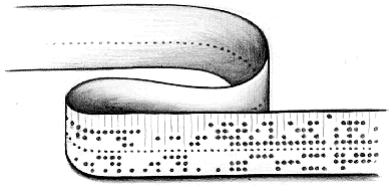

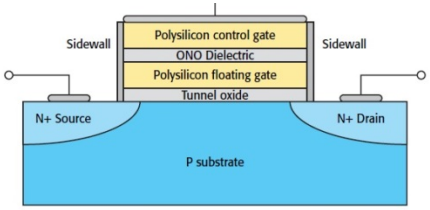

STORAGE MANAGEMENT IN THE ENTERPRISE

MSST.2015 RANDY OLINGER

6/3/2015

WHO IS THIS GUY?

- Currently employed by Optum, a Subsidiary of UnitedHealthGroup.
- 30+ years developing and deploying technology
  - Software development in the early years (hated it)
  - Manufacturing Engineering (loved it)
  - Project and Software development Management (Hated It)
  - Infrastructure Architecture and Engineering (love it)
- Formally Trained as an engineer
- Swore to never work in IT… Until Y2K.

WHO IS OPTUM?

- UnitedHealth Group
  - Over 130,000 employees)
  - $130B in revenue (2014)
- United Healthcare – Health Insurance for groups and individuals. (underwriting, benefits, plans, etc.)
- Optum – Various health services for UHC and others. (Nurse lines, risk abatement, Health Analytics, Pharmacy, Professional Services, Hosting)

THINGS TO BE COVERED… (AGENDA)

(IN NO PARTICULAR ORDER….)

- Enterprise Scale
- Enterprise Storage Architecture
- Getting Big
- Embracing New Stuff
- Three P's
- Objects
- Open Source Approaches
- Q and A

ENTERPRISE SCALE – HUH?

- Not quite HyperScale, whatever that is.
- Probably similar to Cloud Scale, but with more diversity
- 100's of PB of block, File, Object. Big Budgets, aggressive timelines
- You name the brand, we have it
- Embracing more Open Source as a way to break free of lock-in
- Embracing Commodity components to try and bend the cost curve
- Intense regulations (HIPPA, SOX, ARRA, ACA)
- Seperation of payers from consumers

TRADITIONAL ARCHITECTURE

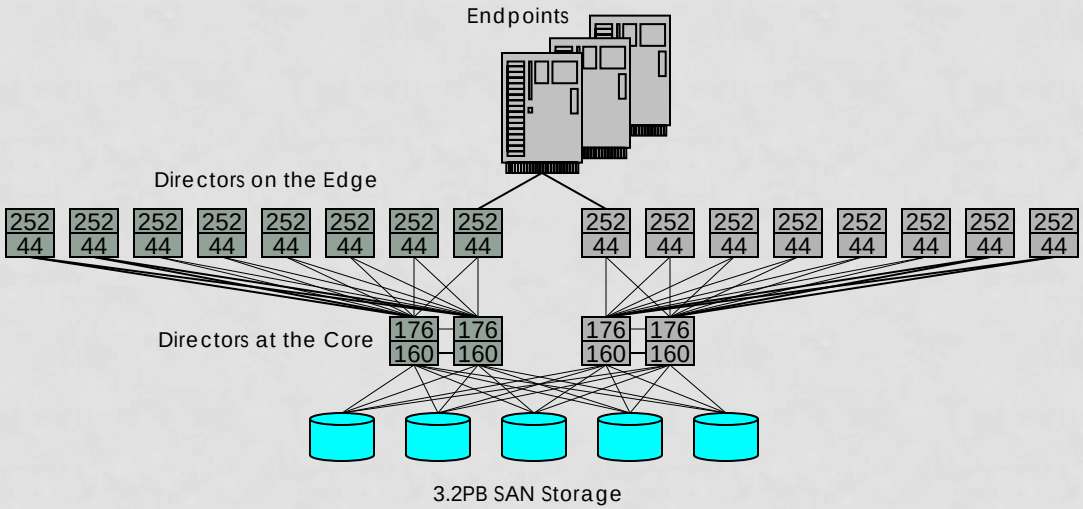

2000 Server SAN Fabric	HOST	ISL
Ports/Connections	2016	352
2x IO Load per Connection	1.3	8
Total IO Load	2621	2816
Sustained IO Load	**1.30**	**1.40**

Specifications

Nbr Attached Hosts	2,016
Nbr Directors	20
Nbr Useable Ports	6,080
Gross Gb/s Thruput	35,840
Total Watts	28,400

Green = Meets our ELR Host Gb/s IO Requirements
Yellow – Provides x2 our Current Host Gb/s IO
Red – Fails to Provide x2 our Current Host Gb/s IO

- Predictable Throughput
- Heavily Engineered and Monitored
- Reliable to 99.9999% (6 nines)
- Expensive, but worth it?
- Does not facilitate sharing

COMMON THEMES

- Organizational
  - Silos – Compute - Storage - Network - Database - Security
- Infrastructure Engineering
  - Each Silo will feather their own bed, we call that optimizing.
  - Need to keep it from getting too project focused
- Conflicts Abound, but can be healthy
- Growing by acquisition
  - 50 transactions per year
  - Sometimes big, Sometimes small
  - Many are on fire (which is why they were acquired)
  - Must be remediated as they can damage the brand.

HOW DO YOU KNOW WHEN YOU ARE GETTING "BIG"?

- Information presented at vendor conferences no longer applies to you
- Capital spending plan needs to specify "All figures are in $M"
- Volume of email from vendors you have never heard of eclipses your attention span
- Constantly fighting over space, power, and installation timelines
- Purchases become more strategic than project oriented.

WHAT TO DO WHEN YOU GET "BIG"

- Of course, automation is key.
- However, don't automate already broken processes, you just go crazy faster.
- Don't be afraid to oversubscribe capacity and performance. Users are hoarders.
- If you have many isolated zones of capacity, develop a scaling plan for new purchases.
  - Buy the bones and install capacity later
  - Squatters Rights in the data Center
- Hold operational staff headcount stable. Expand Engineering and solutions headcount.
- Probably time to "let go" and centralize provisioning and capacity management.

IRREGULAR CAPITAL SPEND

FEAST OR FAMINE

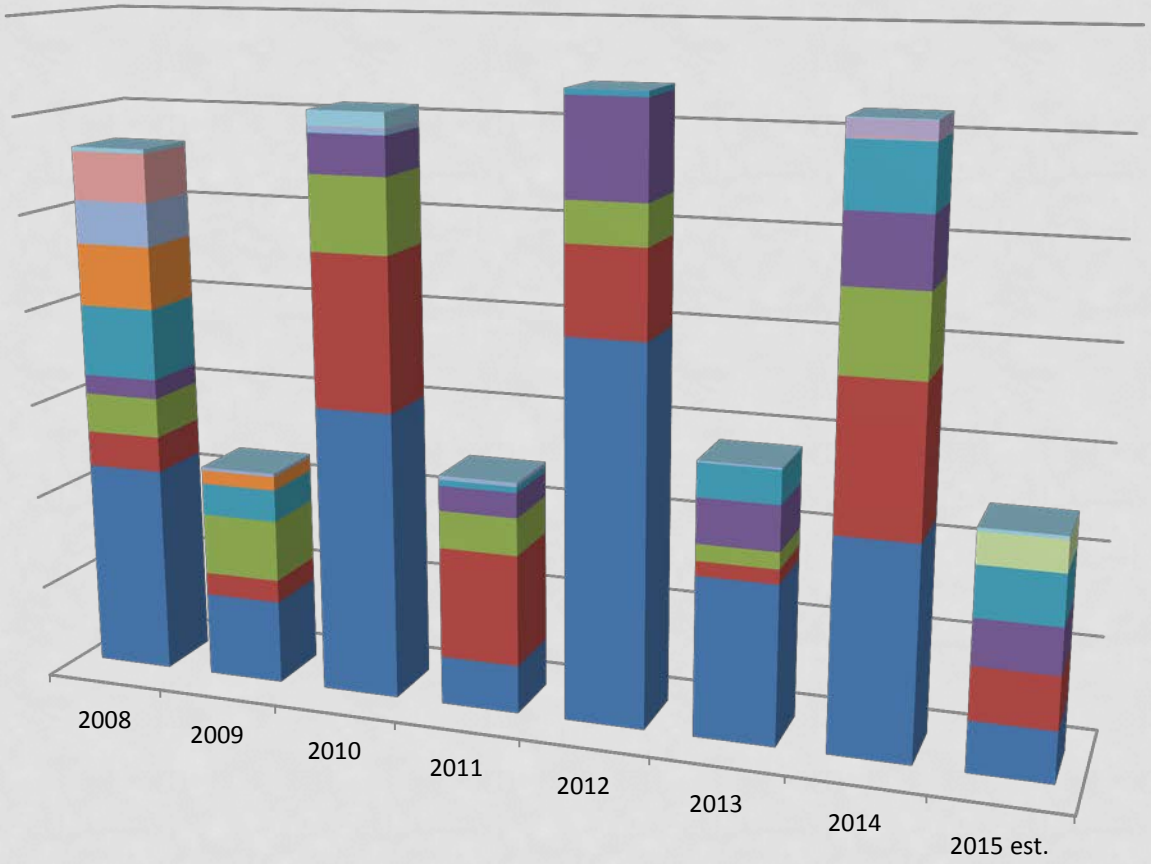

You can be on a 24 month capital schedule, and not even know it...

AUTOMATION…

- We started with our biggest pain point. Self Service ability for "DBA wants more storage"
  - Swap out DBA for VMWare, SAS, e-mail, etc.
  - Turns out that most expansion transactions need very little if any engineering, hard part is already done.
  - 18 day process reduced to 10 minutes
- Doing just part of the job is OK, but probably won't solve the real problem.
- It does, however, clarify where the real problem is.
- We were doing hundreds of these allocations per week, 5 full time people. Reduced to 1.

WHAT IF YOU HAVE TO REDUCE ?

- In Storage? I'll believe it when I see it…
- Actually is a real problem if workload shifts to other hosting facilities or dramatically changes (like losing a customer)
- Make sure you have physical isolation down at the component level, allowing you to re-deploy partial assets. (especially expensive items like flash media)
- Workload mobility can really be an asset here. Virtualized Workloads shine when it comes to elasticity.

WHAT ABOUT NEW STUFF?

- OpenStack
- Do we hate VMWare? Absolutely not. A fine product for the right workloads.
- OpenStack being evaluated from a cost and capability model to see if it can supplement the virtual environment.
- Ultimately it may be more suited for "Cloud Ready" applications
  - No pets allowed, this is cattle country
  - Scale out, resilient, automated deploy and destroy

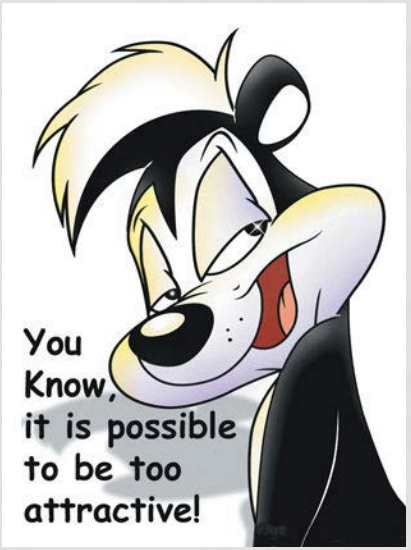

WHAT ABOUT MORE NEW STUFF?

- OpenCompute!
- Do we hate our vendors? Not at all! They just cost too much.
- Looks a lot like HyperConverged offerings in the VMWare space.
- Taking advantage of the commodity Hardware and Software stacks available.
- Trading Capex for Opex, so it only works at scale of hundreds or thousands of servers.
- No great solution for performance storage.

WHAT ABOUT OLD STUFF?

- Tier 1 arrays from established vendors still a viable play
  - Significant competition.
  - Feature rich (Tiers, replication, mobility, non-stop, Thin)
  - Performance Centric with mature tools
- Tier 2 arrays never were a viable play at scale
  - Purchase price usually not much lower than tier 1
  - Scalability limitations
  - Performance and Availability limitations
  - Specific use cases would be the exception, like backup.
- File Servers
  - Yawn…
  - limitations in the NFS and CIFS protocols
  - NFS V4 and CIFS3 are helping

WHAT ABOUT CLOUD STORAGE?

- Very viable for replacing archive systems
- Users should become very familiar with access patterns compared to cost drivers.
  - Some vendors charge per IOP
  - Ingress / Egress charges
  - Snapshots and replication
- Price wars continue to benefit the consumers
- Easy to get in, hard to get out
  - Create your scorched earth plan before putting your first byte in the cloud
  - Remember Nirvanix? Yeah, neither do I...

THE 3 "P"S - THE PERPLEXING PROBLEM OF PRODUCT PROCUREMENT

Product	• This is the process of evaluating the merits of the proposed technology to solve the problem i.e. Requirements • A Bugatti Veyron with a top speed of 267.856 mph isn't any better than a 90 mph SmartCar if the requirement is city driving
Price	• The Price is the Price. • Total Cost of Ownership matters • Beware of the special one time deal
Politics	• Important things that influence a deal that are not directly related to the merits of the Product or the Price • Bias, Personal Preferences, Business Factors, other dealings unrelated to the deal at hand, etc. HiPPO.

ENGINEERING CULTURE SHIFT

- Classic - Ready Ready Ready Aim Aim Aim Aim Aim Ready Ready Aim Ready Aim Aim Fire!
- Iterate – Ready Aim Fire Ready Aim Fire Ready Aim Fire
- Sounds great but difficult to scale. Can lead to Attention Fragmentation (multitasking)
- Adopting Kan Ban as a project discipline
- Depends on accepting failure, and failing fast with minimal investment
- Striving to adopt Lean Startup principles in a large organization

WHAT ABOUT OBJECTS?

- Love the attributes (durability, cost, availability, scalability)
- REST interface is not popular in the enterprise – requires code changes.
- The applications are coming
- Used natively, Object Storage is attractive
  - Erasure coding
  - Data Dispersion
  - Low Cost Hardware
  - Reasonable Throughput

OPEN OPEN OPEN

- Open source projects now the darling of the enterprise
  - Open as in “Free Speech” – Absolutely!
  - Open as in “Free Beer” – When it makes sense
- Can we substitute labor for capital?
  - Open Source projects usually take a different type of engineer to install and manage.
  - Everything is available, from databases , java app servers, service bus, security, communications, project management, business intelligence, you name it.
- Start SMALL and move carefully. Things will go bump in the night, and support may not be available…

OPEN SOURCE SUCCESS - PAAS

- Platform as a Service looks like a clear winner for Open Source
  - Packages available from Ubuntu, Redhat, and others.
  - Suited for Web applications that scale out
  - Ruby, PHP, Java, Python, Go, Node.JS, and many others.
- Not so easy to figure out how to charge customers.
  - Somebody has to pay for it
  - Cloud means metered consumption
  - Pay attention to utilization monitoring
- Still need to solve classic problems like HA, DR, data protection

OPEN SOURCE SUCCESS - RDBMS

- MySQL, MariaDB, Percona – Pick one
- Simple to install, easy to manage
- So cute when they are small… Then they grow up
- Make sure you have tools for query monitoring, Space utilization, Index Management
- Backup is no less a problem than with other common enterprise databases

OPEN SOURCE SUCCESS – BI

- Statistics and Reporting applications have Open Source candidates that have been around for many years.
- Hadoop + MapReduce + R seems to be a good combination
- Visualization packages are getting acquired by the big guys (Jaspersoft, Pentaho)
  - Not a bad thing if it brings legitimacy to the environment
  - Community versions can bring a lot of value
  - Beware the bait and switch

SECURITY

- Too many cases of data theft amongst enterprises lately. Not a good trend.
- C-level executives in all industries are acutely aware of the consequences.
- Despite our best efforts, Health Care fraud is still prevalent, particularly in the public sector.
  - Stolen patient data is unique enabler for this type of criminal activity.
  - Encryption, two-factor, IPS, IDS, forensic Analysis, Application Authorization – All need to be audited.
- Compromise the Backup Administrator

THANK YOU

Contact me at Randy.Olinger@Optum.com
@randyolinger