Data Centre Forecast

How Cloud Computing Impacts Data Center Design

Global Data Center Engineering

Created by:
Global Data Center Engineering

Cloud Computing – What is it?

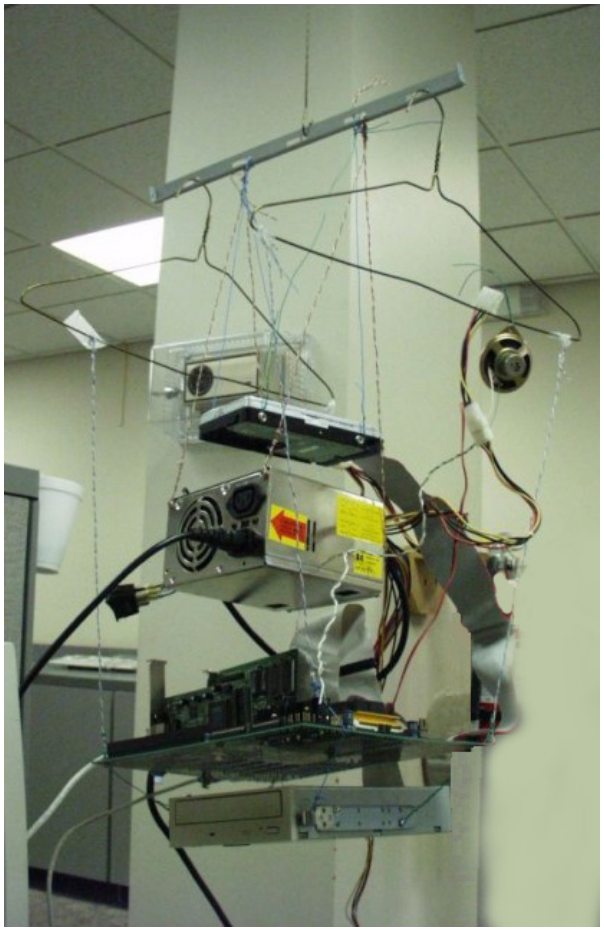

If Your Cloud Computing Initiative Looks like this…

You're Doing It Wrong

Why Cloud Isn't Cloudy

- The mystery of the Cloud is rapidly dissolving
- No, the Cloud is not new, but it is now more accessible to the masses
- Deployment of applications and services happens in minutes instead of days or weeks
- It provides access to what we need, when we need it, and how we want it

Yes, Now – IT as a Service

- Demand for instant answers to complex questions has become the standard
- Enterprise (from tiny to colossal) have low cost access to substantial computing muscle
- Pace of evolution from legacy IT infrastructure continues to increase
- Private, Public, Hybrid and Vertical Clouds are linking, providing clear visibility through fog

The CPU Transistors and Moore's Law at Work

Day	Rate	Total Pay	
1	$0.01	$0.01	
2	$0.02	$0.03	
3	$0.04	$0.07	
4	$0.08	$0.15	
5	$0.16	$0.31	
6	$0.32	$0.63	
7	$0.64	$1.27	Week 1
8	$1.28	$2.55	
9	$2.56	$5.11	
10	$5.12	$10.23	
11	$10.24	$20.47	
12	$20.48	$40.95	
13	$40.96	$81.91	
14	$81.92	$163.83	Week 2
15	$163.84	$327.67	
16	$327.68	$655.35	
17	$655.36	$1,310.71	
18	$1,310.72	$2,621.43	
19	$2,621.44	$5,242.87	
20	$5,242.88	$10,485.75	
21	$10,485.76	$20,971.51	Week 3
22	$20,971.52	$41,943.03	
23	$41,943.04	$83,886.07	
24	$83,886.08	$167,772.15	
25	$167,772.16	$335,544.31	
26	$335,544.32	$671,088.63	
27	$671,088.64	$1,342,177.27	
28	$1,342,177.28	$2,684,354.55	Week 4

$2,684,354.55

Year	Transistor Count	Commercial Comparison (Intel)
1965	280	First IC Etching ~280 transistors
1967	560	
1969	1,120	
1971	2,240	Intel 4004 first single chip CPU
1973	4,480	8008
1975	8,960	8080
1977	17,920	8086
1979	35,840	8088
1981	71,680	80186
1983	143,360	80286 (1982)
1985	286,720	80386 (1986)
1987	573,440	80386 (1987 version)
1989	1,146,880	80486
1991	2,293,760	80486SX
1993	4,587,520	Pentium I (1994)
1995	9,175,040	Pentium Pro
1997	18,350,080	Pentium II
1999	36,700,160	Pentium III
2001	73,400,320	Pentium IV (2000)
2003	146,800,640	Pentium IV
2005	293,601,280	Core2Duo
2007	587,202,560	Core2Quad
2009	1,174,405,120	i7 Extreme
2011	2,348,810,240	Xeon E7 2870
2013	4,697,620,480	Xeon E7 8893
2015	9,395,240,960	<Future>
2017	18,790,481,920	<Future>
2019	37,580,963,840	<Future>

8,388,608x More Powerful than 1965

CPU Transistor Grown and Impact

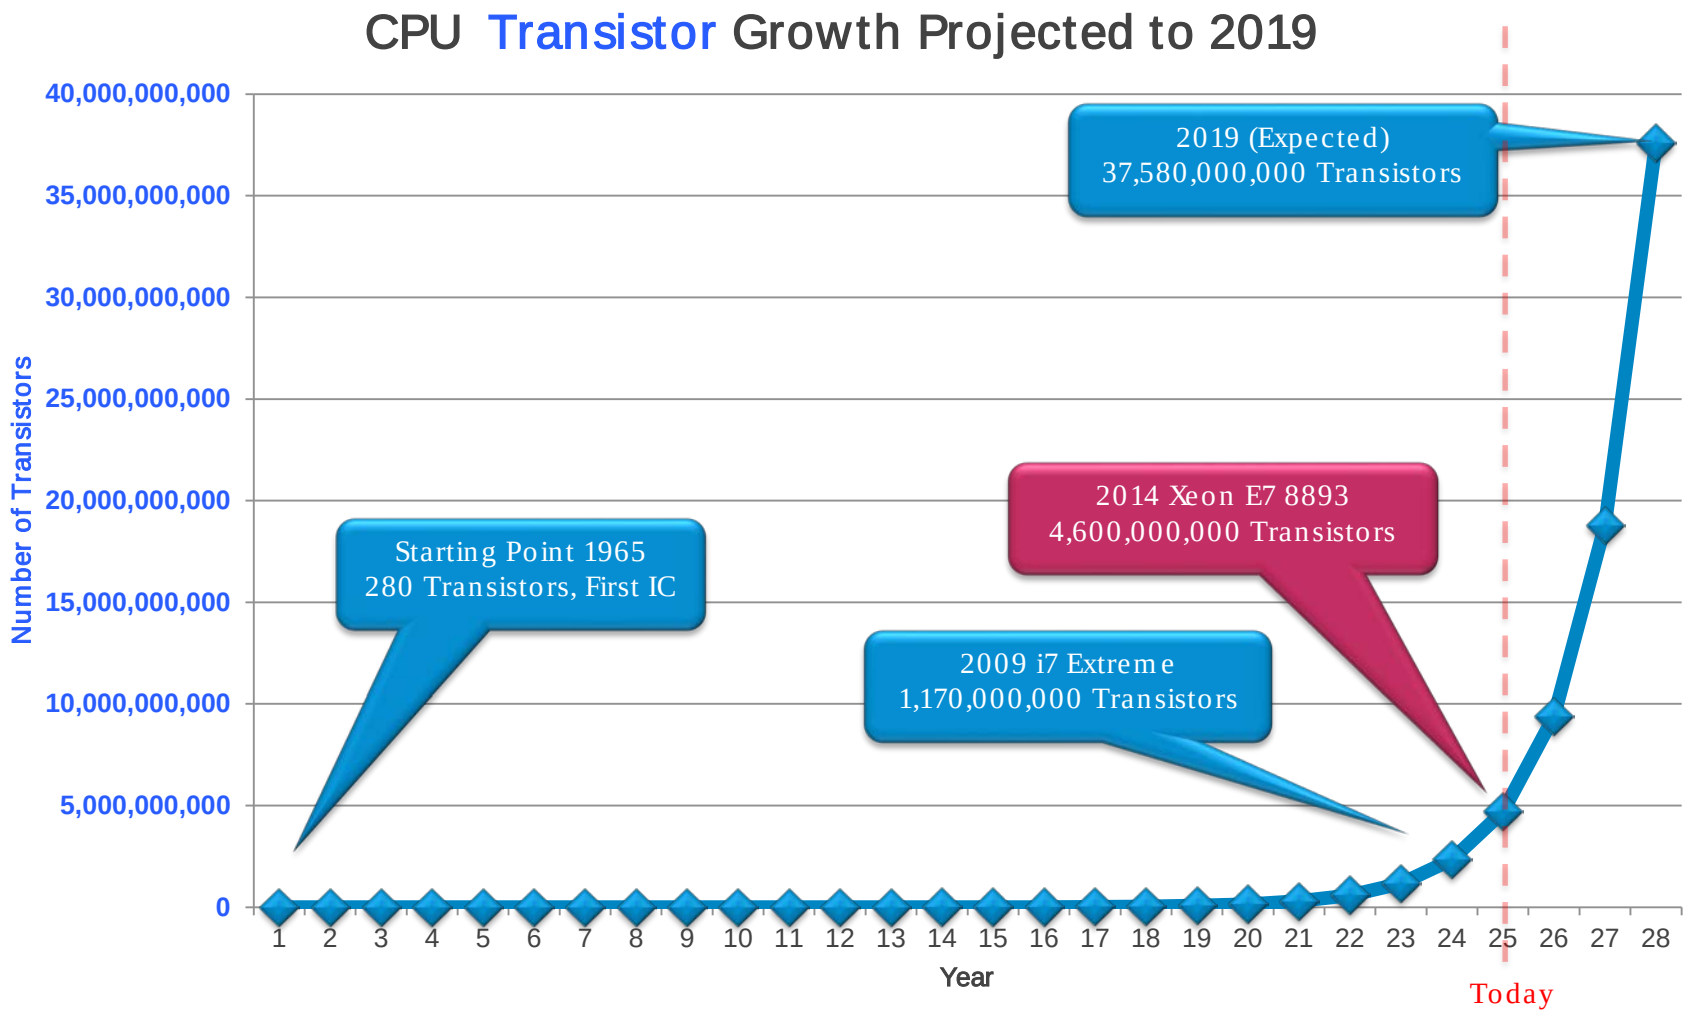

CPU – Bringing It All Together

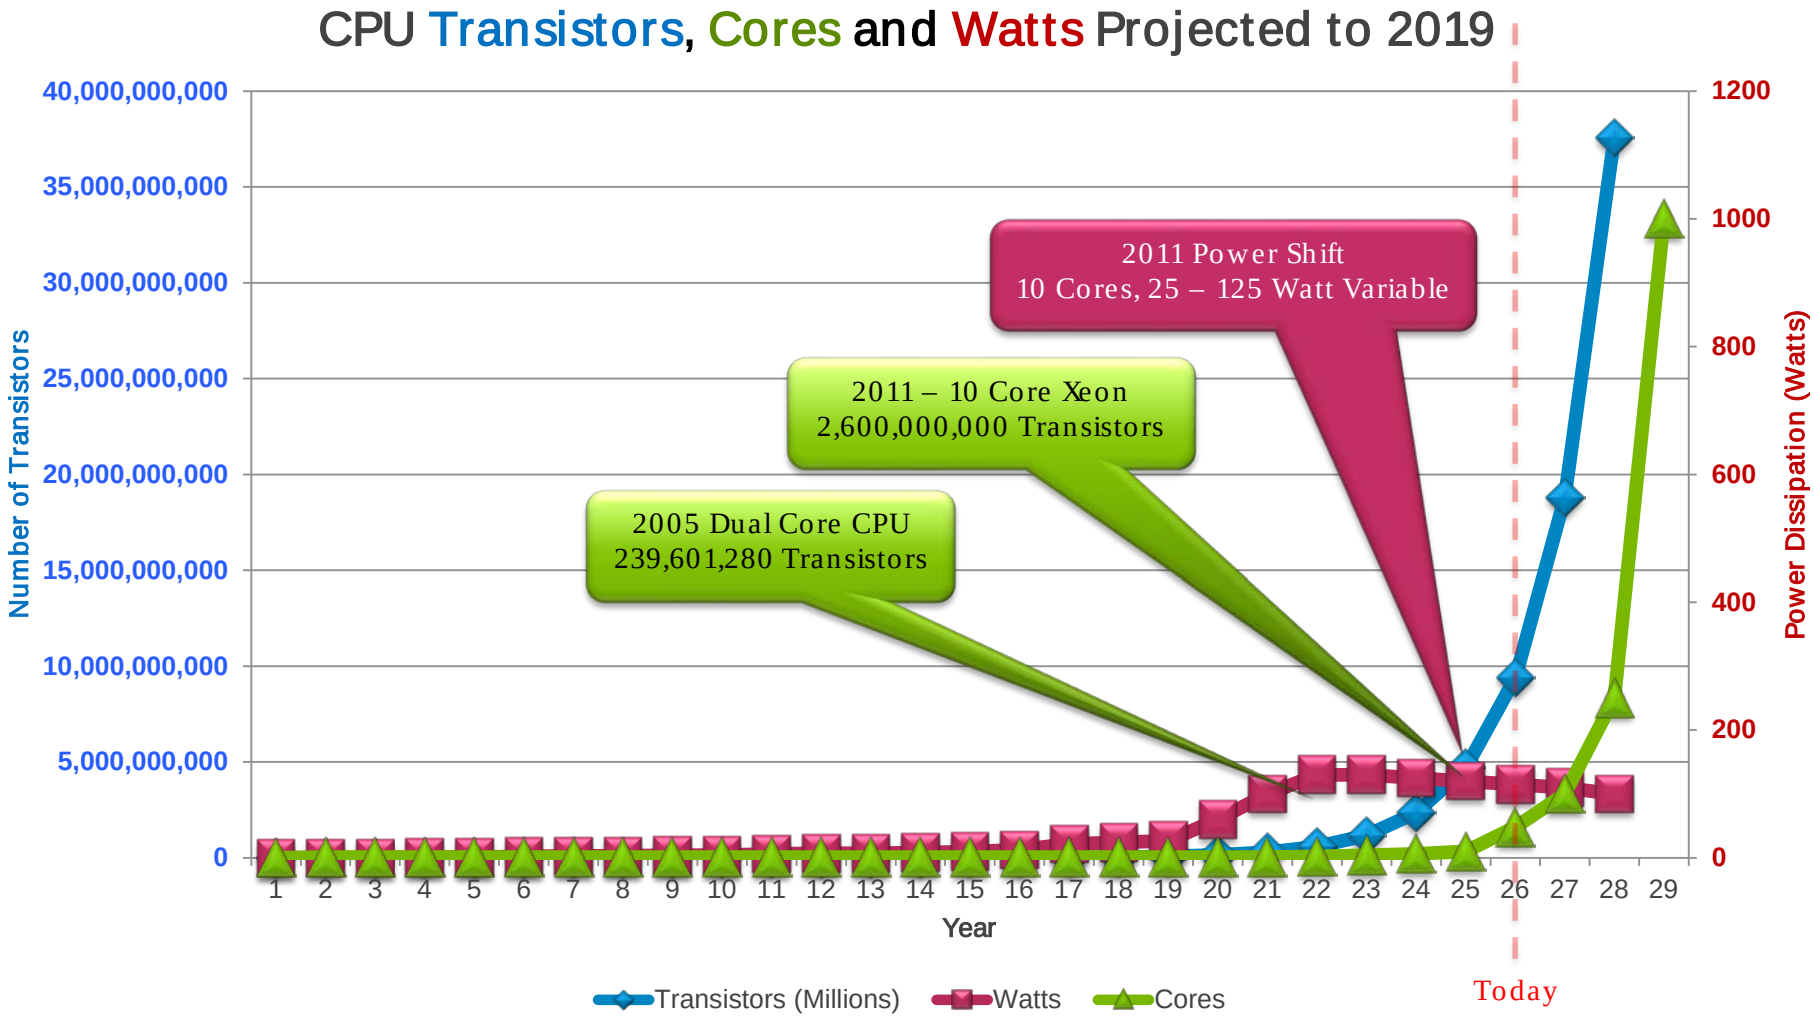

CPU Evolution – Watt Does It Mean

- CPU Innovations Result in:
  - More compute power in less physical space
    - More Cores in a single chip
    - More CPUs supported per board (4 instead of 1 or 2)
  - Improved Energy Efficiency
    - Chips utilize less energy than their predecessors
    - Less heat is generated, requiring less cooling
    - Chips have higher heat tolerance, allowing higher operating temperatures

Space, Heat and Power Dichotomy

1

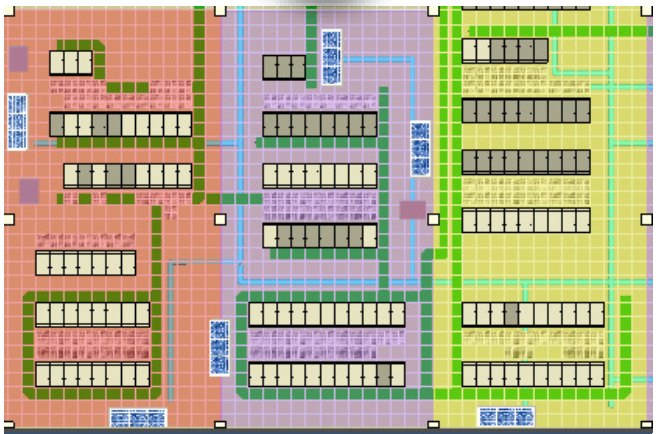

- 600,000 BOPS
- 160 racks
- 1500 servers
- 6400 sq ft
- 560 kW
- 3.5 kW/rack
- 9 Servers per Rack
- TCO / yr: $2.75MM
- TCO / workload: $1810

Don't plan for what you have today…

Transformation Current DC into Private Cloud

2

- 600,000 BOPS
- 16 racks
- 240 servers
- 640 sq ft
- 120 kW
- 7.5 kW/rack
- 15 Servers per Rack
- TCO / yr: $495K
- TCO / workload: $330

Aggressively Optimize first…

Refresh 3-4 YEARS

3

- Highly optimized platform/workload relationship
- Automated provisioning and load distribution
- Scalable, Reliable
- Lifecycle Management

- 600,000 BOPS
- 2 racks
- 40 servers
- 80 sq ft
- 42 kW total
- 21 kW/rack
- 20 Servers per Rack
- TCO / yr: $115K
- TCO / workload: $77

…Design for future and realize returns

Server Trends

- Past
  - In the past, servers and storage were coupled
  - Server consumed 4U (rack Units) or rack space
  - Held one to two CPUs with 2 to 4 Cores
- Present
  - Most servers coupled with boot device
  - Servers consume 1U to 2U rack space
  - Hold one to four CPUs with 2 to 6 Cores
  - Some blade server implementations

Server Trends (Continued)

- Future
  - Microservers (12 + servers in 3U) replace blades
  - Single-Chip Cloud Computers replace DC rooms
  - Storage and boot separates from physical server
  - Memory separates from physical server
  - Processor power becomes 1,000,000 million times more powerful by 2019 than it was in 1980
  - Power to CPU ratio declines significantly as chip cores reach into the thousands on a single CPU

Storage and the Cloud

If your storage grows on trees…

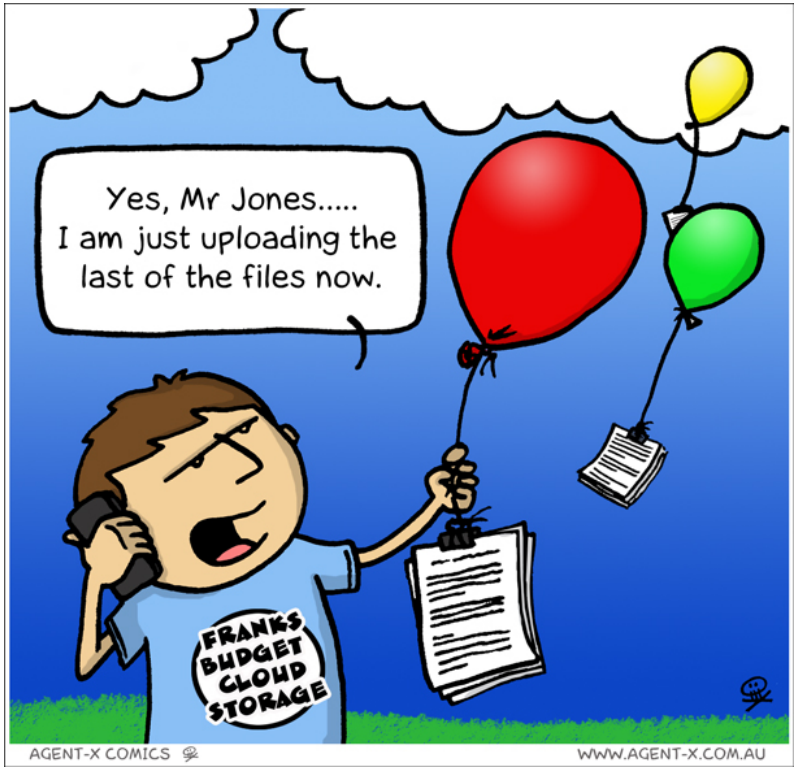

You're Doing It Wrong

The History of Data Storage

- Storage has evolved through a number of stages:
  - Punched Cards & Vacuumed Tubes 1936 (bits)
  - Drum and Core Memory 1951 (Bytes)
  - Disk Can and Floppy Drives 1969 (Kilobytes)
  - Tape and Early Hard Disk 1980 (Megabytes)
  - Hard disk and USB Sticks 1993 (Gigabytes)
  - Modern Hard Disk 2007 (Terabytes)

When Storage Grew on Trees

- Modern base measure for data storage is the Terabyte (1,000 billion bytes or 2^40 bytes = 1,099,511,627,776 bytes)
- Unlike Moore's Law, storage does not follow the same doubling every 2 years.
- Storage does follow a capacity trend:
  - Factor of 10 increase approximately every 14 years
  - That puts petabyte drives on target for 2021

The Direction of Storage Capability

Media Type	Height (mm/in)	Weight (kg/lbs)	Storage Capacity (bytes)	Quantity needed for 1TB	Weight Needed for 1TB (kg/lbs)	Volume to Store	Comparison - Distance	Comparison - Weight
Punched Card	0.178	0.00242	120	9,162,596,898	22,173,484	1630 Kilometers	Hong Kong to Ho Chi Mihn	120 Blue Whales
	0.007	0.005335			48,882,454	994 Miles		
9 Track Tape	12.65	1.1339	178,257,920	6,168	6,994	78 Meter	3 1/4 Tennis Courts end-to-end	Large African Elephant
	0.498	2.5			15,420	256 Feet		
5 1/4 Full Height HDD	82.55	2.2679	640,000,000	1,718	3,896	142 Meter	38 Story Building	Hippopotamus
	3.25	5			8,590	466 Feet		
3 1/2 IN 1TB HDD	26.1	0.622	1,000,000,000,000	1	0.6220	2.6CM	Cheesburger	Hard Back Book 350 Pages
	1.027	1.371			1.3710	1.027IN		
MicroSD	1	0.0002	128,000,000,000	8	0.00160	83.2CM	5 Credit Cards	8 Fingernails (not the pinky)
	0.0393	0.001102311			0.00842	32.75IN		

Since 2007 the ability to generate data has exceeded existing capacity to store it

- A faster and larger storage capacity is required

Storage Trends

- Spinning Hard Disks (HDD) are on the way out
- Advancements in capacity have stalled in recent years as drive manufacturers retool and reinvent for Solid State Drive (SSD)
- Capability for capacity reaching 10's or 100's of Gigabytes is easily achievable
- SSD technology is superior to HDD in every way (Speed, Capacity Potential, Durability)

That's Nice, But How Does This Impact the Cloud, and the DC?

- Cloud will become more about storage capacity than processing power
- In the past, rack space was dominated by servers (4 to 1), in the future, storage racks will outnumber server racks (2 to 1)
- Single Chip Cloud Computing, with potential for 1,000+ cores will replace entire DC halls with 2 or 3 servers (for redundancy)

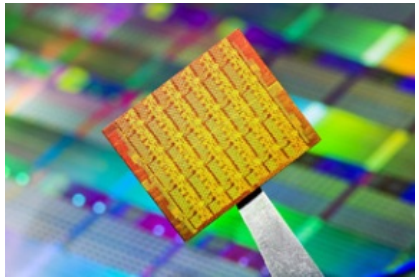

48 Core CPU

DC and Cloud Impact

- Leading driver of data centre build is Cloud service offerings
- CPU technology has started a trend of increased cores, with decreasing energy consumption
- CPUs are more tolerant to higher heat than in the past
- The move from HDD to SSD further improves the state of the DC

DC and Cloud Impact (Continued)

- HDD is one of the most sensitive components to heat in the DC (5c to 55c)
- SSD generates little heat of its own (no moving components) reducing the overall heat generated in the DC, and is highly tolerant of greater operating temperatures (0c to 70c)
- SSD utilizes 2w to 8w during idle and usage time; HDD consumes 10w to 12w
- Storage will take a larger footprint of data centre floor space

DC's – The Next Convenience Store

- Modular and Micro Modular Data Centres are a reality now
- They can be rapidly deployed to any location with sufficient power
- Latency and user demand will

  drive the need for closer

  proximity Data Centres

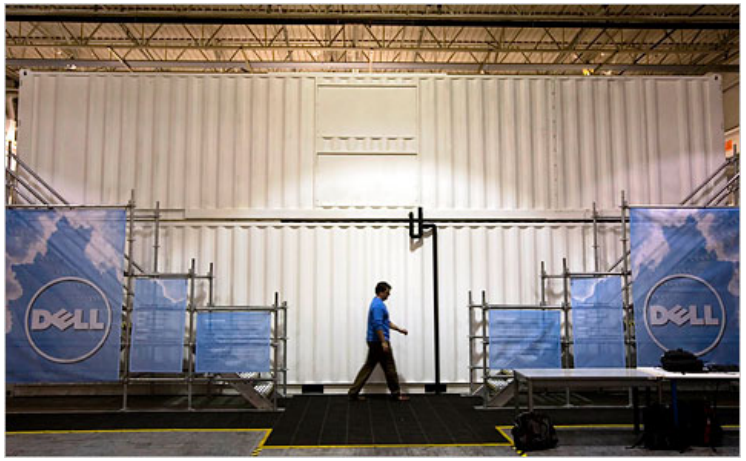

Thank you