Hosted Research Clouds: The Time Is Now

Dan Reed
reed@microsoft.com
Multicore and Scalable Computing Strategist
Managing Director, Cloud Computing Futures

Microsoft

Waves of Innovation

Manycore Processors
Web Services / Mashups
Wi-Fi / Broadband

Fully Productive Computing

XML / SOAP
HTTP / HTML
SMTP

DEVELOPMENT

Hard Engineering → Intellectual Property

Software + Services

Email Clients
Web Browsers

Mouse
GUI
LANs

PC Architecture
MS-DOS

Spreadsheets
Word Processors

ADOPTION

Intellectual Property → Consumer Benefit

PC	APPLICATIONS	INTERNET	CLIENT+CLOUD
Mid ’80s	Late ’80s – Mid ’90s	Mid ’90s	Mid ’00s - future

Science 2020

"In the last two decades advances in computing technology, from processing speed to network capacity and the Internet, have revolutionized the way scientists work.

From sequencing genomes to monitoring the Earth's climate, many recent scientific advances would not have been possible without a parallel increase in computing power - and with revolutionary technologies such as the quantum computer edging towards reality, *what will the relationship between computing and science bring us over the next 15 years?"*

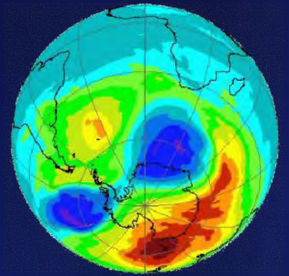

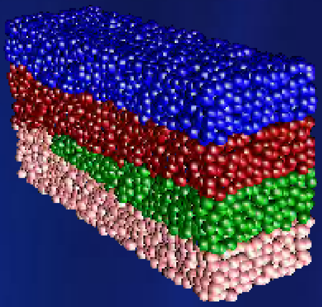

http://research.microsoft.com/towards2020science

Microsoft

Insights: Not Just FLOPS Or Bytes

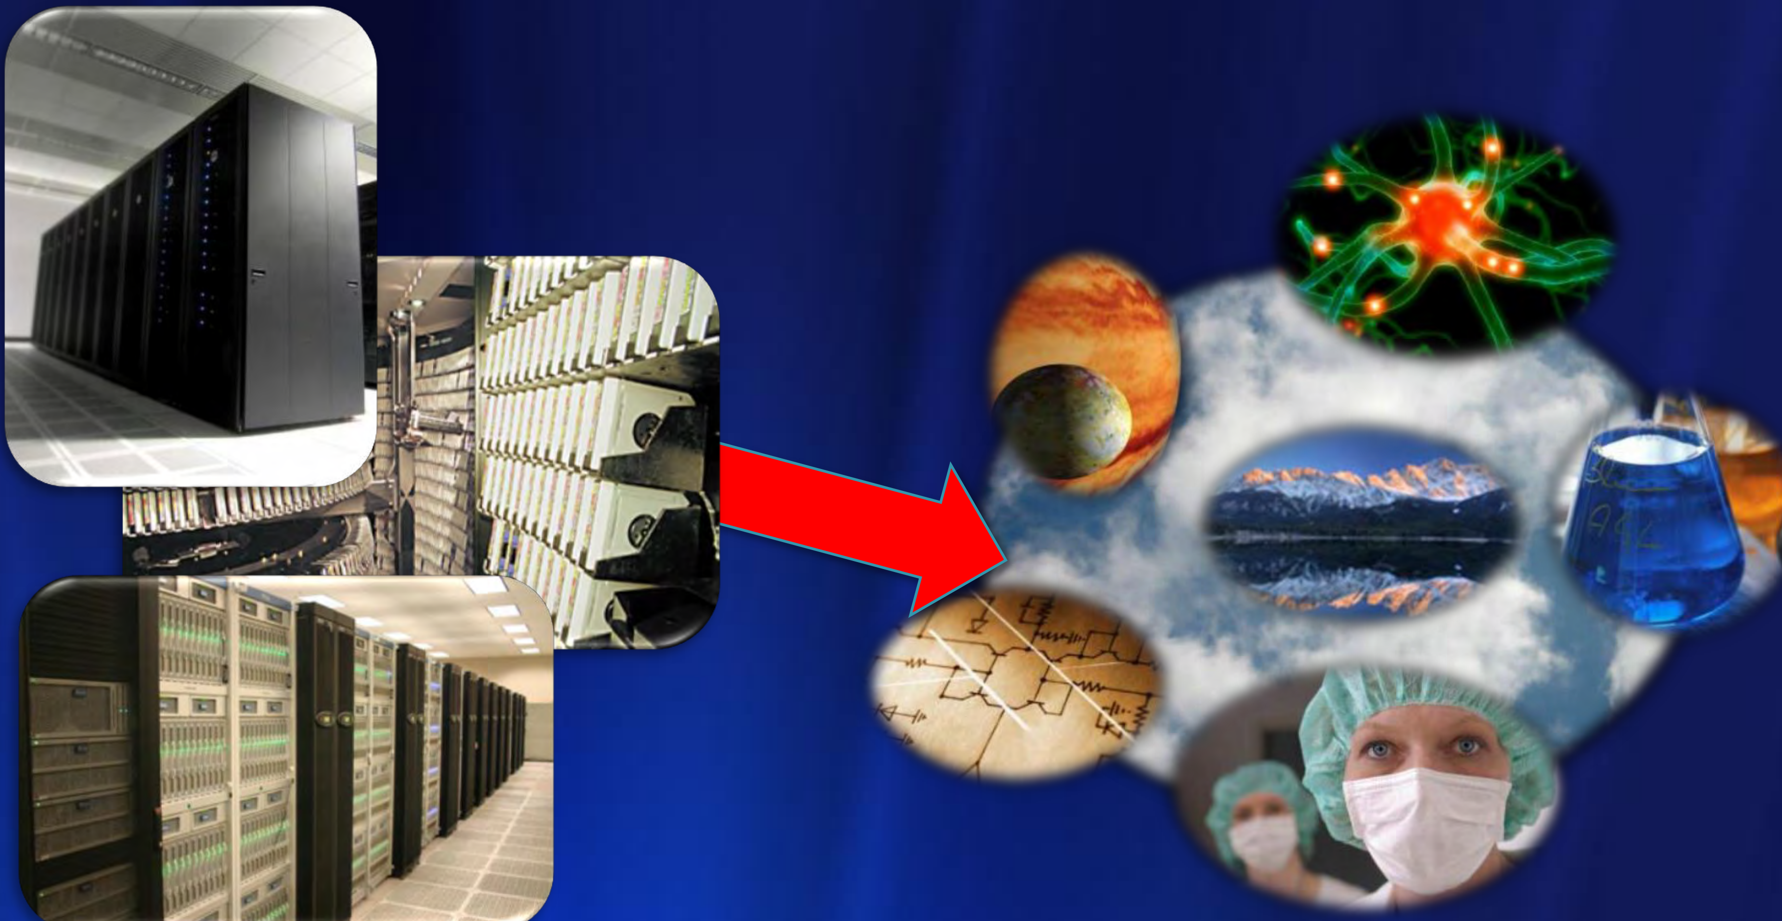

- Software + Data + Scientific Services = Insights

Microsoft

Today's Truisms (2009)

- Bulk computing is almost free
  - … but software and power are not
- Inexpensive sensors are ubiquitous
  - … but scientific data fusion remains difficult
- Moving lots of data is {still} hard
  - … because we're missing trans-terabit/second networks
- People are really expensive!
  - … and robust software remains extremely labor intensive
- Scientific challenges are complex
  - … and social engineering is not our forte
- Our political/technical approaches must change
  - … or we risk solving irrelevant problems

Microsoft

The Pull of Economics …

- Commercial economics
  - Manycore processors and accelerators
  - Software as a service and cloud computing
- They will drive change in technical computing
  - Just as did “killer micros” and inexpensive clusters
- This is a chance to reinvent computational science
  - … at a time of national economic crisis and competition

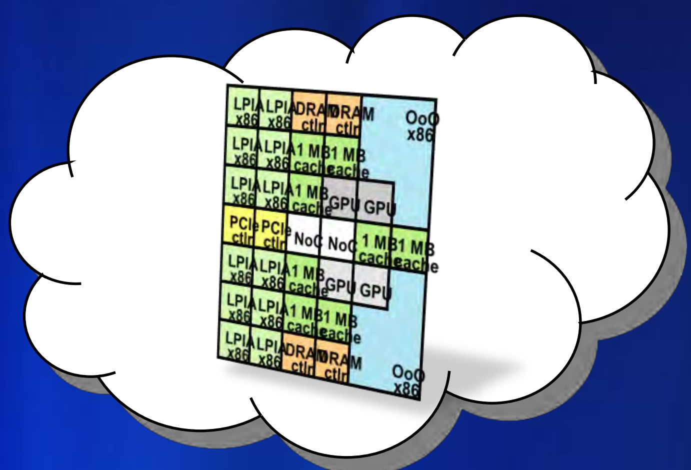

Next-Generation Applications

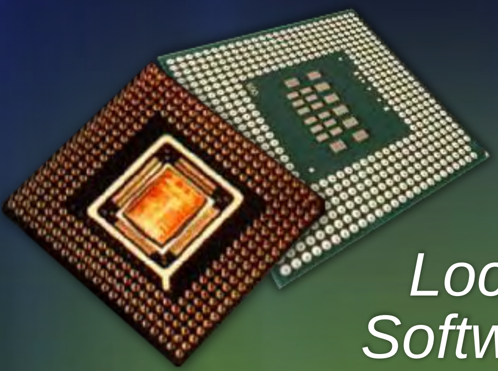

Concurrency Spectrum

Local Software

Software
+
Services

Global Services

Microsoft

Economics Drive Change

- Moore's "Law" favored consumer commodities
  - Economics drove enormous improvements
  - Specialized processors and mainframes faltered
  - The commodity software industry was born
  - Custom HPC hardware largely disappeared
  - Hard to compete against 50%/year improvement
- Implications
  - Consumer product space defines outcomes
    - It does not always go where we hope or expect
  - Research environments track commercial trends
    - Driven by market economics
    - Think about processors, clusters, commodity storage

Can Multicore Supplant Clock Rates?

- Double the number of cores instead of speed
- No, at least without major innovation
  - Sequential code
  - Lack of parallel algorithms
  - Difficult programming
  - Few abstractions

- Parallelism is changing the computing landscape
- If existing applications cannot use large parallelism
  - New applications and systems will arise
    - Software plus services
    - Mobile computing
- Clouds are an “obvious” outcome …

New Software Architecture

Embarrassingly Parallel Processing

- When applications are hosted
  - Even sequential ones are embarrassingly parallel
  - Few dependencies among users
- Moore's benefits accrue to platform owner
  - 2x processors →
    - ½ servers (+ ½ power, space, cooling …)
    - Or 2x service at the same cost
- Tradeoffs not entirely one-sided due to
  - Latency, bandwidth, privacy, off-line considerations
  - Capital investment, security, programming problems

Microsoft

11

The Service Continuum

Software ⟷ Services

- Rich user experiences
- Great offline support
- Security and privacy
- Compliance and regulations
- Customizability

Time-Space Fungibility In the Cloud

Data, data, data

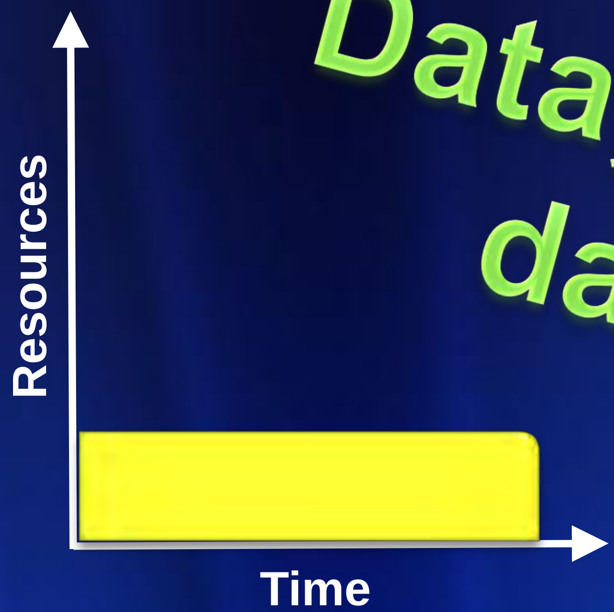

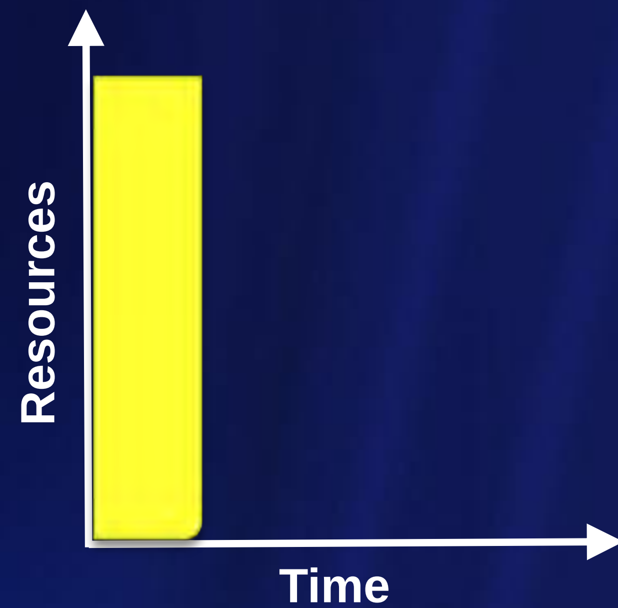

- Time-space fungibility
- Economies of scale

- New science
- Multidisciplinary data fusion

Cloud Opportunities

Ensure Service Continuity in the Cloud

- Move computation to the data
- Address failure and disaster scenarios

Scale On-Demand and Cost Effectively

- Scale data throughput and storage capacity
- Fuse and analyze multidisciplinary data

Support Emerging Applications Rapidly

- Enable rapid development of new web applications
- Easy access to consume multiple data sources

Reduce Infrastructure and Management Costs

- Hardware and software independence
- Lower operational cost of managing data

Cloud Application Frameworks

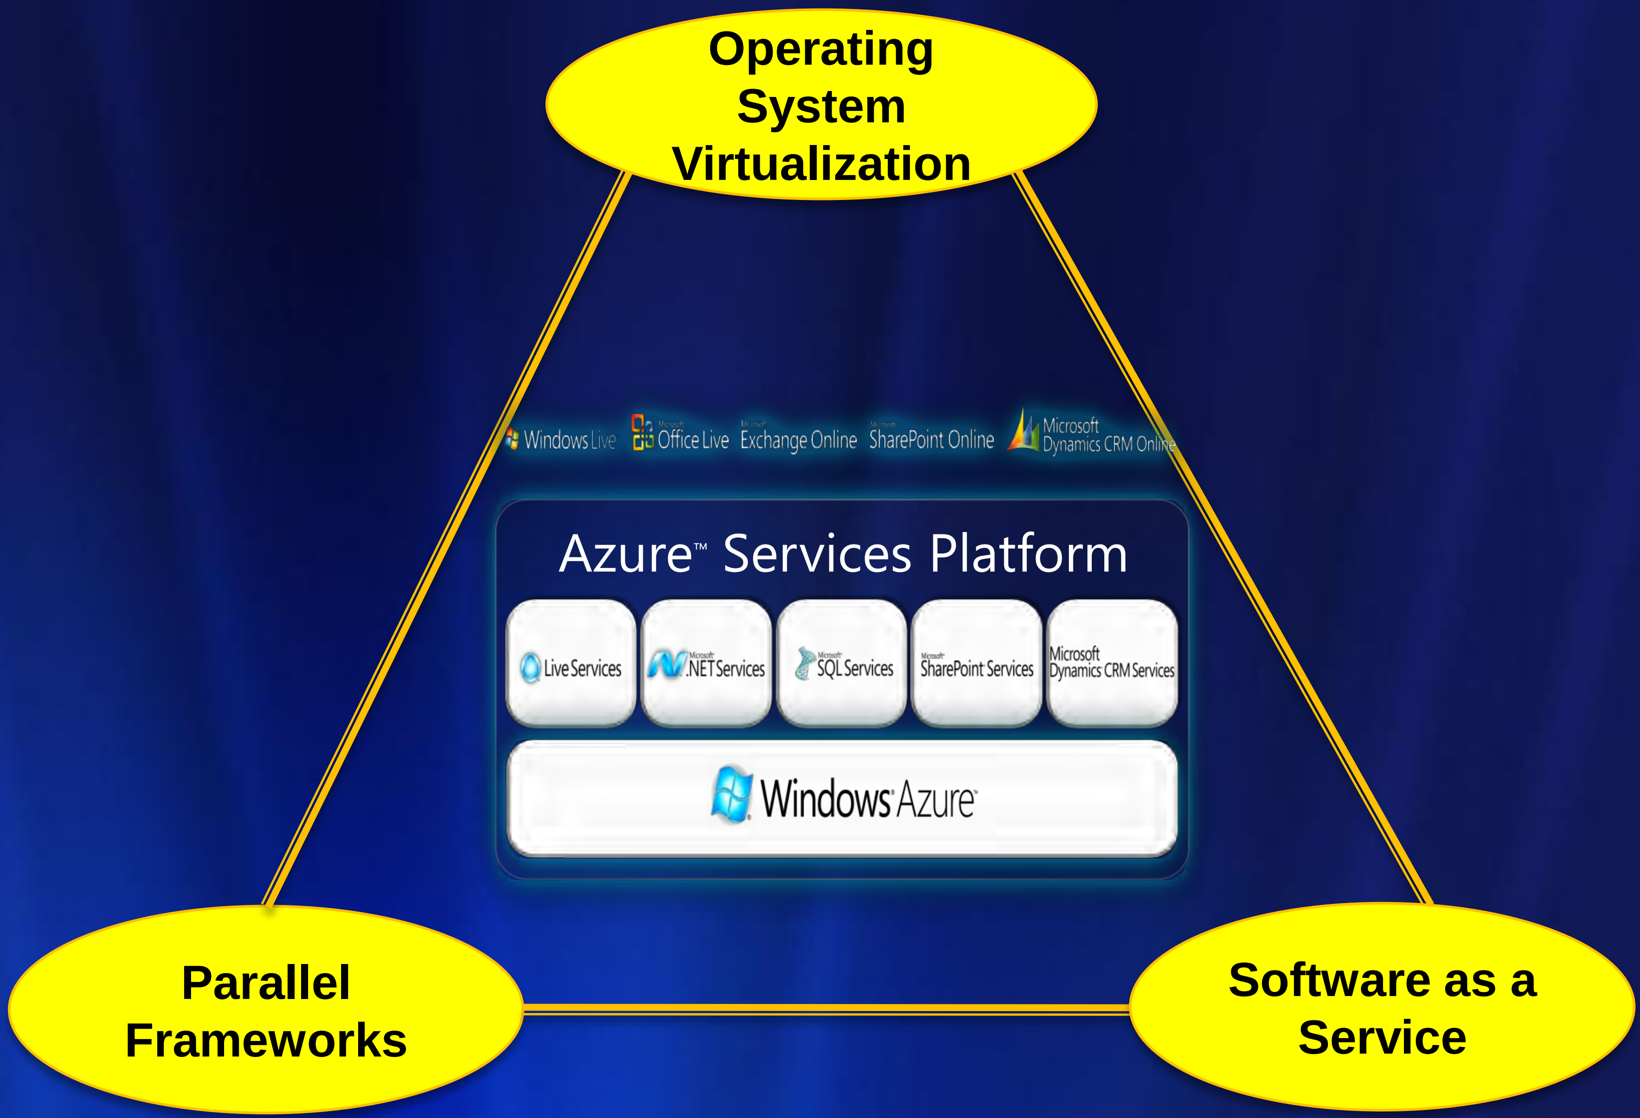

Windows Azure

Extending Windows to the Cloud

- Compute
  - Virtualized compute environment based on Windows Server
- Storage
  - Durable, scalable, and available storage with abstractions
- Management
  - Automated management of the service lifecycle
- Security
  - Authenticated transactions

3 Windows Live | Microsoft Office Live | Microsoft Exchange Online | Microsoft SharePoint Online | Microsoft Dynamics CRM Online

Azure™ Services Platform

2 Live Services | Microsoft .NET Services | Microsoft SQL Services | Microsoft SharePoint Services | Microsoft Dynamics CRM Services

1 Windows Azure™

www.azure.com

Microsoft

Interoperability

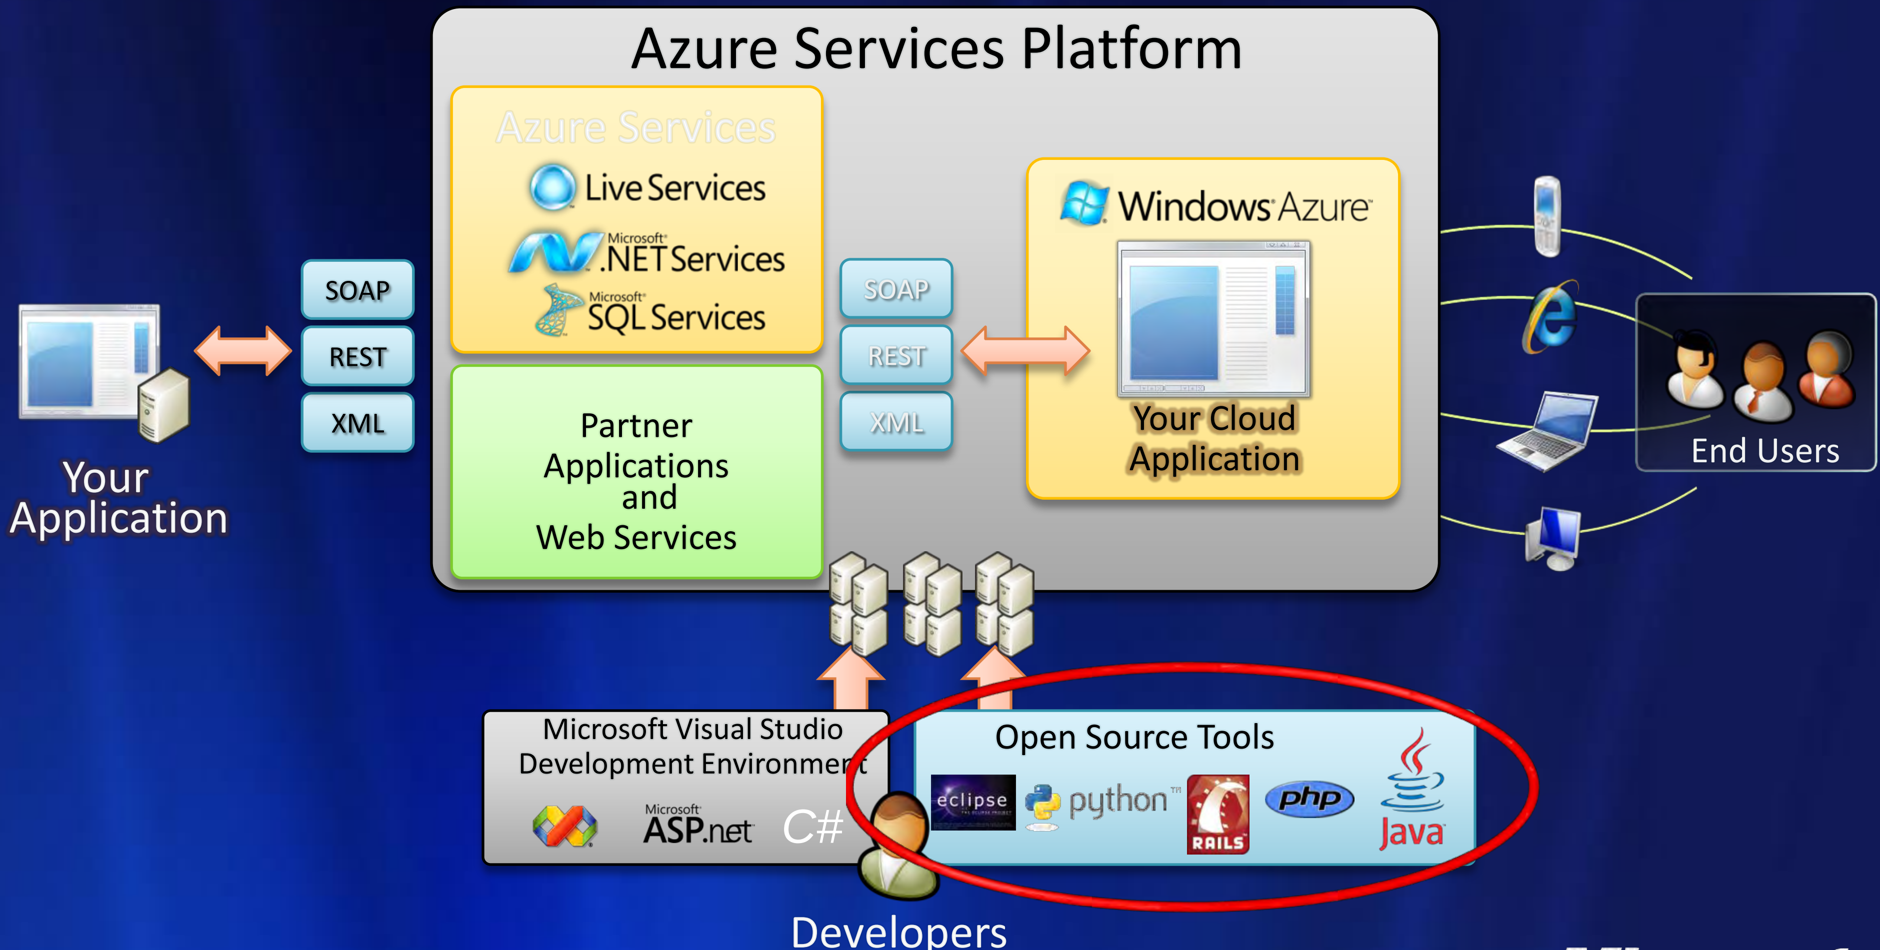

Azure Virtualization Architecture

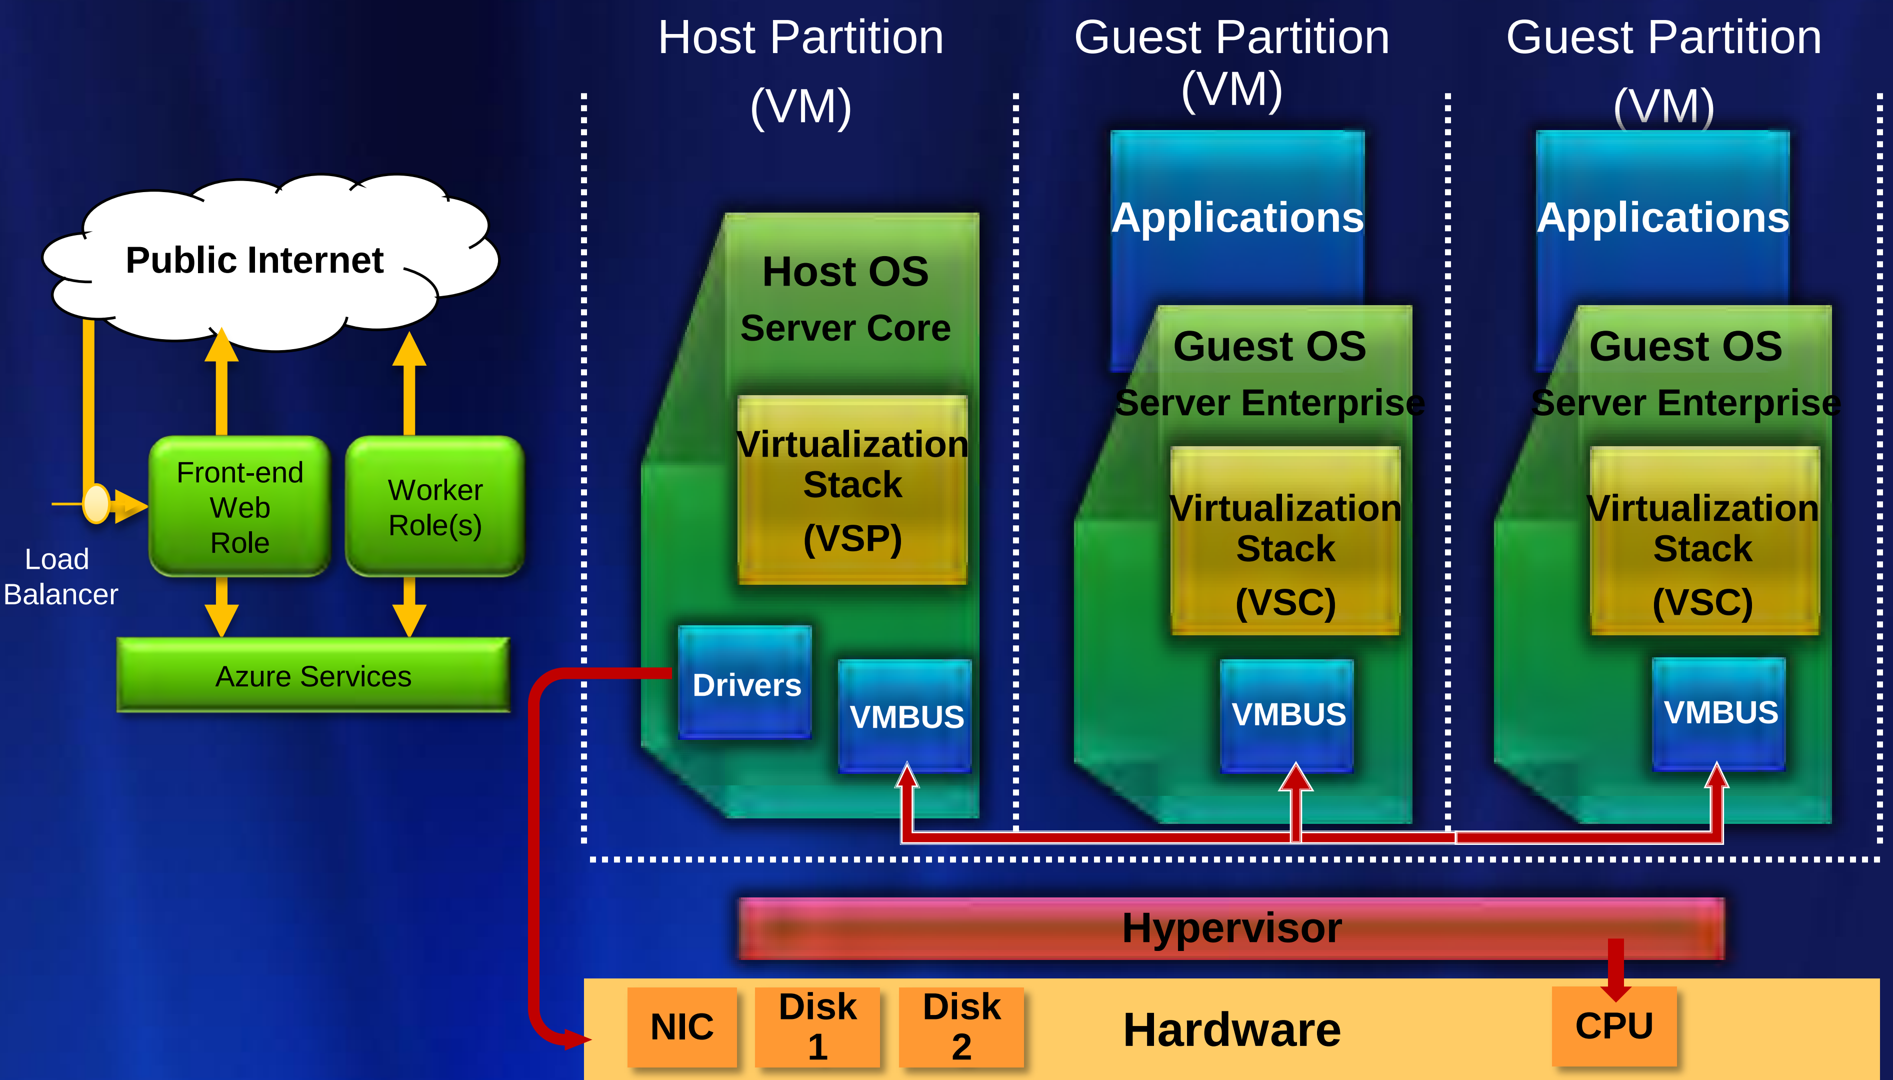

Windows Azure Fabric Controller

Out-of-band communication – hardware control

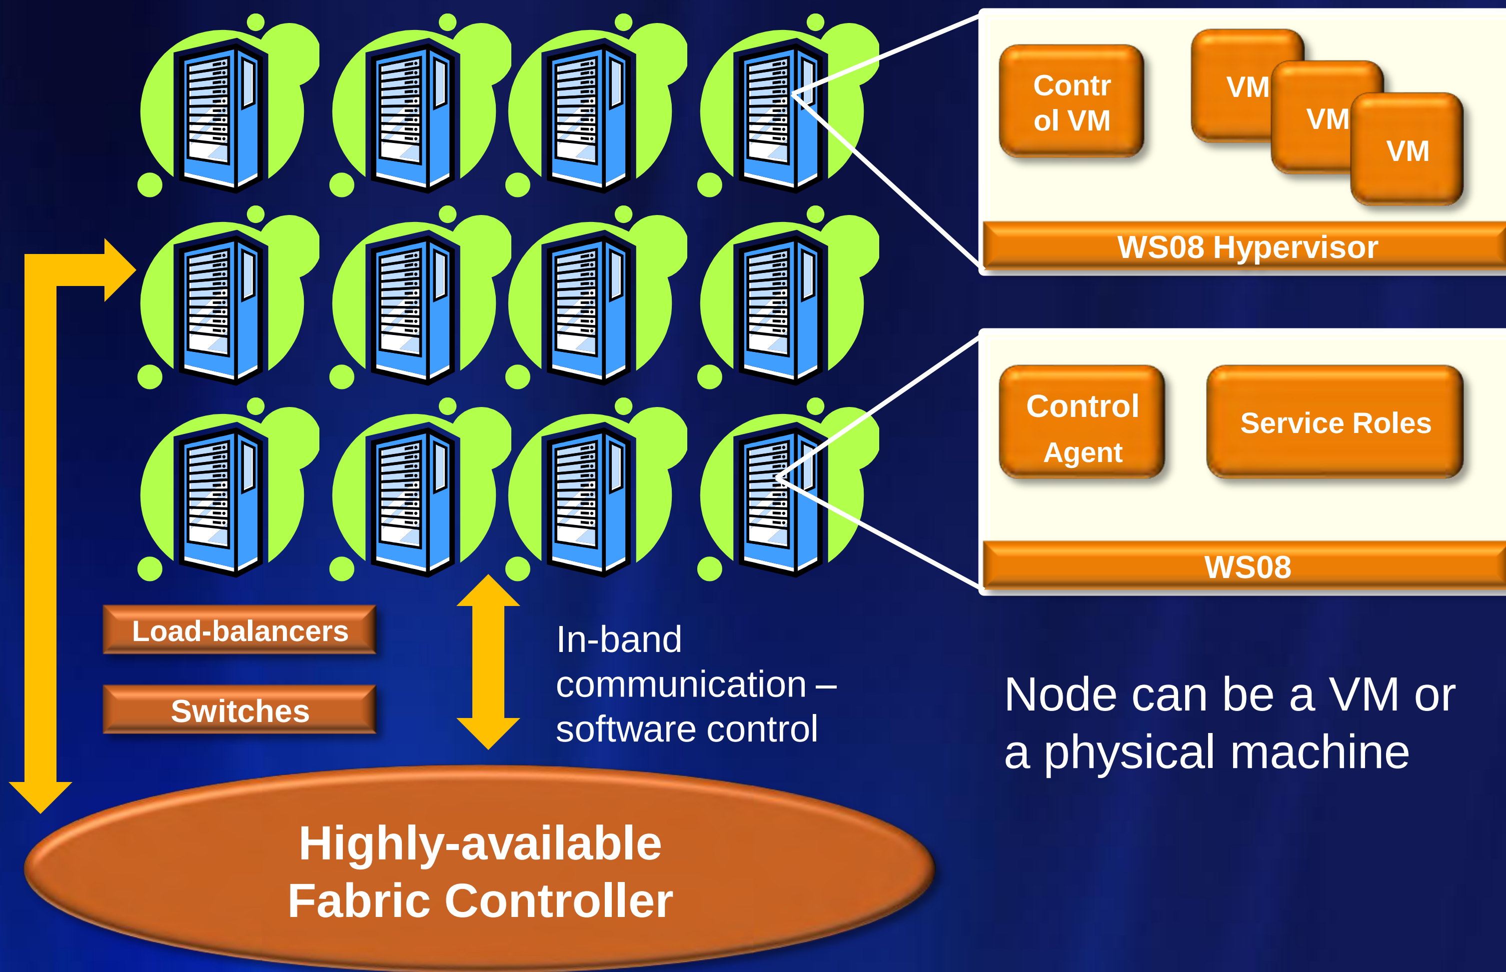

Node can be a VM or a physical machine

The Data Explosion

Experiments

Simulations

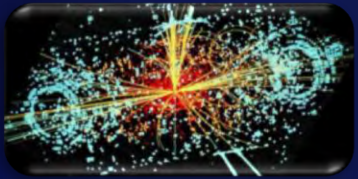

Archives

Literature

Cornell University Library

arXiv.org

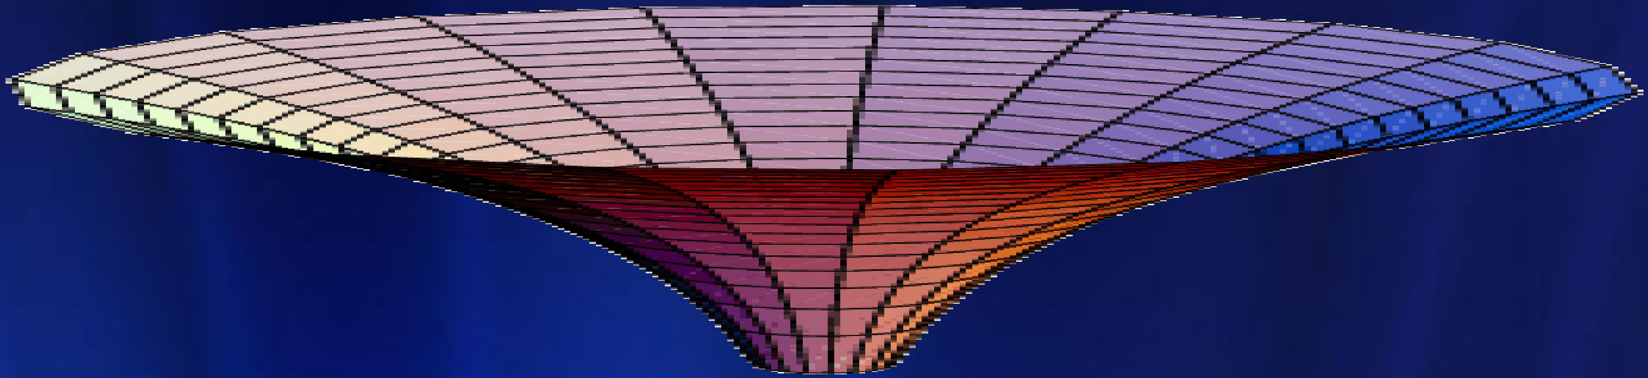

<u>Petabytes</u>

Doubling every 2 years

Microsoft

Social Implications of the Data Deluge

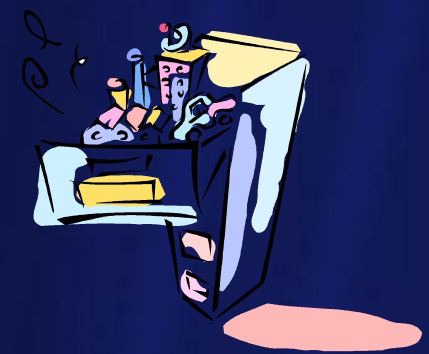

- Hypothesis-driven
  - “I have an idea, let me verify it.”
- Exploratory
  - “What correlations can I glean from everyone’s data?”
- Different tools and techniques
  - Exploratory analysis relies on deep data mining
    - Supervised and unsupervised learning
  - “grep” is not a data mining tool
    - Higher level services can make a difference
- Massive, multidisciplinary data
  - Rising rapidly and at unprecedented scale

Azure Data Storage

- Two levels
  - Basic Azure storage
    - Blobs, tables and queues
    - At least triple replication for reliability
  - SQL database services
    - Most of SQL Server built atop basic storage
- Blobs
  - A simple interface for storing named files and metadata
- Tables (structured storage)
  - A set of entities, which contain a set of properties
  - Relations on “training wheels”
- Queues
  - Reliable storage/delivery of messages for an application

Microsoft

DryadLINQ: From LINQ to Dryad

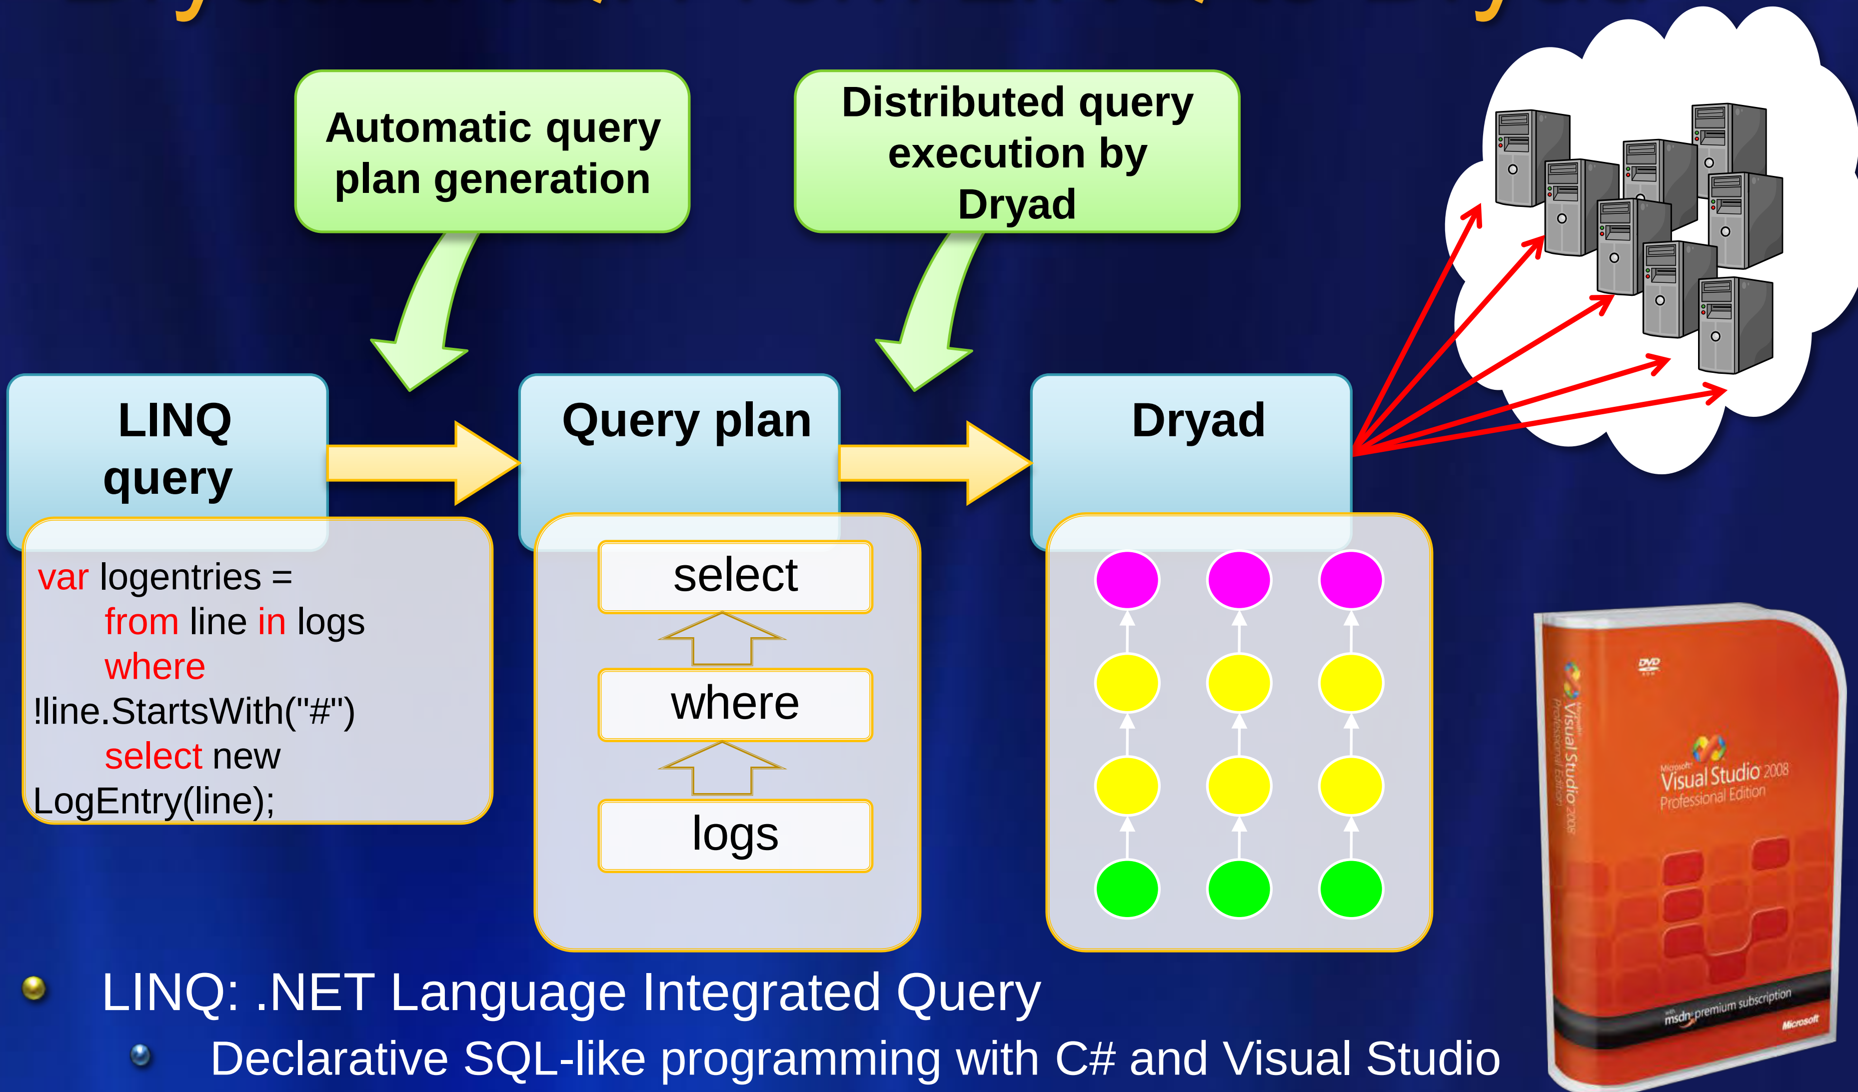

- LINQ: .NET Language Integrated Query
  - Declarative SQL-like programming with C# and Visual Studio
  - Easy expression of data parallelism
  - Elegant and unified data model

Source: Yuan Yu *et al*

Microsoft

Big Doesn't Begin To Describe it

Microsoft

Site Selection Best Practices

Internet Population

Internet Peering/Network

Mobile Users

Power Pricing

Environmental

Construction Costs

Tax Climate

IT Labor Availability

Corporate citizenship

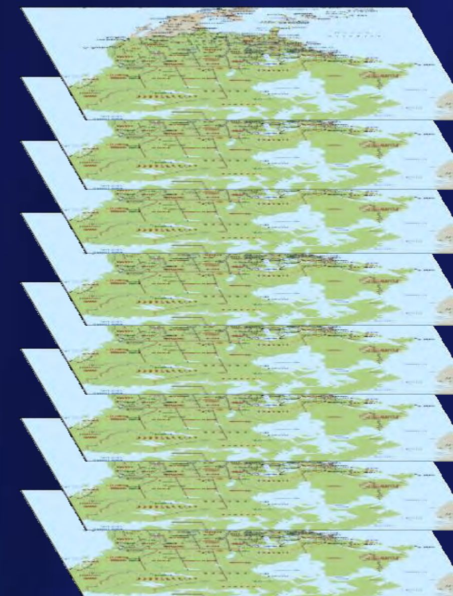

Composite Heat Map

Data Center "PacMan"

- Land - 2%
- Core and shell costs – 9%
- Architectural – 7%
- Mechanical/Electrical – 82%
  - 16% increase/year since 2004

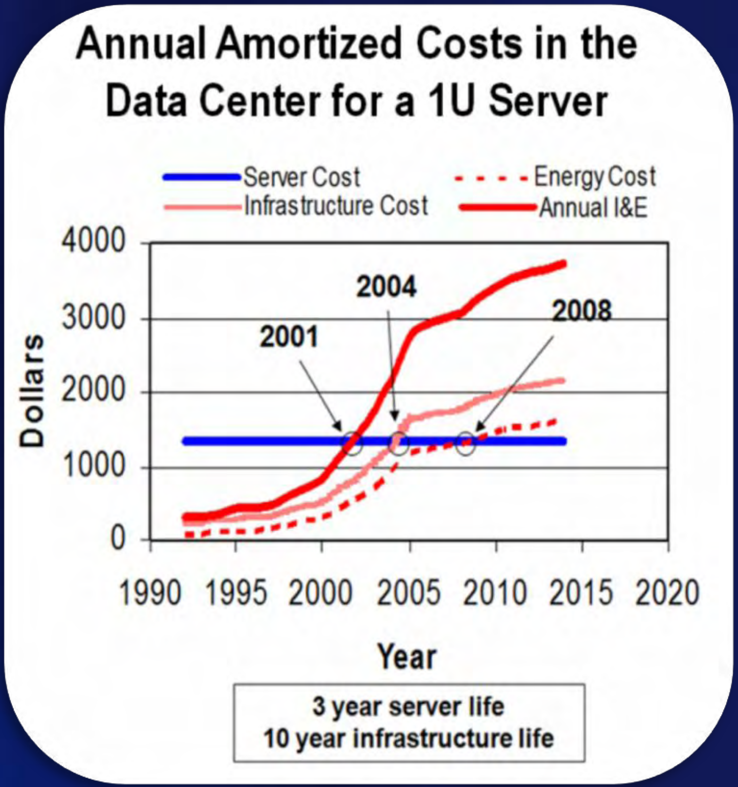

Belady, C., "In the Data Center, Power and Cooling Costs More than IT Equipment it Supports", *Electronics Cooling Magazine* (February 2007)

Source: Mike Manos/Christian Belady

More Similarities Than Differences

- Twins separated at birth
  - Exascale HPC systems
  - Megascale data centers
- Research commonalities
  - Energy efficiency and carbon footprint
  - Reliability and resilience
  - Interconnect and photonics
  - Memory stacking and access
  - Exascale data management
  - Manycore and processor functionality
  - System software and management
  - Programmability and models

Microsoft

Cyberinfrastructure Components

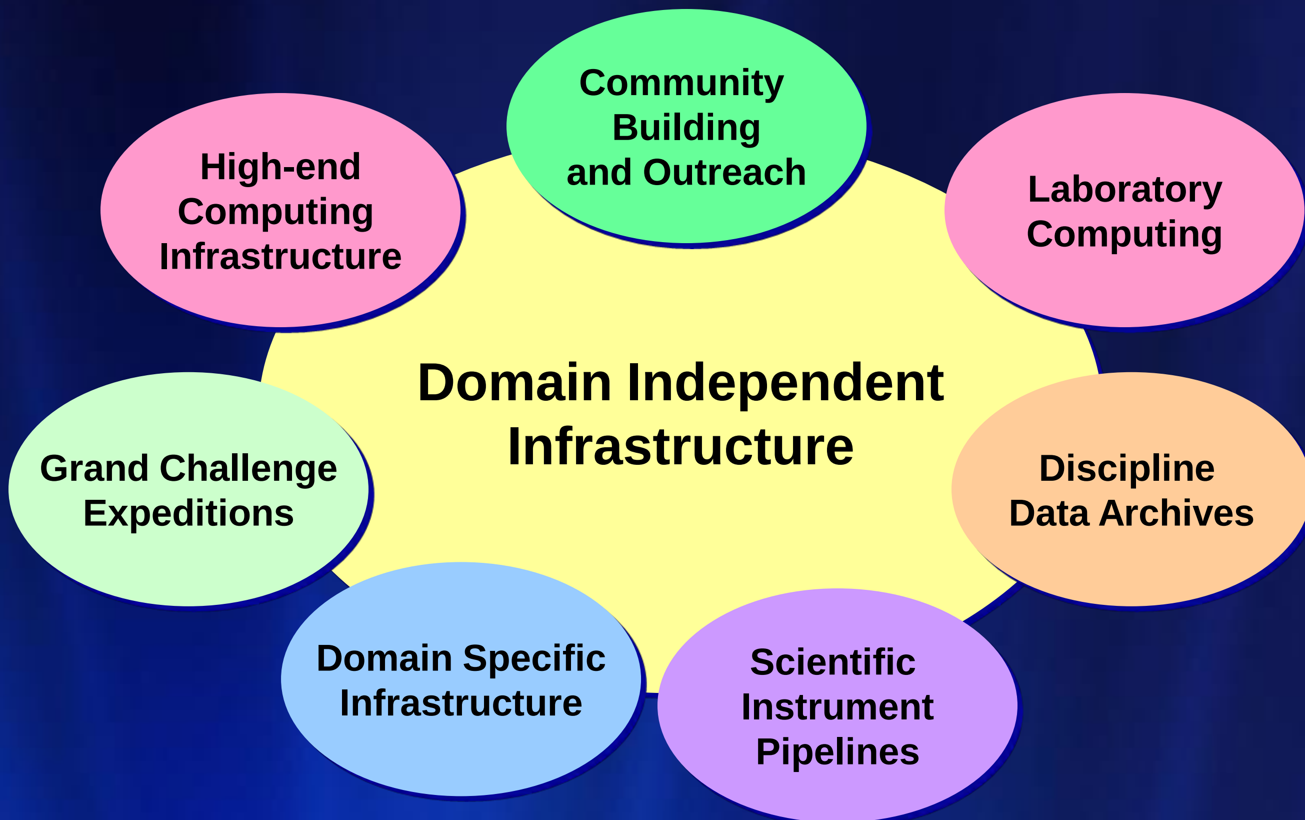

Our Decadal Changes

- Commodity clusters
  - Proliferation of inexpensive hardware
    - “Attack of the Killer Micros”
  - Race for MachoFLOPS
  - Broad base for enabling software
  - Low level programming challenges
- Rise of data
  - Scientific instruments and surveys
  - Storage, management and provenance
  - Data fusion and analysis
- Distributed services
  - Multidisciplinary collaborations
  - Complex stacks and reliability challenges
  - Less broad base for enabling software
  - Multi-organizational social engineering

Microsoft

Research Infrastructure Challenges

- Insatiable demand
  - Cycles, storage, software, support
- Distributed acquisition/deployment
  - Duplicative, non-shared infrastructure
- Distributed cost structures
  - Power, space, staff, staff, hardware
- Long-term sustainability
  - Decades rather than months/years
- The shape of the triangle
  - Apex versus mainstream users

Research Infrastructure Observations

- Business
  - Capital is cheap
  - Labor is expensive
  - Costs are usually explicit …
    - … and had better be lower than revenues!

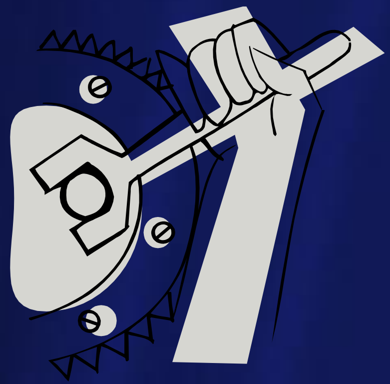

- Academia and government
  - Capital is (seemingly) expensive
  - Labor is (seemingly) cheap
    - Student and faculty time
  - Many costs are implicit …
    - … and often skew realistic assessment

Microsoft

Major DOE Science Resources

- NERSC
  - 38,640 Opteron cores (Cray XT4)
- ORNL
  - 150,152 Opteron 2.3 GHz cores (Cray XT5)
- ANL
  - 163,840 PowerPC 450 cores (IBM Blue Gene/P)

The Computing Continuum

Parameter Studies

Hosted Applications

Data Analysis/Fusion

Desktop Acceleration

Azure

Windows HPC Server 2008

The Computing Pyramid

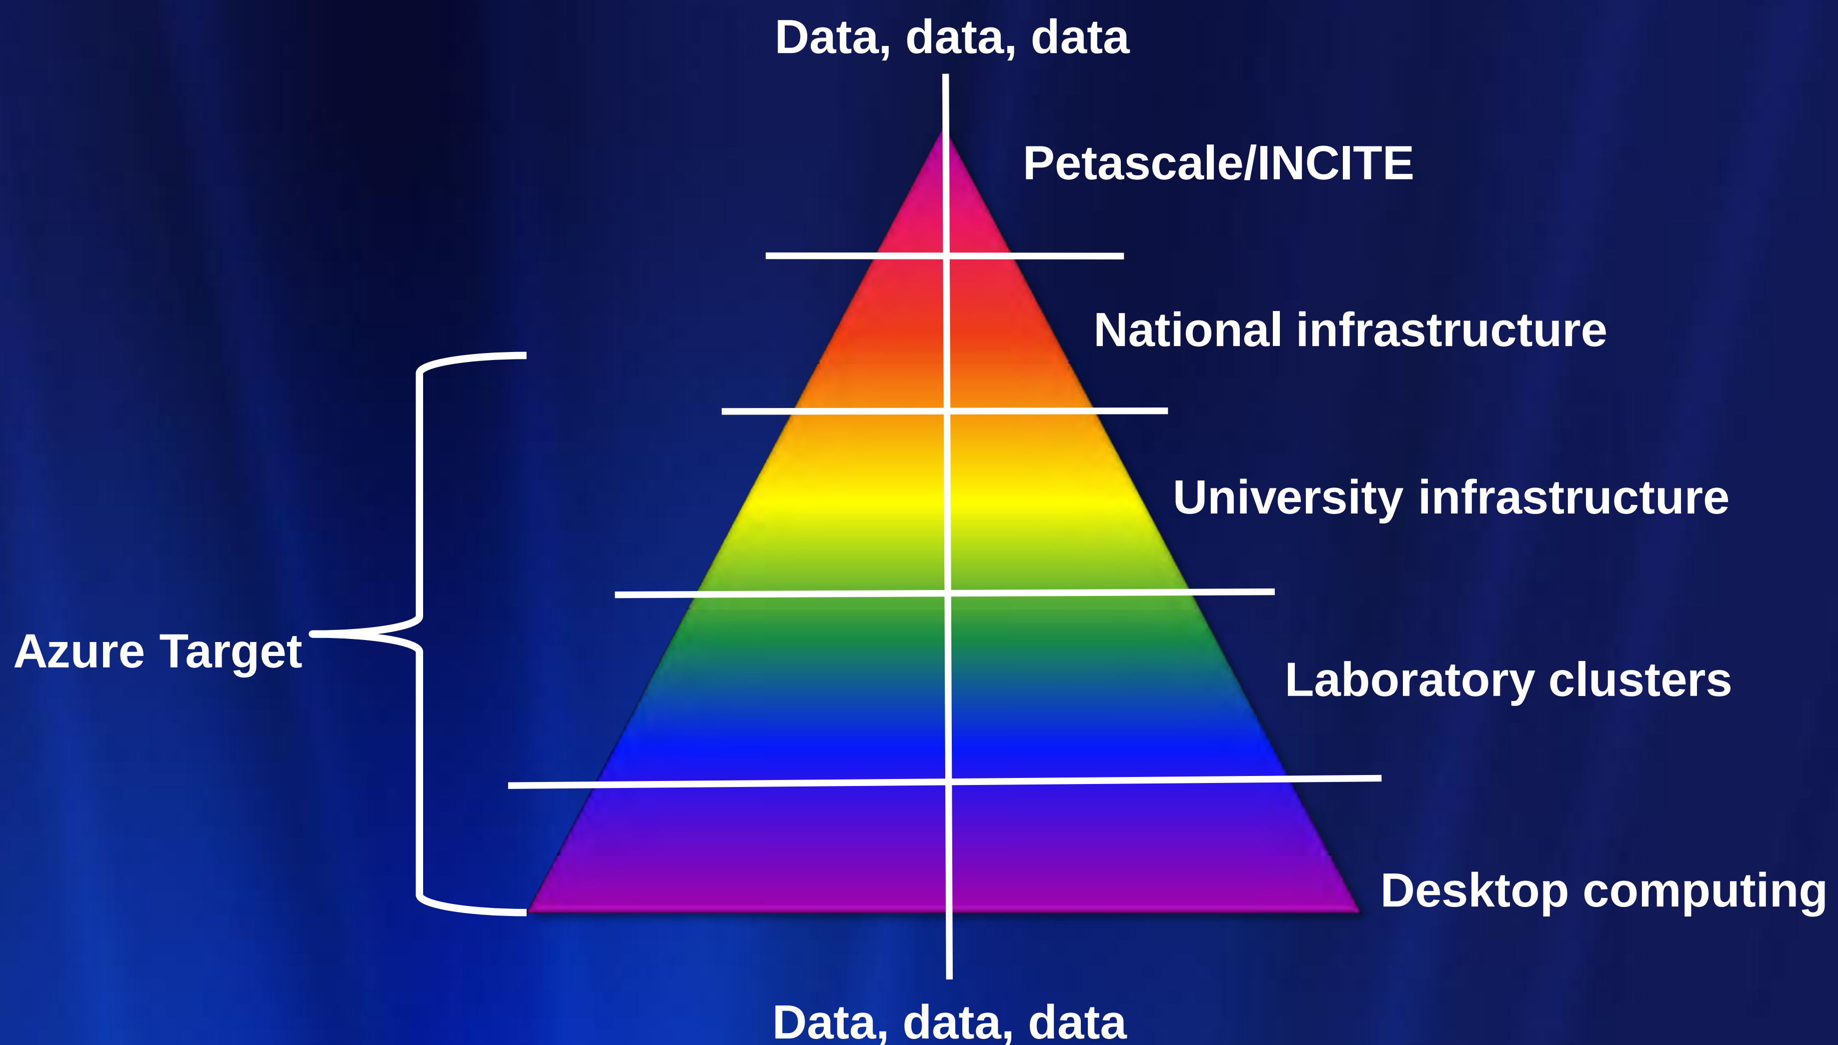

Today Is An Inflection Point

- Economic challenges
  - Research efficiency
  - Infrastructure scaling
- Technology transition
  - Cloud software+services
  - Multicore and storage scaling
- We can change the game, just as we did before
  - Rich cloud services
  - Hosted infrastructure
  - Commodity economics
- Microsoft wants to engage the community
  - Let's work together …

Science In The Clouds

Source: Tony Hey

Microsoft

Your potential. Our passion.™