

DD VE in AWS and VMC

Best Practices Guide

Table of contents

Contents

Table of contents	1
1 Acronyms	2
2 Purpose	2
3 Scope	3
4 General Best Practices	3
4.1 System configuration for DD VE in AWS cloud.....	3
4.2 S3 connectivity for DD VE in AWS cloud.....	3
4.3 System configuration for DD VE in VMC (VMWare Cloud on AWS).....	4
4.4 S3 connectivity for DD VE in VMC	5
4.5 Supportability	5
4.6 ASUP configuration	5
4.7 System headswap.....	5
5 Storage Best Practices	6
5.1 Storage configurations for DD VE on S3	6
5.2 Storage configurations for DD VE on block storage	7
5.3 Moving from evaluation to production.....	8
5.4 Replicating data	8
5.5 Device Mapping.....	9
6 Security Best Practices	9
6.1 Public IP address.....	9
6.2 Default Password.....	9
6.3 User Authentication Methods.....	10
6.4 AWS Security Groups.....	10
6.5 IP Tables feature.....	12
7 Networking Best Practices	13
7.1 VPC Architecture	13
7.2 Public IP Addresses.....	13
7.3 Number of interfaces and IP addresses	13
7.4 Default DHCP configuration	13
7.5 Alias interfaces in DD VE	14
7.6 Attaching an elastic network interface to DD VE	14
7.7 Asymmetric routing.....	15

7.8	Detaching an elastic network interface from DD VE	15
7.9	Bonding.....	15
7.10	VLAN interfaces	16
7.11	Configure NTP server details	16
8	Appendix.....	17
8.1	Ports used by Data Domain Systems for inbound traffic	17
8.2	Ports used by Data Domain Systems for outbound traffic.....	19
9	Reference Documents	20

1 Acronyms

Acronym	Description
AWS	Amazon Web Services Cloud
Bonding	Combining more than one network interface for aggregation or failover purposes.
DDR	Data Domain Replicator
DD VE	Data Domain Virtual Edition. It is the virtual appliance version of the DDR.
MAC address	Media Access Control address. The virtual network adapter will have a MAC address given by AWS.
MTU	Maximum Transmission Unit
NTP	Network Time Protocol
Virtual Network Adapter	Refers to the physical like adapter that is created per VM on the Host.
VLAN	Virtual Local Area Network
VM	Virtual Machine running on the Host. In the context of this document it is DD VE
VMC	VMWare Cloud on AWS

2 Purpose

The purpose of this document is to provide general guidelines for storage, security and networking best practices for running DD VE in AWS and VMC.

3 Scope

The scope of this document is limited to DD VE's storage and networking function in the AWS environment. The extensive storage & networking features of AWS are not in the scope of this document

4 General Best Practices

This section provides information on the best practices to configure DD VE in the Amazon Web Services cloud.

4.1 System configuration for DD VE in AWS cloud

AWS provides a long list of EC2 instance types for various customer's needs. DD VE in AWS supports the following EC2 instance types for the corresponding configuration capacities on S3 and block storage.

Instance type	M4.xlarge	M4.2xlarge	M4.4xlarge
CPU	4	8	16
Memory (GiB)	16	32	64
System disks	250 GiB GP2 Root disk	250 GiB GP2 Root disk	250 GiB GP2 Root disk
	10 GiB GP2 NVRAM disk	10 GiB GP2 NVRAM disk	10 GiB GP2 NVRAM disk
Storage capacity for DD VE on S3	16 TB	32 TB	96 TB

Table 1: System configuration for DD VE in AWS

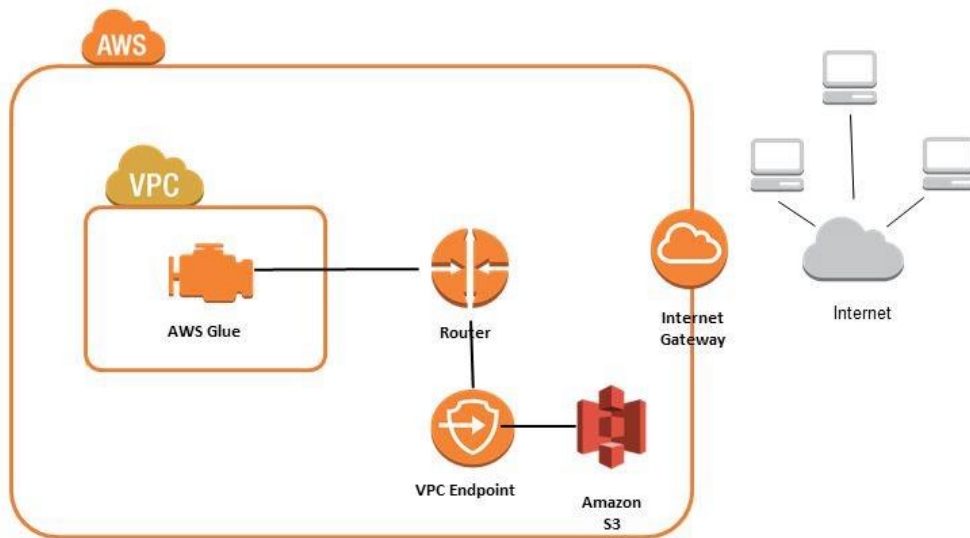
Please note that for DD VE's on block storage, the maximum allowed capacity is 16 TB for all the supported (refer Table 1) instance types.

4.2 S3 connectivity for DD VE in AWS cloud

The DD VE object store feature needs connectivity to its object storage, such as to the S3 bucket. The object store communication is over https, so the outbound

security group setting must allow communication over port 443. There are different ways to enable DD VE connectivity to the object store. Out of the following three we recommend only the third option (Using VPC endpoint).

- Using the public IP from the public subnet: should not be used
- Using NAT (Network Address Translation): If the private subnet is configured to use NAT, then DD VE will be able to communicate to object store over NAT.
- We strongly recommend using VPC endpoint for accessing the Amazon S3. It does not require the DD VE to have a public IP address to communicate to S3, it uses the private IP address instead. (In this case, an internet gateway, NAT, or virtual private gateway are not needed to access S3). This method also allows the traffic to the S3 endpoint to stay within the Amazon network and will be routed internally to S3.



Also, please note that the DD VE instance and the S3 bucket that was created for that instance must be in the same region. The role must be attached to the DD VE instance prior to configuring the object store feature.

4.3 System configuration for DD VE in VMC (VMWare Cloud on AWS)

Instance Type	DD VE Capacity	#vCPU, Memory
*Standard_VMC_16	16 TB	4, 16 GB
*Standard_VMC_32	32 TB	4, 24 GB
*Standard_VMC_96	96 TB	8, 64 GB

Table2: System configuration for DD VE in VMC

* Please note that the instance type names are just logical names given to the corresponding compute resources.

4.4 S3 connectivity for DD VE in VMC

For deploying DD VE in VMware Cloud on AWS (VMC) on S3 object store, you will need to setup a SDDC in VMC. During the setup, you will need to attach the VMWare cloud account with an AWS account/vpc subnet and ensure that both the SDDC and the subnet selected in the AWS account are in the same region.

Create the S3 bucket used by the DD VE in the same region as the SDDC and within the same AWS account. If the DD VE in VMC and the bucket are in different regions, performance may get impacted and additional costs will be incurred. Also, make sure that in the AWS VPC, the S3 endpoint is created. This will ensure that all object store traffic is routed internally within the AWS infrastructure.

4.5 Supportability

AWS EC2 instance does not support interaction with console, but customers can get read-only access to console through the Instance screenshot feature available in AWS.

If you wanted to use ESRS with DD VE in AWS, ESRS gateway needs to be deployed in the cloud.

4.6 ASUP configuration

Set up the following to ensure that autosupport (ASUPs) and alert emails from your system are sent to EMC Data Domain.

- a. Administrator: Enter a password and email address for the Administrator.
- b. Email/Location: Enter the mail server used to send outgoing alert and ASUPs to recipients. Recipients are subscribers to groups. A group named default is created with the email address of two subscribers: the administrator and `autosupportalert@autosupport.datadomain.com`. The Location field is simply for your information, only.
- c. Summary: Review the summary carefully. The default address for alerts and autosupport emails is `autosupportalert@autosupport.datadomain.com`. A detailed autosupport and an alert summary is scheduled to run “daily” at “0600”.

4.7 System headswap

On the target system (system B), before running the headswap command, set the system passphrase to match exactly with the passphrase of the source system (system A). Without this step, the headswap command will fail.

Also, ensure that the system A is powered off before issuing the headswap command on the system B. This is needed to ensure that the bucket gets detached from system A and is available to be attached to system B.

5 Storage Best Practices

There is no need to specify spindle group or change their settings when adding storage. The spindle group assignment is balanced automatically when storage is added. After storage is added, it is recommended to run "storage show all" to verify each data volume has been assigned to different spindle group.

5.1 Storage configurations for DD VE on S3

For AWS, two system disks, an EBS GP2 250 GiB (root disk for DDOS) and an EBS GP2 10GiB (for NVRAM simulation) are needed to deploy the DD VE.

The recommended metadata storage is 10% of the current active tier capacity. Metadata disks should be added incrementally in 1 TiB increments to reach up to the supported system capacity.

DD VE Configuration	Instance Type	Block storage volumes			Object Storage Capacity	Network Interface
		Root Disk	NVRAM Disk	Metadata Disks (Each disk size = 1 TiB)		
16TB	m4.xlarge	GP2 / 250 GiB	GP2 / 10 GiB	GP2 / (1 - 2 Disks)	0 – 16 TB	Default = 1 SRIOV recommended
32TB	m4.2xlarge	GP2 / 250 GiB	GP2 / 10 GiB	GP2 / (1 - 4 Disks)	0 – 32 TB	Default = 1 SRIOV recommended

96TB	m4.4xlarge	GP2 / 250 GiB	GP2 / 10 GiB	GP2 / (1 - 10 Disks)	0 – 96 TB	Default = 1 SRIOV recommended
------	------------	---------------	--------------	----------------------	-----------	-------------------------------------

Table 3: Storage size specification of metadata disks for DD VE on S3

Data storage configuration Notes

- When configuring DD VE on S3 storage for AWS, make sure that the maximum length of the bucket name does not exceed 48 characters.
- Bucket provided during file system creation must be empty, otherwise bucket will not attach to the filesystem and it will not get created.
- When the file system is destroyed, associated bucket is neither deleted nor the objects within are removed, one need to explicitly delete the bucket to avoid cost incurred with the content stored in the object store.

Instance Type	Number of metadata disks (each disk =1 TiB)	Read	Write	Replication In	Replication Out	Combined
m4.xlarge	1	12	36	36	24	36
	2	24	36	36	36	36
m4.2xlarge	1	12	48	48	24	48
	2	24	72	72	48	72
	>=3	40	72	72	72	72
m4.4xlarge	1	12	48	48	24	48
	2	24	96	96	48	96
	>=3	40	144	144	72	144

Table 4: Supported stream counts for DD VE on S3

5.2 Storage configurations for DD VE on block storage

For basic deployment, please use GP2 for root disk, NVRAM simulation and data disks. For deployment without intensive read traffic, ST1 can be used for data disks.

The derive the maximum IOPS, the recommended disk size is 1 TiB for GP2 volumes and 2 TiB for ST1 volumes. Please refer to the link below for more information

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

To get the maximum load balancing under heaviest loads, it is recommended to add multiple disks of same size for higher capacity points. If m4.4xlarge is used as EC2 instance, it is not recommended to use ST1 as data storage, as its performance does not match what m4.4xlarge is able to deliver.

In order to achieve consistent overall performance, please make sure to use the same EBS volume types for data storage. In other words, do not mix the GP2 and ST1 volumes for data storage

Instance Type	DD VE Capacity	Storage type: ST1 Number of Volumes x Capacity (TiB)	Storage type: GP2 Number of Volumes x Capacity (TiB)
m4.xlarge	16TB DD VE	8 x 2	16 x 1 or 8 x 2
m4.2xlarge	16TB DD VE	8 x 2	16 x 1 or 8 x 2
m4.4xlarge	16TB DD VE	Not recommended	16 x 1 or 8 x 2

Table 5: Storage specification for DD VE on block storage

5.3 Moving from evaluation to production

It is strongly recommended to go for fresh deployment instead of upgrading the evaluation version of DD VE.

If in case one want to go with upgrade path, then the recommendation is to destroy the existing file system, delete any smaller data disks (not the root, NVRAM volumes), and configure new volumes as per the recommendations in above sections.

5.4 Replicating data

It is highly recommended and always a best practice to replicate the data into another DD VE in another availability zone (AZ), or DD VE in another region or even to on premises DD VE or DDR.

While replicating to another AZ helps in handling instance failures, but for disaster recovery it is recommended to replicate to another regions or to on premises.

5.5 Device Mapping

When the user attaches a new EBS volume to DD VE, a device name can be specified. By default, AWS provides a name which is sd* or xvd*. The default name can be used as is. To see the device mapping, run the “disk show hardware” CLI command in the DD VE.

```
SE@ip-10-1-0-220## disk show hardware
```

Disk	Slot (pci/idx)	Manufacturer/Model	Firmware	Serial No.	Capacity	Type
dev1	-/a	Virtual BLOCK Device	n/a	(unknown)	250.00 GiB	BLOCK
dev2	-/b	Virtual BLOCK Device	n/a	(unknown)	10.00 GiB	BLOCK
dev3	-/f	Virtual BLOCK Device	n/a	(unknown)	500.00 GiB	BLOCK
dev4	-/g	Virtual BLOCK Device	n/a	(unknown)	500.00 GiB	BLOCK
dev5	-/h	Virtual BLOCK Device	n/a	(unknown)	1000.00 GiB	BLOCK
dev6	-/i	Virtual BLOCK Device	n/a	(unknown)	500.00 GiB	BLOCK
dev7	-/j	Virtual BLOCK Device	n/a	(unknown)	100.00 GiB	BLOCK
dev8	-/k	Virtual BLOCK Device	n/a	(unknown)	500.00 GiB	BLOCK

```
8 drives present.
```

We can use “Slot(pci/idx)” area to map the disk in DD VE(dev*) to the device we see in AWS. If the “Slot(pci/idx)” section is “a”, then in AWS, its corresponding “Block Device” should be “/dev/sda1”. For all other cases, if “Slot(pci/idx)” is X, then in AWS its corresponding “Block Device” should be “/dev/sdX” or “/dev/xvdX”.

6 Security Best Practices

6.1 Public IP address

In order to prevent various brute force attacks on DD VE, it should not be exposed using public IP address.

6.2 Default Password

For DD VE in AWS, the default password for the "sysadmin" account is the instance id. For DD VE in VMC, the default password for the “sysadmin” account is changeme.

These passwords are system generated and assigned respectively during the initial system boot up time. Once you login into the DD VE for the first time, you will be forced to change the default passwords. Please choose a strong password to protect access to your system.

6.3 User Authentication Methods

Following table illustrates the different authentication methods supported by DD VE.

Access Type	Authentication Methods
GUI	username/password X509 certificates
SSH	username/password SSH Keypair
REST Api	username/password X509 certificates

Table 6: Authentication methods supported by DD VE

For better security it is recommended to disable the username/password based user authentication. If the username/password based authentication is desired, it is recommended that a stronger password policy is configured.

6.4 AWS Security Groups

For DD VE in AWS, it is often running in a VPC, the VPC should be configured so that only required and trusted clients have access to the Data Domain system. Security groups in AWS restrict access to an instance based on the

1. Port
2. IP range
3. Security group (its own or another)

Inbound control

The security groups are stateful which means that the responses to the inbound traffic will be allowed to go out regardless of outbound rules. The following are the inbound ports that are allowed for DD VE.

Port	Service	Description
TCP 22	SSH	Used for SSH (CLI) access and for configuring DD VE.
TCP 443	HTTPS	Used for DDSM (GUI) access and for configuring DD VE.
TCP 2049	DD Boost/NFS	Main port used by NFS - can be modified using the nfs set server-port command which requires SE

		mode
TCP 2051	Replication/DD Boost/Optimized Duplication	Used only if replication is configured (run replication show config on Data Domain system to determine). This port can be modified using replication modify.
TCP 3009	SMS (system management)	Used for managing a system remotely using Data Domain System Manager. This port cannot be modified. This port is used only on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the DataDomain System Manager, as the replication partner needs to be added to the Data Domain System Manager

Table 7: Inbound ports allowed for DD VE

Depending on the protocol that is used to backup data to DD VE, additional ports will be allowed with inbound security group rules. For a complete list of all ports allowed for inbound traffic for data domain systems, refer [Inbound Ports Table](#)

Outbound control

As stated earlier the security groups are stateful, which means that if a request is allowed to be sent out of a DD VE, its responses will be allowed regardless of inbound rules. The following are the outbound ports that shall be allowed for DD VE.

Port	Service	Description
UDP 123	NTP	Used by the Data Domain system to synchronize to a time server.
TCP 443	HTTPS	Used for DD VE to be able to communicate with Object store (S3).
TCP 2049	DD Boost/NFS	Main port used by NFS - can be modified using the

		nfs set server-port command which requires SE mode.
TCP 2051	Replication/DD Boost/ Optimized Duplication	Used only if replication is configured (run replication show config on Data Domain system to determine). This port can be modified using replication modify..
TCP 3009	SMS (system management)	Used for managing a system remotely using Data Domain System Manager. This port cannot be modified. This port is used only on Data Domain systems running DDOS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the DataDomain System Manager, as the replication partner needs to be added to the Data Domain System Manager

Table 8: Outbound ports allowed for DD VE

Depending on the other applications/services that are being used, additional ports shall be allowed for outbound security group rules. For a complete list of all ports allowed for outbound traffic for data domain systems, refer [Outbound Ports Table](#)

6.5 IP Tables feature

After protecting the DD VE with secure setup, with in the DD VE we can filter the network traffic that enters by making use of iptables feature. For more information on configuration, please refer to DD 6.1 command reference guide's net filter section.

7 Networking Best Practices

7.1 VPC Architecture

We recommend you use public or private subnet architecture to deploy the DD VE in private subnet. It will secure the DD VEs (VMs) with the appropriate use of various VPC components such as route tables, access control lists, security groups, etc.

7.2 Public IP Addresses

Due to security considerations and in order to protect the DD VE from potential attacks over open internet, the DD VE MUST NOT be exposed using Public IP directly over internet. It is highly recommended that you use VPN connections between different geographical regions (VPCs). For example, the replication between different VPCs, different cloud regions, cloud to on-premise and vice versa can be used via the secure VPN connection.

7.3 Number of interfaces and IP addresses

Deploy DD VE with one network interface. As mentioned by AWS, increasing number of network interfaces will not help in increased bandwidth.

DD VE officially supports 8 interfaces. The first interface is considered as primary and user cannot detach it from the instance.

However, depending on instance type there is a limit in AWS for number of elastic interfaces that can be added to an instance and also on the number of IP addresses that can be assigned to an interface. For more information on the number of elastic network interfaces support please refer to the following link:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html>

For more information on 'how to configure multiple IP address' please refer to the below link:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html>

7.4 Default DHCP configuration

DHCP is by default enabled for up to two interfaces in the DD VE. For the additional interfaces (if exists) either DHCP can be manually enabled, or those interfaces can be configured manually. All the interfaces in DD VE can be configured manually using static IP addresses. However please make sure that the IP addresses are known to corresponding elastic network interfaces in AWS.

7.5 Alias interfaces in DD VE

DD VE supports only one IP address per interface. Having said that, DD VE allows creating number of 'alias' interfaces that can sit over a base (elastic interface) interface. Each alias interface can be configured with appropriate IP address.

We can create the number of alias interfaces in the DD VE on top of a base interface. For more information on how to create alias interface or its configuration, please refer to the 6.1 Admin guide, section "network interface management" -> "Configuring an IP alias".

Every IP address that is configured in the DD VE must be known to AWS, otherwise the routing/switching packets using that IP address in AWS environment will not work.

Which means that we have to specify (either auto assigned or manual) the IP addresses in the AWS environment first and then they can be used to configure the interfaces in the DD VE.

Please refer to the following screen shot.



The secondary IP addresses that are configured in AWS can be used to configure the alias interfaces (that are sitting on top of the corresponding base interface) in the DD VE.

The primary IP address should not be used to configure the alias interfaces. If one wants to configure the primary IP address on the alias interface in DD VE, make sure that...

- 1) The DHCP is not enabled on the base interface. If enabled, it results in getting the primary address assigned to the base interface.
- 2) The primary address is not statically assigned on base interface.

7.6 Attaching an elastic network interface to DD VE

Other than the default interface, there can be situations that might require to add more network interfaces to VMs. DD VE supports the following scenarios for attaching an elastic network interface.

Please refer to the below table:

AWS attach Procedure	Meaning	DD VE support
Hot attach	when the instance is running	Not yet supported
Warm attach	When the instance is stopped	Supported
Cold attach	When the instance being launched	Supported

Table 9: Supported interfacing attachment methods for DD VE in AWS

7.7 Asymmetric routing

As mentioned by AWS documentation, it is worth noting the fact that if more network interfaces are attached from the same subnet, there are chances of encountering networking issues like asymmetric routing. In this case the packet can go out of one interface, but the response can come on a different interface.

DD VE technically accepts such packets, hence we don't foresee any functionality impact.

7.8 Detaching an elastic network interface from DD VE

Just like the case of attaching interface, DD VE does not support detaching an interface while it is running. One must shutdown/stop the DD VE before detaching an interface.

7.9 Bonding

Bonding multiple network interfaces within the DD VE is not supported.

7.10 VLAN interfaces

Although DD VE supports VLAN interfaces, AWS environment does not support VLANs.

7.11 Configure NTP server details

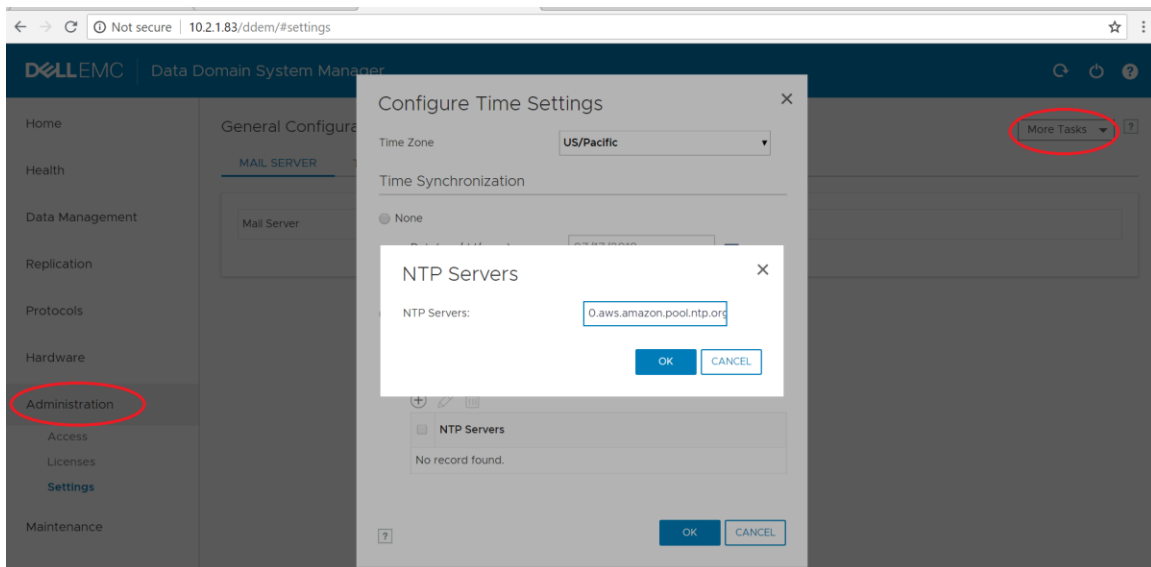
By default, NTP is disabled on the DD VE system. But is important for the DD VE's time to be properly synchronized when running in AWS. Any drift in time may impact secure communication from the object store. Therefore, NTP must be configured for the DD VE that is running in AWS. While performing initial configuration of the DD VE system, enable NTP and configure the NTP server.

According to AWS documentation, if you don't have your own NTP server, use the following NTP server from AWS.

server 0.amazon.pool.ntp.org

Procedure

1. Select the settings under the Administration tab.
2. Select "Configure Time Settings" from the drop down menu of "More Tasks".
3. Select the "Manually Configure" option under NTP and add the NTP servers as 0.amazon.pool.ntp.org



Run the following commands to configure NTP on the DD VE (using CLI)

```
ntp add timeserver 0.amazon.pool.ntp.org
```

```
ntp enable
```

```
ntp sync
```


8 Appendix

8.1 Ports used by Data Domain Systems for inbound traffic

Port	Service	Description
TCP 21	FTP	Port is used for control only if FTP is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 22	SSH	Port is used only if SSH is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 23	Telnet	Port is used only if Telnet is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 80	HTTP	Port is used only if HTTP is enabled (run 'adminaccess show' on the Data Domain system to determine if this is the case).
TCP 111	DDBOOST/ NFS (portmapper)	Used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
UDP111	DDBOOST/ NFS (portmapper)	Used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned
UDP 123	NTP	Port is used only if NTP is enabled on the Data Domain system. Run ntp status to determine if this is the case.
UDP 137	CIFS (NetBIOS Name Service)	Port used by CIFS for NetBIOS name resolution
UDP 138	CIFS (NetBIOS Name Service)	Port used by CIFS for NetBIOS Datagram service
TCP 139	CIFS (NetBIOS Name Service)	Port used by CIFS for session information
UDP 161	SNMP (Query)	Port is used only if SNMP is enabled. Run 'snmp status' to determine if this is the case.
TCP 389	LDAP	LDAP server listens on this port for any LDAP client request. By Default it uses TCP
TCP 443	HTTPS	Port is used only if HTTPS is enabled (run adminaccess show on the Data Domain system to determine if this is the case).

TCP 445	CIFS (Microsoft-DS)	Main port used by CIFS for data transfer.
TCP 2049	DD Boost / NFS	Main port used by NFS. Can be modified via the 'nfs set server-port' command. Command requires SE mode.
TCP 2051	Replication / DD Boost / Optimized Duplication	Port is used only if replication is configured on the Data Domain system. Run replication show config to determine if this is the case. This port can be modified via the replication modify command.
TCP 2052	NFS Mountd / DD BOOST / Optimized Duplication	Main port used by NFS MOUNTD
TCP 3009	SMS (System Management)	Port is used for managing a system remotely using Web Based GUI DD EM (Data Domain Enterprise Manager). This port cannot be modified. This port is only used on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain GUI interface, as the replication partner needs to be added to the DD Enterprise Manager.
TCP 5001	iPerf	Port is default used by iperf. To change the port, it requires -p option from se iperf or port option from the net iperf command. The remote side must listen on the new port.
TCP 5002	Congestion-checker	Port is default used by congestion-checker, when it runs iperf. To change the port the new port needs to be specified in the port option of the net congestion-check command. The remote side must also be listen on the new port. It is available only for DD OS 5.2 and above.

Table 9: Complete list of ports allowed by Data Domain Systems for Inbound Traffic

8.2 Ports used by Data Domain Systems for outbound traffic

Port	Service	Description
TCP 20	FTP	Port is used for data only if FTP is enabled (run <code>adminaccess show</code> on the Data Domain system to determine if this is the case).
TCP 25	SMTP	Used by the Data Domain system to send email autosupports and alerts
UDP/TCP 53	DNS	Port is used by Data Domain system to perform DNS lookups when DNS is configured. Run <code>net show dns</code> to review DNS configuration
TCP 80	HTTP	Used by Data Domain system for uploading log files to Data Domain Support via the support upload command.
UDP123	NTP	Used by the Data Domain system to synchronize to a time server.
UDP 162	SNMP (Trap)	Used by the Data Domain system to send SNMP traps to SNMP host. Use <code>snmp show traphosts</code> to see destination hosts and <code>snmp status</code> to display service status.
TCP 443	HTTPS	Port is used for communicating with Object store (S3).
UDP 514	Syslog	Used by the Data Domain system to send syslog messages, if enabled. Use <code>'log host show'</code> to display destination hosts and service status.
TCP 2051	Replication / OST / Optimized Duplication	Used by Data Domain system only if replication is configured. Use <code>replication show config</code> to determine if this is the case
TCP 3009	SMS (System Management)	Port is used for managing a system remotely using Web Based GUI DD EM (Data Domain Enterprise Manager). This port cannot be modified. This port is only used on Data Domain systems running DD OS 4.7.x or later. This port will also need to be opened if you plan to configure replication from within the Data Domain GUI interface, as the replication partner needs to be added to the DD Enterprise Manager.
TCP 5001	iPerf	Port is default used by iperf. To change the port, it requires <code>-p</code> option from <code>se iperf</code> or <code>port</code> option from the <code>net iperf</code> command. And the remote side must listen on the new port.

TCP 5002	Congestion-checker	Port is default used by congestion-checker, when it runs iperf. To change the port the new port needs to be specified in the port option of the net congestion-check command. The remote side must also be able to listen on the new port. It is available only for DD OS 5.2 and above.
TCP 27000	Avamar client communications with Avamar server	Avamar client network hosts.
TCP 27000	Avamar server communications with Replicator target server (Avamar proprietary communication)	Required if server is used as replicator source
TCP 28001	Avamar client communications with administrator server	Avamar clients required.
TCP 28002	Administrator server communications with Avamar client	Optional for browsing clients and cancelling backups from Avamar administrator management console.
TCP 29000	Avamar client Secure Sockets Layer (SSL) communications with Avamar server	Avamar clients required
TCP 29000	Avamar server SSL communications with Replicator target server	Required if server is replicator source.
TCP 29000	Avamar server SSL communications with Replicator target server	Required if server is replication source.

Table 10: Complete list of ports allowed by Data Domain Systems for Outbound Traffic

9 Reference Documents

Note: Please refer to the latest guides that are available.

6.1 Admin guide

https://support.emc.com/docu85190_Data-Domain-Operating-System-6.1-Administration-Guide.pdf?language=en_US.

6.1 Command reference guide

https://support.emc.com/docu85240_Data-Domain-Operating-System-6.1-Command-Reference-Guide.pdf?language=en_US