




Go further, faster®



Top 10 Things We Learned Implementing OpenStack

Val Bercovici

*Cloud Czar
NetApp*





HELLO

My name is

Val Bercovici

[@valb00](https://twitter.com/valb00)



Solid State
Storage



Cloud
Storage



2012
HADOOP
SUMMIT



What I hope to accomplish today...

Help gain perspective
beyond current hype cycle

&


Make OpenStack mainstream
via *Biz & Tech Context*



Inspiration → 360° Perspective



Customer




Industry



Workloads




Competition



Ecosystem

#10 Cloud Geography

- Silicon Valley(s)
- USA (Middle)
- The rest of the world
- <http://it20.info/2012/09/cloud-and-the-three-it-geographies-silicon-valley-us-and-rest-of-the-world/>




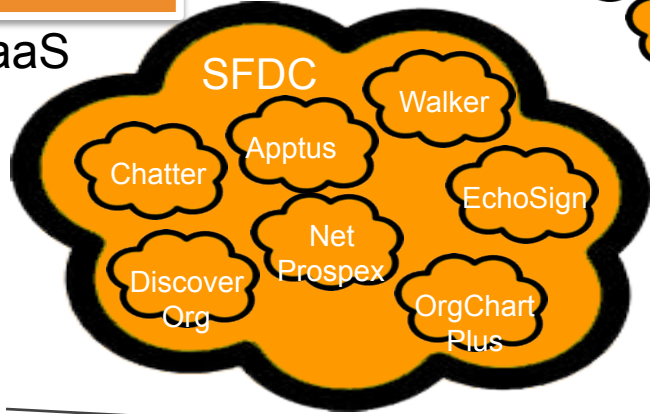


#9 Petri Dish

NetApp®


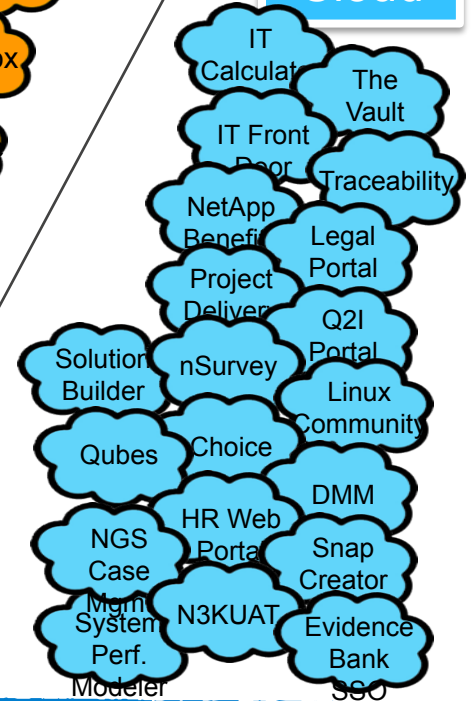
Public Cloud

SaaS



IaaS
PaaS

Dev Cloud



App Cloud



#8 Industry Diversity: Biz → R&D Strategy

Buy & Operate



Deutsche Telekom



Leverage Open Source



HP Enterprise Services

Hyperscalers

Google



amazon

Windows Azure




NetApp®


#7 Tech Diversity: Countervailing Forces



Variety of Consumption Models



Rack – Appliance – Controller – Software – Instance



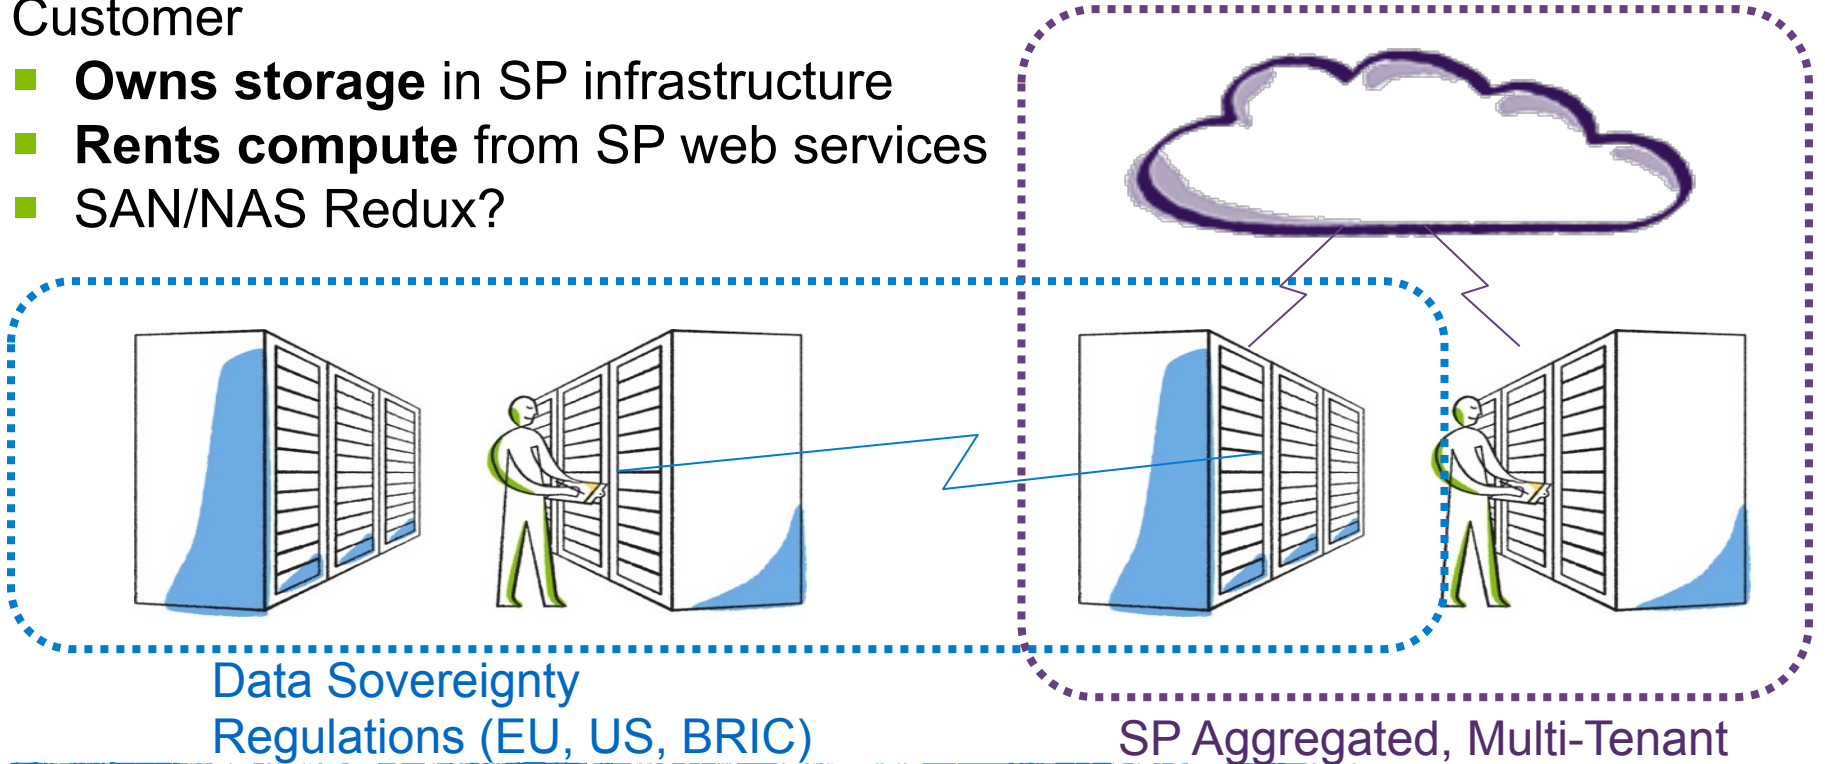


NetApp®

#6 Hybrid Cloud: Storage-Centric

Customer

- Owns storage in SP infrastructure
- Rents compute from SP web services
- SAN/NAS Redux?





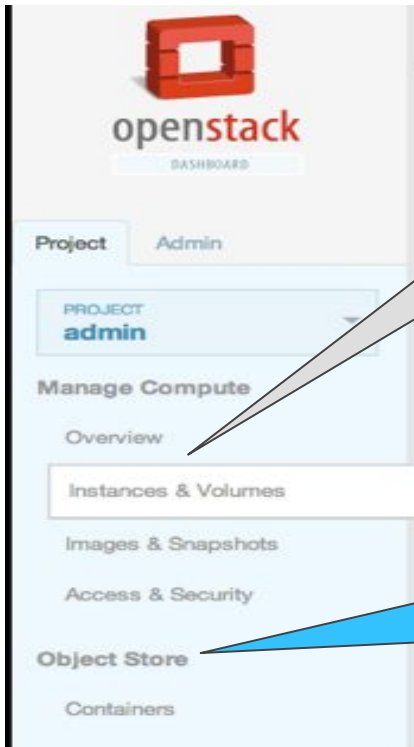
#5 Integration: Continuous vs Staged

- Sushi Fail Whale vs Gourmet Cooking
- Surgery
- Air{Plane|Traffic}
- Automotive Ctrl
- Military
- Nuclear Power
- Wall St






#4 OpenStack Storage Duality



COMPUTE
(Apps)
Abstraction

OBJECT
(Data)
Interface &
Implementation





NetApp®

#3 – Diversity of Dev & Ops

App Platforms



PaaS platforms bundle wide range of data services with pure app middleware features

Amazon


New app platforms are offering flexible data models for:

- Operational Efficiency
- Performance at scale

Data Services




Infra Services



Infra as a Service platforms offer storage management for free/fee



Infra



METADATA = Volumes
DATABASE = Volumes
OBJECT / RING = Containers

#2 Distributed vs Centralized Control



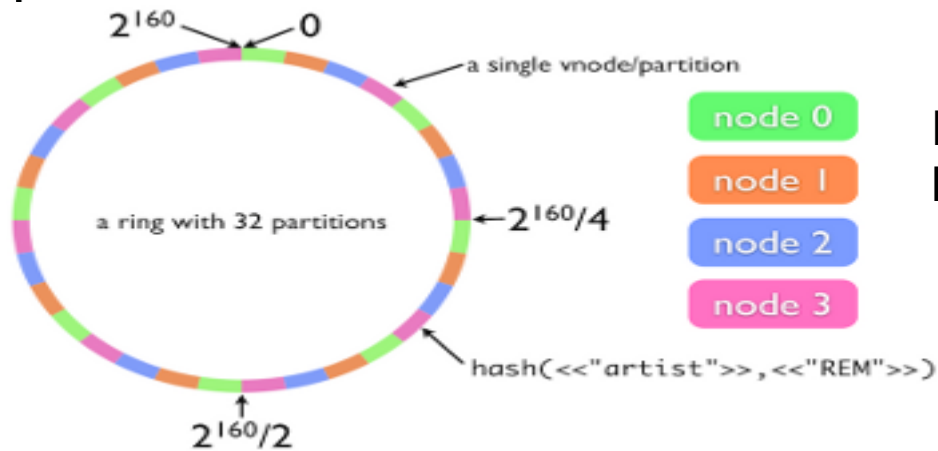



FIG. 1 - Centralized, Decentralized and Distributed Networks

Paul Baran, RAND Corporation (1962)

#1 Resiliency: Erasure Coding vs? RAID



Ideal for Cross-Node Resiliency



Ideal for Disk-shelf Resiliency

Figure 2. An example of one stripe where $k = 4$, $m = 2$ and $w = 4$.

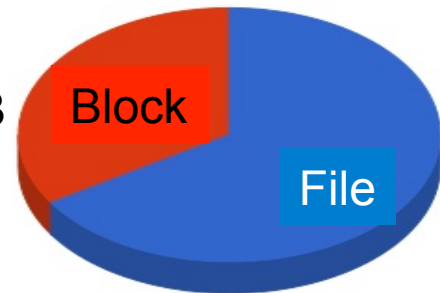


NetApp®

#0 Shares Service

Proposal / Blueprint

IDC
27 ExB
(2012)



Share R&D
between
Guests 1 & 8

File Share Service

Shares Provider

Fin


R&D

Mktg

openstack™
CLOUD SOFTWARE

Nova

Compute Node



https://wiki.openstack.org/wiki/Shares_Service

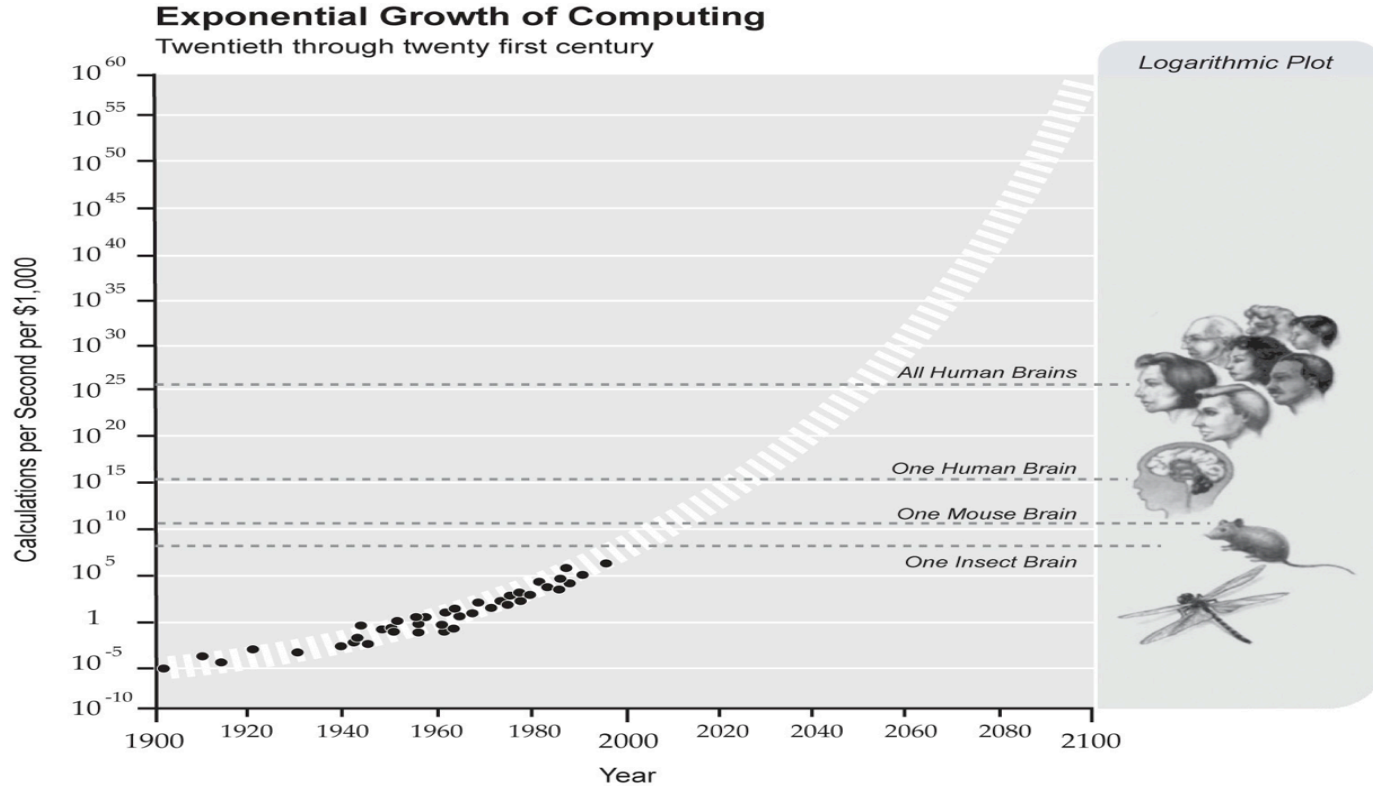


Looking Back in order to
Look *Forward!*



Science (Fiction?)

<http://www.singularity.com/charts/page70.html>





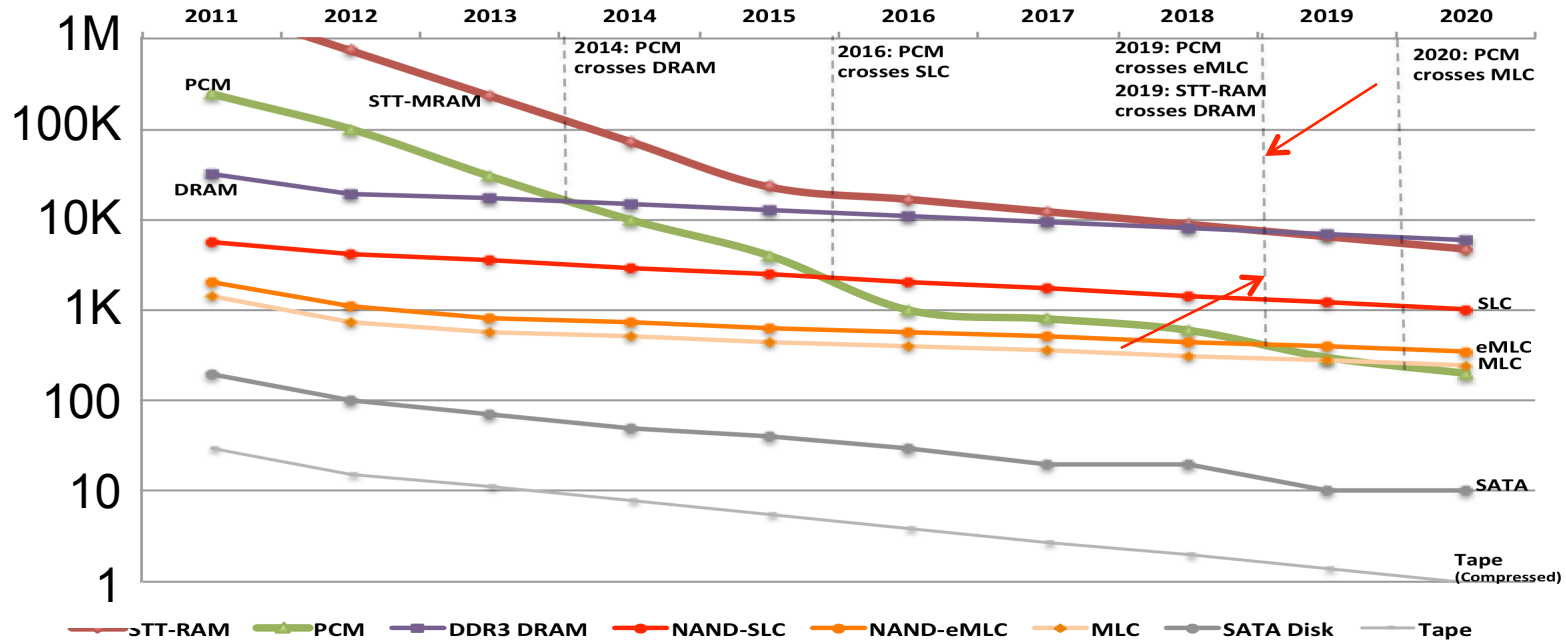
A B C D E F G H ... I?



- Only 17 more letters to go!?
 - We just started the journey... time to consider Khmer
 - it has 74
- What does OpenStack look like in 2 years?
 - Are we building / planning for it?
 - 6 months is a long sprint, but how do we contemplate / incent the projects that take 18?

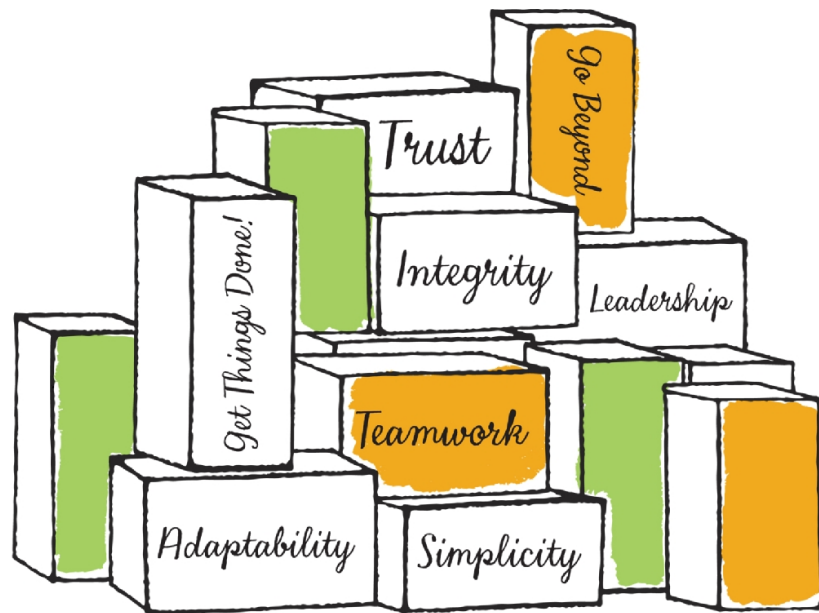


Wither Disk ? and later Wither Flash ??





Thank you





- Austin: The first design summit took place in Austin, TX
- Bexar: The second design summit took place in San Antonio, TX (Bexar county).
- Cactus: Cactus is a city in Texas
- Diablo: Diablo is a city in the bay area near Santa Clara, CA
- Essex: Essex is a city near Boston, MA
- Folsom: Folsom is a city near San Francisco, CA
- Grizzly: Grizzly is an element of the state flag of California (design summit takes place in San Diego, CA)
- Havana: Havana is an unincorporated community in Oregon