2009 System z Expo

October 5 – 9, 2009 – Orlando, FL

Session Title: Cloud Computing 101 What every z Person must know

Session ID: ZDI08

Frank J. De Gilio - degilio@us.ibm.com

Authorized IBM Training

facebook Home Profile Friends Inbox Frank De Gilio Settings Logout

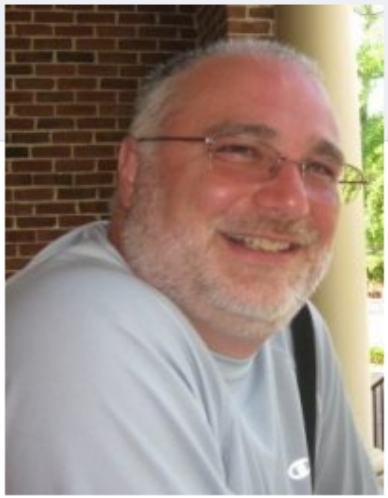

View Photos of Me (46)

Edit My Profile

Frank is just dis guy - you know?

Information

Networks:
Poughkeepsie, NY

Relationship Status:
Married to
Sue Piccone De Gilio

Birthday:
February 3, 1962

Friends

238 friends See All



Frank De Gilio Another Great Day at work! 12 seconds ago clear

Wall Info Photos Compare People Cities Ive ... Notes

What's on your mind?

Attach Share

Options

Frank De Gilio Another Great Day at work!
12 seconds ago · Comment · Like

RECENT ACTIVITY

Frank commented on Susan Buley Baiano's status.

Frank commented on Craig Johnston's photo.

3 more similar stories

Frank De Gilio

Another list

44 ODD Things about you! If you opened this, FILL IT OUT! Learn 44 things about your friends, and let them learn 44 things about you! And if you opened it, you obviously have nothing to do right now, so just copy and paste and fill it out! 1. Do you like blue cheese? Not even at gunpoint! 2...

Yesterday at 8:48pm · Comment · Like · Share

Stephanee Berardi Stratford likes this.

Theresa Hoffman Kennedy number 5 almost made me choke on my ice cream. I was laughing so hard.
Yesterday at 9:26pm · Delete

Nancy DeGilio Lindsay Love the ADD que
about an hour ago · Delete

Nancy DeGilio Lindsay I had to laugh when I saw

View of Cloud Computing

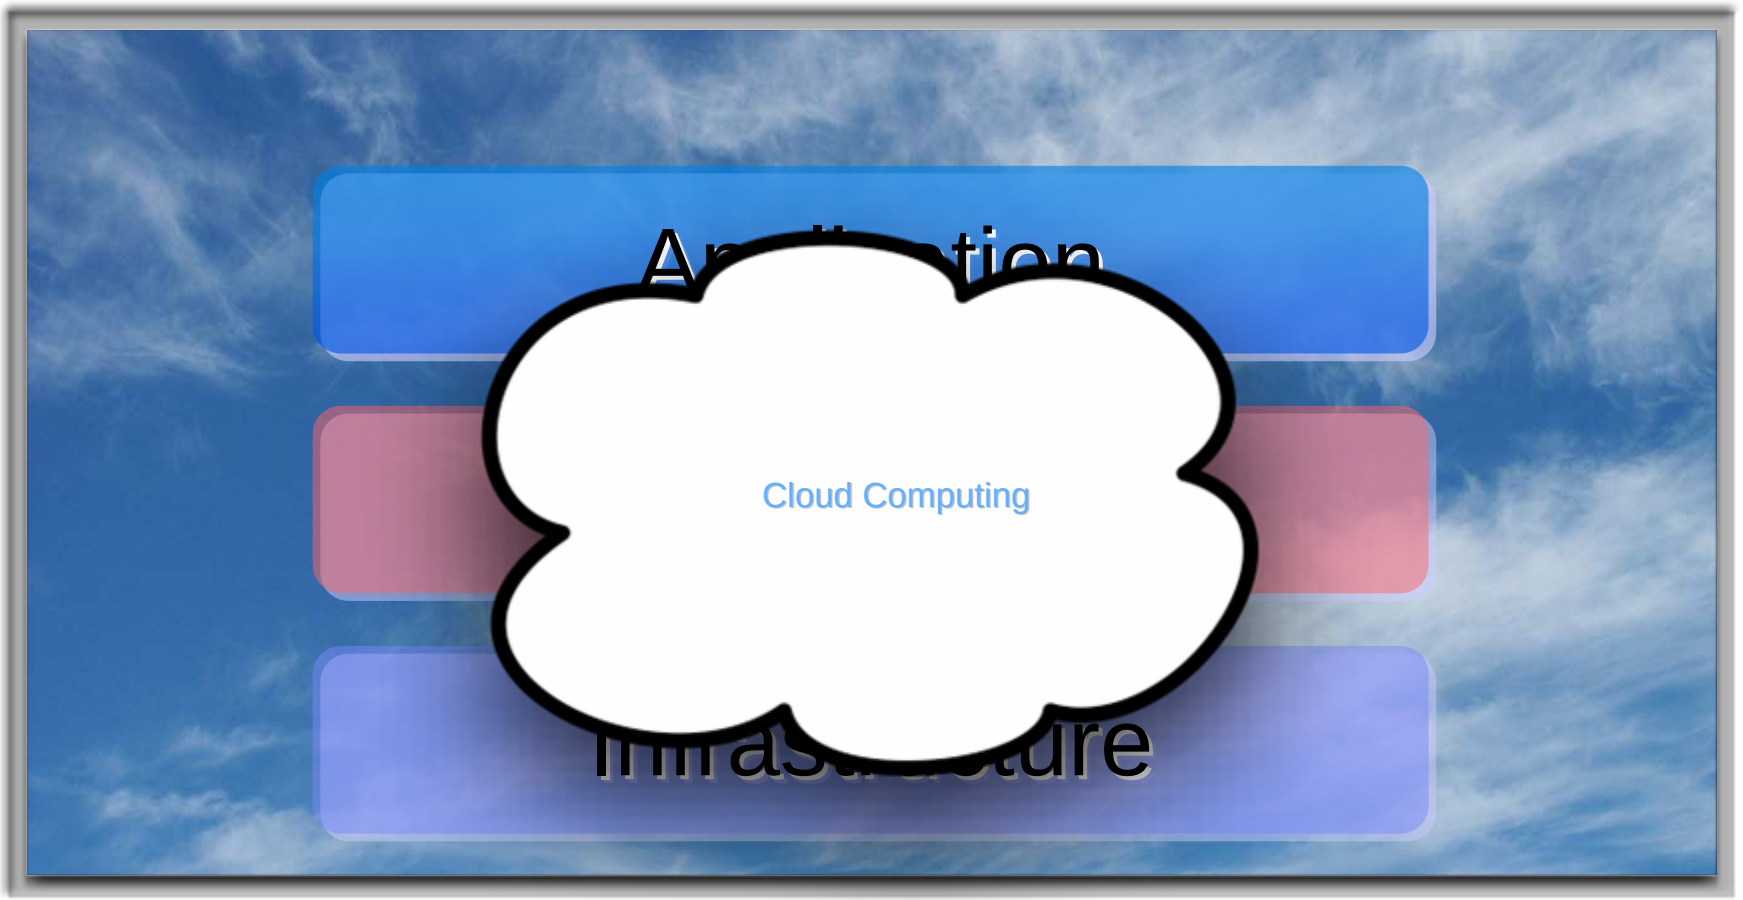

Logical Blocks of Cloud Computing

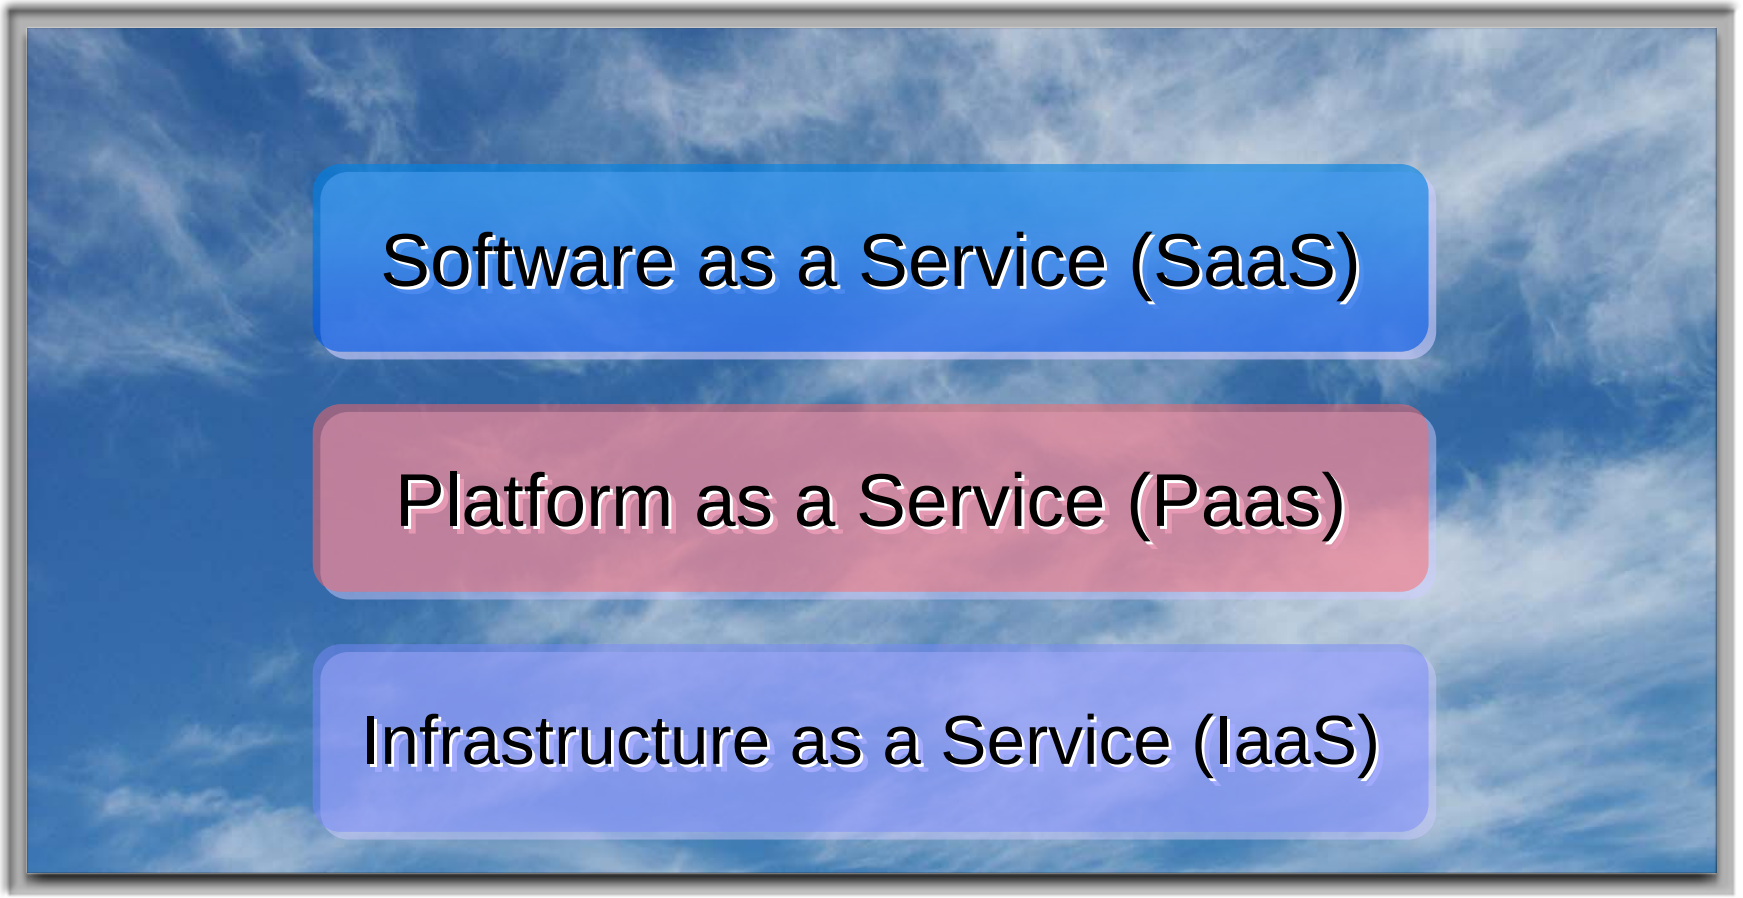

Cloud Computing Genealogy

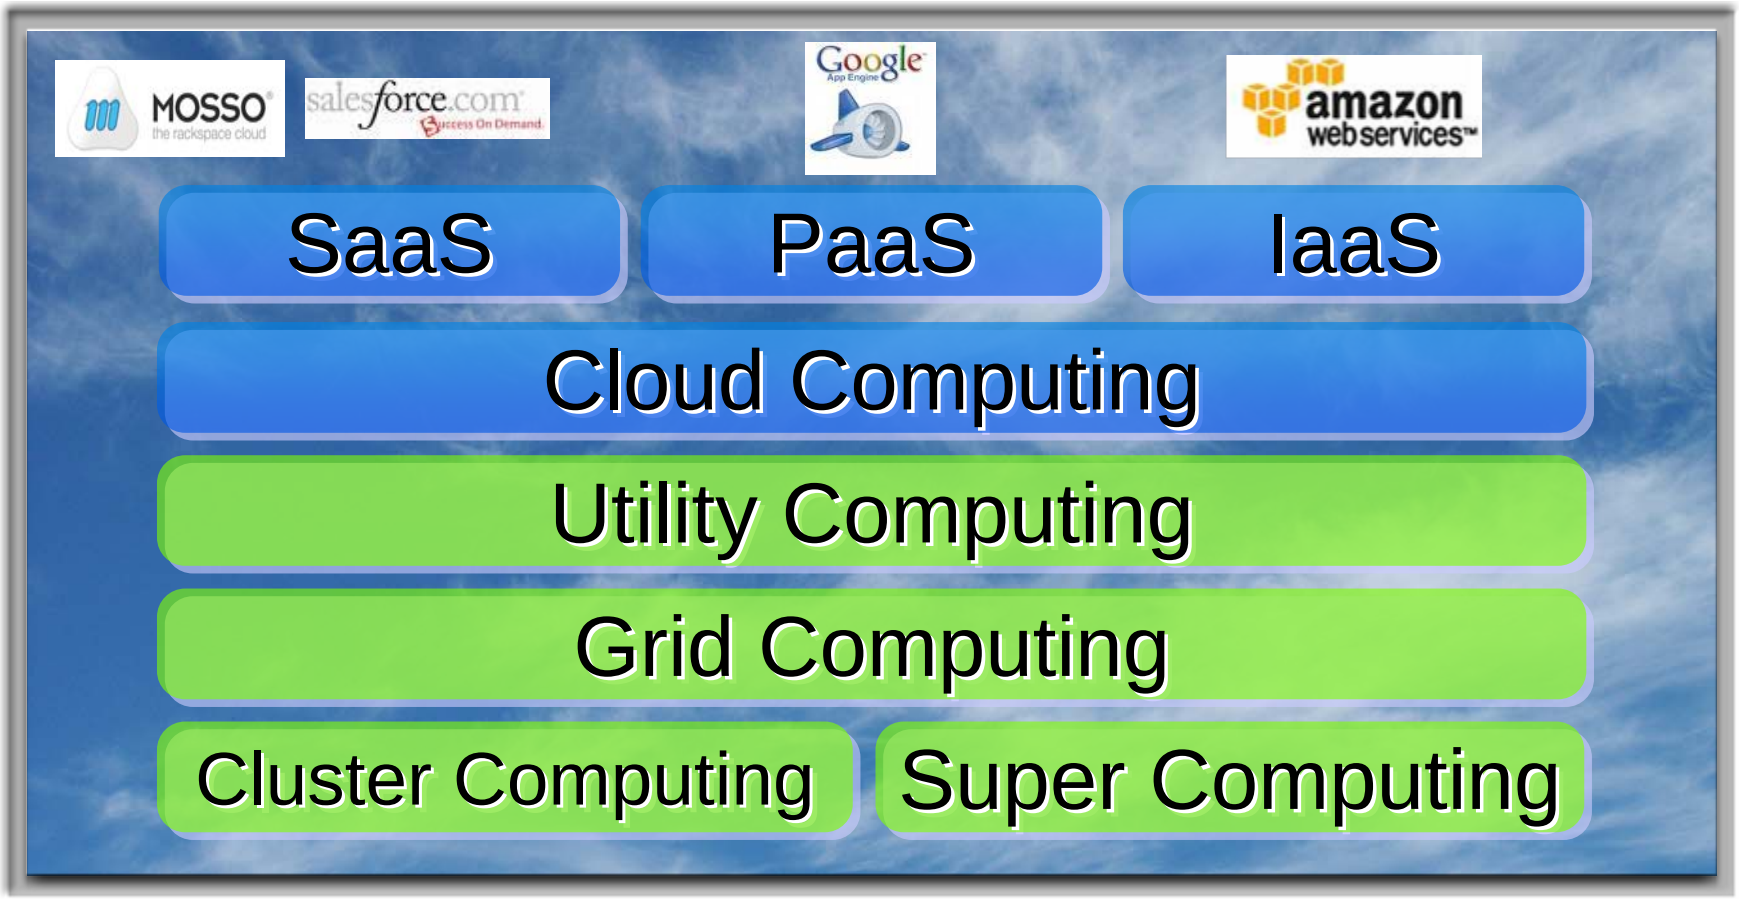

Basic Cloud Definition

- Shared computing resources
  - Virtualized
  - Service
  - Using an API
- The cloud enables a virtual data center.
  - Local
  - Remote
  - Service provider
- Resources are offered as a service
  - As needed
  - Highly scalable
  - On demand

This is a Logical Progression

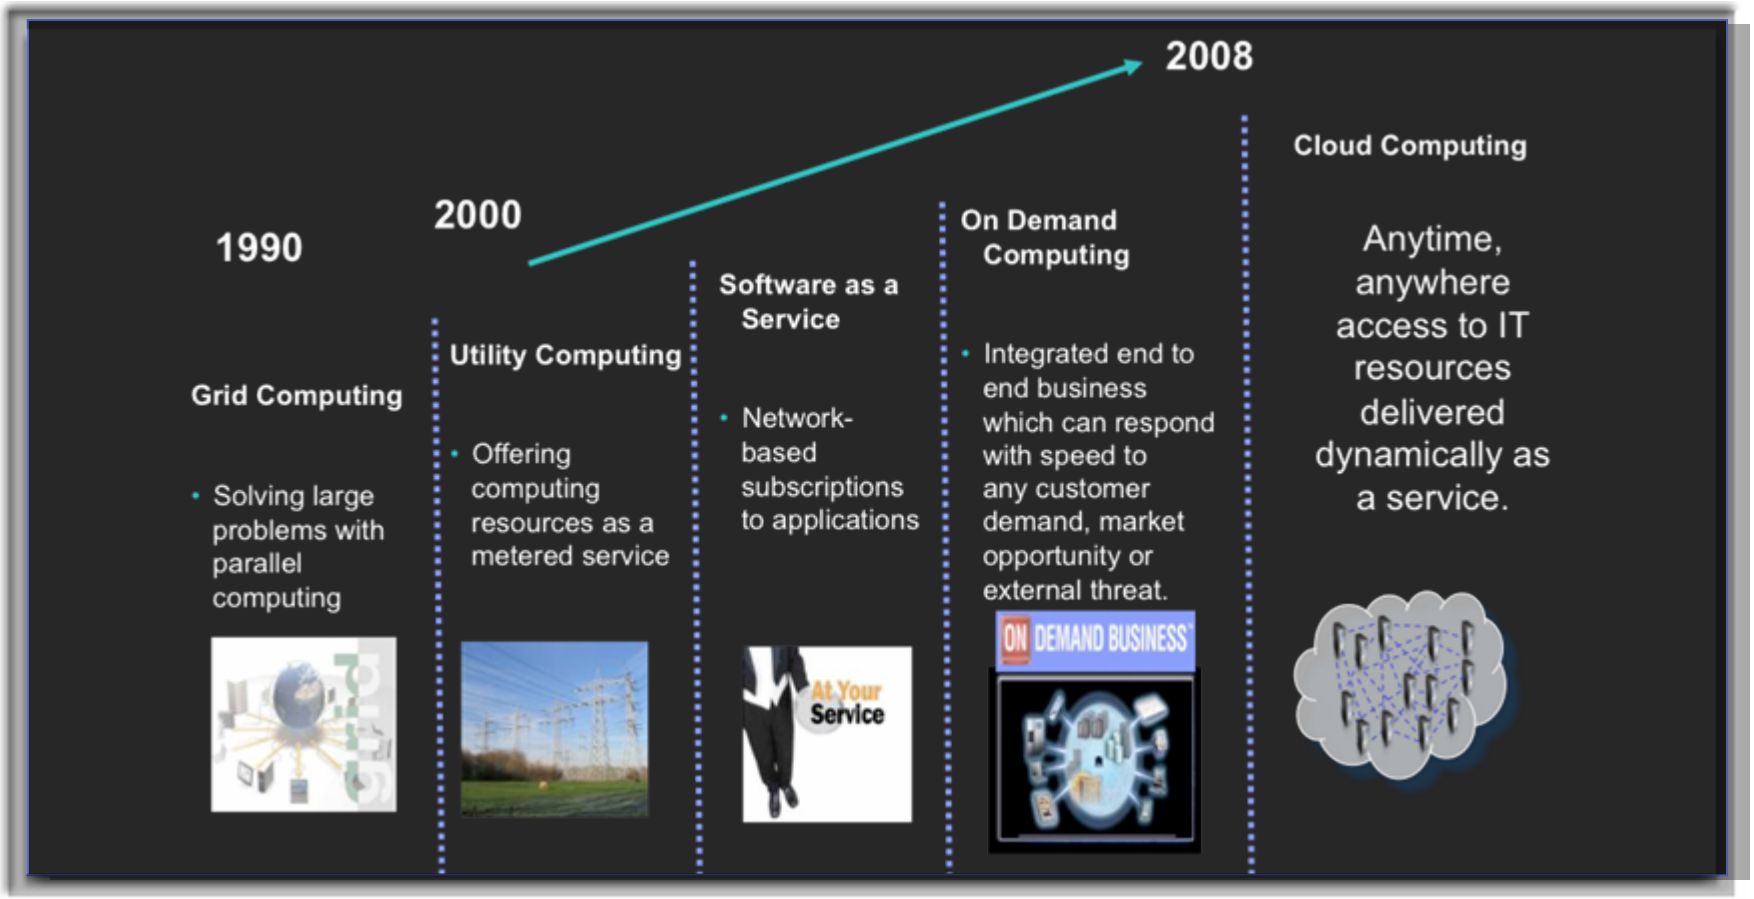

Forerunners to Cloud

- Super Computing - High Cost / High Performance machines, uses specialized hardware to do massive parallel tasks blurring into the cluster space for the last 15 years
- Cluster Computing - Poor man's super computer uses software to fill the gap of non-specialized hardware to behave as a single computer
- Grid Computing - Virtual super computer built of multiple independent clusters not necessarily in one administrative domain
- Utility Computing - computing resources as a metered service, offers low initial cost and scaling to peak demand. Although this has not built on the previous technologies we have seen it grow in parallel to them (particularly grid) and it plays a significant role in Cloud

Cloud Service Models

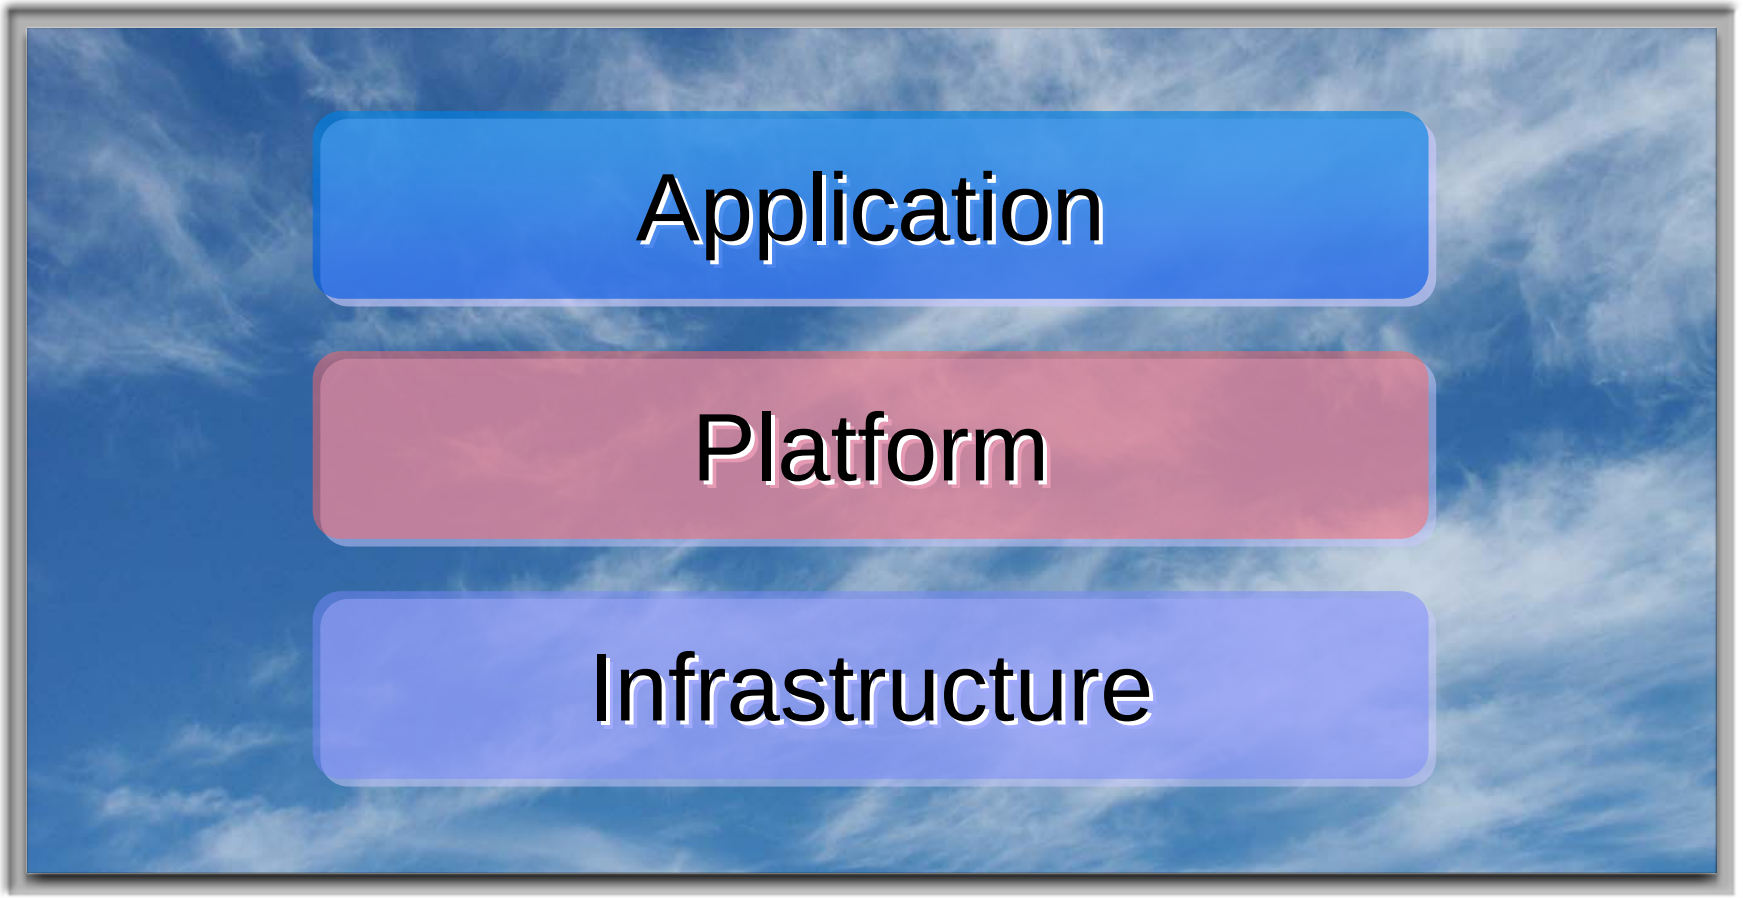

Infrastructure as a Service

- Oldest Piece of Cloud Computing Concept
- Seen in leased servers, virtual machines, and multi tenant web server accounts
  - Older iterations lack management api / tools
  - Scalability on demand
  - Automatic load balancing
- Starting to work well with new technology
- Pioneering the utility / pay as you go model
- Elastic

Technology Enablers for IaaS

- Virtual Appliances with stripped OS and Applications
- Open Virtualization Format (OVF) for interoperability
- Virtualization based on VMWare and Xen (open source)
- Provisioning for OS & Applications
- Automation for creation, cloning, and deletion
- Autonomic scale increase behind load balancers

Platform as a Service

- Infrastructure as a Service with tightly coupled application framework
- Higher level of abstraction with slower adoption
- Best example of this is Google AppEngine, Mosso, etc.
- Provides application environment and application mobility through the platform to provide scale
- This is in an early stage, no real open standards, be concerned about vender lock-in and interoperability

Software as a Service

- Provides applications as the atomic unit. Builds on IaaS and PaaS.
- Good example of this is Salesforce.com running on Google AppEngine.
- Often the PaaS is abstracted away and only the application is exposed to end users.
- Typically the domain of single users and SMB markets.
- Concerns focus on application / data security. SLA guarantees, and future interoperability and migration.
- Providing offerings of applications from the enterprise is a more interesting driver rather than consuming.

Cloud Deployment Models

What Trends Are Driving The Cloud Computing Trend?

Infrastructure Technologies: Virtualization, Automation, SLAs

Application Technologies: Grid, MapReduce, Hadoop, SOA, Web 2.0

Data Intensive Applications: From massively parallel (e.g. Google) to large data files (e.g. You Tube)

Computing & Network Appliances: Special servers designed to handle specific tasks are blurring the lines between Network and Data Center

Open IT: Open Technologies, APIs, protocols, data formats, software platforms / data (e.g. Creative Commons, Open Data License)

Business Agility: Enter new markets. Deploy new application services. Stay ahead of competition.

Broadband: Growth in Internet bandwidth enabling ubiquitous connectivity. Increased reliability and functionality embedded in the network.

Utility Computing: Get as much computing power as you need when you need it, pay for only what you use.

Industrialization of IT: Standardization, and commoditization (e.g email). Falling costs of storage.

Mobility: Explosion of form factors, cell phones/connected devices, Proliferation of sensors

New Business Models: Advertising, Services, Subscription

Web Applications and Platforms: Mashable applications and services built on Web Oriented Architecture (e.g. REST, RSS/ATOM)

Data Center Pressures: Growing costs of power and space, server sprawl

Source: Gartner, Thomas Weisel Partners, Merrill Lynch, IBM AM

Business Drivers for Cloud

- **Deployment Speed**: Get services out as quickly as possible.
- **Businesses want fast prototypes**: Research and Development projects, low priority business applications and collaboration services are all good candidates for the cloud.
- **Lead Times/Costs:** Rapid Provisioning obviates the need for buying servers.
- **Financial Model:** Don't need to tie up capital in IT assets.

Barriers to Adoption

- **Security & Privacy:** Where is my data?
- **Compliance Issues:** Regulations may prohibit the use of clouds for some applications.
- **Reliability:** High availability will be a key concern.
- **Cloud Management:** Service Monitoring / Reporting / Management Technologies immature
- **Cost:** Economies of scale only go so far.
- **Customization:** In the cloud you get what you get!
- **It's New:** We don't know what we don't know.
- **Corporate Culture:** Will going to cloud eliminate my job?

What is Different about Cloud?

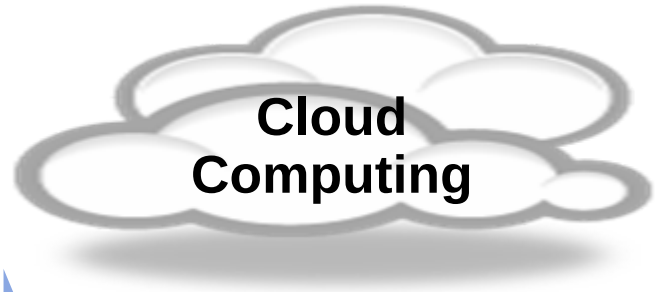

	Traditional Computing		Cloud Computing
Delivery Model	Buy assets and build delivery architecture	→	Buy external **service**
Interface Model	Internal network or intranet	→	Via the **Internet** using standard Internet IFaPs (IP, HTML, HTTP)
Business Model	Pay for fixed assets and administrative overhead	→	Pay directly based on **usage** or indirectly (e.g., subsidized by advertising)
Technical Model	Single tenant	→	Scalable, elastic, dynamic, **multi-tenant**

Source: Gartner

Mapping Workloads to Architectural Affinity

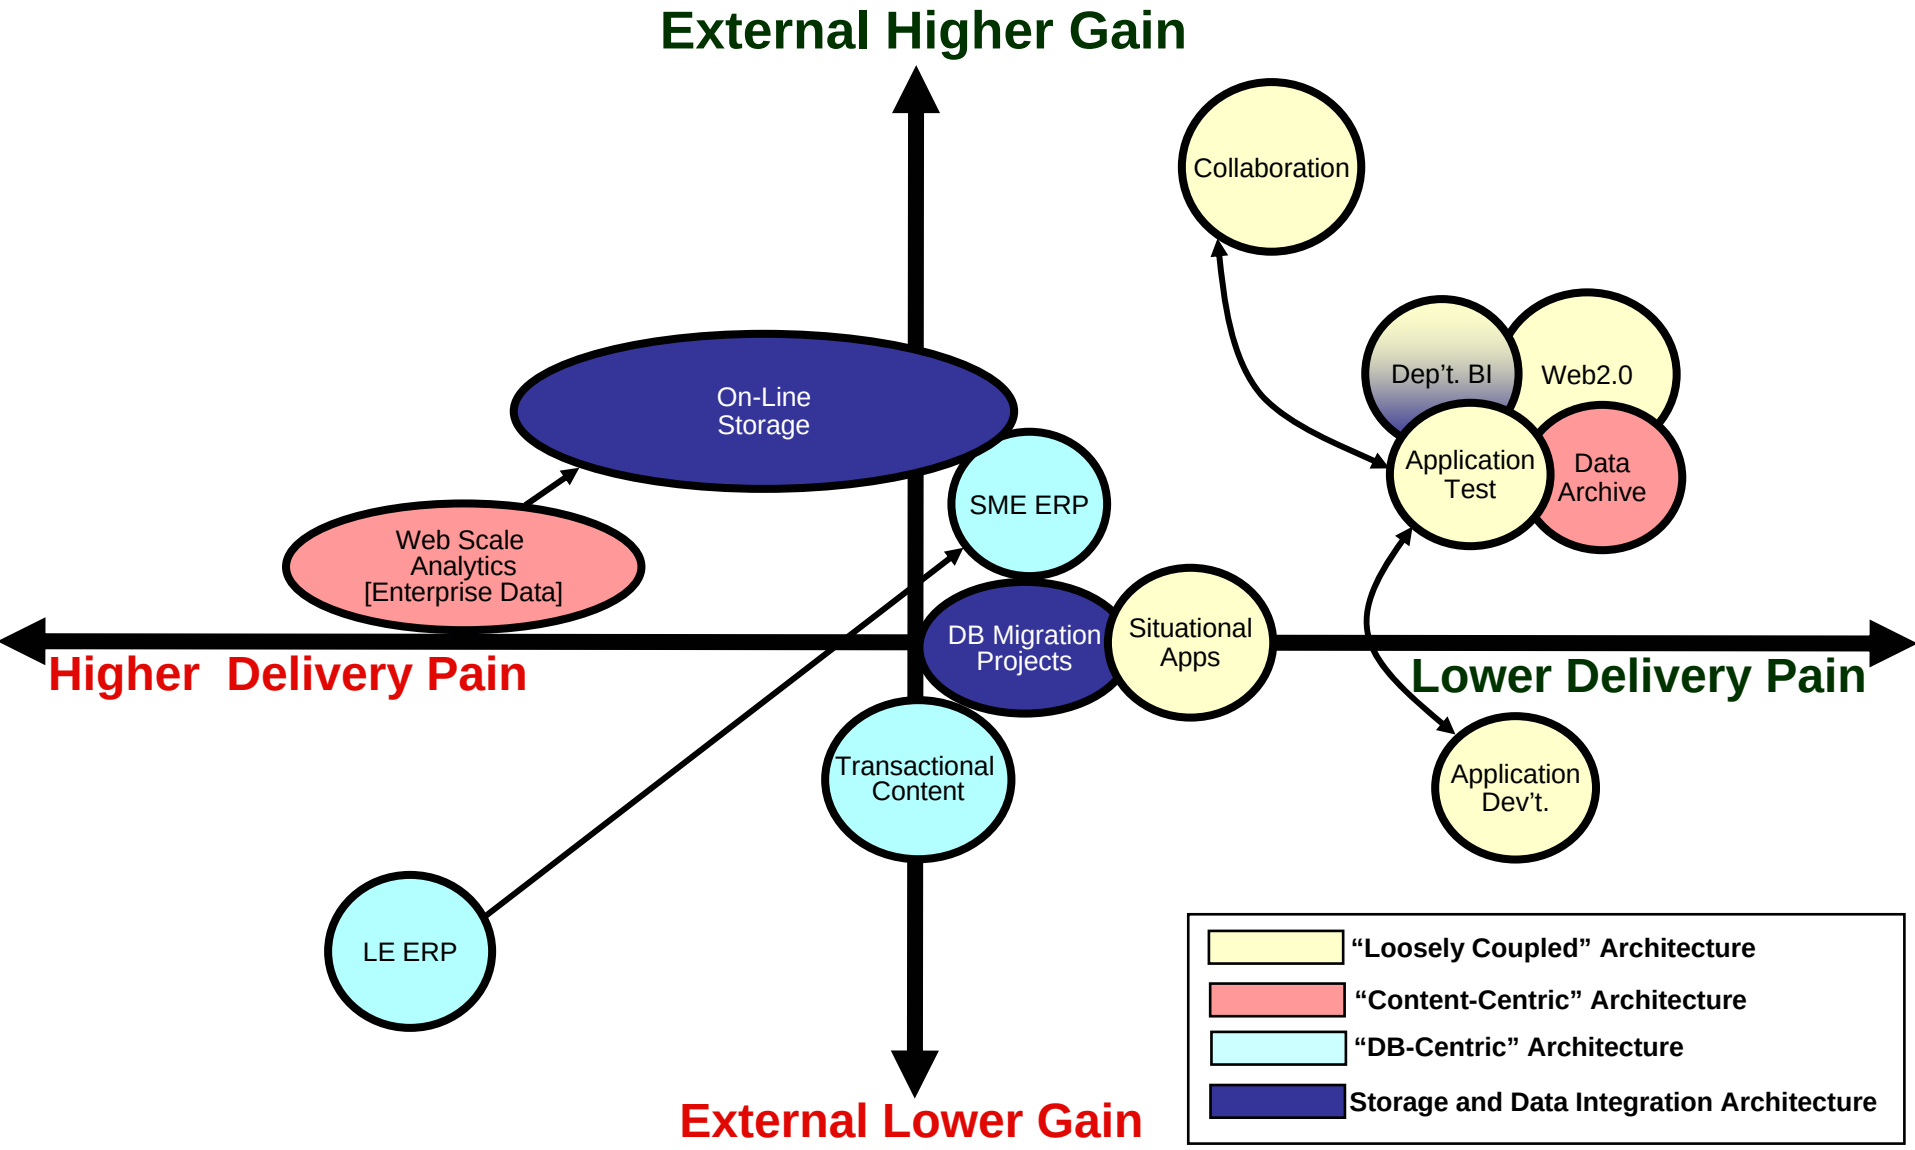

THE TECHNOLOGY CASE FOR IAAS IN THE ENTERPRISE

State of the Yesterday

Division of Responsibilities With Virtualization

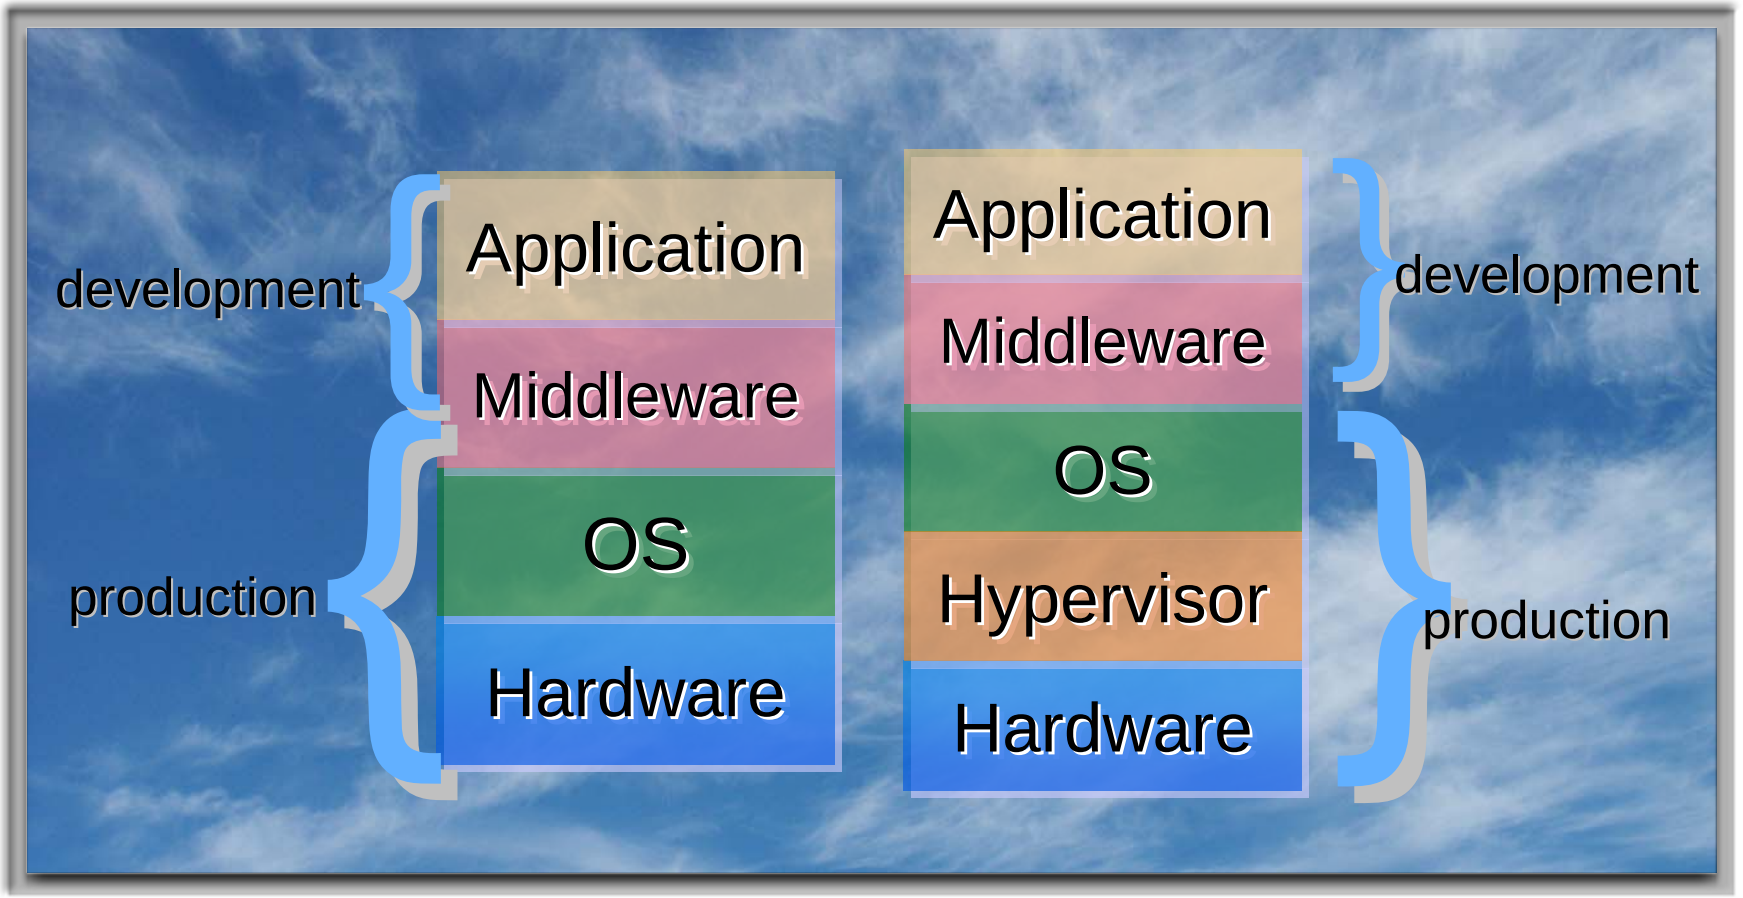

Important Next Steps

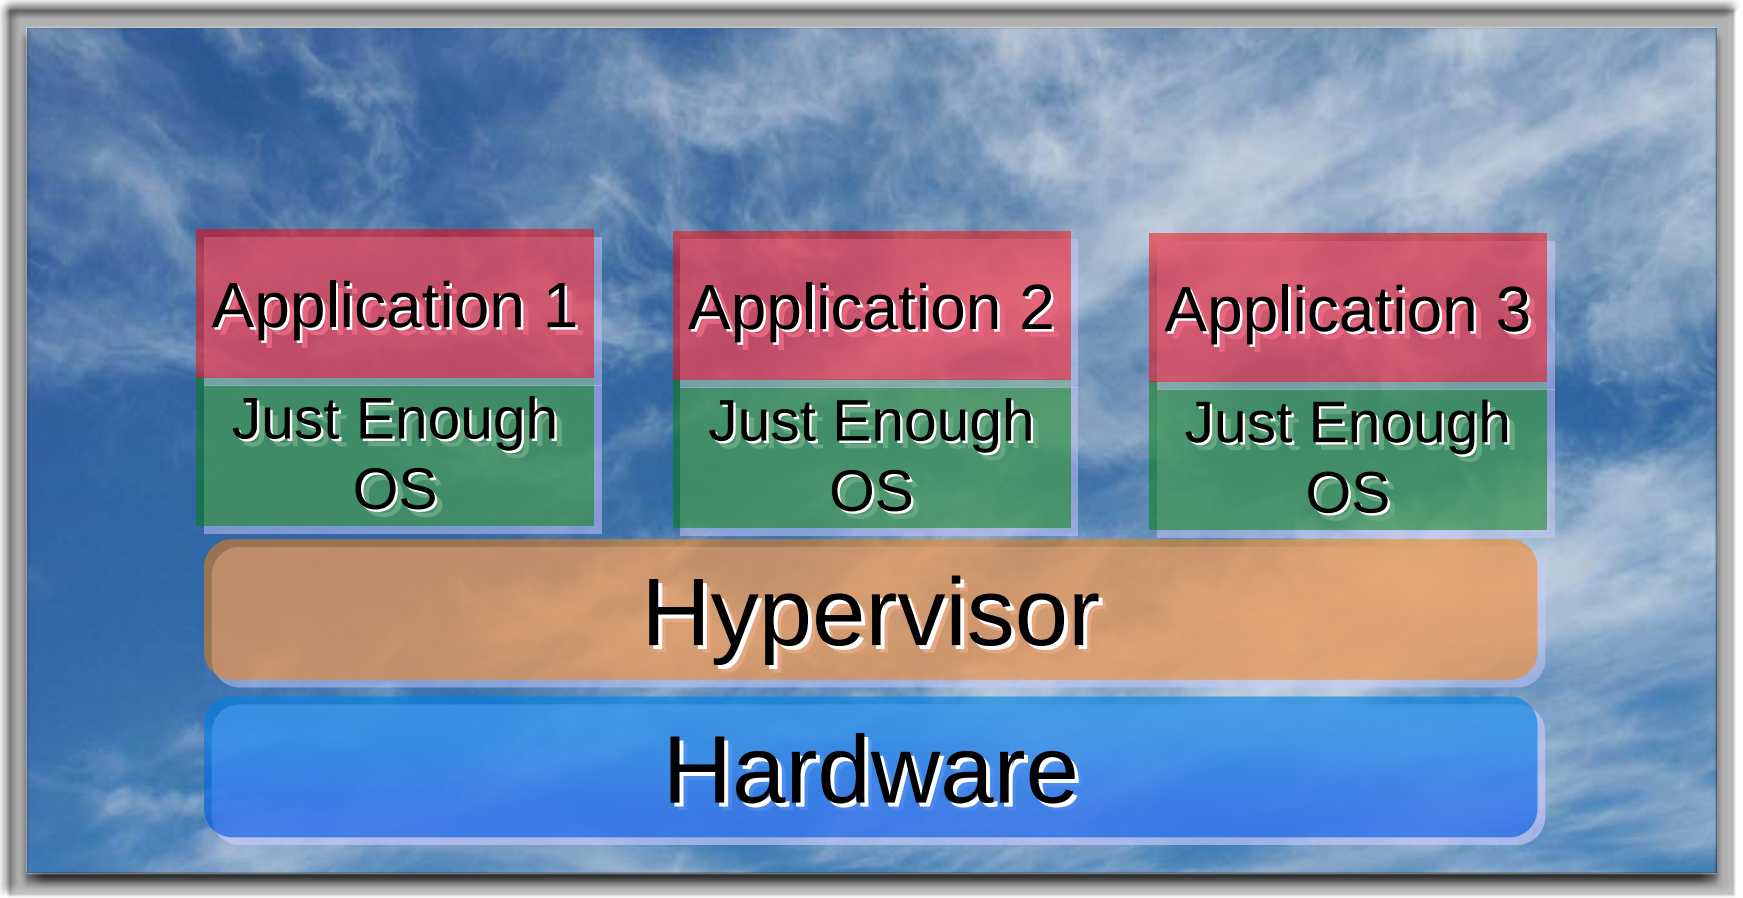

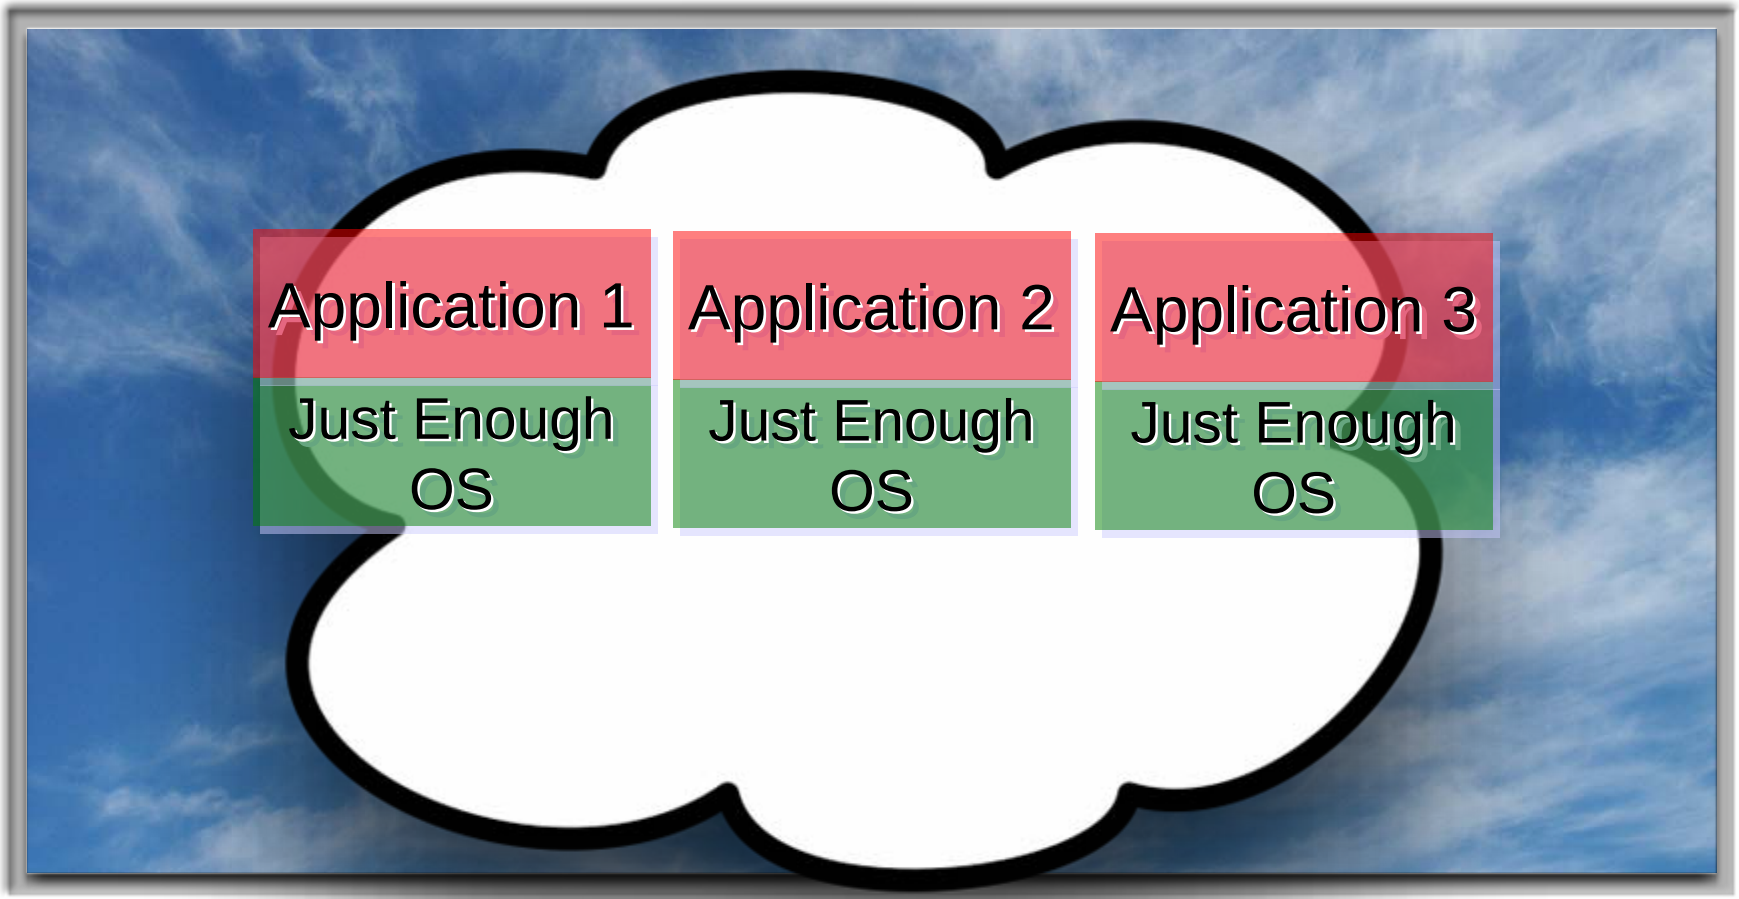

Once the hardware and hypervisor layer are commoditized and standardized we are able to abstract to a known element - the Cloud

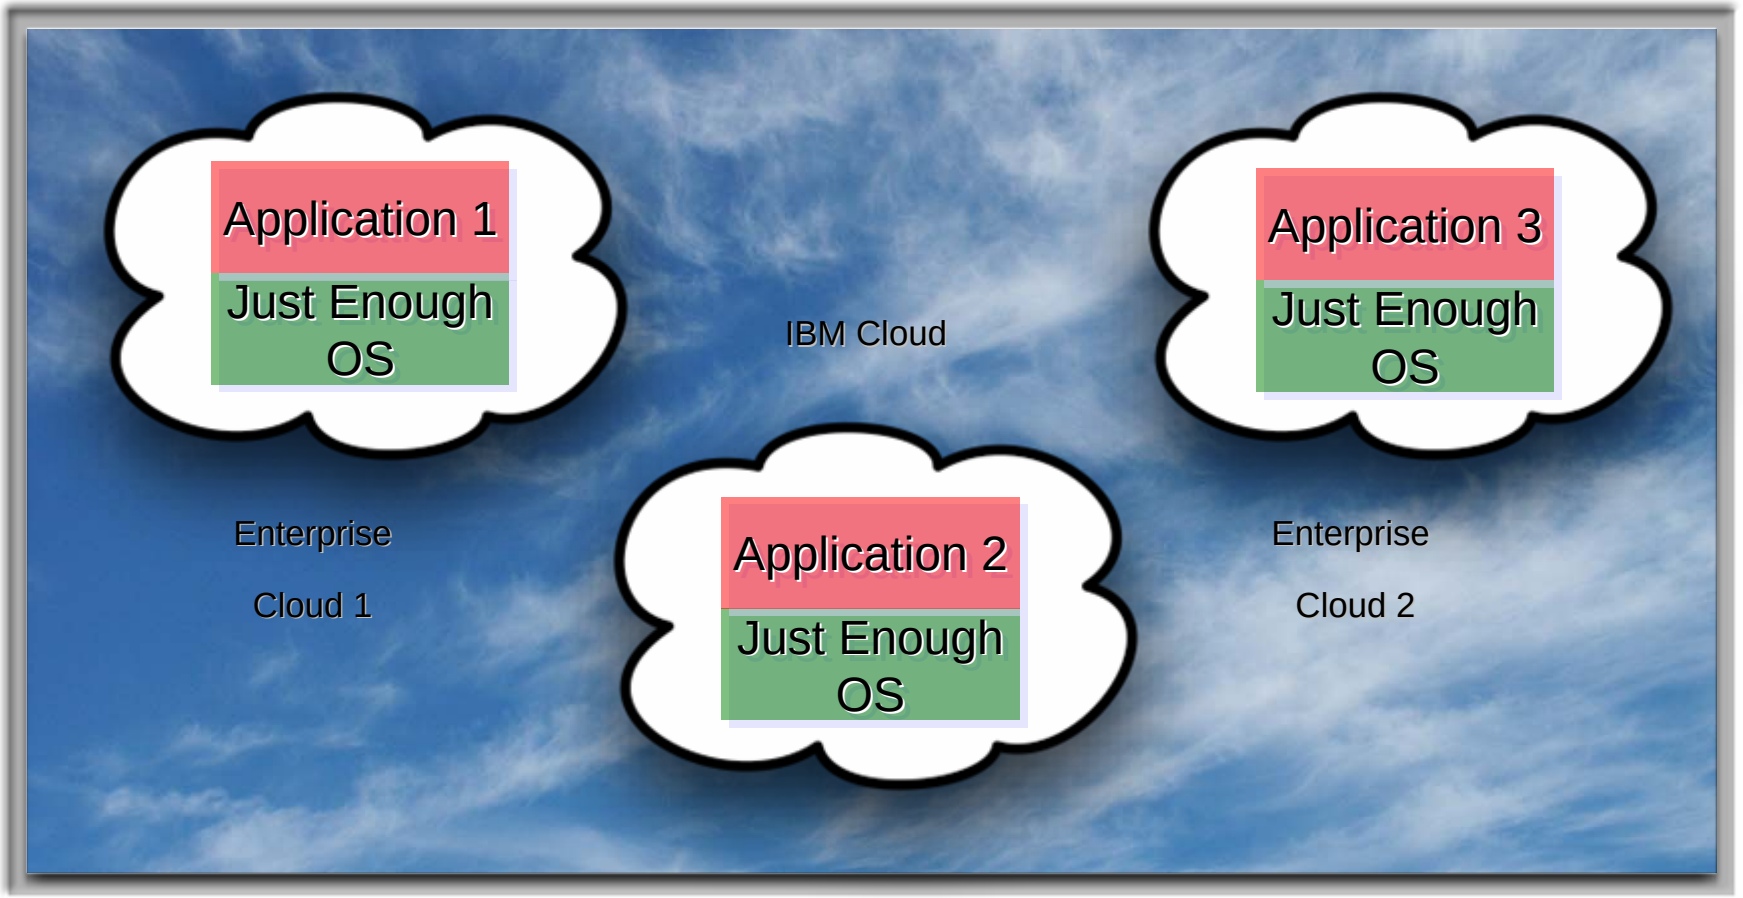

Once standardization permeates the industry the possibilities explode

Homogenization to Enables:

OS Images Good Practices

- Reduce the number of OS's supported
- Select based on Hypervisor compatibility
- Select based on customizability
- Reduce the number of configurations supported for patch and change management
- This applies to middleware and applications as well

Middleware Good Practices

- Support 1 version for production
- Support 1 version for development
- Make middleware cloneable if possible provisionable if required
- Lock in 6 - 12 month iteration
- DB virtualization can be poor
- Consider Multi-tenant instead

Chargeback and Process

- Business processes have some difficulty supporting these models
- Are enterprise IT consumers customers of the IT Dept?
- Is procurement done on individual machines?
- How is procurement and charging done on multi-tenant machines?
- Support model, often full access is allowed on cloud servers. This is a security and process constraint
- Is there a multi-tenant utilization model?
- How is procurement and charging done on multi-tenant machines?

IT Service Management

- Agents
  - Agent proliferation is an issue in virtualized data centers
  - Reduce, reuse (TCA) where possible
- Billing Cycle and machine lifetime may conflict eg. guest exists for less time than chargeable cycle
- CMDB management with short lived virtual guests is difficult
- How will software licenses be managed
- Change Control should be implemented STRICTLY
- Ideally exists at the image master level
- Image instances should be kept ‘vanilla’
- How will ID’s be provisioned and managed?
- Will network provisioning be automatic or manual
- What is the impact to automation

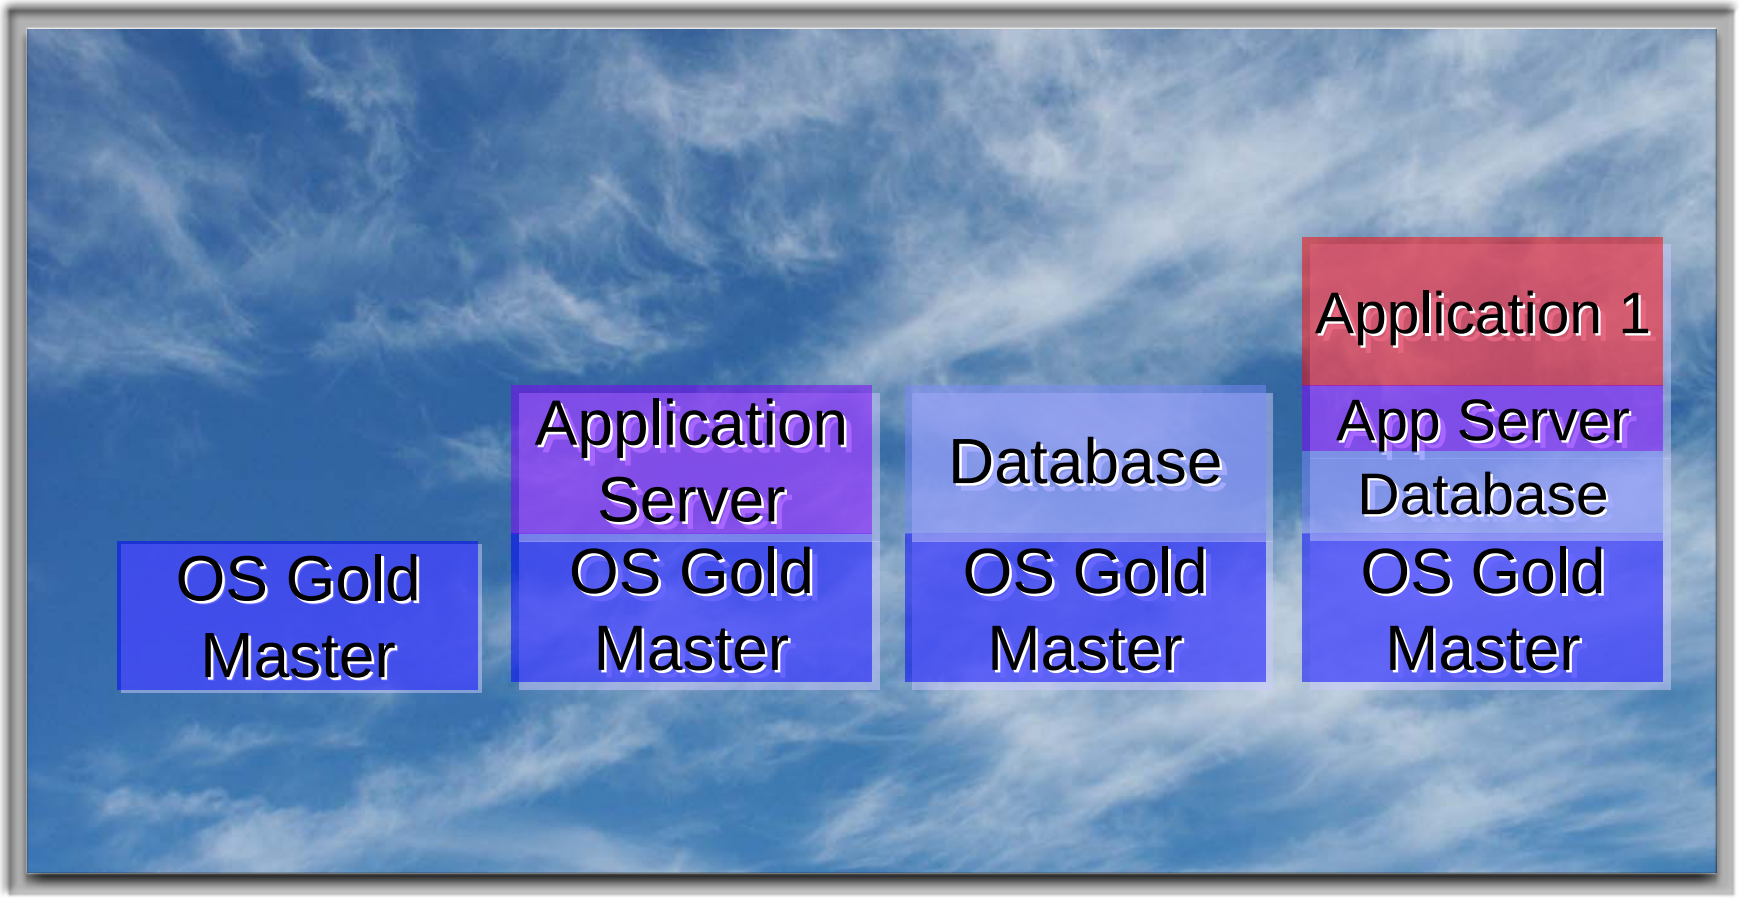

For common image configurations cloning works well. If complex or single instances will be used provisioning is preferred for repeatability and Image Management

Optimal environments to migrate to Cloud Computing

- Development, Test, and QA
- Hadoop / Map Reduce nodes
- On the production side:
  - Software as a Service
  - SOA Applications
- Lightweight Internally supported applications
  - Wiki
  - Blogs
  - Etc.

Dev, Test, & QA Cloud

- Provide self service user portal for resource requests
- Specify
  - OS & Arch
  - Middleware Stack
  - Network requirements
  - Storage (Thin Provisioning and SVC helps)
  - Length of resource allocation
  - SLA vs. Cost
  - Support requirements: self service (root access) vs. help desk

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.