

VBLOCK SIZING

Poorna Chandra Prashanth C
Senior Systems Engineer Analyst
Dell EMC
Poorna.prashanth@Dell.com

Table of Contents

Intended Audience.....	3
Introduction	3
Objective	4
VBlock System Building Blocks.....	4
VBlock System Compute	6
Compute Sizing	7
Data Gathering for Compute	7
Data Consolidation and Analysis for Compute	7
CPU and RAM Requirement calculations.....	8
VBlock System Network.....	9
Storage Sizing for VBlock	11
Data Gathering.....	12
Data Consolidation and Analysis.....	13
Capacity Sizing.....	14
VNX/VMAX sizing	14
XtremIO Sizing.....	15
VSphere Design considerations	16
Compute Virtualization	17
Network Virtualization.....	17
Storage Virtualization	18
References	19

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell EMC's views, processes or methodologies.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.



Intended Audience

This document is targeted toward those who want a preliminary understanding of hardware sizing for VBlock.

Introduction

Data is growing rapidly and companies spend an enormous amount of money on Data Center Infrastructure to support it. There are specific components to process, transfer, and store data in the data center. These components are at different layers which need to be provisioned, monitored, and managed separately. In traditional data centers, a set of resources are devoted to specific computing technology or applications and these support a specific set of workloads and cannot be reconfigured to support varying usage loads.

Converged Infrastructure (CI) is a method to centralize IT resources to consolidate existing systems for increased resource utilization, ease of management, and lower costs significantly. Converged Infrastructure achieves these by pooling compute, storage, and networking resources into single package that can be shared by various applications and can be provisioned, monitored, and managed in a collective manner.

VBlock®, one of today's most advanced Converged Infrastructure systems, integrates best-in-class compute, network, and storage technologies from industry leaders Cisco, EMC, and VMware. It simplifies all aspects of IT enabling organizations to achieve better business outcome faster.

These Integrated stacks are mainly aimed at organizations that want preconfigured and pretested systems for highly virtualized environments.

Objective

Sizing compute, storage, and network resources for a VBlock can be a complex task as VBlock is integrated with multiple vendors at different layers. There are no easy answers. Right Sizing provides more flexibility and the ability to deliver better customer service to the hosted multiple applications and to meet defined service level agreements (SLA).

This paper provides a high-level overview of the basic steps in sizing the hardware for the VBlock infrastructure with best practices for all three layers in VBlock; Compute, Network, and Storage.

Using internal tools or manual calculations of our findings to size primary components of VBlock at all three layers by considering best practices and use cases will help us determine the right components for VBlock.


Sizing results would help Presales and Sales position VBlock proposals better against competing vendors. The results can also play as a use case which captures a customer's environment with specific requests and growth and serve as a road map for a customer's future investments

VBlock System Building Blocks

VBlock system has prebuilt components that are racked, stacked, and cabled together to create a single shippable appliance.

Integrated Building Block Platform

- Cisco UCS Blade and Rack Servers: Compute
- Cisco Nexus and MDS switches: Network
- EMC VNX® and VMAX® arrays: Storage
- VMware vSphere hypervisor: Server Virtualization




For each of these layers, physical and logical components need to be sized to meet customer requirement (Workload, Bandwidth, and Storage) by considering best practices and overhead while provisioning compute, network, and storage resources.

In this document, compute, network and storage is sized separately to identify the right hardware. The consolidated hardware across all 3 layers should determine the VBlock system for immediate and longer term requirements.

We will strive to understand elements and variables at the customer site that decide logical and physical hardware requirements in all three layers of VBlock.

- Compute
 - Determining physical CPU and Memory requirements
 - Determining vCPUs that can be assigned for each CPU or Core
 - Determining VM's required to meet customer requirement
- Network
 - Bandwidth analysis to determine required bandwidth estimates for VBlock environment
- Storage
 - Analysis of current environment to understand:
 - Current running applications
 - IO Bursts and Peak workload
 - Capacity requirement
 - Conclude on storage requirement estimates
 - Decide type of storage in VBlock to meet customer requirement



VBlock System Compute


VBlock systems utilize the Cisco Unified Computing System (UCS) in the compute technology. Its components provide all that is necessary to load and run applications and connect to other parts of the VBlock Systems. Compute components consist of Compute Servers, Server chassis, Fabric Extenders, and Fabric Interconnects.

There are two types of Cisco UCS servers used in the VBlock system models; Blade Servers and Rack Mounted Servers.

Rack Mounted servers are used in the low-end VBlock 200 while Blade Servers are used in high-end VBlock models. VBlock chassis has eight blade slots which can consist of 8 half-width blades, 4 full-width blades, or a combination of both.

Fabric Extenders provide interconnect between blade server and Fabric Interconnect. Each chassis include two Fabric Extenders. Fabric Interconnects come in pairs for redundancy connecting VBlock compute to other parts of the VBlock system.

VBlock system compute consists of both CPU and Memory, depending on the workload and application it runs. We can determine the number of CPU, CPU speed, and Memory while compute sizing and these values can be used to determine the type of server to be used in the VBlock system.



Compute Sizing

Compute sizing in VBlock systems plays a vital role in determining the type and number of servers to be used. It may be rack-mounted servers for VBlock 200 or Blade Servers for higher-end models. The fundamental step in determining the key elements for VBlock compute is to understand the customer environment.

The existing environment can be physical, virtual, or mixed. In this document we focus mainly on Physical to Virtual (P to V) conversion. In virtual-to-virtual conversions we can use VMware Capacity Planner to collect and analyze the data and the result can be used to calculate required compute elements for the VBlock system. This information can be used to decide server requirement for VBlock.

Data Gathering for Compute

The first step in the process is to gather baseline information on the key user groups that have been identified as good candidates for a VBlock environment. The purpose of this step is to understand the performance characteristics of the target users' workload—for instance: What applications do they need? Are the applications more CPU- or memory-intensive? Are there an excessive number of storage operations? What processing ability is required to handle existing workload that is being generated by the end users' activities?

The first step is to determine applications and number of servers targeted to move to the VBlock.

Collect information and background data about existing servers identified to move to the VBlock. It is important to ask the right questions. Context questions can provide valuable insight. Understand how problems in the current environment create gaps in service.

Data can be collected through various methods,

- a) Customer Interaction
- b) Surveys and Questionnaires
- c) Site Visits and Tool-based discovery

Data collection starts with an inventory of the physical systems and capturing data during peak processing period. The more granular the data, the better.

Tools that we can use to collect server data can be:

- a) VMware Capacity Planner
- b) Perfmon for Windows
- c) IOSTAT or MPSTAT or VMSTAT for Linux/Unix

Data Consolidation and Analysis for Compute

The next step is to consolidate and analyze collected data. This would give us the total number of CPU, CPU speed, and memory required for target VBlock Compute. Mitrend is one of the tools which would consolidate and analyze the collected data and provide consolidated results.

Mitrend provides instruction for capturing data and takes captured data as an input. It provides instructions to collect data across applications, servers, networks, and storage boxes, including instructions about collecting data from EMC and non-EMC products; i.e. Windows, UNIX, and ESXi Servers.

We need to record key metrics from the consolidated report. Metrics required in our exercise are:

- Number of CPU in each Server
- CPU peak Utilization for each Server
- Number of Sockets per Server
- Processor type and speed for each server
- Memory configured in each server
- Memory peak utilization

CPU and RAM Requirement calculations

It is important to configure the system with enough resources so that the system meets agreed expectations. All VBlock System components have performance limits. As a rule of thumb, always size with a limit of no more than 70% utilization per component. Consider design for availability so system does not degrade performance in the event of component failure.



- Record the number of CPU for every physical server and calculate the average CPU
- Record CPU peak Utilization for every physical server. Calculate the average CPU peak utilization
- Record memory configured for each server, the average memory configured
- Record memory peak utilization for each server, and the average memory peak utilization

CPU

Average CPU per physical (MHz) x Average CPU chore Count = **Average CPU per physical system**
 $1,000\text{MHz} \times 4 = \mathbf{4,000\text{MHz}}$

Average CPU per physical system x Average peak CPU utilization (percentage) = **Average peak CPU utilization (MHz)**
 $4,000\text{MHz} \times 12\% = \mathbf{480\text{ MHz}}$

Average peak CPU utilization (MHz) x Number of concurrent VM's = **Total peak CPU utilization (MHz)**
 $480\text{MHz} \times 50 = \mathbf{24,000\text{MHz}}$
(Number of virtual machines will be an input from customer)

Consider 70% as threshold

Total CPU required = Average peak CPU utilization (MHz) + (Average peak CPU utilization (MHz) 0.07)

Memory

Average RAM per physical (MB) x Average Peak RAM utilization (percentage) = **Average peak RAM utilization (MB)**

2,000MB x 52% = **1040MB**

Average peak RAM utilization (MB) x Number of concurrent VM's = **Total peak RAM utilization (MB)**

1040MB x 50 = **52,000MB**

Consider 70% as threshold

Total RAM required = Total peak RAM utilization (MB) + (Total peak RAM utilization (MB)*0.07) +

Paging Size

Now that we have calculated **Total CPU requirement** and **Total memory requirement**, we would match the table below to determine the type and number of blades required for VBlock.

Blade Type	Maximum Memory	Maximum # Cores	Processor Speed (GHZ)	Relative Processing Capability*
B200 M3	768 GB	16	2.9	32.48
B420 M3	1536 GB	32	2.7	60.48
B200 M4	768 GB	36	2.3	57.96
C220 M4	768 GB	28	2.6	50.96
C240 M4	768 GB	28	2.6	50.96
B260 M4	1536 GB	30	2.8	58.8
B460 M4	3072 GB	60	2.8	117.6

Always add additional headroom as growth/memory reservations during these calculations. Recognize there may be exceptions/specific requirements which may require additional resources.

Understanding the type of Operating Systems and applications running in the existing environment would help us determine Core to vCPU ratio. The table below would provide insight on those lines.

Standard	Performance	Optimized Performance
1:8 CPU to vCPU ratio	1:4 CPU to vCPU ratio	1:1 CPU to vCPU ratio
0% Memory reserved	50% Memory reserved	100% Memory reserved
CPU and network burst capability	Design for production workloads	Entire server dedicated to specific workload
Lowest price point	Optimization for general workloads	Optimization for large-scale workloads
Value	Balance	Maximum performance

VBlock System Network

The network is a critical element of a VBlock system, providing connectivity within and outside of VBlock. Inside, it provides communication between compute and storage devices and outside it provides communication between VBlock and the site network. VBlock has both LAN and SAN; LAN carries general network and Internet traffic and SAN is dedicated to carry storage-related network traffic.

SANs use Fibre Channel network standards and on physical network layer they use Fibre Channel Protocol (FCP). LANs are Ethernet networks which on physical network layer uses IP.



Vblock System servers have no storage and their OS files are stored on storage devices. Instead of separate IP and SAN switches, Vblock systems can be built with Cisco Unified Switch. A unified switch can perform tasks associated with both IP and Fibre Channel networks. However, a separate LAN and SAN switch would have more ports than a Unified switch. Hence, we have two options to consider for Network in VBlock Systems:

- Unified Switch or
- IP and SAN switch

VCE Vblock™ Systems	Vblock Series 100 BX/DX	Vblock Series 200	Vblock Series 300 EX/FX/GX/HX	Vblock Series 700 MX/LX
Network Environment (Cisco)	Cisco Unified Network (Nexus / MDS)			
Cisco Nexus Aggregation / Uplink	2x Catalyst 3750-24/48	2xNexus 5548UP	2x Nexus 5548UP/5596UP with each 48/96 Ports (Ethernet, FC, FCoE)	
Cisco Nexus 1000V Software Switch	1000v Essentials	1000v Adv.Edition	Nexus 1000v Software Switch (licensed per CPU)	
Cisco MDS FC-SAN	--		2x MDS-9148 with each 48 FC-Ports (optional 2x MDS Director)	
Cisco Network Configuration	ISCSI,NFS Traffic, StackWise pairs	1xCatalyst 3750 for internal Mgmt.	Cisco UCS FI split off internal UCS FCoE traffic to FC/Ethernet	

Nexus 5548 will have up to 48 ports, which is 10 Gigabit Ethernet and FCoE switch offering up to 960 Gbps throughput. Nexus 5596 will have up to 96 ports and can offer up to 1920 Gbps throughput. MDS -9148 is a Multilayer fabric switch which as an aggregated platform, can provide up to 760 Gbps of bandwidth.

Nexus 1000v is a virtual switch




Network design for VBlock would follow the same technical considerations as other infrastructure components. Network connects to all devices in the network and hence, it should have enough ports to support the infrastructure. Network capacity is provisioned along with compute capacity, thus, addition of compute or storage resources will not impact network negatively. However, intelligent distribution of workloads across network is vital to guarantee performance and support SLAs.

Storage Sizing for VBlock

Storage is vital element and plays significant role in VBlock system since servers in VBlock do not have storage. All information including digital files and block data and Operating System-related files are serviced from Storage and hence, servers boot from block level storage from over SAN using Fibre Channel Protocol or iSCSI Protocol.

Both Unified and Block Level storage are used in VBlock systems. Storage products used in VBlock systems are VNXe, VNX EMC Symmetrix® and XtremIO®.



Storage should be sized to handle workload and capacity of the applications for the VBlock Systems. Workload and Capacity requirement decides the optimal storage system for a VBlock. Storage system decides the VBlock system model. Below table can be used to understand type of hardware used currently across all 3 layers in a VBlock.

VBlock Model	VBlock 200	VBlock 300	VBlock 500	VBlock 700
Computer Server Types	Cisco C Series Rack Mounted	Cisco B-series Blade Servers		
Min/Max Servers (Always subject to change)	3/8	4/12	2/128	2/384
Network	Unified	Seggregated or Unified		
Storage	EMC VNX		EMC XtremIO	EMC VMAX

Different types of media are used in storage, like magnetic and solid state drives. For magnetic drives, there is Fibre Channel, SATA, and Serial Attached SCSI drives. EMC storage uses both magnetic and

nonmagnetic drives and moves data across different tiers automatically using Fully Automated Storage Tiering (FAST®) technology. FAST assess the frequency of data access and moves hot data to high performance tiers and cold data to low performance tiers.

It is very important to understand basics of drive and RAID-related technologies and terms before getting into Storage sizing.

Drive types are used based workload and capacity requirement.

- 15K SAS/FC Drives can address approximately 150 IO/s
- 10K SAS/FC Drives can address approximately 150 IO/s
- 7.2K NLSAS Drives can address approximately 150 IO/s
- EFD Drives can address approximately 2500 IO/s

RAID (Redundant Array of Independent Disks) is a data storage virtualization technology that combines multiple physical drives of the same type into single logical units to provide optimized capacity, better performance, and improved data protection. Different RAID types are selected based on workload and capacity requirement preferences.

Protection	Performance	Resiliency
RAID 1	Best	Better
RAID 5	Better	Good
RAID 6	Good	Best

Single host IO can be converted to multiple IO's based on RAID type used. Most reads are addressed with single IO operation. Random write overhead is different for different RAID types; this is called Write Penalty. Write Penalty differs for different RAID types.

RAID Type	Write Penalty	IO conversion
RAID 1	2	2disk writes for every host write
RAID 5	4	2reads, 2writes for every host write
RAID 6	6	3reads, 3writes for every host write

Data Gathering

The first step is to determine the number of applications, servers, and storage targeted to move to the VBlock.

Collect information and background data about existing storage. It is important to ask right questions. Context questions can provide valuable insight. Understand if there are any performance issues in existing storage.

Data can be collected through:

- Customer Interaction
- Tool-based discovery

Below are some of the tools used to collect storage data.

Data Collection Tools		
Host	Tools	Log Collection
Vmware	Vmware Power CLI	Execute task script to gather performance and storage data from vCenter Server
Windows	WMI queries/logman	Collect configuration data from with WMI queries and then uses logman to collect
Storage	Tools	Log Collection
VMAX	Unisphere/storstpd daemon	Gather configuration files (symapi_db.bin or impl.bin) and performance files (btp/ttp) from the storstpd daemon
VNX	Unisphere	Collect SPCollect and NAZ/NAR files from VNX/Clariion
XtremIO	xmcli	Issue XtremIO CLI command xmcli(admin)>create-debug-info to generate set of xtremio log files
HDS 9K	runWin.bat/runUnix.bat	Use runWin.bat/runUnix.bat script from the host that has access to the storage
NetApp	Autosupport	Zip autosupport directory: ETC\$/log/autosupport and the Archive stats directory: ETC\$/log/stats/archive

Data Consolidation and Analysis

The next step would be to consolidate and analyze collected data. Mitrend is one of the tools which would consolidate and analyze the collected data across different platforms and provide consolidated results. These results can be used to understand current configuration and performance of the hardware (Physical/Virtual).

The table below gives an idea about required performance input files of different storage arrays which can be loaded onto Mitrend for consolidation and data analysis.

Storage	Compellent	Dell Equallogic	HDS (99xx, USP, NSC)	HDS (AMS,WMS,HUS)	HDS High End	HP EVA	HP 3PAR	IBMDS (DS 3000,4000,5000)
Performance Data	compellent_db.bak	.csv	Export tool	pfm*.txt	Export tool	evaperf.out	.ins & strfiles	ibmds_perfcollect.bat

Storage	IBM Storage (XIV,SPC,DS)	Isilon	Netapp	VNX File-Celerra	VNX-Clariion	VNX Data Skew	VMAX Symmetrix
Performance Data	IBM TPC Export	isi_gather_perf.py	(/vol0/log/stats/archive)	collect-perfdata.csv.pl	NAR (or NAZ)	FAST VP Pool	ttp/btp/pcf

Once the data is consolidated and analyzed, we need to record key metrics. Metrics to be recorded that are required in our storage sizing exercise are:

- **Peak Workload**
- **Configured Capacity (usable)**
- **Read/Write %**
- **Read/Write IO Size**
- **Cache Hit%**
- **Skew**
- **Sequential/Random**

The matrix above is determined for all storage boxes/Servers in the existing environment. We need to combine IO Profile and skew from all Storage/Servers to derive a consolidated IO Profile and skew with Total Peak Workload and Total Capacity.

Capacity Sizing

For highly performance oriented workloads, XtremIO is the right choice. Meanwhile, VNX or VMAX will be best suited for performance and capacity oriented heavy applications. VNXe will be the right candidate for low performance/capacity.

VNX/VMAX sizing

VNX/VMAX is an Enterprise Storage solution that caters to high workload and capacity requirements.

The goal of sizing is to model the most cost-effective and balanced configuration possible. The Sizing exercise should process any combination of DMX, VMAX, or Unified arrays. The default procedure assumes the existing usable capacity and workload for the source and proposed solution, but growth can be included for capacity, performance, or both. Sizing will be based on FAST/FAST VP which utilizes all three storage tiers (EFD, FC/SAS, SATA/NLSAS) to move hot and cold data across, thereby optimizing drive utilization.


As per EMC best practices, the RAID types below will be used for different tiers.

- **EFD – RAID5**
- **SAS/EFD – RAID 1 or RAID5**
- **SATA/NLSAS – R6**

Total Required drives can be calculated manually by using this formula:

$$\text{Total Number of Drives} = \frac{\text{Total Read IOPS} + (\text{Total Write IOPS} * \text{Write Penalty})}{\text{Disk Speed IOPS}}$$

Skew plays important role in determining the drive configuration. Highly skewed workload would result in a 2Tiered configuration with EFD as one of the tiers and Low skew would result in 3Tiered configuration. If workload and capacity requirement is low with low skew may result in a single tiered configuration.



Skew is calculated by identifying the point at which % of capacity and % of IOPS = 100
If the skew is 70%, that means 70% of IOPS is handled by 30% of the capacity.

Ideally VCE prefers advanced tools such as VMAX and VNX **Sizer** to be used to model the targeted storage array. Determined workload, capacity, skew, and IO Profile can be used as input to Sizer and Sizer would model VNX/VMAX model, drive configuration, Fast Cache for Unified, and Cache and Engine for VMAX using FAST/FAST VP by maintaining 6 '9's availability (99.9999%)

XtremIO Sizing

XtremIO is all-flash storage and uses thin provisioning and on-demand capacity allocation. XtremIO uses data reduction features of inline deduplication and compression to efficiently store data. Necessary metrics needed for XtremIO Sizing are Capacity, workload, and IO profile. The type of application is needed to estimate deduplication and compression factors.

A similar data collection and consolidation method is used to derive required workload, capacity, and IO Profile. Skew will be not necessary for All Flash array solutions. XtremIO is a scale out architecture which uses building blocks called XBrick which consists of Storage Controller, Disk Array Enclosure as fully populated with 25 flash drives (400GB or 800GB), and Infiniband switches which can be clustered together to grow performance and capacity as required.

XBrick comes in 3 variances:

- **With 10TB XBrick**
- **With 20TB XBrick**
- **With 40TB XBrick**

Below is the information which outlines Deduplication and compression ratio for different applications.

Application	Deduplication (x:1)	Compression (x:1)
DB - Oracle	1.10	2.20
DB - SQL	1.10	1.50
VDI - Full Clones	6.10	1.30
VDI - Linked Clones	1.50	1.30
VMware	1.50	1.30
SAP HANA	1.60	1.00

Capacity listed below is as per XBrick configuration by considering **Usable with no data reduction**.

Xbrick Model	Number of SSDs	Raw Capacity	Usable Capacity
10TB	25 x 400GB	10TB	7.58 TiB
20TB	25 X 800GB	20TB	15.16 TiB
40TB	25 x 1.6TB	40TB	30.60 TiB

It is a best practice to always maintain capacity utilization of an XBrick within 80%. For example:

- 10TB XBrick can have up to 7.58TB Usable Capacity
- Based on the IO Profile, if the deduplication is 3:1, the 10TB XBrick can have up to (7.58*3) TB capacity = 22.74 TB Capacity to remain within 80% of Capacity Utilization
- Total Capacity allowed in 10TB XBrick with 3:1 Dedup = 22.74 * 0.8 = 18.19 TB Capacity

XtremIO Performance with **100% Random** with **70/30 Read/Write** ratio can achieve workload as per the table below.

	1 X-Brick	2 X-Bricks	4 X-Bricks	6 X-Bricks
IOPS (70% Read / 30% Write)	150,000	300,000	600,000	900,000

EFD Drives can address approximately 2500 IO/s. There is no RAID Penalty in XtremIO.

VSphere Design considerations

Virtualization is a technology that allows physical resource to be represented as a number of logical resources that can do and behave same like physical resources. In VBlock, virtualization is applied to Compute, Network and Storage.

A virtualized physical computer is called virtual machine (VM). A single physical server can be divided into several VMs which behave like a physical server. Hypervisor is the software which is loaded onto a physical server which creates and manages VMs. It also provisions resources to a VM as per the requirement.

Compute Virtualization

Server virtualization provides resource optimization like load balancing and rapid provisioning. Virtual machine is all about shared CPU and Memory utilization. The number of virtual machines needed to run production applications will be an input from the customer or we can assume 1:1 ratio for physical server to virtual machine.

The number of vCenter instances and number of virtual machines on a single UCS blade server running on VMware ESXi is according to VMware best practices.

The table below offers insight to VMware scalability for CPU and Memory provisioning.

	vSphere 5.5	vSphere 6
Hosts per Cluster	32	64
VMs per Cluster	4,000	8,000
CPUs per Host	320	480
RAM per Host	4 TB	12 TB ¹
VMs per Host	512	1,024
Virtual CPUs per VM	64	128
Virtual RAM per VM	1 TB	4 TB

Virtual CP to Physical CP and is based on type of applications however as a standard for general workload, a good rule of thumb would be to start with:

- **4:1 for virtual CPU to physical CPU oversubscription**
- **1.2:1 for virtual RAM to physical RAM oversubscription**

All virtual machines must be capable of being powered on simultaneously. Unused or unnecessary virtual hardware devices can impact performance and should be disabled. It is a best practice to provide minimum resources to the VM initially and after validation period, resource can be provisioned as per the requirement

Network Virtualization

Network virtualization combines network functionality into software-based virtual network and provides logical extension of physical network. Functionality includes:

- **Network Isolation**
- **Virtual Port Channels for redundancy**
- **VLAN design for multi tenancy**
- **vSphere Standard and Distributed switch**

Each hypervisor has an embedded virtual switch that provides VM-to-VM connection within the same LAN. Virtual Network can isolate tenants from one another or isolate management components from tenants using virtual LANs. VLAN provide for logical groupings of switch ports allowing communications as if all ports were on the same physical LAN.

Using vSphere Standard Switch, each server has its own virtual switch. VSS can route traffic internally between virtual machines and link to external networks.

A vSphere Distributed Switch functions as a single virtual switch across all associated hosts. This function allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts.

Storage Virtualization

Storage virtualization will provide functionality which minimizes system costs, simplifies the process of storage management, and optimize storage utilization.

There are 4 types of storage in VBlock systems:

1. Boot
2. Datastore
3. RDM
4. NAS

Boot device contain ESX operating system.

Datastore is the logical container that holds Virtual Machine and Virtual Disk files. Datastore can be VMFS-based or NFS-based. These datastores can be shared across ESX clusters.

RDM (Raw Device Mapping) provides virtual machines with direct access to the LUN.

NAS (Network Attached Storage). ESX supports NAS through NFS Protocol and CIFS Protocol.

Boot Device sizing considerations

- Separate set of boot device is configured for ESX and Non-ESX server
- For every 32 blades, 20GB of thick LUNs to be configured as boot device
- Boot device can be spread across available drives

Datastore Sizing considerations

- VMFS volumes can be configured up to 64TB with each extent supporting up to 2TB and can have up to 32 extents
- VMFS can support up to 30,000 files and maximum file size for each file can be up to 2TB
- Always consider 20 – 30% additional capacity to accommodate snapshots, .vswp, and .log files
- Consider separate datastores for Production and non-Production
- Be aware of 256 LUN limitations

NAS

- Uses IP for data traffic
- Used for unstructured data
- Required link aggregation data access by many hosts at one time
- Virtual disks created on NFS data stores use a disk format dictated by the NFS server
- NAS accept different IO size from client, however it uses IO Merge to coalesce or break IO size to make it 8K. Hence, IO Size for NAS head is 8K.
- There are two options for NAS storage within VCE solution. Internal is VNX and external can be Isilon®.

Storage sizing for a virtual environment is not much different from sizing for a physical environment. Storage design is an aggregate workload for multiple servers on a single ESXi host. If possible, size each application and their associated server separately and then aggregate the results. VMware Capacity Planner is a tool used to enable faster analysis and consolidation with planning and recommendations.

References

[Student Guide_Architecting a VBlock solution](#)

Dell 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." DELL EMC 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 Dell EMC software described in this publication requires an applicable software license.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.