


# Сетевые системы хранения данных DELL EMC – большие данные и масштабируемые системы

Михаил Владимиров  
Технический  
Консультант

[ISI.ECS@emc.com](mailto:ISI.ECS@emc.com)

# Откуда движется волна?

- Новые источники данных



Internet of Things



Медиа



WEB 3.0



Видеонаблюдение



Темные данные



Умные датчики





Исследования в нефте-  
газовом секторе




Медицинские  
изображения

# Две платформы



Стабильные и надёжные платформы для традиционных задач



Гибкая платформа для мгновенных изменений задач и требований



# ISILON + ECS = от 20ТБ до ∞



## УНИКАЛЬНОЕ РЕШЕНИЕ

для работы с неструктурированными данными

# ECS – “Облачная” СХД из коробки

Traditional/“Platform 2”

Cloud Native/“Platform 3”

Tiered  
Archive



Cloud Backup



Sync  
& Share



Cloud  
Gateway



Cloud Native Apps  
(web/mobile)



IoT



Analytics



SITE 1



SITE 2



SITE 3



*Scale Effortlessly - Store Efficiently - Access Globally*

# ECS в ногу со временем!!!

Решение архивных задач на уровне Geo-scale



- Modern Hyper-scale Cloud Architecture
  - Scales from Petabytes to Exabytes
- Archive data from all sources
  - S3, Swift, Atmos and Centera CAS Object APIs
  - HDFS compatible with Cloudera, Hortonworks, Pivotal etc.
- Break down barriers with Geo-scale data access
- Innovation to enable scalability, efficiency and serviceability!



# EMC Elastic Cloud Storage (ECS)



## U-Series

Available in multiple capacities within a rack  
x86 servers, JBOD DAS, 10GbE connectivity,  
SATA/SAS Disks  
Max 60 disks per DAE per node

## D-Series

Denser Model  
Minimum of eight x86 servers  
Max 98 disks per DAE per node.  
Hyper-scale  
Aggressive seeking lowest \$/GB

# ECS Appliance – Аппаратные составляющие



## Data Network - Dual 10 GBE

1

- Arista Switches
- Each ECS node connects to these switches via dual 10GigE
- Switches connect to customer network (up to 8 uplinks per switch)

## Internal Management Network

2

- Single Gigabit Switch
- Optional uplink to customer network for network based server management

## Servers (ECS nodes)

3

- Intel Dual Quad Core x86 Nodes
  - Dual 4-core IvyBridge (gen1)
  - Dual 6-core Haswell (gen 2)
- 64GB Memory
- 4 nodes in 2U chassis
- OS, HAL, Fabric & ViPR Data Services software package
- Each node connects to a Disk Enclosure via SAS

4

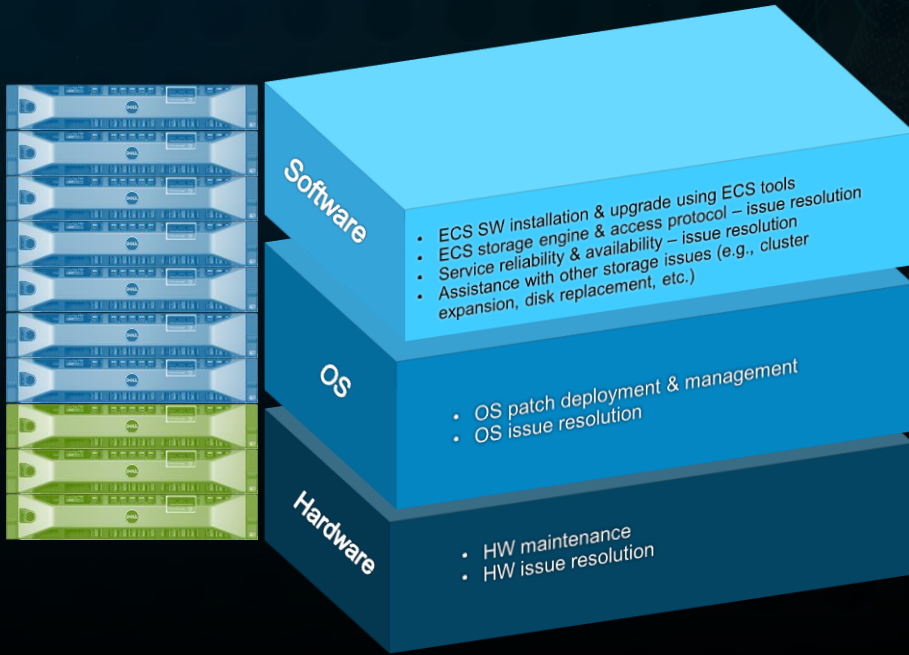
## Disks & Enclosures

- 6TB SATA (gen 1) or 8TB SAS Disks (Gen2, D-Series supports both 8TB and 10 TB )
- Up to 60 or 98 Disks / 4U Enclosure



# ECS Software на серверах x86

Customers EMC



Dell DSS7000

Dell R730xd 13G

Minimum 5 servers

# Объектный доступ

## Namespace / Bucket|Container

- **Namespace** can span multiple instances of physical hardware, and data management functions like data replication and data distribution at object-level granularity.
- Instead of organizing files in a directory hierarchy, object storage systems store files in a flat organization of containers/buckets.



### File storage

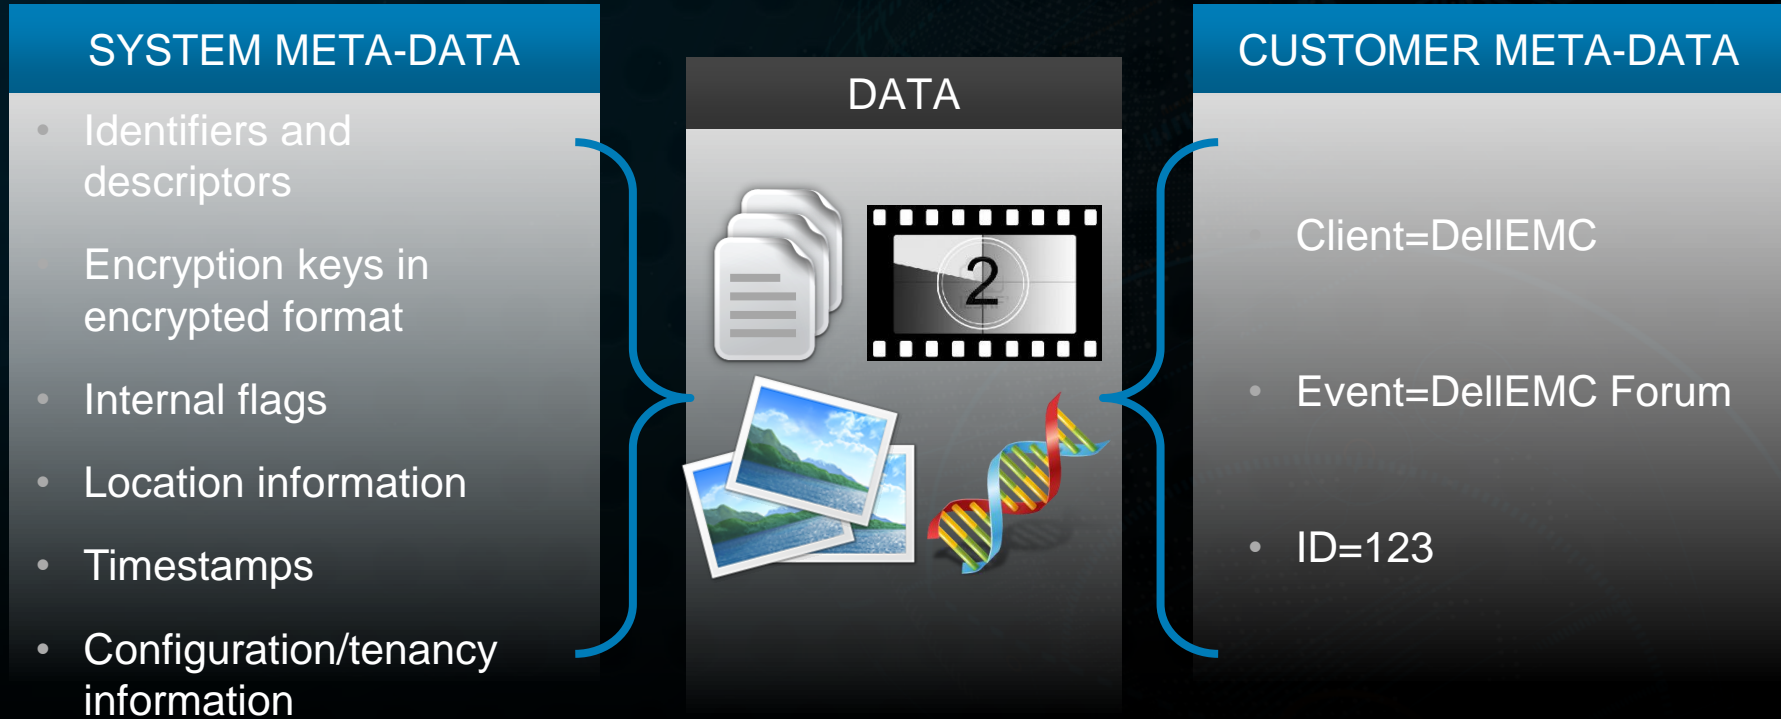
Data stored as 'files' in hierarchically nested 'folders'—ideal for active documents



### Object storage

Data stored as 'objects' in scalable 'buckets'—ideal for unstructured big data, analytics and archiving


# ECS – Как выглядят данные в объектной системе





# Протоколы доступа

- REST:
  - Combination of HTTP methods
  - Amazon S3
  - Openstack Swift
  - Atmos
  - Content Address Storage (CAS)



# Доступ к данным с помощью различных SDK

```
import fileinput
from boto.s3.connection import S3Connection, OrdinaryCallingFormat, Key
```

```
# Считываем параметры доступа из файла, строки добавляем в список
config_file='C:/Users/vladim/Desktop/asd.txt'
config_file_list=[]
with open(config_file) as f:
    for line in f:
        config_file_list.append(line)
```

```
# Устанавливаем соединение с ECS, присваиваем переменным значения списка
host = config_file_list[0].strip()
port = int(config_file_list[1].strip())
secure = port == int(config_file_list[1].strip())
access_key_id = config_file_list[2].strip()
secret_key = config_file_list[3].strip()
debug_level = int(config_file_list[4].strip())
calling_format = OrdinaryCallingFormat()
```

```
s3 = S3Connection(aws_access_key_id=access_key_id, aws_secret_access_key=secret_key, is_secure=secure, port=port,
                 host=host, debug=debug_level,
                 calling_format=calling_format)
```




# Частное облако



- **Lower cost** than public cloud
- Unmatched combination of **storage efficiency** and **data access**
- Anywhere read/write access with strong consistency **simplifies finding and using archived assets**
- No single points of failure increases **availability and performance**
- Universal accessibility **eliminates storage silos and inefficient archiving processes**
- Comprehensive data types satisfy the **broadest range of application needs**



# ECS “Storage Engine” уникальные ВОЗМОЖНОСТИ



# Одновременный много-протокольный доступ

- CAS: SDK v3.1.54



- Swift: byte range update within an object
- Retentions
- Keystone integration: drop-in replacement for OpenStack Swift



- Primary file system with native Ambari 2.2 integration



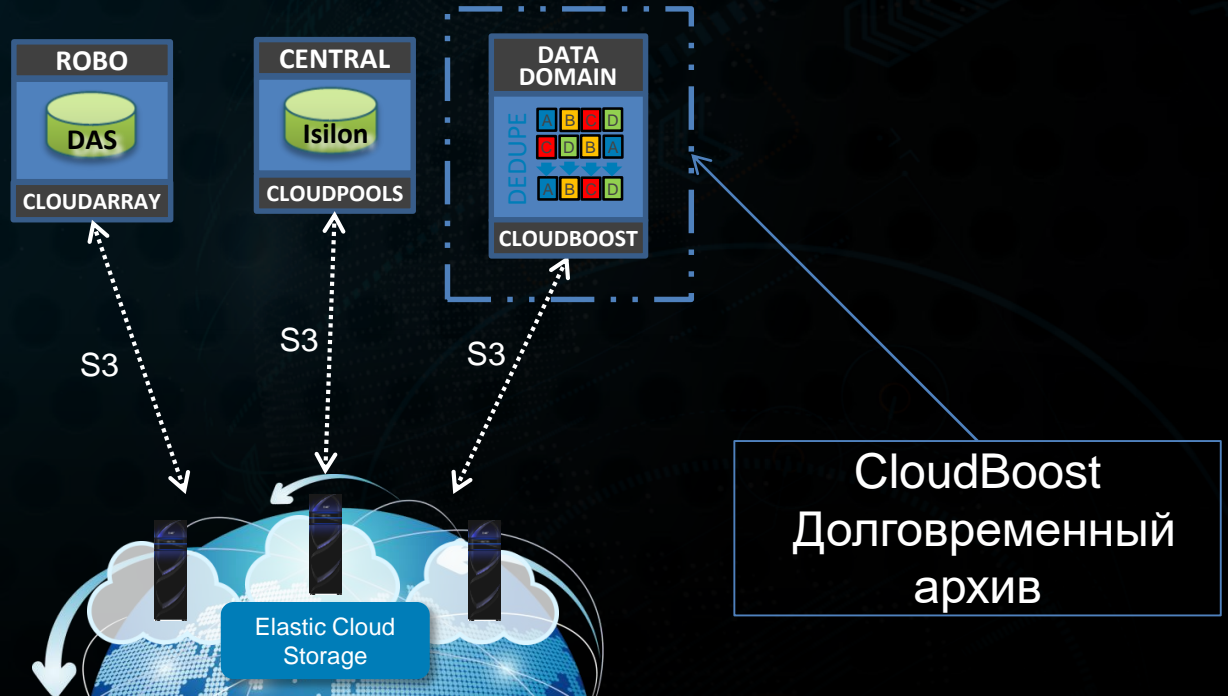
- Byte range updates
- Retentions
- Metadata search extension



- NFS v3
- Global namespace with global locking

# CloudBoost и CloudPools

БЕСШОВНЫЙ АРХИВ ДЛЯ DD, CLOUD ARRAY И ISILON





# Требования регуляторов

## Features

- ✓ Retention policy management
  - ✓ Retention enforcement
  - ✓ Data immutability
  - ✓ Advanced Retention Management
- ✓ Access Locks
  - ✓ Lock/unlock user or bucket

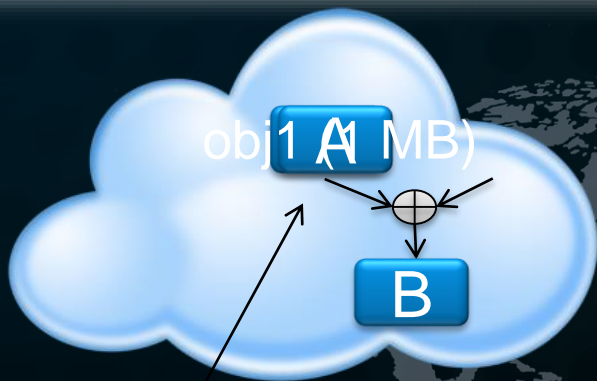
## Key Benefits

- ✓ Meet storage requirements for
  - ✓ SEC 17a-4(f)
  - ✓ CFTC 1.31(b)-(c)
  - ✓ NF Z 42-013



# Гео- распределённое хранение


Высокая надёжность с минимальными накладными расходами



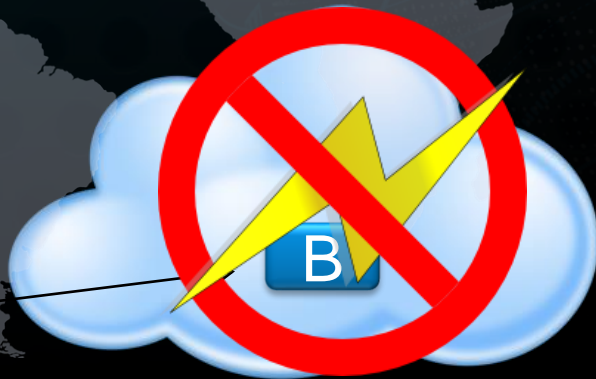
PUT obj1 (1 MB)

PUT obj2

Read the  
XOR construction  
the data from  
Site 1  
container



GET obj1

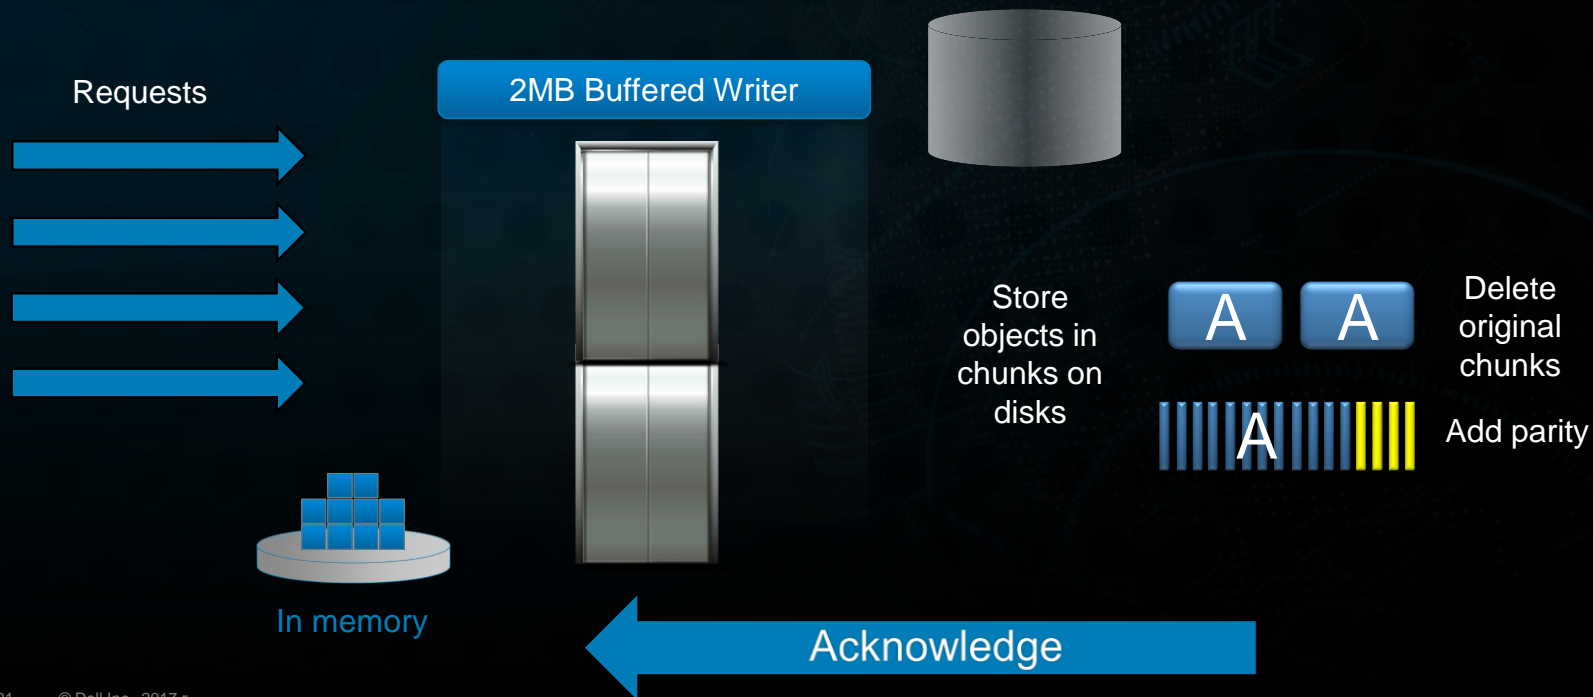


# Эффективное хранение

| Number of Sites in Replication Group | Erasure Code: 12+4<br>Default Use Case | Erasure Code: 10+2<br>Cold Archive Use Case |
|--------------------------------------|--|---|
| 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  |



# Эффективность хранения и производительность



# Поиск по метаданным

## Упрощение работы приложений

Uploading from client

File

Bucket Name

Retention to apply to protect the picture

File size

Image Width

Image Height


GPS Latitude

GPS Longitude

Datetime



Search objects using GPS coordinates, image resolution, ...  
`?query=x-amz-meta-image-gps-latitude>50&...`



Save objects with metadata

# Встроенная поддержка NFS



## Features

- ✓ **Native NFS v3 capability**
- ✓ Rich ACLs
- ✓ **Global namespace**
- ✓ Global locking
- ✓ **Multi-protocol access object, NFS and HDFS**

## Key Benefits

- ✓ Ingest data in native format
- ✓ Requires no change on the application level, accelerating the move to an object platform

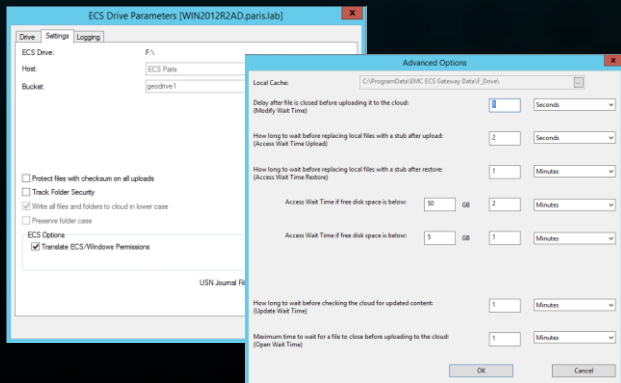

# Поддержка CIFS

## Features

- ✓ S3 API
- ✓ **Caching**
- ✓ Multipart upload and download
- ✓ **Retention & versioning**
- ✓ ACL translation
- ✓ **Client side load balancing**

## Key Benefits

- ✓ Ingest data in native format
- ✓ Requires no change on the application level, accelerating the move to an object platform





A photograph of a Siemens building at dusk. The building has a dark, curved upper section with the word "SIEMENS" in large, white, sans-serif capital letters. Below this is a lighter-colored section with a grid of windows, some of which are illuminated from within. The sky is a deep blue.

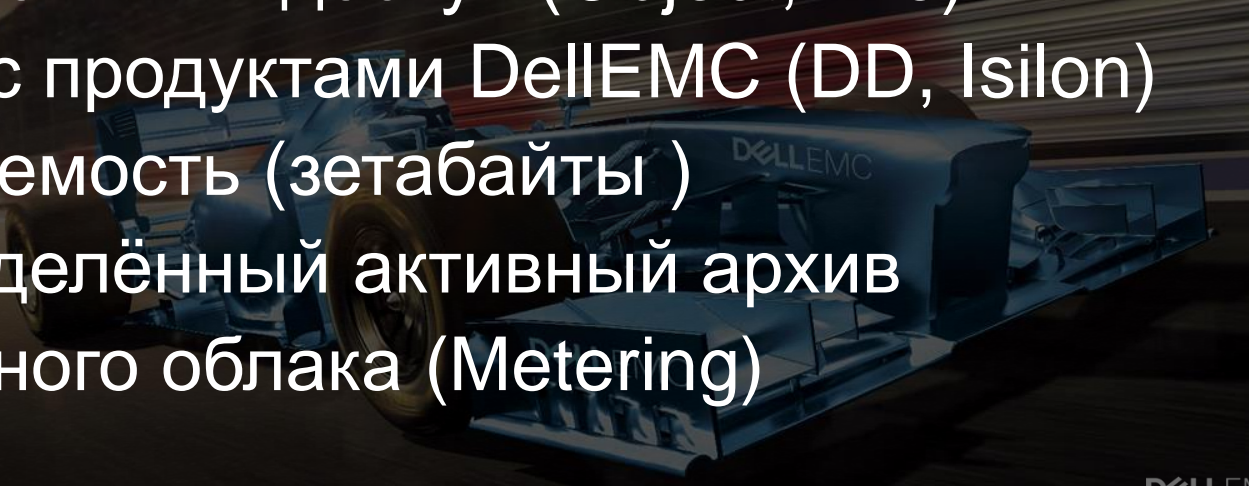
**SIEMENS**

***SMART STORAGE***

- *Sync and Share*
- *IoT and Analytics*
- *Modern App*



# ИТОГИ

- Простое внедрение и настройка
  - Многопротокольный доступ (Object, File)
  - Интеграция с продуктами DellEMC (DD, Isilon)
  - Масштабируемость (зетабайты)
  - Гео- распределённый активный архив
  - Основа частного облака (Metering)
- 



<https://portal.ecstestdrive.com/> -

Публичное облако (идеальный тест)

<https://github.com/EMCECS> -

Примеры, API и т.д.

<https://www.emc.com/collateral/white-papers/h14071-ecs-architectural-guide-wp.pdf> - Архитектура