

# **Centralized vs. Distributed A Great Storage Debate**

**Live Webcast  
September 11, 2018  
10:00 am PT**

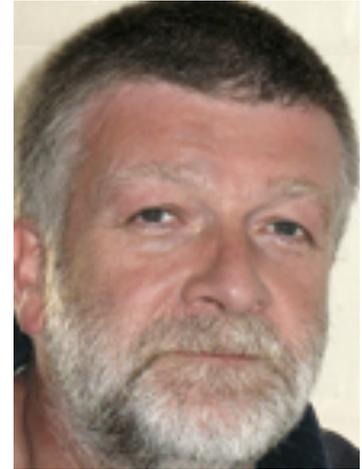
# Today's Presenters



**J Metz**  
**SNIA Board of Directors**  
**Cisco**



**John Kim**  
**SNIA ESF Chair**  
**Mellanox**



**Alex McDonald**  
**SNIA ESF Vice Chair**  
**NetApp**

## SNIA-At-A-Glance



**170**  
industry leading  
organizations



**3,500**  
active contributing  
members



**50,000**  
IT end users & storage  
pros worldwide

Learn more: [snia.org/technical](https://snia.org/technical)



- ◇ The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- ◇ Member companies and individual members may use this material in presentations and literature under the following conditions:
  - ◇ Any slide or slides used must be reproduced in their entirety without modification
  - ◇ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- ◇ This presentation is a project of the SNIA.
- ◇ Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- ◇ The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

**NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

# Today's Agenda

- The Rules of the Debate
- The “Whats and Hows”
  - ◆ Centralized Storage
  - ◆ Distributed Storage
- The Trade-Off Debate



# The Rules of The Debate



No hitting below the belt

- ◆ Spoiler Alert: There is no “winner”
  - ◆ This is all about the “sweet spot”
- ◆ Participants:
  - ◆ Define the technologies
  - ◆ How they work
  - ◆ Discuss the trade-offs

# Storage Has One Job!

- ◆ One main job
  - ◆ “Give me back the correct bit I asked you to hold for me.”
- ◆ Everything we do in storage (including storage networking) is based around completing that job safely, securely, reliably, and without error



You had one job!

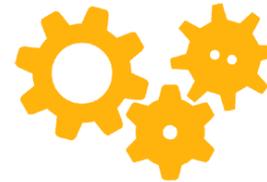
# All Storage

- ◆ Needs to:
  - ◆ Protect data
  - ◆ Keep data secure
  - ◆ Stay within regulatory compliance
  - ◆ Be manageable
  - ◆ Be backed up!
- ◆ May need to:
  - ◆ Be scalable
  - ◆ Be sharable
  - ◆ Be very fast



- › Items to consider in choice of storage
  - ◆ *Access* — what protocols can I use?
  - ◆ *Performance* — will my applications & hence users be happy?
  - ◆ *Availability* — can I tolerate periods without access?
  - ◆ *Capacity* — how big do I need?
  - ◆ *Protection* — how do I ensure my data's integrity?
  - ◆ *Durability* — how long do I need to store my data?
  - ◆ *Security & Privacy* — will sensitive data be OK here?
  - ◆ *Cost* — is it cheaper than the alternatives?
- › Let's discuss some of these

## JUDGING CRITERIA



◆ John Kim

## CENTRALIZED STORAGE

- ◆ Direct Attached Storage (DAS)
  - ◆ Storage directly attached to just one server
- ◆ Storage Area Network (SAN)
  - ◆ Centralized block storage system connected to multiple hosts using networks such as Fibre Channel, iSCSI, NVMe-oF, or InfiniBand
- ◆ Network Attached Storage (NAS)
  - ◆ Centralized or distributed file storage connected to multiple hosts using file protocols, usually using Ethernet networking
- ◆ Hyperconverged Infrastructure (HCI)
  - ◆ Set of servers each with compute and storage resources, often sharing those resources with each other

# First There was Local Storage



- ▶ **Local storage for each server**
  - ◆ Inside the server or directly attached to one server (DAS)
- ▶ **Easy to buy, set up, and consume**
  - ◆ Server vendor/integrator can install
  - ◆ All operating systems/hypervisors can use
  - ◆ No special drivers or networking required
- ▶ **But...**
  - ◆ Inefficient and difficult to manage at scale
  - ◆ Issues with backup, failover, utilization, sharing

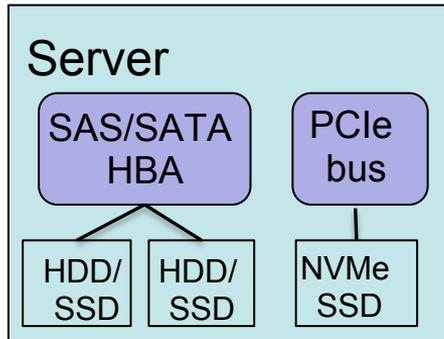
# Then Centralized Storage

- ◆ Consolidate storage into centralized systems
  - ◆ Each supports multiple servers
  - ◆ Connect via PCIe, SAS, SAN, NAS or Object
- ◆ Easier to share and protect data
  - ◆ Higher utilization
  - ◆ Easier backup, recovery, failover, sharing

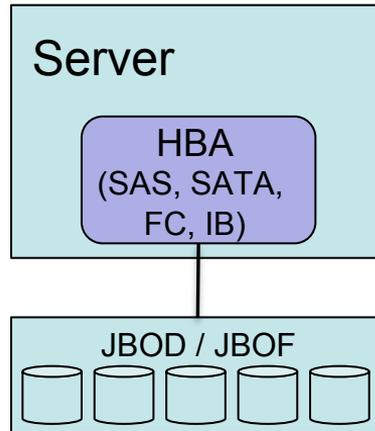


# Comparing Storage Models

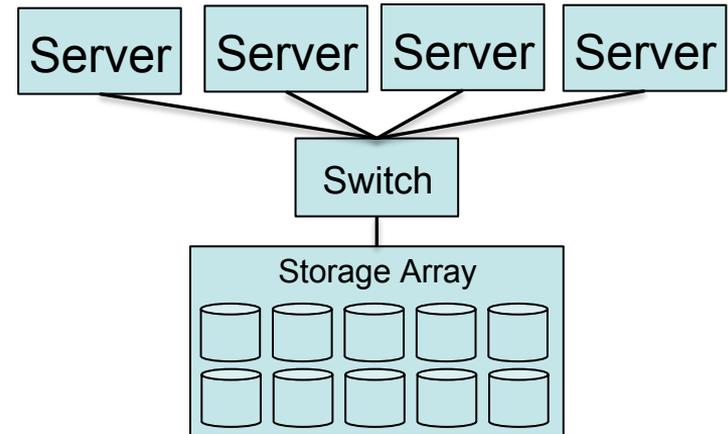
## Internal vs. DAS vs. SAN/NAS



Internal Storage



DAS



SAN / NAS  
(Centralized Storage)

# Comparing Storage

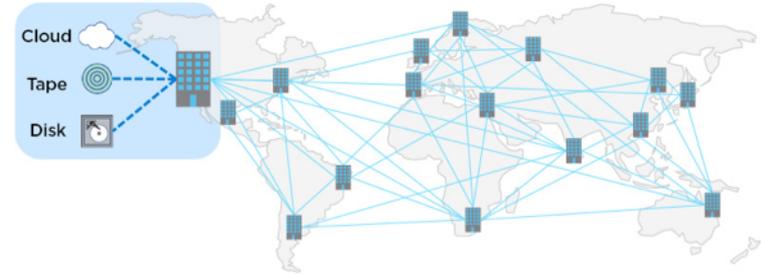
Type	Location	Connection	Sharing Level	Distributed?
Internal	Inside server	SAS/SATA/PCIe	None	No
DAS	Attached to 1 server	SAS/SATA/PCIe, FC, IB	None	No
SAN	Centralized array	FC, Ethernet, IB	Array	Rarely
NAS	Centralized array(s)	Ethernet	File	Sometimes
Object	Multiple arrays	Ethernet	Object	Yes
HCI	In each server, or...	Ethernet	depends	Usually

◇ Alex McDonald

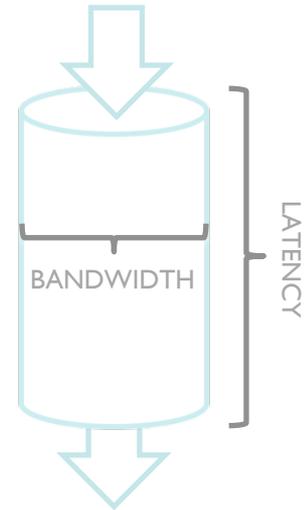
## DISTRIBUTED STORAGE

# Distributed Storage: a definition

- ▶ Difficult to precisely define
  - ◆ Data stored on many systems which behave as a single entity
  - ◆ Geographically or regionally dispersed rather than local to a data center
  - ◆ Accessed over LAN or WAN, commonly Ethernet
  - ◆ Cloudy-ish; often implemented on shared resources
- ▶ Well, I give up...
  - ◆ Not centralized or hyperconverged (HCI)
  - ◆ Scales out (horizontally) rather than up (vertically)



- ◆ Network connectivity & performance criteria
  - ◆ Bandwidth & Latency
    - › “Bandwidth problems can be cured with money. Latency problems are harder because the speed of light is fixed - you can’t bribe God.”
  - ◆ Compute location
    - › Low bandwidth & poor latency tolerable if the compute is next to the data, and we only need to send/receive small amounts
  - ◆ Flash technologies? SSD? NVMe?
    - › Yes; this isn’t just about cheap spinning disk any more
- ◆ Protocols; tend to be application driven
  - ◆ Object type storage (S3, CDMI, Swift)
  - ◆ LAN/WAN protocols (SMB, NFS)
  - ◆ Block (iSCSI)
- ◆ Rule of thumb
  - ◆ The less “cloudy” or “WANny” the access, the less likely the application will tolerate high latency and/or low bandwidth

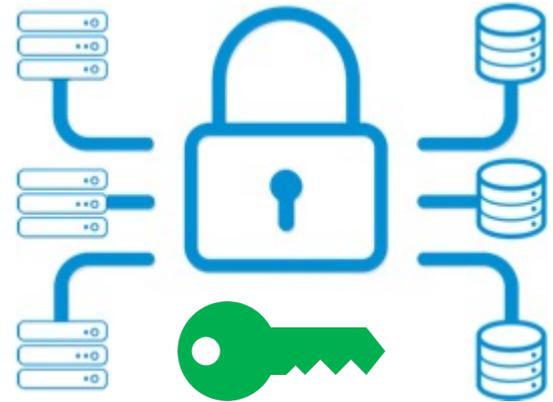


## Security vs. Privacy

- ◆ Security is making sure only the right people/systems have access to the data
- ◆ Privacy ensures that the data isn't misused
  - › Privacy is explored further here:  
<http://sniablog.org/security-gdpr-snia-and-you/>

## Security measures

- ◆ Identification & authentication systems
  - › e.g. Kerberos & NFS, LDAP & SMB
- ◆ End-to-end encryption (including devices)
- ◆ Storing data in the right place & knowing how the data is managed
  - › Replicas, mirroring, cloud brokering, backups can all be in different places and differently secured



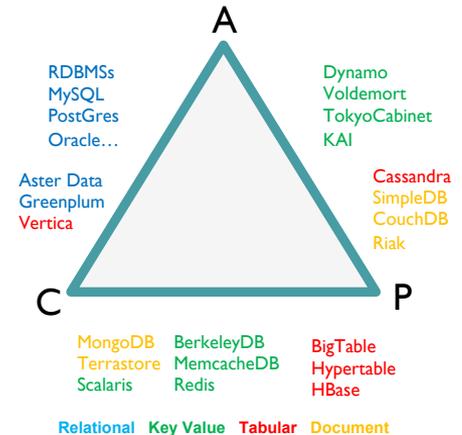
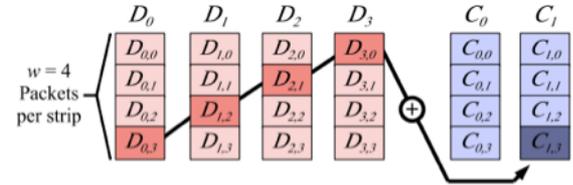
# Scaling out rather than up

- ▶ Capacity can be seen as infinite
  - ◆ “It’s just a matter of cost...”
  - ◆ More capacity tends to exacerbate these issues:
    - ▶ More cold data
    - ▶ Higher bandwidth, especially to distributed storage
    - ▶ Harder to avoid putting compute with the data
    - ▶ Increased data amnesia
    - ▶ Harder systems management problems



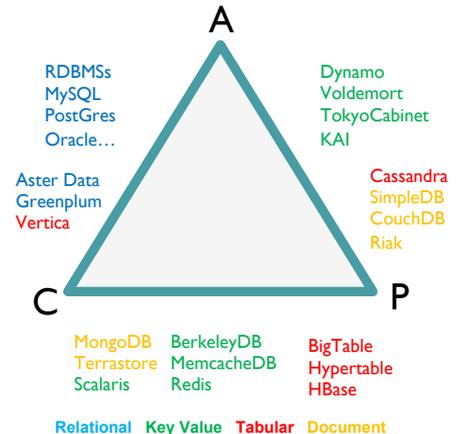
# Protection & Durability

- ▶ Distributed storage uses a variety of techniques
  - ◆ Standard RAID technologies
  - ◆ Mirroring & replication
    - > 2 or 3 location copies
  - ◆ Erasure coding
  - ◆ For a detailed Q&A on these techniques and an on-demand introductory webcast visit:  
<http://sniaesfblog.org/storage-basics-qa-and-no-ones-pride-was-hurt/>
- ▶ Or CAP
  - ◆ Consistent, Available, Partitioned; pick 2



# Future of Distributed Storage

- ▶ Distributed storage offers new & interesting solutions
- ▶ New database technologies
  - ◆ NoSQL, key/value, tabular, document...
- ▶ On-disk compute
  - ◆ Key/value stores directly on the drive
  - ◆ Processing on the drive
    - › Data classification, analysis, automated metadata
  - ◆ Brought together by “consolidating” applications
- ▶ IoT (Internet of Things)
  - ◆ Big data generators
  - ◆ Data at the edge



# Status Check - Midway Summary

- ◆ Centralized
  - ◆ More efficient storage utilization
  - ◆ Simpler storage management
- ◆ Distributed
  - ◆ Scales out, not up
  - ◆ Latency a secondary consideration



# Bring It On!

- › So... what are the trade-offs?



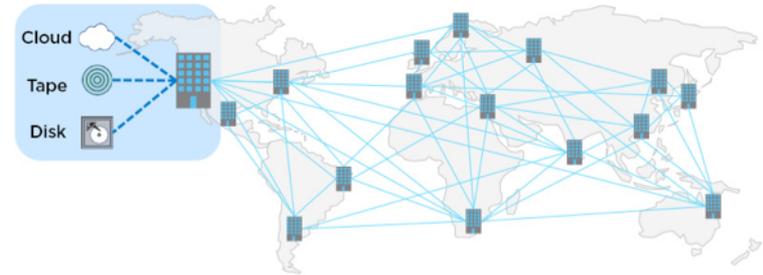
# Is Data Locality Really Important?



- ◆ Centralized Storage
  - ◆ Need servers and storage in same data center
  - ◆ WAN links = too much latency
  - ◆ Install storage near users (i.e. ROBO, cloud)
- ◆ Object and file can support remote access
  - ◆ But then usually set up as distributed storage

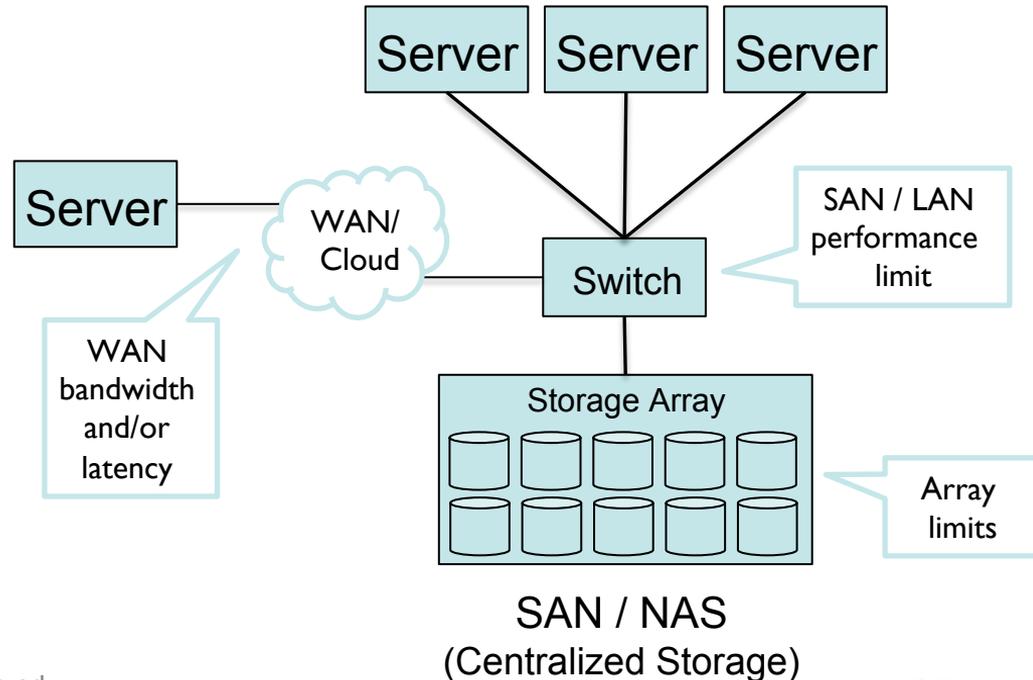
# Is Data Locality Really Important?

- ◇ Distributed Storage
  - ◆ At scale, data locality hard to achieve
  - ◆ Data has mass & inertia
    - › Easiest to process where it's born, centralize the summaries
    - › Partial compute at the edge
  - ◆ New technologies prevent extreme centralization
    - › IoT, blockchain & distributed ledgers, datatypes like video & image, etc.



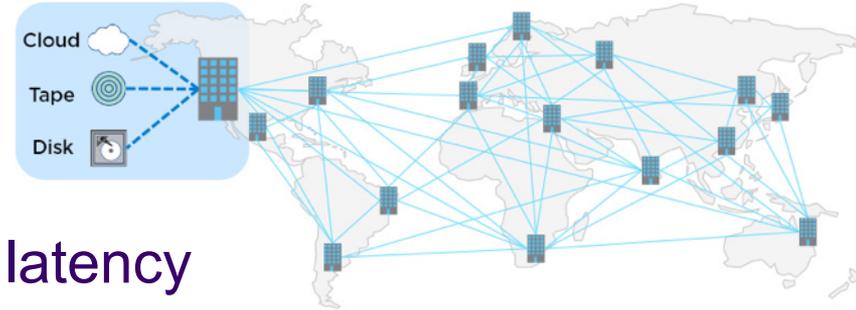
# How to Scale Centralized?

- ▶ Performance scaling
  - ◆ Array performance limits
  - ◆ Network limits
  - ◆ May require locality
- ▶ Capacity Scaling
  - ◆ Adding more arrays
  - ◆ Management burden



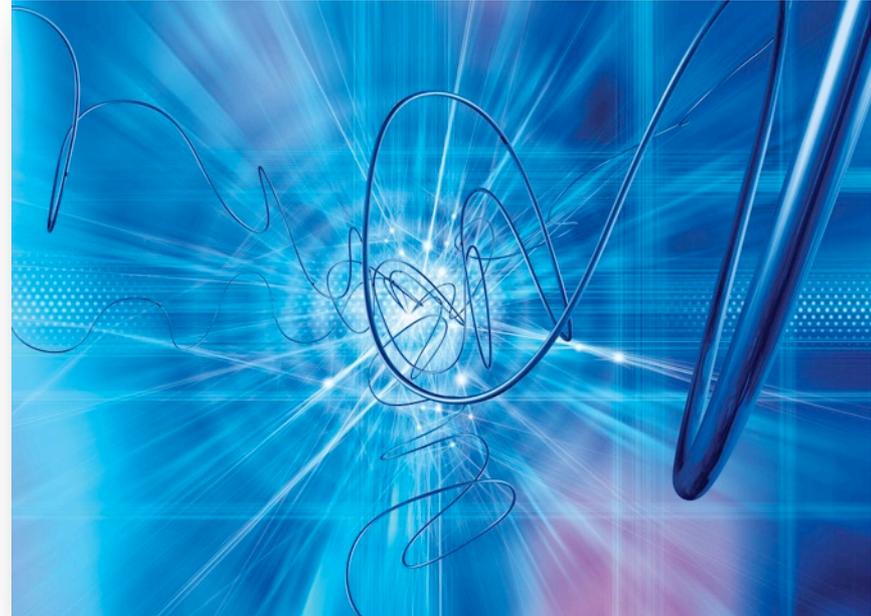
# How to Scale Distributed?

- ▶ Just add more!
- ▶ Limits of scaling may constrain the solution
  - ◆ Economics: cost, bandwidth, latency
  - ◆ Legal: data placement & security
  - ◆ Technical: bandwidth, latency
- ▶ Application plays a part
  - ◆ Not all distributed systems can scale out to infinity
  - ◆ CAP limitations ensure that



# Shared Resources

- ▶ Centralized
  - ◆ Arrays not shared
  - ◆ Network & admins sometimes shared
  - ◆ Might share management tools
- ▶ Different arrays for different workloads
  - ◆ More flexibility in features
  - ◆ Extra management headaches





- › Distributed
  - ◆ Data location is a moveable feast
    - › Backups, mirroring, sharding
  - ◆ Recovery scenarios can be complex
    - › Who & what is impacted by failure & restores?
  - ◆ Fully understand security & privacy
    - › Authentication & authorization
    - › Safe Harbor & GDPR important here
  - ◆ Impacts on performance & capability
    - › “Noisy neighbors”

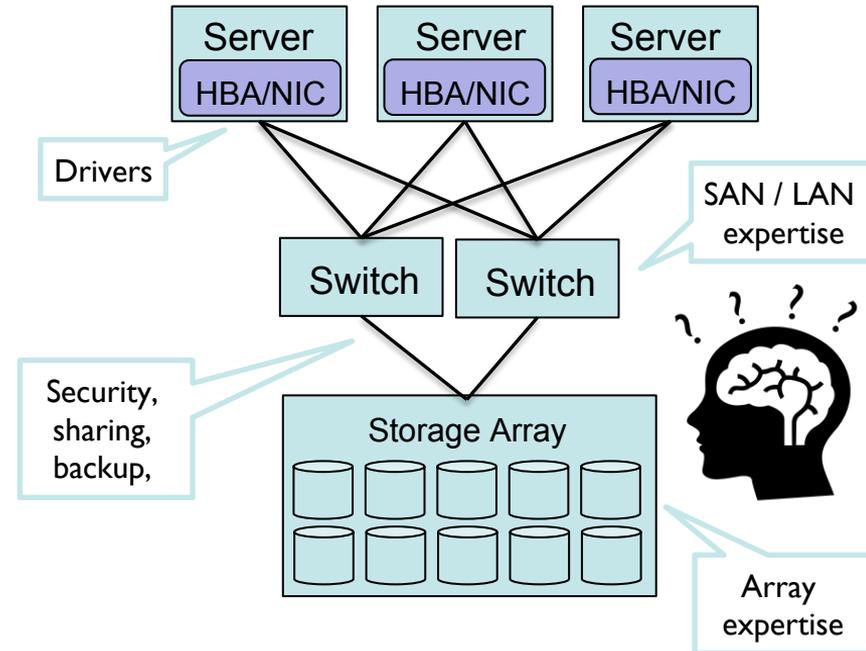
## ◆ Centralized Storage

### ◆ Complex to deploy, manage

- ◆ Need reliable network
- ◆ Might need special drivers
- ◆ Array/network mgmt. skills
- ◆ Security

### ◆ Challenges at large scale

- ◆ Managing many arrays
- ◆ Balancing capacity & workloads
- ◆ May be difficult to automate



- ◆ Distributed Storage
  - ◆ A range of tools
    - › Installation & sizing tools
    - › Capacity, performance, application usage, user usage, chargeback & showback
    - › OpenStack, Docker, Kubernetes...
  - ◆ Offer management consoles & dashboards
    - › Software defined configurations
    - › Compute, network & storage virtualization on one pane of glass
- ◆ New DevOps tools “understand” applications
  - › Ansible, Chef, Puppet...
- ◆ Issues:
  - › Data amnesia; forgetting what was put where it is a big issue
  - › Data migration from system to system can be a challenge
  - › Data can suffer from “container lock in”
  - › Many dashboards are product specific & can be incompatible with each other
  - › Too much choice in DevOps tools?



# What's the Cost/Economic Profile?



## Centralized

- ◆ Usually custom (bespoke) hardware
  - › Dedicated storage platforms
  - › Often uses dedicated network
  - › Less likely to be SDS or cloud
- ◆ More likely to be Cap/Ex
  - › Op/Ex model available through leasing, cloud

# What's the Cost/Economic Profile?

- ◆ Distributed
  - ◆ Cap/Ex or Op/Ex? - “It’s the economy, stupid!”
  - ◆ Cost is a big factor
    - ◆ Consider a longer term cost profile
    - ◆ Largely due to scale
    - ◆ Future unknown, but historical \$/byte cost has fallen pretty consistently
  - ◆ For applications to be of value, their cost components have to be manageable and smaller than the benefits
  - ◆ Pressure of
    - ◆ Systems management costs
    - ◆ New application models (like container, serverless)



# Backup and Data Protection

## ◆ Centralized

- ◆ Easy to backup, fast restores
  - › A big reason to go centralized
- ◆ Usually includes RAID, snapshots, clones
- ◆ Replication and remote backup options
  - › To local system, remote system, or the cloud



Snapshots: point-in-time copies of your data



Storage clones: start identical, change over time

- ◆ Distributed
  - ◆ Backup can be harder
    - › Backup implies a complete redundant copy
    - › Remember CAP & eventually consistent
  - ◆ Durability
    - › Not all data needs to be durable
    - › But when it must be, avoiding “bit rot” & “device obsolescence” requires data to be moved
  - ◆ Long term data retention especially an issue
    - Register for: “The 100 Year Archive Survey Results”  
October 10, 2018  
<https://www.brighttalk.com/webcast/663/335255>

...**backup**

# Debate Summary



- ◆ Centralized makes each array the center of attention
  - ◆ Each array handles backup, security, management
  - ◆ At scale, requires lots of attention, management
- ◆ Distributed spreads performance and capacity across multiple systems
  - ◆ Easy scalability, often lower costs
  - ◆ Security and backup can be more complex
- ◆ Both ways have advantages

- ◆ Other Great Storage Debates
  - ◆ FCoE vs. iSCSI vs. iSER  
<https://www.brighttalk.com/webcast/663/318003>
  - ◆ Fibre Channel vs. iSCSI:  
<https://www.brighttalk.com/webcast/663/297837>
  - ◆ File vs. Block vs. Object Storage:  
<https://www.brighttalk.com/webcast/663/308609>
  - ◆ RoCE vs. iWARP: <https://www.brighttalk.com/webcast/663/329518>
- ◆ On-Demand “Everything You Wanted To Know About Storage But Were Too Proud To Ask” Series
  - ◆ <https://www.snia.org/forums/esf/knowledge/webcasts-topics>

# After This Webcast

- ▶ Please rate this webcast and provide us with feedback
- ▶ This webcast and a PDF of the slides will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand at [www.snia.org/forums/esf/knowledge/webcasts](http://www.snia.org/forums/esf/knowledge/webcasts)
- ▶ A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog: [sniaesfblog.org](http://sniaesfblog.org)
- ▶ Follow us on Twitter @SNIAESF

**Thank You!**