrays

No application impact during migration
between arrays within a data center

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

8

VPLEX Local and/or Metro allows the non-disruptive movement of data from one array to another. This capability can prevent many management issues, such as tech refreshes of arrays.

Here's a quick view of the traditional tech refresh and corresponding data migration scenario. After the new array arrives and is installed and ready, the data source is copied by any of a variety of host- or array-based tools.

Outages are planned over a sequence of weekends to migrate servers to the new storage array, often in batches of associated applications. This is the "Race to Monday," as it's typically done over weekends. This process is risky because it's typically non-reversible.


During this interim period, the customer must manage two separate arrays - remote replication, snapshots, clones, etc.

Ultimately, once all of the data is migrated the old array can finally be decommissioned. The vendor company often allows lease and services concessions to the customer during the migration period. And until that gear arrives at the vendor facilities, the hardware cannot be recycled back into inventory. And of course, until the migration completes, the asset they've paid for cannot be used – the Time-to-Value expectations are not achieved.

VPLEX allows the seamless migration of application data from one array to another. It is not dependent upon array technology; this allows data migration between different array types.

VPLEX Use Cases (Cont'd.)

Storage Pooling



- Pool arrays together
- Reduce cost without increasing risk
- Move devices 100% online between different tiers
- Mirror across arrays for enhanced availability

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


9

VPLEX can transform individual arrays of many vendor and performance types into pools or tiers of storage. The performance needed for certain applications can vary over time. To accommodate this, VPLEX can move the storage from one tier to another without any disruption, and the workload can be placed where it makes sense.

VPLEX Use Cases (Cont'd.)

Availability

- Question:
 - In a worst case disaster scenario, what is the best possible outcome IT could expect?
- Answer:
 - IT continues as if nothing happened



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

10

An unplanned outage is the ever-present risk that organizations have to deal with on a daily basis. And it may seem like there are increasing occurrences where nature is trying to bring your business down. Whether it's a tornado in the midwest, a hurricane in the east, an earthquake in California or a tidal wave and flooding in a coastal country, natural disasters are a real threat that must be planned for.

Equally devastating are unplanned outages caused by human error. A fire in the data center causing the sprinklers to rain array-killing water down on your server racks can be just as destructive to your data center as a natural flood overflowing its banks. And even a mistake by a well meaning IT team member can have disastrous consequences.

Through the use of its continuous data availability, VPLEX provides protection from a data-center-wide failure. This failure can be either planned or unplanned.


Even if there are other disaster recovery methods in place, the work required in order to both failover and fail back is extensive.

VPLEX "Active-Active" technology eliminates the need for extensive planning and testing for these events.

VPLEX Continuous Availability

Continuous Availability = High Availability + Disaster Recovery

Stretched Host Clustering



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

11

There are two strategies for application and Datacenter resiliency:

Active - Passive

and

Active - Active

What dictates the decision to choose one of these topologies? Application RPO/RTO objectives, distance, and bandwidth between sites.

What if you could build redundancy into storage that allowed active-active?

Here are some of continuous availability key benefits:

No more DR testing

Active Assets in both sites

Works with all open applications


No complex failover

Fully automatic recovery (0 RPO, 0 RTO, 0 DTO)

Eliminate both planned and unplanned downtime

Reduce TCO and increase service level

Use Case Example: Site Failure



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

EMC²

12

VPLEX allows for high availability data storage infrastructure as well as for eliminating failover and failback from storage operations. Now you avoid the restart problem and risks.

Stretched VMware Clusters



Automatic load balancing
across sites

* Application Specific

EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

13

A VPLEX Metro supports VMware HA and FT. It also provides VMware DRS and Vmotion integration. In order for the VM's and applications to failover transparently, the data must be shared across cluster nodes. VMware ESXi clustering requires shared storage. To provide non-disruptive movement of virtual machines, VPLEX Metro allows storage to span multiple datacenters. This allows ESXi servers in different failure domains to share access to datastores created on VPLEX Distributed Storage. VPLEX Metro fits perfectly with VPLEX distributed cache coherence for automatic sharing and load balancing.

VPLEX supports VMware vSphere® Storage APIs – Array Integration (VAAI), also referred to as hardware acceleration or hardware offload APIs. The APIs define a set of "storage primitives" that enable the ESXi host to offload certain storage operations to the array, VPLEX, which reduces resource overhead on the ESXi hosts and can significantly improve performance for storage-intensive operations such as storage cloning, zeroing, and so on.



VPLEX and RecoverPoint products together offer continuous availability and operational and disaster recovery. For customers requiring even further levels of availability, VPLEX can be used with RecoverPoint to enable a third site for disaster recovery. This site can be located at any distance, ensuring operational resilience in the event of a regional outage.

Three Site (Synchronous and Asynchronous) Availability Summary:

VPLEX Capabilities

Active-Active

- Transactions must be running and stable at two locations within Metro distances, mostly financial institutions
- Provide Continuous Availability With VPLEX Array Mirroring
- Mirror data from XtremIO to XtremIO to protect against the loss of an array or even an entire site.
- Move Workloads To and From XtremIO Without Disruption
- Application data can be stored on spinning disk and non-disruptively moved to XtremIO when performance requirements demand it
- Expand X-bricks Without Disruption With VPLEX
- Putting VPLEX between hosts and storage allows IT to add XtremIO capacity and performance by adding X-bricks without disrupting application data access

RecoverPoint Capabilities:

Remote replication with recovery to any point-in-time


- Universal replication – any application, any distance, any array using VPLEX, any network
- Supports bi-directional asynchronous replication over a Fibre Channel or IP network over any distance
- Saves 90% of network bandwidth consumption thru snap and replicate capabilities, deduplication and compression

VPLEX MetroPoint

MetroPoint Combines the Full Value of VPLEX and RecoverPoint

MetroPoint provides:

- Single DR copy
- Automatic switchover
- Copies on both VPLEX clusters
- Protection against cluster failure



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

15

In addition to the previously described use with RecoverPoint, VPLEX at GeoSynchrony v5.4 along with RecoverPoint v4.1 SP1 supports a new topology for VPLEX Metro teamed with RecoverPoint protection. This topology is called **MetroPoint**. MetroPoint topology is enabled by combining the features of VPLEX Metro, an active-active multi-site infrastructure, and RecoverPoint, for continuous data replication to the remote third site. Comprehensive data protection continues even under the complete failure of one of the Metro region sites. The simultaneous protection of the Metro region by a distant third site using RecoverPoint provides any Point-in-Time recovery from operational and disaster outages.

MetroPoint topology helps organizations to achieve a high level of continuous availability and data protection that completely closes the RPO/RTO gap.

For the highest level of protection, both VPLEX Metro sites can add RecoverPoint Local protection. This provides additional benefits:


Faster RTO – reading from the local copy will likely be faster than remote copy

Redundant copies – protects against the loss of the remote site or the links to the remote sites

To use the MetroPoint feature, RecoverPoint appliances must be running software version 4.1.1 Patch 1 or later, and VPLEX must be running software version 5.4.1 or later.

MetroPoint Use Case Example: Two-Site Continuous Availability and Local Protection

- Local replication on both sides of VPLEX Metro
- Redundant operational
- Recovery for applications
- Able to survive a cluster failure



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

16

In the two-site topology, distributed devices are protected with an independent local copy at both Metro sites that allows for independent replication at either side of a VPLEX Metro for a distributed volume. Replication continues at one site if an outage occurs at the other site. With VPLEX Metro, production continues to be available from one site after a WAN COM failure or a VPLEX cluster failure when configured with VPLEX Witness. To access copy data, the administrator decides which copy to use. When two VPLEX clusters are not in contact, administrators are recommended to use the copy from the site where the production volume is accessible by the host. Regardless of which copy is used for recovery, replication history is maintained. If recovery is done on the losing site during a VPLEX WAN partition, manual intervention is required to select the site's data to use after the WAN link is restored. The losing site will not be able to create bookmarks while the volumes rebuild, but replication history will not be lost.

RecoverPoint local replication can be added to both sides of a VPLEX Metro to provide redundant operational recovery for the same application. To achieve this, all of the volumes associated with an application must be placed in a MetroPoint group, even though a third site is not configured.

This two-site configuration still requires GeoSynchrony 5.4 and RecoverPoint 4.1 SP1 that support MetroPoint on both VPLEX and RecoverPoint.

Lesson Summary

During this lesson the following topics were covered:

- VPLEX Product Overview
- VPLEX Use Cases
- VPLEX Data Protection



This lesson covered an overview of VPLEX benefits and use cases.

Lesson: VPLEX Architecture

This lesson covers the following topics:


- VPLEX architecture
- VPLEX hardware configuration options
- Hardware redundancy of a VPLEX Cluster
- VPLEX System Software Environment
- Performance and protection options for virtual volumes



This lesson covers an overview of the VPLEX system architecture.

VPLEX Architecture

- Next generation data mobility and access
- Scale out cluster architecture
 - Start small and grow big
- Advanced data caching
 - Improve I/O performance and reduce storage array contention
- Distributed cache coherence
 - Automatic sharing, balancing and failover across VPLEX Engines



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


19

EMC VPLEX is a next-generation architecture for data mobility and information access.

It is based on unique technology that combines scale out clustering and advanced data caching, with the unique distributed cache coherence intelligence to deliver radically new and improved approaches to storage management.

This architecture allows data to be accessed and shared between locations over distance via a distributed federation of storage resources.

VPLEX Local Architecture




VPLEX presents virtualized LUNs to connected hosts from attached physical storage arrays. VPLEX can be broken into two functions. It functions as a host to storage arrays. It also functions as a storage array to hosts.

If you follow the path of storage from the "bottom" up, notice the first layer is the physical storage. VPLEX is zoned and masked to the array much like a normal host. Next is the virtual storage layer. In the context of availability, it is important to note that VPLEX has redundant everything and is battery-backed as well.

VPLEX Local Configuration

- One EMC Rack
- 1, 2 or 4 VPLEX Engines
 - 2 directors per Engine
- Redundant hardware
 - All components are redundant
 - except for the Management Server
- N+1 performance and scaling
 - More ports and engines
- 1+N fault tolerance

4 Engine Configuration



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

21

VPLEX Local consists of a single rack. Most installations are with EMC racks, although custom racks are approvable. The rack contains 1, 2, or 4 VPLEX engines. Each engine consists of 2 directors: A and B. The hardware is fully redundant to survive any single point of failure.

A VPLEX is fully upgradable to the next size. This upgrade can be done without disruption. VPLEX Engines can scale for greater performance and redundancy. In any cluster, the fully redundant hardware can tolerate failure down to a single director remaining with no Data Unavailability or Data Loss condition.

VPLEX Configurations – Non Disruptive Growth



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


EMC

22

VPLEX is configured as a Single, Dual, or Quad Engine Cluster. A VPLEX Local will have one cluster, and a Metro has two clusters connected together. The Engines and associated equipment will come pre-assembled from the factory.

All three configurations contain a single management server. The dual and quad engine configurations have 2 FC Switches for Director connectivity along with UPS's. A VPLEX Cluster can be non-disruptively upgraded from a single engine to a dual or quad.

VPLEX VS2 Engine



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

23

A VPLEX Engine (back view displayed) is divided into 2 directors, A and B. Each director has independent resources such as cache, power, and connectivity. Directors work together in a fully Active-Active configuration. In this type of configuration, each director can service any connected hosts and storage. Each Director has a 32GB SSD drive for local cache.

Management and Power/Fan Redundancy



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²


24

Independent power zones in the data center feed each VPLEX power zone, providing redundant high availability. Each engine is connected to 2 standby power supplies (SPS) that provide a battery backup for cache vaulting in the event of transient site power failure. In single-engine clusters, the management server draws its power directly from the cabinet PDU. In dual- and quad-engine clusters, the management server draws power from UPS-A.

Each VPLEX engine is supported by a pair of standby power supplies (SPS) that provide a hold-up time of five minutes, allowing the system to ride through transient power loss. A single standby power supply provides enough power for the attached engine. Each standby power supply is a FRU and can be replaced with no disruption to the services provided by the system. The recharge time for a standby power supply is up to 5.5 hours. The batteries in the standby power supply are capable of supporting two sequential five-minute outages.

There are two power supplies with fans per director. Both must be removed to pull the director out. To ensure the power supply is completely inserted into the engine, the yellow at the top of the power supply should not be visible.

Standby Power Supply (SPS) LEDs



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


25

Each director is connected to a standby power supply. The standby power supplies should be connected to separate power sources. Generally, the SPS On Battery LED stays on while the SPS units fully charge (which could be a few minutes or a few hours, depending on the state of the battery). If any amber LED not related to the SPS recharge remains on for more than 10 minutes, verify that the components are cabled correctly.

VPLEX I/O Modules



- Slots B2 and A2 (WAN COM) will contain one of the I/O Module types shown:
- Both WAN COM IOMs must be of the same type.



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

26

A VPLEX VS2 Engine has 8 I/O modules, with four allocated to each director. Each director has one four-port 8 Gb/s Fibre Channel I/O module used for front-end or Host connectivity and one four-port 8 Gb/s Fibre Channel I/O module used for back-end or storage array connectivity. Each of these modules has 40 Gb/s effective PCI bandwidth to the CPUs of their corresponding director. A third I/O module, called the WAN COM module, is used for inter-cluster communication. Two variants of this module are offered, one four-port 8 Gb/s Fibre Channel module and one two-port 10 Gb/s Ethernet module. The fourth I/O module provides two ports of 8 Gb/s Fibre Channel connectivity for intra-cluster communication. The fifth I/O module for each director is reserved for future use.


Let's look more closely at these two four-port 8 Gbps Fibre Channel I/O modules, that provide 40 Gbps effective bandwidth: four ports for front-end connectivity and four ports for back-end connectivity. This setup shows us that there is a planned oversubscription within the VPLEX. This is not an issue with proper planning of the connections to the ports and modules.

WAN communication between VPLEX Metro clusters is over Fibre Channel or Gigabit Ethernet. The inter cluster link carries unencrypted user data. To protect the security of the data, secure connections are required between clusters.

Depending on the cluster topology, slots A2 and B2 contain one of the displayed I/O modules (IOMs). Both IOMs must be the same type.

VPLEX Constructs

- Virtual Volume – storage presented to one or more hosts
- (Local) RAID topology applied to one or more extents
- An extent is a contiguous area of a storage volume
- Array volume (LU) presented to VPLEX BE ports



EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

27

Displayed are the basic elements that comprise a virtual volume.

- Storage Volume – the LU presented from an Array to the VPLEX BE ports
- Extent – a slice or portion of the available space from a storage volume – can be the entire storage volume
- Device – (local device) the application of a RAID topology to one or more extents. A device may use another device as an extent. Devices can be made up of any combination of local devices and extents as appropriate for a particular RAID topology/geometry.
- Virtual Volume – created from the top level device. It is the storage or Logical Unit presented to one or more hosts

VPLEX Local Scale Out, Active/Active, Resilient



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

28

VPLEX directors are arranged to work together to provide uniform scalability and data access resiliency in a fully active/active data access configuration. An SSD in each director (56GB) is used for booting, cache vaulting, and log space.

All F/E ports are connected to initiator ports and all B/E ports are connected to the arrays.

Assuming this is a dual fabric SAN configuration you would be presenting at minimum (2) paths per director. You never want to have a single path to a single director in VPLEX.


All Directors must see all storage equally

Active – passive arrays must have redundant fabric links

- (1) active (1) passive ports zoned to each director

Zoning for Active/Passive arrays; make sure the ITL is all in one zone

VPLEX Metro Builds on VPLEX Local




Displayed is a VPLEX Metro HA solution. VPLEX's Access Anywhere technology allows you to export a single virtual volume from both of these VPLEX clusters simultaneously.

In this solution, the physical hosts are only connected to their local VPLEX cluster and no SAN extension between sites is necessary. This helps you to isolate your datacenters and prevent failures in one datacenter from affecting the other.

There are many interesting solutions that can be built atop Metro HA . VMware and Oracle RAC are two examples.

VPLEX Metro Configuration

- Two racks of fully redundant hardware
- 2 to 8 Engines total in a VPLEX Metro
 - Each rack can consist of 1,2 or 4 engines
- FC-WAN 8 Gb Fibre Channel ports
 - 4 per engine / 2 per director



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

30

VPLEX Metro consists of two separate rack sites with less than 5ms round-trip latency. Each rack may contain 1, 2, or 4 engines, with each engine containing two directors. The hardware is fully redundant to survive any single point of failure.

Each director utilizes two 8 Gb FC ports for inter-cluster communication over the FC-WAN.

VPLEX Metro – High Level



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

31

VPLEX Metro consists of two VPLEX clusters connected by inter-cluster links with not more than 5ms Round Trip Time (RTT). VPLEX Metro transparently relocates data and applications over distance, protects the data center against disaster, and enables efficient collaboration between sites. It allows administrators to manage all storage in both data centers from one management interface.

The VPLEX Metro configuration:


- Mirrors stored data to a second site, with full access at near local speeds.

Deploy VPLEX Metro within a data center for:

- Additional virtual storage capabilities beyond that of a VPLEX Local.
- Higher availability.

VPLEX Metro clusters, Cluster-1 and Cluster-2, can be placed in different Datacenters, allowing them to be located in unique failure domains. There is no maximum distance requirement as long as the latency requirement is met.

VPLEX Management IP Infrastructure



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

32

Shown is the high-level architectural view the management connections between the management server and directors. In this picture there are NO internal VPLEX IP switches. The directors are in fact daisy chained together via two redundant Ethernet connections. The management server also connects via two redundant Ethernet connections to the directors in the cluster. The management server is the only VPLEX component that gets configured with a public IP on the data center network. From the data center IP network, the management server can be accessed via SSH or HTTPS.

Lesson Summary

During this lesson the following topics were covered:

- VPLEX architecture
- VPLEX hardware configuration options
- Hardware redundancy of a VPLEX Cluster
- VPLEX System Software Environment
- Performance and protection options for virtual volumes



This lesson covered an overview of the VPLEX system architecture.

Lesson: VPLEX I/O Operations

This lesson covers the following topics:


- Define caching layers within a VPLEX system
- Describe I/O operation within VPLEX Local and Metro
- Describe VPLEX system interaction during reads and writes



This lesson covers an overview of VPLEX I/O operations, including caching, read and write interactions, and component redundancy.

Cache Coherence

- Distributed cache
- Central directory for tracking
- Manages locking and queuing



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

35

Cache coherence creates a consistent global view of a volume.

Distributed cache coherence is maintained using a directory. There is one directory per user volume and each directory is split into chunks (4096 directory entries within each). These chunks exist only if they are populated. There is one directory entry per global cache page, with responsibility for:


- Tracking page owner(s) and remembering the last writer
- Locking and queuing

Every director has a copy of the cache coherence directory. When a host makes a write, the write gets added to a director's local cache. The director then checks the cache coherence table to see if another host has a lock on the write location. If another director does, then the director will wait for the lock to be released.

Once the lock has been released, the director updates the cache coherence directory stating that the block in its local cache is dirty. The director then sends the write to disk.

When another host wants to read from the same block, it first looks in its local cache to see if the data exists. If it does not exist, it will look in the global cache coherence directory to see if another director has the write in its cache. If so, it will go to the director's cache to get the data.

Read Hit Distributed Device



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

36

- A host read request is serviced by a director. The director performs a look up in its local cache. If the data is there it is sent to the host. This is referred to as a local read hit.
- If the data is not in the local cache a look-up is performed in the global cache. If another director, in that cluster only, has the data in its cache, then it is sent via the local com to the servicing director and then to the host. This is referred to as a global read hit. It should be noted that in the case of a VPLEX Metro a read request will never be serviced from another cluster.

Read Miss Distributed Device



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

37

If the data the director is looking for is not in global cache, it is called a global cache miss. During global cache miss, the director must perform a back-end storage read.

1. A Read request is issued to the virtual volume from the host.
2. A look up is performed in the local cache of the ingress director.
3. On a miss, a look up is performed in the global cache.
4. On a miss, the requested data is read from the storage volume into the local cache.
5. The requested data is returned from the local cache to the host.

The diagram illustrates a **Stretched Host Cluster** architecture across two **FAILURE DOMAIN**s, labeled **A** and **B**.

- FAILURE DOMAIN A** contains a **HOST** (represented by a server icon).
- FAILURE DOMAIN B** contains a **HOST** (represented by a server icon) and a **HETEROGENEOUS** storage component (represented by a yellow cylinder).

The two failure domains are connected via a central layer consisting of:

- GLOBAL CACHE** (top section, dark blue bar).
- VIRTUALIZATION** (bottom section, light gray bar).

Connections and Data Flow:

- A green line connects the **HOST** in Failure Domain A to the **GLOBAL CACHE**.
- A green line connects the **HOST** in Failure Domain B to the **GLOBAL CACHE**.
- A green line connects the **HETEROGENEOUS** storage in Failure Domain B to the **GLOBAL CACHE**.
- A dashed line connects the two **HOST**s, labeled **Stretched Host Cluster**.
- A blue box labeled **DATA** is positioned between the **HOST** in Failure Domain B and the **GLOBAL CACHE**.
- A blue box labeled **DATA** is positioned between the **GLOBAL CACHE** and the **HETEROGENEOUS** storage.


Module: VPLEX Product and Technology Details

38

During a write hit:

1. A Write request is issued to the virtual volume from the host.
2. A look up is performed in the local cache of the ingress director.
3. A look up is performed in the global cache. Global cache updates the new location for the data.
4. Any prior data is invalidated in all cache locations.
5. The ingress director transfers the data to the local cache.
6. The new data is written through to the back-end storage.
7. The Write is acknowledged back to the host.

Write Distributed Device



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

EMC²

39

In the example, both sides are active-active. A host will send its write to one director. The director identifies the blocks to be written, and signals the other directors in the VPLEX Metro (both local and remote) that it now owns those blocks in cache. At that point each of them updates their private copy of the cache coherence table, noting which blocks will now be invalid within their own caches. Then both directors write the blocks to cache and through to disk. An acknowledgement is sent back from both sides to the director that sent the write. Lastly, the acknowledgement gets sent to the host. Writes are protected because they are written to disk before they are acknowledged.

Path Redundancy and Failure Handling



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

EMC²

40

Virtual volumes presented out of VPLEX can tolerate port failures by connecting the host to multiple directors and by utilizing multi-pathing software to control the paths. Here virtual volumes are presented out of multiple VPLEX front-end ports on different directors. This yields continuous data availability in the presence of port or director failure.

Path Redundancy Across Engines



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details



41

Virtual volumes presented out of VPLEX can also tolerate entire engine failures by connecting the host to VPLEX front-end ports on different engines. An engine could fail and the host would still be able to access its volumes. It is still best practice to connect the host to one A director and one B director. This will prevent a Data Unavailable situation during a non-disruptive upgrade.

VPLEX Witness



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

EMC²

42

For VPLEX Metro, an optional component called the VPLEX Witness can be deployed at a third location to improve data availability in the presence of cluster failures and inter-cluster communication loss. The VPLEX Witness is implemented as a virtual machine and requires a VMware ESX server for its operation. The customer host must be deployed in a separate failure domain from either VPLEX cluster to eliminate the possibility of a single fault affecting both a cluster and VPLEX Witness. VPLEX Witness connects to both VPLEX clusters over the management IP network.

VPLEX Witness observes the state of the clusters, and thus can distinguish between a outage of the inter-cluster link and a cluster failure. VPLEX Witness uses this information to guide the clusters to either resume or suspend I/O. VPLEX Witness provides seamless zero RTO fail-over for storage volumes in synchronous consistency groups.

Lesson Summary

During this lesson the following topics were covered:

- Define caching layers within a VPLEX system
- Describe I/O operation within VPLEX Local and Metro
- System interaction during reads and writes
- List the different VPLEX system and volume management options



This lesson covered an overview of VPLEX I/O operations, including caching, read and write interactions, and component redundancy.

Lesson: VPLEX Management Options

This lesson covers the following topics:


- VPLEX Management using the CLI
- VPLEX Management using Unisphere for VPLEX



This lesson covers the methods of managing the clusters within VPLEX.

VPLEX Management Server

- Enables management via VPLEX CLI and Unisphere for VPLEX
- Single or multiple (VPLEX Metro) clusters are managed through a single management server.



Management Server Rear View

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

45

A single VPLEX cluster is managed by a management server. In VPLEX Metro configurations, you can manage both clusters from a single management server. The management server coordinates data collection, VPLEX software upgrades, configuration interfaces, diagnostics, event notifications, and some director-to-director communication. The management server also forwards VPLEX Witness traffic between directors in the local cluster and the remote VPLEX Witness server. Redundant internal network IP interfaces connect the management server to the public network. Internally, the management server is on a dedicated management IP network that provides accessibility to all major components in the cluster.

VPLEX CLI Commands

- The VPLEX CLI supports all VPLEX operations
- CLI commands are divided across a hierarchical context tree structure
- Command contexts contain commands that can be accessed only from within that context.
- Global commands can be accessed within any command context



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

46

The CLI is divided into command contexts. The commands under each context are arranged in a hierarchical context tree. These commands can only be executed from the appropriate location in the context tree. Understanding the command context tree is critical to using the VPLEX command line interface effectively.

Except for the 'system-defaults/' sub-context, each of the sub-contexts contains one or more sub-contexts to configure, manage, and display sub-components. Command contexts have commands that can be executed only from that context. The topmost context is the root context, or "/".

Some commands are accessible from all contexts. These are referred to as 'global commands'.

Key Management CLI Root (/) Sub-contexts

Sub Context	Functionality
clusters/	Create and manage links between clusters, devices, extents, system volumes and virtual volumes. Register initiator ports, export target ports, and storage views.
connectivity/	Configure connectivity between back-end storage arrays, front-end hosts, local directors, port-groups and inter-cluster WANs.
distributed-storage/	Create and manage distributed devices and rule sets.
data-migrations/	Create, verify, start, pause, cancel, and resume data migrations of extents or devices.
monitoring/	Create and manage performance monitors.
recoverpoint/	Manage RecoverPoint options.




At the top of the command context hierarchy is the 'root' context. Although there are more, shown here are the key root-level sub-contexts where commands can be accessed to configure, manage and monitor VPLEX clusters, storage and host connectivity.

Using The VPLEX CLI


Log in for VPLEX CLI access

1. Log in using the Management Server 'service' user account
2. Type the **vplexcli** command to connect to the VPLEX command line interface
3. Use the exit command to exit the VPLEX command line interface from any context.



Shown here are the basic operations required to access the VPLEX command line interface.

VPLEX CLI Context Tree



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

49

The VPLEX CLI is based on a tree structure similar to the structure of a Linux file system. Fundamental to the VPLEX CLI is the notion of “object context”. The object context is determined by the current location or pwd within the directory tree of managed objects.

Many VPLEX CLI operations can be performed from the current context. However, some commands may require the user to cd to a different directory before running the command.

Navigating the CLI Context Tree

- Use the **cd** command to navigate between command contexts

```
VPlexcli:/> cd /clusters
VPlexcli:/clusters> cd cluster-1
VPlexcli:/clusters/cluster-1> cd exports
VPlexcli:/clusters/cluster-1/exports> cd initiator-ports
VPlexcli:/clusters/cluster-1/exports/initiator-ports ls

REGISTERED-0x10000000c9c878d4    REGISTERED-ESX31-30
REGISTERED-1ab5p1-0x10000000c9c878d5
UNREGISTERED-0x10000000c9d2f2de  UNREGISTERED-0x10000000c9d2f2df
UNREGISTERED-0x10000000c9d2f331
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

50

The **cd** command is used to navigate between command contexts. The current context is always displayed at the VPLEX command line interface prompt. Use the **cd** command with no arguments or followed by a space and three periods (i.e. **cd ...**) to return to the root context. Use the **cd** command followed by a space and two periods (i.e. **cd ..**) to return to the context immediately above the current context. To navigate directly to a context from any other context, use the **cd** command and specify the absolute context path.

Finding Current Location in the Context Tree

CLI features to help determine current position:

- The command prompt always displays the current context:

```
VPlexcli:/monitoring/directors/director-1-1-B/monitors>
```

- The **ls** command displays the sub-contexts immediately accessible from the current context:

```
VPlexcli:/> ls
clusters      data-migrations      distributed-storage
engines        management-server    monitoring
notifications  system-defaults
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details

51

The VPLEX CLI includes several features to help locate the current position in the context tree and determine what contexts and/or commands are accessible.

Finding The Current Location in the Context Tree (Cont'd.)

- The **ls -l** command displays more information about the current sub-contexts:

```
VPlexcli:/data-migrations> ls -l
Name Description
-----
device-migrations Contains all the device migrations in the system.
extent-migrations Contains all the extent migrations in the system.
```

- For contexts where the next lowest level is a list of individual objects, the 'ls' command displays a list of the objects:

```
VPlexcli:/clusters/cluster-1/exports/ports> ls
P000000003B2017DF-A0-FC00 P000000003B2017DF-A0-FC01
P000000003B2017DF-A0-FC02 P000000003B2017DF-A0-FC03
P000000003B3017DF-B0-FC00 P000000003B3017DF-B0-FC01
P000000003B3017DF-B0-FC02 P000000003B3017DF-B0-FC03
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

52

Of the context tree navigation aids available in the VPLEX CLI, some of the most useful are the **ls** and its variant, the **ls -l** command.

Please refer to the *EMC VPLEX CLI Reference Guide* and the *EMC VPLEX Administration Guide* to find out more about the CLI working environment and CLI command use and syntax.

Unisphere for VPLEX



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Product and Technology Details


EMC²

53

The VPLEX GUI, Unisphere for VPLEX, provides many of the features that the VPLEX CLI provides. The GUI is very easy to navigate and requires no knowledge of VPLEX CLI commands. Operations are accomplished by clicking on VPLEX icons and selecting desired values.

System Status on the navigation bar shows a graphical representation of your system. It allows you to quickly view the status of your system and some of its major components such as directors, storage arrays, and storage views. The cluster display also shows the size of the cluster configuration (single-engine, dual-engine, or quad-engine). Blue lights on the cluster display represent the number of engines in the cluster. For example, a quad-engine configuration display shows four blue lights, a dual-engine display shows two blue lights, and so on. System Status is the default screen when you log into the GUI. Also shown is the **Monitoring** menu. Here, you can monitor VPLEX cluster performance, provisioning jobs status, and general system health details.

Unisphere for VPLEX (Cont'd.)



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Product and Technology Details

EMC

54

The **Provision Storage** menu items assist administrators in provisioning VPLEX-managed storage and storage volumes with the required capacity, local availability, performance and protection options to hosts.


Unisphere for VPLEX (Cont'd.)



Mobility Central selections allow the more detailed setting of data mobility within and across VPLEX clusters. The **Support** menu provides the user with information and help around common system use and system support tasks for the VPLEX system.

VIPR SRM SUITE

Analyze historical performance trends



The Unisphere dashboard is great for real-time analysis, but what about long term reporting? That's where ViPR SRM Suite is one of your best options.

Here's an example of a performance trending report for VPLEX. Performance reporting for element managers like Unisphere for VPLEX typically report on a few key parameters over a short duration of time. ViPR SRM complements element management by tracking performance trends for a broader set of parameters across the data path for hours, days, weeks and months. This help storage teams understand historical behavior to simplify troubleshooting and change planning processes. Here you can see VPLEX front-end and backend cluster throughput (IOPS), and back-end bandwidth (KB/second).

Lesson Summary

During this lesson the following topic was covered:

- VPLEX Management using the CLI
- VPLEX Management using Unisphere for VPLEX



This lesson covered the methods of managing the clusters within VPLEX.

Module Summary

Key points covered in this module:

- List the ways VPLEX fits into the strategy for a virtualized data center
- Describe VPLEX architecture and product capabilities
- List the different VPLEX system management options



This module covered the topics listed here. If interested in learning more about any of these topics please refer to the Education Services catalog to discover what other VPLEX training we offer.

Module: VPLEX Storage Provisioning

Upon completion of this module, you should be able to:

- Explain host-to-storage connectivity through VPLEX
- Describe VPLEX virtual storage constructs
- Provision Virtual Volumes to hosts using
 - Advanced Provisioning
 - EZ Provisioning
 - Integrated storage provisioning (VIAS)
- Describe Storage Volume Encapsulation



This module focuses on managing the storage network using VPLEX.

Lesson: Storage Provisioning

This lesson covers the following topics:

- Configuring host-to-storage connectivity
- VPLEX virtual storage constructs




© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

2

This lesson covers configuring host-to-storage connectivity through VPLEX and understanding the VPLEX virtual storage constructs which support provisioning array storage to hosts.

Host Connectivity Using VPLEX Front-End Ports



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

3

VPLEX uses front-end (FE) Fibre Channel ports, four FE ports per director of a VPLEX Engine, to connect to hosts in a Fibre Channel SAN. In a dual-SAN Fabric, best practice is to cross-connect each director's FE ports into SAN A and SAN B as shown. Hosts HBA ports are connected across both SANs according to the number of ports per HBA. This is the best method to provide redundant pathing between hosts and the SAN Fabrics. In this example, each host uses dual-ported HBAs. To lessen the possibility of a single HBA card failure affecting connectivity to a single SAN Fabric, one port from each physical HBA is connected to each SAN Fabric.

Although the diagram shows a single VPLEX Engine, host connections should span across VPLEX Engines as well as across Directors of a VPLEX Engine as shown.

Determine WWNs of VPLEX FE Ports

- Gather VPLEX FE WWPNS for Zoning to Host Initiators

```
VPlexcli:/> ls -l /engines/**/ports
```

```
/engines/engine-1-1/directors/director-1-1-A/hardware/ports:
```

Name	Address	Role	Port	Status
A0-FC00	0x50001442a0633000	front-end	up	
A0-FC01	0x50001442a0633001	front-end	up	
A0-FC02	0x50001442a0633002	front-end	up	
A0-FC03	0x50001442a0633003	front-end	up	

- Port information output will be grouped by VPLEX engine and director



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

4


Determine the VPLEX front-end and WAN-COM port WWNs for use in configuring SAN connectivity and Zoning to support VPLEX-to-hosts and VPLEX cluster-to-cluster communications. Use the VPLEX CLI **ls -l** command to list the contents of **/engines/**/ports** of all VPLEX engines and directors.

Storage View Establishes Host-to-Storage Connectivity

- A Storage View consists of selected:

- ◆ Registered Initiator
- ▲ VPLEX FE ports
- Virtual volumes

- Objects in a Storage View establish storage communications



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

EMC²


5

A Storage View is a combination of registered Initiators (host HBA ports), VPLEX front-end ports and virtualized storage volumes. It is used to control a single host's access to the selected VPLEX virtual storage volumes. In this example, the VPLEX storage volumes contained in this Storage View are Datastores to three VMware ESX servers that are clustered together and require simultaneous access to all of the Datastores. The standalone server shown is not a server related to the Datastores and therefore is not included in the Storage View that enables access to these two VPLEX storage volumes.

To export VPLEX storage, you must first create a storage view, and then add VPLEX front-end ports and VPLEX volumes to the view. Virtual volumes are not visible to hosts until they are in a storage view with associated ports and initiators

A registered initiator can be in more than one Storage View, and a VPLEX FE port can be in more than one Storage View, the unique/particular combination of a specific <initiator><FE_Port> pair can only be in one Storage View.

VPLEX Virtual Storage Constructs



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

6

Back-end storage arrays are configured to present LUNs to VPLEX through the SAN. Each presented back-end LUN maps to one VPLEX **Storage Volume**. Storage Volumes are initially in the 'unclaimed' state. Unclaimed Storage Volumes may not be used for any purpose within VPLEX other than to create metavolumes that are for system internal use only.

Once a Storage Volume has been claimed within VPLEX, it may be carved into one or more contiguous **Extents**. A single extent may map to an entire Storage Volume. However, it cannot span multiple Storage Volumes.

A VPLEX **Device** is the entity that enables RAID implementation across multiple Extents or other Devices. VPLEX supports RAID-0, RAID-1, and RAID-C. also 1-1 Mapping.

Raid-0 can stripe data across multiple extent/device constructs. When creating a Raid-0 device, if more than 1 extent is chosen, VPLEX creates a Device for each extent chosen and a top level device that is the sum of the size of the extents. An Example would be if three 2GB extents were selected then the top level device would be 6GB. Vplex would stripe data across the "parts" of the top level device. The stripe depth, which is how much data is written to each part of the Top Level Device at a time before moving to the next part, is configurable in the CLI.

Raid-C is similar, but it concatenates the data for each "part" of the Top Level Device.

Raid-1 Mirrors two Device extent/device constructs. The Top Level Device is the same size.

As was just presented, a **Storage View** is the masking construct that controls how one or more VPLEX Virtual Volumes is exposed through VPLEX front-end ports to host initiators.

Once a Storage View is properly configured as described and operational, the host should be able to detect and use **Virtual Volumes** after initiating a bus-scan on its HBAs. Every front-end path to a virtual volume is an active path and the current version of VPLEX presents volume with the product ID 'Invista.' The host requires supported multi-pathing software in a typical high-availability implementation.

Creating a Storage View– Select Initiators and VPLEX FE Ports

The screenshot displays the 'cluster-1: Create Storage View' wizard. The '1 Select Initiators' step is active, showing a list of 'Available Registered Initiators' and a 'Selected Registered Initiators' list. A blue arrow points from the 'Available' list to the 'Selected' list. The '2 Select Ports' step is also visible, showing 'Available Ports' and 'Selected Ports'. A 'Single Path Warning' dialog box is overlaid on the right, with a red arrow pointing to it from the 'Selected Ports' list. The dialog box contains a yellow warning icon and the text: 'Potential single path configuration. Please check that you have at least two initiators and two target-ports for this view.' The wizard interface includes a sidebar with steps 1-6, a main area with 'Available' and 'Selected' lists, and a 'Next >' button. The 'Single Path Warning' dialog has a 'Close' button.

Unisphere for VPLEX is the most efficient and convenient method for creating and managing Storage Views. Use the **Create Storage View** wizard to create a storage view, and add initiators, ports and Virtual Volumes to the view. When you create a storage view with all required components, it automatically becomes active. When a storage view is active, hosts can see the storage and begin I/O to the Virtual Volumes in the Storage View.

NOTE: Although there may be situations that do not require it, far more typical is the need for multiple paths between host initiators and VPLEX FE ports. Ensure that there is path redundancy in the host-to-storage access configuration where this is required.

When initiators are connected to VPLEX FE ports directly or through the SAN, VPLEX automatically discovers them and populates the Initiators View. After discovery the initiators must be registered with VPLEX before they can be added to a Storage View. Registering an initiator is the process of giving a meaningful name (typically the host's DNS name) to the HBA port's pWWN to allow easy identification by an administrator. In the Initiators View, select discovered initiators (prefixed with 'UNREGISTERED_') and then click the Register button. If VPLEX did not automatically discover an initiator it must be manually registered. Important to note is that each port of a host's HBA/HCA must be registered as a separate initiator.

Creating a Storage View– Select Virtual Volumes and Logical LUN Numbers

The screenshot displays the 'cluster-1: Create Storage View' wizard. The left sidebar shows the progress: 1. Select Initiators, 2. Select Ports, 3. Select Virtual Volumes (highlighted), 4. Edit LUNs (optional) (highlighted), 5. Review & Finish, and 6. View Results. The main window is divided into two panes. The left pane, 'Available Unexported Virtual Volumes', lists several volumes with their consistency groups. The right pane, 'Selected Unexported Virtual Volumes', shows the volume 'device_student-1_0...' selected. Below the panes are 'Add', 'Add All', '< Remove', and '< Remove All' buttons. A 'Next >' button is at the bottom. A secondary window titled 'cluster-1: Create Storage View' is overlaid, showing the '4 Edit LUNs (optional)' step. It has radio buttons for 'Auto-assign LUN numbers starting with 0' (selected) and 'Manually enter LUN numbers'. Below, a table shows 'Selected Virtual Volumes' with columns for 'Virtual Volume Name' and 'LUN'. The first row shows 'device_student-1_04D7_1_vol' with LUN '0'. 'Move Up' and 'Move Down' buttons are at the bottom right of this window. A text box explains that LUN numbers are automatically assigned from 0 to 16383 and provides instructions for manual assignment.

© Copyright 2015 EMC Corporation. All rights reserved. Module: VPLEX Storage Provisioning

Select one or more Virtual Volumes to add to the new Storage View. An optional step is to check and assign specific LUN numbers to the volumes that they will present to the host. By default, the Storage View assigns logical LUN numbers beginning with zero. You can change the numbering as required using this wizard step.

Creating a Storage View– Review, Commit and Activate

cluster-1: Create Storage View

- ✓ * 1 Select Initiators
- ✓ * 2 Select Ports
- ✓ 3 Select Virtual Volumes
- ✓ 4 Edit LUNs (optional)
- ✓ * 5 Review & Finish
- 6 View Results

Review the initiators, ports and virtual volumes in the storage view. If satisfied, click Commit.

Note that if the storage view contains all required components, it will automatically become active when you click Commit.

5 Review & Finish

Storage View Name: Student1_SV1

Status (once created): Active

Components:

- ✓ Initiators
- ✓ Ports
- ⓧ ✓ Virtual Volumes

Name	LUN
device_student-1_04D7_1_vol	0

Commit

Expand each field to review members of the group

Review the initiator, VPLEX FE port and Virtual Volume selections for the new Storage Group. You can expand each of the Initiator, Ports and Virtual Volumes fields to review the members. Click **Commit** to create and activate the new Storage View.

Lesson Summary

During this lesson the following topics were covered:

- Configuring host-to-storage connectivity
- VPLEX virtual storage volume constructs



This lesson covered configuring host-to-storage connectivity through VPLEX and the VPLEX virtual storage constructs which support provisioning array storage to hosts.

Lesson: Storage Provisioning Methods

This lesson covers the following topic:

- Preparing storage for host access using
 - Advanced provisioning
 - EZ provisioning
 - VIAS Integrated array service-based provisioning




This lesson covers configuring VPLEX storage into volumes and provisioning and integrating these with the hosts that will access them, using the methods shown here.

VPLEX Storage Provisioning Overview

Three methods of storage provisioning through VPLEX for host access:

- Advanced provisioning
- EZ provisioning
- Integrated array service-based provisioning (VIAS)



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

12


To begin using VPLEX, you must provision storage so that hosts can access that storage. There are three ways to provision storage on VPLEX:

- Advanced provisioning
- EZ provisioning
- Integrated storage provisioning using the VPLEX Integrated Array Service (VIAS)

Note: EMC recommends using the Unisphere for VPLEX GUI to provision storage.

Consistency Group Use Considerations

- Consistency Groups allow a set of volumes to share the same properties
- Distributed Volumes
 - Allow for VPLEX Witness
- Remote Access
- RecoverPoint



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

EMC¹

13

VPLEX consistency groups aggregate volumes to enable the application of a common set of properties to the entire group. Consistency groups aggregate up to 1,000 virtual volumes into a single entity that can be managed easily.

VPLEX supports up to 1024 synchronous consistency groups.

A synchronous consistency group:

- Contains up to 1000 virtual volumes.
- Contains either local or distributed volumes, (but not a mixture of both).
- Contains volumes with either global or local visibility.
- Uses write-through caching (known as synchronous cache mode in the VPLEX user interface).

Consistency Groups in Unisphere

EMC Unisphere for VPLEX v5.4

System Status Monitoring Provision Storage Mobility Central Support

Provision Storage << cluster-1 : Consistency Groups

Find < Previous Next >

Name	Operational Status	Detach Rule	# Volumes	Visibility	Storage at Clusters
cl-2-test	OK	N/A	1	cluster-1	cluster-1
cluster-2a	OK	Winner:cluster-2 (5 seconds)	0	Both Clu..	cluster-2
esx-74	OK	Winner:cluster-2 (5 seconds)	1	Both Clu..	Both Clusters
test	OK	Winner:cluster-1 (5 seconds)	2	Both Clu..	Both Clusters
test33	OK	Winner:cluster-1 (5 seconds)	1	Both Clu..	cluster-1

Create Delete

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

14

The Consistency Groups screen shows all the consistency groups in a cluster. Use this screen to quickly view information about a consistency group. By default, consistency groups are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). To change the direction of the default sort, click on the column header. This view shows both distributed and non-distributed consistency groups.

Advanced Provisioning Overview

- Use the advanced provisioning method to granularly “slice” **Storage Volumes** into **Extents**.
- Use one or more Extents to create a **Device**
- Use Devices to create **Virtual Volumes**
- Present Virtual Volumes to hosts for access

Steps required for Advanced Provisioning:

1. View and claim available storage
2. Create Extents from Storage Volumes
3. Create RAID-0, RAID-1, RAID-C or 1:1 mapping of Extents to Devices
4. Create virtual volumes on the Devices.
5. Register initiators that access VPLEX storage.
6. Create Storage Views that include virtual volumes, initiators, and VPLEX ports
7. Setup host access to the virtual volumes.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Storage Provisioning

15

Advanced provisioning allows the user to slice (use less than the entire capacity) Storage Volumes into Extents. You can then use one or more of these Extents to create Devices, and then virtual volumes on these Devices. Advanced provisioning requires you to perform each of these steps individually, in the order listed in the step-by-step instructions. Use this method when you want to slice Storage Volumes and perform other advanced provisioning tasks such as creating complex Devices. The GUI also provides wizards that allow you to create each of the required storage objects.

Advanced Provisioning - Claim Storage


- Select either Cluster-1 or Cluster-2



Start advanced storage provisioning by selecting the Provision Storage item from the Unisphere command ribbon at the top.

Advanced Provisioning - Claim Storage (Cont'd.)

- Select and Claim Back-end Array Storage Volumes



Navigate to the Claim Storage wizard to begin the process of claiming back-end array LUNs into the VPLEX cluster.

Advanced Provisioning - Claim Storage (Cont'd.)

- Claim Storage wizard: Select the Back-end Array

The screenshot shows a web-based wizard interface for 'cluster-1: Claim Storage'. On the left, a vertical sidebar contains three steps: '* 1 Select Array' (highlighted), '* 2 Specify Claiming Details', and '3 ...'. Below the sidebar, a yellow box contains the text: 'Verify the selected array. If necessary, click the drop down box to select a different array.' The main content area is titled '1 Select Array' and 'Select Array'. It includes the text 'Claim Storage on this Array:' followed by a label 'Array: *' and a dropdown menu. The dropdown menu is open, showing the selected value 'EMC-SYMMETRIX-195901096' and another identical option below it. A 'Next >' button is located at the bottom right of the main content area.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

18

Select the desired back-end arrays and click next.

Advanced Provisioning - Claim Storage (Cont'd.)

- Claim Storage wizard: Specify Claiming Details

cluster-1: Claim Storage

✓ * 1 Select Array

* 2 Specify Claiming Details

* 3 Select Volumes to Claim

* 4 Review Storage to Claim

5 View Results

To customize the claimed storage volume names, select the checkbox and provide a Name Mapping file.

For thinly-provisioned storage volumes, select Use Thin Rebuilds to preserve the thinness of the storage volume during rebuilds.

The Storage type is limited to 15 characters and is included in the storage volume name.

2 Specify Claiming Details

Use Name Mapping File

Found 3 unclaimed storage volumes that do not require a Name Mapping file. To customize the storage volume names, select the Use a Name Mapping file checkbox and provide a Name Mapping file.

☐ Use a Name Mapping file

Select Rebuild Type

☒ Use Thin Rebuilds

☐ Do not use Thin Rebuilds

Select Storage Type (optional)

Enter an Identifier for the storage that indicates the category or type of data to be stored. For example, A - Z, 0 - 9, or Tier1:

Storage Type:

Next >

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

19

Reading the information in the left-hand pane of the wizard details what is being accomplished in this step. In this example, three unclaimed array LUNs exist on the selected back-end array. We named the Storage Volume "Student_SV".

Thin provisioning allows storage to migrate onto a thinly provisioned storage volumes while allocating the minimal amount of thin storage pool capacity. Thinly provisioned storage volumes can be incorporated into RAID 1 mirrors with similar consumption of thin storage pool capacity. VPLEX preserves the unallocated thin pool space of the target storage volume by detecting zeroed data content before writing, and suppressing the write for cases where it would cause an unnecessary allocation. VPLEX requires the user to specify thin provisioning for each back-end storage volume. If a storage volume is **thinly provisioned**, the "thin-rebuild" attribute must be to "**true**" either during or after claiming.

- Claim Storage wizard: Select Volumes to Claim



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

20

Advanced Provisioning - Claim Storage (Cont'd.)

- Claim Storage wizard: Review and Claim Storage Volumes

cluster-1: Claim Storage

✓ * 1 Select Array
✓ * 2 Specify Claiming Details
✓ * 3 Select Volumes to Claim
* 4 Review Storage to Claim
5 View Results

Review the storage volumes to claim.

You can change the name used to claim a storage volume now, or change the name later in the Storage Volume Properties dialog box.

If you enter a name into the Base Name field and click on the Apply button, the name will be prefixed to any unmodified storage volume names.

Click Commit to claim the storage volumes.

4 Review Storage to Claim

Array Name: EMC-SYMMETRIX-195901096 Base Name: Apply

There are 1 unclaimed storage volumes to be claimed:

Storage Volume	Thin Rebuild	To Be Named
VPD83T3:60000970000195901096...	Yes	Student1_SV1

cluster-1: Claim Storage

✓ * 1 Select Array
✓ * 2 Specify Claiming Details
✓ * 3 Select Volumes to Claim
✓ * 4 Review Storage to Claim
5 View Results

5 View Results

Claim Storage Volumes Results

Claimed Storage Volume Name	Result	Details
Student1_SV1	✓	Successfully claimed storage volume.

Commit Close EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

21

Review the storage LUN(s) selected to be claimed and set the name that will be used to identify it. Change the name if desired. Here, we changed the name from its original value to "Student1_SV1". After review click Commit to effect the Storage Volume claim. Finally, view the results and then click Close.

storage-volume claimingwizard Command

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> claimingwizard  
-cluster cluster-1 --set-tier =(Clar0400, L) --thin-rebuild  
Found unclaimed storage-volume VPD83T3:6006016091c50e004f57534d0c17e011  
  vendor DGC : claiming  
and naming Clar0400L_LUN82.  
Found unclaimed storage-volume VPD83T3:6006016091c50e005157534d0c17e011  
  vendor DGC : claiming  
and naming Clar0400L_LUN84.  
Claimed 2 storage-volumes in storage array Clar0400L  
Claimed 2 storage-volumes in total.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Installation and Setup

22

The CLI command `storage-volume claimingwizard` finds unclaimed storage volumes, claims them, and names them appropriately. This command can be used to claim and name many Storage Volumes with a single command.

Storage volumes must be claimed, and optionally named before they can be used in a VPlex cluster. Storage tiers allow the administrator to manage arrays based on price, performance, capacity and other attributes. If a tier ID is assigned, the storage with a specified tier ID can be managed as a single unit. Storage volumes without a tier assignment are assigned a value of 'no tier'.

Optional arguments

`[-c|--cluster] cluster` - Cluster on which to claim storage.

`[-f|--file] file,file...` - List of one or more files containing hints for storage-volume naming, separated by commas.

`[-d|dryRun]` - Do a dry-run only, do not claim and name the storage volumes.


`[-t|--set-tier] list` - Set a storage tier identifier per storage array in the storage-volumenames. Type multiple arrayName, tier-character pairs separated by commas. Storage tier identifiers cannot contain underscores.

`[--force]` - Forces a successful run of the claimingwizard. For use with non-interactive scripts.

`--appc` - Make the specified storage volumes 'application consistent'. Prevents data already on the specified storage volume from being deleted or overwritten.

Advanced Provisioning – Create Extents


- Select Storage Volume and Create Extents



We will now create the next level of volume construct, an extent. Selecting Storage Volumes under Physical Storage will display all Storage Volumes in the selected VPLEX cluster. We selected the Storage Volume we just created, and will create an extent from a 20GB 'slice' of the total 40GB available in the Storage Volume.

Advanced Provisioning – Create Extents (Cont'd.)

- Wizard: Select Storage Volumes and Create Extents



The Storage Volume we selected automatically appears in the Selected Storage Volumes column. Additional volumes may be added if desired.

A Capacity Calculator is available to assist in sizing to ensure achieving the desired final Virtual Device capacity. In this example, we specified that we desired a single 20GB Virtual Device. Since the total capacity of the selected Storage Volume is 40GB, a single extent will be created out of a 20GB “slice” of the 40GB Storage Volume.

Note the instructions available in the left-hand pane. Click Next.

Advanced Provisioning – Create Extents (Cont'd.)

- Create Extents wizard: Specify Extent Capacity / Extent Creation Result

The extent name shows which extent this is among any other Extents that may exist on the Storage Volume

Storage Volumes	Largest Free Chunk	Thin Rebuild	Use	Extent Capacity
Student1_SV1	40.00 GB	true	☐	400

Extent Creation Results:


Extents	Result	Details
extent_Student1_SV1_1	✓	Created Extents

The Extent Capacity listed in the right-most column for the Storage Volume will by default equal the largest chunk of free capacity on the Storage Volume. Since in this example we only desire a single 20GB volume, we'll need to alter the Extent capacity by manually selecting the field and typing in "20G". The Total Extent Capacity and the Total Capacity needed will reflect that a single extent of 20GB will be sliced from the 40GB Storage Volume. Click Finish and then Commit on the screen that follows. Finally, review the results.

Note the name given to the new extent is similar to the Storage Volume name followed by a "_1". This indicates that the new extent is the first and currently only extent sliced from this Storage Volume.

Advanced Provisioning – Create Devices


- Create Devices From Extents



The next step is to take one or more Extents and form a Device with them. We will use the *Create Devices* wizard to accomplish this task. Selecting 'Devices' will display the current Devices allocated in the target VPLEX cluster. Click 'Create' to use the Create Devices wizard to begin creating the new Device.

Advanced Provisioning – Create Devices (Cont'd.)

- Create Devices wizard: Select Device Type / RAID Level Attribute



© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Storage Provisioning

27

The wizard first requires selecting the Device type attribute. Here, we can select a RAID protection or performance attribute or simply specify a one-to-one mapping of one extent used to create a single Device. In this example, we've chosen to create a one-to-one mapping using the extent just created to form the new Device. Click Next.

Advanced Provisioning – Create Devices (Cont'd.)

- Create Devices wizard: Select Extent(s) to Use for New Device



EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

28

Select the Extent to be used for the new Device. In this example, we selected the 'Student1_SV_1' extent created earlier. Click Next.

Advanced Provisioning – Create Virtual Volume

- Create Devices wizard: Create Virtual Volume from Device

cluster-1: Create Devices

✓ * 1 Select Device Type
✓ * 2 Select Extents
* 3 Create Virtual Volumes
* 4 Review & Finish
5 Results

Click Yes to automatically create a virtual volume on each device, or No to create the virtual volume later (for example, if the device will be used to create a distributed device).

3 Create Virtual Volumes

Automatically create a virtual volume on each device?

☒ Yes
☐ No

Next >

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: VPLEX Storage Provisioning

29

Choose to either create a new Virtual Volume from the new Device now, or to wait. You should wait until later to create a Virtual Volume if the new Device will be used to create a Distributed Device. In this example, we are not using a Distributed Device, so we will choose to create the Virtual Volume now. Click Next.

Advanced Provisioning – Create Virtual Volume

- Create Devices wizard: Review and Create New Virtual Volume



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

30

Review the information about the new Device and its Extents, and about the Virtual Volume that will be created from them. Click Commit to create the new Virtual Volume. After creation, review the results. Click Close.

Advanced Provisioning – Map Virtual Volume to Host (Cont'd.)

- Virtual Volume wizard: Allocate Virtual Volume to Storage View

EMC Unisphere for VPLEX v5.4

System Status Monitoring Provision Storage Mobility Central Support

cluster-1 : Storage Views

Name	Operational Status	Active	# Volumes	Initiator	# Ports
Student1_Host	OK	Yes	0	2	2
esx-11_S1	OK	Yes	3	2	2
esx-70SV	OK	Yes	2	2	4
esx_21	OK	Yes	2	2	7
view-3	OK	Yes	0	2	2

Components of Selected Storage View:

Student1_Host

- Initiators
- Ports
- Virtual Volumes

Show: ☐ LUN ☐ VPD ID ☒ Capacity


Add/Remove Virtual Volumes

Select the information columns to be shown in this pane

The final step to use VPLEX storage is to allocate Virtual Volumes to the hosts that will access them. To do this in the Advanced Provisioning method, in the Hosts section, select Storage Views. Then select the host(s) that the Virtual Volume will be mapped to. In this example, we selected a host named 'Student1_Host' to access the new Virtual Device we just prepared for it. Click Add/Remove Virtual Volumes.

Advanced Provisioning – Map Virtual Volume to Host (Cont'd.)

- Virtual Volume wizard: Export Virtual Volume to Storage View



Add the desired Virtual Volumes to the Storage View of the selected host. In this example, we've chosen the Virtual Volume just created. Click OK. Verify that the volume has been successfully added. Click Close. The Virtual Volume is now connected to the specified host and accessible through the host's initiators and VPLEX ports included in the Storage View for this host.

EZ Provisioning Overview

Use the EZ provisioning wizard to provision storage from selected Storage Volumes.

1. Select or creating a consistency group for the volumes
2. Select the type of volume to create (distributed or local)
3. Select protection options, storage, and optionally expose the virtual volume to hosts



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning



33


EZ provisioning is a simple method of provisioning that makes provisioning tasks easy. It is available only in the Unisphere GUI.

EZ provisioning automatically creates a virtual volume with a one-to-one mapping to a selected Storage Volume. Use EZ provisioning to quickly create a virtual volume that uses the entire capacity of the Storage Volume. In EZ provisioning, you select storage arrays and define how you want them to be used, protected, and presented to hosts.

EZ-Provisioning uses the *Provision from Storage Volumes* wizard to guide you through provisioning from selected Storage Volumes on the array. The wizard eliminates the individual steps required to claim storage, create Extents, create Devices, and then virtual volumes on those Devices. The wizard creates one volume at a time, and each volume uses the full capacity of its selected Storage Volume. You can also use this wizard to preserve existing data on a Storage Volume and make that data visible to a host through a virtual volume.

EZ Provisioning Overview (Cont'd.)

- Select an Existing or Create a New Consistency Group



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

34

The first step is to select an existing, or create a new consistency group that is to assume the new storage capacity.

EZ Provisioning Overview (Cont'd.)

- Select Protection Options and Type Volume Name

Provision From Storage Volumes

1 Select Consistency Group

2 **Protection Options**

3 Storage Volumes

4 Expose To Hosts

5 Review

Protection Options: Select high availability options and enter a name for the volume.

Source Cluster: + cluster-1

High Availability: ☐ Enable Mirroring at cluster-1

Volume Name: Student1_newVV

Consistency Group: Student2-PER



cluster-1 cluster-2

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

EMC²

35

Select the source VPLEX cluster that will provide the storage capacity from the back-end arrays connected to it. Also select the appropriate protection and data synchronization attributes for the new storage capacity.

EZ Provisioning Overview (Cont'd.)

- Select the Storage Array and the Desired Volume

Provision From Storage Volumes

1 Select Consistency Group

2 Protection Options

3 Storage Volumes

4 Expose To Hosts

Storage Volumes: Select a storage volume on cluster 'cluster-1'.

Source Storage:

EMC-SYMMETRIX-195901096

Available Storage Volumes

Name	Capacity	Use
Symm1096_00A0	3.00 GB	claimed
Symm1096_00A1	3.00 GB	claimed
Symm1096_00A2	3.00 GB	claimed
Symm1096_00A3	3.00 GB	claimed
VPD83T3:6000097000019590109653303...	3.00 GB	unclaimed
VPD83T3:6000097000019590109653303...	3.00 GB	unclaimed
VPD83T3:6000097000019590109653303...	5.00 GB	unclaimed
VPD83T3:6000097000019590109653303...	5.00 GB	unclaimed

EMC

Select a back-end storage array and LUN connected to the source VPLEX cluster. The selected physical array's LUN data will be copied onto the target volume in the target VPLEX cluster. Back-end array LUNs in either the 'claimed' or 'unclaimed' state may be used.

EZ Provisioning Overview (Cont'd.)

- Select the Storage View

The screenshot shows a web-based provisioning interface. On the left is a vertical sidebar with four steps: 1. Select Consistency Group, 2. Protection Options, 3. Storage Volumes, and 4. Expose To Hosts. Step 4 is highlighted with a white arrow. The main panel is titled 'Expose To Hosts: Expose storage to hosts through storage views.' It contains a section 'Storage View at cluster-1' with a red asterisk. Below this, there are two radio buttons: 'Existing Storage View' (which is selected) and 'None'. To the right of the 'Existing Storage View' radio button is a dropdown menu showing 'Student1_S'.



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

37

Select the hosts connected to each VPLEX cluster that will have access to the new storage being presented.

EZ Provisioning Overview (Cont'd.)

- Review the Accuracy of the Selected Storage Configuration...

The screenshot shows the 'Review' step of the 'Provision From Storage Volumes' process. On the left, a navigation pane lists five steps: 1. Select Consistency Group, 2. Protection Options, 3. Storage Volumes, 4. Expose To Hosts, and 5. Review (which is highlighted). The main area displays a review of the configuration. At the top, it says 'Review: Review your selections. Click Provision Storage to submit the provisioning request.' Below this, the 'Consistency Group' is 'Student2-PER' with a 'Total Capacity to Provision: 3 GB'. A diagram shows two storage clusters, 'cluster-1' and 'cluster-2', each represented by a server icon and a disk icon. Below the diagram, the following details are listed: Volume Name: Student1_newVV, Storage View: Student1_S, Storage Array: EMC-SYMMETRIX-195901096, and Storage Volume: VPD83T3:60000970000195901096533030304134. At the bottom right, there are two buttons: 'Provision Storage' and 'Start Over', both of which are circled in red. The EMC logo is in the bottom right corner. The footer contains the copyright notice '© Copyright 2015 EMC Corporation. All rights reserved.', the module name 'Module: VPLEX Storage Provisioning', and the page number '38'.

Provision From Storage Volumes

1 Select Consistency Group

2 Protection Options

3 Storage Volumes

4 Expose To Hosts

5 Review

Review: Review your selections. Click Provision Storage to submit the provisioning request.

Consistency Group: **Student2-PER**
Total Capacity to Provision: 3 GB

cluster-1 cluster-2

Volume Name: Student1_newVV
Storage View: Student1_S
Storage Array: EMC-SYMMETRIX-195901096
Storage Volume: VPD83T3:60000970000195901096533030304134

Provision Storage Start Over

EMC


© Copyright 2015 EMC Corporation. All rights reserved. Module: VPLEX Storage Provisioning 38

Review your selection for accuracy and then provision the storage as configured using EZ Provisioning. When Provision storage is selected the desired tasks will be completed.

This is a quick method to provision storage to a host from a claimed or unclaimed Storage Volume.

VIAS Provisioning Overview

- The VIAS integrated storage feature uses Array Management Providers (AMPs)
- AMPs allow VPLEX to provision storage directly from storage pools on supported EMC VMAX and VNX arrays



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

39

Using the VPLEX Integrated Services (VIAS) feature, you can create virtual volumes from pre-defined storage pools on supported EMC VNX and VMAX arrays. These storage pools are made visible to VPLEX through Array Management Providers (AMPs). AMPs are back-end array resident and expose native back-end array functionality on supported arrays.

Note: In release v5.4, VIAS uses the Storage Management Initiative Specification (SMI-S) provider to communicate with the arrays that support integrated services to enable

provisioning. For more information about registering AMPs and provisioning from storage pools, refer to the *VPLEX Administration Guide*.

VIAS Provisioning Overview (Cont'd.)

- Storage Pools must already exist on the back-end array
- The AMP must be registered with VPLEX
- Fully understand the back-end array functional limitations and restrictions prior to allocation
- Fully understand how the back-end Storage Pool is configured on the array prior to allocation

Steps required for VIAS Provisioning:

1. Select a Consistency Group for the Virtual Volume or create a new Consistency Group.
2. Add VPLEX volume service attributes and specify the number and size of Virtual Volumes to create. In the case of a distributed volume, select synchronization options and a rule set
3. Select the Storage Pool to be used to create the VPLEX Virtual Volumes and/or mirrors (if applicable).
4. Add the Virtual Volume to a Storage View (optional).
5. Review your selections and submit the provisioning request

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

40

When you create Virtual Volumes using VPLEX Integrated Services, the entire process (from Storage Volume, to Extent, to Device) is done automatically. The encapsulation is a 1-1 relationship of storage objects from the Storage volume to the Virtual volume. In the Virtual Volume and Storage Volume screens, the Integrated Services attribute indicates that these objects were created using VIAS. You can continue to perform all the actions on a VIAS volume that you normally perform on a non-VIAS virtual volume. Note that the underlying Extent and Device do not show the Integrated Services attribute. You can create a maximum of 10 virtual volumes at a time and the minimum capacity of each Storage Volume is 150 MB.

VIAS Storage Provisioning Method (Cont'd.)

- Select Provision from Pools Menu



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning


EMC²

41


Using the VPLEX Integrated Services (VIAS) feature, you can create virtual volumes from pre-defined storage pools on supported EMC VNX and VMAX arrays. These storage pools are made visible to VPLEX through Array Management Providers (AMPs). AMPs are back-end array resident and expose native back-end array functionality on supported arrays.

VIAS Storage Provisioning Method (Cont'd.)

- VIAS wizard: Select Consistency Group



A warning will be displayed if there are no back-end arrays with AMPs registered to a VPLEX cluster:



EMC²

Module: VPLEX Storage Provisioning

42

Select the Consistency Group within the appropriate VPLEX Cluster(s) to use or create a new one. In this example we selected the 'test33' CG. Click Next.

VIAS Storage Provisioning Method (Cont'd.)

- VIAS wizard: Volume and Protection Attributes

EMC Unisphere for VPLEX v5.4

System Status Monitoring Provision Storage Mobility Central Support

Provision From Pools

1 Select Consistency Group

2 **Protection Options**

3 Storage Pools

4 Expose To Hosts

5 Review

6 Results

Protection Options: Select high availability options and configure volumes.

High Availability: ☐ Enable Mirroring at cluster-1

Number Of Volumes: 1

Capacity Per Volume: 5 GB

Total Capacity: 5 GB

Volume Base Name: Student1_VIAS

Consistency Group: test33

cluster-1 cluster-2

Next >

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

EMC²

43

Select the number of volumes to create, the size and a volume name and the protection attributes. In this example, we select the name 'Student1_VIAS'. Click Next.

VIAS Storage Provisioning Method (Cont'd.)

- VIAS wizard: Select the Back-end Array and Storage Pool

EMC Unisphere for VPLEX v5.4

System Status Monitoring Provision Storage Mobility Central Support

Provision From Pools

1 Select Consistency Group

2 Protection Options

3 Storage Pools

Storage Pools: Select a storage pool at cluster 'cluster-1'.

EMC-SYMMETRIX-195901096

Select Array

EMC-SYMMETRIX-195901096

		Free	Total
Test	Thin	19.90 GB	20.00 GB
diskgroup1	Thin	1.19 TB	1.91 TB
diskgroup2	Thin	16.00 GB	828.00 GB
diskgroup3	Thin	47.50 GB	828.00 GB
diskgroup4	Thin	246.00 GB	1005.00 GB

Next >


EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: VPLEX Storage Provisioning 44

Select the back-end array and the Storage Pool. In this example, we selected the 'test' thin-provisioned pool which has 19.90GB of free space. Click Next.

VIAS Storage Provisioning Method (Cont'd.)

- VIAS wizard: Assign the New Virtual Volume to the Selected Host



Assign the new Virtual Volume being created to the desired host(s) Storage View. In this example, we selected the host we created earlier, 'Student1_Host'. Click Next. Review the selections and click Provision Storage.

VIAS Storage Provisioning Method (Cont'd.)

- VIAS wizard: Monitor Provisioning Job Progress Until Complete

The image displays three sequential screenshots of the EMC Unisphere for VPLEX v5.4 interface, illustrating the steps to monitor a provisioning job:

- Step 1:** The 'Provision From Pools' wizard screen. A message states: 'Virtual volumes provisioning job started'. A red box highlights this message, and a red arrow points to the 'View Job Status' link.
- Step 2:** The 'Provision Jobs Status' table. The job 'Provision_3_15-04-15-18-53-34' is shown with a status of 'In Progress'. A red box highlights the 'Status Details' column, which shows 'Exposing 1 storage volume(s) on storage array(s)'.
- Step 3:** The 'Provision Jobs Status' table. The same job is now shown with a status of 'Completed'. A red box highlights the 'Status Details' column, which shows 'Job Completed. Successfully created 1 virtual volume(s)'.

Red arrows connect the 'View Job Status' link in Step 1 to the job entry in Step 2, and then to the job entry in Step 3. A text box on the right says 'Click this link to monitor status' with an arrow pointing to the 'View Job Status' link.

© Copyright 2015 EMC Corporation. All rights reserved. Module: VPLEX Storage Provisioning 46

The VIAS provisioning job will start and use the AMP to instruct the back-end array to create the LUN(s) out of the selected storage pool, allocate and expose them to the VPLEX cluster(s) BE ports. We see in the example that the job progressed and concluded successfully.

Lesson Summary

During this lesson the following topic was covered:

- Preparing storage for host access using
 - Advanced provisioning
 - EZ provisioning
 - VIAS Integrated array service-based provisioning



This lesson covered configuring VPLEX storage into volumes and provisioning and integrating these with the hosts that will access them.

Lesson: Storage Volume Encapsulation

This lesson covers the following topics:


- Storage Volume Encapsulation overview
- Using the 'Storage-Tool Compose' VPLEX CLI command to perform encapsulation
- Generating the appropriate Solve Desktop procedure



This lesson covers claiming Storage Volumes using Encapsulation.

Storage Volume Encapsulation Overview

- Encapsulation claims Storage Volumes while retaining all existing data
- Use to import existing non-VPLEX Volumes into VPLEX
- Feature can be implemented using VPLEX CLI or Unisphere
- CLI allows for scripting



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

49

VPLEX provides the ability to claim back-end Storage Volumes already in use under its control. The process of claiming a Storage Volume while saving existing user data is called storage encapsulation. Any existing back-end user volume may be encapsulated, including non-bootable and host boot image volumes.

Using Storage Volume Encapsulation

- Use the **Storage-Tool Compose** VPLEX CLI command with the appropriate parameters
- Best practice is using the Solve Desktop procedure
- Solve Desktop Procedure varies depending on whether the back-end Storage Volume is a data or boot image LUN

Steps required for Encapsulation:

1. Write down the UID of the existing back-end array LUN
2. Un-mount the back-end Storage Volume from host then detach
3. Claim and Encapsulate the Storage Volume using the VPLEX CLI
4. Provision the new Virtual Volume back to the host and note the new UID under VPLEX
5. Remount the new Virtual Volume to the host

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

50

Storage Volume Encapsulation can be configured using the VPLEX CLI **Storage-Tool Compose** command. Also, it is important to remember that the best practice is to generate the appropriate Solve Desktop procedure to follow and to be fully familiar with it before beginning the operation.

Encapsulation Method using CLI

Use the **Storage-Tool Compose** CLI Command for Encapsulation

- Example of the **Storage-Tool Compose** command and parameters:

```
storage-tool compose 1. -n new_Encap 2. -g raid-0 3. -d VPD83T3:60000970000195901096533030313832  
4. -v /clusters/cluster-1/exports/storage-views/esx_21
```

Argument Definitions:

1. `'-n new_Encap'` Specifies the name of Virtual Volume
2. `'-g raid-0'` Specifies the desired geometry to use `'-d'` [storage-volumes] - Specifies a list of Storage Volumes to be used to build the Virtual Volume.
3. `'-v'` [Storage View name] Specifies the VPLEX storage-view(s) to receive the new virtual-volume. In this example, the Encapsulated volume is being exported to the host 'esx_21' Storage View

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

51

The **Storage Tool Compose** command is used with specific parameters to create a virtual-volume on top of the specified storage volumes, building all intermediate Extents, local, and Distributed Devices as necessary.

Refer to the *VPLEX CLI Reference Guide* for more information about this command and syntax.

Using Storage Volume Encapsulation (Cont'd.)

- The SolVe Desktop procedure will vary depending on the use of the Encapsulated volume

The screenshot displays the EMC² VPlex SolVe Generator interface. On the left, a sidebar shows a tree structure under 'Topic' with 'VPlex Customer Procedures' expanded, and 'Selections' showing 'Procedures: Provision' and 'Provisioning procedures: Encapsulate arrays on ESX with boot from SAN'. The main area shows the generated document content for 'Encapsulating an ESX SAN Boot Disk'. A red arrow points from the 'boot from SAN' selection in the sidebar to the document title. The document content includes assumptions, a note, and a procedure list.

EMC² VPlex SolVe Generator
Solution for Validating your engagement

Topic
VPlex Customer Procedures

Selections
Procedures: Provision
Provisioning procedures: Encapsulate arrays on ESX with boot from SAN

Generated: 4/17/2015 3:50:48 PM
SolVe Generator Updated:

IMPORTANT: To be sure you have the latest SolVe generator application and content, always choose Yes when asked if you want to update the application or repository.

Encapsulating an ESX SAN Boot Disk
This procedure describes the task of encapsulating a SAN boot disk through VPlex in a non-virtualized environment.

Assumptions

- The ESX host is running with a SAN boot disk connected directly (or through a switch) to the storage array.
- At least one virtual machine is running I/Os on the other disks exposed to the ESX host.
- VPlex must be commissioned.
- One new switch (pair of switches if high availability is required) is available for use as a front-end switch.

Note: This document can be used for hosts other than ESX with the difference that virtual machines are not applicable to the other hosts. Instead, we need to create file systems on the disks exposed to the other hosts.

Procedure

- [] Power down the virtual machines on the host by shutting down the operating systems on the virtual machines.
- [] Shut down the ESX host using the command **shutdown -h now**.
- [] Remove the direct connection between ESX server and storage array. If there is a switch between the two, remove zones and disconnect the ESX server from the Switch.
- [] Remove the LUNs from the storage group on the storage array.

Use the most up-to-date SolVe Desktop procedure generator, and be certain to download all VPlex product updates before generating a procedure. When using the SolVe Desktop be certain to select the appropriate options for the particular Storage Volume Encapsulation configuration. In this example, we specified that the Storage Volume we will encapsulate is the host OS boot LUN. Since this host will again boot from this LUN under VPlex, we must be certain that the procedure used will enable this host to boot from the new Virtual Volume when it is represented again after encapsulation.

Lesson Summary

During this lesson the following topics were covered:

- Storage Volume Encapsulation Overview
- Using the 'Storage-Tool Compose' CLI VPLEX CLI command to perform encapsulation
- Generating the appropriate Solve Desktop procedure

This lesson covered claiming Storage Volumes using Encapsulation.

Module: VPLEX Storage Provisioning

Key points covered in this module:

- Configuring host-to-storage connectivity through VPLEX
- VPLEX virtual storage constructs
- Provisioning Virtual Volumes to hosts using
 - Advanced Provisioning
 - EZ Provisioning
 - Integrated storage provisioning (VIAS)
- Storage Volume Encapsulation



Here is a summary of the topics covered in this module.

Labs: VPLEX Health Check Provisioning with VPLEX



The lab covers:

- VPLEX Health Check
- Register Host Initiators
- Provision Storage
- Storage View Creation
- Using Provision from Storage Volume Wizard
- Provisioning using VIAS
- Encapsulation

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Storage Provisioning

EMC²

55

The lab covers Provisioning with VPLEX.

This page intentionally left blank.



Module: Provisioning Distributed Devices

Upon completion of this module, you should be able to:

- Discuss the methods for configuring Distributed Devices
- Create a Distributed Device
- List the different properties of a Consistency Group
- Create Consistency Groups for Distributed Devices
- Discuss the different Detach Rules for a Consistency Group
- Describe the different failure scenarios and their impact



This module focuses on managing the storage network using VPLEX.

Lesson: VPLEX Distributed Device Concepts

This lesson covers the following topics:

- Distributed Devices overview
- Rule-sets to handle inter-cluster failures
- VPLEX Witness overview
- Consistency Groups overview




This lesson introduces the concepts of VPLEX Distributed Devices, rule-sets and VPLEX Witness to handle cluster and inter-cluster communications failures, VPLEX Witness and Consistency Groups.

Distributed Devices Overview

Distributed Devices: VPLEX Metro storage objects created on storage arrays attached to each cluster

- Distributed Devices support Virtual Volumes that can be presented to each host through its own Storage View
- I/O can be issued from either cluster concurrently
- Virtual Volumes placed on Distributed Devices appear as a single shared volume
- Distributed coherent cache preserves data integrity



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

3

Distributed Devices have their underlying storage arrays located at both clusters in a VPLEX Metro system.

Distributed devices support virtual volumes, that are presented to a host through a Storage View. Distributed devices are present at both clusters for simultaneous active/active read/write access. To the host, the Virtual Volumes appear as single volumes located on a single array. VPLEX Access-Anywhere ensures consistency of the data between the clusters.


VPLEX Distributed Devices enable distributed data centers. Some of the benefits of implementing a distributed data center include:

- Increased availability - Both data centers can serve as production workloads, while providing high availability for the other data center.
- Increased asset utilization - Passive data centers can have idle resources.
- Increased performance/locality of data access - Data does not have to be read from the production site as the same data is read/write accessible at both sites.

You can configure up to 8,000 Distributed Devices in a VPLEX system. That is, the total number of distributed Virtual Volumes plus the number of top-level local devices must not exceed 8,000.

Logging Volume

- Keeps track of write I/O during an inter-cluster link outage
- Keeps track of data changes during loss of access to the Distributed Device mirror volume
- Required at each VPLEX cluster before creating a Distributed Device
- Log information used to synchronize mirrors after link or down mirror volume access is restored
- Must be on high-performance storage to keep up with possibly high I/O rates



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

4

Logging volumes are required at each cluster before a Distributed Device can be created. Logging volumes are used to keep track of blocks written during an inter-cluster link failure, or when one leg of a distributed RAID-1 becomes unreachable and then recovers. After a link failure is restored or an unreachable leg recovers, VPLEX uses the information in logging volumes to synchronize the mirrors by sending only changed blocks across the link. After the inter-cluster link or leg is restored, the VPLEX system uses the information in the logging volumes to synchronize the mirrors by sending only changed blocks across the link. Logging volumes also track changes during loss of a volume when that volume is one mirror in a Distributed Device.


During and after link outages, logging volumes are subject to high levels of I/O. Thus, logging volumes must be able to service I/O quickly and efficiently. For more information about logging volume requirements and configuration, please see the *VPLEX Administration Guide*.

Note: Logging is only used with Distributed Devices and not used to optimize re-synchronization of local RAID-1 volumes in a cluster.

Rule-Sets Handle Inter-Cluster Failures

Virtual Volumes to Resume on One Cluster and Suspended On the Other

- Determine which Cluster continues I/O after a failure
 - Inter-cluster link fails
 - Cluster Fails
- Failure Starts delay timer (5 sec default)
 - All I/O is suspended until timer expires
- Rule-set determines “winning cluster”
 - I/O resumes at winning cluster
 - I/O is suspended at losing cluster



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

5

Most I/O workloads require specific sets of Virtual Volumes to resume on one cluster and remain suspended on the other cluster during failures. Rule-sets are predefined rules that determine which cluster continues servicing I/O when connectivity between clusters is lost. If the clusters lose contact with one another or if one cluster fails, rule sets define which

cluster continues operation (the preferred cluster) and which suspends I/O (the non-preferred cluster). This works for most link or cluster failures. In the case where the preferred cluster fails, all I/O is suspended resulting in data unavailability.

When a loss of connectivity first occurs, VPLEX:

Suspends I/O to all Distributed Devices on both clusters, and starts a delay timer.

If connectivity is not restored when the timer expires (default is 5 seconds), VPLEX:

Resumes I/O on the leg of the Distributed Device (the “winning” cluster) as determined by the device’s Rule-set

Keeps I/O suspended on the other leg (the “losing” cluster). When the inter-cluster link is restored, the I/O written to the winning cluster is re-synchronized to the losing cluster.

The detach rule can be changed by selecting the Distributed Device’s supporting device and then selecting a different cluster to detach from. Detach rules can also be modified within the VPLEX CLI by entering `set rule-set-name cluster-<#>-detaches` from within the Distributed Device.

Distributed Device Rule-set Options

Rule-set name	Cluster-1	Cluster-2
cluster-1-detaches	Services I/O	Suspends I/O
cluster-2-detaches	Suspends I/O	Services I/O

Auto-resume determines behavior after connectivity is restored

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

6

EMC recommends that only the two default Rule-sets be applied to Distributed Devices.

The default value of a Distributed Device's Rule-set is determined by the management server on which the device was created. If a device is created on the management server for cluster-1, the default Rule-set for that device is cluster-1-detaches. Rule-sets are located in the distributed-storage/rule-sets context. Rules are located under their Rule-set context.

When cluster connectivity is restored, subsequent behavior is determined by the **auto-resume** attribute for the Distributed Device.


- If auto-resume is set to **false**, the Distributed Device on the losing cluster remains suspended until I/O is manually resumed by the system administrator (using the device resume-link-up command).
- If auto-resume is set to **true**, I/O may start immediately after the link is restored. After I/O is resumed to the mirror leg on the losing cluster, any data written to the device on the winning cluster during the outage is resynchronized from the leg on the winning cluster.

Use the **set** command to configure a device's auto-resume attribute.

VPLEX Witness Benefit

Automatically Maintains Distributed Consistency Group Availability

- Automatically distinguishes between and provides guidance to clusters on how to react to a cluster or inter-cluster communication failure
- Allows VPLEX Metro to provide continuous availability in case of a failure
- The VPLEX Witness VM must be in a third failure domain relative to the VPLEX cluster sites to be an effective advisor to the distributed system
- VPLEX Witness works only in conjunction with distributed Consistency Groups



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

7

VPLEX Witness is an optional component installed as a virtual machine on a customer host. VPLEX Witness helps multi-cluster VPLEX configurations automate the response to cluster failures and inter-cluster link outages.

This slide shows the following features of the VPLEX Cluster Witness:


- The Cluster Witness should be in a “third failure domain”;
- The Cluster Witness Server is installed as a VM;
- The Cluster Witness Server provides guidance to both clusters depending on its observations.

VPLEX Witness observes the state of the clusters, and thus can distinguish between an outage of the inter-cluster link and a cluster failure. VPLEX Witness uses this information to guide the clusters to either resume or suspend I/O. VPLEX Witness works in conjunction with Consistency Groups. VPLEX Witness guidance does not apply to local volumes and distributed volumes that are not members of a Consistency Group.

Note that where VPLEX Witness is used, it will only manage the outcome of Distributed Devices that are in Consistency Groups.

VPLEX Consistency Groups

- Common properties applied to all Virtual Volumes in the group.
- Allows for VPLEX Witness
- Used with RecoverPoint



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices



8


VPLEX Consistency Groups aggregate volumes to enable the application of a common set of properties to the entire group. Consistency groups aggregate up to 1000 Virtual Volumes into a single entity that can be managed as easily as individual volumes. Consistency group detach rules define on which cluster I/O continues during cluster or inter-cluster link failures. The groups work together with VPLEX Witness.

The properties of a Consistency Group are applied to all the Virtual Volumes in the Consistency Group. All Consistency Groups have configurable properties that determine I/O behavior, including:

- cache mode
- visibility
- storage-at-clusters
- local-read-override
- detach-rule
- auto-resume-at-loser
- virtual-volumes
- recoverpoint-enabled

VPLEX supports up to 1024 synchronous Consistency Groups.

Detach Rules for Distributed Consistency Groups



- Detach Rules designates which cluster detaches during a failure
- Three possible values:
 - No Automatic Winner: No winning cluster selected
 - Winner (For Synchronous Consistency Groups) : After the delay the specified cluster is declared winner
- Detach rules apply to all volumes of a Consistency Group and override rules applied to individual volumes

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

9

The Consistency Group Detach Rule designates which cluster detaches if clusters lose connectivity. Possible values:


No-Automatic-Winner — The Consistency Group does not select a winning cluster.
Winner (cluster-name) (delay) — Applicable only to synchronous Consistency Groups. The cluster specified by cluster-name is declared the winner after the inter-cluster link outage lasts more than the number of seconds specified by delay. For a Consistency Group that is participating in RecoverPoint, this value must be set to the cluster at which the RP splitter is running.

If a Consistency Group has a detach-rule configured, the rule applies to all volumes in the Consistency Group, and overrides any rule-sets applied to individual volumes.
(VPLEX Metro and VPLEX Geo configurations only)

Failure Scenario Without VPLEX Witness

Failure Scenario Response: Use Rule-set

- Rule-set Rules:
 - Cluster-1 Detaches
 - Cluster-2 Detaches
 - Delay time to enforce Rule



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

10

For a VPLEX Metro without Cluster Witness, there has to be a method to avoid split-brain scenarios. Each Distributed Device has a Rule-set applied to it.


As discussed, Rule-sets are predefined rules that determine which cluster continues I/O when connectivity between clusters is lost. Rule-sets apply to devices that are not members of Consistency Groups.

As a reminder, a cluster loses connectivity to its peer cluster when:

- The inter-cluster link fails
- The peer cluster fails

When a loss of connectivity occurs, VPLEX starts a delay timer (default 5 sec) and suspends I/O to all Distributed Devices on both clusters. If connectivity is not restored before the timer expires, then the Rule is enforced.

Continuous Availability With VPLEX Witness



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

11

Presented here are the various inter-site failure scenarios and how they each are handled when using VPLEX Witness.

Lesson Summary

During this lesson the following topics were covered:

- Distributed Devices overview
- Rule-sets to handle inter-cluster failures
- VPLEX Witness overview
- Consistency Groups overview



This lesson covered the concepts of VPLEX Distributed Devices, rule-sets and VPLEX Witness to handle cluster and inter-cluster communications failures, VPLEX Witness and Consistency Groups.

Lesson: VPLEX Distributed Devices

This lesson covers the following topics:

- Distributed Device Configuration
- Creating Distributed Devices
- Creating Consistency Groups



This lesson covers the configuration of Distributed Devices along with Consistency Groups.

Configure Distributed Devices

These are the general steps to configure Distributed Devices:

1. Create logging volume
2. Identify claimed storage
3. Create Extents from selected storage
4. Create local devices
5. Create Distributed Device
6. Attach Rule-set/Detach Rules
7. Create Virtual Volumes on the new Distributed Device and export to hosts



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

14

There are many different methods to configure Distributed Devices, including using the VIAS tool. The above steps describe the CLI-based process. Before reviewing these manual steps, we will look at how to configure Distributed Devices in the VPLEX GUI.

To review, Distributed Devices support Virtual Volumes, that are presented to a host through a storage view. From the host, the Virtual Volumes appear as single volumes located on a single array. All Distributed Devices must be associated with a logging volume. During a link outage, the logging volume is used to map the differences between the legs of a DR1. When the link is restored, the legs are resynchronized using the contents of their logging volumes.

All Distributed Devices must have a detach rule-set to determine which cluster continues I/O when connectivity between clusters is lost. For Consistency Groups on Distributed Devices, these are Detach Rules.

Remember that you can configure up to 8,000 Distributed Devices in a VPLEX system. That is, the total number of distributed Virtual Volumes plus the number of top-level local devices must not exceed 8,000.

Create Logging Volume

- Logging volumes created at each cluster during VPLEX Metro system setup
- In Unisphere, the logging volume is automatically selected for a Distributed Device
- To specify a logging volume to be used for a Distributed Device use the VPLEX CLI

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

15

Logging volumes are created during initial VPLEX Metro system setup and are required in each cluster to keep track of any blocks written during a loss of connectivity between clusters. A logging volume must exist in the system to create a Distributed Device. If no logging volume exists in the system, an error message appears to inform you that you cannot create a Distributed Device without a logging volume. In the GUI, the logging volume is automatically selected for the Distributed Device. To specify a logging volume to be used for a Distributed Device, use the VPLEX CLI.

Please refer to the *VPLEX Administration Guide* for more information about logging volume creation and configuration.

Creating Distributed Devices via Unisphere

Using Unisphere to Provision and Manage Distributed Devices

The screenshot displays the EMC Unisphere for VPLEX vs.4 web interface. The top navigation bar includes 'System Status', 'Monitoring', 'Provision Storage' (highlighted with a red box), 'Mobility Central', and 'Support'. Below this, the 'Provision Storage' section is visible, with sub-options for 'Provision from Pools' and 'Provision from Storage Volumes'. A red arrow points from the 'Provision Storage' menu to the 'Distributed' sub-menu. Another red arrow points from the 'Distributed' sub-menu to the 'Distributed Consistency Groups' and 'Distributed Virtual Volumes' options. A third red arrow points from the 'Distributed Consistency Groups' option to the 'Distributed Devices' option. The 'Distributed Devices' option is highlighted with a red box. The interface also shows a 'cluster-2' section with a description of provisioned storage volumes and consistency groups.

EMC Unisphere for VPLEX vs.4

System Status Monitoring Provision Storage Mobility Central Support

Provision from Pools
Provision volumes from storage pools on supported arrays that are visible to VPLEX through Array Management Providers

Provision from Storage Volumes
Provision volumes from storage volumes

Distributed
Provision devices, volumes and consistency groups that will be available to both clusters.

cluster-2
Provision storage volumes, extents, devices, virtual volumes, consistency groups, and storage views that will be available to the cluster. Also provision the initiators that will use the storage and the front-end and HBA ports for exporting VPLEX storage.

EMC Unisphere for VPLEX vs.4

System Status Monitoring Provision Storage Mobility Central Support

Provision Storage > Distributed

Distributed Consistency Groups
Consistency groups allow you to group volumes together and apply a set of properties to the entire group.

Distributed Virtual Volumes
A virtual volume is created on a device or a distributed device, and is presented to a host through a storage view.

Distributed Devices
Distributed devices are configured by using storage from both clusters. A distributed device is created from storage in both clusters.

EMC

© Copyright 2015 EMC Corporation. All rights reserved. Module: Provisioning Distributed Devices 16

Distributed Devices may be created and configured using Unisphere for VPLEX. Shown here is using the Provision Storage menu and the Distributed sub-menus to create the desired Distributed Devices of the various types.

Creating a Distributed Device – Existing Devices

© Copyright 2015 EMC Corporation. All rights reserved. Module: Provisioning Distributed Devices 17

Distributed devices are mirrored between clusters in a VPLEX Metro. From this wizard, devices of the correct or greater capacity must exist at both sites. The distributed RAID-1 device is then created these devices one at each cluster. The local protection is not a consideration. These devices can only be as large as the smaller of the two. Distributed devices are created using the Unisphere for VPLEX Distributed Devices Wizard.

1. Navigate to **Provision Storage>Distributed Storage>Distributed Devices>** and choose the Distributed Device type. Select **Create**. This launches the Wizard. Step 1 allows the user to select a local device to mirror. Select the device and then **Next**.

Creating a Distributed Device (Cont'd.)

The screenshot shows the 'Create Distributed Devices' wizard, Step 2: Select Mirrors. The wizard is titled 'Create Distributed Devices' and has a progress bar on the left with steps: 1. Select Devices to Mirror, 2. Select Mirrors, 3. Set Data Synchronization, 4. Select Consistency Group, 5. Review and Finish, 6. View Results. Step 2 is currently active.

On the left, there is a yellow box with instructions: 'Click Auto-Generate Mirror Mappings to automatically select a mirror for the selected item. Or, manually select the source and target items and then click Add Mirror. Created mappings appear in the Selected Mirrors list below. To remove a mapping from this list, select the mapping and then click Remove Mirror. To save the mirror mappings to a file on your system, click Export Mirror Mappings.'

The main area shows 'Target Cluster: cluster-2'. Below this, it says 'To select appropriate mirrors automatically based on best-fit capacities, click Auto-generate Mirror Mappings. Otherwise, select each source and target to add mirror mappings manually.' There is a button 'Auto-Generate Mirror Mappings'.

Below this, there are two tables. The first table is 'Selected Devices to Mirror' with columns 'Name', 'Capacity', and 'Virt'. It contains one row: 'device_group1-Symm1096_03D7_1...' with a capacity of 4.94 GB. The second table is 'Available Devices to use for Mirroring' with columns 'Name', 'Capacity', and 'Virt'. It contains two rows: 'device_CLARIION3040_LUN_00063_1' and 'device_CLARIION3040_LUN_00072_1', both with a capacity of 4.94 GB.

Below the tables, there is a 'Selected Mirrors' table with columns 'Device' and 'Virtual Volume'. It is currently empty. There is an 'Add Mirror' button and a 'Next >' button.

On the right, there is a dialog box titled '3 Set Data Synchronization'. It has a section 'Specify Device Synchronization' with two radio buttons: 'Synchronize Data' (selected) and 'Do Not Synchronize Data - Advanced Users Only'. The 'Synchronize Data' option has a description: 'This option will synchronize data from the source device to the target device. Select this option to mirror existing devices on another cluster.' The 'Do Not Synchronize Data' option has a description: 'This option will not synchronize data between the devices that do not yet have a virtual volume. Select this option only if you are sure that host-based volume initialization will resolve any inconsistencies between source and target devices. If you are unsure, then select the Synchronize Data option. Selecting this option will not apply if the source device already has a virtual volume. Devices with virtual volumes are automatically synchronized to prevent data loss.'

At the bottom right of the dialog box, there is a 'Next >' button and an 'EMC' logo.

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

18

Step 2 allows for the selection of the volume at the remote cluster. In this example, we find two candidate volumes on cluster-2 that are of similar size to the source volume selected on cluster-1. One or multiple candidate volumes may be selected. Here, we selected a single target volume on cluster-2. Click **Add Mirror**.

In step 3, we are given the opportunity to copy all of the source data to the target volume thus synchronizing the devices between the two clusters. A reason to choose not to perform an initial synchronization is to save time, where the source volume never contained a Virtual Device that was attached to a host. Click **Next**.

Creating a Distributed Device (Cont'd.)

The screenshot shows the 'Create Distributed Devices' wizard window. The left sidebar contains a progress list with steps 1 through 6. Step 4, 'Select Consistency Group or Device Rule Set', is the current step. The main area has three radio buttons: 'Select an existing Consistency Group', 'Create a Consistency Group', and 'Select a Device Rule Set'. The 'Select a Device Rule Set' option is selected. Below it, a 'Rule Set' dropdown menu is open, showing options 'cluster-1-detaches' and 'cluster-2-detaches'. A 'Next >' button is at the bottom right.

Create Distributed Devices

- ✓ 1 Select Devices to Mirror
- ✓ 2 Select Mirrors
- ✓ 3 Set Data Synchronization
- * 4 **Select Consistency Group or Device Rule Set**
- * 5 Review and Finish
- 6 View Results

Virtual volumes will be created on distributed devices that do not currently have a virtual volume on the source device.

Select an existing consistency group, or create a new one, to group an application's volumes together for ease of management.

Select a rule set if you do not want to use a consistency group.

4 Select Consistency Group or Device Rule Set

Select a Consistency Group or a Device Rule Set

☐ Select an existing Consistency Group

Consistency Group: * Select a Consistency Group...

☐ Create a Consistency Group

Consistency Group Name: *

Rule Set: * Select a Rule Set...

You must select a detach rule set to determine which cluster detaches and continues I/O during an inter-cluster link failure.

☒ Select a Device Rule Set

Rule Set: * Select a Rule Set...

You must select a detach rule set to determine which cluster detaches and continues I/O during an inter-cluster link failure.

cluster-1-detaches

cluster-2-detaches

Next >

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

19

The next step allows the user to Create a Consistency Group. Select an existing Consistency Group, or apply a Device Rule-set. In this example, we'll select a Rule-set.

Creating a Distributed Virtual Volume via CLI

```
VPlexcli:/> storage-tool compose
  --name <name>
  -g <raid type>
  -d <Source Storage Volumes VPDs>
  -m <Mirror Storage Volumes VPDs>
  -v <Storage View Name>
  -c <Consistency Group Name>

storage-tool compose --name new-vv -g raid-1
--storage-volumes VPD83T3:60000970000195901096533030344530
-m VPD83T3:600601603e40230007fa1853519fe311 -v my-view -c my-CG
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

20

The **storage-volume compose** command can be used as a simple way to create a Virtual Device. Using this one command, the Distributed Device and Virtual Volume can be created from provisioned Storage Volumes, claimed or un-claimed, from both clusters.

The command allows the user to add the Virtual Volumes created to both a Consistency Group and a storage view. They must exist already. The example displays adding a Virtual Volume named new-vv to a storage view named my-view and a Consistency Group named my-CG.

Creating a Distributed Device via the CLI

- Use the **ll -p **/storage-volumes** command to display the Storage Volumes on each cluster:

```
Vplexcli:/> ll -p ** storage-volumes
```

```
/clusters/cluster-1/storage-elements/storage-volumes:
```

Name	VPD83 ID*	Capacity	Use	Vendor	IO Status	Type	Thin Rebuild
Student1_SV1	VPD83T3:60000970000195901	40G	used	EMC	alive	normal	true
Symm1096_018B	VPD83T3:60000970000195901	40G	used	EMC	alive	normal	true
Symm1096_04D6	VPD83T3:60000970000195901	10G	used	EMC	alive	normal	false
Symm1096_04DA	VPD83T3:60000970000195901	20G	used	EMC	alive	normal	false
Symm1096_04DD	VPD83T3:60000970000195901	20G	used	EMC	alive	normal	false
Symm1096_04DF	VPD83T3:60000970000195901	20G	claimed	EMC	alive	normal	false

```
-- more --(TOP )- [h]elp  
/clusters/cluster-1/storage-elements/storage-volumes:
```

* VPD83 ID field truncated to accommodate font size

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

21

To identify Storage Volumes on each cluster that are available to be combined into the Distributed Device, use the command and the output as shown.

Note: To prevent creating Distributed Devices with unusable “leftover” storage, total capacities of the selected storage at both clusters should be identical.

Create Extents

- From the storage-volumes CLI context, use the **extent create** command to create the extents for the Distributed Device:

```
extent create --storage-volumes <storage-volumes> --size <size> --  
block-offset <offset> --num-extents <integer>
```

Parameter definitions:

--storage-volumes - List of claimed Storage Volumes to be added to the extent.

--size - Size of each extent. If no size is specified, the largest contiguous range of blocks is used to create the specified number of extents.

--block-offset - The block offset on the underlying Storage Volume. If block offset is not specified, it is set automatically.

--num-extents - Number of extents to create per specified storage-volume. If no number is specified, it will automatically create one extent per Storage Volume.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

22

Use the **extent create** command from the VPLEX CLI storage-volumes CLI context to create extents using the Storage Volumes previously identified at each cluster. After completion on the first cluster, navigate to the storage-volumes context on the second cluster and repeat to create one or more extents to be added to the Distributed Device at that cluster.

Identify Extents For Local Devices

- Use the `ll -p **/extents/` command to display the extents created on each cluster:

```
VPlexcli:/> ll -p **/extents/
/clusters/cluster-1/storage-elements/extents:
Name                               StorageVolume Capacity Use
-----
extent_VNX_logging_1              vnx_logging      80G  claimed
extent_VNX_lun0_1                 vnx_lun0          1G   used
extent_Symm1852_21C_1             Symm1852_21C      16G   used
extent_Symm1852_AB4_1             Symm1852_AB4      16G   claimed
extent_Symm1852_AB8_1             Symm1852_AB8      16G   claimed

/clusters/cluster-2/storage-elements/extents:
Name                               StorageVolume      Capacity Use
-----
extent_VNX_logging_2_1            VNX_logging_2      80G  claimed
extent_VNX_Test_Lun_5_1           VNX_Test_Lun_5     10G  claimed
extent_VNX_Test_Lun_6_1           VNX_Test_Lun_6     10G  claimed
extent_VNX_Test_Lun_9_1           VNX_Test_Lun_9     10G  claimed
extent_Cluster_2_ESX_DS_1         Cluster_2_ESX_Datastore 200G claimed
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

23

To identify extents on each cluster that are available to be combined into the local Device, use the command and the output as shown.

Create Local Devices from Extents

- Use the **local-device create** command to create a local device with the specified name on each cluster.

```
local-device create --name name --geometry geometry -extents extents -  
-stripe-depth A multiple of the block size: 4 KB
```

Parameter definitions:

--name - Name for the new device. The name must be unique across all clusters. Devices on different clusters that have the same name cannot be combined into a Distributed Device.

--geometry - Geometry for the new device. Valid values are **raid-0**, **raid-1** or **raid-c**.

--extents - List of pathnames of claimed extents to be added to the device, separated by commas. Extents can also be other local devices (to create a device of devices).

--stripe-depth - Required for devices with a geometry of RAID-0. Specifies the stripe depth in 4KB blocks.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

24

Use the **local-device create** command to create a local device with the specified name using the extents previously identified at each cluster.

Note: The stripe resulting from the '--stripe-depth' parameter is sized using the following formula:

$\text{<stripe-depth> * <the block size on the source storage extents>}$ A depth of 32 means 128 KB (32 x 4 KB) is written to the first disk then the next 128 KB is written to the next disk.

Identify Local Devices for Distributed Devices

- Use the `ll --p ** /devices` command to display the available local devices (with no Virtual Volume configured):

```
VPlexcli:/> ll --p **/devices
/clusters/cluster-1/devices:
Name                Operational Health Block  Block Capacity Geometry Visibility Transfer Virtual
                  Status    State  Count  Size    Capacity Geometry Visibility Size    Volume
-----
TestDevCluster1    ok        ok     4195200 4K      16G     raid-1  local    2M      -
base0              ok        ok     262144 4K      1G      raid-0  local    -       base0_vol
base1              ok        ok     262144 4K      1G      raid-0  local    -       base1_vol
base2              ok        ok     262144 4K      1G      raid-0  local    -       base2_vol
base3              ok        ok     262144 4K      1G      raid-0  local    -       base3_vol

/clusters/cluster-2/devices:
Name                Operational Health Block  Block Capacity Geometry Visibility Transfer Virtual
                  Status    State  Count  Size    Capacity Geometry Visibility Size    Volume
-----
TestDevCluster2    ok        ok     4195200 4K      16G     raid-1  local    2M      -
base01             ok        ok     524288 4K      2G      raid-0  local    -       base01_vol
base02             ok        ok     524288 4K      2G      raid-0  local    -       base02_vol
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

25

To identify Local Devices on each cluster that are available to be combined into the Distributed Device, use the command and the output as shown.

As discussed, Local Devices on each cluster constitute the two “legs” or access points of the Distributed Device. The individual local devices may include any underlying type of Storage Volume or geometry (RAID 0, RAID 1, or RAID C), but they should be of the same capacity. If a Distributed Device is configured with local devices of different capacities:

- The resulting Distributed Device will be only as large as the smallest local device.
- The leftover capacity on the larger device will not be available.

To create a Distributed Device without wasting capacity, choose local devices on each cluster with the same capacity.

Important: If there is pre-existing data on a Storage Volume and the Storage Volume is not claimed as being application-specific, converting an existing local RAID device to a distributed RAID will not initiate a rebuild to copy the data to the other member or “leg”. As a result, the data will exist at only one cluster.

Create the Distributed Device from the Local Device

- Use the **ds dd create** command to create a Distributed Device.

```
ds dd create --name name --devices devices --logging-volume logging-volumes --  
rule-set rule-set
```

Parameter definitions:

--name <name> - The name must be unique across the entire VPLEX configuration.

-- devices - List of pathnames to local devices to add to the Distributed Device. Select devices that have the same capacities. Separate entries in the list by commas.

--logging-volume - List of pathnames to one or more logging volumes to use with the new device.

--rule-set - Attaches the specified rule-set to the device. If no rule-set is specified, the cluster that is local to the management server is assumed to be the winner in the event of an inter-cluster link outage.

The EMC logo is a blue square with the letters "EMC" in white.

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

26

Use the **ds dd create** command to create a Distributed Device using the Local Devices previously identified at each cluster.

Note: For the '--logging-volume' parameter, if no logging volume is specified, a logging volume is automatically selected from any available logging volume that has sufficient space for the required entries. If no available logging volume exists, an error message is returned.

Create Distributed Device

```
VPlexcli:/clusters/cluster-1/devices> ds dd create --name cli_dd
--devices device_Dev_3_c1,device_Test_DD_c2 -r cluster-1-detaches
--source-leg device_Test_DD_c2
VPlexcli:/distributed-storage/distributed-devices/cli_dd> ll
```

Name	Value
-----	-----
auto-resume	false
block-count	5242880
block-size	4K
capacity	20G
clusters-involved	[cluster-1, cluster-2]
geometry	raid-1
health-indications	[rebuilding]
health-state	minor-failure
locality	distributed
operational-status	degraded
rebuild-allowed	true
rule-set-name	cluster-1-detaches
service-status	running



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

27

Shown is an example of using the **ds dd create** command to create a Distributed Device. Also shown is the **ll** command used to list information about Distributed Devices contained in the VPLEX file system area `/distributed-storage/distributed-devices/cli_dd`.

Create Virtual Volumes from Distributed Devices

- Use the **virtual-volume create** command to create a Virtual Volume on a specified Distributed Device:

```
virtual-volume create --device device --set-tier {1|2}
```

Parameter definitions:

--device - Pathname of the device on which the Virtual Volume is configured.

--set-tier - Storage tier for the new Virtual Volume. This value is displayed to the host as part of the Virtual Volume's product ID.

- Use storage tiers to logically group storage. For example, VMAX systems as tier 1 storage and VNX systems as tier 2.
- Use the same tier groupings across third-party arrays as these vary in performance capabilities between vendors and models as well

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

28

Use the **virtual-volume create** command to create a Virtual Device with the specified name using the Distributed Devices previously identified at each cluster.

Allocate the new Virtual Volumes to the appropriate local and remote host Storage Views using the VPLEX CLI or Unisphere GUI.

Virtual Volume Create

```
VPlexcli:/ VPlexcli:/distributed-storage/distributed-devices>
virtual-volume create --device cli_dd

VPlexcli:/clusters/cluster-1/virtual-volumes> show-use-hierarchy -t
cli_dd_vol
virtual-volume: cli_dd_vol (20G, distributed @ cluster-1, unexported)
distributed-device: cli_dd (20G, raid-1)
  distributed-device-component: device_Dev_3_c1 (20G, raid-1, cluster-1)
    extent: extent_60000970000195901096533030344442 (20G)
      storage-volume: VPD83T3:60000970000195901096533030344442 (20G)
        logical-unit: VPD83T3:60000970000195901096533030344442
          storage-array: EMC-SYMMETRIX-195901096

distributed-device-component: device_Test_DD_c2 (20G, raid-1, cluster-2)
  extent: extent_600601603e40230007fa1853519fe311 (20G)
    storage-volume: VPD83T3:600601603e40230007fa1853519fe311 (20G)
      logical-unit: VPD83T3:600601603e40230007fa1853519fe311
        storage-array: EMC-CLARiION-APM00090503040
```

EMC


© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

29

Shown is an example of using the **virtual-volume create** command to create a Virtual Volume. Also shown is the **show-use-hierarchy -t** command used to display the complete usage hierarchy for the specified volume.

The Benefits of Using Consistency Groups for Distributed Devices



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

30

Recall that Consistency groups aggregate volumes together to ensure the common application of a set of properties to the entire group. You should create Consistency Groups for sets of volumes that require the same I/O behavior in the event of a link failure. In the event of a director, cluster, or inter-cluster link failure, Consistency Groups prevents possible data corruption. The optional VPLEX Witness failure recovery semantics apply only to volumes in Consistency Groups. In addition, you can move a Consistency Group from one cluster to another if required.


There are three types of Consistency Groups: **Asynchronous**, **Synchronous** and **RecoverPoint**.

Asynchronous Consistency Groups are used only for VPLEX Geo configurations. Synchronous Consistency Groups provide a convenient way to apply detach rules and properties to a group of volumes at a time; this simplifies system configuration and administration on large systems. Volumes in a synchronous group behave as volumes in a VPLEX Local or Metro, and can have global or local visibility. Synchronous groups can contain local, global, or distributed volumes.

RecoverPoint Consistency Groups contain the Journal and Replica volumes for the corresponding data volumes found in a RecoverPoint-enabled Consistency Group.

VPLEX Witness controls rule sets for Consistency Groups only. Distributed volumes not placed within Consistency Groups will simply apply the detach rules without guidance from VPLEX Witness.

Consistency Groups Visibility



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

EMC²

31

Visibility controls which clusters know about a Consistency Group. By default, a Consistency Group's visibility property is set only to the cluster where the Consistency Group was created. The visibility of the volumes within the Consistency Group must match the visibility of the Consistency Group. Local Consistency Groups with global visibility will always be synchronous.

Creating a Consistency Group

The image displays two screenshots of the 'Create Consistency Group' wizard interface.

Top Screenshot (Step 1: Select Name):

- Wizard steps: 1 Select Name, 2 Select Distributed Setti..., 3 Select Virtual Volumes, 4 Review and Finish, 5 View Results.
- Consistency Group Name: My-CG
- Text: Distributed consistency groups will be synchronous and visible to all clusters. You can add only distributed virtual volumes.

Bottom Screenshot (Step 2: Select Distributed Settings):

- Wizard steps: 1 Select Name, 2 Select Distributed Setti..., 3 Select Virtual Volumes, 4 Review and Finish, 5 View Results.
- Rule Set: Winner:cluster-1 (5 seconds)
- Dropdown menu options: Select a Detach Rule, No Automatic Winner, Winner:cluster-1 (5 seconds), Winner:cluster-2 (5 seconds).
- Text: Select a detach rule for the consistency group, and then click Next to proceed.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

32

A Consistency Group can be created many ways, from CLI or Unisphere. A new group can also be created when creating a Distributed Device or using the VIAS wizard. Displayed in this slide is the Consistency Groups creation wizard from Unisphere for VPLEX.

After providing a group name, a detach rule is chosen. The displayed delay of 5 seconds is the default value. This can only be changed using the CLI. Once a Rule-set is selected a brief explanation is displayed.

Creating a Consistency Group (Contd.)

Create Consistency Group

- ✓ * 1 Select Name
- ✓ 2 Select Distributed Setti...
- 3 Select Virtual Volumes**
- 4 Review and Finish
- 5 View Results

To select virtual volumes, select the volumes in the Available Volumes list, and then click Add to move them to the Selected list. You can also drag-and-drop between lists.

To remove virtual volumes from the list, select the volumes in the Selected list, and then click Remove to move them to the Available list.

3 Select Virtual Volumes

Available Virtual Volumes

Find

Volume Name	Lc
cli_dd_vol	d

Selected Virtual Volumes

Find

Volume Name	Lc
compose_vv	d

Buttons: Add >, Add All >, < Remove, < Remove All

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

33

The next step of the Wizard allows for the selection of Virtual Volumes to be placed in the Consistency Group. The volumes selected will have the properties of the group. An example of these properties are the Detach Rule, Visibility, or RecoverPoint enabled.

These settings can be changed from their default values using the CLI.

Consistency Group Properties

```
VPlexcli:/clusters/cluster-1/consistency-groups> 11 My-CG/
Attributes:
Name                                     Value
-----
active-clusters                         []
cache-mode                             synchronous
detach-rule                             winner cluster-1 after 5s
operational-status                      [(cluster-1,{ summary:: ok, details:: [] })),
(cluster-2,{
                                         summary:: suspended, details:: [] })]
passive-clusters                        []
recoverpoint-enabled                    false
storage-at-clusters                     [cluster-1, cluster-2]
virtual-volumes                         [compose_vv]
visibility                              [cluster-1, cluster-2]
```



Listed are the properties for the Consistency Group "My-CG". Not all the values can be displayed on this slide.

Lesson Summary

During this lesson the following topics were covered:

- Distributed Device configuration
- Creating Distributed Devices
- Benefits of Creating Consistency Groups



This lesson covered Distributed Device configuration using Unisphere and the CLI. It also covered benefits, properties, and creation of Consistency Groups.

Lesson: VPLEX Failure Scenarios

This lesson covers the following topic:

- How VPLEX handles Distributed Device Failures



This lesson covers Distributed Device failure scenarios and how to handle them.

Distributed Device Failure Scenarios

- Inter-cluster WAN failure
- Storage Volume degradation
- Mirror isolation



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

37

Listed here are the three most common failure scenarios that affect normal VPLEX operations. We'll discuss these and how to handle them if they occur.

Inter-cluster WAN Failure

- WAN failure will cause all Distributed Devices to partition
- WAN failure is the most typical problem to affect Distributed Devices



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

38

Local Devices on each VPLEX cluster constitute the two legs of each Distributed Device. An Inter-cluster WAN communication failure will break the ability for Distributed Devices to sustain I/O synchronization on local devices across clusters.

WAN Connectivity Outage

```
VPlexcli:/> connectivity validate-wan-com
```

```
connectivity: NONE
```

```
fc-port-group-3 - FAIL - No connectivity was found from any com port.
```

```
fc-port-group-0 - FAIL - No connectivity was found from any com port.
```

```
fc-port-group-2 - FAIL - No connectivity was found from any com port.
```

```
fc-port-group-1 - FAIL - No connectivity was found from any com port.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

39

From the VPLEX CLI, the displayed command shows that there is no WAN Connectivity from Cluster-1 to Cluster-2. All WAN connections are down.

Distributed Device Properties

```
/distributed-storage/distributed-devices/cli_dd:
```

Name	Value
-----	-----
auto-resume	false
block-count	5242880
block-size	4K
capacity	20G
clusters-involved	[cluster-1, cluster-2]
geometry	raid-1
health-state	ok
locality	distributed
operational-status	ok
rule-set-name	cluster-1-detaches
service-status	running
system-id	cli_dd
virtual-volume	cli_dd_vol



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

40

The displayed properties for a Distributed Device show that the auto-resume property is set to false. This is not the default setting, but the value may have been changed. This will prevent the Virtual Volume on the losing cluster from being accessible after a WAN outage has been corrected. User intervention is required. During this outage the winning cluster Virtual Volume is online and changes are being tracked in the Logging Volume.

Check Status

```
Vplexcli:/distributed-storage/distributed-devices> ll
```

Name	Status	Operational Status	Health State	Auto Resume	Rule-set Name	Transfer
cli_dd	winner-running	stressed	major-failure	false	cluster-1-detaches	128K

The top screenshot shows the 'cluster-1: Storage Views' interface. It features a table with the following data:

Name	Operational Status	Active	# Volumes	# Ports
my-sv	OK	Yes	1	2

The bottom screenshot shows the 'cluster-2: Storage Views' interface. It features a table with the following data:

Name	Operational Status	Active	# Volumes	# Ports
CL-2	Error	No	1	2



Here we see the Storage Views for Cluster-1 and 2. Cluster 1 is the winning cluster in this example. Notice the Operational Status of the 2 Storage Views.

Restoring Connectivity

```
Vplexcli:/> connectivity validate-wan-com
connectivity: FULL

fc-port-group-0 - OK - All expected connectivity is present.
fc-port-group-1 - OK - All expected connectivity is present

Vplexcli:/distributed-storage/distributed-devices> ll

cli_dd    need-resume ok ok      false  cluster-1-detaches
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

42

Now we see that the WAN Connectivity has been restored (after the delay timer expired). From the VPLEX CLI we notice that the Distributed Device needs a "resume" at the losing cluster.

Resume Distributed Device

```
Vplexcli:/clusters/cluster-2/virtual-volumes> device resume-link-up
--virtual-volumes cli_dd_vol -force


Vplexcli:/clusters/cluster-2/virtual-volumes> ll
Name          Status      service    Block    Size  Locality  Device
cli_dd_vol    ok          ok         running  5242880  20G  distributed  cli_dd
```



Here we see the command to resume the Virtual Volume. This can also be done using the device at the losing cluster. This command only works once WAN Connectivity is restored.

Distributed Device Mirror Isolation

- VPLEX periodically (once a minute) monitors if there are any Storage Volumes that have been marked degraded.
- If VPLEX detects degraded Storage Volumes that support a RAID-1 mirror, VPLEX isolates the mirror without causing data unavailability.
- Automatic or manual un-isolation options (automatic is the default setting)



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Provisioning Distributed Devices

44

RAID-1 devices provide redundancy by replicating data across a mirror. In write-through mode, this is achieved by persisting all writes to all mirror legs of a RAID-1 (local or distributed) device to the backend storage-volume, before acknowledging to the host. When a mirror is isolated, new writes are not persisted on the isolated mirror leg. Therefore, a RAID-1 device loses redundancy. If you prefer redundancy over performance enhancement, you might choose to disable the mirror isolation feature.

VPLEX provides the ability to automatically or manually un-isolate a mirror leg. If you choose not to enable the auto-unisolation feature, you must perform a manual un-isolation of the mirror leg.

Lesson Summary

During this lesson the following topics were covered:

- Distributed Device Failure Scenario and Handling



This lesson reviewed some failure scenarios you may encounter when working with Distributed Devices, and steps you can take to resolve the issues observed.

Module: Summary

Key points covered in this module:

- Discuss the methods for configuring Distributed Devices
- Create a Distributed Device
- List the different properties of a Consistency Group
- Create Consistency Groups for Distributed Devices
- Discuss the different Detach Rules for a Consistency Group
- Describe the different failure scenarios and their impact



This module covered configuring remote access, Distributed Devices, and Consistency Groups. We reviewed methods available to configure the appropriate options for your environment, as well as how to handle failure scenarios.

Lab: Configure Distributed Devices



This lab covers the following:

- Configure Distributed Devices
- Configure Consistency Groups
- Present Distributed Devices to Hosts

This lab covers how to configure Distributed Devices.

This slide intentionally left blank.

Module: Volume Management and Protection

Upon completion of this module, you should be able to:

- Explain the possible methods for Virtual Volume expansion and the reasons for choosing them
- Discuss the methods for configuring Virtual Volume protection
- List the VPLEX data migration options and their benefits
- Perform a Migration using Unisphere for VPLEX
- Create a Batch Migration plan and perform it
- Describe the methods of Virtual Volume protection using RecoverPoint

This module focuses on using VPLEX for Data Mobility.

Lesson: Volume Expansion and Protection

This lesson covers the following topics:


- Virtual Volume expansion
- Methods for configuring Virtual Volumes with Local and Remote protection

This lesson covers how to expand a volume using the two available expansion methods, and then how to protect volumes in a VPLEX.

Volume Expansion

2 Methods:

- Storage-Volume
 - VPLEX expands the underlying storage-volume (the corresponding LUNs on the back-end array).
- Concatenation (RAID-C)
 - The virtual volume is expanded by adding only specified extents or devices, as required.



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

3

A VPLEX virtual volume is created on a device or a distributed device, and is presented to a host through a storage view. For a number of reasons, you may want to expand the capacity of a virtual volume. If the volume supports expansion, VPLEX detects the capacity gained by expansion. Then, you identify the available expansion method: either storage-volume (the preferred method) or concatenation (RAID-C expansion). VPLEX can also detect the available expansion method.

You perform volume expansion using a simple, non-disruptive procedure:


1. Expand the storage volume associated with the virtual volume on the underlying storage array.
2. Allow VPLEX to rediscover the underlying storage array.
3. Expand the virtual volume using the CLI or GUI.
4. Possible values for the expansion-method attribute are:

storage-volume - VPLEX makes use of the underlying storage-volume expansion.(the corresponding LUNs on the back-end array).



concatenation (or RAID-C expansion) - The virtual volume is expanded by adding only specified extents or devices, as required.

Expanding virtual volumes using the storage-volume expansion method is only supported when a virtual volume is mapped 1:1 to the underlying storage volume.

Volume Expansion Steps – Storage Volume



② Expand the LUN on the array



EMC²


Module: Volume Management and Protection

4

Before we can expand a volume using the Storage Volume method, we first need to Plan for it. **(1)** To begin, you can list the expandable-capacity attribute (in the CLI) or the "Expandable By" field (in the GUI) to plan capacity of your back-end storage devices. For virtual volumes that can be expanded using the storage-volume method of expansion, the **expandable-capacity/Expandable By** attribute is the capacity added to the back-end storage-volume, but not yet exposed to the host by the virtual volume. A value of zero indicates that there is no expandable-capacity for the volume. Refer the expansion-method attribute to determine if storage-volume based expansion is supported. A non-zero value indicates capacity available to expand the VPLEX virtual volume. Reference the expansion-method attribute to determine if storage-volume based expansion is supported.

(2) Next you need to connect to the array and expand the underlying volume on that array. Once that is done **(3)** we have to rediscover the array from the Physical Storage > Storage Arrays view and confirm. Finally **(4)** we can browse back to the same virtual volume and see that it is now expandable.

Volume Expansion Criteria – Storage Volume



```
VPLEXcli:/clusters/cluster-1/virtual-volumes> expand -v Expansion_1096_018B_1_vol/
Virtual Volume expansion can take some time and once started, cannot be cancelled. Some operations such as
upgrades and data migrations will not be possible during the expansion. In some cases hosts and their
applications may need to be restarted once the expansion has completed.
Do you wish to proceed ? (Yes/No) yes

The expansion of virtual-volume 'Test_volume' has started..
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

5

The storage volume method of expansion supports simple and fast expansion on a variety of device geometries. In order to expand a device or add a target for expansion using the storage-volume expansion method, the VPLEX virtual volume geometry must meet one of the following criteria:


- The virtual volume is mapped 1:1 to the underlying storage volume.
- The virtual volume is a multi-legged RAID-1 or RAID-0 volume, and each of its smallest extents is mapped 1:1 to a back end storage volume.
- The virtual volume is a RAID-C (expansion through the last in-use extent or any following extent, and only through extents 1:1 mapped to storage volumes).
- Volume geometry is a combination of any of the geometries listed previously.

To perform the expansion, we click the "Expand" button in the GUI or run the **virtual-volume expand** CLI command.

RecoverPoint: If a virtual volume you want to expand is a member of a VPLEX RecoverPoint Consistency Group, then you must perform additional steps in order to use the storage volume expansion method. First, remove the volume from RecoverPoint use. Refer to the "Manage VPLEX/RecoverPoint" procedure under RecoverPoint Procedures in SolVe Desktop. Then expand the virtual volume using the storage volume expansion method before adding the volume back to the RecoverPoint Consistency Group you removed it from.

Note: Storage volumes that are not mapped 1:1 onto extents cannot expand the virtual volumes built on them using the **virtual-volume expand** command. To expand a virtual volume whose underlying storage-volume is not mapped 1:1 onto extents, perform an extent migration to an extent that is 1:1 mapped to a storage volume. Alternatively, migrate to a larger device and use the **virtual-volume expand** CLI command to expand the volume to the extra available capacity.

Volume Expansion - Concatenation via GUI



EMC²

Module: Volume Management and Protection

6

The other option for expanding a VPLEX volume is with the Concatenation method where the virtual volume is expanded by adding only

specified extents or devices, as required. To use this method simply click the expand button for a volume. The first step in the wizard is to specify the additional capacity required for the expansion. Once this has been done, choosing **Show Available** will show the list of all available storage volumes with the required capacity.

Continuing to the next step automatically completes the expansion process by creating an extent of the requested additional capacity, and concatenating that extent onto the existing device. It completes the concatenation by converting the extent into a RAID-C device containing the original device (which has been renamed) and the new, requested extent.

Volume Expansion - Concatenation via CLI

```
Vplexcli:/clusters/cluster-1/virtual-volumes/dev_9_vol> expand -v dev_9_vol -e extent_Nav-Cluster1-6gb_1
Vplexcli:/clusters/cluster-1/virtual-volumes/dev_9_vol> ll
```

Name	Value
expandable	true
expandable-capacity	0B
expansion-method	concatenation
expansion-status	-
supporting-device	device_Nav-Cluster1-6gb_1
vpd-id	VPD83T3:60001440000001070633034b593bd6c

```
Vplexcli:/clusters/cluster-1/virtual-volumes/dev_9_vol> cd ..
Vplexcli:/clusters/cluster-1/virtual-volumes> show-use-hierarchy dev_9_vol
storage-view: Nav-Storage_View (cluster-1)
virtual-volume: dev_9_vol (4G, local @ cluster-1, running)
local-device: device_Nav-Cluster1-6gb_1 (4G, raid-c, cluster-1)
local-device: device_Nav-Cluster1-6gb_12015Jun08_130620 (3G, raid-0)
extent: extent_Nav-Cluster1-6gb_1 (3G)
storage-volume: Nav-Cluster1-6gb (6G, 2 slices)
logical-unit: VPD83T3:60000970000195901096533030304237
storage-array: EMC-SYMMETRIX-195901096
extent: extent_Symm1096_00A0_1 (1G)
storage-volume: Symm1096_00A0 (3G)
logical-unit: VPD83T3:60000970000195901096533030304130
storage-array: EMC-SYMMETRIX-195901096
```

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.


Module: Volume Management and Protection

7

The CLI method of expanding a volume is shown here. The size of the extent does not affect performance.

Confirm the expansion by using the **ll** on the root volume, or using drill-down.

Volume Protection - Add a Local Mirror



EMC²

Module: Volume Management and Protection

8

To add a local mirror using Unisphere for VPLEX, navigate to Provision Storage>cluster-x>Devices. Local mirrors are added at the device level. In the first step, you select the device to mirror. This is the source copy.

Add a Local Mirror (Cont'd.)

The screenshot displays the 'cluster-1: Add Local Mirrors' wizard with the following components:

- Step 2: Select Mirrors** (Active):
 - Target Cluster: cluster-1
 - Instructions: To select appropriate mirrors automatically based on best-fit capacities, click Auto-generate Mirror Mappings. Otherwise, select each source and target to add mirror mappings manually.
 - Buttons: Find, Add Mirror, Auto-Generate Mirror Mappings (highlighted with a red box).
 - Selected Mirrors table:

Device	Virtual Volume	Mirror
device-1_1	device-1_1_vol	mirror-1_1
- Step 3: Review and Finish**:
 - Local Mirrors to Add table:

Name	Virtual Volume	Mirror
device-1_1	device-1_1_vol	mirror-1_1
- Step 4: View Results**:
 - Distributed Device Creation Results table:

Device Name	Result	Details
Add local mirror to device-1_1	✓	Success
 - Buttons: Commit, Close

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection


EMC²

9

In Step 2, the target is selected. For a local mirror, you are restricted to the same cluster. If you chose a remote mirror then you would have to select a device from the second site. It is best to choose a device of the same size. Once the mirror is made, any additional space is unusable.

Step 3 allows you to review and commit the changes. Step 4 displays the results. Once the mirror is complete, it will take some time to synchronize the new mirror set. You can review the device or use the Properties page to watch for the rebuild to complete.

Volume Protection - Add a Remote Mirror



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC¹

10

With a Remote Mirror, local devices on each cluster constitute the two legs of a distributed device. The individual local devices may include any underlying type of storage volume or geometry (RAID 0, RAID 1, or RAID C), but they should be of the same capacity.

Add a Remote Mirror (Cont'd.)

Add Remote Mirrors

1 Select Devices to Mirror
2 Select Mirrors
3 Select Consistency Gro...
4 Review and Finish
5 View Results

Click Auto-Generate Mirror Mappings to automatically select a mirror for the selected item. Or, manually select the source and target items and then click Add Mirror.

Created mappings appear in the Selected Mirrors list below. To remove a mapping from this list, select the mapping and then click Remove Mirror. To save the mirror mappings to a file on your system, click Export Mirror Mappings.

Target Cluster: cluster-2

To select appropriate mirrors automatically based on best-fit capacities, click Auto-generate Mirror Mappings. Otherwise, select each source and target to add mirror mappings manually.

Auto-Generate Mirror Mappings

Selected Devices to Mirror: Find [] [] []

Available Devices to use for Mirroring: Find [] [] []

Capacity	Virtual Volume	Name	Capacity

Add Mirror

Selected Mirrors:

Device	Virtual Volume	Mirror
device_new_Encap_c1	new_Encap	device_CLARion3040_LUN_00072_1

Show Properties Links Remove Mirror Remove All Mirrors Export Mirror Mappings

< Back Next > Finish Cancel Help

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

11

If a distributed device is configured with local devices of different capacities, the resulting distributed device will be only as large as the smallest local device; any leftover capacity on the larger device will not be available. To create a distributed device without wasting capacity, choose local devices on each cluster with the same capacity.

Add a Remote Mirror (Cont'd.)

Add Remote Mirrors

- ✓ * 1 Select Devices to Mirror
- ✓ * 2 Select Mirrors
- ✓ * 3 Select Consistency Gro...
- ✓ * 4 Review and Finish
- 5 View Results

Review the devices to create.
Click Commit to create the devices.

Source Cluster: cluster-1
Target Cluster: cluster-2
Consistency Group: esx-74

Remote Mirrors to Add:

Name	Virtual Volume	Mirror
device_new_Encap_c1	new_Encap	device_CLARIION3040_LUN_00072_1

Show Properties Links

< Back Next > Commit Cancel Help

Add Remote Mirrors

- ✓ * 1 Select Devices to Mirror
- ✓ * 2 Select Mirrors
- * 3 Select Consistency Gro...
- * 4 Review and Finish
- 5 View Results

Select an existing consistency group, or create a new one, to group an application's volumes together for ease of management.
Select a rule set if you do not want to use a consistency group.

3 Select Consistency Group or Device Rule Set

Select a Consistency Group or a Device Rule Set

☒ **Select an existing Consistency Group**

Consistency Group: * Select a Consistency Group...

☐ **Create a Consistency Group**

Consistency Group Name: *

Rule Set: * Select a Rule Set...

You must select a detach rule set to determine which cluster link failure.

☐ **Select a Device Rule Set**

Rule Set: * Select a Rule Set...

You must select a detach rule set to determine which cluster link failure.

< Back Next > Finish Cancel Help

Module: Volume Management and Protection

EMC²

12

Before finishing the Protection Wizard, you can choose to add the new Distributed Device to an existing Consistency Group or create a new one. You can also choose to just assign a particular Rule Set to the device instead of using a Consistency Group.

Volume Protection - Attach a Device Mirror

- **device attach-mirror -d <device> -m <context path/mirror> -r <rule_set>**

- Attaches a mirror to another device in a cluster
- Converts a RAID-C or RAID-0 to a RAID-1 device
- Host will be unaware of the change

```
Vplexcli:/clusters/cluster-1/devices> device attach-mirror -d Source_1 -m Mirror_2

Vplexcli:/clusters/cluster-1/devices> show-use-hierarchy -t /clusters/cluster-1/devices/Source_1
local-device: Source_1 (40G, raid-1, cluster-1)
  local-device: Mirror_2 (40G)
    extent: extent_Symm1096_04E2_1 (40G)
      storage-volume: Symm1096_04E2 (40G)
        logical-unit: VPD83T3:60000970000195901096533030344532
          storage-array: EMC-SYMMETRIX-195901096
local-device: Source_12015May12_155911 (40G, raid-0)
  extent: Mirror-A (40G)
    storage-volume: Symm1096_018B (40G)
      logical-unit: VPD83T3:60000970000195901096533030313842
        storage-array: EMC-SYMMETRIX-195901096
```

EMC²

Module: Volume Management and Protection

13

The **device attach mirror** command can be used to attach a leg to a device in the VPLEX. The execution of this command results in a RAID-1 local device or a RAID-1 distributed device.

In this example the <device> is the parent that the <mirror> attaches to. If the mirror is a local extent or device, the device will stay as a RAID-1 local device. If the mirror resides at another cluster, the device will become a distributed device.

When converting a local device to a distributed device, a rule-set can be specified. If a rule-set is not specified, the system will attempt to assign a default rule-set to the device. The system assigned default rule-set will be based on where the parent device resides.

We also use the **show-use-hierarchy** command to display the complete usage hierarchy for a storage element from the top-level element down to the storage-array.

Lesson Summary

During this lesson the following topics were covered:

- Storage-Volume Expansion Method
- Concatenation Expansion Method
- Attaching Local and Remote Mirrors

In this lesson we covered the various methods of expanding a volume on a VPLEX including attaching mirrors.

Lesson: Data Migration


This lesson covers the following topics:

- Current VPLEX data migration options and their benefits
- Performing a Migration using Unisphere for VPLEX
- Creating a Batch Migration plan and performing it

This lesson covers the various options for performing Data Migrations using VPLEX.

Data Mobility

- Moving data with VPLEX
 - Source and target can be local (Local) or remote (Metro)
 - 2 Types:
 - **Extent**
 - Within a Cluster
 - **Device**
 - Within a Cluster or
 - Between Clusters



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

16


The source and target arrays can be in the same data center (VPLEX Local) or in different data centers separated by up to 5ms (VPLEX Metro) or 50ms (VPLEX Geo). With VPLEX, source and target arrays can be heterogeneous. When you use VPLEX to move data, the data retains its original VPLEX volume identifier during and after the mobility operation. There is no change in volume identifiers to affect application cut over. The application continues to use the “same” storage, unaware that it has moved.

There are many types and reasons to move data:

- Move data from a “hot” storage device.
- Move applications from one storage device to another.
- Move operating system files from one storage device to another.
- Consolidate or move data or database instances.
- Move storage infrastructure from one physical location to another. With VPLEX, you no longer need to spend significant time and resources preparing to move data and applications. You don’t have to accept a forced outage and restart the application after the move is completed. Instead, a move can be made instantly between sites, over distance, and the data remains online and available during the move - no outage or downtime is required. Considerations before moving data include the business impact and the type of data to be moved, site locations, total amount of data, and schedules.

Extent Mobility

- Moves extents from one Storage Volume to another, non-disruptively to the host.
- Frees the original extent for reuse.
- Can be used to perform non-disruptive data mobility across or within storage arrays.



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: Volume Management and Protection

17

Extent mobility is an EMC VPLEX mechanism to move all data from a source extent to a target extent. It is completely non-disruptive to any layered devices, and completely transparent to hosts using virtual volumes involving those devices. Over time, storage volumes can become fragmented by adding and deleting extents. Extent mobility can be used to help defrag fragmented storage volumes. Extent mobility cannot occur across clusters.

Device Mobility

- Moves devices from one extent to another, freeing the original. All done non-disruptively to the host.
 - Local or Remote
- Can be used for array migrations



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

18

Unlike data on extents, the data on devices can be moved to other devices within the same cluster OR at remote cluster in a VPLEX Metro. Device migrations between Distributed Devices are not supported.


CLI Mobility Operations

```
VPlexcli:/> dm migration start -n <name> -f <extent/device> -t  
                  <extent/device>
```

```
VPlexcli:/> dm migration commit -m <name> --force
```

```
VPlexcli:/> dm migration clean -m <name> --force
```

```
VPlexcli:/> dm migration remove -m <name> --force
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²


19

There are four basic operations involved in moving extents or devices. They are: start, commit, clean, and remove.

Data mobility is accomplished by creating RAID-1 devices.

- The **start** operation first creates a RAID-1 device on top of the source device. It specifies the source device as one of its legs and the destination device as the other leg. It then copies the source device's data to the destination's device or extent. This operation can be canceled as long as it is not committed.
- The **commit** operation removes the pointer to the source leg. At this point in time the destination device is the only device accessible through the virtual volume.
- The **clean** operation breaks the source device down all the way to the storage volume level. The storage volume is unclaimed after this operation if there are no other Extents configured for this Storage Volume.
- Data mobility operations can also be paused and resumed before the commit operation. It may be beneficial to pause mobility operations during daytime hours.

Extent Mobility Wizard




Extent and device mobility jobs can be created and run from within Mobility Central. Currently the GUI does not support batched mobility. Mobility jobs can be created, removed, canceled, monitored, paused, committed, and throttled.

There are two special options within Mobility Central that can help in performing mobility jobs. When performing *device mobility* there is an option to select and move devices that are associated with a particular virtual volume. When this option is selected, only devices that are connected to the virtual volume are displayed and can be moved.

The optional setting in performing *extent mobility* is the ability to select and move extents associated with a particular storage volume. When this option is selected, only extents that are connected to the specified storage volumes are displayed and can be moved. Source and target devices and extents must be in the same cluster for mobility to work in Mobility Central.

The next few slides will cover the extent mobility process. Launch the wizard by selecting Create Extent Mobility Jobs from the Move Data Within the Local Cluster area of Mobility Central.

Extent Mobility Wizard (Cont'd.)



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

21


Step 1 selects the cluster. This can be any cluster. Step 2 is an optional step; a storage volume can be selected. If you select storage volumes, the extents on those volumes will appear on the next screen.

Extent Mobility Wizard (Cont'd.)



Step 3 selects the Extent(s) to use as the source. Step 4 selects the targets for each designated source. The Auto-generate Mappings button can be used to suggest potential targets. A Target Extent must be the same size or larger than the source. If the target is larger, the space difference cannot be utilized. The target cannot have any virtual volumes.

Extent Mobility Wizard (Cont'd.)



Step 5 configures the job name and sets the rate to perform the migration. The higher the transfer speed, the more system resources are used. This may impact other VPLEX operations. Once the job is configured, press the Start button. Step 6 displays the results.

The **transfer speed** determines the maximum number of bytes of data transferred at a time from the source to the target. The higher the speed, the greater the impact on host I/O. Monitor the mobility job's progress and its effect on host I/O. If the job is progressing too slowly, or I/O is greatly impacted, adjust the transfer speed accordingly. You can change the transfer speed of a job while the job is in the queue or in progress. The change takes effect immediately.


By default, the GUI transfer speed is set to lowest, which translates to 128 KB. This transfer speed provides good throughput while maintaining the best front-end performance in most environments.

The **Transfer size** is the CLI equivalent of the **Transfer Speed** in Unisphere. It is the region of a source element that is temporarily locked, read, and written on the target. Transfer-size can be as small as 40KB, as large as 128 MB, and must be a multiple of 4KB. The default recommended value is 128KB. A larger transfer-size results in higher performance for the migration, but may negatively impact front-end I/O. This is especially true for VPLEX Metro migrations. A smaller transfer-size results in lower performance for the migration, but creates less impact on front-end I/O and response times for hosts. Set a large transfer-size for migrations when the priority is data protection or migration performance. Set a smaller transfer-size for migrations when the priority is front-end storage response time


Transfer Speed (GUI) transfer-size attribute (CLI)

Lowest (default) -----	128 KB
Low -----	2 MB
Medium -----	8 MB
High -----	16 MB
Highest -----	32 MB

Extent Mobility



Monitor Progress



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: Volume Management and Protection

24

Now you can monitor the mobility jobs. You may have up to 25 jobs running at any one time. The GUI will allow you to monitor mobility jobs created with the CLI, including batch jobs.

Once Commit is selected, the job cannot be undone.

Device Mobility Wizard (Cont'd.)




Devices can be migrated either locally or to a remote cluster. Select Move Data Across Clusters from Mobility Central.

Use the following general steps to perform device migrations:

1. Start the migration
2. Monitor the migration's progress
3. Pause, resume, or cancel the migration (optional)
4. Commit the migration. Commit transfers the source virtual volume/device/extent to the target. If the virtual volume on top of a device has a system-assigned default name, committing a device migration renames the virtual volume after the target
5. Clean up (optional)
6. Remove the record of the migration

Device Mobility Wizard (Cont'd.)



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: Volume Management and Protection

26

Step 1 in the Device Mobility Wizard allows for the selection of a cluster. Since the wizard was launched using Across Clusters, the bottom selection will change based on the top selection.

Step 2 is the option to select Virtual Volumes. The devices that host the Virtual Volumes will appear on the next screen.

Device Mobility Wizard (Cont'd.)



Step 3 allows for the selection of devices. The list is based on the cluster selected and the available devices. One or more can be selected. In Step 4, the wizard guides the user through the selection of the "mirror" device at the other cluster. The Auto-generate Mappings button can be used to let the wizard suggest candidates for mapping based on the size of the selected devices. Once a device(s) have been selected, use the Add Mapping button to make device pairs. The selections will drop down to the lower screen, and the Next button can be used to continue.

Device Mobility Wizard (Cont'd.)

Create Device Mobility Jobs

- 1 Select Cluster
- 2 Select Virtual Volumes ...
- 3 Select Source Devices
- 4 Create Source-Target ...
- 5 Set-Up Job & Start Jobs
- 6 View Results

Need help on this step?

Change the mobility job name and transfer speed if desired.

To remove a job, select the job and then click Remove Job(s).

Click Start to create and start the mobility job.

5 Set-Up Job & Start Jobs

Rename Jobs with:

Device Mobility Jobs to be Created and Started:

Device Mobility Job Name	Transfer Speed	Capacity	Source Device Name
mobility-job3	Lowest	20.00 GB	device_Symm1057_04D8_1

Move Data Across Clusters

Moving Data From: Any Cluster To: Any Cluster With Status: Any Status

Mobility Job Name	Type	Status	Percent Done	Source	Target	Source Cluster	Target Cluster
mobility-job3	Device	In Progress	17%	device_Symm1057_04D8_1	device_nnewmapfile7_LU...	cluster-1	cluster-2

1 Selected

[Hide Properties Links](#)

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.


Module: Volume Management and Protection

28

Step 5 allows the user to change the name of the mobility job and define a transfer speed. The default name is displayed, but a more user-friendly name can be given. The transfer speed can be set to a value as previously discussed. The higher the setting the faster the devices are moved, but there is more potential impact to host performance. Click Start to begin the configured Job. Step 6 (not shown) displays the results.

Once started, the mobility job, along with others, can be monitored from Mobility Central. Also, this screen allows for the monitoring of batch migration jobs created via CLI.

Change Mobility Job Properties



EMC²


© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

29

From Mobility Central, select the mobility job link to view the properties. This allows the viewing of the source and target. From this menu the transfer speed can also be adjusted.

Device Mobility Wizard (Cont'd.)



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

30

Once the job transfer is complete it can be committed from Mobility Central. The migration process inserts a temporary RAID 1 structure above the source device/extent with the target device/extent as an out-of-date leg of the RAID 1. The migration can be understood as the synchronization of the out-of-date leg (the target). After the migration is complete, the commit step detaches the source leg of the RAID 1, and removes the RAID 1. The virtual volume, device, or extent is identical to the one before the migration except that the source device/extent is replaced with the target device/extent.

Batched Mobility with CLI

- Allows the batching of extent or device mobility jobs
 - No mix in a single batch
- Enables convenient large scale mobility operations
 - Array replacement or technology refresh
- A plan file is used
- Up to 25 jobs at a time
 - Additional jobs will be queued



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

31

Batched mobility provides the ability to script large-scale mobility operations without having to specify individual extent-by-extent or device-by-device mobility jobs. Batched mobility can only be performed in the CLI.

Batch migrations are run as batch jobs from reusable batch migration plan files. Migration plan files are created using the **create-plan** command. A single batch migration plan can be either for devices or extents, but not both.

Use batch migrations to:

- Retire storage arrays (off-lease arrays) and bring new ones online.
- Migrate devices to a different class of storage array.

There are two additional steps to prepare for a batch migration:

- Create a batch migration plan file using **batch-migrate create-plan**
- Test the batch migration plan file using **batch-migrate check-plan**

The following prerequisites are required for batch migrations:

- The source and targets are both devices or extents. Migrations between devices and extents are not supported.
- Extents must be claimed (extent migrations) or local devices must be configured (device migrations) on the target array.
- The structure of the target is the same as the structure of the source.
- For extent migrations, both source and target extents must be in the same cluster.

Create and Check the Batched Mobility Plan

```
VPlexcli:/clusters/cluster-1/devices> batch-migrate create-plan
--file=migrate.txt -f /clusters/cluster-1/devices/Dev_2,
/clusters/cluster-1/devices/Dev_3,/clusters/cluster-1/devices/Dev_4
-t /clusters/cluster-2/devices/Dev_2a,
/clusters/cluster-2/devices/Dev_3a,
/clusters/cluster-2/devices/Dev_4a --force

Extents matching source pattern: Dev_4, Dev_3, Dev_2
Extents matching target pattern: Dev_4a, Dev_3a, Dev_2a

Creating file /var/log/VPlex/cli/migrate.txt as migration plan file.
Wrote file /var/log/VPlex/cli/migrate.txt. Please review and edit this file, and run this
command in the check-plan phase afterward.

VPlexcli:/clusters/cluster-1/devices> batch-migrate check-plan --file=migrate.txt
Checking migration plan file /var/log/VPlex/cli/migrate.txt.
Plan-check passed.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

32

The **batch-migrate create-plan** command creates a migration plan using the specified sources and targets.

- **--file** - Specify the <migration-filename> filename only if that name is unique in the global namespace. Otherwise, specify a full pathname.
- **--sources** - Pathname(s) of the source extents or devices as a comma-separated list. Each element may contain wildcards. For extent migrations, source extents must be in the same cluster.
- **--targets** - Pathname(s) of the target extents or devices as a comma-separated list. For extent migrations, target extents must be in the same cluster as the source extents.
- **--force** - If a plan file with the same name already exists, forces the old plan to be overwritten.

The **batch-migrate check-plan --file <file-name>** command checks the specified batch migration plan for the following:

- Block-size of source and target extents is equal (4 Kbytes)
- Capacity of target extent is equal or bigger than the source extent's capacity
- Device migrations:
 - Target device has no volumes on it
 - Source device has volumes on it
- Extent migrations:
 - Target extent is claimed and ready for use
 - Source extent is in use

If the migration plan contains errors, a description of the errors is displayed, and the plan check fails.

Start and Cancel Batched Jobs

```
VPlexcli:/clusters/cluster-1/devices> batch-migrate start  
--file=migrate.txt --transfer-size 2M
```

```
Started device migration 'BR0_0'.  
Started device migration 'BR0_1'.  
Started device migration 'BR0_2'.  
Started 3 of 3 migrations.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

33

These commands will start and cancel a batched mobility operation.

Use the **batch-migrate start --transfer-size [40 KB-128 MB] <filename>** command to start the specified batch migration job.

- Transfer-size is the size of the region in cache used to service the migration. The area is globally locked, read at the source, and written at the target.
- Transfer-size can be as small 40KB, as large as 128 MB, and must be a multiple of 4K. The default value is 128 KB.
- A larger transfer-size results in higher performance for the migration, but may negatively impact front-end I/O. This is especially true for VPLEX Metro migrations.
- A smaller transfer-size results in lower performance for the migration, but creates less impact on front-end I/O and response times for hosts.
- Set a large transfer-size for migrations when the priority is data protection or migration performance.
- Set a smaller transfer-size for migrations when the priority is front-end storage response time.

Cancel an active batch migration to return the source volumes to their state prior to the start of the migration. Use the **batch-migrate cancel --file <filename>** command to cancel the specified migration.

Pause and Resume Batched Jobs

```
Vplexcli:/clusters/cluster-1/devices> batch-migrate pause --file=migrate.txt
WARNING: Failed to pause migration BR0_1 : Evaluation of <<dm migration pause
-m /data-migrations/device-migrations/BR0_1>> failed.
Unable to pause the given data migration(s).
Could not pause migration 'BR0_1' because it is not in-progress.
```

Paused 2 of 3 migrations

Job Pause

Job Resume

```
Vplexcli:/clusters/cluster-1/devices> batch-migrate resume --file=migrate.txt
WARNING: Failed to resume migration BR0_1 : Evaluation of <<dm migration
resume -m /data-migrations/device-migrations/BR0_1>> failed.
Unable to resume the given data migration(s).
Could not resume migration 'BR0_1' because it is not paused.
```

Resumed 2 of 3 migrations

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

34

As mentioned previously, active batch migrations (a migration that has been started) can be paused and resumed at a later time.

- Pause an active batch migration to release bandwidth for host I/O during periods of peak traffic.
- Resume the batch migration during periods of low I/O.

Monitor Batched Mobility Jobs

```
Vplexcli:/clusters/cluster-1/devices> batch-migrate summary --file=migrate.txt --verbose
source device  source cluster  target device  target cluster  migration name  status      percentage done  eta
-----
Dev_2          cluster-1        Dev_2a         cluster-2        BR0_0          in-progress  11              20.6min
Dev_3          cluster-1        Dev_3a         cluster-2        BR0_2          in-progress  11              21.8min
Dev_4          cluster-1        Dev_4a         cluster-2        BR0_1          complete    100             -
Processed 3 migrations from batch migration BR0:
  committed: 0
  complete: 1
  in-progress: 2
  queued: 0
  paused: 0
  error: 0
  cancelled: 0
  no-record: 0
```



Use the **batch-migrate summary <filename> --verbose** command to monitor the progress of the specified batch migration.

If more than 25 migrations are active at the same time, they are queued, their status is displayed as in-progress, and percentage-complete is displayed as ?.

Commit, Clean and Remove Batched Jobs

```
VPlexcli:/> batch-migrate commit --file migrate.txt

VPlexcli:/> batch-migrate clean --rename-targets --file migrate.txt
Using migration plan file /temp/migration_plans/migrate.txt for cleanup phase.
0: Deleted source extent
   /clusters/cluster-1/devices/R20061115_Symm2264_010, unclaimed its disks
   Symm2264_010
1: Deleted source extent
   /clusters/cluster-1/extents/R20061115_Symm2264_011, unclaimed its disks
   Symm2264_011

VPlexcli:/> batch-migrate remove /data-migrations/device-migrations --file
migrate.txt
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

36

After the migration is complete, the commit step detaches the source leg of the RAID 1 and then removes the RAID. The virtual volume, device or extent is identical to the one before the migration except that the source device/extent is replaced with the target device/extent. A migration must be committed in order to be cleaned.

When the batch migration is 100% complete, use **batch-migrate commit <filename>**, and then clean and remove commands as with single migration jobs to replicate the volumes on the target devices and remove the volumes from the source devices.

- For device migrations, cleaning dismantles the source device down to its storage volumes. The storage volumes no longer in use are unclaimed.
For device migrations only, use the optional **--rename-target** argument to rename the target device after the source device. If the target device is renamed, the virtual volume on top of it is also renamed - if the virtual volume has a system-assigned default name. Without renaming, the target devices retain their target names, which can make the relationship between volume and device less evident.
- For extent migrations, cleaning destroys the source extent and unclaims the underlying storage-volume if there are no extents on it.

Remove the migration record only if the migration has been committed or canceled. Migration records are in the **/data-migrations/device-migrations** context.

Monitor Batched Jobs in Mobility Central

Move Data Across Clusters

Moving Data From: Any Cluster To: Any Cluster With Status: Any Status

Mobility Job Name	Type	Status	Percent Done	Source	Target
job4_0	Device	Commit Pending	100%	device_Symm1160ENCAPSULATED...	Site2_Dev_6
job3_0	Device	Commit Pending	100%	device_Symm1160ENCAPSULATED...	Site2_Dev_5
BR0_1	Device	In Progress	15%	Dev_3	Dev_3a
BR0_2	Device	In Progress	0%	Dev_4	Dev_4a
BR0_0	Device	In Progress	0%	Dev_2	Dev_2a
O_ED_00F6_1_te2_Dev_8	Device	Committed	100%	device_Symm1160ENCAPSULATED...	Site2_Dev_8
job2_0	Device	Cancelled	100%	device_Symm1160ENCAPSULATED...	Site2_Dev_7
O_ED_00F7_1_te2_Dev_7	Device	Cancelled	100%	device_Symm1160ENCAPSULATED...	Site2_Dev_7

0 Selected [Pause](#) [Resume](#) [Commit](#) [Cancel](#) [Remove](#)

[Create Device Mobility Jobs](#) [Hide Properties Links](#) Last Refresh: 09-12-2012 10:39

User: service Connected to Clusters: cluster-1 Product Type: EMC VPLEX™ Metro 2012-09-12 10:39 (GMT+05:30)

EMC

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

37

Batched Mobility jobs can be monitored in Mobility Central. Batched mobility jobs cannot be started in the GUI. However, they can be paused, cancel, resumed, and committed, and their transfer speed can be changed.

Lab: Mobility with VPLEX



This lab covers the following:

- Moving an Extent
- Moving a Device
- Batch Mobility

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

38

This lab covers mobility with VPLEX.

Lesson Summary

During this lesson the following topics were covered:

- Data Mobility (Extent, Device)
- Batched Mobility

In this lesson we covered the various methods of performing Data Mobility within VPLEX using both the CLI and the GUI.

Lesson: Data Protection with RecoverPoint

This lesson covers the following topics:

- Overview of RecoverPoint functions and topologies
- Methods of VPLEX Volume protection using RecoverPoint
 - Local
 - MetroPoint

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

40

This lesson covers VPLEX integration with EMC RecoverPoint.

RecoverPoint Use Cases

- Disaster Recovery
 - Protection against disasters impacting an entire data center or site
 - Recover inter-dependent applications in CG set with data consistency to any PiT
- Operational Recovery
 - Protect data against errors, disasters, or attacks
 - Recover any application data set to a point in time before an error
- Data Migration
 - Replicate data for test or development purposes
 - Replicate data for technology refresh, data center migration, or remote site implementation.

The Hard Statistics of Outages:

Human Error = 60%


Patches & Updates = 56%

Computer Room Issues = 44%

Power Outage = 29%

Fire, Flood or Explosion = 26%

Earthquakes = 10%



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

41

Vital systems and data need protection against human errors, operational and procedural errors, data corruption, virus attacks, systems failures, utility outages and other outages caused by natural disasters. RecoverPoint's continuous data protection, local or remote, enables recovery to any point in time, providing a finer granularity to achieve desired RPOs.

For mission critical applications (SAP, Oracle, and Microsoft SQL) that require near zero data loss and minimal disruption from down time, the application owners and IT managers have to address how to close the RPO and RTO gap. RecoverPoint is an ideal solution to help address this challenge. Furthermore, RecoverPoint Consistency Groups provide data consistency for recovery, across multiple applications and multiple arrays of heterogeneous types.

Finally, Customers can also replicate data for isolated test and development purposes or replicate data during their technology refresh, or in the case of data center expansion or remote site projects.

RecoverPoint Solution (Cont'd.)

MetroPoint for VPLEX



RecoverPoint 4.1 introduced MetroPoint, a disaster recovery topology for VPLEX Metro, combining RecoverPoint 4.1 and EMC VPLEX 5.4. MetroPoint Topology. This combination provides a seamless, multi-site continuous data protection solution with built-in disaster recovery support. With MetroPoint, the two active-active clusters of a VPLEX Metro configuration are both protected by RecoverPoint.


This arrangement enables data replication from an EMC VPLEX Metro region of two data centers to a third site over distance and provides disaster recovery that can sustain two-site failures.

The MetroPoint topology can maintain up to five copies of data, including one remote copy and one local copy at each VPLEX site. It provides protection for each of the VPLEX Metro sites with continuous local replication via RecoverPoint.

Combining VPLEX Metro with RecoverPoint allows the user to achieve zero Recovery Point Objective (RPO) and any point-in-time (PiT) recovery.

For more information, please see the publication EMC RecoverPoint Deploying with VPLEX Technical Notes.

Two site – VPLEX Local + RecoverPoint Topology



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection


EMC²

43

In this Example we see a two-site simple RecoverPoint – VPLEX Solution where VPLEX only has a single site and uses a two-site RecoverPoint solution to replicate data from the Hopkinton site with VPLEX and RecoverPoint to an off-site DR location with RecoverPoint.

MetroPoint: Continuous Availability + Protection

- 3-site solution for the most demanding environments
- Continuous Availability in two separate regions
- Each region provides DR protection for the other




MetroPoint provides the highest levels of protection and redundancy against full site failures, regional-wide failures, data corruption, and other such events. It combines the full value of VPLEX and RecoverPoint into a 3-site or 4-site topology for continuous availability, disaster recovery, and continuous data protection.

In a MetroPoint topology, a single consistency group spans across three related sites. The consistency group includes an active production copy and a standby production copy. At all times, only the active production copy replicates to the remote site (there is no communication pipe between the standby production copy and the target). The standby production copy will be in marking mode. Using periodic sync-peer-splitter communication with the active production copy, the standby production copy erases delta marker information that is no longer needed.

MetroPoint Consistency Groups

- 3-Cluster Example
- Production volume is a VPLEX distributed device
- CG has 3 copies:
 - 2 Production copies (A+B)
 - Replica copy (C)
- Production copies differ in functionality:
 - Primary (A)
 - Secondary (B)
- Replication is from primary production to replicas (A → C)



EMC²

Module: Volume Management and Protection

45

A VPLEX MetroPoint Consistency Group can contain a maximum of one remote copy and from zero to two local copies. If there is no remote copy, there must be two local copies, one at each side of the VPLEX Metro.

A Distributed Device fracture is defined as when the communication link between the two VPLEX Metro sites is down. When a fracture occurs between the two VPLEX sites, RecoverPoint replication will continue through the VPLEX winner site. The preferred RecoverPoint cluster can be chosen manually in RecoverPoint. Alternatively, the preferred RecoverPoint cluster can be set to follow the VPLEX detach rules for the VPLEX Consistency Group. If the preferred RecoverPoint cluster is not at the VPLEX winning site, a production switchover will occur.

During the fracture, RecoverPoint at the VPLEX winning site will continue to replicate to the remote site. The losing site will not receive any I/Os during the fracture and will not maintain an up-to-date journal. If there is a local copy on the losing site, it will not be updated during the fracture. If there is a local copy at the winning site, it will continue to be updated as usual. Once communication between the two VPLEX clusters is restored, the two VPLEX clusters synchronize their data according to the I/Os that occurred during the fracture. RecoverPoint replication continues during this synchronizing process. When the VPLEX clusters are fully synchronized, if there is a local copy at the losing site, RecoverPoint synchronizes it with the winning site and then resumes replicating to it.

Lesson Summary

During this lesson the following topics were covered:

- RecoverPoint Overview
- VPLEX and RecoverPoint
- MetroPoint

In this lesson we covered how RecoverPoint can get integrated into a VPLEX environment to fully utilize features of both products.

Lesson: Adding RecoverPoint protection to VPLEX


This lesson covers the following topics:

- Configuration of RecoverPoint protection for VPLEX Virtual Volumes

This lesson covers VPLEX integration with EMC RecoverPoint.

Create RecoverPoint Repository Virtual Volume

- 2.86 GB or larger
- Non-Thin Storage Volume
- Raid-1 on separate Arrays



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC

48

A repository volume holds the license key and metadata for a given RPA cluster. One repository volume is required per RPA cluster.

Existing RecoverPoint configurations may already have a repository volume configured. For these configurations, continue to use the existing repository volume.

For new RecoverPoint deployments:

To use VPLEX-hosted storage for a repository volume, select or configure a virtual volume with the following properties:

- 2.86 GB. Capacity beyond 2.86 GB is not used.
- RAID-1 device based on volumes in two different arrays
- Provisioned on non-thin storage
- Provisioned on highly available storage
- Local device only (not a DR1)
- NOT a remote exported volume

Creating RecoverPoint Enabled Consistency Groups

- Unique Groups for RP Cluster and Application Virtual Volumes
- Each must be RP Enabled

RecoverPoint Cluster Group must be Local

cluster-1: Create Consistency Group

1 Select Name and Type

Consistency Group Name: RP-Backup-System

What type of virtual volume would you like to add to the consistency group?

☒ Local: cluster-1

Local consistency groups will be synchronous and visible only to the local cluster unless you select the Global visibility checkbox. You can add local or global virtual volumes.

☐ Global visibility

Consistency groups with global visibility will be synchronous and visible to all clusters.

```
Vplexcli:/clusters/cluster-1/consistency-groups/esx-74> set recoverpoint-enabled true
Vplexcli:/clusters/cluster-2/consistency-groups/esx-74-replica> set recoverpoint-enabled true
```

EMC


Module: Volume Management and Protection

49

Each VPLEX Consistency Group must have the recoverpoint-enabled value set to true. This is the case for all Consistency Groups on VPLEX using RecoverPoint Protection.

Provision Storage to the RPAs

- Register RPA Initiators
- Create Storage View



The RecoverPoint Cluster consists of anywhere from 2-8 RPAs. Each RPA will have 4 Initiators. They must be registered as RecoverPoint initiators and placed into a Storage View.

All the initiators for a RP Cluster must be in the same Storage View. Place all the Virtual Volumes used for RecoverPoint Replication in the View. This includes:

- Repository
- Journals
- Production
- Copies

Import security certificate to VPLEX

```
VPlexcli:/> rp import-certificate
```

This command will cause the VPLEX CLI process to restart if security settings are modified. This will require a new log in from all connected CLI and GUI clients.

To proceed type CONTINUE or hit enter to abort:CONTINUE

Please enter the IPv4 address of the RecoverPoint cluster:10.6.210.85

Use the displayed command to add the Certificate from the RP Cluster to VPLEX.

Adding RecoverPoint to a VPLEX environment

Before:

- Verify version of VPLEX installed supports RP or MetroPoint
- Create RP / MetroPoint Consistency Groups in RecoverPoint

```
Vplexcli: > rp rpa-cluster add --host=10.6.210.87
Enter rpa-cluster administrative password: <admin-password>
Enter rpa-cluster administrative password again: <admin-password>

Vplexcli:/> ll /recoverpoint/rpa-clusters/recoverpoint/rpa-clusters
RPA Host      VPLEX Cluster RPA Site      RPA ID  RPA Version
-----
10.6.210.87   cluster-1     Franklin-1     RPA 1   4.1(d.147)
10.6.211.3    cluster-2     Franklin-2     RPA 1   4.1(d.147)
```


EMC²

Module: Volume Management and Protection

52

After the RecoverPoint clusters start communicating, use the **rp rpa-cluster add** command on the VPLEX and add the clusters that are local to VPLEX Metro to the VPLEX system. Ensure that you specify which VPLEX cluster is the local cluster.

MetroPoint Consistency Group



© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

53

Displayed is an example of a MetroPoint Consistency Group with only local copies shown in Unisphere for RecoverPoint

Lesson Summary

During this lesson the following topics were covered:

- Configuration of RecoverPoint protection for VPLEX Virtual Volumes

In this lesson we covered how RecoverPoint can get integrated into a VPLEX environment to fully utilize features of both products.

Module Summary

In this module we covered:

- Volume Expansion
- Adding Device Mirrors
- Data Migration
 - Extent
 - Device
 - Batch
- Data Protection using VPLEX with RecoverPoint

This module focused on using VPLEX for Data Mobility.

Lab: Configure RecoverPoint Protection



This lab covers the following:

- Create RecoverPoint Storage View
- Create RecoverPoint Consistency Groups
- Add a RecoverPoint Cluster
- Validate the configuration

© Copyright 2015 EMC Corporation. All rights reserved.

Module: Volume Management and Protection

EMC²

56

This lab covers mobility with VPLEX.

Module: VPLEX Monitoring

Upon completion of this module, you should be able to:

- Describe how VPLEX sends events and system reports to an ESRS Gateway
- Describe SNMP integration options
- Set up monitors in VPLEX
- Analyze monitoring data in VPLEX

This module focuses on monitoring a VPLEX environment.

Lesson: VPLEX Support Integration

This lesson covers the following topics:

- Describe how VPLEX sends events and system reports to an ESRS Gateway
- Describe SNMP integration options

This lesson covers how to integrate support options into VPLEX.

Configuring ESRS Gateway

Required information from VPLEX

- VPLEX Model type (VS1 or VS2)
- IP address of the Management Server
- VPLEX Serial Number (Top Level Assembly #)
- Site ID



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring




3

EMC Secure Remote Services (ESRS) is a two-way remote connection between EMC Customer Service and your EMC products and solutions. Available with VPLEX, ESRS offers enhanced or premium warranties and maintenance agreements at no additional cost. ESRS maintains connectivity with your EMC products around the clock and automatically notifies EMC if a problem occurs. If troubleshooting is necessary, and authorized and authenticated EMC Customer Service professional uses the secure connection to establish a remote session to diagnose, and if necessary, to repair your EMC product.

VPLEX is supported on ESRS 1.04.00 under the VS1 product family and requires a Manual Deployment using 1 X IP of the management station.

Call Home and System Reports (SYR)

1. **VPLEX Engine** sends events to the management server
2. **Call Home Listener** determines if the event requires **Call Home**
3. **EMA_Adapator** converts the .TXT files to .XML
4. **ConnectEMC** sends events and SYR reports to the ESRS Gateway



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

4

The processes involved in Call Home and SYR on the director are ECOM, nsfw, and ZPEM. The processes involved on the Management Server are ConnectEMC, Call Home Listener, SYR, EMA_Adapator, and VPlexcli. Nsfw, which is the director firmware, and ZPEM, which is the process that monitors hardware, send event data to their director's system log. ECOM, which also runs on the director, collects events from the system log. VPlexcli runs on the Management Server and pulls events every second from ECOM.

The Call Home Listener process looks at the events and determines which events can trigger a call home. It then places the call home -able events into the **/opt/emc/VPlex/Event_Msg_Folder as .txt** files. The EMA_adapator's job is to take the text files from the Event_Msg_Folder and change them to XML files. The EMA_adapator then places those files into the **/opt/emc/connectemc/poll** directory. The ConnectEMC process picks up the xml event files and send them to ESRS Gateway. If the events are successful in reaching the gateway, they are also copied into the **/opt/emc/connectemc/archive** directory. If the events are unsuccessful, they are placed into the **/opt/emc/connectemc/failed** directory. TCP ports 22, 9010, 443, and 5901 must be open between VPLEX and the ESRS Gateway. The TLA must be set on the cluster in order for the ESRS Gateway to recognize the event coming in.

SNMP Overview

- SNMP can retrieve performance statistics for a single cluster
- One-time SNMP agent configuration is necessary
- Call-home event notifications sent
- Third-party management station often used



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

5

An SNMP Management Station uses UDP for connectivity to the SNMP Agent running on the VPLEX Management Server through port 161 on the Management Server's public IP address. One-time configuration of the SNMP Agent is necessary prior to accessing performance statistics.

The greatest value proposition for SNMP is achieved when using a third-party SNMP Management Station to poll for statistics on a periodic basis. Historical statistics must be tracked in this manner since the VPLEX directors can only hold recent statistics in memory.

SNMP is a protocol that enables the monitoring of network-attached devices in accordance with the variables specified in Management Information Bases (MIBs). MIBs contain the object identifiers (OIDs) of the variables that can be monitored through SNMP in a hierarchical format.

SNMP is supported on VPLEX to retrieve performance statistics for a single cluster. A single SNMP Agent cannot retrieve statistics from both clusters in a VPLEX Metro. SNMP version 2c is supported; version 3 is not supported. SNMP allows two capabilities for monitoring: the ability to actively poll for statistics and to send call-home traps.

SNMP polling is often performed by a third-party management station. The SNMP Management Station periodically sends queries to the SNMP Agent over the public IP network. These management stations can provide historical tracking and alerting capabilities.

SNMP traps are used to send notifications when certain call-home conditions are met. VPLEX does not support trap notifications regarding performance statistics. Configuring an SNMP trap sink will send notifications pertaining to call-home events to a specified remote host. SNMP version 2c is supported for SNMP traps.

There are several statistics in the VPLEX-MIB table that have a static value of zero; these statistics are currently not supported with VPLEX. The MIB table is located on the Management Server in **/opt/emc/Vplex/mibs/VPLEX-MIB.mib**. All statistics are read-only.

SNMP Polling Relative Components

- VPLEX Management Server
 - The SNMP Agent runs here
- SNMP Management Server
 - Can poll the SNMP Agent from the SNMP Management station
 - The SNMP Agent returns stats to the SNMP Management station
- SNMP Agent
 - Polls the directors of a single cluster and returns stats to the SNMP Management station
- VPLEX Directors
 - VPLEX Directors are polled by the SNMP Agent

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

6

There are several components that are relative in the discussion of SNMP for VPLEX.


The VPLEX Management Server is the VPLEX component that is used to manage a VPLEX system. Both the SNMP Agent and information pertaining the MIB structure reside on the VPLEX Management Server. MIBs can be downloaded from **support.emc.com**

The SNMP Management Station is a server that can automatically or manually query the SNMP Agent across port 161 over the public IP network. Performance statistics attained from the directors in a cluster are returned from the SNMP Agent to the Management Station.

The SNMP Agent runs on the Management Server and polls the directors of a single cluster. The Agent returns statistics to the management station.

The VPLEX directors are the components specifically polled by the SNMP Agent. Only Directors in the cluster locally-attached to the Management Server can be polled.

SNMP Operations



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

EMC²

7

SNMP commands executed from an SNMP Management Station commence a series of events on the local VPLEX cluster.

1. An **SNMPGET** command is launched from the SNMP Management station. This command is sent to the public IP address of the VPLEX Management Server over port 161 and forwarded to the SNMP Agent process running on the VPLEX Management Server.
2. The SNMP Agent, serving as the advocate for all director statistics, then communicates with each director in the local cluster only to retrieve the most recent statistic given the OID specified. If the command was an **SNMPGETBULK**, the Agent would retrieve all recent statistics.
3. All statistics are sent back to the Agent.
4. The statistics are then sent to the SNMP Management Station from which the original command was issued.

Supported SNMP Polling Commands

- **snmp-agent configure**
 - Configures the SNMP Agent
 - One time configuration
 - Executed in the VPLEX CLI
- **SNMPGET**
 - Gets the most recent statistics for the OID specified from each director
- **SNMPGETNEXT**
 - Gets the most recent statistics for the next OID from each director
- **SNMPGETBULK**
 - Gets all SNMP statistics from each director



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

8

The relevant commands for SNMP polling are displayed here. The **snmp-agent configure** command initializes the SNMP Agent on the VPLEX Management Server and must be executed prior to usage of SNMP polling.

SNMP commands are often executed from a third-party SNMP management station. The SNMP management station pings the SNMP Agent through the VPLEX management server over the public IP network. The **SNMPGET** command retrieves the most recent statistic from each director for the specified OID. The **SNMPGETNEXT** command returns the most recent statistic for the next OID. The **SNMPGETBULK** command returns all SNMP statistics from each director. SNMPGET, SNMPGETNEXT, and SNMPGETBULK are currently the only SNMP commands that can be issued from a Management Station.

Supported SNMP Trap Commands

- `notifications snmp-trap create`
 - Creates an SNMP trap sink for call-home events
- `set remote host`
 - Configures the trap destination
- `set started true`
 - Begins SNMP trap monitor
- `notifications snmp-trap destroy`
 - Deletes one or multiple SNMP traps



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

9

SNMP Traps are configured in the VPLEX CLI. The notifications **snmp-trap create** command names a construct that will serve as the SNMP sink. The **set remote host** command must then be executed to instruct this sink to send notifications to a specific server. Lastly, the **set started true** command starts the SNMP trap monitor. To delete an SNMP trap, the notifications **snmp-trap destroy** command can be used.

Lesson Summary

During this lesson the following topics were covered:

- VPLEX ESRS Integration
- Using SNMP for Monitoring a VPLEX

In this lesson we covered how to connect your VPLEX to EMC support using the ESRS Gateway as well as how to configure SNMP for monitoring.

Lesson: VPLEX Monitoring

This lesson covers the following topics:

- Use the health-check commands to verify the environment
- Collect system data operations
- View VPLEX performance data



This lesson covers VPLEX monitoring commands. The commands used to monitor metadata volumes, VPLEX objects, directors, storage views, users, clusters, VPN, ports, and distributed devices will be covered. In addition, command options for general and detailed monitoring of storage volumes, extents, local devices, virtual volumes, and storage views will be covered.

Monitor Clusters

```
VPlexcli:/> cluster summary
Clusters:
  Name      Cluster ID  Connected  Expelled  Operational Status  Health State
  -----
cluster-1   1    true      false     ok              ok
cluster-2   2    true      false     ok              ok

Islands:
  Island ID  Clusters
  -----
1 cluster-1, cluster-2

VPlexcli:/> cluster status
Cluster cluster-1
  operational-status: ok
  transitioning-indications:
  transitioning-progress:
  health-state: ok
  health-indications:

Cluster cluster-2
  operational-status: ok
  transitioning-indications:
  transitioning-progress:
  health-state: ok
  health-indications:
```



The **cluster summary** command and the **cluster status** command are good commands to run to monitor VPlex clusters. Connected means that the Management Server is connected to the cluster. The Expelled column checks if the noted cluster has lost its connection to the other clusters. The Operational Status describes how the cluster is functioning. The Health State describes the overall health state of the cluster.

Validate the VPN Status

```
VPlexcli:/> vpn status  
Verifying the VPN status between the management servers...  
IPSEC is UP  
Remote Management Server at IP Address 10.31.25.27 is reachable  
Remote Internal Gateway addresses are reachable  
  
Verifying the VPN status between the management server and the  
cluster witness server...  
IPSEC is UP  
Cluster Witness Server at IP Address 128.221.254.3 is reachable
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

13

The **vpn status** command verifies the VPN connection between management servers.

You can alter the state of the VPN with these commands:

vpn restart

vpn start

vpn stop

Export Port Summary

```
VPlexcli:/> export port summary
```

```
Port health summary(cluster-1):
```

port name	health state	enabled	views	virtual-volumes
P0000000047606330-A0-FC00	healthy	true	1	7
P0000000047606330-A0-FC01	healthy	true	1	3
P0000000047606330-A0-FC02	healthy	true	1	2
P0000000047606330-A0-FC03	healthy	true	0	0
P0000000047706330-B0-FC00	healthy	true	1	7
P0000000047706330-B0-FC01	healthy	true	1	3
P0000000047706330-B0-FC02	healthy	true	1	2
P0000000047706330-B0-FC03	healthy	true	0	0

```
Total 8 ports, 0 unhealthy.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

14

Here we see a simple command to allow you to check the Status of VPlex Front-End Ports.

Director Connectivity

```

Vplexcli:/> director uptime
  Director director-2-1-B: 662 days, 23 hours, 0 minutes, 50 seconds.
  Director director-2-1-A: 662 days, 23 hours, 9 minutes, 48 seconds.
  Director director-1-1-B: 348 days, 20 hours, 42 minutes, 27 seconds.
  Director director-1-1-A: 348 days, 20 hours, 42 minutes, 23 seconds.
Vplexcli:/> connectivity director director-1-1-A
Device VPD83T3:60000970000195901096533030303238 is a default LUN_0.
StorageVolumes discovered - sorted by: name
StorageVolume Name                                WWN/IQN                                LUN                                Ports
-----
VPD83T3:60000970000195901096533030304130  0x50000097310112180  0x0001000000000000  A1-FC00
                                                    0x50000097310112184  0x0001000000000000  A1-FC01

Initiators discovered
Node WWN      Port WWN/IQN      Ports
-----
0x20000000c9d2f331  0x10000000c9d2f331  A0-FC00
0x20000000c9c878d4  0x10000000c9c878d4  A0-FC00

Directors discovered by director-1-1-A, UUID 0x0000000047606330:
Director UUID      Protocol Address      Ports
-----
0x0000000047606368  COMSCSI  0x50001442a0636820  A2-FC00
COMSCSI  0x50001442a0636821  A2-FC01
  
```

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

15

The **director uptime** command displays the amount of time a director has been online.

The **connectivity director <director name>** command displays everything that is connected to a particular director. The objects that are shown are storage volumes, initiators, and other directors.

Monitor Users

- sessions command
 - Displays all active VPLEX Management Server sessions
 - Time the session was started
 - Username using the session
 - Hostname that created the session

```
Vplexcli:/> sessions
Type  Username  Hostname      Port  Creation Time
-----
TELNET_SHELL  service    localhost     4917  Thu Jan 21 16:56:26 UTC 2015
WEB_UI        service    10.12.247.23  54150 Thu Jan 21 15:59:27 UTC 2015
DEFAULT_SHELL -          -            -      Thu May 19 17:17:40 UTC 2011
```



The **sessions** command displays information on users who are logged into the VPLEX Management Console.

Monitor System Volumes

- `ll /**/<cluster>/system-volumes`
 - List the system-volumes on the cluster
 - Identifies if a metadata volume or logging volume is healthy
 - Operational Status and Health State should be ok
 - Active should be true

```

VPlexcli:/> ll /**/system-volumes
/clusters/cluster-1/system-volumes:
Name                               Volume Type  Oper.  Health  Active  Ready  Geometry  Block  Block  Capacity  Slots
-----
POD5_LoggingVol_vol               logging-volume ok      ok      -      -      raid-1    2568000 4K      9.8G      -
Pod5_Meta                         meta-volume  ok      ok      true    true    raid-1    20639744 4K      78.7G    32000
Pod5_Meta_backup_2011Jul20        meta-volume  ok      ok      false   true    raid-1    20639744 4K      78.7G    32000
Pod5_Meta_backup_2011Jul21        meta-volume  ok      ok      false   true    raid-1    20639744 4K      78.7G    32000

/clusters/cluster-2/system-volumes:
Name                               Volume Type  Oper.  Health  Active  Ready  Geometry  Block  Block  Capacity  Slots
-----
POD6_LoggingVol_vol               logging-volume ok      ok      -      -      raid-1    2568000 4K      9.8G      -
Pod6_Meta                         meta-volume  ok      ok      true    true    raid-1    20639744 4K      78.7G    32000
  
```

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

17

A metadata volume can be monitored by listing the system-volumes directory. If the metadata volume is having trouble, it will be shown here. Note that a cluster will not allow configuration changes to the metadata volume if both mirror legs of a configured metadata volume have failed. I/O to existing virtual volumes will not be affected.

Default Report - capacity-arrays

```
VPlxcli:/> report capacity-arrays
Local Storage (in GiB):
  cluster-1:
    EMC-SYMMETRIX-194900994:
      tier:          no-tier
      alloc:         23.775
      unalloc devices: 162.415
      unalloc storage-volumes: 793.579
  cluster-2:
    EMC-CLARiiON-FNM00101100041:
      tier:          no-tier
      alloc:         39.551
      unalloc devices: 34.607
      unalloc storage-volumes: 990.601
Shared Storage (in GiB):
  alloc:         29.663
  unalloc devices: 0.000

Tier summary (local, by cluster):
  cluster-1
    Tier          Allocated (GiB)    Unallocated Devices (GiB)    Unallocated StorageVolumes (GiB)
    no-tier          23.775             162.415             793.579
  cluster-2
    Unallocated          Unallocated
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

18

This monitor generates a capacity report for all the storage in a VPLEX system, grouped by storage arrays. It assumes:

- All storage volumes in a storage array have the same tier value.
- The tier is indicated in the storage-volume name. The tier attribute in the virtual volumes context is ignored.

Tier IDs are required to determine the tier of a storage-volume/storagearray. Storage volumes that do not contain any of the specified IDs are given the tier value 'no-tier'.

The report is separated into two parts:

- **Local storage** refers to storage volumes where the data is physically located at one site only, and includes:
 - Cluster id
 - Storage array
 - The tier of the storage array
 - Allocated - storage that is visible through a view (exported)
 - Unallocated-device: storage capacity that is in devices, but not visible from a view
 - Unallocated-storage: volume capacity in unused storage volumes
- **Shared storage** refers to distributed and remote virtual volumes and includes:
 - Allocated: storage that is visible through a view (exported)
 - Unallocated-device: storage that is in devices, but not visible from a view. For example, a virtual volume that has not been exported or free space in a device that is not part of a virtual volume

Default Report - capacity-clusters

```
VPlexcli:/> report capacity-clusters
```

```
Cluster, Unclaimed disk capacity (GiB), Unclaimed storage_volumes, Claimed  
disk capacity(GiB), Claimed storage_volumes, Used storage-volume capacity  
(GiB), Used storage_volumes, Unexported volume capacity (GiB), Unexported  
volumes, Exported volume capacity (GiB), Exported volumes
```

```
cluster-1, 754.21, 107, 39.37, 6, 261.11, 21, 177.25, 6, 38.61, 7  
cluster-2, 503.54, 94, 487.06, 29, 113.53, 21, 9.89, 2, 59.33, 12
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

19

Information generated from the **report capacity-clusters** monitor includes:

- Unclaimed storage-volume capacity in GB
- Number of unclaimed storage volumes
- Claimed storage-volume capacity in GB
- Number of claimed storage volumes
- Used storage-volume capacity in GB
- Number of used storage volumes
- Unexported virtual volume capacity in GB
- Number of unexported virtual volumes
- Exported virtual volume capacity in GB
- Number of exported virtual volumes

Default Report - capacity-hosts

```
VPlexcli:/> report capacity-hosts
```

```
Cluster, Views, Exported capacity (GiB), Exported volumes
```

```
cluster-1, 2, 38.61, 7
```

```
cluster-2, 1, 59.33, 12
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

20

Information generated by this monitor includes:

- Number of views.
- Total exported capacity in GB.
- Number of exported virtual volumes per cluster.

Monitor VPLEX Object Constructs

- **show-use-hierarchy**
 - Display the complete usage hierarchy for a storage element from the top-level element down to the storage-array.

```
Vplexcli:/clusters/cluster-1/devices> show-use-hierarchy -t device_encap_1
storage-view: Student1_S (cluster-1)
  virtual-volume: device_encap_1_vol (15G, local @ cluster-1, running)
    local-device: device_encap_1 (15G, raid-0, cluster-1)
      extent: extent_encap_1 (15G)
        storage-volume: encap (15G)
          logical-unit: VPD83T3:60000970000195901096533030304239
            storage-array: EMC-SYMMETRIX-195901096
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

21

The **show-use-hierarchy** command is similar to the **drill-down** command in that it drills from the specified target up to the top-level volume and down to the storage-array. The command will detect sliced elements, drill up through all slices and indicate in the output that slices were detected. The original target is highlighted in the output.

You can specify meta, logging and virtual volumes, local and distributed devices, extents, storage-volumes or logical-units on a single command line.

Monitor Storage Volume Details

```
Vplexcli:/> ll **/storage-volumes
```

```
/clusters/cluster-1/storage-elements/storage-volumes:
```

Name	VPD83 ID	Capacity	Use	Vendor	IO	Type	Thin
-----	-----	-----	-----	-----	-----	-----	-----
Encapsulated_001	VPD83T3:94533030314346	4.94G	claimed	EMC	alive	data-protected	false
POD5_LoggingVol_1	VPD83T3:94533030304542	9.8G	used	EMC	alive	normal	false
POD5_LoggingVol_2	VPD83T3:94533030304543	9.8G	used	EMC	alive	normal	false
Symm0994_00B3	VPD83T3:94533030304233	78.7G	used	EMC	alive	normal	true
Symm0994_00B4	VPD83T3:94533030304234	78.7G	used	EMC	alive	normal	true
Symm0994_00BE	VPD83T3:94533030304245	4.94G	used	EMC	alive	normal	false
Symm0994_00BF	VPD83T3:94533030304246	4.94G	used	EMC	alive	normal	false
Symm0994_00C0	VPD83T3:94533030304330	4.94G	used	EMC	alive	normal	false
Symm0994_00C1	VPD83T3:94533030304331	4.94G	used	EMC	alive	normal	false
Symm0994_00C2	VPD83T3:94533030304332	4.94G	used	EMC	alive	normal	false
Symm0994_00C3	VPD83T3:94533030304333	4.94G	used	EMC	alive	normal	false
Symm0994_00C6	VPD83T3:94533030304336	4.94G	claimed	EMC	alive	normal	false

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

EMC²

22

This command provides additional details on storage volumes including: storage volume name and VPD #, capacity of each storage volume, use of each storage volume, the storage volume vendor, the storage volume I/O status, and the storage volume type.

Monitor Storage Volumes

```
Vplexcli:/> storage-volume summary
SUMMARY (cluster-1)
Storage-Volume Summary (no tier)
-----
Health      out-of-date    0
            storage-volumes 138
            unhealthy 0
Vendor      EMC           138
Use         claimed    6
            meta-data  4
            unclaimed  107
            used      21
Capacity    total      1.34T
SUMMARY (cluster-2)
Storage-Volume Summary (no tier)
-----
Health      out-of-date    0
            storage-volumes 148
            unhealthy 0
Vendor      DGC           148
Use         claimed    29
            meta-data  4
            unclaimed  94
```



The **storage-volume summary** command provides information on the number of used, claimed, and unhealthy storage volumes. It also displays the number of metavolumes, the storage volume vendor, and the total capacity of storage volumes.

Monitor Extent Details

```
VPlexcli:/> ll /**/extents
```

```
/clusters/cluster-1/storage-elements/extents:
```

Name	StorageVolume	Capacity	Use
-----	-----	-----	-----
extent_POD5_LoggingVol_1_1	POD5_LoggingVol_1	9.8G	logging
extent_POD5_LoggingVol_2_1	POD5_LoggingVol_2	9.8G	logging
extent_Symm0994Tier-1_00B3_1	Symm0994Tier-1_00B3	78.7G	used
extent_Symm0994Tier-1_00B4_1	Symm0994Tier-1_00B4	78.7G	used
extent_Symm0994Tier-1_00B5_1	Symm0994Tier-1_00B5	4.94G	claimed
extent_Symm0994Tier-1_00B6_1	Symm0994Tier-1_00B6	4.94G	used
extent_Symm0994Tier-1_00B7_1	Symm0994Tier-1_00B7	4.94G	used
extent_Symm0994Tier-1_00B8_1	Symm0994Tier-1_00B8	4.94G	used
extent_Symm0994Tier-1_00B9_1	Symm0994Tier-1_00B9	4.94G	used
extent_Symm0994Tier-1_00BA_1	Symm0994Tier-1_00BA	4.94G	used
extent_Symm0994Tier-1_00BB_1	Symm0994Tier-1_00BB	4.94G	used
extent_Symm0994Tier-1_00BC_1	Symm0994Tier-1_00BC	4.94G	used
extent_Symm0994Tier-1_00BD_1	Symm0994Tier-1_00BD	4.94G	used
extent_Symm0994_00BE_1	Symm0994_00BE	4.94G	used



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

24

This command provides more information on the extents in a particular cluster.

Monitor Extents

```
Vplexcli:/> extent summary
SUMMARY (cluster-1)
Extent Summary
-----
Health extents      21
  out-of-date       0
  unhealthy          0
Use      used      16
  claimed           3
  unclaimed         0
  unusable          0
  logging           2
Capacity      total      260G

SUMMARY (cluster-2)
Extent Summary
-----
Health extents      21
  out-of-date       0
  unhealthy          0
Use      used      19
  claimed           0
  unclaimed         0
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

25

The **extent summary** command displays health state of the extents as well as the number used, claimed, unclaimed, and unusable.

Monitor Local Device Details

Vplexcli:/> **ll **/devices**

/clusters/cluster-1/devices:									
Name	Oper. Status	Health State	Block Count	Block Size	Capacity	Geometry	Visibility	Transfer Size	Virtual Volume

MIGRATE_D-1_00B9_1-1_00BC_1	ok	ok	1296000	4K	4.94G	raid-1	global	2M	device_Symm0994Tier-1_00B9_1_vol
MIGRATE_D_94_00C2_1_mirror1_1	ok	ok	1296000	4K	4.94G	raid-1	local	2M	device_Symm0994_00C2_1_vol
device_Symm0994Tier-1_00B3_1	ok	ok	20640000	4K	78.7G	raid-0	local	-	device_Symm0994Tier-1_00B3_1_vol
device_Symm0994Tier-1_00B4_1	ok	ok	20640000	4K	78.7G	raid-0	local	-	device_Symm0994Tier-1_00B4_1_vol
device_Symm0994Tier-1_00B6_1	ok	ok	1296000	4K	4.94G	raid-0	local	-	device_Symm0994Tier-1_00B6_1_vol
device_Symm0994Tier-1_00B7_1	ok	ok	1296000	4K	4.94G	raid-0	local	-	device_Symm0994Tier-1_00B7_1_vol
device_Symm0994Tier-1_00B8_1	ok	ok	1296000	4K	4.94G	raid-0	global	-	device_Symm0994Tier-1_00B8_1_vol
device_Symm0994Tier-1_00BA_1	ok	ok	1296000	4K	4.94G	raid-0	global	-	Pod5_Symm0994Tier-1_00BA_1_vol
device_Symm0994_00C1_1	ok	ok	2344576	4K	8.94G	raid-c	local	-	device_Symm0994_00C1_1_vol
device_Symm0994_00C3_1	ok	ok	1296000	4K	4.94G	raid-0	local	-	device_Symm0994_00C3_1_vol

/clusters/cluster-2/devices:									
Name	Operational Status	Health State	Block Count	Block Size	Capacity	Geometry	Visibility	Transfer Size	Virtual Volume

device_10	ok	ok	1296000	4K	4.94G	raid-0	local	-	device_10_vol
device_11	ok	ok	1296000	4K	4.94G	raid-0	local	-	-
device_12	ok	ok	2592000	4K	9.89G	raid-0	local	-	-
device_5	ok	ok	1296000	4K	4.94G	raid-0	local	-	-
device_6	ok	ok	1296000	4K	4.94G	raid-0	local	-	device_6_vol
device_7	ok	ok	1296000	4K	4.94G	raid-0	local	-	-
device_8	ok	ok	1296000	4K	4.94G	raid-0	local	-	-
device_9	ok	ok	1296000	4K	4.94G	raid-0	local	-	-



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

26

This command provides additional details on devices, including the operational status and health state, the block count, block size, geometry, and the virtual volume using the device (if any).

Monitor Local Devices

```
VPlexcli:/> local-device summary
SUMMARY (cluster-1)
Device          Summary  (no tier)
-----
Health          devices    10
                unhealthy  0

Visibility      local      10

Capacity        devices w/ space  0
                free capacity  0B
                total capacity 201G

SUMMARY (cluster-2)
Device Summary  (no tier)
-----
Health          devices    14
                unhealthy  0
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

27

The **local-device summary** command displays a summary of local devices. This includes the health state of local devices, the visibility of a device, and the capacity of a device.

Monitor Virtual Volume Details

```
Vplexcli:> ll **/virtual volumes
```

```
/clusters/cluster-1/virtual-volumes:
```

Name	Oper. Status	Health State	Service Status	Block Count	Block Size	Capacity	Locality	Supporting Device
DDVV_Tier-1_00B5_1_vol	ok	ok	running	1296000	4K	4.94G	distributed	DD_Tier-1_00B5_1
DDVV_Tier-1_00BB_1_vol	ok	ok	running	1296000	4K	4.94G	distributed	DD_Tier-1_00BB_1
DDVV_Tier-1_00BD_1_vol	ok	ok	running	1296000	4K	4.94G	distributed	DD_Tier-1_00BD_1
Pod5_remote_Tier-1_vol	ok	ok	running	1296000	4K	4.94G	local	dev_Tier-1_00BA_1
device_Tier-1_00B3_1_vol	ok	ok	unexported	20640000	4K	78.7G	local	dev_Tier-1_00B3_1
device_Tier-1_00B4_1_vol	ok	ok	unexported	20640000	4K	78.7G	local	dev_Tier-1_00B4_1
device_Tier-1_00B6_1_vol	ok	ok	unexported	1296000	4K	4.94G	local	dev_Tier-1_00B6_1
device_Tier-1_00B7_1_vol	ok	ok	running	1296000	4K	4.94G	local	dev_Tier-1_00B7_1
device_Tier-1_00B8_1_vol	ok	ok	running	1296000	4K	4.94G	local	dev_Tier-1_00B8_1
device_Tier-1_00B9_1_vol	ok	ok	running	1296000	4K	4.94G	local	MIGRATE_D_1_00BC_1
device__00C1_1_vol	ok	ok	running	2344576	4K	8.94G	local	device__00C1_1
device__00C2_1_vol	ok	ok	running	1296000	4K	4.94G	local	mirror1_1
device__00C3_1_vol	ok	ok	running	1296000	4K	4.94G	local	device__00C3_1

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

28

This command provides additional details on virtual volumes. These details include virtual volume operational and health state, block count and block size, virtual volume capacity, locality, and virtual volume supporting devices.

Monitor Virtual Volumes

```
VPlexcli:/> virtual-volume summary
Virtual-volume health summary (cluster-1):
    Total 11 virtual-volumes, 0 unhealthy.
Locality summary:
    distributed : 4 virtual-volumes.
    local : 4 virtual-volumes.
    remote : 3 virtual-volumes.
Cache-mode summary:
    synchronous : 11 virtual-volumes.
Total virtual-volume capacity is 341G.

Virtual-volume health summary (cluster-2):
    Total 13 virtual-volumes, 0 unhealthy.
Locality summary:
    distributed : 4 virtual-volumes.
    local : 6 virtual-volumes.
    remote : 3 virtual-volumes.
Cache-mode summary:
    synchronous : 13 virtual-volumes.
Total virtual-volume capacity is 359G.
```

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring



29

The **virtual-volume summary** command displays information about the virtual volumes in a cluster. This information includes the virtual volume health state, locality, cache mode, capacity, and total virtual volume capacity.

Monitor Storage Views

- Identifies volumes exported more than once
- Identifies views which only contain a single port
- Identifies views which are disabled

```
VPlexcli:/> export storage-view checkconfig
Checking cluster cluster-1:
No errors found for cluster cluster-1.
Checking cluster cluster-2:
No errors found for cluster cluster-2.
Volume DDVV_Symm0994Tier-1_00B5_1_vol is exported multiple times:
  view: Win_SV, lun: 3
  view: Pod6_Win, lun: 1
```

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

30

The **export storage-view checkconfig** command will perform a quick configuration verification of VPlex storage views.

Storage View Details

```
Vplexcli:/> ll **/storage-views
/clusters/cluster-1/exports/storage-views:
Name      Operational  initiator-ports  virtual-volumes  port name, enabled, export status
-----
Win_SV     ok           Win-HBA0,        (0,device_Symm...,8.94G),  P000000003CA016DB-A0-FC00,true,ok,
Win-HBA1   (1,device_Symm...,4.94G),  P000000003CA016DB-A1-FC00,true,ok,
           (2,device_Symm...,4.94G),  P000000003CB016DB-B0-FC00,true,ok,
           (3,DDVV_Symm...,4.94G),    P000000003CB016DB-B1-FC00,true,ok
           (4,device_Symm...,4.94G),
           (5,DDVV_Symm...,4.94G),
           (6,device_Symm...,4.94G)

/clusters/cluster-2/exports/storage-views:
Name      Operational  initiator-ports  virtual-volumes  port name, enabled, export status
-----
Pod6_Win  ok           Win6-HBA0,      (0,DDVV_Symm...,4.94G),    P000000003CA01671-A0-FC00,true,ok,
Win6-HBA1 (1,DDVV_Symm...,4.94G),    P000000003CA01671-A1-FC00,true,ok,
           (2,Pod5_remote_Symm...,4.94G), P000000003CB01671-B0-FC00,true,ok,
           (3,device_l0_vol,VPD83T3,4.94G), P000000003CB01671-B1-FC00,true,ok
           (4,device_6_vol,VPD83T3,4.94G),
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

31

This command displays more information on storage views. This includes the operational status of a view, initiator ports, virtual volumes, and VPLEX front-end ports that are part of the view. It also displays a virtual volume's external LUN ID.

Monitor Storage Views

```
VPlexcli:/> export storage-view summary
```

```
View health summary(cluster-1):
```

view name	health-state	exported volumes	ports	registered initiators
POD5_ESX	healthy	0	4	2
Win_SV	healthy	7	4	2

```
Total 2 views, 0 unhealthy.
```

```
View health summary(cluster-2):
```

view name	health-state	exported volumes	ports	registered initiators
Pod6_Win	healthy	12	4	2

```
Total 1 views, 0 unhealthy.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

32

The **export storage-view summary** command lists each view and the number of volumes and initiators that each contains.

Monitor Storage Views (Cont.'d)

```
VPlexcli:/> export storage-view find-unmapped-volumes cluster-2
```

Name	Operational	Health	Service	Block	Block	Capacity
-----	Status	State	Status	Count	Size	-----
DDVV_Symm0994Tier-1_00BD_1_vol	ok	ok	running	1296000	4K	4.94G
device_Symm0994Tier-1_00B8_1_vol	ok	ok	running	1296000	4K	4.94G

```
VPlexcli:/> export storage-view map -v Win_SV
```

```
VPD83T3:6000144000000010a016dbc72f645905 device_Symm0994_00C1_1_vol
VPD83T3:6000144000000010a016dbc72f645900 device_Symm0994_00C2_1_vol
VPD83T3:6000144000000010a016dbc72f6458fb device_Symm0994_00C3_1_vol
VPD83T3:6000144000000010a016dbc72f64595a DDVV_Symm0994Tier-1_00B5_1_vol
VPD83T3:6000144000000010a016dbc72f645948 device_Symm0994Tier-1_00B8_1_vol
VPD83T3:6000144000000010a016dbc72f64598b DDVV_Symm0994Tier-1_00BD_1_vol
VPD83T3:6000144000000010a016dbc72f64594e device_Symm0994Tier-1_00B7_1_vol
```

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

33

The **export storage-view *find-unmapped-volumes* <cluster>** command will display the virtual volumes that are un-exported (not placed in a Storage View).

The **export storage-view map <view>** will display virtual volumes that are part of the storage view and their corresponding local device IDs.

Monitor Distributed Devices

```
VPlexcli:/> ds summary
Device health summary:
    Total 1 devices, 0 unhealthy.
Cluster summary:
    Cluster cluster-2 : 1 distributed devices.
    Cluster cluster-1 : 1 distributed devices.
Capacity summary:
    1 devices have some free capacity.
    60G free capacity of 60G total capacity.

VPlexcli:/> ll **/distributed-devices
/distributed-storage/distributed-devices:
Name                               Status  Operational  Health  Auto  Rule  Transfer
-----
-----
-----
DD_Symm0994Tier-1_00B5_1  running ok          ok      false -      2M
DD_Symm0994Tier-1_00BB_1  running ok          ok      false -      2M
DD_Symm0994Tier-1_00BD_1  running ok          ok      false -      2M
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

34

These commands provide information about distributed devices within a single cluster or two-site installation metro.

Lesson Summary

During this lesson the following topics were covered:

- VPLEX commands used for monitoring

In this lesson we saw a number of commands that are used to monitor the health and status of a VPLEX environment.


Lesson: Configure Monitoring

During this lesson the following topics are covered:

- Create VPLEX monitors in a VPLEX environment
- Use existing monitors in a VPLEX environment

This lesson covers VPLEX monitor functions.

End-to-End Monitoring



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring


37

The performance monitoring dashboard provides a customized view into the performance of your VPLEX system. You decide which aspects of the system's performance to view and compare.

Performance information for the current 5-minute window is displayed as a set of charts, including:

- WAN Link Performance chart
- WAN Link Usage chart
- Back-end Latency chart
- Front-end Latency chart
- Rebuild Status chart
- CPU Utilization chart
- Front-end Aborts chart
- Subpage Writes chart
- Front-End Throughput chart
- Front-End Bandwidth

View the Monitored Statistics



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

38

Here we see the Disk Read Performance report. This shows us various volumes read latency distribution in microseconds on the specified storage volume.

Statistic parameters are listed in the VPLEX CLI Guide.

VPlexcli Pre-Configured Monitors

- Automatic scripted command
 - `report create-monitors`
- Creates three monitors per director
 - `Cluster_n_Dir_nn_diskReportMonitor`
 - `Cluster_n_Dir_nn_portReportMonitor`
 - `Cluster_n_Dir_nn_volumeReportMonitor`
- Output files written to: `/var/log/VPlex/cli/reports/`
- Verify creation of monitors
 - `ll //monitoring/directors/*/monitors`
- In a large system (>1000 volumes) it can take some time to setup the monitors
 - No impact to host performance

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

EMC²

39

The creation of the base monitors has been simplified with the report **create-monitors script** command. This creates three monitors that target different aspects of the VPlex configuration, one for the Disks, one for the Ports and one for the Volumes.

VPlexcli Monitoring Custom Script

Using the `monitors` command script

- Command script used to ease monitor creation. Used as follows:

```
- script -i monitors
- monitors --help
- monitors start -d <duration> -f 10 -s
  "cache.*,directory.*,ramf.*,fe-
  director.*,director.*,storage-volume.write-
latency,storage-volume.read-latency"
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

40

This command creates monitors and file sinks for all directors in the system, for specified duration and polling frequency. It is more suited for focused performance testing rather than general data collection.

CAUTION: Do NOT leave unattended for long periods of time, as files can grow unbounded if left running indefinitely.

Monitor Configuration Steps

1. Determine the type of statistic to collect
`monitor stat-list`
2. Determine how often the monitor should collect statistics
3. Create a monitor
`monitor create`
4. Add one or more sinks to the monitor
`monitor add-sink`
5. Update/Collect statistics
6. Monitor output location:
`/var/log/VPlex/cli/`



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

41

To create and operate a monitor, use the following general steps:

1. Determine the type of statistic to collect from the target object.
 - Use the **monitor stat-list category** and/or the **monitor stat-list *** command to display the statistics to include in the monitor.
 - Note whether the statistic you want to collect requires an argument (port number, volume ID).
2. Determine how often the monitor should collect statistics.
3. Use the **monitor create** command to create a monitor.
4. Use the **monitor add-sink** command to add one or more sinks to the monitor. The monitor begins operation (polling and collecting performance data) when the sink is added to the monitor. To disable automatic polling without deleting the monitor or its sink(s), you can either use the **set** command to change the monitor's period attribute to 0, or change the sink's enabled attribute to false.
5. Use the **monitor collect** command to update/collect statistics immediately without waiting for the monitor's automatic collection.
6. Monitor output. Console sinks will display monitor output on the console. For file sinks, navigate to `/var/log/VPlex/cli/` on the management server and use the `tail -f filename` to display the output or send output to a csv file, open the file in Microsoft Excel and create a chart.

Create a Monitor

```
Vplexcli:/> monitor stat-list
be-prt, cache, cg, com-cluster-io, director, directory, fc-com-port, fe-director, fe-lu, fe-prt,
ip-com-port, ramf, rp-spl-node, rp-spl-vol, storage-volume, virtual-volume, wrt-pacing

Vplexcli:/> monitor create --name TestMonitor --director /engines/engine-1-1/directors/director-1-
1-A --stats fe-prt.* --targets /engines/engine-1-1/directors/director-1-1-A/hardware/ports/A0-
FC0 1
Successfully created 1 monitor(s) out of 1.

Vplexcli:/monitoring/directors/director-1-1-A/monitors> monitor add-file-sink --monitor TestMonitor --
file /var/log/Vplex/cli/TestMonitorOutput.csv

Vplexcli:/monitoring/directors/director-1-1-A/monitors/TestMonitor/sinks> ll file/
/monitoring/directors/director-1-1-A/monitors/TestMonitor/sinks/file:
Name      Value
-----
enabled   true
format    csv
sink-to   /var/log/Vplex/cli/TestMonitorOutput.csv
type      file
```

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

42

Performance monitoring collects and displays statistics to determine how a port or volume is being used, how much I/O is being processed, CPU usage, and so on. The VPLEX collects and displays performance statistics using two user-defined objects, Monitors and Monitor Sinks. Before Creating our monitor, first use the monitor stat-list command to display the available statistics. **Monitors** gather the specified statistics, while **Monitor Sinks** direct the output to the desired destination. Monitor sinks include the console, a file, or a combination of the two. The monitor defines the automatic polling period, the statistics to be collected, and the output of the format. The monitor sinks define the output destination.

Create a Default Monitor Set

```
VPlexcli:/> report create-monitors
Creating monitor diskReportMonitor on Director director-1-1-B monitoring 12 targets, file
/var/log/VPlex/cli/reports/diskReportMonitor_director-1-1-B.csv.
Successfully created 1 monitor(s) out of 1.

time: 0 sec
Creating monitor volumeReportMonitor on Director director-1-1-B monitoring 3 targets, file
/var/log/VPlex/cli/reports/volumeReportMonitor_director-1-1-B.csv.
Successfully created 1 monitor(s) out of 1.

time: 0 sec
Creating monitor portReportMonitor on Director director-1-1-B monitoring 16 targets, file
/var/log/VPlex/cli/reports/portReportMonitor_director-1-1-B.csv.
Successfully created 1 monitor(s) out of 1.

time: 0 sec
Creating monitor diskReportMonitor on Director director-1-1-A monitoring 12 targets, file
/var/log/VPlex/cli/reports/diskReportMonitor_director-1-1-A.csv.
Successfully created 1 monitor(s) out of 1.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

43

Creates three monitors for each director in the system. Monitors are named:

Cluster_n_Dir_nn_diskReportMonitor

Cluster_n_Dir_nn_portReportMonitor

Cluster_n_Dir_nn_volumeReportMonitor

The period attribute for the new monitors is set to 0 (automatic polling is disabled). Use the **report poll-monitors** command to force a poll.

Each monitor has one file sink. The file sinks are enabled. By default, output files are located in /var/log/VPlex/cli/reports/ on the management server. Output filenames are <Monitor name>_<Cluster_n_Dir_nn.csv

Disk report monitors collect:

storage-volume.per-storage-volume-read-latency

storage-volume.per-storage-volume-write-latency

Port report monitors collect:

be-prt.read

be-prt.write

fe-prt.ops

fe-prt.read

fe-prt.write

Volume report monitors collect:

virtual-volume.ops

virtual-volume.read

virtual-volume.write

narasimharao.nelluri@daimler.com

Manually Run Monitor Polling

```
VPlexcli:/> report poll-monitors
```

```
Collecting data for director director-1-1-A monitor director-1-1-A_diskReportMonitor.  
Collecting data for director director-1-1-A monitor director-1-1-A_portReportMonitor.  
Collecting data for director director-1-1-A monitor director-1-1-A_volumeReportMonitor.  
Collecting data for director director-1-1-B monitor director-1-1-B_diskReportMonitor.  
Collecting data for director director-1-1-B monitor director-1-1-B_portReportMonitor.  
Collecting data for director director-1-1-B monitor director-1-1-B_volumeReportMonitor.  
Collecting data for director director-2-1-A monitor director-2-1-A_diskReportMonitor.  
Collecting data for director director-2-1-A monitor director-2-1-A_portReportMonitor.  
Collecting data for director director-2-1-A monitor director-2-1-A_volumeReportMonitor.
```



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

44

The monitors created by the report **create-monitors** command have their period attribute set to 0 seconds (automatic polling is disabled) and one file sink.

Use this command to force an immediate poll and collection of performance data for monitors created by the report create-monitors command.

Output is written to files located in **/var/log/VPlex/cli/reports/** on the management server.

Aggregate the Reports

```
VPlexcli:/> report aggregate-monitors
Aggregating port reports at cluster cluster-2 ...
Aggregating port reports at cluster cluster-1 ...
Aggregating volume reports at cluster cluster-2 ...
Aggregating volume reports at cluster cluster-1 ...
Aggregating disk reports at cluster cluster-2 ...
Aggregating disk reports at cluster cluster-1 ...
```



Module: VPlex Monitoring

45

Here we aggregate the reports generated by the report create-monitors or monitor commands. The reports are aggregated by cluster. An aggregate report is generated for:

- Each cluster
- Each type of per director report in the specified directory

This command assumes that the per director reports have filenames in the format: <report type>ReportMonitor_<director>.csv

All other files in the directory will be ignored. Aggregate report filenames are in the format <report type>Performance_<cluster>.csv

If an aggregate report already exists, it is overwritten.


Output files saved in /var/log/VPlex/cli/reports/

- diskPerformance_cluster-1.csv
- diskPerformance_cluster-2.csv
- portPerformance_cluster-1.csv
- portPerformance_cluster-2.csv
- volumePerformance_cluster-1.csv
- volumePerformance_cluster-2.csv

Add the Monitor Report to the Scheduled Tasks

```
Vplexcli:/> schedule add -t "0 1 * * *" -c report  
poll-monitors
```

```
Vplexcli:/>
```



EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

46

Schedule a job to run at the specified time(s). Example shown will run the report poll-monitors command at 1AM every day.

Required arguments:

- **-t --time <time>** - * Date and time the job executes in crontab-style format enclosed in quote marks. Values are specified in the crontab-style format: <Minute> <Hour> <Day of the Month> <Month> <Day of the week>
 - 1. Minute - 0-59
 - 2. Hour - 0-23
 - 3. Day of the Month - 1-31
 - 4. Month - 1-12, January=1...December=12
 - 5. Day of the week - 0-6, Sunday=0...Saturday=6
- **-c --command <command>** - * The CLI command to be executed at the specified time.

* - Argument is positional.

Perpetual Performance Monitors

- Always on, started from system setup, cannot be modified, disabled, or deleted
- Records a default set of performance categories
- One file per director, local to a cluster
- Files written to `/var/log/VPlex/cli/` and files end in `...PERPETUAL_vplex_sys_perf_mon.log`
- Files capped at 10MB, and rotated up to 10 times
- Files can be copied for offline analysis



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

47

The perpetual performance monitor files are rolled over in a way that is not very user-friendly for this script. The currently open monitor file is capped at 10MB per director, and up to 10 files are stored (.log, .log.1, .log.2, and so on.) Because different I/O activity on the director can result in a different amount of generated performance statistics, each director's file can get out of time sync.

Chances are you might see the most recent perpetual monitor file for one director spans a short time and is quite small, yet another director's currently open file may span several days and approach 10MB.

Perpetual Performance Monitors

```
service@VPLEX-Cluster1: /var/log/VPlex/cli> ls

audit.log_20150205171539
audit.log_20150205174654
audit.log_20150318132657
authentication-service-configuration.log
bole.log_20150207145624
bole.log_20150318132656
CACertSubjectInfo.txt
CallHomeInfo.properties
CallHomeTemplate.xml
capture
capture.log_20150318132656
client.log_20150318132657.9
cluster-witness-config.xml
connections
crontab
cst_conf_bkup_2.0.1.tar
default.cacerts
director-1-1-A_PERPETUAL_vplex_sys_perf_mon_28.zip
director-1-1-A_PERPETUAL_vplex_sys_perf_mon.log
director-1-1-A_PERPETUAL_vplex_sys_perf_mon.log.1
director-1-1-A_PERPETUAL_vplex_sys_perf_mon.log.10
director-1-1-A_PERPETUAL_vplex_sys_perf_mon.log.2
...

HostCertSubjectInfo.txt
installer.log
instance.properties.yy
ipconfig.xml
jaas.config
jaas.config.yy
ListenerEventQParameters.xml
monitor.log_20150207145624
monitor.log_20150207145624.1
monitor.log_20150207145624.2
monitor.log_20150207145624.3
session.log_service_10.12.246.111_W14300554_200210
session.log_service_10.12.246.111_W14352_201503624
session.log_service_10.126.96.23_W1430443_20120329
session.log_service_10.126.96.23_W1430424_20150501
session.log_service_10.126.96.23_W1430502_20150501
session.log_service_10.126.96.23_W1430665_21001038
session.log_service_10.126.96.23_W1430737_20150500
session.log_service_10.126.96.23_W14362883_2015736
session.log_service_10.126.96.23_W1431433_20150154
session.log_service_10.250.42.197_W1430450_2015310
session.log_service_10.251.33.69_W14306906_2015327
```



Here we see the contents of the folder where the perpetual monitors are written. Notice the naming convention of the monitor files.

Lesson Summary

During this lesson the following topics were covered:

- Monitoring in the GUI
- Pre-configured monitors
- How to configure a monitor
- Perpetual Performance monitors

In this lesson we covered the monitors that can be configured or are already pre-configured in a VPLEX.

Lesson: Analyze Monitor Data

During this lesson the following topics are covered:

- Analyze monitor output
- Use tools to review and monitor performance metrics

This lesson covers monitor output analysis.

VPlexcli Monitor Categories

- 15 high-level categories each with subcategories

- List all categories

- `VPlexcli:/> monitor stat-list`

Output:

- be-prt, cache, cg, director, directory, fc-com-port,
 - fe-director, fe-lu, fe-prt, ip-com-port, ramf, rdma,
 - storage-volume, virtual-volume, wrt-pacing

- To see all subcategories:

`VPlexcli:/> monitor stat-list *`

- Monitoring has no impact on host performance

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPlex Monitoring

51

Refer to the *VPLEX CLI Guide, Performance Monitoring* section for a complete listing and explanation of each category.

Monitor stat-list

VPlexcli:/> **monitor stat-list ***

Name	Target	Type	Units
-----	-----	-----	-----
be-prt.read	backend-port	counter	KB/s
be-prt.write	backend-port	counter	KB/s
cache.dirty	n/a	reading	KB
cache.miss	n/a	counter	counts/s
cache.rhit	n/a	counter	counts/s
cache.subpg	n/a	counter	counts/s
cg.closure	consistency-group	bucket	us
cg.delta-util	consistency-group	reading	%
cg.drain-lat	consistency-group	bucket	us
cg.exch-bytes	consistency-group	counter	KB/s
cg.exch-lat	consistency-group	bucket	us
cg.exch-pages	consistency-group	counter	counts/s
cg.input-bytes	consistency-group	counter	KB/s
cg.input-ops	consistency-group	counter	counts/s
cg.inter-closure	consistency-group	bucket	us
cg.outOfDate-counter	consistency-group	counter	counts/s
cg.pipe-util	consistency-group	reading	%

EMC¹

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

52

Statistics are grouped into sub-categories.

Use the **monitor stat-list** command followed by the **<Tab>** key to display the statistics sub-categories. Some stat categories need to be pointed to a particular target. The Type column is discussed on upcoming slides.

Types of Statistics

- **counters** - Monotonically increasing value (analogous to a car's odometer). Counters are used to count bytes, operations, and errors. Often reported as a rate such as counts/second or KB/second.
- **readings** - Instantaneous value (analogous to a car's speedometer). Readings are used to display CPU utilization, memory utilization. Value can change every sample.
- **buckets** - Histogram of 'bucketized' counts. Buckets are used to track latencies, determine median, mode, percentiles, minimums and maximums.

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

53

VPLEX collects and reports three types of statistics:

1. Counters - Monotonically increasing value (analogous to a car's odometer) Counters are used to count bytes, operations, and errors. Often reported as a rate such as counts/second or KB/second.
2. Readings - Instantaneous value (analogous to a car's speedometer) Readings are used to display CPU utilization, memory utilization. Value can change every sample.
3. Period-average - Average of a series calculated over the last sample period. If: $\text{current_reading_sum}$ is the sum of all readings for the particular statistic since the monitor's creation, and $\text{previous_reading_sum}$ is the count of all readings for the statistic since the monitor's creation $\text{period-average} = (\text{current_reading_sum} - \text{previous_reading_sum}) / (\text{current_reading_count} - \text{previous_reading_count})$
4. Buckets - Histogram of bucketized counts. Buckets are used to track latencies, determine median, mode, percentiles, minimums and maximums. Period-averages apply only to latency stats. They are calculated kind of like counters. Reading of two counters at $t=0$ (total latency and total ops), and then another reading at $t=30$. Then some math is done to derive the average latency of all the operations that completed in the time period.

Bucket Statistics

Port FCA13 Bucket Statistics Read Latency in msec

# of issues	28	184	60	3	6
Latency	0-49	50-99	100-149	150-199	Overflow
	10%	66%	21%	1%	2%



© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

54

When using bucket statistics, you track the number of instances or times that a particular event takes place. For example, here we see a read latency for a single port.

Monitor Targets

```
Vplexcli:/> monitor stat-list fe-prt
```

Name	Target	Type	Units
fe-prt.ops	frontend-port	counter	counts/s
fe-prt.read	frontend-port	counter	KB/s
fe-prt.read-lat	frontend-port	bucket	us
fe-prt.write	frontend-port	counter	KB/s
fe-prt.write-lat	frontend-port	bucket	us

```
Vplexcli:/> monitor stat-list cache
```

Name	Target	Type	Units
cache.dirty	n/a	reading	KB
cache.miss	n/a	counter	counts/s
cache.rhit	n/a	counter	counts/s
cache.subpg	n/a	counter	counts/s

EMC²

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring

55

Many statistics require a target port or volume to be specified. Output of the **monitor stat-list** command identifies which statistics need a target defined, and the type of target required when a monitor is created. Some monitor entries require you to specify a target when collecting their data. Use the **monitor stat-list** to determine if targets will be required.

Lab: System Monitoring



This lab covers how to monitor the VPLEX environment.

- System Monitoring
- Configure Monitors

© Copyright 2015 EMC Corporation. All rights reserved.

Module: VPLEX Monitoring



56

This lab covers how to monitor the VPLEX environment.

Lesson Summary

During this lesson the following topics were covered:

- VPLEX monitor output
- Types of statistics generated



In this lesson we demonstrated how to analyze the data that comes out of the VPLEX monitors and how to use it to target your troubleshooting efforts.

Module: Summary

Key points covered in this module:

- Show how VPLEX sends events and system reports to an ESRS Gateway.
- How to perform basic maintenance on various VPLEX aspects.
- How to configure and perform routine monitoring of a VPLEX instance.

This module covered system and performance monitoring. There is also a VPLEX Troubleshooting course and a VPLEX Performance course, each with more detailed instruction on how to monitor, troubleshoot, and increase the performance of the VPLEX.

Course Summary

Key points covered in this course:

- Common VPLEX terms, configuration options, hardware and software architecture
- Installation of a VPLEX Metro configuration
- How data flows within a VPLEX system
- Managing VPLEX using the Unisphere and the CLI
- Provisioning virtual volumes to hosts
- Encapsulating an existing SAN storage volume into VPLEX
- Creating Distributed storage
- Performing VPLEX mobility
- Performing routine VPLEX monitoring tasks



This course covered common VPLEX terms, configuration options, hardware and software architecture, Installation of a VPLEX Metro configuration, How data flows within a VPLEX system, Managing VPLEX using the Unisphere and the CLI, Provisioning virtual volumes to hosts, Encapsulating an existing SAN storage volume into VPLEX, Creating Distributed storage, Performing VPLEX mobility and Performing routine VPLEX monitoring tasks.

The logo for EMC², featuring the letters "EMC" in a large, white, serif font, followed by a superscript "2" and a registered trademark symbol (®). The logo is centered on a solid blue rectangular background.

EMC²®