
CCA
CLOUD COMPUTING ASSOCIATION
&
DCIA
Distributed Computing Industry Association
proudly present:

WELCOME


CLOUD COMPUTING
EAST 2013
Revolutionizing Business Processes in
Government, Healthcare & Financial Services
MAY 19-21, 2013
Boston Marriott Copley
Place, Boston MA

CLOUD COMPUTING DEFINITIONS AND SPECIAL CONSIDERATIONS

Boston University's

- *Dino Konstantopoulos*
- *Alaa Mahmoud*
- *Madhura Pundlik*

What is.. Cloud Computing
CCE2013

May 19, 2013

Dino Konstantopoulos

Topics

1. Cloud as a new computing Model
2. An economic & technical shift based on virtualization
3. The corporate Data Center and the Public Cloud
4. Cloud Myths
5. Cloud visionaries
6. Cloud computing models
7. How Big?
8. The new algorithms & architectures
9. The social dimension of Cloud computing
10. Conclusion & Futures

Just suppose..

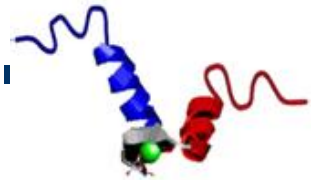

- Tomorrow, scientists announce a breakthrough in DNA Computing
  - Alternative method of computing that takes advantage of the massive parallel processing characteristic of DNA molecules
  - We now have the ability to solve NP-Complete and NP-Hard combinatorial problems in polynomial time

Scared?

- Whether it's DNA computing or quantum computing that ultimately makes this leap, things get really scary when you factor in Moore's law
- Assuming that DNA proteins/qubits are valid analogs for transistors, the first machine capable of breaking all known algorithms will be followed by a version twice as powerful within 18 months

Outcomes

- Scenario 1
  - Nobody needs passwords anymore at all. In a world where encryption is ineffective, the only solution is absolute transparency
  - Every computer recognizes your face, every mouse your palm, every door handle your fingerprints, and *The Internet of Things* is a worldwide Orwellian state where NOTHING you do is unknown
- Scenario 2
  - All secrets become rolling secrets, your social security number changes every 10 seconds
  - And who holds the key?
- Scenario 3
  - Biological computers reduce environmental toxicity
  - Improved energy efficiency allows for a decrease in global energy consumption

Back to reality..

- Your static secrets are still safe..
  - But they have to be >= 8 characters/ digits/#@!*
- But what if I told you that the outcomes we talked about are already here?
  - Your devices already *talk* to the Cloud
  - Your identify is *managed* by your Cloud
  - Your Cloud is *good* for the environment

Part 1

CLOUD AS A NEW COMPUTING MODEL

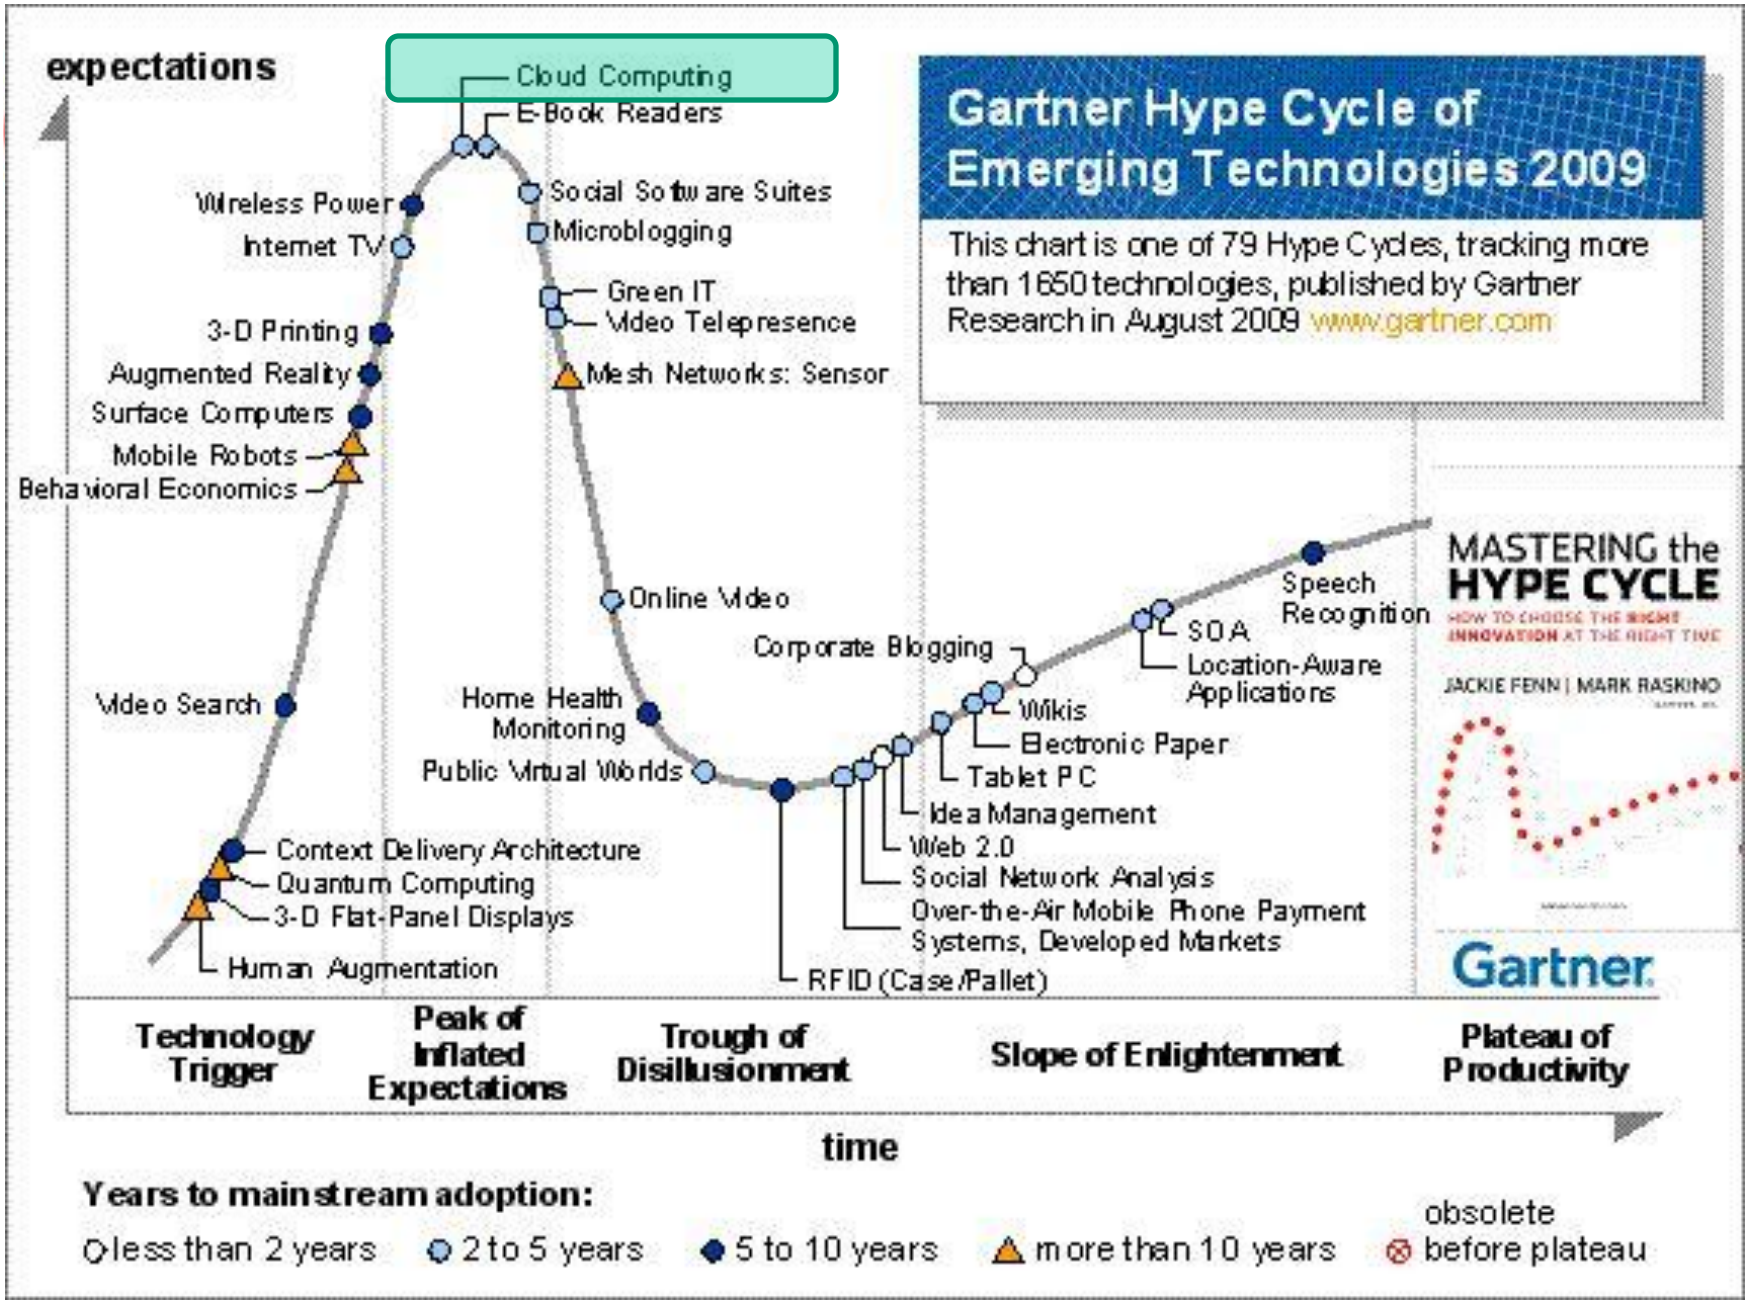
expectations
Cloud Computing
E-Book Readers
Wireless Power
Internet TV
Social Software Suites
Microblogging
Green IT
Video Telepresence
3-D Printing
Augmented Reality
Surface Computers
Mobile Robots
Behavioral Economics
Mesh Networks: Sensor
Online Video
Video Search
Home Health Monitoring
Public Virtual Worlds
Corporate Blogging
Speech Recognition
SOA
Location-Aware Applications
Wikis
Electronic Paper
Tablet PC
Idea Management
Web 2.0
Social Network Analysis
Over-the-Air Mobile Phone Payment Systems, Developed Markets
Context Delivery Architecture
Quantum Computing
3-D Flat-Panel Displays
Human Augmentation
RFID (Case/Pallet)
Gartner Hype Cycle of Emerging Technologies 2009
This chart is one of 79 Hype Cycles, tracking more than 1650 technologies, published by Gartner Research in August 2009 www.gartner.com
MASTERING the HYPE CYCLE
HOW TO CHOOSE THE RIGHT INNOVATION AT THE RIGHT TIME
JACKIE FENN | MARK RASKINO
Gartner
Technology Trigger
Peak of Inflated Expectations
Trough of Disillusionment
Slope of Enlightenment
Plateau of Productivity
time
Years to mainstream adoption:
less than 2 years
2 to 5 years
5 to 10 years
more than 10 years
obsolete before plateau

2010 EMERGING

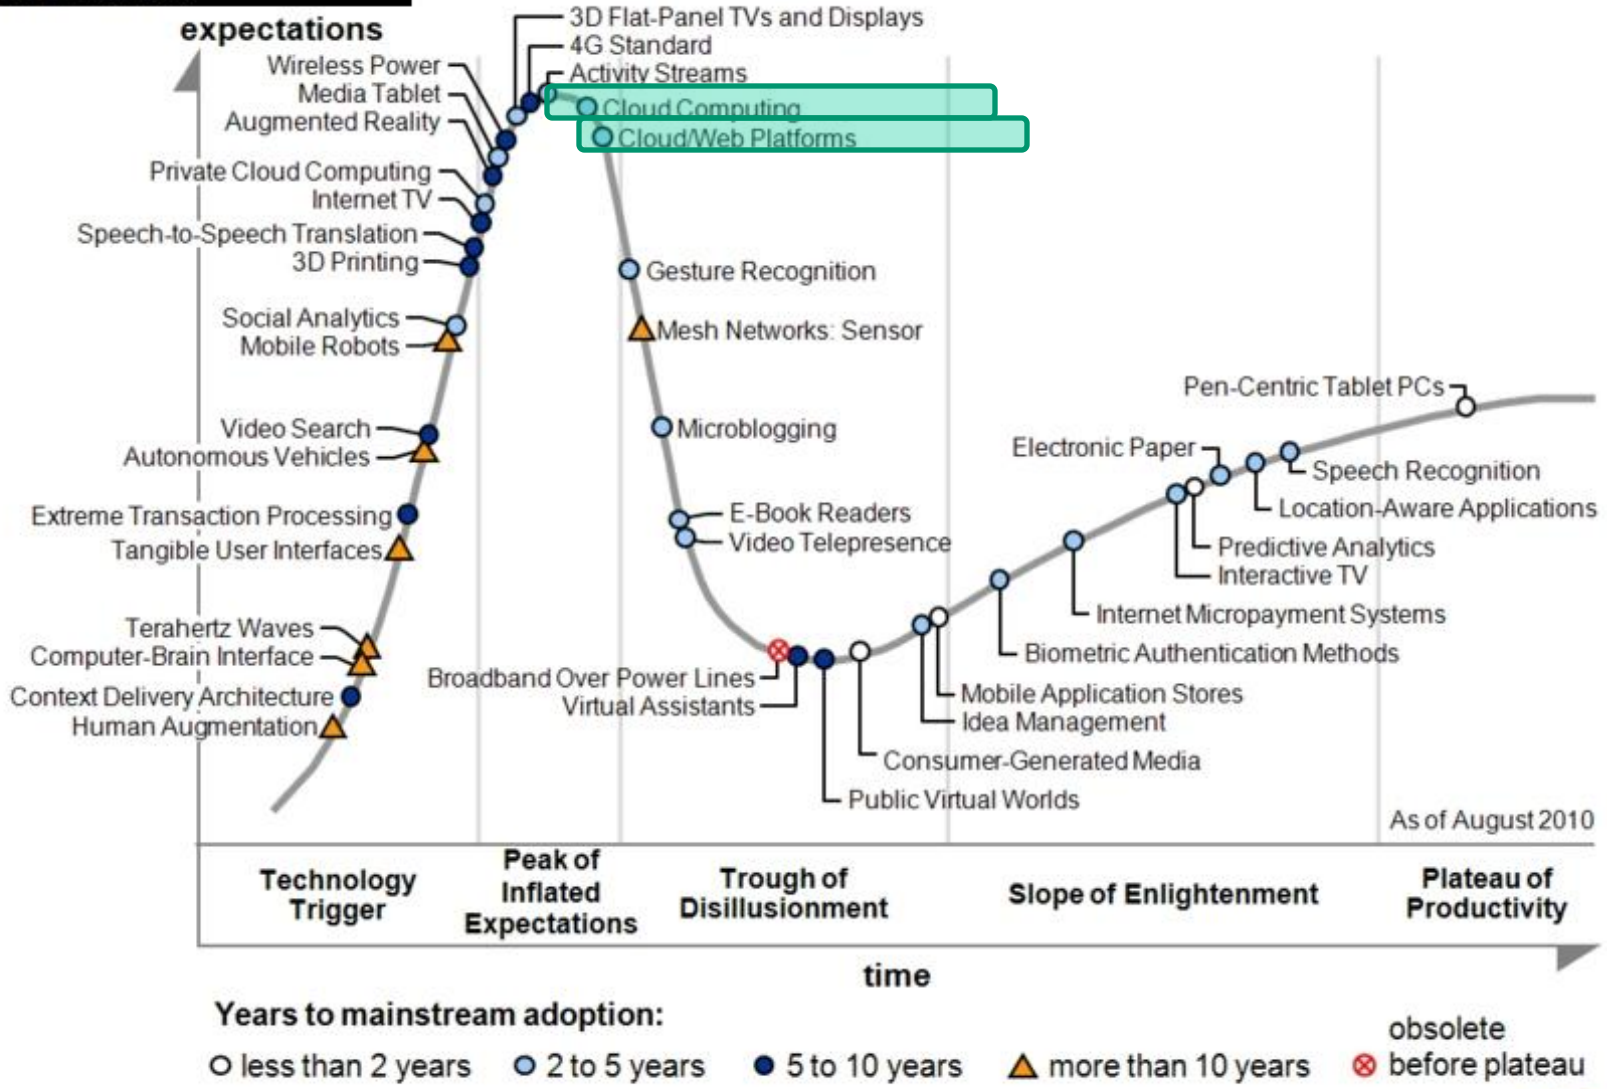

Gartner

Emerging Technologies Hype Cycle 2012

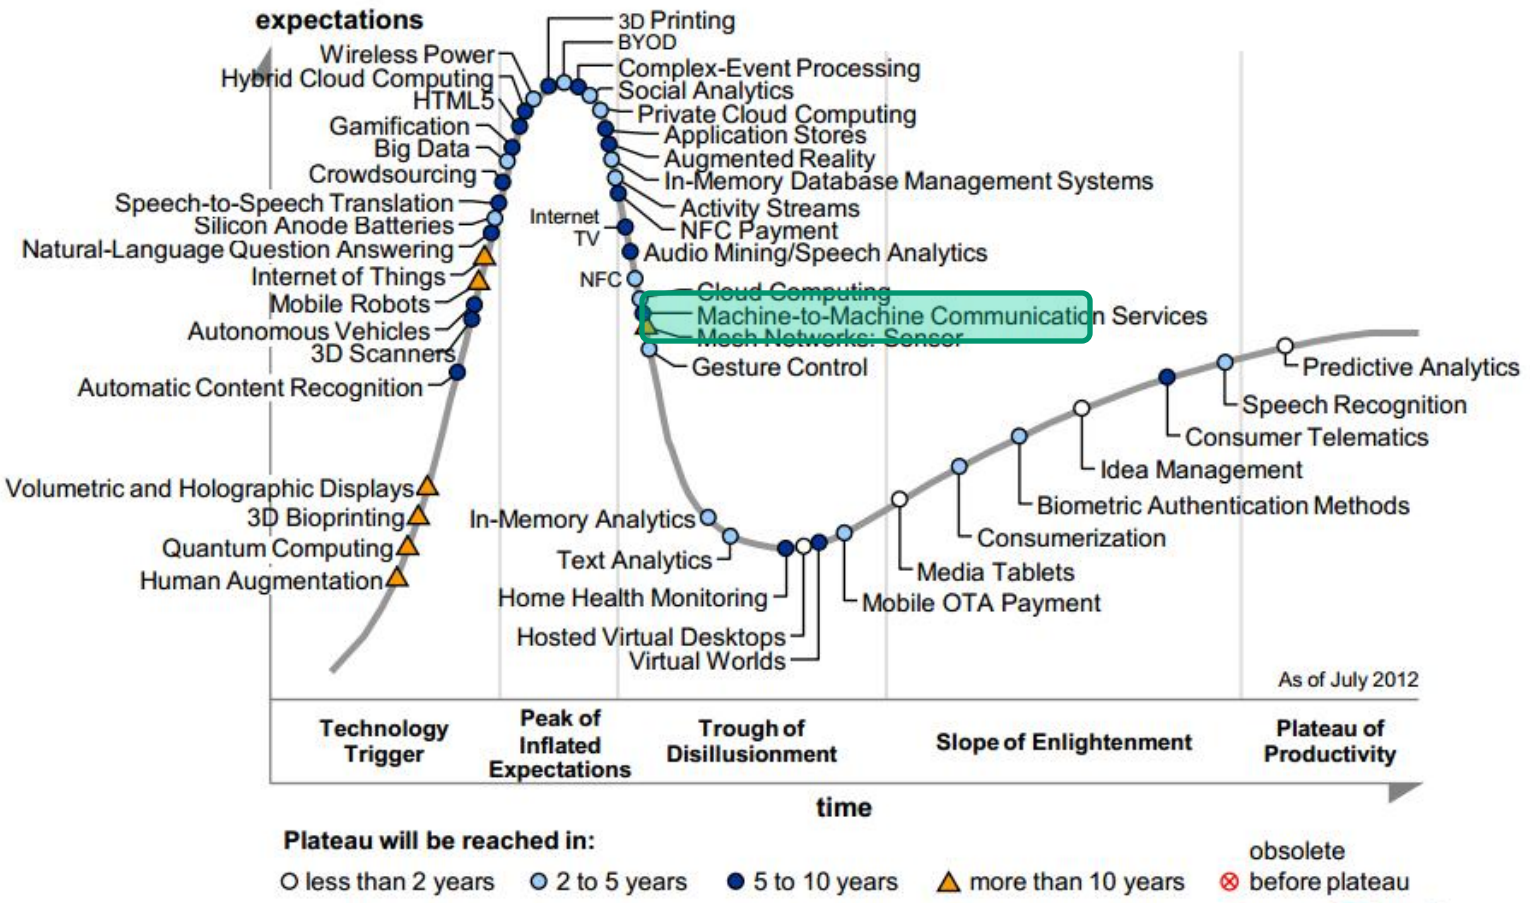

Gartner

A new computing model

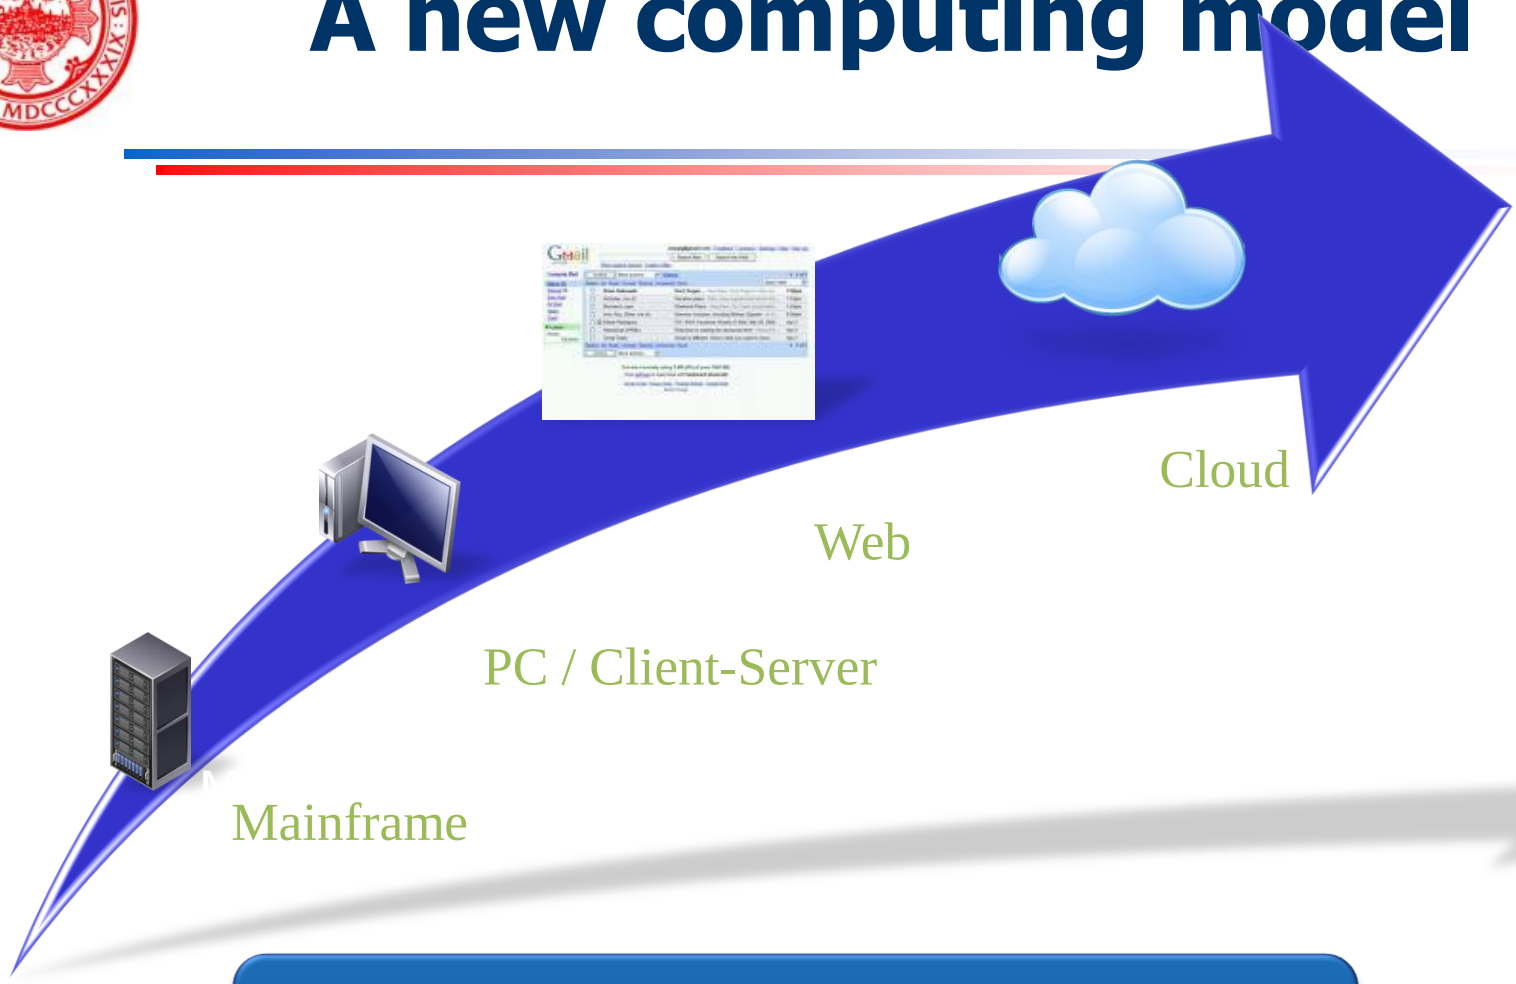

Cloud Computing is transforming the delivery of IT services

Cloud Computing Evolution

Grid Computing

- Solving large problems with parallel computing
- Made mainstream by Globus Alliance

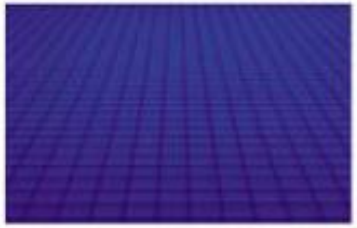

Utility Computing

- Offering computing resources as a metered service
- Introduced in late 1990s

Software as a Service

Network-based subscriptions to applications

Gained momentum in 2001

Cloud Computing

Next-Generation Internet computing

Next-Generation Data Centers

The 5 essential characteristics of the Cloud

Services
Pooled computing resources are consumed as services

Multi-Tenant
Resources are shared among many customers

Consumption-Based
Customers pay only for resources used

Elastic
Resources are provisioned and released in real-time

Internet
Resources are accessed and utilized via the Internet

Part 2

AN ECONOMIC & TECHNICAL SHIFT BASED ON VIRTUALIZATION

The Cloud

- Similar in importance and impact to the transition from mainframe to client/server, Computing is undergoing a shift from Client/Server to the Cloud
- An ***economic*** and ***technical*** shift with great potential to reduce the cost of Information Technology (IT) while improving IT capabilities and stimulating innovation in IT solutions
- Based on OS ***virtualization***

Virtualization, how it came to be

- The greatest failure of Winux, ***ever:***
  - The inability to scale the bandwidth of ***one*** networked application to serve an inflationary burst of clients

Multiple OSs, serving one application

The Hack

- So, instead, companies provide a farm of applications, each one running on a separate OS, with a front-facing load balancer

Wait a second..

- IBM did some research in the 1960s, to build robust ***time-sharing*** systems for its mainframes
  - The CP-40 supported *Dynamic Address Translation*
    - Virtualizes memory, hard disk, etc..
    - Can make it look like ***one app*** owns ***all*** hardware resources, while in reality they are being shared by many apps
    - The hardware's *supervisor state* was virtualized as well, allowing ***multiple OSs*** to run concurrently, too

Hypervisor?

- The VM control program included a *hypervisor-call* handler that intercepts DIAG ("Diagnose") instructions used within a virtual machine
  - Provides non-virtualized execution of file-system access
  - Available for use as a signal to the "host" operating system
  - (Today, this is called paravirtualization)


UNIVERSITAS BOSTONIENSIS
CONDITA MDCCCXXXIX

Hypervisor!


IBM
Introducing the IBM Operating System/2
Version 1.00
(C) Copyright IBM Corp. 1981, 1987. All rights reserved.
IBM
Introducing the IBM Operating System/2
Version 1.00
(C) Copyright IBM Corp. 1981, 1987. All rights reserved.
IBM
Introducing the IBM Operating System/2
Version 1.00
(C) Copyright IBM Corp. 1981, 1987. All rights reserved.

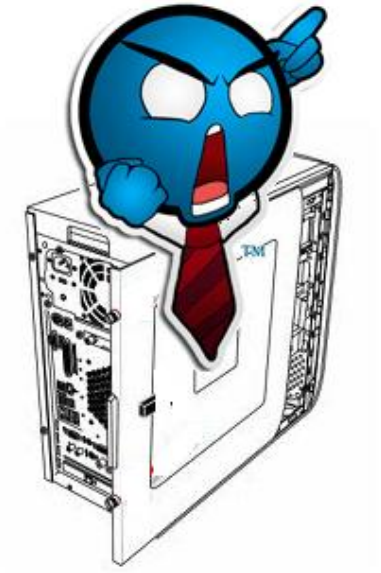

The problem with Intel..

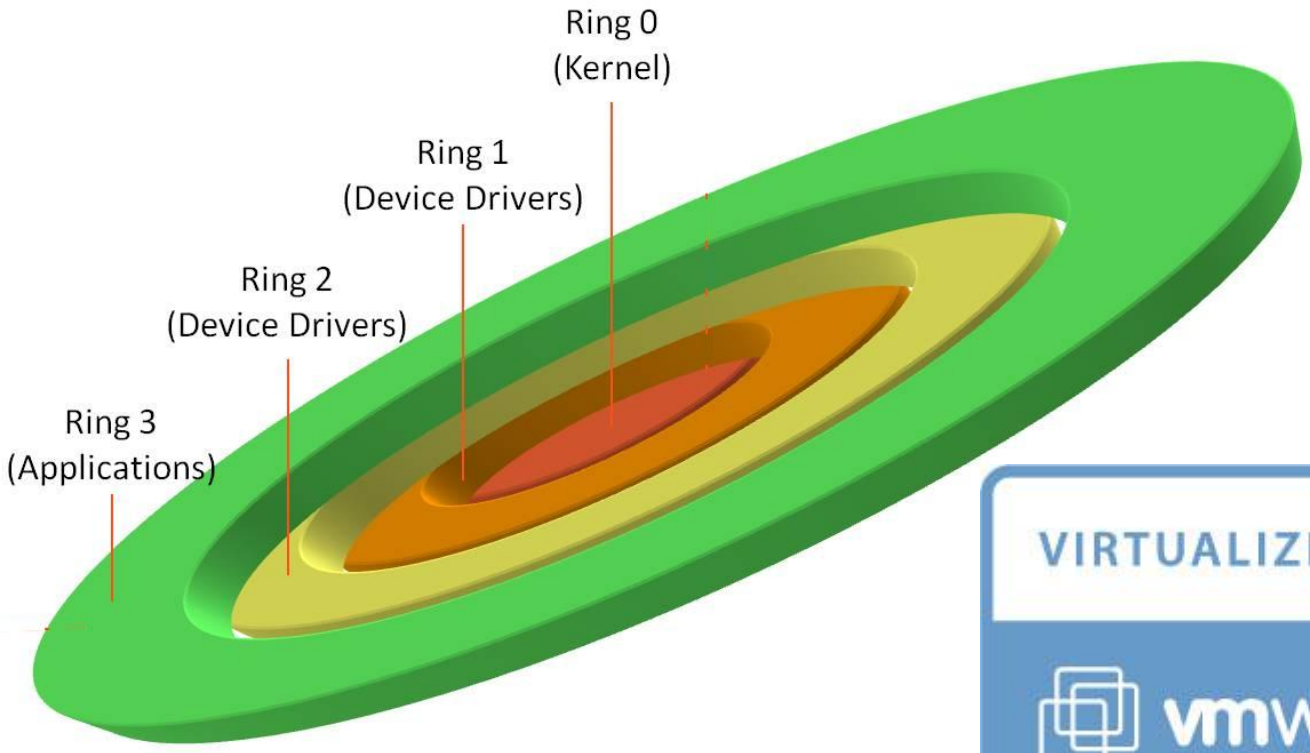

VMWare ESX

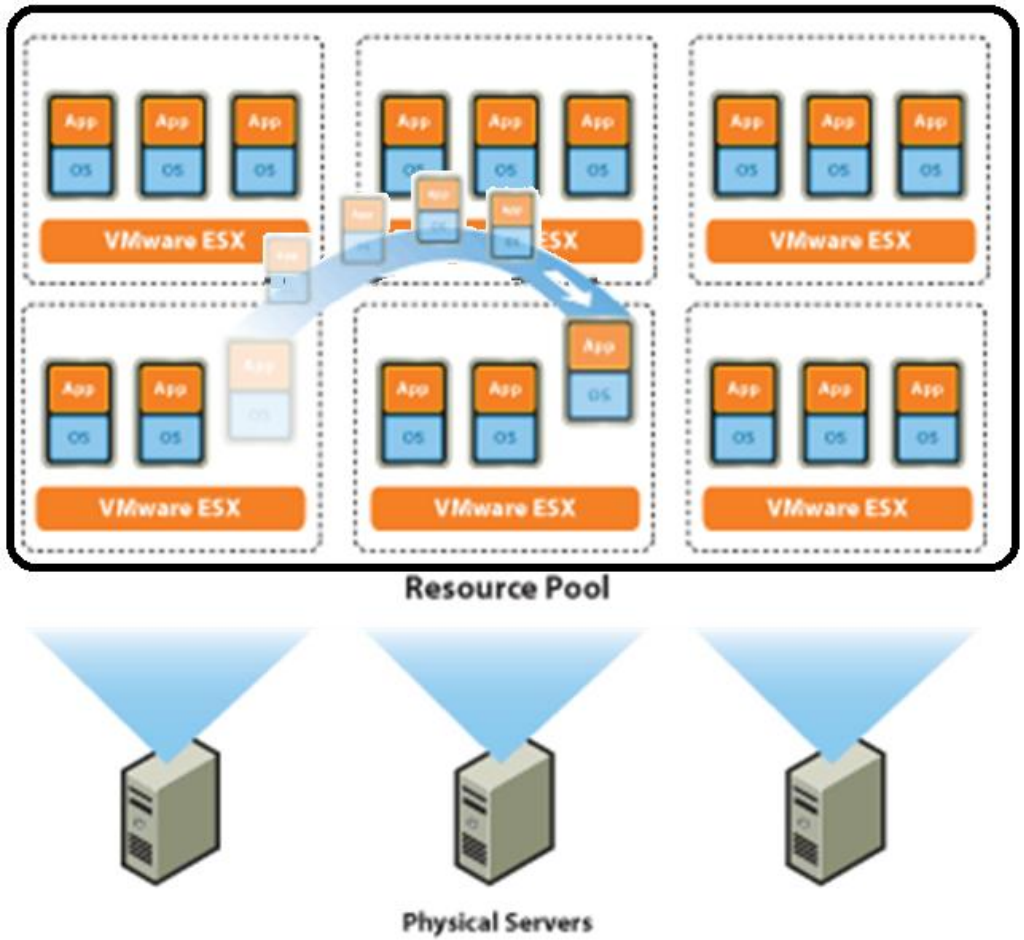

Hypervisor types

Hypervisor Design:

Two approaches

Type 2 Hypervisor

Examples:

- Virtual PC & VirtualBox
- VMware Workstation
- KVM

Type 1 Hypervisor

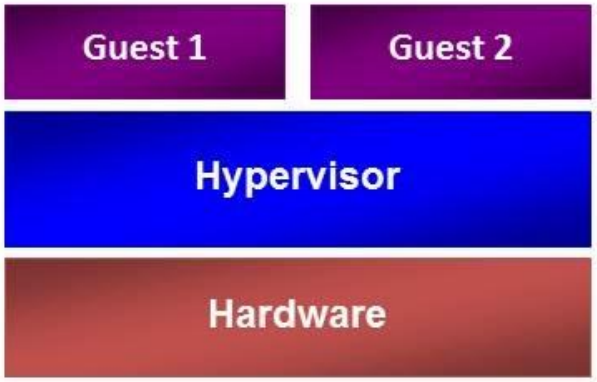

Examples:

- Hyper-V
- Xen
- VMware ESX

Part 3

THE CORPORATE DATA CENTER AND THE PUBLIC CLOUD

The Corporate Data Center: The Private Cloud

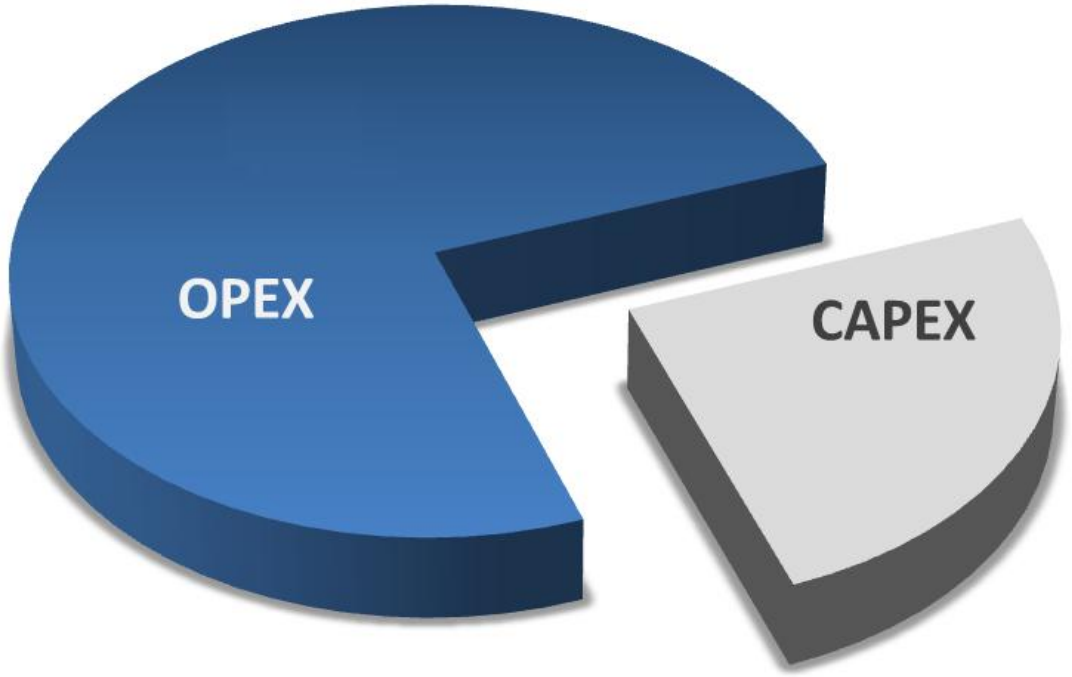

Where the IT budget goes

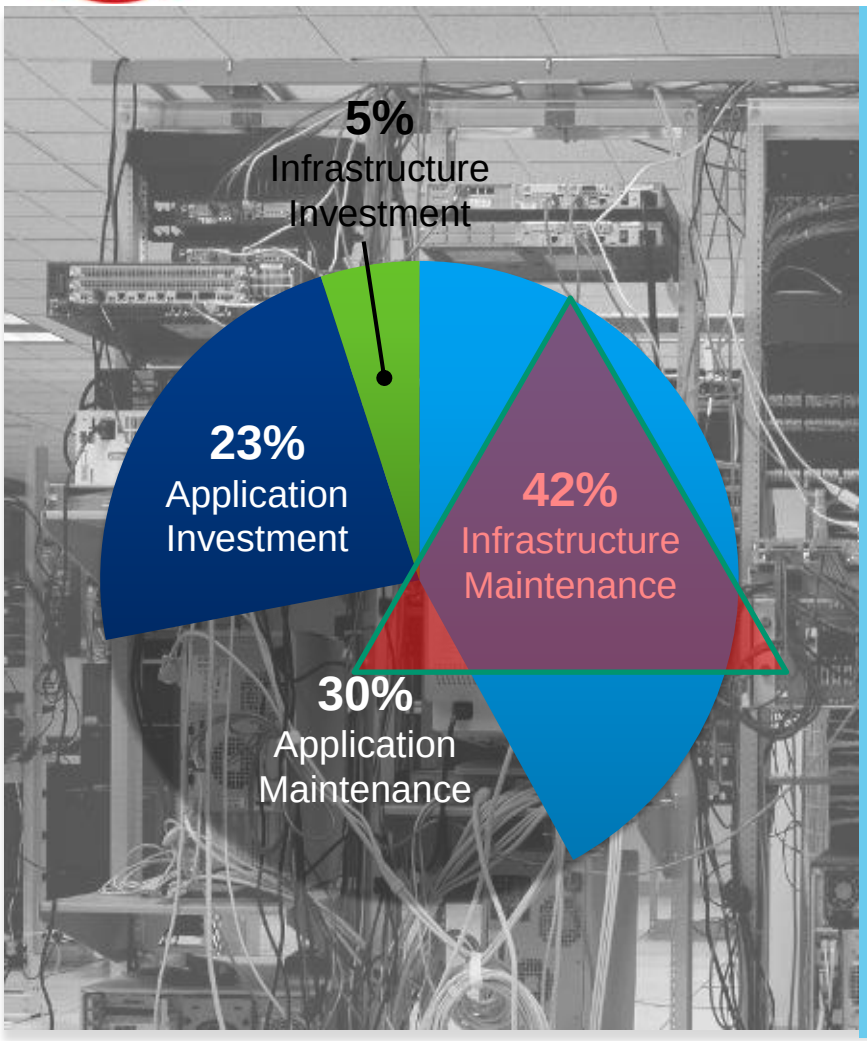

Problem Areas:

- Overwhelming complexity
- Insufficient space & power
- Meeting service levels
- Privacy and security
- Maintaining compliance
- Unfunded mandates
- Adjusting to business needs

Data Centers

- Electricity cost is rapidly rising to become the largest element of total cost of ownership (TCO), currently representing 15%-20%
- Power Usage Effectiveness (PUE) tends to be significantly lower in large facilities than in smaller ones

Data Centers

- **Infrastructure labor costs**
  - While cloud computing significantly lowers labor costs at any scale by automating many repetitive management tasks, larger facilities are able to lower them further than smaller ones..
- **Security and reliability**
  - While often cited as a potential hurdle to public cloud adoption, increased need for security and reliability leads to economies of scale due to the largely fixed level of investment required to achieve operational security and reliability

Data Centers

- **Buying power**
  - Operators of large data centers can get discounts on hardware purchases of up to 30 percent over smaller buyers

The Public Clouds
If you build it, they will come..

Google innovations

- Google realized that the so-called cold aisle in front of the machines could be kept at a relatively balmy 80 degrees or so
  - Workers could wear shorts and T-shirts
- The "hot aisle," a tightly enclosed space where the heat pours from the rear of the servers, could be allowed to hit around 120 degrees
- That heat could be absorbed by coils filled with water, which would then be pumped out of the building and cooled before being circulated back inside

Green Computing

- All of these innovations helped Google achieve unprecedented energy savings
  - Standard measurement of data center efficiency is called ***power usage effectiveness***, or PUE
  - A perfect number is 1.0, meaning all the power drawn by the facility is put to use
  - Experts considered 2.0—indicating half the power is wasted—to be a reasonable number for a data center
  - Google was getting an ***unprecedented 1.2***

Amazon Web Services

- Werner Vogels, joined Amazon in 2004 as director of systems research, became Amazon CTO early in 2005 and later that year was named VP
  - Had a vision of a type of distributed system that relied on inexpensive parts but could scale out infinitely, and not come to a halt if a piece of hardware failed underneath it
- Chris Pinkham was EC2's managing director
- Christopher Brown was lead developer
- Amazon.com (online-merchandising operation) and Amazon EC2 (Cloud) were initially two separate things
- The two pulled it off, and Amazon EC2 was born

Microsoft Azure, June 2012

- Websites allows developers to build sites using ASP.NET, PHP, or Node.js
- Virtual machines let developers migrate applications and infrastructure without changing existing code, and can run both Windows Server and Linux VMs
- Cloud services - Microsoft (PaaS) environment used to create scalable applications and services
- SQL Database, formerly known as SQL Azure
  - Integrates with *Active Directory, Microsoft System Center* and *Hadoop*

The Horseless Carriage

Public Clouds work because..

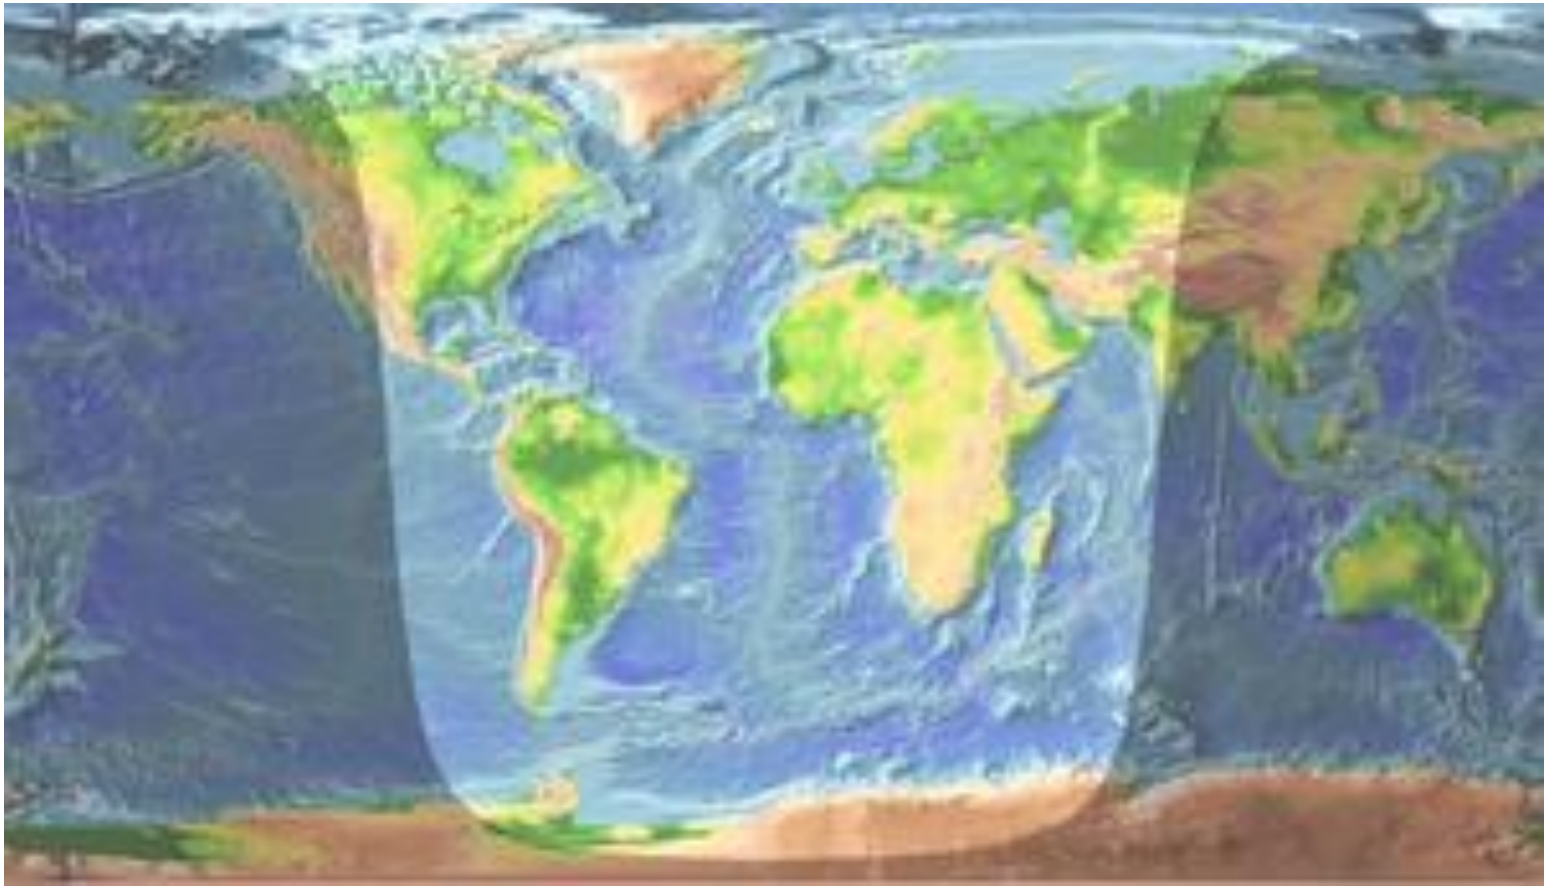

Clouds work because

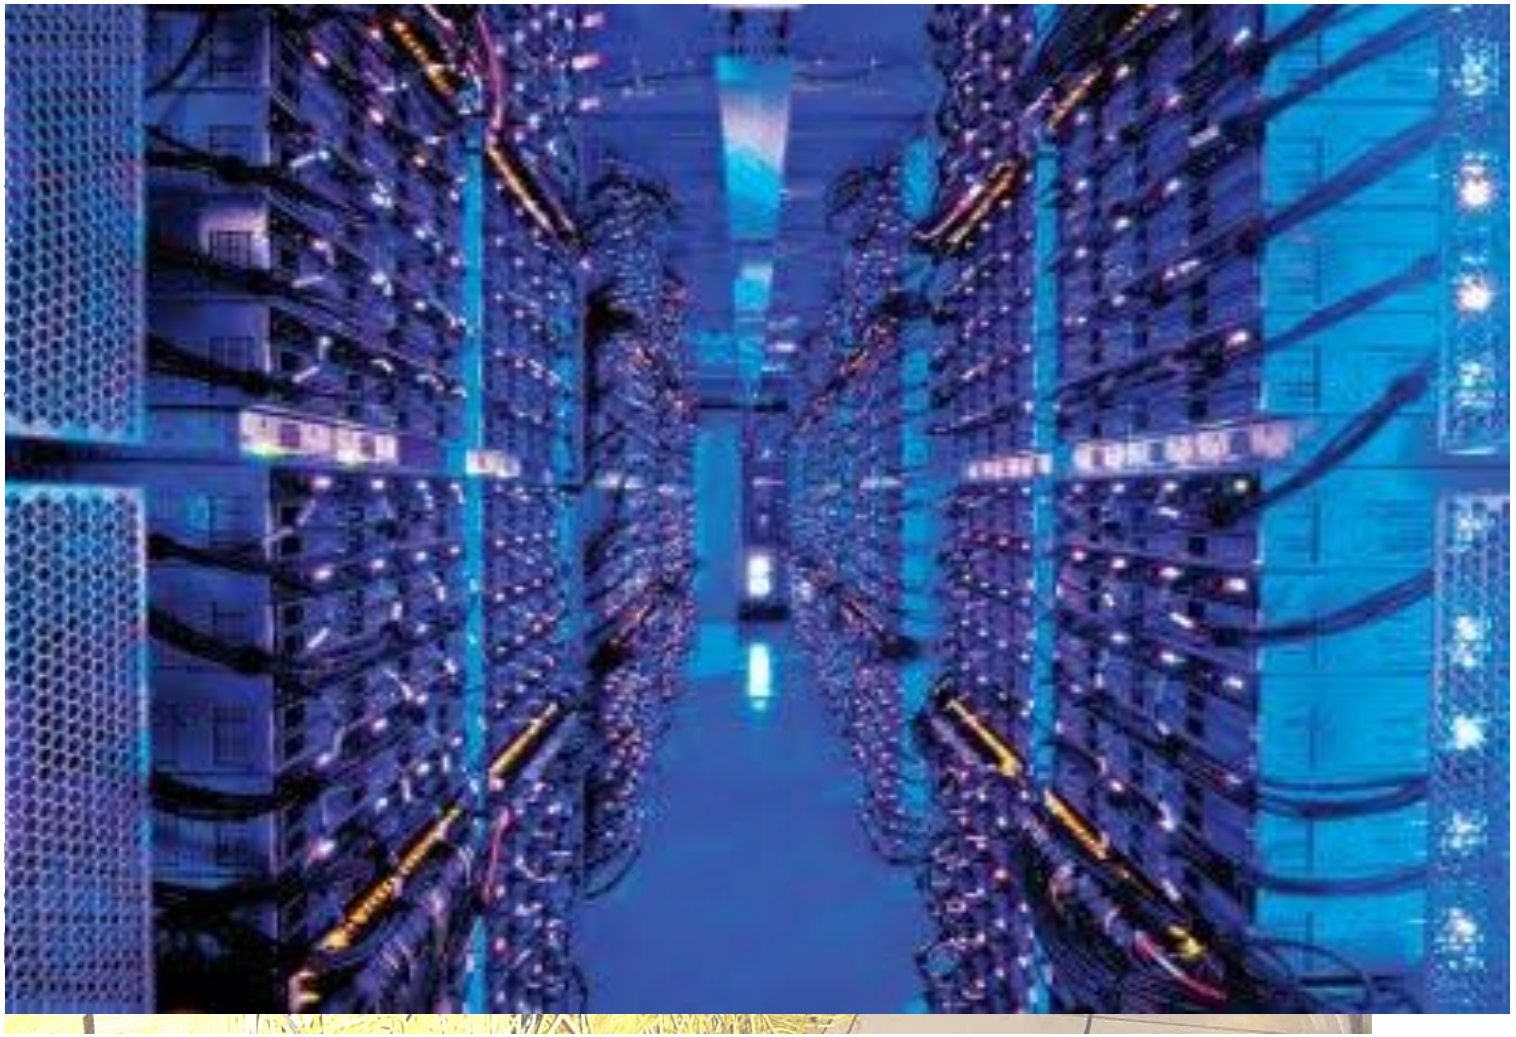

Part 4

CLOUD MYTHS

Myth: Public Cloud is Unsafe

- Public Cloud security is ***more rigorous*** and ***strenuously monitored*** than an enterprise data center's security
  - PCI-compliant operations have been established in Terremark's massive Network Access Point of the Americas data center in downtown Miami
    - Passed stringent DOD Information Assurance Certification and Accreditation Process

Myth: VMs are safe

- One VM can't ***spy*** on another running on the same server, right?
  - Extremely skillful manipulators have been able to draw conclusions about what's going on in a neighboring VM by watching what cache pages get emptied out of host memory after the spying VM has taken its turn using the server core

Myth: Public Cloud Costs More

- If the cloud allows you to avoid making a capital purchase, then it will almost always enjoy a demonstrable cost advantage in the short run
- But how about long term?

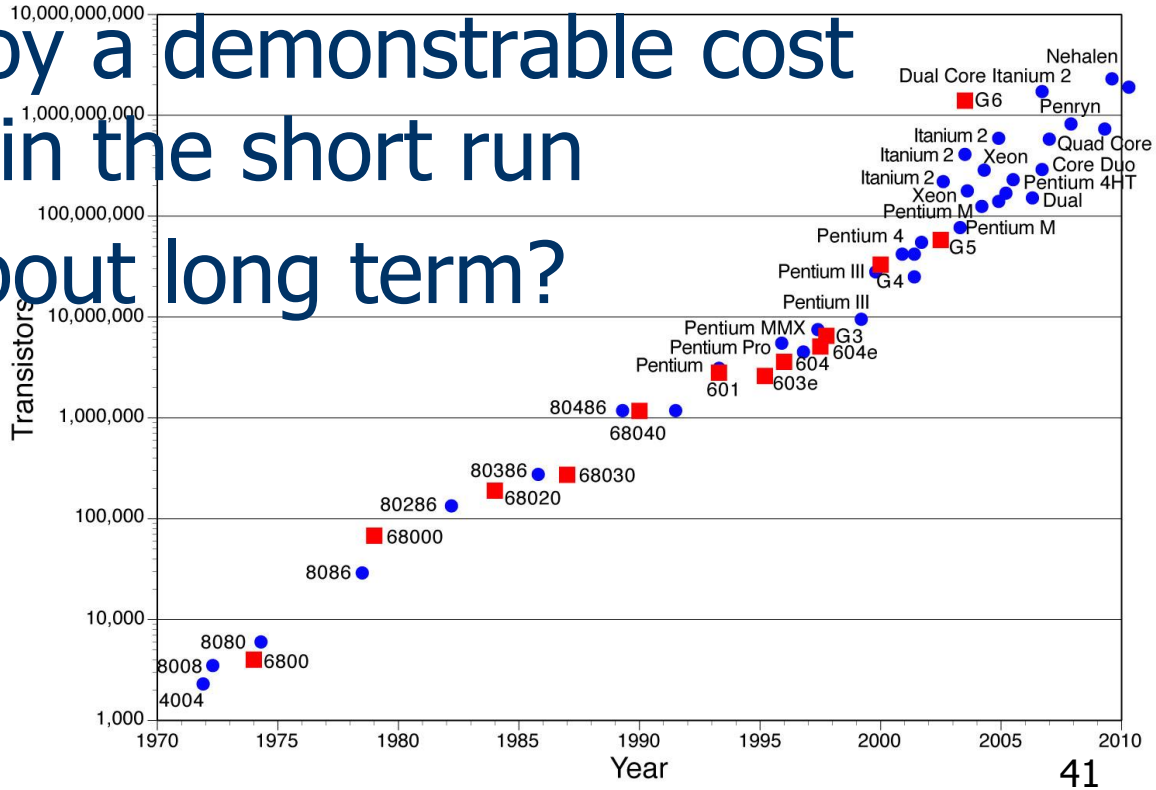

Myth: Pick Google/AWS/Azure & Go

- Public Clouds preconfigure VM templates from micro to small, medium, large and extra-large
  - But nowhere is there a clear definition of these terms..
  - Calculating a comparison of charges from one Public Cloud to another is difficult

Myth: Public Clouds runs Winux

- Public Cloud runs mostly on AMD/Intel commodity servers running the operating systems most common to Intel's x86 instruction set
  - HP is producing data center servers based on Calxeda-designed ***ARM*** chips for telecomms
  - ***SmartOS*** operating system, a derivative of Open Source ***Illumos*** (alternative provider of Solaris)

Myth: Public Cloud is Proprietary

- Proprietary
  - Public Cloud mostly Linux & Windows Server 2008
  - Hypervisors mostly VMWare, Hyper-V, and Xen
- Open Source
  - ***Eucalyptus Systems***, ***OpenStack***, ***CloudStack*** vying to establish a stronghold with developers and users
  - OpenStack's Project Quantum is rich in virtual networking talent as it captures code contributions from HP, Nicira, Cisco, IBM and others

Myth: Cloud Data Centers Are Killing The Environment

- After all, Google ***is*** the biggest single consumer of electricity in the world..
  - The new data centers accelerate power consumption and lead to more carbon dioxide production and environmental degradation
- But ***more*** computing is being done on ***less*** power
  - And Google also built the first data center with a … and opened up about its design

Part 5

CLOUD VISIONARIES


SOCIAL

Part 6

CLOUD COMPUTING MODELS

Evolution of Computing Models

	Technology	Economic	Business Model
Mainframe 	Centralized compute and storage Thin clients	Optimized for efficiency because of the high cost	High up-front costs for hardware and software
Client/ Server 	PCs and servers for distributed compute, storage, and so on	Optimized for agility because of the low cost	Perpetual license for OS and application software
Cloud 	Large DCs, ability to scale, commodity hardware, devices	Efficiency and agility an order of magnitude better	Ability to pay as you go, and only for what you use

S/P/IaaS


SaaS
On-demand applications
• Pay-per-user software
• Fast and easy to deploy
• Managed by the vendor
• Short- or long-term use
http://ww
PaaS
On-demand software platforms
• Build-your-own cloud services
• Scalable test environments
• Platform managed by the vendor—
you manage your applications
IaaS
On-demand data centers
• Pay-per-hour servers, storage, and networking
• Rapid provisioning and massive scalability
• Geographically distributed servers for security

Capturing Cloud Benefits

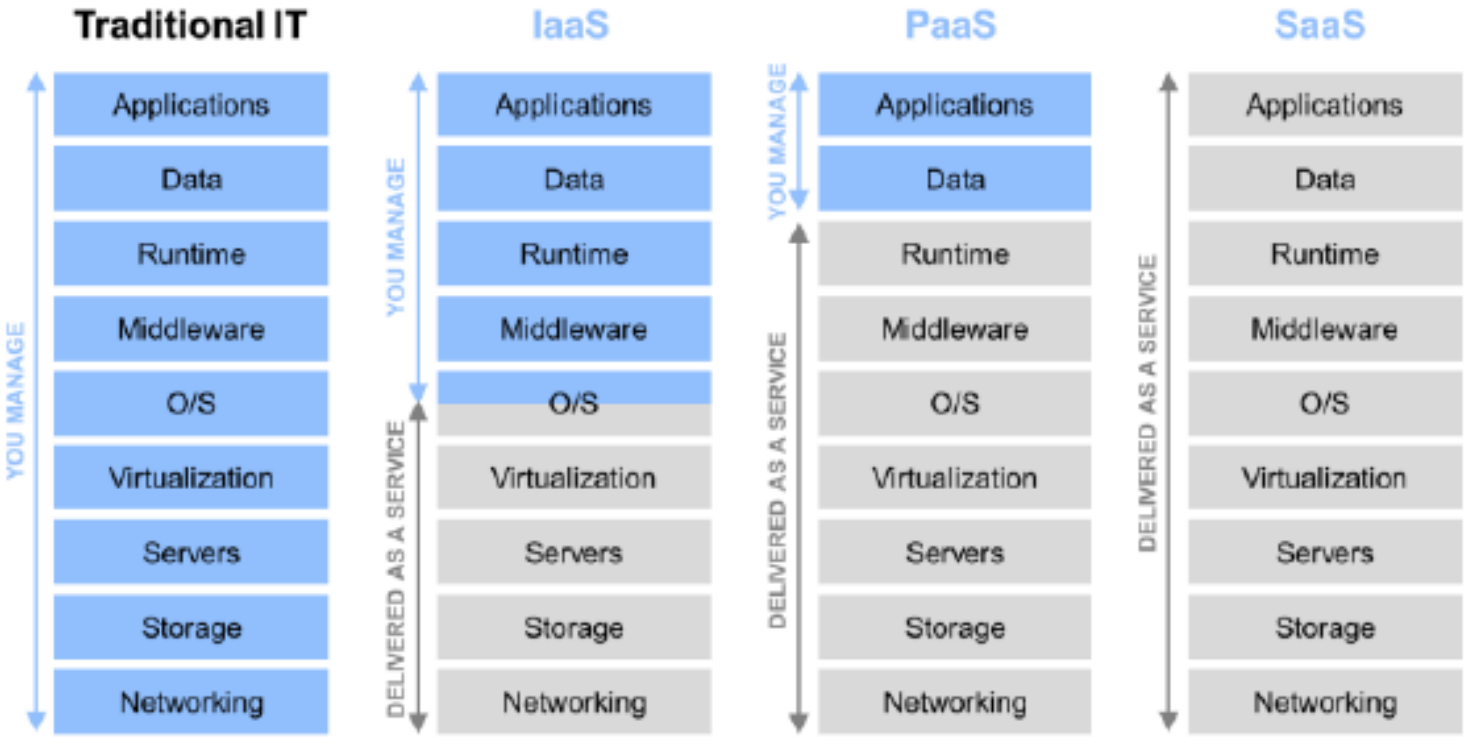

Source: Microsoft.

Cloud-ready workloads

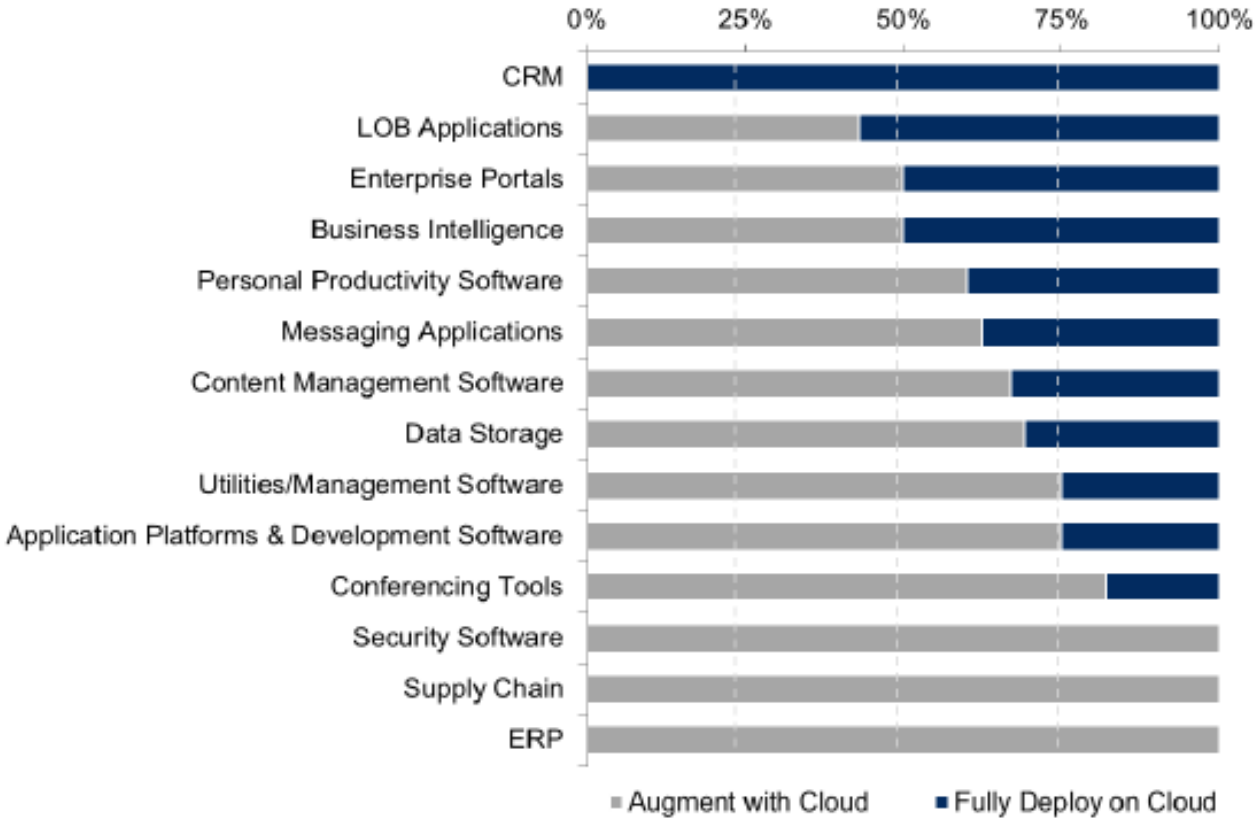

Source: Microsoft survey question "In the next 12-24 months, please indicate if a cloud offering would augment on-premises offering or completely replace it".

Workload patterns optimal for Cloud

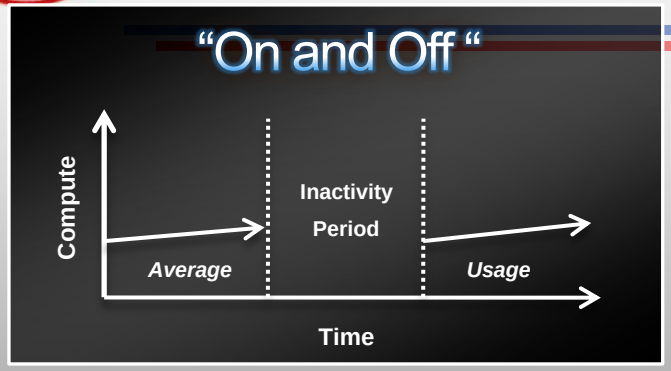

- On & off workloads (e.g. batch job)
- Over provisioned capacity is wasted
- Time to market can be cumbersome

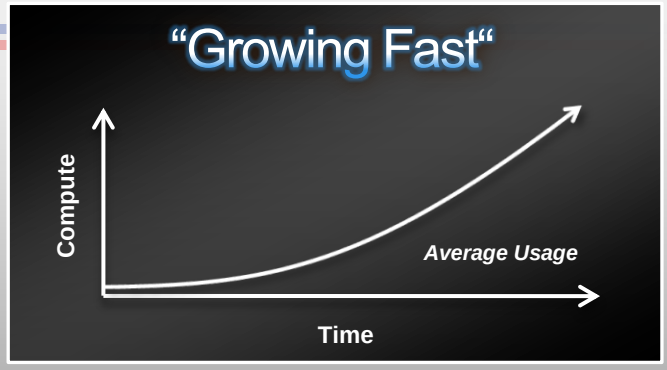

- Successful services needs to grow/scale
- Keeping up w/ growth is big IT challenge
- Complex lead time for deployment

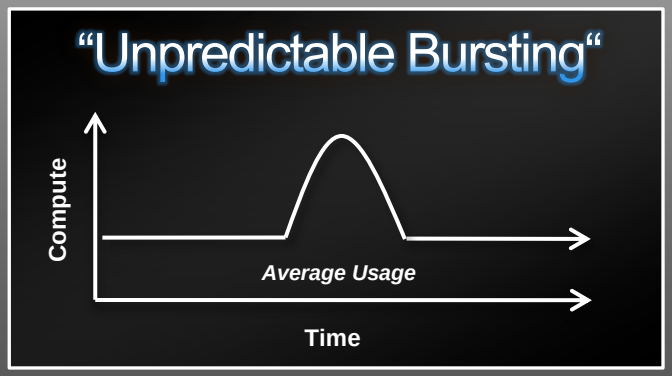

- Unexpected/unplanned peak in demand
- Sudden spike impacts performance
- Can't over provision for extreme cases

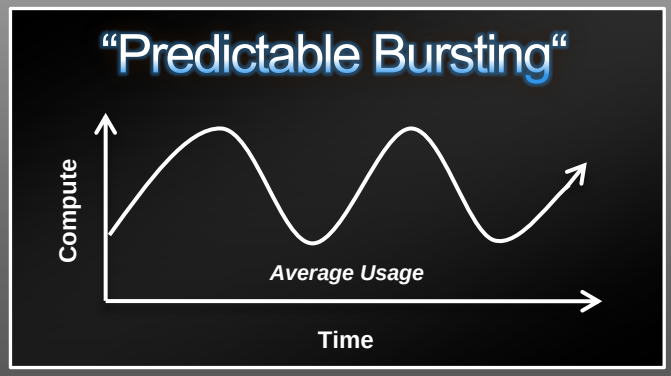

- Services with micro seasonality trends
- Peaks due to periodic increased demand
- IT complexity and wasted capacity

Key Cloud considerations

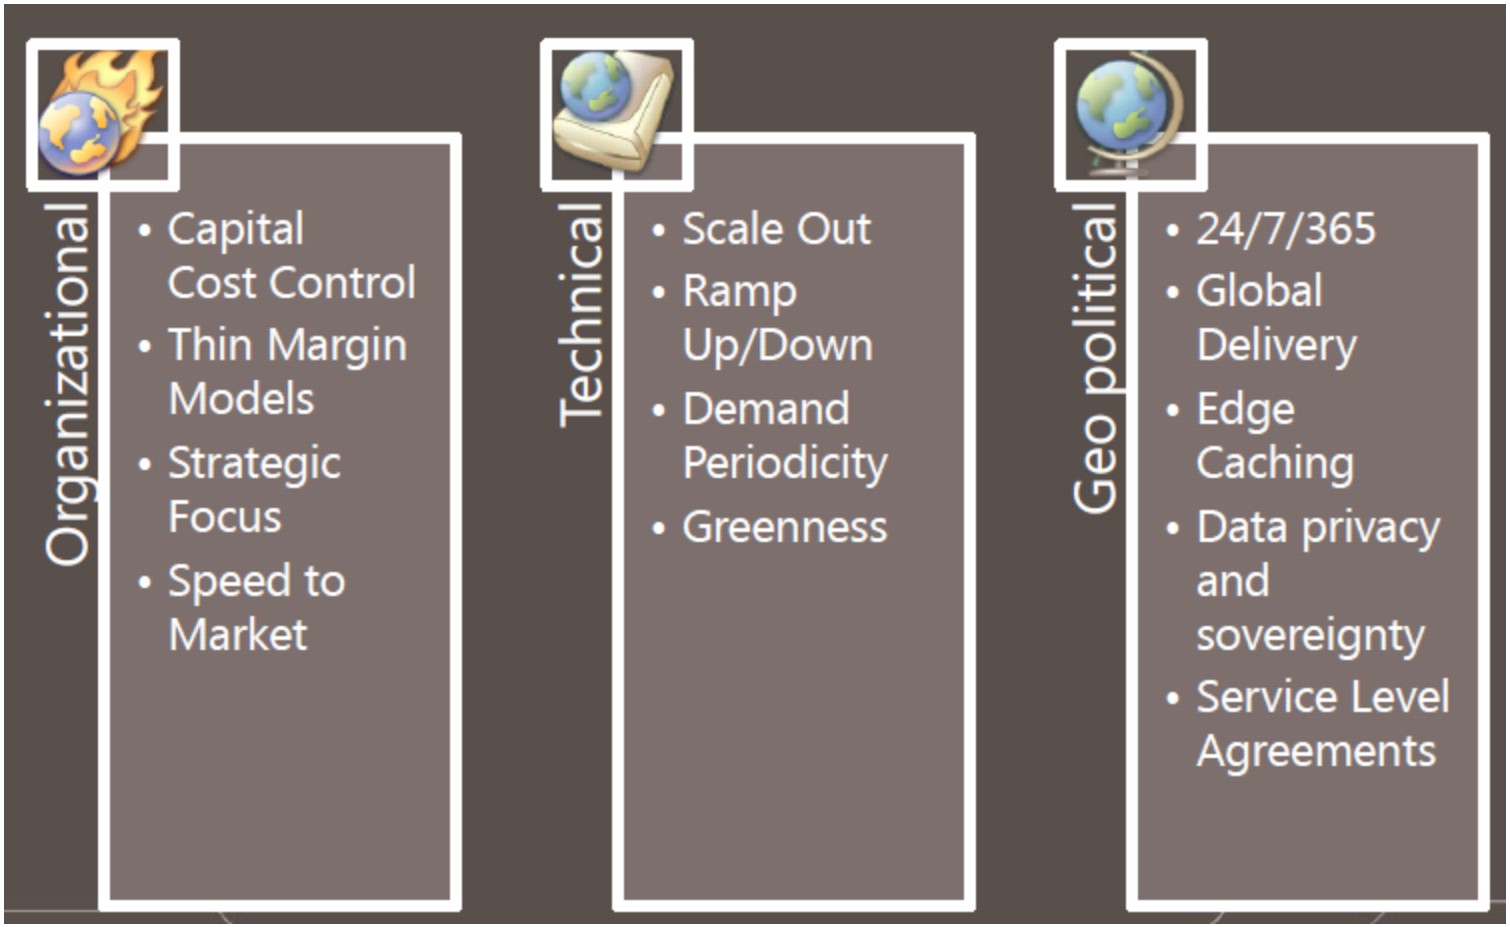

Data Sovereignty, Security

- Does data in the cloud pose a security risk compared to on-premise
- How much does data sovereignty matter
- What data would you never put in the cloud
- Can you trust your offshore vendor
- How about the risk posed by foreign jurisdictions

SLAs

- Monitoring and enforcement is a hard problem
  - US General Services Administration, part of federal government, came up with a RFQ that demands a 99.95% uptime per month

Key sells to your CFO

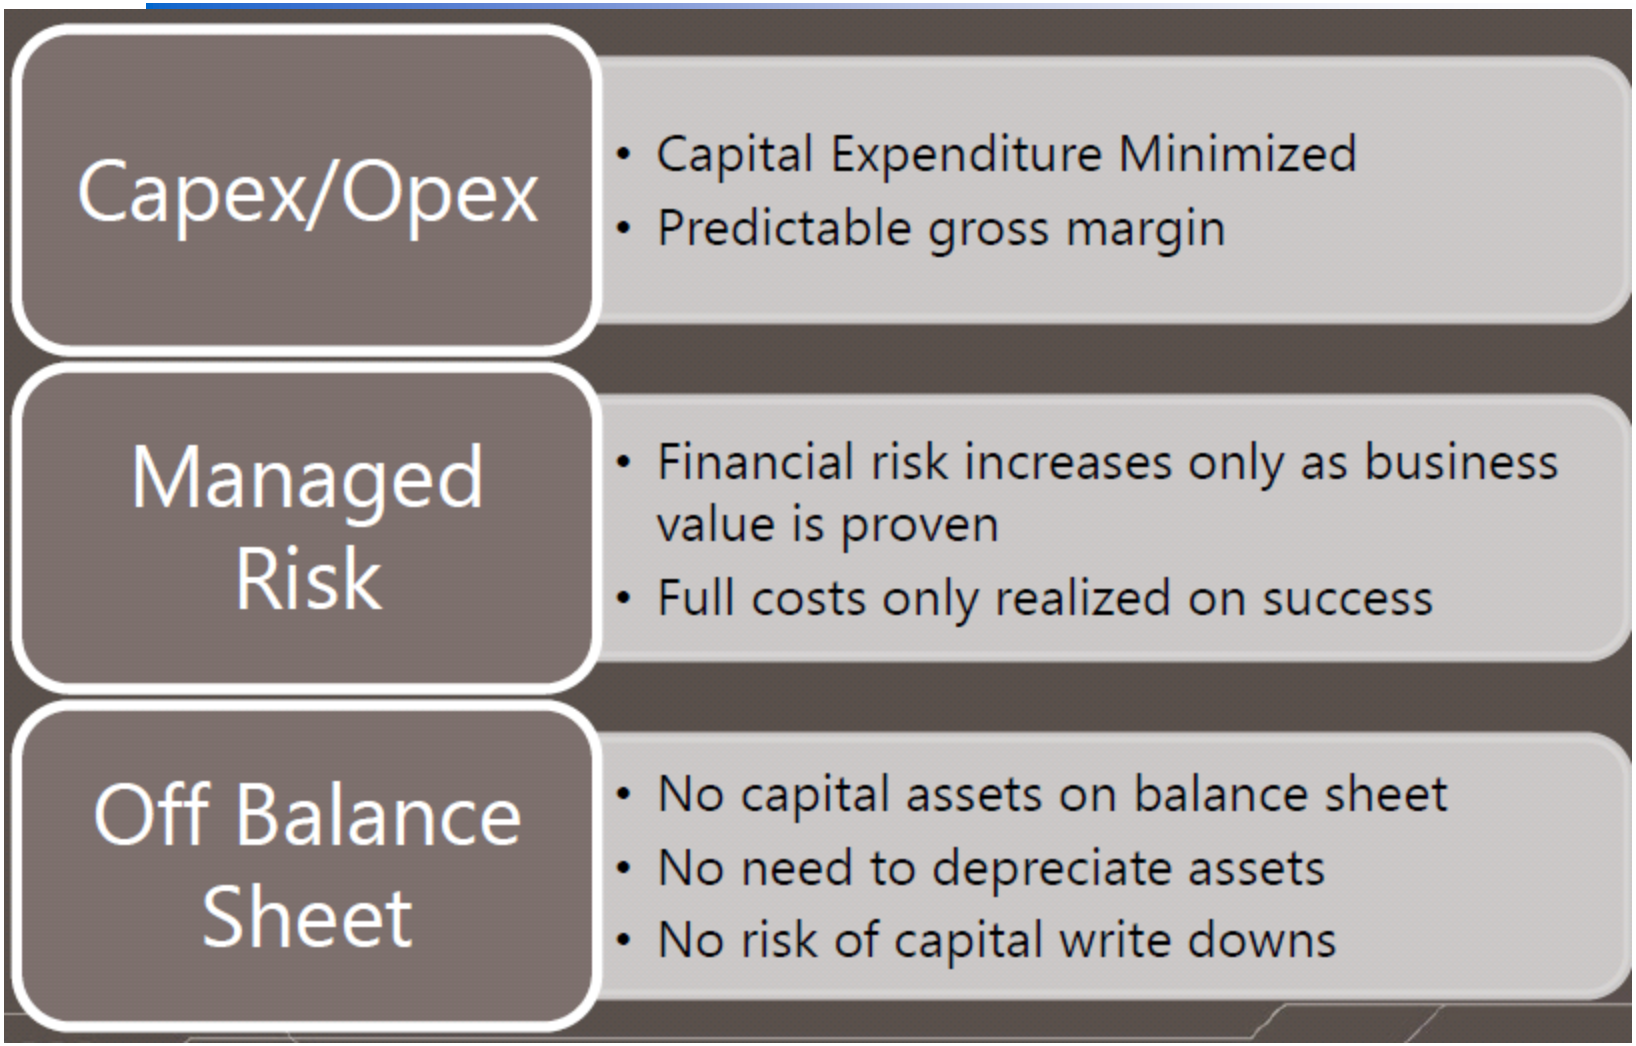

Key sells to your CEO

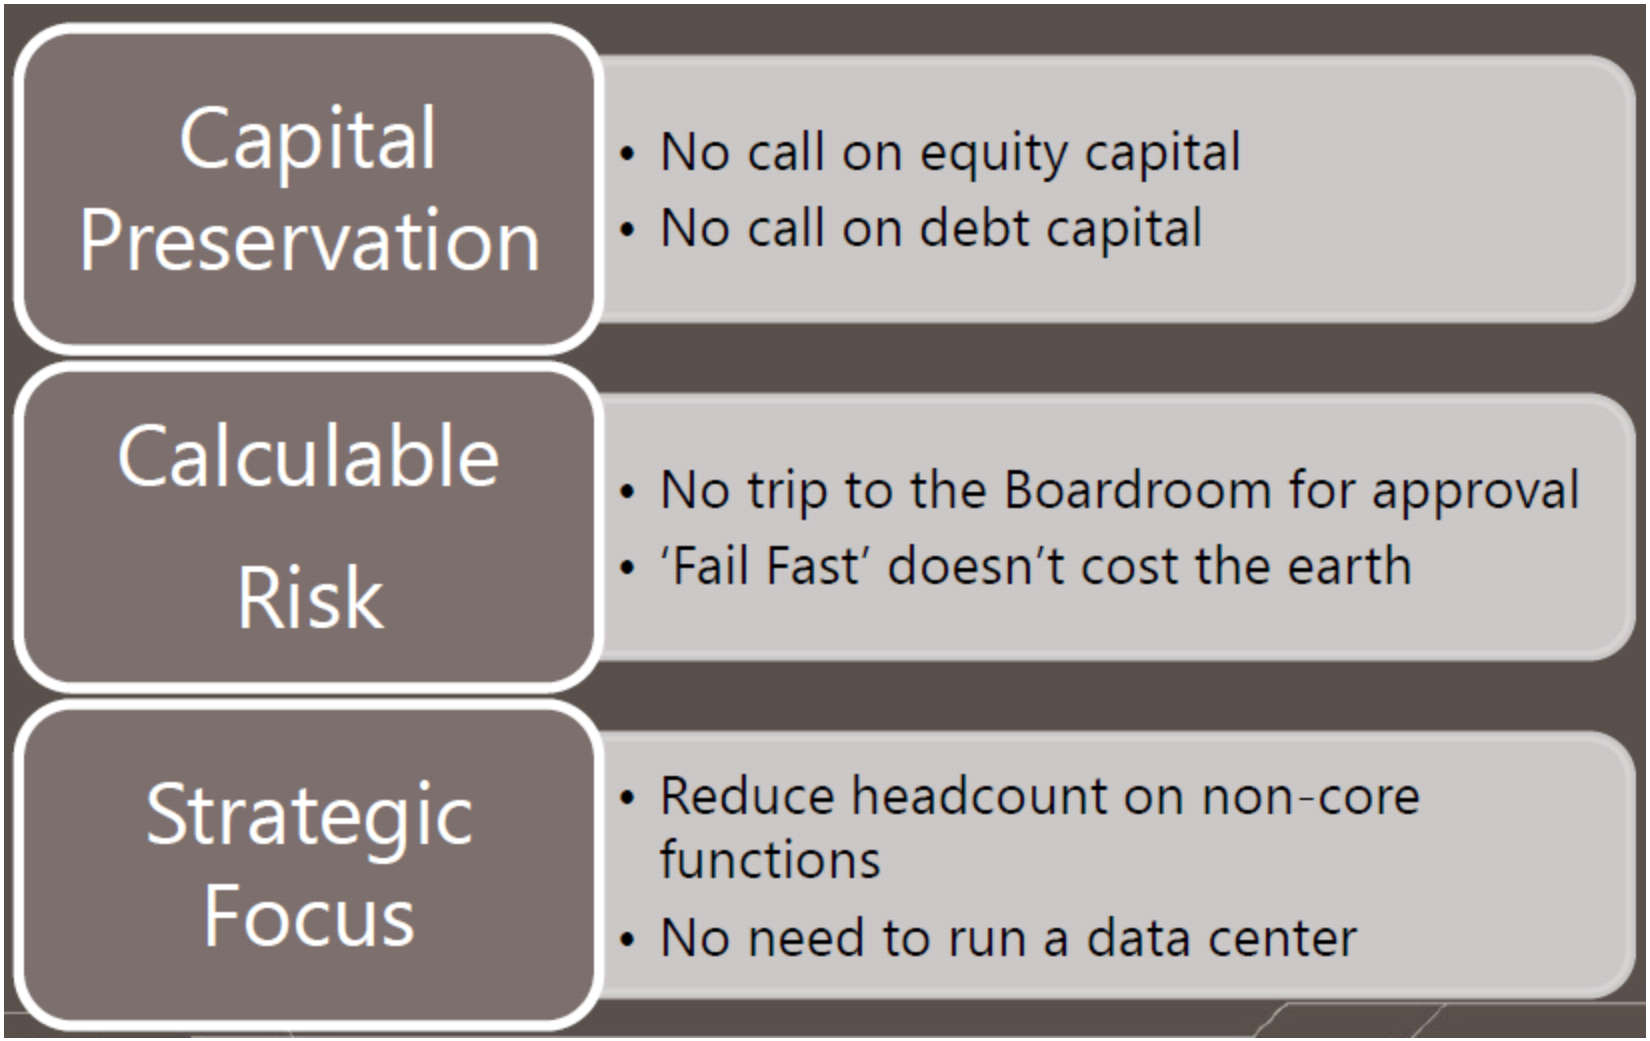

Experience matters

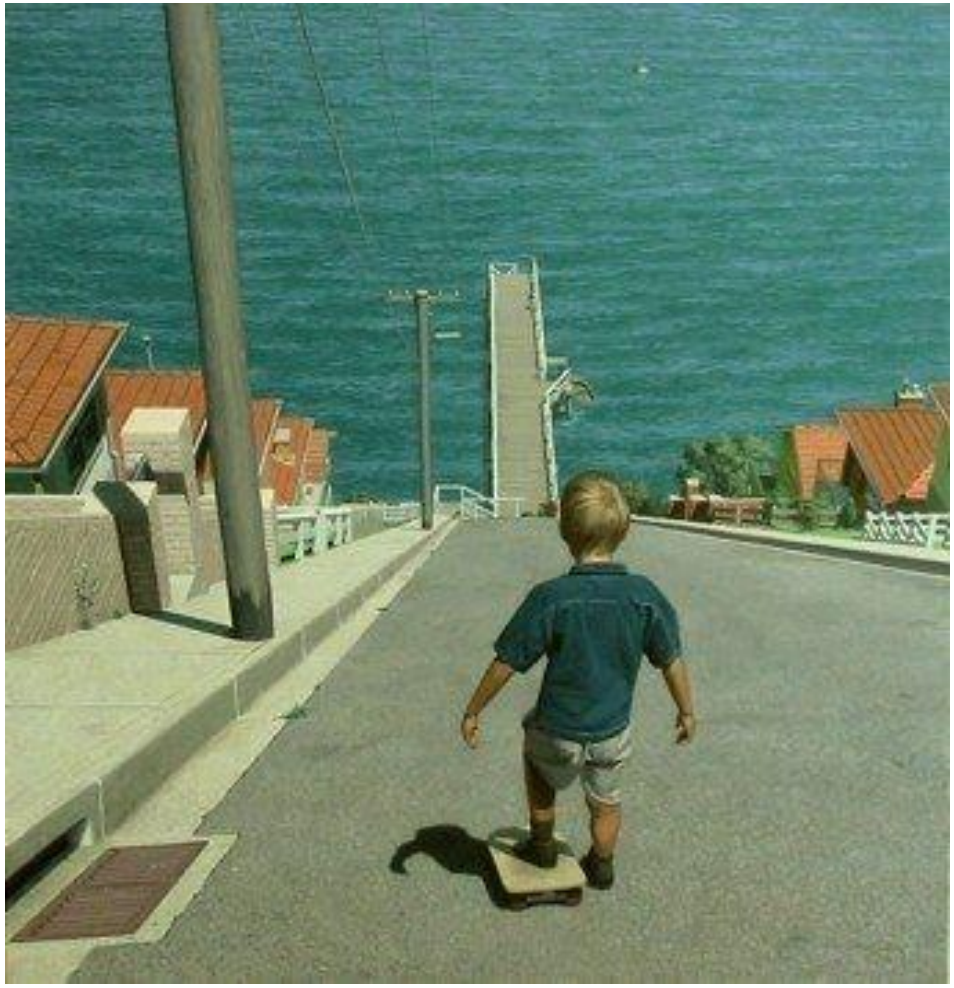

Part 7

HOW BIG IS THE CLOUD?

Cloud Revenues

- ***SaaS***: Easy to measure, SaaS firms have been around for some time. In a new study Forrester Research, a consultancy, estimates that these services generated sales of $11.7 billion in 2010
- ***PaaS***: Easy to measure, since there are only a few providers and their offerings have not really taken off yet. Forrester puts revenues at a mere $311m
- ***IaaS***: Hardest to measure. Firms not reporting revenue numbers..
  - *Amazon* only reveals that it now stores more than 200 billion digital "objects" and has to fulfill nearly 200,000 requests for them per second
  - *Rackspace* says it operates nearly 64,000 servers globally, but notes that only some are used for IaaS

Counting Tanks..

- During the second world war, the allies were worried that a new German tank could keep them from invading Europe. Statisticians were called in to help..
  - They assumed that the Germans, a notoriously methodical lot, had numbered their tanks in the order they were produced..
  - Based on this assumption, they used the serial numbers of captured tanks to estimate the total
  - The number they came up with, 256 a month, was low enough for the allies to go ahead with their plans and turned out to be spot-on. German records showed it to be 255!

Counting AMIs

- Using the Tank approach, Guy Rosen, a blogger, and Cloudkick, a San Francisco start-up which was recently acquired by Rackspace, have come up with a detailed estimate of the size of at least part of Amazon's cloud
- Mr Rosen decrypted the serial numbers of Amazon's "virtual machines", the unit of measurement for buying computing power from the firm
- Alex Polvi, the founder of Cloudkick, then used these serial numbers to calculate the total number of virtual computers plugged in every day
- This number is approaching 90,000 for Amazon's data centers on America's East Coast alone

Counting AMIs

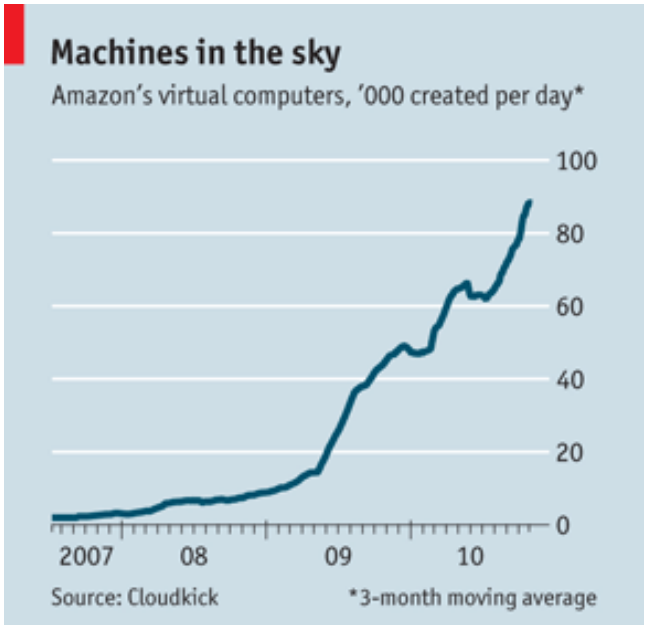

- The results suggest that Amazon's cloud is a bigger business than previously thought..
- Amazon is by far the market leader with a share of between 80% and 90%, according to Mr Bias. Assuming that Cloudkick's and Mr Bias' numbers are correct, revenues generated by computing infrastructure as a service in 2010 may exceed $1 billion

How big is the Cloud?

- How big will it be in ten years?
  - If you count web-based applications and online, Forrester predicts that it will grow to nearly $56 billion by 2020
  - Raw computing services, the core of the cloud, is smaller. Forrester, reckons it will be worth $4 billion in 2020

So, is that all?

- Nooo…
- The heart of cloud computing also revolves around a ***new*** pattern of distributing computing power
  - Very mysterious, very new algorithm
  - In fact, it's the algorithm that made Google what it is..

Part 8

THE NEW ALGORITHMS & ARCHITECTURES

Glass Empire

- In the early 2000s, taking advantage of the failure of some telecom operations, Google began buying up abandoned fiber-optic networks
  - Through acquisition, swaps, and actually laying down thousands of strands, the company has built a mighty empire of glass

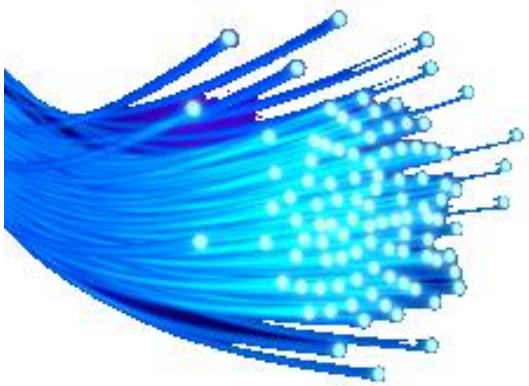

Google File System

- Google has built a software system that allows it to manage its countless servers as if they were one giant entity
  - Its in-house developers can act like puppet masters, dispatching thousands of computers to perform tasks as easily as running a single machine
- In 2002 its scientists created ***Google File System***, smoothly distributes files across many machines

Google MapReduce

- ***MapReduce***, a Google system for writing cloud-based applications, was so successful that an open source version called Hadoop has become an industry standard

Data Storage and Analysis

- While the storage capacities of hard drives have increased massively over the years…
- The rates at which data can be read from drives have not kept up!
  - 1990:
    - Storage: 1,370 MB, xfer: 4.4 MB/s
  - 2010:
    - Storage: 1TB, xfer: 100 MB/s
      - 2.5 hours to read all the data!
      - So, what to do??

We need to read from multiple disks in parallel!

- If we had 100 drives, each holding 1/100th of the data:
  - Read 1TB in under 2 minutes!

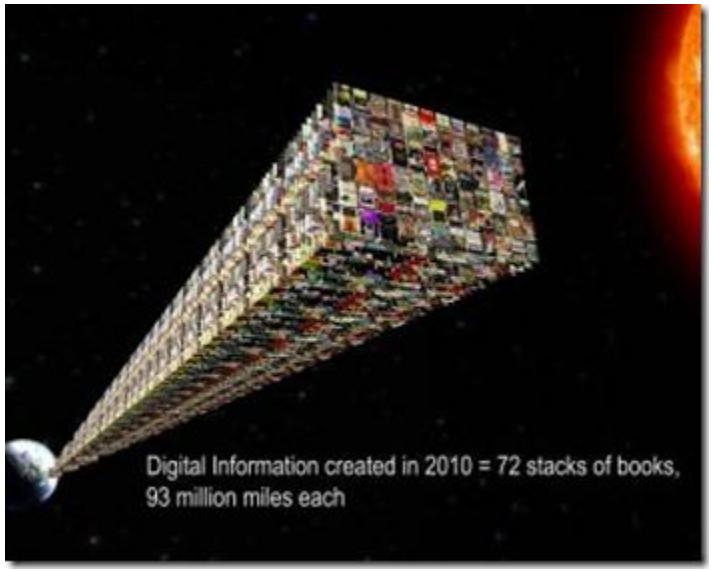

SI decimal prefixes		Binary usage
Name (Symbol)	Value	
kilobyte (kB)	10^3	2^{10}
megabyte (MB)	10^6	2^{20}
gigabyte (GB)	10^9	2^{30}
terabyte (TB)	10^{12}	2^{40}
petabyte (PB)	10^{15}	2^{50}
exabyte (EB)	10^{18}	2^{60}
zettabyte (ZB)	10^{21}	2^{70}
yottabyte (YB)	10^{24}	2^{80}

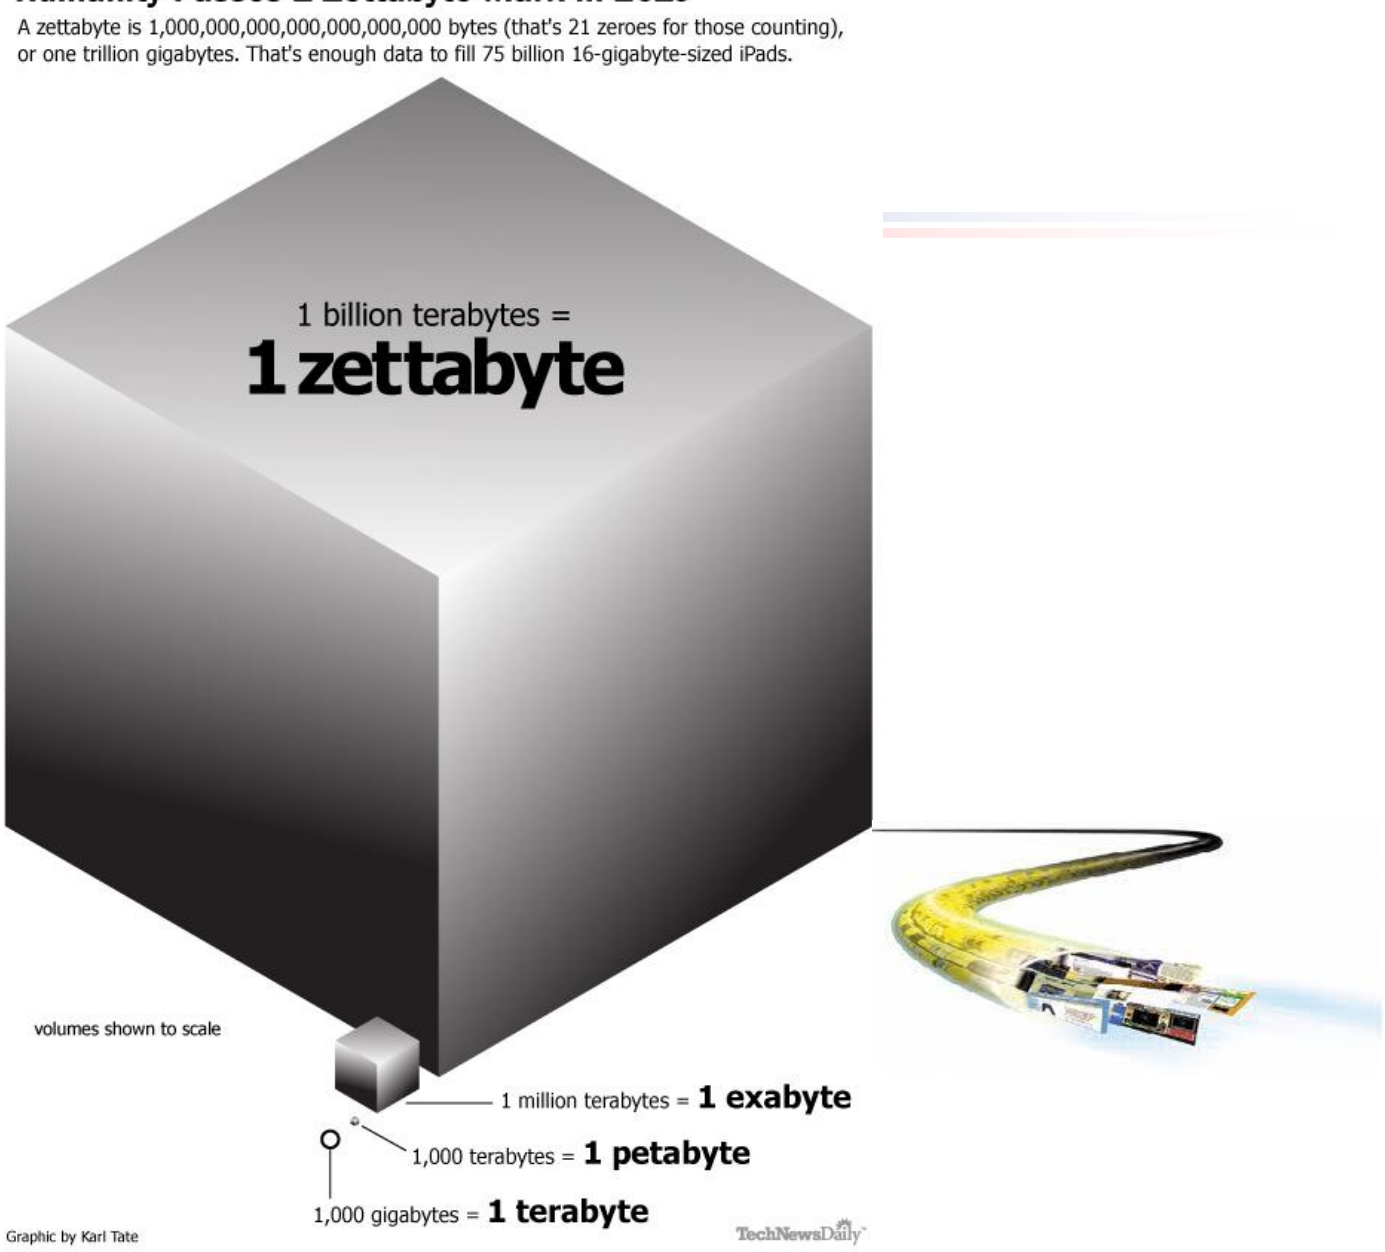
Humanity Passes 1 Zettabyte Mark in 2010
A zettabyte is 1,000,000,000,000,000,000,000 bytes (that's 21 zeroes for those counting), or one trillion gigabytes. That's enough data to fill 75 billion 16-gigabyte-sized iPads.
1 billion terabytes =
1 zettabyte
volumes shown to scale
1 million terabytes = 1 exabyte
1,000 terabytes = 1 petabyte
1,000 gigabytes = 1 terabyte
Graphic by Karl Tate
TechNewsDaily

MapReduce

- Since the data is too huge to move through the network..
  - We leave the data on a distributed File System
  - And we move the computation to the Data!

Hadoop is Open Source MapReduce

- A simple programming model for processing large dataset on large set of computer clusters
  - Batch query processor, with the ability to get results in a reasonable time
  - Gives you the tools to analyze TBs worth of data and the opportunity to innovate with data
- Library deals with messy details
  - Distributing the computation
  - Dealing with hardware failures

The Distributed Computing Challenges that Hadoop solves

- **Cheap nodes fail, especially if you have many**
  - Mean time between failures for 1 node = 3 years
  - Mean time between failures for 1000 nodes = 1 day
  - **Solution**: Build fault-tolerance into system
- **Commodity network = low bandwidth**
  - **Solution**: Push computation to the data
- **Programming distributed systems is hard**
  - **Solution**: Data locality (co-locate data & compute)
  - **Solution**: Gracefully handle partial failure
  - **Solution**: Data-parallel programming model: users write "map" & "reduce" functions, system distributes work and handles faults

Cloud is a different kind of Client/Server

- Client/Server: Moves data to the computation

- Functional style builds the computation at design time, but the data only plugs in at runtime

The shift from enumerable to observable

- Contrasted to an enumerable interface, an observable interface has a life of its own and needs to be waited upon and even throttled at times
  - Think of data being downloaded from the Cloud..
  - Enumerable: Drinking from a cup (pull)
  - Observable: Drinking from a fire hose (push)

Cloud computing Architectures

❖ ***Pull*** Client/Server ➔ ***Push*** Client/Server

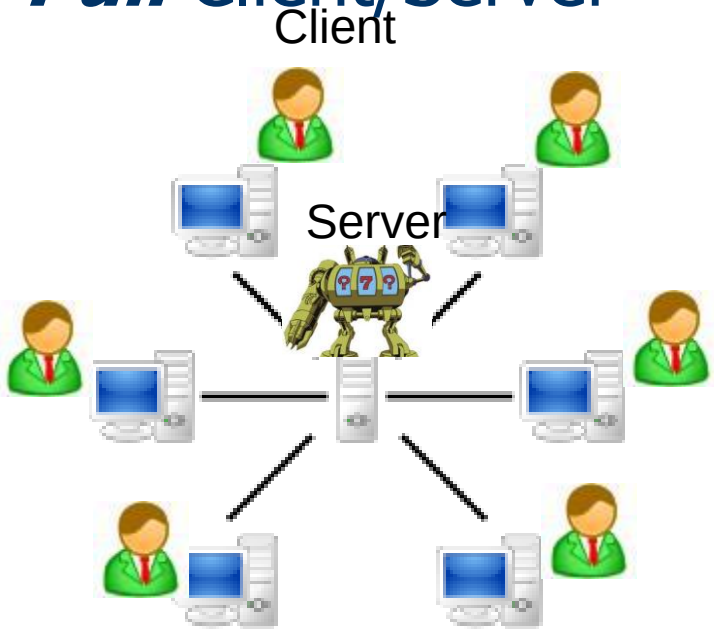

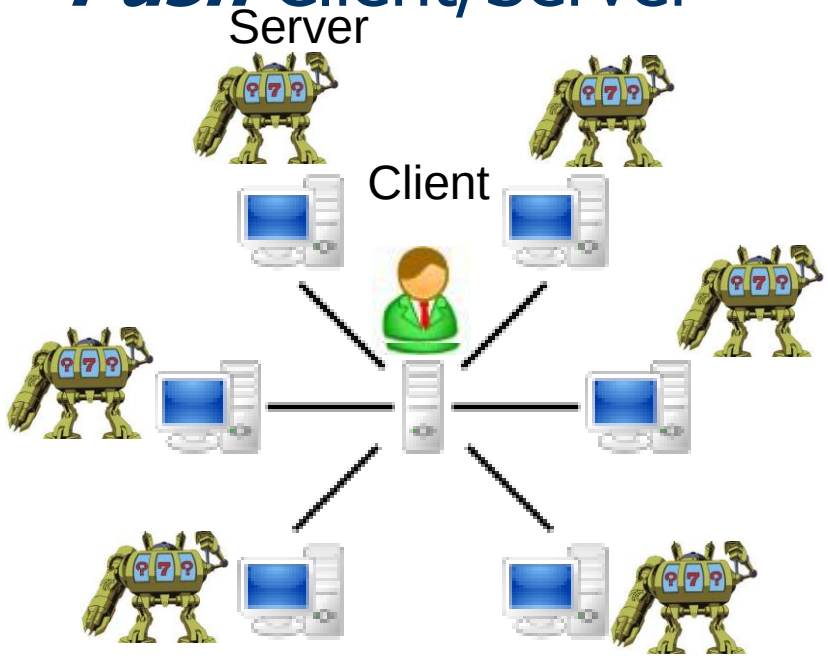

Real-time

Client *pulls* data

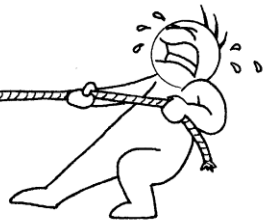

Batch

Server(s) *push* data to client, or client *pushes* computation to server(s)

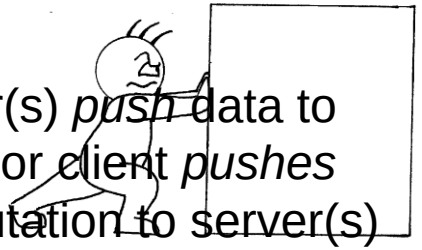

Push of what?

- Dealing with the Future: Pushing functionality
  - Combining event data from multiple sources e.g.: **search results** from Google or **sensor data** from accelerometer, gyro, magnetometer or temperatures
  - Grouping future data e.g. **tweets** by topic or user, or **stock prices** by delta or liquidity
- Dealing with the Past: Pushing data
  - Moving average of a series of values e.g. **stock prices** or **service level agreements** for average latencies or downtime
  - Grouping past data e.g. **tweets** by topic or user, or **stock prices** by delta or liquidity


PAST
FUTURE
copyright © 2007 Bill Frymire

Part 9

THE SOCIAL DIMENSIONS OF CLOUD COMPUTING

Big Data

datasets whose size is beyond the ability of typical database software tools to capture, store, manage, and analyze

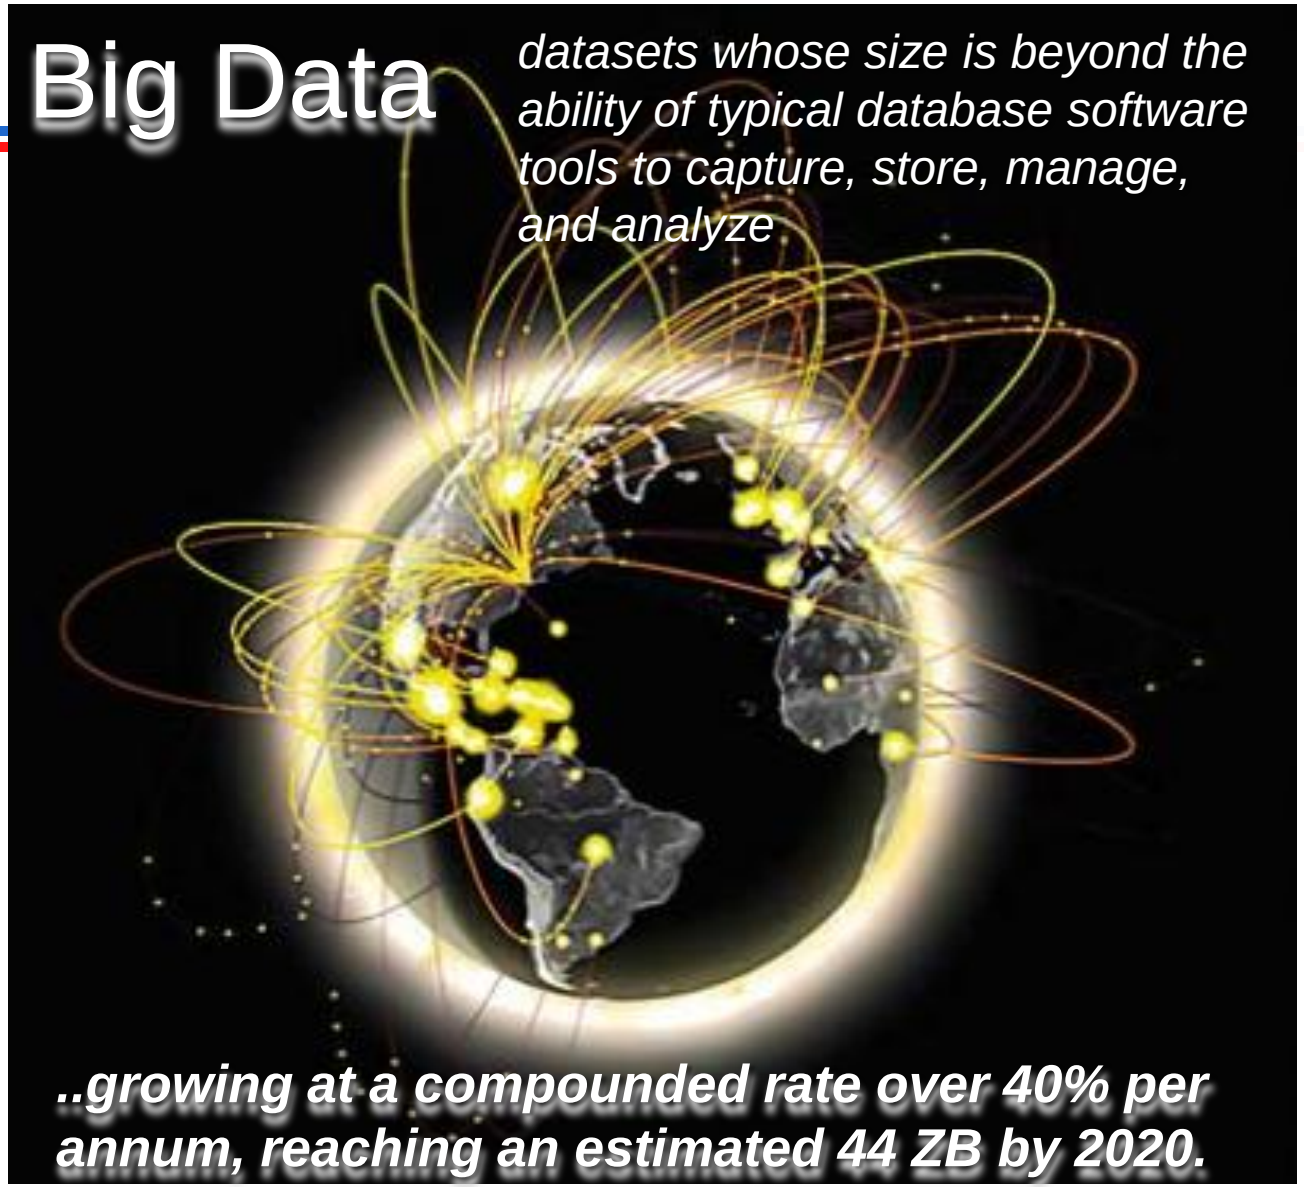

..growing at a compounded rate over 40% per annum, reaching an estimated 44 ZB by 2020.

BigData


THE AVERAGE PERSON TODAY PROCESSES MORE DATA IN A SINGLE DAY THAN A PERSON IN THE 1500'S DID IN AN ENTIRE LIFETIME.
INCEPTION
JULY 16
SOUTH PACIFIC
PALACE
WEST SIDE STORY
BEST MUSICAL REVIVAL
LA CAGE
american IDIOT
PROMISES PROMISES
aerie


EACH OF US NOW LEAVES A TRAIL OF DIGITAL EXHAUST, AN INFINITE STREAM OF PHONE RECORDS, TEXTS, BROWSER HISTORIES, GPS DATA, AND OTHER INFORMATION, THAT WILL LIVE ON FOREVER.
Instead of "find my iPhone," some auto insurance companies are offering a service that may enable parents to "find my teenager." Progressive Insurance, for example, offers the Snapshot, a tracking device that reports on a car's location, acceleration, braking, and distance traveled. Owners who install the device can get a 10 to 15 percent discount on their policy. Privacy activists, however, fear the technology is ripe for abuse.
http://humanfaceofbigdata.com/


DATA TRANSPARENCY, AMPLIFIED BY TWITTER AND FACEBOOK, HAS LED TO MASS MOVEMENTS ON A SCALE NEVER SEEN BEFORE.
From the Greek protests to the Arab Spring, a global citizenry is tapping into social media to create an unstoppable cascade of change in politics and government. In this photo, supporters of Antonis Samaras, the leader of Greece's conservative New Democracy party, wave flags during a pre-election speech in Athens on May 3, 2012.
http://humanfaceofbigdata.com/

Brave New World

- It's a Brave New World out there
  - ***We're lucky*** that our jobs are in the thick of it!
  - ***I'm lucky*** that I get to investigate about it and teach it
- https://sites.google.com/site/bumetcs755syllabus/

Part 10

CONCLUSION & FUTURES

The #1 problem Cloud solves

- Many IT leaders today are faced with the problem that 80% of the budget is spent on keeping the lights on, maintaining existing services and infrastructure

Game Changer: Elasticity

- Renting 1 machine for 1,000 hours will be nearly equivalent to renting 1,000 machines for 1 hour in the cloud
- Being able to both scale up and scale down resource intensity nearly instantly enables a new class of experimentation and entrepreneurship

Game Changer: CapEx Elimination

- Elimination of capital expenditure will lower the risk premium of projects, allowing for more experimentation
- Lowers the costs of starting an operation and lowers the cost of failure and exit

Game Changer: Self Service

- Provisioning servers through a simple web portal rather than through a complex IT procurement and approval chain:
  - Enables rapid provisioning and integration of new services
  - Facilitates consumption model ("if you build it, they will come")

Game Changer: Complexity reduction

- Complexity has been a long standing inhibitor of IT innovation
  - From an end-user perspective ***SaaS*** is setting a new bar for user friendly software
  - From a developer perspective ***PaaS*** greatly simplifies the process of writing new applications
  - From a system administrator perspective, ***IaaS*** greatly simplifies management of IT operations

Public or Private?

- **Public clouds**
  - Have all the same architectural elements as private clouds, but bring massively higher scale to bear on all sources of variability
  - The only way to diversify away:
    - Industry-specific variability
    - Full geographic element of time-of-day variability
  - Bring multi-tenancy benefits into effect
- **Private Clouds**
  - Data sovereignty

Cloud driving CS innovation

- ***Microsoft*** data mining its OS
- ***Facebook*** data mining social data
- ***Google*** data mining the Web
- ***Twitter*** data mining tweets

Long-term migration to Public Cloud

- **Driven by economies of scale**:
  - Larger datacenters can deploy computational resources at much lower cost than smaller ones
  - Demand pooling improves resource utilization
  - Multi-tenancy lowers application maintenance labor costs for large public clouds
  - Unparalleled levels of elasticity and agility that will enable exciting new solutions and applications.

Cloud Future

- **The economic benefit of public cloud will grow over time**
  - Economies of scale will kick in for Public Clouds, the cost premium on private clouds will increase
- **Remaining barriers to cloud adoption will begin to fall**
  - Many technology case studies show that, over time, concerns over issues like compatibility, security, reliability, and privacy will be addressed

Cloud Verticals

- Medical Cloud: Alaa
- Financial Cloud: Madhura
- Government Cloud: Dino

Health Care

Alaa Abou Mahmoud

Healthcare is “dying of thirst in an ocean of data”

90%
of the world’s data was created in the last two years

80%
of the world’s data today is unstructured

1 Trillion
connected devices generate 2.5 quintillion bytes data / day

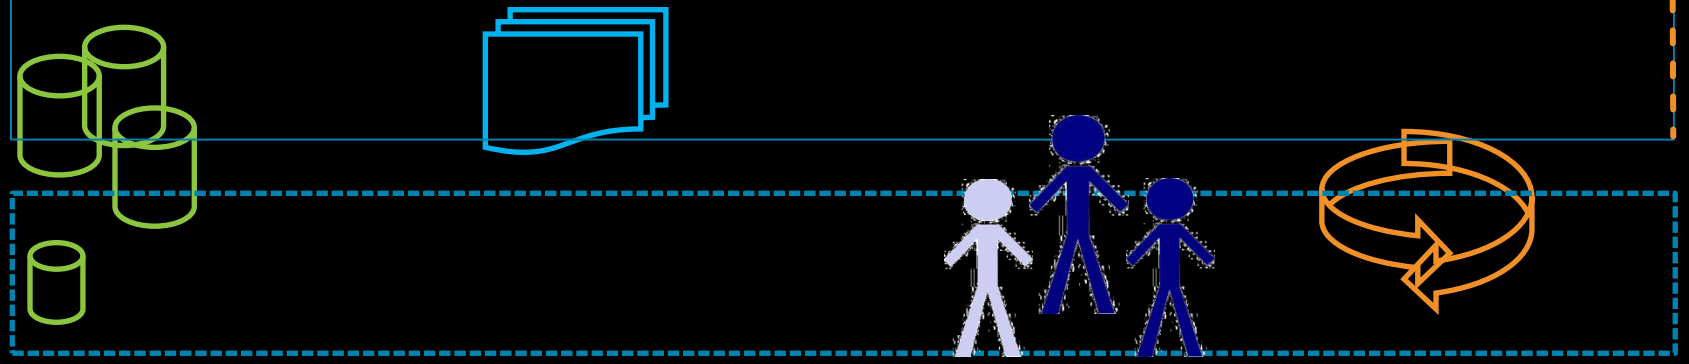

2X
medical information is doubling **every 5 years**

75
new clinical trials start **every day** in the US alone

$750B
or 30 cents of every dollar spent on healthcare in the US is **wasted**

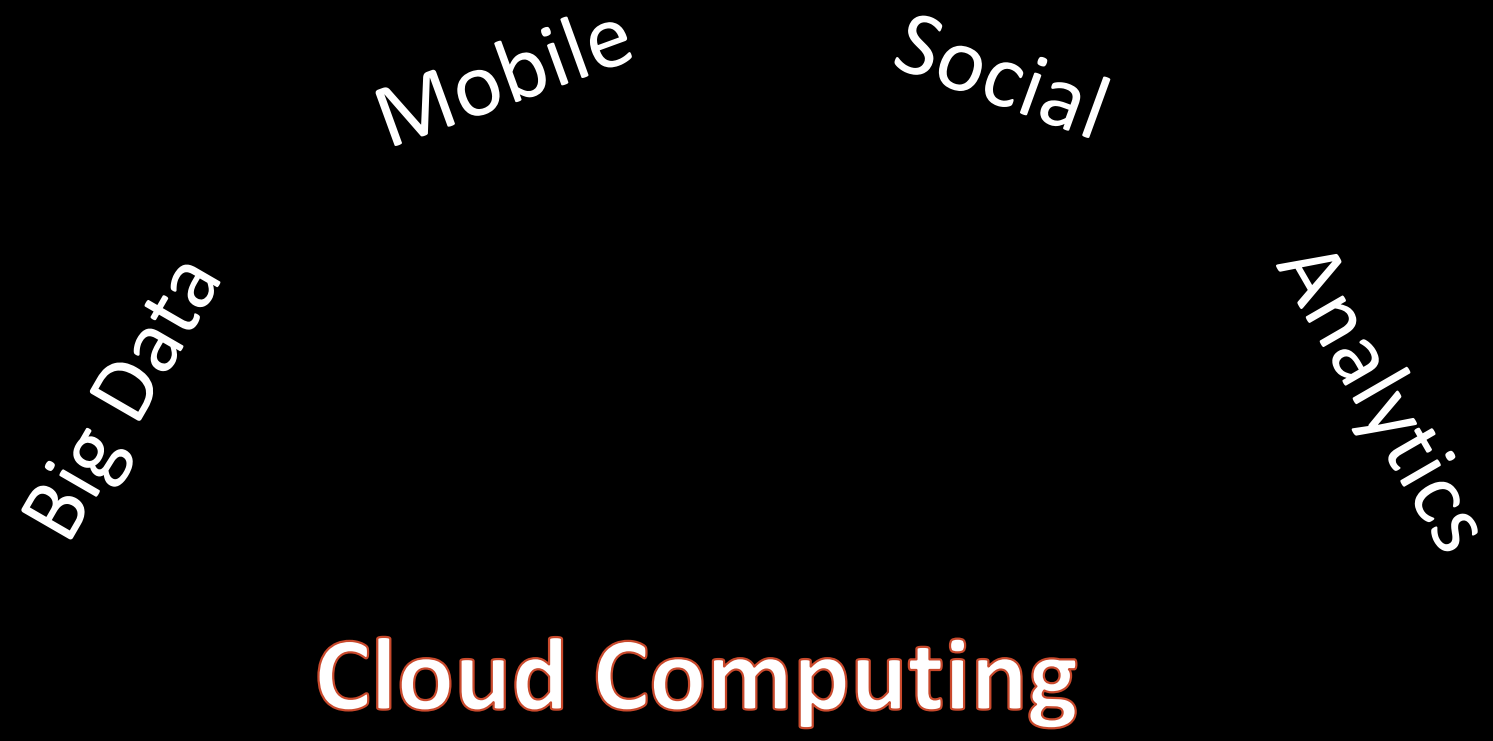
Mobile
Social
Big Data
Analytics
Cloud Computing

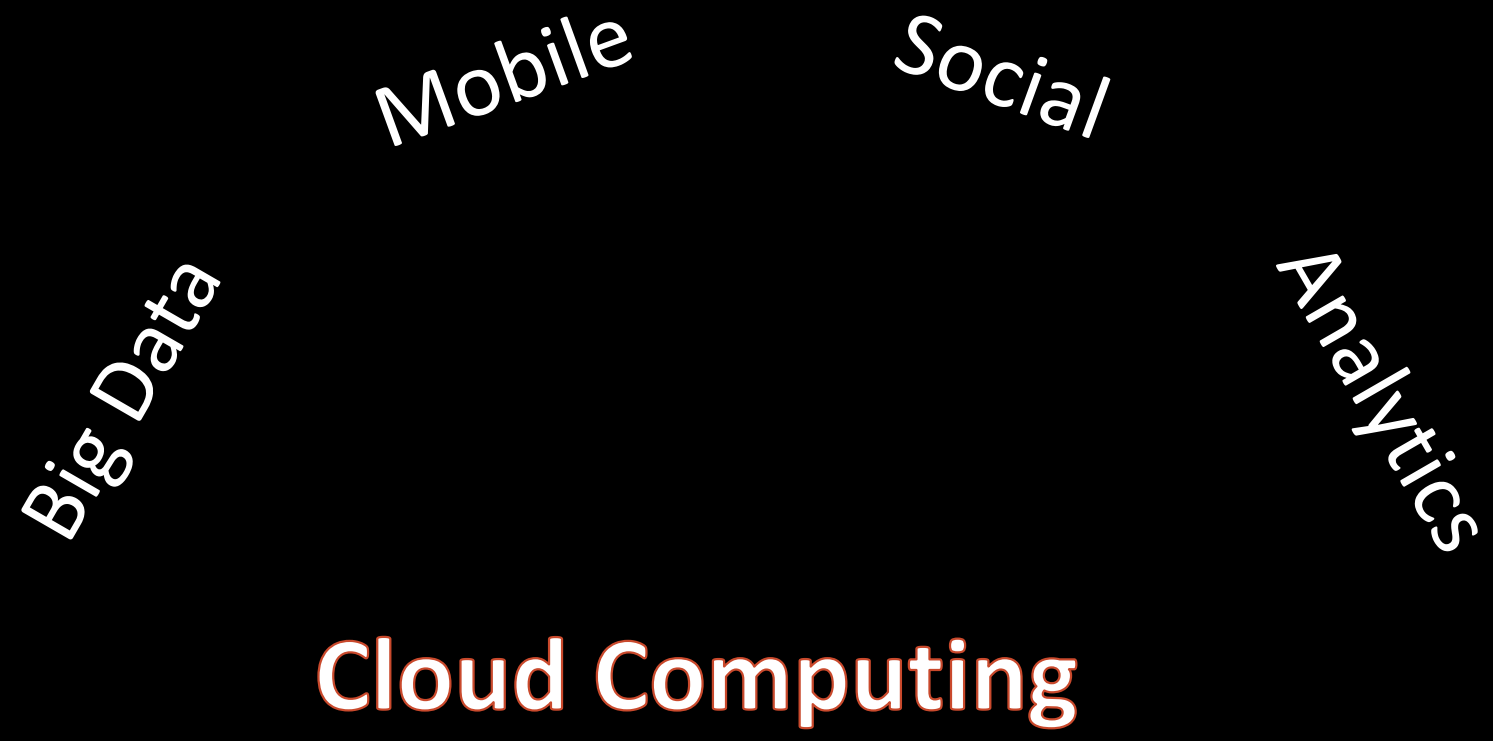
Mobile
Social
Big Data
Analytics
Cloud Computing

Data volume is expanding at an incredible rate
…data will grow 800% in the next five years
…Unstructured data grows 10-50X faster than structured

Data is getting more social. . .
…20M articles on Wikipedia
…30B pieces of Facebook content are shared monthly
…There are 156M public blogs

There are over 2.3B people on the Web today ...
… and a trillion connected objects – cars, appliances, cameras, roadways, pipelines

Big Data: this is just the beginning

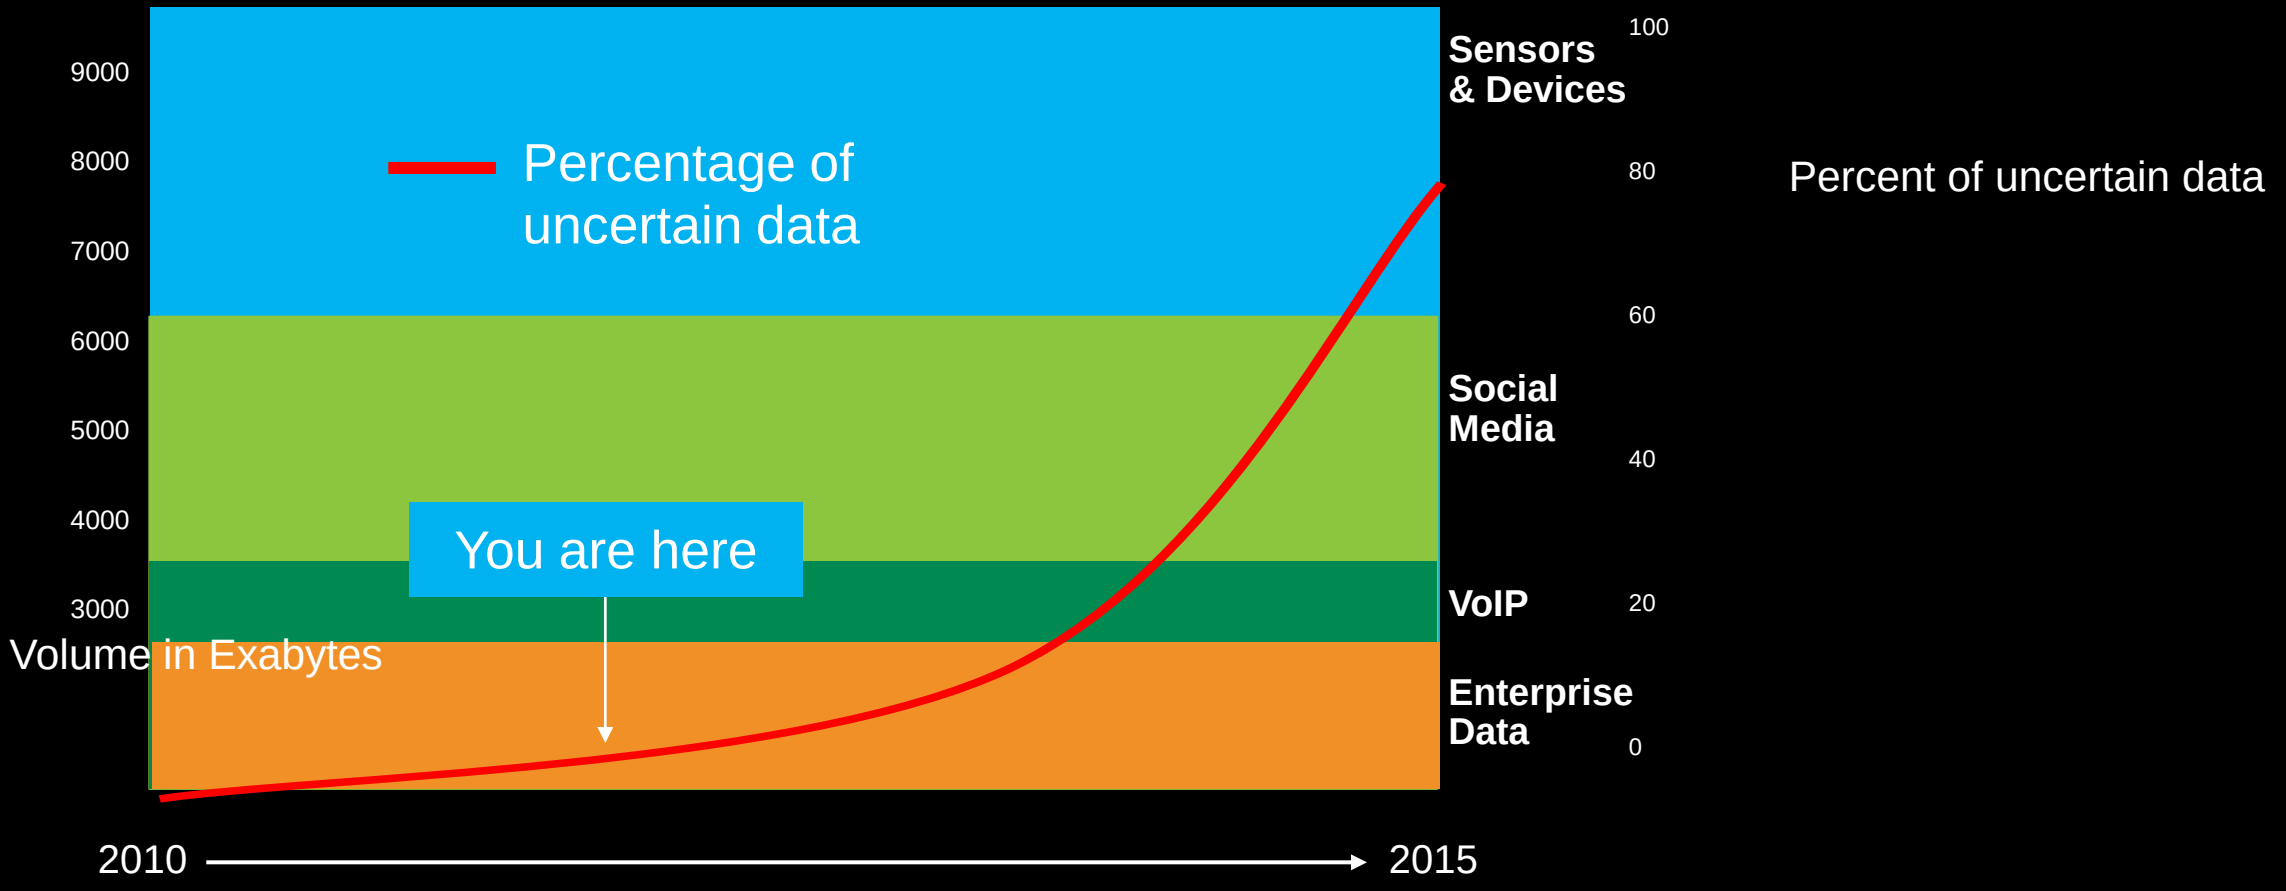

Source: IBM Global Technology Outlook - 2012

Advantage of storing Medical Data in the Cloud

- Cost effective
- Flexible
- Secure
- Better and faster disaster recovery

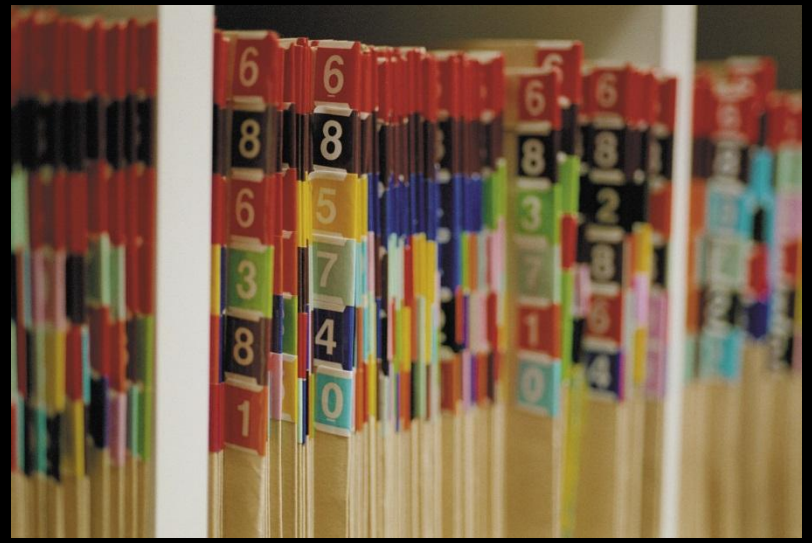

Advantages of storing medical data in the Cloud

- Health care professionals can qualify for government incentives if they implement “meaningful use” of EHR’s under the 2009 Health Information Technology for Economic and Clinical Health (HITECH) Act
- Several cloud services providers are HIPAA compliant which takes the burden of the compliance from the health care provider to the cloud services providers

The cloud is more secure than Enterprise Data Centers

- On-premise environment users experience an average of 61.4 attacks while service provider environment customers suffered and average of 27.8 attacks

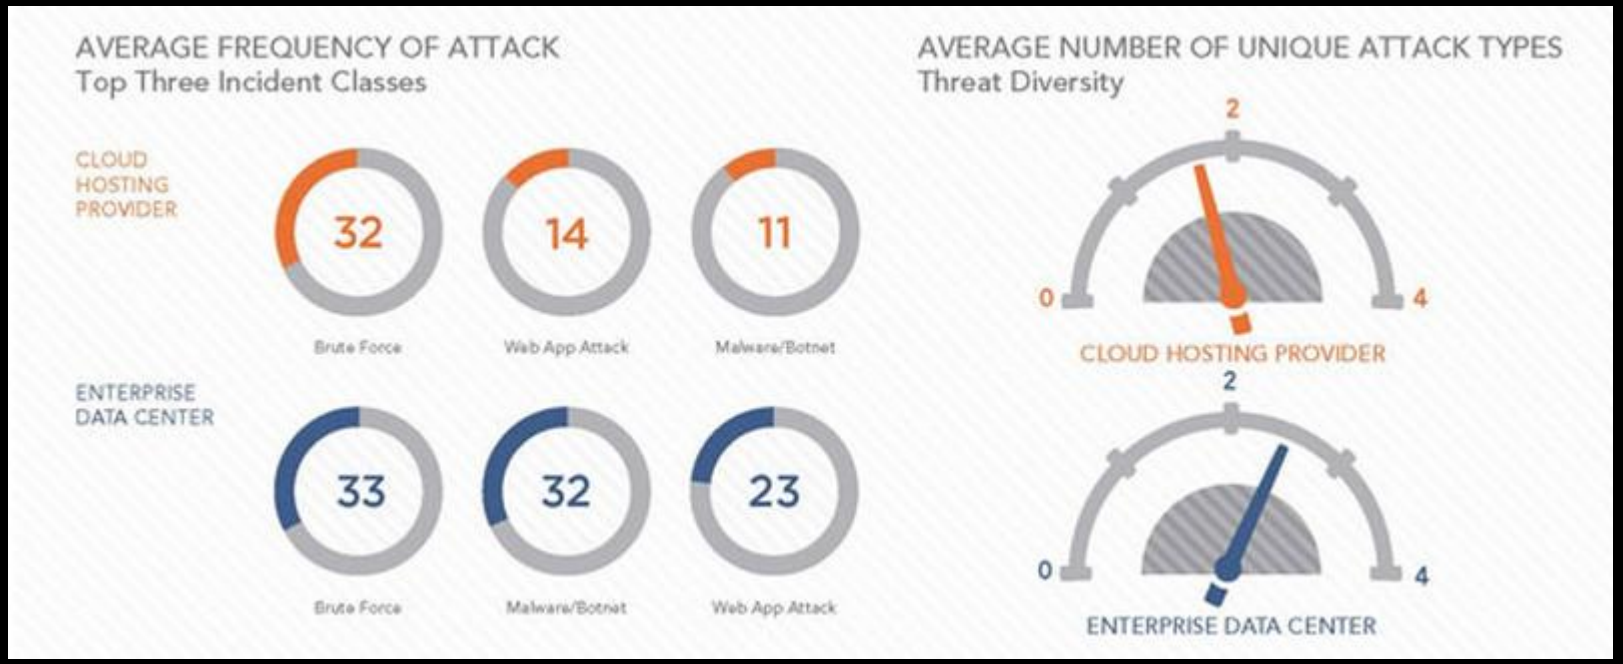

Source: Alert Logic – State of Cloud Security Report 2013

Advantage of storing Medical Data in the Cloud

- Medical imaging data is too big to maintain within the enterprise.
- Health care providers are moving to the cloud to store their medical imaging data with medical imaging data archiving storage providers such as InSite One Inc.

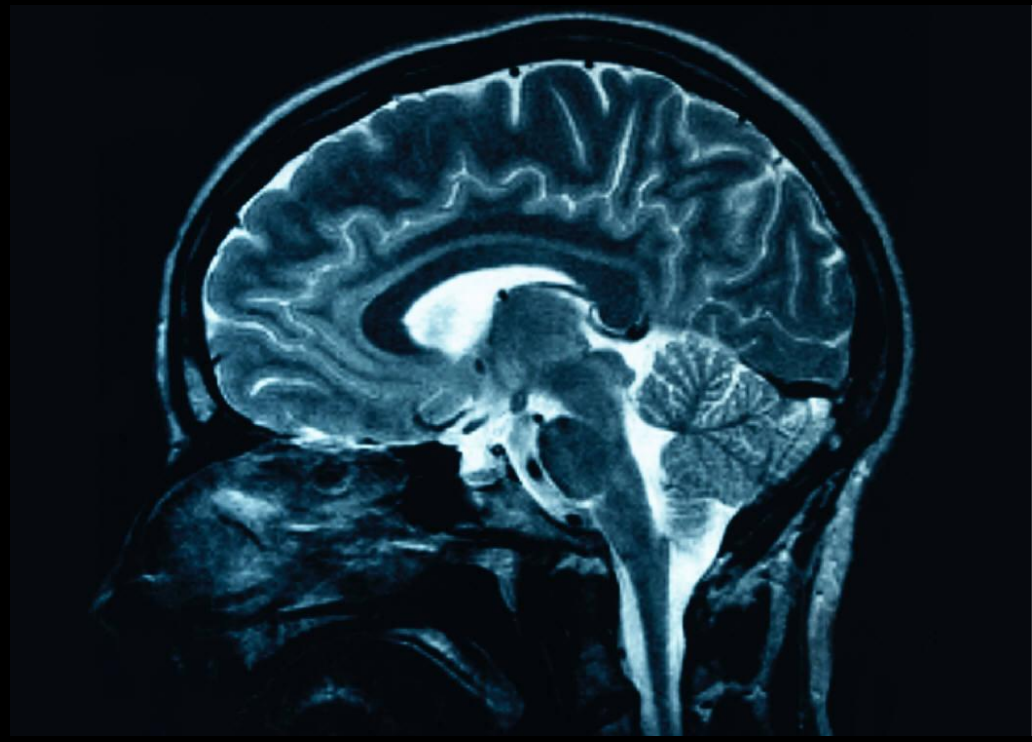

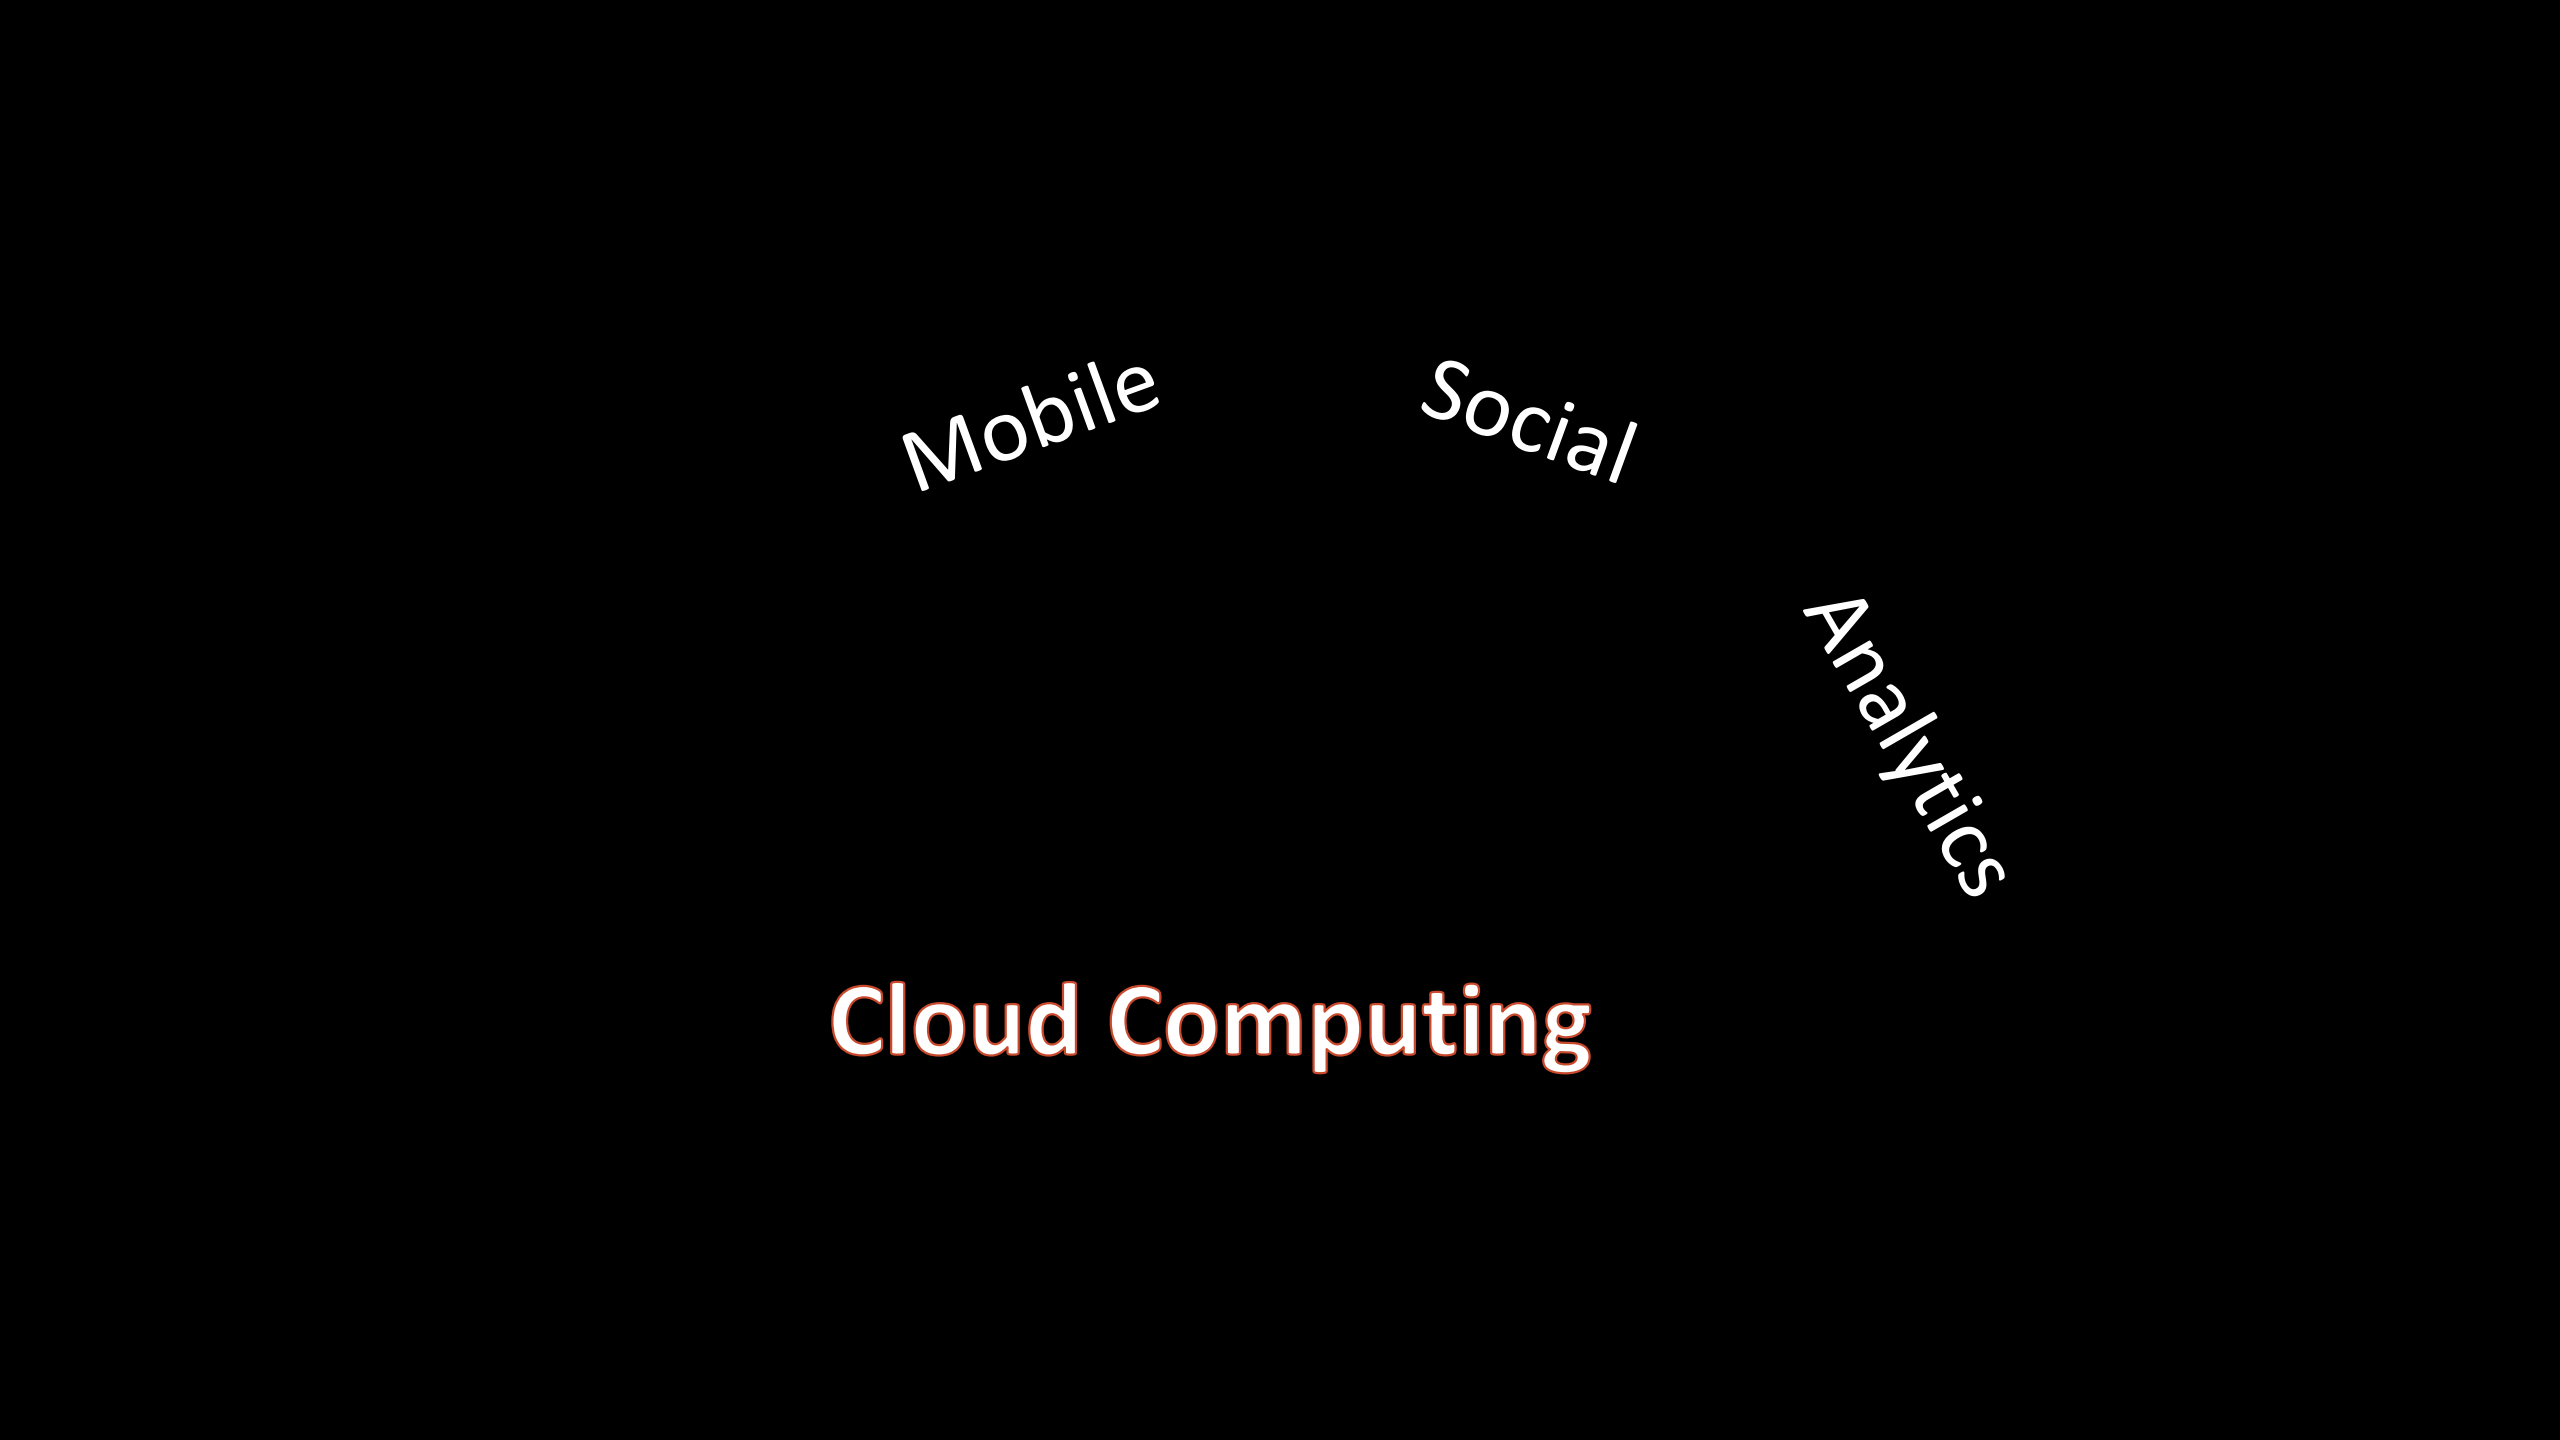
Mobile
Social
Analytics
Cloud Computing

Advantage of storing Medical Data in the Cloud

- The cloud was created to be platform-agnostic.
- It was created with the vision that users will be able to access their data and computing services from anywhere and using any type of device.
- Using the cloud greatly facilitates the usage of mobile devices by health care professionals.
- It improves the security, integrity and the availability of data for its users.
- It also paves the way for better patient-provider collaboration.

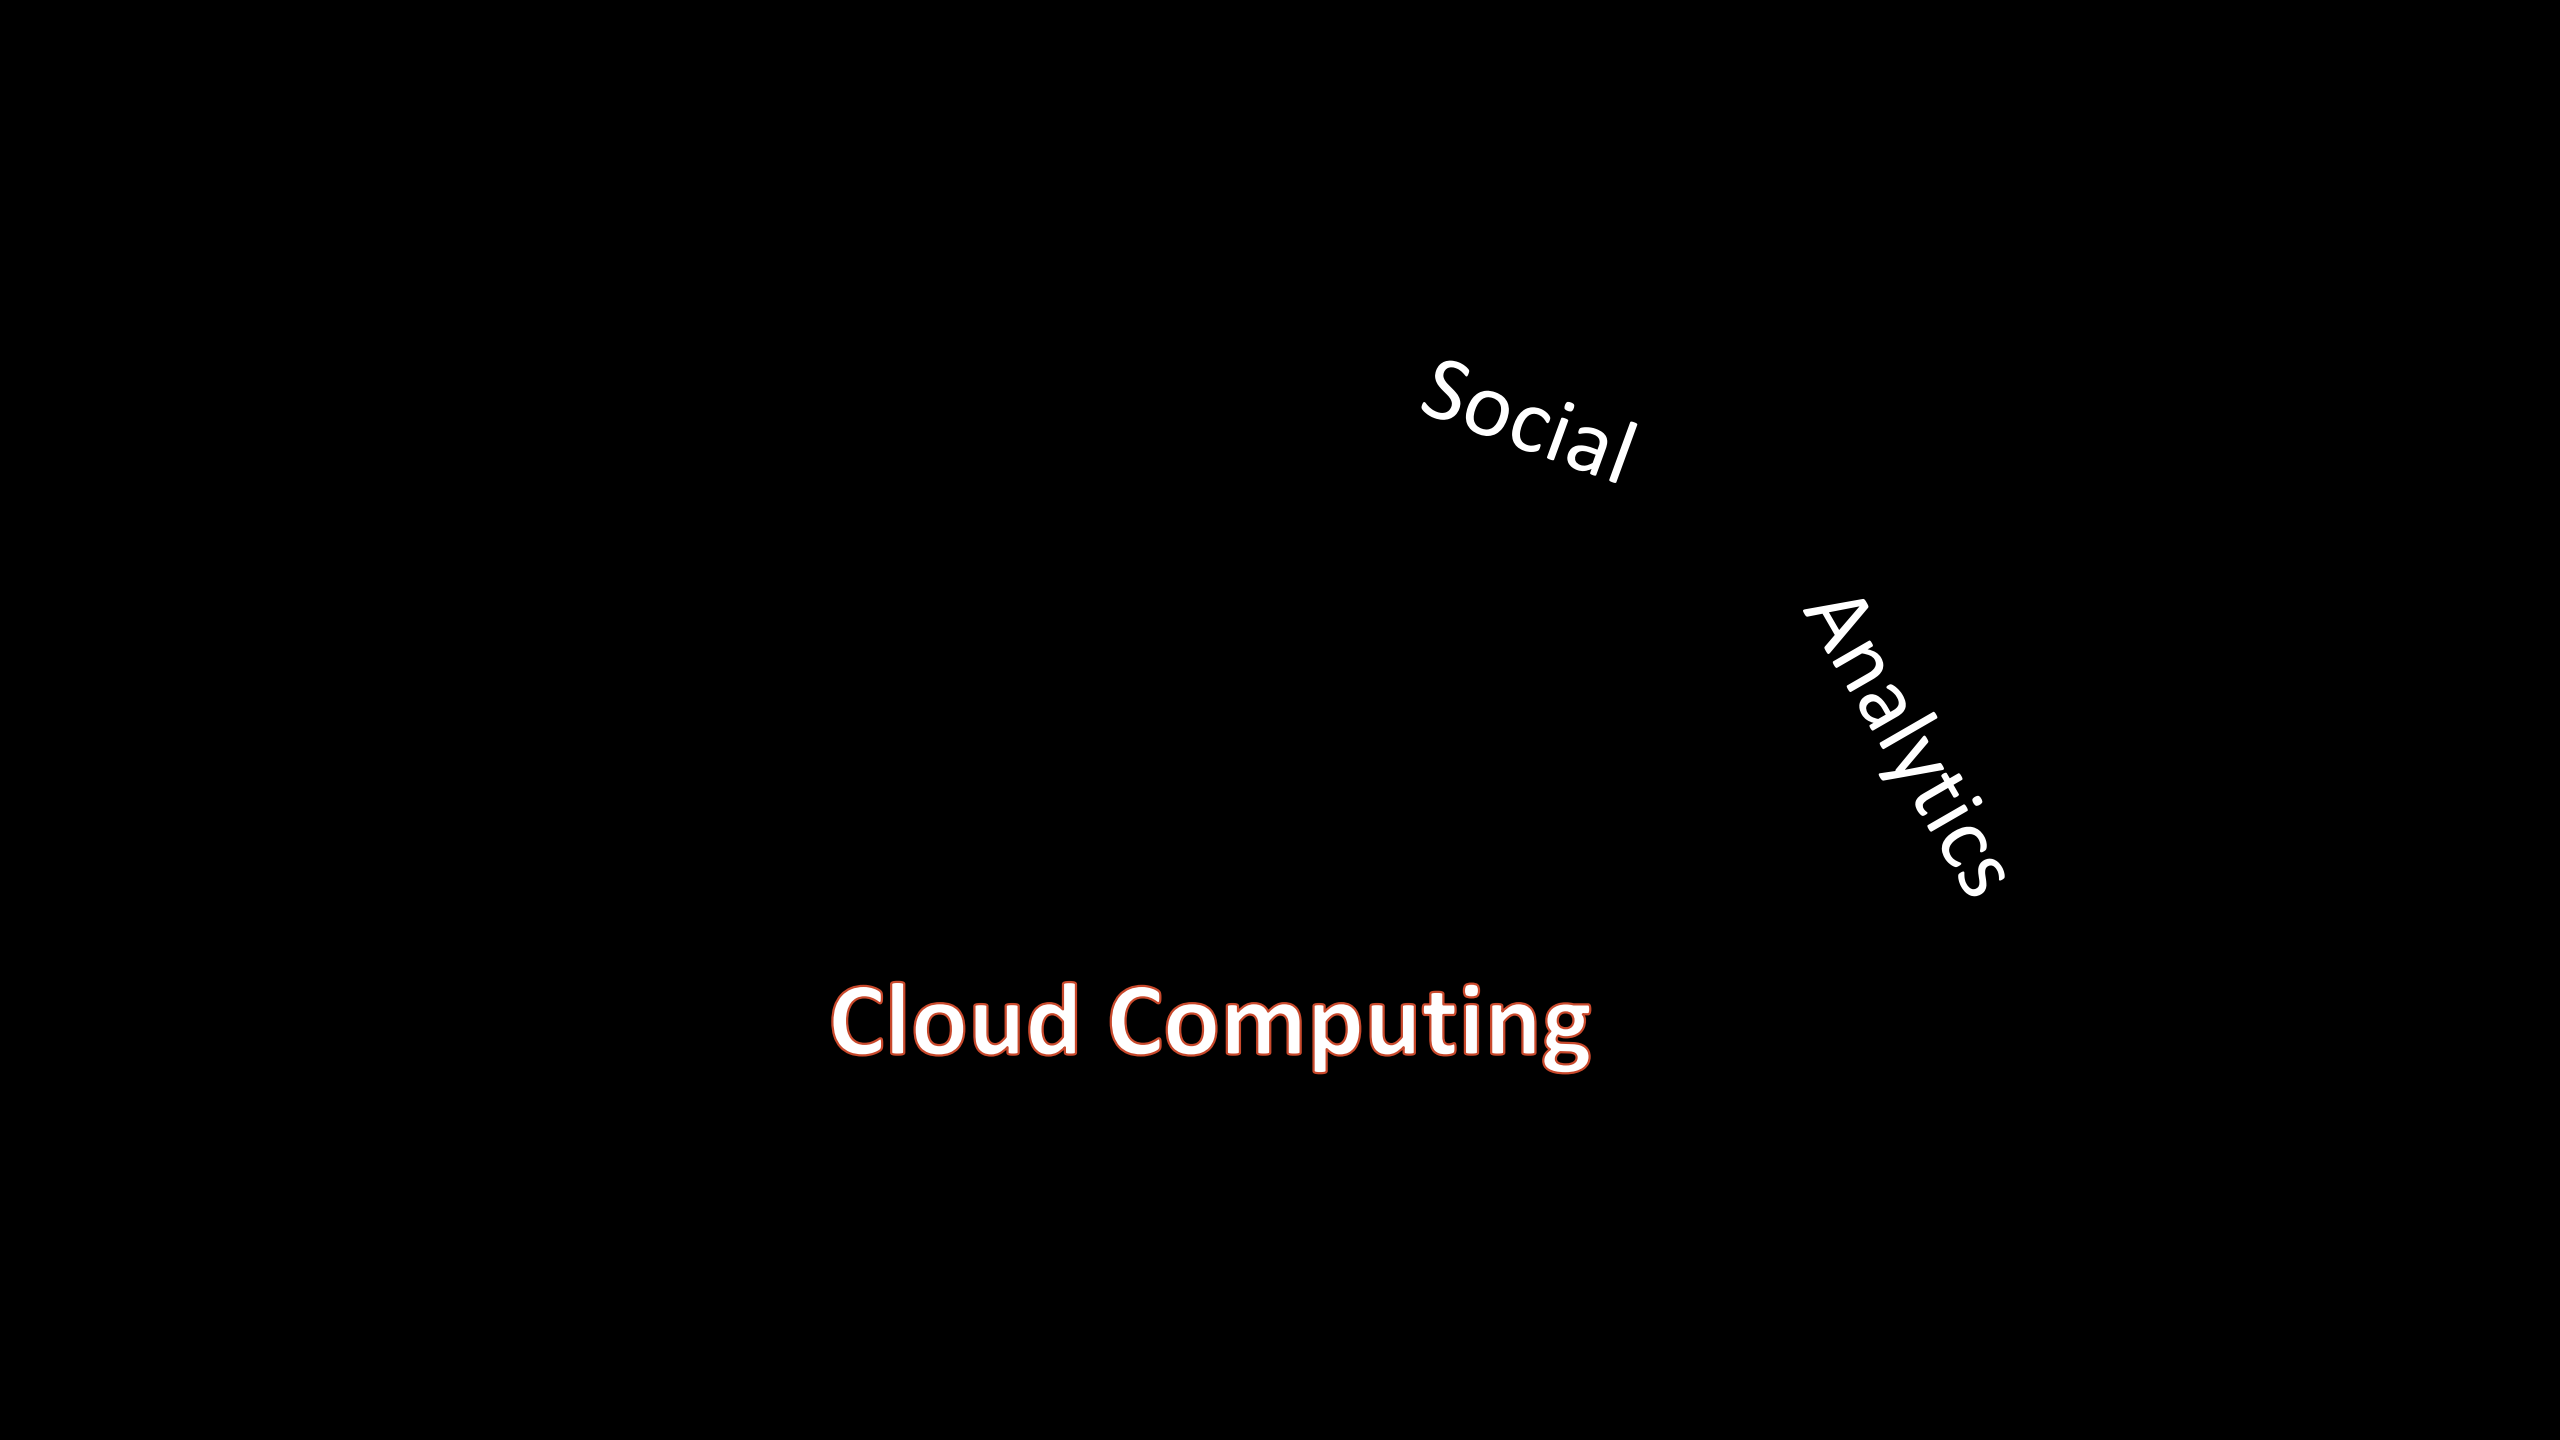
Social
Analytics
Cloud Computing

What does social have to with health care?

- Social isn’t just Facebook and Twitter.
- The social metaphor is being used by enterprises to improve collaboration between employees and also between the enterprise and its customers.
- Social data is playing an increasingly crucial role in the future of all industries, including the health care industry.

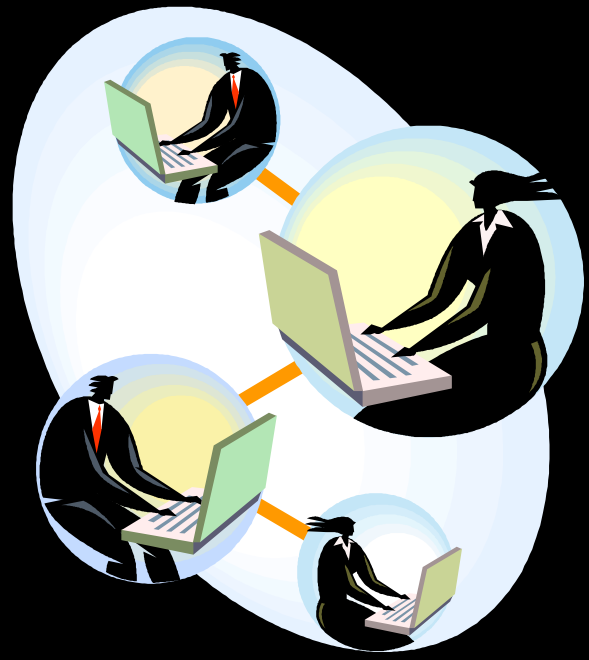

Social data usage in health care

- Social data mining can help predict illnesses before they occur
- Trends in a person's social behavior can give clues for physicians as to whether the patient is taking his/her medication or not
- Social data is being used to predict and track disease transmission and outbreak

Analytics

Cloud Computing

Analytics is the only way to deal with big data

- We have an enormous amount of data to deal with in all industries including the health care industry.
- Data comes from a variety of sources and in countless formats.

The world is getting smarter

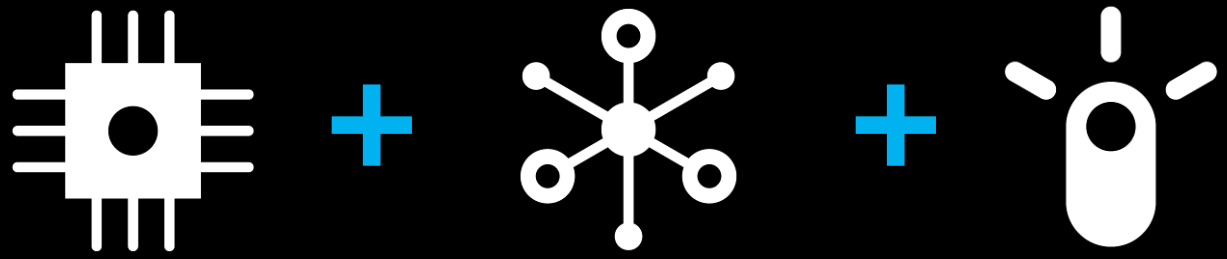

Instrumented

Over 10 billion CPUs were produced in 2008, up 1000% in 8 years.

Interconnected

2 billion people were on the Web by 2011 ... with a trillion connected objects.

Intelligent

Stockholm leverages GPS data to predict traffic – reducing congestions and emissions.

Healthcare is “dying of thirst in an ocean of data”

“Medicine has become too complex. Only about 20% of the knowledge clinicians use today is evidence-based.”

Steven Shapiro
Chief Medical & Scientific Officer
University Pittsburgh Medical Center

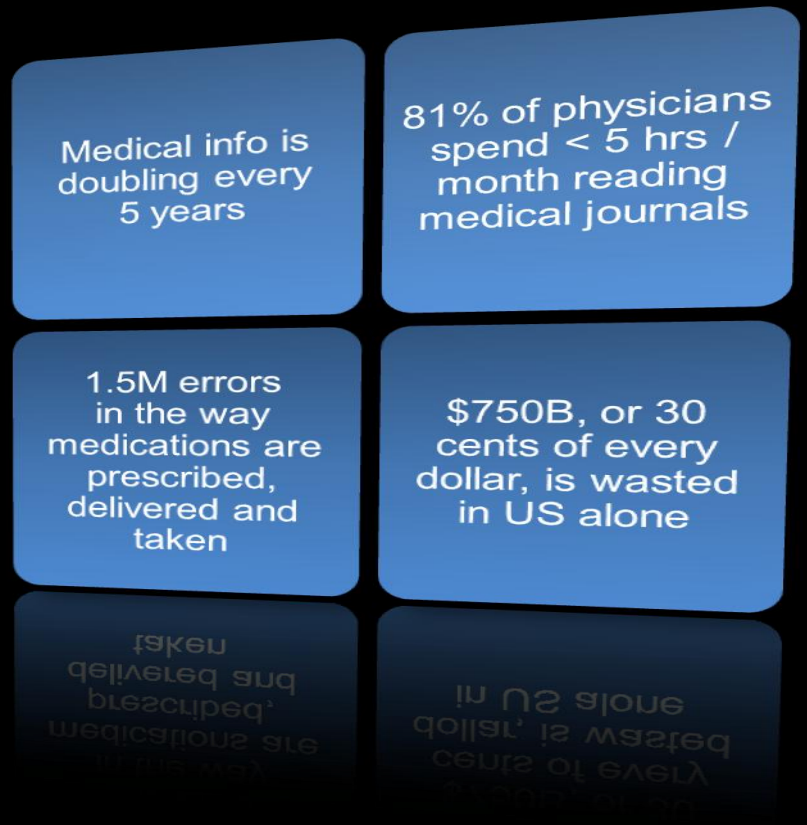

IBM Watson : Using analytics in health care

The 4 V's

Data that drives decisions is growing faster than traditional methods can use it

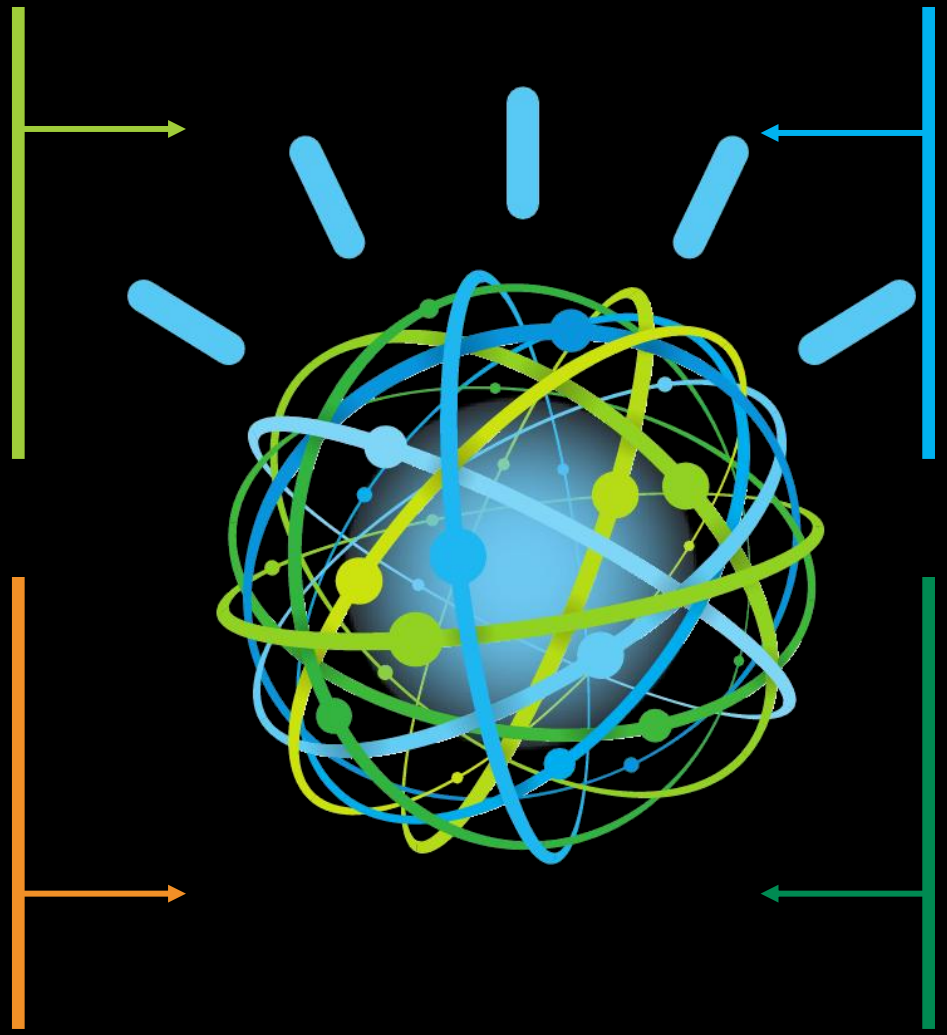

Analytic opportunity

Analytically advanced organizations are 260% more likely to be top performers

Cognitive systems

- Ingest big data as it is generated
- Co-evolve with field
- Keep users one step ahead in the footrace

Best-in class companies

- WellPoint: Largest US health insurance org
- Memorial Sloan Kettering: world's preeminent cancer research org
- Citigroup: world's largest financial services network

IBM Oncology Diagnosis and Treatment Advisor Demonstration

Shows how Watson can assist an Oncologist by:

- Synthesizing disparate data – patient records, clinician notes, test results, pathology reports, etc.
- Identifying missing pieces of data recommending tests with complete transparency
- Suggesting personalized, confidence-weighted, evidence-based options to improve quality of care and patient experience

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria

Symptoms

difficulty swallowing
fever
dry mouth
thirst
anorexia
frequent urination
dizziness

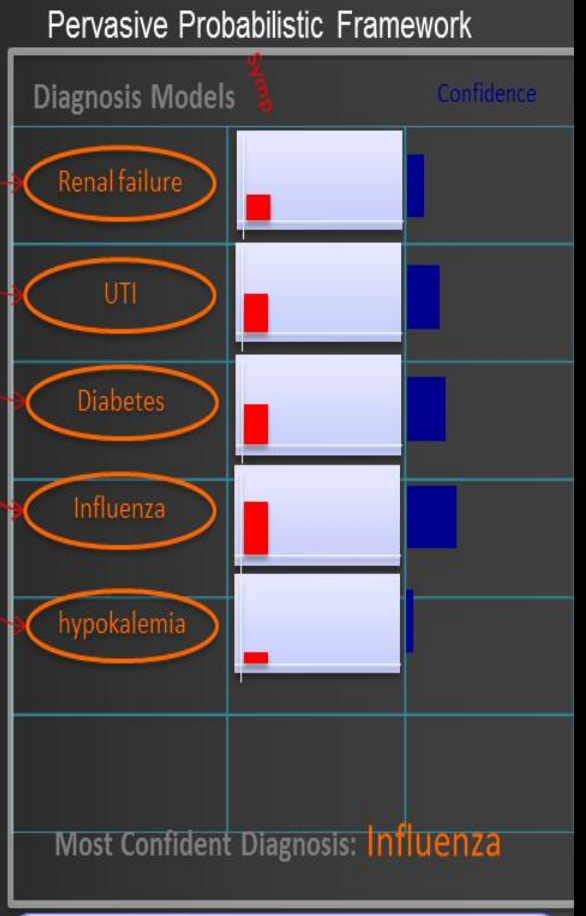

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria

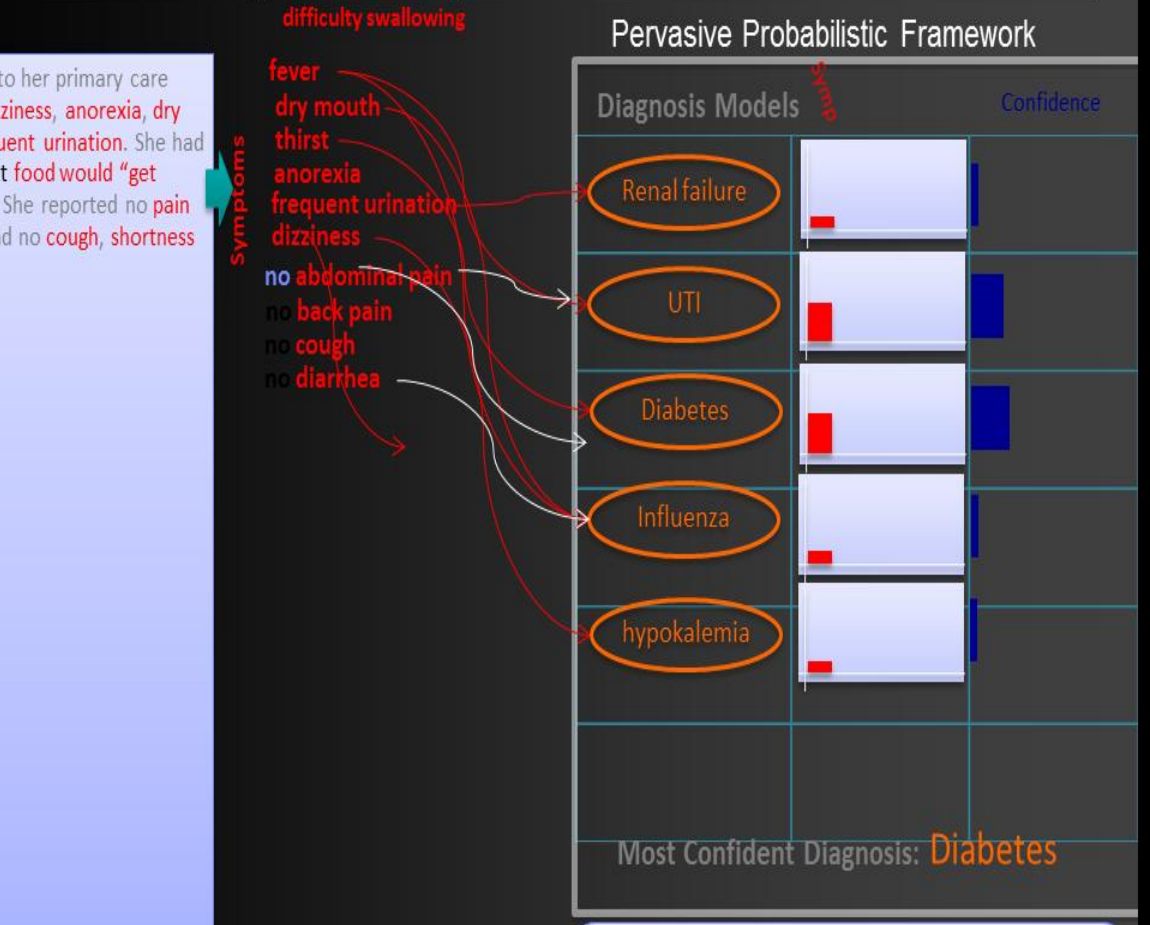

- Identify negative Symptoms

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria

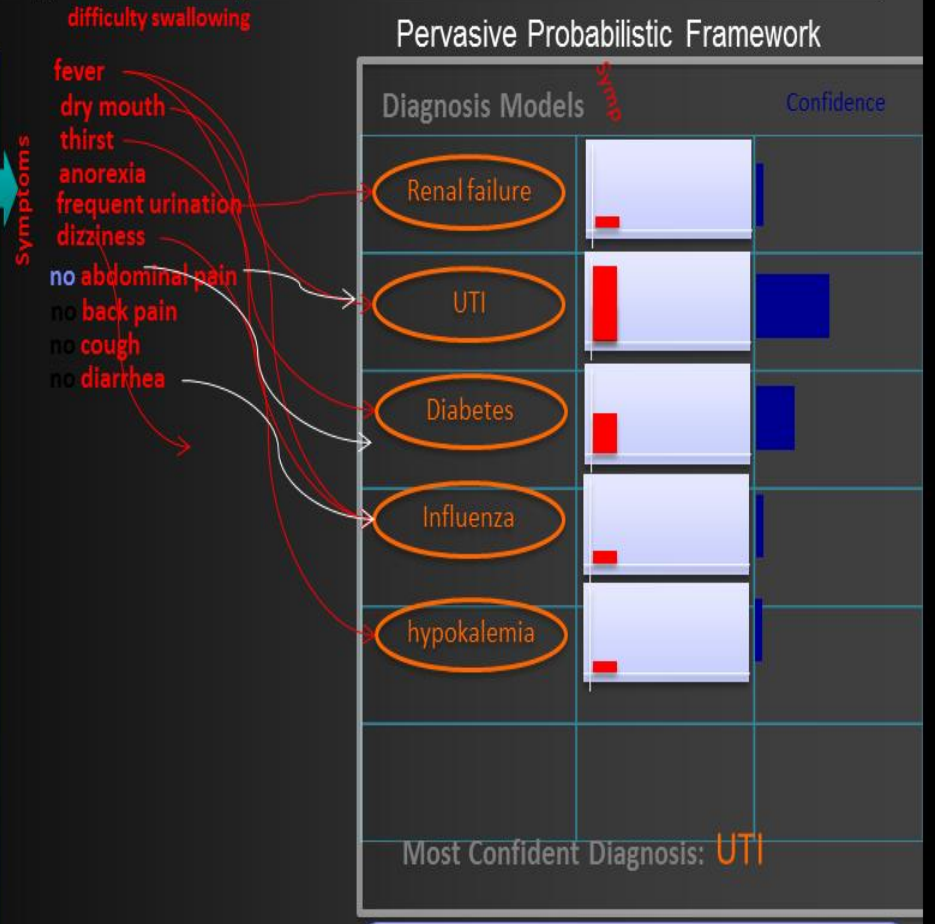

- Identify negative Symptoms
- Reason with mined relations to explain away symptoms (thirst is consistent w/ UTI)

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria. Her family history included oral and bladder cancer in her mother, Graves' disease in two sisters, hemochromatosis in one sister, and idiopathic thrombocytopenic purpura in one sister

Symptoms:
- difficulty swallowing
- fever
- dry mouth
- thirst
- anorexia
- frequent urination
- dizziness
- no abdominal pain
- no back pain
- no cough
- no diarrhea

Family History:
- Oral cancer
- Bladder cancer
- Hemochromatosis
- Purpura
- Graves' Disease (Thyroid Autoimmune)

- Extract Family History
- Use Medical Taxonomies to generalize medical conditions to the granularity used by the models

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria. Her family history included oral and bladder cancer in her mother, Graves' disease in two sisters, hemochromatosis in one sister, and idiopathic thrombocytopenic purpura in one sister. Her history was notable for cutaneous lupus, hyperlipidemia, osteoporosis, frequent urinary tract infections, three uncomplicated cesarean sections, a left oophorectomy for a benign cyst, and primary hypothyroidism, which had been diagnosed a year earlier

Symptoms

- difficulty swallowing
- fever
- dry mouth
- thirst
- anorexia
- frequent urination
- dizziness
- no abdominal pain
- no back pain
- no cough
- no diarrhea

Family History

- Oral cancer
- Bladder cancer
- Hemochromatosis
- Purpura
- Graves' Disease (Thyroid Autoimmune)

Patient History

- cutaneous lupus
- osteoporosis
- hyperlipidemia
- frequent UTI
- hypothyroidism

Pervasive Probabilistic Framework

•Extract Patient History

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria. Her family history included oral and bladder cancer in her mother, Graves' disease in two sisters, hemochromatosis in one sister, and idiopathic thrombocytopenic purpura in one sister. Her history was notable for cutaneous lupus, hyperlipidemia, osteoporosis, frequent urinary tract infections, three uncomplicated cesarean sections, a left oophorectomy for a benign cyst, and primary hypothyroidism, which had been diagnosed a year earlier. Her medications were levothyroxine, hydroxychloroquine, pravastatin, and alendronate.

Symptoms
- difficulty swallowing
- fever
- dry mouth
- thirst
- anorexia
- frequent urination
- dizziness
- no abdominal pain
- no back pain
- no cough
- no diarrhea

Family History
- Oral cancer
- Bladder cancer
- Hemochromatosis
- Purpura
- Graves' Disease (Thyroid Autoimmune)

Patient History
- cutaneous lupus
- osteoporosis
- hyperlipidemia
- frequent UTI
- hypothyroidism

Medications
- alendronate
- pravastatin
- levothyroxine
- hydroxychloroquine

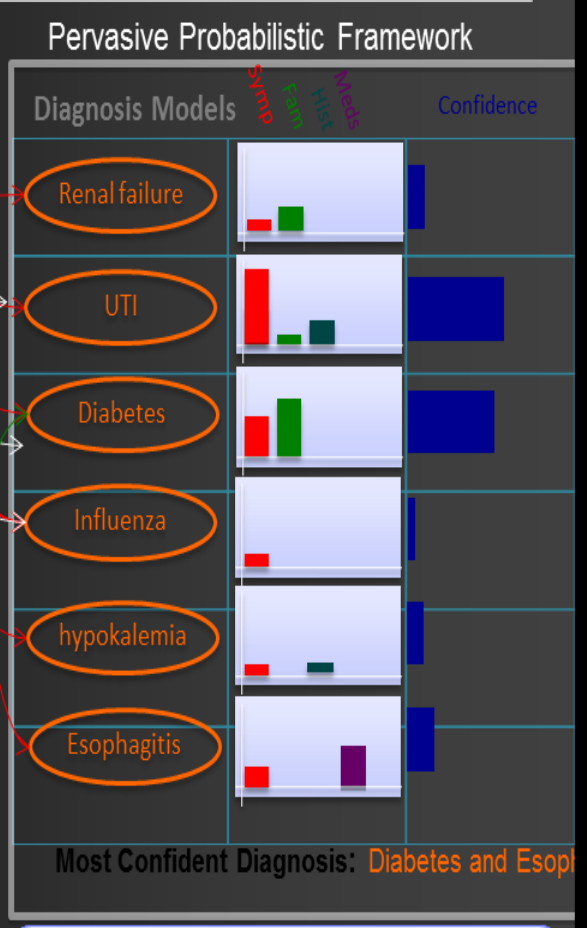

- Extract Patient History
- Extract Medications
- Use database of drug side-effects
- Together, multiple diagnoses may best explain symptoms

Applying Watson to the Real World

Continuous Evidence-Based Diagnostic Analysis

A 58-year-old woman presented to her primary care physician after several days of dizziness, anorexia, dry mouth, increased thirst, and frequent urination. She had also had a fever and reported that food would "get stuck" when she was swallowing. She reported no pain in her abdomen, back, or flank and no cough, shortness of breath, diarrhea, or dysuria. Her family history included oral and bladder cancer in her mother, Graves' disease in two sisters, hemochromatosis in one sister, and idiopathic thrombocytopenic purpura in one sister. Her history was notable for cutaneous lupus, hyperlipidemia, osteoporosis, frequent urinary tract infections, three uncomplicated cesarean sections, a left oophorectomy for a benign cyst, and primary hypothyroidism, which had been diagnosed a year earlier. Her medications were levothyroxine, hydroxychloroquine, pravastatin, and alendronate. A urine dipstick was positive for leukocyte esterase and nitrites. The patient was given a prescription for ciprofloxacin for a urinary tract infection and was advised to drink plenty of fluids. On a follow-up visit with her physician 3 days later, her fever had resolved, but she reported continued weakness and dizziness despite drinking a lot of fluids. She felt better when lying down. Her supine blood pressure was 120/80 mm Hg, and her pulse was 88 beats per minute; on standing, her systolic blood pressure was 84 mm Hg, and her pulse was 92 beats per minute. A urine specimen obtained at her initial presentation had been cultured and grew more than 100,000 colonies of *Escherichia coli*, which is sensitive to ciprofloxacin.

- Extract Patient History
- Extract Medications
- Use database of drug side-effects
- Together, multiple diagnoses may best explain symptoms
- Extract Findings: Confirms that UTI was present

Watson for Healthcare solutions

Watson for Healthcare

TEACH

Enable new methods for teaching & medical training

Ex. Watson Oncology Research Advisor

PRACTICE

Enable research and delivery of evidence based medicine

Ex. Watson Oncology Diagnosis and Treatment Advisor

PAY

Enable the rapid evaluation and pre-auth. of medical treatment

Ex. Watson Utilization Management Advisor

System Specifications

- 2880 Processing Cores
- 90 IBM P750 Servers
- 16 Terabytes Memory (RAM) – 20TB Disk
- 80 Teraflops (80 trillion operations per second)
- Workload Optimized Systems

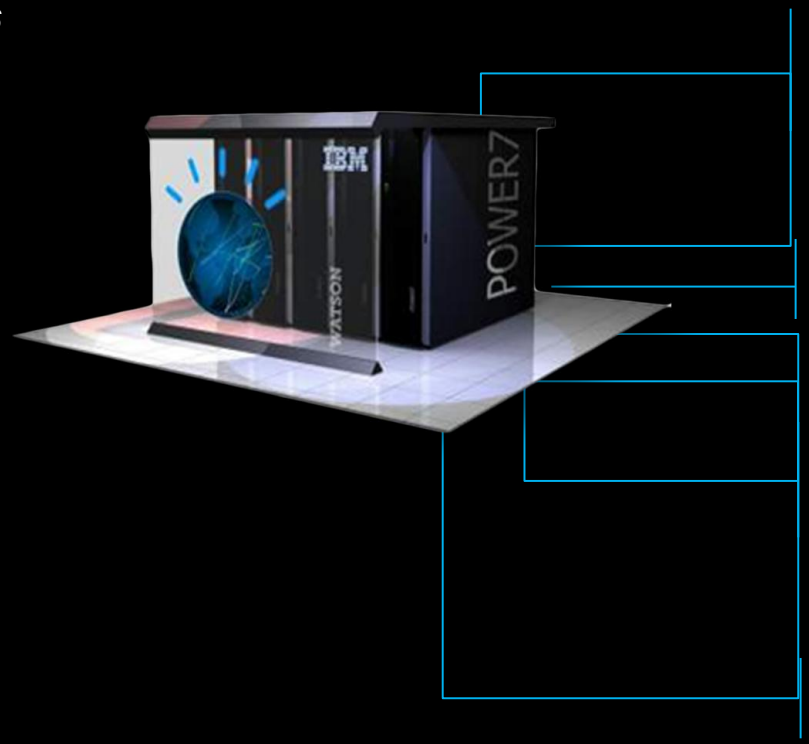

IBM Technology Depth

- Content Analytics
- Business Analytics
- Big Data
- Databases / Data Warehouses

In the past 5 years IBM has spent over $14B in analytical acquisitions and $6B in R&D annually

Watson is delivered as a service for greater access and simplicity

Dynamic Capacity
Automate and control service provisioning

Hybrid Delivery
Extend & integrate on-premise solution with cloud offering

Flexible Consumption
Support alternative delivery and pricing models

Time to Value
Enable incremental automation and business agility

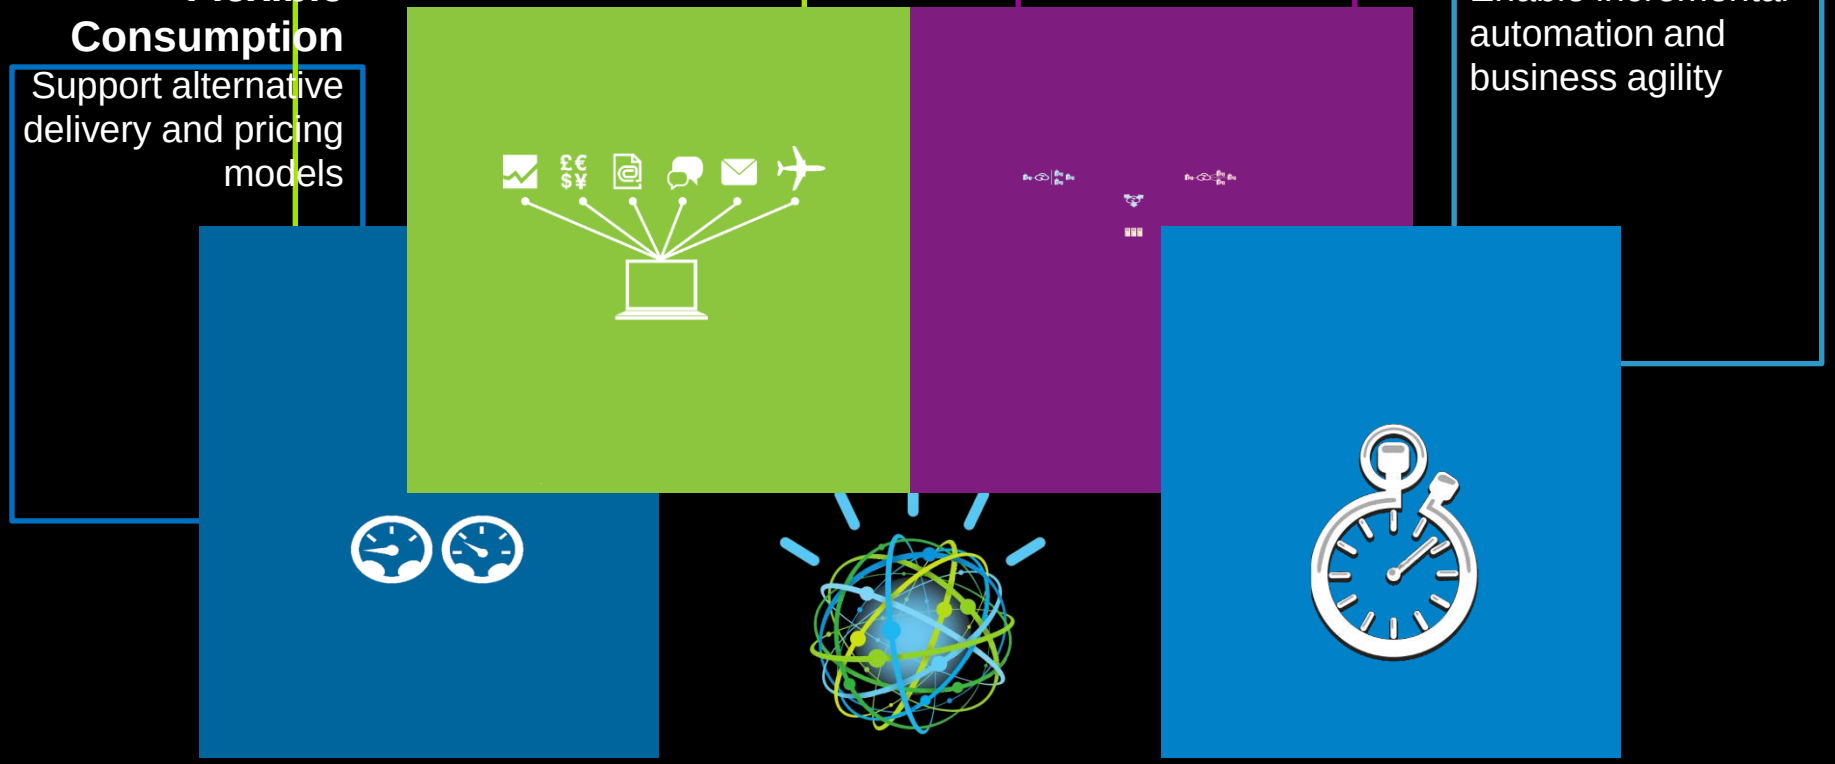

Open Stack

The Open Source Cloud Operating System

- The OpenStack cloud operating system enables enterprises and service providers to offer on-demand computing resources, by provisioning and managing large networks of virtual machines.
- Compute resources are accessible via APIs for developers building cloud applications and via web interfaces for administrators and users.
- The compute architecture is designed to scale horizontally on standard hardware, enabling the cloud economics companies have come to expect.

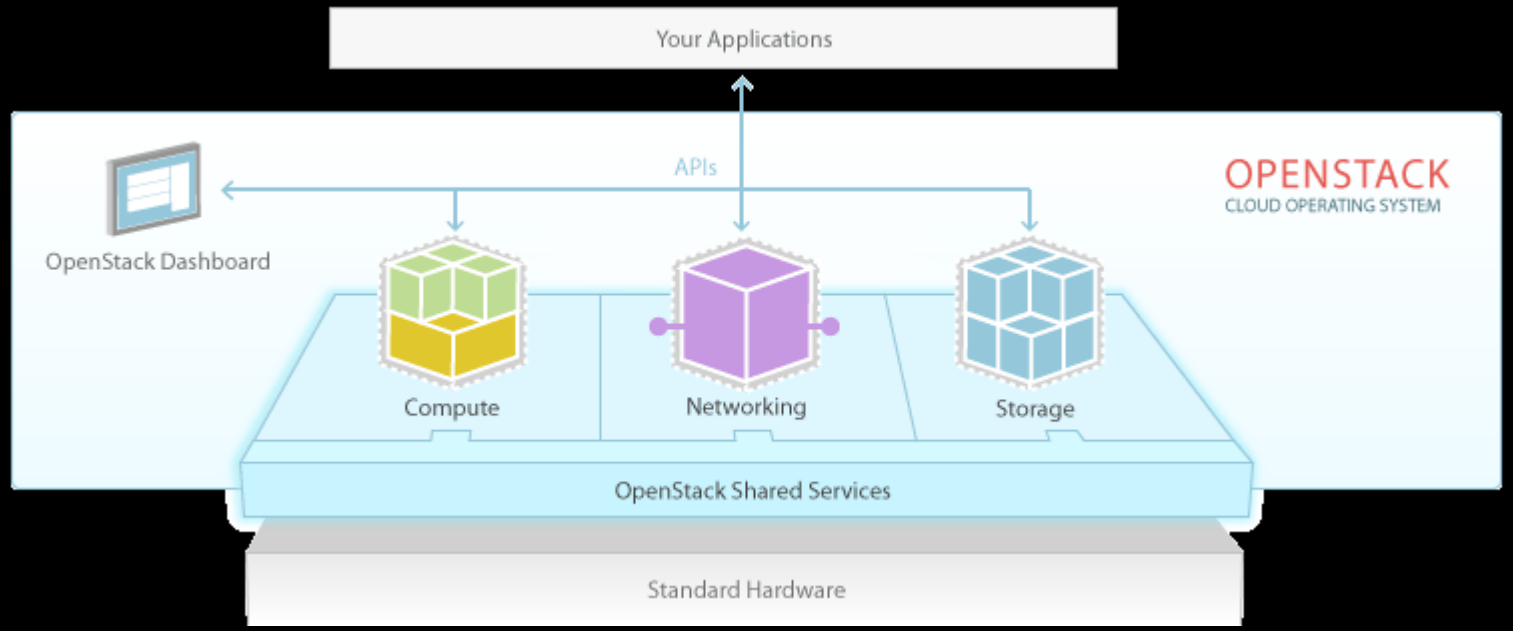

* Source : OpenStack.org

In a world that’s moving rapidly with a ton of medical data to manage, complex problems to resolve and unpredictable need for computing resources, cloud computing stands out as the only IT platform that can allow health care providers to concentrate on advancing their lead in the medical field while having their IT problems taken care of by the ones who know it best.

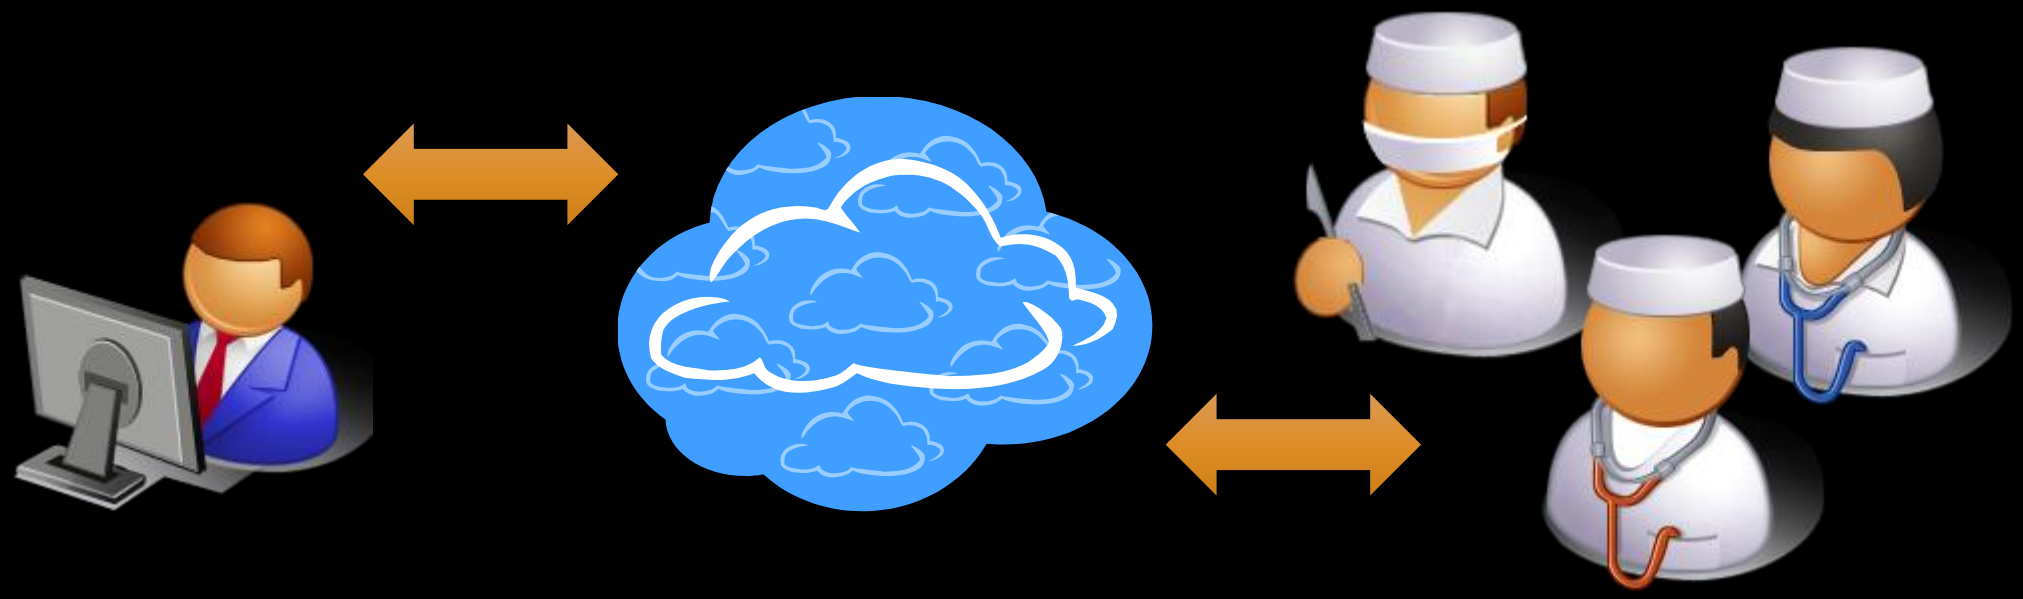

Cloud Computing for Financial Services

Madhura Pundlik

Boston University

May 19, 2013

Why consider the cloud?

Reduce cost

- No need to maintain physical hardware.
- Upfront elimination of fixed costs.
- Convert Capital Expenditure to Operational Expenditure.

Increase flexibility

- Flexibility in the quantum of service used.
- Usage driven to meet organization needs.

Financial Industry - Concerns

Compliance

Availability

Data Privacy

Security

No Standards /
Vendor locking

Current Trends in Financial Industry

- Private and hybrid cloud – a popular choice
- Large banks are using PaaS
- Data Analysis
- Leverage mobile media

Estimated spending on private cloud by financial services worldwide

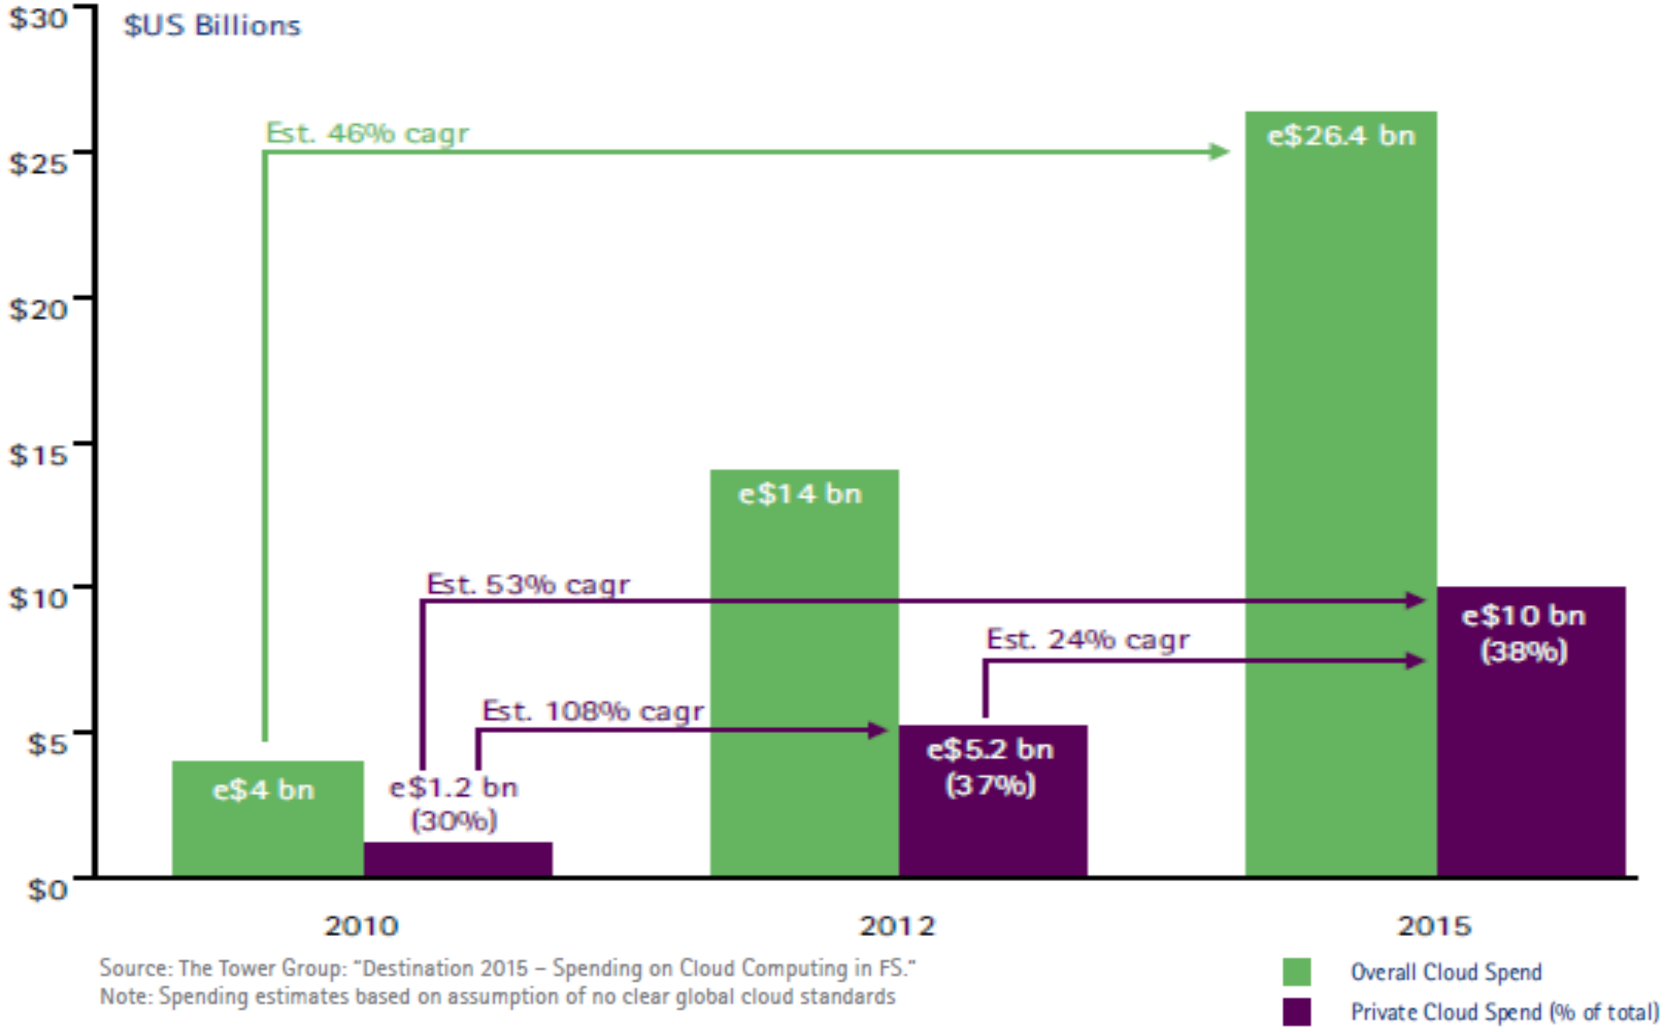

Source: The Tower Group: "Destination 2015 – Spending on Cloud Computing in FS."
Note: Spending estimates based on assumption of no clear global cloud standards

Higher acceptance levels

- Various Wall Street firms are adopting cloud computing.
- Within the next four years the market is projected to become a billion-dollar industry.
- Already have more than 600 million cloud subscribers.
- Debt or equity capital markets, mutual funds, private equity, asset management and banking solutions provided by leading financial institutions worldwide are adopting cloud computing products.

Typical illustrations of Cloud Computing

- Compute Power
  - improved use of infrastructure
- Virtual Desktops
  - getting new desktops up and running faster
- Disaster Recovery
  - duplicate location sitting idle
- Development and Testing
  - churning up the environment takes longer
- Capital expenditure to Operational expenditure
  - extensive procurement process

Typical illustrations of Cloud Computing

- Shared Resources
  - reduce physical IT infrastructure
- Separation of Duties
  - between development staff and third-party contractors
- Seasonal/Peak Bursting
  - able to handle peak workload
- Hybrid Cloud Solutions
- Economies of Scale/Cost Transparency
  - addressing challenge of diminishing customer base

Resultant gain from Cloud deployment

- Shorten Deployment Times
- Burst
- Lower Capex
- Pay-As-You-Go Cost Structure

Cloud – Business Model Innovation

- Main driver for cloud adoption was the creation of new revenue streams.
- Speed service delivery and increase IT efficiencies.
- It is not a one-size-fits-all solution.
- IBM survey addressed three aspects of cloud computing:
  - Sourcing choices
  - Workload considerations
  - Service delivery

IBM Workload Recommendations for public cloud

Barriers \ Push factors	Lower	Higher
Low	• Infrastructure capacity for training • Servers • Application servers • Application streaming	• Conferencing • VOIP infrastructure • WAN capacity • Storage • Service/help desk • Desktop • Data center network • Data mining
High	• CRM/Sales force • Unified communications • Test environment infrastructure • Industry-specific applications • Data warehouses • Development environment tools • Data archiving • Transactional databases • Security	• ERP applications • Email • Data backup • Continuity/DR

IBM Workload Recommendations for private cloud

Barriers \ Push factors	Lower	Higher
Low	• Application servers • WAN capacity • Infrastructure capacity for training • Service/help desk	• **Data mining** • Security • Transactional databases • Continuity/DR • Data archiving • **Data center network** • **Storage** • Unified communications
High	• CRM/Sales force • Servers • Email • Test environment infrastructure • Application streaming • Development environment tools • Conferencing • VOIP infrastructure • Desktop	• ERP applications • Industry-specific applications • Data warehouses • Data backup

Financial Sector and Hybrid Cloud Computing

- Increase business agility.
- Reduce IT services-delivery time.
- Ability to move workloads between clouds.

Indicators of agility

Revenue Growth	Cost Reduction	Risk and Reputation Responsiveness
• Grow existing/new revenue streams • Quickly adapt to market opportunities • Improve customer experience	• Identify operating improvements in business processes • Accelerate operational execution of projects • Rapidly scale costs with top-line demand	• Quickly adapt to new operating requirements • Mitigate impact of unexpected/unlikely events

Indicators of Agile Organizations, Based on McKinsey & Co.'s Business Agility Framework

Hybrid Cloud use cases for Retail Banking

Maintaining Dynamic Highly Available Front-End Customer Web site

- Iaas/Paas reduces infrastructure investment
- Ensures high availability at any demand level
- Improve customer experience
- Focus on providing innovative services

Large and increasing transaction data storage needs

- Reduces development and infrastructure investment of individual banks
- Allows them to cost-effectively store and provide customer transaction history

Increasing High Performance Computing (HPC) needs

- Faster provisioning of resources for risk analysis versus building in-house HPC
- Flexibility to schedule simulations for risk analysis based on business needs, versus waiting for internal HPC to be available

Hybrid Cloud use cases for Capital Markets

Demand spikes at market open and close	• Reduces required capacity of internal datacenters by sourcing capacity during peak load times from cloud. •On-demand scaling enables superior customer experience, thus differentiating the firm and increase customer loyalty while managing operational costs.
Algorithm testing increasing complexity and cost	•Quick response to disaster events by ramping up capacity and scale back when its no longer needed
Platform capability to provide Tools and Data	•Allows insurers to support a large number of affiliated agents and intermediaries •Reduces infrastructure investment needed to support tools for agents and intermediaries, as well as integrate data from back-end systems and pricing engines
Cost of maintaining low-latency trading	•Allows sell-side firms to reduce cost of trading infrastructure when low-latency trading is offered in the cloud.

Hybrid Cloud use cases for Insurance

Highly resilient and elastic infrastructure required for Front-End Insurance-Quote Web Site	• Allows insurers to reduce infrastructure investment • Realizes efficient operations, shorter time to market and cost reduction
Nimbleness – Ability to scale up and down on the fly as needed	• On-demand scaling enable flexible capacity • Quick response to disaster events by ramping up capacity and scale back when its no longer needed
Platform capability to provide Tools and Data	• Allows insurers to support a large number of affiliated agents and intermediaries • Reduces infrastructure investment needed to support tools for agents and intermediaries, as well as integrate data from back-end systems and pricing engines

Financial Institutions path toward cloud computing

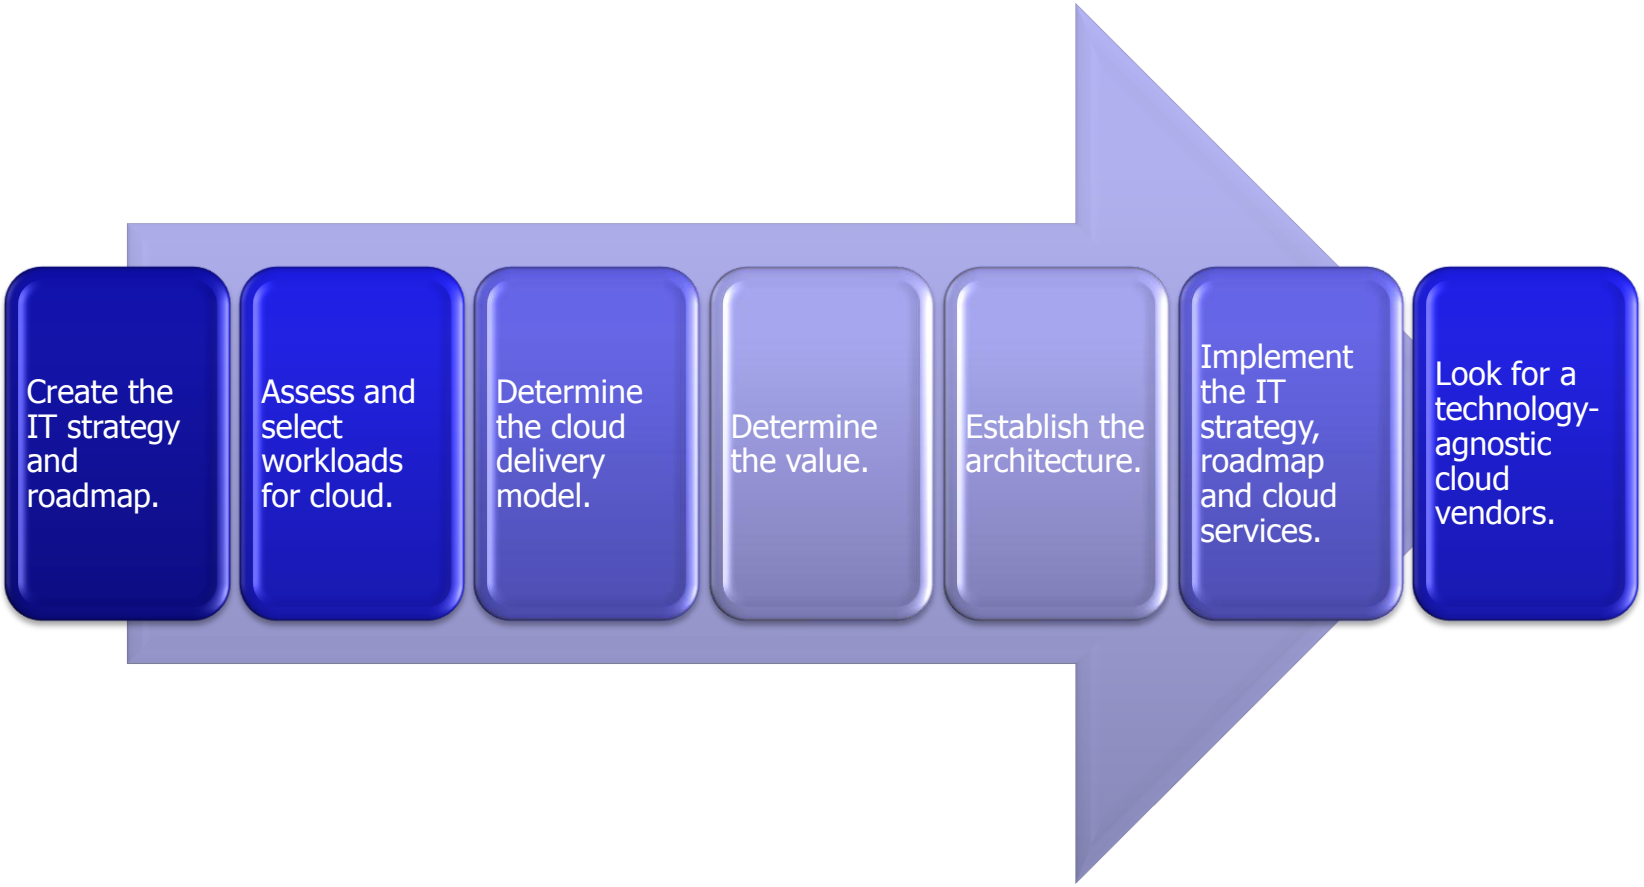

Key internal consideration for using TSPs*

- Develop strong understanding of business and legal risks
- Understand and consciously decide what you are buying
- Know that regulators will want to inspect the cloud
- Know and resolve issues at an early stage
- Have a strong TSP agreement

Cloud computing is still a journey

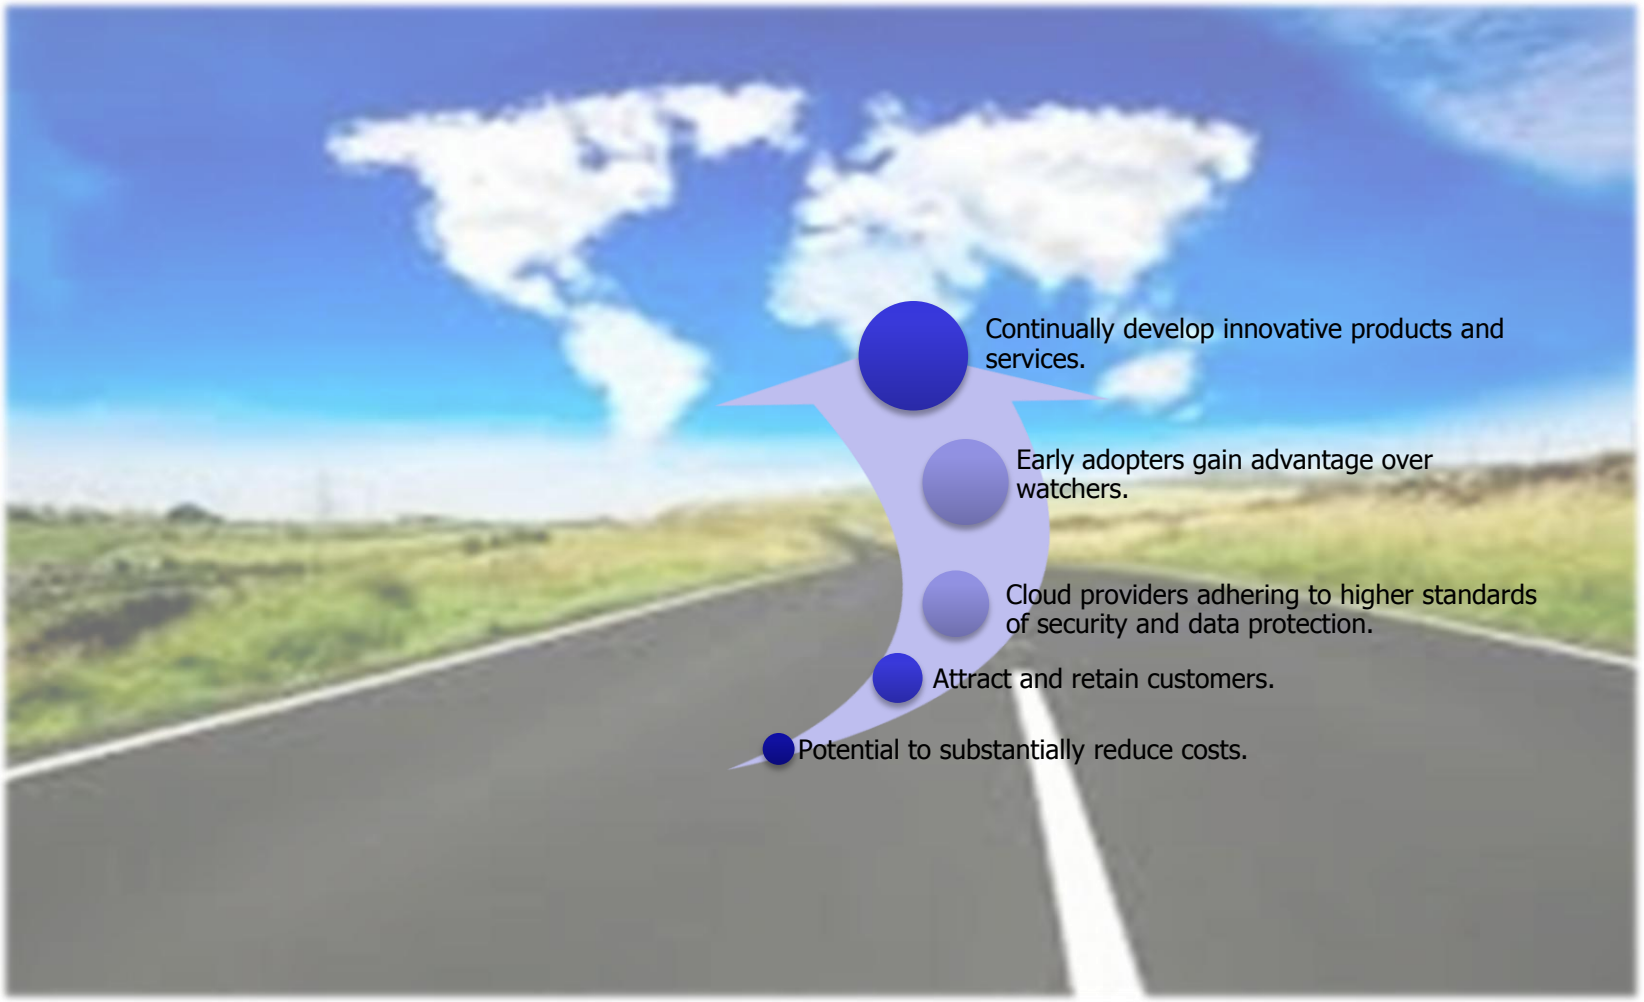

Thank you

References

- http://www.sapient.com/content/dam/sapient/sapientglobalmarkets/pdf/thought-leadership/GM_Cloud_Computing.pdf
- http://digitaljournal.com/article/342859
- http://www.savvis.com/en-us/info_center/documents/fin-casestudy-us-wallstreet.pdf
- https://www.siia.net/index.php?option=com_docman&task=doc_view&gid=2832&Itemid=48
- http://www-05.ibm.com/cz/businesstalks/pdf/WP_Dispelling_the_vapor_around_cloud_computing_in_the_financial_services.pdf
- https://www.vmware.com/files/pdf/VMware-Your-Cloud-in-Finance-Industry-Brief.pdf
- http://about.bloomberglaw.com/practitioner-contributions/are-financial-institutions-ready-for-cloud-computing/
- http://www.accenture.com/us-en/Pages/insight-new-era-banking-cloud-computing.aspx

Cloud Computing in Government
CCE2013

May 19, 2013

Dino Konstantopoulos

Government vision

- *"We are moving forward in the direction of cloud computing in the federal government, and part of that is to ensure we have to make sure we are leveraging those investments across the federal government."*
  - Vivek Kundra, Federal CIO

Cloud is a Bridge

- Cloud is a bridge
  - From the desktop to a world of devices
  - From 9-to-5 government offices to government services ***anywhere*** and ***anytime***

Cloud is an Economy of Scale

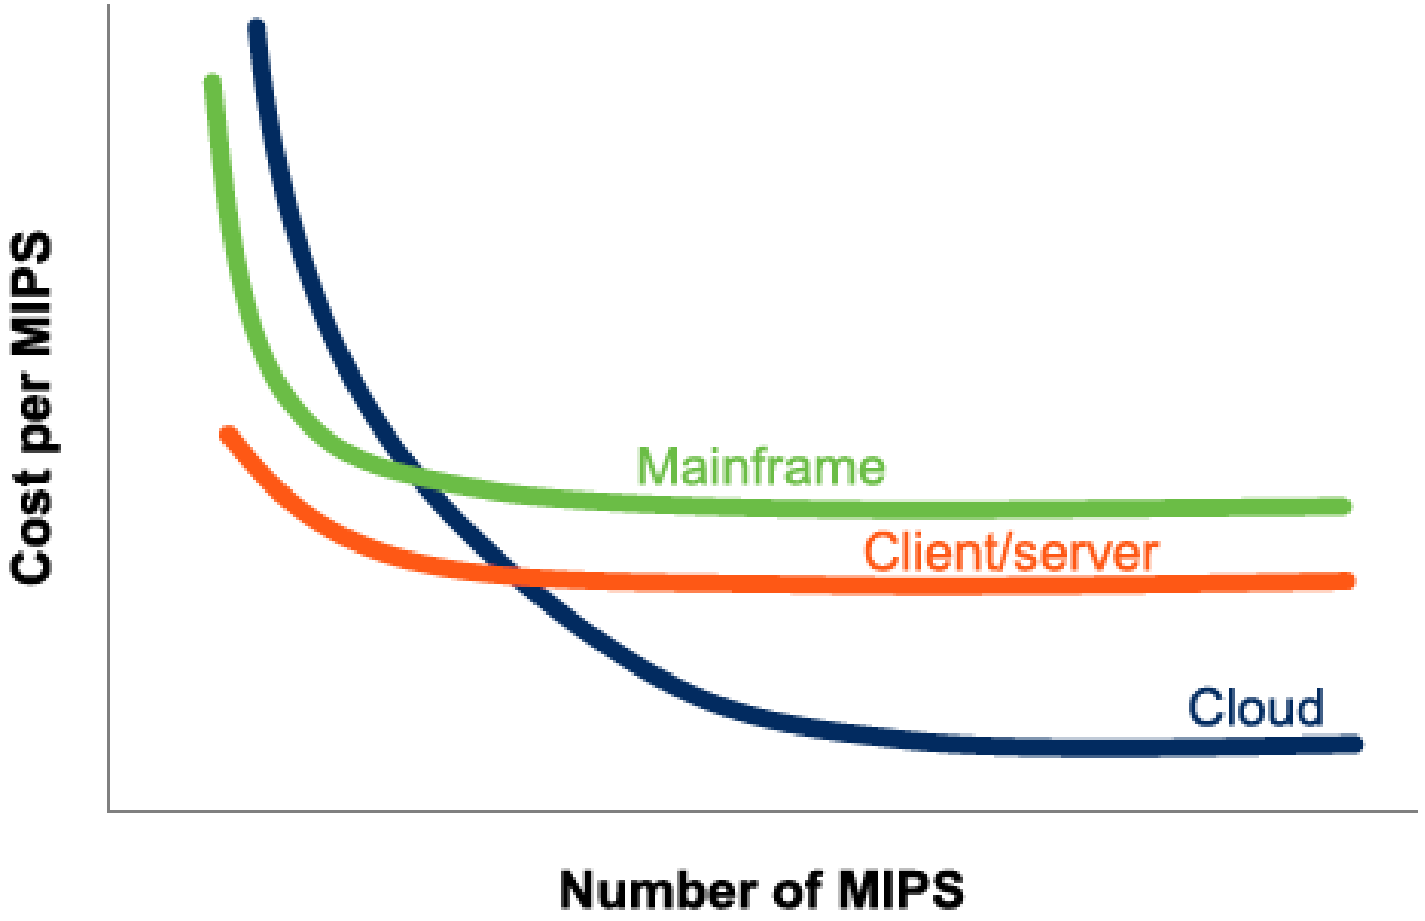

To Cloud or not to Cloud

- The choice to move to the cloud is not an all-or-nothing proposition
  - With different types of cloud offerings, you have flexible options about which services to obtain in the cloud and which to keep on site
  - Your priorities and security requirements determine the level of cloud capabilities to explore

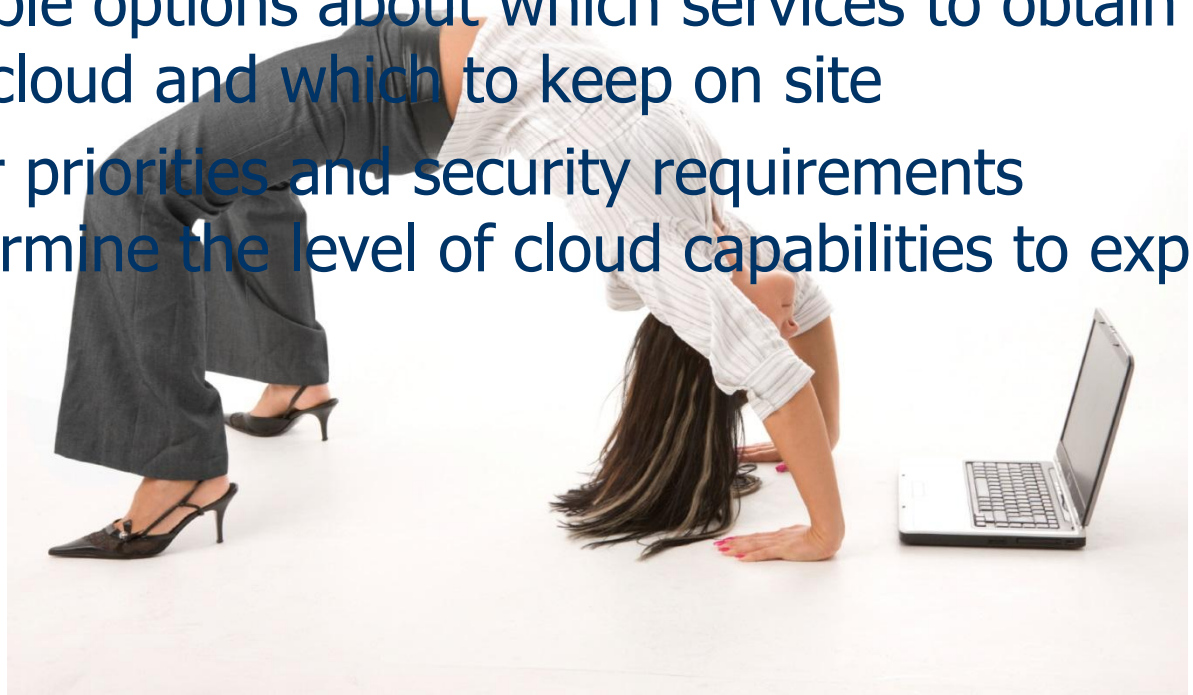

What is Cloud

- A network of ***elastic*** and ***pay-as-you-go*** computing resources that are located just about ***anywhere*** and that can be ***shared***
- Just as cable companies offer consumers a range of services (basic, premium, pay-per-view), cloud computing brings to government flexible choices about computing resources
- Integration of on-premise software assets with Cloud based services gives government ***new choices*** for improving citizen services while balancing System management overhead, cost, and security

Economies of scale and standardization

- Pay-as-you-go benefits are compelling:
  - The federal budget submitted to Congress in February 2010 commits to the use of cloud computing technologies and to a reduction in the number and cost of federal data centers
- By sharing IT services in the cloud, federal, state, and local governments can concentrate on mission-critical needs while achieving savings through economies of scale/standardization
  - e-mail in the cloud is affordable for budget-strapped local governments
  - People have been meeting in the cloud for at least 10 years using hosted conference services
  - NASA uses the expansive storage capacity of the cloud to host unique and large datasets, making them available to any citizen/researcher with a Web browser

Ever changing services

- Always growing number of services and alternative payment models that promise appealing cost savings, security, and flexibility
- Cloud options range from everyday SaaS office & productivity services, like e-mail, calendaring, and collaboration tools, to PaaS infrastructure services that free IT operations from mundane tasks and help reduce capital expenditures
  - System administrators can bring new services and computing capacity online quickly, managing costs as operational expenses. By allowing IT to respond quickly to changes, cloud computing helps administrators manage risks, peak demand, and long-term planning needs
- With cloud computing as part of your IT strategy, you can ***increase*** your data capacity without compromising security or requiring your agency to make heavy infrastructure investments—all while ***lowering*** your total cost of ownership
  - The trick is to find the right balance of on-premise and cloud services for your organization

SaaS, PaaS, IaaS

SaaS

- The cloud hosts the applications you use every day for office productivity, contact management, payment processing…
- SaaS makes use of a cloud computing infrastructure to deliver one application to many users, regardless of their location,
  - Rather than the traditional model of one application per desktop

SaaS

- Software and services on demand debuted in fourth place in the NASCIO 2010 list of top 10 policy and technology priorities for state CIOs
- In the current and future economy, it makes sense
  - It can lower expenses associated with software acquisitions in the near term
  - Longer term, it helps organizations with limited IT resources to deploy and maintain needed software in a timely manner—while, at the same time, reducing energy consumption and expense
  - Klamath County, Oregon, uses an e-mail solution hosted in the cloud. Without the cost savings of the subscription service, the organization could not have afforded to upgrade its straining legacy system—or to improve system performance

PaaS

- Cloud platforms free you to focus on the services you can offer without worrying about or managing the infrastructure needed for those services
  - Operating environment of the cloud with the tools you need on demand to create and host online services, software, Web sites, and mobile applications
  - Concentrate on delivering applications rather than on the underlying infrastructure, which a service provider maintains and updates in its data centers

PaaS

- Use PaaS to create *multi-tenant* applications
  - Services accessed by many users simultaneously
  - s public
    d

IaaS

- IaaS are on-demand data centers
  - Provide compute power, memory, storage, and other services
  - Typically priced per hour, and based on resource consumption
  - You pay only for what you use, and the service provides all the capacity you need
  - You're responsible for monitoring, managing, and patching your on-demand infrastructure

IaaS

- General Services Administ
  hosts a Web site in the clo
  - USA.gov cuts infrastructure cost
  - Challenge: Provide information to public on various government services in a cost-effective way
  - GSA expects to save 90% in operating costs with USA.gov with on-demand scaling

IaaS

- DoD puts compute-intensive applications on the Cloud with RACE
  - Self-service Web portal, with a shopping cart type of experience so DOD members can quickly acquire computing capacity
    - "When you push the button on your shopping cart, within a matter of minutes, whatever you requested gets provisioned, and then you get an e-mail back" confirming the request and providing a host name, password and user ID
  - Original version was for test and development environments. The new release allows DOD users to provision their own production environments
  - RACE is a private cloud that offers similar benefits as a public cloud
    - Sharing resources, and servers are not sitting under tables in your office.
  - During a humanitarian relief effort, a DOD agency could set up a temporary computer collaboration platform with noncoalition partners
    - The computer resources could be acquired quickly and then taken off-line just as quickly when the relief effort concludes
  - A DoD-private Amazon Cloud

Department of State

- Challenge: Unanticipated need for platform to host program ar financial management applicati during disarmament activity
- Solution: Leveraged, highly accessible, secure and cost efficient Platform-as-a-Service without additional IT infrastruct investment

Cloud services provide solutions for temporary need basis without infrastructure investment

Centers for Disease Control

- Challenge: Need effective communication mechanism to inform people of the dangers o H1N1 flu and to contain its spr
- Solution: Free social networkin Software-as-a-Service tool (tw to send and receive updates fr its followers

Cloud services offer tools to gather and diss
public in case of emergencies

ARMY CTSF

- CTSF is the Army's strategic facility responsible for interoperability engineering, executing Army Interoperability Certification (AIC) testing, maintaining configuration control for all operational- through tactical-level information technology/national security systems, and supporting deployed warfighters' digital needs.
- Located at Fort Hood, TX

CTSF Cost comparison

- 72 physical servers @ 20%load
  - Single socket dual-core
- 2 x 500 GB disk drives/server
- 36, 000 GB disk space required
- 8 physical network switches
- 52,984 VA power required
- 180,789 BTU heat output
- 15,07 Tons AC required
- $240,431 hardware acquisition cost
- $41,772 annual power cost
- $23,387 annual AC cost
- $65,159 aggregate operating cost/year

- 1 high-end server
  - 8 socket, quad-core server
- 3 virtual server Array shelves
- 45,000 GB disk space provided
- 1 physical network switch
- 2,912 VA power required
- 9,936 BTU heat output
- 0.83 Tons AC required
- $222,972 hardware acquisition cost
- $2,295 annual power cost
- $3,897 annual AC cost
- $6,193 aggregate operating cost/year

Cost comparison

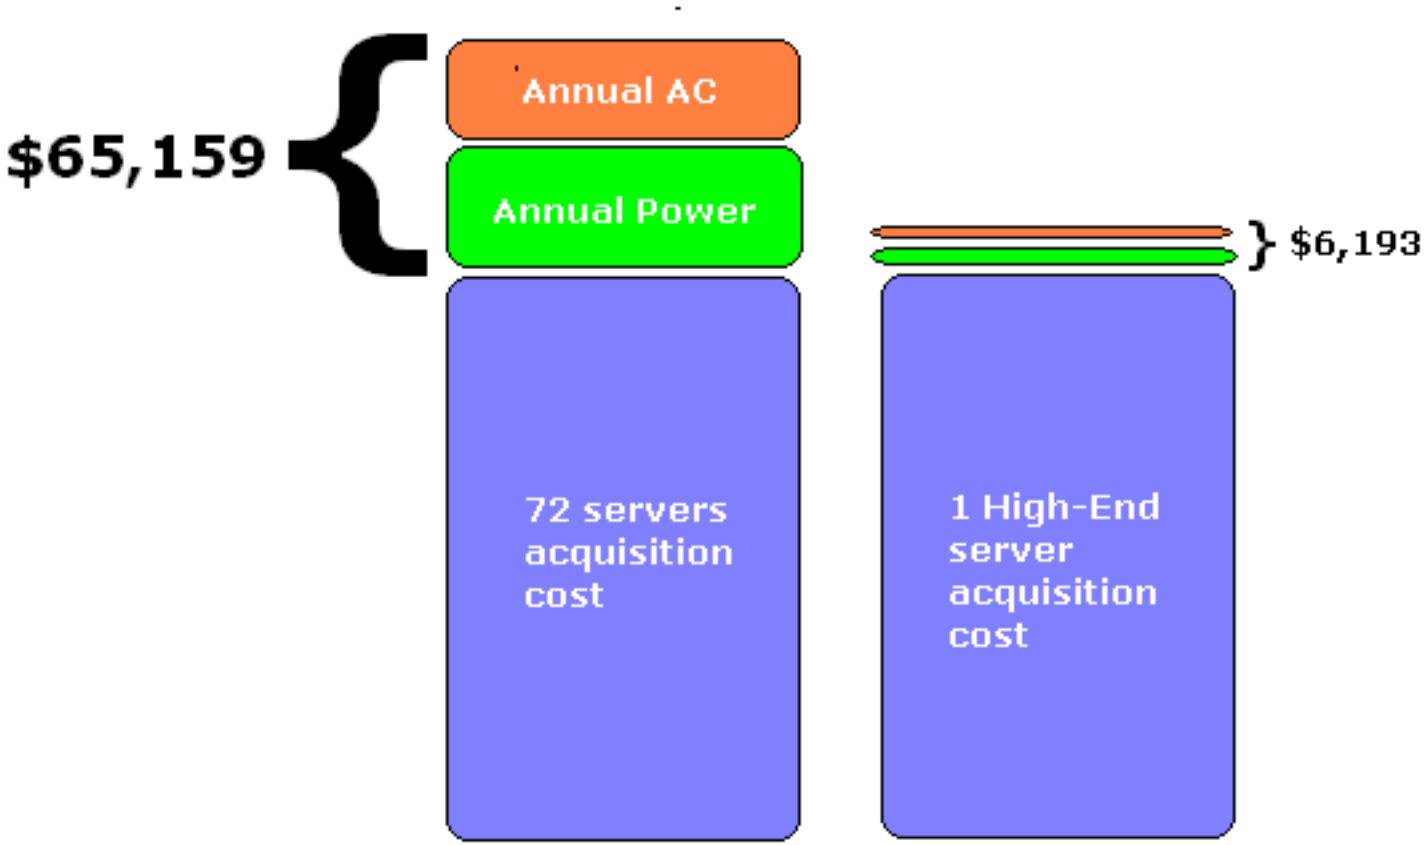

Obama 2012

4Gb/s, 10k requests per second, 2,000 nodes, 3 datacenters, 180TB and 8.5 billion requests. Design, deploy, dismantle in 583 days to elect the President. #madops

- Narwhal
  - Set of services that acted as an interface to a single shared data store for all of the campaign's applications
- Analytics
  - analytics system on EC2 Compute Cluster Eight Extra Large instances, which Amazon targets for high-performance computing jobs

Security

- Many government agencies are entrusted with confidential information and private data
- Cloud computing may seem risky because you cannot secure its perimeter—where are a cloud's boundaries?
- In addition, government agencies must comply with regulatory statutes, such as HIPAA, SOX, and FISMA

Security

- NASA found that cloud computing was as secure as other forms of data sharing and storage currently in use by the federal government
  - By consolidating information and computers onto one platform,
    - Cloud decreases the surface area vulnerable to a cyber attack
    - Reduces the number of systems that would otherwise need to be maintained and monitored independently, making it easier to keep all the computers on a network up-to-date

Security

- Government oversight organizations must set the framework of standards that help agencies weigh the risks and benefits of public cloud scenarios
  - Federal, state, and local agencies vary in their security and regulatory compliance needs, and you know your needs best
  - Agencies with sensitive information and workloads would probably never want all of their data in a public cloud

Private Clouds

- Private clouds offer the scalability and shared resources of cloud computing on your terms—and on your turf, if you can afford it
- To achieve true cloud scalability in a private cloud, you must forecast demand to support the requisite degree of excess capacity and invest accordingly

When to go Private

- A regulatory or security issue prevents you from hosting even encrypted data in a public cloud
- An application requires greater reliability or speed than the Internet
- You want control over your assets, including physical possession of the hardware your data resides on

Ready for Cloud Checklist

1. Justify cloud services
   - Discover how much cloud computing is already taking place in your organization or other agencies
   - Prototype & test
2. Budget for the cloud
   - Predictable budget and plan for IT resources
   - Find ways to share networks, PCs, and services
3. Integrate cloud services
   - Integrate on-premise applications and databases with cloud technologies to offer more or faster services
   - Think big—especially if you're a small agency. Cloud services are massively scalable. Who else might benefit?

Challenges & Obstacles

- **Security & Privacy**
  - Public sector CIOs must be able to report to legislators on how data is being kept private and secure
- **Business continuity**
  - Businesses come and go, and enterprises should ask hard questions about the portability of their data to avoid lock-in or potential loss if the business fails

Challenges & Obstacles

- **Regulatory compliance**
  - Enterprises are accountable for their own data even when it's in a public cloud, and should ensure their providers are ready and willing to undergo audits
- **Data provenance**
  - When selecting a provider, ask where their datacenters are located and if they can commit to specific privacy requirements

Regulatory Compliance

- Sarbanes-Oxley (SOX) Act which affects only public companies
- Payment Card Industry Security Standard (PCI-DSS), which affects any company accepting payment cards
- Federal Health Insurance Portability and Accountability Act (HIPAA) which affects any businesses with even the remotest possibility of touching patient data
- In Europe there is the EU Data Privacy Directive and Canada has an equivalent Personal Information Protection and electronic Documents Act (PIPEDA)

Why Public Cloud in government

- **Supply-side savings**
  - Large-scale data centers lower costs per server
- **Demand-side aggregation**
  - Aggregating demand for computing smoothes overall variability, allowing server utilization rates to increase
- **Multi-tenancy efficiency**
  - When changing to a multitenant application model, increasing the number of users lowers the application management and server cost per user

Conclusion

- Cloud will enable governments to better deliver on some of its key priorities:
  - **Fiscal responsibility**
  - **Better serve its citizens**
  - **Lower emissions**

Public and Private Trends

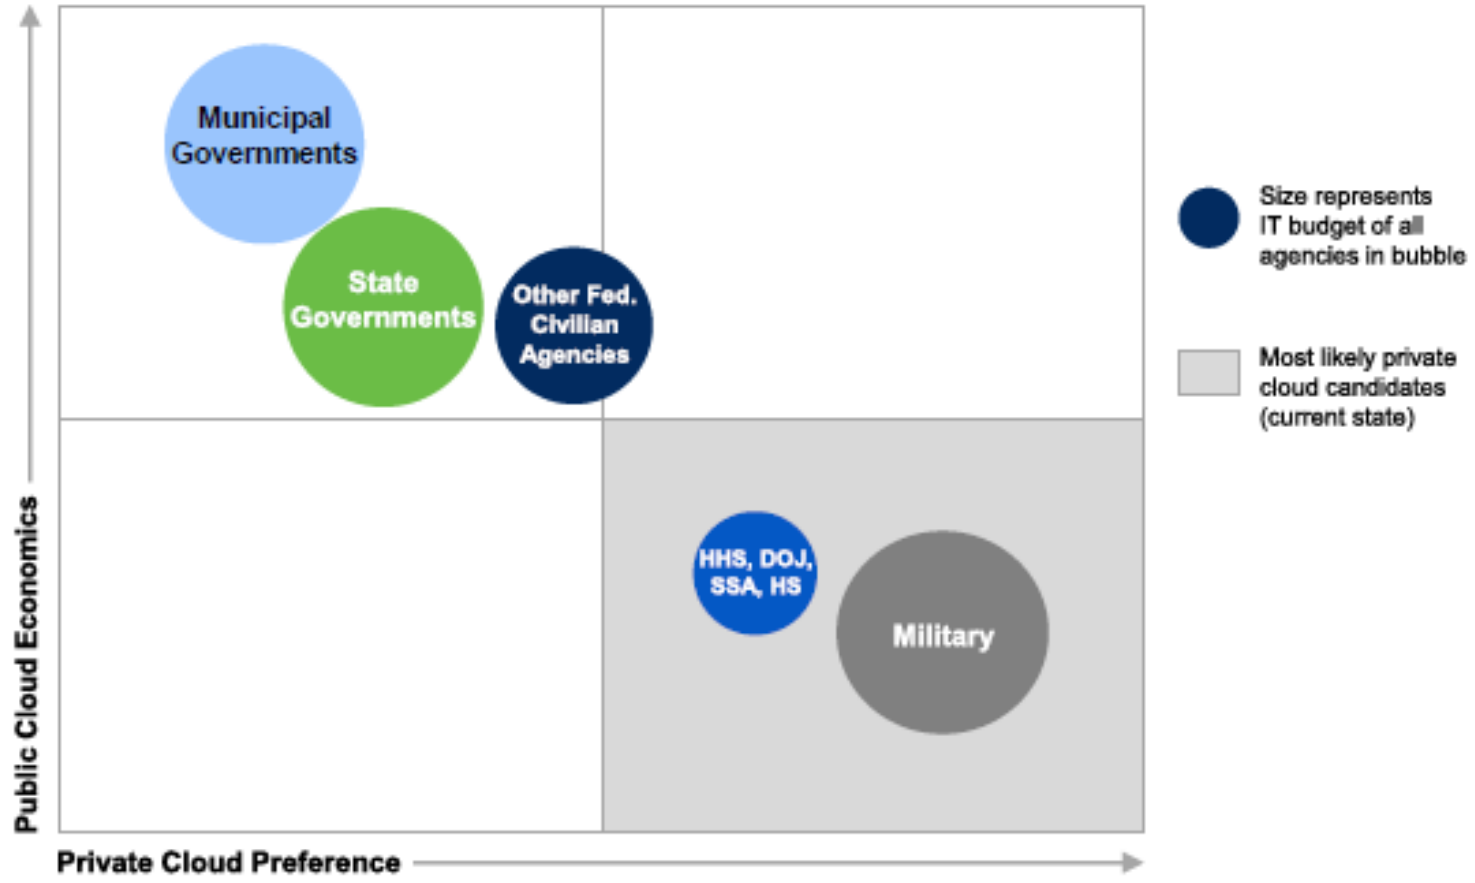

Source: Microsoft.

NIST

- Documents publically available through the NIST Cloud Computing Web site
  - http://www.nist.gov/itl/cloud/index.cfm