

# Behreng Nami

**Kıdemli Sistem Mühendisi / Ürün Müdürü**

- **Seagate Cloud Systems and Solutions**
- **Silver Peak WAN Optimizasyon / Replication Acceleration**

Behreng.nami@vitel.com.tr

Cell: +90 554 792 1141

# Gökhan Sakar

**Bilg.Müh – ODTÜ**

Gokhan.sakar@metu.edu.tr

2020: 7.3 Zettabytes



56%  
of total = in the cloud



90%  
Unstructured data



6.5 Zettabytes  
of unstructured data  
stored in the cloud

# Storage Ara Yüzleri

## BLOCK

DAS

SAN

**Hitachi**

**NetApp**

**EMC**

**Vb.**

## FILE

Single NAS

Cluster NAS'lar

**Qnap**

**EMC Isilon**

**Vb.**

## OBJECT

OBS

**Seagate**

**EMC**

**Fujitsu**

**Vb.**

# Disk, Blok ve Dosyalar

Hard Drive



Blocks

```
110010010100100010  
001010100100100010  
010010010001001010  
100100100010010010  
100100100100100100
```

```
110010010100100010  
001010100100100010  
010010010001001010  
100100100010010010  
100100100100100100
```

```
110010010100100010  
001010100100100010  
010010010001001010  
100100100010010010  
100100100100100100
```

File



# BLOK

Dosya Sunucusu \\fileServer

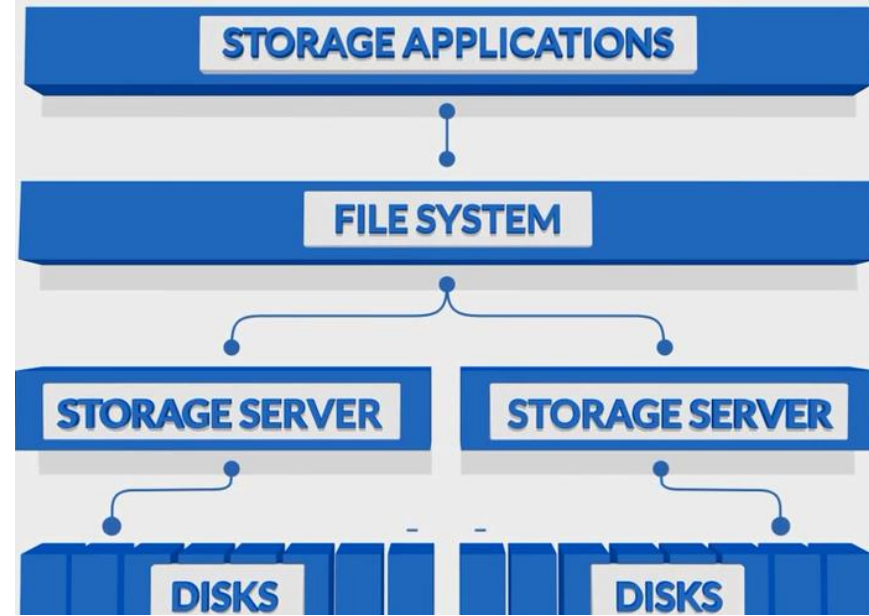
Sanallaştırma Katmanı 'Vmware vb.'

---

Controller 'FC, iSCSI veya SAS'

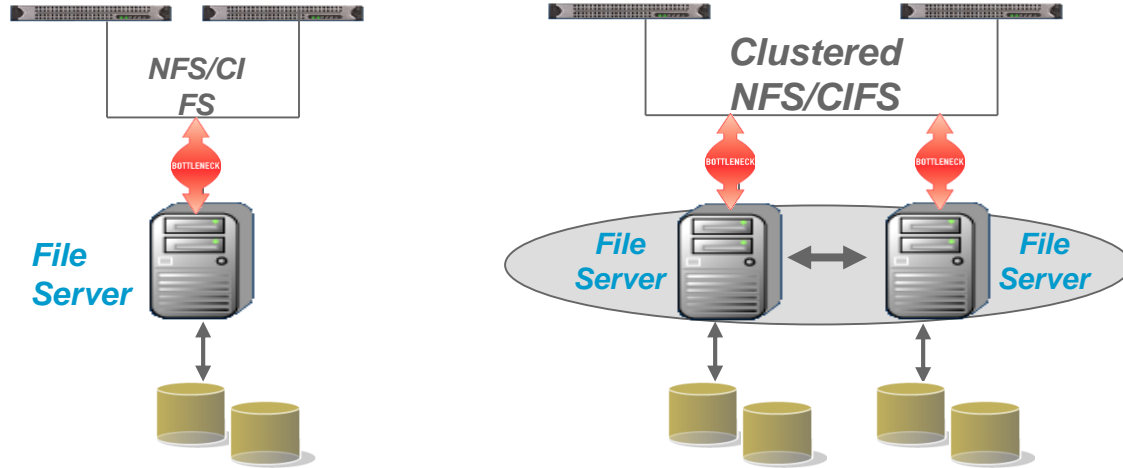
---

Disk Pool



# NAS

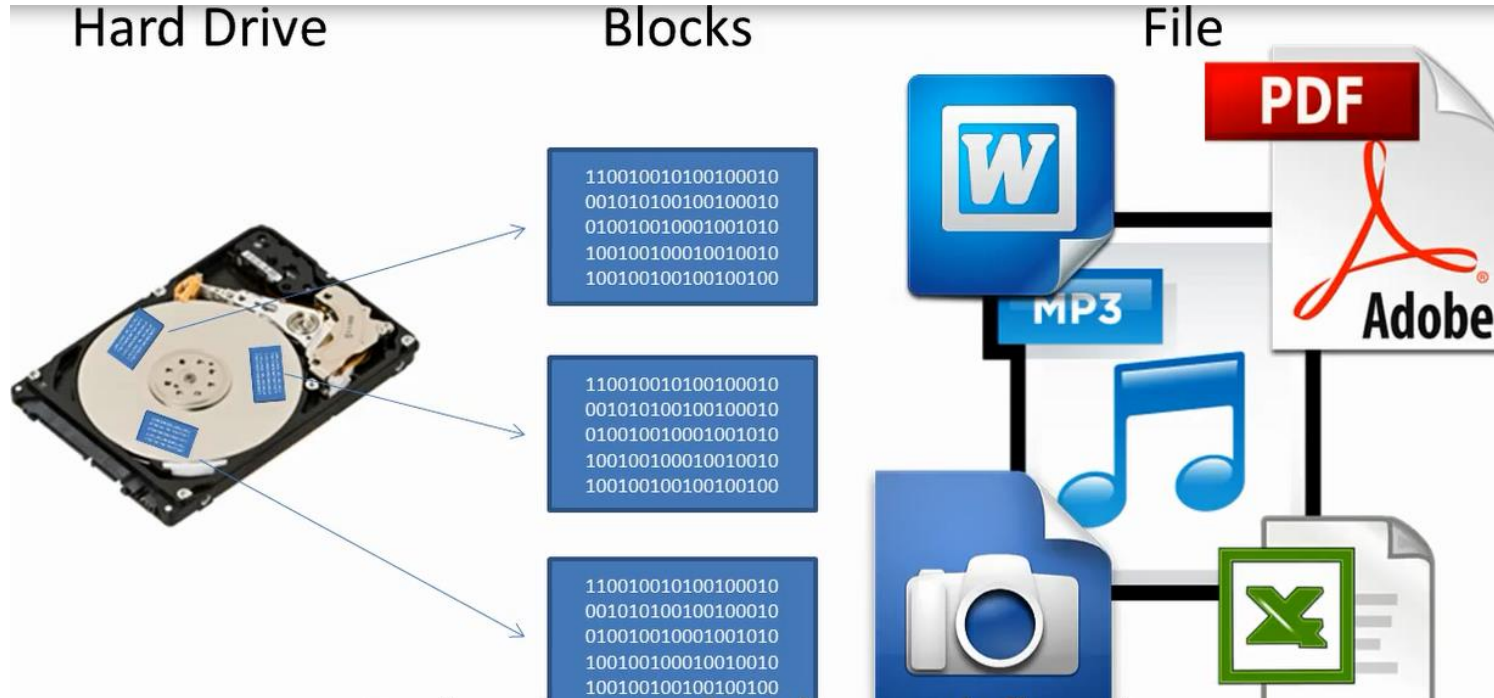
## \\FileServer





**NAS:**  
*Network Attached  
Storage*

**Clustered Storage:**  
*Multiple NAS file servers managed  
as one. Good aggregate  
performance*

# Object Storage



# Object Storage



# Object Storage


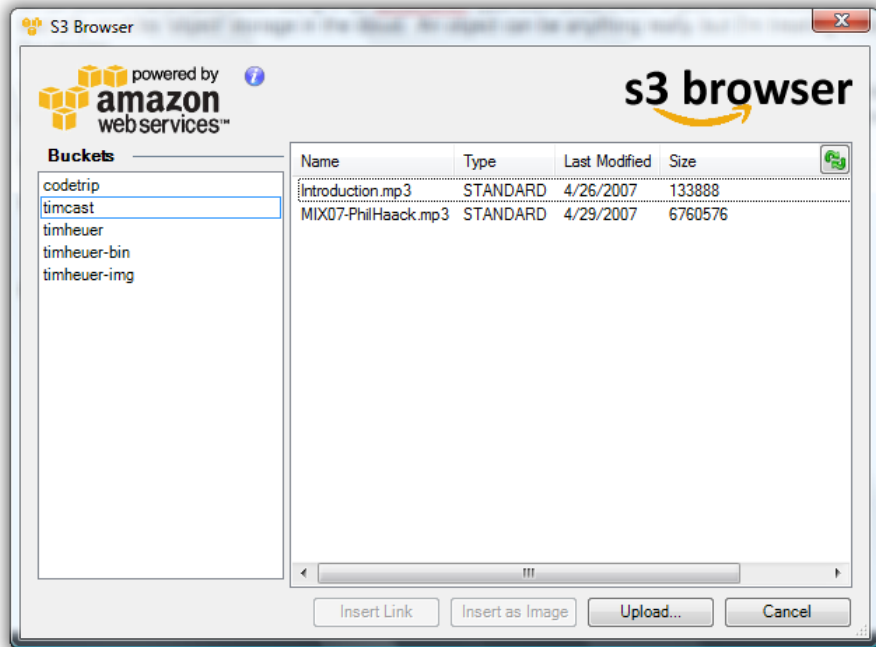
## Object



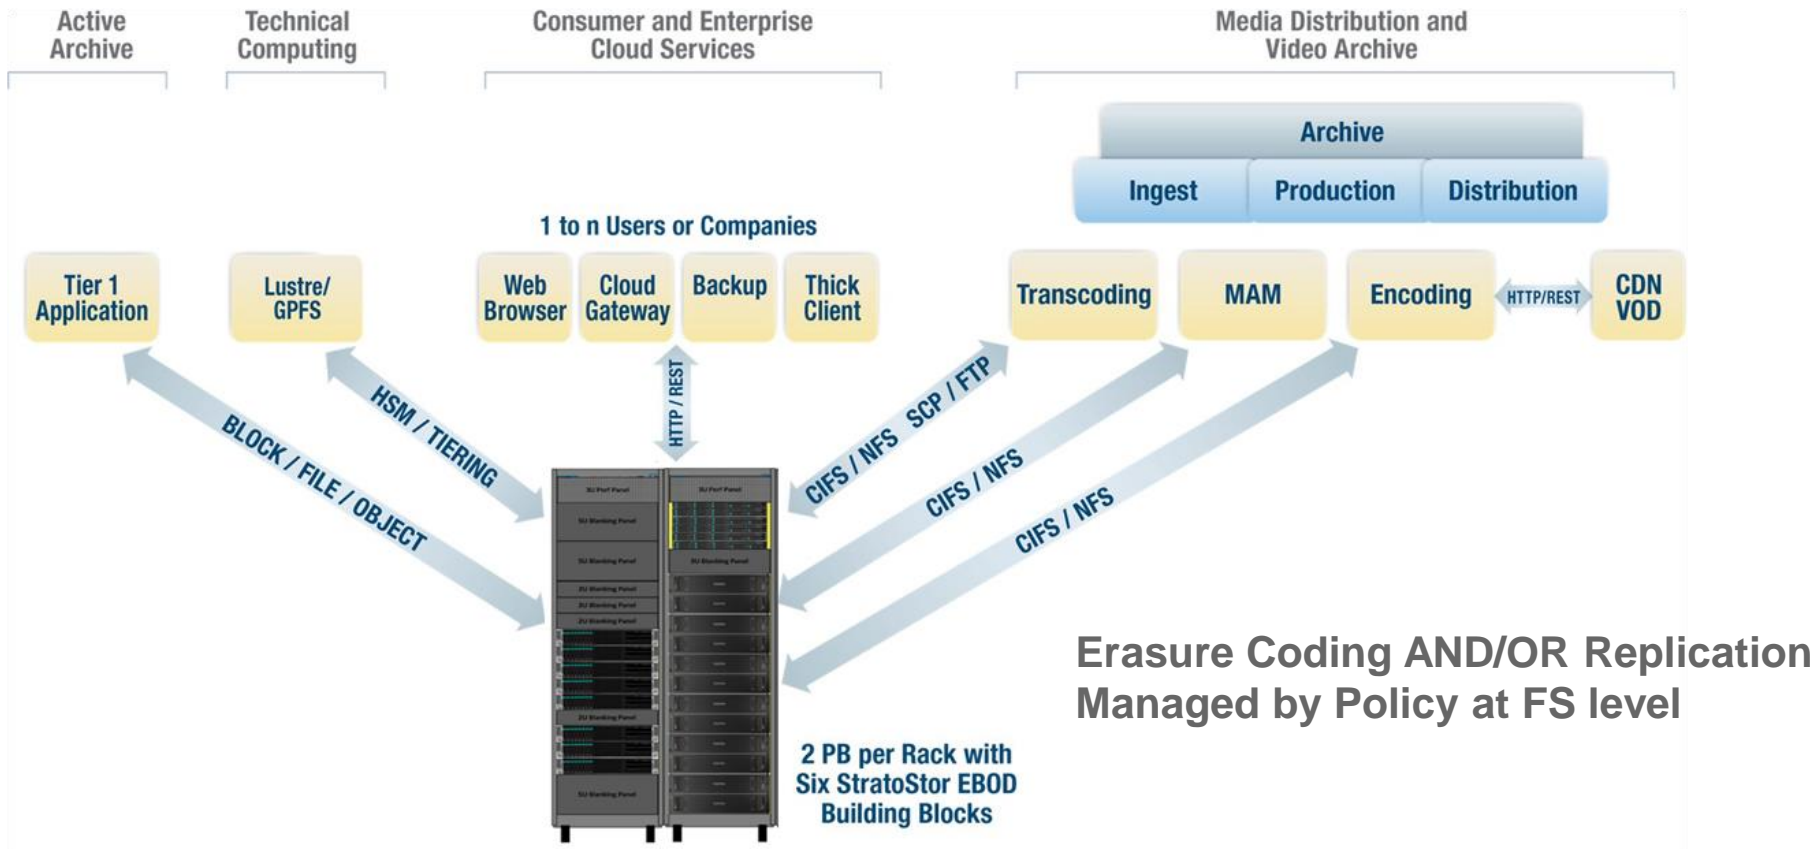
# S3 & SWIFT API for ISV Applications

## REST / HTTP Örneği


- SET
- DELETE
- GET
- PUT



# Seagate Object Storage



# Standard network access with Gateways





Abooooo  
8 ExaByte miiii?

The screenshot shows the VMware vSphere interface for a VM named 'Behreng' (IP: 192.168.2.194). The 'Configuration' tab is active, showing hardware and software settings. The 'Datastores' section is expanded, displaying a table of storage resources. The 'Software Defiend Storage Pool' is highlighted, showing a capacity of 8,388,608,00 TB. The 'Datastore Details' section below shows the server IP (192.168.2.79) and folder path (/opt/CloudGateway/mounts/vol1/).

Identification	Status	Device	Drive Type	Capacity
Disk 1	Normal	HP Serial Attached SCSI Disk (naa.600508b...	Non-SSD	1,81 TB
Disk 2	Normal	HP Serial Attached SCSI Disk (naa.600508b...	Non-SSD	1,82 TB
SATA_1TB	Warning	MSFTiSCSIDisk (naa.60003ff44dc75adcbf...	Non-SSD	1.023,75 GB
Software Defiend Storage Pool	Normal	192.168.2.79:/opt/CloudGateway/mounts/v...	Unknown	8.388.608,00 TB
SSD_1TB	Normal	MSFTiSCSIDisk (naa.60003ff44dc75adcbf...	Non-SSD	899,75 GB

**Datastore Details**

**Software Defiend Storage Pool**

Server: 192.168.2.79      8388608,0      Capacity

Folder: /opt/CloudGateway/mounts/vol1/      0,00 B      Used

8388608,0      Free

# Erasure Coding vs RAID

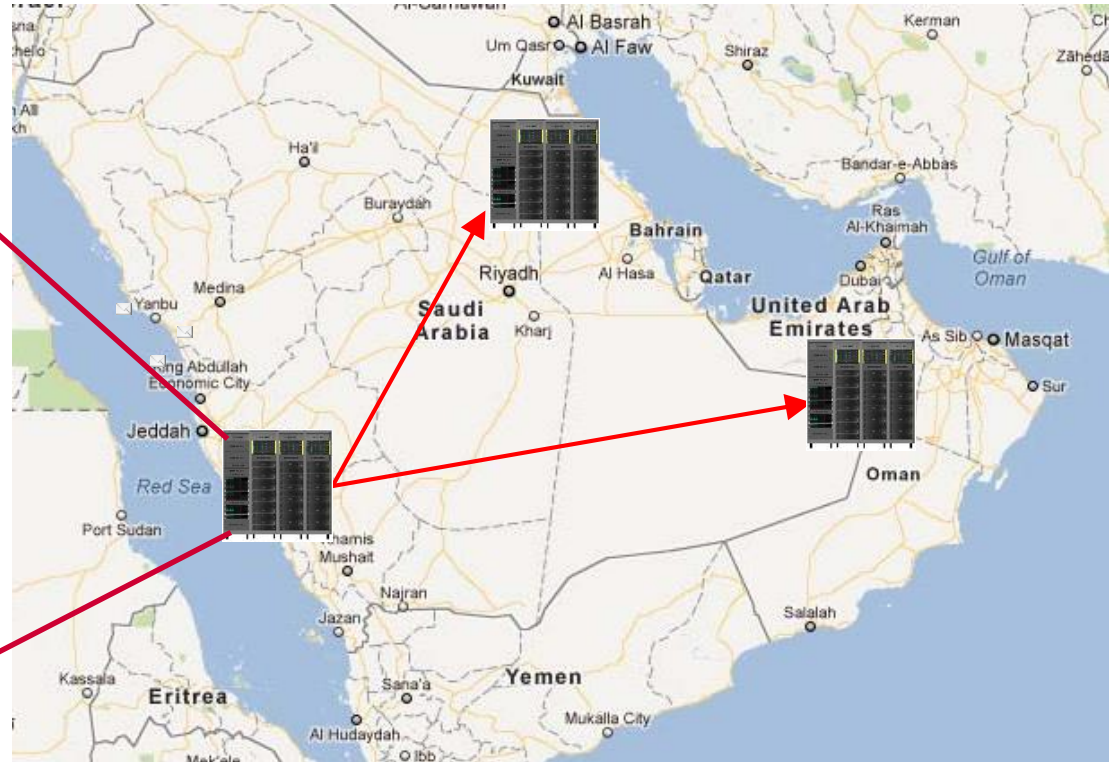
Configuration comparison for 1PB <sup>1</sup>	Raw Storage (no protection)	RAID5 <sup>2</sup>	RAID6 <sup>3</sup>	Replication <sup>4</sup>	Erasure Coding <sup>5</sup>	Deep Archive <sup>6</sup>
Usable capacity (TB)	1,000	1,000	1,000	1,000	1,000	1,000
Redundancy ratio	0	1.2	1.33	3	1.29	1.5
Additional storage for protection (TB)	0	200	333	2,000	286	500
Total raw storage (TB)	1,000	1,200	1,333	3,000	1,286	1,500
Probability of data loss <sup>7</sup> (%)	100%	95%	0.025%	0.015%	<b>0.000015%</b>	<b>0.0000000001%</b>
Durability <sup>8</sup> (%)	0%	5%	99.975%	99.985%	<b>99.999985%</b>	<b>99.9999999999%</b>
Availability	No availability  1 disk lost means the entire disk pool is lost (1 large concat or striped pool of 1000 drives)	Very limited  1 disk max. per disk group. Losing 1 disk group means global downtime (1 large concat or striped pool of 200 RAID5 groups)	Better but risky and not optimal  2 disks max. per disk group. Losing 1 disk group means global downtime (1 large concat or striped pool of 167 RAID6 groups)	Good  2 copies max. (improved with number of copies)	<b>Excellent</b>  <b>Tolerates the loss of up to four servers for any given transaction</b>	<b>Almost Perfect</b>  <b>Tolerates the loss of up to nine servers for any given transaction, and protects against site failure</b>
<sup>1</sup> 1PB=1000TB, Disk size =1TB <sup>2</sup> RAID5 deployed as 5 data disks + 1 parity disk <sup>3</sup> RAID6 deployed as 6 data disks + 2 parity disks		<sup>4</sup> Replication with 3 copies in total <sup>5</sup> Erasure Coding(14,4) <sup>6</sup> Erasure Coding(18,9) 3 sites <sup>7</sup> For 1 year period		<sup>8</sup> The Amazon website describes durability (with respect to an object stored in S3) "as the probability that the object will remain intact and accessible after a period of one year. 100% durability would mean that there's no possible way for the object to be lost, 90% durability would mean there's a 1-in-10 chance, and so forth."		

# Data Availability and Durability

*Available Anytime, Anywhere; can Survive Partial Cluster Failures*



Data Is Replicated...



...And Further Replicated & Distributed

# Object Storage Avantajları

- Sınırsız Ölçeklenme 'Kapasite ve Obje Sayısı'
- Performans
- Multi Tenant Platform
- Yüksek Güvenlik Mimarisi
- %50 düşük TCO
- Yüksek Süreklilik 'Veri Kaybı İhtimali **3.000** Yılda bir kez'

# Use Cases and Verticals

## Enterprise Cloud

Public or Private  
VM Tier 2/3



## Active Archive

Digital Media  
CCTV  
Tape Replacement  
for Backup Solutions



## Distributed Computing

Research and Dev.  
HOME Storage  
with GPFS/Lustre



## Web & Cloud

Website Hosting  
VM Hosting  
Files Storage Services



## Content Distribution

VOD, VDI,  
nDVR, nPVR



Teşekkürler

Sorular ?

LinkedIn: Behreng Nami