

Move Objects to LTFS Tape Using HTTP Web Service Interface

Matt Starr, CTO Spectra Logic

Jeff Braunstein, Developer Evangelist

Data created each year growth
40% annually

2020

$6.6 \times 6.6 = 35$ Zettabytes


2015

$2.8 \times 2.8 = 6.5$ Zettabytes

2009

$1 \times 1 = 800$ Exabytes


Genomics



6 Cows = 1TB data

University of Washington, St.
Louis

Video Data Explosion





What Happens in an Internet Minute?

1,572,877 GB of global IP data transferred¹

10 Million ads displayed²

347,222 Tweets³

3.3 Million pieces of content shared⁴

6.9 Million messages sent⁴

Netflix + Youtube = more than 1/2 of all traffic⁵

\$400 Million during Alibaba peak day sales⁶

438,801 Wiki page views⁷

10 Million WeChat messages at its peak⁸

34.7 Million instant messages (MIM) sent⁸

194,064 app downloads¹⁰

\$133,436 in sales¹¹

31,773 hours of music played¹²

38,194 photos uploaded¹³

57,870 page views¹⁴

100 hours of video uploaded¹⁶

4.1 Million searches¹⁵

138,889 hours of video watched¹⁶

23,148 hours of video watched¹⁷

And Future Growth is Staggering

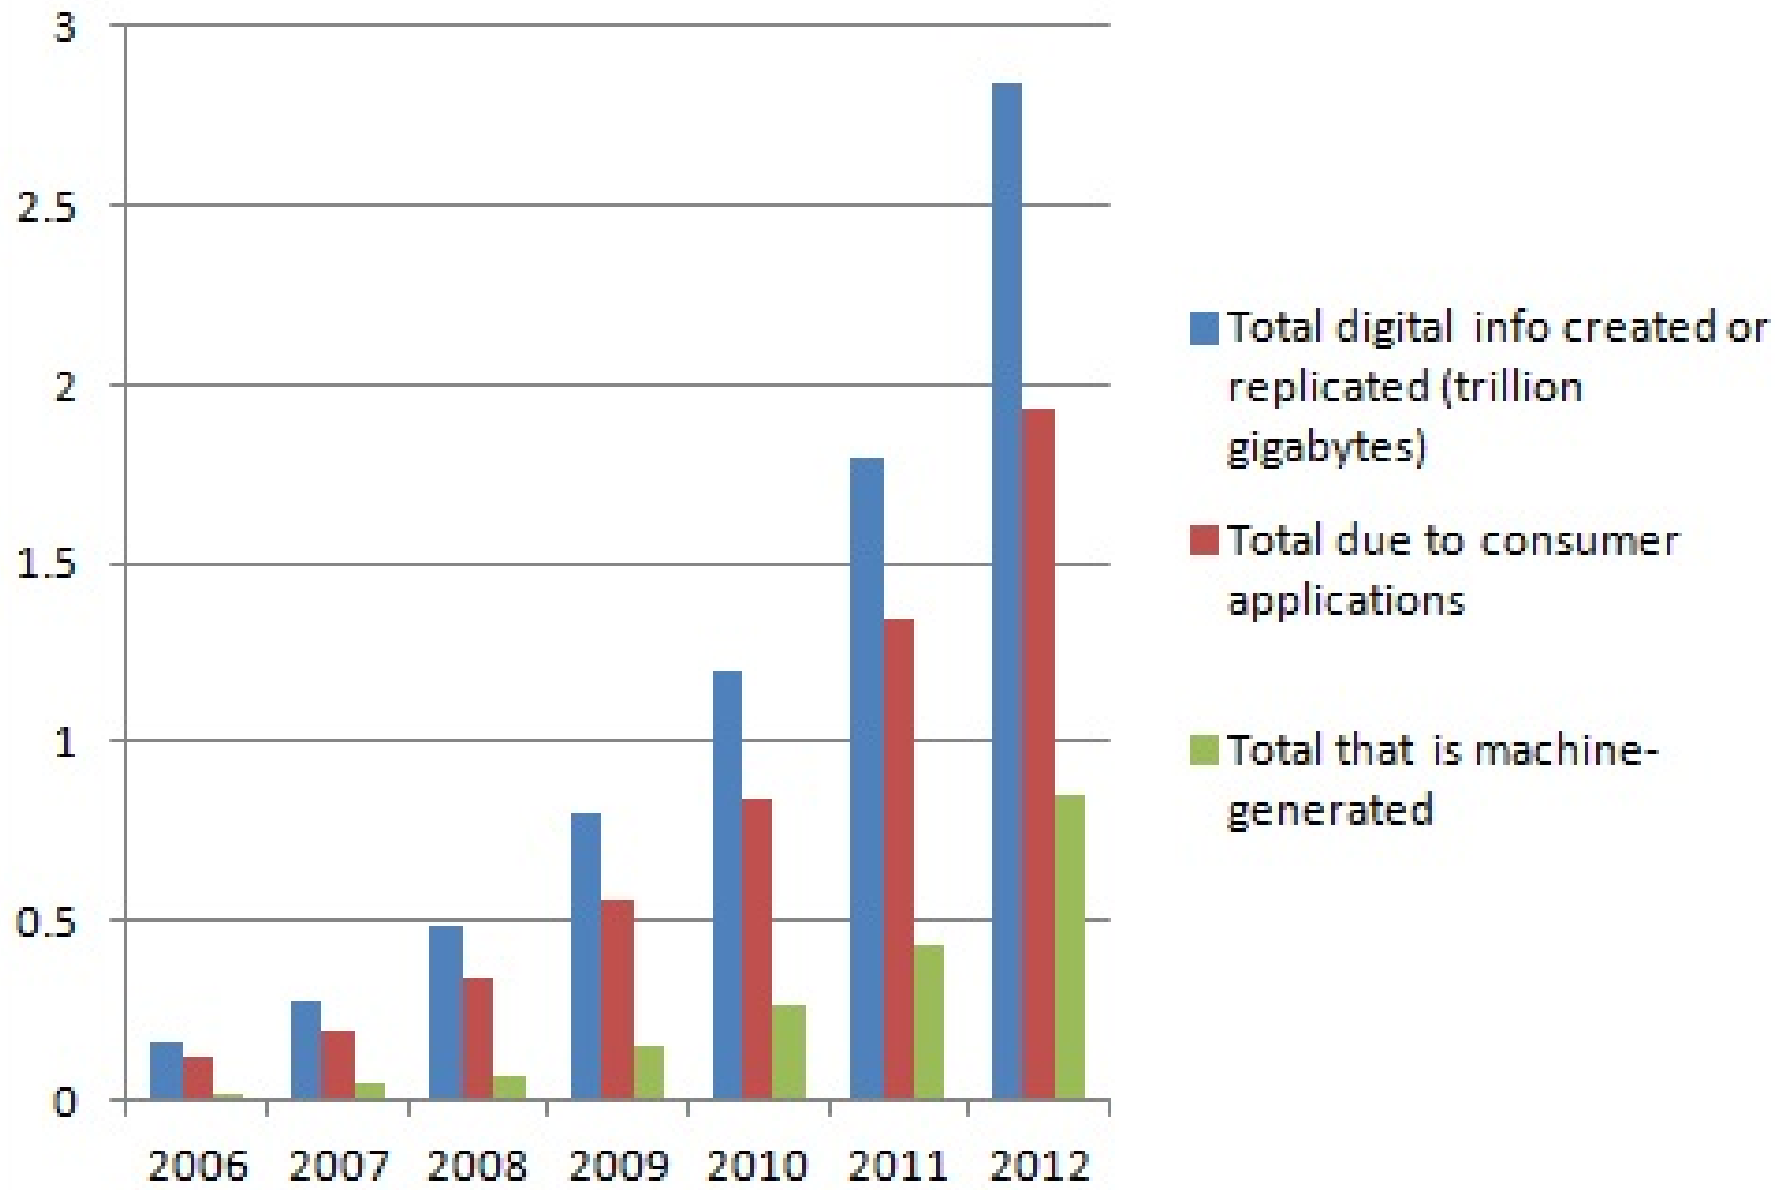


By 2017, mobile traffic will have grown **13X** in just 5 years¹



In 2017, there will be **3X** more connected devices than people on Earth¹

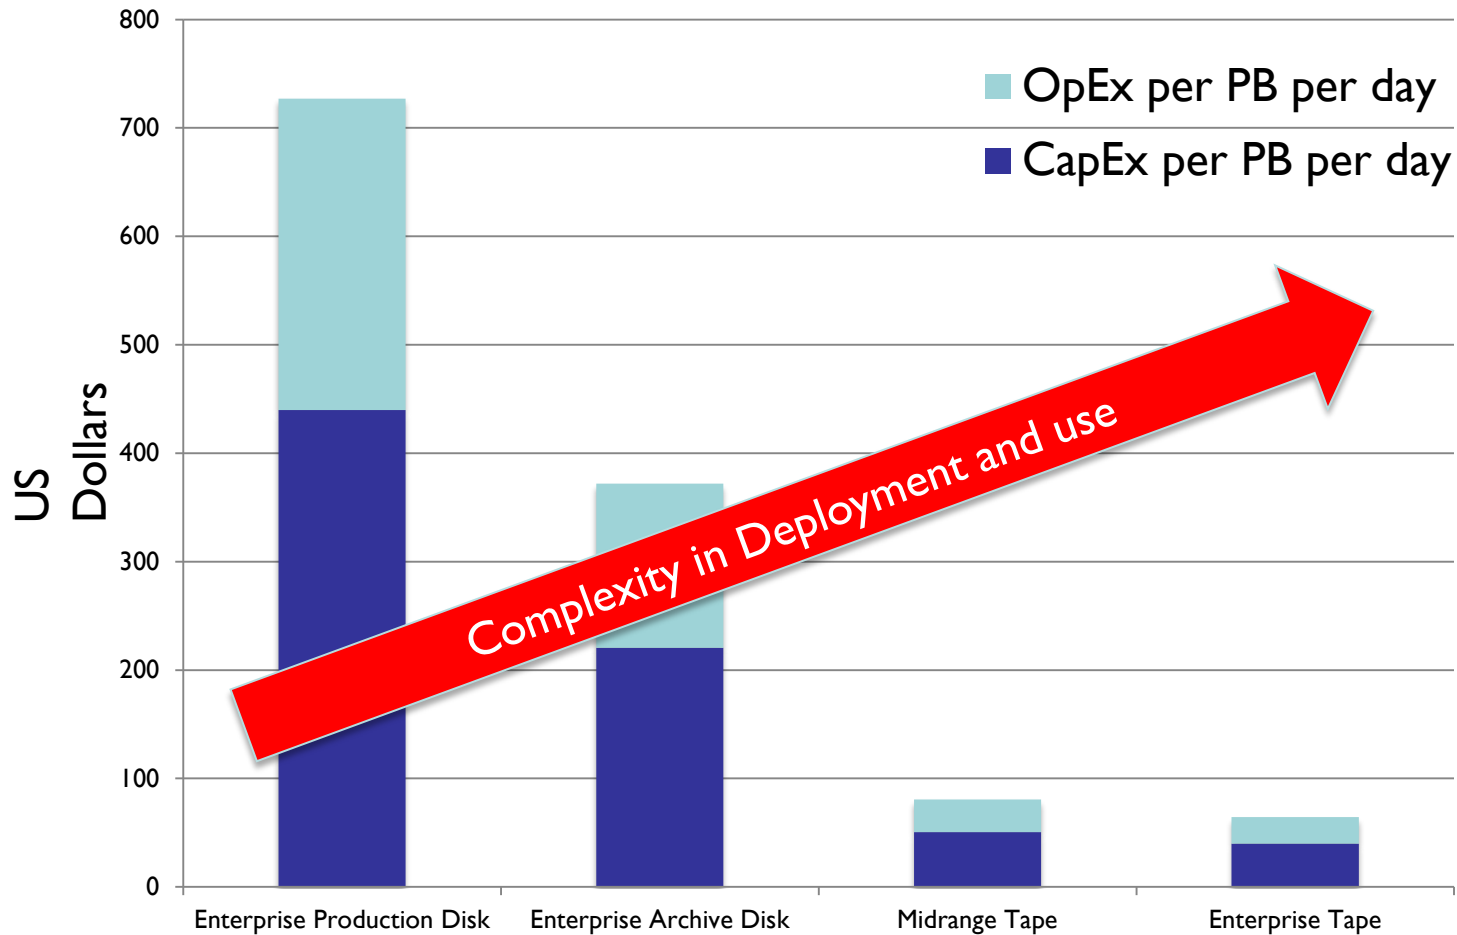
All digital data created reached **4 zettabytes** in 2013¹⁸



Digital System and their growth

- ❑ A \$1000 Genome scan take 780 MB fully compressed
- ❑ A 2011 HiSeq-2000 scanner generate 20TB per month
- ❑ A typical security camera - 105GB of data per day.
 - ❑ City of Riverside, CA has 1000 cameras
- ❑ A digital camera is now 12MP or better
- ❑ A 4K video consumes 4.25 TB per hour
- ❑ Geospatial
 - ❑ An average ingest (rain) of 5 TB a day of new imagery.
 - ❑ 80 TB a day of refined product produced.

Storage Cost Per PB/Day over 5 yrs



What options, cloud?

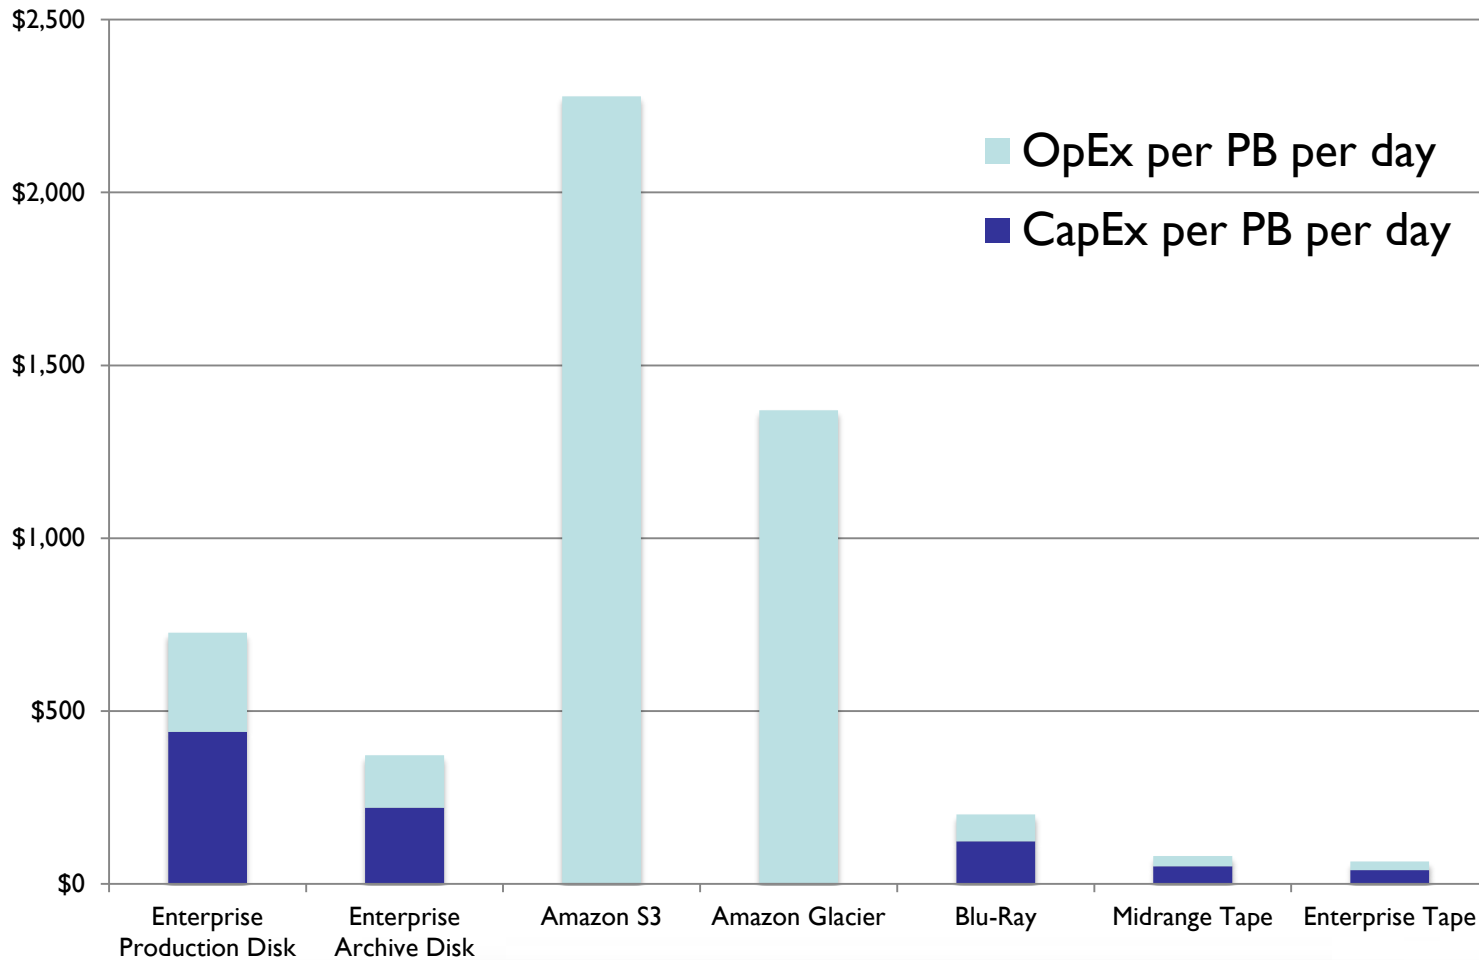
Traditional Storage

- ❑ Disk
 - ❑ Easy to deploy
 - ❑ Expensive to run,
 - ❑ Expensive to buy
- ❑ Tape
 - ❑ Complex to deploy
 - ❑ Inexpensive to run
 - ❑ Inexpensive to buy

Object and Cloud Storage

- ❑ Public Cloud
 - ❑ Easy to deploy
 - ❑ Expensive to run
 - ❑ Little to buy
- ❑ Private Cloud
 - ❑ Complex to Deploy
 - ❑ Expensive to run
 - ❑ Moderate to buy

Pitfalls of Public Cloud (\$/PB/Day)



Amazon

- ❑ Amazon apologizes for cloud outage and permanent data loss
 - ❑ ...After connectivity was restored, the issue of data corruption had to be addressed. By April 24th, the company had restored nearly 99% of their customers' data at which point Amazon "*began forensics on the remaining volumes.*"
- ❑ Amazon has over 1 trillion objects!
- ❑ If they did lose 1% this is 10 billion objects

How'd would this email go over...

From Amazon

This is a notification that your volume vol-13c9647a experienced a failure due to multiple failures of the underlying hardware components and was unable to be recovered. We recommend recovering from your most recent snapshot.

We regret the loss and inconvenience.

Sincerely,

--- Dxxx Dxxxxxxx



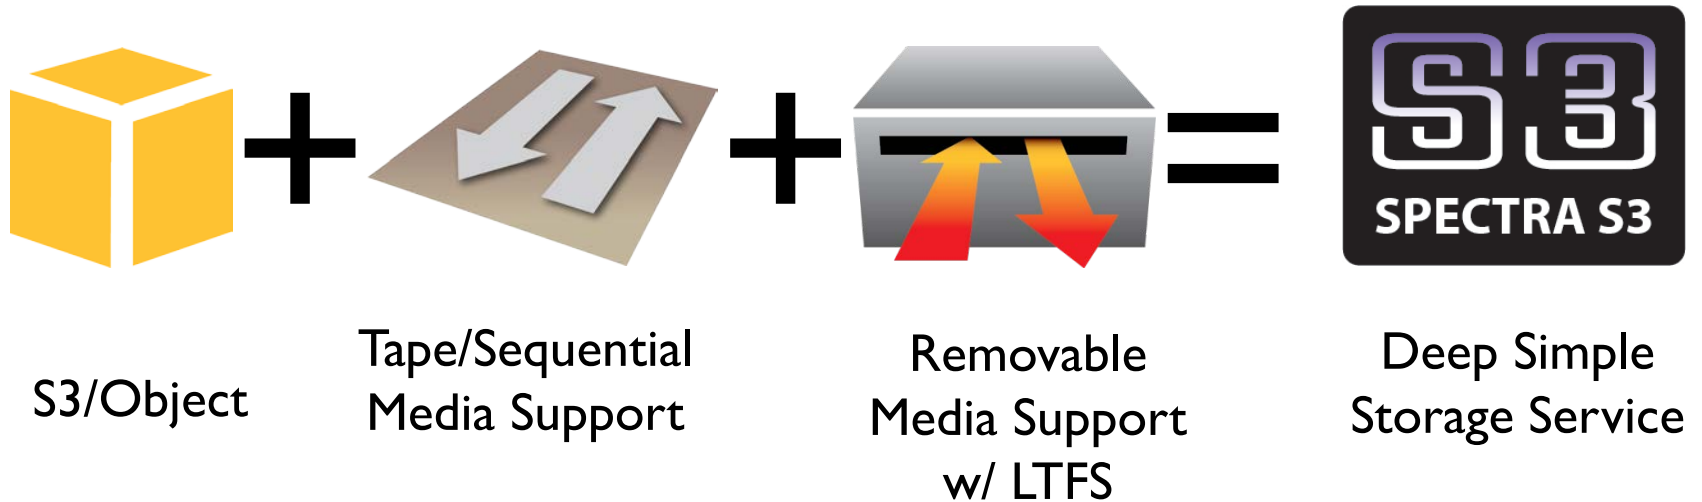
Nirvanix, VC pull the plug

- ❑ On September 16, 2013 Nirvanix notified customers that they had until September 30, 2013 to move their data off of the service at which point Nirvanix would shut down. By September 28, 2013 Nirvanix had closed its doors and shutdown all services
- ❑ The primary lesson to learn is that you, as your organization's data custodian, can outsource the actual data to a cloud provider. But you, not the cloud provider, are ultimately responsible for your organization's data. - Howard Marks

Cloud easy to use, you don't control your data

- ❑ Good points on Cloud Storage
 - ❑ Cloud ready available toolkits
 - ❑ Cloud storage is naturally WAN friend
 - ❑ Cloud storage sets the user expectation
 - ❑ It is put away
- ❑ Dangerous side of Cloud Storage
 - ❑ No proven SLA
 - ❑ Expensive data pipe
 - ❑ Costly get all your data back

What if...



Start with Object, S3 and Rest

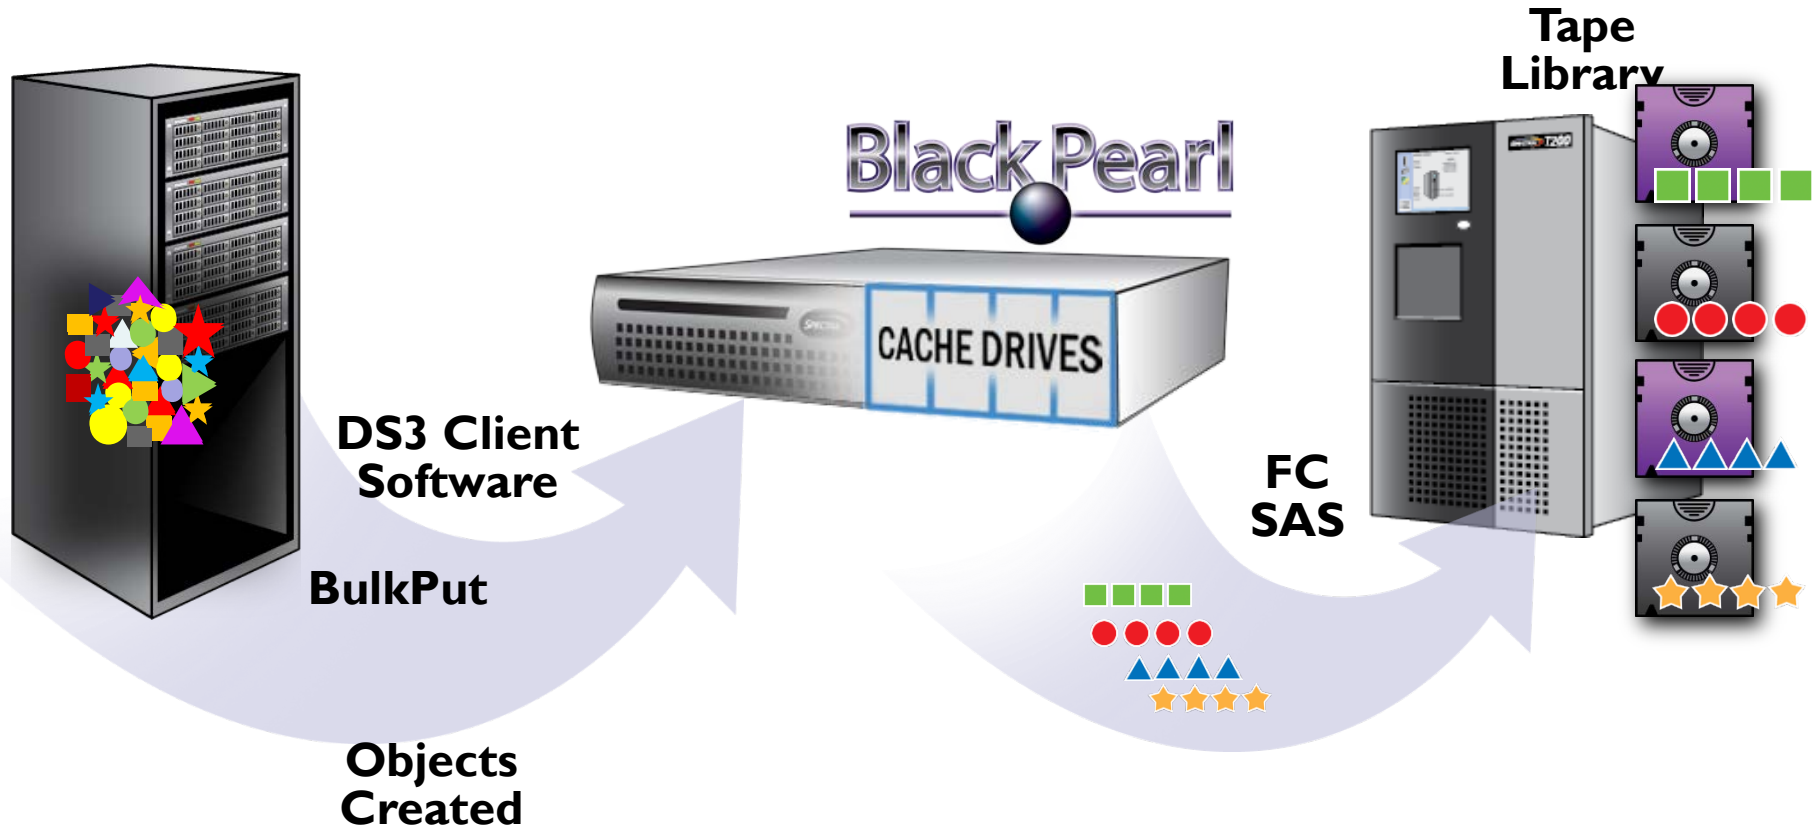
Use Archive media as the target

Spectra S3 facilitates an Objected based Archive

Black Pearl



What does BlackPearl and Spectra S3 do?



What Makes BlackPearl Revolutionary:

**BlackPearl plus
T950 with TSI I150 drives**



\$0.09/GB

7.2 PB Usable

**BlackPearl plus
T950 with LTO6 drives**



\$0.16/GB

2.4 PB Usable

**BlackPearl plus
T380 with LTO6 drives**



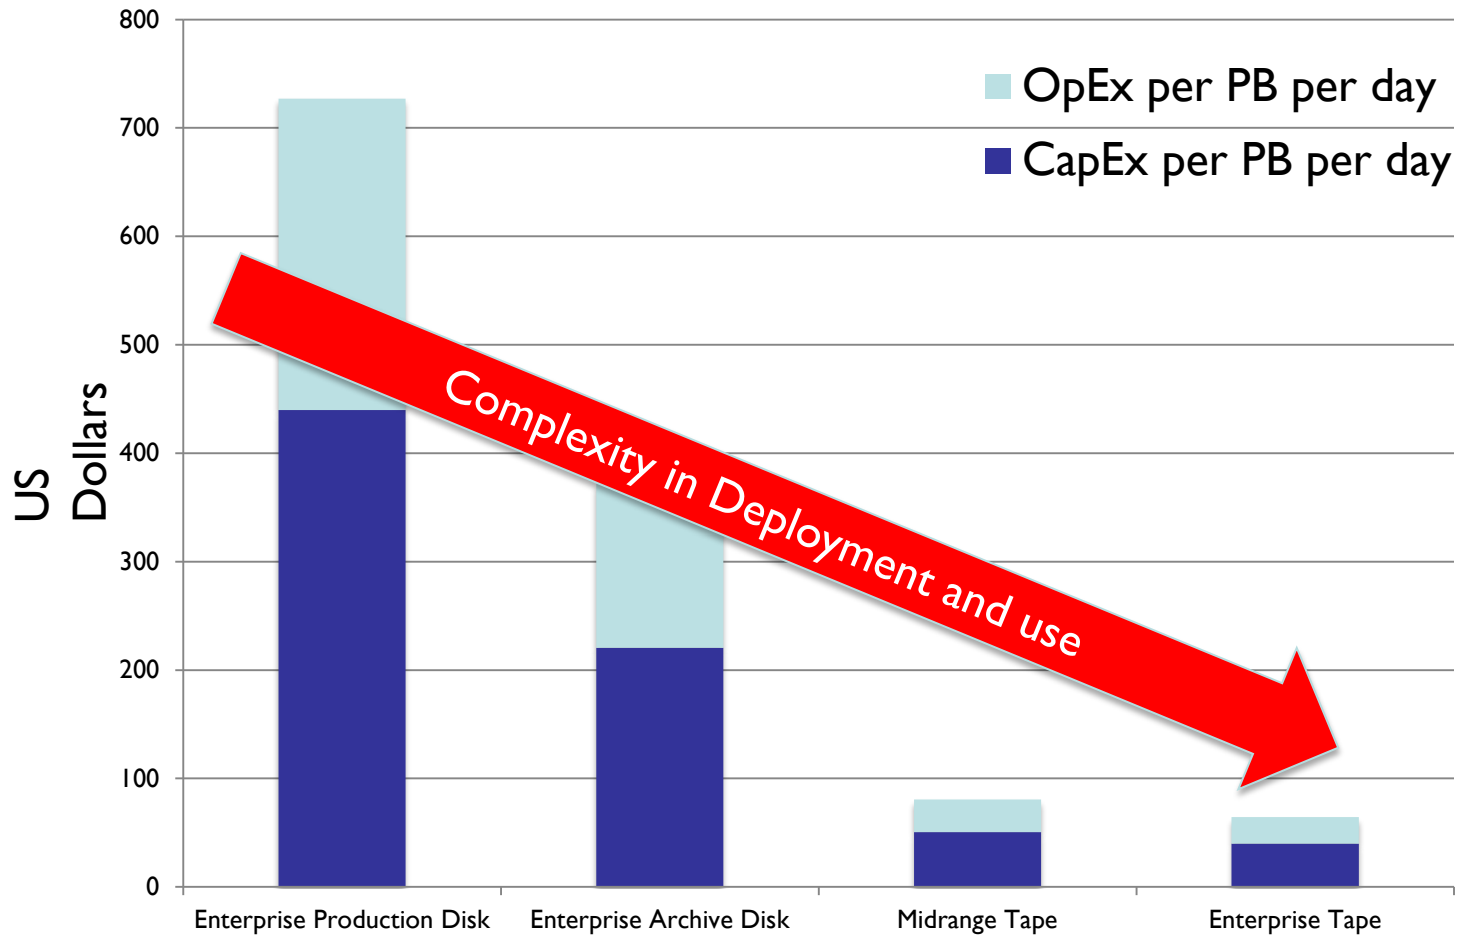
\$0.23/GB

900 TB Usable

Advanced Bucket Management



Storage Cost Per PB/Day over 5 yrs



Data Movement Approaches with Spectra S3:

- ❑ Automated: Application or user built clients
- ❑ Semi-automated: NFI or scripting
- ❑ Manual: Deep Storage Browser or

```
#!/bin/sh
```

```
PATH=$PATH:/Users/matts/Downloads/ds3_java_cli-0.7-3.5-  
SNAPSHOT/bin  
export PATH
```

```
export DS3_ACCESS_KEY=c3BIY3RyYQ==  
export DS3_SECRET_KEY=Vz8m6cLy  
export DS3_ENDPOINT=192.168.56.101:8080
```

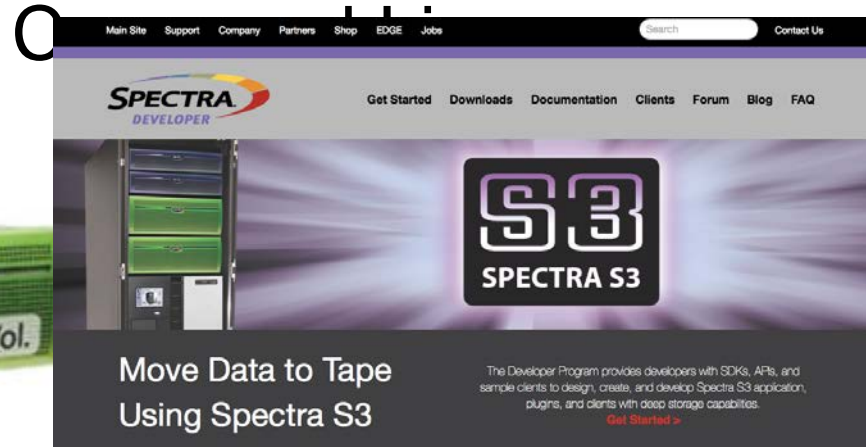
```
DIR_TO_PUT="/Users/matts/archive/Assets"
```

```
ds3_java_cli -http -c put_bucket -b cokemachine  
ds3_java_cli -http -c put_bulk -d $DIR_TO_PUT -b  
cokemachine
```

NFI
Network File Interface

CIFS Vol. or NFS Vol.

Spectra Verde Storage
Capacity up to 1.7PB



Developer Summit

Join us for Spectra Logic's inaugural BlackPearl Developer Summit, Oct. 20, 2015 9:00 AM, WebEx

Quarterly Newsletter

Read the latest quarterly Developer Program newsletter for August 2015

Spectra S3 Developer Portal

Questions

How to develop