Cloud Computing Brought Down to Earth

AKA

Blah blah Cloud Computing blah blah blah

Lots of Cloud hype, legends

"In 1943, IBM President T.J. Watson famously predicted " there is a world market for five computers"."

Did old T.J. really say this? I don't think so.

"Today, computers is (are) everywhere."

Legends

T.J. Watson was chairman and CEO of IBM, until 1956. One of the richest men of his time.

Could he really have been wrong about computers in 1943 (after all there were very, very few in 1943 and they were big and unreliable).

I'll bet he didn't say that - $5.00; but I could be wrong. Place your bets…

Legends

IBM TJ Watson Archives and Kevin Maney's book:

In 1953 on sales trip across the country, selling the IBM 701, … rents for $18,000 a month… we expected to get orders for 5 machines, we came home with orders for 18.

TJ

Arrested for Anti-trust and fined and sentenced to one year in jail (appealed, never served).

WW II puppet and war criminal.

Boy Scout leader.

Clouds

Clouds, Clusters, Grids

This is a computer:

Clouds, Clusters, Grids

This is a computer cluster:

Clouds, Clusters, Grids

These are computer clusters:

Clouds, Clusters, Grids

Pack a cluster into a shipping container:

Clouds, Clusters, Grids

Stack up the shipping containers:

Clouds, Clusters, Grids

And you have warehouse-scale data centers

Clouds, Clusters, Grids

Connect clusters (or warehouses of computers)

Clouds, Clusters, Grids

And you get VOs (virtual organizations), Grids:

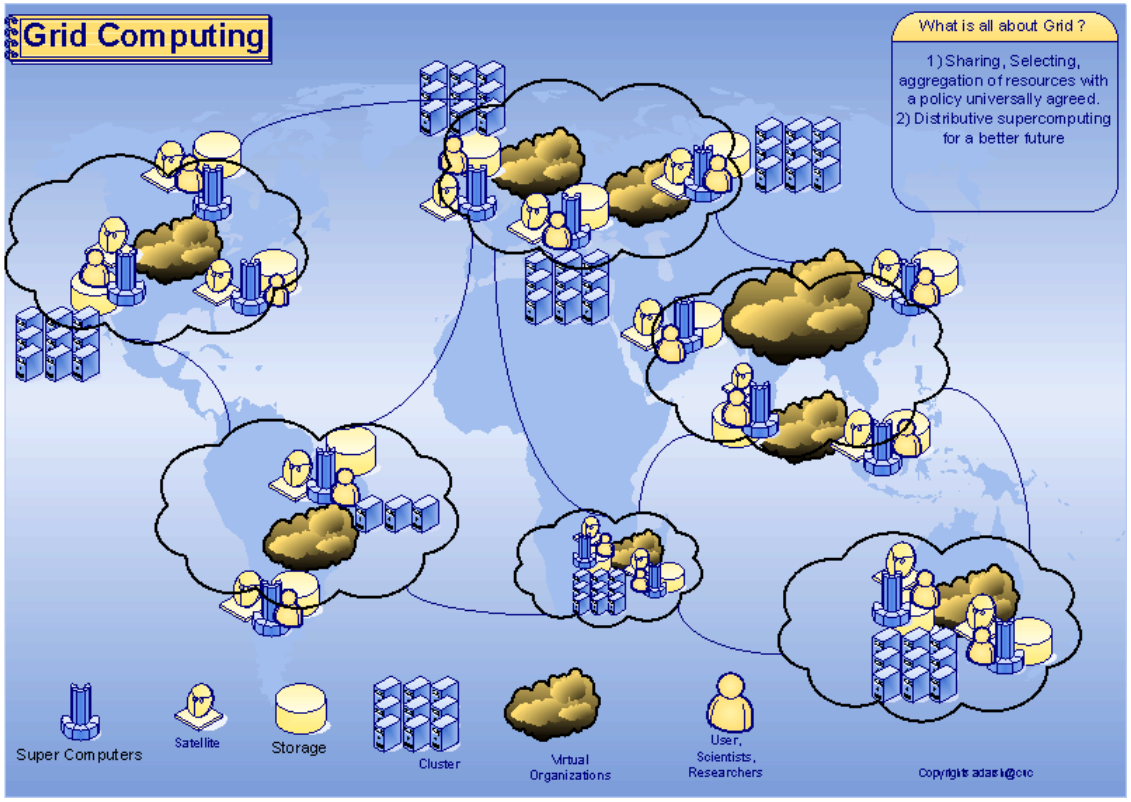

Clouds, Clusters, Grids

Or Clouds:

Amazon

- **North America**
  - **US East (Northern Virginia) Region** EC2 Availability Zones: 5* Launched 2006
- **US West (Northern California) Region**
  - EC2 Availability Zones: 3* Launched 2009
- **US West (Oregon) Region**
  - EC2 Availability Zones: 3 Launched 2011
- **AWS GovCloud (US) Region**
  - EC2 Availability Zones: 2 Launched 2011
- **AWS Edge Locations**
  - Ashburn, VA (2)
  - Dallas/Fort Worth, TX (2)
  - Hayward, CA
  - Jacksonville, FL
  - Los Angeles, CA (2)
  - Miami, FL
  - New York, NY (3)
  - Newark, NJ
  - Palo Alto, CA
  - San Jose, CA
  - Seattle, WA
  - South Bend, IN
  - St. Louis, MO

Amazon

Microsoft

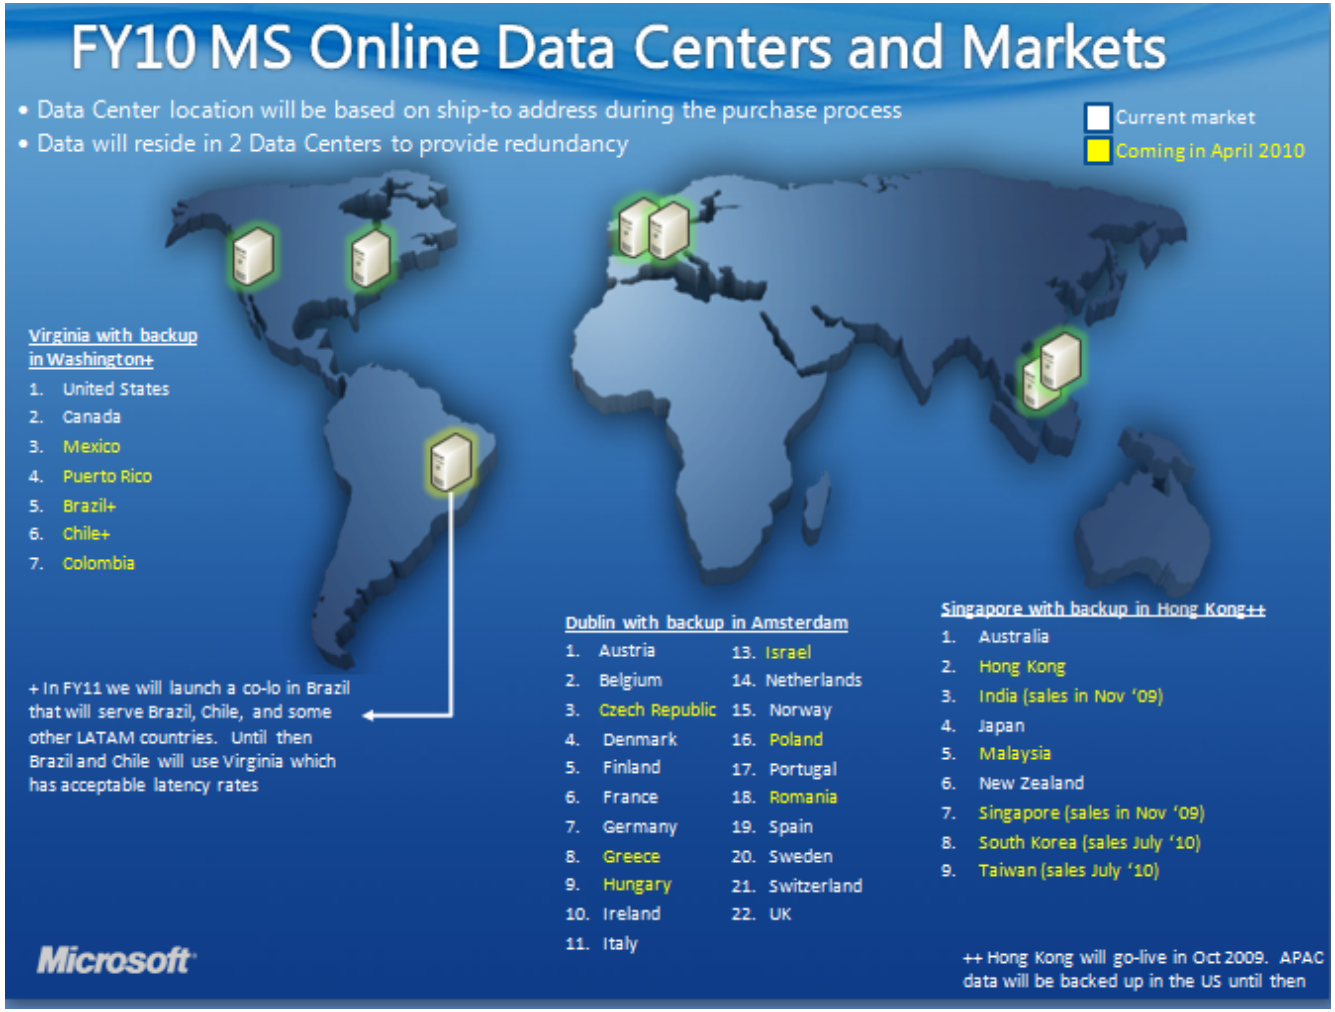
FY10 MS Online Data Centers and Markets
• Data Center location will be based on ship-to address during the purchase process
• Data will reside in 2 Data Centers to provide redundancy
Current market
Coming in April 2010
Virginia with backup in Washington+
1. United States
2. Canada
3. Mexico
4. Puerto Rico
5. Brazil+
6. Chile+
7. Colombia
+ In FY11 we will launch a co-lo in Brazil that will serve Brazil, Chile, and some other LATAM countries. Until then Brazil and Chile will use Virginia which has acceptable latency rates
Dublin with backup in Amsterdam
1. Austria
2. Belgium
3. Czech Republic
4. Denmark
5. Finland
6. France
7. Germany
8. Greece
9. Hungary
10. Ireland
11. Italy
13. Israel
14. Netherlands
15. Norway
16. Poland
17. Portugal
18. Romania
19. Spain
20. Sweden
21. Switzerland
22. UK
Singapore with backup in Hong Kong++
1. Australia
2. Hong Kong
3. India (sales in Nov '09)
4. Japan
5. Malaysia
6. New Zealand
7. Singapore (sales in Nov '09)
8. South Korea (sales July '10)
9. Taiwan (sales July '10)
++ Hong Kong will go-live in Oct 2009. APAC data will be backed up in the US until then
Microsoft

Microsoft

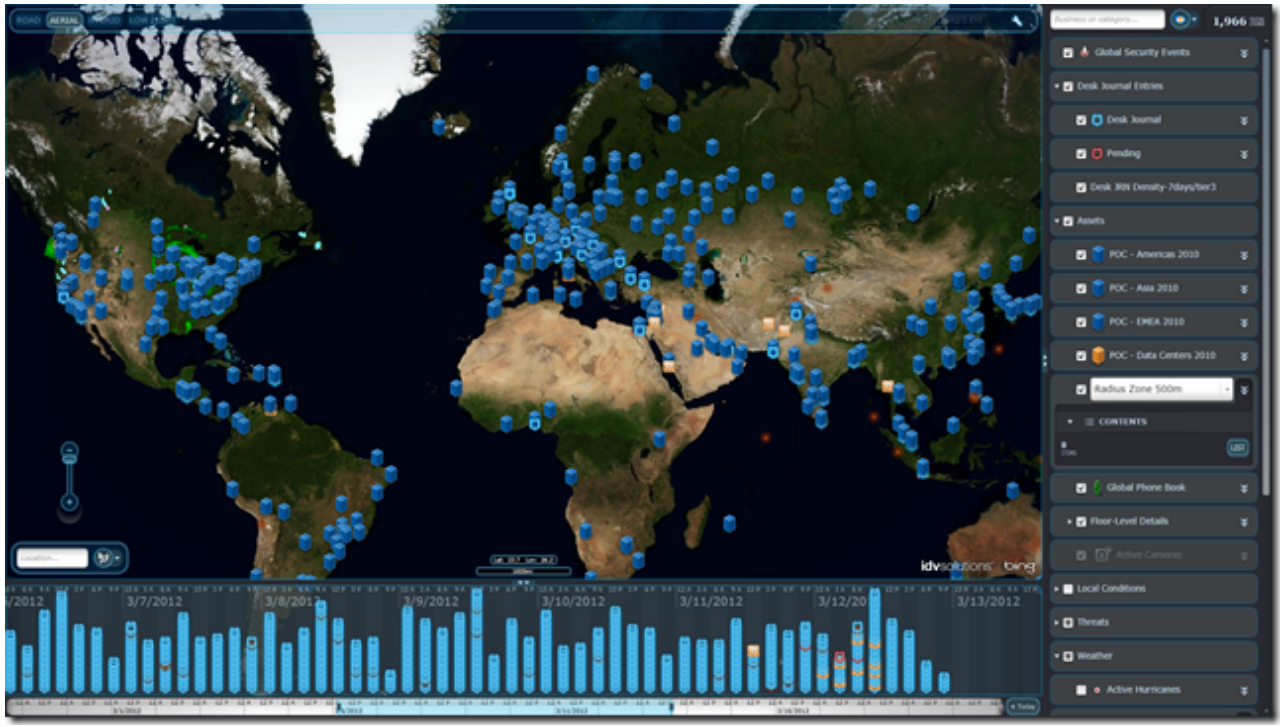
1,966
Global Security Events
Desk Journal Entries
Desk Journal
Pending
Desk JRN Density-7days/tier3
Assets
POC - Americas 2010
POC - Asia 2010
POC - EMEA 2010
POC - Data Centers 2010
Radius Zone 500m
CONTENTS
Global Phone Book
Floor-Level Details
Active Cameras
Local Conditions
Threats
Weather
Active Hurricanes
3/7/2012
3/8/2012
3/9/2012
3/10/2012
3/11/2012
3/12/20
3/13/2012

Microsoft

Multics

Started in 1964, project led by MIT, along with General Electric (later Honeywell) and Bell Labs.

Provide a computing utility – similar to telephone and electricity services. Plug in and pay for what you use.

Many novel ideas.

Died in 1969.

Multics

Plug in with one of these:

Multics

Into one of these:

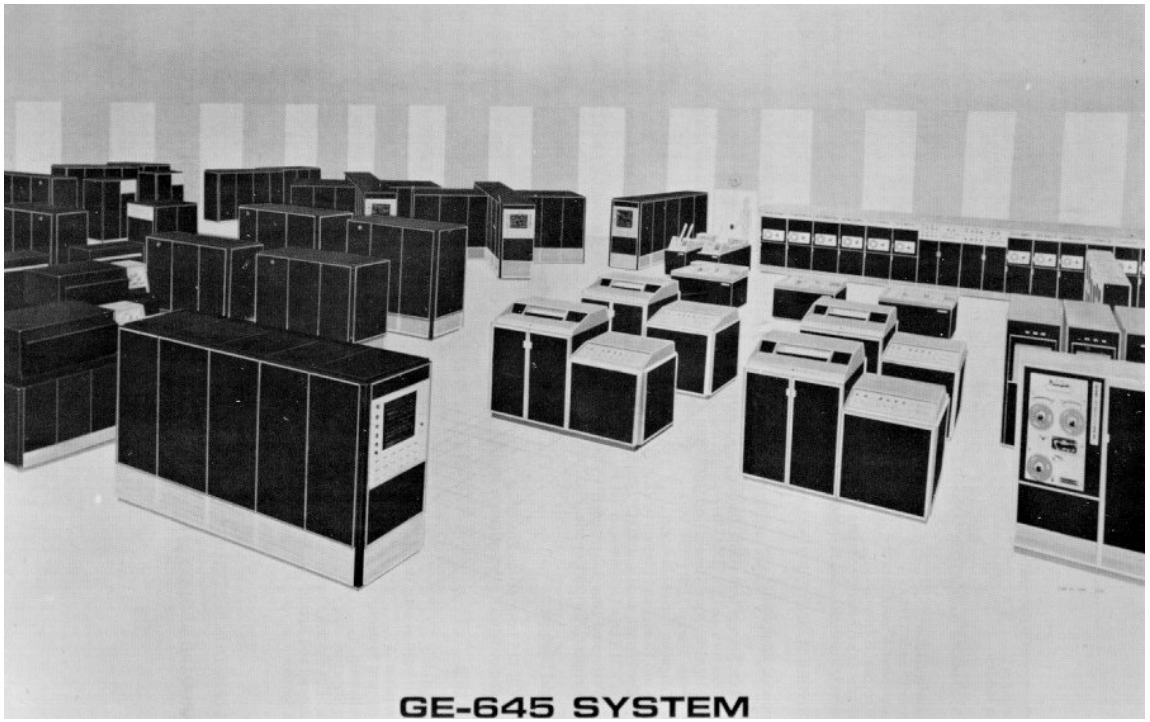

Multics

A nice, comfortable “computer” room:

Multics

Don’t need to buy:

Multics – Unix

All that was left (Ken Thompson, Dennis Richie):

Multics – Unix

GE-635: 512 Kilowords of memory, two accumulator registers, 40-200 thousand instructions per second, $2,900,000

Teletype 110 Baud

Disks – 5 Megabyte

If you weren't rich

Digital Equipment Minicomputers

$18 to $40 thousand (4K words – 128 K Bytes)

Data General Minicomputers

Personal Computers (Altair, Apple, IBM)

Meanwhile


12:49 AM
My Photo
Browser
00:49
My Video
App Market
My Music
Tudou
Camera
Settings
File Browser
Youtube
Picture Cap
Wi-Fi setting


Start
Diane
Michael Lange
18

Meanwhile

Meanwhile

2 to 4 CPU cores
2 Gigabytes memory
32 Gigabytes Flash Storage

16 Mbps Data (Cell)

Clouds

Music, Movies

Apps

Web Sites

Web Services

Utilities

Games

Social Friends

Clouds

Business Models:

Infrastructure as a service (IaaS)

Platform as a service (PaaS)

Software as a service (SaaS)

Storage as a service (STaaS)

Data as a service (DaaS)

Database as a service (DBaaS)

And more

Clouds

Business Models:

Infrastructure as a service (IaaS)

Offer Computers (storage, etc) usually virtualized with Xen, VMWare, KVM

User installs OS image and their software.

(Rackspace, Google Compute engine)

Clouds

Business Models:

Platform as a service (PaaS)

OS, Programming Languages, Databases, Webservers.

Add applications.

Some services scale resources to meet application demands.

(Amazon Elastic Beanstalk, Windows Azure Compute)

Clouds

Business Models:

Software as a service (SaaS)

Provider installed application software.

Cloned virtual machines at runtime,

Load balancing.

(Google Apps, Microsoft Office 365)

Clouds

Applications (on clients) use:

- Ajax (Java Script, XML)
- HTML5
- REST
- SOAP
- SQL or NoSQL DB (Hadoop, Cassandra)

Clouds

What can I do:

- Elastic, on-demand storage
- Dropbox
  - Originally 20K lines of Python, now 100 million users, one billion files stored per day.
  - Has API (multi language interface)

Clouds

Dropbox

Class assignment to encrypt all files dropped into Dropbox folder.

1 – 2 hours of work, 1 to 2 pages of Python

Dropbox has had security vulnerabilities, for example any user could access anyone else's files for a couple of hours one day.

Clouds

A few weeks ago a mathematician (Zach Harris) broke Googles internal use email key (they used a 512 bit DKIM key, since changed).

A little time on a few hundred CPUs, less than $100, could break Google, Amazon, eBay, PayPals keys.

Clouds

- NBC news reports Netflix, with 30 million North American streaming video subscribers uses 34% of all North American downstream internet traffic.
- Netflix demand is very elastic, Saturday and Sunday evenings demand is high, Monday morning, low.
- Netflix abandoned server farms (clusters) and moved to AWS (Amazon). They have used novel methods to serve demand – Cassandra (NoSQL DB) stores and replicates short (a few seconds) pieces of movies, and they are removed very quickly.

Clouds

Zynga with more than 300 million monthly active gamers could not predict usage of new games, and could not purchase, install and manage game servers fast enough for demand.

Zynga has also moved to Amazon Cloud services.

Clouds

Zynga with more than 300 million monthly active gamers could not predict usage of new games, and could not purchase, install and manage game servers fast enough for demand.

Zynga has also moved to Amazon Cloud services.

Clouds

Traditional web service providers are being replaced or forced to provide elastic, scalable, balanced services.

Cloud Problems

Cloud Problems

- 40 Thousand computers in a warehouse with network equipment, backup batteries, large disk arrays, consume significant power. Air conditioning, power conversion from utility providers, are a large overhead.
- Datacenters such as a couple of year old UTA cluster, have a PUE (power, conversion overhead waste) of more than 2.0 (half wasted) Googles newest centers have reached 1.06

Cloud Problems

Privacy

Providers may legally or illegally monitor data and communication from clients. HIPAA, FISMA and other standards may not be followed, or weakened.

Cloud Problems

Data Loss

Providers have lost client data, have corrupted data and in several cases have been compelled by governments to delete users data.

Cloud Problems

While a few providers have either opened their standards or use open source software, most do not. Data and software is very difficult to move to other cloud providers.

Cloud Problems

- Security
  - Security vulnerabilities continue to plague cloud based systems, many due to vulnerable applications (SQL injections, buffer overflows, plain text passwords, etc.) and others due to host OS, database and other host software vulnerabilities. The virtualization layer itself has been attacked (XEN, Virtual Box)

Cloud Problems

- Security
  - Malware has been targeted to cloud based warehouses, and virus/worms may easily spread.

Cloud Problems

Pricing

Providers have a very complex and dynamic pricing scheme. Small (virtual) instance are often as low as $0.02 per hour, but as demand increases from other users, prices change as well. Machines may be “reserved” to guarantee QOS, and datacenter CPU is constantly monitored and used to reevaluate pricing.

Cloud

How can I start (cheaply).

Most cloud service providers will give you a small amount of free service, either in terms of virtual instances (200 hours per month, for example) or hours of use (500 hours). A small amount of raw or database storage is also provided.

Cloud – Case Studies

Amazon installed many computers for ebusiness .COM bust, left with too much capacity.

Sell it! (Rent it)

Small group in Africa build cloud infrastructure.

Cloud – Case Studies

- Amazon - Compute
  - Elastic Compute Cloud (EC2)
  - Elastic Map-Reduce
  - Auto Scale
  - Load balancing

Cloud – Case Studies

- Amazon – Storage
  - RDS – Relational Database services
  - DynamoDB – NoSQL
  - Elastic Cache
  - S3 – Simple storage

Cloud – Case Studies

- Amazon – Manage
  - Elastic Bean Stalk
  - Access Management
  - Billing
  - Cloud Watch

Cloud – Case Studies

Where to put Datacenters

- Amazon – East, West Coast
  - Latency – users disapointed at 120+ ms
- Energy/Power
- Linux->Windows->what software?

Cloud – Case Studies

Intel (X-86)?

ARM?

(140 Watts vs 8)

SPARC, Itanium

Build your own? (Compilers, OS, etc.)

Cloud – Case Studies

Datacenters:

- Rack Mounted – Cases,
  - "Baking Pan"?
- 80 degrees F
- Cool only Hot Spots

Cloud

Skills:

- Software Engineering – designing, implementing and testing a large distributed system
- Distributed Computing
- Parallel Computing
- Databases (Relational and NoSQL)

Cloud

Thank You