

ECS Dedicated Cloud

Technical Overview

August 2017

Revisions

Date	Description
August 7, 2017	Initial release

The information in this publication is provided "as is." Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © August 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [8/10/2017] [Whitepaper] [H15932.2]

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of Contents

Revisions	2
Executive Summary	4
1 Introduction.....	5
1.1 Audience.....	5
1.2 Scope	5
2 Value of ECS Dedicated Cloud	6
3 ECS Dedicated Cloud Overview	7
3.1 Geo-Replication	7
3.1.1 Geo-Data Access (Active-Active).....	8
4 ECS DC Deployment	13
4.1 ECS Traffic Manager (ETM)	14
4.1.1 Load Balancers	14
5 Network Connectivity	20
6 Best Practices	22
7 Support.....	23
8 Conclusion.....	24
9 References	25

Executive Summary

ECS Dedicated Cloud (DC) is a dedicated and customer owned ECS storage managed by Dell EMC and hosted in an off-premise data center such as Virtustream. ECS DC provides an on-demand hosted storage to take advantage of the ECS “exclusive-or” (XOR) feature, which provides increased storage efficiency when three or more sites are employed. It also offers global access and geo-protection.

Cloud computing companies like Virtustream provides an infrastructure as a service as well as cloud computing management software and services to enterprises and companies worldwide. Off-premise hosting companies have data centers or partnerships with service providers located across the world to deliver and manage cloud services. They offer an on-demand, secure, and high performance environment to accommodate customers' requirements.

ECS storage housed in an off-premise data center extends ECS solutions for customers who want to expand to different locations, reduce their storage footprint from their current deployment, and enhanced geo-protection with lower cost. The addition of ECS DC creates a hybrid solution that allows the best of both private cloud control and operation costs of public cloud.

1 Introduction

This document provides a technical overview of Dell EMC® Elastic Cloud Storage Dedicated Cloud (ECS DC). It discusses the value of ECS DC and details on the components required in the deployment of ECS DC hosted in and off-premise data center such as Virtustream and customer's on-premise site(s) are discussed. Best practices for deployment are also described.

1.1 Audience

This whitepaper is intended for Dell EMC field personnel and customers who are interested in understanding the value and technical overview of ECS DC. It describes the components and best practices for deployment of customer's existing ECS storage with ECS DC hosted off-premises.

1.2 Scope

This document focuses primarily on ECS DC. It assumes the readers are well versed on overall ECS architecture covered in the ECS Architecture and Overview whitepaper. It does not cover installation, administration, and upgrade procedures for ECS software or hardware. Refer to ECS Product Documentation for more information on ECS installation and administration.

Updates to this document are done periodically and coincides usually with a major release or new feature and functionality change. To get the latest version of this document, please download from this [link](#).

2 Value of ECS Dedicated Cloud

With the growth of public cloud technologies, customers are eagerly looking for cloud solutions that offer them the best in the terms of security, performance, and lower cost. Private cloud control and on-premise performance and security make it challenging for customers to migrate their data to the cloud. ECS storage hosted off-premises such as at a Virtustream data center extends the capabilities of ECS storage on customer sites. It provides a hybrid cloud solution with the best features of both private and public cloud solutions.

The main value of ECS DC is to improve availability and storage efficiency, and broaden their current ECS deployment to other sites. Other key benefits of implementing ECS DC include optimizing global access by adding another site closer to remote clients and enhancing cloud strategy by providing a hosted site managed by Dell EMC.

ECS DC is targeted for customers interested in the following:

- Expanding their operations to additional data centers closer to other clients who are distant from their other site(s).
- Enhancing geo-protection and global access.
- Growing their existing or planned two site deployment to three sites for increased storage efficiency.
- Desiring to expand but not interested in investing and operating additional data centers.

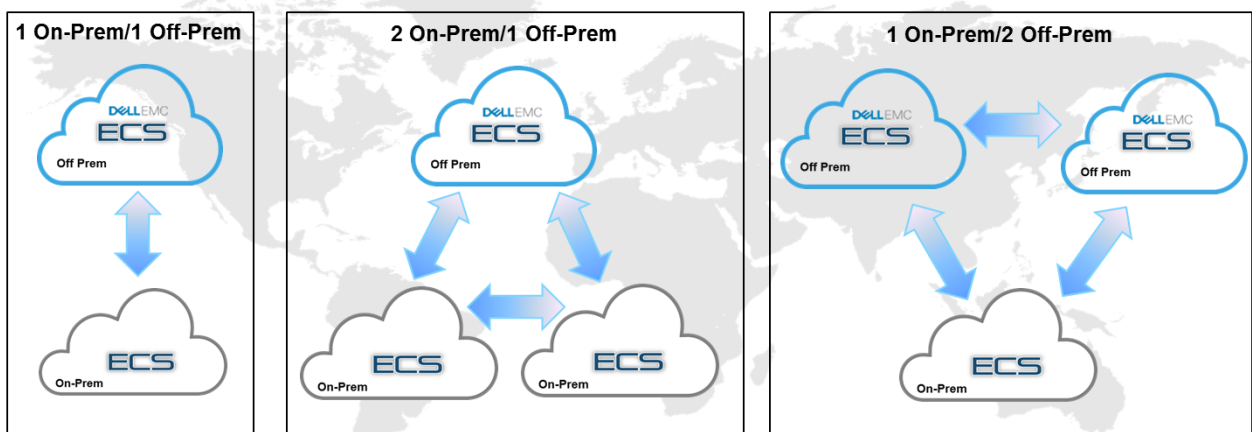
3 ECS Dedicated Cloud Overview

ECS DC is a dedicated ECS storage managed by Dell EMC and hosted in a data center like Virtustream. ECS DC acts as the additional ECS storage site within a replication group. As illustrated in Figure 1, ECS DC is available for new and existing ECS customers who intend to purchase or currently have an ECS on-premise. With the addition of ECS DC, customers receive both the private cloud control and public cloud operation costs. Customers can select from several Virtustream data centers across the United States (i.e. Virginia, Nevada) and Europe (i.e. France, United Kingdom) to host ECS DC. A professional operations team from Dell EMC is responsible for installing, configuring, monitoring and managing the ECS DC rack(s) for customers. Customers retain full ownership and control of the ECS software and hardware. The ECS DC dedicated infrastructure includes customer owned rack with power and designated space, firewall, load balancer, and internet or leased line connectivity.

ECS DC minimum pre-requisites include:

- Greater than 1.9 PB usable capacity (recommended)
- Gen 2 U-Series and D Series Only
- Running ECS version 3.0 (HF1) or later'

Figure 1 - ECS DC Deployment Model Examples

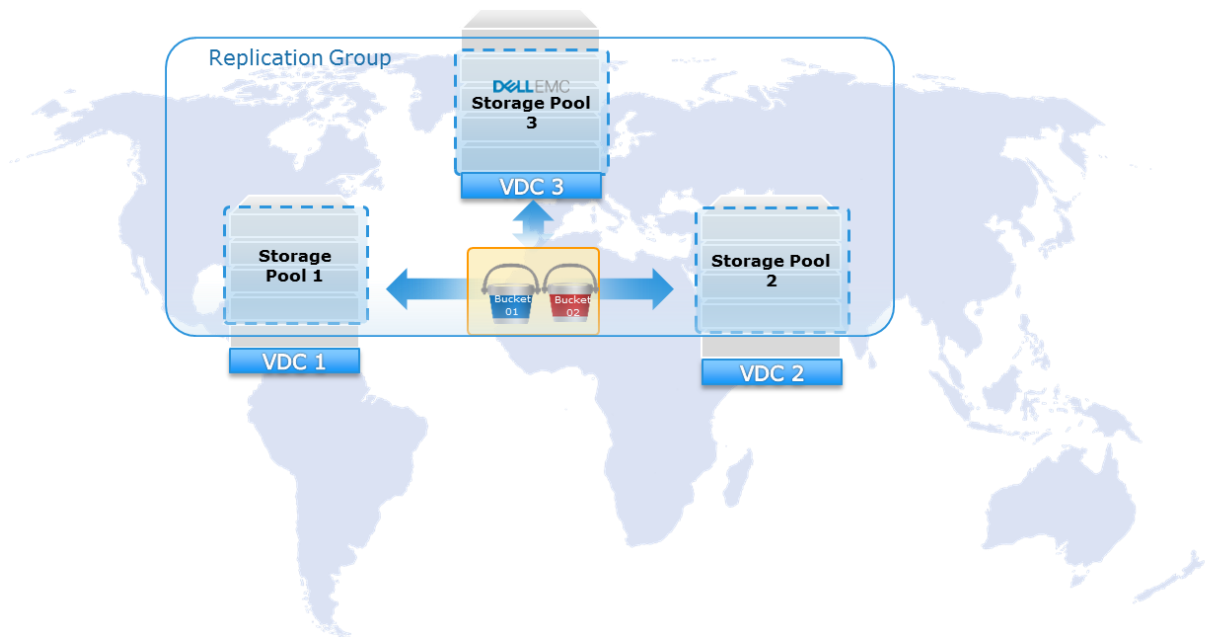


3.1 Geo-Replication

Geo-replication provides enhanced protection against site failures by having multiple copies of the data, i.e., a primary copy of the data at the original site and a secondary copy of the data at remote site. Replication to other site is an asynchronous process and replication is done by site owning the object. Data are added to a queue to be sent over. Multiple worker I/O threads continuously processes the queues until empty. Data is first encrypted (AES256 with Cipher Block Chaining) and then sent to another site via HTTP. Site who receives the replicated data will be responsible for local data protection (erasure coding and/or triple-mirror).

Let's review some of the ECS terminology. A replication group (RG) acts as the logical container that defines what namespace maps to what physical resources, i.e. storage pool(s), to store data. A multi-site RG defines a set of storage pools within which the primary and the secondary copies of data reside. The primary data copy resides at the site that the create operation originated and the other storage pools jointly store a single secondary data copy. For example, in a three or more site geo-replicated deployment, replication is done to only one other site. Figure 2 shows an overview of the mapping from logical model to physical infrastructure, or storage pool. Customer applications access a bucket that belongs to a namespace. One or more namespaces are defined within a replication group that is a logical grouping of storage pools from different virtual data centers (VDC) or sites. In this figure, VDC 1 and VDC 2 are ECS systems on customer premises and VDC 3 is off-premises and managed by Dell EMC.

Figure 2 - Example of Geo-Replicated Deployment (3-Sites) with VDC 3 Managed by Dell EMC



3.1.1 Geo-Data Access (Active-Active)

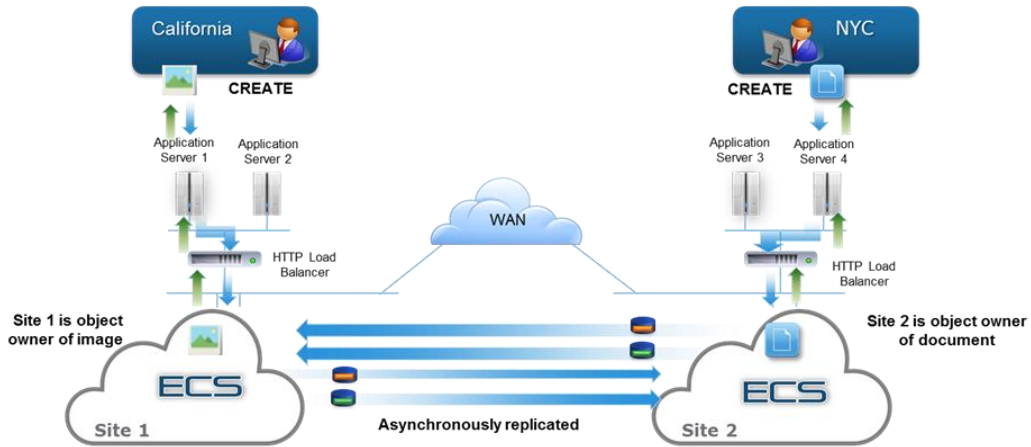
One key feature of ECS is the ability to read and write data from any site within a replication group. In ECS by default, data is replicated asynchronously to one other site within a replication group. The challenge this poses is consistency of data across sites. ECS ensures strong consistency by fetching the metadata from the site that first created the object data. Thus, if an object is created in site 1 and is read from site 2, ECS checks with site 1, who is the object owner, and validates that the copy replicated at site 2 is the latest version of the data. If not, it fetches the data from site 1; otherwise, it uses the data from site 2. Although the following sections shows only two sites, the data path for create, reads and updates in a geo-replicated environment is the same regardless of whether you have 2 or “n” number of sites.

3.1.1.1 Create

The site where the first create of the object is processed becomes the owner of object and in a geo-replicated environment is where the meta-data is checked for consistency during read operations and updated synchronously during an update of object. Figure 3 shows the flow of a create operation for two different

objects in a geo-replicated environment where an image is created in Site 1 and a document is created in Site 2. Data originated on each site is asynchronously replicated to other site. The site where object was created becomes site owner of object.

Figure 3 - Create in a Geo-Replicated Environment

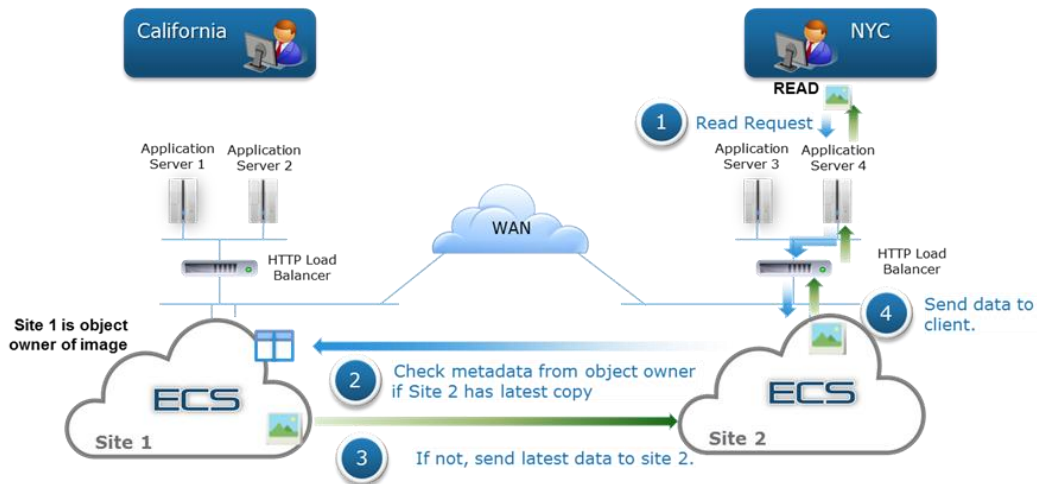


3.1.1.2 Read

The data flow of reads in a geo-replicated environment illustrated in Figure 4 below is as follows:

1. Site 2 does a read request. In this example, Site 1 is the object owner. The object owner is the site who first created object.
2. Read requires checking of the metadata from object owner to validate if Site 2 has the latest copy of data.
3. If not, Site 1 will send the data to Site 2
4. Data is sent to client.

Figure 4 - Read Data Flow in Geo-Replicated Environment



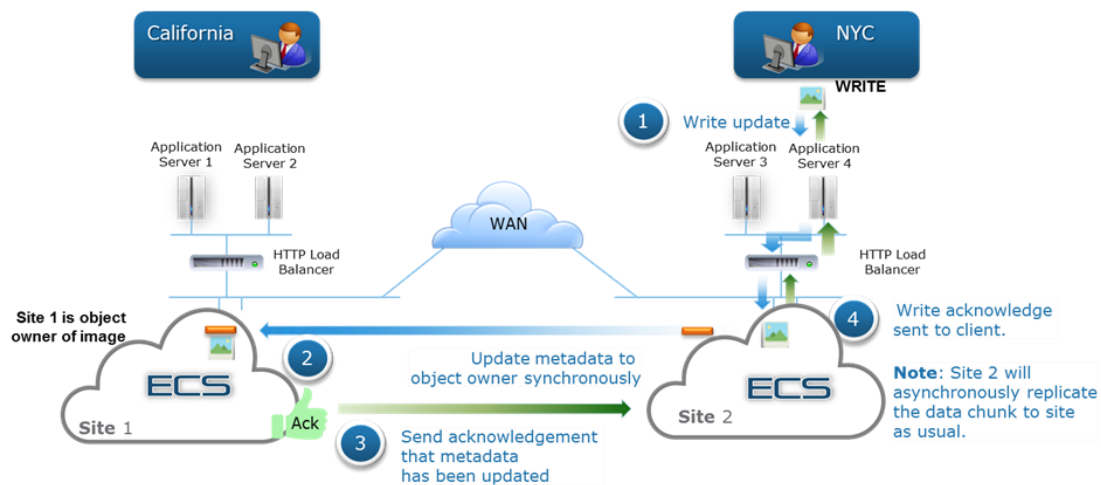
3.1.1.3 Update

The data flow of writes in a geo-replicated environment in which two sites are updating the same object is shown in Figure 5 and described below. In this example, Site 1 created object first and becomes object owner. The data is mirrored or erasure coded and journal is written as usual to local site.

1. Site 2 updates the same file. Site 2, first writes the data locally (data mirrored or erasure coded locally)
2. Site 2 synchronously updates the metadata (journal write) with the object owner, Site 1, and waits for acknowledgement of metadata update from Site 1.
3. Site 1 acknowledges the metadata write to Site 2.
4. Site 2 acknowledges the write to client.

Note: Site 2 asynchronously replicates the data to site 1 (owning site) as usual. If the data is not yet replicated to owning site and owning site wants to read the data, then it will get it from the remote site.

Figure 5 - Update of Same Object Data Flow in Geo-Replicated Environment



In both read and update scenarios in a geo-replicated environment, there is latency involved in reading or updating the metadata from object owner and retrieving data from object owner during a request.


3.1.1.4 Geo-Caching

Customers with multi-site access patterns could experience slow performance if data is always fetched from primary site, where the data was originally written. Consider a geo-replicated environment with Sites 1, 2 and 3. An object, Object A, is written to Site 1 and the secondary copy of the object resides at Site 2. In this scenario, a read for the object at Site 3 needs to fetch the object data from either Site 1 or Site 2 to honor the read. This leads to elongated response times, especially in environments with multi-site access patterns. ECS alleviates the response time impact by using some pre-designated percentage of disk space to cache objects that do not exist locally. For frequently accessed objects this reduces the response time after the initial copy of the object is cached locally. The cache is a Least Recently Used (LRU) implementation and the cache size is adjusted when nodes/disks are added to the storage pool.

3.1.1.5 GEO-XOR in 3 or More Sites

ECS implemented a mechanism in which storage efficiency increases as three or more sites are deployed within a replication group. In a geo-replicated setup with multiple sites, ECS replicates chunks from one site to a remote site for high availability. However, this simple replication can lead to an overhead of disk space. To alleviate this, ECS uses an innovative technique to reduce overhead while preserving high availability features. Consider 3 sites, as pictured in Figure 6, in a multi-site deployment – Site 1, Site 2 and Site 3 (ECS DC). Site 1 has chunk A and Site 2 has chunk B. With simple replication, a secondary copy of chunk A and a secondary copy of chunk B may be placed in Site 3. Since all chunks are of the same size, this will result in a total of 4 x 128MB of space being used to store 2 x 128MB of objects. In this situation ECS can perform an XOR operation of chunk A and chunk B (mathematically, written as $A \oplus B$) to create chunk C in Site 3 and individual secondary copies of chunks A and B are removed. Thus, rather than using 2 x 128MB of space in Site 3, ECS now uses only 128MB. The XOR operation results in a new chunk C of the same size.

Figure 6 - Storage Efficiency Using Mathematical XOR



When one of the sites is inaccessible, for instance, if Site 1 goes down or network connection to Site 1 is broken, ECS can reconstruct chunk A by using chunk B from Site 2 and chunk C ($A \oplus B$) from Site 3. Similarly, if Site 2 goes down, ECS can reconstruct chunk B by using chunk A from Site 1 and chunk C ($A \oplus B$) from Site 3. So, an un-XOR operation is conducted as follows to reconstruct the chunk that is inaccessible: $(C \oplus B=A)$ or $(C \oplus A=B)$.

As the number of linked sites increase, the ECS algorithm is more efficient in reducing the overhead. Table 1 presents information on the storage overhead based on the number of sites for normal erasure coding of 12+4 and cold archive erasure coding of 10+2, and it illustrates how ECS becomes more storage efficient as more sites are linked.

Note: With ECS DC, the option to replicate in all sites should not be enabled since it disables the XOR capability for storage efficiency just described.

Table 1 – Storage Overhead

Number of Sites in Replication Group	Default Use Case (Erasure Code: 12+4)	Cold Archive Use Case (Erasure Code: 10+2)
1	1.33	1.2
2	2.67	2.4
3	2.00	1.8
4	1.77	1.6
5	1.67	1.5
6	1.60	1.44
7	1.55	1.40
8	1.52	1.37

To obtain the lower overhead it is important that the same amount of data be written at each site. Thus, the use of load balancers is essential to the ECS DC deployment, especially in three or more sites deployment, to handle the even distribution across all sites. The next section describes how this is done.

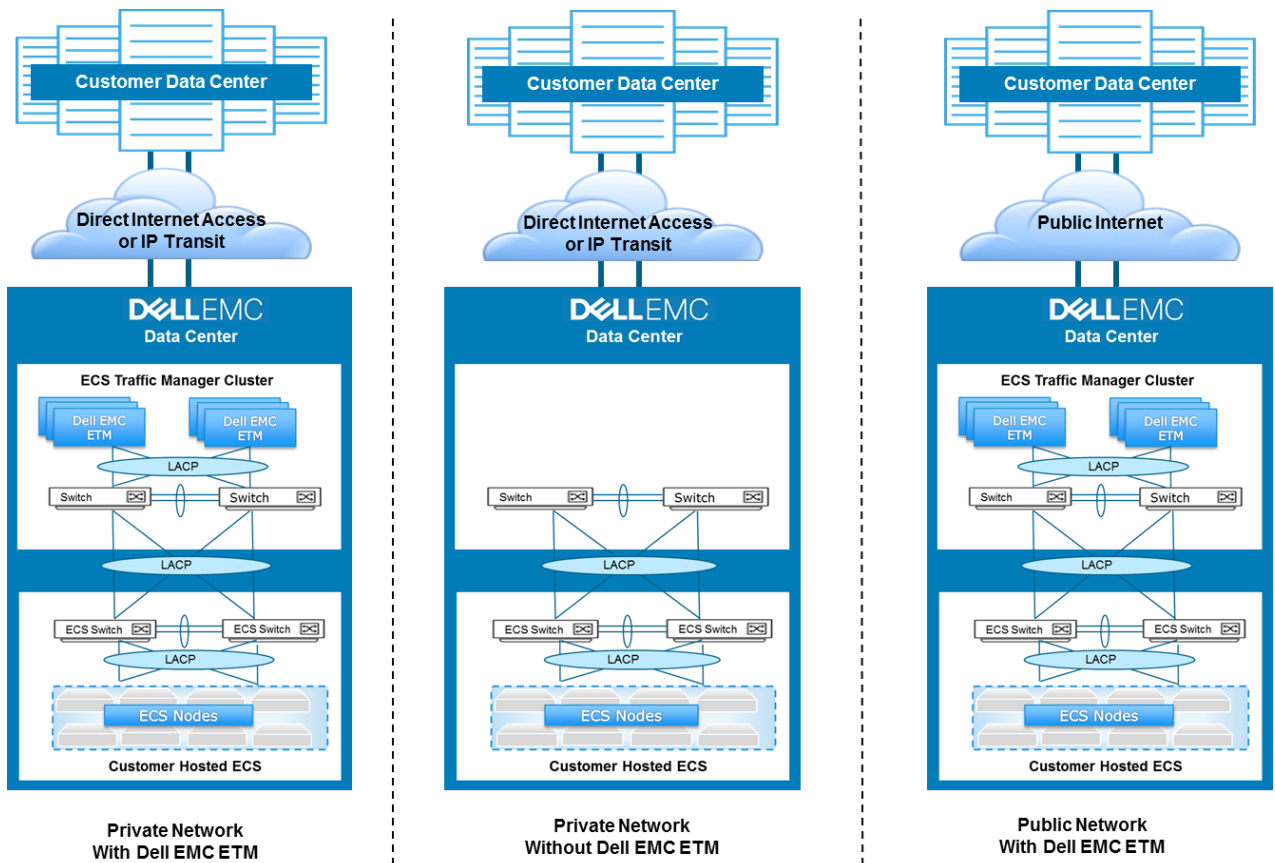
4 ECS DC Deployment

Deployment of ECS DC within Dell EMC data center involves the following:

- **Dell EMC ECS Traffic Manager Cluster (ETM)** – consists of load balancers, traffic shaping, traffic metering, and proxying, firewall, and Dell EMC Secure Remote Services.
- **Network connectivity** – several options are available to customers depending on type of network connection required which includes dedicated, public, or a combination of dedicated and public.

There are several different combinations of ECS DC deployment that can be configured. Depending on the customer requirements, deployments can differ based upon the model of ECS, network connectivity, or inclusion of Dell EMC ECS ETM and other services desired at the hosted site. For instance, three examples deployments shown in Figure 7 illustrate a private customer provided network connection such as direct internet access (DIA) or IP transit with or without the Dell EMC ETM cluster or public internet connection with Dell EMC ETM cluster. Within the data center managed by Dell EMC, there are 10 GbE redundant switches to route traffic between the ETM clusters, Dell EMC managed ECS and customer data center. These switches have Layer 2 to support Link Aggregation Control Protocol (LACP) and Layer 3 to support smart routing such as Border Gateway Protocol (BGP). More details relating to the network connectivity are discussed in a later section.

Figure 7 - Examples of ECS DC Deployment Models within Dell EMC Data Center



4.1 ECS Traffic Manager (ETM)

Each customer gets a minimum of two ETM instances running on diverse hardware in active/passive mode and connected to redundant 10 GbE switches within the data center. There are several components within Dell EMC ETM:

- **Load Balancers** – both global and site load balancers to distribute traffic equally across sites and ECS nodes within a VDC.
- **Traffic Management** – bandwidth throttling or shaping, traffic metering, and additional proxying for routing of different traffic types on different types of network connection.
- **Firewall** – only opens up ports to handle ECS protocols, replication traffic, ESRS and ports required by load balancers.
- **Dell EMC Secure Remote Services (ESRS)** – monitoring, alerting, and troubleshooting of ECS systems at Virtustream data center.

The load balancers and ESRS components are packaged in different Docker (version 1.12 or higher) containers on Linux hosts running Ubuntu version 16.04.1 LTS or higher. Depending on current infrastructure or customer requirements, some of these components can be on same server or different servers. They can be on bare metal or virtual machine. System requirements for these components are highly dependent on incoming connections and request rate and should be sized appropriately based on customers expected workload. Additional Docker containers can be added to support higher load.

4.1.1 Load Balancers

ECS DC deployment with two or more other sites has been carefully architected to maximize storage efficiency and optimize read and write performance. To accomplish this, load balancers are incorporated into the deployment for on-premise ECS and ECS hosted on Virtustream. The goal of the load balancers is to evenly distribute user data across sites for reduced storage footprint and cost. It incorporates a specialized algorithm referred to as “geo-pinning” to optimize reads in an ECS geo-replicated environment and spread creates across sites. Furthermore, it allows access to data from an alternate location in the event of failure. There are two types of load balancers implemented for an ECS hybrid solution:

- **Global Server Load Balancer (GSLB)** – receives requests from clients and distributes requests or workload across SLBs at each site.
- **Site Load Balancer (SLB)** – receives request from GLSB and distributes the request across ECS nodes within the VDC.

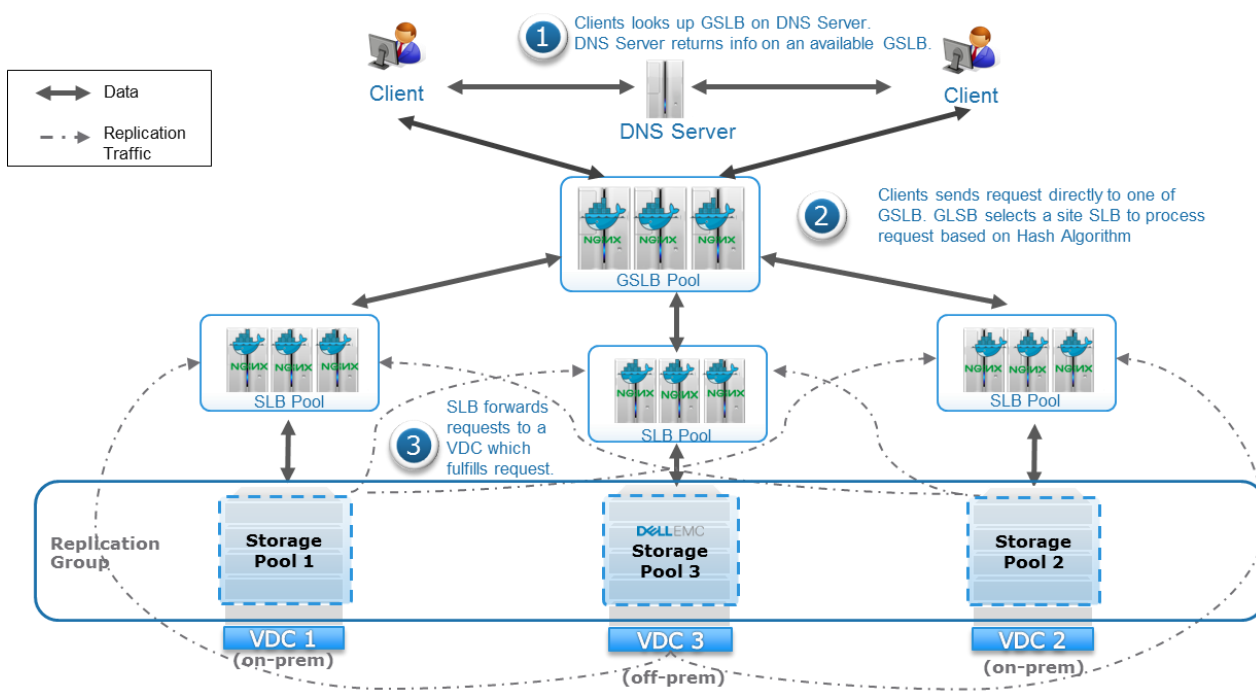
A variant of NGINX is the load balancer utilized for the GSLB and SLB. NGINX is an open source high performance HTTP server software with extended capabilities which includes load balancing, web serving, caching and reverse proxying. The variant of NGINX implemented for ECS DC deployment is OpenResty® (<https://openresty.org>) which integrates NGINX core, LuaJIT (Just-In-Time Lua Compiler, Lua libraries, and other external libraries. It extends the core functionality of NGINX core via the Lua (lightweight scripting language) modules and other third party NGINX modules to provide developers a mechanism to customize a web server best suited for their needs. Lua scripts have been developed to optimize the ECS reads and writes in a geo-replicated environment.

NGINX for ECS DC and on-premise sites are implemented within Docker containers on Linux hosts. Minimum system requirements for server to run load balancer Docker images include: 1 CPU core, 1GB of memory, 10 GB of local disk and running Ubuntu version 16.04.1 or higher.

Note: Depending on version and edition of Ubuntu installed, minimum system requirements may differ.

There would be a pool of SLBs on each VDC site and a pool of GSLBs at the customer's on-premise site. Domain Name Server (DNS) is used by client to do a lookup of the GSLB available in the pool to use. Figure 8 provides a high level view of how the GSLB and SLB are implemented into the deployment of ECS DC with 2 on-premise sites. Data traffic flows between GSLB, SLB and ECS sites and replication traffic flows between ECS and SLB sites.

Figure 8 – Example of ECS DC Deployment (3 Sites) with VDC 3 Managed by Dell EMC



4.1.1.1 Global Server Load Balancer (GSLB)

The GSLB for ECS DC solution is a Docker image containing NGINX, libraries and binaries deployed on Linux hosts preferably on customer's on-premise site. A minimum of two GSLB are recommended for high availability. The features of the GSLB include:

- **Geo-Pinning** - evenly distribute object data by selecting an ECS VDC site based on bucket name and/or object key name.
- **Failure Detection and Auto Failover** – does periodic health checks on SLB and automatically failover when failure is detected.
- **Monitoring** – provides information on state of ECS nodes and some metrics relating to requests processed by GSLB.

DNS is used by clients to do a lookup of available GSLB to send requests to. When one of the GSLB in the pool fails, DNS return information on an up and available GSLB to handle requests. DNS will return GSLB information in a round-robin fashion, by default, among all GSLB within the pool

Geo-Pinning

The GSLB implements a hash algorithm (modulo to “n” sites) referred to as “geo-pinning” to distribute the creation of objects across sites for storage efficiency and direct reads and updates to site owning object to enhance performance. As previously described, the site that originated the creation of object is the site owner of object. In order to maintain strong consistency a read request first checks with site owning object for consistency of data at requesting site. If a request goes to site that does not own the object, response times are prolonged in order to check with owning site first and fetch data if data is not currently cached at remote site. With geo-pinning, the read goes directly to site owner of object to process requests, resulting in improved performance.

[Lua](#), a lightweight scripting language, is embedded within the load balancer configuration file to implement the geo-pinning algorithm. The geo-pinning logic is as follows:

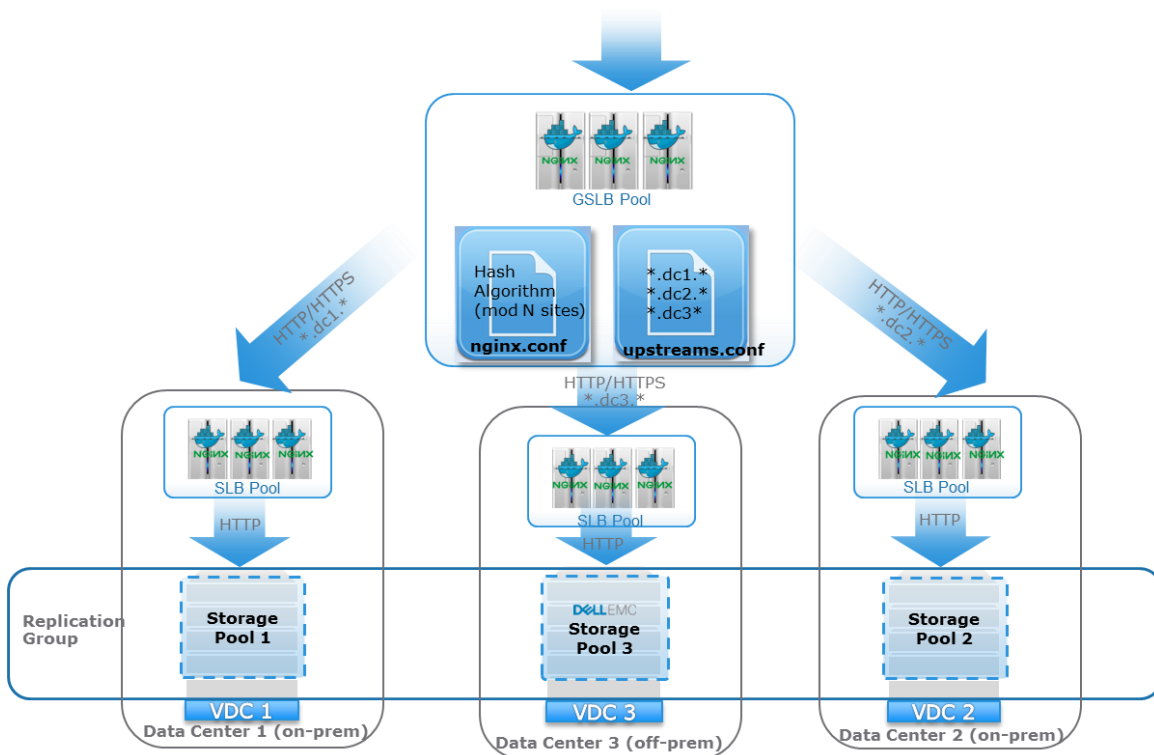
1. Derive a hash ID from the URI and headers of a request. If it is a bucket request, the bucket is the ID. If it is an object request, the key is the ID.
2. Calculate the hash string from the hash ID using SHA1 algorithm.
3. Get first six hexadecimal number of hash string and conduct a modulo operation based on the number of VDCs.

Configuration Files

The NGINX upstream configuration file contains the IPs or DNS names with ECS supported protocol ports (described in Firewall section) for pool of SLB at each site. The NGINX configuration file contains the “geo-pinning” algorithm and contains a Lua script to periodically perform health checks on SLBs at each site for failure detection and auto failover. There is also a directive in the configuration files to handle HTTPS. When a HTTPS request is received by the GSLB, it first decrypts the packet using the SSL certificate. The certificate can be self-signed certificate or signed by a Certificate Authority. As a best practice, using a Certificate Authority to sign certificate is preferred over issuing self-signed certificates. After the HTTP headers are examined and geo-pinning logic is applied, the packet is encrypted again and sent to one of the SLBs specified in the upstream configuration file in a secure channel. Figure 9 illustrates a high level view of

GSLB and how traffic is directed. The response will traverse the same path that the request comes in. For instance, request packet is sent to GSLB 1, then to SLB 1 in data center 1 and then sent to ECS rack 1. Response will be returned from ECS rack 1 to SLB 1 to GSLB 1 and then finally to client. In addition, directives are defined in the configuration file to provide information on state of ECS and type of traffic flowing into load balancer.

Figure 9 – GSLB Deployment Example with VDC 3 Managed by Dell EMC



4.1.1.2 Site Load Balancer (SLB)

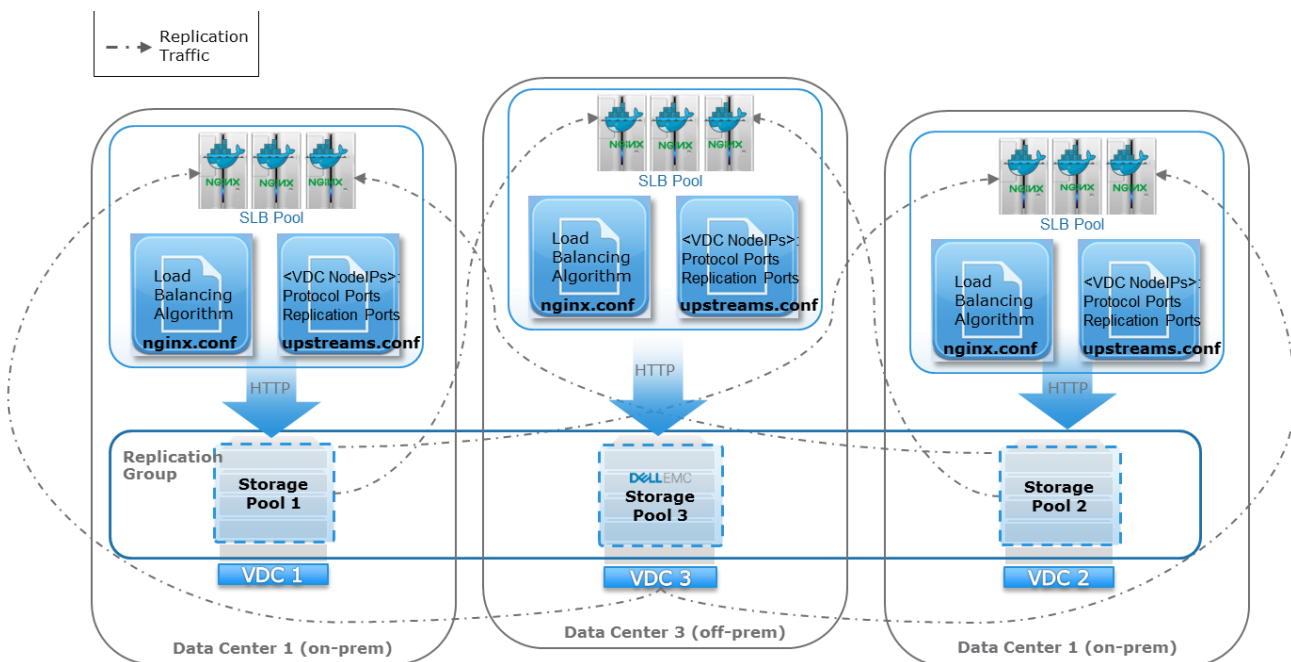
There is a pool of SLBs (minimum 2 SLBs) at each VDC site and goal of SLBs is to receive traffic from GSLB and replication traffic from the other VDC(s) and forward the requests to nodes within the VDC to process as shown in Figure 10. The SLB is also a Docker image containing NGINX, libraries, and binaries deployed on Linux hosts. The SLB features include:

- **Failure Detection and Auto Failover** – conducts periodic health checks on ECS node services and automatically failover when failure is detected
- **SSL Offloading** – handles the encrypting and decrypting traffic sent via SSL
- **Monitoring** - provides some information on state of ECS nodes in addition to metrics relating to requests being processed by SLB.

The upstream configuration file contain all the IP addresses of nodes within the VDC as well as the ports (described in Firewall section) for each of the ECS supported protocols and replication ports. The NGINX configuration file defines the directives on which ports to listen on and where to forward request to. It also performs periodic health checks of the ECS services on the nodes so that traffic is sent to nodes that are up

and available. The health checking is done using a layer 7 ping (i.e. based on the “GET /?ping” S3 API extension). The SLB implements the round-robin load balancing algorithm by default, however, customers may opt to change this based on their requirements. For HTTPS, the SLB are used to offload the SSL traffic and terminated at the SLB in order not to add extra load on ECS nodes to establish SSL sessions. Thus, the certificate is installed and evaluated on SLB and requests are forwarded from SLB to nodes on VDC using non-SSL or HTTP. Similar to GSLB, SLB has directives to provide state of ECS nodes and requests coming in and out of SLB.

Figure 10 - SLB Deployment Examples with VDC 3 Managed by Dell EMC



4.1.1.3 Traffic Management

In some scenarios, throttling bandwidth maybe required to control the amount of data traffic and prevent overloading the system. Linux has a set of tools for managing and controlling the transmission of packets. Linux “tc” (traffic control) and iptables utility can be used to control and shape traffic for specified ports, destination IP and source IPs. These tools are implemented on host running the ETM. Traffic shaping is achieved by applying bandwidth definition on ports and is applied on outgoing or forwarded traffic. Metering of traffic is also available to determine the amount of traffic flowing in and out of the hosted site.

Customers may desire hybrid network connections such as utilizing a dedicated leased line and public internet for the various ECS traffic types (data, replication, or management). In this scenario, additional proxying or proxy servers are utilized to facilitate this configuration.

4.1.1.4 Firewall

On the hosts running GSLB(s) and/or SLB(s), firewalls are set to open up ports related to processing ECS related operations. These ports include the ECS protocol ports supported, replication ports and ESRS ports. Table 2 shows example of ports that are opened. Depending on which ECS protocols are used, the

appropriate ports will be opened. Refer to the [ECS Security Configuration Guide](#) for more detailed information on ECS ports.

Table 2 - Examples of Ports Opened

Type	Transport Protocol or Daemon Service	Port
S3	HTTP	9020
	HTTPS	9021
Atmos	HTTP	9022
	HTTPS	9023
Swift	HTTP	9024
	HTTPS	9025
Replication and ECS nodes to ECS nodes in other sites usage ports. Only on SLB	Geo-Receiver (HTTP)	9094
	Geo Receiver (HTTPS)	9095
	Data (HTTP)	9096
	Data (HTTPS)	9097
	Geo Receiver Communication Manager	9098
ESRS Ports (where ESRS is deployed)	ECS UI https	443
	ECS Management API	9443
	Connect Home support	21, 25

4.1.1.5 Dell EMC Secure Remote Services (ESRS)

Dell EMC operations team will be responsible for monitoring the health of the ECS DC hosted on Virtustream as well as the load balancers (SLB at hosted site). A Docker edition of ESRS is deployed at the Virtustream site for monitoring, alerting, and remote troubleshooting of ECS DC. ESRS provides a secure two-way connection between customers owned Dell EMC equipment and customer support team. This helps to accelerates problem resolution with proactive remote monitoring and repair. The configuration files and notification delivered via ESRS is useful in providing information to engineering and development teams on systems installed at customers' sites.

For information on ESRS Docker Edition requirements, refer to the [ESRS Installation and Operations Guide](#).

Note: The ECS Portal is utilized for administration, traffic monitoring, and management of ECS in all sites. Customers are given administrative permissions to the graphical user interface (portal) of ECS DC.

5 Network Connectivity

ECS DC offers several types of network connections. Current network options offered with ECS DC include:

- **Internet line (Public)** – connection with internet speeds (1 Gb/s)
- **Direct Leased Line (Private)** – dedicated network connection offered by network service provider.
- **Hybrid (Combination of Internet and Direct Leased Line)** – specific ECS traffic types are transmitted to either use Internet and/or Direct Leased Line. For instance, replication traffic uses the direct leased lines and management and normal data traffic uses the internet line.

Figure 11 illustrates an example of direct leased line network connectivity and how the network is configured within the Dell EMC data center. From this example the following can be observed:

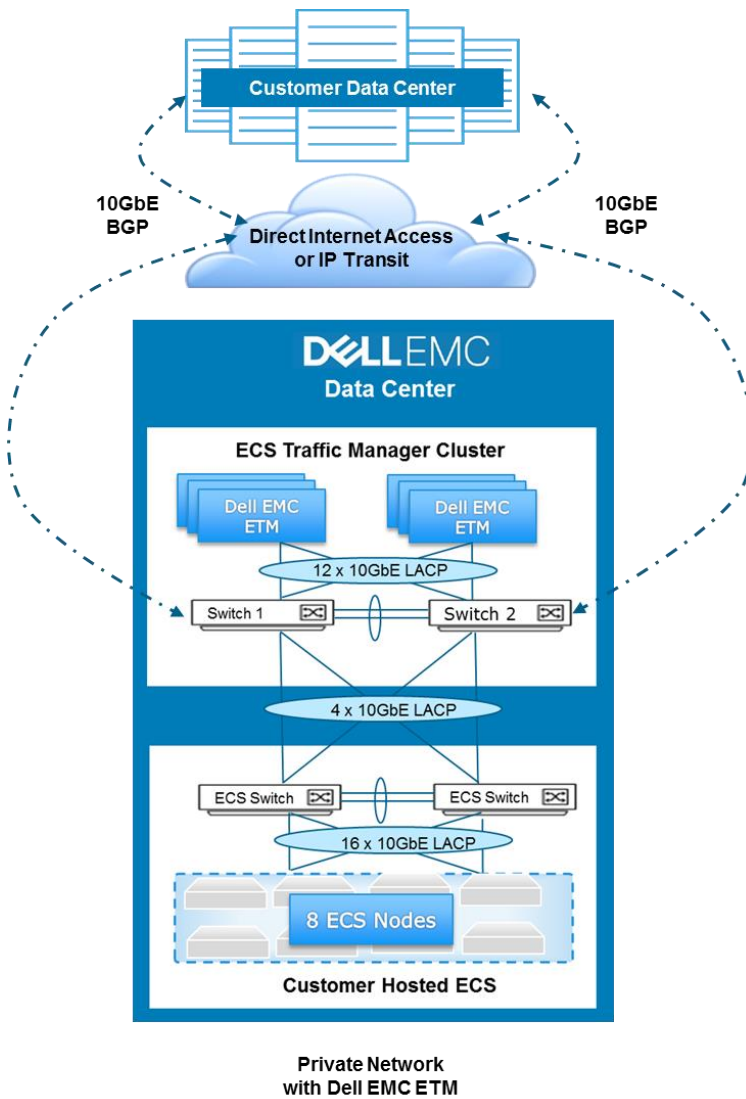
There are two 10 GbE switches connecting the customer data center, ETM cluster, and ECS switches pair. There are two network connections per switch connecting to each of the ECS switches configured using Link Aggregation Control Protocol (LACP) – 4 x 10GbE LACP. These switches are also configured with Border Gateway Protocol (BGP) allowing for intelligent routing and route filtering between the customer data center switches and Dell EMC data center switches.

There are six ETM with each ETM having redundant connections configured using LACP - 12 x 10GbE LACP.

There are eight ECS nodes with each node having 2 NIC interfaces bonded and connected to each ECS switch 10 GbE pair – 16 x 10 GbE LACP.

Since BGP is utilized in this example, it is expected that customer switches within their on-premise data center also needs to have support for BGP.

Figure 11- Example of Direct Leased Line Network Deployment



Network bandwidth or performance is affected by the network connectivity implemented. Customers can decide the best network connectivity solution to meet their needs. As previously mentioned, additional control of the bandwidth is an option for customers who have requirements to throttle network bandwidth for purposes of limiting certain types of ECS traffic at each site or controlling the amount of bandwidth load balancers use within a system. The bandwidth throttling is implemented on hosts running the ETM.

6 Best Practices

When deploying ECS DC with other ECS on-premise sites, there are some best practices to highlight:

- Use 2 or more Load Balancers (SLB and GSLB) implemented for redundancy and to prevent single point of failure.
- The latency between sites should not exceed 1000 ms.
- Use the SLB to terminate SSL connections to reduce the load on the ECS Nodes. If SSL termination is required on ECS nodes themselves, use Layer 4 (tcp) to pass through the SSL traffic to ECS nodes for handling. The certificates need to be installed on the ECS nodes and not on the load balancer.
- Use SSL certificates from a Certificate Authority instead of self-signed certificates for improved security.
- Setup a mechanism or define directives within the configuration files of NGINX to self-monitor the GSLB and/or SLB at customer's on-premise sites.
- On-premise ECS sites are recommended to deploy ESRS or register ECS onto their ESRS gateway. ESRS offers a better customer experience by streamlining the identification, troubleshooting and resolution of customer issues.

7 Support

Initial deployment of the ECS DC and components (i.e. load balancers) deployed at customer on-premise sites are conducted in conjunction with Dell EMC Operations personnel. Customer requirements are gathered and solution is deployed based on the specified requirements. Setup and install of ECS DC and ETM at the Virtustream data center is handled by Dell EMC personnel. Docker images for the load balancers that are run on customer's on-premise sites are provided and Dell EMC personnel will assist in configuration.

Dell EMC customer support team together with a specialized operations team is responsible for proactively supporting and monitoring the ECS DC and ETM components such as load balancers and firewall hosted on the Virtustream data center. Technical issues observed with the components at Virtustream or ECS alerts are automatically triaged to the Global Corporate support team and a service request will be filed and will be resolved or escalated within Dell EMC support or appropriate operations team to address. ECS, SLB, and GSLB on customer's on-premise site are the responsibility of customer to work with Dell EMC customer support to report and address issues based on their support level agreement.

8 Conclusion

ECS DC extends customer's private cloud environment to the public cloud by further enhancing storage utilization, global access, and protection. It also reduces cost and complexity in investing and operating an additional ECS within a customer's data center. Cloud computing companies such as Virtustream provides numerous locations that customers can select from to deploy ECS DC in addition to their other cloud services. As discussed in this whitepaper, deployment of ECS DC requires other components such as load balancers, firewall and ESRS off-premise and/or on customer's sites. ECS DC in conjunction with the customer's on-premise ECS presents a hybrid solution that encompasses the best of both private and public cloud offerings.

9 References

- ECS product documentation at support site or the community links:
 - https://support.emc.com/products/37254_ECS-Appliance-/Documentation/
 - <https://community.emc.com/docs/DOC-53956>
- ECS Architecture and Overview
 - <http://www.emc.com/collateral/white-papers/h14071-ecs-architectural-guide-wp.pdf>
- ECS High Availability Design
 - <http://www.emc.com/collateral/whitepaper/h16344-elastic-cloud-storage-ha-design.pdf>
- ECS Networking and Best Practices
 - <http://www.emc.com/collateral/white-paper/h15718-ecs-networking-bp-wp.pdf>
- ECS with NGINX (OpenResty) Deployment Reference Guide
 - <http://www.emc.com/collateral/white-papers/ecs-with-nginx-deployment-reference-guide.pdf>